Opportunistic Detection Rules

Wenyi Zhang[1], George V. Moustakides[2], H. Vincent Poor[3]
[1]Department of EEIS, University of Science and Technology of China, China
Email: wenyizha@ustc.edu.cn
[2]Department of Electrical and Computer Engineering, University of Patras, Greece
Email: moustaki@upatras.gr
[3]Department of Electrical Engineering, Princeton University, Princeton, NJ, USA
Email: poor@princeton.edu

Abstract— **Opportunistic detection rules (ODRs) are variants of fixed-sample-size detection rules in which the statistician is allowed to make an early decision on the alternative hypothesis opportunistically based on the sequentially observed samples. From a sequential decision perspective, ODRs are also mixtures of one-sided and truncated sequential detection rules. Several key properties of ODRs are established in this paper, in both the asymptotic regime in which the maximum sample size grows without bound, and the finite regime in which the maximum samples size is a fixed finite number. Furthermore, an extended setup, in which the maximum sample size is a random variable following a geometric distribution whose realization is not revealed to the statistician until observing the last sample, is studied.**

I. Introduction and Background

Consider that an array of random variables X_i, $i = 1, 2, \ldots$, is drawn sequentially, in an independent and identically distributed (i.i.d.) manner, from one of two possible distributions p_0 and p_1. A statistician sequentially observes the random variables (called "samples" interchangeably throughout the paper), and following a certain stopping rule, stops at some point and makes a decision on which of the two possible distributions the samples obey. The Wald-Wolfowitz theorem (See, e.g., [1] [2, Sec. 7.6] [3, Thm. 4.7]) is fundamental, asserting that the sequential probability ratio test (SPRT), which sequentially compares the likelihood ratios $\Lambda_k = \prod_{i=1}^{k} \frac{p_1(\mathsf{X}_i)}{p_0(\mathsf{X}_i)}$, $k = 1, 2, \ldots$, against two thresholds $0 < A \leq 1 \leq B < \infty$, and decides hypothesis $\mathcal{H}_0$ (i.e., $\mathsf{X}_\cdot$ obeys p_0) once $\Lambda_k \leq A$ and hypothesis $\mathcal{H}_1$ (i.e., $\mathsf{X}_\cdot$ obeys p_1) once $\Lambda_k \geq B$, is optimal in the sense that, among all possible stopping rules whose false alarm and miss probabilities are no worse than those attained by the SPRT, the SPRT requires the minimum expected stopping time under both hypotheses.

In many applications, it is imperative to attain a small expected stopping time under $\mathcal{H}_1$ (the so-called alternative hypothesis, which usually corresponds to an abnormal condition that requires immediate attention), but of less importance to stop promptly under $\mathcal{H}_0$ (the so-called null hypothesis, which usually corresponds to a normal condition). Under that scenario, the following one-sided SPRT is often considered [4]: Sequentially compare $\Lambda_k = \prod_{i=1}^{k} \frac{p_1(\mathsf{X}_i)}{p_0(\mathsf{X}_i)}$, $k = 1, 2, \ldots$, against a threshold $1 < B < \infty$, and decide hypothesis $\mathcal{H}_1$ once $\Lambda_k \geq B$. The one-sided SPRT stops with probability no greater than $1/B$ under $\mathcal{H}_0$, but stops with probability one under $\mathcal{H}_1$. Furthermore, the one-sided SPRT is optimal in the sense that, among all possible stopping rules whose false alarm probabilities are no worse than that attained by the one-sided SPRT, the one-sided SPRT requires the minimum expected stopping time under $\mathcal{H}_1$ [5, pp. 107-108].

Now, consider imposing a constraint on the maximum number of observed samples; that is, a statistician may at most observe N samples, where N may be a very large, but finite, integer. Such a constraint is reasonable since in any application it is impossible to spend an infinite amount of time to make a decision. For two-sided problems, truncated decision rules were considered in [6], [7] and references therein. For one-sided problems, recently a truncated decision rule termed an opportunistic detection rule (ODR) was considered in [8]: the statistician follows the one-sided SPRT described above, but decides $\mathcal{H}_0$ if the one-sided SPRT has not stopped before observing the last sample X_N.

In this paper, we present results that generalize and improve the ODR in [8], following two directions. In the first direction, we continue to consider the asymptotic regime with N growing without bound, and characterize the tradeoff among the exponents of the error probabilities (i.e., false alarm and miss probabilities) and the normalized expected stopping time under the alternative hypothesis. As an extreme case in the tradeoff, the asymptotic performance of the optimal fixed-sample-size (FSS) decision rule, prescribed by the Stein-Chernoff Lemma, i.e., an error exponent of $D(p_0\|p_1)$, is indeed achievable for any fixed target false alarm probability, with asymptotically vanishing normalized expected stopping time under $\mathcal{H}_1$. Note that the original truncated one-sided SPRT ODR considered in [8], which has a fixed threshold, achieves an exponent of only $C(p_0, p_1)$, the Chernoff information of (p_0, p_1). This comparison clearly indicates that the truncated one-sided SPRT is a strictly suboptimal ODR in an asymptotic sense. Furthermore, we prove that the established achievable tradeoff among the exponents of the error probabilities and the normalized expected stopping time under $\mathcal{H}_1$ is tight. A key idea of the proof makes use of the converse for the channel capacity per unit cost [9] [10].

In the second direction we turn to the finite regime, considering a Bayesian problem formulation. We seek to characterize the optimal ODR that minimizes a Bayesian risk over a finite horizon of fixed N. We show that the Bayesian optimal ODR is a sequence of likelihood ratio threshold tests, consistent with our prior experience in many other sequential decision problems. The Bayesian optimal likelihood ratio threshold tests have time-varying thresholds, which can be conveniently computed through backward recursions. We subsequently investigate an extended setup, in which the

The work of W. Zhang was supported by National Natural Science Foundation of China under grant 61379003, and SRFDP-RGC ERG Joint Research Scheme under grant 20133402140001; the work of H. V. Poor was supported by the U. S. National Science Foundation under Grants DMS-1118605 and ECCS-1343210.

978-1-4799-5186-4/14/$31.00 ©2014 IEEE

maximum sample size is no longer fixed, but is a random variable following a geometric distribution whose realization is revealed to the statistician only upon observing the last sample. This setup is suitable for modeling scenarios in which the observation process is subject to abrupt interruption, or in which an external controller (say a system operator in a smart grid system [11]), in a unanticipated manner, issues a command for prompt decision without further observation. We establish that, interestingly, the Bayesian optimal ODR in that setup is a sequence of likelihood ratio threshold tests with two different thresholds: the first threshold (called the "running threshold"), which is determined by solving a stationary state equation, is used when future samples are still available; the second threshold (called the "terminal threshold"), which is simply the ratio between the priors scaled by costs, is used when the statistician reaches the final sample and thus has to make a decision.

The remaining part of this paper is organized as follows. Section II presents results of the asymptotic analysis. Section III presents the finite-horizon Bayesian risk minimization problem and characterizes its solution. Section IV further treats the Bayesian risk minimization problem under the extended setup of random sample size. Finally, Section V concludes this paper. All the technical derivations are omitted and will be presented in a full version of this paper.

II. Asymptotic Tradeoff between Detection Performance and Stopping Time

In general, an ODR is described by a stopping time T adapted to the filtration generated by X_i, $i = 1, 2, \ldots, N$, and a terminal decision rule D indicating which of $\mathcal{H}_0$ and $\mathcal{H}_1$ the statistician believes the observations obey. Note that when $\mathsf{T} < N$ the decision is definitely $\mathcal{H}_1$, and the terminal decision rule D is invoked only when $\mathsf{T} = N$. For a sequence of ODRs indexed by the maximum sample size $N = 1, 2, \ldots$, we have an asymptotic tradeoff among three performance metrics: the exponential decay rate of the false alarm probability, the exponential decay rate of the miss probability, and the expected stopping time (normalized by N) under $\mathcal{H}_1$. Mathematically a performance tuple $(\Delta_{\mathrm{FA}}, \Delta_{\mathrm{M}}, \eta)$ is achievable if there exists a sequence of ODRs indexed by N, such that $\liminf_{N\to\infty} \frac{-\log P_{\mathrm{FA}}}{N} \geq \Delta_{\mathrm{FA}}$, $\liminf_{N\to\infty} \frac{-\log P_{\mathrm{M}}}{N} \geq \Delta_{\mathrm{M}}$, and $\limsup_{N\to\infty} \frac{T}{N} \leq \eta$, where $T = \mathbb{E}_1[\mathsf{T}]$ is the expected stopping time under $\mathcal{H}_1$.

Furthermore, we may call the closure of the union of achievable tuples under all possible ODRs the *ODR performance region*, which should depend solely upon (p_0, p_1). We denote the ODR performance region by $\mathcal{R}(p_0, p_1)$, which is a subset of $[0, \infty) \times [0, \infty) \times [0, 1] \subset \mathbb{R}^3$.

The following main theorem of this section fully characterizes $\mathcal{R}(p_0, p_1)$.

Theorem 1: The ODR performance region $\mathcal{R}(p_0, p_1)$ is given as follows: for each $0 \leq \eta \leq 1$,

$$\Delta_{\mathrm{FA}} \leq \min\left\{\eta d_1, \sup_{\alpha>0}\left\{\alpha\left[d_1 - \nu(d_0+d_1)\right] - \log \mathbb{E}_0\left[e^{\alpha \log \frac{p_1(\mathsf{X})}{p_0(\mathsf{X})}}\right]\right\}\right\},$$
$$\Delta_{\mathrm{M}} \leq \sup_{\alpha<0}\left\{\alpha\left[d_1 - \nu(d_0+d_1)\right] - \log \mathbb{E}_1\left[e^{\alpha \log \frac{p_1(\mathsf{X})}{p_0(\mathsf{X})}}\right]\right\}, \quad (1)$$

for $0 \leq \nu \leq 1$, where $d_0 = D(p_0\|p_1)$ and $d_1 = D(p_1\|p_0)$.

Theorem 1 is proved in two parts. The achievability part is established by constructing a specific form of ODRs that asymptotically achieve the performance tuple described in Theorem 1. The converse part is established by an argument of contradiction, in which a key idea is information-theoretic, basically asserting that, if the ODR performance region $\mathcal{R}(p_0, p_1)$ can be outperformed, then one can achieve a rate per unit cost higher than the capacity per unit cost [9] for a certain stationary memoryless channel, an impossible task even with feedback and variable-length coding [10].

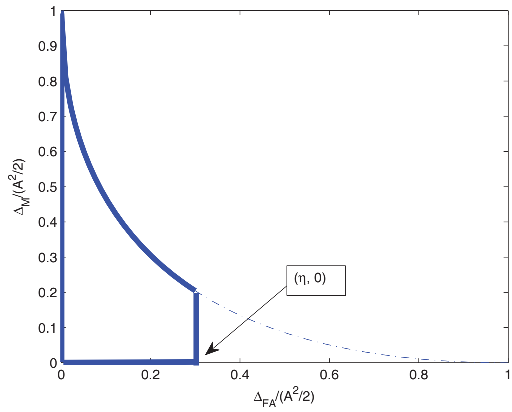

Fig. 1. An illustration of (3).

A. Example: Gaussian Distributions with and without a Drift

To illustrate $\mathcal{R}(p_0, p_1)$, we present an example for the following hypotheses:

$$\mathcal{H}_0 : p_0 \sim \mathcal{N}(0, 1) \quad \text{versus} \quad \mathcal{H}_1 : p_1 \sim \mathcal{N}(A, 1). \quad (2)$$

In this case we have $D(p_0\|p_1) = D(p_1\|p_0) = A^2/2$.

Then, applying Theorem 1, we can obtain the region $\left(\frac{\Delta_{\mathrm{FA}}}{A^2/2}, \frac{\Delta_{\mathrm{M}}}{A^2/2}\right)$ for each $0 \leq \eta \leq 1$, as (cf. Figure 1)

$$\left\{(x, y) : \sqrt{x} + \sqrt{y} \leq 1, 0 \leq x \leq \eta, y \geq 0\right\}. \quad (3)$$

The complete characterization of $\mathcal{R}(p_0, p_1)$ is given by the following and is illustrated in Figure 2.

Corollary 1: For the hypotheses (2), the ODR performance region $\mathcal{R}(p_0, p_1)$ is

$$\left\{\left(\frac{\Delta_{\mathrm{FA}}}{A^2/2}, \frac{\Delta_{\mathrm{M}}}{A^2/2}, \eta\right) = (x, y, z) : \sqrt{x} + \sqrt{y} \leq 1, 0 \leq x \leq z, y \geq 0, 0 \leq z \leq 1\right\}. \quad (4)$$

B. Stein-Chernoff Lemma Revisited

In this subsection, we focus on an extremal case of Theorem 1, in which the false alarm probability is fixed without decreasing toward zero exponentially, or, has an exponent of zero. For this case, Theorem 1 specializes into the following corollary.

Corollary 2: For an arbitrary fixed target false alarm probability $P_{\mathrm{FA}}^* > 0$, there exists a sequence of ODRs such that the miss probability scales toward zero, as N grows without bound, following

$$\lim_{N\to\infty} \frac{-\log P_{\mathrm{M}}}{N} = D(p_0\|p_1), \quad (5)$$

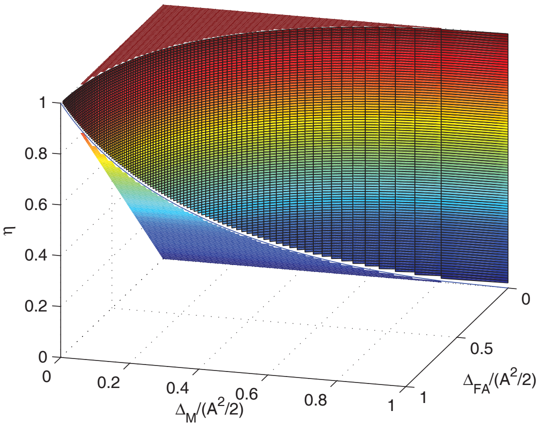

Fig. 2. An illustration of $\mathcal{R}(p_0, p_1)$ in (4).

and that the normalized expected stopping time under $\mathcal{H}_1$, T/N, satisfies

$$\lim_{N\to\infty} \frac{T}{N} = 0. \tag{6}$$

The case of Corollary 2 has been treated in [8], but therein the considered form of ODR is restricted to be a truncated one-sided SPRT, which was shown to behave asymptotically according to the following theorem.

Theorem 2: ([8, Thm. 1, Thm. 2]) For the truncated one-sided SPRT ODR that attains an arbitrary fixed target false alarm probability $0 < P^*_{\mathrm{FA}} \le P_0[p_1(\mathsf{X}) \ge p_0(\mathsf{X})]$, the miss probability scales toward zero as N grows without bound as

$$\lim_{N\to\infty} \frac{-\log P_{\mathrm{M}}}{N} = C(p_0, p_1), \tag{7}$$

where $C(p_0, p_1)$ is the Chernoff information of (p_0, p_1) (see [12] and [13, Ch. 11.9])

$$C(p_0, p_1) = -\inf_{\alpha\in(0,1)} \log \left(\int_{\mathcal{X}} p_0^{\alpha}(x) p_1^{1-\alpha}(x) dx \right), \tag{8}$$

and the normalized expected stopping time under $\mathcal{H}_1$, T/N, satisfies

$$\lim_{N\to\infty} \frac{T}{N} = 0. \tag{9}$$

Comparing Theorem 2 and Corollary 2, we can clearly see that, at asymptotically diminishing sampling cost under $\mathcal{H}_1$, there exists an ODR that outperforms the truncated one-sided SPRT ODR in [8]. Furthermore, the exponent achieved in Corollary 2, $D(p_0\|p_1)$, is exactly that achieved by the optimal FSS decision rule as indicated in the Stein-Chernoff Lemma, but here the corresponding ODR is not FSS, requiring asymptotically diminishing sampling cost under $\mathcal{H}_1$. So in other words, the FSS sampling cost is not fundamental in achieving the Stein-Chernoff Lemma, which appears to be a new and somewhat surprising finding.

III. Bayesian Optimal ODR: Solution Structure and Recursions

In the subsequent part of the paper, we depart from asymptotic analysis and examine the finite regime, wherein we consider Bayesian formulations of the ODR problem. In order to make use of the optimal stopping theory, it turns out to be convenient to formulate the problem in the following way. Consider all stopping times that stop by N, $\mathcal{T}^N$, and terminal decision rules $\mathcal{D} : \mathcal{X}^N \mapsto \{\mathcal{H}_0, \mathcal{H}_1\}$. Note that for ODRs we need to consider only the terminal decision rule, because whenever $\mathsf{T} < N$ the decision is bound to be $\mathcal{H}_1$.

The detection error events can thus be written as

$$\text{False alarm:} \quad \{\mathsf{T} < N\} \cup \{\mathsf{T} = N, D = \mathcal{H}_1\} \text{ w.r.t. } p_0$$
$$\text{Miss:} \quad \{\mathsf{T} = N, D = \mathcal{H}_0\} \text{ w.r.t. } p_1,$$

and the expected stopping time under $\mathcal{H}_1$ is $T = \mathbb{E}_1[\mathsf{T}]$.

We thus formulate the the Bayesian risk as follows:

$$\begin{aligned} \mathcal{J} &= (1-\pi)c_0 P_{\mathrm{FA}} + \pi c_1 P_{\mathrm{M}} + cT \\ &= (1-\pi)c_0 \mathbb{E}_0 \left[\mathbf{1}(\mathsf{T} < N) + \mathbf{1}(\mathsf{T} = N)\mathbf{1}(D = \mathcal{H}_1)\right] \\ &\quad + \pi c_1 \mathbb{E}_1 \left[\mathbf{1}(\mathsf{T} = N)\mathbf{1}(D = \mathcal{H}_0)\right] + c\mathbb{E}_1[\mathsf{T}], \end{aligned} \tag{10}$$

where $0 \le \pi \le 1$ is the prior probability of hypothesis $\mathcal{H}_1$, and $c_0, c_1, c > 0$ are cost assignments. The problem we seek to solve is then to choose a stopping time T and a terminal decision rule D that minimize $\mathcal{J}$, i.e.,

$$\min_{\mathsf{T}\in\mathcal{T}^N, D} \mathcal{J}. \tag{11}$$

The problem (11) can be cast and solved as a Markov optimal stopping problem [3] [14]. For its solution, we have the following main theorem of this section.

Theorem 3: The Bayesian optimal ODR that solves (11) is a sequence of likelihood ratio threshold tests, with time-varying thresholds.

The thresholds τ_k are the solutions of

$$c\lambda + \mathbb{E}_0[h_{k+1}(\lambda \mathsf{L})] = (1-\pi)c_0, \tag{12}$$

for $k = 1, 2, \ldots, N-1$, where the functions $\{h_k\}$ satisfy the following backward recursion relationship:

$$h_{k-1}(\lambda) = \min\{(1-\pi)c_0, c\lambda + \mathbb{E}_0[h_k(\lambda \mathsf{L})]\}, \tag{13}$$

for $k = N, N-1, \ldots, 2$, where $\mathsf{L} = p_1(\mathsf{X})/p_0(\mathsf{X})$ with X following p_0, and $h_N(\lambda) = \min\{(1-\pi)c_0, \pi c_1 \lambda\}$; the terminal threshold τ_N is simply $\frac{(1-\pi)c_0}{\pi c_1}$.

Discussion: Note that the sequence of thresholds, $\{\tau_k\}$, is generally time-varying, due to the finiteness of N. For sufficiently large N, however, the backward recursion (13) becomes asymptotically stationary, and thus the thresholds $\{\tau_k\}$ become approximately constant except for those k's close to N. Furthermore, the thresholds $\{\tau_k\}$ can be computed offline, for some sufficiently large N, and in practice when the maximum sample size N' is smaller than N, only the last N' thresholds are used in the corresponding ODR.

The following algorithm implements the Bayesian optimal ODR under maximum sample size N:

Bayesian Optimal ODR under Maximum Sample Size N

Initial parameters: Hypotheses p_0, p_1 and prior π, maximum sample size N, cost assignments c_0, c_1, c.

Set: A sequence of thresholds $\{\tau_k\}_{k=1}^N$ computed via (12) and $\tau_N = \frac{(1-\pi)c_0}{\pi c_1}$.

Algorithm:

initialize $n = 1$;
while $n \le N$
  do compute Λ_n;
  if $\Lambda_n \ge \tau_n$
    terminate returning $\mathcal{H}_1$;
  else $n = n + 1$;

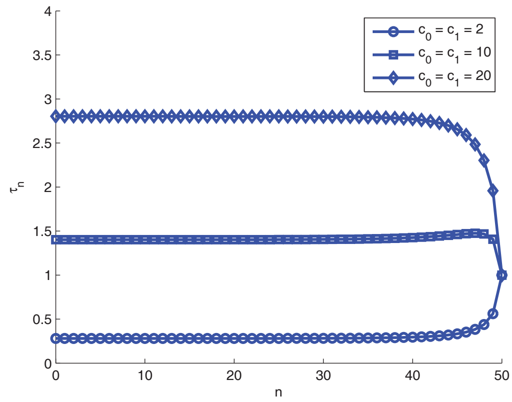

Fig. 3. Numerical results for the thresholds $\{\tau_n\}_{n=1}^N$, for different c_0 and c_1.

end if
end while
terminate returning $\mathcal{H}_0$;

We illustrate the numerical behavior of the optimal ODR by the case study of testing the hypotheses

$$\mathcal{H}_0 : p_0 \sim \mathcal{N}(0, 1) \quad \text{versus} \quad \mathcal{H}_1 : p_1 \sim \mathcal{N}(A, 1), \tag{14}$$

In the presented numerical examples (Figure 3), we set $A = 1$, $\pi = 1/2$, $c = 1$, $N = 50$, and let $c_0 = c_1$ be 2, 10, and 20 respectively. In the plots we display the thresholds $\{\tau_n\}_{n=1}^N$, computed from the preceding algorithm for the Bayesian optimal ODR. From the plots, we observe that the backward recursion quickly leads to stationary thresholds, within ten samples (returning from $n = N$). However, depending upon the values of the tuning parameters (here the effective ones are $c_0 = c_1$), the evolution trend of the thresholds may differ considerably. In the plots, when $c_0 = c_1 = 2$ the sequence $\{\tau_n\}$ increases with n, when $c_0 = c_1 = 20$ it decreases with n, and when $c_0 = c_1 = 10$ there further exists a slight "overshoot" behavior. Intuitively, for small c_0 and c_1, the importance of reducing the expected stopping time under $\mathcal{H}_1$ outweighs that of decreasing the decision error probabilities, and hence it is reasonable to promote early stopping by using lower decision thresholds for early samples; on the contrary, for large c_0 and c_1, the priority is on decreasing the decision error probabilities, and hence it is reasonable to set relatively high decision thresholds for early samples, in order to avoid premature error-prone decisions.

IV. Bayesian Optimal ODR for Random Maximum Sample Size

In this section, we consider an extended setup, in which the maximum sample size is no longer fixed, but a random variable. To motivate this setup, imagine the scenario in which an automated surveillance system keeps monitoring a process and possibly issues a detection of some exceptional condition, and meanwhile a human operator, in a unanticipated manner, may interrupt the system at some random time, asking for a prompt decision without further monitoring.

Specifically, along with the i.i.d. random variable sequence X_i, $i = 1, 2, \ldots$, following p_0 and p_1 with prior probabilities $1 - \pi$ and π respectively, we introduce an additional random variable N, which is independent of $\mathsf{X}_1, \mathsf{X}_2, \ldots$, and follows a geometric distribution with parameter $0 < \epsilon < 1$, i.e., $\Pr[\mathsf{N} = n] = (1-\epsilon)^{n-1}\epsilon$, $n = 1, 2, \ldots$ We view N as the maximum sample size. In the model considered herein, the realization of N is not revealed to the statistician until observing X_N: of course if the statistician has already made his opportunistic detection before observing X_N, there is no need to know about N any more; otherwise if the statistician has reached X_N without a detection yet, then he is required to make his decision immediately with the first N samples at hand, without observing any extra samples.

Due to its geometric distribution, N can be conveniently interpreted as the first time an i.i.d. sequence of Bernoulli trials (with success probability ϵ) returns success. So alternatively N is a stopping time defined as $\mathsf{N} = \min\{n : \mathsf{Z}_n = 1\}$, where Z_n is an i.i.d. sequence of Bernoulli random variables with $\Pr[\mathsf{Z}_1 = 1] = \epsilon$ and $\Pr[\mathsf{Z}_1 = 0] = 1 - \epsilon$. Therefore, for any stopping time T' that is adapted to the filtration generated by $\mathsf{X}_1, \mathsf{X}_2, \ldots$, if we define $\mathsf{T} = \min\{\mathsf{T}', \mathsf{N}\}$, then T is a stopping time adapted to the product filtration generated by $(\mathsf{X}_1, \mathsf{Z}_1), (\mathsf{X}_2, \mathsf{Z}_2), \ldots$ With a thus induced T and an arbitrarily given terminal decision rule D, the statistician declares $\mathcal{H}_1$ if either $\{\mathsf{T} < \mathsf{N}\}$ or $\{\mathsf{T} = \mathsf{N}, D = \mathcal{H}_1\}$ occurs, and declares $\mathcal{H}_0$ if $\{\mathsf{T} = \mathsf{N}, D = \mathcal{H}_0\}$ occurs.

Similar to the problem framework in Section III, we define the Bayesian risk as

$$\mathcal{J} = (1-\pi)c_0 P_{\mathrm{FA}} + \pi c_1 P_{\mathrm{M}} + c\mathbb{E}_1[\mathsf{T}], \tag{15}$$

where $c_0, c_1, c > 0$ are cost assignments, and the problem is to choose T and D to minimize $\mathcal{J}$. Note that here the stopping time is not bounded since N can be arbitrarily large.

After some manipulation, we find that the Bayesian risk (15) is equivalent to the following form:

$$\mathcal{J} = \mathbb{E}_0\left[(1-\epsilon)^{\mathsf{T}} g(\Lambda_\mathsf{T}) + \sum_{n=0}^{\mathsf{T}-1}(1-\epsilon)^n c(\Lambda_n)\right], \tag{16}$$

where

$$g(\lambda) = (1-\pi)c_0 + \frac{\epsilon}{1-\epsilon}\min\{(1-\pi)c_0, \pi c_1 \lambda\}, \tag{17}$$

and $c(\lambda) = c\lambda$, for $\lambda \geq 0$. This is exactly the form that has been treated in [14, 2.14], considering both an instantaneous reward at the stopping time and accumulated sampling costs, with everything discounted by an exponential factor $(1-\epsilon)^k$ at time k.

From [14, Thm. 23], we have that the Bayesian optimal stopping time is given by

$$\mathsf{T} = \min\{n \geq 1 : V(\Lambda_n) = g(\Lambda_n)\}, \tag{18}$$

where $V(\cdot)$ is the solution of

$$V(\lambda) = \min\{g(\lambda), (1-\epsilon)\mathbb{E}_0[V(\lambda\mathsf{L})] + c(\lambda)\}, \tag{19}$$

with $\mathsf{L} = p_1(\mathsf{X})/p_0(\mathsf{X})$, X following p_0. Furthermore, $V(\cdot)$ may be computed as $V(\lambda) = \lim_{n\to\infty} \mathcal{Q}^n g(\lambda)$, with the operator $\mathcal{Q}$ defined by

$$\mathcal{Q}f(\lambda) = \min\{f(\lambda), (1-\epsilon)\mathbb{E}_0[f(\lambda\mathsf{L})] + c(\lambda)\}. \tag{20}$$

The following algorithm implements the Bayesian optimal ODR under a geometrically distributed maximum sample size:

Bayesian Optimal ODR under Geometrically Distributed Maximum Sample Size

Initial parameters: Hypotheses p_0, p_1 and prior π, mean sample size $1/\epsilon$, cost assignments c_0, c_1, c.
Set: A "running" threshold τ_r as the value of λ at the intersection of $g(\lambda)$ and $(1-\epsilon)\mathbb{E}_0[V(\lambda \mathrm{L})] + c(\lambda)$, and a "terminal" threshold τ_t equal to $\frac{(1-\pi)c_0}{\pi c_1}$.
Algorithm:

```
initialize n = 1;
while N has not been revealed
    do compute Λ_n;
    if Λ_n ≥ τ_r
        terminate returning H_1;
    else n = n + 1;
    end if
end while
if Λ_N ≥ τ_t   terminate returning H_1;
else   terminate returning H_0;
end if
```

For the optimal ODR, an interesting property is that it is a two-threshold scheme: the "running" threshold τ_r, which is determined by solving the stationary state equation (19), is used to compare with the likelihood ratio sequence before N, i.e., when future samples are still available; and the "terminal" threshold τ_t, which is simply the ratio between the priors scaled by costs, is used only at the end, i.e., when the statistician is informed that the final sample has been reached and a decision is required immediately. Such a two-threshold scheme is very different from the conventional one-sided and two-sided SPRTs, in which the thresholds are fixed constants throughout.

We use the same case study as that considered in Section III to illustrate the numerical behavior of the optimal ODR under random maximum sample size. Again we set $A = 1$, $\pi = 1/2$, and $c = 1$. For the geometric distribution of N, we set $\epsilon = 0.05$, so that the mean maximum sample size is 20. Note that $g(\lambda)$ is a piecewise linear function of λ with one switching point exactly at $\lambda = \tau_t$; so depending on at which segment the curve $(1-\epsilon)\mathbb{E}_0[V(\lambda \mathrm{L})] + c(\lambda)$ intersects $g(\lambda)$, there are two possible situations: $\tau_r \geq \tau_t$, and $\tau_r < \tau_t$. In Figure 4 we plot the trend of τ_r as $c_0 = c_1$ increases from 0.2 to 16. We observe that τ_r increases with c_0 and c_1, crossing the level of τ_t. Interestingly, the growth trend of τ_r is virtually linear with c_0 and c_1.

V. Conclusion

In this paper, we have formulated the general ODR framework and treated several of its key characteristics, in both asymptotic and finite regimes. An interesting problem beyond the scope of this paper concerns the asymptotic analysis of the Bayesian optimal ODR. In such problems in the sequential analysis literature, one usually proceeds by letting the sampling cost c decrease toward zero in the Bayesian risk; see, e.g., [15, Sec. 13]. For our setup, in order to make the problem meaningful, we may need to tune the growth of the maximum sample size N (or the mean maximum sample size $1/\epsilon$ in the random sample size case in Section IV) accordingly, say, following $O(1/c)$. The interplay between the sampling cost and the maximum sample size thus may exhibit interesting behaviors. Another interesting direction is to develop the ODR framework for continuous-time stochastic processes.

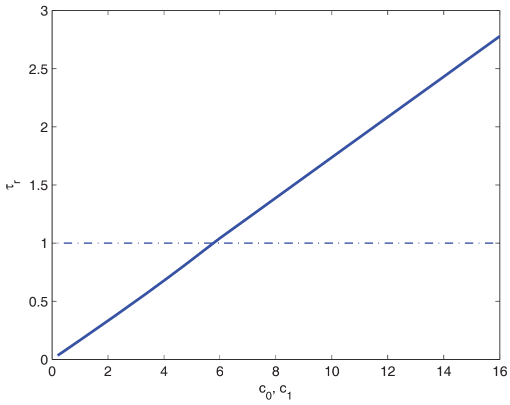

Fig. 4. The running threshold τ_r versus c_0 and c_1. The dash-dot line indicates the terminal threshold τ_t.

References

[1] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," *Ann. Math. Statist.*, vol. 19, pp. 326-339, 1948.
[2] T. S. Ferguson, *Mathematical Statistics: A Decision Theoretic Approach*, New York, NY: Academic Press, 1967.
[3] H. V. Poor and O. Hadjiliadis, *Quickest Detection*, Cambridge, UK: Cambridge University Press, 2009.
[4] G. Fellouris and A. G. Tartakovsky, "Nearly minimax mixture rules for one-sided sequential testing," *Sequent. Anal.*, vol. 31, no. 3, pp. 297-325, 2012.
[5] Y. S. Chow, H. Robbins, and D. Siegmund, *Great Expectations: The Theory of Optimal Stopping*, Boston, MA: Houghton Mifflin, 1971.
[6] J. J. Bussgang and M. B. Marcus, "Truncated sequential hypothesis tests," *IEEE Trans. Inform. Theory*, vol. 13, no. 3, pp. 512-516, Jul. 1967.
[7] S. Tantaratana and H. V. Poor, "Asymptotic efficiencies of truncated sequential tests," *IEEE Trans. Inform. Theory*, vol. 28, no. 6, pp. 911-923, Nov. 1982.
[8] W. Zhang and H. V. Poor, "Opportunistic detection under a fixed-sample-size setting," *IEEE Trans. Inform. Theory*, vol. 59, no. 2, pp. 1107-1114, Feb. 2013.
[9] S. Verdú, "On channel capacity per unit cost," *IEEE Trans. Inform. Theory*, vol. 36, no. 5, pp. 1019-1030, Sep. 1990.
[10] I. Csiszár, "On the capacity of noisy channels with arbitrary signal costs," *Prob. Contr. Inform. Theory*, 2, pp. 283-304, 1973.
[11] E. Hossein, Z. Han and H. V. Poor, *Smart Grid Communications and Networking*, Cambridge, UK: Cambridge University Press, 2012.
[12] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations," *Ann. Math. Statist.*, vol. 23, no. 4, pp. 493-507, 1952.
[13] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., Hoboken, NJ: Wiley, 2006.
[14] A. N. Shiryaev, *Optimal Stopping Rules*, New York, NY: Springer-Verlag, 1978.
[15] H. Chernoff, *Sequential Analysis and Optimal Design*, Philadelphia, PA: SIAM, 1972.