

Active Hypothesis Testing for Anomaly Detection

Kobi Cohen and Qing Zhao, *Fellow, IEEE*

Abstract—The problem of detecting a single anomalous process among a finite number M of processes is considered. At each time, a subset of the processes can be observed, and the observations from each chosen process follow two different distributions, depending on whether the process is normal or abnormal. The objective is a sequential search strategy that minimizes the expected detection time subject to an error probability constraint. This problem can be considered as a special case of active hypothesis testing first considered by Chernoff where a randomized strategy, referred to as the Chernoff test, was proposed and shown to be asymptotically (as the error probability approaches zero) optimal. For the special case considered in this paper, we show that a simple deterministic test achieves asymptotic optimality and offers better performance in the finite regime. We further extend the problem to the case where multiple anomalous processes are present. In particular, we examine the case where only an upper bound on the number of anomalous processes is known.

Index Terms—Sequential detection, anomaly detection, dynamic search, active hypothesis testing, controlled sensing.

I. INTRODUCTION

WE CONSIDER the problem of detecting a single anomalous process among M processes. Borrowing terminologies from target search, we refer to these processes as cells and the anomalous process as the target which can locate in any of the M cells. The decision maker is allowed to search for the target over K cells at a time ($1 \leq K \leq M$). The observations from searching a cell are i.i.d. realizations drawn from two different distributions f and g , depending on whether the target is absent or present. The objective is a sequential search strategy that dynamically determines which cells to search at each time and when to terminate the search so that the expected detection time is minimized under a constraint on the probability of declaring a wrong location of the target.

The problem under study applies to intrusion detection in cyber-systems when an intrusion to a subnet has been detected and the objective is to locate the abnormal component in

the subnet (since the probability of each component being compromised is small, with high probability, there is only one abnormal component). It also finds applications in target search, fraud detection, and spectrum scanning in cognitive radio networks.

A. A Case of Active Hypothesis Testing

The above problem is a special case of the sequential experiment design problem first studied by Chernoff in 1959 [1]. Compared with the classic sequential hypothesis testing pioneered by Wald [2] where the observation model under each hypothesis is predetermined, the sequential design of experiments has a control aspect that allows the decision maker to choose the experiment to be conducted at each time. Different experiments generate observations from different distributions under each hypothesis. Intuitively, as more observations are gathered, the decision maker becomes more certain about the true hypothesis, which in turn leads to better choices of experiments. Chernoff focused on the case of binary hypotheses and showed that a *randomized* strategy, referred to as the Chernoff test, is asymptotically optimal as the maximum error probability diminishes. Specifically, the Chernoff test chooses the current experiment based on a distribution that depends on past actions and observations. Variations and extensions of the problem and the Chernoff test were studied in [3]–[8], where the problem was referred to as controlled sensing for hypothesis testing in [4]–[6] and active hypothesis testing in [7] and [8] (see a more detailed discussion in Section I-C).

It is not difficult to see that the anomaly detection problem considered in this paper is a special case of the active hypothesis testing problem considered in [1], [3]–[5], [7], and [8]. In particular, under each hypothesis that the target is located in a particular cell, the distribution (either f or g) of the next observation depends on the cell chosen to be searched. The Chernoff test and its variations proposed in [3]–[5], [7], and [8] thus directly apply to our problem. However, in contrast to the randomized nature of the Chernoff test and its variations, we show in this paper that a simple *deterministic* test achieves asymptotic optimality and offers better performance in the finite regime.

B. Main Results

Similar to [1], [3]–[5], and [7], we focus on asymptotically optimal policies in terms of minimizing the detection time as the error probability approaches zero. The asymptotic optimality of the Chernoff test as shown in [1] requires that under any experiment, any pair of hypotheses are

Manuscript received March 5, 2014; revised October 8, 2014; accepted December 16, 2014. Date of publication January 6, 2015; date of current version February 12, 2015. This work was supported in part by the U.S. Army Research Laboratory under Grant W911NF1120086 and in part by the National Science Foundation under Grant CCF-1320065 and Grant CNS-1321115. This paper was presented at the 2014 Information Theory and Applications Workshop.

K. Cohen is with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Champaign, IL 61820 USA (e-mail: yscohen@illinois.edu).

Q. Zhao is with the Department of Electrical and Computer Engineering, University of California at Davis, Davis, CA 95616 USA (e-mail: qzhao@ucdavis.edu).

Communicated by V. Saligrama, Associate Editor for Signal Processing.
Digital Object Identifier 10.1109/TIT.2014.2387857

distinguishable (i.e., has positive Kullback-Liebler (KL) divergence). This assumption does not hold in the anomaly detection problem considered in this paper. For instance, under the experiment of searching the i^{th} cell, the hypotheses of the target being in the j^{th} ($j \neq i$) and the k^{th} ($k \neq i$) cells yield the same observation distribution f . Nevertheless, we show in Theorem 2 that the Chernoff test preserves its asymptotic optimality for the problem at hand even without this positivity assumption on all KL divergences. As a result, it serves as a bench mark for comparison.

The Chernoff test, when applied directly to the anomaly detection problem, leads to a randomized cell selection rule: the cells to be searched at the current time are drawn randomly according to a distribution determined by past observations and actions. The main result of this paper is to show that a simple deterministic policy offers the same asymptotic optimality yet with significant performance gain in the finite regime and considerable reduction in implementation complexity. Specifically, under the proposed policy, the selection rule $\phi(n)$ indicating which K cells should be searched at time n is given by:

$$\phi(n) = \begin{cases} (m^{(1)}(n), m^{(2)}(n), \dots, m^{(K)}(n)), & \text{if } D(g||f) \geq \frac{D(f||g)}{(M-1)} \text{ or } K = M \\ (m^{(2)}(n), m^{(3)}(n), \dots, m^{(K+1)}(n)), & \text{if } D(g||f) < \frac{D(f||g)}{(M-1)} \text{ and } K < M \end{cases}$$

where $m^{(i)}(n)$ denotes the cell index with the i^{th} highest sum of log-likelihood ratio (LLR) collected from this cell up to time n , and $D(\cdot||\cdot)$ is the KL divergence between two distributions. Since $D(g||f)$ is the key quantity in the cell selection rule, we refer to the proposed deterministic policy as the DGF policy.

This deterministic selection rule is intuitively satisfying. Consider, for example, $K = 1$. In this case, the DGF policy selects, at each time, either the cell with the largest sum LLRs or the cell with the second largest sum LLRs, depending on the order of $D(g||f)$ and $D(f||g)/(M-1)$. The intuition behind this selection rule is that $D(g||f)$ and $D(f||g)/(M-1)$ determine, respectively, the rates at which the state of the cell with the target and the states of the $M-1$ cells without the target can be accurately inferred. Based on the order of these two rates, the DGF policy aims at identifying either the cell with the target or those $M-1$ cells without the target. The selection rule is thus clear by noticing that searching the cell with the second largest sum LLRs will lead to sufficient exploration of all $M-1$ cells without the target since the less explored cells tend to have higher sum LLRs among these $M-1$ cells. A more detailed discussion of the DGF policy and a rigorous proof of its asymptotic optimality are given in Section III.

We then extend the problem to the case where multiple anomalous processes are present. In particular, we examine the case where only an upper bound on the number of anomalous processes is known. Interestingly, we show that the Chernoff test may not be practically appealing under the latter setting. We thus consider a modified Bayes risk that better captures the design objective of practical systems and

develop a deterministic policy that is again asymptotically optimal.

C. Related Work

Chernoff's pioneer work on active hypothesis testing focuses on sequential binary composite hypothesis testing [1]. The extension to M-ary hypothesis was given by Bessler in [3]. In [5], Nitinawarat et al. considered M-ary active hypothesis testing in both fixed sample size and sequential settings. Under the sequential setting, they developed a modified Chernoff test that is asymptotically optimal without the positivity assumption on all KL divergences as required in [1] and [3]. Furthermore, they examined the asymptotic optimality of the Chernoff test under constraints on decision risks, a stronger condition than the error probability, and developed a modified Chernoff test to meet hard constraints on the decision risks. In [6], a more general model of Markovian Observations and non-uniform control cost was considered. In [7], in addition to the asymptotic optimality adopted by Chernoff in [1], Naghshvar and Javidi examined active sequential hypothesis testing under the notion of non-zero information acquisition rate by letting the number of hypotheses approach infinity and under a stronger notion of asymptotic optimality. They further studied in [8] the roles of sequentiality and adaptivity in active hypothesis testing by characterizing the gain of sequential tests over fixed sample size tests and the gain of closed-loop policies over open-loop policies.

Target search or target whereabouts problems have been widely studied under various scenarios. Results under the sequential setting can be found in [9]–[12], all assuming single process observations (i.e., $K = 1$). Specifically, optimal policies were derived in [9]–[11] for the problem of quickest search over Weiner processes. In [12], an optimal search strategy was established under the constraint that switching to a new process is allowed only when the state of the currently probed process is declared. Optimal policies under general distributions or with general multi-process probing strategies remain an open question. In this paper we address these questions under the asymptotic regime as the error probability approaches zero. Target search with a fixed sample size was considered in [13]–[16]. In [13]–[15], searching in a specific location provides a binary-valued measurement regarding the presence or absence of the target. Similar to this paper, Castanon considered in [16] continuous observations: the observations from a location without the target and with the target have distributions f and g , respectively. Different from this paper where we consider sequential settings and obtain an asymptotically optimal policy that applies to general distributions, [16] focused on the fixed sample size setting and required a symmetry assumption on the distributions (specifically, $f(x) = g(b-x)$ for some b) for the optimality of the proposed policy. The problem of universal outlier hypothesis testing was studied in [17]. Under this setting, a vector of observations containing coordinates with an outlier distribution is observed at each given time. The goal is to detect the coordinates with the outlier distribution based on a sequence of n i.i.d. vectors of observations.

Another set of related work is concerned with sequential detection over multiple independent processes [18]–[32]. In particular, in [22], the problem of identifying the first abnormal sequence among an infinite number of i.i.d. sequences was considered. An optimal cumulative sum (CUSUM) test was established under this setting. Further studies on this model can be found in [23]–[25]. While the objective of finding rare events or a single target considered in [22]–[25] is similar to that of this paper, the main difference is that in [22]–[25] the search is done over an infinite number of i.i.d. processes, where the state of each process (normal or abnormal) is independent of other processes. Under this independence assumption, the structure of the solution is to perform an independent sequential test without memory for each process. At each time when the decision maker decides to switch to a different process, the new process is chosen arbitrarily, and a sequential test starts afresh. In this paper, however, the number of the processes is finite and the number of the abnormal ones is known (or an upper bound is known). As a result, the process states are correlated. Under this model, the selection rule that governs which process to observe at each time is crucial in minimizing the detection delay, whereas in [22]–[25] the order at which the processes are observed is irrelevant. Furthermore, in our model, the sequential tests for the processes have memory. When a process is revisited, all the observations obtained during previous visits are taken into consideration in decision making.

Another related problem considered recently deals with detecting the first disorder of a system involving multiple processes [19], [21], [26], [29]. In this problem, multiple sensors take observations sequentially from the environment and communicate with a fusion center, which determines whether there is a change in the statistical behavior of the observations. The asymptotic optimality of the multi-chart CUSUMs in detecting the first change-point was studied as the mean time between false alarms approaches to infinity. In [19], asymptotic optimality was shown under one-shot schemes, in which the sensors communicate with the fusion center only when they signal an alarm. A Bayesian version of this problem was considered in [21] under the assumption that the fusion center has perfect information about the observations and a priori knowledge of the statistics of the change process. In [26], the problem was examined for the case where an unknown subset of sensors are compromised and a fully distributed low complexity detection scheme was proposed to mitigate the performance degradation and recover the log scaling. In [29], asymptotic optimality of the multi-chart CUSUMs was shown under a coupled system, where observations in one sensor can affect the observations in another. In this paper, however, the goal is to detect the abnormal processes (and not a change point), where the process states are fixed during the detection process.

D. Organization

In Section II we describe the system model and problem formulation. In Section III we propose the deterministic DGF policy and establish its asymptotic optimality. We also

provide a comparison of DGF with the randomized Chernoff test. In Section IV we extend the problem to the case where multiple anomalous processes are present and consider both cases of known and unknown number of anomalous processes. In Section V we provide numerical examples to illustrate the performance of the proposed policy as compared with the Chernoff test. Section VI concludes the paper.

II. SYSTEM MODEL AND PROBLEM FORMULATION

A. System Model

Consider the following anomaly detection problem. A decision maker is required to detect the location of a single anomalous object (referred as a target) located in one of M cells. If the target is in cell m , we say that hypothesis H_m is true. The *a priori* probability that H_m is true is denoted by π_m , where $\sum_{m=1}^M \pi_m = 1$. To avoid trivial solutions, it is assumed that $0 < \pi_m < 1$ for all m .

At each time, only K ($1 \leq K \leq M$) cells can be observed. When cell m is observed at time n , an observation $y_m(n)$ is drawn independently from a distribution in a one-at-a-time manner. If hypothesis m is false, $y_m(n)$ follows distribution $f(y)$; if hypothesis m is true, $y_m(n)$ follows distribution $g(y)$. Let \mathbf{P}_m be the probability measure under hypothesis H_m and \mathbf{E}_m the operator of expectation with respect to the measure \mathbf{P}_m .

We define the stopping rule τ as the time when the decision maker finalizes the search by declaring the location of the target. Let $\delta \in \{1, 2, \dots, M\}$ be a decision rule, where $\delta = m$ if the decision maker declares that H_m is true. Let $\phi(n) \in \{1, 2, \dots, M\}^K$ be a selection rule indicating which K cells are chosen to be observed at time n . The time series vector of selection rules is denoted by $\boldsymbol{\phi} = (\phi(n), n = 1, 2, \dots)$. Let $\mathbf{y}_{\phi(n)}(n)$ be the vector of observations obtain from cells $\phi(n)$ at time n and $\mathbf{y}(n) = \{\phi(t), \mathbf{y}_{\phi(t)}(t)\}_{t=1}^n$ be the set of all cell selections and observations up to time n . A deterministic selection rule $\phi(n)$ at time n is a mapping from $\mathbf{y}(n-1)$ to $\{1, 2, \dots, M\}^K$. A randomized selection rule $\phi(n)$ is a mapping from $\mathbf{y}(n-1)$ to probability mass functions over $\{1, 2, \dots, M\}^K$.

Definition 1: An admissible strategy Γ for the sequential anomaly detection problem is given by the tuple $\Gamma = (\tau, \delta, \boldsymbol{\phi})$.

B. Objective

Let $P_e(\Gamma) = \sum_{m=1}^M \pi_m \alpha_m(\Gamma)$ be the probability of error under strategy Γ , where $\alpha_m(\Gamma) = \mathbf{P}_m(\delta \neq m | \Gamma)$ is the probability of declaring $\delta \neq m$ when H_m is true. Let $\mathbf{E}(\tau | \Gamma) = \sum_{m=1}^M \pi_m \mathbf{E}_m(\tau | \Gamma)$ be the average detection delay under Γ .

We adopt a Bayesian approach as in [1] and [4] by assigning a cost of c for each observation and a loss of 1 for a wrong declaration. The Bayes risk under strategy Γ when hypothesis H_m is true is given by:

$$R_m(\Gamma) \triangleq \alpha_m(\Gamma) + c \mathbf{E}_m(\tau | \Gamma). \quad (1)$$

Note that c represents the ratio of the sampling cost to the cost of wrong detections.

The average Bayes risk is given by:

$$R(\Gamma) = \sum_{m=1}^M \pi_m R_m(\Gamma) = P_e(\Gamma) + c \mathbf{E}(\tau|\Gamma). \quad (2)$$

The objective is to find a strategy Γ that minimizes the Bayes risk $R(\Gamma)$:

$$\inf_{\Gamma} R(\Gamma). \quad (3)$$

C. Notations

Let $\mathbf{1}_m(n)$ be the indicator function, where $\mathbf{1}_m(n) = 1$ if cell m is observed at time n , and $\mathbf{1}_m(n) = 0$ otherwise. Let

$$\ell_m(n) \triangleq \log \frac{g(y_m(n))}{f(y_m(n))}, \quad (4)$$

and

$$S_m(n) \triangleq \sum_{t=1}^n \ell_m(t) \mathbf{1}_m(t) \quad (5)$$

be the log-likelihood ratio (LLR) and the observed sum LLRs of cell m at time n , respectively. We then define $m^{(i)}(n)$ as the index of the cell with the i^{th} highest observed sum LLRs at time n . Let

$$\Delta S(n) \triangleq S_{m^{(1)}(n)}(n) - S_{m^{(2)}(n)}(n) \quad (6)$$

denote the difference between the highest and the second highest observed sum LLRs at time n .

Finally, we define

$$I^*(M, K) \triangleq \begin{cases} D(g||f) + D(f||g), & \text{if } K = M, \\ \max \left[\frac{KD(f||g)}{M-1}, D(g||f) + \frac{(K-1)D(f||g)}{M-1} \right], & \text{if } K < M. \end{cases} \quad (7)$$

In subsequent sections we show that $I^*(M, K)$ plays the role of the rate function, which determines the asymptotically optimal performance of the test. Increasing $I^*(M, K)$ decreases the asymptotic lower bound on the Bayes risk. It is intuitive that $I^*(M, K)$ increases with the observation capability K and decreases with the hypothesis size M .

III. THE DETERMINISTIC DGF POLICY

In this section we propose a deterministic policy, referred to as the DGF policy, to solve (3). Theorem 1 shows that the DGF policy is asymptotically optimal in terms of minimizing the Bayes risk (2) as $c \rightarrow 0$.

A. The DGF Policy

At each time n , the selection rule $\phi(n)$ of the DGF policy chooses cells according to the order of their sum LLRs. Specifically, based on the relative order of $D(g||f)$ and $D(f||g)/(M-1)$, either the K cells with the top K highest sum LLRs or those with the second to the $(K+1)^{\text{th}}$ highest

sum LLRs are chosen, i.e.,¹

$$\phi(n) = \begin{cases} (m^{(1)}(n), m^{(2)}(n), \dots, m^{(K)}(n)), & \text{if } D(g||f) \geq \frac{D(f||g)}{(M-1)} \text{ or } K = M \\ (m^{(2)}(n), m^{(3)}(n), \dots, m^{(K+1)}(n)), & \text{if } D(g||f) < \frac{D(f||g)}{(M-1)} \text{ and } K < M \end{cases} \quad (8)$$

The stopping rule and decision rule under the DGF policy are given by:

$$\tau = \inf \{n : \Delta S(n) \geq -\log c\}, \quad (9)$$

and

$$\delta = m^{(1)}(\tau). \quad (10)$$

The deterministic selection rule of the DGF policy can be intuitively explained as follows. Consider the case where $K = 1$. If cell $m^{(1)}(n)$ is selected at each given time n , the asymptotic detection time approaches $-\log c/D(g||f)$ since the cell with the target (say m) is observed at each given time with high probability (in the asymptotic regime) and the test is finalized once sufficient information is gathered from this cell (for a detailed asymptotic analysis see Appendix A). In this case, $D(g||f)$ determines the asymptotically optimal performance of the test since $\mathbf{E}_m(\ell_m) = D(g||f)$. On the other hand, if cell $m^{(2)}(n)$ is selected at each given time n , the asymptotic detection time approaches $-(M-1) \log c/D(f||g)$ since one of the $M-1$ cells without the target is observed at each given time with high probability and the test is finalized once sufficient information is gathered from all these cells. Since $\mathbf{E}_m(\ell_j) = -D(f||g)$ for all $j \neq m$, the asymptotically optimal performance of the test is determined by $D(f||g)/(M-1)$. Therefore, the selection rule selects the strategy that minimizes the asymptotic detection time according to $\max [D(g||f), D(f||g)/(M-1)]$. When $K > 1$, the rates at which the state of cell m and the states of the rest $M-1$ cells can be accurately inferred are given by $D(g||f) + \frac{(K-1)D(f||g)}{M-1}$ and $\frac{KD(f||g)}{M-1}$, respectively. Since $D(g||f) > D(f||g)/(M-1)$ is equivalent to $D(g||f) + \frac{(K-1)D(f||g)}{M-1} > \frac{KD(f||g)}{M-1}$, the selection rule of DGF is thus clear.

B. Performance Analysis

The following main theorem shows that the DGF policy is asymptotically optimal in terms of minimizing the Bayes risk as c approaches zero:

Theorem 1 (Asymptotic Optimality of the DGF Policy): Let R^* and $R(\Gamma)$ be the Bayes risks under the DGF policy and any other policy Γ , respectively. Then,²

$$R^* \sim \frac{-c \log c}{I^*(M, K)} \sim \inf_{\Gamma} R(\Gamma) \text{ as } c \rightarrow 0. \quad (11)$$

¹Cells with the same sum LLRs can be ordered arbitrarily.

²The notation $f \sim g$ as $c \rightarrow 0$ refers to $\lim_{c \rightarrow 0} f/g = 1$.

Proof: For a detailed proof see Appendix A. We provide here a sketch of the proof. In App. A.1, we show that $\frac{-c \log c}{I^*(M, K)}$ is an asymptotic lower bound on the achievable Bayes risk. Then, we show in App. A.2 that the Bayes risk R^* under the DGF policy approaches the asymptotic lower bound as $c \rightarrow 0$. Specifically, the asymptotic behavior of R^* is established based on Lemma 11 showing that the asymptotic expected detection time approaches $\frac{-\log c}{I^*(M, K)}$, while the error probability is $O(c)$ following Lemma 5.

The basic idea in establishing the asymptotic expected detection time under DGF in Lemma 11 is to upper bound the stopping time τ of DGF by analyzing three last passage times (given in Lemmas 7, 8 and 10). Specifically, if the stopping rule is disregarded and sampling is continued indefinitely, then three last passage times can be defined: τ_1, τ_2, τ_3 , where, roughly speaking, τ_1 is the time when the sum LLRs of the true cell (say m) is the highest among all the cells for all $n \geq \tau_1$; τ_2 is the time when sufficient information for distinguishing hypothesis m from at least one false hypothesis has been gathered; τ_3 is the time when sufficient information for distinguishing hypothesis m from all false hypotheses has been gathered. It should be noted that τ_1, τ_2, τ_3 are not stopping times and the decision maker does not know whether they have arrived (since the true cell is unknown and also τ_1, τ_2, τ_3 depend on the future by definition). However, by the definition of τ_3 (see Definition 7 in Appendix A for details) the actual stopping time τ under DGF is upper bounded by τ_3 (i.e., the decision maker does know that for all $n < \tau$, τ_3 surely has not arrived). As a result, $\mathbf{E}(\tau_3)$ is an upper bound of $\mathbf{E}(\tau)$.

To show the asymptotic behavior of $\mathbf{E}(\tau_3)$, define $n_2 = \tau_2 - \tau_1$ and $n_3 = \tau_3 - \tau_2$. Thus, $\tau_3 = \tau_1 + n_2 + n_3$. Lemma 8 shows that $\mathbf{E}(n_2) \sim -\log c / I^*(M, K)$ as $c \rightarrow 0$. Lemma 7 shows that $\mathbf{E}(\tau_1) / \mathbf{E}(n_2) \rightarrow 0$, i.e., τ_1 does not affect the asymptotic detection time. Note that differing from [5], where only polynomial decay of $\mathbf{P}_m(\tau_1 > n)$ was shown under the extended Chernoff test developed to handle indistinguishable hypotheses under some actions, Lemma 7 shows exponential decay of $\mathbf{P}_m(\tau_1 > n)$ under DGF. Lemma 10 shows that $\mathbf{E}(n_3) / \mathbf{E}(n_2) \rightarrow 0$. Combining Lemmas 7, 8 and 10, we can conclude that $\mathbf{E}(\tau_3) \sim -\log c / I^*(M, K)$. Since the error probability is $O(c)$ following Lemma 5, the proof thus completes by noticing that the upper bound on $c\mathbf{E}(\tau) + P_e$ coincides with the lower bound on the achievable Bayes risk. ■

C. Comparison With the Chernoff Test

Next, we analyze the classic randomized Chernoff test proposed in [1] when it is applied to the anomaly detection problem. We then compare the performance of the proposed DGF policy with the Chernoff test.

1) *The Chernoff Test:* The Chernoff test has a randomized selection rule. Specifically, let $q = (q_1, \dots, q_N)$ be a probability mass function over a set of N available experiments $u = \{u_i\}_{i=1}^N$ that the decision maker can choose from, where q_i is the probability of choosing experiment u_i . For a general M-ary active hypothesis testing problem, the action

at time n under the Chernoff test is drawn from a distribution $q^*(n) = (q_1^*(n), \dots, q_N^*(n))$ that depends on the past actions and observations:

$$q^*(n) = \arg \max_q \min_{j \in \mathcal{M} \setminus \{\hat{i}(n)\}} \sum_{u_i} q_i D(p_{i(n)}^{u_i} \| p_j^{u_i}), \quad (12)$$

where \mathcal{M} is the set of the M hypotheses, $\hat{i}(n)$ is the ML estimate of the true hypothesis at time n based on past actions and observations, and $p_j^{u_i}$ is the observation distribution under hypothesis j when action u_i is taken. The stopping rule and decision rule are given in (9), (10).

It can be shown that when applied to the anomaly detection problem, the Chernoff test works as follows. When $D(g||f) \geq D(f||g)/(M-1)$, the Chernoff test selects cell $m^{(1)}(n)$ and draws the rest $K-1$ cells randomly with equal probability from the remaining $M-1$ cells. When $D(g||f) < D(f||g)/(M-1)$, all K cells are drawn randomly with equal probability from cells $\{m^{(2)}(n), m^{(3)}(n), \dots, m^{(M)}(n)\}$ under the Chernoff test.

Even though the positivity assumption on KL divergences as required in the proof of the asymptotic optimality of the Chernoff test given in [1] no longer holds for the anomaly detection problem, we show in Theorem 2 below that the Chernoff test preserves its asymptotic optimality in this case. Note that in [5], a modified Chernoff test was developed in order to handle indistinguishable hypotheses under some (but not all) actions. The basic idea of the modified test is to replace the action distribution given in (12) with a uniform distribution for a subsequence of time instants that grows at a sublinear rate with time. This subsequence of arbitrary actions are independent of past observations and affects the finite-time performance. In Theorem 2 below we show that this modification is unnecessary for the anomaly detection problem.

Theorem 2: Let R_{CT} and $R(\Gamma)$ be the Bayes risks under the Chernoff test and any other policy Γ , respectively. Then,

$$R_{CT} \sim \frac{-c \log c}{I^*(M, K)} \sim \inf_{\Gamma} R(\Gamma) \text{ as } c \rightarrow 0. \quad (13)$$

Proof: The proof is given in Appendix B and is based on the argument of [5] and the proof of Theorem 1 given in Appendix A. ■

2) *Comparison:* Although both the Chernoff test and the DGF policy are asymptotically optimal, simulation results demonstrate significant performance gain of DGF over the Chernoff test in the finite regime (see Section V). Next, we provide an intuition argument for the better finite-time performance of DGF by drawing an analogy between the anomaly detection problem and the makespan scheduling problem.

Consider the problem of scheduling M jobs over K parallel machines ($K \leq M$). Each job requires a deterministic processing time of T_p time units. The objective is to minimize the makespan which is defined as the completion time of all M jobs. Note that when $K > 1$, processing a job continuously until it is completed can be highly suboptimal since a certain number of machines are left idle when there are less than K unfinished jobs. Note also that keeping machines idle during the scheduling process increases the makespan for all $K \geq 1$. The optimal solution to this problem is given

by the LRPT (the longest remaining processing time first) scheduler [33, Th. 5.2.7] that schedules, at any time n , the K jobs with the longest remaining processing time.

The anomaly detection problem can be viewed as a problem of scheduling $M - 1$ jobs (each being the detection process of distinguishing one of the $M - 1$ false hypotheses from the true hypothesis) over K machines (which is the number of cells that the decision maker can probe simultaneously). Consider first $D(g||f) < D(f||g)/(M - 1)$. In this case, DGF probes cells $(m^{(2)}(n), m^{(3)}(n), \dots, m^{(K+1)}(n))$ at each time, while the Chernoff test selects K cells randomly among the cells $(m^{(2)}(n), m^{(3)}(n), \dots, m^{(M)}(n))$. Both tests terminate once $\Delta S(n) \geq -\log c$ occurs. Assume that hypothesis H_m is true. Roughly speaking, following Lemma 5, once $\Delta S_{m,j} \triangleq S_m(n) - S_j(n) > -\log c$, the decision maker has sufficient evidence to distinguish false hypothesis H_j from the true hypothesis H_m . Except during an asymptotically insignificant initial stage of the detection process, cells $(m^{(2)}, \dots, m^{(M)}(n))$ are the cells without the target (see Lemma 7 for a detailed analysis on the last passage time τ_1 of cell $m^{(1)}(n)$ being the cell with the target for all $n \geq \tau_1$). In this case, cells $(m^{(2)}, \dots, m^{(K+1)}(n))$ as selected by DGF can be viewed as the cells with the longest remaining processing times. The randomized Chernoff test, however, may lead to inefficient exploitation of the probing capacity, as explain above for the makespan scheduling problem. Furthermore, randomly selecting K cells from $(m^{(2)}, \dots, m^{(M)}(n))$ may result in probing a cell whose state can already be inferred with sufficient accuracy (i.e., $\Delta S_{m,j} > -\log c$ as detailed in Appendix A), which can be viewed as scheduling a job that is already completed or equivalently, leaving a machine idle in the makespan problem. Such actions, however, will not occur under DGF. The argument for the case of $D(g||f) > D(f||g)/(M - 1)$ is similar by viewing the problem as scheduling $M - 1$ jobs over $K - 1$ machines. Note that both DGF and the Chernoff test dedicate one machine for probing cell $m^{(1)}(n)$ since under the condition of $D(g||f) > D(f||g)/(M - 1)$, probing the cell with the target is preferred to accelerate the detection process.

IV. EXTENSION TO MULTIPLE ANOMALOUS PROCESSES

In this section we extend the results reported in previous sections to the case where multiple processes are abnormal. In Section IV-A we consider the detection of L abnormal processes, where L is known. In Section IV-B we consider the case where an unknown number $\ell \geq 1$ of abnormal processes are present and only an upper bound $\ell \leq L$ is known.

Throughout this section, we define \mathcal{M}' as the set of all possible combinations of target locations, with cardinality $M' = |\mathcal{M}'|$ (i.e., a set of M' hypotheses, $H_{m'}$, indicating that the locations of all targets are given by the $(m')^{th}$ set in \mathcal{M}') and $\pi_{m'}$ as the *a priori* probability that $H_{m'}$ is true. Here, the decision rule declares a set of target locations (i.e., hypothesis $H_{m'}$) and the error probability under policy Γ is defined as $P_e(\Gamma) = \sum_{m'=1}^{M'} \pi_{m'} \alpha_{m'}(\Gamma)$, where $\alpha_{m'}(\Gamma) = \mathbf{P}_{m'}(\delta \neq H_{m'} | \Gamma)$ is the probability of declaring $\delta \neq H_{m'}$ when $H_{m'}$ is true.

A. Known Number of Abnormal Processes

Consider the case where L abnormal processes are located among the M cells and L is known. In this case, the detection problem involves $M' = \binom{M}{L}$ hypotheses. We show below that a variation of the DGF policy, dubbed the DGF(L) policy, is asymptotically optimal under this setting.

The stopping rule and decision rule under the DGF(L) policy are similar to that under the DGF policy:

$$\tau = \inf \{n : \Delta_L S(n) \geq -\log c\}, \quad (14)$$

where $\Delta_L S(n) \triangleq S_{m^{(L)}(n)}(n) - S_{m^{(L+1)}(n)}(n)$ and

$$\delta = (m^{(1)}(\tau), m^{(2)}(\tau), \dots, m^{(L)}(\tau)). \quad (15)$$

The selection rule under the DGF(L) policy is more involved and depends on the relative order of K and L (or $M - L$). Specifically,

$$\phi(n) = \begin{cases} \phi_g(n), & \text{if } \frac{D(g||f)}{L} \geq \frac{D(f||g)}{M-L}, \\ \phi_f(n), & \text{if } \frac{D(g||f)}{L} < \frac{D(f||g)}{M-L}, \end{cases} \quad (16)$$

where

$$\phi_g(n) = \begin{cases} (m^{(1)}(n), m^{(2)}(n), \dots, m^{(K)}(n)), & \text{if } K \geq L, \\ (m^{(L-K+1)}(n), m^{(L-K+2)}(n), \dots, m^{(L)}(n)), & \text{if } K < L, \end{cases} \quad (17)$$

and

$$\phi_f(n) = \begin{cases} (m^{(M-K+1)}(n), m^{(M-K+2)}(n), \dots, m^{(M)}(n)), & \text{if } K > M - L, \\ (m^{(L+1)}(n), m^{(L+2)}(n), \dots, m^{(L+K)}(n)), & \text{if } K \leq M - L. \end{cases} \quad (18)$$

It is not difficult to see that when $L = 1$, the DGF(L) policy degenerates to the DGF policy.

Next, we analyze the performance of the DGF(L) policy. Let

$$I^*(M, K, L) \triangleq \begin{cases} I_g^*(M, K, L), & \text{if } \frac{D(g||f)}{L} \geq \frac{D(f||g)}{M-L}, \\ I_f^*(M, K, L), & \text{if } \frac{D(g||f)}{L} < \frac{D(f||g)}{M-L}, \end{cases} \quad (19)$$

where

$$I_g^*(M, K, L) \triangleq \begin{cases} D(g||f) + \frac{(K-L)D(f||g)}{M-L}, & \text{if } K \geq L, \\ \frac{KD(g||f)}{L}, & \text{if } K < L, \end{cases} \quad (20)$$

and

$$I_f^*(M, K, L) \triangleq \begin{cases} D(f||g) + \frac{(K-M+L)D(g||f)}{L}, & \text{if } K > M-L, \\ \frac{KD(f||g)}{M-L}, & \text{if } K \leq M-L. \end{cases} \quad (21)$$

The following theorem shows the asymptotically optimal performance of the DGF(L) policy:

Theorem 3: Let R^* and $R(\Gamma)$ be the Bayes risks under the DGF(L) policy and any other policy Γ , respectively. Then,

$$R^* \sim \frac{-c \log c}{I^*(M, K, L)} \sim \inf_{\Gamma} R(\Gamma) \quad \text{as } c \rightarrow 0. \quad (22)$$

Proof: See Appendix C. ■

Note that in the DGF(L) policy, all L targets are declared simultaneously at the termination time of the detection. A modification to DGF(L) leads to a policy where abnormal processes are declared sequentially during the detection. Consider, for example, $K = 1$ and $\frac{D(g||f)}{L} \geq \frac{D(f||g)}{M-L}$. It can be shown (with minor modifications to Theorem 3) that an asymptotically optimal policy is to test the cell with the largest sum LLRs and declare the first target once the largest sum LLRs exceeds the threshold $-\log c$. The same procedure is then applied to the remaining $M - 1$ cells. This repeats until L abnormal processes have been declared, at which point, the detection terminates. The asymptotic expected termination time is given by $-L \log c / D(g||f)$ with $P_e = O(c)$. Even though the total detection time remains the same order as under the DGF(L) policy, this modified version may be more appealing from a practical point of view. In particular, actions can be taken to fix each abnormal process the moment it is identified; the total impact to the system by these L abnormal processes can thus be reduced. If $\frac{D(g||f)}{L} < \frac{D(f||g)}{M-L}$, it can be shown that an asymptotically optimal policy is to test the cell with the smallest sum LLRs and declare the first *normal* process once the smallest sum LLRs drops below $\log c$. The same procedure is then applied to the remaining $M - 1$ processes and is repeated until all $M - L$ objects are declared as normal (thus, the L remaining ones are declared as abnormal). The asymptotic expected termination time is given by $-(M - L) \log c / D(f||g)$ with $P_e = O(c)$. Even though in this case, the modified version also declares all L targets simultaneously at the termination time of the detection, the difference is that this modified version incurs much few switchings among processes than the DGF(L) policy. This may be more advantageous in some practical scenarios when switching among tested processes results in additional cost or delay. To see that the modified version incurs few switchings, we note that the modified version tests the process that the decision maker is most sure to be normal based on past observations while DGF(L) tests the process that the decision maker is least sure to be normal except the L processes currently considered as the targets (see the second line in (18) which shows that DGF(L) chooses the cell with the $(L + 1)^{th}$ largest sum LLRs; the L processes with larger sum LLRs are the current maximum likelihood of the target locations). It should be noted that those modified DGF(L) schemes are expected to achieve the same performance as DGF(L) in both the finite and asymptotic regimes when $K = 1$ following a similar argument as in Section III-C.

B. Unknown Number of Abnormal Processes

In this section we consider the interesting case where the number ℓ of abnormal processes (or targets) is unknown. It is

only known that ℓ is bounded by $1 \leq \ell \leq L$. We consider the case where $K = 1$. We also assume that the number of cells satisfies:

$$M \geq \frac{L(D(g||f) + D(f||g))}{D(g||f)}. \quad (23)$$

Note that (23) implies $\frac{D(g||f)}{L} \geq \frac{D(f||g)}{M-L}$.

Throughout this section, we allow the decision maker to declare the target locations sequentially during the test (similar to the modified DGF(L) policy as discussed at the end of Section IV-A). We refer to the detection time τ_d as the time when the last target has been declared and to the termination time τ as the time when the decision maker terminates the test. Note that $\tau = \tau_d$ when the number ℓ of targets is known (as discussed in previous sections). When ℓ is unknown, however, $\tau_d \leq \tau$ since the decision maker does not know whether it has already identified all targets at time τ_d . In general, the termination time τ increases linearly with M under any policy with $P_e = O(c)$ whenever $\ell < L$. This is due to the fact that even if the ℓ targets have been detected with sufficient reliability, the decision maker must verify whether there are other targets in the remaining $M - \ell$ cells before terminating the test. On the other hand, following the modified DGF(L) policy, one would expect to achieve a detection time τ_d less than $-L \log c / D(g||f)$ for all $\ell \leq L$, which is independent of the number M of total processes.

In scenarios with a large number of processes and $L \ll M$, a policy that focuses on minimizing the termination time τ , which grows linearly with M , may not be practically appealing. It is desirable to have a policy that allows each abnormal process to be identified and fixed as quickly as possible *during* the test. In other words, it is desirable to have a policy that minimizes the detection time τ_d rather than the termination time τ . In this case, even though the test continues after τ_d to ensure there are no other targets, all abnormal processes have been fixed by the detection time τ_d and cease to incur cost to the system. We thus modify the objective function to the following Bayes risk:

$$R(\Gamma) \triangleq P_e(\Gamma) + c\mathbf{E}(\tau_d|\Gamma), \quad (24)$$

and we are interested in finding a strategy Γ that minimizes the Bayes risk (24). This design objective is similar to that considered in [28], [30], and [31].

Before presenting the desired solution for this case, we demonstrate with a specific example that even though the Chernoff test is asymptotically optimal in terms of minimizing the termination time τ , it is highly suboptimal in terms of minimizing the detection time τ_d . Assume that $L = 2$ is the upper bound on the number of targets, which can locate in any of $M = 3$ cells. As a result, the detection problem includes 6 hypotheses, $H_1 = \{1\}$, $H_2 = \{2\}$, $H_3 = \{3\}$, $H_4 = \{1, 2\}$, $H_5 = \{1, 3\}$, $H_6 = \{2, 3\}$. The observation model under every hypothesis and cell selection is given in Table I. Assume that hypothesis H_1 is true and that $\hat{H}(n) = H_1$, where $\hat{H}(n)$ is the ML estimate of the true hypothesis at time n . Consider a deterministic policy that selects the cells according to the order of their sum LLRs and declares an object as target if $S_m(n) > -\log c$ or normal

TABLE I
 OBSERVATION MODEL

	cell 1	cell 2	cell 3
$H_1 = \{1\}$	g	f	f
$H_2 = \{2\}$	f	g	f
$H_3 = \{3\}$	f	f	g
$H_4 = \{1, 2\}$	g	g	f
$H_5 = \{1, 3\}$	g	f	g
$H_6 = \{2, 3\}$	f	g	g

if $S_m(n) < \log c$. This policy achieves $\tau_d \sim -\log c/D(g||f)$ (since cell 1 is first identified as a target with high probability) and $\tau \sim -\log c/D(g||f) - 2\log c/D(f||g)$ (since the number of targets is unknown and $L = 2$, thus the decision maker must continue testing the normal processes before terminating the test). On the other hand, the Chernoff test (which aims to minimize the termination time) will not select cell 1 at time n , since H_4 or H_5 minimizes (12), i.e., $D(p_{H_1}^1||p_{H_4}^1) = D(p_{H_1}^1||p_{H_5}^1) = D(g||g) = 0$. It can be verified that selecting randomly cells 2 or 3 with equal probability $1/(M-1) = 1/2$ maximizes (12) and achieves a rate function $D(f||g)/(M-1) = D(f||g)/2$, which results in $\tau = \tau_d \sim -2\log c/D(f||g)$, which is greater than the detection time under the above deterministic policy. Intuitively speaking, once cells 2,3 are identified as normal, cell 1 is identified as abnormal (because at least 1 target is present). Therefore, the Chernoff test observes cells 2,3 to minimize the termination time τ (by not testing cell 1), while increasing the detection time τ_d .

Next, we present a deterministic policy to minimize the Bayes risk (24). Let $\mathcal{T}(n)$ be the set of cells satisfying $S_m(n) \geq -\log c$ at time n . Define

$$\tilde{m}^{(1)}(n) = \arg \max_{m \notin \mathcal{T}(n)} S_m(n). \quad (25)$$

The selection rule is given by:

$$\phi(n) = \tilde{m}^{(1)}(n). \quad (26)$$

The stopping rule and decision rule are given by:

$$\tau = \inf \{n : |S_m(n)| \geq -\log c \ \forall m\}, \quad (27)$$

and

$$\delta = \mathcal{T}(\tau_d). \quad (28)$$

Note that δ denotes the target locations and the complete set is declared at time τ_d . Since the number of targets is unknown, the decision maker continues taking observations to verify that there is no other target. The test is terminated at time τ .

The following theorem shows the asymptotically optimal performance of the proposed policy:

Theorem 4: Let $\ell \leq L$ be the number of targets, $K = 1$ and assume that (23) holds. Let R^* and $R(\Gamma)$ be the Bayes risks (24) under the proposed policy and any other policy Γ , respectively. Then,

$$R^* \sim \frac{-\ell c \log c}{D(g||f)} \sim \inf_{\Gamma} R(\Gamma) \text{ as } c \rightarrow 0. \quad (29)$$

Proof: See Appendix D. ■


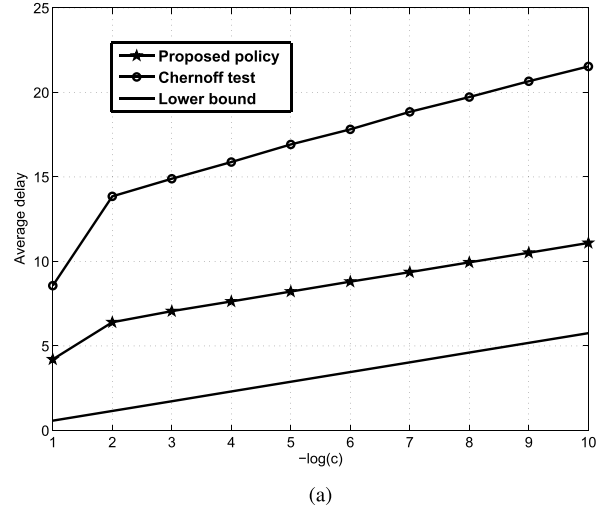


Fig. 1. Performance comparison for $M = 5$, $K = 1$, $\lambda_f = 0.5$, $\lambda_g = 10$. (a) Average sample sizes and the asymptotic lower bound as a function of the cost per observation. (b) The loss in terms of Bayes risk under the DGF policy and the Chernoff test as compared to the asymptotic lower bound. L_{DGF} , L_{Ch} approach 0 as $c \rightarrow 0$.

V. NUMERICAL EXAMPLES

In this section we present numerical examples to illustrate the performance of the proposed deterministic policy as compared to the Chernoff test. We simulated a single anomalous object (i.e., target) located in one of M cells with the following parameters: The *a priori* probability that the target is present in cell m was set to $\pi_m = 1/M$ for all $1 \leq m \leq M$. When cell m is observed at time n , an observation $y_m(n)$ is independently drawn from a distribution $f \sim \exp(\lambda_f)$ or $g \sim \exp(\lambda_g)$, depending on whether the target is absent or present, respectively. It can be verified that:

$$D(g||f) = \log(\lambda_g) - \log(\lambda_f) + \frac{\lambda_f}{\lambda_g} - 1,$$

$$D(f||g) = \log(\lambda_f) - \log(\lambda_g) + \frac{\lambda_g}{\lambda_f} - 1.$$

Let R_{DGF} , R_{Ch} be the Bayes risks under the DGF policy and the Chernoff test, respectively. Let $R_{LB} = \frac{-c \log c}{T^*(M,K)}$ be

TABLE II
VALUES FOR 95% CONFIDENCE LEVEL

$-\log c$	DGF	Chernoff Test
1	$\sigma = 1.84, r = 1.60$	$\sigma = 6.75, r = 1.90$
3	$\sigma = 2.00, r = 1.80$	$\sigma = 7.17, r = 1.95$
5	$\sigma = 2.24, r = 1.90$	$\sigma = 7.75, r = 1.95$


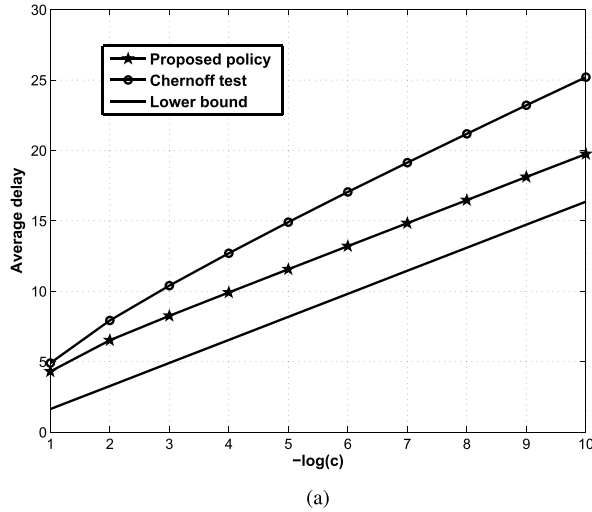


Fig. 2. Performance comparison for $M = 5, K = 2, \lambda_f = 2, \lambda_g = 10$. (a) Average sample sizes and the asymptotic lower bound as a function of the cost per observation. (b) The loss in terms of Bayes risk under the DGF policy and the Chernoff test as compared to the asymptotic lower bound. L_{DGF}, L_{Ch} approach 0 as $c \rightarrow 0$.

the asymptotic lower bound on the Bayes risk as $c \rightarrow 0$. We define:

$$L_{DGF} \triangleq \frac{R_{DGF} - R_{LB}}{R_{LB}},$$

$$L_{Ch} \triangleq \frac{R_{Ch} - R_{LB}}{R_{LB}}.$$

as the relative loss in terms of Bayes risk under the DGF policy and the Chernoff test, respectively, as compared to the asymptotic lower bound. Following Theorems 2, 1, we expect both L_{DGF} and L_{Ch} to approach 0 as $c \rightarrow 0$. L_{DGF} and L_{Ch} serve as performance measures of the tests in the finite regime.





Fig. 3. Performance comparison for $M = 5, K = 2, \lambda_f = 0.5, \lambda_g = 10$. (a) Average sample sizes and the asymptotic lower bound as a function of the cost per observation. (b) The loss in terms of Bayes risk under the DGF policy and the Chernoff test as compared to the asymptotic lower bound. L_{DGF}, L_{Ch} approach 0 as $c \rightarrow 0$.

First, we consider the case where $M = 5$ and $K = 1$. Note that when $K = 1$ and $D(g||f) \geq D(f||g)/(M - 1)$, the Chernoff test coincides with the DGF policy: they both select cell $m^{(1)}(n)$. When $D(g||f) < D(f||g)/(M - 1)$, however, the proposed policy selects cell $m^{(2)}(n)$, while the Chernoff test selects cell $j \neq m^{(1)}(n)$ randomly at each given time n . We set $\lambda_f = 0.5, \lambda_g = 10$ and obtain $D(g||f) \approx 2.05, D(f||g)/(M - 1) \approx 4$. As a result, the Chernoff test and the DGF policy have different cell selection rules. The performance of the Algorithms is presented in Fig. 1(a), 1(b), where 10^7 trials were performed. In Fig. 1(a), the asymptotic lower bound on the expected sample size and the average sample sizes achieved by the algorithms are presented as a function of c (log-scale). In Table II we present the sample standard deviations σ and the standard deviation multipliers r for a 95% confidence intervals $[\bar{\tau} - r\sigma, \bar{\tau} + r\sigma]$, where $\bar{\tau}$ is the average detection delay. In Fig. 1(b), L_{DGF} and L_{Ch} are presented as a function of c . Although both schemes approach the asymptotic lower bound as $c \rightarrow 0$, it can be seen that the

DGF policy significantly outperforms the Chernoff test in the finite regime for all values of c .

Next, we consider the case where $M = 5$ and $K = 2$ (i.e., two cells are observed at a time). In this case, the DGF policy selects cells $m^{(1)}(n)$ and $m^{(2)}(n)$ at each given time n only if $D(g||f) \geq D(f||g)/(M-1)$. Otherwise, it selects cells $m^{(2)}(n)$ and $m^{(3)}(n)$. The Chernoff test selects cells $m^{(1)}(n)$ and $j \neq m^{(1)}(n)$ (randomly) at each given time n only if $D(g||f) \geq D(f||g)/(M-1)$. Otherwise, it selects cells $i, j \neq m^{(1)}(n)$ randomly. First, we set $\lambda_f = 2$, $\lambda_g = 10$ and obtain $D(g||f) \approx 0.8$, $D(f||g)/(M-1) \approx 0.6$. The performance of the algorithms is presented in Fig. 2(a), 2(b). Next, we set $\lambda_f = 0.5$, $\lambda_g = 10$ and obtain $D(g||f) \approx 2.05$, $D(f||g)/(M-1) \approx 4$. The performance of the algorithms in this case is presented in Fig. 3(a), 3(b). In Fig. 2(a), 3(a), the asymptotic lower bound on the expected sample size and the average sample sizes achieved by the algorithms are presented as a function of the cost per observation c . In Fig. 2(b), 3(b), L_{DGF} and L_{Ch} are presented as a function of c . It can be seen that the DGF policy significantly outperforms the Chernoff test in the finite regime for all values of c under all cases. These results demonstrate the advantage of using the deterministic selection rule applied by the DGF policy instead of the randomized Chernoff test for the anomaly detection problem.

VI. CONCLUSION

The problem of detecting a single anomalous process (i.e., target) among M processes (i.e., cells) was investigated. Due to resource constraints, only a subset of the cells can be observed at a time, The objective is a search strategy that minimizes the expected search time subject to an error probability constraint. The observations from searching a cell are realizations drawn from two different distributions f or g , depending on whether the target is absent or present, respectively. A simple deterministic policy was established to solve the Bayesian formulation of the search problem, where a cost of c per observation and a loss of 1 for wrong decisions are assigned. It is shown that the proposed index policy is asymptotically optimal in terms of minimizing the Bayes risk as c approaches zero.

The problem was further extended to handle the case where multiple anomalous processes are present. In particular, the interesting case where only an upper bound on the number of anomalous processes is known was considered. We showed that existing methods may not be practically appealing under the latter setting. Hence, we proposed a modified optimization problem for this case. Asymptotically optimal deterministic policies were developed for these cases as well.

APPENDIX

A. Proof of Theorem 1

In this appendix we prove the asymptotic optimality of the DGF policy as $c \rightarrow 0$. In App. A.1, we show that $\frac{-c \log c}{I^*(M, K)}$ is an asymptotic lower bound on the Bayes risk that can be achieved by any policy Γ . Then, we show in App. A.2 that the Bayes risk R^* under the DGF policy, approaches the

asymptotic lower bound as $c \rightarrow 0$. Specifically, the asymptotic optimality property of DGF is based on Lemma 11, showing that the asymptotic expected search time approaches $\frac{-\log c}{I^*(M, K)}$, while the error probability is $O(c)$ following Lemma 5.

Throughout the appendix we use the following notations: Let

$$N_j(n) \triangleq \sum_{t=1}^n \mathbf{1}_j(t) \quad (30)$$

be the number of times that cell j has been observed up to time n .

We define

$$\Delta S_{m,j}(n) \triangleq S_m(n) - S_j(n), \quad (31)$$

as the difference between the observed sum of LLRs of cells m and j . Let

$$\Delta S_m(n) \triangleq \min_{j \neq m} \Delta S_{m,j}(n). \quad (32)$$

Thus,

$$\Delta S(n) = S_{m^{(1)}(n)}(n) - S_{m^{(2)}(n)}(n) = \max_m \Delta S_m(n). \quad (33)$$

Without loss of generality we prove the theorem when hypothesis m is true. For convenience, we define

$$\tilde{\ell}_k(i) = \begin{cases} \ell_k(i) - D(g||f), & \text{if } k = m, \\ \ell_k(i) + D(f||g), & \text{if } k \neq m. \end{cases} \quad (34)$$

Note that $\tilde{\ell}_k(i)$ is a zero-mean r.v under hypothesis H_m .

1) *The Asymptotic Lower Bound on the Bayes Risk:* The asymptotic lower bound on the Bayes risk is shown in Theorem 5 below and is mainly based on Lemmas 1, 4, provided below. Throughout this section, τ denotes a generic stopping time that can be determined by any policy Γ . In Appendix A.2, however, we will refer to τ as the specific stopping time under the DGF policy. Lemma 1 shows that under \mathbf{P}_m , $\Delta S_m(\tau)$, defined in (32), must be large enough to obtain a sufficiently small error α_m . Lemma 4 implies that τ must be large enough to obtain a sufficiently large $\Delta S_m(\tau)$.

Lemma 1: Assume that $\alpha_j(\Gamma) = O(-c \log c)$ for all $j = 1, \dots, M$. Let $0 < \epsilon < 1$. Then:

$$\mathbf{P}_m(\Delta S_m(\tau) < -(1-\epsilon) \log c \mid \Gamma) = O(-c^\epsilon \log c), \quad (35)$$

for all $m = 1, \dots, M$.

Proof: Note that:

$$\begin{aligned} & \mathbf{P}_m(\Delta S_m(\tau) < -(1-\epsilon) \log c \mid \Gamma) \\ &= \mathbf{P}_m(\Delta S_m(\tau) < -(1-\epsilon) \log c, \delta = m \mid \Gamma) \\ & \quad + \mathbf{P}_m(\Delta S_m(\tau) < -(1-\epsilon) \log c, \delta \neq m \mid \Gamma) \\ & \leq \mathbf{P}_m(\Delta S_m(\tau) < -(1-\epsilon) \log c, \delta = m \mid \Gamma) + \alpha_m(\Gamma) \end{aligned} \quad (36)$$

Note that $\alpha_m(\Gamma) = O(-c \log c)$ as conditioned by the Lemma. Next, we upper bound the term $\mathbf{P}_m(\Delta S_m(\tau) < -(1-\epsilon) \log c, \delta = m \mid \Gamma)$ by changing the measure.

Let R_τ be the subset of the sample space, in which $\Delta S_{m,j}(\tau) < -(1-\epsilon) \log c$ for some $j \neq m$ and H_m is

accepted at time τ . Let $y_k(i)$ be the observation collected from cell k at time i (note that only K observations are obtained at a time. An observation is meaningful only when the process is probed. Otherwise, we can set an arbitrary value). Let $y(\tau) = \{y_1(i), \dots, y_M(i)\}_{i=1}^\tau$ be the set of all the observations up to time τ . Let $\mathcal{N}_k(y(\tau))$ be the set of time indices for the observations $y(\tau)$, containing the time indices in which cell k was probed. Thus, for all $j \neq m$ there exists $G > 0$ such that:

$$\begin{aligned}
& -Gc \log c \geq \mathbf{P}_j (\delta \neq j | \Gamma) \geq \mathbf{P}_j (\delta = m | \Gamma) \\
& \geq \mathbf{P}_j (\Delta S_{m,j}(\tau) \leq -(1-\epsilon) \log c, \delta = m | \Gamma) \\
& = \sum_{\tau=1}^{\infty} \int_{R_\tau} \left[\prod_{i \in \mathcal{N}_m(y(\tau))} f(y_m(i)) \prod_{i \in \mathcal{N}_j(y(\tau))} g(y_j(i)) \right. \\
& \quad \left. \times \prod_{k \neq m, j} \prod_{i \in \mathcal{N}_k(y(\tau))} f(y_k(i)) \right] d\mu(y(\tau)) \\
& = \sum_{\tau=1}^{\infty} \int_{R_\tau} \left[\prod_{i \in \mathcal{N}_m(y(\tau))} \frac{f(y_m(i))}{g(y_m(i))} \prod_{i \in \mathcal{N}_j(y(\tau))} \frac{g(y_j(i))}{f(y_j(i))} \right] \\
& \quad \times \left[\prod_{i \in \mathcal{N}_m(y(\tau))} g(y_m(i)) \prod_{i \in \mathcal{N}_j(y(\tau))} f(y_j(i)) \right. \\
& \quad \left. \times \prod_{k \neq m, j} \prod_{i \in \mathcal{N}_k(y(\tau))} f(y_k(i)) \right] d\mu(y(\tau)) \quad (37) \\
& = \sum_{\tau=1}^{\infty} \int_{R_\tau} \exp\{-\Delta S_{m,j}(\tau)\} \\
& \quad \times \left[\prod_{i \in \mathcal{N}_m(y(\tau))} g(y_m(i)) \prod_{i \in \mathcal{N}_j(y(\tau))} f(y_j(i)) \right. \\
& \quad \left. \times \prod_{k \neq m, j} \prod_{i \in \mathcal{N}_k(y(\tau))} f(y_k(i)) \right] d\mu(y(\tau)) \\
& \geq c^{1-\epsilon} \mathbf{P}_m (\Delta S_{m,j}(\tau) < -(1-\epsilon) \log c, \delta = m | \Gamma).
\end{aligned}$$

Thus,

$$\begin{aligned}
\mathbf{P}_m (\Delta S_{m,j}(\tau) < -(1-\epsilon) \log c, \delta = m | \Gamma) \\
= O(-c^\epsilon \log c) \quad \forall j \neq m. \quad (38)
\end{aligned}$$

As a result, by (32)

$$\begin{aligned}
& \mathbf{P}_m (\Delta S_m(\tau) < -(1-\epsilon) \log c, \delta = m | \Gamma) \\
& \leq \sum_{j \neq m} \mathbf{P}_m (\Delta S_{m,j}(\tau) < -(1-\epsilon) \log c, \delta = m | \Gamma) \\
& = O(-c^\epsilon \log c). \quad (39)
\end{aligned}$$

Finally,

$$\mathbf{P}_m (\Delta S_m(\tau) < -(1-\epsilon) \log c | \Gamma) = O(-c^\epsilon \log c). \quad (40)$$

Lemma 2: Assume that $K < M$ and

$$I^*(M, K) = D(g||f) + \frac{(K-1)D(f||g)}{M-1}.$$

Define the following function:

$$d(t) \triangleq t \left[D(g||f) + \frac{Kt^n - 1}{M-1} D(f||g) \right]. \quad (41)$$

Then, $d(t)$ is monotonically increasing with t .

Proof: Note that $I^*(M, K) = D(g||f) + \frac{(K-1)D(f||g)}{M-1}$ implies:

$$\begin{aligned}
D(g||f) + \frac{(K-1)D(f||g)}{M-1} & \geq \frac{KD(f||g)}{M-1} \\
& \iff D(g||f) \geq \frac{D(f||g)}{M-1}.
\end{aligned}$$

Differentiation $d(t)$ with respect to t yields:

$$\frac{\partial d(t)}{\partial t} = D(g||f) - \frac{D(f||g)}{M-1} \geq 0,$$

which completes the proof. \blacksquare

For the next lemma we define

$$j^*(t) \triangleq \arg \min_{j \neq m} N_j(t) \quad (42)$$

as the cell (except cell m) which has been observed the lowest number of times up to time t and

$$W_m^*(t) \triangleq \sum_{i=1}^t \tilde{\ell}_m(i) \mathbf{1}_m(i) - \sum_{i=1}^t \tilde{\ell}_{j^*(t)}(i) \mathbf{1}_{j^*(t)}(i). \quad (43)$$

Note that $W_m^*(t)$ is a sum of zero-mean r.v. The following lemma shows that $W_m^*(t)$ is sufficiently small. This result will be used in the proof of Lemma 4 to show that $\mathbf{P}_m (\max_{1 \leq t \leq n} \Delta S_m(t) \geq n(I^*(M, K) + \epsilon) | \Gamma) \rightarrow 0$ as $n \rightarrow \infty$.

Lemma 3: For every fixed $\epsilon > 0$ there exist $C > 0$ and $\gamma > 0$ such that

$$\mathbf{P}_m \left(\max_{1 \leq t \leq n} W_m^*(t) \geq n\epsilon | \Gamma \right) \leq C e^{-\gamma n} \quad (44)$$

for all $m = 1, \dots, M$ and for any policy Γ .

Proof: Since that $N_m(t), N_{j^*(t)}(t)$ are r.v, we can upper bound (44) by summing over any possible values that $N_m(t), N_{j^*(t)}(t)$ can take:

$$\begin{aligned}
& \mathbf{P}_m \left(\max_{1 \leq t \leq n} W_m^*(t) \geq n\epsilon | \Gamma \right) \\
& \leq \sum_{t=1}^n \mathbf{P}_m \left(\sum_{r=1}^t \tilde{\ell}_m(r) \mathbf{1}_m(r) - \tilde{\ell}_{j^*(t)}(i) \mathbf{1}_{j^*(t)}(r) \geq n\epsilon | \Gamma \right) \\
& = \sum_{t=1}^n \sum_{i=0}^t \sum_{j=0}^t \mathbf{P}_m \left(\sum_{r=1}^t \tilde{\ell}_m(r) \mathbf{1}_m(r) + \sum_{r=1}^t -\tilde{\ell}_{j^*(t)}(r) \mathbf{1}_{j^*(t)}(r) \right. \\
& \quad \left. \geq n\epsilon, N_m(t) = i, N_{j^*(t)} = j | \Gamma \right) \\
& \leq \sum_{t=1}^n \sum_{i=0}^t \sum_{j=0}^t \left[\mathbf{E}_m \left(e^{s(\tilde{\ell}_m(1)-\epsilon/2)} \right) \right]^i \\
& \quad \times \left[\mathbf{E}_m \left(e^{s(-\tilde{\ell}_{j^*(t)}(1)-\epsilon/2)} \right) \right]^j \times \exp \left\{ -s \frac{\epsilon}{2} (2n - i - j) \right\}, \quad (45)
\end{aligned}$$

for all $s > 0$.

The last inequality follows due to the i.i.d. property of $\ell_k(t)$ across the time series and applying the Chernoff bound for each term in the summation on the RHS of the equality. Note that we used the fact that the measure of the sample space that satisfies $\left\{ \sum_{r=1}^t \tilde{\ell}_m(r) \mathbf{1}_m(r) + \sum_{r=1}^t -\tilde{\ell}_{j^*(t)}(r) \mathbf{1}_{j^*(t)}(r) \geq n\epsilon, N_m(t) = i, N_{j^*(t)} = j \right\}$ under policy Γ is smaller than the measure of the sample space that satisfies $\left\{ \sum_{r=1}^i \tilde{\ell}_m(r) + \sum_{r=1}^j -\tilde{\ell}_{j^*(t)}(r) \geq n\epsilon \right\}$ (which is bounded by the Chernoff bound). This fact follows since that any selection of i, j observations from cells $m, j^*(t)$ (which have i.i.d distributions), yields the same distribution independent of the time they were taken. In particular, the intersection of the sample space $\left\{ \sum_{r=1}^t \tilde{\ell}_m(r) \mathbf{1}_m(r) + \sum_{r=1}^t -\tilde{\ell}_{j^*(t)}(r) \mathbf{1}_{j^*(t)}(r) \geq n\epsilon, \right\}$ and $\{N_m(t) = i, N_{j^*(t)} = j\}$ under policy Γ further decreases the measure.

Clearly, a moment generating function (MGF) is equal to one at $s = 0$. Furthermore, since $\mathbf{E}_m(\tilde{\ell}_m(1) - \epsilon/2) = -\epsilon/2 < 0$ and $\mathbf{E}_m(-\tilde{\ell}_{j^*(t)}(1) - \epsilon/2) = -\epsilon/2 < 0$ are strictly negative, differentiating the MGFs of $\tilde{\ell}_m(1) - \epsilon/2$ and $-\tilde{\ell}_{j^*(t)}(1) - \epsilon/2$ with respect to s yields strictly negative derivatives at $s = 0$. Hence, there exist $s > 0$ and $\gamma' > 0$ such that $\mathbf{E}_m\left(e^{s(\tilde{\ell}_m(1) - \epsilon/2)}\right), \mathbf{E}_m\left(e^{s(-\tilde{\ell}_{j^*(t)}(1) - \epsilon/2)}\right)$ and $e^{-s\epsilon/2}$ are strictly less than $e^{-\gamma'} < 1$. Since $2n - i - j \geq 0$, there exist $C > 0$ and $\gamma > 0$, such that summing over t, i, j yields (44). ■

Lemma 4: For any fixed $\epsilon > 0$,

$$\mathbf{P}_m \left(\max_{1 \leq t \leq n} \Delta S_m(t) \geq n(I^*(M, K) + \epsilon) \mid \Gamma \right) \rightarrow 0 \quad \text{as } n \rightarrow \infty, \quad (46)$$

for all $m = 1, \dots, M$ and for any policy Γ .

Proof: It should be noted that a polynomial decay of a similar condition under a binary composite hypothesis testing was shown in [1, Lemma 5] using a variation of the Kolmogorov's inequality. Here, we use a different approach to show exponential decay of (46). Let

$$\Delta S_m^*(t) \triangleq S_m(t) - S_{j^*(t)}(t).$$

Since $\Delta S_m(t) \leq \Delta S_m^*(t)$ for all m and t , we have:

$$\begin{aligned} \mathbf{P}_m \left(\max_{1 \leq t \leq n} \Delta S_m(t) \geq n(I^*(M, K) + \epsilon) \mid \Gamma \right) \\ \leq \mathbf{P}_m \left(\max_{1 \leq t \leq n} \Delta S_m^*(t) \geq n(I^*(M, K) + \epsilon) \mid \Gamma \right) \end{aligned} \quad (47)$$

Next, we consider three cases:

Case 1 ($K = M$): In this case $I^*(M, K) = D(g||f) + D(f||g)$. Furthermore, note that $N_j(t) = t$, for all j and t . Thus,

$$\begin{aligned} \Delta S_m^*(t) &= W_m^*(t) + t(D(g||f) + D(f||g)) \\ &\leq W_m^*(t) + nI^*(M, K). \end{aligned} \quad (48)$$

Therefore,

$$\Delta S_m^*(t) \geq n(I^*(M, K) + \epsilon)$$

implies

$$W_m^*(t) \geq n\epsilon.$$

Applying Lemma 3 yields:

$$\begin{aligned} \mathbf{P}_m \left(\max_{1 \leq t \leq n} \Delta S_m(t) \geq n(I^*(M, K) + \epsilon) \mid \Gamma \right) \\ \leq \mathbf{P}_m \left(\max_{1 \leq t \leq n} W_m^*(t) \geq n\epsilon \mid \Gamma \right) \\ \leq C e^{-\gamma n} \rightarrow 0 \quad \text{as } n \rightarrow \infty. \end{aligned} \quad (49)$$

Case 2 ($K < M$ and $I^*(M, K) = \frac{KD(f||g)}{M-1}$): Note that:

$$\begin{aligned} \Delta S_m^*(t) &= W_m^*(t) + N_m(t)D(g||f) + N_{j^*(t)}(t)D(f||g) \\ &\leq W_m^*(t) + D(f||g) \left[\frac{N_m(t)}{M-1} + N_{j^*(t)}(t) \right] \end{aligned} \quad (50)$$

The last inequality holds since $\frac{KD(f||g)}{M-1} \geq D(g||f) + \frac{(K-1)D(f||g)}{M-1}$ implies $D(g||f) \leq \frac{D(f||g)}{M-1}$.

Since that $j^*(t) = \arg \min_{j \neq m} N_j(t)$ and $Kt - N_m(t)$ is the total number of observations taken from $M - 1$ cells $j \neq m$, we have:

$$N_{j^*(t)}(t) \leq \frac{Kt - N_m(t)}{M-1} \leq \frac{Kn - N_m(t)}{M-1}. \quad (51)$$

Hence,

$$\begin{aligned} \Delta S_m^*(t) &\leq W_m^*(t) + D(f||g) \frac{Kn}{M-1} \\ &= W_m^*(t) + nI^*(M, K). \end{aligned} \quad (52)$$

Therefore,

$$\Delta S_m^*(t) \geq n(I^*(M, K) + \epsilon)$$

implies

$$W_m^*(t) \geq n\epsilon.$$

The rest of the proof is similar to Case 1.

Case 3 ($K < M$ and $I^*(M, K) = D(g||f) + \frac{(K-1)D(f||g)}{M-1}$): Note that:

$$\begin{aligned} \Delta S_m^*(t) &= W_m^*(t) + N_m(t)D(g||f) + N_{j^*(t)}(t)D(f||g) \\ &\leq W_m^*(t) + N_m(t)D(g||f) + \frac{Kn - N_m(t)}{M-1}D(f||g) \\ &= W_m^*(t) + N_m(t) \left[D(g||f) + \frac{K \frac{n}{N_m(t)} - 1}{M-1} D(f||g) \right] \end{aligned} \quad (53)$$

Since $0 \leq N_m(t) \leq n$, by Lemma 2 we have:

$$\begin{aligned} \Delta S_m^*(t) &\leq W_m^*(t) + n \left[D(g||f) + \frac{K-1}{M-1} D(f||g) \right] \\ &= W_m^*(t) + nI^*(M, K). \end{aligned} \quad (54)$$

The rest of the proof is similar to Case 1. Hence, (46) follows. ■

The following theorem shows that in order to achieve a Bayes risk lower than $\frac{-c \log(c)}{I^*(M, K)}$ under any hypothesis, the risk must be of a greater order than $O(-c \log c)$ under some hypothesis. As a result, it provides a lower bound $\sim \frac{-c \log(c)}{I^*(M, K)}$ on the average Bayes risk (2):

Theorem 5: Any policy Γ that satisfies $R_j(\Gamma) = O(-c \log c)$ for all $j = 1, \dots, M$ must satisfy:

$$R_m(\Gamma) \geq -(1 + o(1)) \frac{c \log(c)}{I^*(M, K)}. \quad (55)$$

for all $m = 1, \dots, M$.

Proof: To show the lower bound on the Bayes risk, we use a similar argument as in [1]. For any $\epsilon > 0$ let $n_c = -(1 - \epsilon) \frac{\log c}{I^*(M, K) + \epsilon}$. Note that

$$\begin{aligned} \mathbf{P}_m(\tau \leq n_c | \Gamma) &= \mathbf{P}_m(\tau \leq n_c, \Delta S_m(\tau) \geq -(1 - \epsilon) \log c | \Gamma) \\ &\quad + \mathbf{P}_m(\tau \leq n_c, \Delta S_m(\tau) < -(1 - \epsilon) \log c | \Gamma) \\ &\leq \mathbf{P}_m\left(\max_{t \leq n_c} \Delta S_m(t) \geq -(1 - \epsilon) \log c | \Gamma\right) \\ &\quad + \mathbf{P}_m(\Delta S_m(\tau) < -(1 - \epsilon) \log c | \Gamma). \end{aligned} \quad (56)$$

The first term in the last inequality approaches zero as $c \rightarrow 0$ by Lemma 4. Next, note that Lemma 1 requires $\alpha_j(\Gamma) = O(-c \log c)$ for all $j = 1, \dots, M$. Since the theorem requires $R_j(\Gamma) = O(-c \log c)$ for all $j = 1, \dots, M$ (and recall that $\alpha_j(\Gamma) \leq R_j(\Gamma)$), we can apply Lemma 1. Thus, the second term in the last inequality approaches zero as $c \rightarrow 0$. As a result, the expected sample size under policy Γ satisfies:

$$\begin{aligned} \mathbf{E}_m(\tau | \Gamma) &\geq \sum_{n=n_c+1}^{\infty} n \mathbf{P}_m(\tau = n | \Gamma) \\ &\geq n_c \mathbf{P}_m(\tau \geq n_c + 1 | \Gamma) \rightarrow n_c \quad \text{as } c \rightarrow 0 \end{aligned} \quad (57)$$

Since $\epsilon > 0$ can be arbitrarily small we have $\mathbf{E}_m(\tau | \Gamma) \geq -(1 + o(1)) \log(c) / I^*(M, K)$. Hence, $R_m(\Gamma) \geq c \mathbf{E}_m(\tau | \Gamma) \geq -(1 + o(1)) c \log(c) / I^*(M, K)$. ■

2) *Asymptotic Optimality of the DGF Policy:* In this section we show that the DGF policy achieves the lower bound on the Bayes risk (55) as $c \rightarrow 0$. We mainly focus on the more interesting case where $K \geq 2$ cells are observed at a time. The case where $K = 1$ is simpler and follows with minor modifications. Below, the proof follows the structure discussed in Section III-B.

Lemma 5: Assume that the DGF policy is implemented. Then, the error probability is upper bounded by:

$$P_e \leq (M - 1)c. \quad (58)$$

Proof: Let $\alpha_{m,j} = \mathbf{P}_m(\delta = j)$ for all $j \neq m$. Thus, $\alpha_m = \sum_{j \neq m} \alpha_{m,j}$. Note that accepting H_j (i.e., $\Delta S_j(n) \geq -\log c$) implies $\Delta S_{j,m} \geq -\log c$. By changing the measure, as in [1, Lemma 3], we can show that for all $j \neq m$ the following holds:

$$\begin{aligned} \alpha_{m,j} &= \mathbf{P}_m(\delta = j) \\ &= \mathbf{P}_m(\Delta S_j(\tau) \geq -\log c) \leq \mathbf{P}_m(\Delta S_{j,m}(\tau) \geq -\log c) \\ &\leq c \mathbf{P}_j(\Delta S_{j,m}(\tau) \geq -\log c) \leq c. \end{aligned} \quad (59)$$

Finally,

$$\alpha_m = \sum_{j \neq m} \alpha_{m,j} \leq (M - 1)c.$$

Hence, (58) follows. ■

Lemma 6: Fix $0 < q < 1$. Then, there exist $C > 0$ and $\gamma > 0$ such that

$$\mathbf{P}_m(S_j(n) \geq S_m(n), N_j(n) \geq qn) \leq C e^{-\gamma n}, \quad (60)$$

and

$$\mathbf{P}_m(S_j(n) \geq S_m(n), N_m(n) \geq qn) \leq C e^{-\gamma n}, \quad (61)$$

hold under any policy for $m = 1, 2, \dots, M$ and $j \neq m$.

Proof [We prove (60)]: Proving (61) applies with minor modifications. Note that we can develop (60) by summing over any possible values that $N_j(n), N_m(n)$ can take (i.e., $N_j(n) = \lceil qn \rceil, \lceil qn \rceil + 1, \dots, n$, and $N_m(n) = 0, \dots, n$). Similar to (45), applying the Chernoff bound and using the i.i.d. property of $\ell_j(t), \ell_m(t)$ across time yield:

$$\begin{aligned} \mathbf{P}_m(S_j(n) \geq S_m(n), N_j(n) \geq qn) &\leq \sum_{r=\lceil qn \rceil}^n \sum_{k=0}^n \mathbf{P}_m\left(\sum_{i=1}^r \ell_j(i) + \sum_{i=1}^k -\ell_m(i) \geq 0\right) \\ &\leq \sum_{r=\lceil qn \rceil}^n \sum_{k=0}^n \left[\mathbf{E}_m\left(e^{s \ell_j(1)}\right)\right]^r \left[\mathbf{E}_m\left(e^{s(-\ell_m(1))}\right)\right]^k \end{aligned} \quad (62)$$

for all $s > 0$.

Note that a moment generating function (MGF) is equal to one at $s = 0$. Furthermore, since $\mathbf{E}_m(\ell_j(1)) = -D(f||g) < 0$ and $\mathbf{E}_m(-\ell_m(1)) = -D(g||f) < 0$ are strictly negative, differentiating the MGFs of $\ell_j(1), \ell_m(1)$ with respect to s yields strictly negative derivatives at $s = 0$. As a result, there exist $s > 0$ and $\gamma_1 > 0$ such that $\mathbf{E}_m(e^{s \ell_j(1)})$, and $\mathbf{E}_m(e^{s(-\ell_m(1))})$ are strictly less than $e^{-\gamma_1} < 1$. Hence, there exist $C > 0$ and $\gamma = \gamma_1 \cdot q > 0$ such that

$$\begin{aligned} \mathbf{P}_m(S_j(n) - S_m(n) \geq 0, N_j(n) \geq qn) &\leq \sum_{r=\lceil qn \rceil}^n e^{-\gamma_1 r} \sum_{k=0}^n e^{-\gamma_1 k} \leq C e^{-\gamma n}. \end{aligned} \quad (63)$$

For the following definition, recall that $S_m(n)$ is a random walk with positive expected increment $\mathbf{E}_m(\ell_m(n)) = D(g||f) > 0$, while $S_j(n)$, for $j \neq m$ is a random walk with negative expected increment $\mathbf{E}_m(\ell_j(n)) = -D(f||g) < 0$. As a result, ultimately, the sample path of $S_m(n)$ will dominate those of $S_j(n), \forall j \neq m$, when n (and also the number of samples taken from cells m or j) is sufficiently large. Below, we define a random time τ_1 , which is the last passage time where the sample path of $S_m(n)$ will dominate those of $S_j(n)$ for all $n \geq \tau_1$, i.e., τ_1 is the last passage time in which $S_m(n)$ crosses $S_j(n)$. It should be noted that τ_1 is not a stopping time (note that τ_1 depends on the future by definition) and the decision maker does not know whether τ_1 has arrived. In Lemma 7 below we show that τ_1 is sufficiently small with high probability. We will use this result later to upper bound the actual stopping time τ under DGF. ■

Definition 2: τ_1 is the smallest integer such that $S_m(n) > S_j(n)$ for all $j \neq m$ for all $n \geq \tau_1$.

Remark 1: In the following lemmas, when we say that the DGF policy is implemented indefinitely we mean that DGF probes the cells indefinitely according to its selection rule, while the stopping rule is disregarded.

Lemma 7: Assume that the DGF policy is implemented indefinitely. Then, there exist $C > 0$ and $\gamma > 0$ such that

$$\mathbf{P}_m(\tau_1 > n) \leq Ce^{-\gamma n}, \quad (64)$$

for $m = 1, 2, \dots, M$.

Proof: We focus on the case where $M > 2$. The case of $M = 2$ is simpler and follows with minor modifications. Note that:

$$\begin{aligned} \mathbf{P}_m(\tau_1 > n) &\leq \mathbf{P}_m\left(\max_{j \neq m} \sup_{t \geq n} (S_j(t) - S_m(t)) \geq 0\right) \\ &\leq \sum_{j \neq m} \sum_{t=n}^{\infty} \mathbf{P}_m(S_j(t) \geq S_m(t)). \end{aligned} \quad (65)$$

Therefore, it suffices to show that there exist $C > 0$ and $\gamma > 0$ such that $\mathbf{P}_m(S_j(n) \geq S_m(n)) \leq Ce^{-\gamma n}$.

Step 1: Bounding each term in the summation on the RHS of (65):

Let

$$\rho = \frac{1}{16(M-2)}.$$

Note that $0 < \rho \leq 1/16$.

Thus,

$$\begin{aligned} \mathbf{P}_m(S_j(n) \geq S_m(n)) &\leq \mathbf{P}_m(S_j(n) \geq S_m(n), N_j(n) < \rho n, N_m(n) < \rho n) \\ &\quad + \mathbf{P}_m(S_j(n) \geq S_m(n), N_j(n) \geq \rho n) \\ &\quad + \mathbf{P}_m(S_j(n) \geq S_m(n), N_m(n) \geq \rho n) \end{aligned} \quad (66)$$

By Lemma 6, there exist $\gamma_1 > 0$ and $D > 0$ such that the second and the third terms on the RHS are upper bounded by $De^{-\gamma_1 n}$. In the case of $K = M$ the first term on the RHS equals zero (since $N_j(n) = N_m(n) = n$ surely). Hence, it remains to show that the first term on the RHS decreases exponentially with n for $K < M$. Note that the event $(N_j(n) < \rho n, N_m(n) < \rho n)$ implies that at least $\tilde{n} = n - N_j(n) - N_m(n) \geq n(1 - 2\rho)$ times cells j, m are not observed. Let $\tilde{N}_r(n)$ be the number of times when cell $r \neq j, m$ has been observed and cells j, m have not been observed up to time n . We refer to each such time as $r_{\neq j, m}$ -**probing time**. There exists a cell $r \neq j, m$ such that $\tilde{N}_r(n) \geq \frac{\tilde{n}}{M-2} = \frac{n(1-2\rho)}{M-2}$. Hence, (66) can be upper bounded by:

$$\begin{aligned} \mathbf{P}_m(S_j(n) \geq S_m(n)) &\leq \sum_{r \neq j, m} \mathbf{P}_m\left(\tilde{N}_r(n) > \frac{n(1-2\rho)}{M-2}, \right. \\ &\quad \left. N_j(n) < \rho n, N_m(n) < \rho n\right) + 2De^{-\gamma_1 n} \end{aligned} \quad (67)$$

It remains to show that each term in the summation on the RHS of (67) decreases exponentially with n .

Step 2: Bounding each term in the summation on the RHS of (67):

Let $\tilde{t}_1^r, \tilde{t}_2^r, \dots, \tilde{t}_{\tilde{N}_r(n)}^r$ be the $r_{\neq j, m}$ -probing time indices and let

$$\zeta \triangleq \frac{1-2\rho}{2(M-2)}.$$

Note that:

- At every $r_{\neq j, m}$ -probing time, $S_j(n) \leq S_r(n)$ or $S_m(n) \leq S_r(n)$ must occur (otherwise, if $S_j(n) > S_r(n)$ and $S_m(n) > S_r(n)$ then j or m are observed).
- In particular, the event $\tilde{N}_r(n) > \frac{n(1-2\rho)}{M-2}$ implies that at time $\tilde{t}_{\zeta n}^r$ the following holds: $S_j(\tilde{t}_{\zeta n}^r) \leq S_r(\tilde{t}_{\zeta n}^r)$ or $S_m(\tilde{t}_{\zeta n}^r) \leq S_r(\tilde{t}_{\zeta n}^r)$ must occur.
- Since $N_r(t)$ is the total number of observations taken from cell r up to time t (during both $r_{\neq j, m}$ -probing times and all other times when cell r was probed), then $N_r(\tilde{t}_{\zeta n}^r) \geq \zeta n$.

Therefore, using the i.i.d. property of the LLRs across time we have³:

$$\begin{aligned} \mathbf{P}_m\left(\tilde{N}_r(n) > \frac{n(1-2\rho)}{M-2}, N_j(n) < \rho n, N_m(n) < \rho n\right) &\leq \sum_{N'_r = \zeta n}^n \mathbf{P}_m\left(\inf_{n' \leq \rho n} \sum_{i=1}^{n'} \ell_j(i) \leq \sum_{i=1}^{N'_r} \ell_r(i)\right) \\ &\quad + \sum_{N'_r = \zeta n}^n \mathbf{P}_m\left(\inf_{n' \leq \rho n} \sum_{i=1}^{n'} \ell_m(i) \leq \sum_{i=1}^{N'_r} \ell_r(i)\right) \\ &\leq \sum_{N'_r = \zeta n}^n \sum_{n'=0}^{\rho n} \mathbf{P}_m\left(\sum_{i=1}^{n'} \ell_j(i) \leq \sum_{i=1}^{N'_r} \ell_r(i)\right) \\ &\quad + \sum_{N'_r = \zeta n}^n \sum_{n'=0}^{\rho n} \mathbf{P}_m\left(\sum_{i=1}^{n'} \ell_m(i) \leq \sum_{i=1}^{N'_r} \ell_r(i)\right) \\ &= \sum_{q=0}^{n-\zeta n} \sum_{n'=0}^{\rho n} \mathbf{P}_m\left(\sum_{i=1}^{n'} \ell_j(i) \leq \sum_{i=1}^{\zeta n+q} \ell_r(i)\right) \\ &\quad + \sum_{q=0}^{n-\zeta n} \sum_{n'=0}^{\rho n} \mathbf{P}_m\left(\sum_{i=1}^{n'} \ell_m(i) \leq \sum_{i=1}^{\zeta n+q} \ell_r(i)\right). \end{aligned} \quad (68)$$

Step 3: Bounding the first term on the RHS of (68):

Note that

$$\begin{aligned} \sum_{i=1}^{\zeta n+q} \ell_r(i) + \sum_{i=1}^{n'} -\ell_j(i) &= \sum_{i=1}^{\zeta n+q} \tilde{\ell}_r(i) + \sum_{i=1}^{n'} -\tilde{\ell}_j(i) - D(f||g)(\zeta n + q - n'). \end{aligned} \quad (69)$$

³For the ease of presentation, throughout the proof we assume that $\zeta n, \rho n$ are integers. This assumption does not affect the exponential decay of the Chernoff bound but only the exact value of $C > 0$ in (64) (since $an - 1 \leq \lfloor an \rfloor \leq \lceil an \rceil \leq an + 1$ holds for all $a \geq 0$ for all $n = 0, 1, \dots$).

and

$$\begin{aligned} \zeta n + q - n' &\geq \zeta n + q - n' - 2(\rho n - n') \\ &= n(\zeta - 2\rho) + q + n' \geq \frac{1}{4(M-2)}n + q + n' \\ &\geq \frac{1}{4(M-2)}(n + q + n'), \end{aligned}$$

for all $n' \leq \rho n$.

Therefore,

$$\sum_{i=1}^{\zeta n+q} \ell_r(i) + \sum_{i=1}^{n'} -\ell_j(i) \geq 0 \quad (70)$$

implies

$$\sum_{i=1}^{\zeta n+q} \tilde{\ell}_r(i) + \sum_{i=1}^{n'} -\tilde{\ell}_j(i) \geq C_1(n + q + n'). \quad (71)$$

where $C_1 = \frac{D(f||g)}{4(M-2)} > 0$.

Applying the Chernoff bound and using the i.i.d. property of $\ell_r(t), \ell_j(t)$ across the time series yield:

$$\begin{aligned} \mathbf{P}_m \left(\sum_{i=1}^{n'} \ell_j(i) \leq \sum_{i=1}^{\zeta n+q} \ell_r(i) \right) \\ \leq \mathbf{P}_m \left(\sum_{i=1}^{\zeta n+q} \tilde{\ell}_r(i) + \sum_{i=1}^{n'} -\tilde{\ell}_j(i) \geq C_1(n + q + n') \right) \\ \leq \left[\mathbf{E}_m \left(e^{s\tilde{\ell}_r(1)} \right) \right]^{\zeta n+q} \left[\mathbf{E}_m \left(e^{s(-\tilde{\ell}_j(1))} \right) \right]^{n'} \\ \times e^{-sC_1(n+q+n')} \\ = \left[\mathbf{E}_m \left(e^{s(\tilde{\ell}_r(1)-C_1)} \right) \right]^{\zeta n+q} \\ \times \left[\mathbf{E}_m \left(e^{s(-\tilde{\ell}_j(1)-C_1)} \right) \right]^{n'} \times e^{-sC_1(n-\zeta n)}. \end{aligned} \quad (72)$$

for all $s > 0$.

Since $\mathbf{E}_m(\tilde{\ell}_r(1) - C_1) = -C_1 < 0$ and $\mathbf{E}_m(-\tilde{\ell}_j(1) - C_1) = -C_1 < 0$ are strictly negative, by applying a similar argument as at the end of the proof of Lemma 6, there exist $s > 0$ and $\gamma_2 > 0$ such that $\mathbf{E}_m \left(e^{s(\tilde{\ell}_r(1)-C_1)} \right)$, $\mathbf{E}_m \left(e^{s(-\tilde{\ell}_j(1)-C_1)} \right)$ and e^{-sC_1} are strictly less than $e^{-\gamma_2} < 1$. Hence,

$$\mathbf{P}_m \left(\sum_{i=1}^{n'} \ell_j(i) \leq \sum_{i=1}^{\zeta n+q} \ell_r(i) \right) \leq e^{-\gamma_2(n+q+n')}. \quad (73)$$

and

$$\begin{aligned} \sum_{q=0}^{n-\zeta n} \sum_{n'=0}^{\rho n} \mathbf{P}_m \left(\sum_{i=1}^{n'} \ell_j(i) \leq \sum_{i=1}^{\zeta n+q} \ell_r(i) \right) \\ \leq e^{-\gamma_2 n} \sum_{q=0}^{n-\zeta n} e^{-\gamma_2 q} \sum_{n'=0}^{\rho n} e^{-\gamma_2 n'} \leq C_2 e^{-\gamma_2 n}, \end{aligned} \quad (74)$$

where $C_2 = (1 - e^{-\gamma_2})^{-2}$.

Step 4: Bounding the second term on the RHS of (68):

Applying the Chernoff bound and using the i.i.d. property of $\ell_r(t), \ell_m(t)$ across the time series yield:

$$\begin{aligned} \mathbf{P}_m \left(\sum_{i=1}^{\zeta n+q} \ell_r(i) + \sum_{i=1}^{n'} -\ell_m(i) \geq 0 \right) \\ \leq \left[\mathbf{E}_m \left(e^{s\ell_r(1)} \right) \right]^{\zeta n+q} \left[\mathbf{E}_m \left(e^{s(-\ell_m(1))} \right) \right]^{n'}. \end{aligned} \quad (75)$$

for all $s > 0$.

Since $\mathbf{E}_m(\ell_r(1)) = -D(f||g) < 0$ and $\mathbf{E}_m(-\ell_m(1)) = -D(g||f) < 0$ are strictly negative, there exist $s > 0$ and $\gamma'_3 > 0$ such that $\mathbf{E}_m \left(e^{s\ell_r(1)} \right)$, $\mathbf{E}_m \left(e^{s(-\ell_m(1))} \right)$ are strictly less than $e^{-\gamma'_3} < 1$. Hence,

$$\mathbf{P}_m \left(\sum_{i=1}^{\zeta n+q} \ell_r(i) + \sum_{i=1}^{n'} -\ell_m(i) \geq 0 \right) \leq e^{-\gamma'_3(\zeta n+q+n')}. \quad (76)$$

Finally, there exists $\gamma_3 = \zeta\gamma'_3 > 0$ such that

$$\begin{aligned} \sum_{q=0}^{n-\zeta n} \sum_{n'=0}^{\rho n} \mathbf{P}_m \left(\sum_{i=1}^{n'} \ell_m(i) \leq \sum_{i=1}^{\zeta n+q} \ell_r(i) \right) \\ \leq e^{-\gamma_3 n} \sum_{q=0}^{n-\zeta n} e^{-\gamma_3 q/\zeta} \sum_{n'=0}^{\rho n} e^{-\gamma_3 n'/\zeta} \leq \frac{e^{-\gamma_3 n}}{(1 - e^{-\gamma_3/\zeta})^2}, \end{aligned} \quad (77)$$

which completes the proof. \blacksquare

It should be noted that differing from [5], where only a polynomial decay of $\mathbf{P}_m(\tau_1 > n)$ was shown to handle indistinguishable hypotheses under some (but not all) actions when applying the extended randomized Chernoff test, Lemma 7 shows exponential decay of $\mathbf{P}_m(\tau_1 > n)$ under DGF.

For the next lemmas we define

$$D'(f||g) \triangleq (K-1)D(f||g)/(M-1). \quad (78)$$

In what follows we define the second random time $\tau_2 \geq \tau_1$. τ_2 can be viewed as the time where sufficient information has been gathered to distinguish hypothesis m from at least one false hypothesis $j \neq m$. We point out that τ_2 is not a stopping time. However, it serves us later in upper bounding the stopping time τ under DGF. Lemma 8 shows that in the asymptotic regime the total time between τ_1 and τ_2 approaches $\sim -\log c/I^*(M, K)$.

Definition 3: τ_2 is defined as follows:

- 1) If $K = M$, τ_2 denotes the smallest integer such that $\sum_{i=\tau_1+1}^n \ell_m(i) \mathbf{1}_m(i) \geq -\frac{D(g||f)}{I^*(M,K)} \log c$ and $\sum_{i=\tau_1+1}^n \ell_{j_n}(i) \mathbf{1}_{j_n}(i) \leq \frac{D(f||g)}{I^*(M,K)} \log c$ for some $j_n \neq m$ for all $n \geq \tau_2 \geq \tau_1$.
- 2) If $K < M$ and $I^*(M, K) = KD(f||g)/(M-1)$, τ_2 denotes the smallest integer such that $\sum_{i=\tau_1+1}^n \ell_{j_n}(i) \mathbf{1}_{j_n}(i) \leq \log c$ for some $j_n \neq m$ for all $n \geq \tau_2 \geq \tau_1$.
- 3) If $K < M$ and $I^*(M, K) = D(g||f) + (K-1)D(f||g)/(M-1)$, τ_2 denotes the smallest integer such that $\sum_{i=\tau_1+1}^n \ell_m(i) \mathbf{1}_m(i) \geq -\frac{D(g||f)}{I^*(M,K)} \log c$ and $\sum_{i=\tau_1+1}^n \ell_{j_n}(i) \mathbf{1}_{j_n}(i) \leq \frac{D(f||g)}{I^*(M,K)} \log c$ for some $j_n \neq m$ for all $n \geq \tau_2 \geq \tau_1$.

Definition 4: $n_2 \triangleq \tau_2 - \tau_1$ denotes the total amount of time between τ_1 and τ_2 .

Lemma 8: Assume that the DGF policy is implemented indefinitely. Then, for every fixed $\epsilon > 0$ there exist $C > 0$ and $\gamma > 0$ such that

$$\mathbf{P}_m(n_2 > n) \leq Ce^{-\gamma n} \quad \forall n > -(1 + \epsilon) \log c / I^*(M, K), \quad (79)$$

for all $m = 1, 2, \dots, M$.

Proof: We prove the lemma for three cases:

Case 1 ($K = M$): In this case $I^*(M, K) = D(g||f) + D(f||g)$ and $\mathbf{1}_k(t) = 1$ for all k, t . Let τ_2^m and τ_2^j for $j \neq m$ be the smallest integers such that $\sum_{i=\tau_1+1}^{\tau_2^m} \ell_m(i) \geq -\frac{D(g||f)}{I^*(M, K)} \log c$ for all $n > \tau_2^m$ and $\sum_{i=\tau_1+1}^n \ell_j(i) \leq \frac{D(f||g)}{I^*(M, K)} \log c$ for all $n > \tau_2^j$ for $j \neq m$, respectively. Similarly, n_2^k denotes the total amount of time between τ_1 and τ_2^k . Clearly, $n_2 \leq \max_k(n_2^k)$. As a result,

$$\mathbf{P}_m(n_2 > n) \leq \sum_{k=1}^M \mathbf{P}_m(n_2^k > n). \quad (80)$$

Thus it remains to show that $\mathbf{P}_m(n_2^k > n)$ decreases exponentially with n . Next, we prove the lemma for cell m . The proof for cell $j \neq m$ follows with minor modifications. Let $\epsilon_1 = D(g||f)\epsilon / (1 + \epsilon) > 0$. Thus,

$$\begin{aligned} & \sum_{i=\tau_1+1}^{t+\tau_1} \ell_m(i) + \frac{D(g||f)}{I^*(M, K)} \log c \\ &= \sum_{i=\tau_1+1}^{t+\tau_1} \tilde{\ell}_m(i) + tD(g||f) + \frac{D(g||f)}{I^*(M, K)} \log c \\ &\geq \sum_{i=\tau_1+1}^{t+\tau_1} \tilde{\ell}_m(i) + t\epsilon_1, \end{aligned} \quad (81)$$

for all $t \geq n > -(1 + \epsilon) \log c / I^*(M, K)$.

As a result,

$$\sum_{i=\tau_1+1}^{t+\tau_1} \ell_m(i) \leq -\frac{D(g||f)}{I^*(M, K)} \log c. \quad (82)$$

implies

$$\sum_{i=\tau_1+1}^{t+\tau_1} \tilde{\ell}_m(i) \leq -t\epsilon_1. \quad (83)$$

Hence, for any $\epsilon > 0$ there exists $\epsilon_1 > 0$ such that

$$\begin{aligned} & \mathbf{P}_m(n_2^m > n) \\ &\leq \mathbf{P}_m\left(\inf_{t \geq n} \sum_{i=\tau_1+1}^{t+\tau_1} \ell_m(i) \leq -\frac{D(g||f)}{I^*(M, K)} \log c\right) \\ &\leq \sum_{t=n}^{\infty} \mathbf{P}_m\left(\sum_{i=\tau_1+1}^{t+\tau_1} \ell_m(i) \leq -\frac{D(g||f)}{I^*(M, K)} \log c\right) \\ &\leq \sum_{t=n}^{\infty} \mathbf{P}_m\left(\sum_{i=\tau_1+1}^{t+\tau_1} -\tilde{\ell}_m(i) \geq t\epsilon_1\right), \end{aligned} \quad (84)$$

for all $t \geq n > -(1 + \epsilon) \log c / I^*(M, K)$.

By applying the Chernoff bound, it can be shown that there exists $\gamma_1 > 0$ such that: $\mathbf{P}_m\left(\sum_{i=\tau_1+1}^{t+\tau_1} -\tilde{\ell}_m(i) \geq t\epsilon_1\right) \leq e^{-\gamma_1 t}$ for all $t \geq n > -(1 + \epsilon) \log c / I^*(M, K)$. Hence, there exist $C_1 > 0$ and $\gamma_1 > 0$ such that $\mathbf{P}_m(n_2^m > n) \leq C_1 e^{-\gamma_1 n}$ for all $n > -(1 + \epsilon) \log c / I^*(M, K)$.

Case 2 ($K < M$ and $I^(M, K) = \frac{KD(f||g)}{M-1}$):* In this case, the cell with the highest index is not observed for all n . As a result, cell m is not observed for all $n \geq \tau_1$ since $S_m(n) > S_j(n)$ for all $j \neq m$ for all $n \geq \tau_1$. Let $N'_j(\tau_1 + t) \triangleq \sum_{i=\tau_1+1}^{\tau_1+t} \mathbf{1}_j(i)$. Let $j^*(\tau_1 + t) = \arg \max_{j \neq m} N'_j(\tau_1 + t)$ be the cell index that was observed the largest number of times since τ_1 has occurred up to time $\tau_1 + t$. Note that if $\sum_{i=\tau_1+1}^{\tau_1+t} \ell_{j^*(\tau_1+t)}(i) \mathbf{1}_{j^*(\tau_1+t)}(i) \leq \log c$ for all $t \geq n$, then $n_2 \leq n$. Hence,

$$\begin{aligned} & \mathbf{P}_m(n_2 > n) \\ &\leq \mathbf{P}_m\left(\sup_{t \geq n} \sum_{i=\tau_1+1}^{\tau_1+t} \ell_{j^*(\tau_1+t)}(i) \mathbf{1}_{j^*(\tau_1+t)}(i) \geq \log c\right). \end{aligned} \quad (85)$$

Note that $N'_{j^*(\tau_1+t)}(\tau_1 + t) \geq \frac{Kt}{M-1}$ (since Kt observations are taken from $M-1$ cells). Thus, for any $\epsilon > 0$ there exists $\epsilon_1 > 0$ such that:

$$\begin{aligned} & \sum_{i=\tau_1+1}^{\tau_1+t} \ell_{j^*(\tau_1+t)}(i) \mathbf{1}_{j^*(\tau_1+t)}(i) - \log c \\ &= \sum_{i=\tau_1+1}^{\tau_1+t} \tilde{\ell}_{j^*(\tau_1+t)}(i) \mathbf{1}_{j^*(\tau_1+t)}(i) \\ &\quad - N'_{j^*(\tau_1+t)}(\tau_1 + t) D(f||g) - \log c \\ &\leq \sum_{i=\tau_1+1}^{\tau_1+t} \tilde{\ell}_{j^*(\tau_1+t)}(i) \mathbf{1}_{j^*(\tau_1+t)}(i) \\ &\quad - \frac{tD(f||g)K}{M-1} \left(1 - \frac{-(M-1) \log c}{tKD(f||g)}\right) \\ &\leq \sum_{i=\tau_1+1}^{\tau_1+t} \tilde{\ell}_{j^*(\tau_1+t)}(i) \mathbf{1}_{j^*(\tau_1+t)}(i) - t\epsilon_1, \end{aligned} \quad (86)$$

for all $t \geq n > -(1 + \epsilon) \log c / I^*(M, K) = -(1 + \epsilon)(M-1) \log c / (KD(f||g))$.

The rest of the proof follows by applying the Chernoff bound.

Case 3 ($K < M$ and $I^(M, K) = D(g||f) + \frac{(K-1)D(f||g)}{M-1}$):* In this case, the cell with the highest index is observed for all n . As a result, cell m is observed for all $n \geq \tau_1$ since $S_m(n) > S_j(n)$ for all $j \neq m$ for all $n \geq \tau_1$. Therefore, a similar argument as in Case 1 applies to cell m . Next, we focus on cell $j^*(\tau_1 + t) \neq m$ as in Case 2. Note that in this case $N_{j^*(\tau_1+t)}(\tau_1 + t) \geq \frac{(K-1)t}{M-1}$ (since t observations are taken from cell m and $(K-1)t$ observations are taken from $M-1$ cells (for $j \neq m$)). Thus, (86) can be developed for this case as well with minor modifications. ■

In what follows we define the dynamic range of the false hypotheses in terms of their sum LLRs. Note that the dynamic range at time τ_2 (which is the time where sufficient information has been gathered to distinguish H_m from at least one false

hypothesis) can be viewed as the total amount of information remains to gather in order to distinguish H_m from all the false hypotheses. Lemma 9 shows that the dynamic range at time τ_2 is sufficiently small.

Definition 5: The dynamic range of the false hypotheses at time t is defined as follows:

$$\text{DR}(t) \triangleq \max_{j \neq m} S_j(t) - \min_{j \neq m} S_j(t). \quad (87)$$

Lemma 9: Assume that the DGF policy is implemented indefinitely. Then, for every fixed $\epsilon_1 > 0$, $\epsilon_2 > 0$ there exist $C > 0$ and $\gamma > 0$ such that

$$\mathbf{P}_m(\text{DR}(\tau_2) > \epsilon_1 n) \leq C e^{-\gamma n} \quad \forall n > -(1 + \epsilon_2) \log c / I^*(M, K), \quad (88)$$

for all $m = 1, 2, \dots, M$.

Proof: Note that

$$\begin{aligned} \mathbf{P}_m(\text{DR}(\tau_2) > \epsilon_1 n) \\ \leq \mathbf{P}_m(\tau_2 > n) + \mathbf{P}_m(\text{DR}(\tau_2) > \epsilon_1 n, \tau_2 \leq n) \end{aligned} \quad (89)$$

Since $\tau_2 = \tau_1 + n_2$, applying Lemmas 7, 8 implies that the first term on the RHS of (89) decreases exponentially with n for all $n > -(1 + \epsilon_2) \log c / I^*(M, K)$ for every fixed $\epsilon_2 > 0$. It remains to show that the second term on the RHS of (89) decreases exponentially with n . Let $\bar{j} = \arg \max_{j \neq m} S_j(\tau_2)$, $\underline{j} = \arg \min_{j \neq m} S_j(\tau_2)$. Let t_0 be the smallest integer such that $S_{\bar{j}}(t) \leq S_{\underline{j}}(t)$ for all $t_0 < t \leq \tau_2$. As a result, $\text{DR}(\tau_2) > \epsilon_1 n$ implies

$$\sum_{t=t_0}^{\tau_2} \ell_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \ell_{\underline{j}} \mathbf{1}_{\underline{j}}(t) > \epsilon_1 n.$$

Note that the second term on the RHS of (89) can be rewritten as:

$$\begin{aligned} \mathbf{P}_m(\text{DR}(\tau_2) > \epsilon_1 n, \tau_2 \leq n) \\ = \mathbf{P}_m(\text{DR}(\tau_2) > \epsilon_1 n, \tau_2 \leq n, t_0 \geq \tau_1) \\ + \mathbf{P}_m(\text{DR}(\tau_2) > \epsilon_1 n, \tau_2 \leq n, t_0 < \tau_1) \end{aligned} \quad (90)$$

Let $\underline{N} = \sum_{t=t_0}^{\tau_2} \mathbf{1}_{\underline{j}}(t)$, $\bar{N} = \sum_{t=t_0}^{\tau_2} \mathbf{1}_{\bar{j}}(t)$.

First, we upper bound the first term on the RHS of (90). Note that for all $\tau_1 \leq t_0 < t \leq \tau_2$, if $\mathbf{1}_{\underline{j}}(t) = 1$ then $\mathbf{1}_{\bar{j}}(t) = 1$ (since $S_m(t) > S_j(t)$ for all $t \geq \tau_1$ and the decision maker observes either the K cells with the top K highest sum LLRs or those with the second to the $(K + 1)^{\text{th}}$ highest sum LLRs). Hence, $\underline{N} \leq \bar{N}$. Thus,

$$\begin{aligned} \sum_{t=t_0}^{\tau_2} \ell_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \ell_{\underline{j}} \mathbf{1}_{\underline{j}}(t) \\ = \sum_{t=t_0}^{\tau_2} \left[\tilde{\ell}_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \tilde{\ell}_{\underline{j}} \mathbf{1}_{\underline{j}}(t) \right] - D(f||g) (\bar{N} - \underline{N}) \\ \leq \sum_{t=t_0}^{\tau_2} \tilde{\ell}_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \tilde{\ell}_{\underline{j}} \mathbf{1}_{\underline{j}}(t) \end{aligned} \quad (91)$$

Similar to (45), applying the Chernoff bound completes the proof for this case.

Next, we upper bound the second term on the RHS of (90). Let $\epsilon_3 \triangleq \frac{\epsilon_1}{4D(f||g)} > 0$. Note that

$$\begin{aligned} \mathbf{P}_m(\text{DR}(\tau_2) > \epsilon_1 n, \tau_2 \leq n, t_0 < \tau_1) \\ \leq \mathbf{P}_m(\tau_1 > \epsilon_3 n) \\ + \mathbf{P}_m(\text{DR}(\tau_2) > \epsilon_1 n, \tau_2 \leq n, t_0 < \tau_1, \tau_1 \leq \epsilon_3 n). \end{aligned} \quad (92)$$

The first term on the RHS of (92) decreases exponentially with n by Lemma 7. Thus, it remains to show that the second term on the RHS of (92) decreases exponentially with n . Note that $\text{DR}(\tau_2) > \epsilon_1 n$ implies

$$\left(\sum_{t=t_0}^{\tau_1} \ell_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \ell_{\underline{j}} \mathbf{1}_{\underline{j}}(t) \right) + \left(\sum_{t=\tau_1+1}^{\tau_2} \ell_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \ell_{\underline{j}} \mathbf{1}_{\underline{j}}(t) \right) > \epsilon_1 n.$$

Therefore, the second term on the RHS of (92) can be rewritten as:

$$\begin{aligned} \mathbf{P}_m(\text{DR}(\tau_2) > \epsilon_1 n, \tau_2 \leq n, t_0 < \tau_1, \tau_1 \leq \epsilon_3 n) \\ \leq \mathbf{P}_m \left(\sum_{t=t_0}^{\tau_1} \ell_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \ell_{\underline{j}} \mathbf{1}_{\underline{j}}(t) > \frac{\epsilon_1 n}{2}, \right. \\ \left. \tau_2 \leq n, t_0 < \tau_1, \tau_1 \leq \epsilon_3 n \right) \\ + \mathbf{P}_m \left(\sum_{t=\tau_1+1}^{\tau_2} \ell_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \ell_{\underline{j}} \mathbf{1}_{\underline{j}}(t) > \frac{\epsilon_1 n}{2}, \right. \\ \left. \tau_2 \leq n, t_0 < \tau_1, \tau_1 \leq \epsilon_3 n \right) \end{aligned} \quad (93)$$

Note that for all $t_0 < \tau_1 + 1 \leq t \leq \tau_2$, if $\mathbf{1}_{\underline{j}}(t) = 1$ then $\mathbf{1}_{\bar{j}}(t) = 1$. As a result, cell \bar{j} is probed more frequently between $\tau_1 + 1$ and τ_2 . Thus, the second term on the RHS of (93) decreases exponentially with n using a similar argument as in (91). Next, it remains to show that the first term on the RHS of (93) decreases exponentially with n . Note that

$$\begin{aligned} \sum_{t=t_0}^{\tau_1} \ell_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \ell_{\underline{j}} \mathbf{1}_{\underline{j}}(t) \\ \leq \sum_{t=t_0}^{\tau_1} \left[\tilde{\ell}_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \tilde{\ell}_{\underline{j}} \mathbf{1}_{\underline{j}}(t) \right] + D(f||g) \tau_1 \\ \leq \sum_{t=t_0}^{\tau_1} \left[\tilde{\ell}_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \tilde{\ell}_{\underline{j}} \mathbf{1}_{\underline{j}}(t) \right] + \frac{\epsilon_1}{4} n \end{aligned} \quad (94)$$

for all $\tau_1 \leq \epsilon_3 n$.

As a result,

$$\sum_{t=t_0}^{\tau_1} \ell_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \ell_{\underline{j}} \mathbf{1}_{\underline{j}}(t) > \frac{\epsilon_1}{2} n \quad (95)$$

implies

$$\sum_{t=t_0}^{\tau_1} \left[\tilde{\ell}_{\bar{j}} \mathbf{1}_{\bar{j}}(t) - \tilde{\ell}_{\underline{j}} \mathbf{1}_{\underline{j}}(t) \right] > \frac{\epsilon_1}{4} n \quad (96)$$

for all $\tau_1 \leq \epsilon_3 n$.

Similar to (45), applying the Chernoff bound completes the proof. \blacksquare

Definition 6: The dynamic range $\text{DR}_j(t)$ for $j \neq m$ at time t is defined as follows:

$$\text{DR}_j(t) \triangleq \max_{k \neq m} S_k(t) - S_j(t). \quad (97)$$

Let

$$\eta \triangleq \begin{cases} \frac{D(f||g)}{I^*(M,K)} \log c, & \text{if } K = M \\ \log c, & \text{if } K < M \text{ and } I^*(M, K) = K \frac{D(f||g)}{M-1} \\ \frac{D'(f||g)}{I^*(M,K)} \log c, & \\ \text{if } K < M \text{ and } I^*(M, K) = D(g||f) + \frac{(K-1)D(f||g)}{M-1} \end{cases} \quad (98)$$

where $D'(f||g)$ is defined in (78).

Definition 7: τ_3^j denotes the smallest integer such that $\sum_{i=\tau_1+1}^n \ell_j(i) \mathbf{1}_j(i) \leq \eta + \text{DR}_j(\tau_1)$ for $j \neq m$ for all $n \geq \tau_3^j \geq \tau_2$. We also define $\tau_3 \triangleq \max_{j \neq m} \tau_3^j$.

Note that $\tau_3^j \geq \tau_2$ by definition (i.e., both τ_2 has passed and the inequality holds for all $n \geq \tau_3^j$).

Remark 2: Using some algebraic manipulations, it can be verified that $\Delta S_{m,j}(n) \geq -\log c$ for all $j \neq m$ for all $n \geq \tau_3^j$. Since $\tau_3 = \max_{j \neq m} \tau_3^j$ we have $\Delta S(n) = S_m(n) - S_{m^{(2)}(n)}(n) \geq -\log c$ for all $n \geq \tau_3$. It should be noted that τ_3 depends on the future and is not a stopping time. The decision maker does not know whether it has arrived. However, it is used to upper bound the actual stopping time under DGF. Note that the stopping time under DGF stops the sampling once $\Delta S(n) \geq -\log c$ first occurs. If $\Delta S(n) \geq -\log c$ first occurs once τ_3 occurs, then $\tau = \tau_3$. Otherwise, $\tau < \tau_3$.

Definition 8: $n_3 \triangleq \tau_3 - \tau_2$ denotes the total amount of time between τ_2 and τ_3 .

Lemma 10: Assume that the DGF policy is implemented indefinitely. Then, for every fixed $\epsilon > 0$ there exist $C > 0$ and $\gamma > 0$ such that

$$\mathbf{P}_m(n_3 > n) \leq C e^{-\gamma n} \quad \forall n > -\epsilon \log c / I^*(M, K), \quad (99)$$

for all $m = 1, 2, \dots, M$.

Proof: Let N_3^j for $j \neq m$ denote the total number of observations, taken from cell j between τ_2 and τ_3^j . Note that $n_3 \leq \sum_{j \neq m} N_3^j$. Thus, it suffices to show that $\mathbf{P}_m(N_3^j > n)$ decreases exponentially with n . Note that

$$\begin{aligned} \mathbf{P}_m(N_3^j > n) &\leq \mathbf{P}_m\left(\text{DR}(\tau_2) > n \frac{D(f||g)}{2}\right) \\ &+ \mathbf{P}_m\left(N_3^j > n \mid \text{DR}(\tau_2) \leq n \frac{D(f||g)}{2}\right). \end{aligned} \quad (100)$$

By Lemma 9, the first term on the RHS of (100) decreases exponentially with n for all $n > -\epsilon \log c / I^*(M, K)$. Thus, it remains to show that the second term on the RHS of (100) decreases exponentially with n .

Let t_1, t_2, \dots denote the time indices when cell j is observed between τ_2 and τ_3^j . Since τ_2 has occurred and $\text{DR}(\tau_2) \leq n \frac{D(f||g)}{2}$, τ_3^j occurs once $\sum_{i=1}^r -\ell_j(t_i) \geq n \frac{D(f||g)}{2}$ holds for

all $r \geq N_3^j$. As a result,

$$\begin{aligned} \mathbf{P}_m\left(N_3^j > n \mid \text{DR}(\tau_2) \leq n \frac{D(f||g)}{2}\right) \\ \leq \mathbf{P}_m\left(\inf_{r > n} \sum_{i=1}^r -\ell_j(t_i) < n \frac{D(f||g)}{2}\right) \\ \leq \sum_{r=n}^{\infty} \mathbf{P}_m\left(\sum_{i=1}^r \tilde{\ell}_j(t_i) > r \frac{D(f||g)}{2}\right). \end{aligned} \quad (101)$$

Thus, it suffices to show that there exists $\gamma > 0$ such that $\mathbf{P}_m\left(\sum_{i=1}^n \tilde{\ell}_j(t_i) > n \frac{D(f||g)}{2}\right) \leq e^{-\gamma n}$. Applying the Chernoff bound and using the i.i.d. property of $\tilde{\ell}_j(t_i)$ completes the proof. ■

Lemma 11: The expected detection time τ under the DGF policy is upper bounded by:

$$\mathbf{E}_m(\tau) \leq -(1 + o(1)) \frac{\log(c)}{I^*(M, K)}, \quad (102)$$

for $m = 1, \dots, M$.

Proof: Note that $\tau \leq \tau_3 = \tau_1 + n_2 + n_3$. Thus, combining Lemmas 7, 8 and 10 completes the proof. ■

Combining Lemmas 5, 11 and Theorem 5 yields the asymptotic optimality property of the DGF policy, presented in Theorem 1.

B. Proof of Theorem 2

Following the same argument as in [5], it suffices to show that $\mathbf{P}_m(\tau_1 > n)$ decreases polynomially with n to prove the theorem. Since $D(g||f) > 0$, the KL divergence between the true hypothesis m and any false hypothesis $j \neq m$ is strictly positive under any observed cell. For the ease of presentation, consider the case where $K = 1$. In the case where $D(g||f) \geq D(f||g)/(M-1)$, the Chernoff test selects $m^{(1)}(n)$ for all n . As a result, exponential decay of $\mathbf{P}_m(\tau_1 > n)$ follows directly from Lemma 7. In the case where $D(g||f) < D(f||g)/(M-1)$, the Chernoff test selects $m^j(n)$ for $j \neq 1$ randomly for all n . As a result, polynomial decay of $\mathbf{P}_m(\tau_1 > n)$ follows by a similar argument as in [5] for the extended Chernoff test. Note that the proof directly applies to the case where $K > 1$ since $\mathbf{P}_m(\tau_1 > n)$ decreases as the number K of observations collected at a time increases. ■

C. Proof of Theorem 3

The proof follows a similar line of arguments as in the proof of Theorem 1. Hence, we provide here only a sketch of the proof. First, similar to Lemma 5, it can be verified that declaring the target locations once $S_{m^{(L)}} - S_{m^{(L+1)}} \geq -\log c$ occurs achieves an error probability $O(c)$. Second, similar to Lemma 11, it can be verified that the detection time approaches $-\log c / I^*(M, K)$. For example, if $\frac{D(g||f)}{L} \geq \frac{D(f||g)}{M-L}$ and $K \geq L$ then all the L targets and a fraction $r = \frac{K-L}{M-L}$ of the false hypotheses are observed at each given time in the asymptotic regime. Therefore, the detection time approaches $\frac{-\log c}{D(g||f) + rD(f||g)}$. Similar arguments apply to the rest of the cases. ■

D. Proof of Theorem 4

The proof follows a similar line of arguments as in the proofs of Theorems 1 and 3. Hence, we provide only a sketch of the proof. With minor modifications to Theorem 3, it can be verified that (29) is the asymptotic lower bound on the Bayes risk when the number of targets ℓ is known, $K = 1$ and (23) holds. Similar to Lemma 5, it can be verified that declaring a target once $S_m(n) \geq -\log c$ occurs achieves an error probability $O(c)$. Following a similar argument as in Lemma 7, it can be verified that the ℓ targets are tested before testing the $M - \ell$ normal processes with high probability in the asymptotic regime. Since the decision maker declares the target locations once $S_m(n) \geq -\log c$ for any m , the Bayes risk approaches (29) as $c \rightarrow 0$. ■

REFERENCES

- [1] H. Chernoff, "Sequential design of experiments," *Ann. Math. Statist.*, vol. 30, no. 3, pp. 755–770, 1959.
- [2] A. Wald, *Sequential Analysis*. New York, NY, USA: Wiley, 1947.
- [3] S. A. Bessler, "Theory and applications of the sequential design of experiments, k-actions and infinitely many experiments: Part I—Theory," Appl. Math. Statist. Lab., Stanford Univ., Stanford, CA, USA, Tech. Rep. 55, 1960.
- [4] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for hypothesis testing," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Mar. 2012, pp. 5277–5280.
- [5] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Trans. Autom. Control*, vol. 58, no. 10, pp. 2451–2464, Oct. 2013.
- [6] S. Nitinawarat and V. V. Veeravalli. (2013). "Controlled sensing for sequential multihypothesis testing with controlled Markovian observations and non-uniform control cost." [Online]. Available: <http://arxiv.org/abs/1310.1844>
- [7] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *Ann. Statist.*, vol. 41, no. 6, pp. 2703–2738, 2013.
- [8] M. Naghshvar and T. Javidi, "Sequentiality and adaptivity gains in active hypothesis testing," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 5, pp. 768–782, Oct. 2013.
- [9] K. S. Ziganirov, "On a problem in optimal scanning," *Theory Probab. Appl.*, vol. 11, no. 2, pp. 294–298, 1966.
- [10] E. M. Klimko and J. Yackel, "Optimal search strategies for Wiener processes," *Stochastic Processes Appl.*, vol. 3, no. 1, pp. 19–33, 1975.
- [11] V. Dragalin, "A simple and effective scanning rule for a multi-channel system," *Metrika*, vol. 43, no. 1, pp. 165–182, 1996.
- [12] L. D. Stone and J. A. Stanshine, "Optimal search using uninterrupted contact investigation," *SIAM J. Appl. Math.*, vol. 20, no. 2, pp. 241–263, 1971.
- [13] K. P. Toggetti, "An optimal strategy for a whereabouts search," *Oper. Res.*, vol. 16, no. 1, pp. 209–211, 1968.
- [14] J. B. Kadane, "Optimal whereabouts search," *Oper. Res.*, vol. 19, no. 4, pp. 894–904, 1971.
- [15] Y. Zhai and Q. Zhao, "Dynamic search under false alarms," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Dec. 2013, pp. 201–204.
- [16] D. A. Castanon, "Optimal search strategies in dynamic hypothesis testing," *IEEE Trans. Syst., Man, Cybern.*, vol. 25, no. 7, pp. 1130–1138, Jul. 1995.
- [17] Y. Li, S. Nitinawarat, and V. V. Veeravalli. (2013). "Universal outlier hypothesis testing." [Online]. Available: <http://arxiv.org/abs/1302.4776>
- [18] H. Li, "Restless watchdog: Selective quickest spectrum sensing in multichannel cognitive radio systems," *EURASIP J. Adv. Signal Process.*, vol. 2009, p. 417457, Sep. 2009.
- [19] O. Hadjiladis, H. Zhang, and H. Vincent Poor, "One shot schemes for decentralized quickest change detection," *IEEE Trans. Inf. Theory*, vol. 55, no. 7, pp. 3346–3359, Jul. 2009.
- [20] Q. Zhao and J. Ye, "Quickest detection in multiple on–off processes," *IEEE Trans. Signal Process.*, vol. 58, no. 12, pp. 5994–6006, Dec. 2010.
- [21] V. Raghavan and V. V. Veeravalli, "Quickest change detection of a Markov process across a sensor array," *IEEE Trans. Inf. Theory*, vol. 56, no. 4, pp. 1961–1981, Apr. 2010.
- [22] L. Lai, H. Vincent Poor, Y. Xin, and G. Georgiadis, "Quickest search over multiple sequences," *IEEE Trans. Inf. Theory*, vol. 57, no. 8, pp. 5375–5386, Aug. 2011.
- [23] M. L. Malloy, G. Tang, and R. D. Nowak, "Quickest search for a rare distribution," in *Proc. 46th Annu. Conf. Inf. Sci. Syst.*, Mar. 2012, pp. 1–6.
- [24] M. L. Malloy and R. D. Nowak. (2012). "Sequential testing for sparse recovery." [Online]. Available: <http://arxiv.org/abs/1212.1801>
- [25] A. Tajar and H. Vincent Poor, "Quick search for rare events," *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4462–4481, Jul. 2013.
- [26] E. Bayraktar and L. Lai. (2013). "Byzantine fault tolerant distributed quickest change detection." [Online]. Available: <http://arxiv.org/abs/1306.2086>
- [27] R. Caromi, Y. Xin, and L. Lai, "Fast multiband spectrum scanning for cognitive radio systems," *IEEE Trans. Commun.*, vol. 61, no. 1, pp. 63–75, Jan. 2013.
- [28] K. Cohen, Q. Zhao, and A. Swami, "Optimal index policies for quickest localization of anomaly in cyber networks," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Dec. 2013, pp. 221–224.
- [29] H. Zhang, O. Hadjiladis, T. Schäfer, and H. Vincent Poor, "Quickest detection in coupled systems," *SIAM J. Control Optim.*, vol. 52, no. 3, pp. 1567–1596, 2014.
- [30] K. Cohen, Q. Zhao, and A. Swami, "Optimal index policies for anomaly localization in resource-constrained cyber systems," *IEEE Trans. Signal Process.*, vol. 62, no. 16, pp. 4224–4236, Aug. 2014.
- [31] K. Cohen and Q. Zhao. (2014). "Asymptotically optimal anomaly detection via sequential testing." [Online]. Available: <http://arxiv.org/abs/1405.4659>
- [32] K. Cohen and Q. Zhao, "Anomaly detection over independent processes: Switching with memory," in *Proc. 52nd Annu. Allerton Conf. Commun., Control, Comput.*, Oct. 2014.
- [33] M. L. Pinedo, *Scheduling: Theory, Algorithms, and Systems*. New York, NY, USA: Springer-Verlag, 2012.

Kobi Cohen received the B.Sc.(cum laude) and Ph.D. degrees in electrical engineering from Bar-Ilan University, Ramat Gan, Israel, in 2007 and 2013, respectively. He is currently a postdoctoral research associate at the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign. From 2012 to 2014, he was with the Department of Electrical and Computer Engineering at the University of California, Davis, as a postdoctoral researcher. His main research interests include decision theory, stochastic optimization, game theory and statistical inference, with applications in large-scale systems, and wireless and wireline networks.

Qing Zhao (F'13) received the Ph.D. degree in Electrical Engineering in 2001 from Cornell University, Ithaca, NY. In August 2004, she joined the Department of Electrical and Computer Engineering at University of California, Davis, where she is currently a Professor. She is also a Professor with the Graduate Group of Applied Mathematics at UC Davis. Her research interests are in the general area of stochastic optimization, decision theory, machine learning, and algorithmic theory in dynamic systems and communication and social-economic networks. Qing Zhao is a Fellow of IEEE. She received the 2010 IEEE Signal Processing Magazine Best Paper Award and the 2000 Young Author Best Paper Award from the IEEE Signal Processing Society. She holds the title of UC Davis Chancellor's Fellow and received the 2014 Outstanding Mid-Career Faculty Research Award and the 2008 Outstanding Junior Faculty Award from the UC Davis College of Engineering. She was a plenary speaker at the 11th IEEE Workshop on Signal Processing Advances in Wireless Communications (SPAWC), 2010. She is also a co-author of two papers that received student paper awards at ICASSP 2006 and the IEEE Asilomar Conference 2006.