

MODEL-BASED RL AS A MINIMALIST APPROACH TO HORIZON-FREE AND SECOND-ORDER BOUNDS

Anonymous authors

Paper under double-blind review

ABSTRACT

Learning a transition model via Maximum Likelihood Estimation (MLE) followed by planning inside the learned model is perhaps the most standard and simplest Model-based Reinforcement Learning (RL) framework. In this work, we show that such a simple Model-based RL scheme, when equipped with optimistic and pessimistic planning procedures, achieves strong regret and sample complexity bounds in online and offline RL settings. Particularly, we demonstrate that under the conditions where the trajectory-wise reward is normalized between zero and one and the transition is time-homogenous, it achieves nearly horizon-free and second-order bounds.

1 INTRODUCTION

The framework of model-based Reinforcement Learning (RL) often consists of two steps: fitting a transition model using data and then performing planning inside the learned model. Such a simple framework turns out to be powerful and has been used extensively in practice on applications such as robotics and control (e.g., (Aboaf et al., 1989; Deisenroth et al., 2011; Venkatraman et al., 2017; Williams et al., 2017; Chua et al., 2018; Kaiser et al., 2019; Yang et al., 2023)).

The simplicity of model-based RL also attracts researchers to analyze its performance in settings such as online RL (Sun et al., 2019) and offline RL (Uehara & Sun, 2021). Mania et al. (2019) showed that this simple scheme — fitting model via data followed by optimal planning inside the model, has a strong performance guarantee under the classic linear quadratic regulator (LQR) control problems. Liu et al. (2023) showed that this simple MBRL framework when equipped with optimism in the face of the uncertainty principle, can achieve strong sample complexity bounds for a wide range of online RL problems with rich function approximation for the models. For offline settings where the model can only be learned from a static offline dataset, Uehara & Sun (2021) showed that MBRL equipped with the pessimism principle can again achieve robust performance guarantees for a large family of MDPs. Ross & Bagnell (2012) showed that in the hybrid RL setting where one has access to both online and offline data, this simple MBRL framework again achieves favorable performance guarantees without any optimism/pessimism algorithm design.

In this work, we do not create new MBRL algorithms, instead, we show that the extremely simple and standard MBRL algorithm — fitting models using Maximum Likelihood Estimation (MLE), followed by optimistic/pessimistic planning (depending on whether operating in online RL or offline RL mode), can already achieve surprising theoretical guarantees. Particularly, we show that under the conditions that trajectory-wise reward is normalized between zero and one, and the transition is time-homogenous, they can achieve nearly horizon-free and instance-dependent regret and sample complexity bounds, in both online and offline RL with non-linear function approximation. Nearly horizon-free bounds mean that the regret or sample complexity bounds have no explicit polynomial dependence on the horizon H . The motivation for studying horizon-free RL is to see if RL problems are harder than bandits due to the longer horizon planning in RL. *Our result here indicates that, even under non-linear function approximation, long-horizon planning is not the bottleneck of achieving statistical efficiency in RL.* For instance-dependent bounds, we focus on second-order bounds. A second-order regret bound scales with respect to the variances of the returns of policies and also directly implies a first-order regret bound which scales with the expected reward of the optimal policy. Thus our instance-dependent bounds can be small under situations such as nearly-deterministic systems or the optimal policy having a small value. When specializing to the case of deterministic ground truth transitions

(but the algorithm does not need to know this a priori), we show that these simple MBRL algorithms demonstrate a faster convergence rate than the worst-case rates. The key message of our work is

Simple and standard MLE-based MBRL algorithms are sufficient for achieving nearly horizon-free and second-order bounds in online and offline RL with function approximation.

We provide a fairly standard analysis to support the above claim. Our analysis follows the standard frameworks of optimism/pessimism in the face of uncertainty. For online RL, we use ℓ_1 Eluder dimension (Liu et al., 2022; Wang et al., 2024), a condition that uses both the MDP structure and the function class, to capture the structural complexity of exploration. For offline RL, we use the similar concentrability coefficient in Ye et al. (2024) to capture the coverage condition of the offline data. The key technique we leverage is the *triangular discrimination* – a divergence that is equivalent to the squared Hellinger distance up to some universal constants. Triangular discrimination was used in contextual bandit and model-free RL for achieving first-order and second-order instance-dependent bounds (Foster & Krishnamurthy, 2021; Wang et al., 2023, 2024). Here we show that it also plays an important role in achieving horizon-free bounds. Our contributions can be summarized as follows.

1. Our results extend the scope of the prior work on horizon-free RL which only applies to tabular MDPs or MDPs with linear functions. Given a finite model class \mathcal{P} (which could be exponentially large), we show that in online RL, the agent achieves an $O\left(\sqrt{(\sum_k \text{VaR}_{\pi^k}) \cdot d_{\text{RL}} \log(KH|\mathcal{P}|/\delta)} + d_{\text{RL}} \log(KH|\mathcal{P}|/\delta)\right)$ regret, where K is the number of episodes, d_{RL} is the ℓ_1 Eluder dimension, VaR_{π^k} is the variance of the total reward of policy π^k learned in episode k and $\delta \in (0, 1)$ denotes the failure probability. Similarly, for offline RL, the agent achieves an $O\left(\sqrt{C^{\pi^*} \text{VaR}_{\pi^*} \log(|\mathcal{P}|/\delta)/K} + C^{\pi^*} \log(|\mathcal{P}|/\delta)/K\right)$ performance gap in finding a comparator policy π^* , where C^{π^*} is the single policy concentrability coefficient over π^* , K denotes the number of offline trajectories, VaR_{π^*} is the variance of the total reward of π^* . For offline RL with finite \mathcal{P} , our result is *completely horizon-free*, not even with $\log H$ dependence.
2. When specializing to MDPs with deterministic ground truth transition (but rewards, and models in the model class could still be stochastic), we show that the same simple MBRL algorithms can adapt to the deterministic environment and achieve a better statistical complexity. For online RL, the regret becomes $O(d_{\text{RL}} \log(KH|\mathcal{P}|/\delta))$, which only depends on the number of episodes K poly-logarithmically. For offline RL, the performance gap to a comparator policy π^* becomes $O\left(C^{\pi^*} \log(|\mathcal{P}|/\delta)/K\right)$, which is tighter than the worst-case $O(1/\sqrt{K})$ rate. All our results can be extended to continuous model class \mathcal{P} using bracket number as the complexity measure.

Overall, our work identifies the *minimalist* algorithms and analysis for nearly horizon-free and instance-dependent (first & second-order) online & offline RL. By saying “*minimalist*” we mean the algorithm designs and analysis are much simpler than previous work on horizon-free and second-order RL.

2 RELATED WORK

Model-based RL. Learning transition models with function approximation and planning with the learned model is a standard approach in RL and control. In the control literature, certainty-equivalence control learns a model from some data and plans using the learned model, which is simple but effective for controlling systems such as Linear Quadratic Regulators (LQRs) (Mania et al., 2019). In RL, such a simple model-based framework has been widely used in theory with rich function approximation, for online RL (Sun et al., 2019; Foster et al., 2021; Song & Sun, 2021; Zhan et al., 2022; Liu et al., 2022, 2023; Zhong et al., 2022), offline RL (Uehara & Sun, 2021), RL with representation learning (Agarwal et al., 2020; Uehara et al., 2021), and hybrid RL using both online and offline data for model fitting (Ross & Bagnell, 2012). Our work builds on the maximum-likelihood estimation (MLE) approach, a standard method for estimating transition models in model-based RL.

Horizon-free and Instance-dependent bounds. Most existing works on horizon-free RL typically focus on tabular settings or linear settings. For instance, Wang et al. (2020) firstly studied horizon-free RL for tabular MDPs and proposed an algorithm that depends on horizon logarithmically. Several follow-up work studied horizon-free RL for tabular MDP with better sample complexity (Zhang et al., 2021a), offline RL (Ren et al., 2021), stochastic shortest path (Tarbouriech et al., 2021) and RL with linear function approximation (Kim et al., 2022; Zhang et al., 2021b; Zhou & Gu, 2022; Di et al., 2023; Zhang et al., 2024b, 2023; Zhao et al., 2023b). Note that all these works have

108 logarithmic dependence on the horizon H . For the tabular setting, recent work further improved
 109 the regret or sample complexity to be completely independent of the horizon (i.e., removing the
 110 logarithmic dependence on the horizon) (Li et al., 2022; Zhang et al., 2022) with a worse dependence
 111 on the cardinality of state and action spaces $|\mathcal{S}|$ and $|\mathcal{A}|$. A recent work Li & Yang (2023) further
 112 improved the dependence on $|\mathcal{S}|$ and $|\mathcal{A}|$. To compare with, we show that simple MBRL algorithms
 113 are already enough to achieve completely horizon-free (i.e., no log dependence) sample complexity
 114 for offline RL when the transition model class is finite, and we provide a simpler approach to achieve
 115 the nearly horizon-free results for tabular MDPs, compared with Zhang et al. (2021a). A recent
 116 work (Huang et al., 2024) also studied the horizon-free and instance-dependent online RL in the
 117 function approximation setting with small Eluder dimensions. They estimated the variances to
 118 conduct variance-weighted regression. To compare, in our online RL part, we use the simple and
 119 standard MLE-based MBRL approach and analysis to get similar guarantees. A more recent work
 120 also studied horizon-free behavior cloning Foster et al. (2024), which is different from our settings.

121 Besides horizon-free RL, another line of work aimed to provide algorithms with instance-dependent
 122 sample complexity/regret bounds, which often enjoy tighter statistical complexity compared with
 123 previous work. To mention a few, Zanette & Brunskill (2019) proposed an EULER algorithm with
 124 an instance-dependent regret which depends on the maximum variance of the policy return over all
 125 policies. Later, Wagenmaker et al. (2022); Wang et al. (2023) proposed algorithms with first-order
 126 regret bounds. A more refined second-order regret bound has been studied. The second-order regret
 127 bound is a well-studied instance-dependent bound in the online learning and bandit literature (Cesa-
 128 Bianchi et al., 2007; Ito et al., 2020; Olkhovskaya et al., 2024), and compared to the bounds that
 129 depend on the maximum variance over all policies, it can be much smaller and it also implies a first-
 130 order regret bound. Zhang et al. (2024a); Zhou et al. (2023) proposed algorithms for tabular MDP with
 131 second-order regret bounds. Zhao et al. (2023b) studied the RL with linear function approximation and
 132 proposed algorithms with both horizon-free and variance-dependent regret bounds. The closest work
 133 to us is Wang et al. (2023, 2024), which used model-free distributional RL methods to achieve first-
 134 order and second-order regret bounds in RL. Their approach relies on distributional RL, a somewhat
 135 non-conventional approach for RL. Their regret bounds have explicit polynomial dependence on the
 136 horizon. Our work focuses on the more conventional model-based RL algorithms and demonstrates
 137 that they are indeed sufficient to achieve horizon-free and second-order regret bounds.

137 3 PRELIMINARIES

138
 139 **Markov Decision Processes.** We consider finite horizon time-homogenous MDP $\mathcal{M} =$
 140 $\{\mathcal{S}, \mathcal{A}, H, P^*, r, s_0\}$ where \mathcal{S}, \mathcal{A} are the state and action space (could be large or even continuous),
 141 $H \in \mathbb{N}^+$ is the horizon, $P^* : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$ is the ground truth transition, $r : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ is the reward
 142 signal which we assume is known to the learner, and s_0 is the fixed initial state.¹ Note that the transi-
 143 tion P^* here is time-homogenous. For notational easiness, we denote $[K - 1] = \{0, 1, \dots, K - 1\}$.

144 We denote π as a deterministic non-stationary policy $\pi = \{\pi_0, \dots, \pi_{H-1}\}$ where $\pi_h : \mathcal{S} \mapsto \mathcal{A}$ maps
 145 from a state to an action. Let Π denote the set of all such policies. $V_h^\pi(s)$ represents the expected
 146 total reward of policy π starting at $s_h = s$, and $Q_h^\pi(s, a)$ is the expected total reward of the process
 147 of executing a at s at time step h followed by executing π to the end. The optimal policy π^* is
 148 defined as $\pi^* = \operatorname{argmax}_\pi V_0^\pi(s_0)$. For notation simplicity, we denote $V^\pi := V_0^\pi(s_0)$. We will denote
 149 $d_h^\pi(s, a)$ as the state-action distribution induced by policy π at time step h . We sometimes will
 150 overload notation and denote $d_h^\pi(s)$ as the corresponding state distribution at h . Sampling $s \sim d_h^\pi$
 151 means executing π starting from s_0 to h and returning the state at time step h .

152 Since we use the model-based approach for learning, we define a general model class $\mathcal{P} \subset \mathcal{S} \times \mathcal{A} \mapsto$
 153 $\Delta(\mathcal{S})$. Given a transition P , we denote $V_{h;P}^\pi$ and $Q_{h;P}^\pi$ as the value and Q functions of policy π
 154 under the model P . Given a function $f : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$, we denote the $(Pf)(s, a) := \mathbb{E}_{s' \sim P(s,a)} f(s')$.
 155 We then denote the *variance induced by one-step transition P and function f* as $(\mathbb{V}_P f)(s, a) :=$
 156 $(Pf^2)(s, a) - (Pf(s, a))^2$ which is equal to $\mathbb{E}_{s' \sim P(s,a)} f^2(s') - (\mathbb{E}_{s' \sim P(s,a)} f(s'))^2$.

157
 158 **Assumptions.** We make the realizability assumption that $P^* \in \mathcal{P}$. We assume that the rewards
 159 are normalized such that $r(\tau) \in [0, 1]$ for any trajectory $\tau := \{s_0, a_0, \dots, s_{H-1}, a_{H-1}\}$ where $r(\tau)$ is

160
 161 ¹For simplicity, we assume initial state s_0 is fixed and known. Our analysis can be easily extended to a
 setting where the initial state is sampled from an unknown fixed distribution.

short for $\sum_{h=0}^{H-1} r(s_h, a_h)$. Note that this setting is more general than assuming each one-step reward is bounded, i.e., $r(s_h, a_h) \in [0, 1/H]$, and allows to represent the sparse reward setting. Without loss of generalizability, we assume $V_{h;P}^\pi(s) \in [0, 1]$, for all $\pi \in \Pi$, $h \in [0, H]$, $P \in \mathcal{P}$, $s \in \mathcal{S}$.

Online RL. For the online RL setting, we focus on the episodic setting where the learner can interact with the environment for K episodes. At episode k , the learner proposes a policy π^k (based on the past interaction history), executes π^k starting from s_0 to time step $H - 1$. We measure the performance of the online learning via *regret*: $\sum_{k=0}^{K-1} (V^{\pi^*} - V^{\pi^k})$. To achieve meaningful regret bounds, we often need additional structural assumptions on the MDP and the model class \mathcal{P} . We use a ℓ_1 Eluder dimension (Liu et al., 2022) as the structural condition due to its ability to capture non-linear function approximators (formal definition will be given in Section 4).

Offline RL. For the offline RL setting, we assume that we have a pre-collected offline dataset $\mathcal{D} = \{\tau^i\}_{i=1}^K$ which contains K trajectories³. For each trajectory, we allow it to potentially be generated by an adversary, i.e., at step h in trajectory k , (i.e., s_h^k), the adversary can select a_h^k based on all history (the past $k - 1$ trajectories and the steps before h within trajectory k) with a fixed strategy, with the only condition that the state transitions follow the underlying transition dynamics, i.e., $s_{h+1}^k \sim P^*(s_h^k, a_h^k)$. We emphasize that \mathcal{D} is not necessarily generated by some offline trajectory distribution. Given \mathcal{D} , we can split the data into HK many state-action-next state (s, a, s') tuples which we can use to learn the transition. To succeed in offline learning, we typically require the offline dataset to have good coverage over some high-quality comparator policy π^* (formal definition of coverage will be given in Section 5). Our goal here is to learn a policy $\hat{\pi}$ that is as good as π^* , and we are interested in the *performance gap* between $\hat{\pi}$ and π^* , i.e., $V^{\pi^*} - V^{\hat{\pi}}$.

Horizon-free and Second-order Bounds. Our goal is to achieve regret bounds (online RL) or performance gaps (offline RL) that are (nearly) horizon-free, i.e., logarithmical dependence on H . In addition to the horizon-free guarantee, we also want our bounds to scale with respect to the variance of the policies. Denote VaR_π as the variance of trajectory reward, i.e., $\text{VaR}_\pi := \mathbb{E}_{\tau \sim \pi} (r(\tau) - \mathbb{E}_{\tau \sim \pi} r(\tau))^2$. Second-order bounds in offline RL scales with VaR_{π^*} – the variance of the comparator policy. Second-order regret bound in online setting scales with respect to $\sqrt{\sum_k \text{VaR}_{\pi^k}}$ instead of \sqrt{K} . Note that in the worst case, $\sqrt{\sum_k \text{VaR}_{\pi^k}}$ scales in the order of \sqrt{K} , but can be much smaller in benign cases such as nearly deterministic MDPs. We also note that second-order regret bound immediately implies first-order regret bound in the reward maximization setting, which scales in the order $\sqrt{KV^{\pi^*}}$ instead of just \sqrt{K} . The first order regret bound $\sqrt{KV^{\pi^*}}$ is never worse than \sqrt{K} since $V^{\pi^*} \leq 1$. Thus, by achieving a second-order regret bound, our algorithm immediately achieves a first-order regret bound.

Additional notations. Given two distributions $p \in \Delta(\mathcal{X})$ and $q \in \Delta(\mathcal{X})$, we denote the triangle discrimination $D_\Delta(p \parallel q) = \sum_{x \in \mathcal{X}} \frac{(p(x) - q(x))^2}{p(x) + q(x)}$, and squared Hellinger distance $\mathbb{H}^2(p \parallel q) = \frac{1}{2} \sum_{x \in \mathcal{X}} (\sqrt{q(x)} - \sqrt{p(x)})^2$ (we replace sum via integral when \mathcal{X} is continuous and p and q are pdfs). Note that D_Δ and \mathbb{H}^2 are equivalent up to universal constants. We will frequently use the following key lemma in (Wang et al., 2024) to control the difference between means of two distributions.

Lemma 1 (Lemma 4.3 in Wang et al. (2024)). *For two distributions $f \in \Delta([0, 1])$ and $g \in \Delta([0, 1])$:*

$$|\mathbb{E}_{x \sim f}[x] - \mathbb{E}_{x \sim g}[x]| \leq 4\sqrt{\text{VaR}_f \cdot D_\Delta(f \parallel g)} + 5D_\Delta(f \parallel g). \quad (1)$$

where $\text{VaR}_f := \mathbb{E}_{x \sim f}(x - \mathbb{E}_{x \sim f}[x])^2$ denotes the variance of the distribution f .

The lemma plays a key role in achieving second-order bounds (Wang et al., 2024). The intuition is the means of the two distributions can be closer if one of the distributions has a small variance. A more naive way of bounding the difference in means is $|\mathbb{E}_{x \sim f}[x] - \mathbb{E}_{x \sim g}[x]| \leq (\max_{x \in \mathcal{X}} |x|) \|f - g\|_1 \lesssim (\max_{x \in \mathcal{X}} |x|) \mathbb{H}(f \parallel g) \lesssim (\max_{x \in \mathcal{X}} |x|) \sqrt{D_\Delta(f \parallel g)}$. Such an approach would have to pay the maximum range $\max_{x \in \mathcal{X}} |x|$ and thus can not leverage the variance VaR_f . In the next sections, we show this lemma plays an important role in achieving horizon-free and second-order bounds.

² $r(\tau) \in [0, 1]$ implies $V_{h;P^*}^\pi(s) \in [0, 1]$. If we do not assume $V_{h;P}^\pi(s) \in [0, 1]$ for all $P \in \mathcal{P}$, we can simply add a filtering step in the algorithm to only choose π, P with $V_{h;P}^\pi(s_0) \in [0, 1]$ to get the same guarantees.

³Our algorithms and analysis can also apply to the case where the dataset only consists of transitions, as they rely solely on transition-level information from the offline dataset

Algorithm 1 Optimistic Model-based RL (O-MBRL)

-
- 1: **Input:** model class \mathcal{P} , confidence parameter $\delta \in (0, 1)$, threshold β .
 - 2: Initialize π^0 , initialize dataset $\mathcal{D} = \emptyset$.
 - 3: **for** $k = 0 \rightarrow K - 1$ **do**
 - 4: Collect a trajectory $\tau = \{s_0, a_0, \dots, s_{H-1}, a_{H-1}\}$ from π^k , split it into tuples of $\{s, a, s'\}$ and add to \mathcal{D} .
 - 5: Construct a version space $\widehat{\mathcal{P}}^k$:

$$\widehat{\mathcal{P}}^k = \left\{ P \in \mathcal{P} : \sum_{s,a,s' \in \mathcal{D}} \log P(s'_i | s_i, a_i) \geq \max_{\tilde{P} \in \mathcal{P}} \sum_{s,a,s' \in \mathcal{D}} \log \tilde{P}(s'_i | s_i, a_i) - \beta \right\}.$$
 - 6: Set $(\pi^k, \widehat{P}^k) \leftarrow \operatorname{argmax}_{\pi \in \Pi, P \in \widehat{\mathcal{P}}^k} V_{0;P}^{\pi}(s_0)$.
 - 7: **end for**
-

4 ONLINE SETTING

In this section, we study the online setting. We present the optimistic model-based RL algorithm (O-MBRL) in Algorithm 1. The algorithm starts from scratch, and iteratively maintains a version space $\widehat{\mathcal{P}}^k$ of the model class using the historical data collected so far. Again the version space is designed such that for all $k \in [0, K - 1]$, we have $P^* \in \widehat{\mathcal{P}}^k$ with high probability. The policy π^k in this case is computed via the optimism principle, i.e., it selects π^k and \widehat{P}^k such that $V_{\widehat{P}^k}^{\pi^k} \geq V^{\pi^*}$.

Note that the algorithm design in Algorithm 1 is not new and in fact is quite standard in the model-based RL literature. For instance, Sun et al. (2019) presented a similar style of algorithm except that they use a min-max GAN style objective for learning models. Zhan et al. (2022) used MLE oracle with optimism planning for Partially observable systems such as Predictive State Representations (PSRs), and Liu et al. (2023) used them for both partially and fully observable systems. However, their analyses do not give horizon-free and instance-dependent bounds. We show that under the structural condition that captures nonlinear function class with small eluder dimensions, Algorithm 1 achieves horizon-free and second-order bounds. Besides, since second-order regret bound implies first-order bound (Wang et al., 2024), our result immediately implies a first-order bound as well.

We first introduce the ℓ_p Eluder dimension as follows.

Definition 1 (ℓ_p Eluder Dimension). $DE_p(\Psi, \mathcal{X}, \epsilon)$ is the eluder dimension for \mathcal{X} with function class Ψ , when the longest ϵ -independent sequence $x^1, \dots, x^L \subseteq \mathcal{X}$ enjoys the length less than $DE_p(\Psi, \mathcal{X}, \epsilon)$. We say that a sequence $x^1, \dots, x^L \subseteq \mathcal{X}$ is ϵ -independent if there exists $g \in \Psi$ such that for all $t \in [L]$, $\sum_{l=1}^{t-1} |g(x^l)|^p \leq \epsilon^p$ and $|g(x^t)| > \epsilon$.

We work with the ℓ_1 Eluder dimension $DE_1(\Psi, \mathcal{S} \times \mathcal{A}, \epsilon)$ with the function class Ψ specified as:

$$\Psi = \{(s, a) \mapsto \mathbb{H}^2(P^*(s, a) \parallel P(s, a)) : P \in \mathcal{P}\}.$$

Remark 1. The ℓ_1 Eluder dimension has been used in previous works such as Liu et al. (2022). We have the following corollary to demonstrate that the ℓ_1 dimension generalizes the original ℓ_2 dimension of Russo & Van Roy (2013), it can capture tabular, linear, and generalized linear models.

Lemma 2 (Proposition 19 in (Liu et al., 2022)). For any $\Psi, \mathcal{X}, \epsilon > 0$, $DE_1(\Psi, \mathcal{X}, \epsilon) \leq DE_2(\Psi, \mathcal{X}, \epsilon)$.

We are ready to present our main theorem for the online RL setting with finite class \mathcal{P} .

Theorem 1 (Main theorem for online setting with finite \mathcal{P}). For any $\delta \in (0, 1)$, let $\beta = 4 \log \left(\frac{K|\mathcal{P}|}{\delta} \right)$, with probability at least $1 - \delta$, Algorithm 1 achieves the following regret bound:

$$\begin{aligned} \sum_{k=0}^{K-1} (V^{\pi^*} - V^{\pi^k}) &\leq O \left(\sqrt{\sum_{k=0}^{K-1} \operatorname{VaR}_{\pi^k} \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(KH|\mathcal{P}|/\delta) \log(KH)} \right. \\ &\quad \left. + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(KH|\mathcal{P}|/\delta) \log(KH) \right). \end{aligned} \quad (2)$$

The above theorem indicates the standard and simple O-MBRL algorithm is already enough to achieve horizon-free and second-order regret bounds: our bound does not have explicit polynomial dependences on horizon H , the leading term scales with $\sqrt{\sum_k \operatorname{VaR}_{\pi^k}}$ instead of the typical \sqrt{K} .

We have the following result about the first-order regret bound.

Corollary 1 (Horizon-free and First-order regret bound). *Let $\beta = 4 \log \left(\frac{K|\mathcal{P}|}{\delta} \right)$, with probability at least $1 - \delta$, Algorithm 1 achieves the following regret bound:*

$$\sum_{k=0}^{K-1} V^{\pi^*} - V^{\pi^k} \leq O\left(\sqrt{KV^{\pi^*} \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(KH|\mathcal{P}|/\delta) \log(KH)}\right) \\ + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(KH|\mathcal{P}|/\delta) \log(KH).$$

Proof. Note that $\text{VaR}_{\pi} \leq V^{\pi} \leq V^{\pi^*}$ where the first inequality is because the trajectory-wise reward is bounded in $[0, 1]$. Therefore, combining with Theorem 1, we directly obtain the first-order result. \square

Note that the above bound scales with respect to $\sqrt{KV^{\pi^*}}$ instead of just \sqrt{K} . Since $V^{\pi^*} \leq 1$, this bound improves the worst-case regret bound when the optimal policy has total reward less than one.⁴

Faster rates for deterministic transitions. When the underlying MDP has deterministic transitions, we can achieve a smaller regret bound that only depends on the number of episodes logarithmically.

Corollary 2 (log K regret bound with deterministic transitions). *When the transition dynamics of the MDP are deterministic, setting $\beta = 4 \log \left(\frac{K|\mathcal{P}|}{\delta} \right)$, w.p. at least $1 - \delta$, Algorithm 1 achieves:*

$$\sum_{k=0}^{K-1} V^{\pi^*} - V^{\pi^k} \leq O\left(DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(KH|\mathcal{P}|/\delta) \log(KH)\right).$$

Extension to infinite class \mathcal{P} . For infinite model class \mathcal{P} , we have a similar result. First, we define the bracketing number of an infinite model class as follows.

Definition 2 (Bracketing Number (Geer, 2000)). *Let \mathcal{G} be a set of functions mapping $\mathcal{X} \rightarrow \mathbb{R}$. Given two functions l, u such that $l(x) \leq u(x)$ for all $x \in \mathcal{X}$, the bracket $[l, u]$ is the set of functions $g \in \mathcal{G}$ such that $l(x) \leq g(x) \leq u(x)$ for all $x \in \mathcal{X}$. We call $[l, u]$ an ϵ -bracket if $\|u - l\| \leq \epsilon$. Then, the ϵ -bracketing number of \mathcal{G} with respect to $\|\cdot\|$, denoted by $\mathcal{N}_{[]}(\epsilon, \mathcal{G}, \|\cdot\|)$ is the minimum number of ϵ -brackets needed to cover \mathcal{G} .*

As an example, according to Appendix C.2 of Uehara & Sun (2021), the bracketing number of Linear Mixture MDPs is $\mathcal{N}_{[]}(\epsilon, \mathcal{P}, \|\cdot\|) = O(1/\epsilon)^d$. We use the bracketing number of \mathcal{P} to denote the complexity of the model class, similar to $|\mathcal{P}|$ in the finite class case. Next, we propose a corollary to characterize the regret with an infinite model class.

Corollary 3 (Regret bound for Algorithm 1 with infinite model class \mathcal{P}). *When the model class \mathcal{P} is infinite, let $\beta = 7 \log(K\mathcal{N}_{[]}((KH|\mathcal{S}|)^{-1}, \mathcal{P}, \|\cdot\|_{\infty})/\delta)$, with probability at least $1 - \delta$, Algorithm 1 achieves the following regret bound:*

$$\sum_{k=0}^{K-1} V^{\pi^*} - V^{\pi^k} \leq O\left(DE_1\left(\Psi, \mathcal{S} \times \mathcal{A}, \frac{1}{KH}\right) \log\left(\frac{KH\mathcal{N}_{[]}((KH|\mathcal{S}|)^{-1}, \mathcal{P}, \|\cdot\|_{\infty})}{\delta}\right) \log(KH)\right) \\ + \sqrt{\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \cdot DE_1\left(\Psi, \mathcal{S} \times \mathcal{A}, \frac{1}{KH}\right) \log\left(\frac{KH\mathcal{N}_{[]}((KH|\mathcal{S}|)^{-1}, \mathcal{P}, \|\cdot\|_{\infty})}{\delta}\right) \log(KH)},$$

where $\mathcal{N}_{[]}((KH|\mathcal{S}|)^{-1}, \mathcal{P}, \|\cdot\|_{\infty})$ is the bracketing number (defined in Definition 2).

A specific example of the infinite model class is the tabular MDP, where \mathcal{P} is the collection of all the conditional distributions over $\mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$. By Corollary 3, we also have a new regret bound for MBRL under the tabular MDP setting, which is nearly horizon-free and second-order.

Example 1 (Tabular MDPs). *When specializing to tabular MDPs, use the fact that tabular MDP has ℓ_2 Eluder dimension being at most $|\mathcal{S}||\mathcal{A}|$ (Section D.1 in Russo & Van Roy (2013)), ℓ_1 dimension is upper bounded by ℓ_2 dimension (Lemma 2), and use the standard ϵ -net argument to show that*

⁴Typically a first-order regret bound makes more sense in the cost minimization setting instead of reward maximization setting. We believe that our results are transferable to the cost-minimization setting.

$\mathcal{N}_{[]}(\epsilon, \mathcal{P}, \|\cdot\|_\infty)$ is upper-bounded by $(c/\epsilon)^{|\mathcal{S}|^2|\mathcal{A}|}$ (e.g., see Uehara & Sun (2021)), we can show that Algorithm 1 achieves the following regret bound for tabular MDP: with probability at least $1 - \delta$,

$$\sum_k V^{\pi^*} - V^{\pi^k} \leq O\left(|\mathcal{S}|^{1.5}|\mathcal{A}|\sqrt{\sum_k \text{VaR}_{\pi^k} \cdot \log\left(\frac{KH|\mathcal{S}|}{\delta}\right)} \log(KH) + |\mathcal{S}|^3|\mathcal{A}|^2 \log\left(\frac{KH|\mathcal{S}|}{\delta}\right) \log(KH)\right).$$

In summary, we have shown that a simple MLE-based MBRL algorithm is enough to achieve nearly horizon-free and second-order regret bounds under non-linear function approximation.

4.1 PROOF SKETCH OF THEOREM 1

Now we are ready to provide a proof sketch of Theorem 1 with the full proof deferred to Appendix D.1. For ease of presentation, we use d_{RL} to denote $\text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH)$, and ignore some log terms.

Overall, our analysis follows the general framework of optimism in the face of uncertainty, but with (1) careful analysis in leveraging the MLE generalization bound and (2) more refined proof in the training-to-testing distribution transfer via Eluder dimension.

By standard MLE analysis, we can show w.p. $1 - \delta$, for all $k \in [K - 1]$, we have $P^* \in \widehat{\mathcal{P}}^k$, and

$$\sum_{i=0}^{k-1} \sum_{h=0}^{H-1} \mathbb{H}^2(P^*(s_h^i, a_h^i) \| \widehat{P}^k(s_h^i, a_h^i)) \leq O(\log(K|\mathcal{P}|/\delta)). \quad (3)$$

From here, trivially applying training-to-testing distribution transfer via the Eluder dimension as previous works (e.g., Wang et al. (2024)) would cause poly-dependence on H . With new techniques detailed in Appendix B, which is one of our technical contributions and may be of independent interest, we can get: there exists a set $\mathcal{K} \subseteq [K - 1]$ such that $|\mathcal{K}| \leq O(d_{\text{RL}} \log(K|\mathcal{P}|/\delta))$, and

$$\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_h \mathbb{H}^2(P^*(s_h^k, a_h^k) \| \widehat{P}^k(s_h^k, a_h^k)) \leq O(d_{\text{RL}} \cdot \log(K|\mathcal{P}|/\delta) \log(KH)). \quad (4)$$

Recall that $(\pi^k, \widehat{P}^k) \leftarrow \arg\max_{\pi \in \Pi, P \in \widehat{\mathcal{P}}^k} V_{0;P}^\pi(s_0)$, with the above realization guarantee $P^* \in \widehat{\mathcal{P}}^k$, we can get the following optimism guarantee: $V_{0;P^*}^* \leq \max_{\pi \in \Pi, P \in \widehat{\mathcal{P}}^k} V_{0;P}^\pi = V_{0;\widehat{P}^k}^{\pi^k}$.

At this stage, one straight-forward way to proceed is to use the standard simulation lemma (Lemma 7):

$$\sum_{k=0}^{K-1} V_{0;\widehat{P}^k}^{\pi^k} - V_{0;P^*}^* \leq \sum_{k=0}^{K-1} \sum_{h=0}^{H-1} \mathbb{E}_{s, a \sim d_h^{\pi^k}} \left[\left| \mathbb{E}_{s' \sim P^*(s, a)} V_{h+1; \widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim \widehat{P}^k(s, a)} V_{h+1; \widehat{P}^k}^{\pi^k}(s') \right| \right]. \quad (5)$$

However, from here, if we naively bound each term on the RHS via $\mathbb{E}_{s, a \sim d_h^{\pi^k}} \|P^*(s, a) - \widehat{P}^k(s, a)\|_1$, which is what previous works such as Uehara & Sun (2021) did exactly, we would end up paying a linear horizon dependence H due to the summation over H on the RHS the above expression. Given the mean-to-variance lemma (Lemma 1), we may consider using it to bound the difference between two means $\mathbb{E}_{s' \sim P^*(s, a)} V_{h+1; \widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim \widehat{P}^k(s, a)} V_{h+1; \widehat{P}^k}^{\pi^k}(s')$. This still can not work if we start from here, because we would eventually get $\sum_k \sum_h \mathbb{E}_{s, a \sim d_h^{\pi^k}} [\mathbb{H}^2(P^*(s, a) \| \widehat{P}^k(s, a))]$ terms, which can not be further upper bounded easily with the MLE generalization guarantee.

To achieve horizon-free and second-order bounds, we need a novel and more careful analysis.

First, we carefully decompose and upper bound the regret in $\tilde{\mathcal{K}} := [K - 1] \setminus \mathcal{K}$ w.h.p. as follows using Bernstein's inequality (for regret in \mathcal{K} we simply upper bound it by $|\mathcal{K}|$)

$$\begin{aligned} & \sum_{k \in \tilde{\mathcal{K}}} \left(V_{0;\widehat{P}^k}^{\pi^k}(s_h^k) - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \right) + \sum_{k \in \tilde{\mathcal{K}}} \left(\sum_{h=0}^{H-1} r(s_h^k, a_h^k) - V_{0;P^*}^* \right) \lesssim \sqrt{\sum_{k \in \tilde{\mathcal{K}}} \sum_h (\mathbb{V}_{P^*} V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k)} \\ & + \sum_{k \in \tilde{\mathcal{K}}} \sum_h \left| \mathbb{E}_{s' \sim \widehat{P}^k(s_h^k, a_h^k)} V_{h+1; \widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*(s_h^k, a_h^k)} V_{h+1; \widehat{P}^k}^{\pi^k}(s') \right| + \sqrt{\sum_k \text{VaR}_{\pi^k} \log(1/\delta)}. \quad (6) \end{aligned}$$

Then, we bound the difference of two means $\mathbb{E}_{s' \sim \widehat{P}^k}(s_h^k, a_h^k) V_{h+1; \widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*}(s_h^k, a_h^k) V_{h+1; \widehat{P}^k}^{\pi^k}(s')$ using variances and the triangle discrimination (see Lemma 1 for more details), together with the fact that $D_\Delta \leq 4\mathbb{H}^2$, and information processing inequality on the squared Hellinger distance, we have

$$\begin{aligned} & \left| \mathbb{E}_{s' \sim \widehat{P}^k}(s_h^k, a_h^k) V_{h+1; \widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*}(s_h^k, a_h^k) V_{h+1; \widehat{P}^k}^{\pi^k}(s') \right| \\ & \leq O\left(\sqrt{(\mathbb{V}_{P^*} V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k)} D_\Delta \left(V_{h+1; \widehat{P}^k}^{\pi^k}(s' \sim P^*(s_h^k, a_h^k)) \parallel V_{h+1; \widehat{P}^k}^{\pi^k}(s' \sim \widehat{P}^k(s_h^k, a_h^k)) \right) \right. \\ & \quad \left. + D_\Delta \left(V_{h+1; \widehat{P}^k}^{\pi^k}(s' \sim P^*(s_h^k, a_h^k)) \parallel V_{h+1; \widehat{P}^k}^{\pi^k}(s' \sim \widehat{P}^k(s_h^k, a_h^k)) \right) \right) \\ & \leq O\left(\sqrt{(\mathbb{V}_{P^*} V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k)} \mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{P}^k(s_h^k, a_h^k) \right) + \mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{P}^k(s_h^k, a_h^k) \right) \right) \end{aligned}$$

where we denote $V_{h+1; \widehat{P}^k}^{\pi^*}(s' \sim P^*(s, a))$ as the distribution of the random variable $V_{h+1; \widehat{P}^k}^{\pi^*}(s')$ with $s' \sim P^*(s, a)$. This is the key lemma used by Wang et al. (2024) to show distributional RL can achieve second-order bounds. We show that this is also crucial for achieving a horizon-free bound.

Then, summing up over k, h , with Cauchy-Schwartz and the MLE generalization bound via Eluder dimension in Equation 4, we have

$$\begin{aligned} & \sum_{k \in \mathcal{K}} \sum_h \left| \mathbb{E}_{s' \sim \widehat{P}^k}(s_h^k, a_h^k) V_{h+1; \widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*}(s_h^k, a_h^k) V_{h+1; \widehat{P}^k}^{\pi^k}(s') \right| \leq O\left(\sum_{k \in \mathcal{K}} \sum_h \mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{P}^k(s_h^k, a_h^k) \right) \right. \\ & \quad \left. + \sqrt{\sum_{k \in \mathcal{K}} \sum_h (\mathbb{V}_{P^*} V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k)} \sum_{k \in \mathcal{K}} \sum_h \mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{P}^k(s_h^k, a_h^k) \right) \right) \\ & \leq O\left(\sqrt{\sum_{k \in \mathcal{K}} \sum_h (\mathbb{V}_{P^*} V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k)} d_{\text{RL}} \log(K|\mathcal{P}|/\delta) \log(KH) + d_{\text{RL}} \log(K|\mathcal{P}|/\delta) \log(KH) \right). \quad (7) \end{aligned}$$

Note that we have $(\mathbb{V}_{P^*} V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k)$ depending on \widehat{P}^k . To get a second-order bound, we need to convert it to the variance under ground truth transition P^* , and we want to do it without incurring any H dependence. *This is another key difference from Wang et al. (2024).*

We aim to replace $(\mathbb{V}_{P^*} V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k)$ by $(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s_h^k, a_h^k)$ which is the variance under P^* (recall that V^π is the value function of π under P^*), and we want to control the difference $(\mathbb{V}_{P^*} (V_{h+1; \widehat{P}^k}^{\pi^k} - V_{h+1}^{\pi^k}))(s_h^k, a_h^k)$. To do so, we need to bound the 2^m moment of the difference $V_{h+1; \widehat{P}^k}^{\pi^k} - V_{h+1}^{\pi^k}$ following the strategy in Zhang et al. (2021a); Zhou & Gu (2022); Zhao et al. (2023b). Let us define the following terms:

$$\begin{aligned} A & := \sum_{k \in \mathcal{K}} \sum_h \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k) \right], C_m := \sum_{k \in \mathcal{K}} \sum_h \left[(\mathbb{V}_{P^*} (V_{h+1; \widehat{P}^k}^{\pi^k} - V_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k) \right], \\ B & := \sum_{k \in \mathcal{K}} \sum_h \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s_h^k, a_h^k) \right], G := \sqrt{A \cdot d_{\text{RL}} \log\left(\frac{K|\mathcal{P}|}{\delta}\right) \log(KH) + d_{\text{RL}} \log\left(\frac{K|\mathcal{P}|}{\delta}\right) \log(KH)}. \end{aligned}$$

With the fact $\mathbb{V}_{P^*}(a+b) \leq 2\mathbb{V}_{P^*}(a) + 2\mathbb{V}_{P^*}(b)$ we have $A \leq 2B + 2C_0$. For C_m , we prove that w.h.p. it has the recursive form $C_m \lesssim 2^m G + \sqrt{\log(1/\delta) C_{m+1} + \log(1/\delta)}$, during which process we also leverage the above Equation 7 and some careful analysis (detailed in Appendix D.1). Then, with the recursion lemma (Lemma 11), we can get $C_0 \lesssim G$, which further gives us

$$A \lesssim B + d_{\text{RL}} \log\left(\frac{K|\mathcal{P}|}{\delta}\right) \log(KH) + \sqrt{A \cdot d_{\text{RL}} \log\left(\frac{K|\mathcal{P}|}{\delta}\right) \log(KH)} \leq O\left(B + d_{\text{RL}} \log\left(\frac{K|\mathcal{P}|}{\delta}\right) \log(KH)\right),$$

where in the last step we use the fact $x \leq 2a + b^2$ if $x \leq a + b\sqrt{x}$. Finally, we note that $B \leq O(\sum_k \text{VaR}_{\pi^k} + \log(1/\delta))$ w.h.p.. Plugging the upper bound of A back into Equation 7 and then to Equation 6, we conclude the proof.

5 OFFLINE SETTING

For the offline setting, we directly analyze the Constrained Pessimism Policy Optimization (CPPO-LR) algorithm (Algorithm 2) proposed by Uehara & Sun (2021). We first explain the algorithm and then present its performance gap guarantee in finding the comparator policy π^* .

Algorithm 2 (Uehara & Sun (2021)) Constrained Pessimistic Policy Optimization with Likelihood-Ratio based constraints (CPPO-LR)

1: **Input:** dataset $\mathcal{D} = \{s, a, s'\}$, model class \mathcal{P} , policy class Π , confidence parameter $\delta \in (0, 1)$, threshold β .

2: Calculate the confidence set based on the offline dataset:

$$\widehat{\mathcal{P}} = \left\{ P \in \mathcal{P} : \sum_{i=1}^n \log P(s'_i | s_i, a_i) \geq \max_{\tilde{P} \in \mathcal{P}} \sum_{i=1}^n \log \tilde{P}(s'_i | s_i, a_i) - \beta \right\}.$$

3: **Output:** $\hat{\pi} \leftarrow \operatorname{argmax}_{\pi \in \Pi} \min_{P \in \widehat{\mathcal{P}}} V_{0;P}^{\pi}(s_0)$.

Algorithm 2 splits the offline trajectory data that contains K trajectories into a dataset of (s, a, s') tuples (note that in total we have $n := KH$ many tuples) which is used to perform maximum likelihood estimation $\max_{P \in \mathcal{P}} \sum_{i=1}^n \log P(s'_i | s_i, a_i)$. It then builds a version space $\widehat{\mathcal{P}}$ which contains models $P \in \mathcal{P}$ whose log data likelihood is not below by too much than that of the MLE estimator. The threshold for the version space is constructed so that with high probability, $P^* \in \widehat{\mathcal{P}}$. Once we build a version space, we perform pessimistic planning to compute $\hat{\pi}$.

We first define the single policy coverage condition as follows.

Definition 3 (Single policy coverage). *Given any comparator policy π^* , denote the data-dependent single policy concentrability coefficient $C_{\mathcal{D}}^{\pi^*}$ as follows:*

$$C_{\mathcal{D}}^{\pi^*} := \max_{h, P \in \mathcal{P}} \frac{\mathbb{E}_{s, a \sim d_h^{\pi^*}} \mathbb{H}^2(P(s, a) \| P^*(s, a))}{1/K \sum_{k=1}^K \mathbb{H}^2(P(s_h^k, a_h^k) \| P^*(s_h^k, a_h^k))}.$$

We assume w.p. at least $1 - \delta$ over the randomness of the generation of \mathcal{D} , we have $C_{\mathcal{D}}^{\pi^*} \leq C^{\pi^*}$.

The existence of C^{π^*} is certainly an assumption. We now give an example in the tabular MDP where we show that if the data is generated from some fixed behavior policy π^b which has non-trivial probability of visiting every state-action pair, then we can show the existence of C^{π^*} .

Example 2 (Tabular MDP with good behavior policy coverage). *If the K trajectories are collected i.i.d. with a fixed behavior policy π^b , and $d_h^{\pi^b}(s, a) \geq \rho_{\min}$, $\forall s, a, h$ (similar to Ren et al. (2021)), then we have: if K is large enough, i.e., $K \geq 2 \log(|\mathcal{S}||\mathcal{A}|H) / \rho_{\min}^2$, w.p. at least $1 - \delta$, $C_{\mathcal{D}}^{\pi^*} \leq 2 / \rho_{\min}$.*

Our coverage definition (Definition 3) shares similar spirits as the one in Ye et al. (2024). It reflects how well the state-action samples in the offline dataset \mathcal{D} cover the state-action pairs induced by the comparator policy π^* . It is different from the coverage definition in Uehara & Sun (2021) in which the denominator is $\mathbb{E}_{s, a \sim d_h^{\pi^b}} \mathbb{H}^2(P(s, a) \| P^*(s, a))$ where π^b is the fixed behavior policy used to collect \mathcal{D} . This definition does not apply in our setting since \mathcal{D} is not necessarily generated by some underlying fixed behavior policy. On the other hand, our horizon-free result does not hold in the setting of Uehara & Sun (2021) where \mathcal{D} is collected with a fixed behavior policy π^b with the concentrability coefficient defined in their way. We leave the derivation of horizon-free results in the setting from Uehara & Sun (2021) as a future work.

Now we are ready to present the main theorem of Algorithm 2 for finite \mathcal{P} , which provides a tighter performance gap than that by Uehara & Sun (2021).

Theorem 2 (Performance gap of Algorithm 2 with finite \mathcal{P}). *For any $\delta \in (0, 1)$, let $\beta = 4 \log(|\mathcal{P}|/\delta)$, w.p. at least $1 - \delta$, Algorithm 2 learns a policy $\hat{\pi}$ that enjoys the following performance gap with respect to any comparator policy π^* :*

$$V^{\pi^*} - V^{\hat{\pi}} \leq O\left(\sqrt{C^{\pi^*} \operatorname{VaR}_{\pi^*} \log(|\mathcal{P}|/\delta)/K} + C^{\pi^*} \log(|\mathcal{P}|/\delta)/K\right).$$

Comparing to the theorem (Theorem 2) of CPPO-LR from Uehara & Sun (2021), our bound has two improvements. First, our bound is horizon-free (not even any $\log(H)$ dependence), while the bound in Uehara & Sun (2021) has $\operatorname{poly}(H)$ dependence. Second, our bound scales with $\operatorname{VaR}_{\pi^*} \in [0, 1]$, which can be small when $\operatorname{VaR}_{\pi^*} \ll 1$. For deterministic system and policy π^* , we have $\operatorname{VaR}_{\pi^*} = 0$ which means the sample complexity now scales at a faster rate C^{π^*}/K . The proof is in Appendix E.1.

We show that the same algorithm can achieve $1/K$ rate when P^* is deterministic (but rewards could be random, and the algorithm does not need to know the condition that P^* is deterministic).

Corollary 4 (C^{π^*}/K performance gap of Algorithm 2 with deterministic transitions). *When the ground truth transition P^* of the MDP is deterministic, for any $\delta \in (0, 1)$, let $\beta = 4 \log(|\mathcal{P}|/\delta)$, w.p. at least $1 - \delta$, Algorithm 2 learns a policy $\hat{\pi}$ that enjoys the following performance gap with respect to any comparator policy π^* :*

$$V^{\pi^*} - V^{\hat{\pi}} \leq O\left(C^{\pi^*} \log(|\mathcal{P}|/\delta)/K\right).$$

For infinite model class \mathcal{P} , we have a similar result in the following corollary.

Corollary 5 (Performance gap of Algorithm 2 with infinite model class \mathcal{P}). *When the model class \mathcal{P} is infinite, for any $\delta \in (0, 1)$, let $\beta = 7 \log(\mathcal{N}_{[]}((KH|\mathcal{S})^{-1}, \mathcal{P}, \|\cdot\|_{\infty})/\delta)$, w.p. at least $1 - \delta$, Algorithm 2 learns a policy $\hat{\pi}$ that enjoys the following PAC bound w.r.t. any comparator policy π^* :*

$$V^{\pi^*} - V^{\hat{\pi}} \leq O\left(\sqrt{\frac{C^{\pi^*} \text{VaR}_{\pi^*} \log(\mathcal{N}_{[]}((KH|\mathcal{S})^{-1}, \mathcal{P}, \|\cdot\|_{\infty})/\delta)}{K}} + \frac{C^{\pi^*} \log(\mathcal{N}_{[]}((KH|\mathcal{S})^{-1}, \mathcal{P}, \|\cdot\|_{\infty})/\delta)}{K}\right),$$

where $\mathcal{N}_{[]}((KH|\mathcal{S})^{-1}, \mathcal{P}, \|\cdot\|_{\infty})$ is the bracketing number defined in Definition 2.

Our next example gives the explicit performance gap bound for tabular MDPs.

Example 3 (Tabular MDPs). *For tabular MDPs, we have $\mathcal{N}_{[]}(\epsilon, \mathcal{P}, \|\cdot\|_{\infty})$ upper-bounded by $(c/\epsilon)^{|\mathcal{S}|^2|\mathcal{A}|}$ (e.g., see Uehara & Sun (2021)). Then with probability at least $1 - \delta$, let $\beta = 7 \log(\mathcal{N}_{[]}((KH|\mathcal{S})^{-1}, \mathcal{P}, \|\cdot\|_{\infty})/\delta)$, Algorithm 2 learns a policy $\hat{\pi}$ satisfying the following performance gap with respect to any comparator policy π^* :*

$$V^{\pi^*} - V^{\hat{\pi}} \leq O\left(|\mathcal{S}| \sqrt{|\mathcal{A}| C^{\pi^*} \text{VaR}_{\pi^*} \log(KH|\mathcal{S}|/\delta)}/K + |\mathcal{S}|^2 |\mathcal{A}| C^{\pi^*} \log(KH|\mathcal{S}|/\delta)/K\right),$$

The closest result to us is from Ren et al. (2021), which analyzes the MBRL for tabular MDPs and obtains a performance gap $\tilde{O}\left(\sqrt{\frac{1}{Kd_m} + \frac{|\mathcal{S}|}{Kd_m}}\right)$, where d_m is the minimum visiting probability for the behavior policy to visit each state and action. Note that their result is not instance-dependent, which makes their gap only $\tilde{O}(1/\sqrt{K})$ even when the environment is deterministic and π^* is deterministic. In a sharp contrast, our analysis shows a better $\tilde{O}(1/K)$ gap under the deterministic environment. Our result would still have the $\log H$ dependence, and we leave getting rid of this logarithmic dependence on the horizon H as an open problem.

6 CONCLUSION

In this work, we presented a minimalist approach for achieving nearly horizon-free and second-order bounds in online and offline RL: simply train transition models via Maximum Likelihood Estimation followed by optimistic or pessimistic planning, depending on whether we operate in the online or offline learning mode. Our horizon-free bounds for function approximation look quite similar to the bounds in Contextual bandits, indicating that the need for long-horizon planning does not make RL harder than CB from a statistical perspective.

Our work has some limitations. First, though our results can achieve *completely horizon free* (not even $\log H$ dependence) for offline RL with finite model classes, in other settings we have $\log H$ dependence. We conjecture that the $\log(H)$ can be eliminated by using more careful analyses in dealing with the learned models' extrapolation errors via eluder dimensions, and using techniques such as peeling/chaining (Dudley, 1978; Zhang, 2006) when tackling infinite model classes. We leave this as an important future direction. Second, our model-based framework cannot capture problems that need to be solved via model-free approaches such as linear MDPs (Jin et al., 2020). An interesting future work is to see if we can extend our analysis to incorporate the function classes with small distributional Eluder dimensions, and if we can develop model-free approaches that can achieve horizon-free and instance-dependent bounds for RL with general function approximation. Finally, the algorithms studied in this work are not computationally tractable. This is due to the need of performing optimism/pessimism planning for exploration. Deriving computationally tractable RL algorithms for the rich function approximation setting is a long-standing question.

REFERENCES

- 540
541
542 Eric W Aboaf, Steven Mark Drucker, and Christopher G Atkeson. Task-level robot learning: Juggling
543 a tennis ball more accurately. In *Proceedings, 1989 International Conference on Robotics and*
544 *Automation*, pp. 1290–1295. IEEE, 1989.
- 545 Alekh Agarwal, Nan Jiang, Sham M Kakade, and Wen Sun. Reinforcement learning: Theory and
546 algorithms. 2019.
- 547 Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. Flambe: Structural complexity
548 and representation learning of low rank mdps. *Advances in neural information processing systems*,
549 33:20095–20107, 2020.
- 550
551 Nicolo Cesa-Bianchi, Yishay Mansour, and Gilles Stoltz. Improved second-order bounds for predic-
552 tion with expert advice. *Machine Learning*, 66:321–352, 2007.
- 553 Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement
554 learning in a handful of trials using probabilistic dynamics models. *Advances in neural information*
555 *processing systems*, 31, 2018.
- 556
557 Marc Deisenroth, Carl Rasmussen, and Dieter Fox. Learning to control a low-cost manipulator using
558 data-efficient reinforcement learning. *Robotics: Science and Systems VII*, 2011.
- 559 Qiwei Di, Jiafan He, Dongruo Zhou, and Quanquan Gu. Nearly minimax optimal regret for learning
560 linear mixture stochastic shortest path. In *International Conference on Machine Learning*, pp.
561 7837–7864. PMLR, 2023.
- 562
563 Richard M Dudley. Central limit theorems for empirical measures. *The Annals of Probability*, pp.
564 899–929, 1978.
- 565 Dylan J Foster and Akshay Krishnamurthy. Efficient first-order contextual bandits: Prediction,
566 allocation, and triangular discrimination. *Advances in Neural Information Processing Systems*, 34:
567 18907–18919, 2021.
- 568
569 Dylan J Foster, Sham M Kakade, Jian Qian, and Alexander Rakhlin. The statistical complexity of
570 interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.
- 571
572 Dylan J Foster, Adam Block, and Dipendra Misra. Is behavior cloning all you need? understanding
573 horizon in imitation learning. *arXiv preprint arXiv:2407.15007*, 2024.
- 574 Sara A Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge university press, 2000.
- 575
576 Jiayi Huang, Han Zhong, Liwei Wang, and Lin Yang. Horizon-free and instance-dependent re-
577 gret bounds for reinforcement learning with general function approximation. In *International*
578 *Conference on Artificial Intelligence and Statistics*, pp. 3673–3681. PMLR, 2024.
- 579 Shinji Ito, Shuichi Hirahara, Tasuku Soma, and Yuichi Yoshida. Tight first-and second-order regret
580 bounds for adversarial linear bandits. *Advances in Neural Information Processing Systems*, 33:
581 2028–2038, 2020.
- 582
583 Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient?
584 *Advances in neural information processing systems*, 31, 2018.
- 585
586 Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement
587 learning with linear function approximation. In *Conference on learning theory*, pp. 2137–2143.
PMLR, 2020.
- 588
589 Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad
590 Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, et al. Model-based
591 reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019.
- 592
593 Yeoneung Kim, Insoon Yang, and Kwang-Sung Jun. Improved regret analysis for variance-adaptive
linear bandits and horizon-free linear mixture mdps. *Advances in Neural Information Processing*
Systems, 35:1060–1072, 2022.

- 594 Shengshi Li and Lin Yang. Horizon-free learning for markov decision processes and games: stochas-
595 tically bounded rewards and improved bounds. In *International Conference on Machine Learning*,
596 pp. 20221–20252. PMLR, 2023.
- 597 Yuanzhi Li, Ruosong Wang, and Lin F Yang. Settling the horizon-dependence of sample complexity
598 in reinforcement learning. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer
599 Science (FOCS)*, pp. 965–976. IEEE, 2022.
- 600
- 601 Qinghua Liu, Alan Chung, Csaba Szepesvári, and Chi Jin. When is partially observable reinforcement
602 learning not scary? In *Conference on Learning Theory*, pp. 5175–5220. PMLR, 2022.
- 603
- 604 Qinghua Liu, Praneeth Netrapalli, Csaba Szepesvari, and Chi Jin. Optimistic mle: A generic model-
605 based algorithm for partially observable sequential decision making. In *Proceedings of the 55th
606 Annual ACM Symposium on Theory of Computing*, pp. 363–376, 2023.
- 607 Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalence is efficient for linear quadratic
608 control. *Advances in Neural Information Processing Systems*, 32, 2019.
- 609
- 610 Julia Olkhovskaya, Jack Mayo, Tim van Erven, Gergely Neu, and Chen-Yu Wei. First-and second-
611 order bounds for adversarial linear contextual bandits. *Advances in Neural Information Processing
612 Systems*, 36, 2024.
- 613 Tongzheng Ren, Jialian Li, Bo Dai, Simon S Du, and Sujay Sanghavi. Nearly horizon-free offline
614 reinforcement learning. *Advances in neural information processing systems*, 34:15621–15634,
615 2021.
- 616
- 617 Stephane Ross and J Andrew Bagnell. Agnostic system identification for model-based reinforcement
618 learning. *arXiv preprint arXiv:1203.1007*, 2012.
- 619
- 620 Daniel Russo and Benjamin Van Roy. Eluder dimension and the sample complexity of optimistic
621 exploration. *Advances in Neural Information Processing Systems*, 26, 2013.
- 622
- 623 Yuda Song and Wen Sun. Pc-mlp: Model-based reinforcement learning with policy cover guided
624 exploration. In *International Conference on Machine Learning*, pp. 9801–9811. PMLR, 2021.
- 625
- 626 Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Model-based
627 rl in contextual decision processes: Pac bounds and exponential improvements over model-free
628 approaches. In *Conference on learning theory*, pp. 2898–2933. PMLR, 2019.
- 629
- 630 Jean Tarbouriech, Runlong Zhou, Simon S Du, Matteo Pirota, Michal Valko, and Alessandro Lazaric.
631 Stochastic shortest path: Minimax, parameter-free and towards horizon-free regret. *Advances in
632 neural information processing systems*, 34:6843–6855, 2021.
- 633
- 634 Masatoshi Uehara and Wen Sun. Pessimistic model-based offline reinforcement learning under partial
635 coverage. *arXiv preprint arXiv:2107.06226*, 2021.
- 636
- 637 Masatoshi Uehara, Xuezhou Zhang, and Wen Sun. Representation learning for online and offline rl
638 in low-rank mdps. *arXiv preprint arXiv:2110.04652*, 2021.
- 639
- 640 Arun Venkatraman, Roberto Capobianco, Lerrel Pinto, Martial Hebert, Daniele Nardi, and J Andrew
641 Bagnell. Improved learning of dynamics models for control. In *2016 International Symposium on
642 Experimental Robotics*, pp. 703–713. Springer, 2017.
- 643
- 644 Andrew J Wagenmaker, Yifang Chen, Max Simchowitz, Simon Du, and Kevin Jamieson. First-order
645 regret in reinforcement learning with linear function approximation: A robust estimation approach.
646 In *International Conference on Machine Learning*, pp. 22384–22429. PMLR, 2022.
- 647
- 648 Kaiwen Wang, Kevin Zhou, Runzhe Wu, Nathan Kallus, and Wen Sun. The benefits of being
649 distributional: Small-loss bounds for reinforcement learning. *Advances in Neural Information
650 Processing Systems*, 36, 2023.
- 651
- 652 Kaiwen Wang, Owen Oertell, Alekh Agarwal, Nathan Kallus, and Wen Sun. More benefits of being
653 distributional: Second-order bounds for reinforcement learning. *arXiv preprint arXiv:2402.07198*,
654 2024.

- 648 Ruosong Wang, Simon S Du, Lin F Yang, and Sham M Kakade. Is long horizon reinforcement learning
649 more difficult than short horizon reinforcement learning? *arXiv preprint arXiv:2005.00527*, 2020.
650
- 651 Grady Williams, Nolan Wagener, Brian Goldfain, Paul Drews, James M Rehg, Byron Boots, and
652 Evangelos A Theodorou. Information theoretic mpc for model-based reinforcement learning. In
653 *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 1714–1721. IEEE,
654 2017.
- 655 Mengjiao Yang, Yilun Du, Kamyar Ghasemipour, Jonathan Tompson, Dale Schuurmans, and Pieter
656 Abbeel. Learning interactive real-world simulators. *arXiv preprint arXiv:2310.06114*, 2023.
657
- 658 Chenlu Ye, Rui Yang, Quanquan Gu, and Tong Zhang. Corruption-robust offline reinforcement
659 learning with general function approximation. *Advances in Neural Information Processing Systems*,
660 36, 2024.
- 661 Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement
662 learning without domain knowledge using value function bounds. In *International Conference on*
663 *Machine Learning*, pp. 7304–7312. PMLR, 2019.
- 664 Wenhao Zhan, Masatoshi Uehara, Wen Sun, and Jason D Lee. Pac reinforcement learning for
665 predictive state representations. *arXiv preprint arXiv:2207.05738*, 2022.
666
- 667 Junkai Zhang, Weitong Zhang, and Quanquan Gu. Optimal horizon-free reward-free exploration for
668 linear mixture mdps. In *International Conference on Machine Learning*, pp. 41902–41930. PMLR,
669 2023.
- 670 Tong Zhang. From ϵ -entropy to kl-entropy: Analysis of minimum information complexity density
671 estimation. *The Annals of Statistics*, pp. 2180–2210, 2006.
- 672 Zihan Zhang, Xiangyang Ji, and Simon Du. Is reinforcement learning more difficult than bandits?
673 a near-optimal algorithm escaping the curse of horizon. In *Conference on Learning Theory*, pp.
674 4528–4531. PMLR, 2021a.
- 675 Zihan Zhang, Jiaqi Yang, Xiangyang Ji, and Simon S Du. Improved variance-aware confidence sets
676 for linear bandits and linear mixture mdp. *Advances in Neural Information Processing Systems*,
677 34:4342–4355, 2021b.
- 678 Zihan Zhang, Xiangyang Ji, and Simon Du. Horizon-free reinforcement learning in polynomial time:
679 the power of stationary policies. In *Conference on Learning Theory*, pp. 3858–3904. PMLR, 2022.
- 680 Zihan Zhang, Yuxin Chen, Jason D Lee, and Simon S Du. Settling the sample complexity of
681 online reinforcement learning. In *The Thirty Seventh Annual Conference on Learning Theory*, pp.
682 5213–5219. PMLR, 2024a.
- 683 Zihan Zhang, Jason D Lee, Yuxin Chen, and Simon S Du. Horizon-free regret for linear markov
684 decision processes. *arXiv preprint arXiv:2403.10738*, 2024b.
- 685 Heyang Zhao, Jiafan He, and Quanquan Gu. A nearly optimal and low-switching algorithm for
686 reinforcement learning with general function approximation. *arXiv preprint arXiv:2311.15238*,
687 2023a.
- 688 Heyang Zhao, Jiafan He, Dongruo Zhou, Tong Zhang, and Quanquan Gu. Variance-dependent regret
689 bounds for linear bandits and reinforcement learning: Adaptivity and computational efficiency. In
690 *The Thirty Sixth Annual Conference on Learning Theory*, pp. 4977–5020. PMLR, 2023b.
- 691 Han Zhong, Wei Xiong, Sirui Zheng, Liwei Wang, Zhaoran Wang, Zhuoran Yang, and Tong Zhang.
692 Gec: A unified framework for interactive decision making in mdp, pomdp, and beyond. *arXiv*
693 *preprint arXiv:2211.01962*, 2022.
- 694 Dongruo Zhou and Quanquan Gu. Computationally efficient horizon-free reinforcement learning for
695 linear mixture mdps. *Advances in neural information processing systems*, 35:36337–36349, 2022.
696
- 697 Runlong Zhou, Zhang Zihan, and Simon Shaolei Du. Sharp variance-dependent bounds in reinforce-
698 ment learning: Best of both worlds in stochastic and deterministic environments. In *International*
699 *Conference on Machine Learning*, pp. 42878–42914. PMLR, 2023.

A SUMMARY OF CONTENTS IN THE APPENDIX

The Appendix is organized as follows.

In Appendix B, we provide some new analyses for Eluder dimension, which we will use for proving the regret bounds for the online RL setting.

In Appendix C, we provide some other supporting lemmas that will be used in our proofs.

In Appendix D, we provide the detailed proofs for the online RL setting (Section 4). Specifically, in Appendix D.1 we give the proof of Theorem 1; in Appendix D.2, we show the proof of Corollary 2; in Appendix D.3, we give the proof of Corollary 3.

In Appendix E, we provide the detailed proofs for the offline RL setting (Section 5). Specifically, in Appendix E.1 we give the proof of Theorem 2; in Appendix E.2, we show the proof of Corollary 4; in Appendix E.3, we give the proof of Corollary 5; in Appendix E.4, we show the proof of the claim in Example 2.

B ANALYSIS REGARDING THE ELUDER DIMENSION

For simplicity, we denote $x_h^k = (s_h^k, a_h^k)$.

First we have two technical lemma. The first lemma bounds the summation of “self-normalization” terms by the Eluder dimension. Our result generalizes the previous result by Zhao et al. (2023a) from the ℓ_2 -Eluder dimension to the ℓ_1 case.

Lemma 3. *Suppose for all $g \in \Psi$, $|g| \leq 1$ and $\lambda > 1$, then we have*

$$\sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, \sup_{g \in \Psi} \frac{|g(x_h^k)|}{\sum_{k'=1}^{k-1} \sum_{h'=1}^H |g(x_{h'}^{k'})| + \sum_{h'=1}^{h-1} |g(x_{h'}^k)| + \lambda} \right\} \leq 12 \log^2(4\lambda KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(8\lambda KH)) + \lambda^{-1}.$$

Proof of Lemma 3. We follow the proof steps of Theorem 4.6 in Zhao et al. (2023a). For simplicity, we use $n = KH$, $i = kH + h$ to denote the indices and denote x_h^k by x_i . Then we need to prove

$$\sum_{i=1}^n \min \left\{ 1, \sup_{g \in \Psi} \frac{|g(x_i)|}{\sum_{t=1}^{i-1} |g(x_t)| + \lambda} \right\} \leq 12 \log^2(4\lambda n) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(8\lambda n)) + \lambda^{-1} \quad (8)$$

Let

$$g_i = \operatorname{argmax}_{g \in \Psi} \frac{|g(x_i)|}{\sum_{j=1}^{i-1} |g(x_j)| + \lambda} \quad (9)$$

For any $1/(\lambda n) \leq \rho \leq 1$ and $1 \leq j \leq \lceil \log(4\lambda n) \rceil$, we define

$$A_\rho^j = \left\{ i \in [n] : 2^{-j} < |g_i(x_i)| \leq 2^{-j+1}, \frac{|g_i(x_i)|}{\sum_{t=1}^{i-1} |g_i(x_t)| + \lambda} \geq \rho/2 \right\}, \quad d_j := DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 2^{-j}). \quad (10)$$

Next we only consider the set A_ρ^j where $|A_\rho^j| > d_j$. We denote $A_\rho^j = \{a_1, \dots, a_A\}$, where $A = |A_\rho^j|$ and $\{a_i\}$ keeps the same order as $\{x_i\}$. Next we do the following constructions. We maintain $k = \lfloor (A-1)/d_j \rfloor$ number of queues Q_1, \dots, Q_k , all of them initialized as emptysets. We put a_1 into Q_1 . For $a_i, i \geq 2$, we put a_i into Q_l , where Q_l is the first queue where a_i is 2^{-j} -independent of all elements in Q_l . Let i_{\max} be the smallest i when we can not put a_i into any existing queue.

We claim that i_{\max} indeed exists, i.e., our construction will stop before we put all elements in A_ρ^j into Q_1, \dots, Q_k . In fact, note the fact that the length of each Q_l is always no more than d_j , which is due to the fact that any 2^{-j} -independent sequence’s length is at most d_j . Meanwhile, since we only have $k = \lfloor (A-1)/d_j \rfloor$, then the amount of elements in $Q_1 \cup \dots \cup Q_k$ will be upper bounded by $k \cdot d_j < A$. That suggests at least one element in A_ρ^j is not contained by $Q_1 \cup \dots \cup Q_k$, i.e., i_{\max} exists.

By the definition of i_{\max} , we know that $a_{i_{\max}}$ is 2^{-j} -dependent to each Q_l . Next we give a bound of A . First, note

$$\sum_{t=1}^{i_{\max}-1} |g_{i_{\max}}(x_t)| \geq \sum_{t \in Q_1 \cup \dots \cup Q_k} |g_{i_{\max}}(a_t)| = \sum_{l=1}^k \sum_{t \in Q_l} |g_{i_{\max}}(a_t)| > k \cdot 2^{-j}, \quad (11)$$

where the first inequality holds since Q_l are the elements that appear before $a_{i_{\max}}$, the second one holds due to the following induction of Eluder dimension: since $a_{i_{\max}}$ is 2^{-j} -dependent to Q_l , then we have

$$\forall g \in \Psi, \sum_{t \in Q_l} |g(a_t)| \leq 2^{-j} \Rightarrow |g(a_{i_{\max}})| \leq 2^{-j}. \quad (12)$$

Therefore, given the fact $|g_{i_{\max}}(a_{i_{\max}})| > 2^{-j}$ (recall the definition of A_ρ^j), we must have $\sum_{t \in Q_l} |g_{i_{\max}}(a_t)| > 2^{-j}$ as well, which suggests the second inequality of Equation 11 holds. Second, we have

$$\sum_{t=1}^{i_{\max}-1} |g_{i_{\max}}(x_t)| \leq 2/\rho \cdot |g_{i_{\max}}(a_{i_{\max}})| \leq 4 \cdot 2^{-j}/\rho, \quad (13)$$

where both inequalities hold due to the definition of A_ρ^j . Combining Equation 11 and Equation 13, we have

$$k < 4/\rho \Rightarrow A \leq 4d_j/\rho + d_j \leq 5d_j/\rho. \quad (14)$$

Therefore, we have that for all $\rho, j, |A_\rho^j| \leq 5d_j/\rho$.

Finally we prove Equation 8. $1/(\lambda n) \leq \rho \leq 1$ and $1 \leq j \leq \lceil \log(4\lambda n) \rceil = J$. Denote

$$A_\rho = \left\{ i \in [n] : \frac{|g_i(x_i)|}{\sum_{t=1}^{i-1} |g_i(x_t)| + \lambda} \geq \rho/2 \right\}. \quad (15)$$

Then it is easy to notice that $|A_\rho| = \sum_j |A_\rho^j| \leq \lceil \log(4\lambda n) \rceil \cdot 5d_j/\rho$, where we use the fact that the Eluder dimension d_j is increasing. Therefore, by the standard peeling technique, we have

$$\begin{aligned} \sum_{i=1}^n \min \left\{ 1, \sup_{g \in \Psi} \frac{|g(x_i)|}{\sum_{t=1}^{i-1} |g(x_t)| + \lambda} \right\} &= \sum_{j \in [\lceil \log(\lambda n) \rceil]} \sum_{i \in A_{2^{-j}} \setminus A_{2^{-j+1}}} + \sum_{j = \lceil \log(\lambda n) \rceil} \sum_{i \in A_{2^{-j+1}}} \\ &\leq \sum_{j \in [\lceil \log(\lambda n) \rceil]} \sum_{i \in A_{2^{-j}} \setminus A_{2^{-j+1}}} 2^{-j-1} + n \cdot 1/(\lambda n) \\ &\leq \sum_{j \in [\lceil \log(\lambda n) \rceil]} \sum_{i \in A_{2^{-j}}} 2^{-j-1} + n \cdot 1/(\lambda n) \\ &\leq \lceil \log(\lambda n) \rceil \cdot \lceil \log(4\lambda n) \rceil \cdot 3d_J + \lambda^{-1}, \end{aligned}$$

which concludes our proof. \square

Next lemma gives a bound to bound the number of episodes where the behavior along these episodes are ‘‘bad’’. Intuitively speaking, our lemma suggests we only have limited number of bad episodes, therefore won’t affect the final performance of our algorithm.

Lemma 4. *Given $\lambda > 1$. There exists at most*

$$13 \log^2(4\lambda KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(8\lambda KH)) \quad (16)$$

number of $k \in [K]$ satisfying the following claim

$$\sup_{g \in \Psi} \frac{\lambda + \sum_{k'=1}^k \sum_{h'=1}^H |g(x_{h'}^{k'})|}{\lambda + \sum_{k'=1}^{k-1} \sum_{h'=1}^H |g(x_{h'}^{k'})|} > 4. \quad (17)$$

Proof of Lemma 4. Note that

$$\begin{aligned} &\sum_{k=1}^K \min \left\{ 2, \log \sup_{g \in \Psi} \frac{\lambda + \sum_{k'=1}^k \sum_{h'=1}^H |g(x_{h'}^{k'})|}{\lambda + \sum_{k'=1}^{k-1} \sum_{h'=1}^H |g(x_{h'}^{k'})|} \right\} \\ &\leq \sum_{k=1}^K \min \left\{ 2, \log \prod_{h=1}^H \sup_{g \in \Psi} \frac{\lambda + \sum_{k'=1}^{k-1} \sum_{h'=1}^H |g(x_{h'}^{k'})| + \sum_{h'=1}^h |g(x_{h'}^k)|}{\lambda + \sum_{k'=1}^{k-1} \sum_{h'=1}^H |g(x_{h'}^{k'})| + \sum_{h'=1}^{h-1} |g(x_{h'}^k)|} \right\} \\ &= \sum_{k=1}^K \min \left\{ 2, \sum_{h=1}^H \log \left(1 + \sup_{g \in \Psi} \frac{|g(x_h^k)|}{\lambda + \sum_{k'=1}^{k-1} \sum_{h'=1}^H |g(x_{h'}^{k'})| + \sum_{h'=1}^{h-1} |g(x_{h'}^k)|} \right) \right\} \end{aligned}$$

$$\begin{aligned}
& \leq \sum_{k=1}^K \sum_{h=1}^H \min \left\{ 2, \sup_{g \in \Psi} \frac{|g(x_h^k)|}{\sum_{k'=1}^{k-1} \sum_{h'=1}^H |g(x_{h'}^{k'})| + \sum_{h'=1}^{h-1} |g(x_{h'}^k)| + \lambda} \right\}, \\
& \leq 2 \sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, \sup_{g \in \Psi} \frac{|g(x_h^k)|}{\sum_{k'=1}^{k-1} \sum_{h'=1}^H |g(x_{h'}^{k'})| + \sum_{h'=1}^{h-1} |g(x_{h'}^k)| + \lambda} \right\} \\
& \leq 24 \log^2(4\lambda KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(8\lambda KH)) + 2\lambda^{-1} \\
& \leq 26 \log^2(4\lambda KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(8\lambda KH)). \tag{18}
\end{aligned}$$

where the first inequality holds since $\sup_g \prod f(g) \leq \prod \sup_g f(g)$, the second one holds since $\log(1+x) \leq x$, the fourth one holds due to Lemma 3. Therefore, there are at most

$$26 \log^2(4\lambda KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(8\lambda KH))/2$$

number of k satisfying

$$\log \sup_{g \in \Psi} \frac{\lambda + \sum_{k'=1}^k \sum_{h'=1}^H |g(x_{h'}^{k'})|}{\lambda + \sum_{k'=1}^{k-1} \sum_{h'=1}^H |g(x_{h'}^{k'})|} > 2,$$

which concludes the proof. \square

We next have the following lemma, which bounds the regret by the Eluder dimension.

Lemma 5 (Theorem 5.3, Wang et al. 2023). *Let $C := \sup_{(s,a) \in \mathcal{S} \times \mathcal{A}, f \in \Psi} |f((s,a))|$ be the envelope. For any sequences $f^{(1)}, \dots, f^{(N)} \subseteq \Psi$, $(s,a)^{(1)}, \dots, (s,a)^{(N)} \subseteq \mathcal{S} \times \mathcal{A}$, let β be a constant such that for all $n \in [N]$ we have, $\sum_{i=1}^{n-1} |f^{(n)}((s,a)^i)| \leq \beta$. Then, for all $n \in [N]$, we have*

$$\sum_{t=1}^n |f^{(t)}((s,a)^t)| \leq \inf_{0 < \epsilon \leq 1} \{DE_1(\Psi, \mathcal{S} \times \mathcal{A}, \epsilon)(2C + \beta \log(C/\epsilon)) + n\epsilon\}.$$

Given Lemma 4 and Lemma Lemma 5, we are able to prove the following key lemma.

Lemma 6 (New Eluder Pigeon Lemma). *Let the event \mathcal{E} be*

$$\mathcal{E} : \forall k \in [K], \sum_{i=1}^{k-1} \sum_{h=1}^H \mathbb{H}^2(\widehat{P}^k(s_h^i, a_h^i) \| P^*(s_h^i, a_h^i)) \leq \eta. \tag{19}$$

Then under event \mathcal{E} , there exists a set $\mathcal{K} \in [K]$ such that

- We have $|\mathcal{K}| \leq 13 \log^2(4\eta KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(8\eta KH))$.
- We have

$$\begin{aligned}
& \sum_{k \in [K] \setminus \mathcal{K}} \sum_{h=1}^H \mathbb{H}^2(P^*(s_h^k, a_h^k) \| \widehat{P}^k(s_h^k, a_h^k)) \\
& \leq \inf_{0 < \epsilon \leq 1} \{DE_1(\Psi, \mathcal{S} \times \mathcal{A}, \epsilon)(2 + 7\eta \log(1/\epsilon)) + KH\epsilon\} \\
& \leq DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH)(2 + 7\eta \log(KH)) + 1, \tag{20}
\end{aligned}$$

where the function class $\Psi = \{(s,a) \mapsto \mathbb{H}^2(P^*(s,a) \| P(s,a)) : P \in \mathcal{P}\}$.

Proof of Lemma 6. We interchangeably use $n = kH + h$ to denote the indices of s_h^k, a_h^k . We set $f^{(n)}((s,a))$ in Lemma 5 as $H^2(P^k(s,a) \| P^*(s,a))$.

First, we prove that the β in Lemma 5 can be selected as 7η under event \mathcal{E} . To show that, let \mathcal{K} denote all the k stated in Lemma 4. Then for all k such that $k+1 \notin \mathcal{K}$, $h = 2, \dots, H$, let $n = kH + h$, we have

$$\sum_{i=0}^{n-1} |f^{(n)}((s,a)^i)| \leq \sum_{i=0}^{kH+H} |f^{(n)}((s,a)^i)|$$

$$\begin{aligned}
&= \left(\lambda + \sum_{i=0}^{kH} |f^{(n)}((s, a)^i)| \right) \cdot \frac{\sum_{i=0}^{kH+H} |f^{(n)}((s, a)^i)| + \lambda}{\sum_{i=0}^{kH} |f^{(n)}((s, a)^i)| + \lambda} - \lambda \\
&\leq \left(\lambda + \sum_{i=0}^{kH} |f^{(n)}((s, a)^i)| \right) \cdot 4 - \lambda \\
&\leq 7\eta,
\end{aligned} \tag{21}$$

where the second inequality holds due to Lemma 4, the last one holds due to the definition of \mathcal{E} . Therefore, we prove our lemma by the conclusion of Lemma 5 with $\beta = 7\eta$. \square

C OTHER SUPPORTING LEMMAS

Lemma 7 (Simulation Lemma (Agarwal et al. (2019))). *We have*

$$V_{0;P^*}^\pi - V_{0;\hat{P}}^\pi \leq \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^\pi} \left[\left| \mathbb{E}_{s' \sim P^*(s,a)} V_{h+1;\hat{P}}^\pi(s') - \mathbb{E}_{s' \sim \hat{P}(s,a)} V_{h+1;\hat{P}}^\pi(s') \right| \right].$$

Lemma 8 (Change of Variance Lemma (Lemma C.5 in Jin et al. (2018))).

$$\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^\pi} \left[(\mathbb{V}_{P^*} V_{h+1;P^*}^\pi)(s, a) \right] = \text{VaR}_\pi.$$

Lemma 9 (Generalization bounds of MLE for finite model class (Theorem E.4 in Wang et al. (2023))). *Let \mathcal{X} be the context/feature space and \mathcal{Y} be the label space, and we are given a dataset $D = \{(x_i, y_i)\}_{i \in [n]}$ from a martingale process: for $i = 1, 2, \dots, n$, sample $x_i \sim \mathcal{D}_i(x_{1:i-1}, y_{1:i-1})$ and $y_i \sim p(\cdot | x_i)$. Let $f^*(x, y) = p(y | x)$ and we are given a realizable, i.e., $f^* \in \mathcal{F}$, function class $\mathcal{F} : \mathcal{X} \times \mathcal{Y} \rightarrow \Delta(\mathbb{R})$ of distributions. Suppose \mathcal{F} is finite. Fix any $\delta \in (0, 1)$, set $\beta = \log(|\mathcal{F}|/\delta)$ and define*

$$\widehat{\mathcal{F}} = \left\{ f \in \mathcal{F} : \sum_{i=1}^n \log f(x_i, y_i) \geq \max_{\tilde{f} \in \mathcal{F}} \sum_{i=1}^n \log \tilde{f}(x_i, y_i) - 4\beta \right\}.$$

Then w.p. at least $1 - \delta$, the following holds:

- (1) The true distribution is in the version space, i.e., $f^* \in \widehat{\mathcal{F}}$.
- (2) Any function in the version space is close to the ground truth data-generating distribution, i.e., for all $f \in \widehat{\mathcal{F}}$

$$\sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} \left[\mathbb{H}^2(f(x, \cdot) \| f^*(x, \cdot)) \right] \leq 22\beta.$$

Lemma 10 (Generalization bounds of MLE for infinite model class (Theorem E.5 in Wang et al. (2023))). *Let \mathcal{X} be the context/feature space and \mathcal{Y} be the label space, and we are given a dataset $D = \{(x_i, y_i)\}_{i \in [n]}$ from a martingale process: for $i = 1, 2, \dots, n$, sample $x_i \sim \mathcal{D}_i(x_{1:i-1}, y_{1:i-1})$ and $y_i \sim p(\cdot | x_i)$. Let $f^*(x, y) = p(y | x)$ and we are given a realizable, i.e., $f^* \in \mathcal{F}$, function class $\mathcal{F} : \mathcal{X} \times \mathcal{Y} \rightarrow \Delta(\mathbb{R})$ of distributions. Suppose \mathcal{F} is finite. Fix any $\delta \in (0, 1)$, set $\beta = \log(\mathcal{N}_{[]}((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_\infty)/\delta)$ (where $\mathcal{N}_{[]}((n|\mathcal{Y}|)^{-1}, \mathcal{F}, \|\cdot\|_\infty)$ is the bracketing number defined in Definition 2) and define*

$$\widehat{\mathcal{F}} = \left\{ f \in \mathcal{F} : \sum_{i=1}^n \log f(x_i, y_i) \geq \max_{\tilde{f} \in \mathcal{F}} \sum_{i=1}^n \log \tilde{f}(x_i, y_i) - 7\beta \right\}.$$

Then w.p. at least $1 - \delta$, the following holds:

- (1) The true distribution is in the version space, i.e., $f^* \in \widehat{\mathcal{F}}$.
- (2) Any function in the version space is close to the ground truth data-generating distribution, i.e., for all $f \in \widehat{\mathcal{F}}$

$$\sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{D}_i} \left[\mathbb{H}^2(f(x, \cdot) \| f^*(x, \cdot)) \right] \leq 28\beta.$$

Lemma 11 (Recursion Lemma). *Let $G > 0$ be a positive constant, $a < G/2$ is also a positive constant, and let $\{C_m\}_{m=0}^{N=\lceil \log_2(\frac{KH}{G}) \rceil}$ be a sequence of positive real numbers satisfying:*

1. $C_m \leq 2^m G + \sqrt{aC_{m+1}} + a$ for all $m \geq 0$,
2. $C_m \leq KH$ for all $m \geq 0$, where $K > 0$ and $H > 0$ are positive constants.

Then, it holds that:

$$C_0 \leq 4G.$$

Proof of Lemma 11. We will prove by induction that for all $m \geq 0$,

$$C_m \leq 2^{m+2}G.$$

Then, for $m = 0$, this would immediately show $C_0 \leq 4G$.

1. The base case $m = N$:

Since $N = \lceil \log_2(\frac{KH}{G}) \rceil$, it is obvious that $2^{N+2}G \geq KH$. Thus, $C_N \leq KH \leq 2^{N+2}G$, the inequality holds for $m = N$.

2. The induction step:

Assume that for some $m \geq 0$, for C_{m+1} , we have:

$$C_{m+1} \leq 2^{m+1+2}G = 2^{m+3}G.$$

Then, we have

$$\begin{aligned} C_m &\leq 2^m G + \sqrt{aC_{m+1}} + a \\ &\leq 2^m G + \sqrt{a2^{m+3}G} + a \\ &\leq 2^m G + \sqrt{\frac{G}{2} \cdot 2^{m+3}G} + \frac{G}{2} \\ &= G \cdot (2^m + 2^{m/2+1} + 2^{-1}) \\ &\leq G \cdot (2^m + 2^{m+1} + 2^m) \\ &= 2^{m+2}G. \end{aligned} \tag{22}$$

Therefore, by induction, we have for all $m \geq 0$,

$$C_m \leq 2^{m+2}G.$$

And the proof follows by setting $m = 0$. \square

D DETAILED PROOFS FOR THE ONLINE SETTING IN SECTION 4

D.1 PROOF OF THEOREM 1

The following is the full proof of Theorem 1.

For notational simplicity, throughout this whole section, we denote

$$\begin{aligned} A &:= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s_h^k, a_h^k) \right] \\ B &:= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s_h^k, a_h^k) \right], \\ C_m &:= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} (V_{h+1}^{\pi^k} - V_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k) \right] \end{aligned}$$

$$G := \sqrt{\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\nabla_{P^*} V_{h+1}^{\pi^k}) (s_h^k, a_h^k) \right]} \cdot \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH) \\ + \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH)$$

$$I_h^k := \mathbb{E}_{s' \sim P^*(s_h^k, a_h^k)} V_{h+1}^{\pi^k}(s') - V_{h+1}^{\pi^k}(s_h^k, a_h^k) \quad (23)$$

We use $\mathbb{I}\{\cdot\}$ to denote the indicator function. We define the following events which we will later show that they happen with high probability.

$$\mathcal{E}_1 := \{ \forall k \in [K-1] : P^* \in \widehat{\mathcal{P}}^k, \text{ and } \sum_{i=0}^{k-1} \sum_{h=0}^{H-1} \mathbb{H}^2(P^*(s_h^i, a_h^i) \| \widehat{P}^k(s_h^i, a_h^i)) \leq 22 \log(K|\mathcal{P}|/\delta) \}, \quad (24)$$

$$\mathcal{E}_2 := \left\{ \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} I_h^k \lesssim \sqrt{\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} (\nabla_{P^*} V_{h+1}^{\pi^k}) (s_h^k, a_h^k) \log(1/\delta) + \log(1/\delta)} \right\}, \quad (25)$$

$$\mathcal{E}_3 := \mathcal{E}_1 \cap \{ \forall m \in [0, \lceil \log_2(\frac{KH}{G}) \rceil] : C_m \lesssim 2^m G + \sqrt{\log(1/\delta) \cdot C_{m+1}} + \log(1/\delta) \}, \quad (26)$$

$$\mathcal{E}_4 := \left\{ \sum_{k=0}^{K-1} \sum_{h=0}^{H-1} \left[(\nabla_{P^*} V_{h+1}^{\pi^k}) (s_h^k, a_h^k) \right] \lesssim \sum_{k=0}^{K-1} \text{VaR}_{\pi^k} + \log(1/\delta) \right\}, \quad (27)$$

$$\mathcal{E}_5 := \left\{ \sum_{k=0}^{K-1} \sum_{h=1}^H r(s_h^k, a_h^k) - \sum_{k=0}^{K-1} V_{0;P^*}^{\pi^k} \lesssim \sqrt{\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \log(1/\delta) + \log(1/\delta)} \right\}, \quad (28)$$

$$\mathcal{E} := \mathcal{E}_2 \cap \mathcal{E}_3 \cap \mathcal{E}_4 \cap \mathcal{E}_5. \quad (29)$$

First, by the realizability assumption, the standard generalization bound for MLE (Lemma 9) with simply setting D_i to be the delta distribution on the realized (s_h^k, a_h^k) pairs, and a union bound over K episodes, we have that w.p. at least $1 - \delta$, for any $k \in [0, K-1]$:

$$(1) P^* \in \widehat{\mathcal{P}}^k;$$

$$(2)$$

$$\sum_{i=0}^{k-1} \sum_{h=0}^{H-1} \mathbb{H}^2(P^*(s_h^i, a_h^i) \| \widehat{P}^k(s_h^i, a_h^i)) \leq 22 \log(K|\mathcal{P}|/\delta). \quad (30)$$

This directly indicates that

$$P(\mathbb{I}\{\mathcal{E}_1\}) \geq 1 - \delta. \quad (31)$$

Under event \mathcal{E}_1 , with the realizability in above (1), and by the optimistic algorithm design $(\pi^k, \widehat{P}^k) \leftarrow \arg\max_{\pi \in \Pi, P \in \widehat{\mathcal{P}}^k} V_{0;P}^{\pi}(s_0)$, for any $k \in [0, K-1]$, we have the following optimism guarantee

$$V_{0;P^*}^* \leq \max_{\pi \in \Pi, P \in \widehat{\mathcal{P}}^k} V_{0;P}^{\pi} = V_{0;\widehat{P}^k}^{\pi^k}.$$

Then, under event \mathcal{E}_1 , we use Lemma 6 and Equation 30 to get the following:

There exists a set $\mathcal{K} \subseteq [K-1]$ such that

- $|\mathcal{K}| \leq 13 \log^2(88 \log(K|\mathcal{P}|/\delta)KH) \cdot \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(176 \log(K|\mathcal{P}|/\delta)KH))$
- And

$$\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \mathbb{H}^2(P^*(s_h^k, a_h^k) \| \widehat{P}^k(s_h^k, a_h^k)) \\ \leq \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot (2 + 154 \log(K|\mathcal{P}|/\delta) \log(KH)) + 1 \\ \lesssim \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH). \quad (32)$$

We upper bound the regret with optimism, and by dividing $k \in [K-1]$ into \mathcal{K} and $[K-1] \setminus \mathcal{K}$ with the assumption that the trajectory-wise cumulative reward is normalized in $[0,1]$, as follows

$$\begin{aligned}
& \sum_{k=0}^{K-1} V_{0;P^*}^* - \sum_{k=0}^{K-1} \sum_{h=1}^H r(s_h^k, a_h^k) \\
& \leq |\mathcal{K}| + \sum_{k \in [K-1] \setminus \mathcal{K}} \left(V_{0;\widehat{P}^k}^{\pi^k} - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \right) \\
& \lesssim \log^2(\log(K|\mathcal{P}|/\delta)KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(\log(K|\mathcal{P}|/\delta)KH)) + \sum_{k \in [K-1] \setminus \mathcal{K}} \left(V_{0;\widehat{P}^k}^{\pi^k} - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \right). \tag{33}
\end{aligned}$$

We then do the following decomposition. Note that for any $k \in [K-1]$, policy π^k is deterministic. We have that for any $k \in [K-1]$

$$\begin{aligned}
& V_{0;\widehat{P}^k}^{\pi^k}(s_0^k) - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \\
& = Q_{0;\widehat{P}^k}^{\pi^k}(s_0^k, a_0^k) - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \\
& = r(s_0^k, a_0^k) + \mathbb{E}_{s' \sim \widehat{P}^k(s_0^k, a_0^k)} V_{1;\widehat{P}^k}^{\pi^k}(s') - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \\
& = \mathbb{E}_{s' \sim \widehat{P}^k(s_0^k, a_0^k)} V_{1;\widehat{P}^k}^{\pi^k}(s') - \sum_{h=1}^H r(s_h^k, a_h^k) \\
& = \mathbb{E}_{s' \sim P^*(s_0^k, a_0^k)} V_{1;\widehat{P}^k}^{\pi^k}(s') - \sum_{h=1}^H r(s_h^k, a_h^k) + \underbrace{\mathbb{E}_{s' \sim \widehat{P}^k(s_0^k, a_0^k)} V_{1;\widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*(s_0^k, a_0^k)} V_{1;\widehat{P}^k}^{\pi^k}(s')}_{I_0^k} \\
& = V_{1;\widehat{P}^k}^{\pi^k}(s_1^k) - \sum_{h=1}^{H-1} r(s_h^k, a_h^k) + \underbrace{\mathbb{E}_{s' \sim P^*(s_0^k, a_0^k)} V_{1;\widehat{P}^k}^{\pi^k}(s') - V_{1;\widehat{P}^k}^{\pi^k}(s_1^k)}_{I_0^k} \\
& \quad + \mathbb{E}_{s' \sim \widehat{P}^k(s_0^k, a_0^k)} V_{1;\widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*(s_0^k, a_0^k)} V_{1;\widehat{P}^k}^{\pi^k}(s'),
\end{aligned}$$

where we use the Bellman equation for several times.

Then, by doing this recursively, we can get for any $k \in [K-1]$

$$\begin{aligned}
& V_{0;\widehat{P}^k}^{\pi^k}(s_h^k) - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \\
& \leq \sum_{h=0}^{H-1} I_h^k + \sum_{h=0}^{H-1} \left| \mathbb{E}_{s' \sim \widehat{P}^k(s_h^k, a_h^k)} V_{h+1;\widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*(s_h^k, a_h^k)} V_{h+1;\widehat{P}^k}^{\pi^k}(s') \right| \tag{34}
\end{aligned}$$

Therefore,

$$\begin{aligned}
& \sum_{k \in [K-1] \setminus \mathcal{K}} \left(V_{0;\widehat{P}^k}^{\pi^k}(s_h^k) - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \right) \\
& \leq \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} I_h^k + \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left| \mathbb{E}_{s' \sim \widehat{P}^k(s_h^k, a_h^k)} V_{h+1;\widehat{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*(s_h^k, a_h^k)} V_{h+1;\widehat{P}^k}^{\pi^k}(s') \right| \tag{35}
\end{aligned}$$

Next we bound $\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} I_h^k$. Note that by Azuma Bernstein's inequality, with probability at least $1 - \delta$

$$\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} I_h^k \leq \sqrt{2 \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} (\mathbb{V}_{P^*} V_{h+1;\widehat{P}^k}^{\pi^k})(s_h^k, a_h^k) \log(1/\delta)} + \frac{2}{3} \log(1/\delta) \tag{36}$$

This directly indicates that

$$P(\mathbb{I}\{\mathcal{E}_2\}) \geq 1 - \delta. \quad (37)$$

Then, we propose the following lemma.

Lemma 12 (Bound of sum of mean value differences for online RL). *Under event \mathcal{E}_1 , we have*

$$\begin{aligned} & \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left| \mathbb{E}_{s' \sim \widehat{\mathcal{P}}^k(s_h^k, a_h^k)} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*(s_h^k, a_h^k)} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s') \right| \\ & \lesssim \sqrt{\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k})(s_h^k, a_h^k) \right] \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K |\mathcal{P}| / \delta) \log(KH)} \\ & \quad + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K |\mathcal{P}| / \delta) \log(KH). \end{aligned}$$

Proof of Lemma 12. Under event \mathcal{E}_1 , we have

$$\begin{aligned} & \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left| \mathbb{E}_{s' \sim \widehat{\mathcal{P}}^k(s_h^k, a_h^k)} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*(s_h^k, a_h^k)} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s') \right| \\ & \leq 4 \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\sqrt{(\mathbb{V}_{P^*} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k})(s_h^k, a_h^k) D_{\Delta} \left(V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s' \sim P^*(s_h^k, a_h^k)) \parallel V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s' \sim \widehat{\mathcal{P}}^k(s_h^k, a_h^k)) \right)} \right] \\ & \quad + 5 \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[D_{\Delta} \left(V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s' \sim P^*(s_h^k, a_h^k)) \parallel V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s' \sim \widehat{\mathcal{P}}^k(s_h^k, a_h^k)) \right) \right] \\ & \leq 8 \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\sqrt{(\mathbb{V}_{P^*} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k})(s_h^k, a_h^k) \mathbb{H}^2 \left(V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s' \sim P^*(s_h^k, a_h^k)) \parallel V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s' \sim \widehat{\mathcal{P}}^k(s_h^k, a_h^k)) \right)} \right] \\ & \quad + 20 \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\mathbb{H}^2 \left(V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s' \sim P^*(s_h^k, a_h^k)) \parallel V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s' \sim \widehat{\mathcal{P}}^k(s_h^k, a_h^k)) \right) \right] \\ & \leq 8 \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\sqrt{(\mathbb{V}_{P^*} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k})(s_h^k, a_h^k) \mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{\mathcal{P}}^k(s_h^k, a_h^k) \right)} \right] \\ & \quad + 20 \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{\mathcal{P}}^k(s_h^k, a_h^k) \right) \right] \\ & \leq 8 \sqrt{\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k})(s_h^k, a_h^k) \right] \cdot \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{\mathcal{P}}^k(s_h^k, a_h^k) \right) \right]} \\ & \quad + 20 \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{\mathcal{P}}^k(s_h^k, a_h^k) \right) \right] \\ & \lesssim \sqrt{\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k})(s_h^k, a_h^k) \right] \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K |\mathcal{P}| / \delta) \log(KH)} \\ & \quad + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K |\mathcal{P}| / \delta) \log(KH), \quad (38) \end{aligned}$$

where in the first inequality, we use Lemma 1 to bound the difference of two means $\mathbb{E}_{s' \sim P^*(s_h^k, a_h^k)} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim \widehat{\mathcal{P}}^k(s_h^k, a_h^k)} V_{h+1; \widehat{\mathcal{P}}^k}^{\pi^k}(s')$ using variances and the triangle discrimination; in the second inequality we use the fact that that triangle discrimination is equivalent to squared Hellinger distance, i.e., $D_{\Delta} \leq 4\mathbb{H}^2$; the third inequality is via data processing inequality on the squared Hellinger distance; the fourth inequality is by the Cauchy–Schwarz inequality; the last inequality holds under \mathcal{E}_1 by Equation 32. \square

The next lemma shows that the event \mathcal{E}_3 happens with high probability.

Lemma 13 (Recursion Event Lemma). *Event \mathcal{E}_3 happens with high probability. Specifically, we have*

$$P(\mathbb{I}\{\mathcal{E}_3\}) \geq 1 - (1 + \lceil \log_2(\frac{KH}{G}) \rceil) \delta. \quad (39)$$

Proof of Lemma 13. Let $\Delta_{h+1}^{\pi^k} := V_{h+1; \widehat{P}^k}^{\pi^k} - V_{h+1}^{\pi^k}$. First, under event \mathcal{E}_1 , with happens with probability at least $1 - \delta$ by Equation 31, and also note that π^k is deterministic for any $k \in [K-1]$, we can prove the following

$$\begin{aligned} & \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[|(\Delta_h^{\pi^k})(s_h^k) - (P^* \Delta_{h+1}^{\pi^k})(s_h^k, a_h^k)| \right] \\ &= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[|(V_{h; \widehat{P}^k}^{\pi^k})(s_h^k) - (P^* V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k) - ((V_h^{\pi^k})(s_h^k) - (P^* V_{h+1}^{\pi^k})(s_h^k, a_h^k))| \right] \\ &= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[|r(s_h^k, a_h^k) + (\widehat{P}^k V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k) - (P^* V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k) - r(s_h^k, a_h^k)| \right] \\ &= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[|(\widehat{P}^k V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k) - (P^* V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k)| \right] \\ &= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\left| \mathbb{E}_{s' \sim P^*(s_h^k, a_h^k)} \left[V_{h+1; \widehat{P}^k}^{\pi^k}(s') \right] - \mathbb{E}_{s' \sim \widehat{P}^k(s_h^k, a_h^k)} \left[V_{h+1; \widehat{P}^k}^{\pi^k}(s') \right] \right| \right] \\ &\lesssim \sqrt{\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{P}^k}^{\pi^k})(s_h^k, a_h^k) \right] \cdot \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH)} \\ &\quad + \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH) \\ &= G \end{aligned} \quad (41)$$

where the first equality is by the definition of $\Delta_{h+1}^{\pi^k}$, the inequality holds under \mathcal{E}_1 by Lemma 12, and the last equality is by definition of A and G .

Under event \mathcal{E}_1 , with probability at least $1 - \lceil \log_2(\frac{KH}{G}) \rceil \delta$, for any $m \in [0, \lceil \log_2(\frac{KH}{G}) \rceil]$

$$\begin{aligned} C_m &= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} (V_{h+1; \widehat{P}^k}^{\pi^k} - V_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k) \right] \\ &= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(P^* (\Delta_{h+1}^{\pi^k})^{2^{m+1}})(s_h^k, a_h^k) - ((P^* (\Delta_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k))^2 \right] \\ &= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\Delta_{h+1}^{\pi^k})^{2^{m+1}}(s_{h+1}^k) \right] - \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[((P^* (\Delta_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k))^2 \right] \\ &\quad + \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left(\mathbb{E}_{s \sim P^*(s_h^k, a_h^k)} \left[(\Delta_{h+1}^{\pi^k})^{2^{m+1}}(s) \right] - (\Delta_{h+1}^{\pi^k})^{2^{m+1}}(s_{h+1}^k) \right) \\ &\leq \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\Delta_h^{\pi^k})^{2^{m+1}}(s_h^k) - ((P^* (\Delta_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k))^2 \right] \\ &\quad + \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left(\mathbb{E}_{s \sim P^*(s_h^k, a_h^k)} \left[(\Delta_{h+1}^{\pi^k})^{2^{m+1}}(s) \right] - (\Delta_{h+1}^{\pi^k})^{2^{m+1}}(s_h^k) \right) \\ &\lesssim \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\Delta_h^{\pi^k})^{2^{m+1}}(s_h^k) - ((P^* (\Delta_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k))^2 \right] + \log(1/\delta) \\ &\quad + \sqrt{\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \mathbb{V}_{P^*} \left((V_{h+1; \widehat{P}^k}^{\pi^k} - V_{h+1}^{\pi^k})^{2^{m+1}} \right) (s_h^k, a_h^k) \log(1/\delta)} \end{aligned}$$

$$\begin{aligned}
&= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\left((\Delta_h^{\pi^k})^{2^m}(s_h^k) + (P^*(\Delta_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k) \right) \cdot \left((\Delta_h^{\pi^k})^{2^m}(s_h^k) - (P^*(\Delta_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k) \right) \right] \\
&\quad + \sqrt{\log(1/\delta) \cdot C_{m+1}} + \log(1/\delta) \\
&= \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\left((\Delta_h^{\pi^k})^{2^m}(s_h^k) + (P^*(\Delta_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k) \right) \cdot \left((\Delta_h^{\pi^k})^{2^m}(s_h^k) - (P^*((\Delta_{h+1}^{\pi^k})^2)^{2^{m-1}})(s_h^k, a_h^k) \right) \right] \\
&\quad + \sqrt{\log(1/\delta) \cdot C_{m+1}} + \log(1/\delta) \\
&\leq \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\left((\Delta_h^{\pi^k})^{2^m}(s_h^k) + (P^*(\Delta_{h+1}^{\pi^k})^{2^m})(s_h^k, a_h^k) \right) \cdot \left((\Delta_h^{\pi^k})^{2^m}(s_h^k) - ((P^*(\Delta_{h+1}^{\pi^k})^2)(s_h^k, a_h^k))^{2^{m-1}} \right) \right] \\
&\quad + \sqrt{\log(1/\delta) \cdot C_{m+1}} + \log(1/\delta) \\
&\leq 2^m \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\left| (\Delta_h^{\pi^k})^2(s_h^k) - ((P^* \Delta_{h+1}^{\pi^k})(s_h^k, a_h^k))^2 \right| \right] + \sqrt{\log(1/\delta) \cdot C_{m+1}} + \log(1/\delta) \\
&= 2^m \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\left| \left((\Delta_h^{\pi^k})(s_h^k) + (P^* \Delta_{h+1}^{\pi^k})(s_h^k, a_h^k) \right) \cdot \left((\Delta_h^{\pi^k})(s_h^k) - (P^* \Delta_{h+1}^{\pi^k})(s_h^k, a_h^k) \right) \right| \right] \\
&\quad + \sqrt{\log(1/\delta) \cdot C_{m+1}} + \log(1/\delta) \\
&\leq 2 \cdot 2^m \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[\left| (\Delta_h^{\pi^k})(s_h^k) - (P^* \Delta_{h+1}^{\pi^k})(s_h^k, a_h^k) \right| \right] + \sqrt{\log(1/\delta) \cdot C_{m+1}} + \log(1/\delta) \\
&\lesssim 2^m G + \sqrt{\log(1/\delta) \cdot C_{m+1}} + \log(1/\delta), \tag{42}
\end{aligned}$$

where in the first inequality we change the index, the second inequality holds with probability at least $1 - \delta$ by Azuma Bernstein's inequality, the third inequality holds because that $E[X^{2^{m-1}}] \geq (E[X])^{2^{m-1}}$ for $m \geq 1$ and $X \geq 0$, the fourth inequality holds by keep using $a^2 - b^2 = (a+b)(a-b)$, then with $E[X^2] \geq E[X]^2$, and the assumption that the trajectory-wise total reward is normalized in $[0, 1]$, the last inequality holds under \mathcal{E}_1 by Equation 41, and we take a union bound to get this hold for all $m \in [0, \lceil \log_2(\frac{KH}{G}) \rceil]$ with probability at least $1 - \lceil \log_2(\frac{KH}{G}) \rceil \delta$ (because for each $m \in [0, \lceil \log_2(\frac{KH}{G}) \rceil]$ we need to apply the Azuma Bernstein's inequality once).

The above reasoning directly implies that

$$P(\mathbb{I}\{\mathcal{E}_3\}) \geq 1 - (1 + \lceil \log_2(\frac{KH}{G}) \rceil) \delta. \tag{43}$$

□

Under the event \mathcal{E}_3 , we prove the following lemma to bound $\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k} | \bar{\mathcal{P}}^k)(s_h^k, a_h^k) \right]$.

Lemma 14 (Variance Conversion Lemma for online RL). *Under event \mathcal{E}_3 , we have*

$$\begin{aligned}
&\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k} | \bar{\mathcal{P}}^k)(s_h^k, a_h^k) \right] \\
&\leq O\left(\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s_h^k, a_h^k) \right] + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH) \right).
\end{aligned}$$

Proof of Lemma 14. Under \mathcal{E}_3 , we have for any $m \in [0, \lceil \log_2(\frac{KH}{G}) \rceil]$

$$C_m \lesssim 2^m G + \sqrt{\log(1/\delta) \cdot C_{m+1}} + \log(1/\delta). \tag{44}$$

Then, by Lemma 11, we have

$$C_0 \lesssim G. \tag{45}$$

Also note that we have $A \leq 2B + 2C_0$ since $\mathbb{V}_{P^*}(a + b) \leq 2\mathbb{V}_{P^*}(a) + 2\mathbb{V}_{P^*}(b)$. Therefore, we have

$$\begin{aligned} A &\leq 2B + 2C_0 \\ &\lesssim B + G \\ &= B + \sqrt{A \cdot \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH)} \\ &\quad + \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH) \end{aligned} \quad (46)$$

Then, with the fact that $x \leq 2a + b^2$ if $x \leq a + b\sqrt{x}$, we have

$$A \leq O\left(B + \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH)\right), \quad (47)$$

which is

$$\begin{aligned} &\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k}) (s_h^k, a_h^k) \right] \\ &\leq O\left(\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k}) (s_h^k, a_h^k) \right] + \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH) \right) \end{aligned} \quad (48)$$

□

By the same reasoning in Lemma 26 of Zhou et al. (2023), we have that with probability at least $1 - \delta$

$$\begin{aligned} \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k}) (s_h^k, a_h^k) \right] &\leq \sum_{k=0}^{K-1} \sum_{h=0}^{H-1} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k}) (s_h^k, a_h^k) \right] \\ &\leq O\left(\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} + \log(1/\delta) \right). \end{aligned} \quad (49)$$

This indicates that

$$P(\mathbb{I}\{\mathcal{E}_4\}) \geq 1 - \delta. \quad (50)$$

We can use the Azuma Bernstein's inequality to get that with probability at least $1 - \delta$:

$$\sum_{k=0}^{K-1} \sum_{h=1}^H r(s_h^k, a_h^k) - \sum_{k=0}^{K-1} V_{0;P^*}^{\pi^k} \lesssim \sqrt{\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \log(1/\delta) + \log(1/\delta)}. \quad (51)$$

This indicates that

$$P(\mathbb{I}\{\mathcal{E}_5\}) \geq 1 - \delta. \quad (52)$$

Then, together with Lemma 13, Equation 31 and Equation 37, we have

$$P(\mathbb{I}\{\mathcal{E}\}) \geq 1 - (5 + \lceil \log_2(\frac{KH}{G}) \rceil) \delta \geq 1 - 5KH\delta. \quad (53)$$

Finally, under event \mathcal{E} , with all the things above (Equation 33, Equation 35, Equation 36, Lemma 12, Lemma 14), we have

$$\begin{aligned} &\sum_{k=0}^{K-1} V_{0;P^*}^{\pi^k} - \sum_{k=0}^{K-1} V_{0;P^*}^{\pi^k} \\ &= \sum_{k=0}^{K-1} V_{0;P^*}^{\pi^k} - \sum_{k=0}^{K-1} \sum_{h=1}^H r(s_h^k, a_h^k) + \sum_{k=0}^{K-1} \sum_{h=1}^H r(s_h^k, a_h^k) - \sum_{k=0}^{K-1} V_{0;P^*}^{\pi^k} \\ &\lesssim |\mathcal{K}| + \sum_{k \in [K-1] \setminus \mathcal{K}} \left(V_{0;P^*}^{\pi^k} - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \right) + \sqrt{\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \log(1/\delta) + \log(1/\delta)} \\ &\lesssim \log^2(\log(K|\mathcal{P}|/\delta)KH) \cdot \text{DE}_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(\log(K|\mathcal{P}|/\delta)KH)) + \sum_{k \in [K-1] \setminus \mathcal{K}} \left(V_{0;P^*}^{\pi^k} - \sum_{h=0}^{H-1} r(s_h^k, a_h^k) \right) \end{aligned}$$

$$\begin{aligned}
& + \sqrt{\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \log(1/\delta) + \log(1/\delta)} \\
& \lesssim \log^2(\log(K|\mathcal{P}|/\delta)KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(\log(K|\mathcal{P}|/\delta)KH)) + \log(1/\delta) \\
& + \sqrt{\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} (\mathbb{V}_{P^*} V_{h+1; \bar{P}^k}^{\pi^k})(s_h^k, a_h^k) \log(1/\delta) + \sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \log(1/\delta)} \\
& + \sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} \left| \mathbb{E}_{s' \sim \bar{P}^k(s_h^k, A^k)} V_{1; \bar{P}^k}^{\pi^k}(s') - \mathbb{E}_{s' \sim P^*(s_h^k, A^k)} V_{1; \bar{P}^k}^{\pi^k}(s') \right| \\
& \lesssim \log^2(\log(K|\mathcal{P}|/\delta)KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(\log(K|\mathcal{P}|/\delta)KH)) + \sqrt{\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \log(1/\delta) + \log(1/\delta)} \\
& + \sqrt{\left(\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} [(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s_h^k, a_h^k)] + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH) \right) \cdot \log(1/\delta)} \\
& + \sqrt{\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} [(\mathbb{V}_{P^*} V_{h+1; \bar{P}^k}^{\pi^k})(s_h^k, a_h^k)] \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH)} \\
& \lesssim \log^2(\log(K|\mathcal{P}|/\delta)KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(\log(K|\mathcal{P}|/\delta)KH)) + \sqrt{\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \log(1/\delta) + \log(1/\delta)} \\
& + \sqrt{\left(\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} [(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s_h^k, a_h^k)] + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH) \right) \cdot \log(1/\delta)} \\
& + \sqrt{\left(\sum_{k \in [K-1] \setminus \mathcal{K}} \sum_{h=0}^{H-1} [(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s_h^k, a_h^k)] + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH) \right)} \\
& \quad \times \sqrt{DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH)} \\
& \lesssim \log^2(\log(K|\mathcal{P}|/\delta)KH) \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/(\log(K|\mathcal{P}|/\delta)KH)) + \sqrt{\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \log(1/\delta) + \log(1/\delta)} \\
& + \sqrt{\left(\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} + \log(1/\delta) + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH) \right) \cdot \log(1/\delta)} \\
& + \sqrt{\left(\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} + \log\left(\frac{1}{\delta}\right) + DE_1\left(\Psi, \mathcal{S} \times \mathcal{A}, \frac{1}{KH}\right) \cdot \log\left(\frac{K|\mathcal{P}|}{\delta}\right) \log(KH) \right)} \\
& \quad \times \sqrt{DE_1\left(\Psi, \mathcal{S} \times \mathcal{A}, \frac{1}{KH}\right) \log\left(\frac{K|\mathcal{P}|}{\delta}\right) \log(KH)} \\
& \leq O\left(\sqrt{\sum_{k=0}^{K-1} \text{VaR}_{\pi^k} \cdot DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH)}\right. \\
& \quad \left. + DE_1(\Psi, \mathcal{S} \times \mathcal{A}, 1/KH) \cdot \log(K|\mathcal{P}|/\delta) \log(KH)\right). \tag{54}
\end{aligned}$$

The final result follows by replacing δ to be $\delta/(5KH)$ to make the event \mathcal{E} happen with probability at least $1 - \delta$.

D.2 PROOF OF COROLLARY 2

Proof of Corollary 2. By Lemma 8, we have

$$\text{VaR}_{\pi^k} = \sum_{h=0}^{H-1} \mathbb{E}_{s, a \sim d_h^{\pi^k}} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s, a) \right] \quad (55)$$

Therefore, when P^* is deterministic, the $\mathbb{E}_{s, a \sim d_h^{\pi^k}} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^k})(s, a) \right]$ terms are all 0 for any $k \in [K-1]$ and $h \in [H-1]$, and then the $\sum_{k=0}^{K-1} \text{VaR}_{\pi^k}$ term in the higher order term in Theorem 1 is 0. \square

D.3 PROOF OF COROLLARY 3

Proof of Corollary 3. We follow the MLE guarantee for the infinite model class in Lemma 10 and the same proof steps in the proof of Theorem 1 in Appendix D.1. \square

E DETAILED PROOFS FOR THE OFFLINE RL SETTING IN SECTION 5

E.1 PROOF OF THEOREM 2

The following is the full proof of Theorem 2.

Proof of Theorem 2. First, by the realizability assumption, the standard generalization bound for MLE (Lemma 9) with simply setting D_i to be the delta distribution on the (s_h^k, a_h^k) pairs in the offline dataset \mathcal{D} , we have that w.p. at least $1 - \delta$:

$$(1) P^* \in \widehat{\mathcal{P}};$$

$$(2)$$

$$\frac{1}{K} \sum_{k=1}^K \sum_{h=0}^{H-1} \mathbb{H}^2(P^*(s_h^k, a_h^k) \| \widehat{P}(s_h^k, a_h^k)) \leq \frac{22 \log(|\mathcal{P}|/\delta)}{K}. \quad (56)$$

Then, with the above realizability in (1), and by the pessimistic algorithm design $\hat{\pi} \leftarrow \arg\max_{\pi \in \Pi} \min_{P \in \widehat{\mathcal{P}}} V_{0;P}^{\pi}(s_0)$, $\widehat{P} \leftarrow \arg\min_{P \in \widehat{\mathcal{P}}} V_{0;P}^{\hat{\pi}}(s_0)$, we have that for any $\pi^* \in \Pi$

$$\begin{aligned} V_{0;P^*}^{\pi^*} - V_{0;\widehat{P}}^{\hat{\pi}} &= V_{0;P^*}^{\pi^*} - V_{0;\widehat{P}}^{\pi^*} + V_{0;\widehat{P}}^{\pi^*} - V_{0;\widehat{P}}^{\hat{\pi}} \\ &\leq V_{0;P^*}^{\pi^*} - V_{0;\widehat{P}}^{\pi^*} + V_{0;\widehat{P}}^{\hat{\pi}} - V_{0;P^*}^{\hat{\pi}} \\ &\leq V_{0;P^*}^{\pi^*} - V_{0;\widehat{P}}^{\pi^*}. \end{aligned} \quad (57)$$

We can then bound $V_{0;P^*}^{\pi^*} - V_{0;\widehat{P}}^{\pi^*}$ using the simulation lemma (Lemma 7):

$$V_{0;P^*}^{\pi^*} - V_{0;\widehat{P}}^{\pi^*} \leq \sum_{h=0}^{H-1} \mathbb{E}_{s, a \sim d_h^{\pi^*}} \left[\left| \mathbb{E}_{s' \sim P^*(s, a)} V_{h+1; \widehat{P}}^{\pi^*}(s') - \mathbb{E}_{s' \sim \widehat{P}(s, a)} V_{h+1; \widehat{P}}^{\pi^*}(s') \right| \right]. \quad (58)$$

Then, we prove the following lemma to bound the RHS of Equation 58.

Lemma 15 (Bound of sum of mean value differences for offline RL). *With probability at least $1 - \delta$, we have*

$$\begin{aligned} &\sum_{h=0}^{H-1} \mathbb{E}_{s, a \sim d_h^{\pi^*}} \left[\left| \mathbb{E}_{s' \sim P^*(s, a)} V_{h+1; \widehat{P}}^{\pi^*}(s') - \mathbb{E}_{s' \sim \widehat{P}(s, a)} V_{h+1; \widehat{P}}^{\pi^*}(s') \right| \right] \\ &\leq 8 \sqrt{\sum_{h=0}^{H-1} \mathbb{E}_{s, a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s, a) \right]} \cdot \frac{22C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K} + \frac{440C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}. \end{aligned}$$

1404 *Proof of Lemma 15.* We have

$$\begin{aligned}
1406 & \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\left| \mathbb{E}_{s' \sim P^*(s,a)} V_{h+1; \widehat{P}}^{\pi^*}(s') - \mathbb{E}_{s' \sim \widehat{P}(s,a)} V_{h+1; \widehat{P}}^{\pi^*}(s') \right| \right] \\
1407 & \leq 4 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\sqrt{(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s,a) D_{\Delta} \left(V_{h+1; \widehat{P}}^{\pi^*}(s' \sim P^*(s,a)) \parallel V_{h+1; \widehat{P}}^{\pi^*}(s' \sim \widehat{P}(s,a)) \right)} \right] \\
1408 & \quad + 5 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[D_{\Delta} \left(V_{h+1; \widehat{P}}^{\pi^*}(s' \sim P^*(s,a)) \parallel V_{h+1; \widehat{P}}^{\pi^*}(s' \sim \widehat{P}(s,a)) \right) \right] \\
1409 & \leq 8 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\sqrt{(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s,a) \mathbb{H}^2 \left(V_{h+1; \widehat{P}}^{\pi^*}(s' \sim P^*(s,a)) \parallel V_{h+1; \widehat{P}}^{\pi^*}(s' \sim \widehat{P}(s,a)) \right)} \right] \\
1410 & \quad + 20 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\mathbb{H}^2 \left(V_{h+1; \widehat{P}}^{\pi^*}(s' \sim P^*(s,a)) \parallel V_{h+1; \widehat{P}}^{\pi^*}(s' \sim \widehat{P}(s,a)) \right) \right] \\
1411 & \leq 8 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\sqrt{(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s,a) \mathbb{H}^2 \left(P^*(s,a) \parallel \widehat{P}(s,a) \right)} \right] + 20 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\mathbb{H}^2 \left(P^*(s,a) \parallel \widehat{P}(s,a) \right) \right] \\
1412 & \hspace{15em} (59)
\end{aligned}$$

1423 where in the first inequality, we use Lemma 1 to bound the difference of two means
1424 $\mathbb{E}_{s' \sim P^*(s,a)} V_{h+1; \widehat{P}}^{\pi^*}(s') - \mathbb{E}_{s' \sim \widehat{P}(s,a)} V_{h+1; \widehat{P}}^{\pi^*}(s')$ using variances and the triangle discrimination; in the
1425 second inequality we use the fact that that triangle discrimination is equivalent to squared Hellinger
1426 distance, i.e., $D_{\Delta} \leq 4\mathbb{H}^2$; the third inequality is via data processing inequality on the squared
1427 Hellinger distance. Next, starting from Equation 59, with probability at least $1 - \delta$, we have

$$\begin{aligned}
1428 & 8 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\sqrt{(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s,a) \mathbb{H}^2 \left(P^*(s,a) \parallel \widehat{P}(s,a) \right)} \right] + 20 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\mathbb{H}^2 \left(P^*(s,a) \parallel \widehat{P}(s,a) \right) \right] \\
1429 & \leq 8 \sqrt{\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s,a) \right] \cdot \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\mathbb{H}^2 \left(P^*(s,a) \parallel \widehat{P}(s,a) \right) \right]} \\
1430 & \quad + 20 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\mathbb{H}^2 \left(P^*(s,a) \parallel \widehat{P}(s,a) \right) \right] \\
1431 & \leq 8 \sqrt{\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s,a) \right] \cdot C^{\pi^*} \frac{1}{K} \sum_{k=1}^K \sum_{h=0}^{H-1} \mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{P}(s_h^k, a_h^k) \right)} \\
1432 & \quad + 20 C^{\pi^*} \frac{1}{K} \sum_{k=1}^K \sum_{h=0}^{H-1} \mathbb{H}^2 \left(P^*(s_h^k, a_h^k) \parallel \widehat{P}(s_h^k, a_h^k) \right) \\
1433 & \leq 8 \sqrt{\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s,a) \right] \cdot \frac{22C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K} + \frac{440C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}}, \quad (60) \\
1434 & \hspace{15em} (60)
\end{aligned}$$

1446 where the first inequality is by the Cauchy–Schwarz inequality; the second inequality is by the
1447 definition of single policy coverage (Definition 3); the last inequality holds with probability at least
1448 $1 - \delta$ with Equation 56. Substituting Equation 60 into Equation 59 ends our proof. \square

1450 We denote \mathcal{E} as the event that Lemma 15 holds. Under the event \mathcal{E} , we prove the following
1451 lemma to bound $\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s,a) \right]$ with $\widetilde{O} \left(\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^*})(s,a) \right] + \right.$
1452 $C^{\pi^*} \log(|\mathcal{P}|/\delta)/K$.

1453 **Lemma 16** (Variance Conversion Lemma for offline RL). *Under event \mathcal{E} , we have*

$$1454 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1; \widehat{P}}^{\pi^*})(s,a) \right] \leq O \left(\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^*})(s,a) \right] + C^{\pi^*} \frac{\log(|\mathcal{P}|/\delta)}{K} \right).$$

Proof of Lemma 16. For notational simplicity, we denote $A := \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^* V_{h+1}^{\pi^*}; \widehat{P}})(s, a) \right]$, and we denote $B := \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^* V_{h+1}^{\pi^*}})(s, a) \right]$, $C := \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^* (V_{h+1}^{\pi^*} - V_{h+1}^{\pi^*})}(s, a) \right]$, then we have

$$A \leq 2B + 2C,$$

since $\mathbb{V}_{P^*}(a+b) \leq 2\mathbb{V}_{P^*}(a) + 2\mathbb{V}_{P^*}(b)$.

Let $\Delta_{h+1}^{\pi^*} := V_{h+1}^{\pi^*} - V_{h+1}^{\pi^*}$. Then, w.p. at least $1 - \delta$, we have

$$\begin{aligned} C &= \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(P^* (\Delta_{h+1}^{\pi^*})^2)(s, a) - (P^* \Delta_{h+1}^{\pi^*})^2(s, a) \right] \\ &= \sum_{h=0}^{H-1} \mathbb{E}_{s \sim d_{h+1}^{\pi^*}} \left[(\Delta_{h+1}^{\pi^*})^2(s) \right] - \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(P^* \Delta_{h+1}^{\pi^*})^2(s, a) \right] \\ &\leq \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\Delta_h^{\pi^*})^2(s) - (P^* \Delta_{h+1}^{\pi^*})^2(s, a) \right] \\ &= \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\left((\Delta_h^{\pi^*})(s) + (P^* \Delta_{h+1}^{\pi^*})(s, a) \right) \cdot \left((\Delta_h^{\pi^*})(s) - (P^* \Delta_{h+1}^{\pi^*})(s, a) \right) \right], \quad (61) \end{aligned}$$

where the first equality is by the definition of variance, the second equality holds as $d_h^{\pi^*}$ is the occupancy measure also generated under P^* , the first inequality is just changing the index, the third equality holds as $a^2 - b^2 = (a+b) \cdot (a-b)$. Starting from Equation 61, we have

$$\begin{aligned} &\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\left((\Delta_h^{\pi^*})(s) + (P^* \Delta_{h+1}^{\pi^*})(s, a) \right) \cdot \left((\Delta_h^{\pi^*})(s) - (P^* \Delta_{h+1}^{\pi^*})(s, a) \right) \right] \\ &\leq 2 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\left| (\Delta_h^{\pi^*})(s) - (P^* \Delta_{h+1}^{\pi^*})(s, a) \right| \right] \\ &= 2 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\left| (V_{h;\widehat{P}}^{\pi^*})(s) - (P^* V_{h+1;\widehat{P}}^{\pi^*})(s, a) - \left((V_h^{\pi^*})(s) - (P^* V_{h+1}^{\pi^*})(s, a) \right) \right| \right] \\ &= 2 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\left| r(s, a) + (\widehat{P} V_{h+1;\widehat{P}}^{\pi^*})(s, a) - (P^* V_{h+1;\widehat{P}}^{\pi^*})(s, a) - r(s, a) \right| \right] \\ &= 2 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\left| (\widehat{P} V_{h+1;\widehat{P}}^{\pi^*})(s, a) - (P^* V_{h+1;\widehat{P}}^{\pi^*})(s, a) \right| \right], \quad (62) \end{aligned}$$

where the inequality holds as the value functions are all bounded by 1 by the assumption that the total reward over any trajectory is bounded by 1, the first equality is by the definition of $\Delta_{h+1}^{\pi^*}$, the second equality is because a is drawn from π^* . Starting from Equation 62, we have

$$\begin{aligned} &2 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\left| (\widehat{P} V_{h+1;\widehat{P}}^{\pi^*})(s, a) - (P^* V_{h+1;\widehat{P}}^{\pi^*})(s, a) \right| \right] \\ &= 2 \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[\left| \mathbb{E}_{s' \sim P^*(s,a)} \left[V_{h+1;\widehat{P}}^{\pi^*}(s') \right] - \mathbb{E}_{s' \sim \widehat{P}(\cdot|s,a)} \left[V_{h+1;\widehat{P}}^{\pi^*}(s') \right] \right| \right] \\ &\leq 16 \sqrt{\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^* V_{h+1;\widehat{P}}^{\pi^*}})(s, a) \right] \cdot \frac{22C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K} + \frac{880C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}} \\ &= 16 \sqrt{A \cdot \frac{22C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K} + \frac{880C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}} \quad (63) \end{aligned}$$

where the inequality holds with probability at least $1 - \delta$ by Lemma 15, and the second equality is by definition of A .

Then combining Equation 61, Equation 62 and Equation 63, we obtain an upper bound for C , which suggests

$$A \leq 2B + 2C$$

$$\leq 2B + \frac{1760C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K} + 32\sqrt{\frac{22C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}} \cdot \sqrt{A}.$$

Then, with the fact that $x \leq 2a + b^2$ if $x \leq a + b\sqrt{x}$, we have

$$A \leq 4B + \frac{3520C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K} + \frac{22528C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K} \leq O\left(B + \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}\right).$$

□

With the above lemmas, we can now prove the final results of Theorem 2. We have that w.p. at least $1 - \delta$

$$\begin{aligned} V_{0;P^*}^{\pi^*} - V_{0;P^*}^{\hat{\pi}} &\leq O\left(\sqrt{A \cdot \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}} + \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}\right) \\ &\leq O\left(\sqrt{\left(B + \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}\right) \cdot \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}} + \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}\right) \\ &\leq O\left(\sqrt{B \cdot \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}} + \sqrt{\frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K} \cdot \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}} + \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}\right) \\ &= O\left(\sqrt{\sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^*})(s, a) \right] \cdot \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}} + \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}\right) \\ &= O\left(\sqrt{\frac{\text{VaR}_{\pi^*} C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}} + \frac{C^{\pi^*} \log(|\mathcal{P}|/\delta)}{K}\right), \end{aligned} \tag{64}$$

where in the last equation we use Lemma 8, and $\text{VaR}_{\pi^*} := \mathbb{E} \left[\left(\sum_{h=0}^{H-1} r(s_h, \pi^*(s_h)) - V_0^{\pi^*} \right)^2 \right]$. □

E.2 PROOF OF COROLLARY 4

Proof of Corollary 4. By Lemma 8, we have

$$\text{VaR}_{\pi^*} = \sum_{h=0}^{H-1} \mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^*})(s, a) \right] \tag{65}$$

Therefore, when P^* is deterministic, the $\mathbb{E}_{s,a \sim d_h^{\pi^*}} \left[(\mathbb{V}_{P^*} V_{h+1}^{\pi^*})(s, a) \right]$ terms are all 0 for any $k \in [K-1]$ and $h \in [H-1]$, and then the VaR_{π^*} term in the higher order term in Theorem 2 is 0. □

E.3 PROOF OF COROLLARY 5

Proof of Corollary 5. This claim follows the proof of Theorem 2, while we take a different choice of β that depends on the bracketing number and follow the MLE guarantee in Lemma 10 for infinite model class. □

E.4 PROOF OF THE CLAIM IN EXAMPLE 2

Proof. Recall that in Definition 3, we have

$$C_{\mathcal{D}}^{\pi^*} := \max_{h, P \in \mathcal{P}} \frac{\mathbb{E}_{s,a \sim d_h^{\pi^*}} \mathbb{H}^2(P(s, a) \parallel P^*(s, a))}{1/K \sum_{k=1}^K \mathbb{H}^2(P(s_h^k, a_h^k) \parallel P^*(s_h^k, a_h^k))}.$$

For each step h , define two distributions, p_h, q_h , where $p_h(s, a) = d_h^{\pi^*}(s, a)$, $q_h(s, a) = \frac{1}{K} \sum_{k=1}^K \mathbb{I}\{(s, a) = (s_h^k, a_h^k)\}$, and we define $f(s, a, P) = \mathbb{H}^2(P(s, a) \parallel P^*(s, a))$, then we have

$$C_{\mathcal{D}}^{\pi^*} = \max_{h, P \in \mathcal{P}} \frac{\mathbb{E}_{s,a \sim p_h} f(s, a, P)}{\mathbb{E}_{s,a \sim q_h} f(s, a, P)}$$

$$\begin{aligned}
1566 & \\
1567 & = \max_{h, P \in \mathcal{P}} \frac{\mathbb{E}_{s, a \sim q_h} \frac{p_h(s, a)}{q_h(s, a)} f(s, a, P)}{\mathbb{E}_{s, a \sim q_h} f(s, a, P)} \\
1568 & \\
1569 & \leq \max_{h, s, a} \frac{p_h(s, a)}{q_h(s, a)} \\
1570 & \\
1571 & \leq \max_{h, s, a} \frac{1}{q_h(s, a)}. \tag{66} \\
1572 & \\
1573 &
\end{aligned}$$

1574 Note that for all h , $\{(s_h^k, a_h^k)\}_{k=1}^K$ are i.i.d. samples drawn from $d_h^{\pi^b}$, therefore, $\mathbb{E}[\mathbb{I}\{(s_h^k, a_h^k) =$
1575 $(s, a)\}] = d_h^{\pi^b}(s, a)$. By Hoeffding's inequality and with a union bound over s, a, h , and for
1576 $K \geq \frac{2 \log(\frac{|S||A|H}{\delta})}{\rho_{\min}^2}$, w.p. at least $1 - \delta$, we have
1577

$$\begin{aligned}
1578 & \\
1579 & q_h(s, a) = \frac{1}{K} \sum_{k=1}^K \mathbb{I}\{(s_h^k, a_h^k) = (s, a)\} \\
1580 & \\
1581 & \geq d_h^{\pi^b}(s, a) - \sqrt{\frac{\log(\frac{|S||A|H}{\delta})}{2K}} \\
1582 & \\
1583 & \geq \frac{d_h^{\pi^b}(s, a)}{2}, \tag{67} \\
1584 & \\
1585 & \\
1586 &
\end{aligned}$$

1587 where in the last inequality we use the assumption that $d_h^{\pi^b}(s, a) \geq \rho_{\min}, \forall s, a, h$, which gives us
1588 $K \geq \frac{2 \log(\frac{|S||A|H}{\delta})}{\rho_{\min}^2} \geq \max_{s, a, h} \frac{2 \log(\frac{|S||A|H}{\delta})}{(d_h^{\pi^b}(s, a))^2}$, so $K \geq \frac{2 \log(\frac{|S||A|H}{\delta})}{(d_h^{\pi^b}(s, a))^2}$ for any s, a, h .
1589

1590 Therefore, with $K \geq \frac{2 \log(\frac{|S||A|H}{\delta})}{\rho_{\min}^2}$, we have that w.p. at least $1 - \delta$
1591

$$1592 \\
1593 C_{\mathcal{D}}^{\pi^*} \leq \max_{h, s, a} \frac{1}{q_h(s, a)} \leq \max_{h, s, a} \frac{2}{d_h^{\pi^b}(s, a)} \leq \frac{2}{\rho_{\min}}. \tag{68} \\
1594 \\
1595$$

□

1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619