
Mixed-Effect Thompson Sampling

Imad Aouali
Criteo AI Lab

Branislav Kveton
Amazon

Sumeet Katariya
Amazon

Abstract

A contextual bandit is a popular framework for online learning to act under uncertainty. In practice, the number of actions is huge and their expected rewards are correlated. In this work, we introduce a general framework for capturing such correlations through a mixed-effect model where actions are related through multiple shared effect parameters. To explore efficiently using this structure, we propose **Mixed-Effect Thompson Sampling** (meTS) and bound its Bayes regret. The regret bound has two terms, one for learning the action parameters and the other for learning the shared effect parameters. The terms reflect the structure of our model and the quality of priors. Our theoretical findings are validated empirically using both synthetic and real-world problems. We also propose numerous extensions of practical interest. While they do not come with guarantees, they perform well empirically and show the generality of the proposed framework.

1 INTRODUCTION

A *contextual bandit* (Slivkins, 2019; Lattimore and Szepesvari, 2019; Li et al., 2010; Chu et al., 2011) is a popular sequential decision-making framework where an *agent* interacts with an environment over n rounds. In each round, the agent observes a *context*, takes an *action*, and receives a *reward* that depends on both the context and the taken action. The goal of the agent is to maximize the expected cumulative reward over n rounds. Since the expected rewards of actions are unknown, the agent must balance between taking the action that maximizes the estimated reward using collected data (*exploitation*), and exploring other actions to improve their estimates (*exploration*). This trade-off is often addressed using either *upper confidence bounds* (UCBs)

(Auer et al., 2002) or *Thompson sampling* (TS) (Thompson, 1933). As an example, in online advertising, contexts can be features of users, actions can be products, and the expected reward can be the click-through rate (CTR).

Efficient exploration in contextual bandits (Langford and Zhang, 2008; Dani et al., 2008; Li et al., 2010; Abbasi-Yadkori et al., 2011; Agrawal and Goyal, 2013) is an important research direction, as their action space is usually huge and naive exploration may lead to suboptimal performance. In this work, we start from the basic observation that the expected rewards of actions in real-world problems are often correlated. To model this phenomenon, we study a structured mixed-effect bandit environment where each *action parameter* depends on one or multiple *effect parameters*. The actions are related through the effect parameters. Therefore, taking an action teaches the agent about its effect parameters, which consequently teaches it about other actions that share the same effect parameters. We present three motivating examples next.

Movie recommendation: Here we want to recommend a movie to a user with the highest expected rating. User j and movie i are represented by vectors x_j (context) and θ_i (action parameter), respectively. The expected rating that user j gives to movie i is $x_j^\top \theta_i$. We assume that the vector x_j is observed. Then the most natural idea is to learn all θ_i individually using classic bandit methods (Li et al., 2010; Chu et al., 2011). This is statistically inefficient since the number of movies is often high. Fortunately, the movies could be organized into L categories and such information can be leveraged to explore efficiently. We present three approaches **(A)**, **(B)** and **(C)** that do this next.

(A) For each category $\ell \in [L]$, a parameter ψ_ℓ is learned online using all interactions with the movies in category ℓ . The parameter ψ_ℓ represents all the movies in category ℓ and is used instead of their individual θ_i . Therefore, this approach has a high bias, as all movies in the same category are assumed to have the same expected rating. This issue can be addressed by a better model. **(B)** We model each movie parameter θ_i as a random variable centered in its category parameter ψ_ℓ . Now movies in the same category no longer have the same expected rating due to the additional uncertainty. Both the category parameters ψ_ℓ and movie parameters θ_i are learned online. The former is learned

using all interactions with the movies in category ℓ , while the latter is learned using all interactions with movie i conditioned on ψ_ℓ . The category parameter ψ_ℓ is learned using more data, which helps to learn θ_i more efficiently. This is a special case of our setting and can also be viewed as extending hierarchical Bayesian bandits (Hong et al., 2022b) to multiple hierarchies. **(C)** The shortcoming of **(B)** is that each movie belongs to a *single category*, which is unrealistic. To address this issue, we allow movies to belong to *multiple categories* and then proceed as in **(B)**. To make the connection with our terminology, the categories $\ell \in [L]$ are the effects, their parameters ψ_ℓ are the effect parameters, and the movie parameters θ_i are the action parameters.

Ad placement: Here the agent selects a list (or *slate*) of M items from a catalog of L items with the objective of maximizing the CTR. We assume that the agent only receives binary bandit feedback that indicates whether the user clicked on *one of the items in the slate* (Dimakopoulou et al., 2019; Rejwan and Mansour, 2020). Again, user j and slate i are represented by x_j (context) and θ_i (action parameter), respectively. The corresponding CTR is $f(x_j^\top \theta_i)$, where f is the sigmoid function. The set of slates (of size $K \approx L^M$) is exponentially large, which makes learning θ_i individually difficult. Fortunately, the slates are related through a much smaller set of items (of size L). Therefore, the slates with common items can teach the agent about each other, which can be used to explore efficiently.

Efficient exploration is achieved by decomposing the parameter of slate i as $\theta_i = \sum_{\ell \in [L]} b_{i,\ell} \psi_\ell + \epsilon_i$. Here $\psi_\ell \in \mathbb{R}^d$ is the parameter of item ℓ and $b_{i,\ell} \in \mathbb{R}$ is a mixing weight that captures position biases. That is, $b_{i,\ell} = 0$ if item ℓ is not in slate i , and $b_{i,\ell}$ is high if item ℓ is ranked high in slate i . This captures the fact that the probability of a click on an item is biased by its position in the slate, and such bias can be estimated offline. Finally, ϵ_i is a random noise that can incorporate uncertainty due to *model misspecification*, for instance due to an estimation error of $b_{i,\ell}$. The benefit of this decomposition is that the parameter of item ℓ , ψ_ℓ , is learned using all interactions with the slates with item ℓ . The slate parameter θ_i is learned using all interactions with slate i conditioned on ψ_ℓ . This is more statistically efficient than learning θ_i individually, which only uses the interactions with slate i .

Drug design: Here the goal is to find the optimal drug design in clinical trials (Durand et al., 2018). Subject j and drug i are represented by vectors x_j and θ_i , respectively, and the expected efficacy of drug i for subject j is $x_j^\top \theta_i$. Again, the most natural idea is to learn all drug parameters θ_i individually. This leads to a statistical inefficiency though when the number of candidate drugs is high. Fortunately, we can leverage the fact that drug candidates in the same trial often share components to explore efficiently. Precisely, a drug is a combination of multiple components, each with a specific dosage. Each component ℓ is represented by a

parameter ψ_ℓ , and the drug parameter θ_i is a *known* combination of the component parameters ψ_ℓ weighted by their dosage. That is, $\theta_i = \sum_{\ell \in [L]} b_{i,\ell} \psi_\ell + \epsilon_i$, where $b_{i,\ell}$ is the dosage of component ℓ in drug i and ϵ_i is a random noise to incorporate uncertainty due to model misspecification. The efficacy of each component has an *effect* on the overall efficacy of the drug and is boosted by the dosage.

In all examples, we assume an underlying structure among the actions, that they are affected by multiple effects. In some problems, it is known how the effect arises. For instance, in the drug design, the actions are the drugs and the effects are their components. The mixing weight that relates an action (drug) to an effect (component) is the dosage of that component in the drug. In other problems, it may not be apparent how the effect arises and this has to be learned. We discuss this in detail in Section 2.6.

We make the following contributions. **1)** We formalize a general mixed-effect bandit framework represented by a two-level graphical model where each action is associated with a d -dimensional parameter that depends on *one or multiple* effect parameters. **2)** We design mixed-effect Thompson sampling (meTS), which leverages this structure to be both statistically and computationally efficient. Despite the complex structure, we show that closed-form posteriors can be derived for Gaussian instances and efficient approximations exist in more general cases. **3)** We prove that the Bayes regret of meTS is bounded by a sum of two terms: one is associated with learning the action parameters and the other quantifies the cost of learning the effect parameters. Both terms reflect the structure of the environment and the quality of priors. **4)** We show empirically that meTS and its variants perform extremely well, and are computationally efficient in both synthetic and real-world problems. **5)** Several extensions of practical interest are given in Appendix E. While they are not analyzed, they enjoy very good empirical performance.

Our setting is more general than previously studied hierarchical models (Section 6) where action parameters are centered at a single latent variable. Thus meTS has a wider range of applications, for which we gave three examples. Our algorithm is general and we provide posterior derivations that are valid for any distribution class. To showcase the generality, we also experiment with meTS on bandit problems with non-linear rewards. This also goes beyond prior works (Section 6) that often considered closed-form Gaussian posteriors only.

2 SETTING

For any positive integer n , we define $[n] = \{1, \dots, n\}$. The i -th coordinate of vector v is v_i . If the vector is already indexed, such as v_j , we write $v_{j,i}$. Let $a_1, \dots, a_n \in \mathbb{R}^d$ be n vectors. We denote by $a = (a_i)_{i \in [n]} \in \mathbb{R}^{nd}$ a vector of length nd obtained by concatenating a_1, \dots, a_n . We use

\otimes to denote the Kronecker product. The derivative of a univariate function f is denoted by \dot{f} .

2.1 Mixed-Effect Bandit

We study a setting where an agent interacts with a *contextual bandit* over n rounds. In round $t \in [n]$, the agent observes context $X_t \in \mathcal{X}$, where $\mathcal{X} \subseteq \mathbb{R}^d$ is a d -dimensional *context space*. After that, it takes an *action* A_t from an *action set* $[K]$, and then observes a *stochastic reward* $Y_t \in \mathbb{R}$ that depends on both X_t and A_t . We consider a structured problem where the expected rewards of actions are correlated. Specifically, each action $i \in [K]$ is associated with an *unknown d -dimensional action parameter* $\theta_{*,i} \in \mathbb{R}^d$. The correlations between the action parameters arise because they are derived from L shared *unknown d -dimensional effect parameters*, $\psi_{*,\ell} \in \mathbb{R}^d$ for $\ell \in [L]$. Specifically, we assume that the action parameter $\theta_{*,i}$ is sampled from the *action prior distribution* $P_{0,i}$ as $\theta_{*,i} \mid \Psi_* \sim P_{0,i}(\cdot \mid \Psi_*)$, where $\Psi_* = (\psi_{*,\ell})_{\ell \in [L]} \in \mathbb{R}^{Ld}$ is a concatenation of the effect parameters. The distribution $P_{0,i}$ can capture sparsity, when $\theta_{*,i}$ depends only on a subset of Ψ_* ; and also incorporate uncertainty due to model misspecification, when $\theta_{*,i}$ is not a deterministic function of Ψ_* . Finally, the effect parameters Ψ_* are sampled from a *joint effect prior* Q_0 , which is known by the agent and represents its initial uncertainty about Ψ_* . In summary, all variables in our environment are generated as

$$\begin{aligned} \Psi_* &\sim Q_0, \\ \theta_{*,i} \mid \Psi_* &\sim P_{0,i}(\cdot \mid \Psi_*), \quad \forall i \in [K], \\ Y_t \mid X_t, \theta_{*,A_t} &\sim P(\cdot \mid X_t; \theta_{*,A_t}), \quad \forall t \in [n], \end{aligned} \quad (1)$$

where $P(\cdot \mid x; \theta_{*,i})$ is the *reward distribution* of action i in context x , which only depends on parameter $\theta_{*,i}$ and the context x . The terminology of effect parameters arises from the fact that $\psi_{*,\ell}$ affect the model parameters $\theta_{*,i}$, which in turn define Y_t . The effects are mixed through the action prior $P_{0,i}$ and hence the name *mixed-effect*.

Our setting can be viewed as a two-level graphical model, where $\psi_{*,1}, \dots, \psi_{*,L}$ are parent nodes and $\theta_{*,1}, \dots, \theta_{*,K}$ are child nodes (Figure 1). The *structure* is represented by missing arrows from parent (effect parameters) to child (action parameters) nodes. A missing arrow from parent $\psi_{*,\ell}$ to child $\theta_{*,i}$ means that action i is independent of the ℓ -th effect. Such models are common in offline learning, for instance QMR-DT (Jaakkola and Jordan, 1999).

All examples in Section 1 can be captured by our model. In movie recommendation, the categories $\ell \in [L]$ and movies $i \in [K]$ would be represented by the effect parameters $\psi_{*,\ell}$ and action parameters $\theta_{*,i}$, respectively. The weight $b_{i,\ell}$ is the relevance of movie i to category ℓ .

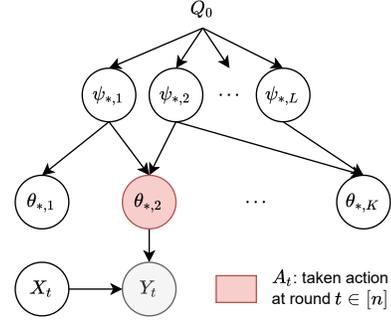


Figure 1: Example of a graphical model induced by (1).

2.2 Notion of Optimality

Let $\Theta_* = (\theta_{*,i})_{i \in [K]} \in \mathbb{R}^{Kd}$ be the concatenation of all action parameters. The expected reward of action $i \in [K]$ in context $x \in \mathcal{X}$ is $r(x, i; \Theta_*) = \mathbb{E}_{Y \sim P(\cdot \mid x; \theta_{*,i})} [Y]$, where r is the *reward function*. Our setting is Bayesian and thus a natural goal for the agent is to minimize its *Bayes regret*

$$\mathcal{BR}(n) = \mathbb{E} \left[\sum_{t=1}^n r(X_t, A_{t,*}; \Theta_*) - r(X_t, A_t; \Theta_*) \right],$$

where $A_{t,*} = \operatorname{argmax}_{i \in [K]} r(X_t, i; \Theta_*)$ is the optimal action in round t . The above expectation is over all random variables in (1). While the Bayes regret is weaker than the frequentist regret, it is a reasonable metric for average performance across multiple instances (Russo and Van Roy, 2014). We present a special case of our setting next.

2.3 Linearity in Effects

A simple yet powerful assumption is that the action prior $P_{0,i}$ is parametrized by a weighted sum of effect parameters

$$\theta_{*,i} \mid \Psi_* \sim P_{0,i} \left(\cdot \mid \sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell} \right), \quad \forall i \in [K],$$

where $b_i = (b_{i,\ell})_{\ell \in [L]} \in \mathbb{R}^L$ are L *known mixing weights* for action i . The effect ℓ on action i is determined by $b_{i,\ell}$. As an example, $b_{i,\ell} = 0$ when action i is independent of effect ℓ . This is an important special case of our setting since additive models are widely used in both theory and practice (McCullagh and Nelder, 1989), as they often lead to closed-form posteriors that are computationally tractable. Next we present two instances of this setting, where $P_{0,i}$ is a multivariate Gaussian with mean $\sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell}$ and covariance $\Sigma_{0,i}$. We defer non-linear effects to Appendix E.3.

2.4 Mixed-Effect Linear Bandit

A natural joint effect prior Q_0 for d -dimensional effect parameters $\psi_{*,\ell}$ is a multivariate Gaussian with mean $\mu_\Psi \in \mathbb{R}^{Ld}$ and covariance $\Sigma_\Psi \in \mathbb{R}^{Ld \times Ld}$. The action prior $P_{0,i}$ is

a Gaussian with mean $\sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell} \in \mathbb{R}^d$ and covariance $\Sigma_{0,i} \in \mathbb{R}^{d \times d}$. This model is a variant of a linear Gaussian model (Koller and Friedman, 2009) and is given by

$$\Psi_* \sim \mathcal{N}(\mu_\Psi, \Sigma_\Psi), \quad (2)$$

$$\theta_{*,i} \mid \Psi_* \sim \mathcal{N}\left(\sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell}, \Sigma_{0,i}\right), \quad \forall i \in [K],$$

$$Y_t \mid X_t, \theta_{*,A_t} \sim \mathcal{N}(X_t^\top \theta_{*,A_t}, \sigma^2), \quad \forall t \in [n],$$

where $\sigma > 0$ is an observation noise. This model reduces to a multi-armed bandit (Appendix B) when $X_t = 1$ for all $t \in [n]$ and $\psi_{*,\ell} \in \mathbb{R}$ for all $\ell \in [L]$.

2.5 Mixed-Effect Generalized Linear Bandit

Here the effect and action parameters are generated as in (2) but the reward Y_t is sampled from a *generalized linear model (GLM)* (McCullagh and Nelder, 1989), which is non-linear. In particular, $P(\cdot \mid X_t; \theta)$ is an exponential-family distribution with mean $f(X_t^\top \theta)$ and the whole model is

$$\Psi_* \sim \mathcal{N}(\mu_\Psi, \Sigma_\Psi), \quad (3)$$

$$\theta_{*,i} \mid \Psi_* \sim \mathcal{N}\left(\sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell}, \Sigma_{0,i}\right), \quad \forall i \in [K],$$

$$Y_t \mid X_t, \theta_{*,A_t} \sim P(\cdot \mid X_t; \theta_{*,A_t}), \quad \forall t \in [n].$$

Let $\text{Ber}(p)$ be a Bernoulli distribution with mean p . One particular choice of a GLM is $f(u) = 1/(1 + \exp(-u))$ and $P(\cdot \mid X_t; \theta) = \text{Ber}(f(X_t^\top \theta))$, which corresponds to a logistic bandit (Filippi et al., 2010).

2.6 Structure Learning

As discussed in Section 1, the structures in (2) and (3) may be intrinsic in some problems, such as drug design. When this is not the case, we propose the following approach to learning a *proxy structure*. For any $i \in [K]$, let θ_i represent an offline estimate of action parameter θ_i . To learn, we fit a Gaussian mixture model (GMM) (Reynolds et al., 2009) with L clusters to $\hat{\theta}_i$. Each cluster $\ell \in [L]$ is represented by its center $\mu_{\psi_\ell} \in \mathbb{R}^d$ and covariance $\Sigma_{\psi_\ell} \in \mathbb{R}^{d \times d}$. These correspond to the mean of the effect parameter ψ_ℓ and its uncertainty. The GMM also outputs the probability that $\hat{\theta}_i$ belongs to cluster ℓ , for all combinations of $i \in [K]$ and $\ell \in [L]$. This probability is the mixing weight $b_{i,\ell}$.

The above procedure is general and can be adapted to any use case. The main challenge is to obtain the offline estimates $\hat{\theta}_i$. This is an offline representation learning problem (Tripuraneni et al., 2021), for which numerous techniques exist. For instance, in our MovieLens experiments in Section 5.2, we use a low-rank factorization of the rating matrix to obtain these offline estimates. This is a strength of our approach. It is highly flexible and can be easily integrated with popular and practical offline learning tools. This can

Algorithm 1 meTS: Mixed-Effect Thompson Sampling.

Input: Joint effect prior Q_0 , action priors P_0 .

Initialize $Q_1 \leftarrow Q_0$ and $P_{1,\cdot} \leftarrow P_0$.

for $t = 1, \dots, n$ **do**

 Sample $\Psi_t \sim Q_t$

for $i = 1, \dots, K$ **do**

 Sample $\theta_{t,i} \sim P_{t,i}(\cdot \mid \Psi_t)$

$\Theta_t \leftarrow (\theta_{t,i})_{i \in [K]}$

$A_t \leftarrow \text{argmax}_{i \in [K]} r(X_t, i; \Theta_t)$

 Receive reward $Y_t \sim P(\cdot \mid X_t; \theta_{*,A_t})$

 Compute new posteriors Q_{t+1} and $P_{t+1,\cdot}$.

be seen as a step towards bridging the gap between offline and online learning.

3 ALGORITHM

We propose a Thompson sampling algorithm (Thompson, 1933; Russo and Van Roy, 2014; Scott, 2010), which is a natural Bayesian solution to our problem. The algorithm is based on hierarchical sampling (Lindley and Smith, 1972), which reflects the structure in our model. Before we present it, we need to introduce additional notation. We denote by $H_t = (X_\ell, A_\ell, Y_\ell)_{\ell \in [t-1]}$ the *history* of all interactions of the agent up to round t , by $S_{t,i} = \{\ell \in [t-1] : A_\ell = i\}$ the rounds where the agent takes action i up to round t , and by $H_{t,i} = (X_\ell, A_\ell, Y_\ell)_{\ell \in S_{t,i}}$ the corresponding history.

Our algorithm meTS is presented in Algorithm 1. Because the effect parameters are shared by all actions, their posteriors are not independent. For this reason, we maintain a single *joint effect posterior*

$$Q_t(\Psi) = \mathbb{P}(\Psi_* = \Psi \mid H_t)$$

for all effect parameters Ψ_* in round t . Moreover, we maintain an *action posterior*

$$P_{t,i}(\theta \mid \Psi) = \mathbb{P}(\theta_{*,i} = \theta \mid H_{t,i}, \Psi_* = \Psi)$$

for each action $i \in [K]$ given $\Psi_* = \Psi$. meTS samples hierarchically as follows. In round t , we first sample effect parameters $\Psi_t \sim Q_t$. Then we sample each action parameter $\theta_{t,i} \sim P_{t,i}(\cdot \mid \Psi_t)$ individually. Note that this is equivalent to sampling from the exact posterior $\mathbb{P}(\theta_{*,i} = \theta \mid H_t)$, since

$$\begin{aligned} \mathbb{P}(\theta_{*,i} = \theta \mid H_t) &= \int_{\Psi} \mathbb{P}(\theta_{i,*} = \theta, \Psi_* = \Psi \mid H_t) d\Psi, \\ &= \int_{\Psi} P_{t,i}(\theta \mid \Psi) Q_t(\Psi) d\Psi. \end{aligned} \quad (4)$$

Finally, we behave optimistically and take the action with the highest expected reward under the posterior-sampled action parameters $\Theta_t = (\theta_{t,i})_{i \in [K]}$.

3.1 Posterior Derivations

The posteriors are computed as follows. We first express the joint effect posterior Q_t as

$$Q_t(\Psi) \propto \prod_{i=1}^K \int_{\theta} \mathcal{L}_{t,i}(\theta) P_{0,i}(\theta | \Psi) d\theta Q_0(\Psi), \quad (5)$$

where $\mathcal{L}_{t,i}(\theta) = \mathbb{P}(H_{t,i} | \theta_{*,i} = \theta) = \prod_{(x,a,y) \in H_{t,i}} P(y | x; \theta)$ is the likelihood of all observations of action i up to round t given $\theta_{*,i} = \theta$. Next, for any action $i \in [K]$, the action posterior $P_{t,i}$ is defined as

$$P_{t,i}(\theta | \Psi) \propto \mathcal{L}_{t,i}(\theta) P_{0,i}(\theta | \Psi). \quad (6)$$

$P_{t,i}$ is similarly sparse to $P_{0,i}$. Specifically, in any round t , $P_{t,i}$ and $P_{0,i}$ are parameterized by the same subset of effect parameters Ψ_* , since $\mathcal{L}_{t,i}(\theta)$ does not depend on Ψ_* .

The joint effect posterior Q_t and action posteriors $P_{t,i}$ have closed forms in Gaussian models, which allows efficient sampling and theoretical analysis. Beyond these, MCMC and variational inference (Doucet et al., 2001) can be used to approximate Q_t and $P_{t,i}$. Next we derive closed-form posteriors for the mixed-effect model with linear rewards in (2) and provide an efficient approximation for the mixed-effect model with non-linear rewards in (3).

3.2 Mixed-Effect Linear Bandit

Let the outer product of contexts corresponding to action i up to round t be $G_{t,i} = \sigma^{-2} \sum_{\ell \in S_{t,i}} X_{\ell} X_{\ell}^{\top}$ and their sum weighted by rewards be $B_{t,i} = \sigma^{-2} \sum_{\ell \in S_{t,i}} Y_{\ell} X_{\ell}$. Both $G_{t,i}$ and $B_{t,i}$ are scaled by the observation noise σ . Using these quantities, the effect posterior is defined as follows.

Proposition 1. *For any round $t \in [n]$, the joint effect posterior is a multivariate Gaussian $Q_t = \mathcal{N}(\bar{\mu}_t, \bar{\Sigma}_t)$, where*

$$\bar{\Sigma}_t^{-1} = \Sigma_{\Psi}^{-1} + \sum_{i=1}^K b_i b_i^{\top} \otimes (\Sigma_{0,i} + G_{t,i}^{-1})^{-1}, \quad (7)$$

$$\bar{\mu}_t = \bar{\Sigma}_t \left(\Sigma_{\Psi}^{-1} \mu_{\Psi} + \sum_{i=1}^K b_i \otimes ((\Sigma_{0,i} + G_{t,i}^{-1})^{-1} G_{t,i}^{-1} B_{t,i}) \right).$$

The effect posterior is additive in individual actions and can be interpreted as follows. Each action is a single noisy observation in its estimate. The *maximum likelihood estimate (MLE)* of the parameter of action i , $G_{t,i}^{-1} B_{t,i}$, contributes to (7) proportionally to its precision, $(\Sigma_{0,i} + G_{t,i}^{-1})^{-1}$. The contribution to the ℓ -th effect parameter is weighted by $b_{i,\ell}$, which is the mixture weight for $\theta_{*,i}$ in (2). Proposition 1 is proved in Appendix C.1.

Note that $G_{t,i}$ in Proposition 1 may not be invertible. We want to stress that the formulas with it are for the ease of exposition only. The reason is that $G_{t,i}^{-1}$ appears after using

the Woodbury matrix identity to invert $\Sigma_{0,i}^{-1} + G_{t,i}$, which is well defined. Precisely, we use that

$$\begin{aligned} \Sigma_{0,i}^{-1} - \Sigma_{0,i}^{-1} (G_{t,i} + \Sigma_{0,i}^{-1})^{-1} \Sigma_{0,i}^{-1} &= (\Sigma_{0,i} + G_{t,i}^{-1})^{-1}, \\ \Sigma_{0,i}^{-1} (G_{t,i} + \Sigma_{0,i}^{-1})^{-1} B_{t,i} &= (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} G_{t,i}^{-1} B_{t,i}. \end{aligned}$$

For consistency, the formulas in Proposition 1 are also used in our experiments, where $(G_{t,i} + 10^{-3} I_d)^{-1}$ replaces $G_{t,i}^{-1}$ for numerical stability. We note that this results in a similar regret to using the correct formulas.

Now we present the action posterior.

Proposition 2. *For any round $t \in [n]$, action $i \in [K]$, and effect parameters Ψ_t , the action posterior is a multivariate Gaussian $P_{t,i}(\cdot | \Psi_t) = \mathcal{N}(\cdot; \tilde{\mu}_{t,i}, \tilde{\Sigma}_{t,i})$, where*

$$\begin{aligned} \tilde{\Sigma}_{t,i}^{-1} &= \Sigma_{0,i}^{-1} + G_{t,i}, \\ \tilde{\mu}_{t,i} &= \tilde{\Sigma}_{t,i} \left(\Sigma_{0,i}^{-1} \sum_{\ell=1}^L b_{i,\ell} \psi_{t,\ell} + B_{t,i} \right). \end{aligned} \quad (8)$$

The action posterior in (8) is a standard multivariate Gaussian posterior whose prior depends on Ψ_t , which is sampled by `meTS`. Proposition 2 is proved in Appendix C.2.

3.3 Mixed-Effect Generalized Linear Bandit

Closed-form posteriors do not exist in this setting and approximations are needed. We opt for a simple scheme that approximates the likelihood $\mathcal{L}_{t,i}(\cdot)$ by a multivariate Gaussian using the Laplace approximation. In particular, since $P(\cdot | X_t; \theta)$ is an exponential-family reward distribution, the log-likelihood for any action $i \in [K]$ is

$$\log \mathcal{L}_{t,i}(\theta) = \sum_{\ell \in S_{t,i}} Y_{\ell} X_{\ell}^{\top} \theta - A(X_{\ell}^{\top} \theta) + C(Y_{\ell}),$$

where C denotes a real function, and A is a twice continuously differentiable function whose derivative is the mean function f , $\dot{A} = f$. Let $\mu_{t,i}^{\text{LAP}}$ and $G_{t,i}^{\text{LAP}}$ be the MLE and the Hessian of $-\log \mathcal{L}_{t,i}(\cdot)$, respectively, defined as

$$\mu_{t,i}^{\text{LAP}} = \operatorname{argmax}_{\theta \in \mathbb{R}^d} \log \mathcal{L}_{t,i}(\theta), \quad (9)$$

$$G_{t,i}^{\text{LAP}} = \sum_{\ell \in S_{t,i}} \dot{f}(X_{\ell}^{\top} \mu_{t,i}^{\text{LAP}}) X_{\ell} X_{\ell}^{\top}.$$

Then the Laplace approximation is

$$\mathcal{L}_{t,i}(\cdot) \approx \mathcal{N}(\cdot; \mu_{t,i}^{\text{LAP}}, (G_{t,i}^{\text{LAP}})^{-1}). \quad (10)$$

Now we plug (10) into (5) and have $Q_t(\cdot) \approx \mathcal{N}(\cdot; \bar{\mu}_t, \bar{\Sigma}_t)$, where $\bar{\mu}_t$ and $\bar{\Sigma}_t$ are computed as in Proposition 1, except that $G_{t,i}^{\text{LAP}}$ and $\mu_{t,i}^{\text{LAP}}$ replace $G_{t,i}$ and $G_{t,i}^{-1} B_{t,i}$, respectively. We plug (10) into (6) and get $P_{t,i}(\cdot | \Psi) \approx \mathcal{N}(\cdot; \tilde{\mu}_{t,i}, \tilde{\Sigma}_{t,i})$, where $\tilde{\mu}_{t,i}$ and $\tilde{\Sigma}_{t,i}$ are computed as in Proposition 2, except

that $G_{t,i}^{\text{LAP}}$ and $G_{t,i}^{\text{LAP}} \mu_{t,i}^{\text{LAP}}$ replace $G_{t,i}$ and $B_{t,i}$, respectively. While these results are a direct consequence of plugging the Laplace approximation (10) into (5) and (6), there is a clear intuition behind them. First, $G_{t,i} \leftarrow G_{t,i}^{\text{LAP}}$ captures the change of curvature due to the non-linearity of the mean function f . Moreover, $G_{t,i}^{-1} B_{t,i} \leftarrow \mu_{t,i}^{\text{LAP}}$ follows from the fact that the MLE of the action parameter $\theta_{*,i}$ in the linear case (Section 2.4) is $G_{t,i}^{-1} B_{t,i}$, and it corresponds to $\mu_{t,i}^{\text{LAP}}$ in the generalized linear case. As discussed before, $G_{t,i}^{\text{LAP}}$ may not be invertible. Therefore, we approximate its inverse in our experiments by $(G_{t,i}^{\text{LAP}} + 10^{-3} I_d)^{-1}$.

3.4 Computational Complexity

The Bayes regret in Section 2.2 does not directly depend on Ψ_* . Thus the benefit of modeling the effect parameters is not immediately clear. It is tempting to marginalize them out, and only maintain a single joint posterior of all action parameters $\Theta_* \in \mathbb{R}^{Kd}$. Although this is feasible, posterior updates would be complex and computationally inefficient when $K \gg L$, which is common in practice.

The main advantage of meTS is that the sampling of effect parameters $\Psi_t \sim Q_t$ allows us to use the conditional independence of actions given Ψ_* , and model $\theta_{*,i} \mid H_{t,i}, \Psi_* = \Psi_t$ individually. This is more computationally efficient than modeling $\Theta_* \mid H_t$ when $K \gg L$. To see this, suppose that all posteriors are multivariate Gaussians (Section 3.2). In this case, $\Theta_* \mid H_t$ requires $\mathcal{O}(K^2 d^2)$ space, due to storing a $Kd \times Kd$ covariance matrix; while meTS requires only $\mathcal{O}((L^2 + K)d^2)$ space, due to storing the covariances of Q_t and $P_{t,i}$. Since the sampling relies on covariance inverses, the time complexity also improves. For the joint posterior, it is $\mathcal{O}(K^3 d^3)$, while it is only $\mathcal{O}((L^3 + K)d^3)$ for meTS.

One can also marginalize out the effect parameters Ψ_* and have K separate posteriors, one for each action parameter $\theta_{*,i}$. While this improves computational efficiency, it does not model that the actions are correlated, since $\theta_{*,i} \mid H_{t,i}$ is modeled instead of $\theta_{*,i} \mid H_t$. This leads to a statistical inefficiency due to the loss of information as the histories of other actions $H_{t,j}$ are discarded. We validate this through theory (Section 4.2) and experiments (Section 5).

4 ANALYSIS

This section is organized as follows. First, we state our regret bound. Second, we discuss how it captures the structure of our problem. Finally, we sketch its proof. We use $\tilde{\mathcal{O}}$ for the big O notation up to polylogarithmic factors.

4.1 Main Result

We analyze meTS in the linear setting in Section 2.4. To ease exposition, we assume that there exist $\sigma_0, \sigma_\Psi, \kappa_x > 0$ such that $\Sigma_{0,i} = \sigma_0^2 I_d$ for all $i \in [K]$, $\Sigma_\Psi = \sigma_\Psi^2 I_{Ld}$, and

$\|X_t\|_2^2 \leq \kappa_x$ for all $t \in [n]$. The last assumption is standard and we relax the remaining two in Appendix D.

Theorem 1. *For any $\delta \in (0, 1)$, the Bayes regret of meTS in the mixed-effect model in Section 2.4 is bounded as*

$$\mathcal{BR}(n) \leq \sqrt{2n(\mathcal{R}^A(n) + \mathcal{R}^E(n)) \log(1/\delta)} + cn\delta, \quad (11)$$

where $c = \sqrt{\frac{2}{\pi} \kappa_x (\sigma_0^2 + \kappa_b \sigma_\Psi^2)} K$, $\kappa_b = \max_{i \in [K]} \|b_i\|_2^2$,

$$\mathcal{R}^A(n) = dKc_A \log\left(1 + \frac{n\kappa_x \sigma_0^2}{d\sigma^2}\right), \quad c_A = \frac{\kappa_x \sigma_0^2}{\log\left(1 + \frac{\kappa_x \sigma_0^2}{\sigma^2}\right)},$$

$$\mathcal{R}^E(n) = dLc_E \log\left(1 + \frac{K\kappa_b \sigma_\Psi^2}{\sigma_0^2 + \frac{\sigma^2}{n\kappa_x}}\right), \quad c_E = \frac{\kappa_x \kappa_b \sigma_\Psi^2 \left(1 + \frac{\kappa_x \sigma_0^2}{\sigma^2}\right)}{\log\left(1 + \frac{\kappa_x \kappa_b \sigma_\Psi^2}{\sigma^2}\right)}.$$

The second term in (11) is constant for $\delta = 1/n$, in which case the above bound is $\tilde{\mathcal{O}}(\sqrt{n})$ and optimal in the horizon n . The main quantities of interest are $\mathcal{R}^A(n)$ and $\mathcal{R}^E(n)$, and they have natural interpretations. $\mathcal{R}^A(n)$ corresponds to the action regression problem: with K parameters of dimension d , prior width σ_0 , maximum context length $\sqrt{\kappa_x}$, and n observations with noise σ . The dependence of $\mathcal{R}^A(n)$ on these quantities is identical to a corresponding linear bandit (Lu and Van Roy, 2019). On the other hand, $\mathcal{R}^E(n)$ corresponds to the effect regression problem: with L parameters of dimension d , prior width σ_Ψ , maximum mixing-weight length $\sqrt{\kappa_b}$, and K actions that can be viewed as observations with noise σ_0 (Section 3.2). The dependence of $\mathcal{R}^E(n)$ on these quantities mimics those in $\mathcal{R}^A(n)$.

To simplify exposition, let $\kappa_x = \kappa_b = \sigma = 1$. Then

$$\mathcal{BR}(n) = \tilde{\mathcal{O}}\left(\sqrt{nd(K\sigma_0^2 + L\sigma_\Psi^2(1 + \sigma_0^2))}\right). \quad (12)$$

Note that $\mathcal{BR}(n)$ decreases when the initial uncertainties σ_0 and σ_Ψ are lower. Also smaller K , L , or d mean fewer parameters to learn and lead to a lower regret. We observe these trends empirically in Appendix F.

Our analysis is under the assumption that the covariances and mixing weights are known. This is typical in Bayesian analyses and represents prior knowledge that reduces regret. Parameter misspecification can be analyzed similarly to Simchowitz et al. (2021). Roughly speaking, if it was $\tilde{\mathcal{O}}(1/n^\alpha)$, the additional regret would be $\tilde{\mathcal{O}}(n^{1-\alpha})$; and thus $\tilde{\mathcal{O}}(\sqrt{n})$ when $\alpha = 0.5$.

4.2 Benefits of Structure

Note that we do not provide a matching lower bound. The only Bayesian lower bound that we know of is $\mathcal{O}(\log^2 n)$ for a K -armed bandit (Theorem 3 of Lai (1987)). Seminal works on Bayes regret minimization (Russo and Van Roy, 2014, 2016) do not match it. Therefore, to argue that our bound reflects the problem structure, we compare meTS to agents that have access to more information or use less

structure. We start with those with more information. Take meTS with known effect parameters Ψ_* . Then $\sigma_\Psi = 0$ in (12) and we obtain a lower regret $\mathcal{BR}(n) = \tilde{O}(\sqrt{ndK\sigma_0^2})$ that does not depend on L . Similarly, take meTS with a perfect linear model, $\theta_{*,i} = \sum_{\ell \in [L]} b_{i,\ell} \psi_{*,\ell}$ for all $i \in [K]$. Then $\sigma_0 = 0$ in (12) and we get a lower regret $\mathcal{BR}(n) = \tilde{O}(\sqrt{ndL\sigma_\Psi^2})$ that does not depend on K . This shows that K in our bound arises due to modeling the stochasticity of action parameters with respect to the effect parameters, incorporated in $\Sigma_{0,i}$.

Next we consider an agent that does not know Ψ_* and also does not model it. Here only Θ_* is learned (Section 3.4) and this is achieved by marginalizing out Ψ_* in (2) as

$$\begin{aligned} \theta_{*,i} &\sim \mathcal{N}\left(\sum_{\ell=1}^L b_{i,\ell} \mu_{\psi_\ell}, \check{\Sigma}_{0,i}\right), \quad \forall i \in [K], \\ Y_t | X_t, \theta_{*,A_t} &\sim \mathcal{N}(X_t^\top \theta_{*,A_t}, \sigma^2), \quad \forall t \in [n], \end{aligned}$$

where $\check{\Sigma}_{0,i} = (\sigma_0^2 + \|b_i\|_2^2 \sigma_\Psi^2) I_d$ is the marginal prior covariance of action i and $\mu_{\psi_\ell} \in \mathbb{R}^d$ is the prior mean of the ℓ -th effect, which satisfies $\mu_\Psi = (\mu_{\psi_\ell})_{\ell \in [L]}$ (Section 2.4). The marginal prior covariance $\check{\Sigma}_{0,i}$ accounts for the uncertainty of the not-modeled Ψ_* weighted by $\|b_i\|_2^2$. The regret of this agent scales as in (12) with $\sigma_\Psi = 0$, except that the maximum prior variance σ_0^2 is replaced with the maximum marginal prior variance $\sigma_0^2 + \sigma_\Psi^2$. This can be proved using the definition of $\check{\Sigma}_{0,i}$ and $\kappa_b = \max_{i \in [K]} \|b_i\|_2^2 = 1$. The resulting regret bound is $\mathcal{BR}(n) = \tilde{O}(\sqrt{ndK(\sigma_0^2 + \sigma_\Psi^2)})$. When $K > L$ up to constants, it can be significantly higher than the regret bound of meTS in (12). The improvement is on the order of $\sqrt{K/L}$ when the effect parameters are much more uncertain than the action ones, $\sigma_\Psi \gg \sigma_0$, which is expected. For instance, in our ad placement example, L is the number of items in the catalog and $K \approx L^M$ is the number of slates of size M . Hence $K/L \approx L^{M-1}$, where we can have $L \approx 10^6$ and $M \approx 10$. This claim is also supported empirically in Section 5.1 and Appendix F, where meTS outperforms classic methods when the effect parameters are more uncertain than the action ones.

4.3 Sketch of the Regret Proof

Now we outline the key technical challenges and novel insights in our proof. Our hierarchical sampling is equivalent to sampling from the exact posterior, which is also a multivariate Gaussian $\mathbb{P}(\theta_{*,i} = \theta | H_t) = \mathcal{N}(\theta; \hat{\mu}_{t,i}, \hat{\Sigma}_{t,i})$ for some $\hat{\mu}_{t,i}$ and $\hat{\Sigma}_{t,i}$. This is because the action posterior $P_{t,i}$ and effect posterior Q_t are Gaussians, and Gaussianity is preserved after marginalization (Koller and Friedman, 2009). Next notice that the context X_t in round t is known, and thus we include it in the history H_t . Now let $\mathbf{A}_t \in \{0, 1\}^K$ and $\mathbf{A}_{t,*} \in \{0, 1\}^K$ be the indicator vectors of the taken action A_t and optimal action $A_{t,*}$, respectively. Moreover, let $\hat{\theta}_t = (X_t^\top \hat{\mu}_{t,i})_{i \in [K]}$ and $\theta_{t,*} = (X_t^\top \theta_{*,i})_{i \in [K]}$. Then the Bayes regret can be decomposed following Russo and

Van Roy (2014) as

$$\begin{aligned} \mathcal{BR}(n) &= \mathbb{E} \left[\sum_{t=1}^n \mathbf{A}_{t,*}^\top \theta_{t,*} - \mathbf{A}_t^\top \theta_{t,*} \right], \\ &= \mathbb{E} \left[\mathbb{E} \left[\mathbf{A}_{t,*}^\top (\theta_{t,*} - \hat{\theta}_t) \mid H_t \right] \right] + \mathbb{E} \left[\mathbb{E} \left[\mathbf{A}_t^\top (\hat{\theta}_t - \theta_{t,*}) \mid H_t \right] \right]. \end{aligned} \quad (13)$$

The identity (13) holds because $\hat{\theta}_t$ is deterministic given H_t (which now includes X_t), and the actions $\mathbf{A}_{t,*}$ and \mathbf{A}_t are i.i.d. given H_t . Conditioned on H_t , $\hat{\theta}_t - \theta_{t,*}$ is a zero-mean Gaussian random vector independent of \mathbf{A}_t , and therefore $\mathbb{E}[\mathbf{A}_t^\top (\hat{\theta}_t - \theta_{t,*}) \mid H_t] = 0$. Thus we only need to bound the first term in (13), which can be further bounded as

$$\mathcal{BR}(n) \leq \sqrt{2n \log(1/\delta)} \sqrt{\mathbb{E} \left[\sum_{t=1}^n \|X_t\|_{\check{\Sigma}_{t,A_t}}^2 \right]} + cn\delta. \quad (14)$$

To bound (14), we need to bound a $\hat{\Sigma}_{t,A_t}$ -norm, while we only know closed forms of $\check{\Sigma}_t$ and $\check{\Sigma}_{t,A_t}$ (Section 3.2). We relate these norms by generalizing the total covariance decomposition in Hong et al. (2022b) to incorporate multiple effect parameters and their mixing weights $b_{i,\ell}$. To account for multiple effects, we represent all effects using a single Ld -dimensional vector $\Psi_* = (\psi_{*,\ell})_{\ell \in [L]}$. Moreover, we let $\Gamma_i = b_i^\top \otimes I_d \in \mathbb{R}^{d \times Ld}$ and observe that

$$\sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell} = \Gamma_i \Psi_*, \quad \forall i \in [K]. \quad (15)$$

The key insight here is that we rewrite $\sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell}$ as a linear function of Ψ_* , $\Gamma_i \Psi_*$, and encode the dependencies between the action parameter $\theta_{*,i}$ and all effect parameters Ψ_* in Γ_i . This reformulation allows us to extend the total covariance decomposition (Lemma 3 in Appendix D.2) as

$$\hat{\Sigma}_{t,i} = \check{\Sigma}_{t,i} + \check{\Sigma}_{t,i} \Sigma_{0,i}^{-1} \Gamma_i \check{\Sigma}_t \Gamma_i^\top \Sigma_{0,i}^{-1} \check{\Sigma}_{t,i}, \quad \forall i \in [K]. \quad (16)$$

The first term in (16) captures uncertainty in $\theta_{*,i} | \Psi_*$. The second captures uncertainty in Ψ_* , weighted by $\check{\Sigma}_{t,i}$, $\Sigma_{0,i}$, and the mixing weights Γ_i for action i . The term $\Gamma_i \check{\Sigma}_t \Gamma_i^\top$ is controlled using the maximum eigenvalue of $\Gamma_i \Gamma_i^\top$, which is at most $\max_{i \in [K]} \|b_i\|_2^2$. Finally, we use the identities in (15) and (16) to bound the $\hat{\Sigma}_{t,A_t}$ -norm in (14). By careful analysis, our regret bound still reflects the structure and captures potential sparsity.

5 EXPERIMENTS

We evaluate meTS on both synthetic and real-world problems. In each plot, we report the average values and their standard errors. Additional experiments are conducted in Appendix F. The code is provided in this [Github repository](#).

5.1 Synthetic Experiments

We start with two synthetic problems: the linear and logistic bandit settings in (2) and (3), respectively. The effect prior

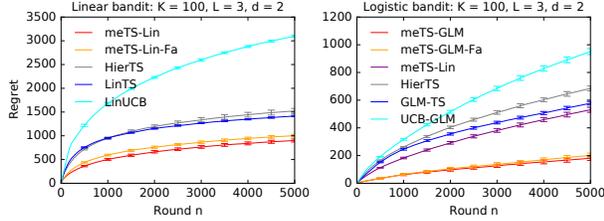


Figure 2: Evaluation on synthetic problems.

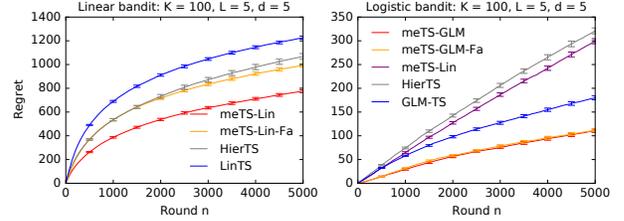


Figure 3: Evaluation on MovieLens problems.

is parameterized by $\mu_\Psi = \mathbf{0}_{Ld}$ and $\Sigma_\Psi = 3I_{Ld}$, the action covariance is $\Sigma_{0,i} = I_d$ for all $i \in [K]$, and the observation noise is $\sigma = 1$. We use this setting since modeling of the effect parameters is the most beneficial when they are more uncertain than the action ones (Section 4.2). The context X_t is sampled uniformly from $[-1, 1]^d$. We run 50 simulations and sample the mixing weights $b_{i,\ell}$ from $[-1, 1]$ in each run.

We consider the following baselines. For the linear setting, we compare meTS-Lin (Section 3.2), LinUCB (Abbasi-Yadkori et al., 2011), LinTS (Agrawal and Goyal, 2013) and HierTS (Hong et al., 2022b). For the logistic setting, we compare meTS-GLM (Section 3.3), meTS-Lin (Section 3.2), UCB-GLM (Li et al., 2017), GLM-TS (Chapelle and Li, 2012) and HierTS (Hong et al., 2022b). GLM-UCB (Filippi et al., 2010) is not included because it has a very high regret. We also include factored approximations of meTS (meTS-Lin-Fa and meTS-GLM-Fa), where the effect parameters are sampled individually (Appendix E.1). This improves the time and space complexities of meTS by L^2 and L , respectively.

All baselines but HierTS ignore the structure. HierTS incorporates the structure similarly to meTS-Lin but only has a single effect parameter with prior $\mathcal{N}(\mathbf{0}_d, 3I_d)$, with the same mean and covariance as the effect parameters of meTS. To compare fairly with LinTS and GLM-TS, their marginal prior mean and covariance are chosen as $\mathbf{0}_d$ and $\check{\Sigma}_{0,i} = \Sigma_{0,i} + \Gamma_i \Sigma_\Psi \Gamma_i^\top$, where $\Gamma_i = b_i^\top \otimes I_d$. This is to account for the uncertainty of the effect parameters despite marginalizing them out.

In Figure 2, we plot the regret in both problems for $n = 5000$, $K = 100$, $L = 3$, and $d = 2$. meTS and its factored variant outperform all baselines that ignore the structure or incorporate it partially. Moreover, meTS-GLM outperforms meTS-Lin in the logistic bandit, which shows the benefit of the approximation in Section 3.3. This attests to the generality and flexibility of meTS and the posterior derivations in Section 3. We also show in Appendix F.1 that a higher K , L , or d leads to a higher regret due to learning more parameters, which is captured by our regret bounds.

5.2 MovieLens Experiments

We study the problem of movie recommendation using the MovieLens 1M dataset (Lam and Herlocker, 2016). This

dataset contains one million ratings given by 6040 users to 3952 movies. We apply low-rank factorization to the rating matrix to obtain 5-dimensional representations: $x_j \in \mathbb{R}^5$ for user $j \in [6040]$ and $\theta_i \in \mathbb{R}^5$ for movie $i \in [3952]$. We use the movies as actions and the context X_t is sampled uniformly from user vectors x_j . We consider both linear and logistic rewards. Given a user x_j , the linear reward for movie θ_i is sampled from $\mathcal{N}(x_j^\top \theta_i, \sigma^2)$ while the logistic reward is sampled from $\text{Ber}(f(x_j^\top \theta_i))$, where f is the sigmoid function. We run 50 simulations with $K = 100$ randomly sampled movies in each run. We compare meTS to most baselines in Section 5.1. We do not include UCB-GLM and GLM-UCB because their regret is very high. In LinTS and GLM-TS, the prior mean of action i is μ and its covariance is $\check{\Sigma}_0 = \text{diag}(v) \in \mathbb{R}^{d \times d}$, where $\mu \in \mathbb{R}^d$ and $v \in \mathbb{R}^d$ are the mean and variance of the movie vectors along all dimensions, respectively.

The mixed-effect structure in (2) and (3) is not available in this problem. Therefore, we use the approach in Section 2.6 to learn it. More precisely, we cluster the movies into $L = 5$ mixture components by training a GMM on the offline action vectors θ_i (Section 2.6). Each cluster center corresponds to an effect parameter mean $\mu_{\psi_\ell} \in \mathbb{R}^d$ and the mixing weight $b_{i,\ell}$ is the probability that movie i belongs to cluster ℓ , as given by the GMM. We set the effect prior covariance as $\Sigma_\Psi = 0.75 \text{diag}((\check{\Sigma}_0)_{\ell \in [L]}) \in \mathbb{R}^{Ld \times Ld}$ and the prior covariance of action i as $\Sigma_{0,i} = 0.25 \check{\Sigma}_0 \in \mathbb{R}^{d \times d}$, where $\check{\Sigma}_0$ is the same as in both LinTS and GLM-TS. This means that the marginal covariance of action i in meTS is $0.25 \check{\Sigma}_0 + 0.75 \Gamma_i \Sigma_\Psi \Gamma_i^\top$, where $\Gamma_i = b_i^\top \otimes I_d$. Therefore, it is on the same order as $\check{\Sigma}_{0,i}$ when $\|b_i\|_2^2 \approx 1$, and meTS is parameterized comparably to LinTS and GLM-TS. At the same time, we also model that the effect parameters are more uncertain than the action ones, since $\Sigma_{0,i} = 0.25 \check{\Sigma}_0$ while $\Sigma_\Psi = 0.75 \text{diag}((\check{\Sigma}_0)_{\ell \in [L]})$.

In Figure 3, we plot the regret for $n = 5000$ rounds. We observe that meTS has the lowest regret, even if the true rewards are not generated from a mixed-effect model. This shows the robustness of meTS to model misspecification, which we further validate in Appendix F.3. It also highlights the flexibility of our framework, where a proxy structure is learned from offline data.

6 RELATED WORK

Thompson sampling (TS) (Thompson, 1933) is a popular exploration algorithm in practice (Chapelle and Li, 2012; Russo et al., 2018). Its first Bayes regret bound was proved by Russo and Van Roy (2014). We apply TS to two-level graphical models with multiple parents. Many recent works (Bastani et al., 2019; Kveton et al., 2021; Basu et al., 2021; Simchowitz et al., 2021; Wan et al., 2021; Hong et al., 2022b; Peleg et al., 2022; Wan et al., 2022; Tomkins et al., 2021) applied TS to the two-level models with a single parent. The main difference in our work is that we consider a mixed-effect model with multiple parents in the contextual setting. Urteaga and Wiggins (2018) proposed TS with a mixture reward distribution. This is very different from the parameter mixtures in our work.

Our analysis (Section 4.3) extends Hong et al. (2022b) to multiple effects in the contextual bandit setting. The main technical challenges are generalizing the regret decomposition in Hong et al. (2022b) to include context and extending their total covariance decomposition to multiple effect parameters with mixing weights. This extension is non-trivial since Hong et al. (2022b) assumed that all action parameters are centered at a single variable. In our setting, this is not true even if we treated all effect parameters as a single vector. Moreover, the action parameters can depend on a small subset of effect parameters, resulting in sparsity that is not captured by their analysis. Although information theory can be used to derive Bayes regret bounds (Russo and Van Roy, 2016; Lu and Van Roy, 2019), we are unaware of any for multiple effects.

We also assume that there exists an underlying structure among the actions. Many such structures have been studied and we review some below. In latent bandits (Maillard and Mannor, 2014; Hong et al., 2020), a single latent variable indexes multiple candidate models. In structured finite-armed bandits (Lattimore and Munos, 2014; Gupta et al., 2018), each arm is associated with a known mean function. The mean functions are parameterized by a shared latent parameter, which is learned. TS was also applied to more complex models, such as graphical models (Yu et al., 2020) and a discretized parameter space (Gopalan et al., 2014). While these frameworks are general, the computational and statistical efficiencies are not guaranteed simultaneously. Meta- and multi-task learning with UCBs have a long history in bandits (Azar et al., 2013; Gentile et al., 2014; Deshmukh et al., 2017; Cella et al., 2020). These works are frequentist, analyze a stronger notion of regret, and often lead to conservative algorithms. In contrast, our approach is Bayesian, we analyze its Bayes regret, and meTS performs well as analyzed without any additional tuning.

Our work is also related to representation learning in multi-task linear bandits (Cella et al., 2022; Hu et al., 2021; Yang et al., 2020). We refer to these works collectively as *representation learning bandits*.

Representation learning bandits. Representation learning bandits can be viewed in our notation as learning $\Theta_* = \Psi_*\Gamma$, where $\Theta_* \in \mathbb{R}^{d \times K}$ is a matrix of action parameters, $\Psi_* \in \mathbb{R}^{d \times L}$ is a matrix of effect parameters, and $\Gamma \in \mathbb{R}^{L \times K}$ is a matrix of mixing weights. Therefore, representation learning bandits assume that the action parameters are a perfect linear combination of effect parameters, $\Sigma_{0,i} = \mathbf{0}_{d,d}$, where $\mathbf{0}_{d,d}$ denotes a $d \times d$ zero matrix. This shows that our setting is more general, since we consider $\Sigma_{0,i} \neq \mathbf{0}_{d,d}$ due to action parameter uncertainty. Consequently, representation learning bandits can have linear regret when $\Sigma_{0,i} \neq \mathbf{0}_{d,d}$, due to model misspecification.

On the other hand, representation learning bandits learn Γ while we assume that it is given. So they can be viewed as more general. Note that the factorization $\Theta_* = \Psi_*\Gamma$ is only beneficial when $d \gg L$. In this case, $\Psi_*\Gamma$ has $dL + KL$ parameters while Θ_* would have $dK \gg dL + KL$. We do not assume that $d \gg L$. In fact, in all of our experiments, $dL + KL > dK$ and thus learning of Θ_* directly is more practical. This is what Lin-TS and GLM-TS already do, and meTS outperforms them in all experiments. Thus representation learning bandits would not be competitive in our setting. This highlights another major difference from representation learning bandits: their L is the number of latent dimensions while ours is the number of effects or clusters. These are two different approaches to modeling, although coinciding algebraically when $\Sigma_{0,i} = \mathbf{0}_{d,d}$.

meTS can be extended to unknown mixing weights. However, this would require solving a matrix factorization problem online, which is expensive and representation learning bandits suffer from the same computational challenge. Since our goal is to design practical and efficient algorithms, we focus on known mixing weights. They are either given or learned offline using off-the-shelf techniques (Section 2.6).

7 CONCLUSIONS

We propose a mixed-effect bandit framework represented by a two-level graphical model where actions can depend on multiple effects. This structure can be used to explore more efficiently, and we design an exploration algorithm meTS that leverages it. meTS performs well on both synthetic and real-world problems when implemented as analyzed. Our algorithmic ideas apply to the general mixed-effect model in (1), although we focus on models in (2) and (3). Our algorithmic and theory foundations open the door to richer models, which we discuss in great detail in Appendix E.

References

Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems 24*, pages 2312–2320, 2011.

- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of the 30th International Conference on Machine Learning*, pages 127–135, 2013.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- Mohammad Gheshlaghi Azar, Alessandro Lazaric, and Emma Brunskill. Sequential transfer in multi-armed bandit with finite set of models. In *Advances in Neural Information Processing Systems 26*, pages 2220–2228, 2013.
- Hamsa Bastani, David Simchi-Levi, and Ruihao Zhu. Meta dynamic pricing: Transfer learning across experiments. *CoRR*, abs/1902.10918, 2019. URL <https://arxiv.org/abs/1902.10918>.
- Soumya Basu, Branislav Kveton, Manzil Zaheer, and Csaba Szepesvari. No regrets for learning the prior in bandits. In *Advances in Neural Information Processing Systems 34*, 2021.
- Christopher M Bishop. *Pattern Recognition and Machine Learning*, volume 4 of *Information science and statistics*. Springer, 2006.
- Leonardo Cella, Alessandro Lazaric, and Massimiliano Pontil. Meta-learning with stochastic linear bandits. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Leonardo Cella, Karim Lounici, and Massimiliano Pontil. Multi-task representation learning with stochastic linear bandits. *arXiv preprint arXiv:2202.10066*, 2022.
- Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. In *Advances in Neural Information Processing Systems 24*, pages 2249–2257, 2012.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.
- Varsha Dani, Thomas Hayes, and Sham Kakade. Stochastic linear optimization under bandit feedback. In *Proceedings of the 21st Annual Conference on Learning Theory*, pages 355–366, 2008.
- Aniket Anand Deshmukh, Urun Dogan, and Clayton Scott. Multi-task learning for contextual bandits. In *Advances in Neural Information Processing Systems 30*, pages 4848–4856, 2017.
- Maria Dimakopoulou, Nikos Vlassis, and Tony Jebara. Marginal posterior sampling for slate bandits. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 2223–2229. International Joint Conferences on Artificial Intelligence Organization, 2019.
- Arnaud Doucet, Nando de Freitas, and Neil Gordon. *Sequential Monte Carlo Methods in Practice*. Springer, New York, NY, 2001.
- Audrey Durand, Charis Achilleos, Demetris Iacovides, Katerina Strati, Georgios D. Mitsis, and Joelle Pineau. Contextual bandits for adapting treatment in a mouse model of de novo carcinogenesis. In *Proceedings of the 3rd Machine Learning for Healthcare Conference*, volume 85, pages 67–82, 2018.
- Sarah Filippi, Olivier Cappé, Aurelien Garivier, and Csaba Szepesvari. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems 23*, pages 586–594, 2010.
- Claudio Gentile, Shuai Li, and Giovanni Zappella. Online clustering of bandits. In *Proceedings of the 31st International Conference on Machine Learning*, pages 757–765, 2014.
- Aditya Gopalan, Shie Mannor, and Yishay Mansour. Thompson sampling for complex online problems. In *Proceedings of the 31st International Conference on Machine Learning*, pages 100–108, 2014.
- Samarth Gupta, Shreyas Chaudhari, Subhojyoti Mukherjee, Gauri Joshi, and Osman Yagan. A unified approach to translate classical bandit algorithms to the structured bandit setting. *CoRR*, abs/1810.08164, 2018. URL <https://arxiv.org/abs/1810.08164>.
- Joey Hong, Branislav Kveton, Manzil Zaheer, Yinlam Chow, Amr Ahmed, and Craig Boutilier. Latent bandits revisited. In *Advances in Neural Information Processing Systems 33*, 2020.
- Joey Hong, Branislav Kveton, Sumeet Katariya, Manzil Zaheer, and Mohammad Ghavamzadeh. Deep hierarchy in bandits. In *Proceedings of the 39th International Conference on Machine Learning*, 2022a.
- Joey Hong, Branislav Kveton, Manzil Zaheer, and Mohammad Ghavamzadeh. Hierarchical Bayesian bandits. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, 2022b.
- Jiachen Hu, Xiaoyu Chen, Chi Jin, Lihong Li, and Liwei Wang. Near-optimal representation learning for linear bandits and linear rl. In *International Conference on Machine Learning*, pages 4349–4358. PMLR, 2021.
- Tommi S. Jaakkola and Michael I. Jordan. Variational probabilistic inference and the qmr-dt network. *J. Artif. Int. Res.*, 10(1):291–322, 1999.
- Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, MA, 2009.
- Branislav Kveton, Mikhail Konobeev, Manzil Zaheer, Chih-Wei Hsu, Martin Mladenov, Craig Boutilier, and Csaba Szepesvari. Meta-Thompson sampling. In *Proceedings of*

- the 38th International Conference on Machine Learning*, 2021.
- Tze Leung Lai. Adaptive treatment allocation and the multi-armed bandit problem. *The Annals of Statistics*, 15(3): 1091–1114, 1987.
- Shyong Lam and Jon Herlocker. MovieLens Dataset. <http://grouplens.org/datasets/movielens/>, 2016.
- John Langford and Tong Zhang. The epoch-greedy algorithm for contextual multi-armed bandits. In *Advances in Neural Information Processing Systems 20*, pages 817–824, 2008.
- Tor Lattimore and Remi Munos. Bounded regret for finite-armed structured bandits. In *Advances in Neural Information Processing Systems 27*, pages 550–558, 2014.
- Tor Lattimore and Csaba Szepesvari. *Bandit Algorithms*. Cambridge University Press, 2019.
- Lihong Li, Wei Chu, John Langford, and Robert Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, 2010.
- Lihong Li, Yu Lu, and Dengyong Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2071–2080, 2017.
- Dennis Lindley and Adrian Smith. Bayes estimates for the linear model. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(1):1–18, 1972.
- Xiuyuan Lu and Benjamin Van Roy. Information-theoretic confidence bounds for reinforcement learning. In *Advances in Neural Information Processing Systems 32*, 2019.
- Odalric-Ambrym Maillard and Shie Mannor. Latent bandits. In *Proceedings of the 31st International Conference on Machine Learning*, pages 136–144, 2014.
- P. McCullagh and J. A. Nelder. *Generalized Linear Models*. Chapman & Hall, 1989.
- Amit Peleg, Naama Pearl, and Ron Meir. Metalearning linear bandits by prior update. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, 2022.
- Idan Rejwan and Yishay Mansour. Top- k combinatorial bandits with full-bandit feedback. In *ALT*, pages 752–776, 2020.
- Douglas A Reynolds et al. Gaussian mixture models. *Encyclopedia of biometrics*, 741(659-663), 2009.
- Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.
- Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of Thompson sampling. *Journal of Machine Learning Research*, 17(68):1–30, 2016.
- Daniel Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on Thompson sampling. *Foundations and Trends in Machine Learning*, 11(1):1–96, 2018.
- Steven Scott. A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26:639 – 658, 2010.
- Max Simchowitz, Christopher Tosh, Akshay Krishnamurthy, Daniel Hsu, Thodoris Lykouris, Miro Dudik, and Robert Schapire. Bayesian decision-making under misspecified priors with applications to meta-learning. In *Advances in Neural Information Processing Systems 34*, 2021.
- Aleksandrs Slivkins. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2): 1–286, 2019.
- William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.
- Sabina Tomkins, Peng Liao, Predrag Klasnja, and Susan Murphy. Intelligentpooling: Practical thompson sampling for mhealth. *Machine learning*, 110(9):2685–2727, 2021.
- Nilesh Tripuraneni, Chi Jin, and Michael Jordan. Provable meta-learning of linear representations. In *International Conference on Machine Learning*, pages 10434–10443. PMLR, 2021.
- Inigo Urteaga and Chris Wiggins. Variational inference for the multi-armed contextual bandit. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, pages 698–706, 2018.
- Runzhe Wan, Lin Ge, and Rui Song. Metadata-based multi-task bandits with Bayesian hierarchical models. In *Advances in Neural Information Processing Systems 34*, 2021.
- Runzhe Wan, Lin Ge, and Rui Song. Towards scalable and robust structured bandits: A meta-learning framework. *CoRR*, abs/2202.13227, 2022. URL <https://arxiv.org/abs/2202.13227>.
- Neil Weiss. *A Course in Probability*. Addison-Wesley, 2005.
- Jiaqi Yang, Wei Hu, Jason D Lee, and Simon S Du. Impact of representation learning in linear bandits. *arXiv preprint arXiv:2010.06531*, 2020.
- Tong Yu, Branislav Kveton, Zheng Wen, Ruiyi Zhang, and Ole Mengshoel. Graphical models meet bandits: A variational Thompson sampling approach. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.

Mixed-Effect Thompson Sampling: Supplementary Materials

ORGANIZATION

The supplementary material is organized as follows. In **Appendix A**, we include a more visual notation and present some preliminary results that we use in our analysis. In **Appendix B**, we introduce the mixed-effect multi-armed bandit setting and provide a closed-form solution for the corresponding effect posterior and action posteriors. In **Appendix C**, we give the derivations of the effect posterior and action posteriors for the mixed-effect linear bandit setting. These proofs can be easily extended to the generalized linear case. In **Appendix D**, we prove an upper bound for Bayes regret of meTS. In **Appendix E**, we discuss in details possible extensions of this work. In **Appendix F**, we present additional experiments.

A PRELIMINARIES

In this section, we include additional notation and provide some basic properties of matrix operations.

A.1 Notation

For any positive integer n , we define $[n] = \{1, 2, \dots, n\}$. We use I_d to denote the identity matrix of dimension $d \times d$. Unless specified, the i -th coordinate of a vector v is v_i . When the vector is already indexed, such as v_j , we write $v_{j,i}$. Similarly, the (i, j) -th entry of a matrix A is $A_{i,j}$. Let $a_1 \in \mathbb{R}^d, \dots, a_n \in \mathbb{R}^d$ be n vectors. We use $A = [a_1, a_2, \dots, a_n] \in \mathbb{R}^{d \times n}$ to denote the $d \times n$ matrix obtained by horizontal concatenation of vectors a_1, \dots, a_n such that the j -th column of A is a_j and its (i, j) -th entry is $A_{i,j} = a_{j,i}$. We also denote by $a = (a_i)_{i \in [n]} \in \mathbb{R}^{nd}$ a vector of length nd obtained by concatenation of vectors a_1, \dots, a_n . $\text{Vec}(\cdot)$ denotes the vectorization operator. For instance, we have that $\text{Vec}([a_1, \dots, a_n]) = (a_i)_{i \in [n]}$. For any matrix $A \in \mathbb{R}^{d \times d}$, we use $\lambda_1(A)$ and $\lambda_d(A)$ to denote the maximum and minimum eigenvalue of A , respectively. Let A_1, \dots, A_n be n matrices of dimension $d \times d$. Then $\text{diag}((A_i)_{i \in [n]}) \in \mathbb{R}^{nd \times nd}$ denotes the block diagonal matrix where A_1, \dots, A_n are the main-diagonal blocks. Similarly, $(A_i)_{i \in [n]} \in \mathbb{R}^{nd \times d}$ is the $nd \times d$ matrix obtained by concatenation of A_1, \dots, A_n . We use \otimes to denote the Kronecker product. Now we provide a more visual presentation of the notation above. Let $a_1 \in \mathbb{R}^d, \dots, a_n \in \mathbb{R}^d$ be n vectors of dimension d , and let $A_1 \in \mathbb{R}^{d \times d}, \dots, A_n \in \mathbb{R}^{d \times d}$ be n matrices of dimension $d \times d$. We have that

$$[a_1, \dots, a_n] = \begin{pmatrix} | & | & \cdots & | \\ a_1 & a_2 & \cdots & a_n \\ | & | & \cdots & | \end{pmatrix} \in \mathbb{R}^{d \times n}, \quad (a_i)_{i \in [n]} = \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} \in \mathbb{R}^{nd},$$

$$\text{diag}((A_i)_{i \in [n]}) = \begin{pmatrix} A_1 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_n \end{pmatrix} \in \mathbb{R}^{nd \times nd}, \quad (A_i)_{i \in [n]} = \begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_n \end{pmatrix} \in \mathbb{R}^{nd \times d}.$$

Finally, let $A_{i,j} \in \mathbb{R}^{d \times d}$, for $i \in [n]$ and $j \in [m]$ be nm matrices of dimensions $d \times d$. We use $(A_{i,j})_{(i,j) \in [n] \times [m]}$ to denote the $nd \times md$ block matrix where $A_{i,j}$ is the (i, j) -th block. We also provide a more visual presentation for this notation.

$$(A_{i,j})_{(i,j) \in [n] \times [m]} = \begin{pmatrix} A_{1,1} & A_{1,2} & \cdots & A_{1,m} \\ A_{2,1} & A_{2,2} & \cdots & A_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ A_{n,1} & A_{n,2} & \cdots & A_{n,m} \end{pmatrix} \in \mathbb{R}^{nd \times md}.$$

A.2 Preliminary Results

In this section, we recall some basic properties of matrix operations.

- (a) **The mixed-product property.** We have that $(A \otimes B)(C \otimes D) = AC \otimes BD$ for any matrices A, B, C, D such that the products AC and BD exist.
- (b) **Transpose.** We have that $(A \otimes B)^\top = A^\top \otimes B^\top$ for any matrices A, B .
- (c) **Vectorization.** Let $A \in \mathbb{R}^{n \times m}, B \in \mathbb{R}^{m \times p}$, then $\text{Vec}(AB) = (I_p \otimes A) \text{Vec}(B) = (B^\top \otimes I_n) \text{Vec}(A)$.
- (d) For any matrix A , we have that $I_1 \otimes A = A$.
- (e) For any matrices A and B , we have that $\lambda_1(A \otimes B) = \lambda_1(A)\lambda_1(B)$.
- (f) For any matrix A and any positive semi-definite matrix B such that the product $A^\top BA$ exists, the following inequality holds $\lambda_1(A^\top BA) \leq \lambda_1(B)\lambda_1(A^\top A)$.

B MIXED-EFFECT MULTI-ARMED BANDIT

We introduce the mixed-effect multi-armed bandit setting. We then provide the effect posterior and action posteriors for this setting. The results below can be derived from Propositions 1 and 2 by setting $d = 1$ and $X_t = 1$ for all $t \in [n]$.

B.1 Mixed-Effect Multi-Armed Bandit

For scalar effect parameters, $\psi_{*,\ell} \in \mathbb{R}$ holds for all $\ell \in [L]$, a natural effect prior Q_0 is a multivariate Gaussian with mean $\mu_\Psi \in \mathbb{R}^L$ and covariance $\Sigma_\Psi \in \mathbb{R}^{L \times L}$. The action prior $P_{0,i}$ is a univariate Gaussian with mean $\sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell} = b_i^\top \Psi_* \in \mathbb{R}$ and variance $\sigma_{0,i}^2 > 0$. This is a non-contextual setting ($X_t = 1$ for any $t \in [n]$) and thus the whole model reads

$$\begin{aligned} \Psi_* &\sim \mathcal{N}(\mu_\Psi, \Sigma_\Psi), \\ \theta_{*,i} \mid \Psi_* &\sim \mathcal{N}(b_i^\top \Psi_*, \sigma_{0,i}^2), \\ Y_t \mid A_t, \theta_* &\sim \mathcal{N}(\theta_{*,A_t}, \sigma^2), \end{aligned} \quad \begin{aligned} \forall i \in [K], \\ \forall t \in [n]. \end{aligned} \quad (17)$$

B.2 Posteriors for Mixed-Effect Multi-Armed Bandit

Fix round $t \in [n]$, and recall that $S_{t,i}$ are the rounds where action i is taken up to round t . We introduce $N_{t,i} = |S_{t,i}|$ as the number of times that action i is taken up to round t and $B_{t,i} = \sum_{\ell \in S_{t,i}} Y_\ell$ as the total reward of action i up to round t . Note that $B_{t,i}$ and weight vectors b_i are unrelated. We derive in (18) and (19) the effect posterior Q_t and the action posteriors $P_{t,i}$ for the model in (17).

Proposition 3. For any round $t \in [n]$, the joint effect posterior is a multivariate Gaussian $Q_t = \mathcal{N}(\bar{\mu}_t, \bar{\Sigma}_t)$, where

$$\bar{\Sigma}_t^{-1} = \Sigma_\Psi^{-1} + \sum_{i \in [K]} \frac{N_{t,i}}{N_{t,i} \sigma_{0,i}^2 + \sigma^2} b_i b_i^\top, \quad \bar{\mu}_t = \bar{\Sigma}_t \left(\Sigma_\Psi^{-1} \mu_\Psi + \sum_{i \in [K]} \frac{B_{t,i}}{N_{t,i} \sigma_{0,i}^2 + \sigma^2} b_i \right). \quad (18)$$

Moreover, for any action $i \in [K]$, and effect parameters Ψ_t , the action posterior is a univariate Gaussian $P_{t,i}(\cdot \mid \Psi_t) = \mathcal{N}(\cdot; \tilde{\mu}_{t,i}, \tilde{\sigma}_{t,i}^2)$, where

$$\tilde{\sigma}_{t,i}^{-2} = \frac{1}{\sigma_{0,i}^2} + \frac{N_{t,i}}{\sigma^2}, \quad \tilde{\mu}_{t,i} = \tilde{\sigma}_{t,i}^2 \left(\frac{\Psi_t^\top b_i}{\sigma_{0,i}^2} + \frac{B_{t,i}}{\sigma^2} \right). \quad (19)$$

The effect posterior is additive in actions and can be interpreted as follows. Each action is a single noisy observation in its estimate. The *maximum likelihood estimate (MLE)* of the expected reward of action i , $B_{t,i}/N_{t,i}$, contributes to (18) proportionally to its precision, $N_{t,i}/(N_{t,i} \sigma_{0,i}^2 + \sigma^2)$. The contribution is weighted by b_i , which are the weights used to generate $\theta_{*,i}$. The action posterior in (19) has a standard form. Note that its prior mean depends on effect parameters Ψ_t , which are sampled.

C POSTERIOR DERIVATIONS FOR MIXED-EFFECT LINEAR BANDIT

Here we provide the derivations of the effect posterior and action posteriors for the setting introduced in Section 2.4 and summarized in (2). Precisely, we present the proof for Proposition 1 in Appendix C.1 and the proof of Proposition 2 in Appendix C.2.

C.1 Effect Posterior Derivation

Proof of Proposition 1 (derivation of Q_t). First, from basic properties of matrix operations in Appendix A.2, we have that $\sum_{\ell \in [L]} b_{i,\ell} \psi_{*,\ell} = \Gamma_i \Psi_*$ where $\Psi_* = (\psi_{*,\ell})_{\ell \in [L]} \in \mathbb{R}^{Ld}$ and $\Gamma_i = b_i^\top \otimes I_d$ (refer to Appendix D.1 for derivation detail). Thus our model can be written as

$$\begin{aligned} \Psi_* &\sim \mathcal{N}(\mu_\Psi, \Sigma_\Psi), \\ \theta_{*,i} \mid \Psi_* &\sim \mathcal{N}(\Gamma_i \Psi_*, \Sigma_{0,i}), & \forall i \in [K], \\ Y_\ell \mid X_\ell, \theta_{*,A_\ell} &\sim \mathcal{N}(X_\ell^\top \theta_{*,A_\ell}, \sigma^2), & \forall \ell \in [t]. \end{aligned} \quad (20)$$

It follows that the joint effect posterior in round t reads

$$Q_t(\Psi) \propto \mathbb{P}(H_t \mid \Psi_* = \Psi) Q_0(\Psi) \stackrel{(i)}{=} \prod_{i \in [K]} \mathbb{P}(H_{t,i} \mid \Psi_* = \Psi) Q_0(\Psi), \quad (21)$$

$$= \prod_{i \in [K]} \int_{\theta_i} \mathbb{P}(H_{t,i}, \theta_{*,i} = \theta_i \mid \Psi_* = \Psi) d\theta_i Q_0(\Psi) = \prod_{i \in [K]} \int_{\theta_i} \mathcal{L}_{t,i}(\theta_i) P_{0,i}(\theta_i \mid \Psi) d\theta_i Q_0(\Psi), \quad (22)$$

$$\propto \prod_{i \in [K]} \int_{\theta_i} \mathcal{L}_{t,i}(\theta_i) \mathcal{N}(\theta_i; \Gamma_i \Psi, \Sigma_{0,i}) d\theta_i \mathcal{N}(\Psi; \mu_\Psi, \Sigma_\Psi),$$

$$= \prod_{i \in [K]} \int_{\theta_i} \left(\prod_{\ell \in S_{t,i}} \mathcal{N}(Y_\ell; X_\ell^\top \theta_i, \sigma^2) \right) \mathcal{N}(\theta_i; \Gamma_i \Psi, \Sigma_{0,i}) d\theta_i \mathcal{N}(\Psi; \mu_\Psi, \Sigma_\Psi). \quad (23)$$

Here (i) follows from the fact that $\theta_{*,i}$ for $i \in [K]$ are conditionally independent given $\Psi_* = \Psi$ and that $H_{t,i}$ depends on Ψ_* only through $\theta_{*,i}$. Now we compute the quantity $\int_{\theta_i} \left(\prod_{\ell \in S_{t,i}} \mathcal{N}(Y_\ell; X_\ell^\top \theta_i, \sigma^2) \right) \mathcal{N}(\theta_i; \Gamma_i \Psi, \Sigma_{0,i}) d\theta_i$ using Lemma 1. Precisely, we obtain that it is proportional to $\mathcal{N}(\Psi; \bar{\mu}_{t,i}, \bar{\Sigma}_{t,i})$ where

$$\begin{aligned} \bar{\Sigma}_{t,i}^{-1} &= \Gamma_i^\top (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} \Gamma_i, \\ \bar{\mu}_{t,i} &= \bar{\Sigma}_{t,i} (\Gamma_i^\top (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} G_{t,i}^{-1} B_{t,i}), \end{aligned}$$

and

$$G_{t,i} = \sigma^{-2} \sum_{\ell \in S_{t,i}} X_\ell X_\ell^\top, \quad B_{t,i} = \sigma^{-2} \sum_{\ell \in S_{t,i}} Y_\ell X_\ell.$$

This means that the effect posterior Q_t is the product of $K + 1$ multivariate Gaussian distributions $\mathcal{N}(\mu_\Psi, \Sigma_\Psi), \mathcal{N}(\bar{\mu}_{t,1}, \bar{\Sigma}_{t,1}), \dots, \mathcal{N}(\bar{\mu}_{t,K}, \bar{\Sigma}_{t,K})$. Thus, the effect posterior Q_t is also a multivariate Gaussian distribution $\mathcal{N}(\bar{\mu}_t, \bar{\Sigma}_t^{-1})$, where

$$\begin{aligned} \bar{\Sigma}_t^{-1} &= \Sigma_\Psi^{-1} + \sum_{i=1}^K \bar{\Sigma}_{t,i}^{-1} = \Sigma_\Psi^{-1} + \sum_{i=1}^K \Gamma_i^\top (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} \Gamma_i, \\ \bar{\mu}_t &= \bar{\Sigma}_t \left(\Sigma_\Psi^{-1} \mu_\Psi + \sum_{i=1}^K \bar{\Sigma}_{t,i}^{-1} \bar{\mu}_{t,i} \right) = \bar{\Sigma}_t \left(\Sigma_\Psi^{-1} \mu_\Psi + \sum_{i=1}^K \Gamma_i^\top (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} G_{t,i}^{-1} B_{t,i} \right). \end{aligned}$$

Moreover, we use the properties in Appendix A.2 and that $\Gamma_i = b_i^\top \otimes I_d$ to rewrite the terms as

$$\begin{aligned} \Gamma_i^\top (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} \Gamma_i &= b_i b_i^\top \otimes (\Sigma_{0,i} + G_{t,i}^{-1})^{-1}, \\ \Gamma_i^\top (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} G_{t,i}^{-1} B_{t,i} &= b_i \otimes ((\Sigma_{0,i} + G_{t,i}^{-1})^{-1} G_{t,i}^{-1} B_{t,i}). \end{aligned}$$

This concludes the proof. \square

To reduce clutter, we fix an action $i \in [K]$ and a round $t \in [n]$ and drop subindexing by i and t in the following lemma. In summary, there exist $i \in [K]$ and $t \in [n]$ such that we have the following correspondences:

$$\Gamma \leftarrow \Gamma_i, \quad \Sigma_0 \leftarrow \Sigma_{0,i}, \quad N \leftarrow N_{t,i}, \quad \theta \leftarrow \theta_i, \quad (X_\ell, Y_\ell)_{\ell \in [N]} \leftarrow (X_\ell, Y_\ell)_{\ell \in \mathcal{S}_{t,i}},$$

Lemma 1 (Gaussian posterior update). *Let $\Gamma \in \mathbb{R}^{d \times Ld}$, $\Sigma_0 \in \mathbb{R}^{d \times d}$, and $\sigma > 0$ then we have that*

$$\int_{\theta} \left(\prod_{\ell=1}^N \mathcal{N}(Y_\ell; X_\ell^\top \theta, \sigma^2) \right) \mathcal{N}(\theta; \Gamma \Psi, \Sigma_0) d\theta \propto \mathcal{N}(\Psi; \mu_N, \Sigma_N).$$

where

$$\begin{aligned} \Sigma_N^{-1} &= \Gamma^\top (\Sigma_0 + G_N^{-1})^{-1} \Gamma, \\ \mu_N &= \Sigma_N (\Gamma^\top (\Sigma_0 + G_N^{-1})^{-1} G_N^{-1} B_N), \end{aligned}$$

and

$$G_N = \sigma^{-2} \sum_{k=1}^N X_k X_k^\top, \quad B_N = \sigma^{-2} \sum_{k=1}^N Y_k X_k.$$

Proof. Let $v = \sigma^{-2}$, $\Lambda_0 = \Sigma_0^{-1}$. We denote the integral in the lemma by $f(\Psi)$. It follows that

$$\begin{aligned} f(\Psi) &= \int_{\theta} \left(\prod_{\ell=1}^N \mathcal{N}(Y_\ell; X_\ell^\top \theta, \sigma^2) \right) \mathcal{N}(\theta; \Gamma \Psi, \Sigma_0) d\theta, \\ &\propto \int_{\theta} \exp \left[-\frac{1}{2} v \sum_{\ell=1}^N (Y_\ell - X_\ell^\top \theta)^2 - \frac{1}{2} (\theta - \Gamma \Psi)^\top \Lambda_0 (\theta - \Gamma \Psi) \right] d\theta, \\ &= \int_{\theta} \exp \left[-\frac{1}{2} \left(v \sum_{\ell=1}^N (Y_\ell^2 - 2Y_\ell \theta^\top X_\ell + (\theta^\top X_\ell)^2) + \theta^\top \Lambda_0 \theta - 2\theta^\top \Lambda_0 \Gamma \Psi + (\Gamma \Psi)^\top \Lambda_0 (\Gamma \Psi) \right) \right] d\theta, \\ &\propto \int_{\theta} \exp \left[-\frac{1}{2} \left(\theta^\top \left(v \sum_{\ell=1}^N X_\ell X_\ell^\top + \Lambda_0 \right) \theta - 2\theta^\top \left(v \sum_{\ell=1}^N Y_\ell X_\ell + \Lambda_0 \Gamma \Psi \right) + (\Gamma \Psi)^\top \Lambda_0 (\Gamma \Psi) \right) \right] d\theta. \end{aligned}$$

To reduce clutter, let

$$\begin{aligned} G_N &= v \sum_{\ell=1}^N X_\ell X_\ell^\top, \quad V_N = (G_N + \Lambda_0)^{-1}, \quad U_N = V_N^{-1}, \\ B_N &= v \sum_{\ell=1}^N Y_\ell X_\ell \quad \text{and} \quad \beta_N = V_N (B_N + \Lambda_0 \Gamma \Psi). \end{aligned}$$

We have that $U_N V_N = V_N U_N = I_d$, and thus

$$\begin{aligned} f(\Psi) &\propto \int_{\theta} \exp \left[-\frac{1}{2} (\theta^\top U_N \theta - 2\theta^\top U_N V_N (B_N + \Lambda_0 \Gamma \Psi) + (\Gamma \Psi)^\top \Lambda_0 (\Gamma \Psi)) \right] d\theta, \\ &= \int_{\theta} \exp \left[-\frac{1}{2} (\theta^\top U_N \theta - 2\theta^\top U_N \beta_N + (\Gamma \Psi)^\top \Lambda_0 (\Gamma \Psi)) \right] d\theta, \\ &= \int_{\theta} \exp \left[-\frac{1}{2} ((\theta - \beta_N)^\top U_N (\theta - \beta_N) - \beta_N^\top U_N \beta_N + (\Gamma \Psi)^\top \Lambda_0 (\Gamma \Psi)) \right] d\theta, \\ &\propto \exp \left[-\frac{1}{2} (-\beta_N^\top U_N \beta_N + (\Gamma \Psi)^\top \Lambda_0 (\Gamma \Psi)) \right], \\ &= \exp \left[-\frac{1}{2} \left(-(B_N + \Lambda_0 \Gamma \Psi)^\top V_N (B_N + \Lambda_0 \Gamma \Psi) + (\Gamma \Psi)^\top \Lambda_0 (\Gamma \Psi) \right) \right], \\ &\propto \exp \left[-\frac{1}{2} (\Psi^\top \Gamma^\top (\Lambda_0 - \Lambda_0 V_N \Lambda_0) \Gamma \Psi - 2\Psi^\top (\Gamma^\top \Lambda_0 V_N B_N)) \right], \\ &= \exp \left[-\frac{1}{2} \Psi^\top \Sigma_N^{-1} \Psi + \Psi^\top \Sigma_N^{-1} \mu_N \right], \end{aligned}$$

where

$$\begin{aligned}\Sigma_N^{-1} &= \Gamma^\top (\Lambda_0 - \Lambda_0 V_N \Lambda_0) \Gamma = \Gamma^\top (\Lambda_0^{-1} + G_N^{-1})^{-1} \Gamma, \\ \Sigma_N^{-1} \mu_N &= (\Gamma^\top \Lambda_0 V_N B_N) = \Gamma^\top (\Sigma_0 + G_N^{-1})^{-1} G_N^{-1} B_N.\end{aligned}\tag{24}$$

We use the Woodbury matrix identity to get the second equalities which concludes the proof. \square

C.2 Action Posterior Derivation for Mixed-Effect Linear Bandit

Proof of Proposition 2 (Derivation of $P_{t,i}$). This proposition is a direct application Lemma 2; in which case we get that the posterior $P_{t,i}$ is a multivariate Gaussian distribution $\mathcal{N}(\tilde{\mu}_{t,i}, \tilde{\Sigma}_{t,i})$, where

$$\begin{aligned}\tilde{\Sigma}_{t,i}^{-1} &= G_{t,i} + \Sigma_{0,i}^{-1}, \\ \tilde{\mu}_{t,i} &= \tilde{\Sigma}_{t,i} \left(B_{t,i} + \Sigma_{0,i}^{-1} \sum_{\ell=1}^L b_{i,\ell} \psi_{t,\ell} \right).\end{aligned}$$

\square

To reduce clutter, we consider a fixed action $i \in [K]$ and round $t \in [n]$, and drop subindexing by t and i in Lemma 2. In summary, there exist $i \in [K]$ and $t \in [n]$ such that we have the following correspondences:

$$b_\ell \leftarrow b_{i,\ell}, \quad \Sigma_0 \leftarrow \Sigma_{0,i}, \quad N \leftarrow N_{t,i}, \quad \theta_* \leftarrow \theta_{*,i}, \quad (X_\ell, Y_\ell)_{\ell \in [N]} \leftarrow (X_\ell, Y_\ell)_{\ell \in S_{t,i}}.$$

Lemma 2. Consider the following model

$$\begin{aligned}\theta_* | \Psi_* &\sim \mathcal{N} \left(\sum_{\ell=1}^L b_\ell \psi_{*,\ell}, \Sigma_0 \right), \\ Y_\ell | X_\ell, \theta_* &\sim \mathcal{N} (X_\ell^\top \theta_*, \sigma^2), \quad \forall \ell \in [N].\end{aligned}$$

Let $H = \{X_1, Y_1, \dots, X_N, Y_N\}$ then we have that $\mathbb{P}(\theta_* = \theta | \Psi_* = \Psi, H) = \mathcal{N}(\theta; \tilde{\mu}_N, \tilde{\Sigma}_N)$, where

$$\begin{aligned}\tilde{\Sigma}_N^{-1} &= \sigma^{-2} \sum_{\ell=1}^N X_\ell X_\ell^\top + \Sigma_0^{-1}, \\ \tilde{\mu}_N &= \tilde{\Sigma}_N \left(\sigma^{-2} \sum_{\ell=1}^N X_\ell Y_\ell + \Sigma_0^{-1} \sum_{\ell=1}^L b_\ell \psi_\ell \right).\end{aligned}$$

Proof. Let $v = \sigma^{-2}$, $\Lambda_0 = \Sigma_0^{-1}$. Then the action posterior decomposes as

$$\begin{aligned}&\mathbb{P}(\theta_* = \theta | \Psi_* = \Psi, H), \\ &\propto \mathbb{P}(H | \Psi_* = \Psi, \theta_* = \theta) \mathbb{P}(\theta_* = \theta | \Psi_* = \Psi), \\ &= \mathbb{P}(H | \theta_* = \theta) \mathbb{P}(\theta_* = \theta | \Psi_* = \Psi), \quad (H \text{ depends on } \Psi_* \text{ only through } \theta_*), \\ &= \prod_{\ell=1}^N \mathcal{N}(Y_\ell; X_\ell^\top \theta, \sigma^2) \mathcal{N}(\theta; \sum_{\ell=1}^L b_\ell \psi_\ell, \Sigma_0), \\ &= \exp \left[-\frac{1}{2} \left(v \sum_{\ell=1}^N (Y_\ell^2 - 2Y_\ell X_\ell^\top \theta + (X_\ell^\top \theta)^2) + \theta^\top \Lambda_0 \theta - 2\theta^\top \Lambda_0 \sum_{\ell=1}^L b_\ell \psi_\ell + \left(\sum_{\ell=1}^L b_\ell \psi_\ell \right)^\top \Lambda_0 \left(\sum_{\ell=1}^L b_\ell \psi_\ell \right) \right) \right], \\ &\propto \exp \left[-\frac{1}{2} \left(\theta^\top \left(v \sum_{\ell=1}^N X_\ell X_\ell^\top + \Lambda_0 \right) \theta - 2\theta^\top \left(v \sum_{\ell=1}^N X_\ell Y_\ell + \Lambda_0 \sum_{\ell=1}^L b_\ell \psi_\ell \right) \right) \right], \\ &\propto \mathcal{N} \left(\theta; \tilde{\mu}_N, \left(\tilde{\Lambda}_N \right)^{-1} \right),\end{aligned}$$

where $\tilde{\Lambda}_N = v \sum_{\ell=1}^N X_\ell X_\ell^\top + \Lambda_0$, and $\tilde{\Lambda}_N \tilde{\mu}_N = v \sum_{\ell=1}^N X_\ell Y_\ell + \Lambda_0 \sum_{\ell=1}^L b_\ell \psi_\ell$. \square

D REGRET PROOFS

In this section, we prove a more general version of Theorem 1. First, we provide a compact formulation of our problem in Appendix D.1. Next, we use the total covariance decomposition to derive the covariance of $\mathbb{P}(\theta_{*,i} = \theta | H_t)$ in Appendix D.2. Finally, we provide preliminary eigenvalue results in Appendix D.3 to proceed with the regret proof in Appendix D.4.

D.1 Problem Reformulation for Regret Analysis

Here, we aim at rewriting our model in a compact form to simplify regret analysis. We first introduce K i.i.d. multivariate Gaussian variables $Z_i \sim \mathcal{N}(0, \Sigma_{0,i})$ for $i \in [K]$, and the following matrix

$$\Psi_{\text{mat},*} = [\psi_{*,1}, \dots, \psi_{*,L}] \in \mathbb{R}^{d \times L}.$$

First, we have that $\text{Vec}(\Psi_{\text{mat},*}) = \Psi_*$ where Ψ_* is defined in Section 2. Moreover notice that $\sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell} = \Psi_{\text{mat},*} b_i$ and thus given matrix $\Psi_{\text{mat},*}$ we have that

$$\theta_{*,i} = \Psi_{\text{mat},*} b_i + Z_i, \quad \forall i \in [K]. \quad (25)$$

We vectorize (25) to obtain

$$\theta_{*,i} = \text{Vec}(\theta_{*,i}) = \text{Vec}(\Psi_{\text{mat},*} b_i + Z_i) = \text{Vec}(\Psi_{\text{mat},*} b_i) + Z_i, \quad (26)$$

where we used that if $X \in \mathbb{R}^d$ (a column vector), then $X = \text{Vec}(X)$ and that $\text{Vec}(\cdot)$ is a linear transformation. Also, we know from (c) in Appendix A.2 that $\text{Vec}(AB) = (B^\top \otimes I_n) \text{Vec}(A)$ for any $A \in \mathbb{R}^{n \times m}$, $B \in \mathbb{R}^{m \times p}$. Therefore,

$$\theta_{*,i} = \Gamma_i \Psi_* + Z_i, \quad (27)$$

where $\Gamma_i = b_i^\top \otimes I_d$ and we used that $\text{Vec}(\Psi_{\text{mat},*}) = \Psi_*$. It follows that

$$\theta_{*,i} | \Psi_* \sim \mathcal{N}(\Gamma_i \Psi_*, \Sigma_{0,i}), \quad (28)$$

This allows us to rewrite our model as a single-parent hierarchical model

$$\begin{aligned} \Psi_* &\sim \mathcal{N}(\mu_\Psi, \Sigma_\Psi), \\ \theta_{*,i} | \Psi_* &\sim \mathcal{N}(\Gamma_i \Psi_*, \Sigma_{0,i}), \\ Y_t | X_t, \theta_{*,A_t} &\sim \mathcal{N}(X_t^\top \theta_{*,A_t}, \sigma^2), \end{aligned} \quad \begin{aligned} & \\ \forall i \in [K], \\ \forall t \in [n]. \end{aligned} \quad (29)$$

D.2 Derivation of $\mathbb{P}(\theta_{*,i} = \theta | H_t)$

Let

$$G_{t,i} = \sigma^{-2} \sum_{\ell \in S_{t,i}} X_\ell X_\ell^\top, \quad B_{t,i} = \sigma^{-2} \sum_{\ell \in S_{t,i}} Y_\ell X_\ell.$$

Lemma 3 (Covariance of $\mathbb{P}(\theta_{*,i} = \theta | H_t)$). *Consider the model in (2) and let $\Psi_* | H_t \sim \mathcal{N}(\bar{\mu}_t, \bar{\Sigma}_t)$, then we have*

$$\hat{\Sigma}_{t,i} = \text{cov}[\theta_{*,i} | H_t] = \tilde{\Sigma}_{t,i} + \tilde{\Sigma}_{t,i} \Sigma_{0,i}^{-1} \Gamma_i \bar{\Sigma}_t \Gamma_i^\top \Sigma_{0,i}^{-1} \tilde{\Sigma}_{t,i}, \quad \forall i \in [K].$$

where $\tilde{\Sigma}_{t,i} = \text{cov}[\theta_{*,i} | \Psi_*, H_t] = (G_{t,i} + \Sigma_{0,i}^{-1})^{-1}$.

Proof. Proposition 2 and the fact that $\sum_{\ell \in [L]} b_{i,\ell} \psi_{*,\ell} = \Gamma_i \Psi_*$ where $\Gamma_i = b_i^\top \otimes I_d$ (Appendix D.1) yield

$$\begin{aligned} \text{cov}[\theta_{*,i} | \Psi_*, H_t] &= (G_{t,i} + \Lambda_{0,i})^{-1} \\ \mathbb{E}[\theta_{*,i} | \Psi_*, H_t] &= \text{cov}[\theta_{*,i} | \Psi_*, H_t] (B_{t,i} + \Lambda_{0,i} \Gamma_i \Psi_*) \end{aligned}$$

First, given H_t , $\text{cov}[\theta_{*,i} | \Psi_*, H_t] = (G_{t,i} + \Lambda_{0,i})^{-1}$ is constant (does not depend on Ψ_*). Thus

$$\mathbb{E}[\text{cov}[\theta_{*,i} | \Psi_*, H_t] | H_t] = \text{cov}[\theta_{*,i} | \Psi_*, H_t] = (G_{t,i} + \Lambda_{0,i})^{-1}.$$

In addition, given H_t , both $(G_{t,i} + \Lambda_{0,i})^{-1}$ and $B_{t,i}$ are constant. Thus

$$\begin{aligned} \text{cov} [\mathbb{E} [\theta_{*,i} | \Psi_*, H_t] | H_t] &= \text{cov} [\text{cov} [\theta_{*,i} | \Psi_*, H_t] \Lambda_{0,i} \Gamma_i \Psi_* | H_t] \\ &= (G_{t,i} + \Lambda_{0,i})^{-1} \Lambda_{0,i} \Gamma_i \text{cov} [\Psi_* | H_t] \Gamma_i^\top \Lambda_{0,i} (G_{t,i} + \Lambda_{0,i})^{-1} \\ &= (G_{t,i} + \Lambda_{0,i})^{-1} \Lambda_{0,i} \Gamma_i \bar{\Sigma}_t \Gamma_i^\top \Lambda_{0,i} (G_{t,i} + \Lambda_{0,i})^{-1}. \end{aligned}$$

Finally, total covariance decomposition (Weiss, 2005) concludes the proof. \square

D.3 Preliminary Eigenvalues Results

Next we present some preliminary upper bounds on the maximum eigenvalues of our covariance matrices.

- **Definitions:** Let $\lambda_{1,0} = \max_{i \in [K]} \lambda_1(\Sigma_{0,i})$, $\lambda_{d,0} = \min_{i \in [K]} \lambda_d(\Sigma_{0,i})$, $\lambda_{1,\Psi} = \lambda_1(\Sigma_\Psi)$, and $\kappa_b = \max_{i \in [K]} \|b_i\|_2^2$.
- **upper bound of $\lambda_1(\Gamma_i \Gamma_i^\top)$:**

$$\lambda_1(\Gamma_i \Gamma_i^\top) \leq \kappa_b, \quad \forall i \in [K]. \quad (30)$$

Similarly, we have that

$$\lambda_1(\Gamma_i^\top \Gamma_i) \leq \kappa_b, \quad \forall i \in [K]. \quad (31)$$

- **upper bound of $\lambda_1(\hat{\Sigma}_{t,i})$:**

$$\lambda_1(\hat{\Sigma}_{t,i}) \leq \lambda_{1,0} + \frac{\lambda_{1,0}^2 \lambda_{1,\Psi} \kappa_b}{\lambda_{d,0}^2}, \quad \forall i \in [K]. \quad (32)$$

- **upper bound of $\lambda_1(\bar{\Sigma}_{n+1}^{-1})$:**

$$\lambda_1(\bar{\Sigma}_{n+1}^{-1}) \leq \frac{1}{\lambda_{1,\Psi}} + K \kappa_b \left(\frac{1}{\lambda_{d,0}} - \frac{1}{\lambda_{1,0}^2 \left(\frac{\kappa_b n}{\sigma^2} + \frac{1}{\lambda_{d,0}} \right)} \right). \quad (33)$$

Proof. We start with (30), we obtain the last inequality as follows. First, recall that $\Gamma_i = b_i^\top \otimes I_d$ for any $i \in [K]$. Thus $\Gamma_i \Gamma_i^\top = (b_i^\top \otimes I_d)(b_i \otimes I_d) = \|b_i\|_2^2 I_d$ for any $i \in [K]$. Then $\lambda_1(\Gamma_i \Gamma_i^\top) = \|b_i\|_2^2 \leq \kappa_b$. The second result follows from the fact that $\lambda_1(\Gamma_i \Gamma_i^\top) = \lambda_1(\Gamma_i^\top \Gamma_i)$.

Now we prove the result in (32). This follows from the expression of $\hat{\Sigma}_{t,i}$ in Lemma 3. Precisely, we have that

$$\hat{\Sigma}_{t,i} = \tilde{\Sigma}_{t,i} + \tilde{\Sigma}_{t,i} \Sigma_{0,i}^{-1} \Gamma_i \bar{\Sigma}_t \Gamma_i^\top \Sigma_{0,i}^{-1} \tilde{\Sigma}_{t,i}, \quad \forall i \in [K].$$

where $\tilde{\Sigma}_{t,i} = (G_{t,i} + \Sigma_{0,i}^{-1})^{-1}$. Thus Weyl's inequality combined with the properties in Appendix A.2 yields that

$$\lambda_1(\hat{\Sigma}_{t,i}) \leq \lambda_1(\tilde{\Sigma}_{t,i}) + \lambda_1(\tilde{\Sigma}_{t,i}) \lambda_1(\Sigma_{0,i}^{-1}) \lambda_1(\Gamma_i \bar{\Sigma}_t \Gamma_i^\top) \lambda_1(\Sigma_{0,i}^{-1}) \lambda_1(\tilde{\Sigma}_{t,i}) \leq \lambda_{1,0} + \frac{\lambda_{1,0}^2 \lambda_{1,\Psi} \kappa_b}{\lambda_{d,0}^2}$$

In the last inequality, we used that $\lambda_1(\Gamma_i \bar{\Sigma}_t \Gamma_i^\top) \leq \lambda_1(\bar{\Sigma}_t) \lambda(\Gamma_i \Gamma_i^\top)$, ((f) in Appendix A.2), $\lambda_1(\Sigma_{0,i}^{-1}) \leq \frac{1}{\lambda_{d,0}}$, and $\lambda_1(\tilde{\Sigma}_{t,i}) \leq \lambda_{1,0}$.

Finally, we prove the result in (33). First, we rewrite the precision matrix of the effect posterior $\bar{\Sigma}_t^{-1}$ using the compact notation introduced in Appendix D.1. Precisely, it follows from (24) that

$$\begin{aligned} \bar{\Sigma}_t^{-1} &\stackrel{(i)}{=} \Sigma_\Psi^{-1} + \sum_{i=1}^K \Gamma_i^\top (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} \Gamma_i \stackrel{(ii)}{=} \Sigma_\Psi^{-1} + \sum_{i=1}^K \Gamma_i^\top (\Sigma_{0,i}^{-1} - \Sigma_{0,i}^{-1} (G_{t,i} + \Sigma_{0,i}^{-1})^{-1} \Sigma_{0,i}^{-1}) \Gamma_i, \\ &\stackrel{(iii)}{=} \Sigma_\Psi^{-1} + \sum_{i=1}^K \Gamma_i^\top (\Sigma_{0,i}^{-1} - \Sigma_{0,i}^{-1} \tilde{\Sigma}_{t,i} \Sigma_{0,i}^{-1}) \Gamma_i. \end{aligned}$$

The equality (i) requires $G_{t,i}$ to be invertible and was only given in the main manuscript to ease the exposition. Note that (i) is obtained by applying the Woodbury matrix identity on (ii). In our proof, we use (ii) and (iii) which are the same; (iii) follows from plugging $\tilde{\Sigma}_{t,i} = (G_{t,i} + \Sigma_{0,i}^{-1})^{-1}$ in (ii). Then we have that

$$\begin{aligned}
 \lambda_1(\bar{\Sigma}_{n+1}^{-1}) &= \lambda_1\left(\Sigma_{\Psi}^{-1} + \sum_{i=1}^K \Gamma_i^{\top} \left(\Sigma_{0,i}^{-1} - \Sigma_{0,i}^{-1} \tilde{\Sigma}_{n+1,i} \Sigma_{0,i}^{-1}\right) \Gamma_i\right) \leq \frac{1}{\lambda_{1,\Psi}} + \sum_{i=1}^K \lambda_1\left(\Gamma_i^{\top} \left(\Sigma_{0,i}^{-1} - \Sigma_{0,i}^{-1} \tilde{\Sigma}_{n+1,i} \Sigma_{0,i}^{-1}\right) \Gamma_i\right), \\
 &\leq \frac{1}{\lambda_{1,\Psi}} + \sum_{i=1}^K \lambda_1(\Gamma_i^{\top} \Gamma_i) \lambda_1\left(\Sigma_{0,i}^{-1} - \Sigma_{0,i}^{-1} \tilde{\Sigma}_{n+1,i} \Sigma_{0,i}^{-1}\right) \leq \frac{1}{\lambda_{1,\Psi}} + \sum_{i=1}^K \kappa_b \left(\lambda_1(\Sigma_{0,i}^{-1}) + \lambda_1\left(-\Sigma_{0,i}^{-1} \tilde{\Sigma}_{n+1,i} \Sigma_{0,i}^{-1}\right)\right), \\
 &\leq \frac{1}{\lambda_{1,\Psi}} + \sum_{i=1}^K \kappa_b \left(\frac{1}{\lambda_{d,0}} - \lambda_d\left(\Sigma_{0,i}^{-1} \tilde{\Sigma}_{n+1,i} \Sigma_{0,i}^{-1}\right)\right) \leq \frac{1}{\lambda_{1,\Psi}} + \sum_{i=1}^K \kappa_b \left(\frac{1}{\lambda_{d,0}} - \lambda_d\left(\Sigma_{0,i}^{-1}\right) \lambda_d\left(\tilde{\Sigma}_{n+1,i}\right) \lambda_d\left(\Sigma_{0,i}^{-1}\right)\right), \\
 &\leq \frac{1}{\lambda_{1,\Psi}} + \sum_{i=1}^K \kappa_b \left(\frac{1}{\lambda_{d,0}} - \frac{1}{\lambda_{1,0}^2} \lambda_d\left(\tilde{\Sigma}_{n+1,i}\right)\right) = \frac{1}{\lambda_{1,\Psi}} + \sum_{i=1}^K \kappa_b \left(\frac{1}{\lambda_{d,0}} - \frac{1}{\lambda_{1,0}^2 \lambda_1\left(G_{n+1,i} + \Sigma_{0,i}^{-1}\right)}\right), \\
 &\leq \frac{1}{\lambda_{1,\Psi}} + \sum_{i=1}^K \kappa_b \left(\frac{1}{\lambda_{d,0}} - \frac{1}{\lambda_{1,0}^2 \left(\frac{\kappa_x n}{\sigma^2} + \frac{1}{\lambda_{d,0}}\right)}\right) = \frac{1}{\lambda_{1,\Psi}} + K \kappa_b \left(\frac{1}{\lambda_{d,0}} - \frac{1}{\lambda_{1,0}^2 \left(\frac{\kappa_x n}{\sigma^2} + \frac{1}{\lambda_{d,0}}\right)}\right).
 \end{aligned}$$

This concludes the proof. \square

D.4 Regret Proof

Here we prove a more general version of Theorem 1 where we do not assume that the covariance matrices $\Sigma_{0,i}$ and Σ_{Ψ} are diagonal. We still assume that there exists $\kappa_x > 0$ such that $\|X_t\|_2^2 \leq \kappa_x$ for any $t \in [n]$.

Theorem 2 (General version of Theorem 1). *For any $\delta \in (0, 1)$, the Bayes regret of meTS in the mixed-effect model in Section 2.4 is bounded as*

$$\mathcal{BR}(n) \leq \sqrt{2n(\mathcal{R}^A(n) + \mathcal{R}^E(n)) \log(1/\delta)} + cn\delta, \quad (34)$$

with $c = \sqrt{\frac{2}{\pi} \kappa_x \left(\lambda_{1,0} + \frac{\lambda_{1,0}^2 \lambda_{1,\Psi} \kappa_b}{\lambda_{d,0}^2}\right) K}$, $\kappa_b = \max_{i \in [K]} \|b_i\|_2^2$, $\lambda_{1,0} = \max_{i \in [K]} \lambda_1(\Sigma_{0,i})$, $\lambda_{d,0} = \min_{i \in [K]} \lambda_d(\Sigma_{0,i})$, $\lambda_{1,\Psi} = \lambda_1(\Sigma_{\Psi})$ and

$$\begin{aligned}
 \mathcal{R}^A(n) &= dKc_A \log\left(1 + \frac{n\kappa_x \lambda_{1,0}}{\sigma^2 d}\right), \quad c_A = \frac{\kappa_x \lambda_{1,0}}{\log\left(1 + \frac{\kappa_x \lambda_{1,0}}{\sigma^2}\right)}, \\
 \mathcal{R}^E(n) &= dLc_E \log\left(1 + K\kappa_b \lambda_{1,\Psi} \left(\frac{1}{\lambda_{d,0}} - \frac{1}{\lambda_{1,0}^2 \left(\frac{\kappa_x n}{\sigma^2} + \frac{1}{\lambda_{d,0}}\right)}\right)\right), \quad c_E = \frac{\kappa_x \kappa_b \lambda_{1,0}^2 \lambda_{1,\Psi} \left(1 + \frac{\kappa_x \lambda_{1,0}}{\sigma^2}\right)}{\lambda_{d,0}^2 \log\left(1 + \frac{\kappa_x \kappa_b \lambda_{1,0}^2 \lambda_{1,\Psi}}{\sigma^2 \lambda_{d,0}^2}\right)}.
 \end{aligned}$$

In particular, the result in Theorem 1 is retrieved when $\lambda_{1,0} = \lambda_{d,0} = \sigma_0^2$, and $\lambda_{\Psi} = \sigma_{\Psi}^2$.

Proof. Consider our model rewritten in (29). As we explained in Section 4.3, the posterior distribution of the action parameter $\theta_{*,i} | H_t$ is a multivariate Gaussian distribution $\mathcal{N}(\hat{\mu}_{t,i}, \hat{\Sigma}_{t,i})$ for some $\hat{\mu}_{t,i} \in \mathbb{R}^d$ and $\hat{\Sigma}_{t,i} \in \mathbb{R}^{d \times d}$ (Lemma 3). Now we let $\theta_{t,*} = (X_t^{\top} \theta_{*,i})_{i \in [K]} \in \mathbb{R}^K$ be the concatenation of the expected rewards of actions in round t . Notice that the context X_t is known in round t , and thus we include it in the history H_t . Then, the joint posterior of the expected rewards, $\theta_{t,*} | H_t$, is also a multivariate Gaussian $\mathcal{N}(\check{\theta}_t, \check{\Sigma}_t)$ for $\check{\theta}_t = (X_t^{\top} \hat{\mu}_{t,i})_{i \in [K]} \in \mathbb{R}^K$ and $\check{\Sigma}_t \in \mathbb{R}^{K \times K}$. This follows from the properties of Gaussian distributions (Koller and Friedman, 2009) and the fact that X_t is now included in H_t . Let $\mathbf{A}_t \in \{0, 1\}^K$ and $\mathbf{A}_{t,*} \in \{0, 1\}^K$ be indicator vectors of the taken action A_t and optimal action $A_{t,*}$, respectively (The vector representations are in bold letters while the integer representations are in regular letters). Then the Bayes regret can be rewritten and consequently decomposed following standard analysis (Russo and Van Roy, 2014) as

$$\begin{aligned}
 \mathcal{BR}(n) &= \mathbb{E} \left[\sum_{t=1}^n X_t^{\top} \theta_{*,A_t,*} - X_t^{\top} \theta_{*,A_t} \right], \\
 &= \mathbb{E} \left[\sum_{t=1}^n \mathbf{A}_{t,*}^{\top} \theta_{t,*} - \mathbf{A}_t^{\top} \theta_{t,*} \right] = \mathbb{E} \left[\mathbb{E} \left[\mathbf{A}_{t,*}^{\top} (\theta_{t,*} - \check{\theta}_t) | H_t \right] \right] + \mathbb{E} \left[\mathbb{E} \left[\mathbf{A}_t^{\top} (\check{\theta}_t - \theta_{t,*}) | H_t \right] \right].
 \end{aligned} \quad (35)$$

This follows from the fact that $\check{\theta}_t = (X_t^\top \hat{\mu}_{t,i})_{i \in [K]}$ is deterministic given H_t (since H_t now includes X_t), and that $\mathbf{A}_{t,*}$ and \mathbf{A}_t are i.i.d. given H_t . Moreover, given H_t , $\check{\theta}_t - \theta_{t,*}$ is a zero-mean multivariate random variable independent of \mathbf{A}_t and thus $\mathbb{E}[\mathbf{A}_t^\top (\check{\theta}_t - \theta_{t,*}) \mid H_t] = 0$. Therefore, we only need to bound the first term in (35). With slight abuse of notation, let \mathcal{A} be the set of all possible indicator vectors of actions $a \in [K]$. Precisely, an action $a \in [K]$ is also represented by an indicator vector $\mathbf{a} \in \mathcal{A} \subset \{0, 1\}^K$ (in bold letter). Then we define the following events

$$E_{t,\mathbf{a}}(\delta) = \left\{ |\mathbf{a}^\top (\theta_{t,*} - \check{\theta}_t)| \leq \sqrt{2 \log(1/\delta)} \|\mathbf{a}\|_{\check{\Sigma}_t} \right\}, \quad \forall \delta \in (0, 1), \forall \mathbf{a} \in \mathcal{A}.$$

Fix history H_t , we split the expectation over the two complementary events $E_{t,\mathbf{A}_{t,*}}(\delta)$ and $\bar{E}_{t,\mathbf{A}_{t,*}}(\delta)$, and use the Cauchy-Schwarz inequality to obtain

$$\mathbb{E}[\mathbf{A}_{t,*}^\top (\theta_{t,*} - \check{\theta}_t) \mid H_t] \leq \sqrt{2 \log(1/\delta)} \mathbb{E}[\|\mathbf{A}_{t,*}\|_{\check{\Sigma}_t} \mid H_t] + \mathbb{E}[\mathbf{A}_{t,*}^\top (\theta_{t,*} - \check{\theta}_t) \mathbb{1}\{\bar{E}_{t,\mathbf{A}_{t,*}}(\delta)\} \mid H_t]. \quad (36)$$

Now the second term in (36) can be bounded as follows. For any $\mathbf{a} \in \mathcal{A}$, let $Z_{\mathbf{a}} = \mathbf{a}^\top (\theta_{t,*} - \check{\theta}_t)$. Then we have that

$$\begin{aligned} \mathbb{E}[\mathbf{A}_{t,*}^\top (\theta_{t,*} - \check{\theta}_t) \mathbb{1}\{\bar{E}_{t,\mathbf{A}_{t,*}}(\delta)\} \mid H_t] &\stackrel{(i)}{=} \mathbb{E}\left[Z_{\mathbf{A}_{t,*}} \mathbb{1}\{|Z_{\mathbf{A}_{t,*}}| > \sqrt{2 \log(1/\delta)} \|\mathbf{A}_{t,*}\|_{\check{\Sigma}_t}\} \mid H_t\right], \\ &\stackrel{(ii)}{\leq} \mathbb{E}\left[|Z_{\mathbf{A}_{t,*}}| \mathbb{1}\{|Z_{\mathbf{A}_{t,*}}| > \sqrt{2 \log(1/\delta)} \|\mathbf{A}_{t,*}\|_{\check{\Sigma}_t}\} \mid H_t\right], \\ &\stackrel{(iii)}{\leq} \sum_{\mathbf{a} \in \mathcal{A}} \mathbb{E}\left[|Z_{\mathbf{a}}| \mathbb{1}\{|Z_{\mathbf{a}}| > \sqrt{2 \log(1/\delta)} \|\mathbf{a}\|_{\check{\Sigma}_t}\} \mid H_t\right], \\ &\stackrel{(iv)}{\leq} \sum_{\mathbf{a} \in \mathcal{A}} \frac{2}{\|\mathbf{a}\|_{\check{\Sigma}_t} \sqrt{2\pi}} \int_{u=\sqrt{2 \log(1/\delta)} \|\mathbf{a}\|_{\check{\Sigma}_t}}^{\infty} u \exp\left[-\frac{u^2}{2\|\mathbf{a}\|_{\check{\Sigma}_t}^2}\right] du, \\ &\stackrel{(v)}{\leq} \sum_{\mathbf{a} \in \mathcal{A}} \|\mathbf{a}\|_{\check{\Sigma}_t} \frac{2}{\sqrt{2\pi}} \int_{u=\sqrt{2 \log(1/\delta)}}^{\infty} u \exp\left[-\frac{u^2}{2}\right] du \stackrel{(vi)}{\leq} \sqrt{\frac{2}{\pi}} \lambda_{\max,t} K \delta. \end{aligned} \quad (37)$$

In (i), we simply rewrite the terms using the random variable $Z_{\mathbf{A}_{t,*}}$. In (ii), we use the fact that $Z_{\mathbf{A}_{t,*}} \leq |Z_{\mathbf{A}_{t,*}}|$. In (iii), we upper bound the expectation over $|Z_{\mathbf{A}_{t,*}}| \mathbb{1}\{|Z_{\mathbf{A}_{t,*}}| > \sqrt{2 \log(1/\delta)} \|\mathbf{A}_{t,*}\|_{\check{\Sigma}_t}\}$ with the sum of the expectations over $|Z_{\mathbf{a}}| \mathbb{1}\{|Z_{\mathbf{a}}| > \sqrt{2 \log(1/\delta)} \|\mathbf{a}\|_{\check{\Sigma}_t}\}$ for $\mathbf{a} \in \mathcal{A}$ since all these random variables are non-negative. Moreover, (iv) follows from the facts that given H_t , $Z_{\mathbf{a}} \sim \mathcal{N}(0, \|\mathbf{a}\|_{\check{\Sigma}_t}^2)$, and that if $Z \sim \mathcal{N}(0, \sigma^2)$, then for any $\epsilon \geq 0$, $\mathbb{P}(|Z| > \epsilon) \leq 2\mathbb{P}(Z > \epsilon)$. In (v), we use the change of variables $u \leftarrow u/\|\mathbf{a}\|_{\check{\Sigma}_t}$. Finally, in (vi), we compute the integral and set $\lambda_{\max,t} = \max_{\mathbf{a} \in \mathcal{A}} \|\mathbf{a}\|_{\check{\Sigma}_t}$. We combine (36) and (37) with the fact that \mathbf{A}_t and $\mathbf{A}_{t,*}$ are i.i.d. given H_t to obtain that

$$\mathbb{E}[\mathbf{A}_{t,*}^\top (\theta_{t,*} - \check{\theta}_t) \mid H_t] \leq \sqrt{2 \log(1/\delta)} \mathbb{E}[\|\mathbf{A}_t\|_{\check{\Sigma}_t} \mid H_t] + \sqrt{\frac{2}{\pi}} \lambda_{\max,t} K \delta. \quad (38)$$

The bound in (38) holds for any history H_t and thus we take an additional expectation and get that

$$\begin{aligned} \mathcal{BR}(n) &= \mathbb{E}\left[\sum_{t=1}^n \mathbf{A}_{t,*}^\top \theta_{t,*} - \mathbf{A}_t^\top \theta_{t,*}\right] \leq \sqrt{2 \log(1/\delta)} \mathbb{E}\left[\sum_{t=1}^n \|\mathbf{A}_t\|_{\check{\Sigma}_t}\right] + \sqrt{\frac{2}{\pi}} \lambda_{\max,t} K n \delta, \\ &\stackrel{(i)}{\leq} \sqrt{2n \log(1/\delta)} \mathbb{E}\left[\sqrt{\sum_{t=1}^n \|\mathbf{A}_t\|_{\check{\Sigma}_t}^2}\right] + \sqrt{\frac{2}{\pi}} \lambda_{\max,t} K n \delta, \\ &\stackrel{(ii)}{\leq} \sqrt{2n \log(1/\delta)} \sqrt{\mathbb{E}\left[\sum_{t=1}^n \|\mathbf{A}_t\|_{\check{\Sigma}_t}^2\right]} + \sqrt{\frac{2}{\pi}} \lambda_{\max,t} K n \delta, \end{aligned}$$

where we use the Cauchy-Schwarz inequality in (i), and (ii) follows from the concavity of the square root. Now note that any $\mathbf{a} \in \mathcal{A}$ is an indicator vector and that $\check{\Sigma}_t$ is the covariance of the joint posterior of the expected rewards $(X_t^\top \theta_{*,a})_{a \in [K]} \mid H_t$. Therefore, for any $\mathbf{a} \in \mathcal{A}$, $\|\mathbf{a}\|_{\check{\Sigma}_t}^2 = \check{\sigma}_a^2$ is the variance of $X_t^\top \theta_{*,a} \mid H_t$. But we know that $\theta_{*,a} \mid H_t$ is a multivariate Gaussian and its covariance is $\hat{\Sigma}_{t,a}$ (Lemma 3). Thus the variance of $X_t^\top \theta_{*,a} \mid H_t$ is $\check{\sigma}_a^2 = X_t^\top \hat{\Sigma}_{t,a} X_t$. It follows that for

any $\mathbf{a} \in \mathcal{A}$, $\|\mathbf{a}\|_{\tilde{\Sigma}_t}^2 = X_t^\top \tilde{\Sigma}_{t,a} X_t = \|X_t\|_{\tilde{\Sigma}_{t,a}}^2$. In particular, $\|\mathbf{A}_t\|_{\tilde{\Sigma}_t}^2 = X_t^\top \tilde{\Sigma}_{t,A_t} X_t$. Combining this with (32) yields that $\lambda_{\max,t} = \max_{\mathbf{a} \in \mathcal{A}} \|\mathbf{a}\|_{\tilde{\Sigma}_t} = \max_{a \in \mathcal{A}} \|X_t\|_{\tilde{\Sigma}_{t,a}} \leq \max_{a \in \mathcal{A}} \sqrt{\lambda_1(\tilde{\Sigma}_{t,a}) \kappa_x} \leq \sqrt{\left(\lambda_{1,0} + \frac{\lambda_{1,0}^2 \lambda_{1,\Psi} \kappa_b}{\lambda_{d,0}^2}\right) \kappa_x}$. Then we let $c = \sqrt{\frac{2}{\pi} \left(\lambda_{1,0} + \frac{\lambda_{1,0}^2 \lambda_{1,\Psi} \kappa_b}{\lambda_{d,0}^2}\right) \kappa_x} K$ which allows us to write

$$\mathcal{BR}(n) \leq \sqrt{2n \log(1/\delta)} \sqrt{\mathbb{E} \left[\sum_{t=1}^n \|X_t\|_{\tilde{\Sigma}_{t,A_t}}^2 \right]} + cn\delta. \quad (39)$$

Now we focus on the the term $\sqrt{\mathbb{E} \left[\sum_{t=1}^n \|X_t\|_{\tilde{\Sigma}_{t,A_t}}^2 \right]}$ that we decompose and bound as

$$\begin{aligned} \|X_t\|_{\tilde{\Sigma}_{t,A_t}}^2 &= \sigma^2 \frac{X_t^\top \tilde{\Sigma}_{t,A_t} X_t}{\sigma^2} \stackrel{(i)}{=} \sigma^2 \left(\sigma^{-2} X_t^\top \tilde{\Sigma}_{t,A_t} X_t + \sigma^{-2} X_t^\top \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \Gamma_{A_t} \bar{\Sigma}_t \Gamma_{A_t}^\top \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} X_t \right), \\ &\stackrel{(ii)}{\leq} c_A \log(1 + \sigma^{-2} X_t^\top \tilde{\Sigma}_{t,A_t} X_t) + c_1 \log(1 + \sigma^{-2} X_t^\top \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \Gamma_{A_t} \bar{\Sigma}_t \Gamma_{A_t}^\top \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} X_t), \end{aligned} \quad (40)$$

where (i) follows from $\tilde{\Sigma}_{t,A_t} = \tilde{\Sigma}_{t,A_t} + \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \Gamma_{A_t} \bar{\Sigma}_t \Gamma_{A_t}^\top \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t}$, and we use the following inequality in (ii)

$$x = \frac{x}{\log(1+x)} \log(1+x) \leq \left(\max_{x \in [0,u]} \frac{x}{\log(1+x)} \right) \log(1+x) = \frac{u}{\log(1+u)} \log(1+x),$$

which holds for any $x \in [0, u]$, where constants c_A and c_1 are derived as

$$c_A = \frac{\kappa_x \lambda_{1,0}}{\log(1 + \sigma^{-2} \kappa_x \lambda_{1,0})}, \quad c_1 = \frac{c_\Psi}{\log(1 + \sigma^{-2} c_\Psi)}, \quad c_\Psi = \frac{\kappa_x \kappa_b \lambda_{1,0}^2 \lambda_{1,\Psi}}{\lambda_{d,0}^2},$$

The derivation of c_A uses that

$$X_t^\top \tilde{\Sigma}_{t,A_t} X_t \leq \lambda_1(\tilde{\Sigma}_{t,A_t}) \|X_t\|^2 \leq \lambda_d^{-1}(\Sigma_{0,A_t}^{-1} + G_{t,A_t}) \kappa_x \leq \lambda_d^{-1}(\Sigma_{0,A_t}^{-1}) \kappa_x = \lambda_1(\Sigma_{0,A_t}) \kappa_x \leq \lambda_{1,0} \kappa_x.$$

The derivation of c_1 follows from

$$\begin{aligned} X_t^\top \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \Gamma_{A_t} \bar{\Sigma}_t \Gamma_{A_t}^\top \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} X_t &\leq \lambda_1^2(\tilde{\Sigma}_{t,A_t}) \lambda_1^2(\Sigma_{0,A_t}^{-1}) \lambda_1(\Gamma_{A_t} \bar{\Sigma}_t \Gamma_{A_t}^\top) \kappa_x \leq \frac{\lambda_1^2(\Sigma_{0,A_t}) \lambda_{1,\Psi} \lambda_1(\Gamma_{A_t} \bar{\Sigma}_t \Gamma_{A_t}^\top) \kappa_x}{\lambda_d^2(\Sigma_{0,A_t})}, \\ &\leq \frac{\lambda_{1,0}^2 \lambda_{1,\Psi} \kappa_b \kappa_x}{\lambda_{d,0}^2}. \end{aligned}$$

The first inequality follows from Weyl's inequality and the fact that $\lambda_1(\bar{\Sigma}_t) \leq \lambda_1(\Sigma_\Psi) = \lambda_{1,\Psi}$ and $\lambda_1(\tilde{\Sigma}_{t,A_t}) \leq \lambda_1(\Sigma_{0,A_t})$. Now we focus on bounding the logarithmic terms in (40).

First Term in (40) We first rewrite this term as

$$\begin{aligned} \log(1 + \sigma^{-2} X_t^\top \tilde{\Sigma}_{t,A_t} X_t) &\stackrel{(i)}{=} \log \det(I_d + \sigma^{-2} \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} X_t X_t^\top \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}}), \\ &= \log \det(\tilde{\Sigma}_{t,A_t}^{-1} + \sigma^{-2} X_t X_t^\top) - \log \det(\tilde{\Sigma}_{t,A_t}^{-1}) = \log \det(\tilde{\Sigma}_{t+1,A_t}^{-1}) - \log \det(\tilde{\Sigma}_{t,A_t}^{-1}), \end{aligned}$$

where (i) follows from the Weinstein–Aronszajn identity. Then we sum over all rounds $t \in [n]$, and get a telescoping that leads to

$$\begin{aligned} \sum_{t=1}^n \log \det(I_d + \sigma^{-2} \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} X_t X_t^\top \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}}) &= \sum_{t=1}^n \log \det(\tilde{\Sigma}_{t+1,A_t}^{-1}) - \log \det(\tilde{\Sigma}_{t,A_t}^{-1}), \\ &= \sum_{t=1}^n \sum_{i=1}^K \log \det(\tilde{\Sigma}_{t+1,i}^{-1}) - \log \det(\tilde{\Sigma}_{t,i}^{-1}) = \sum_{i=1}^K \sum_{t=1}^n \log \det(\tilde{\Sigma}_{t+1,i}^{-1}) - \log \det(\tilde{\Sigma}_{t,i}^{-1}), \\ &= \sum_{i=1}^K \log \det(\tilde{\Sigma}_{n+1,i}^{-1}) - \log \det(\tilde{\Sigma}_{1,i}^{-1}) \stackrel{(i)}{=} \sum_{i=1}^K \log \det(\Sigma_{0,i}^{\frac{1}{2}} \tilde{\Sigma}_{n+1,i}^{-1} \Sigma_{0,i}^{\frac{1}{2}}) \stackrel{(ii)}{\leq} \sum_{i=1}^K d \log \left(\frac{1}{d} \text{Tr}(\Sigma_{0,i}^{\frac{1}{2}} \tilde{\Sigma}_{n+1,i}^{-1} \Sigma_{0,i}^{\frac{1}{2}}) \right) \\ &\leq \sum_{i=1}^K d \log \left(1 + \frac{\kappa_x \lambda_1(\Sigma_{0,i}) n}{\sigma^2 d} \right) \leq K d \log \left(1 + \frac{\kappa_x \lambda_{1,0} n}{\sigma^2 d} \right). \end{aligned}$$

where (i) follows from the fact that $\tilde{\Sigma}_{1,i} = \Sigma_{0,i}$ and we use the inequality of arithmetic and geometric means in (ii).

Second Term in (40) First, we rewrite the covariance matrix of the effect posterior $\bar{\Sigma}_t$ using the compact notation introduced in Appendix D.1. Precisely, it follows from (24) that

$$\begin{aligned} \bar{\Sigma}_t^{-1} &\stackrel{(i)}{=} \Sigma_{\Psi}^{-1} + \sum_{i=1}^K \Gamma_i^{\top} (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} \Gamma_i \stackrel{(ii)}{=} \Sigma_{\Psi}^{-1} + \sum_{i=1}^K \Gamma_i^{\top} (\Sigma_{0,i}^{-1} - \Sigma_{0,i}^{-1} (G_{t,i} + \Sigma_{0,i}^{-1})^{-1} \Sigma_{0,i}^{-1}) \Gamma_i, \\ &\stackrel{(iii)}{=} \Sigma_{\Psi}^{-1} + \sum_{i=1}^K \Gamma_i^{\top} \left(\Sigma_{0,i}^{-1} - \Sigma_{0,i}^{-1} \tilde{\Sigma}_{t,i} \Sigma_{0,i}^{-1} \right) \Gamma_i. \end{aligned} \quad (41)$$

The equality (i) requires $G_{t,i}$ to be invertible and was only given in the main manuscript to ease the exposition. In our proof, we use (ii) and (iii) which are the same; (iii) follows from plugging $\tilde{\Sigma}_{t,i} = (G_{t,i} + \Sigma_{0,i}^{-1})^{-1}$ in (ii). Now let $u = \sigma^{-1} \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} X_t$. Then it follows from (iii) in (41) that

$$\begin{aligned} \bar{\Sigma}_{t+1}^{-1} - \bar{\Sigma}_t^{-1} &= \Gamma_{A_t}^{\top} \left(\Sigma_{0,A_t}^{-1} - \Sigma_{0,A_t}^{-1} (\tilde{\Sigma}_{t,A_t}^{-1} + \sigma^{-2} X_t X_t^{\top})^{-1} \Sigma_{0,A_t}^{-1} - (\Sigma_{0,A_t}^{-1} - \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1}) \right) \Gamma_{A_t}, \\ &= \Gamma_{A_t}^{\top} \left(\Sigma_{0,A_t}^{-1} (\tilde{\Sigma}_{t,A_t} - (\tilde{\Sigma}_{t,A_t}^{-1} + \sigma^{-2} X_t X_t^{\top})^{-1}) \Sigma_{0,A_t}^{-1} \right) \Gamma_{A_t}, \\ &= \Gamma_{A_t}^{\top} \left(\Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} (I_d - (I_d + \sigma^{-2} \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} X_t X_t^{\top} \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}})^{-1}) \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} \Sigma_{0,A_t}^{-1} \right) \Gamma_{A_t}, \\ &= \Gamma_{A_t}^{\top} \left(\Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} (I_d - (I_d + uu^{\top})^{-1}) \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} \Sigma_{0,A_t}^{-1} \right) \Gamma_{A_t}, \\ &\stackrel{(i)}{=} \Gamma_{A_t}^{\top} \left(\Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} \frac{uu^{\top}}{1 + u^{\top}u} \tilde{\Sigma}_{t,A_t}^{\frac{1}{2}} \Sigma_{0,A_t}^{-1} \right) \Gamma_{A_t} = \sigma^{-2} \Gamma_{A_t}^{\top} \left(\Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} \frac{X_t X_t^{\top}}{1 + u^{\top}u} \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \right) \Gamma_{A_t}. \end{aligned} \quad (42)$$

In (i) we use the Sherman-Morrison formula. Moreover, we have that $\|X_t\|^2 \leq \kappa_x$. Therefore,

$$1 + u^{\top}u = 1 + \sigma^{-2} X_t^{\top} \tilde{\Sigma}_{t,A_t} X_t \leq 1 + \sigma^{-2} \kappa_x \lambda_1(\Sigma_{0,A_t}) \leq 1 + \sigma^{-2} \kappa_x \lambda_{1,0} = c_2.$$

This allows us to bound the second logarithmic term in (40) as

$$\begin{aligned} &\log(1 + \sigma^{-2} X_t^{\top} \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \Gamma_{A_t} \bar{\Sigma}_t \Gamma_{A_t}^{\top} \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} X_t), \\ &\stackrel{(i)}{\leq} c_2 \log(1 + c_2^{-1} \sigma^{-2} X_t^{\top} \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \Gamma_{A_t} \bar{\Sigma}_t \Gamma_{A_t}^{\top} \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} X_t), \\ &\stackrel{(ii)}{=} c_2 \log \det(I_{Ld} + c_2^{-1} \sigma^{-2} \tilde{\Sigma}_t^{\frac{1}{2}} \Gamma_{A_t}^{\top} \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} X_t X_t^{\top} \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_t^{\frac{1}{2}} \Gamma_{A_t}), \\ &\stackrel{(iii)}{=} c_2 \left[\log \det(\bar{\Sigma}_t^{-1} + c_2^{-1} \sigma^{-2} \Gamma_{A_t}^{\top} \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} X_t X_t^{\top} \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \Gamma_{A_t}) - \log \det(\bar{\Sigma}_t^{-1}) \right], \\ &\stackrel{(iv)}{\leq} c_2 \left[\log \det(\bar{\Sigma}_t^{-1} + \sigma^{-2} \Gamma_{A_t}^{\top} \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} \frac{X_t X_t^{\top}}{1 + u^{\top}u} \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \Gamma_{A_t}) - \log \det(\bar{\Sigma}_t^{-1}) \right], \\ &\stackrel{(v)}{=} c_2 [\log \det(\bar{\Sigma}_{t+1}^{-1}) - \log \det(\bar{\Sigma}_t^{-1})]. \end{aligned}$$

Here (i) follows from the fact that $\log(1 + x) \leq c_2 \log(1 + c_2^{-1}x)$ for any $x \geq 0$ and $c_2 \geq 1$. In (ii), we use the Weinstein–Aronszajn identity. In (iii), we use the log product formula and the fact that the det is a multiplicative map. In (iv), we use that $c_2^{-1} \leq 1/(1 + u^{\top}u)$. Finally, (v) follows from (42). Now we sum over all rounds and get telescoping

$$\begin{aligned} &\sum_{t=1}^n \log \det(I_{Ld} + \sigma^{-2} \tilde{\Sigma}_t^{\frac{1}{2}} \Gamma_{A_t}^{\top} \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_{t,A_t} X_t X_t^{\top} \tilde{\Sigma}_{t,A_t} \Sigma_{0,A_t}^{-1} \tilde{\Sigma}_t^{\frac{1}{2}} \Gamma_{A_t}), \\ &\leq c_2 [\log \det(\bar{\Sigma}_{n+1}^{-1}) - \log \det(\bar{\Sigma}_1^{-1})] = c_2 \log \det(\Sigma_{\Psi}^{\frac{1}{2}} \bar{\Sigma}_{n+1}^{-1} \Sigma_{\Psi}^{\frac{1}{2}}) \stackrel{(i)}{\leq} c_2 Ld \log \left(\frac{1}{Ld} \text{Tr}(\Sigma_{\Psi}^{\frac{1}{2}} \bar{\Sigma}_{n+1}^{-1} \Sigma_{\Psi}^{\frac{1}{2}}) \right), \\ &\stackrel{(ii)}{\leq} c_2 Ld \log(\lambda_1(\Sigma_{\Psi}^{\frac{1}{2}} \bar{\Sigma}_{n+1}^{-1} \Sigma_{\Psi}^{\frac{1}{2}})) \leq c_2 Ld \log(\lambda_{1,\Psi} \lambda_1(\bar{\Sigma}_{n+1}^{-1})), \\ &\stackrel{(iii)}{\leq} c_2 Ld \log \left(1 + K \kappa_b \lambda_{1,\Psi} \left(\frac{1}{\lambda_{d,0}} - \frac{1}{\lambda_{1,0}^2 \left(\frac{\kappa_x n}{\sigma^2} + \frac{1}{\lambda_{d,0}} \right)} \right) \right), \end{aligned}$$

In (i) we use the inequality of arithmetic and geometric means. In (ii) we bound all eigenvalues in the trace by the maximum eigenvalue. In (iii) we use the result in (33). We combine the upper bounds for both logarithmic terms and get

$$\mathbb{E} \left[\sum_{t=1}^n \|X_t\|_{\Sigma_{t,A_t}}^2 \right] \leq Kdc_A \log \left(1 + \frac{\kappa_x \lambda_{1,0} n}{\sigma^2 d} \right) + Ldc_1 c_2 \log \left(1 + K\kappa_b \lambda_{1,\Psi} \left(\frac{1}{\lambda_{d,0}} - \frac{1}{\lambda_{1,0}^2 \left(\frac{\kappa_x n}{\sigma^2} + \frac{1}{\lambda_{d,0}} \right)} \right) \right).$$

Finally, we set $c_E = c_1 c_2$, which concludes the proof for the general case. To retrieve the result in Theorem 1, we only need to set $\lambda_{1,0} = \lambda_{d,0} = \sigma_0^2$ and $\lambda_{1,\Psi} = \sigma_\Psi^2$ since we assumed that $\Sigma_\Psi = \sigma_\Psi^2 I_{Ld}$ and that $\Sigma_{0,i} = \sigma_0^2 I_d$ for any $i \in [K]$. In that case, the second term simplifies as

$$\begin{aligned} \log \left(1 + K\kappa_b \lambda_{1,\Psi} \left(\frac{1}{\lambda_{d,0}} - \frac{1}{\lambda_{1,0}^2 \left(\frac{\kappa_x n}{\sigma^2} + \frac{1}{\lambda_{d,0}} \right)} \right) \right) &= \log \left(1 + K\kappa_b \sigma_\Psi^2 \left(\frac{1}{\sigma_0^2} - \frac{1}{\sigma_0^4 \left(\frac{\kappa_x n}{\sigma^2} + \frac{1}{\sigma_0^2} \right)} \right) \right), \\ &= \log \left(1 + K\kappa_b \sigma_\Psi^2 \frac{n\kappa_x}{n\kappa_x \sigma_0^2 + \sigma^2} \right). \end{aligned}$$

□

E EXTENSIONS

Here we present and discuss in detail possible extensions of meTS. We start with the factored approximation of the effect posteriors (Appendix E.1) which improves computational efficiency with minimal impact on the empirical regret (Section 5 and Appendix F). We provide closed-form solutions for the factored effect posteriors in all the settings that we consider in this paper. While Algorithm 1 can be applied to the general two-level hierarchical setting introduced in Section 2, we only focused on cases where the dependencies of action parameters with effect parameters are captured through a linear combination in the theoretical analysis and experiments. In Appendix E.2, we provide an extension of our analysis to the case where the weights $b_{i,\ell}$ are replaced by matrices $C_{i,\ell}$. Moreover, in Appendix E.3, we present a way to introduce non-linearity in effects. In Appendix E.4, we motivate *deeper* hierarchies, and provide intuition on the corresponding regret.

E.1 Factored Effect Posteriors

As discussed earlier, the number of actions K is often much larger than the number of effect parameters L . However, L can also be large. In this section, we show how to improve the computational efficiency of meTS using factored distributions (Bishop, 2006). Consider the practical models in (2) and (3) where the effect posterior is a multivariate Gaussian $Q_t = \mathcal{N}(\bar{\mu}_t, \bar{\Sigma}_t)$ (Sections 3.2 and 3.3). Now suppose that it factorizes, that is $Q_t(\Psi) = \prod_{\ell=1}^L Q_{t,\ell}(\psi_\ell)$, where $Q_{t,\ell}$ is the effect posterior of the ℓ -th effect parameter $\psi_{*,\ell}$. Then for any round $t \in [n]$, the effect posterior $Q_{t,\ell}$ is also a multivariate Gaussian $Q_{t,\ell} = \mathcal{N}(\bar{\mu}_{t,\ell}, \bar{\Sigma}_{t,\ell})$, where $\bar{\Sigma}_{t,\ell}$ is the ℓ -th $d \times d$ diagonal block of $\bar{\Sigma}_t$, and $\bar{\mu}_{t,\ell} \in \mathbb{R}^d$ are such that $\bar{\mu}_t = (\bar{\mu}_{t,\ell})_{\ell \in [L]}$. This allows for individual sampling of the effect parameters, which improves the space and time complexity. Next we provide the factored effect posterior for the mixed-effect bandit settings considered in our paper.

Mixed-Effect Linear Bandit Consider the model in (2), we have that for any round $t \in [n]$, the effect posterior $Q_{t,\ell}$ is also a multivariate Gaussian $Q_{t,\ell} = \mathcal{N}(\bar{\mu}_{t,\ell}, \bar{\Sigma}_{t,\ell})$, where

$$\begin{aligned} \bar{\Sigma}_{t,\ell}^{-1} &= \Sigma_{\psi_\ell}^{-1} + \sum_{i \in [K]} b_{i,\ell}^2 (\Sigma_{0,i} + G_{t,i}^{-1})^{-1}, \\ \bar{\mu}_{t,\ell} &= \bar{\Sigma}_{t,\ell} \left(\Sigma_{\psi_\ell}^{-1} \mu_{\psi_\ell} + \sum_{i \in [K]} b_{i,\ell} ((\Sigma_{0,i} + G_{t,i}^{-1})^{-1} G_{t,i}^{-1} B_{t,i}) \right). \end{aligned} \quad (43)$$

Σ_{ψ_ℓ} is the ℓ -th $d \times d$ diagonal block of Σ_Ψ , and $\mu_{\psi_\ell} \in \mathbb{R}^d$ are such that $\mu_\Psi = (\mu_{\psi_\ell})_{\ell \in [L]}$.

Mixed-Effect Generalized Linear Bandit Consider the model in (3), we have that for any round $t \in [n]$, the effect posterior $Q_{t,\ell}$ is also a multivariate Gaussian $Q_{t,\ell} = \mathcal{N}(\bar{\mu}_{t,\ell}, \bar{\Sigma}_{t,\ell})$, where

$$\begin{aligned} \bar{\Sigma}_{t,\ell}^{-1} &= \Sigma_{\psi_\ell}^{-1} + \sum_{i \in [K]} b_{i,\ell}^2 (\Sigma_{0,i} + G_{t,i}^{-1})^{-1}, \\ \bar{\mu}_{t,\ell} &= \bar{\Sigma}_{t,\ell} \left(\Sigma_{\psi_\ell}^{-1} \mu_{\psi_\ell} + \sum_{i \in [K]} b_{i,\ell} ((\Sigma_{0,i} + G_{t,i}^{-1})^{-1} G_{t,i}^{-1} B_{t,i}) \right). \end{aligned} \quad (44)$$

Σ_{ψ_ℓ} is the ℓ -th $d \times d$ diagonal block of Σ_Ψ , and $\mu_{\psi_\ell} \in \mathbb{R}^d$ are such that $\mu_\Psi = (\mu_{\psi_\ell})_{\ell \in [L]}$.

Mixed-Effect Multi-Armed Bandit Consider the model in (17), we have that for any round $t \in [n]$, the effect posterior $Q_{t,\ell}$ is a *univariate* Gaussian $Q_{t,\ell} = \mathcal{N}(\bar{\mu}_{t,\ell}, \bar{\Sigma}_{t,\ell})$, where

$$\begin{aligned}\bar{\sigma}_{t,\ell}^{-2} &= \sigma_{\psi_\ell}^{-2} + \sum_{i \in [K]} b_{i,\ell}^2 \frac{N_{t,i}}{N_{t,i}\sigma_{0,i}^2 + \sigma^2}, \\ \bar{\mu}_{t,\ell} &= \bar{\sigma}_{t,\ell}^2 \left(\sigma_{\psi_\ell}^{-2} \mu_{\psi_\ell} + \sum_{i \in [K]} b_{i,\ell} \frac{B_{t,i}}{N_{t,i}\sigma_{0,i}^2 + \sigma^2} \right).\end{aligned}\quad (45)$$

$\sigma_{\psi_\ell}^2 > 0$ is the ℓ -th diagonal entry of Σ_Ψ , and $\mu_{\psi_\ell} \in \mathbb{R}$ is the ℓ -th entry of μ_Ψ .

Proof. To reduce clutter, we consider a fixed round $t \in [n]$, and drop subindexing by t . It follows that $Q = \mathcal{N}(\bar{\mu}, \bar{\Sigma})$ corresponds to the effect posterior $Q_t = \mathcal{N}(\bar{\mu}_t, \bar{\Sigma}_t)$ for some round t . Here we restrict the family of effect posteriors Q to factored distributions. Precisely, we first partition the elements of $\Psi_* = (\psi_{*,\ell})_{\ell \in [L]}$ into L disjoint d -dimensional groups where each group corresponds to an effect parameter $\psi_{*,\ell}$. We then suppose that Q factorizes across the L effect parameters, that is $Q(\Psi) = \prod_{\ell \in [L]} Q_\ell(\psi_\ell)$, where Q_ℓ are obtained using variational inference techniques (Bishop, 2006) as we show next. First, we know that $Q(\Psi) = \mathcal{N}(\bar{\mu}, \bar{\Sigma})$, where $\bar{\mu} \in \mathbb{R}^{Ld}$ and $\bar{\Sigma} \in \mathbb{R}^{Ld \times Ld}$. We write the mean and covariance by blocks as $\bar{\mu} = (\bar{\mu}_\ell)_{\ell \in [L]}$ and $\bar{\Sigma} = (\bar{\Sigma}_{i,j})_{(i,j) \in [L] \times [L]}$, such that $\bar{\mu}_\ell \in \mathbb{R}^d$ and $\bar{\Sigma}_{i,j} \in \mathbb{R}^{d \times d}$. Now fix $\ell \in [L]$, from known results (Bishop, 2006) the optimal factor Q_ℓ that optimizes the Kullback-Leibler divergence satisfies

$$Q_\ell(\psi_\ell) \propto \exp(\mathbb{E}_{j \neq \ell} [\log Q(\Psi)]), \quad (46)$$

where $\mathbb{E}_{j \neq \ell}[\cdot]$ denotes an expectation with respect to the distributions Q_j such that $j \neq \ell$. Let $\bar{\Lambda} = \bar{\Sigma}^{-1} = (\bar{\Lambda}_{i,j})_{(i,j) \in [L] \times [L]}$, the expectation can be computed as

$$\begin{aligned}Q_\ell(\psi_\ell) &\propto \exp \left(\mathbb{E}_{j \neq \ell} \left[-\frac{1}{2}(\psi_\ell - \bar{\mu}_\ell)^\top \bar{\Lambda}_{\ell,\ell}(\psi_\ell - \bar{\mu}_\ell) + \sum_{j \neq \ell} (\psi_\ell - \bar{\mu}_\ell)^\top \bar{\Lambda}_{\ell,j}(\psi_j - \bar{\mu}_j) \right] \right), \\ &\propto \exp \left(\mathbb{E}_{j \neq \ell} \left[-\frac{1}{2}\psi_\ell^\top \bar{\Lambda}_{\ell,\ell}\psi_\ell + \psi_\ell^\top \bar{\Lambda}_{\ell,\ell}\bar{\mu}_\ell - \psi_\ell^\top \sum_{j \neq \ell} \bar{\Lambda}_{\ell,j}(\mathbb{E}[\psi_j] - \bar{\mu}_j) \right] \right), \\ &\propto \exp \left(\mathbb{E}_{j \neq \ell} \left[-\frac{1}{2}\psi_\ell^\top \bar{\Lambda}_{\ell,\ell}\psi_\ell + \psi_\ell^\top \left(\bar{\Lambda}_{\ell,\ell}\bar{\mu}_\ell - \sum_{j \neq \ell} \bar{\Lambda}_{\ell,j}(\mathbb{E}[\psi_j] - \bar{\mu}_j) \right) \right] \right).\end{aligned}\quad (47)$$

Thus, we have that

$$Q_\ell(\psi_\ell) = \mathcal{N}(\psi_\ell; m_\ell, \bar{\Sigma}_{\ell,\ell}), \quad (48)$$

where $m_\ell = \bar{\mu}_\ell - \bar{\Sigma}_{\ell,\ell} \sum_{j \neq \ell} \bar{\Lambda}_{\ell,j}(\mathbb{E}[\psi_j] - \bar{\mu}_j)$. These solutions are coupled since the optimal factor Q_ℓ depends on the other optimal factors Q_j for $j \neq \ell$. However, we can provide a closed-form solution in this Gaussian case if we set $m_\ell = \mathbb{E}[\psi_\ell] = \bar{\mu}_\ell$ for all $\ell \in [L]$; in which case we get that $Q_\ell(\psi_\ell) = \mathcal{N}(\psi_\ell; \bar{\mu}_\ell, \bar{\Sigma}_{\ell,\ell})$ for all $\ell \in [L]$. To summarize, we showed that if we suppose that the effect posterior factorizes, that is $Q(\Psi) = \prod_{\ell \in [L]} Q_\ell(\psi_\ell)$, and that $Q = \mathcal{N}(\bar{\mu}, \bar{\Sigma})$, then the optimal factors Q_ℓ are also Gaussians $Q_\ell = \mathcal{N}(\bar{\mu}_\ell, \bar{\Lambda}_{\ell,\ell}^{-1})$, where $\bar{\mu}_\ell \in \mathbb{R}^d$ and $\bar{\Lambda}_{\ell,\ell} \in \mathbb{R}^{d \times d}$ are such that $\bar{\mu} = (\bar{\mu}_\ell)_{\ell \in [L]}$ and $\bar{\Lambda} = (\bar{\Lambda}_{i,j})_{(i,j) \in [L] \times [L]}$. Finally, to get the desired results, we simply retrieve the respective $\bar{\mu}_\ell \in \mathbb{R}^d$ and $\bar{\Lambda}_{\ell,\ell} \in \mathbb{R}^{d \times d}$ from the mean and covariance of the exact posterior of either the model in (2), (3) or (17). \square

E.2 Finer Linear Effects

An effective way to capture fine-grained dependencies is to assume that the parameter of action i depends on effect parameters through L known matrices $C_{i,\ell} \in \mathbb{R}^{d \times d}$ as

$$\theta_{*,i} \mid \Psi_* \sim P_{0,i} \left(\cdot \mid \sum_{\ell=1}^L C_{i,\ell} \psi_{*,\ell} \right).$$

This generalizes the setting considered in our analysis, which corresponds to the case where $C_{i,\ell} = b_{i,\ell}I_d$. We first make the observation that $\sum_{\ell=1}^L C_{i,\ell}\psi_{*,\ell} = C_i\Psi_*$, where $C_i = [C_{i,1}, \dots, C_{i,L}] \in \mathbb{R}^{d \times Ld}$. It follows that for any round $t \in [n]$, the joint effect posterior is a multivariate Gaussian $Q_t = \mathcal{N}(\bar{\mu}_t, \bar{\Sigma}_t)$, where

$$\begin{aligned}\bar{\Sigma}_t^{-1} &= \Sigma_{\Psi}^{-1} + \sum_{i=1}^K C_i^{\top} (\Sigma_{0,i} + G_{t,i}^{-1})^{-1} C_i, \\ \bar{\mu}_t &= \bar{\Sigma}_t \left(\Sigma_{\Psi}^{-1} \mu_{\Psi} + \sum_{i=1}^K C_i^{\top} ((\Sigma_{0,i} + G_{t,i}^{-1})^{-1} G_{t,i}^{-1} B_{t,i}) \right).\end{aligned}\quad (49)$$

Moreover, for any round $t \in [n]$ and action $i \in [K]$, the action posterior is a multivariate Gaussian $P_{t,i}(\cdot \mid \Psi_t) = \mathcal{N}(\cdot; \tilde{\mu}_{t,i}, \tilde{\Sigma}_{t,i})$, where

$$\tilde{\Sigma}_{t,i}^{-1} = G_{t,i} + \Sigma_{0,i}^{-1}, \quad \tilde{\mu}_{t,i} = \tilde{\Sigma}_{t,i} \left(B_{t,i} + \Sigma_{0,i}^{-1} \left(\sum_{\ell=1}^L C_{i,\ell} \psi_{t,\ell} \right) \right).\quad (50)$$

Finally, our regret proof extends smoothly leading to a Bayes regret upper bound similar to the one we derived in Theorem 1. The corresponding Bayes regret is given in the following proposition.

Proposition 4. *For any $\delta \in (0, 1)$, the Bayes regret of meTS , for the mixed-effect model in Appendix E.2, is bounded as*

$$\mathcal{BR}(n) \leq \sqrt{2n(\mathcal{R}^A(n) + \mathcal{R}^E(n)) \log(1/\delta)} + cn\delta,$$

where $c = \sqrt{\frac{2}{\pi} \kappa_x (\sigma_0^2 + \kappa_c \sigma_{\Psi}^2)} K$, $\kappa_c = \max_{i \in [K]} \lambda_1(C_i^{\top} C_i)$,

$$\begin{aligned}\mathcal{R}^A(n) &= dKc_A \log \left(1 + \frac{n\kappa_x \sigma_0^2}{d\sigma^2} \right), \quad c_A = \frac{\kappa_x \sigma_0^2}{\log \left(1 + \frac{\kappa_x \sigma_0^2}{\sigma^2} \right)}, \\ \mathcal{R}^E(n) &= dLc_E \log \left(1 + \frac{K\kappa_c \sigma_{\Psi}^2}{\sigma_0^2 + \frac{\sigma^2}{n\kappa_x}} \right), \quad c_E = \frac{\kappa_x \kappa_c \sigma_{\Psi}^2 \left(1 + \frac{\kappa_x \sigma_0^2}{\sigma^2} \right)}{\log \left(1 + \frac{\kappa_x \kappa_c \sigma_{\Psi}^2}{\sigma^2} \right)}.\end{aligned}$$

The interpretation of this result is similar to Theorem 1. The only difference is that sparsity is now captured through κ_c .

E.3 Non-Linear Effects

Here the dependence of effect and action parameters are generated as in (2), except that a non-linear function $g(\cdot)$ is applied to the linear combination $\sum_{\ell=1}^L b_{i,\ell}\psi_{*,\ell}$. An example of g is the sigmoid function, and the whole model is

$$\begin{aligned}\Psi_* &\sim \mathcal{N}(\mu_{\Psi}, \Sigma_{\Psi}), \\ \theta_{*,i} \mid \Psi_* &\sim \mathcal{N} \left(g \left(\sum_{\ell=1}^L b_{i,\ell} \psi_{*,\ell} \right), \Sigma_{0,i} \right), \quad \forall i \in [K], \\ Y_t \mid X_t, \theta_{*,A_t} &\sim \mathcal{N}(X_t^{\top} \theta_{*,A_t}, \sigma^2), \quad \forall t \in [n].\end{aligned}\quad (51)$$

The action posteriors has closed-form solution. Precisely, for any round $t \in [n]$, action $i \in [K]$, and effect parameters Ψ_t , the action posterior is a multivariate Gaussian $P_{t,i}(\cdot \mid \Psi_t) = \mathcal{N}(\cdot; \tilde{\mu}_{t,i}, \tilde{\Sigma}_{t,i})$, where

$$\begin{aligned}\tilde{\Sigma}_{t,i}^{-1} &= \Sigma_{0,i}^{-1} + G_{t,i}, \\ \tilde{\mu}_{t,i} &= \tilde{\Sigma}_{t,i} \left(\Sigma_{0,i}^{-1} g \left(\sum_{\ell=1}^L b_{i,\ell} \psi_{t,\ell} \right) + B_{t,i} \right).\end{aligned}\quad (52)$$

The action posterior is the same as in Proposition 2 except that the prior term $\sum_{\ell=1}^L b_{i,\ell}\psi_{t,\ell}$ is now replaced by $g(\sum_{\ell=1}^L b_{i,\ell}\psi_{t,\ell})$. The effect posterior does not have a closed-form solution and can be approximated using the Laplace approximation similarly to Section 3.3.

E.4 Beyond Two-Level Hierarchies

To motivate deeper hierarchies (Hong et al., 2022a), consider the problem of page construction in movie streaming services where J movies are organized into L categories. First, a category $\ell \in [L]$ is associated with a parameter $\psi_{*,\ell} \in \mathbb{R}^d$. Moreover, each movie $j \in [J]$ is associated with a parameter $\phi_{*,j} \in \mathbb{R}^d$, which is a combination of category parameters $\psi_{*,\ell}$ weighted by scalars that quantify how related is the movie j to each category. Finally, page layouts are actions and they are seen as lists (or slates) of movies. Each page layout $i \in [K]$ is associated with an action parameter $\theta_{*,i}$ which is also a combination of movies parameters $\phi_{*,j}$ weighted by a scalar that quantifies position bias. Precisely, this scalar is set to 0 if the corresponding movie is not present in the page, and it has high value if the movie is placed in a position with high visibility. This setting induces a three-level hierarchical model, for which we give a Gaussian example below.

$$\begin{aligned} \Psi_* &\sim \mathcal{N}(\mu_\Psi, \Sigma_\Psi), \\ \phi_{*,j} \mid \Psi_* &\sim \mathcal{N}\left(\sum_{\ell=1}^L b_{j,\ell} \psi_{*,\ell}, \Sigma_{\phi,j}\right), & \forall j \in [J], \\ \theta_{*,i} \mid \Phi_* &\sim \mathcal{N}\left(\sum_{j=1}^J w_{i,j} \phi_{*,j}, \Sigma_{\theta,i}\right), & \forall i \in [K], \\ Y_t \mid X_t, \theta_{*,A_t} &\sim \mathcal{N}(X_t^\top \theta_{*,A_t}, \sigma^2), & \forall t \in [n], \end{aligned} \quad (53)$$

where $\Psi_* = (\psi_{*,\ell})_{\ell \in [L]} \in \mathbb{R}^{Ld}$ and $\Phi_* = (\phi_{*,j})_{j \in [J]} \in \mathbb{R}^{Jd}$. meTS samples hierarchically as follows. First, we sample Ψ_t from the posterior of $\Psi_* \mid H_t$. We then sample Φ_t from the posterior of $\Phi_* \mid \Psi_*, H_t$. Finally, we sample $\theta_{t,i}$ individually from the posterior of $\theta_{*,i} \mid \Phi_*, H_{t,i}$. We expect the upper bound of the Bayes regret of (53) following our analysis to be decomposed in three terms $\tilde{\mathcal{O}}\left(\sqrt{n(\mathcal{R}^A(n) + \mathcal{R}_1^E + \mathcal{R}_2^E)}\right)$, where $\mathcal{R}^A(n) = \tilde{\mathcal{O}}(Kd)$, $\mathcal{R}_1^E = \tilde{\mathcal{O}}(Jd)$, and $\mathcal{R}_2^E = \tilde{\mathcal{O}}(Ld)$.

F ADDITIONAL EXPERIMENTS

We provide additional experiments where we evaluate meTS using synthetic and real-world problems, and compare it to baselines that either ignore or partially use effect parameters. In each plot, we report the averages and standard errors of the quantities. Both settings are described in Section 5.

F.1 Synthetic Experiments

In Figures 4 and 5, we report regret from 12 experiments with horizon $n = 5000$, where we vary K and d and use both linear and logistic rewards. For the linear setting, we compare meTS-Lin (Section 3.2), LinUCB (Abbasi-Yadkori et al., 2011), LinTS (Agrawal and Goyal, 2013) and HierTS (Hong et al., 2022b). For the logistic setting, we compare meTS-GLM (Section 3.3), meTS-Lin (Section 3.2), UCB-GLM (Li et al., 2017), GLM-TS (Chapelle and Li, 2012) and HierTS (Hong et al., 2022b). We also include the factored approximation of meTS (meTS-Lin-Fa and meTS-GLM-Fa). In all experiments, we observe that meTS-Lin and meTS-Fa outperform other baselines that ignore the effect parameters or incorporate them partially. We also notice that the gain in performance becomes smaller when K/L decreases.

F.2 MovieLens Experiments

We plot the regret of meTS and the baselines up to $n = 5000$ rounds in Figures 6 and 7. We observe that meTS outperforms the other baselines. This is despite the fact that we did not fine-tune the mixing weights, which attests to the robustness of our approach to model misspecification. Similarly to the synthetic problems, we observe that the gap in performance between meTS and other baselines is less significant when K/L is small.

F.3 Robustness to Model Misspecification

We conduct additional synthetic experiments where the hyper-parameters do not match the parameters of the bandit environment to assess the robustness of our approach to misspecification. We provide results for this experiment in Figure 8. Here we consider the setting described in Section 5.1 except that the true hyper-parameters are misspecified as follows. At each run, we sample uniformly 4 misspecification constants c_1, c_2, c_3 , and c_4 from $(0, 2)$ and set the hyper-parameters of meTS-Lin as $c_1 \Sigma_\Psi$, $c_2 \mu_\Psi$, $c_3 \Sigma_{\theta,i}$, and $c_4 b_i$ for any $i \in [K]$; where Σ_Ψ , μ_Ψ , $\Sigma_{\theta,i}$, and b_i for $i \in [K]$ are the true

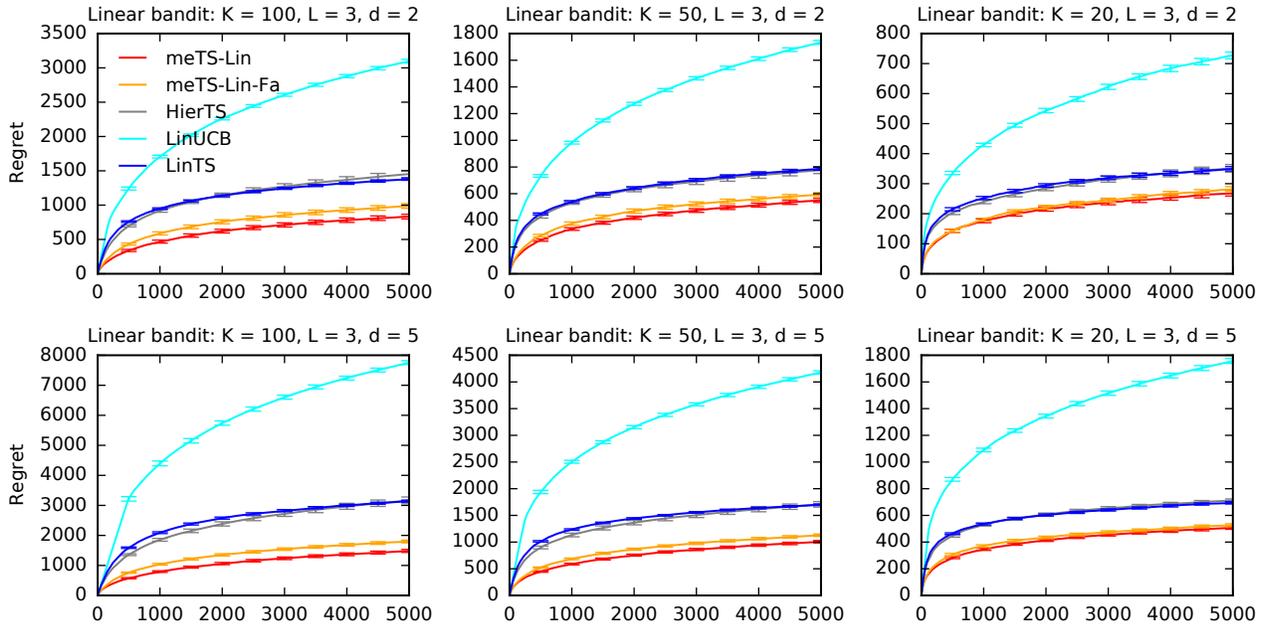


Figure 4: Regret of meTS-Lin on synthetic linear bandit problems with varying feature dimension $d \in \{2, 5\}$ and number of actions $K \in \{20, 50, 100\}$.

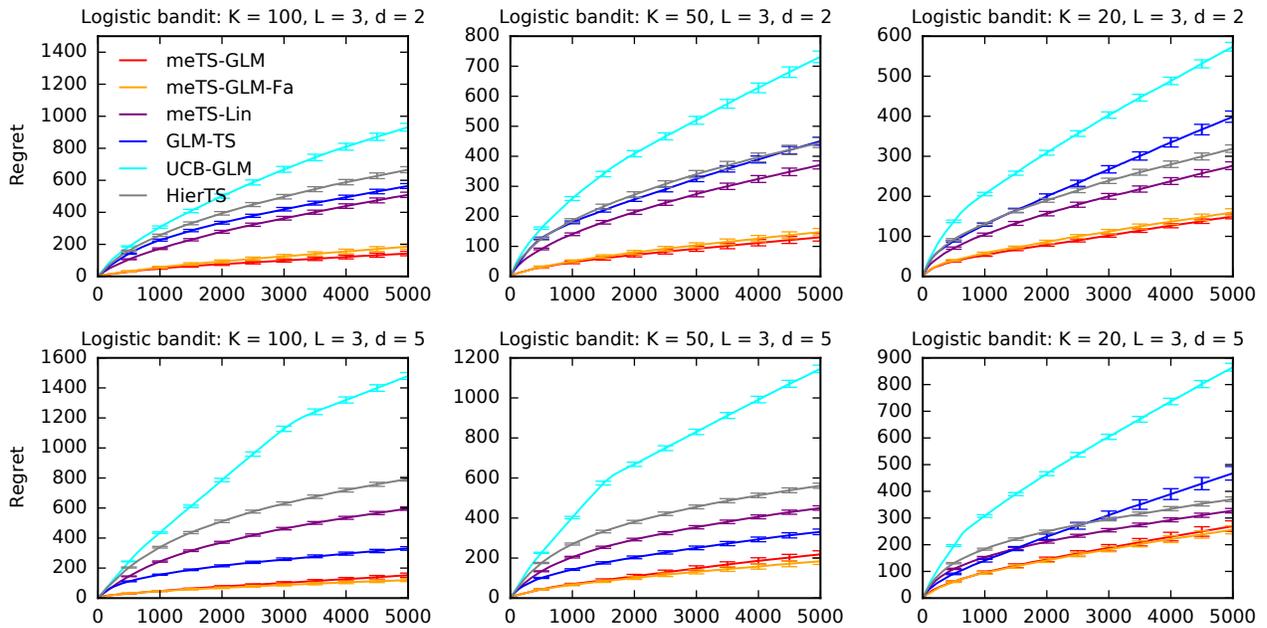


Figure 5: Regret of meTS-GLM on synthetic logistic bandit problems with varying feature dimension $d \in \{2, 5\}$ and number of actions $K \in \{20, 50, 100\}$.

hyper-parameters. Model misspecification is only applied to meTS-Lin and we refer to it as meTS-Lin-mis. We compare it to meTS-Lin and the other baselines, all with the true hyper-parameters. Although the baselines are not misspecified, meTS-Lin-mis still performs better. meTS-Lin-mis also performs similarly to meTS-Lin (with true hyper-parameters).

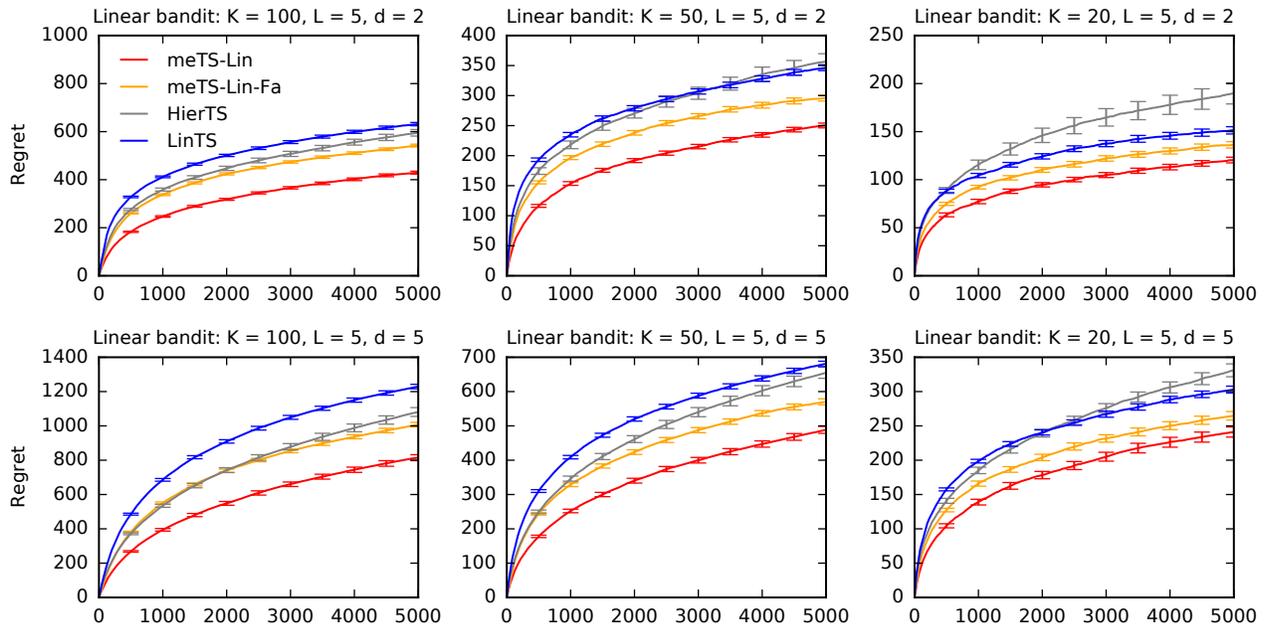


Figure 6: Regret of `meTS-Lin` on the MovieLens dataset with linear rewards and varying feature dimension $d \in \{2, 5\}$ and number of actions $K \in \{20, 50, 100\}$.

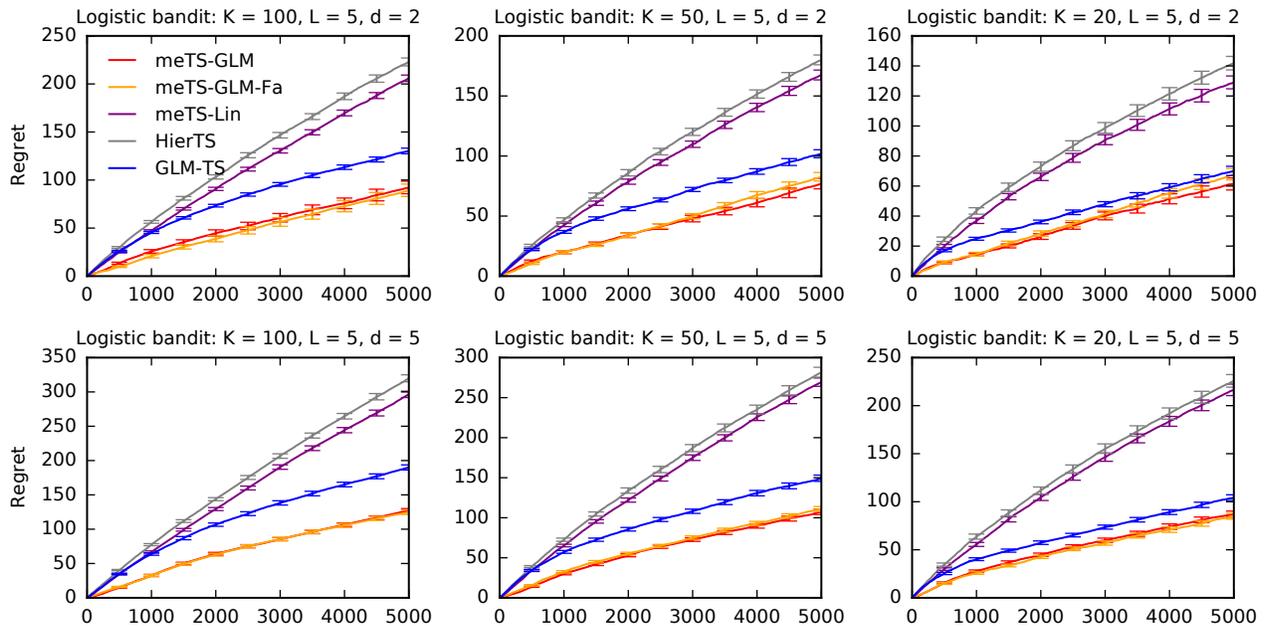


Figure 7: Regret of `meTS-GLM` on the MovieLens dataset with logistic rewards and varying feature dimension $d \in \{2, 5\}$ and number of actions $K \in \{20, 50, 100\}$.

F.4 Effect of Action Uncertainty

As we mentioned in Section 5.1, learning the effect parameters is most beneficial when they are more uncertain than the action parameters. In this section, we support this claim by conducting an experiment where the initial uncertainty of action parameters is greater than the initial uncertainty of the effect parameters. Precisely, we consider the setting described in

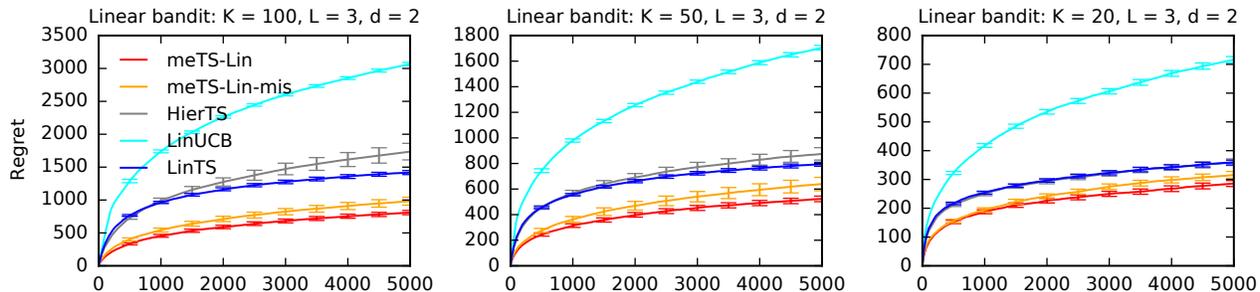


Figure 8: Regret of *misspecified* meTS-Lin on synthetic bandit problems with a varying number of actions K . Here, the *misspecified* meTS, meTS-Lin-mis, is compared to baselines with true hyper-parameters.

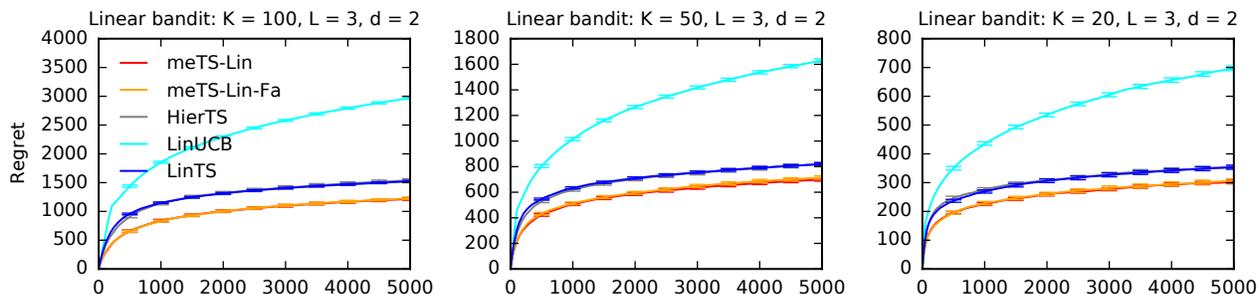


Figure 9: Regret of meTS-Lin on synthetic bandit problems with a varying number of actions K , where the action parameters are more uncertain than the effect parameters.

Section 5.1 except that we set $\Sigma_\Psi = I_{Ld}$ and $\Sigma_{0,i} = 3I_d$ for all $i \in [K]$. We report the results in Figure 9. By comparing Figure 9 to Figure 2, we observe that meTS-Lin still outperforms the baselines but the gap in performance shrinks when the action parameters are more uncertain than the effect parameters.

G SOCIETAL IMPACT

The goal of this work is to develop and analyze practical algorithms for contextual bandits with correlated actions. We are not aware of any potential negative societal impacts of our work since we did not propose any new applications of bandit algorithms than existing ones. A typical application of bandit algorithms is recommendation where the preferred items are shown to users. However, by doing this, the recommender system tends to restrain the user to these preferences which may raise concerns regarding the corresponding societal impact.