

Is Reinforcement Learning More Difficult Than Bandits? A Near-optimal Algorithm Escaping the Curse of Horizon

Zihan Zhang

Tsinghua University

ZIHAN-ZH17@MAILS.TSINGHUA.EDU.CN

Xiangyang Ji

Tsinghua University

XYJI@TSINGHUA.EDU.CN

Simon S. Du

University of Washington

SSDU@CS.WASHINGTON.EDU

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

¹Episodic reinforcement learning and contextual bandits are two widely studied sequential decision-making problems. Episodic reinforcement learning generalizes contextual bandits and is often perceived to be more difficult due to long planning horizon and unknown state-dependent transitions. The current paper shows that the long planning horizon and the unknown state-dependent transitions (at most) pose little additional difficulty on sample complexity.

We consider the episodic reinforcement learning with S states, A actions, planning horizon H , total reward bounded by 1, and the agent plays for K episodes. We propose a new algorithm, **Monotonic Value Propagation (MVP)**, which relies on a new Bernstein-type bonus. Compared to existing bonus constructions, the new bonus is tighter since it is based on a well-designed monotonic value function. In particular, the *constants* in the bonus should be subtly setting to ensure optimism and monotonicity.

We show MVP enjoys an $O\left(\left(\sqrt{SAK} + S^2A\right) \text{poly log}(SAHK)\right)$ regret, approaching the $\Omega\left(\sqrt{SAK}\right)$ lower bound of *contextual bandits* up to logarithmic terms. Notably, this result 1) *exponentially* improves the state-of-the-art polynomial-time algorithms by Dann et al. [2019] and Zanette et al. [2019] in terms of the dependency on H , and 2) *exponentially* improves the running time in [Wang et al. 2020] and significantly improves the dependency on S , A and K in sample complexity.

1. Introduction

Episodic reinforcement learning (RL) and contextual bandits (CB) are two representative sequential decision-making problems. RL is a strict generalization of CB and is often perceived to be much more difficult due to the additional two challenges that are absent in CB: 1) long planning horizon and 2) unknown state-dependent transitions. These two challenges in RL requires the agent to not only consider the immediate reward but also the possible transitions into differing states in the long run. On the other hand, one can view CB as a episodic RL problem with a horizon equal to one.² In CB, it is sufficient to act myopically by choosing the action which maximizes the immediate reward.

1. Accepted for presentation at the Conference on Learning Theory (COLT) 2021

2. See Section 2 for the precise correspondence.

Although RL and CB are widely studied in the literature, somehow surprisingly, the following fundamental problem remains open:

Does episodic reinforcement learning require more samples than contextual bandits?

Here the sample complexity is measured in terms of regret or the number of episodes to learn a near-optimal policy. To put it differently, this question asks whether the long planning horizon and/or the unknown state-dependent transitions pose additional difficulty.

Jiang and Agarwal (2018) conjectured that for tabular, episodic RL problems, under the assumption that the total reward is bounded by 1,³ there exists an $\Omega\left(\frac{SAH}{\epsilon^2}\right)$ PAC learning, or analogically, an $\Omega\left(\sqrt{SAHK}\right)$ regret lower bound, where S is the number of states, A is the number of actions, H is the planning horizon, ϵ is the target sub-optimality and K is the total number of episodes. In contrast, it is well known that for CB, one can achieve an $\tilde{O}\left(\frac{SA}{\epsilon^2}\right)$ PAC learning or an $\tilde{O}\left(\sqrt{SAK}\right)$ regret upper bound.⁴ If this conjecture is true, then there is a formal sample complexity separation between RL and CB.

However, this conjecture was recently refuted by Wang et al. (2020), who presented a new method which enjoys an $O\left(\frac{S^5 A^4 \text{poly log}(HSA/\epsilon)}{\epsilon^3}\right)$ PAC learning upper bound, the first bound that has only a *logarithmic dependency on H* . This encouraging result gives the hope: *episodic reinforcement learning is as easy as contextual bandit* in terms of the sample complexity. Furthermore, this claim would convey a conceptual message in a sense that long planning horizon and unknown state-dependent transitions pose no additional difficulty in sequential decision-making problems.

To formally establish this claim, we need to design an algorithm which enjoys an $O\left(\frac{SA}{\epsilon^2}\right)$ PAC learning and an $O\left(\sqrt{SAK}\right)$ regret upper bounds, which match the sample complexity lower bounds of CB. Ideally, we would also like this algorithm to be computationally efficient. The result in Wang et al. (2020) is still far from this grand goal, as its dependencies on S , A and ϵ are suboptimal and their algorithm runs in exponential time. See Section 3 for more discussions. Indeed, Wang et al. (2020) listed two open problems: 1) to develop an algorithm with sample complexity $\tilde{O}\left(\frac{SA}{\epsilon^2}\right)$ or regret $\tilde{O}\left(\sqrt{SAK}\right)$ and 2) to develop a polynomial-time algorithm whose sample complexity scales logarithmically with H .

1.1. Main Results

In this paper, we take an important step toward this grand goal. We design an upper confidence bound (UCB)-based algorithm, **Monotonic Value Propagation (MVP)**, which enjoys the following sample complexity bounds.

Theorem 1 *Suppose the reward is non-negative and the total reward at every episode is bounded by 1. For any $K \geq 1$ and $\delta \in (0, 1)$, we have that with probability $1 - \delta$, the regret of MVP is bounded by $\text{Regret}(K) = O\left(\left(\sqrt{SAK} + S^2 A\right) \text{poly log}(SAHK/\delta)\right)$.*

Using a standard reduction (see Section 2), we can show that we can find an ϵ -suboptimal policy in $O\left(\left(\frac{SA}{\epsilon^2} + \frac{S^2 A}{\epsilon}\right) \text{poly log}\left(\frac{SAH}{\epsilon\delta}\right)\right)$ episodes.

Our results are significant in the following senses.

3. This assumption is made in order to have a fair comparison with CB. See Section 3 for discussions.

4. Throughout the paper, $\tilde{O}(\cdot)$ omits logarithmic factors.

| Algorithm | Regret | PAC Bound | Poly Time | Non-unif. Reward | Log H |
|---|---|---|-----------|------------------|-------|
| UCBVI-BF Azar et al. (2017) | $\tilde{O}(\sqrt{SAK} + \sqrt{HK} + S^2AH)$ | $\tilde{O}\left(\frac{SA+H}{\epsilon^2} + \frac{S^2AH}{\epsilon}\right)$ | Yes | No | No |
| UBEV ⁵ Dann et al. (2017) | $\tilde{O}(\sqrt{SAH^2K} + S^2AH^2)$ | $\tilde{O}\left(\frac{SAH^2}{\epsilon^2} + \frac{S^2AH^2}{\epsilon}\right)$ | Yes | No | No |
| UCB-Q-Bernstein ⁶ Jin et al. (2018) | $\tilde{O}(\sqrt{SAH^2K} + \sqrt{S^3A^3H^3})$ | $\tilde{O}\left(\frac{SAH^2}{\epsilon^2} + \frac{(SA)^{3/2}H^3}{\epsilon}\right)$ | Yes | No | No |
| ORLC Dann et al. (2019) | $\tilde{O}(\sqrt{SAK} + S^2AH^2)$ | $\tilde{O}\left(\frac{SA}{\epsilon^2} + \frac{S^2AH^2}{\epsilon}\right)$ | Yes | No | No |
| EULER Zanette and Brunskill (2019) | $\tilde{O}(\sqrt{SAK} + S^2AH + S^{3/2}AH^{3/2})$ | $\tilde{O}\left(\frac{SA}{\epsilon^2} + \frac{S^2A+S^{3/2}AH^{3/2}}{\epsilon}\right)$ | Yes | Yes | No |
| UCBADV Zhang et al. (2020a) | $\tilde{O}(\sqrt{SAHK} + S^2A^{3/2}H^6)$ | $\tilde{O}\left(\frac{SAH}{\epsilon^2} + \frac{S^2A^{3/2}H^6}{\epsilon}\right)$ | Yes | No | No |
| Trajectory Synthesis Wang et al. (2020) | - | $\tilde{O}\left(\frac{S^5A^4}{\epsilon^3}\right)$ | No | Yes | Yes |
| MVP This Work | $\tilde{O}(\sqrt{SAK} + S^2A)$ | $\tilde{O}\left(\frac{SA}{\epsilon^2} + \frac{S^2A}{\epsilon}\right)$ | Yes | Yes | Yes |
| CB Lower Bound | $\Omega(\sqrt{SAK})$ | $\Omega\left(\frac{SA}{\epsilon^2}\right)$ | - | - | - |

Table 1: Sample complexity comparisons for state-of-the-art episodic RL algorithms. See Section 3 for discussions on this table. \tilde{O} omits logarithmic factors. **Regret** and **PAC Bound** are sample complexity measures defined in Section 2. **Non-unif. Reward**: Yes means the bound holds under Assumption 1 (allows non-uniformly bounded reward), and No means the bound only holds under Assumption 2. **Poly Time**: Whether the algorithm runs in polynomial time. **Log H**: Whether the sample complexity bound depends logarithmically on H instead of polynomially on H .

1. These bounds match the information theoretical lower bound of CB up to logarithmic factors in the regime where the number of episodes is moderately large, $K = \tilde{\Omega}(S^3A)$ or the target accuracy is moderately small, $\epsilon = \tilde{O}(1/S)$. Our result thus significantly closes the gap between RL and CB.
2. MVP is the first computationally efficient algorithm whose sample complexity scales logarithmically with H , and thus settles the second open problem raised in [Wang et al. \(2020\)](#). Comparing with the state-of-the-art computationally efficient algorithms for episodic RL, e.g., [Azar et al. \(2017\)](#); [Zanette and Brunskill \(2019\)](#); [Dann et al. \(2019\)](#); [Jin et al. \(2018\)](#); [Zhang et al. \(2020a\)](#), our algorithm enjoys an *exponential* improvement in H . Comparing with the algorithm in [Wang et al. \(2020\)](#), our algorithm is exponentially faster and achieves significantly better sample complexity in terms of S, A, ϵ . See Table 1 for more detailed comparisons.

Our algorithm and its analysis rely on the following new ideas.

5. UBEV and ORLC provide a stronger result called mistake-type PAC bounds. For more details, we refer readers to [Dann et al. \(2019\)](#).
6. The model free-algorithms UCB-Q-Bernstein and UCBADV are for the inhomogeneous setting where $P_1(\cdot|s, a), P_2(\cdot|s, a), \dots, P_H(\cdot|s, a)$ are different. This difference necessarily incurs an additional \sqrt{H} factor in the first term and an H factor in the second term in regret. It is still an open problem whether a model-free algorithm can achieve a regret bound with the leading term scales $\tilde{O}(\sqrt{SAK})$.

1. We design a new exploration bonus based on Bernstein bound to ensure optimism. The key insight is that *constants* in the bonus are crucial and helps maintain a monotonic property which helps propagates the optimism from level H to level 1. This property also leads a substantially simpler analysis than those in existing approaches.
2. A crucial step in many UCB-based algorithm, including ours, is bounding the sum of variance of estimated value function across the entire planning horizon. Our technique is to use a higher order expansion to derive a recursive inequality that relates this sum to its higher moments. Importantly, this technique does not use any type of induction from $H, H-1, \dots, 1$, which is used in most previous works and is the main technical barrier to obtain the logarithmic dependency on H .

See Section 4 and Section 5 for more technical expositions.

2. Preliminaries

Notations. Throughout this paper, we use $[N]$ to denote the set $\{1, 2, \dots, N\}$ for $N \in \mathbb{Z}_+$. We use $\mathbf{1}_s$ to denote the one-hot vector whose only non-zero element is in the s -th coordinate. For an event \mathcal{E} , we use $\mathbb{I}[\mathcal{E}]$ to denote the indicator function, i.e., $\mathbb{I}[\mathcal{E}] = 1$ if \mathcal{E} holds and $\mathbb{I}[\mathcal{E}] = 0$ otherwise. For notational convenience, we set $\iota = \ln(2/\delta)$ throughout the paper. For two n -dimensional vectors x and y , we use xy to denote $x^\top y$, use $\mathbb{V}(x, y) = \sum_i x_i y_i^2 - (\sum_i x_i y_i)^2$. In particular, when x is a probability vector, i.e., $x_i \geq 0$ and $\sum_i x_i = 1$, $\mathbb{V}(x, y) = \sum_i x_i (y_i - (\sum_i x_i y_i))^2 = \min_{\lambda \in \mathbb{R}} \sum_i x_i (y_i - \lambda)^2$. We also use x^2 to denote the vector $[x_1^2, x_2^2, \dots, x_n^2]^\top$ for $x = [x_1, x_2, \dots, x_n]^\top$. For two vectors x, y , $x \geq y$ denotes $x_i \geq y_i$ for all $i \in [n]$ and $x \leq y$ denotes $x_i \leq y_i$ for all $i \in [n]$.

Episodic Reinforcement Learning. A finite-horizon stationary Markov Decision Process (MDP) can be described by a tuple $M = (\mathcal{S}, \mathcal{A}, P, R, H, \mu)$. \mathcal{S} is the finite state space with cardinality S . \mathcal{A} is the finite action space with cardinality A . $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition operator which takes a state-action pair and returns a distribution over states. $R : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathbb{R})$ is the reward distribution with a mean function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. $H \in \mathbb{Z}_+$ is the planning horizon (episode length). $\mu \in \Delta(\mathcal{S})$ is the initial state distribution. P, R and μ are unknown.⁷ For notational convenience, we use $P_{s,a}$ and $P_{s',a,s'}$ to denote $P(\cdot|s, a)$ and $P(s'|s, a)$ respectively.

A policy π chooses an action a based on the current state $s \in \mathcal{S}$ and the time step $h \in [H]$. Note even though transition operator and the reward distribution are stationary, i.e., they do not depend on the level $h \in [H]$, the policy can be *non-stationary*, i.e., at different level h , the policy can choose different actions for the same state. Formally, we define $\pi = \{\pi_h\}_{h=1}^H$ where for each $h \in [H]$, $\pi_h : \mathcal{S} \rightarrow \mathcal{A}$ maps a given state to an action. The policy π induces a (random) trajectory $\{s_1, a_1, r_1, s_2, a_2, r_2, \dots, s_H, a_H, r_H\}$, where $s_1 \sim \mu$, $a_1 = \pi_1(s_1)$, $r_1 \sim R(s_1, a_1)$, $s_2 \sim P(\cdot|s_1, a_1)$, $a_2 = \pi_2(s_2)$, etc.

Our target is to find a policy π that maximizes the expected total reward, i.e. $\max_{\pi} \mathbb{E} \left[\sum_{h=1}^H r_h \mid \pi \right]$ where the expectation is over the initial distribution state μ , the transition operator P and the reward distribution R . As for scaling, we make the following assumption about the reward. As we will

7. Some previous works consider the non-stationary MDP where P and R can vary on different $h \in [H]$ (Jin et al., 2018; Zhang et al., 2020a). Non-stationarity will incur an \sqrt{H} factor in the regret, which is necessary. Transforming a regret bound for stationary MDP to that for non-stationary MDP is often straightforward (with an additional \sqrt{H} factor), but not vice-versa, because the main difficulty is how to effectively exploit the stationarity.

discuss in Section 3, this is a more general assumption than the assumption often made in most previous works.

Assumption 1 (Bounded Total Reward) *The reward satisfies that $r_h \geq 0$ for all $h \in [H]$. Besides, for all policy π , $\sum_{h=1}^H r_h \leq 1$ almost surely.*

Given a policy π , a level $h \in [H]$ and a state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, the Q -function is defined as: $Q_h^\pi(s, a) = \mathbb{E} \left[\sum_{h'=h}^H r_{h'} \mid s_h = s, a_h = a, \pi \right]$. Similarly, given a policy π , a level $h \in [H]$, the value function of a given state $s \in \mathcal{S}$ is defined as: $V_h^\pi(s) = \mathbb{E} \left[\sum_{h'=h}^H r_{h'} \mid s_h = s, \pi \right]$. Then Bellman equation establishes the following identities for policy π and $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$: $Q_h^\pi(s, a) = r(s, a) + P_{s,a}^\top V_{h+1}^\pi$ and $V_h^\pi(s) = \max_a Q_h^\pi(s, a)$. Throughout the paper, we let $V_{H+1}(s) = 0$ and $Q_{H+1}(s, a) = 0$ for notational simplicity. We use Q_h^* and V_h^* to denote the optimal Q -function and V -function at level $h \in [H]$, which satisfies for any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, $Q_h^*(s, a) = \max_\pi Q_h^\pi(s, a)$ and $V_h^*(s) = \max_\pi V_h^\pi(s)$.

When $H = 1$, the episodic RL reduces to the problem of finding a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes the expected reward $\max_\pi \mathbb{E}_{s \sim \mu(\cdot), r_{\text{CB}} \sim R(s, \pi(s))} [r_{\text{CB}}]$. This is called the contextual bandit (CB) problem. RL is more difficult than CB as we also need to deal with the long planning horizon H and transition operator P , which are absent in CB. In this paper, we investigate whether these two ingredients incur additional hardness in terms of the sample complexity.

Sample Complexity. In this paper we use two measures to quantify sample complexity. The agent interacts with the environment for K episodes, and it chooses a policy π^k at the k -th episode. The total regret is

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

PAC-RL sample complexity is another measure which counts the number of episodes to find an ϵ -optimal policy π , i.e.,

$$\mathbb{E}_{s_1 \sim \mu} [V_1^*(s_1) - V^\pi(s_1)] \leq \epsilon.$$

As pointed out in Jin et al. (2018), suppose that one has an algorithm that achieves $CK^{1-\alpha}$ regret for some $\alpha \in (0, 1)$ and some C independent of T , by randomly selecting from policy π^k used in K episodes, π satisfies $\mathbb{E}_{s_1 \sim \mu} [V_1^*(s_1) - V^\pi(s_1)] = O(CK^{-\alpha})$. This reduction is often near-optimal to obtain PAC-RL sample complexity guarantee. On the other hand, there is no general near-optimal reduction that transform a PAC-RL bound to a regret bound.

3. Background and Related Work

We mostly focus on papers that are for the episodic RL setting described in Section 2. A summary of the most relevant previous regret and PAC bounds, together with the results proved in this paper is provided in Table 1. We remark that there are also related settings, e.g., infinite-horizon discounted MDP, weakly-communicating MDP, learning with a generative model, etc. These settings are beyond the scope of this paper, though our techniques may be also applied to these settings.

Reward Assumption. In episodic tabular RL, the sample complexity depend on $|\mathcal{S}|$, $|\mathcal{A}|$ and H , all of which are assumed to be finite. For the reward, the widely adopted assumption is $r_h \in [0, 1]$ for all $h \in [H]$, which implies the total reward $\sum_{h=1}^H r_h \in [0, H]$. To have a fair comparison with CB and illustrate the hardness due to the planning horizon and/or unknown transition operator, one should scale down the reward by an H factor such that the total reward is bounded in $[0, 1]$.⁸ This leads to the following assumption.

Assumption 2 (Uniformly Bounded Reward) $r_h \in [0, 1/H]$ for all $h \in [H]$.

Clearly, Assumption 1 is more general than Assumption 2, so any upper bound under Assumption 1, also implies an upper bound under Assumption 2. From practical point of view, as argued in Jiang and Agarwal (2018), since environments under Assumption 1 can have one-step reward as high as a constant, Assumption 1 is more natural in environments with sparse rewards, which are often considered to be hard. From a theoretical point view, to design provably efficient algorithms under Assumption 1 is more difficult, as one needs to consider a more global structure.⁹ The sample complexity bounds in this paper hold under the more general Assumption 1.

Previous Sample Complexity Bounds. There is a long list of sample complexity guarantees for episodic tabular RL (Kearns and Singh, 2002; Brafman and Tennenholtz, 2003; Kakade, 2003; Strehl et al., 2006; Strehl and Littman, 2008; Kolter and Ng, 2009; Bartlett and Tewari, 2009; Jaksch et al., 2010; Szita and Szepesvári, 2010; Lattimore and Hutter, 2012; Osband et al., 2013; Dann and Brunskill, 2015; Azar et al., 2017; Dann et al., 2017; Osband and Van Roy, 2017; Agrawal and Jia, 2017; Jin et al., 2018; Fruit et al., 2018; Talebi and Maillard, 2018; Dann et al., 2019; Dong et al., 2019; Simchowitz and Jamieson, 2019; Russo, 2019; Zhang and Ji, 2019; Cai et al., 2019; Zhang et al., 2020a; Yang et al., 2020; Pacchiano et al., 2020; Neu and Pike-Burke, 2020). There are two popular types of algorithms, model-based algorithms and model-free algorithms. In episodic RL, model-based algorithms' space complexity scales quadratically with S and model-free algorithms and model-free algorithms' space complexity linearly with S . Both types of algorithms often rely on using UCB to ensure optimism and guide exploration. Under Assumption 2, both the state-of-the-art model-based and model-free algorithms achieve regret bounds of the form $\tilde{O}\left(\sqrt{SAK} + \text{poly}(SAH)\right)$. Recently, Zanette and Brunskill (2019) proposed a model-based algorithm which achieves the regret of the same form under Assumption 1. The first term in these bounds matches the lower bound, $\Omega\left(\sqrt{SAK}\right)$ up to logarithmic factors (Bubeck and Cesa-Bianchi, 2012; Dann and Brunskill, 2015; Osband and Roy, 2016). See Table 1 for specific bounds in these works and other related ones.

These bounds become non-trivial (regret bound sub-linear in K or PAC bound smaller than 1) only when $K \gg H$ or $\epsilon \ll \frac{1}{H}$. However, as explained in Jiang and Agarwal (2018), in many scenarios with a long planning horizon such as control, this regime is not interesting, and the more interesting regime is when $K \ll H$ or $\epsilon \gg 1/H$.

The recent work by Wang et al. (2020) bypassed this barrier via a completely different approach and obtained an $\tilde{O}\left(\frac{S^5 A^4}{\epsilon^3}\right)$ PAC-RL sample complexity bound, which is the first bound that scales

8. When comparing with existing algorithms, we also scale down their bounds by an H factor.

9. Under Assumption 1, the reward still satisfies $r_h \in [0, 1]$, so if an algorithms enjoys an sample complexity bound under Assumption 2, scaling up this bound by an H factor for regret or H^2 for PAC bound, one can obtain a bound under Assumption 1. However, this reduction is suboptimal in terms of H , so we display their original results and add a column indicating whether the bound is under Assumption 1 or Assumption 2.

logarithmically with H . They built an ϵ -net over for optimal policies and designed a simulator to evaluate all policies within the ϵ -net. However, their algorithm runs in exponential time and its sample complexity’s dependencies on S , A , ϵ are far from optimal. Furthermore, their work does not rule out the possibility that long planning horizon and/or unknown state-dependent transitions force the agent acquire more samples than CB in terms of S and A to learn a near-optimal policy.

In this work, we follow the conventional UCB-based approach. Our algorithm is computationally efficient and achieves $\tilde{O}\left(\sqrt{SAK} + S^2A\right)$ regret and $\tilde{O}\left(\frac{SA}{\epsilon^2} + \frac{S^2A}{\epsilon}\right)$ PAC-RL bound, which outperform all existing sample complexity bounds, including the additive terms. See Table 1 for more detail.

4. Main Algorithm

In the section, we introduce the **Monotonic Value Prorogation (MVP)** algorithm. The pseudo code is listed in Algorithm 1. The algorithm adopts the doubling update framework proposed in Jaksch et al. (2010). More precisely, we define a trigger set $\mathcal{L} = \{2^{i-1} | 2^{i-1} \leq KH, i = 1, 2, \dots\}$. The algorithm proceeds through epochs where each epoch ends whenever there exists a state-action pair (s, a) such that the number of visits of (s, a) falls into \mathcal{L} . In each epoch, we use the same policy induced by the current estimation of Q -function (cf. Line 9).

We update the empirical reward and transition probability of a state-action pair (s, a) only when the number of visits of (s, a) falls into \mathcal{L} . (cf. Line 14). For the transition probability, we use the standard maximum likelihood estimation. For the reward function, we only use the data collected in the current epoch to calculate the empirical reward. This will simplify the analysis and save a log factor. See Lemma 15 and its proof for more detail.

If in an episode, we update the reward and the transition probability of state-action pair, we will also update the Q -function estimation at the end of this episode. We define the bonus in Equation (1) and our optimistic estimator of Q -function in Equation (2). Note our bonus function only contains three terms. The first term and the third term correspond to the upper confidence bound of transition and the second term corresponds to the upper confidence bound of the reward. The main novelty is that by setting appropriate c_1, c_2, c_3 , the optimism can propagate from level H to level 1 without adding additional terms. We emphasize all previous results that can achieve $O\left(\sqrt{SAK}\right)$ as the first term in the regret bound (cf. Table 1) require more sophisticated bonus constructions. See Section 5 for more technical explanations.

5. Technique Overview

An optimistic algorithm needs to guarantee that (with high probability) the estimated Q -function is always an upper bound of the optimal Q -function, i.e., $Q_h(s, a) \geq Q_h^*(s, a)$ for all $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$. Note this also implies $V_h(s) \geq V_h^*(s)$.¹⁰ model-based algorithms, including ours, use the following estimator for the Q -function

$$Q_h(s, a) = \hat{r}(s, a) + \hat{P}_{s,a} V_{h+1} + b_h(s, a) \quad (3)$$

where b_h is the bonus to guarantee Q_h is an upper bound of Q^* . The main difference among algorithms is the choice of b_h . In the following, we first review existing approaches in constructing b_h

10. In this section, we drop the dependency on k for the ease of presentation.

Algorithm 1 Monotonic Value Propagation (MVP)

```
1: Input: Trigger set  $\mathcal{L} \leftarrow \{2^{i-1}|2^i \leq KH, i = 1, 2, \dots\}$ .  $c_1 = \frac{460}{9}, c_2 = 2\sqrt{2}, c_3 = \frac{544}{9}$ .
2: for  $(s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$  do
3:    $N(s, a) \leftarrow 0; \theta(s, a) \leftarrow 0; n(s, a) \leftarrow 0;$ 
4:    $N(s, a, s') \leftarrow 0; \hat{P}_{s,a,s'} \leftarrow 0, Q_h(s, a) \leftarrow 1; V_h(s) \leftarrow 1.$ 
5: end for
6: for  $k = 1, 2, \dots$  do
7:   for  $h = 1, 2, \dots, H$  do
8:     Observe  $s_h^k$ ;
9:     Take action  $a_h^k = \arg \max_a Q_h(s_h^k, a)$ ;
10:    Receive reward  $r_h^k$  and observe  $s_{h+1}^k$ .
11:    Set  $(s, a, s', r) \leftarrow (s_h^k, a_h^k, s_{h+1}^k, r_h^k)$ .
12:    Set  $N(s, a) \leftarrow N(s, a) + 1, \theta(s, a) \leftarrow \theta(s, a) + r, N(s, a, s') \leftarrow N(s, a, s') + 1.$ 
13:    \\ Update empirical reward and transition probability
14:    if  $N(s, a) \in \mathcal{L}$  then
15:      Set  $\hat{r}(s, a) \leftarrow \mathbb{I}[N(s, a) \geq 2] \frac{2\theta(s, a)}{N(s, a)} + \mathbb{I}[N(s, a) = 1] \theta(s, a)$  and  $\theta(s, a) \leftarrow 0.$ 
16:      Set  $\hat{P}_{s,a,\tilde{s}} \leftarrow N(s, a, \tilde{s})/N(s, a)$  for all  $\tilde{s} \in \mathcal{S}$ .
17:      Set  $n(s, a) \leftarrow N(s, a)$ ;
18:      Set TRIGGERED = TRUE.
19:    end if
20:  end for
21:  \\ Update Q-function
22:  if TRIGGERED then
23:    for  $h = H, H - 1, \dots, 1$  do
24:      for  $(s, a) \in \mathcal{S} \times \mathcal{A}$  do
25:        Set
```

$$b_h(s, a) \leftarrow c_1 \sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}, V_{h+1})\iota}{\max\{n(s, a), 1\}}} + c_2 \sqrt{\frac{\hat{r}(s, a)\iota}{\max\{n(s, a), 1\}}} + c_3 \frac{\iota}{\max\{n(s, a), 1\}}, \quad (1)$$

$$Q_h(s, a) \leftarrow \min\{\hat{r}(s, a) + \hat{P}_{s,a} V_{h+1} + b_h(s, a), 1\}, \quad (2)$$
$$V_h(s) \leftarrow \max_a Q_h(s, a).$$

```
26:   end for
27: end for
28:   Set TRIGGERED = FALSE
29: end if
30: end for
```

and why they failed to obtain the logarithmic dependency on H . Then we introduce our construction of b_h and the corresponding analysis to overcome the barrier.

Main Difficulty. Fix a level h . Suppose the estimator for level $h + 1$ satisfies $Q_{h+1} \geq Q_{h+1}^*$, and this implies $V_{h+1} \geq V_{h+1}^*$. Many previous optimistic algorithms use the following induction

strategy to construct the bonus for level h :

$$\begin{aligned} Q_h(s, a) &= \hat{r}(s, a) + \hat{P}_{s,a} V_{h+1} + b_h(s, a) \\ &\geq \hat{r}(s, a) + \hat{P}_{s,a} V_{h+1}^* + b_h(s, a) \end{aligned} \quad (4)$$

$$= Q_h^*(s, a) + (\hat{P}_{s,a} - P_{s,a}) V_{h+1}^* + (\hat{r}(s, a) - r(s, a)) + b_h(s, a), \quad (5)$$

where the inequality (4) follows from the induction hypothesis $V_{h+1} \geq V_{h+1}^*$ and the last equality follows from Bellman equation. To ensure optimism, existing works design $b_h(s, a)$ to be an upper bound of $(\hat{P}_{s,a} - P_{s,a}) V_{h+1}^* + (\hat{r}(s, a) - r(s, a))$ using concentration inequalities.

The tricky part is in bounding $(\hat{P}_{s,a} - P_{s,a}) V_{h+1}^*$. As discussed in Azar et al. (2017), since one does not know V_{h+1}^* , one has to replace V_{h+1}^* by its estimation V_{h+1} and introduce *additional terms* in $b_h(s, a)$ to ensure optimism. This approach has been used in all previous approaches whose regret bounds' first term is $\tilde{O}(\sqrt{SAK})$ (Azar et al., 2017; Dann et al., 2019; Zanette and Brunskill, 2019; Zhang et al., 2020a).

Unfortunately, the regret induced by the additional terms lead to (at least) a linear dependency on H because in the analyses, one needs to make $\|V_{h+1} - V_{h+1}^*\| = O(\frac{\epsilon}{H})$ so that the final error is $O(\epsilon)$ (via e.g., performance difference lemma (Kakade, 2003)). To make $\|V_{h+1} - V_{h+1}^*\| = O(\frac{\epsilon}{H})$, the sample complexity needs to scale at least linearly with H .

Technique 1: Monotonic Value Propagation. In this work, we do not go through inequality (4) in constructing the bonus. Our main strategy is to view Q_h as a function of the variable V_{h+1} (cf. Equation (3)), which we denote as $Q_h(V_{h+1})$ and we design b_h such that the function $Q_h(\cdot)$ satisfies two principles:¹¹

- **Optimism:** $Q_h(V_{h+1}^*) \geq Q_h^*$;
- **Monotonicity:** For two variables V_{h+1} and V'_{h+1} with $V_{h+1} \geq V'_{h+1}$, $Q_h(V_{h+1}) \geq Q_h(V'_{h+1})$.

If our estimation on Q function satisfies these two properties, under the induction hypothesis that $V_{h+1} \geq V_{h+1}^*$, we have

$$Q_h(V_{h+1}) \geq Q_h(V_{h+1}^*) \geq Q_h^*.$$

While the first principle, optimism, is adopted in most previous algorithms, the second monotonicity principle is new in the literature and we believe this idea can be useful in algorithm design for other RL problems.

Now we instantiate this idea. Recall our estimator defined in Equation (1)-(2)

$$Q_h(s, a) \triangleq \min \left\{ \hat{r}(s, a) + \hat{P}_{s,a} V_{h+1} + c_1 \sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}, V_{h+1})\iota}{\max\{n(s, a), 1\}}} + c_2 \sqrt{\frac{\hat{r}(s, a)\iota}{\max\{n(s, a), 1\}}} + c_3 \frac{\iota}{\max\{n(s, a), 1\}}, 1 \right\}.$$

The optimism principle can be easily implemented using empirical Bernstein inequality (see Lemma 12). For the monotonicity principle, we will carefully tune the *constants* c_1, c_2, c_3 . See Lemma 4 for more details.¹²

11. b_h can depend on V_{h+1} as well.

12. As will be clear in our proof, our actual estimator of Q -function satisfies that $Q_h \geq F_h$ for some function F_h , and F_h satisfies the two principles mentioned above. We do not discuss this subtlety in detail for the ease of presentation.

Technique 2: Bounding the Total Variance via Recursion Using a sequence of fairly standard steps in the literature, we can bound the regret by the square-root of the total variance $\sqrt{\sum_{h=1}^H \mathbb{V}(P_{s_h, a_h}, V_{h+1}^k)}$ along with some other lower order terms. To explain our high-level idea, we present analysis for the total variance in a single episode with estimated value function replaced by the true value function, i.e., $\sum_{h=1}^H \mathbb{V}(P_{s_h, a_h}, V_{h+1}^*)$

$$\begin{aligned}
\sum_{h=1}^H \mathbb{V}(P_{s_h, a_h}, V_{h+1}^*) &= \sum_{h=1}^H (P_{s_h, a_h} (V_{h+1}^*)^2 - (P_{s_h, a_h} V_{h+1}^*)^2) \\
&= \sum_{h=1}^H (P_{s_h, a_h} (V_{h+1}^*)^2 - (V_{h+1}^*(s_{h+1}))^2) + \sum_{h=1}^H ((V_h^*(s_h))^2 - (P_{s_h, a_h} V_{h+1}^*)^2) - (V_1^*(s_1))^2 \\
&\leq \sum_{h=1}^H (P_{s_h, a_h} (V_{h+1}^*)^2 - (V_{h+1}^*(s_{h+1}))^2) + 2 \sum_{h=1}^H (V_h^*(s_h) - Q_h^*(s_h, a_h)) + 2 \sum_{h=1}^H r(s_h, a_h) \\
&\leq \sum_{h=1}^H (P_{s_h, a_h} (V_{h+1}^*)^2 - (V_{h+1}^*(s_{h+1}))^2) + 2 \sum_{h=1}^H (V_h^*(s_h) - Q_h^*(s_h, a_h)) + 2 \quad (6)
\end{aligned}$$

$$\leq \tilde{O} \left(\sqrt{\sum_{h=1}^H \mathbb{V}(P_{s_h, a_h}, (V_{h+1}^*)^2)} + \sum_{h=1}^H (V_h^*(s_h) - Q_h^*(s_h, a_h)) \right). \quad (7)$$

where the first inequality we dropped $V_1^*(s_1)$, the second inequality we used the total reward is bounded by 1 and the last step holds with high probability due to a simple corollary of Freedman's inequality (Freedman, 1975) (see Lemma 13).

We can roughly view the second term in (7) as the regret in this episode. Therefore, Inequality (7) shows the total variance can be bounded by the square-root of the total variance of the *second moment* and the regret. We then apply this argument *recursively*, i.e., $m \geq 1, 2, \dots$, we can bound the total variance of the 2^m -th moment $\sum_{h=1}^H \mathbb{V}(P_{s, a}, (V_{h+1}^*)^{2^m})$ by $\sum_{h=1}^H \mathbb{V}(P_{s, a}, (V_{h+1}^*)^{2^{m+1}})$ and the regret. Also note that $\sum_{h=1}^H \mathbb{V}(P_{s, a}, (V_{h+1}^*)^{2^m})$ is bounded by H almost surely for any m .

Based on the basic lemma below, we can obtain a poly log H bound for $\sum_{h=1}^H \mathbb{V}(P_{s_h, a_h}, V_{h+1}^*)$.

Lemma 2 Let $\lambda_1, \lambda_2, \lambda_4 \geq 0$, $\lambda_3 \geq 1$ and $i' = \log_2(\lambda_1)$. Let $a_1, a_2, \dots, a_{i'}$ be non-negative reals such that $a_i \leq \lambda_1$ and $a_i \leq \lambda_2 \sqrt{a_{i+1}} + 2^{i+1} \lambda_3 + \lambda_4$ for any $1 \leq i \leq i'$. Then we have that $a_1 \leq \max\{(\lambda_2 + \sqrt{\lambda_2^2 + \lambda_4})^2, \lambda_2 \sqrt{8\lambda_3} + \lambda_4\}$.

6. Proof Sketch of Theorem 1

In this section, we present the proof sketch of Theorem 1. We first introduce a few notations: we use $Q_h^k(s, a)$, $V_h^k(s)$ and $\hat{P}_{s, a}^k$ to denote the values of $Q_h(s, a)$, $V_h(s)$ and $\hat{P}_{s, a}$ in the beginning of the k -th episode. Let $n^k(s, a)$, $b_h^k(s, a)$ and $\hat{r}_h^k(s, a)$ denote the value of $\max\{n(s, a), 1\}$, $b_h(s, a)$ and $\hat{r}(s, a)$ in (1) used for computing $Q_h^k(s, a)$. Lastly, we define $V_h^k = [V_h^k(s)]_{s \in \mathcal{S}}^T$ for convenience.

6.1. Proof of Optimism

We define \mathcal{E}_1 to be the event where

$$\left| (\hat{P}_{s,a}^k - P_{s,a}) V_{h+1}^* \right| \leq 2 \sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^*) \iota}{n^k(s, a)}} + \frac{14\iota}{3n^k(s, a)} \quad (8)$$

holds for all $(s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times [H] \times [K]$. We also define \mathcal{E}_2 be the event where

$$\left| \hat{r}_h^k(s, a) - r(s, a) \right| \leq 2 \sqrt{\frac{2\hat{r}_h^k(s, a) \iota}{n^k(s, a)}} + \frac{28\iota}{3n^k(s, a)} \quad (9)$$

holds for any possible $(s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times [H] \times [K]$.

The following lemma shows \mathcal{E}_1 and \mathcal{E}_2 hold with high probability. The analysis will be done assuming the successful event $\mathcal{E}_1 \cap \mathcal{E}_2$ holds in the rest of this section.

Lemma 3 $\mathbb{P}[\mathcal{E}_1 \cap \mathcal{E}_2] \geq 1 - 2SA(\log_2 KH + 1)\delta$.

By our exploration bonus, the Q -function is always optimistic with high probability.

Lemma 4 *Conditioned on $\mathcal{E}_1 \cap \mathcal{E}_2$, $Q_h^k(s, a) \geq Q_h^*(s, a)$ for all $(s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times [H] \times [K]$*

6.2. Bounding the Bellman Error

When the Q -function is optimistic, the major term in the regret of the induced policy is the sum of the Bellman error (see Lemma 7). So we start with a simple bound for the Bellman error induced by the Q -function.

Lemma 5 *With probability $1 - 3S^2AH(\log_2(KH) + 1)\delta$, for any $1 \leq k \leq K$, $1 \leq h \leq H$ and (s, a) , it holds that*

$$\begin{aligned} & Q_h^k(s, a) - r(s, a) - P_{s,a} V_{h+1}^k \\ & \leq \min\{2b_h^k(s, a) + c_4 \sqrt{\frac{\mathbb{V}(P_{s,a}, V_{h+1}^*) \iota}{n^k(s, a)}} + c_5 \sqrt{\frac{S\mathbb{V}(P_{s,a}, V_{h+1}^k - V_{h+1}^*) \iota}{n^k(s, a)}} + c_6 \frac{S\iota}{n^k(s, a)}, 1\} \end{aligned} \quad (10)$$

for some large enough universal constants c_4, c_5 and c_6 .

In the rest of this section, we let $\beta_h^k(s, a)$ be a shorthand of RHS of (10), i.e.,

$$\beta_h^k(s, a) := \max\{2b_h^k(s, a) + c_4 \sqrt{\frac{\mathbb{V}(P_{s,a}, V_{h+1}^*) \iota}{n^k(s, a)}} + c_5 \sqrt{\frac{S\mathbb{V}(P_{s,a}, V_{h+1}^k - V_{h+1}^*) \iota}{n^k(s, a)}} + c_6 \frac{S\iota}{n^k(s, a)}, 1\}. \quad (11)$$

We further define $\tilde{Q}_h^k(s, a) := Q_h^k(s, a) - Q_h^*(s, a)$, $\tilde{V}_h^k(s) = V_h^k(s) - V_h^*(s)$ and $\tilde{V}_h^k = [\tilde{V}_h^k(s)]_{s \in \mathcal{S}}^\top$, so by Lemma 5 and Bellman equation $Q_h^*(s, a) = r(s, a) + P_{s,a} V_{h+1}^*$, we have that with probability $1 - 3S^2AH(\log_2(KH) + 1)\delta$, for all $(s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times [H] \times [K]$:

$$\tilde{V}_h^k(s_h^k) - P_{s_h^k, a_h^k} \tilde{V}_{h+1}^k \leq \tilde{Q}_h^k(s_h^k, a_h^k) - P_{s_h^k, a_h^k} \tilde{V}_{h+1}^k \leq \beta_h^k(s, a). \quad (12)$$

6.3. Regret Analysis

Let \mathcal{K} be the set of indexes of episodes in which no update is triggered. By the update rule, it is obvious that $|\mathcal{K}^C| \leq SA(\log_2(KH) + 1)$. Let $h_0(k)$ be the first time an update is triggered in the k -th episode if there is an update in this episode and otherwise $H + 1$. Define $\mathcal{X}_0 = \{(k, h_0(k)) | k \in \mathcal{K}^C\}$ and $\mathcal{X} = \{(k, h) | k \in \mathcal{K}^C, h_0(k) + 1 \leq h \leq H\}$.

Then we define $\check{V}_h^k(s_h^k, a_h^k) = \mathbb{I}[(k, h) \notin \mathcal{X}] \cdot V_h^k(s_h^k, a_h^k)$. We also set $\check{\beta}_h^k(s_h^k, a_h^k) = \mathbb{I}[(k, h) \notin \mathcal{X}] \cdot \beta_h^k(s_h^k, a_h^k)$ and $\check{r}_h^k = \mathbb{I}[(k, h) \notin \mathcal{X}] \cdot r(s_h^k, a_h^k)$. By Lemma 5, we have that with probability $1 - 3S^2AH(\log_2(KH) + 1)\delta$,

$$\check{V}_h^k(s_h^k, a_h^k) \leq \check{r}_h^k + \check{\beta}_h^k(s_h^k, a_h^k) + P_{s,a} \check{V}_{h+1}^k, \quad (13)$$

for any $(h, k) \notin \mathcal{X}_0$ and

$$\check{V}_h^k(s_h^k, a_h^k) \leq \check{r}_h^k + \check{\beta}_h^k(s_h^k, a_h^k) + P_{s,a} \check{V}_{h+1}^k + 1, \quad (14)$$

for any $(h, k) \in \mathcal{X}_0$.

Remark 6 It is hard to analyze the regret in the episodes not in \mathcal{K} directly since $\mathbb{I}[k \in \mathcal{K}]$ is not measurable in \mathcal{F}_1^k . Instead, we introduce \mathcal{X} and analyze the regret in the steps not in \mathcal{X} because $\mathbb{I}[(k, h) \notin \mathcal{X}]$ is measurable in \mathcal{F}_h^k .

By Lemma 4 and 5, we have that

Lemma 7 With probability at least $1 - 5S^2AH(\log_2(KH) + 1)\delta$,

$$\begin{aligned} \text{Regret}(K) &:= \sum_{k=1}^K \left(V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k) \right) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} - \mathbf{I}_{s_{h+1}^k}) \check{V}_{h+1}^k + \sum_{k=1}^K \sum_{h=1}^H \check{\beta}_h^k(s_h^k, a_h^k) + \sum_{k=1}^K \left(\sum_{h=1}^H \check{r}_h^k - V_1^{\pi^k}(s_1^k) \right) + |\mathcal{K}^C|. \end{aligned} \quad (15)$$

Define $M_1 = \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} - \mathbf{I}_{s_{h+1}^k}) \check{V}_{h+1}^k$, $M_2 = \sum_{k=1}^K \sum_{h=1}^H \check{\beta}_h^k(s_h^k, a_h^k)$ and $M_3 = \sum_{k=1}^K \left(\sum_{h=1}^H \check{r}_h^k - V_1^{\pi^k}(s_1^k) \right)$. We will bound these three terms separately by the lemmas below.

Lemma 8

$$\mathbb{P} \left[|M_1| > 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, \check{V}_{h+1}^k) \iota} + 6\iota \right] \leq 2(\log_2(KH) + 1)\delta. \quad (16)$$

Lemma 9 Define $i_{\max} = \max\{i | 2^{i-1} \leq KH\} = \lfloor \log_2(KH) \rfloor + 1$. With probability $1 - (6S^2AH(\log_2(KH) + 1) + 6(\log_2(KH) + 1)\log_2(H))\delta$,

$$\begin{aligned} M_2 &\leq O \left(\sqrt{SAK i_{\max} \iota} + \sqrt{S^2 A i_{\max} \sqrt{M_2 \iota^{3/2}}} + \sqrt{SA i_{\max} K \iota} + S^2 A \iota \log_2(KH) \right) \\ &\leq O \left(\sqrt{SAK i_{\max} \iota} + S^2 A \iota \log_2(KH) \right). \end{aligned} \quad (17)$$

Lemma 10

$$\mathbb{P} \left[|M_3| > 8\sqrt{K \iota} + 6\iota \right] \leq 2(\log_2(KH) + 2)\delta. \quad (18)$$

Putting All Together By Lemma 7, 8, 9 and 10, we conclude that, with probability $1 - (10S^2AH(\log_2(KH) + 2) + 6(\log_2(KH) + 1)\log_2(KH) + 1)\delta$

$$\begin{aligned}\text{Regret}(K) &\leq M_1 + M_2 + M_3 + |\mathcal{K}^C| \\ &\leq O\left(\sqrt{SAKi_{\max}\iota} + S^2A\iota\log_2(KH) + \sqrt{K\iota} + SA(\log_2(KH) + 1)\right) \\ &= O\left(\sqrt{SAK\log_2(KH)\iota} + S^2A\iota\log_2(KH)\right).\end{aligned}$$

We finish the proof by rescaling δ .

7. Conclusion

In this paper, we gave the first computationally efficient algorithm for tabular, episodic RL whose sample complexity scales logarithmically with H . Furthermore, this algorithm matches the lower bound of a simpler problem, contextual bandits, up to logarithmic factors and an additive S^2A term. One important open problem is how to get rid of the additive S^2A term (also see discussions in Wang et al. (2020)). We remark that in the generative model setting, the optimal sample complexity does not have any additive term (Agarwal et al., 2019; Li et al., 2020).

References

- Alekh Agarwal, Sham Kakade, and Lin F Yang. On the optimality of sparse model-based planning for Markov decision processes. *arXiv preprint arXiv:1906.03804*, 2019.
- Shipra Agrawal and Randy Jia. Optimistic posterior sampling for reinforcement learning: worst-case regret bounds. In *Advances in Neural Information Processing Systems*, pages 1184–1194, 2017.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 263–272. JMLR. org, 2017.
- Peter L Bartlett and Ambuj Tewari. Regal: a regularization based algorithm for reinforcement learning in weakly communicating mdps. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI 2009)*, 2009.
- Ronen I. Brafman and Moshe Tennenholtz. R-max - a general polynomial time algorithm for near-optimal reinforcement learning. *J. Mach. Learn. Res.*, 3(Oct):213–231, March 2003. ISSN 1532-4435.
- Sebastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):QT06,1–7,9–21,23–43,45–65,67–105,107–115,117–127, 2012.
- Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. *arXiv preprint arXiv:1912.05830*, 2019.

- Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2818–2826, 2015.
- Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying PAC and regret: Uniform PAC bounds for episodic reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17, page 5717–5727, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Christoph Dann, Lihong Li, Wei Wei, and Emma Brunskill. Policy certificates: Towards accountable reinforcement learning. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1507–1516, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- Kefan Dong, Yuanhao Wang, Xiaoyu Chen, and Liwei Wang. Q-learning with ucb exploration is sample efficient for infinite-horizon mdp. *arXiv preprint arXiv:1901.09311*, 2019.
- David A Freedman. On tail probabilities for martingales. *the Annals of Probability*, 3(1):100–118, 1975.
- Ronan Fruit, Matteo Pirodda, and Alessandro Lazaric. Near optimal exploration-exploitation in non-communicating markov decision processes. In *Advances in Neural Information Processing Systems*, pages 2994–3004, 2018.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(Apr):1563–1600, 2010.
- Nan Jiang and Alekh Agarwal. Open problem: The dependence of sample complexity lower bounds on planning horizon. In *Conference On Learning Theory*, pages 3395–3398, 2018.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873, 2018.
- Sham M Kakade. *On the sample complexity of reinforcement learning*. PhD thesis, University of London London, England, 2003.
- Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2-3):209–232, 2002.
- J Zico Kolter and Andrew Y Ng. Near-bayesian exploration in polynomial time. In *Proceedings of the 26th annual international conference on machine learning*, pages 513–520, 2009.
- Tor Lattimore and Marcus Hutter. Pac bounds for discounted mdps. In *International Conference on Algorithmic Learning Theory*, pages 320–334. Springer, 2012.
- Gen Li, Yuting Wei, Yuejie Chi, Yuantao Gu, and Yuxin Chen. Breaking the sample size barrier in model-based reinforcement learning with a generative model. *arXiv preprint arXiv:2005.12900*, 2020.
- Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.

- Gergely Neu and Ciara Pike-Burke. A unifying view of optimism in episodic reinforcement learning. *arXiv preprint arXiv:2007.01891*, 2020.
- Ian Osband and Benjamin Van Roy. On lower bounds for regret in reinforcement learning. *ArXiv*, abs/1608.02732, 2016.
- Ian Osband and Benjamin Van Roy. Why is posterior sampling better than optimism for reinforcement learning? In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2701–2710. JMLR. org, 2017.
- Ian Osband, Daniel Russo, and Benjamin Van Roy. (more) efficient reinforcement learning via posterior sampling. In *Advances in Neural Information Processing Systems*, pages 3003–3011, 2013.
- Aldo Pacchiano, Philip Ball, Jack Parker-Holder, Krzysztof Choromanski, and Stephen Roberts. On optimism in model-based reinforcement learning. *arXiv preprint arXiv:2006.11911*, 2020.
- Daniel Russo. Worst-case regret bounds for exploration via randomized value functions. In *Advances in Neural Information Processing Systems*, pages 14433–14443, 2019.
- Max Simchowitz and Kevin G Jamieson. Non-asymptotic gap-dependent regret bounds for tabular mdps. In *Advances in Neural Information Processing Systems*, pages 1153–1162, 2019.
- Alexander L Strehl and Michael L Littman. An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008.
- Alexander L Strehl, Lihong Li, Eric Wiewiora, John Langford, and Michael L Littman. Pac model-free reinforcement learning. In *Proceedings of the 23rd international conference on Machine learning*, pages 881–888. ACM, 2006.
- István Szita and Csaba Szepesvári. Model-based reinforcement learning with nearly tight exploration complexity bounds. In *ICML*, 2010.
- Mohammad Sadegh Talebi and Odalric-Ambrym Maillard. Variance-aware regret bounds for undiscounted reinforcement learning in mdps. *arXiv preprint arXiv:1803.01626*, 2018.
- Ruosong Wang, Simon S Du, Lin F Yang, and Sham M Kakade. Is long horizon reinforcement learning more difficult than short horizon reinforcement learning? *arXiv preprint arXiv:2005.00527*, 2020.
- Kunhe Yang, Lin F Yang, and Simon S Du. Q -learning with logarithmic regret. *arXiv preprint arXiv:2006.09118*, 2020.
- Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pages 7304–7312, 2019.
- Zihan Zhang and Xiangyang Ji. Regret minimization for reinforcement learning by evaluating the optimal bias function. In *Advances in Neural Information Processing Systems*, pages 2823–2832, 2019.

Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. *arXiv preprint arXiv:2004.10019*, 2020a.

Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Model-free reinforcement learning: from clipped pseudo-regret to sample complexity. *arXiv preprint arXiv:2006.03864*, 2020b.

Appendix A. Technical Lemmas

Lemma 11 (Bennet's Inequality) *Let Z, Z_1, \dots, Z_n be i.i.d. random variables with values in $[0, 1]$ and let $\delta > 0$. Define $\mathbb{V}Z = \mathbb{E}[(Z - \mathbb{E}Z)^2]$. Then we have*

$$\mathbb{P} \left[\left| \mathbb{E}[Z] - \frac{1}{n} \sum_{i=1}^n Z_i \right| > \sqrt{\frac{2\mathbb{V}Z \ln(2/\delta)}{n}} + \frac{\ln(2/\delta)}{n} \right] \leq \delta.$$

Lemma 12 (Theorem 4 in Maurer and Pontil (2009)) *Let Z, Z_1, \dots, Z_n ($n \geq 2$) be i.i.d. random variables with values in $[0, 1]$ and let $\delta > 0$. Define $\bar{Z} = \frac{1}{n} \sum_{i=1}^n Z_i$ and $\hat{V}_n = \frac{1}{n} \sum_{i=1}^n (Z_i - \bar{Z})^2$. Then we have*

$$\mathbb{P} \left[\left| \mathbb{E}[Z] - \frac{1}{n} \sum_{i=1}^n Z_i \right| > \sqrt{\frac{2\hat{V}_n \ln(2/\delta)}{n-1}} + \frac{7 \ln(2/\delta)}{3(n-1)} \right] \leq \delta.$$

Lemma 13 (Lemma 10 in Zhang et al. (2020b)) *Let $(M_n)_{n \geq 0}$ be a martingale such that $M_0 = 0$ and $|M_n - M_{n-1}| \leq c$ for some $c > 0$ and any $n \geq 1$. Let $\text{Var}_n = \sum_{k=1}^n \mathbb{E}[(M_k - M_{k-1})^2 | \mathcal{F}_{k-1}]$ for $n \geq 0$, where $\mathcal{F}_k = \sigma(M_1, \dots, M_k)$. Then for any positive integer n , and any $\epsilon, \delta > 0$, we have that*

$$\mathbb{P} \left[|M_n| \geq 2\sqrt{2} \sqrt{\text{Var}_n \ln(1/\delta)} + 2\sqrt{\epsilon \ln(1/\delta)} + 2c \ln(1/\delta) \right] \leq 2(\log_2(\frac{nc^2}{\epsilon}) + 1)\delta.$$

Lemma 2 [Restatement] *Let $\lambda_1, \lambda_2, \lambda_4 \geq 0$, $\lambda_3 \geq 1$ and $i' = \log_2(\lambda_1)$. Let $a_1, a_2, \dots, a_{i'}$ be non-negative reals such that $a_i \leq \lambda_1$ and $a_i \leq \lambda_2 \sqrt{a_{i+1} + 2^{i+1}\lambda_3} + \lambda_4$ for any $1 \leq i \leq i'$. Then we have that $a_1 \leq \max\{(\lambda_2 + \sqrt{\lambda_2^2 + \lambda_4})^2, \lambda_2 \sqrt{8\lambda_3} + \lambda_4\}$.*

Proof Let i_0 be the least integer such that $2^{i_0}\lambda_3 > \lambda_1$ and $i_1 = \max\{i | i \leq i_0, a_i > 2^i\lambda_3\} \cup \{0\}$. Because $\lambda_3 \geq 1$, $i_0 \leq i'$. If $i_1 \leq 1$, then we have $a_2 \leq 4\lambda_3$. Otherwise, by definition, we have

$$2^{i_1}\lambda_3 < a_{i_1} \leq \lambda_2 \sqrt{a_{i_1+1} + 2^{i_1+1}\lambda_3} + \lambda_4 \leq \lambda_2 2^{\frac{i_1+2}{2}} \sqrt{\lambda_3} + \lambda_4,$$

which implies that $(\sqrt{2^{i_1}\lambda_3})^2 < 2\lambda_2 \sqrt{2^{i_1}\lambda_3} + \lambda_4$, and thus

$$2^{i_1}\lambda_3 < a_{i_1} < \bar{a} := (\lambda_2 + \sqrt{\lambda_2^2 + \lambda_4})^2.$$

For $1 \leq i < i_1$, we have that

$$a_i < \lambda_2 \sqrt{a_{i+1} + \bar{a}} + \lambda_4.$$

Because $a_{i_1} < \bar{a}$, we have $a_{i_1-1} < \lambda_2 \sqrt{2\bar{a}} + \lambda_4 \leq \bar{a}$. By induction, we have that $a_2 < \bar{a}$. Therefore, $a_2 \leq \max\{\bar{a}, 4\lambda_3\}$ and $a_1 \leq \max\{\bar{a}, \lambda_2 \sqrt{8\lambda_3} + \lambda_4\}$. The proof is completed. \blacksquare

Appendix B. Missing Proofs in Section 6.1

B.1. Proof of Lemma 3

Proof Next, we will show that \mathcal{E}_1 and \mathcal{E}_2 hold with high probability. For each (s, a) , when $n^k(s, a) = 1$ or 2, (8) and (9) hold trivially. For $n^k(s, a) = 2^i$ with $i \geq 2$, by Lemma 12, we have that

$$\begin{aligned} & \mathbb{P} \left[|(\hat{P}_{s,a}^k - P_{s,a})V_{h+1}^*| > 2\sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^*)\iota}{n^k(s, a)}} + \frac{14\iota}{3n^k(s, a)} \right] \\ & \leq \mathbb{P} \left[|(\hat{P}_{s,a}^k - P_{s,a})V_{h+1}^*| > \sqrt{\frac{2\mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^*)\iota}{n^k(s, a) - 1}} + \frac{7\iota}{3n^k(s, a) - 1} \right] \\ & \leq \delta \end{aligned} \tag{19}$$

and

$$\begin{aligned} & \mathbb{P} \left[|\hat{r}_h^k(s, a) - r(s, a)| > 2\sqrt{\frac{2\hat{r}_h^k(s, a)\iota}{n^k(s, a)}} + \frac{28\iota}{3n^k(s, a)} \right] \\ & \leq \mathbb{P} \left[|\hat{r}_h^k(s, a) - r(s, a)| > 2\sqrt{\frac{\hat{\text{Var}}_h^k(s, a)\iota}{n^k(s, a) - 1}} + \frac{14\iota}{3(n^k(s, a) - 1)} \right] \\ & \leq \delta, \end{aligned} \tag{20}$$

where $\hat{\text{Var}}_h^k(s, a) \leq \hat{r}_h^k(s, a)$ ¹³ is the empirical variance of $R(s, a)$ computed by the $n^k(s, a)$ samples. Via a union bound over all (s, a) and i , we obtain that $\mathbb{P}[\mathcal{E}_1 \cap \mathcal{E}_2] \geq 1 - 2SA(\log_2 KH + 1)\delta$. The proof is completed. \blacksquare

B.2. Proof of Lemma 4

Proof The proof of the two principles, optimism and monotonicity rely on exploiting the properties of the following f defined in the following lemma.

Lemma 14 Let $f : \Delta^S \times \mathbb{R}^S \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ with $f(p, v, n, \iota) = pv + \max \left\{ \bar{c}_1 \sqrt{\frac{\mathbb{V}(p, v)\iota}{n}}, \bar{c}_2 \frac{\iota}{n} \right\}$ with $\bar{c}_1 = \frac{20}{3}$ and $\bar{c}_2 = \frac{400}{9}$. Then f satisfies

1. $f(p, v, n, \iota)$ is non-decreasing in $v(s)$ for all $p \in \Delta^S, \|v\|_\infty \leq 1$ and $n, \iota > 0$;

2. $f(p, v, n, \iota) \geq pv + 2\sqrt{\frac{\mathbb{V}(p, v)\iota}{n}} + \frac{14\iota}{3n}$ for all p, v and $n, \iota > 0$.

13. $\mathbb{E}[(Z - \mathbb{E}[Z])^2] \leq \mathbb{E}[Z]$ for $Z \in [0, 1]$.

The proof of the lemma is straightforward. Note the first property is exactly the monotonicity we want.

Proof To verify the first claim, we fix all other variables but $v(s)$ and view f as a function in $v(s)$. Because the derivative of f in $v(s)$ does not exist only when $c_1 \sqrt{\frac{\mathbb{V}(p,v)\iota}{n}} = c_2 \frac{\iota}{n}$, where the condition has at most two solutions, so it suffices to prove $\frac{\partial f}{\partial v(s)} \geq 0$ when $c_1 \sqrt{\frac{\mathbb{V}(p,v)\iota}{n}} \neq c_2 \frac{\iota}{n}$. Direct computation gives that

$$\begin{aligned} \frac{\partial f}{\partial v(s)} &= p(s) + c_1 \mathbb{I} \left[c_1 \sqrt{\frac{\mathbb{V}(p,v)\iota}{n}} \geq c_2 \frac{\iota}{n} \right] \frac{p(s)(v(s) - pv)\iota}{\sqrt{n\mathbb{V}(p,v)\iota}} \\ &\geq \min\{p(s) + \frac{c_1^2}{c_2} p(s)(v(s) - pv), p(s)\} \\ &\geq p(s) \left(1 - \frac{c_1^2}{c_2}\right) \\ &= 0. \end{aligned} \tag{21}$$

The second claim holds because both $\sqrt{\frac{\mathbb{V}(p,v)\iota}{n}}$ and $\frac{\iota}{n}$ are non-negative. \blacksquare

Recall we chose $c_1 = \frac{460}{9}$, $c_2 = 2\sqrt{2}$ and $c_3 = \frac{544}{9}$. Now we prove $Q_h^k(s, a) \geq Q_h^*(s, a)$ by backward induction conditioned on the event \mathcal{E}_1 and \mathcal{E}_2 hold. Firstly, the conclusion holds for $h = H + 1$ because $Q_{H+1}^* = 0$. For $1 \leq h \leq H$, assuming the conclusion holds for $h + 1$, by (2), we have that

$$\begin{aligned} &Q_h^k(s, a) \\ &= \min\{\hat{r}_h^k(s, a) + \hat{P}_{s,a}^k V_{h+1}^k + b_h^k(s, a), 1\} \\ &\geq \min\{\hat{r}_h^k(s, a) + \hat{P}_{s,a}^k V_{h+1}^k + b_h^k(s, a), Q_h^*(s, a)\} \\ &\geq \min\{\hat{r}_h^k(s, a) + \hat{P}_{s,a}^k V_{h+1}^k + c_1 \sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^k)\iota}{n^k(s, a)}} + c_2 \sqrt{\frac{\hat{r}(s, a)\iota}{n^k(s, a)}} + c_3 \frac{\iota}{n^k(s, a)}, Q_h^*(s, a)\} \end{aligned} \tag{22}$$

$$\geq \min\{r(s, a) + \hat{P}_{s,a}^k V_{h+1}^k + \max\{\bar{c}_1 \sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^k)\iota}{n^k(s, a)}}, \bar{c}_2 \frac{\iota}{n^k(s, a)}\}, Q_h^*(s, a)\} \tag{23}$$

$$\geq \min\{r(s, a) + \hat{P}_{s,a}^k V_{h+1}^* + \max\{\bar{c}_1 \sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^*)\iota}{n^k(s, a)}}, \bar{c}_2 \frac{\iota}{n^k(s, a)}\}, Q_h^*(s, a)\} \tag{24}$$

$$\geq \min\{r(s, a) + \hat{P}_{s,a}^k V_{h+1}^* + 2\sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^*)\iota}{n^k(s, a)}} + \frac{14\iota}{3n^k(s, a)}, Q_h^*(s, a)\} \tag{25}$$

$$\geq \min\{r(s, a) + P_{s,a} V_{h+1}^*, Q_h^*(s, a)\} \tag{26}$$

$$= Q_h^*(s, a).$$

(23) is by the definition of $b_h^k(s, a)$ and $n^k(s, a)$. (24) is by the definition of \mathcal{E}_2 and our choice of

c_1, c_2, c_3 and \bar{c}_1, \bar{c}_2 . (25) is by recognizing $f(\hat{P}_{s,a}^k, V_{h+1}^k, n^k(s, a), \iota) = \hat{P}_{s,a}^k V_{h+1}^k + \max\{\bar{c}_1 \sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^k)\iota}{n^k(s, a)}}, \bar{c}_2 \frac{\iota}{n^k(s, a)}\}$,

then using the first property in Lemma 14 and the induction that $V_{h+1}^k \geq V_{h+1}^*$, (26) is by the second property of Lemma 14 and the definition of \mathcal{E}_1 . ■

Appendix C. Missing Proofs in Section 6.2

C.1. Proof of Lemma 5

Proof It suffices to verify (10) for the first term in **RHS**. Under $\mathcal{E}_1 \cap \mathcal{E}_2$, we have that with probability $1 - SAH(\log_2(KH) + 1)\delta$, for all $(s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times [H] \times [K]$:

$$Q_h^k(s, a) - r(s, a) - P_{s,a}V_{h+1}^k \quad (28)$$

$$\begin{aligned} &\leq \hat{r}_h^k(s, a) - r(s, a) + b_h^k(s, a) + (\hat{P}_{s,a}^k - P_{s,a})(V_{h+1}^k - V_{h+1}^*) + (\hat{P}_{s,a}^k - P_{s,a})V_{h+1}^* \\ &\leq 2b_h^k(s, a) + (\hat{P}_{s,a}^k - P_{s,a})(V_{h+1}^k - V_{h+1}^*) + (\hat{P}_{s,a}^k - P_{s,a})V_{h+1}^*. \end{aligned} \quad (29)$$

Fix s, a, h, k . When $n^k(s, a) = 1$, (10) holds trivially. For $n^k(s, a) = 2^i$ with $i \geq 1$, by Bennet's inequality (see Lemma 11) we have that for each s'

$$\mathbb{P} \left[|\hat{P}_{s,a,s'}^k - P_{s,a,s'}| > \sqrt{\frac{2P_{s,a,s'}\iota}{n^k(s, a)}} + \frac{\iota}{3n^k(s, a)} \right] \leq \delta.$$

So with probability $1 - S\delta$, we have that

$$(\hat{P}_{s,a}^k - P_{s,a})(V_{h+1}^k - V_{h+1}^*) = \sum_{s'} (\hat{P}_{s,a,s'}^k - P_{s,a,s'})(V_{h+1}^k(s') - V_{h+1}^*(s') - P_{s,a}(V_{h+1}^k - V_{h+1}^*)) \quad (30)$$

$$\begin{aligned} &\leq \sum_{s'} \sqrt{\frac{2P_{s,a,s'}\iota}{n^k(s, a)}} |V_{h+1}^k(s') - V_{h+1}^*(s') - P_{s,a}(V_{h+1}^k - V_{h+1}^*)| + \frac{S\iota}{3n^k(s, a)} \\ &\leq \sqrt{\frac{2S\mathbb{V}(P_{s,a}, V_{h+1}^k - V_{h+1}^*)}{n^k(s, a)}} + \frac{S\iota}{3n^k(s, a)}, \end{aligned} \quad (31)$$

where (30) holds because $\sum_{s'} \hat{P}_{s,a,s'}^k = \sum_{s'} P_{s,a,s'} = 1$ and (31) holds by Cauchy-Schwartz inequality. On the other hand, by Bennet's inequality (see Lemma 11) again, we obtain that

$$\mathbb{P} \left[|(\hat{P}_{s,a}^k - P_{s,a})V_{h+1}^*| > \sqrt{\frac{2\mathbb{V}(P_{s,a}, V_{h+1}^*)\iota}{n^k(s, a)}} + \frac{\iota}{3n^k(s, a)} \right] \leq \delta. \quad (32)$$

Combining (29), (31) and (32) and via a union bound over k, h, s, a , we conclude that (10) holds with probability $1 - 3S^2AH(\log_2(KH) + 1)\delta$, and with $c_4 = \sqrt{2}$, $c_5 = \sqrt{2}$ and $c_6 = \frac{2}{3}$. ■

Appendix D. Missing Proofs in Section 6.3

D.1. Proof of Lemma 7

Proof Direct computation gives that

$$\begin{aligned}
\text{Regret}(K) &:= \sum_{k=1}^K \left(V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k) \right) \\
&\leq \sum_{k=1}^K \left(V_1^k(s_1^k) - V_1^{\pi^k}(s_1^k) \right) \\
&= \sum_{k=1}^K \left(\check{V}_1^k(s_1^k) - V_1^{\pi^k}(s_1^k) \right) \\
&= \sum_{k=1}^K \left(\check{V}_1^k(s_1^k) - \sum_{h=1}^H \check{r}_h^k \right) + \sum_{k=1}^K \left(\sum_{h=1}^H \check{r}_h^k - V_1^{\pi^k}(s_1^k) \right) \\
&= \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} - \mathbf{1}_{s_{h+1}^k}) \check{V}_{h+1}^k + \sum_{k=1}^K \sum_{h=1}^H (\check{V}_h^k(s_h^k) - \check{r}_h^k - P_{s_h^k, a_h^k} \check{V}_{h+1}^k) + \sum_{k=1}^K \left(\sum_{h=1}^H \check{r}_h^k - V_1^{\pi^k}(s_1^k) \right) \\
&\leq \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} - \mathbf{1}_{s_{h+1}^k}) \check{V}_{h+1}^k + \sum_{k=1}^K \sum_{h=1}^H \check{\beta}_h^k(s_h^k, a_h^k) + \sum_{k=1}^K \left(\sum_{h=1}^H \check{r}_h^k - V_1^{\pi^k}(s_1^k) \right) + |\mathcal{K}^C|.
\end{aligned} \tag{33}$$

Here the first inequality is due to our optimistic estimation of Q -function, and (33) holds by (13) and (14). ■

D.2. Proof of Lemma 8

Proof We note that M_1 could be viewed as a martingale because \check{V}_{h+1}^k is measurable with respect to \mathcal{F}_h^k where $\mathcal{F}_h^k = \sigma \left(\{s_{h'}^{k'}, a_{h'}^{k'}, r_{h'}^{k'}, s_{h'+1}^{k'}\}_{1 \leq k' < k, 1 \leq h' \leq H} \cup \{s_{h'}^k, a_{h'}^k, r_{h'}^k\}_{1 \leq h' \leq h-1} \cup \{s_h^k, a_h^k\} \right)$, i.e., all past trajectories before s_{h+1}^k is rolled out. To avoid polynomial dependence on H , we use a variance-dependent concentration inequality to bound this term instead of Hoeffding inequality (see Lemma 13). By Lemma 13 with $\epsilon = 1$, we have that

$$\mathbb{P} \left[|M_1| > 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, \check{V}_{h+1}^k) \iota} + 6\iota \right] \leq 2(\log_2(KH) + 1)\delta. \tag{34}$$

To bound M_1 , it suffices to bound $M_4 := \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, \check{V}_{h+1}^k)$. We will deal with this term later. ■

D.3. Proof of Lemma 9

Proof Recall that

$$\beta_h^k(s, a) = O \left(\sqrt{\frac{\mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^k \iota)}{n^k(s, a)}} + \sqrt{\frac{\mathbb{V}(P_{s,a}, V_{h+1}^*)}{n^k(s, a)}} + \sqrt{\frac{S\mathbb{V}(P_{s,a}, V_{h+1}^k - V_{h+1}^*)\iota}{n^k(s, a)}} + \sqrt{\frac{\hat{r}_h^k(s, a)\iota}{n^k(s, a)}} + \frac{S\iota}{n^k(s, a)} \right).$$

By Lemma 12, we have

$$\mathbb{P} \left[\hat{P}_{s,a,s'}^k > \frac{3}{2}P_{s,a,s'} + \frac{4\iota}{3n^k(s, a)} \right] \leq \mathbb{P} \left[\hat{P}_{s,a,s'}^k - P_{s,a,s'} > \sqrt{\frac{2P_{s,a,s'}\iota}{n^k(s, a)}} + \frac{\iota}{3n^k(s, a)} \right] \leq \delta, \quad (35)$$

which implies that, with probability $1 - 2S^2AH(\log_2(KH) + 1)\delta$, it holds that for each k, h

$$\begin{aligned} \mathbb{V}(\hat{P}_{s,a}^k, V_{h+1}^k) &= \sum_{s'} \hat{P}_{s,a,s'}^k \left(V_{h+1}^k(s') - \hat{P}_{s,a}^k V_{h+1}^k \right)^2 \\ &\leq \sum_{s'} \hat{P}_{s,a,s'}^k \left(V_{h+1}^k(s') - P_{s,a} V_{h+1}^k \right)^2 \\ &\leq \sum_{s'} \left(\frac{3}{2}P_{s,a,s'} + \frac{4\iota}{3n^k(s, a)} \right) \cdot \left(V_{h+1}^k(s') - P_{s,a} V_{h+1}^k \right)^2 \\ &\leq \frac{3}{2}\mathbb{V}(P_{s,a}, V_{h+1}^k) + \frac{4S\iota}{3n^k(s, a)}. \end{aligned}$$

Note that $\mathbb{V}(P, X + Y) \leq 2(\mathbb{V}(P, X) + \mathbb{V}(P, Y))$ for any P, X, Y , we then have

$$\beta_h^k(s, a) \leq O \left(\sqrt{\frac{\mathbb{V}(P_{s,a}, V_{h+1}^k \iota)}{n^k(s, a)}} + \sqrt{\frac{S\mathbb{V}(P_{s,a}, V_{h+1}^k - V_{h+1}^*)\iota}{n^k(s, a)}} + \sqrt{\frac{\hat{r}_h^k(s, a)\iota}{n^k(s, a)}} + \frac{S\iota}{n^k(s, a)} \right). \quad (36)$$

Note that under the doubling epoch update framework, despite those episodes in which an update is triggered, the number of visits of (s, a) between the i -th update of $\hat{P}_{s,a}$ and the $i + 1$ -th update of $\hat{P}_{s,a}$ do not exceeds 2^{i-1} . More precisely, recalling the definition of \mathcal{K} , for any (s, a) and any $i \geq 3$, we have

$$\sum_{k=1}^H \sum_{h=1}^H \mathbb{I} \left[(s_h^k, a_h^k) = (s, a), n^k(s, a) = 2^{i-1} \right] \cdot \mathbb{I}[(k, h) \notin \mathcal{K}] \leq 2^{i-1}. \quad (37)$$

Recall $i_{\max} = \max\{i | 2^{i-1} \leq KH\} = \lfloor \log_2(KH) \rfloor + 1$. To facilitate the analysis, we first derive a general deterministic result. Let $w = \{w_h^k \geq 0 | 1 \leq h \leq H, 1 \leq k \leq K\}$ be a group of non-negative weights such that $w_h^k \leq 1$ for any $(k, h) \in [H] \times [K]$ and $w_h^k = 0$ for any $(k, h) \in \mathcal{K}$. Later we will set w_h^k to be the products of $\mathbb{I}[(k, h) \notin \mathcal{K}]$ with $\hat{r}_h^k(s_h^k, a_h^k)$, $\mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^*)$, $\mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k)$ and $\mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k - V_{h+1}^*)$.

We can calculate

$$\sum_{k=1}^K \sum_{h=1}^H \sqrt{\frac{w_h^k}{n^k(s_h^k, a_h^k)}}$$

$$\begin{aligned}
&\leq \sum_{k=1}^K \sum_{h=1}^H \sum_{s,a} \sum_{i=3}^{i_{\max}} \mathbb{I}[(s_h^k, a_h^k) = (s, a), n^k(s, a) = 2^{i-1}] \sqrt{\frac{w_h^k}{2^{i-1}}} + 8SA(\log_2(KH) + 4) \\
&= \sum_{s,a} \sum_{i=3}^{i_{\max}} \frac{1}{\sqrt{2^{i-1}}} \sum_{k=1}^K \sum_{h=1}^H \mathbb{I}[(s_h^k, a_h^k) = (s, a), n^k(s, a) = 2^{i-1}] \sqrt{w_h^k} + 8SA(\log_2(KH) + 4) \\
&\leq \sum_{s,a} \sum_{i=3}^{i_{\max}} \sqrt{\frac{\sum_{k=1}^K \sum_{h=1}^H \mathbb{I}[(s_h^k, a_h^k) = (s, a), n^k(s, a) = 2^{i-1}]}{2^{i-1}}}. \\
&\quad \sqrt{\left(\sum_{k=1}^K \sum_{h=1}^H \mathbb{I}[(s_h^k, a_h^k) = (s, a), n^k(s, a) = 2^{i-1}] w_h^k \right) + 8SA(\log_2(KH) + 4)} \quad (38)
\end{aligned}$$

$$\leq \sqrt{SAi_{\max} \sum_{k=1}^K \sum_{h=1}^H w_h^k + 8SA(\log_2(KH) + 4)}. \quad (39)$$

Here (38) is by Cauchy-Schwarz inequality and (39) is by (37) and Cauchy-Schwarz inequality.

Let $I(k, h)$ be shorthand of $\mathbb{I}[(k, h) \notin \mathcal{X}]$. It is worth noting that by definition, $\sum_{k=1}^K \sum_{h=1}^H |I(k, h) - I(k, h+1)| \leq |\mathcal{K}^C|$. By plugging respectively $w_h^k = I(k, h)\hat{r}_h^k(s_h^k, a_h^k)$, $I(k, h)\mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^*)$, $I(k, h)\mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k)$ and $I(k, h)\mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k - V_{h+1}^*)$ into (39), and recalling (36), we obtain that

$$\begin{aligned}
M_2 &= \sum_{k=1}^K \sum_{h=1}^H \check{\beta}_h^k(s_h^k, a_h^k) \\
&= \sum_{k=1}^K \sum_{h=1}^H \beta_h^k(s_h^k, a_h^k) I(k, h) \quad (40)
\end{aligned}$$

$$\begin{aligned}
&\leq O \left(\sqrt{SAi_{\max}\iota \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k) I(k, h)} + \sqrt{S^2 Ai_{\max}\iota \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k - V_{h+1}^*) I(k, h)} \right) \\
&\quad + O \left(\sqrt{SAi_{\max} \sum_{k=1}^K \sum_{h=1}^H \hat{r}_h^k(s_h^k, a_h^k) I(k, h)\iota} + S^2 A\iota \log_2(KH) \right) \quad (41)
\end{aligned}$$

$$\begin{aligned}
&\leq O \left(\sqrt{SAi_{\max}\iota \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k) I(k, h)} + \sqrt{S^2 Ai_{\max}\iota \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k - V_{h+1}^*) I(k, h)} \right) \\
&\quad + O \left(\sqrt{SAi_{\max} K\iota} + S^2 A\iota \log_2(KH) \right) \quad (42)
\end{aligned}$$

where in (42), we used the following lemma whose proof is deferred to appendix.

Lemma 15 $\sum_{k=1}^K \sum_{h=1}^H \hat{r}_h^k(s_h^k, a_h^k) I(k, h) \leq 2 \sum_{k=1}^K \sum_{h=1}^H r_h^k + 4SA \leq 2K + 4SA$.

Proof For any (k, h) and (k', h') , we define

$$\tilde{w}_h^k(k', h') = \frac{1}{n^k(s_h^k, a_h^k)} \mathbb{I}[(s_h^k, a_h^k) = (s_{h'}^{k'}, a_{h'}^{k'})] \cdot \mathbb{I}[n^{k'}(s_{h'}^{k'}, a_{h'}^{k'}) = 2n^k(s_h^k, a_h^k)] \cdot I(k', h').$$

By the update rule, for each (k', h') pair with $n^{k'}(s_{h'}^{k'}, a_{h'}^{k'}) \geq 2$, $\hat{r}_{h'}^{k'}(s_{h'}^{k'}, a_{h'}^{k'}) = \sum_{k=1}^K \sum_{h=1}^H \tilde{w}_h^k(h', k') r_h^k$. On the other hand, because $\tilde{w}_h^k(h', k') \leq \frac{1}{n^k(s_h^k, a_h^k)}$ for any (k', h') , and $\sum_{k'=1}^K \sum_{h'=1}^H \mathbb{I}[\tilde{w}_h^k(h', k') > 0] \leq 2n^k(s_h^k, a_h^k)$, we have

$$\sum_{k'=1}^K \sum_{h'=1}^H \tilde{w}_h^k(h', k') \leq 2.$$

Therefore, we have

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \hat{r}_h^k(s_h^k, a_h^k) &\leq \sum_{k \in \mathcal{K}} \sum_{h=1}^H \mathbb{I}[n^k(s_h^k, a_h^k) \geq 2] \hat{r}_h^k(s_h^k, a_h^k) + 4SA \\ &\leq 2 \sum_{k=1}^K \sum_{h=1}^H r_h^k + 4SA \\ &\leq 2K + 4SA. \end{aligned}$$

We remark that if we use the standard maximum likelihood estimation, the weight of the a reward would be $1 \cdot 1 + 2 \cdot \frac{1}{2} + 4 \cdot \frac{1}{4} + \dots \approx \log(T)$. However, if we update the empirical reward using the latest half fraction of samples, the weight for each reward is only $2^{i+1} \frac{1}{2^i} \leq 2$. Therefore, we can save a $\log(T)$ factor. ■

Recalling the definition of M_4 , by the fact $\sum_{k=1}^K \sum_{h=1}^H |I(k, h+1) - I(k, h)| \leq |\mathcal{K}^C|$, we have that

$$\begin{aligned} M_4 &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, \check{V}_{h+1}^k) \\ &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k) I(k, h+1) \\ &\geq \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k) I(k, h) - |\mathcal{K}^C|. \end{aligned} \tag{43}$$

We further define $M_5 = \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k - V_{h+1}^*) I(k, h+1)$. Following similar arguments, we have that

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k - V_{h+1}^*) I(k, h) \leq M_5 + |\mathcal{K}^C|. \tag{44}$$

Bounding these two terms is one of the main difficulties in this paper, for which we need to use the recursion-based technique introduced in Section 5. The following two lemmas bound these two terms.

Lemma 16 *With probability $1 - 2(\log_2(KH) + 1) \log_2(KH) \delta$, it holds that*

$$M_4 \leq 2M_2 + 2|\mathcal{K}^C| + 2K + \max\{46\iota, 8\sqrt{(M_2 + |\mathcal{K}^C| + K)\iota} + 6\iota\}. \tag{45}$$

Proof Direct computation gives that

$$\begin{aligned}
M_4 &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, V_{h+1}^k) I(k, h+1) \\
&= \sum_{k=1}^K \sum_{h=1}^H \left(P_{s_h^k, a_h^k} (V_{h+1}^k)^2 - (P_{s_h^k, a_h^k} V_{h+1}^k)^2 \right) I(k, h+1) \\
&= \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (V_{h+1}^k)^2 - (V_{h+1}^k (s_{h+1}^k))^2) I(k, h+1) \\
&\quad + \sum_{k=1}^K \sum_{h=1}^H \left((V_h^k (s_h^k))^2 - (P_{s_h^k, a_h^k} V_{h+1}^k)^2 \right) I(k, h+1) - (V_1^k (s_1^k))^2 \\
&\leq \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (V_{h+1}^k)^2 - (V_{h+1}^k (s_{h+1}^k))^2) I(k, h+1) + 2 \sum_{k=1}^K \sum_{h=1}^H \max\{V_h^k (s_h^k) - P_{s_h^k, a_h^k} V_{h+1}^k, 0\} I(k, h+1) \\
&\leq \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (V_{h+1}^k)^2 - (V_{h+1}^k (s_{h+1}^k))^2) I(k, h+1) + 2 \sum_{k=1}^K \sum_{h=1}^H (r(s_h^k, a_h^k) + \beta_h^k(s_h^k, a_h^k)) I(k, h+1) \\
&\hspace{15cm} (46)
\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (V_{h+1}^k)^2 I(k, h+1) - (V_{h+1}^k (s_{h+1}^k))^2) + 2 \sum_{k=1}^K \sum_{h=1}^H \beta_h^k(s_h^k, a_h^k) I(k, h) + 2|\mathcal{K}^C| + 2K \\
&= \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (V_{h+1}^k)^2 - (V_{h+1}^k (s_{h+1}^k))^2) I(k, h+1) + 2M_2 + 2|\mathcal{K}^C| + 2K. \hspace{10cm} (47)
\end{aligned}$$

Here (46) is by (10) and (47) is by the fact $\sum_{h=1}^H r(s_h^k, a_h^k) \leq 1$.

Define $F(m) = \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (V_{h+1}^k)^{2^m} - (V_{h+1}^k (s_{h+1}^k))^{2^m}) I(k, h+1) = \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (\check{V}_{h+1}^k)^{2^m} - (\check{V}_{h+1}^k (s_{h+1}^k))^{2^m})$ for $1 \leq m \leq \log_2(H)$. Because \check{V}_{h+1}^k is measurable in \mathcal{F}_h^k , $F(m)$ can be viewed as a martingale. For a fixed m , by Lemma 13 with $\epsilon = 1$, we have that for each $m \leq \log_2(H)$,

$$\mathbb{P} \left[|F(m)| > 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, (\check{V}_{h+1}^k)^{2^m}) I(k, h+1)} + 6\iota \right] \leq 2(\log_2(KH) + 1)\delta. \quad (48)$$

Note that

$$\begin{aligned}
\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k}, (\check{V}_{h+1}^k)^{2^m}) &= \sum_{k=1}^K \sum_{h=1}^H \left(P_{s_h^k, a_h^k} (V_{h+1}^k)^{2^{m+1}} - (P_{s_h^k, a_h^k} (V_{h+1}^k)^{2^m})^2 \right) I(k, h+1) \\
&= \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} - \mathbf{1}_{s_{h+1}^k}) (V_{h+1}^k)^{2^{m+1}} I(k, h+1) \\
&\quad + \sum_{k=1}^K \sum_{h=1}^H \left((V_h^k (s_h^k))^{2^{m+1}} - (P_{s_h^k, a_h^k} (V_{h+1}^k)^{2^m})^2 I(k, h+1) \right) - \sum_{k=1}^K (V_1^k (s_1^k))^{2^{m+1}}
\end{aligned}$$

$$\leq F(m+1) + \sum_{k=1}^K \sum_{h=1}^H \left((V_h^k(s_h^k))^{2^{m+1}} - (P_{s_h^k, a_h^k} V_{h+1}^k)^{2^{m+1}} \right) I(k, h+1) \quad (49)$$

$$\leq F(m+1) + 2^{m+1} \sum_{k=1}^K \sum_{h=1}^H \max\{V_h^k(s_h^k) - P_{s_h^k, a_h^k} V_{h+1}^k, 0\} I(k, h+1) \quad (50)$$

$$\leq F(m+1) + 2^{m+1} \sum_{k=1}^K \sum_{h=1}^H \left(r(s_h^k, a_h^k) + \beta_h^k(s_h^k, a_h^k) \right) I(k, h+1)$$

$$\leq F(m+1) + 2^{m+1} \left(\sum_{k=1}^K \sum_{h=1}^H \beta_h^k(s_h^k, a_h^k) I(k, h) + |\mathcal{K}^C| + K \right) \quad (51)$$

$$= F(m+1) + 2^{m+1} (M_2 + |\mathcal{K}^C| + K) \quad (52)$$

Here (49) is by convexity of x^{2^m} and (50) is by the fact $a^x - b^x \leq x \max\{a - b, 0\}$ for $a, b \in [0, 1]$. Via a union bound over $m = 1, 2, \dots, \log_2(KH)$, we have that with probability $1 - 2(\log_2(KH) + 1) \log_2(KH)\delta$,

$$F(m) \leq 2\sqrt{2(F(m+1) + 2^{m+1}(M_2 + |\mathcal{K}^C| + K))\iota} + 6\iota \quad (53)$$

holds for any $1 \leq m \leq \log_2(KH)$. Now we have obtained a recursive formula. In Lemma 2, we obtain the bound for the class of recursive formulas of the same form as (53). The proof of Lemma 2 is deferred to appendix. By (47) and Lemma 2 with parameters $\lambda_1 = KH$, $\lambda_2 = \sqrt{8\iota}$, $\lambda_3 = M_2 + |\mathcal{K}^C| + K$ and $\lambda_4 = 6\iota$, we have that with probability $1 - 2(\log_2(KH) + 1) \log_2(KH)\delta$,

$$M_4 \leq 2M_2 + 2|\mathcal{K}^C| + 2K + F(1) \leq 2M_2 + 2|\mathcal{K}^C| + 2K + \max\{46\iota, 8\sqrt{(M_2 + |\mathcal{K}^C| + K)\iota} + 6\iota\}. \quad (54)$$

■

Lemma 17 *With probability $1 - 2(\log_2(KH) + 1) \log_2(KH)\delta$, it holds that*

$$M_5 \leq 2 \max\{M_2, 1\} + 2|\mathcal{K}^C| + \max\{46\iota, 8\sqrt{(M_2 + |\mathcal{K}^C|)\iota} + 6\iota\}. \quad (55)$$

Proof Recall that $\tilde{V}_{h+1}^k = V_{h+1}^k - V_{h+1}^*$. We compute

$$\begin{aligned} M_5 &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s,a}, \tilde{V}_{h+1}^k) I(k, h+1) \\ &= \sum_{k=1}^K \sum_{h=1}^H \left(P_{s_h^k, a_h^k} (\tilde{V}_{h+1}^k)^2 - (P_{s_h^k, a_h^k} \tilde{V}_{h+1}^k)^2 \right) I(k, h+1) \\ &= \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (\tilde{V}_{h+1}^k)^2 - (\tilde{V}_{h+1}^k (s_{h+1}^k))^2) I(k, h+1) \end{aligned}$$

$$\begin{aligned}
& + \sum_{k=1}^K \sum_{h=1}^H \left((\tilde{V}_h^k(s_h^k))^2 - (P_{s_h^k, a_h^k} \tilde{V}_{h+1}^k)^2 \right) I(k, h+1) - \sum_{k=1}^K (\tilde{V}_1^k(s_1^k))^2 \\
& \leq \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (\tilde{V}_{h+1}^k)^2 - (\tilde{V}_{h+1}^k(s_{h+1}^k))^2) I(k, h+1) + 2 \sum_{k=1}^K \sum_{h=1}^H \max\{\tilde{V}_h^k(s_h^k) - P_{s_h^k, a_h^k} \tilde{V}_{h+1}^k, 0\} I(k, h+1)
\end{aligned} \tag{56}$$

$$\begin{aligned}
& \leq \sum_{k=1}^K \sum_{h=1}^H \left(P_{s_h^k, a_h^k} (\tilde{V}_{h+1}^k)^2 - (\tilde{V}_{h+1}^k(s_{h+1}^k))^2 \right) I(k, h+1) + 2 \sum_{k=1}^K \sum_{h=1}^H \beta_h^k(s_h^k, a_h^k) I(k, h+1) \\
& \leq \sum_{k=1}^K \sum_{h=1}^H \left(P_{s_h^k, a_h^k} (\tilde{V}_{h+1}^k)^2 - (\tilde{V}_{h+1}^k(s_{h+1}^k))^2 \right) I(k, h+1) + 2 \sum_{k=1}^K \sum_{h=1}^H \beta_h^k(s_h^k, a_h^k) I(k, h) + 2|\mathcal{K}^C| \\
& \leq \sum_{k=1}^K \sum_{h=1}^H \left(P_{s_h^k, a_h^k} (\tilde{V}_{h+1}^k)^2 - (\tilde{V}_{h+1}^k(s_{h+1}^k))^2 \right) I(k, h+1) + 2 \max\{M_2, 1\} + 2|\mathcal{K}^C|.
\end{aligned}$$

Here (56) is by (12) Define $\tilde{F}(m) = \sum_{k=1}^K \sum_{h=1}^H (P_{s_h^k, a_h^k} (\tilde{V}_{h+1}^k)^{2^m} - (\tilde{V}_{h+1}^k(s_{h+1}^k))^{2^m}) I(k, h+1)$. Following the same arguments in (48) and (52), we obtain that with probability $1 - 2(\log_2(KH) + 1) \log_2(KH) \delta$,

$$\tilde{F}(m) \leq 2\sqrt{2(\tilde{F}(m+1) + 2^{m+1}(\max\{M_2, 1\} + |\mathcal{K}^C|))\iota} + 6\iota \tag{57}$$

holds for any $1 \leq m \leq \log_2(KH)$. By applying Lemma 2 with $\lambda_1 = KH$, $\lambda_2 = \sqrt{8\iota}$, $\lambda_3 = (\max\{M_2, 1\} + |\mathcal{K}^C|)$ and $\lambda_4 = 6\iota$, we have that with probability $1 - 2(\log_2(KH) + 1) \log_2(KH) \delta$,

$$M_5 \leq 2 \max\{M_2, 1\} + 2|\mathcal{K}^C| + \tilde{F}(1) \leq 2 \max\{M_2, 1\} + \max\{46\iota, 8\sqrt{(M_2 + |\mathcal{K}^C|)\iota} + 6\iota\}. \tag{58}$$

■

Combining (43), (44), (41), (45) and (55), we have that with probability $1 - (6S^2AH(\log_2(KH) + 1) + 6(\log_2(KH) + 1) \log_2(H)) \delta$,

$$M_2 \leq O\left(\sqrt{SAi_{\max}(M_4 + |\mathcal{K}^C|)\iota} + \sqrt{S^2Ai_{\max}(M_5 + |\mathcal{K}^C|)\iota} + \sqrt{SAi_{\max}K\iota} + S^2A\iota \log_2(KH)\right), \tag{59}$$

$$M_4 \leq 2M_2 + 2|\mathcal{K}^C| + 2K + \max\{46\iota, 8\sqrt{(M_2 + 2K)\iota} + 6\iota\}, \tag{60}$$

$$M_5 \leq 2 \max\{M_2, 1\} + 2|\mathcal{K}^C| + \max\{46\iota, \sqrt{M_2\iota} + 6\iota\}. \tag{61}$$

These imply that

$$\begin{aligned}
M_2 & \leq O\left(\sqrt{SAKi_{\max}\iota} + \sqrt{S^2Ai_{\max}\sqrt{M_2\iota^{3/2}}} + \sqrt{SAi_{\max}K\iota} + S^2A\iota \log_2(KH)\right) \\
& \leq O\left(\sqrt{SAKi_{\max}\iota} + S^2A\iota \log_2(KH)\right).
\end{aligned} \tag{62}$$

■

D.4. Proof of Lemma 10

Proof For the term M_3 , we have

$$\begin{aligned}
M_3 &= \sum_{k=1}^K \left(\sum_{h=1}^H \check{r}_h^k - V_1^{\pi^k}(s_1^k) \right) \\
&= \sum_{k=1}^K \sum_{h=1}^H (\check{r}_h^k - r_h^k) + \sum_{k=1}^K \left(\sum_{h=1}^H r_h^k - V_1^{\pi^k}(s_1^k) \right) \\
&\leq \sum_{k=1}^K \sum_{h=1}^H (r(s_h^k, a_h^k) - r_h^k) + \sum_{k=1}^K \left(\sum_{h=1}^H r_h^k - V_1^{\pi^k}(s_1^k) \right). \tag{63}
\end{aligned}$$

For the first term in **RHS** of (63), by Lemma 13, we have that

$$\mathbb{P} \left[\left| \sum_{k=1}^K \sum_{h=1}^H (r(s_h^k, a_h^k) - r_h^k) \right| > 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \text{Var}(s, a) \iota} + 6\iota \right] \leq 2(\log_2(KH) + 1)\delta, \tag{64}$$

where $\text{Var}(s, a) := \mathbb{E} [(R(s, a) - \mathbb{E}[R(s, a)])^2]$. Since for a random variable $Z \in [0, 1]$, $\text{Var}[Z] \leq \mathbb{E}[Z]$, we have

$$\sum_{k=1}^K \sum_{h=1}^H \text{Var}(s, a) \leq \sum_{k=1}^K \sum_{h=1}^H r(s, a) \leq \sum_{k=1}^K \sum_{h=1}^H (r(s_h^k, a_h^k) - r_h^k) + K,$$

Define $\bar{M}_3 := \sum_{k=1}^K \sum_{h=1}^H (r(s_h^k, a_h^k) - r_h^k)$. We then have

$$\mathbb{P} \left[|\bar{M}_3| > 2\sqrt{2(\bar{M}_3 + K)\iota} + 6\iota \right] \leq 2(\log_2(KH) + 1)\delta, \tag{65}$$

which implies that $|\bar{M}_3| \leq 6\sqrt{K\iota} + 21\iota$ with probability at least $1 - 2(\log_2(KH) + 1)\delta$.

As for the second term in **RHS** of (63), we define $Y_k = \sum_{h=1}^H r_h^k - V_1^{\pi^k}(s_1^k)$. Because for each k , $|Y_k| \leq 1$ and $\mathbb{E}[Y_k | \mathcal{F}^{k-1}] = 0$, by Azuma's inequality, we have

$$\mathbb{P} \left[\left| \sum_{k=1}^K Y_k \right| > \sqrt{2K\iota} \right] \leq \delta. \tag{66}$$

Combining (65) with (66), we have that

$$\mathbb{P} \left[|M_3| > 8\sqrt{K\iota} + 6\iota \right] \leq 2(\log_2(KH) + 2)\delta. \tag{67}$$

■