NEAR-OPTIMAL REWARD-FREE EXPLORATION FOR LINEAR MIXTURE MDPS WITH PLUG-IN SOLVER

Anonymous authors
Paper under double-blind review

ABSTRACT

Although model-based reinforcement learning (RL) approaches are considered more sample efficient, existing algorithms are usually relying on sophisticated planning algorithm to couple tightly with the model-learning procedure. Hence the learned models may lack the ability of being re-used with more specialized planners. In this paper we address this issue and provide approaches to learn an RL model efficiently without the guidance of a reward signal. In particular, we take a plug-in solver approach, where we focus on learning a model in the exploration phase and demand that *any planning algorithm* on the learned model can give a near-optimal policy. Specicially, we focus on the linear mixture MDP setting, where the probability transition matrix is a (unknown) convex combination of a set of existing models. We show that, by establishing a novel exploration algorithm, the plug-in approach learns a model by taking $\tilde{O}(d^2H^3/\epsilon^2)$ interactions with the environment and *any* ϵ-optimal planner on the model gives an $O(\epsilon)$-optimal policy on the original model. This sample complexity matches lower bounds for non-plug-in approaches and is *statistically optimal*. We achieve this result by leveraging a careful maximum total-variance bound using Bernstein inequality and properties specified to linear mixture MDP.

1 INTRODUCTION

In reinforcement learning, an agent repeatedly interacts with the unknown environment in order to maximize the cumulative reward. To achieve this goal, an RL algorithm must be equipped with effective exploration mechanisms to learn the unknown environment and find a near-optimal policy. Efficient exploration is critical to the success of reinforcement learning algorithms, which has been widely investigated from both the empirical and the theoretical perspectives (e.g. Stadie et al. (2015); Pathak et al. (2017); Azar et al. (2017); Jin et al. (2018)). Model-based RL is one of the important approaches to solve for the RL environment. In model-based RL, the agent learns the model of the environment and then performs planning in the estimated model. It has been widely applied in many RL scenarios, including both online setting (Kaiser et al., 2019; Luo et al., 2019; Azar et al., 2017) and offline setting (Yu et al., 2020; Kidambi et al., 2020). It is also believed that model-based RL is significantly more sample-efficient than model-free RL, which has been justified by many recent empirical results (e.g. Kaiser et al. (2019); Wang et al. (2019)). Though the theoretical model-based learning in small scale problems has been studied extensively (Azar et al., 2017; Zhou et al., 2020a; Jin et al., 2020a), it is still far from complete, especially with the presence of a function approximator.

As an important implication of model-based approaches, the power of *plug-in* approach have been studied in several works (Cui & Yang, 2020; Agarwal et al., 2020). The idea of plug-in approach is rather simple: We construct an empirical Markov Decision Process (MDP) using maximum likelihood estimate, then return the (approximate) optimal policy with *efficient planning algorithm* in this empirical model. The significance of plug-in approaches is two-folded. For one thing, it preserves an empirical model that keeps the value of the policies, which is of independent interests. For another, the empirical model can be used for any down-stream tasks, which makes the application much more flexible. It is shown that the plug-in approach achieves the minimax sample complexity to compute the ϵ-optimal policies with a generative model in the tabular (Agarwal et al., 2020) and linear settings (Cui & Yang, 2020).

In this paper, we aim to understand the power of plug-in approach in the reward-free exploration with linear function approximation. We study the linear mixture MDPs, where the transition probability kernel is a linear mixture of a number of basis kernels (Ayoub et al., 2020; Zhou et al., 2020b;a). We first build an empirical model with an estimation of the transition dynamics in the exploration phase, and then find a near-optimal policy by planning with the empirical model via arbitrary plug-in solver in the planning phase. Our setting is different from the reward-free exploration with linear function approximation without plug-in model (Wang et al., 2020; Zanette et al., 2020b), in which the agent can directly observe all history samples and design specialized model-free algorithm in the planning phase. Compared to our results, the standard plug-in approach (Agarwal et al., 2020; Cui & Yang, 2020) requires reward signals to guide exploration and may not predict an accurate value function under an arbitrarily given reward function after the exploration phase.

Our results shows that the plug-in approach can achieve near-optimal sample complexity in the reward-free setting. In particular, we proposed a statistically efficient algorithm for reward-free exploration. Our algorithm samples $\tilde{O}(d^2H^4/\epsilon^2)$ trajectories during the exploration phase, which suffices to obtain $O(\epsilon)$-optimal policies for an arbitrary reward function with an ϵ-optimal pluging solver in the planning phase. Here d is the feature dimension, and H is the planning horizon. Furthermore, with a more refined trajectory-wise uncertainty estimation, we further improve the sample complexity bound to $\tilde{O}\left(d^2H^3/\epsilon^2\right)$ in the regime where $d > H$ and $\epsilon \leq H/\sqrt{d}$. This matches the lower bound in Jin et al. (2020a), $\Omega(d^2H^3/\epsilon^2)$, which indicates that our upper bound is near-optimal except for logarithmic factors. To the best of our knowledge, this is the first work that obtains minimax sample complexity bounds for the plug-in approach in reward-free exploration with linear function approximation.

2 Related Work

RL with Linear Function Approximation Reinforcement learning with linear function approximation has been widely studied in the recent few years (e.g. Jiang et al. (2017); Yang & Wang (2019; 2020); Jin et al. (2020b); Modi et al. (2020); Du et al. (2019); Zanette et al. (2020a); Cai et al. (2020); Ayoub et al. (2020); Weisz et al. (2021); Zhou et al. (2020b;a)). The linear mixture MDPs model studied in our work assumes the transition probability function is parameterized as a linear function of a given feature mapping over state-action-next-state triple (Ayoub et al., 2020; Zhou et al., 2020b;a). Based on the Berstein inequality for vector-valued martingales, Zhou et al. (2020a) proposed an efficient algorithm that obtains minimax regret in the regime where $d > H$. Besides linear mixture MDPs, linear MDPs is another category of RL with linear function approximation, which assumes both the transition probability function and reward function are parameterized as a linear function of a given feature mapping over state-action pairs. The algorithms with best regret bound are proposed by Jin et al. (2020b) and Yang & Wang (2020), which study model-free algorithm and model-based algorithm respectively. The minimax regret bound for linear MDPs is still unclear.

Reward-Free Reinforcement Learning In contrast to the standard RL setting, reward-free reinforcement learning separates the exploration problem and the planning problem which allows one to handle them in a theoretically principled way. For tabular setting, reward-free reinforcement learning has been well-exploited in many previous results (Jin et al., 2020a; Kaufmann et al., 2021; Ménard et al., 2020), where the minimax rate is obtained by Ménard et al. (2020). For reward-free exploration with linear function approximation, Wang et al. (2020) proposed the first efficient algorithm that obtains $O(d^2H^6/\epsilon^2)$ sample complexity for linear MDPs. However, their algorithm is model-free in nature and cannot guarantee good performance with any plug-in solver. Besides, their sample complexity bound is not tight, which can be further improved to $O(d^2H^4/\epsilon^2)$ with our technique (Please see Section 4.2 for the detailed discussion).

Plug-in Approach The plug-in approach has been studied in tabular/linear case in restrictive settings. E.g., Agarwal et al. (2020) and Cui & Yang (2020) studied the standard plug-in approach with a generative model, where the algorithm is allowed to query the outcome of any state action pair from an oracle. They showed the plug-in approach also achieved the minimax optimal sample complexity to find an ϵ-optimal policy in both tabular MDPs and linear MDPs. The reward-free algorithms proposed by Jin et al. (2020a) are model-based in nature, thus can be regarded as a solution

in the plug-in solver setting. However, their algorithms are restricted to the tabular case and cannot be applied to the setting with linear function approximation.

3 PRELIMINARIES

3.1 EPISODIC MDPS

We consider the setting of episodic Markov decision processes (MDPs), which can be denoted by a six-tuple $(\mathcal{S}, \mathcal{A}, P, R, H, \nu)$, where $\mathcal{S}$ is the set of states, $\mathcal{A}$ is the action set, P is the transition probability matrix so that $P_h(\cdot|s,a)$ gives the distribution over states if action a is taken on state s at step h, $R_h(s,a)$ is the deterministic reward function of taking action a on state s with support $[0,1]$ in step h, H is the number of steps in each episode, and ν is the distribution of the initial state.

In episode k, the agent starts from an initial state $s_{k,1}$ sampled from the distribution ν. At each step $h \in [H]$, the agent observes the current state $s_{k,h} \in \mathcal{S}$, takes action $a_{k,h} \in \mathcal{A}$, receives reward $R_h(s_{k,h}, a_{k,h})$, and transits to state $s_{k,h+1}$ with probability $P_h(s_{k,h+1}|s_{k,h}, a_{k,h})$. The episodes ends when s_{H+1} is reached.

A deterministic policy π is a collection of H policy functions $\{\pi_h : \mathcal{S} \to \mathcal{A}\}_{h\in[H]}$. We use Π to denote the set of all deterministic policies. For a specific reward function R, we use $V_h^\pi : \mathcal{S} \times R \to \mathbb{R}$ to denote the value function at step h under policy π w.r.t reward R, which gives the expected sum of the remaining rewards received under policy π starting from $s_h = s$, i.e.

$$V_h^\pi(s,R) = \mathbb{E}\left[\sum_{h'=h}^{H} R\left(s_{h'}, \pi_{h'}\left(s_{h'}\right)\right) \mid s_h = s, P\right].$$

Accordingly, we define $Q_h^\pi(s,a,R)$ as the expected Q-value function at step h:

$$Q_h^\pi(s,a,R) = \mathbb{E}\left[R\left(s_h, a_h\right) + \sum_{h'=h+1}^{H} R\left(s_{h'}, \pi_{h'}\left(s_{h'}\right)\right) \mid s_h = s, a_h = a, P\right].$$

We use π_R^* to denote the optimal policy w.r.t. reward R, and we use $V_h^*(\cdot, R)$ and $Q_h^*(\cdot,\cdot,R)$ to denote the optimal value and Q-function under optimal policy π_R^* at step h. We say a policy π is ϵ-optimal w.r.t R if

$$\mathbb{E}\left[\sum_{h=1}^{H} R_h\left(s_h, a_h\right) \mid \pi\right] \geq \mathbb{E}\left[\sum_{h=1}^{H} R_h\left(s_h, a_h\right) \mid \pi_R^*\right] - \epsilon.$$

For the convenience of explanation, we assume the agent always starts from the same state s_1 in each episode. It is straightforward to extend to the case with stochastic initialization, by adding a initial state s_0 with no rewards and only one action a_0, and the transition probability of (s_0, a_0) is the initial distribution μ. We use $P_h V(s,a,R)$ as a shorthand of $\sum_{s'} P_h(s'|s,a)V(s',R)$.

3.2 LINEAR MIXTURE MDPS

We study a special class of MDPs called linear mixture MDPs, where the transition probability kernel is a linear mixture of a number of basis kernels (Ayoub et al., 2020; Zhou et al., 2020b;a). This model is defined as follows in the previous literature: Let $\phi(s,a,s') : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{R}^d$ be a feature mapping satisfying that for any bounded function $V : \mathcal{S} \to [0,1]$ and any tuple $(s,a) \in \mathcal{S} \times \mathcal{A}$, we have

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

Linear mixture MDPs is defined as follows:

Definition 1. *(Ayoub et al., 2020) An MDP is called a linear mixture MDP if there exists a parameter vector $\theta_h \in \mathbb{R}^d$ with $\|\theta_h\|_2 \leq B$ for a constant B and a feature vector $\phi(\cdot,\cdot,\cdot)$ satisfying 1, such that $P_h(s'|s,a) = \theta_h^\top \phi(s,a,s')$ for any state-action-next-state triplet $(s,a,s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ and step $h \in [H]$.*

3.3 REWARD-FREE REINFORCEMENT LEARNING

We study the problem of reward-free exploration with plug-in solver. Our setting is different from the reward-free exploration setting studied in the previous literature. Formally, there are two phases in this setting: exploration phase and planning phase.

During the exploration phase, the agent interacts with the environment for K episodes. In episode k, the agent chooses a policy π_k which induces a trajectory. The agent observes the states and actions $s_{k,1}, a_{k,1}, \cdots, s_{k,H}, a_{k,H}$ as usual, but does not observe any rewards. After K episodes, the agent calculates the estimated model $\{\tilde{P}_h = \tilde{\theta}_h^\top \phi\}_{h\in[H]}$, which will be used in the planning phase to calculate the optimal policy.

During the planning phase, the agent is no longer allowed to interact with the MDP. Also, it cannot directly observe the history samples obtained in the exploration phase. Instead, the agent is given a set of reward function $\{R_h\}_{h\in[H]}$, where $R_h : \mathcal{S} \times \mathcal{A} \to [0, 1]$ is the deterministic reward in step h. We define $\hat{V}_h^{\pi,\tilde{P}}(s, R)$ as the value function of transition $\tilde{P}$ and reward R, i.e.

$$\hat{V}_h^{\pi,\tilde{P}}(s, R) = \mathbb{E}\left[\sum_{h'=h}^{H} R\left(s_{h'}, \pi_{h'}\left(s_{h'}\right)\right) \mid s_h = s, \tilde{P}\right].$$

In the planning phase, the agent calculates the optimal policy $\hat{\pi}_R$ with respect to the reward function R in the estimated model $\tilde{P}$ using any ϵ_{opt}-optimal model-based solver. That is, the returned policy $\hat{\pi}_R$ satisfies:

$$\hat{V}_1^{\star,\tilde{P}}(s_1, R) - \hat{V}_1^{\hat{\pi}_R,\tilde{P}}(s_1, R) \leq \epsilon_{\text{opt}}.$$

The agent's goal is to output an accurate model estimation $\tilde{P}$ after the exploration phase, so that the policy $\hat{\pi}_R$ calculated in planning phase can be $\epsilon + \epsilon_{\text{opt}}$-optimal w.r.t. any reward function R. Compared with the reward-free setting studied in the previous literature (Jin et al., 2020a; Wang et al., 2020; Ménard et al., 2020; Kaufmann et al., 2021), the main difference is that we require that the algorithm maintains an model estimation instead of all the history samples after the exploration phase, and can use any model-based solver to calculate the near-optimal policy in the planning phase.

4 REWARD-FREE RL WITH PLUG-IN SOLVER

In this section, we introduce our reward-free RL algorithm for linear mixture MDPs with plug-in solver.

4.1 ALGORITHM

The exploration phase of the algorithm is presented in Algorithm 1. Recall that for a given value function $\tilde{V}_{k,h+1}$, we have $P_h\tilde{V}_{k,h+1}(s_{k,h}, a_{k,h}) = \theta_h^\top \phi_{\tilde{V}_{k,h+1}}(s_{k,h}, a_{k,h})$ for any k, h. Therefore, $\tilde{V}_{k,h+1}(s_{k,h+1})$ and $\phi_{\tilde{V}_{k,h+1}}(s_{k,h}, a_{k,h})$ can be regarded as the stochastic reward and the linear feature of a linear bandits problem with linear parameter θ_h. We employ the standard least-square regression to learn the underlying parameter θ_h. In each episode, we first update the estimation $\hat{\theta}_k$ based on history samples till episode $k-1$. We define the auxiliary rewards $R_{k,h}$ to guide exploration. We calculate the optimistic Q-function using the parameter estimation $\hat{\theta}_k$, and then execute the greedy policy with respect to the updated Q-function to collect new samples.

The main problem is how to define the exploration-driven reward $R_{k,h}(s, a)$, which measures the uncertainty for state-action pair (s, a) at the current step. In the setting of linear mixture MDPs, the linear feature $\phi_V(s, a)$ is a function of both state-action pair (s, a) and the next-step value function V. This is not a big deal in the standard exploration setting (Ayoub et al., 2020; Zhou et al., 2020a). However, in the reward-free setting, since the reward function is not given beforehand, we need to upper bound the estimation error of value functions for *any* possible rewards. To tackle this problem, we use $\max_{V\in\mathcal{V}} \beta\|\phi_V(s, a)\|_{\Lambda_{k,h}^{-1}} = \max_{V\in\mathcal{V}} \beta\sqrt{\phi_V(s, a)^\top \left(\Lambda_{k,h}\right)^{-1} \phi_V(s, a)}$ as a

Algorithm 1 Reward-free Exploration: Exploration Phase

Input: Failure probability $\delta > 0$ and target accuracy $\epsilon > 0$
$K \leftarrow C_1 d^2 H^4 \log(Hd\delta^{-1}\epsilon^{-1}))/\epsilon^2$ for some constant C_1
$\beta \leftarrow H\sqrt{d\log(4HK\delta^{-1})}$, $\mathcal{V} \leftarrow \{V : \mathcal{S} \to [0, H]\}$
for episode $k = 1, 2, \cdots, K$ **do**
  $Q_{k,H+1}(\cdot,\cdot) = 0$, $V_{k,H+1}(\cdot) = 0$
  for step $h = H, H-1, \cdots, 1$ **do**
    $\Lambda_{k,h} \leftarrow \sum_{t=1}^{k-1} \phi_{t,h}(s_{t,h}, a_{t,h})\phi_{t,h}(s_{t,h}, a_{t,h})^\top + I$
    $\hat{\theta}_{k,h} \leftarrow (\Lambda_{k,h})^{-1} \sum_{t=1}^{k-1} \phi_{t,h}(s_{t,h}, a_{t,h})\tilde{V}_{t,h+1,s,a}(s_{t,h+1})$
    $\tilde{V}_{k,h+1,s,a} \leftarrow \arg\max_{V\in\mathcal{V}} \|\sum_{s'} \phi(s,a,s')V(s')\|_{(\Lambda_{k,h})^{-1}}, \forall s, a$
    $\phi_{k,h}(s,a) \leftarrow \sum_{s'} \phi(s,a,s')\tilde{V}_{k,h,s,a}(s'), \forall s, a$
    $u_{k,h}(s,a) \leftarrow \beta\sqrt{\phi_{k,h}(s,a)^\top (\Lambda_{k,h})^{-1} \phi_{k,h}(s,a)}, \forall s, a$
    Define the exploration-driven reward function $R_{k,h}(s,a) = u_{k,h}(s,a), \forall s, a$
    $Q_{k,h}(s,a) \leftarrow \min\left\{\hat{\theta}_{k,h}^\top \left(\sum_{s'} \phi(s,a,s')V_{k,h+1}(s')\right) + R_{k,h}(s,a) + u_{k,h}(s,a), H\right\}$
    $V_{k,h}(s) \leftarrow \max_{a\in\mathcal{A}} Q_{k,h}(s,a)$, $\pi_{k,h}(s) = \arg\max_{a\in\mathcal{A}} Q_{k,h}(s,a)$
  end for
  for step $h = 1, 2, \cdots, H$ **do**
    Take action $a_{k,h} = \pi_{k,h}(s_{k,h})$ and observe $s_{k,h+1} \sim P_h(s_{k,h}, a_{k,h})$
  end for
end for
Find $\tilde{P}_h$ such that the transition $\tilde{P}_h(\cdot|\cdot,\cdot) = \tilde{\theta}_h^\top \phi(\cdot,\cdot,\cdot)$ is well-defined and $\left\|\tilde{\theta}_h - \hat{\theta}_{K,h}\right\|_{\Lambda_{K,h}} \leq \beta$ for $h \in [H]$
Output: $\{\tilde{P}_h\}_{h=1}^H$

measure of the maximum uncertainty for the state-action pair (s, a), where $\mathcal{V} = \{V : \mathcal{S} \to [0, H]\}$ is the set of all possible value functions, and $\Lambda_{k,h}$ is the summation of all the history samples $\{s_{t,h}, a_{t,h}, s_{t,h+1}\}_{t=1}^{k-1}$ with feature $\phi_{t,h}(s_{t,h}, a_{t,h}) = \max_{V\in\mathcal{V}} \beta\|\phi_V(s,a)\|_{\Lambda_{t,h}^{-1}}$. In each episode, we define $R_{k,h}(s,a) = u_{k,h}(s,a) = \max_{V\in\mathcal{V}} \beta\|\phi_V(s,a)\|_{\Lambda_{k,h}^{-1}}$, where $R_{k,h}(s,a)$ is the exploration-driven reward used to guide exploration, and $u_{k,h}(s,a)$ is the additional bonus term which helps to guarantee that $Q_{k,h}(s,a)$ is an optimistic estimation. After sufficient number of episodes, the uncertainty along all possible trajectories should be low in average, since otherwise the agent would have visited those state-action pairs with large uncertainty as guided by the reward function. Finally, the algorithm returns the model estimation $\{\tilde{P}_h\}_h$ that is well defined (i.e. $\sum_{s'} \tilde{P}_h(s'|s,a) = 1, \tilde{P}_h(s'|s,a) \geq 0, \forall s, a, s'$), and satisfy the constraints $\left\|\tilde{\theta}_h - \hat{\theta}_{K,h}\right\|_{\Lambda_{K,h}} \leq \beta$.

4.2 Analysis

In this subsection, we present the sample complexity bound of our Algorithm 1.

Theorem 1. *After collecting $K = \frac{C_1 d^2 H^4 \log(Hd\delta^{-1}\epsilon^{-1}))}{\epsilon^2}$ trajectories for some constant C_1, Algorithm 1 returns a transition model $\tilde{P}$. Then with probability at least $1-\delta$, for any given reward, a policy returned by any ϵ_{opt}-optimal plug-in solver on $(S, A, \tilde{P}, R, H, \nu)$ is $O(\epsilon + \epsilon_{\text{opt}})$-optimal for the true MDP, (S, A, P, R, H, ν).*

By the reduction from the tabular reward-free RL problem (Jin et al., 2020a), we know that the sample complexity lower bound for reward-free RL in the linear mixture model setting is $\Omega(\frac{d^2H^3}{\epsilon^2})$[1].

[1]The regret lower bound proved by Jin et al. (2020a) has quadratic dependence on H, while the lower bound in our setting is $\Omega(H^3)$. This additional linear dependence is because that we consider time-inhomogeneous MDPs (P_h varies among diffrent step h), instead of time-homogeneous MDPs studied in Jin et al. (2020a). Similar phenomenon has also been observed by many previous works (e.g. Jin et al. (2018); Ménard et al. (2020)).

Compared with this bound, our result matches the lower bound w.r.t the dimension d and the precision ϵ except for logarithmic factors.

There is also a recent paper of Wang et al. (2020) studying reward-free exploration in the setting of linear MDPs. Their sample complexity is $\tilde{O}(d^2H^6/\epsilon^2)$, thus our bound is better than theirs by a factor of H^2. This is achieved by setting the exploration-driven reward $R_{k,h}(s,a) = u_{k,h}(s,a)$ and a more careful analysis of the value estimation error, while they set $R_{k,h}(s,a) = u_{k,h}(s,a)/H$ in their algorithm. Though the setting is different, we find that our parameter choice and the more refined analysis is applicable to their setting, which can help to further improve their bound by a factor of H^2. Besides, the definition of the uncertainty $u_{k,h}(s,a)$ in our algorithm is more complicated than theirs due to the different linear function approximation settings.

Now we briefly explain the main idea in the proof. The detailed explanation is deferred to Appendix A. Firstly, we introduce the value function $\tilde{V}_h^*(s,R)$, which is recursively defined from step $H+1$ to step 1:

$$\tilde{V}_{H+1}^*(s,R) = 0, \forall s \in \mathcal{S} \tag{2}$$

$$\tilde{V}_h^*(s,R) = \max_{a\in\mathcal{A}} \left\{ \min \left\{ R_h(s,a) + P_h \tilde{V}_{h+1}^*(s,a,R), H \right\} \right\}, \forall s \in \mathcal{S}, h \in [H] \tag{3}$$

Compared with the definition of $V_h^*(s,R)$, the main difference is that we take minimization over the value and H at each step. We state the following lemma, which gives an upper bound on the sub-optimality gap of $\hat{\pi}$ in the planning phase.

Lemma 1. *With probability at least $1-\delta$, the sub-optimality gap of the policy $\hat{\pi}_R$ for any reward function R in the planning phase can be bounded by*

$$V_1^*(s_1,R) - V_1^{\hat{\pi}_R}(s_1,R) \le 4\tilde{V}_1^*(s_1,R_K) + \epsilon_{\text{opt}}. \tag{4}$$

This lemma connects the sub-optimality gap with the value function of the auxiliary reward in episode K. So the remaining problem is how to upper bound $\tilde{V}_1^*(s_1,R_K)$. Since the exploration-driven reward R_k is non-increasing w.r.t k, it is not hard to prove that $K\tilde{V}_1^*(s_1,R_K) \le \sum_{k=1}^K \tilde{V}_1^*(s_1,R_k)$. We use the following two lemmas to upper bound $\sum_{k=1}^K \tilde{V}_1^*(s_1,R_k)$.

Lemma 2. *With probability at least $1-\delta$, $\tilde{V}_h^*(s,R_k) \le V_{k,h}(s)$ holds for any $(s,a) \in \mathcal{S}\times\mathcal{A}, h \in [H]$ and $k \in [K]$.*

Lemma 3. *With probability at least $1-\delta$, $\sum_{k=1}^K V_{k,1}(s_1) \le 6dH^2\sqrt{K\log(2HK/\delta)}$.*

With the help of the optimistic bonus term $u_{k,h}$, we can prove that the estimation value $V_{k,h}$ is always optimistic w.r.t $\tilde{V}_h^*$, which is illustrated in Lemma 2. Therefore, we have $\sum_{k=1}^K \tilde{V}_1^*(s_1,R_k) \le \sum_{k=1}^K V_{k,1}(s_1)$. By adapting the regret analysis in the standard RL setting to the reward-free setting, we can upper bound the summation of $V_{k,1}(s_1)$ in Lemma 3.

5 Improving the dependence on H

In Section 4, we propose an efficient algorithm with $O(d^2H^4/\epsilon^2)$ sample complexity. There is still a gap of H between this result and the lower bound. In this section, we close this gap with a careful maximum total-variance bound using Bernstein inequality.

In the previous results studying regert minimization in online RL setting (Azar et al., 2017; Zhou et al., 2020a; Wu et al., 2021), one commonly-used approach to obtain the minimax rate is to upper bound the regret using the total variance of the value function by the Bernstein's concentration inequalities, and finally bound the summation of the one-step transition variance by the law of total variance (Lattimore & Hutter, 2012; Azar et al., 2013; 2017). However, the situation becomes much more complicated in the reward-free setting with linear function approximation. Recall that $\tilde{V}_{k,h+1,s,a}(s')$ defined in Line 9 of Algorithm 1 is the next-step value function that maximizes the uncertainty $\|\phi_V(s,a)\|_{\Lambda_{k,h}^{-1}}$ for state-action pair (s,a). One naive approach is to still use $\max_V \|\phi_V(s,a)\|_{\Lambda_{k,h}^{-1}}$ as the uncertainty measure for (s,a) and upper bound the error rate by the summation of the one-step transition variance of $\tilde{V}_{k,h+1,s_{k,h},a_{k,h}}(s')$. However, we can not upper

bound the variance summation of $\tilde{V}_{k,h+1,s_{k,h},a_{k,h}}(s')$ in H steps by $O(H^2)$ similarly by the law of total variance, since $\{\tilde{V}_{k,h,s_{k,h},a_{k,h}}\}_{h\in[H]}$ is not the value functions induced from the same policy and transition dynamics. To tackle the above problem, we need to define the exploration-driven reward and the confidence bonus for each state-action pair in a more refined way. Our basic idea is to carefully measure the expected total uncertainty along the whole trajectories w.r.t. each MDPs in the confidence set. We will explain the detail in the following subsections.

5.1 ALGORITHM

Algorithm 2 Reward-free Exploration: Exploration Phase

Input: Failure probability $\delta > 0$ and target accuracy $\epsilon > 0$
$K \leftarrow \frac{C_2(d^2H^3+dH^4)\log^3(Hd\delta^{-1}\epsilon^{-1})}{\epsilon^2} + \frac{C_2(d^{2.5}H^2+d^2H^3)\log^3(Hd\delta^{-1}\epsilon^{-1}))}{\epsilon}$ for some $C_2 > 0$
$\hat{\beta} = 16\sqrt{d}\log(32KH/\delta)$, $\check{\beta} = 16d\log(32KH/\delta)$, $\tilde{\beta} = 16H^2\sqrt{d}\log(32KH/\delta)$
Set $\Lambda_{i,k,h} \leftarrow I$, $\hat{\theta}_{i,k,h} \leftarrow \mathbf{0}$ for $k = 1, h \in [H], i = 1, 2, 3, 4, 5$
Set $\mathcal{U}_{1,h}$ to be the set containing all the $\tilde{\theta}_h$ that makes $\tilde{P}_h$ well-defined, $h \in [H]$.
for episode $k = 1, 2, \cdots, K$ **do**
  Calculate $\pi_k, \tilde{\theta}_k, R_k = \arg\max_{\pi,\tilde{\theta}_h\in\mathcal{U}_{k,h},R} V_{k,1}^{\pi,\tilde{P}}(s_1, R)$, where V is defined in Eqn 9.
  for step $h = 1, 2, \cdots, H$ **do**
    Take action according to the policy $\pi_{k,h}$ and observe $s_{k,h+1} \sim P_h(\cdot|s_{k,h}, a_{k,h})$
  end for
  for step $h = 1, 2, \cdots, H$ **do**
    Update $\{\Lambda_{i,k+1,h}\}_{i=1}^5$ using Eqn 63, 64, 67, 68 and 73
    Update the model estimation $\{\hat{\theta}_{i,k+1,h}\}_{i=1}^5$ using Eqn 65, 66, 69, 70 and 74
    Add the constraints (Eqn 5) to the confidence set $\mathcal{U}_{k,h}$, and obtain $\mathcal{U}_{k+1,h}$
  end for
end for
Calculate $\tilde{P}_{K,h}(\cdot|\cdot,\cdot) = \tilde{\theta}_{K,h}^\top \phi(\cdot,\cdot,\cdot)$ for $h \in [H]$
Output: $\left\{\tilde{P}_{K,h}\right\}_{h=1}^H$.

Our algorithm is described in Algorithm 2. At a high level, Algorithm 2 maintains a high-confidence set $\mathcal{U}_{k,h}$ for the real parameter θ_h in each episode k, and calculates an optimistic value function $V_{k,1}^{\pi,\tilde{P}}(s_1, R)$ for any reward function R and transition $\tilde{P}_h = \tilde{\theta}_h^\top \phi$ with $\tilde{\theta}_h \in \mathcal{U}_{k,h}$. Roughly speaking, the value function $V_{k,1}^{\pi,\tilde{P}}(s_1, R)$ measures the expected uncertainty along the whole trajectory induced by policy π in the MDP with transition $\tilde{P}$ and reward R. To collect more "informative" samples and minimize the worst-case uncertainty over all possible transition dynamics and reward function, we calculate $\pi_k = \arg\max_\pi \max_{\tilde{\theta}_h\in\mathcal{U}_{k,h},R} V_{k,1}^{\pi,\tilde{P}}(s_1, R)$, and execute the policy π_k to collect more data in episode k. To ensure that the model estimation $\tilde{\theta}_{k+1,h}$ is close to the true model θ_h w.r.t. features $\phi_V(s, a)$ of different value functions V, we use the samples collected so far to calculate five model estimation $\{\hat{\theta}_{i,k,h}\}_{i=1}^5$ and the corresponding constraints at the end of the episode k. Each constraint is an ellipsoid in the parameter space centered at the parameter estimation $\hat{\theta}_{i,k,h}$ with covariance matrix $\Lambda_{i,k,h}$ and radius β_i, i.e.

$$\|\tilde{\theta}_h - \hat{\theta}_{i,k,h}\|_{\Lambda_{i,k,h}} \leq \beta_i. \tag{5}$$

We update $\mathcal{U}_{k+1,h}$ by adding these constraints to the the confidence set $\mathcal{U}_{k,h}$.

The remaining problems are how to define the value function $V_{k,1}^{\pi,\tilde{P}}(s_1, R)$ that represents the expected uncertainty for policy π, transition $\tilde{P}$ and reward R, and how to update the model estimation $\theta_{i,k,h}$ and the confidence set $\mathcal{U}_{k,h}$.

Uncertainty Measure Instead of using $\max_V \|\phi_V(s, a)\|_{\Lambda_{k,h}^{-1}}$ to measure the maximum uncertainty for the state-action pair (s, a) in Algorithm 1, we separately define the uncertainty along the

trajectories induced by different policy π, reward function R and transition dynamics $\tilde{P}_h = \tilde{\theta}_h^\top \phi$. Specifically, Recall that $\hat{V}_h^{\pi,\tilde{P}}(s,R)$ is the value function of policy π in the MDP model with transition $\tilde{P}$ and reward R. We define the following exploration-driven reward for transition $\tilde{P}$, reward R and policy π:

$$u_{1,k,h}^{\pi,\tilde{P}}(s,a,R) = \hat{\beta} \left\| \sum_{s'} \phi(s,a,s') \hat{V}_{h+1}^{\pi,\tilde{P}}(s',R) \right\|_{(\Lambda_{1,k,h})^{-1}}. \tag{6}$$

Suppose $\tilde{V}_{k,H+1}^{\pi,\tilde{P}}(s,R) = 0, \forall s \in \mathcal{S}$, we define the corresponding value function recursively from step $H+1$ to step 1. That is,

$$\tilde{V}_{k,h}^{\pi,\tilde{P}}(s,R) = \min\left\{ u_{1,k,h}^{\pi,P}(s,\pi_h(s),R) + P_h \tilde{V}_{k,h+1}^{\pi,\tilde{P}}(s,\pi_h(s),R), H \right\}, \forall s \in \mathcal{S}. \tag{7}$$

$\tilde{V}_{k,h}^{\pi,\tilde{P}}(s,R)$ can be regarded as the expected uncertainty along the trajectories induced by policy π for the value function $\hat{V}_h^{\pi,\tilde{P}}(s,R)$. In each episode k, we wish to collect samples by executing the policy π that maximizes the uncertainty measure $\max_{\tilde{\theta} \in \mathcal{U}_{k,h},R} \tilde{V}_{k,1}^{\pi,\tilde{P}}(s,R)$. However, the definition of $\tilde{V}_{k,1}^{\pi,\tilde{P}}(s,R)$ explicitly depends on the real transition dynamics P, which is unknown to the agent. To solve this problem, we construct an optimistic estimation of $\tilde{V}_{k,h}^{\pi,\tilde{P}}(s,R)$. We define $V_{k,h}^{\pi,\tilde{P}}(s,R)$ from step $H+1$ to step 1 recursively. Suppose $V_{k,H+1}^{\pi,\tilde{P}}(s,R) = 0$. We calculate

$$u_{2,k,h}^{\pi,\tilde{P}}(s,a,R) = \hat{\beta} \left\| \sum_{s'} \phi(s,a,s') V_{k,h+1}^{\pi,\tilde{P}}(s',R) \right\|_{(\Lambda_{2,k,h})^{-1}}, \tag{8}$$

$$V_{k,h}^{\pi,\tilde{P}}(s,R) = \min\left\{ u_{1,k,h}^{\pi,\tilde{P}}(s,\pi_h(s),R) + u_{2,k,h}^{\pi,\tilde{P}}(s,\pi_h(s),R) + \tilde{P}_h V_{k,h+1}^{\pi,\tilde{P}}(s,\pi_h(s),R), H \right\}, \tag{9}$$

where $u_{2,k,h}^{\tilde{P},\pi}$ is the confidence bonus which ensures that the optimism $V_{k,h}^{\pi,\tilde{P}}(s,R) \geq \tilde{V}_{k,h}^{\pi,\tilde{P}}(s,R)$ holds with high probability. After calculating $V_{k,h}^{\pi,\tilde{P}}(s,R)$, we take maximization over all $\tilde{\theta}_h \in \mathcal{U}_{k,h}$ and reward function R, and calculate $\pi_k = \arg\max_\pi \max_{\tilde{\theta}_h \in \mathcal{U}_{k,h},R} V_{k,1}^{\pi,\tilde{P}}(s_1,R)$. We execute π_k to collect more samples in episode k.

Weighted Ridge Regression and Confidence Set In Algorithm 2, we maintain five model estimation $\{\hat{\theta}_{i,k,h}\}_{i=1}^5$ and add five constraints to the confidence set $\mathcal{U}_{k,h}$ in each episode. These constraints can be roughly classified into three categories. The first and the third constraints are applied to ensure that $\tilde{P}_{k,h}\hat{V}_{h+1}^{\pi,\tilde{P}}(s,a,R)$ is an accurate estimation of $P_{k,h}\hat{V}_{h+1}^{\pi,\tilde{P}}(s,a,R)$ for any reward function R and $\tilde{P}$ satisfying $\tilde{\theta}_h \in \mathcal{U}_{k,h}$, while the second and the forth constraints are used to guarantee that $\tilde{P}_{k,h}V_{h+1}^{\pi,\tilde{P}}(s,a,R)$ is an accurate estimation of $P_h V_{h+1}^{\pi,\tilde{P}}(s,a,R)$ for any reward function R and $\tilde{P}$ satisfying $\tilde{\theta}_h \in \mathcal{U}_{k,h}$. The last constraint is applied due to technical issue and will be explained in Appendix B.1. In this subsection, we introduce the basic idea behind the construction of the first and the third constraints. The second and the forth constraints follow the same idea but consider the different value function $V_{h+1}^{\pi,\tilde{P}}(s,a,R)$. The formal definition of the parameters $\{\hat{\theta}_{i,k,h}\}_{i=1}^5$ and the constraints are deferred to Appendix B.1.

For notation convenience, we use $V_{k,h}(s)$ as a shorthand of $\hat{V}_{k,h}^{\pi_k,\tilde{P}_k}(s,R_k)$ in this part. The construction of our confidence sets is inspired by a recent algorithm called UCRL-VTR$^+$ proposed by Zhou et al. (2020a). Recall that we use $(s_{k,h}, a_{k,h})$ to denote the state-action pair that the agent encounters at step h in episode k. We use the following ridge regression estimator to calculate the

corresponding $\hat{\theta}_{1,k,h}$ in episode k:

$$\hat{\theta}_{1,k,h} = \arg\min_{\theta\in\mathbb{R}^d} \lambda\|\theta\|_2^2 + \sum_{t=1}^{k-1} \left[\theta^\top \phi_{V_{t,h+1}/\bar{\sigma}_{1,t,h}}(s_{t,h}, a_{t,h}) - V_{t,h+1}(s_{t,h+1}, R_t)/\bar{\sigma}_{1,t,h}\right]^2, \quad (10)$$

where $\bar{\sigma}_{1,k,h}^2 = \max\left\{H^2/d, [\bar{\mathbb{V}}_{k,h}V_{k,h+1}](s_{k,h}, a_{k,h}) + E_{k,h}\right\}$ is an optimistic estimation of the one-step transition variance:

$$\mathbb{V}_h V_{k,h+1}(s_{k,h}, a_{k,h}) = \mathbb{E}_{s'\sim P_h(\cdot|s_{k,h},a_{k,h})}\left[\left(V_{k,h+1}(s') - P_h V_{k,h}(s_{k,h}, a_{k,h})\right)^2\right].$$

In the definition of $\bar{\sigma}_{1,k,h}^2$, $\bar{\mathbb{V}}_{k,h}V_{k,h+1}(s_{k,h}, a_{k,h})$ is an empirical estimation for the variance $\mathbb{V}_h V_{k,h+1}(s_{k,h}, a_{k,h})$, and $E_{k,h}$ is a bonus term defined in Eqn. 57 which ensures that $\bar{\sigma}_{k,h}$ is an optimistic estimation of $\mathbb{V}_h V_{k,h+1}(s_{k,h}, a_{k,h})$. One technical issue here is how to estimate the variance $\mathbb{V}_h V_{k,h+1}(s_{k,h}, a_{k,h})$. Recall that by definition,

$$\mathbb{V}_h V_{k,h+1}(s_{k,h}, a_{k,h}) = P_h V_{k,h+1}^2(s_{k,h}, a_{k,h}) - [P_h V_{k,h+1}(s_{k,h}, a_{k,h})]^2 \quad (11)$$

$$= \theta_h^\top \phi_{V_{k,h+1}^2}(s_{k,h}, a_{k,h}) - [\theta_h^\top \phi_{V_{k,h+1}}(s_{k,h}, a_{k,h})]^2. \quad (12)$$

We use $\tilde{\theta}_{k,h}^\top \phi_{V_{k,h+1}^2}(s_{k,h}, a_{k,h}) - \left[\tilde{\theta}_{k,h}^\top \phi_{V_{k,h+1}}(s_{k,h}, a_{k,h})\right]^2$ as our variance estimator $\bar{\mathbb{V}}_{k,h}V_{k,h+1}(s_{k,h}, a_{k,h})$, where $\tilde{\theta}_{k,h} \in \mathcal{U}_{k,h}$ is the parameter which maximizes the value $V_{k,1}^{\pi,\tilde{P}}(s_1, R)$ in episode k. To ensure that $\bar{\mathbb{V}}_{k,h}V_{k,h+1}(s_{k,h}, a_{k,h})$ is an accurate estimator, we maintain another parameter estimation $\hat{\theta}_{3,k,h}$ using history samples w.r.t. the feature $\phi_{V_{k,h+1}^2}(s_{k,h}, a_{k,h})$.

$$\hat{\theta}_{3,k,h} = \arg\min_{\theta\in\mathbb{R}^d} \lambda\|\theta\|_2^2 + \sum_{t=1}^{k-1} \left[\theta^\top \phi_{V_{t,h+1}^2}(s_{t,h}, a_{t,h}) - V_{t,h+1}^2(s_{t,h+1})\right]^2. \quad (13)$$

After calculating $\hat{\theta}_{1,k,h}$ and $\hat{\theta}_{3,k,h}$, we add the first and the third constraints (Eqn. 5) to the confidence set $\mathcal{U}_{k,h}$, where $\Lambda_{1,k,h}$ and $\Lambda_{3,k,h}$ is the corresponding covariance matrix of all history samples, i.e. $\Lambda_{1,k,h} = \sum_{t=1}^{k-1} \bar{\sigma}_{1,t,h}^{-2} \phi_{V_{t,h+1}}^\top(s_{t,h}, a_{t,h})\phi_{V_{t,h+1}}(s_{t,h}, a_{t,h})$ and $\Lambda_{3,k,h} = \sum_{t=1}^{k-1} \phi_{V_{t,h+1}^2}^\top(s_{t,h}, a_{t,h})\phi_{V_{t,h+1}^2}(s_{t,h}, a_{t,h})$.

5.2 Regret

We present the regret upper bound of Algorithm 2 in Theorem 2. In the regime where $d \geq H$ and $\epsilon \leq H/\sqrt{d}$, we can obtain $\tilde{O}(d^2H^3/\epsilon^2)$ sample complexity upper bound, which matches the sample complexity lower bound except logarithmic factors.

Theorem 2. *After collecting $K = \frac{C_2(d^2H^3+dH^4)\log^3(Hd\delta^{-1}\epsilon^{-1})}{\epsilon^2} + \frac{C_2(d^{2.5}H^2+d^2H^3)\log^3(Hd\delta^{-1}\epsilon^{-1}))}{\epsilon}$ trajectories for some constant C_2, Algorithm 2 returns a transition model $\tilde{P}_K$. Then with probability at least $1-\delta$, for any given reward, a policy returned by any ϵ_{opt}-optimal plug-in solver on $(S, A, \tilde{P}_K, R, H, \nu)$ is $O(\epsilon + \epsilon_{\text{opt}})$-optimal for the true MDP, (S, A, P, R, H, ν).*

6 Conclusion

This paper studies the sample complexity of plug-in solver approach for reward-free reinforcement learning. We propose a statistically efficient algorithm with sample complexity $\tilde{O}\left(d^2H^4/\epsilon^2\right)$. We further refine the complexity by providing an another algorithm with sample complexity $\tilde{O}\left(d^2H^3/\epsilon^2\right)$ in certain parameter regimes. To the best of our knowledge, this is the first minimax sample complexity bound for reward-free exploration with linear function approximation. As a side note, our approaches provide an efficient learning method for the RL model representation, which preserves values and policies for any other down-stream tasks (specified by different rewards).

Our sample complexity bound matches the lower bound only when $d \geq H$ and $\epsilon \leq H/\sqrt{d}$. It is unclear whether minimax rate could be obtained in a more broader parameter regimes. We plan to address this issue in the future work.

REFERENCES

Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *NIPS*, volume 11, pp. 2312–2320, 2011.

Alekh Agarwal, Sham Kakade, and Lin F Yang. Model-based reinforcement learning with a generative model is minimax optimal. In *Conference on Learning Theory*, pp. 67–83. PMLR, 2020.

Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pp. 463–474. PMLR, 2020.

Mohammad Gheshlaghi Azar, Rémi Munos, and Hilbert J Kappen. Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91(3): 325–349, 2013.

Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pp. 263–272. PMLR, 2017.

Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pp. 1283–1294. PMLR, 2020.

Qiwen Cui and Lin F Yang. Is plug-in solver sample-efficient for feature-based reinforcement learning? *arXiv preprint arXiv:2010.05673*, 2020.

Simon S Du, Sham M Kakade, Ruosong Wang, and Lin F Yang. Is a good representation sufficient for sample efficient reinforcement learning? *arXiv preprint arXiv:1910.03016*, 2019.

Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pp. 1704–1713. PMLR, 2017.

Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? *arXiv preprint arXiv:1807.03765*, 2018.

Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration for reinforcement learning. In *International Conference on Machine Learning*, pp. 4870–4879. PMLR, 2020a.

Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020b.

Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, et al. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019.

Emilie Kaufmann, Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Edouard Leurent, and Michal Valko. Adaptive reward-free exploration. In *Algorithmic Learning Theory*, pp. 865–891. PMLR, 2021.

Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. Morel: Model-based offline reinforcement learning. *arXiv preprint arXiv:2005.05951*, 2020.

Tor Lattimore and Marcus Hutter. Pac bounds for discounted mdps. In *International Conference on Algorithmic Learning Theory*, pp. 320–334. Springer, 2012.

Yuping Luo, Huazhe Xu, Yuanzhi Li, Yuandong Tian, Trevor Darrell, and Tengyu Ma. Algorithmic framework for model-based deep reinforcement learning with theoretical guarantees. In *International Conference on Learning Representations*, 2019. URL `https://openreview.net/forum?id=BJe1E2R5KX`.

Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Emilie Kaufmann, Edouard Leurent, and Michal Valko. Fast active learning for pure exploration in reinforcement learning. *arXiv preprint arXiv:2007.13442*, 2020.

Aditya Modi, Nan Jiang, Ambuj Tewari, and Satinder Singh. Sample complexity of reinforcement learning using linearly combined model ensembles. In *International Conference on Artificial Intelligence and Statistics*, pp. 2010–2020. PMLR, 2020.

Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning*, pp. 2778–2787. PMLR, 2017.

Bradly C Stadie, Sergey Levine, and Pieter Abbeel. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.

Ruosong Wang, Simon S Du, Lin F Yang, and Ruslan Salakhutdinov. On reward-free reinforcement learning with linear function approximation. *arXiv preprint arXiv:2006.11274*, 2020.

Tingwu Wang, Xuchan Bao, Ignasi Clavera, Jerrick Hoang, Yeming Wen, Eric Langlois, Shunshi Zhang, Guodong Zhang, Pieter Abbeel, and Jimmy Ba. Benchmarking model-based reinforcement learning. *arXiv preprint arXiv:1907.02057*, 2019.

Gellért Weisz, Philip Amortila, and Csaba Szepesvári. Exponential lower bounds for planning in mdps with linearly-realizable optimal action-value functions. In *Algorithmic Learning Theory*, pp. 1237–1264. PMLR, 2021.

Yue Wu, Dongruo Zhou, and Quanquan Gu. Nearly minimax optimal regret for learning infinite-horizon average-reward mdps with linear function approximation. *arXiv preprint arXiv:2102.07301*, 2021.

Lin Yang and Mengdi Wang. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, pp. 6995–7004. PMLR, 2019.

Lin Yang and Mengdi Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*, pp. 10746–10756. PMLR, 2020.

Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. Mopo: Model-based offline policy optimization. *arXiv preprint arXiv:2005.13239*, 2020.

Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pp. 10978–10989. PMLR, 2020a.

Andrea Zanette, Alessandro Lazaric, Mykel J Kochenderfer, and Emma Brunskill. Provably efficient reward-agnostic navigation with linear value iteration. *arXiv preprint arXiv:2008.07737*, 2020b.

Dongruo Zhou, Quanquan Gu, and Csaba Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. *arXiv preprint arXiv:2012.08507*, 2020a.

Dongruo Zhou, Jiafan He, and Quanquan Gu. Provably efficient reinforcement learning for discounted mdps with feature mapping. *arXiv preprint arXiv:2006.13165*, 2020b.

A Omitted Proof in Section 4

A.1 High-probability Events

We firstly state the following high-probability events.

Lemma 4. *With probability at least $1-\delta/2$, the following inequality holds for any $k \in [K], h \in [H]$:*

$$\|\hat{\theta}_{k,h} - \theta_h\|_{\Lambda_{k,h}} \leq \beta. \tag{14}$$

Proof. For some step $h \in [H]$, the agent selects an action with feature $x_{k,h} = \phi_{k,h}(s_{k,h}, a_{k,h})$. The noise $\eta_{k,h} = \tilde{V}_{k,h,s_{k,h},a_{k,h}}(s') - \theta_h^\top \phi_{k,h}(s_{k,h}, a_{k,h})$ satisfies H-sub-Gaussian. By Lemma 16, the following inequality holds with probability $1 - \frac{\delta}{2H}$:

$$\|\hat{\theta}_{k,h} - \theta_h\|_{\Lambda_{k,h}} \leq H\sqrt{d\log(4HK/\delta)}. \tag{15}$$

By taking union bound over all $h \in [H]$, we can prove the lemma.

□

Lemma 5. *With probability at least* $1 - \delta/2$*, we have*

$$\sum_{k=1}^{K}\sum_{h=1}^{H}\left(P_h V_{k,h+1}(s_{k,h}, a_{k,h}) - V_{k,h+1}(s_{k,h+1})\right) \leq \sqrt{2H^3 K\log(4/\delta)}. \tag{16}$$

Proof. This lemma follows directly by Azuma's inequality. □

During the following analysis, we denote the high-probability events defined in Lemma 4 and Lemma 5 as $\mathcal{E}_1$.

A.2 Proof of Lemma 1

Lemma 6. *(Restatement of Lemma 1) Under event* $\mathcal{E}_1$*, the sub-optimality gap of the policy* $\hat{\pi}_R$ *for any reward function* R *in the planning phase can be bounded by*

$$V_1^*(s_1, R) - V_1^{\hat{\pi}_R}(s_1, R) \leq 4\tilde{V}_1^*(s_1, R_K) + \epsilon_{\text{opt}}. \tag{17}$$

Proof.

$$V_1^*(s_1, R) - V_1^{\hat{\pi}_R}(s_1, R) \tag{18}$$

$$= \left(V_1^*(s_1, R) - \hat{V}_1^{\pi_R^*, \tilde{P}}(s_1, R)\right) + \left(\hat{V}_1^{\hat{\pi}_R, \tilde{P}}(s_1, R) - V_1^{\hat{\pi}_R}(s_1, R)\right) + \left(\hat{V}_1^{\pi_R^*, \tilde{P}}(s_1, R) - \hat{V}_1^{\hat{\pi}_R, \tilde{P}}(s_1, R)\right) \tag{19}$$

$$\leq \left(V_1^*(s_1, R) - \hat{V}_1^{\pi_R^*, \tilde{P}}(s_1, R)\right) + \left(\hat{V}_1^{\hat{\pi}_R, \tilde{P}}(s_1, R) - V_1^{\hat{\pi}_R}(s_1, R)\right) + \epsilon_{\text{opt}}. \tag{20}$$

The inequality is because that the policy $\hat{\pi}_R$ is the ϵ_{opt}-optimal policy in the estimated MDP $\hat{M}$, i.e. $\hat{V}_1^{\pi_R^*, \tilde{P}}(s_1, R) \leq \hat{V}_1^{\hat{\pi}_R, \tilde{P}}(s_1, R) + \epsilon_{\text{opt}}$.

For the notation convenience, for a certain sequence $\{R_h\}_{h=1}^H$, we define the function $W_h(\{R_h\})$ recursively from step $H+1$ to step 1. Firstly, we define $W_{H+1}(\{R_h\}) = 0$. $W_h(\{R_h\})$ is calculated recursively from $W_{h+1}(\{R_h\})$:

$$W(\{R_h\}) = \min\left\{H, R_h + W_{h+1}(\{R_h\})\right\}. \tag{21}$$

We use $\mathrm{traj} \sim (\pi, P)$ to indicate that the trajectory $\{s_h, a_h\}_{h=1}^H$ is sampled from transition P with policy π. For any policy π, we have

$$
\begin{aligned}
&\left|\mathbb{E}_{s_1\sim\mu}\left(\hat{V}_1^{\pi,\tilde{P}}(s_1,R) - V_1^{\pi}(s_1)\right)\right| && (22)\\
=&\left|\mathbb{E}_{\mathrm{traj}\sim(\pi,P)} W_1\left(\left\{(\tilde{P}_h - P_h)\hat{V}_{h+1}^{\pi,\tilde{P}_K}(s_h,a_h,R)\right\}\right)\right| && (23)\\
=&\left|\mathbb{E}_{\mathrm{traj}\sim(\pi,P)} W_1\left(\left\{(\tilde{\theta}_h - \theta_h)\sum_{s'}\phi(s_h,a_h,s')\hat{V}_{h+1}^{\pi,\tilde{P}}(s',R)\right\}\right)\right| && (24)\\
\leq&\mathbb{E}_{\mathrm{traj}\sim(\pi,P)} W_1\left(\left\{\left\|\tilde{\theta}_h - \theta_h\right\|_{\Lambda_{K,h}}\left\|\sum_{s'}\phi(s_h,a_h,s')\hat{V}_{h+1}^{\pi,\tilde{P}}(s',R)\right\|_{(\Lambda_{K,h})^{-1}}\right\}\right) && (25)\\
\leq&\mathbb{E}_{\mathrm{traj}\sim(\pi,P)} W_1\left(\left\{2\beta\left\|\sum_{s'}\phi(s_h,a_h,s')\hat{V}_{h+1}^{\pi,\tilde{P}}(s',R)\right\|_{(\Lambda_{K,h})^{-1}}\right\}\right) && (26)\\
\leq&2\mathbb{E}_{\mathrm{traj}\sim(\pi,P)} W_1\left(\{u_{K,h}(s_h,a_h)\}\right) && (27)\\
=&2\tilde{V}_1^{\pi}(s_1,R_K) && (28)\\
\leq&2\tilde{V}_1^{*}(s_1,R_K). && (29)
\end{aligned}
$$

The second inequality is due to lemma 4 and the definition of $\tilde{\theta}$. Plugging this inequality back to Inq. 18, we can prove the lemma. □

A.3 Proof of Lemma 2

Lemma 7. *(Restatement of Lemma 2) Under event $\mathcal{E}_1$, $\tilde{V}_h^*(s,R_k) \leq V_{k,h}(s)$ holds for any $(s,a) \in \mathcal{S}\times\mathcal{A}, h\in[H]$ and $k\in[K]$. .*

Proof. We prove the lemma by induction. Suppose $\tilde{V}_{h+1}^*(s,R_k) \leq V_{k,h+1}(s)$,

$$
\begin{aligned}
&\tilde{Q}_h^*(s,a,R_k) - Q_{k,h}(s,a) && (30)\\
\leq&-u_{k,h}(s,a) + \left(\theta_{k,h} - \hat{\theta}_{k,h}\right)^\top\cdot\sum_{s'}\phi(s,a,s')V_{k,h}(s') + P_h(\tilde{V}_{h+1}^* - V_{k,h+1})(s,a) && (31)\\
\leq&-u_{k,h}(s,a) + \left(\theta_{k,h} - \hat{\theta}_{k,h}\right)^\top\cdot\sum_{s'}\phi(s,a,s')V_{k,h}(s') && (32)\\
\leq&-u_{k,h}(s,a) + \left\|\theta_{k,h} - \hat{\theta}_{k,h}\right\|_{\Lambda_{k,h}}\left\|\sum_{s'}\phi(s,a,s')V_{k,h}(s')\right\|_{(\Lambda_{k,h})^{-1}} && (33)\\
\leq&-u_{k,h}(s,a) + \left\|\theta_{k,h} - \hat{\theta}_{k,h}\right\|_{\Lambda_{k,h}}\max_{V\in\mathcal{V}}\left\|\sum_{s'}\phi(s,a,s')V(s')\right\|_{(\Lambda_{k,h})^{-1}} && (34)\\
\leq&-u_{k,h}(s,a) + \beta\max_{V\in\mathcal{V}}\left\|\sum_{s'}\phi(s,a,s')V(s')\right\|_{(\Lambda_{k,h})^{-1}} && (35)\\
=&0. && (36)
\end{aligned}
$$

The first inequality is due to induction condition $\tilde{V}_{h+1}^*(s,R_k) \leq V_{k,h+1}(s)$. The last inequality is due to Lemma 4. Since $\tilde{Q}_h^*(s,a) \leq Q_{k,h}(s,a,R_k)$ for any $a\in\mathcal{A}$, we have $\tilde{V}_h^*(s,R_k) \leq V_{k,h}(s)$. □

A.4 Proof of Lemma 3

Lemma 8. *(Restatement of Lemma 3) Under event $\mathcal{E}_1$, $\sum_{k=1}^K V_{k,1}(s_{k,1}) \leq 6dH^2\sqrt{K\log(2HK/\delta)}$.*

Proof.

$$
\begin{aligned}
& Q_{k,h}(s_{k,h}, a_{k,h}) && (37)\\
\leq & R_{k,h}(s_{k,a}, a_{k,h}) + u_{k,h}(s_{k,a}, a_{k,h}) + \hat{P}_{k,h} V_{k,h+1}(s_{k,h}, a_{k,h}) && (38)\\
\leq & 2\beta \left\|\phi_{k,h}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}} + \hat{P}_{k,h} V_{k,h+1}(s_{k,h}, a_{k,h}) && (39)\\
= & 2\beta \left\|\phi_{k,h}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}} + \left(\hat{P}_{k,h} - P_h\right) V_{k,h+1}(s_{k,h}, a_{k,h}) && (40)\\
& + \left(P_h V_{k,h+1}(s_{k,h}, a_{k,h}) - V_{k,h+1}(s_{k,h+1})\right) + V_{k,h+1}(s_{k,h+1}) && (41)\\
\leq & 3\beta \left\|\phi_{k,h}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}} + \left(P_h V_{k,h+1}(s_{k,h}, a_{k,h}) - V_{k,h+1}(s_{k,h+1})\right) + V_{k,h+1}(s_{k,h+1}). && (42)
\end{aligned}
$$

The last inequality is because that

$$
\begin{aligned}
\left(\hat{P}_{k,h} - P_h\right) V_{k,h+1}(s_{k,h}, a_{k,h}) = & (\hat{\theta}_{k,h} - \theta_h)\phi_{V_{k,h+1}}(s_{k,h}, a_{k,h}) && (43)\\
\leq & \left\|\hat{\theta}_{k,h} - \theta_h\right\|_{\Lambda_{k,h}} \left\|\phi_{V_{k,h+1}}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}} && (44)\\
\leq & \beta \left\|\phi_{V_{k,h+1}}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}} && (45)\\
\leq & \beta \left\|\phi_{k,h}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}}. && (46)
\end{aligned}
$$

From Inq 37, we have

$$
\sum_{k=1}^{K} V_{k,1}(s_1) = 3\sum_{k=1}^{K}\sum_{h=1}^{H} \beta \left\|\phi_{k,h}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}} + \sum_{k=1}^{K}\sum_{h=1}^{H} \left(P_h V_{k,h+1}(s_{k,h}, a_{k,h}) - V_{k,h+1}(s_{k,h+1})\right). \tag{47}
$$

For the first term, by Lemma 17, we have for any $h \in [H]$,

$$
\sum_{k=1}^{K} \left\|\phi_{k,h}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}} \leq \sqrt{K \sum_{k=1}^{K} \left\|\phi_{k,h}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}}^2} \leq \sqrt{2dK\log(K)}. \tag{48}
$$

Therefore,

$$
3\sum_{k=1}^{K}\sum_{h=1}^{H} \beta \left\|\phi_{k,h}(s_{k,h}, a_{k,h})\right\|_{\Lambda_{k,h}^{-1}} \leq 3H^2 d\sqrt{2K}\log(2K/\delta). \tag{49}
$$

For the second term, by Lemma 5,

$$
\sum_{k=1}^{K}\sum_{h=1}^{H} \left(P_h V_{k,h+1}(s_{k,h}, a_{k,h}) - V_{k,h+1}(s_{k,h+1})\right) \leq \sqrt{2H^3 K\log(4/\delta)}. \tag{50}
$$

Therefore, we have

$$
\sum_{k=1}^{K} V_{k,1}(s_{k,1}) \leq 6H^2 d\sqrt{K}\log(2HK/\delta). \tag{51}
$$

□

Lemma 9. *Under event* Λ_1, $\tilde{V}_1^*(s_1, R_K) \leq 6dH^2\log(2HK/\delta)/\sqrt{K}$.

Proof. Since $\tilde{V}_1^*(s_{k,1}, R_k) \leq V_{k,1}(s_{k,1})$ by lemma 7, we have

$$
\sum_{k=1}^{K} \tilde{V}_1^*(s_1, R_k) \leq \sum_{k=1}^{K} V_{k,1}(s_1) \leq 6dH^2\sqrt{K}\log(2HK/\delta). \tag{52}
$$

By the definition of R_k, we know that $R_k(s, a)$ is non-increasing w.r.t k. Therefore,

$$
\sum_{k=1}^{K} \tilde{V}_1^*(s_1, R_K) \leq \sum_{k=1}^{K} \tilde{V}_1^*(s_1, R_k) \leq 6dH^2\sqrt{K}\log(2HK/\delta). \tag{53}
$$

The lemma is proved by dividing both sides by K. □

A.5 PROOF OF THEOREM 1

Combining the results in Lemma 9 and Lemma 6, we know that for any reward function R,

$$V_1^*(s_1, R) - V_1^{\hat{\pi}_R}(s_1, R) \leq 24H^3 d \log(2HK/\delta)/\sqrt{K} + \epsilon_{\text{opt}}. \tag{54}$$

Choosing $K = C_1 H^6 d^2/\epsilon^2 \log(Hd\delta^{-1}\epsilon^{-1}))$ for some constant C_1 suffices to guarantee that $V_1^*(s_1, R) - V_1^{\hat{\pi}_R}(s_1, R) \leq \epsilon + \epsilon_{\text{opt}}$.

B OMITTED DETAILS IN SECTION 5

B.1 OMITTED DETAILS OF ALGORITHM 2

In algorithm 2, we maintain five different parameter estimation $\hat{\theta}_{1,k,h}, \hat{\theta}_{2,k,h}, \hat{\theta}_{3,k,h}, \hat{\theta}_{4,k,h}$ and $\hat{\theta}_{5,k,h}$, and the corresponding covariance matrix $\Lambda_{1,k,h}, \Lambda_{2,k,h}, \Lambda_{3,k,h}, \Lambda_{4,k,h}$ and $\Lambda_{5,k,h}$. $\hat{\theta}_{1,k,h}, \hat{\theta}_{2,k,h}$ are the parameter estimation using history samples w.r.t the features of variance-normalized value function $\hat{V}_{k,h+1}^{\pi_k,\tilde{P}_k}/\bar{\sigma}_{1,k,h}$ and $V_{k,h+1}^{\pi_k,\tilde{P}_k}/\bar{\sigma}_{2,k,h}$, respectively. $\hat{\theta}_{3,k,h}, \hat{\theta}_{4,k,h}$ are the parameter estimation using samples w.r.t the features of value function $\left(\hat{V}_{k,h+1}^{\pi_k,\tilde{P}_k}\right)^2$ and $\left(V_{k,h+1}^{\pi_k,\tilde{P}_k}\right)^2$. $\hat{\theta}_{5,k,h}$ is somewhat technical and will be explained later. for any $i \in [5]$ and $h \in [H]$, $\hat{\theta}_{i,1,h}$ are initialized as $\mathbf{0}$, and $\Lambda_{i,1,h}$ are initialized as I in Algorithm 2. $\mathcal{U}_{1,h}$ is the set containing all the $\tilde{\theta}_h$ that makes $\tilde{P}_h$ well-defined, i.e. $\sum_{s'} \tilde{\theta}_h^\top \phi(s'|s,a) = 1$ and $\tilde{\theta}_h^\top \phi(s'|s,a) \geq 0, \forall s, a$.

After observing $\{s_{k,h}, a_{k,h}\}_{h=1}^H$ in episode k, we calculate $\bar{\mathbb{V}}_{1,k,h}(s,a)$ and $\bar{\mathbb{V}}_{2,k,h}(s,a)$ as the corresponding variance in the empirical MDP with transition dynamics $\tilde{P}_k$. We add clipping to control the range of our variance estimator.

$$\bar{\mathbb{V}}_{1,k,h}(s,a) = \tilde{\theta}_{k,h}^\top \sum_{s'} \phi(s,a,s') \left(\hat{V}_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)\right)^2 - \left[\tilde{\theta}_{k,h}^\top \sum_{s'} \phi(s,a,s') \hat{V}_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)\right]^2, \tag{55}$$

$$\bar{\mathbb{V}}_{2,k,h}(s,a) = \tilde{\theta}_{k,h}^\top \sum_{s'} \phi(s,a,s') \left(V_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)\right)^2 - \left[\tilde{\theta}_{k,h}^\top \sum_{s'} \phi(s,a,s') V_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)\right]^2. \tag{56}$$

To guarantee the variance estimation is optimistic, we calculate the confidence bonus $E_{1,k,h}$ and $E_{2,k,h}$ for variance estimation $\bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h})$ and $\bar{\mathbb{V}}_{2,k,h}(s_{k,h}, a_{k,h})$:

$$E_{1,k,h} = \min\left\{H^2, 4H\check{\beta}_k \left\|\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \hat{V}_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)\right\|_{(\Lambda_{1,k,h})^{-1}}\right\} \tag{57}$$

$$+ \min\left\{H^2, 2\tilde{\beta}_k \left\|\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \left(\hat{V}_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)\right)^2\right\|_{(\Lambda_{3,k,h})^{-1}}\right\}, \tag{58}$$

$$E_{2,k,h} = \min\left\{H^2, 4H\check{\beta}_k \left\|\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') V_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)\right\|_{(\Lambda_{2,k,h})^{-1}}\right\} \tag{59}$$

$$+ \min\left\{H^2, 2\tilde{\beta}_k \left\|\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \left(V_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)\right)^2\right\|_{(\Lambda_{4,k,h})^{-1}}\right\}. \tag{60}$$

We add the bonuses to $\bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h})$ and $\bar{\mathbb{V}}_{2,k,h}(s_{k,h}, a_{k,h})$ and maintain the optimistic variance estimation $\bar{\sigma}_{1,k,h}^2$ and $\bar{\sigma}_{2,k,h}^2$:

$$\bar{\sigma}_{1,k,h}^2 = \max\left\{H^2/d, \bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h}) + E_{1,k,h}\right\}, \tag{61}$$

$$\bar{\sigma}_{2,k,h}^2 = \max\left\{H^2/d, \bar{\mathbb{V}}_{2,k,h}(s_{k,h}, a_{k,h}) + E_{2,k,h}\right\}. \tag{62}$$

We use $\bar{\sigma}_{i,k,h}$ to normalize the value obtained in each step. By setting $\lambda = 1$ and solving the ridge regression defined in Eqn. 10, we know that $\hat{\theta}_{i,k+1,h}$ and $\Lambda_{i,k+1,h}$ are updated in the following way:

$$\Lambda_{1,k+1,h} = \Lambda_{1,k,h} + (\bar{\sigma}_{1,k,h})^{-2} \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right)^\top \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right), \tag{63}$$

$$\Lambda_{2,k+1,h} = \Lambda_{1,k,h} + (\bar{\sigma}_{2,k,h})^{-2} \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') V^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right)^\top \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') V^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right), \tag{64}$$

$$\hat{\theta}_{1,k+1,h} = (\Lambda_{1,k+1,h})^{-1} \sum_{t=1}^{k} (\bar{\sigma}_{1,t,h})^{-2} \left(\sum_{s'} \phi(s_{t,h}, a_{t,h}, s') \hat{V}^{\pi_t, \tilde{P}_t}_{t,h+1}(s', R_t) \right) \hat{V}^{\pi_t, \tilde{P}_t}_{t,h+1}(s_{t,h+1}, R_t), \tag{65}$$

$$\hat{\theta}_{2,k+1,h} = (\Lambda_{2,k+1,h})^{-1} \sum_{t=1}^{k} (\bar{\sigma}_{2,t,h})^{-2} \left(\sum_{s'} \phi(s_{t,h}, a_{t,h}, s') V^{\pi_t, \tilde{P}_t}_{t,h+1}(s', R_t) \right) V^{\pi_t, \tilde{P}_t}_{t,h+1}(s_{t,h+1}, R_t). \tag{66}$$

Similarly, we update the estimation w.r.t $\left(\hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k)\right)^2$ and $\left(V^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k)\right)^2$:

$$\Lambda_{3,k+1,h} = \Lambda_{3,k,h} + \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \left(\hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right)^2 \right)^\top \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \left(\hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right)^2 \right), \tag{67}$$

$$\Lambda_{4,k+1,h} = \Lambda_{4,k,h} + \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \left(V^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right)^2 \right)^\top \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \left(V^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right)^2 \right), \tag{68}$$

$$\hat{\theta}_{3,k+1,h} = (\Lambda_{3,k+1,h})^{-1} \sum_{t=1}^{k} \left(\sum_{s'} \phi(s_{t,h}, a_{t,h}, s') \left(\hat{V}^{\pi_t, \tilde{P}_t}_{t,h+1}(s', R_t) \right)^2 \right) \left(\hat{V}^{\pi_t, \tilde{P}_t}_{t,h+1}(s_{t,h+1}, R_t) \right)^2, \tag{69}$$

$$\hat{\theta}_{4,k+1,h} = (\Lambda_{4,k+1,h})^{-1} \sum_{t=1}^{k} \left(\sum_{s'} \phi(s_{t,h}, a_{t,h}, s') \left(V^{\pi_t, \tilde{P}_t}_{t,h+1}(s', R_t) \right)^2 \right) \left(V^{\pi_t, \tilde{P}_t}_{t,h+1}(s_{t,h+1}, R_t) \right)^2. \tag{70}$$

We define $\tilde{Y}_{k,h}$ to be the "value function" with transition $\tilde{P}_k$ and reward $\bar{\mathbb{V}}_{1,k,h}$.

$$\tilde{Y}_{k,H+1}(s) = 0, \tag{71}$$

$$\tilde{Y}_{k,h}(s) = \bar{\mathbb{V}}_{1,k,h}(s, \pi_k(s)) + \tilde{P}_{k,h} \tilde{Y}_{k,h+1}(s, \pi_k(s)). \tag{72}$$

$\hat{\theta}_{5,k,h}$ is the parameter estimation using samples w.r.t $\tilde{Y}_{k,h}(s)$:

$$\Lambda_{5,k+1,h} = \Lambda_{5,k,h} + \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \tilde{Y}_{k,h+1}(s') \right)^\top \left(\sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \tilde{Y}_{k,h+1}(s') \right), \tag{73}$$

$$\hat{\theta}_{5,k+1,h} = (\Lambda_{5,k+1,h})^{-1} \sum_{t=1}^{k} \left(\sum_{s'} \phi(s_{t,h}, a_{t,h}, s') \tilde{Y}_{t,h+1}(s') \right) \tilde{Y}_{t,h+1}(s_{k,h+1}). \tag{74}$$

We update the high-confidence set $\mathcal{U}_{k,h}$ by adding following five constraints:

$$\|\tilde{\theta}_h - \hat{\theta}_{1,k,h}\|_{\Lambda_{1,k,h}} \leq \hat{\beta}, \tag{75}$$

$$\|\tilde{\theta}_h - \hat{\theta}_{2,k,h}\|_{\Lambda_{2,k,h}} \leq \hat{\beta}, \tag{76}$$

$$\|\tilde{\theta}_h - \hat{\theta}_{3,k,h}\|_{\Lambda_{3,k,h}} \leq \tilde{\beta}, \tag{77}$$

$$\|\tilde{\theta}_h - \hat{\theta}_{4,k,h}\|_{\Lambda_{4,k,h}} \leq \tilde{\beta}, \tag{78}$$

$$\|\tilde{\theta}_h - \hat{\theta}_{5,k,h}\|_{\Lambda_{5,k,h}} \leq \tilde{\beta}. \tag{79}$$

Note that we add new constraints to this confidence set $\mathcal{U}_{k,h}$ in each episode, instead of updating $\mathcal{U}_{k,h}$ as the intersection of the new constraints. This operation is designed to ensure that the cardinality of the confidence set $\mathcal{U}_{k,h}$ is non-increasing w.r.t. the episode k, so that $V_{k,1}^{\pi,\tilde{P}}(s_1, R)$ is always non-increasing w.r.t. k.

During the proof, we use $\phi_{V^{\pi,P}}(s,a,R)$ as a shorthand of $\sum_{s'} \phi(s,a,s')V^{\pi,P}(s',R)$.

B.2 High-probability events

For notation convenience, we use $\mathbb{V}_{1,k,h}(s,a)$ and $\mathbb{V}_{2,k,h}(s,a)$ to denote the one-step transition variance with regard to $\hat{V}_{k,h+1}^{\pi_k,P_k}$ and $V_{k,h+1}^{\pi_k,P_k}$, i.e.

$$\mathbb{V}_{1,k,h}(s,a) = \mathbb{E}_{s'\sim P_h(\cdot|s,a)}\left[\left(\hat{V}_{k,h+1}^{\pi_k,P_k}(s',R_k) - P_h\hat{V}_{k,h}^{\pi_k,P_k}(s,a,R_k)\right)^2\right], \tag{80}$$

$$\mathbb{V}_{2,k,h}(s,a) = \mathbb{E}_{s'\sim P_h(\cdot|s,a)}\left[\left(V_{k,h+1}^{\pi_k,P_k}(s',R_k) - P_hV_{k,h}^{\pi_k,P_k}(s,a,R_k)\right)^2\right]. \tag{81}$$

$$\tag{82}$$

Lemma 10. *With probability at least $1-\delta$, the following event holds for any $h \in [H], k \in [K]$:*

$$\left\|\tilde{\theta}_{k,h} - \theta_h\right\|_{\Lambda_{1,k,h}} \leq 2\hat{\beta}, \tag{83}$$

$$\left\|\tilde{\theta}_{k,h} - \theta_h\right\|_{\Lambda_{2,k,h}} \leq 2\hat{\beta}, \tag{84}$$

$$\left|\bar{\mathbb{V}}_{1,k,h}(s_{k,h},a_{k,h}) - \mathbb{V}_{1,k,h}(s_{k,h},a_{k,h})\right| \leq E_{1,k,h}, \tag{85}$$

$$\left|\bar{\mathbb{V}}_{2,k,h}(s_{k,h},a_{k,h}) - \mathbb{V}_{2,k,h}(s_{k,h},a_{k,h})\right| \leq E_{2,k,h}, \tag{86}$$

$$\left\|\tilde{\theta}_{k,h} - \theta_h\right\|_{\Lambda_{5,k,h}} \leq 2\tilde{\beta}. \tag{87}$$

Proof. We firstly prove the first and the third inequality, the second and the fourth inequality can be proved following the same idea. By the definition of $\bar{\mathbb{V}}_{1,k,h}(s_{k,h},a_{k,h})$ and $\mathbb{V}_{1,k,h}(s_{k,h},a_{k,h})$:

$$\left|\bar{\mathbb{V}}_{1,k,h}(s_{k,h},a_{k,h}) - \mathbb{V}_{1,k,h}(s_{k,h},a_{k,h})\right| \tag{88}$$

$$\leq \min\left\{H^2, 2H\left\|\tilde{\theta}_{k,h} - \theta_h\right\|_{\Lambda_{1,k,h}}\left\|\phi_{\hat{V}_{k,h+1}^{\pi_k,P_k}}(s_{k,h},a_{k,h},R_k)\right\|_{(\Lambda_{1,k,h})^{-1}}\right\} \tag{89}$$

$$+ \min\left\{H^2, \left\|\tilde{\theta}_{k,h} - \theta_h\right\|_{\Lambda_{3,k,h}}\left\|\phi_{\left(\hat{V}_{k,h+1}^{\pi_k,P_k}\right)^2}(s_{k,h},a_{k,h},R_k)\right\|_{(\Lambda_{3,k,h})^{-1}}\right\}. \tag{90}$$

Let $x_k = (\bar{\sigma}_{1,k,h})^{-1}\phi_{\hat{V}_{k,h+1}^{\pi_k,P_k}}(s_{k,h},a_{k,h})$, and the noise

$$\eta_k = (\bar{\sigma}_{1,k,h})^{-1}\hat{V}_{k,h+1}^{\pi_k,P_k}(s_{k,h+1},R_k) - (\bar{\sigma}_{1,k,h})^{-1}\left\langle \boldsymbol{\phi}_{\hat{V}_{k,h+1}^{\pi_k,P_k}}(s_{k,h},a_{k,h},R_k), \theta_h^*\right\rangle.$$

Since $\bar{\sigma}_{1,k,h} \geq \sqrt{H^2/d}$, we have $\|x_k\|_2 \leq \sqrt{d}$, $\mathbb{E}\left[\eta_k^2 \mid \mathcal{G}_k\right] \leq d$. By Lemma 18, we have with prob at least $1-\delta/(8H)$, for all $k \leq K$,

$$\left\|\theta_h - \hat{\theta}_{1,k,h}\right\|_{\Lambda_{1,k,h}} \leq 16d\log(32KH/\delta) = \check{\beta}. \tag{91}$$

Similarly, we can prove that

$$\left\|\theta_h - \hat{\theta}_{3,k,h}\right\|_{\Lambda_{3,k,h}} \leq 16H^2\sqrt{d}\log(32KH/\delta) = \tilde{\beta}. \tag{92}$$

Since $\left\|\tilde{\theta}_{k,h} - \hat{\theta}_{1,k,h}\right\|_{\Lambda_{1,k,h}} \leq \hat{\beta} \leq \check{\beta}$ by Inq 75, we have $\left\|\theta_h - \tilde{\theta}_{k,h}\right\|_{\Lambda_{1,k,h}} \leq 2\check{\beta}$. Similarly, since $\left\|\tilde{\theta}_{k,h} - \hat{\theta}_{3,k,h}\right\|_{\Lambda_{3,k,h}} \leq \tilde{\beta}$ by Inq 77, we have $\left\|\theta_h - \tilde{\theta}_{k,h}\right\|_{\Lambda_{3,k,h}} \leq 2\tilde{\beta}$. Plugging the above inequalities back to Inq. 88, we have

$$\left|\bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h}) - \mathbb{V}_{1,k,h}(s_{k,h}, a_{k,h})\right| \leq E_{1,k,h}. \tag{93}$$

The above inequality indicates that $\bar{\sigma}_{1,k,h} \geq \mathbb{V}_{1,k,h}$, which is an optimistic estimation. As a result, the variance of the single-step noise $\eta_k = (\bar{\sigma}_{1,k,h})^{-1}\hat{V}^{\pi_k,P_k}_{k,h+1}(s_{k,h+1}, R_k) - (\bar{\sigma}_{1,k,h})^{-1}\left\langle \boldsymbol{\phi}_{\hat{V}^{\pi_k,P_k}_{k,h+1}}(s_{k,h}, a_{k,h}, R_k), \theta_h \right\rangle$ satisfies $\mathbb{E}\left[\eta_k^2 \mid \mathcal{G}_k\right] \leq 1$. By Lemma 18, we can prove a tighter confidence guarantee for $\hat{\theta}_{1,k,h}$:

$$\left\|\theta_h - \hat{\theta}_{1,k,h}\right\|_{\Lambda_{1,k,h}} \leq \hat{\beta}. \tag{94}$$

Combining with $\left\|\tilde{\theta}_{1,k,h} - \hat{\theta}_{1,k,h}\right\|_{\Lambda_{1,k,h}} \leq \hat{\beta}$ by Inq 75, we have $\left\|\theta_h - \tilde{\theta}_{1,k,h}\right\|_{\Lambda_{1,k,h}} \leq 2\hat{\beta}$.

Now we prove the last inequality in this lemma. Recall that we define

$$\tilde{Y}_{k,H+1}(s) = 0, \tag{95}$$

$$\tilde{Y}_{k,h}(s) = \bar{\mathbb{V}}_{1,k,h}(s, \pi_k(s)) + \tilde{P}_{k,h}\tilde{Y}_{k,h+1}(s, \pi_k(s)). \tag{96}$$

Note that $\tilde{Y}_{k,h}(s) \leq H^2$ by the law of total variance (Lattimore & Hutter, 2012; Azar et al., 2013). Therefore, the variance on the single-step noise in at most H^2. By Lemma 18, we have

$$\left\|\theta_h - \hat{\theta}_{5,k,h}\right\|_{\Lambda_{5,k,h}} \leq \tilde{\beta}. \tag{97}$$

Combining with $\left\|\tilde{\theta}_h - \hat{\theta}_{5,k,h}\right\|_{\Lambda_{5,k,h}} \leq \tilde{\beta}$ by Inq 79, we have $\left\|\theta_h - \tilde{\theta}_{5,k,h}\right\|_{\Lambda_{5,k,h}} \leq 2\tilde{\beta}$. □

Lemma 11. *With probability at least $1 - \delta/2$, we have*

$$\sum_{k=1}^{K}\sum_{h=1}^{H}\left(P_h V^{\pi_k,\tilde{P}_k}_{k,h+1}(s_{k,h}, a_{k,h}, R_k) - V^{\pi_k,\tilde{P}_k}_{k,h+1}(s_{k,h+1}, a_{k,h+1}, R_k)\right) \leq \sqrt{2H^3K\log(8/\delta)}, \tag{98}$$

$$\sum_{k=1}^{K}\sum_{h=1}^{H}\left(P_h \tilde{Y}_{k,h+1}(s_{k,h}, a_{k,h}, R_k) - \tilde{Y}_{k,h+1}(s_{k,h+1}, a_{k,h+1}, R_k)\right) \leq \sqrt{2H^5K\log(8/\delta)}. \tag{99}$$

Proof. This lemma follows directly by Azuma's inequality for martingale difference sequence and union bound. □

During the following analysis, we denote the high-probability events defined in Lemma 10 and Lemma 11 as $\mathcal{E}_2$.

B.3 Proof of Theorem 2

Lemma 12. *Under event $\mathcal{E}_2$, we have*

$$\sum_{k=1}^{K} V^{\pi_k,\tilde{P}_k}_{k,1}(s_1, R_k) \leq O\left(\sqrt{(dH^4 + d^2H^3)K\log^3(HK/\delta))} + (d^{2.5}H^2 + d^2H^3)\log^3(HK/\delta)\right). \tag{100}$$

This lemma can be proved with the technique for regret analysis. We will explain it in Appendix B.4.

Lemma 13. *(Optimism) Under event $\mathcal{E}_2$, for any $h \in [H], k \in [K], s \in \mathcal{S}$, $\pi \in \Pi$, R and $\tilde{P}$ satisfying $\tilde{P}_h = \tilde{\theta}_h^\top \phi$, we have*

$$V_{k,h}^{\pi,\tilde{P}}(s,R) \geq \tilde{V}_{k,h}^{\pi,\tilde{P}}(s,R). \tag{101}$$

Proof. This lemma is proved by induction. Suppose $V_{k,h+1}^{\pi,\tilde{P}}(s,R) \geq \tilde{V}_{k,h+1}^{\pi,\tilde{P}}(s,R)$, we have

$$V_{k,h}^{\pi,\tilde{P}}(s,R) - \tilde{V}_{k,h}^{\pi,\tilde{P}}(s,R) \tag{102}$$

$$\geq u_{2,k,h}^{\pi,\tilde{P}}(s,\pi(s),R) + (\tilde{P}_{k,h} - P_h)V_{k,h+1}^{\pi,\tilde{P}}(s,\pi(s),R) + P_h\left(V_{k,h+1}^{\pi,\tilde{P}} - \tilde{V}_{k,h+1}^{\pi,\tilde{P}}\right)(s,\pi(s),R) \tag{103}$$

$$\geq u_{2,k,h}^{\pi,\tilde{P}}(s,\pi(s),R) + (\tilde{P}_{k,h} - P_h)V_{k,h+1}^{\pi,\tilde{P}}(s,\pi(s),R) \tag{104}$$

$$\geq u_{2,k,h}^{\pi,\tilde{P}}(s,\pi(s),R) - \left\|\tilde{\theta}_{k,h} - \theta_h\right\|_{\Lambda_{2,k,h}} \left\|\sum_{s'} \phi(s,\pi(s),s')V_{k,h+1}^{\pi,\tilde{P}}(s',R)\right\|_{(\Lambda_{2,k,h})^{-1}} \tag{105}$$

$$\geq 0. \tag{106}$$

This indicates that $V_{k,h}^{\pi,\tilde{P}}(s,R) \geq \tilde{V}_{k,h}^{\pi,\tilde{P}}(s,R)$ holds for step h. □

Lemma 14. *under event $\mathcal{E}_2$, we have $V_{K,1}^{\pi_K,\tilde{P}_K}(s_1,R_K) \leq \frac{\sum_{k=1}^K V_{k,1}^{\pi_k,\tilde{P}_k}(s_1,R_k)}{K}$.*

Proof. Firstly, we prove that $V_{k,1}^{\pi,\tilde{P}}(s,R)$ is non-increasing w.r.t k for any fixed $\pi, \tilde{P}, s$ and R. This can be proved by induction. Suppose for any $k_1 \leq k_2$, $V_{k_1,h+1}^{\pi,\tilde{P}}(s,R) \geq V_{k_2,h+1}^{\pi,\tilde{P}}(s,R)$ for any s. Recall that

$$V_{k,h}^{\pi,\tilde{P}}(s,R) = \min\left\{u_{1,k,h}^{\pi,\tilde{P}}(s,\pi_h(s),R) + u_{2,k,h}^{\pi,\tilde{P}}(s,\pi_h(s),R) + \tilde{P}_h V_{k,h+1}^{\pi,\tilde{P}}(s,\pi_h(s),R), H\right\}. \tag{107}$$

Since $\Lambda_{1,k_1,h} \preccurlyeq \Lambda_{1,k_2,h}$, $\Lambda_{2,k_1,h} \preccurlyeq \Lambda_{2,k_2,h}$ and $V_{k_1,h+1}^{\pi,\tilde{P}}(s,R) \geq V_{k_2,h+1}^{\pi,\tilde{P}}(s,R)$ for any s, we can prove that

$$u_{1,k_1,h}^{\pi,\tilde{P}}(s,\pi_h(s),R) \geq u_{1,k_2,h}^{\pi,\tilde{P}}(s,\pi_h(s),R), \tag{108}$$

$$u_{2,k_1,h}^{\pi,\tilde{P}}(s,\pi_h(s),R) \geq u_{2,k_2,h}^{\pi,\tilde{P}}(s,\pi_h(s),R), \tag{109}$$

$$\tilde{P}_h V_{k_1,h+1}^{\pi,\tilde{P}}(s,\pi_h(s),R) \geq \tilde{P}_h V_{k_2,h+1}^{\pi,\tilde{P}}(s,\pi_h(s),R). \tag{110}$$

Therefore, $V_{k_1,h}^{\pi,\tilde{P}}(s,R) \geq V_{k_2,h}^{\pi,\tilde{P}}(s,R)$ holds for step h and any s.

Since the cardinality of the transition set $\mathcal{U}_{k,h}$ is non-increasing (We add more constraints in each episode), we know that $\tilde{\theta}_{k_2,h} \in \mathcal{U}_{k_1,h}$. By the optimality of $\pi_{k_1}, \tilde{\theta}_{k_1}$ and R_{k_1} in episode k_1, we have $V_{k_1,1}^{\pi_{k_1},\tilde{P}_{k_1}}(s_1,R_{k_1}) \geq V_{k_1,1}^{\pi_{k_2},\tilde{P}_{k_2}}(s_1,R_{k_2})$.

Combining the above two inequalities, we have

$$V_{k_1,1}^{\pi_{k_1},\tilde{P}_{k_1}}(s_1,R_{k,1}) \geq V_{k_1,1}^{\pi_{k_2},\tilde{P}_{k_2}}(s_1,R_{k_2}) \geq V_{k_2,1}^{\pi_{k_2},\tilde{P}_{k_2}}(s_1,R_{k_2}). \tag{111}$$

This indicates that the value $V_{k,1}^{\pi_k,\tilde{P}_k}(s_1,R_k)$ is non-increasing w.r.t. k. Therefore,

$$KV_{K,1}^{\pi_K,\tilde{P}_K}(s_1,R_K) \leq \sum_{k=1}^K V_{k,1}^{\pi_k,\tilde{P}_k}(s_1,R_k). \tag{112}$$

□

Lemma 15. *Under event $\mathcal{E}_2$, the sub-optimality gap of the policy $\hat{\pi}_R$ returned in the planning phase for any reward function R can be bounded by*

$$V_1^*(s_1, R) - V_1^{\hat{\pi}_R}(s_1, R) \leq 4V_{K,1}^{\pi_K, \tilde{P}_K}(s_1, R_K) + \epsilon_{\text{opt}}. \tag{113}$$

Proof. Recall that we use $\hat{V}_h^\pi$ to denote the value function of policy π on the estimated model $\tilde{\theta}_K$.

$$V_1^*(s_1, R) - V_1^{\hat{\pi}_R}(s_1, R) \tag{114}$$

$$= \left(V_1^*(s_1, R) - \hat{V}_1^{\pi_R^*, \tilde{P}_K}(s_1, R)\right) + \left(\hat{V}_1^{\hat{\pi}_R, \tilde{P}_K}(s_1, R) - V_1^{\hat{\pi}_R}(s_1)\right) + \left(\hat{V}_1^{\pi_R^*, \tilde{P}_K}(s_1, R) - \hat{V}_1^{\hat{\pi}_R, \tilde{P}_K}(s_1, R)\right) \tag{115}$$

$$\leq V_1^*(s_1, R) - \hat{V}_1^{\pi_R^*, \tilde{P}_K}(s_1, R) + \hat{V}_1^{\hat{\pi}_R, \tilde{P}_K}(s_1, R) - V_1^{\hat{\pi}_R}(s_1, R) + \epsilon_{\text{opt}}. \tag{116}$$

The inequality is because that the policy $\hat{\pi}_R$ is the ϵ_{opt}-optimal policy in the estimated MDP $\hat{M}$.

For notation convenience, we use $\text{traj} \sim (\pi, P)$ to denote that the trajectory $(\{s_h, a_h\}_{h=1}^H)$ is sampled with transition P and policy π. For a certain sequence $\{R_h\}_{h=1}^H$, we define the function $W_h(\{R_h\})$ recursively from step $H+1$ to step 1. Firstly, we define $W_{H+1}(\{R_h\}) = 0$. $W_h(R)$ is calculated from $W_{h+1}(R)$:

$$W_h(\{R_h\}) = \min\left\{H, R_h + W_{h+1}(\{R_h\})\right\}. \tag{117}$$

With the above notation, we can prove that for any policy $\pi \in \Pi$,

$$\left|\hat{V}_1^{\pi, \tilde{P}_K}(s_1, R) - V_1^\pi(s_1)\right| \tag{118}$$

$$\leq \left|\mathbb{E}_{\text{traj}\sim(\pi,P)} W_1\left(\left\{(\tilde{P}_{K,h} - P_h)\hat{V}_{h+1}^{\pi, \tilde{P}_K}(s_h, a_h, R)\right\}\right)\right| \tag{119}$$

$$= \left|\mathbb{E}_{\text{traj}\sim(\pi,P)} W_1\left(\left\{(\tilde{\theta}_{K,h} - \theta_h)\sum_{s'}\phi(s_h, a_h, s')\hat{V}_{h+1}^{\pi, \tilde{P}_K}(s', R)\right\}\right)\right| \tag{120}$$

$$\leq \mathbb{E}_{\text{traj}\sim(\pi,P)} W_1\left(\left\{\left\|\tilde{\theta}_{K,h} - \theta_h\right\|_{\Lambda_{1,K,h}}\left\|\sum_{s'}\phi(s_h, a_h, s')\hat{V}_{h+1}^{\pi, \tilde{P}_K}(s', R)\right\|_{(\Lambda_{1,K,h})^{-1}}\right\}\right) \tag{121}$$

$$\leq \mathbb{E}_{\text{traj}\sim(\pi,P)} W_1\left(\left\{2\hat{\beta}\left\|\sum_{s'}\phi(s_h, a_h, s')\hat{V}_{h+1}^{\pi, \tilde{P}_K}(s', R)\right\|_{(\Lambda_{1,K,h})^{-1}}\right\}\right) \tag{122}$$

$$= 2\mathbb{E}_{\text{traj}\sim(\pi,P)} W_1\left(\left\{u_{1,K,h}^{\pi, \tilde{P}_K}(s_h, a_h, R)\right\}\right) \tag{123}$$

$$= 2\tilde{V}_{K,1}^{\pi, \tilde{P}_K}(s_1, R) \tag{124}$$

$$\leq 2V_{K,1}^{\pi, \tilde{P}_K}(s_1, R) \tag{125}$$

$$\leq 2V_{K,1}^{\pi_K, \tilde{P}_K}(s_1, R_K). \tag{126}$$

The second inequality is due to lemma 10 . The third inequality is due to Lemma 13. The last inequality is due to the optimality of π_K and R_K. Plugging this inequality back to Inq. 114, we can prove the lemma. □

Proof. (Proof of Theorem 2) Combining Lemma 12, Lemma 14 and Lemma 15, we know that $K = \frac{C_2(d^2H^3 + dH^4)\log^3(Hd\delta^{-1}\epsilon^{-1}))}{\epsilon^2} + \frac{C_2(d^{2.5}H^2 + d^2H^3)\log^3(Hd\delta^{-1}\epsilon^{-1})}{\epsilon}$ for some constant C_2 suffices to guarantee that $V_1^*(s_1, R) - V_1^{\hat{\pi}}(s_1, R) \leq \epsilon + \epsilon_{\text{opt}}$. □

B.4 PROOF OF LEMMA 12

Proof. From the standard value decomposition technique (Azar et al., 2017; Jin et al., 2018), we can decompose $V_{k,1}^{\pi_k,\tilde{P}_k}(s_1, R_k)$ into following terms.

$$V_{k,1}^{\pi_k,\tilde{P}_k}(s_1, R_k) \tag{127}$$

$$\leq \sum_{h=1}^{H} u_{1,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k) + \sum_{h=1}^{H} u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k) + \sum_{h=1}^{H} (\tilde{P}_{k,h} - P_h) V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k) \tag{128}$$

$$+ \sum_{h=1}^{H} \left(P_h V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k) - V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h+1}, a_{k,h+1}, R_k) \right) \tag{129}$$

$$\leq \sum_{h=1}^{H} u_{1,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k) \tag{130}$$

$$+ \sum_{h=1}^{H} 2u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k) \tag{131}$$

$$+ \sum_{h=1}^{H} \left(P_h V_{k,1}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k) - V_{k,1}^{\pi_k,\tilde{P}_k}(s_{k,h+1}, a_{k,h+1}, R_k) \right), \tag{132}$$

where the second inequality is because that

$$(\tilde{P}_{k,h} - P_h) V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k) = (\tilde{\theta}_{k,h} - \theta_h) \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') V_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k) \tag{133}$$

$$\leq \left\| \tilde{\theta}_{k,h} - \theta_h \right\|_{\Lambda_{2,k,h}} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') V_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k) \right\|_{\Lambda_{2,k,h}^{-1}} \tag{134}$$

$$\leq u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k). \tag{135}$$

Eqn 132 is a martingale difference sequence. By Lemma 11, the summation over all $k \in [K]$ is at most $\sqrt{2H^3 K \log(4/\delta)}$. We mainly focus on Eqn 130 and Eqn 131.

Upper bound of Eqn 130 Firstly, we bound the summation of $u_{1,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k)$.

$$\sum_{k=1}^{K} \sum_{h=1}^{H} u_{1,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h}, a_{k,h}, R_k) \tag{136}$$

$$= \sum_{k=1}^{K} \sum_{h=1}^{H} \hat{\beta} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') V_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k) \right\|_{\Lambda_{2,k,h}^{-1}} \tag{137}$$

$$= \sum_{k=1}^{K} \sum_{h=1}^{H} \hat{\beta} \bar{\sigma}_{1,k,h} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \frac{\hat{V}_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)}{\bar{\sigma}_{1,k,h}} \right\|_{\Lambda_{1,k,h}^{-1}} \tag{138}$$

$$\leq \hat{\beta} \sqrt{\sum_{k,h} \bar{\sigma}_{1,k,h}^2 \sum_{k,h} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \frac{V_{k,h+1}^{\pi_k,\tilde{P}_k}(s', R_k)}{\bar{\sigma}_{1,k,h}} \right\|_{\Lambda_{1,k,h}^{-1}}^2} \tag{139}$$

$$\leq \hat{\beta} \sqrt{\sum_{k,h} \bar{\sigma}_{1,k,h}^2 \cdot dH \log(K)}. \tag{140}$$

The first inequality is due to Cauchy-Schwarz inequality. The second inequality is due to Lemma 17. By the definition of $\bar{\sigma}^2_{1,k,h}$, we have

$$\bar{\sigma}^2_{1,k,h} \leq H^2/d + E_{1,k,h} + \bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h}). \tag{141}$$

Now we bound $\sum_{k,h} E_{1,k,h}$ and $\sum_{k,h} \bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h})$ respectively. For $\sum_{k,h} E_{1,k,h}$, we have

$$\sum_{k,h} E_{1,k,h} \leq \sum_{h=1}^{H}\sum_{k=1}^{K} 4H\check{\beta} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right\|_{(\Lambda_{1,k,h})^{-1}} \tag{142}$$

$$+ \sum_{h=1}^{H}\sum_{k=1}^{K} 2\tilde{\beta} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \left(\hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right)^2 \right\|_{(\Lambda_{3,k,h})^{-1}}. \tag{143}$$

For the first part,

$$\sum_{h=1}^{H}\sum_{k=1}^{K} 4H\check{\beta} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right\|_{(\Lambda_{1,k,h})^{-1}} \tag{144}$$

$$= 4H \sum_{h=1}^{H}\sum_{k=1}^{K} \check{\beta}\bar{\sigma}_{1,k,h} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k)/\bar{\sigma}_{1,k,h} \right\|_{(\Lambda_{1,k,h})^{-1}} \tag{145}$$

$$\leq 4H\check{\beta} \sqrt{\sum_{k,h} \bar{\sigma}^2_{1,k,h} \sum_{k,h} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \frac{\hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k)}{\bar{\sigma}_{1,k,h}} \right\|^2_{\Lambda^{-1}_{1,k,h}}} \tag{146}$$

$$\leq 4H\check{\beta} \sqrt{\sum_{k,h} \bar{\sigma}^2_{1,k,h} \cdot dH \log(K)}, \tag{147}$$

where the last inequality is due to Lemma 17. Similarly, for the second part, by Lemma 17, we have

$$\sum_{h=1}^{H}\sum_{k=1}^{K} 2\tilde{\beta} \left\| \sum_{s'} \phi(s_{k,h}, a_{k,h}, s') \left(\hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k) \right)^2 \right\|_{(\Lambda_{3,k,h})^{-1}} \leq 2\tilde{\beta}\sqrt{dH^2 K \log(K)}. \tag{148}$$

Note that $\bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h})$ is the empirical variance of $\hat{V}^{\pi_k, \tilde{P}_k}_{k,h+1}(s', R_k)$ with transition $\tilde{P}_k(s'|s_{k,h}, a_{k,h})$. We bound the summation of $\bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h})$ by the law of total variance (Lattimore & Hutter, 2012; Azar et al., 2013).

Recall that we define

$$\tilde{Y}_{k,H+1}(s) = 0, \tag{149}$$

$$\tilde{Y}_{k,h}(s) = \bar{\mathbb{V}}_{1,k,h}(s, \pi_k(s)) + \tilde{P}_{k,h}\tilde{Y}_{k,h+1}(s, \pi_k(s)). \tag{150}$$

By the law of total variance, we have $\tilde{Y}_{k,1}(s) \leq H^2$ holds for any s, a and k. We now bound the difference between $\sum_{k=1}^{K} \tilde{Y}_{k,1}(s_{k,1})$ and $\sum_{k=1}^{K}\sum_{h=1}^{H} \bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h})$.

$$\tilde{Y}_{k,1}(s_{k,1}) - \sum_{h=1}^{H} \bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h}) \tag{151}$$

$$= \tilde{P}_{k,1}\tilde{Y}_{k,2}(s_{k,1}, a_{k,1}) - \sum_{h=2}^{H} \bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h}) \tag{152}$$

$$= \left(\tilde{P}_{k,1} - P_1\right) \tilde{Y}_{k,2}(s_{k,1}, a_{k,1}) + \left(P_1\tilde{Y}_{k,2}(s_{k,1}, a_{k,1}) - \tilde{Y}_{k,2}(s_{k,2})\right) \tag{153}$$

$$+ \left(\tilde{Y}_{k,2}(s_{k,2}) - \sum_{h=2}^{H} \bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h})\right). \tag{154}$$

Therefore, we have

$$\sum_{k=1}^{K} \tilde{Y}_{k,1}(s_{k,1}) - \sum_{k=1}^{K}\sum_{h=1}^{H} \bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h}) \tag{155}$$

$$\leq \sum_{k=1}^{K}\sum_{h=1}^{H} \left(\tilde{P}_{k,h} - P_h\right) \tilde{Y}_{k,h+1}(s_{k,h}, a_{k,h}) \tag{156}$$

$$+ \sum_{h=1}^{H}\sum_{k=1}^{K} \left(P_h \tilde{Y}_{k,h+1}(s_{k,h}, a_{k,h}) - \tilde{Y}_{k,h+1}(s_{k,h+1})\right). \tag{157}$$

For Eqn 157, this term can be regarded as a martingale difference sequence, thus can be bounded by $H^2\sqrt{KH}$ by Lemma 11. For Eqn 156, we can bound this term in the following way:

$$\sum_{k=1}^{K}\sum_{h=1}^{H} \left(\tilde{P}_{k,h} - P_h\right) \tilde{Y}_{k,h+1}(s_{k,h}, a_{k,h}) \tag{158}$$

$$\leq \sum_{k=1}^{K}\sum_{h=1}^{H} \left\|\tilde{\theta}_{k,h} - \theta_h\right\|_{\Lambda_{5,k,h}} \|\phi_{\tilde{Y}}(s_{k,h}, a_{k,h}, R)\|_{\Lambda_{5,k,h}^{-1}} \tag{159}$$

$$\leq \tilde{\beta} \sum_{k=1}^{K}\sum_{h=1}^{H} \|\phi_{\tilde{Y}}(s_{k,h}, a_{k,h}, R)\|_{\Lambda_{5,k,h}^{-1}} \tag{160}$$

$$\leq \tilde{\beta} \sum_{h=1}^{H} \sqrt{K \sum_{k=1}^{K} \|\phi_{\tilde{Y}}(s_{k,h}, a_{k,h}, R)\|^2_{\Lambda_{5,k,h}^{-1}}} \tag{161}$$

$$\leq \tilde{\beta}\sqrt{dH^2K\log(K)}. \tag{162}$$

The second inequality is due to Lemma 10. The third inequality is due to Cauchy-Schwarz inequality. The last inequality is due to Lemma 17.

From the above analysis, we have

$$\sum_{k,h} \bar{\mathbb{V}}_{1,k,h}(s_{k,h}, a_{k,h}) \leq H^2K + \tilde{\beta}\sqrt{dH^2K\log(K)}. \tag{163}$$

Therefore, the summation of $\bar{\sigma}^2_{1,k,h}$ can be bounded as

$$\sum_{k,h} \bar{\sigma}^2_{1,k,h} \tag{164}$$

$$\leq H^3K/d + \sum_{k,h} E_{1,k,h} + \sum_{k,h} \bar{\mathbb{V}}_{1,k,h} \tag{165}$$

$$\leq H^3K/d + 4H\check{\beta}\sqrt{\sum_{k,h} \bar{\sigma}^2_{1,k,h} \cdot dH\log(K)} + 3\tilde{\beta}\sqrt{dH^2K\log(K)} + H^2K. \tag{166}$$

Solving for $\sum_{k,h} \bar{\sigma}^2_{1,k,h}$, we have

$$\sum_{k,h} \bar{\sigma}^2_{1,k,h} \tag{167}$$

$$\leq c_1\left(H^3K/d + H^2K + H^3d^3\log^3(KH/\delta) + \sqrt{H^2d^3K\log^3(KH/\delta)}\right) \tag{168}$$

$$\leq 2c_1\left(H^3K/d + H^2K + H^3d^3\log^3(KH/\delta)\right), \tag{169}$$

where c_1 denote a certain constant. The last inequality is due to $2\sqrt{ab} \leq a+b$ for $a, b \geq 0$. Plugging the above inequality back to Inq 136, we have for a constant c_2,

$$\sum_{k=1}^{K}\sum_{h=1}^{H} u^{\pi_k,\tilde{P}_k}_{1,k,h}(s_{k,h}, a_{k,h}, R_k) \leq c_2\sqrt{(dH^4 + d^2H^3)\,K\log^3(KH/\delta)} + c_2 d^{2.5}H^2\log^3(KH/\delta). \tag{170}$$

Upper bound of Eqn 131 Now we focus on the summation of $u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k)$. The proof follows almost the same arguments as the summation of $u_{1,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k)$, though the upper bound of $\sum_{k=1}^K\sum_{h=1}^H \bar{\mathbb{V}}_{2,k,h}(s_{k,h},a_{k,h})$ is derived in a different way. Following the same proof idea, we can show that

$$\sum_{k=1}^K\sum_{h=1}^H u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k) \le \hat{\beta}\sqrt{\sum_{k,h}\bar{\sigma}_{2,k,h}^2 dH\log(K)}, \tag{171}$$

$$\bar{\sigma}_{2,k,h}^2 \le H^2/d + E_{2,k,h} + \bar{\mathbb{V}}_{2,k,h}(s_{k,h},a_{k,h}), \tag{172}$$

$$\sum_{k=1}^K\sum_{h=1}^H E_{2,k,h} \le 4H\check{\beta}\sqrt{\sum_{k,h}\bar{\sigma}_{2,k,h}^2\cdot dH\log(K)} + 2\tilde{\beta}\sqrt{dH^2K\log(K)}. \tag{173}$$

Combining Inq 172 and Inq 173 and solving for $\sum_{k,h}\bar{\sigma}_{2,k,h}^2$, we have

$$\sum_{k,h}\bar{\sigma}_{2,k,h}^2 \le 2H^3K/d + 16dH^3\check{\beta}^2\log(K) + 4\tilde{\beta}\sqrt{dH^2K\log(K)} + 2\sum_{k,h}\bar{\mathbb{V}}_{2,k,h}(s_{k,h},a_{k,h}) \tag{174}$$

$$\le 6H^3K/d + 16dH^3\check{\beta}^2\log(K) + 4\tilde{\beta}^2d^2/H\log(K) + 2\sum_{k,h}\bar{\mathbb{V}}_{2,k,h}(s_{k,h},a_{k,h}) \tag{175}$$

Plugging this inequality back to Inq 171, we have

$$\sum_{k=1}^K\sum_{h=1}^H u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k) \le \hat{\beta}\sqrt{6H^4K\log(K)} + 4\check{\beta}\hat{\beta}dH^2\log(K) + 2\tilde{\beta}\hat{\beta}d^{3/2}\log(K) \tag{176}$$

$$+\sqrt{32d^2H\log^3(32KH/\delta)\sum_{k,h}\bar{\mathbb{V}}_{2,k,h}(s_{k,h},a_{k,h})}. \tag{177}$$

We now bound the last term in the above equation:

$$\sqrt{32d^2H\log^3(32KH/\delta)\sum_{k,h}\bar{\mathbb{V}}_{2,k,h}(s_{k,h},a_{k,h})} \tag{178}$$

$$\le\sqrt{32d^2H\log^3(32KH/\delta)\sum_{k=1}^K\sum_{h=1}^H\mathbb{E}_{s'\sim\tilde{P}_{k,h}(\cdot|s_{k,h},a_{k,h})}\left[(V_{k,h+1}^{\pi_k,\tilde{P}_k})^2(s',R_k)\right]} \tag{179}$$

$$\le\sqrt{32d^2H^2\log^3(32KH/\delta)\sum_{k=1}^K\sum_{h=1}^H\mathbb{E}_{s'\sim\tilde{P}_{k,h}(\cdot|s_{k,h},a_{k,h})}\left[V_{k,h+1}^{\pi_k,\tilde{P}_k}(s',R_k)\right]} \tag{180}$$

$$\le 4\sqrt{2}\tilde{c}_3d^2H^3\log^3(32KH/\delta) + \frac{1}{\tilde{c}_3H}\sum_{k=1}^K\sum_{h=1}^H\mathbb{E}_{s'\sim\tilde{P}_{k,h}(\cdot|s_{k,h},a_{k,h})}\left[V_{k,h+1}^{\pi_k,\tilde{P}_k}(s',R_k)\right] \tag{181}$$

$$= 4\sqrt{2}\tilde{c}_3d^2H^3\log^3(32KH/\delta) + \frac{1}{\tilde{c}_3H}\sum_{k=1}^K\sum_{h=1}^H V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h+1},R_k) \tag{182}$$

$$+\frac{1}{\tilde{c}_3H}\sum_{k=1}^K\sum_{h=1}^H\left(P_hV_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k) - V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h+1},R_k)\right) \tag{183}$$

$$+\frac{1}{\tilde{c}_3H}\sum_{k=1}^K\sum_{h=1}^H\left(\tilde{P}_{k,h}-P_h\right)V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k), \tag{184}$$

where $\tilde{c}_3 \ge 1$ is a constant to be defined later. The last inequality is due to $2\sqrt{ab} \le a+b$ for any $a,b \ge 0$. Note that by Lemma 10, Eqn 184 is upper bounded by

$\frac{1}{c_3 H}\sum_{k,h} u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k)$. Eqn 183 is a martingale difference sequence, and can be bounded by $\sqrt{HK\log(1/\delta)}$ with high probability. Plugging the above inequality back to Inq 176 and solving for $\sum_{k=1}^{K}\sum_{h=1}^{H} u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k)$, we have

$$\sum_{k=1}^{K}\sum_{h=1}^{H} u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k) \tag{185}$$

$$\leq c_3\left(\sqrt{dH^4K\log^3(HK/\delta)}+(d^{2.5}H^2+d^2H^3)\log^3(HK/\delta)+\frac{1}{\tilde{c}_3 H}\sum_{k=1}^{K}\sum_{h=1}^{H}V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h+1},R_k)\right), \tag{186}$$

where $c_3 \geq 1$ is a constant here. We set $\tilde{c}_3 = c_3$, then we have

$$\sum_{k=1}^{K}\sum_{h=1}^{H} u_{2,k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},a_{k,h},R_k) \tag{187}$$

$$\leq c_3\left(\sqrt{dH^4K\log^3(HK/\delta)}+(d^{2.5}H^2+d^2H^3)\log^3(HK/\delta)\right)+\frac{1}{H}\sum_{k=1}^{K}\sum_{h=1}^{H}V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h+1},R_k), \tag{188}$$

Bounding the summation of $V_{k,1}^{\pi_k,\tilde{P}_k}(s_1,R_k)$ In the above analysis, we derive the upper bound of Eqn 130 and Eqn 131 (Inq 170 and Inq 187). Now we can finally bound $\sum_{k=1}^{K}V_{k,1}^{\pi_k,\tilde{P}_k}(s_{k,1},R_k)$. For a constant c_4,

$$\sum_{k=1}^{K}V_{k,1}^{\pi_k,\tilde{P}_k}(s_{k,1},R_k) \tag{189}$$

$$\leq c_4\sqrt{(dH^4+d^2H^3)K\log^3(HK/\delta))}+c_4(d^{2.5}H^2+d^2H^3)\log^3(HK/\delta)+\frac{1}{H}\sum_{k=1}^{K}\sum_{h=1}^{H}V_{k,h+1}^{\pi_k,\tilde{P}_k}(s_{k,h+1},R_k). \tag{190}$$

Following the same analysis, the above inequality actually also holds for any step $h \in [H]$:

$$\sum_{k=1}^{K}V_{k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},R_k) \tag{191}$$

$$\leq c_4\sqrt{(dH^4+d^2H^3)K\log^3(HK/\delta))}+c_4(d^{2.5}H^2+d^2H^3)\log^3(HK/\delta)+\frac{1}{H}\sum_{k=1}^{K}\sum_{h_1=h}^{H}V_{k,h_1+1}^{\pi_k,\tilde{P}_k}(s_{k,h_1+1},R_k). \tag{192}$$

Define $G = c_4\left(\sqrt{(dH^4+d^2H^3)K\log^3(HK/\delta))}+c_4(d^{2.5}H^2+d^2H^3)\log^3(HK/\delta)\right)$ and $a_h = \sum_{k=1}^{K}V_{k,h}^{\pi_k,\tilde{P}_k}(s_{k,h},R_k)$. The above inequality can be simplified into the following form:

$$a_h \leq G + \frac{1}{H}\sum_{h_1=h}^{H}a_{h_1}, \tag{193}$$

and we have $a_{H+1}=0$. From the elementary calculation, we can prove that $a_1 \leq (1+\frac{1}{H})^H G \leq eG$. Therefore, we have

$$\sum_{k=1}^{K}V_{k,1}^{\pi_k,\tilde{P}_k}(s_{k,1},R_k) \tag{194}$$

$$\leq c_5\left(\sqrt{(dH^4+d^2H^3)K\log^3(HK/\delta))}+(d^{2.5}H^2+d^2H^3)\log^3(HK/\delta)\right). \tag{195}$$

for a constant $c_5 = ec_4$. □

C AUXILIARY LEMMAS

Lemma 16. *(Self-Normalized Bound for Vector-Valued Martingales, Theorem 1 and 2 in Abbasi-Yadkori et al. (2011)) Let $\{F_t\}_{t=0}^{\infty}$ be a filtration. Let $\{\eta_t\}_{t=1}^{\infty}$ be a real-valued stochatic process such that η_t is F_t-measurable and η_t is conditionally R-sub-Gaussian for some $R \geq 0$, i.e.*

$$\forall \lambda \in \mathbb{R}, \mathbb{E}[e^{\lambda \eta_t} | F_{t-1}] \leq \exp\left(\frac{\lambda^2 R^2}{2}\right). \tag{196}$$

Let $\{X_t\}_{t=1}^{\infty}$ be an $\mathbb{R}^d$-valued stochastic process such that X_t is F_{t-1}-measurable. Assume that V is a $d \times d$ positive definite matrix. For any $t \geq 0$, define

$$\bar{V}_t = V + \sum_{s=1}^{t} X_s X_s^\top, S_t = \sum_{s=1}^{t} \eta_s X_s. \tag{197}$$

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

$$\|S_t\|^2_{\bar{V}_t^{-1}} \leq 2R^2 \log\left(\frac{\det\left(\bar{V}_t\right)^{1/2} \det(V)^{-1/2}}{\delta}\right). \tag{198}$$

Further, let $V = I\lambda, \lambda > 0$. Define $Y_t = \langle X_t, \theta_* \rangle + \eta_t$ and assume that $\|\theta_\star\|_2 \leq S$, $\|X_t\|_2 \leq L, \forall t \geq 1$. Then for any $\delta > 0$,with probability at least $1 - \delta$, for all $t \geq 0$, $\theta_\star$ satisfies

$$\left\|\hat{\theta}_t - \theta\right\|_{\bar{V}_t} \leq R\sqrt{d \log\left(\frac{1 + tL^2/\lambda}{\delta}\right)} + \lambda^{1/2} S. \tag{199}$$

Lemma 17. *(Lemma 11 in Abbasi-Yadkori et al. (2011)) Let $\{X_t\}_{t=1}^{\infty}$ be a sequence in $\mathbb{R}^d$, V a $d \times d$ positive definite matrix and define $\bar{V}_t = V + \sum_{s=1}^{t} X_s X_s^\top$. Then, we have that*

$$\log\left(\frac{\det\left(\bar{V}_n\right)}{\det(V)}\right) \leq \sum_{t=1}^{n} \|X_t\|^2_{\bar{V}_{t-1}^{-1}}. \tag{200}$$

Further , if $\|X_t\|_2 \leq L$ for all t, then

$$\sum_{t=1}^{n} \min\left\{1, \|X_t\|^2_{\bar{V}_{t-1}^{-1}}\right\} \leq 2\left(\log\det\left(\bar{V}_n\right) - \log\det V\right) \leq 2\left(d \log\left(\left(\operatorname{trace}(V) + nL^2\right)/d\right) - \log\det V\right), \tag{201}$$

and finally, if $\lambda_{\min}(V) \geq \max(1, L^2)$, then

$$\sum_{t=1}^{n} \|X_t\|^2_{\bar{V}_{t-1}^{-1}} \leq 2 \log \frac{\det\left(\bar{V}_n\right)}{\det(V)}. \tag{202}$$

Lemma 18. *(Bernstein inequality for vector-valued martingales, Theorem 4.1 in Zhou et al. (2020a)) Let $\{\mathcal{G}_t\}_{t=1}^{\infty}$ be a filtration, $\{x_t, \eta_t\}_{t \geq 1}$ a stochastic process so that $x_t \in \mathbb{R}^d$ is $\mathcal{G}_t$-measurable and $\eta_t \in \mathbb{R}$ is $\mathcal{G}_{t+1}$-measurable. Fix $R, L, \sigma, \lambda > 0, \mu^* \in \mathbb{R}^d$. For $t \geq 1$, let $y_t = \langle \mu^*, x_t \rangle + \eta_t$ and suppose that η_t, x_t also satisfy*

$$|\eta_t| \leq R, \mathbb{E}\left[\eta_t \mid \mathcal{G}_t\right] = 0, \mathbb{E}\left[\eta_t^2 \mid \mathcal{G}_t\right] \leq \sigma^2, \|x_t\|_2 \leq L. \tag{203}$$

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

$$\forall t > 0, \left\|\sum_{i=1}^{t} x_i \eta_i\right\|_{z_t^{-1}} \leq \beta_t, \|\mu_t - \mu^*\|_{z_t} \leq \beta_t + \sqrt{\lambda}\|\mu^*\|_2, \tag{204}$$

where for $t \geq 1, \mu_t = Z_t^{-1} b_t, Z_t = \lambda I + \sum_{i=1}^{t} x_i x_i^\top, b_t = \sum_{i=1}^{t} y_i x_i$ and

$$\beta_t = 8\sigma\sqrt{d \log\left(1 + tL^2/(d\lambda)\right) \log\left(4t^2/\delta\right)} + 4R \log\left(4t^2/\delta\right). \tag{205}$$