

---

# Improved Regret Analysis for Variance-Adaptive Linear Bandits and Horizon-Free Linear Mixture MDPs

---

Yeoneung Kim\*  
Gachon University  
yeoneung@gachon.ac.kr

Insoon Yang  
Seoul National University  
insoonyang@snu.ac.kr

Kwang-Sung Jun  
University of Arizona  
kjun@cs.arizona.edu

## Abstract

In online learning problems, exploiting low variance plays an important role in obtaining tight performance guarantees yet is challenging because variances are often not known a priori. Recently, considerable progress has been made by Zhang et al. (2021) where they obtain a variance-adaptive regret bound for linear bandits without knowledge of the variances and a horizon-free regret bound for linear mixture Markov decision processes (MDPs). In this paper, we present novel analyses that improve their regret bounds significantly. For linear bandits, we achieve  $\tilde{O}(\min\{d\sqrt{K}, d^{1.5}\sqrt{\sum_{k=1}^K \sigma_k^2}\} + d^2)$  where  $d$  is the dimension of the features,  $K$  is the time horizon, and  $\sigma_k^2$  is the noise variance at time step  $k$ , and  $\tilde{O}$  ignores polylogarithmic dependence, which is a factor of  $d^3$  improvement. For linear mixture MDPs with the assumption of maximum cumulative reward in an episode being in  $[0, 1]$ , we achieve a horizon-free regret bound of  $\tilde{O}(d\sqrt{K} + d^2)$  where  $d$  is the number of base models and  $K$  is the number of episodes. This is a factor of  $d^{3.5}$  improvement in the leading term and  $d^7$  in the lower order term. Our analysis critically relies on a novel peeling-based regret analysis that leverages the elliptical potential ‘count’ lemma.

## 1 Introduction

In online learning, variance often plays an important role in achieving low regret bounds. For example, for the prediction with expert advice problem, Hazan and Kale [11] proposed an algorithm that achieves a regret bound of  $O(\sqrt{\text{VAR}_K})$  where  $\text{VAR}_K$  is a suitably-defined variance of the loss function up to time step  $K$ , without knowing  $\text{VAR}_K$  ahead of time. The implication is that when the given sequence of loss functions has a small variance, one can perform much better than the previously known regret bound  $O(\sqrt{K})$ . For multi-armed bandits, Audibert et al. [2] proposed an algorithm that achieves regret bounds that depends on the variances of the arms, which means that, again, the regret bound becomes smaller as the variances become smaller.

It is thus natural to obtain similar variance-adaptive bounds for other problems. For example, in  $d$ -dimensional stochastic contextual bandit problems, the optimal worst-case regret bound is  $\tilde{O}(\sigma d\sqrt{K})$  where  $\tilde{O}$  hides polylogarithmic dependencies and  $\sigma^2$  is a uniform upper bound on the noise variance. Following the developments in other online learning problems, it is natural to ask if we can develop a similar variance-adaptive regret bound. The recent work by Zhang et al. [31] has provided an affirmative answer. Their algorithm called VOFUL achieves a regret bound of  $\tilde{O}(d^{4.5}\sqrt{\sum_{k=1}^K \sigma_k^2} + d^5)$  where  $\sigma_k^2$  is the (unknown) noise variance at time step  $k$ . This implies that, indeed, it is possible

---

\*Work done while at Seoul National University.

to adapt to the variance and suffer a much lower regret. Furthermore, they show that a similar variance-adaptive analysis can be used to solve linear mixture Markov decision processes (MDPs) with the *unit cumulative rewards* assumption :

$$\sum_h r_h^k \in [0, 1], \forall k \quad (1)$$

where  $r_h^k$  is the reward received at episode  $k$  and horizon  $h$ . They show a regret bound of  $\tilde{O}(d^{4.5}\sqrt{K} + d^9)$ , which does not depend on the planning horizon length  $H$  up to polylogarithmic factors. We elaborate more on the linear bandit and linear mixture MDP problems in Section 2.

However, the regret rates of these problems have a large gap between the known lower and the upper bounds. For example, in linear bandits, it is well-known that the regret bound has to be  $\Omega(d\sqrt{K})$  [6], which rejects the possibility of obtaining  $o(d\sqrt{\sum_{k=1}^K \sigma_k^2})$ , yet the best upper bound obtained so far is  $O(d^{4.5}\sqrt{\sum_{k=1}^K \sigma_k^2})$ . Thus, the gap is a factor of  $d^{3.5}$ , which is quite large.

In this paper, we reduce such gaps significantly by obtaining much tighter regret upper bounds. Specifically, we show that a slight variation of VOFUL [31] for linear bandits has a regret bound of  $\tilde{O}(\min\{d\sqrt{K}, d^{1.5}\sqrt{\sum_{k=1}^K \sigma_k^2}\})$  without knowledge of the variances. This reduces the gap between the upper and lower bounds to only  $\sqrt{d}$  for the leading term in the regret. Furthermore, we employ a similar technique to show that the algorithm VARLin [31] for linear mixture MDPs with unit cumulative rewards has a regret bound of  $\tilde{O}(d\sqrt{K} + d^2)$ . At the heart of our analysis is a direct peeling of the instantaneous regret terms using an elliptical potential ‘count’ lemma (EPC). EPC bounds, given  $q > 0$ , how many times  $\|x_k\|_{V_{k-1}^{-1}}^2 \geq q$  happens from time  $k = 1$  to  $\infty$  where  $V_{k-1} = \sum_{s=1}^{k-1} x_s x_s^\top$ . Our lemma is an improved and generalized version of [16, Exercise 19.3], which was originally used for improving the regret bound of linear bandit algorithms. We provide the proofs of our main results for linear bandits and linear mixture MDPs in Section 3 and Section 4 respectively. Finally, we conclude the paper with exciting future directions.

**Related work.** There are numerous works on linear bandit problems such as [6, 4, 1, 17] where the information of variance is not used. On the other hand, variance can be exploited to obtain better regret [2]. Recently, works by [31, 33] proposed ways to infuse the variance information in the regret analysis which improves the standard regret bound. Reinforcement learning with linear function approximation has been widely studied to develop efficient learning methods that work for large state-action space [27, 26, 13, 8, 14, 23, 24, 25, 29, 18, 15, 7, 21, 10, 9, 28]. To our knowledge, all aforementioned works derived a regret bound that depends on the planning horizon  $H$  polynomially. It was Zhang et al. [31] who first remove the polynomial dependence of  $H$  in the linear mixture MDP problem, achieving a bound of  $\tilde{O}(d^{4.5}\sqrt{K} + d^9)$ . In contrast, our analysis shows that their algorithm in fact achieves significantly better bound of  $\tilde{O}(d\sqrt{K} + d^2)$ . Note that these results assume the unit cumulative rewards assumption (1) and time-homogeneous transition models. In a similar setup where  $r_{h,k} \in [0, 1]$  with time-inhomogeneous transition models, Zhou et al. [33] achieve the regret  $\tilde{O}(\sqrt{d^2 H} + dH^3\sqrt{HK} + d^2 H^3 + d^3 H^2)$  and show a lower bound of  $\tilde{\Omega}(dH^{3/2}\sqrt{K})$ . These problem setups are incompatible to the setup of [31] and ours.

## 2 Problem Definition

**Notations.** We denote  $d$ -dimensional  $\ell_2$  ball by  $\mathbb{B}_2^d(R) := \{x \in \mathbb{R}^d : \|x\|_2 \leq R\}$  and define  $\mathbb{B}_1^d(R)$  similarly for the  $\ell_1$  ball. Let  $[N] := \{1, 2, \dots, N\}$  for  $N \in \mathbb{N}$ . Given  $\ell \in \mathbb{R}$  and  $x \in \mathbb{R}$ , we define the clipping operator as follows (take  $0/0 = 0$ ):

$$\overline{(x)}_\ell := \min\{|x|, 2^{-\ell}\} \cdot \frac{x}{|x|}. \quad (2)$$

**Linear bandits.** The linear bandit problem has the following protocol. At time step  $k$ , the learner observes an arm set  $\mathcal{X}_k \subseteq \mathbb{B}_2^d(1)$ , chooses an arm  $x_k \in \mathcal{X}_k$ , pulls it. The learner then receives a stochastic reward  $r_k = x_k^\top \theta^* + \epsilon_k$  where  $\theta^* \in \mathbb{B}_2^d(1)$  is an unknown parameter and  $\epsilon_k$  is a zero-mean stochastic noise. Following [31], we assume that (i)  $\forall k \in [K], |r_k| \leq [-\frac{1}{2}, \frac{1}{2}]$  almost surely, (ii)  $\mathbb{E}[\epsilon_k | \mathcal{F}_k] = 0$  where  $\mathcal{F}_k = \sigma(x_1, \epsilon_1, \dots, x_{k-1}, \epsilon_{k-1}, x_k)$ , and (iii)  $\mathbb{E}[\epsilon_k^2 | \mathcal{F}_k] = \sigma_k^2$ .

Note that the bound on  $|r_k|$  implies that  $|\epsilon_k| \leq 1$  almost surely. Our goal is to minimize the regret  $\mathcal{R}^K = \sum_{k=1}^K \max_{x \in \mathcal{X}_k} x^\top \theta^* - x_k^\top \theta^*$ .

**Linear mixture MDPs.** We consider an episodic Markov Decision Process (MDP) with a tuple  $(\mathcal{S}, \mathcal{A}, r(s, a), P(s'|s, a), K, H)$  where  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space,  $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  is the reward function,  $P(s'|s, a)$  is the transition probability,  $K$  is the number of episodes, and  $H$  is the planning horizon. A policy is defined as  $\pi = \{\pi_h : \mathcal{S} \rightarrow \mathcal{D}(\mathcal{A})\}_{h=1}^H$  where  $\mathcal{D}(\mathcal{A})$  is a set of all distributions over  $\mathcal{A}$ . For each episode  $k \in [K]$ , the learner chooses a policy  $\pi^k$ , and then the environment executes  $\pi^k$  on the MDP by successively following  $a_h^k \sim \pi_h^k(s_h^k)$  and  $s_{h+1}^k \sim P(\cdot|s_h^k, a_h^k)$ . Then, the learner observes the rewards  $\{r_h^k \in [0, 1]\}_{k,h}$  and moves onto the next episode. The key modeling assumption of linear mixture MDPs is that the transition probability  $P$  is a linear combination of a known set of models  $\{P^i\}$ , namely,  $P = \sum_{i=1}^d \theta_i^* P^i$  where  $\theta^* \in \mathbb{B}_1^d(1)$  is an unknown parameter. We follow [31] and make the following assumptions:

- The reward at each time step  $h$  and episode  $k$  is  $r_h^k = r(s_h^k, a_h^k)$  for some known function  $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ .
- Unit cumulative rewards:  $\sum_{h=1}^H r_h^k \in [0, 1]$  for any policy  $\pi^k$ .

For a policy  $\pi$ ,  $V_h^\pi(s) := \max_{a \in \mathcal{A}} Q_h^\pi(s, a)$  where  $Q_h^\pi(s, a) = r(s, a) + \mathbb{E}_{s' \sim P(\cdot|s, a)} V_{h+1}^\pi(s')$  and  $V_{H+1}^\pi(s) := 0$ . Denoting  $V^\pi(s_1) = V_1^\pi(s_1)$  and  $V^*(s_1) = V^{\pi^*}(s_1)$ , our goal is to minimize the regret

$$\mathcal{R}^K = \sum_{k=1}^K V^*(s_1^k) - V^k(s_1^k).$$

### 3 Variance-Adaptive Linear Bandits

In this section, we show that VOFUL of Zhang et al. [31] has a tighter regret bound than what was reported in their work. Our version of VOFUL, which we call VOFUL2, has a slightly different confidence set for ease of exposition. Specifically, we use a confidence set that works for every  $\mu \in \mathbb{B}_2^d(2)$  rather than over an  $\epsilon$ -net of  $\mathbb{B}_2^d(2)$  (but we do use an  $\epsilon$ -net for the proof of the confidence set).

The full pseudocode can be found in Algorithm 1. VOFUL2 follows the standard optimism-based arm selection [3, 6, 1]. Let  $\epsilon_s(\theta) := r_s - x_s^\top \theta$  and  $\epsilon_s^2(\theta) := (\epsilon_s(\theta))^2$ . With  $L$  and  $\iota$  defined in Algorithm 1, we define our confidence set after  $k$  time steps as

$$\Theta_k := \cap_{\ell=1}^L \Theta_k^\ell \quad (3)$$

where

$$\Theta_k^\ell := \left\{ \theta \in \mathbb{B}_2^d(1) : \left| \sum_{s=1}^k \overline{(x_s^\top \mu)}_\ell \epsilon_s(\theta) \right| \leq \sqrt{\sum_{s=1}^k \overline{(x_s^\top \mu)}_\ell^2 \epsilon_s^2(\theta) \iota} + 2^{-\ell} \iota, \forall \mu \in \mathbb{B}_2^d(2) \right\}$$

and the clipping operator  $\overline{(z)}_\ell$  is defined in (2). The role of clipping is two-fold: (i) it allows us to factor out  $\sum_{s=1}^k \epsilon_s^2(\theta)$  by  $\sum_s \overline{(x_s^\top \mu)}_\ell^2 \epsilon_s^2(\theta) \leq (2^{-\ell})^2 \sum_{s=1}^k \epsilon_s^2(\theta)$  and (ii) the lower order term is reduced to the order of  $2^{-\ell}$ . Both properties are critical in obtaining variance-adaptive regret bounds as discussed in [31]. The true parameter is contained in our confidence set with high probability as follows.

**Lemma 1.** (Confidence set) Let  $L$ ,  $\iota$ , and  $\delta$  be given as those in Algorithm 1. Then,

$$\mathbb{P}(\mathcal{E}_1 := \{\forall k \in [K], \theta^* \in \Theta_k\}) \geq 1 - \delta.$$

In fact, in our algorithm, we use the confidence set of  $\cap_{s=1}^{k-1} \Theta_s$  at time step  $k$  for a technical reason. VOFUL2 has the following regret bound.

**Theorem 1.** VOFUL2 satisfies, with probability at least  $1 - 2\delta$ ,

$$\mathcal{R}^K = \tilde{O} \left( d^{1.5} \sqrt{\sum_{k=1}^K \sigma_k^2 \ln(1/\delta)} + d^2 \ln(1/\delta) \right)$$

---

**Algorithm 1** VOFUL2

---

- 1: **Initialize:**  $L = 1 \vee \lfloor \log_2(K) \rfloor$  where  $\iota = 128 \ln((12K2^L)^{d+2}/\delta)$  and  $\delta \leq e^{-1}$ .
  - 2: **for**  $k = 1, 2, \dots, K$  **do**
  - 3:   Observe a decision set  $\mathcal{X}_k \subseteq \mathbb{B}_2^d(1)$ .
  - 4:   Compute the optimistic arm as following:  $x_k = \arg \max_{x \in \mathcal{X}_k} \max_{\theta \in \cap_{s=1}^{k-1} \Theta_s} x^\top \theta$  where  $\Theta_s$  is defined in (3).
  - 5:   Receive a reward  $r_k$ .
  - 6: **end for**
- 

where  $\tilde{O}$  hides poly-logarithmic dependence on  $\{d, K, \sum_{k=1}^K \sigma_k^2, \ln(1/\delta)\}$ .

Note that one can also show that VOFUL2 can be slightly modified to achieve the regret bound of  $\tilde{O}(\min\{d\sqrt{K\ln(1/\delta)}, d^{1.5}\sqrt{\sum_{k=1}^K \sigma_k^2 \ln(1/\delta)}\} + d^2 \ln(1/\delta))$ , thus being no worse than OFUL. We postpone the proof of this to Section A.4 to avoid clutter.

**Properties of the confidence sets and implications on the regret.** Before presenting the proof of Theorem 1, we provide some key properties of our confidence set (Lemma 3) and the intuition behind our regret bound. First, let us describe a few preliminaries. Define

$$W_{\ell, k-1}(\mu) := 2^{-\ell} I + \sum_{s=1}^{k-1} \left(1 \wedge \frac{2^{-\ell}}{|x_s^\top \mu|}\right) x_s x_s^\top.$$

Let  $\theta_k$  be the maximizer of the optimization problem at line 4 of Algorithm 1 and define  $\mu_k = \theta_k - \theta^*$ . For brevity, we use a shorthand of

$$W_{\ell, k-1} := W_{\ell, k-1}(\mu_k) = 2^{-\ell} I + \sum_{s=1}^{k-1} \left(1 \wedge \frac{2^{-\ell}}{|x_s^\top \mu_k|}\right) x_s x_s^\top.$$

Finally, we need to define the following event regarding the concentration of the empirical variance around the true variance:

$$\mathcal{E}_2 := \left\{ \forall k \in [K], \sum_{s=1}^k \epsilon_s^2(\theta^*) \leq \sum_{s=1}^k 8\sigma_s^2 + 4 \log\left(\frac{4K(\log_2(K) + 2)}{\delta}\right) \right\},$$

which is true with high probability as follows.

**Lemma 2.** *We have  $\mathbb{P}(\mathcal{E}_2) \geq 1 - \delta$*

*Proof.* The proof is a direct consequence of Lemma 13 in our appendix. □

Let  $\ell_k$  be the integer  $\ell$  such that  $x_k^\top \mu_k \in (2 \cdot 2^{-\ell}, 2 \cdot 2^{-\ell+1}]$  and define  $A_k := \sum_{s=1}^k \sigma_s^2$ . Lemma 3 below states the properties of our confidence set.

**Lemma 3.** *Suppose the events  $\mathcal{E}_1$  and  $\mathcal{E}_2$  are true. Then, for any  $k$  with  $\ell_k = \ell$ ,*

(i) *For some absolute constant  $c_1$ ,*

$$\|\mu_k\|_{W_{\ell, k-1}}^2 \leq 2^{-\ell} \sqrt{128 A_{k-1} \iota} + 11 \cdot 2^{-\ell} \iota \leq c_1 2^{-\ell} (\sqrt{A_{k-1} \iota} + \iota),$$

(ii) *There exists an absolute constant  $c_2$  such that  $x_k^\top \mu_k \leq c_2 \|x_k\|_{W_{\ell, k-1}^{-1}}^2 (\sqrt{A_{k-1} \iota} + \iota)$ .*

The key difference between Lemma 3 and the results of Zhang et al. [31] is that we use the norm notations, although the norm involves a rather complicated matrix  $W_{\ell, k-1}$ . This opens up possibilities of analyzing the regret of VOFUL2 with existing tools such as applying Cauchy-Schwarz inequality and the elliptical potential lemma [1, 5, 16]. In particular, Lemma 3(ii) seems useful because if we had such a result with  $W_{\ell, k-1}$  replaced by  $V_{k-1} = \lambda I + \sum_{s=1}^{k-1} x_s x_s^\top$ , then we would have, ignoring the additive term  $\iota$ ,

$$x_k^\top \mu_k \leq \|x_k\|_{V_{k-1}^{-1}}^2 \sqrt{\sum_{s=1}^{k-1} \sigma_s^2 \iota}.$$

Together with the optimism and the standard elliptical potential lemma (see Section 3.1 for details), this leads to

$$\mathcal{R}^K \leq \sum_{k=1}^K x_k^\top \mu_k \leq c_2 \sum_{k=1}^K \|x_k\|_{V_{k-1}^{-1}}^2 \sqrt{\sum_{s=1}^{k-1} \sigma_s^2 \iota} \leq c_2 \cdot O(d \log(T/d)) \cdot \sqrt{\sum_{s=1}^K \sigma_s^2 \iota}.$$

Since  $\iota$  is linear in  $d$ , we would get the regret bounded by the order of  $d^{1.5} \sqrt{\sum_{k=1}^K \sigma_k^2}$ , roughly speaking. However, the discrepancy between  $W_{\ell, k-1}$  and  $V_{k-1}$  is not trivial to resolve, especially due to the fact that Lemma 3(ii) has  $\mu_k$  on both left and the right hand side. That is,  $\mu_k$  is the key quantity that we need to understand, but we are bounding  $x_k \mu_k$  as a function of  $\mu_k$ . The novelty of our analysis of regret is exactly at relating  $W_{\ell, k-1}$  to  $V_{k-1}$  via a novel peeling-based analysis, which we present below.

### 3.1 Proof of Theorem 1

Throughout the proof, we condition  $\mathcal{E}_1$  and  $\mathcal{E}_2$  where each one is true with probability at least  $1 - \delta$ , as shown in Lemma 1 and 2 respectively. For our regret analysis, it is critical to use Lemma 4 below, which we call the elliptical ‘count’ lemma. This lemma is a generalization of Lattimore and Szepesvári [16, Exercise 19.3], which was originally used therein to improve the dependence of the range of the expected rewards in the regret bound. Similar lemmas have been used in parallel studies [12, 22]. In particular, He et al. [12] employ a lemma similar to elliptical potential count and peeling technique for the regret analysis for the linear MDP as well, which we compare in detail in Section C due to space constraint. We remark that a similar strategy appears in disguise in Russo and Van Roy [20, Proposition 3] as well.

**Lemma 4.** (Elliptical potential count) *Let  $x_1, \dots, x_k \in \mathbb{R}^d$  be a sequence of vectors with  $\|x_s\|_2 \leq X$  for all  $s \in [k]$ . Let  $V_k = \tau I + \sum_{s=1}^k x_s x_s^\top$  for some  $\tau > 0$ . Let  $J \subseteq [k]$  be the set of indices where  $\|x_s\|_{V_{s-1}^{-1}}^2 \geq q$ . Then,*

$$|J| \leq \frac{2}{\ln(1+q)} d \ln \left( 1 + \frac{2/e}{\ln(1+q)} \frac{X^2}{\tau} \right).$$

As the name explains, the lemma above bounds how many times  $\|x_s\|_{V_{s-1}^{-1}}^2$  can go above a given value  $q > 0$ , which is different from existing elliptical potential lemmas that bound the sum of  $\|x_s\|_{V_{s-1}^{-1}}^2$ . Let  $\theta_k$  be the  $\theta$  that maximizes the optimization problem at line 4 of Algorithm 1. We start by the usual optimism-based bounds: due to  $\mathcal{E}_1$ , we have

$$\mathcal{R}^K = \sum_{k=1}^K (\max_{x \in \mathcal{X}_k} (x^\top \theta^* - x^\top \theta_k)) \leq \sum_{k=1}^K (\max_{x \in \mathcal{X}_k, \theta \in \Theta_k} x^\top \theta - x^\top \theta_k) = \sum_{k=1}^K x_k^\top (\theta_k - \theta^*) = \sum_{k=1}^K x_k^\top \mu_k.$$

We now take a peeling-based regret analysis that is quite different from existing analysis techniques:

$$\mathcal{R}^K \leq \sum_{k=1}^K x_k^\top (\theta_k - \theta^*) \leq 2^{-L} K + \sum_{\ell=1}^L 2^{-\ell+2} \sum_{k=1}^K \mathbb{1} \left\{ x_k^\top \mu_k \in (2 \cdot 2^{-\ell}, 2 \cdot 2^{-\ell+1}] \right\},$$

where  $L$  is defined in Algorithm 1. Given  $\ell$  and  $k$ , let  $n_{k,\ell}$  be such that  $\max_{v: k \leq v \leq K, \ell_v = \ell} |x_k^\top \mu_v| \in (2^{-\ell+n}, 2^{-\ell+n+1}]$  if such  $n$  satisfies  $n \geq 1$ . Otherwise, set  $n_{k,\ell} = 0$ , which means  $\max_{v: k \leq v \leq K, \ell_v = \ell} |x_k^\top \mu_v| \leq 2^{-\ell+n+1}$  with  $n = 0$ . We then define  $G_{\ell,n} := \{s \in [K-1] : \ell_s = \ell, n_{s,\ell} = n\}$  and let  $G_{\ell,n}[k] := G_{\ell,n} \cap [k]$ . Then,

$$\sum_{k=1}^K \mathbb{1} \left\{ x_k^\top \mu_k \in (2 \cdot 2^{-\ell}, 2 \cdot 2^{-\ell+1}] \right\} = \sum_{k=1}^K \mathbb{1} \{ \ell_k = \ell \} \leq 1 + \sum_n \sum_{s \in G_{\ell,n}} 1.$$

Letting  $V_{\ell,n,k-1} := 2^{-\ell} I + \sum_{s \in G_{\ell,n}[k-1]} x_s x_s^\top$ , a comparison between two matrices  $W$  and  $V$  is given as follows for every  $v \in \{k, \dots, K\}$ :

$$\begin{aligned} W_{\ell,k-1}(\mu_v) &= 2^{-\ell} I + \sum_{s=1}^{k-1} \left( 1 \wedge \frac{2^{-\ell}}{|x_s^\top \mu_v|} \right) x_s x_s^\top \geq 2^{-\ell} I + \sum_{s \in G_{\ell,n}[k-1]} \left( 1 \wedge \frac{2^{-\ell}}{2^{-\ell+n+1}} \right) x_s x_s^\top \\ &\geq c \cdot 2^{-n} V_{\ell,n,k-1}. \end{aligned} \tag{4}$$

For  $k \in G_{\ell,n}[K-1]$  and  $u = \arg \max_{v:k \leq v \leq K, \ell_v = \ell} |x_k^\top \mu_v|$ , we have

$$\begin{aligned} 2^{-\ell+n} &< |x_k \mu_u| \leq \|x_k\|_{W_{\ell,k-1}^{-1}(\mu_u)} \|\mu_u\|_{W_{\ell,k-1}(\mu_u)} \\ &\leq \|x_k\|_{W_{\ell,k-1}^{-1}(\mu_u)} \|\mu_u\|_{W_{\ell,u-1}(\mu_u)} \quad (u \geq k) \\ &\leq c\sqrt{2^n} \|x_k\|_{V_{\ell,n,k-1}^{-1}} \sqrt{2^{-\ell}(\sqrt{A_K \iota} + \iota)}. \end{aligned} \quad ((4) \text{ and Lemma 3(i)})$$

Consequently,  $\|x_k\|_{V_{\ell,n,k-1}^{-1}}^2 \geq \frac{2^{-\ell+n}}{\sqrt{A_K \iota + \iota}}$ . Thus, using Lemma 4 with  $\tau = 2^{-\ell}$ ,

$$\begin{aligned} 1 + \sum_{n=0}^{\ell} \sum_{k \in G_{\ell,n}} 1 &\leq 1 + \sum_n \sum_{k \in G_{\ell,n}} \mathbb{1} \left\{ \|x_k\|_{V_{\ell,n,k-1}^{-1}}^2 \geq c \frac{2^{-\ell+n}}{\sqrt{A_K \iota + \iota}} \right\} \\ &\leq 1 + c \sum_n 2^{\ell-n} (\sqrt{A_K \iota} + \iota) d \ln \left( 1 + c 4^\ell (\sqrt{A_K \iota} + \iota) \right) \\ &\leq c 2^\ell (\sqrt{A_K \iota} + \iota) d \ln \left( 1 + c 4^\ell (\sqrt{A_K \iota} + \iota) \right). \end{aligned}$$

where we use the fact that  $1/\ln(1+q) \leq c/q$  for an absolute constant  $c$  if  $q$  is bounded by an absolute constant. Finally,

$$\begin{aligned} \sum_{k=1}^K x_k^\top \mu_k &\leq 2^{-L} K + c \sum_{\ell=1}^L 2^{-\ell} 2^\ell (\sqrt{A_K \iota} + \iota) d \ln \left( 1 + c 4^\ell (\sqrt{A_K \iota} + \iota) \right) \\ &= 2^{-L} K + c L (\sqrt{A_K \iota} + \iota) d \ln \left( 1 + c 4^L (\sqrt{A_K \iota} + \iota) \right). \end{aligned}$$

We choose  $L = 1 \vee \lceil \log_2(K) \rceil$ , which leads to  $\mathcal{R}^K \leq c (\sqrt{A_{K-1} \iota} + \iota) d \ln^2 (1 + c K^2 (\sqrt{A_{K-1} \iota} + \iota))$ . This concludes the proof.

## 4 Linear Mixture MDP

As linear bandits and linear mixture MDPs have quite a similar nature, we bring the techniques in our analysis of VOFUL2 to improve the regret bound of VARLin of Zhang et al. [31]. A key feature of linear mixture MDP setting is that one can estimate the upper bound of the variance as it is a quadratic function of  $\theta^*$  while linear bandits do not have a structural assumption on the variance. Thanks to such a structural property, we obtain a slightly better dependence on the dimension  $d$ . The confidence set derived for our proposed algorithm is slightly different from that of VARLin as ours is defined with  $\forall \mu \in \mathbb{B}_1^d(2)$  rather than an  $\epsilon$ -net. Our version of VARLin, which we call VARLin2, is described in Algorithm 2. Given  $s_h^k$  and  $a_h^k$ , let us define  $P_{s_h^k, a_h^k}(V_{h+1}^k) := \mathbb{E}_{s' \sim P_{s_h^k, a_h^k}}[V_{h+1}^k(s')]$  and  $x_{k,h}^m := [P_{s_h^k, a_h^k}^1(V_{h+1}^k)^{2^m}, \dots, P_{s_h^k, a_h^k}^d(V_{h+1}^k)^{2^m}]^\top$  and let  $L$ ,  $\iota$ , and  $\delta$  be given as define Algorithm 2. Let  $\epsilon_{v,u}^m(\theta) := \theta^\top x_{v,u}^m - (V_{u+1}^v(s_{u+1}^v))^{2^m}$  for  $(v, u) \in [K] \times [H]$ ,  $m \in \{0, 1, \dots, L\}$  where  $L$  is defined in Algorithm 2.

We construct our confidence set as

$$\Theta_k := \bigcap_{m=0}^L \bigcap_{i \in [L]} \bigcap_{\ell \in [L]} \Theta_k^{m,i,\ell} \quad (5)$$

where we define  $\Theta_k^{m,i,\ell}$  below, based on the data collected up to episode  $k-1$ . First, let

$$\eta_{k,h}^m := \max_{\theta \in \Theta_k} \{\theta^\top x_{k,h}^{m+1} - (\theta^\top x_{k,h}^m)^2\}$$

$$\text{and } \mathcal{T}_{k,h}^{m,i} := \{(v, u) \in ([k] \times [H]) \cup (\{k\} \times [h]) : \eta_{v,u}^m \in (2^{-i}, 2^{1-i}]\}.$$

We naturally define  $\mathcal{T}_{k,h}^{m,L+1} := \{(v, u) \in \mathcal{T}_{k,h}^{m,i} : \eta_{v,u}^m \leq 2^{-L}\}$ . With  $\iota$  defined in Algorithm 2, define

$$\Theta_{k-1}^{m,i,\ell} := \left\{ \theta \in \mathbb{B}_1^d(1) : \left| \sum_{(v,u) \in \mathcal{T}_{k-1,H}^{m,i}} \overline{((x_{v,u}^m)^\top \mu)_\ell} \epsilon_{v,u}^m(\theta) \right| \leq \right.$$

---

**Algorithm 2** VARLin2

---

```

1: Initialize:  $L = \lceil \log_2 HK \rceil + 1$ ,  $\iota = 3 \ln((2HK)^{2(d+3)}/\delta)$ ,  $\delta \leq e^{-1}$ .
2: for  $k = 1, 2, \dots, K$  do
3:   for  $h = H, \dots, 1$  do
4:     For each  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , define  $Q_h^k(s, a) = \min\{1, r(s, a) + \max_{\theta \in \Theta_{k-1}} \sum_{i=1}^d \theta_i P_{s,a}^i V_{h+1}^k\}$ 
       where  $\Theta_{k-1}$  is defined in Lemma 5
5:     For each state  $s$ ,  $V_h^k(s) = \max_{a \in \mathcal{A}} Q_h^k(s, a)$ .
6:   end for
7:   for  $h = 1, \dots, H$  do
8:     Choose  $a_h^k = \arg \max_{a \in \mathcal{A}} Q_h^k(s_h^k, a)$ .
9:     Observe a reward  $r_h^k$  and the next state  $s_{h+1}^k$ .
10:  end for
11: end for

```

---

$$4 \sqrt{\sum_{(v,u) \in \mathcal{T}_{k-1,H}^{m,i}} \left( (x_{v,u}^m)^\top \mu \right)_\ell^2 \eta_{v,u}^m \iota + 4 \cdot 2^{-\ell} \iota, \forall \mu \in \mathbb{B}_1^d(2)} \quad (6)$$

We show that the confidence set is correct w.h.p. in the following lemma.

**Lemma 5.** (*Confidence set for MDP*)  $\mathbb{P}(\forall k \in [K], \theta^* \in \Theta_k) \geq 1 - \delta$ .

The consequence is that the  $Q$  values computed in VarLin2 is optimistic with high probability due to the following property:

**Lemma 6.** For every  $k \geq 1$ ,  $\theta^* \in \Theta_k \implies \forall h, s, a : Q_h^k(s, a) \geq Q^*(s, a)$ .

Now with the confidence set defined above we state our main result.

**Theorem 2.** With probability at least  $1 - \delta$ ,

$$\mathcal{R}^K = \sum_{k=1}^K [V^*(s_1^k) - V^k(s_1^k)] = \tilde{O}(d\sqrt{K \log^2(1/\delta)} + d^2 \log(1/\delta)).$$

where  $\tilde{O}$  hides poly-logarithmic dependence on  $\{d, K, H, \ln(1/\delta)\}$ .

#### 4.1 Proof of Theorem 2

The main idea of the proof is to infuse a peeling-based argument together with elliptical potential count lemma to both the planning horizon and episode. Noting that the regret of the predicted variance is controlled by the variance of variance, one can expect to reduce the total regret using this information, as done in [31]. We begin by introducing relevant quantities that are parallel with those in linear bandits. Let us first introduce the following lemma.

**Lemma 7.** Let  $x_1, \dots, x_T \in \mathbb{R}^d$  be a sequence of vectors with  $\|x_t\|_2 \leq X$  for all  $t \in [T]$ . Let  $V_t = \lambda I + \sum_{s=1}^t x_s x_s^\top$ . Let  $0 = \tau_0 < \tau_1 < \tau_2 < \dots < \tau_z = T$  where  $\tau_i$  marks the last time step of the  $i$ -th block formed by  $\{\tau_{i-1} + 1, \dots, \tau_i\}$  for all  $i \in \{0, \dots, z\}$ . Let  $\text{anc}(t)$  be the ‘anchor’ of  $t$ , the last time step of the  $i$ -th block such that the  $(i+1)$ -th block contains  $t$ :  $\text{anc}(t) = \max\{\tau_i : i \in \{0, \dots, z\}, \tau_i < t\}$ . Let  $r > 0$ . Define  $J \subseteq [T]$  to be the set of indices  $t$  such that  $\|x_t\|_{V_{\text{anc}(t)}^{-1}}^2 > r \|x_t\|_{V_{t-1}^{-1}}^2$  is true for the first time in the block containing  $t$ . Then,

$$|J| \leq \frac{2}{\ln(r)} d \ln \left( 1 + \frac{2/e}{\ln(r)} \frac{X^2}{\lambda} \right).$$

To keep track of episode-horizon index pairs concisely, we use a *flat index*  $t \in [T]$  where  $T := HK$ . Specifically, an episode  $k$  and a horizon  $h$  corresponds to the flat index  $t = (k-1)H + h$ . Let  $\mathbf{t}(k, h) := (k-1)H + h$ . Let  $\mathbf{k}(t)$  and  $\mathbf{h}(t)$  be the mapping from  $t$  to its corresponding episode and horizon index respectively so that  $k = \mathbf{k}(\mathbf{t}(k, h))$  and  $h = \mathbf{h}(\mathbf{t}(k, h))$ . By taking  $\tau_k$  in Lemma 7 as  $\mathbf{t}(k, H)$ , we have that  $\text{anc}(t) := \mathbf{t}(t-1, H)$ . We define  $\mathcal{T}_t^{m,i} := \{\mathbf{t}(k', h') : (k', h') \in \mathcal{T}_{k,h}^{m,i}\}$  and  $\mu_t^m := \mu_{\mathbf{k}(t), \mathbf{h}(t)}^m$ . Similarly, we define  $x_t^m$ , etc., by replacing the subscript  $k, h$  with  $t$ . Hereafter, any appearance of subscript  $k, h$  can be replaced with  $t$  such that  $t = \mathbf{t}(k, h)$  without changing the meaning.

Given  $m, k$  and  $h$ , we define  $\ell_{k,h}^m$  as the integer  $\ell$  such that  $(x_{k,h}^m)^\top \mu_{k,h}^m \in (2 \cdot 2^{-\ell}, 2 \cdot 2^{-\ell+1}]$  where  $\mu_{k,h}^m := \theta_{k,h}^m - \theta^*$  and  $\theta_{k,h}^m = \arg \max_{\theta \in \Theta_{k-1}} |\{\theta^\top x_t^{m+1} - (\theta^\top x_t^m)^2\}|$ . For simplicity, we abbreviate  $\ell_{k,h}^m$  by  $\ell$ . Define  $W_t^{m,i,\ell}(\mu) := 2^{-\ell}I + \sum_{s \in \mathcal{T}_t^{m,i}} \left(1 \wedge \frac{2^{-\ell}}{|(x_s^m)^\top \mu|}\right) x_s^m (x_s^m)^\top$  and introduce a shorthand  $W_{\text{anc}(t)}^{m,i,\ell} := W_{\text{anc}(t)}^{m,i,\ell}(\mu_t^m)$  as before. With the definition above we have the following:

$$2^{-\ell} \|\mu_t^m\|^2 + \sum_{s \in \mathcal{T}_{\text{anc}(t)}^{m,i}} \overline{((x_s^m)^\top \mu_t^m)_\ell} (x_s^m)^\top \mu_t^m = \|\mu_t^m\|_{W_{\text{anc}(t)}^{m,i,\ell}}^2.$$

We now show the key result of the confidence set of VarLin2 that parallels Lemma 3 for bandits.

**Lemma 8.** Fix  $m \in \{0, \dots, L_m\}$  and  $i \in [L]$ . Let  $t \in \mathcal{T}_T^{m,i}$ . Then, with  $\ell = \ell_t^m$ ,

$$\|\mu_t^m\|_{W_{\text{anc}(t)}^{m,i,\ell}}^2 \leq c_M \cdot \max\{\sqrt{2^{-i}L}, \sqrt{2^{-\ell}L}\}$$

where  $c_M > 0$  is an absolute constant.

What is different from the linear bandit problem is that we do not update  $\theta$  until the planning horizon is over and an additional layer for peeling is imposed on variance. In [31], the authors introduce an indicator  $I_h^k$  to characterize episode-horizon pairs for which growth of the norm of  $\mu$  with respect to  $W_t$  is controlled by the norm with respect to  $W_{\text{anc}(t)}$  with  $d^2$  growth rate, i.e.,

$$I_h^k := \mathbb{1}\{\|\mu_t\|_{W_t^{m,i,\ell}} \leq 4(d+2)^2 \|\mu_t\|_{W_{\text{anc}(t)}^{m,i,\ell}}\}.$$

where  $t = t(k, h)$ . Our novelty novelty lies in being able to replace  $4(d+2)^2$  above by a constant rate  $r$  that is set to 2 later (modulo some differences due to technical reasons). To distinguish we denote such a set by  $I_{k,h}$ . See B.5 for the definition of  $I_{k,h}$  and the proof of the following lemma.

**Lemma 9.**  $\sum_{k=1}^K \sum_{h=1}^{H-1} I_{k,h} - I_{k,h+1} \leq O\left(\frac{d}{\ln(r)}\right) \log(dHK(1 + d^2/\ln(r)))$  for  $r > 0$ .

Note here that once we fix  $r$  such as  $r = 2$ , the bound can be replaced by  $O(d \log^5(dHK))$ . We now use the following regret decomposition due to [31] which just come from replacing  $I_h^k$  by  $I_{k,h}$ .

**Lemma 10.** (Zhang et al. [31])  $\mathcal{R}^K \leq \text{Reg}_1 + \text{Reg}_2 + \text{Reg}_3 + \sum_{k=1}^K \sum_{h=1}^{H-1} (I_{k,h} - I_{k,h+1})$  where  $\text{Reg}_1 = \sum_{k,h} (P_{s_h^k, a_h^k} V_{h+1}^k - V_{h+1}^k(s_{h+1}^k)) I_{k,h}$ ,  $\text{Reg}_2 = \sum_{k,h} (V_h^k(s_h^k) - r_h^k - P_{s_h^k, a_h^k} V_{h+1}^k) I_{k,h}$ , and  $\text{Reg}_3 = \sum_{k=1}^K (\sum_{h=1}^H r_h^k - V_1^{\pi^k}(s_1^k))$ .

Let  $\check{x}_{k,h}^m := x_{k,h}^m I_{k,h}$  and define  $R_m, M_m$  as

$$R_m := \sum_{k,h} (\check{x}_{k,h}^m)^\top \mu_{k,h}^m, \quad \text{and} \quad M_m := \sum_{k,h} (P_{s_h^k, a_h^k} (V_{h+1}^k)^{2^m} - (V_{h+1}^k(s_{h+1}^k))^{2^m}) I_{k,h}.$$

We have that  $\text{Reg}_1 = M_0$  and  $\text{Reg}_2 \leq R_0$  since

$$Q_h^k(s, a) - r(s, a) - P_{s,a} V_{h+1}^k \leq \max_{\theta \in \Theta_k} x_{k,h}^0 (\theta - \theta^*).$$

To proceed, we first note that  $\sum_{k,h} (I_{k,h} - I_{k,h+1})$  and  $\text{Reg}_3$  are bounded by  $O(d \log^5(dHK))$  and  $O(\sqrt{K \log(1/\delta)})$  respectively from Lemma 9 and Lemma 16. Since  $\text{Reg}_1 + \text{Reg}_2 \leq R_0 + M_0$ , it remains to find a bound on  $R_0 + M_0$ . This, however, involves solving a series of recursive inequalities. We leave the details in the appendix and provide a high-level description below.

Let us begin with Lemma 15 in the appendix that shows

$$|M_m| \leq \tilde{O}(\sqrt{M_{m+1}} + d + 2^{m+1}(K + R_0) \log(1/\delta) + \log(1/\delta)) \quad (7)$$

where the RHS is a function of  $\sqrt{M_{m+1}}$  and  $\sqrt{R_0}$ . Taking Proposition 1 below for granted and combining it with the relation from Zhang et al. [31, Eq. (57)] showing

$$\sum_{k,h} \eta_{k,h}^m I_{k,h} \leq M_{m+1} + O(d \log^5(dHK)) + 2^{m+1}(K + R_0) + R_{m+1} + 2R_m,$$

one arrives at  $R_m \leq \tilde{O}(d^{1/2} \sqrt{(M_{m+1} + 2^{m+1}(K + R_0) + R_{m+1} + 2R_m + d)})$ . This bound is the key improvement we obtain via our peeling-based regret analysis. Specifically, the bound on  $R_m$  obtained by [31] has  $d^4$  and  $d^6$  in place of  $d^{1/2}$  and  $d$  above.



We first show how our regret bound helps in obtaining the stated regret bound and then present Proposition 1. Noting that both  $R_L$  and  $M_L$  are trivially bounded by  $HK$ , one can solve the series of inequalities on  $R_m$  and  $|M_m|$  to obtain a bound on  $R_0$ :

$$R_0 \leq \tilde{O}(d^2 \log(1/\delta) + \sqrt{d^2(K + R_0) \log(1/\delta)}). \quad (8)$$

Solving it for  $R_0$ , we obtain  $R_0 \leq \tilde{O}(d^2 \log(1/\delta) + \sqrt{d^2 K \log(1/\delta)})$ . One can now plug in  $R_0$  to the bound (7) and obtain a bound on  $|M_0|$  in a similar way as follows, which concludes the proof:

$$|M_0| \leq \tilde{O}(d\sqrt{K \log^2(1/\delta)} + d^2 \log(1/\delta)).$$

We now show the key proposition that allows us to improve the bound on  $R_m$ . In the paper by Zhang et al. [31],  $d^4$  was derived while we propose the following.

**Proposition 1.** *Let  $\check{\eta}_{k,h}^m := \eta_{k,h}^m I_{k,h}$ . Then, we have*

$$R_m \leq O(d^{0.5} \log^{2.5}(HK) \sqrt{(1 + \sum_{k,h} \check{\eta}_{k,h}^m) \iota \log(d\iota)} + d \log^3(HK) \iota \log(d\iota))$$

*Proof.* Define  $\mathcal{T}^{m,i,\ell} := \{t \in \mathcal{T}_T^{m,i} : (x_t^m)^\top \mu_t^m \in (2 \cdot 2^{-\ell}, 2 \cdot 2^{1-\ell}]\}$  and split the time steps  $\mathcal{T}^{m,i,\ell}$  by

$$\mathcal{T}^{m,i,\ell,\langle 1 \rangle} := \left\{ t \in \mathcal{T}^{m,i,\ell} : \left\| \mu_{k,h}^m \right\|_{W_{\text{anc}(a)}^{m,i,\ell}} \leq c_M \sqrt{2^{-i_\ell}} \right\} \quad \text{and} \quad \mathcal{T}^{m,i,\ell,\langle 2 \rangle} := \mathcal{T}^{m,i,\ell} \setminus \mathcal{T}^{m,i,\ell,\langle 1 \rangle}.$$

Having defined  $I_{k,h}$  with  $r = 2$ , we also denote  $\check{\mathcal{T}}^{m,i,\ell,\langle z \rangle} := \mathcal{T}^{m,i,\ell,\langle z \rangle} \cap \{t \in [T] : I_t = 1\}$ . Now we decompose  $R_m$  as

$$R_m = \sum_{t \in [T]} (\check{x}_t^m)^\top \mu_t^m = \sum_{i,\ell} \sum_{t \in \check{\mathcal{T}}^{m,i,\ell,\langle 1 \rangle}} (\check{x}_t^m)^\top \mu_t^m + \sum_{i,\ell} \sum_{t \in \check{\mathcal{T}}^{m,i,\ell,\langle 2 \rangle}} (\check{x}_t^m)^\top \mu_t^m.$$

Fix  $m, i$  and  $\ell$  and focus on  $\sum_{t \in \check{\mathcal{T}}^{m,i,\ell,\langle z \rangle}} (\check{x}_t^m)^\top \mu_t^m$  for  $z = 1, 2$ . Hereafter, we omit the superscripts and subscripts of  $(m, i, \ell)$  to avoid clutter, unless there is a need. Note that for  $t \in \check{\mathcal{T}}^{(1),n}$  and  $b$  such that  $t < b \in \mathcal{T}^{(1)}$ ,

$$\begin{aligned} W_{\text{anc}(t)}(\mu_b) &= 2^{-\ell} I + \sum_{t' \in \mathcal{T}_{\text{anc}(t)}^{m,i,\ell}} (1 \wedge \frac{2^{-\ell}}{|x_{t'}^\top \mu_b|}) x_{t'} x_{t'}^\top \geq 2^{-\ell} I + \sum_{t' \in \mathcal{T}_{\text{anc}(t)}^{(1),n}} (1 \wedge \frac{2^{-\ell}}{2^{-\ell+n+1}}) x_{t'} x_{t'}^\top \\ &\geq 2^{-\ell} I + 2^{-n-1} \sum_{t' \in \mathcal{T}_{\text{anc}(t)}^{(1),n}} x_{t'} x_{t'}^\top \\ &\geq c 2^{-n} V_{\text{anc}(t)}^{(1),n}. \end{aligned}$$

For the same  $t$ , letting  $b = \arg \max_{t \leq b' \in \check{\mathcal{T}}^{(1)}} |x_{t'}^\top \mu_{b'}|$ ,

$$\begin{aligned} 2^{-\ell+n} &\leq |x_t^\top \mu_b| \leq \|x_t\|_{W_{\text{anc}(t)}^{-1}(\mu_b)} \|\mu_b\|_{W_{\text{anc}(t)}(\mu_b)} \leq \sqrt{2^n} \|x_t\|_{(V_{\text{anc}(t)}^{(1),n})^{-1}} \|\mu_b\|_{W_{\text{anc}(t)}(\mu_b)} \\ &\leq c \sqrt{r} \sqrt{2^n} \|x_t\|_{(V_{t-1}^{(1),n})^{-1}} \sqrt{2^{-i_\ell}} \quad (\text{by } b \in \check{\mathcal{T}}^{(1)}) \end{aligned}$$

This implies that  $\|x_t\|_{(V_{t-1}^{(1),n})^{-1}}^2 \geq c \frac{2^{-2\ell+n}}{r 2^{-i_\ell}}$ . Thus,

$$\begin{aligned} \sum_{t \in \check{\mathcal{T}}^{(1)}} x_t^\top \mu_t &\leq c 2^{-\ell} \sum_{t \in \check{\mathcal{T}}^{(1)}} 1 \leq c 2^{-\ell} \sqrt{|\check{\mathcal{T}}^{(1)}|} \sum_{t \in \check{\mathcal{T}}^{(1)}} 1 \leq c 2^{-\ell} \sqrt{|\check{\mathcal{T}}^{(1)}|} \sum_{n=0}^{\ell} \sum_{t \in \check{\mathcal{T}}^{(1),n}} 1 \\ &\leq c 2^{-\ell} \sqrt{|\check{\mathcal{T}}^{(1)}| \sum_{n=0}^{\ell} \sum_{t \in \check{\mathcal{T}}^{(1),n}} \mathbb{1} \left\{ \|x_t\|_{(V_{t-1}^{(1),n})^{-1}}^2 \geq c \frac{2^{-2\ell+n}}{r 2^{-i_\ell}} \right\}} \\ &\leq c 2^{-\ell} \sqrt{|\check{\mathcal{T}}^{(1)}| \sum_{n=0}^{\ell} \sum_{t \in \check{\mathcal{T}}^{(1),n}} \mathbb{1} \left\{ \|x_t\|_{(V_{t-1}^{(1),n})^{-1}}^2 \geq c \frac{2^{-2\ell+n}}{r 2^{-i_\ell}} \right\}} \\ &\quad (\check{\mathcal{T}}^{(1),n} \subseteq \mathcal{T}^{(1),n}) \end{aligned}$$

$$\begin{aligned}
&\leq c2^{-\ell} \sqrt{|\check{\mathcal{T}}^{(1)}| \sum_{n=0}^{\ell} \frac{r2^{-i_{\ell}}}{2^{-2\ell+n}} d \ln \left( 1 + c \frac{r2^{-i_{\ell}}}{2^{-2\ell+n}2^{-\ell}} \right)} \\
&\leq c\sqrt{dr|\check{\mathcal{T}}^{(1)}|2^{-i_{\ell}} \ln(1 + cr\iota 8^{\ell})} \leq c\sqrt{dr(1 + \sum_{t \in \check{\mathcal{T}}^{(1)}} \eta_t^m) \iota \ln(1 + cr\iota 8^{\ell})}
\end{aligned}$$

where we use the fact that  $|\mathcal{T}^{(1)}| \cdot 2^{-i} \leq O(1 + \sum_{t \in \check{\mathcal{T}}^{(1)}} \eta_t^m)$ , which is straightforward by the definition. The summation over  $\check{\mathcal{T}}^{(2)}$  can be handled in a similar way and the details of the proof is provided in Section B.7 in our appendix.  $\square$

## 5 Conclusion

In this work, we have made significant improvements in the regret upper bounds for linear bandits and linear mixture MDPs by employing a novel peeling-based regret analysis based on the elliptical potential count lemma. Our study opens up numerous future research directions. First, the optimal regret rates are still not identified for these problems. It would be interesting to close the gap between the upper and lower bound. Second, our algorithms are not computationally tractable. We believe computationally tractable algorithms, even at the price of increased regret, may lead to practical algorithms. Finally, characterizing variance-dependent uncertainty in the linear regression setting without prior knowledge of variances is an interesting statistical problem on its own. Identifying novel estimators for it and proving their optimal coverage would be interesting.

## Acknowledgments and Disclosure of Funding

The authors thank Liyu Chen for finding an error in our earlier version. Insoon Yang is supported in part by the National Research Foundation of Korea (MSIT2020R1C1C1009766), the Information and Communications Technology Planning and Evaluation (IITP) grants (MSIT2022-0-00124, MSIT2022-0-00480), and Samsung Electronics. Kwang-Sung Jun is supported by Data Science Academy and Research Innovation & Impact at University of Arizona.

## References

- [1] Y. Abbasi-Yadkori, D. Pal, and C. Szepesvari. Improved Algorithms for Linear Stochastic Bandits. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1–19, 2011.
- [2] J.-Y. Audibert, R. Munos, and C. Szepesvari. Use of variance estimation in the multi-armed bandit problem. *NeurIPS Workshop on On-line Trading of Exploration and Exploitation Workshop*, 2006.
- [3] P. Auer. Using Confidence Bounds for Exploitation-Exploration Trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2002.
- [4] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The Nonstochastic Multiarmed Bandit Problem. *SIAM J. Comput.*, 32(1):48–77, 2003.
- [5] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- [6] V. Dani, T. P. Hayes, and S. M. Kakade. Stochastic Linear Optimization under Bandit Feedback. In *Proceedings of the Conference on Learning Theory (COLT)*, pages 355–366, 2008.
- [7] C. Dann, N. Jiang, A. Krishnamurthy, A. Agarwal, J. Langford, and R. E. Schapire. On Oracle-Efficient PAC RL with Rich Observations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [8] S. S. Du, Y. Luo, R. Wang, and H. Zhang. Provably Efficient Q-learning with Function Approximation via Distribution Shift Error Checking Oracle. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 8058–8068, 2019.

- [9] S. S. Du, J. D. Lee, G. Mahajan, and R. Wang. Agnostic  $q$ -learning with function approximation in deterministic systems: Near-optimal bounds on approximation error and sample complexity. *Advances in Neural Information Processing Systems*, 33, 2020.
- [10] F. Feng, R. Wang, W. Yin, S. S. Du, and L. Yang. Provably Efficient Exploration for Reinforcement Learning Using Unsupervised Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 22492–22504, 2020.
- [11] E. Hazan and S. Kale. Extracting certainty from uncertainty: Regret bounded by variation in costs. *Machine Learning*, 80(2):165–188, 2010.
- [12] J. He, D. Zhou, and Q. Gu. Logarithmic regret for reinforcement learning with linear function approximation. In *International Conference on Machine Learning*, pages 4171–4180. PMLR, 2021.
- [13] N. Jiang, A. Krishnamurthy, A. Agarwal, J. Langford, and R. E. Schapire. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.
- [14] C. Jin, Z. Yang, Z. Wang, and M. I. Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020.
- [15] A. Krishnamurthy, A. Agarwal, and J. Langford. Pac reinforcement learning with rich observations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [16] T. Lattimore and C. Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.
- [17] Y. Li, Y. Wang, and Y. Zhou. Nearly Minimax-Optimal Regret for Linearly Parameterized Bandits. In *Proceedings of the Conference on Learning Theory (COLT)*, pages 2173–2174, 2019.
- [18] D. Misra, M. Henaff, A. Krishnamurthy, and J. Langford. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 6961–6971. PMLR, 2020.
- [19] D. Pollard. Empirical processes: theory and applications. In *NSF-CBMS regional conference series in probability and statistics*, pages i–86. JSTOR, 1990.
- [20] D. Russo and B. Van Roy. Eluder dimension and the sample complexity of optimistic exploration. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2256–2264, 2013.
- [21] W. Sun, N. Jiang, A. Krishnamurthy, A. Agarwal, and J. Langford. Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Proceedings of the Conference on Learning Theory (COLT)*, pages 2898–2933. PMLR, 2019.
- [22] A. J. Wagenmaker, Y. Chen, M. Simchowitz, S. Du, and K. Jamieson. First-order regret in reinforcement learning with linear function approximation: A robust estimation approach. In *International Conference on Machine Learning*, pages 22384–22429. PMLR, 2022.
- [23] R. Wang, S. S. Du, L. Yang, and R. R. Salakhutdinov. On Reward-Free Reinforcement Learning with Linear Function Approximation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 17816–17826, 2020.
- [24] R. Wang, R. Salakhutdinov, and L. F. Yang. Provably Efficient Reinforcement Learning with General Value Function Approximation. *CoRR*, abs/2005.1, 2020.
- [25] Y. Wang, R. Wang, S. S. Du, and A. Krishnamurthy. Optimism in reinforcement learning with generalized linear function approximation. In *International Conference on Learning Representations (ICLR)*, 2020.
- [26] Z. Wen and B. Van Roy. Efficient exploration and value function generalization in deterministic systems. *Advances in Neural Information Processing Systems*, 26, 2013.
- [27] L. Yang and M. Wang. Sample-optimal parametric  $q$ -learning using linearly additive features. In *International Conference on Machine Learning*, pages 6995–7004. PMLR, 2019.

- [28] L. Yang and M. Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*, pages 10746–10756. PMLR, 2020.
- [29] A. Zanette, A. Lazaric, M. Kochenderfer, and E. Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pages 10978–10989. PMLR, 2020.
- [30] Z. Zhang, X. Ji, and S. Du. Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. In *Conference on Learning Theory*, pages 4528–4531. PMLR, 2021.
- [31] Z. Zhang, J. Yang, X. Ji, and S. S. Du. Variance-Aware Confidence Set: Variance-Dependent Bound for Linear Bandits and Horizon-Free Bound for Linear Mixture MDP. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [32] Z. Zhang, Y. Zhou, and X. Ji. Model-free reinforcement learning: from clipped pseudo-regret to sample complexity. In *International Conference on Machine Learning*, pages 12653–12662. PMLR, 2021.
- [33] D. Zhou, Q. Gu, and C. Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*, pages 4532–4576. PMLR, 2021.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
  - (b) Did you describe the limitations of your work? [\[Yes\]](#)
  - (c) Did you discuss any potential negative societal impacts of your work? [\[No\]](#) There is no negative impacts, to our knowledge.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#)
  - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#)
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[N/A\]](#)
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[N/A\]](#)
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[N/A\]](#)
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[N/A\]](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [\[N/A\]](#)
  - (b) Did you mention the license of the assets? [\[N/A\]](#)
  - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#)
5. If you used crowdsourcing or conducted research with human subjects...

- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

# Appendix

## Table of Contents

<b>A</b>	<b>Proofs for VOFUL2</b>	<b>13</b>
A.1	Proof of Lemma 4 . . . . .	13
A.2	Proof of Lemma 1 . . . . .	14
A.3	Proof of Lemma 3 . . . . .	15
A.4	$d\sqrt{K}$ regret bound of VOFUL2 . . . . .	16
A.5	Miscellaneous Lemmas . . . . .	17
<b>B</b>	<b>Proofs for VARLin2</b>	<b>18</b>
B.1	Proof of Lemma 6 . . . . .	18
B.2	Proof of Lemma 5 . . . . .	18
B.3	Proof of Lemma 7 . . . . .	19
B.4	Proof of Lemma 8 . . . . .	20
B.5	The Exact Definition of $I_{k,h}$ and Proof of Lemma 9 . . . . .	20
B.6	Proof of Proposition 1 . . . . .	21
B.7	Proof of Theorem 2 . . . . .	23
B.8	Miscellaneous lemmas . . . . .	24
<b>C</b>	<b>Comparison with He et al. [12]</b>	<b>24</b>

## A Proofs for VOFUL2

### A.1 Proof of Lemma 4

*Proof.* Let  $W_t = V_0 + \sum_{s \in J, s \leq t} x_s x_s^\top$ . Then,

$$\begin{aligned}
 \left( \frac{d\tau + X^2|J|}{d} \right)^d &\geq \left( \frac{\text{tr}(W_t)}{d} \right)^d \\
 &\geq |W_t| && \text{(AM-GM ineq.)} \\
 &= |V_0| \prod_{s \in J} (1 + \|x_s\|_{W_{s-1}}^2) && \text{(rank-1 update equality for det.)} \\
 &\geq |V_0| \prod_{s \in J} (1 + \|x_s\|_{V_{s-1}}^2) && (W_{s-1} \leq V_{s-1}) \\
 &\geq \tau^d 2^{|J|} \\
 \implies |J| &\leq \frac{d}{\ln(2)} \ln \left( 1 + \frac{X^2|J|}{d\tau} \right)
 \end{aligned}$$

Let us generalize it so that we compute the number of times  $\|x_s\|_{V_{t-1}}^2 \geq q$  is true rather than  $\|x_s\|_{V_{t-1}}^2 \geq 1$  in which case we have

$$|J| \leq \frac{d}{\ln(1+q)} \ln \left( 1 + \frac{X^2|J|}{d\tau} \right) =: A \ln(1 + B|J|) \quad (9)$$

We want to solve it for  $|J|$ . We observe the following:

$$|J| \leq A \ln(1 + B|J|) = A \left( \ln \left( \frac{|J|}{2A} \right) + \ln \left( 2A \left( \frac{1}{|J|} + B \right) \right) \right) \quad (10)$$

$$\leq \frac{|J|}{2} + A \ln \left( \frac{2A}{e} \left( \frac{1}{|J|} + B \right) \right) \quad (11)$$

$$\implies |J| \leq 2A \ln \left( \frac{2A}{e} \left( \frac{1}{|J|} + B \right) \right) = \frac{2}{\ln(1+q)} d \ln \left( \frac{2d}{e \ln(1+q)} \left( \frac{1}{|J|} + \frac{X^2}{d\tau} \right) \right) \quad (12)$$

We fix  $c > 0$  and consider two cases:

- Case 1:  $|J| < cd$

In this case, from (9), we have  $|J| \leq \frac{d}{\ln(1+q)} \ln \left( 1 + \frac{cX^2}{\tau} \right)$

- Case 2:  $|J| \geq cd$

In this case, from (12) we have  $|J| \leq \frac{2}{\ln(1+q)} d \ln \left( \frac{2}{e \ln(1+q)} \left( \frac{1}{c} + \frac{X^2}{\tau} \right) \right)$

We set  $c = \frac{2}{e \ln(1+q)}$  to obtain  $|J| \leq \frac{2}{\ln(1+q)} d \ln \left( 1 + \frac{2/e}{\ln(1+q)} \frac{X^2}{\tau} \right)$ . We remark that one can make the constant in front of the log to be  $\frac{d}{\ln(1+q)}$  by plugging this bound into the RHS of (9).  $\square$

## A.2 Proof of Lemma 1

*Proof.* Let  $\epsilon_s := \epsilon_s(\theta^*) = r_s - x^\top \theta^*$ . It suffices to show that the following is true w.p. at least  $1 - \delta$ ,

$$\forall \ell \in [L], k \in [K], \mu \in \mathcal{B}_2^d(2), \left| \sum_{s=1}^k \overline{(x_s^\top \mu)}_\ell \epsilon_s \right| \leq \sqrt{\sum_{s=1}^k \overline{(x_s^\top \mu)}_\ell^2 \epsilon_s^2} + 2^{-\ell}.$$

To show this, we define  $\hat{\mathcal{B}}_\ell$  to be a  $\xi_\ell$ -net over  $\mathbb{B}_2^d(2)$ . with cardinality at most  $\left(\frac{12}{\xi_\ell}\right)^d$ . Such a net exists due to Lemma 4.1 [19]. Let us assume the following event, which happens with probability at least  $1 - 6K \log_2(K) \sum_{\ell=1}^L |\hat{\mathcal{B}}_\ell|$  by Lemma 12:

$$\left| \sum_{s=1}^k \overline{(x_s^\top \mu')}_\ell \epsilon_s \right| \leq 8 \sqrt{\sum_{s=1}^k \overline{(x_s^\top \mu')}_\ell^2 \epsilon_s^2 \ln(1/\delta)} + 16 \cdot 2^{-\ell} \ln(1/\delta). \quad (\mathcal{E})$$

Let us fix  $\ell \in [L]$ ,  $k \in [K]$ , and  $\mu \in \mathbb{B}_2^d(2)$ . Choose  $\mu' \in \hat{\mathcal{B}}_\ell$  such that  $\|\mu - \mu'\|_2 \leq \xi_\ell$ . Then,

$$\begin{aligned} \left| \sum_{s=1}^k \overline{(x_s^\top \mu)}_\ell \epsilon_s \right| &\leq \left| \sum_{s=1}^k (\overline{(x_s^\top \mu)}_\ell - \overline{(x_s^\top \mu')}_\ell) \epsilon_s \right| + \left| \sum_{s=1}^k \overline{(x_s^\top \mu')}_\ell \epsilon_s \right| \\ &\leq \sum_{s=1}^k |\overline{(x_s^\top \mu)}_\ell - \overline{(x_s^\top \mu')}_\ell| + \left| \sum_{s=1}^k \overline{(x_s^\top \mu')}_\ell \epsilon_s \right| \quad (|\epsilon_s| \leq 1) \\ &\stackrel{(a)}{\leq} k \xi_\ell + \left| \sum_{s=1}^k \overline{(x_s^\top \mu')}_\ell \epsilon_s \right| \\ &\leq k \xi_\ell + 8 \sqrt{\sum_{s=1}^k \overline{(x_s^\top \mu')}_\ell^2 \epsilon_s^2 \ln(1/\delta)} + 16 \cdot 2^{-\ell} \ln(1/\delta) \quad (\text{by } (\mathcal{E})) \\ &\leq k \xi_\ell + 8 \sqrt{2 \sum_{s=1}^k \left( \overline{(x_s^\top \mu)}_\ell^2 + \xi_\ell^2 \right) \epsilon_s^2 \ln(1/\delta)} + 16 \cdot 2^{-\ell} \ln(1/\delta) \\ &\leq k \xi_\ell + \xi_\ell \cdot 8 \sqrt{2k \ln(1/\delta)} + 8 \sqrt{2 \sum_{s=1}^k \overline{(x_s^\top \mu)}_\ell^2 \epsilon_s^2 \ln(1/\delta)} + 16 \cdot 2^{-\ell} \ln(1/\delta) \\ &\leq 2^{-\ell} + 2^{-\ell} \cdot 8 \sqrt{2 \ln(1/\delta)} + 8 \sqrt{2 \sum_{s=1}^k \overline{(x_s^\top \mu)}_\ell^2 \epsilon_s^2 \ln(1/\delta)} + 16 \cdot 2^{-\ell} \ln(1/\delta) \\ &\quad (\text{choose } \xi_\ell = 2^{-\ell}/K) \end{aligned}$$

$$\leq 8\sqrt{2 \sum_{s=1}^k \overline{(x_s^\top \mu)}^2 \epsilon_s^2 \ln(1/\delta) + 32 \cdot 2^{-\ell} \ln(1/\delta)} \quad (\text{by } 1 \leq \ln(1/\delta))$$

where (a) follows from the fact that  $|x_s^\top (\mu - \mu')| \leq \epsilon$  and the observation that the clipping operation applied to two real values  $z$  and  $z'$  only makes them closer. It remains to adjust the confidence level. Note that

$$\sum_{\ell=1}^L |\hat{\mathcal{B}}_\ell| K = \sum_{\ell=1}^L (12K2^\ell)^d K \leq 2(12K)^d \cdot \frac{2^{Ld}}{2^d} \cdot K \leq (12K2^L)^{d+1}.$$

Thus,

$$6 \log_2(K) \sum_{\ell=1}^L |\hat{\mathcal{B}}_\ell| K \leq (12K2^L)^{d+2}.$$

Replacing  $\delta$  with  $\delta/(12K2^L)^{d+2}$  and setting  $\iota = 128 \ln((12K2^L)^{d+2}/\delta)$ , we conclude the proof. We remark that we did not optimize the constants in this proof.  $\square$

### A.3 Proof of Lemma 3

*Proof.* Throughout the proof, every clipping operator  $\overline{(\cdot)}$  is a shorthand of  $\overline{(\cdot)}_\ell$ . For (i), we note that

$$\begin{aligned} 2^{-\ell} \|\mu_k\|^2 + \sum_{s=1}^k \overline{(x_s^\top \mu_k)}_\ell x_s^\top \mu_k &= 2^{-\ell} \|\mu_k\|^2 + \sum_{s=1}^k \left( \left( 1 \wedge \frac{2^{-\ell}}{|x_s^\top \mu_k|} \right) x_s \right)^\top \mu_k x_s^\top \mu_k \\ &= \mu_k^\top \left( 2^{-\ell} \lambda I + \sum_{s=1}^k \left( 1 \wedge \frac{2^{-\ell}}{|x_s^\top \mu_k|} \right) x_s x_s^\top \right) \mu_k \\ &= \|\mu_k\|_{W_{\ell, k-1}}^2. \end{aligned}$$

Then,

$$\begin{aligned} &\|\mu_k\|_{(W_{\ell, k-1} - 2^{-\ell} I)}^2 \\ &= \sum_{s=1}^{k-1} \overline{(x_s^\top \mu_k)} x_s^\top \mu_k \\ &= \sum_{s=1}^{k-1} \overline{(x_s^\top \mu_k)} (x_s^\top \theta_k - r_s + r_s - x_s^\top \theta^*) \\ &= \sum_{s=1}^{k-1} \overline{(x_s^\top \mu_k)} (-\epsilon_s(\theta_k) + \epsilon_s(\theta^*)) \\ &\leq \sqrt{\sum_{s=1}^{k-1} \overline{(x_s^\top \mu_k)}^2 \epsilon_s^2(\theta_k) \iota} + 2^{-\ell} \iota + \sqrt{\sum_{s=1}^{k-1} \overline{(x_s^\top \mu_k)}^2 \epsilon_s^2(\theta^*) \iota} + 2^{-\ell} \iota \\ &\stackrel{(a)}{\leq} \sqrt{\sum_{s=1}^{k-1} \overline{(x_s^\top \mu_k)}^2 2(x_s^\top \mu_k)^2 \iota} + 2\sqrt{\sum_{s=1}^{k-1} \overline{(x_s^\top \mu_k)}^2 2\epsilon_s^2(\theta^*) \iota} + 2 \cdot 2^{-\ell} \iota \\ &\leq \sqrt{\sum_{s=1}^{k-1} \overline{(x_s^\top \mu_k)}^2 2(x_s^\top \mu_k)^2 \iota} + 2^{-\ell} \sqrt{4 \left( \sum_{s=1}^{k-1} 8\sigma_s^2 + 4 \ln\left(\frac{4K(\log_2(K) + 2)}{\delta}\right) \right) \iota} + 2 \cdot 2^{-\ell} \iota \quad (\text{By } \mathcal{E}_2) \\ &\leq \sqrt{\sum_{s=1}^{k-1} \overline{(x_s^\top \mu_k)}^2 2(x_s^\top \mu_k)^2 \iota} + 2^{-\ell} \sqrt{32 \sum_{s=1}^{k-1} \sigma_s^2 \iota + 3 \cdot 2^{-\ell} \iota} \\ &\leq \sqrt{4 \sum_{s=1}^{k-1} 2^{-\ell} \overline{(x_s^\top \mu_k)} (x_s^\top \mu_k) \iota} + 2^{-\ell} \sqrt{32 \sum_{s=1}^{k-1} \sigma_s^2 \iota + 3 \cdot 2^{-\ell} \iota} \quad (|x_s^\top \mu_k| \leq 2, \overline{(\cdot)} \leq 2^{-\ell}) \\ &= \sqrt{4 \cdot 2^{-\ell} \|\mu\|_{(W_{\ell, k-1} - 2^{-\ell} I)}^2 \iota} + 2^{-\ell} \sqrt{32 \sum_{s=1}^{k-1} \sigma_s^2 \iota + 3 \cdot 2^{-\ell} \iota} \end{aligned}$$



where (a) follows from  $\epsilon_s^2(\theta_k) = (r_s - x_s^\top \theta_k)^2 = (x_s^\top (\theta_* - \theta_k) + \epsilon_s(\theta_*))^2 \leq 2(x_s^\top \mu_k)^2 + 2\epsilon_s^2$ . We now have  $\|\mu\|_{(W_{\ell,k-1}-2^{-\ell}I)}^2$  on both sides. Using  $X \leq A + \sqrt{BX} \leq A + (B/2) + (X/2) \implies X \leq 2A + B$ , we have

$$\begin{aligned} \|\mu_k\|_{(W_{\ell,k-1}-2^{-\ell}I)}^2 &\leq 2^{-\ell} \sqrt{128 \sum_{s=1}^{k-1} \sigma_s^2 \iota + 8 \cdot 2^{-\ell} \iota} \\ \implies \|\mu_k\|_{W_{\ell,k-1}}^2 &\leq 4 \cdot 2^{-\ell} + 2^{-\ell} \sqrt{128 \sum_{s=1}^{k-1} \sigma_s^2 \iota + 8 \cdot 2^{-\ell} \iota}. \end{aligned}$$

Since  $1 \leq \ln(1/\delta)$ , we have  $4 \cdot 2^{-\ell} \leq 4 \cdot 2^{-\ell} \ln(1/\delta) \leq 2^{-\ell} \iota$ , which concludes the proof of (i).

For (ii), let  $c$  be an absolute constant that may be different every time it is used. We apply Cauchy-Schwarz inequality to obtain

$$\begin{aligned} (x_k^\top \mu_k)^2 &\leq \|x_k\|_{W_{\ell,k-1}^{-1}}^2 \|\mu_k\|_{W_{\ell,k-1}}^2 \\ &\leq \|x_k\|_{W_{\ell,k-1}^{-1}}^2 \cdot c \cdot \left( 2^{-\ell} \sqrt{\sum_{s=1}^{k-1} \sigma_s^2 \iota} + 2^{-\ell} \iota \right) \\ &\leq \|x_k\|_{W_{\ell,k-1}^{-1}}^2 \cdot c \cdot x_k^\top \mu_k \left( \sqrt{\sum_{s=1}^{k-1} \sigma_s^2 \iota} + \iota \right) \quad (2^{-\ell} \leq x_k^\top \mu_k \leq 2^{-\ell+1}) \end{aligned}$$

Dividing both sides by  $x_k^\top \mu_k$  concludes the proof.  $\square$

#### A.4 $d\sqrt{K}$ regret bound of VOFUL2

Let us slightly modify the algorithm so we now add  $\Theta_k^\ell$  with  $\ell = 0$ :

$$\Theta_k = \cap_{\ell=0}^L \Theta_k^\ell.$$

Let us call this algorithm VOFUL3. Note that this slight change will not alter the order of the regret bound of VOFUL2 reported in Theorem 1.

Let  $\lambda > 0$  be an analysis parameter to be determined later. Let  $X_k \in \mathbb{R}^{k \times d}$  be the design matrix where row  $s$  is  $x_s^\top$  and define  $y_k := (r_1, \dots, r_k)^\top$ ,  $\eta_k := (\epsilon_1, \dots, \epsilon_k)^\top$ . Let  $V_k := \lambda I + X_k^\top X_k$  and

$$\hat{\theta}_k := V_k^{-1} X_k^\top y_k \quad (13)$$

We claim that

$$\Theta_k \subseteq \left\{ \theta \in \mathbb{B}_2^d(1) : \|\hat{\theta}_k - \theta\|_{V_k}^2 \leq \beta_k \right\} =: \mathring{\Theta}_k$$

for some  $\beta_k = \tilde{O}(d + \ln(1/\delta))$ . This suffices to show that the VOFUL2 has regret bound of  $\tilde{O}(d\sqrt{K})$  since the proof technique of OFUL [1] can be immediately applied since the UCB computed based on  $\Theta_k$  is bounded above by the UCB computed with  $\mathring{\Theta}_k$ . Thus, VOFUL3 has regret bound of

$$\mathcal{R}^K = \tilde{O} \left( d\sqrt{K \ln(1/\delta)}, d^{1.5} \sqrt{\sum_{k=1}^K \sigma_k^2 \ln(1/\delta)} + d^2 \ln(1/\delta) \right).$$

To see why the claim above is true, let  $\theta \in \Theta_k$ . Then, using  $\epsilon_s^2(\theta) \leq 4$ , we have,  $\forall \mu \in \mathbb{B}_2^d(2)$ ,

$$\left| \sum_s \mu^\top x_s (y_s - x_s^\top \theta) \right| \leq \|\mu\|_{\sum_{s=1}^k x_s x_s^\top} \cdot \sqrt{4\iota} + \iota$$

Let us drop the subscript  $k$  from  $\{\hat{\theta}_k, V_k, X_k, y_k, \eta_k\}$  for brevity. The display above can be rewritten as

$$\begin{aligned} \mu^\top (X^\top y - X^\top X \theta) &= \mu^\top V (\hat{\theta} - \theta) + \lambda \mu^\top \theta \leq \|\mu\|_{X^\top X} \cdot \sqrt{4\iota} + \iota \\ \implies \mu^\top V (\hat{\theta} - \theta) &\leq \|\mu\|_V \cdot \sqrt{4\iota} + \iota + 2\lambda \quad (\|\theta\|_2 \leq 1, |\mu^\top \theta| \leq 2) \end{aligned}$$

We can choose  $\mu = \frac{1}{2}(\hat{\theta} - \theta)$  since  $\frac{1}{2}\|\hat{\theta} - \theta\|_2 \leq \frac{1}{2}(\|\hat{\theta}\|_2 + \|\theta\|_2) \leq 2$  by Lemma 11 and the choice of  $\lambda$  therein. Then,

$$\begin{aligned} \|\hat{\theta} - \theta\|_V^2 &\leq \|\hat{\theta} - \theta\|_V \sqrt{4\iota} + 2\iota + 4\lambda \\ \implies \|\hat{\theta} - \theta\|_V^2 &\leq 8\iota + 8\lambda \quad (\text{AM-GM ineq. on } \|\hat{\theta} - \theta\|_V \sqrt{4\iota}) \\ &= \tilde{O}(d + \ln(1/\delta)) \end{aligned}$$

**Lemma 11.** *Take the assumptions for the linear bandit problem in Section 2. Consider  $\hat{\theta}_k$  defined in (13) with  $\lambda = d \ln(1 + \frac{K}{d}) + 2 \ln(1/\delta)$ . Let  $K \geq (e-1)d$ . Then, with probability at least  $1 - \delta$ , we have,*

$$\forall k \leq K, \|\hat{\theta}_k\|_2 \leq 3.$$

*Proof.* Let us drop the subscript  $k$  from  $\{\hat{\theta}_k, V_k, X_k, y_k, \eta_k\}$ . Then,

$$\|\hat{\theta}\|_2 \leq \|\hat{\theta} - \theta^*\|_2 + \|\theta^*\|_2.$$

Note that  $\|\theta^*\|_2 \leq 1$ . We can further bound the first term:

$$\begin{aligned} \|\hat{\theta} - \theta^*\|_2 &= \|V^{-1}(X^\top \eta - \lambda \theta^*)\|_2 \\ &\leq \|V^{-1}X^\top \eta\|_2 + \lambda \|V^{-1}\theta^*\|_2. \end{aligned}$$

Note that

$$\lambda \|V^{-1}\theta^*\|_2 = \lambda \sqrt{\theta^{*\top} V^{-2} \theta^*} \leq \lambda \sqrt{\theta^{*\top} (\frac{1}{\lambda^2} I) \theta^*} \leq \|\theta^*\|_2 \leq 1.$$

It remains to bound  $\|V^{-1}X^\top \eta\|_2$ . To see this,

$$\begin{aligned} \|V^{-1}X^\top \eta\|_2 &= \sqrt{\eta^\top X V^{-2} X^\top \eta} \\ &= \sqrt{\eta^\top X V^{-1} (\frac{1}{\lambda} \cdot \lambda I) V^{-1} X^\top \eta} \\ &\leq \sqrt{\eta^\top X V^{-1} (\frac{1}{\lambda} \cdot V) V^{-1} X^\top \eta} \\ &= \frac{1}{\sqrt{\lambda}} \|X^\top \eta\|_{V^{-1}} \\ &\leq \frac{1}{\sqrt{\lambda}} \left( \sqrt{d \ln(1 + \frac{k}{d\lambda}) + 2 \ln(1/\delta)} \right) \quad (\text{Abbasi-Yadkori et al. [1, Theorem 1]}) \\ &\stackrel{(a)}{\leq} 1 \end{aligned}$$

where (a) is by choosing  $\lambda = d \ln(1 + \frac{K}{d}) + 2 \ln(1/\delta)$  and then using  $\lambda \geq 1$  (due to  $K \geq (e-1)d$ ).  $\square$

## A.5 Miscellaneous Lemmas

For completeness, we state the lemmas borrowed from prior work.

**Lemma 12.** (Zhang et al. [31, Lemma 9]) *Let  $\{\mathcal{F}_i\}_{i=0}^n$  be a filtration. Let  $\{X_i\}_{i=1}^n$  be a sequence of real-valued random variables such that  $X_i$  is  $\mathcal{F}_i$ -measurable. We assume that  $\mathbb{E}[X_i | \mathcal{F}_{i-1}] = 0$  and that  $|X_i| \leq b$  almost surely. For  $\delta < e^{-1}$ , we have*

$$\mathbb{P} \left( \left| \sum_{i=1}^n X_i \right| \leq 8 \sqrt{\sum_{i=1}^n X_i^2 \ln(1/\delta)} + 16b \ln(1/\delta) \right) \geq 1 - 6\delta \log_2(n)$$

**Lemma 13.** (Zhang et al. [31, Lemma 10]) *Let  $\{\mathcal{F}_i\}_{i=0}^n$  be a filtration. Let  $\{X_i\}_{i=1}^n$  be a sequence of random variables such that  $|X_i| \leq 1$  almost surely, that  $X_i$  is  $\mathcal{F}_i$ -measurable. For every  $\delta \in (0, 1)$ ,*

we have

$$\mathbb{P} \left[ \sum_{i=1}^n X_i^2 \geq \sum_{i=1}^n 8\mathbb{E}[X_i^2 | \mathcal{F}_{i-1}] + 4 \ln \frac{4}{\delta} \right] \leq \left( \lceil \log_2 n \rceil + 1 \right) \delta$$

## B Proofs for VARLin2

Throughout the proof, we use  $c$  as absolute constant that can be different every single time it is used.

### B.1 Proof of Lemma 6

*Proof.* Assume that  $\theta^* \in \Theta_k$  for all  $k \in [K]$ . Since  $\theta^* \in \Theta_k$ ,

$$\begin{aligned} Q_h^k(s, a) &= \min\{r(s, a) + \max_{\theta \in \Theta_k} \sum_{i=1}^d \theta_i P_{s,a}^i V_{h+1}^k\} \\ &\geq \min\{1, r(s, a) + \sum_{i=1}^d \theta_i^* P_{s,a}^i V_{h+1}^k\} \\ &\geq \min\{1, r(s, a) + \sum_{i=1}^d \theta_i^* P_{s,a}^i V_{h+1}^*\} \\ &= Q_h^*(s, a), \end{aligned}$$

so the statement follows.  $\square$

### B.2 Proof of Lemma 5

*Proof.* Similar to the linear bandit case, let  $\hat{\mathcal{B}}_\ell$  be a  $\xi_\ell$ -net over  $\mathbb{B}_1^d(2)$  with cardinality at most  $(\frac{12}{\xi_\ell})^d$  and pick  $\mu \in \mathbb{B}_1^d(2)$  and  $\mu' \in \hat{\mathcal{B}}_\ell$  such that the distance between them is at most  $\xi_\ell$ . Let us define the conditional variance  $\mathbb{V}_{v,u}[\epsilon_{v,u}^m] := \mathbb{V}[\epsilon_{v,u}^m | \mathcal{F}_u^v]$  where  $\mathcal{F}_u^v$  denotes history up to (and including) episode  $v$  and time horizon  $u$ .

Noticing

$$\begin{aligned} \mathbb{V}_{v,u}[\epsilon_{v,u}^m] &= \mathbb{E}_{v,u}[(\epsilon_{v,u}^m)^2] \\ &= ((\theta^*)^\top x_{v,u}^m)^2 - 2((\theta^*)^\top x_{v,u}^m) \cdot \mathbb{E}_{v,u}[(V_{u+1}^v(s_{u+1}^v))^{2^m}] + \mathbb{E}_{v,u}[(V_{u+1}^v(s_{u+1}^v))^{2^m}]^2, \\ \mathbb{E}_{v,u}[(V_{u+1}^v(s_{u+1}^v))^{2^m}] &= (\theta^*)^\top x_{v,u}^m, \text{ and } \mathbb{E}_{v,u}[(V_{u+1}^v(s_{u+1}^v))^{2^m}]^2 = (\theta^*)^\top x_{v,u}^{m+1}, \text{ we have} \\ \mathbb{V}_{v,u}[\epsilon_{v,u}^m] &= (\theta^*)^\top x_{v,u}^{m+1} - ((\theta^*)^\top x_{v,u}^m)^2. \end{aligned}$$

We apply Lemma 14 with  $\epsilon = 1$ ,  $b = 2^{-\ell}$  to obtain

$$\left| \sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \overline{((x_{v,u}^m)^\top \mu')}_\ell \epsilon_{v,u}^m \right| \leq 4 \sqrt{\sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \overline{((x_{v,u}^m)^\top \mu')}_\ell^2 \mathbb{V}[\epsilon_{v,u}^m | \mathcal{F}_u^v] + 4 \cdot 2^{-\ell} \ln(1/\delta)}$$

with probability at least  $1 - \delta(1 + \log_2(HK))$  and repeat the similar procedure by taking the union bound. We drop  $\ell$  from the clipping notation for the sake of brevity.

$$\begin{aligned} & \left| \sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \overline{((x_{v,u}^m)^\top \mu)} \epsilon_{v,u}^m \right| \\ &= \left| \sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \left( \overline{((x_{v,u}^m)^\top \mu)} - \overline{((x_{v,u}^m)^\top \mu')} \right) \epsilon_{v,u}^m \right| + \left| \sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \overline{((x_{v,u}^m)^\top \mu')} \epsilon_{v,u}^m \right| \\ &\leq \sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \left| \overline{((x_{v,u}^m)^\top \mu)} - \overline{((x_{v,u}^m)^\top \mu')} \right| + \left| \sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \overline{((x_{v,u}^m)^\top \mu')} \epsilon_{v,u}^m \right| \quad (|\epsilon_s| \leq 1) \end{aligned}$$

$$\begin{aligned}
&\leq HK\xi_\ell + 4 \sqrt{\sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \left( (x_{v,u}^m)^\top \mu' \right)^2 \mathbb{V}[\epsilon_{v,u}^m \mid \mathcal{F}_u^v] \ln(1/\delta) + 4 \cdot 2^{-\ell} \ln(1/\delta)} \\
&\leq HK\xi_\ell + 4 \sqrt{2 \sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \left\{ \left( (x_{v,u}^m)^\top \mu \right)^2 + \xi_\ell^2 \right\} \mathbb{V}[\epsilon_{v,u}^m \mid \mathcal{F}_u^v] \ln(1/\delta) + 4 \cdot 2^{-\ell} \ln(1/\delta)} \\
&\leq HK\xi_\ell + 4\xi_\ell \sqrt{2HK \ln(1/\delta)} + 4 \sqrt{2 \sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \left( (x_{v,u}^m)^\top \mu \right)^2 \mathbb{V}[\epsilon_{v,u}^m \mid \mathcal{F}_u^v] \ln(1/\delta) + 4 \cdot 2^{-\ell} \ln(1/\delta)} \\
&\leq (4\sqrt{2} + 5) \cdot 2^{-\ell} \ln(1/\delta) + 4 \sqrt{2 \sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \left( (x_{v,u}^m)^\top \mu \right)^2 \mathbb{V}[\epsilon_{v,u}^m \mid \mathcal{F}_u^v] \ln(1/\delta)} \\
&\hspace{20em} (\text{choose } \xi_\ell = 2^{-\ell}/(HK)) \\
&\leq 4 \cdot 2^{-\ell} \ln(1/\delta') + 4 \sqrt{\sum_{(v,u) \in \mathcal{T}_{k,H}^{m,i}} \left( (x_{v,u}^m)^\top \mu \right)^2 \eta_{v,u}^m \ln(1/\delta')} \hspace{2em} (\text{setting } \delta = \delta'^{1/3})
\end{aligned}$$

We then take union bounds over  $m \in \{1, 2, \dots, L\}$ ,  $i, \ell \in [L]$ ,  $k \in [K]$ , and  $\mu' \in \hat{\mathcal{B}}_\ell$ , which invoke applying Lemma 14  $(2HK)^{2(d+2)}$  times. It follows from

$$\sum_{i,\ell,k} |\hat{\mathcal{B}}_\ell| = L_0 L K \sum_{\ell} (HK 2^\ell)^d \leq (HK)^2 (HK)^d \frac{2^{Ld}}{2^d} \leq (HK 2^L)^{d+2} \leq (2HK)^{2(d+2)}.$$

Hence, the display above holds with probability at least  $1 - \delta(1 + \log_2(HK))(2HK)^{2(d+2)} \geq 1 - \delta(2HK)^{2(d+3)}$ . Replacing  $\delta$  with  $1/(2HK)^{2(d+3)} \cdot \delta$  and setting  $\iota = 3 \cdot \ln((2HK)^{2(d+3)}/\delta)$  the result follows.  $\square$

### B.3 Proof of Lemma 7

*Proof.* Lemma 12 of Abbasi-Yadkori et al. [1] shows the following in its proof: Let  $A$ ,  $B$ , and  $C$  be positive semi-definite (PSD) matrices such that  $A = B + C$ . Then we have that

$$\sup_{x \neq 0} \frac{\|x\|_A^2}{\|x\|_B^2} \leq \frac{\det(A)}{\det(B)}.$$

Note that  $V_{\text{anc}(t)}^{-1} = V_{t-1}^{-1} + C$  for some PSD matrix  $C$  (to see this, apply a series of rank-one update formula for the covariance matrix). Thus, we have that if  $t \in J$ , then

$$r < \frac{\|x_t\|_{V_{\text{anc}(t)}^{-1}}^2}{\|x_t\|_{V_{t-1}^{-1}}^2} \leq \frac{|V_{\text{anc}(t)}^{-1}|}{|V_{t-1}^{-1}|} = \frac{|V_{t-1}|}{|V_{\text{anc}(t)}|}.$$

Let  $J[t] = J \cap [t]$  with  $[0] := \emptyset$ . Let  $\text{prev}(t)$  be the time step in  $J$  immediately prior to  $t$ :  $\text{prev}(t) := \max\{s \in \{0\} \cup J : s < t\}$ . Define  $W_t = V_0 + \sum_{s \in J[t]} x_s x_s^\top$ .

$$\begin{aligned}
\left( \frac{d\lambda + X^2|J|}{d} \right)^d &\geq \left( \frac{\text{tr}(W_T)}{d} \right)^d \\
&\geq |W_T| && (\text{AM-GM ineq.}) \\
&= |V_0| \prod_{t \in J} \frac{|W_t|}{|W_{\text{prev}(t)}|} \\
&\geq |V_0| \prod_{t \in J} \frac{|W_t|}{|W_{\text{anc}(t)}|} && (\text{prev}(t) \leq \text{anc}(t)) \\
&= |V_0| r^{|J|} \\
&\geq \lambda^d r^{|J|}
\end{aligned}$$

$$\implies |J| \leq \frac{d}{\ln(r)} \ln \left( 1 + \frac{X^2 |J|}{d\lambda} \right)$$

Then,

$$|J| \leq \frac{d}{\ln(r)} \ln \left( 1 + \frac{X^2 |J|}{d\lambda} \right) =: A \ln(1 + B|J|) \quad (14)$$

We want to solve it for  $|J|$ . Do the following:

$$|J| \leq A \ln(1 + B|J|) = A \left( \ln \left( \frac{|J|}{2A} \right) + \ln \left( 2A \left( \frac{1}{|J|} + B \right) \right) \right) \quad (15)$$

$$\leq \frac{|J|}{2} + A \ln \left( \frac{2A}{e} \left( \frac{1}{|J|} + B \right) \right) \quad (16)$$

$$\implies |J| \leq 2A \ln \left( \frac{2A}{e} \left( \frac{1}{|J|} + B \right) \right) = \frac{2}{\ln(r)} d \ln \left( \frac{2d}{e \ln(r)} \left( \frac{1}{|J|} + \frac{X^2}{d\lambda} \right) \right) \quad (17)$$

We fix  $c > 0$  and consider two cases:

- Case 1:  $|J| < cd$

In this case, from (14), we have  $|J| \leq \frac{d}{\ln(r)} \ln \left( 1 + \frac{cL^2}{\lambda} \right)$

- Case 2:  $|J| \geq cd$

In this case, from (17) we have  $|J| \leq \frac{2}{\ln(r)} d \ln \left( \frac{2}{e \ln(r)} \left( \frac{1}{c} + \frac{X^2}{\lambda} \right) \right)$

We set  $c = \frac{2}{e \ln(r)}$  to obtain  $|J| \leq \frac{2}{\ln(r)} d \ln \left( 1 + \frac{2/e}{\ln(r)} \frac{X^2}{\lambda} \right)$ .  $\square$

#### B.4 Proof of Lemma 8

Recall that  $\theta_t^m = \arg \max_{\theta \in \Theta_{k-1}} |\{\theta x_t^{m+1} - (\theta x_t^m)^2\}|$  and  $\eta_b^m(\theta) = \theta x_b^{m+1} - (\theta x_b^m)^2$ . We have

$$\begin{aligned} \|\mu_t^m\|_{(W_{\text{anc}(t)}^{m,i,\ell} - 2^{-\ell} I)}^2 &= \sum_{b \in \mathcal{T}_{\text{anc}(t)}^{m,i}} \overline{((x_b^m)^\top \mu_t^m)} (x_b^m)^\top \mu_t^m \\ &= \sum_{b \in \mathcal{T}_{\text{anc}(t)}^{m,i}} \overline{((x_b^m)^\top \mu_t^m)} (-\epsilon_b^m(\theta_t^m) + \epsilon_b^m(\theta^*)) \\ &\leq \sqrt{\sum_{b \in \mathcal{T}_{\text{anc}(t)}^{m,i}} \overline{((x_b^m)^\top \mu_t^m)}^2} \eta_b^m(\theta_t^m)_\ell + 4 \cdot 2^{-\ell} \ell + 4 \sqrt{\sum_{b \in \mathcal{T}_{\text{anc}(t)}^{m,i}} \overline{((x_b^m)^\top \mu_t^m)}^2} \eta_b^m(\theta^*)_\ell + 4 \cdot 2^{-\ell} \ell \\ &\leq C \sqrt{\sum_{b \in \mathcal{T}_{\text{anc}(t)}^{m,i}} \overline{((x_b^m)^\top \mu_t^m)}^2} \cdot 2^{-i} \ell + 8 \cdot 2^{-\ell} \ell \\ &\leq C \sqrt{\sum_{b \in \mathcal{T}_{\text{anc}(t)}^{m,i}} \overline{((x_b^m)^\top \mu_t^m)} (x_b^m)^\top \mu_t^m} \cdot 2^{-i} \ell + 8 \cdot 2^{-\ell} \ell \\ &\leq C \|\mu_t^m\|_{W_{\text{anc}(t)}^{m,i,\ell}} \sqrt{2^{-i} \ell} + 8 \cdot 2^{-\ell} \ell \end{aligned}$$

Solving for  $\|\mu_t^m\|_{W_{\text{anc}(t)}^{m,i,\ell}}$ , we get the desired bound.

#### B.5 The Exact Definition of $I_{k,h}$ and Proof of Lemma 9

Fix  $m$  and  $i$ . Let  $t \in \mathcal{T}_T^{m,i}$  and  $\ell = \ell_t^m$ . We recall Lemma 8 that yields

$$\|\mu_t^m\|_{W_{\text{anc}(t)}^{m,i,\ell}} \leq c_M \cdot \max\{\sqrt{2^{-i} \ell}, \sqrt{2^{-\ell} \ell}\}$$

Let us define  $\mathcal{T}^{m,i,\ell} = \{t \in \mathcal{T}_T^{m,i} : (x_t^m)^\top \mu_t^m \in (2 \cdot 2^{-\ell}, 2 \cdot 2^{1-\ell}]\}$  and split the time steps  $\mathcal{T}^{m,i,\ell}$  as follows:

$$\mathcal{T}^{m,i,\ell,\langle 1 \rangle} := \left\{ t \in \mathcal{T}^{m,i,\ell} : \|\mu_t^m\|_{W_{\text{anc}(t)}^{m,i,\ell}} \leq c_M \sqrt{2^{-i} \ell} \right\} \quad \text{and} \quad \mathcal{T}^{m,i,\ell,\langle 2 \rangle} := \mathcal{T}^{m,i,\ell} \setminus \mathcal{T}^{m,i,\ell,\langle 1 \rangle}.$$

Given  $t \in [T]$ , let  $n_t^{m,i,\ell,\langle 1 \rangle}$  be  $n \in \{0, \dots, \ell\}$  such that

$$\max_{t': t \leq t' \in \mathcal{T}^{m,i,\ell,\langle 1 \rangle}} |(x_t^m)^\top \mu_{t'}^m| \in \left( \mathbb{1}\{n \neq 0\} \cdot 2^{-\ell+n}, 2^{-\ell+n+1} \right]$$

where we set  $n = 0$  if the maximum above is less or equal to  $2^{-\ell}$ .

Let  $\mathcal{T}^{m,i,\ell,\langle 1 \rangle,n} = \mathcal{T}^{m,i,\ell,\langle 1 \rangle} \cap \{t : n_t^{m,i,\ell,\langle 1 \rangle} = n\}$  and set  $V_t^{m,i,\ell,\langle 1 \rangle,n} := 2^{-\ell} I + \sum_{b \in \mathcal{T}_t^{m,i,\ell,\langle 1 \rangle,n}} x_b^m (x_b^m)^\top$  with  $\mathcal{T}_t^{m,i,\ell,\langle z \rangle,n} := \{a' \leq a : a' \in \mathcal{T}^{m,i,\ell,\langle z \rangle,n}\}$ .

Define

$$I_t^{m,i,\ell,\langle 1 \rangle,n} := \mathbb{1} \left\{ \forall t' \in \mathcal{T}_t^{m,i,\ell,\langle 1 \rangle,n} \text{ with } k(t') = k(t), \quad \|x_{t'}\|_{(V_{\text{anc}(t')}^{m,i,\ell,\langle 1 \rangle,n})^{-1}}^2 \leq r \|x_{t'}\|_{(V_{t'-1}^{m,i,\ell,\langle 1 \rangle,n})^{-1}}^2 \right\}$$

for some  $r > 1$  to be specified later. With this definition, one can see that  $\sum_{t \in [T]: \text{h}(t) < H} I_t^{m,i,\ell,\langle 1 \rangle,n} - I_{t+1}^{m,i,\ell,\langle 1 \rangle,n}$  is the number of bad episodes where there exists  $h \in [H]$  such that

$$\|x_t\|_{(V_{\text{anc}(t)}^{m,i,\ell,\langle 1 \rangle,n})^{-1}}^2 > r \|x_t\|_{(V_{t-1}^{m,i,\ell,\langle 1 \rangle,n})^{-1}}^2.$$

Define  $I_{k,h}^{m,i,\ell,\langle 1 \rangle,n} := I_{t(k,h)}^{m,i,\ell,\langle 1 \rangle,n}$ . Define similar quantities for  $\langle 2 \rangle$  as well.

Define  $I_t := \prod_{m,i,\ell,n} I_t^{m,i,\ell,\langle 1 \rangle,n} \prod_{m,i,\ell,n} I_t^{m,i,\ell,\langle 2 \rangle,n}$ . We now prove Lemma 9.

Note that

$$\sum_{k=1}^K \sum_{h=1}^{H-1} I_{k,h} - I_{k,h+1} \leq \sum_{k=1}^K \sum_{h=1}^{H-1} \sum_{m \leq L_0} \sum_{i \leq L} \sum_{\ell \leq L} \sum_{z \in [2]} \sum_{n=0}^{\ell} \left( I_{k,h}^{m,i,\ell,\langle z \rangle,n} - I_{k,h+1}^{m,i,\ell,\langle z \rangle,n} \right),$$

We now assume  $z = 1$  without loss of generality. The display above can be written as

$$\sum_{t \in [T]: \text{h}(t) < H} (I_t - I_{t+1}) \leq \sum_{m,i,\ell,n} \sum_{t \in [T]: \text{h}(t) < H} I_t^{m,i,\ell,\langle 1 \rangle,n} - I_{t+1}^{m,i,\ell,\langle 1 \rangle,n}.$$

Recall that the inner sum above is the count of the ‘bad’ episodes. We invoke Lemma 7  $2L_0(L)^3$  times with  $X = \sqrt{d}$  as  $\|x_t\|_1 \leq \sqrt{d}$  to finish the proof as

$$\sum_{k,h} I_{k,h} - I_{k,h+1} \leq O \left( \frac{d}{\ln(r)} \log^5 \left( dHK \left( 1 + \frac{d^2}{\ln(r)} \right) \right) \right)$$

## B.6 Proof of Proposition 1

To proceed we safely choose  $r = 2$  and inherit all notations from Section B.5. Let us define

$$\check{\mathcal{T}}^{m,i,\ell,\langle z \rangle} := \mathcal{T}^{m,i,\ell,\langle z \rangle} \cap \{t \in [T] : I_t = 1\}$$

where  $z \in \{1, 2\}$ . We start with  $R_m$  as

$$\begin{aligned} R_m &= \sum_t (\check{x}_t^m)^\top \mu_t^m \\ &\leq \sum_{t,i,\ell} \sum_{t \in \check{\mathcal{T}}^{m,i,\ell,\langle 1 \rangle}} (\check{x}_t^m)^\top \mu_t^m + \sum_{t,i,\ell} \sum_{t \in \check{\mathcal{T}}^{m,i,\ell,\langle 2 \rangle}} (\check{x}_t^m)^\top \mu_t^m. \end{aligned}$$

Let us fix  $m, i$ , and  $\ell$  and focus on controlling  $\sum_{t \in \check{\mathcal{T}}^{m,i,\ell,\langle z \rangle}} (\check{x}_t^m)^\top \mu_t^m$  for  $z \in [2]$ . Hereafter, we omit the superscripts of  $(m, i, \ell)$  to avoid clutter, unless there is a need.

Note that for  $t \in \check{\mathcal{T}}^{\langle 1 \rangle,n}$  and  $b$  such that  $t \leq b$ ,

$$W_{\text{anc}(t)}(\mu_b) = 2^{-\ell} I + \sum_{t' \in \mathcal{T}_{\text{anc}(t)}^{m,i,\ell}} (1 \wedge \frac{2^{-\ell}}{|x_{t'} \mu_b|}) x_{t'} x_{t'}^\top$$

$$\begin{aligned}
&\geq 2^{-\ell} I + \sum_{t' \in \mathcal{T}_{\text{anc}(t)}^{m, i, \ell, \langle 1 \rangle, n}} (1 \wedge \frac{2^{-\ell}}{|x_{t'} \mu_b|}) x_{t'} x_{t'}^\top \\
&\geq 2^{-\ell} I + \sum_{t' \in \mathcal{T}_{\text{anc}(t)}^{m, i, \ell, \langle 1 \rangle, n}} (1 \wedge \frac{2^{-\ell}}{2^{-\ell+n+1}}) x_{t'} x_{t'}^\top \\
&\geq 2^{-\ell} I + 2^{-n-1} \sum_{t' \in \mathcal{T}_{\text{anc}(t)}^{m, i, \ell, \langle 1 \rangle, n}} x_{t'} x_{t'}^\top \\
&\geq c 2^{-n} V_{\text{anc}(t)}^{\langle 1 \rangle, n}.
\end{aligned}$$

For  $t \in \check{\mathcal{T}}^{\langle 1 \rangle, n}$ , let  $b = \arg \max_{t \leq b' \in \check{\mathcal{T}}^{\langle 1 \rangle}} |x_t^\top \mu_{b'}|$ . Then,

$$\begin{aligned}
2^{-\ell+n} &\leq |x_t^\top \mu_b| && \text{(def'n of } b) \\
&\leq \|x_t\|_{W_{\text{anc}(t)}^{-1}(\mu_b)} \|\mu_b\|_{W_{\text{anc}(t)}(\mu_b)} \\
&\leq \sqrt{2^n} \|x_t\|_{(V_{\text{anc}(t)}^{\langle 1 \rangle, n})^{-1}} \|\mu_b\|_{W_{\text{anc}(b)}(\mu_b)} \\
&\leq c \sqrt{r} \sqrt{2^n} \|x_t\|_{(V_{t-1}^{\langle 1 \rangle, n})^{-1}} \sqrt{2^{-i_\ell}}. && \text{(by } b \in \check{\mathcal{T}}^{\langle 1 \rangle})
\end{aligned}$$

This implies that

$$\|x_t\|_{(V_{t-1}^{\langle 1 \rangle, n})^{-1}}^2 \geq c_{\langle 1 \rangle} \frac{2^{-2\ell+n}}{r 2^{-i_\ell}}.$$

for some absolute constant  $c_{\langle 1 \rangle} > 0$ .

Thus,

$$\begin{aligned}
\sum_{t \in \check{\mathcal{T}}^{\langle 1 \rangle}} x_t \mu_t &\leq c 2^{-\ell} \sum_{t \in \check{\mathcal{T}}^{\langle 1 \rangle}} 1 \\
&\leq c 2^{-\ell} \sqrt{|\check{\mathcal{T}}^{\langle 1 \rangle}| \sum_{t \in \check{\mathcal{T}}^{\langle 1 \rangle}} 1} \\
&\leq c 2^{-\ell} \sqrt{|\check{\mathcal{T}}^{\langle 1 \rangle}| \sum_{n=0}^{\ell} \sum_{t \in \check{\mathcal{T}}^{\langle 1 \rangle, n}} 1} \\
&\leq c 2^{-\ell} \sqrt{|\check{\mathcal{T}}^{\langle 1 \rangle}| \sum_{n=0}^{\ell} \sum_{t \in \check{\mathcal{T}}^{\langle 1 \rangle, n}} \mathbb{I} \left\{ \|x_t\|_{(V_{t-1}^{\langle 1 \rangle, n})^{-1}}^2 \geq c_{\langle 1 \rangle} \frac{2^{-2\ell+n}}{r 2^{-i_\ell}} \right\}} \\
&\leq c 2^{-\ell} \sqrt{|\check{\mathcal{T}}^{\langle 1 \rangle}| \sum_{n=0}^{\ell} \sum_{t \in \mathcal{T}^{\langle 1 \rangle, n}} \mathbb{I} \left\{ \|x_t\|_{(V_{t-1}^{\langle 1 \rangle, n})^{-1}}^2 \geq c_{\langle 1 \rangle} \frac{2^{-2\ell+n}}{r 2^{-i_\ell}} \right\}} && (\check{\mathcal{T}}^{\langle 1 \rangle, n} \subseteq \mathcal{T}^{\langle 1 \rangle, n}) \\
&\leq c 2^{-\ell} \sqrt{|\check{\mathcal{T}}^{\langle 1 \rangle}| \sum_{n=0}^{\ell} \frac{r 2^{-i_\ell}}{2^{-2\ell+n}} d \ln \left( 1 + c \frac{r 2^{-i_\ell}}{2^{-2\ell+n} 2^{-\ell}} \right)} && \text{(Lemma 4)} \\
&\leq c \sqrt{dr |\check{\mathcal{T}}^{\langle 1 \rangle}| 2^{-i_\ell} \ln(1 + cr \iota 8^\ell)} \\
&\leq c \sqrt{dr (1 + \sum_{t \in \check{\mathcal{T}}^{\langle 1 \rangle}} \eta_t) \iota \ln(1 + cr \iota 8^\ell)}
\end{aligned}$$

For the other case involving  $\check{\mathcal{T}}^{\langle 2 \rangle}$ , we use the same logic as above. Let  $t \in \check{\mathcal{T}}^{\langle 1 \rangle, n}$ . Then, for any  $b$  such that  $a \leq b$ , one can show that

$$W_{\text{anc}(t)}(\mu_b) \geq c 2^{-n} V_{\text{anc}(t)}^{\langle 2 \rangle, n}.$$

Then, once again letting  $b = \arg \max_{t < b \in \check{\mathcal{T}}^{\langle 2 \rangle}} |x_t^\top \mu_b|$  for  $t \in \check{\mathcal{T}}^{\langle 2 \rangle, n}$  one can show that

$$2^{-\ell+n} \leq c \sqrt{r} \sqrt{2^n} \|x_t\|_{(V_{t-1}^{\langle 2 \rangle, n})^{-1}} \sqrt{2^{-\ell_\ell}}.$$

This implies that

$$\|x_t\|_{(V_{t-1}^{(2),n})^{-1}}^2 \geq c \frac{2^{-\ell+n}}{r\iota}$$

for some absolute constant  $c > 0$ . Then,

$$\begin{aligned} \sum_{t \in \mathcal{T}_{(2)}} x_t^\top \mu_t &\leq c 2^{-\ell} \sum_{n=0}^{\ell} \sum_{(k,h) \in \mathcal{T}^{(2),n}} 1 \\ &\leq c 2^{-\ell} \sum_{n=0}^{\ell} \sum_{t \in \mathcal{T}^{(2),n}} \mathbb{1} \left\{ \|x\|_{(V_{t-1}^{(2),n})^{-1}}^2 \geq c \frac{2^{n-\ell}}{r\iota} \right\} \\ &\leq c 2^{-\ell} \sum_{n=0}^{\ell} c \frac{r\iota}{2^{n-\ell}} d \ln \left( 1 + c \frac{r\iota}{2^{n-\ell} \cdot 2^{-\ell}} \right) \\ &\leq c d r \iota \ln \left( 1 + c r \iota 4^\ell \right) \end{aligned}$$

Invoking the elliptical potential count lemma together with  $r = 2$ ,

$$R_m \leq c \sum_i^L \sum_\ell^L \left( d^{0.5} \sqrt{(1 + \sum_{t \in \mathcal{T}^{m,i,\ell}} \eta_t^m) \iota \ln(1 + c \iota 8^\ell)} + d \iota \ln(1 + c \iota 8^\ell) \right),$$

which implies that

$$R_m \leq O \left( d^{0.5} \log^{2.5}(HK) \sqrt{\sum_t \check{\eta}_t^m \iota \log(\iota)} + d \log^3(HK) \iota \log(\iota) \right).$$

where  $\check{\eta}_t^m := \eta_t^m I_t$ .

## B.7 Proof of Theorem 2

*Proof.* We continue from the proof in the main paper where it remains to bound  $R_0 + M_0$ . Using the relation (equation (56) and (57) in [31]),

$$\sum_t \check{\eta}_t^m \leq |M_{m+1}| + O(d \log^5(dHK) + 2^{m+1}(K + R_0)) + R_{m+1} + 2R_m,$$

one has, using Proposition 1,

$$R_m \leq O(d^{0.5} \log^{2.5}(dHK) \sqrt{\iota \log(\iota)} \sqrt{|M_{m+1}| + 2^{m+1}(K + R_0) + R_{m+1} + 2R_m} + d \log^5(dHK) \iota \log(\iota))$$

The strategy is to solve the recursive inequalities with respect to  $R_m$  and  $M_m$  to obtain a bound on  $R_0$  and  $M_0$ . By Lemma 15, we have

$$|M_m| \leq O(\sqrt{|M_{m+1}| + 2^{m+1}(K + R_0) \log(1/\delta)} + d^{0.5} \log^{2.5}(dHK) + \log(1/\delta)). \quad (18)$$

With  $b_m := R_m + |M_m|$ , we have

$$b_m \leq \hat{O}(d \log^3(dHK) \sqrt{\log(1/\delta)} \sqrt{b_m + b_{m+1} + 2^{m+1}(K + R_0)} + d^2 \log^7(dHK) \log(1/\delta)).$$

where  $\hat{O}$  ignores doubly logarithmic factors.

We now use Lemma 17 with  $\lambda_1 = HK$ ,  $\lambda_2 = \hat{\Theta}(d \log^3(dHK) \sqrt{\log(1/\delta)})$ ,  $\lambda_3 = (K + R_0)$  and  $\lambda_4 = \hat{\Theta}(d^2 \log^7(dHK) \ln(1/\delta))$  where  $\hat{\Theta}$  ignores doubly logarithmic factors, we obtain

$$R_0 \leq b_0 \leq \tilde{O}(d^2 \log^7(dHK) \log(1/\delta) + \sqrt{(K + R_0) d^2 \log^6(dHK) \log(1/\delta)}).$$

We can solve it for  $R_0$  to obtain  $R_0 \leq \tilde{O}(\sqrt{K d^2 \log^7(dHK) \log(1/\delta)} + d^2 \log^7(dHK) \log(1/\delta))$ .

Next, we apply Lemma 18 to (18) with  $\lambda_2 = \Theta(1)$ ,  $\lambda_3 = (K + R_0) \ln(1/\delta)$ , and  $\lambda_4 = \Theta(d^{0.5} \log^{2.5}(dHK) + \ln(1/\delta))$  to obtain

$$\begin{aligned} |M_0| &\leq O(\sqrt{(K + R_0) \ln(1/\delta)} + d^{0.5} \log^{2.5}(dHK) + \ln(1/\delta)) \\ &\leq \tilde{O}(\sqrt{K \log(1/\delta)} + \sqrt{d^2 \log^7(dHK) \log^2(1/\delta)} + d^{0.5} \log^{2.5}(dHK) + \ln(1/\delta)) \end{aligned}$$



where the last inequality uses  $\sqrt{AB} \leq \frac{A+B}{2}$  to obtain the following:

$$\begin{aligned} K + R_0 &\leq \tilde{O}(K + d^2 \log^7(dHK) \log(1/\delta) + \sqrt{K d^2 \log^7(dHK) \log(1/\delta)}) \\ &\leq \tilde{O}\left(K + d^2 \log^7(dHK) \log(1/\delta) + \frac{1}{2} \cdot K + \frac{1}{2} \cdot d^2 \log^7(dHK) \log(1/\delta)\right). \end{aligned}$$

Altogether, we obtain

$$b_0 = \tilde{O}(\sqrt{K d^2 \log^7(dHK) \log^2(1/\delta)} + d^2 \log^7(dHK) \log(1/\delta))$$

This concludes the proof.  $\square$

## B.8 Miscellaneous lemmas

**Lemma 14.** (Zhang et al. [32, Lemma 11]) Let  $(M_n)_{n \geq 0}$  be a martingale such that  $M_0 = 0$  and  $|M_n - M_{n-1}| \leq b$  almost surely for  $n \geq 1$ . For each  $n \geq 0$ , let  $\mathcal{F}_n = \sigma(M_1, \dots, M_n)$ . Then for any  $n \geq 1$  and  $\epsilon, \delta > 0$ , we have

$$\mathbb{P}\left(|M_n| \geq 2\sqrt{\sum_{i=1}^n \mathbb{E}[(M_i - M_{i-1})^2 | \mathcal{F}_{i-1}] \ln(1/\delta)} + 2\sqrt{\epsilon \ln(1/\delta)} + 2b \ln(1/\delta)\right) \leq 2(\log_2(b^2 n / \epsilon) + 1)\delta$$

**Lemma 15.** (Zhang et al. [31, Lemma 25])

$$|M_m| \leq O\left(\sqrt{M_{m+1} + O(d \log^5(dHK))} + 2^{m+1}(K + R_0) \log(1/\delta) + \log(1/\delta)\right)$$

**Lemma 16.** (Zhang et al. [31, Lemma 6])  $\text{Reg}_3 \leq O(\sqrt{K \log(1/\delta)})$ .

**Lemma 17.** (Zhang et al. [31, Lemma 12]) For  $\lambda_i > 0$ ,  $i \in \{1, 2, 4\}$  and  $\lambda_3 \geq 1$ , let  $\kappa = \max\{\log_2(\lambda_1), 1\}$ . Assume that  $0 \leq a_i \leq \lambda_1$  and  $a_i \leq \lambda_2 \sqrt{a_i + a_{i+1} + 2^{i+1} \lambda_3 + \lambda_4}$  for  $i \in \{1, 2, \dots, \kappa\}$ . Then, we have

$$a_1 \leq 22\lambda_2^2 + 6\lambda_4 + 4\lambda_2 \sqrt{2\lambda_3}$$

**Lemma 18.** (Zhang et al. [30, Lemma 2]) Let  $\lambda_1, \lambda_2, \lambda_4 \geq 0$  and  $\lambda_3 \geq 1$  with  $i' = \log_2(\lambda_1)$ . We have a sequence  $\{a_i\}_i$  for  $i \in \{1, 2, \dots, i'\}$  satisfying  $a_i \leq \lambda_1$  and  $a_i \leq \lambda_2 \sqrt{a_{i+1} + 2^{i+1} \lambda_3 + \lambda_4}$ . Then,

$$a_1 \leq \max\left\{(\lambda_2 + \sqrt{\lambda_2^2 + \lambda_4})^2, \lambda_2 \sqrt{8\lambda_3 + \lambda_4}\right\}$$

## C Comparison with He et al. [12]

At a high-level, He et al. [12] apply a similar strategy to ours. One immediate difference is that we do not incur an extra dependence on  $d$  inside the logarithm, but it could be due to the fact that their lemma is applying the count on a different quantity from ours. Note that our EPC is still novel and tighter than the bound appeared in a concurrent work of Wagenmaker et al. [22, Lemma 6.2] for a large enough  $K$  (i.e., time horizon).

Note that the core of our novelty is the point of view introduced by the the matrix norm with respect to  $W_{\ell, k-1}(\mu)$ . This enables the connection to the elliptical potential lemma and the peeling technique. Such a viewpoint is exactly what the paper of VOFUL [31] did not seem to have realized. Indeed, the proof of VOFUL [31] does not use peeling on  $x_k^\top \mu_k$  as we do.

We also like to highlight that Lemma 7 is still novel and is one of the main contributors to the improved regret bound for the linear mixture MDP.