

Opportunistic Detection Rules: Finite and Asymptotic Analysis

Wenyi Zhang, *Senior Member, IEEE*, George V. Moustakides, *Senior Member, IEEE*,
and H. Vincent Poor, *Fellow, IEEE*

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 results regarding ODRs are established in this paper. In the finite regime, the maximum sample size is modeled either as a fixed finite number, or a geometric random variable with a fixed finite mean. For both cases, the corresponding Bayesian formulations are investigated. The former case is a slight variation of the well-known finite-length sequential hypothesis testing procedure in the literature, whereas the latter case is new, for which the Bayesian optimal ODR is shown to be a sequence of likelihood ratio threshold tests with two different thresholds. A running threshold, which is determined by solving a stationary state equation, is used when future samples are still available, and a terminal threshold (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 immediately. In the asymptotic regime, the tradeoff among the exponents of the (false alarm and miss) error probabilities and the normalized expected stopping time under the alternative hypothesis is completely characterized and proved to be tight, via an information-theoretic argument. Within the tradeoff region, one noteworthy fact is that the performance of the Stein-Chernoff lemma is attainable by ODRs.

Index Terms—Chernoff information, error exponent, fixed-sample-size (FSS) hypothesis testing, opportunistic detection rule (ODR), optimal stopping, sequential hypothesis testing, Stein-Chernoff lemma.

I. INTRODUCTION

IN THIS paper we consider a setup of discriminating two simple hypotheses, as follows. At most N independent

Manuscript received September 23, 2014; revised October 30, 2015; accepted January 19, 2016. Date of publication February 15, 2016; date of current version March 16, 2016. W. Zhang was supported in part by the National Natural Science Foundation of China under Grant 61379003, in part by the National Basic Research Program of China (973 Program) under Grant 2012CB316004, and in part by the SRFDP-RGC ERG Joint Research Scheme through the Specialized Research Fund under Grant 20133402140001. H. V. Poor was supported by the U.S. National Science Foundation under Grant DMS-1118605 and Grant ECCS-1343210.

W. Zhang is with the Department of Electronic Engineering and Information Sciences, University of Science and Technology of China, Hefei 230027, China (e-mail: wenyizha@ustc.edu.cn).

G. V. Moustakides is with the Department of Electrical and Computer Engineering, University of Patras, 26500 Rio, Greece (e-mail: moustaki@upatras.gr).

H. V. Poor is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: poor@princeton.edu).

Communicated by I. Nikiforov, Associate Editor for Detection and Estimation.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2016.2530087

and identically distributed (i.i.d.) random variables (called “samples” interchangeably throughout the paper) X_i , $i = 1, 2, \dots, N$, from one of two distributions with probability densities or mass functions p_0 (the null hypothesis \mathcal{H}_0) and p_1 (the alternative hypothesis \mathcal{H}_1), are drawn sequentially one at a time by a statistician, who is allowed to stop sampling at some time $T < N$ and decide \mathcal{H}_1 , or wait until observing all the N samples and decide either \mathcal{H}_0 or \mathcal{H}_1 . A strategy adopted by the statistician is called an “opportunistic detection rule” (ODR).

To motivate this ODR setup, on one hand, note that it is often imperative to attain a small decision delay under \mathcal{H}_1 (which may correspond to an abnormal condition that requires an immediate attention), but of less importance to stop promptly under \mathcal{H}_0 (which may correspond to a normal condition),¹ on the other hand, the finite maximum sample size is reasonable since in any application it is unlikely to have an infinite number of samples or to spend an infinite amount of time to make a decision.

Such an ODR setup is closely related to both fixed-sample-size (FSS) hypothesis testing and sequential hypothesis testing. Subsequently we briefly discuss the background to place our work into context.

A. Related Work

1) *FSS Hypothesis Testing*: Discriminating two distributions based on an array of i.i.d. samples is of fundamental importance in statistical decision theory. Of particular interest, characterizing the asymptotic performance limit, in terms of achievable exponential decay rates of the (false alarm and miss) error probabilities as the sample size grows without bound, has received significant attention in the literature. The Stein-Chernoff Lemma [1], [2] characterizes the maximum achievable exponential decay rate of the miss probability for any fixed false alarm probability, under a Neyman-Pearson setting. Under a Bayesian setting which seeks to minimize a linear combination of the false alarm and miss probabilities, the Chernoff information [3] measures the maximum achievable exponential decay rate of the Bayesian cost. More generally, the optimal tradeoff between the exponential decay rates of the false alarm probability and the miss probability was originally studied in [4] and later treated in information theory [5], [6]. In the asymptotic analysis of the present paper,

¹Other application scenarios are also possible. For example, consider inspecting a number of products that were made by the same machine, to determine whether they have some common flaw or not. Inspecting a product will destroy it. Therefore, it is desirable to inspect as few products as possible if they do not have the flaw, but immaterial if they do.

we further extend the problem to study the tradeoff among the exponential decay rates of the false alarm probability, the miss probability, and the expected stopping time (normalized by the maximum sample size N) under the alternative hypothesis, because here for ODR it is of interest to reduce the decision delay under the alternative hypothesis, at the cost of sacrificing the decision reliability.

Beyond the scope of the present paper, later developments along this direction include the strong converse [7] which characterizes the exponential decay rate of the rejection probability (i.e., one minus the false alarm probability) when the exponential decay rate of the miss probability exceeds that indicated by the Stein-Chernoff Lemma, and generalization for discriminating two general (non-i.i.d.) sequences of distributions [8], [9].

2) *Sequential Hypothesis Testing*: When there is no limit on the maximum sample size and the statistician is allowed to wait indefinitely before making a decision, the Wald-Wolfowitz Theorem (see, e.g., [10] [11, Sec. 7.6] [12, Th. 4.7]) is fundamental. It asserts that the sequential probability ratio test (SPRT), which sequentially compares the likelihood ratios $\Lambda_k = \prod_{i=1}^k \frac{p_1(X_i)}{p_0(X_i)}$, $k = 1, 2, \dots$, against two thresholds $0 < A \leq 1 \leq B < \infty$, and decides \mathcal{H}_0 once $\Lambda_k \leq A$ or \mathcal{H}_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 times under both hypotheses.²

The requirement in our ODR setup that the statistician stops sampling only when deciding \mathcal{H}_1 is closely related to the so-called power-one or one-sided stopping rules, extensively studied in statistics for testing simple and composite hypotheses; see, e.g., [13]–[16]. A common flavor of these earlier works is the emphasis on fine asymptotic behavior of the expected stopping time, usually in the form of the law of the iterated logarithm. A simple one-sided stopping rule is the one-sided SPRT, which sequentially compares $\Lambda_k = \prod_{i=1}^k \frac{p_1(X_i)}{p_0(X_i)}$, $k = 1, 2, \dots$, against a threshold $1 < B < \infty$, and decides \mathcal{H}_1 once $\Lambda_k \geq B$. Note that, the one-sided SPRT may never stop — indeed, it stops with probability no greater than $1/B$ under \mathcal{H}_0 , but stops with probability one under \mathcal{H}_1 [17]. 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 [18, pp. 107–108].

The requirement in our ODR setup that the statistician must make a decision upon reaching a finite deadline N essentially defines a truncated sequential hypothesis testing problem. The previous study of truncated stopping rules usually considers the performance of composite hypothesis testing with a pair of time-varying threshold sequences that coincide at the truncation point; see, e.g., [19]–[22] and references therein. Meanwhile, truncated stopping rules are

often used as an intermediate step when developing stopping rules without a deadline, by passing to the limit of $N \rightarrow \infty$; see, e.g., [23, Sec. 5.5] [24, Sec. 3.2.2] for the derivation of the SPRT. When combining the one-sided and truncated stopping rules, a specific ODR was recently considered and analyzed in [25]: the statistician follows the one-sided SPRT, but decides \mathcal{H}_0 if the one-sided SPRT has not stopped before observing the last sample X_N . In the finite analysis of the present paper, we develop Bayesian optimal ODRs, for both the case where the maximum sample size is a fixed finite number, and the case where the maximum sample size is a geometric random variable with a fixed finite mean.

In the context of sequential hypothesis testing, asymptotic optimality usually means that the mean (and higher moments, possibly) of the stopping time of a test is close, in a certain asymptotic sense, to that attained by the optimal test with the same error probabilities, which is usually difficult to design analytically or even numerically; see, e.g., [24], [26], and [27].³

B. Overview of Results

In this paper, we establish several results regarding ODRs.

1) *Finite Analysis*: In the finite regime, we consider a Bayesian problem formulation, which seeks to characterize the ODR that minimizes a Bayesian cost as a linear combination of the (false alarm and miss) error probabilities and the expected stopping time. We examine two cases, in which the maximum sample size is modeled either as a fixed finite number, or a geometric random variable with a fixed finite mean, respectively. The former case is a slight variation of the well-known finite-length sequential hypothesis testing procedure in the literature (see, e.g., [23, Sec. 5.5] [24, Sec. 3.2.2]) and is included herein for making the exposition self-contained. The latter case is new. To motivate the model of a random maximum sample size, whose realization is revealed to the statistician only upon observing the last sample, we may consider 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 [29]), in an unanticipated manner, issues a command for a prompt decision without further observation. For this case, we establish that, interestingly, the Bayesian optimal ODR is a sequence of likelihood ratio threshold tests with two different thresholds: a “running threshold”, which is determined by solving a stationary state equation, is used when future samples are still available, and a “terminal threshold” (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.

2) *Asymptotic Analysis*: In the asymptotic regime, we let the maximum sample size N grow without bound, and characterize the tradeoff among the exponential decay rates of the (false alarm and miss) error probabilities and the normalized expected stopping time under the alternative hypothesis. As an

²There is a technical condition that the sum of the false alarm and miss probabilities is no greater than one.

³Relatively few works consider the exponential decay rates of the (false alarm and miss) error probabilities for sequential hypothesis testing. The optimal tradeoff for sequential binary simple hypothesis testing has been characterized in [28], but the definition of the exponential decay rates therein is different from that considered in FSS hypothesis testing. In the present paper, we follow the definition in FSS hypothesis testing.

extreme case in the tradeoff, the asymptotic performance of the optimal fixed-sample-size (FSS) decision rule, described by the Stein-Chernoff Lemma, i.e., an error exponent of $D(p_0 \| p_1)$ for the miss probability, is shown to be achievable for any fixed target false alarm probability, with asymptotically vanishing normalized expected stopping time under \mathcal{H}_1 . The truncated one-sided SPRT ODR considered in [25] is thus suboptimal since it achieves an exponent of only $C(p_0, p_1)$, the Chernoff information of (p_0, p_1) . When establishing the converse part, i.e., proving that the tradeoff is tight, a key idea of the proof makes use of the converse for the channel capacity per unit cost [30], [31].

The remaining part of this paper is organized as follows. Section II presents the finite analysis, characterizing the Bayesian optimal ODRs, for both the case of fixed and geometrically distributed maximum sample sizes. Section III presents results of the asymptotic analysis, fully characterizing the tradeoff among the exponents of the false alarm and miss probabilities and the expected stopping time under \mathcal{H}_1 . Finally, Section IV concludes this paper.

II. FINITE ANALYSIS

In this section, we consider the finite regime, in which the maximum sample size is modeled either as a fixed finite number, or a geometric random variable with a fixed finite mean. For both cases, we investigate the corresponding Bayesian formulations.

A. Case of Fixed Maximum Sample Size

In this case, the maximum sample size is a fixed finite number N . 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 the stopping time $T < N$ the decision is bound to be \mathcal{H}_1 . Therefore, an ODR can be defined as follows.

Definition 1: An opportunistic detection rule consists of a stopping time $T \in \mathcal{T}^N$, and a terminal decision rule $D : \mathcal{X}^N \mapsto \{\mathcal{H}_0, \mathcal{H}_1\}$, such that, when $T < N$, the decision is \mathcal{H}_1 , and when $T = N$, the decision is given by $D(X^N)$. Herein X^N denotes X_1, X_2, \dots, X_N , and thereafter we usually suppress it and simply write $D(X^N)$ as D .

The detection error events can thus be written as

$$\text{False alarm: } \{T < N\} \cup \{T = N, D = \mathcal{H}_1\} \text{ w.r.t. } p_0 \quad (1)$$

$$\text{Miss: } \{T = N, D = \mathcal{H}_0\} \text{ w.r.t. } p_1, \quad (2)$$

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

We thus formulate the Bayesian cost as follows:

$$\begin{aligned} \mathcal{J} &= (1 - \pi)c_0 P_{\text{FA}} + \pi c_1 P_{\text{M}} + cT \\ &= (1 - \pi)c_0 \mathbb{E}_0[\mathbf{1}(T < N) + \mathbf{1}(T = N)\mathbf{1}(D = \mathcal{H}_1)] \\ &\quad + \pi c_1 \mathbb{E}_1[\mathbf{1}(T = N)\mathbf{1}(D = \mathcal{H}_0)] + c\mathbb{E}_1[T] \end{aligned} \quad (3)$$

where $0 \leq \pi \leq 1$ 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.,

$$\min_{T \in \mathcal{T}^N, D} \mathcal{J}. \quad (4)$$

The problem (4) is a slight variation of the well-known finite-length sequential hypothesis testing problem, which is often used as an intermediate step when developing stopping rules without a maximum sample size constraint; see, e.g., [23, Sec. 5.5] or [24, Sec. 3.2.2]. The main difference is that for the ODR problem (4), the decision is one-sided so that only a terminal decision rule upon observing the last sample is needed and that the false alarm event (1) is different from that considered in the literature. On the other hand, the derivation of the one-sided SPRT in [18, pp. 107–108] directly works with the non-truncated case so that it is not applicable to (4) here. Having formulated (4), the remaining analysis is a standard exercise of Markov optimal stopping theory, and we include the solution in the remainder of this subsection for the sake of making the exposition self-contained.

We note the following relationship through change of probability measure that transforms $\mathbb{E}_1[\cdot]$ into $\mathbb{E}_0[\cdot]$:

$$\begin{aligned} \mathbb{E}_1[\mathbf{1}(T = N)\mathbf{1}(D = \mathcal{H}_0)] \\ = \mathbb{E}_0[\Lambda_N \mathbf{1}(T = N)\mathbf{1}(D = \mathcal{H}_0)], \end{aligned} \quad (5)$$

because the miss event $\{T = N, D = \mathcal{H}_0\}$ is \mathcal{F}_N -measurable. On the other hand,

$$\begin{aligned} \mathbb{E}_1[T] &= \mathbb{E}_1\left[\sum_{k=1}^N \mathbf{1}(T \geq k)\right] \\ &= \sum_{k=1}^N \mathbb{E}_1[\mathbf{1}(T \geq k)] \\ &\stackrel{(a)}{=} \sum_{k=1}^N \mathbb{E}_0[\Lambda_{k-1} \mathbf{1}(T \geq k)] \\ &= \mathbb{E}_0\left[\sum_{k=1}^N \Lambda_{k-1} \mathbf{1}(T \geq k)\right] \\ &\stackrel{(b)}{=} \mathbb{E}_0\left[\sum_{k=0}^{T-1} \Lambda_k\right], \end{aligned} \quad (6)$$

where (a) is due to the fact that $\{T \geq k\}$ is \mathcal{F}_{k-1} -measurable, and (b) is due to the fact that $\mathbf{1}(T \geq k) = 0$ for all k 's greater than T . So with (5) and (6), we rewrite the Bayesian cost \mathcal{J} as

$$\begin{aligned} \mathcal{J} &= \mathbb{E}_0\left[(1 - \pi)c_0[\mathbf{1}(T < N) + \mathbf{1}(T = N)\mathbf{1}(D = \mathcal{H}_1)] \right. \\ &\quad \left. + \pi c_1 \Lambda_N \mathbf{1}(T = N)\mathbf{1}(D = \mathcal{H}_0) + c \sum_{k=0}^{T-1} \Lambda_k\right]. \end{aligned} \quad (7)$$

Inspecting (7), it is clear that for any stopping time T , the optimal terminal decision rule is

$$\begin{aligned} D &= \mathcal{H}_1 \quad \text{if } \Lambda_N \geq \frac{(1 - \pi)c_0}{\pi c_1}; \\ \text{and } D &= \mathcal{H}_0 \quad \text{otherwise.} \end{aligned} \quad (8)$$

This is a quite reasonable result in retrospect, since it coincides with the FSS Bayesian optimal decision rule, which should be the case when the statistician has already observed the last sample.

With the optimal terminal decision rule above, we further rewrite the Bayesian cost \mathcal{J} as

$$\mathcal{J} = \mathbb{E}_0 \left[c \sum_{k=0}^{T-1} \Lambda_k + (1 - \pi)c_0 \mathbf{1}(T < N) + \mathbf{1}(T = N) \min\{(1 - \pi)c_0, \pi c_1 \Lambda_N\} \right]. \quad (9)$$

Now let us characterize the stopping time $T \in \mathcal{T}^N$ that minimizes \mathcal{J} . For (9), using backward induction (see, e.g., [12, 3.3.1]), we find that the optimal stopping time is given by

$$T = \min \{1 \leq k \leq N - 1 : h_k(\Lambda_k) = (1 - \pi)c_0\}, \quad (10)$$

and if no such T exists, we set $T = N$. The functions $\{h_k\}$ satisfy a backward recursion as

$$h_{k-1}(\lambda) = \min\{(1 - \pi)c_0, c\lambda + \mathbb{E}_0[h_k(\lambda L)]\}, \quad (11)$$

$k = N, N - 1, \dots, 2$, where $L = p_1(X)/p_0(X)$ with X obeying p_0 , and $h_N(\lambda) = \min\{(1 - \pi)c_0, \pi c_1 \lambda\}$.

From the backward recursion (11), we have that the optimal ODR is indeed a sequence of likelihood ratio threshold tests.

Theorem 1: The Bayesian optimal ODR that solves (4) is a sequence of likelihood ratio threshold tests, with time-varying thresholds. The thresholds τ_k are the solutions of

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

for each $k = 1, 2, \dots, N - 1$, and $\tau_N = \frac{(1-\pi)c_0}{\pi c_1}$.

Proof: Theorem 1 follows from the observations below.

(i) For every k , $h_k(\lambda)$ is monotonically non-decreasing with respect to $\lambda > 0$, and is positive except at $\lambda = 0$. This can be directly verified by induction.

(ii) For every k , $h_k(\lambda)$ is concave and continuous with respect to $\lambda > 0$. This can be shown by noting that expectation and point-wise minimum operations conserve concavity.

(iii) For every k , $h_k(\lambda)$ is equal to $(1 - \pi)c_0$ for any λ no smaller than a certain threshold $\tau_k > 0$, and for any λ smaller than τ_k , $h_k(\lambda)$ is smaller than $(1 - \pi)c_0$. This is equivalent to the property that the curve $c\lambda + \mathbb{E}_0[h_{k+1}(\lambda L)]$ crosses the level $(1 - \pi)c_0$ only once, and thus can be shown by using (i), (ii), and the fact that $c\lambda + \mathbb{E}_0[h_{k+1}(\lambda L)] \rightarrow \infty$ as $\lambda \rightarrow \infty$, for any $c > 0$.

So from (iii) it is clear that the optimal stopping time is given by

$$T = \min \{1 \leq k \leq N - 1 | \Lambda_k \geq \tau_k\}, \quad (13)$$

where the thresholds τ_k are characterized by (12). Together with the optimal terminal decision rule D obtained in (8) we then prove Theorem 1. \square

In summary, the **Algorithm BO-ODR** shown at the top of the next column implements the Bayesian optimal ODR under maximum sample size N .

Algorithm 1 Algorithm BO-ODR: 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  $n \leq N$ 
  do compute  $\Lambda_n$ ;
  if  $\Lambda_n \geq \tau_n$ 
    terminate returning  $\mathcal{H}_1$ ;
  else  $n = n + 1$ ;
  end if
end while
terminate returning  $\mathcal{H}_0$ .
    
```




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

We illustrate the Bayesian optimal ODR by the case study of testing the hypotheses

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

with $A > 0$.

In the presented numerical examples (Figure 1), 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 by **Algorithm BO-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.

B. Case of Geometrically Distributed Maximum Sample Size

In this subsection, we turn to the case in which the maximum sample size is no longer fixed, but is a geometrically distributed random variable N with a fixed finite mean $1/\epsilon$, i.e., $\Pr[N = n] = (1 - \epsilon)^{n-1}\epsilon$, $n = 1, 2, \dots$. We assume that N is independent of X_1, X_2, \dots . 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 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 follows:

$$N = \min\{n : Z_n = 1\}, \quad (15)$$

where Z_n is an i.i.d. sequence of Bernoulli random variables with $\Pr[Z_1 = 1] = \epsilon$ and $\Pr[Z_1 = 0] = 1 - \epsilon$. Therefore, for any stopping time T' that is adapted to the filtration generated by X_1, X_2, \dots , if we define

$$T = \min\{T', N\}, \quad (16)$$

then T is a stopping time adapted to the product filtration generated by $(X_1, Z_1), (X_2, Z_2), \dots$. With a thus induced T and an arbitrary terminal decision rule D , we have the following definition of the ODR.

Definition 2: An ODR for a geometrically distributed maximum sample size N consists of a stopping time T given by (16) and a terminal decision rule $D : \mathcal{X}^N \mapsto \{\mathcal{H}_0, \mathcal{H}_1\}$, such that, the statistician declares \mathcal{H}_1 if either $\{T < N\}$ or $\{T = N, D(X^N) = \mathcal{H}_1\}$ occurs, and declares \mathcal{H}_0 if $\{T = N, D(X^N) = \mathcal{H}_0\}$ occurs. Similar to Definition 1, we may simply write $D(X^N)$ as D , when the context is unambiguous.

Analogous to the problem framework in Section II-A, we define the Bayesian cost as

$$\mathcal{J} = (1 - \pi)c_0 P_{FA} + \pi c_1 P_M + c \mathbb{E}_1[T], \quad (17)$$

where $c_0, c_1, c > 0$ are cost assignments, and the problem then 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.

For the Bayesian cost, we have the following key fact.

Proposition 1: The Bayesian cost \mathcal{J} in (17) can be written in the following form:

$$\begin{aligned} \mathcal{J} = \mathbb{E}_0 \left[\sum_{n=0}^{T-1} (1 - \epsilon)^n c \Lambda_n + (1 - \epsilon)^T \right. \\ \left. \times \left[(1 - \pi)c_0 + \frac{\epsilon}{1 - \epsilon} \min\{(1 - \pi)c_0, \pi c_1 \Lambda_T\} \right] \right]. \end{aligned} \quad (18)$$

Proof: To proceed, consider the conditional Bayesian cost conditioned upon $\{N = n\}$, \mathcal{J}_n ; that is,

$$\mathcal{J} = \sum_{n=1}^{\infty} \Pr[N = n] \mathcal{J}_n. \quad (19)$$

For evaluating the false alarm probability conditioned upon $\{N = n\}$, we note that this conditional event is just $\{T < n\} \cup \{T = n, D = \mathcal{H}_1\}$. So we may write the conditional false alarm probability as

$$P_{FA,n} = \mathbb{E}_0[\mathbf{1}(T < n)] + \mathbb{E}_0[\mathbf{1}(T = n)\mathbf{1}(D = \mathcal{H}_1)]. \quad (20)$$

Similarly, the conditional miss event is $\{T = n, D = \mathcal{H}_0\}$, and we may write the conditional miss probability as

$$\begin{aligned} P_{M,n} &= \mathbb{E}_1[\mathbf{1}(T = n)\mathbf{1}(D = \mathcal{H}_0)] \\ &= \mathbb{E}_0[\Lambda_n \mathbf{1}(T = n)\mathbf{1}(D = \mathcal{H}_0)], \end{aligned} \quad (21)$$

since $\{T = n, D = \mathcal{H}_0\}$ is \mathcal{F}_n -measurable. So we have

$$\begin{aligned} \mathcal{J}_n &= \mathbb{E}_0 [(1 - \pi)c_0 \mathbf{1}(T < n) \\ &+ (1 - \pi)c_0 \mathbf{1}(T = n)\mathbf{1}(D = \mathcal{H}_1) \\ &+ \pi c_1 \Lambda_n \mathbf{1}(T = n)\mathbf{1}(D = \mathcal{H}_0)] + c \mathbb{E}_1[T|N = n]. \end{aligned} \quad (22)$$

Clearly for a given T , for each n , the optimal D that minimizes \mathcal{J}_n is given by

$$\begin{aligned} D = \mathcal{H}_1 \quad \text{if } \Lambda_n \geq \frac{(1 - \pi)c_0}{\pi c_1}; \\ \text{and } D = \mathcal{H}_0 \quad \text{otherwise.} \end{aligned} \quad (23)$$

Since this solution does not depend on n , it is also the unconditional optimal terminal decision rule; that is,

$$D = \mathcal{H}_1 \quad \text{if } \Lambda_N \geq \frac{(1 - \pi)c_0}{\pi c_1}; \quad \text{and } D = \mathcal{H}_0 \quad \text{otherwise.} \quad (24)$$

Now we can rewrite the Bayesian cost given the optimal terminal decision rule as

$$\begin{aligned} \mathcal{J} &= \sum_{n=1}^{\infty} \Pr[N = n] \mathbb{E}_0 [(1 - \pi)c_0 \mathbf{1}(T < n) \\ &+ \mathbf{1}(T = n) \min\{(1 - \pi)c_0, \pi c_1 \Lambda_n\}] + c \mathbb{E}_1[T]. \end{aligned} \quad (25)$$

Noting that

$$\begin{aligned} &\sum_{n=1}^{\infty} \Pr[N = n] \mathbb{E}_0 [(1 - \pi)c_0 \mathbf{1}(T < n)] \\ &= (1 - \pi)c_0 \mathbb{E}_0 \left[\sum_{n=1}^{\infty} \Pr[N = n] \mathbf{1}(T < n) \right] \\ &= (1 - \pi)c_0 \mathbb{E}_0 \left[\sum_{n=T+1}^{\infty} \Pr[N = n] \right] \\ &= (1 - \pi)c_0 \mathbb{E}_0 \left[\sum_{n=T+1}^{\infty} (1 - \epsilon)^{n-1} \epsilon \right] \\ &= (1 - \pi)c_0 \mathbb{E}_0 [(1 - \epsilon)^T], \end{aligned} \quad (26)$$

we have

$$\begin{aligned} \mathcal{J} &= \mathbb{E}_0 \left[(1 - \pi)c_0 (1 - \epsilon)^T \right. \\ &\left. + (1 - \epsilon)^{T-1} \epsilon \min\{(1 - \pi)c_0, \pi c_1 \Lambda_T\} \right] + c \mathbb{E}_1[T]. \end{aligned} \quad (27)$$

Next we evaluate $\mathbb{E}_1[T]$, as follows:

$$\begin{aligned}
\mathbb{E}_1[T] &\stackrel{(a)}{=} \mathbb{E}_1[\min\{T, N\}] \\
&\stackrel{(b)}{=} \sum_{n=1}^{\infty} (1-\epsilon)^{n-1} \epsilon \mathbb{E}_1[\min\{T, n\}] \\
&\stackrel{(c)}{=} \mathbb{E}_1 \left[\sum_{n=1}^T n(1-\epsilon)^{n-1} \epsilon + \sum_{n=T+1}^{\infty} T(1-\epsilon)^{n-1} \epsilon \right] \\
&\stackrel{(d)}{=} \mathbb{E}_1 \left[\frac{1 - (1-\epsilon)^T}{\epsilon} \right] = \mathbb{E}_1 \left[\sum_{n=0}^{T-1} (1-\epsilon)^n \right] \\
&\stackrel{(e)}{=} \mathbb{E}_0 \left[\sum_{n=0}^{T-1} (1-\epsilon)^n \Lambda_n \right], \tag{28}
\end{aligned}$$

where, (a) is due to the fact that T is upper bounded by N , (b) is the total expectation expansion, (c) is due to the fact that the expectation of T is bounded (by the expectation of N), (d) is obtained via algebraic manipulations, and (e) is due to the fact that the event $\{n \leq T-1\}$ is \mathcal{F}_n -measurable.

So, back to the Bayesian cost, we have reached (18) and thus proved Proposition 1. \square

An inspection of (18) reveals that it is exactly the form that has been treated in [32, 2.14], considering both an instantaneous reward at the stopping time and accumulated sampling costs, with everything discounted by an exponential factor $(1-\epsilon)^n$ at time n .

Let us define for $\lambda \geq 0$

$$g(\lambda) = (1-\pi)c_0 + \frac{\epsilon}{1-\epsilon} \min\{(1-\pi)c_0, \pi c_1 \lambda\}, \tag{29}$$

and $c(\lambda) = c\lambda$. First, the relevant regularity conditions [32, eq. (2.168)] hold, namely that $|g(\lambda)|$ is finitely bounded, and that $\mathbb{E}_0[c(\Lambda_n)]$ is finite for every n . So, as a consequence of [32, Th. 23], we have the following result.

Theorem 2: The Bayesian optimal stopping time is given by

$$T = \min\{n \geq 1 : V(\Lambda_n) = g(\Lambda_n)\}, \tag{30}$$

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

$$V(\lambda) = \min\{g(\lambda), (1-\epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)\}, \tag{31}$$

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

$$Qf(\lambda) = \min\{f(\lambda), (1-\epsilon)\mathbb{E}_0[f(\lambda L)] + c(\lambda)\}. \tag{32}$$

The Bayesian optimal stopping time (30) leads to a likelihood ratio threshold test, as given by the following result.

Corollary 1: Define the ‘‘running’’ threshold τ_r as the value of λ at the intersection of $g(\lambda)$ and $(1-\epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)$, which always exists and is unique, and the ‘‘terminal’’ threshold $\tau_t = (1-\pi)c_0/(\pi c_1)$. The Bayesian optimal stopping rule is described by **Algorithm BO-ODR-Geo**.

Proof: It suffices to prove that τ_r always exists and is unique; that is, $g(\lambda)$ and $(1-\epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)$ intersect only once. By induction, it follows that $V(\lambda)$ is a monotonically non-decreasing and concave continuous function of $\lambda > 0$, and

Algorithm 2 Algorithm BO-ODR-Geo: 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: The ‘‘running’’ threshold τ_r and the ‘‘terminal’’ threshold τ_t , according to Corollary 1.

Algorithm:

```

initialize  $n = 1$ ;
while  $N$  has not been revealed
  do compute  $\Lambda_n$ ;
  if  $\Lambda_n \geq \tau_r$ 
    terminate returning  $\mathcal{H}_1$ ;
  else  $n = n + 1$ ;
  end if
end while
if  $\Lambda_N \geq \tau_t$ 
  terminate returning  $\mathcal{H}_1$ ;
else
  terminate returning  $\mathcal{H}_0$ ;
end if

```

that $\lim_{\lambda \rightarrow 0^+} V(\lambda) = 0$. Therefore, $(1-\epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)$ grows without bound as $\lambda \rightarrow \infty$, and hence it must intersect at least once at $g(\lambda)$ over $\lambda > 0$. To prove that the intersection is unique, we note that from (31)

$$\begin{aligned}
V(\lambda) &= \min\{g(\lambda), (1-\epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)\} \\
&\leq (1-\epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda) \\
&\stackrel{(a)}{\leq} (1-\epsilon)V(\lambda \mathbb{E}_0[L]) + c(\lambda) \\
&\stackrel{(b)}{=} (1-\epsilon)V(\lambda) + c(\lambda), \tag{33}
\end{aligned}$$

that is,

$$V(\lambda) \stackrel{(c)}{\leq} \frac{c\lambda}{\epsilon}, \tag{34}$$

wherein, (a) is from the concavity of $V(\cdot)$, (b) is from the fact that $\mathbb{E}_0[L] = 1$, and (c) is from the fact that $c(\lambda) = c\lambda$.

We then consider two cases.

Case 1: $\epsilon > c/(c + \pi c_1)$. In this case, it is impossible for $(1-\epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)$ to intersect $g(\lambda)$ for any $\lambda < \tau_t = (1-\pi)c_0/(\pi c_1)$, or, it is only possible for $(1-\epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)$ to intersect $g(\lambda)$ at some $\lambda \geq \tau_t$, for which $g(\lambda) = (1-\pi)c_0/(1-\epsilon)$ is a horizontal line, and thus the intersection is unique. To see this, note that in the case of $\epsilon > c/(c + \pi c_1)$, we have

$$V(\tau_t) \leq \frac{c\tau_t}{\epsilon} = \frac{c(1-\pi)c_0}{\epsilon\pi c_1} < \frac{(1-\pi)c_0(c + \pi c_1)}{\pi c_1}, \tag{35}$$

and

$$g(\tau_t) = \frac{(1-\pi)c_0}{1-\epsilon} > \frac{(1-\pi)c_0(c + \pi c_1)}{\pi c_1}. \tag{36}$$

Case 2: $\epsilon \leq c/(c + \pi c_1)$. In this case, it is possible for $(1-\epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)$ to intersect $g(\lambda)$ for some $\lambda_1 < \tau_t$. But if this happens, then it is impossible for these two curves to intersect for any other $\lambda > \lambda_1$ and thus the intersection is also unique. To see this, note that the slope




Fig. 2. Illustration of typical situation for τ_r, τ_t , with $c_0 = 10, c_1 = 20$.




Fig. 3. Illustration of typical situation for τ_r, τ_t , with $c_0 = 20, c_1 = 4$.

of $(1 - \epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)$ is always lower bounded by c , while the slope of $g(\lambda)$ for $\lambda < \tau_t$ is $\epsilon\pi c_1/(1 - \epsilon)$, which is no greater than c in the case of $\epsilon \leq c/(c + \pi c_1)$.

Summarizing Cases 1 and 2, we conclude the proof of Corollary 1. \square

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 (31), 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 II-A to illustrate the numerical behavior of the optimal ODR under geometrically distributed 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 $(1 - \epsilon)\mathbb{E}_0[V(\lambda L)] + c(\lambda)$ intersects $g(\lambda)$, there are two possible situations, as illustrated in Figures 2 and 3, respectively. In the former, $\tau_r \geq \tau_t$, and in the latter, $\tau_r < \tau_t$. 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 .

III. ASYMPTOTIC ANALYSIS

In the previous section, we have focused on the Bayesian optimal stopping rules, which yield ODRs that minimize corresponding Bayesian costs. In this section, we turn to the asymptotic regime, letting the maximum sample size N grow without bound. The performance metrics here are, instead of the Bayesian cost, the exponential decay rates of the (false alarm and miss) error probabilities and the expected stopping time under \mathcal{H}_1 .




Fig. 4. The running threshold τ_r versus c_0 and c_1 (herein $c_0 = c_1$). The dash-dot line indicates the terminal threshold τ_t .

In general, for a sequence of ODRs indexed by the maximum sample size $N = 1, 2, \dots$, 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 $(\Delta_{\text{FA}}, \Delta_{\text{M}}, \eta)$ is achievable if there exists a sequence of ODRs indexed by N , such that

$$\liminf_{N \rightarrow \infty} \frac{-\log P_{\text{FA}}}{N} \geq \Delta_{\text{FA}}, \quad (37)$$

$$\liminf_{N \rightarrow \infty} \frac{-\log P_{\text{M}}}{N} \geq \Delta_{\text{M}}, \quad (38)$$

$$\limsup_{N \rightarrow \infty} \frac{T}{N} \leq \eta, \quad (39)$$

where $T = \mathbb{E}_1[T]$ 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 theorem is the main result of our asymptotic analysis, which fully characterizes $\mathcal{R}(p_0, p_1)$.

Theorem 3: The ODR performance region $\mathcal{R}(p_0, p_1)$ is given as follows: for each $0 \leq \eta \leq 1$,

$$\begin{aligned} \Delta_{\text{FA}} &\leq \min \left\{ \eta d_1, \sup_{\alpha > 0} \left\{ \alpha [d_1 - \nu(d_0 + d_1)] \right. \right. \\ &\quad \left. \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \right\}, \\ \Delta_{\text{M}} &\leq \sup_{\alpha < 0} \left\{ \alpha [d_1 - \nu(d_0 + d_1)] \right. \\ &\quad \left. - \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\}, \end{aligned} \quad (40)$$

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

Theorem 3 is proved in two parts. The achievability part is established by constructing a specific form of ODRs that asymptotically achieve the performance tuple as described in Theorem 3. 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 [30] for a certain stationary memoryless channel, an impossible task even with feedback and variable-length coding [31]. The detailed steps of the proof are given in Sections III-C through III-E. In the next two subsections we provide some illustration and discussion of Theorem 3.

A. Case Study: Gaussian Distributions With and Without a Drift

To illustrate the ODR performance region in Theorem 3, we present a case study for the following hypotheses:

$$\mathcal{H}_0 : p_0 \sim \mathcal{N}(0, 1) \quad \text{versus} \quad \mathcal{H}_1 : p_1 \sim \mathcal{N}(A, 1), \quad (41)$$

with $A > 0$. In this case we have $D(p_0 \| p_1) = D(p_1 \| p_0) = A^2/2$.

Then, applying Theorem 3, we can obtain the (normalized) region $\left(\frac{\Delta_{\text{FA}}}{A^2/2}, \frac{\Delta_{\text{M}}}{A^2/2} \right)$ for every fixed $0 \leq \eta \leq 1$, as

$$\left\{ \left(\frac{\Delta_{\text{FA}}}{A^2/2}, \frac{\Delta_{\text{M}}}{A^2/2} \right) = (x, y) : \right. \\ \left. \sqrt{x} + \sqrt{y} \leq 1, 0 \leq x \leq \eta, y \geq 0 \right\}, \quad (42)$$

illustrated in Figure 5.

The complete characterization of $\mathcal{R}(p_0, p_1)$ is given by the following corollary and illustrated in Figure 6.

Corollary 2: For the hypotheses (41), the ODR performance region $\mathcal{R}(p_0, p_1)$ is

$$\left\{ \left(\frac{\Delta_{\text{FA}}}{A^2/2}, \frac{\Delta_{\text{M}}}{A^2/2}, \eta \right) = (x, y, z) : \right. \\ \left. \sqrt{x} + \sqrt{y} \leq 1, 0 \leq x \leq z, y \geq 0, 0 \leq z \leq 1 \right\}. \quad (43)$$

B. Stein-Chernoff Lemma Revisited

In this subsection, we focus on an extremal case of Theorem 3, in which the false alarm probability is fixed without decreasing toward zero exponentially, or, has an exponent




Fig. 5. An illustration of (42).




Fig. 6. An illustration of $\mathcal{R}(p_0, p_1)$ in (43).

of zero. For this case, Theorem 3 specializes into the following corollary.

Corollary 3: For an arbitrary fixed target false alarm probability $P_{\text{FA}}^* > 0$, among all sequences of ODRs such that the normalized expected stopping time under \mathcal{H}_1 , T/N , satisfies $\lim_{N \rightarrow \infty} T/N = 0$, the maximum achievable error exponent of P_{M} is

$$\lim_{N \rightarrow \infty} \frac{-\log P_{\text{M}}}{N} = D(p_0 \| p_1). \quad (44)$$

A similar situation has been treated in [25], wherein the considered form of ODR is restricted to be a truncated one-sided SPRT, that is,

$$\hat{\mathcal{H}} = \mathbf{1} \left(\bigcup_{k=1}^N \{\Lambda_k \geq B\} \right), \quad (45)$$

where $\mathbf{1}(\cdot)$ is the indicator function. It was shown therein that the above form of ODR behaves asymptotically according to the following theorem.

Theorem 4 ([25, Ths. 1 and 2]): For the truncated one-sided SPRT ODR of the form (45) that attains an arbitrary fixed target false alarm probability $0 < P_{\text{FA}}^* \leq P_0[p_1(X) \geq p_0(X)]$, the miss probability scales toward zero as N grows without bound as

$$\lim_{N \rightarrow \infty} \frac{-\log P_{\text{M}}}{N} = C(p_0, p_1), \quad (46)$$

where $C(p_0, p_1)$ is the Chernoff information of (p_0, p_1) (see [3] and [33, Ch. 11.9])

$$C(p_0, p_1) = - \inf_{\alpha \in (0,1)} \log \left(\int_{\mathcal{X}} p_0^\alpha(x) p_1^{1-\alpha}(x) dx \right), \quad (47)$$

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

$$\lim_{N \rightarrow \infty} \frac{T}{N} = 0. \quad (48)$$

Comparing Theorem 4 and Corollary 3, we can conclude that, among all sequences of ODRs such that $\lim_{N \rightarrow \infty} T/N = 0$, there exist ODRs that achieve a larger error exponent of P_M than that achieved by the truncated one-sided SPRT ODR in [25].⁴ The error exponent achieved in Corollary 3, $D(p_0 \| p_1)$, 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, only 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.

C. Proof of Corollary 3

Let us represent an ODR in a general form beyond that in (45) as

$$\hat{\mathcal{H}}(\underline{f}, \underline{B}) = \mathbf{1} \left(\bigcup_{k=1}^N \{f_k(X_1, \dots, X_k) \geq B_k\} \right), \quad (49)$$

where \underline{f} is a sequence of processing functions, and \underline{B} is a sequence of thresholds. Note that $\hat{\mathcal{H}}(\underline{f}, \underline{B})$ also includes as a special case the FSS Neyman-Pearson decision rules, which have \underline{f} as likelihood ratios, and \underline{B} as $B_k = \infty$ for all k except for $k = N$.

There are several ways of constructing an ODR that achieves the asymptotic performance in Corollary 3. Here we give a proof based on a simple idea of two-stage ODRs, which may not be the most sensible choice for finite N in practice but is sufficient for proving the asymptotic result and is quite convenient to analyze. That is, we restrict the sequence of processing functions, \underline{f} , to be likelihood ratios, and among the elements of the threshold sequence \underline{B} , we only let two of them be finite, i.e.,

$$B_M = e^{\tau_M}, \quad B_N = e^{\tau_N}, \quad \text{and } B_k = \infty \text{ for } k \neq M, N, \quad (50)$$

where $M < N$ corresponds to an early stage at which an opportunistic decision may be made, and $e^{\tau_M, N}$ are the thresholds for the two stages. We identify such two-stage ODRs with the designation 2-ODR, and sometimes represent them with the notation $\hat{\mathcal{H}}_2(M, \tau_M, \tau_N)$. So in words, for 2-ODRs there is only one opportunity (upon observing the first M samples) to stop early.

⁴It should be noted that this comparison is based on the footing of $\lim_{N \rightarrow \infty} T/N = 0$. It is possible that for finer asymptotic behaviors, the comparison becomes more delicate; for example, one may ask whether the conclusion still holds if one focuses on sequences of ODRs such that $T = O(N^\alpha)$ for some $0 < \alpha < 1$, and the result is unknown.

In the proof of Corollary 3, we fix $M/N = \epsilon > 0$, and let $\tau_M = -M [D(p_0 \| p_1) - \delta]$ and $\tau_N = -N [D(p_0 \| p_1) - \delta]$ for some small $\delta > 0$. The miss probability is thus bounded as

$$\begin{aligned} P_M &= P_1 [\Lambda_M < e^{\tau_M}, \Lambda_N < e^{\tau_N}] \\ &\leq P_1 [\log \Lambda_N < \tau_N] \\ &\leq \exp \left\{ \inf_{0 \leq \alpha \leq 1} \left[-\alpha N (D(p_0 \| p_1) - \delta) \right. \right. \\ &\quad \left. \left. + \log \mathbb{E}_1 \left[e^{-\alpha \log \Lambda_N} \right] \right] \right\} \\ &= \exp \left\{ -N \sup_{0 \leq \alpha \leq 1} \left[\alpha D(p_0 \| p_1) - \alpha \delta \right. \right. \\ &\quad \left. \left. - \log \mathbb{E}_1 \left[e^{-\alpha \log p_1(X)/p_0(X)} \right] \right] \right\} \\ &\leq \exp \{-N [D(p_0 \| p_1) - \delta]\}, \end{aligned} \quad (51)$$

by letting $\alpha = 1$.

The false alarm probability is bounded as

$$\begin{aligned} P_{FA} &= P_0 [\Lambda_M \geq e^{\tau_M} \text{ or } \Lambda_N \geq e^{\tau_N}] \\ &\leq P_0 [\log \Lambda_M \geq \tau_M] + P_0 [\log \Lambda_N \geq \tau_N] \end{aligned} \quad (52)$$

due to the union bound. From the weak law of large numbers, both probabilities in (52) approach zero for any fixed $\delta > 0$, as N grows without bound. Thus P_{FA} can be ensured to be arbitrarily small as N grows without bound.

Regarding the expected stopping time under \mathcal{H}_1 , we have

$$\begin{aligned} T &= M \cdot P_1 [\Lambda_M \geq e^{\tau_M}] + N \cdot P_1 [\Lambda_M < e^{\tau_M}] \\ &\leq M + N \cdot \exp \{-M [D(p_0 \| p_1) - \delta]\}. \end{aligned} \quad (53)$$

So it follows that

$$\frac{T}{N} \leq \frac{M}{N} + \exp \{-M [D(p_0 \| p_1) - \delta]\}, \quad (54)$$

which converges as $N \rightarrow \infty$ to $\epsilon = M/N$. So the proof of Corollary 3 is completed by letting $\delta \rightarrow 0$ and $\epsilon \rightarrow 0$. \square

D. The $(\Delta_M = 0, \Delta_{FA} = \eta D(p_1 \| p_0))$ Corner Point

According to Theorem 3, a boundary of $\mathcal{R}(p_0, p_1)$ is given by $\Delta_{FA} = \eta D(p_1 \| p_0)$ for every $0 \leq \eta \leq 1$. This extremal case corner point of $(\Delta_M = 0, \Delta_{FA} = \eta D(p_1 \| p_0))$ specializes into the following corollary.

Corollary 4: For an arbitrary fixed target miss probability $P_M^* > 0$, there exists a sequence of ODRs such that when the normalized expected stopping time under \mathcal{H}_1 , T/N , satisfies

$$\lim_{N \rightarrow \infty} \frac{T}{N} = \eta \leq 1, \quad (55)$$

the false alarm probability satisfies

$$\lim_{N \rightarrow \infty} \frac{-\log P_{FA}}{N} = \eta D(p_1 \| p_0). \quad (56)$$

Furthermore, no ODR may achieve a larger exponent for P_{FA} under the constraint of (55) on T .

Proof of the Achievability Part: To prove the existence of ODRs that achieve (55) and (56), consider 2-ODRs $\{\hat{\mathcal{J}}_2(M, \tau_M, \tau_N)\}$, in which we set

$$\begin{aligned} M &= \eta N, \quad \tau_M = M [D(p_1 \| p_0) - \delta], \\ \tau_N &= N [D(p_1 \| p_0) - \delta], \end{aligned} \quad (57)$$

for some small $\delta > 0$. The miss probability P_M satisfies

$$\begin{aligned} P_M &= P_1 [\log \Lambda_M < \tau_M, \log \Lambda_N < \tau_N] \\ &\leq P_1 [\log \Lambda_N < \tau_N], \end{aligned} \quad (58)$$

which can be ensured to be arbitrarily small as N grows without bound, for any fixed $\delta > 0$. For T , we have

$$\begin{aligned} \frac{T}{N} &= \frac{M}{N} \cdot P_1 [\log \Lambda_M \geq \tau_M] + P_1 [\log \Lambda_M < \tau_M] \\ &\leq \eta + P_1 [\log \Lambda_M < \tau_M], \end{aligned} \quad (59)$$

which converges to η as N grows without bound, for any fixed $\delta > 0$.

For the false alarm probability P_{FA} , we have

$$\begin{aligned} P_{FA} &= P_0 [\log \Lambda_M \geq \tau_M \text{ or } \log \Lambda_N \geq \tau_N] \\ &\leq P_0 [\log \Lambda_M \geq \tau_M] + P_0 [\log \Lambda_N \geq \tau_N] \\ &\leq \exp \left\{ -M \sup_{0 \leq \alpha \leq 1} \left[\alpha D(p_1 \| p_0) - \alpha \delta \right. \right. \\ &\quad \left. \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right] \right\} \\ &\quad + \exp \left\{ -N \sup_{0 \leq \alpha \leq 1} \left[\alpha D(p_1 \| p_0) - \alpha \delta \right. \right. \\ &\quad \left. \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right] \right\} \\ &\leq \exp \{-\eta N [D(p_1 \| p_0) - \delta]\} \\ &\quad + \exp \{-N [D(p_1 \| p_0) - \delta]\}, \end{aligned} \quad (60)$$

by setting $\alpha = 1$ in both exponents. Herein, the first term dominates the exponential decay behavior as N grows without bound. So the achievability part of Corollary 4 is established by letting $\delta \rightarrow 0$.

Proof of the Converse Part: To prove that there are no ODRs that outperform the asymptotic performance specified in (55) and (56), we use the argument of contradiction, which borrows ideas from the information-theoretic analysis of the channel capacity per unit cost [30].

Here we briefly describe the channel capacity per unit cost problem, simplified for our problem setup. Consider a stationary memoryless channel. Let the channel input alphabet \mathcal{S} consist of two letters, s_0 and s_1 , whose corresponding conditional output distributions are $p_0(x)$ and $p_1(x)$, respectively. We assume that the cost of using s_0 as the channel input is zero, and that of using s_1 is one. As established in [30, Th. 3], the channel capacity per unit cost of this channel is given by $\mathbf{C} = D(p_1 \| p_0)$. Furthermore, from [31], \mathbf{C} remains the channel capacity per unit cost even in the presence of feedback and variable-length coding.

Then, for the channel model above, consider the following encoding/decoding scheme. Denote the size of the message set by M , in which a message is selected uniformly at random for

transmission, and introduce a parameter N . A ‘‘root’’ codebook is constructed as a collection of M different $M \times N$ matrices, wherein the m th message corresponds to an $M \times N$ codeword matrix whose m th-row elements are all s_1 , and the remaining elements in the matrix are all s_0 .

When no feedback is available, the root codebook is the actual codebook used for transmission [30, pp. 1023–1024]. That is, once a message m is selected, the encoder transmits the corresponding matrix, row by row. The decoder conducts a binary hypothesis test for each received row; that is, for the i th row, the decoder decides either $r_i = 0$ (i.e., s_0 has been sent through that row) or $r_i = 1$ (i.e., s_1 has been sent through that row). Assuming that message m is sent, a decoding error occurs if either $r_m = 0$, or for any of $i \neq m$, $r_i = 1$. According to the Stein-Chernoff Lemma, for any fixed probability of $\Pr[r_m = 0 | m \text{ sent}] = \epsilon$ and any fixed $\delta > 0$, we can achieve

$$\Pr[r_i = 1 | m \text{ sent}] \leq \exp \{-N [D(p_1 \| p_0) - \delta]\}, \quad (61)$$

as N grows sufficiently large, for each $i \neq m$. Hence from the union bound, the decoding error probability with message m sent is upper bounded by

$$\begin{aligned} \Pr[\text{some } m' \neq m \text{ declared} | m \text{ sent}] \\ \leq \epsilon + M \cdot \exp \{-N [D(p_1 \| p_0) - \delta]\}. \end{aligned} \quad (62)$$

So by choosing the coding rate appropriately as long as $(\log M)/N < D(p_1 \| p_1) - \delta$, and then by letting ϵ and δ approach zero, the decoding error probability can be made arbitrarily close to zero. Noting that the total cost of transmitting the $M \times N$ matrix codeword is N , the achieved rate per unit cost thus can be made arbitrarily close to $\mathbf{C} = D(p_1 \| p_1)$.

When feedback is available, while receiving each row of the codeword matrix, the decoder can operate an ODR, so as to permit early termination with p_1 declared. Hence we can have an adaptive transmission scheme as follows (see Figure 7 for an illustration),

- 1) Set $i = 1$.
- 2) The encoder transmits the elements of the i th row of the corresponding codeword matrix in the root codebook, one by one; meanwhile the decoder performs an ODR with maximum sample size N .
- 3) The decoder informs the encoder through the feedback link of its decision immediately when the decision is made using the ODR.
- 4) The encoder stops transmitting its current row once receiving the decision from the decoder, increases i by one (unless $i = M$ already), and goes to Step 2).
- 5) If $i = M$, then the encoder halts; the decoder declares the decoded message to be \hat{m} if there is only one row index \hat{m} whose ODR detects p_1 , and for all other cases (no such \hat{m} exists or more than one such \hat{m} exist) the decoder arbitrarily makes a declaration of the decoded message.

We note that the above adaptive transmission scheme is feasible due to the availability of feedback, because then the encoder and the decoder can maintain synchronization regarding which row is being sent across the channel, even under ODR with variable stopping times.

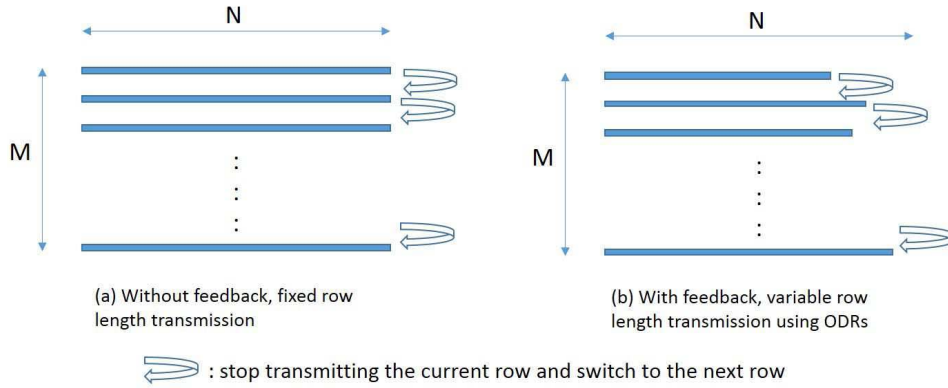


Fig. 7. Illustration of the transmission schemes without and with feedback.

Now let us suppose that there exist ODRs that outperform the asymptotic performance in Corollary 4; that is, at least for certain $0 < \eta < 1$, for an arbitrarily small but fixed target miss probability $P_M^* > 0$, there exist ODRs such that when the normalized expected stopping time under \mathcal{H}_1 , T/N , satisfies

$$\lim_{N \rightarrow \infty} \frac{T}{N} = \eta \leq 1, \quad (63)$$

and the false alarm probability satisfies

$$\liminf_{N \rightarrow \infty} \frac{-\log P_{FA}}{N} > \eta D(p_1 \| p_0). \quad (64)$$

Let the decoder in the adaptive transmission scheme use the ODRs that we have supposed to exist satisfying (63) and (64). For such η , we have that the expected cost of sending a codeword is $\eta N + o(N)$ with $o(N)/N \rightarrow 0$ as $N \rightarrow \infty$, noting that transmitting all the rows, other than the one corresponding to the message index, incurs zero cost. On setting the probability $\Pr[r_m = 0 | m \text{ sent}] = \epsilon > 0$ arbitrarily, according to (64) we have

$$\Pr[r_i = 1 | m \text{ sent}] \leq \exp\{-N\Delta\}, \quad (65)$$

for some $\Delta > \eta D(p_1 \| p_0)$, as N grows without bound, for each $i \neq m$. Hence from the same union bounding technique as that used in (62), the size of the message set can be made arbitrarily close to $\log M = N\Delta$, for achieving arbitrarily small decoding error probability as $N \rightarrow \infty$. Consequently, The achieved rate per unit cost is up to $N\Delta/(\eta N + o(N)) > N\eta D(p_1 \| p_0)/(\eta N + o(N)) \rightarrow D(p_1 \| p_0) = \mathbf{C}$, as $N \rightarrow \infty$. Therefore we encounter a contradiction since \mathbf{C} cannot be outperformed by any coding scheme even in the presence of feedback and variable-length coding [31], and hence the supposed ODRs cannot exist. This establishes the converse part of Corollary 4. \square

E. Completing the Proof of Theorem 3

Having established the two extremal cases in Corollaries 3 and 4, in this subsection we complete the proof of Theorem 3.

Proof of the Converse Part: We prove that no ODRs may outperform the performance region in Theorem 3. First, note that for any $0 \leq \eta \leq 1$, the pair (Δ_{FA}, Δ_M) has to be

bounded by

$$\begin{aligned} \Delta_{FA} &\leq \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\ &\quad \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\}, \\ \Delta_M &\leq \sup_{\alpha < 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\ &\quad \left. - \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\}, \end{aligned} \quad (66)$$

for $0 \leq \nu \leq 1$, because that is the performance boundary achieved by the FSS likelihood ratio threshold test rules, under sample size N .

Second, according to Corollary 4, for any $0 \leq \eta \leq 1$, Δ_{FA} has to satisfy

$$\Delta_{FA} \leq \eta D(p_1 \| p_0). \quad (67)$$

The converse part of Theorem 3 thus follows from combining (66) and (67).

Proof of the Achievability Part: We prove that ODRs attaining the performance region in Theorem 3 exist. For this, using 2-ODRs $\{\hat{\mathcal{H}}_2(M, \tau_M, \tau_N)\}$ with parameters

$$\begin{aligned} M &= \eta N \\ \frac{\tau_M}{M} &= D(p_1 \| p_0) - \mu [D(p_0 \| p_1) + D(p_1 \| p_0)] \\ \frac{\tau_N}{N} &= D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)], \end{aligned} \quad (68)$$

for $0 < \mu, \nu < 1$, we have the following achievable performance tuple:

$$\begin{aligned} \Delta_{FA} &= \min \left\{ \sup_{\alpha > 0} \left\{ \alpha \eta [D(p_1 \| p_0) - \mu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \right. \\ &\quad \left. \left. - \eta \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\}, \right. \\ &\quad \left. \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \right. \\ &\quad \left. \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \right\} \\ \Delta_M &= \max \left\{ \sup_{\alpha < 0} \left\{ \alpha \eta [D(p_1 \| p_0) - \mu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \right. \end{aligned}$$

$$\begin{aligned}
& -\eta \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right], \\
& \sup_{\alpha < 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\}, \quad (69)
\end{aligned}$$

and $\lim_{N \rightarrow \infty} T/N = \eta$.

We need to prove that the above region (69) contains the region described in Theorem 3. For any fixed $0 \leq \eta \leq 1$, denote the value of ν that solves

$$\begin{aligned}
& \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} = \eta D(p_1 \| p_0) \quad (70)
\end{aligned}$$

by ν^* . What needs to be proved then is that, for any $\nu \geq \nu^*$, there exists a μ such that

$$\begin{aligned}
& \eta \cdot \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - \mu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \\
& \geq \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\}, \text{ and} \\
& \eta \cdot \sup_{\alpha < 0} \left\{ \alpha [D(p_1 \| p_0) - \mu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \\
& \leq \sup_{\alpha < 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\}. \quad (71)
\end{aligned}$$

This is because, if the two inequalities in (71) hold, then we have that the tuple

$$\begin{aligned}
\Delta_{\text{FA}} &= \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \\
\Delta_{\text{M}} &= \sup_{\alpha < 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \quad (72)
\end{aligned}$$

is achievable, for any $\nu^* \leq \nu < 1$, and thus the region (69) contains the region described in Theorem 3. But clearly letting μ be sufficiently close to zero suffices to satisfy (71), because

$$\begin{aligned}
& \eta \cdot \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - 0 \cdot [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} = \eta D(p_1 \| p_0) \\
& = \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - \nu^* \cdot [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right\}
\end{aligned}$$

$$\begin{aligned}
& - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \Big\} \\
& \geq \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - \nu \cdot [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \quad (73)
\end{aligned}$$

for any $\nu \geq \nu^*$, and

$$\begin{aligned}
& \eta \cdot \sup_{\alpha < 0} \left\{ \alpha [D(p_1 \| p_0) - 0 \cdot [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \\
& = 0 \leq \sup_{\alpha < 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\}. \quad (74)
\end{aligned}$$

This thus completes the achievability part of Theorem 3.

Discussion: A pivotal operating point for FSS decision rules is that when

$$\begin{aligned}
& \sup_{\alpha > 0} \left\{ \alpha [D(p_1 \| p_0) - \nu \cdot [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_0 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \\
& = \sup_{\alpha < 0} \left\{ \alpha [D(p_1 \| p_0) - \nu [D(p_0 \| p_1) + D(p_1 \| p_0)]] \right. \\
& \left. - \log \mathbb{E}_1 \left[e^{\alpha \log p_1(X)/p_0(X)} \right] \right\} \quad (75)
\end{aligned}$$

holds, and then their common value is exactly the Chernoff information of (p_0, p_1) , $C(p_0, p_1)$, and can be equivalently expressed as

$$C(p_0, p_1) = - \inf_{\alpha \in (0,1)} \log \left(\int_{\mathcal{X}} p_0^\alpha(x) p_1^{1-\alpha}(x) dx \right), \quad (76)$$

an expression which has been used in Theorem 4. From the above proof of Theorem 3, an immediate consequence is that ODRs may achieve the operating point of $\Delta_{\text{FA}} = \Delta_{\text{M}} = C(p_0, p_1)$ if and only if $\eta \geq C(p_0, p_1)/D(p_1 \| p_0)$.

IV. CONCLUSION

In this paper, we have formulated the general ODR framework and treated several of its key characteristics. We have considered both finite and asymptotic problems. In the finite regime, we have established Bayesian optimal ODRs for the case of a fixed maximum sample size, and the case of a geometrically distributed maximum sample size. For the latter, the Bayesian optimal ODR is a likelihood ratio threshold test with two thresholds. In the asymptotic regime, as the maximum sample size grows without bound, we have completely characterized the tradeoff among the exponents of the (false alarm and miss) error probabilities and the normalized expected stopping time under the alternative hypothesis.

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 cost; see, e.g., [26, Sec. 13]. For our setup, in order to make the problem meaningful, we need to further tune the growth of the maximum sample size N (or the mean maximum sample size $1/\epsilon$ in the case of geometrically distributed maximum sample size) accordingly, say, following $O(1/c)$, and the interplay between the sampling cost and the maximum sample size may exhibit interesting behavior.

REFERENCES

- [1] H. Chernoff, "Large-sample theory: Parametric case," *Ann. Math. Statist.*, vol. 27, no. 1, pp. 1–22, 1956.
- [2] C. Stein, "Information and comparison of experiments," unpublished.
- [3] 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.
- [4] W. Hoeffding, "Asymptotically optimal tests for multinomial distributions," *Ann. Math. Statist.*, vol. 36, no. 2, pp. 369–400, 1965.
- [5] I. Csiszár and G. Longo, "On the error exponent for source coding and testing simple hypotheses," *Stud. Sci. Math. Hungarica*, vol. 6, pp. 181–191, 1971.
- [6] R. E. Blahut, "Hypothesis testing and information theory," *IEEE Trans. Inf. Theory*, vol. 20, no. 4, pp. 405–417, Jul. 1974.
- [7] T. S. Han and K. Kobayashi, "The strong converse theorem for hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 35, no. 1, pp. 178–180, Jan. 1989.
- [8] T. S. Han, "Hypothesis testing with the general source," *IEEE Trans. Inf. Theory*, vol. 46, no. 7, pp. 2415–2427, Nov. 2000.
- [9] H. Nagaoka and M. Hayashi, "An information-spectrum approach to classical and quantum hypothesis testing for simple hypotheses," *IEEE Trans. Inf. Theory*, vol. 53, no. 2, pp. 534–549, Feb. 2007.
- [10] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," *Ann. Math. Statist.*, vol. 19, no. 3, pp. 326–339, 1948.
- [11] T. S. Ferguson, *Mathematical Statistics: A Decision Theoretic Approach Paperback*. New York, NY, USA: Academic, 1967.
- [12] H. V. Poor and O. Hadjiladis, *Quickest Detection*. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [13] R. H. Farrell, "Asymptotic behavior of expected sample size in certain one sided tests," *Ann. Math. Statist.*, vol. 35, no. 1, pp. 36–72, 1964.
- [14] D. Siegmund, "The variance of one-sided stopping rules," *Ann. Math. Statist.*, vol. 40, no. 3, pp. 1074–1077, 1969.
- [15] H. Robbins and D. Siegmund, "The expected sample size of some tests of power one," *Ann. Statist.*, vol. 2, no. 3, pp. 415–436, 1974.
- [16] T. L. Lai, "Power-one tests based on sample sums," *Ann. Statist.*, vol. 5, no. 5, pp. 866–880, 1977.
- [17] G. Fellouris and A. G. Tartakovsky, "Nearly minimax one-sided mixture-based sequential tests," *Sequential Anal., Design Methods Appl.*, vol. 31, no. 3, pp. 297–325, 2012.
- [18] Y. S. Chow, H. Robbins, and D. Siegmund, *Great Expectations: The Theory of Optimal Stopping*, Boston, MA, USA: Houghton Mifflin, 1971.
- [19] T. W. Anderson, "A modification of the sequential probability ratio test to reduce the sample size," *Ann. Math. Statist.*, vol. 31, no. 1, pp. 165–197, 1960.
- [20] J. J. Busgang and M. Marcus, "Truncated sequential hypothesis tests," *IEEE Trans. Inf. Theory*, vol. 13, no. 3, pp. 512–516, Jul. 1967.
- [21] S. Tantarata and H. V. Poor, "Asymptotic efficiencies of truncated sequential tests," *IEEE Trans. Inf. Theory*, vol. 28, no. 6, pp. 911–923, Nov. 1982.
- [22] D. Siegmund, *Sequential Analysis: Tests and Confidence Intervals*. New York, NY, USA: Springer-Verlag, 1985.
- [23] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Belmont, MA, USA: Athena Scientific, 2000.
- [24] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Change-point Detection*. Boca Raton, FL, USA: CRC Press, 2015.
- [25] W. Zhang and H. V. Poor, "Opportunistic detection under a fixed-sample-size setting," *IEEE Trans. Inf. Theory*, vol. 59, no. 2, pp. 1107–1114, Feb. 2013.
- [26] H. Chernoff, *Sequential Analysis and Optimal Design*. Philadelphia, PA, USA: SIAM, 1972.
- [27] V. P. Draglia, A. G. Tartakovsky, and V. V. Veeravalli, "Multihypothesis sequential probability ratio tests. I. Asymptotic optimality," *IEEE Trans. Inf. Theory*, vol. 45, no. 7, pp. 2448–2461, Nov. 1999.
- [28] Y. Polyanskiy and S. Verdú, "Binary hypothesis testing with feedback," presented at the Inf. Theory Appl. Workshop (ITA), San Diego, CA, USA, Feb. 2011.
- [29] E. Hossain, Z. Han, and H. V. Poor, *Smart Grid Communications and Networking*. Cambridge, U.K.: Cambridge Univ. Press, 2012.
- [30] S. Verdú, "On channel capacity per unit cost," *IEEE Trans. Inf. Theory*, vol. 36, no. 5, pp. 1019–1030, Sep. 1990.
- [31] I. Csiszár, "On the capacity of noisy channels with arbitrary signal costs," *Problem Control Inf. Theory*, vol. 2, nos. 3–4, pp. 283–304, 1973.
- [32] A. N. Shiryaev, *Optimal Stopping Rules*. New York, NY, USA: Springer-Verlag, 1978.
- [33] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Hoboken, NJ, USA: Wiley, 2006.

Wenyi Zhang (S'00–M'07–SM'11) is with the faculty of the Department of Electronic Engineering and Information Science, University of Science and Technology of China. Prior to that, he was affiliated with the Communication Science Institute, University of Southern California, as a postdoctoral research associate, and with Qualcomm Incorporated, Corporate Research and Development. He studied at Tsinghua University and obtained his Bachelor's degree in Automation in 2001; he studied at the University of Notre Dame, Indiana, USA, and obtained his Master's and Ph.D. degrees, both in Electrical Engineering, in 2003 and 2006, respectively. His research interests include wireless communications and networking, information theory, and statistical signal processing.

George V. Moustakides (M'82–SM'97) was born in Drama, Greece, in 1955. He received the diploma in Electrical and Mechanical Engineering from the National Technical University of Athens, Greece, in 1979, the M.Sc in Systems Engineering from the Moore School of Electrical Engineering, University of Pennsylvania, Philadelphia, in 1980, and the Ph.D in Electrical Engineering and Computer Science from Princeton University, Princeton NJ, in 1983. Since 2007 he is a professor with the department of Electrical and Computer Engineering, University of Patras, Patras, Greece. Prof. Moustakides also held long-term appointments as Junior and Senior Researcher with INRIA, France and as Professor with the University of Thessaly, Greece. He also held numerous appointments as Visiting Scholar and Adjunct Professor with Princeton University, University of Pennsylvania, Columbia University, University of Maryland, Georgia Institute of Technology, University of Southern California, University of Illinois at Urbana-Champaign and Rutgers University. He has served as Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY. His interests include Sequential Detection and Statistical Signal Processing.

H. Vincent Poor (S'72–M'77–SM'82–F'87) received the Ph.D. degree in electrical engineering and computer science from Princeton University in 1977. From 1977 until 1990, he was on the faculty of the University of Illinois at Urbana-Champaign. Since 1990 he has been on the faculty at Princeton, where he is the Dean of Engineering and Applied Science, and the Michael Henry Strater University Professor of Electrical Engineering. He has also held visiting appointments at several other institutions, most recently at Imperial College and Stanford. His research interests are in the areas of information theory, stochastic analysis and statistical signal processing, and their applications in wireless networks and related fields. Among his publications in these areas is the recent book *Mechanisms and Games for Dynamic Spectrum Allocation* (Cambridge University Press, 2014).

Dr. Poor is a member of the National Academy of Engineering and the National Academy of Sciences, and is a foreign member of Academia Europaea and the Royal Society. He is also a fellow of the American Academy of Arts and Sciences, the Royal Academy of Engineering (U. K.), and the Royal Society of Edinburgh. In 1990, he served as President of the IEEE Information Theory Society, in 2004–07 as the Editor-in-Chief of these *Transactions*, and in 2009 as General Co-chair of the IEEE International Symposium on Information Theory, held in Seoul, South Korea. He received a Guggenheim Fellowship in 2002 and the IEEE Education Medal in 2005. Recent recognition of his work includes the 2014 URSI Booker Gold Medal, the 2015 EURASIP Athanasios Papoulis Award, the 2016 John Fritz Medal, and honorary doctorates from Aalborg University, Aalto University, HKUST, and the University of Edinburgh.