
Adaptive Exploration in Linear Contextual Bandit

Botao Hao

Princeton University

Tor Lattimore

Deepmind

Csaba Szepesvári

Deepmind and University of Alberta

Abstract

Contextual bandits serve as a fundamental model for many sequential decision making tasks. The most popular theoretically justified approaches are based on the optimism principle. While these algorithms can be practical, they are known to be sub-optimal asymptotically. On the other hand, existing asymptotically optimal algorithms for this problem do not exploit the linear structure in an optimal way and suffer from lower-order terms that dominate the regret in all practically interesting regimes. We start to bridge the gap by designing an algorithm that is asymptotically optimal and has good finite-time empirical performance. At the same time, we make connections to the recent literature on when exploration-free methods are effective. Indeed, if the distribution of contexts is well behaved, then our algorithm acts mostly greedily and enjoys sub-logarithmic regret. Furthermore, our approach is adaptive in the sense that it automatically detects the nice case. Numerical results demonstrate significant regret reductions by our method relative to several baselines.

1 INTRODUCTION

Stochastic contextual linear bandits, the problem we consider, is interesting due to its rich structure and also because of its potential applications, e.g., in

online recommendation systems [2, 18]. In this paper we propose a new algorithm for this problem that is asymptotically optimal, computationally efficient and empirically well-behaved in finite-time regimes. As a consequence of asymptotic optimality, the algorithm adapts to easy instances where it achieves sub-logarithmic regret.

Popular approaches for regret minimisation in contextual bandits include ϵ -greedy [15], explicit optimism-based algorithms [9, 20, 7, 1], and implicit ones, such as Thompson sampling [3]. Although these algorithms enjoy near-optimal worst-case guarantees and can be quite practical, they are known to be arbitrarily suboptimal in the asymptotic regime, even in the non-contextual linear bandit [16].

We propose an optimisation-based algorithm that estimates and tracks the optimal allocation for each context/action pair. This technique is most well known for its effectiveness in pure exploration [6, 12, 10, and others]. The approach has been used in regret minimisation in linear bandits with fixed action sets [16] and structured bandits [8]. The last two articles provide algorithms for the non-contextual case and hence cannot be applied directly to our setting. More importantly, however, the algorithms are not practical. The first algorithm uses a complicated three-phase construction that barely updates its estimates. The second algorithm is not designed to handle large action spaces and has a ‘lower-order’ term in the regret that depends linearly on the number of actions and dominates the regret in all practical regimes. This lower-order term is not merely a product of the analysis, but also reflected in the experiments (see Section 5.4 for details).

The most closely related work is by Ok et al. [19] who study a reinforcement learning setting. A stochastic contextual bandit can be viewed as a Markov decision

process where the state represents the context and the transition is independent of the action. The structured nature of the mentioned paper means our setting is covered by their algorithm. Again, however, the algorithm is too general to exploit the specific structure of the contextual bandit problem. Their algorithm is asymptotically optimal, but suffers from lower-order terms that are linear in the number of actions and dominate the regret in all practically interesting regimes. In contrast, our algorithm is asymptotically optimal, but also practical in finite-horizon regimes, as will be demonstrated by our experiments.

The contextual linear bandit also serves as an interesting example where the asymptotics of the problem are not indicative of what should be expected in finite-time (see the second scenario in Section 5.2). This is in contrast to many other bandit models where the asymptotic regret is also roughly optimal in finite time [17]. There is an important lesson here. Designing algorithms that optimize for the asymptotic regret may make huge sacrifices in finite-time.

Another interesting phenomenon is related to the idea of ‘natural exploration’ that occurs in contextual bandits [5, 13]. A number of authors have started to investigate the striking performance of greedy algorithms in contextual bandits. In most bandit settings the greedy policy does not explore sufficiently and suffers linear regret. In some contextual bandit problems, however, the changing features ensure the algorithm cannot help but explore. Our algorithm and analysis highlights this effect (see Section 3.1 for details). If the context distribution is sufficiently rich, then the algorithm is eventually almost completely greedy and enjoys sub-logarithmic regret. As opposed to the cited previous works, our algorithm achieves this under the cited favourable conditions *while* at the same time it satisfies the standard optimality guarantees when the favourable conditions do not hold. As another contribution, we prove that algorithms based on optimism, similarly to the new algorithm, also enjoy sub-logarithmic regret in the rich-context distribution setting (Theorem 3.9), and hence differences appear in lower order terms only between these algorithms.

The rest of the paper is organized as follows. We first introduce the problem setting (Section 2), which we follow by presenting our asymptotic lower bound

(Section 3). Section 4 introduces our new algorithm, which is claimed to match the lower bound. A proof sketch of this claim is presented in the same section. Section 5 presents experiments to illuminate the behaviour of the new algorithm in comparison to its strongest competitors. Section 6 discusses remaining notable open questions.

Notation Let $[n] = \{1, 2, \dots, n\}$. For a vector x and positive semidefinite matrix A we let $\|x\|_A = \sqrt{x^\top A x}$. The cardinality of a set \mathcal{A} is denoted by $|\mathcal{A}|$.

2 PROBLEM SETTING

We consider the stochastic K -armed contextual linear bandit with a horizon of n rounds and M possible contexts. The assumption that the contexts are discrete cannot be dropped but as we shall at least M will not play an important role in the regret bounds. This assumption would hold for example in a recommender system if users are clustered into finitely many groups. For each context $m \in [M]$ there is a known feature/action set $\mathcal{A}^m \subset \mathbb{R}^d$ with $|\mathcal{A}^m| = K$. The interaction protocol is as follows. First the environment samples a sequence of independent contexts $(c_t)_{t=1}^n$ from an unknown distribution p over $[M]$ and each context is assumed to appear with positive probability. At the start of round t the context c_t is revealed to the learner, who may use their observations to choose an action $X_t \in \mathcal{A}_t = \mathcal{A}^{c_t}$. The reward is

$$Y_t = \langle X_t, \theta \rangle + \eta_t,$$

where $(\eta_t)_{t=1}^n$ is a sequence of independent standard Gaussian random variables and $\theta \in \mathbb{R}^d$ is an unknown parameter. The Gaussian assumption can be relaxed to conditional sub-Gaussian assumption for the regret upper bound, but is necessary for the regret lower bound. Throughout, we consider a frequentist setting in the sense that θ is fixed. For simplicity, we assume each \mathcal{A}^m spans \mathbb{R}^d and $\|x\|_2 \leq 1$ for all $x \in \cup_m \mathcal{A}^m$.

The performance metric is the cumulative expected regret, which measures the difference between the expected cumulative reward collected by the omniscient policy that knows θ and the learner’s expected cumulative reward. The optimal arm associated with context m is $x_m^* = \arg\max_{x \in \mathcal{A}^m} \langle x, \theta \rangle$. Then the expected cumulative regret of a policy π when facing

the bandit determined by θ is

$$R_\theta^\pi(n) = \mathbb{E} \left[\sum_{t=1}^n \langle x_{c_t}^*, \theta \rangle - \sum_{t=1}^n Y_t \right].$$

Note that this cumulative regret also depends on the context distribution p and action sets. They are omitted from the notation to reduce clutter and because there will never be ambiguity.

3 ASYMPTOTIC LOWER BOUND

We investigate the fundamental limit of linear contextual bandit by deriving its instance-dependent asymptotic lower bound. First, we define the class of policies that are taken into consideration.

Definition 3.1 (Consistent Policy). A policy π is called consistent if the regret is subpolynomial for any bandit in that class and all context distributions:

$$R_\theta^\pi(n) = o(n^\varepsilon), \text{ for all } \varepsilon > 0 \text{ and all } \theta \in \mathbb{R}^d. \quad (3.1)$$

The next lemma is the key ingredient in proving the asymptotic lower bound. Given a context m and $x \in \mathcal{A}^m$ let $\Delta_x^m = \langle x_m^* - x, \theta \rangle$ be the suboptimality gap. Furthermore, let $\Delta_{\min} = \min_{m \in [M]} \min_{x \in \mathcal{A}^m, \Delta_x^m > 0} \Delta_x^m$.

Lemma 3.2. Assume that $p(m) > 0$ for all $m \in [M]$ and that x_m^* is uniquely defined for each context m and let π be consistent. Then for sufficiently large n the expected covariance matrix

$$\bar{G}_n = \mathbb{E} \left[\sum_{t=1}^n X_t X_t^\top \right], \quad (3.2)$$

is invertible. Furthermore, for any context m and any arm $x \in \mathcal{A}^m$,

$$\limsup_{n \rightarrow \infty} \log(n) \|x - x_m^*\|_{\bar{G}_n^{-1}}^2 \leq \frac{(\Delta_x^m)^2}{2}. \quad (3.3)$$

The proof is deferred to Appendix A.1 in the supplementary material. Intuitively, the lemma shows that any consistent policy must collect sufficient statistical evidence at confidence level $1 - 1/n$ that suboptimal arms really are suboptimal. This corresponds to ensuring that the width of an appropriate confidence interval $\sqrt{2 \log(n)} \|x - x_m^*\|_{\bar{G}_n^{-1}}$ is approximately smaller than the sub-optimality gap Δ_x^m .

Theorem 3.3 (Asymptotic Lower Bound). Under the same conditions as Lemma 3.2,

$$\liminf_{n \rightarrow \infty} \frac{R_\theta^\pi(n)}{\log(n)} \geq \mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M), \quad (3.4)$$

where $\mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M)$ is defined as the optimal value of the following optimisation problem:

$$\inf_{\alpha_{x,m} \in [0, \infty]} \sum_{m=1}^M \sum_{x \in \mathcal{A}^m} \alpha_{x,m} \Delta_x^m \quad (3.5)$$

subject to the constraint that for any context m and suboptimal arm $x \in \mathcal{A}^m$,

$$x^\top \left(\sum_{m=1}^M \sum_{x \in \mathcal{A}^m} \alpha_{x,m} x x^\top \right)^{-1} x \leq \frac{(\Delta_x^m)^2}{2}. \quad (3.6)$$

Given the result in Lemma 3.2, the proof of Theorem 3.3 follows exactly the same idea of the proof of Corollary 2 in [16] and thus is omitted here. Later on we will prove a matching upper bound in Theorem 4.3 and argue that our asymptotic lower bound is sharp.

Remark 3.4. In the above we adopt the convention that $\infty \times 0 = 0$ so that $\alpha_{x,m} \Delta_x^m = 0$ whenever $\Delta_x^m = 0$. The inverse of a matrix with infinite entries is defined by passing to the limit in the obvious way, and is not technically an inverse.

Remark 3.5. Let us denote $\{\alpha_{x,m}^*\}_{x \in \mathcal{A}^m, m \in [M]}$ as an optimal solution to the above optimisation problem. It serves as the *optimal allocation rule* for each arm such that the cumulative regret is minimized subject to the width of the confidence interval of each sub-optimal arm is small. Specifically, $\alpha_{x,m}^* \log(n)$ can be interpreted as the approximate optimal number of times arm x should be played having observed context m .

Remark 3.6. Our lower bound may also be derived from a more general bound of Ok et al. [19], since a stochastic contextual bandit can be viewed as a kind of Markov decision process. We use an alternative proof technique and the two lower bound statements have different forms. The proof is included for completeness.

Example 3.7. When $M = 1$ and $\mathcal{A}^1 = \{e_1, \dots, e_d\}$ is the standard basis vectors, the problem reduces to classical multi-armed bandit and $\mathcal{C}(\theta, \mathcal{A}^1) = \sum_{x \in \mathcal{A}^1, \Delta_x > 0} 2/\Delta_x$, which matches the well-known asymptotic lower bound by [14].

The constant $\mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M)$ depends on both the unknown parameter θ and the action sets $\mathcal{A}^1, \dots, \mathcal{A}^M$, but *not* the context distribution p . In this sense there is a certain discontinuity in the hardness measure \mathcal{C} as a function of the context distribution. More precisely, problems where $p(m)$ is arbitrarily close to zero may have different regret asymptotically than the problem obtained by removing context m entirely. Clearly as $p(m)$ tends to zero the m th context is observed with vanishingly small probability in finite time and hence the asymptotically optimal regret may not be representative of the finite-time hardness.

3.1 Sub-logarithmic regret

Our matching upper and lower bounds reveal the interesting phenomenon that if the action sets satisfy certain conditions, then sub-logarithmic regret is possible. Consider the scenario that the set of optimal arms $\{x_1^*, \dots, x_M^*\}$ spans \mathbb{R}^d ¹. Let $\Lambda \in \mathbb{R}$ be a large constant to be defined subsequently and for each context m and arm $x \in \mathcal{A}^m$, let $\alpha_{x,m}$ be 0 if $x \neq x_m^*$, and be Λ if else. Then,

$$\sum_{m=1}^M \sum_{x \in \mathcal{A}^m} \alpha_{x,m} x x^\top = \Lambda \sum_{m=1}^M x_m^* x_m^{*\top}. \quad (3.7)$$

Since the set of optimal arms spans \mathbb{R}^d it holds that for any context m and arm $x \in \mathcal{A}^m$,

$$x^\top \left(\sum_{m=1}^M x_m^* x_m^{*\top} \right)^{-1} x < \infty. \quad (3.8)$$

Combining Eq. (3.7) and Eq. (3.8),

$$\begin{aligned} x^\top \left(\sum_{m=1}^M \sum_{x \in \mathcal{A}^m} \alpha_{x,m} x x^\top \right)^{-1} x \\ = \Lambda^{-1} x^\top \left(\sum_{m=1}^M x_m^* x_m^{*\top} \right)^{-1} x. \end{aligned}$$

Hence, the constraint in Eq. (3.6) is satisfied for sufficiently large Λ . Since with this choice of $(\alpha_{x,m})$ we have $\sum_{m=1}^M \sum_{x \in \mathcal{A}^m} \alpha_{x,m} \Delta_x^m = 0$, it follows that

¹This condition is both sufficient and necessary. More precisely, sub-logarithmic regret is possible if and only if the optimal arms span the space that is spanned by all the available actions. Since in this paper we assume the action set spans \mathbb{R}^d for simplicity, the two conditions are equivalent.

$\mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M) = 0$. Therefore our upper bound will show that when the set of optimal actions $\{x_1^*, \dots, x_M^*\}$ spans \mathbb{R}^d our new algorithm satisfies

$$\liminf_{n \rightarrow \infty} \frac{R_\theta^\pi(n)}{\log(n)} = 0.$$

Remark 3.8. The choice of $\alpha_{x,m}$ above shows that when $\{x_1^*, \dots, x_M^*\}$ span \mathbb{R}^d , then an asymptotically optimal algorithm only needs to play suboptimal arms sub-logarithmically often, which means the algorithm is eventually very close to the greedy algorithm. Bastani et al. [5], Kannan et al. [13] also investigate the striking performance of greedy algorithms in contextual bandits. However, [5] assume the covariate diversity on the context distribution while [13] assume the context is artificially perturbed with noise – these assumptions make these works brittle. In addition, [5] only provide a rate-optimal algorithm while our algorithm is optimal in constants (see Theorem 4.3 for details).

As claimed in the introduction, we also prove that algorithms based on optimism can enjoy bounded regret when the set of optimal actions spans the space of all actions. The proof of the following theorem is given in Appendix B.7.

Theorem 3.9. Consider the policy π that plays optimistically by

$$X_t = \operatorname{argmax}_{x \in \mathcal{A}^{c_t}} \langle \hat{\theta}_{t-1}, x \rangle + \|x\|_{G_t^{-1} \beta_t}^{1/2}.$$

Suppose that θ is such that $\{x_1^*, \dots, x_M^*\}$ spans \mathbb{R}^d . Then, for suitable $(\beta_t)_{t=1}^n$ with $\beta_t = O(d \log(t))$, it holds that $\limsup_{n \rightarrow \infty} R_\theta^\pi(n) < \infty$.

Note, the choice of (β_t) for which the above theorem holds also guarantees the standard $\tilde{O}(d\sqrt{n})$ minimax bound for this algorithm, showing that LinUCB can adapt online to this nice case.

4 OPTIMAL ALLOCATION MATCHING

The instance-dependent asymptotic lower bound provides an optimal allocation rule. However, the optimal allocation $\{\alpha_{x,m}^*\}_{x,m}$ depends on the unknown sub-optimality gap. In this section, we present a novel matching algorithm that simultaneously estimates the unknown parameter θ using least squares and updates the allocation rule.

4.1 Algorithm

Let $N_x(t) = \sum_{s=1}^t \mathbb{I}(X_s = x)$ be the number of pulls of arm x after round t and $G_t = \sum_{s=1}^t X_s X_s^\top$. The least squares estimator is $\hat{\theta}_t = G_t^{-1} \sum_{s=1}^t X_s Y_s$. For each context m the estimated sub-optimality gap of arm $x \in \mathcal{A}^m$ is $\hat{\Delta}_x^m(t) = \max_{y \in \mathcal{A}^m} \langle y - x, \hat{\theta}_t \rangle$ and the estimated optimal arm is $\hat{x}_m^*(t) = \operatorname{argmax}_{x \in \mathcal{A}^m} \langle x, \hat{\theta}_t \rangle$. The minimum nonzero estimated gap is

$$\hat{\Delta}_{\min}(t) = \min_{m \in [M]} \min_{x \in \mathcal{A}^m, \hat{\Delta}_x^m(t) > 0} \hat{\Delta}_x^m(t).$$

Next, we define a similar optimisation problem as in (3.5) but with a different normalisation.

Definition 4.1. Let $f_{n,\delta}$ be the constant given by

$$f_{n,\delta} = 2(1 + 1/\log(n)) \log(1/\delta) + cd \log(d \log(n)), \quad (4.1)$$

where c is an absolute constant. We write $f_n = f_{n,1/n}$. For any $\tilde{\Delta} \in [0, \infty)^{|\cup_m \mathcal{A}^m|}$ define $T(\tilde{\Delta})$ as a solution of the following optimisation problem:

$$\min_{(T_x^m)_{x,m \in [0,\infty]}} \sum_{m=1}^M \sum_{x \in \mathcal{A}^m} T_x^m \tilde{\Delta}_x^m, \quad (4.2)$$

subject to

$$\|x\|_{H_T^{-1}}^2 \leq \frac{\Delta_x^2}{f_n}, \forall x \in \mathcal{A}^m, m \in [M].$$

and that $H_T = \sum_{m=1}^M \sum_{x \in \mathcal{A}^m} T_x^m x x^\top$ is invertible.

If $\tilde{\Delta}$ is an estimate of Δ , we call the solution $T(\tilde{\Delta})$ an *approximated allocation rule* in contrast to the *optimal allocation rule* defined in Remark 3.5. Our algorithm alternates between exploration and exploitation, depending on whether or not all the arms have satisfied the approximated allocation rule. We are now ready to describe the algorithm, which starts with a brief initialisation phase.

Initialisation In the first d rounds the algorithm chooses any action X_t in the action set such that X_t is not in the span of $\{X_1, \dots, X_{t-1}\}$. This is always possible by the assumption that \mathcal{A}^m spans \mathbb{R}^d for all contexts m . At the end of the initialisation phase G_t is guaranteed to be invertible.

Main phase In each round after the initialisation phase the algorithm checks if the following criterion

holds for any $x \in \mathcal{A}^{c_t}$:

$$\|x\|_{G_{t-1}^{-1}}^2 \leq \max \left\{ \frac{(\hat{\Delta}_{\min}(t-1))^2}{f_n}, \frac{(\hat{\Delta}_x^{c_t}(t-1))^2}{f_n} \right\}. \quad (4.3)$$

The algorithm exploits if Eq. (4.3) holds and explores otherwise, as explained below.

Exploitation. The algorithm exploits by taking the greedy action:

$$X_t = \operatorname{argmax}_{x \in \mathcal{A}^{c_t}} x^\top \hat{\theta}_{t-1}. \quad (4.4)$$

Exploration. The algorithm explores when Eq. (4.3) does not hold. This means that some actions have not been explored sufficiently. There are two cases to consider. First, when there exists an arm $x' \in \mathcal{A}^{c_t}$ such that

$$N_{x'}(t-1) < \min(T_{x'}^{c_t}(\hat{\Delta}(t-1)), f_n/\hat{\Delta}_{\min}^2(t-1)),$$

the algorithm then computes two actions

$$\begin{aligned} b_1 &= \operatorname{argmin}_{x \in \mathcal{A}^{c_t}} \frac{N_x(t-1)}{\min(T_x^{c_t}(\hat{\Delta}(t-1)), f_n/\hat{\Delta}_{\min}^2(t-1))} \\ b_2 &= \operatorname{argmin}_{x \in \mathcal{A}^{c_t}} N_x(t-1). \end{aligned} \quad (4.5)$$

Let $s(t)$ be the number of exploration rounds defined in Algorithm 1. If $N_{b_2}(t-1) \leq \varepsilon_t s(t)$ the algorithm plays arm $X_t = b_2$ – a form of forced exploration. Otherwise the algorithm plays arm $X_t = b_1$. Finally, rounds where an $x' \in \mathcal{A}^{c_t}$ with the required property does not exist are called *wasted*. In these rounds the algorithm acts optimistically as LinUCB [1]:

$$X_t = \operatorname{argmax}_{x \in \mathcal{A}^{c_t}} x^\top \hat{\theta}_{t-1} + \sqrt{f_{n,1/(s(t))^2}} \|x\|_{G_{t-1}^{-1}}, \quad (4.6)$$

where $f_{n,1/(s(t))^2}$ is defined in Eq. (4.1). The complete algorithm is presented in Algorithm 1.

Remark 4.2. The naive forced exploration can be improved by calculating a barycentric spanner [4] for each action set and then playing the least played action in the spanner. In normal practical setups this makes very little difference, where the forced exploration plays a limited role. For finite-time worst-case analysis, however, it may be crucial, since otherwise the regret may depend linearly on the number of actions, while using the spanner guarantees the forced exploration is sample efficient.

4.2 Asymptotic Upper Bound

Our main theorem states that Algorithm 1 is asymptotically optimal under mild assumptions.

Theorem 4.3. Suppose that $T_x^m(\Delta)$ is uniquely defined and $T_x^m(\cdot)$ is continuous at Δ for all contexts m and actions $x \in \mathcal{A}^m$. Then the policy π_{oam} proposed in Algorithm 1 with $\varepsilon_t = 1/\log(\log(t))$ satisfies

$$\limsup_{n \rightarrow \infty} \frac{R_{\theta}^{\pi_{\text{oam}}}(n)}{\log(n)} \leq \mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M). \quad (4.7)$$

Together with the asymptotic lower bound in Theorem 3.3, we can argue that optimal-allocation matching algorithm is asymptotically optimal and the lower bound in Eq. (3.4) is sharp.

Remark 4.4. The assumption that $T_x^m(\cdot)$ is continuous at Δ is used to ensure the stability of our algorithm. We prove that the uniqueness assumption actually implies continuity (Lemma C.5 in the supplementary material) and thus the continuity assumption could be omitted. There are, however, certain corner cases where uniqueness does not hold. For example when $\theta = (1, 0)^\top$, $\mathcal{A} = \{(1, 0), (0, 1), (0, -1)\}$.

4.3 Proof Sketch

The complete proof is deferred to Appendix A.2 in the supplementary material. At a high level the analysis of the optimisation-based approach consists of three parts. (1) Showing that the algorithm's estimate of the true parameter is close to the truth in finite time. (2) Showing that the algorithm subsequently samples arms approximately according to the unknown optimal allocation and (3) Showing that the greedy action when arms have been sampled sufficiently according to the optimal allocation is optimal with high probability. Existing optimisation-based algorithms suffer from dominant 'lower-order' terms because they use simple empirical means for Part (1), while here we use the data-efficient least-squares estimator.

We let $\text{Explore} = \text{F-Explore} \cup \text{UW-Explore} \cup \text{W-Explore}$ be the set of exploration rounds, decomposed into disjoint sets of forced exploration ($X_t = b_1$), unwasted exploration ($X_t = b_2$) and wasted exploration (LinUCB), and let Exploit be the set of exploitation rounds.

Algorithm 1 Optimal Allocation Matching

Input: exploration parameter ε_t , exploration counter $s(d) = 0$.

initialisation

for $t = 1$ **to** d **do**

 Observe an action set \mathcal{A}^{c_t} , pull arm X_t such that X_t is not in the span of $\{X_1, \dots, X_{t-1}\}$.

end

for $t = d + 1$ **to** n **do**

 Observe an action set \mathcal{A}^{c_t} and solve the optimisation problem (4.2) based on the estimated gap $\hat{\Delta}(t-1)$.

if $\|x\|_{G_{t-1}^{-1}}^2 \leq \max\{\frac{\hat{\Delta}_{\min}^2(t-1)}{f_n}, \frac{(\hat{\Delta}_x^{c_t}(t-1))^2}{f_n}\}, \forall x \in \mathcal{A}^{c_t}$, **then**

exploitation

 Pull arm $X_t = \arg\max_{x \in \mathcal{A}^{c_t}} x^\top \hat{\theta}_{t-1}$.

else

exploration

$s(t) = s(t-1) + 1$

if $N_x(t-1) \geq \min(T_x(\hat{\Delta}(t-1)), f_n/(\hat{\Delta}_{\min}(t-1))^2), \forall x \in \mathcal{A}^{c_t}$, **then**

 Pull arm according to LinUCB in (4.6).

else

 Calculate b_1, b_2 as in Eq. (4.5).

if $N_{b_2}(t-1) \leq \varepsilon_t s(t-1)$ **then**

 Pull arm $X_t = b_2$.

else

 Pull arm $X_t = b_1$.

end

end

end

 Update $\hat{\theta}_t, \hat{\Delta}_x^{c_t}(t), \hat{\Delta}_{\min}(t)$.

end

Regret while exploiting The criterion in Eq. (4.3) guarantees that the greedy action is optimal with high probability in exploitation rounds. To see this, note that if t is an exploitation round, then the sub-optimality gap of greedy action X_t satisfies the following with high probability:

$$\Delta_{X_t}^{c_t} \lesssim \sqrt{\frac{\log(\log(n))}{1 \vee N_{X_t}(t-1)}} < \Delta_{\min}.$$

Since the instantaneous regret either vanishes or is larger than Δ_{\min} , we have

$$\mathbb{E} \left[\sum_{t \in \text{Exploit}} \Delta_t^{c_t} \right] = o(\log(n)).$$

Regret while exploring Based on the design of our algorithm, the regret while exploring is decomposed into three terms,

$$\mathbb{E} \left[\sum_{t \in \text{Explore}} \Delta_{X_t}^{c_t} \right] = \mathbb{E} \left[\sum_{t \in \text{F-Explore}} \Delta_{X_t}^{c_t} \right] + \mathbb{E} \left[\sum_{t \in \text{W-Explore}} \Delta_{X_t}^{c_t} \right] + \mathbb{E} \left[\sum_{t \in \text{UW-Explore}} \Delta_{X_t}^{c_t} \right].$$

Shortly we argue that the regret incurred in $\text{W-Explore} \cup \text{UW-Explore}$ is at most logarithmic and hence the regret in rounds associated with forced exploration is sub-logarithmic:

$$\mathbb{E} \left[\sum_{t \in \text{F-Explore}} \Delta_{X_t}^{c_t} \right] = O(\varepsilon_n |\text{Explore}|) = o(\log(n)).$$

The regret in W-Explore is also sub-logarithmic. To see this, we first argue that $|\text{W-Explore}| = O(|\text{UW-Explore}|)$ since each context has positive probability. Combining with the fact that $|\text{UW-Explore}|$ is logarithmic in n and the regret of LinUCB is square root in time horizon,

$$\mathbb{E} \left[\sum_{t \in \text{W-Explore}} \Delta_t^{c_t} \right] = o(\log(n)).$$

The regret in UW-Explore is logarithmic in n with the asymptotically optimal constant using the definition of the optimal allocation:

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t \in \text{UW-Explore}} \Delta_t^{c_t} \right]}{\log(n)} = \mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M).$$

Of course many details have been hidden here, which are covered in detail in the supplementary material.

5 EXPERIMENTS

In this section, we first compare our proposed algorithm and LinUCB [1] on some specific problem instances to showcase their strengths and weaknesses. We examine OSSB [8] on instances with large action sets to illustrate its weakness due to not using the linear structure everywhere. Since Combes et al. [8] demonstrated that OSSB dominates the algorithm of Lattimore and Szepesvári [16], we omit this algorithm from our experiments. In the end, we include the comparison with LinTS [3]. Some additional

experiments are deferred to Appendix D in the supplementary material.

To save computation, we follow the lazy-update approach, similar to that proposed in Section 5.1 of [1]: The idea is to resolve the optimisation problem (4.2) whenever $\det(G_t)$ increases by a constant factor $(1 + \zeta)$ and in all scenarios we choose (the arbitrary value) $\zeta = 0.1$. All codes were written in Python. To solve the convex optimisation problem (4.2), we use the CVXPY library [11].

5.1 Fixed Action Set

Finite-armed linear bandits with fixed action set are a special case of linear contextual bandits. Let $d = 2$ and let the true parameter be $\theta = (1, 0)^\top$. The action set $\mathcal{A} = \{x_1, x_2, x_3\}$ is fixed and $x_1 = (1, 0)^\top$, $x_2 = (0, 1)^\top$, $x_3 = (1 - u, 5u)^\top$. We consider $u = \{0.1, 0.2\}$. By construction, x_1 is the optimal arm. From Figure 1, we observe that LinUCB suffers significantly more regret than our algorithm. The reason is that if u is very small, then x_1 and x_3 point in almost the same direction and so choosing only these arms does not provide sufficient information to quickly learn which of x_1 or x_3 is optimal. On the other hand, x_2 and x_1 point in very different directions and so choosing x_2 allows a learning agent to quickly identify that x_1 is in fact optimal. LinUCB stops pulling x_2 once it is optimistic and thus fails to find the right balance between information and reward. Our algorithm, however, takes this into consideration by tracking the optimal allocation ratios.




Figure 1: Fixed action set. The results are averaged over 100 realisations. Here and also later, the shaded areas show the standard errors.

5.2 Changing Action Set

We consider a simple but representative case when there are only two action sets \mathcal{A}^1 and \mathcal{A}^2 available.

Scenario One. In each round, \mathcal{A}^1 is drawn with probability 0.3 while \mathcal{A}^2 is drawn with probability




Figure 2: Changing action sets. The left panel is for scenario one and the right panel is for scenario two. The results are averaged over 100 realisations.

0.7. Set \mathcal{A}^1 contains $x_1^1 = (1, 0, 0)^\top$, $x_2^1 = (0, 1, 0)^\top$, and $x_3^1 = (0.9, 0.5, 0)^\top$, while set \mathcal{A}^2 contains $x_1^2 = (0, 1, 0)^\top$, $x_2^2 = (0, 0, 1)^\top$, and $x_3^2 = (0, 0.5, 0.9)^\top$. The true parameter θ is $(1, 0, 1)^\top$. From the left panel of Figure 2, we observe that LinUCB, while starts better, eventually again suffers more regret than our algorithm.

Scenario Two. In each round, \mathcal{A}^1 is drawn with probability 0.99, while \mathcal{A}^2 is drawn with probability 0.01. Set \mathcal{A}^1 contains three actions: $x_1^1 = (1, 0)^\top$, $x_2^1 = (0, 1)^\top$, $x_3^1 = (0.9, 0.5)^\top$, while set \mathcal{A}^2 contains three actions: $x_1^2 = (0, 1)^\top$, $x_2^2 = (-1, 0)^\top$, $x_3^2 = (-1, 0)$. Apparently, x_1^1 and x_2^1 are the optimal arms for each action set and they span \mathbb{R}^2 . Based on the allocation rule in Section 3.1, the algorithm is advised to pull actions x_1^1 and x_2^1 very often based on asymptotics. However, since the probability that \mathcal{A}^2 is drawn is extremely small, we are very likely to fall back to wasted exploration and use LinUCB to explore. Thus, in the short term, our algorithm will suffer from the drawback that optimistic algorithms also suffer from and what is described in Section 5.1. Although, the asymptotics will eventually “kick in”, it may take extremely long time to see the benefits of this and the algorithm’s finite-time performance will be poor. Indeed, this is seen on the right panel of Figure 2, which shows that in this case LinUCB and our algorithm are nearly indistinguishable.

5.3 Sublinear/Bounded Regret

Earlier we have argued that when the optimal arms of all action sets span \mathbb{R}^d , our algorithm achieves sub-logarithmic regret. Here, we experimentally study this interesting case. We consider $M = 2$. In each round, \mathcal{A}^1 is drawn with probability 0.8 while \mathcal{A}^2 is drawn with probability 0.2 and the true parameter θ is $(1, 0)^\top$. Set \mathcal{A}^1 contains three actions: $x_1^1 = (1, 0)^\top$, $x_2^1 = (0, 1)^\top$, $x_3^1 = (0.9, 0.5)^\top$, while set \mathcal{A}^2 con-




Figure 3: The left panel is for bounded regret and right panel is for large action space. The results are averaged over 100 realisations.

tains three actions: $x_1^2 = (0, 1)^\top$, $x_2^2 = (-1, 0)^\top$, $x_3^2 = (-1, 0)$. As discussed before, x_1^1 and x_2^1 are the optimal arms for each action set and they span \mathbb{R}^2 . The results are shown in the left subpanel of Figure 3. The regret of our algorithm appears to have stopped growing after a short period of increase. In line with Theorem 3.9, LinUCB is seen to achieve bounded regret in this problem.

5.4 Large, Fixed Action Set

We let $d = 2$ and $\theta = (1, 0)^\top$. We generate 100 uniformly distributed on the d -dimensional unit sphere (fixed action set). The results are shown in the right subfigure of Figure 3. When the action space is large, OSSB suffers significantly large regret and becomes unstable due to not using the linear structure everywhere. The regret of (the theoretically justified version of) LinTS is also very large due to the unnecessary variance factor required by its theory.

6 DISCUSSION

We presented a new optimisation-based algorithm for linear contextual bandits that is asymptotically optimal and adapts to both the action sets and unknown parameter. The new algorithm enjoys sub-logarithmic regret when the collection of optimal actions spans \mathbb{R}^d , a property that we also prove for optimism-based approaches. There are many open questions. A natural starting point is to prove near-minimax optimality of the new algorithm, possibly with minor modifications. Our work also highlights the dangers of focusing too intensely on asymptotics, which for contextual bandits hide completely the dependence on the context distribution. This motivates the intriguing challenge to understand the finite-time instance-dependent regret. Another open direction is to consider the asymptotics when the context space is continuous, which has not seen any attention.

Acknowledgements

Csaba Szepesvári gratefully acknowledges funding from the Canada CIFAR AI Chairs Program, Amii and NSERC.

References

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- [2] Deepak Agarwal, Bee-Chung Chen, Pradheep Elango, Nitin Motgi, Seung-Taek Park, Raghu Ramakrishnan, Scott Roy, and Joe Zachariah. Online models for content optimization. In *Advances in Neural Information Processing Systems*, pages 17–24, 2009.
- [3] Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
- [4] Baruch Awerbuch and Robert Kleinberg. Online linear optimization and adaptive routing. *Journal of Computer and System Sciences*, 74(1):97–114, 2008.
- [5] Hamsa Bastani, Mohsen Bayati, and Khashayar Khosravi. Mostly exploration-free algorithms for contextual bandits. *arXiv preprint arXiv:1704.09011*, 2017.
- [6] H. P. Chan and T. L. Lai. Sequential generalized likelihood ratios and adaptive treatment allocation for optimal sequential selection. *Sequential Analysis*, 25:179–201, 2006.
- [7] Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.
- [8] Richard Combes, Stefan Magureanu, and Alexandre Proutiere. Minimal exploration in structured stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 1763–1771, 2017.
- [9] Varsha Dani, Thomas P. Hayes, and Sham M. Kakade. Stochastic linear optimization under bandit feedback. In Rocco A. Servedio and Tong Zhang, editors, *21st Annual Conference on Learning Theory - COLT*, pages 355–366, 2008.
- [10] Rémy Degenne, Wouter M Koolen, and Pierre Ménard. Non-asymptotic pure exploration by solving games. *arXiv preprint arXiv:1906.10431*, 2019.
- [11] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- [12] A. Garivier and E. Kaufmann. Optimal best arm identification with fixed confidence. In V. Feldman, A. Rakhlin, and O. Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 998–1027, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR.
- [13] Sampath Kannan, Jamie H Morgenstern, Aaron Roth, Bo Waggoner, and Zhiwei Steven Wu. A smoothed analysis of the greedy algorithm for the linear contextual bandit problem. In *Advances in Neural Information Processing Systems*, pages 2227–2236, 2018.
- [14] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [15] John Langford and Tong Zhang. The epoch-greedy algorithm for contextual multi-armed bandits. In *Proceedings of the 20th International Conference on Neural Information Processing Systems*, pages 817–824, 2007.
- [16] Tor Lattimore and Csaba Szepesvári. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Artificial Intelligence and Statistics*, pages 728–737, 2017.
- [17] Tor Lattimore and Csaba Szepesvári. Bandit algorithms. *preprint*, 2019.
- [18] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web, WWW ’10*, pages 661–670, 2010.
- [19] Jungseul Ok, Alexandre Proutiere, and Damiianos Tranos. Exploration in structured rein-

- forcement learning. In *Advances in Neural Information Processing Systems*, pages 8874–8882, 2018.
- [20] Paat Rusmevichientong and John N Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.
- [21] Alexandre B. Tsybakov. *Introduction to Non-parametric Estimation*. Springer, 1st edition, 2008. ISBN 0387790519, 9780387790510.
- [22] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.

Supplement to “Adaptive Exploration in Linear Contextual Bandit”

In Section A, we provide main proofs for asymptotic lower bound and upper bound. In Section B, we prove several main lemmas. In Section C, some supporting lemmas are presented for the sake of completeness.

A Proofs of Asymptotic Lower and Upper Bounds

First of all, we define the sub-optimal action set as $\mathcal{A}_-^m = \mathcal{A}^m \setminus \{x : \Delta_x^m = 0\}$ and denote $\mathcal{A} = \cup_{m=1}^M \mathcal{A}^m$ and $\mathcal{A}_- = \cup_{m=1}^M \mathcal{A}_-^m$.

A.1 Proof of Lemma 3.2

The proof idea follows if \bar{G}_n is not sufficiently large in every direction, then some alternative parameters are not sufficiently identifiable.

Step One. We fix a consistent policy π and fix a context $m \in [M]$ as well as a sub-optimal arm $x \in \mathcal{A}_-^m$. Consider another parameter $\tilde{\theta} \in \mathbb{R}^d$ such that it is close to θ but x_m^* is not the optimal arm in bandit $\tilde{\theta}$ for action set \mathcal{A}^m . Specifically, we construct

$$\tilde{\theta} = \theta + \frac{H(x - x_m^*)}{\|x - x_m^*\|_H^2} (\Delta_x^m + \varepsilon),$$

where $H \in \mathbb{R}^{d \times d}$ is some positive semi-definite matrix and $\varepsilon > 0$ is some absolute constant that will be specified later. Since the sub-optimality gap $\tilde{\Delta}_{x_m^*}^m$ satisfies

$$\langle x - x_m^*, \tilde{\theta} \rangle = \langle x - x_m^*, \theta \rangle + \Delta_x^m + \varepsilon = \varepsilon > 0, \quad (\text{A.1})$$

it ensures that x_m^* is ε -suboptimal in bandit $\tilde{\theta}$.

We define $T_x(n) = \sum_{t=1}^n \mathbb{I}(X_t = x)$ and let \mathbb{P} and $\tilde{\mathbb{P}}$ be the measures on the sequence of outcomes $(X_1, Y_1, \dots, X_n, Y_n)$ induced by the interaction between the policy and the bandit θ and $\tilde{\theta}$ respectively. By the definition of \bar{G}_n in (3.2), we have

$$\begin{aligned} \frac{1}{2} \|\theta - \tilde{\theta}\|_{\bar{G}_n}^2 &= \frac{1}{2} (\theta - \tilde{\theta})^\top \bar{G}_n (\theta - \tilde{\theta}) \\ &= \frac{1}{2} (\theta - \tilde{\theta})^\top \mathbb{E} \left[\sum_{x \in \mathcal{A}} T_x(n) x x^\top \right] (\theta - \tilde{\theta}) \\ &= \frac{1}{2} \sum_{x \in \mathcal{A}} \mathbb{E} [T_x(n)] \langle x, \theta - \tilde{\theta} \rangle^2. \end{aligned}$$

Applying the Bretagnolle-Huber inequality in Lemma C.1 and divergence decomposition lemma in Lemma C.2, it holds that for any event \mathcal{D} ,

$$\frac{1}{2} \|\theta - \tilde{\theta}\|_{\bar{G}_n}^2 = \text{KL}(\mathbb{P}, \tilde{\mathbb{P}}) \geq \log \left(\frac{1}{2(\mathbb{P}(\mathcal{D}) + \tilde{\mathbb{P}}(\mathcal{D}^c))} \right). \quad (\text{A.2})$$

Step Two. In the following, we start to derive a lower bound of $R_\theta^\pi(n)$,

$$\begin{aligned} R_\theta^\pi(n) &= \mathbb{E} \left[\sum_{t=1}^n \langle x_{c_t}^* - X_t, \theta \rangle \right] = \mathbb{E} \left[\sum_{m=1}^M \sum_{t:c_t=m} \langle x_m^* - X_t, \theta \rangle \right] \\ &\geq \mathbb{E} \left[\sum_{t:c_t=m} \langle x_m^* - X_t, \theta \rangle \right] = \mathbb{E} \left[\sum_{t:c_t=m} \Delta_{X_t}^m \right] \\ &\geq \Delta_{\min} \mathbb{E} \left[\sum_{t:c_t=m} \mathbb{I}(X_t \neq x_m^*) \right] = \Delta_{\min} \mathbb{E} \left[\sum_{t=1}^n \mathbb{I}(c_t = m) - \sum_{t=1}^n \mathbb{I}(c_t = m) \mathbb{I}(X_t = x_m^*) \right], \end{aligned}$$

where the first inequality comes from the fact that $\langle x_m^* - X_t, \theta \rangle \geq 0$ for all $m \in [M]$. Define the event \mathcal{D} as follows,

$$\mathcal{D} = \left\{ \sum_{t=1}^n \mathbb{I}(c_t = m) \mathbb{I}(X_t = x_m^*) \leq \frac{1}{2} \sum_{t=1}^n \mathbb{I}(c_t = m) \right\}. \quad (\text{A.3})$$

When event \mathcal{D} holds, we will only pull at most half of total rounds for the optimal action of action set m . Then it holds that

$$\begin{aligned} R_\theta^\pi(n) &\geq \Delta_{\min} \mathbb{E} \left[\left(\sum_{t=1}^n \mathbb{I}(c_t = m) - \sum_{t=1}^n \mathbb{I}(c_t = m) \mathbb{I}(X_t = x_m^*) \right) \mathbb{I}(\mathcal{D}) \right] \\ &\geq \Delta_{\min} \mathbb{E} \left[\frac{1}{2} \sum_{t=1}^n \mathbb{I}(c_t = m) \mathbb{I}(\mathcal{D}) \right]. \end{aligned}$$

Define another event \mathcal{B} as follows,

$$\mathcal{B} = \left\{ \frac{1}{2} \sum_{t=1}^n \mathbb{I}(c_t = m) \geq \frac{np_m}{2} - \delta/2 \right\}, \quad (\text{A.4})$$

where $\delta > 0$ will be chosen later and p_m is the probability that the environment picks context m . From the definition of c_t , we have $\mathbb{E}[\sum_{t=1}^n \mathbb{I}(c_t = m)] = np_m$. By the standard Hoeffding's inequality [22], it holds that

$$\mathbb{P} \left(\frac{1}{2} \sum_{t=1}^n \mathbb{I}(c_t = m) - \frac{np_m}{2} \geq -\frac{\delta}{2} \right) \geq 1 - \exp(-\frac{2\delta^2}{n}),$$

which implies

$$\mathbb{P}(\mathcal{B}^c) \leq \exp(-2\delta^2/n).$$

By the definition of events \mathcal{D}, \mathcal{B} in (A.3), (A.4), we have

$$\begin{aligned} R_\theta^\pi(n) &\geq \Delta_{\min} \mathbb{E} \left[\frac{1}{2} \sum_{t=1}^n \mathbb{I}(c_t = m) \mathbb{I}(\mathcal{D}) \mathbb{I}(\mathcal{B}) \right] \\ &\geq \Delta_{\min} \mathbb{E} \left[\left(\frac{1}{2} np_m - \frac{\delta}{2} \right) \mathbb{I}(\mathcal{D}) \mathbb{I}(\mathcal{B}) \right] \\ &= \Delta_{\min} \left(\frac{1}{2} np_m - \frac{\delta}{2} \right) \mathbb{P}(\mathcal{D} \cap \mathcal{B}) \\ &\geq \Delta_{\min} \left(\frac{1}{2} np_m - \frac{\delta}{2} \right) (\mathbb{P}(\mathcal{D}) - \mathbb{P}(\mathcal{B}^c)). \end{aligned}$$

Letting $\delta = np_m/2$, we have

$$R_\theta^\pi(n) \geq \Delta_{\min} \frac{np_m}{4} \left(\mathbb{P}(\mathcal{D}) - \exp(-\frac{np_m^2}{2}) \right). \quad (\text{A.5})$$

On the other hand, we let $\tilde{\mathbb{E}}$ is taken with respect to probability measures $\tilde{\mathbb{P}}$. Then $R_\theta^\pi(n)$ can be lower bounded as follows,

$$\begin{aligned} R_\theta^\pi(n) &= \tilde{\mathbb{E}} \left[\sum_{m=1}^M \sum_{t=1}^n \mathbb{I}(c_t = m) \tilde{\Delta}_{X_t}^m \right] \\ &\geq \tilde{\mathbb{E}} \left[\sum_{t=1}^n \mathbb{I}(c_t = m) \mathbb{I}(X_t = x_m^*) \right] \tilde{\Delta}_{x_m^*}^m, \end{aligned}$$

where we throw out all the sub-optimality gap terms except $\tilde{\Delta}_{x_m^*}^m$. Using the fact that $\tilde{\Delta}_{x_m^*}^m$ is ε -suboptimal, it holds that

$$\begin{aligned}
R_\theta^\pi(n) &\geq \varepsilon \tilde{\mathbb{E}} \left[\left(\sum_{t=1}^n \mathbb{I}(c_t = m) \mathbb{I}(X_t = x_m^*) \right) \mathbb{I}(\mathcal{D}^c) \right] \\
&> \varepsilon \tilde{\mathbb{E}} \left[\frac{1}{2} \sum_{t=1}^n \mathbb{I}(c_t = m) \mathbb{I}(\mathcal{D}^c) \right] \\
&\geq \varepsilon \tilde{\mathbb{E}} \left[\frac{1}{2} \sum_{t=1}^n \mathbb{I}(c_t = m) \mathbb{I}(\mathcal{D}^c) \mathbb{I}(\mathcal{B}) \right] \\
&\geq \varepsilon \left(\frac{np_m}{2} - \frac{\delta}{2} \right) \tilde{\mathbb{P}}(\mathcal{D}^c \cap \mathcal{B}) \\
&\geq \varepsilon \left(\frac{np_m}{2} - \frac{\delta}{2} \right) (\tilde{\mathbb{P}}(\mathcal{D}^c) - \tilde{\mathbb{P}}(\mathcal{B}^c)) \\
&\geq \varepsilon \left(\frac{np_m}{2} - \frac{\delta}{2} \right) (\tilde{\mathbb{P}}(\mathcal{D}^c) - \exp(-\frac{2\delta^2}{n})) \\
&= \varepsilon \frac{np_m}{4} \tilde{\mathbb{P}}(\mathcal{D}^c) - \varepsilon \frac{np_m}{4} \exp(-\frac{np_m^2}{2}).
\end{aligned} \tag{A.6}$$

Now we have derived the lower bounds (A.5)(A.6) for $R_\theta^\pi(n)$, $R_\theta^\pi(n)$ respectively.

Step Three. Combining the lower bounds of $R_\theta^\pi(n)$ and $R_\theta^\pi(n)$ together, it holds that

$$R_\theta^\pi(n) + R_\theta^\pi(n) \geq \frac{np_m}{4} \left(\mathbb{P}(\mathcal{D}) \Delta_{\min} + \tilde{\mathbb{P}}(\mathcal{D}^c) \varepsilon \right) - \frac{np_m}{4} \exp(-\frac{np_m^2}{2}) (\varepsilon + \Delta_{\min}).$$

Letting $\varepsilon \leq \Delta_{\min}$, we have

$$R_\theta^\pi(n) + R_\theta^\pi(n) \geq \varepsilon \frac{np_m}{4} \left(\mathbb{P}(\mathcal{D}) + \tilde{\mathbb{P}}(\mathcal{D}^c) \right) - \frac{np_m}{4} \exp(-\frac{np_m^2}{2}) 2\Delta_{\min}.$$

This implies

$$\frac{R_\theta^\pi(n) + R_\theta^\pi(n)}{\varepsilon np_m/4} + \frac{1}{\varepsilon} \exp(-\frac{np_m^2}{2}) 2\Delta_{\min} \geq \mathbb{P}(\mathcal{D}) + \tilde{\mathbb{P}}(\mathcal{D}^c). \tag{A.7}$$

Plugging (A.7) into (A.2), we have

$$\begin{aligned}
\frac{1}{2} \|\theta - \tilde{\theta}\|_{G_n}^2 &\geq \log \left(\frac{1}{2(\mathbb{P}(\mathcal{D}) + \tilde{\mathbb{P}}(\mathcal{D}^c))} \right) \\
&\geq \log \left(\frac{1}{\frac{R_\theta^\pi(n) + R_\theta^\pi(n)}{\varepsilon np_m/8} + \frac{1}{\varepsilon} \exp(-\frac{np_m^2}{2}) 4\Delta_{\min}} \right) \\
&= \log \left(\frac{n}{\frac{R_\theta^\pi(n) + R_\theta^\pi(n)}{\varepsilon p_m/8} + \frac{n}{\varepsilon} \exp(-\frac{np_m^2}{2}) 4\Delta_{\min}} \right) \\
&= \log(n) - \log \left(\frac{R_\theta^\pi(n) + R_\theta^\pi(n)}{\varepsilon p_m/8} + \frac{4n}{\varepsilon} \exp(-\frac{np_m^2}{2}) \Delta_{\min} \right).
\end{aligned}$$

Dividing by $\log(n)$ for both sides, we reach

$$\frac{\|\theta - \tilde{\theta}\|_{G_n}^2}{2 \log(n)} \geq 1 - \frac{\log \left(\frac{R_\theta^\pi(n) + R_\theta^\pi(n)}{\varepsilon p_m/8} + \frac{4n}{\varepsilon} \exp(-\frac{np_m^2}{2}) \Delta_{\min} \right)}{\log(n)}.$$

From the definition of consistent policies (3.1), it holds that

$$\limsup_{n \rightarrow \infty} \frac{\log(R_\theta^\pi(n) + R_\theta^\pi(n))}{\log(n)} \leq 0.$$

In addition, by using the fact that $\lim_{n \rightarrow \infty} n \exp(-n) = 0$, it follows that

$$\liminf_{n \rightarrow \infty} \frac{\|\theta - \tilde{\theta}\|_{\bar{G}_n}^2}{2 \log(n)} \geq 1. \quad (\text{A.8})$$

Step Four. Let's denote

$$\rho_n(H) = \frac{\|x - x_m^*\|_{\bar{G}_n^{-1}}^2 \|x - x_m^*\|_{H \bar{G}_n H}^2}{\|x - x_m^*\|_H^4}.$$

Then we can rewrite

$$\frac{1}{2} \|\theta - \tilde{\theta}\|_{\bar{G}_n}^2 = \frac{(\Delta_x^m + \varepsilon)^2}{2 \|x - x_m^*\|_{\bar{G}_n^{-1}}^2} \rho_n(H).$$

Plugging this into (A.8) and letting ε to zero, we see that

$$\liminf_{n \rightarrow \infty} \frac{\rho_n(H)}{\|x - x_m^*\|_{\bar{G}_n^{-1}}^2 \log(n)} \geq \frac{2}{(\Delta_x^m)^2}. \quad (\text{A.9})$$

Now, we consider the following lemma, extracted from the proof of Theorem 25.1 of the book by [17]. The detailed proof is deferred to Section B.6.

Lemma A.1. Let $\{G_n\}_{n \geq 0}$ be a sequence of $d \times d$ positive definite matrices, $s \in \mathbb{R}^d$. For H positive semi-definite $d \times d$ matrix such that $\|s\|_H > 0$ and $n \geq 0$, let $\rho_n(H) = \frac{\|s\|_{G_n^{-1}}^2 \|s\|_{H G_n H}^2}{\|s\|_H^4}$. Assume that $\liminf_{n \rightarrow \infty} \frac{\lambda_{\min}(G_n)}{\log(n)} > 0$ and that for some $c > 0$,

$$\liminf_{n \rightarrow \infty} \frac{\rho_n(H)}{\|s\|_{G_n^{-1}}^2 \log(n)} \geq c. \quad (\text{A.10})$$

Then, $\limsup_{n \rightarrow \infty} \log(n) \|s\|_{G_n^{-1}}^2 \leq 1/c$.

The proof of $\liminf_{n \rightarrow \infty} \frac{\lambda_{\min}(G_n)}{\log(n)} > 0$ could refer Appendix C in [16]. Clearly, this lemma with $G_n = \bar{G}_n$, $c = 2/(\Delta_x^m)^2$, $H = \lim_{n \rightarrow \infty} \bar{G}_n^{-1}/\|\bar{G}_n^{-1}\|$ and $s = x - x_m^*$ gives the desired statement. ■

A.2 Proof of Theorem 4.3: Asymptotic Upper Bound

We write $\Delta_{\max} = \max_{x,m} \Delta_x^m$ and abbreviate $R(n) = R_\theta^\pi(n)$. From the design of the initialisation, G_t is guaranteed to be invertible since each \mathcal{A}^m is assumed to span \mathbb{R}^d . The regret during the initialisation is at most $d\Delta_{\max} \approx o(\log(n))$ and thus we ignore the regret during initialisation in the following.

First, we introduce a refined concentration inequality for the least square estimator constructed by adaptive data. The proof could refer to the proof of Theorem 8 in [16].

Lemma A.2. Suppose for $t \geq d$, G_t is invertible. For any $\delta \in (0, 1)$, we have

$$\mathbb{P}\left(\exists t \geq d, \exists x \in \mathcal{A}, \text{ such that } |\langle x, \hat{\theta}_t \rangle - \langle x, \theta \rangle| \geq \|x\|_{G_t^{-1}} f_{n,\delta}^{1/2}\right) \leq \delta,$$

and

$$f_{n,\delta} = 2\left(1 + \frac{1}{\log(n)}\right) \log(1/\delta) + cd \log(d \log(n)), \quad (\text{A.11})$$

where $c > 0$ is some universal constant. We write $f_n = f_{n,1/n}$ for short.

Let us define the event \mathcal{B}_t as follows

$$\mathcal{B}_t = \left\{ \exists t \geq d, \exists x \in \mathcal{A}, \text{ such that } |x^\top \hat{\theta}_t - x^\top \theta| \geq \|x\|_{G_t^{-1}} f_n^{1/2} \right\}. \quad (\text{A.12})$$

From Lemma A.2, we have $\mathbb{P}(\mathcal{B}_t) \leq 1/n$ by choosing $\delta = 1/n$. We decompose the cumulative regret with respect to event \mathcal{B}_t as follows,

$$\begin{aligned} R(n) &= \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x^{c_t} \mathbb{I}(X_t = x) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x^{c_t} \mathbb{I}(X_t = x, \mathcal{B}_t) \right] + \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x^{c_t} \mathbb{I}(X_t = x, \mathcal{B}_t^c) \right]. \end{aligned} \quad (\text{A.13})$$

To bound the first term in (A.13), we observe that

$$\begin{aligned} & \limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x^{c_t} \mathbb{I}(X_t = x, \mathcal{B}_t) \right]}{\log(n)} \\ &= \limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t=1}^n \Delta_{X_t}^{c_t} \mathbb{I}(\mathcal{B}_t) \right]}{\log(n)} \leq \limsup_{n \rightarrow \infty} \frac{\Delta_{\max} \sum_{t=1}^n \mathbb{P}(\mathcal{B}_t)}{\log(n)} = \limsup_{n \rightarrow \infty} \frac{\Delta_{\max} \sum_{t=1}^n \frac{1}{n}}{\log(n)} \\ &= \limsup_{n \rightarrow \infty} \frac{\Delta_{\max}}{\log(n)} = 0. \end{aligned} \quad (\text{A.14})$$

To bound the second term in (A.13), we define the event \mathcal{D}_{t,c_t} as follows,

$$\mathcal{D}_{t,c_t} = \left\{ \forall x \in \mathcal{A}^{c_t}, \|x\|_{G_t^{-1}}^2 \leq \max \left\{ \frac{(\hat{\Delta}_{\min}(t))^2}{f_n}, \frac{(\Delta_x^{c_t}(t))^2}{f_n} \right\} \right\}. \quad (\text{A.15})$$

When \mathcal{D}_{t,c_t} occurs, the algorithm exploits at round t . Otherwise, the algorithm explores at round t . We decompose the second term in (A.13) as the exploitation regret and exploration regret:

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x^{c_t} \mathbb{I}(X_t = x, \mathcal{B}_t^c) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x^{c_t} \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}) \right] + \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x^{c_t} \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c) \right]. \end{aligned} \quad (\text{A.16})$$

We bound those two terms in Lemmas A.3-A.4 respectively.

Lemma A.3. The exploitation regret satisfies

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}) \right]}{\log(n)} = 0 \quad (\text{A.17})$$

Lemma A.4. The exploration regret satisfies

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c) \right]}{\log(n)} \leq \mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M), \quad (\text{A.18})$$

where $\mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M)$ is defined in Theorem 3.3.

Combining Lemmas A.3-A.4 together, we reach our conclusion. ■

B Proofs of Several lemmas

B.1 Proof of Lemma A.3: Exploitation Regret

When \mathcal{B}_t^c defined in (A.12) occurs, we have

$$\max_{x \in \mathcal{A}} |\langle \hat{\theta}_t - \theta, x \rangle| \leq \|x\|_{G_t^{-1}} f_n^{1/2}. \quad (\text{B.1})$$

When $\mathcal{D}_{t,m}$ defined in (A.15) occurs, we have

$$\|x\|_{G_t^{-1}}^2 \leq \max \left\{ \frac{\hat{\Delta}_{\min}^2(t)}{f_n}, \frac{(\hat{\Delta}_x^m(t))^2}{f_n} \right\} = \frac{(\hat{\Delta}_x^m(t))^2}{f_n}, \quad (\text{B.2})$$

holds for any action $x \in \mathcal{A}^m$ and $\hat{\Delta}_x^m(t) > 0$. If $x_m^* = \hat{x}_m^*(t)$, there is no regret occurred. Otherwise, putting (B.1) and (B.2) together with the optimal action x_m^* , it holds that

$$|\langle \hat{\theta}_t - \theta, x_m^* \rangle| \leq \|x_m^*\|_{G_t^{-1}} f_n^{1/2} \leq \hat{\Delta}_{x_m^*}^m(t). \quad (\text{B.3})$$

We decompose the sub-optimality gap of $\hat{x}_m^*(t)$ as follows,

$$\begin{aligned} & \langle x_m^*, \theta \rangle - \langle \hat{x}_m^*(t), \theta \rangle \\ &= \langle x_m^*, \theta - \hat{\theta}_t \rangle + \langle x_m^*, \hat{\theta}_t \rangle - \langle \hat{x}_m^*(t), \theta - \hat{\theta}_t \rangle - \langle \hat{x}_m^*(t), \hat{\theta}_t \rangle \\ &= \langle x_m^*, \theta - \hat{\theta}_t \rangle - \hat{\Delta}_{x_m^*}^m(t) + \langle \hat{x}_m^*(t), \hat{\theta}_t - \theta \rangle \\ &\leq \langle \hat{x}_m^*(t), \hat{\theta}_t - \theta \rangle. \end{aligned} \quad (\text{B.4})$$

For each $x \in \mathcal{A}$, we define

$$\tau_x = \min \left\{ N : \forall t \geq d, \mathcal{D}_{t,c_t} \text{ occurs}, N_x(t) \geq N, \text{ implies } |\langle \hat{\theta}_t - \theta, x \rangle| \leq \frac{\Delta_{\min}}{2} \right\}. \quad (\text{B.5})$$

When $N_{\hat{x}_m^*(t)}(t) \geq \tau_{\hat{x}_m^*(t)}$, it holds that

$$|\langle \hat{\theta}_t - \theta, \hat{x}_m^*(t) \rangle| \leq \frac{\Delta_{\min}}{2}.$$

Together with (B.4), we have

$$\langle x_m^*, \theta \rangle - \langle \hat{x}_m^*(t), \theta \rangle \leq \frac{\Delta_{\min}}{2}.$$

Combining this with the fact that the instantaneous regret either vanishes or is larger than Δ_{\min} , it indicates $x_m^* = \hat{x}_m^*(t)$. Therefore, we can decompose the exploitation regret with respect to event $\{N_{\hat{x}_m^*(t)}(t) \geq \tau_{\hat{x}_m^*(t)}\}$ as follows,

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_{-}^{c_t}} \Delta_x^{c_t} \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}) \right] \\ &\leq \mathbb{E} \left[\sum_{m=1}^M \sum_{t=1}^n \sum_{x \in \mathcal{A}_{-}^m} \Delta_x^m \mathbb{I} \left(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,m}, N_{\hat{x}_m^*(t)}(t) \geq \tau_{\hat{x}_m^*(t)} \right) \right] \\ &+ \mathbb{E} \left[\sum_{m=1}^M \sum_{t=1}^n \sum_{x \in \mathcal{A}_{-}^m} \Delta_x^m \mathbb{I} \left(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,m}, N_{\hat{x}_m^*(t)}(t) < \tau_{\hat{x}_m^*(t)} \right) \right]. \end{aligned} \quad (\text{B.6})$$

During exploiting the algorithm always executes the greedy action. When $x_m^* = \hat{x}_m^*(t)$ the first term in (B.6) results in no regret. For the second term in (B.6), we have

$$\begin{aligned}
& \mathbb{E} \left[\sum_{m=1}^M \sum_{t=1}^n \sum_{x \in \mathcal{A}_-^m} \Delta_x^m \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,m}, N_{\hat{x}_m^*(t)} < \tau_{\hat{x}_m^*(t)}) \right] \\
& \leq \mathbb{E} \left[\sum_{m=1}^M \sum_{t=1}^n \mathbb{I}(\mathcal{B}_t^c, \mathcal{D}_{t,m}, N_{\hat{x}_m^*(t)}(t) < \tau_{\hat{x}_m^*(t)}) \right] \Delta_{\max} \\
& \leq \sum_{m=1}^M \sum_{x \in \mathcal{A}} \mathbb{E}(\tau_x) \Delta_{\max} \leq \sum_{x \in \mathcal{A}} \mathbb{E}[\tau_x] \Delta_{\max}.
\end{aligned} \tag{B.7}$$

It remains to bound $\mathbb{E}[\tau_x]$ for any $x \in \mathcal{A}$. Let

$$\Lambda = \min \left\{ \lambda \geq 1 : \forall t \geq d, |\langle \hat{\theta}_t - \theta, x \rangle| \leq \|x\|_{G_t^{-1}} f_{n,1/\lambda}^{1/2} \right\}.$$

From the definition of τ_x in (B.5), we have

$$\tau_x \leq \max \left\{ N : (f_{n,1/\lambda}/N)^{1/2} \geq \frac{\Delta_{\min}}{2} \right\},$$

which implies $\tau_x \leq 4f_{n,1/\lambda}/\Delta_{\min}^2$. From Lemma A.2, we know that $\mathbb{P}(\Lambda \geq \lambda) \leq 1/\lambda$, which implies $\mathbb{E}[\log \Lambda] \leq 1$. Overall,

$$\mathbb{E}[\tau_x] \leq \frac{4\mathbb{E}[f_\Lambda]}{\Delta_{\min}^2} \leq \frac{8(1 + 1/\log(n)) + 4cd \log(d \log(n))}{\Delta_{\min}^2}. \tag{B.8}$$

Combining (B.6)-(B.8) together, we reach

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}^{c_t}} \Delta_x \mathbb{I}(x_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}) \right]}{\log(n)} \\
& \leq \limsup_{n \rightarrow \infty} \frac{|\mathcal{A}| \Delta_{\max} (8(1 + 1/\log(n)) + 4cd \log(d \log(n)))}{\Delta_{\min}^2 \log(n)} = 0.
\end{aligned} \tag{B.9}$$

This ends the proof. ■

B.2 Proof of Lemma A.4: Exploration Regret

If all the actions $x \in \mathcal{A}$ satisfy

$$N_x(t) \geq \min \left\{ f_n / \hat{\Delta}_{\min}^2(t), T_x(\hat{\Delta}(t)) \right\}, \tag{B.10}$$

the following holds using Lemma C.4,

$$\|x\|_{G_t^{-1}}^2 \leq \max \left\{ \frac{\hat{\Delta}_{\min}^2(t)}{f_n}, \frac{(\hat{\Delta}_x^{c_t}(t))^2}{f_n} \right\}, \text{ for any } x \in \mathcal{A}.$$

In other words, this implies if there exists an action x such that (B.10) does not hold, e.g. \mathcal{D}_{t,c_t}^c occurs, there must exist an action $x' \in \mathcal{A}$ (x and x' may not be the identical) satisfying

$$N_{x'}(t) \leq \min \left\{ f_t / \hat{\Delta}_{\min}^2(t), T_{x'}(\hat{\Delta}(t)) \right\}.$$

Based on the criterion in Algorithm 1, we should explore. However, if x' does not belong to \mathcal{A}^{c_t} and all the actions within \mathcal{A}^{c_t} have been explored sufficiently according to the approximation optimal allocation, this

exploration is interpreted as “wasted”. To alleviate the regret of the wasted exploration, the algorithm acts optimistically as LinUCB.

Let’s define a set that records the index of action sets that has not been fully explored until round t ,

$$\mathcal{M}_t = \left\{ m : \exists x \in \mathcal{A}^m, N_x(t) \leq \min\{f_n/\widehat{\Delta}_{\min}^2(t), T_x(\widehat{\Delta}(t))\} \right\}. \quad (\text{B.11})$$

When \mathcal{D}_{t,c_t}^c occurs, it means that $\mathcal{M}_t \neq \emptyset$. If \mathcal{D}_{t,c_t}^c occurs but c_t does not belong to \mathcal{M}_t , the algorithm suffers a wasted exploration. We decompose the exploration regret according to the fact if c_t belongs to \mathcal{M}_t ,

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c) \right] \\ = & \underbrace{\mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c, c_t \in \mathcal{M}_t) \right]}_{R_{\text{ue}}: \text{unwasted exploration}} \\ & + \underbrace{\mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c, c_t \notin \mathcal{M}_t) \right]}_{R_{\text{we}}: \text{wasted exploration}}. \end{aligned} \quad (\text{B.12})$$

We will bound the unwasted exploration regret and wasted exploration regret in the following two lemmas respectively.

Lemma B.1. The regret during the unwasted explorations satisfies

$$\limsup_{n \rightarrow \infty} \frac{R_{\text{ue}}}{\log(n)} \leq \mathcal{C}(\theta, \mathcal{A}_1, \dots, \mathcal{A}_M). \quad (\text{B.13})$$

The detailed proof is deferred to Section B.3.

Lemma B.2. The regret during the wasted explorations satisfies

$$\limsup_{n \rightarrow \infty} \frac{R_{\text{we}}}{\log(n)} = 0. \quad (\text{B.14})$$

The detailed proof is deferred to Section B.5.

Putting (B.12)-(B.14) together, we reach

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c) \right]}{\log(n)} \leq \mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M),$$

which ends the proof. ■

B.3 Proof of Lemma B.1: Unwasted Exploration

First, we derive a lower bound for each $N_x(t)$ during the unwasted exploration. Denote $s(t)$ as the number of rounds for unwasted explorations until round t . Indeed, forced exploration can guarantee a lower bound for $N_x(t)$: $\min_{x \in \mathcal{A}} N_x(t) \geq \varepsilon_t s(t)/2$. We prove this by the contradiction argument. Assume this is not true. There may exist $s(t)/2$ rounds $\{t_1, \dots, t_{s(t)/2}\} \subset \{1, \dots, t\}$ such that $\min_{x \in \mathcal{A}} N_x(t) \leq \varepsilon_t s(t)$. After $|\mathcal{A}|$

such rounds, we have $\min_x N_x(t)$ is incremented by at least 1 which implies $\min_x N_x(t) \geq s(t)/(2|\mathcal{A}|)$. If $\varepsilon_t \leq 1/|\mathcal{A}|$, it leads to the contradiction. This is satisfied when t is large since $\varepsilon_t = 1/\log(\log t)$.

Second, we set $\beta_n = 1/\log(\log(n))$ and define

$$\zeta = \min \left\{ s : \forall t \geq s, \forall x \in \mathcal{A}, \text{ such that } |\langle x, \hat{\theta}_t \rangle - \langle x, \theta \rangle| \leq \beta_n \right\}. \quad (\text{B.15})$$

Then we decompose the regret during unwasted explorations with respect to event $\{s(t) \geq \zeta\}$ as follows,

$$\begin{aligned} R_{\text{ue}} &= \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_{-}^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c, c_t \in \mathcal{M}_t) \right] \\ &= \underbrace{\mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_{-}^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c, s(t) \geq \zeta, c_t \in \mathcal{M}_t) \right]}_{I_1} \\ &\quad + \underbrace{\mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_{-}^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c, s(t) < \zeta, c_t \in \mathcal{M}_t) \right]}_{I_2}. \end{aligned} \quad (\text{B.16})$$

To bound I_2 , we have

$$I_2 = \mathbb{E} \left[\sum_{t=1}^n \Delta_{X_t} \mathbb{I}(\mathcal{B}_t^c, \mathcal{D}_{t,c_t}^c, c_t \in \mathcal{M}_t, s(t) < \zeta) \right] \leq \Delta_{\max} \mathbb{E} \left[\sum_{t=1}^n \mathbb{I}(s(t) < \zeta, c_t \in \mathcal{M}_t, \mathcal{D}_{t,c_t}^c) \right] \leq \Delta_{\max} \mathbb{E}[\zeta].$$

It remains to bound $\mathbb{E}[\zeta]$. Let's define

$$\Lambda = \min \left\{ \lambda : \forall t : \mathcal{D}_{t,c_t}^c, \forall x \in \mathcal{A}, s(t) \geq s, \text{ such that } |\langle x, \hat{\theta}_t \rangle - \langle x, \theta \rangle| \leq \left(\frac{2}{\varepsilon_t s(t)} f_{n,1/\lambda} \right)^{1/2} \right\}.$$

From the definition of ζ in (B.15), we have

$$\zeta \leq \max \left\{ s : \left(\frac{f_{n,1/\lambda}}{\varepsilon_t s} \right)^{1/2} \geq \beta_n \right\},$$

which implies

$$\zeta \leq \frac{2f_{n,1/\Lambda}}{\varepsilon_t \beta_n^2}. \quad (\text{B.17})$$

In addition, we define

$$\Lambda' = \min \left\{ \lambda : \forall t \geq d, \forall x \in \mathcal{A}, \text{ such that } |\langle x, \hat{\theta}_t \rangle - \langle x, \theta \rangle| \leq \|x\|_{G_t^{-1}} f_{n,1/\lambda}^{1/2} \right\}.$$

Using the lower bound of $N_x(t)$, it holds that

$$\|x\|_{G_t^{-1}}^2 \leq \frac{1}{N_x(t)} \leq \frac{2}{\varepsilon_t s(t)}.$$

By Lemma A.2, we have

$$\mathbb{P} \left(\Lambda \geq \frac{1}{\delta} \right) \leq \mathbb{P} \left(\Lambda' \geq \frac{1}{\delta} \right) \leq \delta,$$

which implies that $\mathbb{E}[\log \Lambda] \leq 1$. From (B.17),

$$\mathbb{E}[\zeta] \leq \frac{2(1 + 1/\log(n)) + cd \log(\log(d \log(n)))}{\varepsilon_n \beta_n^2}. \quad (\text{B.18})$$

From (B.18), we have

$$\limsup_{n \rightarrow \infty} \frac{I_2}{\log(n)} \leq \limsup_{n \rightarrow \infty} \frac{\Delta_{\max} \mathbb{E}[\zeta]}{\log(n)} = 0, \quad (\text{B.19})$$

since β_n and ε_n are both sub-logarithmic. It remains to bound I_1 . When $s(t) \geq \zeta$, from the definition of ζ in (B.15) we have

$$\langle x, \hat{\theta}_t \rangle - \langle x, \theta \rangle \leq \beta_n,$$

holds for any $x \in \mathcal{A}$. For each $m \in [M]$, we have

$$\begin{aligned} \hat{\Delta}_{x_m^*}(t) &= \langle \hat{\theta}_t, \hat{x}_m^*(t) \rangle - \langle \hat{\theta}_t, x_m^* \rangle \\ &= \langle \hat{\theta}_t, \hat{x}_m^*(t) \rangle - \langle \theta, \hat{x}_m^*(t) \rangle - \langle \hat{\theta}_t, x_m^* \rangle + \langle \theta, x_m^* \rangle - \langle \theta, x_m^* \rangle + \langle \theta, \hat{x}_m^*(t) \rangle \\ &\leq 2\beta_t - \Delta_{\min}. \end{aligned}$$

When n is sufficiently large, it holds that $\beta_n \leq \Delta_{\min}/2$. This implies $\hat{\Delta}_{x_m^*}(t) = 0$ such that $x_m^* = \hat{x}_m^*(t)$ for all $t : s(t) > \zeta$. For notation simplicity, we denote $\mathcal{E}_t = \mathcal{B}_t^c \cap \mathcal{D}_{t,c_t}^c \cap \{s(t) \geq \zeta\} \cap \{c_t \in \mathcal{M}_t\}$. When \mathcal{E}_t occurs, the algorithm is in the unwasted exploration stage and $x_m^* = \hat{x}_m^*(n)$.

When \mathcal{D}_{t,c_t}^c occurs and $c_t \in \mathcal{M}_t$, there exists $x' \in \mathcal{A}^{c_t}$ such that $N_{x'}(t) \leq \min(f_n/\hat{\Delta}_{\min}^2(t), T_{x'}(\hat{\Delta}(t)))$. From the design of Algorithm 1, it holds that

- If $x = b_1$, then $N_x(t) \leq \min(f_n/\hat{\Delta}_{\min}^2(t), T_x(\hat{\Delta}(t)))$.
- If $x = b_2$, then $N_x(t) = \min_{x \in \mathcal{A}^{c_t}} N_x(t) \leq \min(f_n/\hat{\Delta}_{\min}^2(t), T_{x'}(\hat{\Delta}(t)))$.

Since the algorithm either pulls b_1 or b_2 in the unwasted exploration, it implies an upper bound for $s(t)$:

$$s(t) \leq \sum_{x \in \mathcal{A}^{c_t}} N_x(t) \leq |\mathcal{A}| \max_x \min(f_n/\hat{\Delta}_{\min}^2(t), T_x(\hat{\Delta}(t))). \quad (\text{B.20})$$

Let Λ be the random variable given by

$$\Lambda = \min \left\{ \lambda : \max_{x \in \mathcal{A}} |\langle x, \hat{\theta}_t - \theta \rangle| \leq \|x\|_{G_t^{-1}} f_{n,1/\lambda}^{1/2} \text{ for all } t \in [n] \right\},$$

where $f_{n,1/\lambda}$ is defined in Eq. (A.11). By the concentration inequality Lemma A.2, for any $\lambda \geq 1$,

$$\mathbb{P}(\Lambda \geq \lambda) \leq 1/\lambda. \quad (\text{B.21})$$

Hence the event $F = \{\Lambda \geq n\}$ satisfies $\mathbb{P}(F) \leq 1/n$. Denote $\alpha_x^m(\Delta) = T_x^m(\Delta)/f_n$ where $T_x^m(\Delta)$ is the solution of optimisation problem in Definition 4.1 with true Δ . Given $v > 0$ let

$$v(\delta) = \sup \left\{ \|\alpha(\Delta) - \alpha(\tilde{\Delta})\|_{\infty} : \|\tilde{\Delta} - \Delta\|_{\infty} \leq \delta \right\},$$

where $\alpha(\Delta) = \{\alpha_x^m(\Delta)\}_{x \in \mathcal{A}^m, m \in [M]}$. By continuity assumption of α at Δ we have $\lim_{\delta \rightarrow 0} v(\delta) = 0$. Moreover, let's define

$$\tau_{\delta} = \min \left\{ t : \max_{x \in \mathcal{A}} |\langle x, \hat{\theta}_s - \theta \rangle| \leq \delta/2 \text{ for all } x \in \mathcal{A} \text{ and } s \geq t \right\}.$$

Since $N_x(t) \geq \varepsilon_n s(t)/2$,

$$\max_{x \in \mathcal{A}} |\langle x, \hat{\theta}_t - \theta \rangle| \leq \sqrt{\frac{2f_{n,\Lambda}}{\varepsilon_n s(t)}}.$$

Therefore the number of exploration steps at time τ_δ is bounded by $s(\tau_\delta) \leq 8f_{n,1/\Lambda}\varepsilon_n^{-1}\delta^{-2}$.

Let $(\delta_n)_{n=1}^\infty$ be a sequence with $\lim_{n \rightarrow \infty} \delta_n = 0$ and $\log(\log(n))/\delta_n^2 = o(\log(n))$. I_{11} decomposed as

$$\begin{aligned} I_{11} &= \mathbb{E} \left[\sum_{t=1}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{E}_t) \right] \\ &\leq \mathbb{E} [s(\tau_{\delta_n})] + \mathbb{E} \left[\sum_{t=\tau_{\delta_n}}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{E}_t) \right]. \end{aligned} \quad (\text{B.22})$$

The first term in (B.22) is bounded by

$$\mathbb{E} [s(\tau_{\delta_n})] \leq \frac{8}{\varepsilon_n \delta_n^2} \mathbb{E} [f_{n,1/\Lambda}] = o(\log(n)),$$

where we used the assumption on (δ_n) and the fact that $\mathbb{E} [f_{n,1/\Lambda}] = O(\log \log(n))$. By the continuity assumption, the following statement holds

$$\begin{aligned} \sum_{t=\tau_{\delta_n}+1}^n \mathbb{I}(X_t = x, \mathcal{E}_t) &\leq \varepsilon_n s(n) + f_n \min \left(1/\widehat{\Delta}_{\min}^2(n), \alpha_x^{c_t}(\widehat{\Delta}(n))/2 \right) \\ &\leq \varepsilon_n s(n) + f_n \min \left(\frac{1}{\widehat{\Delta}_{\min}^2(n)}, (\alpha_x^{c_t}(\Delta) + v(\delta_n))/2 \right). \end{aligned} \quad (\text{B.23})$$

The second term in (B.22) is bounded by

$$\begin{aligned} &\mathbb{E} \left[\sum_{t=\tau_{\delta_n}}^n \sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \mathbb{I}(X_t = x, \mathcal{E}_t) \right] \\ &\leq \mathbb{E} \left[\sum_{m=1}^M \sum_{x \in \mathcal{A}_-^m} \Delta_x \sum_{t=1}^n \mathbb{I}(X_t = x, \mathcal{E}_t) \right] \\ &\leq \mathbb{E} \left[\sum_{m=1}^M \sum_{x \in \mathcal{A}_-^m} \Delta_x \varepsilon_n s(n) \mathbb{I}(\mathcal{E}_n) \right] + \mathbb{E} \left[\sum_{m=1}^M \sum_{x \in \mathcal{A}_-^m} \Delta_x f_n (\alpha_x^m(\Delta) + v(\delta_n))/2 \mathbb{I}(\mathcal{E}_n) \right]. \end{aligned}$$

To bound the second term, we take the limit as n tends to infinity and the fact that $\lim_{n \rightarrow \infty} v(\delta_n) = 0$ and $f_n \sim 2 \log(n)$ shows that

$$\limsup_{n \rightarrow \infty} \frac{1}{\log(n)} \mathbb{E} \left[\sum_{m=1}^M \sum_{x \in \mathcal{A}_-^m} \Delta_x f_n (\alpha_x^m(\Delta) + v(\delta_n))/2 \mathbb{I}(\mathcal{E}_n) \right] \leq \mathcal{C}(\theta, \mathcal{A}^1, \dots, \mathcal{A}^M). \quad (\text{B.24})$$

We bound the first term in the following lemma. The detailed proofs are deferred to Section B.4.

Lemma B.3. The regret contributed by the forced exploration satisfies

$$\limsup_{n \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{x \in \mathcal{A}_-^{c_t}} \Delta_x \varepsilon_n s(n) \mathbb{I}(\mathcal{E}_n) \right]}{\log(n)} = 0.$$

This ends the proof. ■

B.4 Proof of Lemma B.3: Forced Exploration Regret

By the upper bound of unwasted exploration counter $s(n)$ in (B.20), it holds that

$$\begin{aligned} \sum_{m=1}^M \sum_{x \in \mathcal{A}_-^m} \Delta_x^m \varepsilon_n s(n) \mathbb{I}(\mathcal{E}_n) &\leq \sum_{m=1}^M \sum_{x \in \mathcal{A}_-^m} \Delta_x^m \varepsilon_n |\mathcal{A}| \max_x \min(f_n / \widehat{\Delta}_{\min}^2(n), T_x(\widehat{\Delta}(n))) \mathbb{I}(\mathcal{E}_n) \\ &\leq \varepsilon_n |\mathcal{A}| \sum_{m=1}^M \sum_{x \in \mathcal{A}_-^m} \Delta_x^m f_n / \widehat{\Delta}_{\min}(n) \mathbb{I}(\mathcal{E}_n). \end{aligned}$$

When event \mathcal{E}_n occurs,

$$\begin{aligned} \max_{x \neq \widehat{x}_m^*(n)} \frac{(\Delta_x^m)^2}{(\widehat{\Delta}_x(n))^2} &\leq \max_{x \neq \widehat{x}_m^*(n)} \frac{(\Delta_x^m)^2}{(\Delta_x^m - 2\beta_n)^2} \\ &= \max_{x \neq \widehat{x}_m^*(n)} \left(1 + \frac{4(\Delta_x^m - \beta_n)\beta_n}{(\Delta_x^m - 2\beta_n)^2} \right) \leq 1 + \frac{16\beta_n}{\Delta_{\min}}, \end{aligned}$$

For any $x \in \mathcal{A}^m$,

$$\widehat{\Delta}_{\min}(n) \geq \frac{1}{1 + 16\beta_n / \Delta_{\min}} \Delta_{\min}. \quad (\text{B.25})$$

Since $\varepsilon_n = 1/(\log \log(n))$, we have

$$\limsup_{n \rightarrow \infty} \frac{\sum_{x \in \mathcal{A}_-} \Delta_x \varepsilon_n s(n) \mathbb{I}(\mathcal{E})}{\log(n)} = 0. \quad (\text{B.26})$$

This ends the proof. ■

B.5 Proof of Lemma B.2: Wasted Exploration

First, we define

$$\mathcal{F}_s = \left\{ \exists t \geq d, \exists x : \langle x, \widehat{\theta}_t \rangle - \langle x, \theta \rangle \geq \|x\|_{G_t^{-1}} f_{n,1/s^2}^{1/2} \right\}, \quad (\text{B.27})$$

where $f_{n,1/s^2}$ is defined in Lemma A.2. From Lemma A.2, we also have $\mathbb{P}(\mathcal{F}_s) \leq 1/s^2$. Let $s'(t), s(t)$ be the number of rounds for wasted explorations, unwasted explorations until round t accordingly, and x_t^* is the optimal arm at round t . We decompose the regret as follows

$$R_{\text{we}} \leq \underbrace{\mathbb{E} \left[\sum_{t \in \text{wasted}} \mathbb{I}(\mathcal{F}_{s'(t)}) \Delta_{\max} \right]}_{I_1} + \underbrace{\mathbb{E} \left[\sum_{t \in \text{unwasted}} \mathbb{I}(\mathcal{F}_{s'(t)}^c) \langle x_t^* - X_t, \theta \rangle \right]}_{I_2}. \quad (\text{B.28})$$

To bound I_1 , we have

$$I_1 \leq \sum_{s=1}^n \mathbb{P}(\mathcal{F}_s) \Delta_{\max} \leq \sum_{s=1}^n \frac{1}{s^2} \Delta_{\max} = (2 - \frac{1}{n}) \Delta_{\max}. \quad (\text{B.29})$$

To bound I_2 , let's denote $\widetilde{\theta}_t$ as the optimistic estimator. Following the standard one step regret decomposition (See the proof of Theorem 19.2 in [17] for details), it holds that

$$\begin{aligned} \langle x_t^* - X_t, \theta \rangle &= \langle x_t^*, \theta \rangle - \langle X_t, \theta \rangle \\ &\leq \langle X_t, \widetilde{\theta}_t \rangle - \langle X_t, \theta \rangle \\ &= \langle X_t, \widehat{\theta}_t - \theta \rangle + \langle X_t, \widetilde{\theta}_t - \widehat{\theta}_t \rangle. \end{aligned}$$

When $\mathcal{F}_{s'(t)}^c$ occurs, we have

$$\langle X_t, \hat{\theta}_t - \theta \rangle \leq \|X_t\|_{G_t^{-1}} f_{n,1/(s'(t)^2)}, \langle X_t, \tilde{\theta}_t - \hat{\theta}_t \rangle \leq \|X_t\|_{G_t^{-1}} f_{n,1/(s'(t)^2)}.$$

Putting the above results together, we have

$$\langle x_t^* - X_t, \theta \rangle \leq 2\|X_t\|_{G_t^{-1}} f_{n,1/(s'(t)^2)}^{1/2}.$$

Applying Lemma C.3, we can bound I_2 as follows

$$\begin{aligned} I_2 &\leq \mathbb{E} \left[2f_{n,1/(s'(t)^2)}^{1/2} \sum_{t \in \text{wasted}} \|X_t\|_{G_t^{-1}} \right] \\ &\leq \mathbb{E} \left[2f_{n,1/(s'(t)^2)}^{1/2} \sqrt{2s'(n)d \log \left(\frac{s'(n) + d}{d} \right)} \right]. \end{aligned} \quad (\text{B.30})$$

Recall that $p_{\min} = \min_m p_m$ be the minimum probability that each action set arrives. It is easy to see $\mathbb{P}(c_t \in \mathcal{M}_t | \mathcal{D}_{t,c_t}^c) = \mathbb{P}(c_t \in \mathcal{M}_t | \mathcal{M}_t \neq \emptyset) = \sum_{m \in \mathcal{M}_t} p_m \geq p_{\min}$. We bound $s'(n)$ by $s(n)$ as follows

$$\begin{aligned} \mathbb{E}[s'(n)] &= \mathbb{E} \left[\sum_{t=1}^n \mathbb{I}(\mathcal{D}_{t,c_t}^c, c_t \notin \mathcal{M}_t) \right] = \sum_{t=1}^n \mathbb{P}(\mathcal{D}_{t,c_t}^c) \mathbb{P}(c_t \notin \mathcal{M}_t | \mathcal{D}_{t,c_t}^c) \\ &\leq \frac{1}{p_{\min}} \sum_{t=1}^n \mathbb{P}(\mathcal{D}_{t,c_t}^c) \mathbb{P}(c_t \in \mathcal{M}_t | \mathcal{D}_{t,c_t}^c) \\ &= \frac{1}{p_{\min}} \mathbb{E} \left[\sum_{t=1}^n \mathbb{I}(\mathcal{D}_{t,c_t}^c, c_t \in \mathcal{M}_t) \right] = \frac{1}{p_{\min}} \mathbb{E}[s(n)]. \end{aligned} \quad (\text{B.31})$$

Putting (B.29)-(B.31) together, The regret in the wasted exploration can be upper bounded by

$$R_{\text{we}} \leq (2 - \frac{1}{n}) \Delta_{\max} + \frac{2}{p_{\min}} \sqrt{2d \log \left(\frac{s(n)/p_{\min} + d}{d} \right) f_{n,(p_{\min}/s(n))^2} s(n)/p_{\min}}, \quad (\text{B.32})$$

where $f_{n,(p_{\min}/s(n))^2}$ is defined in (A.2).

Next, we recall the upper bound (B.20) for the number of pulls in unwasted exploration,

$$\begin{aligned} s(n) &\leq |\mathcal{A}| \max_x \min \left\{ f_n / \hat{\Delta}_{\min}(n), T_x(\hat{\Delta}(n)) \right\} \\ &\leq |\mathcal{A}| f_n / \hat{\Delta}_{\min}(n). \end{aligned}$$

From (B.25), we have

$$\hat{\Delta}_{\min}(n) \geq \frac{1}{1 + \delta_n} \Delta_{\min} \geq \frac{\Delta_{\min}^2}{\Delta_{\min} + 16\beta_n},$$

where $\beta_n = 1/\log(\log(n))$. Overall, we see $s(n) \leq \mathcal{O}(\log(n))$. Plugging this into (B.32), we reach

$$\limsup_{n \rightarrow \infty} \frac{R_{\text{we}}}{\log(n)} = 0.$$

This ends the proof. ■

B.6 Proof of Lemma A.1

First, we start by the following claim:

Claim B.4. Assume H_n is a sequence of $d \times d$ positive definite matrices such that $H_n \rightarrow H$ and H is positive semidefinite. Then, $HH_n^{-1}H \rightarrow H$ as $n \rightarrow \infty$.

Proof. Without loss of generality, we can assume that H is given in the block matrix form

$$H = \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix}$$

where A is a nonsingular $m \times m$ matrix with $m > 0$. (If $m = 0$, H is the all zero matrix and the claim trivially holds.) Consider the same block partitioning of H_n :

$$H_n = \begin{pmatrix} A_n & B_n \\ B_n^\top & D_n \end{pmatrix},$$

where A_n is thus also an $m \times m$ matrix. Clearly, $A = \lim_{n \rightarrow \infty} A_n$ and A_n is nonsingular (or H_n would be singular), while $B_n \rightarrow B$ and $D_n \rightarrow D$ where all entries in B and D are zero. Then, as is well known,

$$H_n^{-1} = \begin{pmatrix} A_n^{-1} + A_n^{-1}B_nS_n^{-1}B_n^\top A_n^{-1} & -A_n^{-1}B_nS_n^{-1} \\ -S_n^{-1}B_n^\top A_n^{-1} & S_n^{-1} \end{pmatrix}.$$

where $S_n = D_n - B_n^\top A_n^{-1}B_n$ is the Schur-complement of block D_n of matrix H_n . Note that

$$HH_n^{-1}H = \begin{pmatrix} A(A_n^{-1} + A_n^{-1}B_nS_n^{-1}B_n^\top A_n^{-1})A & 0 \\ 0 & 0 \end{pmatrix}.$$

Since the matrix inverse is continuous if the limit is nonsingular, $A_n^{-1} \rightarrow A^{-1}$. Clearly, it suffices to show that $A_n^{-1} + A_n^{-1}B_nS_n^{-1}B_n^\top A_n^{-1} \rightarrow A^{-1}$. Hence, it remains to check that $A_n^{-1}B_nS_n^{-1}B_n^\top A_n^{-1} \rightarrow 0$. This follows because $B_n \rightarrow B$ and $D_n \rightarrow D$ and $S_n \rightarrow D - B^\top A^{-1}B = 0$ where $D = 0$ and $B = 0$. \square

Proof of Lemma A.1. Let $L = \limsup_{n \rightarrow \infty} \log(n) \|s\|_{G_n^{-1}}^2$. We need to prove that $L \leq 1/c$. Without loss of generality, assume that $L > 0$ (otherwise there is nothing to be proven) and that for some H positive semidefinite matrix, $\zeta \in \mathbb{R}$ and $\kappa \in \mathbb{R} \cup \{\infty\}$, (i) $\log(n) \|s\|_{G_n^{-1}}^2 \rightarrow L$; (ii) $H_n = G_n^{-1} / \|G_n^{-1}\| \rightarrow H$; (iii) $\lambda_{\min}(G_n) / \log(n) \rightarrow \zeta > 0$ and (iv) $\frac{\rho_n(H)}{\log(n) \|s\|_{G_n^{-1}}^2} \rightarrow \kappa \geq c$. We claim that $\|s\|_H > 0$, hence $\rho_n(H)$ is well-defined and in particular $\rho_n(H) \rightarrow 1$ as $n \rightarrow \infty$. If this was true, then the proof was ready since

$$L = \lim_{n \rightarrow \infty} \frac{\log(n) \|s\|_{G_n^{-1}}^2}{\rho_n(H)} = \frac{1}{\lim_{n \rightarrow \infty} \frac{\rho_n(H)}{\log(n) \|s\|_{G_n^{-1}}^2}} = 1/\kappa \leq 1/c.$$

Hence, it remains to show the said claim. We start by showing that $\|s\|_H > 0$. For this note that $\|G_n^{-1}\| = 1/\lambda_{\min}(G_n)$ and hence

$$\|s\|_{\frac{G_n^{-1}}{\|G_n^{-1}\|}}^2 = \frac{\lambda_{\min}(G_n)}{\log(n)} \|s\|_{G_n^{-1}}^2 \log(n).$$

Taking the limit of both sides, we get $\|s\|_H^2 \rightarrow \zeta L > 0$. Now,

$$\rho_n(H) = \frac{\|s\|_{G_n^{-1}}^2 \|s\|_{HG_nH}^2}{\|s\|_H^4} = \frac{\|s\|_{H_n}^2 \|s\|_{HH_n^{-1}H}^2}{\|s\|_H^4} \xrightarrow{n \rightarrow \infty} \frac{\|s\|_H^2 \|s\|_H^2}{\|s\|_H^4} = 1,$$

where we used Claim B.4. \square

B.7 Proof of Theorem 3.9

Suppose that $\{x_m^* : m \in [M]\}$ spans \mathbb{R}^d . Recall that LinUCB chooses

$$X_t = \operatorname{argmax}_{x \in \mathcal{A}^c_t} \langle x, \hat{\theta}_{t-1} \rangle + \|x\|_{G_t^{-1}} \beta_t^{1/2},$$

where $\beta_t = O(d \log(t))$ is chosen so that

$$\mathbb{P} \left(\|\hat{\theta}_t - \theta\|_{G_t} \geq \beta_t \right) \leq 1/t^3,$$

which is known to be possible [17, §20]. Define F_t to be the event that $\|\hat{\theta}_t - \theta\|_{G_t} \geq \beta_t$. Then the instantaneous pseudo-regret of LinUCB is bounded by

$$\Delta_t \leq \mathbf{1}_{F_t} + \langle x_t^* - X_t, \theta \rangle \leq \mathbf{1}_{F_t} + 2\beta_t^{1/2} \|X_t\|_{G_t^{-1}} \leq \mathbf{1}_{F_t} + 2\sqrt{\beta_t \|G_t^{-1}\|},$$

where the matrix norm is the operator name (in this case, maximum eigenvalue). Let $\tau = 1 + \max\{t : F_t \text{ holds}\}$, which satisfies $\mathbb{E}[\tau] = O(1)$. The cumulative regret after τ is bounded almost surely by

$$\sum_{t=\tau}^n \langle x_t^* - X_t, \theta \rangle = O(\sqrt{n \log(n)}),$$

where the Big-Oh hides constants that only depend on the dimension. Hence all optimal arms are played linearly often after τ , which by the assumption that $\{x_m^* : m \in [M]\}$ spans \mathbb{R}^d implies that $\|G_t^{-1}\| = O(1/t)$. Hence the instantaneous regret for times $t \geq \tau$ satisfies

$$\Delta_t = O\left(\sqrt{\frac{\beta_t}{t}}\right).$$

Since $\Delta_t \in \{0\} \cup [\Delta_{\min}, 1]$, it follows that the regret vanishes once $\Delta_t < \Delta_{\min}$. But by the previous argument and the assumption on β_t we have for $t \geq \tau$ that

$$\Delta_t \leq 2\sqrt{\beta_t \|G_t^{-1}\|} = O\left(\sqrt{\frac{\log(t)}{t}}\right).$$

Hence for sufficiently large t (independent of n) the regret vanishes, which completes the proof.

C Supporting Lemmas

Lemma C.1 (Bretagnolle-Huber Inequality). Let \mathbb{P} and $\tilde{\mathbb{P}}$ be two probability measures on the same measurable space (Ω, \mathcal{F}) . Then for any event $\mathcal{D} \in \mathcal{F}$,

$$\mathbb{P}(\mathcal{D}) + \tilde{\mathbb{P}}(\mathcal{D}^c) \geq \frac{1}{2} \exp\left(-\text{KL}(\mathbb{P}, \tilde{\mathbb{P}})\right), \quad (\text{C.1})$$

where \mathcal{D}^c is the complement event of \mathcal{D} ($\mathcal{D}^c = \Omega \setminus \mathcal{D}$) and $\text{KL}(\mathbb{P}, \tilde{\mathbb{P}})$ is the KL-divergence between \mathbb{P} and $\tilde{\mathbb{P}}$, which is defined as $+\infty$, if \mathbb{P} is not absolutely continuous with respect to $\tilde{\mathbb{P}}$, and is $\int_{\Omega} d\mathbb{P}(\omega) \log \frac{d\mathbb{P}}{d\tilde{\mathbb{P}}}(\omega)$ otherwise.

The proof can be found in the book of Tsybakov [21]. When $\text{KL}(\mathbb{P}, \tilde{\mathbb{P}})$ is small, we may expect the probability measure \mathbb{P} is close to the probability measure $\tilde{\mathbb{P}}$. Note that $\mathbb{P}(\mathcal{D}) + \mathbb{P}(\mathcal{D}^c) = 1$. If $\tilde{\mathbb{P}}$ is close to \mathbb{P} , we may expect $\mathbb{P}(\mathcal{D}) + \tilde{\mathbb{P}}(\mathcal{D}^c)$ to be large.

Lemma C.2 (Divergence Decomposition). Let \mathbb{P} and $\tilde{\mathbb{P}}$ be two probability measures on the sequence $(A_1, Y_1, \dots, A_n, Y_n)$ for a fixed bandit policy π interacting with a linear contextual bandit with standard Gaussian noise and parameters θ and $\tilde{\theta}$ respectively. Then the KL divergence of \mathbb{P} and $\tilde{\mathbb{P}}$ can be computed exactly and is given by

$$\text{KL}(\mathbb{P}, \tilde{\mathbb{P}}) = \frac{1}{2} \sum_{x \in \mathcal{A}} \mathbb{E}[T_x(n)] \langle x, \theta - \tilde{\theta} \rangle^2, \quad (\text{C.2})$$

where \mathbb{E} is the expectation operator induced by \mathbb{P} .

This lemma appeared as Lemma 15.1 in the book of Lattimore and Szepesvári [17], where the reader can also find the proof.

Lemma C.3. Let $\{X_t\}_{t=1}^\infty$ be a sequence in \mathbb{R}^d satisfying $\|X_t\|_2 \leq 1$ and $G_t = \sum_{s=1}^t X_s X_s^\top$. Suppose that $\lambda_{\min}(G_d) \geq c$ for some strictly positive c . For all $n > 0$, it holds that

$$\sum_{t=d+1}^n \|X_t\|_{G_t^{-1}} \leq \sqrt{2nd \log\left(\frac{d+n}{d}\right)}.$$

Lemma C.4. Let $\varepsilon > 0$ and denote $T(\hat{\Delta}(n)) \in \mathbb{R}^{|\mathcal{A}|}$ as the solution of the optimisation problem defined in Definition 4.1. Then we define

$$S_\varepsilon(\hat{\Delta}(n)) = \min \left\{ \varepsilon f_n, T(\hat{\Delta}(n)) \right\}.$$

Then for all $x \in \mathcal{A}$,

$$\|x\|_{H_{S_\varepsilon(\hat{\Delta}(n))}^{-1}}^2 \leq \max \left\{ \frac{\varepsilon^2}{f_n}, \frac{\hat{\Delta}_x^2(n)}{f_n} \right\}.$$

This is Lemma 17 in the book of Lattimore and Szepesvári [16], where the reader can also find the proof.

Lemma C.5. Suppose that $T_x^m(\cdot)$ is uniquely defined at Δ . Then it is continuous at Δ .

Proof. Suppose it is not continuous. Then there exists a sequence $(\Delta_n)_{n=1}^\infty$ with $\lim_{n \rightarrow \infty} \|\Delta_n - \Delta\| = 0$ and for which $\lim_{n \rightarrow \infty} T_x^m(\Delta_n) \neq T_x^m(\Delta)$ for some m and $x \in \mathcal{A}^m$. Since $\Delta_n \rightarrow \Delta$ it follows that for sufficiently large n the optimal actions with respect to Δ_n are the same as Δ . Hence, for sufficiently large n , by the definition of the optimisation problem,

$$T_{x_m^*}^m(\Delta_n) = \infty = T_{x_m^*}^m(\Delta).$$

Therefore there exists a context m and suboptimal action $x \neq x_m^*$ such that $\lim_{n \rightarrow \infty} T_x^m(\Delta_n) \neq T_x^m(\Delta)$. It is easy to check that the value of the optimisation problem is continuous. Specifically, that

$$\lim_{n \rightarrow \infty} \sum_{m=1}^M \sum_{x \in \mathcal{A}^m} T_x^m(\Delta_n) = \sum_{m=1}^M \sum_{x \in \mathcal{A}^m} T_x^m(\Delta).$$

Hence $\limsup_{n \rightarrow \infty} T_x^m(\Delta_n) < \infty$ for $x \neq x_m^*$. Therefore a compactness argument shows there exists a cluster point S of the allocation $(T(\Delta_n))_{n=1}^\infty$ with $S_x^m \neq T_x^m(\Delta)$ for some m and $x \neq x_m^*$. And yet by the previous display

$$\sum_{m=1}^M \sum_{x \in \mathcal{A}^m} S_x^m = \sum_{m=1}^M \sum_{x \in \mathcal{A}^m} T_x^m(\Delta).$$

Since the constraints of the optimisation problem are continuous it follows that S also satisfies the constraints in the optimisation problem and so $S \neq T(\Delta)$ is another optimal allocation, contradicting uniqueness. Therefore $T_x^m(\cdot)$ is continuous at Δ . \square




Figure 4: The left panel is for small size action set and the right panel is for randomly generated θ . The results are averaged over 100 realisations.

D Additional Experiments

In this section, we consider two more experiment settings in Figure 4.

1. Small size action set. We conduct the experiments with the number of action set equal to 5. Comparing with large size action set (Section 5.4), we found that OAM still outperforms OSSB but the improvement is smaller, as one might expect.

2. Randomly generated θ . For each replication, θ is randomly generated from multivariate normal with variance 10 and we normalise θ such that its ℓ_2 norm is 1. OAM still outperforms OSSB for randomly generated θ . In addition, we compare with the heuristic LinTS (remove all the variance blowup factors and use a Gaussian prior). We find that the heuristic LinTS enjoys the best performance by a modest margin. Analysing heuristic LinTS, however, remains a fascinating open problem. As far as we are aware, it is not known whether or not it even achieves sublinear regret in the worst case.