
Near-Optimal Randomized Exploration for Tabular Markov Decision Processes

Zhihan Xiong^{1*}
zhihanx@cs.washington.edu

Ruoqi Shen^{1*}
shenr3@cs.washington.edu

Qiwen Cui^{1*}
qwcui@cs.washington.edu

Maryam Fazel²
mfazel@uw.edu

Simon S. Du¹
ssdu@cs.washington.edu

¹ Paul G. Allen School of Computer Science & Engineering, University of Washington
² Department of Electrical & Computer Engineering, University of Washington

Abstract

We study algorithms using randomized value functions for exploration in reinforcement learning. This type of algorithms enjoys appealing empirical performance. We show that when we use 1) a single random seed in each episode, and 2) a Bernstein-type magnitude of noise, we obtain a worst-case $\tilde{O}\left(H\sqrt{SAT}\right)$ regret bound for episodic time-inhomogeneous Markov Decision Process where S is the size of state space, A is the size of action space, H is the planning horizon and T is the number of interactions. This bound polynomially improves all existing bounds for algorithms based on randomized value functions, and for the first time, matches the $\Omega\left(H\sqrt{SAT}\right)$ lower bound up to logarithmic factors. Our result highlights that randomized exploration can be near-optimal, which was previously achieved only by optimistic algorithms. To achieve the desired result, we develop 1) a new clipping operation to ensure both the probability of being optimistic and the probability of being pessimistic are lower bounded by a constant, and 2) a new recursive formula for the absolute value of estimation errors to analyze the regret.

1 Introduction

This paper concerns learning in tabular Markov Decision Processes (MDP), arguably the most fundamental model for reinforcement learning (RL). Existing algorithms that achieve the near-optimal minimax $\tilde{O}\left(H\sqrt{SAT}\right)$ regret bound are based on the principle of *Optimism in the face of Uncertainty* (OFU), such as upper confidence bound (UCB) [7, 53, 16, 57, 55].¹ Here S is the number of states, A is the number of actions, H is the planning horizon, and T is the total number of interactions between the agent and the environment.

*Equal Contribution

¹This bound is for time-inhomogeneous MDP with each reward bounded by 1 and T is sufficiently large.

Another broad category is algorithms with randomized exploration such as Thompson Sampling [36, 5, 39]. These algorithms inject (carefully tuned) random noise to value function to encourage exploration. UCB-type algorithms enjoy well-established theoretical guarantees but suffer from difficult implementation since an upper confidence bound is usually infeasible for many practical models like neural networks. Instead, practitioners prefer randomized exploration such as noisy networks in [19], and algorithms with randomized exploration have been widely used in practice [37, 13, 11, 35]. However, how to design randomized exploration algorithms in a principled way and perform randomized exploration optimally is far from clear. While randomized exploration can have great performance in practice, theoretically, the best known worst-case regret bound for algorithms with randomized exploration is $\tilde{O}\left(H^2 S \sqrt{AT}\right)$ [2], which is far worse than that of the UCB-type algorithms. In this paper, we introduce a new randomized exploration algorithm and show it enjoys a near-optimal $\tilde{O}\left(H \sqrt{SAT}\right)$ worst-case regret bound, thus closing the gap. Our work sheds new light on randomized exploration on both the algorithmic side and the theoretical side.

Our Contributions. Our contributions are summarized below:

- We propose a new algorithm, Single Seed Randomization (SSR), which incorporates a crucial algorithmic idea: using a single random seed for the entire episode, in contrast to previous methods of randomized exploration which use one seed for each time step. SSR is able to explore more efficiently than previous methods by avoiding having noise at different time steps canceling with each other. Theoretically, we show, thanks to this new idea, if one uses a Hoeffding-type magnitude of noise, SSR achieves an $\tilde{O}\left(H^{1.5} \sqrt{SAT}\right)$ regret bound, improving upon the best existing result on randomized exploration algorithm [2].
- We further design a new Bernstein-type magnitude of noise for our algorithm, and achieve an $\tilde{O}\left(H \sqrt{SAT}\right)$ regret bound, resolving an open problem raised in [2]. To our knowledge, this is the first time that a Bernstein-type bound is used in randomized exploration. More importantly, our upper bound matches the $\Omega\left(H \sqrt{SAT}\right)$ minimax lower bound up to logarithmic factors.

We note that our goal is not to show randomized exploration is better than optimistic algorithms [7] in the tabular setting. Instead, we aim to provide a solid theoretical understanding of a practically relevant algorithm. Indeed, understanding randomized exploration itself is an important theoretical research direction and has attracted much interest in the community [3, 4, 6, 41, 52, 48, 2, 36, 39, 37, 35].

Main Challenge and Technical Overview. Besides the aforementioned algorithmic ideas (single random seed and Bernstein-type magnitude of noise), we also need additional ideas in analysis to prove the desired regret bound. The main challenge is that unlike UCB-type algorithms, the estimated value in algorithm with randomized exploration, is not an upper bound of the true optimal value. This leads to the failure of directly utilizing their analysis, which only need to analyze the one-sided error in estimation. We instead work on the *absolute value* of the estimation error, whose analysis is more complicated than that for the one-sided error in UCB-type algorithms. Working with absolute value forces us to ensure that both the probability that the estimated value is optimistic and the probability that the estimated value is pessimistic are lower bounded. However, the clipping strategy in existing algorithm cannot maintain pessimism. To tackle with this issue, we develop a new clipping method. Below we list our technical contributions.

1. First, we propose a new clipping strategy to constrain the estimated value function (cf. Eqn. (4)). Previous clipping strategies in [52, 2] are based on uncertainty and can only maintain optimism. Our clipping strategy directly works on the value function, which is similar to those used in UCB-type algorithms [7, 22, 57]. Our clipping strategy can maintain both the optimism and pessimism. In addition, the number of times that the clipping is used can still be bounded.
2. Second, we prove that the single seed randomization ensures that the estimated value function can both be optimistic or pessimistic with constant probability at all states and timesteps. This is stronger than previous randomized exploration algorithms that are only shown to be optimistic at the initial state with constant probability. With this property, we can then bound the difference between the optimal value function and estimated value function from both above and below, which results in a bound on its absolute value. See Section 5.1, Appendix C and Appendix D.

3. Third, we prove a novel recursion argument on the absolute value of the policy estimation error. As mentioned in [2], the recursion in UCB-type algorithms can not be directly utilized because our estimated value function is not a high-probability upper bound of the true optimal value function. With the bound of absolute value, we are able to prove new recursion formulas and together we can control the policy estimation error. See Section 5.2 and Appendix E.

4. At last, we bound the sum of variance in a novel manner. In [7], the UCB-type estimation guarantees that the policy estimation error is always positive so the difference of the variance can be directly bounded. We generalize the argument to the absolute value of the estimation error to bound the sum of variance. See Section 5.3.1 and Appendix G.

2 Related Work

In this section we review existing provably efficient algorithms for tabular MDP. There is a long list of sample complexity guarantees for tabular MDP [27, 9, 25, 43, 44, 28, 8, 21, 45, 30, 36, 14, 7, 15, 38, 5, 22, 20, 46, 16, 18, 42, 41, 54, 12, 56, 51, 40, 34, 55, 49, 2, 41, 5, 17, 33, 31]. The state-of-the-art methods are based on upper confidence bound (UCB) [7, 53, 16, 57, 55, 33, 31]. For the setting considered in this paper where the transition is time-inhomogeneous and the reward is bounded by 1, one can achieve an $\tilde{O}\left(H\sqrt{SAT}\right)$ in the regime where T is sufficiently large.

Algorithms with randomized exploration have been proved to enjoy favorable regret bounds in bandit problems [29, 3, 26, 10, 4]. In certain settings, randomized exploration can match the worst-case regret bound of UCB-based approaches and achieve nearly minimax optimal regret bounds [24, 4]. However, for RL, existing theory for randomized exploration are far from optimal [2, 41, 5, 50, 52]. For the setting considered in this paper, the sharpest existing regret bound among algorithms with randomized exploration is $\tilde{O}\left(H^2S\sqrt{AT}\right)$ proved in [2]. Our paper closes this gap and thus deepens our understanding about randomized exploration.

3 Preliminaries

We consider time-inhomogeneous finite-horizon MDP $M = (H, \mathcal{S}, \mathcal{A}, P, R, s_1)$, where $|\mathcal{S}| = S$ and $|\mathcal{A}| = A$. Here, $\mathcal{S} = \{1, \dots, S\}$ is the finite state space. $\mathcal{A} = \{1, \dots, A\}$ is the finite action space. H is the length of an episode. For convenience, we take s_1 to be the fixed initial state, although a more general initial distribution will not change the conclusion. $P : \mathcal{S} \times \mathcal{A} \times [H] \rightarrow \Delta(\mathcal{S})$ is the transition function, where if the agent stays at state s and takes action a at time h , it transits to state s' with probability $P_{h,s,a}(s') \in [0, 1]$. $R : \mathcal{S} \times \mathcal{A} \times [H] \rightarrow [0, 1]$ is the reward function, where if the agent stays at s and takes action a at time h , it will receive reward $r_{h,s,a} \in [0, 1]$ such that $\mathbb{E}[r_{h,s,a}] = R_{h,s,a}$.

A deterministic policy for such a MDP is defined as a tuple $\pi = (\pi_1, \dots, \pi_H)$, where $\pi_h : \mathcal{S} \mapsto \mathcal{A}$. The associated value function at state $s \in \mathcal{S}$ and level $h \in \{1, \dots, H\}$ is recursively defined as

$$V_h^\pi(s) = R_{h,s,\pi_h(s)} + \sum_{s' \in \mathcal{S}} P_{h,s,\pi_h(s)}(s') V_{h+1}^\pi(s').$$

For convenience, we set $V_{H+1}^\pi = \mathbf{0} \in \mathbb{R}^S$. The corresponding optimal value function is $V_h^*(s) = \max_{\pi \in \Pi} V_h^\pi(s)$, where Π is the set of all possible deterministic policies. For a particular algorithm Alg, let π^k denote the policy that Alg employs during episode k . Then, the regret of running Alg on MDP M for K episodes is defined as

$$\text{Reg}(M, K, \text{Alg}) = \sum_{k=1}^K \left(V_1^*(s_1) - V_1^{\pi^k}(s_1) \right). \quad (1)$$

Note that the regret, $\text{Reg}(M, K, \text{Alg})$, is a random variable due to randomness in state transition and the algorithm, Alg. In this paper, we show the regret of our proposed algorithm can be upper bounded with high probability, and the upper bound matches the known lower bound up to logarithmic factors.

To facilitate our later analysis, we introduce some notations for empirical estimation. At episode k , we collect a trajectory $(s_1^k, a_1^k, r_1^k, \dots, s_H^k, a_H^k, r_H^k)$ as specified in Algorithm 1. Let $n_k(h, s, a) =$

Algorithm 1: Single Seed Randomization (SSR)

Input: H, S, A , perturbation type $\text{ty} \in \{\text{Ho}, \text{Be}\}$

- 1 **for** episodes $k = 1, 2, \dots, K$ **do**
- 2 Sample $\hat{z}_k \sim \mathcal{N}(0, 1)$
- 3 Define terminal value function $\bar{Q}_{H+1,k} = \mathbf{0} \in \mathbb{R}^{SA}$ and $\bar{V}_{H+1,k} = \mathbf{0} \in \mathbb{R}^S$
- 4 **for** time periods $h = H, \dots, 1$ **do**
- 5 $\bar{Q}_{h,k}(s, a) \leftarrow \hat{R}_{h,s,a}^k + \langle \hat{P}_{h,s,a}^k, \bar{V}_{h+1,k} \rangle + \sigma_{\text{ty}}^k(h, s, a) \hat{z}_k$; // $\sigma_{\text{ty}}^k(h, s, a)$ is defined in (5) and (6).
- 6 Define $\bar{V}_{h,k}(s) = \text{clip}_{2(H-h+1)}(\max_{a \in \mathcal{A}} \bar{Q}_{h,k}(s, a))$ for all $s \in \mathcal{S}$
- 7 **end**
- 8 Agent takes actions $a_h^k = \text{argmax}_{a \in \mathcal{A}} \bar{Q}_{h,k}(s_h^k, a)$ throughout this episode
- 9 Observe data $s_1^k, a_1^k, r_1^k, \dots, s_H^k, a_H^k, r_H^k$ and compute $\hat{R}_{h,s,a}^{k+1}, \hat{P}_{h,s,a}^{k+1}$ and $n_{k+1}(h, s, a)$ for all $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$
- 10 **end**

$\sum_{l=1}^{k-1} \mathbb{1}\{(s_h^l, a_h^l) = (s, a)\}$ be the number of times action a is taken at state s and time h before episode k , where $\mathbb{1}\{\cdot\}$ is the indicator function. We define

$$\hat{R}_{h,s,a}^k = \frac{\sum_{l=1}^{k-1} \mathbb{1}\{(s_h^l, a_h^l) = (s, a)\} r_{h,s_h^l, a_h^l}^l}{n_k(h, s, a) + 1}, \quad (2)$$

$$\hat{P}_{h,s,a}^k(s') = \frac{\sum_{l=1}^{k-1} \mathbb{1}\{(s_h^l, a_h^l, s_{h+1}^l) = (s, a, s')\}}{n_k(h, s, a) + 1}. \quad (3)$$

Then, define empirical MDP based on our observation and estimation before episode k as the tuple $\hat{M}^k = (H, \mathcal{S}, \mathcal{A}, \hat{P}^k, \hat{R}^k, s_1)$. Since $\hat{P}_{h,s,a}^k$ is not a valid distribution over \mathcal{S} , for being rigorous, we can imagine there is an additional virtual absorbing state that every state will transit to with remaining probability.

In addition to the above notations, let $\tilde{O}(\cdot)$, $\tilde{\Theta}(\cdot)$ and $\tilde{\Omega}(\cdot)$ be asymptotic notations ignoring all poly-logarithmic terms. For distribution $D \in \Delta^{\mathcal{S}}$ and value function $V \in \mathbb{R}^{\mathcal{S}}$, let $\mathbb{V}(D, V)$ denote the variance of V under distribution D , which is defined as $\mathbb{V}(D, V) = \sum_{s \in \mathcal{S}} D(s) (V(s) - \langle D, V \rangle)^2$. For constant $a > 0$, we define the corresponding clipping function as $\text{clip}_a(\cdot) = \max\{-a, \min\{a, \cdot\}\}$. Immediately we have $|\text{clip}_a(x)| \leq a$ for any $a > 0$. We introduce the definitions of other notations when used. In appendix, we summarize the notations and definitions used in this paper.

4 Main Results

4.1 Algorithm

The main contribution of this paper is that we show algorithm with randomized value functions can achieve regret that matches the known lower bound $\Omega\left(H\sqrt{SAT}\right)$ [21, 17] up to logarithmic factors in the tabular setting. To facilitate exploration, this type of algorithms uses random value perturbation instead of deterministic bonus. The algorithm we consider is summarized in Algorithm 1. In our algorithm, SSR, the random perturbation ensures that optimism/pessimism can be obtained with constant probability in each episode. Moreover, randomized value function has its origin from posterior sampling for reinforcement learning (Thompson sampling). The randomized perturbation can be interpreted as approximate sampling from the posterior distribution of the value function on randomized training data [41].

We first give an overview of SSR. In Algorithm 1, the policy used at episode k is computed using the empirical MDP, $\hat{M}^k = (H, \mathcal{S}, \mathcal{A}, \hat{P}^k, \hat{R}^k, s_1)$, which is based on observation and estimation before episode k . However, instead of directly choosing optimal policy for \hat{M}^k , we add a small random perturbation when computing the value of each state and action pair. To be more precise, at each

episode k , we first estimate the reward and transition function for each state s and action a based on (2) and (3). Then, we compute the value function for state s and action a ,

$$\bar{Q}_{h,k}(s, a) \leftarrow \hat{R}_{h,s,a}^k + \left\langle \hat{P}_{h,s,a}^k, \bar{V}_{h+1,k} \right\rangle + \sigma_{\text{ty}}^k(h, s, a) \hat{z}_k.$$

Here, $\hat{z}_k \sim \mathcal{N}(0, 1)$ is a standard Gaussian random variable sampled once every episode. The magnitude of the perturbation, σ_{ty}^k depends on how many samples $n_k(h, s, a)$ we have observed and how confident we are on the estimations $\hat{R}_{h,s,a}^k$ and $\hat{P}_{h,s,a}^k$. We will discuss more about the choice of the magnitude later in this section.

In order to prevent estimated value function from behaving badly, we add a clipping to the value function:

$$\bar{V}_{h,k}(s) = \text{clip}_{2(H-h+1)} \left(\max_{a \in \mathcal{A}} \bar{Q}_{h,k}(s, a) \right) \quad (4)$$

As our analysis will show, this kind of clipping can bound the value function, maintain optimism and pessimism and also guarantee that clipping will not happen for a lot of times. The constant 2 (instead of 1) plays a crucial role because it means the value function grows at an additive rate of 2 from $h = H$ to $h = 1$. If we do not consider the added noise, then the value function should at most grow 1 at each timestep because the reward is at most 1. For our clipping technique, if a clip is triggered, there exists a timestep such that the added noise is more than 1, which is equivalent to a small number of visits (cf. Definition 23 and Lemma 8). As our later analysis will show, the clipping only affects the lower-order term and will not compromise the long-term performance of the algorithm. Finally, after computing the value function and clipping, SSR chooses the action a_h^k that maximizes $\bar{Q}_{h,k}(s_h^k, a)$ at each time step, $h = 1, \dots, H$, throughout the episode.

Note that from a Bayesian perspective, when there is no clipping, in Algorithm 1, $\bar{Q}_{h,k}$ follows distribution

$$\bar{Q}_{h,k}(s, a) \mid \bar{V}_{h+1,k} \sim \mathcal{N} \left(\hat{R}_{h,s,a}^k + \left\langle \hat{P}_{h,s,a}^k, \bar{V}_{h+1,k} \right\rangle, (\sigma_{\text{ty}}^k(h, s, a))^2 \right).$$

This resembles posterior sampling because when estimating some parameter $\theta^* \sim \mathcal{N}(0, \beta^2)$ based on noisy observations $\theta_1, \dots, \theta_n \sim \mathcal{N}(\theta, \beta^2)$, the posterior distribution of θ^* given $\{\theta_i\}_{i=1}^n$ is $\theta^* \mid \{\theta_i\}_{i=1}^n \sim \mathcal{N} \left(\frac{1}{n+1} \sum_{i=1}^n \theta_i, \frac{\beta^2}{n+1} \right)$. Although exact posterior sampling may not be possible in complex reinforcement learning settings, in SSR, $\sigma_{\text{ty}}^k(h, s, a)$ is chosen at scale $\tilde{\Theta} \left(1/\sqrt{n_k(h, s, a)} \right)$ and therefore can be interpreted as doing approximate posterior sampling. Moreover, SSR can be viewed as a variant of Randomized Least Square Value Iteration (RLSVI). The major differences are at the clipping function and a single random seed used in each episode instead of different random seeds at different tuples (h, s, a) . We will discuss more about the choice of the random seed later in this section. We refer to [37] and [41] for a more detailed discussion on the relationship among RLSVI, posterior sampling and randomized value function.

In the following paragraphs, we discuss in more details about the three major algorithmic innovations:

Single Random Seed in Each Episode. SSR is similar to the algorithms analyzed in [41] and [2]. The major difference is that in the algorithm we propose, we use a single random seed \hat{z}_k to generate the perturbations for all time steps $h = 1, \dots, H$ in an episode k .

When using different random seeds in an episode, the algorithm can be optimistic in some time step while being pessimistic in others. Then, the effects of the perturbations at different time steps will cancel with each other. As a result, to ensure sufficient exploration, the magnitude of the perturbation has to be large. This issue was also pointed out in [2, 1].

A large perturbation magnitude can increase the instability of the algorithm and worsen the algorithm's performance. When a single random seed is used, a small perturbation magnitude is enough to guarantee that the algorithm is optimistic with constant probability in any episode. We are able to show that using a single random seed can significantly increase the stability of the algorithm and therefore enjoy much smaller regret. Coincidentally, [48] also uses a similar single randomization in bandit problems to build a near-optimal randomized exploration algorithm and our work can be treated as its natural extension to RL problems.

Clipping. To obtain a tight regret bound, the estimated value function needs to be well bounded. In [41], no clipping is used and the estimated value function is at the order of $\tilde{O}(H^{5/2}S)$, which results in a suboptimal regret bound. Generally there are two types of clipping methods. The first one is uncertainty-based, i.e. the value is clipped to $H - h + 1$ at timestep h whenever the uncertainty is large [52, 2]. However this type of clipping cannot maintain pessimism which is critical in our analysis. The other kind of clipping is value-based, mostly in UCB-type algorithms [23]. These algorithms truncate estimated value greater than a certain threshold, i.e. $H - h + 1$ at time step h . The problem here is that the number of clippings cannot be bounded because if the true value function is close to $H - h + 1$ at timestep h , the clipping will happen with some constant probability.

Our clipping method leverages both type of clipping methods in the existing literature. Though our clipping is based on the value function, we show that whenever the clip is triggered, the estimation error must be large, which implies that the uncertainty at that state is large. This clipping method inherits the desired properties from both uncertainty-based and value-based clipping, i.e. the optimism/pessimism is maintained and the number of clippings can be bounded.

Magnitude of Perturbation. A large magnitude of perturbation can encourage exploration, but at the same time increase instability. In our algorithm, the magnitudes are chosen as the smallest values so that the algorithm can be optimistic with constant probability. Since the value function can roughly be bounded by $O(H)$, a naive choice of the perturbation magnitude can be $\Theta\left(\frac{H}{\sqrt{n_k(h, s, a)}}\right)$. In this way, by Hoeffding’s inequality, as long as the random Gaussian variable sampled \hat{z}_k is bigger than a constant, which happens with constant probability, the estimated value function will be optimistic. By similar reasoning, we can see that the estimated value function will also be pessimistic with constant probability.

To make the magnitude even smaller, inspired by [7] who showed one can use an (empirical) Bernstein’s inequality to derive a sharp exploration bonus for UCB-based algorithms, we propose a new choice of perturbation magnitude based on Bernstein’s inequality. The Bernstein-based perturbation uses the empirical variance of the value function, which makes it smaller than the Hoeffding-based one mostly, but still maintains optimism with constant probability.

In our paper, we study both types of magnitudes. In particular, we show that the regret of SSR based on Bernstein’s inequality matches the known lower bound $\Omega\left(H\sqrt{SAT}\right)$. Following are the two choices:

$$\sigma_{\text{Ho}}^k(h, s, a) = H\sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1}} + \frac{H}{n_k(h, s, a) + 1}, \quad (5)$$

$$\sigma_{\text{Be}}^k(h, s, a) = \sqrt{\frac{16\mathbb{V}\left(\hat{P}_{h,s,a}^k, \bar{V}_{k,h+1}\right) \log(2HSAk^2)}{n_k(h, s, a) + 1}} + \frac{65H \log(2HSAk^2)}{n_k(h, s, a) + 1} + \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1}}, \quad (6)$$

where subscript “Ho” represents that the perturbation is based on Hoeffding’s inequality and “Be” represents Bernstein’s inequality, correspondingly. Here, for proof convenience, $\hat{P}_{h,s,a}^k$ is defined by replacing the denominator in $\hat{P}_{h,s,a}^k$ by $\max\{n_k(h, s, a), 1\}$. To clarify, when subscript “ty” is used, which stands for “type” as a placeholder for “Ho” or “Be”, it means that there is no need to write two copies of expressions for Hoeffding-based and Bernstein-based noises separately.

Practical Considerations. Here, we explain why randomized exploration is widely used in practice and why our algorithmic formulation practically has advantage over UCB-type algorithms. In randomized exploration, there are usually two important components: (1) the algorithm (e.g., Algorithm 1) and (2) the noise magnitude (σ_{ty}). In practice, the main advantage of randomized exploration lies in the algorithm component. The generalization from the tabular setting to the function approximation setting is straightforward: one can just add a random regularization term in the value estimation step, whose details can be found in [35]. On the other hand, the generalization of optimistic algorithms from the tabular setting to the function approximation setting is more non-trivial because it often requires an explicit construction of the confidence set. For the second component, although generalizing our strategy of tuning noise magnitude to the real-world function approximation setting is indeed not straightforward, it is often set as a hyper-parameter in practice.

4.2 Regret Analysis

We analyze the regret, defined in (1), of our algorithm **SSR** using both types of perturbations. Our main theorems are presented in Theorem 1 and 2. In particular, Theorem 2 shows **SSR** with Bernstein-based perturbation can achieve the regret that matches the known lower bound $\Omega\left(H\sqrt{SAT}\right)$ up to logarithmic factors. We sketch the proof of Theorem 1 and Theorem 2 in Section 5.

Theorem 1. *If the Hoeffding-type noise (5) is used, then for any MDP $M = (H, \mathcal{S}, \mathcal{A}, P, R, s_1)$, with probability at least $1 - \delta$, Algorithm 1 satisfies*

$$\text{Reg}(M, K, \text{SSR}_{\text{Ho}}) \leq \tilde{O}\left(H^{1.5}\sqrt{SAT} + H^4S^2A\right).$$

In particular, when $T \geq \tilde{\Omega}\left(H^5S^3A\right)$, it holds that $\text{Reg}(M, K, \text{SSR}_{\text{Ho}}) \leq \tilde{O}\left(H^{1.5}\sqrt{SAT}\right)$.

Theorem 2. *If the Bernstein-type noise (6) is used, then when $T \geq \tilde{\Omega}\left(H^5S^2A\right)$, for any MDP $M = (H, \mathcal{S}, \mathcal{A}, P, R, s_1)$, with probability at least $1 - \delta$, Algorithm 1 satisfies*

$$\text{Reg}(M, K, \text{SSR}_{\text{Be}}) \leq \tilde{O}\left(H\sqrt{SAT} + H^4S^2A\right).$$

In particular, if we further have $T \geq \tilde{\Omega}\left(H^6S^3A\right)$, it then holds that $\text{Reg}(M, K, \text{SSR}_{\text{Be}}) \leq \tilde{O}\left(H\sqrt{SAT}\right)$.

We give a brief comparison between **SSR** and other related works. [41] shows that RLSVI, an algorithm similar to **SSR**, can achieve $\tilde{O}\left(H^{2.5}S^{1.5}\sqrt{AT}\right)$ regret in expectation over the randomness of MDP and the algorithm. In [2], an improved high probability regret bound $\tilde{O}\left(H^2S\sqrt{AT}\right)$ is proposed, which is the sharpest bound for randomized algorithms prior to this work. Our paper closes the gap between those previous bounds and the lower bound in tabular setting.

We also run numerical simulations to empirically compare **SSR** and RLSVI in the deep-sea environment, which is commonly used as a benchmark to test an algorithm's ability to explore. The results show that **SSR** significantly outperforms RLSVI as predicted by our regret analysis. More details about our experiment can be found in Appendix J.

5 Proof Outline

In this section, we present an proof outline of Theorem 1 and 2. Since their proofs follow the same framework, we will present an unified outline and explain the individual steps particularly for each case when necessary. The details of complete proof are deferred to the appendix.

Notation For the ease of exposure, we will use a simplified notations during this sketch. Specifically, let $x = (h, s, a)$ and $x_h^k = (h, s_h^k, a_h^k)$.

5.1 Concentration and Optimism/Pessimism

We start by introducing a set of MDPs $\mathcal{M}_{\text{ty}}^k$ as a confidence set such that the empirical MDP \hat{M}^k belongs to it with high probability, meaning that we have a good estimation of the true MDP. Specifically, with $M' = (H, \mathcal{S}, \mathcal{A}, P', R', s_1)$, we define

$$\mathcal{M}_{\text{ty}}^k := \left\{ M' : \forall x = (h, s, a), |(R'_x - R_x) + \langle P'_x - P_x, V_{h+1}^* \rangle| \leq \sqrt{e_{\text{ty}}^k(x)} \right\},$$

where $\sqrt{e_{\text{Ho}}^k(x)} = \sigma_{\text{Ho}}^k(x)$ and $\sqrt{e_{\text{Be}}^k(x)} \approx \sigma_{\text{Be}}^k(x)$.

Define the event $\mathcal{C}_{\text{ty}}^k := \left\{ \hat{M}^k \in \mathcal{M}_{\text{ty}}^k \right\}$. Then, by applying Hoeffding's inequality or Bernstein's inequality, for both types of perturbation, it is possible to show that

$$\sum_{k=1}^{\infty} \mathbb{P}\left(\left(\mathcal{C}_{\text{ty}}^k\right)^c\right) = \sum_{k=1}^{\infty} \mathbb{P}\left(\hat{M}^k \notin \mathcal{M}_{\text{ty}}^k\right) \leq \frac{\pi^2}{3}.$$

Since the value function is bounded in $[0, H]$, this inequality tells us that the regret incurred by bad estimation is at most $\tilde{O}(H)$. To be precise, it holds with high probability that

$$\sum_{k=1}^K \mathbb{1} \left\{ (\mathcal{C}_{\text{ty}}^k)^c \right\} \left(V_1^* - V_{1,k}^{\pi^k} \right) (s_1^k) \leq \tilde{O}(H). \quad (7)$$

Then, to better control the estimated value function, we need it to be bounded, which requires us to clip it. Specifically, we will use two crucial properties of our clipping method. First, if $\bar{Q}_{h,k}(s, a) \geq Q_h^*(s, a), \forall (s, a) \in \mathcal{S} \times \mathcal{A}$, then we have $\bar{V}_{h,k}(s) \geq V_h^*(s), \forall s \in \mathcal{S}$. Similarly if $\bar{Q}_{h,k}(s, a) \leq Q_h^*(s, a), \forall (s, a) \in \mathcal{S} \times \mathcal{A}$, then we have $\bar{V}_{h,k}(s) \leq V_h^*(s), \forall s \in \mathcal{S}$.

In addition, we can prove that whenever a clip is triggered for s_h^k , we have $n_k(h, s_h^k, a_h^k) \leq \alpha_k$ with $\alpha_k = \tilde{O}(H^2)$. As a result, it is possible to show that the total regret incurred by clipping is at most $\tilde{O}(H^4 SA)$, which is a lower-order term when T is sufficiently large. That is, let $\mathcal{E}_{H,k}^{\text{cum}}$ denote the event that there is no clipping during episode k . Then, it holds with high probability that²

$$\sum_{k=1}^K \mathbb{1} \left\{ \mathcal{C}_{\text{ty}}^k \cap (\mathcal{E}_{H,k}^{\text{cum}})^c \right\} \left(V_1^* - V_{1,k}^{\pi^k} \right) (s_1^k) \leq \tilde{O}(H^4 S^2 A). \quad (8)$$

As claimed before, because of the randomness in Gaussian noise, our algorithm SSR will encourage exploration and it takes effect when there is no clipping and the estimation is not too bad. In other words, it can be optimistic. However, also because of this randomness, its optimism only holds in a probabilistic sense. In precise, it is possible to show that

$$\mathbb{P} \left(\bar{V}_{h,k}(s) \geq V_h^*(s), \forall h \in [H], s \in \mathcal{S} \mid \mathcal{C}_{\text{ty}}^k \right) \geq C_{\text{ty}}, \quad (9)$$

where the value of constant C_{ty} depends on the type of noise we choose. Meanwhile, we can also prove a very similar probabilistic pessimism, which means to have $\bar{V}_{h,k}(s) \leq V_h^*(s), \forall h \in [H], s \in \mathcal{S}$ with constant probability. The property of optimism and pessimism will help us upper bound the absolute value of $V_1^*(s_1^k) - \bar{V}_{1,k}(s_1^k)$, which will be discussed soon.

5.2 Regret Decomposition

Now, given equations (7) and (8), we can see that for each episode k , it only remains to bound $\mathbb{1} \left\{ \mathcal{C}_{\text{ty}}^k \cap \mathcal{E}_{H,k}^{\text{cum}} \right\} \left(V_1^* - V_{1,k}^{\pi^k} \right) (s_1^k)$. Technically, the further defined the good event \mathcal{G}_k will help make $\bar{V}_{h,k}$ better-behaved. Its precise definition will be given in the appendix. Therefore, it is sufficient to bound $\mathbb{1} \left\{ \mathcal{G}_k \right\} \left(V_1^* - V_{1,k}^{\pi^k} \right) (s_1^k)$, which means to have

$$\text{Reg}(M, K, \text{SSR}_{\text{ty}}) \leq \sum_{k=1}^K \mathbb{1} \left\{ \mathcal{G}_k \right\} \underbrace{\left(V_1^* - \bar{V}_{1,k} \right)}_{\text{pessimism}} + \underbrace{\left(\bar{V}_{1,k} - V_{1,k}^{\pi^k} \right)}_{\text{estimation error}} (s_1^k) + \tilde{O}(H^4 SA). \quad (10)$$

To proceed, we need to define two auxiliary value functions $\underline{V}_{h,k}$ and $\bar{\bar{V}}_{h,k}$, which are obtained by virtually running policy π^k on some deliberately perturbed MDPs. In particular, they are designed such that $\underline{V}_{h,k} \leq \bar{V}_{h,k} \leq \bar{\bar{V}}_{h,k}$ holds under the good event \mathcal{G}_k .

Pessimism Term Here, as a technical novelty, we bound the pessimism term's absolute value. Meanwhile, different from [52] and [2], by applying both optimism and pessimism, we do not resort to an independent copy of the perturbed MDP to bound the pessimism term and give a conceptually simpler analysis. In particular, by defining $C_1 = 1 / \min \{ C_{\text{Ho}}, C_{\text{Be}} \}$, it is possible to show that

$$\mathbb{1} \left\{ \mathcal{G}_k \right\} \left| V_{h,k}^*(s_h^k) - \bar{V}_{h,k}(s_h^k) \right| \leq \mathbb{1} \left\{ \mathcal{G}_k \right\} C_1 \left(\left| \bar{\bar{V}}_{h,k}^{\pi^k}(s_h^k) - V_{h,k}^{\pi^k}(s_h^k) \right| + \left| \underline{V}_{h,k}(s_h^k) - V_{h,k}^{\pi^k}(s_h^k) \right| \right). \quad (11)$$

The full proof is given in Appendix under Lemma 15.

²Technically, this is not precisely how we bound the regret incurred by clipping, but it aligns better with the intuition. Full technical details can be found in Appendix.

Estimation Error Term The sum of pessimism term and estimation error term can be further bounded via the techniques of recursion used in [7]. However, we want to emphasize the difference that in their algorithm, the estimated value is optimistic with high probability, which makes $\bar{V}_{h,k}(s_h^k) - V_h^*(s_h^k)$ always positive. Instead, since our optimism only holds with constant probability, we use absolute value to keep the estimation error terms positive. As a result, we show that

$$\left| \bar{V}_{1,k} - V_{1,k}^{\pi^k} \right| (s_1^k) + \left| \overline{\bar{V}}_{1,k}^{\pi^k} - V_{1,k}^{\pi^k} \right| (s_1^k) + \left| \underline{V}_{1,k}^{\pi^k} - V_{1,k}^{\pi^k} \right| (s_1^k) \lesssim e^{3C} \sum_{h=1}^H \left(L \sigma_{\text{ty}}^k(x_h^k) + \mathcal{M}_{h,k} \right), \quad (12)$$

where L denotes some poly-logarithmic term and $\mathcal{M}_{h,k}$ denotes some martingale difference sequence term at period h , episode k . The full proof is given in Appendix under Lemma 19,

5.3 Combining Different Terms

By combining equations (10), (11) and (12) and applying concentration inequalities to MDP $\mathcal{M}_{h,k}$, it is possible to show that

$$\text{Reg}(M, K, \text{SSR}_{\text{ty}}) \leq e^{3C_1} \sum_{k=1}^K \sum_{h=1}^H \mathbb{1}\{\mathcal{G}_k\} L \sigma_{\text{ty}}^k(x_h^k) + \tilde{O}\left(H\sqrt{T} + H^4 S^2 A\right). \quad (13)$$

Then, a final high-probability regret bound can be obtained by summing each individual terms over k, h separately. It is well-known among literature that

$$\sum_{k=1}^K \sum_{h=1}^H \sqrt{\frac{1}{n_k(x_h^k) + 1}} \leq \tilde{O}\left(\sqrt{HSA T}\right), \quad \sum_{k=1}^K \sum_{h=1}^H \frac{1}{n_k(x_h^k) + 1} \leq \tilde{O}(HSA). \quad (14)$$

Recall the definition of σ_{Ho}^k in equation (5). By using these two inequalities, the bound in equation (13) can be made explicit if we use Hoeffding-type noise. As a result, we have

$$\text{Reg}(M, K, \text{SSR}_{\text{Ho}}) \leq \tilde{O}\left(H^{1.5}\sqrt{SAT} + H^4 S^2 A\right).$$

5.3.1 Bound on Sum of Variance

Analyses become more involved when Bernstein-type noise is used. Specifically, notice that inequalities in (14) cannot directly be used to bound $\sum_{k,h} \mathbb{V}\left(\tilde{P}_{x_h^k}^k, \bar{V}_{h+1,k}\right)$. Here, we apply some techniques developed in [7]. However, since the optimism only holds with constant probability, the details for specific terms are quite different.

For the ease of exposure, we will ignore all constants and define $\hat{\mathbb{V}}_{h,k}^* = \mathbb{V}\left(\tilde{P}_{x_h^k}^k, V_h^*\right)$, $\hat{\mathbb{V}}_{h,k} = \mathbb{V}\left(\tilde{P}_{x_h^k}^k, \bar{V}_{h,k}\right)$. Then, by using Cauchy-Schwartz inequality and equation (14), we can get

$$U \stackrel{\text{def}}{=} \sum_{k,h} \mathbb{1}\{\mathcal{G}_k\} \sqrt{\frac{L}{n_k(x_h^k) + 1}} \left(\sqrt{\hat{\mathbb{V}}_{h,k}^*} + \sqrt{\hat{\mathbb{V}}_{h,k}} \right) \leq \sqrt{\tilde{O}(HSA) \sum_{k,h} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h,k}^* + \hat{\mathbb{V}}_{h,k} \right)} \quad (15)$$

Here, note that $U \approx \sum_{k,h} \sigma_{\text{Be}}^k(x_h^k)$. Then, after some steps of algebra, it is possible to show that

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h,k}^* + \hat{\mathbb{V}}_{h,k} \right) &\leq \tilde{O}(HT + H^2U) && \text{(When } T \geq \tilde{\Omega}(H^5 S^2 A)\text{)} \\ \implies U &\leq \tilde{O}\left(\sqrt{HSA(HT + H^2U)}\right) \leq \tilde{O}\left(H\sqrt{SAT} + H^{1.5}\sqrt{U}\right). && \text{(By using equation (15))} \end{aligned}$$

Now, we can see that $\sum_{k,h} \sigma_{\text{Be}}^k(x_h^k) \approx U \leq \tilde{O}\left(H\sqrt{SAT}\right)$ satisfies this inequality. Finally, by plugging this result back into equation (13), we can have

$$\text{Reg}(M, K, \text{SSR}_{\text{Be}}) \leq \tilde{O}\left(H\sqrt{SAT} + H^4 S^2 A\right),$$

which matches the known lower bound when $T \geq \tilde{\Omega}(H^6 S^3 A)$.

6 Conclusion

We gave a new algorithm with randomized exploration, SSR, for tabular MDP, which enjoys a near-optimal $\tilde{O}\left(H\sqrt{SAT}\right)$ regret bound in the time-homogeneous model. Previously, near-optimal regret bounds can only be achieved by optimistic algorithms. Our result also highlights the importance of using a single random seed for the entire episode and using the variance information in tuning the magnitude of noise (cf. Bernstein’s inequality).

One important open problem is whether randomized exploration can achieve a horizon-free regret bound in the time-homogeneous model where the transition is the same at different levels [53, 49, 55]. Another possible future direction is to consider whether the sub-optimal lower order terms $\tilde{O}\left(H^4S^2A\right)$ can be further improved to relax the current requirement $T \geq \tilde{\Omega}\left(H^6S^3A\right)$ for being near-optimal.

Acknowledgements

We sincerely thank Jinglin Chen and Chao Qin for pointing out mistakes in the initial draft of this paper. We are grateful for their careful reading and insightful discussion. This work was supported in part by NSF TRIPODS II-DMS 2023166, NSF CCF 2007036, NSF IIS 2110170, NSF DMS 2134106, NSF CCF 2212261, NSF IIS 2143493, NSF CCF 2019844.

References

- [1] Marc Abeille and Alessandro Lazaric. Linear thompson sampling revisited. In *Artificial Intelligence and Statistics*, pages 176–184. PMLR, 2017.
- [2] Priyank Agrawal, Jinglin Chen, and Nan Jiang. Improved worst-case regret bounds for randomized least-squares value iteration. *Thirty-fifth AAAI conference on artificial intelligence*, 2021.
- [3] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*, pages 39–1. JMLR Workshop and Conference Proceedings, 2012.
- [4] Shipra Agrawal and Navin Goyal. Near-optimal regret bounds for thompson sampling. *J. ACM*, 64(5), September 2017.
- [5] 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.
- [6] Shipra Agrawal and Randy Jia. Posterior sampling for reinforcement learning: worst-case regret bounds. In *Advances in Neural Information Processing Systems*, pages 1184–1194, 2017.
- [7] 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.
- [8] Peter L Bartlett and Ambuj Tewari. Regal: A regularization based algorithm for reinforcement learning in weakly communicating mdps. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 35–42. AUAI Press, 2009.
- [9] 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.
- [10] Sébastien Bubeck and Che-Yu Liu. Prior-free and prior-dependent regret bounds for thompson sampling. In *2014 48th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–9. IEEE, 2014.
- [11] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.

- [12] Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. *arXiv preprint arXiv:1912.05830*, 2019.
- [13] Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pages 2249–2257, 2011.
- [14] Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2818–2826, 2015.
- [15] 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.
- [16] 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.
- [17] Omar Darwiche Domingues, Pierre Ménard, Emilie Kaufmann, and Michal Valko. Episodic reinforcement learning in finite mdps: Minimax lower bounds revisited. In *Algorithmic Learning Theory*, pages 578–598. PMLR, 2021.
- [18] 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.
- [19] Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Matteo Hessel, Ian Osband, Alex Graves, Volodymyr Mnih, Remi Munos, Demis Hassabis, et al. Noisy networks for exploration. In *International Conference on Learning Representations*, 2018.
- [20] Ronan Fruit, Matteo Pirotta, and Alessandro Lazaric. Near optimal exploration-exploitation in non-communicating markov decision processes. In *Advances in Neural Information Processing Systems*, pages 2994–3004, 2018.
- [21] Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(Apr):1563–1600, 2010.
- [22] 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.
- [23] Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. *arXiv preprint arXiv:1907.05388*, 2019.
- [24] Tianyuan Jin, Pan Xu, Jieming Shi, Xiaokui Xiao, and Quanquan Gu. Mots: Minimax optimal thompson sampling. *arXiv preprint arXiv:2003.01803*, 2020.
- [25] Sham M Kakade. *On the sample complexity of reinforcement learning*. PhD thesis, University of London London, England, 2003.
- [26] Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *International conference on algorithmic learning theory*, pages 199–213. Springer, 2012.
- [27] Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2-3):209–232, 2002.
- [28] 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.
- [29] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [30] Tor Lattimore and Marcus Hutter. Pac bounds for discounted mdps. In *International Conference on Algorithmic Learning Theory*, pages 320–334. Springer, 2012.

- [31] Gen Li, Laixi Shi, Yuxin Chen, Yuantao Gu, and Yuejie Chi. Breaking the sample complexity barrier to regret-optimal model-free reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- [32] Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- [33] Pierre Menard, Omar Darwiche Domingues, Xuedong Shang, and Michal Valko. Ucb momentum q-learning: Correcting the bias without forgetting. *arXiv preprint arXiv:2103.01312*, 2021.
- [34] Gergely Neu and Ciara Pike-Burke. A unifying view of optimism in episodic reinforcement learning. *arXiv preprint arXiv:2007.01891*, 2020.
- [35] Ian Osband, John Aslanides, and Albin Cassirer. Randomized prior functions for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 8617–8629, 2018.
- [36] 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.
- [37] Ian Osband, Daniel Russo, Zheng Wen, and Benjamin Van Roy. Deep exploration via randomized value functions. *arXiv preprint arXiv:1703.07608*, 2017.
- [38] 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.
- [39] Ian Osband, Benjamin Van Roy, and Zheng Wen. Generalization and exploration via randomized value functions. *arXiv preprint arXiv:1402.0635*, 2014.
- [40] 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.
- [41] Daniel Russo. Worst-case regret bounds for exploration via randomized value functions. In *Advances in Neural Information Processing Systems*, pages 14433–14443, 2019.
- [42] 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.
- [43] 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.
- [44] 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.
- [45] István Szita and Csaba Szepesvári. Model-based reinforcement learning with nearly tight exploration complexity bounds. In *ICML*, 2010.
- [46] Mohammad Sadegh Talebi and Odalric-Ambrym Maillard. Variance-aware regret bounds for undiscounted reinforcement learning in mdps. *arXiv preprint arXiv:1803.01626*, 2018.
- [47] Tian Tan, Zhihan Xiong, and Vikranth R Dwaracherla. Parameterized indexed value function for efficient exploration in reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5948–5955, 2020.
- [48] Sharan Vaswani, Abbas Mehrabian, Audrey Durand, and Branislav Kveton. Old dog learns new tricks: Randomized ucb for bandit problems. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, pages 1988–1998, 2020.
- [49] 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.

- [50] Ziping Xu and Ambuj Tewari. Worst-case regret bound for perturbation based exploration in reinforcement learning. *Ann Arbor*, 1001:48109, 2019.
- [51] Kunhe Yang, Lin F Yang, and Simon S Du. Q -learning with logarithmic regret. *arXiv preprint arXiv:2006.09118*, 2020.
- [52] Andrea Zanette, David Brandfonbrener, Emma Brunskill, Matteo Pirodda, and Alessandro Lazaric. Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*, pages 1954–1964, 2020.
- [53] 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.
- [54] 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.
- [55] Zihan Zhang, Xiangyang Ji, and Simon S Du. Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. *arXiv preprint arXiv:2009.13503*, 2020.
- [56] Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. *arXiv preprint arXiv:2004.10019*, 2020.
- [57] Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Model-free reinforcement learning: from clipped pseudo-regret to sample complexity. *arXiv preprint arXiv:2006.03864*, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [No]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Contents

1	Introduction	1
2	Related Work	3
3	Preliminaries	3
4	Main Results	4
4.1	Algorithm	4
4.2	Regret Analysis	7
5	Proof Outline	7
5.1	Concentration and Optimism/Pessimism	7
5.2	Regret Decomposition	8
5.3	Combining Different Terms	9
5.3.1	Bound on Sum of Variance	9
6	Conclusion	10
A	Table of Notations	16
B	Good Events	17
C	Optimism	21
C.1	Hoeffding-type Noise	21
C.2	Bernstein-type Noise	22
D	Pessimism	24
D.1	Hoeffding-type Noise	24
D.2	Bernstein-type Noise	25
E	Regret Decomposition	26
E.1	Pessimism Term	27
E.2	Estimation Error Term	28
E.3	Combining Estimation and Pessimism Terms	31
F	Bounds on Individual Terms	33
F.1	Bounds on Martingale Difference	33
F.2	Bounds on Lower-order Terms	34
G	Bounds on Sum of Variance	34
G.1	Auxiliary Lemmas	37
H	Proof of the Main Theorems	39

A Table of Notations

Symbol	Meaning
\mathcal{S}	The state space
\mathcal{A}	The action space
S	Size of state space
A	Size of action space
H	The length of horizon
K	The total number of episodes
T	The total number of steps, $T = HK$
π^k	The greedy policy generated in Algorithm 1 at episode k
$R_{h,s,a}$	Expected reward function at (h, s, a)
$P_{h,s,a}(s')$	Transition probability
M	Underlying true MDP, $M = (H, \mathcal{S}, \mathcal{A}, R, P, s_1)$
$n_k(h, s, a)$	$\sum_{k'=1}^{k-1} \mathbb{1}\{(s_h^{k'}, a_h^{k'}) = (s, a)\}$
$\hat{R}_{h,s,a}^k$	Estimated reward function, $\frac{1}{n_k(h,s,a)+1} \sum_{l=1}^{k-1} \mathbb{1}\{(s_h^l, a_h^l) = (s, a)\} r_{h,s_h^l,a_h^l}^l$
$\hat{P}_{h,s,a}^k(s')$	Estimated transition kernel, $\frac{1}{n_k(h,s,a)+1} \sum_{l=1}^{k-1} \mathbb{1}\{(s_h^l, a_h^l, s_{h+1}^l) = (s, a, s')\}$
$\tilde{P}_{h,s,a}^k(s')$	Estimated transition probability with a slightly different denominator, $\frac{1}{\max\{n_k(h,s,a), 1\}} \sum_{l=1}^{k-1} \mathbb{1}\{(s_h^l, a_h^l, s_{h+1}^l) = (s, a, s')\}$
\hat{M}^k	Estimated MDP, $\hat{M}^k = (H, \mathcal{S}, \mathcal{A}, \hat{P}, \hat{R}, s_1)$
$\gamma_{\text{ty}}^k(h, s, a)$	$\sigma_{\text{ty}}^k(h, s, a) \sqrt{\log(40k^4)}$
\hat{z}_k	Perturbation's single random source during episode k from a standard Gaussian, $\hat{z}_k \sim \mathcal{N}(0, 1)$
$w_{\text{ty}}^k(h, s, a)$	Noise of type "ty", $w_{\text{ty}}^k(h, s, a) = \sigma_{\text{ty}}^k(h, s, a) \hat{z}_k$
$\underline{w}_{\text{ty}}^k(h, s, a)$	$-\gamma_{\text{ty}}^k(h, s, a)$
$\overline{w}_{\text{ty}}^k(h, s, a)$	$\gamma_{\text{ty}}^k(h, s, a)$
$\overline{M}_{\text{ty}}^k$	Perturbed estimated MDP with ty-type noise, $\overline{M}^k = (H, \mathcal{S}, \mathcal{A}, \hat{P}, \hat{R} + w_{\text{ty}}^k, s_1)$
$\underline{M}_{\text{ty}}^k$	Negatively perturbed MDP, $\underline{M}_{\text{ty}}^k = (H, \mathcal{S}, \mathcal{A}, \hat{P}, \hat{R} + \underline{w}_{\text{ty}}^k, s_1)$
$\overline{M}_{\text{ty}}^k$	Positively perturbed MDP, $\overline{M}_{\text{ty}}^k = (H, \mathcal{S}, \mathcal{A}, \hat{P}, \hat{R} + \overline{w}_{\text{ty}}^k, s_1)$
$V_h^* / V_{h,k}^*$	Optimal value function at step h for true MDP M
$V_h^{\pi^k} / V_{h,k}^{\pi^k}$	Value function at step h by running policy π^k on true MDP M
$\overline{Q}_{h,k}$	Q -value function obtained by running Algorithm 1
$\overline{V}_{h,k}$	Value function obtained by running policy π^k on \overline{M}^k with a clipping of threshold $2(H - h + 1)$
$\underline{V}_{h,k}$	Value function obtained by running policy π^k on \underline{M}^k with a clipping of threshold $2(H - h + 1)$
$\overline{\overline{V}}_{h,k}$	Value function obtained by running policy π^k on $\overline{\overline{M}}^k$ with a clipping of threshold $2(H - h + 1)$
$\mathcal{R}_{h,s,a}^k$	$\hat{R}_{h,s,a}^k - R_{h,s,a}$
$\mathcal{P}_{h,s,a}^k$	$\langle \hat{P}_{h,s,a}^k - P_{h,s,a}, V_{h+1}^* \rangle$
\mathcal{H}_h^k	The historical observations and actions till time h in episode k , $\{(s_l^j, a_l^j, r_l^j) : j \leq k \text{ and } l \leq H \text{ if } j < k, \text{ else } l \leq h\}$
$\overline{\mathcal{H}}_h^k$	The historical observations and actions till time h and episode k , plus the randomness in episode k , $\mathcal{H}_h^k \cup \{\hat{z}_k\}$
$\mathbb{V}(P, V)$	Variance of $V \in \mathbb{R}^{\mathcal{S}}$ under distribution $P \in \Delta^{\mathcal{S}}$, $\sum_{s \in \mathcal{S}} P(s)(V(s) - \langle P, V \rangle)^2$
α_k	$200H^2 \log(2HSAk^2) \log(40k^4)$

$\sigma_{\text{ty}}^k(h, s, a)$	Magnitude of perturbation. $\text{ty} \in \{\text{Ho}, \text{Be}\}$
ty	Reserved subscript for denoting perturbation type, $\text{ty} \in \{\text{Ho}, \text{Be}\}$, where ‘‘Ho’’ denotes Hoeffding-type and ‘‘Be’’ denotes Bernstein-type
$\sigma_{\text{Ho}}^k(h, s, a)$	$H \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1} + \frac{H}{n_k(h, s, a) + 1}}$
$\sigma_{\text{Be}}^k(h, s, a)$	$\sqrt{\frac{16\mathbb{V}(\tilde{P}_{h, s, a}^k, \bar{V}_{k, h+1}) \log(2HSAk^2)}{n_k(h, s, a) + 1} + \frac{65H \log(2HSAk^2)}{n_k(h, s, a) + 1} + \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1}}$
$\sqrt{e_{\text{Ho}}^k(h, s, a)}$	$H \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1} + \frac{H}{n_k(h, s, a) + 1}}$
$\sqrt{e_{\text{Be}}^k(h, s, a)}$	$\sqrt{\frac{6\mathbb{V}(\tilde{P}_{h, s, a}^k, V_{h+1}^*) \log(2HSAk^2)}{n_k(h, s, a) + 1} + \frac{9H \log(2HSAk^2)}{n_k(h, s, a) + 1} + \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1}}$
C_1	$\frac{1}{\Phi(1.5) - \Phi(1)}$

B Good Events

Definition 1. Let $M' = (H, S, \mathcal{A}, P', R', s_1)$. We define the following confidence sets for both Bernstein-type and Hoeffding-type noise

$$\mathcal{M}_{\text{ty}}^k = \left\{ M' : \forall (h, s, a), |(R'_{h, s, a} - R_{h, s, a}) + \langle P'_{h, s, a} - P_{h, s, a}, V_{h+1}^* \rangle| \leq \sqrt{e_{\text{ty}}^k(h, s, a)} \right\}, \quad (16)$$

where the confidence widths are set as

$$\begin{aligned} \sqrt{e_{\text{Be}}^k(h, s, a)} &= \sqrt{\frac{6\mathbb{V}(\tilde{P}_{h, s, a}^k, V_{h+1}^*) \log(2HSAk^2)}{n_k(h, s, a) + 1}} \\ &\quad + \frac{9H \log(2HSAk^2)}{n_k(h, s, a) + 1} + \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1}}, \end{aligned} \quad (17)$$

$$\sqrt{e_{\text{Ho}}^k(h, s, a)} = H \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1} + \frac{H}{n_k(h, s, a) + 1}}. \quad (18)$$

We also define two events \mathcal{E}_k^1 and \mathcal{E}_k^2 as the following:

$$\mathcal{E}_k^1 = \left\{ \left| \hat{R}_{h, s, a}^k - R_{h, s, a} \right| \leq \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1} + \frac{1}{n_k(h, s, a) + 1}}, \forall (h, s, a) \right\}, \quad (19)$$

$$\begin{aligned} \mathcal{E}_k^2 &= \left\{ \left| \langle \hat{P}_{h, s, a}^k - P_{h, s, a}, V_{h+1}^* \rangle \right| \leq \sqrt{\frac{6\mathbb{V}(\tilde{P}_{h, s, a}^k, V_{h+1}^*) \log(2HSAk^2)}{n_k(h, s, a) + 1}} \right. \\ &\quad \left. + \frac{8H \log(2HSAk^2)}{n_k(h, s, a) + 1}, \forall (h, s, a) \right\}. \end{aligned} \quad (20)$$

We have the following lemmas about concentration of events.

Lemma 1. For fixed (k, h, s, a) , let $n = n_k(h, s, a)$. Then, if $n \geq 1$, for any fixed $\delta > 0$, we have

$$\mathbb{P} \left(\left| \hat{R}_{h, s, a}^k - R_{h, s, a} \right| \geq \sqrt{\frac{\log(2/\delta)}{n+1} + \frac{1}{n+1}} \right) \leq \delta.$$

Proof. Let $\hat{R}_{h, s, a}^k = \frac{1}{n+1} \sum_{i=1}^n r^{(h, s, a), i}$, where $r^{(h, s, a), i} \sim \mathcal{R}_{h, s, a}$ are i.i.d. samples. By definition of the MDP, we have $\mathbb{E} [r^{(h, s, a), i}] = R_{h, s, a}$. Then, notice that

$$\hat{R}_{h, s, a}^k = \frac{1}{n+1} \sum_{i=1}^n r^{(h, s, a), i} = \frac{1}{n} \sum_{i=1}^n r^{(h, s, a), i} - \frac{1}{n(n+1)} \sum_{i=1}^n r^{(h, s, a), i}.$$

Since the reward is assumed to be bounded in $[0, 1]$, we have $\frac{1}{n(n+1)} \sum_{i=1}^n r_{(h,s,a),i} \leq \frac{1}{n+1}$. Then, for fixed $\delta > 0$, we have

$$\begin{aligned}
& \mathbb{P} \left(\left| \hat{R}_{h,s,a}^k - R_{h,s,a} \right| \geq \sqrt{\frac{\log(2/\delta)}{n+1}} + \frac{1}{n+1} \right) \\
&= \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n r_{(h,s,a),i} - R_{h,s,a} - \frac{1}{n(n+1)} \sum_{i=1}^n r_{(h,s,a),i} \right| \geq \sqrt{\frac{\log(2/\delta)}{n+1}} + \frac{1}{n+1} \right) \\
&\leq \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n r_{(h,s,a),i} - R_{h,s,a} \right| + \frac{1}{n+1} \geq \sqrt{\frac{\log(2/\delta)}{n+1}} + \frac{1}{n+1} \right) \quad (\text{By triangle inequality}) \\
&\leq \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n r_{(h,s,a),i} - R_{h,s,a} \right| \geq \sqrt{\frac{\log(2/\delta)}{2n}} \right) \quad (\text{Since } n+1 \leq 2n \text{ for } n \geq 1) \\
&\leq \delta. \quad (\text{By standard Hoeffding's inequality})
\end{aligned}$$

□

Lemma 2. For fixed (k, h, s, a) , let $n = n_k(h, s, a)$ and $V \in \mathbb{R}^S$ be some non-negative value function such that $\|V\|_\infty \leq H$. Then, if $n \geq 1$, for any fixed $\delta > 0$, we have

$$\mathbb{P} \left(\left| \left\langle \hat{P}_{h,s,a}^k - P_{h,s,a}, V \right\rangle \right| \geq H \sqrt{\frac{\log(2/\delta)}{n+1}} + \frac{H}{n+1} \right) \leq \delta, \quad (21)$$

$$\mathbb{P} \left(\left| \left\langle \hat{P}_{h,s,a}^k - P_{h,s,a}, V \right\rangle \right| \geq \sqrt{\frac{6\mathbb{V}(\tilde{P}_{h,s,a}^k, V) \log(2/\delta)}{n+1}} + \frac{8H \log(2/\delta)}{n+1} \right) \leq \delta. \quad (22)$$

Proof. For fixed (h, s, a) , we generate n i.i.d. samples of $s_{(h,s,a),i} \sim P_{h,s,a}$ and consider $V(s_{(h,s,a),i})$. Then, by taking $n_k(h, s, a) = n$, we have

$$\left\langle \hat{P}_{h,s,a}^k, V \right\rangle = \frac{1}{n} \sum_{i=1}^n V(s_{(h,s,a),i}) - \frac{1}{n(n+1)} \sum_{i=1}^n V(s_{(h,s,a),i}).$$

The first result in equation (21) can be proved very similarly as Lemma 1 using Hoeffding's inequality by simply replacing the upper bound of 1 in reward by H .

Then, for second result, we first consider $n \geq 2$. For some $\delta > 0$, define

$$b_{(h,s,a),n} = \sqrt{\frac{2\mathbb{V}(\tilde{P}_{h,s,a}^k, V) \log(2/\delta)}{n-1}} + \frac{7H \log(2/\delta)}{3(n-1)} + \frac{H}{n+1}.$$

By noticing that $F(s) \leq H$ and applying similar technique in proof of Lemma 1, we have

$$\begin{aligned}
& \mathbb{P} \left(\left| \left\langle \hat{P}_{h,s,a}^k - P_{h,s,a}, V \right\rangle \right| \geq b_{(h,s,a),n} \right) \\
&\leq \mathbb{P} \left(\left| \left\langle \tilde{P}_{h,s,a}^k - P_{h,s,a}, V \right\rangle \right| \geq \sqrt{\frac{2\mathbb{V}(\tilde{P}_{h,s,a}^k, V) \log(2/\delta)}{n-1}} + \frac{7H \log(2/\delta)}{3(n-1)} \right) \\
&\leq \delta. \quad (\text{By Lemma 31, the empirical Bernstein's inequality})
\end{aligned}$$

Then, since $3(n-1) \geq n+1$ when $n \geq 2$, we can easily check that

$$b_{(h,s,a),n} \leq \sqrt{\frac{6\mathbb{V}(\tilde{P}_{h,s,a}^k, V) \log(2/\delta)}{n+1}} + \frac{8H \log(2/\delta)}{n+1}.$$

Finally, since $\|V\|_\infty \leq H$, when $n = 1$, we trivially have

$$\left| \left\langle \hat{P}_{h,s,a}^k - P_{h,s,a}, V \right\rangle \right| \leq H \leq \sqrt{\frac{6\mathbb{V}\left(\tilde{P}_{h,s,a}^k, V\right) \log(2/\delta)}{n+1} + \frac{8H \log(2/\delta)}{n+1}}.$$

Therefore, we can conclude that

$$\mathbb{P}\left(\left|\left\langle \hat{P}_{h,s,a}^k - P_{h,s,a}, V \right\rangle\right| \geq \sqrt{\frac{6\mathbb{V}\left(\tilde{P}_{h,s,a}^k, V\right) \log(2/\delta)}{n+1} + \frac{8H \log(2/\delta)}{n+1}}\right) \leq \delta.$$

□

Lemma 3. $\sum_{k=1}^{\infty} \mathbb{P}\left((\mathcal{E}_k^1)^c\right) \leq \frac{\pi^2}{6}$.

Proof. Let $n = n_k(h, s, a)$. Then, for some fixed (h, s, a) , $n \geq 1$ and $\delta_n > 0$, by Lemma 1, we have

$$\mathbb{P}\left(\left|\hat{R}_{h,s,a}^k - R_{h,s,a}\right| \geq \sqrt{\frac{\log(2/\delta_n)}{n+1}} + \frac{1}{n+1}\right) \leq \delta_n.$$

Therefore, by taking $\delta_n = \frac{1}{HSAn^2}$, a union bound will give us

$$\sum_{n=1}^{\infty} \sum_{h,s,a} \mathbb{P}\left(\left\{\left|\hat{R}_{h,s,a}^k - R_{h,s,a}\right| \geq \sqrt{\frac{\log(2HSAn^2)}{n+1}} + \frac{1}{n+1}\right\}\right) \leq \sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{\pi^2}{6}.$$

Therefore, we have

$$\begin{aligned} & \sum_{k=1}^{\infty} \mathbb{P}\left(\exists (h, s, a) : n_k(h, s, a) > 0, \left|\hat{R}_{h,s,a}^k - R_{h,s,a}\right| \geq \sqrt{\frac{\log(2HSAn_k(h, s, a)^2)}{n_k(h, s, a) + 1}} + \frac{1}{n_k(h, s, a) + 1}\right) \\ & \leq \frac{\pi^2}{6}. \end{aligned}$$

Since the MDP is time-inhomogeneous, each (h, s, a) can only be visited at most once during one episode, which implies $n_k(h, s, a) \leq k$. Therefore, we have

$$\sqrt{\frac{\log(2HSAn_k(h, s, a))}{n_k(h, s, a) + 1}} + \frac{1}{n_k(h, s, a) + 1} \leq \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1}} + \frac{1}{n_k(h, s, a) + 1}$$

and thus the proof is complete. □

Lemma 4. $\sum_{k=1}^{\infty} \mathbb{P}\left((\mathcal{E}_k^2)^c\right) \leq \frac{\pi^2}{6}$.

Proof. This proof will be very similar to proof of Lemma 3. In specific, for fixed (h, s, a) , let $n = n_k(h, s, a) \geq 1$. Then, for any $\delta_n > 0$, since $\|V_{h+1}^*\|_\infty \leq H$, by Lemma 2, we have

$$\mathbb{P}\left(\left|\left\langle \hat{P}_{h,s,a}^k - P_{h,s,a}, V_{h+1}^* \right\rangle\right| \geq \sqrt{\frac{6\mathbb{V}\left(\tilde{P}_{h,s,a}^k, V_{h+1}^*\right) \log(2/\delta_n)}{n+1} + \frac{8H \log(2/\delta_n)}{n+1}}\right) \leq \delta_n.$$

Therefore, by taking $\delta_n = \frac{1}{HSAn^2}$ and applying a similar union bound argument used in the proof of Lemma 3, we can conclude $\sum_{k=1}^{\infty} \mathbb{P}\left((\mathcal{E}_k^2)^c\right) \leq \frac{\pi^2}{6}$. □

We further define the event $\mathcal{C}_{\text{ty}}^k = \left\{\hat{M}^k \in \mathcal{M}_{\text{ty}}^k\right\}$. With what we have proved above, it will be straightforward to show the following results about $\mathcal{C}_{\text{ty}}^k$.

Lemma 5. $\sum_{k=1}^{\infty} \mathbb{P}((\mathcal{C}_{\text{Be}}^k)^c) = \sum_{k=1}^{\infty} \mathbb{P}(\hat{M}^k \notin \mathcal{M}_{\text{Be}}^k) \leq \frac{\pi^2}{3}$

Proof. We can easily notice $\mathcal{E}_k^1 \cap \mathcal{E}_k^2 \implies \hat{M}^k \in \mathcal{M}_{\text{Be}}^k$, which implies $\hat{M}^k \notin \mathcal{M}_{\text{Be}}^k \implies (\mathcal{E}_k^1)^c \cup (\mathcal{E}_k^2)^c$. The first result then follows straightforwardly by applying Lemma 3 and Lemma 4. \square

Lemma 6. $\sum_{k=1}^{\infty} \mathbb{P}((\mathcal{C}_{\text{Ho}}^k)^c) = \sum_{k=1}^{\infty} \mathbb{P}(\hat{M}^k \notin \mathcal{M}_{\text{Ho}}^k) \leq \frac{\pi^2}{3}$.

Proof. Similarly, for fixed (h, s, a) , we generate n i.i.d. samples $s_{(h,s,a),i} \sim P_{h,s,a}$ and $r_{(h,s,a),i} \sim \mathcal{R}_{h,s,a}$ for $i = 1, \dots, n$ respectively. Define $Y_{(h,s,a),i} = r_{(h,s,a),i} + V_{h+1}^*(s_{(h,s,a),i})$ and we have $\mathbb{E}[Y_{(h,s,a),i}] = R_{h,s,a} + \langle P_{h,s,a}, V_{h+1}^* \rangle$.

By definition of MDP, we know that $Y_{(h,s,a),i} \leq H$. Thus, we can use an argument similar to the proof of Lemma 1. In specific, let $n = n_k(h, s, a)$ and for $\delta_n > 0$, we have

$$\begin{aligned} & \mathbb{P}\left(\left|\frac{1}{n+1} \sum_{i=1}^n Y_{(h,s,a),i} - \mathbb{E}[Y_{(h,s,a),i}]\right| \geq H \sqrt{\frac{\log(2/\delta_n)}{n+1}} + \frac{H}{n+1}\right) \\ &= \mathbb{P}\left(\left|\left(\hat{R}_{h,s,a}^k - R_{h,s,a}\right) + \left\langle \hat{P}_{h,s,a}^k - P_{h,s,a}, V_{h+1}^* \right\rangle\right| \geq H \sqrt{\frac{\log(2/\delta_n)}{n+1}} + \frac{H}{n+1}\right) \\ &\leq \delta_n. \end{aligned}$$

Then, we can take $\delta_n = \frac{1}{HSA n^2}$ and apply a similar union bound argument in used in the proof of Lemma 3. As a result, we can obtain

$$\sum_{k=1}^{\infty} \mathbb{P}(\hat{M}^k \notin \mathcal{M}_{\text{Ho}}^k) \leq \frac{\pi^2}{6} \leq \frac{\pi^2}{3}.$$

\square

We can also have well-behaved bounds on magnitude of noise and estimated value functions.

Definition 2. We define $w_{\text{ty}}^k(h, s, a) = \sigma_{\text{ty}}^k(h, s, a) \hat{z}_k$ and $\gamma_{\text{ty}}^k(h, s, a) = \sigma_{\text{ty}}^k(h, s, a) \sqrt{\log(40k^4)}$. We define the event \mathcal{E}_k^w as

$$\mathcal{E}_k^w = \{\forall(h, s, a), |w_{\text{ty}}^k(h, s, a)| \leq \gamma_{\text{ty}}^k(h, s, a)\}.$$

Lemma 7. $\sum_{k=1}^K \mathbb{P}((\mathcal{E}_k^w)^c) \leq \frac{\pi^2}{3}$ regardless the type of noise we choose.

Proof. For any $k \in [K]$, by the tail bound of Gaussian distribution,

$$\mathbb{P}\left(|\hat{z}_k| \geq \sqrt{\log(40k^4)}\right) \leq 2 \exp\left(-\frac{\log(40k^4)}{2}\right) \leq \frac{2}{k^2}.$$

Summing over $k \in [K]$,

$$\sum_{k=1}^K \mathbb{P}\left((\mathcal{E}_k^w)^c\right) = \sum_{k=1}^K \mathbb{P}\left(|\hat{z}_k| \geq \sqrt{\log(40k^4)}\right) \leq \sum_{k=1}^{\infty} \frac{2}{k^2} \leq \frac{\pi^2}{3}.$$

Note that this result does not depend on the type of noise we choose. \square

Now, we define the following good events that hold with high probability and will be used throughout the whole proof.

Definition 3 (Good events \mathcal{G}_k). Let $\mathcal{G}_{k,\text{ty}} = \{\mathcal{C}_{\text{ty}}^k \cap \mathcal{E}_k^w\}$.

The subscript ‘‘ty’’ will be ignored later since it is clear from the context.

Definition 4. With $\alpha_k = 200H^2 \log(2HSAk^2) \log(40k^4)$, we define events $\mathcal{E}_{h,k}^{\text{th}}$ and $\mathcal{E}_{h,k}^{\text{cum}}$ as

$$\mathcal{E}_{h,k}^{\text{th}} = \{n_k(h, s_h^k, a_h^k) \geq \alpha_k\}, \quad \mathcal{E}_{h,k}^{\text{cum}} = \bigcap_{i=1}^h \mathcal{E}_{i,k}^{\text{th}}. \quad (23)$$

We will show that under events \mathcal{E}_k^w , $\mathcal{E}_{h,k}^{th}$ and $\hat{M}^k \in \mathcal{M}_{\text{ty}}^k$, no clipping happens on s_h^k .

Lemma 8. *Assume that \mathcal{E}_k^w , $\mathcal{E}_{h,k}^{th}$ and $\hat{M}^k \in \mathcal{M}_{\text{ty}}^k$ hold. Then, regardless the type of noise we choose, it holds that*

$$|\bar{Q}_{h,k}(s_h^k, a_h^k)| \leq 2(H - h + 1),$$

which immediately tells us that no clipping is triggered for any (s_h^k, a_h^k) .

Proof. We have that

$$\bar{Q}_{h,k}(s_h^k, a_h^k) = \hat{R}_{h,s_h^k,a_h^k}^k + \left\langle \hat{P}_{h,s_h^k,a_h^k}^k, \bar{V}_{h+1,k} \right\rangle + \sigma_{\text{ty}}^k(h, s_h^k, a_h^k) \hat{z}_k.$$

As we have $|\bar{V}_{h+1,k}| \leq 2(H - h)$ by clipping and $\hat{R}_{h,s_h^k,a_h^k}^k \in [0, 1]$, we only need to show that $\sigma_{\text{ty}}^k(h, s_h^k, a_h^k) \hat{z}_k \leq 1$. Under event \mathcal{E}_k^w , we have $|\sigma_{\text{ty}}^k(h, s_h^k, a_h^k) \hat{z}_k|$ is bounded by $\gamma_{\text{ty}}^k(h, s_h^k, a_h^k) = \sigma_{\text{ty}}^k(h, s_h^k, a_h^k) \sqrt{\log(40k^4)}$. Note that we have $\bar{V}_{h+1,k}(s) \in [-2H, 2H]$ by clipping for any $s \in \mathcal{S}$. Thus, by Lemma 32, we have $\mathbb{V}(\hat{P}_{h,s,a}^k, \bar{V}_{h+1,k}) \leq 4H^2$ for any (h, s, a) .

By taking $\alpha_k = 200H^2 \log(2HSAk^2) \log(40k^4)$ and referring to the definitions of $\sigma_{\text{Be}}^k(h, s, a)$ in Equation (6), we can check that

$$\begin{aligned} & \gamma_{\text{Be}}^k(h, s_h^k, a_h^k) \\ &= \sigma_{\text{Be}}^k(h, s_h^k, a_h^k) \sqrt{\log(40k^4)} \\ &= \left(\sqrt{\frac{16\mathbb{V}(\hat{P}_{h,s_h^k,a_h^k}^k, \bar{V}_{k,h+1}) \log(2HSAk^2)}{n_k(h, s_h^k, a_h^k) + 1}} + \frac{65H \log(2HSAk^2)}{n_k(h, s_h^k, a_h^k) + 1} \right) \sqrt{\log(40k^4)} \\ & \quad + \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s_h^k, a_h^k) + 1}} \cdot \sqrt{\log(40k^4)} \\ &\leq \left(\sqrt{\frac{64H^2 \log(2HSAk^2)}{\alpha_k}} + \frac{65H \log(2HSAk^2)}{\alpha_k} + \sqrt{\frac{\log(2HSAk^2)}{\alpha_k}} \right) \sqrt{\log(40k^4)} \\ & \quad \text{(Event } \mathcal{E}_k^{th} \text{ implies } n_k(h, s_h^k, a_h^k) \geq \alpha_k) \\ &\leq \sqrt{\frac{64}{200}} + \frac{65}{200H} + \sqrt{\frac{1}{200H^2}} \leq 1. \end{aligned}$$

Thus, we have $\gamma_{\text{Be}}^k(h, s, a) \leq 1$ and we can similarly check that $\gamma_{\text{Ho}}^k(h, s, a) \leq 1$. As a result, we have

$$|\bar{Q}_{h,k}(s_h^k, a_h^k)| \leq 2(H - h + 1),$$

which completes the proof. \square

C Optimism

Let \mathcal{H}_h^k denote the history trajectory, which is defined as

$$\mathcal{H}_h^k = \{(s_l^j, a_l^j, r_l^j) : j \leq k \text{ and } l \leq H \text{ if } j < k, \text{ else } l \leq h\}. \quad (24)$$

We will prove that for both types of noise, $\bar{V}_{h,k}$ is optimistic with constant probability under certain conditions.

C.1 Hoeffding-type Noise

Lemma 9. *Condition on history \mathcal{H}_H^{k-1} , if $\mathcal{G}_{k,\text{Ho}}$ holds and Hoeffding-based noise is applied, then $\bar{V}_{h,k}$ is optimistic with constant probability for any $h \in [H]$. Specifically, we have*

$$\mathbb{P}(\bar{V}_{h,k}(s) \geq V_h^*(s), \forall h \in [H], s \in \mathcal{S} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_{k,\text{Ho}}) \geq \Phi(1.9) - \Phi(1) := C_{\text{Ho}}.$$

Proof. We will show that if $\hat{z}_k \geq 1$, then for all $h \in [H]$ and $s \in \mathcal{S}$, we have $\bar{V}_{h,k}(s) \geq V_h^*(s)$. The proof will use induction and the argument is true for $h = H + 1$ as $\bar{V}_{H+1,k}(s) = V_{H+1}^*(s) = 0$. Suppose the argument is true for timestep $h + 1$ and for timestep h we have

$$\begin{aligned}
\bar{V}_{h,k}(s) &= \text{clip}_{2(H-h+1)} \left(\max_{a \in \mathcal{A}} \bar{Q}_{h,k}(s, a) \right) \\
&\geq \min \left\{ 2(H-h+1), \max_{a \in \mathcal{A}} \bar{Q}_{h,k}(s, a) \right\} \\
&\geq \min \left\{ (H-h+1), \bar{Q}_{h,k}(s, \pi_h^*(s)) \right\} \\
&\geq \min \left\{ (H-h+1), \hat{R}_{h,s,\pi_h^*(s)}^k + \left\langle \hat{P}_{h,s,\pi_h^*(s)}^k, \bar{V}_{h+1,k} \right\rangle + \sigma_{\text{ty}}^k(h, s, \pi_h^*(s)) \hat{z}_k \right\} \\
&\geq \min \left\{ (H-h+1), \hat{R}_{h,s,\pi_h^*(s)}^k + \left\langle \hat{P}_{h,s,\pi_h^*(s)}^k, V_{h+1,k}^* \right\rangle + \sigma_{\text{ty}}^k(h, s, \pi_h^*(s)) \hat{z}_k \right\} \\
&\hspace{15em} \text{(Inductive hypothesis)} \\
&\geq \min \left\{ (H-h+1), R_{h,s,\pi_h^*(s)}^k + \left\langle P_{h,s,\pi_h^*(s)}^k, V_{h+1,k}^* \right\rangle \right\} \\
&\hspace{15em} \text{(Since } \hat{M}^k \in \mathcal{M}_{\text{Ho}}^k \text{ inferred by } \mathcal{G}_{k,\text{Ho}} \text{ and } \hat{z}_k \geq 1) \\
&\geq \min \left\{ (H-h+1), Q_h^*(s, \pi_h^*(s)) \right\} \\
&\geq V_h^*(s).
\end{aligned}$$

Then by induction we have that the optimism is achieved for all $h \in [H]$ and $s \in \mathcal{S}$ simultaneously. Meanwhile, as stated in Definition 2, we have $\hat{z}_k \leq \sqrt{\log(40k^4)}$ under event \mathcal{E}_k^w and numerically, $\sqrt{\log(40k^4)} \geq 1.9$. Therefore, the probability that $\hat{z}_k \geq 1$ under \mathcal{E}_k^w , inferred by $\mathcal{G}_{k,\text{Ho}}$, is at least

$$\mathbb{P}(\hat{z}_k \geq 1 \mid \mathcal{H}_H^{k-1}, \mathcal{G}_{k,\text{Ho}}) = \frac{\Phi(1.9) - \Phi(1)}{\Phi(1.9) - \Phi(-1.9)} \geq \Phi(1.9) - \Phi(1) := C_{\text{Ho}}.$$

Thus, we can conclude that

$$\mathbb{P}(\bar{V}_{h,k}(s) \geq V_h^*(s), \forall h \in [H], s \in \mathcal{S} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_{k,\text{Ho}}) \geq C_{\text{Ho}}.$$

□

C.2 Bernstein-type Noise

The following proof of optimism applies some techniques used in [55]. We first present a technical lemma.

Lemma 10. *Let $f_z : \Delta^S \times \mathbb{R}_+^S \times \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$ with $f_z(p, v, n, L) = \frac{n}{n+1} \langle p, v \rangle + \max \left\{ 4\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}}, \frac{64HL}{n+1} \right\} \cdot z$ for some constant $H > 0$ and $z \in \mathbb{R}$. Then, f_z satisfies*

(i) $f_z(p, v, n, L)$ is non-decreasing in $v(s)$ for all $p \in \Delta^S$, $\|v\|_\infty \leq 2H$, $L > 0$, $n \geq 3$ and $z \in [-1.5, 1.5]$

(ii) $f_z(p, v, n, L) \geq \frac{n}{n+1} \langle p, v \rangle + \left(3\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} + \frac{8HL}{n+1} \right) \cdot z$ for $z \in [1, 1.5]$.

(iii) $f_z(p, v, n, L) \leq \frac{n}{n+1} \langle p, v \rangle + \left(3\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} + \frac{8HL}{n+1} \right) \cdot z$ for $z \in [-1.5, -1]$.

Proof. It is obvious that $f_z(p, v, n, L)$ is continuous in $v(s)$ and not differentiable at only one point where $4\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} = \frac{64HL}{n+1}$. Therefore, to prove statement (i), we only need to show that $\frac{\partial f_z(p, v, n, L)}{\partial v(s)} \geq 0$. Specifically, we have

$$\frac{\partial f_z(p, v, n, L)}{\partial v(s)} = \frac{n}{n+1} \cdot p(s) + \mathbb{1} \left\{ 4\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} \geq \frac{64HL}{n+1} \right\} \frac{4p(s)(v(s) - \langle p, v \rangle)L}{\sqrt{(n+1)\mathbb{V}(p,v)L}} \cdot z$$

$$\begin{aligned}
&\stackrel{(a)}{\geq} \frac{n}{n+1} \cdot p(s) + \mathbf{1} \left\{ 4\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} \geq \frac{64HL}{n+1} \right\} \frac{-8HL}{\sqrt{(n+1)\mathbb{V}(p,v)L}} \cdot z \\
&\stackrel{(b)}{\geq} p(s) \left(\frac{n}{n+1} - \frac{z}{2} \right) \\
&\geq 0.
\end{aligned}$$

Here, The inequality (a) holds because $\|v\|_\infty \leq 2H$ and v is non-negative, which means to have $v(s) - \langle p, v \rangle \geq -2H$. The inequality (b) above holds because when the condition inside indicator $\mathbf{1}\{\cdot\}$ holds, we will have $\sqrt{(n+1)\mathbb{V}(p,v)L} \geq 16HL$. The last inequality holds because we have $n \geq 3$ and $z \leq 1.5$. Therefore, $f_z(p, v, n, L)$ is non-decreasing in $v(s)$.

For the statement (ii), we consider two cases. First, when $4\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} \geq \frac{64HL}{n+1}$ holds, we have $\frac{8HL}{n+1} \leq \frac{1}{2}\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}}$, which means to have

$$\frac{n}{n+1} \langle p, v \rangle + \left(3\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} + \frac{8HL}{n+1} \right) \cdot z \leq \frac{n}{n+1} \langle p, v \rangle + \frac{7z}{2} \sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} \leq f_z(p, v, n, L).$$

When $4\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} \leq \frac{64HL}{n+1}$ holds, we have $3\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} \leq \frac{48HL}{n+1}$, which similarly leads to

$$\frac{n}{n+1} \langle p, v \rangle + \left(3\sqrt{\frac{\mathbb{V}(p,v)L}{n+1}} + \frac{8HL}{n+1} \right) \cdot z \leq f_z(p, v, n, L).$$

The state (iii) can be shown similarly and thus the proof is complete. \square

Lemma 11. *Condition on history \mathcal{H}_H^{k-1} , if $\mathcal{G}_{k, \text{Ho}}$ holds and Bernstein-based noise is applied, then $\bar{V}_{h,k}$ is optimistic with constant probability for any $h \in [H]$. Specifically, we have*

$$\mathbb{P}(\bar{V}_{h,k}(s) \geq V_h^*(s), \forall h \in [H], s \in \mathcal{S} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_{k, \text{Be}}) \geq \Phi(1.5) - \Phi(1) := C_{\text{Be}}$$

Proof. Similar to what we have discussed in the proof of Lemma 9, under event \mathcal{E}_k^w , we have $\hat{z}_k \in [1, 1.5]$ with probability at least $\Phi(1.5) - \Phi(1) = C_{\text{Be}}$. Then, we will show that $\bar{Q}_{h,k}(s, a) \geq Q_h^*(s, a)$ for any h with arbitrary s, a and $\hat{z}_k \in [1, 1.5]$. The proof will use induction. For simplicity, let $L = \log(2HSAk^2)$.

For $h = H + 1$, the inequality holds trivially because both sides are 0. Then, by assuming $\bar{Q}_{h+1,k}(s, a) \geq Q_h^*(s, a)$ for any (s, a) such that $n_k(h, s, a) \geq 3$, we have

$$\begin{aligned}
\bar{Q}_{h,k}(s, a) &= \hat{R}_{h,s,a}^k + \langle \hat{P}_{h,s,a}^k, \bar{V}_{h+1,k} \rangle + \sigma_{\text{Be}}^k(h, s, a) \hat{z}_k \\
&\geq R_{h,s,a} + \langle \hat{P}_{h,s,a}^k, \bar{V}_{h+1,k} \rangle + \left(4\sqrt{\frac{\mathbb{V}(\hat{P}_{h,s,a}^k, \bar{V}_{h+1,k})L}{n_k(h, s, a) + 1}} + \frac{64HL}{n_k(h, s, a) + 1} \right) \cdot \hat{z}_k \\
&\quad \text{(Replace } \hat{R}_{h,s,a}^k \text{ by } R_{h,s,a} \text{ through applying event } \mathcal{E}_k^1 \text{ defined in (19))} \\
&\geq R_{h,s,a} + \langle \hat{P}_{h,s,a}^k, \bar{V}_{h+1,k} \rangle + \max \left\{ 4\sqrt{\frac{\mathbb{V}(\hat{P}_{h,s,a}^k, \bar{V}_{h+1,k})L}{n_k(h, s, a) + 1}}, \frac{64HL}{n_k(h, s, a) + 1} \right\} \cdot \hat{z}_k \\
&\stackrel{(a)}{\geq} R_{h,s,a} + \langle \hat{P}_{h,s,a}^k, V_{h+1}^* \rangle + \max \left\{ 4\sqrt{\frac{\mathbb{V}(\hat{P}_{h,s,a}^k, V_{h+1}^*)L}{n_k(h, s, a) + 1}}, \frac{64HL}{n_k(h, s, a) + 1} \right\} \cdot \hat{z}_k \\
&\geq R_{h,s,a} + \langle \hat{P}_{h,s,a}^k, V_{h+1}^* \rangle + \left(3\sqrt{\frac{\mathbb{V}(\hat{P}_{h,s,a}^k, V_{h+1}^*)L}{n_k(h, s, a) + 1}} + \frac{8HL}{n_k(h, s, a) + 1} \right) \cdot \hat{z}_k \\
&\quad \text{(By applying statement (ii) of Lemma 10)}
\end{aligned}$$

$$\begin{aligned}
&\geq R_{h,s,a} + \langle \hat{P}_{h,s,a}^k, V_{h+1}^* \rangle + \sqrt{\frac{6\mathbb{V}(\tilde{P}_{h,s,a}^k, V_{h+1}^*)}{n_k(h,s,a)+1}} + \frac{8HL}{n_k(h,s,a)+1} \quad (\text{Since } \hat{z}_k \geq 1) \\
&\geq R_{h,s,a} + \langle P_{h,s,a}, V_{h+1}^* \rangle \quad (\text{By applying event } \mathcal{E}_k^2 \text{ defined in (20)}) \\
&= Q_h^*(s,a).
\end{aligned}$$

Here, the above inequality (a) holds by applying inductive hypothesis and statement (i) in Lemma 10. It is applicable because when \mathcal{E}_k^w holds, and by the clipping function, $\|\bar{V}_{h+1,k}\|_\infty \leq 2H$. When $n_k(h,s,a) < 3$, $\bar{Q}_{h,k}(s,a) \geq Q_h^*(s,a)$ holds trivially because $Q_h^*(s,a) \leq H$ by definition. Therefore, the induction is complete.

Now, for arbitrary (k, h, s) , set $a = \operatorname{argmax}_{a \in \mathcal{A}} \bar{Q}_{h,k}(s, a)$ and we have

$$\begin{aligned}
\bar{V}_{h,k}(s) &= \operatorname{clip}_{2(H-h+1)} \left(\max_{a \in \mathcal{A}} \bar{Q}_{h,k}(s, a) \right) \\
&\geq \min \left\{ 2(H-h+1), \max_{a \in \mathcal{A}} \bar{Q}_{h,k}(s, a) \right\} \\
&\geq \min \left\{ (H-h+1), \bar{Q}_{h,k}(s, \pi_h^*(s)) \right\} \\
&\geq \min \left\{ (H-h+1), Q_{h,k}^*(s, \pi_h^*(s)) \right\} \\
&\geq V_{h,k}^*(s)
\end{aligned}$$

□

D Pessimism

Similar to what we have proved in Section C, in this section we will prove that for both types of noise, $\bar{V}_{h,k}$ is pessimistic with constant probability under certain conditions.

D.1 Hoeffding-type Noise

Lemma 12. *Condition on history \mathcal{H}_H^{k-1} , if $\mathcal{G}_{k,\text{Ho}}$ holds and Hoeffding-based noise is applied, then $\bar{V}_{h,k}$ is optimistic with constant probability for any $h \in [H]$. Specifically, we have*

$$\mathbb{P}(\bar{V}_{h,k}(s) \leq V_h^*(s), \forall h \in [H], s \in \mathcal{S} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_{k,\text{Ho}}) \geq \Phi(1.9) - \Phi(1) := C_{\text{Ho}}.$$

Proof. We will show that if $\hat{z}_k \leq -1$, then for all $h \in [H]$ and $s \in \mathcal{S}$, we have $\bar{V}_{h,k}(s) \leq V_h^*(s)$. The proof will use induction and the argument is true for $h = H+1$ as $\bar{V}_{H+1,k}(s) = V_{H+1}^*(s) = 0$. Suppose the argument is true for timestep $h+1$ and we consider timestep h . Set $a = \operatorname{argmax}_{a \in \mathcal{A}} \bar{Q}_{h,k}(s, a)$.

$$\begin{aligned}
\bar{V}_{h,k}(s) &= \operatorname{clip}_{2(H-h+1)} (\bar{Q}_{h,k}(s, a)) \\
&\leq \max \left\{ -2(H-h+1), \bar{Q}_{h,k}(s, a) \right\} \\
&\leq \max \left\{ -(H-h+1), \bar{Q}_{h,k}(s, a) \right\} \\
&\leq \max \left\{ -(H-h+1), \hat{R}_{h,s,a}^k + \langle \hat{P}_{h,s,a}^k, \bar{V}_{h+1,k} \rangle + \sigma_{\text{ty}}^k(h, s, a) \hat{z}_k \right\} \\
&\leq \max \left\{ -(H-h+1), \hat{R}_{h,s,a}^k + \langle \hat{P}_{h,s,a}^k, V_{h+1,k}^* \rangle + \sigma_{\text{ty}}^k(h, s, a) \hat{z}_k \right\} \\
&\hspace{15em} (\text{Induction Hypothesis}) \\
&\leq \max \left\{ -(H-h+1), R_{h,s,a}^k + \langle P_{h,s,a}^k, V_{h+1,k}^* \rangle \right\} \quad (\text{Since } \hat{M}^k \in \mathcal{M}_{\text{Ho}}^k \text{ and } \hat{z}_k \leq -1) \\
&\leq \max \left\{ -(H-h+1), Q_h^*(s, a) \right\} \\
&\leq \max \left\{ -(H-h+1), \max_{a \in \mathcal{A}} Q_h^*(s, a) \right\} \\
&\leq V_h^*(s).
\end{aligned}$$

Then by induction we have that the optimism is achieved for all $h \in [H]$ and $s \in \mathcal{S}$ simultaneously. By using argument similar to the proof of Lemma 9, we can see that when $\hat{z}_k \leq -1$, we have $\bar{V}_{h,k}(s) \leq V_h^*(s)$ and this hold simultaneously for any $h \in [H]$, $s \in \mathcal{S}$. Furtherm as stated in Definition 2, we have $|\hat{z}_k| \leq \sqrt{\log(40k^4)}$ under event \mathcal{E}_k^w and numerically, $\sqrt{\log(40k^4)} \geq 1.9$. Therefore, the probability that $\hat{z}_k \leq -1$ under \mathcal{E}_k^w is at least

$$\mathbb{P}(\hat{z}_k \leq -1 \mid \mathcal{H}_H^{k-1}, \mathcal{G}_{k, \text{Ho}}) = \frac{\Phi(1.9) - \Phi(1)}{\Phi(1.9) - \Phi(-1.9)} \geq \Phi(1.9) - \Phi(1) = C_{\text{Ho}}.$$

Thus, we can conclude that

$$\mathbb{P}(\bar{V}_{h,k}(s) \leq V_h^*(s), \forall h \in [H], s \in \mathcal{S} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_{k, \text{Ho}}) \geq C_{\text{Ho}}.$$

□

D.2 Bernstein-type Noise

Lemma 13. *Condition on history \mathcal{H}_H^{k-1} , if $\mathcal{G}_{k, \text{Ho}}$ holds and Bernstein-based noise is applied, then $\bar{V}_{h,k}$ is pessimistic with constant probability for any $h \in [H]$. Specifically, we have*

$$\mathbb{P}(\bar{V}_{h,k}(s) \leq V_h^*(s), \forall h \in [H], s \in \mathcal{S} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_{k, \text{Be}}) \geq C_{\text{Be}}$$

Proof. Similar to what we have discussed in the proof of Lemma 12, under event \mathcal{E}_k^w , we have $\hat{z}_k \in [-1.5, -1]$ with probability at least $\Phi(1.5) - \Phi(1) = C_{\text{Be}}$. Then, we will show that $\bar{Q}_{h,k}(s, a) \leq Q_h^*(s, a)$ for any h with arbitrary s, a and $\hat{z}_k \in [-1.5, -1]$. The proof will go by induction. For simplicity, let $L = \log(2HSAk^2)$.

For $h = H + 1$, the inequality holds trivially because both sides are 0. Then, by assuming $\bar{Q}_{h+1,k}(s, a) \leq Q_h^*(s, a)$ for any (s, a) such that $n_k(h, s, a) \geq 3$, we have

$$\begin{aligned} \bar{Q}_{h,k}(s, a) &= \hat{R}_{h,s,a}^k + \left\langle \hat{P}_{h,s,a}^k, \bar{V}_{h+1,k} \right\rangle + \sigma_{\text{Be}}^k(h, s, a) \hat{z}_k \\ &\leq R_{h,s,a} + \left\langle \hat{P}_{h,s,a}^k, \bar{V}_{h+1,k} \right\rangle - \left(4 \sqrt{\frac{\mathbb{V}(\tilde{P}_{h,s,a}^k, \bar{V}_{h+1,k}) L}{n_k(h, s, a) + 1}} + \frac{64HL}{n_k(h, s, a) + 1} \right) \\ &\quad \text{(Replace } \hat{R}_{h,s,a} \text{ by } R_{h,s,a} \text{ through applying event } \mathcal{E}_k^1 \text{ defined in (19))} \\ &\stackrel{\text{(a)}}{\leq} R_{h,s,a} + \left\langle \hat{P}_{h,s,a}^k, V_{h+1}^* \right\rangle - \max \left\{ 4 \sqrt{\frac{\mathbb{V}(\tilde{P}_{h,s,a}^k, V_{h+1}^*) L}{n_k(h, s, a) + 1}}, \frac{64HL}{n_k(h, s, a) + 1} \right\} \\ &\leq R_{h,s,a} + \left\langle \hat{P}_{h,s,a}^k, V_{h+1}^* \right\rangle - \left(3 \sqrt{\frac{\mathbb{V}(\tilde{P}_{h,s,a}^k, V_{h+1}^*) L}{n_k(h, s, a) + 1}} + \frac{8HL}{n_k(h, s, a) + 1} \right) \\ &\quad \text{(By applying statement (iii) of Lemma 10)} \\ &\leq R_{h,s,a} + \left\langle \hat{P}_{h,s,a}^k, V_{h+1}^* \right\rangle - \sqrt{\frac{6\mathbb{V}(\tilde{P}_{h,s,a}^k, V_{h+1}^*)}{n_k(h, s, a) + 1}} - \frac{8HL}{n_k(h, s, a) + 1} \\ &\leq R_{h,s,a} + \left\langle P_{h,s,a}, V_{h+1}^* \right\rangle \quad \text{(By applying event } \mathcal{E}_k^2 \text{ defined in (20))} \\ &= Q_h^*(s, a). \end{aligned}$$

Here, the above inequality (a) holds by applying inductive hypothesis and statement (i) in Lemma 10. It is applicable because when \mathcal{E}_k^w holds, by the clipping function, $\|\bar{V}_{h+1,k}\|_\infty \leq 2H$. When $n_k(h, s, a) < 3$, $\bar{Q}_{h,k}(s, a) \leq Q_h^*(s, a)$ holds trivially because $0 \leq Q_h^*(s, a) \leq H$ by definition. Therefore, the induction is complete.

Now, for arbitrary (k, h, s) , set $a = \operatorname{argmax}_{a \in \mathcal{A}} \bar{Q}_{h,k}(s, a)$ and we have

$$\bar{V}_{h,k}(s) = \operatorname{clip}_{2(H-h+1)}(\bar{Q}_{h,k}(s, a))$$

$$\begin{aligned}
&\leq \max \{-2(H-h+1), \bar{Q}_{h,k}(s,a)\} \\
&\leq \max \{-(H-h+1), \bar{Q}_{h,k}(s,a)\} \\
&\leq \max \{-(H-h+1), Q_{h,k}^*(s,a)\} \\
&\leq \max \{-(H-h+1), Q_{h,k}^*(s, \pi_h^*(s))\} \\
&\leq V_{h,k}^*(s).
\end{aligned}$$

□

E Regret Decomposition

In this section, we prove the multiple lemmas necessary for bounding the regret. The regret is mainly composed of two terms, the pessimism term and the estimation error term. The pessimism term, $V_{1,k}^*(s_1^k) - \bar{V}_{1,k}(s_1^k)$, measures how much regret is due to the value the algorithm uses, $\bar{V}_{1,k}$, is smaller than the true value, $V_{1,k}^*(s_1^k)$. The estimation error term, $\bar{V}_{1,k}(s_1^k) - V_{1,k}^{\pi^k}(s_1^k)$ measure how much regret is due to the value, $\bar{V}_{1,k}$, does not estimate $V_{1,k}^{\pi^k}(s_1^k)$, the true value of the policy π^k accurately.

We first introduce a few definitions key to this section. In this section, we omit k if it is clear from the context. Let $a_h^k = \pi_h^k(s_h^k)$ unless specified otherwise.

Definition 5. Let $\mathcal{P}_{h,s,a}^k = \langle \hat{P}_{h,s,a}^k - P_{h,s,a}, V_{h+1}^* \rangle$ and $\mathcal{R}_{h,s,a}^k = \hat{R}_{h,s,a}^k - R_{h,s,a}$.

Definition 6 ($\underline{M}_{\text{ty}}^k$ and $\underline{V}_{h,k}$). Given history \mathcal{H}_H^{k-1} (defined in equation (24)), \hat{P}^k and \hat{R}^k , we define $\underline{w}_{\text{ty}}^k(h, s, a) = -\gamma_{\text{ty}}^k(h, s, a)$ and $\underline{V}_{h,k}$ be the value function obtained by running policy π^k on the MDP $\underline{M}_{\text{ty}}^k = (H, \mathcal{S}, \mathcal{A}, \hat{P}^k, \hat{R}^k + \underline{w}_{\text{ty}}^k, s_1^k)$ plus a magnitude clipping with threshold $2(H-h+1)$.

Definition 7 (\bar{M}_{ty}^k and $\bar{V}_{h,k}$). Given history \mathcal{H}_H^{k-1} (defined in equation (24)), \hat{P}^k and \hat{R}^k , we define $\bar{w}_{\text{ty}}^k(h, s, a) = \gamma_{\text{ty}}^k(h, s, a)$ and $\bar{V}_{h,k}$ be the value function obtained by running policy π^k on the MDP $\bar{M}_{\text{ty}}^k = (H, \mathcal{S}, \mathcal{A}, \hat{P}^k, \hat{R}^k + \bar{w}_{\text{ty}}^k, s_1^k)$ plus a magnitude clipping with threshold $2(H-h+1)$.

Similar to Lemma 8, we can also show that under good event \mathcal{G}_k and $\mathcal{E}_{h,k}^{\text{th}}$, no clipping happens on s_h^k for $\underline{V}_{h,k}(s_h^k)$ and $\bar{V}_{h,k}(s_h^k)$.

Lemma 14. Under the good event \mathcal{G}_k , we have $\underline{V}_{h,k}(s) \leq \bar{V}_{h,k}(s) \leq \bar{\bar{V}}_{h,k}(s)$ for all $h \in [H]$, $s \in \mathcal{S}$.

Proof. This is an immediate result by noticing that under good event \mathcal{G}_k , we have $\underline{w}_{\text{ty}}^k(h, s, a) \leq w_{\text{ty}}^k(h, s, a) \leq \bar{w}_{\text{ty}}^k(h, s, a)$ for all $h \in [H]$ and $s \in \mathcal{S}$. □

Definition 8. Define $\underline{\delta}_h^\pi(s_h)$, $\bar{\delta}_h^\pi(s_h)$, $\bar{\bar{\delta}}_h^\pi(s_h)$, $\delta_h^\pi(s_h)$, $\underline{\delta}_h(s_h)$, $\bar{\delta}_h(s_h)$ and $\bar{\bar{\delta}}_h(s_h)$ as

$$\begin{aligned}
\underline{\delta}_h^\pi(s_h) &= \underline{V}_h(s_h) - V_h^\pi(s_h), \\
\bar{\delta}_h^\pi(s_h) &= \bar{V}_h(s_h) - V_h^\pi(s_h), \\
\bar{\bar{\delta}}_h^\pi(s_h) &= \bar{\bar{V}}_h(s_h) - V_h^\pi(s_h), \\
\delta_h^\pi(s_h) &= V_h^*(s_h) - V_h^\pi(s_h), \\
\underline{\delta}_h(s_h) &= \underline{V}_h(s_h) - V_h^*(s_h), \\
\bar{\delta}_h(s_h) &= \bar{V}_h(s_h) - V_h^*(s_h), \\
\bar{\bar{\delta}}_h(s_h) &= \bar{\bar{V}}_h(s_h) - V_h^*(s_h).
\end{aligned}$$

Definition 9. We denote the history trajectory $\bar{\mathcal{H}}_h^k = \mathcal{H}_h^k \cup \{\hat{z}_k\}$. With filtration sets $\{\bar{\mathcal{H}}_h^k\}_{h,k}$, we define the following sequences:

$$M_{\delta_h(s_h)} = \mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{\text{cum}}\} [(P_{h,s_h,a_h}, \delta_{h+1}) - \delta_{h+1}(s_{h+1})],$$

where $\delta \in \{\underline{\delta}^\pi, \bar{\delta}^\pi, \bar{\bar{\delta}}^\pi, \delta^\pi, \underline{\delta}, \bar{\delta}, \bar{\bar{\delta}}\}$. We will show the sequences are martingales in Lemma 22.

Finally, the regret can be decomposed as

$$\begin{aligned}
& \text{Regret}(M, K, \text{SSR}_{\text{ty}}) \\
&= \sum_{k=1}^K \left(V_1^*(s_1^k) - V_{1,k}^{\pi^k}(s_1^k) \right) \\
&= \sum_{k=1}^K \mathbb{1}(\mathcal{C}_{\text{ty}}^k) \left(V_1^*(s_1^k) - V_{1,k}^{\pi^k}(s_1^k) \right) + \sum_{k=1}^K \mathbb{1}\left(\left(\mathcal{C}_{\text{ty}}^k\right)^c\right) \left(V_1^*(s_1^k) - V_{1,k}^{\pi^k}(s_1^k) \right) \\
&= \sum_{k=1}^K \mathbb{1}\{\mathcal{C}_{\text{ty}}^k\} \left(\underbrace{V_{1,k}^*(s_1^k) - \bar{V}_{1,k}(s_1^k)}_{\text{pessimism term} = -\bar{\delta}_{1,k}(s_1^k)} + \underbrace{\bar{V}_{1,k}(s_1^k) - V_{1,k}^{\pi^k}(s_1^k)}_{\text{estimation error term} = \bar{\delta}_{1,k}^{\pi^k}(s_1^k)} \right) \\
&\quad + \underbrace{\sum_{k=1}^K \mathbb{1}\left(\left(\mathcal{C}_{\text{ty}}^k\right)^c\right) \left(V_1^*(s_1^k) - V_{1,k}^{\pi^k}(s_1^k) \right)}_{(a)}.
\end{aligned}$$

By Lemma 5 and 6, we know that

$$\mathbb{E} \left[\sum_{k=1}^K \mathbb{1}\left(\left(\mathcal{C}_{\text{ty}}^k\right)^c\right) \right] = \sum_{k=1}^K \mathbb{P}\left(\left(\mathcal{C}_{\text{ty}}^k\right)^c\right) \leq \sum_{k=1}^{\infty} \mathbb{P}\left(\left(\mathcal{C}_{\text{ty}}^k\right)^c\right) \leq \frac{\pi^2}{3}.$$

Therefore, by standard Hoeffding's inequality, it holds with probability at least $1 - \delta$ that

$$\sum_{k=1}^K \mathbb{1}\left(\left(\mathcal{C}_{\text{ty}}^k\right)^c\right) \leq \frac{\pi^2}{3} + \sqrt{\frac{\log(1/\delta)}{2K}}.$$

Since the value functions of true MDP is bounded in $[0, H]$, with probability at least $1 - \delta$, we have

$$(a) \leq H \sum_{k=1}^K \mathbb{1}\left(\left(\mathcal{C}_{\text{ty}}^k\right)^c\right) \leq \frac{\pi^2 H}{3} + H \sqrt{\frac{\log(1/\delta)}{2K}} = \tilde{O}(H).$$

Further, notice that the good event $\mathcal{G}_k = \mathcal{C}_{\text{ty}}^k \cap \mathcal{E}_k^w$ and by Lemma 7, we have $\sum_{k=1}^{\infty} \mathbb{P}\left(\left(\mathcal{E}_k^w\right)^c\right) \leq \frac{\pi^2}{3}$. Therefore, we can similarly address the regret incurred by $\left(\mathcal{E}_k^w\right)^c$ as the bound for term (a). As a result, it will be sufficient to only consider $\mathbb{1}\{\mathcal{G}_k\} \left(V_{1,k}^*(s_1^k) - V_{1,k}^{\pi^k}(s_1^k) \right)$ when bounding pessimism and estimation error terms. That is, with probability at least $1 - \delta$, it holds that

$$\text{Regret}(M, K, \text{SSR}_{\text{ty}}) \leq \sum_{k=1}^K \mathbb{1}\{\mathcal{G}_k\} \left(\left| \bar{\delta}_{1,k}(s_1^k) \right| + \left| \bar{\delta}_{1,k}^{\pi^k}(s_1^k) \right| \right) + \tilde{O}(H). \quad (25)$$

Then, we decompose the estimation error term in Section E.2. We decompose the pessimism term in Section E.1. We combine the decomposition of the pessimism term and the estimation error term in Section E.3.

E.1 Pessimism Term

Lemma 15. Let $C_1 = \max \left\{ \frac{1}{\Phi(1.9) - \Phi(1)}, \frac{1}{\Phi(1.5) - \Phi(1)} \right\} = \frac{1}{\Phi(1.5) - \Phi(1)} \approx 10.9$. Then, for any h, k, s_h^k and the type of noise we used, under the good event \mathcal{G}_k , the following bound holds,

$$\mathbb{1}\{\mathcal{G}_k\} \left| \bar{\delta}_{h,k}(s_h^k) \right| \leq \mathbb{1}\{\mathcal{G}_k\} C_1 \left(\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| + \left| \delta_{h,k}^{\pi^k}(s_h^k) \right| \right). \quad (26)$$

Proof. Let \mathcal{O}_k be the event that $\bar{V}_{h,k}(s) \geq V_h^*(s)$ simultaneously for all $s \in \mathcal{S}$ and $h \in [H]$. By Lemma 9 and 11, we know that $\mathbb{P}(\mathcal{O}_k \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k) \geq \min \{ \Phi(1.9) - \Phi(1), \Phi(1.5) - \Phi(1) \} = \Phi(1.5) - \Phi(1)$, which means $\frac{1}{\mathbb{P}(\mathcal{O}_k)} \leq C_1$ regardless the type of noise used.

The definition of \mathcal{O}_k implies $V_h^* \leq \mathbb{E} [\bar{V}_{h,k} \mid \mathcal{O}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k]$. Meanwhile, notice that

$$\begin{aligned}
& \mathbb{1}\{\mathcal{G}_k\} \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \underline{V}_{h,k} \right) \\
&= \mathbb{1}\{\mathcal{G}_k\} \mathbb{P}(\mathcal{O}_k \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k) \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{O}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \underline{V}_{h,k} \right) \\
&\quad + \mathbb{1}\{\mathcal{G}_k\} \mathbb{P}((\mathcal{O}_k)^c \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k) \underbrace{\left(\mathbb{E} [\bar{V}_{h,k} \mid (\mathcal{O}_k)^c, \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \underline{V}_{h,k} \right)}_{(a) \geq 0} \\
&\geq \mathbb{1}\{\mathcal{G}_k\} \mathbb{P}(\mathcal{O}_k \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k) \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{O}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \underline{V}_{h,k} \right) \\
&\implies \mathbb{1}\{\mathcal{G}_k\} \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{O}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \underline{V}_{h,k} \right) \leq \mathbb{1}\{\mathcal{G}_k\} C_1 \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \underline{V}_{h,k} \right).
\end{aligned}$$

Here, we have term (a) ≥ 0 since $\underline{V}_{h,k} \leq \bar{V}_{h,k}$ under event \mathcal{G}_k , by Lemma 14.

Therefore, we have

$$\begin{aligned}
\mathbb{1}\{\mathcal{G}_k\} (V_h^*(s_h^k) - \bar{V}_{h,k}(s_h^k)) &\leq \mathbb{1}\{\mathcal{G}_k\} \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{O}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k] (s_h^k) - \bar{V}_{h,k}(s_h^k) \right) \\
&\leq \mathbb{1}\{\mathcal{G}_k\} \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{O}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k] (s_h^k) - \underline{V}_{h,k}(s_h^k) \right) \\
&\leq \mathbb{1}\{\mathcal{G}_k\} C_1 \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k] (s_h^k) - \underline{V}_{h,k}(s_h^k) \right). \quad (27)
\end{aligned}$$

We can similarly use constant probability pessimism shown in Lemma 12 and 13. In particular, let \mathcal{N}_k be the event that $\bar{V}_{h,k}(s) \leq V_h^*(s)$ for all $s \in \mathcal{S}$ and $h \in [H]$. Then, we have

$$\begin{aligned}
& \mathbb{1}\{\mathcal{G}_k\} \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \bar{\bar{V}}_{h,k} \right) \\
&= \mathbb{1}\{\mathcal{G}_k\} \mathbb{P}(\mathcal{N}_k \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k) \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{N}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \bar{\bar{V}}_{h,k} \right) \\
&\quad + \mathbb{1}\{\mathcal{G}_k\} \mathbb{P}((\mathcal{N}_k)^c \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k) \underbrace{\left(\mathbb{E} [\bar{V}_{h,k} \mid (\mathcal{N}_k)^c, \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \bar{\bar{V}}_{h,k} \right)}_{(b) \leq 0} \\
&\leq \mathbb{1}\{\mathcal{G}_k\} \mathbb{P}(\mathcal{N}_k \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k) \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{N}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \bar{\bar{V}}_{h,k} \right) \\
&\implies \mathbb{1}\{\mathcal{G}_k\} \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{N}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \bar{\bar{V}}_{h,k} \right) \geq \mathbb{1}\{\mathcal{G}_k\} C_1 \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k] - \bar{\bar{V}}_{h,k} \right).
\end{aligned}$$

Thus, we have

$$\begin{aligned}
\mathbb{1}\{\mathcal{G}_k\} (V_h^*(s_h^k) - \bar{V}_{h,k}(s_h^k)) &\geq \mathbb{1}\{\mathcal{G}_k\} \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{N}_k, \mathcal{H}_H^{k-1}, \mathcal{G}_k] (s_h^k) - \bar{\bar{V}}_{h,k}(s_h^k) \right) \\
&\geq \mathbb{1}\{\mathcal{G}_k\} C_1 \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k] (s_h^k) - \bar{\bar{V}}_{h,k}(s_h^k) \right). \quad (28)
\end{aligned}$$

Since good event \mathcal{G}_k implies $\underline{V}_{h,k} \leq \bar{V}_{h,k} \leq \bar{\bar{V}}_{h,k}$ by Lemma 14, the RHS of (27) is non-negative and the RHS of (28) is non-positive. Therefore, we can then conclude

$$\begin{aligned}
& \mathbb{1}\{\mathcal{G}_k\} |V_h^*(s_h^k) - \bar{V}_{h,k}(s_h^k)| \\
&\leq \mathbb{1}\{\mathcal{G}_k\} C_1 \left(\left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k] (s_h^k) - \underline{V}_{h,k}(s_h^k) \right) - \left(\mathbb{E} [\bar{V}_{h,k} \mid \mathcal{H}_H^{k-1}, \mathcal{G}_k] (s_h^k) - \bar{\bar{V}}_{h,k}(s_h^k) \right) \right) \\
&= \mathbb{1}\{\mathcal{G}_k\} C_1 \left(\bar{\bar{V}}_{h,k}(s_h^k) - \underline{V}_{h,k}(s_h^k) \right) \\
&\leq \mathbb{1}\{\mathcal{G}_k\} C_1 \left(\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| + \left| \underline{\delta}_{h,k}^{\pi^k}(s_h^k) \right| \right).
\end{aligned}$$

□

E.2 Estimation Error Term

We first bound the estimation error of $\bar{\bar{V}}$, which can be regarded as the optimistic estimate used in UCB-type algorithms. For convenience, we will ignore notation $\mathbb{1}\{\mathcal{G}_k\}$ in this section since all statements are proved under the good event \mathcal{G}_k .

Lemma 16. *With probability at least $1 - \delta$, for all (k, h, s_h^k) , under the good event \mathcal{G}_k it holds that*

$$\begin{aligned} & \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| \\ & \leq \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left(\left| \mathcal{P}_{h,s_h^k,a_h^k}^k + \mathcal{R}_{h,s_h^k,a_h^k}^k + \bar{w}_{\text{ty}}^k(h, s_h^k, a_h^k) \right| + \mathcal{M}_{\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right|} + \mathcal{M}_{\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right|} + \frac{2SH^2L}{n_k(h, s_h^k, a_h^k)} \right) \\ & + \mathbb{1}\{\mathcal{E}_{h+1,k}^{cum}\} \left(\frac{C_1}{H} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \frac{H+1+C_1}{H} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \frac{1}{H} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right) \\ & + \mathbb{1}\{\mathcal{E}_{h,k}^{cum} \cap (\mathcal{E}_{h+1,k}^{th})^c\} \left(\frac{C_1}{H} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \frac{H+1+C_1}{H} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \frac{1}{H} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right), \end{aligned}$$

where $L = \log(2HS^2AK/\delta)$.

Proof. Since both \bar{V} and V^{π^k} are obtained by choosing actions based on policy π^k under event \mathcal{G}_k , we have

$$\begin{aligned} & \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| \\ & = \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left| \bar{V}_{h,k}(s_h^k) - V_{h,k}^{\pi^k}(s_h^k) \right| \\ & = \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left| \bar{Q}_{h,k}(s_h^k, a_h^k) - Q_{h,k}^{\pi^k}(s_h^k, a_h^k) \right| \quad (\text{Since no clipping under } \mathcal{E}_{h,k}^{cum} \text{ for } \bar{V}_{h,k}(s_h^k)) \\ & = \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left| \hat{R}_{h,s_h^k,a_h^k}^k - R_{h,s_h^k,a_h^k} + \bar{w}_{\text{ty}}^k(h, s_h^k, a_h^k) + \langle \hat{P}_{h,s_h^k,a_h^k}^k, \bar{V}_{h+1,k} \rangle - \langle P_{h,s_h^k,a_h^k}^k, V_{h+1,k}^{\pi^k} \rangle \right| \\ & = \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left| \hat{R}_{h,s_h^k,a_h^k}^k - R_{h,s_h^k,a_h^k} + \bar{w}_{\text{ty}}^k(h, s_h^k, a_h^k) + \langle \hat{P}_{h,s_h^k,a_h^k}^k, \bar{V}_{h+1,k} \rangle - \langle P_{h,s_h^k,a_h^k}^k, V_{h+1,k}^{\pi^k} \rangle \right| \\ & \quad + \left| \langle \hat{P}_{h,s_h^k,a_h^k}^k - P_{h,s_h^k,a_h^k}^k, V_{h+1,k}^* \rangle - \langle \hat{P}_{h,s_h^k,a_h^k}^k - P_{h,s_h^k,a_h^k}^k, V_{h+1,k}^{\pi^k} \rangle \right| \\ & \leq \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left(\left| \mathcal{P}_{h,s_h^k,a_h^k}^k + \mathcal{R}_{h,s_h^k,a_h^k}^k + \bar{w}_{\text{ty}}^k(h, s_h^k, a_h^k) \right| + \left| \langle P_{h,s_h^k,a_h^k}^k, \bar{V}_{h+1,k} - V_{h+1,k}^{\pi^k} \rangle \right| \right) \\ & \quad + \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left| \langle \hat{P}_{h,s_h^k,a_h^k}^k - P_{h,s_h^k,a_h^k}^k, \bar{V}_{h+1,k} - V_{h+1,k}^* \rangle \right| \\ & = \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left(\left| \mathcal{P}_{h,s_h^k,a_h^k}^k + \mathcal{R}_{h,s_h^k,a_h^k}^k + \bar{w}_{\text{ty}}^k(h, s_h^k, a_h^k) \right| + \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \mathcal{M}_{\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right|} \right) \\ & \quad + \mathbb{1}\{\mathcal{E}_{h,k}^{cum}\} \left| \langle \hat{P}_{h,s_h^k,a_h^k}^k - P_{h,s_h^k,a_h^k}^k, \bar{V}_{h+1,k} - V_{h+1,k}^* \rangle \right|. \end{aligned}$$

For the last term, we use Lemma 33 and then for $L = \log(2HS^2AK/\delta)$, with probability at least $1 - \delta$, we have

$$\begin{aligned} & \left| \langle \hat{P}_{h,s_h^k,a_h^k}^k - P_{h,s_h^k,a_h^k}^k, \bar{V}_{h+1,k} - V_{h+1,k}^* \rangle \right| \\ & \leq \sum_{s_{h+1} \in \mathcal{S}} \left| \hat{P}_{h,s_h^k,a_h^k}^k(s_{h+1}) - P_{h,s_h^k,a_h^k}^k(s_{h+1}) \right| \left| \bar{V}_{h+1,k}(s_{h+1}) - V_{h+1,k}^*(s_{h+1}) \right| \\ & \leq \sum_{s_{h+1} \in \mathcal{S}} \left(2\sqrt{\frac{P_{h,s_h^k,a_h^k}^k(s_{h+1})L}{n_k(h, s_h^k, a_h^k)}} + \frac{4L}{3n_k(h, s_h^k, a_h^k)} \right) \left| \bar{\delta}_{h+1,k}(s_{h+1}) \right| \\ & = \sum_{s_{h+1}: P_{h,s_h^k,a_h^k}^k(s_{h+1})n_k(h, s_h^k, a_h^k) \geq 4LH^2} 2P_{h,s_h^k,a_h^k}^k(s_{h+1}) \sqrt{\frac{L}{P_{h,s_h^k,a_h^k}^k(s_{h+1})n_k(h, s_h^k, a_h^k)}} \left| \bar{\delta}_{h+1,k}(s_{h+1}) \right| \\ & \quad + \sum_{s_{h+1}: P_{h,s_h^k,a_h^k}^k(s_{h+1})n_k(h, s_h^k, a_h^k) < 4LH^2} 2\sqrt{\frac{LP_{h,s_h^k,a_h^k}^k(s_{h+1})n_k(h, s_h^k, a_h^k)}{n_k(h, s_h^k, a_h^k)^2}} \left| \bar{\delta}_{h+1,k}(s_{h+1}) \right| \\ & \quad + \frac{4SHL}{3n_k(h, s_h^k, a_h^k)} \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{s_{h+1} \in \mathcal{S}} P_{h, s_h^k, a_h^k}(s_{h+1}) \frac{1}{H} \left| \bar{\delta}_{h+1, k}(s_{h+1}) \right| + \frac{4SHL + 2SH^2\sqrt{L}}{3n_k(h, s_h^k, a_h^k)} \\
&\leq \frac{1}{H} \left| \bar{\delta}_{h+1, k}(s_{h+1}^k) \right| + \mathcal{M}_{|\bar{\delta}_{h, k}(s_h^k)|} + \frac{2SH^2L}{n_k(h, s_h^k, a_h^k)} \\
&\leq \frac{1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{1}{H} \left| \bar{\delta}_{h+1, k}(s_{h+1}^k) \right| + \mathcal{M}_{|\bar{\delta}_{h, k}(s_h^k)|} + \frac{2SH^2L}{n_k(h, s_h^k, a_h^k)} \\
&\hspace{15em} \text{(By triangle inequality)} \\
&\leq \frac{1 + C_1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{C_1}{H} \left| \underline{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \mathcal{M}_{|\bar{\delta}_{h, k}(s_h^k)|} + \frac{2SH^2L}{n_k(h, s_h^k, a_h^k)}. \\
&\hspace{15em} \text{(By using Lemma 15)}
\end{aligned}$$

Combining the above two arguments, we can prove the argument:

$$\begin{aligned}
&\mathbb{1} \{ \mathcal{E}_{h, k}^{cum} \} \left| \bar{\delta}_{h, k}^{\pi^k}(s_h^k) \right| \\
&\leq \mathbb{1} \{ \mathcal{E}_{h, k}^{cum} \} \left(\left| \mathcal{P}_{h, s_h^k, a_h^k}^k + \mathcal{R}_{h, s_h^k, a_h^k}^k + \bar{w}_{\text{ty}}^k(h, s_h^k, a_h^k) \right| + \mathcal{M}_{|\bar{\delta}_{h, k}^{\pi^k}(s_h^k)|} + \mathcal{M}_{|\bar{\delta}_{h, k}(s_h^k)|} + \frac{2SH^2L}{n_k(h, s_h^k, a_h^k)} \right) \\
&\quad + \mathbb{1} \{ \mathcal{E}_{h, k}^{cum} \} \left(\frac{H+1+C_1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{C_1}{H} \left| \underline{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| \right).
\end{aligned}$$

Then, the proof is complete by noticing that $\mathcal{E}_{h+1, k}^{cum} = \mathcal{E}_{h, k}^{cum} \cap \mathcal{E}_{h+1, k}^{th}$. \square

Lemma 17. *With probability at least $1 - \delta$, for all (k, h, s_h^k) , under good event \mathcal{G}_k it holds that*

$$\begin{aligned}
&\mathbb{1} \{ \mathcal{E}_{h, k}^{cum} \} \left| \underline{\delta}_{h, k}^{\pi^k}(s_h^k) \right| \\
&\leq \mathbb{1} \{ \mathcal{E}_{h, k}^{cum} \} \left(\left| \mathcal{P}_{h, s_h^k, a_h^k}^k + \mathcal{R}_{h, s_h^k, a_h^k}^k + \underline{w}_{\text{ty}}^k(h, s_h^k, a_h^k) \right| + \mathcal{M}_{|\underline{\delta}_{h, k}^{\pi^k}(s_h^k)|} + \mathcal{M}_{|\underline{\delta}_{h, k}(s_h^k)|} + \frac{2SH^2L}{n_k(h, s_h^k, a_h^k)} \right) \\
&\quad + \mathbb{1} \{ \mathcal{E}_{h+1, k}^{cum} \} \left(\frac{C_1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{H+1+C_1}{H} \left| \underline{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| \right) \\
&\quad + \mathbb{1} \{ \mathcal{E}_{h, k}^{cum} \cap (\mathcal{E}_{h+1, k}^{th})^c \} \left(\frac{C_1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{H+1+C_1}{H} \left| \underline{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| \right).
\end{aligned}$$

Proof. The proof exactly follows the proof of Lemma 16. \square

Lemma 18. *With probability at least $1 - \delta$, for all (k, h, s_h^k) , under good event \mathcal{G}_k it holds that*

$$\begin{aligned}
&\mathbb{1} \{ \mathcal{E}_{h, k}^{cum} \} \left| \bar{\delta}_{h, k}^{\pi^k}(s_h^k) \right| \\
&\leq \mathbb{1} \{ \mathcal{E}_{h, k}^{cum} \} \left(\left| \mathcal{P}_{h, s_h^k, a_h^k}^k + \mathcal{R}_{h, s_h^k, a_h^k}^k + w_{\text{ty}}^k(h, s_h^k, a_h^k) \right| + \mathcal{M}_{|\bar{\delta}_{h, k}^{\pi^k}(s_h^k)|} + \mathcal{M}_{|\bar{\delta}_{h, k}(s_h^k)|} + \frac{2SH^2L}{n_k(h, s_h^k, a_h^k)} \right) \\
&\quad + \mathbb{1} \{ \mathcal{E}_{h+1, k}^{cum} \} \left(\frac{C_1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{C_1}{H} \left| \underline{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{H+1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| \right) \\
&\quad + \mathbb{1} \{ \mathcal{E}_{h, k}^{cum} \cap (\mathcal{E}_{h+1, k}^{th})^c \} \left(\frac{C_1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{C_1}{H} \left| \underline{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + \frac{H+1}{H} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| \right).
\end{aligned}$$

Proof. The proof exactly follows the proof of Lemma 16. \square

Lemma 19. *With probability at least $1 - \delta$, for all (k, i, s_i^k) , under good event \mathcal{G}_k it holds that*

$$\begin{aligned}
&\mathbb{1} \{ \mathcal{E}_{i, k}^{cum} \} \left(\left| \bar{\delta}_{i, k}^{\pi^k}(s_i^k) \right| + \left| \bar{\delta}_{i, k}^{\pi^k}(s_i^k) \right| + \left| \underline{\delta}_{i, k}^{\pi^k}(s_i^k) \right| \right) \\
&\leq 3e^{3C_1} \left(\sum_{h=i}^H \sqrt{e_{\text{ty}}^k(h, s_h^k, a_h^k)} + \sum_{h=i}^H \gamma_{\text{ty}}^k(h, s_h^k, a_h^k) + \sum_{h=i}^H \frac{SH^2L}{n_k(h, s_h^k, a_h^k)} \right)
\end{aligned}$$

$$\begin{aligned}
& + e^{3C_1} \sum_{h=i+1}^H \mathbf{1} \left\{ (\mathcal{E}_{i+1,k}^{th})^c \right\} \left(\left| \bar{\delta}_{i+1,k}^{\pi^k}(s_{i+1}^k) \right| + \left| \bar{\delta}_{i+1,k}^{\pi^k}(s_{i+1}^k) \right| + \left| \underline{\delta}_{i+1,k}^{\pi^k}(s_{i+1}^k) \right| \right) \\
& + \sum_{h=i}^H \left(1 + \frac{3C_1}{H} \right)^{h-1} \mathbf{1} \left\{ \mathcal{E}_{h,k}^{cum} \right\} \mathcal{M}_{h,k},
\end{aligned}$$

$$\begin{aligned}
\text{where } \mathcal{M}_{h,k} = & \mathcal{M}_{\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right|} + \mathcal{M}_{\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right|} + \mathcal{M}_{\left| \underline{\delta}_{h,k}^{\pi^k}(s_h^k) \right|} \\
& + \mathcal{M}_{\left| \underline{\delta}_{h,k}^{\pi^k}(s_h^k) \right|} + \mathcal{M}_{\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right|} + \mathcal{M}_{\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right|}.
\end{aligned}$$

Proof. By summing results in Lemma 16, Lemma 17 and Lemma 18, we have

$$\begin{aligned}
& \mathbf{1} \left\{ \mathcal{E}_{h,k}^{cum} \right\} \left(\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| + \left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| + \left| \underline{\delta}_{h,k}^{\pi^k}(s_h^k) \right| \right) \\
\leq & \mathbf{1} \left\{ \mathcal{E}_{h+1,k}^{cum} \right\} \left(1 + \frac{3C_1}{H} \right) \left(\left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \left| \underline{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right) \\
& + |w_{\text{ty}}^k(h, s_h^k, a_h^k)| + |\bar{w}_{\text{ty}}^k(h, s_h^k, a_h^k)| + |w_{\text{ty}}^k(h, s_h^k, a_h^k)| + \frac{6SH^2L}{n_k(h, s_h^k, a_h^k)} + \mathbf{1} \left\{ \mathcal{E}_{h,k}^{cum} \right\} \mathcal{M}_{h,k} \\
& + \mathbf{1} \left\{ \mathcal{E}_{h,k}^{cum} \cap (\mathcal{E}_{h+1,k}^{th})^c \right\} \left(1 + \frac{3C_1}{H} \right) \left(\left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \left| \underline{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right) \\
& + 3 \left| \mathcal{P}_{h,s_h^k,a_h^k}^k + \mathcal{R}_{h,s_h^k,a_h^k}^k \right| \\
\stackrel{(i)}{\leq} & \mathbf{1} \left\{ \mathcal{E}_{h+1,k}^{cum} \right\} \left(1 + \frac{3C_1}{H} \right) \left(\left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \left| \underline{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right) \\
& + 3\sqrt{e_{\text{ty}}^k(h, s_h^k, a_h^k)} + 3\gamma_{\text{ty}}^k(h, s_h^k, a_h^k) + \frac{6SH^2L}{n_k(h, s_h^k, a_h^k)} + \mathbf{1} \left\{ \mathcal{E}_{h,k}^{cum} \right\} \mathcal{M}_{h,k} \\
& + \mathbf{1} \left\{ (\mathcal{E}_{h+1,k}^{th})^c \right\} \left(1 + \frac{3C_1}{H} \right) \left(\left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \left| \underline{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right)
\end{aligned}$$

Here, the inequality (i) above holds because of two reasons. Firstly, under event \mathcal{G}_k , we have $|w_{\text{ty}}^k(h, s_h^k, a_h^k)| \leq |\bar{w}_{\text{ty}}^k(h, s_h^k, a_h^k)|$ and

$$\begin{aligned}
\left| \mathcal{P}_{h,s_h^k,a_h^k}^k + \mathcal{R}_{h,s_h^k,a_h^k}^k \right| & = \left| \left\langle \hat{P}_{h,s_h^k,a_h^k}^k - P_{h,s_h^k,a_h^k}, V_{h+1}^* \right\rangle + \left(\hat{R}_{h,s_h^k,a_h^k}^k - R_{h,s_h^k,a_h^k} \right) \right| \\
& \hspace{15em} \text{(By Definition 5)} \\
& \leq \sqrt{e_{\text{ty}}^k(h, s_h^k, a_h^k)}. \hspace{10em} \text{(Under event } \mathcal{G}_k, \hat{M} \in \mathcal{M}_{\text{ty}}^k)
\end{aligned}$$

Then, the proof is complete by using this recursion from $h = i$ to $h = H$ and utilizing the fact that $\left(1 + \frac{3C_1}{H} \right)^H \leq e^{3C_1}$. \square

E.3 Combining Estimation and Pessimism Terms

Lemma 20. *With probability at least $1 - \delta$, it holds that*

$$\begin{aligned}
\text{Regret}(M, K, \text{SSR}_{\text{ty}}) & \leq \mathbf{1} \left\{ \mathcal{G}_k \right\} 3C_1 e^{3C_1} \sum_{k=1}^K \sum_{h=i}^H \left(\sqrt{e_{\text{ty}}^k(h, s_h^k, a_h^k)} + \gamma_{\text{ty}}^k(h, s_h^k, a_h^k) \right) \\
& \quad + \tilde{O} \left(H^4 S^2 A + H\sqrt{T} \right).
\end{aligned}$$

Proof. Recall in equation (25), with probability at least $1 - \delta$, we have

$$\text{Regret}(M, K, \text{SSR}_{\text{ty}})$$

$$\begin{aligned}
&\leq \sum_{k=1}^K \mathbb{1}\{\mathcal{G}_k\} \left(\left| \bar{\delta}_{1,k}^{\pi^k}(s_1^k) \right| + \left| \bar{\delta}_{1,k}^{\pi^k}(s_1^k) \right| \right) + \tilde{O}(H) \\
&\leq \sum_{k=1}^K \mathbb{1}\{\mathcal{G}_k\} C_1 \left(\left| \bar{\delta}_{1,k}^{\pi^k}(s_1^k) \right| + \left| \bar{\delta}_{1,k}^{\pi^k}(s_1^k) \right| + \left| \underline{\delta}_{1,k}^{\pi^k}(s_1^k) \right| \right) + \tilde{O}(H) \quad (\text{By using Lemma 15}) \\
&= \sum_{k=1}^K \mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{1,k}^{cum}\} C_1 \left(\left| \bar{\delta}_{1,k}^{\pi^k}(s_1^k) \right| + \left| \bar{\delta}_{1,k}^{\pi^k}(s_1^k) \right| + \left| \underline{\delta}_{1,k}^{\pi^k}(s_1^k) \right| \right) + \tilde{O}(H) \\
&\quad + \sum_{k=1}^K \mathbb{1}\{(\mathcal{E}_{1,k}^{th})^c\} \left(\left| \bar{\delta}_{1,k}^{\pi^k}(s_1^k) \right| + \left| \bar{\delta}_{1,k}^{\pi^k}(s_1^k) \right| + \left| \underline{\delta}_{1,k}^{\pi^k}(s_1^k) \right| \right) \\
&\leq \mathbb{1}\{\mathcal{G}_k\} 3C_1 e^{3C_1} \left(\sum_{k=1}^K \sum_{h=i}^H \sqrt{e_{\text{ty}}^k(h, s_h^k, a_h^k)} + \sum_{k=1}^K \sum_{h=i}^H \gamma_{\text{ty}}^k(h, s_h^k, a_h^k) + \sum_{k=1}^K \sum_{h=i}^H \frac{SH^2L}{n_k(h, s_h^k, a_h^k)} \right) \\
&\quad + \sum_{k=1}^K \sum_{h=i}^H (1 + \frac{3C_1}{H})^{h-1} \mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{h,k} + \tilde{O}(H) \\
&\quad + \sum_{k=1}^K \sum_{h=1}^H \mathbb{1}\{(\mathcal{E}_{h,k}^{th})^c\} \left(\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| + \left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| + \left| \underline{\delta}_{h,k}^{\pi^k}(s_h^k) \right| \right) \quad (\text{By using Lemma 19}) \\
&\stackrel{(i)}{\leq} \mathbb{1}\{\mathcal{G}_k\} 3C_1 e^{3C_1} \sum_{k=1}^K \sum_{h=i}^H \left(\sqrt{e_{\text{ty}}^k(h, s_h^k, a_h^k)} + \gamma_{\text{ty}}^k(h, s_h^k, a_h^k) \right) + \tilde{O}(H^3 S^2 A + H\sqrt{T}) \\
&\quad + \tilde{O}(H) \sum_{k=1}^K \sum_{h=1}^H \mathbb{1}\{(\mathcal{E}_{h,k}^{th})^c\} \\
&\leq \mathbb{1}\{\mathcal{G}_k\} 3C_1 e^{3C_1} \sum_{k=1}^K \sum_{h=i}^H \left(\sqrt{e_{\text{ty}}^k(h, s_h^k, a_h^k)} + \gamma_{\text{ty}}^k(h, s_h^k, a_h^k) \right) + \tilde{O}(H^4 S^2 A + H\sqrt{T}). \\
&\hspace{15em} (\text{By using Lemma 21})
\end{aligned}$$

The inequality (i) above holds for two reasons. First, it uses Lemma 22 and 24. Second, by our clipping threshold, we know that $\left(\left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| + \left| \bar{\delta}_{h,k}^{\pi^k}(s_h^k) \right| + \left| \underline{\delta}_{h,k}^{\pi^k}(s_h^k) \right| \right) \leq \hat{O}(H)$. \square

Lemma 21 (Lemma 20 in [2]).

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{1}\{(\mathcal{E}_{h,k}^{th})^c\} \leq \tilde{O}(H^3 SA).$$

Proof. It holds that

$$\begin{aligned}
\sum_{k=1}^K \sum_{h=1}^H \mathbb{1}\{(\mathcal{E}_{h,k}^{th})^c\} &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{1}\{n_k(h, s_h^k, a_h^k) \leq \alpha_k\} \\
&\leq \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{h=1}^H \alpha_k \\
&\leq 200H^3 SA \log(2HSAK^2) \log(40K^4) \quad (\text{By our choice of } \alpha_k) \\
&= \tilde{O}(H^3 SA).
\end{aligned}$$

\square

F Bounds on Individual Terms

F.1 Bounds on Martingale Difference

Lemma 22. For $i \in [H]$, the sequences starting from 0 and with difference between two consecutive terms given by $\mathbb{1}\{\mathcal{G}_k\}\mathcal{M}_{h,k}$ for $h = i, \dots, H, k = 1, \dots, K$ are martingales with respect to filtration $\{\overline{\mathcal{H}}_h^k\}_{h=i, \dots, H, k=1, \dots, K}$. Moreover, for any $\delta' > 0$, with probability at least $1 - \delta'$, for any $i \in [H]$, the following hold,

$$\left| \sum_{k=1}^K \sum_{h=i}^H \left(1 + \frac{3C_1}{H}\right)^h \mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{h,k} \right| = \tilde{O}\left(H\sqrt{T}\right).$$

Proof. We first show the sequence starting from 0 and with difference between two consecutive terms given by $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \left(1 + \frac{3C_1}{H}\right)^h \mathcal{M}_{|\overline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$ is a martingale sequence. For any $h \in \{i, \dots, H\}$ and $k \in [K]$,

$$\begin{aligned} & \mathbb{E} \left[\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\overline{\delta}_{h,k}^{\pi^k}(s_h^k)|} \mid \overline{\mathcal{H}}_h^k \right] \\ &= \mathbb{E} \left[\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \left(\left\langle P_{h,s_h^k, a_h^k}, \left| \overline{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right\rangle - \left| \overline{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right) \mid \overline{\mathcal{H}}_h^k \right] = 0. \end{aligned}$$

Similarly, we have $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\overline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$, $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\underline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$, $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\overline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$, $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\underline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$ are martingale difference sequences. As $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{h,k}$ is the sum of several martingale difference sequences, it is a martingale difference sequence.

Next, we bound $\left| \mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\overline{\delta}_{h,k}^{\pi^k}(s_h^k)|} \right|$. When $h = H$, $\left| \mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\overline{\delta}_{h,k}^{\pi^k}(s_h^k)|} \right| = 0$. When \mathcal{G}_k holds, for $h < H$ and any state x ,

$$\begin{aligned} \left| \overline{\delta}_{h+1,k}^{\pi^k}(x) \right| &= \left| \overline{V}_{h+1}(x) - V_{h+1}^{\pi}(x) \right| = \left| \langle P_{h+2,x,\pi(x)}, \overline{V}_{h+2} - V_{h+2}^{\pi} \rangle + w_{\text{ty}}^k(h+1, x, \pi(x)) \right| \\ &\leq \langle P_{h+2,x,\pi(x)}, |\overline{V}_{h+2} - V_{h+2}^{\pi}| \rangle + |w_{\text{ty}}^k(h+1, x, \pi(x))| \end{aligned}$$

By our choice of α_k , when \mathcal{G}_k holds, $|w_{\text{ty}}^k(h, s, a)| \leq \gamma_{\text{ty}}^k(h, s, a) \leq 1$ for all k, h, s, a as shown in Lemma 8. Then, by expanding $\left| \overline{\delta}_{h+1,k}^{\pi^k}(x) \right| = \left| \overline{V}_{h+1}(x) - V_{h+1}^{\pi}(x) \right|$ recursively from $h+1$ to H , we have

$$\left| \mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\overline{\delta}_{h,k}^{\pi^k}(s_h^k)|} \right| \leq 2H \gamma_{\text{ty}}^k(h, s, a) \leq 2H.$$

Similarly, we have the bound on $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\overline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$, $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\underline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$, $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\overline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$, $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\underline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$ and $\mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{|\underline{\delta}_{h,k}^{\pi^k}(s_h^k)|}$.

As a result, $\left| \mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \left(1 + \frac{3C_1}{H}\right)^h \mathcal{M}_{h,k} \right|$ is bounded by $12e^{3C_1}H$. By Azuma-Hoeffding inequality, with probability at least $1 - \delta'$, we have

$$\left| \sum_{k=1}^K \sum_{h=i}^H \left(1 + \frac{3C_1}{H}\right)^h \mathbb{1}\{\mathcal{G}_k \cap \mathcal{E}_{h,k}^{cum}\} \mathcal{M}_{h,k} \right| \leq \tilde{O} \left(\sqrt{\sum_{k=1}^K \sum_{h=i}^H H^2} \right) = \tilde{O} \left(H\sqrt{T} \right).$$

□

F.2 Bounds on Lower-order Terms

The following two lemmas are standard results in literature and we present their proofs here for completeness.

Lemma 23.

$$\sum_{k=1}^K \sum_{h=1}^H \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s_h^k, a_h^k) + 1}} \leq \tilde{O}(\sqrt{HSAT}).$$

Proof. Let $L = \log(2HSAK^2)$. Then, it can be bounded as

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s_h^k, a_h^k) + 1}} &\leq \sqrt{L} \sum_{k=1}^K \sum_{h=1}^H \sqrt{\frac{1}{n_k(h, s_h^k, a_h^k) + 1}} \\ &= \sqrt{L} \sum_{h,s,a} \sum_{n=1}^{n_K(h,s,a)} \sqrt{\frac{1}{n+1}} \\ &\leq \sqrt{L} \sum_{h,s,a} \int_0^{n_K(h,s,a)} \sqrt{\frac{1}{x}} dx \\ &\leq 2\sqrt{L} \sum_{h,s,a} \sqrt{n_K(h,s,a)} \\ &\leq 2\sqrt{L} \cdot \sqrt{HSA \sum_{h,s,a} n_K(h,s,a)} \\ &\hspace{15em} \text{(By Cauchy-Schwartz inequality)} \\ &= \tilde{O}(\sqrt{HSAT}). \hspace{5em} \text{(Since } \sum_{h,s,a} n_K(h,s,a) = T) \end{aligned}$$

□

Lemma 24.

$$\sum_{k=1}^K \sum_{h=1}^H \frac{\log(2HSAk^2)}{n_k(h, s_h^k, a_h^k) + 1} \leq \tilde{O}(HSA).$$

Proof. Let $L = \log(2HSAK^2)$. Then, it can be bounded as

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \frac{\log(2HSAk^2)}{n_k(h, s_h^k, a_h^k) + 1} &\leq L \sum_{k=1}^K \sum_{h=1}^H \frac{1}{n_k(h, s_h^k, a_h^k) + 1} \\ &= L \sum_{h,s,a} \sum_{n=1}^{n_K(h,s,a)} \frac{1}{n+1} \\ &\leq L \sum_{h,s,a} \log(n_K(h,s,a)) \hspace{2em} \text{(Since } \sum_{n=1}^N \frac{1}{n} \leq \log(N) + 1) \\ &\leq LHS A \cdot \max_{h,s,a} \log(n_K(h,s,a)) \\ &\leq LHS A \log(T) \\ &= \tilde{O}(HSA). \end{aligned}$$

□

G Bounds on Sum of Variance

When we use the Bernstein-type noise, the regret analysis needs to bound the sum of variance. This proof applies some techniques developed in [7]. However, since our optimism only holds with

constant probability instead of deterministically, the details are quite different. For simplicity, we first define

$$\begin{aligned}\hat{\mathbb{V}}_{h+1,k}^* &= \mathbb{V}\left(\tilde{P}_{h,s_h^k,a_h^k}^k, V_{h+1}^*\right), & \mathbb{V}_{h+1,k}^* &= \mathbb{V}\left(P_{h,s_h^k,a_h^k}, V_{h+1}^*\right), \\ \hat{\mathbb{V}}_{h+1,k} &= \mathbb{V}\left(\tilde{P}_{h,s_h^k,a_h^k}^k, \bar{V}_{h+1,k}\right), & \bar{\mathbb{V}}_{h+1,k} &= \mathbb{V}\left(P_{h,s_h^k,a_h^k}, \bar{V}_{h+1,k}\right), \\ \mathbb{V}_{h+1,k}^{\pi^k} &= \mathbb{V}\left(P_{h,s_h^k,a_h^k}, V_{h+1,k}^{\pi^k}\right), \\ U_{h,k,1} &= \sqrt{\frac{\hat{\mathbb{V}}_{h+1,k}^* \log(2HSAK^2)}{n_k(h, s_h^k, a_h^k) + 1}}, & U_{h,k,2} &= \sqrt{\frac{\hat{\mathbb{V}}_{h+1,k} \log(2HSAK^2)}{n_k(h, s_h^k, a_h^k) + 1}},\end{aligned}$$

We will first give a full proof of the bound on sum of variance and then present all the auxiliary lemmas in Section G.1.

Lemma 25. *Let $U_{h,k} = U_{h,k,1} + U_{h,k,2}$. For any $\delta > 0$, with probability at least $1 - \delta$, when $T \geq \Omega(H^5 S^2 A)$, it holds that*

$$\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} \leq \tilde{O}\left(H\sqrt{SAT}\right).$$

Proof. First, we have

$$\begin{aligned}\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} &\leq \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \sqrt{\frac{\log(2HSAK^2)}{n_k(h, s_h^k, a_h^k) + 1}} \left(\sqrt{\hat{\mathbb{V}}_{h+1,k}^*} + \sqrt{\hat{\mathbb{V}}_{h+1,k}}\right) \\ &\leq \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \sqrt{\frac{\log(2HSAK^2)}{n_k(h, s_h^k, a_h^k) + 1}} \cdot \sqrt{2} \sqrt{\hat{\mathbb{V}}_{h+1,k}^* + \hat{\mathbb{V}}_{h+1,k}} \\ &\hspace{15em} (\text{Since } \sqrt{a} + \sqrt{b} \leq \sqrt{2(a+b)} \text{ for } a, b \geq 0) \\ &\leq \sqrt{2} \sqrt{\left(\sum_{k=1}^K \sum_{h=1}^{H-1} \frac{\log(2HSAK^2)}{n_k(h, s_h^k, a_h^k) + 1}\right) \left(\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} (\hat{\mathbb{V}}_{h+1,k}^* + \hat{\mathbb{V}}_{h+1,k})\right)} \\ &\hspace{15em} (\text{By Cauchy-Schwartz inequality}) \\ &\leq \sqrt{\tilde{O}(HSA) \left(\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} (\hat{\mathbb{V}}_{h+1,k}^* + \hat{\mathbb{V}}_{h+1,k})\right)},\end{aligned}\tag{29}$$

where the last inequality above applies Lemma 24.

We will then bound the two sums of variance separately. Specifically, by applying Lemma 26 and Lemma 28, we have with probability at least $1 - \delta/3$,

$$\begin{aligned}&\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \hat{\mathbb{V}}_{h+1,k}^* \\ &= \frac{3}{2} \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \mathbb{V}_{h+1,k}^{\pi^k} + \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h+1,k}^* - \frac{3}{2} \mathbb{V}_{h+1,k}^{\pi^k}\right) \\ &\leq \tilde{O}\left(HT + H^2\sqrt{T} + H^3 + H^3 S^2 A + H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \delta_{h+1,k}^{\pi^k} (s_{h+1}^k)\right).\end{aligned}\tag{31}$$

By similarly applying Lemma 26 and Lemma 29, we have with probability at least $1 - \delta/3$,

$$\begin{aligned}&\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \hat{\mathbb{V}}_{h+1,k} \\ &= \frac{3}{2} \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \mathbb{V}_{h+1,k}^{\pi^k} + \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h+1,k} - \frac{3}{2} \mathbb{V}_{h+1,k}^{\pi^k}\right)\end{aligned}\tag{32}$$

$$\leq \tilde{O} \left(HT + H^2\sqrt{T} + H^3 + H^3S^2A + H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right), \quad (33)$$

By combining equations (31) and (33), we have

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h+1,k}^* + \hat{\mathbb{V}}_{h+1,k} \right) \\ & \leq \tilde{O} \left(HT + H^2\sqrt{T} + H^3S^2A + H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\delta_{h+1,k}^{\pi^k}(s_{h+1}^k) + \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right) \right). \end{aligned} \quad (34)$$

Then, by referring to definitions of $\sqrt{e_{\text{Be}}^k(h, s_h^k, a_h^k)}$ and $\gamma_{\text{Be}}^k(h, s_h^k, a_h^k)$, with probability at least $1 - \delta/3$, we have

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\delta_{h+1,k}^{\pi^k}(s_{h+1}^k) + \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right) \\ & \leq \sum_{h=1}^{H-1} \left(\sum_{k=1}^K \mathbb{1}\{\mathcal{G}_k\} \left(\left| \delta_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| \right) \right) \\ & \leq \mathbb{1}\{\mathcal{G}_k\} 3C_1 e^{3C_1} H \sum_{k=1}^K \sum_{h=1}^H \left(\sqrt{e_{\text{ty}}^k(h, s_h^k, a_h^k)} + \gamma_{\text{ty}}^k(h, s_h^k, a_h^k) \right) + \tilde{O} \left(H^5S^2A + H^2\sqrt{T} \right) \\ & \hspace{15em} \text{(By referring to the proof of Lemma 20)} \\ & \leq \tilde{O} \left(H^5S^2A + H^2\sqrt{T} + \sqrt{H^3SAT} + H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} \right) \quad \text{(By Lemma 23 and 24)} \\ & \leq \tilde{O} \left(H^5S^2A + H^2\sqrt{SAT} + H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} \right). \end{aligned} \quad (35)$$

By plugging equation (35) into equation (34), we can have

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h+1,k}^* + \hat{\mathbb{V}}_{h+1,k} \right) \\ & \leq \tilde{O} \left(HT + H^2\sqrt{T} + H^3S^2A + H^6S^2A + H^3\sqrt{SAT} + H^2 \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} \right) \\ & \leq \tilde{O} \left(HT + H^3\sqrt{SAT} + H^6S^2A + H^2 \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} \right) \\ & \leq \tilde{O} \left(HT + H^2 \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} \right) \quad \text{(When } T \geq \Omega(H^5S^2A)\text{)} \end{aligned}$$

Now, by plugging the above result into equation (29), when $T \geq \Omega(H^5S^2A)$, it holds that

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} & \leq \sqrt{\tilde{O} \left(HSA \left(HT + H^2 \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} \right) \right)} \\ & \leq \tilde{O} \left(H\sqrt{SAT} + H^{1.5} \sqrt{\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k}} \right). \end{aligned}$$

It is easy to check that the above inequality implies $\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} U_{h,k} \leq \tilde{O} \left(H\sqrt{SAT} \right)$ and thus the proof is complete. \square

G.1 Auxiliary Lemmas

The lemmas used for proving Lemma 25 are presented as the following.

Lemma 26 (Lemma 8 in [7]). *For any $\delta > 0$, with probability at least $1 - \delta$, it holds that*

$$\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}(\mathcal{G}_k) \mathbb{V}_{h+1,k}^{\pi^k} \leq \tilde{O}\left(HT + H^2\sqrt{T} + H^3\right).$$

Lemma 27. *For any $\delta > 0$, with probability at least $1 - \delta$, for any $k \in [K]$, $h \in [H]$, it holds that*

$$\begin{aligned} \hat{\mathbb{V}}_{h+1,k}^* &\leq \frac{3}{2}\mathbb{V}_{h+1,k}^* + \frac{2H^2S \log(2HS^2AK/\delta)}{n_k(h, s_h^k, a_h^k)}, \\ \hat{\mathbb{V}}_{h+1,k} &\leq \frac{3}{2}\bar{\mathbb{V}}_{h+1,k} + \frac{2H^2S \log(2HS^2AK/\delta)}{n_k(h, s_h^k, a_h^k)}. \end{aligned}$$

Proof. The proof apply some techniques in [55]. Fix some $\delta > 0$ and let $L = \log(2HS^2AK/\delta)$ for simplicity. First, by Lemma 30, for some tuple (k, h, s, a, s') , we have

$$\begin{aligned} &\mathbb{P}\left(\tilde{P}_{h,s,a}^k(s') \geq \frac{3}{2}P_{h,s,a}(s') + \frac{2L}{n_k(h, s, a)}\right) \\ &\leq \mathbb{P}\left(\tilde{P}_{h,s,a}^k(s') - P_{h,s,a}(s') \geq \sqrt{\frac{2P_{h,s,a}(s')L}{n_k(h, s, a)}} + \frac{L}{n_k(h, s, a)}\right) \\ &\hspace{15em} (\text{Since } a + b \geq 2\sqrt{ab} \text{ for } a, b \geq 0) \\ &\leq \frac{\delta}{HS^2AK}. \end{aligned}$$

Then, a union bound says that its complement holds for any (k, h, s, a, s') with probability at least $1 - \delta$. Thus, we have

$$\begin{aligned} \hat{\mathbb{V}}_{h+1,k}^* &= \sum_{s' \in \mathcal{S}} \tilde{P}_{h,s_h^k, a_h^k}^k(s') \left(V_{h+1}^*(s') - \left\langle \tilde{P}_{h,s_h^k, a_h^k}^k, V_{h+1}^* \right\rangle\right)^2 \\ &\leq \sum_{s' \in \mathcal{S}} \tilde{P}_{h,s_h^k, a_h^k}^k(s') \left(V_{h+1}^*(s') - \left\langle P_{h,s_h^k, a_h^k}, V_{h+1}^* \right\rangle\right)^2 \\ &\hspace{15em} (\text{Since } \mathbb{E}[X] \text{ is the minimizer of } \min_x \mathbb{E}[(X - x)^2]) \\ &\leq \sum_{s' \in \mathcal{S}} \left(\frac{3}{2}P_{h,s,a}(s') + \frac{2L}{n_k(h, s, a)}\right) \left(V_{h+1}^*(s') - \left\langle P_{h,s_h^k, a_h^k}, V_{h+1}^* \right\rangle\right)^2 \\ &\leq \frac{3}{2}\mathbb{V}_{h+1,k}^* + \frac{2H^2S \log(2HS^2AK/\delta)}{n_k(h, s_h^k, a_h^k)}. \end{aligned}$$

For $\hat{\mathbb{V}}_{h+1,k}$, we just need to follow a similar argument and thus the proof is complete. \square

Lemma 28. *For any $\delta > 0$, with probability at least $1 - \delta$, it holds that*

$$\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h+1,k}^* - \frac{3}{2}\mathbb{V}_{h+1,k}^{\pi^k}\right) \leq \tilde{O}\left(H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \delta_{h+1,k}^{\pi^k}(s_{h+1}^k) + H^2\sqrt{T} + H^3S^2A\right).$$

Proof. We begin by applying Lemma 27. Thus, with probability at least $1 - \frac{\delta}{2}$, we have

$$\begin{aligned} &\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h+1,k}^* - \frac{3}{2}\mathbb{V}_{h+1,k}^{\pi^k}\right) \\ &\leq \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\frac{3}{2}\mathbb{V}_{h+1,k}^* - \frac{3}{2}\mathbb{V}_{h+1,k}^{\pi^k}\right) + \frac{4H^2S \log(4HS^2AK/\delta)}{3n_k(h, s_h^k, a_h^k)} \quad (\text{By Lemma 27}) \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\frac{3}{2} \mathbb{V}_{h+1,k}^* - \frac{3}{2} \mathbb{V}_{h+1,k}^{\pi^k} \right) + \tilde{O}(H^3 S^2 A) && \text{(By Lemma 24)} \\
&\leq \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \frac{3}{2} \mathbb{E}_{s' \sim P_{h, s_h^k, a_h^k}} \left[\left(V_{h+1}^*(s') \right)^2 - \left(V_{h+1, k}^{\pi^k}(s') \right)^2 \right] + \tilde{O}(H^3 S^2 A) \\
& && \text{(Since } V_{h+1, k}^{\pi^k} \leq V_{h+1}^*) \\
&\leq 3H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \mathbb{E}_{s' \sim P_{h, s_h^k, a_h^k}} \left[\delta_{h+1, k}^{\pi^k}(s') \right] + \tilde{O}(H^3 S^2 A) \\
& && \text{(Since } V_{h+1, k}^{\pi^k} \leq V_{h+1}^* \leq H \text{ and } a^2 - b^2 = (a+b)(a-b)) \\
&\leq \tilde{O} \left(H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \delta_{h+1, k}^{\pi^k}(s_{h+1}^k) + H^2 \sqrt{T} + H^3 S^2 A \right). && \text{(By Lemma 22)}
\end{aligned}$$

The last line above holds because by Lemma 22, with probability at least $1 - \frac{\delta}{2}$, we have

$$\left| \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\mathbb{E}_{s' \sim P_{h, s_h^k, a_h^k}} \left[\delta_{h+1, k}^{\pi^k}(s') \right] - \delta_{h+1, k}^{\pi^k}(s_{h+1}^k) \right) \right| \leq \tilde{O}(H\sqrt{T}).$$

□

Lemma 29. For any $\delta > 0$, with probability at least $1 - \delta$, it hold that

$$\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h+1, k} - \frac{3}{2} \mathbb{V}_{h+1, k}^{\pi^k} \right) \leq \tilde{O} \left(H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left| \bar{\delta}_{h+1, k}^{\pi^k}(s_{h+1}^k) \right| + H^2 \sqrt{T} + H^3 S^2 A \right).$$

Proof. Similarly, we begin by applying Lemma 27 and with probability at least $1 - \frac{\delta}{3}$, we have

$$\begin{aligned}
&\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{\mathbb{V}}_{h+1, k} - \frac{3}{2} \mathbb{V}_{h+1, k}^{\pi^k} \right) \\
&\leq \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\frac{3}{2} \bar{\mathbb{V}}_{h+1, k} - \frac{3}{2} \mathbb{V}_{h+1, k}^{\pi^k} \right) + \frac{4H^2 S \log(6HS^2 AK/\delta)}{3n_k(h, s_h^k, a_h^k)} \\
&\leq \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\frac{3}{2} \bar{\mathbb{V}}_{h+1, k} - \frac{3}{2} \mathbb{V}_{h+1, k}^{\pi^k} \right) + \tilde{O}(H^3 S^2 A) && \text{(By Lemma 24)} \\
&= \underbrace{\frac{3}{2} \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\left\langle P_{h, s_h^k, a_h^k}, (\bar{\mathbb{V}}_{h+1, k})^2 \right\rangle - \left\langle P_{h, s_h^k, a_h^k}, (V_{h+1, k}^{\pi^k})^2 \right\rangle \right)}_{(a)} \\
&\quad + \underbrace{\frac{3}{2} \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\left\langle P_{h, s_h^k, a_h^k}, V_{h+1, k}^{\pi^k} \right\rangle^2 - \left\langle P_{h, s_h^k, a_h^k}, \bar{\mathbb{V}}_{h+1, k} \right\rangle^2 \right)}_{(b)} + \tilde{O}(H^3 S^2 A).
\end{aligned}$$

(By definition of variance)

We will bound (a) and (b) separately. For term (a), with probability at least $1 - \frac{\delta}{3}$, we have

$$\begin{aligned}
(a) &= \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left\langle P_{h, s_h^k, a_h^k}, (\bar{\mathbb{V}}_{h+1, k})^2 - (V_{h+1, k}^{\pi^k})^2 \right\rangle \\
&\leq \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left\langle P_{h, s_h^k, a_h^k}, \left| \bar{\mathbb{V}}_{h+1, k} - V_{h+1, k}^{\pi^k} \right| \left| \bar{\mathbb{V}}_{h+1, k} + V_{h+1, k}^{\pi^k} \right| \right\rangle \\
& && \text{(Since } a^2 - b^2 = (a+b)(a-b))
\end{aligned}$$

$$\begin{aligned}
&\leq 3H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left\langle P_{h,s_h^k,a_h^k}, \left| \bar{V}_{h+1,k} - V_{h+1,k}^{\pi^k} \right| \right\rangle \quad (\text{Since } \|\bar{V}_{h+1,k}\|_\infty \leq 2H \text{ under } \mathcal{G}_k) \\
&\leq 3H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \tilde{O}\left(H^2\sqrt{T}\right). \quad (\text{By Lemma 22})
\end{aligned}$$

For term (b), with probability at least $1 - \frac{\delta}{3}$, we have

$$\begin{aligned}
(\text{b}) &= \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left\langle P_{h,s_h^k,a_h^k}, V_{h+1,k}^{\pi^k} + \bar{V}_{h+1,k} \right\rangle \left\langle P_{h,s_h^k,a_h^k}, V_{h+1,k}^{\pi^k} - \bar{V}_{h+1,k} \right\rangle \\
&\leq 3H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left\langle P_{h,s_h^k,a_h^k}, \left| V_{h+1,k}^{\pi^k} - \bar{V}_{h+1,k} \right| \right\rangle \\
&\leq 3H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + \tilde{O}\left(H^2\sqrt{T}\right). \quad (\text{By Lemma 22})
\end{aligned}$$

Therefore, in summary, we have with probability at least $1 - \delta/2$,

$$\sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left(\hat{V}_{h+1,k} - \frac{3}{2} V_{h+1,k}^{\pi^k} \right) \leq \tilde{O}\left(H \sum_{k=1}^K \sum_{h=1}^{H-1} \mathbb{1}\{\mathcal{G}_k\} \left| \bar{\delta}_{h+1,k}^{\pi^k}(s_{h+1}^k) \right| + H^2\sqrt{T} + H^3 S^2 A \right).$$

□

H Proof of the Main Theorems

In this section, we state and prove our two main theorems.

Theorem 1. *If the Hoeffding-type noise is used, then for any MDP $M = (H, \mathcal{S}, \mathcal{A}, P, R, s_1)$, for any $\delta > 0$, with probability at least $1 - \delta$, Algorithm 1 satisfies*

$$\text{Reg}(M, K, \text{SSR}_{\text{Ho}}) \leq \tilde{O}\left(H^{1.5}\sqrt{SAT} + H^4 S^2 A\right).$$

In particular, when $T \geq \tilde{\Omega}(H^5 S^3 A)$, it holds that $\text{Reg}(M, K, \text{SSR}_{\text{Ho}}) \leq \tilde{O}\left(H^{1.5}\sqrt{SAT}\right)$.

Proof. By using the result of Lemma 20, under Hoeffding-type noise, with probability at least $1 - \delta$, we have

$$\begin{aligned}
&\text{Reg}(M, K, \text{SSR}_{\text{Ho}}) \\
&\leq \mathbb{1}\{\mathcal{G}_k\} 3C_1 e^{3C_1} \sum_{k=1}^K \sum_{h=1}^H \left(\sqrt{e_{\text{Ho}}^k(h, s_h^k, a_h^k)} + \gamma_{\text{Ho}}^k(h, s_h^k, a_h^k) \right) + \tilde{O}\left(H^4 S^2 A + H\sqrt{T}\right) \\
&\leq 6C_1 e^{3C_1} \sum_{k=1}^K \sum_{h=1}^{H-1} \left(H \sqrt{\frac{\log(2HSAk^2)}{n_k(h, s, a) + 1}} + \frac{H}{n_k(h, s, a) + 1} \right) + \tilde{O}\left(H^4 S^2 A + H\sqrt{T}\right) \\
&= \tilde{O}\left(H^{1.5}\sqrt{SAT} + H^4 S^2 A\right).
\end{aligned}$$

Here, the second inequality is from the definitions of $\sqrt{e_{\text{Ho}}^k(h, s_h^k, a_h^k)}$ and $\gamma_{\text{Ho}}^k(h, s_h^k, a_h^k)$, and the last step is from Lemma 23 and 24.

□

Theorem 2. *For Bernstein-type noise and $T \geq \tilde{\Omega}(H^5 S^2 A)$, then for any MDP $M = (H, \mathcal{S}, \mathcal{A}, P, R, s_1)$, for any $\delta > 0$, with probability at least $1 - \delta$, Algorithm 1 satisfies*

$$\text{Reg}(M, K, \text{SSR}_{\text{Be}}) \leq \tilde{O}\left(H\sqrt{SAT} + H^4 S^2 A\right).$$

In particular, if we further have $T \geq \tilde{\Omega}(H^6 S^3 A)$, it then holds that $\text{Reg}(M, K, \text{SSR}_{\text{Be}}) \leq \tilde{O}\left(H\sqrt{SAT}\right)$.

Proof. Similar to the proof of Theorem 1, under Bernstein-type noise, it holds with probability at least $1 - \frac{\delta}{2}$ that

$$\begin{aligned}
& \text{Reg}(M, K, \text{SSR}_{\text{Be}}) \\
& \leq \mathbb{1}\{\mathcal{G}_k\} 3C_1 e^{3C_1} \sum_{k=1}^K \sum_{h=i}^H \left(\sqrt{e_{\text{Be}}^k(h, s_h^k, a_h^k)} + \gamma_{\text{Be}}^k(h, s_h^k, a_h^k) \right) + \tilde{O}\left(H^4 S^2 A + H\sqrt{T}\right) \\
& \leq \tilde{O}\left(\sum_{k=1}^K \sum_{h=1}^{H-1} \left(\mathbb{1}\{\mathcal{G}_k\} U_{h,k} + \sqrt{\frac{\log(2HS^2 AK^2)}{n_k(h, s, a) + 1}} + \frac{H}{n_k(h, s, a) + 1} \right)\right) + \tilde{O}(H^4 S^2 A + H\sqrt{T}) \\
& = \tilde{O}(H\sqrt{SAT} + H^4 S^2 A),
\end{aligned}$$

where the last step is from Lemma 25. □

I Technical Lemmas

Lemma 30 (Bennet's Inequality). *Let Z_1, \dots, Z_n be i.i.d. random variables bounded in $[0, 1]$. Then, for any $\delta > 0$, we have*

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

Lemma 31 (from [32]). *Let Z_1, \dots, Z_n with $n \geq 2$ be i.i.d. random variables bounded in $[0, H]$. 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, for any $\delta > 0$, we have*

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

Lemma 32. *Let X be arbitrary random variable bounded in $[a, b]$ for some $a, b \in \mathbb{R}$. Then, we have $\text{Var}(X) \leq \frac{(b-a)^2}{4}$.*

Lemma 33. *For any $\delta > 0$, with probability at least $1 - \delta$, it holds for all k, h, s, a, s' that*

$$\left| \hat{P}_{h,s,a}^k(s') - P_{h,s,a}(s') \right| \leq \sqrt{\frac{4P_{h,s,a}(s')(1 - P_{h,s,a}(s')) \log(2HS^2 AK/\delta)}{n_k(h, s, a) + 1}} + \frac{3 \log(2HS^2 AK/\delta)}{n_k(h, s, a) + 1}.$$

Proof. Let $\delta' = \frac{\delta}{HS^2 AK}$ and fix (k, h, s, a, s') such that $n_k(h, s, a) \geq 1$. Then, we have

$$\begin{aligned}
& \mathbb{P}\left(\left|\hat{P}_{h,s,a}^k(s') - P_{h,s,a}(s')\right| \geq \sqrt{\frac{4P_{h,s,a}(s')(1 - P_{h,s,a}(s')) \log(2/\delta')}{n_k(h, s, a) + 1}} + \frac{3 \log(2/\delta')}{n_k(h, s, a) + 1}\right) \\
& \leq \mathbb{P}\left(\left|\tilde{P}_{h,s,a}^k(s') - P_{h,s,a}(s')\right| \geq \sqrt{\frac{4P_{h,s,a}(s')(1 - P_{h,s,a}(s')) \log(2/\delta')}{n_k(h, s, a) + 1}} + \frac{3 \log(2/\delta') - 1}{n_k(h, s, a) + 1}\right) \\
& \leq \mathbb{P}\left(\left|\tilde{\tilde{P}}_{h,s,a}^k(s') - P_{h,s,a}(s')\right| \geq \sqrt{\frac{4P_{h,s,a}(s')(1 - P_{h,s,a}(s')) \log(2/\delta')}{n_k(h, s, a) + 1}} + \frac{2 \log(2/\delta')}{n_k(h, s, a) + 1}\right) \\
& \leq \mathbb{P}\left(\left|\tilde{\tilde{\tilde{P}}}_{h,s,a}^k(s') - P_{h,s,a}(s')\right| \geq \sqrt{\frac{2P_{h,s,a}(s')(1 - P_{h,s,a}(s')) \log(2/\delta')}{n_k(h, s, a)}} + \frac{\log(2/\delta')}{n_k(h, s, a)}\right) \\
& \hspace{15em} (\text{Since } n + 1 \leq 2n \text{ for } n \geq 1) \\
& \leq \delta' = \frac{\delta}{HS^2 AK}. \hspace{10em} (\text{By Lemma 30, the Bennet's inequality})
\end{aligned}$$

Then, the proof is complete by taking a union bound over all possible (k, h, s, a, s') . □

J Numeric Simulations

In this section, we empirically compare RLSVI [41], UCBVI [7] and our algorithm SSR on the famous deep sea environment, which is a tabular environment frequently used to test an algorithm’s ability to do efficient exploration [35, 37, 47].

Deep sea, as shown in Figure 1, is a grid-like deterministic environment with $N \times N$ cell states, action space $\{0, 1\}$ and action mask $M_{ij} \sim \text{Bernoulli}(0.5)$, $(i, j) \in \mathcal{S}$, whose values are sampled when initializing the environment. At each cell (i, j) , Action M_{ij} represents going “right”, which leads the agent to the lower right cell, and $1 - M_{ij}$ represents going “left”, which leads the agent to the lower left cell. An episode of this environment will end after N steps. When going “left” or going “right” at the off-diagonal, the agent will receive 0 reward; when going “right” along the diagonal before reaching the lower right corner, the agent will receive negative reward $-\frac{0.01}{N}$. Finally, when reaching the lower right corner, depending on the environment initialization, the agent will either receive reward $+1$ or -1 . In our experiment, we set this to $+1$, which results in an obvious optimal policy “always going right” with total reward 0.99 per episode.

The experiment results are shown in Figure 2.³ From the plots, we can see that in both settings, SSR performs significantly better than RLSVI as predicted by our theory. Specifically, because of the instability incurred by the independent random seeds and large perturbation magnitude, RLSVI almost never reaches the lower right corner in both settings and thus incurs linear regret. On the other hand, SSR obtains a much lower sub-linear regret because it can explore consistently with the single random seed.

Meanwhile, in both settings, SSR performs comparably with the UCBVI, which is expected since both algorithms achieve the minimax lower bound and our analysis does not indicate that one is better than the other.

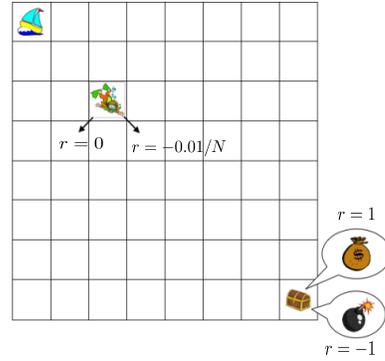


Figure 1: An example deep sea environment with $N = 8$ [37].

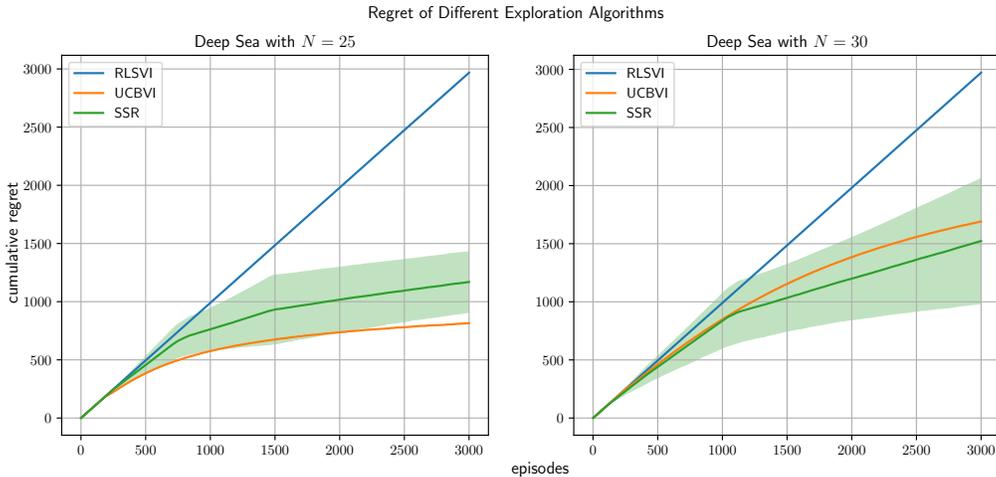


Figure 2: Empirical evaluation of RLSVI, UCBVI and SSR in deep sea environments with $N = 25$ and $N = 30$. The results are averaged over 10 repeated trials and the shaded area represents the standard deviation. For simplicity, we use Hoeffding-type bonus for both UCBVI and SSR.

³Bonuses for all three algorithms are scaled down from the theoretical values by a factor of 7×10^4 since without scaling, none of them can learn anything even in the deep sea with $N = 5$.

Finally, we also do an ablation study to show that the better performance of SSR over the RLSVI indeed comes from the single seed randomization instead of smaller noise magnitude. In particular, we run both algorithms in a deep sea environment with $N = 25$ and apply the same noise magnitude, whose results are shown in Figure 3. We can see that although using the same noise magnitude, SSR still significantly outperforms RLSVI.

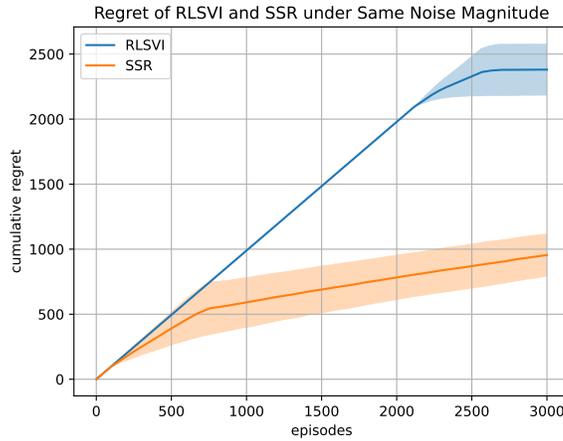


Figure 3: Empirical evaluation of RLSVI and SSR in deep sea environments with $N = 25$, where both algorithms use the same noise magnitude.