

E-detectors: A Nonparametric Framework for Sequential Change Detection

JAEHYEOK SHIN*, AADITYA RAMDAS, AND ALESSANDRO RINALDO

Abstract

Sequential change detection is a classical problem with a variety of applications. However, the majority of prior work has been parametric, for example, focusing on exponential families. We develop a fundamentally new and general framework for sequential change detection when the pre- and post-change distributions are nonparametrically specified (and thus composite). Our procedures come with clean, nonasymptotic bounds on the average run length (frequency of false alarms). In certain nonparametric cases (like sub-Gaussian or sub-exponential), we also provide near-optimal bounds on the detection delay following a changepoint. The primary technical tool that we introduce is called an *e-detector*, which is composed of sums of e-processes—a fundamental generalization of nonnegative supermartingales—that are started at consecutive times. We first introduce simple Shiryaev-Roberts and CUSUM-style e-detectors, and then show how to design their mixtures in order to achieve both statistical and computational efficiency. Our e-detector framework can be instantiated to recover classical likelihood-based procedures for parametric problems, as well as yielding the first change detection method for many nonparametric problems. As a running example, we tackle the problem of detecting changes in the mean of a bounded random variable without i.i.d. assumptions, with an application to tracking the performance of a basketball team over multiple seasons.

KEYWORDS AND PHRASES: Sequential change detection, Nonparametric, Shiryaev-Roberts, CUSUM.

1. INTRODUCTION

Suppose we observe sequentially a stream of random variables X_1, X_2, \dots , whose marginal distributions may change at some unknown time, or *changepoint*, ν . To take one concrete example that we generalize later, denote the data stream by $X_1, \dots, X_\nu \sim P_{\mu_0}$ and $X_{\nu+1}, X_{\nu+2}, \dots \sim P_\mu$ where P_{μ_0} and P_μ are some probability distributions with parameters $\mu_0, \mu \in \mathbb{R}$, respectively. Let \mathbb{P}_ν and \mathbb{E}_ν denote probability and expectation, respectively, with respect to the distribution of the entire infinite data stream, when the change occurs at time ν . If there is no change, we think of ν as being equal to ∞ , and we let \mathbb{P}_∞ and \mathbb{E}_∞ refer to the corresponding probability and expectation.

We are concerned with designing sequential changepoint detection procedures to determine, at each time, whether a changepoint has occurred in the near past that (i) provide non-asymptotic false alarm guarantees, (ii) allow for non-parametric classes of pre- and post-change distributions and (iii) are computationally efficient. Formally, a sequential changepoint detection algorithm consists of a data-dependent stopping rule $N^* \geq 1$. If not specified explicitly, the underlying filtration with respect to which N^* is defined is assumed to be the natural filtration generated by the data stream X_1, X_2, \dots , but in some cases it is beneficial to coarsen the filtration. If the stopping time N^* is finite,

then we declare that a changepoint has been detected, in the sense that sufficient evidence has accumulated to support the hypothesis that the data generating distribution has changed. If the algorithm never stops, we set $N^* = \infty$ and no changepoint is proclaimed. To evaluate the performances of detection algorithms, we can check how quickly the algorithm can detect a change in the distribution while controlling the frequency of false alarms.

To control false alarms, we adopt the average run length (ARL) metric [23], defined by

$$\text{ARL} := \mathbb{E}_\infty N^*. \tag{1.1}$$

We say that the “ARL is controlled at level α ” if $\mathbb{E}_\infty N^* \geq 1/\alpha$. An equivalent error metric is the False Alarm Rate (FAR), which is the reciprocal of the ARL, and we would like to ensure that the FAR is at most α , and we call a sequential change detection procedure as “valid” if it satisfies the above constraint.

A widely-used measure of the speed of detection after a changepoint is the worst average delays conditioned on the least favorable observations before the change [18], or conditioned on the event that the algorithm stops after the changepoint [25]. These are defined by

$$\mathcal{J}_L(N^*) := \sup_{\nu \geq 0} \text{esssup } \mathbb{E}_\nu [[N^* - \nu]_+ | \mathcal{F}_\nu], \tag{1.2}$$

*Corresponding author.

$$\mathcal{J}_P(N^*) := \sup_{\nu \geq 0} \mathbb{E}_\nu[N^* - \nu \mid N^* > \nu], \quad (1.3)$$

where the subscripts indicate the authors, and it is known that $\mathcal{J}_P(N^*) \leq \mathcal{J}_L(N^*)$ [22].

Our implicit objective is to (approximately) minimize $\mathcal{J}_L(N^*)$ or $\mathcal{J}_P(N^*)$ while guaranteeing that the ARL is controlled at a prespecified level $\alpha \in (0, 1)$. We differ from other work in that our focus is on nonparametric and composite pre- and post-change distributions, as well as on deriving nonasymptotic guarantees. In several such settings, it is a priori unclear how to define *any* valid sequential change detection algorithm, let alone an optimal one. Accordingly, we first address the design problem, and then move to questions involving (approximate) optimality for detection delays.

1.1 Prior Work and Our Contributions

If both pre- and post-change parameters μ_0, μ are known and the distributions have densities p_{μ_0} and p_μ with respect to some common reference measure, then the CUSUM procedure [23] has been known to achieve the optimal worst average delay (exactly for $\mathcal{J}_L(N^*)$ and asymptotically for $\mathcal{J}_P(N^*)$ as $\alpha \rightarrow 0$) among all procedures controlling ARL at the same level [18, 21, 31, 14]. Recall that the CUSUM procedure is defined by the stopping time $N_{\text{CU}}^* := \inf\{n \geq 1 : M_n^{\text{CU}} \geq c_\alpha^{\text{CU}}\}$, where $c_\alpha^{\text{CU}} > 0$ is a constant chosen so that $\mathbb{E}_\infty(N_{\text{CU}}^*) = 1/\alpha$, and the test statistic M_n^{CU} is defined by the recursive formula

$$M_n^{\text{CU}} = \frac{p_\mu(X_n)}{p_{\mu_0}(X_n)} \cdot \max\{M_{n-1}^{\text{CU}}, 1\}, \quad M_0^{\text{CU}} := 0. \quad (1.4)$$

The Shiryaev–Roberts (SR) procedure [37, 32] is defined by the stopping time $N_{\text{SR}}^* := \inf\{n \geq 1 : M_n^{\text{SR}} \geq c_\alpha^{\text{SR}}\}$, where $c_\alpha^{\text{SR}} > 0$ is a constant chosen so that $\mathbb{E}_\infty(N_{\text{SR}}^*) = 1/\alpha$, and the test statistic M_n^{SR} is obtained recursively as

$$M_n^{\text{SR}} = \frac{p_\mu(X_n)}{p_{\mu_0}(X_n)} \cdot [M_{n-1}^{\text{SR}} + 1], \quad M_0^{\text{SR}} := 0. \quad (1.5)$$

Unlike CUSUM, the SR procedure does not achieve exact minimax optimality for $\mathcal{J}_L(N^*)$. However, the SR procedure and its generalized versions enjoy strong *asymptotic* optimality guarantees [26, 27, 41, 39].

The literature sometimes assumes that the pre-change distribution is known or can be approximated with a high precision by using the previous history. However, the post-change distribution is typically unknown and is assumed to belong to a family of distributions $\mathcal{P} := \{p_\mu : \mu \in \Theta\}$. In this case, one natural approach would be to replace the unknown parameter μ with an estimator $\hat{\mu}$. If we use the maximum likelihood estimator (MLE), then we obtain the CUSUM procedure based on the generalized likelihood ratio (GLR) rule [e.g., see 1, 49, 38]. For those not familiar with sequential change detection, [15] provides a good overview.

Limitations of Prior Work Though the usage of GLR statistic for the sequential change detection problem has a long history and often yields good empirical performance, the current literature has two main limitations.

First, most existing methodology relies on parametric assumptions on the family of distributions (e.g.: exponential families), both pre-change and post-change. There have been attempts to move away from this setting, and we will discuss these later. However, a general framework for deriving sequential change detection procedures in general nonparametric or composite settings has not been previously presented in the generality that we do here.

Second, the study of statistical properties has typically focused on the asymptotic regime of $\alpha \rightarrow 0$, unless the GLR statistic is defined on a well-separated post-change parameter space. In this paper, we guarantee nonasymptotic control on the ARL at a prespecified level (such as $\alpha = 0.001$). In fact, in many settings considered, we do not know of any existing method to guarantee (even asymptotic) ARL control.

(We think that in theory and practice, the first is a bigger issue than the second. Luckily, our solution for the first automatically handles the second. Indeed, in the composite and nonparametric settings we consider, it is quite unclear how to control the ARL in any asymptotic sense.)

Finally, a direct online implementation of the GLR rule is infeasible since the memory and computation time both increase at least linearly as $n \rightarrow \infty$. One natural approach to tackle this online implementation issue is to use window-limited versions of the GLR rule [49, 13]. For instance, a simple form of the window-limited GLR rule can be defined by, at each time n , computing $\hat{\mu}$ over only times $n - W$ to n for a properly chosen window size $W > 0$. However, the optimal choice of window size W has been studied only in the asymptotic setting ($\alpha \rightarrow 0$). For a fixed α , the optimal window size depends on the difference between the pre- and post-change distributions, which is unknown.

Our Contributions We present a general framework for sequential change detection, focusing on (parametric and nonparametric) settings with composite pre- and post-change distributions, and nonasymptotic guarantees on the ARL:

1. We introduce the concept of an *e-detector* that underlies our construction of sequential change detection procedures. The e-detector utilizes a generalization of the underlying martingale structure of likelihood ratios in classical sequential change detection procedures. This e-detector framework is applicable in nonparametric settings including sub-Gaussian, sub-exponential, and bounded random variables, among many others. In such settings, there is no common reference measure and likelihood ratios cannot be directly defined, thus composite nonnegative supermartingales, or more generally “e-processes”, must be employed in their place.

2. Despite handling composite pre- and post-change distributions, even without an iid (independent and identically distributed) assumption on the data, our e-CUSUM and e-SR sequential change detection procedures based on e-detectors can always nonasymptotically control the ARL at level α .
3. Nonasymptotic bounds on the worst average delay are derived in special cases for nonparametric distributions with exponential tail decay (such as sub-Gaussian or sub-exponential), and they match the rate of known lower bounds for exponential families as $\alpha \rightarrow 0$.
4. Computationally feasible algorithms are presented, in order to run our procedures in an online fashion without windowing. Practical strategies to choose hyperparameters are discussed. These are based on an adaptive mixture method, with the number of mixture components growing slowly over time.

Our procedures have natural gambling interpretations, and our work can be viewed as setting the foundations for a game-theoretic approach to sequential change detection. Before discussing the general framework in detail, in the following subsection, we present a motivating real-world example involving bounded random variables to illustrate how our nonparametric framework can be easily used in settings in which it is nontrivial to apply other common methods.

1.2 Example: A Changepoint in Cleveland Cavaliers 2011–2018

The Cleveland Cavaliers are an American professional basketball team. We use the Cavaliers' game point records over 2010–11 to 2017–18 NBA seasons to illustrate how our proposed nonparametric sequential change detection algorithm can be applied to detect an interesting changepoint in the Cavaliers' recent history.¹

The left plot in Figure 1 shows the difference between the scores of the Cavaliers and those of their opposing teams (also known as Plus-Minus) in all the games from the 2010–11 to the 2017–18 regular seasons. Each red line refers to the yearly average difference score in each season. Roughly, this value shows how well the Cavaliers performed against their opponent in terms of scoring. Typically, if a seasonal average is positive (or negative) then we may say that the Cavaliers showed a strong (or poor) performance in the corresponding season. The right plot shows a changepoint detected in early 2015; NBA fans may recall one major cause of the sharp improvement — LeBron James returned to the Cavaliers in 2014. However, how can we detect such a change on the fly by only tracking the Plus-Minus for each game?

This type of question fits well into the sequential change detection framework. Let X_1, X_2, \dots be the sequence of

Plus-Minus stats we observe sequentially. After observing a poor performance of the Cavaliers in 2010–11 season, we define the Cavaliers' pre-change distribution on the Plus-Minus stats as follows: the average Plus-Minus of the team is less than or equal to $\mu_0 := -1$. Now, we are interested in detecting a meaningful performance improvement on the fly by defining the post-change distribution as follows: the average Plus-Minus of the team is greater than $\mu_1 := 1$. Here, the gap $|\mu_1 - \mu_0|$ between averages of Plus-Minus in pre- and post-changes refers to the degree of improvement we consider as a significant one.

Although the formulation of the problem is simple as described above, it is still nontrivial to fit this problem into commonly used sequential change detection procedures for the following reasons. First, it is not easy to choose a proper parametric model to fit observed Plus-Minus stats since they are integer-valued samples with varying mean and variance over seasons as Figure 1 illustrates. Second, even if we can choose a proper model, it is difficult to find a threshold to detect the changepoint since we are interested in detecting any changes larger than $|\mu_1 - \mu_0|$ instead of a fixed post-change. Many common methods have been relying on a high-quality simulator or large enough sample history for pre-change observations to compute a valid threshold. For this example, however, it is hard to access such tools since the Cavaliers' overall Plus-Minus stats are difficult to model directly, and it is tricky to justify using existing records to get a valid threshold as the team's overall performance varies a lot around the 2010–11 season.

Based on the introduced framework of the sequential change detection procedure using e-detectors, we detour difficulties in the commonly used methods illustrated above as follows. First, due to the nonparametric nature of the new framework, we do not need to choose any parametric model to fit the data. Instead, we simply assume the absolute value of each Plus-Minus stat is bounded by a large number — in this example, we set 80 as the boundary. Though we set a conservatively large boundary, our detection procedure is variance-adaptive so that we can detect the changepoint efficiently without specifying the variance of observations. Second, the nonasymptotic analysis of the new framework makes it possible to choose an explicit detection boundary, which is equal to $\log(1/\alpha)$, to build a sequential change detection procedure controlling the ARL by $1/\alpha$ for any given $\alpha \in (0, 1)$. In this example, we choose $\alpha = 10^{-3}$ to make ARL is larger than at least 10 regular seasons of games. The right plot in Figure 1 shows the log of e-detectors on which we build the sequential change detection procedure. The red horizontal line corresponds to the detection boundary given by $\log(1/\alpha)$. We can check the procedure detects the changepoint of the Plus-Minus stat of the Cavaliers in the middle of the 2014–15 season. See Section 5.2 for the detailed explanation about how we construct the sequential change detection procedure based on the general methodology we introduce in the paper.

¹The R code to reproduce all the plots and simulation results of the paper is available at https://github.com/shinjaehyeok/e_detector_paper

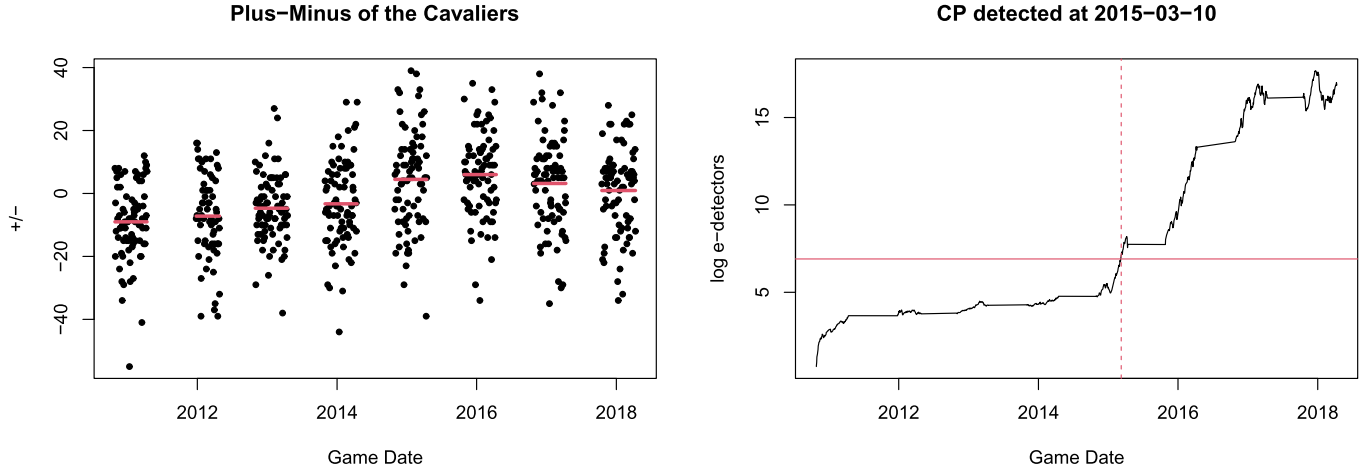


Figure 1: *Left: Plus-Minus of the Cavaliers from 2010–11 to 2017–18 seasons. Each horizontal red line corresponds to the seasonal average. Right: The sample path of (the logarithm of) one of our e-detectors. The horizontal red line is the threshold equal to $\log(1/\alpha)$ controlling the ARL by $1/\alpha = 1000$. In this example, the procedure detects the changepoint in the Plus-Minus of the Cavaliers at the end of 2014–15 season.*

Paper Outline The rest of the paper is organized as follows. In Section 2, we introduce a general framework about how to build composite, nonparametric, and nonasymptotic sequential change detection procedures using e-detectors. Section 3 extends the previous framework to the case where we have a set of e-detectors and explain how to use a mixture method to combine multiple e-detectors effectively. In Section 4, we introduce an exponential structure of e-detectors that makes it possible to design a near-optimal detection procedure with an explicit upper bound on worst average delays. Based on the proposed framework, Section 5 presents two canonical examples of Bernoulli (parametric) and bounded random variables (nonparametric) cases with real data applications of the Cavaliers 2011–2018 statistics. We conclude with a discussion, and defer proofs to the supplement.

2. NONPARAMETRIC SEQUENTIAL CHANGE DETECTION USING E-DETECTORS

2.1 Problem Setup

Let \mathcal{P} denote the set of possible pre-change distributions, which could in general be a nonparametric class. We do not assume the observations in the sequence to be independent or identically distributed: each $P \in \mathcal{P}$ is a distribution over an infinite sequence of random variables.

We will assume throughout that up to the unknown changepoint ν , the observations X_1, \dots, X_ν follow a distribution $P \in \mathcal{P}$. The remaining observations $X_{\nu+1}, X_{\nu+2}, \dots$ are drawn from a distribution Q in a class of post-change distributions \mathcal{Q} . In this case, we let $\mathbb{P}_{P,\nu,Q}$, $\mathbb{E}_{P,\nu,Q}$ and $\mathbb{V}_{P,\nu,Q}$ denote probability, expectation and variance operators over the entire data stream.

If there never is a change, we will use the notation $\mathbb{P}_{P,\infty}$, $\mathbb{E}_{P,\infty}$ and $\mathbb{V}_{P,\infty}$. Also, if a change occurs at the beginning ($\nu = 0$) then we use $\mathbb{P}_{0,Q}$, $\mathbb{E}_{0,Q}$ and $\mathbb{V}_{0,Q}$. Note that technically $\mathbb{P}_{0,Q} = \mathbb{P}_{Q,\infty}$, but we use the former to denote that Q is a post-change distribution and a changepoint has occurred at the very start, and the latter to denote that Q is a pre-change distribution and a changepoint never occurs.

Let $\mathcal{F} := \{\mathcal{F}_n\}_{n \geq 0}$ be a filtration where we let $\mathcal{F}_0 := \{\emptyset, \Omega\}$ for simplicity. Let $M := \{M_n\}_{n \geq 0}$ be a nonnegative adapted process with respect to the filtration \mathcal{F} . If required, we define $M_\infty := \limsup_{n \rightarrow \infty} M_n$ and $\mathcal{F}_\infty := \sigma(\bigcup_{n \geq 0} \mathcal{F}_n)$ [see, for example, 4]. It is common to consider the natural filtration $\mathcal{F}_n := \sigma(X_1, \dots, X_n)$, but there are situations where restricting the filtration could be advantageous (for example, when there are nuisance parameters). There are yet other situations when enlarging the filtration with external randomness can be useful.

Let \mathcal{T} denote the set of all stopping times with respect to \mathcal{F} , but we will later see that it typically suffices to consider finite stopping times, or those with finite expectation.

Remark 2.1. In our paper, the changepoint ν and the post-change observations also do not need to be independent of the pre-change observations, and they do not need to be identically distributed. In other words, ν could be a stopping time, and Q could itself depend on the pre-change data. It may be helpful to imagine an adversary who adaptively decides at each step whether or not to “stop” the pre-change data. If they choose to stop at time ν , there are two options for post-change points: (a) then can pick a distribution Q , draw a sequence Y_1, Y_2, \dots , from Q , and reveal $X_{\nu+i} = Y_i$ sequentially, or (b) the pick a distribution Q and draw $X_{\nu+1}, X_{\nu+2}, \dots$ from $Q \mid \mathcal{F}_\nu$. In situations without a common reference measure, setting (b) may be

tricky to formally define (especially if the pre-change data have probability zero under Q), so one may think of setting (a). All of our results on ARL do not require any assumptions on the changepoint ν or post-change distribution of the data. When analyzing detection delay, we will typically assume that $X_{\nu+1}, X_{\nu+2}, \dots$ are independent of the pre-change data, and are drawn from Q as if the time ν was reset to zero; we believe this can be relaxed in future work. But for much of the paper, it is also okay to assume that the post-change data is drawn a data-dependent Q (or drawn from $Q | \mathcal{F}_\nu$).

With the appropriate definitions and setup in place, we can now define our central concept, an e-detector.

2.2 What Is an E-detector?

Definition 2.2 (\mathcal{P} -e-detector). The process M is called an *e-detector* with respect to the class of pre-change distributions \mathcal{P} if it satisfies the property

$$\mathbb{E}_{P,\infty}[M_\tau] \leq \mathbb{E}_{P,\infty}[\tau], \quad \forall \tau \in \mathcal{T}, \quad \forall P \in \mathcal{P}. \quad (2.1)$$

For brevity, we refer to M as “an e-detector for \mathcal{P} ”, or if \mathcal{P} is understood from context, then simply “an e-detector”. If a stopping time τ has a nonzero probability of being infinite, then inequality (2.1) is trivially satisfied. Thus, the condition is only really required to hold for stopping times with finite expectation under some P . This latter set of stopping times depends on \mathcal{P} , and so in order to not complicate notation, we continue to simply consider all stopping times \mathcal{T} . (Also note that the condition can only be satisfied if process M is integrable under any $P \in \mathcal{P}$, so this is implicitly assumed to be the case in what follows.)

By the linearity of expectation and Tonelli’s theorem, an average (or “a mixture”) of e-detectors is also an e-detector. More formally, if $\{M^a\}_{a \in A}$ is a set of e-detectors (where a is a tuning parameter), then so is $\int M^a d\mu(a)$ for any fixed probability distribution μ over A . Later in this paper, we will in particular use finite mixtures of the form $(M^1 + M^2 + \dots + M^K)/K$ in order to adapt to the unknown post-change distribution. (In fact, we will develop more sophisticated mixtures whose number of components grows slowly with time.) For later reference, we state the above as a proposition:

Proposition 2.3. *Let $\{M^a\}_{a \in A}$ be a set of e-detectors. Then for any probability measure μ on A , the mixture of e-detectors, $\int M^a d\mu(a)$ forms a valid e-detector.*

An e-detector provides a quantification of evidence for whether a changepoint has occurred or not, and may be continuously monitored, stopped and easily interpreted – e.g, a steep and steady increase of the process in recent times should be taken as an indication that a change has taken place. The following theorem shows how one can immediately obtain a sequential change detection procedure from an e-detector M .

Theorem 2.4. *For any $\alpha \in (0, 1)$ and e-detector M , if we declare a changepoint at the stopping time*

$$N^* := \inf\{n \geq 1 : M_n \geq 1/\alpha\}, \quad (2.2)$$

then we have

$$\inf_{P \in \mathcal{P}} \mathbb{E}_{P,\infty} N^* \geq 1/\alpha. \quad (2.3)$$

That is, the sequential change detection procedure in (2.2) controls the ARL at level α .

The informal proof is one line long: dropping subscripts, the definition of an e-detector implies that $\mathbb{E}N^* \geq \mathbb{E}M_{N^*}$, but by definition of N^* , we know that $M_{N^*} \geq 1/\alpha$ if N^* is finite. (If N^* is not almost surely finite, the claim holds anyway.) The full proof is in Appendix A.1.

2.3 Constructing an E-detector Based on a Sequence of E-processes

The central building block of our e-detector is called an e-process. E-processes are newly-developed tools that have been shown to play a fundamental role in sequential hypothesis testing, especially in composite, nonparametric settings. E-processes are generalizations of nonnegative martingales and supermartingales, and in particular, *e-processes are nonparametric and composite generalizations of likelihood ratios*. They have strong game-theoretic roots, and have found utility in the meta-analysis, as well as for the purposes of anytime-valid inference in the presence of continuous monitoring [28, 29, 7, 11, 12]. The properties of e-processes have not yet been explored in changepoint analysis, and we undertake this effort here.

To understand their definition, we briefly forget about sequential change detection and consider testing the null hypothesis that $X_1, X_2, \dots \sim P$ for some $P \in \mathcal{P}$. An e-process for \mathcal{P} , called a \mathcal{P} -e-process, is a sequence of nonnegative random variables $(E_t)_{t \geq 1}$ such that for any $P \in \mathcal{P}$ and any stopping time τ , we have $\mathbb{E}_P[E_\tau] \leq 1$. As before, the underlying filtration can be that of the data or that of (E_t) , or some enlargement of these. The value of E_t measures evidence against the null (larger values, more evidence). A level- α sequential test can be obtained by rejecting the null as soon as E_t exceeds $1/\alpha$; this is a consequence of Ville’s inequality [42, 11].

As a result of the optional stopping theorem, nonnegative \mathcal{P} -martingales (i.e., the process is a nonnegative P -martingales simultaneously for every $P \in \mathcal{P}$) and \mathcal{P} -supermartingales are examples of e-processes. However, e-processes are a distinct and more general class of processes. In fact, there exist natural classes \mathcal{P} for which the only \mathcal{P} -martingales are constants, and the only \mathcal{P} -supermartingales are decreasing sequences, but there are e-processes for \mathcal{P} that can increase to infinity when the data are not from \mathcal{P} . See [29] for one such example, arising from sequentially testing exchangeability of a binary sequence, and [6] for another example arising from testing log-concavity.

In this subsection, we show how to leverage e-processes to build e-detectors.

Definition 2.5 (*e_j-process*). For any $j \geq 1$, $\Lambda^{(j)} := \{\Lambda_n^{(j)}\}_{n \geq 1}$ is called an e_j -process for \mathcal{P} if it is a nonnegative adapted process such that $\Lambda_1^{(j)} = \dots = \Lambda_{j-1}^{(j)} = 1$, and

$$\sup_{P \in \mathcal{P}} \mathbb{E}_{P, \infty} [\Lambda_\tau^{(j)} \mid \mathcal{F}_{j-1}] \leq 1, \quad \forall \tau \in \mathcal{T}. \tag{2.4}$$

When $j = 1$, the e_j -process is simply a standard e-process, which tests whether the data distribution is different from the proclaimed null (pre-change) distribution. For $j > 1$, each e_j -process can be viewed as an e-process that begins at time j , and tests whether the data occurring after time j are well explained by the null hypothesis.

Definition 2.6 (*SR and CUSUM e-detectors*). Based on a sequence of e-processes $\{\Lambda^{(j)}\}_{j \geq 1}$, define SR and CUSUM e-detectors M^{SR} and M^{CU} , respectively by $M_0^{\text{SR}} = M_0^{\text{CU}} := 0$ and for each $n \geq 1$,

$$M_n^{\text{SR}} := \sum_{j=1}^n \Lambda_n^{(j)}, \quad \text{and} \quad M_n^{\text{CU}} := \max_{j \in [n]} \Lambda_n^{(j)}. \tag{2.5}$$

It is not hard to check that the above processes are indeed e-detectors, meaning they satisfy (2.1).

Remark 2.7. If each e-process starts with an initial weight smaller than 1 such that the sum of all initial weights is less than or equal to 1, then corresponding SR and CUSUM procedures control the probability of the false alarm, $\sup_{P \in \mathcal{P}} \mathbb{P}_{P, \infty}(\tau < \infty) \leq \alpha$, which is a much more stringent error metric than the ARL. The price to pay is that if there is a change at time ν , then the detection delay will no longer be independent of ν , and will typically increase logarithmically with ν (so that the worst average delay is unbounded).

2.4 Constructing Computationally Efficient E-detectors Using Baseline Increments

In general, it may take $O(n)$ time to update the aforementioned e-detectors at time n . In order to construct e-detectors that can be updated online in sublinear time and memory (or even near-constant time and memory), it turns out to be computationally convenient to use a common “baseline” increment in order to build the underlying e_j -processes, as we do below. Effectively, this amounts to using e_j -processes that are \mathcal{P} -supermartingales, which is a special case of particular interest.

Definition 2.8 (*Baseline increment*). A nonnegative, adapted process $L := \{L_n\}_{n \geq 1}$ is called a baseline increment if for each $n \geq 1$, we have

$$\sup_{P \in \mathcal{P}} \mathbb{E}_{P, \infty} [L_n \mid \mathcal{F}_{n-1}] \leq 1. \tag{2.6}$$

It is easy to check that if L^1 and L^2 are baseline increments, and A^1 and A^2 are nonnegative and predictable processes (meaning that A_n^1 and A_n^2 are both \mathcal{F}_{n-1} -measurable) such that $A^1 + A^2$ is strictly positive, then the mixture $(A^1 L^1 + A^2 L^2) / (A^1 + A^2)$ also forms a baseline increment. In short, “predictable mixtures” retain the baseline increment property.

Comparing (2.6) to (2.4), we see that a baseline increment L is not itself an e_j -process, because the expectation in (2.6) applies only at fixed times n , with the conditioning being on the previous step $n - 1$, but (2.4) calculates expectations at any stopping time, and conditions on $j - 1$. It is best to think about the baseline increment as the multiplicative increment that forms the e_j -process, as follows.

Definition 2.9 (*Baseline e_j -process*). For a given baseline increment $L := \{L_n\}_{n \geq 1}$, we define the corresponding “baseline e_j -process” $\Lambda^{(j)}$, for each $j, n \in \mathbb{N}$, as below:

$$\Lambda_n^{(j)} := \begin{cases} 1 & \text{if } n < j \\ \prod_{i=j}^n L_i & \text{otherwise.} \end{cases} \tag{2.7}$$

Under any pre-change distribution $P \in \mathcal{P}$, each $\Lambda^{(j)}$ is a nonnegative supermartingale by definition of the baseline increment L_i . Therefore, a straightforward application of the optional stopping theorem implies that each $\Lambda^{(j)}$ satisfies condition (2.4), and thus is a valid e_j -process.

As an example of a baseline e_j -process, consider the case where we have iid observations X_1, X_2, \dots from a distribution p_θ parameterized by $\theta \in \Theta$, and the pre-change distribution is given by θ_0 . Then, for any post-change distribution p_{θ_1} with $\theta_1 \neq \theta_0$, the likelihood ratio between two distributions, $L_n := p_{\theta_1}(X_n) / p_{\theta_0}(X_n)$ yields a baseline increment process with the inequality in (2.6) being replaced by the equality. Then, each $\Lambda_n^{(j)}$ is the likelihood ratio based on X_j, \dots, X_n . Further, instead of using a fixed post-change parameter θ_1 , we can also plug-in a running MLE or any other online nonanticipating estimator that is based on the previous history \mathcal{F}_{n-1} only, say $\hat{\theta}_{n-1}$, into the likelihood ratio. In this case, although the value $L_n := p_{\hat{\theta}_{n-1}}(X_n) / p_{\theta_0}(X_n)$ at time n of the resulting process may depend on the previous history \mathcal{F}_{n-1} , the inequality in (2.7) will be satisfied as an equality, yielding a valid baseline e_j -processes.

Remark 2.10. While baseline increment processes provide a natural and computationally convenient way to construct e_j -process, we emphasize that any e-detector, even one that does not use baseline increments, will automatically control the ARL by Theorem 2.4. To elaborate, baseline e_j -processes are composite \mathcal{P} -supermartingales (meaning P -supermartingales for every $P \in \mathcal{P}$), but there exist other \mathcal{P} -e-processes—which are not \mathcal{P} -supermartingales—that naturally arise and these can be used to form e-detectors; for example using universal inference [47, Section 8].

Definition 2.11 (*Baseline SR and CUSUM e-detectors*). When an SR or CUSUM e-detector is constructed using a

sequence of baseline e_j-processes (2.7), we call it a “baseline SR or CUSUM e-detector”.

Each baseline SR or CUSUM e-detector can be computed recursively like their classical analogs:

$$M_n^{\text{SR}} = L_n \cdot [M_{n-1}^{\text{SR}} + 1], \tag{2.8}$$

$$M_n^{\text{CU}} = L_n \cdot \max\{M_{n-1}^{\text{CU}}, 1\}, \tag{2.9}$$

with $M_0^{\text{SR}} = M_0^{\text{CU}} = 0$ for each $n \in \mathbb{N}$. The above computational benefit is the primary reason to consider baseline e-detectors, but as mentioned in Remark 2.10 and when introducing e-processes, more general e-detectors are sometimes necessary for certain classes \mathcal{P} .

We briefly verify below that the processes M^{SR} and M^{CU} defined above are valid e-detectors. Indeed, for any stopping time τ and pre-change distribution $P \in \mathcal{P}$, if $\mathbb{P}_{P,\infty}(\tau = \infty) > 0$ then the condition of the e-detector in (2.1) holds trivially. If not, then τ is finite almost surely, and we have by linearity of expectation and the tower rule:

$$\begin{aligned} \mathbb{E}_{P,\infty} M_\tau^{\text{CU}} &\leq \mathbb{E}_{P,\infty} M_\tau^{\text{SR}} \\ &= \mathbb{E}_{P,\infty} \sum_{j=1}^{\infty} \Lambda_\tau^{(j)} \mathbb{1}(j \leq \tau) \\ &= \sum_{j=1}^{\infty} \mathbb{E}_{P,\infty} \Lambda_\tau^{(j)} \mathbb{1}(j \leq \tau) \\ &= \sum_{j=1}^{\infty} \mathbb{E}_{P,\infty} [\mathbb{1}(j \leq \tau) \mathbb{E}_{P,\infty} [\Lambda_\tau^{(j)} \mid \mathcal{F}_{j-1}]] \\ &\leq \sum_{j=1}^{\infty} \mathbb{E}_{P,\infty} \mathbb{1}(j \leq \tau) = \mathbb{E}_{P,\infty} \tau, \end{aligned}$$

where the first inequality comes from the nonnegativity of e-processes, and the second inequality comes from the definition of the e-process $\Lambda^{(j)}$ for each $j \geq 1$. Note that this proof is also applicable to general SR and CUSUM e-detectors.

2.5 Sequential Change Detection Procedures by Thresholding E-detectors

The value of any e-detector process, like M^{SR} or M^{CU} , is directly interpretable without specifying an explicit threshold: a larger value signals an accumulation of evidence of a changepoint. These can be monitored and adaptively stopped. Nevertheless, to explicitly control the ARL at level α , we define SR and CUSUM-style change detection procedures, called “e-SR” and “e-CUSUM” procedures as follows.

Definition 2.12 (e-SR and e-CUSUM procedures). Given SR and CUSUM e-detectors M^{SR} and M^{CU} , define e-SR and e-CUSUM procedures by the stopping times

$$N_{\text{SR}}^* := \inf\{n \geq 1 : M_n^{\text{SR}} \geq 1/\alpha\}, \tag{2.10}$$

$$N_{\text{CU}}^* := \inf\{n \geq 1 : M_n^{\text{CU}} \geq c_\alpha\}, \tag{2.11}$$

where c_α is a constant chosen to control the ARL of the e-CUSUM procedure by $1/\alpha$ for some $\alpha \in (0, 1)$. By Theorem 2.4, $1/\alpha$ is a valid choice for c_α .

We note $c_\alpha = 1/\alpha$ may be a very conservative choice for the e-CUSUM procedure. Indeed, suppose we use the trivial e-processes, that is, we set $\Lambda_n^{(j)} := 1$ for all j, n . In this setting, the SR e-detector is given by $M_n^{\text{SR}} = n$ for each n . In contrast, the CUSUM e-detector, M_n^{CU} is equal to 1 for all n . Therefore, any valid threshold we can choose for the e-SR procedure must be larger than $\lceil 1/\alpha \rceil$. On the other hand, any threshold above 1 makes $N_{\text{CU}}^* = \infty$, which of course controls ARL by $1/\alpha$, but the true ARL is much above the target. Building from this trivial example, it is possible to construct nontrivial examples in which letting $\alpha \rightarrow 0$ makes the gap between tight thresholds of e-SR and e-CUSUM procedures arbitrarily different.

Remark 2.13. Unless we assume the pre-change distribution is time-stationary, known and parametric, or we can access a good sample of the pre-change distribution or large enough historical data, computing a tight or even valid threshold c_α can be a challenging task. In the application sections below, we will mainly deal with non-stationary pre-change distributions where pre-change observations may not be identically distributed and thus all observations before the changepoint may follow different distributions. In this case, setting $c_\alpha = 1/\alpha$ seems to be the only reasonable choice, and we recommend using the e-SR procedure rather than e-CUSUM since if we use the same threshold for both procedures, the former always detects the changepoint faster than the latter while provably controlling the ARL at the same level.

2.6 Some Nontrivial Instantiations of E-detectors

E-detectors can be thought of as a general reduction of change detection to sequential testing. Given the recent advances in nonparametric, composite sequential testing using nonnegative supermartingales and e-processes, our e-detectors now make new classes of change detection problems possible. We detail below some interesting nontrivial examples of change detection problems that can now be solved using e-detectors.

Example 1 (When Likelihood Ratios Are Well-defined). Consider first the parametric case, when \mathcal{P}, \mathcal{Q} have a common reference measure and likelihood ratios are well-defined. When the pre- and post-change distributions are known (meaning \mathcal{P}, \mathcal{Q} are singletons), then the standard likelihood-ratio based CUSUM and SR processes from (1.4) and (1.5) are both e-detectors. If \mathcal{Q} is composite, then taking a mixture likelihood ratio yields e-detectors (using either a non-anticipating predictable mixture or using a fixed mixture distribution). If \mathcal{P} is also composite, but maximum likelihood estimation is efficient over \mathcal{P} , then e-processes can be constructed that take the ratio of mixture likelihoods over

the alternative to maximum likelihood under the null, as done in universal inference [47, Section 8] or in [40]. If the “reverse information projection” (RIPr) is computable (analytically or numerically), then one can use the method of [7], though there are some subtleties: by default the method produces e-values over blocks of observations, which can be multiplied across independent blocks to produce a nonnegative supermartingale (and thus e-process) to be used within our framework. But sometimes, the sequence of RIPr’s over increasing sample sizes (nested blocks) automatically produces an e-process, and when this happens, it is more powerful than universal inference (see [7, 30] for details).

Example 2 (Change in Distribution). In this example, \mathcal{P} is the set of all iid product distributions over infinite sequences (or its convex closure, the set of all exchangeable distributions), so $\mathcal{P} = \{\mu^\infty \text{ for some probability distribution } \mu\}$. The conformal sequential change detection procedures by [44, 43] are designed to test deviations from exchangeability, meaning that they develop a test martingale (and thus an e-process) for \mathcal{P} , meaning that their procedure fits neatly into our framework. Importantly, their filtration is restricted, and is smaller than the natural filtration of the data. This allows nonparametric martingales to exist, but the e-detector property only holds with respect to a smaller class of stopping times. Nevertheless, thresholding our e-detector at $1/\alpha$ still controls the ARL at level α , as the latter property is independent of the filtration used to construct the e-detector. As a side note, if one wanted to construct an e-detector that was valid at *all* stopping times with respect to the natural filtration of the data, e-detectors based on martingales provably do not suffice, but e-detectors based on e-processes can be constructed using the techniques from [29] (at least for categorical distributions).

Example 3 (Nonparametric Two-sample Testing). Suppose we have two streams of (general multivariate) data: $X_1, X_2, \dots \sim P_X$ and $Y_1, Y_2, \dots \sim P_Y$. For simplicity below, assume that at time t , we observe one point from each stream (X_t, Y_t) . Before the changepoint (if one exists), the distributions of X_t and Y_t are equal, meaning that $\mathcal{P} = \{(P_X, P_Y)^\infty : P_X = P_Y\}$. This is a very nonparametric class, since it specifies nothing about the distributions except for the fact that they are equal before the changepoint. After the changepoint, the streams have different distributions (maybe the distribution of X changes, or that of Y changes, or both), thus $\mathcal{Q} = \{(P_X, P_Y)^\infty : P_X \neq P_Y\}$. For this very general nonparametric two-sample testing setup, [34] construct test martingales for \mathcal{P} that are provably consistent against \mathcal{Q} under minimal assumptions (in particular not requiring any minimum separation between the different distributions after the changepoint, since the tests automatically adapt to the closeness of the unknown alternative). These test martingales fit seamlessly into our e-detector framework, yielding new and practicable e-detectors for detecting a change from homogeneity to non-homogeneity between the streams.

Example 4 (Nonparametric Independence Testing). In this problem setting, we observe a pair of random variables $(X_t, Y_t) \sim P_{XY}$ at each step t , where X_t, Y_t can each lie in a general space. Before the changepoint (if one exists), the data are independent, meaning that $\mathcal{P} = \{P_{XY}^\infty : P_{XY} = P_X \times P_Y\}$. Beyond saying that the joint distribution factorizes into the product of marginals, there is no further structure assumed, making this a rich nonparametric composite class. As before, after the changepoint, (X, Y) become dependent, meaning that $\mathcal{Q} = \{P_{XY}^\infty : P_{XY} \neq P_X \times P_Y\}$. For this general nonparametric independence testing problem, [24] construct test martingales for \mathcal{P} that are provably consistent against \mathcal{Q} under minimal, weak assumptions (as before, not requiring any separation). Again as before, (in the testing problem) the power of these tests automatically adapts to the difficulty of the unknown alternative. When plugged into our framework, it delivers a novel e-detector for a change from independence to dependence.

We briefly remark that in the two preceding examples (homogeneity and independence), we can move past the iid assumption. The same methods work even when the distribution is allowed to drift within \mathcal{P} before the changepoint, and drift within \mathcal{Q} after the changepoint. We refer to the original aforementioned papers for details.

Example 5 (Log-concavity). Here, the data before the changepoint comes from a log-concave distribution (in a general dimension $d \geq 1$), so $\mathcal{P} = \{\mu^\infty : \mu \text{ has a log-concave Lebesgue density}\}$. This is a rich, nonparametric, shape-constrained class. The post-change class of distributions \mathcal{Q} consists of, for example, any distribution that has a nonzero KL-divergence and Hellinger distance from every distribution in \mathcal{P} . For testing \mathcal{P} against \mathcal{Q} , [6] show that there exists no nontrivial nonnegative supermartingales, but they design a powerful (and computationally efficient) e-process using universal inference. When plugged into our e-detector, this yields a nontrivial procedure that can detect a deviation from log-concavity.

Example 6 (Symmetry). Suppose $\mathcal{P} = \{\mu^\infty : \mu \text{ is symmetric around } 0\}$ consists of the set of all distributions (in a general dimension $d \geq 1$) that are symmetric around the origin, while \mathcal{Q} consists its complement (that is, distributions which are not symmetric around the origin). [28] characterize all processes that are nonnegative martingales for \mathcal{P} . When used with our e-detector, these provide a clean way to detect a change from symmetry to asymmetry.

Example 7 (Change in Mean). Suppose $\mathcal{P}^C = \{\mu^\infty : \mathbb{E}_{X \sim \mu}[X] \leq 0, X \text{ satisfies } C\}$ consists of the set of all univariate distributions with mean less than or equal to zero satisfying some constraint C , while $\mathcal{Q}^C = \{\mu^\infty : \mathbb{E}_{X \sim \mu}[X] > 0, X \text{ satisfies } C\}$ consists of those with positive mean. [11] provides a large variety of nonnegative supermartingales under various conditions C , such as when X are subGaussian, or bounded from above, or bounded from below, or have

only two or three moments; see also [45] for heavy-tailed supermartingales. These can be plugged into the e-detector to yield new nonparametric schemes for changes in mean.

Example 8 (Huber-robust Change Detection). As a final example, suppose we wish to detect a change in mean of heavy-tailed data (as above). But now, suppose that an adversary can also arbitrarily corrupt an ϵ fraction of the data. [46] develop Huber-robust supermartingales for this setting, which can be plugged into e-detectors to yield a valid e-detector in the presence of adversarial corruptions.

Note that in Examples 3, 4, 5 and 6, if \mathcal{P} and \mathcal{Q} were swapped, the testing problem is much harder, and we are not aware of any powerful test or change detection method. However, if one was simply interested in detecting a change *in homogeneity* or *in dependence*, i.e. from some distribution in $\mathcal{P} \cup \mathcal{Q}$ to some other one, there are two possible change detection methods that come to mind. First, an e-detector based on the conformal change detection methods in Example 2 would detect any change from any distribution to any other, though choosing the conformity score may be tricky. As a second and more direct option, one can choose a measure of homogeneity or dependence (like the kernel maximum mean discrepancy or energy distance, or the Hilbert Schmidt independence criterion or distance covariance), construct a confidence sequence for that measure (see [19] for several specific, tight, constructions), and plug it into the recent change detection scheme of [35].

Finally, it is worth noting that it is possible to define e-detectors in cases where there is no common reference measure amongst the pre-change and post-change distributions, and thus no easily-defined likelihood ratio process, and also when there is no nontrivial martingale that can be constructed. This is precisely the utility of e-processes, which are nonparametric and composite generalizations of likelihood ratios. We develop one more interesting nonparametric example (not described above) in the simulations section: detecting change in mean of a bounded random variable.

2.7 Warm-up: Bounds on Worst Average Delays for Baseline E-detectors, when Q is Known

Recall that our objective is to minimize worst average delays $\mathcal{J}_L(N^*)$ or $\mathcal{J}_P(N^*)$, given above in (1.2) and (1.3) respectively, while controlling the ARL $\mathbb{E}_{P,\infty}(N^*) \geq 1/\alpha$. Note that worst average delays in (1.2) and (1.3) were defined for a fixed pair of pre- and post-change distributions implicitly. In our setting where the pre-change distribution space could be composite, we take an additional supremum over all pre-change distributions when defining both worst average delays for each fixed post-change distribution.

To derive bounds on the worst average delays, we further assume that the post-change observations $X_{\nu+1}, X_{\nu+2}, \dots$ are independent of the pre-change observations and form a

strongly stationary process. That is, we assume that, for any finite subset $I \subset \mathbb{N}$ and any $j \in \mathbb{N}$, the joint distributions of $\{X_{\nu+i}\}_{i \in I}$ and $\{X_{\nu+i+j}\}_{i \in I}$ are equal to each other. About the underlying baseline increment, we further assume that there exist a function f and an integer $m \geq 0$ such that $L_n = f(X_n, X_{n-1}, \dots, X_{n-m})$ for each n . In this warm-up section, assume that we know the post-change distribution Q . Then, as we shall soon see, an optimal choice for L_n would set $m = 0$, but if m is a strictly positive number then we implicitly assume that we can access m observations $X_0, X_{-1}, \dots, X_{1-m}$ from the pre-change distribution in order to build sequential change detection procedures. Under these conditions, the following theorem provides analytically more tractable upper bounds on worst average delays for e-SR and e-CUSUM procedures.

Proposition 2.14. *For a given $\alpha \in (0, 1)$, let N_{SR}^* and N_{CU}^* be e-SR and e-CUSUM procedures using baseline e-detectors. Under the settings described above, their worst average delays are upper bounded as*

$$\mathcal{J}_P(N_{SR}^*) \leq \mathcal{J}_L(N_{SR}^*) \leq \mathbb{E}_{0,Q} N_{1/\alpha} + m, \tag{2.12}$$

$$\mathcal{J}_P(N_{CU}^*) \leq \mathcal{J}_L(N_{CU}^*) \leq \mathbb{E}_{0,Q} N_{c_\alpha} + m, \tag{2.13}$$

respectively, where N_c is the stopping time defined for any $c > 1$ as $N_c := \inf\{n \geq 1 : \sum_{i=1}^n \log L_i \geq \log(c)\}$, and $c_\alpha \leq 1/\alpha$ is any threshold that ensures the e-CUSUM procedure (2.11) has an ARL no smaller than $1/\alpha$. Furthermore, if the post-change observations are iid and each L_n is a function of X_n only (i.e. $m = 0$) with $\mathbb{E}_{0,Q} \log L_1 > 0$, then

$$\mathbb{E}_{0,Q} N_c \leq \frac{\log(c)}{\mathbb{E}_{0,Q} \log L_1} + \frac{\mathbb{V}_{0,Q} \log L_1}{[\mathbb{E}_{0,Q} \log L_1]^2} + 1. \tag{2.14}$$

The proof can be found in Appendix A.1.

Remark 2.15. The stopping time $N_{1/\alpha}$ delivers a level- α sequential test for the null hypothesis $H_0 : P \in \mathcal{P}$. On the other hand, the stopping time N_{c_α} may not necessarily control the type-1 error by α since the threshold c_α can be significantly smaller than $1/\alpha$ as discussed earlier.

The expected stopping time in the upper bounds (2.12) and (2.13) depends on the parameter $m \geq 0$. When Q is known, $m = 0$ suffices because (2.14) suggests that L_i should simply be chosen to maximize $\mathbb{E}_{0,Q} \log L_1$, which is identical to the log-optimality criterion used for testing \mathcal{P} against Q (as discussed in many recent works, like [33, 7, 48]).

In applications when Q is unknown, the underlying baseline increment may require a long enough sample history (large m) in order to achieve a reasonably small expected stopping time by “learning” Q or using an empirical distribution plug-in for Q . Then, the above results suggest that a reasonable way to choose the window size m is to pick the one minimizing the upper bound on the worst average delays. However, since the optimal choice of the window size should depend on the unknown post-change Q , it remains

difficult to minimize the upper bound directly. In our simulations, we often encounter cases where a larger window size is better. In this case, we would choose a window size as large as possible while keeping the procedure computationally tractable. However, the right way to handle unknown Q is dealt with in detail next.

3. COMBINING BASELINE E-DETECTORS USING THE METHOD OF MIXTURES

In the previous section, we discussed how one can construct a valid e-detector and derive upper bounds on worst average delays. However, in most composite and nonparametric sequential change detection scenarios, there is no single optimal e-detector but instead there are often several applicable e-detectors to choose from. In this section, we introduce a practicable and computationally efficient strategy to construct a good e-detector for minimizing the upper bound on worst average delays in Proposition 2.14, especially for the upper bound (2.14) in the $m = 0$ case.

In detail, suppose we have a set of baseline increments $\{L^\lambda\}_{\lambda \in \Pi}$ parametrized by $\lambda \in \Pi$. Then, under the additional condition assumed in the second part of Proposition 2.14 (namely, that $m = 0$ and the post-change observations are from an iid sequence), an ideal choice of the parameter λ^{op} for a post-change distribution Q is given by

$$\lambda^{\text{op}}(Q) = \arg \max_{\lambda \in \Pi} \mathbb{E}_{0,Q} \log L_1^{(\lambda)}, \tag{3.1}$$

minimizing the first term of the upper bound in (2.14), which often becomes a leading term especially for small enough α . In turn, this term is inversely proportional to

$$\mathbb{E}_{0,Q} \log L_1^{(\lambda^{\text{op}})} := D(Q||\mathcal{P}), \tag{3.2}$$

where the second argument \mathcal{P} in $D(Q||\mathcal{P})$ explicitly refers to the dependency of the baseline increment $L^{\lambda^{\text{op}}}$ to the class of pre-change distributions \mathcal{P} . For the rest of the paper, we will assume that the set of baseline increments, $\{L^\lambda\}_{\lambda \in \Pi}$ is rich enough such that $D(Q||\mathcal{P}) > 0$ for all $Q \in \mathcal{Q}$. As we observe later, in many canonical cases, we have $D(Q||\mathcal{P}) = \inf_{P \in \mathcal{P}} \text{KL}(Q||P)$ where $\text{KL}(Q||P)$ is the Kullback-Leibler (KL) divergence from Q to P .

Generally, computing λ^{op} is not feasible since it depends on the unknown post-change distribution Q . Next, we show how to build a mixture of baseline e-detectors that can detect the changepoint nearly as quickly as the one with λ^{op} , when known lower and upper bounds λ_L and λ_U on λ^{op} are available.

Notice that an average of e-detectors is also a valid e-detector, in the sense of satisfying condition (2.1). Therefore, for any mixing distribution W supported on $[\lambda_L, \lambda_U]$, we can define mixtures of e-SR and e-CUSUM procedures by following stopping times:

$$N_{\text{mSR}}^* := \inf \left\{ n \geq 1 : \int \sum_{j=1}^n \prod_{i=j}^n L_i^{(\lambda)} dW(\lambda) \geq 1/\alpha \right\}, \tag{3.3}$$

$$N_{\text{mCU}}^* := \inf \left\{ n \geq 1 : \int \max_{j \in [n]} \prod_{i=j}^n L_i^{(\lambda)} dW(\lambda) \geq c_\alpha \right\}, \tag{3.4}$$

where $c_\alpha > 1$ is a fixed constant which controls the ARL for some $\alpha \in (0, 1)$. Since the mixture of e-CUSUM procedure is based on a valid e-detector, we can always set the threshold c_α to be equal to $1/\alpha$ as same as the threshold of the mixture of e-SR procedures.

Remark 3.1. Instead of using mixtures, one may be tempted to consider swapping the above integral with a supremum over $\lambda \in [\lambda_L, \lambda_U]$. However, this does not in general yield a valid e-detector.

3.1 Computational and Analytical Aspects of Mixtures of Baseline E-detectors

Though any mixing distribution yields a valid e-detector, for computational efficiency, we only consider discrete mixtures where the support of mixing distribution has at most countably many elements. To be specific, let $\{\omega_k\}_{k \geq 1}$ be a set of nonnegative mixing weights with $\sum_{k \geq 1} \omega_k = 1$ and let $\{\lambda_k\}_{k \geq 1}$ be the corresponding supporting set. For ease of notation, we denote $L^{(\lambda_k)} := L(k)$ for each $k \geq 1$. Based on the set of nonnegative mixing weights and the corresponding set of baseline increments, we define mixtures of SR and CUSUM e-detectors as $M_0^{\text{mSR}} = M_0^{\text{mCU}} := 1$, and for each $n \in \mathbb{N}$,

$$M_n^{\text{mSR}} := \sum_{k=1}^\infty \omega_k \sum_{j=1}^n \prod_{i=j}^n L_i(k) := \sum_{k=1}^\infty \omega_k M_n^{\text{SR}}(k), \tag{3.5}$$

$$M_n^{\text{mCU}} = \sum_{k=1}^\infty \omega_k \max_{j \in [n]} \prod_{i=j}^n L_i(k) := \sum_{k=1}^\infty \omega_k M_n^{\text{CU}}(k). \tag{3.6}$$

Let $K := |\{k : \omega_k > 0\}|$ be the number of nonzero mixing weights.

Finite Mixtures If $K < \infty$, we may for simplicity assume that the first K weights $\omega_1, \dots, \omega_K$ are the only nonzero values. In this case, we can compute mixtures of SR and CUSUM e-detectors by

$$M_n^{\text{mSR}} = \sum_{k=1}^K \omega_k M_n^{\text{SR}}(k), \quad \text{and} \quad M_n^{\text{mCU}} = \sum_{k=1}^K \omega_k M_n^{\text{CU}}(k), \tag{3.7}$$

where $M_n^{\text{SR}}(k)$ and $M_n^{\text{CU}}(k)$ are computed recursively as

$$M_n^{\text{SR}}(k) = L_n(k) \cdot [M_{n-1}^{\text{SR}}(k) + 1], \tag{3.8}$$

$$M_n^{\text{CU}}(k) = L_n(k) \cdot \max\{M_{n-1}^{\text{CU}}(k), 1\} \tag{3.9}$$

with $M_0^{\text{SR}}(k) = M_0^{\text{CU}}(k) = 0$ for each $k \in [K]$ and $n \in \mathbb{N}$. Therefore, if each computation of $L_n(k)$ has constant time and space complexities, then the evaluation of mixtures of SR and CUSUM e-detectors at each time n requires $O(K)$ time and space complexity.

Infinite Mixtures, Scheduling Functions and Adaptive Re-weighting If $K = \infty$ or if K is to be chosen adaptively as an increasing function of n we modify our strategy as follows. We first choose an increasing function $K : \mathbb{N} \rightarrow \mathbb{N}$, and let $K^{-1} : \mathbb{N} \rightarrow \mathbb{N}$ be the generalized inverse function of K defined by $K^{-1}(k) := \inf\{j \geq 1 : K(j) \geq k\}$ for each $k \in \mathbb{N}$. Note that K^{-1} is also an increasing function. We call such function K as a *scheduling function*. We intentionally overload notation: in what follows, $K(n)$ plays the same role as the constant K in the case of finite support. Note that $K^{-1}(k) \leq n$ for any $k \leq K(n)$; we will use this simple fact below when defining nested summations.

Based on a scheduling function K and its generalized inverse K^{-1} , we define *adaptive SR and CUSUM e-detectors*, M_n^{aSR} and M_n^{aCU} , respectively, as

$$M_n^{\text{aSR}} = \sum_{k=1}^{K(n)} \omega_k \sum_{j=K^{-1}(k)}^n \gamma_j \prod_{i=j}^n L_i(k) := \sum_{k=1}^{K(n)} \omega_k M_n^{\text{SR}}(k), \tag{3.10}$$

$$M_n^{\text{aCU}} = \sum_{k=1}^{K(n)} \omega_k \max_{K^{-1}(k) \leq j \leq n} \gamma_j \prod_{i=j}^n L_i(k) := \sum_{k=1}^{K(n)} \omega_k M_n^{\text{CU}}(k), \tag{3.11}$$

where each $\gamma_j := 1/\sum_{k=1}^{K(j)} \omega_k \geq 1$ is the *adaptively re-weighting factor* at time j , ensuring that the mixing weights always sum to one at each time. Here, we restrict not only the space over the index k from $[1, \infty]$ to $[1, K(n)]$ but also the space over the index j from $[1, n]$ to $[K^{-1}(k), n]$. This choice makes it possible to compute both M_n^{aSR} and M_n^{aCU} efficiently since each $M_n^{\text{SR}}(k)$ and $M_n^{\text{CU}}(k)$ have following recursive representations

$$M_n^{\text{SR}}(k) = L_n(k) \cdot [M_{n-1}^{\text{SR}}(k) + \gamma_n], \tag{3.12}$$

$$M_n^{\text{CU}}(k) = L_n(k) \cdot \max\{M_{n-1}^{\text{CU}}(k), \gamma_n\}, \tag{3.13}$$

for each $n \geq K^{-1}(k)$ and $M_n^{\text{aSR}}(k) = M_n^{\text{aCU}}(k) = 0$ for all $n = 0, 1, \dots, K^{-1}(k) - 1$. Therefore, if each computation of $L_n(k)$ has constant time and space complexities then the computations of adaptive SR and CUSUM e-detectors at each time n have $O(K(n))$ time and space complexities as well. For the purpose of implementing an online algorithm, we are typically interested in the case $K(n) = O(\log(n))$.

Remark 3.2. Both mixtures of SR and CUSUM e-detectors can be viewed as special cases of their adaptive counterparts where the scheduling function K is understood as a constant function. In this case, we have $\gamma_j = \sum_{k=1}^K \omega_k = 1$ for each j , and thus $M_n^{\text{aSR}} = M_n^{\text{mSR}}$ and $M_n^{\text{aCU}} = M_n^{\text{mCU}}$ for each $n \in \mathbb{N}$.

Unlike finite mixtures, the mixing distribution deployed in the adaptive SR and CUSUM e-detectors vary over time. Hence, we cannot simply apply Proposition 2.3 to check whether this adaptive scheme yields valid e-detectors. The

following proposition formally states the validity of adaptive SR and CUSUM e-detectors. The proof can be found in Appendix A.2.

Proposition 3.3. *For any mixing weights $\{\omega_k\}_{k \in \mathbb{N}}$ and a scheduling function K , adaptive SR and CUSUM e-detectors defined in (3.10) and (3.11) form valid e-detectors satisfying the condition (2.1).*

Now, based on M_n^{aCU} and M_n^{aSR} , the adaptive e-SR and e-CUSUM procedures are defined by the stopping times:

$$N_{\text{aSR}}^* := \inf\{n \geq 1 : M_n^{\text{aSR}} \geq 1/\alpha\}, \tag{3.14}$$

$$N_{\text{aCU}}^* := \inf\{n \geq 1 : M_n^{\text{aCU}} \geq c_\alpha\}, \tag{3.15}$$

where $\alpha \in (0, 1)$ is a fixed constant and c_α is a positive value controlling ARL of the adaptive e-CUSUM procedure by $1/\alpha$. Similar to the usual e-CUSUM procedure case we discussed before, we can always set $c_\alpha = 1/\alpha$. In this case, from the fact $N_{\text{aSR}}^* \leq N_{\text{aCU}}^*$, which is implied by $M_n^{\text{aSR}} \geq M_n^{\text{aCU}}$, we have

$$\mathbb{E}_{P, \infty} N_{\text{aCU}}^* \geq \mathbb{E}_{P, \infty} N_{\text{aSR}}^* \geq 1/\alpha, \tag{3.16}$$

where the last inequality comes from Theorem 2.4 with the fact that M^{aSR} is a valid e-detector. However, the threshold c_α for the adaptive e-CUSUM procedure can be chosen to be a significantly smaller value if we have enough knowledge about the pre-change distribution, as discussed in Section 2.5.

3.2 Worst Average Delay Analysis for Adaptive Mixtures of E-detectors

We now derive general upper bounds on worst average delays of the adaptive e-SR and e-CUSUM procedures. As we did before, we further assume that post-change observations $X_{\nu+1}, X_{\nu+2}, \dots$ are independent of the pre-change observations and form a strong stationary process. Also, we further assume that there exist a function f_k and an integer $m \geq 0$ such that $L_n(k) = f_k(X_n, X_{n-1}, \dots, X_{n-m})$ for each k and n . Again, if m is a strictly positive number then we implicitly assume that there exist m observations $X_0, X_{-1}, \dots, X_{1-m}$ from the pre-change distribution we can use to build sequential change detection procedures. Under this additional condition for worst average delay analysis, the following theorem provides analytically more tractable upper bounds on worst average delays for N_{aSR}^* and N_{aCU}^* .

Theorem 3.4. *Under additional conditions described above, worst average delays for N_{aSR}^* and N_{aCU}^* can be upper bounded as follows:*

$$\mathcal{J}_P(N_{\text{aSR}}^*) \leq \mathcal{J}_L(N_{\text{aSR}}^*) \leq \min_{j>m} [\mathbb{E}_{0, Q} N_{1/\alpha}(j) + j - 1], \tag{3.17}$$

$$\mathcal{J}_P(N_{\text{aCU}}^*) \leq \mathcal{J}_L(N_{\text{aCU}}^*) \leq \min_{j>m} [\mathbb{E}_{0, Q} N_{c_\alpha}(j) + j - 1], \tag{3.18}$$

where, for $j \in \mathbb{N}$ and $c > 0$, $N_c(j)$ is the stopping time

$$N_c(j) := \inf \left\{ n \geq 1 : \sum_{k=1}^{K(j)} \omega_k \prod_{i=1}^n L_i(k) \geq c \right\}. \quad (3.19)$$

Here, c_α is the same threshold used to build the adaptive e-CUSUM procedure in (3.15). Note that for mixtures of the SR and CUSUM e-detectors where the scheduling function K is a constant function, the stopping times in the upper bounds do not depend on the index j , and thus the upper bounds can be reduced as

$$\mathcal{J}_P(N_{\text{mSR}}^*) \leq \mathcal{J}_L(N_{\text{mSR}}^*) \leq \mathbb{E}_{0,Q} N_{1/\alpha} + m, \quad (3.20)$$

$$\mathcal{J}_P(N_{\text{mCU}}^*) \leq \mathcal{J}_L(N_{\text{mCU}}^*) \leq \mathbb{E}_{0,Q} N_{c_\alpha} + m, \quad (3.21)$$

where, for $c > 0$, N_c is the stopping time

$$N_c := \inf \left\{ n \geq 1 : \sum_{k=1}^K \omega_k \prod_{i=1}^n L_i(k) \geq c \right\}. \quad (3.22)$$

The proof of upper bounds on worst average delays can be found in Appendix A.2.

Unlike the baseline e-detector case, however, due to mixing weights, it is nontrivial to get further simplified upper bounds on worst average delays as we did in Section 2.7. Next, we present specific adaptive e-SR and e-CUSUM procedures based on exponential baseline increments where we can compute both procedures efficiently and derive upper bounds on worst average delays in explicit forms.

4. EXPONENTIAL BASELINE E-DETECTORS AND THEIR MIXTURES

Building upon recent advances in time uniform concentration inequalities and sequential testing developed in [11] and [36], below we consider an exponential structure on baseline e-detectors. We show that, in this setting, it is possible to approximate the ‘‘oracle’’ e-SR and e-CUSUM procedures based on the knowledge of the optimal (but unknown) λ^{op} by adaptive procedures built using a mixture of carefully chosen set of baseline increments $\{L^{\lambda_k}\}_{k \geq 1}$ with mixing weights $\{\omega_k\}_{k \geq 1}$.

To be specific, assume there exists an extended real-valued convex function ψ on \mathbb{R} that is finite and strictly convex on a set $\Pi \subset \mathbb{R}$ containing 0 in its interior Π° . Furthermore, assume ψ is continuously differentiable on Π° with $\nabla\psi(0) = 0 = \psi(0)$. Then define the ‘‘exponential baseline increment’’ as follows.

Definition 4.1 (Exponential baseline increment). For each $n \in \mathbb{N}$ and $\lambda \in \Pi$, define

$$L_n^\lambda = \exp\{\lambda s(X_n) - \psi(\lambda)v(X_n)\}, \quad (4.1)$$

where s is a real-valued function and v is a positive function on the sample space. $L^\lambda := \{L_n^\lambda\}_{n \geq 1}$ is called an exponential baseline increment if it satisfies condition (2.6) in Definition 2.8.

Above, s and v are mnemonics for sum and variance. For each $Q \in \mathcal{Q}$, define

$$\mu(Q) := \mathbb{E}_{0,Q} s(X_1), \quad \sigma^2 := \mathbb{E}_{0,Q} v(X_1) \quad \text{and} \quad \Delta^{\text{op}}(Q) := \frac{\mu(Q)}{\sigma^2(Q)}, \quad (4.2)$$

where we assume that all expectations are finite. The following proposition provides an explicit expression for $D(Q||\mathcal{P}) := \max_{\lambda \in \Pi} \mathbb{E}_{0,Q} \log L_1^\lambda$ and a sufficient condition to have $D(Q||\mathcal{P}) > 0$ when the underlying baseline increments have the form specified in (4.1).

Proposition 4.2. For a fixed $Q \in \mathcal{Q}$, suppose there exist $\lambda^{\text{op}} \in \Pi^\circ$ such that $\Delta^{\text{op}}(Q) = \nabla\psi(\lambda^{\text{op}})$. Then,

$$D(Q||\mathcal{P}) = \mathbb{E}_{0,Q} \log L_1^{\lambda^{\text{op}}} = \psi^*(\Delta^{\text{op}}(Q))\sigma^2(Q) \geq 0, \quad (4.3)$$

where ψ^* is the convex conjugate of ψ . Thus, if $\Delta^{\text{op}}(Q) \neq 0$, we have $D(Q||\mathcal{P}) > 0$.

The proof of Proposition 4.2 can be found in Appendix B.1. For the rest of the section, we assume that

$$\begin{aligned} \{\lambda \in \mathbb{R} : \Delta^{\text{op}}(Q) = \nabla\psi(\lambda), Q \in \mathcal{Q}\} &\subset \Pi, \\ \Delta^{\text{op}}(Q) &\neq 0, \quad \forall Q \in \mathcal{Q}. \end{aligned}$$

Then, Proposition 4.2 implies $D(Q||\mathcal{P}) > 0$ for all $Q \in \mathcal{Q}$. Also, for ease of notation, we will drop the dependency of Q from related parameters and simply write $\lambda^{\text{op}}, \mu, \sigma^2$ and Δ .

The exponential structure of the baseline increment in (4.1) results in a simple form of λ^{op} such that

$$\lambda^{\text{op}} = (\nabla\psi)^{-1}(\Delta^{\text{op}}) = \nabla\psi^*(\Delta^{\text{op}}), \quad (4.4)$$

where the second equality comes from the fact that $\lambda = \nabla\psi^* \circ \nabla\psi(\lambda)$ for each $\lambda \in \Pi$. Although λ^{op} still depends on the unknown post-change distribution Q via Δ^{op} , in many cases, we can find upper and lower bounds on Δ^{op} . In this section, we explain how to use the knowledge of the range of Δ^{op} to build a mixture of exponential baseline e-detectors that has explicit upper bounds on worst average delays.

4.1 Separated Pre- and Post-change Distributions

Suppose we have knowledge of upper and lower bounds on the parameter Δ^{op} given in (4.2), i.e. we know a pair (Δ_L, Δ_U) such that $\Delta_L < \Delta^{\text{op}} < \Delta_U$. It then follows that $\lambda_L < \lambda^{\text{op}} < \lambda_U$, where $\lambda_L = \nabla\psi^*(\Delta_L)$, $\lambda^{\text{op}} = \nabla\psi^*(\Delta^{\text{op}})$ and $\lambda_U = \nabla\psi^*(\Delta_U)$. To simplify presentation, we only consider the *one-sided* and *well-separated* case: $0 < \lambda_L < \lambda_U$.

Let $1/\alpha$ be the target level of the ARL control for a fixed $\alpha \in (0, 1)$. Let $\{L(k)\}_{k \in [K]}$ and $\{\omega_k\}_{k \in [K]}$ be K exponential baseline increments and mixing weights whose specific values will be defined later in this subsection. Since each $L_n(k)$ is a function of the n -th observation X_n for each $k \in [K]$, Theorem 3.4 implies that, if the post-change observations form a strong stationary process then the worst

average delays for mixtures of e-SR and e-CUSUM procedures, N_{mSR}^* and N_{mCU}^* can be upper bounded by $\mathbb{E}_{0,Q}N_{1/\alpha}$ and $\mathbb{E}_{0,Q}N_{c_\alpha}$, respectively, where N_c is the stopping time defined in (3.22). Furthermore, as we can always set the threshold for the e-CUSUM procedure to be $c_\alpha \leq 1/\alpha$, we have $\mathbb{E}_{0,Q}N_{c_\alpha} \leq \mathbb{E}_{0,Q}N_{1/\alpha}$. Therefore, in this subsection, we construct a set of baseline increments for which we can derive a tight bound on $\mathbb{E}_{0,Q}N_{1/\alpha}$.

Algorithm 1 describes our methodology for computing mixtures of e-SR procedures in detail. The inputs to the algorithm are the upper and lower bounds Δ_U and Δ_L on Δ^{op} and the maximal number of baselines processes K_{max} . Mixture of e-CUSUM procedures can be executed similarly by replacing Line 9 by

$$M_n^{\text{CU}}(k) \leftarrow \exp\{\lambda_k s(X_n) - \psi_k v(X_n)\} \cdot \max\{M_{n-1}^{\text{CU}}(k), 1\}. \quad (4.5)$$

Also, for the mixture of e-CUSUM procedures, we can replace the threshold $1/\alpha$ with a smaller value c_α if we have enough information about the pre-change distribution. For both e-SR and e-CUSUM, at each time n , updates of mixtures of e-detectors have $O(K_\alpha)$ time and space complexities, which do not depend on n .

Algorithm 1 relies critically on the function `computeBaseline` in Line 1, which returns a set of parameters and weights to compute a mixture of e-detectors along with a threshold value $g_\alpha > 0$ that will appear in the upper bound on worst average delays given in Theorem 4.3 that will be explained below. The details of `computeBaseline` are fairly technical and are given in Algorithm 3 in Appendix B.1.

In the main result of this section, we provide bounds on ARL and worst average delays for the mixtures of e-CP procedures obtained with Algorithm 1 that is a function of the parameter λ^{op} and the threshold value g_α . The proof can be found in Appendix B.1.

Theorem 4.3. *Let N_{mSR}^* and N_{mCU}^* be the stopping times corresponding to the mixtures of e-SR and e-CUSUM procedures in Algorithm 1 and its variant, respectively. Then, both procedures control the ARL by $1/\alpha$. If we further assume that the post-change observations $X_{\nu+1}, X_{\nu+2}, \dots$ are iid samples from a post-change distribution Q , then the worst average delays for N_{mSR}^* and N_{mCU}^* can be bounded as*

$$\begin{aligned} & \max\{\mathcal{J}_L(N_{\text{mSR}}^*), \mathcal{J}_L(N_{\text{mCU}}^*)\} \\ & \leq \frac{g_\alpha}{D(Q||\mathcal{P})} + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda^{\text{op}})}]}{[D(Q||\mathcal{P})]^2} + 1. \end{aligned} \quad (4.6)$$

The same bound holds also for $\mathcal{J}_P(N_{\text{mSR}}^*)$ and $\mathcal{J}_P(N_{\text{mCU}}^*)$.

In Proposition B.2 in Appendix B.1, we show that if the number of baseline processes K_{max} in Algorithm 1 is chosen large enough, then the quantity g_α returned by `computeBaseline` is at most

$$\inf_{\eta > 1} \eta \left[\log(1/\alpha) + \log \left(1 + \left[\log_\eta \frac{\psi^*(\Delta_U)}{\psi^*(\Delta_L)} \right] \right) \right], \quad (4.7)$$

which can be easily evaluated numerically. Expression (B.15) in Appendix B.1 provides a precise formula for how large K_{max} needs to be in order for the above bound to be in effect. In most practical cases, $K_{\text{max}} = 1000$ is a large enough choice. Also, in many canonical examples we will present later, if we choose large enough K_{max} satisfying the condition (B.15) then the first term $\frac{g_\alpha}{D(Q||\mathcal{P})}$ of the upper bound of worst average delays in Theorem 4.3 become a leading term. In this case, from the inequality (4.7), we can check that this leading term is $O(\log(1/\alpha)/D(Q||\mathcal{P}))$ as $\alpha \rightarrow 0$.

Remark 4.4. If there is only one pre-change distribution P and one post-change distribution Q , both from a natural univariate exponential family, then their likelihood ratio forms an exponential baseline increment. In this case, the above upper bound becomes $O(\log(1/\alpha)/\text{KL}(Q||P))$ as $\alpha \rightarrow 0$, matching the rate of the known lower bounds [18].

The bound on worst average delays in Theorem 4.3 is obtained by analyzing an auxiliary stopping time

$$\bar{N}_g := \inf \left\{ n \geq 1 : \sup_{\lambda \in (\lambda_L, \lambda_U)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g \right\}, \quad g > 1. \quad (4.8)$$

Using the same arguments as in the proof of Proposition 2.14, we immediately have that if the post-change observations are iid from Q , then for any $g > 1$,

$$\mathbb{E}_{0,Q}\bar{N}_g \leq \frac{g}{D(Q||\mathcal{P})} + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda^{\text{op}})}]}{[D(Q||\mathcal{P})]^2} + 1. \quad (4.9)$$

The bound (4.6) is finally established by showing that the stopping time \bar{N}_{g_α} obtained by using the threshold g_α produced by Algorithm 3 is a deterministic upper bound to the stopping times N_{c_α} and $N_{1/\alpha}$ corresponding to mixtures of SR and CUSUM e-detectors. In detail, it holds that for any stream of observations X_1, X_2, \dots ,

$$N_{c_\alpha} \leq N_{1/\alpha} \leq \bar{N}_{g_\alpha}. \quad (4.10)$$

This nontrivial result is formally stated in Lemma B.1 in Appendix B.1. Its proof leverages geometric arguments used in [36, Theorem 2] to analyze sequential generalized likelihood ratio tests.

4.2 Non-separated Pre- and Post-change Distributions

The previous subsection discussed how to build mixtures of e-SR and e-CUSUM procedures with an explicit upper bound on worst average delays when we have known and positive boundary values, λ_L and λ_U on the unknown λ^{op} via the knowledge of $\Delta_L < \Delta^{\text{op}} < \Delta_U$. However, in many cases, we may not be fully certain about the boundary values. In this subsection, we discuss how we can generalize the previous argument to the no separation case whereby we

Algorithm 1: Pseudo-code of the mixture of e-SR procedures.

Input: ARL parameter $\alpha \in (0, 1)$, Boundary values $0 < \Delta_L < \Delta_U$,
 Maximum number of baselines $K_{\max} \in \mathbb{N}$.
Output: Stopping time N_{mSR}^* of the mixture of e-SR procedures.
Data: Data stream X_1, X_2, \dots (observed sequentially)
 1 $\{\lambda_0, \lambda_1, \dots, \lambda_{K_\alpha}\}, \{\omega_0, \omega_1, \dots, \omega_{K_\alpha}\}, g_\alpha \leftarrow \text{computeBaseline}(\alpha, \Delta_L, \Delta_U, K_{\max})$
 2 **for** $k = 0, 1, \dots, K_\alpha$ **do**
 3 $M_0^{\text{SR}}(k) \leftarrow 0, \psi_k \leftarrow \psi(\lambda_k)$
 4 $M_0^{\text{mSR}} \leftarrow 0, n \leftarrow 0$
 5 **while** $M_n^{\text{mSR}} < 1/\alpha$ **do**
 6 $n \leftarrow n + 1$
 7 Observe X_n
 8 **for** $k = 0, 1, \dots, K_\alpha$ **do**
 9 $M_n^{\text{SR}}(k) \leftarrow \exp\{\lambda_k s(X_n) - \psi_k v(X_n)\} \cdot [M_{n-1}^{\text{SR}}(k) + 1]$
 10 $M_n^{\text{mSR}} \leftarrow \sum_{k=0}^{K_\alpha} \omega_k M_n^{\text{SR}}(k)$
 11 $N_{\text{mSR}}^* \leftarrow n$
 12 **return** The stopped time N_{mSR}^*

only know the sign of $\lambda^{\text{op}} (> 0)$ but do not have specific boundary values.

Recall that, for the well-separated case, we calibrated the mixtures of finitely many exponential baseline e-detectors using the stopping time \bar{N}_{g_α} in (4.8), which is in turn based on the maximum of underlying baseline increments over the known upper and lower bounds of λ^{op} . Since we no longer have knowledge of the boundary values λ_L and λ_U , we may use similar stopping times where the range of maximum and the threshold slowly increase over time. In this case, we need an infinite sequence of baseline procedures $\{L(k)\}_{k \in \mathbb{N}}$ and mixing weights $\{\omega_k\}_{k \in \mathbb{N}}$ to build adaptive e-SR and e-CUSUM procedures.

The bound in Theorem 3.4 along with the fact $\gamma_j \geq 1$ for all $j \in \mathbb{N}$ implies that, for any given scheduling function $K : \mathbb{N} \rightarrow \mathbb{N}$, if the post-change observations form a strong stationary process then worst average delays for adaptive e-SR and e-CUSUM procedures can be upper bounded by $\min_{j \geq 1} [\mathbb{E}_{0,Q} N_{1/\alpha}(j) + j - 1]$ and $\min_{j \geq 1} [\mathbb{E}_{0,Q} N_{c_\alpha}(j) + j - 1]$, respectively, where we recall that $N_c(j)$ is defined for $c > 0$ by

$$N_c(j) := \inf \left\{ n \geq 1 : \sum_{k=1}^{K(j)} \omega_k \prod_{i=1}^n L_i(k) \geq c \right\}. \quad (4.11)$$

Again, since we can set the threshold for the e-CUSUM procedure in such a manner that $c_\alpha \leq 1/\alpha$ (so that $\mathbb{E}_{0,Q} N_{c_\alpha}(j) \leq \mathbb{E}_{0,Q} N_{1/\alpha}(j)$), in this subsection, we focus on constructing a set of baseline increments on which we can derive a tight upper bound on $\min_{j \geq 1} [\mathbb{E}_{0,Q} N_{1/\alpha}(j) + j - 1]$.

To derive the set of baseline increments, we use a time-varying boundary function g . Here, we intentionally overload notation: the constant g in the previous subsection for the well-separation case can be viewed as a constant function g in what follows. Let $g : [1, \infty) \rightarrow [0, \infty)$ be a nonnegative

and nondecreasing continuous function such that the mapping $t \mapsto g(t)/t$ is nonincreasing and $\lim_{t \rightarrow \infty} g(t)/t = 0$. For a chosen positive number $\Delta_0 > 0$, let

$$D_0 := \psi^*(\Delta_0) \text{ and } V_0 := \inf\{t \geq 1 : D_0 \geq g(t)/t\}.$$

Now, for any fixed $\eta > 1$ and $j \in \mathbb{N}$, define $\Delta_1 > \Delta_2 > \dots$ as positive solutions of the equations

$$\psi^*(\Delta_k) = \frac{g(V_0 \eta^k)}{V_0 \eta^k}, \quad k = 0, 1, 2, \dots \quad (4.12)$$

Finally, based on the sequence $\{\Delta_k\}_{k \geq 0}$, define

$$\lambda_k := \nabla \psi^*(\Delta_k),$$

and set $\omega_0 := \alpha^{-1} e^{-g(V_0)} \mathbb{1}(g(V_0) > v_{\min} D_0)$, $\omega_k := \alpha^{-1} e^{-g(V_0 \eta^k)}/\eta$ for each $k \in \mathbb{N}$ where $v_{\min} := \min_x v(x)$, recalling the function v from Definition 4.1.

Based on the quantities defined above, we can construct the stopping time $N_{1/\alpha}(j)$ for each j . The following lemma shows that we can upper bound the stopping time $N_{1/\alpha}(j)$ with another stopping time $\bar{N}_g(j)$ from which we can derive an explicit upper bound on its expected stopping time.

Lemma 4.5. *For any fixed $j \geq 1$, $\Delta_0 > 0$, and tuning parameter $\eta > 1$, let $N_{1/\alpha}(j)$ be the stopping time based on the parameters defined above. Then, we have*

$$N_{1/\alpha}(j) \leq \bar{N}_g(j), \quad (4.13)$$

where $\bar{N}_g(j)$ is a stopping time defined by

$$\bar{N}_g(j) := \inf \left\{ n \geq 1 : \sup_{\lambda \in (\lambda_{K(j)}, \lambda_0)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g(V_0 \eta^{K(j)}) \right\}. \quad (4.14)$$

Note that the chosen set of weights $\{\omega_k\}_{k \geq 0}$ yields valid adaptive e-SR and e-CUSUM procedures if

$$e^{-g(V_0)} \mathbb{1}(g(V_0) > v_{\min} D_0) + \sum_{k=1}^{\infty} e^{-g(V_0 \eta^k)/\eta} \leq \alpha. \quad (4.15)$$

Once the above condition is satisfied, we can use the worst average delay analysis in Section 3.2 with the bound in Lemma 4.5 to get an explicit upper bound on the worst average delay of the adaptive e-SR and e-CUSUM procedures.

In detail, let j^{op} be the smallest integer satisfying $\lambda_{K(j^{\text{op}})} < \lambda^{\text{op}}$ and set $K^{\text{op}} := K(j^{\text{op}})$. If we also have $\lambda^{\text{op}} < \lambda_0$ then Lemma 4.5 implies

$$\mathbb{E}_{0,Q} N_{1/\alpha}(j^{\text{op}}) \leq \mathbb{E}_{0,Q} \bar{N}_g(j^{\text{op}}) \leq \mathbb{E}_{0,Q} N_{\text{op}}, \quad (4.16)$$

where the stopping time N_{op} is defined by

$$N_{\text{op}} := \inf \left\{ n \geq 1 : \sum_{i=1}^n \log L_i^{(\lambda^{\text{op}})} \geq g(V_0 \eta^{K^{\text{op}}}) \right\}, \quad (4.17)$$

and the expectation $\mathbb{E}_{0,Q} N_{\text{op}}$ is typically on the order of $g(V_0 \eta^{K^{\text{op}}})/D(Q|\mathcal{P})$. Based on this observation, in the rest of this subsection, we introduce a practical and interpretable way to choose a boundary function g and related tuning parameters which minimize the leading term $g(V_0 \eta^{K^{\text{op}}})$ while satisfying the condition (4.15) on the set of mixing weights.

First note that, although we have no bounds on Δ^{op} in the no separation case, we can still choose Δ_L and Δ_0 with $\Delta_L < \Delta_0$ as tuning parameters that represent our initial guess on the range of the unknown Δ^{op} . Since it is possible that the unknown parameter Δ^{op} of the post-change distribution is outside of the boundary (Δ_L, Δ_0) , instead of assigning the entire α to the inside of the guessed interval, we split it into two parts by $r\alpha$ and $(1-r)\alpha$, respectively where $r \in (0, 1)$ is another tuning parameter called the importance weight. Roughly speaking, larger r implies we make a higher bet on that the unknown Δ^{op} is inside of our chosen boundaries (Δ_L, Δ_0) .

Now, given tuning parameters Δ_L, Δ_0 and r , we compute the set of $\{g_{r\alpha}, K_L, \eta\}$ by executing the function `computeBaseline`, just like in Algorithm 1, except that α is replaced by $r\alpha$. Then, we can extend the boundary function g to accommodate the case in which the unknown Δ^{op} is not inside the initial interval we had guessed. To be specific, we use the boundary function

$$t \in [1, \infty) \mapsto g(t) := g_{r\alpha} + s\eta \log \left(1 + \log_{\eta} \left(\frac{t}{V_0 \eta^{K_L}} \vee 1 \right) \right), \quad (4.18)$$

where $V_0 := g_{r\alpha}/D_0$ and $s > 1$ is a constant obtained as the

solution of the equation

$$\begin{aligned} \zeta(s) - 1 &:= \sum_{k=1}^{\infty} \frac{1}{(1+k)^s} \\ &= e^{g_{r\alpha}/\eta} [\alpha - \{e^{-g_{r\alpha}} \mathbb{1}(g_{r\alpha} > D_0) + K_L e^{-g_{r\alpha}/\eta}\}]. \end{aligned} \quad (4.19)$$

Note that the right hand side of the above equation is approximately equal to $(1-r)\alpha e^{g_{r\alpha}/\eta}$. Therefore,

$$s \approx \zeta^{-1}(1 + [1-r]\alpha e^{g_{r\alpha}/\eta}). \quad (4.20)$$

In Algorithm 2, we provide the detailed steps for the adaptive e-SR procedure based on the boundary function in (4.18). The algorithm can be easily modified for the adaptive e-CUSUM procedure by replacing the update in Line 14 with the rule

$$M_n^{\text{CU}}(k) \leftarrow \exp\{\lambda_k s(X_n) - \psi_k v(X_n)\} \cdot \max\{M_{n-1}^{\text{CU}}(k), \gamma\}. \quad (4.21)$$

Also, for the adaptive e-CUSUM procedure, we can replace the threshold $1/\alpha$ with a smaller value c_α if we have enough information about the pre-change distribution.

In term of computational complexity, in Algorithm 2 we set the scheduling function $K : \mathbb{N} \rightarrow \mathbb{N}$ as

$$K(n) := K_L + \lceil m \log_{\eta} n \rceil, \quad (4.22)$$

where $m \geq 1$ is a tuning parameter. Therefore, for both adaptive e-SR and e-CUSUM procedures, updates of statistics have $O(m \log_{\eta} n)$ time and space complexities at each time n . Although it is not a fully online algorithm, logarithm time and space complexities make it feasible to run adaptive e-SR and e-CUSUM procedures in most practical online settings.

From Section 3.1, we know that both procedures control the ARL by $1/\alpha$. The following theorem introduces explicit bounds on the worst average delays for both procedures.

Corollary 4.6. *Let N_{aSR}^* and N_{aCU}^* be stopping times corresponding to the adaptive e-SR procedures in Algorithm 2 and its e-CUSUM variant, respectively. Then, both procedures control ARL by $1/\alpha$. If we further assume that post-change observations $X_{\nu+1}, X_{\nu+2}, \dots$ are iid samples from a post-change distribution then the worst average delays for N_{aSR}^* and N_{aCU}^* can be upper bounded as*

$$\begin{aligned} &\max\{\mathcal{J}_L(N_{\text{aSR}}^*), \mathcal{J}_L(N_{\text{aCU}}^*)\} \\ &\leq \begin{cases} \frac{g_{r\alpha}}{D(Q|\mathcal{P})} \frac{\psi^*(\Delta^{\text{op}})}{\psi^*(\Delta_0)} + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda_0)}]}{[D(Q|\mathcal{P})]^2} [\frac{\psi^*(\Delta^{\text{op}})}{\psi^*(\Delta_0)}]^2 + 1 & \text{if } \Delta^{\text{op}} \geq \Delta_0 \\ \frac{g_{r\alpha}}{D(Q|\mathcal{P})} + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda^{\text{op}})}]}{[D(Q|\mathcal{P})]^2} + 1 & \text{if } \Delta^{\text{op}} \in (\Delta_L, \Delta_0) \\ \frac{g_{r\alpha} + s\eta \log(1 + K^{\text{op}} - K_L)}{D(Q|\mathcal{P})} + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda^{\text{op}})}]}{[D(Q|\mathcal{P})]^2} + [\frac{\psi^*(\Delta_L)}{\psi^*(\Delta^{\text{op}})} \frac{g_{r\alpha} + s\eta \log(1 + K^{\text{op}} - K_L)}{g_{r\alpha}}]^{1/m} & \text{if } \Delta^{\text{op}} \leq \Delta_L. \end{cases} \end{aligned} \quad (4.23)$$

Algorithm 2: Pseudo-code of the adaptive e-SR procedures.

Input: ARL parameter $\alpha \in (0, 1)$, Tuning parameters $\Delta_L < \Delta_0$, importance weight $r \in (0, 1)$, scheduling parameter $m \geq 1$, Number of baselines for the well-separated regime $K_0 \in \mathbb{N}$.

Output: Stopping time N_{aSR}^* of the adaptive e-SR procedure.

Data: Data stream X_1, X_2, \dots (observed sequentially)

- 1 Obtain $\{\lambda_0, \lambda_1, \dots, \lambda_{K_L}\}$, $\{\omega_0, \omega_1, \dots, \omega_{K_L}\}$, $\{g_{r\alpha}, K_L, \eta, W\}$ by executing `computeBaseline`($r\alpha, \Delta_L, \Delta_0, K_0$) in Algorithm 3.
- 2 $s \leftarrow \zeta^{-1}(1 + [\alpha - W]e^{g_{r\alpha}/\eta}) / * \zeta(s) \approx 1 + [1 - r]\alpha e^{g_{r\alpha}/\eta} \quad */$
- 3 $M_0^{\text{SR}}(k) \leftarrow 0$, $\psi_k \leftarrow \psi(\lambda_k)$, $\omega_k \leftarrow \omega_k \frac{W}{\alpha}$, $\forall k = 0, 1, \dots, K_L$
- 4 $M_0^{\text{aSR}} \leftarrow 0$, $\gamma \leftarrow 1 / \sum_{k=0}^{K_L} \omega_k$, $n \leftarrow 0$
- 5 **while** $M_n^{\text{aSR}} < 1/\alpha$ **do**
- 6 $n \leftarrow n + 1$
- 7 /* Occasionally add a new baseline increment. */
- 8 $K_n \leftarrow K_L + \lceil m \log_\eta n \rceil$
- 9 **if** $K_n > K_{n-1}$ **then**
- 10 **for** $k = K_{n-1} + 1, \dots, K_n$ **do**
- 11 Compute Δ_k as the solution of $\psi^*(z) = \frac{g_k}{V_0 \eta^k}$ with respect to $z(> 0)$, where $V_0 := g_{r\alpha}/D_0$ and
- 12 $g_k := g_{r\alpha} + s\eta \log(1 + k - K_L)$.
- 13 $M_{n-1}^{\text{aSR}}(\mu_k) \leftarrow 0$, $\psi_k \leftarrow \psi(\lambda_k)$, $\omega_k \leftarrow \alpha^{-1} e^{-g_k/\eta}$
- 14 $\gamma \leftarrow (1/\gamma + \sum_{k=K_{n-1}}^{K_n} \omega_k)^{-1}$
- 15 Observe X_n
- 16 $M_n^{\text{SR}}(k) \leftarrow \exp\{\lambda_k s(X_n) - \psi_k v(X_n)\} \cdot [M_{n-1}^{\text{SR}}(k) + \gamma]$, $\forall k = 0, 1, \dots, K_n$
- 17 $M_n^{\text{aSR}} \leftarrow \sum_{k=0}^{K_n} \omega_k M_n^{\text{SR}}(k)$
- 18 $N_{\text{aSR}}^* \leftarrow n$
- 19 **return** The stopped time N_{aSR}^*

Note that $\eta, s > 1$ and $r \in (0, 1)$ do not depend on the unknown Δ^{op} .

5. APPLICATION TO REAL DATA AND SIMULATION STUDY

5.1 Bernoulli Random Variables with Dependent, Time-varying Means

Winning Rates of the Cavaliers To illustrate how sequential change detection procedures based on e-detectors work, we revisit the example of the Cleveland Cavaliers, an American professional basketball team introduced in Section 1.2. Instead of using Plus-Minus stats, in this example, we are monitoring the performance of the Cavaliers by keeping track of wins and losses over all the games. Let $X_1, X_2, \dots \in \{0, 1\}$ be the sequence of win indicators during 2010–11 to 2017–18 regular seasons, where $X_i = 1$ if the Cavaliers won game i . Though Figure 2 presents monthly and seasonal averages for the purpose of visualization, we use the underlying binary sequence to build a sequential change detection procedure.

Modeling Winning Probabilities as a Dependent Sequence of Bernoullis To detect a significant improvement of the performance of the Cavaliers, we assume that before an unknown changepoint $\nu \in \mathbb{N} \cup \{\infty\}$, the conditional average of

winning probability given the sample history is less than or equal to $p_0 := 0.49$. That is, under any pre-change distribution P we have $p_n := \mathbb{E}_{P, \infty}[X_n | \mathcal{F}_{n-1}] \leq p_0$. (For simplicity, \mathcal{F} is taken to be the natural filtration of the data.) Thus, the pre-change class of distributions is

$$\mathcal{P} := \{(p_1, p_2, \dots) : p_i \leq p_0, \forall i \geq 1\},$$

where we parameterize each distribution P over binary sequences by the sequence of conditional probabilities.

Our objective is to build mixtures of e-SR and e-CUSUM procedures tuned to quickly detect any significantly improved win rate larger than $q_0 := 0.51$ after the changepoint. This can be modeled by assuming that after some changepoint ν , the distribution Q is such that $\mathbb{E}_{P, \nu, Q}[X_n | \mathcal{F}_{n-1}, n > \nu] := q_n \geq q_0$. Thus, we may think of the post-change class of distributions as being

$$\mathcal{Q} := \{(q_1, q_2, \dots) : q_i \geq q_0, \forall i \geq 1\}.$$

In particular, this formalization allows for the winning probabilities to fluctuate over time before and after the changepoint (accounting for factors like form, injuries, etc.).

Deriving Exponential Baseline Processes For each $\lambda > 0$, define a baseline increment process $L^{(\lambda)} := \{L_n^{(\lambda)}\}_{n \geq 1}$ as

$$L_n^{(\lambda)} := \exp\{\lambda(X_n - p_0) - \psi_B(\lambda)\}, \quad (5.1)$$

where $\psi_B(\lambda) := \log(1 - p_0 + p_0 e^\lambda) - \lambda p_0$ is the Bernoulli cumulant generating function. Note that each $L^{(\lambda)}$ is a valid baseline increment as it satisfies the inequality (2.6) in Definition 2.8. That is, under any pre-change distribution P , we have

$$\begin{aligned} & \mathbb{E}_{P,\infty}[L_n^{(\lambda)} \mid \mathcal{F}_{n-1}] \\ &= \mathbb{E}_{P,\infty}[\exp\{\lambda X_n - \log(1 - p_0 + p_0 e^\lambda)\} \mid \mathcal{F}_{n-1}] \\ &= \frac{\mathbb{E}_{P,\infty}[e^{\lambda X_n} \mid \mathcal{F}_{n-1}]}{1 - p_0 + p_0 e^\lambda} = \frac{1 - p_n + p_n e^\lambda}{1 - p_0 + p_0 e^\lambda} \leq 1, \quad \forall \lambda \geq 0. \end{aligned}$$

To derive exponential baseline processes, we first consider a simplified post-change distribution Q where each post-change observation is identically distributed with $\mathbb{E}_{0,Q}[X_1] := q \geq q_0 > p_0$. In this case, the optimal choice of $\lambda \geq 0$ given by

$$\lambda^{\text{op}} := \arg \max_{\lambda \geq 0} \mathbb{E}_{0,Q} \exp\{\lambda(X_1 - p_0) - \psi_B(\lambda)\}. \quad (5.2)$$

Since the baseline increment has the exponential structure, by Proposition 4.2, we have that

$$D(Q \parallel \mathcal{P}) := \mathbb{E}_{0,Q} \log L_1^{(\lambda^{\text{op}})} = \psi_B^*(q - p_0) = \text{KL}(q \parallel p_0), \quad (5.3)$$

where $\text{KL}(q \parallel p_0)$ is the Kullback-Leibler (KL) divergence of Bernoulli distributions with parameters q and p_0 written as

$$\text{KL}(q \parallel p_0) := q \log \frac{q}{p_0} + (1 - q) \log \frac{1 - q}{1 - p_0}, \quad (5.4)$$

for $q, p_0 \in (0, 1)$. The appearance of the KL divergence in (5.3) is not a coincidence as the baseline increment can be viewed as a re-parametrized likelihood ratios between two Bernoulli processes. However, the simple geometric structure of the baseline increment make it possible to utilize a prior knowledge about the post-change distribution via Algorithm 1 and 2.

For instance, suppose we know upper and lower bounds of conditional means of the post-change distribution as $q_n \in (q_L, q_U), \forall n > \nu$. Let N_{mSR}^* and N_{mCU}^* be stopping times of mixtures of e-SR and e-CUSUM procedures in Algorithm 1. In this case, derived sequential change detection procedures do not rely on a specific choice of a post-change distribution $Q \in \mathcal{Q}$. However, these procedures can still perform almost as well as the one optimized to a specific choice of the post-change distribution within the same range (q_L, q_U) . Typically, if the post-change observations are iid samples from a post-change distribution Q with $\mathbb{E}_{0,Q}[X_1] := q \in (q_L, q_U)$, then by Theorem 4.3, the worst average delays have the following explicit bound:

$$\begin{aligned} & \max\{\mathcal{J}_L(N_{\text{mSR}}^*), \mathcal{J}_L(N_{\text{mCU}}^*)\} \\ & \leq \frac{g_\alpha}{\text{KL}(q \parallel p_0)} + \frac{q(1 - q)[\log(\frac{1 - p_0}{p_0} \frac{q}{1 - q})]^2}{[\text{KL}(q \parallel p_0)]^2} + 1. \end{aligned}$$

Typically for small $\alpha \ll 1$, from Proposition B.2, we can simplify the above upper bound as

$$\begin{aligned} & \max\{\mathcal{J}_L(N_{\text{mSR}}^*), \mathcal{J}_L(N_{\text{mCU}}^*)\} \\ & \lesssim \frac{\inf_{\eta > 1} \eta [\log(1/\alpha) + \log(1 + \lceil \log_\eta \frac{\text{KL}(q_U \parallel p_0)}{\text{KL}(q_L \parallel p_0)} \rceil)]}{\text{KL}(q \parallel p_0)}, \end{aligned}$$

which matches the rate of the worst average delays, $O(\log(1/\alpha)/\text{KL}(q \parallel p_0))$ of the oracle sequential change detection procedure as $\alpha \rightarrow 0$.

Implementation of Algorithm 1 and Its Results The lower bound q_L can be chosen as $q_0 = 0.51$ since it is the minimum winning rate we consider as a significant improvement from before the changepoint, when the rates are upper bounded by $p_0 = 0.49$. We can also safely assume that the win rate cannot be too high given the competitiveness of the NBA, so that the improved win rates cannot be larger than 0.9. In our framework, these considerations can be encoded by setting $\Delta_L := q_0 - p_0 = 0.02$ and $\Delta_U := 0.41$ as input parameters of Algorithm 1. As in Section 1.2, we set $\alpha := 10^{-3}$ to ensure that the ARL is at least $1/\alpha := 10^3$, which is more than the total number of games over 12 years of regular seasons. Finally, we set the maximum number of baselines $K_{\text{max}} := 1000$. In fact, the `computeBaseline` function of Algorithm 3 returns only 69 baseline processes, and thus the resulting mixtures of e-SR and e-CUSUM procedures of Algorithm 1 can be computed efficiently in an online fashion.

The right plot in Figure 2 presents the log e-detector values using mixtures of e-SR (red) and e-CUSUM (green) procedures. Although there were a few months in which monthly win rates were higher than p_0 , overall log e-detectors remained at a stable level over the first four seasons. However, after the 2014–15 season starts, the log e-detectors increase rapidly and both procedures detect a changepoint during the 2014–15 season, which is the season that marked the return of LeBron James to the Cavaliers.

5.2 Mean-shift Detection in General Bounded Random Variables

Plus-Minus of the Cavaliers Revisited We return to the Cavaliers 2011–2018 example from Section 1.2. Let $\tilde{X}_1, \tilde{X}_2, \dots$ be the sequence of Plus-Minus stats from each game. We assume that the average Plus-Minus of the team is less than or equal to $\mu_{<} := -1$ before the changepoint (if any), while after the changepoint it is greater than $\mu_{>} := 1$. Here, the gap $|\mu_{>} - \mu_{<}|$ between averages of Plus-Minus in pre- and post-changes refers to the degree of improvement we consider as significant.

For convenience, we first normalize the observed sequence. We assume that the absolute value of each Plus-Minus is bounded by 80, meaning that no team beats another by over 80 points (such an extreme game has never

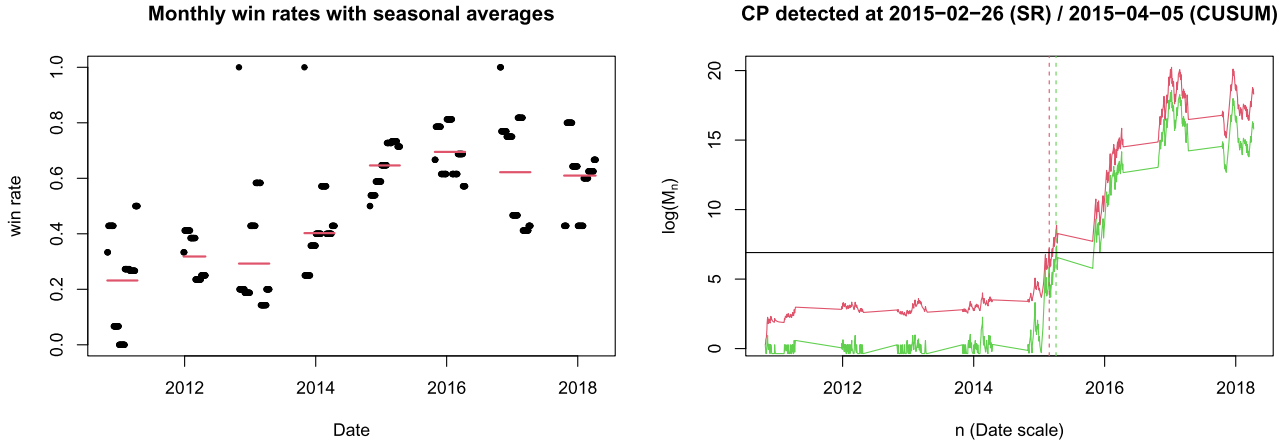


Figure 2: Left: Monthly win rates of the Cavaliers from 2010–11 to 2017–18 seasons (the raw data is Bernoulli, which is harder to visualize). Each red line corresponds to the seasonal average. Right: Paths of log e-detectors (SR: red; CUSUM: green). The horizontal line is the threshold (common to both procedures) equal to $\log(1/\alpha)$, ensuring that the ARL is at least $1/\alpha = 10^3$, larger than the number of games in 12 seasons (82 per season). The e-SR procedure detects a changepoint during the 2014–15 season.

happened in NBA history). Accordingly, define the normalized Plus-Minus, $X_n := (\tilde{X}_n + 80)/160 \in [0, 1]$ for each n . Then, the pre-change observations have conditional mean at most $m := (\mu_{<} + 80)/160 = 0.494$ and the minimum gap to detect is equal to $\delta := |\mu_{>} - \mu_{<}|/160 = 0.0125$.

Modeling Plus-Minus Stats as a Bounded Sequence with Time-varying, Dependent Means After the normalization above, the Plus-Minus stats form sequence of bounded random variables X_1, X_2, \dots on $[0, 1]$. Each observation may have different distribution (due to seasonal effects, injuries, form, etc.), but we assume that all observations before an unknown changepoint ν have a mean less than or equal to a known boundary $m \in (0, 1)$, when conditioned on the past sample history. That is, under any pre-change distribution P , we have $\mu_n := \mathbb{E}_{P,\infty}[X_n | \mathcal{F}_{n-1}] \leq m, \forall n \geq 1$. In other words, we use

$$\mathcal{P} := \{P : \mu_n \leq m, \forall n \geq 1\},$$

where other characteristics about P (outside of its sequence of conditional means) are irrelevant. But after the changepoint, all observations have (conditional) mean larger than the boundary m with the minimum gap equal to δ . Thus,

$$\mathcal{Q} := \{P : \mu_n \geq m + \delta, \forall n \geq 1\},$$

To build an e-SR or e-CUSUM procedure, we need to choose a baseline increment. To derive it, we first consider a simplified setting where both pre- and post-change observations are independently and identically distributed with $\mathbb{E}_{P,\infty}[X] \leq m$ and $\mathbb{E}_{0,Q}[X] \geq m + \delta$, respectively. In this simplified case, we simply refer P and Q to marginal pre- and post-change distributions and \mathcal{P} and \mathcal{Q} to their collections. Then, define $\text{KL}_{\text{inf}}(Q; m) := \inf_{P \in \mathcal{P}} \text{KL}(Q, P)$ to be

the smallest KL divergence between Q and \mathcal{P} . It is known (see, e.g., [9, 10]) that KL_{inf} has the following variational representation:

$$\begin{aligned} \text{KL}_{\text{inf}}(Q, m) &= \sup_{\lambda \in (0,1)} \mathbb{E}_{0,Q} \log \left(1 + \lambda \left(\frac{X}{m} - 1 \right) \right) \\ &=: D(Q || \mathcal{P}). \end{aligned} \tag{5.5}$$

Accordingly, for each $\lambda \in (0, 1)$, define the baseline increment $L^\lambda := \{L_n\}_{n \geq 1}$ as

$$L_n^\lambda := 1 + \lambda \left(\frac{X_n}{m} - 1 \right), \tag{5.6}$$

for each $n \in \mathbb{N}$. Though the baseline increment above has been derived in the simplified iid setting, it can be checked that L^λ is also a valid baseline increment for the general time-varying, dependent means case since it is nonnegative whenever $X_n, m \in [0, 1]$ as assumed in our setup, and for each pre-change distribution $P \in \mathcal{P}$, we have

$$\mathbb{E}_{P,\infty}[L_n^\lambda | \mathcal{F}_{n-1}] = 1 + \lambda \left(\frac{\mathbb{E}_{P,\infty}[X_n | \mathcal{F}_{n-1}]}{m} - 1 \right) \leq 1,$$

where the inequality comes from the condition $\mu_n \leq m$ for any pre-change distribution.

Interestingly, the baseline increments in (5.6) correspond to rescaled increments of the capital process used in [48] to design test martingales for confidence sequences of means of bounded random variables. Though the expressions are essentially identical, ours was obtained via a variational representation of the KL divergence between distributions of bounded random variables, while the derivation presented in [48] is based on a betting interpretation of hypothesis testing.

For any $Q \in \mathcal{Q}$, let λ^{op} be the optimal choice of $\lambda \in [0, 1]$ given by

$$\lambda^{\text{op}} = \arg \max_{\lambda \in [0,1]} \mathbb{E}_{0,Q} \log \left(1 + \lambda \left(\frac{X}{m} - 1 \right) \right). \quad (5.7)$$

Unfortunately, it is typically difficult to compute the optimal λ^{op} since it depends on the unknown post-change distribution Q in a complicated way. In this case, we use a sub-exponential lower bound from [5, 12], given by

$$\begin{aligned} \tilde{L}_n^\lambda &:= \exp \left\{ \lambda \left(\frac{X_n}{m} - 1 \right) - \psi_E(\lambda) \left(\frac{X_n}{m} - 1 \right)^2 \right\} \\ &\leq 1 + \lambda \left(\frac{X_n}{m} - 1 \right) = L_n^\lambda, \end{aligned} \quad (5.8)$$

where $\psi_E(\lambda) := -\log(1 - \lambda) - \lambda$ for $\lambda \in (0, 1)$. For each $\lambda \in (0, 1)$, the process \tilde{L}^λ is itself a valid exponential baseline increment with $s(x) := x/m - 1$ and $v(x) := (x/m - 1)^2$.

The lower bound in (5.8) also implies the lower bound

$$\text{KL}_{\text{inf}}(Q, m) \geq \sup_{\lambda \in [0,1]} \{ \lambda \mu - \psi_E(\lambda) \sigma^2 \} = \sigma^2 \psi_E^*(\Delta^{\text{op}}), \quad (5.9)$$

where $\psi_E^*(u) := u - \log(1 + u)$ is the convex conjugate of ψ_E , while μ, σ^2 and Δ^{op} from Section 4 are:

$$\begin{aligned} \mu &:= \mathbb{E}_{0,Q} s(X) = \frac{\mathbb{E}_{0,Q} X - m}{m}, \\ \sigma^2 &:= \mathbb{E}_{0,Q} v(X) = \frac{\mathbb{E}_{0,Q} (X - m)^2}{m^2}, \\ \Delta^{\text{op}} &:= \frac{\mu}{\sigma^2} = \frac{m[\mathbb{E}_{0,Q} X - m]}{\mathbb{E}_{0,Q} (X - m)^2}. \end{aligned}$$

Noting that $\psi_E^*(u) \approx u^2/2$ for small u , we see that for small $\Delta^{\text{op}} \ll 1$, one has

$$\text{KL}_{\text{inf}}(Q, m) \gtrsim \frac{[\mathbb{E}_{0,Q} X - m]^2}{2\mathbb{E}_{0,Q} (X - m)^2}. \quad (5.10)$$

Note that the oracle Δ^{op} depends on the unknown post-change distribution only via first and second moments. Therefore, in contrast to the original set of baseline increments $\{L^\lambda\}_{\lambda \in (0,1)}$, the exponential baseline increments $\{\tilde{L}^\lambda\}_{\lambda \in (0,1)}$ that lower bound them allow us to more easily set a range (Δ_L, Δ_U) to build mixtures of the e-SR and e-CUSUM procedures. For example, if we assume that the post-change distribution has mean at least $m + \delta$ for a positive δ then we can upper and lower bound Δ^{op} by

$$\Delta_L := \frac{m\delta}{(1 - m)^2} \leq \Delta^{\text{op}} \leq \frac{m(1 - m)}{\delta^2} =: \Delta_U. \quad (5.11)$$

Now, given Δ_L and Δ_U , we can use Algorithm 1 to run the mixture of e-SR or e-CUSUM procedure to detect the changepoint based on the exponential baseline increments

$\{\tilde{L}^\lambda\}_{\lambda \in (0,1)}$. It is also straightforward to build the corresponding mixtures of e-SR and e-CUSUM procedures for the original baseline increment $\{L^\lambda\}_{\lambda \in (0,1)}$ which is always more sample-efficient.

Implementation of Algorithm 1 and Its Results Recall that in the plus-minus stats running example, we use pre-change mean $m = 0.494$ and the minimum gap $\delta = 0.0125$, which bounds Δ^{op} by $\Delta_L := 0.024$ and $\Delta_U := \frac{m(1-m)}{\delta^2} = 1600$. As before, we choose $\alpha = 10^{-3}$ to make the ARL larger than 12 regular seasons and set the maximum number of baselines $K_{\text{max}} = 1000$. Based on these parameters, we can build mixtures of e-SR and e-CUSUM procedures. Though the difference between Δ_L and Δ_U may seem to be large, the actual number of baselines returned by the function `computeBaseline` in Algorithm 1 is 190, which is small enough to update the procedure efficiently on the fly.

Figure 3 shows e-detectors (left) and their logarithms (right). The horizontal line corresponds to the detection boundary given by $1/\alpha$ (left) and $\log \alpha^{-1}$ (right). Similar to the winning rate example, the log e-detectors remained stable during the first four regular seasons, although the difference between SR and CUSUM e-detectors is larger than before. After 2014–15 season started, both e-detectors increased rapidly, and the e-SR procedure detects a changepoint during the 2014–15 season, but e-CUSUM detects the changepoint only in the following season (as expected, since both procedures use the same threshold).

5.3 Simulation-based Comparison with Parametric Methods

In the Bernoulli example of Section 5.1, we showed that, in the simple i.i.d Bernoulli setup, our mixtures of e-SR and e-CUSUM procedures match the rate of the worst average delays $O(\log(1/\alpha)/\text{KL}(q||p_0))$ of the oracle sequential change detection procedure as $\alpha \rightarrow 0$. In this subsection, we conduct a simulation study to compare the efficiency of our e-SR procedure with the oracle CUSUM procedure with the exact threshold [21, 31], given by

$$N_{\text{CU}}^* := \inf \left\{ n \geq 1 : \max_{0 \leq j < n} \sum_{i=j}^n \log \frac{f_{p^*}(X_i)}{f_{p_0}(X_i)} \geq c_\alpha^* \right\}, \quad (5.12)$$

where p^* is the true post-change distribution parameter (hence the oracle designation) and c_α^* is the value of the threshold so that the ARL is exactly $1/\alpha$. It is well known that the oracle CUSUM procedure (with the appropriate choice of the stopping threshold that controls ARL exactly) minimizes the worst average delay.

We also compare our method with a version of the GLR




Figure 3: Left: *E*-detectors (SR: red; CUSUM: green) over eight seasons (82 games per season). Right: Logarithm of *e*-detectors against date (the sharp rise of *e*-SR is simply due to the log scale). In both plots, horizontal lines are thresholds equal to $1/\alpha$ (left) and $\log(1/\alpha)$ (right) controlling the ARL by $1/\alpha = 10^3$. The *e*-SR procedure detects a change during the 2014–15 season, while *e*-CUSUM takes longer (as expected).

procedure based on the stopping time

$$\begin{aligned}
 N_{\text{GL}}^* &:= \inf \left\{ n \geq 1 : \max_{0 \leq j < n} \sup_p \sum_{i=j}^n \log \frac{f_p(X_i)}{f_{p_0}(X_i)} \geq c_\alpha \right\} \\
 &= \inf \left\{ n \geq 1 : \max_{0 \leq j < n} \sum_{i=j}^n \log \frac{f_{\hat{p}_{j:n}}(X_i)}{f_{p_0}(X_i)} \geq c_\alpha \right\}, \tag{5.13}
 \end{aligned}$$

where each $\hat{p}_{j:n}$ is the MLE of the post-change parameter and the exact threshold c_α is tuned to control ARL exactly at $1/\alpha$ (this is typically only possible in such simple parametric settings, either by analytic derivations or simulations). Unlike the oracle CUSUM procedure, the GLR procedure does not have an iterative update rule, as we need to recompute the MLE of the post-change parameter at each time. As a result, its computational cost at time n is $O(n)$, which makes an online implementation very costly. In practice, we may want to use a window-limited GLR procedure to overcome the computational challenge. However, in our study, we deploy the GLR procedure to avoid the additional challenge of picking a window size.

Simulation Details Throughout this simulation, we draw pre-change observations as iid Bernoulli random variables with $p_0 = 0.5$, and post-change observations using $p_1 = 0.6$. For non-oracle methods, we will only assume that the post-change parameter is known to be in the interval $[0.51, 0.99]$. The *e*-SR and *e*-CUSUM procedures in Algorithm 1 will use this range to set $\Delta_L := 0.01$ and $\Delta_U := 0.49$. For the GLR procedure, the MLE of the post-change parameter is

$$\hat{p}_{j:n} := \min \{ \max \{ \bar{X}_{j:n}, 0.51 \}, 0.99 \}, \tag{5.14}$$

where $\bar{X}_{j:n}$ is the sample average over last $n - j + 1$ observations. Our ARL target is equal to $1/\alpha := 500$, and each

simulation is repeated 5000 times to estimate average delays. The time of the changepoint ν varies in $\{0, 100, \dots, 500\}$. For simplicity, each run will end no later than time $n = 1000$.

For the oracle CUSUM, GLR, and *e*-CUSUM procedures (but not *e*-SR), we use the same simulation setup to find the exact threshold value that controls the ARL exactly $1/\alpha$ for each method. For the *e*-SR procedure, we simply use the universal threshold of $\log(1/\alpha)$ to demonstrate its efficiency. In particular, practitioners are not required to resort to expensive simulation to identify a good threshold value.

Figure 4 shows the average delays of oracle CUSUM, GLR, *e*-CUSUM and *e*-SR procedures for each changepoint $\nu \in \{0, 100, \dots, 500\}$ (each experiment has only one change at ν). The vertical bar on each point represents 95% confidence interval of the average delay. As the theory guarantees, the oracle CUSUM procedure with an exact threshold results in the smallest worst average delay (91.3 ± 1.7), while surprisingly the GLR procedure with an exact threshold shows the largest worst average delay (123.7 ± 2.7) despite its high computational complexity. The two *e*-detectors perform reasonably well, and the *e*-SR detector in particular performs quite favorably overall despite using its conservative $\log(1/\alpha)$ threshold.

As the value of the changepoint approaches the ARL target of 500, the average delays tend to decrease sharply for *e*-SR procedure, even falling below the oracle CUSUM method. This is plausibly because *e*-SR sums the underlying *e*-processes, in contrast to the other CUSUM-style procedures which take the maximum of the underlying *e*-processes. We may also intuitively expect the oracle CUSUM delay to be relatively flat across ν because it is known to be minimax optimal, in the sense of minimizing the worst case delay, and minimax procedures often have constant “risk profiles”. Proving these fine-grained behaviors is beyond the scope of the current paper.




Figure 4: Average detection delay for each changepoint $\nu = 0, 100, \dots, 500$ (each experiment has exactly one changepoint at ν). Three of the methods use an exact threshold calculated via simulation (only possible in this simple, parametric example). Only the Oracle CUSUM method knows the post-change distribution. Even though e-SR uses the conservative $\log(1/\alpha)$ threshold, its detection delay is excellent, often even better than the Oracle CUSUM method (which has optimal average worst-case (across ν) delay).




Figure 5: Pre-change false alarm rates of detection methods for each changepoint $\nu = 0, 100, \dots, 500$ (each experiment has exactly one changepoint at ν). Three of the methods use an exact threshold calculated via simulation (only possible in this simple, parametric example). Only the Oracle CUSUM method knows the post-change distribution exactly. e-SR has the smallest false alert ratio (defined in the text).

Figure 5 illustrates the “pre-change false alarm rate”: the fraction of simulation runs in which the detection procedures stopped *before* the changepoint at ν (if a change occurs at $\nu = 0$, then it is zero by definition). This is not a common metric, since we provably control ARL at the target level. However, it is an interesting metric, so we plot it. As the time of the true changepoint ν becomes closer to the ARL target of 500, false alarm rates increase across all methods. We notice that the e-SR procedure with the $\log(1/\alpha)$ threshold results in the smallest false alarm rates in most cases.

6. DISCUSSION

6.1 Game-theoretic Interpretation of an E-detector

We briefly mention here a game-theoretic interpretation of an e-detector along the lines of the game-theoretic interpretations of martingales and supermartingales as the wealth of a gambler playing a fair game (well known since the time of [42]). We first summarize the game-theoretic interpretation of a \mathcal{P} -e-process, as described in [29].

The standard game-theoretic setup of [33] involves three players: a forecaster, a skeptic, and reality. The forecaster claims at the beginning that \mathcal{P} is a plausible model for the yet-to-be-observed data; meaning that the observations are in accordance with (or generated by) some $P \in \mathcal{P}$. The skeptic plays (in parallel) a family of games indexed by $P \in \mathcal{P}$ against nature and begins with one dollar in each game. The objective of the skeptic in the P -th game is to sequentially test whether P is a good explanation for the data by betting against P . At each time step, the skeptic places fair bets (relative to P , in the P -th game) about the next outcome. Then nature reveals the next outcome, and the skeptic’s wealth in every game is updated. The magnitude of the skeptic’s wealth in the P -th game is direct evidence against P being a good explanation; the higher the wealth, the more unlikely the data came from P . Thus in each game, the gambler places different bets, but nature’s moves (the outcomes) are identical across all games. The skeptic’s overall evidence against \mathcal{P} is measured by their *worst* wealth across all the games. If this evidence exceeds $1/\alpha$, it means that the skeptic multiplied their initial capital by at least $1/\alpha$ in every game, and if we reject \mathcal{P} when this happens, Ville’s inequality implies that we have a valid level- α sequential test.

Since our e-detectors are constructed to be cumulative sums of e-processes started at consecutive times, their game-theoretic interpretation builds on the aforementioned one. Informally, the forecaster not only claims that the data sequence follows \mathcal{P} from the start, but that this will not change after some amount of time. The skeptic now wishes to detect a change, if one occurs, as soon as possible. To accomplish this task, the skeptic is provided with one extra dollar every day that they invest (using a \mathcal{P} e-process) into testing whether the data from that day onwards is still explained well by \mathcal{P} . E-detectors use the wealth in all these games

(one against each $P \in \mathcal{P}$, starting at each time) as a measure of evidence against the forecaster’s claims. The SR e-detector uses the sum (across time) of the minimum wealth (across P at each time), though it could use the amount that this wealth exceeds n , which is the total dollar amount invested up to time n . The CUSUM e-detector uses the maximum (across time) of the minimum wealth (across P). These are only two ways of constructing e-detectors, and we leave other constructions to future work.

6.2 Viewing Lorden’s Reduction to Sequential Testing as an E-detector

Lorden [18] proposed a simple method to construct a change detection method with ARL control via a reduction to sequential testing. We describe this below, first defining a sequential test formally.

A *sequential test* ϕ is a mapping from increasing amounts of data to a sequence of zeros and ones, where a one represents a rejection of the null hypothesis, and a zero means “continue collecting data”. Formally, define the decision at time t as $\phi_t : \mathcal{X}^t \rightarrow \{0, 1\}$, and let $\phi := \{\phi_t\}_{t \geq 0}$ be the collection of such decisions made one at a time based on the first t datapoints, with $\phi_0 = 0$ by default. The sequence of tests ϕ is called a level- α sequential test for \mathcal{P} if

$$\sup_{P \in \mathcal{P}} P(\exists t \geq 1 : \phi_t(X_1, \dots, X_t) = 1) \leq \alpha,$$

i.e. if the probability of ever falsely rejecting the null is at most α . By convention, if $\phi_t = 1$, we set $\phi_s = 1$ for $s > t$. This is equivalent to requiring that, for each $P \in \mathcal{P}$,

$$P(\phi_\tau(X_1, \dots, X_\tau) = 1) \leq \alpha \text{ for any stopping time } \tau.$$

Let $\phi^{(j)}$ denote a sequential test is started at time j ; that is, for $\phi^{(j)}$, the first observation is actually X_j (and not X_1), but the test itself can depend on the first $j - 1$ observations (for example, we can choose our betting strategy based the first $j - 1$ points, even though our betting score will only be evaluated from time j onwards). Note that $\phi^{(1)}$ is simply a standard sequential test as defined above. Ideally, these tests are powerful against alternative \mathcal{Q} .

Lorden’s change detection procedure is simple and works as follows. At each time t , start a new sequential test $\phi^{(t)}$, in addition to the ones that are already running. In other words, consider a sequence of level- α sequential tests $\phi^{(1)}, \phi^{(2)}, \dots$, starting at consecutive times. Lorden declares a change if any of those sequential tests rejects the null \mathcal{P} :

$$N^{\text{Lorden}} := \inf \left\{ n \geq 1 : \max_{1 \leq j \leq n} \phi_n^{(j)}(X_{n-j+1}, \dots, X_n) = 1 \right\}.$$

Lorden proved that this method controls the ARL at $1/\alpha$ if the data are iid, and if the same test $\phi^{(j)}$ is deployed at each j (i.e. apart from the delayed start, the tests are identical).

We first observe that Lorden’s method is a special case of an e-detector. Indeed, with each level- α sequential test $\phi \equiv$

$\{\phi_t\}_{t \geq 0}$, we can associate an e-process $\Lambda_t^{\text{Lorden}} := \frac{1(\phi_t=1)}{\alpha} = \frac{\phi_t}{\alpha}$. Note that Λ only takes on two values: $0, 1/\alpha$, and when it reaches the latter, it stays there. Furthermore, note that $\mathbb{E}[\Lambda_\tau^{\text{Lorden}}] \leq 1$ for any stopping time τ , which makes it an e-process as claimed.

Last, note that if we form a ‘‘Lorden e-detector’’ using either the SR or CUSUM methods in (2.5), then both e-detectors start out at 0, and either e-detector jumps to level $1/\alpha$ if and only if one of the (delayed start) sequential tests rejects the null, and further our e-detector declares a changepoint at exactly the same instant that Lorden’s does. Thus, Lorden’s procedure can be subsumed within our e-detector framework without any loss of generality or performance.

There are two benefits to viewing Lorden’s method as an e-detector. First, we can dispense with both the aforementioned conditions that Lorden required to prove ARL control: the iid assumption, and the condition that the underlying tests $\phi^{(j)}$ are identical across j . Indeed, our main result guarantees ARL control for any e-detector no matter what the underlying e_j -processes are, or whether the data are iid or not.

Second, this viewpoint allows us to see why e-detectors could have much smaller detection delay than Lorden’s method (that is, Lorden’s e-detector). In Lorden’s e-detector, there is no sharing of evidence across different e_j -processes: each sequential test acts alone without help from the others, and we need a single e_j -process to reach $1/\alpha$ before we can declare a change. When a general (say SR-type) e-detector crosses $1/\alpha$, the reason it does so will usually be because of a collaborative effect across various e-processes (caused by the nontrivial sum of e_j -processes in the definition of the e-detector), none of which have yet individually reached $1/\alpha$. This will happen much sooner than any individual one reaches that threshold, causing an earlier detection than Lorden’s method. In fact, every level- α sequential test in some sense must be based on threshold an e-process at level $1/\alpha$ [28] and using our e-detector with those underlying e-processes will be more statistically efficient (shorter delay) than directly using Lorden’s reduction.

For the sake of future reference, we summarize the above observations below in a ‘‘generalized Lorden’s lemma’’, whose proof follows immediately from the properties of an e-detector and the discussion above.

Lemma 6.1 (Generalized Lorden’s Lemma). *Suppose the data initially come from a distribution in the pre-change class \mathcal{P} and, if a change occurs, they later come from a distribution in the post-change class \mathcal{Q} (note the lack of any iid assumption). For each j , let $\phi^{(j)}$ denote a (one-sided) level- α sequential test for \mathcal{P} against \mathcal{Q} that is started at time j (but need not be identical or related in any way to any other $\phi^{(k)}$, for $k \neq j$). If we declare a change at the first time when any one of these sequential tests rejects the null, the resulting change detection procedure has ARL at most $1/\alpha$. Further, this generalization of Lorden’s change detector is a special case of an e-detector.*

6.3 Future Directions

There remain a whole host of follow-up directions; we mention only a few below. First, our sequential change detection framework can be straightforwardly generalized to the multi-stream setting where we are monitoring a large number of data streams. In the classical parametric setting, minimum or summation of local CUSUM statistics for multi-stream data were proposed and their asymptotic optimality was studied [8, 20]. Since either minimum or scaled summation (average) of e-detectors also forms a valid e-detector, we can apply the framework in this paper to the multi-stream setting seamlessly. It is interesting to investigate how the framework can be even further generalized to structural multi-stream settings [50, 51, 2].

The kernel sequential change detection is an important class of sequential change detection methods [3, 16]. It is an interesting open direction how to instantiate existing kernel-based methods into our general framework to make it possible to analyze kernel sequential change detection algorithms in a nonasymptotic way. As it currently stands, neither framework is more general than the other, because the kernel methods often assume iid data before the changepoint, while we abstain from such strong assumptions.

Last, throughout this paper, we have only focused on detecting *whether* a changepoint happened or not but have not dealt with inferential questions surrounding *when* the change occurred. Future work could study how to perform such inference with e-detectors in our nonparametric settings, either online or post-hoc.

7. SUMMARY

We have presented a general framework for sequential change detection based on a new concept called e-detectors. The proposed framework is nonparametric as it does not rely on a parametric assumption on the data-generating distribution (though, when such assumptions are made, we recover well-known parametric methods as special cases). Also, the framework comes with nonasymptotic guarantees, since every component of the framework can be chosen and analyzed explicitly without any asymptotic approximations. By introducing additional structures such as baseline increments and exponential e-detectors on the top of the general framework, we can construct computationally and statistically efficient online algorithms that have explicit upper bounds on worst average delays. Finally, through examples involving Bernoulli and bounded random variables, we explained how one can apply the presented framework in practical settings, with NBA data serving as a running case study.

APPENDIX A. MAIN PROOFS

A.1 Proofs for Statements in Section 2

Proof of Theorem 2.4 (ARL control). For any given $\alpha \in (0, 1)$ let N^* be the stopping time defined by

$$N^* := \inf\{n \geq 1 : M_n \geq 1/\alpha\}. \quad (\text{A.1})$$

For any pre-change distribution $P \in \mathcal{P}$, we may assume $N^* < \infty$ with probability one under P , without loss of generality. (If not, we have $\mathbb{E}_{P,\infty} N^* = \infty$, which immediately proves the claim.) Then, from the definition of an e-detector, we have

$$\mathbb{E}_{P,\infty}[M_{N^*}] \leq \mathbb{E}_{P,\infty} N^*, \quad (\text{A.2})$$

which implies that, for any $P \in \mathcal{P}$,

$$\begin{aligned} \mathbb{E}_{P,\infty} N^* &\geq \mathbb{E}_{P,\infty}[M_{N^*}] \\ &= \mathbb{E}_{P,\infty}[M_{N^*} \mathbb{1}(N^* < \infty)] \\ &\geq \mathbb{E}_{P,\infty}\left[\frac{1}{\alpha} \mathbb{1}(N^* < \infty)\right] \quad (\text{by definition of } N^*) \\ &= \frac{1}{\alpha} \mathbb{P}_{P,\infty}(N^* < \infty) = \frac{1}{\alpha}, \end{aligned}$$

as desired. \square

Proof of Proposition 2.14 (Bounds on worst average delay). We first prove (2.12). Note that

$$\begin{aligned} N_{\text{SR}}^* &= \inf \left\{ n \geq 1 : \sum_{j=1}^n \prod_{i=j}^n L_i \geq 1/\alpha \right\} \\ &\leq \min_{j \geq 1} \inf \left\{ n \geq j : \prod_{i=j}^n L_i \geq 1/\alpha \right\} := \min_{j \geq 1} N_j. \end{aligned}$$

For each fixed $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$, since $N_j \perp\!\!\!\perp \mathcal{F}_\nu$ for all $j > \nu + m$, we have

$$\begin{aligned} &\mathbb{E}_{P,\nu,Q}[[N_{\text{SR}}^* - \nu]_+ | \mathcal{F}_\nu] \\ &\leq \min_{j \geq 1} \mathbb{E}_{P,\nu,Q}[[N_j - \nu]_+ | \mathcal{F}_\nu] \\ &\leq \min_{j > \nu + m} \mathbb{E}_{P,\nu,Q}[N_j - \nu | \mathcal{F}_\nu] \quad (\text{since } N_j \geq j > \nu + m) \\ &= \min_{j > \nu + m} \mathbb{E}_{P,\nu,Q}[N_j] - \nu \quad (\text{since } N_j \perp\!\!\!\perp \mathcal{F}_\nu, \forall j > \nu + m) \\ &= \min_{j > \nu + m} \mathbb{E}_{P,\nu,Q} \left[\inf \left\{ n \geq j : \prod_{i=j}^n L_i \geq 1/\alpha \right\} \right] - \nu \\ &= \min_{j - \nu > m} \mathbb{E}_{0,Q} \left[\inf \left\{ n \geq j - \nu : \prod_{i=j-\nu}^n L_i \geq 1/\alpha \right\} \right] \\ &= \min_{j' > m} \mathbb{E}_{0,Q} \left[\inf \left\{ n \geq j' : \prod_{i=j'}^n L_i \geq 1/\alpha \right\} \right] \quad (\text{by } j' := j - \nu) \\ &= \min_{j' > m} \mathbb{E}_{0,Q} \left[\inf \left\{ n \geq 1 : \prod_{i=1}^n L_i \geq 1/\alpha \right\} \right] + j' - 1 \\ &= \min_{j' > m} \mathbb{E}_{0,Q} N_{1/\alpha} + j' - 1 \\ &= \mathbb{E}_{0,Q} N_{1/\alpha} + m, \end{aligned}$$

where the third equality comes from the fact that the distribution of $X_{j-m}, X_{j-m+1}, \dots$ under $\mathbb{P}_{P,\nu,Q}$ is equal to the one of $X_{j-\nu-m}, X_{j-\nu-m+1}, \dots$ under $\mathbb{P}_{0,Q}$ provided that

$j - m > \nu$, and the fifth equality is based on the strong stationarity of the post-change observations. Since the last term does not depend on P, ν or \mathcal{F}_ν , we obtain the claimed result

$$\begin{aligned} \mathcal{J}_P(N_{\text{SR}}^*) &\leq \mathcal{J}_L(N_{\text{SR}}^*) \\ &= \sup_{P \in \mathcal{P}, \nu \geq 0} \text{esssup} \mathbb{E}_{P,\nu,Q}[[N_{\text{SR}}^* - \nu]_+ | \mathcal{F}_\nu] \\ &\leq \mathbb{E}_{0,Q} N_{1/\alpha} + m, \end{aligned}$$

as desired.

To prove (2.13), first note that

$$\begin{aligned} N_{\text{CU}}^* &= \inf \left\{ n \geq 1 : \max_{j \in [n]} \prod_{i=j}^n L_i \geq c_\alpha \right\} \\ &\leq \min_{j \geq 1} \inf \left\{ n \geq j : \prod_{i=j}^n L_i \geq c_\alpha \right\} \\ &:= \min_{j \geq 1} N_j. \end{aligned}$$

The remaining part of the proof for (2.13) is followed by the same argument for (2.12).

Finally, to prove (2.14), it is enough to show the following inequality holds:

$$\mathbb{E}_{0,Q} N(g) \leq \frac{g}{\mathbb{E}_{0,Q} \log L_1} + \frac{\mathbb{V}_{0,Q} \log L_1}{[\mathbb{E}_{0,Q} \log L_1]^2} + 1, \quad \forall g > 0, \quad (\text{A.3})$$

where $N(g) := \inf\{n \geq 1 : \sum_{i=1}^n \log L_i \geq g\}$ for each $g > 0$. The proof of the above upper bound (A.3) is based on the Lorden's inequality [17] which can be stated as follows:

Fact A.1 (Lorden's inequality [17]). *Suppose X_1, X_2, \dots are i.i.d. samples with $\mathbb{E}X_1 = \mu > 0$ and $\mathbb{E}X_1^2 < \infty$. For each $g > 0$, set $N(g) := \inf\{n : S_n := \sum_{i=1}^n X_i \geq g\}$ and $R_g := S_{N(g)} - g$. Then, the following inequality holds:*

$$\sup_{g > 0} \mathbb{E}[R_g] \leq \frac{\mathbb{E}X_1^2}{\mu} = \mu + \frac{\sigma^2}{\mu}, \quad (\text{A.4})$$

where $\sigma^2 = \mathbb{V}X_1$.

Now, to prove the upper bound (A.3), fix a constant $g > 0$. Since $\mathbb{E}_{0,Q} \log L_1 > 0$, we have $\mathbb{E}_{0,Q} N(g) < \infty$. Therefore, by Wald's equation,

$$\mathbb{E}_{0,Q} \log L_1 \mathbb{E}_{0,Q} N(g) = \mathbb{E}_{0,Q} \left[\sum_{i=1}^{N(g)} \log L_i \right]. \quad (\text{A.5})$$

For each $g > 0$, set $R_g := \sum_{i=1}^{N(g)} \log L_i - g$. Then, from the Lorden's inequality, we have

$$\begin{aligned} \mathbb{E}_{0,Q} \log L_1 \mathbb{E}_{0,Q} N(g) &= \mathbb{E}_{0,Q} \left[\sum_{i=1}^{N(g)} \log L_i \right] \\ &\leq c + \sup_{g > 0} \mathbb{E}_{0,Q}[R_g] \end{aligned}$$

$$\leq g + \frac{\mathbb{V}_{0,Q} \log L_1}{\mathbb{E}_{0,Q} \log L_1} + \mathbb{E}_{0,Q} \log L_1, \quad (\text{A.6})$$

where the first inequality comes from the definition of $N(g)$. By multiplying $1/\mathbb{E}_{0,Q} \log L_1$ on both sides of the inequality (A.6), we have the claimed upper bound, completing the proof. \square

A.2 Proofs for Statements in Section 3

Proof of Proposition 3.3 (Validity of adaptive e-detectors). To see adaptive SR and CUSUM e-detectors are actually valid e-detectors, first note that $M_n^{\text{aCU}} \leq M_n^{\text{aSR}}$ for each $n \in \mathbb{N}$. Therefore, it is enough to show that $\mathbb{E}_{P,\infty} M_\tau^{\text{aSR}} \leq \mathbb{E}_{P,\infty} \tau$ for any stopping time τ and pre-change distribution $P \in \mathcal{P}$. If $\mathbb{P}_{P,\infty}(\tau = \infty) > 0$ then the above inequality holds trivially. Otherwise if $\mathbb{P}_{P,\infty}(\tau = \infty) = 0$, we have that

$$\begin{aligned} \mathbb{E}_{P,\infty} M_\tau^{\text{aSR}} &= \mathbb{E}_{P,\infty} \sum_{k=1}^{K(\tau)} \omega_k \sum_{j=K^{-1}(k)}^{\tau} \gamma_j \prod_{i=j}^{\tau} L_i(k) \\ &:= \mathbb{E}_{P,\infty} \sum_{k=1}^{K(\tau)} \omega_k \sum_{j=K^{-1}(k)}^{\tau} \gamma_j \Lambda_\tau^{(j)}(k) \\ &= \mathbb{E}_{P,\infty} \sum_{j=1}^{\tau} \sum_{k=1}^{K(j)} \gamma_j \omega_k \Lambda_\tau^{(j)}(k) \\ &= \sum_{j=1}^{\infty} \sum_{k=1}^{K(j)} \gamma_j \omega_k \mathbb{E}_{P,\infty} \Lambda_\tau^{(j)}(k) \mathbb{1}(j \leq \tau) \\ &= \sum_{j=1}^{\infty} \sum_{k=1}^{K(j)} \gamma_j \omega_k \mathbb{E}_{P,\infty} \mathbb{1}(j \leq \tau) \mathbb{E}_{P,\infty} [\Lambda_\tau^{(j)}(k) \mid \mathcal{F}_{j-1}] \\ &\leq \sum_{j=1}^{\infty} \sum_{k=1}^{K(j)} \gamma_j \omega_k \mathbb{E}_{P,\infty} \mathbb{1}(j \leq \tau) \\ &= \sum_{j=1}^{\infty} \mathbb{E}_{P,\infty} \mathbb{1}(j \leq \tau) \\ &= \mathbb{E}_{P,\infty} \tau, \end{aligned}$$

as desired. Above, the sole inequality follows since each $\Lambda^{(j)}(k)$ is an e_j -process, and the following equality invokes the definition of γ_j for each j . \square

Proof of Theorem 3.4 (Delay bounds for adaptive e-detectors). We first prove the upper bound for the adaptive e-SR procedure in (3.17). Note that

$$\begin{aligned} N_{\text{aSR}}^* &= \inf \left\{ n \geq 1 : \sum_{k=1}^{K(n)} \omega_k \sum_{j=K^{-1}(k)}^n \gamma_j \prod_{i=j}^n L_i(k) \geq 1/\alpha \right\} \\ &\leq \inf \left\{ n \geq 1 : \sum_{k=1}^{K(n)} \omega_k \sum_{j=K^{-1}(k)}^n \prod_{i=j}^n L_i(k) \geq 1/\alpha \right\} \end{aligned}$$

$$\begin{aligned} &= \inf \left\{ n \geq 1 : \sum_{j=1}^n \sum_{k=1}^{K(j)} \omega_k \prod_{i=j}^n L_i(k) \geq 1/\alpha \right\} \\ &\leq \min_{j \geq 1} \inf \left\{ n \geq j : \sum_{k=1}^{K(j)} \omega_k \prod_{i=j}^n L_i(k) \geq 1/\alpha \right\} \\ &:= \min_{j \geq 1} N_j, \end{aligned}$$

where the first inequality follows because $\gamma_j \geq 1$ for each j . For each fixed $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$, since $N_j \perp\!\!\!\perp \mathcal{F}_\nu$ for all $j > \nu + m$, we have

$$\begin{aligned} \mathbb{E}_{P,\nu,Q} [[N_{\text{aSR}}^* - \nu]_+ \mid \mathcal{F}_\nu] &\leq \min_{j > \nu + m} \mathbb{E}_{P,\nu,Q} [N_j - \nu \mid \mathcal{F}_\nu] \quad (\text{since } N_j \geq j > \nu + m) \\ &= \min_{j > \nu + m} \mathbb{E}_{P,\nu,Q} \left[\inf \left\{ n \geq j : \sum_{k=1}^{K(j)} \omega_k \prod_{i=j}^n L_i(k) \geq 1/\alpha \right\} \right] - \nu \\ &\quad (\text{since } N_j \perp\!\!\!\perp \mathcal{F}_\nu, \forall j > \nu + m) \\ &= \min_{j - \nu > m} \mathbb{E}_{0,Q} \left[\inf \left\{ n \geq j - \nu : \sum_{k=1}^{K(j)} \omega_k \prod_{i=j-\nu}^n L_i(k) \geq 1/\alpha \right\} \right] \\ &= \min_{j' > m} \mathbb{E}_{0,Q} \left[\inf \left\{ n \geq j' : \sum_{k=1}^{K(j'+\nu)} \omega_k \prod_{i=j'}^n L_i(k) \geq 1/\alpha \right\} \right] \\ &\quad (\text{by setting } j' := j - \nu) \\ &= \min_{j' > m} \mathbb{E}_{0,Q} \left[\inf \left\{ n \geq 1 : \sum_{k=1}^{K(j'+\nu)} \omega_k \prod_{i=1}^n L_i(k) \geq 1/\alpha \right\} \right] \\ &\quad + j' - 1 \\ &\leq \min_{j' > m} \mathbb{E}_{0,Q} \left[\inf \left\{ n \geq 1 : \sum_{k=1}^{K(j')} \omega_k \prod_{i=1}^n L_i(k) \geq 1/\alpha \right\} \right] \\ &\quad + j' - 1 \\ &= \min_{j' > m} \mathbb{E}_{0,Q} N_{1/\alpha}(j') + j' - 1, \end{aligned}$$

where the second equality comes from the fact that the distribution of $X_{j-m}, X_{j-m+1}, \dots$ under $\mathbb{P}_{P,\nu,Q}$ is equal to the one of $X_{j-\nu-m}, X_{j-\nu-m+1}, \dots$ under $\mathbb{P}_{0,Q}$ provided by $j - m > \nu$ and the forth equality is based on the strong stationarity of post-change observations. Since the last term does not depend on P, ν nor \mathcal{F}_ν , we have the claimed result:

$$\begin{aligned} \mathcal{J}_P(N_{\text{aSR}}^*) &\leq \mathcal{J}_L(N_{\text{aSR}}^*) \\ &= \sup_{P \in \mathcal{P}, \nu \geq 0} \text{esssup} \mathbb{E}_{P,\nu,Q} [[N_{\text{aSR}}^* - \nu]_+ \mid \mathcal{F}_\nu] \\ &\leq \min_{j > m} [\mathbb{E}_{0,Q} N_{1/\alpha}(j) + j - 1], \end{aligned}$$

as desired.

To prove the adaptive e-CUSUM procedure case in (3.18),

first note that

$$\begin{aligned}
 N_{\text{aCU}}^* &= \inf \left\{ n \geq 1 : \sum_{k=1}^{K(n)} \omega_k \max_{K^{-1}(k) \leq j \leq n} \gamma_j \prod_{i=j}^n L_i(k) \geq c_\alpha \right\} \\
 &\leq \inf \left\{ n \geq 1 : \max_{j \leq n} \sum_{k=1}^{K(j)} \omega_k \prod_{i=j}^n L_i(k) \geq c_\alpha \right\} \\
 &\leq \min_{j \geq 1} \inf \left\{ n \geq j : \sum_{k=1}^{K(j)} \omega_k \prod_{i=j}^n L_i(k) \geq c_\alpha \right\} \\
 &:= \min_{j \geq 1} N_j,
 \end{aligned}$$

where the first inequality comes from $\gamma_j \geq 1$ with the fact $K^{-1}(k) \leq j$ for any $k \leq K(j)$. The remaining part of the proof of (3.18) follows the same argument used to obtain (3.17). \square

APPENDIX B. REMAINING PROOFS

B.1 Proofs for Statements in Section 4

Proof of Proposition 4.2. To simplify notation, we drop the subscripts $\{0, Q\}$ and 1. The claim reduces to proving that

$$\mathbb{E} \log L^{(\lambda^{\text{op}})} = \psi^*(\Delta^{\text{op}}) \sigma^2, \quad (\text{B.1})$$

where $\log L^{(\lambda^{\text{op}})} := \lambda^{\text{op}} s(X) - \psi(\lambda^{\text{op}}) v(X)$ and $\Delta^{\text{op}} = \nabla \psi(\lambda^{\text{op}})$. To prove the equality, first note that from the definition of the convex conjugate ψ^* of ψ , we have

$$\psi^*(\Delta^{\text{op}}) := \sup_{\lambda \in \Pi} \{\lambda \Delta^{\text{op}} - \psi(\lambda)\} = \lambda^{\text{op}} \Delta^{\text{op}} - \psi(\lambda^{\text{op}}), \quad (\text{B.2})$$

which implies

$$\begin{aligned}
 \mathbb{E} \log L^{(\lambda^{\text{op}})} &= \lambda^{\text{op}} \mu - \psi(\lambda^{\text{op}}) \sigma^2 = \sigma^2 [\lambda^{\text{op}} \Delta^{\text{op}} - \psi(\lambda^{\text{op}})] \\
 &= \sigma^2 \psi^*(\Delta^{\text{op}}),
 \end{aligned}$$

as desired. \square

Proof of Theorem 4.3. We first recall the definition of the stopping time in (4.8):

$$\bar{N}_g := \inf \left\{ n \geq 1 : \sup_{\lambda \in (\lambda_L, \lambda_U)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g \right\}, \quad g > 0. \quad (\text{B.3})$$

Then, the same argument used in the proof Proposition 2.14 immediately implies that, if the post-change observations are i.i.d. from Q , then, for any $g > 0$,

$$\mathbb{E}_{0,Q} \bar{N}_g \leq \frac{g}{D(Q||\mathcal{P})} + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda^{\text{op}})}]}{[D(Q||\mathcal{P})]^2} + 1. \quad (\text{B.4})$$

The claim of the theorem follows from Lemma B.1, whose statement and proof are given below. \square

Lemma B.1. *Let $N_{1/\alpha}$ and N_{c_α} be stopping times where the underlying mixing weights $\{\omega_k\}$ and parameters of baseline increments $\{\lambda_k\}$ are chosen via Algorithm 3. Let \bar{N}_{g_α} be the stopping time defined in (4.8) with the threshold given by Algorithm 3. Then, for any stream of observations X_1, X_2, \dots ,*

$$N_{c_\alpha} \leq N_{1/\alpha} \leq \bar{N}_{g_\alpha}, \quad (\text{B.5})$$

deterministically, provided that $1 < c_\alpha < 1/\alpha$.

Proof of Lemma B.1 and Algorithm 3. Throughout this proof, we set $D_L := \psi^*(\Delta_L) < \psi^*(\Delta_U) =: D_U$. The first inequality $N_{c_\alpha} \leq N_{1/\alpha}$ follows directly from the definition of the stopping time in (3.22) along with the condition that $c_\alpha \leq 1/\alpha$. To prove the second inequality $N_{1/\alpha} \leq \bar{N}_{g_\alpha}$, we will exploit on general geometric construction introduced in [36] to analyze the performance of sequential generalized likelihood ratio tests. To that effect, set $\hat{\mu}_n := S_n/V_n$. Then, for each fixed $\lambda > 0$ such that $\Delta = \nabla \psi(\lambda)$, Proposition 4.2 and the identity $\lambda = \nabla \psi^*(\Delta)$ imply that the function $\hat{\mu}_n \mapsto V_n^{-1} \sum_{i=1}^n \log L_i^{(\lambda)} = \lambda \hat{\mu}_n - \psi(\lambda) = \lambda(\hat{\mu}_n - \Delta) + \psi^*(\Delta)$ is a mapping of $\hat{\mu}_n$ into the tangent line of the function $z \mapsto \psi^*(z)$ at $z = \Delta$. Next, set $V_U := g_\alpha/D_U$ and $V_L := g_\alpha/D_L$, and define the set

$$R := \{(z, y) \in [0, \infty)^2 : y \leq \psi^*(z)\}. \quad (\text{B.6})$$

Then, the stopping event of \bar{N}_{g_α} can be expressed as

$$\begin{aligned}
 &\left\{ \exists n \geq 1 : \sup_{\lambda \in (\lambda_L, \lambda_U)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g_\alpha \right\} \\
 &= \left\{ \exists n \geq 1 : V_n < V_U, \left(\hat{\mu}_n, \frac{g_\alpha}{V_n} \right) \in H(\Delta_U) \right\} \\
 &\quad \cup \left\{ \exists n \geq 1 : V_n \geq V_U, \left(\hat{\mu}_n, \frac{g_\alpha}{V_n} \right) \in R \setminus H(\Delta_L) \right\} \\
 &= \left\{ \exists n \geq 1 : V_n < V_U, \sum_{i=1}^n \log L_i^{(\lambda_U)} \geq g_\alpha \right\} \\
 &\quad \cup \left\{ \exists n \geq 1 : V_n \geq V_U, \sup_{\lambda > \lambda_L} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g_\alpha \right\},
 \end{aligned}$$

$H(\Delta_U)$ and $H(\Delta_L)$ are half spaces contained in and tangent to R at $(\Delta_U, g_\alpha/V_U)$ and $(\Delta_L, g_\alpha/V_L)$, respectively. See Figure 6 for an illustration of the stopping event of \bar{N}_{g_α} .

Note that the first decomposition part $\{\exists n \geq 1 : V_n < V_U, \sum_{i=1}^n \log L_i^{(\lambda_U)} \geq g_\alpha\}$ is nonempty only if $V_U > \min_x v(x) := v_{\min}$, which is equivalent to $g_\alpha > v_{\min} D_U$. For the second part, a straightforward extension of Lemma 1 in the appendix of [36] implies that, for any fixed $\eta > 1$, the second part can be further decomposed by sets of simple events as follows:

$$\left\{ \exists n \geq 1 : V_n \geq V_U, \sup_{\lambda > \lambda_L} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g_\alpha \right\}$$




Figure 6: Illustration of the stopping event of \bar{N}_{g_α} defined in (4.8), and related regions $H(\Delta_U)$, $H(\Delta_L)$ and R . The stopping time \bar{N}_{g_α} is the first time when $(\hat{\mu}_n, g_\alpha/V_n)$ is located in one of the colored areas.

$$\begin{aligned} & \subset \bigcup_{k=1}^{K(\eta)-1} \left\{ \exists n \geq 1: V_n \in [V_U \eta^{k-1}, V_U \eta^k], \sum_{i=1}^n \log L_i^{(\lambda_k)} \geq g_\alpha/\eta \right\} \\ & \cup \left\{ \exists n \geq 1: V_n \geq V_U \eta^{K(\eta)-1}, \sum_{i=1}^n \log L_i^{(\lambda_{K(\eta)})} \geq g_\alpha/\eta \right\}, \end{aligned} \quad (\text{B.7})$$

where $K(\eta)$ is a positive integer defined by

$$K(\eta) := \left\lceil \log_\eta \left(\frac{D_U}{D_L} \right) \right\rceil, \quad (\text{B.8})$$

and, for $k = 1, \dots, K(\eta) - 1$, λ_k is given by $\lambda_k := \psi^*(\Delta_k)$, with Δ_k the solution with respect to $z > 0$ of the equation

$$\psi^*(z) = \frac{D_U}{\eta^k}, \quad (\text{B.9})$$

while $\Delta_{K(\eta)} := \Delta_L$. It can be checked that $\lambda_U := \lambda_0 > \lambda_1 > \lambda_2 > \dots > \lambda_{K(\eta)} = \lambda_L$. Decomposing the stopping event of \bar{N}_{g_α} , we can lower bound the stopping time \bar{N}_{g_α} for any $\eta > 1$ as:

$$\begin{aligned} \bar{N}_{g_\alpha} &= \inf \left\{ n \geq 1: \sup_{\lambda \in (\lambda_L, \lambda_U)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g_\alpha \right\} \\ &\geq \inf \left\{ n \geq 1: e^{-g_\alpha} \prod_{i=1}^n L_i^{(\lambda_0)} \mathbb{1}(g_\alpha > v_{\min} D_U) \right. \\ &\quad \left. + \sum_{k=1}^{K(\eta)} e^{-g_\alpha/\eta} \prod_{i=1}^n L_i^{(\lambda_k)} \geq 1 \right\} \\ &= \inf \left\{ n \geq 1: \alpha^{-1} e^{-g_\alpha} \mathbb{1}(g_\alpha > v_{\min} D_U) \right. \\ &\quad \times \exp\{\lambda_0 S_n - \psi(\lambda_0) V_n\} \\ &\quad \left. + \sum_{k=1}^{K(\eta)} \alpha^{-1} e^{-g_\alpha/\eta} \exp\{\lambda_k S_n - \psi(\lambda_k) V_n\} \geq 1/\alpha \right\} \end{aligned}$$

$$\begin{aligned} &:= \inf \left\{ n \geq 1: \sum_{k=0}^{K(\eta)} \omega_k(\eta) \exp\{\lambda_k S_n - \psi(\lambda_k) V_n\} \geq 1/\alpha \right\} \\ &:= N(\eta), \end{aligned}$$

where $\omega_0(\eta) := \alpha^{-1} e^{-g_\alpha} \mathbb{1}(g_\alpha > v_{\min} D_U)$ and $\omega_k(\eta) := \alpha^{-1} e^{-g_\alpha/\eta}$ for each $\eta > 1$ and $k = 1, \dots, K(\eta)$. As a quick remark, note that $\omega_0(\eta)$ does not depend on η and each $\omega_k(\eta)$ in fact does not depend on the index k but we use this notation just for consistency.

Finally, set $\eta_\alpha := (\frac{D_U}{D_L})^{1/K_\alpha}$ where K_α is the integer defined in (B.11) of Algorithm 3. It can be easily checked that $K(\eta_\alpha) = K_\alpha$. Therefore, once we choose $\eta = \eta_\alpha$ then, from the definition of the mixing weights in Algorithm 3, we conclude that $N(\eta_\alpha) \geq N_{1/\alpha}$ provided that

$$W := e^{-g_\alpha} \mathbb{1}(g_\alpha > v_{\min} D_U) + K_\alpha e^{-g_\alpha/\eta_\alpha} \leq \alpha, \quad (\text{B.13})$$

where the constant g_α is the constant used in Algorithm 3 and given by

$$\begin{aligned} g_\alpha &:= \inf \left\{ g > \log(1/\alpha) : e^{-g} \mathbb{1}(g > v_{\min} D_U) \right. \\ &\quad \left. + \min_{k \in [K_{\max}]} k \exp \left\{ -g \left(\frac{D_U}{D_L} \right)^{-1/k} \right\} \leq \alpha \right\}. \end{aligned} \quad (\text{B.14})$$

By the definition of K_α , we can immediately check that the inequality (B.13) holds, which proves the claimed inequality $N_{1/\alpha} \leq N(\eta_\alpha) \leq \bar{N}_{g_\alpha}$, as desired. \square

We conclude this section with a formal proof of the validity of the upper bound in (4.7).

Proposition B.2. *Let K_{\max} be a large enough integer such that*

$$\begin{aligned} K_\alpha &= \arg \min_{k \in \{1, \dots, K_{\max}\}} k \exp \left\{ -g_\alpha \left(\frac{D_U}{D_L} \right)^{-1/k} \right\} \\ &= \arg \min_{k \in \mathbb{N}} k \exp \left\{ -g_\alpha \left(\frac{D_U}{D_L} \right)^{-1/k} \right\}. \end{aligned} \quad (\text{B.15})$$

Then, the quantity g_α Algorithm 1 specified in (B.10) is such that

$$g_\alpha < \inf_{\eta > 1} \eta \left[\log(1/\alpha) + \log \left(1 + \left\lceil \log_\eta \frac{\psi^*(\Delta_U)}{\psi^*(\Delta_L)} \right\rceil \right) \right]. \quad (\text{B.16})$$

Proof of Proposition B.2. Once K_{\max} is large enough to satisfy (B.15) then the constant g_α can be written as

$$\begin{aligned} g_\alpha &= \inf \left\{ g > \log(1/\alpha) : e^{-g} \mathbb{1}(g > v_{\min} D_U) \right. \\ &\quad \left. + \inf_{\eta > 1} \left[\log_\eta \left(\frac{D_U}{D_L} \right) \right] e^{-g/\eta} \leq \alpha \right\}, \end{aligned} \quad (\text{B.17})$$

Algorithm 3: Pseudo-code of computeBaseline function.

Input: ARL parameter $\alpha \in (0, 1)$, Boundary values $0 < \Delta_L < \Delta_U$,
Maximum number of baselines $K_{\max} \in \mathbb{N}$.

Output: Parameters of baseline increments $\lambda_U = \lambda_0 > \lambda_1 > \dots > \lambda_{K_\alpha} = \lambda_L$,
Mixing weights $\omega_0, \omega_1, \dots, \omega_{K_\alpha} \in [0, 1]$,
Auxiliary values used to compute baseline increments $\{g_\alpha, K_\alpha, \eta_\alpha, W\}$.

- 1 Compute parameters for boundary values by $\lambda_L := \nabla\psi^*(\Delta_L)$ and $\lambda_U := \nabla\psi^*(\Delta_U)$.
/* If the separation is large enough, use a single baseline. */
- 2 **if** $\log(1/\alpha) \leq v_{\min}\psi^*(\Delta_L)$ **then**
- 3 Set $K_\alpha := 1$, $\lambda_1 := \lambda_L$ and $\omega_1 := 1$
- 4 **return** Parameter λ_1 and mixing weight ω_1
- 5 Compute the threshold $g_\alpha > \log(1/\alpha)$ given by

$$g_\alpha := \inf \left\{ g > \log(1/\alpha) : e^{-g} \mathbb{1}(g > v_{\min} D_U) + \min_{k \in [K_{\max}]} k \exp \left\{ -g \left(\frac{D_U}{D_L} \right)^{-1/k} \right\} \leq \alpha \right\}. \quad (\text{B.10})$$
- (See Algorithm 4 for an explicit way to compute it)
- 6 Compute the number of baselines $K_\alpha \in \mathbb{N}$ by

$$K_\alpha = \arg \min_{k \in [K_{\max}]} k \exp \left\{ -g_\alpha \left(\frac{D_U}{D_L} \right)^{-1/k} \right\}, \quad (\text{B.11})$$
- where $D_L := \psi^*(\Delta_L) < \psi^*(\Delta_U) =: D_U$.
- 7 Compute the spacing parameter $\eta_\alpha := (D_U/D_L)^{1/K_\alpha}$.
/* Compute parameters of baseline increments and mixing weights */
- 8 Set $\lambda_0 := \lambda_U$ and $\lambda_{K_\alpha} := \lambda_L$.
- 9 **if** $K_\alpha \geq 2$ **then**
- 10 **for** $k = 1, \dots, K_\alpha - 1$ **do**
- 11 Compute Δ_k as the solution of the equation $\psi^*(z) = D_U \eta^{-k}$ with respect to $z > 0$.
- 12 Compute the k -th parameter as $\lambda_k := \nabla\psi^*(\Delta_k)$.
- 13 Set $W := e^{-g_\alpha} \mathbb{1}(g_\alpha > v_{\min} D_U) + K_\alpha e^{-g_\alpha/\eta_\alpha}$ and compute mixing weights by

$$\omega_0 = W^{-1} e^{-g_\alpha} \mathbb{1}(g_\alpha > v_{\min} D_U) \text{ and } \omega_k = W^{-1} e^{-g_\alpha/\eta_\alpha}, \quad \forall k \in [K_\alpha]. \quad (\text{B.12})$$
- 14 **return** $\{\lambda_0, \lambda_1, \dots, \lambda_{K_\alpha}\}, \{\omega_0, \omega_1, \dots, \omega_{K_\alpha}\}, \{g_\alpha, K_\alpha, \eta_\alpha, W\}$

because the following equality holds for each $g > 0$ and $D_U > D_L$.

$$\min_{k \in \mathbb{N}} k \exp \left\{ -g \left(\frac{D_U}{D_L} \right)^{-1/k} \right\} = \inf_{\eta > 1} \left[\log_\eta \left(\frac{D_U}{D_L} \right) \right] e^{-g/\eta}. \quad (\text{B.18})$$

Finally, to prove the claimed bound in (B.16), first note that from (B.17), we have $g_\alpha \leq g(\eta)$ where $g(\eta) > 0$ is given by

$$g(\eta) := \inf \left\{ g > \log(1/\alpha) : e^{-g} + \left[\log_\eta \left(\frac{D_U}{D_L} \right) \right] e^{-g/\eta} \leq \alpha \right\} \\ \leq \eta \left[\log(1/\alpha) + \log \left(1 + \left[\log_\eta \frac{D_U}{D_L} \right] \right) \right],$$

for each $\eta > 1$. By taking infimum over $\eta > 1$, we get the upper bound in (B.16), as desired. \square

Proof of Lemma 4.5. The proof of Lemma 4.5 is similar to the one of Lemma B.1 except the previous threshold g_α being replaced with $g(V_0 \eta^{K(j)})$. Also note that, in this proof,

the terms Δ_0 and V_0 play a similar role of Δ_U and V_U in the previous proof of Lemma B.1.

As same as the Lemma B.1 case, set $S_n := \sum_{i=1}^n s(X_i)$, $V_n := \sum_{i=1}^n v(X_i)$, and $\hat{\mu}_n := S_n/V_n$. Then, for each fixed $\lambda > 0$ with $\Delta = \nabla\psi(\lambda)$, the function $\hat{\mu}_n \mapsto V_n^{-1} \sum_{i=1}^n \log L_i^{(\lambda)} = \lambda \hat{\mu}_n - \psi(\lambda) = \lambda(\hat{\mu}_n - \Delta) + \psi^*(\Delta)$ is a mapping of $\hat{\mu}_n$ to the tangent line of the function $z \mapsto \psi^*(z)$ at $z = \Delta$. Now, since the boundary function g is non-decreasing, the stopping event of $\bar{N}_g(j)$ can be bounded as

$$\left\{ \exists n \geq 1 : \sup_{\lambda \in (\lambda_{K(j)}, \lambda_0)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g(V_0 \eta^{K(j)}) \right\} \\ \subset \left\{ \exists n \geq 1 : V_n < V_0, \sum_{i=1}^n \log L_i^{(\lambda_0)} \geq g(V_0) \right\} \\ \cup \left\{ \exists n \geq 1 : V_n \in [V_0, V_0 \eta^{K(j)}], \sup_{\lambda \in (\lambda_{K(j)}, \lambda_0)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g(V_n) \right\} \\ \cup \left\{ \exists n \geq 1 : V_n \geq V_0 \eta^{K(j)}, \sum_{i=1}^n \log L_i^{(\lambda_{K(j)})} \geq g(V_0 \eta^{K(j)}) \right\}$$




Figure 7: Illustration of the stopping event of $\bar{N}_g(j)$ defined in (4.14), and related regions $H(\Delta_0)$, $H(\Delta_{K(j)})$ and R . The stopping time $\bar{N}_g(j)$ is the first time when $(\hat{\mu}_n, g(\bar{V}_n^{(j)})/V_n)$ is located in one of the colored areas, where $\bar{V}_n^{(j)} := \max\{V_0, \min\{V_n, V_0\eta^{K(j)}\}\}$.

$$\begin{aligned}
 &= \left\{ \exists n \geq 1 : V_n < V_0, \left(\hat{\mu}_n, \frac{g(V_0)}{V_n} \right) \in H(\Delta_0) \right\} \\
 &\cup \left\{ \exists n \geq 1 : V_n \in [V_0, V_0\eta^{K(j)}], \left(\hat{\mu}_n, \frac{g(V_n)}{V_n} \right) \in R \right\} \\
 &\cup \left\{ \exists n \geq 1 : V_n \geq V_0\eta^{K(j)}, \left(\hat{\mu}_n, \frac{g(V_0\eta^{K(j)})}{V_n} \right) \in H(\Delta_{K(j)}) \right\} \\
 &= \left\{ \exists n \geq 1 : V_n < V_0, \sum_{i=1}^n \log L_i^{(\lambda_0)} \geq g(V_0) \right\} \\
 &\cup \left\{ \exists n \geq 1 : V_n \in [V_0, V_0\eta^{K(j)}], \sup_{\lambda \in (\lambda_{K(j)}, \lambda_0)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g(V_n) \right\} \\
 &\cup \left\{ \exists n \geq 1 : V_n \geq V_0\eta^{K(j)}, \sum_{i=1}^n \log L_i^{(\lambda_{K(j)})} \geq g(V_0\eta^{K(j)}) \right\},
 \end{aligned}$$

where $V_0 := \inf\{t \geq 1 : D_0 \geq g(t)/t\}$ and the set R is defined by

$$R := \{(z, y) \in [0, \infty)^2 : y \leq \psi^*(z)\}, \quad (\text{B.19})$$

and $H(\Delta_0)$ and $H(\Delta_{K(j)})$ are half spaces contained in and tangent to R at $(\Delta_0, \frac{g(V_0)}{V_0})$ and $(\Delta_L, \frac{g(V_0\eta^{K(j)})}{V_0\eta^{K(j)}})$, respectively. See Figure 7 for an illustration of the upper bound of the stopping event of $\bar{N}_g(j)$.

Note that the first decomposition part $\{\exists n \geq 1 : V_n < V_0, \sum_{i=1}^n \log L_i^{(\lambda_0)} \geq g(V_0)\}$ is nonempty only if $V_0 > v_{\min}$, which is equivalent to $g(V_0) > v_{\min}D_0$. For the second part, a straightforward extension of Lemma 1 in the appendix of [36] implies that, for any fixed $\eta > 1$, the second part can be further decomposed by sets of simple events as follows:

$$\begin{aligned}
 &\left\{ \exists n \geq 1 : V_n \in [V_0, V_0\eta^{K(j)}], \sup_{\lambda \in (\lambda_{K(j)}, \lambda_0)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g(V_n) \right\} \\
 &\subset \bigcup_{k=1}^{K(j)} \left\{ \exists n \geq 1 : V_n \in [V_0\eta^{k-1}, V_0\eta^k], \right.
 \end{aligned}$$

$$\left. \sum_{i=1}^n \log L_i^{(\lambda_k)} \geq g(V_0\eta^k)/\eta \right\}, \quad (\text{B.20})$$

where each λ_k is given by $\lambda_k := \psi^*(\Delta_k)$ and each Δ_k is the solution of the equation (4.12) for $k = 1, \dots, K(j)$.

From this decomposition of the stopping event of $\bar{N}_g(j)$, for any fixed $\eta > 1$, we can lower bound the stopping time $\bar{N}_g(j)$ as follows:

$$\begin{aligned}
 \bar{N}_g(j) &= \inf \left\{ n \geq 1 : \sup_{\lambda \in (\lambda_{K(j)}, \lambda_0)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g(V_0\eta^{K(j)}) \right\} \\
 &\geq \inf \left\{ n \geq 1 : e^{-g(V_0)} \prod_{i=1}^n L_i^{(\lambda_0)} \mathbb{1}(g(V_0) > v_{\min}D_0) \right. \\
 &\quad \left. + \sum_{k=1}^{K(j)} e^{-g(V_0\eta^k)/\eta} \prod_{i=1}^n L_i^{(\lambda_k)} \geq 1 \right\} \\
 &= \inf \left\{ n \geq 1 : \sum_{k=0}^{K(j)} \omega_k \prod_{i=1}^n L_i^{(\lambda_k)} \geq 1/\alpha \right\} \\
 &= N_{1/\alpha}(j),
 \end{aligned}$$

which proves the claimed inequality. \square

Proof of Corollary 4.6. From Theorem 3.4 and Lemma 4.5, worst average delays of adaptive e-SR and e-CUSUM procedures are upper bounded by $\max_{\nu \geq 0} \min_{j \geq 1} \mathbb{E}_{0,Q} N_G(j + \nu) + j - 1$ where $\bar{N}_g(j)$ is a stopping time defined by

$$\bar{N}_g(j) := \inf \left\{ n \geq 1 : \sup_{\lambda \in (\lambda_{K(j)}, \lambda_0)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g(V_0\eta^{K(j)}) \right\}, \quad (\text{B.21})$$

for each $j \geq 1$. Now, let us first consider the case $\lambda^{\text{op}} \geq \lambda_0$. In this case, we use the following simple upper bound:

$$\begin{aligned}
 \min_{j \geq 1} [\mathbb{E}_{0,Q} N_G(j) + j - 1] &\leq \mathbb{E}_{0,Q} N_G(1) \\
 &= \mathbb{E}_{0,Q} \inf \left\{ n \geq 1 : \sup_{\lambda \in (\lambda_{K(1)}, \lambda_0)} \sum_{i=1}^n \log L_i^{(\lambda)} \geq g(V_0\eta^{K(1)}) \right\} \\
 &\leq \mathbb{E}_{0,Q} \inf \left\{ n \geq 1 : \sum_{i=1}^n \log L_i^{(\lambda_0)} \geq g_{r\alpha} \right\}.
 \end{aligned}$$

Since $\mathbb{E}_{0,Q} \log L_1^{(\lambda_0)} = \sigma^2(\lambda_0\Delta^{\text{op}} - \psi(\lambda_0)) \geq \sigma^2(\lambda_0\Delta_0 - \psi(\lambda_0)) = \sigma^2\psi^*(\Delta_0)$, by the same argument of Proposition 2.14, the last term above can be further upper bounded by

$$\frac{g_{r\alpha}}{D(Q|\mathcal{P})} \frac{\psi^*(\Delta^{\text{op}})}{\psi^*(\Delta_0)} + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda_0)}]}{[D(Q|\mathcal{P})]^2} \left[\frac{\psi^*(\Delta^{\text{op}})}{\psi^*(\Delta_0)} \right]^2 + 1. \quad (\text{B.22})$$

Since we are in the case $\lambda^{\text{op}} \geq \lambda_0$, we have $\frac{\psi^*(\Delta^{\text{op}})}{\psi^*(\Delta_0)} \geq 1$, which can be understood as a measure of inefficiency due to the misspecified upper bound of the oracle λ^{op} .

Now, consider the case $\lambda^{\text{op}} < \lambda_0$ where we correctly specified the upper bound. In this case, let j^{op} be the smallest integer satisfying $\lambda_{K(j^{\text{op}})} := \lambda_{K^{\text{op}}} < \lambda^{\text{op}} < \lambda_0$. Then, we can further upper bound the worst average delays by

$$\begin{aligned} & \min_{j \geq 1} [\mathbb{E}_{0,Q} N_G(j) + j - 1] \\ & \leq \mathbb{E}_{0,Q} \inf \left\{ n \geq 1 : \sum_{i=1}^n \log L_i^{(\lambda^{\text{op}})} \geq g(V_0 \eta^{K^{\text{op}}}) \right\} + j^{\text{op}} - 1 \\ & := \mathbb{E}_{0,Q} N_{\text{op}} + j^{\text{op}} - 1. \end{aligned}$$

By Equation (4.9), we have the following intermediate upper bound on the worst average delays,

$$\mathbb{E}_{0,Q} N_{\text{op}} + j^{\text{op}} - 1 \leq \frac{g(V_0 \eta^{K^{\text{op}}})}{D(Q||P)} + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda^{\text{op}})}]}{[D(Q||P)]^2} + j^{\text{op}}. \quad (\text{B.23})$$

Note that if $j^{\text{op}} = 1$ then $\lambda_{K(1)} = \lambda_L < \lambda^{\text{op}}$. Thus, in this case, we also correctly specified the lower bound, and the above bound is reduced to the same upper bound on the worst average delays in Theorem 4.3 of the well-separation case except the ARL parameter α being replaced by $r\alpha$.

Finally, to get an explicit upper bound on j^{op} for the case $j^{\text{op}} > 1$, first note that, from the definition of $\lambda_{K(j^{\text{op}}-1)}$ with the fact $\lambda^{\text{op}} < \lambda_1^{(j^{\text{op}}-1)} \Leftrightarrow \Delta^{\text{op}} < \Delta_{K(j^{\text{op}}-1)}$, we have

$$\frac{g(V_0 \eta^{K(j^{\text{op}}-1)})}{V_0 \eta^{K(j^{\text{op}}-1)}} = \psi^*(\Delta_{K(j^{\text{op}}-1)}) > \psi^*(\Delta^{\text{op}}). \quad (\text{B.25})$$

Also, the condition $K(j) \geq K_L + m \log_\eta j$ implies

$$j \leq \left[\frac{\eta^{-K_L}}{V_0} V_0 \eta^{K(j)} \right]^{1/m}, \quad (\text{B.26})$$

for each $j \geq 1$. By combining two inequalities above, we

have

$$\begin{aligned} j^{\text{op}} - 1 & \leq \left[\frac{\eta^{-K_L}}{V_0} V_0 \eta^{K(j^{\text{op}}-1)} \right]^{1/m} \\ & < \left[\frac{1}{V_0 \eta^{K_L}} \frac{g(V_0 \eta^{K(j^{\text{op}}-1)})}{\psi^*(\Delta^{\text{op}})} \right]^{1/m} \\ & \leq \left[\frac{\psi^*(\Delta_L)}{\psi^*(\Delta^{\text{op}})} \frac{g(V_0 \eta^{K^{\text{op}}})}{g_{r\alpha}} \right]^{1/m}. \end{aligned}$$

In sum, by combining all bounds above, we have

$$\begin{aligned} & \min_{j \geq 1} [\mathbb{E}_{0,Q} N_G(j) + j - 1] \\ & \leq \begin{cases} \frac{g_{r\alpha}}{D(Q||P)} \frac{\psi^*(\Delta^{\text{op}})}{\psi^*(\Delta_0)} \\ + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda_0)}]}{[D(Q||P)]^2} \left[\frac{\psi^*(\Delta^{\text{op}})}{\psi^*(\Delta_0)} \right]^2 + 1 & \text{if } \lambda^{\text{op}} \geq \lambda_0 \\ \frac{g_{r\alpha}}{D(Q||P)} + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda^{\text{op}})}]}{[D(Q||P)]^2} + 1 & \text{if } \lambda^{\text{op}} \in (\lambda_L, \lambda_0) \\ \frac{g_{r\alpha} + s\eta \log(1 + K^{\text{op}} - K_L)}{D(Q||P)} \\ + \frac{\mathbb{V}_{0,Q}[\log L_1^{(\lambda^{\text{op}})}]}{[D(Q||P)]^2} \\ + \left[\frac{\psi^*(\Delta_L)}{\psi^*(\Delta^{\text{op}})} \frac{g_{r\alpha} + s\eta \log(1 + K^{\text{op}} - K_L)}{g_{r\alpha}} \right]^{1/m} & \text{if } \lambda^{\text{op}} \leq \lambda_L, \end{cases} \end{aligned}$$

as desired. \square

APPENDIX C. AN EXPLICIT WAY TO COMPUTE THE THRESHOLD IN ALGORITHM 3

The following pseudo-code describes how to compute the threshold g_α defined by

$$\begin{aligned} g_\alpha & := \inf \left\{ g > \log(1/\alpha) : e^{-g} \mathbb{1}(g > v_{\min} D_U) \right. \\ & \quad \left. + \min_{k \in [K_{\max}]} k \exp \left\{ -g \left(\frac{D_U}{D_L} \right)^{-1/k} \right\} \leq \alpha \right\}. \quad (\text{C.1}) \end{aligned}$$

Algorithm 4: Pseudo-code of computeThreshold function.

Input: ARL parameter $\alpha \in (0, 1)$, Boundary values $0 < \Delta_L < \Delta_U$,
Maximum number of baselines $K_{\max} \in \mathbb{N}$, Tolerance $\epsilon > 0$.

Output: Threshold $g_\alpha > 0$ that is defined and used in Algorithm 3.

1 Set $D_L := \psi^*(\Delta_L) < \psi^*(\Delta_U) =: D_U$ and define a function f on \mathbb{R}_+ as

$$f(g) := \min_{k \in [K_{\max}]} k \exp \left\{ -g \left(\frac{D_U}{D_L} \right)^{-1/k} \right\}. \quad (\text{B.24})$$

2 **if** $f(v_{\min} D_U) \leq \alpha$ **then**

3 Compute $g_\alpha := \inf \{ g \in (\log(1/\alpha), v_{\min} D_U) : f(g) \leq \alpha \}$ by using the bisection method to the function $g \mapsto f(g) - \alpha$ with endpoints $\{\log(1/\alpha), v_{\min} D_U\}$ and tolerance ϵ .

4 **else**

5 Compute $g_\alpha := \inf \{ g \in (v_{\min} D_U, \frac{D_U}{D_L} \log(2/\alpha)) : e^{-g} + f(g) \leq \alpha \}$ by using the bisection method to the function $g \mapsto e^{-g} + f(g) - \alpha$ with endpoints $\{v_{\min} D_U, \frac{D_U}{D_L} \log(2/\alpha)\}$.

6 **return** $g_\alpha > 0$

FUNDING

A. Ramdas was partially supported by NSF grants IIS-2229881 and DMS-2310718. J. Shin and A. Rinaldo were partially supported by NSF grant DMS-EPSC 2015489.

Accepted 12 October 2023

REFERENCES

- [1] BARNARD, G. A. (1959). Control charts and stochastic processes. *Journal of the Royal Statistical Society: Series B (Methodological)* **21**(2) 239–257.
- [2] CHEN, Y., WANG, T. and SAMWORTH, R. J. (2022). High-dimensional, multiscale online changepoint detection. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **84**(1) 234–266. <https://doi.org/10.1111/rssb.12447>. MR4400396
- [3] DESOBRY, F., DAVY, M. and DONCARLI, C. (2005). An online kernel change detection algorithm. *IEEE Transactions on Signal Processing* **53**(8) 2961–2974. <https://doi.org/10.1109/TSP.2005.851098>. MR2169647
- [4] DURRETT, R. (2019) *Probability: Theory and Examples* **49**. Cambridge University Press. <https://doi.org/10.1017/9781108591034>. MR3930614
- [5] FAN, X., GRAMA, I. and LIU, Q. (2015). Exponential inequalities for martingales with applications. *Electronic Journal of Probability* **20**, 1. <https://doi.org/10.1214/EJP.v20-3496>. MR3311214
- [6] GANGRADE, A., RINALDO, A. and RAMDAS, A. (2023). A Sequential Test for Log-Concavity. *arXiv preprint arXiv:2301.03542*.
- [7] GRÜNWARD, P., DE HEIDE, R. and KOOLEN, W. (2023). Safe testing. *Journal of the Royal Statistical Society, Series B*. **195** 47–63. <https://doi.org/10.1016/j.jspi.2017.09.014>. MR3760837
- [8] HADJILIADIS, O., ZHANG, H. and POOR, H. V. (2009). One shot schemes for decentralized quickest change detection. *IEEE Transactions on Information Theory* **55**(7) 3346–3359. <https://doi.org/10.1109/TIT.2009.2021311>. MR2598025
- [9] HONDA, J. and TAKEMURA, A. (2010). An Asymptotically Optimal Bandit Algorithm for Bounded Support Models. In *Conference on Learning Theory* 67–79. Citeseer.
- [10] HONDA, J. and TAKEMURA, A. (2015). Non-asymptotic analysis of a new bandit algorithm for semi-bounded rewards. *Journal of Machine Learning Research* **16** 3721–3756. MR3450551
- [11] HOWARD, S. R., RAMDAS, A., MCAULIFFE, J. and SEKHON, J. (2020). Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys* **17** 257–317. <https://doi.org/10.1214/18-PS321>. MR4100718
- [12] HOWARD, S. R., RAMDAS, A., MCAULIFFE, J. and SEKHON, J. (2021). Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics* **49**(2) 1055–1080. <https://doi.org/10.1214/20-aos1991>. MR4255119
- [13] LAI, T. L. (1995). Sequential changepoint detection in quality control and dynamical systems. *Journal of the Royal Statistical Society: Series B (Methodological)* **57**(4) 613–644. MR1354072
- [14] LAI, T. L. (1998). Information bounds and quick detection of parameter changes in stochastic systems. *IEEE Transactions on Information Theory* **44**(7) 2917–2929. <https://doi.org/10.1109/18.737522>. MR1672051
- [15] LAI, T. L. (2001). Sequential analysis: some classical problems and new challenges. *Statistica Sinica* **11**(2) 303–351. MR1844531
- [16] LI, S., XIE, Y., DAI, H. and SONG, L. (2015). M-Statistic for Kernel Change-Point Detection. In *Advances in Neural Information Processing Systems* **28**.
- [17] LORDEN, G. (1970). On excess over the boundary. *The Annals of Mathematical Statistics* **41** 520–527. <https://doi.org/10.1214/aoms/1177697092>. MR0254981
- [18] LORDEN, G. (1971). Procedures for reacting to a change in distribution. *The Annals of Mathematical Statistics* **42** 1897–1908. <https://doi.org/10.1214/aoms/1177693055>. MR0309251
- [19] MANOLE, T. and RAMDAS, A. (2023). Martingale methods for sequential estimation of convex functionals and divergences. *IEEE Transactions on Information Theory* **69**(7) 4641–4658. MR4613565
- [20] MEI, Y. (2010). Efficient scalable schemes for monitoring a large number of data streams. *Biometrika* **97**(2) 419–433. <https://doi.org/10.1093/biomet/asq010>. MR2650748
- [21] MOUSTAKIDES, G. V. (1986). Optimal stopping times for detecting changes in distributions. *The Annals of Statistics* **14**(4) 1379–1387. <https://doi.org/10.1214/aos/1176350164>. MR0868306
- [22] MOUSTAKIDES, G. V. (2008). Sequential change detection revisited. *The Annals of Statistics* **36**(2) 787–807. <https://doi.org/10.1214/009053607000000938>. MR2396815
- [23] PAGE, E. S. (1954). Continuous inspection schemes. *Biometrika* **41**(1/2) 100–115. <https://doi.org/10.1093/biomet/41.1-2.100>. MR0088850
- [24] PODKOPAEV, A., BLÖBAUM, P., KASIVISWANATHAN, S. P. and RAMDAS, A. (2023). Sequential Kernelized Independence Testing. In: *International Conference on Machine Learning*.
- [25] POLLAK, M. and SIEGMUND, D. (1975). Approximations to the expected sample size of certain sequential tests. *The Annals of Statistics* **3**(6) 1267–1282. MR0403114
- [26] POLLAK, M. (1985). Optimal detection of a change in distribution. *The Annals of Statistics* **13**(1) 206–227. <https://doi.org/10.1214/aos/1176346587>. MR0773162
- [27] POLUNCHENKO, A. S. and TARTAKOVSKY, A. G. (2010). On optimality of the Shiryaev-Roberts procedure for detecting a change in distribution. *The Annals of Statistics* **38**(6) 3445–3457. <https://doi.org/10.1214/09-AOS775>. MR2766858
- [28] RAMDAS, A., RUF, J., LARSSON, M. and KOOLEN, W. (2020). Admissible anytime-valid sequential inference must rely on nonnegative martingales. *arXiv preprint arXiv:2009.03167*.
- [29] RAMDAS, A., RUF, J., LARSSON, M. and KOOLEN, W. M. (2022). Testing exchangeability: Fork-convexity, supermartingales and e-processes. *International Journal of Approximate Reasoning*. **141** 83–109. <https://doi.org/10.1016/j.ijar.2021.06.017>. MR4364897
- [30] RAMDAS, A., GRÜNWARD, P., VOVK, V. and SHAFER, G. (2023). Game-theoretic statistics and safe anytime-valid inference. *Statistical Science* **38**(4) 576–601. <https://doi.org/10.1214/23-sts894>. MR4665027
- [31] RITOV, Y. (1990). Decision theoretic optimality of the CUSUM procedure. *The Annals of Statistics* **18**(3) 1464–1469. <https://doi.org/10.1214/aos/1176347761>. MR1062720
- [32] ROBERTS, S. (1966). A comparison of some control chart procedures. *Technometrics* **8**(3) 411–430. <https://doi.org/10.2307/1266688>. MR0196887
- [33] SHAFER, G. and VOVK, V. (2019) *Game-Theoretic Foundations for Probability and Finance* **455**. John Wiley & Sons. <https://doi.org/10.1002/0471249696>. MR1852450
- [34] SHEKHAR, S. and RAMDAS, A. (2023). Nonparametric two sample testing by betting. *IEEE Transactions on Information Theory*. <https://doi.org/10.1109/TIT.2023.3305867>
- [35] SHEKHAR, S. and RAMDAS, A. (2023). Sequential change detection via backward confidence sequences. In *International Conference on Machine Learning*.
- [36] SHIN, J., RAMDAS, A. and RINALDO, A. (2021). Nonparametric iterated-logarithm extensions of the sequential generalized likelihood ratio test. *IEEE Journal on Selected Areas in Information Theory*.
- [37] SHIRYAEV, A. N. (1963). On optimum methods in quickest detection problems. *Theory of Probability & Its Applications* **8**(1) 22–46. MR0155708
- [38] SIEGMUND, D. and VENKATRAMAN, E. (1995). Using the generalized likelihood ratio statistic for sequential detection of a changepoint. *The Annals of Statistics* **23** 255–271. <https://doi.org/10.1214/aos/1176324466>. MR1331667
- [39] TARTAKOVSKY, A., NIKIFOROV, I. and BASSEVILLE, M. (2014) *Sequential Analysis: Hypothesis Testing and Changepoint Detec-*

- tion. CRC Press. [MR3241619](#)
- [40] TARTAKOVSKY, A. G. (2014). Nearly optimal sequential tests of composite hypotheses revisited. *Proceedings of the Steklov Institute of Mathematics* **287**(1) 268–288. <https://doi.org/10.1134/S0081543814080161>. [MR3484334](#)
- [41] TARTAKOVSKY, A. G., POLLAK, M. and POLUNCHENKO, A. S. (2012). Third-order Asymptotic Optimality of the Generalized Shiryaev–Roberts Changepoint Detection Procedures. *Theory of Probability & Its Applications* **56**(3) 457–484. <https://doi.org/10.4213/tvp4406>. [MR3136464](#)
- [42] VILLE, J. (1939). Étude critique de la notion de collectif. *Monographies des Probabilités* **3** 144. [MR3533075](#)
- [43] VOVK, V. (2021). Testing randomness online. *Statistical Science* **36**(4) 595–611. <https://doi.org/10.1214/20-sts817>. [MR4323055](#)
- [44] VOVK, V., GAMMERMAN, A. and SHAFER, G. (2005) *Algorithmic Learning in a Random World*. Springer Science & Business Media. [MR2161220](#)
- [45] WANG, H. and RAMDAS, A. (2023). Catoni-style confidence sequences for heavy-tailed mean estimation. *Stochastic Processes and Applications* **163** 168–202. <https://doi.org/10.1016/j.spa.2023.05.007>. [MR4610125](#)
- [46] WANG, H. and RAMDAS, A. (2023). Huber-robust confidence sequences. In *International Conference on Artificial Intelligence and Statistics* 9662–9679. PMLR.
- [47] WASSERMAN, L., RAMDAS, A. and BALAKRISHNAN, S. (2020). Universal inference. *Proceedings of the National Academy of Sciences* **117**(29) 16880–16890. <https://doi.org/10.1073/pnas.1922664117>. [MR4242731](#)
- [48] WAUDBY-SMITH, I. and RAMDAS, A. (2023). Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B: Statistical Methodology*.
- [49] WILLSKY, A. and JONES, H. (1976). A generalized likelihood ratio approach to the detection and estimation of jumps in linear systems. *IEEE Transactions on Automatic Control* **21**(1) 108–112. <https://doi.org/10.1109/tac.1976.1101146>. [MR0396019](#)
- [50] XIE, Y. and SIEGMUND, D. (2013). Sequential multi-sensor change-point detection. In *2013 Information Theory and Applications Workshop (ITA)* 1–20. IEEE. <https://doi.org/10.1214/13-AOS1094>. [MR3099117](#)
- [51] ZOU, S., VEERAVALLI, V. V., LI, J. and TOWSLEY, D. (2019). Quickest detection of dynamic events in networks. *IEEE Transactions on Information Theory* **66**(4) 2280–2295. <https://doi.org/10.1109/TIT.2019.2948350>. [MR4087685](#)

Jaehyeok Shin. Google, USA. E-mail address: shinjaehyeok@google.com

Aaditya Ramdas. Carnegie Mellon University, USA. E-mail address: aramdas@cmu.edu

Alessandro Rinaldo. Carnegie Mellon University, USA. E-mail address: arinaldo@cmu.edu