

Empirical Bayesian Multi-Bandit Learning

XIA JIANG, Institute of Science and Technology for Brain-inspired Intelligence, Fudan University, China
RONG J.B. ZHU*, PolicyLab, Children's Hospital of Philadelphia, US

Multi-task learning in contextual bandits has attracted significant research interest due to its potential to enhance decision-making across multiple related tasks by leveraging shared structures while accommodating task-specific heterogeneity. In this article, we propose a novel hierarchical Bayesian framework for learning across multiple bandit instances. This framework captures both heterogeneity and dependence among bandit instances through a hierarchical Bayesian model, enabling effective information sharing while accommodating instance-specific variations. Unlike existing methods that overlook the learning of the covariance structure across bandits, we introduce an empirical Bayesian approach to estimate the covariance matrix of the prior distribution. This enhances the practicality and flexibility of learning across multi-bandits. Building on this approach, we develop two efficient algorithms, EBM-TS (empirical Bayesian multi-bandit Thompson Sampling) and EBM-UCB (empirical Bayesian multi-bandit Upper Confidence Bound), which incorporate the estimated prior into the decision-making process. We establish the frequentist regret upper bounds for the proposed algorithms, providing theoretical guarantees for learning across multiple bandit instances. Extensive experiments on synthetic and real-world datasets demonstrate the superior performance of our proposed methods, particularly in complex environments. Our algorithms achieve lower cumulative regret compared to existing approaches across a range of settings, highlighting their ability to effectively balance exploration and exploitation while leveraging shared information across bandits.

JAIR Associate Editor: Marek Petrik

JAIR Reference Format:

Xia Jiang and Rong J.B. Zhu. 2026. Empirical Bayesian Multi-Bandit Learning. *Journal of Artificial Intelligence Research* 87, Article 17 (October 2026), 31 pages. doi: [10.1613/jair.1.20966](https://doi.org/10.1613/jair.1.20966)

1 Introduction

The contextual bandit problem provides a fundamental framework for analyzing decision-making in uncertain environments (Li et al., 2010; Agrawal and Goyal, 2013; Abeille and Lazaric, 2017; Bouneffouf et al., 2020; Bastani et al., 2021). It involves a decision-maker who, at each time step, observes a context vector describing the current state of the environment before making a choice among available arms. At each time step, the decision-maker must choose one of several available arms (or actions) based on the observed context. The chosen arm yields a reward, which is typically a function of both the context and the arm taken. The objective is to maximize cumulative reward over time by balancing exploration — trying new arms to gather information — and exploitation — choosing arms that have historically yielded high rewards (Auer et al., 2002; Bubeck et al., 2012; Agrawal and Goyal, 2017). The adaptability of contextual bandits to changing environments and their ability to optimize decisions based on contextual information make them a powerful tool for decision-making. Contextual bandits have been widely studied and applied in various domains, including personalized recommendation systems (Li

*Corresponding Author.

Authors' Contact Information: Xia Jiang, ORCID: [0009-0009-6805-2820](https://orcid.org/0009-0009-6805-2820), Institute of Science and Technology for Brain-inspired Intelligence, Fudan University, Shanghai, China; Rong J.B. Zhu, ORCID: [0000-0002-1893-2901](https://orcid.org/0000-0002-1893-2901), zhur1@chop.edu, PolicyLab, Children's Hospital of Philadelphia, Philadelphia, US.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).
doi: [10.1613/jair.1.20966](https://doi.org/10.1613/jair.1.20966)

et al., 2010; Yang et al., 2020; Guo et al., 2020), dynamic pricing (Misra et al., 2019; Mueller et al., 2019), and healthcare (Woodroffe, 1979; Mate et al., 2020).

Despite the extensive research on contextual bandits, most existing work focuses on single-bandit scenarios (Li et al., 2010; Agrawal and Goyal, 2013), where the decision-making process is confined to a single bandit instance. More recently, Lee et al. (2023) study context-lumpable contextual bandits, where contexts can be grouped according to a shared latent reward structure. However, learning across multiple bandit instances is common in practice. For example, in movie recommendation systems (Qin et al., 2014), there are numerous movies waiting to be recommended to different users. These users can be viewed as separate bandit instances, with each movie representing an arm. While certain movies may appeal similarly across users, the actual outcomes can vary significantly due to differences in users' preferences and characteristics. In such settings, traditional models that assume shared parameters across arms fall short, as they fail to capture user-specific variations and the nuanced influence of each movie on individual outcomes. Treating users as homogeneous overlooks the inherent heterogeneity among individuals. Conversely, modeling each user independently fails to leverage the similarities across users and leads to poor data usage, resulting in inefficient learning and suboptimal performance. This dichotomy highlights the necessity for a multi-task learning approach (Soare et al., 2014; Deshmukh et al., 2017; Wan et al., 2021; Fang and Tao, 2015; Su et al., 2024; Hong et al., 2023) that can effectively capture both the similarities and differences among multiple bandit instances. By doing so, such an approach can enhance learning efficiency and improve decision-making in complex, heterogeneous environments.

In the realm of multi-task learning, various strategies have been proposed to address the challenges of learning across multiple tasks. One common approach builds on meta-learning (Wan et al., 2021; Kveton et al., 2021), where each task is associated with a task-specific parameter vector, and these parameters are typically drawn from a common distribution, allowing for the transfer of knowledge across tasks through a hierarchical structure (Hong et al., 2023, 2022). However, applying this approach in the context of multi-armed learning is challenging, as the task-specific parameter is shared across all arms within a task, implying inherent correlations among the arms of the same task. To address this issue, prior work (Xu and Bastani, 2025; Huang et al., 2025; Cella and Pontil, 2021) imposed a sparse heterogeneity assumption on the arm parameters. While this assumption aims to capture both within-task heterogeneity and cross-task correlations among arms, it can result in poor performance when the assumption does not hold — highlighting the limited generalizability of such methods.

In this paper, we introduce a hierarchical Bayesian model designed to capture both the diversity and the shared structure across multiple bandit instances. Each arm is modeled with its own parameter vector, reflecting the natural heterogeneity among arms, while these parameters are connected through a shared prior distribution that enables knowledge transfer across tasks. This shared prior allows the model to learn from related bandits, improving decision-making even when data from individual instances is limited.

The hierarchical framework assumes a Gaussian prior and updates parameter beliefs through Bayes' theorem. To make the model adaptive and data-driven, we estimate the covariance matrix of the prior and the noise variance using an empirical Bayesian approach, allowing the shared structure to emerge naturally from the data rather than being fixed in advance.

Building on this foundation, we design two exploration strategies — Thompson Sampling (TS) and Upper Confidence Bound (UCB) — that utilize the estimated posterior distributions. These lead to two algorithms: empirical Bayesian multi-bandit Thompson Sampling (EBMTS) and empirical Bayesian multi-bandit Upper Confidence Bound (EBMUCB). We further derive the frequentist regret bounds for these algorithms, which explicitly illustrate how the incorporation of the prior influences learning performance. Through extensive experiments on both synthetic and real-world datasets, we demonstrate that our methods consistently achieve superior performance and greater robustness relative to existing approaches.

In summary, our key contributions are as follows:

- We introduce a hierarchical Bayesian model that captures shared structure across multiple bandit instances. This framework accommodates both heterogeneity and correlations among bandit instances, enabling efficient information sharing while preserving instance-specific flexibility.
- We develop an empirical Bayesian approach to learn the shared structure across multiple bandit instances, with a key contribution being the estimation of the covariance matrix — an essential yet often overlooked component. By employing the thresholded covariance matrix estimator (Bickel and Levina, 2008), our approach offers an automatic, data-driven way to uncover meaningful correlations across bandit instances. This method preserves strong correlations while eliminating weaker ones, effectively discarding those that do not contribute meaningfully to across-instance learning. We also propose computational techniques that enhance scalability and make our algorithms practical for large-scale applications.
- Building on our empirical Bayesian framework, we develop two algorithms — EBMTS and EBMUCB — that integrate Thompson Sampling and UCB-based exploration strategies for multi-bandit problems. These algorithms are evaluated across diverse datasets, demonstrating strong effectiveness and robustness in various application settings.
- We further establish a rigorous theoretical foundation by deriving an upper bound on the frequentist regret. Unlike most existing studies on multi-bandit problems (Hong et al., 2022; Aouali et al., 2023), which primarily analyze Bayesian regret, we explicitly characterize the estimation error and establish its upper bound with the incorporation of prior information — a topic that has rarely been addressed in the multi-bandit literature. Building on this analysis, Theorem 4.1 provides the resulting frequentist regret bound, thereby filling an important gap in the multi-bandit literature.

The remainder of this article is organized as follows. Section 2 provides a comprehensive review of the literature on contextual bandits and multi-task learning. Section 3 introduces the multi-bandit model and formulates the problem. Section 4 details the estimation procedure. Section 5 presents the proposed EBMTS and EBMUCB algorithms, developed based on the estimation results. Section 6 provides the frequentist regret bounds for EBMTS and EBMUCB. Section 7 reports experimental results on both synthetic and real-world datasets. Finally, Section 8 concludes the article.

2 Related Works

Our work is closely related to Xu and Bastani (2025), who propose a robust method for high-dimensional bandit problems under the assumption that the parameters of different bandits are sparse deviations from a shared global parameter. Huang et al. (2025) further improve the approach while maintaining the sparse heterogeneity assumption. In contrast, our framework removes the sparsity requirement and instead assumes that arm parameters are drawn from a shared prior distribution, enabling more flexible and adaptive learning of common structures across bandits.

Bayesian methods are particularly well-suited for multi-task learning, as they naturally incorporate prior knowledge and provide principled uncertainty quantification. Hong et al. (2022) propose a hierarchical Thompson Sampling algorithm (HierTS) that achieves substantially lower regret than methods that ignore shared structure. However, their approach does not address the estimation of the prior covariance matrix — a key component that remains an open challenge. Wan et al. (2021) introduce a multi-task Thompson Sampling (MTTS) framework based on metadata, along with several computational techniques and efficient variants to improve scalability. Unlike our method, MTTS leverages metadata to capture shared information rather than imposing structural assumptions on the model parameters. Both HierTS and MTTS fall under the paradigm of task-relation learning, where inter-task relationships are explicitly modeled. In contrast, our approach captures shared structure through a common prior, without explicitly modeling task dependencies.

Multi-task learning has been explored in various bandit-related domains. For example, [Gabillon et al. \(2011\)](#) and [Scarlett et al. \(2019\)](#) study multi-task approaches for best-arm identification, while [Swersky et al. \(2013\)](#) extend Bayesian optimization — a framework for automatic hyperparameter tuning — by incorporating multi-task Gaussian processes, substantially accelerating the optimization process compared to single-task methods. However, these studies address problems different from ours, as we focus on cumulative regret minimization.

3 Problem Formulation

We consider a contextual bandit problem with N bandit instances, each associated with K arms. The decision-making process unfolds over n time steps. At each time step t , a new individual with a d -dimensional context vector $\mathbf{x}_t$ arrives at one of the N bandit instances, determined by a random variable Z_t that follows a categorical distribution with probabilities p_j for $j \in [N]$. The context is drawn from an unknown distribution $\mathbb{P}(\mathbf{x})$. Based on this context vector and historical information, the agent chooses an arm from the K available arms of the selected bandit. The reward $y_{k,j,t}$ corresponding to the chosen arm is then obtained. We assume that the reward is linearly structured, that is, the reward for pulling arm k at time t in instance j is given by

$$y_{k,j,t} = \mathbf{x}_t^\top \boldsymbol{\beta}_{k,j} + \epsilon_{k,j,t}, \quad (1)$$

where $\boldsymbol{\beta}_{k,j} \in \mathbb{R}^d$ is the unknown parameter vector for arm k in instance j , and $\epsilon_{k,j,t}$ is the noise term, which follows a normal distribution with zero mean and standard deviation σ_k , i.e., $\epsilon_{k,j,t} \sim \mathcal{N}(0, \sigma_k^2)$. Herein, we account for the heteroscedasticity of arm-specific noise. In the experiments, the noise variance of each individual arm is estimated exclusively using the data collected by that arm itself, thereby ensuring the independence across different arms.

Following the setting in [Xu and Bastani \(2025\)](#), we consider that individuals arrive at different bandit instances with varying probabilities. These arrival probabilities determine the number of samples that each bandit instance receives. We consider the following two different settings:

- Data-balanced setting: The arrival probabilities p_j are approximately equal for all bandit instances j . Therefore, each bandit receives a similar number of arrivals over time.
- Data-poor setting: One or more bandit instances have significantly lower arrival probabilities compared to others. This setting is particularly useful for evaluating the algorithm's ability to handle cold start problems and leverage multi-task learning to improve efficiency.

To enable joint learning across the N bandit instances, we impose a structured prior on the parameters $\boldsymbol{\beta}_{k,j}$ associated with each arm k in bandit instance j . Specifically, we assume that the parameters $\boldsymbol{\beta}_{k,j}$ follow a normal distribution centered around a shared parameter vector $\boldsymbol{\beta}_{k0}$:

$$\boldsymbol{\beta}_{k,j} \sim \mathcal{N}(\boldsymbol{\beta}_{k0}, \Sigma_k),$$

where $\boldsymbol{\beta}_{k0}$ represents the shared parameter vector across all N bandit instances and Σ_k is the covariance matrix capturing the variability of the parameters for arm k . The deviation $\boldsymbol{\beta}_{k,j} - \boldsymbol{\beta}_{k0}$ quantifies the heterogeneity among bandit instances, allowing each instance to have its own unique characteristics while still benefiting from the shared structure.

Our modeling approach differs from [Xu and Bastani \(2025\)](#) in two key aspects. First, our setting ensures that the shared parameter $\boldsymbol{\beta}_{k0}$ is identifiable, in contrast to [Xu and Bastani \(2025\)](#), where the global parameter is not identifiable. This identifiability enhances the robustness of our estimation procedure, as non-identifiable models can be highly sensitive to the choice of hyperparameters — making them difficult to tune in practice. Second, unlike the sparse heterogeneity assumption adopted in [Xu and Bastani \(2025\)](#) — where only a few components of $\boldsymbol{\beta}_{k,j} - \boldsymbol{\beta}_{k0}$ are assumed to be nonzero — we assume that $\boldsymbol{\beta}_{k,j} - \boldsymbol{\beta}_{k0}$ follows a zero-mean normal distribution. In our simulation studies, our method consistently outperformed the RMBANDIT algorithm from [Xu and Bastani \(2025\)](#),




Fig. 1. Graphical model of our hierarchical Bayesian framework. The shared β_{k0} serves as the prior mean for each instance-specific $\beta_{k,j}$, while the deviation $\beta_{k,j} - \beta_{k0}$ characterizes how each bandit deviates from the shared structure.

even in settings that satisfy the sparse heterogeneity assumption. This observation is interesting. One possible explanation is that RMBANDIT may not fully exploit the underlying sparsity structure, whereas a better-calibrated prior may allow our method to more effectively borrow information from similar tasks, even when the sparse heterogeneity assumption holds.

To evaluate the performance of our sequential decision-making policy, we use cumulative expected regret, a standard metric in the analysis of contextual bandit problems. This metric quantifies the performance gap between our policy and an optimal policy with perfect knowledge of the underlying arm parameters. Specifically, at each time step t , upon observing the bandit Z_t , we define an optimal policy π_t^* that knows the true arm parameters $\{\beta_{k,Z_t}\}_{k \in [K]}$ in advance and always selects the arm with the highest expected reward:

$$\pi_t^* = \arg \max_{k \in [K]} \mathbf{x}_t^\top \beta_{k,Z_t},$$

where $\mathbf{x}_t$ is the context vector observed at time t . Let π_t denote the arm selected by the algorithm at time step t given the bandit Z_t . The expected regret at time t for bandit instance Z_t is then

$$r_t = \mathbb{E} \left[\mathbf{x}_t^\top \beta_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \beta_{\pi_t, Z_t} \right].$$

The cumulative regret over n time steps is the sum of the per-step regrets:

$$R_n = \sum_{t=1}^n \mathbb{E} \left[\mathbf{x}_t^\top \beta_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \beta_{\pi_t, Z_t} \right].$$

We also consider the instance-specific cumulative regret for each bandit instance j , defined as

$$R_{j,n} = \sum_{t=1}^n \mathbb{E} \left[\mathbb{I}(Z_t = j) \left(\mathbf{x}_t^\top \beta_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \beta_{\pi_t, Z_t} \right) \right],$$

where $\mathbb{I}(\cdot)$ is the indicator function and $n_j = \sum_{t=1}^n \mathbb{I}(Z_t = j)$ is the number of arrivals at bandit j . The instance-specific cumulative regret provides insight into our policy's performance for each bandit, especially in data-poor settings, and reflects how effectively the algorithm leverages shared information from other instances. It is worth noting that [Xu and Bastani \(2025\)](#) define regret based on the oracle policy. In contrast, we define regret with

respect to the actual policy π_t^* . A comparison of the two approaches is provided in Appendix F, where empirical results show very similar performance.

4 Estimation

We propose a hierarchical Bayesian approach for our model. Specifically, we introduce a prior distribution over the shared parameter β_{k0} and the individual deviations $\beta_{k,j} - \beta_{k0}$. This hierarchical structure facilitates information sharing across bandit instances while accommodating instance-specific variations. The hierarchical model is formally defined as follows:

$$\begin{aligned} y_{k,j,t} \mid \beta_{k,j} &\sim \mathcal{N}(\mathbf{x}_t^\top \beta_{k,j}, \sigma_k^2); \\ \beta_{k,j} \mid \beta_{k0} &\sim \mathcal{N}(\beta_{k0}, \Sigma_k); \\ \beta_{k0} &\sim \mathcal{N}(\mathbf{0}, \lambda^{-1} \mathbf{I}). \end{aligned}$$

In Section 4.1, we present the estimators $\hat{\beta}_{k,j}$ for $\beta_{k,j}$ along with their corresponding covariance matrices. In Section 4.2, we focus on predicting the expected payoff of arm k in bandit j for a new context $\mathbf{x}_t$, denoted by $\mu_{k,j,t} = \mathbf{x}_t^\top \beta_{k,j}$, and quantifying its associated uncertainty. In Section 4.3, we describe the estimation of the variance parameters Σ_k and σ_k^2 .

4.1 Estimation of $\beta_{k,j}$ and its Uncertainty

We introduce the estimator $\hat{\beta}_{k,j,t}$ for $\beta_{k,j}$ and quantify its uncertainty under the assumption that Σ_k and σ_k^2 are known. Let $\mathbb{T}_{k,j,t}$ denote the set of time steps when arm k of bandit j is pulled before time step t , and define $T_{k,j,t} = |\mathbb{T}_{k,j,t}|$ as its cardinality. Let $\mathbf{y}_{k,j,t} = (y_{k,j,s})_{s \in \mathbb{T}_{k,j,t}}^\top$ be the column vector of observed rewards, $\boldsymbol{\epsilon}_{k,j,t} = (\epsilon_{k,j,s})_{s \in \mathbb{T}_{k,j,t}}^\top$ the corresponding reward noise, and $\mathbf{X}_{k,j,t} = (\mathbf{x}_s)_{s \in \mathbb{T}_{k,j,t}}^\top$ a $T_{k,j,t} \times d$ matrix of the associated context vectors. The model (1) can be rewritten in matrix form as

$$\mathbf{y}_{k,j,t} = \mathbf{X}_{k,j,t} \beta_{k,j} + \boldsymbol{\epsilon}_{k,j,t}.$$

Noting that both $\beta_{k,j}$ and the noise are random, the covariance matrix $\mathbf{V}_{k,j,t}$ of the reward vector $\mathbf{y}_{k,j,t}$ is given by

$$\mathbf{V}_{k,j,t} = \mathbf{X}_{k,j,t} \Sigma_k \mathbf{X}_{k,j,t}^\top + \sigma_k^2 \mathbf{I}_{T_{k,j,t}},$$

where $\mathbf{I}_{T_{k,j,t}}$ is the identity matrix of size $T_{k,j,t}$. In this expression, the first term $\mathbf{X}_{k,j,t} \Sigma_k \mathbf{X}_{k,j,t}^\top$ captures the variability due to $\beta_{k,j}$, while the second term $\sigma_k^2 \mathbf{I}_{T_{k,j,t}}$ accounts for the reward noise.

With these notations in place, we now describe the estimation process in detail. Our procedure unfolds in three steps to effectively utilize the hierarchical structure of the model and ensure accurate arm parameter estimation. First, we use the instance-specific data to estimate the instance-specific $\beta_{k,j}$, conditioned on the shared parameter β_{k0} . This conditional expectation, denoted as $\tilde{\beta}_{k,j,t}$, is expressed as a linear function of β_{k0} . In the second step, we estimate the shared parameter β_{k0} using the aggregated data from all instances, denoted as $\hat{\beta}_{k0,t}$. Finally, we substitute the estimated $\hat{\beta}_{k0,t}$ into $\tilde{\beta}_{k,j,t}$ to obtain the final arm parameter estimate $\hat{\beta}_{k,j,t}$. We now proceed to derive the estimation results.

Step 1: Conditional Expectation of $\beta_{k,j}$. Given the rewards $\mathbf{y}_{k,j,t}$ that have been obtained and β_{k0} , the posterior distribution of $\beta_{k,j}$ can be expressed as

$$\mathbb{P}(\beta_{k,j} \mid \mathbf{y}_{k,j,t}, \beta_{k0}) \propto \exp \left\{ -\frac{1}{2\sigma_k^2} \left\| \mathbf{y}_{k,j,t} - \mathbf{X}_{k,j,t} \beta_{k,j} \right\|_2^2 - \frac{1}{2} \left(\beta_{k,j} - \beta_{k0} \right)^\top \Sigma_k \left(\beta_{k,j} - \beta_{k0} \right) \right\}.$$

Direct calculation shows that

$$\boldsymbol{\beta}_{k,j} \mid \mathbf{y}_{k,j,t}, \boldsymbol{\beta}_{k0} \sim \mathcal{N}(\tilde{\boldsymbol{\beta}}_{k,j,t}, \sigma_k^2 \tilde{\mathbf{C}}_{k,j,t}),$$

where

$$\begin{aligned} \tilde{\boldsymbol{\beta}}_{k,j,t} &= \sigma_k^2 \tilde{\mathbf{C}}_{k,j,t} \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\beta}_{k0} + \tilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t}^\top \mathbf{y}_{k,j,t}; \\ \tilde{\mathbf{C}}_{k,j,t} &= (\mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t} + \sigma_k^2 \boldsymbol{\Sigma}_k^{-1})^{-1}. \end{aligned} \quad (2)$$

Step 2: Estimation of $\boldsymbol{\beta}_{k0}$. Now we derive the posterior of $\boldsymbol{\beta}_{k0}$ given $\{\mathbf{y}_{k,j,t}\}_{j \in [N]}$. We can rewrite the hierarchical Bayesian model as follows:

$$\begin{aligned} \mathbf{y}_{k,j,t} \mid \boldsymbol{\beta}_{k0} &\sim \mathcal{N}(\mathbf{X}_{k,j,t} \boldsymbol{\beta}_{k0}, \mathbf{V}_{k,j,t}); \\ \boldsymbol{\beta}_{k0} &\sim \mathcal{N}(\mathbf{0}, \lambda^{-1} \mathbf{I}). \end{aligned}$$

Given the rewards $\{\mathbf{y}_{k,j,t}\}_{j \in [N]}$, the posterior distribution of $\boldsymbol{\beta}_{k0}$ can be expressed as

$$\mathbb{P}(\boldsymbol{\beta}_{k0} \mid \{\mathbf{y}_{k,j,t}\}_{j \in [N]}) \propto \exp \left\{ -\frac{1}{2} \sum_{j=1}^N (\mathbf{y}_{k,j,t} - \mathbf{X}_{k,j,t} \boldsymbol{\beta}_{k0})^\top \mathbf{V}_{k,j,t}^{-1} (\mathbf{y}_{k,j,t} - \mathbf{X}_{k,j,t} \boldsymbol{\beta}_{k0}) - \frac{\lambda}{2} \boldsymbol{\beta}_{k0}^\top \boldsymbol{\beta}_{k0} \right\}.$$

Then, we have

$$\boldsymbol{\beta}_{k0} \mid \{\mathbf{y}_{k,j,t}\}_{j \in [N]} \sim \mathcal{N}(\hat{\boldsymbol{\beta}}_{k0,t}, \boldsymbol{\Phi}_{k0,t}),$$

where

$$\begin{aligned} \hat{\boldsymbol{\beta}}_{k0,t} &= \boldsymbol{\Phi}_{k0,t} \sum_{j=1}^N \mathbf{X}_{k,j,t}^\top \mathbf{V}_{k,j,t}^{-1} \mathbf{y}_{k,j,t}; \\ \boldsymbol{\Phi}_{k0,t} &= \left(\sum_{j=1}^N \mathbf{X}_{k,j,t}^\top \mathbf{V}_{k,j,t}^{-1} \mathbf{X}_{k,j,t} + \lambda \mathbf{I} \right)^{-1}. \end{aligned}$$

Step 3: Estimation of $\boldsymbol{\beta}_{k,j}$. To derive the posterior expectation and variance of the arm parameter vector $\boldsymbol{\beta}_{k,j}$ given the observed rewards $\{\mathbf{y}_{k,\ell,t}\}_{\ell \in [N]}$, we utilize the law of total expectation and the law of total variance (Hong et al., 2022), and obtain

$$\boldsymbol{\beta}_{k,j} \mid \{\mathbf{y}_{k,\ell,t}\}_{\ell \in [N]} \sim \mathcal{N}(\hat{\boldsymbol{\beta}}_{k,j,t}, \mathbf{C}_{k,j,t}),$$

where

$$\begin{aligned} \hat{\boldsymbol{\beta}}_{k,j,t} &= \sigma_k^2 \tilde{\mathbf{C}}_{k,j,t} \boldsymbol{\Sigma}_k^{-1} \hat{\boldsymbol{\beta}}_{k0,t} + \tilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t}^\top \mathbf{y}_{k,j,t}; \\ \mathbf{C}_{k,j,t} &= \sigma_k^2 \tilde{\mathbf{C}}_{k,j,t} + \sigma_k^4 \tilde{\mathbf{C}}_{k,j,t} \boldsymbol{\Sigma}_k^{-1} \boldsymbol{\Phi}_{k0,t} \boldsymbol{\Sigma}_k^{-1} \tilde{\mathbf{C}}_{k,j,t}. \end{aligned} \quad (3)$$

Calculating the posterior variance in hierarchical Bayesian models can be computationally intensive, especially when dealing with large datasets. To address this challenge, we employ the Woodbury matrix identity. In (3), noting that both $\tilde{\mathbf{C}}_{k,j,t}$ and $\boldsymbol{\Sigma}_k$ are $d \times d$ matrices, it is easy to compute when d is not too large. The difficulty in computing $\mathbf{C}_{k,j,t}$ mainly lies in $\boldsymbol{\Phi}_{k0,t}$, where $\mathbf{V}_{k,j,t}$ is an $T_{k,j,t} \times T_{k,j,t}$ matrix. As the amount of data increases gradually, the computational difficulty also increases. Recall that $\mathbf{V}_{k,j,t} = \mathbf{X}_{k,j,t} \boldsymbol{\Sigma}_k \mathbf{X}_{k,j,t}^\top + \sigma_k^2 \mathbf{I}_{T_{k,j,t}}$. Applying the Woodbury matrix identity, we have

$$\mathbf{V}_{k,j,t}^{-1} = \sigma_k^{-2} \mathbf{I}_{T_{k,j,t}} - \sigma_k^{-2} \mathbf{X}_{k,j,t} \tilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t}^\top. \quad (4)$$

In (4), we transform the problem of inverting a $T_{k,j,t} \times T_{k,j,t}$ matrix into the problem of inverting a $d \times d$ matrix, which significantly accelerates the calculation of the posterior variance, making our estimation process more computationally tractable.

4.1.1 Estimation error bound analysis. Now we present the upper bound of our estimation error. Prior to this, it is necessary to propose the following assumption, which is not exceptional and has been widely adopted in existing literature (Xu and Bastani, 2025; Hong et al., 2022; Agrawal and Goyal, 2013).

ASSUMPTION 1. The ground truth $\beta_{k,j}$ is bounded in Euclidean norm, i.e., $\|\beta_{k,j}\|_2 \leq c_\beta$, for all $k \in [K], j \in [N]$. The context $\mathbf{x}$ is bounded in Euclidean norm, i.e., $d^{-1}\|\mathbf{x}\|_2^2 \leq c_x$, for all $\mathbf{x}$. The eigenvalues of covariance matrix Σ_k^{-1} are bounded, i.e., $0 < \lambda_d \leq \lambda(\Sigma_k^{-1}) \leq \lambda_1$, for all $k \in [K]$. The noise variances of all arms are identical, i.e., $\sigma_k^2 = \sigma^2$, for all $k \in [K]$ (This assumption is made solely for the simplicity of the proof and does not need to be satisfied in subsequent experiments.).

It is worth noting that our assumption regarding the boundedness of $\mathbf{x}$ requires that its infinity norm be bounded (i.e., $\|\mathbf{x}\|_\infty < \infty$). Since the ℓ_2 and ℓ_∞ norms are related by the inequality $\|\mathbf{x}\|_\infty \leq \|\mathbf{x}\|_2 \leq \sqrt{d} \cdot \|\mathbf{x}\|_\infty$, this assumption directly yields the requisite bounds utilized throughout the subsequent technical analysis.

THEOREM 4.1. Under Assumption 1, the estimator $\hat{\beta}_{k,j,t}$ that incorporates prior information satisfies the following inequality with probability at least $1 - \delta$, for any fixed $t \geq 1$ and $\mathbf{x} \in \mathbb{R}^d$:

$$\begin{aligned} \|\tilde{\mathbf{C}}_{k,j,t}^{-\frac{1}{2}} (\hat{\beta}_{k,j,t} - \beta_{k,j})\|_2 &\leq \sigma \alpha_t(\delta), \\ |\mathbf{x}^\top (\hat{\beta}_{k,j,t} - \beta_{k,j})| &\leq \alpha_t(\delta) \|\mathbf{x}\|_{\mathbf{C}_{k,j,t}}, \end{aligned}$$

where

$$\alpha_t(\delta) = 2\sqrt{dm_1 \log\left(\frac{m_2 + tc_x}{m_3\delta}\right)} + Nm_4 = O(\sqrt{d \log(t/\delta)}),$$

$$m_1 = \max\{\lambda_1/\lambda, 1\}, m_2 = \max\{\lambda, \sigma^2\lambda_1\}, m_3 = \sqrt{\lambda\lambda_d\sigma^2/2} \text{ and } m_4 = c_\beta\sqrt{\lambda_1}.$$

Theorem 4.1 provides an explicit characterization of the upper bound for the estimation error under the incorporation of prior information. Most existing studies on multi-bandit problems (Hong et al., 2022; Aouali et al., 2023) focus on analyzing Bayesian regret as the primary performance metric. Building on Theorem 4.1, we derive the frequentist regret in Section 6 and fill this research gap by establishing rigorous theoretical results for the frequentist upper bound.

We briefly outline the proof for Theorem 4.1 here; full details are provided in Appendix A. The estimation error is decomposed into two terms:

$$\tilde{\mathbf{C}}_{k,j,t}^{-\frac{1}{2}} (\hat{\beta}_{k,j,t} - \beta_{k,j}) = \sigma^2 \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \Sigma_k^{-1} (\hat{\beta}_{k,0,t} - \beta_{k,j}) + \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \mathbf{X}_{k,j,t} \epsilon_{k,j,t},$$

where the first term represents the deviation due to the prior-informed estimator, and the second term captures the randomness introduced by observational noise. The noise term is bounded using the theoretical result established in Abbasi-Yadkori et al. (2011), stated in Lemma E.1. The primary challenge in establishing Theorem 4.1 lies in handling the first term, which involves the prior-informed estimation error $\hat{\beta}_{k,0,t}$. Its L_2 -norm can be bounded by that of the prior-induced error, $\hat{\beta}_{k,0,t} - \beta_{k,j}$, which consists of a term involving the matrix norms of the true parameters across all bandits and a noise term. We carefully examine the boundedness of these norms and apply the theoretical results from Abbasi-Yadkori et al. (2011) to establish the bound.

4.2 Prediction of $\mu_{k,j,t}$ under Context $\mathbf{x}_t$ and its Uncertainty

From (3), the predicted reward for a new context $\mathbf{x}_t$ is given by

$$\hat{\mu}_{k,j,t} = \mathbf{x}_t^\top \hat{\beta}_{k,j,t} = \sigma_k^2 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1} \hat{\beta}_{k,0,t} + \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t}^\top \mathbf{y}_{k,j,t}. \quad (5)$$

Now we measure the uncertainty of $\hat{\mu}_{k,j,t}$ in (5) for efficient exploration. We measure the uncertainty by the mean squared error $\mathbb{E}[(\hat{\mu}_{k,j,t} - \mu_{k,j,t})^2 \mid \{\mathbf{y}_{k,j,t}\}_{j \in [N]}]$, where $\mu_{k,j,t} = \mathbf{x}_t^\top \boldsymbol{\beta}_{k,j}$. Then we have

$$\mathbb{E}[(\hat{\mu}_{k,j,t} - \mu_{k,j,t})^2 \mid \{\mathbf{y}_{k,j,t}\}_{j \in [N]}] = \mathbf{x}_t^\top \mathbf{C}_{k,j,t} \mathbf{x}_t =: \tau_{k,j,t}^2. \quad (6)$$

4.3 Estimation of Σ_k and σ_k^2

Both the estimator $\hat{\mu}_{k,j,t}$ and its uncertainty $\tau_{k,j,t}^2$ depend on Σ_k and σ_k^2 . When these parameters are unknown, they can be estimated from the data. Substituting the estimates of Σ_k and σ_k^2 into the expressions for $\hat{\mu}_{k,j,t}$ and $\tau_{k,j,t}^2$ yields the empirical Bayesian estimators. This approach is referred to as empirical Bayesian multi-bandit learning.

We estimate σ_k^2 by

$$\hat{\sigma}_{k,t}^2 = \frac{\sum_{j=1}^N \|\mathbf{y}_{k,j,t} - \mathbf{X}_{k,j,t} \hat{\boldsymbol{\beta}}_{k,j,t}\|_2^2}{\max\{\sum_{j=1}^N T_{k,j,t} - d - 1, 1\}}.$$

Estimating Σ_k involves the problem of covariance matrix estimation. A common approach is to use the sample covariance matrix. However, when the dimensionality of the variables increases with the number of bandit instances N , the estimation becomes a high-dimensional problem. To address this issue, we adopt the covariance matrix estimation method proposed by [Bickel and Levina \(2008\)](#), which is based on the following sparsity assumption: Σ_k belongs to a class of covariance matrices defined by

$$C_q(c_0(d), M, M_0) = \left\{ \Sigma : \sigma_{ii} \leq M, \sum_{j=1}^d |\sigma_{ij}|^q \leq c_0(d), \forall i; \lambda_{\min}(\Sigma) \geq M_0 > 0 \right\},$$

where $0 \leq q < 1$. Denote $\hat{\boldsymbol{\beta}}_{k,j,t}^{\text{ols}} = (\mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t})^{-1} \mathbf{X}_{k,j,t}^\top \mathbf{y}_{k,j,t}$. We obtain the sample covariance matrix

$$\mathbf{S}_{k,t} = \frac{1}{N} \sum_{j=1}^N \hat{\boldsymbol{\beta}}_{k,j,t}^{\text{ols}} \hat{\boldsymbol{\beta}}_{k,j,t}^{\text{ols}\top} - \frac{1}{N^2} \sum_{j=1}^N \hat{\boldsymbol{\beta}}_{k,j,t}^{\text{ols}} \sum_{j=1}^N \hat{\boldsymbol{\beta}}_{k,j,t}^{\text{ols}\top}.$$

Therefore, the threshold matrix estimator is defined as

$$\hat{\Sigma}_{k,t} = (s_{ij} \mathbb{I}(|s_{ij}| \geq \gamma)),$$

where s_{ij} denotes the element in the i -th row and j -th column of the matrix $\mathbf{S}_{k,t}$, and $\mathbb{I}(\cdot)$ is the indicator function, which takes the value 1 if the condition is satisfied and 0 otherwise. We follow the approach of [Bickel and Levina \(2008\)](#) to select the thresholding parameter γ . In their work, [Bickel and Levina \(2008\)](#) establish the following consistency result.

LEMMA 4.2. Assume that $\Sigma_k \in C_q(c_0(d), M, M_0)$ and $\frac{\log d}{N} = o(1)$, then we have that

$$\|\hat{\Sigma}_{k,t} - \Sigma_k\| = O_p\left(c_0(d) (N^{-1} \log d)^{(1-q)/2}\right).$$

The use of the thresholded covariance matrix estimator is a key component of our approach. It provides an automatic, data-driven approach for identifying meaningful correlations across bandit instances. Broadly speaking, the method retains strong correlations between instances, while setting weak correlations to zero — effectively removing those that are insufficient to support across-instance learning.

5 Algorithms

We propose two algorithms for Empirical Bayesian Multi-Bandit (ebm). Based on the posterior distribution of the arm parameters derived in Section 4.1, we first introduce a sampling-based algorithm that leverages the posterior of $\beta_{k,j}$. We refer to this algorithm as EBM_{TS}, which stands for the posterior sampling strategy. It is worth noting, however, that EBM_{TS} may not represent true posterior sampling, as the prior parameters Σ_k and σ_k^2 are estimated using a frequentist approach. The second algorithm follows the UCB framework and builds on the ReUCB algorithm introduced in [Zhu and Kveton \(2022a,b\)](#). We refer to this variant as EBM_{UCB}, which incorporates both random effects and contextual information in a unified framework.

Algorithm 1 EBM_{TS} and EBM_{UCB} in Empirical Bayesian Multi-Bandit.

```

1: Input: hyperparameters  $\lambda, a$ .
2: for  $t = 1, 2, \dots, n$  do
3:   Observe an arrival at instance  $Z_t = j$ .
4:   Observe context  $\mathbf{x}_t$ .
5:   for  $k = 1, 2, \dots, K$  do
6:     Obtain  $\hat{\mu}_{k,j,t}$  from (5) and  $\hat{\tau}_{k,j,t}^2$  from (6).
7:     Define
        EBMTS:  $U_{k,j,t} = \mathbf{x}_t^\top \check{\beta}_{k,j,t}$  with  $\check{\beta}_{k,j,t} \sim \mathcal{N}(\hat{\beta}_{k,j,t}, \alpha_t^2(\delta) \mathbf{C}_{k,j,t})$ , OR
        EBMUCB:  $U_{k,j,t} = \hat{\mu}_{k,j,t} + \alpha_t(\delta) \tau_{k,j,t}$ .
8:   end for
9:   if  $t \leq K$  then  $\pi_t \leftarrow t$  else  $\pi_t \leftarrow \arg \max_{k \in [K]} U_{k,j,t}$ .
10:  Pull action  $\pi_t$  of bandit  $j$  and observe reward  $y_{\pi_t,j,t}$ .
11:  Update  $\hat{\beta}_{\pi_t,\ell,t}$  and  $\mathbf{C}_{\pi_t,\ell,t}$  for  $\ell \in [N]$ , and update  $\hat{\Sigma}_{\pi_t}, \hat{\sigma}_{\pi_t}^2$ .
12: end for

```

We summarize both EBM_{TS} and EBM_{UCB} in Algorithm 1. Given a new context $\mathbf{x}_t$ at time step t , we apply (5) and (6) to compute the predicted reward:

$$\hat{\mu}_{k,j,t} = \mathbf{x}_t^\top \hat{\beta}_{k,j,t},$$

and its corresponding uncertainty

$$\tau_{k,j,t}^2 = \mathbf{x}_t^\top \mathbf{C}_{k,j,t} \mathbf{x}_t.$$

We now describe the two algorithms in detail.

EBM_{TS}. We presented the posterior distribution in (3), which is $\beta_{k,j} \mid \{y_{k,\ell,t}\}_{\ell \in [N]} \sim \mathcal{N}(\hat{\beta}_{k,j,t}, \mathbf{C}_{k,j,t})$. This allows us to sample directly from the posterior distribution. At time step t , the sampling-based selection rule is given by

$$\pi_{j,t} = \arg \max_{k \in [K]} \mathbf{x}_t^\top \check{\beta}_{k,j,t} \text{ with } \check{\beta}_{k,j,t} \sim \mathcal{N}(\hat{\beta}_{k,j,t}, \alpha_t^2(\delta) \mathbf{C}_{k,j,t}),$$

It should be noted that $\alpha_t(\delta)$ is of the order of $\sqrt{\log t}$. Therefore, in practical experiments, we set $\alpha_t(\delta) = a\sqrt{\log t}$, where $a > 0$ is a tunable hyperparameter to control the degree of exploration.

EBM_{UCB}. We adopt a UCB-based exploration strategy to address the uncertainty in estimating $\hat{\mu}_{k,j,t}$ for sequential decision-making. The exploration bonus for arm k of bandit j at time step t is given by $\alpha_t(\delta) \tau_{k,j,t}$. Accordingly, the selection rule of EBM_{UCB} at time step t is

$$\pi_{j,t} = \arg \max_{k \in [K]} U_{k,j,t} \text{ with } U_{k,j,t} = \hat{\mu}_{k,j,t} + \alpha_t(\delta) \tau_{k,j,t},$$

where, similar to EBM-TS, in practical experiments, we set the same $\alpha_t(\delta)$ as above.

6 Regret Analysis

In this section, we present the frequentist cumulative regret upper bounds for the algorithms EBMUCB and EBM-TS, which differ from the Bayesian regret bound reported in previous literature (Hong et al., 2022).

THEOREM 6.1. *For any $\delta > 0$, the overall cumulative regret of EBMUCB is at most*

$$\begin{aligned} R_n &\leq 2\alpha_n(\delta) \sqrt{ndK \left[c_1 \sum_{j=1}^N \log(1 + c_2 n_j) + c_3 \log(1 + c_4 N) \right]} + 2\sqrt{dc_x} c_\beta K N n \delta \\ &= O \left(d^{3/2} \sqrt{nK \log \frac{n}{\delta} \left(\sum_{j=1}^N \log n_j + d \log N \right)} \right), \end{aligned}$$

where $c_1 = \frac{\lambda_d^{-1} dc_x}{\log(1 + \sigma^{-2} \lambda_d^{-1} dc_x)} = O(d)$, $c_2 = \sigma^{-2} \lambda_d^{-1} c_x$, $c_3 = \frac{\lambda_d^{-2} \lambda_1^{-1} dc_x (1 + \sigma^{-2} \lambda_d^{-1} dc_x)}{\log(1 + \sigma^{-2} \lambda_d^{-2} \lambda_1^{-1} dc_x)} = O(d^2)$, and $c_4 = \lambda_1 \lambda^{-1}$.

THEOREM 6.2. *For any $\delta > 0$, the overall cumulative regret of EBM-TS is at most*

$$\begin{aligned} R_n &\leq 2C\alpha'_n(\delta) \sqrt{ndK \left[c_1 \sum_{j=1}^N \log(1 + c_2 n_j) + c_3 \log(1 + c_4 N) \right]} + 2\sqrt{dc_x} c_\beta K N (n\delta + \pi^2/6) \\ &= O \left(d^2 \sqrt{nK \log \frac{n}{\delta} \log n \left(\sum_{j=1}^N \log n_j + d \log N \right)} \right), \end{aligned}$$

where $C = \max_{t \geq 1} \left[\frac{2}{\frac{1}{2\sqrt{2\pi e}} - \frac{1}{t^2}} \right] + 1$, $\alpha'_n(\delta) = (\gamma_n + 1)\alpha_n(\delta)$, and $\gamma_n = \sqrt{2d \log(dn^2)}$.

Theorems 6.1 and 6.2 establish the frequentist upper bounds on regret, where each component admits a clear interpretation: $\alpha_n(\delta)$ represents the contribution from the estimation error bound; $\sqrt{c_1 ndK \sum_{j=1}^N \log(1 + c_2 n_j)}$ reflects the regret associated with learning instance-specific parameters; and $\sqrt{c_3 ndK \log(1 + c_4 N)}$ captures the regret incurred in learning the arm-level prior parameters. These results show that the cost of learning the shared prior grows only logarithmically with the number of instances, while the instance-specific learning cost depends on the individual sample sizes through $\sum_{j=1}^N \log(1 + c_2 n_j)$. This is consistent with our empirical results, where information shared through the common prior improves performance across instances, particularly when some instances have limited data.

Compared with the bound in Theorem 6.1 for EBMUCB, the regret bound in Theorem 6.2 for EBM-TS includes an additional factor, γ_n , arising from sampling error, reflecting the inherent variability introduced by its sampling-based nature.

Compared with the Bayesian regret upper bound for multi-bandit problems (Hong et al., 2022), our frequentist regret bounds exhibit a similar structural form. A key difference, however, lies in the scaling: the Bayesian bound scales with $d^{\frac{3}{2}}$ because Bayesian regret inherently accounts only for sampling error and disregards estimation error. In contrast, our algorithm EBM-TS, when analyzed under the frequentist regret framework, must account for this estimation error. This necessity introduces an additional factor of $\sqrt{d}$, leading to a higher-order dependence on d . Despite this, our bounds remain competitive. Crucially, our algorithms exhibit superior empirical performance, which we attribute to their ability to leverage shared structure for more efficient learning.

The proofs of Theorems 6.2 and 6.1 follow three main steps. First, we decompose the cumulative regret into two components: the estimation error and the variance summation term, $\mathcal{V}_n$, defined as $\mathcal{V}_n = \sum_{t=1}^n \mathbf{x}_t^\top \mathbf{C}_{\pi_t, Z_t, t} \mathbf{x}_t$. Next, we apply Theorem 4.1 to control the estimation error. Finally, Lemmas D.1 and D.2 in Appendix D are used to bound the variance term $\mathcal{V}_n$.

7 Experiments

In this section, we present experiments on both synthetic and real-world datasets to evaluate the performance of our proposed algorithms. In all experiments, we set $\lambda = 0.001$ and $a = 0.1$ for both EBMTS and EBMUCB. A sensitivity analysis of these hyperparameters is provided in Appendix G. The results show that the performance is robust to their values within reasonable ranges. In practice, these hyperparameters can be tuned if necessary. To ensure robustness, each experiment is repeated with 100 different random seeds. All computations are performed on a computer equipped with an Apple M3 processor, 24 GB of memory, and OS X operating system.


We compare our methods against the following baseline algorithms:

- **RMBANDIT** (Xu and Bastani, 2025): It is designed for multi-task learning under the assumption of sparse heterogeneity across bandit instances. For hyperparameters, we take $\eta_0 = \eta_{1,0} = 0.2$, $h = 15$ and $q = 50$, the same as (Xu and Bastani, 2025).
- **OLSBANDIT** (Goldenshluger and Zeevi, 2013): It uses ordinary least squares (OLS) regression to estimate the parameters for each bandit instance independently, without leveraging any shared structure. For hyperparameters, we take $h = 15$ and $q = 1$, the same as Xu and Bastani (2025).
- **LINTS** (Agrawal and Goyal, 2013): It is based on a sampling strategy but it does not perform multi-task learning.
- **LINUCB** (Li et al., 2010): It uses a UCB-based exploration but does not perform multi-task learning.


7.1 Experiments on Synthetic Data

We generate synthetic data to simulate a multi-bandit environment using the hierarchical Bayesian model. The shared parameters β_{k0} are drawn from a normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$, and the instance-specific parameters $\beta_{k,j}$ are then drawn from $\mathcal{N}(\beta_{k0}, \Sigma_k)$, where the covariance matrices Σ_k are constructed as $\mathbf{b}\mathbf{b}^\top + \mathbf{I}$, with $\mathbf{b} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Context vectors $\mathbf{x}_t$ are drawn from a mixture of Gaussian distributions. Specifically, each element of $\mathbf{x}_t$ is independently sampled from $\mathcal{N}(-1, 1)$ with probability 0.5, and from $\mathcal{N}(1, 1)$ with probability 0.5. We use Gaussian noise with a standard deviation of $\sigma_k = 1$ for all arms. In the data-balanced setting, we assign equal sampling probabilities with $p_j = 1/N$ for all j . In the data-poor setting, we set $p_1 = 0.1p_i$ and $p_i = p_j$ for all $i, j \geq 2$, creating a scenario in which the first bandit receives significantly fewer samples than the others.

In Figure 2a, we present the cumulative regret and the instance-specific cumulative regret for the first bandit as an example in the data-balanced setting, respectively. Similarly, in Figure 2b, we show the corresponding results for the data-poor setting, where the first bandit receives fewer samples. The plots in Figure 2 demonstrate that our algorithms, EBMTS and EBMUCB, consistently achieve significantly lower cumulative regret compared to the baseline methods. By effectively leveraging shared information across bandit instances, our algorithms exhibit faster convergence and reduced regret over time in both settings. In contrast, LINTS, LINUCB, and OLSBANDIT do not share information across bandit instances, resulting in higher cumulative regret that continues to grow over time due to less efficient exploration and exploitation. Although RMBANDIT incorporates information sharing, its performance remains suboptimal. It initially emphasizes exploration, leading to a substantial accumulation of regret in the early stages. It is worth noting that the cumulative regret curves under the data-poor and data-balanced settings are nearly identical in shape. Each arm shares a common prior across instances, while the data-poor instance contributes negligibly to the aggregate regret. Thus, the total regret is dominated by the data-rich instances, leaving its overall shape largely unchanged.



(a) Cumulative regret in the data-balanced setting.



(b) Cumulative regret in the data-poor setting.

Fig. 2. Performance under $N = 10$, $K = 5$, and $d = 3$.

Our model framework is built on Gaussian assumptions regarding the coefficients across arms. We now investigate whether these Gaussian assumptions restrict its applicability in broader settings. Specifically, we consider a sparse heterogeneity assumption, similar to the setup in [Xu and Bastani \(2025\)](#), under which our hierarchical model is misspecified. Figure 3 reports the cumulative regret results. The results show that our algorithms, **EBM_{TS}** and **EBM_{UCB}**, continue to outperform **RMBANDIT** even under the sparse heterogeneity assumption. This demonstrates that our methods can effectively exploit shared information for joint learning and capture contextual signals, despite model misspecification. Overall, these results highlight the robustness and broad applicability of our approach.




Fig. 3. Performance under sparse heterogeneity.

We also evaluate the performance of the algorithms under various parameter settings, including different context distributions, numbers of bandit instances (N), numbers of arms (K), and dimensions of the context vectors (d). The corresponding results are presented in Figure 4.


In Figure 4a, we present the performance of the algorithms across different context distributions, focusing on the uniform and mixed Gaussian settings. The uniform distribution, where each context element is sampled uniformly from $[-1, 1]$, represents a simple continuous case that satisfies the covariate diversity condition discussed in Bastani et al. (2021). Under such conditions, exploration requirements can be significantly reduced – potentially leading to near-exploration-free behavior. In contrast, the mixed Gaussian distribution is more complex and violates the covariate diversity condition. In this setting, EBM-TS and EBM-UCB achieve notably lower cumulative regret, demonstrating their robustness in handling complex and heterogeneous contextual information.

The effects of other key parameters – including the number of bandit instances (N), the number of arms (K), and the context dimensionality (d) – are presented in Figures 4b-4d, respectively. As N increases, the number of tasks grows, but so does the pool of shared information. Our algorithms effectively exploit this shared structure, consistently outperforming the baselines. Similarly, as K and d increase – making the learning problem more challenging – EBM-TS and EBM-UCB maintain strong performance. Between the two, EBM-TS and EBM-UCB exhibit comparable average cumulative regret, though EBM-UCB tends to show lower variability across trials. This stability is evident in the box plots, where EBM-UCB displays narrower interquartile ranges, indicating more consistent performance.

7.2 Experiments on Real-World Data

SARCOS Dataset. The SARCOS dataset¹ addresses a multi-output learning problem for modeling the inverse dynamics of a SARCOS anthropomorphic robot with seven degrees of freedom. Each sample includes 21 input features – comprising seven joint positions, seven joint velocities, and seven joint accelerations. This dataset has been widely used in the literature (Balduzzi and Ghifary, 2015; Zhang and Yang, 2021), and contains 44,484 training examples and 4,449 test examples. We treat each output (i.e., degree of freedom) as a separate arm. To

¹<https://gaussianprocess.org/gpml/data/>

Fig. 4. Performance under different context distributions, N , K , and d .

simulate a multi-task bandit environment, we first apply linear regression to the test dataset to estimate the model parameters and their variances. These estimates are then used to generate task-specific parameters for the $N = 30$ bandit instances we consider. Finally, we evaluate the performance of our algorithms on the training dataset to assess their effectiveness.

Activity Recognition Dataset. The UCI dataset titled "Activity recognition with healthy older people using a batteryless wearable sensor"² is designed to monitor the activities of healthy elderly individuals with the aim of reducing the occurrence of harmful events, such as falls. The dataset provides 60 *.csv files for room 1 and 28 *.csv files for room 2. Each *.csv file includes eight features from the W2ISP (Wearable Wireless Identification and Sensing Platform) sensor and the RFID (Radio Frequency Identification) reader, along with the label for each record. The labels indicate activities such as sitting on the bed, sitting on a chair, lying in bed, and walking. We

²<https://archive.ics.uci.edu/dataset/427/activity+recognition+with+healthy+older+people+using+a+batteryless+wearable+sensor>

focus on Room 1 due to its larger data volume. Through one-hot encoding of categorical features and principal component analysis (PCA) for dimensionality reduction on the features, an 8-dimensional feature set is finally obtained. After filtering out files with insufficient data, we retain 16 files, each treated as a separate task. The activity labels are interpreted as arms in the bandit problem framework. For each file, we split the data into training and testing sets using a 30%-70% ratio. A hierarchical Bayesian model is fitted to the training data to estimate the underlying environment parameters, which are then used to evaluate the algorithms on the corresponding test sets.

MovieLens 10M Dataset. The MovieLens dataset, which is widely used in the research of contextual bandits (Cella et al., 2020; Christakopoulou and Banerjee, 2018; Hong et al., 2023; Wan et al., 2021), comes in various sizes to accommodate different research needs. The MovieLens 10M dataset³ (Harper and Konstan, 2015), contains 10 million ratings for 10,677 movies from 69,878 users. As a first step, we complete the sparse rating matrix using singular value decomposition (SVD) with rank $d = 10$. More specifically, let $\mathbf{R}$ denote the rating matrix, where the element at the i -th row and j -th column represents the rating given by user i to movie j . Applying SVD to $\mathbf{R}$ yields the approximation $\mathbf{R} \approx \mathbf{U}\mathbf{V}^T$. Here, the i -th row of $\mathbf{U}$, denoted as $\mathbf{u}_i$, represents the features of user i , while the j -th row of $\mathbf{V}$, denoted as $\mathbf{v}_j$, represents the features of movie j . The rating of movie j by user i is then obtained by the dot product $\mathbf{u}_i^T \mathbf{v}_j$. Then we apply a Gaussian mixture model (GMM) with $K = 10$ clusters to the rows of $\mathbf{V}$. We set the prior parameters to the center and covariance estimated by the Gaussian Mixture Model (GMM). And we generate the parameter vectors for each task to simulate similar tasks. We set the number of tasks, N , to be 10.

Letter Recognition Dataset. The letter recognition dataset⁴ (Slate, 1991) is a multi-class classification dataset containing handwritten capital letters from various writers. It has been widely used in prior studies (Deshmukh et al., 2017; Wang et al., 2019). The dataset consists of 20,000 samples, each representing a capital letter from the English alphabet. Each sample is described by 16 numerical features that capture various attributes of the letter's shape and structure. We split it into training and test sets with a ratio of 30% for training and 70% for testing. For each category (arm), we fit a linear regression model on the training data to estimate the corresponding coefficients and variances. These estimates are then used to generate parameter vectors for each task, enabling the simulation of related tasks. In this experiment, we set $N = 30$.

We evaluate the performance of our algorithms, EBMTS and EBMUCB, on these four real-world datasets to assess their practical effectiveness and generalizability, as shown in Figure 5. Across diverse application domains, both algorithms consistently outperform the baseline methods. This superior performance highlights their ability to efficiently share information across related tasks, adapt to heterogeneous contextual structures, and maintain robustness under varying data complexities. These results demonstrate the practical value of our framework for real-world decision-making scenarios.

8 Conclusion

In this article, we introduced a novel hierarchical Bayesian framework that captures shared structures across multiple bandit instances while accounting for instance-specific variations. Building on this framework, we developed an empirical Bayesian multi-task approach for contextual bandits, addressing the limitations of existing methods and resulting in more efficient algorithms — EBMTS and EBMUCB — suitable for real-world applications.

Unlike prior methods that rely on the sparse heterogeneity assumption (Xu and Bastani, 2025), our approach imposes no such restrictions, allowing for more flexible learning across bandit instances. In addition, unlike previous hierarchical models (Hong et al., 2022; Wan et al., 2021), we explicitly estimate the prior covariance

³<https://grouplens.org/datasets/movielens/10m/>

⁴<https://archive.ics.uci.edu/dataset/59/letter+recognition>




Fig. 5. Performance on real-world datasets.

matrix, providing both practical guidance for implementation and deeper insight into the shared structure among arms. Our theoretical analysis, based on frequentist regret, rigorously characterizes the performance of the proposed algorithms, while extensive experiments on synthetic and real-world datasets demonstrate that **EBMTS** and **EBMUCB** consistently outperform existing methods in cumulative and instance-specific regret. These results highlight the ability of our framework to handle complex, heterogeneous environments, bridging the gap between theoretical guarantees and practical applicability.

However, there are several limitations to our work that merit discussion. First, our paper assumes Gaussian distributions and linear models for analytical convenience. While the framework could be extended to more general distributions and complex models, doing so would require approximate methods for posterior sampling. Second, although our framework effectively leverages shared information across tasks, it assumes that relationships between tasks are relatively consistent. This assumption may not hold in dynamic environments with rapidly changing task-specific dynamics. Extending our approach to such dynamic settings represents an important

direction for future research. Third, our theoretical analysis relies on the prior covariance matrix being full-rank. When $\text{rank}(\Sigma_k) = r < d$, the arm parameters lie in an r -dimensional subspace, with no variation along the remaining $d - r$ null-space directions. Extending the analysis to this low-rank setting would reveal the potential benefits of exploiting the underlying low-rank structure. Finally, the scalability of our methods may be challenging in large-scale environments, where high-dimensional feature spaces or a large number of bandit instances can lead to substantial computational overhead due to the complexity of managing and updating hierarchical structures and prior covariance matrices. More efficient implementations will need to be explored in future work.

Acknowledgments

The authors would like to thank the Associate Editor, Marek Petrik, and the three anonymous reviewers for their invaluable comments, which greatly improved this manuscript. Zhu's work was partially supported by the National Institutes of Health (NIH) Intramural Research Program under grant number: OT2OD040373. The findings and conclusions presented in this paper are those of the author(s) and do not necessarily reflect the views of the NIH or the U.S. Department of Health and Human Services.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24.
- Abeille, M. and Lazaric, A. (2017). Linear thompson sampling revisited. In *Artificial Intelligence and Statistics*, pages 176–184. PMLR.
- Agrawal, S. and Goyal, N. (2013). Thompson sampling for contextual bandits with linear payoffs. In *International conference on machine learning*, pages 127–135. PMLR.
- Agrawal, S. and Goyal, N. (2017). Near-optimal regret bounds for thompson sampling. *Journal of the ACM (JACM)*, 64(5):1–24.
- Aouali, I., Kveton, B., and Katariya, S. (2023). Mixed-effect thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, volume 206, pages 2087–2115. PMLR.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256.
- Balduzzi, D. and Ghifary, M. (2015). Compatible value gradients for reinforcement learning of continuous deep policies. *arXiv preprint arXiv:1509.03005*.
- Bastani, H., Bayati, M., and Khosravi, K. (2021). Mostly exploration-free algorithms for contextual bandits. *Management Science*, 67(3):1329–1349.
- Bickel, P. J. and Levina, E. (2008). Covariance regularization by thresholding. *Annals of statistics*, 36(6):2577–2604.
- Bouneffouf, D., Rish, I., and Aggarwal, C. (2020). Survey on applications of multi-armed and contextual bandits. In *2020 IEEE congress on evolutionary computation (CEC)*, pages 1–8. IEEE.
- Bubeck, S., Cesa-Bianchi, N., et al. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122.
- Cella, L., Lazaric, A., and Pontil, M. (2020). Meta-learning with stochastic linear bandits. In *37th International Conference on Machine Learning, ICML 2020*, volume 119, pages 1337–1347. Proceedings of Machine Learning Research (PMLR).
- Cella, L. and Pontil, M. (2021). Multi-task and meta-learning with sparse linear bandits. In *Uncertainty in Artificial Intelligence*, pages 1692–1702. PMLR.
- Christakopoulou, K. and Banerjee, A. (2018). Learning to interact with users: A collaborative-bandit approach. In *Proceedings of the 2018 SIAM International Conference on Data Mining*, pages 612–620. SIAM.

- Deshmukh, A. A., Dogan, U., and Scott, C. (2017). Multi-task learning for contextual bandits. *Advances in neural information processing systems*, 30.
- Fang, M. and Tao, D. (2015). Active multi-task learning via bandits. In *Proceedings of the 2015 SIAM International Conference on Data Mining*, pages 505–513. SIAM.
- Gabillon, V., Ghavamzadeh, M., Lazaric, A., and Bubeck, S. (2011). Multi-bandit best arm identification. *Advances in Neural Information Processing Systems*, 24.
- Goldenshluger, A. and Zeevi, A. (2013). A linear response bandit problem. *Stochastic Systems*, 3(1):230–261.
- Guo, D., Ktena, S. I., Myana, P. K., Huszar, F., Shi, W., Tejani, A., Kneier, M., and Das, S. (2020). Deep bayesian bandits: Exploring in online personalized recommendations. In *Proceedings of the 14th ACM conference on recommender systems*, pages 456–461.
- Harper, F. M. and Konstan, J. A. (2015). The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19.
- Hong, J., Kveton, B., Zaheer, M., and Ghavamzadeh, M. (2022). Hierarchical bayesian bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 7724–7741. PMLR.
- Hong, J., Kveton, B., Zaheer, M., Katariya, S., and Ghavamzadeh, M. (2023). Multi-task off-policy learning from bandit feedback. In *International Conference on Machine Learning*, pages 13157–13173. PMLR.
- Huang, X., Xu, K., Lee, D., Hassani, H., Bastani, H., and Dobriban, E. (2025). Optimal multitask linear regression and contextual bandits under sparse heterogeneity. *Journal of the American Statistical Association*, pages 1–14.
- Kveton, B., Konobeev, M., Zaheer, M., Hsu, C.-w., Mladenov, M., Boutilier, C., and Szepesvari, C. (2021). Meta-thompson sampling. In *International Conference on Machine Learning*, pages 5884–5893. PMLR.
- Lee, C.-W., Liu, Q., Abbasi Yadkori, Y., Jin, C., Lattimore, T., and Szepesvári, C. (2023). Context-lumpable stochastic bandits. *Advances in Neural Information Processing Systems*, 36:77915–77941.
- Li, L., Chu, W., Langford, J., and Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670.
- Mate, A., Killian, J., Xu, H., Perrault, A., and Tambe, M. (2020). Collapsing bandits and their application to public health intervention. *Advances in Neural Information Processing Systems*, 33:15639–15650.
- Misra, K., Schwartz, E. M., and Abernethy, J. (2019). Dynamic online pricing with incomplete information using multiarmed bandit experiments. *Marketing Science*, 38(2):226–252.
- Mueller, J. W., Syrgkanis, V., and Taddy, M. (2019). Low-rank bandit methods for high-dimensional dynamic pricing. *Advances in Neural Information Processing Systems*, 32.
- Qin, L., Chen, S., and Zhu, X. (2014). Contextual combinatorial bandit and its application on diversified online recommendation. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, pages 461–469. SIAM.
- Scarlett, J., Bogunovic, I., and Cevher, V. (2019). Overlapping multi-bandit best arm identification. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 2544–2548. IEEE.
- Slate, D. (1991). Letter Recognition. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5ZP40>.
- Soare, M., Alsharif, O., Lazaric, A., and Pineau, J. (2014). Multi-task linear bandits. In *NIPS2014 workshop on transfer and multi-task learning: theory meets practice*.
- Su, Y., Lu, H., Li, Y., Liu, L., Bi, S., Chi, E. H., and Chen, M. (2024). Multi-task neural linear bandit for exploration in recommender systems. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5723–5730.
- Swersky, K., Snoek, J., and Adams, R. P. (2013). Multi-task bayesian optimization. *Advances in neural information processing systems*, 26.
- Wan, R., Ge, L., and Song, R. (2021). Metadata-based multi-task bandits with bayesian hierarchical models. *Advances in Neural Information Processing Systems*, 34:29655–29668.

- Wang, L., Bai, Y., Bhalla, A., and Joachims, T. (2019). Batch learning from bandit feedback through bias corrected reward imputation. In *ICML Workshop on Real-World Sequential Decision Making*.
- Woodroffe, M. (1979). A one-armed bandit problem with a concomitant variable. *Journal of the American Statistical Association*, 74(368):799–806.
- Xu, K. and Bastani, H. (2025). Multitask learning and bandits via robust statistics. *Management Science*, 71(9):7752–7773.
- Yang, L., Liu, B., Lin, L., Xia, F., Chen, K., and Yang, Q. (2020). Exploring clustering of bandits for online recommendation system. In *Proceedings of the 14th ACM Conference on Recommender Systems*, pages 120–129.
- Zhang, Y. and Yang, Q. (2021). A survey on multi-task learning. *IEEE transactions on knowledge and data engineering*, 34(12):5586–5609.
- Zhu, R. and Kveton, B. (2022a). Random effect bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 3091–3107. PMLR.
- Zhu, R. and Kveton, B. (2022b). Robust contextual linear bandits. *arXiv preprint arXiv:2210.14483*.

A Proof of Theorem 4.1

In this section, we provide the detailed proof of Theorem 4.1. We begin by introducing two supporting lemmas. Lemma A.1 decomposes the estimation error $\hat{\beta}_{k,j,t} - \beta_{k,j}$ into two components: the first corresponds to the error induced by the prior estimation $\hat{\beta}_{k0,t}$, and the second arises from the observation noise. Lemma A.2 reformulates the term $\mathbf{X}_{k,j,t}^\top \mathbf{V}_{k,j,t}^{-1} \mathbf{X}_{k,j,t}$ used in the construction of the prior-informed estimator $\hat{\beta}_{k0,t}$.

LEMMA A.1. *The estimation error $\hat{\beta}_{k,j,t} - \beta_{k,j}$ can be explicitly decomposed into two interpretable components: the first captures the discrepancy arising from the prior-informed estimator, while the second reflects the uncertainty due to observational noise, as shown below.*

$$\hat{\beta}_{k,j,t} - \beta_{k,j} = \underbrace{\sigma^2 \tilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1} (\hat{\beta}_{k0,t} - \beta_{k,j})}_{\text{Error from the Prior-Informed Estimator}} + \underbrace{\tilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t} \epsilon_{k,j,t}}_{\text{Error Induced by Observational Noise}}.$$

PROOF. Recalling that $\mathbf{y}_{k,j,t} = \mathbf{X}_{k,j,t} \beta_{k,j} + \epsilon_{k,j,t}$, and substituting this expression into $\hat{\beta}_{k,j,t}$, we obtain

$$\begin{aligned} \hat{\beta}_{k,j,t} &= \sigma^2 \tilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1} \hat{\beta}_{k0,t} + \tilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t}^\top \mathbf{y}_{k,j,t} \\ &= \beta_{k,j} + \sigma^2 \tilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1} (\hat{\beta}_{k0,t} - \beta_{k,j}) + \tilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t}^\top \epsilon_{k,j,t}, \end{aligned}$$

where the last equality is due to $\tilde{\mathbf{C}}_{k,j,t} = (\mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t} + \sigma^2 \Sigma_k^{-1})^{-1}$. $\square$

LEMMA A.2. *We have*

$$\mathbf{X}_{k,j,t}^\top \mathbf{V}_{k,j,t}^{-1} \mathbf{X}_{k,j,t} = \Sigma_k^{-1} - \sigma^2 \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1}.$$

PROOF. From (4), we have

$$\begin{aligned} \mathbf{X}_{k,j,t}^\top \mathbf{V}_{k,j,t}^{-1} \mathbf{X}_{k,j,t} &= \sigma^{-2} \mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t} - \sigma^{-2} \mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t} \tilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t} \\ &= \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t} \\ &= \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,j,t} (\mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t} + \sigma^2 \Sigma_k^{-1} - \sigma^2 \Sigma_k^{-1}) \\ &= \Sigma_k^{-1} - \sigma^2 \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1}, \end{aligned}$$

where the second equality is from $\tilde{\mathbf{C}}_{k,j,t} = (\mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t} + \sigma^2 \Sigma_k^{-1})^{-1}$. $\square$

Proof of Theorem 4.1. By the Cauchy-Schwarz inequality, we have

$$\left| \mathbf{x}^\top \left(\hat{\boldsymbol{\beta}}_{k,j,t} - \boldsymbol{\beta}_{k,j} \right) \right| \leq \sigma^{-1} \|\mathbf{C}_{k,j,t}^{\frac{1}{2}} \mathbf{x}\|_2 \cdot \|\widetilde{\mathbf{C}}_{k,j,t}^{-\frac{1}{2}} \left(\hat{\boldsymbol{\beta}}_{k,j,t} - \boldsymbol{\beta}_{k,j} \right)\|_2. \quad (7)$$

Now we derive the upper bound for $\|\widetilde{\mathbf{C}}_{k,j,t}^{-\frac{1}{2}} \left(\hat{\boldsymbol{\beta}}_{k,j,t} - \boldsymbol{\beta}_{k,j} \right)\|_2$. According to Lemma A.1, the estimation error is decomposed into the following two components

$$\hat{\boldsymbol{\beta}}_{k,j,t} - \boldsymbol{\beta}_{k,j} = \sigma^2 \widetilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1} \left(\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j} \right) + \widetilde{\mathbf{C}}_{k,j,t} \mathbf{X}_{k,j,t} \boldsymbol{\epsilon}_{k,j,t}.$$

By employing the triangle inequality, $\|\widetilde{\mathbf{C}}_{k,j,t}^{-\frac{1}{2}} \left(\hat{\boldsymbol{\beta}}_{k,j,t} - \boldsymbol{\beta}_{k,j} \right)\|_2$ can be upper bounded as follows:

$$\|\widetilde{\mathbf{C}}_{k,j,t}^{-\frac{1}{2}} \left(\hat{\boldsymbol{\beta}}_{k,j,t} - \boldsymbol{\beta}_{k,j} \right)\|_2 \leq \|\sigma^2 \widetilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \Sigma_k^{-1} \left(\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j} \right)\|_2 + \|\widetilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \mathbf{X}_{k,j,t}^\top \boldsymbol{\epsilon}_{k,j,t}\|_2. \quad (8)$$

We first analyze the first term on the right-hand side of (8). Through algebraic manipulation and noting that $\sigma^2 \Sigma_k^{-1} \widetilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1} \preceq \Sigma_k^{-1}$, we have

$$\begin{aligned} \|\sigma^2 \widetilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \Sigma_k^{-1} \left(\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j} \right)\|_2^2 &= \left(\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j} \right)^\top \sigma^4 \Sigma_k^{-1} \widetilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1} \left(\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j} \right) \\ &\leq \left(\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j} \right)^\top \sigma^2 \Sigma_k^{-1} \left(\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j} \right) \\ &\leq \sigma^2 \lambda_1 \|\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j}\|_2^2, \end{aligned} \quad (9)$$

where the last equality is due to Assumption 1. Let $\widetilde{\mathbf{X}}_{k,j,t} = \mathbf{V}_{k,j,t}^{-\frac{1}{2}} \mathbf{X}_{k,j,t}$, $\widetilde{\mathbf{y}}_{k,j,t} = \mathbf{V}_{k,j,t}^{-\frac{1}{2}} \mathbf{y}_{k,j,t}$, $\widetilde{\boldsymbol{\epsilon}}_{k,j,t} = \mathbf{V}_{k,j,t}^{-\frac{1}{2}} \boldsymbol{\epsilon}_{k,j,t}$, and $\mathbf{M} = \sum_{j=1}^N \widetilde{\mathbf{X}}_{k,j,t}^\top \widetilde{\mathbf{X}}_{k,j,t} + \lambda \mathbf{I}$. We obtain

$$\begin{aligned} \hat{\boldsymbol{\beta}}_{k0,t} &= \mathbf{M}^{-1} \sum_{j=1}^N \widetilde{\mathbf{X}}_{k,j,t}^\top (\widetilde{\mathbf{X}}_{k,j,t} \boldsymbol{\beta}_{k,j} + \widetilde{\boldsymbol{\epsilon}}_{k,j,t}) \\ &= \sum_{s=1}^N \mathbf{M}^{-1} \widetilde{\mathbf{X}}_{k,s,t}^\top \widetilde{\mathbf{X}}_{k,s,t} \boldsymbol{\beta}_{k,s} + \mathbf{M}^{-1} \sum_{j=1}^N \widetilde{\mathbf{X}}_{k,j,t}^\top \widetilde{\boldsymbol{\epsilon}}_{k,j,t}, \end{aligned}$$

It follows

$$\|\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j}\|_2 \leq N c_\beta + \frac{1}{\sqrt{\lambda}} \left\| \mathbf{M}^{-\frac{1}{2}} \sum_{j=1}^N \widetilde{\mathbf{X}}_{k,j,t}^\top \widetilde{\boldsymbol{\epsilon}}_{k,j,t} \right\|_2,$$

where the inequality arises from the fact that, for all $s \in [N]$, the eigenvalues of $\widetilde{\mathbf{X}}_{k,s,t}^\top \widetilde{\mathbf{X}}_{k,s,t}$ are smaller than those of $\sum_{j=1}^N \widetilde{\mathbf{X}}_{k,j,t}^\top \widetilde{\mathbf{X}}_{k,j,t}$. By applying Lemma E.1, we obtain the following inequality, which holds with probability at least $1 - \delta/2$:

$$\|\hat{\boldsymbol{\beta}}_{k0,t} - \boldsymbol{\beta}_{k,j}\|_2 \leq N c_\beta + \sqrt{d/\lambda \log \left(\frac{\lambda + t c_x}{\lambda \delta/2} \right)}. \quad (10)$$

We next analyze the second term on the right-hand side of (8) using Lemma E.1. Noting that $\lambda_d \leq \lambda(\Sigma_k^{-1}) \leq \lambda_1$, we obtain the following inequality, which holds with probability at least $1 - \delta/2$:

$$\|\widetilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \mathbf{X}_{k,j,t}^\top \boldsymbol{\epsilon}_{k,j,t}\|_2 \leq \sqrt{\sigma^2 d \log \left(\frac{\sigma^2 \lambda_1 + t c_x}{\sigma^2 \lambda_d \delta/2} \right)}. \quad (11)$$

By combining (9), (10), and (11), and substituting the resulting expressions into (8), the following holds with probability at least $1 - \delta$:

$$\begin{aligned} \|\tilde{\mathbf{C}}_{k,j,t}^{-\frac{1}{2}} (\hat{\boldsymbol{\beta}}_{k,j,t} - \boldsymbol{\beta}_{k,j})\|_2 &\leq \sqrt{\sigma^2 d \lambda_1 / \lambda \log \left(\frac{\lambda + t c_x}{\lambda \delta / 2} \right)} + \sqrt{\sigma^2 d \log \left(\frac{\sigma^2 \lambda_1 + t c_x}{\sigma^2 \lambda_d \delta / 2} \right)} + \sigma N c_\beta \sqrt{\lambda_1} \\ &\leq 2 \sqrt{\sigma^2 d \max\{\lambda_1 / \lambda, 1\} \log \left(\frac{\max\{\lambda, \sigma^2 \lambda_1\} + t c_x}{\sqrt{\lambda \lambda_d} \sigma^2 \delta / 2} \right)} + \sigma N c_\beta \sqrt{\lambda_1}. \end{aligned}$$

Substituting it into (7) completes the proof.

B Proof of Theorem 6.1

Proof. The U-value for arm k of bandit j at time step t is defined as

$$U_{k,j,t} = \mathbf{x}_t^\top \hat{\boldsymbol{\beta}}_{k,j,t} + \alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{k,j,t}}.$$

Note that at each time step, the EBMUCB algorithm selects the arm with the highest U-value. We define the event ξ_t as

$$\xi_t = \left\{ \forall k \in [K], \forall j \in [N] : \left| \mathbf{x}_t^\top (\hat{\boldsymbol{\beta}}_{k,j,t} - \boldsymbol{\beta}_{k,j}) \right| \leq \alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{k,j,t}} \right\}.$$

The regret R_n can be decomposed as

$$\begin{aligned} R_n &= \sum_{t=1}^n \mathbb{E} \left[\mathbb{I}(\xi_t) \left(\mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t, Z_t} \right) \right] + \sum_{t=1}^n \mathbb{E} \left[\mathbb{I}(\bar{\xi}_t) \left(\mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t, Z_t} \right) \right] \\ &\leq \sum_{t=1}^n \mathbb{E} \left[\left(\mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t, Z_t} \right) \mid \xi_t \right] + 2\sqrt{d c_x c_\beta} \sum_{t=1}^n \mathbb{P}(\bar{\xi}_t). \end{aligned} \quad (12)$$

In the right-hand side of (12), the upper bound of R_n consists of two terms. The first term represents the gap between the optimal arm and the chosen arm, conditioned on the event ξ_t . The second term corresponds to the probability that ξ_t does not hold. We first focus on the first term in the decomposition of (12), with each subterm in this summation given as follows:

$$\begin{aligned} \mathbb{E} \left[\left(\mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t, Z_t} \right) \mid \xi_t \right] &\leq \mathbb{E} \left[\left(\mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t} - U_{\pi_t^*, Z_t, t} + U_{\pi_t, Z_t, t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t, Z_t} \right) \mid \xi_t \right] \\ &\leq 2 \mathbb{E} \left[\alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{\pi_t, Z_t, t}} \mid \xi_t \right] \\ &\leq 2\alpha_n(\delta) \mathbb{E} \left[\sqrt{\mathbf{x}_t^\top \mathbf{C}_{\pi_t, Z_t, t} \mathbf{x}_t} \mid \xi_t \right], \end{aligned}$$

where the first inequality follows from the fact that $U_{\pi_t, Z_t, t} \geq U_{\pi_t^*, Z_t, t}$ when selecting arm π_t , and the second inequality holds because, for any arm k conditioned on ξ_t , $0 \leq U_{k,j,t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{k,j} \leq 2\alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{k,j,t}}$. As a result, we sum over all possible values of t ,

$$\begin{aligned} \sum_{t=1}^n \mathbb{E} \left[\left(\mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t, Z_t} \right) \mid \xi_t \right] &\leq 2\alpha_n(\delta) \sum_{t=1}^n \mathbb{E} \left[\sqrt{\mathbf{x}_t^\top \mathbf{C}_{\pi_t, Z_t, t} \mathbf{x}_t} \mid \xi_t \right] \\ &\leq 2\alpha_n(\delta) \sqrt{n \mathbb{E} \left[\sum_{t=1}^n \mathbf{x}_t^\top \mathbf{C}_{\pi_t, Z_t, t} \mathbf{x}_t \mid \xi_t \right]}, \end{aligned}$$

where the last inequality follows from the Cauchy–Schwarz inequality and the concavity of the square root, which implies that, for non-negative random variables X_i , $\sum_{i=1}^n \mathbb{E}(\sqrt{X_i}) \leq \sqrt{n \mathbb{E}(\sum_{i=1}^n X_i)}$.

From (3), $\mathbf{x}_t^\top \mathbf{C}_{k,j,t} \mathbf{x}_t$ can be decomposed into two terms as follows

$$\mathbf{x}_t^\top \mathbf{C}_{k,j,t} \mathbf{x}_t = \sigma^2 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t} \mathbf{x}_t + \sigma^4 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t} \Sigma_k^{-1} \Phi_{k0,t} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,j,t} \mathbf{x}_t. \quad (13)$$

The detailed derivation of the upper bound for the summation of the first term in (13) is provided in Lemma D.1, which yields

$$\sigma^2 \sum_{t=1}^n \mathbf{x}_t^\top \tilde{\mathbf{C}}_{\pi_t, Z_t, t} \mathbf{x}_t \leq c_1 d K \sum_{j=1}^N \log(1 + c_2 n_j). \quad (14)$$

For the second term in (13), Lemma D.2 implies

$$\sum_{t=1}^n \sigma^4 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{\pi_t, Z_t, t} \Sigma_{\pi_t}^{-1} \Phi_{\pi_t 0, t} \Sigma_{\pi_t}^{-1} \tilde{\mathbf{C}}_{\pi_t, Z_t, t} \mathbf{x}_t \leq c_3 d K \log(1 + c_4 N). \quad (15)$$

Next, we bound the probability that ξ_t does not hold,

$$\mathbb{P}(\bar{\xi}_t) \leq KN\delta. \quad (16)$$

Combining (14), (15) and (16), we can derive the upper bound of (12):

$$R_n \leq 2\alpha_n(\delta) \sqrt{ndK \left[c_1 \sum_{j=1}^N \log(1 + c_2 n_j) + c_3 \log(1 + c_4 N) \right]} + 2\sqrt{dc_x c_\beta} KNn\delta.$$

C Proof of Theorem 6.2

Proof. We define two following events $\xi_t^{(1)}$ and $\xi_t^{(2)}$ as:

$$\begin{aligned} \xi_t^{(1)} &= \left\{ \forall k \in [K], \forall j \in [N] : \left| \mathbf{x}_t^\top (\hat{\boldsymbol{\beta}}_{k,j,t} - \boldsymbol{\beta}_{k,j}) \right| \leq \alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{k,j,t}} \right\} \\ \xi_t^{(2)} &= \left\{ \forall k \in [K], \forall j \in [N] : \left| \mathbf{x}_t^\top (\check{\boldsymbol{\beta}}_{k,j,t} - \hat{\boldsymbol{\beta}}_{k,j,t}) \right| \leq \alpha_t(\delta) \gamma_t \|\mathbf{x}_t\|_{\mathbf{C}_{k,j,t}} \right\}, \end{aligned}$$

where $\gamma_t = \sqrt{2d \log dt^2}$. The event $\xi_t^{(2)}$ is introduced because, as shown in Lemma C.1, the error between the sampled variable $\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{k,j,t}$ and its mean $\mathbf{x}_t^\top \hat{\boldsymbol{\beta}}_{k,j,t}$ is bounded by $\alpha_t(\delta) \gamma_t \|\mathbf{x}_t\|_{\mathbf{C}_{k,j,t}}$ with probability of at least $1 - 1/t^2$ for all $k \in [K]$ and $j \in [N]$. The cumulative regret can be decomposed into

$$R_n \leq \sum_{t=1}^n \mathbb{E} \left[\left(\mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t, Z_t} \right) \mid \xi_t^{(1)}, \xi_t^{(2)} \right] + 2\sqrt{dc_x c_\beta} \sum_{t=1}^n \left(\mathbb{P}(\bar{\xi}_t^{(1)}) + \mathbb{P}(\bar{\xi}_t^{(2)}) \right).$$

Denote the set of saturated arms (Agrawal and Goyal, 2013) as,

$$\mathcal{S}(t) = \{k \in [K] : \Delta_{k,t} > (\gamma_t + 1) \alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{k, Z_t, t}}\},$$

where $\Delta_{k,t} = \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{k, Z_t}$. Let $\pi_t^\dagger$ represent the unsaturated arm with the smallest $\|\mathbf{x}_t\|_{\mathbf{C}_{k, Z_t, t}}$ norm, i.e.,

$$\pi_t^\dagger = \arg \min_{k \notin \mathcal{S}(t)} \|\mathbf{x}_t\|_{\mathbf{C}_{k, Z_t, t}}.$$

The existence of an unsaturated arm $\pi_t^\dagger$ is guaranteed since $\pi_t^* \notin \mathcal{S}(t)$. When both $\xi_t^{(1)}$ and $\xi_t^{(2)}$ hold, the suboptimality can be expressed as follows:

$$\begin{aligned} \Delta_{\pi_t, t} &= \Delta_{\pi_t^\dagger, t} + \mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^\dagger, Z_t} - \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t, Z_t} \\ &\leq \Delta_{\pi_t^\dagger, t} + \left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^\dagger, Z_t, t} + (\gamma_t + 1)\alpha_t(\delta) \|\mathbf{x}_t\|_{C_{\pi_t^\dagger, Z_t, t}} \right) - \left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t, Z_t, t} - (\gamma_t + 1)\alpha_t(\delta) \|\mathbf{x}_t\|_{C_{\pi_t, Z_t, t}} \right), \end{aligned}$$

Since we select arm π_t at time step t , i.e., $\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t, Z_t, t} > \mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^\dagger, Z_t, t}$, it follows

$$\begin{aligned} \Delta_{\pi_t, t} &\leq \Delta_{\pi_t^\dagger, t} + (\gamma_t + 1)\alpha_t(\delta) \|\mathbf{x}_t\|_{C_{\pi_t^\dagger, Z_t, t}} + (\gamma_t + 1)\alpha_t(\delta) \|\mathbf{x}_t\|_{C_{\pi_t, Z_t, t}} \\ &\leq (\gamma_t + 1)\alpha_t(\delta) \left(2\|\mathbf{x}_t\|_{C_{\pi_t^\dagger, Z_t, t}} + \|\mathbf{x}_t\|_{C_{\pi_t, Z_t, t}} \right), \end{aligned}$$

where the second inequality results from $\pi_t^\dagger \notin \mathcal{S}(t)$. Note that $\|\mathbf{x}_t\|_{C_{\pi_t, Z_t, t}} \geq \|\mathbf{x}_t\|_{C_{\pi_t^\dagger, Z_t, t}}$ with constant probability, i.e.,

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{x}_t\|_{C_{\pi_t, Z_t, t}} \mid \xi_t^{(1)}, \xi_t^{(2)} \right] &\geq \mathbb{E} \left[\|\mathbf{x}_t\|_{C_{\pi_t, Z_t, t}} \mid \pi_t \notin \mathcal{S}(t), \xi_t^{(1)}, \xi_t^{(2)} \right] \cdot \mathbb{P} \left(\pi_t \notin \mathcal{S}(t) \mid \xi_t^{(1)}, \xi_t^{(2)} \right) \\ &\geq \left(\frac{1}{2\sqrt{2\pi e}} - \frac{1}{t^2} \right) \mathbb{E} \left[\|\mathbf{x}_t\|_{C_{\pi_t^\dagger, Z_t, t}} \mid \xi_t^{(1)}, \xi_t^{(2)} \right], \end{aligned}$$

The detailed proof for the lower bound of $\mathbb{P} \left(\pi_t \notin \mathcal{S}(t) \mid \xi_t^{(1)}, \xi_t^{(2)} \right)$ is provided in Lemma C.3, and the last inequality follows from the definition of $\pi_t^\dagger$ as the unsaturated arm with the smallest $\|\mathbf{x}_t\|_{C_{k, Z_t, t}}$. As a result,

$$\begin{aligned} \mathbb{E} \left[\Delta_{\pi_t, t} \mid \xi_t^{(1)}, \xi_t^{(2)} \right] &\leq (\gamma_t + 1)\alpha_t(\delta) \mathbb{E} \left[2\|\mathbf{x}_t\|_{C_{\pi_t^\dagger, Z_t, t}} + \|\mathbf{x}_t\|_{C_{\pi_t, Z_t, t}} \mid \xi_t^{(1)}, \xi_t^{(2)} \right] + 2\sqrt{dc_x c_\beta} NK \left(\delta + \frac{1}{t^2} \right) \\ &\leq \left(\frac{2}{\frac{1}{2\sqrt{2\pi e}} - \frac{1}{t^2}} + 1 \right) (\gamma_t + 1)\alpha_t(\delta) \mathbb{E} \left[\|\mathbf{x}_t\|_{C_{\pi_t^\dagger, Z_t, t}} \mid \xi_t^{(1)}, \xi_t^{(2)} \right] + 2\sqrt{dc_x c_\beta} NK \left(\delta + \frac{1}{t^2} \right). \end{aligned}$$

Let $C = \max_{t \geq 1} \left[\frac{2}{\frac{1}{2\sqrt{2\pi e}} - \frac{1}{t^2}} + 1 \right] + 1$, then sum all t , we have

$$R_n \leq C(\gamma_n + 1)\alpha_n(\delta) \sqrt{n \mathbb{E} \left[\sum_{t=1}^n \|\mathbf{x}_t\|_{C_{\pi_t, Z_t, t}}^2 \mid \xi_t^{(1)}, \xi_t^{(2)} \right]} + 2\sqrt{dc_x c_\beta} KN(n\delta + \pi^2/6).$$

Combining the upper bound in Lemma D.1 and Lemma D.2, we obtain the final result.

LEMMA C.1. *The sampled variable $\mathbf{x}^\top \check{\boldsymbol{\beta}}_{k, j, t}$ in the algorithm EBMTS satisfies the following bound on the sampling error:*

$$\left| \mathbf{x}^\top \left(\check{\boldsymbol{\beta}}_{k, j, t} - \hat{\boldsymbol{\beta}}_{k, j, t} \right) \right| \leq \alpha_t(\delta) \sqrt{2d \log d / \delta'} \|\mathbf{x}\|_{C_{k, j, t}},$$

with probability at least $1 - \delta'$.

PROOF. Note that $\check{\boldsymbol{\beta}}_{k, j, t} \sim \mathcal{N} \left(\hat{\boldsymbol{\beta}}_{k, j, t}, \alpha_t^2(\delta) C_{k, j, t} \right)$, which implies $C_{k, j, t}^{-\frac{1}{2}} \left(\check{\boldsymbol{\beta}}_{k, j, t} - \hat{\boldsymbol{\beta}}_{k, j, t} \right) / \alpha_t(\delta) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. Applying the anti-concentration inequality from Lemma E.5 for the standard normal distribution, we have that, for

every $t \in [n]$, the following holds with probability at least $1 - \delta'$:

$$\begin{aligned} \left| \mathbf{x}^\top \left(\check{\boldsymbol{\beta}}_{k,j,t} - \hat{\boldsymbol{\beta}}_{k,j,t} \right) \right| &= \left| \mathbf{x}^\top \mathbf{C}_{k,j,t}^{-\frac{1}{2}} \mathbf{C}_{k,j,t}^{-\frac{1}{2}} \left(\check{\boldsymbol{\beta}}_{k,j,t} - \hat{\boldsymbol{\beta}}_{k,j,t} \right) \right| \\ &\leq \left\| \mathbf{C}_{k,j,t}^{-\frac{1}{2}} \left(\check{\boldsymbol{\beta}}_{k,j,t} - \hat{\boldsymbol{\beta}}_{k,j,t} \right) \right\|_2 \|\mathbf{x}\|_{\mathbf{C}_{k,j,t}} \\ &\leq \alpha_t(\delta) \sqrt{2d \log d / \delta'} \|\mathbf{x}\|_{\mathbf{C}_{k,j,t}}, \end{aligned}$$

where the first inequality follows from the Cauchy-Schwarz inequality. We now provide a detailed explanation of the second inequality. Let $\mathbf{C}_{k,j,t}^{-\frac{1}{2}} \left(\check{\boldsymbol{\beta}}_{k,j,t} - \hat{\boldsymbol{\beta}}_{k,j,t} \right) / \alpha_t(\delta) =: (u_1, u_2, \dots, u_d)$, where each u_i is a standardized random variable with unit norm, and the u_i are independent of one another. We have

$$\begin{aligned} \mathbb{P} \left(\left\| \frac{\mathbf{C}_{k,j,t}^{-\frac{1}{2}} \left(\check{\boldsymbol{\beta}}_{k,j,t} - \hat{\boldsymbol{\beta}}_{k,j,t} \right)}{\alpha_t(\delta)} \right\|_2 \geq \sqrt{2d \log d / \delta'} \right) &= \mathbb{P} \left(\sqrt{u_1^2 + \dots + u_d^2} \geq \sqrt{2d \log d / \delta'} \right) \\ &\leq \mathbb{P} \left(\exists i \in [d], u_i \geq \sqrt{2 \log d / \delta'} \right) \\ &\leq d \mathbb{P} \left(u_1 \geq \sqrt{2 \log d / \delta'} \right) \leq \delta'. \end{aligned}$$

□

LEMMA C.2. For any $t \geq 1$, conditional on $\xi_t^{(1)}$ holding, there exists a constant probability that the sampled variable $\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t}$ of the optimal arm serves as an upper confidence bound, i.e.,

$$\mathbb{P} \left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t} > \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t, t} \mid \xi_t^{(1)} \right) \geq \frac{1}{2\sqrt{2\pi e}}.$$

PROOF. In the EBM algorithm, $\check{\boldsymbol{\beta}}_{k,j,t}$ is sampled from a normal distribution with mean $\hat{\boldsymbol{\beta}}_{k,j,t}$. Using this, the probability that the sampled variable $\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t}$ serves as an upper bound for the true reward is lower-bounded by

$$\begin{aligned} \mathbb{P} \left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t} > \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t, t} \mid \xi_t^{(1)} \right) &= \mathbb{P} \left(\frac{\mathbf{x}_t^\top (\check{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t} - \hat{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t})}{\alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{\pi_t^*, Z_t, t}}} > \frac{\mathbf{x}_t^\top (\boldsymbol{\beta}_{\pi_t^*, Z_t, t} - \hat{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t})}{\alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{\pi_t^*, Z_t, t}}} \mid \xi_t^{(1)} \right) \\ &\geq \mathbb{P} \left(\frac{\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t} - \mathbf{x}_t^\top \hat{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t}}{\alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{\pi_t^*, Z_t, t}}} > 1 \mid \xi_t^{(1)} \right) \geq \frac{1}{2\sqrt{2\pi e}}, \end{aligned}$$

Here, the first inequality follows from the fact that, under $\xi_t^{(1)}$,

$$\frac{\mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*, Z_t, t} - \mathbf{x}_t^\top \hat{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t}}{\alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{\pi_t^*, Z_t, t}}} \leq \frac{\alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{\pi_t^*, Z_t, t}}}{\alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{\pi_t^*, Z_t, t}}} = 1.$$

The second inequality follows from the anti-concentration inequality for the standard normal distribution (see Lemma E.5) and from the fact that

$$\frac{\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t} - \mathbf{x}_t^\top \hat{\boldsymbol{\beta}}_{\pi_t^*, Z_t, t}}{\alpha_t(\delta) \|\mathbf{x}_t\|_{\mathbf{C}_{\pi_t^*, Z_t, t}}} \sim \mathcal{N}(0, 1).$$

□

LEMMA C.3. For any $t \geq 1$, conditional on $\xi_t^{(1)}$ and $\xi_t^{(2)}$ holding, there exists a constant probability that the chosen arm π_t is not a saturated arm, i.e.,

$$\mathbb{P}\left(\pi_t \notin \mathcal{S}(t) \mid \xi_t^{(1)}, \xi_t^{(2)}\right) \geq \frac{1}{2\sqrt{2\pi e}} - \frac{1}{t^2}.$$

PROOF. Recall the definition of the set of saturated arms:

$$\mathcal{S}(t) = \{k \in [K] : \Delta_{k,t} > (\gamma_t + 1)\alpha_t(\delta)\|\mathbf{x}_t\|_{\mathbf{C}_{k,Z_t,t}}\},$$

and note that the selected arm $\pi_t = \arg \max_{k \in [K]} \mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{k,Z_t,t}$ is chosen to maximize the sampled estimated reward. Consequently, $\pi_t \notin \mathcal{S}(t)$ if $\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} > \mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{k,Z_t,t}$ for all saturated arms $k \in \mathcal{S}(t)$, meaning that the sampled estimated reward for the optimal arm exceeds those of all saturated arms. This statement can be easily verified via the contrapositive: if $\pi_t \in \mathcal{S}(t)$, then there exists at least one arm in $\mathcal{S}(t)$ — specifically π_t itself — such that

$$\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} \leq \mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t,Z_t,t}.$$

It follows

$$\mathbb{P}\left(\pi_t \notin \mathcal{S}(t) \mid \xi_t^{(1)}, \xi_t^{(2)}\right) \geq \mathbb{P}\left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} > \mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{k,Z_t,t}, \forall k \in \mathcal{S}(t) \mid \xi_t^{(1)}, \xi_t^{(2)}\right). \quad (17)$$

When both events $\xi_t^{(1)}$ and $\xi_t^{(2)}$ hold, for $k \in \mathcal{S}(t)$, we have

$$\begin{aligned} \mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{k,Z_t,t} &\leq \mathbf{x}_t^\top \boldsymbol{\beta}_{k,Z_t,t} + (\gamma_t + 1)\alpha_t(\delta)\|\mathbf{x}_t\|_{\mathbf{C}_{k,Z_t,t}} \\ &< \mathbf{x}_t^\top \boldsymbol{\beta}_{k,Z_t,t} + \Delta_{k,t} = \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*,Z_t,t}. \end{aligned}$$

It follows

$$\begin{aligned} \mathbb{P}\left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} > \mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{k,Z_t,t}, \forall k \in \mathcal{S}(t) \mid \xi_t^{(1)}, \xi_t^{(2)}\right) &\geq \mathbb{P}\left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} > \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*,Z_t,t} \mid \xi_t^{(1)}, \xi_t^{(2)}\right) \\ &\geq \mathbb{P}\left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} > \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*,Z_t,t} \mid \xi_t^{(1)}\right) - \mathbb{P}\left(\overline{\xi_t^{(2)}}\right), \end{aligned} \quad (18)$$

where the last inequality follows from

$$\begin{aligned} \mathbb{P}\left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} > \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*,Z_t,t} \mid \xi_t^{(1)}\right) &= \mathbb{P}\left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} > \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*,Z_t,t} \mid \xi_t^{(1)}, \xi_t^{(2)}\right) \mathbb{P}\left(\xi_t^{(2)}\right) \\ &\quad + \mathbb{P}\left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} > \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*,Z_t,t} \mid \overline{\xi_t^{(1)}}, \xi_t^{(2)}\right) \mathbb{P}\left(\overline{\xi_t^{(2)}}\right) \\ &\leq \mathbb{P}\left(\mathbf{x}_t^\top \check{\boldsymbol{\beta}}_{\pi_t^*,Z_t,t} > \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_t^*,Z_t,t} \mid \xi_t^{(1)}, \xi_t^{(2)}\right) + \mathbb{P}\left(\overline{\xi_t^{(2)}}\right). \end{aligned}$$

By combining (17) and (18) and applying Lemma C.2, we have the conclusion

$$\mathbb{P}\left(\pi_t \notin \mathcal{S}(t) \mid \xi_t^{(1)}, \xi_t^{(2)}\right) \geq \frac{1}{2\sqrt{2\pi e}} - \frac{1}{t^2}.$$

□

D Useful Lemmas

LEMMA D.1. For fixed j and k , we have the following bound

$$\sigma^2 \sum_{t=1:\pi_t=k, Z_t=j} \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t} \mathbf{x}_t \leq c_1 d \log(1 + c_2 n_j),$$

where $c_1 = \frac{\lambda_d^{-1} d c_x}{\log(1 + \sigma^{-2} \lambda_d^{-1} d c_x)}$, $c_2 = \sigma^{-2} \lambda_d^{-1} c_x$ and n_j is the total time steps of bandit j .

PROOF. For every individual term in the summation, we have

$$\begin{aligned} \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t} \mathbf{x}_t &\leq \tilde{c}_1 \log(1 + \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t} \mathbf{x}_t) \\ &= \tilde{c}_1 \log \det \left(\mathbf{I}_d + \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \mathbf{x}_t \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \right), \end{aligned} \quad (19)$$

where the inequality is from the fact: for $x \in [0, u]$,

$$x = \frac{x}{\log(1+x)} \log(1+x) \leq \frac{u}{\log(1+u)} \log(1+x).$$

The constant $\tilde{c}_1$ in (19) is derived from eigenvalue bounds of covariance matrix,

$$\mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t} \mathbf{x}_t \leq \lambda_{\max}(\tilde{\mathbf{C}}_{k,j,t}) d c_x \leq \lambda_{\min}^{-1}(\sigma^2 \Sigma_k^{-1}) d c_x \leq \sigma^{-2} \lambda_d^{-1} d c_x,$$

following that $\tilde{c}_1$ is

$$\tilde{c}_1 = \frac{\sigma^{-2} \lambda_d^{-1} d c_x}{\log(1 + \sigma^{-2} \lambda_d^{-1} d c_x)}.$$

For the log-determinant term in (19), the matrix determinant property follows

$$\begin{aligned} \log \det \left(\mathbf{I}_d + \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \mathbf{x}_t \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \right) &= \log \det \left(\tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \left(\tilde{\mathbf{C}}_{k,j,t}^{-1} + \mathbf{x}_t \mathbf{x}_t^\top \right) \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \right) \\ &= \log \det \left(\tilde{\mathbf{C}}_{k,j,t}^{-1} + \mathbf{x}_t \mathbf{x}_t^\top \right) - \log \det \left(\tilde{\mathbf{C}}_{k,j,t}^{-1} \right). \end{aligned}$$

Then we have

$$\begin{aligned} \sum_{t=1:\pi_t=k, Z_t=j} \log \det \left(\mathbf{I}_d + \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \mathbf{x}_t \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \right) &= \log \det \left(\tilde{\mathbf{C}}_{k,j,n+1}^{-1} \right) - \log \det \left(\tilde{\mathbf{C}}_{k,j,1}^{-1} \right) \\ &= \log \det \left(\sigma^{-2} \Sigma_k^{\frac{1}{2}} \tilde{\mathbf{C}}_{k,j,n+1}^{-1} \Sigma_k^{\frac{1}{2}} \right) \\ &\leq d \log \frac{1}{\sigma^2 d} \text{trace}(\Sigma_k^{\frac{1}{2}} \tilde{\mathbf{C}}_{k,j,n+1}^{-1} \Sigma_k^{\frac{1}{2}}), \end{aligned}$$

where the first inequality is from the trace-determinant inequality. Noting $\tilde{\mathbf{C}}_{k,j,t}^{-1} = \mathbf{X}_{k,j,t}^\top \mathbf{X}_{k,j,t} + \sigma^2 \Sigma_k^{-1}$, we have

$$\begin{aligned} \text{trace}(\Sigma_k^{\frac{1}{2}} \tilde{\mathbf{C}}_{k,j,n+1}^{-1} \Sigma_k^{\frac{1}{2}}) &= \text{trace}(\Sigma_k^{\frac{1}{2}} \left(\mathbf{X}_{k,j,n+1}^\top \mathbf{X}_{k,j,n+1} + \sigma^2 \Sigma_k^{-1} \right) \Sigma_k^{\frac{1}{2}}) \\ &= \sigma^2 d + \sum_{t=1:\pi_t=k, Z_t=j}^n \mathbf{x}_t^\top \Sigma_k \mathbf{x}_t \\ &\leq \sigma^2 d + \lambda_d^{-1} d c_x n_j. \end{aligned}$$

Thus, we have the bound

$$\sum_{t=1:\pi_t=k, Z_t=j} \log \det \left(\mathbf{I}_d + \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \mathbf{x}_t \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t}^{\frac{1}{2}} \right) \leq d \log(1 + c_2 n_j).$$

where $c_2 = \sigma^{-2}\lambda_d^{-1}c_x$. Letting $c_1 = \sigma^2\tilde{c}_1$ completes the proof. $\square$

LEMMA D.2. For fixed k , we have following bound:

$$\sum_{t=1:\pi_t=k} \sigma^4 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \Phi_{k0,t} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t \leq c_3 d \log(1 + c_4 N),$$

where $c_3 = c_5 c_6$, $c_5 = \frac{\lambda_d^{-2} \lambda_1^2 \lambda^{-1} d c_x}{\log(1 + \sigma^{-2} \lambda_d^{-2} \lambda_1^2 \lambda^{-1} d c_x)}$, $c_6 = 1 + \sigma^{-2} \lambda_d^{-1} d c_x$, and $c_4 = \lambda_1 \lambda^{-1}$.

PROOF. The proof of this lemma follows the same structure as that of Lemma D.1, differing only in a few technical details. First, each term in the summation can be bounded by

$$\sigma^2 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \Phi_{k0,t} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t \leq \tilde{c}_5 \log \left(1 + \sigma^2 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \Phi_{k0,t} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t / c_6 \right),$$

where the constant $\tilde{c}_5 = \frac{\sigma^{-2} \lambda_d^{-2} \lambda_1^2 \lambda^{-1} d c_x}{\log(1 + \sigma^{-2} \lambda_d^{-2} \lambda_1^2 \lambda^{-1} d c_x)}$ is derived as follows:

$$\begin{aligned} \sigma^2 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \Phi_{k0,t} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t &\leq \sigma^2 \lambda_{\max}^2(\tilde{\mathbf{C}}_{k,Z_t,t}) \lambda_{\max}^2(\Sigma_k^{-1}) \lambda_{\max}(\Phi_{k0,t}) d c_x \\ &\leq \sigma^{-2} \lambda_d^{-2} \lambda_1^2 \lambda^{-1} d c_x. \end{aligned}$$

Applying the inequality that $\log(1 + x) \leq c \log(1 + x/c)$ for any $x \geq 0$ and constant $c \geq 1$, and letting $c_6 = 1 + \sigma^{-2} \lambda_d^{-1} d c_x > 1$, we obtain

$$\begin{aligned} \log \left(1 + \sigma^2 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \Phi_{k0,t} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t \right) &\leq c_6 \log \left(1 + \sigma^2 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \Phi_{k0,t} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t / c_6 \right) \\ &\leq c_6 \log \det \left(\mathbf{I}_d + \sigma^2 \Phi_{k0,t}^{\frac{1}{2}} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \Phi_{k0,t}^{\frac{1}{2}} / c_6 \right) \\ &= c_6 \left[\log \det \left(\Phi_{k0,t}^{-1} + \sigma^2 \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} / c_6 \right) - \log \det \left(\Phi_{k0,t}^{-1} \right) \right], \end{aligned}$$

where the second inequality is from the matrix determinant property. Using the result of Lemma D.3, we have

$$\log \left(1 + \sigma^2 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \Phi_{k0,t} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t \right) \leq c_6 \left[\log \det \left(\Phi_{k0,s}^{-1} \right) - \log \det \left(\Phi_{k0,t}^{-1} \right) \right],$$

where s is the next time step choosing arm k . Summing this term over all time steps t in which arm k is chosen, we have

$$\begin{aligned} \sum_{t=1:\pi_t=k} \log \left(1 + \sigma^2 \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \Phi_{k0,t} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t \right) &\leq c_6 \left[\log \det \left(\Phi_{k0,n+1}^{-1} \right) - \log \det \left(\lambda \mathbf{I}_d \right) \right] \\ &\leq c_6 d \log \frac{1}{d} \text{trace} \left(\lambda^{-1} \Phi_{k0,n+1}^{-1} \right) \\ &\leq c_6 d \log(1 + \lambda_1 \lambda^{-1} N), \end{aligned}$$

where the second inequality follows from the trace-determinant inequality, and the third inequality follows from the conclusion of Lemma A.2, specifically,

$$\begin{aligned} \text{trace} \left(\Phi_{k0,n+1}^{-1} \right) &= \sum_{j=1}^N \text{trace} \left(\mathbf{X}_{k,j,n+1}^\top \mathbf{V}_{k,j,n+1}^{-1} \mathbf{X}_{k,j,n+1} \right) + \lambda d \\ &\leq \sum_{j=1}^N \text{trace} \left(\Sigma_k^{-1} \right) + \lambda d \leq \lambda_1 d N + \lambda d. \end{aligned}$$

Let $c_5 = \sigma^2 \tilde{c}_5$ and $c_4 = \lambda_1 \lambda^{-1}$. This completes the proof. $\square$

LEMMA D.3. Let s and t denote two consecutive time steps at which the same arm k is selected, with $s > t$. Then, we have

$$\sigma^2 \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} / c \leq \Phi_{k0,s}^{-1} - \Phi_{k0,t}^{-1},$$

where $c = 1 + \sigma^{-2} \lambda_d^{-1} d c_x$.

PROOF. Since the matrix $\Phi_{k0,t}^{-1}$ is updated only when arm k is selected, most terms in $\Phi_{k0,s}^{-1}$ and $\Phi_{k0,t}^{-1}$ remain identical. This leads to substantial cancellation when taking their difference. Specifically, at time step t , we handle with Z_t -th bandit, and if arm k is chosen, only the term $\mathbf{X}_{k,Z_t,t}^\top \mathbf{V}_{k,Z_t,t}^{-1} \mathbf{X}_{k,Z_t,t}$ in $\Phi_{k0,t}^{-1}$ is updated. Consequently, their difference simplifies to the following expression:

$$\begin{aligned} \Phi_{k0,s}^{-1} - \Phi_{k0,t}^{-1} &= \mathbf{X}_{k,Z_t,s}^\top \mathbf{V}_{k,Z_t,s}^{-1} \mathbf{X}_{k,Z_t,s} - \mathbf{X}_{k,Z_t,t}^\top \mathbf{V}_{k,Z_t,t}^{-1} \mathbf{X}_{k,Z_t,t} \\ &= \sigma^2 \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} - \sigma^2 \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1} \\ &= \sigma^2 \Sigma_k^{-1} \left[\tilde{\mathbf{C}}_{k,Z_t,t} - \left(\tilde{\mathbf{C}}_{k,Z_t,t}^{-1} + \mathbf{x}_t \mathbf{x}_t^\top \right)^{-1} \right] \Sigma_k^{-1} \\ &= \sigma^2 \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t}^{\frac{1}{2}} \left[\mathbf{I}_d - \left(\mathbf{I}_d + \tilde{\mathbf{C}}_{k,Z_t,t}^{\frac{1}{2}} \mathbf{x}_t \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t}^{\frac{1}{2}} \right)^{-1} \right] \tilde{\mathbf{C}}_{k,Z_t,t}^{\frac{1}{2}} \Sigma_k^{-1}. \end{aligned}$$

where the second equality is based on the result of Lemma A.2. Letting $\mathbf{v} = \tilde{\mathbf{C}}_{k,Z_t,t}^{\frac{1}{2}} \mathbf{x}_t$, we have

$$\Phi_{k0,s}^{-1} - \Phi_{k0,t}^{-1} = \sigma^2 \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t}^{\frac{1}{2}} \left[\mathbf{I}_d - (\mathbf{I}_d + \mathbf{v} \mathbf{v}^\top)^{-1} \right] \tilde{\mathbf{C}}_{k,Z_t,t}^{\frac{1}{2}} \Sigma_k^{-1}.$$

Note that $(\mathbf{I}_d + \mathbf{v} \mathbf{v}^\top)^{-1} = \mathbf{I}_d - \mathbf{v} (1 + \mathbf{v}^\top \mathbf{v})^{-1} \mathbf{v}^\top$, and that $1 + \mathbf{v}^\top \mathbf{v}$ admits the following upper bound:

$$1 + \mathbf{v}^\top \mathbf{v} = 1 + \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,j,t} \mathbf{x}_t \leq 1 + \sigma^{-2} \lambda_d^{-1} d c_x.$$

We then derive the following lower bound for $\Phi_{k0,s}^{-1} - \Phi_{k0,t}^{-1}$, as follows

$$\Phi_{k0,s}^{-1} - \Phi_{k0,t}^{-1} \geq \frac{\sigma^2}{1 + \sigma^{-2} \lambda_d^{-1} d c_x} \Sigma_k^{-1} \tilde{\mathbf{C}}_{k,Z_t,t} \mathbf{x}_t \mathbf{x}_t^\top \tilde{\mathbf{C}}_{k,Z_t,t} \Sigma_k^{-1}.$$

Let $c = 1 + \sigma^{-2} \lambda_d^{-1} d c_x$. This completes the proof. $\square$

E Auxiliary Lemmas

LEMMA E.1. (Abbasi-Yadkori et al., 2011) Let $\{\mathcal{F}_t\}_{t=0}^\infty$ be a filtration. Let $\{\eta_t\}_{t=1}^\infty$ be a real-valued stochastic process such that η_t is $\mathcal{F}_t$ -measurable and η_t is conditionally R -sub-Gaussian for some $R \geq 0$. Let $\{\mathbf{x}_t\}_{t=1}^\infty$ be an $\mathbb{R}^d$ -valued stochastic process such that $\mathbf{x}_t$ is $\mathcal{F}_{t-1}$ -measurable and $\mathbf{x}_t$ is bounded in Euclidean norm, i.e., $\|\mathbf{x}_t\|_2 \leq L$ for some constant $L > 0$ and all $t \geq 1$. Assume that $\mathbf{V}$ is a $d \times d$ positive definite matrix with eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_d$ sorted in descending order: $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d > 0$. For any $t \geq 0$, define

$$\bar{\mathbf{V}}_t = \mathbf{V} + \sum_{s=1}^t \mathbf{x}_s \mathbf{x}_s^\top \quad \text{and} \quad \mathbf{S}_t = \sum_{s=1}^t \eta_s \mathbf{x}_s.$$

Then, for any $\delta > 0$, with probability at least $1 - \delta$, for $d \geq 2$ and all $t \geq 0$,

$$\|\mathbf{S}_t\|_{\bar{\mathbf{V}}_t}^2 \leq R^2 d \log \left(\frac{\lambda_1 + t L^2 / d}{\lambda_d \delta} \right).$$

LEMMA E.2 (WOODBURY MATRIX IDENTITY). For matrices $\mathbf{A}$, $\mathbf{U}$, $\mathbf{C}$ and $\mathbf{V}$ of appropriate dimensions which $\mathbf{A}$ and $\mathbf{C}$ are invertible, the inverse of the matrix sum $\mathbf{A} + \mathbf{UCV}$ can be computed as

$$(\mathbf{A} + \mathbf{UCV})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{U}(\mathbf{C}^{-1} + \mathbf{VA}^{-1}\mathbf{U})^{-1}\mathbf{VA}^{-1}.$$

LEMMA E.3 (MATRIX DETERMINANT PROPERTY). For any $\mathbf{x} \in \mathbb{R}^d$, we have

$$\det(\mathbf{I}_d + \mathbf{x}\mathbf{x}^\top) = \det(\mathbf{I}_d + \mathbf{x}\mathbf{x}^\top).$$

LEMMA E.4 (TRACE-DETERMINANT INEQUALITY). For any positive definite matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$, we have

$$\log \det(\mathbf{A}) \leq d \log \frac{1}{d} \text{trace}(\mathbf{A}).$$

LEMMA E.5 (ANTI-CONCENTRATION INEQUALITY). For a standard Gaussian distributed random variable Z with mean 0 and variance 1, it holds that for any x that

$$\frac{1}{\sqrt{2\pi}} \frac{x}{x^2 + 1} \exp(-x^2/2) \leq \mathbb{P}(Z > x) \leq \exp(-x^2/2).$$





Fig. 6. Performance under $N = 10$, $K = 5$, $d = 3$.

F Regret Performance under the Oracle Policy

Xu and Bastani (2025) define regret based on the oracle policy. Specifically, at each time step t , when $Z_t = j \in [N]$, we define an optimal policy $\pi_{j,t}^*$ that knows the true arm parameters $\{\boldsymbol{\beta}_{k,j}\}_{k \in [K]}$ of bandit j in advance and always selects the arm with the highest expected reward. That is,

$$\pi_{j,t}^* = \arg \max_{k \in [K]} \mathbf{x}_t^\top \boldsymbol{\beta}_{k,j},$$

where $\mathbf{x}_t$ is the context vector observed at time t . The expected regret at time t for bandit instance j is defined as the difference between the expected reward of the optimal arm chosen by $\pi_{j,t}^*$ and that of the arm chosen by our policy $\pi_{j,t}$: $r_{j,t} = \mathbb{E} \left[\left(\mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_{j,t}^*,j} - \mathbf{x}_t^\top \boldsymbol{\beta}_{\pi_{j,t},j} \right) \mid Z_t = j \right]$. Given that $\mathbb{P}(Z_t = j) = p_j$, the law of total expectation

Fig. 7. Performance under $N = 10$, $K = 5$, $d = 3$.

implies that the overall expected regret at time t can be written as $r_t = \sum_{j=1}^N p_j r_{j,t}$. The cumulative regret over n time steps is then $R_n = \sum_{t=1}^n r_t$, and the instance-specific cumulative regret for each bandit instance j , given by $R_{j,n} = \sum_{t=1}^n p_j r_{j,t}$. We evaluate the performance of the cumulative regret in a simulation study, and the results are reported in Figure 6, under the same model setting as in Figure 2. From Figure 6, the performance is very similar to that shown in Figure 2.

G Sensitivity to Hyperparameters

We examine the sensitivity of EBMTS and EBMUCB to the hyperparameters a and λ . We vary $a \in \{0.05, 0.1, 0.2\}$ and $\lambda \in \{0.0005, 0.001, 0.01\}$ on a grid and report the cumulative regret at the final time step as a function of each hyperparameter, while fixing the other at its default value. As shown in Figure 7, we rank the candidate configurations by cumulative regret in ascending order. Both algorithms remain robust across the considered hyperparameter values and consistently outperform the baseline.

Received 06 November 2025; accepted 22 September 2026