
Almost Minimax Optimal Best Arm Identification in Piecewise Stationary Linear Bandits

Yunlong Hou

Department of Mathematics
National University of Singapore
yhou@u.nus.edu

Vincent Y. F. Tan

Department of Mathematics
Department of Electrical and Computer Engineering
National University of Singapore
vtan@nus.edu.sg

Zixin Zhong

Data Science and Analytics Thrust
Hong Kong University of Science and Technology (Guangzhou)
zixinzhong@hkust-gz.edu.cn

Abstract

We propose a *novel* piecewise stationary linear bandit (PSLB) model, where the environment randomly samples a context from an unknown probability distribution at each changepoint, and the quality of an arm is measured by its return averaged over all contexts. The contexts and their distribution, as well as the changepoints are unknown to the agent. We design *Piecewise-Stationary ε -Best Arm Identification⁺* ($\text{PS}\varepsilon\text{BAI}^+$), an algorithm that is guaranteed to identify an ε -optimal arm with probability $\geq 1 - \delta$ and with a minimal number of samples. $\text{PS}\varepsilon\text{BAI}^+$ consists of two subroutines, $\text{PS}\varepsilon\text{BAI}$ and $\text{NAIVE } \varepsilon\text{-BAI}$ ($\text{N}\varepsilon\text{BAI}$), which are executed in parallel. $\text{PS}\varepsilon\text{BAI}$ actively detects changepoints and aligns contexts to facilitate the arm identification process. When $\text{PS}\varepsilon\text{BAI}$ and $\text{N}\varepsilon\text{BAI}$ are utilized judiciously in parallel, $\text{PS}\varepsilon\text{BAI}^+$ is shown to have a finite expected sample complexity. By proving a lower bound, we show the expected sample complexity of $\text{PS}\varepsilon\text{BAI}^+$ is optimal up to a logarithmic factor. We compare $\text{PS}\varepsilon\text{BAI}^+$ to baseline algorithms using numerical experiments which demonstrate its efficiency. Both our analytical and numerical results corroborate that the efficacy of $\text{PS}\varepsilon\text{BAI}^+$ is due to the delicate change detection and context alignment procedures embedded in $\text{PS}\varepsilon\text{BAI}$.

1 Introduction

In stochastic *multi-armed bandits* (MABs), an agent interacts with the environment at each time step. The agent pulls an arm and observes the corresponding return provided by the environment. The classical MAB framework assumes a *stationary* environment where the expected return of each arm remains unchanged over time. However, we usually face ever-changing environments in real life. For instance, in investment option selection and portfolio management, fund managers want to select a subset of good candidate portfolios. However, the market may be bullish, bearish, or in some other state. The transition between these states can be well-modelled as being stochastic. We wish to select portfolios that yield the best long-term option under such a piecewise stationary environment. Further examples such as one based on agriculture in the face of stochastically changing weather patterns are discussed in detail in Appendix A. These motivate us to formulate and investigate a *piecewise stationary linear bandit* (PSLB) model.

Our PSLB model is equipped with an *arm set* \mathcal{X} , a *context set* Θ and a deterministic but unknown sequence of *changepoints* \mathcal{C} . At each changepoint, the environment samples a context $\theta \in \Theta$ from an

unknown probability distribution P_θ , and the returns of arms may change when the context changes. The return of each arm under each context is determined by its feature $x \in \mathcal{X}$ and the context θ . In particular, the expected return of an arm is the weighted sum $\mu_x = \mathbb{E}_{P_\theta}[x^\top \theta]$. While the sequence of changepoints, as well as the distribution and latent vectors of contexts are not known, the agent samples an arm and observes the corresponding return at each time step so that it can identify the best arm $\arg \max_x \mu_x$ up to some tolerance ε with probability $\geq 1 - \delta$ and with as few samples as possible. The agent’s behavior does not affect the sequence of contexts that is drawn from P_θ .

Main Contributions. We are the *first* to study the fixed-confidence *best arm identification* (BAI) problem in *piecewise stationary bandits* (PSB). Given $\delta > 0$, we say the arm with the highest expected return μ^* is the *best*, and an arm is ε -optimal if its expected return is at least $\mu^* - \varepsilon$. We seek to design an (ε, δ) -PAC algorithm which can identify an ε -optimal arm with probability $\geq 1 - \delta$ in as few time steps as possible, i.e., with minimal sample complexity.

Our **first** contribution concerns the formulation of a *novel* PSLB model, where we measure the quality of an arm x according to its *expected return* $\mu_x = \mathbb{E}_{\theta \sim P_\theta}[x^\top \theta]$ for the following reasons. Consider that an arm is measured by its average return across time, which is a generalization of the definition in *stationary bandits* (SB). A notable feature of PSB models is that the context changes as time evolves, and hence the arm’s average return across time also changes, in general. Hence, we aim to identify an arm whose *average return across contexts* is high, and benefits the agent for interacting with the environment in the *long run* after the arm identification task. We are thus inspired to introduce the distribution of contexts under the PSLB model, define the expected return μ_x for each arm, and use this ensemble (non-time varying) statistic as the benchmark for what we seek to learn. The BAI task using this statistic is meaningful but challenging, as the agent needs to reliably estimate the context vectors, changepoints, and context distribution.

Secondly, we propose PIECEWISE-STATIONARY ε -BAI⁺ (PS ε BAI⁺), an algorithm designed to tackle the BAI problem in our PSLB model. We prove that it is (ε, δ) -PAC and bound its sample complexity. PS ε BAI⁺ samples arms according to a suitably defined G-optimal allocation, and runs two algorithms: NAÏVE ε -BAI (N ε BAI) and PS ε BAI as subroutines in parallel to achieve efficiency and attain a bound on the sample complexity in expectation.

- Being a baseline but naïve algorithm, the complexity of N ε BAI grows linearly with the maximum length between two changepoints L_{\max} , motivating us to design a more efficient algorithm, PS ε BAI, to reduce the impact of L_{\max} .
- PS ε BAI is equipped with two delicately designed subroutines LINEAR-CHANGE DETECTION (LCD) and LINEAR-CONTEXT ALIGNMENT (LCA) to actively detect changepoints and align contexts with those observed in the previous time steps respectively. Concretely, **in terms of the design**, PS ε BAI determines whether samples from two intervals are under the identical context via a sliding window mechanism, and detects changepoints and aligns contexts accordingly; this facilitates the estimation of context vectors and their distribution. Combining these elements into the design of PS ε BAI and analyzing them requires some care. **On the theoretical side**, we prove PS ε BAI identifies an ε -optimal arm faster than N ε BAI with high probability. The success of PS ε BAI relies on the LCD and LCA subroutines, while a minor drawback is that they require a non-vanishing failure probability budget which does not allow us to bound the sample complexity of PS ε BAI in expectation. To achieve a complete theoretical understanding, we delicately design the PS ε BAI⁺ algorithm whose efficiency is inherited via the LCD and LCA procedures in PS ε BAI as well as the effective utilization of running PS ε BAI and N ε BAI in parallel.

Thirdly, we derive a lower bound on the complexity of any (ε, δ) -PAC algorithm in PSLB models. To derive this bound, we first lower bound the complexity of an algorithm when the contextual information (and changepoints) are available, and then quantify the number of arm samples (and realized contexts) required to reliably infer an ε -best arm. We compare the upper bound of PS ε BAI⁺ and this generic lower bound in several instances. The matching (up to logarithmic terms) of bounds illustrate that our PS ε BAI⁺ algorithm is almost asymptotically optimal.

Lastly, we demonstrate the efficiency of PS ε BAI⁺ with numerical experiments. The first half of our experiment shows that PS ε BAI⁺ is (ε, δ) -PAC and with significantly lower sample complexity compared to N ε BAI, corroborating our theoretical findings. In the second half, we compare PS ε BAI⁺ to N ε BAI, and two other benchmarks DISTRIBUTION ε -BAI (D ε BAI) and D ε BAI $_\beta$. While contexts and changepoints are not available to PS ε BAI⁺ and N ε BAI, they are observed by D ε BAI and D ε BAI $_\beta$. Nevertheless, PS ε BAI⁺ is still competitive compared to D ε BAI and D ε BAI $_\beta$ in our

empirical experiments. Hence, both experiments justify the necessity of the change detection and context alignment procedures for boosting the learning of contexts and their distributions, as well as the identification of the best arm. We also show empirically that misspecifications to the knowledge of L_{\min} and L_{\max} do not affect the performance of our algorithm significantly.

Related work. *Best arm identification* (BAI) and *regret minimization* (RM) are two fundamental problems in multi-armed bandits. In stationary linear bandits, [1, 2] focus on BAI and [3, 4] aim to solve the RM problem. An efficient algorithm can choose the G-optimal allocation or $\mathcal{X}\mathcal{Y}$ -allocation rule to quickly identify a good arm [1, 5]. Besides, [6, 7] focus on ε -optimal arm identification, which is a generalization of the standard BAI problem.

The BAI and RM problems are also studied in the *non-stationary bandits* (NSB), where in contrast to the SB model, the context varies with time [8–10]. NSB can be largely divided into two classes: the drifting bandit (DB) model, where the context changes at each step [8], and PSB, where the context changes less frequently [10]. [11] provides an extensive discussion on the definition of NSB models.

On one hand, the RM problem has been investigated extensively in DB models [12, 13]. On the other hand, concerning the BAI task in DB models, [14] investigated BAI with a *fixed-horizon*, where the best arm has the highest average return over this horizon; [15] assumes the best arm remains unchanged after certain time step and explores the *fixed-confidence* setting. Besides, when the contextual information in NSB models is available, NSB models are known as *contextual bandit* (CB) models [16–18]. [16] showed that the contextual information accelerates the best arm identification process. More discussions on DB and CB models are presented in Appendix C.

Moreover, the context can drift dramatically in a DB model while it remains unvarying in a SB model. Straddling between the DB and SB models, PSB models assume there is an interim stationary interval between each two consecutive changepoints, where the context remains unchanged. The context changes can be characterized in different ways and affect the performance of proposed algorithms in PSB models. For instance, a changepoint signals the return of at least one arm shift as in [10], and indicates the best arm changes as in [19].

In PSB models, a large body of works focus on RM. While [10, 20–22] equip their algorithms with changepoint detection techniques to handle the context changes, [23] actively checks the quality of each arm. However, there is no existing work on the fixed-confidence BAI problem in a PSB model. To fill this gap in the literature, we design $\text{PS}\varepsilon\text{BAI}^+$ for ε -optimal arm identification in our proposed PSLB model. We show $\text{PS}\varepsilon\text{BAI}^+$ is almost asymptotically optimal by comparing its complexity to a generic lower bound of all algorithms, and validate its efficiency through numerical experiments.

2 Problem Setup

For $m \in \mathbb{N}$, let $[m] := \{1, 2, \dots, m\}$. For a finite set S , let Δ_S denote the $|S|$ -dim probability simplex on S . Let $A(\mathbf{q}) := \sum_{x \in \mathcal{X}} q_x x x^\top$ be the matrix induced by $\mathbf{q} \in \Delta_{\mathcal{X}}$ with $\mathcal{X} \subset \mathbb{R}^d$. An instance of *piecewise stationary linear bandit* is a tuple $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$. Specifically, $x \in \mathbb{R}^d$ is an arm (vector) and the arm set $\mathcal{X} \subset \mathbb{R}^d$ is composed of $|\mathcal{X}| = K$ arms that spans \mathbb{R}^d . The latent vector matrix $\Theta = (\theta_1^*, \dots, \theta_N^*) \in \mathbb{R}^{d \times N}$ contains N latent column vectors where the j^{th} column θ_j^* is associated with context $j \in [N]$. For the sake of normalization, we assume $|x^\top \theta_j^*| \leq 1$ for all $x \in \mathcal{X}, j \in [N]$. Let P_θ denote the distribution (probability mass function) of the latent vectors and $p_j = P_\theta[\theta_j^*]$. We represent the probabilities of latent vectors as $\mathbf{p} = (p_1, \dots, p_N) \in \Delta_N$. The fixed but unknown sequence of changepoints $\mathcal{C} := (c_1, c_2, \dots)$ is an sequence of increasing positive integers $1 = c_1 < c_2 < \dots$, characterizing all the changepoints (time steps).

The *return* of arm x under latent context j is a random variable $Y = x^\top \theta_j^* + \eta$, where η is a zero-mean random variable (noise) supported on $[-1, 1]$, and the *expected return* of arm x is $\mu_x := \mathbb{E}_{\theta \sim P_\theta}[x^\top \theta] = \sum_{j=1}^N P_\theta[\theta_j^*] x^\top \theta_j^*$. The *best arm*, which we assume is unique, is denoted as $x^* := \arg \max_{x \in \mathcal{X}} \mu_x$ with mean $\mu^* := \mu_{x^*}$. Given a slackness parameter $\varepsilon > 0$, we define the set of ε -*best arms* $\mathcal{X}_\varepsilon := \{x \in \mathcal{X} : \mu_x \geq \mu^* - \varepsilon\}$. For each pair of arms $(x, \tilde{x}) \in \mathcal{X}^2$, define the *contextual mean gap* between x and \tilde{x} under latent context j as $\Delta_j(x, \tilde{x}) := (x - \tilde{x})^\top \theta_j^*$ and the *mean gap* between x and \tilde{x} as $\Delta(x, \tilde{x}) := \mu_x - \mu_{\tilde{x}}$.

Given $l \in \mathbb{N}$, the interval $(c_l, \dots, c_{l+1} - 1)$ is known as the l^{th} *stationary segment* and its length is $L_l := c_{l+1} - c_l$. We assume $L_{\min} \leq L_l \leq L_{\max}$. Let $l_t := \max\{l : c_l \leq t\}$ denote the number of

stationary segments up to time step t . At time step $t \in [T]$ (see Dynamics 1),

(i) If $t \in \mathcal{C}$, the environment samples a latent vector $\theta_{j_t}^*$ according to P_θ , that is, it generates a (latent) *context sample* with P_θ ; otherwise the latent vector remains unchanged, i.e., $\theta_{j_t}^* = \theta_{j_{t-1}}^*$. The contexts $\{\theta_{j_{c_l}}^*\}_{l \in \mathbb{N}}$ are sampled i.i.d. from P_θ at each changepoint $\{c_l\}_{l \in \mathbb{N}}$.

(ii) The agent pulls an arm x_t and observes the stochastic return $Y_{t,x_t} = x_t^\top \theta_{j_t}^* + \eta_t$, where η_t is drawn independently from a distribution supported on $[-1, 1]$.

The agent uses an *online* algorithm $\pi := \{(\pi_t, \tau_t, r_t)\}_{t \in \mathbb{N}}$ to actively interact with the instance Λ and only has access to the arm set \mathcal{X} , number of latent vectors N , the bounds on the length of each segment L_{\min} and L_{\max} , the slackness parameter ε , and the confidence parameter δ .

- sampling rule $\pi_t : \mathcal{H}_t^\pi := ((x_s^\pi, Y_{s,x_s^\pi}))_{s \in [t-1]} \rightarrow \mathcal{X}$ samples an arm x_t^π based on the observation history \mathcal{H}_t^π and observe the corresponding random return Y_{t,x_t^π} ;

- stopping rule $\tau_t : \mathcal{H}_{t+1}^\pi \rightarrow \{\text{STOP}, \text{CONTINUE}\}$ decides whether to stop or continue to execute given the observation history \mathcal{H}_{t+1}^π . The stopping time under algorithm π is denoted as τ^π ;

- recommendation rule $r_\tau : \mathcal{H}_{\tau+1}^\pi \rightarrow \mathcal{X}$ recommends an arm \hat{x}^π based on $\mathcal{H}_{\tau+1}^\pi$ upon termination.

The stopping time τ^π is the *sample complexity* of the algorithm π under instance Λ . The *expected sample complexity* is $\mathbb{E}[\tau^\pi]$, where the expectation is taken w.r.t. the random returns, the realization of the contexts governed by the latent vector distribution P_θ , and the randomness of the algorithm π .

An algorithm π is (ε, δ) -PAC (Probably Approximately-Correct) if

$$\mathbb{P}[\hat{x}^\pi \in \mathcal{X}_\varepsilon] \geq 1 - \delta.$$

Our overarching goal in this paper is to devise an (ε, δ) -PAC algorithm that minimizes τ^π with high probability (w.h.p.) and in expectation.

3 Algorithms

3.1 A Naïve Baseline: N ε BAI

We first devise the NAÏVE ε -BEST ARM IDENTIFICATION (or N ε BAI) algorithm (presented in Algorithm 2). In the design of N ε BAI, only the choice of confidence radius $\tilde{\rho}_t$ takes the potential context changes into consideration. Even though N ε BAI does not attempt to detect potential changes in the context, it can identify an ε -optimal arm w.h.p. and is with a finite expected sample complexity.

Proposition 3.1. Let $\Delta_{\min} = \min_{x \neq x^*} \Delta(x^*, x)$,

$$T_V^N = \frac{d}{(\varepsilon + \Delta_{\min})^2} \ln \frac{1}{\delta} \quad \text{and} \quad T_D^N = \frac{L_{\max}}{(\varepsilon + \Delta_{\min})^2} \ln \frac{1}{\delta}.$$

The N ε BAI algorithm is (ε, δ) -PAC and its expected sample complexity is $\tilde{O}(T_V^N + T_D^N)$.

The upper bound in Proposition 3.1 (also see Appendix F) consists of two main terms. (i) As N ε BAI samples arms according to the G-optimal allocation (see Appendix D), the amount of samples needed to estimate the average of latent vectors $\sum_{s=1}^t \theta_{j_s}^*/t$ contributes to T_V^N . (ii) T_D^N quantifies how fast $\sum_{s=1}^t \theta_{j_s}^*/t$ converges to the expectation of context vectors $\sum_{j=1}^N p_j \theta_j^*$.

The sample complexity of N ε BAI grows linearly with L_{\max} , but we surmise that the sample complexity of a close-to-optimal algorithm should have a *reduced dependence on L_{\max}* .

3.2 Piecewise-Stationary ε -Best Arm Identification

The algorithm PIECEWISE-STATIONARY ε -BEST ARM IDENTIFICATION (or PS ε BAI) is presented in Algorithm 1. By using a sliding window mechanism, PS ε BAI actively detects the changepoints and aligns the current latent context with contexts observed in the previous time steps via LINEAR-CHANGE DETECTION (or LCD) and LINEAR-CONTEXT ALIGNMENT (or LCA), which are presented in Algorithms 3 and 4 (see App. D.2.2), respectively. PS ε BAI consists of three phases:

(i) *Exploration phase (Exp)*: Estimate latent vectors and their distribution P_θ (Lines 8 to 11 and 25);

(ii) *Change Detection phase (CD)*: Detect changepoints (Lines 12 to 16);

(iii) *Context Alignment phase (CA)*: Evaluate the current context and align it with the contexts observed in previous time steps (Lines 17 to 21).

At time t , we estimate Θ and \mathbf{p} with $\hat{\Theta}_t = (\hat{\theta}_{t,1}, \dots, \hat{\theta}_{t,N})$ and $\hat{\mathbf{p}}_t = (\hat{p}_{t,1}, \dots, \hat{p}_{t,N})^\top$, respectively.¹ We denote the empirical mean gap between x and \tilde{x} under context j as $\hat{\Delta}_{t,j}(x, \tilde{x}) := (x - \tilde{x})^\top \hat{\theta}_{t,j}$.

PS ε BAI first computes the G-optimal allocation [1] λ^* on the arm set \mathcal{X} and its maximum possible stopping time τ^* (Line 2). It initializes CD_{sample} and CA_{id}. CD_{sample} collects samples to detect changepoints and CA_{id} maintains a dictionary of {latent context index : identification samples} pairs (Line 3);² CA_{id}[j] is the sequence of CD samples used to *identify* latent context j . It also initializes $\mathcal{T}_{t,j}$, the collection of time indices in $[t]$ in the Exp phases under estimated context j . Define

$$\mathcal{T}_t = \bigcup_{j \in [N]} \mathcal{T}_{t,j}, \quad T_t = |\mathcal{T}_t|, \quad T_{t,j} = |\mathcal{T}_{t,j}|, \quad \forall j \in [N]. \quad (3.1)$$

Algorithm 1 PIECEWISE-STATIONARY ε -BEST ARM IDENTIFICATION (PS ε BAI)

<p>1: Input: arm set \mathcal{X}, size of the set of latent vectors N, bounds on the segment lengths L_{\min} and L_{\max}, slackness parameter ε, confidence parameter δ, sampling parameter γ and window size w, threshold b.</p> <p>2: Initialize: Compute the G-optimal allocation λ^* and $\tau^* = \frac{38400 \ln(80) N L_{\max}}{\varepsilon^2} \ln \frac{N^2 K L_{\max}}{\delta \varepsilon^2}$.</p> <p>3: Set CD_{sample} = [], CA_{id} = { }. Set $t_{\text{CD}} = +\infty$.</p> <p>4: Set $\mathcal{T}_{t,j} = \emptyset$ and initialize $\mathcal{T}_t, T_{t,j}, T_t$ with (3.1) for all $t \leq \tau^*, j \in [N]$.</p> <p>5: Sample $\frac{w}{2}$ arms $\{x_s\}_{s=1}^{\frac{w}{2}} \sim \lambda^*$ and observe the associated returns $\{Y_{s,x_s}\}_{s=1}^{\frac{w}{2}}, t = \frac{w}{2}, t_{\text{CA}} = \frac{w}{2}$.</p> <p>6: CA_{id} = $\{1 : [(x_s, Y_{s,x_s})]_{s=1}^{\frac{w}{2}}\}, \hat{j}_t = 1$.</p> <p>7: while $t \leq \tau^*$ do</p> <p>8: $t = t + 1$</p> <p>9: Sample an arm $x_t \sim \lambda^*$ and observe return Y_{t,x_t}.</p> <p>10: if $\text{mod}(t - t_{\text{CA}}, \gamma) \neq 0$ then</p> <p>11: Update $\hat{j}_t = \hat{j}_{t-1}, \mathcal{T}_{t,\hat{j}_t} = \mathcal{T}_{t-1,\hat{j}_t} \cup \{t\}, \mathcal{T}_{t,j} = \mathcal{T}_{t-1,j}$ for $j \neq \hat{j}_t$.</p> <p>12: else</p> <p>13: CD_{sample} = CD_{sample} + $[(x_t, Y_{t,x_t})]$.</p>	<p>14: Update $\hat{j}_t = \hat{j}_{t-1}, \mathcal{T}_{t,j} = \mathcal{T}_{t-1,j}$ for all $j \in [N]$.</p> <p>15: if $\text{CD}_{\text{sample}} \geq w$ then</p> <p>16: if LCD($\mathcal{X}, w, b, \text{CD}_{\text{sample}}[-w :]$) then</p> <p>17: CD_{sample} = [].</p> <p>18: $t = t + \frac{w}{2}, t_{\text{CA}} = t, t_{\text{CD}} = +\infty$.</p> <p>19: $\hat{j}_t, \text{CA}_{\text{id}} = \text{LCA}(\mathcal{X}, w, b, \text{CA}_{\text{id}})$.</p> <p>20: if $\hat{j}_t = N + 1$ then break.</p> <p>21: Revert $\mathcal{T}_{t,j} = \mathcal{T}_{t - \frac{w(\gamma+1)}{2}, j}$ for all $j \in [N]$.</p> <p>22: end if</p> <p>23: end if</p> <p>24: end if</p> <p>25: Update the estimates using (3.1), (3.2) and (3.3).</p> <p>26: if Condition (3.4) is met and $t_{\text{CD}} = +\infty$ then</p> <p>27: Record $\hat{x}_\varepsilon = \arg \max_{x \in \mathcal{X}} x^\top \hat{\Theta}_t \hat{\mathbf{p}}_t$.</p> <p>28: $t_{\text{CD}} = \text{CD}_{\text{sample}}$.</p> <p>29: else if $t_{\text{CD}} = \text{CD}_{\text{sample}} - \frac{w}{2}$ then</p> <p>30: Recommend arm \hat{x}_ε.</p> <p>31: break</p> <p>32: end if</p> <p>33: end while</p>
--	---

It then collects $\frac{w}{2}$ samples and stores them in CA_{id}, which is then used to identify the first latent context (Lines 5 to 6).

In the Exp phase, PS ε BAI **firstly** samples an arm x_t with λ^* and observes the return $Y_{t,x_t} = x_t^\top \theta_{t,\hat{j}_t}^* + \eta_t$ (Line 9). It **then** updates the estimated context index and time collectors (Line 11). It also updates the estimates of value and probability of each context j (Line 25) with

$$\hat{\theta}_{t,j} = \frac{1}{T_{t,j}} \sum_{s \in \mathcal{T}_{t,j}} A(\lambda^*)^{-1} x_s Y_{s,x_s} \quad \text{and} \quad \hat{p}_{t,j} = \frac{T_{t,j}}{T_t}, \quad (3.2)$$

¹The empirical latent vector-probability pairs $[(\hat{\theta}_{t,j}, \hat{p}_{t,j})]_{j=1}^N$ can only approximate the unknown pairs $[(\theta_{\sigma(j)}^*, p_{\sigma(j)})]_{j=1}^N$ up to a permutation $\sigma : [N] \rightarrow [N]$, which is determined by the occurrence order of latent vectors. Thus we assume the latent vectors appear in order of increasing indices.

²A sample in CD_{sample} is a CD sample; A dictionary has a pairing structure {key:value} and dictionary[key] = value. $[a_i]_{i=1}^n$ denotes a sequence of elements a_1, \dots, a_n .

where $\hat{\theta}_{t,j} = \mathbf{0}$ if $T_{t,j} = 0$. We define α_t , ξ_t , $\beta_{t,j}$, and $\hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x})$ in Appendix D.2.1. For each pairs of arms (x, \tilde{x}) , the confidence radius of $\Delta(x, \tilde{x})$ at time step t is

$$\rho_t(x, \tilde{x}) := 2(\alpha_t + \xi_t) + \sum_{j=1}^N \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) + \zeta_t(x, \tilde{x})|; \quad (3.3)$$

PS ε BAI actively enters the CD phase every γ time steps (Line 12). It **firstly** adds a CD sample to $\text{CD}_{\text{sample}}$ (Line 13). **Next**, if there are sufficient CD samples (Line 15), the LCD subroutine (presented in Algorithm 3) is called and utilizes the most recent w CD samples to check whether a changepoint just occurred (Line 16). PS ε BAI steps into the CA phase if a changepoint is detected, and skips the CA phase otherwise, which is illustrated by Figures 1(b) and 1(a) respectively.

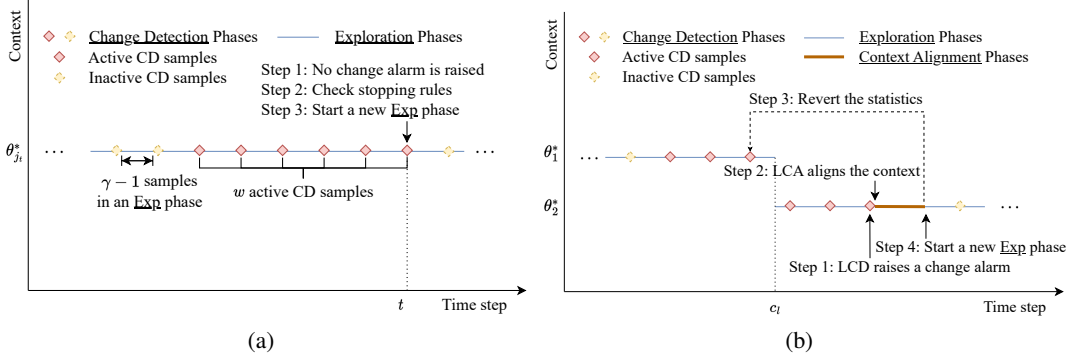


Figure 1: (a) No change alarm is raised during a stationary segment. The active CD samples are the input to the LCD subroutine at current time step t . (b) A changepoint is detected by LCD, followed by a CA phase and a statistics reversion step.

In the CA phase, PS ε BAI **starts** by resetting $\text{CD}_{\text{sample}}$ (Line 17), updating the count of time steps and recording the ending time of this CA phase (Line 18). **Thereafter**, the CA subroutine (presented in Algorithm 4) is invoked, which estimates the current latent context index \hat{j}_t and updates CA_{id} (Line 19). If $\hat{j}_t = N + 1$, i.e., PS ε BAI identifies $N + 1$ latent contexts, which is incorrect under instance Λ , it terminates and fails to identify an ε -optimal arm (Line 20). **Lastly**, all empirical statistics are reverted to those from $(w(\gamma + 1)/2)$ time steps ago, i.e., the most recent $(w\gamma/2)$ samples in the Exp phases are abandoned (Line 21).

The stopping rule is described in Lines 26 to 32. **(I)** If the following condition is satisfied (Line 26):

$$\min_{x: x \neq x_t^*} \hat{\Delta}_t(x_t^*, x) - \rho_t(x_t^*, x) \geq -\varepsilon \quad \text{and} \quad T_t \geq \frac{2L_{\max}}{9} \ln \left(\frac{2}{\delta_{d,T_t}} \right) \quad (3.4)$$

where the empirical mean gap $\hat{\Delta}_t(x_t^*, x) := (x_t^* - x)^\top \hat{\Theta}_t \hat{\mathbf{p}}_t$ and $x_t^* := \arg \max_{x \in \mathcal{X}} x^\top \hat{\Theta}_t \hat{\mathbf{p}}_t$, PS ε BAI records the arm with the highest empirical mean as \hat{x}_ε and the number of CD samples t_{CD} (Lines 27 and 28) but does not terminate immediately. Besides, a mild forced arm pull procedure is in the second line of (3.4), which is inspired by Lemma E.1 and to ensure the performance of PS ε BAI. **(II)** PS ε BAI will execute for another $(w\gamma/2)$ time steps in which $w/2$ CD samples are collected; if no changepoint is detected with these $w/2$ CD samples, the recorded arm \hat{x}_ε is recommended and PS ε BAI terminates (Lines 29 to 30). Part **(II)** of the stopping rule assures PS ε BAI does not terminate when a changepoint has occurred but has not been detected, as PS ε BAI may fail to identify an ε -optimal arm otherwise.

We remark that even though PS ε BAI uses the knowledge of L_{\max} , our experiments show that the performance of PS ε BAI is robust to small misspecifications in L_{\max} (see Appendix O.2). Furthermore, the computational complexity of PS ε BAI is computed in detail in Appendix D.4. The derived computational complexity indicates the proposed algorithm depends in a natural manner on the problem parameters such as d , K , N , and γ . Lastly, thanks to the LCD and LCA subroutines, a slightly modified variant of PS ε BAI can also solve the “ ε -Best Arm Tuple identification problem”, which aims to identify an ε -best arm *under each context*; see Appendix Q for details.

3.2.1 Theoretical guarantee of $\text{PS}_\varepsilon\text{BAI}$

To facilitate the analysis of $\text{PS}_\varepsilon\text{BAI}$, we propose the following assumptions. Note that our $\text{PS}_\varepsilon\text{BAI}$ algorithm may still succeed to identify an ε -optimal arm w.h.p. when the assumptions do not hold.

Assumption 1 (Distinguishability Condition). *The agent can choose w, γ and b such that (1) $2b \leq \Delta_c$ where $\Delta_c := \min_{\theta_j^* \neq \theta_j^*} \max_{x \in \mathcal{X}} |x^\top (\theta_j^* - \theta_j^*)|$ is the minimum gap between two contexts; and (2) $3w\gamma \leq L_{\min}$. A possible choice is*

$$b = \frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}} + \sqrt{\left(\frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}}\right)^2 + \frac{24d}{w} \ln \frac{2}{\delta_{\text{FAE}}}} \quad \text{where} \quad \delta_{\text{FAE}} = \frac{\gamma\delta}{4(\tau^*)^2 K}. \quad (3.5)$$

This assumption guarantees (i) $\text{PS}_\varepsilon\text{BAI}$ will not abandon all samples during the reversion procedure (Line 21 of Algorithm 1); (ii) each two latent vectors can be distinguished if the window size w is sufficiently large (e.g., $L_{\min}/6$). We clarify that this assumption is only for the rigor of theoretical guarantees and it holds provided that each stationary segment is sufficiently long; this is a feature of PSB models and similar assumptions are also present in existing works for their analyses [21, 10, 22]. We demonstrate the robustness of $\text{PS}_\varepsilon\text{BAI}$ to these parameters using experiments in Section 6.

Theorem 3.2. *Define the context distribution estimation (DE) hardness parameter*

$$H_{\text{DE}}(x_\varepsilon, x) := \frac{L_{\max}}{(\Delta(x^*, x) + \varepsilon)^2} \bar{H}(x_\varepsilon, x)$$

where $\bar{H}(x_\varepsilon, x) := \left(\sum_{j=1}^N \sqrt{\min\{16p_j, 1/4\}} |\Delta_j(x_\varepsilon, x) + \varepsilon|\right)^2$. Under Assumption 1, with probability at least $1 - \delta$, $\text{PS}_\varepsilon\text{BAI}$ identifies an ε -optimal arm and its sample complexity is

$$\tilde{O} \left(\underbrace{\max_{\substack{x_\varepsilon \in \mathcal{X}_\varepsilon \\ x \neq x_\varepsilon, x^*}} \frac{d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{1}{\delta}}_{T_V(x)} + \underbrace{H_{\text{DE}}(x_\varepsilon, x) \ln \frac{1}{\delta}}_{T_D(x_\varepsilon, x)} + \underbrace{\frac{NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{1}{\delta}}_{T_R(x)} \right). \quad (3.6)$$

The upper bound comprises three terms which serve distinct purposes:

(i) *Latent vector estimation (VE)*: $\tilde{O}(T_V(x))$ quantifies the bulk of samples needed to obtain a good estimate of latent context vectors such that the returns of x_ε and x can be distinguished, where x_ε is an ε -best arm and $x \notin \mathcal{X}_\varepsilon$ is a suboptimal arm. $T_V(x)$ recovers the sample complexity in the stationary linear bandits in [1], indicating that $\text{PS}_\varepsilon\text{BAI}$ estimates latent vectors efficiently.

(ii) *Context distribution estimation (DE)*: $\tilde{O}(T_D(x_\varepsilon, x))$ characterizes the bulk of samples needed to learn the distribution of latent context vectors.

(iii) *Residual estimation (RE)*: $\tilde{O}(T_R(x))$ counts the remaining samples needed for VE and DE, in addition to $\tilde{O}(T_V + T_D)$.

Besides, the max operator is applied to exclude all suboptimal arms. We also see that $T_V(x)$ and $T_D(x_\varepsilon, x)$ are similar to T_V^N and T_D^N in Proposition 3.1 respectively.

Firstly, the bound in (3.6) implies that, in an instance with smaller *relaxed mean gap* $\Delta(x^*, x) + \varepsilon$, $\text{PS}_\varepsilon\text{BAI}$ terminates after a larger number of time steps; in other words, it is more difficult to identify an ε -optimal arm. In difficult instances with small $\Delta(x^*, x) + \varepsilon$, the different orders of this term in $T_V(x)$, $T_D(x_\varepsilon, x)$ and $T_R(x)$ indicate that, $T_R(x)$ is small compared to $T_V(x)$ and $T_D(x_\varepsilon, x)$.

Secondly, DE solely utilizes context samples generated with P_θ and they are generated *only* at changepoints in \mathcal{C} , while all the observations in Exp phases facilitate VE. From this perspective, there are less samples that can be used for DE than for VE as $\text{PS}_\varepsilon\text{BAI}$ processes, and hence $T_D(x_\varepsilon, x)$ is supposed to be with larger order than $T_V(x)$.

Moreover, for the purpose of DE, $\text{PS}_\varepsilon\text{BAI}$ needs to observe context samples at $\tilde{O}\left(\frac{\bar{H}(x_\varepsilon, x)}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{1}{\delta}\right) = \tilde{O}\left(\frac{H_{\text{DE}}(x_\varepsilon, x)}{L_{\max}} \ln \frac{1}{\delta}\right)$ changepoints where L_{\max} is the maximum length of a stationary segment, leading us to $T_D(x_\varepsilon, x)$. Close examination of the definition of $\bar{H}(x_\varepsilon, x)$ reveals that *both* the vectors and their probabilities influence the number of samples needed for DE. The comparison between $T_D(x_\varepsilon, x)$ and T_D^N in Proposition 3.1 clearly indicates that $\text{PS}_\varepsilon\text{BAI}$ mitigates the influence of L_{\max} by detecting changepoints and aligning the detected context with observed ones, while $N_\varepsilon\text{BAI}$ does not do so.

3.3 $\text{PS}\varepsilon\text{BAI}^+ = \text{PS}\varepsilon\text{BAI} \cup \text{N}\varepsilon\text{BAI}$

We have provided a *high-probability* result for $\text{PS}\varepsilon\text{BAI}$ in Theorem 3.2. The design of $\text{PS}\varepsilon\text{BAI}$ (Line 7 of Algorithm 1) indicates that $\text{PS}\varepsilon\text{BAI}$ will not recommend any arm if it does not terminate at time τ^* . This result is nontrivial, as the high-probability result in Theorem 3.2 depends on the success of change detection (Algorithm 3) and context alignment (Algorithm 4), which requires a non-vanishing failure probability (e.g., $\delta/2$). Thus, we cannot derive an upper bound on the expected sample complexity of $\text{PS}\varepsilon\text{BAI}$. We devise a solution by designing the Piecewise-Stationary E-Best Arm Identification⁺ ($\text{PS}\varepsilon\text{BAI}^+$) algorithm with a simple but effective trick.

The $\text{PS}\varepsilon\text{BAI}^+$ algorithm samples *one* arm with the G-optimal allocation λ^* at each time step, with which Algorithms 1 and 2 are executed in parallel (detailed in Algorithm 5). This is feasible since $\text{PS}\varepsilon\text{BAI}$ and $\text{N}\varepsilon\text{BAI}$ algorithms have the same sampling rule.

Theorem 3.3. *The $\text{PS}\varepsilon\text{BAI}^+$ algorithm is (ε, δ) -PAC and its expected sample complexity is*

$$\tilde{O}\left(\min\left\{\max_{x_\varepsilon \in \mathcal{X}_\varepsilon, x \neq x_\varepsilon, x^*} T_V(x) + T_D(x_\varepsilon, x) + T_R(x), T_V^N + T_D^N\right\}\right).$$

$\text{PS}\varepsilon\text{BAI}^+$ inherits the superiority of $\text{PS}\varepsilon\text{BAI}$ to adapt to the piecewise stationary environment, and employs the stopping rule of $\text{N}\varepsilon\text{BAI}$ to maintain a finite expected sample complexity. As a result, the expected complexity of $\text{PS}\varepsilon\text{BAI}^+$ in Theorem 3.3 is of the same order as the high-probability one of $\text{PS}\varepsilon\text{BAI}$ in Theorem 3.2 and is not larger than the complexity of $\text{N}\varepsilon\text{BAI}$ in Proposition 3.1. We show how our results particularize to the *stationary* linear bandits BAI problem, as well as additional discussions on the upper bound, in Appendix P.

4 Lower Bound on the Sample Complexity

Given $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$, define the alternative instance $\Lambda' = (\mathcal{X}, \Theta', P_{\theta'}, \mathcal{C})$ w.r.t. Λ , where $\Theta' = (\theta'_1, \dots, \theta'_n) \in \mathbb{R}^{d \times N}$, $P_{\theta'}[\theta'_j] = P_\theta[\theta_j^*]$, and there exists $x \in \mathcal{X} \setminus \mathcal{X}_\varepsilon$, such that $x_\varepsilon^\top \mathbb{E}_{\theta' \sim P_{\theta'}}[\theta'] < x^\top \mathbb{E}_{\theta' \sim P_{\theta'}}[\theta'] - \varepsilon$ for all $x_\varepsilon \in \mathcal{X}_\varepsilon$. Let $\text{Alt}_\Theta(\Lambda)$ be set of all alternative instances (w.r.t. Λ).

Theorem 4.1. *For all (ε, δ) -PAC algorithm π , there exists an instance $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$ such that*

$$\mathbb{E}[\tau_\pi] \geq \max\left\{T_\varepsilon(\Lambda) \ln \frac{1}{2.4\delta}, c_{N_C}\right\},$$

where

$$T_\varepsilon(\Lambda)^{-1} := \max_{\{v_j\}_{j=1}^N} \min_{\Lambda' \in \text{Alt}_\Theta(\Lambda)} \sum_{j,x} p_j v_{j,x} \frac{(x^\top (\theta_j^* - \theta'_j))^2}{2}, \quad \text{and}$$

$$N_C := \max_{x \neq x^*} \frac{\sum_j p_j (\Delta_j(x^*, x) + \varepsilon)^2}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{1}{4\delta}.$$

Recall that c_{N_C} is the N_C -th changepoint in the changepoint sequence \mathcal{C} , which is lower bounded by $N_C L_{\min}$ and is $N_C L_{\max}$ in the worst case. To derive the lower bound in Theorem 4.1, we investigate two environments different from the one defined in Section 2 (and as in Dynamics 1):

- Dynamics 2: the agent observes the index of current context j_t (i.e., contextual linear bandits);
- Dynamics 3: the agent observes the changepoints in \mathcal{C} and context vector $\theta_{j_t}^*$'s, and hence she solely needs to estimate the distribution of contexts.

We bound the sample complexity of an (ε, δ) -BAI algorithm in Dynamics 2 and 3 respectively, which when combined, yield the lower bound in Theorem 4.1; this is detailed in Appendix M.

Note that $T_\varepsilon(\Lambda)^{-1}$ in the lower bound generalizes [16] to the setting of linear bandits. In addition, Theorem 4.1 can be reduced to a bound in stationary linear bandits with one latent vector [24] (see the discussion leading to (M.15)).

5 On the Asymptotic Optimality of $\text{PS}\varepsilon\text{BAI}^+$

To illustrate the efficiency of our $\text{PS}\varepsilon\text{BAI}^+$ algorithm, we compare the upper bound on its expected sample complexity in Theorem 3.3 and the generic lower bound in Theorem 4.1 under specific instances below and in Appendix N. We also gain further insight into our $\text{PS}\varepsilon\text{BAI}^+$ algorithm.

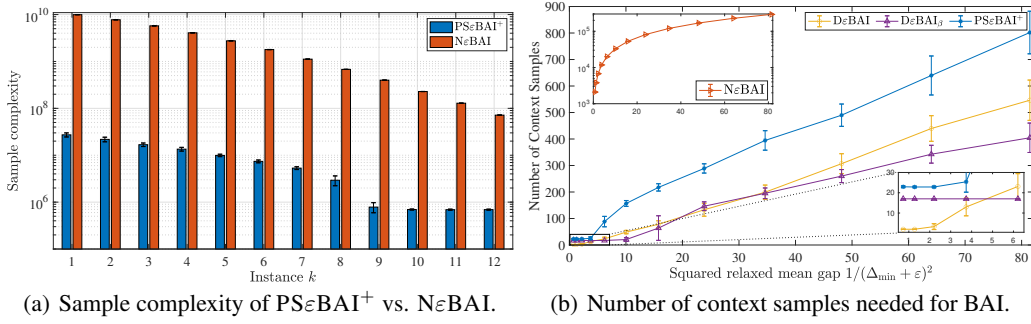


Figure 2: Experimental results

Example 1. Instance $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$ is with (i) $2d - 1$ arms: $x_{(1)} = \mathbf{e}_1, x_{(i)} = \mathbf{e}_i, x_{(d+i-1)} = \mathbf{e}_1 \cos \phi + \mathbf{e}_i \sin \phi$ for all $i \in \{2, \dots, d\}$ where $\phi \in [0, \pi/4)$, (ii) $2d - 2$ contexts: $\theta_{j\pm}^* = \mathbf{e}_1 \cos \phi \pm \mathbf{e}_{j+1} \sin \phi$ for all $j \in [d - 1]$, (iii) Context distribution: $p_j = 1/N$ for all $j \in [N]$.

Under the instance defined in Example 1, $x_{(i)}$ for all $i \neq 1$ is inferior to $x_{(1)}$ under all contexts and $x_{(i+d)}$ for all $i \in [d - 1]$ is marginally better than $x_{(1)}$ by $1 - \cos \phi$ only under context θ_{i+}^* and $\Delta(x_{(1)}, x_{(i+d)}) = \cos \phi - \cos^2 \phi$. We expect PSεBAI⁺ to discover this feature of the instances and quickly identify an ε -optimal arm with a coarse estimation of the context distribution.

Corollary 5.1. For the instance defined in Example 1, we have $H_{\text{DE}}(x_\varepsilon, x) = \tilde{O}(NL_{\max})$ for all $(x_\varepsilon, x) \in \mathcal{X}_\varepsilon \times (\mathcal{X} \setminus \mathcal{X}_\varepsilon)$. In addition, if $\varepsilon < (\cos \phi)(1 - \cos \phi)$, we have

$$\frac{\mathbb{E}[\tau]^*}{\ln(1/\delta)} \in \tilde{\Theta} \left((1 + f(\phi)) \cdot \frac{d}{(\Delta_{x_{(1)}, x_{(d+1)}} + \varepsilon)^2} \right), \quad (5.1)$$

where $\mathbb{E}[\tau]^*$ is the minimal expected sample complexity over all (ε, δ) -PAC algorithms and $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfies $f(\phi) \rightarrow 0$ as $\phi \rightarrow 0^+$. The upper bound in (5.1) is achieved by PSεBAI⁺.

The order of H_{DE} in Corollary 5.1 indicates that $T_{\text{D}}(x_{(1)}, x_{(d+1)}) = \tilde{O}(NL_{\max} \ln(1/\delta))$ is not dominating the sample complexity of PSεBAI⁺, suggesting that a coarse estimation of the context distribution is sufficient when ϕ is small. In other words, PSεBAI⁺ exploits the feature of instances and utilize samples mostly for estimate context vectors, which is again expected.

Corollary 5.1 implies that under such instances, the upper and lower bounds on the sample complexity of PSεBAI⁺ match up to logarithmic factors, that is, the performance PSεBAI⁺ is near optimal. Besides, the bound of NεBAI in Proposition 3.1 is with an extra additive term L_{\max} compared to the lower bound in Corollary 5.1, illustrating that NεBAI is suboptimal and again emphasizing the significance of detecting changes and aligning contexts for PSεBAI⁺ to reduce the impact of L_{\max} .

6 Numerical Experiments

We now evaluate the empirical performance of PSεBAI⁺. We utilize the instance defined in Example 1 with $d = 2, \phi = \pi/8$. We generate a changepoint sequence \mathcal{C} such that $c_{l+1} = c_l + L_l$ with $L_{\min} = 3 \times 10^4, L_{\max} = 5 \times 10^4, \mathbb{P}[L_l = L_{\min}] = 0.8, \mathbb{P}[L_l = L_{\max}] = 0.2$, and fix it throughout the whole set of experiments. We set the confidence parameter $\delta = 0.05$ and vary the slackness parameter ε from 0.04 to 0.6 (i.e., $\varepsilon = 0.03 \times 1.35^k$ for $k \in [12]$). We set $\gamma = 6$, the window size $w = L_{\min}/(3\gamma)$ and compute b via (3.5) in Assumption 1.³ For each choice of algorithm and instance, we run 20 independent trials. All the code to reproduce our experiments can be found at <https://github.com/Y-Hou/BAI-in-PSLB.git>.

We **first** compare PSεBAI⁺ and NεBAI. Both algorithms succeed to identify an ε -optimal arm, while empirically, the complexity of PSεBAI⁺ is $\leq 1\%$ of that of NεBAI. The empirical averages and standard deviations of the sample complexities of both algorithms are presented in Figure 2(a).

³We clarify that (3.5) in Assumption 1 is only for our theoretical guarantees. In practice, our algorithm has shown robustness w.r.t. the parameters and we can safely neglect the constants in the formula.

Figure 2(a) illustrates that empirically, the termination and arm recommendation of $\text{PS}\varepsilon\text{BAI}^+$ are determined by the execution of $\text{PS}\varepsilon\text{BAI}$ as a subroutine, suggesting that in Theorem 3.3, the first term resulting from $\text{PS}\varepsilon\text{BAI}$ actually determines the complexity of $\text{PS}\varepsilon\text{BAI}^+$.

Next, we test the efficacy of $\text{PS}\varepsilon\text{BAI}^+$ to learn and exploit the latent vectors and the distribution of contexts. Specifically, we run $\text{PS}\varepsilon\text{BAI}^+$ and $\text{N}\varepsilon\text{BAI}$ under Dynamics 1 where neither the index nor the vector of current context is visible to the agent. We also run benchmark algorithms: $\text{D}\varepsilon\text{BAI}$ and its variant $\text{D}\varepsilon\text{BAI}_\beta$ under Dynamics 3 where context vectors and changepoints are all observed; these two algorithms are detailed and analyzed in Appendix O.4.

As the changepoint sequence \mathcal{C} is fixed in a given instance and $t/L_{\max} \leq l_t \leq t/L_{\min}$ for all $t \in \mathbb{N}$, we regard the number of context samples l_τ as a proxy of the sample complexity τ . We present the number of context samples need by $\text{PS}\varepsilon\text{BAI}^+$, $\text{N}\varepsilon\text{BAI}$, $\text{D}\varepsilon\text{BAI}$ and $\text{D}\varepsilon\text{BAI}_\beta$ for arm identification w.r.t. $1/(\Delta_{\min} + \varepsilon)^2$ in Figure 2(b).

Figure 2(b) contains three messages. First, the complexity of $\text{PS}\varepsilon\text{BAI}^+$ scales as $1/(\Delta_{\min} + \varepsilon)^2$, corroborating Theorem 3.3. Second, although $\text{PS}\varepsilon\text{BAI}^+$ has access to neither context vectors nor changepoints, it needs roughly the same number of context samples as $\text{D}\varepsilon\text{BAI}$ and $\text{D}\varepsilon\text{BAI}_\beta$, suggesting that it is competitive compared to these algorithms that have oracle information about the environment. Third, $\text{D}\varepsilon\text{BAI}_\beta$ uses the confidence radius in (3.3) and terminates with fewer context samples compared to $\text{D}\varepsilon\text{BAI}$, implying that the confidence radius is well-designed.

Furthermore, when ε decreases from 0.03×1.35^{12} to 0.03×1.35^9 , the complexity of $\text{PS}\varepsilon\text{BAI}$ almost remains unchanged while that of $\text{N}\varepsilon\text{BAI}$ increases rapidly as presented in instances 9 to 12 in Figure 2(a). Meanwhile, the number of context samples need by two algorithms are shown to be with the same pattern in Figure 2(b). This contrast indicates that the cost of distribution estimation (T_D in (3.6)) for $\text{PS}\varepsilon\text{BAI}^+$ has been significantly minimized compared to $\text{N}\varepsilon\text{BAI}$.

To summarize, we emphasize that the empirical superiority of $\text{PS}\varepsilon\text{BAI}^+$ over $\text{N}\varepsilon\text{BAI}$ implies that *the efficacy of $\text{PS}\varepsilon\text{BAI}^+$ is inherited from $\text{PS}\varepsilon\text{BAI}$* . Our experiments show that actively exploiting the context information, via changepoint detection and context alignment (as in $\text{PS}\varepsilon\text{BAI}^+$ and $\text{PS}\varepsilon\text{BAI}$) facilitates identifying the ε -optimal arm efficiently.

Similar to many existing algorithms in piecewise-stationary bandits [21, 10, 22], our algorithm requires Assumption 1 and the knowledge of L_{\max} . These may not be available in practice. Thus, we conduct more experiments in Appendix O.2 to exhibit the robustness of $\text{PS}\varepsilon\text{BAI}^+$. Specifically, in Appendix O.2, we conduct experiments for the case in which L_{\max} is *misspecified*. In Appendix O.3, we alter the change detection frequency γ so that w and b change accordingly. In both sets of experiments, the overall sample complexity of $\text{PS}\varepsilon\text{BAI}^+$ does not vary significantly and retains its superiority over $\text{N}\varepsilon\text{BAI}$. We conclude that $\text{PS}\varepsilon\text{BAI}^+$ is robust to slight misspecifications in these parameters, as long as Assumption 1 is not severely violated. Please refer to Appendix O for further details and experiments.

7 Conclusion and Future Work

We proposed a novel PSLB model and designed the $\text{PS}\varepsilon\text{BAI}^+$ algorithm to identify an ε -optimal arm with probability $\geq 1 - \delta$. The efficacy of $\text{PS}\varepsilon\text{BAI}^+$ has been demonstrated both empirically and theoretically. We argued that this is due to the embedded change detection and context alignment procedures. There are several directions for further exploration.

Firstly, our $\text{PS}\varepsilon\text{BAI}^+$ algorithm provides a fairly general *framework* for algorithm design. For instance, in addition to utilization the G-optimal allocation to sample arms as in $\text{PS}\varepsilon\text{BAI}^+$, the $\mathcal{X}\mathcal{Y}$ -allocation and adaptive $\mathcal{X}\mathcal{Y}$ -allocation [1] can also be considered. In other words, our $\text{PS}\varepsilon\text{BAI}^+$ algorithm can be generalized to form an entire class of algorithms for BAI in PSLB models. In addition, deriving instance-dependent guarantees is also of great interest.

Secondly, most of the literature on piecewise-stationary bandits [21, 10, 22] make assumptions to provide theoretical guarantees. It would be interesting to remove or reduce these assumptions under our ε -BAI problem setup, and yet still be able to provide similar theoretical guarantees.

Finally, we believe that it is possible to adapt our $\text{PS}\varepsilon\text{BAI}^+$ algorithm to the fixed-budget setting, i.e., to identify an ε -optimal arm with high probability in a fixed time horizon in PSLB models.

Acknowledgements: This work is funded by the Singapore Ministry of Education AcRF Tier 2 grant (A-8000423-00-00) and Tier 1 grants (A-8000189-01-00 and A-8000980-00-00).

References

- [1] Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. In *Proceedings of the 27th Advances in Neural Information Processing Systems*, volume 27, pages 828–836, 2014.
- [2] Junwen Yang and Vincent Tan. Minimax optimal fixed-budget best arm identification in linear bandits. *Proceedings of the 36th Advances in Neural Information Processing Systems*, 35:12253–12266, 2022.
- [3] Yasin Abbasi-yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Proceedings of the 24th Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011.
- [4] Marc Abeille and Alessandro Lazaric. Linear thompson sampling revisited. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 176–184. PMLR, 2017.
- [5] Tanner Fiez, Lalit Jain, Kevin G Jamieson, and Lillian Ratliff. Sequential experimental design for transductive linear bandits. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Proceedings of the 32nd Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [6] Marc Jourdan, Rémy Degenne, and Emilie Kaufmann. An ϵ -best-arm identification algorithm for fixed-confidence and beyond, November 2023. arXiv:2305.16041.
- [7] Rémy Degenne and Wouter M Koolen. Pure exploration with multiple correct answers. *Proceedings of the 32nd Advances in Neural Information Processing Systems*, 32, 2019.
- [8] Omar Besbes, Yonatan Gur, and Assaf Zeevi. Stochastic multi-armed-bandit problem with non-stationary rewards. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Proceedings of the 27th Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [9] Aurélien Garivier and Eric Moulines. On upper-confidence bound policies for switching bandit problems. In Jyrki Kivinen, Csaba Szepesvári, Esko Ukkonen, and Thomas Zeugmann, editors, *The 22nd International conference on Algorithmic learning theory*, pages 174–188, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [10] Yang Cao, Zheng Wen, Branislav Kveton, and Yao Xie. Nearly optimal adaptive procedure with change detection for piecewise-stationary bandit. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, pages 418–427. PMLR, 2019.
- [11] Yueyang Liu, Xu Kuang, and Benjamin Van Roy. A definition of non-stationary bandits, Feb 2023. arXiv:2302.12202.
- [12] Wang Chi Cheung, David Simchi-Levi, and Ruihao Zhu. Learning to optimize under non-stationarity. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1079–1087. PMLR, 2019.
- [13] Jonathan Shapiro, Carlos M Carvalho, and Pradeep Ravikumar. Thompson sampling in switching environments with bayesian online change point detection. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics*, pages 442–450. Microtome Publishing, 2013.
- [14] Zhihan Xiong, Romain Camilleri, Maryam Fazel, Lalit Jain, and Kevin Jamieson. A/B testing and best-arm identification for linear bandits with robustness to non-stationarity. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics, PMLR 238*, pages 1585–1593. PMLR, 2023.

- [15] Robin Allesiardo, Raphaël Féraud, and Odalric-Ambrym Maillard. The non-stationary stochastic multi-armed bandit problem. *International Journal of Data Science and Analytics*, 3:267–283, 2017.
- [16] Masahiro Kato and Kaito Ariu. The role of contextual information in best arm identification, 2021. arXiv:2106.14077.
- [17] Yoan Russac, Christina Katsimerou, Dennis Bohle, Olivier Cappé, Aurélien Garivier, and Wouter M Koolen. A/B/n testing with control in the presence of subpopulations. *Proceedings of the 34th Advances in Neural Information Processing Systems*, 34:25100–25110, 2021.
- [18] Andrea Zanette, Kefan Dong, Jonathan N Lee, and Emma Brunskill. Design of experiments for stochastic contextual linear bandits. *Proceedings of the 34th Advances in Neural Information Processing Systems*, 34:22720–22731, 2021.
- [19] Yasin Abbasi-Yadkori, András György, and Nevena Lazić. A new look at dynamic regret for non-stationary stochastic bandits. *Journal of Machine Learning Research*, 24(288):1–37, 2023.
- [20] Jia Yuan Yu and Shie Mannor. Piecewise-stationary bandit problems with side observations. In *Proceedings of the 26th International Conference on Machine Learning, ICML '09*, page 1177–1184. Association for Computing Machinery, 2009.
- [21] Fang Liu, Joohyun Lee, and Ness Shroff. A change-detection based framework for piecewise-stationary multi-armed bandit problem. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [22] Lilian Besson, Emilie Kaufmann, Odalric-Ambrym Maillard, and Julien Seznec. Efficient change-point detection for tackling piecewise-stationary bandits. *The Journal of Machine Learning Research*, 23(1):3337–3376, 2022.
- [23] Peter Auer, Pratik Gajane, and Ronald Ortner. Adaptively tracking the best bandit arm with an unknown number of distribution changes. In *Proceedings of the 32nd Conference on Learning Theory*, pages 138–158. PMLR, 2019.
- [24] Yassir Jedra and Alexandre Proutiere. Optimal best-arm identification in linear bandits. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Proceedings of the 34th Advances in Neural Information Processing Systems*, volume 33, pages 10007–10017. Curran Associates, Inc., 2020.
- [25] Peng Zhao, Lijun Zhang, Yuan Jiang, and Zhi-Hua Zhou. A simple approach for non-stationary linear bandits. In *The 23rd International Conference on Artificial Intelligence and Statistics*, pages 746–755. PMLR, 2020.
- [26] Masahiro Kato, Masaaki Imaizumi, Takuya Ishihara, and Toru Kitagawa. Asymptotically optimal fixed-budget best arm identification with variance-dependent bounds, 2023. arXiv:2302.02988.
- [27] Chao Qin and Daniel Russo. On adaptivity and confounding in contextual bandit experiments. In *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2021.
- [28] Yasin Abbasi-Yadkori, Peter Bartlett, Victor Gabillon, Alan Malek, and Michal Valko. Best of both worlds: Stochastic & adversarial best-arm identification. In *Proceedings of the 31st Conference on Learning Theory*, pages 918–949. PMLR, 2018.
- [29] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.
- [30] Kwang-Sung Jun, Kevin G Jamieson, Robert D Nowak, and Xiaojin Zhu. Top arm identification in multi-armed bandits with batch arm pulls. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 139–148, 2016.
- [31] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17:1–42, 2016.

- [32] Marc Jourdan and Rémy Degenne. Choosing answers in epsilon-best-answer identification for linear bandits. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 10384–10430. PMLR, 17–23 Jul 2022.
- [33] Michael J. Todd. *Minimum-Volume Ellipsoids: Theory and Algorithms*. SIAM, 2016.

Appendices

The contents of the appendices are organized as follows:

- In Appendix A, we provide further motivating examples for our problem.
- In Appendix B, we discuss the limitations of our method.
- In Appendix C, we review more related works on drifting and contextual bandits.
- In Appendix D, we provide more details about our algorithms
 - Appendix D.1: pseudo-code of $N\varepsilon$ BAI in Algorithm 2.
 - Appendix D.2: more details of $PS\varepsilon$ BAI including the precise definition of the confidence radius ρ_t and details of LCD and LCA subroutines.
 - Appendix D.3: pseudo-code of $PS\varepsilon$ BAI⁺ in Algorithm 5.
 - Appendix D.4: the computational complexity of $PS\varepsilon$ BAI in Algorithm 1.
- In Appendix E: we provide a useful lemma for estimating the expected return of any arm in linear bandits when the sampling rule is according to the G-optimal allocation.
- In Appendix F: we prove the upper bound on the expected sample complexity of $N\varepsilon$ BAI.
- In Appendices G to K: we detail the analysis of $PS\varepsilon$ BAI, i.e., we provide the proof of the upper bound on its complexity in Theorem 3.2.
 - Appendix G: outline of the proof.
 - Appendix H: analysis of the Change Detection (CD) and Context Alignment (CA) procedures.
 - Appendix I: analysis of the estimation error by decomposing it into three terms: Vector-Estimation Error (VE), Distribution-Estimation Error (DE), and Residual Estimation Error (RE).
 - Appendix J: proof of Theorem 3.2 based on the analysis above.
 - Appendix K: proof of technical lemmas that are utilized in the analysis of $PS\varepsilon$ BAI.
- In Appendix L: we prove the upper bound on the expected sample complexity of $PS\varepsilon$ BAI⁺ based on the analysis of $N\varepsilon$ BAI and $PS\varepsilon$ BAI.
- In Appendix M: we derive the lower bound on the expected sample complexity of any algorithm, i.e., proof Theorem 4.1.
- In Appendix N: we provide more examples to compare the derived upper and lower bounds on the expected sample complexity, and illustrate the efficacy of our $PS\varepsilon$ BAI⁺ algorithm.
- In Appendix O: we provide more details of numerical experiments.
- In Appendix P: we provide more discussions on
 - the related methods for BAI in nonstationary bandits,
 - the instance-dependent upper bound,
 - the connection between the piecewise-stationary linear bandits model to the stationary linear bandits model,
 - the special case where $N = 1$.
- In Appendix Q: we provide analytical results on the “Best Arm Tuple Identification Problem”.

A Further Motivating Examples

We elaborate on the some concrete real-life examples that motivate our problem setup of identifying the ensemble best arm in piecewise-stationary linear bandits.

In scenarios such as investment option selection and portfolio management also mentioned by [10, 20], there is a multitude of options for fund managers to choose from and typically, they want to find, in the initial pure exploration process, a small subset of candidate portfolios (or even the “best” portfolio) based on various economic indicators and the market performance of individual stocks before further exploitation. In a bearish market, more portfolios tend to incur losses; while in a bullish market, more portfolios tend to generate gains. The transition between these two contexts can be effected by stochastic factors, e.g., the weather, or the outbreak of a pandemic, making the market conditions (contexts) stochastic. In the face of these uncertainties (in the contexts and rewards), we wish to design and analyze algorithms that selected portfolio to yield the best long-term option under such a piecewise-stationary environment.

Crop rotation is another example. Since crop yields can be influenced by various factors, such as weather conditions (analogous to our stochastically generated contexts), selecting the most suitable crop to grow and harvest from is crucial. Given several candidate crops, crops of similar types (e.g., potatoes and sweet potatoes) are correlated as they tend to favor similar conditions, thus, they can be

modelled by bandits with a linear reward structure. Contextual factors, like weather conditions, are well-modelled as being stochastic. A fixed weather condition will last a period of time and it will not change suddenly. Domain knowledge from historical data/records provides us with prior knowledge on L_{\min} and L_{\max} . These observations dovetail with our model. Our objective is to choose the crop that offers the near-highest yield potential over a long time period (an ensemble ε -best arm) and is adaptable to local environment factors, such as weather patterns.

B Limitations

Similar to existing works on piecewise stationary bandits [21, 10, 22], we introduced some assumptions to provide theoretical guarantees for $\text{PS}\varepsilon\text{BAI}^+$ although $\text{PS}\varepsilon\text{BAI}^+$ has shown to be robust even in the absence of these assumptions. Since an algorithm may not need to differentiate two contexts with close latent vectors for identifying an ε -optimal arm, we surmise it is possible to weaken these assumptions. For this purpose, we will consider clustering the contexts into a few classes based on the distances between their latent vectors and design an algorithm that only aims to detect the change of context class, instead of the change of context.

C More Related Work

In **drifting bandits** (DB), the regrets of algorithms are affected by the level of nonstationarity of the environment, which can be measured by various quantities, such as the total variation of the context sequence and the number of time steps when the return of at least one arm changes [23, 25].

In **contextual bandits** (CB), where the contextual information is visible to the agent, [16, 26, 17, 27] aim to identify the best arm with the assumption that the context changes at every time step according to a fixed distribution, and the return of an arm is averaged across all contexts. We see that the context distribution is involved to measure the quality of arms. However, while the agent in CB models can observe the context information, the agent has no access to the contexts but still aims to identify an ε -optimal arm in our piecewise stationary linear bandit (PSLB) model.

In **adversarial bandits**, existing works pertaining to the BAI problem only explored the fixed-horizon setting [28]. Due to the difference between the fixed-confidence and fixed-horizon setting and the difference between adversarial and piecewise stationary bandits, their results cannot be trivially extended to solve the fixed-confidence BAI problem in piecewise stationary bandits.

D More Details of Algorithms $\text{N}\varepsilon\text{BAI}$, $\text{PS}\varepsilon\text{BAI}$ and $\text{PS}\varepsilon\text{BAI}^+$

All proposed algorithms make use of the well known G-optimal allocation (design) [1, 29], which is widely used in the linear bandits literature. The G-optimal allocation minimizes the maximal mean-squared prediction error in all directions [1]. Given an arm set \mathcal{X} , the G-optimal allocation λ^* is a distribution over the arm set, which is the minimizer of $g(\lambda) = \max_{x \in \mathcal{X}} \|x\|_{A(\lambda)^{-1}}^2$, where $A(\lambda) = \sum_{x \in \mathcal{X}} \lambda(x)xx^\top$ and $\lambda \in \Delta_{\mathcal{X}}$. Interested readers may refer to Chapter 21 in [29].

D.1 The NAÏVE ε -BEST ARM IDENTIFICATION ($\text{N}\varepsilon\text{BAI}$) Algorithm

As presented in Algorithm 2, $\text{N}\varepsilon\text{BAI}$ samples an arm with the G-optimal allocation at each time steps; its stopping rule is grounded on the property of G-optimal allocation (see Lemma E.1) and affected by the maximum length of a single stationary segment L_{\max} .

Algorithm 2 NAÏVE ε -BEST ARM IDENTIFICATION (N ε BAI)

- 1: **Input:** the arm set \mathcal{X} , the phase length bounds L_{\min}, L_{\max} , the slackness parameter ε , the confidence parameter δ .
- 2: **Initialize:** Compute the G-optimal allocation λ^* .
- 3: Compute $C_3 = \sum_{n=1}^{\infty} n^{-3}$ and $t^* = 3d \ln(6dKC_3/\delta)$.
- 4: Sample t^* arms $\{x_s\}_{s=1}^{t^*} \sim \lambda^*$ and observe the associated returns $\{Y_{s,x_s}\}_{s=1}^{t^*}$. Let $t = t^*$.
- 5: Compute

$$\begin{aligned}\tilde{\theta}_t &= \frac{1}{t} \sum_{s=1}^t A(\lambda^*)^{-1} x_s Y_{s,x_s}, \quad \hat{x}_t = \arg \max_{x \in \mathcal{X}} x^\top \tilde{\theta}_t, \\ \tilde{\rho}_t &= \sqrt{\frac{8L_{\max}}{t} \ln \frac{4KC_3 t^3}{\delta}} + 5\sqrt{\frac{d}{t} \ln \frac{4KC_3 t^3}{\delta}}.\end{aligned}\tag{D.1}$$

- 6: **while** $\hat{x}_t^\top \tilde{\theta}_t - \tilde{\rho}_t + \varepsilon < \max_{x \neq \hat{x}_t} x^\top \tilde{\theta}_t + \tilde{\rho}_t$ **do**
 - 7: Sample an arm $x_t \sim \lambda^*$ and observe return Y_{t,x_t} and let $t = t + 1$.
 - 8: Update $\tilde{\theta}_t, \hat{x}_t$ and $\tilde{\rho}_t$ with (D.1).
 - 9: **end while**
 - 10: Recommend arm $\hat{x}_\varepsilon = \hat{x}_t$.
-

D.2 Details about PS ε BAI

D.2.1 Confidence radius utilized in PS ε BAI

For each pair of arms (x, \tilde{x}) , the confidence radius of $\Delta(x, \tilde{x})$ at time step t is

$$\rho_t(x, \tilde{x}) := 2(\alpha_t + \xi_t) + \sum_{j=1}^N \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) + \zeta_t(x, \tilde{x})|,$$

where $a^{\text{clip}_2} := \min\{\max\{a, -2\}, 2\}$ denotes the value of a that is clipped to the interval $[-2, 2]$ and

$$\begin{aligned}\alpha_t &:= 5\sqrt{\frac{d}{T_t} \ln \frac{2}{\delta_{v,T_t}}}, \quad \xi_t := 25\sqrt{2} \frac{NL_{\max}}{T_t} \ln \frac{2}{\delta_{m,T_t}}, \\ \beta_{t,j} &:= \min \left\{ \frac{5}{2} \sqrt{\frac{2\phi_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}}, 1 \right\}, \\ \phi_{t,j} &:= \min \left\{ 4 \max \left\{ \hat{\rho}_{t,j}, \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \right\}, \frac{1}{4} \right\}, \\ \delta_{v,T_t} &= \frac{\delta}{15KT_t^3}, \quad \delta_{m,T_t} = \frac{\delta}{15KNT_t^3}, \quad \delta_{d,T_t} = \frac{\delta}{15NT_t^3}.\end{aligned}$$

$\zeta_t(x, \tilde{x})$ minimizes the last summation. In the final theoretical upper bound on the sample complexity, we take $\zeta_t(x, \tilde{x}) = \varepsilon$ for simplicity. In the experiment, we utilize

$$\zeta_t(x, \tilde{x}) = \arg \min_{\zeta_t(x, \tilde{x}) \in \mathbb{R}} \sum_{j=1}^N \beta_{t,j} (\hat{\Delta}_{t,j}(x, \tilde{x}) + \zeta_t(x, \tilde{x}))^2 = -\frac{\sum_{j=1}^N \beta_{t,j} \hat{\Delta}_{t,j}(x, \tilde{x})}{\sum_{j=1}^N \beta_{t,j}}$$

for a simple and effective analytic expression in the experiment.

D.2.2 Subroutines of PS ε BAI

In the LCD subroutine (Algorithm 3), two estimates $\tilde{\theta}_1$ and $\tilde{\theta}_2$ of the latent vectors are independently computed from the first and second halves of the input CD samples (Line 2). LCD only raises an alarm of changepoint (Line 4) when the difference between $\tilde{\theta}_1$ and $\tilde{\theta}_2$ is sufficiently large (Line 3) and indicates a changepoint occurs w.h.p.

The LCA subroutine (Algorithm 4) estimates the latent vector of current context and updates the dictionary CA_{id} . Specifically, it **firstly** samples $w/2$ samples with λ^* (Line 2). It **then** checks

Algorithm 3 LINEAR-CHANGE DETECTION (LCD)

- 1: **Input:** arm set \mathcal{X} , detection window w , threshold b and w arms with the associated observations $[(\tilde{x}_1, Y_{s, \tilde{x}_1}), \dots, (\tilde{x}_w, Y_{w, \tilde{x}_w})]$.
 - 2: Compute $\tilde{\theta}_1 = \frac{2}{w} \sum_{s=1}^{\frac{w}{2}} A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s}$ and $\tilde{\theta}_2 = \frac{2}{w} \sum_{s=\frac{w}{2}+1}^w A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s}$.
 - 3: **if** $\exists x \in \mathcal{X}, s.t. |x^\top (\tilde{\theta}_2 - \tilde{\theta}_1)| > b$ **then**
 - 4: **Return** True
 - 5: **else**
 - 6: **Return** False
 - 7: **end if**
-

Algorithm 4 LINEAR-CONTEXT ALIGNMENT (LCA)

- 1: **Input:** arm set \mathcal{X} , detection window w , threshold b and context ids CA_{id} .
 - 2: Sample $\frac{w}{2}$ arms $\{\tilde{x}_s\}_{s=1}^{\frac{w}{2}} \sim \lambda^*$ and observe the associated returns $\{Y_{s, \tilde{x}_s}\}_{s=1}^{\frac{w}{2}}$, where $\{(x_s, Y_{s, x_s})\}_{s=t-\frac{w}{2}+1}^t = \{\tilde{x}_s, Y_{s, \tilde{x}_s}\}_{s=1}^{\frac{w}{2}}$.
 - 3: **for** $j \in CA_{\text{id}}$ **do**
 - 4: **if** not LCD($\mathcal{X}, w, b, [(\tilde{x}_s, Y_{s, \tilde{x}_s})]_{s=1}^{\frac{w}{2}} + CA_{\text{id}}[j]$) **then**
 - 5: $CA_{\text{id}}[j] = [(\tilde{x}_s, Y_{s, \tilde{x}_s})]_{s=1}^{\frac{w}{2}}$.
 - 6: **Return** j, CA_{id} .
 - 7: **end if**
 - 8: **end for**
 - 9: index = len(CA_{id}) + 1.
 - 10: $CA_{\text{id}}[\text{index}] = [(\tilde{x}_s, Y_{s, \tilde{x}_s})]_{s=1}^{\frac{w}{2}}$.
 - 11: **Return** index, CA_{id} .
-

whether the current latent context has been visited in previous time steps by scanning through CA_{id} (Line 3). In order to learn if the current context can be aligned with context j , the LCD subroutine is called with the $w/2$ recent samples and $CA_{\text{id}}[j]$ as input:

(i) If the current context is aligned with $CA_{\text{id}}[j]$ (Line 4), the current latent context is j w.h.p. Thus the LCA subroutine updates $CA_{\text{id}}[j]$ and returns the index j (Lines 5 to 6), which would be \hat{j}_t in Line 19 of Algorithm 1. Note that in Lines 10 and 11 of Algorithm 1, all the collected samples will be assigned to context \hat{j}_t (until the next changing alarm) and will be used to estimate $\theta_{\hat{j}_t}^*$ and all p_j 's. Therefore, *aligning* contexts allows PS ϵ BAI to make good use of observation history.

(ii) If the current context is not aligned with any observed context, i.e., it has not been visited, CA_{id} gets extended with a new index-samples pair and is returned along with the new index (Lines 9 to 11).

D.3 The PIECEWISE-STATIONARY ε -BEST ARM IDENTIFICATION⁺ (PS ε BAI⁺) Algorithm

To help understand PS ε BAI⁺ algorithm, we highlight the **differences** between PS ε BAI and PS ε BAI⁺, and the **differences** between subroutines LINEAR-CONTEXT ALIGNMENT (LCA) and LINEAR-CONTEXT ALIGNMENT⁺ (LCA⁺).

Algorithm 5 PIECEWISE-STATIONARY ε -BEST ARM IDENTIFICATION⁺ (PS ε BAI⁺)

- 1: **Input:** arm set \mathcal{X} , number of latent vectors N , bounds on the segment lengths L_{\min} and L_{\max} , slackness parameter ε , confidence parameter δ , sampling parameter γ , window size w and threshold b .
 - 2: **Initialize:** Compute the G-optimal allocation λ^* and $\tau^* = \frac{38400 \ln(80) N L_{\max}}{\varepsilon^2} \ln \frac{N^2 K L_{\max}}{\delta \varepsilon^2}$.
 - 3: Set $\text{CD}_{\text{sample}} = []$, $\text{CA}_{\text{id}} = \{ \}$, $t_{\text{CD}} = +\infty$. Set $\mathcal{T}_{t,j} = \emptyset$ and initialize $\mathcal{T}_t, T_{t,j}, \hat{T}_t$ with (3.2) for all $t \leq \tau^*$, $j \in [N]$.
 - 4: Compute $C_3 = \sum_{n=1}^{\infty} n^{-3}$ and $t^* = 3d \ln(6dKC_3/\delta)$.
 - 5: Sample $w/2$ arms $\{x_s\}_{s=1}^{w/2} \sim \lambda^*$ and observe the associated returns $\{Y_{s,x_s}\}_{s=1}^{w/2}$.
 - 6: Set $t = \frac{w}{2}$, $t_{\text{CA}} = w/2$, $\text{CA}_{\text{id}} = \{1 : [(x_s, Y_{s,x_s})]_{s=1}^{w/2}\}$, $\hat{j}_t = 1$.
 - 7: Set $\hat{x}_t = 0$, $\hat{\rho}_t = 2\varepsilon$ and $\hat{x}_\varepsilon = -\infty$. Compute $\hat{\theta}_t$ with (D.1).
 - 8: **while True do**
 - 9: $t = t + 1$
 - 10: Sample an arm $x_t \sim \lambda^*$ and observe return Y_{t,x_t} .
 - 11: Compute $\tilde{\theta}_t, \tilde{x}_t, \tilde{\rho}_t = \text{EU}(t, \tilde{\theta}_{t-1}, t^*, C_3)$.
 - 12: **if** $\tilde{x}_t^\top \tilde{\theta}_t - \tilde{\rho}_t + \varepsilon \geq \max_{x \neq \hat{x}_t} x^\top \tilde{\theta}_t + \tilde{\rho}_t$ **then**
 - 13: $\hat{x}_\varepsilon = \tilde{x}_t$.
 - 14: **Break**
 - 15: **end if**
 - 16: **if** $t > \tau^*$ **then**
 - 17: **Continue**
 - 18: **end if**
 - 19: **if** $\text{mod}(t - t_{\text{CA}}, \gamma) \neq 0$ **then**
 - 20: Update $\hat{j}_t = \hat{j}_{t-1}$, $\mathcal{T}_{t,\hat{j}_t} = \mathcal{T}_{t-1,\hat{j}_t} \cup \{t\}$, $\mathcal{T}_{t,j} = \mathcal{T}_{t-1,j}$ for $j \neq \hat{j}_t$.
 - 21: **else**
 - 22: $\text{CD}_{\text{sample}} = \text{CD}_{\text{sample}} + [(x_t, Y_{t,x_t})]$.
 - 23: **if** $|\text{CD}_{\text{sample}}| \geq w$ **then**
 - 24: **if** $\text{LCD}(\mathcal{X}, w, b, \text{CD}_{\text{sample}}[-w :]) \text{ then}$
 - 25: $\text{CD}_{\text{sample}} = []$.
 - 26: $t = t + \frac{w}{2}$, $t_{\text{CA}} = t$, $t_{\text{CD}} = +\infty$.
 - 27: $\hat{j}_t, \text{CA}_{\text{id}} = \text{LCA}^+(\mathcal{X}, w, b, \text{CA}_{\text{id}})$
 - 28: **if** $\hat{j}_t = N + 1$ **then break.**
 - 29: Revert $\mathcal{T}_{t,j} = \mathcal{T}_{t - \frac{w(\gamma+1)}{2}, j}$ for all $j \in [N]$.
 - 30: **end if**
 - 31: **end if**
 - 32: **end if**
 - 33: **if** $\hat{x}_\varepsilon \neq -\infty$ **then**
 - 34: **Break**
 - 35: **end if**
 - 36: Update the empirical estimates by (3.2) and (3.3).
 - 37: **if** condition (3.4) is met and $t_{\text{CD}} = +\infty$ **then**
 - 38: Record $\hat{x}_\varepsilon = \arg \max_{x \in \mathcal{X}} x^\top \hat{\Theta}_t \hat{\mathbf{p}}_t$.
 - 39: $t_{\text{CD}} = |\text{CD}_{\text{sample}}|$.
 - 40: **else if** $t_{\text{CD}} = |\text{CD}_{\text{sample}}| - \frac{w}{2}$ **then**
 - 41: $\hat{x}_\varepsilon = \hat{x}_\varepsilon$
 - 42: **Break**
 - 43: **end if**
 - 44: **end while**
 - 45: **Recommend** $\hat{x}_\varepsilon = \hat{x}_\varepsilon$.
-

Algorithm 6 ESTIMATE UPDATE (EU)

- 1: **Input:** time step t , vector $\tilde{\theta}_{t-1}$, threshold t^* and constant C_3 .
 - 2: Compute $\tilde{\theta}_t = \frac{1}{t} \left[(t-1)\tilde{\theta}_{t-1} + A(\lambda^*)^{-1} x_t Y_{s,x_t} \right]$, $\dot{x}_t = 0$, $\tilde{\rho}_t = 2\varepsilon$.
 - 3: **if** $t \geq t^*$ **then**
 - 4: Compute $\dot{x}_t = \arg \max_{x \in \mathcal{X}} x^\top \tilde{\theta}_t$, $\tilde{\rho}_t = \sqrt{\frac{8L_{\max}}{t} \ln \frac{4KC_3 t^3}{\delta}} + 5\sqrt{\frac{d}{t} \ln \frac{4KC_3 t^3}{\delta}}$.
 - 5: **end if**
 - 6: **Return** $\tilde{\theta}_t$, \dot{x}_t , $\tilde{\rho}_t$.
-

Algorithm 7 LINEAR-CONTEXT ALIGNMENT⁺ (LCA⁺)

- 1: **Input:** arm set \mathcal{X} , detection window w , threshold b , context ids CA_{id} , time step t , vector $\tilde{\theta}_t$, threshold t^* and constant C_3 .
 - 2: Set $s = 0$, $\dot{x}_\varepsilon = -\infty$.
 - 3: **for** $s \leq w/2$ **do**
 - 4: $s = s + 1$, $t = t + 1$.
 - 5: Sample an arm $\tilde{x}_s \sim \lambda^*$ and observe return Y_{s,\tilde{x}_s} .
 - 6: Set $(x_t, Y_{t,x_t}) = (\tilde{x}_s, Y_{s,\tilde{x}_s})$.
 - 7: Compute $\tilde{\theta}_t$, \dot{x}_t , $\tilde{\rho}_t = \text{EU}(t, \tilde{\theta}_{t-1}, t^*, C_3)$.
 - 8: **if** $\dot{x}_t^\top \tilde{\theta}_t - \tilde{\rho}_t + \varepsilon \geq \max_{x \neq \dot{x}_t} x^\top \tilde{\theta}_t + \tilde{\rho}_t$ **then**
 - 9: $\dot{x}_\varepsilon = \dot{x}_t$, $\text{index} = \text{len}(\text{CA}_{\text{id}}) + 1$.
 - 10: **Return** index , CA_{id} , \dot{x}_ε , $\tilde{\theta}_t$.
 - 11: **end if**
 - 12: **end for**
 - 13: **for** $j \in \text{CA}_{\text{id}}$ **do**
 - 14: **if** not $\text{LCD}(w, b, [(\tilde{x}_s, Y_{s,\tilde{x}_s})]_{s=1}^{\frac{w}{2}}, \text{CA}_{\text{id}}[j])$ **then**
 - 15: $\text{CA}_{\text{id}}[j] = [(\tilde{x}_s, Y_{s,\tilde{x}_s})]_{s=1}^{\frac{w}{2}}$.
 - 16: **Return** j , CA_{id} .
 - 17: **end if**
 - 18: **end for**
 - 19: $\text{index} = \text{len}(\text{CA}_{\text{id}}) + 1$.
 - 20: $\text{CA}_{\text{id}}[\text{index}] = [(\tilde{x}_s, Y_{s,\tilde{x}_s})]_{s=1}^{\frac{w}{2}}$.
 - 21: **Return** index , CA_{id} , \dot{x}_ε , $\tilde{\theta}_t$.
-

D.4 Computational Complexity of $\text{PS}_\varepsilon\text{BAI}$

We provide analysis for the computational complexity of $\text{PS}_\varepsilon\text{BAI}$ in this subsection.

We consider the number of these operations: arithmetic (addition and multiplication) operations, logic operations and comparison operations and we also regard $\ln(\cdot)$, $\sqrt{\cdot}$ and sampling from a distribution as one step of operation.

We decompose the main loop of Algorithm 1 as follows, where the lines with $O(1)$ operations are omitted:

- Exploration phase (Lines 8 to 11): $O(1)$.
- Change Detection phase (Lines 12 to 16):
 - The LCD subroutine in Line 16 needs $O(wd^2 + Kd)$ operations.
- Context Alignment phase (Lines 17 to 21):
 - The LCA subroutine in Line 19 needs $O((wd^2 + Kd)N)$ operations, as the LCD subroutine will be invoked N times in the worst case in Line 4 of LCA .
 - The reversion procedure in Line 21 needs $O(\gamma wd)$ operations.
- The updating procedure and the stopping rule checking (Lines 25 to 32):
 - Updating $\tilde{\theta}_{t,j}$ needs $O(d^2)$ operations, as we need to incorporate the latest sample into the estimate.

- Updating $\hat{p}_{t,j}, j \in [N]$ requires N operations.
- Updating the confidence radius and the empirical best arm need $O(KN)$ operations.
- The stopping condition in Equation (3.4) requires $O(K)$ operations.

We remark that some intermediate results can be stored to avoid repeated computations and we believe the algorithm is efficient overall. In particular, since the reward structure is linear, there are K arms, and N contexts, $O(d^2)$, $O(K)$, and $O(N)$ operations probably cannot be avoided.

From the above analysis, in the decreasing order of the number of operations:

- The Change Detection phase requires the most budget as it will be invoked every γ time steps.
- The Context Alignment phase requires the second many operations. While it requires many operations when it is implemented, it will not be called during a stationary segment.
- The updating procedure uses small portion of the operations, as w is usually much greater than K and N .
- The exploration phase demands constant operations in each loop.

E Auxiliary results

Lemma E.1. *Let X_s be the arm drawn with the G-optimal allocation λ^* at time step s and $Y_{s,x_s} = x_s^\top \theta_{j_s}^* + \eta_s$ be the corresponding return with $\mathbb{E}\eta_s = 0$, $|x_s^\top \theta_{j_s}^*| \leq 1$ and $|\eta_s| \leq 1$. Then we have*

$$\mathbb{P} \left[\frac{1}{n} \left| \sum_{s=1}^n x^\top A(\lambda^*)^{-1} x_s Y_{s,x_s} - x^\top \theta_{j_s}^* \right| \geq \epsilon \right] \leq \delta$$

for all $\epsilon \geq \frac{d}{n} \ln \frac{2}{\delta} + \sqrt{\left(\frac{d}{n} \ln \frac{2}{\delta}\right)^2 + \frac{4d}{n} \ln \frac{2}{\delta}}$. In addition, if $n \geq \frac{d}{4} \ln \frac{2}{\delta}$, we have

$$\mathbb{P} \left[\frac{1}{n} \left| \sum_{s=1}^n x^\top A(\lambda^*)^{-1} x_s Y_{s,x_s} - x^\top \theta_{j_s}^* \right| \geq 5 \sqrt{\frac{d}{n} \ln \frac{2}{\delta}} \right] \leq \delta$$

Proof. Note that the arms are selected according to the G-optimal design, and thus $\mathbb{E}_{x \sim \lambda^*} [xx^\top] = A(\lambda^*)$. Therefore, for all $s \in [n]$,

$$\mathbb{E} [x^\top A(\lambda^*)^{-1} x_s Y_{s,x_s} - x^\top \theta_{j_s}^*] = 0 \quad (\text{E.1})$$

and

$$\begin{aligned} & |x^\top A(\lambda^*)^{-1} x_s Y_{s,x_s} - x^\top \theta_{j_s}^*| \\ & \leq |x^\top A(\lambda^*)^{-1} x_s x_s^\top \theta_{j_s}^*| + |x^\top A(\lambda^*)^{-1} x_s \eta_s| + 1 \\ & \leq |x^\top A(\lambda^*)^{-1} x_s| |x_s^\top \theta_{j_s}^*| + |x^\top A(\lambda^*)^{-1} x_s| |\eta_s| + 1 \\ & \leq 2|x^\top A(\lambda^*)^{-1} x_s| + 1 \\ & \leq 2\|x\|_{A(\lambda^*)^{-1}} \|x_s\|_{A(\lambda^*)^{-1}} + 1 \\ & \leq 3d \end{aligned} \quad (\text{E.2})$$

where we make use of the property of the G-optimal allocation in the last inequality

$$\|x\|_{A(\lambda^*)^{-1}}^2 \leq d, \quad \forall x \in \mathcal{X}. \quad (\text{E.3})$$

Additionally,

$$\begin{aligned} & \mathbb{E} \left[(x^\top A(\lambda^*)^{-1} x_s Y_{s,x_s})^2 \right] \\ & = \mathbb{E} \left[(x^\top A(\lambda^*)^{-1} x_s (x_s^\top \theta_{j_s}^* + \eta_s))^2 \right] \\ & = \mathbb{E} \left[(x^\top A(\lambda^*)^{-1} x_s x_s^\top \theta_{j_s}^*)^2 \right] + \mathbb{E} \left[(x^\top A(\lambda^*)^{-1} x_s \eta_s)^2 \right] \\ & \quad + 2\mathbb{E} [x^\top A(\lambda^*)^{-1} x_s x_s^\top \theta_{j_s}^* \cdot x^\top A(\lambda^*)^{-1} x_s \eta_s] \\ & \stackrel{(a)}{=} \mathbb{E} \left[(x^\top A(\lambda^*)^{-1} x_s)^2 (x_s^\top \theta_{j_s}^*)^2 \right] + \mathbb{E} \left[\eta_s^2 (x^\top A(\lambda^*)^{-1} x_s)^2 \right] \end{aligned} \quad (\text{E.4})$$

$$\begin{aligned}
&\stackrel{(b)}{\leq} 2\mathbb{E} \left[\left(x^\top A(\lambda^*)^{-1} x_s \right)^2 \right] \\
&\stackrel{(c)}{=} 2 x^\top A(\lambda^*)^{-1} x \\
&\stackrel{(d)}{\leq} 2d
\end{aligned}$$

where we make use of the fact that η_t is zero-mean and is independent of other random variables in (a); $|x_s^\top \theta_{j_s}^*| \leq 1$ and $\mathbb{P}[\eta_s^2 \leq 1] = 1$ in (b); $x_s \sim \lambda^*$ in (c); and the property of the G-optimal allocation (E.3) in (d).

According to the Bernstein's inequality with (E.1), (E.2) and (E.4),

$$\begin{aligned}
&\mathbb{P} \left[\frac{1}{n} \left| \sum_{s=1}^n x^\top A(\lambda^*)^{-1} x_s Y_{s,x_s} - x^\top \theta_{j_s}^* \right| \geq \epsilon \right] \\
&\leq 2 \exp \left(-\frac{\frac{1}{2}(n\epsilon)^2}{n \cdot 2d + dn\epsilon} \right) \\
&= 2 \exp \left(-\frac{n\epsilon^2}{4d + 2d\epsilon} \right).
\end{aligned}$$

In order to upper bound the error probability by δ , we need

$$\begin{aligned}
&2 \exp \left(-\frac{n\epsilon^2}{4d + 2d\epsilon} \right) \leq \delta \\
\Rightarrow \quad \epsilon &\geq \frac{d}{n} \ln \frac{2}{\delta} + \sqrt{\left(\frac{d}{n} \ln \frac{2}{\delta} \right)^2 + \frac{4d}{n} \ln \frac{2}{\delta}}.
\end{aligned}$$

In addition, if $n \geq \frac{d}{4} \ln \frac{2}{\delta}$, we have

$$5 \sqrt{\frac{d}{n} \ln \frac{2}{\delta}} = \frac{5}{2} \max \left\{ \frac{d}{n} \ln \frac{2}{\delta}, \sqrt{\frac{4d}{n} \ln \frac{2}{\delta}} \right\} \geq \frac{d}{n} \ln \frac{2}{\delta} + \sqrt{\left(\frac{d}{n} \ln \frac{2}{\delta} \right)^2 + \frac{4d}{n} \ln \frac{2}{\delta}}.$$

This finishes the proof. \square

F Analysis of N_ϵ BAI

To analyze the theoretical performance of N_ϵ BAI, we first show that it can identify an ϵ -optimal arm with probability $1 - \delta$ and derive a high-probability upper bound on its stopping time in Lemma F.1.

Lemma F.1 (High-probability upper bound of N_ϵ BAI). *With probability $1 - \delta$, the N_ϵ BAI algorithm identifies an ϵ -optimal arm after at most*

$$\tilde{O} \left(\frac{L_{\max} + d}{(\epsilon + \Delta_{\min})^2} \ln \frac{1}{\delta} \right)$$

time steps, where $\Delta_{\min} = \min_{i \neq i^} \Delta(x^* - x_i)$.*

Next, we prove that after a sufficiently large number of time steps, the probability that N_ϵ BAI does not terminate is small in Lemma F.2.

Lemma F.2. *Let*

$$T_0 = \frac{768(8L_{\max} + 25d)}{(\epsilon + \Delta_{\min})^2} \ln \frac{768KC_3(8L_{\max} + 25d)}{(\epsilon + \Delta_{\min})^2 \delta}$$

with $C_3 = \sum_{n=1}^{\infty} n^{-3}$. For $t \geq T_0$, the probability that N_ϵ BAI does not terminate before t time steps is $\frac{1}{(\alpha-1)C_3(t/2)^2}$.

Lastly, we apply Lemmas F.1 and F.2 to prove Proposition 3.1.

Proposition 3.1. Let $\Delta_{\min} = \min_{x \neq x^*} \Delta(x^*, x)$,

$$T_V^N = \frac{d}{(\varepsilon + \Delta_{\min})^2} \ln \frac{1}{\delta} \quad \text{and} \quad T_D^N = \frac{L_{\max}}{(\varepsilon + \Delta_{\min})^2} \ln \frac{1}{\delta}.$$

The $N\varepsilon$ BAI algorithm is (ε, δ) -PAC and its expected sample complexity is $\tilde{O}(T_V^N + T_D^N)$.

The detailed analysis is presented as below.

F.1 Detailed analysis of $N\varepsilon$ BAI

Proof of Lemma F.1. The result is to be proven with the following procedures.

First step: Prove $\left| x^\top \mathbb{E}_{\theta \sim P_\theta} \theta - x^\top \tilde{\theta}_t \right|$ can be bounded with high probability.

Let $\tilde{\theta}_t = \frac{1}{t} \sum_{s=1}^t \theta_{j_s}^*$. (i) For all $\varepsilon > 0$, we have,

$$\begin{aligned} \mathbb{P} \left[\left| x^\top \mathbb{E}_{\theta \sim P_\theta} \theta - x^\top \tilde{\theta}_t \right| > \varepsilon \right] &= \mathbb{P} \left[\left| x^\top \mathbb{E}_{\theta \sim P_\theta} \theta - x^\top \left(\sum_{l=1}^{N_t} L_l \theta^* j_{c_l} \right) \right| > t\varepsilon \right] \\ &= \mathbb{P} \left[\left| \sum_{l=1}^{N_t} L_l (x^\top \mathbb{E}_{\theta \sim P_\theta} \theta - L_l x^\top \theta^* j_{c_l}) \right| > t\varepsilon \right] \stackrel{(a)}{\leq} 2 \exp \left(-\frac{t^2 \varepsilon^2}{2 \sum_{l=1}^{N_t} (2L_l)^2} \right) \\ &\stackrel{(b)}{\leq} 2 \exp \left(-\frac{t\varepsilon^2}{8L_{\max}} \right). \end{aligned}$$

Since $|x^\top \mathbb{E}_{\theta \sim P_\theta} \theta| \leq 1$ and $|x^\top \theta^* j_{c_l}| \leq 1$ for all l , we obtain (a) by applying Hoeffding's inequality. We obtain (b) from the fact that $\sum_{l=1}^{N_t} L_l = t$ and $L_l \leq L_{\max}$ for all l . In other words, for all $\delta > 0$,

$$\mathbb{P} \left[\left| x^\top \mathbb{E}_{\theta \sim P_\theta} \theta - x^\top \tilde{\theta}_t \right| > \sqrt{\frac{8L_{\max}}{t} \ln \frac{2}{\delta}} \right] \leq \delta.$$

(ii) Besides, according to Lemma E.1, if $t \geq \frac{d}{4} \ln \frac{2}{\delta}$, we have

$$\mathbb{P} \left[\left| x^\top \tilde{\theta}_t - x^\top \bar{\theta}_t \right| \geq 5\sqrt{\frac{d}{t} \ln \frac{2}{\delta}} \right] = \mathbb{P} \left[\frac{1}{t} \left| \sum_{s=1}^t x^\top A(\lambda^*)^{-1} x_s Y_{s, x_s} - x^\top \theta_{j_s}^* \right| \geq 5\sqrt{\frac{d}{t} \ln \frac{2}{\delta}} \right] \leq \delta.$$

(iii) For all $t > 0$, all $x \in \mathcal{X}$, since

$$\left| x^\top \mathbb{E}_{\theta \sim P_\theta} \theta - x^\top \tilde{\theta}_t \right| \leq \left| x^\top \mathbb{E}_{\theta \sim P_\theta} \theta - x^\top \bar{\theta}_t \right| + \left| x^\top \bar{\theta}_t - x^\top \tilde{\theta}_t \right|,$$

we have

$$\mathbb{P} \left[\left| x^\top \mathbb{E}_{\theta \sim P_\theta} \theta - x^\top \tilde{\theta}_t \right| > \tilde{\rho}_t(\alpha) \right] \leq \frac{\delta}{KC_\alpha t^\alpha} \quad \text{when} \quad t \geq \frac{d}{4} \ln \frac{4KC_\alpha t^\alpha}{\delta},$$

where

$$\tilde{\rho}_t(\alpha) = \sqrt{\frac{8L_{\max}}{t} \ln \frac{4KC_\alpha t^\alpha}{\delta}} + 5\sqrt{\frac{d}{t} \ln \frac{4KC_\alpha t^\alpha}{\delta}}, \quad C_\alpha = \sum_{n=1}^{\infty} n^{-\alpha}, \quad \alpha > 2.$$

For simplicity, we set $\alpha = 3$ and write $\tilde{\rho}_t(3)$ as $\tilde{\rho}_t$ in the following analysis. We now show that $t \geq \frac{d}{4} \ln \frac{4KC_\alpha t^\alpha}{\delta}$ holds when $t \geq 3d \ln \frac{6dKC_3}{\delta}$ with the following lemma (the proof of which is presented by the end of Appendix F.1).

Lemma F.3. Let $a, b, c > 0$ and $ac \geq 1/2$, then

$$t \geq \max\{2a \ln b, 4ac \ln(2ac)\} \Rightarrow t \geq a \ln(bt^c).$$

With

$$a = \frac{d}{4}, \quad b = \frac{4KC_\alpha}{\delta} = \frac{4KC_3}{\delta}, \quad c = \alpha = 3,$$

Lemma F.3 implies that

$$t \geq 4ac \ln(2abc) = 3d \ln \frac{6dKC_3}{\delta} \Rightarrow t \geq \max\{2a \ln b, 4ac \ln(2ac)\} \Rightarrow t \geq a \ln(bt^c).$$

Define event $\tilde{E}_t = \bigcap_{x \in \mathcal{X}} \left\{ \left| x^\top \mathbb{E}_{\theta \sim P_\theta} \theta - x^\top \tilde{\theta}_t \right| \leq \tilde{\rho}_t(\alpha) \right\}$ for $t \geq t^*$ and $\tilde{E} = \bigcap_{t \geq t^*} \tilde{E}_t$. We have

$$\mathbb{P} \left[\tilde{E}_t^c \right] \leq \sum_{x \in \mathcal{X}} \frac{\delta}{KC_\alpha t^\alpha} = \frac{\delta}{C_\alpha t^\alpha}, \quad \mathbb{P} \left[\tilde{E}^c \right] \leq \sum_{t=t^*}^{\infty} \mathbb{P} \left[\tilde{E}_t^c \right] \leq \sum_{t=1}^{\infty} \frac{\delta}{C_\alpha t^\alpha} = \delta.$$

We assume \tilde{E} holds in the second and third steps in this proof.

Second step: Prove N_ε BAI recommends an ε -optimal arm when it stops.

If the algorithm terminates and returns $\hat{x}_\varepsilon \neq x^*$, we have

$$\hat{x}_\varepsilon^\top \tilde{\theta}_t - \tilde{\rho}_t + \varepsilon \geq \max_{x \neq \hat{x}_\varepsilon} x^\top \tilde{\theta}_t + \tilde{\rho}_t.$$

Since

$$\hat{x}_\varepsilon^\top \mathbb{E}_{\theta \sim P_\theta} \theta + \varepsilon \geq \hat{x}_\varepsilon^\top \tilde{\theta}_t - \tilde{\rho}_t + \varepsilon \quad \text{and} \quad \max_{x \neq \hat{x}_\varepsilon} x^\top \tilde{\theta}_t + \tilde{\rho}_t \geq \max_{x \neq \hat{x}_\varepsilon} x^\top \mathbb{E}_{\theta \sim P_\theta} \theta \geq x^{*\top} \mathbb{E}_{\theta \sim P_\theta} \theta,$$

we have

$$\hat{x}_\varepsilon^\top \mathbb{E}_{\theta \sim P_\theta} \theta + \varepsilon \geq x^{*\top} \mathbb{E}_{\theta \sim P_\theta} \theta,$$

implying that $\hat{x}_\varepsilon \in \mathcal{X}_\varepsilon$. Hence, we always have $\hat{x}_\varepsilon \in \mathcal{X}_\varepsilon$.

Third step: Derive the stopping time of N_ε BAI.

(i) Let $x_{(2)} = \arg \max_{x \neq x^*} x^\top \mathbb{E}_{\theta \sim P_\theta} \theta$ and $\Delta_{\min} = \min_{x \neq x^*} \Delta(x^* - x) = (x^* - x_{(2)})^\top \mathbb{E}_{\theta \sim P_\theta} \theta$. We apply the following lemma to show that N_ε BAI stops when $\tilde{\rho}_t$ is sufficiently small.

Lemma F.4 ([30, Lemma 3]). *Denote by \hat{i} the index of the item with empirical mean is i -th largest: i.e., $\hat{w}(\hat{1}) \geq \dots \geq \hat{w}(\hat{L})$. Assume that the empirical means of the arms are controlled by ε : i.e., $\forall i, |\hat{w}(i) - w(i)| < \varepsilon$. Then,*

$$\forall i, w(i) - \varepsilon \leq \hat{w}(\hat{i}) \leq w(i) + \varepsilon.$$

Lemma F.4 implies that

$$\begin{aligned} \hat{x}_t^\top \tilde{\theta}_t - \tilde{\rho}_t + \varepsilon &\geq x^{*\top} \mathbb{E}_{\theta \sim P_\theta} \theta - 2\tilde{\rho}_t + \varepsilon = x_{(2)}^\top \mathbb{E}_{\theta \sim P_\theta} \theta - 2\tilde{\rho}_t + \varepsilon + \Delta_{\min} \quad \text{and} \\ \max_{x \neq \hat{x}_t} x^\top \tilde{\theta}_t + \tilde{\rho}_t &\leq x_{(2)}^\top \mathbb{E}_{\theta \sim P_\theta} \theta + 2\tilde{\rho}_t. \end{aligned}$$

When $\tilde{\rho}_t \leq (\varepsilon + \Delta_{\min})/4$, we have

$$\hat{x}_t^\top \tilde{\theta}_t - \tilde{\rho}_t + \varepsilon \geq x_{(2)}^\top \tilde{\theta}_t + \frac{\varepsilon + \Delta_{\min}}{2} \geq x_{(2)}^\top \tilde{\theta}_t + 2\tilde{\rho}_t \geq \max_{x \neq \hat{x}_t} x^\top \tilde{\theta}_t + \tilde{\rho}_t,$$

which will lead to the termination of N_ε BAI.

(ii) According to the definition of $\tilde{\rho}_t$, we have

$$\begin{aligned} \tilde{\rho}_t(\alpha) &\leq (\varepsilon + \Delta_{\min})/4 \\ \Leftrightarrow \sqrt{\frac{8L_{\max}}{t} \ln \frac{4KC_\alpha t^\alpha}{\delta}} + 5\sqrt{\frac{d}{t} \ln \frac{4KC_\alpha t^\alpha}{\delta}} &\leq (\varepsilon + \Delta_{\min})/4 \\ \Leftrightarrow \sqrt{\frac{1}{t} \ln \frac{4KC_\alpha t^\alpha}{\delta}} \left(\sqrt{8L_{\max}} + 5\sqrt{d} \right) &\leq (\varepsilon + \Delta_{\min})/4 \\ \Leftrightarrow t &\geq \frac{\left(\sqrt{8L_{\max}} + 5\sqrt{d} \right)^2}{(\varepsilon + \Delta_{\min})^2 / 16} \ln \frac{4KC_\alpha t^\alpha}{\delta} \\ \Leftrightarrow t &\geq \frac{32(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2} \ln \frac{4KC_\alpha t^\alpha}{\delta} \end{aligned}$$

$$\Leftrightarrow t \geq \frac{192(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2} \ln \frac{768KC_3(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2 \delta}.$$

We apply Lemma F.3 to invert the last line above. With

$$a = \frac{32(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2}, \quad b = \frac{4KC_\alpha}{\delta}, \quad c = \alpha = 3,$$

Lemma F.3 indicates that

$$\begin{aligned} t &\geq 4ac \ln(2abc) = \frac{384(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2} \ln \frac{768KC_3(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2 \delta} \\ &\Rightarrow t \geq \max\{2a \ln b, 4ac \ln(2ac)\} \Rightarrow t \geq a \ln(bt^c). \end{aligned}$$

Altogether, we show that with probability $1 - \delta$, $\text{N}\varepsilon\text{BAI}$ identifies an ε -optimal arm after at most

$$\frac{384(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2} \ln \frac{768KC_3(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2 \delta}$$

time steps. □

Proof of Lemma F.2. To begin with, let

$$T_0 = \frac{768(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2} \ln \frac{768KC_3(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2 \delta}$$

For any $T \geq T_0$, let $\bar{T} = \lceil T/2 \rceil$.

(I) If $\text{N}\varepsilon\text{BAI}$ terminates within \bar{T} time steps, there is nothing to prove.

(II) Assume $\text{N}\varepsilon\text{BAI}$ does not terminate within \bar{T} time steps. According to the analysis in the proof of Lemma F.1, if $\text{N}\varepsilon\text{BAI}$ does not terminate before T time steps, i.e., the stopping time τ satisfies

that $\tau \geq \bar{T}$, then event $\bigcap_{t=\bar{T}+1}^{T-1} \tilde{E}_t$ does not hold. This indicates

$$\begin{aligned} \mathbb{P}[\tau \geq T | \tau \geq \bar{T}] &\leq \mathbb{P}\left[\bigcup_{t=\bar{T}+1}^{T-1} \tilde{E}_t^c\right] \leq \sum_{t=\bar{T}}^{T-1} \mathbb{P}[\tilde{E}_t^c] \leq \sum_{t=\bar{T}}^{T-1} \frac{\delta}{C_\alpha t^\alpha} \leq \frac{\delta}{C_\alpha} \int_{\bar{T}}^T t^{-\alpha} dt \\ &\leq \frac{\delta}{(\alpha-1)C_\alpha} \left(\frac{1}{(\bar{T})^{\alpha-1}} - \frac{1}{T^\alpha} \right) \stackrel{(a)}{\leq} \frac{\delta}{(\alpha-1)C_\alpha (T/2)^{\alpha-1}}. \end{aligned}$$

(a) results from the fact that $\bar{T} \geq T/2$.

Altogether, for $T \geq T_0$, the probability $\text{N}\varepsilon\text{BAI}$ does not terminate before T time steps is

$$\frac{\delta}{(\alpha-1)C_\alpha (T/2)^{\alpha-1}}.$$

□

Proof of Proposition 3.1. First, Tonelli's Theorem implies that

$$\mathbb{E}[\tau] = \mathbb{E}\left[\int_0^\tau 1 dx\right] = \mathbb{E}\left[\int_0^{+\infty} \mathbb{I}(\tau > x) dx\right] = \int_0^{+\infty} \mathbb{E}[\mathbb{I}(\tau > x)] dx = \int_0^{+\infty} \mathbb{P}(\tau > x) dx.$$

Next, we apply Lemma F.2 to show

$$\begin{aligned} \mathbb{E}\tau &\leq T_0 + \mathbb{E}[\tau | \tau \geq T_0 + 1] \cdot \mathbb{P}[\tau \geq T_0 + 1] \leq T_0 + \int_{T_0}^{+\infty} \mathbb{P}(\tau \geq x) dx \\ &\leq T_0 + \int_{T_0}^{+\infty} \frac{\delta}{(\alpha-1)C_\alpha (x/2)^{\alpha-1}} dx = T_0 + \frac{\delta}{(\alpha-1)(\alpha-2)(T_0/2)^{\alpha-2}}. \quad (\text{F.1}) \end{aligned}$$

Besides, Lemma F.1 indicates that $\text{N}\varepsilon\text{BAI}$ can identify an ε -optimal arm with probability $1 - \delta$. □

Proof of Lemma F.3. For $x > 0$, let

$$f(x) = x - a \ln(bx^c) = x - a \ln b - ac \ln x = x \left(1 - \frac{a \ln b}{x} - \frac{ac \ln x}{x} \right).$$

Since

$$f(x) \geq 0 \Leftrightarrow \begin{cases} \frac{a \ln b}{x} \leq \frac{1}{2} \\ \frac{ac \ln x}{x} \leq \frac{1}{2} \end{cases} \Leftrightarrow \begin{cases} x \geq 2a \ln b \\ x \geq 2ac \ln x \end{cases}.$$

Let $d = 2ac$, $x = d \ln(dy)$, then

$$x \geq 2ac \ln x = d \ln x \Leftrightarrow d \ln(dy) \geq d \ln(d \ln(dy)) \Leftrightarrow \ln(dy) \geq \ln(d \ln(dy)) \Leftrightarrow y \geq \ln(dy).$$

Since $z^{0.4} \geq \ln z$ for all $z \geq 1$, we have $\ln(dy) \leq (dy)^{0.4} \leq \sqrt{dy}$ and

$$y \geq \ln(dy) \Leftrightarrow y \geq \sqrt{dy} \Leftrightarrow y \geq d$$

when $yd \geq 1$. Hence, when $ac \geq 1/2$, $y \geq d = 2ac \geq 1$, we have $y \geq \ln(dy)$. Furthermore, for x such that $x \geq \max\{2a \ln b, 4ac \ln(2ac)\}$, we have $f(x) \geq 0$. □

G Proof Outline of Theorem 3.2

A proof outline of Theorem 3.2 is provided in this section. It consists of three steps:

Step 1: $\text{PS}\varepsilon\text{BAI}$ (Algorithm 1) depends on the success of change detection and context alignment (Algorithm 3 and 4). We firstly upper bound the failure probability of these two subroutines via Lemma G.1, Lemma G.2, Lemma G.3 and summarized in Lemma G.4. More details about these two subroutines are provided in Appendix D.2.2 and the proof of the Lemmas are postponed to Appendix H.

Step 2: Subsequently, conditioned on their success, we provide a theoretical guarantee for the choice of the confidence radius $\rho_t(x, \tilde{x})$ for $\Delta(x, \tilde{x})$ at time step t in Lemma G.5. The proof is detailed in Appendix I.

Step 3: Lastly, utilizing the above elements, we provide a sufficient condition for the stopping rule, and upper bound the number of time steps in Exp phases T_τ via Lemma G.6 whose proof is presented in Appendix J. As the total number of time steps τ is upper bounded by a constant multiple of T_τ , the high-probability upper bound on the stopping time τ is obtained. This finishes the proof of Theorem 3.2 (please refer to Appendix J).

Step 1: We borrow the terminology in hypothesis testing to define the errors. Let the null hypothesis be: *the algorithm has undergone a changepoint within the last w CD samples*. For Algorithm 3, we will characterize the type I error: *the algorithm has experienced a changepoint within the last w CD samples but it fails to raise a changing alarm*, and the type II error: *a changepoint has not occurred but Algorithm 3 raises a changing alarm*. We refer to the event Failed Alarm (FA) and the event False Alarm Error (FAE) as that a type I error occurs and that a type II error occurs, respectively:

$$\begin{aligned} \text{FA}_l &:= \{\text{a type I error occurs at changepoint } c_l\}, l \in \mathbb{N}, \\ \text{FAE}_t &:= \{\text{a type II error occurs time step } t\}, t \in \mathcal{T}_{\text{FAE}}, \\ \text{FA} &:= \bigcup_{\{l: c_l \leq \tau\}} \text{FA}_l \quad \text{and} \quad \text{FAE} := \bigcup_{t \in \mathcal{T}_{\text{FAE}}} \text{FAE}_t, \end{aligned}$$

where $\mathcal{T}_{\text{FAE}} := \{t : t \text{ in CD phase}, t \leq \tau, [t - w\gamma, t] \cap \mathcal{C} = \emptyset\}$ and τ is the stopping time. In term of Algorithm 4, define the event $\text{MI}_l := \{\text{Algorithm 4 misidentifies } \theta_{j_t}^*\}$, $l \in \mathbb{N}$ and $\text{MI} := \bigcup_{\{l: c_l \leq \tau\}} \text{MI}_l$. Define the good event

$$\text{Good} := \{\text{FA}^c \cap \text{FAE}^c \cap \text{MI}^c\}. \tag{G.1}$$

We upper bound the failure probability of Goodin Lemma G.1, Lemma G.2, Lemma G.3 and conclude the results in Lemma G.4.

Lemma G.1. For any $\delta_{\text{FAE}} \in (0, 1)$, with $b \geq \frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}} + \sqrt{\left(\frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}}\right)^2 + \frac{24}{w} d \ln \frac{2}{\delta_{\text{FAE}}}}$, LCD makes no false alarm before the stopping time τ with probability at least

$$\mathbb{P}[\text{FAE}^c] \geq 1 - \frac{\tau}{\gamma} K \delta_{\text{FAE}}.$$

Lemma G.1 can also be stated as, fix b ,

$$\mathbb{P}[\text{FAE}^c] \geq 1 - \frac{\tau}{\gamma} 2K \exp\left(-\min\left\{\frac{3wb}{20d}, \frac{wb^2}{150d}\right\}\right),$$

which indicates we can always decrease the error probability by enlarging the window size w .

Assume c_l is observable, i.e., $\theta_{j_{c_l}}^* \neq \theta_{j_{c_l+1}}^*$, and denote \hat{c}_l as the alarm time of c_l from Algorithm 3.

Lemma G.2. Conditional on FAE^c , for any $\delta_{\text{FA}} \in (0, 1)$, if $\frac{\Delta_c - b}{2} \geq \frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}} + \sqrt{\left(\frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}}\right)^2 + \frac{4d}{w} \ln \frac{2}{\delta_{\text{FA}}}}$,

$$\mathbb{P}\left[c_l \leq \hat{c}_l \leq c_l + \frac{w\gamma}{2} \mid \hat{c}_l \geq c_l\right] \geq 1 - \delta_{\text{FA}}.$$

In addition, $\mathbb{P}[\text{FA} \mid \text{FAE}^c] \leq l_\tau \delta_{\text{FA}}$.

Lemma G.2 guarantees a prompt change alarm will be raised within $\frac{w\gamma}{2}$ time steps after a changepoint occurs. This also explains why $\frac{w\gamma}{2}$ samples are abandoned at line 16 of Algorithm 1 so that $\mathcal{T}_{t,j}$ only keeps samples from context j .

Lemma G.3. Conditional on FAE^c and FA^c , based on the conditions in Lemma G.1 and Lemma G.2,

$$\mathbb{P}[\text{MI} \mid \text{FAE}^c, \text{FA}^c] \leq l_\tau [(N-1)\delta_{\text{FA}} + K\delta_{\text{FAE}}].$$

The intuition of Lemma G.3 is simple: change alarm should not be raised when the samples are from the same context and, change alarms should be raised when the samples are from the other $(N-1)$ contexts. The error made by Algorithm 3 and 4 can be concluded as follows

Lemma G.4. Assume the instance satisfies Assumption 1,

$$\mathbb{P}[\text{Good}] \geq 1 - \frac{\delta}{2\tau^*} \geq 1 - \frac{\delta}{2}.$$

Step 2: To give an upper bound on T_τ , we firstly prove that the estimate of $\Delta(x, \tilde{x})$ at time t is within distance $\rho_t(x, \tilde{x})$ from the ground truth w.h.p. Define the event CI as the estimates of all the mean gaps locate in the high-probability Confidence Interval

$$\text{CI}_t := \{|\hat{\Delta}_t(x, \tilde{x}) - \Delta(x, \tilde{x})| \leq \rho_t(x, \tilde{x}), \forall x, \tilde{x} \in \mathcal{X}\}$$

and $\text{CI} := \bigcap_t \text{CI}_t$, where ρ_t is defined in (3.3).

Lemma G.5. If $T_t \geq \frac{2L_{\max}}{9} \ln \frac{2}{\delta_{d, T_t}}$,

$$\mathbb{P}[\text{CI}_t \mid \text{Good}] \geq 1 - (K\delta_{v, T_t} + N\delta_{d, T_t} + KN\delta_{m, T_t}).$$

In addition,

$$\mathbb{P}[\text{CI} \mid \text{Good}] \geq 1 - \frac{\delta}{2}.$$

In addition to an estimation over the latent vector which is sufficient under the stationary case (α_t in ρ_t), we also need to control the derivation between the empirical distribution $\hat{\mathbf{p}}_t$ and the ground truth \mathbf{p} (β_t in ρ_t), as well as the interactions between the latent vectors and the distribution (ξ_t in ρ_t). This is reflected by $K\delta_{v, T_t}$, $N\delta_{d, T_t}$ and $KN\delta_{m, T_t}$ which bound the the failure probability of Vector Estimation, Distribution Estimation and Residual Estimation, respectively.

Step 3:

Lemma G.6. Conditional on Good and CI, the recommended arm $\hat{x}_\varepsilon \in \mathcal{X}_\varepsilon$ and when Algorithm 1 terminates, the order of T_t is upper bounded by (3.6).

The detailed proof is presented in Appendix J.

H Analysis of PS ϵ BAI: Change Detection and Context Alignment

As the output of Algorithms 3 and 4 depends on random samples, the output may not meet our expectation with some probabilities. We will characterize the probabilities of the three “bad” events: FAE in Appendix H.1, FA in Appendix H.2 and MI in Appendix H.3; and finally upper bound the failure probability of the good event Good in this section.

Lemma G.4. *Assume the instance satisfies Assumption 1,*

$$\mathbb{P}[\text{Good}] \geq 1 - \frac{\delta}{2\tau^*} \geq 1 - \frac{\delta}{2}.$$

Some notations are introduced first

- $\theta_{1:\infty} := (\theta_{j_t}^*)_{t=1}^\infty$ is the latent vector sequence.
- \hat{c}_l indicates the time step when we raise alarm for the l^{th} changepoint.

According to the dynamics (see Dynamics 1) of the problem, at a changepoint $c_l \in \mathcal{C}$, the next latent vector will be sampled from P_θ . Therefore, it may happen that $\theta_{j_{c_l}}^* = \theta_{j_{c_l-1}}^*$, i.e. this changepoint c_l is hidden and these two consecutive stationary segments are observed as one stationary segment. In order to upper bound the errors made by Algorithm 3, we make the following assumption, which yields the worst-case scenario in terms of the error probability.

Assumption 2. *Given $\theta_{1:\infty}$ and the changepoint sequence \mathcal{C} , $\theta_{j_{c_l}}^* \neq \theta_{j_{c_l+1}}^*, \forall l \in \mathbb{N}$.*

In this section, we do not consider the randomness of $\theta_{1:\infty}$, i.e., we consider a realization of the sequence.⁴ The results do not involve any parameter depending on the specific sequence and thus can be applied to any latent vector sequence.

H.1 False Alarm Error

Lemma G.1. *For any $\delta_{\text{FAE}} \in (0, 1)$, with $b \geq \frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}} + \sqrt{\left(\frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}}\right)^2 + \frac{24}{w} d \ln \frac{2}{\delta_{\text{FAE}}}}$, LCD makes no false alarm before the stopping time τ with probability at least*

$$\mathbb{P}[\text{FAE}^c] \geq 1 - \frac{\tau}{\gamma} K \delta_{\text{FAE}}.$$

Proof of Lemma G.1. Note that within a stationary segment of length t , the observations are i.i.d., and thus the time segments between two consecutive false alarm time $\{\hat{c}_l - \hat{c}_{l-1}\}_{l=1}^t$ (where $\hat{c}_0 = 0$) are i.i.d. i.e., it is a renewal process. Thus, it is sufficient to bound the error probability under the stationary case, i.e. $\mathcal{C} = \emptyset$ and $\theta_{1:\infty} = (\theta_{j_t}^* = \theta_{j_1}^*)_{t=1}^\infty$.

Assume that we are under a stationary segment of length t . We wish to upper bound the error $\mathbb{P}[\text{FAE}] = \mathbb{P}[\hat{c}_1 \leq t]$. This is given by Lemma H.1.

Lemma H.1. *Given w CD samples from the same context $[(\tilde{x}_s, Y_{s\gamma}, \tilde{x}_s)]_{s=1}^w$ in a CD phase at time step t , Algorithm 3 makes a false alarm error with probability upper bounded by $K \delta_{\text{FAE}}$, i.e.*

$$\mathbb{P}[\text{FAE}_t] \leq K \delta_{\text{FAE}}$$

$$\text{if } b \geq \frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}} + \sqrt{\left(\frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}}\right)^2 + \frac{24}{w} d \ln \frac{2}{\delta_{\text{FAE}}}}.$$

Assume γ divides t . In the case of a stationary segment, the input samples to Algorithm 3 are $\left\{[(x_{s\gamma}, Y_{s\gamma}, x_{s\gamma})]_{s=k}^{k+w-1}\right\}_{k=1}^{\frac{t}{\gamma}}$ if no changing alarm is raised (We ignore the initialization at Line 5 in Algorithm 1 for simplicity). Denote $\Gamma_i(t) := \{k\gamma : k \equiv i \pmod{w}, k\gamma \leq t\}$, $i = 0, 1, \dots, w-1$. Given any $i = 0, \dots, w-1$, for any $k_1 \neq k_2 \in \Gamma_i(t)$, the two list of samples, $[(x_{s\gamma}, Y_{s\gamma}, x_{s\gamma})]_{s=k_1}^{k_1+w-1}$ and $[(x_{s\gamma}, Y_{s\gamma}, x_{s\gamma})]_{s=k_2}^{k_2+w-1}$, do not overlap (thus, they are independent). Therefore, $\mathbb{P}[k] := \mathbb{P}[\hat{c}_1 =$

⁴This can also be taken as, we are conditioning on $\theta_{1:\infty}$, as the realization of the sequence is independent from the behavior of any agent/algorithm.

$k\gamma + i, \hat{c}_1 \in \Gamma_i(t)$ is a geometric distribution with parameter upper bounded by $K\delta_{FAE}$. We have

$$\mathbb{P}[\hat{c}_1 \in \Gamma_i(t)] \leq 1 - (1 - K\delta_{FAE})^{|\Gamma_i(t)|}.$$

Hence, the cumulative false alarm error is

$$\begin{aligned} \mathbb{P}[FAE] &= \mathbb{P}[\hat{c}_1 \leq t] \\ &= \sum_{i=0}^{w-1} \left[1 - (1 - K\delta_{FAE})^{|\Gamma_i(t)|} \right] \\ &\leq \sum_{i=0}^{w-1} |\Gamma_i(t)| K\delta_{FAE} \\ &\leq \frac{t}{\gamma} K\delta_{FAE}. \end{aligned}$$

□

Proof of Lemma H.1. According to Algorithm 3, given w CD samples from the same context $[(\tilde{x}_s, Y_{s, \tilde{x}_s})]_{s=1}^w$ and $x \in \mathcal{X}$, we need to bound

$$\mathbb{P} \left[\left| \sum_{s=1}^{\frac{w}{2}} x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - \sum_{s=\frac{w}{2}+1}^w x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} \right| \geq \frac{w}{2} b \right].$$

The proof is similar to Lemma E.1. For any $s \in [\frac{w}{2}]$ and $\tilde{s} = s + \frac{w}{2}$

$$\mathbb{E} [x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - x^\top A(\lambda^*)^{-1} \tilde{x}_{\tilde{s}} Y_{\tilde{s}, \tilde{x}_{\tilde{s}}}] = 0$$

and

$$\begin{aligned} &|x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - x^\top A(\lambda^*)^{-1} \tilde{x}_{\tilde{s}} Y_{\tilde{s}, \tilde{x}_{\tilde{s}}}| \\ &\leq |x^\top A(\lambda^*)^{-1} \tilde{x}_s \tilde{x}_s^\top \theta_{j_s}^*| + |x^\top A(\lambda^*)^{-1} \tilde{x}_s \eta_s| + |x^\top A(\lambda^*)^{-1} \tilde{x}_{\tilde{s}} \tilde{x}_{\tilde{s}}^\top \theta_{j_{\tilde{s}}}^*| + |x^\top A(\lambda^*)^{-1} \tilde{x}_{\tilde{s}} \eta_{\tilde{s}}| \\ &\leq 4d. \end{aligned}$$

By making use of (E.4),

$$\begin{aligned} &\mathbb{E} \left[(x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - x^\top A(\lambda^*)^{-1} \tilde{x}_{\tilde{s}} Y_{\tilde{s}, \tilde{x}_{\tilde{s}}})^2 \right] \\ &= \mathbb{E} \left[(x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s})^2 \right] + \mathbb{E} \left[(x^\top A(\lambda^*)^{-1} \tilde{x}_{\tilde{s}} Y_{\tilde{s}, \tilde{x}_{\tilde{s}}})^2 \right] \\ &\quad - 2\mathbb{E} [x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} \cdot x^\top A(\lambda^*)^{-1} \tilde{x}_{\tilde{s}} Y_{\tilde{s}, \tilde{x}_{\tilde{s}}}] \\ &\leq 4d + 2\mathbb{E} [|x^\top \theta_{j_s}^* \cdot x^\top \theta_{j_{\tilde{s}}}^*|] \\ &\leq 6d. \end{aligned}$$

According to the Bernstein's inequality,

$$\begin{aligned} &\mathbb{P} \left[\left| \sum_{s=1}^{\frac{w}{2}} (x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - x^\top A(\lambda^*)^{-1} \tilde{x}_{s+\frac{w}{2}} Y_{s+\frac{w}{2}, \tilde{x}_{s+\frac{w}{2}}}) \right| \geq \frac{w}{2} \epsilon \right] \\ &\leq 2 \exp \left(-\frac{\frac{1}{2} \left(\frac{w}{2} \epsilon \right)^2}{\frac{w}{2} \cdot 6d + \frac{4d}{3} \frac{w}{2} \epsilon} \right) \\ &= 2 \exp \left(-\frac{\frac{w}{2} \epsilon^2}{12d + \frac{8d}{3} \epsilon} \right). \end{aligned}$$

In order to upper bound the error probability by δ_{FAE} , we need

$$2 \exp \left(-\frac{\frac{w}{2} \epsilon^2}{12d + \frac{8d}{3} \epsilon} \right) \leq \delta_{FAE}$$

$$\Rightarrow \epsilon \geq \frac{4d}{3 \cdot \frac{w}{2}} \ln \frac{2}{\delta_{\text{FAE}}} + \sqrt{\left(\frac{4d}{3 \cdot \frac{w}{2}} \ln \frac{2}{\delta_{\text{FAE}}}\right)^2 + 12 \cdot \frac{2}{w} d \ln \frac{2}{\delta_{\text{FAE}}}}.$$

By the choice of $b \geq \frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}} + \sqrt{\left(\frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}}\right)^2 + \frac{24}{w} d \ln \frac{2}{\delta_{\text{FAE}}}}$, we have

$$\mathbb{P} \left[\left| \frac{2}{w} \sum_{s=1}^{\frac{w}{2}} x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - \frac{2}{w} \sum_{s=\frac{w}{2}+1}^w x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} \right| \geq b \right] \leq \delta_{\text{FAE}}.$$

A union bound over the K arms yields the final result. \square

H.2 Failed Alarm

Lemma G.2. *Conditional on FAE^c , for any $\delta_{\text{FA}} \in (0, 1)$, if $\frac{\Delta_c - b}{2} \geq \frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}} + \sqrt{\left(\frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}}\right)^2 + \frac{4d}{w} \ln \frac{2}{\delta_{\text{FA}}}}$,*

$$\mathbb{P} \left[c_l \leq \hat{c}_l \leq c_l + \frac{w\gamma}{2} \mid \hat{c}_l \geq c_l \right] \geq 1 - \delta_{\text{FA}}.$$

In addition, $\mathbb{P} [\text{FA} \mid \text{FAE}^c] \leq l_\tau \delta_{\text{FA}}$.

Proof of Lemma G.2. Conditioned on FAE^c , the detection at the changepoints is independent. Thus we can assume there is only one changepoint c_1 within a certain number of consecutive time steps.

Algorithm 3 is given w CD samples which are collected under two different context. Without loss of generality, we assume the sample selected at time c_1 is among the CD samples (otherwise, we can regard the time step of the first sample from the second context as c_1). We wish to bound

$$\mathbb{P} \left[c_1 \leq \hat{c}_1 \leq c_1 + \frac{\gamma w}{2} \mid \hat{c}_1 \geq c_1 \right]$$

which is the probability of the event that, after a changepoint occurs, a changing alarm will be raised within $\frac{w\gamma}{2}$ time steps (or $\frac{w}{2}$ CD samples). Here, $\hat{c}_1 \geq c_1$ can be guaranteed as we are conditioning on that there is no false alarm error (FAE^c).

The event $c_1 \leq \hat{c}_1 \leq c_1 + \frac{\gamma w}{2}$ indicates, at least one of the CD sample list in

$$\left\{ \left[(x_{c_1+(s-k)\gamma}, Y_{c_1+(s-k)\gamma, x_{c_1+(s-k)\gamma}}) \right]_{s=1}^w \right\}_{k=\frac{w}{2}+1}^w$$

will trigger the changing alarm in Algorithm 3. In particular, in Lemma H.2, we consider the failed arm probability when the CD samples are composed by exactly half samples from each contexts,

$$\left[(x_{c_1+(s-1-\frac{w}{2})\gamma}, Y_{c_1+(s-1-\frac{w}{2})\gamma, x_{c_1+(s-1-\frac{w}{2})\gamma}}) \right]_{s=1}^w.$$

Lemma H.2. *Given w CD samples from the two different contexts $[(\tilde{x}_s, Y_{s, \tilde{x}_s})]_{s=1}^w$ in a CD phase, where $[(\tilde{x}_s, Y_{s, \tilde{x}_s})]_{s=1}^{\frac{w}{2}}$ is from latent vector θ_j^* and $[(\tilde{x}_s, Y_{s, \tilde{x}_s})]_{s=\frac{w}{2}+1}^w$ is from latent vector θ_j^* , Algorithm 3 raises a changing alarm with probability lower bounded by*

$$1 - \delta_{\text{FA}}$$

if $\frac{\Delta_c - b}{2} \geq \frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}} + \sqrt{\left(\frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}}\right)^2 + \frac{4d}{w} \ln \frac{2}{\delta_{\text{FA}}}}$ where $\Delta_c := \max_{x \in \mathcal{X}} |x^\top (\theta_j^* - \theta_j^*)|$ and it is assumed to be greater than b .

By making use of Lemma H.2,

$$\begin{aligned} & \mathbb{P} \left[c_1 \leq \hat{c}_1 \leq c_1 + \frac{w\gamma}{2} \mid \hat{c}_1 \geq c_1 \right] \\ &= \mathbb{P} \left[\exists k \in \left\{ \frac{w}{2} + 1, \dots, w \right\}, \right. \end{aligned}$$

$$\left. \text{LCD} \left(w, b, [(x_{c_1+(s-k)\gamma}, Y_{c_1+(s-k)\gamma, x_{c_1+(s-k)\gamma}})]_{s=1}^w \right) = \text{True} \mid \hat{c}_1 \geq c_1 \right]$$

$$\begin{aligned} &\geq \mathbb{P} \left[\text{LCD} \left(w, b, [(x_{c_1+(s-1-\frac{w}{2})\gamma}, Y_{c_1+(s-1-\frac{w}{2})\gamma, x_{c_1+(s-1-\frac{w}{2})\gamma})}]_{s=1}^w \right) = \text{True} \mid \hat{c}_1 \geq c_1 \right] \\ &\geq 1 - \delta_{\text{FA}} \end{aligned}$$

where $\Delta_c = \max_{x \in \mathcal{X}} |x^\top (\theta_j^* - \theta_j^*)|$ and θ_j^*, θ_j^* are the latent vectors. Hence,

$$\mathbb{P} [\text{FA} | \text{FAE}^c] = 1 - \mathbb{P} [\text{FA}^c | \text{FAE}^c] \leq 1 - (1 - \delta_{\text{FA}})^{l_\tau} \leq l_\tau \delta_{\text{FA}}$$

□

Proof of Lemma H.2. According to the design of LCD, there exists an $x \in \mathcal{X}$,

$$\begin{aligned} &\mathbb{P} \left[\left| \frac{2}{w} \sum_{s=1}^{\frac{w}{2}} x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - \frac{2}{w} \sum_{s=\frac{w}{2}+1}^w x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} \right| \geq b \right] \\ &\geq 1 - \mathbb{P} \left[\left| \sum_{s=1}^{\frac{w}{2}} x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - \sum_{s=\frac{w}{2}+1}^w x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - \frac{w}{2} x^\top (\theta_j^* - \theta_j^*) \right| \right. \\ &\quad \left. \geq \left| \frac{w}{2} b - \frac{w}{2} x^\top (\theta_j^* - \theta_j^*) \right| \right] \\ &= 1 - \mathbb{P} \left[\left| \sum_{s=1}^{\frac{w}{2}} (x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - x^\top \theta_j^*) - \sum_{s=\frac{w}{2}+1}^w (x^\top A(\lambda^*)^{-1} \tilde{x}_s Y_{s, \tilde{x}_s} - x^\top \theta_j^*) \right| \right. \\ &\quad \left. \geq \left| \frac{w}{2} b - \frac{w}{2} x^\top (\theta_j^* - \theta_j^*) \right| \right] \\ &\geq 1 - \delta_{\text{FA}} \end{aligned}$$

where the last inequality holds as $\frac{\Delta_c - b}{2} \geq \frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}} + \sqrt{\left(\frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}} \right)^2 + \frac{4d}{w} \ln \frac{2}{\delta_{\text{FA}}}}$. □

H.3 Context Alignment

Lemma G.3. *Conditional on FAE^c and FA^c , based on the conditions in Lemma G.1 and Lemma G.2,*

$$\mathbb{P} [\text{MI} | \text{FAE}^c, \text{FA}^c] \leq l_\tau [(N-1)\delta_{\text{FA}} + K\delta_{\text{FAE}}].$$

Proof of Lemma G.3. The error of the context alignment procedure can be derived from the FAE and FA analyses. Conditioned on $\text{FAE}^c, \text{FA}^c$ and that the previous $l-1$ contexts are correctly identified, i.e., $\bigcap_{k=1}^{l-1} \text{MI}_k^c$, we have the following statements.

Firstly, according to Lemma H.1, the change alarm at Line 4 in Algorithm 4 will not be triggered with probability at least $(1 - K\delta_{\text{FAE}})$ if the current context is context j . This error has been taken into account in the FAE (see Remark H.3).

Secondly, if the current context is not j (which will occur at most $N-1$ times), a change alarm will be raised with probability at least $1 - \delta_{\text{FA}}$ by Lemma G.2.

Therefore, the error probability at c_l is upper bounded by

$$\mathbb{P} [\text{MI}_l | \text{FAE}^c, \text{FA}^c, \bigcap_{k=1}^{l-1} \text{MI}_k^c] \leq (N-1)\delta_{\text{FA}} + K\delta_{\text{FAE}}$$

and by a union bound, the cumulative error probability bounded by

$$\begin{aligned} \mathbb{P} [\text{MI} | \text{FAE}^c, \text{FA}^c] &= 1 - \mathbb{P} [\text{MI}^c | \text{FAE}^c, \text{FA}^c] \\ &\leq 1 - (1 - (N-1)\delta_{\text{FA}} - K\delta_{\text{FAE}})^{l_\tau} \\ &\leq l_\tau ((N-1)\delta_{\text{FA}} + K\delta_{\text{FAE}}). \end{aligned} \tag{H.1}$$

□

Remark H.3. *Note that (1) we will only bound the FAE when the CD samples are from the same context; (2) in the analysis of FAE, we bound the error probability on the whole horizon up to time t*

as we assume there is no changepoint. In particular, there will be unused and redundant $\frac{w}{2}K\delta_{\text{FAE}}$ errors budget for FAE at each changepoint, which accumulates to $l_\tau \frac{w}{2}K\delta_{\text{FAE}}$ before time step τ . Therefore, the second error term in (H.1), $l_\tau K\delta_{\text{FAE}}$, can be covered by the unused $l_\tau \frac{w}{2}K\delta_{\text{FAE}}$ error budget from FAE, and thus it can be neglected in (H.1).

Proof of Lemma G.4. According to Lemma G.1, G.2, G.3, Assumption 1 and Remark H.3, given a time τ , the total failure probability is upper bounded by

$$\begin{aligned} \mathbb{P}[\text{Good}^c] &= \mathbb{P}[\text{FAE} \cup \text{FA} \cup \text{MI}] & (\text{H.2}) \\ &\leq \frac{\tau}{\gamma}K\delta_{\text{FAE}} + l_\tau\delta_{\text{FA}} + l_\tau \cdot (N-1)\delta_{\text{FA}} \\ &= \frac{\tau}{\gamma}K\delta_{\text{FAE}} + l_\tau N\delta_{\text{FA}}. \end{aligned}$$

In particular, when τ is upper bounded by τ^* in Line 2 of Algorithm 1

$$\tau \leq \tau^* = c_0 \frac{NL_{\max}}{\varepsilon^2} \ln \frac{N^2KL_{\max}/\varepsilon^2}{\delta}$$

where $c_0 = 3 \cdot 6400 \ln 6400$. The number of changepoints till τ is upper bounded by $l_{\tau^*} \leq \frac{\tau^*}{L_{\min}}$.

By Assumption 1, when $b = \frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}} + \sqrt{\left(\frac{8d}{3w} \ln \frac{2}{\delta_{\text{FAE}}}\right)^2 + \frac{24}{w}d \ln \frac{2}{\delta_{\text{FAE}}}}$ and $\delta_{\text{FAE}} = \frac{\gamma\delta}{4(\tau^*)^2K}$, the conditions of Lemma G.1 are met. And we can upper bound $\frac{\tau}{\gamma}K\delta_{\text{FAE}}$ by $\frac{\delta}{4\tau^*} \leq \frac{\delta}{4}$.

By making use of Assumption 1 and setting $\delta_{\text{FA}} = \frac{\delta}{4Nl_{\tau^*}\tau^*}$, we have $\delta_{\text{FA}} > \delta_{\text{FAE}}$ and

$$\begin{aligned} \frac{\Delta_c - b}{2} &\geq \frac{b}{2} \geq \frac{4d}{3w} \ln \frac{2}{\delta_{\text{FAE}}} + \sqrt{\left(\frac{4d}{3w} \ln \frac{2}{\delta_{\text{FAE}}}\right)^2 + \frac{6}{w}d \ln \frac{2}{\delta_{\text{FAE}}}} \\ &\geq \frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}} + \sqrt{\left(\frac{d}{w} \ln \frac{2}{\delta_{\text{FA}}}\right)^2 + \frac{4d}{w} \ln \frac{2}{\delta_{\text{FA}}}}. \end{aligned}$$

Thus, Lemma G.2 can be applied. $l_{\tau^*}N\delta_{\text{FA}}$ can be upper bounded by $\frac{\delta}{4\tau^*} \leq \frac{\delta}{4}$.

Therefore, according to (H.2), $\mathbb{P}[\text{Good}] \geq 1 - \frac{\delta}{2\tau^*} \geq 1 - \frac{\delta}{2}$. □

I Analysis of PS ε BAI: Estimation Error

Lemma G.5 is proved in this section. We will upper bound these three error terms (see (I.1)): VE error in App I.1, DE error in App I.2, RE error in App I.3, and finally we will prove Lemma G.5 which upper bounds the failure probability of CI at the end of this section.

Lemma G.5. If $T_t \geq \frac{2L_{\max}}{9} \ln \frac{2}{\delta_{d,T_t}}$,

$$\mathbb{P}[\text{CI}_t | \text{Good}] \geq 1 - (K\delta_{v,T_t} + N\delta_{d,T_t} + KN\delta_{m,T_t}).$$

In addition,

$$\mathbb{P}[\text{CI} | \text{Good}] \geq 1 - \frac{\delta}{2}.$$

Given any two arms $x, \tilde{x} \in \mathcal{X}$, by the triangular inequality, the deviation between $\hat{\Delta}_t(x, \tilde{x})$ and $\Delta(x, \tilde{x})$ can be upper bounded by three terms: the Vector-Estimation Error (VE) term, the Distribution-Estimation Error (DE) term, and the Residual Estimation Error (RE) term:

$$\begin{aligned} &|\hat{\Delta}_t(x, \tilde{x}) - \Delta(x, \tilde{x})| & (\text{I.1}) \\ &= |(x - \tilde{x})^\top \hat{\Theta}_t \hat{\mathbf{p}}_t - (x - \tilde{x})^\top \Theta \mathbf{p}| \end{aligned}$$

$$\begin{aligned}
&\leq \underbrace{\left| \sum_{j=1}^N \left(\hat{\Delta}_{t,j}(x, \tilde{x}) - \Delta_j(x, \tilde{x}) \right) \hat{p}_{t,j} \right|}_{\text{VE term}} + \underbrace{\left| \sum_{j=1}^N \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) (\hat{p}_{t,j} - p_j) \right|}_{\text{DE term}} \\
&\quad + \underbrace{\left| \sum_{j=1}^N \left(\Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) \right) (\hat{p}_{t,j} - p_j) \right|}_{\text{RE term}},
\end{aligned}$$

where $a^{\text{clip}_2} := \text{clip}_2(a) := \min\{\max\{a, -2\}, 2\}$ is a shorthand notation for the value of a that is clipped to the interval $[-2, 2]$. The reason the value $\hat{\Delta}_{t,j}(x, \tilde{x})$ is clipped is the ground truth $|\Delta_{t,j}(x, \tilde{x})| \leq 2$.

Recall the event CI

$$\begin{aligned}
\text{CI}_t &:= \{|\hat{\Delta}_t(x, \tilde{x}) - \Delta(x, \tilde{x})| \leq \rho_t(x, \tilde{x}), \forall x, \tilde{x} \in \mathcal{X}\}, \\
\text{CI} &:= \bigcap_t \text{CI}_t,
\end{aligned}$$

and the confidence radius

$$\rho_t(x, \tilde{x}) := 2(\alpha_t + \xi_t) + \sum_{j=1}^N \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) + \zeta_t(x, \tilde{x})|, \tag{I.2}$$

$$\text{where } \alpha_t := 5\sqrt{\frac{d}{T_t} \ln \frac{2}{\delta_{v,T_t}}}, \quad \beta_{t,j} := \frac{5}{2}\sqrt{\frac{2\phi_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}},$$

$$\xi_t := 25\sqrt{2} \frac{N L_{\max}}{T_t} \ln \frac{2}{\delta_{m,T_t}}, \quad \phi_{t,j} := \min \left\{ 4 \max \left\{ \hat{p}_{t,j}, \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \right\}, \frac{1}{4} \right\},$$

and $\zeta_t(x, \tilde{x}) \in \mathbb{R}$ can be any value. In particular, it can be the value that minimizes $\sum_{j=1}^N \beta_{t,j} |\hat{\Delta}_{t,j}(x, \tilde{x}) + \zeta_t(x, \tilde{x})|$ or it can be taken as ε . For simplicity, we will take $\zeta_t(x, \tilde{x}) = \arg \min_{\zeta_t(x, \tilde{x}) \in \mathbb{R}} \sum_{j=1}^N \beta_{t,j} (\hat{\Delta}_{t,j}(x, \tilde{x}) + \zeta_t(x, \tilde{x}))^2 = -\frac{\sum_{j=1}^N \beta_{t,j} \hat{\Delta}_{t,j}(x, \tilde{x})}{\sum_{j=1}^N \beta_{t,j}}$ for a simple and effective analytic expression in the algorithm.

I.1 Vector-Estimation Error

For the VE term $\left| \sum_{j=1}^N \left(\hat{\Delta}_{t,j}(x, \tilde{x}) - \Delta_j(x, \tilde{x}) \right) \hat{p}_{t,j} \right|$, note that

$$\left| \sum_{j=1}^N \left(\hat{\Delta}_{t,j}(x, \tilde{x}) - \Delta_j(x, \tilde{x}) \right) \hat{p}_{t,j} \right| \leq |x^\top (\hat{\Theta}_t - \Theta) \hat{\mathbf{p}}_t| + |\tilde{x}^\top (\hat{\Theta}_t - \Theta) \hat{\mathbf{p}}_t|$$

Lemma I.1. Given $x \in \mathcal{X}$ and $T_t \geq \frac{d}{4} \ln \frac{2}{\delta_{v,T_t}}$,

$$\mathbb{P} \left[|x^\top (\hat{\Theta}_t - \Theta) \hat{\mathbf{p}}_t| \geq \alpha_t | \text{Good} \right] \leq \delta_{v,T_t}$$

Proof of Lemma I.1. By the definition of the estimators in (3.2),

$$\begin{aligned}
x^\top (\hat{\Theta}_t - \Theta) \hat{\mathbf{p}}_t &= \sum_{j=1}^N x^\top (\hat{\theta}_{t,j} - \theta_j^*) \hat{\mathbf{p}}_{t,j} \\
&= \sum_{j=1}^N x^\top \left(\frac{1}{T_{t,j}} \sum_{s \in \mathcal{T}_{t,j}} A(\lambda^*)^{-1} x_s Y_{s,x_s} - \theta_j^* \right) \frac{T_{t,j}}{T_t} \\
&= \frac{1}{T_t} \sum_{j=1}^N \sum_{s \in \mathcal{T}_{t,j}} x^\top A(\lambda^*)^{-1} x_s Y_{s,x_s} - x^\top \theta_j^*
\end{aligned}$$

By Lemma E.1,

$$\mathbb{P} \left[|x^\top (\hat{\Theta}_t - \Theta) \hat{\mathbf{p}}_t| \geq \alpha_t | \theta_{1:\infty}, \text{Good} \right] \leq \delta_{v, T_t}$$

By the property of conditional probability, we have the desired result. \square

Therefore, conditional on Good, with probability at least $1 - K\delta_{v, T_t}$, $\text{VE} \leq 2\alpha_t$ for any $x, \tilde{x} \in \mathcal{X}$.

I.2 Distribution-Estimation Error

Lemma I.2. *Given Good and $T_t \geq \frac{2L_{\max}}{9} \ln \frac{2}{\delta_{d, T_t}}$, for any $j \in [N]$,*

$$\mathbb{P} \left[|\hat{p}_{t,j} - p_j| \geq \beta_{t,j} | \text{Good} \right] \leq \delta_{d, T_t}.$$

Additionally, with probability $1 - N\delta_{d, T_t}$,

$$\left| \sum_{j=1}^N \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) (\hat{p}_{t,j} - p_j) \right| \leq \sum_{j=1}^N \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x})| + \zeta_t(x, \tilde{x})$$

where $\zeta_t(x, \tilde{x})$ can be any value.

Proof of Lemma I.2. Given any $j \in [N]$, and the stationary segment l , we denote $X_{j,l} := \hat{L}_l \mathbb{1}\{\theta_{j_l}^* = \theta_j^*\} - p_j \hat{L}_l$, where \hat{L}_l is the total length of the Exp phases in the l th stationary segment. Note that $\mathbb{E}[X_{j,l} | \{\hat{L}_l\}_{l=1}^{l_t}, \text{Good}] = 0$, $|X_{j,l}| \leq L_{\max}$ a.s. and $\text{Var}[X_{j,l} | \{\hat{L}_l\}_{l=1}^{l_t}, \text{Good}] = p_j(1-p_j) \hat{L}_l^2 \leq p_j(1-p_j) \hat{L}_l L_{\max}$. By Bernstein's inequality,

$$\begin{aligned} & \mathbb{P} \left[\left| \sum_{l=1}^{l_t} X_{j,l} \right| \geq \epsilon | \{\hat{L}_l\}_{l=1}^{l_t}, \text{Good} \right] \leq 2 \exp \left(- \frac{\epsilon^2}{2 \sum_{l=1}^{l_t} \text{Var}[X_{j,l} | \text{Good}] + \frac{2}{3} L_{\max} \epsilon} \right) \\ \Rightarrow & \mathbb{P} \left[\left| \frac{1}{T_t} \sum_{l=1}^{l_t} X_{j,l} \right| \geq \epsilon | \{\hat{L}_l\}_{l=1}^{l_t}, \text{Good} \right] \leq 2 \exp \left(- \frac{T_t^2 \epsilon^2}{2 \sum_{l=1}^{l_t} \text{Var}[X_{j,l} | \text{Good}] + \frac{2}{3} L_{\max} T_t \epsilon} \right) \end{aligned}$$

As $T_{t,j} = \sum_{l=1}^{l_t} \hat{L}_l \mathbb{1}\{\theta_{j_l}^* = \theta_j^*\}$, $T_t = \sum_{l=1}^{l_t} \hat{L}_l$ and $\sum_{l=1}^{l_t} \hat{L}_l = T_t$, we have

$$\begin{aligned} \mathbb{P} \left[|\hat{p}_{t,j} - p_j| \geq \epsilon | \{\hat{L}_l\}_{l=1}^{l_t}, \text{Good} \right] & \leq 2 \exp \left(- \frac{T_t^2 \epsilon^2}{2 \sum_{l=1}^{l_t} \text{Var}[X_{j,l} | \{\hat{L}_l\}_{l=1}^{l_t}, \text{Good}] + \frac{2}{3} L_{\max} T_t \epsilon} \right) \\ & \leq 2 \exp \left(- \frac{T_t^2 \epsilon^2}{2 \sum_{l=1}^{l_t} p_j(1-p_j) \hat{L}_l L_{\max} + \frac{2}{3} L_{\max} T_t \epsilon} \right) \\ & \leq 2 \exp \left(- \frac{T_t \epsilon^2}{2 p_j(1-p_j) L_{\max} + \frac{2}{3} L_{\max} \epsilon} \right) \end{aligned}$$

As the last bound is independent of $\{\hat{L}_l\}_{l=1}^{l_t}$, we have

$$\mathbb{P} \left[|\hat{p}_{t,j} - p_j| \geq \epsilon | \text{Good} \right] \leq 2 \exp \left(- \frac{T_t \epsilon^2}{2 p_j(1-p_j) L_{\max} + \frac{2}{3} L_{\max} \epsilon} \right)$$

If we want to upper bound the above by $\delta_{d, T_t} \in (0, 1)$, i.e.,

$$2 \exp \left(- \frac{T_t \epsilon^2}{2 p_j(1-p_j) L_{\max} + \frac{2}{3} L_{\max} \epsilon} \right) \leq \delta_{d, T_t}$$

we require

$$\epsilon \geq \frac{\frac{1}{3} L_{\max} \ln \frac{2}{\delta_{d, T_t}} + \sqrt{\left(\frac{1}{3} L_{\max} \ln \frac{2}{\delta_{d, T_t}} \right)^2 + 2 T_t p_j(1-p_j) L_{\max} \ln \frac{2}{\delta_{d, T_t}}}}{T_t} \quad (\text{I.3})$$

In particular,

$$\epsilon = \tilde{\beta}_{t,j} := \frac{5}{2} \max \left\{ \frac{L_{\max}}{3T_t} \ln \frac{2}{\delta_{d,T_t}}, \sqrt{\frac{2p_j(1-p_j)L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \right\}$$

satisfies the condition, we have

$$\mathbb{P} \left[|\hat{p}_{t,j} - p_j| \geq \tilde{\beta}_{t,j} \mid \text{Good} \right] \leq \delta_{d,T_t}. \quad (\text{I.4})$$

A problem here is we do not have access to p_j during the dynamics, therefore, we adopt Lemma K.1 to further upper bound $\tilde{\beta}_{t,j}$ by $\beta_{t,j}$.

Lemma K.1. *Given any $j \in [N]$ and $T_t \geq \frac{2L_{\max}}{9} \ln \frac{2}{\delta_{d,T_t}}$, $\tilde{\beta}_{t,j} \leq \beta_{t,j}$.*

Therefore,

$$\mathbb{P} \left[|\hat{p}_{t,j} - p_j| \geq \beta_{t,j} \mid \text{Good} \right] \leq \delta_{d,T_t}.$$

In addition, with probability at least $1 - N\delta_{d,T_t}$, we can upper bound DE term as:

$$\begin{aligned} \left| \sum_{j=1}^N \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x})(\hat{p}_{t,j} - p_j) \right| &= \left| \sum_{j=1}^N \left(\hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) + \zeta_t(x, \tilde{x}) \right) (\hat{p}_{t,j} - p_j) \right| \\ &\leq \sum_{j=1}^N \beta_{t,j} \left| \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) + \zeta_t(x, \tilde{x}) \right| \end{aligned} \quad (\text{I.5})$$

where $\zeta_t(x, \tilde{x}) \in \mathbb{R}$ can be any value. \square

I.3 Residual-Estimation Error

RE is composed by the product of two deviations, i.e., $\left(\Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) \right)$ and $(\hat{p}_{t,j} - p_j)$, as the time step t becomes large, we expect it will converge to zero fast. Thus, it is sufficient to have a coarse estimation of it.

Lemma I.3. *For any $x \in \mathcal{X}$, conditional on that $|\hat{p}_{t,j} - p_j| \leq \beta_{t,j}, \forall j \in [N]$,*

$$\mathbb{P} \left[\left| \sum_{j=1}^N \left(\Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) \right) (\hat{p}_{t,j} - p_j) \right| \geq \xi_t \mid \text{Good} \right] \leq N\delta_{m,T_t}$$

where $\xi_t := 25\sqrt{2} \frac{NL_{\max}}{T_t} \ln \frac{2}{\delta_{m,T_t}}$.

Proof of Lemma I.3. According to Lemma E.1, for any $x \in \mathcal{X}, j \in [N]$ we have

$$\mathbb{P} \left[|x^\top (\theta_j^* - \hat{\theta}_{t,j})| \geq 5\sqrt{\frac{d}{T_{t,j}} \ln \frac{2}{\delta_{m,T_t}}} \mid \text{Good} \right] \leq \delta_{m,T_t} \quad (\text{I.6})$$

if $T_{t,j} \geq \frac{d}{4} \ln \frac{2}{\delta_{m,T_t}}$. Note the fact that $\Delta_j(x, \tilde{x}) \in [-2, 2]$, hence,

$$\left| \Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) \right| \leq \left| \Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}(x, \tilde{x}) \right| \leq |x^\top (\theta_j^* - \hat{\theta}_{t,j})| + |\tilde{x}^\top (\theta_j^* - \hat{\theta}_{t,j})|. \quad (\text{I.7})$$

And $|x^\top (\theta_j^* - \hat{\theta}_{t,j})| \leq 3d$ with probability 1. Denote

$$\psi_{t,j,1} := \max \left\{ \hat{p}_{t,j}, \frac{d}{4T_t} \ln \frac{2}{\delta_{m,T_t}} \right\}, \quad \psi_{t,j,2} := \max \left\{ \hat{p}_{t,j}, \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \right\}. \quad (\text{I.8})$$

We have⁵

$$\begin{aligned}
& \left| \sum_{j=1}^N \left(\Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) \right) (\hat{p}_{t,j} - p_j) \right| \\
& \leq \sum_{j: \psi_{t,j,1} > \hat{p}_{t,j}} \left| \left(\Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) \right) \beta_{t,j} \right| + \sum_{\substack{j: \psi_{t,j,1} = \hat{p}_{t,j}, \\ \psi_{t,j,2} > \hat{p}_{t,j}}} \left| \left(\Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j,2}^{\text{clip}_2}(x, \tilde{x}) \right) \beta_{t,j} \right| \\
& \quad + \sum_{j: \psi_{t,j,2} = \hat{p}_{t,j}} \left| \left(\Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) \right) \beta_{t,j} \right| \\
& \stackrel{(a)}{\leq} \sum_{j: \psi_{t,j,1} > \hat{p}_{t,j}} \left| \left(\Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) \right) \beta_{t,j} \right| + \sum_{\substack{j: \psi_{t,j,1} = \hat{p}_{t,j}, \\ \psi_{t,j,2} > \hat{p}_{t,j}}} \left| \left(\Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j,2}^{\text{clip}_2}(x, \tilde{x}) \right) \beta_{t,j} \right| \\
& \quad + \sum_{j: \psi_{t,j,2} = \hat{p}_{t,j}} |x^\top (\theta_j^* - \hat{\theta}_{t,j}) \beta_{t,j}| + |\tilde{x}^\top (\theta_j^* - \hat{\theta}_{t,j}) \beta_{t,j}| \tag{I.9} \\
& \stackrel{(b)}{\leq} \sum_{j: \psi_{t,j,1} > \hat{p}_{t,j}} 4 \cdot \frac{5}{2} \cdot \frac{5\sqrt{2} L_{\max}}{2 T_t} \ln \frac{2}{\delta_{d,T_t}} \\
& \quad + \sum_{\substack{j: \psi_{t,j,1} = \hat{p}_{t,j}, \\ \psi_{t,j,2} > \hat{p}_{t,j}}} \min \left\{ 4, 2 \cdot 5 \sqrt{\frac{d}{T_{t,j}} \ln \frac{2}{\delta_{m,T_t}}} \right\} \cdot \frac{5}{2} \cdot \frac{5\sqrt{2} L_{\max}}{2 T_t} \ln \frac{2}{\delta_{d,T_t}} \\
& \quad + \sum_{j: \psi_{t,j,3} = \hat{p}_{t,j}} 2 \cdot 5 \sqrt{\frac{d}{T_{t,j}} \ln \frac{2}{\delta_{m,T_t}}} \cdot \frac{5}{2} \sqrt{\frac{8\hat{p}_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \\
& \leq \sum_{j: \psi_{t,j,1} > \hat{p}_{t,j}} 25\sqrt{2} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \\
& \quad + \sum_{\substack{j: \psi_{t,j,1} = \hat{p}_{t,j}, \\ \psi_{t,j,2} > \hat{p}_{t,j}}} \min \left\{ 4, 10 \sqrt{\frac{d}{T_{t,j}} \ln \frac{2}{\delta_{m,T_t}}} \right\} \cdot \frac{25\sqrt{2} L_{\max}}{4 T_t} \ln \frac{2}{\delta_{d,T_t}} \\
& \quad + \sum_{j: \psi_{t,j,2} = \hat{p}_{t,j}} 50\sqrt{2} \frac{\sqrt{d} L_{\max}}{T_t} \ln \frac{2}{\delta_{m,T_t}} \tag{I.10} \\
& \leq 50\sqrt{2} \frac{N L_{\max}}{T_t} \ln \frac{2}{\delta_{m,T_t}} = 2\xi_t
\end{aligned}$$

where (a) adopts (I.7) and (b) is obtained by the following derivations:

(1) $\psi_{t,j,1} > \hat{p}_{t,j}$ indicates $\hat{p}_{t,j} < \frac{d}{4T_t} \ln \frac{2}{\delta_{m,T_t}} < \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}$, and thus $\beta_{t,j} \leq \frac{5}{2} \cdot \frac{5\sqrt{2} L_{\max}}{2 T_t} \ln \frac{2}{\delta_{d,T_t}}$. In addition,⁶

$$\left| \Delta_j(x, \tilde{x}) - \hat{\Delta}_{t,j}^{\text{clip}_2}(x, \tilde{x}) \right| \leq 4.$$

(2) $\psi_{t,j,1} = \hat{p}_{t,j}$ indicates $\hat{p}_{t,j} \leq \frac{d}{4T_t} \ln \frac{2}{\delta_{m,T_t}}$ and $\psi_{t,j,2} > \hat{p}_{t,j}$ implies $\hat{p}_{t,j} < \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}$, thus (I.6) can be applied and $\beta_{t,j} \leq \frac{5}{2} \cdot \frac{5\sqrt{2} L_{\max}}{2 T_t} \ln \frac{2}{\delta_{d,T_t}}$.

⁵The last inequality can be loose. When each context has been visited at least once, the first summation in (I.10) can be ignored. For empirical performance, (I.10) should be used, while ξ is adopted for theoretical guarantees.

⁶This demonstrates the usefulness of the clipping technique. Without the clipping technique, we can only get an upper bound of $6d$ by utilizing (E.2) which introduces another d factor in ξ_t .

(3) $\psi_{t,j,2} = \hat{p}_{t,j}$ indicates $\hat{p}_{t,j} \geq \frac{25}{16} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \geq \frac{d}{4T_t} \ln \frac{2}{\delta_{m,T_t}}$, thus (I.6) can be applied and $\beta_{t,j} \leq \frac{5}{2} \sqrt{\frac{8\hat{p}_{t,j}L_{\max}}{T_t} \ln \frac{2}{\delta}}$.
Therefore, conditional on Good, with probability at least $1 - KN\delta_{m,T_t}$, $\text{RE} \leq 2\xi_t$. \square

Proof of Lemma G.5. By Lemma I.1, I.2 and I.3, conditional on Good, with probability at least $1 - (K\delta_{v,T_t} + N\delta_{d,T_t} + KN\delta_{m,T_t})$, $|\hat{\Delta}_t(x, \tilde{x}) - \Delta(x, \tilde{x})| \leq \rho_t(x, \tilde{x})$. This finishes the proof of the first result in Lemma G.5.

We then accumulates all the error probabilities. By the choices of $\delta_{v,T_t} = \frac{\delta}{15KT_t^3}$, $\delta_{d,T_t} = \frac{\delta}{15NT_t^3}$, $\delta_{m,T_t} = \frac{\delta}{15KNT_t^3}$,

$$\begin{aligned} \sum_{T_t=1}^{\infty} K\delta_{v,T_t} + N\delta_{d,T_t} + KN\delta_{m,T_t} &= \sum_{T_t=1}^{\infty} K \cdot \frac{\delta}{15KT_t^3} + N \cdot \frac{\delta}{15NT_t^3} + KN \cdot \frac{\delta}{15KNT_t^3} \\ &= \sum_{T_t=1}^{\infty} \frac{3\delta}{15T_t^3} \\ &\leq \frac{\delta}{4} \end{aligned}$$

Note that there is an reversion step at Line 18 Algorithm 1, for each $T \in \mathbb{N}$, there are at most two time steps $t_1 < t_2$ with $T_{t_1} = T_{t_2}$. Therefore,

$$\sum_{t=1}^{\infty} K\delta_{v,T_t} + N\delta_{d,T_t} + KN\delta_{m,T_t} \leq 2 \sum_{T_t=1}^{\infty} K\delta_{v,T_t} + N\delta_{d,T_t} + KN\delta_{m,T_t} \leq \frac{\delta}{2}$$

This proves the second statement. \square

J Upper Bound of $\text{PS}_{\varepsilon}\text{BAI}$: Proof of Theorem 3.2

Theorem 3.2 is proved in this section by three steps.

- Firstly, we show that the recommended arm \hat{x}_{ε} is an ε -best arm upon the termination of the algorithm in Lemma J.1.
- Secondly, we present a sufficient condition for the termination of the algorithm in terms of T_t in Lemma G.6.
- Lastly, we show that T_{τ} is bounded by a constant fraction of τ and thus, obtain an upper bound on τ .

Step 1:

According to Lemma G.5, CI_t holds with probability $1 - (K\delta_{v,T_t} + N\delta_{d,T_t} + KN\delta_{m,T_t})$. Conditional on the good event Good in (G.1) and the event CI_t (the mean gaps are well approximated at time step t), we expect the recommended arm will be an ε -optimal arm. This is formalized in the following lemma.

Lemma J.1. *Conditional on Good and CI_t , if the algorithm stops at time step t , the recommended arm $\hat{x}_{\varepsilon} = x_t^*$ is an ε -best arm.*

Proof of Lemma J.1. If $\hat{x}_{\varepsilon} = x^*$, the lemma holds.

If $\hat{x}_{\varepsilon} \neq x^*$, according to the termination condition (3.4)

$$\min_{x: x \neq \hat{x}_{\varepsilon}} \hat{\Delta}_t(\hat{x}_{\varepsilon}, x) - \rho_t(\hat{x}_{\varepsilon}, x) > -\varepsilon$$

we have

$$\Delta(\hat{x}_{\varepsilon}, x^*) \geq \hat{\Delta}_t(\hat{x}_{\varepsilon}, x^*) - \rho_t(\hat{x}_{\varepsilon}, x^*) \geq \min_{x: x \neq \hat{x}_{\varepsilon}} \hat{\Delta}_t(\hat{x}_{\varepsilon}, x) - \rho_t(\hat{x}_{\varepsilon}, x) > -\varepsilon$$

where the first inequality holds due to CI_t . This indicates $\Delta(x^*, \hat{x}_{\varepsilon}) \leq \varepsilon$ and thus the recommended arm $\hat{x}_{\varepsilon} = x_t^*$ is an ε -best arm. \square

Step 2:

Lemma G.6. *Conditional on Good and CI, the recommended arm $\hat{x}_\varepsilon \in \mathcal{X}_\varepsilon$ and when Algorithm 1 terminates, the order of T_t is upper bounded by (3.6).*

Proof of Lemma G.6. By Lemma J.1, $\hat{x}_\varepsilon \in \mathcal{X}_\varepsilon$. Note that for any $x \neq \hat{x}_\varepsilon$, when

$$\begin{cases} 2\rho_t(x, \hat{x}_\varepsilon) \leq \Delta(x, \hat{x}_\varepsilon) + \varepsilon, & \hat{x}_\varepsilon \neq x^* \text{ and } x = x^*, \\ \rho_t(\hat{x}_\varepsilon, x) + \rho_t(x^*, x) \leq \Delta(x^*, x) + \varepsilon, & \text{otherwise.} \end{cases} \quad (\text{J.1})$$

By utilizing CI_t, we have

$$\begin{cases} \hat{\Delta}_t(\hat{x}_\varepsilon, x^*) - \rho_t(\hat{x}_\varepsilon, x^*) \geq \hat{\Delta}_t(x^*, \hat{x}_\varepsilon) - \rho_t(\hat{x}_\varepsilon, x^*) \geq \Delta_t(x^*, \hat{x}_\varepsilon) - 2\rho_t(x^*, \hat{x}_\varepsilon) \geq -\varepsilon, \\ \hat{x}_\varepsilon \neq x^* \text{ and } x = x^*, \\ \hat{\Delta}_t(\hat{x}_\varepsilon, x) - \rho_t(\hat{x}_\varepsilon, x) \geq \hat{\Delta}_t(x^*, x) - \rho_t(\hat{x}_\varepsilon, x) \geq \Delta(x^*, x) - \rho_t(x^*, x) - \rho_t(\hat{x}_\varepsilon, x) \geq -\varepsilon, \\ \text{otherwise.} \end{cases}$$

Therefore, if (J.1) hold for all $x \neq \hat{x}_\varepsilon$, the algorithm must have stopped, i.e., it is sufficient to have the following for any $x \neq \hat{x}_\varepsilon$:

$$\begin{cases} \rho_t(x^*, \hat{x}_\varepsilon) \leq \frac{\Delta(x^*, \hat{x}_\varepsilon) + \varepsilon}{2}, & \hat{x}_\varepsilon \neq x^* \text{ and } x = x^*, \\ \rho_t(\hat{x}_\varepsilon, x) \leq \frac{\Delta(x^*, x) + \varepsilon}{2}, \rho_t(x^*, x) \leq \frac{\Delta(x^*, x) + \varepsilon}{2}, & \text{otherwise.} \end{cases}$$

Lemma J.2 gives an upper bound on the total number of time steps in Exp phases at which (J.1) must hold for any two arms $x_\varepsilon \in \mathcal{X}_\varepsilon, x \in \mathcal{X} \setminus \{x_\varepsilon\}$.

Lemma J.2. *For any two arms (x_ε, x) , where $x \neq x_\varepsilon, x^*$, when*

$$\begin{aligned} T_t &= \frac{6400 \ln 6400 \cdot d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{Kd / (\Delta(x^*, x) + \varepsilon)^2}{\delta} \\ &\quad + 6400 \ln 6400 \cdot \text{H}_{\text{DE}}(x_\varepsilon, x) \ln \frac{N\text{H}_{\text{DE}}(x_\varepsilon, x)}{\delta} \\ &\quad + \frac{3200\sqrt{2} \ln 3200\sqrt{2} \cdot NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{KNL_{\max}}{\Delta(x^*, x) + \varepsilon} \\ &= \tilde{O} \left(\frac{d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{1}{\delta} + \text{H}_{\text{DE}}(x_\varepsilon, x) \ln \frac{1}{\delta} + \frac{NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{1}{\delta} \right) \end{aligned} \quad (\text{J.2})$$

where

$$\text{H}_{\text{DE}}(x_\varepsilon, x) := \frac{L_{\max}}{(\Delta(x^*, x) + \varepsilon)^2} \left(\sum_{j=1}^N \sqrt{\min \left\{ 16p_j, \frac{1}{4} \right\}} |\Delta_j(x_\varepsilon, x) + \varepsilon| \right)^2,$$

we must have

$$\rho_t(x_\varepsilon, x) \leq \frac{\Delta(x^*, x) + \varepsilon}{2}$$

By taking the maximum of (J.2) over all ε -best arms and the suboptimal arms, we get

$$\begin{aligned} \tilde{T} &:= \max_{x_\varepsilon \in \mathcal{X}_\varepsilon, x \neq x_\varepsilon, x^*} \frac{6400 \ln 6400 \cdot d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{Kd / (\Delta(x^*, x) + \varepsilon)^2}{\delta} \\ &\quad + 6400 \ln 6400 \cdot \text{H}_{\text{DE}}(x_\varepsilon, x) \ln \frac{N\text{H}_{\text{DE}}(x_\varepsilon, x)}{\delta} \\ &\quad + \frac{3200\sqrt{2} \ln 3200\sqrt{2} \cdot NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{KNL_{\max}}{\Delta(x^*, x) + \varepsilon} \\ &= \tilde{O} \left(\max_{x_\varepsilon \in \mathcal{X}_\varepsilon, x \neq x_\varepsilon, x^*} \frac{d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{1}{\delta} + \text{H}_{\text{DE}}(x_\varepsilon, x) \ln \frac{1}{\delta} + \frac{NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{1}{\delta} \right). \end{aligned} \quad (\text{J.3})$$

□

Proof of Lemma J.2. Let $c_v, c_d, c_m \in (0, 1)$ be constants, satisfying $c_v + c_d + 2c_m = 1$. By the definition of $\rho_t(x_\varepsilon, x)$ in (3.3), it is sufficient to have

$$\begin{cases} 2\alpha_t \leq c_v \frac{\Delta(x^*, x) + \varepsilon}{2} \\ \sum_{j=1}^N \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x_\varepsilon, x) + \zeta_t(x_\varepsilon, x)| \leq (c_d + c_m) \frac{\Delta(x^*, x) + \varepsilon}{2} \\ 2\xi_t \leq c_m \frac{\Delta(x^*, x) + \varepsilon}{2} \end{cases}$$

where we take $\zeta_t(x_\varepsilon, x) = \varepsilon$ for theoretical simplicity, but we adopt $\zeta_t(x, \tilde{x}) = \arg \min_{\zeta_t(x, \tilde{x}) \in \mathbb{R}} \sum_{j=1}^N \beta_{t,j} (\hat{\Delta}_{t,j}(x, \tilde{x}) + \zeta_t(x, \tilde{x}))^2 = -\frac{\sum_{j=1}^N \beta_{t,j} \hat{\Delta}_{t,j}(x, \tilde{x})}{\sum_{j=1}^N \beta_{t,j}}$ in the algorithm, which still enjoys the theoretical guarantee.

VE Term According to the definition of α_t , we need to bound:

$$\begin{aligned} \alpha_t &\leq c_v \frac{\Delta(x^*, x) + \varepsilon}{4} \\ \Leftrightarrow 5\sqrt{\frac{d}{T_t} \ln \frac{2}{\delta_{v, T_t}}} &\leq c_v \frac{\Delta(x^*, x) + \varepsilon}{4} \end{aligned}$$

By a coarse estimation, it is sufficient to have

$$T_t \geq \frac{\bar{c}_v d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{Kd / (\Delta(x^*, x) + \varepsilon)^2}{\delta} = O\left(\frac{d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{Kd / (\Delta(x^*, x) + \varepsilon)^2}{\delta}\right)$$

where $\bar{c}_v = \frac{6400}{c_v^2} \ln \frac{6400}{c_v^2}$.

DE Term The difficulty lies at the DE term.

$$\begin{aligned} &\sum_{j=1}^N \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x_\varepsilon, x) + \zeta_t(x_\varepsilon, x)| \\ &\leq \sum_{j=1}^N \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x_\varepsilon, x) + \varepsilon| \\ &\leq \sum_{j:\psi_{t,j} > \hat{p}_{t,j}} \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x_\varepsilon, x) + \varepsilon| \\ &\quad + \sum_{j:\psi_{t,j} = \hat{p}_{t,j}} \beta_{t,j} |\Delta_j(x_\varepsilon, x) + \varepsilon| + \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x_\varepsilon, x) - \Delta_j(x_\varepsilon, x)| \\ &\leq \sum_{j:\psi_{t,j} > \hat{p}_{t,j}} \beta_{t,j} (2 + \varepsilon) + \sum_{j:\psi_{t,j} = \hat{p}_{t,j}} \beta_{t,j} |\hat{\Delta}_{t,j}^{\text{clip}_2}(x_\varepsilon, x) - \Delta_j(x_\varepsilon, x)| \\ &\quad + \sum_{j:\psi_{t,j} = \hat{p}_{t,j}} \beta_{t,j} |\Delta_j(x_\varepsilon, x) + \varepsilon| \\ &\leq 2\xi_t + \sum_{j:\psi_{t,j} = \hat{p}_{t,j}} \beta_{t,j} |\Delta_j(x_\varepsilon, x) + \varepsilon| \end{aligned}$$

We will upper bound $2\xi_t$ by $c_m \frac{\Delta(x^*, x) + \varepsilon}{2}$ first; the second term will be included in the RE term and analyzed later. Therefore, it is sufficient to have

$$\sum_{j:\psi_{t,j} = \hat{p}_{t,j}} \beta_{t,j} |\Delta_j(x_\varepsilon, x) + \varepsilon| \leq c_d \frac{\Delta(x^*, x) + \varepsilon}{2} \quad (\text{J.4})$$

By the definition of $\beta_{t,j}$ in (I.2) and Lemma K.3, we get

$$\beta_{t,j} = \frac{5}{2} \sqrt{\frac{2\phi_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \leq \frac{5}{2} \sqrt{\frac{2 \min \left\{ 16p_j, \frac{1}{4} \right\} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \quad (\text{J.5})$$

Lemma K.3. *If $\psi_{t,j} = \hat{p}_{t,j}$, then*

$$\phi_{t,j} \leq \min \left\{ 16p_j, \frac{1}{4} \right\}.$$

Therefore, in order for (J.4) to hold, it is sufficient to have

$$\begin{aligned} & \sum_{j:\psi_{t,j,2}=\hat{p}_{t,j}} \frac{5}{2} \sqrt{\frac{2 \min \left\{ 16p_j, \frac{1}{4} \right\} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} |\Delta_j(x_\varepsilon, x) + \varepsilon| \leq c_d \frac{\Delta(x^*, x) + \varepsilon}{2} \\ \Leftrightarrow & \frac{5}{2} \sqrt{\frac{2L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \sum_{j:\psi_{t,j,2}=\hat{p}_{t,j}} \sqrt{\min \left\{ 16p_j, \frac{1}{4} \right\}} |\Delta_j(x_\varepsilon, x) + \varepsilon| \leq c_d \frac{\Delta(x^*, x) + \varepsilon}{2} \\ \Leftarrow & \frac{5}{2} \sqrt{\frac{2L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \sum_{j=1}^N \sqrt{\min \left\{ 16p_j, \frac{1}{4} \right\}} |\Delta_j(x_\varepsilon, x) + \varepsilon| \leq c_d \frac{\Delta(x^*, x) + \varepsilon}{2} \\ \Leftarrow & T_t \geq \bar{c}_d \text{H}_{\text{DE}}(x_\varepsilon, x) \ln \frac{\text{NH}_{\text{DE}}(x_\varepsilon, x)}{\delta} = O \left(\text{H}_{\text{DE}}(x_\varepsilon, x) \ln \frac{\text{NH}_{\text{DE}}(x_\varepsilon, x)}{\delta} \right) \end{aligned}$$

where $\bar{c}_d = \frac{100}{c_d^2} \ln \frac{100}{c_d^2}$.

RE Term By the definition of ξ_t in (I.2), it is sufficient to have

$$\begin{aligned} & 25\sqrt{2} \frac{NL_{\max}}{T_t} \ln \frac{2}{\delta_{m,T_t}} \leq c_m \frac{\Delta(x^*, x) + \varepsilon}{2} \\ \Leftarrow & T_t \geq \frac{\bar{c}_m \cdot NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{KNL_{\max}}{\delta} = O \left(\frac{NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{KNL_{\max}}{\delta} \right) \end{aligned}$$

where $\bar{c}_m = \frac{400\sqrt{2}}{c_m} \ln \frac{400\sqrt{2}}{c_m}$.

By taking $c_v = \frac{3}{4}$, $c_d = \frac{1}{8}$ and $c_m = \frac{1}{8}$, the upper bound on T_t can be concluded as

$$\begin{aligned} T_t &= \frac{6400 \ln 6400 \cdot d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{Kd / (\Delta(x^*, x) + \varepsilon)^2}{\delta} + 6400 \ln 6400 \cdot \text{H}_{\text{DE}}(x_\varepsilon, x) \ln \frac{\text{NH}_{\text{DE}}(x_\varepsilon, x)}{\delta} \\ &\quad + \frac{3200\sqrt{2} \ln 3200\sqrt{2} \cdot NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{KNL_{\max}}{\delta} \\ &= O \left(\frac{d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{Kd / (\Delta(x^*, x) + \varepsilon)^2}{\delta} \right. \\ &\quad \left. + \text{H}_{\text{DE}}(x_\varepsilon, x) \ln \frac{\text{NH}_{\text{DE}}(x_\varepsilon, x)}{\delta} + \frac{NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{KNL_{\max}}{\delta} \right) \\ &= \tilde{O} \left(\frac{d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{1}{\delta} + \text{H}_{\text{DE}}(x_\varepsilon, x) \ln \frac{1}{\delta} + \frac{NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{1}{\delta} \right). \end{aligned}$$

□

Step 3:

Theorem 3.2. *Define the context distribution estimation (DE) hardness parameter*

$$\text{H}_{\text{DE}}(x_\varepsilon, x) := \frac{L_{\max}}{(\Delta(x^*, x) + \varepsilon)^2} \bar{H}(x_\varepsilon, x)$$

where $\bar{H}(x_\varepsilon, x) := \left(\sum_{j=1}^N \sqrt{\min\{16p_j, 1/4\}} |\Delta_j(x_\varepsilon, x) + \varepsilon| \right)^2$. Under Assumption 1, with probability at least $1 - \delta$, $\text{PS}\varepsilon\text{BAI}$ identifies an ε -optimal arm and its sample complexity is

$$\tilde{O} \left(\underbrace{\max_{\substack{x_\varepsilon \in \mathcal{X}_\varepsilon \\ x \neq x_\varepsilon, x^*}} \frac{d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{1}{\delta}}_{T_V(x)} + \underbrace{H_{\text{DE}}(x_\varepsilon, x)}_{T_D(x_\varepsilon, x)} \ln \frac{1}{\delta} + \underbrace{\frac{NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{1}{\delta}}_{T_R(x)} \right). \quad (3.6)$$

Proof of Theorem 3.2. By Assumption 1, and the choice of the parameters, Lemma G.4 indicates Good is guaranteed with probability at least $1 - \frac{\delta}{2}$.

Note that some arm pulls are not counted in T_t :

- In every γ time steps, Algorithm 1 enters the CD phase to select a CD sample.
- When a changepoint is detected, Algorithm 1 steps into the CA phase, where $\frac{w}{2}$ arms are sampled in Algorithm 4.
- During the CA phase, we abandon $\frac{w\gamma}{2}$ samples.

Therefore, when t is large (or after the first stationary segment)

$$t \leq T_t \cdot \frac{\gamma}{\gamma - 1} \frac{L_{\min}}{L_{\min} - w\gamma - \frac{w}{2}} \leq T_t \cdot \frac{\gamma}{\gamma - 1} \frac{L_{\min}}{L_{\min} - \frac{3w\gamma}{2}}. \quad (\text{J.6})$$

As $\gamma \geq 2$ and $w\gamma$ are set to be upper bounded by $\frac{L_{\min}}{3}$, so the fractions above is upper bounded by an absolute constant 4, e.g., with $\gamma = 2$ and $w = \frac{L_{\min}}{6}$, the constant is 4. Therefore, t is of the same order as T_t . By Lemma G.6, (3.6) holds with probability $1 - \delta$. \square

K Analysis of $\text{PS}\varepsilon\text{BAI}$: Technical Lemmas

Lemma K.1. Given any $j \in [N]$ and $T_t \geq \frac{2L_{\max}}{9} \ln \frac{2}{\delta_{d, T_t}}$, $\tilde{\beta}_{t, j} \leq \beta_{t, j}$.

Proof of Lemma K.1. Recall

$$\phi_{t, j} = \min \left\{ 4 \max \left\{ \hat{p}_{t, j}, \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d, T_t}} \right\}, \frac{1}{4} \right\}$$

The proof is scheduled in 2 steps.

Step 1: Upper bound $\sqrt{\frac{2p_j(1-p_j)L_{\max}}{T_t} \ln \frac{2}{\delta_{d, T_t}}}$.

When $\frac{L_{\max}}{3T_t} \ln \frac{2}{\delta_{d, T_t}} \leq \sqrt{\frac{2p_j(1-p_j)L_{\max}}{T_t} \ln \frac{2}{\delta_{d, T_t}}}$, we have

$$\tilde{\beta}_{t, j} = \frac{5}{2} \sqrt{\frac{2p_j(1-p_j)L_{\max}}{T_t} \ln \frac{2}{\delta_{d, T_t}}}$$

As $p_j(1-p_j) \leq \frac{1}{4}$, $\forall p_j \in (0, 1)$, this naïve bound gives $\tilde{\beta}_{t, j} \leq \frac{5}{2} \sqrt{\frac{2 \cdot \frac{1}{4} \cdot L_{\max}}{T_t} \ln \frac{2}{\delta_{d, T_t}}}$.

If $T_{t, j} > 0$, according to Lemma K.2,

$$\begin{aligned} \tilde{\beta}_{t, j} &= \frac{5}{2} \sqrt{\frac{2p_j(1-p_j)L_{\max}}{T_t} \ln \frac{2}{\delta_{d, T_t}}} \\ &\leq \frac{5}{2} \sqrt{\frac{2p_j T_{t, j} L_{\max}}{T_t T_{t, j}} \ln \frac{2}{\delta_{d, T_t}}} \\ &\leq \frac{5}{2} \sqrt{\frac{2\hat{p}_{t, j} L_{\max}}{T_t} \cdot 4 \max \left\{ 1, \frac{25}{4} \frac{L_{\max}}{T_{t, j}} \ln \frac{2}{\delta_{d, T_t}} \right\} \ln \frac{2}{\delta_{d, T_t}}}, \end{aligned}$$

thus we have

$$\tilde{\beta}_{t,j} \leq \frac{5}{2} \sqrt{\frac{2\phi_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \quad (\text{K.1})$$

If $T_{t,j} = 0$, i.e., the latent context j has never been observed and $\hat{p}_{t,j} = 0$, we have

$$\begin{aligned} \tilde{\beta}_{t,j} &= \frac{5}{2} \sqrt{\frac{2p_j(1-p_j)L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \\ &\leq \frac{5}{2} \sqrt{\frac{2\tilde{\beta}_{t,j}(1-\tilde{\beta}_{t,j})L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \\ \Rightarrow \tilde{\beta}_j(t, \delta) &\leq \frac{\frac{25}{8} L_{\max} \ln \frac{2}{\delta_{d,T_t}}}{T_t + \frac{25}{8} L_{\max} \ln \frac{2}{\delta_{d,T_t}}} \leq \frac{25}{8} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \end{aligned}$$

Take the trivial bound into consideration and observe that $\phi_{t,j} = \min\{\frac{25L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}, \frac{1}{4}\}$ and that

$$\frac{25}{8} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \leq \frac{25\sqrt{2}L_{\max}}{2T_t} \ln \frac{2}{\delta_{d,T_t}} = \frac{5}{2} \sqrt{\frac{2\phi_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}}.$$

Thus, (K.1) also holds for $T_{t,j} = 0$.

To conclude, we have

$$\begin{aligned} \tilde{\beta}_{t,j} &= \frac{5}{2} \max \left\{ \frac{L_{\max}}{3T_t} \ln \frac{2}{\delta_{d,T_t}}, \sqrt{\frac{2p_j(1-p_j)L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \right\} \\ &\leq \frac{5}{2} \max \left\{ \frac{L_{\max}}{3T_t} \ln \frac{2}{\delta_{d,T_t}}, \sqrt{\frac{2\phi_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \right\} \end{aligned}$$

Step 2: Simplify the bound. Recall that

$$\phi_{t,j} = \min \left\{ 4 \max \left\{ \hat{p}_{t,j}, \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \right\}, \frac{1}{4} \right\}$$

As $T_t \geq \frac{2L_{\max}}{9} \ln \frac{2}{\delta_{d,T_t}}$, we have

$$\frac{L_{\max}}{3T_t} \ln \frac{2}{\delta_{d,T_t}} \leq \sqrt{\frac{2 \cdot \frac{1}{4} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}}$$

According to the definition of ϕ , if $\hat{p}_{t,j} \geq \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}$, we have

$$\frac{L_{\max}}{3T_t} \ln \frac{2}{\delta_{d,T_t}} \leq \frac{5\sqrt{2}L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \leq \sqrt{\frac{2 \cdot 4\hat{p}_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}}$$

If $\hat{p}_{t,j} \leq \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}$, we have

$$\frac{L_{\max}}{3T_t} \ln \frac{2}{\delta_{d,T_t}} \leq \frac{25\sqrt{2}L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} = \sqrt{\frac{2L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \cdot 25 \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}$$

Thus

$$\max \left\{ \frac{L_{\max}}{3T_t} \ln \frac{2}{\delta_{d,T_t}}, \sqrt{\frac{2\phi_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \right\} = \sqrt{\frac{2\phi_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}}$$

This gives us the desired result. \square

Lemma K.2. When $\tilde{\beta}_{t,j} = \frac{5}{4} \sqrt{\frac{2p_j(1-p_j)L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}}$ and $T_{t,j} > 0$, we have

$$\frac{p_j}{T_{t,j}} \leq \frac{4}{T_t} \max \left\{ 1, \frac{25}{16} \frac{L_{\max}}{T_{t,j}} \ln \frac{2}{\delta_{d,T_t}} \right\}$$

Proof of Lemma K.2.

$$\begin{aligned}
\frac{p_j}{T_{t,j}} &= \frac{p_j T_t}{T_{t,j} T_t} \\
&\leq \frac{\hat{p}_{t,j} T_t + \frac{5}{4} \sqrt{2p_j(1-p_j)L_{\max} T_t \ln \frac{2}{\delta_{d,T_t}}}}{T_{t,j} T_t} \\
&= \frac{1}{T_t} + \frac{5}{4} \sqrt{\frac{p_j}{T_{t,j}}} \cdot \sqrt{\frac{1}{T_{t,j}} - \frac{p_j}{T_{t,j}}} \cdot \sqrt{\frac{2L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \\
\Rightarrow \frac{p_j}{T_{t,j}} &\leq \frac{4}{T_t} \max \left\{ 1, \frac{25 L_{\max}}{16 T_{t,j}} \ln \frac{2}{\delta_{d,T_t}} \right\}
\end{aligned}$$

□

Lemma K.3. *If $\psi_{t,j} = \hat{p}_{t,j}$, then*

$$\phi_{t,j} \leq \min \left\{ 16p_j, \frac{1}{4} \right\}.$$

Proof of Lemma K.3. Recall to the definition of $\phi_{t,j}$ in (I.2) and $\psi_{t,j}$ in (I.8),

$$\phi_{t,j} = \min \left\{ 4 \max \left\{ \hat{p}_{t,j}, \frac{25 L_{\max}}{4 T_t} \ln \frac{2}{\delta_{d,T_t}} \right\}, \frac{1}{4} \right\}, \quad \psi_{t,j} = \max \left\{ \hat{p}_{t,j}, \frac{25 L_{\max}}{4 T_t} \ln \frac{2}{\delta_{d,T_t}} \right\},$$

we have

$$\hat{p}_{t,j} \geq \frac{25 L_{\max}}{4 T_t} \ln \frac{2}{\delta_{d,T_t}}, \quad \phi_{t,j} = \min \left\{ 4\hat{p}_{t,j}, \frac{1}{4} \right\}$$

We only need to upper bound $\hat{p}_{t,j}$.

By (I.4),

$$\begin{aligned}
\psi_{t,j} = \hat{p}_{t,j} &\leq p_j + \frac{5}{2} \max \left\{ \frac{L_{\max}}{3T_t} \ln \frac{2}{\delta_{d,T_t}}, \sqrt{\frac{2p_j(1-p_j)L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \right\} \\
\Rightarrow \psi_{t,j} &\leq p_j + \frac{5}{2} \max \left\{ \frac{4}{75} \psi_{t,j}, \sqrt{p_j(1-p_j) \frac{8}{25} \psi_{t,j}} \right\} \\
&= p_j + \max \left\{ \frac{4}{15} \psi_{t,j}, \sqrt{2p_j(1-p_j) \psi_{t,j}} \right\} \\
\Rightarrow \psi_{t,j} &\leq 4p_j
\end{aligned}$$

Thus $\hat{p}_{t,j} = \psi_{t,j} \leq 4p_j$, and $\phi_{t,j} \leq \min \{16p_j, \frac{1}{4}\}$. □

L Upper Bound of $\text{PS}_{\varepsilon}\text{BAI}^+$: Proof of Theorem 3.3

Theorem 3.3. *The $\text{PS}_{\varepsilon}\text{BAI}^+$ algorithm is (ε, δ) -PAC and its expected sample complexity is*

$$\tilde{O} \left(\min \left\{ \max_{x_{\varepsilon} \in \mathcal{X}_{\varepsilon}, x \neq x_{\varepsilon}, x^*} T_V(x) + T_D(x_{\varepsilon}, x) + T_R(x), T_V^N + T_D^N \right\} \right).$$

Proof. We let τ_1 denote the stopping time of $\text{PS}_{\varepsilon}\text{BAI}$, τ_2 denote the stopping time of $\text{N}_{\varepsilon}\text{BAI}$, and τ denote the stopping time of $\text{PS}_{\varepsilon}\text{BAI}^+$ when an identical sequence of samples and returns are applied to them. By the design of these algorithms, we have

$$\tau \leq \min\{\tau_1, \tau_2\}.$$

Part I: Prove that $\mathbb{P}(\hat{x}_{\varepsilon} \notin \mathcal{X}_{\varepsilon}) \leq \delta$.

Let event $\mathcal{E}_1 = 1\{\tau_1 < \tau_2\} \cap 1\{\text{PS}_{\varepsilon}\text{BAI} \text{ returns arm } x_{\varepsilon} \text{ when it terminates}\}$. Recall that \hat{x}_{ε} denotes the recommended arm of $\text{N}_{\varepsilon}\text{BAI}$ when it terminates and \hat{x}_{ε} denotes the recommended arm of

$\text{PS}_{\varepsilon}\text{BAI}^+$ when it terminates. The design of algorithms indicates that:

$$\hat{x}_{\varepsilon} = x_{\varepsilon} \text{ when } \mathcal{E}_1 \text{ occurs, and } \hat{x}_{\varepsilon} = \hat{x}_{\varepsilon} \text{ when } \mathcal{E}_1^c \text{ occurs.}$$

Moreover, with the performance guarantees of $\text{N}_{\varepsilon}\text{BAI}$ and $\text{PS}_{\varepsilon}\text{BAI}$ (shown in Proposition 3.1 and Theorem 3.2 individually), we have

$$\begin{aligned} \mathbb{P}[\hat{x}_{\varepsilon} \notin \mathcal{X}_{\varepsilon}] &= \mathbb{P}[\hat{x}_{\varepsilon} \notin \mathcal{X}_{\varepsilon} | \mathcal{E}_1] \cdot \mathbb{P}[\mathcal{E}_1] + \mathbb{P}[\hat{x}_{\varepsilon} \notin \mathcal{X}_{\varepsilon} | \mathcal{E}_1^c] \cdot \mathbb{P}[\mathcal{E}_1^c] \\ &= \mathbb{P}[x_{\varepsilon} \notin \mathcal{X}_{\varepsilon} | \mathcal{E}_1] \cdot \mathbb{P}[\mathcal{E}_1] + \mathbb{P}[\hat{x}_{\varepsilon} \notin \mathcal{X}_{\varepsilon} | \mathcal{E}_1^c] \cdot \mathbb{P}[\mathcal{E}_1^c] \\ &\leq \delta \cdot \mathbb{P}[\mathcal{E}_1] + \delta \cdot \mathbb{P}[\mathcal{E}_1^c] = \delta \cdot (\mathbb{P}[\mathcal{E}_1] + \mathbb{P}[\mathcal{E}_1^c]) = \delta. \end{aligned}$$

Part II: Derive the expected sample complexity of $\text{PS}_{\varepsilon}\text{BAI}^+$.

We first derive the conditional expectation of the stopping time of $\text{PS}_{\varepsilon}\text{BAI}$. By the same argument as in (J.6), $T_t \leq t \leq 4T_t$ for any t . Therefore, it is sufficient to upper bound $\mathbb{E}[T_{\tau_1}]$ and $\mathbb{E}[\tau_1]$ can be upper bounded by $4\mathbb{E}[T_{\tau_1}]$.

According to Lemma G.6 and Lemma G.6, when both events Good and CI occur, $\text{PS}_{\varepsilon}\text{BAI}$ terminates with T_{τ_1} satisfying $T_{\tau_1} \leq \tilde{T}$, where \tilde{T} is defined in (J.3) and as below:

$$\begin{aligned} \tilde{T} := & \max_{x_{\varepsilon} \in \mathcal{X}_{\varepsilon}, x \neq x_{\varepsilon}, x^*} \frac{6400 \ln 6400 \cdot d}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{Kd / (\Delta(x^*, x) + \varepsilon)}{\delta} \\ & + 6400 \ln 6400 \cdot \text{H}_{\text{DE}}(x_{\varepsilon}, x) \ln \frac{N\text{H}_{\text{DE}}(x_{\varepsilon}, x)}{\delta} \\ & + \frac{3200\sqrt{2} \ln 3200\sqrt{2} \cdot NL_{\max}}{\Delta(x^*, x) + \varepsilon} \ln \frac{KNL_{\max}}{\Delta(x^*, x) + \varepsilon}. \end{aligned}$$

Given time step t with $T_t \geq 2\tilde{T}$ (thus $t \geq 8\tilde{T}$), we will upper bound $\mathbb{P}[\tau_1 > t]$ in the following. We firstly upper bound $\mathbb{P}[T_{\tau_1} > T_t]$. Note that

$$\begin{aligned} \mathbb{P}[T_{\tau_1} > T_t] &= \mathbb{P}\left[T_{\tau_1} > T_t | T_{\tau_1} \geq \frac{T_t}{2} + 1\right] \mathbb{P}\left[T_{\tau_1} \geq \frac{T_t}{2} + 1\right] \\ &\quad + \mathbb{P}\left[T_{\tau_1} > T_t | T_{\tau_1} < \frac{T_t}{2} + 1\right] \mathbb{P}\left[T_{\tau_1} < \frac{T_t}{2} + 1\right] \\ &\leq \mathbb{P}\left[T_{\tau_1} > T_t | T_{\tau_1} \geq \frac{T_t}{2} + 1\right]. \end{aligned}$$

If $\text{PS}_{\varepsilon}\text{BAI}$ fails to terminate with $T_{\tau_1} \leq \frac{T_t}{2}$, it implies that event $\text{Good} \cap (\bigcap_{s \in \mathcal{I}} \text{CI}_s)$ fails, where $\mathcal{I} = \{s : \frac{T_t}{2} + 1 \leq T_s \leq T_t\}$. Hence,

$$\begin{aligned} \mathbb{P}[T_{\tau_1} > T_t] &\leq \mathbb{P}\left[T_{\tau_1} > T_t \mid T_{\tau_1} \geq \frac{T_t}{2} + 1\right] \\ &\leq \mathbb{P}\left[\text{Good}^c \cup \left(\bigcup_{s \in \mathcal{I}} \text{CI}_s^c\right)\right] \\ &\leq \mathbb{P}[\text{Good}^c] + \mathbb{P}\left[\text{Good} \cap \left(\bigcup_{s \in \mathcal{I}} \text{CI}_s^c\right)\right] \\ &\leq \mathbb{P}[\text{Good}^c] + \mathbb{P}\left[\left(\bigcup_{s \in \mathcal{I}} \text{CI}_s^c\right) \mid \text{Good}\right] \\ &\stackrel{(a)}{\leq} \frac{\delta}{2\tau^*} + 2 \sum_{T_s = T_t/2 + 1}^{T_t} (K\delta_{v, T_s} + N\delta_{d, T_s} + KN\delta_{m, T_s}) \\ &\stackrel{(b)}{\leq} \frac{\delta}{2\tau^*} + 2 \sum_{T_s = T_t/2 + 1}^{T_t} \frac{\delta}{5T_s^3} \leq \frac{\delta}{2\tau^*} + 2 \int_{T_t/2}^{T_t} \frac{\delta}{5x^3} dx = \frac{\delta}{2\tau^*} + \frac{3\delta}{5T_t^2}, \end{aligned}$$

where (a) is obtained by applying Lemma G.5 to time steps in $\mathcal{I} \cap \mathcal{T}_t$ (i.e., the Exp time steps in \mathcal{I}) and note that there are at most two time steps $t_1 < t_2$ with $T_{t_1} = T_{t_2}$ due to the reversion step; (b) is derived by substituting $\delta_{v,T_t} = \frac{\delta}{15KT_t^3}$, $\delta_{d,T_t} = \frac{\delta}{15NT_t^3}$, $\delta_{m,T_t} = \frac{\delta}{15KNT_t^3}$ which are used to define the confidence radius ρ in (3.3). By using the fact that $4T_{\tau_1} \geq \tau_1$, for any $T_t \geq 2\tilde{T}$, we have

$$\begin{aligned} \mathbb{P}[\tau_1 > 4T_t] &\leq \mathbb{P}[T_{\tau_1} > T_t] \leq \frac{\delta}{2\tau^*} + \frac{3\delta}{5T_t^2} \\ \Rightarrow \mathbb{P}[\tau_1 > 4T_t + i] &\leq \frac{\delta}{2\tau^*} + \frac{3\delta}{5T_t^2}, \quad i = 0, 1, 2, 3 \\ \Rightarrow \mathbb{P}[\tau_1 > t] &\leq \frac{\delta}{2\tau^*} + \frac{48\delta}{5(t-4)^2}, \quad \forall 8\tilde{T} \leq t \leq \tau^*. \end{aligned}$$

This indicates for $t \geq 8\tilde{T}$, the probability that $\text{PS}\varepsilon\text{BAI}$ does not stop after t time steps is $\frac{\delta}{2\tau^*} + \frac{48\delta}{5(t-4)^2}$. Therefore, by Tonelli's Theorem, we have

$$\begin{aligned} \mathbb{E}[\tau_1 | 8\tilde{T} < \tau_1 \leq \tau^*] &\leq \sum_{t=8\tilde{T}}^{\tau^*-1} \mathbb{P}[\tau_1 > t] \leq \sum_{t=8\tilde{T}}^{\tau^*-1} \left(\frac{\delta}{2\tau^*} + \frac{48\delta}{5(t-4)^2} \right) \\ &\leq 1 + \int_{8\tilde{T}-1}^{\tau^*-1} \frac{48\delta}{5(t-4)^2} dt \leq 2. \end{aligned} \quad (\text{L.1})$$

Next, we bound $\mathbb{E}\tau$, the expected sample complexity of $\text{PS}\varepsilon\text{BAI}^+$. $\mathbb{E}\tau$ can be decomposed as below:

$$\mathbb{E}\tau \leq 8\tilde{T} + \mathbb{E}[\tau | 8\tilde{T} < \tau \leq \tau^*] + \mathbb{E}[\tau | \tau > \tau^*]. \quad (\text{L.2})$$

Since $\tau = \min\{\tau_1, \tau_2\}$ and $\mathbb{P}[\min\{\tau_1, \tau_2\} > t] \leq \mathbb{P}[\tau_1 > t]$ for all t , we have

$$\begin{aligned} \mathbb{E}[\tau | 8\tilde{T} < \tau \leq \tau^*] &= \sum_{t=8\tilde{T}}^{\tau^*-1} \mathbb{P}[\tau > t] = \sum_{t=8\tilde{T}}^{\tau^*-1} \mathbb{P}[\min\{\tau_1, \tau_2\} > t] \\ &\leq \sum_{t=8\tilde{T}}^{\tau^*-1} \mathbb{P}[\tau_1 > t] = \mathbb{E}[\tau_1 | 8\tilde{T} < \tau_1 \leq \tau^*]. \end{aligned} \quad (\text{L.3})$$

Besides, since $\text{PS}\varepsilon\text{BAI}$ will terminate after τ^* time steps, i.e., $\mathbb{P}(\tau_1 \leq \tau^*) = 1$. We have

$$\mathbb{E}[\tau | \tau \geq \tau^*] = \sum_{t=\tau^*}^{\infty} \mathbb{P}[\tau \geq t] \leq 1 + \sum_{t=\tau^*+1}^{\infty} \mathbb{P}[\tau_2 \geq t] = 1 + \mathbb{E}[\tau_2 | \tau_2 > \tau^*]. \quad (\text{L.4})$$

Lemma F.1 indicates that $\mathbb{P}[\tau_2 \geq t] \leq \frac{\delta}{(\alpha-1)C_3(t/2)^2}$ for all $t \geq T_0$, where

$$T_0 = \frac{768(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2} \ln \frac{768KC_3(8L_{\max} + 25d)}{(\varepsilon + \Delta_{\min})^2 \delta}, \quad C_3 = \sum_{n=1}^{\infty} n^{-3}.$$

Since the order of T_0 is apparently larger than τ^* , we have $T_0 \leq \tau^*$. Then, with the same method as in the proof of Proposition 3.1, the conditional sample complexity of $\text{N}\varepsilon\text{BAI}$ can be bounded as follows:

$$\mathbb{E}[\tau_2 | \tau_2 > \tau^*] \leq \int_{\tau^*}^{+\infty} \mathbb{P}(\tau_2 \geq x) dx \leq \int_{\tau^*}^{+\infty} \frac{\delta}{(\alpha-1)C_3(x/2)^2} dx = \frac{\delta}{C_3\tau^*}. \quad (\text{L.5})$$

Substituting terms in (L.2) with (L.1), (L.3), (L.4) and (L.5), we have

$$\mathbb{E}\tau \leq 8\tilde{T} + 3 + \frac{\delta}{C_3\tau^*}.$$

Besides,

$$\mathbb{E}\tau \leq \mathbb{E}\tau_2 \stackrel{(a)}{\leq} T_0 + \frac{\delta}{(\alpha-1)(\alpha-2)(T_0/2)^{\alpha-2}},$$

Where (a) is shown in (F.1).

Altogether, we have

$$\mathbb{E}\tau \leq \min \left\{ 8\tilde{T} + 3 + \frac{\delta}{C_{3\tau^*}}, 8\tilde{T} + 3 + \frac{\delta}{C_{3\tau^*}} \right\}.$$

□

M Analysis of Lower Bound: Proof of Theorem 4.1

The dynamics for under our PSLB model is displayed in Dynamics 1.

Dynamics 1 Dynamics for piecewise-stationary linear bandits

- 1: The instance: $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$
 - 2: **while** the algorithm does not stop at time step t **do**
 - 3: **if** $t \in \mathcal{C}$ **then**
 - 4: The environment samples $\theta_{j_t}^* \sim P_\theta$
 - 5: **else**
 - 6: $\theta_{j_t}^* = \theta_{j_{t-1}}^*$ (the environment does not change)
 - 7: **end if**
 - 8: The agent samples an arm x_t based on the history up to time $t - 1$.
 - 9: The reward $Y_{t,x_t} = x_t^\top \theta_{j_t}^* + \eta_t$ is revealed to the agent.
 - 10: **end while**
 - 11: Recommend an ε -best arm \hat{x}_ε .
-

To derive the lower bound in Theorem 4.1, we investigate two environments different from the one defined in Section 2 (and as in Dynamics 1):

- Dynamics 2: the agent observes the index of current context j_t , and the environment reduces to contextual linear bandits; the definition of Dynamics 2 and the lower bound under it are detailed in Appendix M.1.
- Dynamics 3: the agent observes the changepoints in \mathcal{C} and context vector $\theta_{j_t}^*$'s, and hence she solely needs to estimate the distribution of contexts; the definition of Dynamics 3 and the lower bound under it are detailed in Appendix M.2.

We **first** derive a lower bound for (ε, δ) -BAI algorithms in Dynamics 2, that is, in contextual linear bandits.

Corollary M.1. *For any (ε, δ) -PAC algorithm π , there exists an instance $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$ with Dynamics 2 such that*

$$\mathbb{E}[\tau] \geq T_\varepsilon(\Lambda) \log \frac{1}{2.4\delta}.$$

$$\text{where } T_\varepsilon(\Lambda)^{-1} = \max_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \min_{\Lambda' \in \text{Alt}_\Theta(\Lambda)} \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \frac{(x^\top (\theta_j^* - \theta'_j))^2}{2}$$

is defined in Theorem 4.1. In addition, when $|\mathcal{X}_\varepsilon| = 1$,

$$T_\varepsilon(\Lambda) = \min_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \max_{x \neq x^*} \frac{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2}{(\Delta(x^*, x) + \varepsilon)^2}.$$

This lower bound generalizes the result of [16] to the linear bandit setting.

We **next** study Dynamics 3. In this setting, the agent solely needs to estimate the distribution of contexts P_θ with context samples. Once the agent obtains a good estimate of P_θ , she can identify an ε -optimal arm w.h.p. Hence, the lower bound on the complexity of an (ε, δ) -BAI algorithm is the product of the minimum length of a stationary segment and the minimum number of context samples/changepoints needed for distribution estimation.

Corollary M.2. *For any (ε, δ) -PAC algorithm π , there exists an instance $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$ with Dynamics 3 such that $\mathbb{E}\tau \geq c_{N_c}$, where c_{N_c} is the N_c^{th} changepoint in the changepoint sequence \mathcal{C} .*

Altogether, the sample complexity of an (ε, δ) -BAI algorithm in Dynamics 2 and 3 build up the lower bound in Theorem 4.1.

M.1 Lower Bound for ε -BAI in Contextual Linear Bandits

In this section, we consider a sub-problem where the index of the current context j_t is revealed at each time step t . The original piecewise-stationary problem becomes a contextual problem whose dynamics is presented in Algorithm 2. As more information is provided to the agent, the lower

Dynamics 2 Dynamics for contextual linear bandits

- 1: The instance: $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$
 - 2: **while** the algorithm does not stop at time step t **do**
 - 3: **if** $t \in \mathcal{C}$ **then**
 - 4: The environment samples $\theta_{j_t}^* \sim P_\theta$
 - 5: **else**
 - 6: $\theta_{j_t}^* = \theta_{j_{t-1}}^*$ (the environment does not change)
 - 7: **end if**
 - 8: Reveal j_t to the agent.
 - 9: The agent samples an arm x_t based on the history up to time $t - 1$.
 - 10: The reward $Y_{t,x_t} = x_t^\top \theta_{j_t}^* + \eta_t$ is revealed to the agent.
 - 11: **end while**
 - 12: Recommend an ε -best arm \hat{x}_ε .
-

bound for this sub-problem is smaller than the one for the original problem.

Corollary M.1. *For any (ε, δ) -PAC algorithm π , there exists an instance $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$ with Dynamics 2 such that*

$$\mathbb{E}[\tau] \geq T_\varepsilon(\Lambda) \log \frac{1}{2.4\delta}.$$

$$\text{where } T_\varepsilon(\Lambda)^{-1} = \max_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \min_{\Lambda' \in \text{Alt}_\Theta(\Lambda)} \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \frac{(x^\top (\theta_j^* - \theta'_j))^2}{2}$$

is defined in Theorem 4.1. In addition, when $|\mathcal{X}_\varepsilon| = 1$,

$$T_\varepsilon(\Lambda) = \min_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \max_{x \neq x^*} \frac{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)}^2}{(\Delta(x^*, x) + \varepsilon)^2}.$$

For simplicity, we consider the noise model is the Clipped Gaussian Distribution $CN(1)$, i.e., $\eta \sim CN(1)$ or $P_\eta = CN(1)$.

Definition M.3. *A random variable x follows the Clipped Gaussian Distribution with parameter σ , denoted by $x \sim CN(\sigma)$, if it has the probability distribution function*

$$f(x) = \frac{1}{(2\Phi(\sigma) - 1) \cdot \sqrt{2\pi\sigma^2}} \exp\left(-\frac{x^2}{2\sigma^2}\right), \quad \forall x \in [-1, 1]$$

where $\Phi(x)$ is the cumulative distribution function of the standard Gaussian distribution.

Some notations are introduced here:

- Given an instance $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$, define the alternative instance $\Lambda' = (\mathcal{X}, \Theta', P_{\theta'}, \mathcal{C})$ with respect to Λ , where $\Theta' = (\theta'_1, \dots, \theta'_n) \in \mathbb{R}^{d \times N}$ and $P_{\theta'}[\theta'_j] = P_\theta[\theta_j^*]$, s.t. there $\exists x \in \mathcal{X} \setminus \mathcal{X}_\varepsilon, \forall x_\varepsilon \in \mathcal{X}_\varepsilon, \text{s.t.}, x_\varepsilon^\top \mathbb{E}_{\theta' \sim P_{\theta'}} < x^\top \mathbb{E}_{\theta \sim P_\theta} - \varepsilon$. We denote the set containing all the alternative instance (w.r.t. Λ) as $\text{Alt}_\Theta(\Lambda)$.
- $\mathcal{H}_t = (x_s, Y_{s,x_s}, j_s)_{s=1}^{t-1}$ is the observation history up to but not include time t .
- $N_{j,x}(t) = \sum_{s=1}^t \mathbb{1}\{x_s = x, j_s = j\}$ is the number of times arm in which x is sampled under context j . And $N_j(t) := \sum_{x \in \mathcal{X}} N_{j,x}(t)$ is the number of times in which context j appears.
- $\text{kl}(p, q) = \text{KL}(\text{Bern}(p), \text{Bern}(q))$ is the KL divergence between two Bernoulli distributions with parameters p and q .

Therefore, the probability of the observation history $\mathcal{H}_{t+1} = (x_s, Y_{s,x_s}, j_s)_{s=1}^t$ is

$$P_{\pi\Lambda} \left[(x_s, Y_{s,x_s}, j_s)_{s=1}^t \right] = \prod_{l=1}^{l_t} P_{\theta}[\theta_{j_{c_l}}^*] \prod_{s=0}^{L_l-1} \pi(x_{c_l+s} | \mathcal{H}_{c_l+s}) P_{\eta}(Y_{c_l+s, x_{c_l+s}} | x_{c_l+s}, \theta_{j_{c_l}}^*)$$

Then the log-likelihood between the two instance up to time t , given the observed data, is

$$\begin{aligned} & L_t \left[(x_s, Y_{s,x_s}, j_s)_{s=1}^t \right] \\ &= \log \frac{P_{\pi\Lambda} \left[(x_s, Y_{s,x_s}, j_s)_{s=1}^t \right]}{P_{\pi\Lambda'} \left[(x_s, Y_{s,x_s}, j_s)_{s=1}^t \right]} \\ &\stackrel{(a)}{=} \log \left(\frac{\prod_{l=1}^{l_t} P_{\theta}[\theta_{j_{c_l}}^*] \prod_{s=0}^{L_l-1} \pi(x_{c_l+s} | \mathcal{H}_{c_l+s}) P_{\eta}(Y_{c_l+s, x_{c_l+s}} | x_{c_l+s}, \theta_{j_{c_l}}^*)}{\prod_{l=1}^{l_t} P_{\theta'}[\theta_{j_{c_l}}'] \prod_{s=0}^{L_l-1} \pi(x_{c_l+s} | \mathcal{H}_{c_l+s}) P_{\eta}(Y_{c_l+s, x_{c_l+s}} | x_{c_l+s}, \theta_{j_{c_l}}')} \right) \\ &= \log \left(\frac{\prod_{l=1}^{l_t} \prod_{s=0}^{L_l-1} P_{\eta}(Y_{c_l+s, x_{c_l+s}} - x_{c_l+s}^{\top} \theta_{j_{c_l}}^*)}{\prod_{l=1}^{l_t} \prod_{s=0}^{L_l-1} P_{\eta}(Y_{c_l+s, x_{c_l+s}} - x_{c_l+s}^{\top} \theta_{j_{c_l}}')} \right) \\ &= \prod_{l=1}^{l_t} \sum_{s=0}^{L_l-1} \log \left(\frac{\exp(-\eta_{c_l+s}^2/2)}{\exp(-(\eta'_{c_l+s})^2/2)} \right) \\ &\stackrel{(b)}{=} \sum_{l=1}^{l_t} \sum_{s=0}^{L_l-1} \frac{-(Y_{c_l+s, x_{c_l+s}} - x_{c_l+s}^{\top} \theta_{j_{c_l}}^*)^2 + (Y_{c_l+s, x_{c_l+s}} - x_{c_l+s}^{\top} \theta_{j_{c_l}}')^2}{2} \\ &\stackrel{(c)}{=} \sum_{l=1}^{l_t} \sum_{s=0}^{L_l-1} \frac{x_{c_l+s}^{\top} \vartheta_{j_{c_l}} (2\eta_{c_l+s} + x_{c_l+s}^{\top} \vartheta_{j_{c_l}})}{2} \end{aligned}$$

where (a) utilizes $P_{\theta}[\theta_{j_{c_l}}^*] = P_{\theta'}[\theta_{j_{c_l}}']$, (b) makes use of the relationship between the arm and observation and in (c) $\vartheta_{j_{c_l}} := \theta_{j_{c_l}}^* - \theta_{j_{c_l}}'$. Thus, the expectation of the log-likelihood is (in the following, the expectation is taken under instance Λ and algorithm π)

$$\begin{aligned} \mathbb{E}[L_t] &= \mathbb{E} \left[\sum_{l=1}^{l_t} \sum_{s=0}^{L_l-1} \frac{x_{c_l+s}^{\top} \vartheta_{j_{c_l}} (2\eta_{c_l+s} + x_{c_l+s}^{\top} \vartheta_{j_{c_l}})}{2} \right] \\ &= \mathbb{E} \left[\sum_{l=1}^{l_t} \sum_{s=0}^{L_l-1} \frac{\vartheta_{j_{c_l}}^{\top} x_{c_l+s} x_{c_l+s}^{\top} \vartheta_{j_{c_l}}}{2} \right] \\ &= \frac{1}{2} \mathbb{E} \left[\sum_{l=1}^{l_t} \vartheta_{j_{c_l}}^{\top} \left(\sum_{s=0}^{L_l-1} x_{c_l+s} x_{c_l+s}^{\top} \right) \vartheta_{j_{c_l}} \right] \\ &= \frac{1}{2} \mathbb{E} \left[\sum_{x \in \mathcal{X}} \mathbb{1}\{x_{c_l+s} = x\} \sum_{j=1}^N \mathbb{1}\{j_{c_l} = j\} \sum_{l=1}^{l_t} \vartheta_{j_{c_l}}^{\top} \left(\sum_{s=0}^{L_l-1} x_{c_l+s} x_{c_l+s}^{\top} \right) \vartheta_{j_{c_l}} \right] \\ &= \frac{1}{2} \sum_{x \in \mathcal{X}} \sum_{j=1}^N \mathbb{E}[N_{j,x}(t)] \vartheta_j^{\top} x x^{\top} \vartheta_j \\ &= \frac{1}{2} \mathbb{E}[t] \sum_{j=1}^N \frac{\mathbb{E}[N_j(t)]}{\mathbb{E}[t]} \sum_{x \in \mathcal{X}} \frac{\mathbb{E}[N_{j,x}(t)]}{\mathbb{E}[N_j(t)]} \vartheta_j^{\top} x x^{\top} \vartheta_j \end{aligned} \tag{M.8}$$

As the lengths of all the stationary phases are upper bounded, i.e., $L_l \leq L_{\max}, \forall l \in \mathbb{N}$, then by Wald's Lemma,

$$\frac{\mathbb{E}[N_j(t)]}{\mathbb{E}[t]} = \frac{\mathbb{E}[t] P_{\theta}[\theta_j^*]}{\mathbb{E}[t]} = P_{\theta}[\theta_j^*] = p_j \tag{M.9}$$

The above also holds for $t = \tau$ where τ is a stopping time. According to Lemma 19 in [31],

$$\mathbb{E}[L_\tau] \geq \sup_{\mathcal{E} \in \mathcal{F}_\tau} \text{kl}(P_{\pi\Lambda}[\mathcal{E}], P_{\pi\Lambda'}[\mathcal{E}]) \quad (\text{M.10})$$

where $\mathcal{F}_\tau = \sigma(\mathcal{H}_{\tau+1})$. In addition, let $\mathcal{E} = \{\text{the recommended arm } \hat{x}_\varepsilon \notin \mathcal{X}_\varepsilon\}$, as the algorithm π is δ -PAC, then

$$\text{kl}(P_{\pi\Lambda}[\mathcal{E}], P_{\pi\Lambda'}[\mathcal{E}]) \geq \text{kl}(1 - \delta, \delta) \geq \log \frac{1}{2.4\delta} \quad (\text{M.11})$$

From (M.8), (M.9), (M.10) and (M.11), we conclude that

$$\begin{aligned} & \frac{1}{2} \mathbb{E}[\tau] \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} \frac{\mathbb{E}[N_{j,x}(\tau)]}{\mathbb{E}[N_j(\tau)]} \vartheta_j^\top x x^\top \vartheta_j \geq \log \frac{1}{2.4\delta} \\ \Rightarrow & \min_{\Lambda' \in \text{Alt}_\Theta(\Lambda)} \frac{1}{2} \mathbb{E}[\tau] \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} \frac{\mathbb{E}[N_{j,x}(\tau)]}{\mathbb{E}[N_j(\tau)]} \vartheta_j^\top x x^\top \vartheta_j \geq \log \frac{1}{2.4\delta} \\ \Rightarrow & \max_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \min_{\Lambda' \in \text{Alt}_\Theta(\Lambda)} \frac{1}{2} \mathbb{E}[\tau] \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \vartheta_j^\top x x^\top \vartheta_j \geq \log \frac{1}{2.4\delta} \\ \Leftrightarrow & \mathbb{E}[\tau] \geq T_\varepsilon(\Lambda) \log \frac{1}{2.4\delta} \end{aligned}$$

where

$$T_\varepsilon(\Lambda)^{-1} = \max_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \min_{\Lambda' \in \text{Alt}_\Theta(\Lambda)} \frac{1}{2} \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \vartheta_j^\top x x^\top \vartheta_j.$$

The solution to the above optimization problem is in general intractable, even for the stationary case [32]. We can establish a connection of the above problem to the stationary ε -best identification problem in linear bandits when we assume the change of the latent vectors are the same, i.e., $\vartheta_1 = \dots = \vartheta_N$.

M.1.1 Connection with the Stationary Case

In the alternative instance $\text{Alt}_\Theta(\Lambda)$, $\exists x \in \mathcal{X} \setminus \mathcal{X}_\varepsilon$, for any arm $x_\varepsilon \in \mathcal{X}_\varepsilon$, s.t.

$$\begin{aligned} & \sum_{j=1}^N p_j x_\varepsilon^\top \theta'_j + \varepsilon < \sum_{j=1}^N p_j x^\top \theta'_j \quad (\text{M.12}) \\ \Leftrightarrow & \sum_{j=1}^N p_j (x_\varepsilon - x)^\top \theta'_j + \varepsilon < 0 \\ \Leftrightarrow & - \sum_{j=1}^N p_j (x_\varepsilon - x)^\top \theta'_j + \sum_{j=1}^N p_j (x_\varepsilon - x)^\top \theta_j^* > \sum_{j=1}^N p_j (x_\varepsilon - x)^\top \theta_j^* + \varepsilon \\ \Leftrightarrow & \sum_{j=1}^N p_j (x_\varepsilon - x)^\top \vartheta_j > \sum_{j=1}^N p_j (x_\varepsilon - x)^\top \theta_j^* + \varepsilon = \Delta(x_\varepsilon, x) + \varepsilon \end{aligned}$$

Thus,

$$\begin{aligned} & \text{Alt}_\Theta(\Lambda) \\ & = \left\{ (\mathcal{X}, \Theta', P_\theta, \mathcal{C}) : \theta'_j = \theta_j^* - \vartheta_j, \exists x \in \mathcal{X} \setminus \mathcal{X}_\varepsilon, \sum_{j=1}^N p_j (x_\varepsilon - x)^\top \vartheta_j > \Delta(x_\varepsilon, x) + \varepsilon, \forall x_\varepsilon \in \mathcal{X}_\varepsilon \right\}. \end{aligned}$$

Define $\text{Alt}(\Lambda)_{\text{restricted}} \subset \text{Alt}_\Theta(\Lambda)$, with the additional constraint that $\vartheta_1 = \vartheta_2 = \dots = \vartheta_N$. Then we have

$$\text{Alt}(\Lambda)_{\text{restricted}}$$

$$= \left\{ (\mathcal{X}, \Theta', P_\theta, \mathcal{C}) : \theta'_j = \theta_j^* - \vartheta_1, \exists x \in \mathcal{X} \setminus \mathcal{X}_\varepsilon, \sum_{j=1}^N p_j (x_\varepsilon - x)^\top \vartheta_j > \Delta(x_\varepsilon, x) + \varepsilon, \forall x_\varepsilon \in \mathcal{X}_\varepsilon \right\}.$$

Note that in stationary linear bandits, the instance can be characterized by the arm set $\mathcal{X} \subset \mathbb{R}^d$ and the latent vector $\theta \in \mathbb{R}^d$. Define the alternative instance in linear bandits [32] for the instance $(\mathcal{X}, \mathbb{E}_{\theta \sim P_\theta} \theta)$

$$\begin{aligned} & \text{Alt}((\mathcal{X}, \mathbb{E}_{\theta \sim P_\theta} \theta))_{\text{stationary}} \\ &= \left\{ (\mathcal{X}, \theta') : \theta' = \mathbb{E}_{\theta \sim P_\theta} \theta - \vartheta_1, (x_\varepsilon - x)^\top \vartheta_1 > (x_\varepsilon - x)^\top \mathbb{E}_{\theta \sim P_\theta} \theta + \varepsilon, \forall x_\varepsilon \in \mathcal{X}_\varepsilon \right\}. \end{aligned}$$

Note that

$$\begin{aligned} & \max_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \min_{\Lambda' \in \text{Alt}_\Theta(\Lambda)} \frac{1}{2} \mathbb{E}[\tau] \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \vartheta_j^\top x x^\top \vartheta_j \\ & \leq \max_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \min_{\Lambda' \in \text{Alt}(\Lambda)_{\text{restricted}}} \frac{1}{2} \mathbb{E}[\tau] \vartheta_1^\top \left(\sum_{x \in \mathcal{X}} \left(\sum_{j=1}^N p_j v_{j,x} \right) x x^\top \right) \vartheta_1 \\ & \stackrel{(a)}{=} \max_{\bar{v} \in \Delta_{\mathcal{X}}} \min_{\Lambda' \in \text{Alt}(\Lambda)_{\text{restricted}}} \frac{1}{2} \mathbb{E}[\tau] \vartheta_1^\top \left(\sum_{x \in \mathcal{X}} \bar{v}_i x x^\top \right) \vartheta_1 \\ & = \max_{\bar{v} \in \Delta_{\mathcal{X}}} \min_{\Lambda' \in \text{Alt}((\mathcal{X}, \mathbb{E}_{\theta \sim P_\theta} \theta))_{\text{stationary}}} \frac{1}{2} \mathbb{E}[\tau] \vartheta_1^\top \left(\sum_{x \in \mathcal{X}} \bar{v}_i x x^\top \right) \vartheta_1 \end{aligned}$$

where in (a) we denote $\bar{v} := \sum_{j=1}^N p_j v_j$ as a mixture of $\{v_j\}_{j=1}^N$. In other words, the max-min problem becomes the one for the ε -best arm identification problem in stationary linear bandits. According to [32], the solution to the last optimization problem above is in general intractable.

However, the optimization problem can be simplified under some simple cases, e.g., the set of ε -best arm is a singleton [32]. In the next subsection, a lower bound for the original problem with \mathcal{X}_ε being a singleton will be derived.

M.1.2 A Simple Case: $|\mathcal{X}_\varepsilon| = 1$

Assume that the set of ε -best arm is a singleton, i.e., $\mathcal{X}_\varepsilon = \{x^*\}$, we will solve the original optimization problem:

$$\min_{\Lambda' \in \text{Alt}_\Theta(\Lambda)} \frac{1}{2} \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \vartheta_j^\top x x^\top \vartheta_j$$

we extend the procedures in [1] to the piecewise-stationary setup.

Lemma M.4.

$$\min_{\Lambda' \in \text{Alt}_\Theta(\Lambda)} \frac{1}{2} \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \vartheta_j^\top x x^\top \vartheta_j \geq \frac{1}{2} \min_{x \neq x^*} \frac{(\Delta(x^*, x))^2}{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)}^2}$$

Proof. Note that Λ' differs from Λ in the context matrix Θ . By the derivations in (M.12), there exists arm $x \neq x^*$, s.t.,

$$\sum_{j=1}^N p_j (x^* - x)^\top \vartheta_j > \Delta(x_\varepsilon, x) + \varepsilon$$

Therefore, the optimization problem becomes

$$\begin{aligned} & \min_{\vartheta_1, \dots, \vartheta_N} \frac{1}{2} \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \vartheta_j^\top x x^\top \vartheta_j \\ & \text{s.t. } \exists x \in \mathcal{X} \setminus \{x^*\}, \sum_{j=1}^N p_j (x^* - x)^\top \vartheta_j \geq \Delta(x^*, x) + \varepsilon + a =: \Delta_{a+\varepsilon}(x^*, x) \quad (\text{M.13}) \end{aligned}$$

where $a > 0$. The Lagrangian function for this problem is

$$L(\vartheta_1, \dots, \vartheta_N, \lambda) = \frac{1}{2} \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \vartheta_j^\top x x^\top \vartheta_j + \lambda \left(- \sum_{j=1}^N p_j (x^* - x)^\top \vartheta_j + \Delta_{a+\varepsilon}(x^*, x) \right)$$

Then

$$\frac{\partial L}{\partial \vartheta_j} = 0 \Rightarrow p_j \sum_{x \in \mathcal{X}} v_{j,x} x x^\top \vartheta_j - \lambda p_j (x^* - x) = 0 \Rightarrow A(v_j)^{1/2} \vartheta_j = \lambda A(v_j)^{-1/2} (x^* - x)$$

$$\frac{\partial L}{\partial \lambda} = 0 \Rightarrow \sum_{j=1}^N p_j (x^* - x)^\top \vartheta_j = \Delta_{a+\varepsilon}(x^*, x)$$

This indicates that

$$\begin{aligned} \sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \vartheta_j^\top x x^\top \vartheta_j &= \sum_{j=1}^N p_j \|\vartheta_j\|_{A(v_j)}^2 \\ &\geq \frac{\left(\sum_{j=1}^N p_j \|\vartheta_j\|_{A(v_j)} \|x^* - x\|_{A(v_j)^{-1}} \right)^2}{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2} \\ &= \frac{\left(\sum_{j=1}^N p_j \vartheta_j^\top (x^* - x) \right)^2}{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2} \\ &= \frac{(\Delta_{a+\varepsilon}(x^*, x))^2}{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2} \end{aligned}$$

Let $a \rightarrow 0$, we have

$$\sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} \vartheta_j^\top x x^\top \vartheta_j \geq \frac{(\Delta(x^*, x) + \varepsilon)^2}{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2}$$

Due to (M.13), we only require there exists $x \neq x^*$, such that the constraint is satisfied, therefore,

$$\frac{1}{2} \sum_{j=1}^N p_j \|\vartheta_j\|_{A(v_j)}^2 \geq \frac{1}{2} \min_{x \neq x^*} \frac{(\Delta(x^*, x) + \varepsilon)^2}{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2}$$

□

By Lemma M.4, the stopping time can be lower bounded as

$$\mathbb{E}[\tau] \geq 2 \log \frac{1}{2.4\delta} \min_{\{v_j \in \mathbf{\Delta}_x\}_{j=1}^N} \max_{x \neq x^*} \frac{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2}{(\Delta(x^*, x) + \varepsilon)^2}$$

This lower bound indicates that, (1) a good algorithm should actively detects and makes use of the contextual information to facilitate the arm identification process. (2) our lower bound extends the result of [16] to the linear bandits case.

The above lower bound can be further lower bounded if we restrict $v_1 = \dots = v_N$.

Lemma M.5. Let $\text{SPD}(d) := \{A : A \in \mathbb{R}^{d \times d}, A > 0\}$ denote the set of SPD matrices of dimension $d \times d$. Given any $x \in \mathbb{R}^d$, define the function $f : \text{SPD}(d) \rightarrow \mathbb{R}$, $f(A) = x^\top A^{-1} x$, then f is convex.

Proof of Lemma M.5. Given any $A, B \in \text{SPD}(d)$, define

$$g : \{t \in \mathbb{R} : A + tB \in \text{SPD}(d)\} \rightarrow \mathbb{R}, g(t) = x^\top (A + tB)^{-1} x$$

It is suffice to prove g is convex.

$$g'(t) = -x^\top (A + tB)^{-1} B (A + tB)^{-1} x$$

$$g''(t) = 2x^\top (A + tB)^{-1} B (A + tB)^{-1} B (A + tB)^{-1} x$$

$$= 2(x^\top (A + tB)^{-1} B) (A + tB)^{-1} (B(A + tB)^{-1} x) \geq 0$$

Therefore, g is convex. \square

Given Lemma M.5, we have

$$\begin{aligned} \sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2 &= (x^* - x)^\top \left(\sum_{j=1}^N p_j A(v_j)^{-1} \right) (x^* - x) \\ &\geq (x^* - x)^\top \left(\sum_{j=1}^N p_j A(v_j) \right)^{-1} (x^* - x) \\ &= (x^* - x)^\top \left(\sum_{j=1}^N p_j \sum_{x \in \mathcal{X}} v_{j,x} x x^\top \right)^{-1} (x^* - x) \\ &= (x^* - x)^\top \left(\sum_{x \in \mathcal{X}} \left(\sum_{j=1}^N p_j v_{j,x} \right) x x^\top \right)^{-1} (x^* - x) \\ &= (x^* - x)^\top A(\bar{v})^{-1} (x^* - x) \end{aligned}$$

where $\bar{v} := \sum_{j=1}^N p_j v_j \in \Delta_{\mathcal{X}}$ and the inequality becomes equality when $v_1 = \dots = v_n$. Thus,

$$\begin{aligned} \mathbb{E}[\tau] &\geq 2 \log \frac{1}{2.4\delta} \min_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \max_{x \neq x^*} \frac{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2}{(\Delta(x^*, x) + \varepsilon)^2} \\ &\geq 2 \log \frac{1}{2.4\delta} \min_{\bar{v} \in \Delta_{\mathcal{X}}} \max_{x \neq x^*} \frac{\|x^* - x\|_{A(\bar{v})^{-1}}^2}{(\Delta(x^*, x) + \varepsilon)^2}. \end{aligned} \quad (\text{M.14})$$

The last term mimics the lower bound in the stationary linear bandits with the latent vector $\mathbb{E}_{\theta \sim P_\theta} \theta$.

In addition, if we let $p_1 = 1$ and $p_j = 0$ for all $j \neq 1$ in the instance, our lower bound can be simplified to the lower bound in stationary linear bandits with latent vector θ_1^* [24]

$$\begin{aligned} \mathbb{E}[\tau] &\geq 2 \log \frac{1}{2.4\delta} \min_{\{v_j \in \Delta_{\mathcal{X}}\}_{j=1}^N} \max_{x \neq x^*} \frac{\sum_{j=1}^N p_j \|x^* - x\|_{A(v_j)^{-1}}^2}{(\Delta(x^*, x) + \varepsilon)^2} \\ &= 2 \log \frac{1}{2.4\delta} \min_{v_1 \in \Delta_{\mathcal{X}}} \max_{x \neq x^*} \frac{\|x^* - x\|_{A(v_1)^{-1}}^2}{(\Delta_1(x^*, x) + \varepsilon)^2}. \end{aligned} \quad (\text{M.15})$$

M.2 Lower Bound on the Number of changepoints

In this section, we consider an even easier problem: all the contextual information is known to the agent, except for the distribution P_θ . The dynamics is displayed in Dynamics 3.

Dynamics 3 Dynamics for an easier problem

- 1: The instance: $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$
 - 2: **while** the algorithm does not stop at time step t **do**
 - 3: **if** $t \in \mathcal{C}$ **then**
 - 4: The agent acknowledges $t \in \mathcal{C}$.
 - 5: The environment samples $\theta_{j_t}^* \sim P_\theta$
 - 6: **else**
 - 7: $\theta_{j_t}^* = \theta_{j_{t-1}}^*$ (the environment does not change)
 - 8: **end if**
 - 9: The agent observes $\theta_{j_t}^*$
 - 10: **end while**
 - 11: Recommend an ε -best arm \hat{x}_ε .
-

We are going to consider the alternative instances with respect to the distribution on θ , i.e., $\Lambda' = (\mathcal{X}, \Theta, P'_\theta, \mathcal{C})$. where $\exists x \in \mathcal{X} \setminus \mathcal{X}_\varepsilon, \forall x_\varepsilon \in \mathcal{X}_\varepsilon, s.t., x_\varepsilon^\top \mathbb{E}_{\theta \sim P'_\theta} < x^\top \mathbb{E}_{\theta \sim P'_\theta} - \varepsilon$. Denote the set of all the alternative instance (w.r.t. Λ) as $\text{Alt}_P(\Lambda)$. According to the Pinsker's inequality, let $\mathcal{E} := \{\text{the recommended arm } \hat{x}_\varepsilon \notin \mathcal{X}_\varepsilon\}$, then for any δ -PAC algorithm π ,

$$\begin{aligned}
& P_{\pi\Lambda}[\mathcal{E}] + P_{\pi\Lambda'}[\mathcal{E}^c] \geq \frac{1}{2} \exp(-\text{KL}(P_{\pi\Lambda}, P_{\pi\Lambda'})), \quad \forall \Lambda' \in \text{Alt}_P(\Lambda) \\
\Rightarrow & 4\delta \geq \exp(-\text{KL}(P_{\pi\Lambda}, P_{\pi\Lambda'})), \quad \forall \Lambda' \in \text{Alt}_P(\Lambda) \\
\Rightarrow & \text{KL}(P_{\pi\Lambda}, P_{\pi\Lambda'}) \geq \ln \frac{1}{4\delta}, \quad \forall \Lambda' \in \text{Alt}_P(\Lambda) \\
\Rightarrow & \min_{\Lambda' \in \text{Alt}_P(\Lambda)} \text{KL}(P_{\pi\Lambda}, P_{\pi\Lambda'}) \geq \ln \frac{1}{4\delta} \\
\Rightarrow & \mathbb{E}[l_t] \min_{\Lambda' \in \text{Alt}_P(\Lambda)} \text{KL}(P_\theta, P'_\theta) \geq \ln \frac{1}{4\delta} \\
\Rightarrow & \mathbb{E}[l_t] \geq \frac{1}{\min_{\Lambda' \in \text{Alt}_P(\Lambda)} \text{KL}(P_\theta, P'_\theta)} \ln \frac{1}{4\delta} \tag{M.16}
\end{aligned}$$

We will give an upper bound on the denominator. Note that

$$\min_{\Lambda' \in \text{Alt}_P(\Lambda)} \text{KL}(P_\theta, P'_\theta) = \min_{x \in \mathcal{X} \setminus \mathcal{X}_\varepsilon} \min_{\Lambda_x} \text{KL}(P_\theta, P_\theta^x) \tag{M.17}$$

where Λ_x is an alternative instance with distribution P_θ^x s.t. $x^\top \mathbb{E}_{\theta \sim P_\theta^x} \theta - \varepsilon > x_\varepsilon^\top \mathbb{E}_{\theta \sim P_\theta^x} \theta, \forall x_\varepsilon \in \mathcal{X}_\varepsilon$.

Given any $x \notin \mathcal{X}_\varepsilon$, we denote the shorthand notation $q_j = P_\theta^x[\theta_j^*]$. We have

$$\sum_{j=1}^N q_j \Delta_j(x, x_\varepsilon) \geq \varepsilon + a, \text{ for any } a > 0, x_\varepsilon \in \mathcal{X}_\varepsilon$$

Fix $a > 0$ which is sufficiently small, we have the following optimization problem:

$$\begin{aligned}
& \min_{q \in \Delta^N} \sum_{j=1}^N p_j \ln \frac{p_j}{q_j} \\
& s.t. \quad \sum_{j=1}^N q_j = 1 \\
& \quad \sum_{j=1}^N q_j \Delta_j(x_\varepsilon, x) + \varepsilon + a \leq 0, \forall x_\varepsilon \in \mathcal{X}_\varepsilon
\end{aligned}$$

If such alternative distribution q exists, let L denote the augmented Lagrangian function

$$L(q, \lambda, \{\lambda_{x_\varepsilon}\}_{x_\varepsilon \in \mathcal{X}_\varepsilon}) = \sum_{j=1}^N p_j \ln \frac{p_j}{q_j} + \lambda \left(\sum_{j=1}^N q_j - 1 \right) + \sum_{x_\varepsilon \in \mathcal{X}_\varepsilon} \lambda_{x_\varepsilon} (q_j \Delta_j(x_\varepsilon, x) + \varepsilon + a)$$

By the KKT conditions, we have:

$$\frac{\partial L}{\partial q_j} = \frac{-p_j}{q_j} + \lambda + \sum_{x_\varepsilon \in \mathcal{X}_\varepsilon} \lambda_{x_\varepsilon} \Delta_j(x_\varepsilon, x) = 0, \quad \forall j \in [N] \tag{M.18}$$

$$\frac{\partial L}{\partial \lambda} = \sum_{j=1}^N q_j - 1 = 0$$

$$\frac{\partial L}{\partial \lambda_{x_\varepsilon}} = \sum_{j=1}^N q_j \Delta_j(x_\varepsilon, x) + \varepsilon + a = 0, \quad \forall x_\varepsilon \in \mathcal{X}_\varepsilon$$

or $\lambda_{x_\varepsilon} = 0$ for some $x_\varepsilon \in \mathcal{X}_\varepsilon$, which indicates some conditions are not satisfied. The equations above give

$$q_j = \frac{p_j}{1 + \sum_{x_\varepsilon \in \mathcal{X}_\varepsilon} \lambda_{x_\varepsilon} (\Delta_j(x_\varepsilon, x) + \varepsilon + a)}, \quad \forall j \in [N] \tag{M.19}$$

By solving

$$\sum_{j=1}^N \frac{p_j}{1 + \sum_{x_\varepsilon \in \mathcal{X}_\varepsilon} \lambda_{x_\varepsilon} (\Delta_j(x_\varepsilon, x) + \varepsilon + a)} - 1 = 0 \quad (\text{M.20})$$

which is a polynomial with $N - 1$ degree. Thus, an explicit solution for any instance is not applicable. However, there are some cases where we can get an estimate of λ_{x_ε}

Lemma M.6. Denote $j_0 = \arg \min_{j \in [N]} \Delta_j(x^*, x)$ and $j_1 = \arg \max_{j \in [N]} \Delta_j(x^*, x)$. When $|\mathcal{X}_\varepsilon| = 1$, $-\Delta_{j_0}(x^*, x) - 2\varepsilon > \Delta_{j_1}(x^*, x) > \varepsilon$ and $\sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon + a)^3 \leq 0$, the solution to (M.20) is upper bounded by

$$\lambda_{i^*} \leq \frac{\Delta(x^*, x) + \varepsilon + a}{\sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon + a)^2}$$

Proof of Lemma M.6. When the ε -best arm set is a singleton, (M.20) becomes (where a is sufficiently small)

$$\sum_{j=1}^N \frac{p_j}{1 + \lambda_{x^*} (\Delta_j(x^*, x) + \varepsilon + a)} - 1 = 0$$

As (M.19) is a probability distribution, we need $1 + \lambda_{x^*} (\Delta_{j_0}(x^*, x) + \varepsilon + a) \geq 0$, which indicates $\lambda_{x^*} \leq \frac{-1}{\Delta_{j_0}(x^*, x) + \varepsilon + a}$. By the condition on j_0 and j_1 , $0 < \lambda_{x^*} (\Delta_j(x^*, x) + \varepsilon + a) < 1, \forall j \in [N]$.

Therefore, by using $\frac{1}{1+x} = 1 - x + x^2 - x^3 + \frac{x^4}{1+x}$ for $|x| < 1$, we can expand the equation as follows

$$\begin{aligned} 1 &= \sum_{j=1}^N \frac{p_j}{1 + \lambda_{x^*} (\Delta_j(x^*, x) + \varepsilon + a)} \\ &\geq \sum_{j=1}^N p_j \left((1 - \lambda_{x^*} (\Delta_j(x^*, x) + \varepsilon + a) + (\lambda_{x^*} (\Delta_j(x^*, x) + \varepsilon + a))^2 \right. \\ &\quad \left. - (\lambda_{x^*} (\Delta_j(x^*, x) + \varepsilon + a))^3 \right) \\ \Leftrightarrow 0 &\geq -\lambda_{x^*}^2 \sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon + a)^3 + \lambda_{x^*} \sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon + a)^2 \\ &\quad - \sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon + a) \\ \Rightarrow \lambda_{x^*} &\leq \frac{\sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon + a)}{\sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon + a)^2} = \frac{\Delta(x^*, x) + \varepsilon + a}{\sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon + a)^2}. \end{aligned}$$

□

In general, when we get the $\{\lambda_{x_\varepsilon}\}_{x_\varepsilon \in \mathcal{X}_\varepsilon}$ and plug it in (M.19), the alternative distribution q is obtained. Finally, we let $a \rightarrow 0$.

This gives the lower bound for the number of changepoints that need to be observed. A coarse estimation of the KL divergence without solving (M.20) can be done as follows

$$\begin{aligned} \sum_{j=1}^N p_j \ln \frac{p_j}{q_j} &= \sum_{j=1}^N p_j \ln \left(1 + \sum_{x_\varepsilon \in \mathcal{X}_\varepsilon} \lambda_{x_\varepsilon} (\Delta_j(x_\varepsilon, x) + \varepsilon) \right) \\ &\leq \sum_{j=1}^N p_j \sum_{x_\varepsilon \in \mathcal{X}_\varepsilon} \lambda_{x_\varepsilon} (\Delta_j(x_\varepsilon, x) + \varepsilon) \\ &= \sum_{x_\varepsilon \in \mathcal{X}_\varepsilon} \lambda_{x_\varepsilon} (\Delta(x_\varepsilon, x) + \varepsilon). \end{aligned}$$

Note that the solution of λ_{x_ε} depends on $\Delta(x_\varepsilon, x)$ (e.g. Lemma M.6), so the final solution is of order $(\Delta(x_\varepsilon, x) + \varepsilon)^2$ for a given $x \notin \mathcal{X}$. This is the solution to the inside minimization problem in (M.17). The final lower bound will be

$$\begin{aligned} \min_{\Lambda' \in \text{Alt}_P(\Lambda)} \text{KL}(P_\theta, P_{\theta'}) &= \min_{x \in \mathcal{X} \setminus \mathcal{X}_\varepsilon} \min_{\Lambda_x} \text{KL}(P_\theta, P_{\theta'}^x) \\ &\leq \min_{x \notin \mathcal{X}_\varepsilon} \sum_{x_\varepsilon \in \mathcal{X}_\varepsilon} \lambda_{x_\varepsilon} (\Delta(x_\varepsilon, x) + \varepsilon). \end{aligned}$$

The lower bound is

$$\mathbb{E}[l_t] \geq \max_{x \notin \mathcal{X}_\varepsilon} \frac{1}{\sum_{x_\varepsilon \in \mathcal{X}_\varepsilon} \lambda_{x_\varepsilon} (\Delta(x_\varepsilon, x) + \varepsilon)} \ln \frac{1}{4\delta}$$

where λ_{x_ε} is the solution to (M.20).

With the setup in Lemma M.6, we have

$$\begin{aligned} \min_{\Lambda' \in \text{Alt}_P(\Lambda)} \text{KL}(P_\theta, P_{\theta'}) &= \min_{x \in \mathcal{X} \setminus \mathcal{X}_\varepsilon} \min_{\Lambda_x} \text{KL}(P_\theta, P_{\theta'}^x) \\ &\leq \min_{x \neq x_\varepsilon} \frac{(\Delta(x^*, x) + \varepsilon)^2}{\sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon)^2} \end{aligned}$$

and the lower bound is

$$\mathbb{E}[l_t] \geq \max_{x \neq x_\varepsilon} \frac{\sum_{j=1}^N p_j (\Delta_j(x^*, x) + \varepsilon)^2}{(\Delta(x^*, x) + \varepsilon)^2} \ln \frac{1}{4\delta}.$$

Remark M.7. We give some comments on the existence and uniqueness of the solution to (M.20).

Existence:

- It is possible that (M.20) does not have a solution. For instance, consider a three-arm instance: $x_{(1)} = (1, 0.5), x_{(2)} = (0.5, 1), x_{(3)} = (0.6, 0.6), \theta_1^* = (1, 0), \theta_2^* = (0, 1), P_\theta = (0.5, 0.5), \varepsilon = 0.1$. We have $x_{(1)}^\top \mathbb{E}_{\theta \sim P_\theta} = x_{(2)}^\top \mathbb{E}_{\theta \sim P_\theta} = 0.75$ and $x_{(3)}^\top \mathbb{E}_{\theta \sim P_\theta} = 0.6$, thus $x_{(3)}$ is not an ε -best arm. Furthermore, there does not exist an alternative distribution q , such that $x_{(3)}$ is the best arm and neither $x_{(1)}, x_{(2)}$ is ε -best. Under such case, the lower bound on $\mathbb{E}[l_t]$ is 0 (we regard $\min_{x \in \emptyset} f(x) = +\infty$ by convention).
- It is possible that (M.20) does not have a solution and it is unnecessary to estimate P_θ . For instance, consider a two-arm instance: $x_{(1)} = (1, 0.5), x_{(2)} = (0.5, 0.1), \theta_1^* = (1, 0), \theta_2^* = (0, 1), P_\theta = (0.5, 0.5), \varepsilon = 0.1$. Arm $x_{(1)}$ is better than arm $x_{(2)}$ under all contexts. Therefore, no matter what P_θ is, arm $x_{(1)}$ is the ε -best arm. Under such cases, there may exist an algorithm and it is sufficient for it to determine the best arm if the context vectors are well-approximated. There is no need to estimate P_θ .
- A necessary condition for the existence of the solution of (M.20) is: there exists $x \notin \mathcal{X}_\varepsilon$, for any $x_\varepsilon \in \mathcal{X}_\varepsilon$, there $\exists j(x_\varepsilon) \in [N]$, s.t. $\Delta_j(x_\varepsilon, x) < -\varepsilon$. In other words, for each ε -best arm x_ε , there is at least one context in which the alternative arm x is better than x_ε by at least ε .
- A sufficient condition for the existence of the solution of (M.20) is: there exists $x \notin \mathcal{X}_\varepsilon$ and $j \in [N]$, s.t. $\Delta_j(x_\varepsilon, x) < -\varepsilon, \forall x_\varepsilon \in \mathcal{X}_\varepsilon$. In other words, x is better than any ε -best arm under context j . In the alternative instance, we can lift $P'_\theta[\theta_j^*]$ close to 1 so that arm x becomes the ε -best arm and $\mathcal{X}'_\varepsilon \cap \mathcal{X}_\varepsilon = \emptyset$.

Uniqueness:

- The uniqueness of the solution is not guaranteed, as the KKT conditions (M.18) is only a necessary condition for the solution. We need to look for the solution that minimizes the KL divergence.
- A sufficient condition for the uniqueness of the solution is (if it exists, which indicates $\exists j \in [N]$, s.t. $\Delta_j(x^*, x) < -\varepsilon$): $|\mathcal{X}_\varepsilon| = 1$. Specifically, denote $f(\lambda_{x^*}) := \sum_{j=1}^N \frac{p_j}{1 + \lambda_{x^*} (\Delta_j(x^*, x) + \varepsilon + a)}$, we have

$$\frac{\partial f}{\partial \lambda_{x^*}} = \sum_{j=1}^N \frac{-p_j (1 + \lambda_{x^*} (\Delta_j(x^*, x) + \varepsilon + a))}{1 + \lambda_{x^*} (\Delta_j(x^*, x) + \varepsilon + a)}$$

$$\frac{\partial^2 f}{\partial \lambda_{x^*}^2} = \sum_{j=1}^N \frac{2p_j(1 + \lambda_{x^*}(\Delta_j(x^*, x) + \varepsilon + a))^2}{1 + \lambda_{x^*}(\Delta_j(x^*, x) + \varepsilon + a)} > 0$$

As $f(0) = 1$, $\frac{\partial f}{\partial \lambda_{x^*}}(0) = -1$ and $\frac{\partial^2 f}{\partial \lambda_{x^*}^2} > 0$, so there is exactly 1 solution $\lambda_{x^*} > 0$.

N More Examples and Details

In this section, we **firstly** provide one more example to illustrate the tightness of our derived upper bound lower bounds, indicating the efficiency of our $\text{PS}\varepsilon\text{BAI}^+$ algorithm. In addition, we can observe how the upper and lower bounds are affected by the level of piecewise non-stationarity and whether our $\text{PS}\varepsilon\text{BAI}^+$ algorithm can reduce the influence manifested by L_{\max} . **After that**, we present a proof sketch of the results in Corollaries 5.1 and N.1 in Appendix N.1.

Example 2. Instance $\Lambda = (\mathcal{X}, \Theta, P_\theta, \mathcal{C})$ is with

- d arms: $x_{(i)} = \mathbf{e}_i, i \in [d]$.
- $N = d$ contexts: $\theta_1^* = (a, 0, 0, \dots, 0)^\top$, $\theta_2^* = (a, b, b, \dots, b)^\top - b\mathbf{e}_j, j \geq 2$, where $b > a > \varepsilon, b - a > \varepsilon$
- Context distribution: $p_j = p, j \geq 2$ and $p_1 = 1 - (N - 1)p$, where $p \in (0, \frac{a - \varepsilon}{(N - 2)b})$.

Under Example 2, we have

$$\begin{aligned} \Delta(x_{(1)}, x_{(i)}) &= (1 - (N - 2)p) \cdot a - (N - 2)p \cdot (b - a) > \varepsilon \\ \Delta_j(x_{(1)}, x_{(i)}) &= -b + a < -\varepsilon, \quad i \geq 2, j \neq 1, i \end{aligned}$$

Thus, (1) $x_{(1)}$ is the unique ε -best arm; (2) $\{x_{(i)}\}_{i \geq 2}$ are equivalent, and $\Delta_{\min} := \Delta(x_{(1)}, x_{(i)})$; (3) for any $i \geq 2$, $x_{(i)}$ can be an ε -optimal arm under some alternative distributions.

Corollary N.1. Firstly, for the instances defined in Example 2, we have

$$\text{H}_{\text{DE}} \leq \frac{16(N - 2)L_{\max}}{(\Delta_{\min} + \varepsilon)^2} \left((a + \varepsilon)^2 + (b - a - \varepsilon)^2 \right),$$

and the sample complexity of the $\text{PS}\varepsilon\text{BAI}^+$ is tight up to (NL_{\max}/L_{\min}) and logarithmic factors. We also further observe some specific instances: (i) when $p \rightarrow 0^+$, with $\Delta_{\min} = \min_{x \neq x^*} \Delta(x^*, x)$, we have

$$\begin{aligned} \frac{\mathbb{E}[\tau]^*}{\ln(1/\delta)} &\in \tilde{O} \left(\min \left\{ \frac{d}{(\Delta_{\min} + \varepsilon)^2} + \frac{NL_{\max}}{\Delta_{\min} + \varepsilon}, \frac{L_{\max} + d}{(\Delta_{\min} + \varepsilon)^2} \right\} \right) \\ &\quad \cap \Omega \left\{ \max \left\{ \frac{d}{(\Delta_{\min} + \varepsilon)^2}, \frac{L_{\min}(b - a - \varepsilon)}{\Delta_{\min} + \varepsilon} \right\} \right\}. \end{aligned}$$

(ii) When $p \rightarrow \left(\frac{a - \varepsilon}{(N - 2)b}\right)^-$ and $(a + \varepsilon)^2 + (b - a - \varepsilon)^2 = \Omega(1)$, we have $\text{H}_{\text{DE}} = \frac{NL_{\max}}{(\Delta(x_{(1)}, x_{(i)}) + \varepsilon)^2}$ and

$$\frac{\mathbb{E}[\tau]^*}{\ln(1/\delta)} \in \tilde{O} \left(\min \left\{ \text{H}_{\text{DE}}, \frac{d + L_{\max}}{\varepsilon^2} \right\} \right) \cap \Omega \left(\frac{d + L_{\min}}{\varepsilon^2} \right).$$

The upper bounds are achieved by the $\text{PS}\varepsilon\text{BAI}^+$ algorithm.

We can observe from Corollary N.1 that

- In case (i), when $p \rightarrow 0^+$, $p_1 \rightarrow 1 - \frac{(N - 1)(a - \varepsilon)}{(N - 2)b}$ and $p_j \rightarrow \left(\frac{a - \varepsilon}{(N - 2)b}\right)$, and thus the instance tends to be non-stationary and $\Delta_{\min} \rightarrow \varepsilon$. We will obtain

$$\frac{\mathbb{E}\tau}{\ln(1/\delta)} \in \tilde{\Theta} \left(\frac{d}{(\Delta_{\min} + \varepsilon)^2} \right),$$

indicating that our algorithm can also reduce the impact of L_{\max} .

- In case (ii), the upper and lower bounds are with the same order, and the difference is solely manifested by a additive term $L_{\max} - L_{\min}$, suggesting that $\text{PS}_{\varepsilon}\text{BAI}^+$ is near optimal and again, $\text{PS}_{\varepsilon}\text{BAI}^+$ mitigates the impact of L_{\max} .

N.1 Analysis of examples

Recall the instance $\Lambda = (\mathcal{X}, \Theta, P_{\theta}, \mathcal{C})$ in Example 1:

Example 1. Instance $\Lambda = (\mathcal{X}, \Theta, P_{\theta}, \mathcal{C})$ is with (i) $2d - 1$ arms: $x_{(1)} = \mathbf{e}_1, x_{(i)} = \mathbf{e}_i, x_{(d+i-1)} = \mathbf{e}_1 \cos \phi + \mathbf{e}_i \sin \phi$ for all $i \in \{2, \dots, d\}$ where $\phi \in [0, \pi/4]$, (ii) $2d - 2$ contexts: $\theta_{j\pm}^* = \mathbf{e}_1 \cos \phi \pm \mathbf{e}_{j+1} \sin \phi$ for all $j \in [d - 1]$, (iii) Context distribution: $p_j = 1/N$ for all $j \in [N]$.

Corollary 5.1. For the instance defined in Example 1, we have $\text{H}_{\text{DE}}(x_{\varepsilon}, x) = \tilde{O}(NL_{\max})$ for all $(x_{\varepsilon}, x) \in \mathcal{X}_{\varepsilon} \times (\mathcal{X} \setminus \mathcal{X}_{\varepsilon})$. In addition, if $\varepsilon < (\cos \phi)(1 - \cos \phi)$, we have

$$\frac{\mathbb{E}[\tau]^*}{\ln(1/\delta)} \in \tilde{\Theta} \left((1 + f(\phi)) \cdot \frac{d}{(\Delta_{x_{(1)}, x_{(d+1)}} + \varepsilon)^2} \right), \quad (5.1)$$

where $\mathbb{E}[\tau]^*$ is the minimal expected sample complexity over all (ε, δ) -PAC algorithms and $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfies $f(\phi) \rightarrow 0$ as $\phi \rightarrow 0^+$. The upper bound in (5.1) is achieved by $\text{PS}_{\varepsilon}\text{BAI}^+$.

Proof of Corollary 5.1. As $\mu_{x_{(1)}} = \cos \phi, \mu_{x_{(i)}} = 0, \mu_{x_{(d+i)}} = \cos^2 \phi$ for $i = 2, \dots, d$ and $\varepsilon \in (1 - \cos \phi, \cos \phi)$, $x_{(1)}$ is the best arm and $x_{(i)}, i = 2, \dots, d$ are not ε -best arms and $x_{(d+i)}, i = 2, \dots, d$ are ε -best arms. $\Delta_{\min} = \cos \phi - \cos^2 \phi$.

When $\varepsilon < \cos \phi - \cos^2 \phi = \Delta_{\min}$, $x_{(1)}$ is the unique ε -best arm. By solving (M.14), we see that there exists a continuous function $f(\phi)$ such that $f(\phi) \rightarrow 0$ as $\phi \rightarrow 0$ and we can get the lower bound for the VE term as

$$\Omega \left((1 + f(\phi)) \cdot \frac{d}{(\Delta(x_{(1)}, x_{(d+1)}) + \varepsilon)^2} \ln \frac{1}{\delta} \right). \quad (\text{N.1})$$

As $(x_{(1)} - x_{(i)})^{\top} \theta_{j\pm}^* > 0$ for $i = 2, \dots, d, j \in [d - 1]$, $x_{(i)}$ cannot be an ε -best arm under any alternative distribution, we only need to consider $x_{(d+i)}, i = 2, \dots, d$, which are equivalent. Given any $x_{(d+i)}$, by solving (M.20), we can upper bound

$$\lambda_{x_{(1)}} \leq \frac{\Delta_{\min} + \varepsilon + \sin^2 \phi}{-(\Delta_{i+}(x_{(1)}, x_{(d+i)}) + \varepsilon)(\Delta_{i-}(x_{(1)}, x_{(d+i)}) + \varepsilon)},$$

and thus the number of change points or context samples can be lower bounded as

$$\begin{aligned} \mathbb{E}[l_{\tau}] &\geq \frac{-(\Delta_{i+}(x_{(1)}, x_{(d+i)}) - \varepsilon)(\Delta_{i-}(x_{(1)}, x_{(d+i)}) + \varepsilon)}{(\Delta_{\min} + \varepsilon + \sin^2 \phi)(\Delta_{\min} + \varepsilon)} \\ &= \frac{(1 - \cos \phi - \varepsilon)(1 + \cos \phi - 2 \cos^2 \phi + \varepsilon)}{(\Delta_{\min} + \varepsilon + \sin^2 \phi)(\Delta_{\min} + \varepsilon)} = \frac{1 - \cos \phi - \varepsilon}{\Delta_{\min} + \varepsilon} = O(1). \end{aligned}$$

According to Theorem 3.3, the upper bound is

$$\begin{aligned} &\tilde{O} \left(\frac{d}{(\Delta(x_{(1)}, x_{(d+1)}) + \varepsilon)^2} \ln \frac{1}{\delta} + \text{H}_{\text{DE}}(\Delta(x_{(1)}, x_{(d+1)})) \ln \frac{1}{\delta} + \frac{NL_{\max}}{\Delta(x_{(1)}, x_{(d+1)}) + \varepsilon} \ln \frac{1}{\delta} \right) \\ &= \tilde{O} \left(\frac{d}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} + \text{H}_{\text{DE}}(\Delta(x_{(1)}, x_{(d+1)})) \ln \frac{1}{\delta} + \frac{NL_{\max}}{\Delta_{\min} + \varepsilon} \ln \frac{1}{\delta} \right) \end{aligned}$$

where

$$\text{H}_{\text{DE}}(\Delta(x_{(1)}, x_{(d+1)})) = \begin{cases} 16NL_{\max}, & \varepsilon \geq 1 - \cos \phi \\ \frac{16NL_{\max}(\Delta_{\min} + \varepsilon + 2(1 - \cos \phi)/N)^2}{(\Delta_{\min} + \varepsilon)^2}, & \varepsilon < 1 - \cos \phi \end{cases}$$

As $\Delta_{\min} = \cos \phi - \cos^2 \phi = \cos \phi(1 - \cos \phi)$, thus $1 - \cos \phi = \frac{\Delta_{\min}}{\cos \phi} \leq \sqrt{2}\Delta_{\min}$, and $\text{H}_{\text{DE}}(\Delta(x_{(1)}, x_{(d+1)})) < 144NL_{\max}$ for any choice of ε .

Hence, the upper bound is

$$\begin{aligned}\mathbb{E}[\tau] &= \tilde{O}\left(\frac{d}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} + NL_{\max} \ln \frac{1}{\delta} + \frac{NL_{\max}}{\Delta_{\min} + \varepsilon} \ln \frac{1}{\delta}\right) \\ &= \tilde{O}\left(\frac{d}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} + \frac{NL_{\max}}{\Delta_{\min} + \varepsilon} \ln \frac{1}{\delta}\right).\end{aligned}$$

As $\phi \rightarrow 0$, the first term (the VE term) dominates and it matches the lower bound in (N.1). Therefore, the upper bound is asymptotically tight up to logarithmic terms. \square

Corollary N.1. *Firstly, for the instances defined in Example 2, we have*

$$H_{\text{DE}} \leq \frac{16(N-2)L_{\max}}{(\Delta_{\min} + \varepsilon)^2} \left((a + \varepsilon)^2 + (b - a - \varepsilon)^2 \right),$$

and the sample complexity of the $\text{PS}\varepsilon\text{BAI}^+$ is tight up to (NL_{\max}/L_{\min}) and logarithmic factors. We also further observe some specific instances: (i) when $p \rightarrow 0^+$, with $\Delta_{\min} = \min_{x \neq x^*} \Delta(x^*, x)$, we have

$$\begin{aligned}\frac{\mathbb{E}[\tau]^*}{\ln(1/\delta)} &\in \tilde{O}\left(\min\left\{\frac{d}{(\Delta_{\min} + \varepsilon)^2} + \frac{NL_{\max}}{\Delta_{\min} + \varepsilon}, \frac{L_{\max} + d}{(\Delta_{\min} + \varepsilon)^2}\right\}\right) \\ &\quad \cap \Omega\left\{\max\left\{\frac{d}{(\Delta_{\min} + \varepsilon)^2}, \frac{L_{\min}(b - a - \varepsilon)}{\Delta_{\min} + \varepsilon}\right\}\right\}.\end{aligned}$$

(ii) When $p \rightarrow \left(\frac{a-\varepsilon}{(N-2)b}\right)^-$ and $(a + \varepsilon)^2 + (b - a - \varepsilon)^2 = \Omega(1)$, we have $H_{\text{DE}} = \frac{NL_{\max}}{(\Delta(x_{(1)}, x_{(i)}) + \varepsilon)^2}$ and

$$\frac{\mathbb{E}[\tau]^*}{\ln(1/\delta)} \in \tilde{O}\left(\min\left\{H_{\text{DE}}, \frac{d + L_{\max}}{\varepsilon^2}\right\}\right) \cap \Omega\left(\frac{d + L_{\min}}{\varepsilon^2}\right).$$

The upper bounds are achieved by the $\text{PS}\varepsilon\text{BAI}^+$ algorithm.

Proof of Corollary N.1. Under Example 2, we have

$$\begin{aligned}\Delta(x_{(1)}, x_{(i)}) &= (1 - (N-2)p) \cdot a - (N-2)p \cdot (b-a) > \varepsilon \\ \Delta_j(x_{(1)}, x_{(i)}) &= -b + a < -\varepsilon, \quad i \geq 2, j \neq 1, i\end{aligned}$$

Thus, (1) $x_{(1)}$ is the unique ε -best arm; (2) $\{x_{(i)}\}_{i \geq 2}$ are equivalent, and $\Delta_{\min} := \Delta(x_{(1)}, x_{(i)})$; (3) for any $i \geq 2$, $x_{(i)}$ can be an ε -best arm under some alternative distributions.

For any $i \geq 2$, by solving (M.20) with $x_\varepsilon = x_{(1)}$, $x = x_{(i)}$, we obtain $\lambda_{x_{(1)}} = \frac{\Delta(x_{(1)}, x_{(i)}) + \varepsilon}{(a + \varepsilon)(b - a - \varepsilon)}$ and the alternative distribution

$$P'_\theta[\theta_j^*] = \frac{p_j}{1 + \lambda_{x_{(1)}}(\Delta_j(x_{(1)}, x_{(i)}) + \varepsilon)}.$$

The lower bound on the expected number of changepoints is

$$\mathbb{E}[l_\tau] \geq \frac{(a + \varepsilon)(b - a - \varepsilon)}{(\Delta(x_{(1)}, x_{(i)}) + \varepsilon)^2} \ln \frac{4}{\delta}.$$

Thus the time complexity is lower bounded by

$$\mathbb{E}[\tau] \geq L_{\min} \cdot \frac{(a + \varepsilon)(b - a - \varepsilon)}{(\Delta(x_{(1)}, x_{(i)}) + \varepsilon)^2} \ln \frac{4}{\delta}.$$

Furthermore, by solving (M.14), we obtain the lower bound on the expected sample complexity when the context index j_t is revealed:

$$\begin{aligned}\mathbb{E}[\tau] &\geq 2 \log \frac{1}{2.4\delta} \min_{v \in \Delta_\kappa} \max_{i \neq 1} \frac{\frac{1}{v_1} + \frac{1}{v_i}}{(\Delta(x_{(1)}, x_{(i)}) + \varepsilon)^2} \geq 2 \log \frac{1}{2.4\delta} \frac{(\sqrt{d-1} + 1)^2}{(\Delta(x_{(1)}, x_{(i)}) + \varepsilon)^2} \\ &\geq \frac{d}{(\Delta(x_{(1)}, x_{(i)}) + \varepsilon)^2} \cdot 2 \log \frac{1}{2.4\delta}.\end{aligned}$$

We get the lower bound on the expected sample complexity:

$$\mathbb{E}[\tau] \geq \max \left\{ \frac{d}{(\Delta_{\min} + \varepsilon)^2} \cdot 2 \log \frac{1}{2.4\delta}, \frac{L_{\min}(a + \varepsilon)(b - a - \varepsilon)}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{4}{\delta} \right\}. \quad (\text{N.2})$$

According to Theorem 3.3, the upper bound on the expected sample complexity is

$$\begin{aligned} & \tilde{O} \left(\min \left\{ \frac{d}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} + H_{\text{DE}} \ln \frac{1}{\delta} + \frac{NL_{\max}}{\Delta_{\min} + \varepsilon} \ln \frac{1}{\delta}, \frac{L_{\max} + d}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} \right\} \right) \\ & \leq \tilde{O} \left(\min \left\{ \frac{d}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} + \frac{NL_{\max} \left((a + \varepsilon)^2 + (b - a - \varepsilon)^2 \right)}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} \right. \right. \\ & \quad \left. \left. + \frac{NL_{\max}}{\Delta_{\min} + \varepsilon} \ln \frac{1}{\delta}, \frac{L_{\max} + d}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} \right\} \right) \end{aligned} \quad (\text{N.3})$$

where we utilize

$$\begin{aligned} H_{\text{DE}} &= \frac{L_{\max}}{(\Delta_{\min} + \varepsilon)^2} \left(\sqrt{\min \left\{ 16p_j, \frac{1}{4} \right\}} |a + \varepsilon| + \sqrt{\min \left\{ 16p, \frac{1}{4} \right\}} |a + \varepsilon| \right. \\ & \quad \left. + (N - 2) \sqrt{\min \left\{ 16p, \frac{1}{4} \right\}} |b - a - \varepsilon| \right)^2 \\ & \leq \frac{16(N - 2)L_{\max}}{(\Delta_{\min} + \varepsilon)^2} \left((a + \varepsilon)^2 + (b - a - \varepsilon)^2 \right). \end{aligned} \quad (\text{N.4})$$

By comparing the lower bound in (N.2) and the upper bound (N.3), we conclude that

- the sample complexity of $\text{PS}\varepsilon\text{BAI}^+$ is tight up to $\frac{NL_{\max}}{L_{\min}}$ and logarithmic factors.
- When the mean gap is small, the sample complexity of $\text{PS}\varepsilon\text{BAI}^+$ is dominated by the former term, i.e., the design of $\text{PS}\varepsilon\text{BAI}$.
- When $p \rightarrow 0^+$, then $p_1 \rightarrow 1, p_j \rightarrow 0$ for $j \geq 2$, so the instance tends to be stationary and $\Delta_{\min} \rightarrow a$. The lower bound in (N.2) becomes

$$\mathbb{E}[\tau] \geq \max \left\{ \frac{d}{(\Delta_{\min} + \varepsilon)^2} \cdot 2 \log \frac{1}{2.4\delta}, \frac{L_{\min}(b - a - \varepsilon)}{\Delta_{\min} + \varepsilon} \ln \frac{4}{\delta} \right\}$$

and the upper bound in (N.3) turns into

$$\tilde{O} \left(\min \left\{ \frac{d}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} + \frac{NL_{\max}}{\Delta_{\min} + \varepsilon} \ln \frac{1}{\delta}, \frac{L_{\max} + d}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} \right\} \right).$$

The vector estimation term dominates the sample complexity and our upper bound is tight.

- When $p \rightarrow \left(\frac{a - \varepsilon}{(N - 2)b} \right)^-$, then $p_1 \rightarrow 1 - \frac{(N - 1)(a - \varepsilon)}{(N - 2)b}$ and $p_j \rightarrow \left(\frac{a - \varepsilon}{(N - 2)b} \right)$, so the instance tends to be non-stationary and $\Delta_{\min} \rightarrow \varepsilon$. The lower bound in (N.2) becomes

$$\mathbb{E}[\tau] \geq \max \left\{ \frac{d}{4\varepsilon^2} \cdot 2 \log \frac{1}{2.4\delta}, \frac{L_{\min}(a + \varepsilon)(b - a - \varepsilon)}{4\varepsilon^2} \ln \frac{4}{\delta} \right\}$$

and the upper bound in (N.3) turns into

$$\tilde{O} \left(\min \left\{ \frac{d}{4\varepsilon^2} \ln \frac{1}{\delta} + H_{\text{DE}} \cdot \ln \frac{1}{\delta} + \frac{NL_{\max}}{2\varepsilon} \ln \frac{1}{\delta}, \frac{L_{\max} + d}{4\varepsilon^2} \ln \frac{1}{\delta} \right\} \right)$$

where H_{DE} is upper bounded by (N.4).

- (1) If $\left((a + \varepsilon)^2 + (b - a - \varepsilon)^2 \right) = O(4\varepsilon^2 + NL_{\max})$, DE is independent of ε and VE increases as ε decreases, thus the expected sample complexity is dominated by vector estimation term when ε is small;
- (2) If $\left((a + \varepsilon)^2 + (b - a - \varepsilon)^2 \right) = \Omega(1)$, the expected sample complexity is dominated by

distribution estimation term.
Under both scenarios, our upper bound is tight. □

O Experimental Details

For the computation of the G-optimal allocation, we adopt the Wolfe–Atwood Algorithm with the Kumar–Yildirim start introduced in [33], where the input to the function is the arm set and the output is the G-optimal allocation. All experiments are conducted via MATLAB R2021b on a MacBook Pro with Apple M1 Pro chip and 16 GB memory.

To shorten the execution time,

- Before $\text{PS}_\varepsilon\text{BAI}$ stops, $\text{PS}_\varepsilon\text{BAI}$ and $\text{N}_\varepsilon\text{BAI}$ are conducted in parallel (i.e., we run $\text{PS}_\varepsilon\text{BAI}^+$). After $\text{PS}_\varepsilon\text{BAI}$ stops, $\text{N}_\varepsilon\text{BAI}$ continues and is run in a batch manner, i.e., we sample L_{\min} samples according to the G-optimal allocation one time and update the statistics. As the sample complexity of $\text{N}_\varepsilon\text{BAI}$ is of order at least 10^7 and each stationary segment is of order 10^4 , the effect of this batch sampling procedure can be largely ignored.
- $\text{D}_\varepsilon\text{BAI}$ and $\text{D}_\varepsilon\text{BAI}_\beta$ are both conducted in segments, because the latent context vector can be observed by the two algorithms and the latent vector does not change within a stationary segment.
- As the experiment in Section 6 illustrates that $\text{PS}_\varepsilon\text{BAI}$ outperforms $\text{N}_\varepsilon\text{BAI}$ and dominates the $\text{PS}_\varepsilon\text{BAI}^+$ algorithm, we run $\text{PS}_\varepsilon\text{BAI}$ instead of $\text{PS}_\varepsilon\text{BAI}^+$ for the addition experiments in Section O.2 and Section O.3.

To increase the robustness of the algorithm, the window size for LCD is doubled. Note that this will only influence an absolute constant in the sample complexity of the proposed algorithm and the order of the sample complexity remains.

O.1 Modification of Confidence Radii in $\text{PS}_\varepsilon\text{BAI}$ and $\text{PS}_\varepsilon\text{BAI}^+$

During the proof of the upper bound, we have relaxed the absolute constants in the confidence radii to simplify the proof and increase the readability. In the experiments, we utilize the tighter confidence radii to gain better empirical performance. Note that when these tighter confidence radii are utilized by our algorithm, it still enjoys the current theoretical guarantee. The choice of confidence radii are as follows:

- α_t : according to Lemma I.1 and Lemma E.1, α can be tightened to be

$$\alpha_t^{\text{alg}} = \frac{d}{T_t} \ln \frac{2}{\delta_{v,T_t}} + \sqrt{\left(\frac{d}{T_t} \ln \frac{2}{\delta_{v,T_t}}\right)^2 + \frac{4d}{T_t} \ln \frac{2}{\delta_{v,T_t}}} \leq 5\sqrt{\frac{d}{T_t} \ln \frac{2}{\delta_{v,T_t}}}.$$

- $\beta_{t,j}$: according to (I.3) and Lemma K.1, $\beta_{t,j}$ can be tightened to be

$$\beta_{t,j}^{\text{alg}} = \min \left\{ \frac{\frac{1}{3}L_{\max} \ln \frac{2}{\delta_{d,T_t}} + \sqrt{\left(\frac{1}{3}L_{\max} \ln \frac{2}{\delta_{d,T_t}}\right)^2 + \frac{2\phi_{t,j}L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}}}{T_t}, 1 \right\},$$

$$\text{where } \phi_{t,j} := \min \left\{ 4 \max \left\{ \hat{p}_{t,j}, \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \right\}, \frac{1}{4} \right\}.$$

- ξ_t : instead of using $25\sqrt{2}\frac{NL_{\max}}{T_t} \ln \frac{2}{\delta_{m,T_t}}$, we turn to bound the residual error by (I.9):

$$\begin{aligned} \xi_t^{\text{alg}} &= \sum_{j:\psi_{t,j,1} > \hat{p}_{t,j}} 4 \cdot \beta_{t,j}^{\text{alg}} + \sum_{\substack{j:\psi_{t,j,1} = \hat{p}_{t,j}, \\ \psi_{t,j,2} > \hat{p}_{t,j}}} \min \left\{ 4, 2\xi_{t,j}^{\text{alg}} \right\} \cdot \beta_{t,j}^{\text{alg}} \\ &\quad + \sum_{j:\psi_{t,j,3} = \hat{p}_{t,j}} 10\sqrt{\frac{d}{T_{t,j}} \ln \frac{2}{\delta_{m,T_t}}} \cdot \beta_{t,j}^{\text{alg}} \end{aligned}$$




Figure 3: Misspecified L_{\max} and L_{\min} .

where

$$\psi_{t,j,1} := \max \left\{ \hat{p}_{t,j}, \frac{d}{4T_t} \ln \frac{2}{\delta_{m,T_t}} \right\}, \quad \psi_{t,j,2} := \max \left\{ \hat{p}_{t,j}, \frac{25 L_{\max}}{4 T_t} \ln \frac{2}{\delta_{d,T_t}} \right\},$$

$$\tilde{\xi}_{t,j}^{\text{alg}} = \frac{d}{T_{t,j}} \ln \frac{2}{\delta_{m,T_t}} + \sqrt{\left(\frac{d}{T_{t,j}} \ln \frac{2}{\delta_{m,T_t}} \right)^2 + \frac{4d}{T_{t,j}} \ln \frac{2}{\delta_{m,T_t}}} \leq 5 \sqrt{\frac{d}{T_{t,j}} \ln \frac{2}{\delta_{m,T_t}}}.$$

$\tilde{\xi}_{t,j}^{\text{alg}}$ is obtained in a similar manner as α_t^{alg} . ξ_t^{alg} characterizes the confidence radii of each context at a finer level.

These finer confidence bounds can save a constant of 2.5 when t is large.

O.2 Misspecified L_{\max} and L_{\min}

As $\text{PS}\varepsilon\text{BAI}^+$ requires the knowledge of L_{\max} , we empirically test the robustness of the algorithm towards L_{\max} on the instances in section 6. We run $\text{PS}\varepsilon\text{BAI}$ with $\tilde{L}_{\min} = \nu L_{\min}$ and $\tilde{L}_{\max} = \nu L_{\max}$ where $\nu = 0.8$ or 1.2 . An ε -best arm is recommended in all experiments. The sample complexities are presented in Figure 3. The result indicates that $\text{PS}\varepsilon\text{BAI}$ (thus $\text{PS}\varepsilon\text{BAI}^+$) is robust towards the knowledge of L_{\max} and its superiority over $\text{N}\varepsilon\text{BAI}$ is maintained.

O.3 Robustness towards w and b

According to the distinguishability condition (Assumption 1), point (2) indicates we can set $w = \frac{L_{\min}}{3\gamma}$, where γ is the change detection frequency, and (3.5) indicates w and b are coupled. We denote $\tilde{w} := \frac{L_{\min}}{18}$.

To exam the robustness towards w and b , we choose to vary the choice of γ , thus, w and b will change accordingly. Specifically, we select $\gamma \in \{2, 3, 6, 12\}$ and the corresponding $w \in \{3\tilde{w}, 2\tilde{w}, \tilde{w}, \frac{\tilde{w}}{2}\}$. The other parameters remain unchanged. The experiment result is presented in Figure 4. When $\gamma = 18$, $w = \frac{\tilde{w}}{3}$, Assumption 1 is severely violated and the result is not desirable.

The result indicates that, while smaller γ and greater w can result in slightly greater sample complexity, the overall sample complexity does not vary much and the superiority of our algorithm over the naive uniform sampling algorithm $\text{N}\varepsilon\text{BAI}$ is maintained. We conclude that our algorithm is robust against these choices, as long as Assumption 1 is not severely violated.




Figure 4: Robustness towards w and b .

O.4 Benchmarks: $D\varepsilon$ BAI and $D\varepsilon$ BAI $_\beta$

We present the Distribution ε -Best Arm Identification ($D\varepsilon$ BAI) in Algorithm 8 and its variant $D\varepsilon$ BAI $_\beta$ in Algorithm 9. According to Dynamics 3, the agent has access to the context vector $\theta_{j_t}^*$ and the changepoint $t \in \mathcal{C}$. Thus only the distribution P_θ remains unknown and the agent needs to estimate it via the observed contexts.

Algorithm 8 DISTRIBUTION ε -BEST ARM IDENTIFICATION ($D\varepsilon$ BAI)

- 1: **Input:** the arm set \mathcal{X} , the latent vector matrix Θ , the slackness parameter ε , the confidence parameter δ .
- 2: **Initialize:** Compute the G-optimal allocation λ^* .
- 3: Compute $C_3 = \sum_{n=1}^{\infty} n^{-3}$.
- 4: Observe $\theta_{j_1}^*$.
- 5: Compute

$$\ddot{p}_{t,j} = \sum_{l_s=1}^{l_t} \frac{\mathbb{1}\{c_{l_s} = c_j\}}{l_t}, \quad \ddot{x}_t = \arg \max_{x \in \mathcal{X}} x^\top \Theta \ddot{p}_t \quad (\text{O.1})$$

$$\ddot{\beta}_t = \sqrt{\frac{1}{2l_t} \ln \frac{2C_3 N l_t^3}{\delta}}, \quad \ddot{\rho}_t(\ddot{x}_t, x) = \ddot{\beta}_t \min_{\zeta(\ddot{x}_t, x) \in \mathbb{R}} \sum_{j=1}^N |\Delta_j(\ddot{x}_t, x) + \zeta(\ddot{x}_t, x)|$$

- 6: **while** $\min_{x \neq \ddot{x}_t} (\ddot{x}_t - x)^\top \Theta \ddot{p}_t - \ddot{\rho}_t < -\varepsilon$ or $t \notin \mathcal{C}$ **do**
 - 7: Observe $\theta_{j_t}^*$ and $\mathbb{1}\{t \in \mathcal{C}\}$ update $t = t + 1$.
 - 8: Update \ddot{x}_t and $\ddot{\rho}_t$ with (O.1) and update l_t if $t \in \mathcal{C}$.
 - 9: **end while**
 - 10: Recommend arm $\ddot{x}_\varepsilon = \ddot{x}_t$.
-

The design is straightforward except for $\zeta(\ddot{x}_t, x)$. It minimizes the summation $\zeta(\ddot{x}_t, x) := \arg \min_{y \in \mathbb{R}} \sum_{j=1}^N |\Delta_j(\ddot{x}_t, x) + \zeta(\ddot{x}_t, x)|$. This trick has also been utilized in the design of $\text{PS}\varepsilon$ BAI (see (I.5)). It helps to better exploit the structure of the latent vectors Θ and facilitates the identification process. For easy implementation, we choose a proxy $\zeta(\ddot{x}_t, x) = \arg \min_{y \in \mathbb{R}} \sum_{j=1}^N (\Delta_j(\ddot{x}_t, x) + \zeta(\ddot{x}_t, x))^2 = -\frac{1}{N} \sum_{j=1}^N \Delta_j(\ddot{x}_t, x)$

Algorithm 9 DISTRIBUTION ε -BEST ARM IDENTIFICATION- β ($D\varepsilon\text{BAI}_\beta$)

- 1: **Input:** the arm set \mathcal{X} , the latent vector matrix Θ , the slackness parameter ε , the confidence parameter δ .
- 2: **Initialize:** Compute the G-optimal allocation λ^* .
- 3: Compute $C_3 = \sum_{n=1}^{\infty} n^{-3}$.
- 4: Observe $\theta_{j_1}^*$.
- 5: Compute

$$\hat{p}_{t,j}^\circ = \sum_{l_s=1}^{l_t} \frac{L_{l_s} \mathbb{1}\{c_{l_s} = c_j\}}{t}, \quad \hat{x}_t = \arg \max_{x \in \mathcal{X}} x^\top \Theta \hat{p}_t^\circ, \quad (\text{O.2})$$

$$\hat{\beta}_{t,j}^\circ := \min \left\{ \frac{\frac{1}{3} L_{\max} \ln \frac{2}{\delta_t} + \sqrt{\left(\frac{1}{3} L_{\max} \ln \frac{2}{\delta_t}\right)^2 + \frac{2\phi_{t,j} L_{\max}}{t} \ln \frac{2}{\delta_t}}}{t}, 1 \right\},$$

$$\hat{\phi}_{t,j}^\circ := \min \left\{ 4 \max \left\{ \hat{p}_{t,j}^\circ, \frac{25}{4} \frac{L_{\max}}{t} \ln \frac{2}{\delta_t} \right\}, \frac{1}{4} \right\}, \quad \delta_t = \frac{\delta}{C_3 N t^3},$$

$$\hat{\rho}_t^\circ(\hat{x}_t, x) = \min_{\zeta(\hat{x}_t, x) \in \mathbb{R}} \sum_{j=1}^N |\Delta_j(\hat{x}_t, x) + \zeta(\hat{x}_t, x)| \hat{\beta}_{t,j}^\circ.$$

- 6: **while** $\min_{x \neq \hat{x}_t} (\hat{x}_t - x)^\top \Theta \hat{p}_t^\circ - \hat{\rho}_t^\circ < -\varepsilon$ **do**
 - 7: Observe $\theta_{j_t}^*$ and $\mathbb{1}\{t \in \mathcal{C}\}$ and let $t = t + 1$.
 - 8: Update \hat{x}_t and $\hat{\rho}_t^\circ$ with (O.2).
 - 9: **end while**
 - 10: Recommend arm $\hat{x}_\varepsilon = \hat{x}_t$.
-

$D\varepsilon\text{BAI}_\beta$ utilizes the techniques we have used to bound the deviation between the true distribution p_j and the estimated ones $\hat{p}_{t,j}^\circ$ for all $j \in [N]$. Similarly, we let $\zeta(\hat{x}, x) = \arg \min_{\zeta(\hat{x}, x) \in \mathbb{R}} \sum_{j=1}^N \hat{\beta}_{t,j}^\circ (\Delta_j(\hat{x}_t, x) + \zeta(\hat{x}_t, x))^2 = -\frac{\sum_{j=1}^N \hat{\beta}_{t,j}^\circ \Delta_j(\hat{x}_t, x)}{\sum_{j=1}^N \hat{\beta}_{t,j}^\circ}$ for efficient computing.

As the theoretical guarantee of $D\varepsilon\text{BAI}_\beta$ can be derived following a similar manner as the proof of the DE term of $\text{PS}\varepsilon\text{BAI}$ in Appendix I.2 and in the proof of Lemma G.6, for simplicity, we just present the result and omit the proof here here.

Theorem O.1. $D\varepsilon\text{BAI}_\beta$ identifies an ε -best arm within

$$\tilde{O} \left(\max_{x_\varepsilon \in \mathcal{X}, x \neq x_\varepsilon, x^*} \frac{L_{\max} \hat{H}^2(x_\varepsilon, x)}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} \right),$$

time steps with probability at least $1 - \delta$ and in expectation, where

$$\hat{H}(x_\varepsilon, x) := \min_{\zeta(\hat{x}_\varepsilon, x) \in \mathbb{R}} \sum_{j=1}^N \sqrt{\min\{16p_j, \frac{1}{4}\}} |\Delta_j(\hat{x}_\varepsilon, x) + \zeta(\hat{x}_\varepsilon, x)|.$$

We present a theorem, along with a proof sketch, for the theoretical guarantee of $D\varepsilon\text{BAI}$ below.

Theorem O.2. $D\varepsilon\text{BAI}$ identifies an ε -best arm within

$$\tilde{O} \left(\max_{x_\varepsilon \in \mathcal{X}, x \neq x_\varepsilon, x^*} \frac{L_{\max} \ddot{H}^2(x_\varepsilon, x)}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} \right),$$

time steps with probability at least $1 - \delta$ and in expectation, where $\ddot{H}(x_\varepsilon, x) := \min_{\zeta(\ddot{x}_\varepsilon, x) \in \mathbb{R}} \sum_{j=1}^N |\Delta_j(\ddot{x}_\varepsilon, x) + \zeta(\ddot{x}_\varepsilon, x)|$.

Proof. For the sake of conciseness and simplicity, only a proof sketch is provided for this Theorem. Similar results can be found in the referred contents. The proof is composed by 4 steps.

- Step 1: bound the deviation of $P_\theta[\theta_j^*]$ by Lemma O.3 and Remark O.4.
- Step 2: prove the recommended arm is an ε -best arm. Conditional on Step 1, we can prove this following the procedures in Lemma J.1.
- Step 3: give a sufficient condition and then obtain a high probability upper bound. This can be seen from (J.1) and by solving

$$\begin{aligned} \check{p}_t(\check{x}_t, x) &= \check{\beta}_t \sum_{j=1}^N |\Delta_j(\check{x}_t, x) + \zeta(\check{x}_t, x)| \leq \frac{\Delta_{\min} + \varepsilon}{2} \\ \Rightarrow l_t &= \tilde{O} \left(\frac{\check{H}^2(x_\varepsilon, x)}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} \right). \\ \Rightarrow t &= \tilde{O} \left(\frac{L_{\max} \check{H}^2(x_\varepsilon, x)}{(\Delta_{\min} + \varepsilon)^2} \ln \frac{1}{\delta} \right). \end{aligned}$$

The high probability result is obtained by maximizing the above over $x_\varepsilon \in \mathcal{X}, x \neq x_\varepsilon, x^*$.

- Step 4: the expected result can be derived in a similar method as in the proof of $N\varepsilon$ BAI in Appendix F. The expected sample complexity and the high-probability sample complexity is of the same order.

□

Lemma O.3. *With probability at least $1 - \delta$, $\sup_{j \in [N]} |p_j - \check{p}_{t,j}| \leq \check{\beta}_t$ for all $t \in \mathbb{N}$, where $\check{\beta}_t = \sqrt{\frac{2}{l_t} \ln \frac{2}{\delta_{l_t}}}$ and $\delta_{l_t} = \frac{\delta}{C_3 l_t^3}$.*

Proof. We may define the cumulative distribution function (CDF) and the empirical CDF for P_θ as $F_\theta(j) = \sum_{k=1}^j P_\theta[\theta_j^*]$ and $\check{F}_t(j) = \sum_{l_s=1}^{l_t} \mathbb{1}\{c_{l_s} = c_j, l_s \leq j\}$ respectively for $j \in [N]$. According to the DKW inequality, we have

$$\mathbb{P} \left[\sup_{j \in [N]} |\check{F}_t(j) - F(j)| \geq \epsilon \right] \leq 2 \exp(-2l_t \epsilon^2).$$

By the implication

$$\begin{aligned} \epsilon &\leq |p_j - \check{p}_{t,j}| = |F(j) - F(j-1) - \check{F}_t(j) + \check{F}_t(j-1)| \\ &\leq |F(j) - \check{F}_t(j)| + |F(j-1) - \check{F}_t(j-1)| \\ \Rightarrow |F(j) - \check{F}_t(j)| &\geq \frac{\epsilon}{2} \quad \text{or} \quad |F(j-1) - \check{F}_t(j-1)| \geq \frac{\epsilon}{2}, \end{aligned}$$

we have that

$$\begin{aligned} \mathbb{P} \left[\sup_{j \in [N]} |p_j - \check{p}_{t,j}| \geq \epsilon \right] &\leq \mathbb{P} \left[\sup_{j \in [N]} |\check{F}_t(j) - F(j)| \geq \frac{\epsilon}{2} \right] \leq 2 \exp \left(-2l_t \left(\frac{\epsilon}{2} \right)^2 \right) \\ &= 2 \exp \left(-\frac{l_t \epsilon^2}{2} \right). \end{aligned}$$

This is equivalent to

$$\mathbb{P} \left[\sup_{j \in [N]} |p_j - \check{p}_{t,j}| \geq \check{\beta}_t \right] \leq \delta_{l_t}.$$

A union bound gives an upper bound of the failure probability $\sum_{l_t=1}^{\infty} \delta_{l_t} = \delta$. □

Remark O.4. *The use of the DKW inequality gives a union bound over the deviation of the distribution estimation all contexts. This avoid the N factor in the logarithm in $\check{\beta}_t$ and is beneficial when the number of contexts N is large.*

When there are a few contexts, it is also possible to directly bound the deviation for each context via Azuma-Hoeffding's inequality, i.e.,

$$\mathbb{P} \left[|p_j - \check{p}_{t,j}| \leq \check{\beta}_t \right] \leq \delta_t$$

where $\check{\beta}_t := \sqrt{\frac{1}{2l_t} \ln \frac{2C_3 N l_t^3}{\delta}}$. A union bound gives that with probability at least $1 - \delta$, $|p_j - \check{p}_{t,j}| \leq \check{\beta}_t$ for all $t \in \mathbb{N}, j \in [N]$. As this new confidence radius is the same as the one in Lemma O.3 up to constant and logarithmic terms, the sample complexity should be of the same order.

We adopt this confidence radius $\check{\beta}_t := \sqrt{\frac{1}{2l_t} \ln \frac{2C_3 N l_t^3}{\delta}}$ in the experiment.

P Additional Discussions

In this section, we provide additional discussions on

- the related methods for BAI in the nonstationary bandits literature in subsection P.1,
- the upper on the sample complexity in Theorem 3.2 and Theorem 3.3 in subsection P.2,
- the connection between the piecewise-stationary linear bandits model to the stationary linear bandits model in subsection P.3,
- the special case where the number of context $N = 1$ in subsection P.4,

P.1 Discussion on the Related Methods in Nonstationary Bandits

To the best of our knowledge, there is limited literature investigating the BAI in the nonstationary bandits setup (there is comparatively a much richer literature for regret minimization in non-stationary environments): [14] studies BAI in the *fixed-horizon* setup and [15] assumes the *best arm remains unchanged* after certain time step and explores the fixed-confidence setting. Both of the works are not directly comparable to our proposed piecewise-stationary setup.

Additionally, we provide strong baselines algorithms $D\varepsilon\text{BAI}$ and $D\varepsilon\text{BAI}_\beta$ in Section 6, which are detailed in Section O.4. These two baselines are given oracle information about the context (see Dynamics 3), while $\text{PS}\varepsilon\text{BAI}^+$ is not. As indicated by Figure 2(b), $\text{PS}\varepsilon\text{BAI}^+$ is competitive compared to these strong baselines and is much better than the naive approach.

P.2 More Discussions on the Upper Bound

Firstly, we emphasize that the sample complexity in Theorem 3.2 and Theorem 3.3 are instance-dependent, in particular, the term H_{DE} characterizes the difficulty in estimating the distribution of the context. Furthermore, the presence of the gaps $\Delta(x, x^*)$ also underscores that our upper bounds are functions of the instance.

Secondly, our algorithm adopts the G-optimal design, which is minimax optimal for the BAI in standard linear bandits problem [1, 2]. Although we have not proved it, we believe that the sample complexity of the proposed algorithm may not be instance-dependent optimal.

Lastly, we provide some remarks on the instance-dependent optimal algorithms:

- (1) The G-optimal design is the cornerstone for the more efficient and adaptive rules like \mathcal{XY} -allocation and \mathcal{XY} -Adaptive Allocation in [1]. Therefore, our algorithm provides a framework for more sophisticated algorithms in the piecewise-stationary linear bandits problem with the BAI task. Empirically, one can attempt to replace the G-optimal design by the \mathcal{XY} -allocation during implementation, which can possibly yield empirical benefits.
- (2) The current lower bound is established with two simpler problems (see Section 4), which is sufficient to show our algorithm is minimax optimal. However, a tighter instance-dependent lower bound based on the original problem may be required if one wishes to show an algorithm is instance-dependent optimal. This can be challenging because the distribution of the arms, the contexts and the changepoints are unknown, and the characterization of the "alternative instance" (i.e., the ε -best arm set is changed) is difficult.

To conclude, we believe an instance-dependent optimal algorithm is appealing and can lead to future research.

P.3 Connections to the Stationary Bandits

Firstly, we would like to clarify that, in the piecewise-stationary linear bandits model, the instance tends towards a stationary one when $\max_{j \in [N]} p_j \rightarrow 1$, instead of the scenario in which $L_{\max} \rightarrow \infty$.

When L_{\min} and L_{\max} are large, estimating each latent vector becomes easier since there are more observations in a single stationary segment. However, the overall reward of an arm also depends on P_θ , the distribution of latent context vectors, which is also assumed to be unknown in the setup. In order to estimate P_θ , a learning agent can get one (unobservable) sample from P_θ only at the (unobservable) change points. In other words, L_{\max} characterizes the ‘‘sparsity’’ of the context samples: the larger L_{\max} is, the sparser the context samples are. A larger L_{\max} indicates that samples from P_θ are likely to be generated less frequently, which will result in the increase of the overall sample complexity. Consider an extreme example of Dynamics 3 where the context vector is revealed at every time step and L_{\min} and L_{\max} are very large. If at least l_τ context samples from P_θ are required to identify the best arm, then a sample complexity of $l_\tau L_{\min}$ is unavoidable in our setup under this extreme case.

Secondly, while $\max_{j \in [N]} p_j$ is important in characterizing the (non)stationarity of the instance, we would like to emphasize that **both** the distribution P_θ and latent context vectors are essential in characterizing the sample complexity, as shown in the term $\bar{H}(x_\epsilon, x)$ in Theorem 3.2. To further illustrate this, we provide a simple but illuminating instance as follows: The instance is composed by $N = 3$ contexts and $K = 2$ arms, and Δ_j denotes the mean gap between arm 1 and arm 2 under context $j \in [N]$ (to be specified below) and $\Delta := \sum_{j \in [N]} p_j \Delta_j$ denotes the weighted mean gap between arm 1 and arm 2. The probability of context 1 is $p_1 = \max_{j \in [N]} p_j = 1 - p_2 - p_3$, and p_1 is close to 1. Let \hat{x} denote the empirical version of the statistics x in the following arguments.

In such a case, consider this question: *if an algorithm has only detected context 1 in the first t time steps and t is large, should it stop or not?*

As no change point has occurred, the current dynamics behaves similarly to that of a stationary bandit environment up till now. Thus the empirical probabilities of the 3 contexts are $\hat{p}_1 = 1, \hat{p}_2 = \hat{p}_3 = 0$ and the empirical mean gap under context 1 $\hat{\Delta}_1$ is close to Δ_1 . Consider the following two cases:

- Case 1: $\Delta_1 = p_2$ and $\Delta_2 = \Delta_3 = -p_1$. The true weighted mean gap between arm 1 and arm 2 is $\Delta = \sum_{j \in [N]} p_j \Delta_j = -p_1 p_3 < 0$, and thus arm 2 is the best arm. This indicates that, in order to estimate Δ_2 and Δ_3 , an algorithm needs to observe contexts 2 and 3 despite their small probabilities. But currently $\hat{p}_2 = \hat{p}_3 = 0, \hat{\Delta} = \hat{\Delta}_1 \approx \Delta_1 > 0$, so the algorithm should not terminate.
- Case 2: $\Delta_1 = 1$ and $\Delta_2, \Delta_3 \in [-2, 2]$. Thus $\hat{\Delta} \approx \Delta \geq p_1 - 2p_2 - 2p_3 > 0$. This indicates there is no need to observe contexts 2 and 3, because arm 1 is the best even in the worst case ($\Delta_2 = \Delta_3 = -2$). In this case, as long as the algorithm gets a good estimate of Δ_1 , it can confidently terminate.

As these two cases indicate, in addition to P_θ , the latent vectors (which determine the means of the arms) are equally important for the stopping time. Our algorithm takes both factors into account, as reflected by the summation term in equation (3.4) and $\hat{\Delta}_t(x_t^*, x)$ in (3.5) for the algorithm design, and by $\bar{H}(x_\epsilon, x)$ for the theoretical upper bound.

We **conclude** that

- The unknown distribution P_θ and the latent vectors **jointly** determine the sample complexity. $\bar{H}(x_\epsilon, x)$ in Theorem 3.2 characterizes their roles, where the contexts with larger probabilities have commensurately greater influence on the sample complexity of the algorithm.
- When $\max_{j \in [N]} p_j \rightarrow 1$ with other parameters fixed, the piecewise-stationary linear bandit instance reduces to a stationary one. $H_{DE}(x_\epsilon, x)$ in Theorem 3.2 becomes $L_{\max}/4$, which implies that a **constant** number of change points need to be observed, and thus $T_D(x_\epsilon, x)$ can be regarded as a constant. The dominant term in our upper bound, the $T_V(x)$ term, recovers the upper bound in the stationary bandits, that is, $\max_{x \neq x^*} \frac{d}{(\Delta(x^*, x) + \epsilon)^2} \ln(1/\delta)$ [32].

P.4 Special Case: $N = 1$

We only consider $N > 1$ as the bandit model is only **truly** piecewise-stationary when $N > 1$. Thus, we did not specifically derive upper and lower bounds for the extreme case where $N = 1$. Nevertheless, our analysis can cover this case due to the following reasons:

- The (minimax) lower bound in Theorem 4.1 is not directly applicable since it is not derived for this extreme instance. However, if we step back to the analysis equation (M.16), the feasible set in the optimization problem is empty, and thus the lower bound on $\mathbb{E}[l_t]$ (N_C in Theorem 4.1) is 0. In this case, the lower bound in Theorem 4.1 reduces to the lower bound in the stationary bandits.
- Regarding the upper bound in Theorem 3.2, when $N = 1$, we have $p_1 = 1$ and $H_{DE} = \frac{L_{\max}}{4}$. In this case, as H_{DE} does not depend on the mean gaps and the mean gaps are among the fundamental quantities to characterize the difficulty of an instance, the T_D term can be regarded as a constant. Therefore, the dominant term in the upper bound is T_V . This recovers the bound in the stationary setup.
- In addition, as we assume N is an input to our algorithm, when we are aware $N = 1$, we can actually adopt the algorithms for the stationary setting, e.g., the G -allocation or $\mathcal{X}\mathcal{Y}$ -allocation rule [1].

Q ε -Best Arm Tuple Identification Problem

In the main paper, we consider the identification of an ε -best arm in terms of the “ensemble” quality $\mu_x := \mathbb{E}_{\theta \sim P_\theta} [x^\top \theta] = \sum_{j=1}^N P_\theta[\theta_j^*] x^\top \theta_j^*$. The curious reader may wonder whether we can identify the ε -best arm tuple $\mathcal{X}_\varepsilon^{\text{tuple}} := (x_1^\varepsilon, \dots, x_N^\varepsilon)$, where x_j^ε is an ε -best arm under context j , i.e., $x_j^\varepsilon \in \{x : x^\top \theta_j^* \geq \max_{x \in \mathcal{X}} x^\top \theta_j^* - \varepsilon\}$. Given the tools we developed in this manuscript, we answer this problem in the affirmative.

Intuitions: Let $x_j^* := \arg \max_{x \in \mathcal{X}} x^\top \theta_j^*$ and $\Delta_j^* := \min_{x \neq x_j^*} \Delta_j(x_j^*, x)$. We expect to have an upper bound taking the form of

$$\tilde{O} \left(\max_{j \in [N]} \frac{L_{\max}}{p_j} \cdot \frac{d}{(\Delta_j^* + \varepsilon)^2 L_{\min}} \ln \frac{1}{\delta} \right)$$

where $O(\frac{d}{(\Delta_j^* + \varepsilon)^2 L_{\min}} \ln \frac{1}{\delta})$ is an upper bound on the number of context samples j and context j will occur once among every $\frac{1}{p_j}$ contexts in expectation.

Analysis of the problem: The change detection and context alignment subroutines are only effective within τ^* time steps (Line 2 in Algorithm 1). However, if a context is with small occurrence probability p_j , it may not appear before τ^* time steps.

Regarding this scenario, we only expect to obtain a high-probability upper bound on the sample complexity, whereas the expected sample complexity requires more techniques beyond our parallel execution trick (which is used to design $\text{PS}\varepsilon\text{BAI}^+$) and is an interesting direction left for future research.

Besides, it is more feasible to consider the identification of ε -best arms under contexts with high occurrence probability, i.e., we aim to identify

$$\mathcal{X}_{\varepsilon, \bar{p}}^{\text{tuple}} := \{x_j^\varepsilon : p_j > \bar{p}\}$$

where \bar{p} is a threshold on the occurrence probability of contexts. Let $\Theta_{\bar{p}} = \{\theta_j^* : p_j > \bar{p}\}$ denote those context with high occurrence probability.

Goal: We aim to devise an algorithm with the minimum sample complexity (arm pulls) to ascertain either (1) an ε -best arm under context j , or (2) $\theta_j^* \notin \Theta_{\bar{p}}$.

With the above goal, we expect to identify an ε -best arm for all j with $\theta_j^* \in \Theta_{\bar{p}}$; for those j with $\theta_j^* \notin \Theta_{\bar{p}}$, either an ε -best arm is identified or $\theta_j^* \notin \Theta_{\bar{p}}$ is ascertained.

We propose Algorithm 10 for this problem, which is quite similar to Algorithm 1, except for a few changes:

Algorithm 10 PIECEWISE-STATIONARY ε -BEST ARM TUPLE IDENTIFICATION

- 1: **Input:** arm set \mathcal{X} , size of the set of latent vectors N , bounds on the segment lengths L_{\min} and L_{\max} , slackness parameter ε , confidence parameter δ , sampling parameter γ and window size w , threshold b , probability threshold \bar{p} .
- 2: **Initialize:** Compute the G-optimal allocation λ^* and $\tau^* = \frac{38400 \ln(80) N L_{\max}}{\varepsilon^2} \ln \frac{N^2 K L_{\max}}{\delta \varepsilon^2}$ and Flag = $[N]$ and Hold = $[\]$ and Output = $[\]$.
- 3: Set $\text{CD}_{\text{sample}} = [\]$, $\text{CA}_{\text{id}} = \{ \}$. Set $t_{\text{CD}} = +\infty$.
- 4: Set $\mathcal{T}_{t,j} = \emptyset$ and initialize $\mathcal{T}_t, T_{t,j}, T_t$ with (3.1) for all $t \leq \tau^*, j \in [N]$.
- 5: Sample $\frac{w}{2}$ arms $\{x_s\}_{s=1}^{\frac{w}{2}} \sim \lambda^*$ and observe the associated returns $\{Y_{s,x_s}\}_{s=1}^{\frac{w}{2}}, t = \frac{w}{2}, t_{\text{CA}} = \frac{w}{2}$.
- 6: $\text{CA}_{\text{id}} = \{1 : [(x_s, Y_{s,x_s})]_{s=1}^{\frac{w}{2}}\}, \hat{j}_t = 1$.
- 7: **while** $t \leq \tau^*$ and Flag $\neq \emptyset$ **do**
- 8: $t = t + 1$
- 9: Sample an arm $x_t \sim \lambda^*$ and observe return Y_{t,x_t} .
- 10: **if** $\text{mod}(t - t_{\text{CA}}, \gamma) \neq 0$ **then**
- 11: Update $\hat{j}_t = \hat{j}_{t-1}, \mathcal{T}_{t,\hat{j}_t} = \mathcal{T}_{t-1,\hat{j}_t} \cup \{t\}, \mathcal{T}_{t,j} = \mathcal{T}_{t-1,j}$ for $j \neq \hat{j}_t$.
- 12: **else**
- 13: $\text{CD}_{\text{sample}} = \text{CD}_{\text{sample}} + [(x_t, Y_{t,x_t})]$.
- 14: Update $\hat{j}_t = \hat{j}_{t-1}, \mathcal{T}_{t,j} = \mathcal{T}_{t-1,j}$ for all $j \in [N]$.
- 15: **if** $|\text{CD}_{\text{sample}}| \geq w$ **then**
- 16: **if** $\text{LCD}(\mathcal{X}, w, b, \text{CD}_{\text{sample}}[-w :]) \text{ then}$
- 17: $\text{CD}_{\text{sample}} = [\]$.
- 18: $t = t + \frac{w}{2}, t_{\text{CA}} = t, t_{\text{CD}} = +\infty$.
- 19: $\hat{j}_t, \text{CA}_{\text{id}} = \text{LCA}(\mathcal{X}, w, b, \text{CA}_{\text{id}})$.
- 20: **if** $\hat{j}_t = N + 1$ **then break**.
- 21: Revert $\mathcal{T}_{t,j} = \mathcal{T}_{t - \frac{w(\gamma+1)}{2}, j}$ for all $j \in [N]$.
- 22: **end if**
- 23: **end if**
- 24: **end if**
- 25: Update the estimates with (3.1), (3.2) and (Q.1).
- 26: **if** $\exists j \in \text{Flag}$ condition (Q.2) is met for j and $t_{\text{CD}} = +\infty$ **then**
- 27: **for** $j \in \text{Flag}$ **do**
- 28: **if** $\min_{x: x \neq x_{t,j}^*} \hat{\Delta}_{t,j}(x_{t,j}^*, x) - \alpha_{t,j} \geq -\varepsilon$ **then**
- 29: Record j and $\hat{x}_{j,\varepsilon} := \arg \max_{x \in \mathcal{X}} x^\top \hat{\theta}_{t,j}$ in Hold.
- 30: Flag = Flag $\setminus \{j\}$ and $t_{\text{CD}} = |\text{CD}_{\text{sample}}|$.
- 31: **else if** $\hat{p}_j + \beta_{t,j} < \bar{p}$ **then**
- 32: Record j and $\hat{x}_{j,\varepsilon} = \text{NAN}$ in Hold.
- 33: Flag = Flag $\setminus \{j\}$ and $t_{\text{CD}} = |\text{CD}_{\text{sample}}|$.
- 34: **end if**
- 35: **end for**
- 36: **else if** $t_{\text{CD}} = |\text{CD}_{\text{sample}}| - \frac{w}{2}$ **then**
- 37: **for** $(j, \hat{x}_{j,\varepsilon})$ in Hold **do**
- 38: Output $[j] = \hat{x}_{j,\varepsilon}$
- 39: **end for**
- 40: **end if**
- 41: **end while**
- 42: Recommend Output

- On Line 7, the algorithm stops when an ε -best arm or $p_j \leq \bar{p}$ is identified for all contexts.
- On Line 25, it computes the confidence radii for the arms in context j_t as well as the confidence radii for the occurrence probabilities

$$\alpha_{t,j} = \frac{d}{n} \ln \frac{2}{\delta_{v,T_t}} + \sqrt{\left(\frac{d}{n} \ln \frac{2}{\delta_{v,T_t}}\right)^2 + \frac{4d}{n} \ln \frac{2}{\delta_{v,T_t}}}, \quad (\text{Q.1})$$

$$\beta_{t,j} := \min \left\{ \frac{5}{2} \sqrt{\frac{2\phi_{t,j} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}}, 1 \right\},$$

$$\text{where } \phi_{t,j} := \min \left\{ 4 \max \left\{ \hat{p}_{t,j}, \frac{25}{4} \frac{L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}} \right\}, \frac{1}{4} \right\},$$

• On Line 26 to 40, the stopping rule is changed:

- Stopping condition **(I)** on Line 26

$$\left(\min_{x: x \neq x_{t,j}^*} \hat{\Delta}_{t,j}(x_{t,j}^*, x) - \alpha_{t,j} \geq -\varepsilon \quad \text{or} \quad \hat{p}_j + \beta_{t,j} < \bar{p} \right) \quad (\text{Q.2})$$

$$\text{and } T_t \geq (2L_{\max}/9) \ln(2/\delta_{d,T_t})$$

where $x_{t,j}^* := \arg \max_{x \in \mathcal{X}} x^\top \hat{\theta}_{t,j}$.

- Lines 27 to 35: identify an ε -best arm or ascertain $p_j \leq \bar{p}$ among the remaining contexts, and record these observations in Hold for easy access in stopping condition **(II)**.
- Lines 36 to 40: the recommended ε -best arm is recorded, where we adopt NAN to flag those contexts with small occurrence probabilities.

Theorem Q.1. *Given an instance Λ , with probability at least $1 - \delta$, Algorithm 10 can recommend an ε -best arm for all context $j \in \Theta_{\bar{p}}$ with sample complexity*

$$\max \left\{ \tilde{O} \left(\max_{\theta_j^* \in \Theta_{\bar{p}}} \max \left\{ L_{\max}, \frac{d}{(\Delta_j^* + \varepsilon)^2} \right\} \cdot \frac{\ln(1/\delta)}{p_j} \right), \tilde{O} \left(\max_{\theta_j^* \notin \Theta_{\bar{p}/2}} \frac{\min \{p_j, \frac{1}{64}\} L_{\max} \ln(1/\delta)}{(p_j - \bar{p})^2} \right) \right\},$$

$$\tilde{O} \left(\max_{\theta_j^* \in \Theta_{\bar{p}/2} \setminus \Theta_{\bar{p}}} \min \left\{ \max \left\{ L_{\max}, \frac{d}{(\Delta_j^* + \varepsilon)^2} \right\} \frac{\ln(1/\delta)}{p_j}, \frac{\min \{p_j, \frac{1}{64}\} L_{\max} \ln(1/\delta)}{(p_j - \bar{p})^2} \right\} \right),$$

which can be simplified as

$$\tilde{O} \left(\max \left\{ L_{\max}, \frac{d}{\varepsilon^2} \right\} \cdot \frac{\ln(1/\delta)}{\bar{p}} \right).$$

Proof of Theorem Q.1. We provide a concise proof sketch for this theorem.

By adapting the stopping rule for best arm identification in stationary linear bandits to ε -best arm identification, we observe that the number of arm pulls needed for ε -best arm identification under context j is

$$T_{t,j} = \tilde{O} \left(\frac{d}{(\Delta_j^* + \varepsilon)^2} \ln \frac{1}{\delta} \right).$$

The remaining problem is to determine how many context samples/changepoints are needed such that the above number of arm samples can be achieved.

Recall that $\hat{p}_{t,j} = \frac{T_{t,j}}{T_t}$ (3.2) and Lemma I.2, we have

$$\mathbb{P} [|T_{t,j} - T_t p_j| \geq T_t \beta_{t,j} | \text{Good}] \leq \delta_{d,T_t}.$$

In addition, we have an upper bound on $\beta_{t,j}$ as in (J.5), so there would be sufficient arm samples for ε -best arm identification under context j if

$$T_t \cdot \left(p_j - 2 \cdot \frac{5}{2} \sqrt{\frac{2 \min \{16p_j, \frac{1}{4}\} L_{\max}}{T_t} \ln \frac{2}{\delta_{d,T_t}}} \right) \geq \tilde{O} \left(\frac{d}{(\Delta_j^* + \varepsilon)^2} \ln \frac{1}{\delta} \right)$$

By solving this inequality in terms of T_t , we obtain

$$T_t = \tilde{O} \left(\max \left\{ L_{\max}, \frac{d}{(\Delta_j^* + \varepsilon)^2} \right\} \cdot \frac{\ln(1/\delta)}{p_j} \right).$$

As we aim to identify the ε -best arms under each context $\theta_j^* \in \Theta_{\bar{p}}$, the upper bound is at least

$$T_t = \tilde{O} \left(\max_{\theta_j^* \in \Theta_{\bar{p}}} \max \left\{ L_{\max}, \frac{d}{(\Delta_j^* + \varepsilon)^2} \right\} \cdot \frac{\ln(1/\delta)}{p_j} \right) \quad (\text{Q.3})$$

even if the set $\Theta_{\bar{p}}$ is given.

Similarly, for those contexts with small occurrence probability $\theta_j^* \notin \Theta_{\bar{p}}$, by Lemma I.2 and (J.5), we can identify them when

$$\begin{aligned} p_j + 2 \cdot \frac{5}{2} \sqrt{\frac{2 \min \{16p_j, \frac{1}{4}\} L_{\max}}{T_t} \ln \frac{2}{\delta_{d, T_t}}} &\leq \bar{p} \\ \Rightarrow T_t = \tilde{O} \left(\frac{\min \{p_j, \frac{1}{64}\} L_{\max} \ln(1/\delta)}{(p_j - \bar{p})^2} \right). \end{aligned}$$

Careful readers may notice that the denominator depends on a ‘‘probability gap’’, which can be very small. In practice, if we can actually identify the ε -best arm in those contexts, it is also acceptable. Therefore, we instead only choose not to identify the ε -best arm in those contexts $\theta_j^* \in \Theta_{\bar{p}/2}$, which yields an upper bound

$$T_t = \tilde{O} \left(\max_{\theta_j^* \notin \Theta_{\bar{p}/2}} \frac{\min \{p_j, \frac{1}{64}\} L_{\max} \ln(1/\delta)}{(p_j - \bar{p})^2} \right). \quad (\text{Q.4})$$

In this case, the bound is at most $\tilde{O} \left(\frac{L_{\max} \ln(1/\delta)}{\bar{p}} \right)$. For the rest contexts $\theta_j^* \notin \Theta_{\bar{p}/2} \setminus \Theta_{\bar{p}}$, either identifying the ε -best arm or ascertaining its small occurrence probability suffices, that is,

$$T_t = \tilde{O} \left(\max_{\theta_j^* \in \Theta_{\bar{p}/2} \setminus \Theta_{\bar{p}}} \min \left\{ \max \left\{ L_{\max}, \frac{d}{(\Delta_j^* + \varepsilon)^2} \right\} \frac{\ln(1/\delta)}{p_j}, \frac{\min \{p_j, \frac{1}{64}\} L_{\max} \ln(1/\delta)}{(p_j - \bar{p})^2} \right\} \right). \quad (\text{Q.5})$$

The above bound is upper bounded by $\tilde{O} \left(\max \left\{ L_{\max}, \frac{d}{\varepsilon^2} \right\} \cdot \frac{\ln(1/\delta)}{\bar{p}} \right)$.

By taking the maximum of (Q.3), (Q.4) and (Q.5), we can obtain a high-probability problem-dependent upper bound on the sample complexity. In addition, we can get a high-probability problem-independent upper bound

$$\tilde{O} \left(\max \left\{ L_{\max}, \frac{d}{\varepsilon^2} \right\} \cdot \frac{\ln(1/\delta)}{\bar{p}} \right).$$

By setting the threshold \bar{p} carefully, the algorithm is guaranteed to terminate before τ^* given the good event Good (G.1). \square

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Empirical experiments are conducted to show that Assumption 1 is needed for theoretical proof and the proposed algorithm is robust to mild violations of the assumptions. The limitations are discussed in Appendix B.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The problem setting is clearly introduced in Section 2 and all theoretical results are accompanied with proofs or proof outlines.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The procedures of the algorithms are clearly presented. All experiment details are provided in Section 6 and Appendix O.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code is released. Please refer to Section 6.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Please refer to Section 6 and Appendix O.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Error bars (standard deviations) are included in all experiment results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.

- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Please see Appendix O.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: The authors conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The work is mainly theoretical, thus there is no identifiable societal impact.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is

being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.

- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The work is mainly theoretical.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The MATLAB code for the computation of G-optimal allocation used existing code, and has been mentioned and cited in Appendix O.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.

- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.