2009

Sequential Detection and Identification of a Change in the Distribution of a Markov-Modulated Random Sequence

Savas Dayanik

Christian L. Goulding
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/fnce_papers

Part of the [Finance Commons](#), and the [Finance and Financial Management Commons](#)

Recommended Citation

Dayanik, S., & Goulding, C. (2009). Sequential Detection and Identification of a Change in the Distribution of a Markov-Modulated Random Sequence. *IEEE Transactions on Information Theory*, 55 (7), 3323-3345. <http://dx.doi.org/10.1109/TIT.2009.2021382>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/fnce_papers/347
For more information, please contact repository@pobox.upenn.edu.

Sequential Detection and Identification of a Change in the Distribution of a Markov-Modulated Random Sequence

Abstract

The problem of detection and identification of an unobservable change in the distribution of a random sequence is studied via a hidden Markov model (HMM) approach. The formulation is Bayesian, on-line, discrete-time, allowing both single- and multiple- disorder cases, dealing with both independent and identically distributed (i.i.d.) and dependent observations scenarios, allowing for statistical dependencies between the change-time and change-type in both the observation sequence and the risk structure, and allowing for general discrete-time disorder distributions. Several of these factors provide useful new generalizations of the sequential analysis theory for change detection and/or hypothesis testing, taken individually. In this paper, a unifying framework is provided that handles each of these considerations not only individually, but also concurrently. Optimality results and optimal decision characterizations are given as well as detailed examples that illustrate the myriad of sequential change detection and identification problems that fall within this new framework.

Disciplines

Finance | Finance and Financial Management

DETECTION AND IDENTIFICATION OF AN UNOBSERVABLE CHANGE IN THE DISTRIBUTION OF A MARKOV-MODULATED RANDOM SEQUENCE

SAVAS DAYANIK AND CHRISTIAN GOULDING

ABSTRACT. The problem of detection and diagnosis of an unobservable change in the distribution of a random sequence is studied via a hidden Markov model approach. The formulation is Bayesian, on-line, discrete-time, allowing both single- and multiple- disorder cases, dealing with both i.i.d. and dependent observations scenarios, allowing for statistical dependencies between the change-time and change-type in both the observation sequence and the risk structure, and allowing for general discrete-time disorder distributions. Several of these factors provide useful new generalizations of the sequential analysis theory for change detection and/or hypothesis testing, taken individually. In this paper, a unifying framework is provided that handles each of these considerations not only individually, but also concurrently. Optimality results and optimal decision characterizations are given as well as detailed examples that illustrate the myriad of sequential change detection and diagnosis problems that fall within this new framework.

1. INTRODUCTION, CONTRIBUTIONS, AND RELATED WORK

Suppose that we observe sequentially the random variables $X_1, X_2, \dots$ whose finite-dimensional distribution changes at an unobservable disorder time T due to an unobservable cause A , which represents one of several competing risks. Our objective is to detect quickly this disorder and determine accurately its cause based only on the observation sequence.

We approach this problem by modeling in a Bayesian framework the disorder time T , its cause A , and the *distribution* of the observation sequence $X := \{X_t; t \geq 1\}$, as functionals of an underlying hidden Markov chain Y . We derive an optimal solution and useful characteristics of its structure. We show how the classical sequential change detection and hypothesis-testing problems and recent extensions fit this formulation. Furthermore, we illustrate how this framework provides a foundation for the study of an expansive range of new sequential analysis problems.

This problem has been studied extensively in very restricted forms under headings such as sequential change detection and sequential (multiple) hypothesis testing, employing a variety of approaches, with numerous cited applications, including fault detection and isolation in industrial processes, target detection and identification in national defense, pattern

2000 *Mathematics Subject Classification.* Primary 62L10; Secondary 62L15, 62C10, 60G40.

Key words and phrases. sequential change detection, sequential hypothesis testing, hidden Markov models, optimal stopping.

recognition and machine learning, radar and sonar signal processing, seismology, speech and image processing, biomedical signal processing, finance, and insurance. Foundational studies include the non-Bayesian (minimax) formulation proposed by Lorden [21] and the Bayesian formulation proposed by Shirayev [27] for sequential change detection and the papers of Wald and Wolfowitz [28] and Arrow, Blackwell, and Girshick [1] on sequential hypothesis testing. We refer the reader to Basseville and Nikiforov [3], Dragalin, Tartakovsky, and Veeravalli [10, 11], and Lai [17, 19] and the references therein for detailed background on these topics, recent developments, and discussion of applications. Despite the progress that has been made in these areas, the literature offers few non-asymptotic optimality results and provides for the most part very limited models of the general problem, and hence many important considerations have remained open. In this paper we provide a comprehensive formulation that addresses such considerations.

The contributions of this paper and its connection to the literature are detailed below. The organization of the paper is outlined in Section 1.6.

1.1. General disorder distributions. Almost all Bayesian approaches to sequential change detection and/or sequential hypothesis testing have been limited to the (zero-modified) geometric prior distribution for the disorder time in discrete-time settings (or its continuous-time analog of the exponential distribution for continuous-time frameworks). This is due primarily to the memoryless property of the geometric distribution which lends readily to a Markovian sufficient statistic that can be used, for example, in an optimal stopping framework. A geometric disorder time can be a reasonable assumption in some cases, such as in modeling the lifetime of a highly reliable system. However, in many cases the disorder distribution of interest cannot be approximated adequately by a geometric distribution. In fact, the disorder is often the result of several competing risks, and hence the disorder time depends on some underlying root cause, which cannot be captured in general with the geometric assumption. Such dependencies arise frequently in a variety of applications; see Crowder [8] for an overview. In this paper we provide an optimal solution for more general disorder distributions through the use of phase-type representations. See Examples 1, 2, and 3 of Section 2 for more details.

1.2. Statistically dependent disorder time and its cause. Just as it is natural to expect consequences of the disorder to depend on its cause, it is also of interest to diagnose the underlying cause of the disorder in addition to detecting when it happens. Yet, the literature has been sparse along this direction. The first results regarding the extension of the sequential change detection problem to include the diagnosis task are given by Nikiforov [25] and Lai [18] in a non-Bayesian framework. Dayanik, Goulding, and Poor [9] study the extension in a Bayesian framework, albeit under the assumption of statistical *independence* between the disorder and its cause. However, not only may the disorder be statistically dependent on its cause, but also the detection delay cost (as well as the false alarm and

misdiagnosis costs) may depend on the cause of the disorder. In this paper we incorporate such dependencies and their expanded risk structure in a Bayesian framework for the joint problem of detection and diagnosis; see Remark 1 and Examples 2 and 4 of Section 2.

1.3. Multiple regime changes. The on-line change detection and hypothesis testing literature has focused on models that recognize only a *single abrupt* change in the distribution of observations. This restriction to a single change (i.e., a “two-regime” model) is oversimplifying and inappropriate for virtually all real applications. In addition, the restriction to an *abrupt* change can be unsuitable since *gradual* changes in a system’s performance are often more realistic. The solution we present in this paper applies, with or without these restrictions, to the separate problems of on-line change detection and hypothesis testing as well as the joint problem of detection and diagnosis; see Example 4 of Section 2 for more details.

1.4. Markov channels with noise. The Bayesian on-line change detection and hypothesis testing literature has focused much attention on the case of i.i.d. observations within each regime. An exception is a paper by Yakir [29] that gives a solution for the quickest detection of a single abrupt unobservable change in the transition matrix of a Markov chain based on sequential observations of the states of the chain, where the change time has a (zero-modified) geometric prior distribution. Our model includes that solution as a special case but also allows for generalizations beyond its scope, such as “noisy” observations of the chain (see Example 5 of Section 2), general disorder distributions, multiple regime changes, and several ultimate regimes.

1.5. Hidden Markov models. The observation sequence X together with its modulating unobservable Markov chain Y forms a doubly stochastic process known as a hidden Markov model. Such models have appeared in an increasing variety of applications, such as digital communications, economics, molecular biology, speech recognition, etc. However, the main focus has been on the inference techniques for the hidden state, while the detection and diagnosis problem has received much less attention except in simple models. An exception is a recent paper by Fuh [12] that studies the detection problem in a non-Bayesian framework. In the current paper, we solve the joint detection and diagnosis problem for a hidden Markov model in a Bayesian framework.

1.6. Organization of the paper. In Section 2 we formulate precisely the problem and illustrate its flexibility on a variety of examples. In Section 3 we carry out a Bayesian posterior analysis and derive an optimal solution to the problem. In Section 4 we provide numerical illustrations of the optimal sequential decision strategy. Proofs of selected results are deferred to the appendix.

2. PROBLEM DESCRIPTION

Let $Y = \{Y_t; t \geq 0\}$ be a time-homogeneous Markov chain on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with finite state space $\mathcal{Y}$, initial state distribution η , and one-step transition matrix P . Suppose that $\mathcal{Y}_1, \dots, \mathcal{Y}_a$ are a closed (but not necessarily irreducible) mutually disjoint subsets of the state space $\mathcal{Y}$, and the collection $\mathcal{Y}_0 = \mathcal{Y} \setminus \bigcup_{k=1}^a \mathcal{Y}_k$ of the remaining states does not contain any closed sets of $\mathcal{Y}$; in other words, every state in $\mathcal{Y}_0$ is transient. The submatrices of P corresponding to subsets $\mathcal{Y}_0, \mathcal{Y}_1, \dots, \mathcal{Y}_a$ are denoted by $P_0, P_1, \dots, P_a$, respectively, and the submatrices $Q_1, \dots, Q_a$ denote the transition probabilities from $\mathcal{Y}_0$ into $\mathcal{Y}_1, \dots, \mathcal{Y}_a$, respectively:

$$(2.1) \quad P = \begin{array}{c} \begin{array}{ccccc} & \mathcal{Y}_0 & \mathcal{Y}_1 & \mathcal{Y}_2 & \dots & \mathcal{Y}_a \\ \begin{bmatrix} P_0 & Q_1 & Q_2 & & Q_a \\ & P_1 & & & \\ & & P_2 & \dots & \\ & & \vdots & \ddots & \\ & & & & P_a \end{bmatrix} & \begin{array}{c} \mathcal{Y}_0 \\ \mathcal{Y}_1 \\ \mathcal{Y}_2 \\ \vdots \\ \mathcal{Y}_a \end{array} \end{array} \end{array}$$

Since Y has finitely many states, it leaves the transient set $\mathcal{Y}_0$ eventually and is absorbed into one of the closed sets $\mathcal{Y}_1, \dots, \mathcal{Y}_a$ almost surely. Let

$$T := \min\{t \geq 0; Y_t \notin \mathcal{Y}_0\} \quad \text{and} \quad A := \arg\{j = 1, \dots, a; Y_T \in \mathcal{Y}_j\} \quad \text{on } \{T < \infty\}$$

be the absorption time and the closed set that absorbs the Markov chain Y , respectively.

Let $\mathbb{P}(y, dx)$, $y \in \mathcal{Y}$ be probability measures on some common measurable space $(\mathcal{X}, \mathcal{X})$. They always admit density functions $f(y, x)$, $y \in \mathcal{Y}$ with respect to some sigma-finite measure $\nu(dx)$ on $(\mathcal{X}, \mathcal{X})$; for example, we can take $\nu(dx) = \sum_{y \in \mathcal{Y}} \mathbb{P}(y, dx)$. Let $X = \{X_t; t \geq 1\}$ be a new stochastic process on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$ such that

$$\mathbb{P}\{Y_0 = y_0, Y_1 = y_1, \dots, Y_t = y_t, X_1 \in \mathcal{X}_1, \dots, X_t \in \mathcal{X}_t\} = \eta(y_0) \prod_{n=1}^t P(y_{n-1}, y_n) \mathbb{P}(y_n, \mathcal{X}_n)$$

for every $t \geq 1$, $y_0, y_1, \dots, y_t \in \mathcal{Y}$, and $\mathcal{X}_1, \dots, \mathcal{X}_t \in \mathcal{X}$. Namely, given $Y_1, \dots, Y_t$, the random variables $X_1, \dots, X_t$ are independent with conditional distributions $\mathbb{P}(Y_1, dx), \dots, \mathbb{P}(Y_t, dx)$ on $(\mathcal{X}, \mathcal{X})$, respectively; in other words, Y modulates the distribution of the process X .

However, the Markov chain Y is unobservable in general. *The problem is to detect the absorption time T of the hidden Markov chain Y into one of the closed sets $\mathcal{Y}_1, \dots, \mathcal{Y}_a$ as soon as it happens and, at the same time, to identify the absorbing set A as accurately as possible based only on the observations of the process X .*

More precisely, if $\mathbb{F} = \{\mathcal{F}_t\}_{t \geq 0}$ denotes the observation filtration with

$$(2.2) \quad \mathcal{F}_0 = \{\emptyset, \Omega\} \quad \text{and} \quad \mathcal{F}_t = \sigma\{X_1, \dots, X_t\}, \quad t \geq 1,$$

then a *decision rule* (τ, δ) is by definition a pair consisting of an $\mathbb{F}$ -stopping time τ that indicates the absorption time and a $\{1, \dots, a\}$ -valued $\mathcal{F}_\tau$ -measurable random variable δ that identifies the absorbing set. We denote by Δ the collection of decision rules and define the *Bayes risk*

$$(2.3) \quad u(\tau, \delta) := \mathbb{E} \left[\sum_{t=0}^{\infty} c(Y_t) 1_{\{t < \tau\}} + 1_{\{\tau < \infty\}} C(Y_\tau, \delta) \right] \quad \text{for every } (\tau, \delta) \in \Delta$$

as the sum of the expected detection delay cost $\mathbb{E} [\sum_{t=0}^{\infty} c(Y_t) 1_{\{t < \tau\}}]$ and the expected identification cost $\mathbb{E} [1_{\{\tau < \infty\}} C(Y_\tau, \delta)]$ for some known penalty functions $c : \mathcal{Y} \mapsto [0, \infty)$ and $C : \mathcal{Y} \times \{1, \dots, a\} \mapsto [0, \infty)$. Then the objective is (i) to calculate the minimum Bayes risk

$$u^* := \inf_{(\tau, \delta) \in \Delta} u(\tau, \delta),$$

and (ii) to find a decision rule $(\tau, \delta) \in \Delta$ whose Bayes risk $u(\tau, \delta)$ attains the minimum u^* , if it exists.

Remark 1. Perhaps it is more natural to express the expected cost of delay in detecting the time T that the Markov chain Y exits the set $\mathcal{Y}_0$ by $\mathbb{E} [\sum_{t=0}^{\infty} c(Y_t) 1_{\{T \leq t < \tau\}}]$. However, this follows easily from our definition $\mathbb{E} [\sum_{t=0}^{\infty} c(Y_t) 1_{\{t < \tau\}}]$ of expected delay cost if we set $c(y) = 0$ for every $y \in \mathcal{Y}_0$, because $Y_t \in \mathcal{Y}_0$ on $\{t < T\}$. In fact, if

$$c(y) = \begin{cases} 0, & y \in \mathcal{Y}_0, \\ c, & y \in \mathcal{Y} \setminus \mathcal{Y}_0, \end{cases} \quad \text{and} \quad C(y, j) = \begin{cases} 1, & y \in \mathcal{Y}_0, 1 \leq j \leq a, \\ \tilde{C}(i, j), & y \in \mathcal{Y}_i, 1 \leq i \leq a, 1 \leq j \leq a \end{cases}$$

for some known constant $c > 0$ and function $\tilde{C} : \{1, \dots, a\} \times \{1, \dots, a\} \mapsto [0, \infty)$ satisfying $\tilde{C}(i, i) = 0$ for every $i \in \{1, \dots, a\}$, then the Bayes risk $u(\tau, \delta)$ of every decision rule $(\tau, \delta) \in \Delta$ simplifies to

$$(2.4) \quad \mathbb{P}\{\tau < T\} + c \mathbb{E}[(\tau - T)^+] + \mathbb{E} [1_{\{T \leq \tau < \infty\}} \tilde{C}(A, \delta)],$$

which can be regarded as a sum of expected false-alarm, detection-delay, and misdiagnosis penalties, respectively.

Under this Bayes risk in (2.4), Dayanik et al. [9] study the joint detection and identification of a disorder time with zero-modified geometric prior distribution and a disorder type with a finite prior distribution *independent* of the disorder time. More precisely, they assume that one of a alternatives, labeled $1, \dots, a$, occur with probabilities $\alpha = (\alpha_1, \dots, \alpha_a)$ after a sudden change at some unobservable time that is zero with some probability q and has otherwise a geometric distribution with “success” parameter p independently of the alternative. This change or disorder time is said to have *zero-modified geometric distribution with parameters q and p* . In terms of the notation and Markov model of disorder above, Dayanik et al. [9] have a hidden Markov chain Y on $\mathcal{Y} = \{0, 1, \dots, a\}$ with initial state distribution $\eta =$

$(1 - q, q\alpha_1, \dots, q\alpha_a)$ and one-step transition function

$$P = \begin{bmatrix} 1-p & p\alpha_1 & \dots & p\alpha_a \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{bmatrix},$$

where T is the first time the chain enters into one of the closed sets $\mathcal{Y}_i = \{i\}$, $i = 1, \dots, a$, while $A = Y_T$. It is easy to check that the disorder time T and disorder type A are independent and have the prescribed distributions. Dayanik et al. [9] describe a solution for this special case and Bayes risk in (2.4).

Unfortunately, the solution obtained by Dayanik et al. [9] does not extend when the disorder time T and change type A are *not* statistically independent. For example, suppose that the change type has distribution α on $\{1, \dots, a\}$ as before, but for each $i = 1, \dots, a$, the disorder time has *conditionally* zero-modified geometric distribution with parameters q_i and p_i given that $A = i$. This example falls within the framework of the current paper: if Y is the Markov chain on $\mathcal{Y} = \{(1, 0), (2, 0), \dots, (a, 0), (1, 1), (2, 1), \dots, (a, 1)\}$ with initial state distribution

$$\eta = ((1 - q_1)\alpha_1, \dots, (1 - q_a)\alpha_a, q_1\alpha_1, \dots, q_a\alpha_a)$$

and one-step transition function

$$P = \left[\begin{array}{ccc|ccc} 1-p_1 & & & p_1 & & \\ & \ddots & & & \ddots & \\ & & 1-p_a & & p_a & \\ \hline & & & 1 & & \\ & & & & \ddots & \\ & & & & & 1 \end{array} \right],$$

and if $\mathcal{Y}_0 = \{(i, 0) \mid 1 \leq i \leq a\}$ and $\mathcal{Y}_i = \{(i, 1)\}$ for $i = 1, \dots, a$, then $T = \min\{t \geq 0; Y_t \notin \mathcal{Y}_0\}$ and $A = Y_T$ have the prescribed joint distribution.

Various other examples are given below to illustrate further the versatility of the new framework for detection and diagnosis of an unobservable change. Briefly, they illustrate (i) how to work with phase-type distributions, which can approximate arbitrarily well any prior distribution of a disorder time, (ii) how to model jointly a disorder time and its cause when they are not statistically independent, (iii) how to incorporate cyclic shifts in the proneness of a physical system to change, (iv) how to detect and diagnose the inception of a particular regime when there are multiple serial regimes and/or several scenarios of regime succession, (v) how to detect the onset of a gradual change, and (vi) how to detect a change in the

transition matrix of a Markov chain, even with “noisy” observations. The solution of the general problem will be presented in Section 3 after these motivational examples.

Example 1 (Phase-type distributed disorder times). Suppose that exactly one of the states of the Markov chain Y is absorbing, and the others are transient. If $a = 1$, and $\mathcal{Y}_1$ contains only the absorbing state, then $\mathcal{Y}_0 = \mathcal{Y} \setminus \mathcal{Y}_1$ contains transient states, and $T = \min\{t \geq 0; Y_t \notin \mathcal{Y}_0\}$ is said to have a discrete *phase-type distribution with representation* (η_0, P_0) , where η_0 is the restriction of the initial state distribution η of Y to the transient states in $\mathcal{Y}_0$; see Neuts [23, 24], Latouche and Ramaswami [20], Bobbio et al. [5]. If we denote by

$$(2.5) \quad I, \quad e, \quad e_1, \quad e_2, \quad \dots$$

an identity matrix, a column vector whose every entry is one, and unit column vectors, all of whose dimensions will be clear from the context, then

$$(2.6) \quad \mathbb{P}\{T = 0\} = 1 - \eta_0 e \quad \text{and} \quad \mathbb{P}\{T = t\} = \eta_0 P_0^{t-1} Q_1, \quad t \geq 1.$$

The matrix $I - P_0$ is invertible, and

$$\sum_{t=0}^{\infty} P_0^t = (I - P_0)^{-1}, \quad Q_1 = e - P_0 e = (I - P_0)e, \quad \text{and} \quad (I - P_0)^{-1} Q_1 = e;$$

see, for example, Çinlar [7, Chapter 6]. Therefore,

$$\mathbb{P}\{T < \infty\} = 1 - \eta_0 e + \sum_{t=1}^{\infty} \eta_0 P_0^{t-1} Q_1 = 1 - \eta_0 e + \eta_0 (I - P_0)^{-1} Q_1 = 1 - \eta_0 e + \eta_0 e = 1,$$

$$\mathbb{P}\{T > t\} = \mathbb{P}\{Y_t \in \mathcal{Y}_0\} = \eta_0 \sum_{n=t+1}^{\infty} P_0^{n-1} Q_1 = \eta_0 P_0^t (I - P_0)^{-1} Q_1 = \eta_0 P_0^t e, \quad t \geq 0,$$

$$\mathbb{E}T = \sum_{t=0}^{\infty} \mathbb{P}\{T > t\} = \sum_{t=0}^{\infty} \eta_0 P_0^t e = \eta_0 \sum_{t=0}^{\infty} P_0^t e = \eta_0 (I - P_0)^{-1} e < \infty.$$

Every discrete distribution with finitely many atoms is a phase-type distribution. Suppose, for example, that $(p_1, \dots, p_s, p_0)$ is a probability distribution on $\mathcal{Y} := \{1, \dots, s, 0\}$ for some integer $s \geq 0$. Let $(Y_t)_{t \geq 0}$ be a Markov chain on the state space $\mathcal{Y}$ with initial state distribution $\eta = (p_1, \dots, p_s, p_0)$ and one-step transition matrix P all of whose entries are zero except $P(s, s-1) = \dots = P(1, 0) = P(0, 0) = 1$; see Figure 1. Then the first exit time $T = \inf\{n \geq 0 : Y_n \notin \mathcal{Y}_0\}$ of Y from $\mathcal{Y}_0 := \{1, \dots, s\}$ has distribution $(p_1, \dots, p_s, p_0)$ on $\mathcal{Y}$. Such a distribution is a natural candidate as a prior distribution for the time of a disorder in the arrival/demand process for services or products that have an expiration date, such as the arrivals to motor vehicle service offices for the renewal of driver licenses expiring on a fixed date or tax-software sales between its release time and April 15th of each year.

Example 2 (Statistically dependent disorder time and disorder cause). In many situations it is natural to expect that the disorder time is determined partly by its cause. A typical

FIGURE 1. The state-diagram of the hidden Markov chain representing discrete distribution $(p_1, \dots, p_s, p_0)$ on $\{1, \dots, s, 0\}$ as a discrete phase-type distribution.

example in reliability is a mechanical system whose failure is a function of competing risks, such as corrosion and random vibration. If failure occurs within a relatively short time, then its cause may be most likely from random vibration. Indeed, it is reasonable to expect that the tendency of failure due to random vibration stays constant over time whereas the tendency of corrosion-based failure increases with age; hence, the failure time and its cause are statistically dependent. We can incorporate such dependencies as follows.

Suppose that we are given the joint distribution

$$\mathbb{P}\{T = t, A = i\}, \quad t \geq 0, \quad i = 1, \dots, a$$

of the change time T and its cause A . We can construct the state-space $\mathcal{Y} = \bigcup_{i=0}^a \mathcal{Y}_i$, initial state distribution η , and one-step transition matrix P of a Markov chain Y that models the random variables $T = \min\{t \geq 0; Y_t \notin \mathcal{Y}_0\}$ and $A = \arg\{i = 1, \dots, a; Y_T \in \mathcal{Y}_i\}$ as follows:

- (i) Compute the marginal distribution of A ; i.e.,

$$\alpha_i := \mathbb{P}\{A = i\} = \sum_{t=0}^{\infty} \mathbb{P}\{T = t, A = i\}, \quad i = 1, \dots, a.$$

- (ii) Find a phase-type representation $(\eta^{(i)}, P^{(i)})$ for the conditional change-time distribution

$$\mathbb{P}\{T = t \mid A = i\}, \quad t \geq 0 \quad \text{for each } i = 1, \dots, a.$$

A phase-type distribution can be approximated to any distribution; see Johnson and Taaffe [16] and Asmussen et al. [2]. See Horváth and Telek [15] for an algorithm designed for fitting discrete phase-type distributions. See also Horton and Isensee [14] for discussion of algorithms used in approximating discrete phase-type distributions.

- (iii) Let the set $\mathcal{Y}_i := \{i\}$ represent cause i , let the set $\mathcal{Y}_0^{(i)}$ consist of the states corresponding to each phase-type representation (i.e., one state for each element of $\eta^{(i)}$) for each $i = 1, \dots, a$, such that each of these sets are mutually disjoint, and define the set $\mathcal{Y}_0 := \mathcal{Y}_0^{(1)} \cup \dots \cup \mathcal{Y}_0^{(a)}$.
- (iv) Then form the initial state distribution and one-step transition matrix, respectively, as

$$\eta = (\alpha_1 \eta^{(1)}, \dots, \alpha_a \eta^{(a)}, \alpha_1 (1 - \eta^{(1)} e), \dots, \alpha_a (1 - \eta^{(a)} e))$$

and

$$P = \left[\begin{array}{ccc|ccc} \mathcal{Y}_0^{(1)} & \cdots & \mathcal{Y}_0^{(a)} & \mathcal{Y}_1 & \cdots & \mathcal{Y}_a \\ P^{(1)} & & & (I - P^{(1)})e & & \\ & \ddots & & & \ddots & \\ & & P^{(a)} & & & (I - P^{(a)})e \\ \hline & & & 1 & & \\ & & & & \ddots & \\ & & & & & 1 \end{array} \right] \begin{bmatrix} \mathcal{Y}_0^{(1)} \\ \vdots \\ \mathcal{Y}_0^{(a)} \\ \mathcal{Y}_1 \\ \vdots \\ \mathcal{Y}_a \end{bmatrix}.$$

Compare with the dependency formulation of Remark 1; see also Example 4 for an illustration of this type of approach.

Example 3 (Cyclic disorder-time distributions). Cyclic disorder-time distributions arise in many change-detection applications in which observations are modulated by regular cycles of behavior. A basic example is machinery monitoring during daily peak and off-peak utilization sub-periods, where the disorder tendency may be, for example, higher during peak sub-periods.

Suppose that every period is divided into three sub-periods, labeled 1, 2, and 3. A disorder may have happened already in the past with probability p_0 . Given that it has not happened yet, it happens with probability $\theta \in [0, 1]$ some time in the next period. Conditionally on the disorder happening in a given period, it happens in sub-period 1, 2, or 3 with probabilities p_1 , p_2 , and p_3 , respectively, for some $0 \leq p_1, p_2, p_3 \leq 1$ such that $p_1 + p_2 + p_3 = 1$. If time is measured in the number of “sub-periods,” then the disorder time T has distribution

$$(2.7) \quad \mathbb{P}\{T = 0\} = p_0 \quad \text{and} \quad \mathbb{P}\{T = 3n + k\} = (1 - p_0)(1 - \theta)^n \theta p_k, \quad n \geq 0, \quad k = 1, 2, 3.$$

This distribution is the same as that of the exit time of a suitable finite-state Markov chain. More precisely, let $(Y_t)_{t \geq 0}$ be a Markov chain on the state space $\mathcal{Y} = \{1, 2, 3, 0\}$ with initial state distribution $\eta \equiv (1 - p_0, 0, 0, p_0)$ and one-step transition matrix

$$P = \begin{bmatrix} 0 & 1 - \tilde{p}_1 & 0 & \tilde{p}_1 \\ 0 & 0 & 1 - \tilde{p}_2 & \tilde{p}_2 \\ 1 - \tilde{p}_3 & 0 & 0 & \tilde{p}_3 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \text{with} \quad P_0 = \begin{bmatrix} 0 & 1 - \tilde{p}_1 & 0 \\ 0 & 0 & 1 - \tilde{p}_2 \\ 1 - \tilde{p}_3 & 0 & 0 \end{bmatrix}, \quad Q_1 = \begin{bmatrix} \tilde{p}_1 \\ \tilde{p}_2 \\ \tilde{p}_3 \end{bmatrix},$$

whose entries are expressed in terms of the *hazard probabilities*

$$\tilde{p}_1 := \theta p_1, \quad \tilde{p}_2 := \frac{\theta p_2}{1 - \theta p_1}, \quad \tilde{p}_3 := \frac{\theta p_3}{1 - \theta p_1 - \theta p_2};$$

see also Figure 2(a). If $\mathcal{Y}_0 := \{1, 2, 3\}$, then the exit time $T := \inf\{t \geq 0 : Y_t \notin \mathcal{Y}_0\}$ of the Markov chain $(Y_t)_{t \geq 0}$ from $\mathcal{Y}_0$ is the same as its entrance time to the absorbing state 0, and its distribution given by (2.6) can be shown to be the same as that in (2.7).

FIGURE 2. The state-diagrams of the Markov chain representing cyclic distributions (a) in (2.7) and (b) in (2.8) as a discrete phase-type distributions.

In general, suppose that a period consists of s sub-periods, labeled $1, 2, \dots, s$. A disorder may happen in a given period with probability θ given that it has not yet happened. Conditionally on that it has happened some time in a given period, it happens in sub-period $1, \dots, s$ respectively with probabilities $p_1, \dots, p_s$. If the probability that the disorder has already happened in the past is p_0 , then the distribution of the disorder time T is

$$(2.8) \quad \mathbb{P}\{T = 0\} = p_0, \quad \mathbb{P}\{T = sn + k\} = (1 - p_0)(1 - \theta)^n \theta p_k, \quad n \geq 0, 1 \leq k \leq s,$$

If $\mathcal{Y} := \{1, 2, \dots, s, 0\}$ and $\mathcal{Y}_0 := \{1, 2, \dots, s\}$, then this is the same as the distribution of the exit time $T := \inf\{t \geq 0 : Y_t \notin \mathcal{Y}_0\}$ from $\mathcal{Y}_0$ of a Markov chain $(Y_t)_{t \geq 0}$ with state space $\mathcal{Y}$, initial state distribution $\eta = (1 - p_0, 0, \dots, 0, p_0)$, and one-step transition matrix P whose nonzero entries $P(0, 0) = 1$, $P(s, 0) = 1 - P(s, 1) = \tilde{p}_s$, $P(i, 0) = 1 - P(i, i + 1) = \tilde{p}_i$, $1 \leq i \leq s$ are expressed in terms of the hazard probabilities

$$\tilde{p}_i := \frac{\theta p_i}{1 - \theta \sum_{k=1}^{i-1} p_k}, \quad 1 \leq i \leq s$$

of the distribution $(p_0, p_1, \dots, p_s)$ on $\{0, 1, \dots, s\}$; see Figure 2(b) for the state diagram. Note that if $s = 1$, then we must have $p_1 = 1$, and the distribution in (2.8) of the disorder time T becomes the zero-modified geometric distribution with success probability θ and probability mass p_0 at 0.

Example 4 (Detection and diagnosis when there are multiple regime changes). The sequential change detection and hypothesis testing literature deals almost exclusively with random sequences that undergo only one regime shift. Accordingly, the natural goal in these studies is to detect this sole change. However, it is frequently the case in real applications that the

random sequence of interest may actually be subject to several successive regime shifts, and the goal is to detect the occurrence of a particular shift, especially the first or last shift. For instance, more than one characteristic of the sequence may be subject to change and each of these changes need not occur simultaneously. Moreover, the occurrence of a change in one characteristic does not necessarily preclude the other characteristics from changing in the future, resulting in additional changes to the distribution of the ensuing observations. So, if we are interested in detecting when at least one change has occurred, then we must take into account the potential time-inhomogeneity of the observation process beyond the first change whenever our decision is delayed. That is, the law of observations may continue to change beyond the first disorder while we are still waiting to raise the alarm. This is especially relevant if we are also interested in diagnosing the cause of the first change since the diagnosis task does not discourage waiting and is likely to take more observations than the detection task alone. Likewise, similar considerations are needed if we are interested in detecting the last change (e.g., when the random sequence reaches stationarity). The following cases illustrate how this departure from the “standard” assumption of a single change falls within the framework of the current paper.

Quickest detection of the k th out of n successive regime shifts. Consider an observation process $X := \{X_t; t \geq 1\}$ that progresses through a series of n different probability regimes whose durations form an unobservable sequence of nonnegative a.s.-finite random variables $S_1, S_2, \dots, S_n$. Following the n th regime change the sequence X enters its final regime and remains there forever. Let the beginning of each regime epoch be denoted by the times

$$T_0 \equiv 0 \quad \text{and} \quad T_i := S_1 + \dots + S_i, \quad i = 1, \dots, n.$$

Suppose that conditionally on the disorder times $T_1, T_2, \dots, T_n$, the random variables $X_1, X_2, \dots$ are independent and have a common probability density, one of $f(0, \cdot), \dots, f(n, \cdot)$, that depends on the index of the most recent disorder:

$$\overbrace{\underbrace{X_1, \dots, X_{T_1-1}}_{\sim f(0, \cdot)} \underbrace{X_{T_1}, \dots, X_{T_2-1}}_{\sim f(1, \cdot)} \dots \underbrace{X_{T_{n-1}}, \dots, X_{T_n-1}}_{\sim f(n-1, \cdot)} \underbrace{X_{T_n}, X_{T_n+1}, \dots}_{\sim f(n, \cdot)}}^{\text{independent given } T_1, T_2, \dots, T_n}$$

For a fixed k in $\{1, \dots, n\}$, the objective is to detect the k th change out of the n successive regime changes as quickly as possible with a low frequency of false alarm based only on observations of the process X .

Specifically, suppose that the regime durations $S_1, \dots, S_n$ are independent random variables and that the duration S_i of the i th regime has the zero-modified geometric distribution with parameters q_i and p_i , $i = 1, \dots, n$. Let $Y = \{Y_t; t \geq 0\}$ be the counting process

$$Y_t := \sum_{i=1}^n 1_{[0, t]}(T_i), \quad t \geq 0,$$

representing the regime the process X is in at time t . The process Y is unobservable and is a Markov chain with state space $\mathcal{Y} := \{0, 1, 2, \dots, n\}$, initial state distribution η given by

$$\eta(y) = (1 - q_{y+1}) \prod_{1 \leq i \leq y} q_i, \quad y \in \mathcal{Y}$$

and one-step transition function P given by

$$P(y, y+k) = \begin{cases} 0, & \text{if } k < 0 \\ 1 - p_{y+1}, & \text{if } k = 0 \\ p_{y+1}(1 - q_{y+k+1}) \prod_{2 \leq i \leq k} q_{y+i}, & \text{if } 1 \leq k \leq n - y \end{cases}, \quad y \in \mathcal{Y},$$

where $q_{n+1} := p_{n+1} := 0$ and $\prod_{i \in \emptyset} \equiv 1$.

If we partition $\mathcal{Y}$ into the sets $\mathcal{Y}_0 := \{0, 1, 2, \dots, k-1\}$ and $\mathcal{Y}_1 := \mathcal{Y} \setminus \mathcal{Y}_0$, then this model fits the framework of the current paper, where the time of the k th change is given by

$$T := \min\{t \geq 0 \mid Y_t \notin \mathcal{Y}_0\} = T_k,$$

and index A equals 1 on $\{T < \infty\}$. Note that we can set

$$0 = c(0) = c(1) = \dots = c(k-1) < c(k) < \dots < c(n) < \infty$$

to incorporate an increasing delay penalty for each successive regime encountered after the k th change.

This multiple-change formulation is useful in a variety of applications. For example, suppose that the observation process X consists of measurements of specific protein levels present in a sequence of blood samples and that based on these samples we are to detect the onset of a certain epoch in the progression of a disease in order to initiate an appropriate treatment (or to discontinue treatment). As another example, suppose we monitor the output of a deteriorating system with three working states denoted by “good=0,” “fair=1,” and “poor=2,” and a failed state “fail=3.” In such a system, the tendency for the system to deteriorate can often accelerate as the state of the system worsens. That is, starting with a system in good condition, the durations S_1 , S_2 , and S_3 could be stochastically decreasing; i.e., $p_1 < p_2 < p_3$. In this case the increasing tendency of occurrence of additional changes beyond the first change can be a significant barrier to applying a “two-regime” model for detecting when the system first falls below good condition.

Detection of the onset of a gradual change. Suppose that when the disorder happens it does not completely manifest itself in the observation sequence but instead it gradually works its way into the system. This can be seen as a special case of the multiple change formulation above, where we are interested in the first change of many successive changes that occur quickly but with gradual effects. For instance, in the example of quickest detection of the k th out of n regime changes above, set p_1 to a very small value and $p_2, \dots, p_n$ to relatively large values and set $f(i-1, \cdot) \approx f(i, \cdot)$, $1 \leq i \leq n$, but with relatively dissimilar

$f(0, \cdot)$ and $f(n, \cdot)$, such that once the first change happens the second through n th changes occur quickly thereafter but with only gradual changes to the distribution of observations for each regime change.

Detection and diagnosis of drift and/or volatility changes in a random walk.

Consider a stochastic process $Z := \{Z_t; t \geq 0\}$ that satisfies the recursive relationship

$$Z_t = Z_{t-1} + \mu + \sigma W_t, \quad t \geq 1,$$

where the stochastic process $W = \{W_t; t \geq 1\}$ is an i.i.d. mean-zero variance-one white noise sequence with distribution $w(\cdot)$ independent of the random variable Z_0 . The process Z is called a random walk with drift μ and volatility σ . Suppose that the drift μ is initially equal to μ_0 and then at some unobservable time $T^{(\mu)}$ the drift changes to μ_1 . Similarly, suppose that the volatility σ is initially equal to σ_0 and then at some other unobservable time $T^{(\sigma)}$ it changes to σ_1 . That is,

$$Z_t - Z_{t-1} = (\mu_0 1_{\{t < T^{(\mu)}\}} + \mu_1 1_{\{t \geq T^{(\mu)}\}}) + (\sigma_0 1_{\{t < T^{(\sigma)}\}} + \sigma_1 1_{\{t \geq T^{(\sigma)}\}}) W_t, \quad t \geq 1.$$

Accordingly, conditionally on the disorder times $T^{(\mu)}$ and $T^{(\sigma)}$, the increments process $X = \{X_t; t \geq 1\}$, defined by $X_t := Z_t - Z_{t-1}$, $t \geq 1$, forms a sequence of independent random variables with conditional mean function $\mu(t) := \mathbb{E}\{X_t | T^{(\mu)}, T^{(\sigma)}\}$ and conditional variance function $\sigma^2(t) := \text{Var}\{X_t | T^{(\mu)}, T^{(\sigma)}\}$, where

$$(\mu(t), \sigma^2(t)) = \begin{cases} (\mu_0, \sigma_0^2) & \text{on } \{t < T^{(\mu)} \wedge T^{(\sigma)}\} \\ (\mu_1, \sigma_0^2) & \text{on } \{T^{(\mu)} \leq t < T^{(\sigma)}\} \\ (\mu_0, \sigma_1^2) & \text{on } \{T^{(\sigma)} \leq t < T^{(\mu)}\} \\ (\mu_1, \sigma_1^2) & \text{on } \{T^{(\mu)} \vee T^{(\sigma)} \leq t\} \end{cases}, \quad t \geq 1.$$

The problem is to detect a change immediately and determine its cause based only on sequential observations of the increments $Z_t - Z_{t-1} =: X_t$, $t \geq 1$.

The process X goes through three different regimes: in the first regime (possibly of zero duration) neither the drift nor volatility parameter has changed, in the second regime (possibly of zero duration) exactly one of either the drift or volatility parameter has changed, and in the final regime both parameters have changed. Let the durations of the first and second regimes be denoted by

$$S_1 := T^{(\mu)} \wedge T^{(\sigma)} \quad \text{and} \quad S_2 := (T^{(\mu)} \vee T^{(\sigma)}) - S_1,$$

respectively, and let the ordering of the parameter changes be captured by the index

$$A := 1_{\{T^{(\mu)} < T^{(\sigma)}\}} + 2 \cdot 1_{\{T^{(\mu)} > T^{(\sigma)}\}} + 3 \cdot 1_{\{T^{(\mu)} = T^{(\sigma)}\}},$$

where the events $\{A = 1\}$, $\{A = 2\}$, and $\{A = 3\}$ mean, respectively, that the drift parameter changes before the volatility parameter, the drift parameter changes after the volatility parameter, and both the drift and volatility parameters change simultaneously.

For example, suppose that given A the change times $T^{(\mu)}$ and $T^{(\sigma)}$ are conditionally independent zero-modified geometric random variables. Then, the durations S_1 and S_2 and the index A are *dependent* random variables such that

$$\mathbb{P}\{A = i\} = \alpha_i, \quad i = 1, 2, 3, \quad \text{where} \quad \alpha_1 + \alpha_2 + \alpha_3 = 1,$$

and conditionally on the index A , the durations S_1 and S_2 are independent with distributions

$$\begin{aligned} S_1 &\sim \text{zmgeom}(q_1, p_{11}) \quad \text{and} \quad S_2 \sim \text{geom}(p_{12}) && \text{on } \{A = 1\}, \\ S_1 &\sim \text{zmgeom}(q_2, p_{21}) \quad \text{and} \quad S_2 \sim \text{geom}(p_{22}) && \text{on } \{A = 2\}, \\ S_1 &\sim \text{zmgeom}(q_3, p_{31}) \quad \text{and} \quad S_2 \equiv 0 && \text{on } \{A = 3\}, \end{aligned}$$

where $\text{zmgeom}(q_i, p_{i1})$ and $\text{geom}(p_{i2})$ denote, respectively, the zero-modified geometric and geometric distribution, whose parameters are determined by the distributions of $T^{(\mu)}$ and $T^{(\sigma)}$. Define the stochastic process $Y = \{Y_t; t \geq 0\}$ by the sequence of random triples

$$Y_t := (1_{\{t \geq T^{(\mu)}\}}, 1_{\{t \geq T^{(\sigma)}\}}, A), \quad t \geq 0,$$

where the first component of each triple denotes the index (0 or 1) of the drift parameter, the second component corresponds to the index (0 or 1) of the volatility parameter, and the third component indicates whether the drift changes first (1), the volatility changes first (2), or both change at the same time (3). Then the process Y is a (hidden) Markov chain on the state space

$$\mathcal{Y} := \{(0, 0, 1), (0, 0, 2), (0, 0, 3), (1, 0, 1), (1, 1, 1), (0, 1, 2), (1, 1, 2), (1, 1, 3)\}$$

with initial distribution

$$\eta = (\alpha_1(1 - q_1), \alpha_2(1 - q_2), \alpha_3(1 - q_3), \alpha_1 q_1, 0, \alpha_2 q_2, 0, \alpha_3 q_3)$$

and one-step transition matrix

$$P = \begin{array}{c} \begin{array}{ccccccc} (0, 0, 1) & (0, 0, 2) & (0, 0, 3) & (1, 0, 1) & (1, 1, 1) & (0, 1, 2) & (1, 1, 2) & (1, 1, 3) \end{array} \\ \left[\begin{array}{cccccc|cc|cc} 1 - p_{11} & & & & & & & & & & & \\ & 1 - p_{21} & & & & & & & & & & \\ & & 1 - p_{31} & & & & & & & & & \\ \hline & & & 1 - p_{12} & p_{12} & & & & & & & \\ & & & & 1 & & & & & & & \\ \hline & & & & & 1 - p_{22} & p_{22} & & & & & \\ & & & & & & 1 & & & & & \\ \hline & & & & & & & & & & 1 & \end{array} \right] \begin{array}{l} (0, 0, 1) \\ (0, 0, 2) \\ (0, 0, 3) \\ (1, 0, 1) \\ (1, 1, 1) \\ (0, 1, 2) \\ (1, 1, 2) \\ (0, 1, 3) \end{array} \end{array}.$$

To fit the formulation to the aforementioned objective, let $\mathcal{Y}_0 := \{(0, 0, 1), (0, 0, 2), (0, 0, 3)\}$, $\mathcal{Y}_1 := \{(1, 0, 1), (1, 1, 1)\}$, $\mathcal{Y}_2 := \{(0, 1, 2), (1, 1, 2)\}$, and $\mathcal{Y}_3 := \{(1, 1, 3)\}$, and define

$$T := \min\{t \geq 0 \mid Y_t \notin \mathcal{Y}_0\} = T^{(\mu)} \wedge T^{(\sigma)}.$$

Then the relation $A = \arg\{i = 1, 2, 3; Y_T \in \mathcal{Y}_i\}$ holds on $\{T < \infty\}$, which completes the formulation.

Note that this formulation allows for an extended parameter structure such that the drift parameter value after the change can also depend on whether it changes first or second (and similarly for volatility). So, not only does this formulation handle multiple changes, but also it incorporates dependencies between the change time and change type at the same time. This generalization also carries over to the cost structure since the delay cost can depend on the order of the changes as well.

Many other structural change detection and diagnosis models can be formulated within this framework, such as shifts in the parameters governing an auto-regressive or moving-average time series model or regression models with unobservable structural breaks. See Hackl and Westlund [13] for an annotated bibliography of studies on statistical analysis and detection of structural changes.

Example 5 (Quickest detection of an unobservable change in the distribution of a Markov chain). Consider a finite-state Markov chain $M = \{M_t; t \geq 0\}$ whose initial distribution and one-step transition matrix change suddenly at some unobservable time T . Suppose that, conditionally on the disorder time T , the Markov chain M is time-homogeneous before time T with initial distribution μ and one-step transition matrix Q and is time-homogeneous thereafter with initial distribution ρ and one-step transition matrix R :

$$\overbrace{M_1, M_2, \dots, M_{T-1}, M_T}^{\text{Markov chain under } \mu, Q} \quad \text{and} \quad \overbrace{M_T, M_{T+1}, \dots}^{\text{Markov chain under } \rho, R}$$

For example, consider a communications channel where a digital signal is transmitted according to a Markov chain M and the problem is to detect quickly when the underlying one-step transition matrix has changed from Q to R based only on observations received from the transmission.

Yakir [29] gives an optimal decision rule for the corresponding quickest detection problem under Shiryaev's Bayes risk in (2.4) with a zero-modified geometric prior distribution for the change time. However, that result requires perfect state knowledge, and hence, it does not extend when there are imperfect observations of the Markov chain; i.e., noisy channels, which are present in virtually all real applications. Also, it does not extend when the disorder time is not zero-modified geometric nor when the one-step transition matrix can undergo multiple successive changes and/or have several alternative change types. In contrast, the framework of the current paper allows for each of these considerations. We illustrate below how to handle the situation of noisy observations.

Noisy channels. Consider a communications channel that carries bits (0 or 1) transmitted according to the Markov chain M on the state space $\{0, 1\}$ whose initial distribution μ and one-step transition matrix Q changes abruptly to ρ and R , respectively, at disorder time

T . Let the process $X = \{X_t; t \geq 0\}$ denote the sequence of observations received via this channel.

In a *binary symmetric channel* the bits are transmitted accurately with probability $1 - \beta$ and are otherwise flipped (i.e., 1 becomes 0 and 0 becomes 1) with probability β . That is,

$$\begin{aligned}\mathbb{P}\{X_t = 0 \mid M_t = 0\} &= \mathbb{P}\{X_t = 1 \mid M_t = 1\} = 1 - \beta \quad \text{and} \\ \mathbb{P}\{X_t = 1 \mid M_t = 0\} &= \mathbb{P}\{X_t = 0 \mid M_t = 1\} = \beta\end{aligned}$$

where β is the (symmetric) bit error rate.

Let us define the process Y by

$$Y_t := (M_t, 1_{[0,t]}(T)), \quad t \geq 0$$

on the state space $\mathcal{Y} := \{(0, 0), (1, 0), (0, 1), (1, 1)\}$ with the partition $\mathcal{Y}_0 := \{(0, 0), (1, 0)\}$ and $\mathcal{Y}_1 := \{(0, 1), (1, 1)\}$. Suppose the disorder time T has zero-modified geometric distribution with parameters θ_0 and θ , and

$$\begin{aligned}\mathbb{P}\{M_0 = m_0, M_1 = m_1, \dots, M_n = m_n, T = t\} \\ = \begin{cases} \theta_0 \rho(m_0) \prod_{1 \leq k \leq n} R(m_{k-1}, m_k), & t = 0 \\ (1 - \theta_0)(1 - \theta)^{t-1} \theta \mu(m_0) \prod_{1 \leq k \leq t-1} Q(m_{k-1}, m_k) \prod_{t \leq l \leq n} R(m_{l-1}, m_l), & t \geq 1 \end{cases}.\end{aligned}$$

It can be shown that the process Y is a Markov chain with initial distribution $\eta := ((1 - \theta_0)\mu, \theta_0\rho)$ and one-step transition matrix

$$P := \begin{bmatrix} \mathcal{Y}_0 & \mathcal{Y}_1 \\ (1 - \theta)Q & \theta R \\ 0 & R \end{bmatrix} \begin{matrix} \mathcal{Y}_0 \\ \mathcal{Y}_1 \end{matrix},$$

where Q and R are known transition matrices on the state space $\{0, 1\}$ of the Markov chain M . In terms of the hidden Markov chain Y the distribution of the modulated observation process X follows

$$f(y, x) = (1 - \beta) 1_{\{y(1)=x\}} + \beta 1_{\{y(1) \neq x\}}, \quad y \in \mathcal{Y}, \quad x \in \{0, 1\}.$$

Setting $c(y) = c 1_{\{y(2)=1\}}$ for some constant $c > 0$ and using the Bayes risk in (2.4) gives Shiryaev's quickest detection formulation for the *binary symmetric channel*. When the bit error rate β is zero the observation process X coincides with the underlying state of the Markov chain M and we can reduce the state-space of the model to recover the state-dependent solution presented by Yakir [29].

Note that with simple modifications of the conditional density functions of the process X , the bit error rate can be asymmetric in general and allowed to take different values before

and after the disorder time: for example,

$$f(y, x) = (1 - \beta_y) 1_{\{y(1)=x\}} + \beta_y 1_{\{y(1) \neq x\}}, \quad y \in \mathcal{Y}, \quad x \in \{0, 1\},$$

where $\beta_y \in [0, 1], y \in \mathcal{Y}$. A variety of other noisy channels can be accommodated. For example, suppose we have a *binary erasure channel* in which the bits are either received correctly or are converted to an erasure denoted by the symbol $\diamond$. Again, we need only to modify appropriately the conditional density functions of the observation process X on its expanded range $\{0, 1, \diamond\}$. See MacKay [22, Chapter 9] for background on communication over a noisy channel.

3. POSTERIOR ANALYSIS AND SOLUTION

For every $t \geq 0$, let $\Pi_t = (\Pi_t(y), y \in \mathcal{Y})$ be the row vector of the posterior probabilities

$$\Pi_t(y) := \mathbb{P}\{Y_t = y \mid \mathcal{F}_t\}, \quad y \in \mathcal{Y}$$

that the hidden Markov chain Y is in state $y \in \mathcal{Y}$ at time t given the history $\mathcal{F}_t$ of the observation process X . For every vector $g = (g(y), y \in \mathcal{Y})$, we shall denote by $\text{diag}(g)$ the diagonal matrix on $\mathcal{Y}$ whose y th diagonal entry is $g(y)$. For every $f : \mathcal{Y} \times \mathcal{X} \mapsto \mathbb{R}$, we define $\text{diag}(f)(x)$ at every $x \in \mathcal{X}$ by setting $g(y) = f(y, x), y \in \mathcal{Y}$ above.

Proposition 2. (i) If for every $t \geq 1$ we define $n_t : \mathcal{X}^t \times \mathcal{Y} \mapsto \mathbb{R}_+$ by

$$n_t(x_1, \dots, x_t, y) := \sum_{y_0, y_1, \dots, y_{t-1} \in \mathcal{Y}} \left(\eta(y_0) \prod_{k=1}^{t-1} P(y_{k-1}, y_k) f(y_k, x_k) \right) P(y_{t-1}, y) f(y, x_t),$$

then it is the joint probability density function of $X_1, \dots, X_t, Y_t$ on $(\mathcal{X}^t \times \mathcal{Y}, \mathcal{X}^t \otimes 2^{\mathcal{Y}})$ with respect to the product of t copies of the measure ν on $(\mathcal{X}, \mathcal{X})$ and the counting measure on $(\mathcal{Y}, 2^{\mathcal{Y}})$, and

$$(3.1) \quad \Pi_t(y) = \frac{n_t(X_1, \dots, X_t, y)}{\sum_{y' \in \mathcal{Y}} n_t(X_1, \dots, X_t, y')}, \quad y \in \mathcal{Y}, \quad t \geq 1.$$

(ii) For every nonnegative function $w : \mathcal{X} \mapsto \mathbb{R}_+$, we have

$$(3.2) \quad \mathbb{E}[w(X_{t+1}) \mid \mathcal{F}_t] = \int_{\mathcal{X}} w(x) \Pi_t P f(x) \nu(dx), \quad t \geq 0;$$

namely, the conditional probability density function of X_{t+1} given $\mathcal{F}_t = \sigma\{X_1, \dots, X_t\}$ with respect to the sigma-finite measure ν on $(\mathcal{X}, \mathcal{X})$ is

$$\Pi_t P f(x) = \sum_{y, y' \in \mathcal{Y}} \Pi_t(y) P(y, y') f(y', x), \quad x \in \mathcal{X}.$$

(iii) The process $\{\Pi_t, \mathcal{F}_t; t \geq 0\}$ is a Markov process on the state space $\mathcal{P} \triangleq \{\pi \in [0, 1]^{|\mathcal{Y}|}; \sum_{y \in \mathcal{Y}} \pi(y) = 1\}$, with dynamics

$$(3.3) \quad \Pi_{t+1} = \frac{\Pi_t P \text{diag}(f)(X_{t+1})}{\Pi_t P f(X_{t+1})}, \quad t \geq 0.$$

(iv) For every nonnegative function $w : \mathcal{P} \mapsto \mathbb{R}_+$, we have $\mathbb{E}[w(\Pi_{t+1}) \mid \mathcal{F}_t] = (\mathbb{T}w)(\Pi_t)$, $t \geq 0$, where

$$(3.4) \quad (\mathbb{T}w)(\pi) := \int_{\mathcal{X}} w \left(\frac{\pi P \text{diag}(f)(x)}{\pi P f(x)} \right) \pi P f(x) \nu(dx), \quad \pi \in \mathcal{P}.$$

(v) We have $\mathbb{E} \sum_{t \geq 0} \sum_{y \in \mathcal{Y}_0} \Pi_t(y) = \mathbb{E}T \leq \sum_{y, y' \in \mathcal{Y}_0} (I - P_0)^{-1}(y, y') < \infty$.

(vi) For every closed subset $\bar{\mathcal{Y}} \subset \mathcal{Y}$, the process $\{\sum_{y \in \bar{\mathcal{Y}}} \Pi_t(y), \mathcal{F}_t; t \geq 0\}$ is a submartingale.

(vii) The process $\{\sum_{y \in \mathcal{Y}_0} \Pi_t(y), \mathcal{F}_t; t \geq 0\}$ is a supermartingale.

The proofs are given in the appendix. The next proposition shows that the original problem reduces to an optimal stopping problem for the Markov process Π . It also identifies an optimal terminal decision rule $\delta(\cdot)$. It remains to calculate the minimum Bayes risk and find an optimal alarm time, if it exists.

Proposition 3. *If we define for every $\pi \in \mathcal{P}$*

$$(3.5) \quad g(\pi) := \sum_{y \in \mathcal{Y}} \pi(y)c(y) \equiv \pi c, \quad h(\pi) := \min_{1 \leq j \leq a} h(\pi, j) := \sum_{y \in \mathcal{Y}} \pi(y)C(y, j) \equiv (\pi C)(j),$$

then the Bayes risk in (2.3) can be expressed as

$$(3.6) \quad u(\tau, \delta) = \mathbb{E} \left[\sum_{t=0}^{\tau-1} g(\Pi_t) + 1_{\{\tau < \infty\}} h(\Pi_\tau, \delta) \right] \quad \text{for every } (\tau, \delta) \in \Delta.$$

If $\delta(t)$ is a $\{1, \dots, a\}$ -valued random variable such that

$$\delta(t) \in \arg \min \{h(\Pi_t, j); 1 \leq j \leq a\} \quad \text{for every } t \geq 0,$$

then $(\tau, \delta(\tau)) \in \Delta$ for every a.s. finite $\mathbb{F}$ -stopping time τ , and $u(\tau, \delta) \geq u(\tau, \delta(\tau))$ for every decision rule $(\tau, \delta) \in \Delta$. Therefore, the minimum Bayes risk is $u^ = \inf_{\tau \in \mathbb{F}, \tau < \infty \text{ a.s.}} u(\tau, \delta(\tau)) = v(\eta)$, where*

$$(3.7) \quad v(\pi) := \inf_{\substack{\tau \in \mathbb{F}, \\ \tau < \infty \text{ a.s.}}} \mathbb{E}_\pi \left[\sum_{t=0}^{\tau-1} g(\Pi_t) + 1_{\{\tau < \infty\}} h(\Pi_\tau) \right], \quad \pi \in \mathcal{P}$$

is the value function of an optimal stopping problem for the posterior probability process Π with running and terminal cost functions $g(\cdot)$ and $h(\cdot)$ in (3.5), respectively, and $\mathbb{E}_\pi$ is the expectation with respect to $\mathbb{P}$ given that $\Pi_0 = \pi$.

In the remainder of this section, we will try to solve the problem in (3.7) and identify an optimal stopping time if it exists. Let us start by defining

$$\begin{aligned}
 (3.8) \quad Z_t &:= -\sum_{n=0}^{t-1} g(\Pi_n) - h(\Pi_t), & 0 \leq t < \infty, \\
 C_n &:= \{\tau \in \mathbb{F} \mid n \leq \tau < \infty \text{ a.s., and } \mathbb{E}Z_\tau^- < \infty\}, & n \geq 0, \\
 C_n^N &:= \{\tau \wedge N \mid \tau \in C_n\}, & 0 \leq n \leq N, \ N \geq 0, \\
 \gamma_n &:= \operatorname{ess\,sup}_{\tau \in C_n} \mathbb{E}[Z_\tau \mid \mathcal{F}_n], & n \geq 0, \\
 \gamma_n^N &:= \operatorname{ess\,sup}_{\tau \in C_n^N} \mathbb{E}[Z_\tau \mid \mathcal{F}_n], & 0 \leq n \leq N, \ N \geq 0, \\
 -v_n &:= \sup_{\tau \in C_n} \mathbb{E}Z_\tau, & n \geq 0, \\
 -v_n^N &:= \sup_{\tau \in C_n^N} \mathbb{E}Z_\tau, & 0 \leq n \leq N, \ N \geq 0.
 \end{aligned}$$

Then for every a.s. finite $\mathbb{F}$ -stopping time τ , the expectation $\mathbb{E}Z_\tau = -\mathbb{E}Z_\tau^-$ exists, and

$$(3.9) \quad -u^* = \sup_{\substack{\tau \in \mathbb{F} \\ \tau < \infty \text{ a.s.}}} \mathbb{E}Z_\tau = \sup_{\tau \in C_0} \mathbb{E}Z_\tau \equiv -v_0.$$

The left-hand side of the second equality is greater than or equal to the right-hand side. For the reverse inequality, it is enough to note that for every a.s. finite $\tau \in \mathbb{F}$, the $\mathbb{F}$ -stopping time $\tilde{\tau} := \tau 1_{\{Z_0 \leq \mathbb{E}Z_\tau\}}$ belongs to C_0 because

$$-\mathbb{E}Z_{\tilde{\tau}}^- = \mathbb{E}Z_{\tilde{\tau}} = \max\{Z_0, \mathbb{E}Z_\tau\} \geq Z_0 = h(\Pi_0) > -\infty,$$

and $\mathbb{E}Z_\tau \leq \max\{Z_0, \mathbb{E}Z_\tau\} = \mathbb{E}Z_{\tilde{\tau}}$. Chow et al. [6, Theorems 3.2 and 4.1] show that

$$-v_n = \mathbb{E}\gamma_n, \ n \geq 0 \quad \text{and} \quad -v_n^N = \mathbb{E}\gamma_n^N, \ 0 \leq n \leq N, \ N \geq 0;$$

the processes $(\gamma_n)_{n \geq 0}$ and $(\gamma_n^N)_{0 \leq n \leq N}$ for every $N \geq 0$ satisfy

$$(3.10) \quad \begin{aligned} \gamma_n &= \max\{Z_n, \mathbb{E}[\gamma_{n+1} \mid \mathcal{F}_n]\}, \quad n \geq 0, \\ \gamma_N^N &= Z_N, \quad \gamma_n^N = \max\{Z_n, \mathbb{E}[\gamma_{n+1}^N \mid \mathcal{F}_n]\}, \quad 0 \leq n \leq N-1, \end{aligned}$$

respectively. Because $(C_n^N)_{N \geq n}$ is increasing, so is $(\gamma_n^N)_{N \geq 0}$. Therefore, $\lim_{N \rightarrow \infty} \gamma_n^N$ exists for every $n \geq 0$, and by the next proposition, it coincides with γ_n .

Proposition 4. *For every $n \geq 0$, we have $\gamma_n = \lim_{N \rightarrow \infty} \gamma_n^N$ almost surely.*

This proposition and the monotone convergence theorem imply that $v_0 = \lim_{N \rightarrow \infty} v_0^N$; namely, v_0^N is a “good” approximate solution of (3.9) for every large N . In the sequel, we derive a bound on the error $|v_0 - v_0^N|$ that gives a finite N that is large enough to attain any specified level of accuracy.

The next proposition shows that the successive approximations $(\gamma_n^N)_{0 \leq n \leq N}$, $N \geq 0$ of $(\gamma_n)_{n \geq 0}$ can be calculated by using an operator $\mathbb{M}$ acting on the bounded functions $w : \mathcal{P} \mapsto \mathbb{R}$ according to

$$(3.11) \quad (\mathbb{M}w)(\pi) := \min\{h(\pi), g(\pi) + (\mathbb{T}w)(\pi)\} \quad \text{for every } \pi \in \mathcal{P},$$

where $\mathbb{T}$ is given by (3.4). Because the function h is bounded, by (3.11) the function $\mathbb{M}w$ from $\mathcal{P}$ into $\mathbb{R}$ is also bounded, and $\mathbb{M}^2w := \mathbb{M}(\mathbb{M}w)$ is well-defined. Similarly, $\mathbb{M}^nw := \mathbb{M}^{n-1}(\mathbb{M}w)$ makes sense for every $n \geq 1$, where we define $\mathbb{M}^0w := w$.

Proposition 5. *We have $\gamma_n^N = -\sum_{k=0}^{n-1} g(\Pi_k) - (\mathbb{M}^{N-n}h)(\Pi_n)$ for every $0 \leq n \leq N$, $N \geq 0$.*

This proposition implies that the random variables

$$(3.12) \quad -v_0^N = \mathbb{E}\gamma_0^N = -(\mathbb{M}^N h)(\Pi_0), \quad N \geq 0 \quad \text{and} \quad -v_0 = \lim_{N \rightarrow \infty} \mathbb{E}\gamma_0^N = \lim_{N \rightarrow \infty} (\mathbb{M}^N h)(\Pi_0)$$

are in fact functions of the initial state Π_0 of the posterior probability process Π . Let $(\Omega, \mathcal{F}, \mathbb{P}_\pi)$ be a probability space on which $\Pi = \{\Pi_n; n \geq 0\}$ is by definition the Markov process with transition semigroup $\mathbb{T}$ in (3.4) and $\mathbb{P}_\pi\{\Pi_0 = \pi\} = 1$. Inspired by (3.12) let $v_n^N : \mathcal{P} \mapsto \mathbb{R}$, $0 \leq n \leq N$ for every $N \geq 0$ and $v_n : \mathcal{P} \mapsto \mathbb{R}$, $n \geq 0$ be the functions

$$-v_n^N(\pi) := \sup_{\tau \in C_n^N} \mathbb{E}_\pi Z_\tau \quad \text{and} \quad -v_n(\pi) := \sup_{\tau \in C_n} \mathbb{E}_\pi Z_\tau,$$

respectively. Then by Proposition 5 we have

$$(3.13) \quad v_0^N(\pi) = (\mathbb{M}^N h)(\pi) \quad \text{and} \quad v_0(\pi) = \lim_{N \rightarrow \infty} (\mathbb{M}^N h)(\pi) \quad \text{for every } \pi \in \mathcal{P}.$$

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 prove that $v_0(\cdot)$ satisfies the *optimality equation*

$$(3.14) \quad v_0(\pi) = (\mathbb{M}v_0)(\pi) = \min\{h(\pi), g(\pi) + (\mathbb{T}v_0)(\pi)\} \quad \text{for every } \pi \in \mathcal{P}.$$

Lemma 6. *If $w : \mathcal{P} \mapsto \mathbb{R}$ is a bounded concave function, then so is $(\mathbb{T}w)(\cdot)$.*

If $w : \mathcal{P} \mapsto \mathbb{R}$ is bounded and concave, then so is $\pi \mapsto (\mathbb{M}w)(\pi) = \min\{h(\pi), g(\pi) + (\mathbb{T}w)(\pi)\}$ because the mappings $h(\pi) = \min_{1 \leq j \leq a} (\pi C)(j)$, $g(\pi) = \pi c$, and $(\mathbb{T}w)(\pi)$ are concave by the previous lemma, and the pointwise minimum of concave functions is always concave. Therefore, repeated applications of the operator $\mathbb{M}$ to a bounded concave function also produce concave functions, and together with (3.13) this proves the next proposition.

Proposition 7. *The mappings $\pi \mapsto v_0^N(\pi)$, $N \geq 0$ and $\pi \mapsto v_0(\pi)$ from $\mathcal{P}$ into $\mathbb{R}$ are concave.*

In the remainder, we suppose that

$$(3.15) \quad \begin{aligned} c(y) &> 0 && \text{for every } y \in \mathcal{Y} \setminus \mathcal{Y}_0 \equiv \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_a, \\ C(y, j) &= 0 && \text{for every } y \in \mathcal{Y}_j, j = 1, \dots, a, \end{aligned}$$

and we define

$$c_p := \min\{c(y); y \in \mathcal{Y} \setminus \mathcal{Y}_0\} > 0,$$

which is positive by (3.15) and since the state-space $\mathcal{Y}$ of the hidden Markov chain Y is finite. The next theorem hints a useful numerical method that can calculate the value function $v_0(\cdot)$ at any desired level of accuracy:

For any given $\varepsilon > 0$ one can choose $N = N(\varepsilon)$ such that the rightmost bound in (3.16) is less than ε , and then calculate $v_0^N(\cdot)$ on $\mathcal{P}$ by applying N times the operator $\mathbb{M}$ to the function $h(\cdot)$. Then (3.16) guarantees that the difference between $v_0(\pi)$ and $v_0^N(\pi)$ is less than or equal to ε for every $\pi \in \mathcal{P}$.

Recall from (v) of Proposition 2 that the bound on the right-hand side of (3.16) is always finite. Here and elsewhere, $\|w\|$ denotes $\sup_{\pi} |w(\pi)|$ over the domain of the function $w(\cdot)$.

Theorem 8. *The sequence $(v_0^N(\pi))_{N \geq 0}$ converges to $v_0(\pi)$ as $N \rightarrow \infty$ uniformly in $\pi \in \mathcal{P}$. More precisely,*

$$(3.16) \quad v_0(\pi) \leq v_0^N(\pi) \leq v_0(\pi) + \frac{\|h\|}{N} \left(\frac{\|h\|}{c_p} + \sum_{y, y' \in \mathcal{Y}_0} (I - P_0)^{-1}(y, y') \right), \quad \pi \in \mathcal{P}, \quad N \geq 1.$$

The next two results are needed in showing the existence of an optimal alarm time. The continuity of the mapping $\pi \mapsto v_0^N(\pi)$ follows from the special form of the operator $\mathbb{M}$, whereas the continuity of $v_0(\cdot) = \lim_{N \rightarrow \infty} v_0^N(\cdot)$ is a corollary of Propositions 8 and 9.

Proposition 9. *For every $N \geq 0$, the function $v_0^N : \mathcal{P} \mapsto \mathbb{R}_+$ is continuous.*

Corollary 10. *The function $v_0 : \mathcal{P} \mapsto \mathbb{R}$ is continuous.*

Let us next show the existence of an optimal alarm time and describe its structure, as well as, its relation to successive approximations $(v_0^N)_{N \geq 0}$. We define

$$\begin{aligned} \Gamma_N &:= \{\pi \in \mathcal{P} \mid v_0^N(\pi) = h(\pi)\}, \quad \Gamma_N^{(j)} := \{\pi \in \mathcal{P} \mid v_0^N(\pi) = \bar{h}(\pi, j)\}, \quad 1 \leq j \leq a, \quad N \geq 0, \\ \Gamma &:= \{\pi \in \mathcal{P} \mid v_0(\pi) = h(\pi)\}, \quad \Gamma^{(j)} := \{\pi \in \mathcal{P} \mid v_0(\pi) = \bar{h}(\pi, j)\}, \quad 1 \leq j \leq a; \end{aligned}$$

in Theorem 13 we will show that the first time the posterior probability process Π enters the *stopping region* Γ gives an optimal alarm time. Therefore, the structure of the set Γ , identified carefully by the next proposition, will be useful to implement the optimal terminal decision rule and interpret the numerical results.

For example, Proposition 11 implies that upon entrance into the stopping region Γ , one finds herself in one of a subregions, $\Gamma^{(1)}, \dots, \Gamma^{(a)}$, associated respectively with closed subsets $\mathcal{Y}_1, \dots, \mathcal{Y}_a$ of the state-space $\mathcal{Y}$ of the hidden Markov chain Y ; moreover, upon stopping it is optimal to declare that Y must have made in the past a transition into the associated subset of its state-space. Because Γ is determined by $v_0(\cdot)$, which is obtained approximately with the sequence $(v_0^N)_{N \geq 0}$, the next proposition also provides the relation between stopping regions of the original and truncated problems.

Proposition 11. *For every $1 \leq j \leq a$, $(\Gamma_N^{(j)})_{N \geq 0}$ is a decreasing sequence of non-empty closed convex subsets of $\mathcal{P}$, and*

$$\begin{aligned} \Gamma_0^{(j)} &\supseteq \Gamma_1^{(j)} \supseteq \cdots \supseteq \Gamma^{(j)} \supseteq \{\pi \in \mathcal{P} \mid h(\pi, j) \leq \min\{h(\pi), g(\pi)\}\} \supseteq \{e_y \mid y \in \mathcal{Y}_j\}, \\ \Gamma &= \cap_{N=1}^{\infty} \Gamma_N = \cup_{j=1}^a \Gamma^{(j)}, \quad \Gamma_N = \cup_{j=1}^a \Gamma_N^{(j)}, \quad N \geq 0, \quad \Gamma^{(j)} = \cap_{N=1}^{\infty} \Gamma_N^{(j)}, \quad 1 \leq j \leq a. \end{aligned}$$

Therefore, $\mathcal{P} = \Gamma_0 \supseteq \Gamma_1 \supseteq \cdots \supseteq \Gamma \supsetneq \{e_y \mid y \in \mathcal{Y}_1 \cup \cdots \cup \mathcal{Y}_N\}$, and every $\Gamma^{(j)}$, $1 \leq j \leq a$ is a nonempty closed convex subset of $\mathcal{P}$.

The next lemma follows immediately from Proposition 5 and (3.13).

Lemma 12. *For every $n \geq 0$ and $N \geq n$, we have $\gamma_n^N = -\sum_{k=0}^{n-1} g(\Pi_k) - v_0^{N-n}(\Pi_n)$ and $\gamma_n = -\sum_{k=0}^{n-1} g(\Pi_k) - v_0(\Pi_n)$.*

Theorem 13. *The $\mathbb{F}$ -stopping time $\sigma := \inf\{t \geq 0 \mid \Pi_t \in \Gamma\}$ solves optimally the problem $-v_0(\pi) = \sup_{\tau \in \mathbb{F}} \mathbb{E}_{\pi} Z_{\tau}$ for every $\pi \in \mathcal{P}$. The stopped process $\{\gamma_{\sigma \wedge n}, \mathcal{F}_n; n \geq 0\}$ is a martingale, and $\mathbb{E}_{\pi} \sigma < \infty$ for every $\pi \in \mathcal{P}$.*

Relaxation of the positive-cost assumption. Let us now remove the condition “ $c(y) > 0$ for every $y \in \mathcal{Y} \setminus \mathcal{Y}_0$ ” in (3.15) and assume that $c(y) \geq 0$ for every $y \in \mathcal{Y}$ instead. In this generality, the stopping time σ of Theorem 13 is not always optimal. For example, if $c(\cdot) \equiv 0$, and hidden states of the Markov chain Y are identifiable from the observations, then indefinitely long sampling does not cost and drives the terminal decision cost to zero; therefore, there is no a.s. finite stopping rule in this case. However, the value function $v_0(\cdot)$ can always be calculated as the next theorem describes, and an a.s. finite nearly optimal stopping rule can always be found. Let us define

$$c_{\varepsilon} := c(y) + \varepsilon \quad \text{for every } y \in \mathcal{Y} \text{ and } \varepsilon > 0,$$

and denote by the $Z_n\langle\varepsilon\rangle$, $C_n\langle\varepsilon\rangle$, $C_n^N\langle\varepsilon\rangle$, $\gamma_n\langle\varepsilon\rangle$, $\gamma_n^N\langle\varepsilon\rangle$, $v_n\langle\varepsilon\rangle$, $v_n^N\langle\varepsilon\rangle$ the variables defined in (3.8) after replacing $c(\cdot)$ with $c_{\varepsilon}(\cdot)$. Then $Z_n\langle\varepsilon\rangle = Z_n - \varepsilon n$ for every $n \geq 0$ and $\varepsilon > 0$, and because $c_{\varepsilon}(y) > 0$ for every $y \in \mathcal{Y} \setminus \mathcal{Y}_0$, all of the previous results apply to the solution of the optimal stopping problem

$$(3.17) \quad -v_0\langle\varepsilon\rangle = \sup_{\tau \in C_0\langle\varepsilon\rangle} \mathbb{E} Z_{\tau}\langle\varepsilon\rangle \quad \text{for every } \varepsilon > 0.$$

Theorem 14. *For every $n \geq 0$, we have $\gamma_n = \lim_{\varepsilon \searrow 0} \uparrow \gamma_n\langle\varepsilon\rangle$ almost surely, and $v_n = \lim_{\varepsilon \searrow 0} \uparrow v_n\langle\varepsilon\rangle$.*

As a consequence, if $\varepsilon > 0$ is sufficiently small, then the stopping time $\sigma(\varepsilon)$ defined as in Theorem 13 for (3.17) will be nearly optimal for the original problem.

4. NUMERICAL ILLUSTRATIONS

In this section we extend the examples described in Section 2 by demonstrating numerically and graphically the optimal solution developed in Section 3 for some particular parameter values. Following the construction given in Dayanik et al. [9, Sections 5.3-5.4] we give graphical representations of the simplex $\mathcal{P}$ and its resident features such as sample paths of Π and optimal stopping region Γ for models with state space of size $|\mathcal{Y}| = 3$ or $|\mathcal{Y}| = 4$.

4.1. Second-order phase-type disorder time with two alternatives. Here we revisit the basic independence/dependence examples that were introduced and compared in Remark 1 of Section 2. First we consider two simple cases in which the disorder time T and its cause A are *independent*. In the first independence case, T has a zero-modified geometric prior distribution with parameters $q_1 := 0.04$ and $p_1 := 0.05$, A has prior distribution $\alpha := (0.5, 0.5)$, and A and T are independent. Plugging these values into the corresponding formulation from Remark 1 we have initial distribution $\eta := (0.96, 0.02, 0.02)$ and one-step transition matrix

$$P := \begin{bmatrix} 0.95 & 0.025 & 0.025 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

for the Markov chain Y on state space $\mathcal{Y} := \{0, 1, 2\}$ with partition $\mathcal{Y}_i := \{i\}, i = 0, 1, 2$. To complete the formulation we set the cost parameters and observation probability laws:

$$c := \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}, \quad C := \begin{bmatrix} 20 & 20 \\ 0 & 10 \\ 10 & 0 \end{bmatrix}, \quad f := \begin{bmatrix} 0.25 & 0.25 & 0.25 & 0.25 \\ 0.40 & 0.30 & 0.20 & 0.10 \\ 0.10 & 0.20 & 0.30 & 0.40 \end{bmatrix}.$$

In words, the cost structure consists of a 1-unit penalty for each period of alarm delay, a 20-unit penalty for a false-alarm (of either type), and a 10-unit mis-diagnosis penalty for an incorrect terminal decision. The observations can take one of four distinct values each equally likely before the change occurs and then after the change the probabilities are skewed either right or left determined by $A = 1$ or $A = 2$, respectively.

The optimal stopping region for this first case is displayed graphically in Figure 3(a). The triangle represents a linear mapping from $\mathbb{R}^3$ into $\mathbb{R}^2$ of the simplex $\mathcal{P} = \{\pi = (\pi(0), \pi(1), \pi(2)) \in [0, 1]^3 \mid \pi(0) + \pi(1) + \pi(2) = 1\}$ with extreme points $e_0 = (1, 0, 0)$, $e_1 = (0, 1, 0)$, and $e_2 = (0, 0, 1)$ forming its corners, where each point π in the triangle represents a convex combination of the corners e_0 , e_1 , and e_2 with weights proportional to π 's (Euclidean) distances to the triangle edge opposite of each respective corner. On this figure, the shaded regions comprise the stopping region. A sample path of $(\Pi_t)_{t=0}^\sigma$ is shown for a particular realization of Y and X , which shows how the optimal sequential decision strategy can be implemented by tracking the path of Π with this representation and raising an alarm as soon as it enters the stopping region for the first time. In this instance its

entrance is into the subregion labeled by $\Gamma^{(2)}$, so the optimal terminal diagnosis decision is that the second alternative has happened: $\delta = 2$.

FIGURE 3. Graphical comparison of the optimal stopping regions $\Gamma = \Gamma^{(1)} \cup \Gamma^{(2)}$ when A and T are independent ((a), (b)) and when they are dependent (c). Sample paths of $(\Pi_t)_{t=0}^\sigma$ are shown for which (a) $T = 6, A = 2$, (b) $T = 13, A = 2$, and (c) $T = 8, A = 2$.

In the second independence case, we keep all of the above parameters the same except for T which instead has a zero-modified geometric prior distribution with parameters $q_2 := 0.04$ and $p_2 := 0.15$. The optimal stopping region for this second case is displayed graphically in Figure 3(b). This figure is to be interpreted in the same way as Figure 3(a).

Next we consider a related dependence case such that T is *conditionally* zero-modified geometric with parameters q_i and p_i given $A = i$, $i = 1, 2$, where $q_1 := 0.04$, $p_1 := 0.05$ and $q_2 := 0.04$, $p_2 := 0.15$ are as above. Again A has distribution $\alpha := (0.5, 0.5)$, but here it is not independent of T . Plugging these values into the *dependence* formulation from Remark

1 we have initial distribution $\eta := (0.48, 0.48, 0.02, 0.02)$ and one-step transition matrix

$$P := \begin{bmatrix} 0.95 & 0 & 0.05 & 0 \\ 0 & 0.85 & 0 & 0.15 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

for the Markov chain Y on state space $\mathcal{Y} := \{(1, 0), (2, 0), (1, 1), (2, 1)\}$ with partition $\mathcal{Y}_0 := \{(1, 0), (2, 0)\}$, $\mathcal{Y}_i := \{(i, 1)\}$, $i = 1, 2$. To complete the formulation we set the cost parameters and observation probability laws in correspondence with the preceding independence examples:

$$c := \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}, \quad C := \begin{bmatrix} 20 & 20 \\ 20 & 20 \\ 0 & 10 \\ 10 & 0 \end{bmatrix}, \quad f := \begin{bmatrix} 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.40 & 0.30 & 0.20 & 0.10 \\ 0.10 & 0.20 & 0.30 & 0.40 \end{bmatrix}.$$

The optimal stopping region for this dependence case is displayed graphically in Figure 3(c). The tetrahedron depicted in this figure represents a linear mapping from $\mathbb{R}^4$ into $\mathbb{R}^3$ of the simplex $\mathcal{P} = \{\pi = (\pi(1, 0), \pi(2, 0), \pi(1, 1), \pi(2, 1)) \in [0, 1]^4 \mid \pi(1, 0) + \pi(2, 0) + \pi(1, 1) + \pi(2, 1) = 1\}$ with extreme points $e_{(1,0)} = (1, 0, 0, 0)$, $e_{(2,0)} = (0, 1, 0, 0)$, $e_{(1,1)} = (0, 0, 1, 0)$, and $e_{(2,1)} = (0, 0, 0, 1)$ forming its corners. Again, the shaded regions comprise the stopping region and a sample path of Π is shown for a particular realization of Y and X .

4.2. Two-sub-period cyclic model with two alternatives. Here we revisit the cyclic model introduced in Example 3. In that example we assumed an ultimate i.i.d. sequence upon disorder; i.e., that the cyclic nature of observations ceased. This was not required, but was used only for clarity of exposition in that discussion. Here we examine the natural case where cycles continue after the change with different laws. That is, suppose that we collect observations during alternating peak and off-peak sub-periods, and that at some unobservable disorder time T the distribution of observations changes to one of two alternative regimes, but the cyclic peak/off-peak behavior persists. For example, take each period to be a day and take daytime and nighttime as its sub-periods, where peak activity occurs during the daytime.

In particular, we formulate the problem via a Markov chain Y with state space $\mathcal{Y} := \{(0, 0), (1, 0), (0, 1), (1, 1), (0, 2), (1, 2)\}$, initial distribution $\eta := (0, 0.99, 0, 0.005, 0, 0.005)$,

and one-step transition matrix

$$P := \begin{array}{c} \begin{array}{cc|cc|cc} (0,0) & (1,0) & (0,1) & (1,1) & (0,2) & (1,2) \\ \hline 0 & 0.98 & 0 & 0.01 & 0 & 0.01 \\ 0.88 & 0 & 0.06 & 0 & 0.06 & 0 \\ \hline & & 0 & 1 & & \\ & & 1 & 0 & & \\ \hline & & & & 0 & 1 \\ & & & & 1 & 0 \end{array} \begin{array}{l} (0,0) \\ (1,0) \\ (0,1) \\ (1,1) \\ (0,2) \\ (1,2) \end{array} \end{array}$$

for the partition $\mathcal{Y}_i := \{(0, i), (1, i)\}$, $i = 0, 1, 2$, where first component of each state denotes the time of day (0 for daytime and 1 for nighttime) and the second component of each state indicates the regime (0 for the initial regime and 1, 2 for the alternative regimes). We set the cost structure to reflect our belief that delays during peak sub-periods (e.g., daytime) are more costly than in off-peak sub-periods (e.g., nighttime), together with asymmetric penalties for misdiagnosis and observation probability laws given by

$$c := \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0.5 \\ 1 \\ 0.5 \end{bmatrix}, \quad C := \begin{bmatrix} 20 & 20 \\ 20 & 20 \\ 0 & 3 \\ 0 & 3 \\ 5 & 0 \\ 5 & 0 \end{bmatrix}, \quad f := \begin{bmatrix} 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.40 & 0.30 & 0.20 & 0.10 \\ 0.40 & 0.30 & 0.20 & 0.10 \\ 0.10 & 0.20 & 0.30 & 0.40 \\ 0.10 & 0.20 & 0.30 & 0.40 \end{bmatrix}.$$

In words, the cost structure consists of a 1-unit penalty for each peak sub-period of alarm delay, a 0.5-unit penalty for each off-peak sub-period of alarm delay, a 20-unit penalty for a false-alarm (of either type), and 3-unit and 5-unit mis-diagnosis penalties for incorrect terminal decisions, respectively. The observations can take one of four distinct values each equally likely before the change occurs and then after the change the probabilities are skewed either right or left determined by $A = 1$ or $A = 2$, respectively.

By construction, the underlying Markov chain Y cycles through peak/off-peak sub-periods (beginning at off-peak): $Y_t \in \{(0, i) \mid i = 1, 2, 3\}$ if t is odd (peak), $Y_t \in \{(1, i) \mid i = 1, 2, 3\}$ if t is even (off-peak). This special structure yields

$$\sum_{y \in \mathcal{Y}: y(1)=0} \Pi_t(y) = 1 \quad \text{or} \quad \sum_{y \in \mathcal{Y}: y(1)=1} \Pi_t(y) = 1, \quad t \geq 0.$$

Hence, the additional knowledge of the sub-period Y reduces the dimension of the state space and we can consider a different stopping rule for peak and off-peak observation sub-periods instead. Figure 4(a) and Figure 4(b) represent the optimal stopping regions for observations during peak and off-peak sub-periods, respectively. The light-shaded regions of Figure 4(a) correspond to the stopping region for peak observations (e.g., daytime measurements), while the dark-shaded regions of Figure 4(b) correspond to the off-peak stopping region (e.g.,

nighttime measurements). On each of these figures is the same sample path with alternating dark and light shaded points representing whether that value is updated in an off-peak or peak sub-period, respectively. Whenever a peak point enters a peak stopping region (light, daytime) or an off-peak point enters an off-peak stopping region (dark, nighttime), it is optimal to raise an alarm. This joint strategy is depicted in Figure 4(c) by merging the two scenarios onto a single plot.

Figure 5 depicts similar scenarios with all parameters as above except with the following changes to the off-peak misdiagnosis costs: $C((1,0),1) = C((1,0),2) = 17$, $C((1,1),2) = 5$, and $C((1,2),1) = 3$.

FIGURE 4. Optimal stopping region $\Gamma = \Gamma^{(1)} \cup \Gamma^{(2)}$ for a cyclic model with day/night cycles corresponding to (a) daytime measurements, (b) nighttime measurements, and (c) composite stopping region formed by combining sub-figures (a) and (b). A sample path of $(\Pi_t)_{t=0}^\sigma$ is shown for which $T = 11$, $A = 1$.

FIGURE 5. Similar to Figure 4 but with different nighttime misdiagnosis penalties. A sample of $(\Pi_t)_{t=0}^\sigma$ is shown for which $T = 9$, $A = 1$.

4.3. **Quickest detection of one of two successive regime changes.** Consider the following three-regime model. Let $\eta := (0.96, 0.03, 0.01)$ and

$$P := \begin{bmatrix} 0.8 & 0.1 & 0.1 \\ 0 & 0.9 & 0.1 \\ 0 & 0 & 1 \end{bmatrix}, \quad c := \begin{bmatrix} 0 \\ 1 \\ 3 \end{bmatrix}, \quad C := \begin{bmatrix} 20 \\ 0 \\ 0 \end{bmatrix}, \quad f := \begin{bmatrix} 0.75 & 0.25 \\ 0.5 & 0.5 \\ 0.25 & 0.75 \end{bmatrix}.$$

For $\mathcal{Y} := \{0, 1, 2\}$, $\mathcal{Y}_0 := \{0\}$, and $\mathcal{Y}_1 := \{1, 2\}$, we are interested in detecting the first of the two regime changes. This formulation not only penalizes for detection delay, but also incorporates a greater delay penalty for each period of alarm delay after the *second* change. The optimal stopping region for this problem is depicted in Figure 6(a).

Now consider a second related model with all parameters the same except for

$$c := \begin{bmatrix} 0 \\ 0 \\ 3 \end{bmatrix} \quad \text{and} \quad C := \begin{bmatrix} 20 \\ 10 \\ 0 \end{bmatrix}.$$

For $\mathcal{Y} := \{0, 1, 2\}$, $\mathcal{Y}_0 := \{0, 1\}$, and $\mathcal{Y}_1 := \{2\}$, we are interested in detecting the second of the two regime changes. This formulation not only penalizes for false-alarm, but also penalizes more substantially if the *first* regime change has not yet happened upon alarm. The optimal stopping region for this problem is depicted in Figure 6(b).

FIGURE 6. Optimal stopping region for the quickest detection of (a) the first of two successive regime changes and (b) the last of two successive regime changes. A sample path is shown on each sub-figure for which (a) $T = 11$ and (b) $T = 10$.

4.4. Random walk simple disorder model. Here we continue with the random walk disorder model introduced in Example 4. We assume that the white noise sequence has the standard Normal distribution, and we consider the case where $\mu_0 := 0$, $\mu_1 := 1$, $\sigma_0^2 := 4$, and $\sigma_1^2 := 9$. Hence, the increments process X has initially the $N(0, 4)$ (Normal mean-zero, variance-four) distribution, then the distribution changes to either $N(1, 4)$ or $N(0, 9)$ for some random duration, and then finally it changes to $N(1, 9)$. In terms of state space of the model in Example 4, we have

$$f(y, x) = \left\{ \begin{array}{ll} (8\pi)^{-1/2} e^{-x^2/8}, & y \in \{(0, 0, 1), (0, 0, 2), (0, 0, 3)\} \\ (8\pi)^{-1/2} e^{-(x-1)^2/8}, & y = (1, 0, 1) \\ (18\pi)^{-1/2} e^{-x^2/18}, & y = (0, 1, 2) \\ (18\pi)^{-1/2} e^{-(x-1)^2/18}, & y \in \{(1, 1, 1), (1, 1, 2), (1, 1, 3)\} \end{array} \right\}.$$

Now, suppose that $T^{(\mu)}$ and $T^{(\sigma)}$ are i.i.d. $\text{geom}(0.1)$. Then, it can be shown that $S_1 \sim \text{geom}(0.19)$, $S_2 \sim \text{geom}(0.1)$, A has distribution $\alpha_1 = \alpha_2 = 0.09/0.19$ and $\alpha_3 = 0.01/0.19$, and that A and T are independent. In terms of the formulation and parameters of Example 4 we have $p_{11} = p_{21} = p_{31} = 0.19$, $p_{12} = p_{22} = 0.1$, and $q_1 = q_2 = q_3 = 0$. So, this completely specifies the initial distribution η and one-step transition matrix P given in that example.

To complete the formulation, let us set the delay costs to

$$c(y) = \left\{ \begin{array}{ll} 0, & y \in \{(0, 0, 1), (0, 0, 2), (0, 0, 3)\} \\ 1, & y = (1, 0, 1) \\ 1, & y = (0, 1, 2) \\ 2, & y \in \{(1, 1, 1), (1, 1, 2), (1, 1, 3)\} \end{array} \right\}$$

and the rows of the misdiagnosis cost matrix C to

$$C(y, \cdot) = \left\{ \begin{array}{ll} (20, 20, 20), & y \in \{(0, 0, 1), (0, 0, 2), (0, 0, 3)\} \\ (0, 10, 15), & y = (1, 0, 1) \\ (10, 0, 15), & y = (0, 1, 2) \\ (15, 15, 0), & y \in \{(1, 1, 1), (1, 1, 2), (1, 1, 3)\} \end{array} \right\}.$$

In words, we have a 1-unit penalty for each period of delay beyond the first change (of either type), a 2-unit penalty for each period beyond the second change. A false alarm costs 20, a misdiagnosis of the second regime costs 10, and a misdiagnosis of the final regime costs 15.

For this special scenario we can reduce the dimension of the state space from eight states to four states as follows. Let $(0, 0)$ represent the collection $\{(0, 0, 1), (0, 0, 2), (0, 0, 3)\}$ and let $(1, 1)$ represent the collection $\{(1, 1, 1), (1, 1, 2), (1, 1, 3)\}$. Also, let $(1, 0)$ denote $(1, 0, 1)$ and $(0, 1)$ denote $(0, 1, 2)$. Each element y in the collection $(0, 0)$ has the same density function $f(y, x)$ and the same cost functions $c(y) = 0$ and $C(y, j) = 20, 1 \leq j \leq 3$. Similarly, each element y in the collection $(1, 1)$ has the same density function $f(y, x)$ and the same cost

functions $c(y) = 2$ and $C(y, \cdot) = (15, 15, 0)$. Moreover, the transition probabilities in and out of each collection of states $(0, 0)$ and $(1, 1)$ simplify to

$$P := \begin{array}{cccc} & (0, 0) & (1, 0) & (0, 1) & (1, 1) \\ \begin{bmatrix} 0.81 & 0.09 & 0.09 & 0.01 \\ & 0.9 & & 0.1 \\ & & 0.9 & 0.1 \\ & & & 1 \end{bmatrix} & (0, 0) \\ & (1, 0) \\ & (0, 1) \\ & (1, 1) \end{array}.$$

Since the running cost function $g(\pi)$ and terminal cost functions $\bar{h}(\pi, j)$, $1 \leq j \leq 3$ are each linear in π , the optimality equation remains valid if we take the collections $(0, 0)$ and $(1, 1)$ as individual states since

$$\begin{aligned} \Pi_t((0, 0)) &= \Pi_t((0, 0, 1)) + \Pi_t((0, 0, 2)) + \Pi_t((0, 0, 3)) \\ \text{and } \Pi_t((1, 1)) &= \Pi_t((1, 1, 1)) + \Pi_t((1, 1, 2)) + \Pi_t((1, 1, 3)), \quad t \geq 0. \end{aligned}$$

That is, we can take the Markov chain Y on the state space $\mathcal{Y} := \{(0, 0), (1, 0), (0, 1), (1, 1)\}$ with transition matrix P given above, initial distribution $\eta := (1, 0, 0, 0)$, and partition $\mathcal{Y}_0 := \{(0, 0)\}$, $\mathcal{Y}_1 := \{(1, 0)\}$, $\mathcal{Y}_2 := \{(0, 1)\}$, and $\mathcal{Y}_3 := \{(1, 1)\}$. The corresponding optimal decision region is depicted in Figure 7(a).

4.5. Finite-state Markov channel disorder problem. Continuing Example 5 we solve the problem of quickest detection of an unobservable change in the one-step transition matrix of a binary Markov chain with initial distribution $(0.5, 0.5)$ from Q to R , where

$$Q := \begin{bmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{bmatrix} \quad \text{and} \quad R := \begin{bmatrix} 0.6 & 0.4 \\ 0.4 & 0.6 \end{bmatrix},$$

with a geometric ($\theta := 0.1$) prior change time distribution. This corresponds to a scenario where the successive zero and one bits are independent before the change but have a slight tendency of “streaks” after the change. That is, after the change, each bit is 50% more likely to be followed by a bit of the same type than by the other bit type. So let $\mathcal{Y} := \{(0, 0), (1, 0), (0, 1), (1, 1)\}$, $\eta := (0.5, 0.5, 0, 0)$, and

$$P := \begin{bmatrix} 0.45 & 0.45 & 0.06 & 0.04 \\ 0.45 & 0.45 & 0.04 & 0.06 \\ 0 & 0 & 0.6 & 0.4 \\ 0 & 0 & 0.4 & 0.6 \end{bmatrix},$$

$$c := \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}, \quad C := \begin{bmatrix} 20 \\ 20 \\ 0 \\ 0 \end{bmatrix}, \quad f := \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \\ 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix}.$$

Hence, this is a simple binary symmetric channel with $\beta := 0.1$ both before and after the change. See Figure 7(b) for a depiction of the optimal stopping region for this problem and a sample path for a particular realization of Y and X . The edge of the tetrahedron from $e_{(0,0)}$ to $e_{(0,1)}$ corresponds to transmission of a zero bit, while the edge from $e_{(1,0)}$ to $e_{(1,1)}$ corresponds to transmission of a one bit. Hence, as each observation is collected, if it is a zero then the sample path tends to “walk” along $e_{(0,0)}$ to $e_{(0,1)}$ as the probability that the change has happened increases with time. The path moves along the edge but not exactly on the edge due to the noise in each observation. Similarly, if the observation is a one, then the sample path tends to “walk” along $e_{(1,0)}$ to $e_{(1,1)}$. This behavior is exhibited in the sample path of Figure 7(b) as the first three observations are zeroes while the fourth and fifth are ones and the last five observations are all zeroes. This final “streak” of zeroes is significant enough in this instance to signal an alarm.

FIGURE 7. Optimal stopping region for (a) a random walk simple disorder model and (b) the quickest detection of an unobservable change in the transition matrix of a noisy Markov channel. A sample path is shown on each sub-figure for which (a) $T = 2$, $A = 1$ and (b) $T = 5$.

ACKNOWLEDGMENT

The research of Savas Dayanik was supported by the Air Force Office of Scientific Research, under grant AFOSR-FA9550-06-1-0496.

APPENDIX A. PROOFS OF SELECTED RESULTS

Proof of Proposition 2. For every $\mathcal{F} = \{(X_1, \dots, X_t) \in \mathcal{B}\} \in \mathcal{F}_t$, because

$$\begin{aligned} \mathbb{P}(\mathcal{F}) &= \sum_{y_0, y_1, \dots, y_t \in \mathcal{Y}} \mathbb{P}\{Y_0 = y_0, Y_1 = y_1, \dots, Y_t = y_t, (X_1, \dots, X_t) \in \mathcal{B}\} \\ &= \int \cdots \int_{\mathcal{B}} \sum_{y_0, y_1, \dots, y_t \in \mathcal{Y}} \left(\eta(y_0) \prod_{k=1}^t P(y_{k-1}, y_k) f(y_k, x_k) \right) \nu(dx_1) \cdots \nu(dx_t), \end{aligned}$$

we have for every $y \in \mathcal{Y}$ that

$$\begin{aligned} \int_{\mathcal{F}} \Pi_t(y) d\mathbb{P} &= \mathbb{P}\{Y_t = y, (X_1, \dots, X_t) \in \mathcal{B}\} \\ &= \sum_{y_0, y_1, \dots, y_{t-1} \in \mathcal{Y}} \mathbb{P}\{Y_0 = y_0, Y_1 = y_1, \dots, Y_{t-1} = y_{t-1}, Y_t = y, (X_1, \dots, X_t) \in \mathcal{B}\} \\ &= \int \cdots \int_{\mathcal{B}} \sum_{y_0, y_1, \dots, y_{t-1} \in \mathcal{Y}} \left(\eta(y_0) \prod_{k=1}^{t-1} P(y_{k-1}, y_k) f(y_k, x_k) \right) P(y_{t-1}, y) f(y, x_t) \nu(dx_1) \cdots \nu(dx_t), \\ &= \int \cdots \int_{\mathcal{B}} \frac{\sum_{y_0, y_1, \dots, y_{t-1} \in \mathcal{Y}} \left(\eta(y_0) \prod_{k=1}^{t-1} P(y_{k-1}, y_k) f(y_k, x_k) \right) P(y_{t-1}, y) f(y, x_t)}{\sum_{y_0, y_1, \dots, y_t \in \mathcal{Y}} \eta(y_0) \prod_{k=1}^t P(y_{k-1}, y_k) f(y_k, x_k)} \\ &\quad \sum_{y_0, y_1, \dots, y_t \in \mathcal{Y}} \left(\eta(y_0) \prod_{k=1}^t P(y_{k-1}, y_k) f(y_k, x_k) \right) \nu(dx_1) \cdots \nu(dx_t) \\ &= \int_{\mathcal{F}} \frac{n_t(X_1, \dots, X_t, y)}{\sum_{y_t \in \mathcal{Y}} n_t(X_1, \dots, X_t, y_t)} d\mathbb{P}, \end{aligned}$$

which completes the proof of (3.1). By using (3.1), we can write $\Pi_{t+1}(y)$ as

$$\frac{n_{t+1}(X_1, \dots, X_{t+1}, y)}{\sum_{y' \in \mathcal{Y}} n_{t+1}(X_1, \dots, X_{t+1}, y')} = \frac{\sum_{y' \in \mathcal{Y}} n_t(X_1, \dots, X_t, y') P(y', y) f(y, X_{t+1})}{\sum_{y' \in \mathcal{Y}} \sum_{y'' \in \mathcal{Y}} n_t(X_1, \dots, X_t, y'') P(y'', y') f(y', X_{t+1})},$$

and dividing both numerator and denominator by $\sum_{y''' \in \mathcal{Y}} n_t(X_1, \dots, X_t, y''')$ gives

$$\Pi_{t+1}(y) = \frac{\sum_{y' \in \mathcal{Y}} \Pi_t(y') P(y', y) f(y, X_{t+1})}{\sum_{y' \in \mathcal{Y}} \sum_{y'' \in \mathcal{Y}} \Pi_t(y'') P(y'', y') f(y', X_{t+1})} = \frac{(\Pi_t P)(y) f(y, X_{t+1})}{\sum_{y' \in \mathcal{Y}} (\Pi_t P)(y') f(y', X_{t+1})}, \forall y \in \mathcal{Y},$$

which is the same as (3.3). On the other hand,

$$\begin{aligned}
& \int_{\mathcal{F}} \mathbb{E}[w(X_{t+1}) \mid \mathcal{F}_t] d\mathbb{P} = \int_{(X_1, \dots, X_t) \in \mathcal{B}} w(X_{t+1}) d\mathbb{P} \\
&= \int_{\mathcal{B} \times \mathcal{X}} w(x_{t+1}) \sum_{y_0, y_1, \dots, y_{t+1} \in \mathcal{Y}} \left(\eta(y_0) \prod_{k=1}^{t+1} P(y_{k-1}, y_k) f(y_k, x_k) \right) \nu(dx_1) \cdots \nu(dx_{t+1}) \\
&= \int_{\mathcal{B} \times \mathcal{X}} w(x_{t+1}) \sum_{y_{t+1} \in \mathcal{Y}} \sum_{y_t \in \mathcal{Y}} n_t(x_1, \dots, x_t, y_t) P(y_t, y_{t+1}) f(y_{t+1}, x_{t+1}) \nu(dx_1) \cdots \nu(dx_{t+1}) \\
&= \int_{\mathcal{B}} \int_{\mathcal{X}} w(x_{t+1}) \sum_{y_{t+1} \in \mathcal{Y}} \sum_{y_t \in \mathcal{Y}} n_t(x_1, \dots, x_t, y_t) P(y_t, y_{t+1}) f(y_{t+1}, x_{t+1}) \nu(dx_{t+1}) \nu(dx_1) \cdots \nu(dx_t) \\
&= \int_{\mathcal{B}} \int_{\mathcal{X}} w(x_{t+1}) \sum_{y_{t+1} \in \mathcal{Y}} \sum_{y_t \in \mathcal{Y}} \frac{n_t(x_1, \dots, x_t, y_t)}{\sum_{y \in \mathcal{Y}} n_t(x_1, \dots, x_t, y)} P(y_t, y_{t+1}) f(y_{t+1}, x_{t+1}) \nu(dx_{t+1}) \\
&\quad \sum_{y \in \mathcal{Y}} n_t(x_1, \dots, x_t, y) \nu(dx_1) \cdots \nu(dx_t) \\
&= \int_{\mathcal{F}} \left(\int_{\mathcal{X}} w(x_{t+1}) \sum_{y_{t+1} \in \mathcal{Y}} \sum_{y_t \in \mathcal{Y}} \Pi_t(y_t) P(y_t, y_{t+1}) f(y_{t+1}, x_{t+1}) \nu(dx_{t+1}) \right) d\mathbb{P} \\
&= \int_{\mathcal{F}} \left(\int_{\mathcal{X}} w(x_{t+1}) \Pi_t P f(x_{t+1}) \nu(dx_{t+1}) \right) d\mathbb{P},
\end{aligned}$$

which proves (3.2). Finally, for every nonnegative function $w : \mathcal{P} \mapsto \mathbb{R}_+$, we have

$$\begin{aligned}
\mathbb{E}[w(\Pi_{t+1}) \mid \mathcal{F}_t] &= \mathbb{E} \left[w \left(\frac{\Pi_t P \text{diag}(f)(X_{t+1})}{\Pi_t P f(X_{t+1})} \right) \mid \mathcal{F}_t \right] \\
&= \int_{\mathcal{X}} w \left(\frac{\Pi_t P \text{diag}(f)(x_{t+1})}{\Pi_t P f(x_{t+1})} \right) \Pi_t P f(x_{t+1}) \nu(dx_{t+1}) \equiv (\mathbb{T}w)(\Pi_t).
\end{aligned}$$

Since the right-hand side is a Borel function of Π_t , this establishes $\mathbb{E}[w(\Pi_{t+1}) \mid \mathcal{F}_t] = \mathbb{E}[w(\Pi_{t+1}) \mid \Pi_t] = (\mathbb{T}w)(\Pi_t)$ and the Markov property of the process $\{\Pi_t, \mathcal{F}_t; t \geq 0\}$.

In part (v), $\mathbb{E} \sum_{y' \in \mathcal{Y}_0} \Pi_t(y') = \mathbb{E} \sum_{y' \in \mathcal{Y}_0} 1_{\{y'\}}(Y_t) = \mathbb{P}\{T > t\} = \sum_{y, y' \in \mathcal{Y}_0} \Pi_0(y) P_0^t(y, y')$ for every $t \geq 0$, and $\mathbb{E} \sum_{t \geq 0} \sum_{y' \in \mathcal{Y}_0} \Pi_t(y) = \mathbb{E}T = \sum_{y, y' \in \mathcal{Y}_0} \Pi_0(y) \sum_{t \geq 0} P_0^t(y, y')$, which equals $\sum_{y, y' \in \mathcal{Y}_0} \Pi_0(y) (I - P_0)^{-1}(y, y') \leq \sum_{y, y' \in \mathcal{Y}_0} (I - P_0)^{-1}(y, y') < \infty$, because $\mathcal{Y}_0$ is finite and contains only transient states of Y ; see for example Çinlar [7, Chapter 6].

In part (vi), for every $t \geq 0$, the sum $\sum_{y \in \bar{\mathcal{Y}}} \Pi_t(y)$ is bounded and $\mathcal{F}_t$ -measurable by definition, and $\mathbb{E}[\sum_{y \in \bar{\mathcal{Y}}} \Pi_{t+1}(y) \mid \mathcal{F}_t] = \mathbb{E}[\mathbb{P}\{Y_{t+1} \in \bar{\mathcal{Y}} \mid \mathcal{F}_{t+1}\} \mid \mathcal{F}_t] = \mathbb{P}\{Y_{t+1} \in \bar{\mathcal{Y}} \mid \mathcal{F}_t\} \geq \mathbb{P}\{Y_t \in \bar{\mathcal{Y}} \mid \mathcal{F}_t\} = \sum_{y \in \bar{\mathcal{Y}}} \Pi_t(y)$, where the inequality follows since $\bar{\mathcal{Y}}$ is closed; i.e., $\{Y_{t+1} \in \bar{\mathcal{Y}}\} \supset \{Y_t \in \bar{\mathcal{Y}}\}$.

Part (vii) follows from application of part (vi) to the closed set $\mathcal{Y} \setminus \mathcal{Y}_0$ and by the fact that $\sum_{y \in \mathcal{Y}_0} \Pi_t(y) = 1 - \sum_{y \in \mathcal{Y} \setminus \mathcal{Y}_0} \Pi_t(y)$. $\square$

Proof of Proposition 3. For every $t \geq 0$ and a.s. finite stopping time $\tau \in \mathbb{F}$, we have

$$\mathbb{E} [c(Y_t) 1_{\{t < \tau\}}] = \sum_{y \in \mathcal{Y}} c(y) \mathbb{E} [1_{\{Y_t=y\}} 1_{\{t < \tau\}}] = \mathbb{E} \left[\sum_{y \in \mathcal{Y}} \Pi_t(y) c(y) 1_{\{t < \tau\}} \right] = \mathbb{E} [g(\Pi_t) 1_{\{t < \tau\}}].$$

Moreover, for every $\{1, \dots, a\}$ -valued random variable $\delta \in \mathcal{F}_\tau$, we get

$$\begin{aligned} \mathbb{E} [1_{\{\tau < \infty\}} C(Y_\tau, \delta)] &= \sum_{t=0}^{\infty} \mathbb{E} [1_{\{\tau=t\}} C(Y_t, \delta)] = \sum_{t=0}^{\infty} \sum_{y \in \mathcal{Y}} \mathbb{E} [1_{\{Y_t=y\}} \overbrace{1_{\{\tau=t\}} C(y, \delta)}^{\in \mathcal{F}_t}] \\ &= \sum_{t=0}^{\infty} \sum_{y \in \mathcal{Y}} \mathbb{E} [\Pi_t(y) 1_{\{\tau=t\}} C(y, \delta)] = \sum_{t=0}^{\infty} \mathbb{E} \left[1_{\{\tau=t\}} \sum_{y \in \mathcal{Y}} \Pi_t(y) C(y, \delta) \right] = \mathbb{E} [1_{\{\tau < \infty\}} h(\Pi_\tau, \delta)]. \end{aligned}$$

This completes the proof of (3.6). Since $\delta(\tau) 1_{\{\tau=t\}} = \delta(t) 1_{\{\tau=t\}} \in \mathcal{F}_t$ for every $t \geq 0$, we have $(\tau, \delta(\tau)) \in \Delta$ for every a.s. finite $\mathbb{F}$ -stopping time τ . Moreover, for every $(\tau, \delta) \in \Delta$, (3.6) implies

$$u(\tau, \delta) = \mathbb{E} \left[\sum_{t=0}^{\tau-1} g(\Pi_t) + 1_{\{\tau < \infty\}} h(\Pi_\tau, \delta) \right] \geq \mathbb{E} \left[\sum_{t=0}^{\tau-1} g(\Pi_t) + 1_{\{\tau < \infty\}} h(\Pi_\tau) \right] = u(\tau, \delta(\tau)),$$

and $u^* = \inf_{(\tau, \delta) \in \Delta} u(\tau, \delta) = \inf_{\tau \in \mathbb{F}} u(\tau, \delta(\tau)) = v(\eta)$ in terms of the value function $v(\cdot)$ of the optimal stopping problem in (3.7). $\square$

Proof of Proposition 4. Let $\gamma'_n := \lim_{N \rightarrow \infty} \gamma_n^N$. Since Z_n is integrable, and $\gamma'_n \geq \gamma_n^N = Z_n$, taking the limit as $N \rightarrow \infty$ in the third equation of (3.10) and the monotone convergence theorem give $\gamma'_n = \max\{Z_n, \mathbb{E}[\gamma'_{n+1} \mid \mathcal{F}_n]\}$ for every $n \geq 0$. Hence, $(\gamma'_n)_{n \geq 0}$ is an $\mathbb{F}$ -supermartingale majorizing $(Z_n)_{n \geq 0}$. Because $\gamma_n^N \leq \gamma_n$ for every $N \geq n$, we also have $\gamma'_n \leq \gamma_n$. For the converse inequality, it is enough to prove that $\gamma'_n \geq \mathbb{E}[Z_\tau \mid \mathcal{F}_n]$ for every $\tau \in C_n$. Fix $\tau \in C_n$ and $F \in \mathcal{F}_n$. Then the $\mathbb{F}$ -supermartingale property of $(\gamma'_n)_{n \geq 0}$ implies

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

for every $m \geq n$. Because $\gamma'_k \geq Z_k$ for every $k \geq 0$, we have $\gamma'_\tau \geq Z_\tau$, $\gamma'_m \geq -(\gamma'_m)^-$, and

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

However, $\gamma'_m \geq Z_m$ implies $(\gamma'_m)^- \leq Z_m^- = \sum_{k=0}^{m-1} g(\Pi_k) + h(\Pi_m) \leq Z_\tau^- + \sup_{\pi \in \mathcal{P}} |h(\pi)|$ on the event $\{\tau > m\}$. Since τ is a.s. finite, and Z_τ^- has finite expectation, we have $\int_{\{\tau > m\}} (\gamma'_m)^- d\mathbb{P} \leq \int_{\{\tau > m\}} (Z_\tau^- + \sup_{\pi \in \mathcal{P}} |h(\pi)|) d\mathbb{P} \rightarrow 0$ as $m \rightarrow \infty$. Finally, limit as $m \rightarrow \infty$ of the displayed inequality and dominated convergence ($Z_\tau = -Z_\tau^-$ is integrable) imply that $\int_F \gamma'_n d\mathbb{P} \geq \int_{F \cap \{n \leq \tau < \infty\}} Z_\tau d\mathbb{P} = \int_F Z_\tau d\mathbb{P}$, and this completes the proof. $\square$

Proof of Proposition 5. Fix any $N \geq 0$. The equality is true for $n = N$. Suppose it holds for some $1 \leq n \leq N$. Then $\gamma_{n-1}^N = \max\{Z_{n-1}, \mathbb{E}[\gamma_n^N \mid \mathcal{F}_{n-1}]\}$ equals

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

where the second equality follows from (iv) of Proposition 2, and this completes the proof. $\square$

Proof of Lemma 6. Fix $\pi, \bar{\pi} \in \mathcal{P}$, $\lambda \in [0, 1]$, and $\bar{\lambda} := 1 - \lambda$. Then

$$\begin{aligned} & (\mathbb{T}w)(\lambda\pi + \bar{\lambda}\bar{\pi}) = \int_{\mathcal{X}} w \left(\frac{(\lambda\pi + \bar{\lambda}\bar{\pi})P \operatorname{diag}(f)(x)}{(\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)} \right) (\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)\nu(dx) \\ &= \int_{\mathcal{X}} w \left(\frac{\lambda}{(\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)} \pi P \operatorname{diag}(f)(x) + \frac{\bar{\lambda}}{(\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)} \bar{\pi} P \operatorname{diag}(f)(x) \right) \\ & \quad (\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)\nu(dx) \\ &= \int_{\mathcal{X}} w \left(\frac{\lambda\pi Pf(x)}{(\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)} \cdot \frac{\pi P \operatorname{diag}(f)(x)}{\pi Pf(x)} + \frac{\bar{\lambda}\bar{\pi} Pf(x)}{(\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)} \cdot \frac{\bar{\pi} P \operatorname{diag}(f)(x)}{\bar{\pi} Pf(x)} \right) \\ & \quad (\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)\nu(dx) \\ &\geq \int_{\mathcal{X}} \left[\frac{\lambda\pi Pf(x)}{(\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)} \cdot w \left(\frac{\pi P \operatorname{diag}(f)(x)}{\pi Pf(x)} \right) + \frac{\bar{\lambda}\bar{\pi} Pf(x)}{(\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)} \cdot w \left(\frac{\bar{\pi} P \operatorname{diag}(f)(x)}{\bar{\pi} Pf(x)} \right) \right] \\ & \quad (\lambda\pi + \bar{\lambda}\bar{\pi})Pf(x)\nu(dx) \\ &= \lambda \int_{\mathcal{X}} w \left(\frac{\pi P \operatorname{diag}(f)(x)}{\pi Pf(x)} \right) \pi Pf(x)\nu(dx) + \bar{\lambda} \int_{\mathcal{X}} w \left(\frac{\bar{\pi} P \operatorname{diag}(f)(x)}{\bar{\pi} Pf(x)} \right) \bar{\pi} Pf(x)\nu(dx) \\ &= \lambda(\mathbb{T}w)(\pi) + \bar{\lambda}(\mathbb{T}w)(\bar{\pi}), \end{aligned}$$

which proves the concavity of $\pi \mapsto (\mathbb{T}w)(\pi)$. $\square$

Proof of Theorem 8. Let us fix an arbitrary $\pi \in \mathcal{P}$. It is enough to prove the displayed inequalities. Because $C_0 \supseteq C_0^N$, we have $v_0(\pi) = -\sup_{\tau \in C_0} \mathbb{E}_{\pi} Z_{\tau} \leq -\sup_{\tau \in C_0^N} \mathbb{E}_{\pi} Z_{\tau} = v_0^N(\pi)$, which is the first inequality. Let us prove the second inequality. Because $0 \geq -v_0(\pi) \geq -h(\pi) > -\infty$, for every $\varepsilon > 0$ there exists some stopping time $\tau_{\varepsilon} \equiv \tau_{\varepsilon}(\pi) \in C_0$

such that

$$\begin{aligned}
 -v_0(\pi) - \varepsilon &\leq \mathbb{E}_\pi \left[-\sum_{k=0}^{\tau_\varepsilon-1} g(\Pi_k) - h(\Pi_{\tau_\varepsilon}) \right] \leq \left[-\sum_{k=0}^{(\tau_\varepsilon \wedge N)-1} g(\Pi_k) - h(\Pi_{\tau_\varepsilon \wedge N}) \right] \\
 &\quad + \mathbb{E}_\pi [h(\Pi_{\tau_\varepsilon \wedge N}) - h(\Pi_{\tau_\varepsilon})] \leq -v_0^N(\pi) + \mathbb{E}_\pi [1_{\{\tau_\varepsilon > N\}} (h(\Pi_{\tau_\varepsilon \wedge N}) - h(\Pi_{\tau_\varepsilon}))] \\
 (A.1) \quad &\leq -v_0^N(\pi) + \|h\| \mathbb{P}_\pi \{\tau_\varepsilon > N\} \leq -v_0^N(\pi) + \frac{\|h\|}{N} \mathbb{E}_\pi \tau_\varepsilon.
 \end{aligned}$$

From the first inequality, it also follows that

$$\begin{aligned}
 -\|h\| - \varepsilon &\leq -v_0(\pi) - \varepsilon \leq \mathbb{E}_\pi \left[-\sum_{k=0}^{\tau_\varepsilon-1} g(\Pi_k) - h(\Pi_{\tau_\varepsilon}) \right] \leq -\mathbb{E}_\pi \sum_{k=0}^{\tau_\varepsilon-1} \sum_{y \in \mathcal{Y}} \Pi_k(y) c(y) \\
 &\leq -\mathbb{E}_\pi \sum_{k=0}^{\tau_\varepsilon-1} \sum_{y \in \mathcal{Y} \setminus \mathcal{Y}_0} \Pi_k(y) c(y) \leq -c_p \mathbb{E}_\pi \sum_{k=0}^{\tau_\varepsilon-1} \left(1 - \sum_{y \in \mathcal{Y}_0} \Pi_k(y) \right) \\
 &= -c_p \mathbb{E}_\pi \tau_\varepsilon + c_p \mathbb{E}_\pi \sum_{k=0}^{\tau_\varepsilon-1} \sum_{y \in \mathcal{Y}_0} \Pi_k(y) \leq -c_p \mathbb{E}_\pi \tau_\varepsilon + c_p \mathbb{E}_\pi \sum_{k=0}^{\infty} \sum_{y \in \mathcal{Y}_0} \Pi_k(y) \\
 &\leq -c_p \mathbb{E}_\pi \tau_\varepsilon + c_p \sum_{y, y' \in \mathcal{Y}_0} (I - P_0)^{-1}(y, y')
 \end{aligned}$$

because of (v) of Proposition 2, and we obtain $\mathbb{E}_\pi \tau_\varepsilon \leq [(\|h\| + \varepsilon)/c_p] + \sum_{y, y' \in \mathcal{Y}_0} (I - P_0)^{-1}(y, y')$. Using this bound in (A.1) gives

$$v_0^N(\pi) \leq v_0(\pi) + \varepsilon + \frac{\|h\|}{N} \mathbb{E}_\pi \tau_\varepsilon \leq v_0(\pi) + \varepsilon + \frac{\|h\|}{N} \left(\frac{\|h\| + \varepsilon}{c_p} + \sum_{y, y' \in \mathcal{Y}_0} (I - P_0)^{-1}(y, y') \right),$$

and because $\varepsilon > 0$ is arbitrary, the second inequality in (3.16) follows. $\square$

Proof of Proposition 9. Recall that $v_0^N(\cdot) = \mathbb{M}^N h(\cdot)$, and $h(\cdot)$ is continuous. Therefore, it is enough to prove that $\mathbb{M}$ preserves continuity, which is immediate if $\mathbb{T}$ does, where $\mathbb{T}$ is defined by (3.4). Note that the mappings $\pi \mapsto (\pi P \text{diag}(f)(x))/\pi P f(x)$ and $\pi \mapsto \pi P f(x)$ from $\mathcal{P}$ into $\mathbb{R}_+$ are continuous for every $x \in \mathcal{X}$, and that

$$0 \leq \pi P f(x) = \sum_{y, y' \in \mathcal{Y}} \pi(y) P(y, y') f(y', x) \leq \sum_{y, y' \in \mathcal{Y}} \pi(y) f(y', x) = \sum_{y' \in \mathcal{Y}} f(y', x).$$

Since $\sum_{y' \in \mathcal{Y}} f(y', x)$ is $\nu(dx)$ -integrable, the dominated convergence theorem yields that $\mathbb{T}$ preserves continuity. $\square$

Proof of Corollary 10. For every $N \geq 0$, the mapping $v_0^N : \mathcal{P} \mapsto \mathbb{R}$ is continuous, and $(v_0^N)_{N \geq 0}$ converges to v_0 uniformly as $N \rightarrow \infty$. $\square$

Proof of Proposition 11. Fix $1 \leq j \leq a$. Recall that $\hbar(\cdot, j) \geq v_0^N(\cdot) \geq v_0^{N+1}(\cdot)$. If $\pi \in \Gamma_{N+1}^{(j)}$, then $v_0^{N+1}(\pi) = \hbar(\pi, j)$, and therefore, $v_0^N(\pi) = \hbar(\pi, j)$ and $\pi \in \Gamma_N^{(j)}$. Hence, $\Gamma_{N+1}^{(j)} \subseteq \Gamma_N^{(j)}$.

for every $N \geq 0$. Because $v_0(\cdot) = \lim_{N \rightarrow \infty} \downarrow v_0^N(\cdot)$, we have $\Gamma^{(j)} \subseteq \Gamma_N^{(j)}$, $N \geq 0$ by a similar argument. Now suppose that $\bar{h}(\pi, j) \leq \min\{h(\pi), g(\pi)\}$ for some $\pi \in \mathcal{P}$. Then

$$\bar{h}(\pi, j) \geq v_0(\pi) = \min\{h(\pi), g(\pi) + (\mathbb{T}v_0)(\pi)\} \geq \min\{h(\pi), g(\pi)\} \geq \bar{h}(\pi, j)$$

implies $v_0(\pi) = \bar{h}(\pi, j)$ and $\pi \in \Gamma^{(j)}$. Finally, for every $y \in \mathcal{Y}_j$ we have $\bar{h}(e_y, j) = (e_y C)(j) = C(y, j) = 0 \leq \min\{h(e_y), g(e_y)\}$, and $\{\pi \in \mathcal{P} \mid \bar{h}(\pi, j) \leq \min\{h(\pi), g(\pi)\}\} \supseteq \{e_y \mid y \in \mathcal{Y}_j\}$.

Because the functions $v_0^N(\cdot)$, $v_0(\cdot)$, and $\bar{h}(\cdot, j)$ for every $N \geq 0$ and $1 \leq j \leq a$ are continuous, the sets $\Gamma_N^{(j)}$, $N \geq 0$ are closed. To prove that $\Gamma_N^{(j)}$ is convex, take $\pi, \bar{\pi} \in \Gamma_N^{(j)}$, $\lambda \in [0, 1]$, and set $\bar{\lambda} := 1 - \lambda$. Because $v_0^N(\cdot)$ is concave, and $\pi \mapsto \bar{h}(\pi, j) = (\pi C)(j)$ is a linear mapping, we have $\lambda v_0^N(\pi) + \bar{\lambda} v_0^N(\bar{\pi}) \leq v_0^N(\lambda\pi + \bar{\lambda}\bar{\pi}) \leq \bar{h}(\lambda\pi + \bar{\lambda}\bar{\pi}, j) = \lambda\bar{h}(\pi, j) + \bar{\lambda}\bar{h}(\bar{\pi}, j) = \lambda v_0^N(\pi) + \bar{\lambda} v_0^N(\bar{\pi})$, which implies that $\lambda\pi + \bar{\lambda}\bar{\pi} \in \Gamma_N^{(j)}$; i.e., $\Gamma_N^{(j)}$ is convex.

Clearly, $\Gamma^{(j)} \subseteq \cap_{N=1}^{\infty} \Gamma_N^{(j)}$. To show the reverse inequality, take any $\pi \in \cap_{N=1}^{\infty} \Gamma_N^{(j)}$. Then $v_0^N(\pi) = \bar{h}(\pi, j)$ for every $N \geq 0$, and the limit as $N \rightarrow \infty$ gives $v_0(\pi) = \bar{h}(\pi, j)$; i.e., $\pi \in \Gamma^{(j)}$. The proof of $\Gamma = \cap_{N=1}^{\infty} \Gamma_N$ is similar; and $\Gamma = \cup_{j=1}^a \Gamma^{(j)}$ and $\Gamma_N = \cup_{j=1}^a \Gamma_N^{(j)}$ are obvious, which also imply that $\Gamma_0 \supseteq \Gamma_1 \supseteq \dots \supseteq \Gamma$.

Earlier we proved that $\Gamma^{(j)} \supseteq \{e_y \mid y \in \mathcal{Y}_j\}$ for every $1 \leq j \leq a$. Therefore, $\Gamma = \cup_{j=1}^a \Gamma^{(j)} \supseteq \{e_y \mid y \in \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_a\}$. Now we shall prove that the set on the left is strictly larger. Fix any $1 \leq j \leq a$, $y \in \mathcal{Y}_j$, and $y' \in \mathcal{Y} \setminus \{y\}$. Let us define

$$\pi_\lambda := \lambda e_{y'} + (\bar{\lambda}) e_y \quad \text{for every } \lambda \in \left(0, \frac{c(y)}{C(y', j) + c(y)}\right).$$

Because $c(y'') > 0$ for every $y'' \in \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_a$, the interval above is non-empty, and every λ in that interval satisfies $(1 - \lambda)c(y) \geq \lambda C(y', j) = \lambda C(y', j) + (\bar{\lambda})C(y, j) = \bar{h}(\pi_\lambda, j)$. Hence, we have $h(\pi_\lambda) \leq \bar{h}(\pi_\lambda, j) \leq (\bar{\lambda})c(y) \leq \lambda c(y') + (\bar{\lambda})c(y) = g(\pi_\lambda) \leq g(\pi_\lambda) + (\mathbb{T}v_0)(\pi_\lambda)$. Because $v_0(\cdot)$ satisfies the optimality equation in (3.14), this implies that $v_0(\pi_\lambda) = h(\pi_\lambda)$. Therefore, $\Gamma \ni \pi_\lambda \notin \{e_y \mid y \in \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_a\}$. $\square$

Proof of Theorem 13. Because on the event $\{\sigma > n\}$, we have $\Pi_n \notin \Gamma$; i.e., $v_0(\Pi_n) < h(\Pi_n)$, and $\gamma_n = -\sum_{k=0}^{n-1} g(\Pi_k) - v_0(\Pi_n) > -\sum_{k=0}^{n-1} g(\Pi_k) - h(\Pi_n) = Z_n$ by Lemma 12. Then the dynamic programming equation in (3.10) implies that $\gamma_n = \mathbb{E}[\gamma_{n+1} \mid \mathcal{F}_n]$ on $\{\sigma > n\}$. Therefore, for every $n \geq 0$, we have

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

which proves that $\{\gamma_{\sigma \wedge n}, \mathcal{F}_n; n \geq 0\}$ is a martingale. Therefore, $-v_0 = \mathbb{E}\gamma_0 = \mathbb{E}\gamma_{\sigma \wedge n} = \mathbb{E}\gamma_\sigma 1_{\{\sigma \leq n\}} + \mathbb{E}\gamma_n 1_{\{\sigma > n\}} = \mathbb{E}Z_\sigma 1_{\{\sigma \leq n\}} + \mathbb{E}\gamma_n 1_{\{\sigma > n\}}$ for every $n \geq 0$. Because $\gamma_n \leq 0$ for every $n \geq 0$, taking limits supremum as $n \rightarrow \infty$ of both sides and Fatou's lemma imply that

$$(A.2) \quad -v_0 \leq \mathbb{E}Z_\sigma 1_{\{\sigma < \infty\}} + \mathbb{E}(\limsup_{n \rightarrow \infty} \gamma_n) 1_{\{\sigma = \infty\}}.$$

Observe that

$$\begin{aligned}\gamma_n &= -\sum_{k=0}^{n-1} g(\Pi_k) - v_0(\Pi_n) \leq -\sum_{k=0}^{n-1} \sum_{y \in \mathcal{Y}} \Pi_k(y) c(y) \leq -\sum_{k=0}^{n-1} \sum_{y \in \mathcal{Y} \setminus \mathcal{Y}_0} \Pi_k(y) c(y) \\ &\leq -c_p \sum_{k=0}^{n-1} \sum_{y \in \mathcal{Y} \setminus \mathcal{Y}_0} \Pi_k(y) = -c_p \sum_{k=0}^{n-1} \left(1 - \sum_{y \in \mathcal{Y}_0} \Pi_k(y) \right) = -c_p n + \sum_{k=0}^{n-1} \sum_{y \in \mathcal{Y}_0} \Pi_k(y),\end{aligned}$$

and because $\sum_{k=0}^{\infty} \sum_{y \in \mathcal{Y}_0} \Pi_k(y) < \infty$ by (v) of Proposition 2, taking limit supremum as $n \rightarrow \infty$ of both sides gives that $\limsup_{n \rightarrow \infty} \gamma_n = -\infty$ almost surely. Using this in (A.2) and the fact that $-\infty < -\|h\| \leq -v_0$ imply that $\sigma < \infty$ almost surely, and $-\infty < -v_0 \leq \mathbb{E}Z_\sigma = -\mathbb{E}Z_\sigma^-$. Therefore, $\sigma \in C_0$, and $-v_0 \leq \mathbb{E}Z_\sigma \leq \sup_{\tau \in C_0} \mathbb{E}Z_\tau = -v_0$; i.e., $\mathbb{E}Z_\sigma = -v_0$, which proves the optimality of σ in C_0 . Finally,

$$\begin{aligned}\infty > \mathbb{E}Z_\sigma^- &= \mathbb{E} \left[\sum_{k=0}^{\sigma-1} g(\Pi_k) + h(\Pi_\sigma) \right] \geq \mathbb{E} \sum_{k=0}^{\sigma-1} \sum_{y \in \mathcal{Y}} \Pi_k(y) c(y) \geq \mathbb{E} \sum_{k=0}^{\sigma-1} \sum_{y \in \mathcal{Y} \setminus \mathcal{Y}_0} \Pi_k(y) c(y) \\ &\geq c_p \mathbb{E} \sum_{k=0}^{\sigma-1} \sum_{y \in \mathcal{Y} \setminus \mathcal{Y}_0} \Pi_k(y) = c_p \mathbb{E} \sum_{k=0}^{\sigma-1} \left(1 - \sum_{y \in \mathcal{Y}_0} \Pi_k(y) \right) = c_p \mathbb{E}\sigma - c_p \mathbb{E} \sum_{k=0}^{\sigma-1} \sum_{y \in \mathcal{Y}_0} \Pi_k(y) \\ &\geq c_p \mathbb{E}\sigma - c_p \mathbb{E} \sum_{k=0}^{\infty} \sum_{y \in \mathcal{Y}_0} \Pi_k(y),\end{aligned}$$

where the second term is finite by (v) of Proposition 2 (which also justifies the split of the expectation in passing to the second equality), and therefore, $\mathbb{E}\sigma < \infty$. $\square$

Proof of Theorem 14. Because $c_\varepsilon(y) > 0$ for every $y \in \mathcal{Y} \setminus \mathcal{Y}_0$, all of the previous results apply to the solution of the problem $-v_0\langle \varepsilon \rangle = \sup_{\tau \in C_0\langle \varepsilon \rangle} \mathbb{E}Z_\tau\langle \varepsilon \rangle$ for every $\varepsilon > 0$. Particularly,

$$(A.3) \quad \gamma_n\langle \varepsilon \rangle = \max\{Z_n\langle \varepsilon \rangle, \mathbb{E}[\gamma_{n+1}\langle \varepsilon \rangle \mid \mathcal{F}_n]\}, \quad n \geq 0.$$

Clearly, $(\gamma_n\langle \varepsilon \rangle)_{\varepsilon > 0}$ is decreasing, and the limit

$$\bar{\gamma}_n := \lim_{\varepsilon \searrow 0} \uparrow \gamma_n\langle \varepsilon \rangle \quad \text{for every } n \geq 0$$

exists. Because $\gamma_n\langle \varepsilon \rangle \geq Z_n\langle \varepsilon \rangle - \varepsilon n$, and Z_n is integrable, the monotone convergence theorem after taking limit as $\varepsilon \searrow 0$ in (A.3) gives

$$(A.4) \quad \bar{\gamma}_n = \max\{Z_n, \mathbb{E}[\bar{\gamma}_{n+1} \mid \mathcal{F}_n]\}, \quad n \geq 0.$$

If $\tau \in C_n\langle \varepsilon \rangle$, then τ is an $\mathbb{F}$ -stopping time such that (i) $n \leq \tau < \infty$ almost surely, and (ii) $\infty > \mathbb{E}(Z_\tau\langle \varepsilon \rangle)^- = \mathbb{E}Z_\tau^- + \varepsilon \mathbb{E}\tau \geq \mathbb{E}Z_\tau^-$; therefore, τ also belongs to C_n , and

$$\gamma_n\langle \varepsilon \rangle = \text{ess sup}_{\tau \in C_n\langle \varepsilon \rangle} \mathbb{E}[Z_\tau\langle \varepsilon \rangle \mid \mathcal{F}_n] \leq \text{ess sup}_{\tau \in C_n\langle \varepsilon \rangle} \mathbb{E}[Z_\tau \mid \mathcal{F}_n] \leq \text{ess sup}_{\tau \in C_n} \mathbb{E}[Z_\tau \mid \mathcal{F}_n] = \gamma_n, \quad n \geq 0;$$

taking limit as $\varepsilon \searrow 0$ gives $\bar{\gamma}_n \leq \gamma_n$ almost surely for every $n \geq 0$. For the proof of the reverse inequality, it is enough to show that $\bar{\gamma}_n \geq \mathbb{E}[Z_\tau \mid \mathcal{F}_n]$ for every $\tau \in C_n$.

Note that, by (A.4) the process $(\bar{\gamma}_k)_{k \geq 0}$ is an $\mathbb{F}$ -supermartingale dominating $(Z_k)_{k \geq 0}$. Therefore, for every $\tau \in C_n$, $F \in \mathcal{F}_n$, and $m \geq n$, we have

$$\begin{aligned}
 (\text{A.5}) \quad \int_F \bar{\gamma}_n d\mathbb{P} &= \int_{F \cap \{\tau=n\}} \bar{\gamma}_\tau d\mathbb{P} + \int_{F \cap \{\tau>n\}} \bar{\gamma}_n d\mathbb{P} \geq \int_{F \cap \{\tau=n\}} \bar{\gamma}_\tau d\mathbb{P} + \int_{F \cap \{\tau>n\}} \bar{\gamma}_{n+1} d\mathbb{P} \\
 &= \int_{F \cap \{n \leq \tau \leq n+1\}} \bar{\gamma}_\tau d\mathbb{P} + \int_{F \cap \{\tau>n+1\}} \bar{\gamma}_{n+1} d\mathbb{P} \geq \cdots \geq \int_{F \cap \{n \leq \tau \leq m\}} \bar{\gamma}_\tau d\mathbb{P} + \int_{F \cap \{\tau>m\}} \bar{\gamma}_m d\mathbb{P} \\
 &\geq \int_{F \cap \{n \leq \tau \leq m\}} Z_\tau d\mathbb{P} + \int_{F \cap \{\tau>m\}} (\bar{\gamma}_m)^- d\mathbb{P}.
 \end{aligned}$$

Because $\tau \in C_n$, we have $\tau < \infty$ almost surely, and $\mathbb{E}|Z_\tau| = -\mathbb{E}Z_\tau^- < \infty$. Therefore, $\lim_{m \rightarrow \infty} \int_{F \cap \{n \leq \tau \leq m\}} Z_\tau d\mathbb{P} = \int_{F \cap \{n \leq \tau < \infty\}} Z_\tau d\mathbb{P} = \int_F Z_\tau d\mathbb{P}$. Moreover, since $\bar{\gamma}_m \geq Z_m$,

$$\begin{aligned}
 (\bar{\gamma}_m)^- &\leq Z_m^- = \sum_{k=0}^{m-1} g(\Pi_k) + h(\Pi_m) \leq \sum_{k=0}^{\tau-1} g(\Pi_k) + h(\Pi_m) \leq \sum_{k=0}^{\tau-1} g(\Pi_k) - h(\Pi_\tau) + h(\Pi_m) \\
 &\leq Z_\tau^- + \|h\| \quad \text{on the event } \{\tau > m\},
 \end{aligned}$$

and $\liminf_{m \rightarrow \infty} \int_{\{\tau>m\}} (\bar{\gamma}_m)^- d\mathbb{P} \leq \liminf_{m \rightarrow \infty} \left(\int_{\{\tau>m\}} Z_\tau^- d\mathbb{P} + \|h\| \mathbb{P}\{\tau > m\} \right) = 0$. Then taking limit infimum as $m \rightarrow \infty$ in (A.5) gives $\int_F \bar{\gamma}_n d\mathbb{P} \geq \int_F Z_\tau d\mathbb{P}$, which implies that $\bar{\gamma}_n \geq \mathbb{E}[Z_\tau \mid \mathcal{F}_n]$ almost surely and completes the proof of $\gamma_n = \bar{\gamma}_n \equiv \lim_{\varepsilon \searrow 0} \uparrow \gamma_n \langle \varepsilon \rangle$. Finally, $-v_0 = \mathbb{E}\gamma_0 = \lim_{\varepsilon \searrow 0} \mathbb{E}\gamma_0 \langle \varepsilon \rangle = \lim_{\varepsilon \searrow 0} \uparrow v_n \langle \varepsilon \rangle$ follows from the monotone convergence theorem. $\square$

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] S. Asmussen, O. Nerman, and M. Olsson. Fitting phase-type distributions via the EM algorithm. *Scand. J. Stat.*, 23(4):419–441, 1996.
- [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. Blackwell and M. A. Girshick. *Theory of Games and Statistical Decisions*. Dover Publications Inc., New York, 1979. Reprint of the 1954 edition.
- [5] A. Bobbio, A. Horváth, M. Scarpa, and M. Telek. Acyclic discrete phase type distributions: properties and a parameter estimation algorithm. *Perform. Eval.*, 54(1):1–32, 2003.
- [6] Y. S. Chow, H. Robbins, and D. Siegmund. *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin Co., Boston, Mass., 1971.
- [7] E. Çinlar. *Introduction to Stochastic Processes*. Prentice-Hall Inc., Englewood Cliffs, N.J., 1975.
- [8] M. J. Crowder. *Classical Competing Risks*. Chapman & Hall/CRC Press, London, 2001.
- [9] S. Dayanik, C. Goulding, and H. V. Poor. Bayesian sequential change diagnosis. <http://www.princeton.edu/~sdayanik/diagnosis.pdf>, August 2006.
- [10] 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.

- [11] 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.
- [12] C.-D. Fuh. SPRT and CUSUM in hidden Markov models. *Ann. Statist.*, 31(3):942–977, 2003.
- [13] P. Hackl and A. H. Westlund. Statistical analysis of "structural change": An annotated bibliography. *Empirical Economics*, 14(2):167–92, 1989. available at <http://ideas.repec.org/a/spr/empeco/v14y1989i2p167-92.html>.
- [14] G. Horton and C. Isensee. Approximation of discrete phase-type distributions. *Proc. of 38th Ann. Simul. Symposium*, 2005.
- [15] A. Horváth and M. Telek. Phfit: A general phase-type fitting tool. *Proc. of 12th Performance TOOLS*, pages 82–91, 2002.
- [16] M. A. Johnson and M. R. Taaffe. The denseness of phase distributions. Res. Memo. 88-20, School of Ind. Eng., Purdue Univ., West Lafayette, IN, 1988.
- [17] T. L. Lai. Sequential change-point detection in quality control and dynamical systems. *J. Royal Statistical Society Series B*, 57(4):613–658, 1995.
- [18] T. L. Lai. Sequential multiple hypothesis testing and efficient fault detection-isolation in stochastic systems. *IEEE Trans. Inform. Theory*, 46(2):595–608, 2000.
- [19] 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.
- [20] G. Latouche and V. Ramaswami. *Introduction to Matrix Analytic Methods in Stochastic Modeling*. ASA-SIAM Series on Statistics and Applied Probability. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1999.
- [21] G. Lorden. Procedures for reacting to a change in distribution. *Ann. Math. Statist.*, 42:1897–1908, 1971.
- [22] D. J. C. Mackay. *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, June 2002.
- [23] M. F. Neuts. *Liber Amicorum Prof. Emeritus H. Florin*, chapter Probability distributions of phase type, pages 173–206. University of Louvain, 1975.
- [24] M. F. Neuts. *Matrix-Geometric Solutions in Stochastic Models*. Dover Publications Inc., New York, 1994. An algorithmic approach, Corrected reprint of the 1981 original.
- [25] I. V. Nikiforov. A generalized change detection problem. *IEEE Trans. Inform. Theory*, 41(1):171–187, 1995.
- [26] A. N. Shiryaev. Optimal methods in quickest detection problems. *Teor. Veroyatnost. i Primenen.*, 8:26–51, 1963.
- [27] A. N. Shiryaev. *Optimal Stopping Rules*. Springer-Verlag, New York, 1978. Translated from the Russian by A. B. Aries, Applications of Mathematics, Vol. 8.
- [28] A. Wald and J. Wolfowitz. Bayes solutions of sequential decision problems. *Ann. Math. Statistics*, 21:82–99, 1950.
- [29] B. Yakir. Optimal detection of a change in distribution when the observations form a Markov chain with a finite state space. In *Change-Point Problems (South Hadley, MA, 1992)*, volume 23 of *IMS Lecture Notes Monogr. Ser.*, pages 346–358. Inst. Math. Statist., Hayward, CA, 1994.

(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, S. Dayanik: `sdayanik@princeton.edu`

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