

---

# Computationally Efficient Horizon-Free Reinforcement Learning for Linear Mixture MDPs

---

**Dongruo Zhou**

Department of Computer Science  
University of California, Los Angeles  
Los Angeles, CA 90095  
drzhou@cs.ucla.edu

**Quanquan Gu**

Department of Computer Science  
University of California, Los Angeles  
Los Angeles, CA 90095  
qgu@cs.ucla.edu

## Abstract

Recent studies have shown that episodic reinforcement learning (RL) is not more difficult than contextual bandits, even with a long planning horizon and unknown state transitions. However, these results are limited to either tabular Markov decision processes (MDPs) or computationally inefficient algorithms for linear mixture MDPs. In this paper, we propose the first computationally efficient horizon-free algorithm for linear mixture MDPs, which achieves the optimal  $\tilde{O}(d\sqrt{K} + d^2)$  regret up to logarithmic factors. Our algorithm adapts a weighted least square estimator for the unknown transitional dynamic, where the weight is both *variance-aware* and *uncertainty-aware*. When applying our weighted least square estimator to heterogeneous linear bandits, we can obtain an  $\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + d)$  regret in the first  $K$  rounds, where  $d$  is the dimension of the context and  $\sigma_k^2$  is the variance of the reward in the  $k$ -th round. This also improves upon the best-known algorithms in this setting when  $\sigma_k^2$ 's are known.

## 1 Introduction

How to design efficient algorithms is a central problem for reinforcement learning (RL). Here, the *efficiency* includes both *statistical efficiency*, which requires the RL algorithm enjoy a low regret/polynomial sample complexity for finding the near-optimal policy, and *computational efficiency*, which expects the RL algorithm have polynomial running time. When restricting to episodic RL with total reward upper bounded by 1<sup>1</sup>, a longstanding question is whether episodic RL is statistically and computationally more difficult than contextual bandits (Jiang and Agarwal, 2018), since episodic RL can be seen as an extension of contextual bandits to have a long planning horizon and unknown state transition. For tabular RL, this questions has been fully resolved by a line of works (Wang et al., 2020; Zhang et al., 2021b; Li et al., 2022; Zhang et al., 2022), which propose various horizon-free algorithms. Here we say an algorithm is horizon-free if its regret/sample complexity has at most a polylogarithmic dependence on the planning horizon  $H$ . In particular, (Zhang et al., 2021b) proposed the first computationally efficient algorithm for tabular RL whose regret enjoys a polylogarithmic dependence on the planning horizon, and (Zhang et al., 2022) further removed the polylogarithmic dependence on the planning horizon.

For RL with *function approximation* to deal with large state space, Zhang et al. (2021c); Kim et al. (2021) have made some progress towards horizon-free RL for a class of MDPs called *linear mixture MDPs* (Jia et al., 2020; Ayoub et al., 2020; Zhou et al., 2021a), whose transition dynamic can be represented as a linear combination of  $d$  basic transition models. More specifically, Zhang et al.

---

<sup>1</sup>See Assumption 3.1 for a detailed description.

(2021c) proposed a VARLin algorithm for linear mixture MDPs with an  $\tilde{O}(d^{4.5}\sqrt{K} + d^9)$  regret for the first  $K$  episodes, and Kim et al. (2021) proposed a VARLin2 algorithm with an improved regret  $\tilde{O}(d\sqrt{K} + d^2)$ . However, neither algorithm is computationally efficient, because both of them need to work with nonconvex confidence sets and do not provide a polynomial-time algorithm to solve the maximization problem over these sets.

So the following question remains open:

*Can we design computationally efficient horizon-free RL algorithms when function approximation is employed?*

In this paper, we answer the above question affirmatively for linear function approximation by proposing the first computationally efficient horizon-free RL algorithm for linear mixture MDPs. Our contributions are summarized as follows.

- As a warm-up, we consider the heterogeneous linear bandits where the variances of rewards in each round are different. Such a setting can be regarded as a special case of linear mixture MDPs. We propose a computationally efficient algorithm  $\text{WeightedOFUL}^+$  and prove that in the first  $K$ -rounds, the regret of  $\text{WeightedOFUL}^+$  is  $\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + dR + d)$ , where  $\sigma_k^2$  is the variance of the reward in the  $k$ -th round and  $R$  is the magnitude of the reward noise. Our regret is *variance-aware*, i.e., it only depends on the summation of variances and does not have a  $\sqrt{K}$  term. This directly improves the  $\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + \sqrt{dK} + d)$  regret achieved by Zhou et al. (2021a).
- For linear mixture MDPs, when the total reward for each episode is upper bounded by 1, we propose a  $\text{HF-UCRL-VTR}^+$  algorithm and show that it has an  $\tilde{O}(d\sqrt{K} + d^2)$  regret for the first  $K$  episodes, where  $d$  is the number of basis transition dynamic. Our  $\text{HF-UCRL-VTR}^+$  is *computationally efficient, horizon-free and near-optimal*, as it matches the regret lower bound proved in our paper up to logarithmic factors. Our regret is strictly better than the regret attained in previous works (Zhou et al., 2021a; Zhang et al., 2021c; Kim et al., 2021).
- At the core of both  $\text{WeightedOFUL}^+$  and  $\text{HF-UCRL-VTR}^+$  is a carefully designed *weighted linear regression* estimator, whose weight is both *variance-aware* and *uncertainty-aware*, in contrast to previous weighted linear regression estimator that is only *variance-aware* (Zhou et al., 2021a). For linear mixture MDPs, we further propose a HOME that constructs the variance-uncertainty-aware weights for high-order moments of the value function, which is pivotal to obtain a horizon-free regret for linear mixture MDPs.

For a better comparison between our results and previous results, we summarize these results in Table 1. It is evident that our results improve upon all previous results in the respective settings<sup>2</sup>.

**Notation.** We use lower case letters to denote scalars, and use lower and upper case bold face letters to denote vectors and matrices respectively. We denote by  $[n]$  the set  $\{1, \dots, n\}$ , by  $\overline{[n]}$  the set  $\{0, \dots, n-1\}$ . For a vector  $\mathbf{x} \in \mathbb{R}^d$  and a positive semi-definite matrix  $\Sigma \in \mathbb{R}^{d \times d}$ , we denote by  $\|\mathbf{x}\|_2$  the vector’s Euclidean norm and define  $\|\mathbf{x}\|_\Sigma = \sqrt{\mathbf{x}^\top \Sigma \mathbf{x}}$ . For  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ , let  $\mathbf{x} \odot \mathbf{y}$  be the Hadamard (componentwise) product of  $\mathbf{x}$  and  $\mathbf{y}$ . For two positive sequences  $\{a_n\}$  and  $\{b_n\}$  with  $n = 1, 2, \dots$ , we write  $a_n = O(b_n)$  if there exists an absolute constant  $C > 0$  such that  $a_n \leq Cb_n$  holds for all  $n \geq 1$  and write  $a_n = \Omega(b_n)$  if there exists an absolute constant  $C > 0$  such that  $a_n \geq Cb_n$  holds for all  $n \geq 1$ . We use  $\tilde{O}(\cdot)$  to further hide the polylogarithmic factors. We use  $\mathbb{1}\{\cdot\}$  to denote the indicator function. For  $a, b \in \mathbb{R}$  satisfying  $a \leq b$ , we use  $[x]_{[a,b]}$  to denote the truncation function  $x \cdot \mathbb{1}\{a \leq x \leq b\} + a \cdot \mathbb{1}\{x < a\} + b \cdot \mathbb{1}\{x > b\}$ .

## 2 Related work

In this section, we will review prior works that are most relevant to ours.

<sup>2</sup>The only exception is that for heterogeneous linear bandits, our algorithm needs to know the noise variance, while Zhang et al. (2021c); Kim et al. (2021) do not. But their algorithms are computationally inefficient. We plan to extend our algorithm to deal with unknown variance in the future work.

Table 1: Comparisons of regrets for linear bandits and linear mixture MDPs.

Algorithm	Regret	Assumption	Computationally Efficient?
OFUL (Abbasi-Yadkori et al., 2011)	$\tilde{O}(d\sqrt{K})$	-	Yes
WeightedOFUL (Zhou et al., 2021a)	$\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + \sqrt{dK} + d)$	Known variance	Yes
VOFUL (Zhang et al., 2021c)	$\tilde{O}(d^{4.5}\sqrt{\sum_{k=1}^K \sigma_k^2} + d^5)$	Unknown variance	No
VOFUL2 (Kim et al., 2021)	$\tilde{O}(d^{1.5}\sqrt{\sum_{k=1}^K \sigma_k^2} + d^2)$	Unknown variance	No
WeightedOFUL <sup>+</sup> (Theorem 4.1)	$\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + d)$	Known variance	Yes
UCRL-VTR (Jia et al., 2020; Ayoub et al., 2020)	$\tilde{O}(d\sqrt{H^3K})$	Homogeneous, $\sum_h r_h \leq H$	Yes
UCRL-VTR <sup>+</sup> (Zhou et al., 2021a)	$\tilde{O}(\sqrt{d^2H^3} + dH^4\sqrt{K} + d^2H^3 + d^3H^2)$	Inhomogeneous, $\sum_h r_h \leq H$	Yes
VARLin (Zhang et al., 2021c)	$\tilde{O}(d^{4.5}\sqrt{K} + d^9)$	Homogeneous, $\sum_h r_h \leq 1$	No
VARLin2 (Kim et al., 2021)	$\tilde{O}(d\sqrt{K} + d^2)$	Homogeneous, $\sum_h r_h \leq 1$	No
HF-UCRL-VTR <sup>+</sup> (Theorem 5.1)	$\tilde{O}(d\sqrt{K} + d^2)$	Homogeneous, $\sum_h r_h \leq 1$	Yes
Lower bound (Theorem 5.3)	$\Omega(d\sqrt{K})$	-	-

**Heterogeneous linear bandits.** Linear bandits have been studied for a long time. Most of existing works focus on the homogeneous linear bandits where the noise distributions at different round are identical (Auer, 2002; Chu et al., 2011; Li et al., 2010; Dani et al., 2008; Abbasi-Yadkori et al., 2011; Li et al., 2019a,b). Recently, a series of works focus on the heterogeneous linear bandits where the noise distribution changes over time. Lattimore et al. (2015) assumed that the noise distribution is Bernoulli and proposed an algorithm with an  $\tilde{O}(d\sqrt{K})$  regret. Kirschner and Krause (2018) assumed that the noise at  $k$ -th round is  $\sigma_k^2$ -sub-Gaussian, and they proposed a weighted ridge regression-based algorithm with an  $\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2})$  regret. A recent line of works assume the variance of the noise at  $k$ -th round is bounded by  $\sigma_k^2$ . Under this assumption, Zhang et al. (2021c) proposed a VOFUL algorithm with an  $\tilde{O}(d^{4.5}\sqrt{\sum_{k=1}^K \sigma_k^2} + d^5)$  regret. Kim et al. (2021) proposed a VOFUL2 algorithm with an  $\tilde{O}(d^{1.5}\sqrt{\sum_{k=1}^K \sigma_k^2} + d^2)$  regret. Both VOFUL and VOFUL2 do not need to know the variance information. However, they are computationally inefficient since they need to work with nonconvex confidence sets defined by a series of second-order constraints, and do not propose a polynomial-time algorithm to solve the maximization problem over these sets. With the variance information, Zhou et al. (2021a) proposed a computationally efficient WeightedOFUL with an  $\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + \sqrt{dK} + d)$  regret. Our work is under the same assumptions as Zhou et al. (2021a) and improves the regret for linear bandits.

**Horizon-free tabular RL.** RL is widely believed to be harder than contextual bandits problem due to its long planning horizon and the unknown state transitions. For tabular RL, under the assumption that the total reward obtained by any policy is upper bounded by 1, Jiang and Agarwal (2018) conjectured that any algorithm to find an  $\epsilon$ -optimal policy needs to have a polynomial dependence on the planning horizon  $H$  in the sample complexity. Such a conjecture was firstly refuted by Wang et al. (2020) by proposing a horizon-free algorithm with an  $\tilde{O}(|\mathcal{S}|^5|\mathcal{A}|^4\epsilon^{-2}\text{polylog}(H))$  sample complexity that depends on  $H$  polylogarithmically, where  $\epsilon$  is the target sub-optimality of the policy,  $\mathcal{S}$  is the state space and  $\mathcal{A}$  is the action space. Zhang et al. (2021b) proposed a near-optimal algorithm with an improved regret  $O((\sqrt{|\mathcal{S}||\mathcal{A}|K} + |\mathcal{S}|^2|\mathcal{A}|)\text{polylog}(H))$  and sample complexity. Similar

regret/sample complexity guarantees with a polylogarithmic  $H$  dependence have also been established under different RL settings (Zhang et al., 2020; Ren et al., 2021; Tarbouriech et al., 2021). Recently Li et al. (2022); Zhang et al. (2022) further proposed algorithms with  $H$ -independent regret/sample complexity guarantees. However, all the above works are limited to tabular RL. Our work proposes an algorithm with a regret bound that depends on  $H$  polylogarithmically for linear mixture MDPs, which extends these horizon-free tabular RL algorithms.

**RL with linear function approximation.** Recent years have witnessed a trend on RL with linear function approximation (e.g., Jiang et al., 2017; Dann et al., 2018; Yang and Wang, 2019a; Jin et al., 2020; Wang et al., 2019; Du et al., 2019; Sun et al., 2019; Zanette et al., 2020a,b; Weisz et al., 2020; Yang and Wang, 2019b; Modi et al., 2020; Jia et al., 2020; Ayoub et al., 2020; Zhou et al., 2021b). All these works assume that the MDP enjoys some linear representation and propose different statistical and computational complexities which depend on the dimension of the linear representation. Among these assumptions, our work falls into the category of *linear mixture MDP* which assumes that the transition dynamic can be represented as a linear combination of several basis transition probability functions (Yang and Wang, 2019b; Modi et al., 2020; Jia et al., 2020; Ayoub et al., 2020; Zhou et al., 2021b). Previous algorithms for linear mixture MDPs either suffer from a polynomial dependence on the episode horizon  $H$  (Yang and Wang, 2019b; Modi et al., 2020; Jia et al., 2020; Ayoub et al., 2020; Zhou et al., 2021b; Cai et al., 2019; He et al., 2021; Zhou et al., 2021a) or do not have a computationally efficient implementation (Zhang et al., 2021c; Kim et al., 2021). Our work achieves the best of both worlds for the first time.

### 3 Preliminaries

#### 3.1 Heterogeneous linear bandits

We consider the same heterogeneous linear bandits as studied in Zhou et al. (2021a). Let  $\{\mathcal{D}_k\}_{k=1}^\infty$  be decision sets that are fixed. At each round  $k$ , the agent selects an action  $\mathbf{a}_k \in \mathcal{D}_k$  satisfying  $\|\mathbf{a}_k\|_2 \leq A$ , then receives a reward  $r_k$  provided by the environment. Specifically,  $r_k$  is generated by  $r_k = \langle \boldsymbol{\theta}^*, \mathbf{a}_k \rangle + \epsilon_k$ , where  $\boldsymbol{\theta}^* \in \mathbb{R}^d$  is an unknown vector, and  $\epsilon_k$  is a random noise satisfying

$$\forall k, |\epsilon_k| \leq R, \mathbb{E}[\epsilon_k | \mathbf{a}_{1:k}, \epsilon_{1:k-1}] = 0, \mathbb{E}[\epsilon_k^2 | \mathbf{a}_{1:k}, \epsilon_{1:k-1}] \leq \sigma_k^2,$$

where  $\sigma_k$  is an upper bound of the variance of the noise  $\epsilon_k$  that are observable to the agent. We assume that  $\sigma_k$  is  $(\mathbf{a}_{1:k}, \epsilon_{1:k-1})$ -measurable. The agent aims to minimize the *pseudo-regret* defined as follows:

$$\text{Regret}(K) = \sum_{k=1}^K [\langle \mathbf{a}_k^*, \boldsymbol{\theta}^* \rangle - \langle \mathbf{a}_k, \boldsymbol{\theta}^* \rangle], \text{ where } \mathbf{a}_k^* = \underset{\mathbf{a} \in \mathcal{D}_k}{\text{argmax}} \langle \mathbf{a}, \boldsymbol{\theta}^* \rangle.$$

#### 3.2 Episodic reinforcement learning

We also study RL with linear function approximation for episodic linear mixture MDPs. We introduce the necessary definitions of MDPs here. The reader can refer to Puterman (2014) for more details.

**Episodic MDP.** We denote a homogeneous, episodic MDP by a tuple  $M = M(\mathcal{S}, \mathcal{A}, H, r, \mathbb{P})$ , where  $\mathcal{S}$  is the state space and  $\mathcal{A}$  is the action space,  $H$  is the length of the episode,  $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  is the deterministic reward function, and  $\mathbb{P}$  is the transition probability function. For the sake of simplicity, we restrict ourselves to countable state space and finite action space. A policy  $\pi = \{\pi_h\}_{h=1}^H$  is a collection of  $H$  functions, where each of them maps a state  $s$  to an action  $a$ .

**Value function and regret.** For  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , we define the action-value function  $Q_h^\pi(s, a)$  and (state) value function  $V_h^\pi(s)$  as follows:

$$Q_h^\pi(s, a) = \mathbb{E} \left[ \sum_{h'=h}^H r(s_{h'}, a_{h'}) \middle| s_h = s, a_h = a, s_{h'} \sim \mathbb{P}(\cdot | s_{h'-1}, a_{h'-1}), a_{h'} = \pi_{h'}(s_{h'}) \right],$$

$$V_h^\pi(s) = Q_h^\pi(s, \pi_h(s)), V_{H+1}^\pi(s) = 0.$$

The optimal value function  $V_h^*(\cdot)$  and the optimal action-value function  $Q_h^*(\cdot, \cdot)$  are defined by  $V_h^*(s) = \sup_\pi V_h^\pi(s)$  and  $Q_h^*(s, a) = \sup_\pi Q_h^\pi(s, a)$ , respectively. For any function  $V : \mathcal{S} \rightarrow \mathbb{R}$ ,

we introduce the following shorthands to denote the conditional variance of  $V$  at  $\mathbb{P}(\cdot|s, a)$ :

$$[\mathbb{P}V](s, a) = \mathbb{E}_{s' \sim \mathbb{P}(\cdot|s, a)} V(s'), \quad [\mathbb{V}V](s, a) = [\mathbb{P}V^2](s, a) - ([\mathbb{P}V](s, a))^2,$$

where  $V^2$  stands for the function whose value at  $s$  is  $V^2(s)$ . Using this notation, the Bellman equations for policy  $\pi$  and the Bellman optimality equation can be written as

$$Q_h^\pi(s, a) = r(s, a) + [\mathbb{P}V_{h+1}^\pi](s, a), \quad Q_h^*(s, a) = r(s, a) + [\mathbb{P}V_{h+1}^*](s, a).$$

The goal is to minimize the  $K$ -episode regret defined as follows:

$$\text{Regret}(K) = \sum_{k=1}^K [V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k)].$$

In this paper, we focus on proving high probability bounds on the regret  $\text{Regret}(K)$ .

In this work we make the following assumptions. The first assumption assumes that for any policy, the accumulated reward of an episode is upper bounded by 1, which has been considered in previous works (Krishnamurthy et al., 2016; Jiang and Agarwal, 2018). The accumulated reward assumption ensures that the only factor that can affect the final statistical complexity is the planning difficulty brought by the episode length, rather than the scale of the reward.

**Assumption 3.1** (Bounded total reward). For any policy  $\pi$ , let  $\{s_h, a_h\}_{h=1}^H$  be any states and actions satisfying  $a_h = \pi_h(s_h)$  and  $s_{h+1} \sim \mathbb{P}(\cdot|s_h, a_h)$ . Then we have  $0 \leq \sum_{h=1}^H r(s_h, a_h) \leq 1$ .

Next assumption assumes that the transition dynamic enjoys a linearized representation w.r.t. some feature mapping. We define the *linear mixture MDPs* (Jia et al., 2020; Ayoub et al., 2020; Zhou et al., 2021b) as follows.

**Assumption 3.2** (Linear mixture MDP).  $M$  is an episodic  $B$ -bounded linear mixture MDP, such that there exists a vector  $\theta^* \in \mathbb{R}^d$  and  $\phi(\cdot|\cdot, \cdot)$  such that  $\mathbb{P}(s'|s, a) = \langle \phi(s'|s, a), \theta^* \rangle$  for any state-action-next-state triplet  $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ . Meanwhile,  $\|\theta^*\|_2 \leq B$  and for any bounded function  $V : \mathcal{S} \rightarrow [0, 1]$  and any tuple  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , we have

$$\|\phi_V(s, a)\|_2 \leq 1, \text{ where } \phi_V(s, a) = \sum_{s' \in \mathcal{S}} \phi(s'|s, a)V(s').$$

Lastly, for any  $V : \mathcal{S} \rightarrow [0, 1]$ ,  $\phi_V$  can be calculated efficiently within  $\mathcal{O}$  time.

**Remark 3.3.** A key property of linear mixture MDP is that for any function  $V : \mathcal{S} \rightarrow \mathbb{R}$  and any state-action pair  $(s, a)$ , the conditional expectation of  $V$  over  $\mathbb{P}(\cdot|s, a)$  is a linear function of  $\theta^*$ , i.e.,  $[\mathbb{P}V](s, a) = \langle \phi_V(s, a), \theta^* \rangle$ . Meanwhile, the conditional variance of  $V$  over  $\mathbb{P}(\cdot|s, a)$  is a quadratic function of  $\theta^*$ , i.e.,  $[\mathbb{V}V](s, a) = \langle \phi_{V^2}(s, a), \theta^* \rangle - [\langle \phi_V(s, a), \theta^* \rangle]^2$ .

**Remark 3.4.** For a general class of  $\phi$ ,  $\phi_V(s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$  can be computed efficiently for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  if  $V : \mathcal{S} \rightarrow \mathbb{R}$  can be computed efficiently. For instance,  $\phi(s'|s, a) = \mathbf{e}_{s', s, a} \in \mathbb{R}^{|\mathcal{S}|^2|\mathcal{A}|}$  and  $\phi(s'|s, a) = \psi(s') \odot \xi(s, a)$ , where  $\psi, \xi$  are two sub feature functions. More discussions are referred to Zhou et al. (2021a).

## 4 Computationally efficient variance-aware linear bandits

In this section, we propose our algorithm  $\text{WeightedOFUL}^+$  in Algorithm 1 for the heterogeneous linear bandits introduced in Section 3.1.  $\text{WeightedOFUL}^+$  adopts the *weighted ridge regression* estimator used by  $\text{WeightedOFUL}$  (Zhou et al., 2021a), but uses a refined weight. It first computes a weighted estimate of  $\theta^*$ , denoted by  $\hat{\theta}_k$ , based on previous contexts and rewards, where the weights  $\bar{\sigma}_k$  are computed by the noise variance  $\sigma_k$ . Then  $\text{WeightedOFUL}^+$  constructs the confidence set of  $\theta^*$ , denoted by  $\hat{\mathcal{C}}_k$ , estimates the reward  $\langle \mathbf{a}, \theta \rangle$  for  $\theta \in \hat{\mathcal{C}}_k$ , and selects the arm that maximizes the estimated reward optimistically. The selection rule of  $\mathbf{a}_k$  is identical to selecting the best arm w.r.t. to their upper confidence bound, i.e.,  $\mathbf{a}_k \leftarrow \arg\max_{\mathbf{a} \in \mathcal{D}_k} \langle \mathbf{a}, \hat{\theta}_k \rangle + \hat{\beta}_k \|\mathbf{a}\|_{\hat{\Sigma}_k^{-1}}$  (Li et al., 2010).  $\text{WeightedOFUL}^+$  constructs  $\bar{\sigma}_k^2$  as the maximization of the variance, a constant, and the uncertainty  $\|\mathbf{a}_k\|_{\hat{\Sigma}_k^{-1}}$ , as defined in (4.1). Note that Zhou et al. (2021a) proposed a *variance-aware* weight  $\bar{\sigma}_k = \max\{\sigma_k, \alpha\}$  for heterogeneous linear bandits. The additional uncertainty term in our weight

---

**Algorithm 1** WeightedOFUL<sup>+</sup>


---

**Require:** Regularization parameter  $\lambda > 0$ , and  $B$ , an upper bound on the  $\ell_2$ -norm of  $\theta^*$ , confidence radius  $\hat{\beta}_k$ , variance parameters  $\alpha, \gamma$

- 1:  $\hat{\Sigma}_1 \leftarrow \lambda \mathbf{I}, \hat{\mathbf{b}}_1 \leftarrow \mathbf{0}, \hat{\theta}_1 \leftarrow \mathbf{0}, \hat{\beta}_1 = \sqrt{\lambda} B$
- 2: **for**  $k = 1, \dots, K$  **do**
- 3:   Let  $\hat{\mathcal{C}}_k \leftarrow \{\theta : \|\hat{\Sigma}_k^{1/2}(\theta - \hat{\theta}_k)\|_2 \leq \hat{\beta}_k\}$ , observe  $\mathcal{D}_k$
- 4:   Set  $(\mathbf{a}_k, \theta_k) \leftarrow \operatorname{argmax}_{\mathbf{a} \in \mathcal{D}_k, \theta \in \hat{\mathcal{C}}_k} \langle \mathbf{a}, \theta \rangle$
- 5:   Observe  $(r_k, \sigma_k)$ , set  $\bar{\sigma}_k$  as

$$\bar{\sigma}_k \leftarrow \max\{\sigma_k, \alpha, \gamma \|\mathbf{a}_k\|_{\hat{\Sigma}_k^{-1}}^{1/2}\} \quad (4.1)$$

- 6:    $\hat{\Sigma}_{k+1} \leftarrow \hat{\Sigma}_k + \mathbf{a}_k \mathbf{a}_k^\top / \bar{\sigma}_k^2, \hat{\mathbf{b}}_{k+1} \leftarrow \hat{\mathbf{b}}_k + r_k \mathbf{a}_k / \bar{\sigma}_k^2, \hat{\theta}_{k+1} \leftarrow \hat{\Sigma}_{k+1}^{-1} \hat{\mathbf{b}}_{k+1}$
- 7: **end for**

---

enables us to build a tighter confidence set  $\hat{\mathcal{C}}_k$  since the uncertainty of arms can be leveraged by a tighter Bernstein-type concentration inequality (See Section 4.1 for more details). Meanwhile, we notice that He et al. (2022) proposed a pure *uncertainty-aware* weight  $\bar{\sigma}_k = \max\{\alpha, \gamma \|\mathbf{a}_k\|_{\hat{\Sigma}_k^{-1}}^{1/2}\}$  to deal with the corruption in contextual linear bandits, which serves a different purpose compared to our setting (there is no corruption in our bandit model).

The following theorem gives the regret bound of WeightedOFUL<sup>+</sup>.

**Theorem 4.1.** Let  $0 < \delta < 1$ . Suppose that for all  $k \geq 1$  and all  $\mathbf{a} \in \mathcal{D}_k$ ,  $\langle \mathbf{a}, \theta^* \rangle \in [-1, 1]$ ,  $\|\theta^*\|_2 \leq B$  and  $\{\hat{\beta}_k\}_{k \geq 1}$  are set to

$$\begin{aligned} \hat{\beta}_k &= 12\sqrt{d \log(1 + kA^2/(\alpha^2 d\lambda)) \log(32(\log(\gamma^2/\alpha) + 1)k^2/\delta)} \\ &\quad + 30 \log(32(\log(\gamma^2/\alpha) + 1)k^2/\delta) R/\gamma^2 + \sqrt{\lambda} B. \end{aligned} \quad (4.2)$$

Then with probability at least  $1 - \delta$ , the regret of WeightedOFUL<sup>+</sup> is bounded by

$$\operatorname{Regret}(K) \leq 4d\iota + 4d\gamma^2 \hat{\beta}_K \iota + 4\hat{\beta}_K \sqrt{\sum_{k=1}^K \sigma_k^2 + K\alpha^2 d\iota}, \quad (4.3)$$

where  $\iota = \log(1 + KA^2/(d\lambda\alpha^2))$ . Moreover, setting  $\alpha = 1/\sqrt{K}, \gamma = R^{1/2}/d^{1/4}$  and  $\lambda = d/B^2$  yields a high probability regret  $\operatorname{Regret}(K) = \tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + dR + d)$ .

**Remark 4.2.** Treating  $R$  as a constant, the regret of WeightedOFUL<sup>+</sup> becomes  $\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + d)$ .

It strictly outperforms the  $\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + \sqrt{dK} + d)$  regret achieved in Zhou et al. (2021a).

Compared with the  $\tilde{O}(d^{4.5}\sqrt{\sum_{k=1}^K \sigma_k^2} + d^5)$  regret by VOFUL (Zhang et al., 2021c) and the  $\tilde{O}(d^{1.5}\sqrt{\sum_{k=1}^K \sigma_k^2} + d^2)$  regret by VOFUL2 (Kim et al., 2021), the regret of WeightedOFUL<sup>+</sup> has a better dependence on  $d$ , and WeightedOFUL<sup>+</sup> is computationally efficient. It is worth noting that both VOFUL and VOFUL2 do not need to know the variance  $\sigma_k^2$  while our algorithm does.

Whether there exists an algorithm that can achieve the  $\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + \sqrt{dK} + d)$  regret without knowing the variance information remains an open problem.

#### 4.1 Proof sketch

We give a proof sketch of Theorem 4.1 along with two key lemmas that are pivotal to obtain the improved regret. To begin with, we first show that  $\theta^*$  belongs to confidence balls centering at  $\hat{\theta}_k$  with radius  $\hat{\beta}_k$ . This can be proved by the following lemma, which is an improved version of the Bernstein-type self-normalized martingale inequality proposed by Zhou et al. (2021a).

**Lemma 4.3.** Let  $\{\mathcal{G}_k\}_{k=1}^\infty$  be a filtration, and  $\{\mathbf{x}_k, \eta_k\}_{k \geq 1}$  be a stochastic process such that  $\mathbf{x}_k \in \mathbb{R}^d$  is  $\mathcal{G}_k$ -measurable and  $\eta_k \in \mathbb{R}$  is  $\mathcal{G}_{k+1}$ -measurable. Let  $L, \sigma, \lambda, \epsilon > 0, \boldsymbol{\mu}^* \in \mathbb{R}^d$ . For  $k \geq 1$ , let

$y_k = \langle \boldsymbol{\mu}^*, \mathbf{x}_k \rangle + \eta_k$  and suppose that  $\eta_k, \mathbf{x}_k$  also satisfy

$$\mathbb{E}[\eta_k | \mathcal{G}_k] = 0, \mathbb{E}[\eta_k^2 | \mathcal{G}_k] \leq \sigma^2, |\eta_k| \leq R, \|\mathbf{x}_k\|_2 \leq L. \quad (4.4)$$

For  $k \geq 1$ , let  $\mathbf{Z}_k = \lambda \mathbf{I} + \sum_{i=1}^k \mathbf{x}_i \mathbf{x}_i^\top$ ,  $\mathbf{b}_k = \sum_{i=1}^k y_i \mathbf{x}_i$ ,  $\boldsymbol{\mu}_k = \mathbf{Z}_k^{-1} \mathbf{b}_k$ , and

$$\begin{aligned} \beta_k &= 12\sqrt{\sigma^2 d \log(1 + kL^2/(d\lambda)) \log(32(\log(R/\epsilon) + 1)k^2/\delta)} \\ &\quad + 24 \log(32(\log(R/\epsilon) + 1)k^2/\delta) \max_{1 \leq i \leq k} \{|\eta_i| \min\{1, \|\mathbf{x}_i\|_{\mathbf{Z}_{i-1}^{-1}}\}\} + 6 \log(32(\log(R/\epsilon) + 1)k^2/\delta)\epsilon. \end{aligned}$$

Then, for any  $0 < \delta < 1$ , we have with probability at least  $1 - \delta$  that,

$$\forall k \geq 1, \left\| \sum_{i=1}^k \mathbf{x}_i \eta_i \right\|_{\mathbf{Z}_k^{-1}} \leq \beta_k, \|\boldsymbol{\mu}_k - \boldsymbol{\mu}^*\|_{\mathbf{Z}_k} \leq \beta_k + \sqrt{\lambda} \|\boldsymbol{\mu}^*\|_2,$$

Note that  $\hat{\boldsymbol{\theta}}_k$  can be regarded as  $\boldsymbol{\mu}_{k-1}$  in Lemma 4.3 with  $\epsilon = R/\gamma^2$ ,  $\mathbf{x}_k = \mathbf{a}_k/\bar{\sigma}_k$ ,  $y_k = r_k/\bar{\sigma}_k$ ,  $\eta_k = \epsilon_k/\bar{\sigma}_k$ ,  $\mathbf{Z}_k = \hat{\boldsymbol{\Sigma}}_{k+1}$  and  $\boldsymbol{\mu}^* = \boldsymbol{\theta}^*$ . With the help of the weight  $\bar{\sigma}_k$ , the variance of  $\eta_k$  is upper bounded by 1 (since  $\bar{\sigma}_k \geq \sigma_k$ ) and  $|\eta_k| \min\{1, \|\mathbf{x}_k\|_{\mathbf{Z}_k^{-1}}\} \leq |\epsilon_k| \|\mathbf{a}_k\|_{\hat{\boldsymbol{\Sigma}}_k^{-1}}/\bar{\sigma}_k^2 \leq R/\gamma^2$  (since  $\bar{\sigma}_k^2 \geq \gamma^2 \|\mathbf{a}_k\|_{\hat{\boldsymbol{\Sigma}}_k^{-1}}$ ). Therefore, by Lemma 4.3, w.h.p.  $\boldsymbol{\theta}^* \in \hat{\mathcal{C}}_k$ . Following the standard procedure to bound the regret of the optimistic algorithm (Abbasi-Yadkori et al., 2011), we have

$$\text{Regret}(K) = \sum_{k=1}^K \langle \mathbf{a}_k^* - \mathbf{a}_k, \boldsymbol{\theta}^* \rangle \leq 2 \sum_{k=1}^K \min\{1, \hat{\beta}_k \|\mathbf{a}_k\|_{\hat{\boldsymbol{\Sigma}}_k^{-1}}\}. \quad (4.5)$$

Next lemma gives an upper bound of (4.5).

**Lemma 4.4.** Let  $\{\sigma_k, \hat{\beta}_k\}_{k \geq 1}$  be a sequence of non-negative numbers,  $\alpha, \gamma > 0$ ,  $\{\mathbf{a}_k\}_{k \geq 1} \subset \mathbb{R}^d$  and  $\|\mathbf{a}_k\|_2 \leq A$ . Let  $\{\bar{\sigma}_k\}_{k \geq 1}$  and  $\{\hat{\boldsymbol{\Sigma}}_k\}_{k \geq 1}$  be (recursively) defined as follows:  $\hat{\boldsymbol{\Sigma}}_1 = \lambda \mathbf{I}$ ,

$$\forall k \geq 1, \bar{\sigma}_k = \max\{\sigma_k, \alpha, \gamma \|\mathbf{a}_k\|_{\hat{\boldsymbol{\Sigma}}_k^{-1}}^{1/2}\}, \hat{\boldsymbol{\Sigma}}_{k+1} = \hat{\boldsymbol{\Sigma}}_k + \mathbf{a}_k \mathbf{a}_k^\top / \bar{\sigma}_k^2.$$

Let  $\iota = \log(1 + KA^2/(d\lambda\alpha^2))$ . Then we have

$$\sum_{k=1}^K \min\left\{1, \hat{\beta}_k \|\mathbf{a}_k\|_{\hat{\boldsymbol{\Sigma}}_k^{-1}}\right\} \leq 2d\iota + 2 \max_{k \in [K]} \hat{\beta}_k \gamma^2 d\iota + 2\sqrt{d\iota} \sqrt{\sum_{k=1}^K \hat{\beta}_k^2 (\sigma_k^2 + \alpha^2)}.$$

By Lemma 4.4, the regret of WeightedOFUL<sup>+</sup> can be bounded by

$$\text{Regret}(K) = \tilde{O}(d + \sqrt{d} \hat{\beta}_K (\sqrt{\sum_{k=1}^K \sigma_k^2 + K\alpha^2} + \sqrt{d}\gamma^2)) \quad (4.6)$$

with  $\hat{\beta}_K = \tilde{O}(\sqrt{d} + R/\gamma^2 + \sqrt{\lambda}B)$ , which finishes the proof.

Here we compare the regret of WeightedOFUL<sup>+</sup> and the regret of WeightedOFUL (Zhou et al., 2021a) (which chooses the weight  $\bar{\sigma}_k = \max\{\sigma_k, \alpha\}$ ) to see where the improvement comes from. In particular, Zhou et al. (2021a) proved the regret of WeightedOFUL as follows

$$\text{Regret}(K) = \tilde{O}(d + \sqrt{d} \hat{\beta}_K^\alpha \sqrt{\sum_{k=1}^K \sigma_k^2 + K\alpha^2}) \quad (4.7)$$

with  $\hat{\beta}_K^\alpha = \tilde{O}(\sqrt{d} + R/\alpha + \sqrt{\lambda}B)$ . At the first glance, both (4.6) and (4.7) have a  $\sqrt{K\alpha^2}$  term. However, thanks to the design of  $\bar{\sigma}_k$  and Lemma 4.3, the  $\sqrt{K\alpha^2}$  term in (4.6) can be shaved by choosing a small enough  $\alpha$ , while this term in (4.7) cannot be shaved due to the existence of a  $R/\alpha$  term in  $\hat{\beta}_K^\alpha$ .

## 5 Computationally efficient horizon-free RL for linear mixture MDPs

In this section, we propose a horizon-free RL algorithm HF-UCRL-VTR<sup>+</sup> in Algorithm 2. We leave the discussion of the computational complexity of HF-UCRL-VTR<sup>+</sup> to Appendix B.

HF-UCRL-VTR<sup>+</sup> follows the *value targeted regression (VTR)* framework proposed by Jia et al. (2020); Ayoub et al. (2020) to learn the linear mixture MDP. In detail, following the observation in

---

**Algorithm 2** HF-UCRL-VTR<sup>+</sup>


---

**Require:** Regularization parameter  $\lambda$ , an upper bound  $B$  of the  $\ell_2$ -norm of  $\theta^*$ , confidence radius

- $\{\hat{\beta}_k\}_{k \geq 1}$ , level  $M$ , variance parameters  $\alpha, \gamma$ ,  $[\overline{M}] = \{0, \dots, M-1\}$
- 1: For  $m \in [\overline{M}]$ , set  $\hat{\theta}_{1,m} \leftarrow \mathbf{0}$ ,  $\tilde{\Sigma}_{0,H+1,m} \leftarrow \lambda \mathbf{I}$ ,  $\tilde{\mathbf{b}}_{0,H+1,m} \leftarrow \mathbf{0}$ . Set  $V_{1,H+1}(\cdot) \leftarrow 0$
  - 2: **for**  $k = 1, \dots, K$  **do**
  - 3:   **for**  $h = H, \dots, 1$  **do**
  - 4:     Set  $Q_{k,h}(\cdot, \cdot) \leftarrow \left[ r(\cdot, \cdot) + \langle \hat{\theta}_{k,0}, \phi_{V_{k,h+1}}(\cdot, \cdot) \rangle + \hat{\beta}_k \left\| \tilde{\Sigma}_{k,0}^{-1/2} \phi_{V_{k,h+1}}(\cdot, \cdot) \right\|_2 \right]_{[0,1]}$
  - 5:     Set  $\pi_h^k(\cdot) \leftarrow \operatorname{argmax}_{a \in \mathcal{A}} Q_{k,h}(\cdot, a)$
  - 6:     Set  $V_{k,h}(\cdot) \leftarrow \max_{a \in \mathcal{A}} Q_{k,h}(\cdot, a)$
  - 7:   **end for**
  - 8:   Receive  $s_1^k$ . For  $m \in [\overline{M}]$ , set  $\tilde{\Sigma}_{k,1,m} \leftarrow \tilde{\Sigma}_{k-1,H+1,m}$
  - 9:   **for**  $h = 1, \dots, H$  **do**
  - 10:     Take action  $a_h^k \leftarrow \pi_h^k(s_h^k)$ , receive  $s_{h+1}^k \sim \mathbb{P}(\cdot | s_h^k, a_h^k)$ .
  - 11:     For  $m \in [\overline{M}]$ , denote  $\phi_{k,h,m} = \phi_{V_{k,h+1}^{2m}}(s_h^k, a_h^k)$ .
  - 12:     Set  $\{\bar{\sigma}_{k,h,m}\}_{m \in [\overline{M}]} \leftarrow \text{Algorithm 3}(\{\phi_{k,h,m}, \hat{\theta}_{k,m}, \tilde{\Sigma}_{k,h,m}, \hat{\Sigma}_{k,m}\}_{m \in [\overline{M}]}, \hat{\beta}_k, \alpha, \gamma)$
  - 13:     For  $m \in [\overline{M}]$ , set  $\tilde{\Sigma}_{k,h+1,m} \leftarrow \tilde{\Sigma}_{k,h,m} + \phi_{k,h,m} \phi_{k,h,m}^\top / \bar{\sigma}_{k,h,m}^2$
  - 14:     For  $m \in [\overline{M}]$ , set  $\tilde{\mathbf{b}}_{k,h+1,m} \leftarrow \tilde{\mathbf{b}}_{k,h,m} + \phi_{k,h,m} V_{k,h+1}^{2m}(s_{h+1}^k) / \bar{\sigma}_{k,h,m}^2$
  - 15:   **end for**
  - 16:   For  $m \in [\overline{M}]$ , set  $\hat{\Sigma}_{k+1,m} \leftarrow \tilde{\Sigma}_{k,H+1,m}$ ,  $\hat{\mathbf{b}}_{k+1,m} \leftarrow \tilde{\mathbf{b}}_{k,H+1,m}$ ,  $\hat{\theta}_{k+1,m} \leftarrow \hat{\Sigma}_{k+1,m}^{-1} \hat{\mathbf{b}}_{k+1,m}$
  - 17: **end for**
- 

**Remark 3.3,** VTR estimates  $\theta^*$  by solving a regression problem over predictors/contexts  $\phi_{k,h,0} = \phi_{V_{k,h+1}}(s_h^k, a_h^k)$  and responses  $V_{k,h+1}(s_{h+1}^k)$ . Specifically, HF-UCRL-VTR<sup>+</sup> takes the estimate  $\hat{\theta}_{k,0}$  as the solution to the following weighted regression problem:

$$\hat{\theta}_{k,0} = \operatorname{argmin}_{\theta \in \mathbb{R}^d} \lambda \|\theta\|_2^2 + \sum_{j=1}^{k-1} \sum_{h=1}^H [\langle \phi_{j,h,0}, \theta \rangle - V_{j,h+1}(s_{h+1}^j)]^2 / \bar{\sigma}_{j,h,0}^2, \quad (5.1)$$

where  $\bar{\sigma}_{j,h,0}$  is the upper bound of the conditional variance  $[\mathbb{V}V_{k,h+1}](s_h^k, a_h^k)$ . Such a weighted regression scheme has been adapted by UCRL-VTR<sup>+</sup> in Zhou et al. (2021a). With  $\hat{\theta}_{k,0}$ , HF-UCRL-VTR<sup>+</sup> then constructs the optimistic estimates  $Q_{k,h}$  (resp.  $V_{k,h}$ ) of the optimal value functions  $Q_h^*$  (resp.  $V_h^*$ ) and takes actions optimistically. Note that  $\hat{\theta}_{k,0}$  is updated at the end of each episode. We highlight several improved algorithm designs of HF-UCRL-VTR<sup>+</sup> compared to UCRL-VTR<sup>+</sup> as follows.

**Improved weighted linear regression estimator.** HF-UCRL-VTR<sup>+</sup> sets  $\bar{\sigma}_{j,h,0}$  similar to the weight used in WeightedOFUL<sup>+</sup>. Assuming that the conditional variance  $[\mathbb{V}V_{k,h+1}](s_h^k, a_h^k)$  can be computed for any value function  $V$  and state action pair  $(s, a)$ , then the weight can be set as

$$\bar{\sigma}_{k,h,0}^2 = \max\{[\mathbb{V}V_{k,h+1}](s_h^k, a_h^k), \alpha^2, \gamma^2 \|\tilde{\Sigma}_{k,h,0}^{-1/2} \phi_{k,h,0}\|_2\},$$

where  $\tilde{\Sigma}_{k,h,0}$  is the weighted sample covariance matrix of  $\phi_{k,h,0}$  up to  $k$ -th episode and  $h$ -th stage. However, the true variance is not accessible since  $\mathbb{P}(\cdot | s, a)$  is unknown. Therefore, HF-UCRL-VTR<sup>+</sup> replaces  $[\mathbb{V}V_{k,h+1}](s_h^k, a_h^k)$  with its estimate  $[\bar{\mathbb{V}}_{k,0} V_{k,h+1}](s_h^k, a_h^k)$  and an error bound  $E_{k,h,0}$  satisfying  $[\bar{\mathbb{V}}_{k,0} V_{k,h+1}](s_h^k, a_h^k) + E_{k,h,0} \geq [\mathbb{V}V_{k,h+1}](s_h^k, a_h^k)$  with high probability. Thanks to the fact that  $[\mathbb{V}V_{k,h+1}](s_h^k, a_h^k)$  is a quadratic function of  $\theta^*$  as illustrated in Remark 3.3,  $[\bar{\mathbb{V}}_{k,0} V_{k,h+1}](s_h^k, a_h^k)$  can be estimated as follows:

$$[\bar{\mathbb{V}}_{k,0} V_{k,h+1}](s_h^k, a_h^k) = [\langle \phi_{k,h,1}, \hat{\theta}_{k,1} \rangle]_{[0,1]} - [\langle \phi_{k,h,0}, \hat{\theta}_{k,0} \rangle]_{[0,1]}^2, \quad (5.2)$$

where  $\hat{\theta}_{k,1}$  is the solution to some regression problem over predictors/contexts  $\phi_{k,h,1} = \phi_{V_{k,h+1}^2}(s_h^k, a_h^k)$  and responses  $V_{k,h+1}^2(s_{h+1}^k)$ .

**Higher-order moment regression.** To obtain a better estimate, it is natural to set  $\hat{\theta}_{k,1}$  as the solution to the weighted regression problem on  $\phi_{k,h,1}$  and  $V_{k,h+1}^2(s_{h+1}^k)$  with weight  $\bar{\sigma}_{k,h,1}$ . Here  $\bar{\sigma}_{k,h,1}$  is



---

**Algorithm 3** High-order moment estimator (HOME)
 

---

**Require:** Features  $\{\phi_{k,h,m}\}_{m \in \overline{[M]}}$ , vector estimators  $\{\hat{\theta}_{k,m}\}_{m \in \overline{[M]}}$ , covariance matrix  $\{\tilde{\Sigma}_{k,h,m}, \hat{\Sigma}_{k,m}\}_{m \in \overline{[M]}}$ , confidence radius  $\hat{\beta}_k, \alpha, \gamma$

- 1: **for**  $m = 0, \dots, M - 2$  **do**
- 2:   Set  $[\bar{V}_{k,m} V_{k,h+1}^{2^m}](s_h^k, a_h^k) \leftarrow [\langle \phi_{k,h,m+1}, \hat{\theta}_{k,m+1} \rangle]_{[0,1]} - [\langle \phi_{k,h,m}, \hat{\theta}_{k,m} \rangle]_{[0,1]}^2$
- 3:   Set  $E_{k,h,m} \leftarrow \min \{1, 2\hat{\beta}_k \|\tilde{\Sigma}_{k,m}^{-1/2} \phi_{k,h,m}\|_2\} + \min \{1, \hat{\beta}_k \|\hat{\Sigma}_{k,m+1}^{-1/2} \phi_{k,h,m+1}\|_2\}$
- 4:   Set  $\bar{\sigma}_{k,h,m}^2 \leftarrow \max \{[\bar{V}_{k,m} V_{k,h+1}^{2^m}](s_h^k, a_h^k) + E_{k,h,m}, \alpha^2, \gamma^2 \|\tilde{\Sigma}_{k,h,m}^{-1/2} \phi_{k,h,m}\|_2\}$
- 5: **end for**
- 6: Set  $\bar{\sigma}_{k,h,M-1}^2 \leftarrow \max \{1, \alpha^2, \gamma^2 \|\tilde{\Sigma}_{k,h,M-1}^{-1/2} \phi_{k,h,M-1}\|_2\}$

**Ensure:**  $\{\bar{\sigma}_{k,h,m}\}_{m \in \overline{[M]}}$

---

constructed in a similar way to  $\bar{\sigma}_{k,h,0}$ , which relies on the conditional variance of  $[\mathbb{V}V_{k,h+1}^2](s_h^k, a_h^k)$ . By repeating this process, we recursively estimate the conditional  $2^m$ -th moment of  $V_{k,h+1}$  by its variance, which is the conditional  $2^{m+1}$ -th moment of  $V_{k,h+1}$ . It is worth noting that the idea of high-order recursive estimation has been used in Li et al. (2020) and later in Zhang et al. (2021b,c) to achieve horizon-free regret/sample complexity guarantees. Similar recursive analysis also appeared in Lattimore and Hutter (2012).

The estimated conditional moment  $[\bar{V}_{k,m} V_{k,h+1}^{2^m}](s_h^k, a_h^k)$  relies on  $\hat{\theta}_{k,m+1}$  and  $\hat{\theta}_{k,m}$ , and  $\langle \phi_{k,h,m+1}, \hat{\theta}_{k,m+1} \rangle$  serves as the estimate of the higher-moment  $[\mathbb{V}V_{k,h+1}^{2^{m+1}}](s_h^k, a_h^k)$ . The detailed constructions for the high-order moment estimator are summarized in Algorithm 3.

We provide the regret bound for HF-UCRL-VTR<sup>+</sup> here.

**Theorem 5.1.** Set  $M = \log(3KH)/\log 2$ . For any  $\delta > 0$ , set  $\{\hat{\beta}_k\}_{k \geq 1}$  as

$$\begin{aligned} \hat{\beta}_k &= 12\sqrt{d \log(1 + kH/(\alpha^2 d \lambda))} \log(32(\log(\gamma^2/\alpha) + 1)k^2 H^2/\delta) \\ &\quad + 30 \log(32(\log(\gamma^2/\alpha) + 1)k^2 H^2/\delta)/\gamma^2 + \sqrt{\lambda} B, \end{aligned} \quad (5.3)$$

then with probability at least  $1 - (2M + 1)\delta$ , the regret of Algorithm 2 is bounded by

$$\begin{aligned} \text{Regret}(K) &\leq 12(8d\iota + 8\hat{\beta}_K \gamma^2 d\iota + 4\hat{\beta}_K \sqrt{d\iota} \sqrt{Md\iota/2 + KH\alpha^2} + \sqrt{Md\iota} \zeta + \zeta) \\ &\quad + 864 \max\{8\hat{\beta}_K^2 d\iota, \zeta\} + Md\iota/2 + [\sqrt{2 \log(1/\delta)} + 32 \max\{2\hat{\beta}_K \sqrt{2d\iota}, \sqrt{\zeta}\}] \sqrt{K}, \end{aligned}$$

where  $\iota = \log(1 + KH/(d\lambda\alpha^2))$ ,  $\zeta = 4 \log(4 \log(KH)/\delta)$ . Moreover, setting  $\alpha = \sqrt{d/(KH)}$ ,  $\gamma = 1/d^{1/4}$  and  $\lambda = d/B^2$  yields a high-probability regret  $\text{Regret}(K) = \tilde{O}(d\sqrt{K} + d^2)$ .

**Remark 5.2.** The regret of HF-UCRL-VTR<sup>+</sup> is strictly better than that of VOFUL  $\tilde{O}(d^{4.5}\sqrt{K} + d^9)$  Zhang et al. (2021c), and it matches the regret of VOFUL2 (Kim et al., 2021). More importantly, HF-UCRL-VTR<sup>+</sup> is computationally efficient, while there is no efficient implementation of VOFUL/VOFUL2.

Next theorem provides the regret lower bound and suggests that the regret obtained by HF-UCRL-VTR<sup>+</sup> is near-optimal. The lower bound is proved by constructing hard-instances of linear mixture MDPs following Zhou et al. (2021b,a); Zhang et al. (2021a).

**Theorem 5.3.** Let  $B > 1$ . Then for any algorithm, when  $K \geq \max\{3d^2, (d-1)/(192(B-1))\}$ , there exists a  $B$ -bounded linear mixture MDP satisfying Assumptions 3.1 and 3.2 such that its expected regret  $\mathbb{E}[\text{Regret}(K)]$  is lower bounded by  $d\sqrt{K}/(16\sqrt{3})$ .

**Remark 5.4.** When specialized to tabular MDPs where  $d = |\mathcal{S}^2|\mathcal{A}|$ , HF-UCRL-VTR<sup>+</sup> yields a horizon-free regret  $\tilde{O}(|\mathcal{S}^2|\mathcal{A}|\sqrt{K} + |\mathcal{S}^4|\mathcal{A}|^2)$ . Although the regret does not match the near-optimal result  $\tilde{O}(\sqrt{|\mathcal{S}||\mathcal{A}|K} + |\mathcal{S}^2|\mathcal{A}|)$  (Zhang et al., 2021b), it is not surprising since HF-UCRL-VTR<sup>+</sup> is designed for a more general MDP class. We leave the design of algorithms that achieve near-optimal regret for both linear mixture MDPs and tabular MDPs simultaneously as a future work.

We conduct some numerical experiments to suggest the validity of HF-UCRL-VTR<sup>+</sup> in Appendix A.

## 6 Conclusion

In this work, we propose a new weighted linear regression estimator that adapts *variance-uncertainty-aware* weights, which can be applied to both heterogeneous linear bandits and linear mixture MDPs. For heterogeneous linear bandits, our WeightedOFUL<sup>+</sup> algorithm achieves an  $\tilde{O}(d\sqrt{\sum_{k=1}^K \sigma_k^2} + d)$  regret in the first  $K$  rounds. For linear mixture MDPs, our HF-UCRL-VTR<sup>+</sup> algorithm achieves the near-optimal  $\tilde{O}(d\sqrt{K} + d^2)$  regret. Both of our algorithms are computationally efficient and yield the state-of-the-arts regret results.

## Acknowledgments and Disclosure of Funding

We thank the anonymous reviewers for their helpful comments. DZ and QG are partially supported by the National Science Foundation CAREER Award 1906169 and the Sloan Research Fellowship. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

## References

- ABBASI-YADKORI, Y., PÁL, D. and SZEPESVÁRI, C. (2011). Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*.
- AUER, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research* **3** 397–422.
- AYOUB, A., JIA, Z., SZEPESVARI, C., WANG, M. and YANG, L. F. (2020). Model-based reinforcement learning with value-targeted regression. *arXiv preprint arXiv:2006.01107*.
- AZUMA, K. (1967). Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal, Second Series* **19** 357–367.
- CAI, Q., YANG, Z., JIN, C. and WANG, Z. (2019). Provably efficient exploration in policy optimization. *arXiv preprint arXiv:1912.05830*.
- CHU, W., LI, L., REYZIN, L. and SCHAPIRE, R. (2011). Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*.
- DANI, V., HAYES, T. P. and KAKADE, S. M. (2008). Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*.
- DANN, C., JIANG, N., KRISHNAMURTHY, A., AGARWAL, A., LANGFORD, J. and SCHAPIRE, R. E. (2018). On oracle-efficient pac rl with rich observations. In *Advances in neural information processing systems*.
- DU, S. S., KAKADE, S. M., WANG, R. and YANG, L. F. (2019). Is a good representation sufficient for sample efficient reinforcement learning? In *International Conference on Learning Representations*.
- DZHAPARIDZE, K. and VAN ZANTEN, J. (2001). On bernstein-type inequalities for martingales. *Stochastic processes and their applications* **93** 109–117.
- FAN, X., GRAMA, I. and LIU, Q. (2017). Martingale inequalities of type dzhaparidze and van zanten. *Statistics* **51** 1200–1213.
- HE, J., ZHOU, D. and GU, Q. (2021). Logarithmic regret for reinforcement learning with linear function approximation. In *International Conference on Machine Learning*. PMLR.
- HE, J., ZHOU, D., ZHANG, T. and GU, Q. (2022). Nearly optimal algorithms for linear contextual bandits with adversarial corruptions. *arXiv preprint arXiv:2205.06811*.

- JIA, Z., YANG, L., SZEPESVARI, C. and WANG, M. (2020). Model-based reinforcement learning with value-targeted regression. In *LADC*.
- JIANG, N. and AGARWAL, A. (2018). Open problem: The dependence of sample complexity lower bounds on planning horizon. In *Conference On Learning Theory*. PMLR.
- JIANG, N., KRISHNAMURTHY, A., AGARWAL, A., LANGFORD, J. and SCHAPIRE, R. E. (2017). Contextual decision processes with low Bellman rank are PAC-learnable. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org.
- JIN, C., ALLEN-ZHU, Z., BUBECK, S. and JORDAN, M. I. (2018). Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*.
- JIN, C., YANG, Z., WANG, Z. and JORDAN, M. I. (2020). Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*.
- KIM, Y., YANG, I. and JUN, K.-S. (2021). Improved regret analysis for variance-adaptive linear bandits and horizon-free linear mixture mdps. *arXiv preprint arXiv:2111.03289* .
- KIRSCHNER, J. and KRAUSE, A. (2018). Information directed sampling and bandits with heteroscedastic noise. In *Conference On Learning Theory*.
- KRISHNAMURTHY, A., AGARWAL, A. and LANGFORD, J. (2016). Pac reinforcement learning with rich observations. *Advances in Neural Information Processing Systems* **29**.
- LATTIMORE, T., CRAMMER, K. and SZEPESVÁRI, C. (2015). Linear multi-resource allocation with semi-bandit feedback. In *Advances in Neural Information Processing Systems*.
- LATTIMORE, T. and HUTTER, M. (2012). PAC bounds for discounted MDPs. In *International Conference on Algorithmic Learning Theory*. Springer.
- LI, G., WEI, Y., CHI, Y., GU, Y. and CHEN, Y. (2020). Breaking the sample size barrier in model-based reinforcement learning with a generative model. *Advances in neural information processing systems* **33** 12861–12872.
- LI, L., CHU, W., LANGFORD, J. and SCHAPIRE, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*.
- LI, Y., WANG, R. and YANG, L. F. (2022). Settling the horizon-dependence of sample complexity in reinforcement learning. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE.
- LI, Y., WANG, Y. and ZHOU, Y. (2019a). Nearly minimax-optimal regret for linearly parameterized bandits. In *Conference on Learning Theory*.
- LI, Y., WANG, Y. and ZHOU, Y. (2019b). Tight regret bounds for infinite-armed linear contextual bandits. *arXiv preprint arXiv:1905.01435* .
- MODI, A., JIANG, N., TEWARI, A. and SINGH, S. (2020). Sample complexity of reinforcement learning using linearly combined model ensembles. In *International Conference on Artificial Intelligence and Statistics*. PMLR.
- PUTERMAN, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- REN, T., LI, J., DAI, B., DU, S. S. and SANGHAVI, S. (2021). Nearly horizon-free offline reinforcement learning. *Advances in neural information processing systems* **34**.
- STREHL, A. L. and LITTMAN, M. L. (2008). An analysis of model-based interval estimation for Markov decision processes. *Journal of Computer and System Sciences* **74** 1309–1331.
- SUN, W., JIANG, N., KRISHNAMURTHY, A., AGARWAL, A. and LANGFORD, J. (2019). Model-based RL in contextual decision processes: PAC bounds and exponential improvements over model-free approaches. In *Conference on Learning Theory*. PMLR.

- TARBOURIECH, J., ZHOU, R., DU, S. S., PIROTTA, M., VALKO, M. and LAZARIC, A. (2021). Stochastic shortest path: Minimax, parameter-free and towards horizon-free regret. *Advances in Neural Information Processing Systems* **34**.
- WANG, R., DU, S. S., YANG, L. F. and KAKADE, S. M. (2020). Is long horizon reinforcement learning more difficult than short horizon reinforcement learning? *arXiv preprint arXiv:2005.00527*.
- WANG, Y., WANG, R., DU, S. S. and KRISHNAMURTHY, A. (2019). Optimism in reinforcement learning with generalized linear function approximation. *arXiv preprint arXiv:1912.04136*.
- WEISZ, G., AMORTILA, P. and SZEPESVÁRI, C. (2020). Exponential lower bounds for planning in MDPs with linearly-realizable optimal action-value functions. *arXiv preprint arXiv:2010.01374*.
- YANG, L. and WANG, M. (2019a). Sample-optimal parametric Q-learning using linearly additive features. In *International Conference on Machine Learning*.
- YANG, L. F. and WANG, M. (2019b). Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. *arXiv preprint arXiv:1905.10389*.
- ZANETTE, A., BRANDFONBRENER, D., BRUNSKILL, E., PIROTTA, M. and LAZARIC, A. (2020a). Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*.
- ZANETTE, A., LAZARIC, A., KOCHENDERFER, M. and BRUNSKILL, E. (2020b). Learning near optimal policies with low inherent Bellman error. *arXiv preprint arXiv:2003.00153*.
- ZHANG, W., ZHOU, D. and GU, Q. (2021a). Reward-free model-based reinforcement learning with linear function approximation. *Advances in Neural Information Processing Systems* **34**.
- ZHANG, Z., DU, S. S. and JI, X. (2020). Nearly minimax optimal reward-free reinforcement learning. *arXiv preprint arXiv:2010.05901*.
- ZHANG, Z., JI, X. and DU, S. (2021b). Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. In *Conference on Learning Theory*. PMLR.
- ZHANG, Z., JI, X. and DU, S. S. (2022). Horizon-free reinforcement learning in polynomial time: the power of stationary policies. *arXiv preprint arXiv:2203.12922*.
- ZHANG, Z., YANG, J., JI, X. and DU, S. S. (2021c). Improved variance-aware confidence sets for linear bandits and linear mixture mdp. *Advances in Neural Information Processing Systems* **34**.
- ZHANG, Z., ZHOU, Y. and JI, X. (2021d). Model-free reinforcement learning: from clipped pseudo-regret to sample complexity. In *International Conference on Machine Learning*. PMLR.
- ZHOU, D., GU, Q. and SZEPESVARI, C. (2021a). Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*. PMLR.
- ZHOU, D., HE, J. and GU, Q. (2021b). Provably efficient reinforcement learning for discounted mdps with feature mapping. In *International Conference on Machine Learning*. PMLR.

## 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]** Our paper proposes a computational efficient and horizon-free RL algorithm for linear mixture MDPs.
  - (b) Did you describe the limitations of your work? **[Yes]** See Remark 4.2
  - (c) Did you discuss any potential negative societal impacts of your work? **[N/A]**
  - (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] See Assumptions 3.1 and 3.2.
  - (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix D and E.
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]

## A Experiments

In this section, we conduct some numerical experiments to validate our theory. The code and data for our experiments can be found on Github<sup>3</sup>.

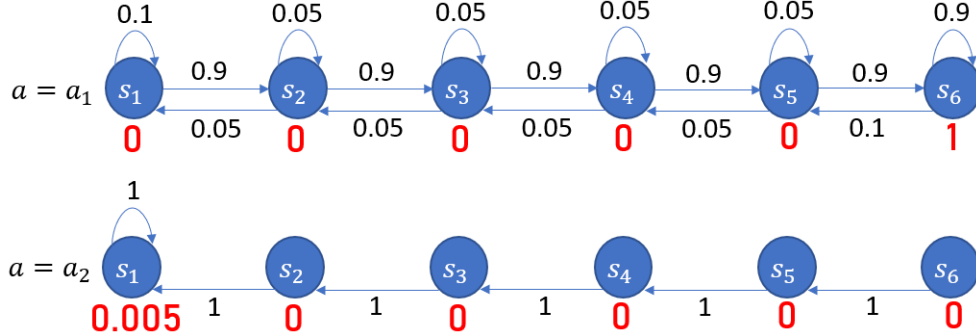


Figure 1: Transition dynamic and reward of RiverSwim. The black numbers denote the transition probability, the red numbers denote the reward. The first row represents the transition probability with action  $a_1$ , and the second row represents the transition probability with action  $a_2$ .

**Environment** We follow the RiverSwim environment introduced by Strehl and Littman (2008) and considered by Ayoub et al. (2020). The state space is  $\mathcal{S} = \{s_1, \dots, s_6\}$ , the action space is  $\mathcal{A} = \{a_1, a_2\}$ . The transition dynamic and reward are demonstrated by Figure 1. The length of the episode  $H = 100, 500$ . The initial state is  $s_1$ . The RiverSwim MDP is a linear mixture MDP with  $d = 3$  and a feature mapping  $\phi(s'|s, a) \in \mathbb{R}^3$ . Here

$$\begin{aligned} \phi(s'|s_1, a_1) &= \mathbb{1}(s' = s_1)(0, 1, 1) + \mathbb{1}(s' = s_2)(1, 0, 0) \\ \forall 2 \leq i \leq 5, \phi(s'|s_i, a_1) &= \mathbb{1}(s' = s_{i-1})(0, 0, 1) + \mathbb{1}(s' = s_i)(0, 1, 0) + \mathbb{1}(s' = s_{i+1})(1, 0, 0) \\ \phi(s'|s_6, a_1) &= \mathbb{1}(s' = s_5)(0, 1, 1) + \mathbb{1}(s' = s_6)(1, 0, 0) \\ \phi(s'|s_1, a_2) &= \mathbb{1}(s' = s_1)(1, 1, 1) \\ \forall 2 \leq i \leq 6, \phi(s'|s_i, a_2) &= \mathbb{1}(s' = s_{i-1})(1, 1, 1) \end{aligned}$$

The transition dynamic  $\mathbb{P} = \langle \phi/3, \theta \rangle$  with  $\theta = (2.7, 0.15, 0.15)$ .

**Benchmarks** We compare the following algorithms.

- *HF-UCRL-VTR<sup>+</sup>* with  $M = 4$  (level parameter of *HOME*)
- *HF-UCRL-VTR<sup>+</sup>* with  $M = 2$  (level parameter of *HOME*)
- *UCRL-VTR<sup>+</sup>* (Zhou et al., 2021a)
- *UCRL-VTR* (Jia et al., 2020; Ayoub et al., 2020)
- *Q-learning* (Jin et al., 2018)
- *Random*: Randomly select actions

The hyperparameters of each algorithm are grid searched for the best performance. For *HF-UCRL-VTR<sup>+</sup>* with  $M = 4$ , *HF-UCRL-VTR<sup>+</sup>* with  $M = 2$ , *UCRL-VTR<sup>+</sup>* and *UCRL-VTR*, we set the regularization parameter  $\lambda = 0.001$ , and the confidence radius parameter  $\beta_k = 1$ . For *HF-UCRL-VTR<sup>+</sup>* with  $M = 4$ , *HF-UCRL-VTR<sup>+</sup>* with  $M = 2$ , we set the variance parameters  $\alpha = 0.01, \gamma = 0.5$ . For *Q-learning*, we set its confidence radius parameter  $\beta_k = 0.005$ .

**Results** We compare each algorithm with their average accumulated reward, which is the summation of received rewards divided by the number of current episode. We run each algorithm for 500 episodes. Each algorithm has 10 independent runs. We plot their mean values and error bar ([mean-stand error, mean+stand error]) w.r.t. number of episodes. The results for  $H = 100$  and  $H = 500$  are plotted in Figure 2a and 2b.

<sup>3</sup><https://github.com/uclaml/HF-UCRL-VTR>

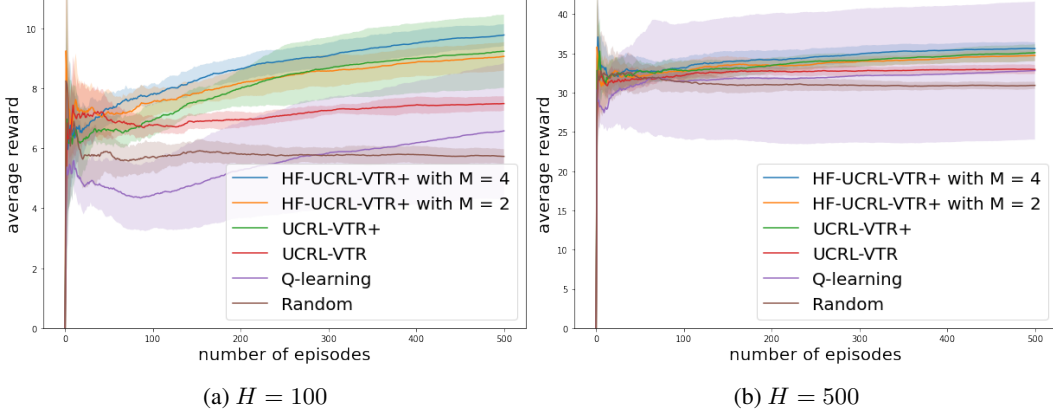


Figure 2: Average accumulated reward of benchmarks.

From Figure 2a and 2b we can draw the following conclusions: (1) HF-UCRL-VTR<sup>+</sup> with  $M = 4$  outperforms HF-UCRL-VTR<sup>+</sup> with  $M = 2$ , which validates the advantage of the HOME estimator. (2) HF-UCRL-VTR<sup>+</sup> and UCRL-VTR<sup>+</sup> outperforms UCRL-VTR, which suggests that the Bernstein-type bonus is useful. (3) All UCRL-VTR-based algorithms outperform the tabular Q-learning. The average reward of UCRL-VTR-based algorithms is generally higher than that of Q-learning, and the variance of UCRL-VTR-based algorithms is also smaller than that of Q-learning. This suggests the validity of the value-targeted regression, as well as the use of the feature  $\phi$ .

## B Computational complexity of HF-UCRL-VTR<sup>+</sup>

At each episode  $k$  and each stage  $h$ , HF-UCRL-VTR<sup>+</sup> needs to compute  $\{\phi_{k,h,m}\}_{m \in \overline{[M]}}$  and  $\{\bar{\sigma}_{k,h,m}\}_{m \in \overline{[M]}}$ , and update  $\{\tilde{\Sigma}_{k,h+1,m}\}_{m \in \overline{[M]}}$ . According to Algorithm 3,  $\{\bar{\sigma}_{k,h,m}\}_{m \in \overline{[M]}}$  can be computed in  $O(Md^2)$  time as they only require the computation of the inner-product between vectors and the inner-product between an inversion of matrix and a vector. For  $\{\phi_{k,h,m}\}_{m \in \overline{[M]}}$ , they can be computed within  $O(\mathcal{O} \cdot M)$  time. Finally, to selection actions based on  $\pi_h^k$ , HF-UCRL-VTR<sup>+</sup> needs to compute  $|\mathcal{A}|$  number of action-value function  $Q_{k,h}$ , while each of them needs to compute the  $\phi_V(\cdot, \cdot)$  within  $\mathcal{O}$  time and the inner-product between an inversion of a matrix and a vector by  $O(d^2)$  time. Therefore, the total amount of time HF-UCRL-VTR<sup>+</sup> takes is  $O(KHMd^2 + KHAM + |\mathcal{A}|KHO + |\mathcal{A}|KHd^2)$ .

**Remark B.1.** Compared with HF-UCRL-VTR<sup>+</sup>, VOFUL/VOFUL2 (Zhang et al., 2021c; Kim et al., 2021) need to compute the upper bounds of moments as the maximum of the quadratic function  $[\langle \phi_{k,h,m+1}, \theta \rangle]_{[0,1]} - [\langle \phi_{k,h,m}, \theta \rangle]_{[0,1]}^2$  over a series of implicit confidence sets, which are not implementable.

## C Proof of the Bernstein-type concentration inequality (Lemma 4.3)

We prove Lemma 4.3 here. First, we provide some new concentration inequalities here, which will be used in the following proof.

**Lemma C.1.** Let  $M, v > 0$  be constants. Let  $\{x_i\}_{i=1}^n$  be a stochastic process,  $\mathcal{G}_i = \sigma(x_1, \dots, x_i)$  be the  $\sigma$ -algebra of  $x_1, \dots, x_i$ . Suppose  $\mathbb{E}[x_i | \mathcal{G}_{i-1}] = 0$ ,  $|x_i| \leq M$  and  $\sum_{i=1}^n \mathbb{E}[x_i^2 | \mathcal{G}_{i-1}] \leq v^2$  almost surely. Then for any  $\delta, \epsilon > 0$ , let  $\iota = \sqrt{\log((\log(M/\epsilon) + 2)/\delta)}$ , we have

$$\mathbb{P}\left(\sum_{i=1}^n x_i \leq 3\iota v + 9\iota^2 \max_{i \in [n]} |x_i| + 4\iota^2 \epsilon\right) > 1 - \delta.$$

*Proof.* For simplicity, let

$$S_n = \sum_{i=1}^n x_i, \bar{S}_n = \sqrt{\sum_{i=1}^n \mathbb{E}[x_i^2 | \mathcal{G}_{i-1}]}, \bar{x}_n = \max_{i \in [n]} |x_i|.$$

We have  $\bar{S}_n < v^2, \bar{x}_n \leq M$ . Set threshold  $\epsilon > 0$ . Set the following sets

$$\Gamma_0 = [0, \epsilon), \Gamma_i = [\epsilon \cdot e^{i-1}, \epsilon \cdot e^i), i = 1, \dots, I = \lfloor \log(M/\epsilon) \rfloor + 1.$$

Then there exist  $0 \leq i \leq I$  such that  $\bar{x}_n \in \Gamma_i$ . Let  $\iota = 3\sqrt{\log(1/\delta)}$  and  $\zeta = \iota^2$ . Taking union bound, we have

$$\mathbb{P}\left(S_n > \iota v + \zeta \bar{x}_n + \eta\right) \leq \sum_{i=0}^I \mathbb{P}\left(S_n > \iota v + \zeta \bar{x}_n + \eta, \bar{S}_n \leq v, \bar{x}_n \in \Gamma_i\right),$$

For each  $0 \leq i \leq I$ , we have

$$\begin{aligned} & \mathbb{P}\left(S_n > \iota v + \zeta \bar{x}_n + \eta, \bar{S}_n < v, \bar{x}_n \in \Gamma_i\right) \\ & \leq \mathbb{P}\left(S_n > \max\{\iota v, \zeta e^{i-1} \epsilon\}, \bar{S}_n < v, \bar{x}_n < e^i \epsilon\right) \\ & \leq \mathbb{P}\left(S_n > \max\{\iota v, \zeta e^{i-1} \epsilon\}, \bar{S}_n^2 + \sum_{k=1}^n x_k^2 \mathbf{1}\{|x_k| > e^i \epsilon\} < v^2\right) \\ & \leq \exp\left(\frac{-\max\{\iota v, \zeta e^{i-1} \epsilon\}^2}{2(v^2 + e^i \epsilon \cdot \max\{\iota v, \zeta e^{i-1} \epsilon\}/3)}\right) \\ & \leq \exp\left(\frac{-1}{2(1/\iota^2 + 1/\zeta)}\right) \\ & \leq \delta, \end{aligned} \tag{C.1}$$

where the first inequality holds since for any  $0 \leq i \leq I$  and  $\bar{x}_n \in \Gamma_i$ , we have

$$\iota v + \zeta \bar{x}_n + \eta > \iota v + \zeta \bar{x}_n + \zeta \epsilon / e \geq \max\{\iota v, \zeta e^{i-1} \epsilon\},$$

the second one holds since  $\bar{x}_n < e^i \epsilon$  implies  $\mathbf{1}\{|x_k| > e^j \epsilon\} = 0$  for each  $k \in [n]$ , the third one holds due to Lemma F.3, the fourth one holds by basic calculation, the last one holds since  $\iota = 3\sqrt{\log(1/\delta)}$ . Therefore, taking summation for (C.1) over all  $0 \leq i \leq I$ , our statement holds.  $\square$

**Lemma C.2.** Let  $\{x_i \geq 0\}_{i \geq 1}$  be a stochastic process,  $\{\mathcal{G}_i\}_{i \geq 1}$  be a filtration satisfying  $x_i$  is  $\mathcal{G}_i$ -measurable. Let  $M, v > 0$ , and  $x_i \leq M, \sum_{i=1}^n \mathbb{E}[x_i | \mathcal{G}_{i-1}] \leq v$  almost surely. Then for any  $\delta, \epsilon > 0$ , let  $\iota = \log(4(\log(M/\epsilon) + 2)/\delta)$ , we have

$$\mathbb{P}\left(\sum_{i=1}^n x_i \leq 4\iota v + 11\iota \max_{1 \leq i \leq n} x_i + 4\iota \epsilon\right) > 1 - \delta.$$

*Proof.* For simplicity, let

$$S_n = \sum_{i=1}^n x_i, \bar{S}_n = \sum_{i=1}^n \mathbb{E}[x_i | \mathcal{G}_{i-1}], \bar{x}_n = \max_{i \in [n]} |x_i|.$$

We have  $\bar{S}_n \leq v$  and  $\bar{x}_n \leq M$ . Set threshold  $\epsilon > 0$ . Set the following sets

$$\Gamma_0 = [0, \epsilon), \Gamma_i = [\epsilon \cdot e^{i-1}, \epsilon \cdot e^i), i = 1, \dots, I = \lfloor \log(M/\epsilon) \rfloor + 1.$$

Then there exist  $0 \leq i \leq I$  such that  $\bar{x}_n \in \Gamma_i$ . Let  $\iota = 4 \log(4/\delta), \zeta = 11 \log(4/\delta), \eta = 4 \log(4/\delta) \epsilon$ . By union bound we have

$$\mathbb{P}\left(S_n > \iota v + \zeta \bar{x}_n + \eta\right) \leq \sum_{i=0}^I \mathbb{P}\left(S_n > \iota v + \zeta \bar{x}_n + \eta, \bar{S}_n \leq v, \bar{x}_n \in \Gamma_i\right),$$



For each  $0 \leq i \leq I$ , we have

$$\begin{aligned}
& \mathbb{P}\left(S_n > \iota v + \zeta \bar{x}_n + \eta, \bar{S}_n \leq v, \bar{x}_n \in \Gamma_i\right) \\
& \leq \mathbb{P}\left(S_n > \iota v + \zeta \bar{x}_n + \eta, \bar{S}_n < \log(4/\delta)(v + e^i \epsilon), \bar{x}_n < e^i \epsilon\right) \\
& = \mathbb{P}\left(\sum_{k=1}^n x_k \mathbb{1}\{x_k \leq e^i \epsilon\} > \iota v + \zeta \bar{x}_n + \eta, \bar{S}_n < \log(4/\delta)(v + e^i \epsilon), \bar{x}_n < e^i \epsilon\right) \\
& \leq \mathbb{P}\left(\sum_{k=1}^n x_k \mathbb{1}\{x_k \leq e^i \epsilon\} \geq 4 \log(4/\delta)(v + e^i \epsilon), \bar{S}_n \leq \log(4/\delta)(v + e^i \epsilon)\right) \\
& \leq \mathbb{P}\left(\sum_{k=1}^n x_k \mathbb{1}\{x_k \leq e^i \epsilon\} \geq 4 \log(4/\delta)(v + e^i \epsilon), \sum_{k=1}^n \mathbb{E}[x_k \mathbb{1}\{x_k \leq e^i \epsilon\} | \mathcal{G}_{k-1}] \leq \log(4/\delta)(v + e^i \epsilon)\right) \\
& = \mathbb{P}\left(\sum_{k=1}^n \frac{x_k \mathbb{1}\{x_k \leq e^i \epsilon\}}{e^i \epsilon} \geq 4 \log(4/\delta) \frac{v + e^i \epsilon}{e^i \epsilon}, \sum_{k=1}^n \mathbb{E}\left[\frac{x_k \mathbb{1}\{x_k \leq e^i \epsilon\}}{e^i \epsilon} \middle| \mathcal{G}_{k-1}\right] \leq \log(4/\delta) \frac{v + e^i \epsilon}{e^i \epsilon}\right) \\
& \leq \delta. \tag{C.2}
\end{aligned}$$

Here, the second inequality holds since for any  $0 \leq i \leq I$  and  $\bar{x}_n \in \Gamma_i$ , we have

$$\iota v + \zeta \bar{x}_n + \eta = 4 \log(4/\delta)v + 11 \log(4/\delta)\bar{x}_n + 4 \log(4/\delta)\epsilon \geq 4 \log(4/\delta)(v + e^i \epsilon).$$

The third inequality holds since  $x_k \geq x_k \mathbb{1}\{x_k \leq e^i \epsilon\}$ . The last inequality holds by Lemma F.4 with  $c = (v + e^i \epsilon)/(e^i \epsilon) > 1$ . Therefore, taking summation for (C.2) over all  $0 \leq i \leq I$ , our statement holds.  $\square$

Now we begin to prove Lemma 4.3. We first define the following notations:

$$\mathbf{d}_1 = \mathbf{0}, \mathbf{d}_k = \sum_{i=1}^k \mathbf{x}_i \eta_i, w_k = \|\mathbf{x}_k\|_{\mathbf{Z}_{k-1}^{-1}}, \mathcal{E}_k = \{1 \leq s \leq k, \|\mathbf{d}_s\|_{\mathbf{Z}_s^{-1}} \leq \beta_s\}, \tag{C.3}$$

where  $k \geq 1$  and we define  $\beta_1 = 0$ . Recalling that  $x_k$  is  $\mathcal{G}_k$ -measurable and  $\eta_k$  is  $\mathcal{G}_{k+1}$ -measurable, we find that  $\mathbf{d}_k, \mathbf{Z}_k, \mathcal{E}_k$  and  $w_k$  is  $\mathcal{G}_k$ -measurable. Recall that  $\beta_t$  is defined as follows:

$$\begin{aligned}
\beta_t &= 12\sqrt{\sigma^2 d \log(1 + tL^2/(d\lambda)) \log(32(\log(R/\epsilon) + 1)t^2/\delta)} \\
&\quad + 24 \log(32(\log(R/\epsilon) + 1)t^2/\delta) \max_{1 \leq i \leq t} \{\eta_i | \min\{1, w_i\}\} + 6 \log(32(\log(R/\epsilon) + 1)t^2/\delta)\epsilon.
\end{aligned}$$

The proof of Lemma 4.3 follows the proof strategy in Dani et al. (2008); Zhou et al. (2021a). Briefly speaking, we can bound the self-normalized martingale  $\|\sum_{i=1}^k \mathbf{x}_i \eta_i\|_{\mathbf{Z}_k^{-1}}^2$  by the following two terms,

$$\sum_{i=1}^k \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathbb{1}\{\mathcal{E}_{i-1}\}, \text{ and } \sum_{i=1}^k \frac{\eta_i^2 w_i^2}{1 + w_i^2}.$$

Next two lemmas bound these terms by  $\beta_t$  separately. The main difference between the following lemmas and their counterparts in Dani et al. (2008); Zhou et al. (2021a) is that our  $\beta_t$  is also a *random variable*, which requires us to use some advanced concentration inequalities (Lemmas C.1 and C.2) rather than vanilla ones.

**Lemma C.3.** Let  $\mathbf{d}_i, w_i, \mathcal{E}_i$  be defined in (C.3). Then, with probability at least  $1 - \delta/2$ , simultaneously for all  $k \geq 1$  it holds that

$$\sum_{i=1}^k \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathbb{1}\{\mathcal{E}_{i-1}\} \leq 3\beta_k^2/4.$$

*Proof.* We have

$$\left| \frac{2\mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1 + w_i^2} \mathbb{1}\{\mathcal{E}_{i-1}\} \right| \leq \frac{2\|\mathbf{x}_i\|_{\mathbf{Z}_{i-1}^{-1}} [\|\mathbf{d}_{i-1}\|_{\mathbf{Z}_{i-1}^{-1}} \mathbb{1}\{\mathcal{E}_{i-1}\}]}{1 + w_i^2} \leq \frac{2w_i \beta_{i-1}}{1 + w_i^2} \leq \min\{1, 2w_i\} \beta_{i-1}, \tag{C.4}$$

where the first inequality holds due to Cauchy-Schwarz inequality, the second inequality holds due to the definition of  $\mathcal{E}_{i-1}$ , the last inequality holds by algebra. Fix  $t > 0$ . For simplicity, let  $\ell_i$  denote

$$\ell_i = \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{\beta_{i-1}(1+w_i^2)} \mathbb{1}\{\mathcal{E}_{i-1}\}. \quad (\text{C.5})$$

We are preparing to apply Lemma C.1 to  $\{\ell_i\}_i$  and  $\{\mathcal{G}_i\}_i$ . First note that  $\mathbb{E}[\ell_i|\mathcal{G}_i] = 0$ . Meanwhile, by (C.4) we have  $|\ell_i| \leq R$  and

$$|\ell_i| \leq 2|\eta_i| \min\{1, w_i\}. \quad (\text{C.6})$$

We also have

$$\begin{aligned} \sum_{i=1}^t \mathbb{E}[\ell_i^2|\mathcal{G}_i] &\leq \sigma^2 \sum_{i=1}^t \left( \frac{2\mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{1+w_i^2} \mathbb{1}\{\mathcal{E}_{i-1}\} / \beta_{i-1} \right)^2 \\ &\leq 4\sigma^2 \sum_{i=1}^t [\min\{1, w_i\}]^2 \\ &\leq 4\sigma^2 \sum_{i=1}^t \min\{1, w_i^2\} \\ &\leq 8\sigma^2 d \log(1 + tL^2/(d\lambda)), \end{aligned} \quad (\text{C.7})$$

where the first inequality holds since  $\mathbb{E}[\eta_i^2|\mathcal{G}_i] \leq \sigma^2$ , the second inequality holds due to (C.4), the third inequality holds again since  $\{\beta_i\}_i$  is increasing, the last inequality holds due to Lemma F.5. Let  $\iota_t = \sqrt{\log(4(\log(R/\epsilon) + 2)t^2/\delta)}$ . Therefore, by (C.6) and (C.7), using Lemma C.1, we know that with probability at least  $1 - \delta/(4t^2)$ , we have

$$\sum_{i=1}^t \ell_i \leq 3\iota_t \sqrt{8\sigma^2 d \log(1 + tL^2/(d\lambda))} + 18\iota_t^2 \max_{1 \leq i \leq t} \{|\eta_i| \min\{1, w_i\}\} + 4\iota_t^2 \epsilon \leq \frac{3}{4}\beta_t,$$

where the last inequality holds due to the definition of  $\beta_t$ . Therefore, using the fact  $\beta_{i-1} \leq \beta_t$ , we have

$$\sum_{i=1}^t \frac{2\eta_i \mathbf{x}_i^\top \mathbf{Z}_{i-1}^{-1} \mathbf{d}_{i-1}}{(1+w_i^2)} \mathbb{1}\{\mathcal{E}_{i-1}\} \leq \beta_t \sum_{i=1}^t \ell_i \leq 3\beta_t^2/4. \quad (\text{C.8})$$

Taking union bound for (C.8) from  $t = 1$  to  $\infty$  and using the fact that  $\sum_{t=1}^{\infty} t^{-2} < 2$  finishes the proof.  $\square$

We also need the following lemma.

**Lemma C.4.** Let  $w_i$  be defined in (C.3). Then, with probability at least  $1 - \delta/2$ , simultaneously for all  $k \geq 1$  it holds that

$$\sum_{i=1}^k \frac{\eta_i^2 w_i^2}{1+w_i^2} \leq \beta_k^2/4.$$

*Proof.* For simplicity, let

$$\ell_i = \frac{\eta_i^2 w_i^2}{1+w_i^2}.$$

Fix  $t$ . Then by the fact that  $\mathbb{E}[\eta_i^2|\mathcal{G}_i] \leq \sigma^2$  and Lemma F.5, we have

$$\sum_{i=1}^t \mathbb{E}[\ell_i|\mathcal{G}_i] = \sum_{i=1}^t \mathbb{E} \left[ \frac{\eta_i^2 w_i^2}{1+w_i^2} \middle| \mathcal{G}_i \right] \leq \sigma^2 \sum_{i=1}^t \min\{1, w_i^2\} \leq 2\sigma^2 d \log(1 + tL^2/(d\lambda)). \quad (\text{C.9})$$

Furthermore, we have  $|\ell_i| \leq R^2$  and

$$\ell_i \leq |\eta_i|^2 \min\{1, w_i\}^2.$$

Let  $\iota_t = \log(32(\log(R/\epsilon) + 1)t^2/\delta)$ . Then by Lemma C.2, we have with probability at least  $1 - \delta/(4t^2)$ ,

$$\sum_{i=1}^t \ell_i \leq 8\iota_t \sigma^2 d \log(1 + tL^2/(d\lambda)) + 11\iota_t \max_{1 \leq i \leq t} \{|\eta_i| \min\{1, w_i\}\}^2 + 4\iota_t \epsilon^2 \leq \beta_t^2/4. \quad (\text{C.10})$$

where the last inequality holds due to the definition of  $\beta_t$ . Taking union bound for (C.10) from  $t = 1$  to  $\infty$  and using the fact that  $\sum_{t=1}^{\infty} t^{-2} < 2$  finishes the proof.  $\square$

With above two lemmas, we are ready to prove Lemma 4.3.

*Proof of Lemma 4.3.* The proof is nearly the same as the proof of Theorem 2 in Zhou et al. (2021a), only to replace Lemmas 13 and 14 in Zhou et al. (2021a) with Lemma C.3 and C.4.  $\square$

## D Proof of main results in Section 4

In this section, we prove Lemma 4.4 and Theorem 4.1.

### D.1 Proof of Lemma 4.4

We restate Lemma 4.4 here with a different set of notations, as it will be used in the later proof for RL as well.

**Lemma D.1** (Restatement of Lemma 4.4). Let  $\{\sigma_k, \beta_k\}_{k \geq 1}$  be a sequence of non-negative numbers,  $\alpha, \gamma > 0$ ,  $\{\mathbf{x}_k\}_{k \geq 1} \subset \mathbb{R}^d$  and  $\|\mathbf{x}_k\|_2 \leq L$ . Let  $\{\mathbf{Z}_k\}_{k \geq 1}$  and  $\{\bar{\sigma}_k\}_{k \geq 1}$  be recursively defined as follows:  $\mathbf{Z}_1 = \lambda \mathbf{I}$ ,

$$\forall k \geq 1, \bar{\sigma}_k = \max\{\sigma_k, \alpha, \gamma \|\mathbf{x}_k\|_{\mathbf{Z}_k^{-1}}^{1/2}\}, \mathbf{Z}_{k+1} = \mathbf{Z}_k + \mathbf{x}_k \mathbf{x}_k^\top / \bar{\sigma}_k^2.$$

Let  $\iota = \log(1 + KL^2/(d\lambda\alpha^2))$ . Then we have

$$\sum_{k=1}^K \min\left\{1, \beta_k \|\mathbf{x}_k\|_{\mathbf{Z}_k^{-1}}\right\} \leq 2d\iota + 2 \max_{k \in [K]} \beta_k \gamma^2 d\iota + 2\sqrt{d\iota} \sqrt{\sum_{k=1}^K \beta_k^2 (\sigma_k^2 + \alpha^2)}.$$

*Proof of Lemma D.1.* We decompose the set  $[K]$  into a union of two disjoint subsets  $[K] = \mathcal{I}_1 \cup \mathcal{I}_2$ .

$$\mathcal{I}_1 = \left\{k \in [K] : \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}} \geq 1\right\}, \mathcal{I}_2 = [K] \setminus \mathcal{I}_1. \quad (\text{D.1})$$

Then the following upper bound of  $|\mathcal{I}_1|$  holds:

$$|\mathcal{I}_1| = \sum_{k \in \mathcal{I}_1} \min\left\{1, \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}^2\right\} \leq \sum_{k=1}^K \min\left\{1, \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}^2\right\} \leq 2d\iota, \quad (\text{D.2})$$

where the first inequality holds since  $\|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}} \geq 1$  for  $k \in \mathcal{I}_1$ , the third inequality holds due to Lemma F.5 together with the fact  $\|\mathbf{x}_k / \bar{\sigma}_k\|_2 \leq L/\alpha$ . Therefore, we have

$$\begin{aligned} & \sum_{k \in [K]} \min\left\{1, \beta_k \|\mathbf{x}_k\|_{\mathbf{Z}_k^{-1}}\right\} \\ &= \sum_{k \in \mathcal{I}_1} \min\left\{1, \bar{\sigma}_k \beta_k \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}\right\} + \sum_{k \in \mathcal{I}_2} \min\left\{1, \bar{\sigma}_k \beta_k \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}\right\} \\ &\leq \left[\sum_{k \in \mathcal{I}_1} 1\right] + \sum_{k \in \mathcal{I}_2} \beta_k \bar{\sigma}_k \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}} \\ &\leq 2d\iota + \sum_{k \in \mathcal{I}_2} \beta_k \bar{\sigma}_k \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}, \end{aligned} \quad (\text{D.3})$$

where the first inequality holds since  $\min\{1, x\} \leq 1$  and also  $\min\{1, x\} \leq x$ , the second inequality holds due to (D.2). Next we further bound the second summation term in (D.3). We decompose  $\mathcal{I}_2 = \mathcal{J}_1 \cup \mathcal{J}_2$ , where

$$\mathcal{J}_1 = \left\{ k \in \mathcal{I}_2 : \bar{\sigma}_k = \sigma_k \cup \bar{\sigma}_k = \alpha \right\}, \quad \mathcal{J}_2 = \left\{ k \in \mathcal{I}_2 : \bar{\sigma}_k = \gamma \sqrt{\|\mathbf{x}_k\|_{\mathbf{Z}_k^{-1}}} \right\}.$$

For the summation over  $\mathcal{J}_1$ , we have

$$\begin{aligned} \sum_{k \in \mathcal{J}_1} \beta_k \bar{\sigma}_k \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}} &\leq \sum_{k \in \mathcal{J}_1} \beta_k (\sigma_k + \alpha) \min\{1, \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}\} \\ &\leq \sum_{k=1}^K \beta_k (\sigma_k + \alpha) \min\{1, \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}\} \\ &\leq \sqrt{2 \sum_{k=1}^K (\sigma_k^2 + \alpha^2) \beta_k^2} \sqrt{\sum_{k=1}^K \min\{1, \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}\}^2} \\ &\leq 2 \sqrt{\sum_{k=1}^K \beta_k^2 (\sigma_k^2 + \alpha^2)} \sqrt{d}, \end{aligned} \quad (\text{D.4})$$

where the first inequality holds since  $\bar{\sigma}_k \leq \sigma_k + \alpha$  for  $k \in \mathcal{J}_1$  and  $\|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}} \leq 1$  since  $k \in \mathcal{J}_1 \subseteq \mathcal{I}_2$ , the third one holds due to Cauchy-Schwarz inequality and the last one holds due to Lemma F.5. Next we bound the summation over  $\mathcal{J}_2$ . First, for  $k \in \mathcal{J}_2$ , we have the following equation:

$$\bar{\sigma}_k = \gamma^2 \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}.$$

Then we can bound the summation over  $\mathcal{J}_2$  as follows:

$$\sum_{k \in \mathcal{J}_2} \beta_k \bar{\sigma}_k \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}} = \gamma^2 \cdot \sum_{k \in \mathcal{J}_2} \beta_k \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}^2 = \gamma^2 \cdot \sum_{k=1}^K \beta_k \min\{1, \|\mathbf{x}_k / \bar{\sigma}_k\|_{\mathbf{Z}_k^{-1}}^2\} \leq 2 \max_{k \in [K]} \beta_k \gamma^2 d, \quad (\text{D.5})$$

where the inequality holds due to Lemma F.5. Substituting (D.4) and (D.5) into (D.3) completes the proof.  $\square$

## D.2 Proof of Theorem 4.1

*Proof of Theorem 4.1.* By the assumption on  $\epsilon_k$ , we know that

$$\begin{aligned} |\epsilon_k / \bar{\sigma}_k| &\leq R / \alpha, \\ |\epsilon_k / \bar{\sigma}_k| \cdot \min\{1, \|\mathbf{a}_k / \bar{\sigma}_k\|_{\hat{\Sigma}_k^{-1}}\} &\leq R \|\mathbf{a}_k\|_{\hat{\Sigma}_k^{-1}} / \bar{\sigma}_k^2 \leq R / \gamma^2, \\ \mathbb{E}[\epsilon_k | \mathbf{a}_{1:k}, \epsilon_{1:k-1}] &= 0, \quad \mathbb{E}[(\epsilon_k / \bar{\sigma}_k)^2 | \mathbf{a}_{1:k}, \epsilon_{1:k-1}] \leq 1, \quad \|\mathbf{a}_k / \bar{\sigma}_k\|_2 \leq A / \alpha, \end{aligned}$$

Therefore, setting  $\mathcal{G}_k = \sigma(\mathbf{a}_{1:k}, \epsilon_{1:k-1})$ , and using that  $\sigma_k$  is  $\mathcal{G}_k$ -measurable, applying Theorem 4.3 to  $(\mathbf{x}_k, \eta_k) = (\mathbf{a}_k / \bar{\sigma}_k, \epsilon_k / \bar{\sigma}_k)$  with  $\epsilon = R / \gamma^2$ , we get that with probability at least  $1 - \delta$ , for all  $k \geq 1$ ,

$$\begin{aligned} \|\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}^*\|_{\hat{\Sigma}_k} &\leq 12 \sqrt{d \log(1 + kA^2 / (\alpha^2 d \lambda)) \log(32(\log(\gamma^2 / \alpha) + 1)k^2 / \delta)} \\ &\quad + 30 \log(32(\log(\gamma^2 / \alpha) + 1)k^2 / \delta) R / \gamma^2 + \sqrt{\lambda} B \\ &:= \hat{\beta}_k, \end{aligned} \quad (\text{D.6})$$

Let  $\mathcal{E}_{D.6}$  denote the event such that (D.6) happens for all  $k \geq 1$ . Then, on the event  $\mathcal{E}_{D.6}$ , we have

$$\langle \mathbf{a}_k^*, \boldsymbol{\theta}^* \rangle \leq \langle \mathbf{a}_k^*, \hat{\boldsymbol{\theta}}_k \rangle + \|\mathbf{a}_k^*\|_{\hat{\Sigma}_k^{-1}} \|\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}^*\|_{\hat{\Sigma}_k} \leq \langle \mathbf{a}_k^*, \hat{\boldsymbol{\theta}}_k \rangle + \hat{\beta}_k \|\mathbf{a}_k^*\|_{\hat{\Sigma}_k^{-1}} \leq \langle \mathbf{a}_k^*, \hat{\boldsymbol{\theta}}_k \rangle + \hat{\beta}_k \|\mathbf{a}_k\|_{\hat{\Sigma}_k^{-1}}, \quad (\text{D.7})$$

where the last inequality holds due to the selection of  $\mathbf{a}_k$ . By (D.7), we have

$$\langle \mathbf{a}_k^*, \boldsymbol{\theta}^* \rangle - \langle \mathbf{a}_k, \boldsymbol{\theta}^* \rangle \leq \|\mathbf{a}_k\|_{\hat{\Sigma}_k^{-1}} (\hat{\beta}_k + \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_k\|_{\hat{\Sigma}_k}) \leq 2\hat{\beta}_k \|\mathbf{a}_k\|_{\hat{\Sigma}_k^{-1}}, \quad (\text{D.8})$$

where the first inequality holds due to Cauchy-Schwarz inequality, the second one holds due to event  $\mathcal{E}_{D,6}$ . Meanwhile, we have  $0 \leq \langle \mathbf{a}_k^*, \boldsymbol{\theta}^* \rangle - \langle \mathbf{a}_k, \boldsymbol{\theta}^* \rangle \leq 2$ . Thus, substituting (D.8) into (D.7) and summing up (D.7) for  $k \in [K]$ , we have

$$\text{Regret}(K) = \sum_{k=1}^K [\langle \mathbf{a}_k^*, \boldsymbol{\theta}^* \rangle - \langle \mathbf{a}_k, \boldsymbol{\theta}^* \rangle] \leq 2 \sum_{k=1}^K \min \left\{ 1, \widehat{\beta}_K \|\mathbf{a}_k\|_{\widehat{\Sigma}_k^{-1}} \right\}. \quad (\text{D.9})$$

Finally, to bound (D.9), by Lemma D.1 we have

$$\text{Regret}(K)/2 \leq 2d\iota + 2d\gamma^2 \widehat{\beta}_K \iota + \widehat{\beta}_K \sqrt{2 \sum_{k=1}^K \sigma_k^2 + 2K\alpha^2 \sqrt{2d\iota}},$$

which completes our proof.  $\square$

## E Proof of main results in Section 5

For  $k \in [K]$ ,  $h \in [H]$ , let  $\mathcal{F}_{k,h}$  be the  $\sigma$ -algebra generated by the random variables representing the state-action pairs up to and including those that appear stage  $h$  of episode  $k$ . That is,  $\mathcal{F}_{k,h}$  is generated by

$$\begin{aligned} & s_1^1, a_1^1, \dots, s_h^1, a_h^1, \dots, s_H^1, a_H^1, \\ & s_1^2, a_1^2, \dots, s_h^2, a_h^2, \dots, s_H^2, a_H^2, \\ & \vdots \\ & s_1^k, a_1^k, \dots, s_h^k, a_h^k. \end{aligned}$$

Note that, by construction,

$$\bar{V}_{k,h} V_{k,h+1}^{2^m}(s_h^k, a_h^k), E_{k,h,m}, \bar{\sigma}_{k,h,m}, \tilde{\Sigma}_{k,h+1,m}, \phi_{k,h,m}$$

are  $\mathcal{F}_{k,h}$ -measurable, and  $Q_{k,h}, V_{k,h}, \pi_h^k, \widehat{\Sigma}_{k,m}, \widehat{\mathbf{b}}_{k,m}, \widehat{\boldsymbol{\theta}}_{k,m}$  are  $\mathcal{F}_{k-1,H}$  measurable.

First we propose a lemma that suggests that the vector  $\boldsymbol{\theta}^*$  lies in a series of confidence sets, which further implies that the difference between the estimated high-order moment and true moment can be bounded by error terms  $E_{k,h,m}$ .

**Lemma E.1.** Set  $\{\widehat{\beta}_k\}_{k \geq 1}$  as (5.3), then, with probability at least  $1 - M\delta$ , we have for any  $k \in [K]$ ,  $h \in [H]$ ,  $m \in \overline{[M]}$ ,

$$\left\| \widehat{\Sigma}_{k,m}^{1/2} (\widehat{\boldsymbol{\theta}}_{k,m} - \boldsymbol{\theta}^*) \right\|_2 \leq \widehat{\beta}_k, \left| [\bar{V}_{k,m} V_{k,h+1}^{2^m}](s_h^k, a_h^k) - [\mathbb{V} V_{k,h+1}^{2^m}](s_h^k, a_h^k) \right| \leq E_{k,h,m}.$$

### E.1 Proof of Lemma E.1

We start with the following lemma.

**Lemma E.2.** Let  $V_{k,h}, \widehat{\Sigma}_{k,m}, \widehat{\boldsymbol{\theta}}_{k,m}, \phi_{k,h,m}, \bar{V}_{k,m}$  be defined in Algorithm 2. For any  $k \in [K]$ ,  $h \in [H]$ ,  $m \in \overline{[M]}$ , we have

$$\begin{aligned} & \left| \mathbb{V} V_{k,h+1}^{2^m}(s_h^k, a_h^k) - \bar{V}_{k,m} V_{k,h+1}^{2^m}(s_h^k, a_h^k) \right| \\ & \leq \min \left\{ 1, \left\| \widehat{\Sigma}_{k,m+1}^{-1/2} \phi_{k,h,m+1} \right\|_2 \left\| \widehat{\Sigma}_{k,m+1}^{1/2} (\widehat{\boldsymbol{\theta}}_{k,m+1} - \boldsymbol{\theta}^*) \right\|_2 \right\} \\ & \quad + \min \left\{ 1, 2 \left\| \widehat{\Sigma}_{k,m}^{-1/2} \phi_{k,h,m} \right\|_2 \left\| \widehat{\Sigma}_{k,m}^{1/2} (\widehat{\boldsymbol{\theta}}_{k,m} - \boldsymbol{\theta}^*) \right\|_2 \right\}. \end{aligned}$$

*Proof.* The proof is nearly identical to the proof of Lemma 15 in Zhou et al. (2021a). The only difference is to replace the upper bound  $H$  with 1,  $V_{k,h+1}$  with  $V_{k,h+1}^{2^m}$ ,  $\widehat{\Sigma}_{k,h}$  with  $\widehat{\Sigma}_{k,m+1}$ ,  $\widehat{\boldsymbol{\theta}}_{k,h}$  with  $\widehat{\boldsymbol{\theta}}_{k,m+1}$ ,  $\widehat{\Sigma}_{k,h}$  with  $\widehat{\Sigma}_{k,m}$ ,  $\widehat{\boldsymbol{\theta}}_{k,h}$  with  $\widehat{\boldsymbol{\theta}}_{k,m}$ , and  $\boldsymbol{\theta}_h^*$  with  $\boldsymbol{\theta}^*$ .  $\square$

*Proof of Lemma E.1.* First we recall the definition of  $\bar{\sigma}_{k,h,m}$ :

$$\begin{aligned}\bar{\sigma}_{k,h,m}^2 &= \max \left\{ [\bar{\mathbb{V}}_{k,m} V_{k,h+1}^{2m}](s_h^k, a_h^k) + E_{k,h,m}, \alpha^2, \gamma^2 \left\| \widehat{\Sigma}_{k,m}^{-1/2} \phi_{k,h,m} \right\|_2 \right\}, \\ \bar{\sigma}_{k,h,M-1}^2 &= \max \left\{ 1, \alpha^2, \gamma^2 \left\| \widehat{\Sigma}_{k,M-1}^{-1/2} \phi_{k,h,M-1} \right\|_2 \right\},\end{aligned}$$

We prove the statement by induction. We define the following quantities. For simplicity, let  $\widehat{\mathcal{C}}_{k,m}$  be defined as

$$\widehat{\mathcal{C}}_{k,m} := \left\{ \boldsymbol{\theta} : \left\| \widehat{\Sigma}_{k,m}^{1/2} (\widehat{\boldsymbol{\theta}}_{k,m} - \boldsymbol{\theta}) \right\|_2 \leq \widehat{\beta}_k \right\}.$$

For each  $m$ , let

$$\begin{aligned}\mathbf{x}_{k,h,m} &= \bar{\sigma}_{k,h,m}^{-1} \phi_{k,h,m}, \\ \eta_{k,h,m} &= \bar{\sigma}_{k,h,m}^{-1} \mathbb{1}\{\boldsymbol{\theta}^* \in \widehat{\mathcal{C}}_{k,m} \cap \widehat{\mathcal{C}}_{k,m+1}\} [V_{k,h+1}^{2m}(s_{h+1}^k) - \langle \phi_{k,h,m}, \boldsymbol{\theta}^* \rangle], \\ \eta_{k,h,M-1} &= \bar{\sigma}_{k,h,M-1}^{-1} [V_{k,h+1}^{2M-1}(s_{h+1}^k) - \langle \phi_{k,h,M-1}, \boldsymbol{\theta}^* \rangle] \\ \mathcal{G}_{k,h} &= \mathcal{F}_{k,h}, \\ \boldsymbol{\mu}^* &= \boldsymbol{\theta}^*.\end{aligned}$$

We have

$$\mathbb{E}[\eta_{k,h,m} | \mathcal{G}_{k,h}] = 0, \quad \|\mathbf{x}_{k,h,m}\|_2 \leq \bar{\sigma}_{k,h,m}^{-1} \leq 1/\alpha, \quad |\eta_{k,h,m}| \leq 1/\alpha.$$

Note that  $\mathbb{1}\{\boldsymbol{\theta}^* \in \widehat{\mathcal{C}}_{k,m} \cap \widehat{\mathcal{C}}_{k,m+1}\}$  is  $\mathcal{G}_{k,h}$ -measurable, then we have bound the variance of  $\eta_{k,h,m}$  as follows: for  $m \in \overline{[M-1]}$ ,

$$\begin{aligned}\mathbb{E}[\eta_{k,h,m}^2 | \mathcal{G}_{k,h}] &= \bar{\sigma}_{k,h,m}^{-2} \mathbb{1}\{\boldsymbol{\theta}^* \in \widehat{\mathcal{C}}_{k,m} \cap \widehat{\mathcal{C}}_{k,m+1}\} [\mathbb{V} V_{k,h+1}^{2m}](s_h^k, a_h^k) \\ &\leq \bar{\sigma}_{k,h,m}^{-2} \mathbb{1}\{\boldsymbol{\theta}^* \in \widehat{\mathcal{C}}_{k,m} \cap \widehat{\mathcal{C}}_{k,m+1}\} \left[ \bar{\mathbb{V}}_{k,m} V_{k,h+1}^{2m}(s_h^k, a_h^k) \right. \\ &\quad \left. + \min \left\{ 1, \left\| \widehat{\Sigma}_{k,m+1}^{-1/2} \phi_{k,h,m+1} \right\|_2 \left\| \widehat{\Sigma}_{k,m+1}^{1/2} (\widehat{\boldsymbol{\theta}}_{k,m+1} - \boldsymbol{\theta}^*) \right\|_2 \right\} \right. \\ &\quad \left. + \min \left\{ 1, 2 \left\| \widehat{\Sigma}_{k,m}^{-1/2} \phi_{k,h,m} \right\|_2 \left\| \widehat{\Sigma}_{k,m}^{1/2} (\widehat{\boldsymbol{\theta}}_{k,m} - \boldsymbol{\theta}^*) \right\|_2 \right\} \right] \\ &\leq \bar{\sigma}_{k,h,m}^{-2} \left[ \bar{\mathbb{V}}_{k,m} V_{k,h+1}^{2m}(s_h^k, a_h^k) + \min \left\{ 1, \widehat{\beta}_k \left\| \widehat{\Sigma}_{k,m+1}^{-1/2} \phi_{k,h,m+1} \right\|_2 \right\} \right. \\ &\quad \left. + \min \left\{ 1, 2 \widehat{\beta}_k \left\| \widehat{\Sigma}_{k,m}^{-1/2} \phi_{k,h,m} \right\|_2 \right\} \right] \\ &\leq 1.\end{aligned}$$

For  $m = M - 1$ , we directly have  $\mathbb{E}[\eta_{k,h,M-1}^2 | \mathcal{G}_{k,h}] \leq 1$ . Meanwhile, for any  $m \in \overline{[M]}$ , we have

$$|\eta_{k,h,m}| \max \left\{ 1, \|\mathbf{x}_{k,h,m}\|_{\widehat{\Sigma}_{k,h-1,m}^{-1}} \right\} \leq \bar{\sigma}_{k,h,m}^{-2} \|\phi_{k,h,m}\|_{\widehat{\Sigma}_{k,h-1,m}^{-1}} \leq 1/\gamma^2.$$

Now, let  $y, \mathbf{Z}, \mathbf{b}, \boldsymbol{\mu}, \epsilon$  defined in Theorem 4.3 be set as follows:

$$\begin{aligned}y_{k,h,m} &= \langle \boldsymbol{\mu}^*, \mathbf{x}_{k,h,m} \rangle + \eta_{k,h,m} \\ \mathbf{Z}_{k,m} &= \lambda \mathbf{I} + \sum_{i=1}^k \sum_{h=1}^H \mathbf{x}_{i,h,m} \mathbf{x}_{i,h,m}^\top = \widehat{\Sigma}_{k+1,m} \\ \mathbf{b}_{k,m} &= \sum_{i=1}^k \sum_{h=1}^H \mathbf{x}_{i,h,m} y_{i,h,m} \\ \boldsymbol{\mu}_{k,m} &= \mathbf{Z}_{k,m}^{-1} \mathbf{b}_{k,m}, \\ \epsilon &= 1/\gamma^2.\end{aligned}$$

Then, by Theorem 4.3, for each  $m \in \overline{[M]}$ , with probability at least  $1 - \delta$ ,  $\forall k \in [K + 1]$ ,

$$\begin{aligned} \|\boldsymbol{\mu}_{k-1,m} - \boldsymbol{\theta}^*\|_{\widehat{\boldsymbol{\Sigma}}_{k,m}} &\leq 12\sqrt{d \log(1 + kH/(\alpha^2 d\lambda)) \log(32(\log(\gamma^2/\alpha) + 1)k^2 H^2/\delta)} \\ &\quad + 30 \log(32(\log(\gamma^2/\alpha) + 1)k^2 H^2/\delta)/\gamma^2 + \sqrt{\lambda}B \\ &= \widehat{\beta}_k. \end{aligned} \quad (\text{E.1})$$

Denote the event that (E.1) happens for all  $k$  and  $m$  as  $\mathcal{E}$ . Conditioned on  $\mathcal{E}$ , we have the following observations:

- For  $k = 1$ ,  $m \in \overline{[M]}$ , by the definitions of  $\widehat{\boldsymbol{\theta}}_{1,m}$ ,  $\widehat{\boldsymbol{\Sigma}}_{1,m}$ , we have  $\|\boldsymbol{\theta}^* - \widehat{\boldsymbol{\theta}}_{1,m}\|_{\widehat{\boldsymbol{\Sigma}}_{1,m}} = \|\boldsymbol{\theta}^*\|_{\boldsymbol{\Lambda}} \leq \sqrt{\lambda}B = \widehat{\beta}_1$ , which implies

$$\boldsymbol{\theta}^* \in \widehat{\mathcal{C}}_{1,m}. \quad (\text{E.2})$$

- For  $k \in [K]$  and  $m = M - 1$ , we directly have  $\boldsymbol{\mu}_{k,M-1} = \widehat{\boldsymbol{\theta}}_{k+1,M-1}$ , which implies

$$\boldsymbol{\theta}^* \in \widehat{\mathcal{C}}_{k+1,M-1}. \quad (\text{E.3})$$

- For  $k \in [K]$  and  $m \in \overline{[M - 1]}$ , we have

$$\boldsymbol{\theta}^* \in \widehat{\mathcal{C}}_{k,m} \cap \widehat{\mathcal{C}}_{k,m+1} \Rightarrow y_{k,h,m} = \bar{\sigma}_{k,h,m}^{-1} V_{k,h+1}^{2^m}(s_{h+1}^k) \Rightarrow \boldsymbol{\mu}_{k,m} = \widehat{\boldsymbol{\theta}}_{k+1,m} \Rightarrow \boldsymbol{\theta}^* \in \widehat{\mathcal{C}}_{k+1,m}. \quad (\text{E.4})$$

Therefore by induction based on the initial conditions (E.2) and (E.3) and induction rule (E.4), we have for  $k \in [K]$  and  $m \in \overline{[M]}$ ,  $\boldsymbol{\theta}^* \in \widehat{\mathcal{C}}_{k,m}$ . Lastly, conditioned on  $\mathcal{E}$ , by Lemma E.1, we have for all  $k \in [K]$ ,  $h \in [H]$ ,  $m \in \overline{[M]}$ ,

$$|[\bar{\mathbb{V}}_{k,m} V_{k,h+1}^{2^m}](s_h^k, a_h^k) - [\mathbb{V} V_{k,h+1}^{2^m}](s_h^k, a_h^k)| \leq E_{k,h,m}. \quad \square$$

## E.2 Proof of Theorem 5.1

Let  $\mathcal{E}_{E.1}$  denote the event described by Lemma E.1. We have the following lemmas.

**Lemma E.3.** Let  $Q_{k,h}$ ,  $V_{k,h}$  be defined in Algorithm 2. Then, on the event  $\mathcal{E}_{E.1}$ , for any  $s, a, k, h$  we have that  $Q_h^*(s, a) \leq Q_{k,h}(s, a)$ ,  $V_h^*(s) \leq V_{k,h}(s)$ .

*Proof.* The proof is nearly identical to the proof of Lemma 19 in Zhou et al. 2021a. The only difference is to replace the event used in Lemma 19 in Zhou et al. (2021a) with the event defined by Lemma E.1. Besides, we replace  $\widehat{\mathcal{C}}_{k,h}$  with  $\widehat{\mathcal{C}}_{k,0}$ ,  $\widehat{\boldsymbol{\theta}}_{k,h}$  with  $\widehat{\boldsymbol{\theta}}_{k,0}$ ,  $\widehat{\boldsymbol{\Sigma}}_{k,h}$  with  $\widehat{\boldsymbol{\Sigma}}_{k,0}$ ,  $\mathbb{P}_h$  with  $\mathbb{P}$  and  $\boldsymbol{\theta}_h^*$  with  $\boldsymbol{\theta}$ .  $\square$

**Lemma E.4.** Let  $V_{k,h}$ ,  $\widehat{\boldsymbol{\Sigma}}_{k,0}$ ,  $\boldsymbol{\phi}_{k,h,0}$  be defined in Algorithm 2,  $\widehat{\beta}_k$  be defined in (5.3). Then on the event  $\mathcal{E}_{E.1}$ , for any  $k \in [K]$ ,  $h \in [H]$ , we have

$$V_{k,h}(s_h^k) - r(s_h^k, a_h^k) - [\mathbb{P} V_{k,h+1}](s_h^k, a_h^k) \leq 2 \min\{1, \widehat{\beta}_k \left\| \widehat{\boldsymbol{\Sigma}}_{k,0}^{-1/2} \boldsymbol{\phi}_{k,h,0} \right\|_2\} \quad (\text{E.5})$$

*Proof.* For any  $k \in [K]$  and  $h \in [H]$ , we have

$$\begin{aligned} &V_{k,h}(s_h^k) - r(s_h^k, a_h^k) - [\mathbb{P} V_{k,h+1}](s_h^k, a_h^k) \\ &\leq \langle \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k), \widehat{\boldsymbol{\theta}}_{k,0} \rangle + \widehat{\beta}_k \left\| \widehat{\boldsymbol{\Sigma}}_{k,0}^{-1/2} \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k) \right\|_2 - \langle \boldsymbol{\phi}_{V_{k,h+1}}(s_h^k, a_h^k), \boldsymbol{\theta}^* \rangle \\ &\leq \left\| \widehat{\boldsymbol{\Sigma}}_{k,0}^{1/2} (\widehat{\boldsymbol{\theta}}_{k,0} - \boldsymbol{\theta}^*) \right\|_2 \left\| \widehat{\boldsymbol{\Sigma}}_{k,0}^{-1/2} \boldsymbol{\phi}_{k,h,0} \right\|_2 + \widehat{\beta}_k \left\| \widehat{\boldsymbol{\Sigma}}_{k,0}^{-1/2} \boldsymbol{\phi}_{k,h,0} \right\|_2 \\ &\leq 2\widehat{\beta}_k \left\| \widehat{\boldsymbol{\Sigma}}_{k,0}^{-1/2} \boldsymbol{\phi}_{k,h,0} \right\|_2, \end{aligned} \quad (\text{E.6})$$

where the first inequality holds due to the definition of  $V_{k,h}$ , the second one holds due to Cauchy-Schwarz inequality and the last one holds since on  $\mathcal{E}_{E.1}$ ,  $\widehat{\boldsymbol{\theta}}_{k,0} \in \widehat{\mathcal{C}}_{k,0}$ . Meanwhile, we have  $V_{k,h}(s_h^k) - r(s_h^k, a_h^k) - [\mathbb{P} V_{k,h+1}](s_h^k, a_h^k) \leq 2$ . Therefore, combining two bounds completes the proof.  $\square$

Recall  $\iota = \log(1 + KH/(d\lambda\alpha^2))$ ,  $\zeta = 4 \log(4 \log(KH)/\delta)$ . For any  $k \in [K]$ ,  $h \in [H]$  we define the following indicator function  $I_h^k$  such that

$$I_h^k := \mathbb{1} \left\{ \forall m \in [\overline{M}], \det(\widehat{\Sigma}_{k,m}^{-1/2}) / \det(\widetilde{\Sigma}_{k,h,m}^{-1/2}) \leq 4 \right\}.$$

Clearly  $I_h^k$  is  $\mathcal{F}_{k,h}$ -measurable and monotonically decreasing. For  $m \in [\overline{M}]$ , we also define the following quantities:

$$R_m = \sum_{k=1}^K \sum_{h=1}^H I_h^k \min \left\{ 1, \widehat{\beta}_k \left\| \widehat{\Sigma}_{k,m}^{-1/2} \phi_{k,h,m} \right\|_2 \right\}, \quad (\text{E.7})$$

$$A_m = \sum_{k=1}^K \sum_{h=1}^H I_h^k [\mathbb{P}V_{k,h+1}^{2m}](s_h^k, a_h^k) - V_{k,h+1}^{2m}(s_{h+1}^k), \quad (\text{E.8})$$

$$S_m = \sum_{k=1}^K \sum_{h=1}^H I_h^k [\mathbb{V}V_{k,h+1}^{2m}](s_h^k, a_h^k), \quad (\text{E.9})$$

$$G = \sum_{k=1}^K (1 - I_H^k). \quad (\text{E.10})$$

We aim to bound  $R_m, A_m, S_m, G$  by the following lemmas.

**Lemma E.5.** Let  $\gamma, \alpha$  be defined in Algorithm 2,  $\{R_m, S_m\}_{m \in [\overline{M}]}$  be defined in (E.7) and (E.9).

Then for  $m \in [\overline{M} - 1]$ , we have

$$R_m \leq \min \{ 4d\iota + 4\widehat{\beta}_K \gamma^2 d\iota + 2\widehat{\beta}_K \sqrt{d\iota} \sqrt{S_m + 4R_m + 2R_{m+1} + KH\alpha^2}, KH \}. \quad (\text{E.11})$$

For  $R_{M-1}$ , we have  $R_{M-1} \leq KH$ .

*Proof.* For  $(k, h)$  where  $I_h^k = 1$ , by Lemma F.6 we have

$$\left\| \widehat{\Sigma}_{k,m}^{-1/2} \phi_{k,h,m} \right\|_2 \leq \left\| \widetilde{\Sigma}_{k,h,m}^{-1/2} \phi_{k,h,m} \right\|_2 \cdot \sqrt{\frac{\det \widehat{\Sigma}_{k,m}^{-1/2}}{\det \widetilde{\Sigma}_{k,h,m}^{-1/2}}} \leq 2 \left\| \widetilde{\Sigma}_{k,h,m}^{-1/2} \phi_{k,h,m} \right\|_2, \quad (\text{E.12})$$

then substituting (E.12) into (E.7), we have

$$R_m \leq 2 \sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, I_h^k \widehat{\beta}_k \left\| \widetilde{\Sigma}_{k,h,m}^{-1/2} \phi_{k,h,m} \right\|_2 \right\}. \quad (\text{E.13})$$

(E.13) can be bounded by Lemma D.1, with  $\beta_{k,h} = I_h^k \widehat{\beta}_k$ ,  $\bar{\sigma}_{k,h} = \bar{\sigma}_{k,h,m}$ ,  $\mathbf{a}_{k,h} = \phi_{k,h,m}$  and  $\widehat{\Sigma}_{k,h} = \widetilde{\Sigma}_{k,h,m}$ . We have

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, I_h^k \widehat{\beta}_k \left\| \widetilde{\Sigma}_{k,h,m}^{-1/2} \phi_{k,h,m} \right\|_2 \right\} \\ & \leq 2d\iota + 2\widehat{\beta}_K \gamma^2 d\iota + 2\widehat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H I_h^k [\mathbb{V}V_{k,h+1}^{2m}](s_h^k, a_h^k) + E_{k,h,m}} + KH\alpha^2 \sqrt{d\iota} \\ & \leq 2d\iota + 2\widehat{\beta}_K \gamma^2 d\iota + 2\widehat{\beta}_K \sqrt{\sum_{k=1}^K \sum_{h=1}^H I_h^k [\mathbb{V}V_{k,h+1}^{2m}](s_h^k, a_h^k) + 2E_{k,h,m}} + KH\alpha^2 \sqrt{d\iota} \end{aligned}$$

Note that

$$\sum_{k=1}^K \sum_{h=1}^H I_h^k E_{k,h,m} = \sum_{k=1}^K \sum_{h=1}^H I_h^k \min \left\{ 1, 2\widehat{\beta}_k \left\| \widetilde{\Sigma}_{k,m}^{-1/2} \phi_{k,h,m} \right\|_2 \right\}$$



$$\begin{aligned}
& + \sum_{k=1}^K \sum_{h=1}^H I_h^k \min \left\{ 1, \widehat{\beta}_k \left\| \widehat{\Sigma}_{k,m+1}^{-1/2} \phi_{k,h,m+1} \right\|_2 \right\} \\
& \leq 2R_m + R_{m+1},
\end{aligned}$$

by the definition of  $R_m$  in (E.7). This completes our proof.  $\square$

**Lemma E.6** (Lemma 25, Zhang et al. 2021c). Let  $\{S_m, A_m\}_{m \in \overline{[M]}}$  be defined in (E.8) and (E.9),  $G$  be defined in (E.10). Then on the event  $\mathcal{E}_{E,1}$ , for  $m \in \overline{[M-1]}$ , we have

$$S_m \leq |A_{m+1}| + G + 2^{m+1}(K + 2R_0). \quad (\text{E.14})$$

*Proof.* The proof follows the proof of Lemma 25 in Zhang et al. 2021c. We have

$$\begin{aligned}
S_m &= \sum_{k=1}^K \sum_{h=1}^H I_h^k [\mathbb{P}V_{k,h+1}^{2^{m+1}}(s_h^k, a_h^k) - ([\mathbb{P}V_{k,h+1}^{2^m}](s_h^k, a_h^k))^2] \\
&= \sum_{k=1}^K \sum_{h=1}^H I_h^k [\mathbb{P}V_{k,h+1}^{2^{m+1}}(s_h^k, a_h^k) - V_{k,h+1}^{2^{m+1}}(s_{h+1}^k)] + I_h^k [V_{k,h}^{2^{m+1}}(s_h^k) - ([\mathbb{P}V_{k,h+1}^{2^m}](s_h^k, a_h^k))^2] \\
&\quad + \sum_{k=1}^K \sum_{h=1}^H I_h^k (V_{k,h+1}^{2^{m+1}}(s_{h+1}^k) - V_{k,h}^{2^{m+1}}(s_h^k)) \\
&\leq A_{m+1} + \sum_{k=1}^K \sum_{h=1}^H I_h^k [V_{k,h}^{2^{m+1}}(s_h^k) - ([\mathbb{P}V_{k,h+1}^{2^m}](s_h^k, a_h^k))^2] + \sum_{k=1}^K I_{h_k}^k V_{k,h_k+1}^{2^{m+1}}(s_{h_k+1}^k),
\end{aligned}$$

where  $h_k$  is the largest index such that  $I_{h_k}^k = 1$ . Note that if  $h_k < H$ , we have  $I_{h_k}^k V_{k,h_k+1}^{2^{m+1}}(s_{h_k+1}^k) \leq 1 = 1 - I_H^k$ , and if  $h_k = H$ , we have  $I_{h_k}^k V_{k,h_k+1}^{2^{m+1}}(s_{h_k+1}^k) = 0 = 1 - I_H^k$ . Thus, for both cases, we have

$$S_m \leq A_{m+1} + \sum_{k=1}^K (1 - I_H^k) + \sum_{k=1}^K \sum_{h=1}^H I_h^k [V_{k,h}^{2^{m+1}}(s_h^k) - ([\mathbb{P}V_{k,h+1}^{2^m}](s_h^k, a_h^k))^2]. \quad (\text{E.15})$$

For the third term in (E.15), we have

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h=1}^H I_h^k [V_{k,h}^{2^{m+1}}(s_h^k) - ([\mathbb{P}V_{k,h+1}^{2^m}](s_h^k, a_h^k))^2] \\
& \leq \sum_{k=1}^K \sum_{h=1}^H I_h^k [V_{k,h}^{2^{m+1}}(s_h^k) - ([\mathbb{P}V_{k,h+1}](s_h^k, a_h^k))^{2^{m+1}}] \\
& = \sum_{k=1}^K \sum_{h=1}^H I_h^k (V_{k,h}(s_h^k) - [\mathbb{P}V_{k,h+1}](s_h^k, a_h^k)) \prod_{i=0}^m (V_{k,h}^{2^i}(s_h^k) + [\mathbb{P}V_{k,h+1}](s_h^k, a_h^k)^{2^i}) \\
& \leq 2^{m+1} \sum_{k=1}^K \sum_{h=1}^H I_h^k \max\{V_{k,h}(s_h^k) - [\mathbb{P}V_{k,h+1}](s_h^k, a_h^k), 0\} \\
& \leq 2^{m+1} \sum_{k=1}^K \sum_{h=1}^H I_h^k [r(s_h^k, a_h^k) + 2 \min\{1, \widehat{\beta}_k \left\| \widehat{\Sigma}_{k,0}^{-1/2} \phi_{k,h,0} \right\|_2\}] \\
& \leq 2^{m+1}(K + 2R_0),
\end{aligned} \quad (\text{E.16})$$

where the first inequality holds by recursively applying  $\mathbb{E}X^2 \geq (\mathbb{E}X)^2$ , the second one holds since  $V_{k,h} \in [0, 1]$ , the third one holds due to Lemma E.4, the last one holds due to the definition of  $R_0$  (E.7). Substituting (E.16) into (E.15) completes our proof.  $\square$

**Lemma E.7** (Lemma 25, Zhang et al. 2021c). Let  $\{A_m, S_m\}_{m \in \overline{[M]}}$  be defined in (E.8) and (E.9). Then we have  $\mathbb{P}(\mathcal{E}_{E,7}) > 1 - M\delta$ , where

$$\mathcal{E}_{E,7} := \{\forall m \in \overline{[M]}, |A_m| \leq \min\{\sqrt{2\zeta S_m} + \zeta, KH\}\}. \quad (\text{E.17})$$

*Proof.* The proof follows the proof of Lemma 25 in Zhang et al. (2021c). We use Lemma F.2 for each  $m$ . Let  $x_{k,h} = I_h^k[\mathbb{P}V_{k,h+1}^{2^m}(s_h^k, a_h^k) - V_{k,h+1}^{2^m}(s_{h+1}^k)]$ , then we have  $\mathbb{E}[x_{k,h}|\mathcal{F}_{k,h}] = 0$  and  $\mathbb{E}[x_{k,h}^2|\mathcal{F}_{k,h}] = I_h^k[\mathbb{V}V_{k,h+1}^{2^m}(s_h^k, a_h^k)]$ . Therefore, for each  $m \in \overline{[M]}$ , with probability at least  $1 - \delta$ , we have

$$A_m = \sum_{k=1}^K \sum_{h=1}^H x_{k,h} \leq \sqrt{2\zeta \sum_{k=1}^K \sum_{h=1}^H I_h^k[\mathbb{V}V_{k,h+1}^{2^m}(s_h^k, a_h^k)]} + \zeta.$$

Taking union bound over  $m \in \overline{[M]}$  completes the proof.  $\square$

**Lemma E.8.** Let  $G$  be defined in (E.10). Then we have  $G \leq Mdt/2$ .

*Proof.* Recall the definition of  $I_H^k$ , by the fact that  $\det(\widehat{\Sigma}_{k+1,m}^{-1/2}) < \det(\widetilde{\Sigma}_{k,H,m}^{-1/2})$ , we have

$$\begin{aligned} (1 - I_H^k) = 1 &\Leftrightarrow \exists m \in \overline{[M]}, \det(\widehat{\Sigma}_{k,m}^{-1/2}) / \det(\widetilde{\Sigma}_{k,H,m}^{-1/2}) > 4 \\ &\Rightarrow \exists m \in \overline{[M]}, \det(\widehat{\Sigma}_{k,m}^{-1/2}) / \det(\widehat{\Sigma}_{k+1,m}^{-1/2}) > 4. \end{aligned}$$

Let  $\mathcal{D}_m$  denote the indices  $k$  such that

$$\mathcal{D}_m := \left\{ k \in [K] : \det(\widehat{\Sigma}_{k+1,m}) / \det(\widehat{\Sigma}_{k,m}) > 16 \right\}.$$

Then we have  $G \leq |\cup_{m=0}^{M-1} \mathcal{D}_m| \leq \sum_{m=0}^{M-1} |\mathcal{D}_m|$ . For each  $m$ , we have

$$2|\mathcal{D}_m| < \sum_{k \in \mathcal{D}_m} \log 16 < \sum_{k \in \mathcal{D}_m} \log(\det(\widehat{\Sigma}_{k+1,m}) / \det(\widehat{\Sigma}_{k,m})) \leq \sum_{k=1}^K \log(\det(\widehat{\Sigma}_{k+1,m}) / \det(\widehat{\Sigma}_{k,m})).$$

Furthermore, by the facts that  $\log \det(\widehat{\Sigma}_{K+1,m}) \leq d \log[\text{tr}(\widehat{\Sigma}_{K+1,m})/d]$  and  $\text{tr}(\widehat{\Sigma}_{K+1,m}) \leq \text{tr}(\lambda \mathbf{I}) + \sum_{k,h} \|\phi_{k,h,m}\|_2^2 / \bar{\sigma}_{k,h,m}^2 \leq d\lambda + KH/\alpha^2$ , we have

$$\begin{aligned} \sum_{k=1}^K \log(\det(\widehat{\Sigma}_{k+1,m}) / \det(\widehat{\Sigma}_{k,m})) &= \log(\det(\widehat{\Sigma}_{K+1,m}) / \det(\widehat{\Sigma}_{1,m})) \\ &\leq d(\log(\lambda + KH/d\alpha^2) - \log \lambda). \end{aligned}$$

Therefore,  $|\mathcal{D}_m|$  can be upper bounded by

$$|\mathcal{D}_m| \leq d/2 \cdot \log(1 + KH/(d\lambda\alpha^2)) = d/2 \cdot \iota.$$

Taking the summation over  $m$  gives the upper bound of  $G$ .  $\square$

Finally, we define the event  $\mathcal{E}_{E.18}$  as

$$\mathcal{E}_{E.18} := \left\{ \sum_{k=1}^K \left( \sum_{h=1}^H r(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k) \right) \leq \sqrt{2K \log(1/\delta)} \right\}. \quad (\text{E.18})$$

By Azuma-Hoeffding inequality (Lemma F.1) we have

$$\mathbb{P}(\mathcal{E}_{E.18}) > 1 - \delta. \quad (\text{E.19})$$

With all above lemmas, we are ready to prove Theorem 5.1.

*Proof of Theorem 5.1.* All the following proofs are conditioned on  $\mathcal{E}_{E.1} \cap \mathcal{E}_{E.7} \cap \mathcal{E}_{E.18}$ , which happens with probability at least  $1 - (2M + 1)\delta$  by union bound and the probabilities of individual events  $\mathcal{E}_{E.1}$ ,  $\mathcal{E}_{E.7}$ ,  $\mathcal{E}_{E.18}$  specified in Lemma E.1, Lemma E.7 and (E.19). First we have  $\text{Regret}(K) \leq \sum_{k=1}^K [V_{k,1}(s_1^k) - V_1^{\pi^k}(s_1^k)]$  by Lemma E.3. Next we have

$$\sum_{k=1}^K V_{k,1}(s_1^k)$$

$$\begin{aligned}
&= \sum_{k=1}^K \sum_{h=1}^H [I_h^k [V_{k,h}(s_h^k) - V_{k,h+1}(s_{h+1}^k)] + (1 - I_h^k) [V_{k,h}(s_h^k) - V_{k,h+1}(s_{h+1}^k)]] \\
&= \sum_{k=1}^K \left[ \sum_{h=1}^H I_h^k r(s_h^k, a_h^k) + \sum_{h=1}^H I_h^k [V_{k,h}(s_h^k) - r(s_h^k, a_h^k) - [\mathbb{P}V_{k,h+1}](s_h^k, a_h^k)] \right. \\
&\quad \left. + \sum_{h=1}^H I_h^k [\mathbb{P}V_{k,h+1}](s_h^k, a_h^k) - V_{k,h+1}(s_{h+1}^k) \right] + \sum_{k=1}^K \sum_{h=1}^H (1 - I_h^k) [V_{k,h}(s_h^k) - V_{k,h+1}(s_{h+1}^k)] \\
&\leq \sum_{k=1}^K \left[ \sum_{h=1}^H r(s_h^k, a_h^k) + \sum_{h=1}^H I_h^k [V_{k,h}(s_h^k) - r(s_h^k, a_h^k) - [\mathbb{P}V_{k,h+1}](s_h^k, a_h^k)] \right. \\
&\quad \left. + \sum_{h=1}^H I_h^k [\mathbb{P}V_{k,h+1}](s_h^k, a_h^k) - V_{k,h+1}(s_{h+1}^k) \right] + \sum_{k=1}^K (1 - I_{h_k}^k) V_{k,h_k}(s_{h_k}^k),
\end{aligned}$$

where  $h_k$  is the smallest index such that  $I_{h_k}^k = 0$ . Then we have

$$\begin{aligned}
&\text{Regret}(K) \\
&\leq \underbrace{\sum_{k=1}^K \left( \sum_{h=1}^H r(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k) \right)}_{I_1} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H I_h^k [V_{k,h}(s_h^k) - r(s_h^k, a_h^k) - [\mathbb{P}V_{k,h+1}](s_h^k, a_h^k)]}_{I_2} \\
&\quad + \underbrace{\sum_{k=1}^K \sum_{h=1}^H I_h^k [[\mathbb{P}V_{k,h+1}](s_h^k, a_h^k) - V_{k,h+1}(s_{h+1}^k)]}_{A_0} + \underbrace{\sum_{k=1}^K (1 - I_{h_k}^k)}_G \\
&\leq \sqrt{2K \log(1/\delta)} + 2R_0 + |A_0| + G. \tag{E.20}
\end{aligned}$$

The bound of  $I_1$  comes from  $\mathcal{E}_{E.18}$ . For  $I_2$ , by Lemma E.4 we have

$$I_2 \leq 2 \sum_{k=1}^K \sum_{h=1}^H I_h^k \min\{1, \hat{\beta}_k \|\widehat{\Sigma}_{k,0}^{-1/2} \phi_{k,h,0}\|_2\} = 2R_0.$$

Next we bound  $2R_0 + |A_0|$  in (E.20). Substituting (E.14) in Lemma E.6 into (E.17) in Lemma E.7, we have

$$\begin{aligned}
|A_m| &\leq \sqrt{2\zeta(|A_{m+1}| + G + 2^{m+1}(K + 2R_0))} + \zeta \\
&\leq \sqrt{2\zeta} \sqrt{|A_{m+1}| + 2^{m+1}(K + 2R_0)} + \sqrt{2\zeta G} + \zeta \tag{E.21}
\end{aligned}$$

Substituting (E.14) in Lemma E.6 into (E.11) in Lemma E.5, we have

$$\begin{aligned}
R_m &\leq 4d\iota + 4\hat{\beta}_K \gamma^2 d\iota + 2\hat{\beta}_K \sqrt{d\iota} \sqrt{|A_{m+1}| + G + 2^{m+1}(K + 2R_0)} + 4R_m + 2R_{m+1} + KH\alpha^2 \\
&\leq 2\hat{\beta}_K \sqrt{d\iota} \sqrt{|A_{m+1}| + 2^{m+1}(K + 2R_0)} + 4R_m + 2R_{m+1} + \\
&\quad + \underbrace{4d\iota + 4\hat{\beta}_K \gamma^2 d\iota + 2\hat{\beta}_K \sqrt{d\iota} \sqrt{G + KH\alpha^2}}_{I_c} \tag{E.22}
\end{aligned}$$

Calculating (E.21) + 2×(E.22) and using  $\sqrt{a} + \sqrt{b} \leq \sqrt{2(a+b)}$ , we have

$$\begin{aligned}
&|A_m| + 2R_m \\
&\leq 2I_c + \sqrt{2\zeta G} + \zeta + \sqrt{2} \max\{4\hat{\beta}_K \sqrt{d\iota}, \sqrt{2\zeta}\} \\
&\quad \cdot \sqrt{|A_{m+1}| + 2^{m+1}(K + 2R_0) + 4R_m + 2R_{m+1} + |A_{m+1}| + 2^{m+1}(K + 2R_0)} \\
&\leq 2I_c + \sqrt{2\zeta G} + \zeta + 2 \max\{4\hat{\beta}_K \sqrt{d\iota}, \sqrt{2\zeta}\} \\
&\quad \cdot \sqrt{|A_m| + 2R_m + |A_{m+1}| + 2R_{m+1} + 2^{m+1}(K + |A_0| + 2R_0)}
\end{aligned}$$

Then by Lemma F.7 with  $a_m = |A_m| + 2R_m \leq 3KH$  and  $M = \log(3KH)/\log 2$ ,  $|A_0| + 2R_0$  can be bounded as

$$\begin{aligned}
& |A_0| + 2R_0 \\
& \leq 22 \cdot 8 \max\{2\widehat{\beta}_K^2 dt, \zeta\} + 6 \cdot (2I_c + \sqrt{2\zeta G} + \zeta) + 8 \max\{4\widehat{\beta}_K \sqrt{dt}, \sqrt{2\zeta}\} \sqrt{2(K + |A_0| + 2R_0)} \\
& \leq 176 \max\{2\widehat{\beta}_K^2 dt, \zeta\} + 6(8dt + 8\widehat{\beta}_K \gamma^2 dt + 4\widehat{\beta}_K \sqrt{dt} \sqrt{G + KH\alpha^2} + \sqrt{2\zeta G} + \zeta) \\
& \quad + 16 \max\{2\widehat{\beta}_K \sqrt{2dt}, \sqrt{\zeta}\} \sqrt{K} + 16 \max\{2\widehat{\beta}_K \sqrt{2dt}, \sqrt{\zeta}\} \sqrt{|A_0| + 2R_0}. \tag{E.23}
\end{aligned}$$

By the fact  $x \leq a\sqrt{x} + b \Rightarrow x \leq 2a^2 + 2b$ , (E.23) implies

$$\begin{aligned}
|A_0| + 2R_0 & \leq 864 \max\{8\widehat{\beta}_K^2 dt, \zeta\} + 12(8dt + 8\widehat{\beta}_K \gamma^2 dt + 4\widehat{\beta}_K \sqrt{dt} \sqrt{G + KH\alpha^2} + \sqrt{2\zeta G} + \zeta) \\
& \quad + 32 \max\{2\widehat{\beta}_K \sqrt{2dt}, \sqrt{\zeta}\} \sqrt{K}. \tag{E.24}
\end{aligned}$$

Finally, substituting (E.24) into (E.20) and bounding  $G$  by Lemma E.8, the regret is bounded as

$$\begin{aligned}
\text{Regret}(K) & \leq 12(8dt + 8\widehat{\beta}_K \gamma^2 dt + 4\widehat{\beta}_K \sqrt{dt} \sqrt{Mdt/2 + KH\alpha^2} + \sqrt{Mdt\zeta} + \zeta) \\
& \quad + 864 \max\{8\widehat{\beta}_K^2 dt, \zeta\} + Mdt/2 + [\sqrt{2\log(1/\delta)} + 32 \max\{2\widehat{\beta}_K \sqrt{2dt}, \sqrt{\zeta}\}] \sqrt{K},
\end{aligned}$$

which completes our proof.  $\square$

### E.3 Proof of Theorem 5.3

We have the following lemma to lower bound the regret for linear bandits.

**Lemma E.9** (Lemma 25, Zhou et al. 2021a). Fix a positive real  $0 < \delta \leq 1/3$ , and positive integers  $K, d$  and assume that  $K \geq d^2/(2\delta)$ . Let  $\Delta = \sqrt{\delta/K}/(4\sqrt{2})$  and consider the linear bandit problems  $\mathcal{L}_\mu$  parameterized with a parameter vector  $\mu \in \{-\Delta, \Delta\}^d$  and action set  $\mathcal{A} = \{-1, 1\}^d$  so that the reward distribution for taking action  $\mathbf{a} \in \mathcal{A}$  is a Bernoulli distribution  $B(\delta + \langle \mu^*, \mathbf{a} \rangle)$ . Then for any bandit algorithm  $\mathcal{B}$ , there exists a  $\mu^* \in \{-\Delta, \Delta\}^d$  such that the expected pseudo-regret of  $\mathcal{B}$  over first  $K$  steps on bandit  $\mathcal{L}_{\mu^*}$  is lower bounded as follows:

$$\mathbb{E}_{\mu^*} \text{Regret}(K) \geq \frac{d\sqrt{K\delta}}{8\sqrt{2}}.$$

*Proof of Theorem 5.3.* The linear mixture MDP instance is similar to the MDP instances considered in Zhou et al. (2021b,a); Zhang et al. (2021a). The state space  $\mathcal{S} = \{x_1, x_2, x_3\}$ . The action space  $\mathcal{A} = \{\mathbf{a}\} = \{-1, +1\}^{d-1}$ . The reward function satisfies  $r(x_1, \mathbf{a}) = r(x_2, \mathbf{a}) = 0$  and  $r(x_3, \mathbf{a}) = 1/H$ . The transition probability satisfies  $\mathbb{P}(x_2|x_1, \mathbf{a}) = 1 - (\delta + \langle \mu, \mathbf{a} \rangle)$  and  $\mathbb{P}(x_3|x_1, \mathbf{a}) = \delta + \langle \mu, \mathbf{a} \rangle$ , where  $\delta = 1/6$  and  $\mu \in \{-\Delta, \Delta\}^{d-1}$  with  $\Delta = \sqrt{\delta/K}/(4\sqrt{2})$ .

First, similar to the proof in (Section E.1, Zhou et al. 2021a), we can verify that when  $K \geq (d-1)/(192(B-1))$ , our instance is a  $B$ -bounded linear mixture MDP with  $\mathbb{P}(s'|s, \mathbf{a}) = \langle \phi(s'|s, \mathbf{a}), \boldsymbol{\theta} \rangle$ , where

$$\begin{aligned}
\phi(s'|s, \mathbf{a}) & = \begin{cases} (\alpha(1-\delta), -\beta\mathbf{a}^\top)^\top, & s = x_1, s' = x_2; \\ (\alpha\delta, \beta\mathbf{a}^\top)^\top, & s = x_1, s' = x_3; \\ (\alpha, \mathbf{0}^\top)^\top, & s \in \{x_2, x_3\}, s' = s; \\ \mathbf{0}, & \text{otherwise.} \end{cases} \\
\boldsymbol{\theta} & = (1/\alpha, \mu^\top/\beta)^\top.
\end{aligned}$$

where  $\alpha = \sqrt{1/(1 + \Delta(d-1))}$ ,  $\beta = \sqrt{\Delta/(1 + \Delta(d-1))}$ .

Second, our instance can be regarded as a linear bandit instance with a Bernoulli reward distribution  $B(\delta + \langle \boldsymbol{\theta}, \mathbf{a} \rangle)$ . Therefore, the lower bound of regret for linear mixture MDP directly follows the regret for linear bandits in Lemma E.9, by picking  $\mu = \mu^*$  and  $\delta = 1/6$ .  $\square$

## F Auxiliary lemmas

**Lemma F.1** (Azuma-Hoeffding inequality, Azuma 1967). Let  $M > 0$  be a constant. Let  $\{x_i\}_{i=1}^n$  be a stochastic process,  $\mathcal{G}_i = \sigma(x_1, \dots, x_i)$  be the  $\sigma$ -algebra of  $x_1, \dots, x_i$ . Suppose  $\mathbb{E}[x_i | \mathcal{G}_{i-1}] = 0$ ,  $|x_i| \leq M$  almost surely. Then, for any  $0 < \delta < 1$ , we have

$$\mathbb{P}\left(\sum_{i=1}^n x_i \leq M\sqrt{2n \log(1/\delta)}\right) > 1 - \delta.$$

**Lemma F.2** (Lemma 11, Zhang et al. 2021d). Let  $M > 0$  be a constant. Let  $\{x_i\}_{i=1}^n$  be a stochastic process,  $\mathcal{G}_i = \sigma(x_1, \dots, x_i)$  be the  $\sigma$ -algebra of  $x_1, \dots, x_i$ . Suppose  $\mathbb{E}[x_i | \mathcal{G}_{i-1}] = 0$ ,  $|x_i| \leq M$  and  $\mathbb{E}[x_i^2 | \mathcal{G}_{i-1}] < \infty$  almost surely. Then, for any  $\delta, \epsilon > 0$ , we have

$$\begin{aligned} \mathbb{P}\left(\left|\sum_{i=1}^n x_i\right| \leq 2\sqrt{2 \log(1/\delta) \sum_{i=1}^n \mathbb{E}[x_i^2 | \mathcal{G}_{i-1}] + 2\sqrt{\log(1/\delta)}\epsilon + 2M \log(1/\delta)}\right) \\ > 1 - 2(\log(M^2 n / \epsilon^2) + 1)\delta. \end{aligned}$$

**Lemma F.3** (Unbounded Freedman's inequality, Dzhaparidze and Van Zanten 2001; Fan et al. 2017). Let  $\{x_i\}_{i=1}^n$  be a stochastic process,  $\mathcal{G}_i = \sigma(x_1, \dots, x_i)$  be the  $\sigma$ -algebra of  $x_1, \dots, x_i$ . Suppose  $\mathbb{E}[x_i | \mathcal{G}_{i-1}] = 0$  and  $\mathbb{E}[x_i^2 | \mathcal{G}_{i-1}] < \infty$  almost surely. Then, for any  $a, v, y > 0$ , we have

$$\mathbb{P}\left(\sum_{i=1}^n x_i > a, \sum_{i=1}^n (\mathbb{E}[x_i^2 | \mathcal{G}_{i-1}] + x_i^2 \mathbf{1}\{|x_i| > y\}) < v^2\right) \leq \exp\left(\frac{-a^2}{2(v^2 + ay/3)}\right).$$

**Lemma F.4** (Lemma 8, Zhang et al. 2021c). Let  $\{x_i \geq 0\}_{i \geq 1}$  be a stochastic process,  $\{\mathcal{G}_i\}_{i \geq 1}$  be a filtration satisfying  $x_i$  is  $\mathcal{G}_i$ -measurable. We also have  $|x_i| \leq 1$ . Then for any  $c \geq 1$  we have

$$\mathbb{P}\left(\exists n, \sum_{i=1}^n x_i \geq 4c \log(4/\delta), \sum_{i=1}^n \mathbb{E}[x_i | \mathcal{G}_{i-1}] \leq c \log(4/\delta)\right) \leq \delta.$$

**Lemma F.5** (Lemma 11, Abbasi-Yadkori et al. 2011). For any  $\lambda > 0$  and sequence  $\{\mathbf{x}_k\}_{k=1}^K \subset \mathbb{R}^d$  for  $k \in [K]$ , define  $\mathbf{Z}_k = \lambda \mathbf{I} + \sum_{i=1}^{k-1} \mathbf{x}_i \mathbf{x}_i^\top$ . Then, provided that  $\|\mathbf{x}_k\|_2 \leq L$  holds for all  $k \in [K]$ , we have

$$\sum_{k=1}^K \min\{1, \|\mathbf{x}_k\|_{\mathbf{Z}_k}^2\} \leq 2d \log(1 + KL^2/(d\lambda)).$$

**Lemma F.6** (Lemma 12, Abbasi-Yadkori et al. 2011). Suppose  $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{d \times d}$  are two positive definite matrices satisfying  $\mathbf{A} \succeq \mathbf{B}$ , then for any  $\mathbf{x} \in \mathbb{R}^d$ ,  $\|\mathbf{x}\|_{\mathbf{A}} \leq \|\mathbf{x}\|_{\mathbf{B}} \cdot \sqrt{\det(\mathbf{A})/\det(\mathbf{B})}$ .

**Lemma F.7** (Lemma 12, Zhang et al. 2021c). Let  $\lambda_1, \lambda_2, \lambda_4 > 0$ ,  $\lambda_3 \geq 1$  and  $\kappa = \max\{\log_2 \lambda_1, 1\}$ . Let  $a_1, \dots, a_\kappa$  be non-negative real numbers such that  $a_i \leq \min\{\lambda_1, \lambda_2 \sqrt{a_i + a_{i+1}} + 2^{i+1} \lambda_3 + \lambda_4\}$  for any  $1 \leq i \leq \kappa$ . Let  $a_{\kappa+1} = \lambda_1$ . Then we have  $a_1 \leq 22\lambda_2^2 + 6\lambda_4 + 4\lambda_2 \sqrt{2\lambda_3}$ .