

Finite-time Analysis of Globally Nonstationary Multi-Armed Bandits

Junpei Komiyama

Stern School of Business, New York University, USA

JUNPEI@KOMIYAMA.INFO

Edouard Fouché

Institute for Program Structures and Data Organization, Karlsruhe Institute of Technology, Germany

EDOUARD.FOUCHE@GMAIL.COM

Junya Honda

*Department of Systems Science, Graduate School of Informatics, Kyoto University, Japan
Center for Advanced Intelligence Project, RIKEN, Japan*

HONDA@I.KYOTO-U.AC.JP

Editor: Aurelien Garivier

Abstract

We consider nonstationary multi-armed bandit problems where the model parameters of the arms change over time. We introduce the adaptive resetting bandit (ADR-bandit), a bandit algorithm class that leverages adaptive windowing techniques from literature on data streams. We first provide new guarantees on the quality of estimators resulting from adaptive windowing techniques, which are of independent interest. Furthermore, we conduct a finite-time analysis of ADR-bandit in two typical environments: an abrupt environment where changes occur instantaneously and a gradual environment where changes occur progressively. We demonstrate that ADR-bandit has nearly optimal performance when abrupt or gradual changes occur in a coordinated manner that we call global changes. We demonstrate that forced exploration is unnecessary when we assume such global changes. Unlike the existing nonstationary bandit algorithms, ADR-bandit has optimal performance in stationary environments as well as nonstationary environments with global changes. Our experiments show that the proposed algorithms outperform the existing approaches in synthetic and real-world environments.

Keywords: multi-armed bandits, adaptive windows, nonstationary bandits, changepoint detection, sequential learning.

1. Introduction

1.1 Motivations

The multi-armed bandit (MAB; Thompson (1933); Robbins (1952)) is a fundamental model capturing the dilemma between exploration and exploitation in sequential decision making. This problem involves K arms (i.e., possible actions). The decision-maker selects a set of arms at each time step and observes a corresponding reward. The goal of the decision-maker is to maximize the cumulative reward over time. The performance of a bandit algorithm is usually measured via the metric of “regret”: the difference between the obtained rewards and the rewards one would have obtained by choosing the best arms. Minimizing the regret corresponds to maximizing the expected reward.

MAB has been used to solve numerous problems, such as experimental clinical design (Thompson, 1933), online recommendations (Li et al., 2010), online advertising (Chapelle and Li, 2011; Komiyama et al., 2015), and stream monitoring (Fouché et al., 2019). The most widely studied version of this model is the stochastic MAB, which assumes that the reward for each arm is drawn from an unknown but fixed distribution. In the stochastic MAB, several algorithms, such as the upper confidence bound (UCB; Lai and Robbins, 1985; Auer et al., 2002) and Thompson sampling (TS; Thompson, 1933) are known to have $\Theta(K \log T / \Delta_{\min})$ regret,¹ which is optimal (Lai and Robbins, 1985). While it is reasonable to assume in some cases that the reward-generating process does not change, as in these algorithms, the distribution of rewards may change over time in many applications. To observe this, we consider the following two examples:

Example 1 (Online advertising) A website has several advertisement slots. Based on each user’s query, the website decides which ads to display from a set of candidates (i.e., “relevant advertisements”). Some advertisements are more appealing to a user than others. Each advertisement is associated with a click-through rate (CTR), the number of clicks per view. Websites receive revenue from clicks on advertisements; thus, maximizing the CTR maximizes the revenue. This problem is structured as the bandit problem, with advertisements and clicks as arms and rewards. However, it is well known that the CTR of some advertisements may change over time for several reasons, such as seasonality or changing user interests. In this case, naively applying a stochastic MAB algorithm leads to suboptimal rewards.

Example 2 (Predictive maintenance) Correlation often results from physical relationships, for example, between the temperature and pressure of a fluid. A change in these correlations often signals a shift in the system’s state (like a fluid solidifying) or a failure or degradation in equipment (such as a leak). In the context of large-scale factory monitoring, keeping track of these correlations can help anticipate issues, thus reducing maintenance costs. Nevertheless, constantly updating the full correlation matrix is computationally unfeasible due to the data’s high-dimensionality and dynamic nature. A more efficient solution consists of updating only a few elements of the matrix based on a notion of utility (e.g., high correlation values). The system must minimize the monitoring cost while maximizing the total utility in a possibly nonstationary environment. In other words, correlation monitoring can be considered an instance of the bandit problem (Fouché et al., 2019), in which pairs of sensors and correlation coefficients correspond to arms and rewards.

In such settings, the reward may evolve over time (i.e., it is nonstationary²). The nonstationary MAB (NS-MAB) describes a class of MAB algorithms addressing this particular setting. Most of the NS-MAB algorithms rely on passive forgetting methods based on a sliding window (Garivier and Moulines, 2011) or fixed-time resetting (Gur et al., 2014). Recent work has proposed more sophisticated change detection mechanisms based on adaptive windows (Srivastava et al., 2014) or sequential likelihood ratio tests (Besson et al., 2022). However, the existing methods have several drawbacks, as we describe later.

1. The value Δ_{\min} is a distribution-dependent constant quantifying the hardness of the problem instance.
 2. The use of the term “nonstationary” in the bandit literature is different from the literature on time series analysis. We formally define stationary and nonstationary streams in Definition 1.




Figure 1: In the Bioliq dataset, each arm corresponds to the correlation coefficient between a pair of sensors. As we can see, the rewards between different arms tend to change coordinately. We describe this dataset in details in Section 6.

1.2 Challenges in nonstationary bandits

Although change detectors can help MAB algorithms adapt to changing rewards, they often come with costs. Let us consider the case of an abruptly changing environment, where the reward distributions change drastically at some time steps. Previous work by Garivier and Moulines (2011) indicated that the regret of any K -armed bandit algorithm for such a case is³ $\Omega(\sqrt{T})$. This finding implies that the performance of NS-MAB algorithms in stationary (i.e., non-changing) environments is inferior to the performance of standard stochastic MAB algorithms, such as the UCB and TS, because $O(\sqrt{T})$ is much larger than $O(K \log T / \Delta_{\min})$ given a moderate value of Δ_{\min} . Virtually all NS-MAB algorithms conduct $O(\sqrt{T})$ forced exploration for all the arms, which is the leading factor of regret—no matter what changepoint detection algorithm is used.

Another drawback of most existing methods is that they require several parameters that are highly specific to the problem and require unreasonable environmental assumptions, such as the number of changes or an estimation of the amount of “nonstationarity” in the stream. If such parameters are not set correctly, the actual performance may deviate widely from the given theoretical bounds.

This paper solves these issues by introducing adaptive resetting bandit (ADR-bandit) algorithms, a new class of bandit algorithms. Several algorithms, such as UCB and TS, have been established in the stationary case; therefore, we aim to extend them without

3. This holds even when the change is sufficiently large, regardless of the value of a distribution-dependent constant Δ_{\min} . See details in Theorem 13, Corollary 14, and Remark 17 in Garivier and Moulines (2011).

Table 1: The table below compares the achievable performance of the existing NS-MAB algorithms, ADR-bandit, and stationary MAB algorithms (e.g., TS and UCB). The value d defines the speed of gradual changes. ADR-bandit has optimal performance in stationary environments as well as nonstationary environments with global changes. Most existing NS-MAB algorithms only handle abrupt changes and do not provide any performance guarantees under gradual changes. Performance bounds under the gradual changes are provided in Besbes et al. (2019); Allesiardo and Féraud (2015); Wei et al. (2016); Wei and Srivastava (2018); Trovò et al. (2020); Krishnamurthy and Gopalan (2021). Here, \tilde{O} is a Landau notation that ignores a polylogarithmic factor.

	Stationary	Abrupt	Gradual
Existing NS-MAB	$\tilde{O}(\sqrt{T})$	$\tilde{O}(\sqrt{T})$	$\tilde{O}(T^{1-d/3})$
ADR-bandit	$\mathbf{O}(\log T/\Delta_{\min})$	$\tilde{O}(\sqrt{T})$ (Under GC)	$\tilde{O}(T^{1-d/3})$ (Under GC)
Stationary MAB	$O(\log T/\Delta_{\min})$	$O(T)$	$O(T)$

introducing any forced exploration. For this purpose, we combine them with an adaptive windowing technique. For example, ADR-TS, which is an instance of the ADR-bandit, combines the adaptive windowing with TS. Our method deals with a subclass of changes that we call *global changes*. Intuitively, if all arms change in a coordinated manner, we can avoid forced exploration to improve performance. This type of change is natural in the predictive maintenance of Example 2, as illustrated in Figure 1, where the nonstationarity often results from changes in the entire system.

Our analysis shows that the proposed method has optimal performance for stationary streams and comparable bounds with existing NS-MAB algorithms for abruptly changing and gradually changing environments under the global change assumption. Table 1 compares the bounds from the proposed method with the existing types of NS-MAB algorithms.

1.3 Contributions

We articulate our contributions as follows. First, we provide an extensive analysis of adaptive windowing techniques. We focus, on the adaptive windowing (ADWIN) algorithm (Bifet and Gavaldà, 2007), which is still considered state-of-the-art in the data stream literature. The ADWIN algorithm performs well regarding various types of streams (Gama et al., 2014). Bifet and Gavaldà (2007) provided false positive and false negative rate bounds.⁴ However, existing analyses are not sufficient for our aim. For the analysis of bandit algorithms using ADWIN, we need an estimate on the accuracy of the estimator $\hat{\mu}_t$. Thus, we conduct a finite-time analysis on the estimation error $|\mu_t - \hat{\mu}_t|$ (Section 3). As a by-product, this analysis explains why adaptive windowing methodologies perform well in many stream learning problems by bounding the error for abrupt and gradual changes.

4. Although, rigorously speaking, their analysis is not correct, as we will discuss later.

After formalizing the MAB problem (Section 4), we introduce the ADR (ADaptive Resetting) bandit algorithm (Section 5). Our study builds on a recent paper by Fouché et al. (2019) that applies the MAB to sensor streams. Although Fouché et al. (2019) proposed a way to combine the stationary bandit algorithms (e.g., UCB and TS) with adaptive windows, they did not provide any theoretical guarantees for the nonstationary case. We slightly modify their framework to enable rigorous analysis. Our work bridges the divide between these methods’ practical application and theoretical comprehension. We believe our analysis and characterization of global changes present novel contributions to the literature on nonstationary bandit algorithms.

Finally, we demonstrate the performance of our proposed method concerning synthetic and real-world environments (Section 6). The proposed method outperforms the existing nonstationary bandit methods in stationary streams and abruptly and gradually changing environments.

2. Related Work

2.1 Change detection

Change detection is an important problem in data mining. The goal is to detect when the statistical properties (e.g., the mean) of a stream of observations change. Such changes are commonly attributed to the phenomenon known as *concept drift* (Gama et al., 2014; Lu et al., 2019), which characterizes unforeseeable changes in the underlying data distribution. Detecting such changes is crucial to virtually any monitoring tasks in streams, such as controlling the performance of online machine learning algorithms (Bifet and Gavaldà, 2007) or detecting correlation changes (Seliniotaki et al., 2014).

There are numerous methods to detect changes. The fundamental idea is to measure whether the estimated parameters of the current data distribution (e.g., its mean) have changed at any time. In other words, change detection is about separating the signal from the noise (Gama and Castillo, 2006).

Changepoint detection approaches can be classified into three major categories (cf. Table II in Gama et al. (2014)). These categories are sequential analysis approaches (e.g., CUSUM (Page, 1954) and its variant Page-Hinkley (PH) testing (Hinkley, 1971)), statistical process control (e.g., the Drift Detection Method (DDM) (Gama et al., 2004) and its variants), and monitoring two distributions (e.g., ADWIN in Bifet and Gavaldà, 2007).

In this work, we primarily consider ADWIN (Bifet and Gavaldà, 2007) because it works with any bounded distribution and has a good affinity with online learning analyses. Moreover, ADWIN monitors the mean from a sequence of observations over a sliding window of adaptive size. When ADWIN detects a change between two subwindows, the oldest observations are discarded. Otherwise, the window grows indefinitely. The success of this approach is due to the quasi-absence of parameters, making it highly adaptive.

Gonçalves et al. (2014) empirically compared ADWIN with other drift detectors. They found that ADWIN is one of the fastest detectors and is one of the only methods to provide false positive and false negative (fp/fn) rate guarantees. Many of the existing methods do not provide any performance guarantees, which is the case for DDM (Gama et al., 2004), EDDM (Baena-Garcia et al., 2006), and ECDD (Ross et al., 2012) (per Blanco et al., 2015). Although we believe that one can provide fp/fn rate guarantees for the statistical process

control and CUSUM (Page, 1954) approaches by choosing the appropriate parameters, limited discussions are available on the theoretical properties of these algorithms. For a thorough history and comparison of change-detection methods, we refer to the surveys by Gama et al. (2014); Krawczyk et al. (2017).

2.2 Nonstationary bandits

In comparison to the change detection literature, theoretical performance guarantees (i.e., regret bounds) are much more emphasized in the MAB literature. Traditionally, NS-MAB algorithms are divided into two categories. *Active* methods actively seek to detect changes, and *passive* methods do not. Another type of algorithms, known as adversarial bandits (Auer et al., 1995), such as Exp3, can deal with changing environments. However, the guarantee of adversarial bandit algorithms is limited when no arm is consistently good. In the following paragraphs, we discuss the related work on active and passive methods.

Active methods: Hartland et al. (2007) proposed Adapt-EvE, which combines the PH test with UCB, and Mellor and Shapiro (2013) suggested changepoint TS. However, these two studies do not provide any regret bounds. Liu et al. (2018) proposed CUSUM-UCB and PH-UCB, which combine the UCB algorithm with a CUSUM-based (or PH-based) resetting and forced exploration. They derived a regret bound of $\tilde{O}(\sqrt{MT})$ in an abruptly changing environment with M known change points. Cao et al. (2019) suggested the M-UCB algorithm, which combines UCB with an adaptive-window-based resetting method with a regret bound of $\tilde{O}(\sqrt{MT})$ for an abruptly changing environment. Allesiaro and Féraud (2015) proposed Exp3.R, which combines Exp3 (Auer et al., 1995) with changepoint detection and provides a $\tilde{O}(M\sqrt{T})$ regret bound. Although the assumptions are slightly different among algorithms, many of the regret bounds are of the order of $\tilde{O}(\sqrt{MT})$ concerning the number of change points M and the number of time steps T . Auer et al. (2019) provided an epoch-based bandit algorithm that does not require the knowledge of M . Moreover, Sez nec et al. (2020) applied the adaptive window to the rotting bandit problem where the reward of the arms only decreases. Krishnamurthy and Gopalan (2021) proposed an elimination-based algorithm and analyzed its performance in a two-armed gradually changing environment. They derived a similar regret bound $O(T^{1-d/3})$ as this paper, where d is the change-speed parameter that is the same as ours. Besson et al. (2022) proposed GLR-kLUCB, a combination of KL-UCB (Garivier and Cappé, 2011) with the Bernoulli Generalized Likelihood Ratio (GLR) test (Maillard, 2017) and forced exploration for an abruptly changing environment. They discussed global and local resets and provided a regret bound of $\tilde{O}(\sqrt{MT})$.⁵ One of the closest works to ours is Mukherjee and Maillard (2019), where they proposed algorithms UCBL-CPD and ImpCPD by combining UCB with a change point detector. They considered the global changes in the abrupt environment with the same spirit as ours to avoid forced exploration. Still, their main goal is to characterize detailed regret bounds for this environment, while ours is to cover both the global abrupt and global gradual environments by a single algorithm. To be more specific, they obtain both distribution-dependent and independent bounds using a confidence interval established through the Laplace method. This approach circumvents the need for union

5. They also discussed dependence on the distribution-dependent constants (Corollary 6 therein).

bounds, leading to a reduced confidence level.⁶ In terms of distribution-independent regret, ImpCPD has $O(\sqrt{MT})$ distribution-independent regret bound, while UCBL-CPD and our algorithm have $\tilde{O}(\sqrt{MT})$ regret including the polylogarithmic factor in T . On the other hand, most of the effort of this paper is devoted to the analysis for the gradual changes. We think that the application of the Laplace method could potentially improve polylogarithmic factors of our bounds, but we opted to use confidence bounds based on union bounds for simplicity to focus on simultaneously handling two environments.

Passive methods: Kocsis and Szepesvári (2006) proposed Discounted UCB (D-UCB). Garivier and Moulines (2011) suggested Sliding Window UCB (SW-UCB) and analyzed D-UCB to demonstrate that the two algorithms have a $\tilde{O}(\sqrt{MT})$ regret bound for abruptly changing environments. Besbes et al. (2019) considered the case of limited variation V and proposed the Rexp3 algorithm with the regret bound $\Theta(V^{1/3}T^{2/3})$. Wei et al. (2016) generalized the analysis of Besbes et al. (2019) to the case of intervals where each interval has limited variation. Wei and Srivastava (2018) proposed the LM-DSEE and SW-UCB# algorithms with a $\tilde{O}(\sqrt{MT})$ regret bound for an abruptly changing environment. The latter algorithm adopts an adaptive window with a limited length. They also analyzed the case of a slowly changing environment and derived a regret bound of SW-UCB#. Chen et al. (2019) proposed a very involved algorithm and stated that it can adapt to abrupt changes as well as gradual changes with a $\Theta(\min(\sqrt{MT}, V^{1/3}T^{2/3}))$ regret bound. Trovò et al. (2020) proposed the Sliding Window TS (SW-TS) algorithm and provided a $\tilde{O}(\sqrt{MT})$ regret bound for an abruptly changing environment. They also provided a distribution-dependent regret bound for a gradually changing environment.

Our approach is an active one; however, it does strikingly differ from the existing methods. Our algorithm does not sacrifice performance in the stationary case, whereas almost all existing NS-MAB algorithms are exclusively designed for nonstationary environments. In real-world settings, users may not know whether a given stream of data must be considered stationary or not. Similarly, the nature of the stream may also change over time. In such settings, our approach is advantageous, as our experiments show.

Although the analysis that follows is quite involved, our algorithms are conceptually simple and aimed at practical use. In comparison, our competitors tend to require more computation and rely on parameters that are difficult to set.

3. Analysis of ADWIN

This section analyzes ADWIN (Bifet and Gavaldà, 2007). This algorithm monitors at any time t an estimate of the mean $\hat{\mu}_t$ from a single stream of univariate observations.

3.1 Data streams

We assume that each observation $x_t \in [0, 1]$ at any time step t is drawn from some distribution with the mean μ_t . The value μ_t is not known to the algorithm, and ADWIN estimates it from a sequence of (possibly noisy) observations $S : (x_1, x_2, \dots, x_T)$. The goal of ADWIN

6. In particular, confidence levels required in Theorems 1 and 2 therein are $1/T$ to the number of rounds T . This is superior to the T^3 level of confidence required in our paper.




Figure 2: We classify the four types of concepts drifts (abrupt, recurring, gradual, incremental) into two types of changes: abrupt and gradual changes. Here, dashed lines represent μ_t , whereas solid lines represent x_t .

is to maintain an estimator of μ_t at any t based on past observations. Due to *concept drift* (Gama et al., 2014), the mean μ_t may change over time, so the task is not trivial.

The data stream literature (Gama et al., 2014) identifies the four main categories of concept drifts: abrupt, recurring, gradual, and incremental drifts. Figure 2 indicates that the mean μ_t (the dashed line in the figure) changes *abruptly* and stays the same for some time with abrupt and recurring drifts, whereas μ_t changes *gradually* over time with gradual and incremental drifts. The actual observations x_t can be arbitrarily noisy. We are only interested in the change in the mean μ_t ; thus, we simplify our analysis to the two types of changes.⁷ In addition, we also consider stationary streams where the mean μ_t never changes over time. In summary, we consider three types of streams:

Definition 1 (Stationary, abruptly changing, and gradually changing streams)

1. A stream is static if $\mu_t = \mu$ for all $t \in [T]$ and some $\mu \in [0, 1]$.
2. A stream is abruptly-changing if $\mu_t = \mu_{t+1}$ except for changepoints $\mathcal{T}_C \subset [T]$.
3. A stream is gradually-changing if $|\mu_{t+1} - \mu_t| \leq b$ for all $t \in [T]$ and some constant $b \in (0, 1)$.

These definitions refer to the mean μ_t . The observations $\{x_1, x_2, \dots\}$ can be noisy.

3.2 The adaptive window algorithm

Algorithm 1 is the pseudo-code for ADWIN (Bifet and Gavaldà, 2007). The law of large numbers implies that using more observations helps estimate the parameter. However, the nature of older observations might be different from more recent observations. To determine a good trade-off between these two effects, ADWIN maintains a window $W(t)$ of past time and discards data points outside the window.

⁷ This categorization matches the MAB literature (Wei and Srivastava, 2018; Trovò et al., 2020).

Algorithm 1 ADWIN

Require: A univariate stream of values $S : (x_1, x_2, \dots) \in [0, 1]$, confidence level $\delta \in (0, 1)$.

- 1: $W(1) = \{\}$
 - 2: **for** $t = 1, 2, \dots$ **do**
 - 3: $W(t+1) = W(t) \cup \{t\}$
 - 4: **while** $|\hat{\mu}_{W_1} - \hat{\mu}_{W_2}| \geq \epsilon_{\text{cut}}^\delta$ holds for some split $W(t+1) = W_1 \cup W_2$ **do**
 - 5: $W(t+1) = W_2$.
 - 6: **end while**
 - 7: **end for**
-




Figure 3: Terminology of changepoint and detection time. A changepoint of an abruptly changing stream is a time step t where $\mu_t \neq \mu_{t+1}$. A detection time is the time step where ADWIN shrinks the window.

We omit the index (t) of $W(t)$ when the time step of interest is clear. At each time step t , ADWIN receives a new observation x_t and extends the window W to include the observation. For the current window W , we let $\hat{\mu}_W = \frac{1}{|W|} \sum_{t \in W} x_t$ be the corresponding empirical mean. For each new observation, ADWIN tests whether the mean of the underlying distribution has changed. If we can split W into two consecutive disjoint subwindows $W_1 \cup W_2 = W$ whose empirical means are significantly different (i.e., by some threshold ϵ_{cut}), then ADWIN discards W_1 (i.e., ADWIN *shrinks* the window):

$$\epsilon_{\text{cut}}^\delta = \sqrt{\frac{1}{2|W_1|} \log\left(\frac{1}{\delta}\right)} + \sqrt{\frac{1}{2|W_2|} \log\left(\frac{1}{\delta}\right)}, \quad (1)$$

where $|W|$ is the cardinality of a set W . The threshold ϵ_{cut} is based on Hoeffding’s inequality⁸. At each time step, ADWIN checks every possible split of W .

The following definition formalizes the notion of a window and the shrinking of the current window. See also Figure 3 for illustration.

Definition 2 (Detection times) *The time step t is a detection time of ADWIN if $|W(t+1)| \leq |W(t)|$. In addition, ADWIN “shrinks” the window at time t if time step t is a*

8. Note that ϵ_{cut} above is slightly different from the original ADWIN where the harmonic mean of $|W_1|$ and $|W_2|$ is used.

detection time. We define the breakpoint as the last round of W_2 from the split $W = W_1 \cup W_2$ (Figure 3). We let \mathcal{T}_d be the set of all detection times.

Remark 3 Unlike the set of changepoints \mathcal{T}_c (Definition 1) that is only defined for an abruptly changing stream and is independent of any underlying algorithm, \mathcal{T}_d is defined for both abrupt stream and gradual stream. The detection time is a random variable that is defined via ADWIN.

3.2.1 BOUND FOR THE ESTIMATOR OF THE MEAN μ_t

Bifet and Gavaldà (2007) derived a bound on the false positive and false negative rates by replacing δ with $\delta' = \delta/|W(t)|$, to account for multiple tests⁹. However, the existing analysis is incomplete because it implicitly assumes that the window size $|W(t)|$ is deterministic.

In this paper, we use δ instead of δ' because it clarifies the analysis. The value of δ is in accordance with the possible multiplicity. With this aim, we introduce the notion of the window set and split and then introduce a confidence bound that holds with high probability.

Definition 4 (Window set) *The window set \mathcal{W} is the set of all the segments, which are the candidates of the current window W .*

$$\mathcal{W} = \{W' : W' = \{t', t' + 1, \dots, t''\}, 1 \leq t' < t'' < T\}. \quad (2)$$

Definition 5 (Split and estimator in a window) *Letting $W' = \{t', t' + 1, \dots, t''\}$, a split W_1, W_2 of a window $W' \in \mathcal{W}$ is defined as two disjoint subsets of W' such that $W_1 = \{t', t' + 1, \dots, m\}$ and $W_2 = \{m + 1, m + 2, \dots, t''\}$ for some $m \in W' \setminus \{t''\}$. For a window $W' \in \mathcal{W}$,*

$$\mu_{W'} = \frac{1}{|W'|} \sum_{s \in W'} \mu_s, \quad (3)$$

and its empirical estimator is

$$\hat{\mu}_{W'} = \frac{1}{|W'|} \sum_{s \in W'} x_s. \quad (4)$$

Proposition 1 (Uniform Hoeffding bound for the window set) *Let $p > 0$ be arbitrary. With probability $1 - 2/T^p$, we have the following:*

$$\forall W' \in \mathcal{W}, \quad |\mu_{W'} - \hat{\mu}_{W'}| \leq \sqrt{\frac{\log(T^{2+p})}{2|W'|}}. \quad (5)$$

Proposition 1 is derived using the Hoeffding inequality (Lemma 25, in Appendix C) over all possible $|W| \leq T^2$ windows. The bound above holds regardless of the randomness of the current window $W = W(t)$. Moreover, it holds for any W and any split $W_1 \cup W_2 = W$. We typically set $p = 1$ because this ensures a uniform bound with probability $1 - O(1/T)$, and an event with probability $O(1/T)$ is usually negligible. Our results are based on this bound.

⁹. There are $|W(t)| - 1$ ways to split $W(t) = W_1 \cup W_2$.

3.2.2 BOUND FOR THE TOTAL ERROR OF THE MEAN ESTIMATOR

In the following, we provide a bound on the following total error of the estimator $\hat{\mu}_W$:

Definition 6 (Total error) *The total error of estimator $\hat{\mu}_W$ is defined as follows:*

$$\text{Err}(T) = \sum_{t=1}^T |\hat{\mu}_{W(t)} - \mu_t|. \quad (6)$$

We first derive the error bound of ADWIN for a stationary stream, which directly follows from the false positive rate (Eq. (5)).

Theorem 7 (ADWIN in stationary environments) *Let the stream be stationary. Then, for ADWIN with $\delta = 1/T^3$, we have the following:*

$$\mathbb{E}[\text{Err}(T)] \leq \tilde{O}(\sqrt{T}). \quad (7)$$

Proof of Theorem 7 Equation (5) with $p = 1$, together with Definition 2 implies that no shrink occurs with probability $1 - 2/T$. Therefore,

$$\mathbb{E}[\text{Err}(T)] \leq \underbrace{T \times \frac{2}{T}}_{\text{case where at least one shrink occurs}} + \underbrace{\sum_{t=1}^T \sqrt{\frac{\log(T^3)}{2t}}}_{\text{case where no shrink occurs}} \quad (8)$$

$$\leq 1 + \sum_{t=1}^T \sqrt{\frac{\log(T^3)}{2t}} \quad (9)$$

$$\leq \sqrt{6T(\log T)} + 1, \quad \left(\text{by } \sum_{t=1}^T (1/\sqrt{t}) \leq 2\sqrt{T} \right) \quad (10)$$

which completes the proof. ■

The $\tilde{O}(\sqrt{T})$ error is the optimal rate because we can only identify the true value of μ_W up to $O(\sqrt{1/|W|})$. The error per time step $|\mu_t - \mu_W|$ is at least $\Omega(1/\sqrt{|W|}) = \Omega(1/\sqrt{t})$, and the total error is $\Omega(\sum_t 1/\sqrt{t}) = \Omega(\sqrt{T})$.

Theorem 7 states that ADWIN can learn from a stationary environment without any unnecessary shrinking. This statement contrasts with such methods as periodic resetting or fixed-size windowing algorithms which discard their entire memory after a fixed period. Having derived the learnability of ADWIN for a stationary stream, we are now interested in the property of ADWIN in the face of a nonstationary stream.

3.3 Analysis of ADWIN for abrupt changes

This section derives an error bound of ADWIN in the face of an abrupt change. As informally discussed in Bifet and Gavaldà (2007), ADWIN is able to detect abrupt changes if the changes are infrequent and gaps are detectable. Our results here are even more robust. Somewhat surprisingly, the bound in Theorem 8 does not depend on the detectability of the change. No matter how large or small the changes are and in which interval the changes occur, the algorithm's performance is bounded in terms of the number of changes.

Theorem 8 (Error bound of ADWIN under abrupt changes) *Let the environment be abrupt with M changepoints. The total error of ADWIN with $\delta = 1/T^3$ is bounded as*

$$\mathbb{E}[\text{Err}(T)] = \tilde{O}(\sqrt{MT}). \quad (11)$$

Proof Sketch Let $c(t)$ be the number of time steps after the last changepoint. First, we demonstrate that $|\mu_W - \mu_t| = \tilde{O}(1/\sqrt{c(t)})$ because ADWIN otherwise would shrink the window further (i.e., the event \mathcal{B} in the proof). Second, the window size $|W|$ is also bounded below. Given a sufficiently high confidence parameter, $|W| = O(c(t))$ and thus $|\mu_W - \hat{\mu}_W| = \tilde{O}(1/\sqrt{c(t)})$. Combining these two yields the bound of

$$|\hat{\mu}_W - \mu_t| \leq |\mu_W - \mu_t| + |\mu_W - \hat{\mu}_W| = \tilde{O}(1/\sqrt{c(t)}). \quad (12)$$

Using the Cauchy-Schwarz inequality yields the bound of Eq. (11). ■

The formal proof is found in Appendix D.1.

Remark 9 Theorem 8 implies the optimality of ADWIN under abrupt drift. To observe this, assume that a changepoint exists every T/M time steps. As discussed in Theorem 7, the optimal rate of error for each interval between changepoints is $\tilde{\Theta}(\sqrt{T/M})$, and the total error should be $\tilde{\Theta}(M \times \sqrt{T/M}) = \tilde{\Theta}(\sqrt{MT})$, which matches Theorem 8.

3.4 Analysis of ADWIN for gradual changes

In this section, we analyze ADWIN for a gradually changing stream, where the mean μ_t changes slowly with constant b (Definition 1). We consider the error for $b = T^{-d}$ by following the framework of Wei and Srivastava (2018).

Theorem 10 (Error bound of ADWIN under gradual changes) *Assume that there exists $d \in (0, 3/2)$ such that, the stream is gradually changing with its constant satisfies $b \leq T^{-d}$. Then, the performance of ADWIN with $\delta \leq 1/T^3$ is bounded as follow:*

$$\mathbb{E}[\text{Err}(T)] = \tilde{O}(T^{1-d/3}). \quad (13)$$

Proof Sketch We establish two lemmas in the appendix. Lemma 26 states that the drift $|\mu_s - \mu_t|$ for any two time steps s, t in the current window $W(t)$ is bounded by $\tilde{O}(bN + \sqrt{1/N})$ for any $N \geq |W(t)|$. Moreover, Lemma 27 states that the window is likely to grow until $|W| = O(b^{-2/3})$. Combining these two lemmas yields $|\mu_s - \mu_t| = O(b^{1/3}) = O(T^{-d/3})$. ■

The formal proof is found in Appendix D.4. Theorem 10 states that if the change is slow compared with the current scale of interest T , then the error per time step $\text{Err}(T)/T$ approaches zero.

In this section, we have bounded the total error for streams with abrupt changes (Section 3.3) as well as for streams with gradual changes (Section 3.4). This concludes our discussion on ADWIN. In subsequent sections, we consider the idea of combining ADWIN with the MAB setting where multiple streams are involved, and only a selected subset of streams are observable.

Algorithm 2 Thompson sampling (TS)

Require: Set of arms $[K]$.

- 1: **Initialize:** $W(t) = \emptyset$
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: **for** $i = 1, \dots, K$ **do**
 - 4: $S_i(t) = \sum_{s \in W(t)} \mathbf{1}[i = I(s)]x_{i,s}$, $N_i(t) = \sum_{s \in W(t)} \mathbf{1}[i = I(s)]$.
 - 5: $\theta_{i,t} \sim \text{Beta}(S_i(t) + 1, N_i(t) - S_i(t) + 1)$.
 - 6: **end for**
 - 7: Play arm $I(t) := \arg \max_i \theta_{i,t}$. ▷ Posterior sampling
 - 8: Receive reward $x(t)$.
 - 9: Update window $W(t + 1) = W(t) \cup \{t\}$.
 - 10: **end for**
-

Algorithm 3 Kullback–Leibler UCB (KL-UCB).

Require: Set of arms $[K]$.

- 1: **Initialize:** $W(t) = \emptyset$
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: **for** $i = 1, \dots, K$ **do**
 - 4: $S_i(t) = \sum_{s \in W(t)} \mathbf{1}[i = I(s)]x_{i,s}$, $N_i(t) = \sum_{s \in W(t)} \mathbf{1}[i = I(s)]$.
 - 5: $\hat{\mu}_{i,W(t)} = S_i(t)/N_i(t)$ where $0/0 = 1$.
 - 6: $U_i(t) = \max\{q \in [0, 1] : N_i(t)d_{\text{KL}}(\hat{\mu}_{i,W(t)}, q) \leq \log(t/N_i(t))\}$. ▷ KL-UCB index
 - 7: **end for**
 - 8: Play arm $I(t) := \arg \max_i U_i(t)$. ▷ Arm with the largest UCB index
 - 9: Receive reward $x(t)$.
 - 10: Update window $W(t + 1) = W(t) \cup \{t\}$.
 - 11: **end for**
-

4. Multi-armed Bandits

Up to now, we have studied the single-stream learning problem, in which every values of the stream are observed. From this section, we study the Multi-Armed Bandit (MAB) problem (Thompson, 1933; Robbins, 1952). This problem involves K streams, and a forecaster can only observe one of these streams at each time step. We first formalize the problem setting we consider.

Let there be K arms. At each time step $t = 1, \dots, T$, the forecaster selects an arm $I(t) \in [K]$, and then receives a reward $x_{I(t),t}$ from the selected arm. We also use the term “environment” to describe K streams that generates rewards $x_{i,t}$. Let $(\mu_{i,t})_{t=[T]}$ be the i -th stream. The stationary (resp. abruptly-changing and gradually-changing) environment is defined as a set of K stationary (resp. abruptly-changing and gradually-changing) streams, respectively, where these streams are defined in Definition 1.

The goal of the forecaster is to maximize the sum of the rewards of the selected arms by using a good algorithm, and the performance of a bandit algorithm is usually measured by the regret, which is defined as the difference between the expected reward of the best arm and the expected reward of the arms selected by the algorithm. That is,

$$\text{reg}(t) = \max_{i \in [K]} \mu_{i,t} - \mu_{I(t),t} \quad (14)$$

and

$$\text{Reg}(T) = \sum_{t=1}^T \text{reg}(t). \quad (15)$$

In the K -armed bandit problem, Thompson sampling (Thompson, 1933, TS) and the upper confidence bound (Lai and Robbins, 1985; Auer et al., 2002, UCB) are widely known algorithms that balance exploration and exploitation. Among several variants of UCB, the one that uses KL divergence, which is called KL-UCB (Garivier and Cappé, 2011), is known to have a state-of-the-art performance in stationary environments. TS and KL-UCB in our notation are described in Algorithms 2 and 3, respectively. Here, $\mathbf{1}[\mathcal{A}] = 1$ if \mathcal{A} or 0 otherwise.

Under a changing environment, the performance of these algorithms is no longer guaranteed. The next section extends the MAB framework for the case of abruptly-changing and gradually-changing streams, by combining adaptive window technology with bandit algorithms.

These algorithms are designed to deal with stationary environments: The regret of KL-UCB and TS is bounded as follows:

Definition 11 (Stationary regret) *For a stationary environment (i.e., $\mu_{i,t} = \mu_i$), for ease of discussion, we assume $\mu_1 > \mu_2 > \dots > \mu_K$. Let the suboptimality gap be $\Delta_i = \mu_1 - \mu_i$. A bandit algorithm has a logarithmic stationary regret if a universal constant C^{st} exists such that*

$$\mathbb{E}[\text{Reg}(T)] \leq C^{\text{st}} \sum_{i \neq 1} \frac{\log T}{\Delta_i}, \quad (16)$$

where $\text{Reg}(T)$ denotes the regret when we run the bandit algorithm.

In the literature of multi-armed bandit problem, this bound is often referred to as the distribution-dependent bound in the sense that it depends on Δ_i . The inverse of Δ_i defines the hardness of the instance. Note that the logarithmic bound and its first-order dependence on Δ_i^{-1} is optimal (Lai and Robbins, 1985; Auer et al., 2002; Kaufmann et al., 2012; Agrawal and Goyal, 2013).

4.1 Drift-tolerant regret

While the stationary regret is the most widely studied measure of the performance of stationary bandits, it cannot characterize the performance of a bandit algorithm under a changing environment. The following extends this to the environments with drifts.

Definition 12 (Drift-tolerant regret) *Assume a nonstationary environment that is abruptly or gradually changing. Let*

$$\Delta_i = \mu_{1,1} - \mu_{i,1} \tag{17}$$

be the gap at $t = 1$. Let

$$\epsilon(t) = \max_{s \leq t} \max_i |\mu_{i,s} - \mu_{i,1}|, \tag{18}$$

which is the maximum drift of the arms by time step t . For $c > 0$, let

$$\text{Reg}_{\text{tr}}(T, c) := \sum_t (\text{reg}(t) - c\epsilon(t))_+,$$

where $(x)_+ = \max(x, 0)$. Namely, $\text{Reg}_{\text{tr}}(T, c)$ is the regret where the regret proportional to drift is tolerated. A bandit algorithm has logarithmic drift-tolerant regret if a factor $C^{\text{dt}} = O(1)$ exists such that

$$\mathbb{E}[\text{Reg}_{\text{tr}}(T, C^{\text{dt}})] \leq C^{\text{dt}} \sum_{i \neq 1} \frac{\log T}{\Delta_i}. \tag{19}$$

Remark 13 Definition 12 is a generalization of the distribution-dependent regret in the literature of stationary bandits that allows a drift proportional to $\epsilon(t)$. Letting $\epsilon(t) = 0$ immediately derives the stationary regret of Definition 11.

As the following lemma characterizes, Thompson sampling (TS, Algorithm 2) is drift-tolerant.

Lemma 14 *Assume that rewards $x_{1,t}, x_{2,t}, \dots$ are binary (i.e., 0 or 1). Then, TS has logarithmic drift-tolerant regret.*

The formal proof is in Appendix E.2. Note that it is not very difficult to derive the drift-tolerant regret of KL-UCB by following similar steps as the proof of Lemma 14.

Although the drift-tolerant regret characterizes the performance of TS and KL-UCB under nonstationary environments, these algorithms are not good enough to deal with nonstationarity. This is because $\epsilon(t)$ increases over time and can lead to a meaningless bound of $\Theta(T)$ when $\epsilon(t)$ reaches $\Theta(1)$. A capable nonstationary bandit algorithm should be able to forget past data when it identifies a substantial drift. In the following section, we will merge these base-bandit algorithms with ADWIN.

5. Analysis of MAB for Globally Nonstationary Environments

This section proposes the combination of adaptive windowing and bandit algorithms.

5.1 Adaptive resetting bandit algorithm

Algorithm 4 ADS-bandit

Require: Set of arms $[K]$, confidence level δ , base-bandit algorithm

- 1: Initialize the base-bandit algorithm.
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: $(I(t), X(t)) = \text{BASE-BANDIT}(W(t))$ \triangleright Do one time step of the base-bandit algorithm
 - 4: **if** there exists a split $W(t+1) = W_1 \cup W_2$ such that $|\hat{\mu}_{i,W_1} - \hat{\mu}_{i,W_2}| \geq \epsilon_{\text{cut}}^\delta$ **then**
 - 5: Update the window $W(t+1) = W_2$ of the base-bandit algorithm.
 - 6: **end if**
 - 7: **end for**
-

Algorithm 5 ADR-bandit

Require: Set of arms $[K]$, confidence level δ , monitoring parameter $N \in \mathbb{N}$, base-bandit algorithm

- 1: Initialize base-bandit algorithm.
 - 2: **for** $l = 1, 2, \dots, \lceil \log_2(T/(KN) + 1) \rceil$ **do**
 - 3: **for** $t = (2^{l-1} - 1)KN + 1, \dots, \min\{(2^l - 1)KN, T\}$ **do**
 - 4: **if** $l \geq 2$ and $t = 0 \pmod{K}$ **then**
 - 5: Pull $i^{(l-1)}$. \triangleright Monitoring arm of the previous block.
 - 6: **else if** $l \geq 2$ and $t = 1 \pmod{K}$ and $t \geq (2^{l-1} - 2)KN + 1$ **then**
 - 7: Pull $i^{(l)}$. \triangleright Monitoring arm of the current block.
 - 8: **else**
 - 9: $(I(t), X(t)) = \text{BASE-BANDIT}(W(t))$ \triangleright Do one step of the base-bandit algorithm
 - 10: **end if**
 - 11: **if** there exists $W(t+1) = W_1 \cup W_2$ such that $|\hat{\mu}_{i,W_1} - \hat{\mu}_{i,W_2}| \geq \epsilon_{\text{cut}}^\delta$ **then**
 - 12: Reset the entire algorithm with $T := T - t$.
 - 13: **end if**
 - 14: **if** $t = KN$ **then** \triangleright The end of the first block $l = 1$.
 - 15: Set $i^{(1)} = \arg \max_{i \in [K]} N_i^{(1)}$ \triangleright See Eq. (20) for $N_i^{(1)}$.
 - 16: **else if** $l \geq 2$ and $t = (2^{l-1} - 2)KN$ **then** \triangleright Before the beginning of the last subblock of the block l .
 - 17: Set $i^{(l)} = \arg \max_{i \in [K]} N_i^{(l)}$. \triangleright See Eq. (20) for $N_i^{(l)}$
 - 18: **end if**
 - 19: **end for**
 - 20: **end for**
-




Figure 4: Illustration of the blocks in Algorithm 5. Here, “ $i^{(l)}$ ” for each block l represents the monitoring arm, which is drawn at least N times in each subblock. In the last subblock, two monitoring arms i^{l-1} and i^l are drawn at least N times, for the aim of monitoring consistency.

ADS-bandit (adaptive shrinking bandit; Algorithm 4) combines adaptive windows with a *base-bandit* algorithm such as TS or KL-UCB. In this algorithm, each arm has an associated adaptive window change detector. When such a detector identifies a split in the window W into W_1 and W_2 , it follows the ADWIN procedure (as explained in Section 3) of reducing the window size to retain W_2 .

We explain the steps of Algorithm 4 with TS as a base-bandit algorithm. Before the first time step, the base-bandit algorithm is initialized (Line 1 of Algorithm 2). At each time step $t = 1, 2, \dots$, Line 3 of Algorithm 4 runs an iteration of the base-bandit (Line 2 of Algorithm 2) to obtain $I(t)$ and corresponding reward $x(t) = x_{I(t),t}$. Afterward, Line 4 of Algorithm 4 checks whether a change occurred for each arm. If at least one change is detected, it shrinks window $W(t+1)$ to W_2 and drops the memory of W_1 .

Although we will demonstrate the superior empirical performance of Algorithm 4 in Section 6, it is highly nontrivial to derive a regret bound of Algorithm 4. The ADR-bandit (adaptive resetting bandit; Algorithm 5) addresses the two primary issues for the analysis as follows: First, the updated window W_2 after a shrink can include some amount of observations before the changepoint¹⁰. To address this issue, ADR-bandit resets the entire window for each detection of a changepoint. Second, it is hard to guarantee a consistent estimation of a single arm. For example, in a data stream with an abrupt change, if only arm 1 is selected prior to the change and only arm 2 is selected after the change, the algorithm is incapable of identifying this change. Although this particular scenario is unlikely to happen, it poses a challenge to the analysis of the algorithm. To address this issue, ADR-bandit introduces the *monitoring arm* that we continue to draw with some frequency.

ADR-bandit (Algorithm 5) works as follows. After each observation, if one of the change point detectors identifies a split in the window W into W_1 and W_2 , it resets the entire algorithm (Line 12) to start with $W(t+1) = \emptyset$. Moreover, Algorithm 5 divides the rounds into blocks $l = 1, 2, 3, \dots$ (Line 2). Each block consists of $O(2^{l-1})$ subblocks, and each subblock consists of KN rounds. Before the beginning of the last subblock of each l , it

10. We can cope with this problem in the case of single-stream ADWIN (Section 3.3) since μ_W eventually converges to μ_t as W grows. However, under bandit feedback, our algorithms focus on optimal arms and draw suboptimal arms less often. Thus, bounding the error of the estimators for suboptimal arms is hard if the current window W still includes observations before the changepoint.

determines the monitoring arm $i^{(l)}$ (Line 17), and then $i^{(l)}$ is drawn once in each K rounds during the next block. Here, monitoring arm is chosen to be $\arg \max_i N_i^{(l)}$, where

$$N_i^{(l)} = \begin{cases} |\{s \leq KN : I(s) = i\}| & l = 1 \\ |\{(2^{l-1} - 1)KN + 1 \leq s \leq (2^l - 2)KN : I(s) = i, \text{Base-bandit algorithm is called}\}| & l \geq 2 \end{cases}, \quad (20)$$

be the number of draws of arm i in the current block l . Intuitively speaking $N_i^{(l)}$ is the number of rounds in which the arm i is chosen by the base bandit algorithm.

5.2 Properties of Algorithm 5

As shown in Figure 4, the duration of two monitoring arms $i^{(l-1)}, i^{(l)}$ are overlapped so that the following properties are satisfied:

Definition 15 (Monitoring consistency) *ADR-bandit with any base-bandit algorithm has the following two properties:*

- Let $t \geq KN + 1$ (i.e., any round in $l \geq 2$). Then, at least one of the arms $\{i^{(l-1)}, i^{(l)}\}$ satisfies the following: This arm was drawn at least N times before round t and will be drawn at least N times within the next KN rounds.
- For any block $l = 1, 2, \dots$, there exists at least one arm that is drawn at least N times for each subblock in l .

The first property in Definition 15 is used for the abrupt case, and the second property in Definition 15 is used for the gradual case.

Moreover, the following theorem states that ADR-bandit is designed so that it does not compromise the performance of the base-bandit algorithm when no reset takes place.

Lemma 16 (Regret due to monitoring) *Assume that we run Algorithm 5 with a base-bandit algorithm with $K \geq 3$. Assume that Algorithm 5 has not reset itself until round S . Let*

$$\mathcal{T}^{\text{base}} = \{t \in [S] : \text{base-bandit algorithm is evoked}\}$$

be the set of rounds where the base-bandit algorithm is evoked (i.e., when Line 9 in Algorithm 5 is executed). Then, the following holds:

$$\sum_{t=1}^S \text{reg}(t) \leq 4 \left(\sum_{t \in \mathcal{T}^{\text{base}}} \text{reg}(t) \right) + \sum_{t=1}^S 3\epsilon(t). \quad (21)$$

The proof of Lemma 16 is found in Appendix E.3. The first term is a constant multiplication of the base-bandit regret. As a result, it can fully utilize the logarithmic drift-tolerant regret of the base-bandit algorithm, which we formalized in the following Lemma.

Lemma 17 (Constant multiplication) *Assume that we run ADR-bandit algorithm with a base-bandit algorithm that has logarithmic drift-tolerant regret (Definition 12). Assume that no reset has occurred¹¹ until round S . Then, there exists a constant $C^{\text{dt}} = O(1)$ that is independent of S , and the regret up to S is bounded as*

$$\mathbb{E}[\text{Reg}(S)] \leq C^{\text{dt}} \left(\sum_{i \neq 1} \frac{\log T}{\Delta_i} + \mathbb{E} \left[\sum_{t=1}^S \epsilon(t) \right] \right). \quad (22)$$

Proof of Lemma 17 Definition 12 and Lemma 16 imply

$$\mathbb{E} \left[\sum_{t \leq S} (\text{reg}(t) - C^{\text{dt}} \epsilon(t))_+ \right] \leq \mathbb{E} \left[\sum_{t \leq T} (\text{reg}(t) - C^{\text{dt}} \epsilon(t))_+ \right] \leq C^{\text{dt}} \sum_{i \neq 1} \frac{\log T}{\Delta_i},$$

which implies Eq. (22). ■

5.3 Regret bounds of ADR-bandit algorithms

Having defined the ADR-bandit algorithm and the related properties, we state the main results regarding the performance of these algorithms in stationary, abruptly, or gradually changing environments.

Theorem 18 (Regret bound of ADR-bandit, stationary case) *Let the environment be stationary. Let the base-bandit algorithm has a logarithmic stationary regret. Then, the regret of Algorithm 5 with $\delta = 1/T^3$ is bounded as follows:*

$$\mathbb{E}[\text{Reg}(T)] \leq C^{\text{st}} \sum_i \frac{\log T}{\Delta_i} + O(1). \quad (23)$$

Theorem 18 states that changepoint detectors do not deteriorate the performance of the base-bandit algorithm. The proof directly follows from the following facts. First, adaptive windows do not make a false reset with at least a probability of $1 - O(K/T)$. Second, the regret of the ADR-bandit algorithm is equal to the base-bandit algorithm when no reset occurs (Lemma 17). For completeness, we provide the proof of Theorem 18 in Appendix E.1. This result has a striking difference from most existing nonstationary bandit algorithms, which have the cost of a forced exploration of $\Omega(\sqrt{T})$. In the following sections, we derive the regret bound in abruptly and gradually changing environments.

We next derive a regret bound for an abruptly changing environment. We first define the global changes that we assume on the nature of the streams (Definition 19) and then state the regret bound.

Definition 19 (Globally abrupt changes) *Let M be the number of change points and $\mathcal{T}_c = \{T_1, T_2, \dots, T_M\} = \{t : \exists i \in [K] \mu_{i,t} \neq \mu_{i,t+1}\}$ be the set of changepoints. Moreover, $(T_0, T_{M+1}) = (0, T)$. The m -th changepoint is a global change with C^{ab} if*

$$\max_{i,j \in [K], t \in \mathcal{T}_c} \frac{|\mu_{j,t} - \mu_{j,t+1}|}{|\mu_{i,t} - \mu_{i,t+1}|} \leq C^{\text{ab}} \quad (24)$$

11. Note that S is a stopping time (random variable).

is finite.

Intuitively, C^{ab} is the maximum ratio of changes among arms. Assuming it as a constant indicates that any modification in one arm is proportionally mirrored in all the other arms.

Definition 20 (Detectability) *For the m -th changepoint, let $\epsilon_m = \min_i |\mu_{i,m} - \mu_{i,m+1}|$. We let $U(\epsilon) = (\log(T^3))/(2\epsilon^2)$. The m -th changepoint is detectable if $T_m - T_{m-1}, T_{m+1} - T_m \geq 48KU(\epsilon_m)$.*

The detectability assumption is essentially the same as those in the literature of active nonstationary bandit algorithms, such as Cao et al. (2019) and Besson et al. (2022).

Theorem 21 (Regret bound of ADR-bandit under abrupt changes) *Let the base-bandit algorithm have a logarithmic drift-tolerant regret bound (Definition 12). Let the environment have M detectable globally abrupt changes (Definitions 19 and 20) with $C^{\text{ab}} > 0$. Then, the regret of the Algorithm 5 is bounded as follows:*

$$\mathbb{E}[\text{Reg}(T)] = \tilde{O}(\sqrt{MKT}), \quad (25)$$

if $\delta = 1/T^3$, $N \geq 16U(\epsilon_m)$, $KN \leq (T_{m+1} - T_m)/3$ holds for all m .

Proof Sketch The proof sketch is as follows.

- With a high probability, the algorithm resets itself within $16KU(\Delta_m)$ time steps after m -th changepoint there exists a corresponding reset.
- By the sublinear drift-tolerant regret of the algorithm the expected regret between m -th and $(m+1)$ -th changepoint is $\tilde{O}(\sqrt{K(T_{m+1} - T_m)})$, where we have used the standard discussion of distribution-independent regret (c.f., p32 in Bubeck, 2010) that applies the Cauchy-Schwarz inequality.
- The desired bound is yielded by summing the regret over all change points and another application of the Cauchy-Schwarz inequality on $\sum_m (T_m - T_{m-1}) \leq T$.

The formal proof is found in Appendix E.4. ■

Theorem 21 implies that the ADR-bandit learns the environment when the changes are infrequent (i.e., $M = o(T)$) and the changes are detectable. This bound matches many existing active algorithms, such as those in Cao et al. (2019); Besson et al. (2022).

We next analyze the performance of ADR-bandit with a gradually changing environment. We employ a technique similar to that in Section 3.4 for single-stream ADWIN. We first describe the global changes in a gradually changing environment and then state a regret bound.

Assumption 22 (Globally gradual changes) *A set of K streams has globally gradual changes if a constant $C^{\text{gr}} \in (0, 1]$ exists such that*

$$|\mu_{i,t} - \mu_{i,s}| \geq C^{\text{gr}} |\mu_{j,t} - \mu_{j,s}| \quad (26)$$

holds for any two arms $i, j \in [K]$ and any two time steps $t, s \in [T]$.

Intuitively, Assumption 22 states that all arms are drifting in a coordinated manner. The drift on the mean of an arm is proportional to the drift on the means of other arms.

Theorem 23 (Regret bound of the ADR-bandit under gradual changes) *Let the base-bandit algorithm have a logarithmic drift-tolerant regret (Definition 12) and have monitoring consistency (Definition 15). Let the environment be gradual with change speed b and d be such that $b = T^{-d}$. Let the changes be global (Assumption 22). Then the regret of Algorithm 5 with $\delta = 1/T^3$, $N = \tilde{\Theta}((bK)^{-2/3})$ is bounded as:*

$$\mathbb{E}[\text{Reg}(T)] = \tilde{O}\left(\sqrt{KT^{1-d/3}}\right). \quad (27)$$

Proof Sketch The proof sketch is as follows.

- Using Lemma 34, the current window grows until $|W| = O(b^{-2/3})$ with a high probability, which implies that the number of resets (= detection times) M_d is at most $\tilde{O}(Tb^{2/3}) = \tilde{O}(T^{1-2d/3})$.
- With a high probability, $|\mu_t - \mu_W| = \tilde{O}(bKN) + \tilde{O}\left(\sqrt{\log T/N}\right)$ always holds (Lemma 33). This term is minimized when $N = \tilde{\Theta}((bK)^{-2/3})$ and $|\mu_t - \mu_W| = \tilde{O}((bK)^{1/3})$.
- Similarly to abrupt case, the regret between the m -th and $(m+1)$ -th resets is bounded in expectation as $\tilde{O}(\sqrt{K(T_{d,m+1} - T_{d,m})})$ plus the drift term of $\sum_t \tilde{O}((bK)^{1/3})$.
- By using the Cauchy-Schwarz inequality on $\sum_{m=1}^{M_d} T_{d,m+1} - T_{d,m} = T$ and $M_d = \tilde{O}(T^{1-2d/3})$, we have $\sum_{m=1}^{M_d} \tilde{O}(\sqrt{K(T_{d,m+1} - T_{d,m})}) = \tilde{O}(\sqrt{M_dKT}) = \tilde{O}(\sqrt{KT^{1-d/3}})$. Moreover, $\sum_t \tilde{O}((bK)^{1/3}) = T\tilde{O}((bK)^{1/3}) = \tilde{O}(K^{1/3}T^{1-d/3})$.

The formal proof is in Appendix E.5. ■

Theorem 23 states that ADR-bandit learns the environment when the change is sufficiently slow (i.e., $b = o(1)$). Unlike most nonstationary bandit algorithms, our algorithm design does not require prior knowledge of b .

Remark 24 (Optimality of the regret bound) The rate matches the $\tilde{\Theta}(V^{1/3}T^{2/3})$ lower bound of Besbes et al. (2019) up to a logarithmic order. Under a gradually changing environment, the total variation is $V := \sum_t |\mu_{i,t} - \mu_{i,t+1}| = bT = O(T^{1-d})$; thus, $V^{1/3}T^{2/3} = O(T^{1-d/3})$. A recent work by Krishnamurthy and Gopalan (2021) shows the optimality of this rate for gradually changing streams.¹² Note that, unlike Besbes et al. (2019); Krishnamurthy and Gopalan (2021), our guarantee in the nonstationary setting is limited to the case of global changes.

In summary, ADR-bandit adapts to both the abruptly and gradually changing environments under mild knowledge on the speed of changes if the base-bandit satisfies has a logarithmic drift-tolerant regret bound. We demonstrate that TS has this bound and consider that it is not very difficult to derive the same bound for KL-UCB.

¹² The result from Trovò et al. (2020) has a better rate of $\tilde{O}(T^{1-d})$. However, they require an additional assumption (Assumption 2, therein) that essentially states that the two arms are not too close to distinguish for a long time.

6. Experiments

This section reports the empirical performance of our proposed algorithms in simulations. We release the source code to reproduce our experiments on GitHub¹³.

6.1 Environments

The rewards from arm i are drawn from a Bernoulli distribution with parameter μ_i that we determine for each arm. We define the following synthetic environments:

- **Stationary:** We define a stationary environment (i.e., with no change) with 100 arms where the mean $\mu_i = (101 - i)/100$ never changes over time, i.e., the rewards of arms are evenly distributed between 0 and 1.
- **Gradual:** We define a gradually changing environment with 100 arms with $\mu_{i,t} = \frac{(T-t+1)}{T}\mu_{i,1} + \frac{t-1}{T}(1 - \mu_{i,1})$, where $\mu_{i,1} = (101 - i)/100$. In this environment, the mean of each arm is the same as the stationary environment at time $t = 1$, but evolves gradually toward $1 - \mu_{i,1}$ as t increases up to T .
- **Abrupt:** We define an abruptly changing environment with 100 arms and $\mu_{i,1} = (101 - i)/100$ at $t = 1$. Unlike the stationary environment, $\mu_{i,t} = 1 - \mu_{i,1}$ between round $t' = T/3$ and round $t'' = 2T/3$. In other words, the rewards of every arm change abruptly two times at $t = T/3, 2T/3$.
- **Abrupt local:** We define an environment with 100 arms in which only a portion of the arms change abruptly, i.e., the global change assumption does not hold. As in the other environment, $\mu_{i,1} = (101 - i)/100$ at $t = 1$, but for $i \in [1, 10]$, $\mu_{i,t} = 0.5$ between round $t' = T/3$ and round $t'' = 2T/3$. In other words, only the top-10 arms change abruptly two times, and the other arms are stationary.

The stationary, abrupt, and gradual settings are in line with the assumptions from our algorithms (ADS-TS, ADR-TS, ADR-KL-UCB) as the changes are global. The abrupt local setting violates these assumptions because only a subset of the arms change.

In addition, we consider the following real-world scenarios:

- **Bioliq:** This dataset was released by Fouché et al. (2019). The rewards are generated from a stream of measurements between 20 sensors in a power plant for a duration of 1 week. The authors considered the task of monitoring the high correlation between those sensors and wanted to use bandit algorithms for efficient monitoring. Each pair of sensors is seen as an arm (i.e., there are $20 * 19 / 2 = 190$ arms) and the reward is 1 in case the correlation is greater than some threshold over the last 1000 measurements, otherwise 0. Figure 1 is the reward distribution for this environment; as we can see, arms tend to change periodically together.
- **Zozo** (Saito et al., 2021) is a real-world environment where the rewards are generated from an ad recommender on an e-commerce website. We use the data generated by

13. <https://github.com/edouardfouche/G-NS-MAB>

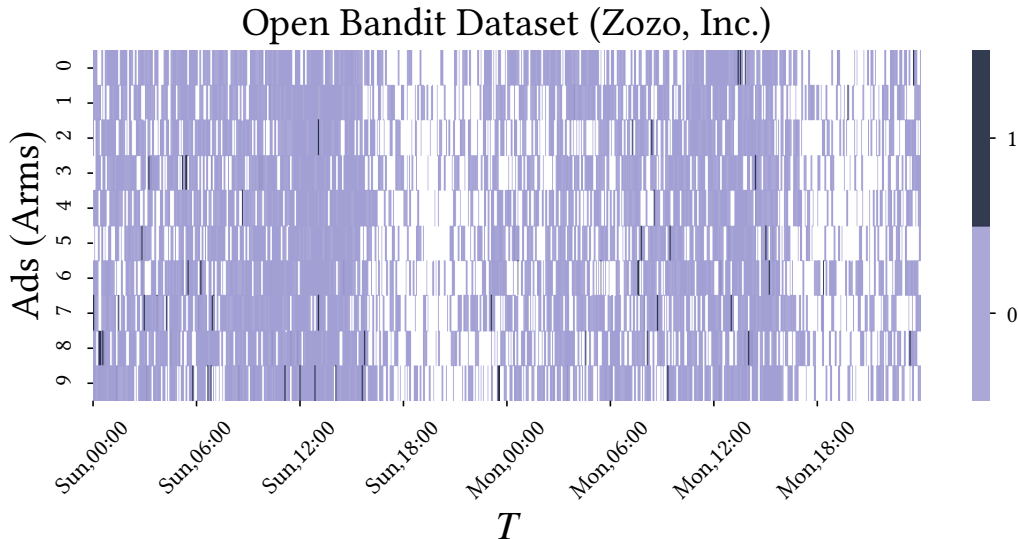


Figure 5: Rewards in the Zozo dataset. We chose ten ads (= arms, y -axis) among 80 ads. Dark blue and light blue indicate the reward 1 and 0, respectively. White indicates that the corresponding ad is not shown at that time step (no feedback). The ratio of reward 1 to reward 0 is between 0.5% to 3%. We only show the first two days, but the dataset consists of one week of data.

the uniform recommender from the duration of a week and adopted an unbiased offline evaluator (Li et al., 2011) for our experiment¹⁴. There are 80 different ads. Due to the sparseness of the rewards, we picked 10 arms. Since only a few of such ads are presented at each round to each user, we concatenated together all the ads presented within a second and assigned a reward of 1 to ads on which at least one user clicked. We assigned a reward of 0 to ads that were presented, but no user clicked on them. We skipped the round whenever the selected was not presented at all. Figure 5 shows the reward distribution of the dataset.

All the results are averaged over 100 runs. Note that these datasets match with our motivational examples (Examples 1 and 2).

6.2 Bandit algorithms

We compared the following bandit algorithms:

- **ADS-TS** is the ADS-bandit algorithm (Algorithm 4) with TS (Algorithm 2) as a base-bandit algorithm.
- **ADR-TS** and **ADR-KL-UCB** are the ADR-bandit algorithm (Algorithm 5) with TS (Algorithm 2) and KL-UCB (Algorithm 3) as a base-bandit algorithm, respectively.

14. See also <https://github.com/st-tech/zr-obp>

- Passive algorithms: **RExp3** (Besbes et al., 2019) is an adversarial bandit algorithm. Discounted UCB (**D-UCB**) (Garivier and Moulines, 2011) and Sliding Window TS (**SW-TS**) (Trovò et al., 2020), **SW-UCB#** (Wei and Srivastava, 2018) (comes in two variants: for Abrupt (A) and Gradual (G) changes) are bandit algorithms with finite memory.
- Active algorithms: **GLR-klUCB** (Besson et al., 2022) is an likelihood-ratio based change detector. It comes in two variants: for local and global changes, and we show the results of the latter. **M-UCB** (Cao et al., 2019) is a bandit algorithm using another change detector as ours. These algorithms involve forced exploration (uniform exploration of size $O(\sqrt{T})$) unlike our algorithms. **UCBL-CPD** and **ImpCPD** (Mukherjee and Maillard, 2019) are algorithms for global abrupt changes that are also free from forced exploration.

The original implementation of ADWIN requires $O(|W|K)$ times at each round, which can be large in the case of stationary environments. Bifet and Gavaldà (2007) introduced a more computationally efficient version of ADWIN, named ADWIN2, which only checks a logarithmic number of such splits. We adopt ADWIN2 in our simulations. Algorithmic hyperparameters are described in Section A.

6.3 Experimental results

First, Figure 6 compares ADR-bandit and existing algorithms in the stationary, abrupt, and gradual environments. In the stationary environment, ADR-TS has very low regret, which is consistent with the fact that its regret is logarithmic. In the two nonstationary environments with global changes, ADR-TS outperforms the other algorithms. Second, Figure 7 compares these algorithms in the abrupt local environment. Although the ADR-TS has linear regret as it sometimes do not detect these two local changes, it still outperforms the other algorithms. Finally, Figure 8 compares these algorithms against the Bioliq and Zozo datasets. The proposed algorithm consistently outperforms existing nonstationary bandits algorithms.

We provide in Section B additional simulations on the effect of resetting (compared with the shrinking of the original ADWIN), the sensitivity ADR-TS to the value of parameter δ , and results with $L > 1$. We find that ADR-TS is not very sensitive to the choice of parameter δ , as long as it is moderately small.

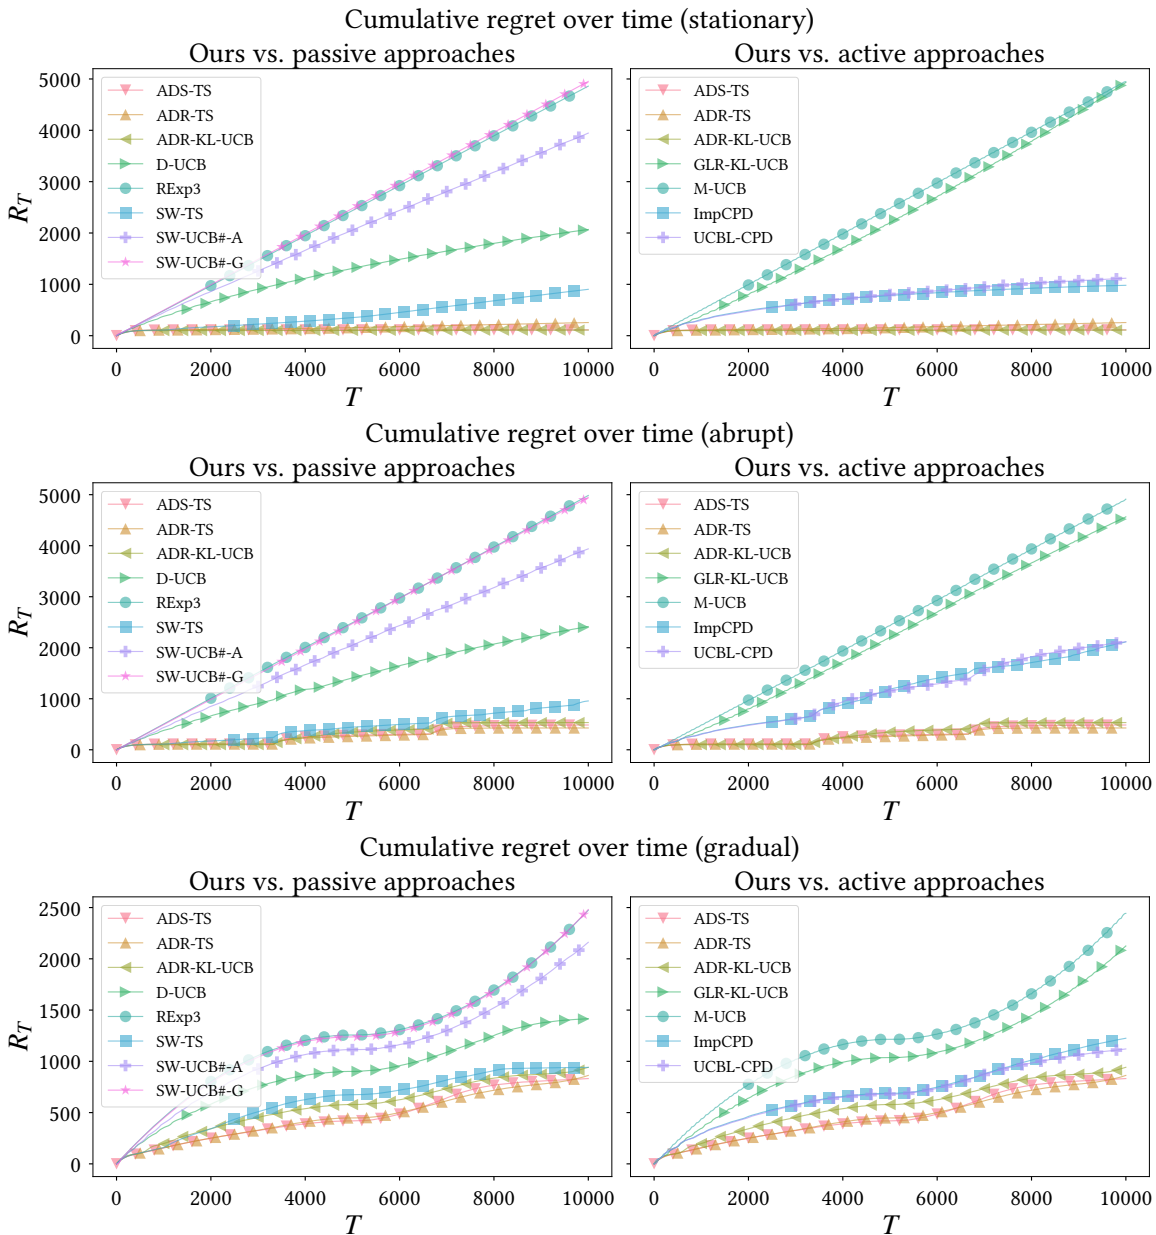


Figure 6: Regret of algorithms in stationary, abrupt, and gradual environments. Smaller regret (R_T : y-axis) indicates better performance.

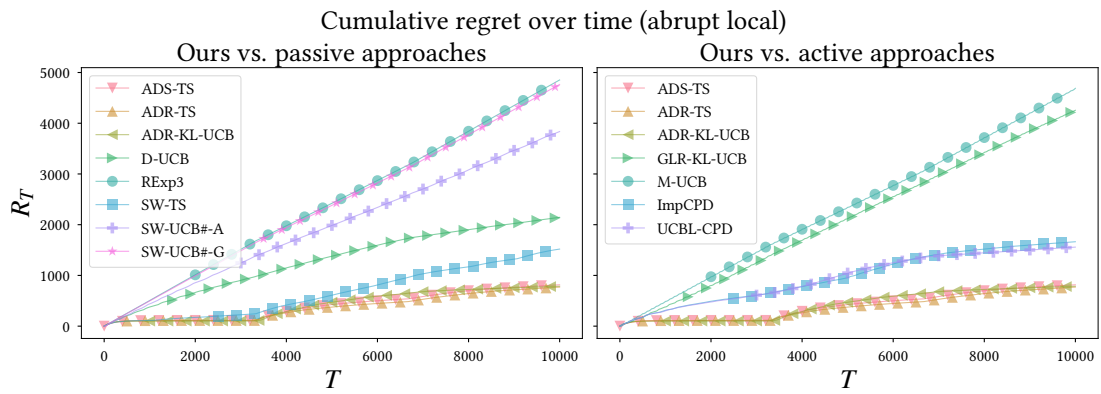


Figure 7: Regret of algorithms in abrupt local environments. Smaller regret (R_T : y-axis) indicates better performance. ADR-TS has linear regret after two local changes. However, it still outperforms algorithms that require uniform exploration.

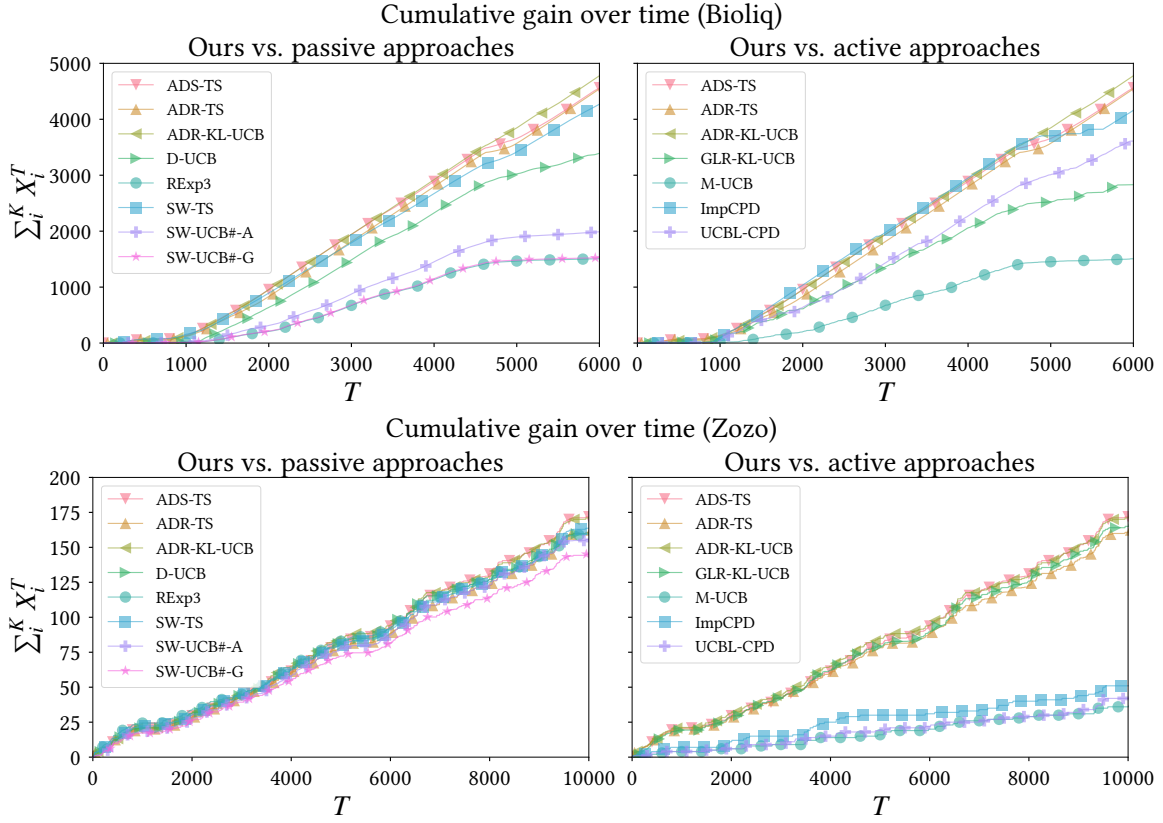


Figure 8: Performance of algorithms based on two real-world datasets (Bioliq and Zozo). Unlike synthetic environments, we do not have ground truth of $\mu_{i,t}$ and thus we report cumulative reward. Larger cumulative reward (y-axis) indicates better performance. M-UCB tends to have $\gamma \geq 1$ in our setting, which results in a uniform exploration.

7. Conclusions

We have expanded stationary bandit algorithms to nonstationary settings without requiring forced exploration. To this aim, we first analyzed the theoretical property of adaptive windows in a single-stream setting (Section 3). After that, we combined bandit algorithms (Section 4) with adaptive windowing by introducing ADR-bandit. Unlike existing algorithms, ADR-bandit does not act on the base-bandit algorithm unless change points are detected, and thus, it does not compromise the performance in a stationary environment. We have demonstrated its ability to detect global changes in Section 5 and tested it in simulated and real-world settings through experiments in Section 6.

Acknowledgments

We would like to express our sincere gratitude to the reviewers and the editors for their valuable comments and suggestions. Their insights and expertise have greatly contributed to improving the quality of our work. EF was supported by German Research Foundation (DFG) Research Training Group GRK 2153: “Energy Status Data – Informatics Methods for its Collection, Analysis and Exploitation” and by the Baden-Württemberg Foundation via the Elite Program for Postdoctoral Researchers. JH was supported by JSPS, KAKENHI Grant Number JP21K11747, Japan.

References

- Shipra Agrawal and Navin Goyal. Further optimal regret bounds for thompson sampling. In *AISTATS*, volume 31 of *JMLR Workshop and Conference Proceedings*, pages 99–107, 2013.
- Robin Allesiardo and Raphaël Féraud. EXP3 with drift detection for the switching bandit problem. *Proceedings of the 2015 IEEE International Conference on Data Science and Advanced Analytics, DSAA 2015*, (August), 2015.
- Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 322–331, Oct 1995.
- Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2-3):235–256, 2002.
- Peter Auer, Pratik Gajane, and Ronald Ortner. Adaptively tracking the best bandit arm with an unknown number of distribution changes. In *COLT*, volume 99, pages 138–158, 2019.
- Manuel Baena-Garcia, José del Campo-Ávila, Raúl Fidalgo, Albert Bifet, R Gavalda, and R Morales-Bueno. Early drift detection method. In *Fourth international workshop on knowledge discovery from data streams*, volume 6, pages 77–86, 2006.
- Omar Besbes, Yonatan Gur, and Assaf Zeevi. Optimal exploration–exploitation in a multi-armed bandit problem with non-stationary rewards. *Stochastic Systems*, 9(4):319–337, 2019. doi: 10.1287/stsy.2019.0033.
- Lilian Besson, Emilie Kaufmann, Odalric-Ambrym Maillard, and Julien Seznec. Efficient change-point detection for tackling piecewise-stationary bandits. *Journal of Machine Learning Research*, 23(77):1–40, 2022.
- Albert Bifet and Ricard Gavaldà. Learning from time-changing data with adaptive windowing. In *SDM*, pages 443–448. SIAM, 2007.
- Isvani Inocencio Frías Blanco, José del Campo-Ávila, Gonzalo Ramos-Jiménez, Rafael Morales Bueno, Agustín Alejandro Ortiz Díaz, and Yailé Caballero Mota. Online and Non-Parametric Drift Detection Methods Based on Hoeffding’s Bounds. *IEEE Trans. Knowl. Data Eng.*, 27(3):810–823, 2015.
- Sébastien Bubeck. *Bandits Games and Clustering Foundations*. Theses, Université des Sciences et Technologie de Lille - Lille I, June 2010. URL <https://theses.hal.science/tel-00845565>.
- Yang Cao, Zheng Wen, Branislav Kveton, and Yao Xie. Nearly optimal adaptive procedure with change detection for piecewise-stationary bandit. In *AISTATS*, volume 89, pages 418–427, 2019.

- Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In *NIPS*, pages 2249–2257, 2011.
- Yifang Chen, Chung-Wei Lee, Haipeng Luo, and Chen-Yu Wei. A new algorithm for non-stationary contextual bandits: Efficient, optimal and parameter-free. In *COLT*, volume 99, pages 696–726, 2019.
- Edouard Fouché, Junpei Komiyama, and Klemens Böhm. Scaling multi-armed bandit algorithms. In *KDD*, pages 1449–1459. ACM, 2019.
- João Gama and Gladys Castillo. Learning with local drift detection. In *ADMA*, volume 4093 of *Lecture Notes in Computer Science*, pages 42–55. Springer, 2006.
- João Gama, Pedro Medas, Gladys Castillo, and Pedro Pereira Rodrigues. Learning with drift detection. In *SBIA*, volume 3171 of *Lecture Notes in Computer Science*, pages 286–295. Springer, 2004.
- João Gama, Indre Zliobaite, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM Comput. Surv.*, 46(4):44:1–44:37, 2014.
- Aurélien Garivier and Olivier Cappé. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *COLT*, volume 19 of *JMLR Proceedings*, pages 359–376, 2011.
- Aurélien Garivier and Eric Moulines. On upper-confidence bound policies for switching bandit problems. In *Algorithmic Learning Theory - 22nd International Conference, ALT*, volume 6925 of *Lecture Notes in Computer Science*, pages 174–188. Springer, 2011.
- Paulo Mauricio Gonçalves, Silas Garrido Teixeira de Carvalho Santos, Roberto Souto Maior de Barros, and Davi Carnauba de Lima Vieira. A comparative study on concept drift detectors. *Expert Syst. Appl.*, 41(18):8144–8156, 2014.
- Yonatan Gur, Assaf J. Zeevi, and Omar Besbes. Stochastic multi-armed-bandit problem with non-stationary rewards. In *NIPS*, pages 199–207, 2014.
- Cédric Hartland, Nicolas Baskiotis, Sylvain Gelly, Michèle Sebag, and Olivier Teytaud. Change Point Detection and Meta-Bandits for Online Learning in Dynamic Environments. In *CAP 2007 : 9è Conférence francophone sur l'apprentissage automatique*, pages 237–250, July 2007.
- D. V. Hinkley. Inference about the change-point from cumulative sum tests. *Biometrika*, 58(3):509–523, 1971. ISSN 00063444.
- Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *ALT*, volume 7568 of *Lecture Notes in Computer Science*, pages 199–213. Springer, 2012.
- Levente Kocsis and Csaba Szepesvári. Discounted ucb. In *2nd PASCAL Challenges Workshop*, pages 784–791, 2006.

- Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Optimal Regret Analysis of Thompson Sampling in Stochastic Multi-armed Bandit Problem with Multiple Plays. In *ICML*, volume 37, pages 1152–1161, 2015.
- Bartosz Krawczyk, Leandro L. Minku, João Gama, Jerzy Stefanowski, and Michal Wozniak. Ensemble learning for data stream analysis: A survey. *Inf. Fusion*, 37:132–156, 2017.
- Ramakrishnan Krishnamurthy and Aditya Gopalan. On slowly-varying non-stationary bandits. *CoRR*, abs/2110.12916, 2021. URL <https://arxiv.org/abs/2110.12916>.
- T.L. Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4 – 22, 1985.
- Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *WWW*, pages 661–670. ACM, 2010.
- Lihong Li, Wei Chu, John Langford, and Xuanhui Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *WSDM*, pages 297–306. ACM, 2011.
- Fang Liu, Joohyun Lee, and Ness B. Shroff. A change-detection based framework for piecewise-stationary multi-armed bandit problem. In *AAAI*, pages 3651–3658, 2018.
- Jie Lu, Anjin Liu, Fan Dong, Feng Gu, João Gama, and Guangquan Zhang. Learning under concept drift: A review. *IEEE Trans. Knowl. Data Eng.*, 31(12):2346–2363, 2019.
- Odalric-Ambrym Maillard. Boundary crossing for general exponential families. In *ALT*, volume 76 of *Proceedings of Machine Learning Research*, pages 151–184, 2017.
- Joseph Charles Mellor and Jonathan Shapiro. Thompson sampling in switching environments with bayesian online change detection. In *AISTATS*, volume 31, pages 442–450, 2013.
- Subhojyoti Mukherjee and Odalric-Ambrym Maillard. Distribution-dependent and time-uniform bounds for piecewise i.i.d bandits. *CoRR*, abs/1905.13159, 2019.
- E. S. Page. Continuous inspection schemes. *Biometrika*, 41(1/2):100–115, 1954. ISSN 00063444.
- Herbert Robbins. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.*, 58(5):527–535, 09 1952.
- Gordon J. Ross, Niall M. Adams, Dimitris K. Tasoulis, and David J. Hand. Exponentially weighted moving average charts for detecting concept drift. *Pattern Recognit. Lett.*, 33(2):191–198, 2012.
- Yuta Saito, Shunsuke Aihara, Megumi Matsutani, and Yusuke Narita. Open bandit dataset and pipeline: Towards realistic and reproducible off-policy evaluation. In *Proceedings of the NeurIPS Track on Datasets and Benchmarks*, 2021.

- Aleka Seliniotaki, George Tzagkarakis, Vassilis Christophides, and Panagiotis Tsakalides. Stream correlation monitoring for uncertainty-aware data processing systems. In *IISA*, pages 342–347. IEEE, 2014.
- Julien Seznec, Pierre Ménard, Alessandro Lazaric, and Michal Valko. A single algorithm for both restless and rested rotting bandits. In *AISTATS*, volume 108 of *Proceedings of Machine Learning Research*, pages 3784–3794, 2020.
- Vaibhav Srivastava, Paul Reverdy, and Naomi Ehrich Leonard. Surveillance in an abruptly changing world via multiarmed bandits. In *CDC*, pages 692–697. IEEE, 2014.
- William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Francesco Trovò, Marcello Restelli, and Nicola Gatti. Sliding-Window Thompson Sampling for Non-Stationary Settings. *J. Artif. Intell. Res.*, 68:311–364, 2020.
- Chen-Yu Wei, Yi-Te Hong, and Chi-Jen Lu. Tracking the best expert in non-stationary stochastic environments. In *NIPS*, pages 3972–3980, 2016.
- Lai Wei and Vaibhav Srivastava. On abruptly-changing and slowly-varying multiarmed bandit problems. In *ACC*, pages 6291–6296. IEEE, 2018.

Table 2: List of tested hyperparameters of the algorithms. Bold letters indicate the ones reported in this paper. We set the confidence level of UCBL-CPD and ImpCPD to be the same as ADR-TS.

Algorithm(s)	Hyperparameters
ADS-TS and ADR-TS, ADR-KL-UCB	$\delta = 10^{-1}, 10^{-2}, \mathbf{10^{-3}}, 10^{-4}, 10^{-5}, 10^{-6}, 10^{-8}, 10^{-12}, 10^{-15}$
RExp3	$\Delta_T = 100, 500, \mathbf{1000}, 5000$
SW-TS	$W = 100, 500, \mathbf{1000}, 5000$
D-UCB	$\gamma = 0.7, 0.8, \mathbf{0.9}, 0.99, 0.999$
SW-UCB#-A	$\nu = \mathbf{0.1}, 0.2$ and $\lambda = 12.3$
SW-UCB#-G	$\kappa = \mathbf{0.1}, 0.2$ and $\lambda = 4.3$
GLR-klUCB	$\alpha = \sqrt{kA \log(T)/T}$ and $\delta = 1/\sqrt{T}$
M-UCB	$w = \mathbf{1000}, 5000$ and $M = \mathbf{10}, 100$
UCBL-CPD	$\delta = \mathbf{10^{-3}}$
ImpCPD	$\gamma = \mathbf{0.05}, \psi\epsilon_m^2 = \mathbf{10^3}$

Appendix A. Hyperparameters in Experiments

The hyperparameters of the algorithms are shown in Table 2.

Appendix B. Additional Experiments

This section reports the results of additional experiments.

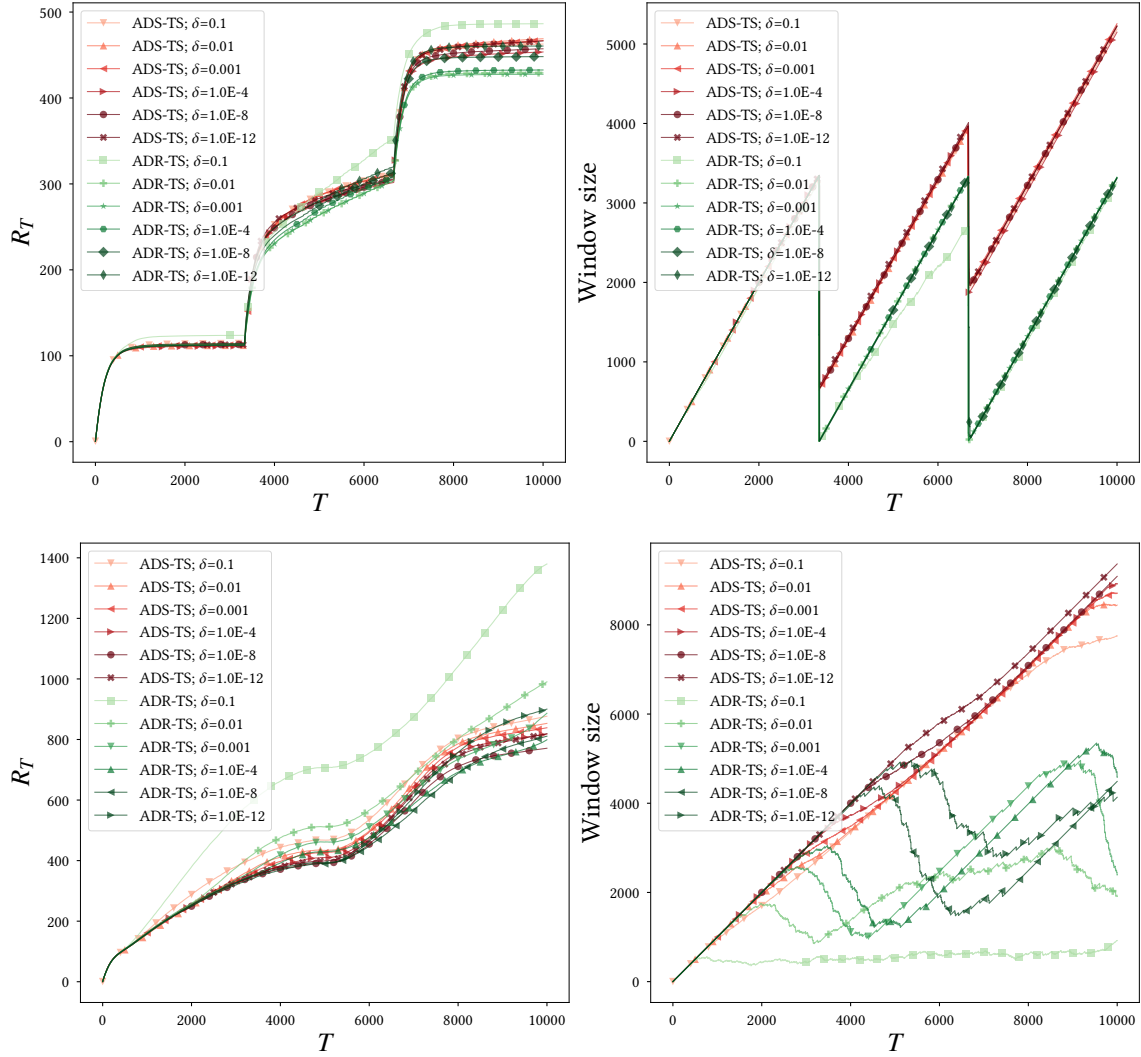


Figure 9: Comparison of ADR-TS and ADS-TS in synthetic environments. Left: Smaller regret (R_T : y-axis) indicates better performance. Right: Size of the window over time. The performance gap between ADR-TS and ADS-TS is very small. The effect of hyperparameter δ in change point detection is not very large, except for a very large or very small value of δ .

Appendix C. Lemmas

This section describes the lemmas that are used in the proofs of this paper.

The Hoeffding inequality, which is one of the most well-known concentration inequalities, provides a high-probability bound of the sum of bounded independent random variables.

Lemma 25 (Azuma-Hoeffding inequality) *Let x_1, x_2, \dots, x_n be martingale random variables in $[0, 1]$ with their conditional mean $\mu_m = \mathbb{E}[x_m | x_1, x_2, \dots, x_{m-1}]$. Let $\bar{x} = (1/n) \sum_{t=1}^n x_t$ and $\bar{\mu} = (1/n) \sum_{t=1}^n \mu_t$. Then,*

$$\Pr \left[\bar{x} - \bar{\mu} \geq \sqrt{\frac{\log(1/\delta)}{2n}} \right] \leq \delta, \quad (28)$$

$$\Pr \left[\bar{x} - \bar{\mu} \leq -\sqrt{\frac{\log(1/\delta)}{2n}} \right] \leq \delta. \quad (29)$$

Moreover, taking a union bound over the two inequalities yields

$$\Pr \left[|\bar{x} - \bar{\mu}| \geq \sqrt{\frac{\log(1/\delta)}{2n}} \right] \leq 2\delta. \quad (30)$$

Appendix D. Proofs of ADWIN

We denote $\mathcal{A}, \mathcal{B} = \mathcal{A} \cap \mathcal{B}$ for two events \mathcal{A}, \mathcal{B} .

D.1 Proof of Theorem 8

The overall idea here is as follows. With sufficiently long time after a changepoint, we can expect that ADWIN shrinks the window. However, there might be some time steps left in the current window $W(t)$ even if a shrink occurs. Still, we can show that

$$|\hat{\mu}_{W(t)} - \mu_t| \leq |\mu_{W(t)} - \mu_t| + |\hat{\mu}_{W(t)} - \mu_{W(t)}| \quad (31)$$

$$\leq \tilde{O} \left(\sqrt{\frac{1}{c(t)}} \right) + \tilde{O} \left(\sqrt{\frac{1}{c(t)}} \right) \quad (32)$$

where $c(t)$ is the number of time steps after the changepoint.¹⁵ Events \mathcal{C} and \mathcal{D} in the following corresponds to the bounds of $|\mu_{W(t)} - \mu_t|$ and $|\hat{\mu}_{W(t)} - \mu_{W(t)}|$, respectively.

Proof of Theorem 8

By Eq. (5), event

$$\mathcal{B} = \bigcap_{W' \in \mathcal{W}} \left\{ |\mu_{W'} - \hat{\mu}_{W'}| \leq \sqrt{\frac{\log(T^3)}{2|W'|}} \right\} \quad (33)$$

holds with probability at least $1 - 2/T$.

¹⁵. A formal definition of $c(t)$ is given in Eq. (34).

For each time step $t \in [T]$, let

$$c(t) = t - \max_{s < t, s \in \mathcal{T}_c} s. \quad (34)$$

Namely, $c(t)$ is the number of the time steps since the latest changepoint. Moreover, for $c(t) > 2$, let

$$\Delta(c) = 4\sqrt{\frac{\log(T^3)}{c-2}} \quad (35)$$

and $\Delta(c) = 1$ for $c \leq 2$. Let

$$\mathcal{C} = \bigcap_t \{|\mu_{W(t)} - \mu_t| < \Delta(c(t))\}. \quad (36)$$

In the following we show \mathcal{C} holds under \mathcal{B} . If $|W(t)| + 1 \leq c(t)$, then all the time steps in $W(t)$ are after the last change point, which implies $\mu_{W(t)} = \mu_t$ and thus \mathcal{C} . Otherwise, let

$$W_2 = \{t - c(t)/2 + 2, t - c(t)/2 + 3, \dots, t\} \quad (\text{i.e. the last } c(t)/2 - 1 \text{ time steps}) \quad (37)$$

$$W_1 = W(t-1) \setminus W_2 \quad (38)$$

Clearly $|W_1|, |W_2| \geq c(t)/2 - 1$ and $\mu_{W_2} = \mu_t$. ADWIN at time step $t-1$ shrinks the window until

$$|\hat{\mu}_{W_1}| - |\hat{\mu}_{W_2}| \leq \sqrt{\frac{\log(T^3)}{2|W_1|}} + \sqrt{\frac{\log(T^3)}{2|W_2|}} \quad (39)$$

holds. Let $n_1, n_2 = |W_1|, |W_2|$.

$$\sqrt{\frac{\log(T^3)}{2|W_1|}} + \sqrt{\frac{\log(T^3)}{2|W_2|}} \geq |\hat{\mu}_{W_1}| - |\hat{\mu}_{W_2}| \quad (\text{by Eq. (39)}) \quad (40)$$

$$\geq |\mu_{W_1}| - |\mu_{W_2}| - \sqrt{\frac{\log(T^3)}{2|W_1|}} - \sqrt{\frac{\log(T^3)}{2|W_2|}} \quad (\text{by } \mathcal{B}) \quad (41)$$

$$= |\mu_{W_1}| - |\mu_t| - \sqrt{\frac{\log(T^3)}{2|W_1|}} - \sqrt{\frac{\log(T^3)}{2|W_2|}} \quad (\text{by } |W_2| + 1 \leq c(t)) \quad (42)$$

and thus

$$|\mu_{W(t)} - \mu_t| \leq |\mu_{W_1} - \mu_t| \quad (43)$$

$$\leq 2\sqrt{\frac{\log(T^3)}{2|W_1|}} + 2\sqrt{\frac{\log(T^3)}{2|W_2|}} \quad (\text{by Eq. (42)}) \quad (44)$$

$$\leq 4\sqrt{\frac{\log(T^3)}{c-2}}. \quad (\text{by } |W_1|, |W_2| \geq (c-2)/2) \quad (45)$$

In summary, \mathcal{B} implies \mathcal{C} .

Note that under event \mathcal{B} , ADWIN never makes a false shrink (i.e., a shrink when $\mu_{W_1} = \mu_{W_2}$). This implies that between two changepoints, a shrink that makes $|W(t)| + 1 < c(t)$ occurs at most once, which leads to the event

$$\mathcal{D} = \bigcap_{n \in \mathbb{N}} \left\{ \sum_t \mathbf{1}\{|W(t)| = n \cap |W(t)| + 1 < c(t)\} \leq 2M \right\}. \quad (46)$$

By using this, the total error is bounded as

$$\mathbb{E}[\text{Err}(T)] \quad (47)$$

$$\leq T \Pr[\mathcal{B}^c] + \mathbb{E}[\text{Err}(T) \mathbf{1}[\mathcal{B}, \mathcal{C}, \mathcal{D}]] \quad (\mathcal{B} \text{ implies } \mathcal{C} \text{ and } \mathcal{D}) \quad (48)$$

$$\leq 2 + \mathbb{E}[\text{Err}(T) \mathbf{1}[\mathcal{B}, \mathcal{C}, \mathcal{D}]] \quad (49)$$

$$= 2 + \mathbb{E} \left[\sum_t |\mu_t - \hat{\mu}_W| \mathbf{1}[\mathcal{B}, \mathcal{C}, \mathcal{D}] \right] \quad (50)$$

$$\leq 2 + \mathbb{E} \left[\sum_t (|\mu_t - \mu_W| + |\hat{\mu}_W - \mu_W|) \mathbf{1}[\mathcal{B}, \mathcal{C}, \mathcal{D}] \right] \quad (51)$$

$$\leq 2 + \sum_t \Delta(c(t)) + \mathbb{E} \left[\sum_t |\hat{\mu}_W - \mu_W| \mathbf{1}[\mathcal{B}, \mathcal{C}, \mathcal{D}] \right] \quad (\text{by } \mathcal{C}) \quad (52)$$

$$\leq 2 + \sum_t \Delta(c(t)) + O(M) + \sum_t \sqrt{\frac{\log(T^3)}{2c(t)}} \quad (53)$$

$$+ \mathbb{E} \left[\sum_n \sum_t \mathbf{1}[|W(t)| = n, |W(t)| + 1 < c(t)] \sqrt{\frac{\log(T^3)}{2n}} \mathbf{1}[\mathcal{D}] \right] \quad (\text{by } \mathcal{B}) \quad (54)$$

$$\leq O(M) + \sum_t O \left(\sqrt{\frac{\log(T)}{c(t)}} \right) + \mathbb{E} \left[\sum_n \sum_t \mathbf{1}[|W(t)| = n, |W(t)| + 1 < c(t)] \sqrt{\frac{\log(T^3)}{2n}} \mathbf{1}[\mathcal{D}] \right]. \quad (55)$$

Here

$$\sqrt{\frac{\log(T)}{c(t)}} \leq \sqrt{T} \times \sqrt{\sum_t \left(\frac{\log(T)}{c(t)} \right)^2} \quad (\text{by Cauchy-Schwarz inequality}) \quad (56)$$

$$= \tilde{O} \left(\sqrt{MT} \right). \quad (c(t) = n \text{ holds at most } M + 1 \text{ times for each } n) \quad (57)$$

Another application of the Cauchy-Schwarz inequality, combined with \mathcal{D} , yields

$$\sum_n \sum_t \mathbf{1}[|W(t)| = n, |W(t)| \leq c(t)] \sqrt{\frac{\log(T^3)}{2n}} \mathbf{1}[\mathcal{D}] = \tilde{O} \left(\sqrt{MT} \right), \quad (58)$$

which completes the proof. \blacksquare

We next discuss the error bound under gradual drift. The following Lemmas 26 and 27 characterize the accuracy of estimator $\hat{\mu}_W$ under gradual drift. These lemmas are used in the proof of Lemma Theorem 10.

D.2 Lemma 26

Lemma 26 *Let the stream be gradual with the speed of the change b . Let the parameter of ADWIN be $\delta = 1/T^3$. Then, with a probability of at least $1 - 2/T$,*

$$|\mu_s - \mu_{s'}| \leq 3bN + \tilde{O}\left(\sqrt{\frac{1}{N}}\right) \quad (59)$$

for any two time steps s, s' in the current window $W = W(t)$, from which it easily follows that

$$|\mu_t - \mu_W| \leq 3bN + \tilde{O}\left(\sqrt{\frac{1}{N}}\right). \quad (60)$$

Lemma 26 is a strong characterization because $\mu_s - \mu_{s'}$ does not depend on the window size $|W|$: No matter how long the current window is, $\mu_s - \mu_{s'}$ and thus $\mu_t - \mu_W$ is bounded in terms of change speed b : In other words, if $\mu_t - \mu_W$ is larger, ADWIN shrinks the window. Note that the argument here is more general than the (informal) discussion in Bifet and Gavaldà (2007) for gradual drift. The discussion in Bifet and Gavaldà (2007) is limited to the case of linear drift, whereas our result holds for arbitrary drift as long as it moves at most b per time step.

Proof of Lemma 26 We first consider the case $|W| = CN$ for some integer $C \in \mathbb{N}^+$. We decompose W into C subwindows of equal size N and let W_c be the c -th subwindow for $c \in [C]$. For $c \in [C] \setminus \{1\}$, let $W_{:c}$ be the joint subwindow of W before W_c . Namely, $W_{:c} = W_1 \cup W_2 \cup \dots \cup W_{c-1}$. The fact that the window grows to size W without a detection implies that each split $W_{:c} \cup W_c$ satisfies

$$|\hat{\mu}_{W_{:c}} - \hat{\mu}_{W_c}| \leq 2\sqrt{\frac{\log(T^3)}{N}}. \quad (61)$$

Let $c \in [C]$ be arbitrary. By recursively applying Eq.(61) we have

$$|\hat{\mu}_W - \hat{\mu}_{W_c}| = \left| \frac{1}{C}\hat{\mu}_{W_c} + \frac{C-1}{C}\hat{\mu}_{W_{:c}} - \hat{\mu}_{W_c} \right| \quad (62)$$

$$\leq |\hat{\mu}_{W_{:c}} - \hat{\mu}_{W_c}| + \frac{2}{C}\sqrt{\frac{\log(T^3)}{N}} \quad (\text{by Eq.(61)}) \quad (63)$$

$$\leq |\hat{\mu}_{W_{:c-1}} - \hat{\mu}_{W_c}| + \left(\frac{2}{C-1} + \frac{2}{C}\right)\sqrt{\frac{\log(T^3)}{N}} \quad (64)$$

$$\dots \quad (65)$$

$$\leq |\hat{\mu}_{W_{:c}} - \hat{\mu}_{W_c}| + \sum_{c'=c+1}^C \frac{2}{c'}\sqrt{\frac{\log(T^3)}{N}} \quad (66)$$

$$\leq \left(1 + \sum_{c'=c+1}^C \frac{2}{c'}\right)\sqrt{\frac{\log(T^3)}{N}} \quad (\text{by Eq.(61)}) \quad (67)$$

$$\leq (\log T)\sqrt{\frac{\log(T^3)}{N}}, \quad (68)$$

which implies for any $c, c' \in [C]$ we have

$$|\hat{\mu}_{W_c} - \hat{\mu}_{W_{c'}}| \leq 2(\log T) \sqrt{\frac{\log(T^3)}{N}}. \quad (69)$$

By Eq. (5) we have

$$\begin{aligned} |\mu_{W_c} - \hat{\mu}_{W_c}| &\leq \sqrt{\frac{\log(T^3)}{N}} \\ |\mu_{W_{c'}} - \hat{\mu}_{W_{c'}}| &\leq \sqrt{\frac{\log(T^3)}{N}}. \end{aligned} \quad (70)$$

By the fact that μ_t moves bN within a subwindow of size w , for any $s \in W_c, s' \in W_{c'}$, we have

$$\begin{aligned} |\mu_{W_c} - \mu_s| &\leq bN \\ |\mu_{W_{c'}} - \mu_{s'}| &\leq bN. \end{aligned} \quad (71)$$

By using these, we have

$$|\mu_s - \mu_{s'}| \leq |\mu_{W_c} - \mu_s| + |\mu_{W_{c'}} - \mu_{s'}| + |\mu_{W_c} - \hat{\mu}_{W_c}| + |\mu_{W_{c'}} - \hat{\mu}_{W_{c'}}| + |\hat{\mu}_{W_c} - \hat{\mu}_{W_{c'}}| \quad (72)$$

$$\leq 2bN + 2(1 + \log T) \sqrt{\frac{\log(T^3)}{N}}. \quad (\text{by Eq. (69), (70), (71)}) \quad (73)$$

The general case of $|W| = CN + n$ for $n \in \{0, 1, \dots, N - 1\}$ is easily proven by replacing $2bN$ with $3bN$ since μ_t drift at most bN in n time steps. \blacksquare

D.3 Lemma 27

Lemma 26 characterizes the accuracy of the estimator. However, Lemma 26 only holds when window size $|W| \geq N$. When ADWIN shrinks the window very frequently, we cannot guarantee the quality of the estimator $\hat{\mu}_W$. The following Lemma 27 states that this is not the case: With high probability, the current window grows until $|W| = O(b^{-2/3})$.

Lemma 27 (Bound on erroneous shrinking) *Let $C_1 = \tilde{O}(1)$ be a sufficiently large value that is later defined in Eq. (86). Let the parameter of ADWIN be $\delta = 1/T^3$. Let the drift speed b be such that $C_1 b^{-2/3} \leq T$. Let*

$$\mathcal{P}(t) = \bigcup_{W_1, W_2: W(t) = W_1 \cup W_2} \left\{ |W_1| \leq C_1 b^{-2/3}, |W_2| \leq C_1 b^{-2/3}, |\hat{\mu}_{W_1} - \hat{\mu}_{W_2}| \geq \epsilon_{\text{cut}}^\delta \right\}. \quad (74)$$

Let

$$\mathcal{P} = \bigcup_{t \in [T]} \mathcal{P}(t). \quad (75)$$

Then,

$$\Pr[\mathcal{P}] \leq 2C_1 T^{-1}. \quad (76)$$

Proof of Lemma 27 Let $C_1 = \tilde{O}(1)$ that we define later in Eq. (86). Let

$$\mathcal{W}_{C_1} = \{W' \in \mathcal{W} : |W'| \leq C_1 b^{-2/3}\} \quad (77)$$

be the subset of the window set \mathcal{W} with their sizes at most $C_1 b^{-2/3}$. It is easy to show that $|\mathcal{W}_{C_1}| \leq TC_1 b^{-2/3}$. Let d be such that $T^{-d} = b$. Similarly to Eq. (5), by the union bound of the Hoeffding bound over all windows of \mathcal{W}_{C_1} , with probability at least

$$1 - \frac{2}{T^{2+d}} \times TC_1 b^{-2/3} = 1 - 2C_1 T^{-1} b^{1/3} \geq 1 - 2C_1 T^{-1}, \quad (78)$$

we have

$$\begin{aligned} |\mu_{W_1} - \hat{\mu}_{W_1}| &\leq \sqrt{\frac{\log(T^{2+d})}{2|W_1|}} \\ |\mu_{W_2} - \hat{\mu}_{W_2}| &\leq \sqrt{\frac{\log(T^{2+d})}{2|W_2|}} \end{aligned} \quad (79)$$

holds for all t and any split $W_1 \cup W_2 = W(t) : |W_1|, |W_2| \leq C_1 b^{-2/3}$. Let $N = C_1 b^{-2/3}$. By definition of gradually changing stream,

$$|\mu_{W_1} - \mu_{W_2}| \leq 2bN. \quad (80)$$

In the following, we show that $|W_1| \leq C_1 b^{-2/3}$, $|W_2| \leq C_1 b^{-2/3}$, $|\hat{\mu}_{W_1} - \hat{\mu}_{W_2}| \geq \epsilon_{\text{cut}}^\delta$ never occur under Eq. (79) by using a proof-by-contradiction argument. Namely,

$$\sqrt{\frac{\log(T^3)}{2|W_1|}} + \sqrt{\frac{\log(T^3)}{2|W_2|}} \leq |\hat{\mu}_{W_1} - \hat{\mu}_{W_2}| \quad (\text{when } |\hat{\mu}_{W_1} - \hat{\mu}_{W_2}| \geq \epsilon_{\text{cut}}^\delta) \quad (81)$$

$$\leq |\mu_{W_1} - \mu_{W_2}| + |\hat{\mu}_{W_1} - \mu_{W_1}| + |\hat{\mu}_{W_2} - \mu_{W_2}| \quad (\text{by triangular inequality}) \quad (82)$$

$$\leq 2bN + |\hat{\mu}_{W_1} - \mu_{W_1}| + |\hat{\mu}_{W_2} - \mu_{W_2}| \quad (\text{by (80)}) \quad (83)$$

$$\leq 2bN + \sqrt{\frac{\log(T^{2+d})}{2|W_1|}} + \sqrt{\frac{\log(T^{2+d})}{2|W_2|}} \quad (\text{by (79)}) \quad (84)$$

which implies

$$\frac{\sqrt{3} - \sqrt{2+d}}{\sqrt{2}} \left(\sqrt{\frac{\log T}{|W_1|}} + \sqrt{\frac{\log T}{|W_2|}} \right) \leq 2bN, \quad (85)$$

which does not hold for

$$N = b^{-2/3} \underbrace{\left(\frac{\sqrt{3} - \sqrt{2+d}}{2\sqrt{2}} \sqrt{\log T} \right)^{2/3}}_{=: C_1}, \quad |W_1|, |W_2| \leq N. \quad (86)$$

In summary, with probability $1 - 2C_1 T^{-1}$, we have \mathcal{P}^c . ■

D.4 Proof of Theorem 10

Let d be such that $b = T^{-d}$. Let C_1 be the value defined in Eq. (86) and $N_1 = C_1 b^{-2/3}$.

By using Lemma 27, we bound the number of shrinks.

Lemma 28 (Number of shrinks) *Let*

$$\mathcal{T}_d^{\text{sm1}} := \{t \in \mathcal{T}_d, |W_2(t)| < N_1\}. \quad (87)$$

Then, under \mathcal{P}^c , $|\mathcal{T}_d^{\text{sm1}}| \leq T/N_1$ holds.

Proof of Lemma 28 \mathcal{P}^c implies that for any shrink $|W_1| \geq N_1$ or $|W_2| \geq N_1$ holds. A shrink of the latter case is not included in $\mathcal{T}_d^{\text{sm1}}$. A shrink of the former case reduces the size of the window at least N_1 and cannot occur more than T/N_1 times. \blacksquare

Proof of Theorem 10 Let $c^{\text{br}}(t) = t - \max_{s < t, s \in \mathcal{T}_d} s$ be the number of time steps after the last detection time. We have

$$\text{Err}(T) \leq T \mathbf{1}[\mathcal{P}] + \sum_{n=1}^{N_1} \sum_{t=1}^T |\mu_t - \hat{\mu}_W| \mathbf{1}[c^{\text{br}}(t) = n, \mathcal{P}^c] + \sum_{t=1}^T |\mu_t - \hat{\mu}_W| \mathbf{1}[c^{\text{br}}(t) \geq N_1, \mathcal{P}^c]. \quad (88)$$

By Lemma 27, the expectation of the first term of Eq. (88) is bounded as

$$\mathcal{P}[T \mathbf{1}[\mathcal{P}]] = T \times 2C_1 T^{-1} = \tilde{O}(1). \quad (89)$$

We next bound the second term of Eq. (88). By Eq. (5),

$$\forall_{t \in [T]} \left\{ |\mu_t - \hat{\mu}_W| \leq \sqrt{\frac{\log(T^3)}{2(c^{\text{br}}(t) - 1)}} + b c^{\text{br}}(t) \right\} \quad (90)$$

holds with probability $1 - 2/T$. The expectation of the second term is bounded as:

$$\sum_{n=1}^{N_1} \sum_{t=1}^T \Pr[|\mu_t - \hat{\mu}_W| \mathbf{1}[c(t) = n, \mathcal{X}^c]] \quad (91)$$

$$\leq (|\mathcal{T}_d^{\text{sm1}}| + 1) \sum_{n=1}^{N_1} \left(\sqrt{\frac{\log(T^3)}{2n}} + bn \right) + T \times \frac{2}{T} \quad (92)$$

(by Eq. (90) and the fact that $c^{\text{br}}(t) = n$ for each n occurs at most once between two detection times)

$$\leq \left(\frac{T}{N_1} + 1 \right) \sum_{n=1}^{N_1} \left(\sqrt{\frac{\log(T^3)}{2n}} + bn \right) + 2 \quad (93)$$

$$\text{(by Lemma 28)} \quad (94)$$

$$\leq \tilde{O} \left(\frac{T}{N_1} \times (\sqrt{N_1} + bN_1^2) \right) = O(T^{1-d/3}) \quad (95)$$

where we used $b, N_1 = T^{-d}, \tilde{O}(T^{(2d)/3})$ in the last transformation. Moreover, by using Lemma 26, the expectation of the third term is bounded as

$$\sum_{t=1}^T \mathbb{E}[|\mu_t - \hat{\mu}_W| \mathbf{1}[c(t) \geq N_1]] \leq T \times \left(3bN_1 + \tilde{O}\left(\sqrt{\frac{1}{N_1}}\right) \right) = \tilde{O}(T^{1-d/3}), \quad (96)$$

which completes the proof. \blacksquare

Appendix E. Proofs of MAB

We first clarify the notation: In bandit streams, let

$$W^i = \{s \in W : i = I(s)\}. \quad (97)$$

Let

$$\mu_{i,W} = \sum_{s \in W^i} \mu_{i,s} \quad (98)$$

and its empirical estimate be

$$\hat{\mu}_{i,W} = \sum_{s \in W^i} x_{i,s}. \quad (99)$$

In the bandit setting, only one arm $I(t) \in [K]$ is observable at round t , and thus $W^i \subset W$. Still, the following inequality holds with probability $1 - 2/T^p$:

$$\forall W' \in \mathcal{W} \quad |\mu_{i,W'} - \hat{\mu}_{i,W'}| \leq \sqrt{\frac{\log(T^{2+p})}{2|(W')^i}}. \quad (100)$$

Eq. (100), which is the same form as Eq. (5), holds because it corrects T^2 multiplications. The union bound of Eq. (100) over all arms holds with $1 - K/T^p$.

E.1 Proof of Theorem 18

Proof of Theorem 18 Under stationary environment, Eq. (100) with $p = 1$ implies that ADR-bandit never reset itself with high probability. The regret is bounded by

$$\text{Reg}(T) \leq \underbrace{T \times \frac{2K}{T}}_{\text{regret of false reset}} + \underbrace{\sum_i C^{\text{st}} \frac{\log T}{\Delta_i}}_{\text{regret of ADR-bandit, Lemma 17}} \quad (101)$$

$$\leq C^{\text{st}} \sum_i \frac{\log T}{\Delta_i} + O(1). \quad (102)$$

\blacksquare

E.2 Analysis of Thompson sampling under drift

This section derives a proof of Lemma 14. Without loss of generality, we define $\mu_{1,1} \geq \mu_{2,1} \geq \dots \geq \mu_{K,1}$ at round 1. For ease of notation, let $\mu_i := \mu_{i,1}$ and $\epsilon = \epsilon(T)$.

In the following, we derive the drift-tolerant regret of TS.

Lemma 29 (Number of suboptimal draws) *Let $c > 0$ be arbitrary. Assume that we run TS. The following inequality holds*

$$\sum_{t \in [T]} \Pr[\theta_{I(t),t} \leq \mu_{1,1} - \epsilon(t) - c] = O\left(\frac{1}{c^2}\right). \quad (103)$$

Proof of Lemma 29

We have,

$$\begin{aligned} \sum_t \mathbf{1}[\theta_{I(t),t} \leq \mu_{1,1} - \epsilon(t) - c] &= \sum_{n=0}^T \sum_{t=1}^T \mathbf{1}[\theta_{I(t),t} \leq \mu_{1,1} - \epsilon(t) - c, N_1(t) = n] \quad (104) \\ &= \sum_{n=0}^T \sum_{m=1}^T \mathbf{1} \left[\sum_t \mathbf{1}[\theta_{I(t),t} \leq \mu_{1,1} - \epsilon(t) - c, N_1(t) = n] \geq m \right]. \quad (105) \end{aligned}$$

Here, if

$$\left\{ \theta_1(t) > \mu_{1,1} - \epsilon(t) - c, \max_{i \neq 1} \theta_{i,t} \leq \mu_{1,1} - \epsilon(t) - c, N_1(t) = n \right\} \quad (106)$$

then arm 1 is drawn. This fact implies that event $\{\sum_{t=1}^T \mathbf{1}[\theta_{I(t),t} \leq \mu_{1,1} - \epsilon(t) - c, N_1(t) = n] \geq m\}$ requires that

$$\theta_1(t) \leq \mu_{1,1} - \epsilon(t) - c$$

holds in the first m rounds of the subsequence $\tau_1, \tau_2, \tau_3, \dots, \tau_m = \{t : \max_{i \neq 1} \theta_{i,t} \leq \mu_{1,1} - \epsilon(t) - c, N_1(t) = n\}$. Letting $S_\beta(x; \mu, n)$ be the survival function¹⁶ of Beta($1 + \mu n, 1 + (1 - \mu)n$), we have

$$\sum_{n=0}^T \mathbb{E} \left[\sum_{m=1}^T \sum_{t=1}^T \mathbf{1}[\theta_{I(t),t} \leq \mu_{1,1} - \epsilon(t) - c, N_1(t) = n] \geq m \right] \quad (107)$$

$$\leq \sum_{n=0}^T \mathbb{E} \left[\sum_{m=1}^T \prod_{k=1}^m \Pr[\theta_{1,\tau_k} \leq \mu_{1,1} - \epsilon(\tau_k) - c | \hat{\mu}_1^{(n)}] \right] \quad (108)$$

$$= \sum_{n=0}^T \mathbb{E} \left[\sum_{m=1}^T \left(S_\beta(\mu_{1,1} - \epsilon(\tau_1) - c; \hat{\mu}_1^{(n)}, n) \right)^m \right] \quad (\epsilon(\tau_k) \text{ is nondecreasing in } k) \quad (109)$$

$$\leq \sum_{n=0}^T \mathbb{E} \left[\frac{1 - S_\beta(\mu_{1,1} - \epsilon(\tau_1) - c; \hat{\mu}_1^{(n)}, n)}{S_\beta(\mu_{1,1} - \epsilon(\tau_1) - c; \hat{\mu}_1^{(n)}, n)} \right] \quad (110)$$

$$= O(1/c^2). \quad (\text{by Lemma 31}) \quad (111)$$

16. The probability $\theta \sim \text{Beta}(1 + \mu n, 1 + (1 - \mu)n)$ exceeds x .

where, to apply Lemma 31, we used the fact that $\hat{\mu}_1^{(n)}$ is the mean of n independent samples from Bernoulli distributions with means $\mu_{1,1} - \epsilon(\tau_1)$ or more. \blacksquare

E.2.1 PROOF OF LEMMA 14

Let $(x)_+ = \max(x, 0)$. We have,

$$\text{Reg}_{\text{tr}}(T, 3) \tag{112}$$

$$= \sum_t (\text{reg}(t) - 3\epsilon(t))_+ \tag{113}$$

$$= \sum_t \sum_{i \in [K]} (\Delta_i - 3\epsilon(t))_+ \mathbf{1}[I(t) = i] \tag{114}$$

$$\leq \sum_t \sum_{i \in [K]} (\Delta_i - 3\epsilon(t))_+ \left(\mathbf{1} \left[I(t) = i, \theta_{I(t),t} < \frac{\mu_1 + \mu_i}{2} \right] + \mathbf{1} \left[I(t) = i, \theta_{I(t),t} \geq \frac{\mu_1 + \mu_i}{2} \right] \right) \tag{115}$$

$$= \sum_t \sum_{i \in [K]} (\Delta_i - 3\epsilon(t))_+ \left(\mathbf{1} \left[I(t) = i, \theta_{I(t),t} < \mu_1 - \frac{\Delta_i}{2} \right] + \mathbf{1} \left[I(t) = i, \theta_{I(t),t} \geq \mu_i + \frac{\Delta_i}{2} \right] \right) \tag{116}$$

$$\leq \sum_t \sum_{i \in [K]} (\Delta_i - 3\epsilon(t))_+ \left(\mathbf{1} \left[I(t) = i, \theta_{I(t),t} < \mu_1 - \epsilon(t) - \frac{\Delta_i}{6} \right] + \mathbf{1} \left[I(t) = i, \theta_{I(t),t} \geq \mu_i + \epsilon(t) + \frac{\Delta_i}{6} \right] \right). \tag{117}$$

(by $\Delta_i \geq 3\epsilon(t)$ or $(\Delta_i - 3\epsilon(t))_+ = 0$)

The first term is bounded in expectation as

$$\sum_t \sum_{i \in [K]} \Delta_i \mathbb{E} \left[\mathbf{1} \left[I(t) = i, \theta_{I(t),t} \leq \mu_1 - \epsilon(t) - \frac{\Delta_i}{6} \right] \right] = \sum_{i \in [K]} \Delta_i \times O \left(\frac{1}{\Delta_i^2} \right) = \sum_{i \in [K]} O \left(\frac{1}{\Delta_i} \right),$$

where we used Lemma 29 in the first transformation with $c = \Delta_i/3$.

We next bound the second term.

$$\sum_t \sum_{i \in [K]} \Delta_i \mathbf{1} \left[I(t) = i, \theta_{I(t),t} > \mu_1 + \epsilon(t) + \frac{\Delta_i}{6} \right] \tag{118}$$

$$\leq \sum_{i \in [K]} \left(\frac{144 \log T}{\Delta_i} + \sum_t \Delta_i \mathbf{1} \left[\theta_{i,t} \geq \mu_i + \epsilon(t) + \frac{\Delta_i}{6}, N_i(t) \geq \frac{144 \log T}{\Delta_i^2} \right] \right). \tag{119}$$

Here,

$$\sum_t \Delta_i \mathbf{1} \left[\theta_{i,t} \geq \mu_i + \epsilon(t) + \frac{\Delta_i}{6}, N_i(t) \geq \frac{144 \log T}{\Delta_i^2} \right] \tag{120}$$

$$\leq \sum_t \Delta_i \mathbf{1} \left[\hat{\mu}_{i,t} \geq \mu_i + \epsilon(t) + \frac{\Delta_i}{12}, N_i(t) \geq \frac{144 \log T}{\Delta_i^2} \right]$$

$$+ \sum_t \Delta_i \mathbf{1} \left[\theta_{i,t} \geq \mu_i + \epsilon(t) + \frac{\Delta_i}{6}, \hat{\mu}_{i,t} \leq \mu_i + \epsilon(t) + \frac{\Delta_i}{12}, N_i(t) \geq \frac{144 \log T}{\Delta_i^2} \right]. \quad (121)$$

By using the Hoeffding inequality, the first term of Eq. (121) is bounded as

$$\sum_t \Delta_i \mathbf{1} \left[\hat{\mu}_{i,t} \geq \mu_i + \epsilon(t) + \frac{\Delta_i}{12}, N_i(t) \geq \frac{144 \log T}{\Delta_i^2} \right] \quad (122)$$

$$\leq \sum_N T \exp \left(-2 \times \frac{144 \log T}{\Delta_i^2} \times \left(\frac{\Delta_i}{12} \right)^2 \right) \quad (123)$$

$$\leq T^2 \exp(-2 \log T) = 1, \quad (124)$$

where we have used the Hoeffding inequality and the fact that $\hat{\mu}_{i,t}$ is the mean of $N_i(t)$ samples with each mean no more than $\mu_i + \epsilon(t)$. By using Lemma 32, the second term of Eq. (121) is bounded as

$$\begin{aligned} \sum_t \Delta_i \mathbb{E} \left[\theta_{i,t} \geq \mu_i + \epsilon(t) + \frac{\Delta_i}{6}, \hat{\mu}_{i,t} \leq \mu_i + \epsilon(t) + \frac{\Delta_i}{12}, N_i(t) \geq \frac{144 \log T}{\Delta_i^2} \right] \\ \leq \sum_{n,t} \exp \left(-2 \times \frac{144 \log T}{\Delta_i^2} \times \left(\frac{\Delta_i}{12} \right)^2 \right) \leq 1. \end{aligned} \quad (125)$$

E.2.2 LEMMAS FOR THOMPSON SAMPLING

The following lemma is a version of Lemma 4 in Agrawal and Goyal (2013).

Lemma 30 *Let $\hat{\mu}_1^{(n)}$ is the mean of n independent samples that are drawn from Bernoulli distributions with their means no less than $\mu_1 - \epsilon(t)$. Then,*

$$\mathbb{E} \left[\frac{1}{S_\beta(\mu_1 - \epsilon(t) - c; \hat{\mu}_1^{(n)}, n)} \right] \leq \begin{cases} 1 + \frac{3}{c} & n < \frac{8}{c} \\ 1 + \Theta \left(e^{-c^2 n/2} + \frac{1}{(n+1)c^2} e^{-D_i n} + \frac{1}{e^{c^2 n/4} - 1} \right) & n \geq \frac{8}{c} \end{cases}, \quad (126)$$

where $D_i = 2c^2$.

Lemma 31 *Let $\hat{\mu}_1^{(n)}$ is the mean of n independent samples that are drawn from Bernoulli distributions with their means no less than $\mu_1 - \epsilon(t)$. Then, the following inequality holds:*

$$\sum_{n=0}^T \mathbb{E} \left[\frac{1 - S_\beta(\mu_1 - \epsilon(t) - c; \hat{\mu}_1^{(n)}, n)}{S_\beta(\mu_1 - \epsilon(t) - c; \hat{\mu}_1^{(n)}, n)} \right] \quad (127)$$

$$\leq \frac{24}{c^2} + \sum_{n=0}^T \Theta \left(e^{-c^2 n/2} + \frac{1}{(n+1)c^2} e^{-D_i n} + \frac{1}{e^{c^2 n/4} - 1} \right) \quad (128)$$

$$= O \left(\frac{1}{c^2} \right). \quad (129)$$

Proof of Lemma 31 The proof of Lemma 31 is straightforward from Lemma 30 by following similar steps to the problem-independent bound of Theorem 2 in Agrawal and Goyal

(2013). ■

The following lemma is well-known in the Beta posterior.

Lemma 32

$$\Pr[\theta_{i,t} > \hat{\mu}_{i,t} + c \mid \hat{\mu}_{i,t}, N_i(t) = N] \leq \exp(-2Nc^2) \quad (130)$$

for any $c > 0$.

Proof of Lemma 32

$$\begin{aligned} & \Pr[\theta_{i,t} > \hat{\mu}_{i,t} + c \mid \hat{\mu}_{i,t}, N_i(t) = N] \\ & \leq e^{-Nd_{\text{KL}}(\hat{\mu}_{i,t}, \hat{\mu}_{i,t} + c)} \quad (\text{Lemma 3 in Agrawal and Goyal (2013)}) \\ & \leq e^{-Nc^2}. \quad (\text{by Pinsker's inequality}) \end{aligned} \quad (131)$$
■

E.3 Regret due to Monitoring

Proof of Lemma 16 We call each consecutive KN rounds a “subblock”. For each block $l = 1, 2, \dots$, the number of subblocks is at most 2^{l-1} . The l -th monitoring arm $i^{(l)}$ is drawn during $2^l + 1$ subblocks, and at each subblock it is drawn for N rounds (Figure 4).

The goal of this theorem is bound the ratio between the regret during monitoring rounds divided by the regret during base-bandit rounds. Let

$$\mathcal{T}_l^{\text{monitor}} = \{t : \text{arm } i^{(l)} \text{ is drawn as the } l\text{-th monitoring arm}\}, \quad (132)$$

as illustrated in Figure 4. By definition,

$$\sum_t \text{reg}(t) = \sum_{t \in \mathcal{T}^{\text{base}}} \text{reg}(t) + \sum_l \sum_{t \in \mathcal{T}_l^{\text{monitor}}} \text{reg}(t) = \sum_l \left(\sum_{t \in \mathcal{T}^{\text{base}} \cap \mathcal{T}_l} \text{reg}(t) + \sum_{t \in \mathcal{T}_l^{\text{monitor}}} \text{reg}(t) \right), \quad (133)$$

where \mathcal{T}_l is the set of rounds in block l . First, for each monitoring arm $i^{(l)}$, we bound the ratio between the number of draws of the monitoring arms divided by the number of draws of base-bandit arms. Namely,

$$\frac{\sum_t \mathbf{1}[t \in \mathcal{T}_l^{\text{monitor}}]}{\sum_t \mathbf{1}[t \in \mathcal{T}^{\text{base}} \cap \mathcal{T}_l] \mathbf{1}[I(t) = i^{(l)}]}, \quad (134)$$

for each l . First, we consider the ratio of the regret due to the first monitoring arm $i^{(1)}$. The arm $i^{(1)}$ is the most drawn arm in the first subblock, which is drawn at least $KN/K = N$ rounds at $l = 1$ by the base-bandit algorithm. We draw $i^{(1)}$ for $2N$ times during block $l = 2$ as a monitoring arm. Therefore, Eq. (134) for $l = 1$ is at most two.

In the following, we consider the Eq. (134) for $l \geq 2$. Since the arm most frequently pulled in the first $(2^l - 2)N$ subblocks is chosen as $i^{(l)}$, it has been drawn at least $(2^l - 2)N$

times by the base-bandit algorithm before the block l . We draw this arm for $(2^l + 1)N$ times as a monitoring arm. Therefore, the ratio is at most

$$\frac{(2^l + 1)N}{(2^l - 2)N} \leq 3. \quad (135)$$

In summary, Eq. (134) is at most 3. In the following, we bound the total regret due to the regret during rounds $\mathcal{T}^{\text{base}}$. Namely,

$$\sum_{t \leq S} \text{reg}(t) = \sum_l \left(\sum_{t \in \mathcal{T}^{\text{base}} \cap \mathcal{T}_l} \text{reg}(t) + \sum_{t \in \mathcal{T}_l^{\text{monitor}}} \text{reg}(t) \right) \quad (\text{by Eq. (133)}) \quad (136)$$

$$= \sum_l \left(\sum_{t \in \mathcal{T}^{\text{base}} \cap \mathcal{T}_l} \Delta_{I(t)} + \sum_{t \in \mathcal{T}_l^{\text{monitor}}} \Delta_{i^{(l)}} \right) + \sum_t \epsilon(t) \quad (137)$$

$$\leq \sum_l \left(\sum_{t \in \mathcal{T}^{\text{base}} \cap \mathcal{T}_l} \Delta_{I(t)} + 3 \sum_{t \in \mathcal{T}^{\text{base}} \cap \mathcal{T}_l} \mathbf{1}[I(t) = i^{(l)}] \Delta_{i^{(l)}} \right) + \sum_t \epsilon(t) \quad (\text{by Eq. (134)}) \quad (138)$$

$$\leq 4 \sum_l \sum_{t \in \mathcal{T}^{\text{base}} \cap \mathcal{T}_l} \Delta_{I(t)} + \sum_t \epsilon(t) \quad (139)$$

$$= 4 \sum_{t \in \mathcal{T}^{\text{base}}} \Delta_{I(t)} + \sum_t \epsilon(t) \quad (140)$$

$$\leq 4 \sum_{t \in \mathcal{T}^{\text{base}}} \text{reg}(t) + 5 \sum_t \epsilon(t). \quad (141)$$

Let N_l^{base} be the number of draws of arm $i^{(l)}$ at block l . By the fact that the regret due to draws of arm $i^{(l)}$ at round t is at most $\Delta_{i^{(l)}} + \epsilon(t)$ and one arm is drawn for each round, the regret due to monitoring rounds is at most

$$3 \sum_{l \geq 1} N_l^{\text{base}} \Delta_{i^{(l)}} + \sum_t \epsilon(t) \leq 3 \sum_{t \in \mathcal{T}^{\text{base}}} \text{reg}(t) + \sum_t 2\epsilon(t), \quad (142)$$

by using the fact that the regret during the base-bandit rounds is at least $N_l^{\text{base}} \Delta_{i^{(l)}} - \sum_t \epsilon(t)$. The total regret, which is the sum of that of base-bandit rounds and monitoring rounds, is bounded as

$$\sum_{t=1}^S \text{reg}(t) \leq 4 \left(\sum_{t \in \mathcal{T}^{\text{base}}} \text{reg}(t) \right) + \sum_{t=1}^S 3\epsilon(t) \quad (143)$$

which concludes the proof. \blacksquare

E.4 Regret in abrupt environment: proof of Theorem 21

Proof of Theorem 21 Similar to the case of single stream, we define detection times, which is the subset of rounds where the reset occurs. Let $\mathcal{T}_d = \{t \in [T] : |W(t+1)| = 0\}$

be the set of detection times, and $M_d = |\mathcal{T}_d|$ be the number of detection times¹⁷. Let $T_{d,m}$ be the m -th element of \mathcal{T}_d . For convenience, let $(T_{d,0}, T_{d,M_d+1}) = (0, T)$. We denote $\mathcal{S}_m = \{T_{d,m} + 1, T_{d,m} + 2, \dots, T_{d,m+1}\}$ for $m \in \{0, \dots, M_d\}$, which is the interval between m -th and $(m + 1)$ -th detection times. We write $S_m = |\mathcal{S}_m|$.

We define an event

$$\mathcal{V} = \{\forall m \in [M] \exists s \in \mathcal{T}_d 0 \leq s - T_m \leq 16KU(\epsilon_m)\} \cap \{M_d = M\}. \quad (144)$$

Event \mathcal{V} states that for each changepoint $T_m \in \mathcal{T}_c$ (indexed by $m = 1, 2, \dots, M = |\mathcal{T}_c|$), there exists a corresponding detection time $T_{d,m}$ within $16KU(\epsilon_m)$ time steps. We first show that \mathcal{V} holds with high probability by induction. Assume that for every changepoint m' up to $1, 2, \dots, m - 1$, there exists unique detection time $s_{m'}$ such that $0 \leq s_{m'} - T_{m'} \leq 16KU(\epsilon_m)$. We show that $0 \leq s_m - T_m \leq 16KU(\epsilon_m)$. During $t \leq T_m$ (these time steps are between m -th and $(m - 1)$ -th changepoint), $\mu_{i,t}$ stays the same. We denote $\mu_i = \mu_{i,t}$ during these time steps. Eq. (100) implies that, with probability $1 - 2K/T$, for any split $W(t) = W_1 \cup W_2$,

$$\forall_i |\mu_i - \hat{\mu}_{i,W_1}| \leq \sqrt{\frac{\log(T^3)}{2|W_{i,1}|}} \quad (145)$$

$$\forall_i |\mu_i - \hat{\mu}_{i,W_2}| \leq \sqrt{\frac{\log(T^3)}{2|W_{i,2}|}} \quad (146)$$

where $W_{i,1} = \{t \in W_1 : i \in I(t)\}$ and $W_{i,2} = \{t \in W_2 : i \in I(t)\}$. This implies ADR-bandit with $\delta = T^3$ never makes a split before the m -th changepoint (i.e., $0 \leq s_m - T_m$). Let $s = T_m + 16KU(\epsilon_m)$. Assume that there is no detection between time step $T_{m-1} + 16KU(\epsilon_m)$ and $T_m + 16KU(\epsilon_m)$. Then for a split $W(s) = W_1 \cup W_2$, $W_1 = W \cap [T_m]$, $W_2 = W \setminus W_1$, we have

$$|W_1| \geq (T_m - T_{m-1} - 16KU(\epsilon_m)), \quad |W_2| \geq 16KU(\epsilon_m), \quad (147)$$

where we have used $s_{m-1} - T_{m-1} \leq 16KU(\epsilon_m)$ in the first inequality. By assumption of Theorem 21, the first inequality implies

$$|W_1| \geq (T_m - T_{m-1} - 16KU(\epsilon_m)) \geq 3KN - 16KU(\epsilon_m) \geq 2KN. \quad (148)$$

By the first property the monitoring consistency (Definition 15), there exists an arm $i \in [K]$ (i.e., monitoring arm $i^{(l)}$) such that

$$|W_{i,1}|, |W_{i,2}| \geq 16U(\epsilon_m). \quad (149)$$

Again by Eq. (100) we have

$$\forall_i |\mu_{i,T_m} - \hat{\mu}_{i,W_1}| \leq \sqrt{\frac{\log(T^3)}{2|W_{i,1}|}} \quad (150)$$

17. Remember the difference between changepoints \mathcal{T} and detection times \mathcal{T}_d . The former is defined on an abrupt environment, whereas the latter is defined for the ADR-bandit algorithm (and thus the latter is a random variable).

$$\forall_i |\mu_{i,T_m+1} - \hat{\mu}_{i,W_2}| \leq \sqrt{\frac{\log(T^3)}{2|W_{i,2}|}}. \quad (151)$$

By definition of the global changepoint,

$$|\mu_{i,T_m} - \mu_{i,T_m+1}| \geq \epsilon_m = \sqrt{\frac{\log(T^3)}{2U(\epsilon_m)}}. \quad (152)$$

Combining these yields

$$|\hat{\mu}_{i,W_1} - \hat{\mu}_{i,W_2}| \geq |\mu_{i,T_m} - \mu_{i,T_m+1}| - |\mu_{i,T_m} - \hat{\mu}_{i,W_1}| - |\mu_{i,T_m+1} - \hat{\mu}_{i,W_2}| \quad (\text{by triangular inequality}) \quad (153)$$

$$\geq \sqrt{\frac{\log(T^3)}{2U(\epsilon_m)}} - |\mu_{i,T_m} - \hat{\mu}_{i,W_1}| - |\mu_{i,T_m+1} - \hat{\mu}_{i,W_2}| \quad (154)$$

$$\geq 4 \max \left(\sqrt{\frac{\log(T^3)}{2|W_{i,1}|}}, \sqrt{\frac{\log(T^3)}{2|W_{i,2}|}} \right) - |\mu_{i,T_m} - \hat{\mu}_{i,W_1}| - |\mu_{i,T_m+1} - \hat{\mu}_{i,W_2}| \quad (\text{by Eq. (149)}) \quad (155)$$

$$\geq \sqrt{\frac{\log(T^3)}{2|W_{i,1}|}} + \sqrt{\frac{\log(T^3)}{2|W_{i,2}|}}. \quad (\text{by Eq. (150), (151)}) \quad (156)$$

which implies that ADR-bandit resets the window at round s . In summary, under Eq. (100) with $p = 1$, \mathcal{V} holds. Therefore,

$$\Pr[\mathcal{V}] \geq 1 - 2K/T. \quad (157)$$

In the following, we bound the regret. Let

$$\text{Reg}_m = \sum_{t=T_{d,m}}^{T_{d,m+1}} \text{reg}(t). \quad (158)$$

Namely, Reg_m corresponding to the regret between m -th and $(m+1)$ -th detection times. The regret is decomposed as

$$\begin{aligned} \text{Reg}(T) &\leq T\mathbf{1}[\mathcal{V}^c] + \mathbf{1}[\mathcal{V}]\text{Reg}(T) \\ &\leq T\mathbf{1}[\mathcal{V}^c] + \mathbf{1}[\mathcal{V}]\text{Reg}(T) \\ &= T\mathbf{1}[\mathcal{V}^c] + \mathbf{1}[\mathcal{V}] \sum_{m=0}^M \text{Reg}_m. \quad (\mathcal{V} \text{ implies } |\mathcal{T}_d| = M) \end{aligned} \quad (159)$$

The regret in the case of \mathcal{V}^c is bounded as

$$T \Pr[\mathcal{V}^c] \leq 2K. \quad (160)$$

Under \mathcal{V} ,

$$T_{d,m+1} - T_{d,m} \leq T_{m+1} - T_m + 16KU(\epsilon_m). \quad (161)$$

Let $\mu_{i,m}$ be the mean of i -th arm between m -th changepoint and $(m+1)$ -th changepoint. Let $\Delta_{i,m} = \max_j \mu_{j,m} - \mu_i \geq 0$ be the corresponding gap.

By the definition of abrupt changes, $\epsilon(t) \leq C^{\text{ab}}\epsilon_m$ after the m -th changepoint and 0 before the changepoint. By Lemma 17 and Eq. (161), we have

$$\mathbb{E} \left[\mathbf{1}[\mathcal{V}] \sum_{m=0}^M \text{Reg}_m \right] \leq \sum_{m=0}^M \sum_{i:\Delta_{i,m}>0} \left(C^{\text{dt}} \frac{\log T}{\Delta_{i,m}} + 2C^{\text{ab}}\epsilon_m \times 16KU(\epsilon_m) \right). \quad (162)$$

Moreover, letting $N_{i,m} = \sum_{t=T_{d,m}+1}^{T_{d,m+1}} \mathbf{1}[I(t) = i]$, we also have

$$\mathbf{1}[\mathcal{V}] \text{Reg}_m \leq \Delta_{i,m} N_{i,m} + 2C^{\text{ab}}\epsilon_m \times 16KU(\epsilon_m),$$

and thus

$$\mathbb{E} \left[\mathbf{1}[\mathcal{V}] \sum_{m=0}^M \text{Reg}_m \right] \leq \sum_{m=0}^M \sum_{i:\Delta_{i,m}>0} \left(\Delta_{i,m} \mathbb{E}[N_{i,m}] + 2C^{\text{ab}}\epsilon_m \times 16KU(\epsilon_m) \right) \quad (163)$$

Taking the minimum of Eq. (162) and Eq. (163), it holds that

$$\mathbb{E} \left[\mathbf{1}[\mathcal{V}] \sum_{m=0}^M \text{Reg}_m \right] \leq \mathbb{E} \left[\sum_{m=0}^M \sum_{i:\Delta_{i,m}>0} \left(\max(1, C^{\text{dt}}) R_{i,m} + 2C^{\text{ab}}\epsilon_m \times 16KU(\epsilon_m) \right) \right], \quad (164)$$

where $R_{i,m} = \min(\Delta_{i,m} \mathbb{E}[N_{i,m}], \frac{\log T}{\Delta_{i,m}})$.

In the following, we bound the two terms of Eq. (164) separately.

The first term of Eq. (164) is bounded by using standard discussion of distribution-independent regret as follows. We have

$$\sum_{m=0}^M \sum_i R_{i,m} = \sum_{m=0}^M \sum_{i:\Delta_{i,m}>0} \sqrt{\Delta_{i,m} R_{i,m}} \sqrt{(\Delta_{i,m})^{-1} R_{i,m}} \quad (165)$$

$$\leq \sqrt{\log T} \sum_{m=0}^M \sum_{i:\Delta_{i,m}>0} \sqrt{(\Delta_{i,m})^{-1} R_{i,m}} \quad (166)$$

$$\leq \sqrt{\log T} \sum_{m=0}^M \sum_{i:\Delta_{i,m}>0} \sqrt{N_{i,m}} \quad (167)$$

$$\leq \sqrt{\log T} \sum_{m=0}^M \sqrt{K(T_{m+1} - T_m)}$$

$$\text{(by the Cauchy-Schwarz inequality and } \sum_i N_{i,m} = (T_{d,m+1} - T_{d,m})) \quad (168)$$

$$\leq \sqrt{\log T} \sqrt{KMT}.$$

$$\text{(by the Cauchy-Schwarz inequality and } \sum_m (T_{m+1} - T_m) = T) \quad (169)$$

The second term of Eq. (164) is bounded by a similar technique as follows.

$$\sum_{m=0}^M C^{\text{ab}} \epsilon_m 16KU(\epsilon_m) \quad (170)$$

$$= \sum_{m=0}^M C^{\text{ab}} \epsilon_m \min \left(16KU(\epsilon_m), \frac{T_{m+1} - T_m}{3} \right) \quad (171)$$

(by $(T_{m+1} - T_m) \geq 48KU(\epsilon_m)$ and $\min(a, b) = a$ if $a \leq b$)

$$= \sum_{m=0}^M C^{\text{ab}} \min \left(\frac{16K \log(T^3)}{\epsilon_m}, \epsilon_m \frac{T_{m+1} - T_m}{3} \right) \quad (172)$$

(by definition of $U(\epsilon_m)$ in Definition 20)

$$\leq \sum_{m=0}^M C^{\text{ab}} \sqrt{16K(\log(T^3))(T_{m+1} - T_m)} \quad (\text{by } \min(a, b) \leq \sqrt{ab}) \quad (173)$$

$$= C^{\text{ab}} \sqrt{16K(\log(T^3))(M+1)T} \quad (174)$$

(by the Cauchy-Schwarz inequality and $\sum_m (T_{m+1} - T_m) = T$)

$$= \tilde{O}(\sqrt{MKKT}). \quad (175)$$

■

E.5 Regret in gradual environment: proof of Theorem 23

We first state Lemmas 33 and 34, then go to the proof of Theorem 23. The high-level implications of these lemmas are as follows. Lemma 33 limits the possible amount of drift such that no reset occurs. Lemma 34 bounds the number of resets as $M_d := |\mathcal{T}_d| = \tilde{O}(T^{1-2d/3})$.

Lemma 33 *With probability at least $1 - 2K/T$, for any $N < |W(t)|, t \in [T], i \in [K]$ and $s, s' \in W(t)$*

$$|\mu_{i,s} - \mu_{i,s'}| \leq \frac{1}{C^{\text{gr}}} \left(3bKN + \tilde{O} \left(\sqrt{1/N} \right) \right) \quad (176)$$

holds.

Lemma 33 is a version of Lemma 26 for the bandit setting. This case is much more challenging mainly because the monitoring arm may change among blocks.

Proof of Lemma 33 We use a tuple (l, c) to represent the c -th subblock of the l -th block for $l = 1, 2, \dots$, and $c = 1, 2, \dots, 2^{l-1}$, that is, the window consisting of $(KN(2^{l-1} + c - 2) + 1, \dots, KN(2^{l-1} + c - 1))$ -th rounds after the last reset. We write \underline{t}_l (resp. \bar{t}_l) for the first (resp. last) round of the l -th block, that is, $\underline{t}_l = KN(2^{l-1} - 1) + 1$ and $\bar{t}_l = KN(2^l - 1)$. The window $W_{:(l,c)}$ consists of all subblocks before $W_{(l,c)}$ (not including $W_{(l,c)}$). Similarly, subwindow $W_{(l,c):(l,c')}$ for $c < c'$ denotes the joint window consisting of $W_{(l,c)}, W_{(l,c+1)}, \dots, W_{(l,c'-1)}$.

Fix an arbitrary $l \in \mathbb{N}$. The second property of the monitoring consistency (Definition 15) implies the following: Assume that no reset occurred up to the l -th block. There exists an arm that is drawn at least N times for each subblock $c = 1, 2, \dots, 2^{l-1}$ in the l -th block. Moreover, this arm is drawn at least N times in the final subblock of the $(l-1)$ -th block.

By the property above and the fact that no reset occurs up to subblock (l, c) , there exists i_l such that for any $l \in \mathbb{N}$ and $c \in [2^{l-1}]$

$$|\hat{\mu}_{i_l, W_{(l,c)}} - \hat{\mu}_{i_l, W_{:(l,c)}}| \leq 2\sqrt{\frac{\log(T^3)}{2N}}, \quad (177)$$

$$|\hat{\mu}_{i_l, W_{(l,1):(l,c)}} - \hat{\mu}_{i_l, W_{:(l,1)}}| \leq 2\sqrt{\frac{\log(T^3)}{2N}} \quad (178)$$

because otherwise a reset should occur. Then, for any $l \geq 2$ and $2 \leq c \leq 2^{l-1}$ we have

$$\begin{aligned} & |\hat{\mu}_{i_l, W_{(l,c)}} - \hat{\mu}_{i_l, W_{(l,1)}}| \\ & \leq |\hat{\mu}_{i_l, W_{:(l,c)}} - \hat{\mu}_{i_l, W_{:(l,1)}}| + 2\sqrt{\frac{\log(T^3)}{2N}} \quad (\text{by (177)}) \\ & = \left| \frac{N_{i_l, W_{(l,1):(l,c)}} \hat{\mu}_{i_l, W_{(l,1):(l,c)}} + (N_{i_l, W_{:(l,c)}} - N_{i_l, W_{(l,1):(l,c)}}) \hat{\mu}_{i_l, W_{:(l,1)}}}{N_{i_l, W_{:(l,c)}}} - \hat{\mu}_{i_l, W_{(l,1)}} \right| + 2\sqrt{\frac{\log(T^3)}{2N}} \\ & \leq \left| \frac{N_{i_l, W_{(l,1):(l,c)}} \hat{\mu}_{i_l, W_{:(l,1)}} + (N_{i_l, W_{:(l,c)}} - N_{i_l, W_{(l,1):(l,c)}}) \hat{\mu}_{i_l, W_{:(l,1)}}}{N_{i_l, W_{:(l,c)}}} - \hat{\mu}_{i_l, W_{(l,1)}} \right| + 4\sqrt{\frac{\log(T^3)}{2N}} \quad (\text{by (178)}) \\ & = \left| \hat{\mu}_{i_l, W_{:(l,1)}} - \hat{\mu}_{i_l, W_{(l,1)}} \right| + 4\sqrt{\frac{\log(T^3)}{2N}} \\ & \leq 6\sqrt{\frac{\log(T^3)}{2N}} \quad (\text{by (177)}). \end{aligned} \quad (179)$$

Also, for $c = 1$, (179) is trivial. For $l = 1$, it is also trivial since $c = 1$ must hold from $c \leq 2^{l-1}$. By following the same discussion, we also have

$$|\hat{\mu}_{i_l, W_{(l,c)}} - \hat{\mu}_{i_l, W_{(l,2^{l-1})}}| \leq 6\sqrt{\frac{\log(T^3)}{2N}}. \quad (180)$$

By Eq. (100) with $p = 1$ we have

$$|\mu_{i, W_{l,c}} - \hat{\mu}_{i, W_{l,c}}| \leq \sqrt{\frac{\log(T^3)}{2N}} \quad (181)$$

for all $l \in \mathbb{N}, c \in 2^{n-1}$ with probability at least $1 - 2K/T$. By the fact that μ_t moves at most bKN within a subblock of size KN ,

$$\begin{aligned} |\mu_{i, W_{(l,c)}} - \mu_{i, t}| & \leq bKN \quad \text{for round } t \text{ in subblock } W_{l,c} \\ |\mu_{i, W_{(l-1, 2^{l-2})}} - \mu_{i, W_{(l,1)}}| & \leq bKN. \end{aligned} \quad (182)$$

Now, let $W_{l,c}$ (resp. $W_{l',c'}$) be the subwindow that s -th (resp., s' -th) round belongs to. Here, we assume without loss of generality that $s < s'$. Then we have

$$|\mu_{i,s} - \mu_{i,s'}|$$

$$\begin{aligned}
 &\leq |\mu_{i,s} - \mu_{i,\bar{t}_l}| + |\mu_{i,\bar{t}_l} - \mu_{i,t_{l+1}}| + |\mu_{i,t_{l+1}} - \mu_{i,s'}| \\
 &\leq |\mu_{i,s} - \mu_{i,\bar{t}_l}| + b + |\mu_{i,t_{l+1}} - \mu_{i,s'}| \tag{183}
 \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{1}{C^{\text{gr}}} |\mu_{i,s} - \mu_{i,\bar{t}_l}| + b + |\mu_{i,t_{l+1}} - \mu_{i,s'}| \\
 &\leq \frac{1}{C^{\text{gr}}} (|\mu_{i,W_{(l,c)}} - \mu_{i,W_{(l,2^{l-1})}}| + 2bKN) + b + |\mu_{i,t_{l+1}} - \mu_{i,s'}| \quad (\text{by (182)}) \\
 &\leq \frac{1}{C^{\text{gr}}} \left(8\sqrt{\frac{\log(T^3)}{2N}} + 2bKN \right) + b + |\mu_{i,t_{l+1}} - \mu_{i,s'}|, \quad (\text{by (180) and (181)}) \tag{184}
 \end{aligned}$$

and recursively applying this transformation for $l, l+1, l+2, \dots, l'-1$, we have

$$\begin{aligned}
 &|\mu_{i,s} - \mu_{i,s'}| \\
 &\leq \frac{l' - l}{C^{\text{gr}}} \left(8\sqrt{\frac{\log(T^3)}{2N}} + 2bKN \right) + b(l' - l) + |\mu_{i,t_{l'}} - \mu_{i,s'}| \\
 &\leq \frac{l' - l + 1}{C^{\text{gr}}} \left(8\sqrt{\frac{\log(T^3)}{2N}} + 2bKN \right) + b(l' - l),
 \end{aligned}$$

where we obtain the last inequality by applying the same transformation as that from (183) to (184). We obtain the desired result since $l' \leq \log_2 T = O(\log T) = \tilde{O}(1)$. \blacksquare

Lemma 34 *There exists an event \mathcal{X} that holds with probability at least*

$$\Pr[\mathcal{X}] \leq \tilde{O}(KT^{-1}) \tag{185}$$

such that, under \mathcal{X}^c ,

$$M_d \leq T/(C_1 b^{-2/3}) = \tilde{O}(T^{1-2d/3}) \tag{186}$$

holds.

Lemma 34 is a version of Lemma 27 for the bandit setting. In the following, we derive Lemma 34 for completeness. The steps are very similar to the proof of Lemma 27.

Proof of Lemma 34 Let

$$\mathcal{X}_j(t) = \bigcup_{W_1, W_2: W(t)=W_1 \cup W_2, j \in [K]} \left\{ |W_1| \leq C_1 b^{-2/3}, |W_2| \leq C_1 b^{-2/3}, |\hat{\mu}_{j,W_1} - \hat{\mu}_{j,W_2}| \geq \epsilon_{\text{cut}}^\delta \right\} \tag{187}$$

where $C_1 = \tilde{O}(1)$ is a factor that is specified in Eq. (200). Let $\mathcal{X} = \bigcup_{t \in [T], j \in [K]} \mathcal{X}_j(t)$. Let

$$\mathcal{W}_{C_1} = \{W_0 \in \mathcal{W} : |W_0| \leq C_1 b^{-2/3}\} \tag{188}$$

be the set of windows of size at most $C_1 b^{-2/3}$. It is easy to show that $|\mathcal{W}_{C_1}| \leq TC_1 b^{-2/3}$. For a fixed window W and $i \in [K]$, Hoeffding inequality states that

$$|\mu_{i,W} - \hat{\mu}_{i,W}| > \sqrt{\frac{\log(\eta^{-1})}{2|W_i|}}, \tag{189}$$

occurs with probability at most 2η . By the union bound of Eq. (189) with $\eta^{-1} = T^{2+d}$ over all i and all windows of \mathcal{W}_{C_1} and all K arms, with probability at least

$$1 - \frac{2K}{T^{2+d}} \times TC_1 b^{-2/3} = 1 - 2C_1 K T^{-1} b^{1/3} \geq 1 - 2C_1 K T^{-1}, \quad (\text{by } T^{-d} = b) \quad (190)$$

In summary,

$$\bigcap_{i, W \in \mathcal{W}_{C_1}} |\mu_{i,W} - \hat{\mu}_{i,W}| \leq \sqrt{\frac{\log(T^{2+d})}{2|W_i|}} \quad (191)$$

holds with probability at least $1 - 2C_1 K T^{-1}$. In the following, we show that \mathcal{X} never occurs under the event of Eq. (191). In particular we use the proof-by-contradiction argument on the split event $|\hat{\mu}_{j,W_1} - \hat{\mu}_{j,W_2}| \geq \epsilon_{\text{cut}}^\delta$. Eq. (191) implies

$$\begin{aligned} |\mu_{j,W_1} - \hat{\mu}_{j,W_1}| &\leq \sqrt{\frac{\log(T^{2+d})}{2|W_{j,1}|}} \\ |\mu_{j,W_2} - \hat{\mu}_{j,W_2}| &\leq \sqrt{\frac{\log(T^{2+d})}{2|W_{j,2}|}} \end{aligned} \quad (192)$$

for all j, t and any split $W_1 \cup W_2 = W(t) : |W_1|, |W_2| \leq C_1 b^{-2/3}$. Let $N = C_1 b^{-2/3}$. By the definition of the gradually changing stream,

$$|\mu_{j,W_1} - \mu_{j,W_2}| \leq 2bN. \quad (193)$$

Assuming that $|\hat{\mu}_{j,W_1} - \hat{\mu}_{j,W_2}| \geq \epsilon_{\text{cut}}^\delta$, it holds that

$$\epsilon_{\text{cut}}^\delta = \sqrt{\frac{\log(T^3)}{2|W_{j,1}|}} + \sqrt{\frac{\log(T^3)}{2|W_{j,2}|}} \quad (\text{by the definition of } \epsilon_{\text{cut}}^\delta \text{ in (1)}) \quad (194)$$

$$\leq |\hat{\mu}_{j,W_1} - \hat{\mu}_{j,W_2}| \quad (195)$$

$$\leq |\mu_{j,W_1} - \mu_{j,W_2}| + |\hat{\mu}_{j,W_1} - \mu_{j,W_1}| + |\hat{\mu}_{j,W_2} - \mu_{j,W_2}| \quad (196)$$

(by triangular inequality)

$$\leq 2bN + |\hat{\mu}_{j,W_1} - \mu_{j,W_1}| + |\hat{\mu}_{j,W_2} - \mu_{j,W_2}| \quad (\text{by (193)}) \quad (197)$$

$$\leq 2bN + \sqrt{\frac{\log(T^{2+d})}{2|W_{j,1}|}} + \sqrt{\frac{\log(T^{2+d})}{2|W_{j,2}|}}. \quad (\text{by (192)}) \quad (198)$$

Eq. (194) \leq Eq. (198) is equivalent to

$$\frac{\sqrt{3} - \sqrt{2+d}}{\sqrt{2}} \left(\sqrt{\frac{\log T}{|W_{j,1}|}} + \sqrt{\frac{\log T}{|W_{j,2}|}} \right) \leq 2bN, \quad (199)$$

which cannot hold for

$$N = b^{-2/3} \underbrace{\left(\frac{\sqrt{3} - \sqrt{2+d}}{2\sqrt{2}} \sqrt{\log T} \right)^{2/3}}_{=: C_1}, \quad |W_{j,1}|, |W_{j,2}| \leq N. \quad (200)$$

Therefore, $\mathcal{X}_j(t)$ never occurs under the event of Eq. (191). \blacksquare

Proof of Theorem 23 Let $R(t)$ be the most recent reset before t , or $R(t) = 0$ if no reset has occurred yet. Let

$$\epsilon_R(t) := \max_{s:R(t)<s<t} \max_i |\mu_{i,s} - \mu_{i,R(t)+1}|, \quad (201)$$

which is the amount of drift in view of the current ADR-bandit. By Eq. (18) and Lemma 33, the event

$$\mathcal{Y} = \bigcap_{t \in [T]} \left\{ \epsilon_R(t) \leq (C^{\text{gr}})^{-1} \left(3bKN + \tilde{O} \left(\sqrt{1/N} \right) \right) \right\} \quad (202)$$

holds with probability at least $1 - 2K/T$.

We assume $\mathcal{X}^c \cap \mathcal{Y}$, because it holds with probability $1 - O(K/T)$ and thus the regret under $\mathcal{X} \cup \mathcal{Y}^c$ is negligible. Eq. (186) states that event \mathcal{X}^c implies that $M_d = O(T^{1-2d/3})$. Event \mathcal{Y} with $N = \tilde{\Theta}((bK)^{-2/3})$ implies that

$$\epsilon_R(t) = \tilde{O} \left((bK)^{1/3} \right) \quad (203)$$

for $t \in [T]$.

By Eq. (186) we see that $M_d \leq T/(C_1 b^{-2/3})$ under \mathcal{X}^c . Letting $T_{d,0} = 1$ and $T_{d,m} = T$ for $m > M_d$, we first have

$$\mathbb{E} [\text{Reg}(T)] \leq \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\sum_{t=T_{d,m}+1}^{T_{d,m+1}} \text{reg}(t) \right] + \mathbb{E} [\mathbf{1}[\mathcal{X}^c] \text{Reg}(T)] \quad (204)$$

$$\leq \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\sum_{t=T_{d,m}+1}^{T_{d,m+1}} \text{reg}(t) \right] + O(1). \quad (205)$$

Let $N_{i,m} = \sum_{t=T_{d,m}+1}^{T_{d,m+1}} \mathbf{1}[I(t) = i]$ be the number of draw on arm i between the m -th and $(m+1)$ -th reset, and $\Delta_{i,m} = \max_j \mu_{j,T_{d,m}+1} - \mu_{i,T_{d,m}+1}$ be the gap at the first round after the m -th reset. Let \mathcal{H}_m be the history until the m -th reset. Then the regret between the m -th and $(m+1)$ -th reset is bounded by Lemma 17 and we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=T_{d,m}+1}^{T_{d,m+1}} \text{reg}(t) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\sum_{t=T_{d,m}+1}^{T_{d,m+1}} \text{reg}(t) \middle| \mathcal{H}_m \right] \right] \end{aligned} \quad (206)$$

$$\leq \mathbb{E} \left[\sum_i \left(\min \left\{ \Delta_{i,m} \mathbb{E}[N_{i,m} | \mathcal{H}_m], O \left(\frac{\log T}{\Delta_{i,m}} \right) \right\} + \mathbb{E} \left[N_{i,m} \max_t \epsilon_R(t) \middle| \mathcal{H}_m \right] \right) \right]. \quad (207)$$

Here, the summation over m for the first term is transformed as follows.

$$\begin{aligned}
 & \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\sum_i \min \left\{ \Delta_{i,m} \mathbb{E}[N_{i,m} | \mathcal{H}_m], O \left(\frac{\log T}{\Delta_{i,m}} \right) \right\} \right] \\
 & \leq \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\sum_i O \left(\sqrt{\mathbb{E}[N_{i,m} | \mathcal{H}_m] \log T} \right) \right] \quad (\text{by } \min(a, b) \leq \sqrt{ab}) \tag{208}
 \end{aligned}$$

$$\leq \sum_{m=1}^{T/(C_1 b^{-2/3})} \sum_i O \left(\sqrt{\mathbb{E}[\mathbb{E}[N_{i,m} | \mathcal{H}_m]] \log T} \right) \tag{209}$$

$$= \sum_{m=1}^{T/(C_1 b^{-2/3})} \sum_i O \left(\sqrt{\mathbb{E}[N_{i,m}] \log T} \right) \tag{210}$$

$$\leq O \left(\sqrt{T (T/(C_1 b^{-2/3})) K \log T} \right) = \tilde{O}(T b^{1/3} \sqrt{K}), \tag{211}$$

where, in the second and third inequalities, we have applied Jensen's inequality on $T/(C_1 b^{-2/3}) \times K$ elements such that $\sum_{m=1}^{T/(C_1 b^{-2/3})} \sum_i \mathbb{E}[N_{i,m}] = T$ holds. For the summation over the second term we have

$$\begin{aligned}
 & \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\sum_i \mathbb{E} \left[N_{i,m} \max_t \epsilon_R(t) \middle| \mathcal{H}_m \right] \right] \\
 & \leq \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\mathbf{1}[\mathcal{Y}] \sum_i \mathbb{E} \left[N_{i,m} \max_t \epsilon_R(t) \middle| \mathcal{H}_m \right] \right] + \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\mathbf{1}[\mathcal{Y}^c] \sum_i \mathbb{E} \left[N_{i,m} \max_t \epsilon_R(t) \middle| \mathcal{H}_m \right] \right] \tag{212}
 \end{aligned}$$

$$\begin{aligned}
 & \leq \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\sum_i \mathbb{E} \left[N_{i,m} \max_t \tilde{O}((bK)^{1/3}) \middle| \mathcal{H}_m \right] \right] + \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\mathbf{1}[\mathcal{Y}^c] \sum_i \mathbb{E} [N_{i,m} | \mathcal{H}_m] \right] \\
 & \quad (\text{by Eq. (203)}) \tag{213}
 \end{aligned}$$

$$= \tilde{O}((bK)^{1/3}) \sum_{m=1}^{T/(C_1 b^{-2/3})} \sum_i \mathbb{E} [N_{i,m}] + \sum_{m=1}^{T/(C_1 b^{-2/3})} \mathbb{E} \left[\mathbf{1}[\mathcal{Y}^c] \sum_i \mathbb{E} [N_{i,m} | \mathcal{H}_m] \right] \tag{214}$$

$$= \tilde{O}(T(bK)^{1/3}) + O(1). \tag{215}$$

In summary,

$$\mathbb{E} [\text{Reg}(T)] \leq \tilde{O}(T b^{1/3} \sqrt{K}) + \tilde{O}(T(bK)^{1/3}) + O(1) = \tilde{O}(T b^{1/3} \sqrt{K}). \tag{216}$$

■