
Improved No-Regret Algorithms for Stochastic Shortest Path with Linear MDP

Liyu Chen¹ Rahul Jain¹ Haipeng Luo¹

Abstract

We introduce two new no-regret algorithms for the stochastic shortest path (SSP) problem with a linear MDP that significantly improve over the only existing results of (Vial et al., 2021). Our first algorithm is computationally efficient and achieves a regret bound $\tilde{O}(\sqrt{d^3 B_*^2 T_* K})$, where d is the dimension of the feature space, B_* and T_* are upper bounds of the expected costs and hitting time of the optimal policy respectively, and K is the number of episodes. The same algorithm with a slight modification also achieves logarithmic regret of order $\mathcal{O}\left(\frac{d^3 B_*^4}{c_{\min}^2 \text{gap}_{\min}} \ln^5 \frac{dB_* K}{c_{\min}}\right)$, where gap_{\min} is the minimum sub-optimality gap and c_{\min} is the minimum cost over all state-action pairs. Our result is obtained by developing a simpler and improved analysis for the finite-horizon approximation of (Cohen et al., 2021) with a smaller approximation error, which might be of independent interest. On the other hand, using variance-aware confidence sets in a global optimization problem, our second algorithm is computationally inefficient but achieves the first “horizon-free” regret bound $\tilde{O}(d^{3.5} B_* \sqrt{K})$ with no polynomial dependency on T_* or $1/c_{\min}$, almost matching the $\Omega(dB_* \sqrt{K})$ lower bound from (Min et al., 2021).

1. Introduction

We study the stochastic shortest path (SSP) model, where a learner attempts to reach a goal state while minimizing her costs in a stochastic environment. SSP is a suitable model for many real-world applications, such as games, car navigation, robotic manipulation, etc. Online reinforcement learning in SSP has received great attention recently. In this setting, learning proceeds in K episodes over a Markov Decision Process (MDP). In each episode, starting from a

fixed initial state, the learner sequentially takes an action, incurs a cost, and transits to the next state until reaching the goal state. The performance of the learner is measured by her regret, the difference between her total costs and that of the optimal policy. SSP is a strict generalization of the heavily-studied finite-horizon reinforcement learning problem, where the learner is guaranteed to reach the goal state after a fixed number of steps.

Modern reinforcement learning applications often need to handle a massive state space, in which function approximation is necessary. There is huge progress in the study of linear function approximation, for both the finite-horizon setting (Ayoub et al., 2020; Jin et al., 2020b; Yang & Wang, 2020; Zanette et al., 2020a;b; Zhou et al., 2021a) and the infinite-horizon setting (Wei et al., 2021b; Zhou et al., 2021a;b). Recently, Vial et al. (2021) took the first step in considering linear function approximation for SSP. They study SSP defined over a *linear MDP*, and proposed a computationally inefficient algorithm with regret $\tilde{O}(\sqrt{d^3 B_*^3 K/c_{\min}})$, as well as another efficient algorithm with regret $\tilde{O}(K^{5/6})$ (omitting other dependency). Here, d is the dimension of the feature space, B_* is an upper bound on the expected costs of the optimal policy, and c_{\min} is the minimum cost across all state-action pairs. Later, Min et al. (2021) study a related but different SSP problem defined over a *linear mixture MDP* and achieve a $\tilde{O}(dB_*^{1.5} \sqrt{K}/c_{\min})$ regret bound. Despite leveraging the advances from both the finite-horizon and infinite-horizon settings, results above are still far from optimal in terms of regret guarantee or computational efficiency, demonstrating the unique challenge of SSP problems.

In this work, we further extend our understanding of SSP with linear function approximation (more specifically, with linear MDPs). Our contributions are as follows:

- In Section 3, we first propose a new analysis for the finite-horizon approximation of SSP introduced in (Cohen et al., 2021), which is much simpler and achieves a smaller approximation error. Our analysis is also *model agnostic*, meaning that it does not make use of the modeling assumption and can be applied to both the tabular setting and function approximation settings. Combining this new analysis with a simple finite-horizon algorithm similar to that of (Jin et al.,

¹University of Southern California. Correspondence to: Liyu Chen <liyuc@usc.edu>.

2020b), we achieve a regret bound of $\tilde{O}(\sqrt{d^3 B_*^2 T_* K})$, with $T_* \leq B_*/c_{\min}$ being an upper bound of the hitting time of the optimal policy, which strictly improves over that of (Vial et al., 2021). Notably, unlike their algorithm, ours is computationally efficient without any extra assumption.

- In Section 3.3, we further show that the same algorithm above with a slight modification achieves a logarithmic instance-dependent expected regret bound of order $\mathcal{O}\left(\frac{d^3 B_*^4}{c_{\min}^2 \text{gap}_{\min}} \ln^5 \frac{dB_* K}{c_{\min}}\right)$ where gap_{\min} is some sub-optimality gap. As far as we know, this is the first logarithmic regret bound for SSP (with or without function approximation). We also establish a lower bound of order $\Omega\left(\frac{dB_*^2}{\text{gap}_{\min}}\right)$, which further advances our understanding for this problem even though it does not exactly match our upper bound.
- To remove the undesirable T_* dependency in our instance-independent bound, in Section 4, we further develop a computationally inefficient algorithm that makes use of certain variance-aware confidence sets in a global optimization problem and achieves $\tilde{O}(d^{3.5} B_* \sqrt{K})$ regret. Importantly, this bound is horizon-free in the sense that it has no polynomial dependency on T_* or $\frac{1}{c_{\min}}$ even in the lower order terms. Moreover, this almost matches the best known lower bound $\Omega(dB_* \sqrt{K})$ from (Min et al., 2021).

Techniques Our results are built upon several technical innovations. First, as mentioned, we develop an improved analysis for the finite-horizon approximation of (Cohen et al., 2021), which might be of independent interest. The key idea is to directly bound the total approximation error with respect to the regret bound of the finite-horizon algorithm, instead of analyzing the estimation precision for each state-action pair as done in (Cohen et al., 2021).

Second, to obtain the logarithmic bound in Section 3, we note that it is not enough to simply combine the aforementioned finite-horizon approximation and the existing logarithmic regret results for the finite-horizon setting such as (He et al., 2021), since the sub-optimality gap obtained in this way is in terms of the finite-horizon counterpart instead of the original SSP and could be substantially smaller. We resolve this issue via a longer horizon in the approximation and a careful two-stage analysis.

Finally, our horizon-free result in Section 4 is obtained by a novel combination of several ideas, including the global optimization algorithm of (Zanette et al., 2020b; Wei et al., 2021b), the variance-aware confidence sets of (Zhang et al., 2021) (for a related but different setting with linear mixture MDPs), an improved analysis of the variance-aware confidence sets (Kim et al., 2021), and finally a new clip-

ping trick and new update conditions that we propose. Our analysis does not require the recursion-based technique of (Zhang et al., 2020a) (for the tabular case), nor estimating higher order moments of value functions as in (Zhang et al., 2021) (for linear mixture MDPs), which might also be of independent interest.

Related work Regret minimization of SSP under stochastic costs has been well studied in the tabular setting (that is, no function approximation) (Tarbouriech et al., 2020; Cohen et al., 2020; 2021; Tarbouriech et al., 2021; Chen et al., 2021a; Jafarnia-Jahromi et al., 2021). There are also several works (Rosenberg & Mansour, 2020; Chen et al., 2021b; Chen & Luo, 2021) considering the more challenging setting with adversarial costs (which is beyond the scope of this work).

Beyond linear function approximation, in the finite-horizon setting researchers also start considering theoretical guarantees for general function approximation (Wang et al., 2020; Ishfaq et al., 2021; Kong et al., 2021). The study for SSP, which again is a strict generalization of the finite-horizon problems and might be a better model for many applications, falls behind in this regard, motivating us to explore in this direction with the goal of providing a more complete picture at least for linear function approximation.

The use of variance information is crucial in obtaining optimal regret bounds in MDPs. This dates back to the work of (Lattimore & Hutter, 2012) for the discounted setting, which has been significantly extended to the finite-horizon setting (Azar et al., 2017; Jin et al., 2018; Zanette & Brunskill, 2019; Zhang et al., 2020a;b). Constructing variance-aware confidence sets for linear bandits and linear mixture MDPs has also gained recent attention (Zhou et al., 2021a; Zhang et al., 2021; Kim et al., 2021). We are among the first to do so for linear MDPs (a concurrent work (Wei et al., 2021a) also does so but for a completely different purpose of improving robustness against corruption).

Logarithmic gap-dependent bounds have been shown in different settings; see for example (Jaksch et al., 2010; Simchowitz & Jamieson, 2019; Jin et al., 2021; He et al., 2021), but to our knowledge, we are the first to show similar bounds for SSP.

2. Preliminary

An SSP instance is defined by an MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, s_{\text{init}}, g, c, P)$. Here, \mathcal{S} is the state space, \mathcal{A} is the (finite) action space (with $A = |\mathcal{A}|$), $s_{\text{init}} \in \mathcal{S}$ is the initial state, $g \notin \mathcal{S}$ is the goal state, $c : \mathcal{S} \times \mathcal{A} \rightarrow [c_{\min}, 1]$ is the cost function with some global lower bound $c_{\min} \geq 0$, and $P = \{P_{s,a}\}_{(s,a) \in \mathcal{S} \times \mathcal{A}}$ with $P_{s,a} \in \Delta_{\mathcal{S}_+}$ is the transition function, where \mathcal{S}_+ is a shorthand for $\mathcal{S} \cup \{g\}$ and $\Delta_{\mathcal{S}_+}$ is

the simplex over \mathcal{S}_+ .

The learning protocol is as follows: the learner interacts with the environment for $K \geq 2$ episodes. In each episode, the learner starts in initial state s_{init} , sequentially takes an action, incurs a cost, and transits to the next state. An episode ends when the learner reaches the goal state g . We denote by (s_t, a_t, s'_t) the t -th state-action-state triplet observed among all episodes, so that $s'_t \sim P_{s_t, a_t}$ for each t , and $s'_t = s_{t+1}$ unless $s'_t = g$ (in which case $s_{t+1} = s_{\text{init}}$). Also denote by T the total number of steps in K episodes.

Learning objective The learner's goal is to learn a policy that reaches the goal state with minimum costs. Formally, a (stationary and deterministic) policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ is a mapping that assigns an action $\pi(s)$ to each state $s \in \mathcal{S}$. We say π is *proper* if following π (that is, taking action $\pi(s)$ whenever in state s) reaches the goal state with probability 1. Given a proper policy π , we define its value function and action-value function as follows:

$$V^\pi(s) = \mathbb{E} \left[\sum_{i=1}^I c(s_i, \pi(s_i)) \middle| P, s_1 = s \right],$$

$$Q^\pi(s, a) = c(s, a) + \mathbb{E}_{s' \sim P_{s, a}} [V^\pi(s')],$$

where the expectation in V^π is with respect to the randomness of next states $s_{i+1} \sim P_{s_i, \pi(s_i)}$ and the number of steps I before reaching the goal g . Let Π be the set of proper policies. We make the basic assumption that Π is non-empty. Under this assumption, there exists an optimal proper policy π^* , such that $V^{\pi^*}(s) = \min_a Q^{\pi^*}(s, a)$, and $V^{\pi^*}(s) = \min_{\pi \in \Pi} V^\pi(s)$ for all s (Bertsekas & Yu, 2013). We use V^* and Q^* as shorthands for V^{π^*} and Q^{π^*} . The formal goal of the learner is then to minimize her regret against π^* , that is, the difference between her total costs and that of the optimal proper policy, defined as

$$R_K = \sum_{t=1}^T c(s_t, a_t) - K \cdot V^*(s_{\text{init}}).$$

We also define $R_K = \infty$ if $T = \infty$.

Linear SSP In the so-called tabular setting, the state space is assumed to be small, and algorithms with computational complexity and regret bound depending on $S = |\mathcal{S}|$ are acceptable. To handle a potentially massive state space, however, we consider the same linear function approximation setting of (Vial et al., 2021), where the MDP enjoys a linear structure in both the transition and cost functions (known as linear or low-rank MDP).

Assumption 1 (Linear SSP). *For some $d \geq 2$, there exist known feature maps $\{\phi(s, a)\}_{(s, a) \in \mathcal{S} \times \mathcal{A}} \subseteq \mathbb{R}^d$, unknown parameters $\theta^* \in \mathbb{R}^d$ and $\{\mu(s')\}_{s' \in \mathcal{S}_+} \subseteq \mathbb{R}^d$, such that for*

any $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $s' \in \mathcal{S}_+$, we have:

$$c(s, a) = \phi(s, a)^\top \theta^*, \quad P_{s, a}(s') = \phi(s, a)^\top \mu(s').$$

Moreover, we assume $\|\phi(s, a)\|_2 \leq 1$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, $\|\theta^*\|_2 \leq \sqrt{d}$, and $\|\int h(s') d\mu(s')\|_2 \leq \sqrt{d} \|h\|_\infty$ for any $h \in \mathbb{R}^{\mathcal{S}_+}$.

We refer the reader to (Vial et al., 2021) and references therein for justification on this widely-used structural assumption (especially on the last few norm constraints). Under Assumption 1, by definition we have $Q^*(s, a) = \phi(s, a)^\top w^*$, where $w^* = \theta^* + \int V^*(s') d\mu(s') \in \mathbb{R}^d$, that is, Q^* is also linear in the features.

Key parameters and notations Two extra parameters that play a key role in our analysis are: $B_\star = \max_s V^*(s)$, the maximum cost of the optimal policy starting from any state, and $T_\star = \max_s T^{\pi^*}(s)$, the maximum hitting time of the optimal policy starting from any state, where $T^\pi(s)$ is the expected number of steps before reaching the goal if one follows policy π starting from state s . By definition, we have $T_\star \leq B_\star / c_{\min}$.

For simplicity, we assume that B_\star , T_\star , and c_{\min} are known to the learner for most discussions, and defer to the appendix what we can achieve when some of these parameters are unknown. We also assume $B_\star > 1$ and $c_{\min} > 0$ by default (and will discuss the case $c_{\min} = 0$ for specific algorithms if modifications are needed).

For $n \in \mathbb{N}_+$, we define $[n] = \{1, \dots, n\}$. For any $l \leq r$, we define $[x]_{[l, r]} = \min\{\max\{x, l\}, r\}$ as the projection of x onto the interval $[l, r]$. The notation $\tilde{O}(\cdot)$ hides all logarithmic terms including $\ln K$ and $\ln \frac{1}{\delta}$ for some confidence level $\delta \in (0, 1)$.

3. An Efficient Algorithm for Linear SSP

In this section, we introduce a computationally efficient algorithm for linear SSP. In Section 3.1, we first develop an improved analysis for the finite-horizon approximation of (Cohen et al., 2021). Then in Section 3.2, we combine this approximation with a simple finite-horizon algorithm, which together achieves $\tilde{O}(\sqrt{d^3 B_\star^2 T_\star K})$ regret. Finally, in Section 3.3, we further obtain a logarithmic regret bound via a slightly modified algorithm and a careful two-stage analysis.

3.1. Finite-Horizon Approximation of SSP

Finite-horizon approximation has been frequently used in solving SSP problems (Chen et al., 2021b; Chen & Luo, 2021; Cohen et al., 2021; Chen et al., 2021a). In particular, Cohen et al. (2021) proposed a black-box reduction from SSP to a finite-horizon MDP, which achieves minimax opti-

mal regret bound in the tabular case when combining with a certain finite-horizon algorithm. We will make use of the same algorithmic reduction in our proposed algorithm, but with an improved analysis.

Specifically, for an SSP instance $\mathcal{M} = (\mathcal{S}, \mathcal{A}, s_{\text{init}}, g, c, P)$, define its finite-horizon MDP counterpart as $\widetilde{\mathcal{M}} = (\mathcal{S}_+, \mathcal{A}, \widetilde{c}, c_f, \widetilde{P}, H)$, where $\widetilde{c}(s, a) = c(s, a)\mathbb{I}\{s \neq g\}$ is the extended cost function, $c_f(s) = 2B_*\mathbb{I}\{s \neq g\}$ is the terminal cost function (more details to follow), $\widetilde{P} = \{P_{s,a}\}_{(s,a) \in \mathcal{S} \times \mathcal{A}} \cup \{P_{g,a}\}_{a \in \mathcal{A}}$ with $P_{g,a}(s') = \mathbb{I}\{s' = g\}$ is the extended transition function, and H is a horizon parameter. Assume the access to a corresponding finite-horizon algorithm \mathfrak{A} which learns through a certain number of “intervals” following the protocol below. At the beginning of an interval m , the learner \mathfrak{A} is first reset to an arbitrary state s_1^m . Then, in each step $h = 1, \dots, H$ within this interval, \mathfrak{A} decides an action a_h^m , transits to $s_{h+1}^m \sim \widetilde{P}_{s_h^m, a_h^m}$, and suffers cost $\widetilde{c}(s_h^m, a_h^m)$. At the end of the interval, the learner suffers an additional terminal cost $c_f(s_{H+1}^m)$, and then moves on to the next interval.

With such a black-box access to \mathfrak{A} , the reduction of (Cohen et al., 2021) is depicted in Algorithm 1. The algorithm partitions the time steps into intervals of length $H \geq 4T_* \ln(4K)$ (such that π^* reaches g within H steps with high probability). In each step, the algorithm follows \mathfrak{A} in a natural way and feeds the observations to \mathfrak{A} (Line 7 and Line 9). If the goal state is not reached within an interval, \mathfrak{A} naturally enters the next interval with the initial state being the current state (Line 10). Otherwise, if the goal state is reached within some interval, we keep feeding g and zero cost to \mathfrak{A} until it finishes the current interval (Line 8 and Line 9), and after that, the next interval corresponds to the beginning of the next episode of the original SSP problem (Line 1).

Analysis Cohen et al. (2021) showed that in this reduction, the regret R_K of the SSP problem is very close to the regret of \mathfrak{A} in the finite-horizon MDP $\widetilde{\mathcal{M}}$. Specifically, define $\widetilde{R}_{M'} = \sum_{m=1}^{M'} (\sum_{h=1}^H c_h^m + c_f(s_{H+1}^m) - V_1^*(s_1^m))$ as the regret of \mathfrak{A} over the first M' intervals of $\widetilde{\mathcal{M}}$ (note the inclusion of the terminal costs), where V_1^* is the optimal value function of the first layer of $\widetilde{\mathcal{M}}$ (see Appendix B.1 for the formal definition). Denote by M the final (random) number of intervals created during the K episodes. Then Cohen et al. (2021) showed the following (a proof is included in Appendix B.2 for completeness).

Lemma 1. *Algorithm 1 ensures $R_K \leq \widetilde{R}_M + B_*$.*

This lemma suggests that it remains to bound the number of intervals M . The analysis of Cohen et al. (2021) does so by marking state-action pairs as “known” or “unknown” based on how many times they have been visited, and showing that in each interval, the learner either reaches an “unknown”

Algorithm 1 Finite-Horizon Approximation of SSP from (Cohen et al., 2021)

Input: Algorithm \mathfrak{A} for finite-horizon MDP $\widetilde{\mathcal{M}}$ with horizon $H \geq 4T_* \ln(4K)$.

Initialize: interval counter $m \leftarrow 1$.

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

```

1   Set  $s_1^m \leftarrow s_{\text{init}}$ .
2   while  $s_1^m \neq g$  do
3       Feed initial state  $s_1^m$  to  $\mathfrak{A}$ .
4       for  $h = 1, \dots, H$  do
5           Receive action  $a_h^m$  from  $\mathfrak{A}$ .
6           if  $s_h^m \neq g$  then
7               Play action  $a_h^m$ , observe cost  $c_h^m =$ 
8                $c(s_h^m, a_h^m)$  and next state  $s_{h+1}^m$ .
9           else Set  $c_h^m = 0$  and  $s_{h+1}^m = g$ .
10          Feed  $c_h^m$  and  $s_{h+1}^m$  to  $\mathfrak{A}$ .
11      Set  $s_1^{m+1} = s_{H+1}^m$  and  $m \leftarrow m + 1$ .
```

state-action pair or with high probability reaches the goal state. This analysis requires \mathfrak{A} to be “admissible” (defined through a set of conditions) and also heavily makes use of the tabular setting to keep track of the status of each state-action pair, making it hard to be directly generalized to function approximation settings. Furthermore, it also introduces T_* dependency in the upper bound of M , since the total cost for an interval where an “unknown” state-action pair is visited is trivially bounded by $H = \Omega(T_*)$.

Instead, we propose the following simple and improved analysis. The idea is to separate intervals into “good” ones within which the learner reaches the goal state, and “bad” ones within which the learner does not. Then, our key observation is that the regret in each bad interval is at least B_* — this is because the learner’s cost is at least $2B_*$ in such intervals by the choice of the terminal cost c_f , and the optimal policy’s expected cost is at most B_* . Therefore, if \mathfrak{A} is a no-regret algorithm, the number of bad intervals has to be small. More formally, based on this idea we can bound M directly in terms of the regret guarantee of \mathfrak{A} without requiring any extra properties from \mathfrak{A} , as shown in the following lemma.

Theorem 1. *Suppose that \mathfrak{A} enjoys the following regret guarantee with certain probability: $\widetilde{R}_m = \tilde{O}(\gamma_0 + \gamma_1 \sqrt{m})$ for some problem-dependent coefficients γ_0 and γ_1 (that are independent of m) and any number of intervals $m \leq M$. Then, with the same probability, the number of intervals created by Algorithm 1 satisfies $M = \tilde{O}\left(K + \frac{\gamma_1^2}{B_*^2} + \frac{\gamma_0}{B_*}\right)$.*

Proof. For any finite $M_\dagger \leq M$, we will show $M_\dagger = \tilde{O}\left(K + \frac{\gamma_1^2}{B_*^2} + \frac{\gamma_0}{B_*}\right)$, which then implies that M has to be finite and is upper bounded by the same quantity. To

do so, we define the set of good intervals $\mathcal{C}_g = \{m \in [M_\dagger] : s_{H+1}^m = g\}$ where the learner reaches the goal state, and also the total costs of the learner in interval m of $\widetilde{\mathcal{M}}$: $C^m = \sum_{h=1}^H c_h^m + c_f(s_{H+1}^m)$. By definition and the guarantee of \mathfrak{A} , we have

$$\begin{aligned} \widetilde{R}_{M_\dagger} &= \sum_{m \in \mathcal{C}_g} (C^m - V_1^*(s_1^m)) + \sum_{m \notin \mathcal{C}_g} (C^m - V_1^*(s_1^m)) \\ &\leq \widetilde{\mathcal{O}}\left(\gamma_0 + \gamma_1 \sqrt{M_\dagger}\right). \end{aligned} \quad (1)$$

Next, we derive lower bounds on $\sum_{m \in \mathcal{C}_g} (C^m - V_1^*(s_1^m))$ and $\sum_{m \notin \mathcal{C}_g} (C^m - V_1^*(s_1^m))$ respectively. First note that by Lemma 17 and $H \geq 4T_* \ln(4K)$, we have that π^* reaches the goal within H steps with probability at least $1 - 1/2K$. Therefore, executing π^* in an episode of $\widetilde{\mathcal{M}}$ leads to at most $B_* + \frac{2B_*}{2K} \leq \frac{3}{2}B_*$ costs in expectation, which implies $V_1^*(s) \leq \frac{3}{2}B_*$ for any s . By $|\mathcal{C}_g| \leq K$, we thus have

$$\sum_{m \in \mathcal{C}_g} (C^m - V_1^*(s_1^m)) \geq -\frac{3}{2}B_*K.$$

On the other hand, for $m \notin \mathcal{C}_g$, we have $C^m \geq 2B_*$ due to the terminal cost $c_f(s_{H+1}^m) = 2B_*$, and thus

$$\sum_{m \notin \mathcal{C}_g} (C^m - V_1^*(s_1^m)) \geq \frac{B_*}{2}(M_\dagger - |\mathcal{C}_g|) \geq \frac{B_*}{2}(M_\dagger - K).$$

Combining the two lower bounds above with Eq. (1), we arrive at $\frac{B_*}{2}M_\dagger \leq \widetilde{\mathcal{O}}(\gamma_0 + \gamma_1 \sqrt{M_\dagger}) + 2B_*K$. By Lemma 28, this implies $M_\dagger = \widetilde{\mathcal{O}}\left(K + \frac{\gamma_1^2}{B_*} + \frac{\gamma_0}{B_*}\right)$, finishing the proof. \square

Now plugging in the bound on M in Theorem 1 into Lemma 1, we immediately obtain the following corollary on a general regret bound for the finite-horizon approximation.

Corollary 2. *Under the same condition of Theorem 1, Algorithm 1 ensures $R_K = \widetilde{\mathcal{O}}\left(\gamma_1 \sqrt{K} + \frac{\gamma_1^2}{B_*} + \gamma_0 + B_*\right)$ (with the same probability stated in Theorem 1).*

Proof. Combining Lemma 1 and Theorem 1, we have $R_K \leq \widetilde{R}_M + B_* \leq \widetilde{\mathcal{O}}(\gamma_1 \sqrt{M} + \gamma_0 + B_*) \leq \widetilde{\mathcal{O}}\left(\gamma_1 \sqrt{K} + \frac{\gamma_1^2}{B_*} + \gamma_1 \sqrt{\frac{\gamma_0}{B_*}} + \gamma_0 + B_*\right)$. Further realizing $\gamma_1 \sqrt{\frac{\gamma_0}{B_*}} \leq \frac{1}{2}\left(\frac{\gamma_1^2}{B_*} + \gamma_0\right)$ by AM-GM inequality proves the statement. \square

Note that the final regret bound completely depends on the regret guarantee of the finite horizon algorithm \mathfrak{A} . In particular, in the tabular case, if we apply a variant of EB-SSP (Tarbouriech et al., 2021) that achieves $\widetilde{R}_m =$

$\widetilde{\mathcal{O}}(B_* \sqrt{SAm} + B_* S^2 A)$ (note the lack of polynomial dependency on H),¹ then Corollary 2 ensures that $R_K = \widetilde{\mathcal{O}}(B_* \sqrt{SAK} + B_* S^2 A)$, improving the results of (Cohen et al., 2021) and matching the best existing bounds of (Tarbouriech et al., 2021; Chen et al., 2021a); see Appendix B.5 for more details. This is not achievable by the analysis of (Cohen et al., 2021) due to the T_* dependency in the lower order term mentioned earlier.

More importantly, our analysis is *model agnostic*: it only makes use of the regret guarantee of the finite-horizon algorithm, and does not leverage any modeling assumption on the SSP instance. This enables us to directly apply our result to settings with function approximation. In Appendix B.6, we provide an example for SSP with a linear mixture MDP, which gives a regret bound $\widetilde{\mathcal{O}}(B_* \sqrt{dT_* K} + B_* d \sqrt{K})$ via combining Corollary 2 and the near optimal finite-horizon algorithm of (Zhou et al., 2021a).

3.2. Applying an Efficient Finite-Horizon Algorithm for Linear MDPs

Similarly, if there were a horizon-free algorithm for finite-horizon linear MDPs, we could directly combine it with Algorithm 1 and obtain a T_* -independent regret bound. However, to our knowledge, this is still open due to some unique challenge for linear MDPs.

Nevertheless, even combining Algorithm 1 with a horizon-dependent linear MDP algorithm already leads to significant improvement over the state-of-the-art for linear SSP. Specifically, the finite-horizon algorithm \mathfrak{A} we apply is a variant of LSVI-UCB (Jin et al., 2020b), which performs Least-Squares Value Iteration with an optimistic modification. The pseudocode is shown in Algorithm 2. Utilizing the fact that action-value functions are linear in the features for a linear MDP, in each interval m , we estimate the parameters $\{w_h^m\}_{h=1}^H$ of these linear functions by solving a set of least square linear regression problems using all observed data (Line 1), and we encourage exploration by subtracting a bonus term $\beta_m \|\phi(s, a)\|_{\Lambda_m^{-1}}$ in the definition of $\widehat{Q}_h^m(s, a)$ (Line 2). Then, we simply act greedily with respect to the truncated action-value estimates $\{Q_h^m\}_h$ (Line 3). Clearly, this is an efficient algorithm with polynomial (in d, H, m and A) time complexity for each interval m .

We refer the reader to (Jin et al., 2020b) for more explanation of the algorithm, and point out three key modifications we make compared to their version. First, Jin et al. (2020b) maintain a separate covariance matrix Λ_h^m for each layer h using data only from layer h , while we only maintain a single covariance matrix Λ^m using data across all layers (Line 3). This is possible (and resulting in a better regret

¹This variant is equivalent to applying EB-SSP on a homogeneous finite-horizon MDP.

Algorithm 2 Finite-Horizon Linear-MDP Algorithm

Parameters: $\lambda = 1$, $\beta_m = 50dB\sqrt{\ln(16BmHd/\delta)}$ where δ is the failure probability and $B \geq 1$.

Initialize: $\Lambda_1 = \lambda I$.

for $m = 1, \dots, M$ **do**

 Define $V_{H+1}^m(s) = c_f(s)$.

for $h = H, \dots, 1$ **do**

 Compute

$$w_h^m = \Lambda_m^{-1} \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'}(c_{h'}^{m'} + V_{h+1}^m(s_{h'+1}^{m'})),$$

 where $\phi_h^m = \phi(s_h^m, a_h^m)$.

 Define $\phi(g, a) = 0$ and

$$\widehat{Q}_h^m(s, a) = \phi(s, a)^\top w_h^m - \beta_m \|\phi(s, a)\|_{\Lambda_m^{-1}}$$

$$Q_h^m(s, a) = [\widehat{Q}_h^m(s, a)]_{[0, B]}$$

$$V_h^m(s) = \min_a Q_h^m(s, a)$$

for $h = 1, \dots, H$ **do**

 Play $a_h^m = \operatorname{argmin}_a Q_h^m(s_h^m, a)$, suffer c_h^m , and transit to s_{h+1}^m .

 Compute $\Lambda_{m+1} = \Lambda_m + \sum_{h=1}^H \phi_h^m \phi_h^{m\top}$.

bound) since the transition function is the same in each layer of \widetilde{M} . Another modification is to define $V_{H+1}^m(s)$ as $c_f(s)$ simply for the purpose of incorporating the terminal cost. Finally, we project the action-value estimates onto $[0, B]$ for some parameter B similar to (Vial et al., 2021) (Line 2). In the main text we simply set $B = 3B_*$, and the upper bound truncation at B has no effect in this case. However, this projection will become important when learning without the knowledge of B_* (see Appendix B.4).

We show the following regret guarantee of Algorithm 2 following the analysis of (Vial et al., 2021) (see Appendix B.3).

Lemma 2. *With probability at least $1 - 4\delta$, Algorithm 2 with $B = 3B_*$ ensures $\widetilde{R}_m = \widetilde{O}(\sqrt{d^3 B_*^2 H m} + d^2 B_* H)$ for any $m \leq M$.*

Applying Corollary 2 we then immediately obtain the following new result for linear SSP.

Theorem 3. *Applying Algorithm 1 with $H = 4T_* \ln(4K)$ and \mathfrak{A} being Algorithm 2 with $B = 3B_*$ to the linear SSP problem ensures $R_K = \widetilde{O}(\sqrt{d^3 B_*^2 T_* K} + d^3 B_* T_*)$ with probability at least $1 - 4\delta$.*

There is some gap between our result above and the existing lower bound $\Omega(dB_*\sqrt{K})$ for this problem (Min et al., 2021). In particular, the dependency on T_* inherited from the H dependency in Lemma 2 is most likely unnecessary. Never-

theless, this already strictly improves over the best existing bound $\widetilde{O}(\sqrt{d^3 B_*^2 K/c_{\min}})$ from (Vial et al., 2021) since $T_* \leq B_*/c_{\min}$. Moreover, our algorithm is computationally efficient, while the algorithms of Vial et al. (2021) are either inefficient or achieve a much worse regret bound such as $\widetilde{O}(K^{5/6})$ (unless some strong assumptions are made). This improvement comes from the fact that our algorithm uses non-stationary policies (due to the finite-horizon approximation), which avoids the challenging problem of solving the fixed point of some empirical Bellman equation. This also demonstrates the power of finite-horizon approximation in solving SSP problems. On the other hand, obtaining the same regret guarantee by learning stationary policies only is an interesting future direction.

Learning without knowing B_* or T_* Note that the result of Theorem 3 requires the knowledge of B_* and T_* . Without knowing these parameters, we can still efficiently obtain a regret bound of order $\widetilde{O}(\sqrt{d^3 B_*^3 K/c_{\min}} + d^3 B_*^2/c_{\min})$, matching the bound of (Vial et al., 2021) achieved by their inefficient algorithm. See Appendix B.4 for details.

3.3. Logarithmic Regret

Many optimistic algorithms attain a more favorable regret bound of the form $C \ln K$, where C is an instance dependent constant usually inversely proportional to some gap measure; see e.g. (Jaksch et al., 2010) for the infinite-horizon setting and (Simchowitz & Jamieson, 2019) for the finite-horizon setting. In this section, we show that a slight modification of our algorithm also leads to an expected regret bound that is polylogarithmic in K and inversely proportional to $\operatorname{gap}_{\min} = \min_{s,a:\operatorname{gap}(s,a)>0} \operatorname{gap}(s, a)$ with $\operatorname{gap}(s, a) = Q^*(s, a) - V^*(s)$.²

The high-level idea is as follows. The first observation is that similarly to a recent work by He et al. (2021), we can show that our Algorithm 2 obtains a gap-dependent logarithmic regret bound $\widetilde{O}(\frac{\ln m}{\operatorname{gap}'_{\min}})$ for the finite-horizon problem. The caveat is that $\operatorname{gap}'_{\min}$ here is naturally defined using the optimal value and action-value functions V_h^* and Q_h^* for the finite-horizon MDP (which is different for each layer h); more specifically, $\operatorname{gap}'_{\min} = \min_{s,a,h:\operatorname{gap}_h(s,a)>0} \operatorname{gap}_h(s, a)$ where $\operatorname{gap}_h(s, a) = Q_h^*(s, a) - V_h^*(s)$. The difference between $\operatorname{gap}_{\min}$ and $\operatorname{gap}'_{\min}$ can in fact be significant; see Appendix B.7 for an example where $\operatorname{gap}'_{\min}$ is arbitrarily smaller than $\operatorname{gap}_{\min}$.

To get around this issue, we set H to be a larger value of order $\widetilde{O}(\frac{B_*}{c_{\min}})$ and perform the following two-stage analysis. For the first $H/2$ layers, we are able to show

²Note that for our definition of regret, a polylogarithmic bound is only possible in expectation, because even if the learner always executes π^* , the deviation of her total costs from $KV^*(s_{\text{init}})$ is already of order \sqrt{K} .

$Q_h^*(s, a) \approx Q^*(s, a)$ and thus $\text{gap}_h(s, a) \approx \text{gap}(s, a)$, leading to a $\tilde{O}(\frac{\ln m}{\text{gap}_{\min}})$ bound on the regret suffered for these layers. Then, for the last $H/2$ layers, we further consider two cases: if the learner’s policy for the first $H/2$ layers are nearly optimal, then the probability of not reaching the goal within the first $H/2$ layers is very low by the choice of H , and thus the costs suffered in the last $H/2$ layers are negligible; otherwise, we simply bound the costs using the number of times the learner takes a non-near-optimal action in the first $H/2$ layers, which is again shown to be of order $\tilde{O}(\frac{\ln m}{\text{gap}_{\min}})$.

One final detail is to carefully control the regret under some failure event that happens with a small probability (recall that we are aiming for an expected regret bound; see Footnote 2). This is necessary since in SSP the learner’s cost under such events could be unbounded in the worst case. To resolve this issue, we make a slight modification to Algorithm 1 and occasionally restart \mathfrak{A} whenever the number of total intervals reaches some multiple of a threshold; see Algorithm 7 in the appendix. This finally leads to our main result summarized in the following theorem (whose proof is deferred to Appendix B.8).

Theorem 4. *There exist b' and δ such that applying Algorithm 7 with horizon $H = \frac{b'B_*}{c_{\min}} \ln(\frac{dB_*K}{c_{\min}})$ and \mathfrak{A} being Algorithm 2 (with $B = 3B_*$ and failure probability δ) ensures $\mathbb{E}[R_K] = \mathcal{O}\left(\frac{d^3 B_*^4}{c_{\min}^2 \text{gap}_{\min}} \ln^5 \frac{dB_*K}{c_{\min}}\right)$.*

As far as we know, this is the first polylogarithmic bound for any SSP problem. Our result also indicates that the instance-dependent quantities of SSP can be well preserved after using some finite-horizon approximation.

Lower bounds To better understand instance-dependent regret bounds for this problem, we further show the following lower bound.

Theorem 5. *For any algorithm \mathfrak{A} , there exists a linear SSP instance with $d \geq 2$ and $B_* \geq 1$ such that $\mathbb{E}_{\mathfrak{A}}[R_K] = \Omega(dB_*^2/\text{gap}_{\min})$.*

This lower bound exhibits a relatively large gap from our upper bound. One important question is whether the $\frac{1}{c_{\min}}$ dependency in the upper bound is really necessary, which we leave as a future direction.

4. An Inefficient Horizon-Free Algorithm

Recall that the dominating term of the regret bound shown in Theorem 3 depends on T_* , which is most likely unnecessary. Due to the lack of a horizon-free algorithm for finite-horizon linear MDPs (which, as discussed, would have addressed this issue), in this section we propose a different approach leading to a computationally inefficient algorithm with a regret bound that is horizon-free (that is, no polynomial

Algorithm 3 Variance-Aware Global OPTimization with Optimism (VA-GOPO)

Initialize: $t = t' = 1, k = 1, s_1 = s_{\text{init}}, B_1 = 1$.

Define: $s'_0 = g$ and $V_t = V_{w_t, B_t}$.

while $k \leq K$ **do**

```

1   if  $s'_{t-1} = g$  or Eq. (4) holds or  $V_t(s_t) = 2B_t$  then
2       while True do
3           Compute  $w_t = \text{argmin}_{w \in \Omega_t(w, B_t)} V_{w, B_t}(s_t)$ 
              (see Eq. (2) and Eq. (3) for definitions).
4           if  $V_t(s_t) > B_t$  then  $B_t \leftarrow 2B_t$ ; else break.
          Record the most recent update time  $t' \leftarrow t$ .
        else  $(w_t, B_t) = (w_{t-1}, B_{t-1})$ .
          Take action  $a_t = \text{argmin}_a \phi(s_t, a)^\top w_t$ , suffer cost
           $c_t = c(s_t, a_t)$ , and transits to  $s'_t$ .
          if  $s'_t = g$  then  $s_{t+1} = s_{\text{init}}, k \leftarrow k + 1$ ; else  $s_{t+1} = s'_t$ .
          Increment time step  $t \leftarrow t + 1$ .

```

dependency on T_*) but has a worse dependency on d .

As stated in previous work for the tabular setting (Cohen et al., 2020; 2021; Tarbouriech et al., 2021; Chen et al., 2021a), achieving a horizon-free regret bound requires constructing variance-aware confidence sets on the transition functions. While this is straightforward in the tabular case, it is much more challenging with linear function approximation. Zhou et al. (2021a); Zhang et al. (2021) construct variance-aware confidence sets for linear mixture MDPs, but we are not aware of similar results for linear MDPs since they impose extra challenges. Our algorithm VA-GOPO, shown in Algorithm 3, is the first one to successfully make use of these ideas.

VA-GOPO follows a similar framework of the ELEANOR algorithm of (Zanette et al., 2020b) (for the finite-horizon setting) and the FOPO algorithm of (Wei et al., 2021b) (for the infinite-horizon setting) — they all maintain an estimate w_t of the true weight vector w^* (recall $Q^*(s, a) = \phi(s, a)^\top w^*$), found by optimistically minimizing the value of the current state s_t (roughly $\min_a \phi(s_t, a)^\top w_t$) over a confidence set of w_t , and then simply act according to $\text{argmin}_a \phi(s_t, a)^\top w_t$. The main differences are the construction of the confidence set and the conditions under which w_t is updated, which we explain in detail below.

Confidence Set For a parameter $B > 0$ and a weight vector $w \in \mathbb{R}^d$, inspired by (Zhang et al., 2021) we define a variance-aware confidence set for time step t as

$$\Omega_t(w, B) = \bigcap_{j \in \mathcal{J}_B} \Omega_t^j(w, B), \quad (2)$$

where $\mathcal{J}_B = \{\lceil \log_2 \epsilon \rceil, \dots, \lceil \log_2(6\sqrt{dB}) \rceil\}$ with $\epsilon = \frac{c_{\min}}{150d^3K}$, and $\Omega_t^j(w, B) = \mathbb{B}(3\sqrt{dB}) \cap$

$$\left\{ w' : \forall \nu \in \mathcal{G}_{\epsilon/t}(6\sqrt{dB}), \left| \sum_{i < t} \text{clip}_j(\phi_i^\top \nu) \epsilon_{V_{w,B}}^i(w') \right| \leq \sqrt{\sum_{i < t} \text{clip}_j^2(\phi_i^\top \nu) \eta_{V_{w,B}}^i(w') \iota_{B,t} + B2^j \iota_{B,t}} \right\}, \quad (3)$$

with $\mathbb{B}(r) = \{x \in \mathbb{R}^d : \|x\|_2 \leq r\}$ being the d -dimensional L_2 -ball of radius r , $\mathcal{G}_\xi(r) = \{\xi n, n \in \mathbb{Z}\}^d \cap \mathbb{B}(r)$ being the ξ -net of $\mathbb{B}(r)$, $\text{clip}_j(x) = [x]_{[-2^j, 2^j]}$ (recall $[x]_{[l,r]} = \min\{\max\{x, l\}, r\}$), ϕ_i being a shorthand of $\phi(s_i, a_i)$, $V_{w,B}(s) = \min_a [\phi(s, a)^\top w]_{[0, 2B]}$ (and $V_{w,B}(g) = 0$), $\epsilon_V^i(w') = \phi_i^\top w' - c_i - V(s'_i)$, $\eta_V^i(w') = \epsilon_V^i(w')^2$, and finally $\iota_{B,t} = 2^{11} d \ln \frac{48dBt}{\epsilon \delta}$ for some failure probability δ . The key difference between our confidence set and that of (Zhang et al., 2021) is in the definition of $\epsilon_V^i(w')$ and $\eta_V^i(w')$ due to the different structures between linear MDPs and linear mixture MDPs. In particular, we note that the value function V (more formally $V_{w,B}$) in our definitions is itself defined with respect to another weight vector w .

With this confidence set, when VA-GOPO decides to update w_t , it searches over all w such that $w \in \Omega_t(w, B_t)$ and finds the one that minimizes the value at the current state $V_{w,B_t}(s_t)$ (Line 2). Here, B_t is a running estimate of B_* . VA-GOPO maintains the inequality $V_t(s_t) \leq B_t$ during the update by doubling the value of B_t and repeating Line 2 whenever this is violated (Line 3). Note that the constraint $w \in \Omega_t(w, B_t)$ is in a sense *self-referential* — we consider w within a confidence set defined in terms of w itself, which is an important distinction compared to (Zhang et al., 2021) and is critical for linear MDPs.

To provide some intuition on our confidence set, denote V_{w_t, B_t} by V_t and $\Omega_t(w_t, B_t)$ by Ω_t . Note that if we ignore the dependency between V_t and $\{\phi_i\}_{i < t}$ (an issue that will eventually be addressed by some covering arguments), then $\{\epsilon_{V_t}^i(w')\}_{i < t}$ forms a martingale sequence when $w' = \tilde{w}_t \triangleq \theta^* + \int V_t(s') d\mu(s')$, and thus the inequality in Eq. (3) holds with high probability by some Bernstein-style concentration inequality (Lemma 36). Formally, this allows us to show the following.

Lemma 3. *With probability at least $1 - \delta$, $\tilde{w}_t \in \Omega_t, \forall t \geq 1$.*

Since w_t is also in Ω_t , the difference between $\phi(s, a)^\top w_t$ and $c(s, a) + \mathbb{E}_{s' \sim P_{s,a}}[V_t(s')] = \phi(s, a)^\top \tilde{w}_t$ is controlled by the size of the confidence set Ω_t , which is overall shrinking and thus making sure that w_t is getting closer and closer to w^* . In addition, we also show that V_t is optimistic at state s_t whenever an update is performed and that B_t never overestimates B_* significantly.

Lemma 4. *With probability at least $1 - \delta$, we have $V_t(s_t) \leq$*

$V^*(s_t)$ if an update (Line 2) is performed at time step t , and $B_t \leq 2B_*$ for all t .

Update Conditions VA-GOPO updates w_t whenever one of the three conditions in Line 1 is triggered. The first condition $s'_{t-1} = g$ simply indicates that the current step is the start of a new episode. The second condition is

$$\exists j \in \mathcal{J}_{B_t}, \nu \in \mathcal{G}_{\epsilon/t}(6\sqrt{dB_t}) : \Phi_t^j(\nu) > 8d^2 \Phi_{t'}^j(\nu), \quad (4)$$

where t' is the most recent update time step (Line 4) and $\Phi_t^j(\nu) = \sum_{i < t} f_j(\phi_i^\top \nu) + 2^j \|\nu\|_2^2$ with $f_j(x) = \text{clip}_j(x)x$. This lazy update condition makes sure that the algorithm does not update w_t too often (see Lemma 27) while still enjoying a small enough estimation error. The last condition $V_{t'}(s_t) = 2B_t$ (we call it *overestimate condition*) tests whether the current state has an overestimated value (note that $2B_t$ is the maximum value of $V_{t'}$ due to the truncation in its definition). This condition helps remove a $d^{1/4}$ factor in the regret bound without using some complicated ideas as in previous works; see Appendix C.4 for more explanation.

Regret Guarantee We prove the following regret guarantee for VA-GOPO, and provide a proof sketch in Appendix 4.1 followed by full proofs in the rest of Appendix C.

Theorem 6. *With probability at least $1 - 6\delta$, Algorithm 3 ensures $R_K = \tilde{O}(d^{3.5} B_* \sqrt{K} + d^7 B_*^2)$.*

Ignoring the lower order term, our bound is (potentially) suboptimal only in the d -dependency compared to the lower bound $\Omega(dB_* \sqrt{K})$ (Min et al., 2021). We note again that this is the first horizon-free regret bound for linear SSP: it does not have any polynomial dependency on T_* or $\frac{1}{c_{\min}}$ even in lower order terms. Furthermore, VA-GOPO also does not require the knowledge of B_* or T_* . For simplicity, we have assumed $c_{\min} > 0$. However, even when $c_{\min} = 0$, we can obtain essentially the same bound by running the same algorithm on a modified cost function (Appendix A).

4.1. Proof Sketch of Theorem 6

We focus on deriving the dominating term and ignore the lower order terms. For notational convenience, define $\iota_t = \iota_{B_t, t}$, $\mathcal{J}_t = \mathcal{J}_{B_t}$, $P_t = P_{s_t, a_t}$, and $C_t = \sum_{i=1}^t c(s_i, a_i)$. We divide the whole learning process into epochs indexed by l , and a new epoch begins whenever w_t is recomputed. Denote by $t_l + 1$ the first time step in epoch l , and for a quantity, function or set f_t indexed by time step t , we define $f_l = f_{t_l+1}$. Denote by l_t the epoch time step t belongs to, and we often ignore the subscript t when there is no confusion. Clearly, $V_t = V_l$, and similarly for $w_l, \tilde{w}_l, \iota_l, \Omega_l$ (ignoring the dependency on t for l).

By some straightforward calculation, we decompose the

regret as follows:

$$\begin{aligned}
 R_K \leq & \underbrace{\sum_{t=1}^T [V_l(s'_t) - P_t V_l]}_{\text{DEVIATION}} + \underbrace{\sum_{t=1}^T |\phi_t^\top (\tilde{w}_l - w_l)|}_{\text{ESTIMATION-ERR}} \\
 & + \underbrace{\sum_{l=1}^L \left(V_l(s_{t_{l+1}}) - V_l(s'_{t_{l+1}}) \right)}_{\text{SWITCHING-COST}} - K \cdot V^*(s_{\text{init}}).
 \end{aligned}$$

We bound each of these terms as follows.

Bounding DEVIATION This term is a sum of martingale difference sequence and is of order $\tilde{O}(\sqrt{\sum_{t=1}^T \mathbb{V}(P_t, V_l)})$. We then show that $\sum_{t=1}^T \mathbb{V}(P_t, V_l) \lesssim B_* C_T + B_*$. ESTIMATION-ERR (see Lemma 21).

Bounding ESTIMATION-ERR Here the variance-aware confidence set Ω_t comes into play. We define $\nu_t = \operatorname{argmax}_{\nu = \tilde{w}_l - w, w \in \Omega_l} |\phi_t^\top \nu|$. By $w_l \in \Omega_l$, we have $|\phi_t^\top (\tilde{w}_l - w_l)| \leq |\phi_t^\top \nu_t|$. Thus, it suffices to bound $\sum_{t=1}^T |\phi_t^\top \nu_t|$. As in (Kim et al., 2021), the main idea is to bound the matrix norm of ν_t w.r.t some special matrix by a variance-aware term, and then apply the elliptical potential lemma on $\{\phi_t\}_t$. We define a special covariance matrix $W_{j_t, l}(\nu) = 2^j I + \sum_{i < t} \min\{1, 2^j / |\phi_i^\top \nu|\} \phi_i \phi_i^\top$, such that $\Phi_t^j(\nu) = \|\nu\|_{W_{j_t, l}(\nu)}^2$. For any epoch l , $j \in \mathcal{J}_l$ and $\nu = \tilde{w}_l - \hat{w}$ with $\hat{w} \in \Omega_l$, we have the following key inequality (see Lemma 24):

$$\|\nu\|_{W_{j_t, l}(\nu)}^2 \lesssim 2^j \sqrt{\sum_{i \leq t} \mathbb{V}(P_i, V_l) \iota_l}. \quad (5)$$

One important step is thus to bound $\sum_{i \leq t_l} \mathbb{V}(P_i, V_l)$. Note that this term has a similar form of $\sum_{t=1}^T \mathbb{V}(P_t, V_l)$, and by a similar analysis (see Lemma 23):

$$\sum_{t \leq t_l} \mathbb{V}(P_t, V_l) \lesssim B_* C_{t_l} + B_* \sum_{t \leq t_l} |\phi_t^\top \nu'_t|. \quad (6)$$

where $\nu'_t = \operatorname{argmax}_{\nu = \tilde{w}_l - w, w \in \Omega_l} |\phi_t^\top \nu|$ (note that here l is fixed and independent of t). Define $j'_t \in \mathcal{J}_l$ such that $|\phi_t^\top \nu'_t| \in (2^{j'_t-1}, 2^{j'_t}]$. By Eq. (5):

$$\begin{aligned}
 |\phi_t^\top \nu'_t| & \lesssim \|\phi_t\|_{W_{j'_t, l}^{-1}(\nu'_t)} \|\nu'_t\|_{W_{j'_t, l}(\nu'_t)} \\
 & \lesssim \|\phi_t\|_{W_{j'_t, l}^{-1}(\nu'_t)} \sqrt{|\phi_t^\top \nu'_t|} \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_l) \iota_l}. \quad (7)
 \end{aligned}$$

Solving for $|\phi_t^\top \nu'_t|$ and by $\sum_{t \leq t_l} \|\phi_t\|_{W_{j'_t, l}^{-1}(\nu'_t)}^2 = \tilde{O}(d)$

(similar to elliptical potential lemma), we get

$$\sum_{t \leq t_l} |\phi_t^\top \nu'_t| = \tilde{O} \left(d \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_l) \iota_l} \right).$$

Plugging this back to Eq. (6) and solving a quadratic inequality, we get: $\sum_{i \leq t_l} \mathbb{V}(P_i, V_l) \lesssim B_* C_{t_l}$ (Lemma 23). Now by an analysis similar to Eq. (7) (Lemma 22):

$$\begin{aligned}
 \sum_{t=1}^T |\phi_t^\top \nu_t| & \lesssim d^2 \sum_{t=1}^T \|\phi_t\|_{W_{j_t, l}^{-1}(\nu_t)}^2 \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_l) \iota_l} \\
 & \lesssim d^{3.5} \sqrt{B_* C_T},
 \end{aligned}$$

where $j_t \in \mathcal{J}_t$ such that $|\phi_t^\top \nu_t| \in (2^{j_t-1}, 2^{j_t}]$. The extra d^2 factor is from the inequality $\Phi_t^j(\nu) \leq 8d^2 \Phi_l^j(\nu)$.

Bounding SWITCHING-COST By considering each condition of starting a new epoch, we show that SWITCHING-COST = $\tilde{O}(dB_* - L')$, where L' is the number of epochs started by triggering the overestimate condition; see Appendix C.3. We provide more tuition on including the overestimate condition in Appendix C.4. In short, it removes a factor of $d^{1/4}$ in the dominating term without incorporating unpractical decision sets as in previous works.

Putting Everything Together Combining the bounds above, we get $R_K = C_T - KV^*(s_{\text{init}}) \lesssim d^{3.5} \sqrt{B_* C_T}$. Solving a quadratic inequality w.r.t $\sqrt{C_T}$, we have $C_T \lesssim B_* K$. Plugging this back, we obtain $R_K \lesssim d^{3.5} B_* \sqrt{K}$.

5. Conclusion

In this work, we make significant progress towards better understanding of linear function approximation in the challenging SSP model. Two algorithms are proposed: the first one is efficient and achieves a regret bound strictly better than (Vial et al., 2021), while the second one is inefficient but achieves a horizon-free regret bound. In developing these results, we also propose several new techniques that might be of independent interest, especially the new analysis for the finite-horizon approximation of (Cohen et al., 2021).

A natural future direction is to close the gap between existing upper bounds and lower bounds in this problem, especially with an efficient algorithm. Another interesting direction is to study SSP with adversarially changing costs under linear function approximation.

Acknowledgments HL is supported by NSF Award IIS-1943607 and a Google Faculty Research Award. RJ's research was supported by the NSF awards ECCS-1810447 and CCF-1817212, and by ONR award N00014-20-1-2258.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 24:2312–2320, 2011.
- Ayoub, A., Jia, Z., Szepesvari, C., Wang, M., and Yang, L. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pp. 463–474. PMLR, 2020.
- Azar, M. G., Osband, I., and Munos, R. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pp. 263–272. PMLR, 2017.
- Bertsekas, D. P. and Yu, H. Stochastic shortest path problems under weak conditions. *Lab. for Information and Decision Systems Report LIDS-P-2909, MIT*, 2013.
- Chen, L. and Luo, H. Finding the stochastic shortest path with low regret: The adversarial cost and unknown transition case. In *International Conference on Machine Learning*, 2021.
- Chen, L., Jafarnia-Jahromi, M., Jain, R., and Luo, H. Implicit finite-horizon approximation and efficient optimal algorithms for stochastic shortest path. *Advances in Neural Information Processing Systems*, 2021a.
- Chen, L., Luo, H., and Wei, C.-Y. Minimax regret for stochastic shortest path with adversarial costs and known transition. In *Conference on Learning Theory*, pp. 1180–1215. PMLR, 2021b.
- Cohen, A., Kaplan, H., Mansour, Y., and Rosenberg, A. Near-optimal regret bounds for stochastic shortest path. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 8210–8219. PMLR, 2020.
- Cohen, A., Efroni, Y., Mansour, Y., and Rosenberg, A. Minimax regret for stochastic shortest path. *Advances in Neural Information Processing Systems*, 2021.
- He, J., Zhou, D., and Gu, Q. Logarithmic regret for reinforcement learning with linear function approximation. In *International Conference on Machine Learning*, pp. 4171–4180. PMLR, 2021.
- Ishfaq, H., Cui, Q., Nguyen, V., Ayoub, A., Yang, Z., Wang, Z., Precup, D., and Yang, L. F. Randomized exploration for reinforcement learning with general value function approximation. *International Conference on Machine Learning*, 2021.
- Jafarnia-Jahromi, M., Chen, L., Jain, R., and Luo, H. Online learning for stochastic shortest path model via posterior sampling. *arXiv preprint arXiv:2106.05335*, 2021.
- Jaksch, T., Ortner, R., and Auer, P. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pp. 4863–4873, 2018.
- Jin, C., Jin, T., Luo, H., Sra, S., and Yu, T. Learning adversarial Markov decision processes with bandit feedback and unknown transition. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 4860–4869, 2020a.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020b.
- Jin, T., Huang, L., and Luo, H. The best of both worlds: stochastic and adversarial episodic mdps with unknown transition. *Advances in Neural Information Processing Systems*, 2021.
- Kim, Y., Yang, I., and Jun, K.-S. Improved regret analysis for variance-adaptive linear bandits and horizon-free linear mixture mdps. *arXiv preprint arXiv:2111.03289*, 2021.
- Kong, D., Salakhutdinov, R., Wang, R., and Yang, L. F. Online sub-sampling for reinforcement learning with general function approximation. *arXiv preprint arXiv:2106.07203*, 2021.
- Lattimore, T. and Hutter, M. Pac bounds for discounted mdps. In *International Conference on Algorithmic Learning Theory*, pp. 320–334. Springer, 2012.
- Lattimore, T. and Szepesvári, C. *Bandit algorithms*. Cambridge University Press, 2020.
- Min, Y., He, J., Wang, T., and Gu, Q. Learning stochastic shortest path with linear function approximation. *arXiv preprint arXiv:2110.12727*, 2021.
- Rosenberg, A. and Mansour, Y. Stochastic shortest path with adversarially changing costs. *arXiv preprint arXiv:2006.11561*, 2020.
- Shani, L., Efroni, Y., Rosenberg, A., and Mannor, S. Optimistic policy optimization with bandit feedback. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 8604–8613, 2020.
- Simchowitz, M. and Jamieson, K. G. Non-asymptotic gap-dependent regret bounds for tabular mdps. *Advances in Neural Information Processing Systems*, 32:1153–1162, 2019.

- Tarbouriech, J., Garcelon, E., Valko, M., Pirotta, M., and Lazaric, A. No-regret exploration in goal-oriented reinforcement learning. In *International Conference on Machine Learning*, pp. 9428–9437. PMLR, 2020.
- Tarbouriech, J., Zhou, R., Du, S. S., Pirotta, M., Valko, M., and Lazaric, A. Stochastic shortest path: Minimax, parameter-free and towards horizon-free regret. *Advances in Neural Information Processing Systems*, 2021.
- Vial, D., Parulekar, A., Shakkottai, S., and Srikant, R. Regret bounds for stochastic shortest path problems with linear function approximation. *arXiv preprint arXiv:2105.01593*, 2021.
- Wang, R., Salakhutdinov, R., and Yang, L. F. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. *Advances in Neural Information Processing Systems*, 2020.
- Wei, C.-Y., Dann, C., and Zimmert, J. A model selection approach for corruption robust reinforcement learning. *arXiv preprint arXiv:2110.03580*, 2021a.
- Wei, C.-Y., Jahromi, M. J., Luo, H., and Jain, R. Learning infinite-horizon average-reward mdps with linear function approximation. In *International Conference on Artificial Intelligence and Statistics*, pp. 3007–3015. PMLR, 2021b.
- Yang, L. and Wang, M. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*, pp. 10746–10756. PMLR, 2020.
- Zanette, A. and Brunskill, E. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 7304–7312, 2019.
- Zanette, A., Brandfonbrener, D., Brunskill, E., Pirotta, M., and Lazaric, A. Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*, pp. 1954–1964. PMLR, 2020a.
- Zanette, A., Lazaric, A., Kochenderfer, M., and Brunskill, E. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pp. 10978–10989. PMLR, 2020b.
- Zhang, Z., Ji, X., and Du, S. S. Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. In *Conference On Learning Theory*, 2020a.
- Zhang, Z., Zhou, Y., and Ji, X. Almost optimal model-free reinforcement learning via reference-advantage decomposition. In *Advances in Neural Information Processing Systems*, volume 33, pp. 15198–15207. Curran Associates, Inc., 2020b.
- Zhang, Z., Yang, J., Ji, X., and Du, S. S. Variance-aware confidence set: Variance-dependent bound for linear bandits and horizon-free bound for linear mixture mdp. *Advances in Neural Information Processing Systems*, 2021.
- Zhou, D., Gu, Q., and Szepesvari, C. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*, pp. 4532–4576. PMLR, 2021a.
- Zhou, D., He, J., and Gu, Q. Provably efficient reinforcement learning for discounted mdps with feature mapping. In *International Conference on Machine Learning*, pp. 12793–12802. PMLR, 2021b.

A. Preliminary

Extra Notations in Appendix For a function $X : \mathcal{S}_+ \rightarrow \mathbb{R}$ and a distribution $P \in \Delta_{\mathcal{S}_+}$, we define $PX = \mathbb{E}_{S \sim P}[X(S)]$ and $\mathbb{V}(P, X) = \text{VAR}_{S \sim P}[X(S)]$.

Cost Perturbation for $c_{\min} = 0$ We follow the receipt in (Vial et al., 2021, Appendix A.3) to deal with zero costs: the main idea is to run the SSP algorithm with perturbed cost $c_\epsilon(s, a) = c(s, a) + \epsilon$ for some $\epsilon > 0$, which is equivalent to solving a different SSP instance $\mathcal{M}_\epsilon = (\mathcal{S}, \mathcal{A}, s_{\text{init}}, g, c_\epsilon, P)$. Let $\theta_\epsilon^* = \theta^* + \epsilon \sum_{s'} \mu(s')$. Then, $c_\epsilon(s, a) = \phi(s, a)^\top \theta_\epsilon^*$. Therefore, \mathcal{M}_ϵ is also a linear SSP with $c_{\min} = \epsilon$ (up to some a small constant, since $c_\epsilon(s, a)$ can be as large as $1 + \epsilon$). Denote by V_ϵ^* the optimal value function in \mathcal{M}_ϵ , and define $R'_K = \sum_t c_\epsilon(s_t, a_t) - KV_\epsilon^*(s_{\text{init}})$ as the regret in \mathcal{M}_ϵ . We have $V_\epsilon^*(s) \leq V^{\pi^*}(s) + \epsilon T_\star \leq B_\star + \epsilon T_\star$, and

$$R_K = \sum_{t=1}^T c(s_t, a_t) - KV^*(s_{\text{init}}) \leq \sum_{t=1}^T c_\epsilon(s_t, a_t) - KV_\epsilon^*(s_{\text{init}}) + K(V_\epsilon^*(s_{\text{init}}) - V^*(s_{\text{init}})) \leq R'_K + \epsilon T_\star K.$$

Therefore, by running an SSP algorithm on perturbed cost c_ϵ , we recover its regret guarantee with $c_{\min} \leftarrow \epsilon$, $B_\star \leftarrow B_\star + \epsilon T_\star$, and an addition bias $\epsilon T_\star K$ in regret.

B. Omitted Details for Section 3

Notations For $\widetilde{\mathcal{M}}$, denote by $V_h^\pi(s)$ the the expected cost of executing policy π starting from state s in layer h , and by π^m the policy executed in interval m (for example, $\pi^m(s, h) = \text{argmin}_a Q_h^m(s, a)$ in Algorithm 2). For notational convenience, define $P_h^m = \widetilde{P}_{s_h^m, a_h^m}$, and $w_h^* = \theta^* + \int V_{h+1}^*(s') d\mu(s')$ for $h \in [H]$ such that $Q_h^*(s, a) = \phi(s, a)^\top w_h^*$. Define indicator $\mathbb{I}_s(s') = \mathbb{I}\{s = s'\}$, and auxiliary feature $\phi(g, a) = \mathbf{0} \in \mathbb{R}^d$ for all $a \in \mathcal{A}$, such that $\widetilde{c}(s, a) = \phi(s, a)^\top \theta^*$ and $\widetilde{P}_{s, a} V = \phi(s, a)^\top \int V(s') d\mu(s')$ for any $s \in \mathcal{S}_+$, $a \in \mathcal{A}$ and $V : \mathcal{S}_+ \rightarrow \mathbb{R}$ with $V(g) = 0$. Finally, for Algorithm 2, define stopping time $\overline{M} = \inf_m \{m \leq M, \exists h \in [H] : \widehat{Q}_h^m(s_h^m, a_h^m) > Q_h^m(s_h^m, a_h^m)\}$, which is the number of intervals until finishing K episodes or upper bound truncation on Q estimate is triggered.

B.1. Formal Definition of Q_h^* and V_h^*

It is not hard to see that we can define Q_h^* and V_h^* recursively without resorting to the definition of $\widetilde{\mathcal{M}}$:

$$Q_h^*(s, a) = \widetilde{c}(s, a) + \widetilde{P}_{s, a} V_{h+1}^*, \quad V_h^*(s) = \min_a Q_h^*(s, a),$$

with $Q_{H+1}^*(s, a) = c_f(s)$ for all (s, a) .

B.2. Proof of Lemma 1

Proof. Denote by \mathcal{I}_k the set of intervals in episode k , and by m_k the first interval in episode k . We bound the regret in episode k as follows: by Lemma 17 and $H \geq 4T_\star \ln(4K)$, we have the probability that following π^* takes more than H steps to reach g in $\widetilde{\mathcal{M}}$ is at most $\frac{1}{2K}$. Therefore,

$$V_1^{\pi^*}(s) \leq V^{\pi^*}(s) + 2B_\star P(s_{H+1} \neq g | \pi^*, P, s_1 = s) \leq V^{\pi^*}(s) + \frac{B_\star}{K}.$$

Thus,

$$\begin{aligned} \sum_{m \in \mathcal{I}_k} \sum_{h=1}^H c_h^m - V^{\pi^*}(s_1^{m_k}) &\leq \sum_{m \in \mathcal{I}_k} \sum_{h=1}^H c_h^m - V_1^{\pi^*}(s_1^{m_k}) + \frac{B_\star}{K} \\ &= \sum_{m \in \mathcal{I}_k} \left(\sum_{h=1}^H c_h^m - V_1^{\pi^*}(s_1^m) \right) + \sum_{m \in \mathcal{I}_k} V_1^{\pi^*}(s_1^m) - V_1^{\pi^*}(s_1^{m_k}) + \frac{B_\star}{K} \\ &\leq \sum_{m \in \mathcal{I}_k} \left(\sum_{h=1}^H c_h^m + c_f(s_{H+1}^m) - V_1^{\pi^*}(s_1^m) \right) + \frac{B_\star}{K}. \\ (V_1^*(s) \leq V_1^{\pi^*}(s) \text{ and } \sum_{m \in \mathcal{I}_k} V_1^{\pi^*}(s_1^m) - V_1^{\pi^*}(s_1^{m_k}) &\leq 2B_\star(|\mathcal{I}_k| - 1) = \sum_{m \in \mathcal{I}_k} c_f(s_{H+1}^m)) \end{aligned}$$

Summing terms above over $k \in [K]$ and by the definition of R_K, \tilde{R}_M we obtain the desired result. \square

B.3. Proof of Lemma 2

We first bound the error of one-step value iteration w.r.t \hat{Q}_h^m and V_{h+1}^m , which is essential to our analysis.

Lemma 5. *For any $B \geq \max\{1, \max_s c_f(s)\}$, with probability at least $1 - \delta$, we have $0 \leq \tilde{c}(s, a) + \tilde{P}_{s,a} V_{h+1}^m - \hat{Q}_h^m(s, a) \leq 2\beta_m \|\phi(s, a)\|_{\Lambda_m^{-1}}$ and $V_h^m(s) \leq V_h^*(s)$ for any $m \in \mathbb{N}_+, h \in [H]$.*

Proof. Define $\tilde{w}_h^m = \theta^* + \int V_{h+1}^m(s') d\mu(s')$, so that $\phi(s, a)^\top \tilde{w}_h^m = \tilde{c}(s, a) + \tilde{P}_{s,a} V_{h+1}^m$. Then,

$$\begin{aligned} \tilde{w}_h^m - w_h^m &= \Lambda_m^{-1} \left(\Lambda_m \tilde{w}_h^m - \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'} (c_{h'}^{m'} + V_{h+1}^m(s_{h'+1}^{m'})) \right) \\ &= \lambda \Lambda_m^{-1} \tilde{w}_h^m + \underbrace{\Lambda_m^{-1} \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'} (P_{h'}^{m'} V_{h+1}^m - V_{h+1}^m(s_{h'+1}^{m'}))}_{\epsilon_h^m}. \end{aligned}$$

By $V_{h+1}^m(s) \leq B$ and Lemma 31, we have with probability at least $1 - \delta$, for any $m, h \in [H]$:

$$\|\epsilon_h^m\|_{\Lambda_m} \leq 2B \sqrt{\frac{d}{2} \ln \left(\frac{mH + \lambda}{\lambda} \right)} + \ln \frac{\mathcal{N}_\varepsilon}{\delta} + \frac{\sqrt{8mH\varepsilon}}{\sqrt{\lambda}}, \quad (8)$$

where \mathcal{N}_ε is the ε -cover of the function class of V_{h+1}^m with $\varepsilon = \frac{1}{mH}$. Note that $V_{h+1}^m(s)$ is either $c_f(s)$ or

$$V_{h+1}^m(s) = \left[\min_a \phi(s, a)^\top w - \beta_m \sqrt{\phi(s, a)^\top \Gamma \phi(s, a)} \right]_{[0, B]},$$

for some PSD matrix Γ such that $\frac{1}{\lambda + mH} \leq \lambda_{\min}(\Gamma) \leq \lambda_{\max}(\Gamma) \leq \frac{1}{\lambda}$ by the definition of Λ_m^{-1} , and for some $w \in \mathbb{R}^d$ such that $\|w\|_2 \leq \lambda_{\max}(\Gamma) \times mH \times \sup_{s,a} \|\phi(s, a)\|_2 \times (B + 1) \leq \frac{mH}{\lambda} (B + 1)$ by the definition of w_h^m . We denote by \mathcal{V} the function class of V_{h+1}^m . Now we apply Lemma 32 to \mathcal{V} with $\alpha = (w, \Gamma)$, $n = d^2 + d$, $D = mH\sqrt{d}(B + 1)/\lambda \geq \max\{\frac{mH}{\lambda}(B + 1), \sqrt{d/\lambda^2}\}$ (note that $|\Gamma_{i,j}| \leq \|\Gamma\|_F = \sqrt{\sum_{i=1}^d \lambda_i^2(\Gamma)} \leq \sqrt{d/\lambda^2}$), and $L = \beta_m \sqrt{\lambda + mH}$, which is given by $||x|_{[0, B]} - |y|_{[0, B]}\| \leq |x - y|$ (Vial et al., 2021, Claim 2) and the following calculation: for any $\Delta w = \varepsilon e_i$ for some $\varepsilon \neq 0$,

$$\frac{1}{|\varepsilon|} |(w + \Delta w)^\top \phi(s, a) - w^\top \phi(s, a)| = |e_i^\top \phi(s, a)| \leq \|\phi(s, a)\| \leq 1,$$

and for any $\Delta \Gamma = \varepsilon e_i e_j^\top$,

$$\begin{aligned} & \frac{1}{|\varepsilon|} \left| \beta_m \sqrt{\phi(s, a)^\top (\Gamma + \Delta \Gamma) \phi(s, a)} - \beta_m \sqrt{\phi(s, a)^\top \Gamma \phi(s, a)} \right| \\ & \leq \beta_m \frac{|\phi(s, a)^\top e_i e_j^\top \phi(s, a)|}{\sqrt{\phi(s, a)^\top \Gamma \phi(s, a)}} \quad (\sqrt{u+v} - \sqrt{u} \leq \frac{|v|}{\sqrt{u}}) \\ & \leq \beta_m \frac{|\phi(s, a)^\top (\frac{1}{2} e_i e_i^\top + \frac{1}{2} e_j e_j^\top) \phi(s, a)|}{\sqrt{\phi(s, a)^\top \Gamma \phi(s, a)}} \quad (|ab| \leq \frac{1}{2}(a^2 + b^2)) \\ & \leq \beta_m \frac{\phi(s, a)^\top \phi(s, a)}{\sqrt{\phi(s, a)^\top \Gamma \phi(s, a)}} \leq \frac{\beta_m}{\sqrt{\lambda_{\min}(\Gamma)}} \leq \beta_m \sqrt{\lambda + mH}. \end{aligned}$$

Lemma 32 then implies $\ln \mathcal{N}_\varepsilon \leq (d^2 + d) \ln \frac{32d^{2.5} B m^2 H^2 \beta_m}{\lambda \varepsilon}$. Plugging this back, we get

$$\|\epsilon_h^m\|_{\Lambda_m} \leq \frac{\beta_m}{2}. \quad (9)$$

Moreover, $\|\tilde{w}_h^m\|_{\Lambda_m^{-1}} \leq \|\tilde{w}_h^m\|_2 / \sqrt{\lambda} \leq \sqrt{d/\lambda}(1+B)$. Thus,

$$\|w_h^m - \tilde{w}_h^m\|_{\Lambda_m} \leq \lambda \|\tilde{w}_h^m\|_{\Lambda_m^{-1}} + \|\epsilon_h^m\|_{\Lambda_m} \leq \beta_m.$$

Therefore, $\tilde{c}(s, a) + \tilde{P}_{s,a} V_{h+1}^m - \hat{Q}_h^m(s, a) = \phi(s, a)^\top (\tilde{w}_h^m - w_h^m) + \beta_m \|\phi(s, a)\|_{\Lambda_m^{-1}} \in [0, 2\beta_m \|\phi(s, a)\|_{\Lambda_m^{-1}}]$ by $\phi(s, a)^\top (\tilde{w}_h^m - w_h^m) \in [-\|\phi(s, a)\|_{\Lambda_m^{-1}} \|w_h^m - \tilde{w}_h^m\|_{\Lambda_m}, \|\phi(s, a)\|_{\Lambda_m^{-1}} \|w_h^m - \tilde{w}_h^m\|_{\Lambda_m}]$, and the first statement is proved. For any $m \in \mathbb{N}_+$, we prove the second statement by induction on $h = H+1, \dots, 1$. The base case $h = H+1$ is clearly true by $V_{h+1}^m(s) = V_{h+1}^*(s) = c_f(s)$. For $h \leq H$, we have by the induction step:

$$\hat{Q}_h^m(s, a) \leq \tilde{c}(s, a) + \tilde{P}_{s,a} V_{h+1}^m \leq \tilde{c}(s, a) + \tilde{P}_{s,a} V_{h+1}^* \leq Q_h^*(s, a).$$

Thus, $V_h^m(s) \leq \min_a \max\{0, \hat{Q}_h^m(s, a)\} \leq \min_a Q_h^*(s, a) = V_h^*(s)$. \square

Next, we prove a general regret bound, from which [Lemma 2](#) is a direct corollary.

Lemma 6. Assume $c_f(s) \leq H$. Then with probability at least $1 - 2\delta$, [Algorithm 2](#) ensures for any $M' \leq \bar{M}$

$$\tilde{R}_{M'} = \tilde{O}\left(\sqrt{d^3 B^2 H M'} + d^2 B H\right).$$

Proof. Define $c_{H+1}^m = c_f(s_{H+1}^m)$. Note that for $m < \bar{M}$, we have $V_h^m(s_h^m) = \max\{0, \hat{Q}_h^m(s_h^m, a_h^m)\}$, and with probability at least $1 - \delta$,

$$\begin{aligned} \sum_{h=1}^{H+1} c_h^m - V_1^*(s_1^m) &\leq \sum_{h=1}^{H+1} c_h^m - V_1^m(s_1^m) \leq \sum_{h=1}^{H+1} c_h^m - \hat{Q}_1^m(s_1^m, a_1^m) \leq \sum_{h=2}^{H+1} c_h^m - P_1^m V_2^m + 2\beta_m \|\phi(s_1^m, a_1^m)\|_{\Lambda_m^{-1}} \\ &\leq \sum_{h=2}^{H+1} c_h^m - V_2^m(s_2^m) + (\mathbb{I}_{s_2^m} - P_2^m) V_2^m + 2\beta_m \|\phi(s_1^m, a_1^m)\|_{\Lambda_m^{-1}} \\ &\leq \dots \leq \sum_{h=1}^H \left((\mathbb{I}_{s_{h+1}^m} - P_{h+1}^m) V_{h+1}^m + 2\beta_m \|\phi(s_h^m, a_h^m)\|_{\Lambda_m^{-1}} \right). \end{aligned} \quad (c_{H+1}^m = V_{H+1}^m(s_{H+1}^m))$$

([Lemma 5](#))

Therefore, by [Lemma 18](#) and [Lemma 38](#), with probability at least $1 - \delta$:

$$\begin{aligned} \tilde{R}_{M'} &\leq \tilde{R}_{M'-1} + H \leq \sum_{m=1}^{M'-1} \sum_{h=1}^H \left((\mathbb{I}_{s_{h+1}^m} - P_{h+1}^m) V_{h+1}^m + 2\beta_m \|\phi(s_h^m, a_h^m)\|_{\Lambda_m^{-1}} \right) + H \\ &= \tilde{O}\left(\sqrt{d^3 B^2 H M'} + d^2 B H\right). \end{aligned}$$

\square

We are now ready to prove [Lemma 2](#).

Proof of Lemma 2. Note that when $B = 3B_*$, $V_h^m(s) \leq V_h^*(s) \leq 3B_* = B$ by [Lemma 5](#). Thus, $\bar{M} = M$, and the statement directly follows from [Lemma 6](#) with $M' = \bar{M}$. \square

B.4. Learning without Knowing B_* or T_*

In this section, we develop a parameter-free algorithm that achieves $\tilde{O}(\sqrt{d^3 B_*^3 K / c_{\min}} + d^3 B_*^2 / c_{\min})$ regret without knowing B_* or T_* , which matches the best bound and knowledge of parameters of ([Vial et al., 2021](#)) while being computationally efficient under the most general assumption. Here we apply the finite-horizon approximation with zero terminal costs, and develop a new analysis on this approximation.

Algorithm 4 Adaptive Finite-Horizon Approximation of SSP

Input: upper bound estimate B and function $U(B)$ from Lemma 8.

Initialize: \mathfrak{A} an instance of finite-horizon algorithm with horizon $\lceil \frac{10B}{c_{\min}} \ln(8BK) \rceil$.

Initialize: $m = 1, m' = 0, k = 1, s = s_{\text{init}}$.

while $k \leq K$ **do**

 Execute \mathfrak{A} for H steps starting from state s and receive s_{H+1}^m .

if $s_{H+1}^m = g$ **then** $k \leftarrow k + 1, s \leftarrow s_{\text{init}}$; **else** $m' \leftarrow m' + 1, s \leftarrow s_{H+1}^m$.

if $m' > U(B)$ or \mathfrak{A} detects $B < B_*$ **then**
 $B \leftarrow 2B$.

 Initialize \mathfrak{A} as an instance of finite-horizon algorithm with horizon $\lceil \frac{10B}{c_{\min}} \ln(8BK) \rceil$.

 $m' \leftarrow 0$.

 $m \leftarrow m + 1$.

Finite-Horizon Approximation of SSP with Zero Terminal Costs To avoid knowledge of B_* or T_* , we apply finite-horizon approximation with zero terminal costs and horizon of order $\tilde{\mathcal{O}}(\frac{B}{c_{\min}})$ for some estimate B of B_* , that is, running Algorithm 1 with $c_f(s) = 0$ and $H = \tilde{\mathcal{O}}(\frac{B}{c_{\min}})$. We show that in this case there is an alternative way to bound the regret R_K by \tilde{R}_M , and there is a tighter bound on the total number of intervals M when $B \geq B_*$.

Lemma 7. Algorithm 1 with $c_f(s) = 0$ ensures $R_K \leq \tilde{R}_M + B_* \sum_{m=1}^M \mathbb{I}\{s_{H+1}^m \neq g\}$.

Proof. Denote by \mathcal{I}_k the set of intervals in episode k . We have:

$$\begin{aligned} R_K &= \sum_{k=1}^K \left(\sum_{m \in \mathcal{I}_k} \sum_{h=1}^H c_h^m - V^*(s_{\text{init}}) \right) = \sum_{k=1}^K \left(\sum_{m \in \mathcal{I}_k} \left(\sum_{h=1}^H c_h^m - V_1^*(s_1^m) \right) + \sum_{m \in \mathcal{I}_k} V_1^*(s_1^m) - V^*(s_{\text{init}}) \right) \\ &\leq \tilde{R}_M + B_* \sum_{m=1}^M \mathbb{I}\{s_{H+1}^m \neq g\}. \end{aligned} \quad (V_1^*(s) \leq V^*(s) \leq B_* \text{ by } c_f(s) = 0)$$

□

Lemma 8. Suppose when $B \geq B_*$, \mathfrak{A} with horizon $H = \lceil \frac{10B}{c_{\min}} \ln(8BK) \rceil$ ensures $\max\{\tilde{R}_{M'}, \sum_{m=1}^{M'} (V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m))\} = \tilde{\mathcal{O}}(\gamma_0(B) + \gamma_1(B)\sqrt{M'})$ for any $M' \leq M$ with probability at least $1 - \delta$, where γ_0, γ_1 are functions of B and are independent of M' . Then Algorithm 1 with $c_f(s) = 0$ ensures with probability at least $1 - 4\delta$,

$$\sum_{m=1}^M \mathbb{I}\{s_{H+1}^m \neq g\} = \tilde{\mathcal{O}} \left(\gamma_0(B)/B + \gamma_1(B)^2/B^2 + \gamma_1(B)\sqrt{K}/B \right) \triangleq U(B).$$

Proof. For any finite $M' \leq M$, we will show $\sum_{m=1}^{M'} \mathbb{I}\{s_{H+1}^m \neq g\} = \tilde{\mathcal{O}}(\gamma_0(B)/B + \gamma_1(B)^2/B^2 + \gamma_1(B)\sqrt{K}/B)$, which then implies that $\sum_{m=1}^M \mathbb{I}\{s_{H+1}^m \neq g\}$ has to be finite and is upper bounded by the same quantity. Define $\tilde{V}_1^\pi(s) = \mathbb{E}[\sum_{h=1}^{H/2} c(s_h, a_h) | \pi, P, s_1 = s]$ as the expected cost for the first $H/2$ layers and \tilde{V}_1^* as the optimal value function for the first $H/2$ layers. By (Chen et al., 2021a, Lemma 1) and $B \geq B_*$, we have $V^*(s) - \tilde{V}_1^*(s) \in [0, \frac{1}{4K}]$ and $V^*(s) - V_1^*(s) \in [0, \frac{1}{4K}]$. Moreover, when $s_{H+1}^m \neq g$, we have $\sum_{h>H/2} c_h^m \geq 2B$. Denote by $P_m(\cdot)$ the conditional probability of certain event conditioning on the history before interval m . Then with probability at least $1 - \delta$,

$$\begin{aligned} 2B \sum_{m=1}^{M'} P_m(s_{H+1}^m \neq g) + \sum_{m=1}^{M'} (\tilde{V}_1^{\pi^m}(s_1^m) - \tilde{V}_1^*(s_1^m)) &\leq \frac{M'}{2K} + \sum_{m=1}^{M'} (V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m)) \\ (2B \sum_{m=1}^{M'} P_m(s_{H+1}^m \neq g) + \tilde{V}_1^{\pi^m}(s_1^m)) &\leq V_1^{\pi^m}(s_1^m) \text{ and } V_1^*(s) \leq V^*(s) \leq \tilde{V}_1^*(s) + \frac{1}{4K} \\ &\leq \frac{1}{2K} \sum_{m=1}^{M'} \mathbb{I}\{s_{H+1}^m \neq g\} + \tilde{\mathcal{O}} \left(\gamma_0(B) + \gamma_1(B)\sqrt{M'} \right) \quad (M' \leq K + \sum_{m=1}^{M'} \mathbb{I}\{s_{H+1}^m \neq g\} \text{ and guarantee of } \mathfrak{A}) \\ &\leq \frac{1}{K} \sum_{m=1}^{M'} P_m(s_{H+1}^m \neq g) + \tilde{\mathcal{O}} \left(\gamma_0(B) + \gamma_1(B)\sqrt{M'} \right). \end{aligned} \quad (\text{Lemma 39})$$

Then by $\tilde{V}_1^{\pi^m}(s_1^m) \geq \tilde{V}_1^*(s_1^m)$ and reorganizing terms, we get $\sum_{m=1}^{M'} P_m(s_{H+1}^m \neq g) = \tilde{\mathcal{O}}(\gamma_0(B)/B + \gamma_1(B)\sqrt{M'}/B)$. Again by Lemma 39, we have with probability at least $1 - \delta$:

$$\sum_{m=1}^{M'} \mathbb{I}\{s_{H+1}^m \neq g\} = \tilde{\mathcal{O}}\left(\sum_{m=1}^{M'} P_m(s_{H+1}^m \neq g)\right) = \tilde{\mathcal{O}}\left(\gamma_0(B)/B + \gamma_1(B)\sqrt{M'}/B\right).$$

By $M' \leq K + \sum_{m=1}^{M'} \mathbb{I}\{s_{H+1}^m \neq g\}$ and solving a quadratic inequality w.r.t $\sqrt{\sum_{m=1}^{M'} \mathbb{I}\{s_{H+1}^m \neq g\}}$, we get $\sum_{m=1}^{M'} \mathbb{I}\{s_{H+1}^m \neq g\} = \tilde{\mathcal{O}}(\gamma_0(B)/B + \gamma_1(B)^2/B^2 + \gamma_1(B)\sqrt{K}/B)$. Thus, we also get the same bound for $\sum_{m=1}^M \mathbb{I}\{s_{H+1}^m \neq g\}$. \square

Remark 1. Note that the result of Lemma 8 is similar to (Tarbouriech et al., 2020, Lemma 7), which also shows that the number of “bad” intervals is of order $\tilde{\mathcal{O}}(\sqrt{K})$. However, their result is derived by explicitly analyzing the transition confidence sets, while we only make use of the regret guarantee of the finite-horizon algorithm. Thus, our approach is again model-agnostic and directly applicable to linear function approximation while their result is not.

Note that Lemma 7 and Lemma 8 together implies a $\tilde{\mathcal{O}}(\sqrt{K})$ regret bound when $B \geq B_*$. Moreover, since the total number of “bad” intervals is of order $\tilde{\mathcal{O}}(\sqrt{K})$, we can properly bound the cost of running finite-horizon algorithm with wrong estimates on B_* . We now present an adaptive version of finite-horizon approximation of SSP (Algorithm 4) which does not require the knowledge of B_* or T_* . The main idea is to perform finite-horizon approximation with zero costs, and maintain an estimate B of B_* . The learner runs a finite-horizon algorithm with horizon of order $\tilde{\mathcal{O}}(\frac{B}{c_{\min}})$. Whenever \mathfrak{A} detects $B \leq B_*$, or the number of “bad” intervals is more than expected (Line 1), it doubles the estimate B and start a new instance of finite-horizon algorithm with the updated estimate. The guarantee of Algorithm 4 is summarized in the following theorem.

Theorem 7. Suppose \mathfrak{A} takes an estimate B as input, and when $B < B_*$, it has some probability of detecting the anomaly (the event $B < B_*$) and halts. Define stopping time $\bar{M}' = \min\{M, \inf_m \{\text{anomaly detected in episode } m\}\}$, and suppose for any $B \geq 1$, \mathfrak{A} with horizon $H = \lceil \frac{10B}{c_{\min}} \ln(8BK) \rceil$ ensures $\max\{\bar{R}_{M'}, \sum_{m=1}^{M'} V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m)\} = \tilde{\mathcal{O}}(\gamma_0(B) + \gamma_1(B)\sqrt{M'})$ for any $M' \leq \bar{M}'$, where $\gamma_0(B)/B, \gamma_1(B)/B$ are non-decreasing w.r.t B . Then, Algorithm 4 ensures $R_K = \tilde{\mathcal{O}}(\gamma_0(B_*) + \gamma_1(B_*)\sqrt{K} + \gamma_1(B_*)^2/B_*)$ with probability at least $1 - 4\delta$.

Proof. We divide the learning process into epochs indexed by ϕ based on the update of B , so that $B_1 = B$ (the input value) and $B_{\phi+1} = 2B_\phi$. Let $\phi^* = \min_\phi \{B_\phi \geq B_*\}$. Define the regret in epoch ϕ as $\bar{R}_\phi = C_\phi - \sum_{k \in \mathcal{K}_\phi} V^*(s_1^{\phi,k})$, where C_ϕ is the total costs suffered in epoch ϕ , \mathcal{K}_ϕ is the set of episodes overlapped with epoch ϕ , and $s_1^{\phi,k}$ is the initial state in episode k and epoch ϕ (note that an episode can overlap with multiple epochs). Clearly, $\sum_\phi |\mathcal{K}_\phi| \leq K + \phi^* \leq K + \mathcal{O}(\log_2 B_*)$. Note that \mathfrak{A} satisfies the assumptions in Lemma 8, since no anomaly will be detected when $B \geq B_*$. Thus in epoch ϕ^* , no new epoch will be started by Lemma 8. Moreover, by Lemma 7 and $B_{\phi^*} \leq 2B_*$, the regret is bounded by:

$$\bar{R}_{\phi^*} = \tilde{\mathcal{O}}\left(\gamma_0(B_*) + \gamma_1(B_*)\sqrt{K + U(B_*)} + B_*U(B_*)\right) = \tilde{\mathcal{O}}\left(\gamma_0(B_*) + \gamma_1(B_*)\sqrt{K} + \gamma_1(B_*)^2/B_*\right).$$

For $\phi < \phi^*$, by the conditions of starting a new epoch, the number of intervals that does not reach the goal is upper bounded by $U(B_\phi)$ and the total number of intervals in epoch ϕ is upper bounded by $K + U(B_\phi)$. Thus by Lemma 7 and the guarantee of \mathfrak{A} ,

$$\bar{R}_\phi = \tilde{\mathcal{O}}\left(\gamma_0(B_\phi) + \gamma_1(B_\phi)\sqrt{K + U(B_\phi)} + B_*U(B_\phi)\right) = \tilde{\mathcal{O}}\left(\gamma_0(B_*) + \gamma_1(B_*)\sqrt{K} + \gamma_1(B_*)^2/B_*\right),$$

where the last equality is by the fact that $\gamma_0(B), \gamma_1(B)$ and $U(B)$ are non-decreasing w.r.t B . Thus,

$$\begin{aligned} R_K &= \sum_\phi C_\phi - \sum_{k=1}^K V^*(s_{\text{init}}) = \sum_\phi \bar{R}_\phi + \sum_\phi \sum_{k \in \mathcal{K}_\phi} V^*(s_1^{\phi,k}) - \sum_{k=1}^K V^*(s_{\text{init}}) \\ &= \tilde{\mathcal{O}}\left(\gamma_0(B_*) + \gamma_1(B_*)\sqrt{K} + \gamma_1(B_*)^2/B_*\right). \end{aligned}$$

\square

Algorithm 5 MVP+
Input: an estimate B such that $B \geq B_*$.

Initialize: $n(s, a), n(s, a, s'), Q_h(s, a), V_h(s), V_{H+1}(s) = 0$ for $(s, a) \in \mathcal{S}_+ \times \mathcal{A}, s' \in \mathcal{S}_+, h \in [H]$.

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

 for $h = 1, \dots, H$ **do**

 Take action $a_t = \operatorname{argmin}_a Q_h(s_t, a)$, incur cost $\tilde{c}(s_t, a_t)$ and transit to $s'_t \sim \tilde{P}_{s_t, a_t}$.

 $n(s_t, a_t) \leftarrow n(s_t, a_t) + 1, n(s_t, a_t, s'_t) \leftarrow n(s_t, a_t, s'_t) + 1$.

 if $n(s, a) = 2^j$ for some $j \in \mathbb{N}$ **then**

 for $h = H, \dots, 1$ **do**

 for $(s, a) \in \mathcal{S} \times \mathcal{A}$ **do**

 $\iota_{s,a} \leftarrow 20 \ln \frac{2SA n(s,a)}{\delta}, \bar{P}_{s,a}(s') \leftarrow \frac{n(s,a,s')}{\max\{1, n(s,a)\}}$ for all $s' \in \mathcal{S}_+$.

 $b_h(s, a) \leftarrow \max \left\{ 7 \sqrt{\frac{\mathbb{V}(\bar{P}_{s,a}, V_{h+1}) \iota_{s,a}}{\max\{1, n(s,a)\}}}, \frac{49B \iota_{s,a}}{\max\{1, n(s,a)\}} \right\}$.

 $Q_h(s, a) \leftarrow \max\{0, \tilde{c}(s, a) + \bar{P}_{s,a} V_{h+1} - b_h(s, a)\}$.

 $V_h(s) = \operatorname{argmin}_a Q_h(s, a)$.

Theorem 8. Applying Algorithm 4 with Algorithm 2 as \mathfrak{A} to the linear SSP problem ensures $R_K = \tilde{O}(\sqrt{d^3 B_*^3 K / c_{\min}} + d^3 B_*^2 / c_{\min})$ with probability at least $1 - 4\delta$.

Proof. Note that $\bar{M}' = \bar{M}$ for Algorithm 2, and Lemma 6 ensures that Algorithm 2 satisfies assumptions of Theorem 7 with $\gamma_0(B) = d^2 B H$ and $\gamma_1(B) = \sqrt{d^3 B^2 H}$, where $H = \lceil \frac{10B}{c_{\min}} \ln(8BK) \rceil$ (bounding $\sum_{m=1}^{M'} (V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m))$) is straightforward following similar arguments in Lemma 6). Then by Theorem 7, we have: $R_K = \tilde{O}(\sqrt{d^3 B_*^3 K / c_{\min}} + d^3 B_*^2 / c_{\min})$. \square

Remark 2. Comparing the bound achieved by Theorem 8 with that of Theorem 3, we see that $\frac{B_*}{c_{\min}}$ is in place of T_* , making it a worse bound since $T_* \leq \frac{B_*}{c_{\min}}$. Previous works in SSP (Cohen et al., 2021; Tarbouriech et al., 2021; Chen et al., 2021a) suggest that algorithms that obtain a bound with dependency on $\frac{B_*}{c_{\min}}$ is easier to be made parameter-free compared to those with dependency on T_* . Our findings in this section are consistent with that in previous works.

B.5. Horizon-Free Regret in the Tabular Setting with Finite-Horizon Approximation

Here we present a finite-horizon algorithm (Algorithm 5) that achieves $\tilde{R}_m = \tilde{O}(B_* \sqrt{SAm} + B_* S^2 A)$ and thus gives $R_K = \tilde{O}(B_* \sqrt{SAK} + B_* S^2 A)$ when combining with Corollary 2. For simplicity we assume that the cost function is known. We can think of Algorithm 5 as a variant of EB-SSP, which is applied on a finite-horizon MDP with state space $\mathcal{S} \times [H]$ and the transition is shared across layers. Note that due to the loop-free structure of the MDP, the value iteration converges in one sweep. Thus, skewing the empirical transition as in (Tarbouriech et al., 2021) is unnecessary. Then by the analysis of EB-SSP and the fact that transition data is shared across layers, we obtain the same regret guarantee $\tilde{R}_m = \tilde{O}(B_* \sqrt{SAm} + B_* S^2 A)$ (it is not hard to see that the algorithm achieves anytime regret since its updates on parameters are independent of K).

B.6. Application to Linear Mixture MDP

In this section, we provide a direct application of our finite-horizon approximation to the linear mixture MDP setting. We first introduce the problem setting of linear mixture SSP following (Min et al., 2021).

Assumption 2 (Linear Mixture SSP). *The number of states and actions are finite: $|\mathcal{S} \times \mathcal{A}| < \infty$. For some $d \geq 2$, there exist a known cost function $c : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, a known feature map $\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S}_+ \rightarrow \mathbb{R}^d$, and an unknown vector $\theta^* \in \mathbb{R}^d$ with $\|\theta^*\|_2 \leq \sqrt{d}$, such that:*

- for any $(s, a), s' \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}_+$, we have $P_{s,a}(s') = \langle \phi(s'|s, a), \theta^* \rangle$;
- for any bounded function $F : \mathcal{S} \rightarrow [0, 1]$, we have $\|\phi_F(s, a)\|_2 \leq \sqrt{d}$, where $\phi_F(s, a) = \sum_{s'} \phi(s'|s, a) F(s') \in \mathbb{R}^d$.

Algorithm 6 UCRL-VTR-SSP

Initialize: $\lambda = 1, \widehat{\Sigma}_1, \widetilde{\Sigma}_1 = \lambda I, \widehat{b}_1, \widetilde{b}_1, \widehat{\theta}_1, \widetilde{\theta}_1 = \mathbf{0}$.

Define: $\widehat{\beta}_m = 8\sqrt{d \ln(1 + dmH/\lambda) \ln(4m^2 H^2/\delta)} + 4\sqrt{d} \ln(4m^2 H^2/\delta) + \sqrt{\lambda d}$.

Define: $\widetilde{\beta}_m = 72B_*^2 \sqrt{d \ln(1 + 81dmHB_*^4/\lambda) \ln(4m^2 H^2/\delta)} + 36B_*^2 \ln(4m^2 H^2/\delta) + \sqrt{\lambda d}$.

Define: $\check{\beta}_m = 8d\sqrt{\ln(1 + dmH/\lambda) \ln(4m^2 H^2/\delta)} + 4\sqrt{d} \ln(4m^2 H^2/\delta) + \sqrt{\lambda d}$.

for $m = 1, \dots, M$ **do**
for $h = H, \dots, 1$ **do**

$$\begin{aligned} Q_h^m(\cdot, \cdot) &= \widetilde{c}(\cdot, \cdot) + \langle \widehat{\theta}_m, \phi_{V_{h+1}^m}(\cdot, \cdot) \rangle - \widehat{\beta}_m \left\| \phi_{V_{h+1}^m}(\cdot, \cdot) \right\|_{\widehat{\Sigma}_m^{-1}}, \text{ where } V_{H+1}^m(s) = 2B_* \mathbb{I}\{s \neq g\}. \\ V_h^m(\cdot) &= \min_a [Q_h^m(\cdot, a)]_{[0, 3B_*]}. \end{aligned}$$

for $h = 1, \dots, H$ **do**

 Take action $a_h^m = \operatorname{argmin}_a Q_h^m(s_h^m, a)$, suffer cost c_h^m , and transit to next state s_{h+1}^m .

$$\nu_h^m = \left[\left\langle \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m), \widetilde{\theta}_m \right\rangle \right]_{[0, 9B_*^2]} - \left[\left\langle \phi_{V_{h+1}^m}(s_h^m, a_h^m), \widehat{\theta}_m \right\rangle \right]_{[0, 3B_*]}.$$

$$E_h^m = \min \left\{ 9B_*^2, \widetilde{\beta}_m \left\| \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m) \right\|_{\widetilde{\Sigma}_m^{-1}} \right\} + \min \left\{ 9B_*^2, 6B_* \check{\beta}_m \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) \right\|_{\widehat{\Sigma}_m^{-1}} \right\}.$$

$$\bar{\sigma}_h^m = \sqrt{\max\{9B_*^2/d, \nu_h^m + E_h^m\}}.$$

$$\widehat{\Sigma}_{m+1} = \widehat{\Sigma}_m + \sum_{h=1}^H (\bar{\sigma}_h^m)^{-2} \phi_{V_{h+1}^m}(s_h^m, a_h^m) \phi_{V_{h+1}^m}(s_h^m, a_h^m)^\top.$$

$$\widetilde{\Sigma}_{m+1} = \widetilde{\Sigma}_m + \sum_{h=1}^H \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m) \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m)^\top.$$

$$\widehat{b}_{m+1} = \widehat{b}_m + \sum_{h=1}^H (\bar{\sigma}_h^m)^{-2} V_{h+1}^m(s_{h+1}^m) \phi_{V_{h+1}^m}(s_h^m, a_h^m).$$

$$\widetilde{b}_{m+1} = \widetilde{b}_m + \sum_{h=1}^H V_{h+1}^m(s_{h+1}^m)^2 \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m).$$

$$\widehat{\theta}_{m+1} = (\widehat{\Sigma}_{m+1})^{-1} \widehat{b}_{m+1}, \widetilde{\theta}_{m+1} \leftarrow \widetilde{\Sigma}_{m+1}^{-1} \widetilde{b}_{m+1}.$$

We also assume B_* is known and $c_{\min} > 0$. Define $\widetilde{c}(s, a) = c(s, a) \mathbb{I}\{s \neq g\}$, $\widetilde{P} = \{P_{s,a}\}_{(s,a) \in S \times \mathcal{A}} \cup \{P_{g,a}\}_{a \in \mathcal{A}}$ with $P_{g,a}(s') = \mathbb{I}\{s' = g\}$ as before, and $\phi(s'|g, a) = \mathbb{I}\{s' = g\} \sum_{s''} \phi(s''|s_{\text{init}}, a)$. Note that by the definitions above, $\widetilde{P}_{s,a} F = \langle \phi_F(s, a), \theta^* \rangle$. Also define total costs $C_{M'} = \sum_{m=1}^{M'} \sum_{h=1}^H c_h^m$ for any $M' \in \mathbb{N}_+$. With our approximation scheme, it suffices to provide a finite-horizon algorithm. We start by stating the regret guarantee of the proposed finite-horizon algorithm (Algorithm 6).

Theorem 9. *Algorithm 6 ensures $\widetilde{R}_{M'} = \widetilde{O}(B_* \sqrt{dM'H} + B_* d \sqrt{M'} + B_* d^2 H + B_* d^{2.5})$ for any $M' \in \mathbb{N}_+$ with probability at least $1 - 5\delta$.*

Combining Algorithm 6 with our finite-horizon approximation, we get the following regret guarantee on linear mixture SSP.

Theorem 10. *Applying Algorithm 1 with $H = \lceil 4T_* \ln(4K) \rceil$ and Algorithm 6 as \mathcal{A} to the linear mixture SSP problem ensures $R_K = \widetilde{O}(B_* \sqrt{dT_* K} + B_* d \sqrt{K} + B_* d^2 T_* + B_* d^{2.5})$ with probability at least $1 - 5\delta$.*

Proof. This directly follows from Theorem 9 and Corollary 2 with $\gamma_0 = B_* d^2 H$ and $\gamma_1 = B_* \sqrt{dH} + B_* d$. \square

Note that our bound strictly improves over that of (Min et al., 2021), and it is minimax optimal when $d \geq T_*$. Now we introduce the proposed finite-horizon algorithm, which is a variant of (Zhou et al., 2021a, Algorithm 2). The high level idea is to construct Bernstein-style confidence sets on transition function and then compute value function estimate through empirical value iteration with bonus. We summarize the ideas in Algorithm 6. Before proving Theorem 9, we need the following key lemma regarding the confidence sets on transition function.

Lemma 9. *With probability at least $1 - 3\delta$, we have for all $m \in \mathbb{N}_+$, $\left\| \theta^* - \widehat{\theta}_m \right\|_{\widehat{\Sigma}_m} \leq \widehat{\beta}_m$ and $|\nu_h^m - \mathbb{V}(P_h^m, V_{h+1}^m)| \leq E_h^m$.*

Proof. For the first statement, we first prove that $\left\| \theta^* - \widehat{\theta}_m \right\|_{\widehat{\Sigma}_m} \leq \check{\beta}_m$ and $\left\| \theta^* - \widetilde{\theta}_m \right\|_{\widetilde{\Sigma}_m} \leq \widetilde{\beta}_m$ for $m \in \mathbb{N}_+$. We adopt the indexing by t in Section 2: for a given time step $t = (m-1)H + h$ that corresponds to (m, h) , that is, the h -th step in the m -th interval, define $\bar{\sigma}_t = \bar{\sigma}_h^m$, $V_t = V_{h+1}^m$, $\nu_t = \nu_h^m$, and $E_t = E_h^m$. We apply Lemma 33 with $\mathcal{F}_t = \sigma(s_{1:t}, a_{1:t})$,

$x_t = \bar{\sigma}_t^{-1} \phi_{V_t}(s_t, a_t)$, $y_t = \bar{\sigma}_t^{-1} V_t(s'_t)$, $\mu^* = \theta^*$, $\eta_t = \bar{\sigma}_t^{-1} (V_t(s'_t) - \langle \phi_{V_t}(s_t, a_t), \theta^* \rangle)$. Then, we have $Z_t = \widehat{\Sigma}_t$, $\mu_t = \widehat{\theta}_t$, where $\widehat{\Sigma}_t = \lambda I + \sum_{i=1}^t \bar{\sigma}_i^{-2} \phi_{V_i}(s_i, a_i) \phi_{V_i}(s_i, a_i)^\top$, $\widehat{\theta}_t = \widehat{\Sigma}_t^{-1} \widehat{b}_t$, and $\widehat{b}_t = \sum_{i=1}^t \bar{\sigma}_i^{-2} \phi_{V_i}(s_i, a_i) V_i(s'_i)$. Moreover,

$$|\eta_t| \leq R = \sqrt{d}, \mathbb{E}[\eta_t^2 | \mathcal{G}_t] \leq \sigma^2 = d, \|x_t\|_2 \leq L = d, \|\mu^*\|_2 = \|\theta^*\|_2 \leq \sqrt{d}.$$

Therefore, with probability at least $1 - \delta$, for any $t = (m-1)H$ for some $m \in \mathbb{N}_+$, which corresponds to $(m-1, H)$:

$$\left\| \widehat{\theta}_m - \theta^* \right\|_{\widehat{\Sigma}_m} \leq 8d \sqrt{\ln(1 + d^2 t / (d\lambda)) \ln(4t^2 / \delta)} + 4\sqrt{d} \ln(4t^2 / \delta) + \sqrt{\lambda d} \leq \check{\beta}_m.$$

Next, we apply Lemma 33 with $\mathcal{F}_t = \sigma(s_{1:t}, a_{1:t})$, $x_t = \phi_{V_t^2}(s_t, a_t)$, $y_t = V_t^2(s'_t)$, $\mu^* = \theta^*$, $\eta_t = V_t^2(s'_t) - \langle \phi_{V_t^2}(s_t, a_t), \theta^* \rangle$. Then, we have $Z_t = \widetilde{\Sigma}_t$, $\mu_t = \widetilde{\theta}_t$, where $\widetilde{\Sigma}_t = \lambda I + \sum_{i=1}^t \phi_{V_i^2}(s_i, a_i) \phi_{V_i^2}(s_i, a_i)^\top$, $\widetilde{\theta}_t = \widetilde{\Sigma}_t^{-1} \widetilde{b}_t$, and $\widetilde{b}_t = \sum_{i=1}^t \phi_{V_i^2}(s_i, a_i) V_i^2(s'_i)$. Moreover,

$$|\eta_t| \leq R = 9B_*^2, \mathbb{E}[\eta_t^2 | \mathcal{G}_t] \leq \sigma^2 = 81B_*^4, \|x_t\|_2 \leq L = 9B_*^2 \sqrt{d}, \|\mu^*\|_2 = \|\theta^*\|_2 \leq \sqrt{d}.$$

Therefore, with probability at least $1 - \delta$, for any $t = (m-1)H$ for some $m \in \mathbb{N}_+$, which corresponds to $(m-1, H)$:

$$\left\| \widetilde{\theta}_m - \theta^* \right\|_{\widetilde{\Sigma}_m} \leq 72B_*^2 \sqrt{d \ln(1 + 81tB_*^4 d / (d\lambda)) \ln(4t^2 / \delta)} + 36B_*^2 \ln(4t^2 / \delta) + \sqrt{\lambda d} \leq \check{\beta}_m.$$

Conditioned on the event $\mathcal{C} = \left\{ \left\| \theta^* - \widehat{\theta}_m \right\|_{\widehat{\Sigma}_m} \leq \check{\beta}_m, \left\| \theta^* - \widetilde{\theta}_m \right\|_{\widetilde{\Sigma}_m} \leq \check{\beta}_m, \forall m \in \mathbb{N}_+ \right\}$, we have for t corresponding to (m, h) :

$$\begin{aligned} & |\nu_t - \mathbb{V}(P_t, V_t)| \\ & \leq \left| \left[\left\langle \phi_{V_t^2}(s_t, a_t), \widetilde{\theta}_m \right\rangle \right]_{[0, 9B_*^2]} - \left\langle \phi_{V_t^2}(s_t, a_t), \theta^* \right\rangle \right| + \left| \left[\left\langle \phi_{V_t}(s_t, a_t), \widehat{\theta}_m \right\rangle \right]_{[0, 3B_*]}^2 - \left\langle \phi_{V_t}(s_t, a_t), \theta^* \right\rangle^2 \right| \\ & \leq \min \left\{ 9B_*^2, \left| \left\langle \phi_{V_t^2}(s_t, a_t), \widetilde{\theta}_m - \theta^* \right\rangle \right| \right\} + \min \left\{ 9B_*^2, 6B_* \left| \left\langle \phi_{V_t}(s_t, a_t), \widehat{\theta}_m - \theta^* \right\rangle \right| \right\} \\ & \leq \min \left\{ 9B_*^2, \left\| \phi_{V_t^2}(s_t, a_t) \right\|_{\widetilde{\Sigma}_m^{-1}} \left\| \widetilde{\theta}_m - \theta^* \right\|_{\widetilde{\Sigma}_m} \right\} + \min \left\{ 9B_*^2, 6B_* \left\| \phi_{V_t}(s_t, a_t) \right\|_{\widehat{\Sigma}_m^{-1}} \left\| \widehat{\theta}_m - \theta^* \right\|_{\widehat{\Sigma}_m} \right\} \\ & \leq \min \left\{ 9B_*^2, \check{\beta}_m \left\| \phi_{V_t^2}(s_t, a_t) \right\|_{\widetilde{\Sigma}_m^{-1}} \right\} + \min \left\{ 9B_*^2, 6B_* \check{\beta}_m \left\| \phi_{V_t}(s_t, a_t) \right\|_{\widehat{\Sigma}_m^{-1}} \right\} = E_t. \end{aligned}$$

Thus the second statement is proved. Now we show that $\left\| \theta^* - \widehat{\theta}_m \right\|_{\widehat{\Sigma}_m} \leq \widehat{\beta}_m$. We conditioned on event \mathcal{C} , and apply Lemma 33 with $\mathcal{F}_t = \sigma(s_{1:t}, a_{1:t})$, $x_t = \bar{\sigma}_t^{-1} \phi_{V_t}(s_t, a_t)$, $y_t = \bar{\sigma}_t^{-1} V_t(s'_t)$, $\mu^* = \theta^*$, $\eta_t = \bar{\sigma}_t^{-1} (V_t(s'_t) - \langle \phi_{V_t}(s_t, a_t), \theta^* \rangle)$. Then, we have $Z_t = \widehat{\Sigma}_t$, $\mu_t = \widehat{\theta}_t$. Moreover, $|\eta_t| \leq R = \sqrt{d}$, $\|x_t\|_2 \leq L = d$, and for t corresponding to (m, h) ,

$$\mathbb{E}[\eta_t^2 | \mathcal{G}_t] = \bar{\sigma}_t^{-2} \mathbb{V}(P_t, V_t) \leq \bar{\sigma}_t^{-2} (\nu_t + E_t) \leq 1.$$

Therefore, with probability at least $1 - \delta$, for any $t = (m-1)H$ for some $m \in \mathbb{N}_+$, which corresponds to $(m-1, H)$:

$$\left\| \widehat{\theta}_m - \theta^* \right\|_{\widehat{\Sigma}_m} \leq 8\sqrt{d \ln(1 + dt/\lambda) \ln(4t^2 / \delta)} + 4\sqrt{d} \ln(4t^2 / \delta) + \sqrt{\lambda d} \leq \widehat{\beta}_m.$$

This completes the proof. \square

We are now ready to prove Theorem 9.

Proof of Theorem 9. We condition on the event of Lemma 9, Lemma 10 and Lemma 11, which happens with probability at

least $1 - 4\delta$. We decompose the regret as follows: with probability at least $1 - \delta$,

$$\begin{aligned}
 \tilde{R}_{M'} &= \sum_{m=1}^{M'} \left(\sum_{h=1}^H c_h^m + c_f(s_{H+1}^m) - V_1^*(s_1^m) \right) \leq \sum_{m=1}^{M'} \left(\sum_{h=1}^H c_h^m + c_f(s_{H+1}^m) - V_1^m(s_1^m) \right) & (\text{Lemma 10}) \\
 &= \sum_{m=1}^{M'} \sum_{h=1}^H (c_h^m + V_{h+1}^m(s_{h+1}^m) - V_h^m(s_h^m)) & (c_f = V_{H+1}^m) \\
 &\leq \sum_{m=1}^{M'} \sum_{h=1}^H \left(V_{h+1}^m(s_{h+1}^m) - P_h^m V_{h+1}^m + \langle \theta^* - \hat{\theta}_m, \phi_{V_{h+1}^m}(s_h^m, a_h^m) \rangle + \hat{\beta}_m \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) \right\|_{\hat{\Sigma}_m^{-1}} \right) \\
 &\quad \left(V_h^m(s_h^m) \geq Q_h^m(s_h^m, a_h^m) = c_h^m + \langle \hat{\theta}_m, \phi_{V_{h+1}^m}(s_h^m, a_h^m) \rangle - \hat{\beta}_m \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) \right\|_{\hat{\Sigma}_m^{-1}} \right) \\
 &\leq \tilde{\mathcal{O}} \left(\sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H \mathbb{V}(P_h^m, V_{h+1}^m)} + B_\star + \sum_{m=1}^{M'} \sum_{h=1}^H \hat{\beta}_m \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) \right\|_{\hat{\Sigma}_m^{-1}} \right). & (\text{Lemma 38, Cauchy-Schwarz inequality, and Lemma 9})
 \end{aligned}$$

The first term is of order $\tilde{\mathcal{O}}(\sqrt{B_\star^2 M' + B_\star C_{M'}})$ by Lemma 11. For the third term, define $\mathcal{I} = \left\{ (m, h) \in [M'] \times [H] : \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_m^{-1}} \geq 1 \right\}$ and $\hat{\mathcal{I}} = \{m \in [M'] : \det(\hat{\Sigma}_{m+1}) > 2 \det(\hat{\Sigma}_m)\}$. Then,

$$\begin{aligned}
 &\sum_{m=1}^{M'} \sum_{h=1}^H \hat{\beta}_m \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) \right\|_{\hat{\Sigma}_m^{-1}} = \sum_{m=1}^{M'} \sum_{h=1}^H \hat{\beta}_m \bar{\sigma}_h^m \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_m^{-1}} \\
 &\leq \tilde{\mathcal{O}} \left(\sum_{(m,h) \in \mathcal{I}} B_\star d \right) + \sum_{m=1}^{M'} \sum_{h=1}^H \hat{\beta}_m \bar{\sigma}_h^m \min \left\{ 1, \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_m^{-1}} \right\} \quad (\hat{\beta}_m = \tilde{\mathcal{O}}(\sqrt{d}) \text{ and } V_{h+1}^m = \mathcal{O}(B_\star)) \\
 &\stackrel{(i)}{=} \tilde{\mathcal{O}} \left(B_\star d^2 H + \sum_{m \in \hat{\mathcal{I}}} B_\star d H + \sum_{m=1}^{M'} \sum_{h=1}^H \hat{\beta}_m \bar{\sigma}_h^m \min \left\{ 1, \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_{m+1}^{-1}} \right\} \right) \\
 &= \tilde{\mathcal{O}} \left(B_\star d^2 H + \hat{\beta}_{M'} \sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H (\bar{\sigma}_h^m)^2} \sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H \min \left\{ 1, \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_{m+1}^{-1}}^2 \right\}} \right) \\
 &\quad (|\hat{\mathcal{I}}| = \tilde{\mathcal{O}}(d) \text{ and Cauchy-Schwarz inequality}) \\
 &= \tilde{\mathcal{O}} \left(B_\star d^2 H + d \sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H (\bar{\sigma}_h^m)^2} \right), \quad (\hat{\beta}_{M'} = \tilde{\mathcal{O}}(\sqrt{d}) \text{ and Lemma 29})
 \end{aligned}$$

where in (i) we apply $\hat{\beta}_m \bar{\sigma}_h^m = \tilde{\mathcal{O}}(B_\star d)$, Lemma 30, and:

$$\begin{aligned}
 |\mathcal{I}| &= \sum_{m=1}^{M'} \sum_{h=1}^H \mathbb{I} \left\{ \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_m^{-1}}^2 \geq 1 \right\} \leq \sum_{m=1}^{M'} \sum_{h=1}^H \min \left\{ 1, \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_m^{-1}}^2 \right\} \\
 &\leq |\hat{\mathcal{I}}| H + \sqrt{2} \sum_{m=1}^{M'} \sum_{h=1}^H \min \left\{ 1, \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_{m+1}^{-1}}^2 \right\} & (\text{Lemma 30}) \\
 &= \tilde{\mathcal{O}}(dH). & (|\hat{\mathcal{I}}| = \tilde{\mathcal{O}}(d) \text{ and Lemma 29})
 \end{aligned}$$

It remains to bound $\sum_{m=1}^{M'} \sum_{h=1}^H (\bar{\sigma}_h^m)^2$. Note that

$$\begin{aligned}
 \sum_{m=1}^{M'} \sum_{h=1}^H (\bar{\sigma}_h^m)^2 &\leq \frac{9B_\star^2 M' H}{d} + \sum_{m=1}^{M'} \sum_{h=1}^H (v_h^m + E_h^m) \\
 &\leq \frac{9B_\star^2 M' H}{d} + \sum_{m=1}^{M'} \sum_{h=1}^H (\mathbb{V}(P_h^m, V_{h+1}^m) + 2E_h^m) \quad (\text{Lemma 9}) \\
 &\leq \frac{9B_\star^2 M' H}{d} + \tilde{\mathcal{O}} \left(B_\star^2 M' + B_\star C_{M'} + B_\star^2 d \sqrt{M' H} + B_\star d^{3/2} \sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H (\bar{\sigma}_h^m)^2 + B_\star^2 d^2 H} \right). \quad (\text{Lemma 11 and Lemma 12})
 \end{aligned}$$

By Lemma 28 and $\sqrt{M' H} \leq M'/d + dH$, we get $\sum_{m=1}^{M'} \sum_{h=1}^H (\bar{\sigma}_h^m)^2 = \tilde{\mathcal{O}}(\frac{B_\star^2 M' H}{d} + B_\star^2 M' + B_\star C_{M'} + B_\star^2 d^3 + B_\star^2 d^2 H)$. Putting everything together, we get:

$$\begin{aligned}
 \tilde{R}_{M'} &= \tilde{\mathcal{O}} \left(\sqrt{B_\star^2 M' + B_\star C_{M'}} + B_\star d^2 H + d \sqrt{\frac{B_\star^2 M' H}{d} + B_\star^2 M' + B_\star C_{M'} + B_\star^2 d^3 + B_\star^2 d^2 H} \right) \\
 &= \tilde{\mathcal{O}} \left(B_\star \sqrt{d M' H} + B_\star d \sqrt{M'} + d \sqrt{B_\star C_{M'}} + B_\star d^2 H + B_\star d^{2.5} \right).
 \end{aligned}$$

Now by $\tilde{R}_{M'} = C_{M'} - M' V_1^\star(s_1^m)$ and Lemma 28, we get: $C_{M'} = \tilde{\mathcal{O}}(B_\star M')$. Plugging this back, we get $\tilde{R}_{M'} = \tilde{\mathcal{O}}(B_\star \sqrt{d M' H} + B_\star d \sqrt{M'} + B_\star d^2 H + B_\star d^{2.5})$. \square

Lemma 10. Conditioned on the event of Lemma 9, $Q_h^m(s, a) \leq \tilde{c}(s, a) + \tilde{P}_{s,a} V_{h+1}^m$ and $V_h^m(s) \leq V_h^\star(s) \leq 3B_\star$.

Proof. Note that by Lemma 9:

$$\begin{aligned}
 \left\langle \hat{\theta}_m, \phi_{V_{h+1}^m}(s, a) \right\rangle - \hat{\beta}_m \left\| \phi_{V_{h+1}^m}(s, a) \right\|_{\tilde{\Sigma}_m^{-1}} &= \tilde{P}_{s,a} V_{h+1}^m + \left\langle \hat{\theta}_m - \theta^\star, \phi_{V_{h+1}^m}(s, a) \right\rangle - \hat{\beta}_m \left\| \phi_{V_{h+1}^m}(s, a) \right\|_{\tilde{\Sigma}_m^{-1}} \\
 &\leq \tilde{P}_{s,a} V_{h+1}^m + \left\| \hat{\theta}_m - \theta^\star \right\|_{\tilde{\Sigma}_m} \left\| \phi_{V_{h+1}^m}(s, a) \right\|_{\tilde{\Sigma}_m^{-1}} - \hat{\beta}_m \left\| \phi_{V_{h+1}^m}(s, a) \right\|_{\tilde{\Sigma}_m^{-1}} \leq \tilde{P}_{s,a} V_{h+1}^m.
 \end{aligned}$$

The first statement then follows from the definition of Q_h^m . For any $m \in \mathbb{N}_+$, we prove the second statement by induction on $h = H + 1, \dots, 1$. The base case $h = H + 1$ is clearly true by the definition of V_{H+1}^m . For $h \leq H$, note that $Q_h^m(s, a) \leq \tilde{c}(s, a) + \tilde{P}_{s,a} V_{h+1}^m \leq c(s, a) + \tilde{P}_{s,a} V^\star \leq Q^\star(s, a)$ by the induction step and the first statement. Thus, $V_h^m(s) \leq \max\{0, \min_a Q_h^m(s, a)\} \leq V^\star(s)$. \square

Lemma 11. Conditioned on the event of Lemma 10, with probability at least $1 - \delta$, $\sum_{m=1}^{M'} \sum_{h=1}^H \mathbb{V}(P_h^m, V_{h+1}^m) = \tilde{\mathcal{O}}(B_\star^2 M' + B_\star C_{M'})$ for any $M' \in \mathbb{N}_+$.

Proof. Conditioned on the event of Lemma 10, we have with probability at least $1 - \delta$:

$$\begin{aligned}
 \sum_{m=1}^{M'} \sum_{h=1}^H \mathbb{V}(P_h^m, V_{h+1}^m) &= \sum_{m=1}^{M'} \sum_{h=1}^H P_h^m (V_{h+1}^m)^2 - (P_h^m V_{h+1}^m)^2 \\
 &= \sum_{m=1}^{M'} \sum_{h=1}^H (P_h^m (V_{h+1}^m)^2 - V_{h+1}^m (s_{h+1}^m)^2) + \sum_{m=1}^{M'} \sum_{h=1}^H (V_{h+1}^m (s_{h+1}^m)^2 - V_h^m (s_h^m)^2) + \sum_{m=1}^{M'} \sum_{h=1}^H (V_h^m (s_h^m)^2 - (P_h^m V_{h+1}^m)^2) \\
 &\stackrel{(i)}{=} \tilde{\mathcal{O}} \left(\sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H \mathbb{V}(P_h^m, (V_{h+1}^m)^2) + B_\star^2 M' + B_\star C_{M'}} \right) \stackrel{(ii)}{=} \tilde{\mathcal{O}} \left(B_\star \sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H \mathbb{V}(P_h^m, V_{h+1}^m) + B_\star^2 M' + B_\star C_{M'}} \right).
 \end{aligned}$$

Here, (ii) is by Lemma 34, and (i) is by Lemma 38, $V_{H+1}^m(s) \leq 2B_\star$ and:

$$\begin{aligned} \sum_{m=1}^{M'} \sum_{h=1}^H V_h^m(s_h^m)^2 - (P_h^m V_{h+1}^m)^2 &= \sum_{m=1}^{M'} \sum_{h=1}^H (V_h^m(s_h^m, a_h^m) + P_h^m V_{h+1}^m)(V_h^m(s_h^m) - P_h^m V_{h+1}^m) \\ &= \sum_{m=1}^{M'} \sum_{h=1}^H (V_h^m(s_h^m, a_h^m) + P_h^m V_{h+1}^m)(\max\{0, Q_h^m(s_h^m, a_h^m)\} - P_h^m V_{h+1}^m) \leq 6B_\star C_M. \\ &\quad (0 \leq V_h^m(s) \leq 3B_\star \text{ and } Q_h^m(s_h^m, a_h^m) \leq c_h^m + P_h^m V_{h+1}^m \text{ by Lemma 10}) \end{aligned}$$

By Lemma 28, we get $\sum_{m=1}^{M'} \sum_{h=1}^H \mathbb{V}(P_h^m, V_{h+1}^m) = \tilde{\mathcal{O}}(B_\star^2 M' + B_\star C_M)$. \square

Lemma 12. $\sum_{m=1}^{M'} \sum_{h=1}^H E_h^m = \tilde{\mathcal{O}}\left(B_\star^2 d \sqrt{M'H} + B_\star d^{3/2} \sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H (\bar{\sigma}_h^m)^2} + B_\star^2 d^2 H\right)$ for any $M' \in \mathbb{N}_+$.

Proof. Note that:

$$\begin{aligned} \sum_{m=1}^{M'} \sum_{h=1}^H E_h^m &= \sum_{m=1}^{M'} \sum_{h=1}^H \min \left\{ 9B_\star^2, \tilde{\beta}_m \left\| \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m) \right\|_{\tilde{\Sigma}_{m-1}^{-1}} \right\} + \min \left\{ 9B_\star^2, 6B_\star \check{\beta}_m \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) \right\|_{\tilde{\Sigma}_{m-1}^{-1}} \right\} \\ &\leq \sum_{m=1}^{M'} \sum_{h=1}^H \tilde{\beta}_m \min \left\{ 1, \left\| \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m) \right\|_{\tilde{\Sigma}_{m-1}^{-1}} \right\} + 6B_\star \sum_{m=1}^{M'} \sum_{h=1}^H \check{\beta}_m \bar{\sigma}_h^m \min \left\{ 1, \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\tilde{\Sigma}_{m-1}^{-1}} \right\}. \\ &\quad (\tilde{\beta}_m \geq 9B_\star^2 \text{ and } \check{\beta}_m \bar{\sigma}_h^m \geq 3B_\star) \end{aligned}$$

For the first sum, define $\tilde{\mathcal{I}} = \{m \in [M'] : \det(\tilde{\Sigma}_{m+1}) > 2 \det(\tilde{\Sigma}_m)\}$. Then by Lemma 30,

$$\begin{aligned} &\sum_{m=1}^{M'} \sum_{h=1}^H \tilde{\beta}_m \min \left\{ 1, \left\| \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m) \right\|_{\tilde{\Sigma}_{m-1}^{-1}} \right\} \\ &\leq \tilde{\mathcal{O}} \left(\sum_{m \in \tilde{\mathcal{I}}} B_\star^2 \sqrt{dH} \right) + \sqrt{2} \sum_{m \notin \tilde{\mathcal{I}}} \sum_{h=1}^H \tilde{\beta}_m \min \left\{ 1, \left\| \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m) \right\|_{\tilde{\Sigma}_{m+1}^{-1}} \right\} \quad (\tilde{\beta}_m = \tilde{\mathcal{O}}(B_\star^2 \sqrt{d})) \\ &= \tilde{\mathcal{O}} \left(B_\star^2 d^{3/2} H + \tilde{\beta}_{M'} \sqrt{M'H \sum_{m=1}^{M'} \sum_{h=1}^H \min \left\{ 1, \left\| \phi_{(V_{h+1}^m)^2}(s_h^m, a_h^m) \right\|_{\tilde{\Sigma}_{m+1}^{-1}}^2 \right\}} \right). \\ &\quad (|\tilde{\mathcal{I}}| = \tilde{\mathcal{O}}(d) \text{ and Cauchy-Schwarz inequality}) \\ &= \tilde{\mathcal{O}} \left(B_\star^2 d^{3/2} H + B_\star^2 d \sqrt{M'H} \right). \quad (\text{Lemma 29}) \end{aligned}$$

For the second sum, similarly define $\hat{\mathcal{I}} = \{m \in [M'] : \det(\hat{\Sigma}_{m+1}) > 2 \det(\hat{\Sigma}_m)\}$. Then,

$$\begin{aligned} &6B_\star \sum_{m=1}^{M'} \sum_{h=1}^H \check{\beta}_m \bar{\sigma}_h^m \min \left\{ 1, \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_{m-1}^{-1}} \right\} \\ &\leq \tilde{\mathcal{O}} \left(\sum_{m \in \hat{\mathcal{I}}} B_\star^2 dH \right) + 6\sqrt{2} B_\star \check{\beta}_M \sum_{m \notin \hat{\mathcal{I}}} \sum_{h=1}^H \bar{\sigma}_h^m \min \left\{ 1, \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_{m+1}^{-1}} \right\} \quad (\check{\beta}_m \bar{\sigma}_h^m = \tilde{\mathcal{O}}(B_\star d)) \\ &= \tilde{\mathcal{O}} \left(B_\star^2 d^2 H + B_\star d \sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H (\bar{\sigma}_h^m)^2} \sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H \min \left\{ 1, \left\| \phi_{V_{h+1}^m}(s_h^m, a_h^m) / \bar{\sigma}_h^m \right\|_{\hat{\Sigma}_{m+1}^{-1}}^2 \right\}} \right). \\ &\quad (|\hat{\mathcal{I}}| = \tilde{\mathcal{O}}(d) \text{ and Cauchy-Schwarz inequality}) \\ &= \tilde{\mathcal{O}} \left(B_\star^2 d^2 H + B_\star d^{3/2} \sqrt{\sum_{m=1}^{M'} \sum_{h=1}^H (\bar{\sigma}_h^m)^2} \right). \quad (\text{Lemma 29}) \end{aligned}$$

\square

B.7. An instance of SSP with $\text{gap}'_{\min} \ll \text{gap}_{\min}$

Consider an SSP with four states $\{s_0, s_1, s_2, s_3\}$ and two actions $\{a_1, a_2\}$. At s_0 , we have $c(s_0, a) = 0$ and $P(s_1|s_0, a) = p$, $P(s_0|s_0, a) = 1 - p$ for $a \in \{a_1, a_2\}$ and some $p > 0$. At s_1 , we have $c(s_1, a_1) = 0$, $c(s_1, a_2) = \epsilon$, and $P(s_2|s_1, a_1) = 1$, $P(s_3|s_1, a_2) = 1$. At s_2 , we have $c(s_2, a_1) = c(s_2, a_2) = 1$ and $P(g|s_2, a) = q$, $P(s_1|s_2, a) = 1 - q$ for any a and some $q \in (0, 1)$. At s_3 , we have $c(s_3, a) = 0$ and $P(g|s_3, a) = 1$ for $a \in \{a_1, a_2\}$. The role of s_0 here is to create the possibility that the learner will visit state s_1 at any time step. Then under our finite-horizon approximation, we have

$$\text{gap}'_{\min} \leq \min_{(s,a): \text{gap}_H(s,a) > 0} \text{gap}_H(s,a) = c(s_1, a_2) - c(s_1, a_1) = \epsilon.$$

On the other hand, when $\frac{1}{q} > \epsilon$, $\text{gap}_{\min} = Q^*(s_1, a_1) - V^*(s_1) = \frac{1}{q} - \epsilon$, and $\frac{1}{q}$ can be arbitrarily large.

B.8. Omitted Details in Section 3.3

We first prove a lemma bounding $Q_h^*(s, a) - Q_h^m(s, a)$ and another lemma on regret decomposition w.r.t the gap functions $\text{gap}_h(s, a)$ in $\tilde{\mathcal{M}}$.

Lemma 13. *Suppose $B = 3B_*$. With probability at least $1 - \delta$, for all $m \in \mathbb{N}_+$, $h \in [H]$, and $(s, a) \in \mathcal{S}_+ \times \mathcal{A}$, Algorithm 2 ensures:*

$$0 \leq Q_h^*(s, a) - \widehat{Q}_h^m(s, a) \leq \tilde{P}_{s,a}(V_{h+1}^* - V_{h+1}^m) + 2\beta_m \|\phi(s, a)\|_{\Lambda_m^{-1}}.$$

Proof. Note that:

$$\begin{aligned} w_h^* - w_h^m &= \Lambda_m^{-1} \left(\lambda I + \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'} \phi_{h'}^{m'\top} \right) w_h^* - \Lambda_m^{-1} \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'} (c_{h'}^{m'} + V_{h+1}^m(s_{h'+1}^{m'})) \\ &= \lambda \Lambda_m^{-1} w_h^* + \Lambda_m^{-1} \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'} [c_{h'}^{m'} + P_{h'}^{m'} V_{h+1}^*] - \Lambda_m^{-1} \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'} (c_{h'}^{m'} + V_{h+1}^m(s_{h'+1}^{m'})) \\ &= \lambda \Lambda_m^{-1} w_h^* + \Lambda_m^{-1} \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'} P_{h'}^{m'} [V_{h+1}^* - V_{h+1}^m] + \epsilon_h^m \\ &\quad \text{(Define } \epsilon_h^m = \Lambda_m^{-1} \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'} [P_{h'}^{m'} V_{h+1}^m - V_{h+1}^m(s_{h'+1}^{m'})]) \\ &= \lambda \Lambda_m^{-1} w_h^* + \Lambda_m^{-1} \sum_{m'=1}^{m-1} \sum_{h'=1}^H \phi_{h'}^{m'} \phi_{h'}^{m'\top} \int (V_{h+1}^*(s') - V_{h+1}^m(s')) d\mu(s') + \epsilon_h^m \\ &= \lambda \Lambda_m^{-1} w_h^* + \int (V_{h+1}^*(s') - V_{h+1}^m(s')) d\mu(s') - \lambda \Lambda_m^{-1} \int (V_{h+1}^*(s') - V_{h+1}^m(s')) d\mu(s') + \epsilon_h^m. \end{aligned}$$

Therefore,

$$\begin{aligned} Q_h^*(s, a) - \widehat{Q}_h^m(s, a) &= \phi(s, a)^\top (w_h^* - w_h^m) + \beta_m \|\phi(s, a)\|_{\Lambda_m^{-1}} \\ &\leq \underbrace{\lambda \phi(s, a)^\top \Lambda_m^{-1} w_h^*}_{\xi_1} + \underbrace{P_{s,a}(V_{h+1}^* - V_{h+1}^m) - \lambda \phi(s, a)^\top \Lambda_m^{-1} \int (V_{h+1}^*(s') - V_{h+1}^m(s')) d\mu(s')}_{\xi_2} \\ &\quad + \underbrace{\phi(s, a)^\top \epsilon_h^m}_{\xi_3} + \beta_m \|\phi(s, a)\|_{\Lambda_m^{-1}}. \end{aligned}$$

For ξ_1 , note that $\|w_h^*\|_2 = \|\theta^* + \int V_{h+1}^*(s') d\mu(s')\|_2 \leq (1 + 3B_*)\sqrt{d}$ by $V_{h+1}^*(s) \leq V^*(s) + 2B_* \leq 3B_*$ for any $s \in \mathcal{S}$, $h \in [H]$. Therefore, by the Cauchy-Schwarz inequality,

$$|\xi_1| \leq \|\phi(s, a)\|_{\Lambda_m^{-1}} \|\lambda w_h^*\|_{\Lambda_m^{-1}} \leq \|\phi(s, a)\|_{\Lambda_m^{-1}} \sqrt{\lambda} \|w_h^*\|_2 \leq \frac{\beta_m}{4} \|\phi(s, a)\|_{\Lambda_m^{-1}},$$

where the second inequality is by $\lambda_{\max}(\Lambda_m^{-1}) \leq \frac{1}{\lambda}$. Similarly, for ξ_2 ,

$$\begin{aligned} |\xi_2| &\leq \|\phi(s, a)\|_{\Lambda_m^{-1}} \left\| \lambda \int (V_{h+1}^*(s') - V_{h+1}^m(s')) d\mu(s') \right\|_{\Lambda_m^{-1}} && \text{(Cauchy-Schwarz inequality)} \\ &\leq \sqrt{\lambda} \|\phi(s, a)\|_{\Lambda_m^{-1}} \left\| \int (V_{h+1}^*(s') - V_{h+1}^m(s')) d\mu(s') \right\|_2 && (\lambda_{\max}(\Lambda_m^{-1}) \leq \frac{1}{\lambda}) \\ &\leq 3B_* \sqrt{\lambda d} \|\phi(s, a)\|_{\Lambda_m^{-1}} \leq \frac{\beta_m}{4} \|\phi(s, a)\|_{\Lambda_m^{-1}}. && (V_{h+1}^*(s) - V_{h+1}^m(s) \in [0, 3B_*] \text{ for any } s \in \mathcal{S}) \end{aligned}$$

For ξ_3 , by Eq. (9), $\|\epsilon_h^m\|_{\Lambda_m} \leq \frac{\beta_m}{2}$ with probability at least $1 - \delta$. Thus, $|\xi_3| \leq \|\phi(s, a)\|_{\Lambda_m^{-1}} \|\epsilon_h^m\|_{\Lambda_m} \leq \frac{\beta_m}{2} \|\phi(s, a)\|_{\Lambda_m^{-1}}$. To conclude, we have for all $m, h, (s, a)$:

$$0 \leq Q_h^*(s, a) - \widehat{Q}_h^m(s, a) \leq \widetilde{P}_{s,a}(V_{h+1}^* - V_{h+1}^m) + 2\beta_m \|\phi(s, a)\|_{\Lambda_m^{-1}}.$$

This completes the proof. \square

Lemma 14. *With probability at least $1 - \delta$, $\sum_{m=1}^{M'} V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m) \leq 2 \sum_{m=1}^{M'} \sum_{h=1}^H \text{gap}_h(s_h^m, a_h^m) + \mathcal{O}(B_* H \ln(M'/\delta))$ for any given $M' \in \mathbb{N}_+$.*

Proof. By the extended value difference lemma (Shani et al., 2020, Lemma 1):

$$\begin{aligned} V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m) &= \mathbb{E} \left[\sum_{h=1}^H \sum_a (\pi^m(a|s_h^m) - \widetilde{\pi}^*(a|s_h^m)) Q_h^*(s_h^m, a) \middle| \pi^m \right] \\ &= \mathbb{E} \left[\sum_{h=1}^H Q_h^*(s_h^m, a_h^m) - V_h^*(s_h^m) \middle| \pi^m \right] = \mathbb{E} \left[\sum_{h=1}^H \text{gap}_h(s_h^m, a_h^m) \middle| \pi^m \right], \end{aligned}$$

where $\widetilde{\pi}^*$ is the optimal policy of $\widetilde{\mathcal{M}}$. Therefore, by Lemma 39 and $\text{gap}_h(s, a) = \mathcal{O}(B_*)$, with probability at least $1 - \delta$,

$$\sum_{m=1}^{M'} V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m) \leq 2 \sum_{m=1}^{M'} \sum_{h=1}^H \text{gap}_h(s_h^m, a_h^m) + \mathcal{O}\left(B_* H \ln \frac{M'}{\delta}\right).$$

This completes the proof. \square

The next lemma provides an upper bound on the sum of gap functions satisfying some constraints. We denote by \mathcal{F}_h^m the interaction history up to (s_h^m, a_h^m) in $\widetilde{\mathcal{M}}$.

Lemma 15. *Suppose $B = 3B_*$, $\{z_h^m\}_{m=1}^{M'}$ are indicator functions such that $z_h^m \in \mathcal{F}_h^m$ for some $M' \in \mathbb{N}_+$, $h \in [H]$, and define $M_z = \sum_{m=1}^{M'} z_h^m$. Then with probability at least $1 - \delta$, Algorithm 2 ensures*

$$\sum_{m=1}^{M'} z_h^m \sum_{h'=h}^H \text{gap}_{h'}^m = \mathcal{O}\left(\sqrt{d^3 B_*^2 H M_z} \ln \frac{dB_* M' H}{\delta} + d^2 B_* H \ln^{1.5} \frac{dB_* M' H}{\delta}\right).$$

Proof. Denote by m_i the i -th interval among $[M']$ such that $z_h^{m_i} = 1$. Then,

$$\begin{aligned} &\sum_{i=1}^{M_z} \sum_{h'=h}^H Q_{h'}^*(s_{h'}^{m_i}, a_{h'}^{m_i}) - V_{h'}^*(s_{h'}^{m_i}) + \sum_{i=1}^{M_z} \sum_{h'=h}^H V_{h'}^*(s_{h'}^{m_i}) - V_{h'}^{m_i}(s_{h'}^{m_i}) \\ &= \sum_{i=1}^{M_z} \sum_{h'=h}^H Q_{h'}^*(s_{h'}^{m_i}, a_{h'}^{m_i}) - Q_{h'}^{m_i}(s_{h'}^{m_i}, a_{h'}^{m_i}) \quad (Q_{h'}^{m_i}(s_{h'}^{m_i}, a_{h'}^{m_i}) = V_{h'}^{m_i}(s_{h'}^{m_i}) \text{ by Lemma 5 and } B = 3B_*) \\ &\leq \sum_{i=1}^{M_z} \sum_{h'=h}^H P_{h'}^{m_i}(V_{h'+1}^* - V_{h'+1}^{m_i}) + 2 \sum_{i=1}^{M_z} \sum_{h'=h}^H \beta_{m_i} \|\phi_{h'}^{m_i}\|_{\Lambda_{m_i}^{-1}} && \text{(Lemma 13)} \\ &= \sum_{i=1}^{M_z} \sum_{h'=h}^H (V_{h'+1}^*(s_{h'+1}^{m_i}) - V_{h'+1}^{m_i}(s_{h'+1}^{m_i})) + \sum_{i=1}^{M_z} \sum_{h=h'}^H \left(\epsilon_{h'}^{m_i} + 2\beta_{m_i} \|\phi_{h'}^{m_i}\|_{\Lambda_{m_i}^{-1}} \right), \end{aligned}$$

where $\epsilon_h^{m_i} = P_h^{m_i}(V_{h+1}^* - V_{h+1}^{m_i}) - (V_{h+1}^*(s_{h+1}^{m_i}) - V_{h+1}^{m_i}(s_{h+1}^{m_i}))$. Reorganizing terms, and by $V_{H+1}^* = V_{H+1}^m = c_f$, $V_{h+1}^m(s) \leq V_{h+1}^*(s)$ (Lemma 5), we get:

$$\begin{aligned} \sum_{i=1}^{M_z} \sum_{h'=h}^H \text{gap}_{h'}^{m_i} &= \sum_{i=1}^{M_z} \sum_{h'=h}^H Q_{h'}^*(s_{h'}^{m_i}, a_{h'}^{m_i}) - V_{h'}^*(s_{h'}^{m_i}) \leq \sum_{i=1}^{M_z} \sum_{h=h'}^H \left(\epsilon_h^{m_i} + 2\beta_{m_i} \|\phi_h^{m_i}\|_{\Lambda_{m_i}^{-1}} \right) \\ &= \sum_{m=1}^{M'} \sum_{h'=h}^H z_h^m \epsilon_{h'}^m + 2 \sum_{i=1}^{M_z} \sum_{h=h'}^H \beta_{m_i} \|\phi_h^{m_i}\|_{\Lambda_{m_i}^{-1}}. \end{aligned}$$

For the first term, by $z_h^m \epsilon_{h'}^m \in \mathcal{F}_{h'+1}^m$ for $h' \geq h$ and Lemma 38, with probability at least $1 - \delta$,

$$\begin{aligned} \sum_{m=1}^{M'} \sum_{h'=h}^H z_h^m \epsilon_{h'}^m &\leq 3 \sqrt{\sum_{m=1}^{M'} \sum_{h'=h}^H \mathbb{E}[(z_h^m \epsilon_{h'}^m)^2 | \mathcal{F}_{h'}^m]} + \mathcal{O}\left(B_* \ln \frac{B_* M' H}{\delta}\right) \\ &= \mathcal{O}\left(B_* \sqrt{H M_z \ln \frac{B_* M' H}{\delta}} + B_* \ln \frac{B_* M' H}{\delta}\right). \quad (z_h^m \in \mathcal{F}_{h'}^m \text{ and } |\epsilon_{h'}^m| = \mathcal{O}(B_*)) \end{aligned}$$

For the second term, by Lemma 18,

$$\sum_{i=1}^{M_z} \sum_{h=h'}^H \beta_{m_i} \|\phi_h^{m_i}\|_{\Lambda_{m_i}^{-1}} = \mathcal{O}\left(\sqrt{d^3 B_*^2 H M_z} \ln \frac{d B_* M' H}{\delta} + d^2 B_* H \ln^{1.5} \frac{d B_* M' H}{\delta}\right).$$

Plugging these back completes the proof. \square

We are now ready to prove a bound on $\sum_m V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m)$, which is the key to proving Theorem 4.

Lemma 16. For any $M' \geq 3$, Algorithm 2 with $B = 3B_*$ and $H \geq \lceil \frac{35B_*}{c_{\min}} \ln(8B_* M' H) \rceil$ for some horizon H ensures with probability at least $1 - 3\delta - 1/4B_* M' H$, $\sum_{m=1}^{M'} V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m) = \mathcal{O}\left(\frac{d^3 B_*^4}{c_{\min}^2 \text{gap}_{\min}} \ln^5(dB_* M' H/\delta)\right)$.

Proof. First note that $V_h^{\pi^*}(s) \leq 3B_*$ for any $s \in \mathcal{S}, h \in [H]$. Thus, the expected hitting time of π^* in $\widetilde{\mathcal{M}}$ is at most $\frac{3B_*}{c_{\min}}$ starting from any state and layer. Without loss of generality, we assume that H is an even integer. Note that $\widetilde{\mathcal{M}}$ can be treated as an SSP instance where the learner teleports to the goal state at the $(H+1)$ -th step. Thus by Lemma 17 and $H \geq \frac{35B_*}{c_{\min}} \ln(8B_* M' H)$, when $h \leq \frac{H}{2} + 1$, $P(s_{H+1} \neq g | s_h = s, \pi^*) \leq \frac{1}{4B_* M' H}$ for any state s , and for any $h \leq \frac{H}{2}$:

$$Q_h^*(s, a) - Q^*(s, a) \leq Q_h^{\pi^*}(s, a) - Q^*(s, a) = P_{s,a}(V_{h+1}^{\pi^*} - V^*) \leq 2B_* \max_s P(s_{H+1} \neq g | \pi^*, s_{h+1} = s) \leq \frac{1}{2M'H}.$$

It also implies $|\text{gap}_h(s, a) - \text{gap}(s, a)|$ for $h \leq \frac{H}{2}$, since:

$$|\text{gap}_h(s, a) - \text{gap}(s, a)| \leq |Q_h^*(s, a) - Q^*(s, a)| + |V_h^*(s) - V^*(s)| \leq \frac{1}{2M'H} + \max_a |Q_h^*(s, a) - Q^*(s, a)| \leq \frac{1}{M'H}.$$

Define $\text{gap}_h^m = \text{gap}_h(s_h^m, a_h^m)$ and a threshold $\eta = \frac{3}{M'H}$. By Lemma 14, it suffices to bound $\sum_{m=1}^{M'} \sum_{h=1}^H \text{gap}_h^m$. Note that

$$\begin{aligned} \sum_{m=1}^{M'} \sum_{h=1}^H \text{gap}_h^m &\leq \sum_{m=1}^{M'} \sum_{h=1}^H \text{gap}_h^m \mathbb{I}\{\text{gap}_h^m > \eta\} + \mathcal{O}\left(\sum_{m=1}^{M'} \sum_{h=1}^H \frac{B_*}{M'H}\right) \\ &\leq \sum_{m=1}^{M'} \sum_{h \leq H/2} \text{gap}_h^m \mathbb{I}\{\text{gap}_h^m > \eta\} + \sum_{m=1}^{M'} \sum_{h > H/2} \text{gap}_h^m + \mathcal{O}(B_*). \end{aligned}$$

For the first term, define $N = \lceil \log_2(\frac{3B_*+1}{\eta}) \rceil = \mathcal{O}(\ln(B_* M' H))$, and

$$n^* = \min \left\{ n \in [N] : \exists (s', a'), h' \leq \frac{H}{2} \text{ such that } \text{gap}_{h'}(s', a') \in (\eta 2^{n-1}, \eta 2^n) \right\}.$$

Then by the definition of n^* and $|\text{gap}(s, a) - \text{gap}_h(s, a)| \leq \frac{1}{M'H}$ for $h \leq \frac{H}{2}$, there exist $(s', a'), h' \leq \frac{H}{2}$ such that

$$\text{gap}_{\min} \leq \text{gap}(s', a') \leq \text{gap}_{h'}(s', a') + \frac{1}{M'H} \leq \eta^{2n^*} + \frac{1}{M'H} \leq \frac{4}{3} \cdot \eta^{2n^*}. \quad (10)$$

Moreover, for each $n \in \mathbb{N}$ and $h \leq \frac{H}{2}$, define $z_h^m = \mathbb{I}\{\text{gap}_h^m > \eta^{2n}\}$. Then by Lemma 15, with probability at least $1 - \frac{\delta}{2(n+1)^2}$,

$$\eta^{2n} M_z \leq \sum_{m=1}^{M'} z_h^m \text{gap}_h^m = \mathcal{O} \left(\sqrt{d^3 B_*^2 H M_z} \ln \frac{dB_* M' H(n+1)}{\delta} + d^2 B_* H \ln^{1.5} \frac{dB_* M' H(n+1)}{\delta} \right),$$

where $M_z = \sum_{m=1}^{M'} z_h^m$. Solving a quadratic inequality w.r.t $\sqrt{M_z}$ gives:

$$\sum_{m=1}^{M'} \mathbb{I}\{\text{gap}_h^m > \eta^{2n}\} = \mathcal{O} \left(\frac{d^3 B_*^2 H}{\eta^2 4^n} \ln^2 \frac{dB_* M' H(n+1)}{\delta} + \frac{d^2 B_* H}{\eta^{2n}} \ln^{1.5} \frac{dB_* M' H(n+1)}{\delta} \right). \quad (11)$$

By a union bound, Eq. (11) holds for all $n \in \mathbb{N}$ simultaneously with probability at least $1 - \delta$. Therefore, the first term is bounded as follows:

$$\begin{aligned} & \sum_{m=1}^{M'} \sum_{h \leq H/2} \text{gap}_h^m \mathbb{I}\{\text{gap}_h^m > \eta\} \\ &= \sum_{m=1}^{M'} \sum_{h \leq H/2} \sum_{n=n^*}^N \text{gap}_h^m \mathbb{I}\{\text{gap}_h^m \in (\eta^{2n-1}, \eta^{2n}]\} \leq \sum_{m=1}^{M'} \sum_{h \leq H/2} \sum_{n=n^*}^N \eta^{2n} \mathbb{I}\{\text{gap}_h^m > \eta^{2n-1}\} \\ &= \mathcal{O} \left(\sum_{h \leq H/2} \sum_{n=n^*}^N \left(\frac{d^3 B_*^2 H}{\eta^{2n}} \ln^2 \frac{dB_* M' H(n+1)}{\delta} + d^2 B_* H \ln^{1.5} \frac{dB_* M' H(n+1)}{\delta} \right) \right) \quad (\text{Eq. (11)}) \\ &= \mathcal{O} \left(\frac{d^3 B_*^2 H^2}{\eta^{2n^*}} \ln^3 \frac{dB_* M' H}{\delta} + d^2 B_* H^2 \ln^{2.5} \frac{dB_* M' H}{\delta} \right) \quad (N = \mathcal{O}(\ln(B_* M' H))) \\ &= \mathcal{O} \left(\frac{d^3 B_*^2 H^2}{\text{gap}_{\min}} \ln^3 \frac{dB_* M' H}{\delta} + d^2 B_* H^2 \ln^{2.5} \frac{dB_* M' H}{\delta} \right). \quad (\text{Eq. (10)}) \end{aligned}$$

For the second term, note that:

$$\sum_{m=1}^{M'} \sum_{h > H/2} \text{gap}_h^m \leq \underbrace{\sum_{m=1}^{M'} \sum_{h > H/2} \text{gap}_h^m \mathbb{I}\left\{\exists h \leq \frac{H}{2} : \text{gap}_h^m > \eta\right\}}_{\xi_1} + \underbrace{\sum_{m=1}^{M'} \sum_{h > H/2} \text{gap}_h^m \mathbb{I}\left\{\forall h \leq \frac{H}{2} : \text{gap}_h^m \leq \eta\right\}}_{\xi_2}.$$

For ξ_1 , define $z_{\frac{H}{2}+1}^m = \mathbb{I}\{\exists h \leq \frac{H}{2} : \text{gap}_h^m > \eta\}$ and $M_z = \sum_{m=1}^{M'} z_{\frac{H}{2}+1}^m$. Then by Lemma 15, with probability at least $1 - \delta$,

$$\xi_1 = \sum_{m=1}^{M'} z_{\frac{H}{2}+1}^m \sum_{h > H/2} \text{gap}_h^m = \mathcal{O} \left(\sqrt{d^3 B_*^2 H M_z} \ln \frac{dB_* M' H}{\delta} + d^2 B_* H \ln^{1.5} \frac{dB_* M' H}{\delta} \right).$$

It suffices to bound M_z . Note that by the definition of n^* , we have $\min_{s, a, h \leq H/2, \text{gap}_h(s, a) > \eta} \text{gap}_h(s, a) \in (\eta^{2n^*-1}, \eta^{2n^*}]$. Thus, by Eq. (11),

$$\begin{aligned} M_z &= \sum_{m=1}^{M'} \mathbb{I}\left\{\exists h \leq \frac{H}{2} : \text{gap}_h^m > \eta\right\} \leq \sum_{m=1}^{M'} \sum_{h \leq H/2} \mathbb{I}\{\text{gap}_h^m > \eta\} \leq \sum_{m=1}^{M'} \sum_{h \leq H/2} \mathbb{I}\{\text{gap}_h^m > \eta^{2n^*-1}\} \\ &= \mathcal{O} \left(\frac{d^3 B_*^2 H^2}{\eta^2 4^{n^*-1}} \ln^2 \frac{dB_* M' H n^*}{\delta} + \frac{d^2 B_* H^2}{\eta^{2n^*-1}} \ln^{1.5} \frac{dB_* M' H n^*}{\delta} \right). \end{aligned}$$

Plugging this back and by Eq. (10), we get:

$$\xi_1 = \mathcal{O} \left(\frac{d^3 B_*^2 H^{1.5}}{\text{gap}_{\min}} \ln^2 \frac{dB_* M' H}{\delta} + d^2 B_* H \ln^{1.5} \frac{dB_* M' H}{\delta} \right).$$

For ξ_2 , denote by $\tilde{\pi}^m$ the near-optimal policy ‘‘closest’’ to π^m , such that:

$$\tilde{\pi}^m(s, h) = \begin{cases} \pi^m(s, h), & h \leq H/2 \text{ and } \text{gap}_h(s, \pi^m(s, h)) \leq \eta, \\ \pi^*(s, h), & h \leq H/2 \text{ and } \text{gap}_h(s, \pi^m(s, h)) > \eta, \\ \pi^*(s, h), & h > H/2. \end{cases}$$

Note that $\text{gap}_h(s, \tilde{\pi}^m(s, h)) \leq \eta$ for all s, h . By the extended value difference lemma (Shani et al., 2020, Lemma 1), $V_h^{\tilde{\pi}^m}(s) - V_h^*(s) = \mathbb{E}[\sum_{h'=h}^H \text{gap}_{h'}(s_{h'}, a_{h'}) | s_h = s, \tilde{\pi}^m] \leq \frac{3}{M'} \leq B_*$ for all s, h by $M' \geq 3$. Therefore, $V_h^{\tilde{\pi}^m}(s) \leq 4B_*$ for all s, h . Denote by \mathcal{F}_m the interaction history before interval m . Then, $\pi^m, \tilde{\pi}^m \in \mathcal{F}_m$, and

$$\begin{aligned} & P \left(\sum_{h>H/2} \text{gap}_h^m \mathbb{I} \{ \forall h \leq H/2 : \text{gap}_h^m \leq \eta \} = 0 \mid \pi^m, \mathcal{F}_m \right) \\ & \geq P \left(\exists h \leq H/2, \text{gap}_h^m > \eta \text{ or } \forall h \leq H/2, \text{gap}_h^m \leq \eta, s_{H/2+1} = g \mid \pi^m, \mathcal{F}_m \right) \\ & = P \left(\exists h \leq H/2, \tilde{\pi}^m(s_h^m, h) \neq \pi^m(s_h^m, h) \text{ or } \forall h \leq H/2, \tilde{\pi}^m(s_h^m, h) = \pi^m(s_h^m, h), s_{H/2+1} = g \mid \pi^m, \mathcal{F}_m \right) \\ & = P \left(\exists h \leq H/2, \tilde{\pi}^m(s_h^m, h) \neq \pi^m(s_h^m, h) \text{ or } \forall h \leq H/2, \tilde{\pi}^m(s_h^m, h) = \pi^m(s_h^m, h), s_{H/2+1} = g \mid \tilde{\pi}^m, \mathcal{F}_m \right) \\ & \geq P \left(s_{H/2+1} = g \mid \tilde{\pi}^m, \mathcal{F}_m \right) \geq 1 - \frac{1}{4B_* M' H}, \end{aligned}$$

where in the last inequality we apply Lemma 17, the fact that $V_h^{\tilde{\pi}^m}(s) \leq 4B_*$ for all s, h , and $H \geq \frac{35B_*}{c_{\min}} \ln(8B_* M' H)$. Now by Lemma 14 and $H = \lceil \frac{35B_*}{c_{\min}} \ln(8B_* M' H) \rceil$, we have:

$$\begin{aligned} \sum_{m=1}^{M'} V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m) & \leq 2 \sum_{m=1}^{M'} \sum_{h=1}^H \text{gap}_h(s_h^m, a_h^m) + \mathcal{O}(B_* H \ln(M'/\delta)) \\ & = \mathcal{O} \left(\frac{d^3 B_*^2 H^2}{\text{gap}_{\min}} \ln^3 \frac{dB_* M' H}{\delta} + d^2 B_* H^2 \ln^{2.5}(dB_* M' H/\delta) \right) \\ & = \mathcal{O} \left(\frac{d^3 B_*^4}{c_{\min}^2 \text{gap}_{\min}} \ln^5(dB_* M' H/\delta) \right). \end{aligned}$$

□

We are now ready to prove Theorem 4.

Proof of Theorem 4. First note that for a given $H \geq 4T_* \ln(4K)$, by Lemma 2 and Theorem 1, we have: $M = \tilde{\mathcal{O}}(K + d^3 H)$ with probability at least $1 - 4\delta$ for some $\delta > 0$ when running Algorithm 1 with Algorithm 2 and horizon H . That is, there exist $b > 0$ and constant $p \geq 1$ such that $M \leq b(K + d^3 H) \ln^p(dB_* H K/\delta)$. Now let $M' = b(K + d^3 H) \ln^p(dB_* H K/\delta)$. To obtain the regret bound in Lemma 16, it suffices to have $H \geq \frac{35B_*}{c_{\min}} \ln(8B_* M' H)$. Plugging in the definition of M' and by $x > \ln x$ for $x > 0$, it suffices to have $H = \frac{b' B_*}{c_{\min}} \ln(\frac{dB_* K}{\delta c_{\min}})$ for some constant $b' > 0$. To conclude, we have $M \leq M'$ with probability at least $1 - 4\delta$ when running Algorithm 1 with Algorithm 2 and horizon $H = \frac{b' B_*}{c_{\min}} \ln(\frac{dB_* K}{\delta c_{\min}})$. Moreover, with probability at least $1 - 3\delta - 1/4B_* M' H$, we have $\sum_{m=1}^{\min\{M, M'\}} V_1^{\pi^m}(s_1^m) - V_1^*(s_1^m) = \mathcal{O}(\frac{d^3 B_*^4}{c_{\min}^2 \text{gap}_{\min}} \ln^5(dB_* M' H/\delta))$. To obtain an expected regret bound, we further need to bound the cost under the low probability ‘‘bad’’ event. We make the following modification to Algorithm 1: whenever the counter $m = n \cdot M'$ for some $n \in \mathbb{N}_+$, we restart Algorithm 2. Ideas above are summarized in Algorithm 7. Now consider running Algorithm 7 with Algorithm 2, horizon $H = \frac{b' B_*}{c_{\min}} \ln(\frac{dB_* K}{\delta c_{\min}})$, failure probability $\delta = \frac{1}{4M' H}$, and restart threshold

Algorithm 7 Finite-Horizon Approximation of SSP from (Cohen et al., 2021)

Input: Algorithm \mathfrak{A} for finite-horizon MDP $\widetilde{\mathcal{M}}$ with horizon $H \geq 4T_\star \ln(4K)$ and restart threshold M' .

Initialize: interval counter $m \leftarrow 1$.

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

```

1   Set  $s_1^m \leftarrow s_{\text{init}}$ .
2   while  $s_1^m \neq g$  do
3       Feed initial state  $s_1^m$  to  $\mathfrak{A}$ .
4       for  $h = 1, \dots, H$  do
5           Receive action  $a_h^m$  from  $\mathfrak{A}$ .
6           if  $s_h^m \neq g$  then
7               Play action  $a_h^m$ , observe cost  $c_h^m = c(s_h^m, a_h^m)$  and next state  $s_{h+1}^m$ .
8           else Set  $c_h^m = 0$  and  $s_{h+1}^m = g$ .
9           Feed  $c_h^m$  and  $s_{h+1}^m$  to  $\mathfrak{A}$ .
10          Set  $s_1^{m+1} = s_{H+1}^m$  and  $m \leftarrow m + 1$ .
          if  $m = n \cdot M'$  for some  $n \in \mathbb{N}_+$  then Reinitialize  $\mathfrak{A}$ .
```

M' . By the choice of M' , we have $P(M > M') \leq 4\delta$. By a recursive argument, we have $P(M > n \cdot M') \leq (4\delta)^n$ for $n \in \mathbb{N}_+$. We have by Lemma 1 and Lemma 16:

$$\begin{aligned} \mathbb{E}[R_K] &\leq \mathbb{E}[\widetilde{R}_M] + B_\star \leq \mathbb{E}[\widetilde{R}_{\min\{M, M'\}}] + \mathbb{E}[\max\{0, M - M'\}(H + 2B_\star)] + B_\star \\ &= \mathcal{O}\left(\frac{d^3 B_\star^4}{c_{\min}^2 \text{gap}_{\min}} \ln^5(dB_\star M' H)\right) = \mathcal{O}\left(\frac{d^3 B_\star^4}{c_{\min}^2 \text{gap}_{\min}} \ln^5 \frac{dB_\star K}{c_{\min}}\right), \end{aligned}$$

where we apply

$$\begin{aligned} \mathbb{E}[\max\{0, M - M'\}(H + 2B_\star)] &\leq \sum_{n=1}^{\infty} P(M \in (nM', (n+1)M']) \cdot nM'(H + 2B_\star) \\ &\leq \sum_{n=1}^{\infty} n \cdot P(M > nM') M'(H + 2B_\star) \leq \sum_{n=1}^{\infty} n(4\delta)^n M'(H + 2B_\star) \leq \frac{16\delta M'(H + 2B_\star)}{1 - 4\delta} = \mathcal{O}(1). \end{aligned}$$

This completes the proof. □

B.9. Extra Lemmas for Section 3

Lemma 17. (Rosenberg & Mansour, 2020, Lemma 6) *Let π be a policy with expected hitting time at most τ starting from any state. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, π takes no more than $4\tau \ln \frac{2}{\delta}$ steps to reach the goal state.*

Lemma 18. *For an arbitrary set of intervals $\mathcal{I} \subseteq [M']$ for some $M' \in \mathbb{N}_+$, we have:*

$$\sum_{m \in \mathcal{I}} \sum_{h=1}^H \beta_m \|\phi_h^m\|_{\Lambda_m^{-1}} = \mathcal{O}\left(\sqrt{d^3 B^2 H |\mathcal{I}|} \ln \frac{dB M' H}{\delta} + d^2 B H \ln^{1.5} \frac{dB M' H}{\delta}\right).$$

Proof. We bound the sum by considering two cases:

$$\begin{aligned}
 \sum_{m \in \mathcal{I}} \sum_{h=1}^H \beta_m \|\phi_h^m\|_{\Lambda_m^{-1}} &\leq \beta_{M'} \sum_{m \in \mathcal{I}: \det(\Lambda_{m+1}) \leq 2 \det(\Lambda_m)} \sum_{h=1}^H \|\phi_h^m\|_{\Lambda_m^{-1}} + \beta_{M'} \sum_{m \in \mathcal{I}: \det(\Lambda_{m+1}) > 2 \det(\Lambda_m)} \sum_{h=1}^H \|\phi_h^m\|_{\Lambda_m^{-1}} \\
 &\leq \sqrt{2} \beta_{M'} \sum_{m \in \mathcal{I}} \sum_{h=1}^H \|\phi_h^m\|_{\Lambda_{m+1}^{-1}} + \mathcal{O}(\beta_{M'} d \ln(M'H/\lambda)H) \\
 &\quad (2\Lambda_m \succcurlyeq \Lambda_{m+1} \text{ by Lemma 30, and } \det(\Lambda_{M'})/\det(\Lambda_0) \leq ((\lambda + M'H)/\lambda)^d) \\
 &= \mathcal{O}\left(\beta_{M'} \sqrt{H|\mathcal{I}| \sum_{m \in \mathcal{I}} \sum_{h=1}^H \|\phi_h^m\|_{\Lambda_{m+1}^{-1}}^2} + \beta_{M'} dH \ln(M'H)\right) \quad (\text{Cauchy-Schwarz inequality}) \\
 &= \mathcal{O}\left(\sqrt{d^3 B^2 H |\mathcal{I}|} \ln \frac{dB M' H}{\delta} + d^2 B H \ln^{1.5} \frac{dB M' H}{\delta}\right). \\
 &\quad ((\text{Jin et al., 2020b, Lemma D.2}), \lambda = 1, \text{ and definition of } \beta_{M'})
 \end{aligned}$$

□

B.10. Proof of Theorem 5

Proof. Define $\delta = \frac{1}{3}$, $\Delta = \frac{\sqrt{\delta/K}}{8\sqrt{2}}$ and assume $K \geq \frac{d^2}{2\delta}$. Consider a family of SSP parameterized by $\rho \in \{-\Delta, \Delta\}^d$ with action set $\mathcal{A} = \{-1, 1\}^d$. For the SSP instance parameterized by ρ , it consists of two states $\{s_0, s_1\}$. The transition probabilities are as follows:

$$\begin{aligned}
 P(s_1|s_0, a) &= 1 - \delta - \langle \rho, a \rangle, & P(g|s_0, a) &= \delta + \langle \rho, a \rangle, \\
 P(s_1|s_1, a) &= 1 - 1/B_\star, & P(g|s_1, a) &= 1/B_\star,
 \end{aligned}$$

and the cost function is $c(s, a) = \mathbb{I}\{s = s_1\}$. The SSP instance above can be represented as a linear SSP of dimension $d + 2$ as follows: define $\alpha = \sqrt{\frac{1}{1+\Delta d}}$, $\beta = \sqrt{\frac{\Delta}{1+\Delta d}}$,

$$\phi(s, a) = \begin{cases} [\alpha, \beta a^\top, 0]^\top, & s = s_0 \\ [0, 0, 1]^\top, & s = s_1 \end{cases} \quad \mu(s') = \begin{cases} [(1-\delta)/\alpha, -\rho^\top/\beta, 1-1/B_\star]^\top, & s' = s_1 \\ [\delta/\alpha, \rho^\top/\beta, 1/B_\star]^\top, & s' = g \end{cases}$$

and $\theta^\star = [0, 0, 1]$. Note that it satisfies $c(s, a) = \phi(s, a)^\top \theta^\star$, $P(s'|s, a) = \phi(s, a)^\top \mu(s')$, $\|\phi(s, a)\|_2 \leq 1$, and $\|\theta^\star\|_2 \leq 1 \leq \sqrt{d+2}$. Moreover, for any function $h: \mathcal{S}_+ \rightarrow \mathbb{R}$, we have:

$$\sum_{s'} h(s') \mu(s') = \begin{bmatrix} h(s_1)(1-\delta)\sqrt{1+\Delta d} + h(g)\delta\sqrt{1+\Delta d} \\ (h(g) - h(s_1))\rho\sqrt{(1+\Delta d)/\Delta} \\ h(s_1)(1-1/B_\star) + h(g)/B_\star \end{bmatrix}.$$

Note that when $K \geq \frac{d^2}{2\delta}$, $\Delta d \leq \frac{\delta}{8} = \frac{1}{24}$, and

$$\begin{aligned}
 (h(s_1)(1-\delta)\sqrt{1+\Delta d} + h(g)\delta\sqrt{1+\Delta d})^2 &\leq \|h\|_\infty^2 (1+\Delta d) \leq \frac{25}{24} \|h\|_\infty^2, \\
 \left\| (h(g) - h(s_1))\rho\sqrt{(1+\Delta d)/\Delta} \right\|_2^2 &\leq 4 \|h\|_\infty^2 \Delta d (1+\Delta d) \leq \frac{25}{24} \|h\|_\infty^2, \\
 (h(s_1)(1-1/B_\star) + h(g)/B_\star)^2 &\leq \|h\|_\infty^2.
 \end{aligned}$$

Thus, we have $\|\sum_{s'} h(s') \mu(s')\|_2 \leq \|h\|_\infty \sqrt{d+2}$ by $d \geq 2$, and the SSP instance satisfies [Assumption 1](#). The regret is bounded as follows: let a_k denote the first action taken by the learner in episode k . Then for any $\rho \in \{-\Delta, \Delta\}^d$, the

expected cost of taking action a as the first action is $C_\rho(a) = B_\star(1 - \delta - \langle \rho, a \rangle)$.

$$\begin{aligned} \mathbb{E}_\rho[R_K] &= \sum_{k=1}^K \mathbb{E}_\rho \left[C_\rho(a_k) - \min_a C_\rho(a) \right] = B_\star \sum_{k=1}^K \mathbb{E}_\rho \left[\max_a \langle \rho, a \rangle - \langle \rho, a_k \rangle \right] \\ &= 2B_\star \Delta \sum_{k=1}^K \mathbb{E}_\rho \left[\sum_{j=1}^d \mathbb{I}\{\text{sgn}(\rho_j) \neq \text{sgn}(a_{k,j})\} \right] = 2B_\star \Delta \sum_{j=1}^d \mathbb{E}_\rho[N_j(\rho)], \end{aligned}$$

where we define $N_j(\rho) = \sum_{k=1}^K \mathbb{I}\{\text{sgn}(\rho_j) \neq \text{sgn}(a_{k,j})\}$, and \mathbb{E}_ρ is the expectation w.r.t the SSP instance parameterized by ρ . Let ρ^j denote the vector that differs from ρ at its j -th coordinate only. Then, we have $N_j(\rho^j) + N_j(\rho) = K$, and for a fixed j ,

$$\begin{aligned} 2 \sum_{\rho} \mathbb{E}_\rho[R_K] &= \sum_{\rho} (\mathbb{E}_\rho[R_K] + \mathbb{E}_{\rho^j}[R_K]) = 2B_\star \Delta \sum_{\rho} \sum_{j=1}^d (K + \mathbb{E}_\rho[N_j(\rho)] - \mathbb{E}_{\rho^j}[N_j(\rho)]) \\ &\geq 2B_\star \Delta \sum_{\rho} \sum_{j=1}^d \left(K - K \sqrt{2\text{KL}(P_\rho, P_{\rho^j})} \right), \end{aligned}$$

where P_ρ is the joint probability of K trajectories induced by the interactions between the learner and the SSP parameterized by ρ , and in the last inequality we apply Pinsker's inequality to obtain:

$$|\mathbb{E}_\rho[N_j(\rho)] - \mathbb{E}_{\rho^j}[N_j(\rho)]| \leq K \|P_\rho - P_{\rho^j}\|_1 \leq K \sqrt{2\text{KL}(P_\rho, P_{\rho^j})}.$$

By the divergence decomposition lemma (see e.g. (Lattimore & Szepesvári, 2020, Lemma 15.1)), we further have

$$\begin{aligned} \text{KL}(P_\rho, P_{\rho^j}) &= \sum_{k=1}^K \mathbb{E}_\rho \left[\text{KL}(\text{Bernoulli}(\delta + \langle a_k, \rho \rangle), \text{Bernoulli}(\delta + \langle a_k, \rho^j \rangle)) \right] \\ &= \sum_{k=1}^K \mathbb{E}_\rho \left[\frac{2 \langle a_k, \rho - \rho^j \rangle^2}{\delta + \langle a_k, \rho \rangle} \right] \leq \frac{16K\Delta^2}{\delta}, \end{aligned} \quad (d\Delta \leq \delta/2)$$

where in the second last inequality we apply $\text{KL}(\text{Bernoulli}(a), \text{Bernoulli}(b)) \leq 2(a-b)^2/a$ when $a \leq 1/2, a+b \leq 1$, which is true when $\delta \leq 1/3, d\Delta \leq \delta/2$. Substituting these back, we get:

$$2 \sum_{\rho} \mathbb{E}_\rho[R_K] \geq 2B_\star \Delta \sum_{\rho} \sum_{j=1}^d \left(K - K \sqrt{32K\Delta^2/\delta} \right) = \Omega \left(\sum_{\rho} B_\star d \sqrt{\delta K} \right). \quad (12)$$

Now note that $\text{gap}(s_1, a) = 0$ for all a . Define $a^\star = \text{argmax}_a \langle \rho, a \rangle$. Then for any $a \neq a^\star$,

$$Q^\star(s_0, a) - V^\star(s_0) = (1 - \delta - \langle \rho, a \rangle)B_\star - (1 - \delta - \langle \rho, a^\star \rangle)B_\star = B_\star \langle \rho, a^\star - a \rangle \geq 2B_\star \Delta.$$

Thus, $\text{gap}_{\min} = 2B_\star \Delta$. By $\sqrt{K} = \frac{\sqrt{\delta}}{8\sqrt{2}\Delta}$ and Eq. (12), we get:

$$\sum_{\rho} \mathbb{E}_\rho[R_K] = \Omega \left(\sum_{\rho} B_\star d \sqrt{\delta K} \right) = \Omega \left(\sum_{\rho} \frac{dB_\star \delta}{\Delta} \right) = \Omega \left(\sum_{\rho} \frac{dB_\star^2}{\text{gap}_{\min}} \right).$$

Selecting ρ^\star which maximizes $\mathbb{E}_\rho[R_K]$, we get: $\mathbb{E}_{\rho^\star}[R_K] = \Omega \left(\frac{dB_\star^2}{\text{gap}_{\min}} \right)$. \square

C. Omitted Details for Section 4

Notations Define $Q_t(s, a) = \phi(s, a)^\top w_t$ such that $a_t = \text{argmin}_a Q_t(s_t, a)$, and operator $U_B : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $U_B w = \theta^\star + \int V_{w,B}(s') d\mu(s')$. Define $\iota_t = \iota_{B_t, t}, \mathcal{J}_t = \mathcal{J}_{B_t}, P_t = P_{s_t, a_t}, C_t = \sum_{i=1}^t c(s_i, a_i)$, and $\mathcal{J} = \mathcal{J}_{2B_\star}$. By Lemma 4, $\mathcal{J}_t \subseteq \mathcal{J}$ for any $t \in [T]$.

For notational convenience, we divide the whole learning process into epochs indexed by l , and a new epoch begins whenever w_t is recomputed. Denote by $t_l + 1$ the first time step in epoch l , and for a quantity, function or set f_t indexed by time step t , we define $f_l = f_{t_l+1}$. Denote by l_t the epoch time step t belongs to, and we often ignore the subscript t when there is no confusion. Clearly, $V_t = V_l$, and similarly for $w_l, \tilde{w}_l, \iota_l, \Omega_l$ (ignoring the dependency on t for l). With this notation setup, we define L' as the number of epochs that starts by the overestimate condition, that is, $L' = |\{l > 1 : V_{l-1}(s'_{t_l}) = 2B_{l-1}\}|$. Also define $\nu_t = \operatorname{argmax}_{\nu = \tilde{w}_l - w, w \in \Omega_l} |\phi_t^\top \nu|$ and a special covariance matrix $W_{j,t}(\nu) = 2^j I + \sum_{i < t} \min\{1, 2^j / |\phi_i^\top \nu|\} \phi_i \phi_i^\top$. Note that $\Phi_t^j(\nu) = \|\nu\|_{W_{j,t}(\nu)}^2$.

Assumption For simplicity, we assume that $\{\phi(s, a)\}_{(s,a) \in \mathcal{S} \times \mathcal{A}}$ spans \mathbb{R}^d . It implies that if $\phi(s, a)^\top v = \phi(s, a)^\top w$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, then $v = w$.

Truncating the Interaction for Technical Issue An important question in SSP is whether the algorithm halts in finite number of steps. To overcome some technical issues, we first assume that the algorithm halts after T' steps for an arbitrary $T' \in \mathbb{N}_+$, even if the goal state is not reached. Specifically, we redefine the notation T to be the minimum between the number of steps taken by the learner in K episodes and T' , that is, $T = T'$ if the learner does not finish K episodes in T' steps. We also redefine R_K under the new definition of T , and the true regret now becomes $\lim_{T' \rightarrow \infty} R_K$. The implication under truncation is that s'_T may not be g , and $T \leq T'$. In [Appendix C.3](#), we prove a regret bound on R_K independent of T' . Thus, the proven regret bound is also an upper bound of the true regret, as it is a valid upper bound of $\lim_{T' \rightarrow \infty} R_K$.

Below we provide detailed proofs of lemmas and the main theorem.

C.1. Proof of [Lemma 3](#)

We will prove a more general statement, from which [Lemma 3](#) is a directly corollary.

Lemma 19. *With probability at least $1 - \delta$, for any $t \in \mathbb{N}_+$, $B \in \{2^i\}_{i \in \mathbb{N}}$, and $w \in \mathbb{B}(3\sqrt{dB})$, we have $U_B w \in \Omega_t(w, B)$.*

Proof. For each $t \in \mathbb{N}_+$, $B \in \{2^i\}_{i \in \mathbb{N}}$, $w \in \mathcal{G}_{\epsilon/t}(3\sqrt{dB})$, $j \in \mathcal{J}_B$, $\nu \in \mathcal{G}_{\epsilon/t}(6\sqrt{dB})$, by [Lemma 36](#), we have with probability at least $1 - 6\delta' \log_2 t$ with $\delta' = \delta / (24t^2 \log_2^2(2B) \log_2(t) |\mathcal{J}_B| (12\sqrt{dB}t/\epsilon)^{2d})$:

$$\begin{aligned} & \left| \sum_{i < t} \operatorname{clip}_j(\phi_i^\top \nu) \epsilon_{V_{w,B}^i}(U_B w) \right| = \left| \sum_{i < t} \operatorname{clip}_j(\phi_i^\top \nu) (P_i V_{w,B} - V_{w,B}(s'_i)) \right| \\ & \leq 8 \sqrt{\sum_{i < t} \operatorname{clip}_j^2(\phi_i^\top \nu) \eta_{V_{w,B}^i}^i(U_B w) \ln \frac{1}{\delta'}} + 32B2^j \ln \frac{1}{\delta'} \leq \sqrt{\sum_{i < t} \operatorname{clip}_j^2(\phi_i^\top \nu) \eta_{V_{w,B}^i}^i(U_B w) \frac{\iota_{B,t}}{3}} + \frac{B}{2} 2^j \iota_{B,t}. \end{aligned} \quad (13)$$

Taking a union bound, [Eq. \(13\)](#) holds for any $t, B \in \{2^i\}_{i \in \mathbb{N}}$, $w \in \mathcal{G}_{\epsilon/t}(3\sqrt{dB})$, $j \in \mathcal{J}_B$, $\nu \in \mathcal{G}_{\epsilon/t}(6\sqrt{dB})$ with probability at least $1 - \delta$.

Now for any $t \in \mathbb{N}_+$, $B \in \{2^i\}_{i \in \mathbb{N}}$, $w \in \mathbb{B}(3\sqrt{dB})$, there exist $w' \in \mathcal{G}_{\epsilon/t}(3\sqrt{dB})$ such that $\|w - w'\|_\infty \leq \frac{\epsilon}{t}$. Also define $V = V_{w,B}$, $V' = V_{w',B}$, $\tilde{w} = U_B w$, and $\tilde{w}' = U_B w'$. Note that

$$\|V - V'\|_\infty \leq \max_{s,a} |\phi(s, a)^\top (w - w')| \leq \sqrt{d} \|w - w'\|_\infty \leq \frac{\sqrt{d}\epsilon}{t}, \quad (14)$$

$$\|\tilde{w} - \tilde{w}'\|_2 = \left\| \int (V(s') - V'(s')) d\mu(s') \right\|_2 \leq \sqrt{d} \|V - V'\|_\infty \leq \frac{d\epsilon}{t}. \quad (15)$$

Thus, we have for any $j \in \mathcal{J}_B, \nu \in \mathcal{G}_{\epsilon/t}(6\sqrt{dB})$:

$$\begin{aligned}
 & \left| \sum_{i < t} \text{clip}_j(\phi_i^\top \nu) \epsilon_V^i(\tilde{w}) \right| = \left| \sum_{i < t} \text{clip}_j(\phi_i^\top \nu) (\phi_i^\top \tilde{w} - c_i - V(s'_i)) \right| \\
 & \leq \left| \sum_{i < t} \text{clip}_j(\phi_i^\top \nu) (\phi_i^\top \tilde{w}' - c_i - V'(s'_i)) \right| + \left| \sum_{i < t} \text{clip}_j(\phi_i^\top \nu) \phi_i^\top (\tilde{w} - \tilde{w}') \right| + \left| \sum_{i < t} \text{clip}_j(\phi_i^\top \nu) (V(s'_i) - V'(s'_i)) \right| \\
 & \leq \sqrt{\sum_{i < t} \text{clip}_j^2(\phi_i^\top \nu) (\phi_i^\top \tilde{w}' - c_i - V'(s'_i))^2} \frac{\ell_{B,t}}{3} + \frac{B}{2} 2^j \ell_{B,t} + 2^{j+1} d\epsilon \quad (\text{Eq. (13), Eq. (14), and Eq. (15)}) \\
 & \leq \sqrt{\sum_{i < t} \text{clip}_j^2(\phi_i^\top \nu) \eta_V^i(\tilde{w}) \ell_{B,t}} + \sqrt{\sum_{i < t} \text{clip}_j^2(\phi_i^\top \nu) (\phi_i^\top (\tilde{w}' - \tilde{w}))^2 \ell_{B,t}} + \sqrt{\sum_{i < t} \text{clip}_j^2(\phi_i^\top \nu) (V'(s'_i) - V(s'_i))^2 \ell_{B,t}} \\
 & \quad + \frac{B}{2} 2^j \ell_{B,t} + 2^{j+1} d\epsilon \quad ((a+b+c)^2 \leq 3(a^2+b^2+c^2) \text{ and } \sqrt{a+b} \leq \sqrt{a} + \sqrt{b}) \\
 & \leq \sqrt{\sum_{i < t} \text{clip}_j^2(\phi_i^\top \nu) \eta_V^i(\tilde{w}) \ell_{B,t}} + \frac{B}{2} 2^j \ell_{B,t} + 4 \cdot 2^j d\epsilon \ell_{B,t} \leq \sqrt{\sum_{i < t} \text{clip}_j^2(\phi_i^\top \nu) \eta_V^i(\tilde{w}) \ell_{B,t}} + B 2^j \ell_{B,t}. \\
 & \quad (\text{Eq. (14), Eq. (15), and } 8d\epsilon \leq 1)
 \end{aligned}$$

Moreover, $\tilde{w} \in \mathbb{B}(3\sqrt{dB})$ by $\|V_{w,B}\|_\infty \leq 2B$. Thus, $U_B w \in \Omega_t(w, B)$ for any $t \in \mathbb{N}_+, B \in \{2^i\}_{i \in \mathbb{N}}$, and $w \in \mathbb{B}(3\sqrt{dB})$, and the statement is proved. \square

Proof of Lemma 3. This directly follows from Lemma 19 by $w_t \in \mathbb{B}(3\sqrt{dB}_t)$, $V_t = V_{w_t, B_t}$, and $\tilde{w}_t = \theta^* + \int V_t(s') d\mu(s') = U_{B_t} w_t$. \square

C.2. Proof of Lemma 4

Lemma (restatement of Lemma 4). *With probability at least $1 - \delta$, $V_l(s_{t_l+1}) \leq V^*(s_{t_l+1})$ for any epoch l and $B_t \leq 2B_*$.*

Proof. For the first statement, note that any epoch l , by Lemma 20, there exists $w_l^\infty \in \mathbb{B}(3\sqrt{dB}_l)$ such that $w_l^\infty = U_{B_l} w_l^\infty$ and $V_{w_l^\infty, B_l}(s) \leq V^*(s)$. Therefore, $w_l^\infty \in \Omega_l(w_l^\infty, B_l)$, and $V_l(s_{t_l+1}) = V_{w_l, B_l}(s_{t_l+1}) \leq V_{w_l^\infty, B_l}(s_{t_l+1}) \leq V^*(s_{t_l+1})$ by the definition of w_l . The second statement is a direct corollary of the first statement and how B_t is updated. \square

Lemma 20. *For any $B > 0$, there exists $w \in \mathbb{B}(3\sqrt{dB})$ such that $w = U_B w$, and $V_{w,B}(s) \leq V^*(s)$.*

Proof. Define $w^1 = \mathbf{0} \in \mathbb{R}^d$, and $w^{n+1} = U_B w^n$. We prove by induction that $\phi(s, a)^\top (w^{n+1} - w^n) \geq 0$ and $\phi(s, a)^\top w^n \leq Q^*(s, a)$. The base case $n = 1$ is clearly true. Now for $n > 1$, assume that we have $\phi(s, a)^\top (w^n - w^{n-1}) \geq 0$ and $\phi(s, a)^\top w^{n-1} \leq Q^*(s, a)$. Then, $\phi(s, a)^\top (w^{n+1} - w^n) = P_{s,a}(V_{w^n, B} - V_{w^{n-1}, B}) \geq 0$ and $\phi(s, a)^\top w^n = c(s, a) + P_{s,a} V_{w^{n-1}, B} \leq c(s, a) + P_{s,a} V^* \leq Q^*(s, a)$. Therefore, the sequence $\{\phi(s, a)^\top w^n\}_{n=1}^\infty$ is non-decreasing and bounded, and thus converges. Since $\{\phi(s, a)\}_{(s,a) \in \mathcal{S} \times \mathcal{A}}$ spans \mathbb{R}^d , the limit $w^\infty = \lim_{n \rightarrow \infty} w^n$ exists and $w^\infty = U_B w^\infty$. Moreover, $w^\infty \in \mathbb{B}(3\sqrt{dB})$ by $\|V_{w^\infty, B}\|_\infty \leq 2B$ and $V_{w^\infty, B}(s) \leq V^*(s)$ since $\phi(s, a)^\top w^\infty = \lim_{n \rightarrow \infty} \phi(s, a)^\top w^n \leq Q^*(s, a)$. This completes the proof. \square

C.3. Proof of Theorem 6

Proof. We decompose the regret as follows:

$$R_K = \sum_{t=1}^T c_t - K \cdot V^*(s_{\text{init}}) = \sum_{l=1}^L \left(\sum_{t=t_l+1}^{t_{l+1}} c_t - V_l(s_{t_l+1}) \right) + \sum_{l=1}^L V_l(s_{t_l+1}) - K \cdot V^*(s_{\text{init}}).$$

For the first term, for a fixed epoch l , define $\chi_\tau = \sum_{t=\tau}^{t_{l+1}} c_t - V_l(s_\tau)$ for $\tau \in \{t_l + 1, \dots, t_{l+1}\}$ and $\chi_{t_{l+1}+1} = -V_l(s'_{t_{l+1}})$. Note that within epoch l , we have $V_l(s_\tau) = [Q_l(s_\tau, a_\tau)]_{[0, \infty)} \geq Q_l(s_\tau, a_\tau) = \phi_\tau^\top w_l$. Thus, for $\tau \in \{t_l + 1, \dots, t_{l+1}\}$,

$$\begin{aligned} \chi_\tau &= \sum_{t=\tau}^{t_{l+1}} c_t - V_l(s_\tau) \leq \sum_{t=\tau+1}^{t_{l+1}} c_t + c_\tau - \phi_\tau^\top w_l \\ &= \sum_{t=\tau+1}^{t_{l+1}} c_t - V_l(s'_\tau) + (V_l(s'_\tau) - P_\tau V_l) + \phi_\tau^\top (\tilde{w}_l - w_l) && (c_\tau + P_\tau V_l = \phi_\tau^\top \tilde{w}_l) \\ &= \chi_{\tau+1} + (V_l(s'_\tau) - P_\tau V_l) + \phi_\tau^\top (\tilde{w}_l - w_l) \\ &\leq \dots \leq -V_l(s'_{t_{l+1}}) + \sum_{t=\tau}^{t_{l+1}} (V_l(s'_t) - P_t V_l) + \sum_{t=\tau}^{t_{l+1}} \phi_t^\top (\tilde{w}_l - w_l). \end{aligned}$$

Therefore, we have:

$$\begin{aligned} R_K &= \sum_{l=1}^L \chi_{t_{l+1}} + \sum_{l=1}^L V_l(s_{t_{l+1}}) - K \cdot V^*(s_{\text{init}}) \\ &\leq \sum_{l=1}^L \sum_{t=t_{l+1}}^{t_{l+1}} [(V_l(s'_t) - P_t V_l) + \phi_t^\top (\tilde{w}_l - w_l)] + \sum_{l=1}^L (V_l(s_{t_{l+1}}) - V_l(s'_{t_{l+1}})) - K \cdot V^*(s_{\text{init}}). \end{aligned}$$

We first bound the switching costs, that is, the last two terms above. We consider three cases based on how an epoch starts: define $\mathcal{L}_1 = \{l : s'_{t_l} = g\}$, $\mathcal{L}_2 = \{l > 1 : \exists j \in \mathcal{J}_l, \nu \in \mathcal{G}_{\epsilon/(t_l+1)}(6\sqrt{d}B_{l-1}), \Phi_{t_l+1}^j(\nu) > 8d^2\Phi_{t_{l-1}+1}^j(\nu)\}$, and $\mathcal{L}_3 = \{l > 1 : V_{l-1}(s'_{t_l}) = 2B_{l-1}\}$. Then,

$$\begin{aligned} &\sum_{l=1}^L (V_l(s_{t_{l+1}}) - V_l(s'_{t_{l+1}})) - K \cdot V^*(s_{\text{init}}) \\ &= \underbrace{\sum_{l \in \mathcal{L}_1} V_l(s_{t_{l+1}}) - K \cdot V^*(s_{\text{init}})}_{\xi_1} + \underbrace{\sum_{l \in \mathcal{L}_2} V_l(s_{t_{l+1}})}_{\xi_2} + \underbrace{\sum_{l \in \mathcal{L}_3} V_l(s_{t_{l+1}}) - \sum_{l=1}^L V_l(s'_{t_{l+1}})}_{\xi_3}. \end{aligned}$$

Note that $\xi_1 \leq 0$ since for $l \in \mathcal{L}_1$, $V_l(s_{t_{l+1}}) = V_l(s_{\text{init}}) \leq V^*(s_{\text{init}})$ by Lemma 4. For ξ_2 , note that $|\mathcal{L}_2| = \tilde{\mathcal{O}}(d)$ by Lemma 27. Thus, $\xi_2 = \tilde{\mathcal{O}}(dB_*)$ by $\|V_l\|_\infty \leq 4B_*$ (Lemma 4). For ξ_3 , note that for each $l \in \mathcal{L}_3$, $V_l(s_{t_{l+1}}) - V_{l-1}(s'_{t_l}) \leq B_l - 2B_{l-1} \leq 2B_* \mathbb{I}\{B_l \neq B_{l-1}\} - 1$ by $V_l(s_{t_{l+1}}) \leq B_l \leq 2B_*$ and $B_l \geq 1$. Thus, $\xi_3 \leq \tilde{\mathcal{O}}(B_*) - L'$, by $|\mathcal{L}_3| = L'$ and $\sum_{l=1}^L \mathbb{I}\{B_l \neq B_{l-1}\} = \mathcal{O}(\log_2 B_*)$. Therefore, with probability at least $1 - 5\delta$,

$$\begin{aligned} R_K &\leq \sum_{l=1}^L \sum_{t=t_{l+1}}^{t_{l+1}} [(V_l(s'_t) - P_t V_l) + \phi_t^\top (\tilde{w}_l - w_l)] + \tilde{\mathcal{O}}(dB_* - L') \\ &= \tilde{\mathcal{O}} \left(\sqrt{\sum_{t=1}^T \mathbb{V}(P_t, V_t)} + \sum_{t=1}^T |\phi_t^\top \nu_t| + dB_* - L' \right) && (\text{Lemma 38, } w_l \in \Omega_l, \text{ and definition of } \nu_t) \\ &= \tilde{\mathcal{O}} \left(\sqrt{B_*^2 L' + B_* C_T + B_* \sum_{t=1}^T |\phi_t^\top \nu_t| + \sum_{t=1}^T |\phi_t^\top \nu_t| + dB_* - L'} \right) && (\text{Lemma 21}) \\ &= \tilde{\mathcal{O}} \left(\sqrt{B_* C_T} + \sum_{t=1}^T |\phi_t^\top \nu_t| + dB_* + B_*^2 \right) && (\sqrt{x+y} \leq \sqrt{x} + \sqrt{y} \text{ and } \sqrt{ab} \leq \frac{a+b}{2}) \\ &\leq \tilde{\mathcal{O}} \left(d^{3.5} \sqrt{B_* C_T} + d^{3.5} \sqrt{B_* \epsilon T} + d^5 B_*^2 \right) + 65d^{2.5} \epsilon T && (\text{Lemma 22}) \\ &\leq \tilde{\mathcal{O}} \left(d^{3.5} \sqrt{B_* C_T} + d^5 B_*^2 \right) + \frac{C_T}{2K}. && (\text{definition of } \epsilon \text{ and } c_{\min} T \leq C_T) \end{aligned}$$

By $R_K = C_T - K \cdot V^*(s_{\text{init}})$ and Lemma 28 with $x = C_T$ (we also bound T by C_T/c_{\min} in logarithmic terms), we get $C_T = \tilde{\mathcal{O}}(B_*K + d^7B_* + d^5B_*^2)$. Plugging this back, we obtain

$$R_K = \tilde{\mathcal{O}}\left(d^{3.5}B_*\sqrt{K} + d^7B_*^2\right).$$

This completes the proof. \square

C.4. Intuition for Overestimate Condition

Now we provide more reasonings on including the overestimate condition. Similar to (Zanette et al., 2020b; Wei et al., 2021b), we incorporate global optimism at the starting state of each epoch via solving an optimization problem. This is different from many previous work (Jin et al., 2020b; Vial et al., 2021) that adds bonus terms to ensure local optimism over all states. The advantage of global optimism is that it avoids using a larger function class of Q_t, V_t for the bonus terms, which reduces the order of d in the regret bound. However, this improvement also requires $\|V_t\|_\infty$ is of order B_* . In (Zanette et al., 2020b), they directly enforcing this constraint, which is not practical under large state space as we may need to iterate over all state-action pairs to check this constraint.

Here we take a new approach: we first enforce a bound on $\|V_t\|_\infty$ by direct truncation. However, the upper bound truncation on V_t may break the analysis. To resolve this, we start a new epoch whenever V_t is overestimated by a large amount. By the objective of the optimization problem, $V_t(s_t)$ will not be overestimated in the new epoch. Hence, the upper bound truncation will not be triggered. Moreover, the overestimate of V_t cancels out the switching cost in this case as in previous discussion.

The disadvantage of the overestimation condition is that we may update policy at every time step in the worst case. If we remove this condition, $\|V_t\|_\infty = \tilde{\mathcal{O}}(\sqrt{dB_*})$ by the norm constraint on w_t , which brings back an extra \sqrt{d} factor. However, we only recompute policy for $\mathcal{O}(K + d \ln T)$ times in this case.

C.5. Extra Lemmas for Section 4

Lemma 21. *With probability at least $1 - \delta$, $\sum_{t=1}^T \mathbb{V}(P_t, V_t) = \tilde{\mathcal{O}}\left(dB_*^2 + B_*^2L' + B_*C_T + B_*\sum_{t=1}^T |\phi_t^\top \nu_t|\right)$.*

Proof. Note that when $V_l(s_t) = 0$, $V_l(s_t) - P_tV_l \leq 0$. Otherwise, $Q_l(s_t, a) > 0$ for any a and $V_l(s_t) \leq Q_l(s_t, a_t)$. Thus, $V_l(s_t)^2 - (P_tV_l)^2 = (V_l(s_t) + P_tV_l)(V_l(s_t) - P_tV_l) \leq (V_l(s_t) + P_tV_l)|Q_l(s_t, a_t) - P_tV_l|$. Then with probability at least $1 - \delta$,

$$\begin{aligned} \sum_{t=1}^T \mathbb{V}(P_t, V_t) &= \sum_{t=1}^T (P_tV_t^2 - V_t^2(s'_t)) + \sum_{t=1}^T (V_t^2(s'_t) - V_t^2(s_t)) + \sum_{t=1}^T (V_t^2(s_t) - (P_tV_t)^2) \\ &\stackrel{(i)}{=} \tilde{\mathcal{O}}\left(\sqrt{\sum_{t=1}^T \mathbb{V}(P_t, V_t^2) + dB_*^2 + B_*^2L'} + \sum_{t=1}^T (V_l(s_t) + P_tV_l)|c_t + \phi_t^\top(w_l - \tilde{w}_l)|\right) \\ &\stackrel{(ii)}{=} \tilde{\mathcal{O}}\left(B_*\sqrt{\sum_{t=1}^T \mathbb{V}(P_t, V_t) + dB_*^2 + B_*^2L'} + B_*C_T + B_*\sum_{t=1}^T |\phi_t^\top \nu_t|\right), \end{aligned}$$

where in (i) we apply Lemma 38, $V_l(s_t)^2 - (P_tV_l)^2 \leq (V_l(s_t) + P_tV_l)|Q_l(s_t, a_t) - P_tV_l|$, $Q_l(s_t, a_t) = \phi_t^\top w_l$, $c_t + P_tV_l = \phi_t^\top \tilde{w}_l$, and we bound the term $\sum_{t=1}^T V_t^2(s'_t) - V_t^2(s_t) = \sum_{l=1}^L V_l^2(s'_{t_{l+1}}) - V_l^2(s_{t_{l+1}})$ as follows: we consider four cases based on how epoch l ends:

1. $s'_{t_{l+1}} = g$, then $V_l^2(s'_{t_{l+1}}) - V_l^2(s_{t_{l+1}}) \leq 0$.
2. $V_l(s'_{t_{l+1}}) = 2B_l$; this happens L' times and the sum of these terms is of order $\tilde{\mathcal{O}}(B_*^2L')$.
3. Triggered by Eq. (4). By Lemma 27, this happens at most $\tilde{\mathcal{O}}(d)$ times and the sum of these terms is of order $\tilde{\mathcal{O}}(dB_*^2)$.
4. $l = L$ is the last epoch. This happens only once and the term is bounded by $\mathcal{O}(B_*^2)$.

In (ii), we apply Lemma 34, $w_l \in \Omega_l$, definition of ν_t , and $\|V_l\|_\infty = \mathcal{O}(B_*)$ by Lemma 4. Solving a quadratic inequality w.r.t $\sqrt{\sum_{t=1}^T \mathbb{V}(P_t, V_l)}$, we have:

$$\sum_{t=1}^T \mathbb{V}(P_t, V_l) = \tilde{\mathcal{O}} \left(dB_*^2 + B_*^2 L' + B_* C_T + B_* \sum_{t=1}^T |\phi_t^\top \nu_t| \right).$$

This completes the proof. \square

Lemma 22. *With probability at least $1 - 4\delta$, $\sum_{t=1}^T |\phi_t^\top \nu_t| \leq \tilde{\mathcal{O}} \left(d^{3.5} \sqrt{B_* C_T} + d^{3.5} \sqrt{B_* \epsilon T} + d^5 B_*^2 \right) + 65d^{2.5} \epsilon T$.*

Proof. Define $u_t = \operatorname{argmax}_{t \leq t' \leq T} |\phi_{t'}^\top \nu_{t'}|$, $V_{j,t} = 2^j I + \sum_{i \in \mathcal{T} \cap [t-1]: |\phi_i^\top \nu_{u_t}| \leq 2^j} \phi_i \phi_i^\top$, and j_t such that $|\phi_{j_t}^\top \nu_{u_t}| \in (2^{j_t-1}, 2^{j_t}]$. Also define $\mathcal{T} = \{t \in [T] : \exists j \in \mathcal{J}_t, |\phi_j^\top \nu_{u_t}| \in (2^{j-1}, 2^j]\}$. Note that when $t \notin \mathcal{T}$, $|\phi_t^\top \nu_t| \leq \epsilon$. Then, for any $t \in \mathcal{T}$:

$$\begin{aligned} |\phi_t^\top \nu_{u_t}| &\leq \|\phi_t\|_{W_{j_t, u_t}^{-1}(\nu_{u_t})} \|\nu_{u_t}\|_{W_{j_t, u_t}(\nu_{u_t})} \\ &\stackrel{(i)}{\leq} 2\sqrt{2}d \|\phi_t\|_{W_{j_t, u_t}^{-1}(\nu_{u_t})} \|\nu_{u_t}\|_{W_{j_t, u_t}(\nu_{u_t})} + \tilde{\mathcal{O}} \left(\sqrt{2^{j_t} \left(\frac{d^3 B_* \epsilon}{u_t} \right)} \right) + \sqrt{2^{j_t+5} d^{2.5} \epsilon} \\ &\stackrel{(ii)}{\leq} 2\sqrt{2}d \|\phi_t\|_{V_{j_t, t}^{-1}} \sqrt{2^{j_t} \left(\sqrt{\sum_{i \leq t_{u_t}} \mathbb{V}(P_i, V_{i, u_t})} + \sqrt{dB_* l_{u_t}} + dB_*^2 \right)} + \tilde{\mathcal{O}} \left(\sqrt{2^{j_t} \left(\frac{d^3 B_* \epsilon}{u_t} \right)} \right) + \sqrt{2^{j_t+5} d^{2.5} \epsilon}, \\ &= \tilde{\mathcal{O}} \left(d \|\phi_t\|_{V_{j_t, t}^{-1}} \sqrt{2^{j_t} \left(\sqrt{d^4 B_*^3 + dB_* C_T + dB_* \epsilon T} + \sqrt{dB_* l_{u_t}} + dB_*^2 \right)} + \sqrt{2^{j_t} \left(\frac{d^3 B_* \epsilon}{u_t} \right)} \right) + \sqrt{2^{j_t+5} d^{2.5} \epsilon} \end{aligned} \quad (\text{Lemma 23})$$

where in (i) we define $\bar{\nu}_{u_t} \in \mathcal{G}_{\epsilon/u_t}(6\sqrt{d}B_{l_{u_t}})$ such that $\|\nu_{u_t} - \bar{\nu}_{u_t}\|_\infty \leq \frac{\epsilon}{u_t}$ and apply

$$\begin{aligned} \|\nu_{u_t}\|_{W_{j_t, u_t}(\nu_{u_t})}^2 &= \Phi_{u_t}^{j_t}(\nu_{u_t}) = 2^{j_t} \|\nu_{u_t}\|_2^2 + \sum_{i < u_t} f_{j_t}(\phi_i^\top \nu_{u_t}) \\ &\leq 2^{j_t} \|\bar{\nu}_{u_t}\|_2^2 + \sum_{i < u_t} f_{j_t}(\phi_i^\top \bar{\nu}_{u_t}) + 2^{j_t} \left(\|\nu_{u_t}\|_2^2 - \|\bar{\nu}_{u_t}\|_2^2 \right) + 2^{j_t+1} \sum_{i < u_t} |\phi_i^\top (\nu_{u_t} - \bar{\nu}_{u_t})| \quad (f_{j_t} \text{ is } (2 \cdot 2^{j_t})\text{-Lipschitz}) \\ &\leq 8d^2 \left(2^{j_t} \|\bar{\nu}_{u_t}\|_2^2 + \sum_{i \leq t_{u_t}} f_{j_t}(\phi_i^\top \bar{\nu}_{u_t}) \right) + \frac{12 \cdot 2^{j_t} dB_{l_{u_t}} \epsilon}{u_t} + 2^{j_t+1} \sqrt{d} \epsilon \quad (\nu_{u_t}, \bar{\nu}_{u_t} \in \mathbb{B}(6\sqrt{d}B_{l_{u_t}})) \\ &\leq 8d^2 \left(2^{j_t} \|\nu_{u_t}\|_2^2 + \sum_{i \leq t_{u_t}} f_{j_t}(\phi_i^\top \nu_{u_t}) \right) + \tilde{\mathcal{O}} \left(2^{j_t} \left(\frac{d^3 B_* \epsilon}{u_t} \right) \right) + 2^{j_t+5} d^{2.5} \epsilon, \quad (\nu_{u_t}, \bar{\nu}_{u_t} \in \mathbb{B}(6\sqrt{d}B_{l_{u_t}})) \end{aligned}$$

and in (ii) we apply Lemma 24 and:

$$W_{j_t, u_t}(\nu_{u_t}) \succcurlyeq W_{j_t, t}(\nu_{u_t}) = 2^{j_t} I + \sum_{i < t} \min\{1, 2^{j_t}/|\phi_i^\top \nu_{u_t}|\} \phi_i^\top \phi_i \stackrel{(i)}{\succcurlyeq} 2^{j_t} I + \sum_{i \in \mathcal{T} \cap [t-1]: |\phi_i^\top \nu_{u_t}| \leq 2^{j_t}} \phi_i \phi_i^\top = V_{j_t, t}.$$

Here, (i) is by $|\phi_i^\top \nu_{u_t}| \leq |\phi_i^\top \nu_{u_t}|$ by the definition of u_t . Reorganizing terms by $|\phi_t^\top \nu_{u_t}| \in (2^{j_t-1}, 2^{j_t}]$, we have for $t \in \mathcal{T}$:

$$|\phi_t^\top \nu_t| \leq |\phi_t^\top \nu_{u_t}| = \tilde{\mathcal{O}} \left(d^2 \|\phi_t\|_{V_{j_t, t}^{-1}}^2 \left(\sqrt{dB_* C_T + dB_* \epsilon T} + d^2 B_*^2 \right) + \frac{d^3 B_* \epsilon}{u_t} \right) + 64d^{2.5} \epsilon.$$

Finally, note that:

$$\begin{aligned} \sum_{t \in \mathcal{T}} |\phi_t^\top \nu_t| &= \tilde{\mathcal{O}} \left(\sum_{t \in \mathcal{T}} d^2 \|\phi_t\|_{V_{j_t, t}^{-1}}^2 \left(\sqrt{dB_* C_T + dB_* \epsilon T} + d^2 B_*^2 \right) + \sum_{t=1}^T \frac{d^3 B_* \epsilon}{t} \right) + 64d^{2.5} \epsilon T \\ &= \tilde{\mathcal{O}} \left(\sqrt{dB_*} \sum_{t \in \mathcal{T}} \mathbb{I} \left\{ \|\phi_t\|_{V_{j_t, t}^{-1}}^2 \geq 1 \right\} + d^2 \sum_{t \in \mathcal{T}} \min \left\{ 1, \|\phi_t\|_{V_{j_t, t}^{-1}}^2 \right\} \left(\sqrt{dB_* C_T + dB_* \epsilon T} + d^2 B_*^2 \right) + d^3 B_* \epsilon \right) + 64d^{2.5} \epsilon T. \end{aligned}$$

The first term is bounded by

$$\begin{aligned} \sqrt{dB_*} \sum_{t \in \mathcal{T}} \mathbb{I} \left\{ \|\phi_t\|_{V_{j_t, t}^{-1}}^2 \geq 1 \right\} &\leq \sqrt{dB_*} \sum_{t \in \mathcal{T}} \min \left\{ 1, \|\phi_t\|_{V_{j_t, t}^{-1}}^2 \right\} \\ &= \sqrt{dB_*} \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}} \mathbb{I} \{j_t = j\} \min \left\{ 1, \|\phi_t\|_{V_{j, t}^{-1}}^2 \right\} = \tilde{\mathcal{O}} \left(d^{1.5} B_* \right), \quad (\text{Lemma 29}) \end{aligned}$$

For the second term:

$$\begin{aligned} &d^2 \sum_{t \in \mathcal{T}} \min \left\{ 1, \|\phi_t\|_{V_{j_t, t}^{-1}}^2 \right\} \left(\sqrt{dB_* C_T + dB_* \epsilon T} + d^2 B_*^2 \right) \\ &= \tilde{\mathcal{O}} \left(d^2 \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}} \mathbb{I} \{j_t = j\} \min \left\{ 1, \|\phi_t\|_{V_{j, t}^{-1}}^2 \right\} \left(\sqrt{dB_* C_T + dB_* \epsilon T} + d^2 B_*^2 \right) \right) \\ &\stackrel{(i)}{=} \tilde{\mathcal{O}} \left(\sum_{j \in \mathcal{J}} d^3 \left(\sqrt{dB_* C_T + dB_* \epsilon T} + d^2 B_*^2 \right) \right) = \tilde{\mathcal{O}} \left(d^3 \left(\sqrt{dB_* C_T + dB_* \epsilon T} + d^2 B_*^2 \right) \right). \end{aligned}$$

where in (i) we apply Lemma 29. Putting everything together, we get:

$$\sum_{t=1}^T |\phi_t^\top \nu_t| \leq \sum_{t \in \mathcal{T}} |\phi_t^\top \nu_t| + \epsilon T \leq \tilde{\mathcal{O}} \left(d^{3.5} \sqrt{B_* C_T} + d^{3.5} \sqrt{B_* \epsilon T} + d^5 B_*^2 \right) + 65d^{2.5} \epsilon T.$$

This completes the proof. \square

Lemma 23. *With probability at least $1 - 3\delta$, $\sum_{i \leq t_l} \mathbb{V}(P_i, V_l) = \tilde{\mathcal{O}} \left(d^3 B_*^3 + B_* C_{t_l} + B_* \epsilon t_l \right)$.*

Proof. Note that when $V_l(s_i) = 0$, $V_l(s_i) - P_i V_l \leq 0$. Otherwise, $Q_l(s_i, a) > 0$ for any a and $V_l(s_i) \leq Q_l(s_i, a_i)$. Therefore, $V_l^2(s_i) - (P_i V_l)^2 = (V_l(s_i) + P_i V_l)(V_l(s_i) - P_i V_l) \leq (V_l(s_i) + P_i V_l) |Q_l(s_i, a_i) - P_i V_l|$. Then with probability at least $1 - \delta$,

$$\begin{aligned} \sum_{i \leq t_l} \mathbb{V}(P_i, V_l) &= \sum_{i \leq t_l} P_i (V_l)^2 - (P_i V_l)^2 \\ &= \sum_{i \leq t_l} (P_i (V_l)^2 - V_l^2(s'_i)) + \sum_{i \leq t_l} (V_l^2(s'_i) - V_l^2(s_i)) + \sum_{i \leq t_l} (V_l^2(s_i) - (P_i V_l)^2) \\ &\stackrel{(i)}{=} \tilde{\mathcal{O}} \left(\sqrt{d \sum_{i \leq t_l} \mathbb{V}(P_i, V_l^2)} + dB_*^2 + B_*^2 + \sum_{i \leq t_l} (V_l(s_i) + P_i V_l) |c_i + \phi_i^\top (w_l - \tilde{w}_l)| \right) \\ &\stackrel{(ii)}{=} \tilde{\mathcal{O}} \left(\sqrt{dB_*} \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_l)} + dB_*^2 + B_* C_{t_l} + B_* \sum_{i \leq t_l} |\phi_i^\top (w_l - \tilde{w}_l)| \right). \end{aligned}$$

In (i) we apply Lemma 25, $\sum_{i \leq t_l} (V_l^2(s'_i) - V_l^2(s_i)) \leq \sum_{i \leq t_l} (V_l^2(s_{i+1}) - V_l^2(s_i)) = \tilde{\mathcal{O}}(B_*^2)$, $V_l^2(s_i) - (P_i V_l)^2 \leq (V_l(s_i) + P_i V_l) |Q_l(s_i, a_i) - P_i V_l|$, $Q_l(s_i, a_i) = \phi_i^\top w_l$, and $c_i + P_i V_l = \phi_i^\top \tilde{w}_l$. In (ii) we apply Lemma 34. For $t \leq t_l$, define $\nu'_t = \arg \max_{\nu = \tilde{w}_l - w, w \in \Omega_l} |\phi_t^\top \nu|$. Then by $w_l \in \Omega_l$ and the definition of ν'_t , we have $|\phi_t^\top (w_l - \tilde{w}_l)| \leq |\phi_t^\top \nu'_t|$.

Now it suffices to bound $\sum_{t \leq t_l} |\phi_t^\top \nu'_t|$. Define $\mathcal{T} = \{t \leq t_l : \exists j \in \mathcal{J}_t, |\phi_t^\top \nu'_t| \in (2^{j-1}, 2^j]\}$ and for $t \in \mathcal{T}$, define $j_t \in \mathcal{J}_t$ such that $|\phi_t^\top \nu'_t| \in (2^{j_t-1}, 2^{j_t}]$. Note that when $t \notin \mathcal{T}$, $|\phi_t^\top \nu'_t| \leq \epsilon$. Also define $V_{j,t} = 2^j I + \sum_{i \in \mathcal{T} \cap [t-1]: |\phi_i^\top \nu'_i| \leq 2^j} \phi_i \phi_i^\top$. Then, for any $t \in \mathcal{T}$, with probability at least $1 - 2\delta$:

$$|\phi_t^\top \nu'_t| \leq \|\phi_t\|_{W_{j_t,t}^{-1}(\nu'_t)} \|\nu'_t\|_{W_{j_t,t}(\nu'_t)} \leq \|\phi_t\|_{V_{j_t,t}^{-1}} \sqrt{2^{j_t} \left(\sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l} + \sqrt{d} B_\star \iota_l + dB_\star^2 \right)},$$

where in the last inequality we apply [Lemma 24](#) and:

$$W_{j_t,t}(\nu'_t) = 2^{j_t} I + \sum_{i \leq t_l} \min\{1, 2^{j_t} / |\phi_i^\top \nu'_i|\} \phi_i \phi_i^\top \stackrel{(i)}{\succcurlyeq} 2^{j_t} I + \sum_{i \in \mathcal{T}: |\phi_i^\top \nu'_i| \leq 2^{j_t}} \phi_i \phi_i^\top = V_{j_t,t}.$$

Here, (i) is by $|\phi_i^\top \nu'_i| \leq |\phi_i^\top \nu'_i|$ by the definition of ν'_i . Reorganizing terms by $|\phi_t^\top \nu'_t| \in (2^{j_t-1}, 2^{j_t}]$, we have:

$$\begin{aligned} \sum_{t \in \mathcal{T}} |\phi_t^\top \nu'_t| &= \tilde{\mathcal{O}} \left(\sum_{t \in \mathcal{T}} \|\phi_t\|_{V_{j_t,t}^{-1}}^2 \left(\sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l} + \sqrt{d} B_\star \iota_l + dB_\star^2 \right) \right) \\ &= \tilde{\mathcal{O}} \left(\sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}} \mathbb{I}\{j_t = j\} \|\phi_t\|_{V_{j,t}^{-1}}^2 \left(\sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l} + \sqrt{d} B_\star \iota_l + dB_\star^2 \right) \right) \\ &\stackrel{(i)}{=} \tilde{\mathcal{O}} \left(\sum_{j \in \mathcal{J}} d \left(\sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l} + \sqrt{d} B_\star \iota_l + dB_\star^2 \right) \right) = \tilde{\mathcal{O}} \left(d \left(\sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l} + \sqrt{d} B_\star \iota_l + dB_\star^2 \right) \right), \end{aligned}$$

where in (i) we apply

$$\sum_{t \in \mathcal{T}} \mathbb{I}\{j_t = j\} \|\phi_t\|_{V_{j,t}^{-1}}^2 = \text{tr} \left(V_{j,l}^{-1} \sum_{t \in \mathcal{T}} \mathbb{I}\{j_t = j\} \phi_t \phi_t^\top \right) \leq \text{tr} \left(V_{j,l}^{-1} V_{j,l} \right) = d.$$

Putting everything together and by $\sum_{t \leq t_l} |\phi_t^\top \nu'_t| \leq \sum_{t \in \mathcal{T}} |\phi_t^\top \nu'_t| + \epsilon t_l$, we have:

$$\begin{aligned} \sum_{i \leq t_l} \mathbb{V}(P_i, V_i) &= \tilde{\mathcal{O}} \left(\sqrt{d} B_\star \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l} + dB_\star^2 + B_\star C_{t_l} + B_\star \left(d^{2.5} B_\star^2 + d^{1.5} \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l} + \epsilon t_l \right) \right) \\ &= \tilde{\mathcal{O}} \left(d^{1.5} B_\star \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l} + d^{2.5} B_\star^3 + B_\star C_{t_l} + B_\star \epsilon t_l \right). \end{aligned}$$

Solving a quadratic inequality w.r.t $\sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l}$, we have $\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) = \tilde{\mathcal{O}} \left(d^3 B_\star^3 + B_\star C_{t_l} + B_\star \epsilon t_l \right)$. \square

Lemma 24. *With probability at least $1 - 2\delta$, for any epoch l , $j \in \mathcal{J}_l$, and $\nu = \tilde{w}_l - \hat{w}$ with $\hat{w} \in \Omega_l$,*

$$\|\nu\|_{W_{j,l}(\nu)}^2 = \mathcal{O} \left(2^j \left(\sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_i) \iota_l} + \sqrt{d} B_\star \iota_l + dB_\star^2 \right) \right).$$

Proof. Define $\epsilon_l^i(w) = \epsilon_{V_l}^i(w) = \phi_i^\top w - c_i - V_l(s'_i)$ and $\eta_l^i(w) = \eta_{V_l}^i(w)$. Note that with probability at least $1 - 2\delta$:

$$\begin{aligned}
 \|\nu\|_{W_{j,l}(\nu)}^2 - 2^j I &= \sum_{i \leq t_l} \text{clip}_j(\phi_i^\top \nu) \phi_i^\top \nu = \sum_{i \leq t_l} \text{clip}_j(\phi_i^\top \nu) (\epsilon_l^i(\tilde{w}_l) - \epsilon_l^i(\hat{w})) \\
 &\leq \sqrt{\sum_{i \leq t_l} \text{clip}_j^2(\phi_i^\top \nu) \eta_l^i(\tilde{w}_l) \iota_l} + \sqrt{\sum_{i \leq t_l} \text{clip}_j^2(\phi_i^\top \nu) \eta_l^i(\hat{w}) \iota_l} + 2B_l 2^j \iota_l && \text{(Lemma 3 and } \hat{w} \in \Omega_l) \\
 &\leq 3 \sqrt{\sum_{i \leq t_l} \text{clip}_j^2(\phi_i^\top \nu) \eta_l^i(\tilde{w}_l) \iota_l} + \sqrt{2 \sum_{i \leq t_l} \text{clip}_j^2(\phi_i^\top \nu) (\phi_i^\top \nu)^2 \iota_l} + 2B_l 2^j \iota_l \\
 &\hspace{15em} (\phi_i^\top \nu = \epsilon_l^i(\tilde{w}_l) - \epsilon_l^i(\hat{w}) \text{ and } (a+b)^2 \leq 2a^2 + 2b^2) \\
 &= \tilde{\mathcal{O}} \left(2^j \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_l) \iota_l} + \sqrt{2^j \sqrt{dB_\star} \sum_{i \leq t_l} \text{clip}_j(\phi_i^\top \nu) (\phi_i^\top \nu) \iota_l} + B_\star 2^j \iota_l \right) \\
 &\hspace{15em} \text{(Lemma 25, } \text{clip}_j(\cdot) \leq 2^j, B_l \leq 2B_\star, \text{ and } |\phi_i^\top \nu| \leq 12\sqrt{dB_\star}) \\
 &= \tilde{\mathcal{O}} \left(2^j \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_l) \iota_l} + \sqrt{2^j \sqrt{dB_\star} \|\nu\|_{W_{j,l}(\nu)}^2 \iota_l} + B_\star 2^j \iota_l \right).
 \end{aligned}$$

Solving a quadratic inequality, we get $\|\nu\|_{W_{j,l}(\nu)}^2 = \mathcal{O} \left(2^j \sqrt{\sum_{i \leq t_l} \mathbb{V}(P_i, V_l) \iota_l} + \sqrt{dB_\star} 2^j \iota_l + 2^j dB_\star^2 \right)$. \square

Lemma 25. *With probability at least $1 - \delta$, for any epoch l , $\sum_{i=1}^{t_l} (P_i V_l - V_l(s'_i))^2 = \tilde{\mathcal{O}}(\sum_{i=1}^{t_l} \mathbb{V}(P_i, V_l) + dB_\star^2)$ and $\sum_{i=1}^{t_l} P_i V_l^2 - V_l^2(s'_i) = \tilde{\mathcal{O}} \left(\sqrt{d \sum_{i=1}^{t_l} \mathbb{V}(P_i, V_l^2)} + dB_\star^2 \right)$.*

Proof. For any $t \in \mathbb{N}_+$, $B \in \{2^i\}_{i=1}^{\lceil \log_2 B_\star \rceil}$, and $w \in \mathcal{G}_{\epsilon/t}(3\sqrt{dB})$, define $X_i = (\phi_i^\top U_B w - c_i - V_{w,B}(s'_i))^2 = (P_i V_{w,B} - V_{w,B}(s'_i))^2$ and \mathbb{E}_i as the conditional expectation conditioned on the interaction history $(s_1, a_1, \dots, s_i, a_i)$. Note that $\mathbb{E}_i[X_i] = \mathbb{V}(P_i, V_{w,B})$ and $|X_i| \leq 4B^2$. Then by Lemma 37 with $\lambda = \frac{1}{4B^2}$, with probability at least $1 - \delta'$ with $\delta' = \delta / (8(t \log_2(2B))^2 (6\sqrt{dB}t/\epsilon)^d)$, we have:

$$\sum_{i=1}^t (X_i - \mathbb{V}(P_i, V_{w,B})) \leq \lambda \sum_{i=1}^t \mathbb{E}_i[X_i^2] + \frac{\ln(1/\delta')}{\lambda} \leq \sum_{i=1}^t \mathbb{V}(P_i, V_{w,B}) + \tilde{\mathcal{O}}(dB_\star^2).$$

Reorganizing terms and by a union bound, we have with probability at least $1 - \delta/2$, for any $t \in \mathbb{N}_+$, $B \in \{2^i\}_{i=1}^{\lceil \log_2 B_\star \rceil}$, and $w \in \mathcal{G}_{\epsilon/t}(3\sqrt{dB})$:

$$\sum_{i=1}^t (P_i V_{w,B} - V_{w,B}(s'_i))^2 = \sum_{i=1}^t X_i \leq 2 \sum_{i=1}^t \mathbb{V}(P_i, V_{w,B}) + \tilde{\mathcal{O}}(dB_\star^2). \quad (16)$$

Moreover, for any $t \in \mathbb{N}_+$, $B \in \{2^i\}_{i=1}^{\lceil \log_2 B_\star \rceil}$, and $w \in \mathcal{G}_{\epsilon/t}(3\sqrt{dB})$, by Lemma 38, with probability at least $1 - \delta'$:

$$\sum_{i=1}^t P_i V_{w,B}^2 - V_{w,B}^2(s'_i) = \tilde{\mathcal{O}} \left(\sqrt{\sum_{i=1}^t \mathbb{V}(P_i, V_{w,B}^2) \ln \frac{1}{\delta'}} + B_\star^2 \ln \frac{1}{\delta'} \right) = \tilde{\mathcal{O}} \left(\sqrt{d \sum_{i=1}^t \mathbb{V}(P_i, V_{w,B}^2)} + dB_\star^2 \right). \quad (17)$$

Then again by a union bound, the equation above holds with probability at least $1 - \delta/2$ for any $t \in \mathbb{N}_+$, $B \in \{2^i\}_{i=1}^{\lceil \log_2 B_\star \rceil}$, and $w \in \mathcal{G}_{\epsilon/t}(3\sqrt{dB})$.

Now for any epoch l , pick $w'_l \in \mathcal{G}_{\epsilon/t_l}(3\sqrt{dB_l})$ such that $\|w'_l - w_l\|_\infty \leq \epsilon/t_l$. Also define $V'_l = V_{w'_l, B_l}$ and $\tilde{w}'_l = U_{B_l} w'_l$. Then similar to Eq. (14) and Eq. (15), we have

$$\|V_l - V'_l\|_\infty \leq \sqrt{d\epsilon}/t_l, \quad \|\tilde{w}_l - \tilde{w}'_l\|_2 \leq d\epsilon/t_l. \quad (18)$$

For the first statement:

$$\begin{aligned}
 \sum_{i=1}^{t_l} (P_i V_l - V_l(s'_i))^2 &= \sum_{i=1}^{t_l} (\phi_i^\top \tilde{w}_l - c_i - V_l(s'_i))^2 \\
 &\leq 3 \sum_{i=1}^{t_l} ((\phi_i^\top \tilde{w}'_l - c_i - V'_l(s'_i))^2 + (V_l(s'_i) - V'_l(s'_i))^2 + (\phi_i^\top (\tilde{w}_l - \tilde{w}'_l))^2) \\
 &\leq \tilde{O} \left(\sum_{i=1}^{t_l} \mathbb{V}(P_i, V'_l) + dB_\star^2 \right) + \frac{6d^2 \epsilon^2}{t_l} \quad (\text{Eq. (16) and Eq. (18)}) \\
 &= \tilde{O} \left(\sum_{i=1}^{t_l} \mathbb{V}(P_i, V_l) + \sum_{i=1}^{t_l} \mathbb{V}(P_i, V'_l - V_l) + dB_\star^2 \right) = \tilde{O} \left(\sum_{i=1}^{t_l} \mathbb{V}(P_i, V_l) + dB_\star^2 \right). \\
 &(\text{VAR}[X + Y] \leq 2\text{VAR}[X] + 2\text{VAR}[Y], \mathbb{V}(P_i, V'_l - V_l) \leq \|V'_l - V_l\|_\infty^2, \text{Eq. (18), and } d\epsilon \leq 1)
 \end{aligned}$$

For the second statement,

$$\begin{aligned}
 \sum_{i=1}^{t_l} P_i (V_l)^2 - V_l^2(s'_i) &= \sum_{i=1}^{t_l} (P_i (V'_l)^2 - V_l^2(s'_i)) + \sum_{i=1}^{t_l} (P_i (V_l)^2 - P_i (V'_l)^2) + \sum_{i=1}^{t_l} (V_l^2(s'_i) - V_l^2(s'_i)) \\
 &\leq \tilde{O} \left(\sqrt{d \sum_{i=1}^{t_l} \mathbb{V}(P_i, V_l^2) + dB_\star^2} \right) + 4B_\star \sum_{i=1}^{t_l} \|V_l - V'_l\|_\infty \\
 &\quad (\text{Eq. (17) and } \max\{\|V_l\|_\infty, \|V'_l\|_\infty\} \leq 4B_\star) \\
 &\leq \tilde{O} \left(\sqrt{d \sum_{i=1}^{t_l} \mathbb{V}(P_i, V_l^2) + dB_\star^2} + \sqrt{d \sum_{i=1}^{t_l} \mathbb{V}(P_i, V_l^2 - V_l'^2) + \sqrt{d} B_\star \epsilon} \right) \\
 &\quad (\text{VAR}[X + Y] \leq 2\text{VAR}[X] + 2\text{VAR}[Y], \sqrt{x+y} \leq \sqrt{x} + \sqrt{y}, \text{ and Eq. (18)}) \\
 &= \tilde{O} \left(\sqrt{d \sum_{i=1}^{t_l} \mathbb{V}(P_i, V_l^2) + dB_\star^2} \right). \quad (\text{Eq. (18) and } \epsilon \leq 1)
 \end{aligned}$$

Thus, the second statement is proved. \square

For the next lemma, we define the following auxiliary function:

$$g_j(x) = \begin{cases} x^2, & |x| \leq 2^j, \\ 2^{j+1}x - 4^j, & x > 2^j \\ -2^{j+1}x - 4^j, & x < -2^j \end{cases}$$

Note that $g_j(x)$ is convex and $f_j(x) \leq g_j(x) \leq 2f_j(x)$.

Lemma 26. For $\lambda \in (0, 1]$, $g_j(\lambda x) \geq \lambda^2 g_j(x)$.

Proof. Let $\ell = 2^j$. When $|\lambda x| \leq \ell$, we have: $g_j(\lambda x) = \lambda^2 x^2 \geq \lambda^2 g_j(x)$. When $\lambda x > \ell$ (arguments are similar for $\lambda x < -\ell$), we have $x > \ell$, and

$$\begin{aligned}
 g_j(\lambda x) - \lambda^2 g_j(x) &= 2\ell\lambda x - \ell^2 - \lambda^2(2\ell x - \ell^2) = 2\ell\lambda x(1 - \lambda) - \ell^2(1 - \lambda^2) \\
 &= (1 - \lambda)\ell(2\lambda x - (1 + \lambda)\ell) \geq 0.
 \end{aligned}$$

\square

Lemma 27. Fix $2^j \geq \epsilon > 0$. Let $x_1, \dots, x_t \in \mathbb{B}(1)$. If there exists $0 = \tau_0 < \tau_1 < \dots < \tau_z = t$ such that for each $1 \leq \zeta \leq z$, there exists $\nu_\zeta \in \mathbb{B}(B) \setminus \mathbb{B}(\epsilon)$ for some $B > \epsilon$ such that

$$\sum_{i=1}^{\tau_\zeta} f_j(x_i^\top \nu_\zeta) + 2^j \|\nu_\zeta\|_2^2 > 8d^2 \left(\sum_{i=1}^{\tau_{\zeta-1}} f_j(x_i^\top \nu_\zeta) + 2^j \|\nu_\zeta\|_2^2 \right) \quad (19)$$

Then, $z = \tilde{\mathcal{O}}(d)$.

Proof. Note that when Eq. (19) holds:

$$\begin{aligned} \sum_{i=1}^{\tau_\zeta} g_j(x_i^\top \nu_\zeta) + 2^j \|\nu_\zeta\|_2^2 &\geq \sum_{i=1}^{\tau_\zeta} f_j(x_i^\top \nu_\zeta) + 2^j \|\nu_\zeta\|_2^2 > 8d^2 \left(\sum_{i=1}^{\tau_{\zeta-1}} f_j(x_i^\top \nu_\zeta) + 2^j \|\nu_\zeta\|_2^2 \right) \\ &\geq 4d^2 \left(\sum_{i=1}^{\tau_{\zeta-1}} g_j(x_i^\top \nu_\zeta) + 2^j \|\nu_\zeta\|_2^2 \right). \end{aligned} \quad (20)$$

Thus, it suffices to bound the number times Eq. (20) holds. Define $E_t(\nu) = \sum_{i=1}^t g_j(x_i^\top \nu) + 2^j \|\nu\|_2^2$. Clearly E_t is convex since g_j is convex, and $E_t(\nu) \in [2^j \epsilon^2, 2^j B^2 + 2t2^j B]$ for $\nu \in \mathbb{B}(B) \setminus \mathbb{B}(\epsilon)$. Define:

$$\Lambda = \{i \in \mathbb{Z} : \lceil \log_2(2^j \epsilon^2) \rceil \leq i \leq \lceil \log_2(2^j B^2 + 2t2^j B) \rceil\}.$$

For each ζ , there exists $i_\zeta \in \Lambda$ such that $E_{\tau_{\zeta-1}}(\nu_\zeta) \in (2^{i_\zeta-1}, 2^{i_\zeta}]$. Define $D_{t',i} = \{\nu \in \mathbb{B}(B) : E_{t'}(\nu) \leq 2^i\}$. Note that $\nu_\zeta \in D_{\tau_{\zeta-1}, i_\zeta}$, and $D_{t',i}$ is a symmetric convex set since E_t is a convex function and $E_t(\nu) = E_t(-\nu)$. By Lemma 26, we have $E_{\tau_\zeta}(\nu_\zeta/d) \geq \frac{1}{d^2} E_{\tau_\zeta}(\nu_\zeta) > 4E_{\tau_{\zeta-1}}(\nu_\zeta) > 2^{i_\zeta}$. Therefore, $\nu_\zeta/d \notin D_{\tau_\zeta, i_\zeta}$, which means that in the direction of ν_ζ , the intercept of D_{τ_ζ, i_ζ} is at most $1/d$ times of that of $D_{\tau_{\zeta-1}, i_\zeta}$. By Lemma 35, we have: $\text{Vol}(D_{\tau_\zeta, i_\zeta}) \leq \frac{6}{7} \text{Vol}(D_{\tau_{\zeta-1}, i_\zeta})$. Note that when $\|\nu\|_2 \leq 2^j$, we have $E_t(\nu) \leq (t+2^j) \|\nu\|_2^2$. Therefore, when $\|\nu\|_2 \leq \epsilon' = \sqrt{2^j/(t+2^j)}\epsilon$, we have $E_t(\nu) \leq 2^j \epsilon^2$. Therefore, $\text{Vol}(D_{t,i}) \geq \text{Vol}(\mathbb{B}(\epsilon'))$ for $i \in \Lambda$. Due to the fact that $D_{t,i}$ is decreasing in t , we have

$$z = \mathcal{O}(|\Lambda| \log_{7/6}(\text{Vol}(\mathbb{B}(B))/\text{Vol}(\mathbb{B}(\epsilon')))) = \tilde{\mathcal{O}}(d).$$

This completes the proof. \square

D. Auxiliary Lemmas

Lemma 28. *If $x \leq (a\sqrt{x} + b) \ln^p(cx)$ for some $a, b, c > 0$ and absolute constant $p \geq 1$, then $x = \tilde{\mathcal{O}}(a^2 + b)$.*

Proof. First note that $x \leq 2b \ln^p(cx)$ implies $x \leq 2b(2p)^p \sqrt{cx}$ by $\ln x \leq x$ for $x > 0$, which gives $x \leq 4(2p)^{2p} b^2 c$. Plugging this back, we get $x \leq 2b \ln^p(4(2p)^{2p} b^2 c^2)$. Therefore, $x > 2b \ln^p(4(2p)^{2p} b^2 c^2)$ implies $x > 2b \ln^p(cx)$. Next, note that $x \leq 2a\sqrt{x} \ln^p(cx)$ implies $x \leq 2ac^{1/4} (4p)^p x^{3/4}$ by $\ln x \leq x$ for $x > 0$, which gives $x \leq 16(4p)^{4p} a^4 c$. Plugging this back, we get $x \leq 2a\sqrt{x} \ln^p(16(4p)^{4p} a^4 c^2)$, which gives $x \leq 4a^2 \ln^{2p}(16(4p)^{4p} a^4 c^2)$. Therefore, $x > 2a\sqrt{x} \ln^p(16(4p)^{4p} a^4 c^2)$ implies $x > 2a\sqrt{x} \ln^p(cx)$. Thus, $x > 4a^2 \ln^{2p}(16(4p)^{4p} a^4 c^2) + 2b \ln^p(4(2p)^{2p} b^2 c^2)$ implies $\frac{x}{2} > a\sqrt{x} \ln^p(cx)$ and $\frac{x}{2} > b \ln^p(cx)$, which implies $x > (a\sqrt{x} + b) \ln^p(cx)$. Taking the contrapositive, the statement is proved. \square

Lemma 29. (Abbasi-Yadkori et al., 2011, Lemma 11) *Let $\{X_i\}_{i=1}^\infty$ be a sequence in \mathbb{R}^d , V a $d \times d$ positive definite matrix, and define $V_n = V + \sum_{i=1}^n X_i X_i^\top$. Then, $\sum_{i=1}^n \min\{1, \|X_i\|_{V_{i-1}}^2\} \leq 2 \ln \frac{\det(V_n)}{\det(V)}$ for any $n \geq 1$.*

Lemma 30. (Abbasi-Yadkori et al., 2011, Lemma 12) *Let A, B be positive semi-definite matrices such that $A \succcurlyeq B$. Then, we have $\sup_{x \neq 0} \frac{x^\top A x}{x^\top B x} \leq \frac{\det(A)}{\det(B)}$.*

Lemma 31. (Wei et al., 2021b, Lemma 11) *Let $\{x_t\}_{t=1}^\infty$ be a martingale sequence on state space \mathcal{X} w.r.t a filtration $\{\mathcal{F}_t\}_{t=0}^\infty$, $\{\phi_t\}_{t=1}^\infty$ be a sequence of random vectors in \mathbb{R}^d so that $\phi_t \in \mathcal{F}_{t-1}$ and $\|\phi_t\| \leq 1$, $\Lambda_t = \lambda I + \sum_{s=1}^{t-1} \phi_s \phi_s^\top$, and $\mathcal{V} \subseteq \mathbb{R}^{\mathcal{X}}$ be a set of functions defined on \mathcal{X} with \mathcal{N}_ε as its ε -covering number w.r.t the distance $\text{dist}(v, v') = \sup_x |v(x) - v'(x)|$ for some $\varepsilon > 0$. Then for any $\delta > 0$, we have with probability at least $1 - \delta$, for all $t > 0$ and $v \in \mathcal{V}$ so that $\sup_x |v(x)| \leq B$:*

$$\left\| \sum_{s=1}^{t-1} \phi_s (v(x_s) - \mathbb{E}[v(x_s) | \mathcal{F}_{s-1}]) \right\|_{\Lambda_t^{-1}}^2 \leq 4B^2 \left[\frac{d}{2} \ln \left(\frac{t+\lambda}{\lambda} \right) + \ln \frac{\mathcal{N}_\varepsilon}{\delta} \right] + \frac{8t^2 \varepsilon^2}{\lambda}.$$

Lemma 32. (Wei et al., 2021b, Lemma 12) *Let \mathcal{V} be a class of mappings from \mathcal{X} to \mathbb{R} parameterized by $\alpha \in [-D, D]^n$. Suppose that for any $v \in \mathcal{V}$ (parameterized by α) and $v' \in \mathcal{V}'$ (parameterized by α'), the following holds:*

$$\sup_{x \in \mathcal{X}} |v(x) - v'(x)| \leq L \|\alpha - \alpha'\|_1.$$

Then, $\ln \mathcal{N}_\varepsilon \leq n \ln \left(\frac{2DLn}{\varepsilon} \right)$, where \mathcal{N}_ε is the ε -covering number of \mathcal{V} with respect to the distance $\text{dist}(v, v') = \sup_{x \in \mathcal{X}} |v(x) - v'(x)|$.

Lemma 33. (Zhou et al., 2021a, Theorem 4.1) Let $\{\mathcal{F}_t\}_{t=1}^\infty$ be a filtration, $\{x_t, \eta_t\}_{t \geq 1}$ a stochastic process so that $x_t, \eta_t \in \mathbb{R}^d$ and $x_t \in \mathcal{F}_t, \eta_t \in \mathcal{F}_{t+1}$. Moreover, define $y_t = \langle \mu^*, x_t \rangle + \eta_t$ and we have:

$$|\eta_t| \leq R, \mathbb{E}[\eta_t | \mathcal{F}_t] = 0, \mathbb{E}[\eta_t^2 | \mathcal{F}_t] \leq \sigma^2, \|x_t\|_2 \leq L.$$

Then with probability at least $1 - \delta$, we have for any $t \geq 1$:

$$\left\| \sum_{i=1}^t x_i \eta_i \right\|_{Z_t^{-1}} \leq \beta_t, \|\mu_t - \mu^*\|_{Z_t} \leq \beta_t + \sqrt{\lambda} \|\mu^*\|_2,$$

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

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

Lemma 34. (Chen et al., 2021a, Lemma 30) For any two random variables X, Y , we have:

$$\text{VAR}[XY] \leq 2\text{VAR}[X] \|Y\|_\infty^2 + 2(\mathbb{E}[X])^2 \text{VAR}[Y].$$

Consequently, $\|X\|_\infty \leq C \implies \text{VAR}[X^2] \leq 4C^2 \text{VAR}[X]$.

Lemma 35. (Zhang et al., 2021, Lemma 16) Let D be a bounded symmetric convex subset of \mathbb{R}^d with $d \geq 2$. Suppose $u \in \partial D$, that is, u is on the boundary of D , and D' is another bounded symmetric convex set such that $D \subseteq D'$ and $d \cdot u \in \partial D'$. Then $\text{Vol}(D') \geq \frac{7}{6} \text{Vol}(D)$, where $\text{Vol}(S)$ is the volume of the set S .

Lemma 36. (Zhang et al., 2021, Theorem 4) Let $\{X_i\}_{i=1}^n$ be a martingale difference sequence and $|X_i| \leq b$ almost surely. Then for $\delta < e^{-1}$, we have with probability at least $1 - 6\delta \log_2 n$,

$$\left| \sum_{i=1}^n X_i \right| \leq 8 \sqrt{\sum_{i=1}^n X_i^2 \ln \frac{1}{\delta}} + 16b \ln \frac{1}{\delta}.$$

Lemma 37. (Jin et al., 2020a, Lemma 9) Let $\{X_i\}_{i=1}^n$ be a martingale difference sequence adapted to the filtration $\{\mathcal{F}_i\}_{i=0}^n$, and $X_i \leq B$ almost surely for some $B > 0$. Then, for any $\lambda \in [0, 1/B]$, with probability at least $1 - \delta$:

$$\sum_{i=1}^n X_i \leq \lambda \sum_{i=1}^n \mathbb{E}[X_i^2 | \mathcal{F}_{i-1}] + \frac{\ln(1/\delta)}{\lambda}.$$

Lemma 38. Let $\{X_i\}_{i=1}^\infty$ be a martingale difference sequence adapted to the filtration $\{\mathcal{F}_i\}_{i=0}^\infty$ and $|X_i| \leq B$ for some $B > 0$. Then with probability at least $1 - \delta$, for all $n \geq 1$ simultaneously,

$$\left| \sum_{i=1}^n X_i \right| \leq 3 \sqrt{\sum_{i=1}^n \mathbb{E}[X_i^2 | \mathcal{F}_{i-1}] \ln \frac{4B^2 n^3}{\delta}} + 2B \ln \frac{4B^2 n^3}{\delta}.$$

Proof. For each $n \geq 1$, applying Lemma 37 to $\{X_i\}_{i=1}^n$ and $\{-X_i\}_{i=1}^n$ with each $\lambda \in \Lambda = \left\{ \frac{1}{B2^i} \right\}_{i=0}^{\lceil \log_2 n \rceil}$, we have with probability at least $1 - \frac{\delta}{2n^2}$, for any $\lambda \in \Lambda$,

$$\left| \sum_{i=1}^n X_i \right| \leq \lambda \sum_{i=1}^n \mathbb{E}[X_i^2 | \mathcal{F}_{i-1}] + \frac{\ln \frac{4Bn^3}{\delta}}{\lambda}, \quad (21)$$

Note that there exists $\lambda^* \in \Lambda$ such that $\lambda^* / \min \left\{ 1/B, \sqrt{\frac{\ln(4Bn^3/\delta)}{\sum_{i=1}^n \mathbb{E}[X_i^2 | \mathcal{F}_{i-1}]}} \right\} \in (\frac{1}{2}, 1]$. Plugging λ^* into Eq. (21), we get

$|\sum_{i=1}^n X_i| \leq 3 \sqrt{\sum_{i=1}^n \mathbb{E}[X_i^2 | \mathcal{F}_{i-1}] \ln \frac{4Bn^3}{\delta}} + 2B \ln \frac{4Bn^3}{\delta}$. By a union bound over n , the statement is proved. \square

Lemma 39. (Cohen et al., 2020, Lemma D.4) and (Cohen et al., 2021, Lemma E.2) Let $\{X_i\}_{i=1}^\infty$ be a sequence of random variables w.r.t to the filtration $\{\mathcal{F}_i\}_{i=0}^\infty$ and $X_i \in [0, B]$ almost surely. Then with probability at least $1 - \delta$, for all $n \geq 1$ simultaneously:

$$\sum_{i=1}^n \mathbb{E}[X_i | \mathcal{F}_{i-1}] \leq 2 \sum_{i=1}^n X_i + 4B \ln \frac{4n}{\delta},$$

$$\sum_{i=1}^n X_i \leq 2 \sum_{i=1}^n \mathbb{E}[X_i | \mathcal{F}_{i-1}] + 8B \ln \frac{4n}{\delta}.$$