Nonparametric Iterated-Logarithm Extensions of the Sequential Generalized Likelihood Ratio Test

Jaehyeok Shin, Aaditya Ramdas, and Alessandro Rinaldo

***Abstract*—We develop a nonparametric extension of the sequential generalized likelihood ratio (GLR) test and corresponding time-uniform confidence sequences for the mean of a univariate distribution. By utilizing a geometric interpretation of the GLR statistic, we derive a simple analytic upper bound on the probability that it exceeds any prespecified boundary; these are intractable to approximate via simulations due to infinite horizon of the tests and the composite nonparametric nulls under consideration. Using time-uniform boundary-crossing inequalities, we carry out a unified nonasymptotic analysis of expected sample sizes of one-sided and open-ended tests over nonparametric classes of distributions (including sub-Gaussian, sub-exponential, sub-gamma, and exponential families). Finally, we present a flexible and practical method to construct time-uniform confidence sequences that are easily tunable to be uniformly close to the pointwise Chernoff bound over any target time interval.**

***Index Terms*—Sequential analysis, testing, maximum likelihood detection, error probability.**

I. Introduction

CONSIDER the following setup for open-ended sequential testing: we observe i.i.d. data sequentially from an infinite stream $X_1, X_2, \ldots$ generated by an unknown distribution P over the real line with finite first moment and belonging to large nonparametric class of distributions. We wish to test a one-sided hypothesis about its mean $\mu := \int x \mathrm{d}P(x)$ by deciding, at each time point, whether to reject the null hypothesis or instead to continue sampling, possibly indefinitely. With a slight abuse of notation, we denote with $\mathbb{P}_\mu$ and $\mathbb{E}_\mu$ the probability and expectation when the mean of the data generating distribution is equal to μ. (In many parametric models, one can safely assume that there is a one to one map between $\mathbb{P}_\mu$ and μ and, therefore, that the mean value parametrization is well-defined, but in the nonparametric settings considered in this paper there could be more than one distribution with the same mean μ; as discussed in Remark 2 below, this does not affect the validity of our results.) For a fixed $\alpha \in (0,1)$, we are concerned with developing level α open-ended tests —also known as tests with power one [1]—for the one-sided hypothesis problem about μ of the form

$$H_0 : \mu \le \mu_0 \ \text{ vs. } \ H_1 : \mu > \mu_1, \ \text{ for some } \mu_1 \ge \mu_0. \quad (1)$$

Formally, a level α open-ended test consists of a stopping time N with respect to the natural filtration generated by the data, which satisfies the constraints

$$\begin{aligned} &\text{Bounded type-1 error:} && \mathbb{P}_\mu(N < \infty) \le \alpha, \ \text{ if } \mu \le \mu_0, \\ &\text{Asymptotically power one:} && \mathbb{P}_\mu(N < \infty) = 1, \ \text{ if } \mu > \mu_1. \quad (2) \end{aligned}$$

At each time n, we either stop and reject the null if $N = n$ or continue sampling if $N > n$. In particular, when $\mu \le \mu_0$, we never stop with probability at least $1 - \alpha$. The case of $\mu_0 = \mu_1$ in which there is no separation between the null and alternative hypothesis will be of special interest.

The possibility of sampling indefinitely is characteristic of tests of power one [2], and stems from allowing for arbitrarily small signal strengths (or, equivalently of no separation between the null and the alternative). This might initially be viewed as an undesirable property. However, open-ended tests are typically adaptive to the underlying signal strength and will stop early when μ is much larger than μ_1. Indeed, open-ended tests not only have practical applications such as post-marketing drug and vaccine safety surveillance [3] but also serve as building blocks of more complicated sequential analyses. For instance, if we want to relax the power one constraint to a level $\beta \in (0,1)$ of type-2 error control, we can introduce another stopping time M corresponding to a level β open-ended test with swapped null and alternative hypotheses. Then, the minimum of the two stopping times N and M can be used as a sequential testing procedure that simultaneously controls type-1 and type-2 error at level α and β, respectively. In this case, with probability 1 the procedure will stop in finite time under both the null and alternative hypotheses.

The expected sample size $\mathbb{E}_\mu N$ under an alternative distribution with mean $\mu > \mu_1$ is a traditional and widely used measure to quantify the performance of a sequential testing procedure satisfying the error constraints in (2); see, e.g., [2], [4]–[6]. In particular, the smaller the expected time to rejection under the alternative, the better the test. In parametric settings, if the testing problem is simple, i.e., if it takes the form

$$H_0 : \mu = \mu_0 \ \text{ vs } \ H_1 : \mu = \mu_1 \quad (3)$$

for some $\mu_0 \neq \mu_1$, then the optimal testing procedure is the sequential probability ratio test (SPRT), originally put forward

Manuscript received October 15, 2020; revised March 19, 2021 and May 7, 2021; accepted May 8, 2021. Date of publication May 17, 2021; date of current version June 21, 2021. The work of Aaditya Ramdas was supported in part by the NSF DMS under Grant 1916320, and in part by the NSF CAREER under Grant 1945266. The work of Alessandro Rinaldo was supported by the NSF DMS under Grant 2015489. *(Corresponding author: Jaehyeok Shin.)*

Jaehyeok Shin and Alessandro Rinaldo are with the Department of Statistics and Data Science, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: shinjaehyeok@cmu.edu; arinaldo@cmu.edu).

Aaditya Ramdas is with the Department of Statistics and Data Science and the Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: aramdas@cmu.edu).

This article has supplementary downloadable material available at https://doi.org/10.1109/JSAIT.2021.3081105, provided by the authors.

Digital Object Identifier 10.1109/JSAIT.2021.3081105

by Wald [7] and further studied by Wald and Wolfowitz [4]. In detail, based on observations $X_1, X_2, \ldots$, the one-sided SPRT is defined by the stopping time

$$N := \inf\{n \geq 1 : \log \mathrm{L}_n(\mu_1, \mu_0) \geq A\}, \tag{4}$$

where $A \geq 0$ is an appropriately chosen threshold and $\mathrm{L}_n(\mu_1, \mu_0)$ is the likelihood ratio (LR) statistic given by

$$\mathrm{L}_n(\mu_1, \mu_0) := \prod_{i=1}^{n} \frac{p_{\mu_1}(X_i)}{p_{\mu_0}(X_i)}. \tag{5}$$

Here p_{μ_1} and p_{μ_0} are the probability densities functions (with respect to a common dominating measure) of the data generating distributions under the alternative and null hypothesis, respectively. If the threshold A is chosen to satisfy the constraint in (2), then the one-sided SPRT is optimal in the sense that it minimizes $\mathbb{E}_{\mu_1} N$, the expected sample size for a rejection under the alternative, among all test satisfying the error constraints; see, e.g., [4], [8].

More generally, one may wish to design a sequential test that satisfies the error constraints (2) while nearly minimizing the expected sample size uniformly over many possible null and alternatives distributions. However, it is well known that in this composite setting the SPRT does not yield such a guarantee [9], [10].

A natural way to extend the SPRT to accommodate composite null and alternative hypotheses in parametric settings is to use the generalized likelihood ratio (GLR) statistic:

$$\mathrm{GL}_n(\mu_1, \mu_0) := \frac{\sup_{\mu > \mu_1 \text{ or } \mu \leq \mu_0} \prod_{i=1}^{n} p_\mu(X_i)}{\sup_{\mu \leq \mu_0} \prod_{i=1}^{n} p_\mu(X_i)}. \tag{6}$$

(Note that the above GLR statistic cannot be smaller than one by definition.) The one-sided sequential GLR (SGLR) test can then be defined by the stopping time

$$N_g := \inf\{n \geq 1 : \log \mathrm{GL}_n(\mu_1, \mu_0) \geq g_\alpha(n)\}, \tag{7}$$

where $g_\alpha : \mathbb{N} \to [0, \infty)$ is a *boundary function,* appropriately chosen to ensure that the error constraints in (2), namely $\sup_{\mu \leq \mu_0} \mathbb{P}_\mu(N_g < \infty) \leq \alpha$ and $\inf_{\mu > \mu_1} \mathbb{P}_\mu(N_g < \infty) = 1$, are fulfilled.

For exponential families, Farrell [5] derived a sharp asymptotic lower bound on the expected sample size of any composite sequential test satisfying the constraint (2) in the moderate confidence regime in which the testing level α is fixed but μ_1 approaches μ_0. The author further proposed a procedure to threshold the GLR statistic that attains this lower bound in the limit as the gap $|\mu_1 - \mu_0| \to 0$, but did not provide an explicit boundary function g_α. Lorden [11] obtained an explicit boundary and nonasymptotic bounds on the testing errors and the expected sample sizes for well-separated alternatives ($\mu_1 > \mu_0$); see Section III-A.

Remark 1: Throughout the paper, we will follow the convention from the sequential analysis literature of using the term asymptotic to describe a vanishing separation between the null and the alternative hypotheses or a vanishing value of α. Accordingly, we will say that a result holds nonasymptotically when it holds for all finite values of μ_0, μ_1 or α and not only in the limit.

In the more challenging non-separated case ($\mu_1 = \mu_0$), virtually all of the existing SGLR tests are designed under parametric assumptions. In this case, one may hope to calibrate a boundary to a desired level α using simulations, but this is a non-trivial task for one-sided, open-ended tests because the type-1 error guarantee must hold for an infinite time horizon. Further, in nonparametric settings, the choice of which distribution to use is itself not obvious. Our analytic boundaries solve these issues. Also, most older works only deliver asymptotic analyses of error bounds and of expected sample sizes in the high-confidence regime where $\alpha \to 0$; see, e.g., [10], [12]. In contrast, our boundaries allow for a thorough nonasymptotic analysis in this setting.

It is also important to point out that the recent literature on best-arm identification has produced several time-uniform law of iterated logarithm ("finite LIL") bounds that allow for nonasymptotic analyses of type-1 error bounds with explicit boundary functions; see [13]–[17]. However, most existing boundaries are of a rigid form and are difficult to tune to be as tight as possible over any target time interval. Further, the expected sample sizes of corresponding testing procedures have been studied mostly in the asymptotic framework of the high confidence regime in which $\alpha \to 0$ but μ is fixed. As a result, nonasymptotic expected sample size analyses for the "moderate confidence regime" (fixed α, $\mu \to \mu_0$) are not yet thoroughly studied beyond the (sub-)Gaussian case. In this paper, we bring to bear and sharpen tools from this line of work and apply them in novel ways to the problem of designing SGLR tests.

Below we outline our main contributions, which advance both the theory and practice of sequential testing based on the GLR statistic. A technical summary of our results can be also found in Table I in the Discussion Section V.

1) We design new SGLR tests that satisfy the error constraints in (2) with an explicit boundary function that are applicable in nonparametric settings in which a likelihood function is not available. Specifically, we present a unified analysis of sequential testing for sub-Gaussian, sub-exponential, and exponential family distributions (among others) via a new geometric interpretation of GLR statistics.
2) We derive novel nonasymptotic bounds on the sample size for the SGLR tests that hold both in expectation and with high probability and are valid under any alternative. Though our results apply to nonparametric families of distributions, the bounds match in rate the known lower bounds for exponential family distributions in the moderate confidence regime where α is fixed and $|\mu_1 - \mu_0| \to 0$.
3) Leveraging the duality between sequential tests and confidence sequences [15, Sec. 6], we develop a flexible method to construct confidence sequences which can be easily tuned to be uniformly close to the fixed-sample Chernoff bound on prespecified time intervals.

The rest of the paper is organized as follows. In Section II, we introduce the sub-ψ_M family of distributions, a nonparametric generalization of the exponential family which includes sub-Gaussian, sub-exponential, and exponential family distributions as special cases. Section III presents the sequential GLR-like (SGLR-like) test, which is a nonparametric counterpart of the SGLR test for the sub-ψ_M family of distributions. We then derive nonasymptotic bounds on expected sample sizes, which demonstrate that the proposed SGLR-like test can detect the alternative signal in a sample efficient way. In Section IV, we introduce a flexible method to build anytime-valid confidence sequences that can be tuned to be close to the pointwise Chernoff bound on target time intervals. We conclude with a brief summary of our contribution and discussions on future directions. In the interest of space, we defer proofs and simulations to the supplement.

II. GLR Statistic for the Exponential Family and Its Nonparametric Extension

Before presenting our main results in full generality and in order to build some intuition for our results, we first review the GLR statistic in the standard setting of exponential families of distributions. Consider a natural exponential family of distributions with densities of the form

$$p_\theta(x) = \exp\{\theta x - B(\theta)\}, \quad \theta \in \Theta \subset \mathbb{R}, \tag{8}$$

with respect to a reference Borel measure ν on the real line, where $\Theta \subset \{\theta \in \mathbb{R} : \int e^{\theta x}\nu(\mathrm{d}x) < \infty\}$ is the natural parameter space and $B \colon \Theta \to \mathbb{R}$ is a strictly convex function given by $\theta \mapsto B(\theta) = \int e^{\theta x}\nu(\mathrm{d}x)$. Throughout, we assume Θ to be nonempty and open. For each natural parameter $\theta \in \Theta$, let $\mu = \mu(\theta)$ be the corresponding mean parameter

$$\mu = \int x p_\theta(x)\nu(\mathrm{d}x).$$

It is well known [18] that, under the stated assumptions, there is a one-to-one correspondence between natural and mean parameters via $\mu = \nabla B(\theta)$, where throughout the manuscript, for a differentiable univariate function f, ∇f will refer to its derivative function. Therefore, we can reparameterize the exponential family based on the mean parameter space $M := \{\nabla B(\theta) : \theta \in \Theta\}$, so that, each point $\mu \in M$ will uniquely identify the density $p_\mu := p_{(\nabla B)^{-1}(\mu)}$.

When the samples $X_1, X_2, \ldots$ are i.i.d. from a distribution in the family, for any $\mu_0, \mu_1 \in M$, the likelihood ratio (LR) statistic based on first n samples is defined by

$$\mathrm{L}_n(\mu_1, \mu_0) := \prod_{i=1}^{n} \frac{p_{\mu_1}(X_i)}{p_{\mu_0}(X_i)}. \tag{9}$$

For fixed choices of μ_0 and μ_1, it is well-known that the normalized log LR statistics can be expressed as a function of the sample mean $\bar{X}_n = \frac{1}{n}\sum_{i=1}^{n} X_i$ in two equivalent forms as follows:

$$\frac{1}{n}\log \mathrm{L}_n(\mu_1, \mu_0) = \mathrm{KL}(\bar{X}_n, \mu_0) - \mathrm{KL}(\bar{X}_n, \mu_1) \tag{10}$$

$$= \mathrm{KL}(\mu_1, \mu_0) + \nabla_z \mathrm{KL}(z, \mu_0)|_{z=\mu_1}(\bar{X}_n - \mu_1), \tag{11}$$

where $\mathrm{KL}(z_1, z_2)$ is the Kullback-Leibler (KL) divergence from p_{z_2} to p_{z_1}, for $z_1, z_2 \in M$. For the completeness, we derive the above identities in Appendix A in the supplementary material. From the first expression in (10), we see that the normalized LR statistic is equal to the difference between the KL divergences from the distribution in the family parametrized by the sample mean $\bar{X}_n$ to the one corresponding to the null and alternative hypotheses. Perhaps more importantly for our derivations below, the second expression in (11) shows that the normalized log LR statistic, as a function of $\bar{X}_n$, is also the tangent line to the KL divergence function $z \mapsto \mathrm{KL}(z, \mu_0)$ at $z = \mu_1$. See Figure 1 for an illustration.

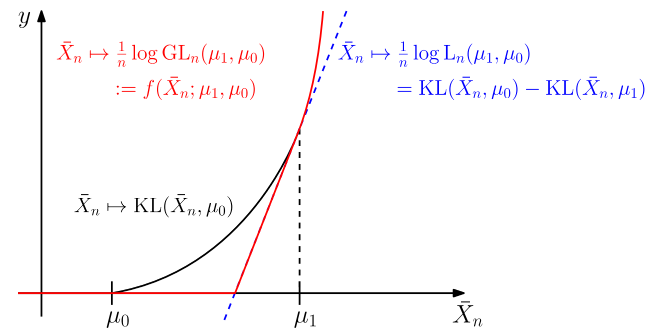

Fig. 1. Illustration of normalized log LR and GLR statistics for exponential family distributions. The dashed blue line corresponds to the normalized log LR statistic as a function of the sample mean $\bar{X}_n$ which is tangent to the KL divergence function at μ_1. The red line shows the normalized log GLR statistic which is equal to the KL divergence for $\bar{X}_n > \mu_1$ and its "clipped" tangent line for $\bar{X}_n \le \mu_1$, respectively.

Now, recall that we are concerned with one-sided composite testing problem

$$H_0 : \mu \le \mu_0 \quad \text{vs} \quad H_1 : \mu > \mu_1, \tag{12}$$

where $\mu_1 \ge \mu_0 \in M$. From the expression of the LR statistic in (10), it can be easily shown that the corresponding GLR statistic is given by

$$\mathrm{GL}_n(\mu_1, \mu_0) = \sup_{z > \mu_1} \mathrm{L}_n(z, \mu_0) \vee 1.$$

Using the alternative expression of the LR statistic in (11), we conclude that the normalized log GLR statistic can be written as $\frac{1}{n}\log \mathrm{GL}_n(\mu_1, \mu_0) := f(\bar{X}_n; \mu_1, \mu_0)$, where

$$f(z; \mu_1, \mu_0) := \begin{cases} \left[\mathrm{KL}(\mu_1, \mu_0) + \nabla_z \mathrm{KL}(z, \mu_0)|_{z=\mu_1}(z - \mu_1)\right]_+ & \text{if } z \le \mu_1 \\ \mathrm{KL}(z, \mu_0) & \text{if } z > \mu_1 \end{cases}. \tag{13}$$

Figure 1 depicts the normalized log LR (dashed blue line) and GLR statistics (red line) based on the first n samples, as functions of $\bar{X}_n$. In particular, the normalized log GLR statistic is equal to the KL divergence between $p_{\bar{X}_n}$ and p_{μ_0} when $\bar{X}_n > \mu_1$ and to its tangent line at μ_1, clipped at zero, otherwise.

The above expression for the normalized GLR statistic and its simple but revealing geometric interpretation provide the conceptual underpinning and intuition for much of the contributions made in this article. In the next subsection, we introduce a class of distributions, called sub-ψ distributions [19], [20] that exhibit analogous GLR statistic for the testing problem at hand and thus can be viewed as a natural nonparametric generalization of the exponential family distributions.

A. Extensions to the Sub-ψ_M Family Distributions

In this section we will assume throughout some familiarity with basic concepts from convex analysis; see, e.g., [21]. Let M be an open, convex subset of $\mathbb{R}$ such that, for each $\mu \in M$, there exists an extended real-valued convex function ψ_μ that is finite and strictly convex on a common closed supporting set $\Lambda \subset \mathbb{R}$ and differentiable on its nonempty interior Λ^o containing 0, with $\psi_\mu(0) = 0$ and $\nabla\psi_\mu(0) = \mu$. We also assume that, for each $\mu \in M$, the convex conjugate of ψ_μ, denoted with ψ_μ^*, is finite and differentiable on M.

For any $\mu \in M$, a collection $\mathcal{P}$ of probability distributions over the real line is said to be a *sub-ψ_μ family* if, for each $P \in \mathcal{P}$, $\text{supp}(P) \subset \overline{M}$, $\mathbb{E}_{X\sim P}[X] = \mu$ and

$$\log \mathbb{E}_{X\sim P}\left[e^{\lambda X}\right] \leq \psi_\mu(\lambda), \quad \forall\lambda \in \Lambda. \tag{14}$$

We will denote such class with $\mathcal{P}_{\psi_\mu}$. Notice that, for each $\mu \in M$, $\mathcal{P}_{\psi_\mu}$ is a nonparametric statistical model and the set M plays the role of all possible mean parameters of interest. Finally, we will write

$$\mathcal{P}_{\psi_M} := \bigcup_{\mu\in M} \mathcal{P}_{\psi_\mu}$$

for the collection of all sub-ψ_μ families as μ ranges in M, which we will then refer to as a *sub-ψ_M family of distributions.*

We will further require that a sub-ψ_M family satisfies the following *order-preserving* property for the Bregman divergences arising from the conjugate functions ψ_μ^*.

Definition 1: For each $\mu \in M$, let $D_{\psi_\mu^*}(\cdot,\cdot)$ be the Bregman divergence with respect to ψ_μ^*. We say a sub-ψ_M family of distributions has an order-preserving class of Bregman divergences if

$$D_{\psi_{z_0}^*}(\mu_1, z_0) \geq D_{\psi_{\mu_0}^*}(\mu_1, \mu_0) \tag{15}$$

holds for any $z_0, \mu_0, \mu_1 \in M$ such that $\mu_1 \leq \mu_0 \leq z_0$ or $z_0 \leq \mu_0 \leq \mu_1$.

At a high-level, the order-preserving property implies that the Bregman divergence expresses a natural ordering of a sub-ψ_M family with respect to the mean parametrization. This ordering is naturally suited to handle the hypothesis testing problem we are studying. As an important special case, sub-Gaussian distributions with a common variance parameter σ^2 form a sub-ψ_M family with an order-preserving class of Bregman divergence where $M = \Lambda = \mathbb{R}, \psi_\mu(\lambda) = \lambda\mu + \frac{\sigma^2}{2}\lambda^2$ and $D_{\psi_{\mu_0}^*}(\mu_1, \mu_0) = \frac{1}{2\sigma^2}(\mu_1 - \mu_0)^2$. Another important example of a sub-ψ_M family of distributions with an order-preserving class of the Bregman divergences is the family of Bernoulli distributions where we have $M = (0, 1)$, $\Lambda = \mathbb{R}$, $\psi_\mu(\lambda) = \log(1-\mu+\mu e^\lambda)$ and $D_{\psi_{\mu_0}^*}(\mu_1, \mu_0) = \text{KL}(\mu_1, \mu_0) = \mu_1 \log(\frac{\mu_1}{\mu_0}) + (1 - \mu_1)\log(\frac{1-\mu_1}{1-\mu_0})$. These two examples are representative distributions belonging to the following two large nonparametric classes of distributions.

1) *Additive sub-ψ distributions:* A sub-ψ_M family is said to be additive if for any $\mu \in M$,

$$\psi_\mu(\lambda) - \lambda\mu = \psi_0(\lambda) := \psi(\lambda) \qquad \text{for all } \lambda \in \mathbb{R}. \tag{16}$$

It can be checked that sub-Gaussian and sub-exponential distributions are instances of additive sub-ψ distributions. For each additive sub-ψ distribution, the Bregman divergence can be expressed as $D_{\psi_{\mu_0}^*}(\mu_1, \mu_0) = \psi^*(\mu_1 - \mu_0)$ for each $\mu_1, \mu_0 \in M$.

2) *Exponential family-like (EF-like) sub-B distributions:* A sub-ψ_M family of distributions is called an EF-like sub-B family if there exists an extended real-valued convex function B which is finite, strictly convex and differentiable on $\Lambda := (\nabla B)^{-1}(M)$ such that, for each $\mu \in M$,

$$\psi_\mu(\lambda) = B\big(\lambda + \theta_\mu\big) - B\big(\theta_\mu\big), \ \forall\lambda \in \mathbb{R}, \tag{17}$$

where $\theta_\mu := (\nabla B)^{-1}(\mu)$. All exponential family distributions and sub-Gaussian distributions are instances of EF-like sub-B distributions. In fact, for exponential family of distributions, it is immediate to see that each ψ_μ is the logarithm of the moment generating function. For each EF-like sub-B distribution, the Bregman divergence can be written as $D_{\psi_{\mu_0}^*}(\mu_1, \mu_0) = D_{B^*}(\mu_1.\mu_0)$, which, for a class of exponential families with densities satisfying (8) is equal to the dual of the KL divergence; see, e.g., [22].

From the expressions of the Bregman divergence described above, it follows that all additive sub-ψ and EF-like sub B families have classes of Bregman divergences satisfying the order-preserving property. For completeness, this fact and related properties of sub-ψ_M family of distributions are proven in Appendix A in the supplementary material.

We now describe how to construct LR- and GLR-like statistics for the nonparametric class of distributions $\mathcal{P}_M$ in ways that mirror exactly the derivation of the LR and GLR statistics in exponential families, as described in the previous section. We will assume throughout that $X_1, X_2, \ldots$ is a sequence of independent, though not necessarily identically distributed, random variables with the same but unknown finite mean μ, each drawn from a distribution belonging to a common sub-class $\mathcal{P}_{\psi_\mu}$ of a sub-ψ_M family of distributions.

Remark 2: Slightly overloading notation, we denote with $\mathbb{P}_\mu$ and $\mathbb{E}_\mu$ the probability and expectation for the stochastic process $(X_n)_{n\in\mathbb{N}}$. That is, every statement written with respect to $\mathbb{P}_\mu$ and $\mathbb{E}_\mu$ refers to the case where each independent observation of the underlying stochastic process $(X_n)_{n\in\mathbb{N}}$ has a distribution in $\mathcal{P}_{\psi_\mu}$ with the same mean μ. We emphasize that, because of the nonparametric nature of our models, the observations need not be identically distributed. Further, since there could be more than one distribution with mean μ, every statement related to $\mathbb{P}_\mu$ and $\mathbb{E}_\mu$ should be understood as a reference to any possible sub-ψ_μ distribution of $(X_n)_{n\in\mathbb{N}}$.

Now, consider the following test for the mean:

$$H_0 : \mu = \mu_0 \ \text{ vs } \ H_1 : \mu = \mu_1, \tag{18}$$

for some $\mu_1, \mu_0 \in M$. Let $\lambda_1 := \nabla\psi_{\mu_0}^*(\mu_1)$, and define

$$\mathrm{L}_n(\mu_1, \mu_0) := \exp\big\{n\big[\lambda_1\bar{X}_n - \psi_{\mu_0}(\lambda_1)\big]\big\}. \tag{19}$$

The above expression has the same form of the likelihood ration statistics (9) for parametric exponential families. Thus, with a slight abuse of notation, we refer to $\mathrm{L}_n(\mu_1, \mu_0)$ as the *LR-like statistic* for the above simple hypothesis testing based on first n samples.

Just like in the case of an exponential family, for a sub-ψ_M family of distributions, the normalized log LR-like statistics can be re-written as

$$\frac{1}{n}\log \mathrm{L}_n(\mu_1, \mu_0) = D_{\psi^*_{\mu_0}}(\bar{X}_n, \mu_0) - D_{\psi^*_{\mu_0}}(\bar{X}_n, \mu_1) \tag{20}$$

$$= D_{\psi^*_{\mu_0}}(\mu_1, \mu_0) + \nabla_z D_{\psi^*_{\mu_0}}(z, \mu_0)|_{z=\mu_1}(\bar{X}_n - \mu_1), \tag{21}$$

where $D_{\psi^*_{\mu_0}}(z_1, z_2)$ is the Bregman divergence from z_2 to z_1 with respect to $\psi^*_{\mu_0}$ for $z_1, z_2 \in M$. See Appendix A in the supplementary material for details. Thus, similarly to the exponential family case, the normalized LR-like statistic is equal to the difference between two Bregman divergences, one from each hypothesis, to the sample mean. Also, from the second expression in (21), we can check that the normalized log LR-like statistic is given by the tangent line to the Bregman divergence function $z \mapsto D_{\psi^*_{\mu_0}}(z, \mu_0)$ at $z = \mu_1$. This is not a coincidence. Recall that every exponential family distribution is an EF-like sub-B distribution. In this case, the corresponding LR and LR-like statistics are equal to each other since, for any $\mu_0, \mu_1 \in M$,

$$D_{\psi^*_z}(\mu_1, \mu_0) = D_{B^*}(\mu_1, \mu_0) = \mathrm{KL}(\mu_1, \mu_0), \quad \forall z \in M. \tag{22}$$

Finally, based on the definition of the LR-like statistic, for the one-sided testing problem

$$H_0 : \mu \le \mu_0 \quad \text{vs} \quad H_1 : \mu > \mu_1,$$

with $\mu_1 \ge \mu_0 \in M$, the *GLR-like statistic* is defined by

$$\mathrm{GL}_n(\mu_1, \mu_0) := \sup_{z > \mu_1} \mathrm{L}_n(z, \mu_0) \vee 1. \tag{23}$$

From the expressions of LR statistic in (21), we can derive a closed form expression for the normalized log GLR-like statistic $\frac{1}{n}\log \mathrm{GL}_n(\mu_1, \mu_0) := f(\bar{X}_n; \mu_1, \mu_0)$, where

$$f(z; \mu_1, \mu_0) := \begin{cases} \Big[D_{\psi^*_{\mu_0}}(\mu_1, \mu_0) \\ \quad + \nabla_z D_{\psi^*_{\mu_0}}(z, \mu_0)\,|_{z=\mu_1}\,(z - \mu_1)\Big]_+ & \text{if } z \le \mu_1 \\ D_{\psi^*_{\mu_0}}(z, \mu_0) & \text{if } z > \mu_1 \end{cases} \tag{24}$$

This expression matches exactly the one for the normalized GLR statistics in exponential families given in (13) and in fact recovers it as a special case since for exponential families the Bregman divergence correspond to the KL divergence.

See Figure 2 for an illustration of the relationship between normalized log LR-like and GLR-like statistics based on first n samples which is identical to the exponential family case in Figure 1. The dashed blue line corresponds to the normalized log LR-like statistic as a function of the sample mean $\bar{X}_n$ which is tangent to the Bregman divergence function at μ_1. The red line shows the normalized log GLR-like statistic which is equal to the Bregman divergence for $\bar{X}_n > \mu_1$ and its "clipped" tangent line for $\bar{X}_n \le \mu_1$, respectively. Since both LR- and GLR-like statistics depend only on the sample mean $\bar{X}_n$ and not on the entire history, we can update both test statistics (at each step) using constant time and memory. This property makes it possible to run sequential tests and confidence sequences introduced below in a fully online fashion.

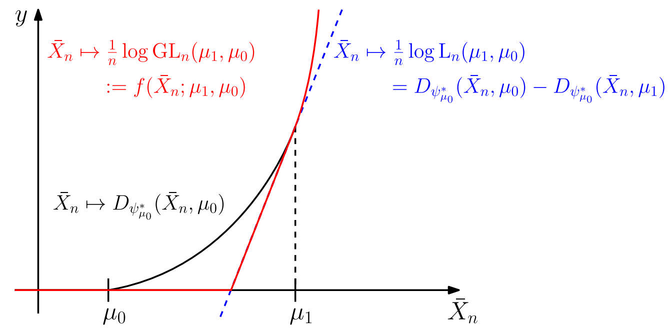

Fig. 2. Illustration of normalized log LR-like and GLR-like statistics for the sub-ψ_M family of distributions. The dashed blue line corresponds to the normalized log LR-like statistic as a function of the sample mean $\bar{X}_n$ which is tangent to the Bregman divergence function at μ_1. The red line shows the normalized log GLR-like statistic which is equal to the Bregman divergence for $\bar{X}_n > \mu_1$ and its "clipped" tangent line for $\bar{X}_n \le \mu_1$, respectively.

III. SGLR-Like Test for Sub-ψ_M Family Distributions

We now describe sequential tests based on the GLR-like statistics for one-sided hypotheses

$$H_0 : \mu \le \mu_0 \quad \text{vs} \quad H_1 : \mu > \mu_1, \tag{25}$$

for some fixed $\mu_1 \ge \mu_0 \in M$.

Consider the settings described in the previous section and let $\alpha \in (0, 1)$ be fixed. A level α SGLR-like test is defined by a stopping time of the form

$$N_{\mathrm{GL}} := \inf\{n \ge 1 : \log \mathrm{GL}_n(\mu_1, \mu_0) \ge g_\alpha(n)\}, \tag{26}$$

where g_α is a given, positive function on $[1, \infty)$. In particular, if $\mu_1 = \mu_0$ then the SGLR-like test can be simplified to

$$N_{\mathrm{GL}} = \inf\Big\{n \ge 1 : \bar{X}_n \ge \mu_0, nD_{\psi^*_{\mu_0}}(\bar{X}_n, \mu_0) \ge g_\alpha(n)\Big\}. \tag{27}$$

At each time point, we check if the stopping criteria are met and, if so, we reject the null hypothesis. In general, by the law of large numbers, if $g_\alpha(n)/n \to 0$ as $n \to \infty$, then the power one guarantee $\mathbb{P}_\mu(N_{\mathrm{GL}} < \infty) = 1$ can be easily satisfied for each $\mu > \mu_1$. However, it is a nontrivial task to design a proper boundary function g_α satisfying the type-1 error control

$$\sup_{\mu \le \mu_0} \mathbb{P}_\mu(N_{\mathrm{GL}} < \infty) \le \alpha. \tag{28}$$

To tackle this challenge, we develop a general methodology to bound the boundary crossing probability of the event in (26). This result plays a key role in building the SGLR-like test.

Theorem 1: Let the boundary function $g : [1, \infty) \to [0, \infty)$ satisfy the following conditions:

1) g is nonnegative and nondecreasing;
2) the mapping $t \mapsto g(t)/t$ is nonincreasing on $[1, \infty)$ and $\lim_{t\to\infty} g(t)/t = 0$.

Then, for any sub-ψ_M family of distributions, the crossing probability under the null is such that

$$\sup_{\mu\leq\mu_0} \mathbb{P}_\mu(\exists n \geq 1 : \log \mathrm{GL}_n(\mu_1,\mu_0) \geq g(n)) \tag{29}$$

$$\leq \begin{cases} e^{-g(1)} & \text{if } D_{\psi^*_{\mu_0}}(\mu_1,\mu_0) \geq g(1), \\ \inf_{\eta>1} \sum_{k=1}^{K_\eta} e^{-g(\eta^k)/\eta} & \text{otherwise,} \end{cases} \tag{30}$$

where, for any $\eta > 1$, $K_\eta \in \mathbb{N} \cup \{0,\infty\}$ is defined by

$$K_\eta := \inf\left\{k \in \{0\} \cup \mathbb{N} : D_{\psi^*_{\mu_0}}(\mu_1,\mu_0) \geq \frac{g(\eta^k)}{\eta^k}\right\}. \tag{31}$$

The second boundary value for the probability of the curve-crossing event in the expression (30) is obtained using a technique known as "stitching" or "peeling" (see, e.g., [15]) that is designed to derive uniform probabilistic guarantees over an infinite time horizon. Roughly speaking, the time course is divided into geometrically spaced time epochs, and in each of them, a line crossing inequality is derived. The final bound is obtained by appropriately stitching together these separate linear boundaries. The parameter $\eta > 1$ is the ratio of adjacent "intrinsic time" (accumulated variance) epochs used in the stitching process. A smaller value of η yields a better approximation of the curve-crossing events but requires controlling a larger number (namely, K_η) of line-crossing events. The optimal choice of η minimizing the bound achieves the optimal balance for this trade-off. The infimum over $\eta > 1$ in the bound explicitly shows how to choose the best η. The stitching construction is only required when $D_{\psi^*_{\mu_0}}(\mu_1,\mu_0) \geq g(1)$. In the other case, it is sufficient to use a single line-crossing inequality. See Appendix B in the supplementary material for a detailed explanation and a formal proof.

Remark 3: In the above theorem, we follow the convention that $\inf\emptyset = \infty$. From the two conditions assumed for the boundary function g, we can check that $K_\eta = 0$ if and only if $D_{\psi^*_{\mu_0}}(\mu_1,\mu_0) \geq g(1)$ and $K_\eta = \infty$ if and only if $\mu_1 = \mu_0$ for each $\eta > 1$.

Remark 4: If the mapping $\mu \to \psi_\mu(\lambda)$ is concave for each fixed λ, it is possible to extend Theorem 1 to a broader class of stochastic processes. In detail, let $(X_i)_{i\in\mathbb{N}}$ be a real-valued process adapted to an underlying filtration $(\mathcal{F}_i)_{i\in\{0\}\cup\mathbb{N}}$ such that, conditioned on $\mathcal{F}_{i-1}$, each X_i follows a sub-ψ_{μ^i} distribution with $\mu^i := \mathbb{E}[X_i \mid \mathcal{F}_{i-1}]$. That is, for each $i \in \mathbb{N}$ and $\lambda \in \mathbb{R}$,

$$\log \mathbb{E}\left[e^{\lambda X_i} \mid \mathcal{F}_{i-1}\right] \leq \psi_{\mu^i}(\lambda) \quad a.s. \tag{32}$$

Let $\mathcal{P}_{\mu_0}$ be the set of probability distributions of $(X_i)_{i\in\mathbb{N}}$ such that $\frac{1}{n}\sum_{i=1}^n \mu^i := \bar{\mu}_n \leq \mu_0$ for all n. Then, it is possible to show that

$$\sup_{P\in\mathcal{P}_{\mu_0}} P(\exists n \geq 1 : \log \mathrm{GL}_n(\mu_1,\mu_0) \geq g(n)) \tag{33}$$

$$\leq \begin{cases} e^{-g(1)} & \text{if } D_{\psi^*_{\mu_0}}(\mu_1,\mu_0) \geq g(1) \\ \inf_{\eta>1} \sum_{k=1}^{K_\eta} e^{-g(\eta^k)/\eta} & \text{otherwise.} \end{cases} \tag{34}$$

Note that the concavity condition for sub-ψ_μ function is satisfied by all additive sub-ψ distributions and many important exponential families with discrete support, including Bernoulli, Poisson, Geometric and Negative binomial with known number of failures (Appendix B in the supplementary material).

In the following subsections we will deploy Theorem 1 to develop SGLR-like tests for the one-sided hypothesis (25). We will analyze separately the case in which the null and alternative hypotheses are well-separated ($\mu_1 > \mu_0$), and the case of no separation ($\mu_1 = \mu_0$).

A. SGLR-Like Tests for Well-Separated Alternatives

In this subsection, we focus on the scenario where μ_1 is strictly larger than μ_0. In many applications, this separation condition can be derived from prior knowledge of the underlying distribution or from requirements on the minimal effect sizes one seeks to detect. Furthermore, even if we intend to run an open-ended testing procedure, there might be an upper limit on the sample size due to time and budget constraints on the experiment. In this case, the separation of null and alternative hypotheses can be imposed indirectly by the upper limit because if the null and alternative hypotheses are too close to each other then no fixed-level test can detect such a small separation given the upper limit on the sample size; see [23]. In Remark 6 below, we will present a natural way of choosing the boundary of the alternative space for μ_1 given an upper limit on sample size $n_{\max}$ for the sequential testing procedure proposed in this subsection.

Now, choosing a constant boundary $g \in \mathbb{R}^+$, Theorem 1 immediately yields that

$$\sup_{\mu\leq\mu_0} \mathbb{P}_\mu(\exists n \geq 1 : \log \mathrm{GL}_n(\mu_1,\mu_0) \geq g)$$

$$\leq \begin{cases} e^{-g} & \text{if } D_1 \geq g, \\ \inf_{\eta>1} \left\lceil \log_\eta\left(\frac{g}{D_1}\right)\right\rceil e^{-g/\eta} & \text{otherwise,} \end{cases} \tag{35}$$

where $D_1 := D_{\psi^*_{\mu_0}}(\mu_1,\mu_0)$.

Remark 5: For the constant boundary case, the term in the upper bound that depends on the infimum over $\eta > 1$ can be rewritten as

$$\inf_{\eta>1} \left\lceil \log_\eta\left(\frac{g}{D_1}\right)\right\rceil e^{-g/\eta} = \inf_{k\in\mathbb{N}} k \exp\left\{-g\left(\frac{D_1}{g}\right)^{1/k}\right\}. \tag{36}$$

The right expression can be evaluated efficiently since it optimizes over integers, not reals.

For any given level $\alpha \in (0,1]$, let $g_\alpha(\mu_1,\mu_0) > 0$ be a constant boundary value that makes the right hand side of (35) equal to α. The boundary value $g_\alpha(\mu_1,\mu_0)$ can be numerically computed using equation (36). Alternatively, as shown in Appendix C in the supplementary material, $g_\alpha(\mu_1,\mu_0)$ can be upper bounded as follows:

$$g_\alpha(\mu_1,\mu_0) \leq \inf_{\eta>1}\left\{\eta \log\left(\frac{1}{\alpha}\left[1 + 2\log_\eta\left(\frac{\eta\sqrt{\eta}}{\alpha D_1 \log\eta} \vee 1\right)\right]\right)\right\}. \tag{37}$$

This bound is slightly loose but useful for the numerical computation of $g_\alpha(\mu_1, \mu_0)$.

Based on $g_\alpha(\mu_1, \mu_0)$, let $N_{\rm GL}(g_\alpha, \mu_1, \mu_0)$ be the stopping time of the SGLR-like test defined by

$$N_{\rm GL}(g_\alpha, \mu_1, \mu_0) := \inf\{n \geq 1 : \log \mathrm{GL}_n(\mu_1, \mu_0) \geq g_\alpha(\mu_1, \mu_0)\}. \tag{38}$$

By Theorem 1, we can check that $N_{\rm GL}(g_\alpha, \mu_1, \mu_0)$ induces a valid level α test. Furthermore, from Lorden's inequality [24], we can derive a nonasymptotic upper bound on the expected sample size under any alternative distribution, as shown next.

Theorem 2: Let $N_{\rm GL}(g_\alpha, \mu_1, \mu_0)$ be the stopping time of the SGLR-like test defined in (38). Then, $\sup_{\mu \leq \mu_0} \mathbb{P}_\mu(N_{\rm GL}(g_\alpha, \mu_1, \mu_0) < \infty) \leq \alpha$. Furthermore, if the observations $X_1, X_2, \ldots$ are i.i.d. then, for any $\mu \geq \mu_1 \in M$, we have

$$\mathbb{E}_\mu\left[N_{\rm GL}(g_\alpha, \mu_1, \mu_0)\right] \leq \frac{g_\alpha(\mu_1, \mu_0)}{D_{\psi^*_{\mu_0}}(\mu, \mu_0)} + \left[\frac{\sigma_\mu \nabla \psi^*_{\mu_0}(\mu)}{D_{\psi^*_{\mu_0}}(\mu, \mu_0)}\right]^2 + 1, \tag{39}$$

where $\sigma_\mu^2 := \sup_{P \in \psi_\mu} \int (x-\mu)^2 \mathrm{d}P$ is the maximal variances of all the probability distributions in the sub-ψ_μ class which can be upper bounded as $\sigma_\mu^2 \leq \nabla^2 \psi_\mu(0) = 1/\nabla^2 \psi^*_\mu(\mu)$.

In the special but highly relevant case of σ^2-sub-Gaussian distributions, the upper bound on the expected sample size in (39) can be written as

$$\mathbb{E}_\mu\left[N_{\rm GL}(g_\alpha, \mu_1, \mu_0)\right] \leq \frac{2\sigma^2\left[g_\alpha(\mu_1, \mu_0) + 2\right]}{(\mu - \mu_0)^2} + 1. \tag{40}$$

The proof of Theorem 2 can be found in Appendix B in the supplementary material.

Under the same setting of well-separated hypotheses, Lorden [11] investigated the properties of the SGLR test for exponential family distributions with a constant boundary. For any constant $g > 0$, Lorden proved that

$$\begin{aligned} &\sup_{\mu \leq \mu_0} \mathbb{P}_\mu(n \geq 1 : \log \mathrm{GL}_n(\mu_1, \mu_0) \geq g) \\ &\leq \begin{cases} e^{-g} & \text{if } \mathrm{KL}(\mu_1, \mu_0) \geq g, \\ \left(1 + \frac{g}{\mathrm{KL}(\mu_1, \mu_0)}\right) e^{-g} & \text{otherwise.} \end{cases} \end{aligned} \tag{41}$$

In fact, the above bound can be immediately extended to hold also over sub-ψ_M families by replacing the KL divergence term $\mathrm{KL}(\mu_1, \mu_0)$ with the Bregman divergence $D_{\psi^*_{\mu_0}}(\mu_1, \mu_0)$.

To compare our result with Lorden's bound, let $g^L_\alpha(\mu_1, \mu_0)$ be the smallest boundary value one can obtain from Lorden's bound in (41) for any given α. In general, neither our choice $g_\alpha(\mu_1, \mu_0)$ nor Lorden's choice $g^L_\alpha(\mu_1, \mu_0)$ dominates the other one. However, in the moderate confidence regime where α is fixed and $\mu_1 \to \mu_0$, our bound increases at an exponentially slower rate compared to the one stemming from Lorden's $g^L_\alpha(\mu_1, \mu_0)$. In detail, we can check that in the moderate confidence regime we have

$$g^L_\alpha(\mu_1, \mu_0) = O\Big(\log\Big(1/D_{\psi^*_{\mu_0}}(\mu_1, \mu_0)\Big)\Big), \tag{42}$$

$$g_\alpha(\mu_1, \mu_0) = O\Big(\log\log\Big(1/D_{\psi^*_{\mu_0}}(\mu_1, \mu_0)\Big)\Big), \tag{43}$$

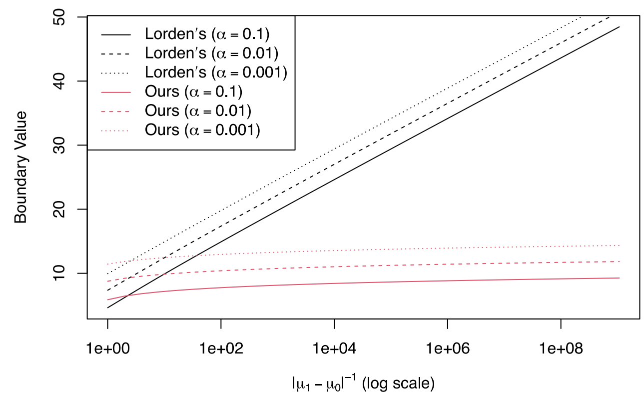

Fig. 3. Boundary values $g^L_\alpha(\mu_1, \mu_0)$ (black lines) and $g_\alpha(\mu_1, \mu_0)$ (red lines) for level α GLR tests based on Lorden's (41) upper bounds grow significantly faster than our (35) upper bounds.

where the log-log dependency makes it possible to apply our testing method to cases in which the separation is exponentially small. This is well illustrated in Fig. 3 by comparing the curves $|\mu_1 - \mu_0|^{-1} \mapsto g^L_\alpha$ and $|\mu_1 - \mu_0|^{-1} \mapsto g_\alpha$ for a normal distributions with $\sigma = 1$ (notice that the separation is shown on a log-scale). From the plot, we can check that, even for the exponentially small separation (i.e., $|\mu_1 - \mu_0| \approx 10^{-10}$), the boundary value remains at a practical level.

The log-log dependency of the boundary value $g_\alpha(\mu_1, \mu_0)$ and the corresponding upper bound on the expected sample size in Theorem 2 are sharp since, in the worst case where $\mu = \mu_1$, the bounds on the expected sample size can be shown to be asymptotically tight in exponential families of distributions. In detail, for an exponential family of distributions parameterized by the mean parameter $\mu \in M$ and a given level $\alpha \in (0, 1]$, let N be a stopping time such that $\mathbb{P}_{\mu_0}(N < \infty) \leq \alpha$ for a fixed $\mu_0 \in M$. Then, by Farrell's theorem on open-ended tests [2], [5], the expected sample size under alternative sequences can be lower bounded as:

$$\begin{aligned} \limsup_{\mu_1 \to \mu_0} \frac{(|\mu_1 - \mu_0|^2 / 2\sigma^2_{\mu_0}) \mathbb{E}_{\mu_1} N}{\log\log(1/|\mu_1 - \mu_0|)} &\geq \mathbb{P}_{\mu_0}(N = \infty) \\ &\geq 1 - \alpha, \end{aligned} \tag{44}$$

where $\sigma^2_{\mu_0} := \nabla^2 B(\theta_{\mu_0})$ is the variance of the probability distribution at $\mu = \mu_0$. Using the Taylor series expansion of B^* and the fact that $\mathrm{KL}(\mu_1, \mu_0) = D_{B^*}(\mu_1, \mu_0)$, we see

$$\begin{aligned} \mathrm{KL}(\mu_1, \mu_0) &= D_{B^*}(\mu_1, \mu_0) \\ &= B^*(\mu_1) - B^*(\mu_0) - \nabla B^*(\mu_0)(\mu_1 - \mu_0) \\ &= \frac{1}{2} \nabla^2 B^*(\mu_0)(\mu_1 - \mu_0)^2 + o\Big(|\mu_1 - \mu_0|^2\Big) \\ &= \frac{1}{2\sigma^2_{\mu_0}}(\mu_1 - \mu_0)^2 + o\Big(|\mu_1 - \mu_0|^2\Big), \end{aligned}$$

where the last equality is due to the fact that $\nabla^2 B^*(\mu_0) = [\nabla^2 B(\theta_{\mu_0})]^{-1}$. Farrell's lower bound in (44) then yields

$$\begin{aligned} \limsup_{\mu_1 \to \mu_0} \frac{\mathrm{KL}(\mu_1, \mu_0) \mathbb{E}_{\mu_1} N}{\log\log(1/\mathrm{KL}(\mu_1, \mu_0))} &\geq \mathbb{P}_{\mu_0}(N = \infty) \\ &\geq 1 - \alpha. \end{aligned} \tag{45}$$

Therefore, we conclude that the log-log dependence in (43) and in the corresponding upper bound on the expected sample size in (35) cannot be avoided.

The upper bound on the expected sample size in Theorem 2 is not fully adaptive to the underlying true but unknown mean parameter μ since the boundary value $g_\alpha(\mu_1, \mu_0)$ depends on the boundary of the alternative space μ_1 instead of each alternative mean value μ itself. Although we can make the separation gap $|\mu_1 - \mu_0|$ between null and alternative spaces double-exponentially small while maintaining the threshold $g_\alpha(\mu_1, \mu_0)$ at a practical level, it is an interesting and practically relevant problem to design a SGLR-like test whose expected sample size is fully adaptive to the unknown alternative distribution. In the next subsection, we address this question using the boundary-crossing probability bound in Theorem 1.

B. SGLR-Like Tests With No Separation

Below we consider the case of no separation (i.e., $\mu_0 = \mu_1$) between the null and alternative spaces in the one-sided test:

$$H_0 : \mu \le \mu_0 \quad \text{vs} \quad H_1 : \mu > \mu_0. \tag{46}$$

From the order-preserving property of the Bregman divergence, the log GLR-like statistic based on the first n observations can be written as $nD_{\psi^*_{\mu_0}}(\bar{X}_n, \mu_0)\mathbb{1}(\bar{X}_n \ge \mu_0)$ and the upper bound on the boundary crossing probability from Theorem 1 is of the form

$$\begin{aligned}\sup_{\mu \le \mu_0} \quad & \mathbb{P}_\mu\Big(\exists n \ge 1 : \bar{X}_n \ge \mu_0, nD_{\psi^*_{\mu_0}}(\bar{X}_n, \mu_0) \ge g(n)\Big) \\ & \le \inf_{\eta > 1} \sum_{k=1}^{\infty} \exp\Big\{-g(\eta^k)/\eta\Big\}. \end{aligned} \tag{47}$$

Since the right hand side now involves an infinite sum, we cannot use a constant function g for the boundary function. In fact, from the law of iterated logarithm, we know that the boundary function g must increase at least at a log-log scale as the number of samples n goes to infinity to get a nontrivial bound on the crossing probability.

For simplicity, in order to build SGLR-like tests at level $\alpha \in (0, 1]$, we will consider the boundary function

$$g^c_\alpha(n) := c\big[\log(1/\alpha) + 2\log\big(\log_c cn\big)\big], \tag{48}$$

where $c > 1$ is a fixed constant. Then, using (47),

$$\begin{aligned}\sup_{\mu \le \mu_0} \ & \mathbb{P}_\mu\Big(\exists n \ge 1 : \bar{X}_n \ge \mu_0, \\ & \qquad nD_{\psi^*_{\mu_0}}(\bar{X}_n, \mu_0) \ge c\big[\log(1/\alpha) + 2\log\big(\log_c cn\big)\big]\Big) \\ & \le \inf_{\eta > 1} \sum_{k=1}^{\infty} \exp\Big\{-\frac{c}{\eta}\log(1/\alpha)\Big\} \frac{1}{\big(1 + k\log_c \eta\big)^{2c/\eta}} \\ & \le \alpha \sum_{k=1}^{\infty} \frac{1}{(1+k)^2} = \alpha\left(\frac{\pi^2}{6} - 1\right) \le \alpha. \end{aligned} \tag{49}$$

Hence, a level α open-ended sequential testing procedure for (46) can be obtained based on the stopping time

$$\begin{aligned} N_{\mathrm{GL}}\big(g^c_\alpha, \mu_0\big) := \inf\Big\{& n \ge 1 : \bar{X}_n \ge \mu_0, nD_{\psi^*_{\mu_0}}\big(\bar{X}_n, \mu_0\big) \\ & \ge c\big[\log(1/\alpha) + 2\log\big(\log_c cn\big)\big]\Big\}. \end{aligned} \tag{50}$$

Unlike the scenario analyzed in the previous subsection, however, Lorden's inequality is no longer applicable to the non-constant, non-linear boundary g^c_α. Below we present novel, nonasymptotic probabilistic bounds on the stopping time $N_{\mathrm{GL}}(g^c_\alpha, \mu_0)$ under the alternative, which we can then use to provide a high-probability bound on the sample size. We remark that the finite-sample feature of our bounds sets it apart from other results in the literature, which for the most part have relied on asymptotic arguments.

To derive the desired probabilistic bounds, we rely on a notion of divergence between two sub-ψ_M distributions having different mean parameters $\mu > \mu_0$ given by

$$D^*_{\psi^*}(\mu, \mu_0) := \max\Big\{D_{\psi^*_{\mu_0}}\big(z^*, \mu_0\big), D_{\psi^*_\mu}\big(z^*, \mu\big)\Big\}, \tag{51}$$

where $z^* = z^*(\mu, \mu_0)$ is the minimizer of the function

$$z \in M \mapsto f(z) := \max\Big\{D_{\psi^*_{\mu_0}}(z, \mu_0), D_{\psi^*_\mu}(z, \mu)\Big\}, \tag{52}$$

which is also equal to the unique solution of the equation $D_{\psi^*_{\mu_0}}(z, \mu_0) = D_{\psi^*_\mu}(z, \mu)$ with respect to $z \in M$. For the exponential family case, the divergence (51) was introduced by Wong [25] to characterize the asymptotic behavior of the SGLR test under the alternative. For a concrete example, if the two distribution belongs to a sub-Gaussian class with common variance parameter σ^2 then

$$D^*_{\psi^*}(\mu, \mu_0) = \frac{1}{4}D_{\psi^*_{\mu_0}}(\mu, \mu_0) = \frac{1}{8\sigma^2}(\mu - \mu_0)^2,$$

while for Bernoulli distributions, we instead have

$$\begin{aligned}\frac{1}{2}(\mu - \mu_0)^2 & \le D^*_{\psi^*}(\mu, \mu_0) \\ & \le \mu\log\left(\frac{\mu}{\mu_0}\right) + (1-\mu)\log\left(\frac{1-\mu}{1-\mu_0}\right). \end{aligned}$$

Note that the $D^*_{\psi^*}$ divergence satisfies $D^*_{\psi^*}(\mu_0, \mu_0) = 0$ and $D^*_{\psi^*}(\mu, \mu_0) = D^*_{\psi^*}(\mu_0, \mu)$ for any $\mu, \mu_0 \in M$. Using $D^*_{\psi^*}$, we can formulate the following high probability bound on the stopping time $N_{\mathrm{GL}}(g^c_\alpha, \mu_0)$ under the alternative.

Theorem 3: For any fixed $\mu > \mu_0 \in M$, $\delta \in (0, 1]$ and $c > 1$, define $T_{\mathrm{high}}(\delta)$ by

$$T_{\mathrm{high}}(\delta) := \inf\left\{t \ge 1 : \frac{c\big[\log(1/\delta) + 2\log\big(\log_c ct\big)\big]}{D^*_{\psi^*}(\mu, \mu_0)} \le t\right\}. \tag{53}$$

Then, for any $\delta \in (0, \alpha)$, we have

$$\mathbb{P}_\mu\big(N_{\mathrm{GL}}\big(g^c_\alpha, \mu_0\big) \le T_{\mathrm{high}}(\delta)\big) \ge 1 - \delta. \tag{54}$$

Also, $T_{\mathrm{high}}(\delta)$ can be upper bounded by $\max(1, A)$, where

$$A = \frac{2c}{D^*_\mu}\log(1/\delta) + \frac{2c}{D^*_\mu}\log\left(2\log_c\left(\frac{2c^2}{\log c}\right) + 2\big[\log_c(1/D^*_\mu)\big]_+\right), \tag{55}$$

with $D^*_\mu := D^*_{\psi^*}(\mu, \mu_0)$.

The proof of Theorem 3 can be found in Appendix B in the supplementary material. As a consequence, we have the following upper bound on the expected sample size:

$$\mathbb{E}_\mu\left[N_{\mathrm{GL}}(g_\alpha^c,\mu_0)\right] \le 1 + \frac{2c}{D_\mu^*}\log(1/\alpha) + \frac{2c}{D_\mu^*}\log\left(2\log_c\left(\frac{2c^{2.5}}{\log c}\right) + 2\left[\log_c\left(1/D_\mu^*\right)\right]_+\right), \tag{56}$$

Note that both bounds, in probability and in expectation, are fully adaptive to the true alternative parameter μ via the divergence $D_\mu^* := D_{\psi_{\mu_0}^*}^*(\mu,\mu_0)$. For exponential family distributions, the above upper bound matches Farrell's optimal rate given in (44) up to a constant factor under the moderate confidence regime in which $\mu \to \mu_0$ and α is fixed.

IV. Confidence Sequences via GLR-Like Statistics

Previously, we discussed how to build open-ended sequential testing procedure for the one-sided testing problem based on the general upper bound on the boundary crossing probability given in Theorem 1. Relying on the same bound, we can construct confidence sequences for μ. A level α confidence sequence for μ is a sequence of sets $\{\mathrm{CI}_n\}_{n\in\mathbb{N}}$ such that

$$\mathbb{P}_\mu(\forall n \ge 1 : \mu \in \mathrm{CI}_n) \ge 1-\alpha, \quad \forall \mu \in M. \tag{57}$$

For any chosen boundary g and mapping $\mu_0 \mapsto \mu_1(\mu_0)$ with $\mu_1 > \mu_0$, each CI_n is defined by

$$\mathrm{CI}_n := \left\{\mu_0 \in M : \log \mathrm{GL}_n(\mu_1,\mu_0) < g(n), \text{ or } \inf_{\eta>1}\sum_{k=1}^{1\vee K_\eta}\exp\left\{-g\left(\eta^k\right)/\eta\right\} > \alpha\right\}, \tag{58}$$

where K_η is given in (31). Of course, different choices of the boundary function g and of the mapping $\mu_0 \mapsto \mu_1(\mu_0)$ result in confidence sequences of different shapes. For any such choice, the above construction of CI_n is guaranteed to yield a valid level α confidence sequence since

$$\begin{aligned}
\mathbb{P}_\mu(\forall n \ge 1 : \mu \in \mathrm{CI}_n) &= 1 - \mathbb{P}_\mu(\exists n \ge 1 : \mu \notin \mathrm{CI}_n)\\
&\ge 1 - \mathbb{P}_\mu(\exists n \ge 1 : \log \mathrm{GL}_n(\mu_1(\mu),\mu) \ge g(n))\\
&\quad \cdot \mathbb{1}\left(\inf_{\eta>1}\sum_{k=1}^{1\vee K_\eta}\exp\left\{-g(\eta^k)/\eta\right\} \le \alpha\right)\\
&\ge 1 - \left(\inf_{\eta>1}\sum_{k=1}^{1\vee K_\eta}\exp\left\{-g(\eta^k)/\eta\right\}\right)\\
&\quad \cdot \mathbb{1}\left(\inf_{\eta>1}\sum_{k=1}^{1\vee K_\eta}\exp\left\{-g(\eta^k)/\eta\right\} \le \alpha\right)\\
&\ge 1-\alpha,
\end{aligned}$$

for each $\mu \in M$ where the second-to-last inequality comes from the bound in Theorem 1. In this section, we present a slightly generalized version of Equation (35) to obtain confidence sequences that are valid over $\mathbb{N}$ and uniformly close to the Chernoff bound on chosen finite time intervals.

A. Confidence Sequence Uniformly Close to the Chernoff Bound on Finite Time Intervals

Recall that for each n and $\alpha \in (0,1]$, the (pointwise) Chernoff bound for sub-ψ_M distributions takes the form

$$\mathbb{P}_\mu\left(\bar{X}_n \ge \mu_\epsilon\right) \le e^{-nD_{\psi_\mu^*}(\mu_\epsilon,\mu)}, \tag{59}$$

for any fixed $\mu_\epsilon > \mu \in M$. Since the mapping $z \mapsto D_{\psi_\mu^*}(z,\mu)$ is increasing on $[\mu,\infty)\cap M$, by letting $g := nD_{\psi_\mu^*}(\mu_\epsilon,\mu)$, the Chernoff bound can be written as

$$\mathbb{P}_\mu\left(\bar{X}_n \ge \mu,\ nD_{\psi_\mu^*}(\bar{X}_n,\mu) \ge g\right) \le e^{-g}, \tag{60}$$

for any fixed $\mu \in M$.

As discussed in Section III, from the law of iterated logarithm, we know that the above bound cannot be extended to an anytime-valid bound. That is, there is no time-independent boundary having a nontrivial upper bound on the boundary-crossing probability

$$\mathbb{P}_\mu\left(\exists n \ge 1 : \bar{X}_n \ge \mu,\ nD_{\psi_\mu^*}(\bar{X}_n,\mu) \ge g\right), \tag{61}$$

for all $\mu \in M$. However, in virtually all practical cases, there exists an effective limit on the duration of the experiment and therefore on the sample size. And conversely, we often impose a minimum sample size requirement to, e.g., meet prespecified accuracy requirements, or to stabilize the experiment or to take account seasonality effects. As a result, in many situations, practitioners may find it desirable to use anytime confidence sequences that are uniformly close to the optimal pointwise Chernoff bound on a prespecified finite time interval $[n_{\min}, n_{\max}]$. We next describe such a guarantee.

Theorem 4: For any $g > 0$ and $\mu_0 \in M$, define

$$n_0 := \inf\left\{n \ge 1 : \sup_{z\in M, z>\mu_0} D_{\psi_{\mu_0}^*}(z,\mu_0) \ge g/n\right\}. \tag{62}$$

For any $n_0 \le n_{\min} < n_{\max} \in \mathbb{N}$, let $\mu_1 < \mu_2$ be solutions on $(\mu_0,\infty)\cap M$ of the equations

$$\frac{g}{n_{\max}} = D_{\psi_{\mu_0}^*}(\mu_1,\mu_0) \quad \text{and} \quad \frac{g}{n_{\min}} = D_{\psi_{\mu_0}^*}(\mu_2,\mu_0) \tag{63}$$

respectively. Then, the boundary crossing probability for the likelihood ratio process is such that

$$\mathbb{P}_{\mu_0}\left(\exists n \ge 1 : \sup_{z\in(\mu_1,\mu_2)} \log(\mathrm{L}_n(z,\mu_0) \vee 1) \ge g\right) \le e^{-g}\mathbb{1}(n_{\min} > n_0) + \inf_{\eta>1}\left\lceil \log_\eta\left(\frac{n_{\max}}{n_{\min}}\right)\right\rceil e^{-g/\eta}. \tag{64}$$

Remark 6: If $n_0 = n_{\min} < n_{\max}$, we can get rid of μ_2 in the supremum in the boundary crossing probability in (64). In this case, the above inequality is reduced to the usual constant boundary crossing inequality from Section III-A with a specifically chosen alternative μ_1 given by $\frac{g}{n_{\max}} = D_{\psi_{\mu_0}^*}(\mu_1,\mu_0)$. From the perspective of the duality between sequential tests and confidence sequences, this observation tells us that if we have a practical upper limit of the sequential testing procedure described in Section III-A, it is natural to use the value μ_1 given by $\frac{g_\alpha}{n_{\max}} = D_{\psi_{\mu_0}^*}(\mu_1,\mu_0)$ as the boundary of the alternative space where $\alpha \in (0,1]$ is the target level of the test and

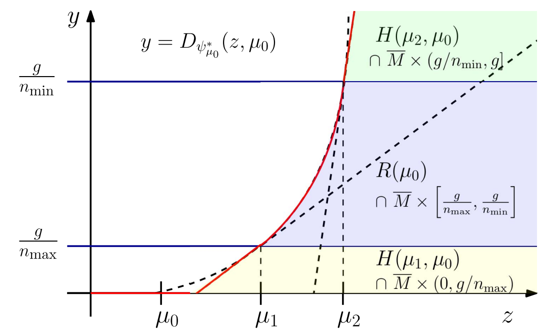

Fig. 4. Illustration of the boundary-crossing events and related regions $H(\mu_1, \mu_0), H(\mu_2, \mu_0)$ and $R(\mu_0)$ in (65).

g_α is the constant boundary which makes the right hand side of (64) equal to α.

Note that the event on the left hand side can be written as

$$\begin{aligned}
&\left\{\exists n \geq 1 : \left(\bar{X}_n, \frac{g}{n}\right) \in H(\mu_2, \mu_0) \cap R(\mu_0) \cap H(\mu_1, \mu_0)\right\} \\
&= \left\{\exists n \in [1, n_{\min}) : \left(\bar{X}_n, \frac{g}{n}\right) \in H(\mu_2, \mu_0)\right\} \\
&\cup \left\{\exists n \in [n_{\min}, n_{\max}] : \left(\bar{X}_n, \frac{g}{n}\right) \in R(\mu_0)\right\} \\
&\cup \left\{\exists n \in (n_{\max}, \infty) : \left(\bar{X}_n, \frac{g}{n}\right) \in H(\mu_1, \mu_0)\right\}, \qquad (65)
\end{aligned}$$

where the set $R(\mu_0)$ is defined by

$$R(\mu_0) := \left\{(z, y) \in \overline{M} \times [0, \infty) : y \leq D_{\psi^*_{\mu_0}}(z, \mu_0), z \geq \mu_0\right\}, \qquad (66)$$

and $H(\mu_1, \mu_0)$ and $H(\mu_2, \mu_0)$ are halfspaces contained in and tangent to $R(\mu_0)$ at $(\mu_1, g/n_{\max})$ and $(\mu_2, g/n_{\min})$, respectively. See Figure 4 for an illustration of $H(\mu_1, \mu_0), H(\mu_2, \mu_0)$ and $R(\mu_0)$.

In particular, on the time interval $[n_{\min}, n_{\max}]$, the boundary-crossing event can be written as

$$\left\{\exists n \in [n_{\min}, n_{\max}] : n\bar{X}_n \geq \mu_0, D_{\psi^*_{\mu_0}}(\bar{X}_n, \mu_0) \geq g\right\}. \qquad (67)$$

Therefore, we can check that the confidence sequence based on Theorem 4 is anytime-valid and uniformly close to the pointwise Chernoff bound on the time interval $[n_{\min}, n_{\max}]$.

More generally, we may have a sequence of time intervals $\{[n^{(k)}_{\min}, n^{(k)}_{\max}]\}_{k=1}^K$ on which we want the confidence sequence to be uniformly closer to the pointwise Chernoff bound. In this case, we can extend the bound in Theorem 4 as follows.

Corollary 1: Choose a sequence of boundary values $\{g^{(k)}\}_{k=1}^K$, $\mu_0 \in M$ and sequence of time intervals $\{[n^{(k)}_{\min}, n^{(k)}_{\max}]\}_{k=1}^K$ with $n_0 := n^{(0)}_{\max} \leq n^{(1)}_{\min} < n^{(1)}_{\max} \leq \cdots \leq n^{(K)}_{\min} < n^{(K)}_{\max}$. For each $k \in [K]$, let $\mu^{(k)}_1 < \mu^{(k)}_2$ be solutions on $(\mu_0, \infty) \cap M$ to the equations:

$$\frac{g^{(k)}}{n^{(k)}_{\max}} = D_{\psi^*_{\mu_0}}(\mu^{(k)}_1, \mu_0), \quad \text{and} \quad \frac{g^{(k)}}{n^{(k)}_{\min}} = D_{\psi^*_{\mu_0}}(\mu^{(k)}_2, \mu_0). \qquad (68)$$

Then, the boundary crossing probability can be bounded as

$$\begin{aligned}
&\mathbb{P}_{\mu_0}\left(\exists n \geq 1, \exists k \in [K] : \sup_{z \in (\mu^{(k)}_1, \mu^{(k)}_2)} \log(\mathrm{L}_n(z, \mu_0) \vee 1) \geq g^{(k)}\right) \\
&\leq \sum_{k=1}^K \left[e^{-g^{(k)}}\left[\mathbb{1}(n^{(k)}_{\min} > n^{(k-1)}_{\max})\right] + \inf_{\eta > 1}\left\lceil \log_\eta\left(n^{(k)}_{\max}/n^{(k)}_{\min}\right)\right\rceil e^{-g^{(k)}/\eta}\right]. \qquad (69)
\end{aligned}$$

Remark 7: If $n^{(k-1)}_{\max} = n^{(k)}_{\min}$ for each $k \in [K]$, the inequality (69) can be viewed as a piecewise constant boundary crossing probability of the form

$$\begin{aligned}
&\mathbb{P}_{\mu_0}\left(\exists n \geq 1 : \sup\nolimits_{z \geq \mu^{(K)}_1} \log(\mathrm{L}_n(z, \mu_0) \vee 1) \geq g_c(n)\right) \\
&\quad \leq \sum\nolimits_{k=1}^K \left[\inf\nolimits_{\eta > 1}\left\lceil \log_\eta\left(\frac{n^{(k)}_{\max}}{n^{(k)}_{\min}}\right)\right\rceil e^{-g^{(k)}/\eta}\right], \qquad (70)
\end{aligned}$$

where $g_c(n)$ is a piecewise constant function defined by

$$\begin{aligned}
g_c(n) := &\sum_{k=1}^K \min\left\{g^{(k-1)}, g^{(k)}\right\} \mathbb{1}\left(n = n^{(k)}_{\min}\right) \\
&+ g^{(k)} \mathbb{1}\left(n \in (n^{(k)}_{\min}, n^{(k)}_{\max})\right) \\
&+ \min\left\{g^{(k)}, g^{(k+1)}\right\} \mathbb{1}\left(n = n^{(k)}_{\max}\right). \qquad (71)
\end{aligned}$$

Here, we set $g^{(0)} := g^{(1)}$ and $g^{(K+1)} := g^{(K)}$. Note that if all boundary values $\{g^{(k)}\}_{k=1}^K$ are equal to each other, then the above inequality recovers the constant boundary-crossing inequality in Section III-A with a specifically chosen alternative $\mu^{(K)}_1$. In fact, the right hand side of the above inequality provides a sharper bound, since it allows us to choose different η for each k.

In what follows, we focus on the single time interval case in Theorem 4 for simplicity, but the arguments can be straightforwardly extended to the case of multiple time intervals in Corollary 1.

To convert the upper bound on the boundary-crossing probability in (64), for a given $\alpha \in (0, 1]$, let g_α be the constant boundary which makes the upper bound in (64) less than or equal to α. The boundary value g_α can be efficiently computed using the fact

$$\inf_{\eta > 1}\left\lceil \log_\eta\left(\frac{n_{\max}}{n_{\min}}\right)\right\rceil e^{-g/\eta} = \inf_{k \in \mathbb{N}} k \exp\left\{-g\left(\frac{n_{\min}}{n_{\max}}\right)^{1/k}\right\}. \qquad (72)$$

As shown in the previous section, in the moderate confidence regime where α is fixed and $n_{\max}/n_{\min} \to \infty$, we can check that $g_\alpha = O(\log\log(n_{\max}/n_{\min}))$. Therefore, in practice, we can still compute a boundary value g_α even for an exponential large $n_{\max}/n_{\min}$. Also, since g_α scales with respect to the ratio between two ends points of the time interval instead of its length, the confidence sequence can be designed to be uniformly close to the pointwise Chernoff bound on a time interval of exponentially large length. This property is of course especially useful when designing a large-scale experimentation.

For a given g_α, the corresponding confidence sequence is

$$\mathrm{CI}_n := \begin{cases}
\left\{\mu_0 \in M : \begin{array}{l} L_2(\bar{X}_n - \mu_2) \\ < \frac{g_\alpha}{n} - \frac{g_\alpha}{n_{\min}} \end{array}\right\} & \text{if } n \in [1, n_{\min}), \\
\left\{\mu_0 \in M : \begin{array}{l} D_{\psi^*_{\mu_0}}(\bar{X}_n, \mu_0) \\ < \frac{g_\alpha}{n} \text{ or } \bar{X} < \mu_0 \end{array}\right\} & \text{if } n \in [n_{\min}, n_{\max}], \\
\left\{\mu_0 \in M : \begin{array}{l} L_1(\bar{X}_n - \mu_1) \\ < \frac{g_\alpha}{n} - \frac{g_\alpha}{n_{\max}} \end{array}\right\} & \text{if } n \in (n_{\max}, \infty),
\end{cases} \qquad (73)$$

where $L_i := \nabla_z D_{\psi^*_{\mu_0}}(z, \mu_0) \mid_{z=\mu_i}$ for each $i = 1, 2$. Recall that both μ_1 and μ_2 depend on μ_0 via (63).

As a sanity check, we can verify that $[\bar{X}_n, \infty) \subset \mathrm{CI}_n$ for each $n \in \mathbb{N}$, since $\nabla_z D_{\psi^*_{\mu_0}}(z, \mu_0) > 0$ for any $z > \mu_0$. We can also check that the width of the confidence sequence (the distance from the left end point to the sample mean) does not shrink to zero even if $n \to \infty$ which implies that the confidence sequence can be loose outside of the target time interval $[n_{\min}, n_{\max}]$. Therefore, in practice, we recommend using the time interval of large enough size to cover the entire intended duration of experimentation.

Note that, for additive sub-ψ classes, μ_1 and μ_2 have the following simple relationships:

$$\mu_1 = \mu_0 + \psi_+^{*-1}\left(\frac{g_\alpha}{n_{\max}}\right) := \mu_0 + \Delta_1 \tag{74}$$

$$\mu_2 = \mu_0 + \psi_+^{*-1}\left(\frac{g_\alpha}{n_{\min}}\right) := \mu_0 + \Delta_2, \tag{75}$$

where ψ_+^{*-1} is the inverse function of $z \mapsto \psi^*(z)\mathbb{1}(z \geq 0)$. We can also check $L_i = \nabla\psi^*(\Delta_i)$ for each $i = 1, 2$.

Using this observation, we can derive confidence sequence for additive sub-ψ families in explicit form:

$$\mathrm{CI}_n := \begin{cases} \left(\bar{X}_n - \Delta_2 - L_2^{-1}\left(\frac{g_\alpha}{n} - \frac{g_\alpha}{n_{\min}}\right), \infty\right) & \text{if } n \in [1, n_{\min}), \\ \left(\bar{X}_n - \psi_+^{*-1}\left(\frac{g_\alpha}{n}\right), \infty\right) & \text{if } n \in [n_{\min}, n_{\max}], \\ \left(\bar{X}_n - \Delta_1 - L_1^{-1}\left(\frac{g_\alpha}{n} - \frac{g_\alpha}{n_{\max}}\right), \infty\right) & \text{if } n \in (n_{\max}, \infty). \end{cases} \tag{76}$$

For sub-Gaussian distributions with a parameter σ^2, this reduces to

$$\mathrm{CI}_n := \begin{cases} \left(\bar{X}_n - \sigma\sqrt{\frac{2g_\alpha}{n_{\min}}}\left[\frac{1}{2}\left(\frac{n_{\min}}{n} + 1\right)\right], \infty\right) & \text{if } n \in [1, n_{\min}), \\ \left(\bar{X}_n - \sigma\sqrt{\frac{2g_\alpha}{n}}, \infty\right) & \text{if } n \in [n_{\min}, n_{\max}], \\ \left(\bar{X}_n - \sigma\sqrt{\frac{2g_\alpha}{n_{\max}}}\left[\frac{1}{2}\left(\frac{n_{\max}}{n} + 1\right)\right], \infty\right) & \text{if } n \in (n_{\max}, \infty). \end{cases} \tag{77}$$

Remark 8: For additive sub-ψ distributions, the sequence of intervals defined in (76) can be applied to the time-varying mean case described in Remark 4. That is, for any chosen $\alpha \in (0, 1]$, the confidence sequence in (76) satisfies

$$\mathbb{P}(\exists n \geq 1 : \bar{\mu}_n \in \mathrm{CI}_n) \geq 1 - \alpha, \tag{78}$$

where $(\bar{\mu}_n)_{n \geq 1}$ is the sequence of running averages of conditional means, defined by

$$\bar{\mu}_n := \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\left[X_i \mid \mathcal{F}_{i-1}\right]. \tag{79}$$

In general, there is no closed form expression for the confidence sequence. However, if a sub-ψ_M family of distributions has order-preserving Bregman divergences, then for any given data, the mapping $\mu_0 \mapsto D_{\psi^*_{\mu_0}}(\bar{X}_n, \mu_0)\mathbb{1}(\bar{X}_n \geq \mu_0)$ is nonincreasing. Therefore, on the target time interval $[n_{\min}, n_{\max}]$, the confidence sequence given in (73) is an open-interval and it can be efficiently computed by binary search.

Outside of the target time interval, however, the confidence sequence is not necessarily an open interval. To avoid this potentially undesirable feature, below we introduce a sufficient condition under which we can guarantee that an EF-like sub-B family of distributions admits a confidence sequence consisting of open intervals.

Proposition 1: For a given EF-like sub-B family of distributions, suppose ∇B is a convex function. Then, for any data, the mapping $\mu_0 \mapsto \mathrm{L}_n(\mu_1, \mu_0)$ is nonincreasing on $(-\infty, \bar{X}_n] \cap M$ where μ_1 is a function of μ_0, and any $d > 0$, as defined by the solution of the equation

$$D_{\psi^*_{\mu_0}}(\mu_1, \mu_0)\mathbb{1}(\mu_1 \geq \mu_0) = d. \tag{80}$$

Consequently, the corresponding confidence sequence in (73) is an open-interval for each n.

The proof of Proposition 1 can be found in Appendix C in the supplementary material. For example, Poisson, Exponential, and Negative binomial (with a known number of failures) distributions are sub-classes of EF-like sub-B families satisfying the condition in Proposition 1, and thus we can efficiently compute confidence sequences by using binary search. On the other hand, Bernoulli distribution does not satisfies the condition, and thus we need to use a grid-search to compute the confidence sequence on the outside of the target time interval.

B. Tighter Confidence Sequences via Discrete Mixtures

Confidence sequences based on nonnegative mixture of martingales have been extensively studied [7], [15], [16], [26]–[28]. However, as discussed in [15], different choices of mixing methods yield different boundaries, and each confidence sequence is typically tighter and looser on some time intervals than others, so there is no time-uniformly dominating one.

From this point of view, it is an interesting question to choose a proper mixing method to obtain a confidence sequence with a desired shape, satisfying application-specific constraints. This subsection explains how we can use confidence sequences based on GLR-like statistics in the previous subsection to design discrete mixture-based confidence sequences that are almost uniformly close to the Chernoff bound on finite target time intervals.

The confidence sequences in the previous subsection are built on the constant boundary-crossing probability in Theorem 4, which is based on a GLR-like curve-crossing time. In this subsection, we will demonstrate that, for any GLR-like curve-crossing time, we can also build a discrete mixture of martingales such that the corresponding crossing time is always smaller than or equal to the GLR-like one. Therefore, we can always construct a tighter confidence sequence by using the discrete mixture-crossing time compared to the GLR-like one. Consequently, the shape of the obtained discrete mixture martingale based confidence sequence is dominated by the GLR-like based one, whose overall shape can be tailored to the specific application at hand.

To elaborate, for any $\alpha \in (0, 1]$ and target time interval $[n_{\min}, n_{\max}]$, let g_α be a positive value such that the upper bound in (64) is less than or equal to α. If $n_0 < n_{\min} = n_{\max}$ then both GLR-like and discrete mixture-crossing events are equal to the line-crossing event given by

$$\left\{\exists n \geq 1 : \begin{array}{l} D_{\psi^*_{\mu_0}}(\mu_1, \mu_0) \\ \quad +\nabla_z D_{\psi^*_{\mu_0}}(z, \mu_0)\mid_{z=\mu_1} (\bar{X}_n - \mu_1) \geq \frac{g_\alpha}{n} \end{array}\right\}, \quad (81)$$

which is equal to the pointwise Chernoff bound at $n = n_{\min} = n_{\max}$, and thus g_α can be chosen as $\log(1/\alpha)$. Therefore, in the rest of this subsection, we only consider the nontrivial case $n_0 \leq n_{\min} < n_{\max}$, where g_α is a positive value such that

$$e^{-g_\alpha} \mathbb{1}(n_{\min} > n_0) + \inf_{\eta > 1} \left\lceil \log_\eta \left(\frac{n_{\max}}{n_{\min}} \right) \right\rceil e^{-g_\alpha/\eta} \leq \alpha. \quad (82)$$

For a fixed g_α, let $\eta_\alpha > 1$ be the value attaining the infimum in the LHS of the above inequality. From the equivalent expression of $\inf_{\eta>1}\lceil \log_\eta(\frac{n_{\max}}{n_{\min}})\rceil e^{-g_\alpha/\eta}$ in (72), we have $\eta_\alpha = (n_{\max}/n_{\min})^{1/K_\alpha}$ where

$$K_\alpha := \arg\min_{k \in \mathbb{N}} \ k \exp\left\{ -g_\alpha \left(\frac{n_{\min}}{n_{\max}} \right)^{1/k} \right\}. \quad (83)$$

Next, for each $k = 0, 1, \ldots, K_\alpha$, let z_k be the function of μ_0 defined as the solution of the following equation:

$$D_{\psi^*_{\mu_0}}(z_k, \mu_0) = \frac{g_\alpha}{n_{\min} \eta_\alpha^k}, \quad z_k > \mu_0. \quad (84)$$

Finally, for given n samples, define a nonnegative random variable $M_n(\mu_0; \alpha)$ by

$$\begin{aligned} M_n(\mu_0; \alpha) &:= e^{-g_\alpha} \mathbb{1}(n_{\min} > n_0) \mathrm{L}_n(z_0(\mu_0), \mu_0) \\ &\quad + e^{-g_\alpha/\eta_\alpha} \sum_{k=1}^{K_\alpha} \mathrm{L}_n(z_k(\mu_0), \mu_0). \end{aligned} \quad (85)$$

Above, each $\mathrm{L}_n(z_k(\mu_0), \mu_0)$ is the LR-like statistic for $H_0 : \mu = \mu_0$ vs $H_1 : \mu = z_k(\mu_0)$, given by

$$\mathrm{L}_n(z_k(\mu_0), \mu_0) = \exp\left\{ n \left[\frac{g_\alpha}{n_{\min} \eta_\alpha^k} + L_k \left[\bar{X}_n - z_k(\mu_0) \right] \right] \right\}, \quad (86)$$

where $L_k := \nabla D_{\psi^*_{\mu_0}}(z, \mu_0) \mid_{z = z_k(\mu_0)}$ for each $k = 0, 1, \ldots, K_\alpha$. In particular, for additive sub-ψ family, each $\mathrm{L}_n(z_k(\mu_0), \mu_0)$ can be expressed as

$$\mathrm{L}_n(z_k(\mu_0), \mu_0) = \exp\left\{ n \left[\frac{g_\alpha}{n_{\min} \eta_\alpha^k} + L_k \left[\bar{X}_n - \mu_0 - \Delta_k \right] \right] \right\}, \quad (87)$$

where $\Delta_k := \psi_+^{*-1}(\frac{g_\alpha}{n_{\min}\eta_\alpha^k})$ and $L_k := \nabla \psi^*(\Delta_k))$.

Now, let $M_0(\mu_0; \alpha) := e^{-g_\alpha} \mathbb{1}(n_{\min} > n_0) + K_\alpha e^{-g_\alpha/\eta_\alpha} \in (0, \alpha]$. Then, under any sub-$\psi_{\mu_0}$ distribution, $\{M_n(\mu_0; \alpha)/M_0(\mu_0; \alpha)\}_{n \geq 0}$ is a nonnegative supermartingale with respect to the natural filtration. Therefore, by letting

$$\mathrm{CI}_n^M := \{\mu_0 \in M : M_n(\mu_0; \alpha)/M_0(\mu_0; \alpha) < 1/\alpha\},$$

we have a discrete mixture confidence sequence which is uniformly tighter than the GLR-like one.

Corollary 2: For any $\alpha \in (0, 1]$, the sequence of intervals $\{\mathrm{CI}_n^M\}_{n \in \mathbb{N}}$ is a valid level α confidence sequence satisfying

$$\mathbb{P}_\mu \left(\forall n : \mu \in \mathrm{CI}_n^M \right) \geq 1 - \alpha, \quad \forall \mu \in M. \quad (88)$$

Furthermore, we have $\mathrm{CI}_n^M \subset \mathrm{CI}_n$ for each $n \in \mathbb{N}$ where $\{\mathrm{CI}_n\}_{n \in \mathbb{N}}$ is the confidence sequence based on GLR-like statistics in (73) which is uniformly close to the pointwise Chernoff bound on the target time interval $[n_{\min}, n_{\max}]$.

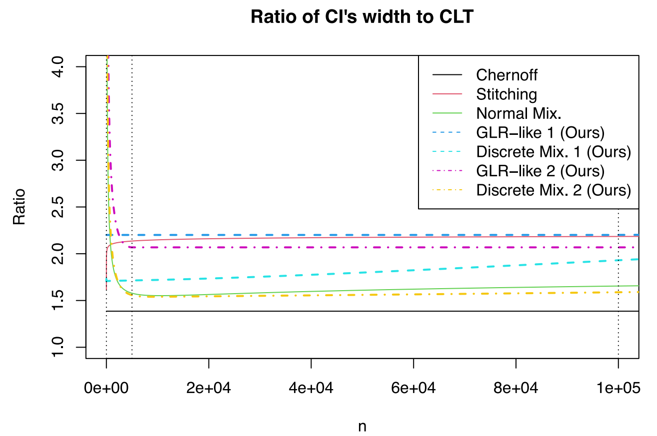

Fig. 5. Ratio of widths of the confidence intervals to the pointwise and asymptotically valid normal confidence intervals based on the central limit theorem. The black solid line corresponds to the pointwise and nonasymptotically valid Chernoff bound. Red and green solid lines come from the stitching and normal mixture method in [1], [15]. The rest of lines are based on our GLR-like confidence sequences for sub-Gaussian distributions in (77) and their discrete mixture counterparts with different choices of target time intervals ($[1, 10^5]$ and $[5 \times 10^3, 4 \times 10^5]$). See Section IV-C for the details of these confidence sequences.

Since g_α and K_α do not depend on μ_0, using Proposition 1 we can check that each CI_n^M is an open interval for every additive sub-ψ family and EF-like sub-B family of distributions with convex ∇B. Therefore, we can efficiently compute each confidence interval by binary-search. However, for general sub-ψ_M family of distributions, we may need to rely on grid-search methods to compute each confidence interval.

C. Examples of Sub-Gaussian Confidence Sequences

To illustrate the practicality of confidence sequences based on GLR-like statistics and corresponding discrete mixtures, in Figure 5, we calculate ratios of widths of confidence intervals to the pointwise and asymptotically valid normal confidence intervals based on the central limit theorem.[1] Here, we use the sub-Gaussian with $\sigma = 1$ for simplicity and set $\alpha = 0.025$.

The black solid line corresponds to the pointwise, nonasymptotically valid Chernoff bound. Red and green solid lines come from the stitching and normal mixture methods in [1], [15] where each confidence intervals for stitching $\mathrm{CI}_n^{\mathrm{ST}}$ and normal mixture method $\mathrm{CI}_n^{\mathrm{NM}}$ is given by

$$\mathrm{CI}_n^{\mathrm{ST}} := \left(\bar{X}_n - \frac{1.7}{\sqrt{n}} \sqrt{\log\log(2n) + 0.72 \log\left(\frac{5.2}{\alpha} \right)}, \infty \right)$$

$$\mathrm{CI}_n^{\mathrm{NM}} := \left(\bar{X}_n - \sqrt{2\left(\frac{1}{n} + \frac{\rho}{n^2} \right) \log\left(\frac{1}{2\alpha} \sqrt{\frac{n + \rho}{\rho} + 1} \right)}, \infty \right),$$

where we set $\rho = 1260$ by following the setting in [15, Fig. 9]. Note that the original normal mixture confidence interval in [1] did not have a closed form expression. In this subsection, we use the explicit closed form upper bound in [15] for simplicity. The rest of lines are based on our

[1]The R code to reproduce all the plots and simulation results of the paper is available on the repository https://github.com/shinjaehyeok/SGLRT_paper.

TABLE I
SUMMARY OF BOUNDARY CROSSING PROBABILITIES AND EXPECTED SAMPLE SIZES. ($D_1 \coloneqq D_{\psi^*_{\mu_0}}(\mu_1,\mu_0)$, $D_\mu \coloneqq D_{\psi^*_{\mu_0}}(\mu,\mu_0)$, $D_1^* \coloneqq D^*_{\psi^*_{\mu_0}}(\mu_1,\mu_0)$ AND $D_\mu^* \coloneqq D^*_{\psi^*_{\mu_0}}(\mu_1,\mu_0)$)

Task	Boundary crossing probability	Sample size
Test ($\mu_1 > \mu_0$) (Sec. III-A)	$\sup_{\mu \le \mu_0} \mathbb{P}_\mu(\exists n \ge 1 : \log \mathrm{GL}_n(\mu_1,\mu_0) \ge g)$ $\le e^{-g}\mathbb{1}(D_1 \ge g) + \inf_{\eta>1}\left\lceil \log_\eta\left(\frac{g}{D_1}\right)\right\rceil e^{-g/\eta}\mathbb{1}(D_1 < g)$	$O\left(\frac{\log\log(1/D_1)}{D_\mu}\right)$
Test ($\mu_1 = \mu_0$) (Sec. III-B)	$\sup_{\mu \le \mu_0} \mathbb{P}_\mu\left(\exists n \ge 1 : \bar{X}_n \ge \mu_0, nD_{\psi^*_{\mu_0}}(\bar{X}_n,\mu_0) \ge g(n)\right)$ $\le \inf_{\eta>1}\sum_{k=1}^{\infty}\exp\left\{-g(\eta^k)/\eta\right\}$	$O\left(\frac{\log\log(1/D_\mu^*)}{D_\mu^*}\right)$
Confidence sequence (Sec. IV)	$\mathbb{P}_{\mu_0}\left(\exists n \ge 1 : \sup_{z\in(\mu_1,\mu_2)} \log(\mathrm{L}_n(z,\mu_0) \vee 1) \ge g\right)$ $\le e^{-g}\mathbb{1}(n_{\min} > n_0) + \inf_{\eta>1}\left\lceil \log_\eta\left(\frac{n_{\max}}{n_{\min}}\right)\right\rceil e^{-g/\eta}.$	not applicable

GLR-like confidence sequences for sub-Gaussian distributions in (77) and their discrete mixture counterparts. For 'GLR-like 1' and 'Discrete Mixture 1' lines, we set $[n_{\min}, n_{\max}] = [1, 10^5]$. For 'GLR-like 2' and 'Discrete Mixture 2', we use $[n_{\min}, n_{\max}] = [5 \times 10^3, 4 \times 10^5]$. Vertical dotted lines are corresponding to the lines $n = 1, 5 \times 10^3, 10^5$. We can check both GLR-like confidence sequences are uniformly close to the pointwise Chernoff bound on their target time intervals and, as Corollary 2 tells, each discrete mixture counterpart has uniformly smaller widths of confidence intervals on its target time interval.

V. DISCUSSION

We have presented nonasymptotic analyses of sequential tests and confidence sequences based on GLR-like statistics, which can be viewed as a nonparametric generalization of the GLR statistic. Our main contribution is to provide a unified nonasymptotic framework for the sub-ψ_M family of distributions by leveraging a novel geometrical interpretation of GLR-like statistics and of the corresponding time-uniform concentration inequalities. In Table I, we provide a technical summary of our results by displaying the boundary crossing probabilities for the sequential tests and confidence sequences developed in the paper along with the corresponding expected sample sizes for the moderate confidence regime (fixed α, $\mu, \mu_1 \to \mu_0$). In the table, we set $D_1 \coloneqq D_{\psi^*_{\mu_0}}(\mu_1,\mu_0)$, $D_\mu \coloneqq D_{\psi^*_{\mu_0}}(\mu,\mu_0)$ and similarly for $D_1^* \coloneqq D^*_{\psi^*_{\mu_0}}(\mu_1,\mu_0)$ and $D_\mu^* \coloneqq D^*_{\psi^*_{\mu_0}}(\mu_1,\mu_0)$.

There remain several important open problems. First, although we have mainly focused on the case where each sample in a data stream $X_1, X_2, \ldots,$ is an independent random variable with the same mean, as we discussed in Remarks 4 and 8, some of the presented analyses can be naturally extended to a real-valued process $(X_i)_{i\in\mathbb{N}}$ adapted to a filtration $(\mathcal{F}_i)_{i\in\{0\}\cup\mathbb{N}}$ where the conditional expectation $\mu_i \coloneqq \mathbb{E}[X_i \mid \mathcal{F}_{i-1}]$ can vary over time. However, except for additive sub-ψ classes, it is unclear how to extend the expected sample size analysis of SGLR-like tests and the construction of the confidence sequences to the time-varying mean case. For example, the nonasymptotic upper bound on the expected sample size of SGLR-like test in Theorem 2 is only applicable to the i.i.d. random variables. Also, except for the additive sub-ψ case, the confidence sequence in Section IV is not applicable in time-varying settings. Generalizing these analyses to a more flexible nonparametric setting is an important open direction.

Second, the SGLR-like tests derived here and the corresponding confidence sequences are applicable only to the univariate case in which the underlying data stream is a real-valued sequence. It is a natural to inquire whether one can generalize the SGLR-like tests and confidence sequences to multiple sources based multivariate data streams. If we have multiple independent univariate data stream then there is a simple method to combine upper bounds on boundary crossing probabilities for univariate data stream in Theorem 1 into a multivariate one. To be specific, let K be the number of independent univariate data streams. For each $a \in [K]$, let $\{X^a_{N_a(t)}\}_{t\ge 0}$ be a sequence of independent observations from a sub-ψ_{μ^a} distribution where $N_a(t)$ is the number of sample from the a-th distribution at time $t \ge 0$. We assume $N_a(t) \ge 1$ for each $a \in [K]$ and $t \ge 0$. Define $\mathrm{GL}^a_t(\mu^a_1, \mu^a_0)$ be the GLR-like statistic based on $N_a(t)$ samples from the a-th distribution up to time t. Then, from Theorem 1, we can find a boundary function g^a_α such that the following inequality holds for all $\alpha \in (0, 1]$ and for each $a \in [K]$:

$$\mathbb{P}_0\left(\exists t \ge 0 : \log \mathrm{GL}^a_t(\mu^a_1, \mu^a_0) \ge g^a_\alpha(N_a(t))\right) \le \alpha, \quad (89)$$

where $\mathbb{P}_0$ is a null distribution of K independent data streams. Suppose each boundary function can be decomposed as $g^a_\alpha(n) = f^a(n) + h^a(\alpha)$ where f^a is a nonnegative function on $\mathbb{N}$ which does not depends on α and

h^a is a nonnegative and nonincreasing function on $(0,1]$ such that $\lim_{\alpha\to 0} h^a(\alpha) = \infty$. Then, for each $\epsilon > 0$, we have the following upper bound on the boundary-crossing probability for the multiple sources of univariate data streams:

$$\mathbb{P}_0\left(\exists t \geq 0 : \sum_{a=1}^{K} \log \mathrm{GL}_t^a(\mu_1^a, \mu_0^a) \geq \sum_{a=1}^{K} f^a(N_a(t)) + \epsilon\right) \leq \mathbb{P}_0\left(\sum_{a=1}^{K} h^a(U_a) \geq \epsilon\right), \tag{90}$$

where each U_a is an independent Uniform$[0,1]$ random variable. Since the right hand side only depends on K independent uniformly distributed random variables, we can evaluate the probability of right hand side via a simple Monte Carlo simulation. See Appendix C in the supplementary material for the derivation of the above inequality. It is an interesting open question whether we can derive similar upper bound for general multiple sources of multivariate data streams.

References

[1] H. Robbins, "Statistical methods related to the law of the iterated logarithm," *Ann. Math. Stat.*, vol. 41, no. 5, pp. 1397–1409, 1970.

[2] H. Robbins and D. Siegmund, "The expected sample size of some tests of power one," *Ann. Stat.*, vol. 2, no. 3, pp. 415–436, 1974.

[3] M. Kulldorff, R. L. Davis, M. Kolczak, E. Lewis, T. Lieu, and R. Platt, "A maximized sequential probability ratio test for drug and vaccine safety surveillance," *Sequential Anal.*, vol. 30, no. 1, pp. 58–78, 2011.

[4] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," *Ann. Math. Stat.*, vol. 19, no. 3, pp. 326–339, 1948.

[5] R. H. Farrell, "Asymptotic behavior of expected sample size in certain one sided tests," *Ann. Math. Stat.*, vol. 35, pp. 36–72, Mar. 1964.

[6] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Changepoint Detection*. Boca Raton, FL, USA: CRC, 2014.

[7] A. Wald, "Sequential tests of statistical hypotheses," *Ann. Math. Stat.*, vol. 16, no. 2, pp. 117–186, 1945.

[8] Y. Chow, H. Robbins, and D. Siegmund, *Great Expectations: The Theory of Optimal Stopping*. Boston, MA, USA: Houghton Mifflin, 1971.

[9] K. Jack and W. Lionel, "Some properties of generalized sequential probability ratio tests," *Ann. Math. Stat.*, vol. 28, no. 1, pp. 57–74, 1957.

[10] T. L. Lai, "Nearly optimal sequential tests of composite hypotheses," *Ann. Stat.*, vol. 16, no. 2, pp. 856–886, 1988.

[11] G. Lorden, "Open-ended tests for Koopman-Darmois families," *Ann. Stat.*, vol. 1, no. 4, pp. 633–643, 1973.

[12] H. Chernoff, "Sequential tests for the mean of a normal distribution," in *Proc. 4th Berkeley Symp. Math. Stat. Probab.*, vol. 1, 1961, pp. 79–91.

[13] K. Jamieson, M. Malloy, R. Nowak, and S. Bubeck, "lil' UCB: An optimal exploration algorithm for multi-armed bandits," in *Proc. 27th Conf. Learn. Theory*, vol. 35, 2014, pp. 423–439.

[14] S. Zhao, E. Zhou, A. Sabharwal, and S. Ermon, "Adaptive concentration inequalities for sequential decision problems," in *Advances in Neural Information Processing Systems*. Red Hook, NY, USA: Curran, 2016, pp. 1343–1351.

[15] S. R. Howard, A. Ramdas, J. McAuliffe, and J. Sekhon, "Time-uniform, nonparametric, nonasymptotic confidence sequences," *Ann. Stat.*, vol. 49, no. 2, pp. 1055–1080, 2021.

[16] E. Kaufmann and W. Koolen, "Mixture martingales revisited with applications to sequential tests and confidence intervals," 2018, arXiv:1811.11419.

[17] A. Garivier and E. Kaufmann, "Nonasymptotic sequential tests for overlapping hypotheses applied to near-optimal arm identification in bandit models," *Sequential Anal.*, vol. 40, no. 1, pp. 61–96, 2021.

[18] O. E. Barndorff-Nielsen, *Information and Exponential Families: In Statistical Theory*. Chichester, U.K.: Wiley, 2014.

[19] S. R. Howard, A. Ramdas, J. McAuliffe, and J. Sekhon, "Time-uniform Chernoff bounds via nonnegative supermartingales," *Probab. Surveys*, vol. 17, pp. 257–317, May 2020.

[20] J. Shin, A. Ramdas, and A. Rinaldo, "On the bias, risk and consistency of sample means in multi-armed bandits," 2019, arXiv:1902.00746.

[21] R. T. Rockafellar, *Convex Analysis*, vol. 36. Princeton, NJ, USA: Princeton Univ. Press, 1997.

[22] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh, "Clustering with Bregman divergences," *J. Mach. Learn. Res.*, vol. 6, no. 58, pp. 1705–1749, 2005.

[23] T. L. Lai and M.-C. Shih, "Power, sample size and adaptation considerations in the design of group sequential clinical trials," *Biometrika*, vol. 91, no. 3, pp. 507–528, 2004.

[24] G. Lorden, "On excess over the boundary," *Ann. Math. Stat.*, vol. 41, pp. 520–527, Apr. 1970.

[25] S. P. Wong, "Asymptotically optimum properties of certain sequential tests," *Ann. Math. Stat.*, vol. 39, no. 4, pp. 1244–1263, 1968.

[26] D. A. Darling and H. Robbins, "Confidence sequences for mean, variance, and median," *Proc. Nat. Acad. Sci. United States Amer.*, vol. 58, no. 1, pp. 66–68, 1967.

[27] T. L. Lai, "On confidence sequences," *Ann. Stat.*, vol. 4, no. 2, pp. 265–280, 1976.

[28] V. H. Peña, M. J. Klass, and T. L. Lai, "Pseudo-maximization and self-normalized processes," *Probab. Surveys*, vol. 4, pp. 172–192, 2007.

[29] J. Ville, *Étude Critique de la Notion de Collectif*. Paris, France: Gauthier-Villars, 1939.

[30] R. Durrett, *Probability: Theory and Examples*, vol. 49. Cambridge, U.K.: Cambridge Univ. Press, 2019.

[31] T. Lattimore and C. Szepesvári, *Bandit Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2020.