

Fine-Grained Gap-Dependent Bounds for Tabular MDPs via Adaptive Multi-Step Bootstrap

Haike Xu

Tsinghua University

XHK18@MAILS.TSINGHUA.EDU.CN

Tengyu Ma

Stanford University

TENGYUMA@STANFORD.EDU

Simon S. Du

University of Washington

SSDU@CS.WASHINGTON.EDU

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

This paper presents a new model-free algorithm for episodic finite-horizon Markov Decision Processes (MDP), **Adaptive Multi-step Bootstrap (AMB)**, which enjoys a stronger gap-dependent regret bound. The first innovation is to estimate the optimal Q -function by combining an optimistic bootstrap with an **adaptive** multi-step Monte Carlo rollout. The second innovation is to select the action with the largest confidence interval length among admissible actions that are not dominated by any other actions. We show when each state has a unique optimal action, **AMB** achieves a gap-dependent regret bound that only scales with the sum of the inverse of the sub-optimality gaps. In contrast, [Simchowitz and Jamieson \(2019\)](#) showed all upper-confidence-bound (UCB) algorithms suffer an additional $\Omega\left(\frac{S}{\Delta_{\min}}\right)$ regret due to over-exploration where Δ_{\min} is the minimum sub-optimality gap and S is the number of states. We further show that for general MDPs, **AMB** suffers an additional $\frac{|Z_{\text{mul}}|}{\Delta_{\min}}$ regret, where Z_{mul} is the set of state-action pairs (s, a) 's satisfying a is a *non-unique optimal action* for s . We complement our upper bound with a lower bound showing the dependency on $\frac{|Z_{\text{mul}}|}{\Delta_{\min}}$ is unavoidable for any consistent algorithm. This lower bound also implies a separation between reinforcement learning and contextual bandits.

Keywords: reinforcement learning, Markov Decision Process, gap-dependent bounds

1. Introduction

In reinforcement learning (RL), an agent iteratively interacts with an unknown environment with the goal of maximizing the reward. The state-of-the-art algorithms and analyses for tabular Markov Decision Process (MDP) achieve regret bounds that scale with \sqrt{K} where K is the number of episodes. These regret bounds hold for worst-case MDPs and are conservative—if a specific problem instance has benign structures, a much smaller regret is possible. One such structure is a nontrivial sub-optimality gap for the optimal Q -function—for every state s , the best action (or the set of best actions) is better than other actions by a margin. This structure exists in many real-world scenarios such as board games (tic-tac-toe, Chess) and Atari games (e.g., Freeway) ([Mnih et al., 2013](#)).

Researchers have extensively studied leveraging the suboptimality gap in the contextual bandits, which is a simplification of RL with horizon $H = 1$. It is well-known that the standard upper-confidence-bound (UCB) algorithm can achieve an optimal $O\left(\left(\sum_{\substack{(s,a) \in \mathcal{S} \times \mathcal{A}: \\ \Delta(s,a) > 0}} \frac{1}{\Delta(s,a)}\right) \log K\right)$ gap-dependent regret bound ([Bubeck and Cesa-Bianchi, 2012](#); [Lattimore and Szepesvári, 2020](#);

Slivkins, 2019). Here, \mathcal{S} is the state space with $|\mathcal{S}| = S$, \mathcal{A} is the action space with $|\mathcal{A}| = A$, and $\Delta(s, a)$ is the suboptimality gap of action a at the state s (that is, the advantage function at (s, a)). Notably, this regret only scales with $\log K$ instead of \sqrt{K} as in the formulation without the gap condition. One fruitful direction is to develop similar gap-dependent regret bounds for RL.

Previous gap-dependent RL regret bounds are mostly asymptotic (Jaksch et al., 2010; Tewari and Bartlett, 2008; Ok et al., 2018). Recently, Simchowit and Jamieson (2019); Lykouris et al. (2019); Yang et al. (2020) have developed non-asymptotic gap-dependent regret bounds for tabular MDPs. In particular, Simchowit and Jamieson (2019) showed that a UCB-based algorithm can achieve an

$$\tilde{O} \left(\left(\sum_{\substack{(s,a,h) \in \mathcal{S} \times \mathcal{A} \times [H]: \\ \Delta_h(s,a) > 0}} \frac{1}{\Delta_h(s,a)} + \frac{|Z_{\text{opt}}|}{\Delta_{\min}} + S^2 A \right) \text{poly}(H) \log K \right)^1 \quad (1.1)$$

regret bound where $\Delta_h(s, a)$ is the sub-optimality gap of state-action pair (s, a) at the h -th level (that is, h -th step), Δ_{\min} is the smallest gap among all state-action pairs at all levels, and Z_{opt} is the set of all optimal state-action pairs which satisfies $S \leq |Z_{\text{opt}}| \leq SA$. Comparing with gap-dependent regret bound of contextual bandits, there is an additional $|Z_{\text{opt}}|/\Delta_{\min}$ term.

Interestingly, Simchowit and Jamieson (2019) constructed an intriguing example in which $A = 2$, $H = 2$ and every state has a unique optimal action. They proved that all UCB algorithms will suffer an $\Omega(S/\Delta_{\min})$ regret on this example. See Section 1.2 for more expositions. One open question asked by Simchowit and Jamieson (2019) is

Can we develop a non-UCB algorithm whose regret does not depend on S/Δ_{\min} ?

The answer to this question has an important conceptual message. Recall in the contextual bandits setting, the regret does not depend on S/Δ_{\min} . Therefore, if the answer to the above question is negative, it demonstrates a formal *separation between contextual bandits and RL*.² On the other hand, if the answer is positive, then RL may not be more difficult than contextual bandits in terms of the gap-dependent regret.

1.1. Our Contributions

In this paper, we give both positive and negative results.

An Improved Algorithm. First, we design a new algorithm, Adaptive Multi-step Bootstrap (AMB), which enjoys the following gap-dependent regret guarantee.

Theorem 1.1 *For fixed K , AMB algorithm enjoys a gap-dependent regret upper bound with high probability*

$$\tilde{O} \left(\left(\sum_{\substack{(s,a,h) \in \mathcal{S} \times \mathcal{A} \times [H]: \\ \Delta_h(s,a) > 0}} \frac{1}{\Delta_h(s,a)} + \frac{|Z_{\text{mul}}|}{\Delta_{\min}} + SA \right) \text{poly}(H) \log K \right)$$

2. For the worst-case regret bound, it is still unclear whether there is a separation between contextual bandits and RL. See Jiang and Agarwal (2018); Wang et al. (2020); Zhang et al. (2020a).

where Z_{mul} is the set of state-action pairs (s, a) 's satisfying a is a non-unique optimal action for s .

The main difference between our bound and those in [Simchowitz and Jamieson \(2019\)](#) is about the second term in (1.1)—ours scales with $|Z_{\text{mul}}|/\Delta_{\min}$ whereas theirs scales with $|Z_{\text{opt}}|/\Delta_{\min}$, although our H dependency is worse than theirs. The following corollary illustrates the main improvement of our result in the special case where every state has a unique optimal action, that is, $|Z_{\text{mul}}| = 0$ but $|Z_{\text{opt}}| = S$.

Corollary 1.2 *For a fixed K , if every state of a MDP has a unique optimal action, then AMB enjoys a gap-dependent regret bound with high probability*

$$\tilde{O} \left(\left(\sum_{\substack{(s,a,h) \in \mathcal{S} \times \mathcal{A} \times [H]: \\ \Delta_h(s,a) > 0}} \frac{1}{\Delta_h(x,a)} + SA \right) \text{poly}(H) \log(K) \right).$$

In this case, the regret of the UCB-based algorithm in [Simchowitz and Jamieson \(2019\)](#) has an S/Δ_{\min} term because $|Z_{\text{opt}}| = S$, but ours does not. Therefore, when Δ_{\min} is small, the improvement of our bound is significant. More importantly, this improvement is not only from a better analysis, but also from fundamental algorithmic innovations. [Simchowitz and Jamieson \(2019\)](#) show that $\Omega(S/\Delta_{\min})$ regret is necessary for all UCB algorithms. AMB, instead, bypasses this technical barrier by considering both upper and lower confidence bounds of the Q -values (instead of only upper bounds as in UCB).

Another advantage is that our algorithm is model-free, which is more memory- and time-efficient than the model-based algorithms in [Simchowitz and Jamieson \(2019\)](#).³ Comparing with the previous $\tilde{O}(SA \cdot \text{poly}(H) \cdot (\log K)/\Delta_{\min})$ model-free gap-dependent regret bound in [Yang et al. \(2020\)](#), ours is more fine-grained as ours depends on the sum of the inverse of gaps and $|Z_{\text{mul}}|$, which in many instances are significantly tighter.

We note that our algorithm also enjoys a worst-case regret bound that scales with \sqrt{K} . See [Corollary B.10](#) in [Section B](#) for the formal statement and proofs.

A New Lower Bound. Now we turn to the negative result. Note for some MDPs, the quantity $|Z_{\text{mul}}|$ can be as large as SA . The next natural question is whether the dependency on $|Z_{\text{mul}}|/\Delta_{\min}$ is necessary. Our negative result shows that this is unavoidable.

Theorem 1.3 (Informal) *There is no algorithm ALG that can achieve a regret such that for all MDP M and K approaching infinity,*

$$\mathbb{E} [\text{Regret}_K(M, \text{ALG})] = O \left(\left(\sum_{\substack{(s,a,h) \in \mathcal{S} \times \mathcal{A} \times [H]: \\ \Delta_h(s,a) > 0}} \frac{1}{\Delta_h(x,a)} \right) \log(K) \right)$$

3. For tabular MDPs, a model-free algorithm's space complexity scales at most linearly in S , and a model-based algorithm's space complexity scales quadratically with S ([Strehl et al., 2006](#); [Jin et al., 2018](#)).

This lower bound shows that it is not possible to achieve a regret bound that solely depends on the sum of the inverse of the gaps. This lower bound also conveys a conceptual message that there is a separation between RL and contextual bandits because we know the UCB algorithm can achieve a regret bound that solely depends on $O\left(\sum \frac{1}{\Delta(s,a)}\right)$. As will be clear in Section 5, the transition operator in MDP allows us to construct harder problem instance that cannot be constructed in contextual bandits.

1.2. Main Challenges and Technique Overview

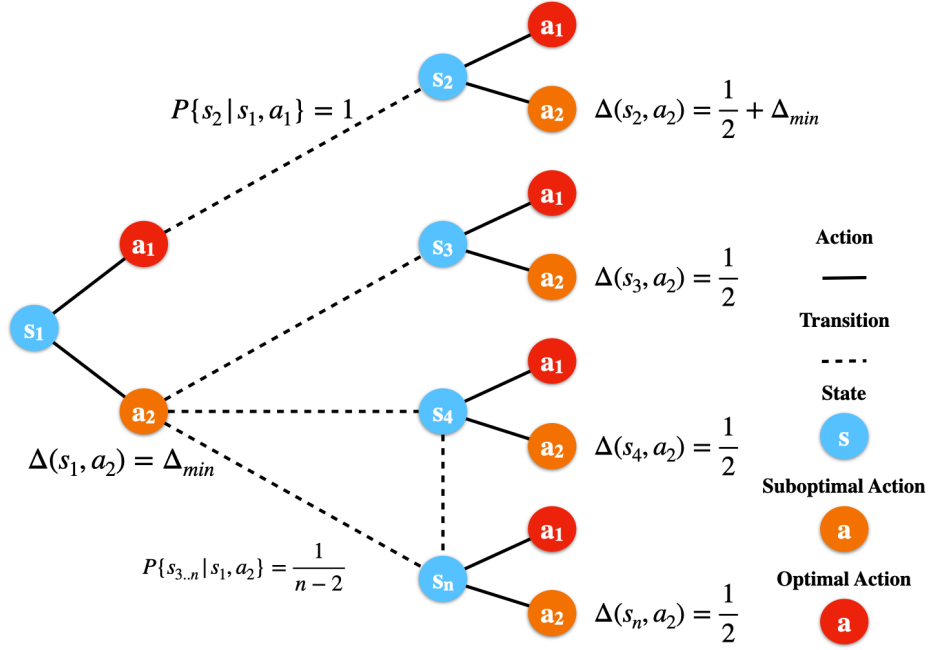


Figure 1: A simplified $H = 2$ hard instance that make all UCB algorithms incur an $\Omega(S/\Delta_{\min})$ regret. There are n states and two actions, a_1 and a_2 . a_1 is always the optimal action. The reward distribution satisfies $\mathcal{R}(s_2, a_1) = 1/2 + \Delta_{\min} + \eta$ and $r(s_i, a_1) = 1/2 + \eta$ for $i = 3, \dots, n$, where η is a zero-mean noise with variance 1. All other state-action pairs have reward 0. $\Delta(s_1, a_2) = \Delta_{\min}$, $\Delta(s_2, a_2) = 1/2 + \Delta_{\min}$, and $\Delta(s_i, a_2) = 1/2$ for $i = 3, \dots, n$. s_1 is the starting state. (s_1, a_1) will transit to s_2 deterministically, and (s_1, a_2) will transit to s_3, \dots, s_n , each with probability $\frac{1}{n-2}$.

1.2.1. THE HARD EXAMPLE IN SIMCHOWITZ AND JAMIESON (2019)

We first review the intuition about the hard instance in Simchowitz and Jamieson (2019) (cf. Figure 1) that makes all UCB algorithms suffer an $\Omega(S/\Delta_{\min})$ regret. In the hard instance, there are two actions a_1 and a_2 . At the starting state s_1 , a_1 is the optimal action with $Q^*(s_1, a_1) = 1/2 + \Delta_{\min}$, and a_2 is the suboptimal action with $Q^*(s_1, a_2) = 1/2$. In order to find the optimal action a_1 , the agent needs to estimate $Q^*(s_1, a_2)$ within Δ_{\min} error. To estimate $Q^*(s_1, a_2) =$

$\frac{1}{n-2} \sum_{i=3}^n V^*(s_i)$, all UCB algorithms rely on *optimistic bootstrap*, i.e., they maintain exploration bonuses, $b(s_3), \dots, b(s_n)$ for $V^*(s_3), \dots, V^*(s_n)$, and the over-estimation of $Q^*(s_1, a_2)$ will have a term $\approx \frac{1}{n-2} \sum_{i=3}^n b(s_i)$. To make this term smaller than Δ_{\min} , these algorithms essentially need $b(s_i) = O(\Delta_{\min})$ for all $i = 3, \dots, n$, which leads to an $\Omega(S/\Delta_{\min})$ regret.

1.2.2. GAP-DEPENDENT UPPER BOUND

Monte Carlo V.S. Optimistic Bootstrap To bypass the $\Omega(S/\Delta_{\min})$ lower bound in [Simchowitz and Jamieson \(2019\)](#)'s example, our main technique is to *collapse paths*. Notice we only need to pay $\tilde{O}(S)$ regret to find the optimal action a_1 , for s_3, \dots, s_n because these states have an $\Omega(1)$ gap. Now, to estimate $Q^*(s_1, a_2)$, instead of using the optimistic bootstrap in UCB algorithms, we use *Monte Carlo*. Since we know the optimal policy π^* for $\{s_3, \dots, s_n\}$, just by executing π^* , we can directly estimate $Q^*(s_1, a_2)$. This estimator only needs to pay $\tilde{O}(1/\Delta_{\min})$ regret in order to estimate $Q^*(s_1, a_2)$ within error $O(\Delta_{\min})$, in sharp contrast to UCB algorithms which need to pay $\Omega(S/\Delta_{\min})$ regret. The intuition is that Monte Carlo is estimating the mean of *one* random variable, whereas the optimistic bootstrap needs to estimate the means of S random variables.

This example shows the power of Monte Carlo for the scenario when subsequent states' optimal actions are known. This observation natural leads to a new estimator for Q^* , which *adaptively* combines optimistic bootstrap and Monte Carlo. We note that algorithmically, combining bootstrap and Monte Carlo is not new (see e.g., [Sutton and Barto \(1998\)](#)). However, to our knowledge, our algorithm is the first that adaptively combines optimistic bootstrap and Monte Carlo, and enjoys provable theoretical gains. See Section 4 for more details. Besides this new estimator, we also need some additional technical ingredients to obtain the improved gap-dependent bound.

Maximal Confidence Interval The estimator requires to identify a set of states whose the best action has been found. Identifying the best action inevitably involves the *action elimination* operation. Unfortunately, the existing UCB algorithms have no such operation. Our algorithm maintains an upper and a lower bound of each Q^* value. Importantly, at each episode, we select the action that 1) has not been eliminated and 2) has the largest uncertainty (measured by the difference between the upper and the lower confidence bound). Our action selection scheme is crucial because it has been shown in [Lykouris et al. \(2019\)](#) that the naive action selection scheme, randomly sampling one action from the remaining action set, suffers an exponential regret. Also note that selecting the action based on UCB may not work, because UCB never chooses actions which do not have the highest optimistic value, which makes their confidence bounds not tight. However, the action elimination operation requires accurate estimation for all un-eliminated actions.

1.2.3. GAP-DEPENDENT LOWER BOUND

The construction of our hard instance for $\Omega(SA/\Delta_{\min})$ lower bound relies on simple intuition: *A tabular MDP can simulate a multi-armed bandits with SA arms*, in which $\sum_{\Delta_h(s,a) > 0} \frac{1}{\Delta_h(s,a)}$ is much smaller than $\frac{SA}{\Delta_{\min}}$. Therefore, we can construct a multi-armed bandits example that has $\Omega(SA)$ arms with gap Δ_{\min} , while the MDP has only a few ($O(\log S)$) state-action pairs gap Δ_{\min} . This is in sharp contrast to contextual bandits with S states and A actions, which *cannot* simulate a multi-armed bandits with SA arms.

2. Related Work

There is a long line of results about worst-case regret bound of tabular RL. An incomplete list includes [Kearns and Singh \(2002\)](#); [Brafman and Tennenholtz \(2003\)](#); [Strehl et al. \(2006\)](#); [Jaksch et al. \(2010\)](#); [Dann and Brunskill \(2015\)](#); [Azar et al. \(2017\)](#); [Dann et al. \(2017\)](#); [Jin et al. \(2018\)](#); [Dann et al. \(2019\)](#); [Zhang et al. \(2020b\)](#); [Yang et al. \(2020\)](#); [Wang et al. \(2020\)](#); [Zhang et al. \(2020a\)](#). Algorithmically, we use the same step size for the model-free update in [Jin et al. \(2018\)](#). The state-of-the-art result by [Zhang et al. \(2020a\)](#) showed one can achieve $\tilde{O}\left(\sqrt{SAK} + S^2A\right)$ regret bound.⁴ Contextual bandits can be viewed as an episodic RL problem with $H = 1$, and its worst-case regret bound is $\Theta(\sqrt{SAK})$. Till today, it is still unclear whether there is a separation between RL and contextual bandits for the worst-case regret bound.

When there is a strictly positive sub-optimality gap, it is possible to achieve $\log K$ -type regret bounds. This type of results have been widely studied in the bandit literature. In RL, earlier work obtained asymptotic logarithmic regret bounds [Auer and Ortner \(2007\)](#); [Tewari and Bartlett \(2008\)](#). Recently, non-asymptotic logarithmic regret bounds were obtained ([Jaksch et al., 2010](#); [Simchowitz and Jamieson, 2019](#); [Yang et al., 2020](#)). Specially, [Jaksch et al. \(2010\)](#) developed a model-based algorithm, and their bound depends with the policy gap instead of the action gap studied in this paper. [Simchowitz and Jamieson \(2019\)](#) extended the model-based algorithm by [Zanette and Brunskill \(2019\)](#) and obtained logarithmic regret bounds. [Yang et al. \(2020\)](#) showed the model-free algorithm, the optimistic Q -learning algorithm by [Jin et al. \(2018\)](#) enjoyed a logarithmic regret. More recently, logarithmic regret bounds are obtained in linear function approximation settings ([He et al., 2020](#)). Lastly, [Ok et al. \(2018\)](#) derived problem-specific $\log K$ -type lower bounds for both structured and unstructured MDPs.

3. Preliminary

We denote a tabular episodic Markov Decision Process (MDP) by $M = (\mathcal{S}, \mathcal{A}, H, \mathcal{R}, \mathcal{P}, \mu)$ where \mathcal{S} is the state space with $|\mathcal{S}| = S$, \mathcal{A} is the action space with $|\mathcal{A}| = A$, H is the episode length (horizon), $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the reward distribution, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition probability distribution, and $\mu \in \Delta(\mathcal{S})$ is the initial state probability distribution. To streamline our analysis, we make a standard assumption for episodic settings that \mathcal{S} can be partitioned into disjoint sets \mathcal{S}_h , $h \in [H]$, such that $\mathcal{P}(\cdot | s, a)$ is supported on \mathcal{S}_{h+1} whenever $s \in \mathcal{S}_h$.⁵ A deterministic policy, π , assigns an action for each state, and can be seen as a function $\pi : \mathcal{S} \rightarrow \mathcal{A}$. Playing a policy π on a MDP M will induce a trajectory: $s_1, a_1, r_1, s_2, a_2, r_2, \dots, s_H, a_H, r_H$, where $s_1 \sim \mu$, $a_1 = \pi(s_1)$, $r_1 \sim \mathcal{R}(s_1, a_1)$, $s_2 \sim \mathcal{P}(s_1, a_1)$, etc.

For a given policy π , at each level $h = 1, \dots, H$, we define the the value function $V_h^\pi : \mathcal{S}_h \rightarrow \mathcal{R}$ and the Q -function $Q_h^\pi : \mathcal{S}_h \times \mathcal{A} \rightarrow \mathcal{R}$ as

$$V_h^\pi(x) = \mathbb{E}^\pi \left[\sum_{h'=h}^H r(s_{h'}, a_{h'}) | s_h = x \right], \quad Q_h^\pi(x, a) = \mathbb{E}^\pi \left[\sum_{h'=h}^H r(s_{h'}, a_{h'}) | s_h = x, a_h = a \right]$$

4. Their result holds for the setting where the reward is non-negative and the total reward is bounded by 1. This is a fair scaling when comparing with contextual bandits. See more expositions in [Jiang and Agarwal \(2018\)](#).

5. One can always augment the state space of the original episodic MDP to satisfy this assumption. The augmented state space is H times larger than the original one.

For simplicity, we define $V_0^\pi = \mathbb{E}[V_1^\pi(s_1)]$ to denote the value of a policy π . We use π^* to denote the optimal policy and $a^*(x)$ to denote the optimal action at state x (arbitrarily break tie if there are multiple optimal actions). This implies $\pi^*(x) = a^*(x)$. We write $V_h^*(x)$ in short for $V_h^{\pi^*}(x)$ and $Q_h^*(x, a)$ for $Q_h^{\pi^*}(x, a)$.

The agent interacts with the environment for K episodes. On each episode $k \in [1, K]$, the agent uses a policy π_k . We use cumulative simple regret $\text{Regret}_K = \sum_{k=1}^K V_0^* - V_0^{\pi_k}$ to measure the performance.

We focus on gap-dependent regret. For $(x, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$, the gap is defined as: $\Delta_h(x, a) = V_h^*(x) - Q_h^*(x, a)$. Note the optimal action at a state has the gap equals to zero. Following [Simchowitz and Jamieson \(2019\)](#), we let $Z_{\text{opt}}^h(x)$ denote the set of optimal actions for a state x on level $h \in [H]$, i.e., $Z_{\text{opt}}^h(x) = \{a \in \mathcal{A} : \Delta_h(x, a) = 0\}$. We use $Z_{\text{opt}} = \{(h, x, a) \mid \Delta_h(x, a) = 0\}$ to denote the set of optimal state-action pairs. We also define the local minimal gap: $\Delta_{h,\min}(x) = \min_{a \neq a^*(x)} \Delta_h(x, a)$ which should be 0 if $|Z_{\text{opt}}^h(x)| > 1$, and global minimal gap: $\Delta_{\min} = \min_{(x,a,h) \in \mathcal{S} \times \mathcal{A} \times [H] \wedge \Delta_h(x,a) > 0} \Delta_h(x, a)$. Our paper gives a fine-grained characterization of gap-dependent bounds, which rely on the following set

$$Z_{\text{mul}} = \left\{ (h, x, a) \mid \Delta_h(x, a) = 0 \wedge |Z_{\text{opt}}^h(x)| > 1 \right\}. \quad (3.1)$$

This is the set of state-action pairs whose states have multiple optimal actions. Note we always have $|Z_{\text{mul}}| \leq |Z_{\text{opt}}|$. Furthermore, under the following assumption, $|Z_{\text{mul}}| = 0$ whereas $|Z_{\text{opt}}| = S$. In the analysis below, we may drop the subscript h for some quantities because, by our assumption, any chosen state x implicitly contains the information of the level h which it belongs to.

Assumption 3.1 (Unique Optimal Action) *We say a MDP satisfies the unique optimal action assumption if for any state $x \in S$, it has a unique optimal action, i.e. $\forall x \in S, h \in [H], |Z_{\text{opt}}^h(x)| = 1$.*

4. Algorithm and Analysis Sketch

We will first describe the main algorithm and then in subsection 4.1 we will provide a proof sketch. Pseudocodes are listed in Algorithm 1.

Our algorithm maintains valid upper bounds and lower bounds of the Q -function at every episode k , denoted by $\overline{Q}_k(x, a)$ and $\underline{Q}_k(x, a)$, respectively. Given these bounds, for every state x , it maintains a set of candidate optimal actions, denoted by $A_k(x)$, by eliminating every action a whose Q -value upper bound is lower than another action's lower bound. Once only a single action survives for a state x , that is, $|A_k(x)| = 1$, we know that we have found the optimal action, and we call the state x a “decided” state. Otherwise we call x an “undecided” state. Let $G_k = \{x \mid |A_k(x)| = 1\}$ represent the subset of all decided states.

The key idea of the paper is to construct the upper and lower bounds of the Q -function by splitting the Q -function into two parts: the rewards from the decided states and those from the undecided states. This allows us to estimate the former part with more accurate sampling. Concretely, given the decided states G_k at episode k , we have

$$Q^*(x, a) = Q_k^{*b}(x, a) + Q_k^{*r}(x, a) \quad (4.1)$$

where $Q_k^{*b}(x, a)$ is the expected reward received by playing π^* after (x, a) until arriving a state that does not belong to G_k , and the $Q_k^{*r}(x, a)$ is the expected reward of the rest of the steps after seeing

Algorithm 1: Adaptive Multi-step Bootstrap (AMB)

Input: $\delta \in (0, 1/3)$ (failure probability), $H, \mathcal{A}, \mathcal{S}, K \geq 1$

1 $\forall x, a, \bar{Q}_0(x, a) \leftarrow H, \underline{Q}_0(x, a) \leftarrow 0, G_1 = \emptyset$, and $A_1(x) \leftarrow \mathcal{A}. \forall k, \bar{V}_k(\perp) = \underline{V}_k(\perp) = 0$

2 $\forall k$, let $\alpha_k = \frac{H+1}{H+k}$.

3 **for** $k = 1, 2, \dots$ **do**

4 **Collect data:**

5 Rollout from a random initial state $s_{k,1} \sim \mu$ using the policy π_k , defined as

$$\pi_k(x) \triangleq \begin{cases} \operatorname{argmax}_{a \in A_k(x)} \bar{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a) & \text{if } |A_k(x)| > 1 \\ \text{the element in } A_k(x) & \text{if } |A_k(x)| = 1 \end{cases},$$

and obtain an episode $s_{k,1}, \dots, s_{k,H}$.

6 **Update Q-function:**

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

8 **if** $s_{k,h} \notin G_k$ **then**

9 Let $n = n_k(s_{k,h}, a_{k,h})$ be the number of visits to $(s_{k,h}, a_{k,h})$.

10 Suppose $s_{k,h'}$ be the first state in the episode after $s_{k,h}$ that is not in G_k . (If such a state does not exist, let $h' = H+1$ and $s_{k,h'} = \perp$.)

11 Set the bonus $b_n = c\sqrt{H^3 \log(SAK/\delta)/n}$ for some universal constant c .

12 Let $\hat{Q}_k^{*b}(s_{k,h}, a_{k,h}) = \sum_{h \leq i < h'} r_{k,i}$

13 $\bar{Q}_k(s_{k,h}, a_{k,h}) = \min \left\{ H, (1 - \alpha_n) \bar{Q}_{k-1}(s_{k,h}, a_{k,h}) + \alpha_n \left(\hat{Q}_k^{*b}(s_{k,h}, a_{k,h}) + \bar{V}_{k-1}(s_{k,h'}) + b_n \right) \right\}$

14 $\underline{Q}_k(s_{k,h}, a_{k,h}) = \max \left\{ 0, (1 - \alpha_n) \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) + \alpha_n \left(\hat{Q}_k^{*b}(s_{k,h}, a_{k,h}) + \underline{V}_{k-1}(s_{k,h'}) - b_n \right) \right\}$

15 $\bar{V}_k(s_{k,h}) = \max_{a \in A_k(s_{k,h})} \bar{Q}_k(s_{k,h}, a)$

16 $\underline{V}_k(s_{k,h}) = \max_{a \in A_k(s_{k,h})} \underline{Q}_k(s_{k,h}, a)$

17 **end**

18 **end**

19 **for** $(x, a) \in \{\mathcal{S} \times A_k(x)\} \setminus \{(s_{k,h}, a_{k,h})\}_{h=1}^H$ **do**

20 $\bar{Q}_k(x, a) = \bar{Q}_{k-1}(x, a), \underline{Q}_k(x, a) = \underline{Q}_{k-1}(x, a), \bar{V}_k(x) = \bar{V}_{k-1}(x)$, and

$\underline{V}_k(x) = \underline{V}_{k-1}(x)$.

21 **end**

22 **Eliminate the sub-optimal actions:**

23 $\forall x \in \mathcal{S}$, set $A_{k+1}(x) = \{a \in A_k(x) : \bar{Q}_k(x, a) \geq \underline{V}_k(x)\}$

24 Let $G_{k+1} = \{x \in \mathcal{S} : |A_{k+1}(x)| = 1\}$.

25 **end**

any state that does not belong to G_k . Formally, suppose the state x is on the level h , and after taking the optimal actions, we arrive at the sequence of states x_{h+1}, \dots, x_H . Let h' be the smallest index (that is at least $h+1$) such that $x_{h'} \notin G_k$, then we can decompose $Q^*(x, a)$ into the sum of the

following two quantities:

$$Q_k^{*b}(x, a) \triangleq \mathbb{E} \left[\sum_{\ell=h}^{h'-1} r(x_\ell, a^*(x_\ell)) \right] \quad \text{and} \quad Q_k^{*r}(x, a) \triangleq \mathbb{E} [V^*(x_{h'})]. \quad (4.2)$$

For $Q_k^{*b}(x, a)$, the summation of observed empirical rewards can serve as an unbiased estimate, because we have taken the optimal action for states $a^*(x_\ell)$ for $h + 1 \leq \ell < h'$:

$$\hat{Q}_k^{*b}(s_{k,h}, a_{k,h}) = \sum_{\ell=h}^{h'-1} r_{k,\ell} \quad (4.3)$$

On the other hand, for $Q_k^{*r}(x, a)$, we can use the exiting V -values estimates on $x_{h'}$ to perform the bootstrapping, similarly to standard Bellman updates (Szepesvári, 2010). We will add a reward bonus term to counterbalance the stochasticity introduced in the estimation (4.2), so that finally we maintain valid upper and lower bounds in the sense that $\underline{Q}_k(x, a) \leq Q^*(x, a) \leq \overline{Q}_k(x, a)$ and $\underline{V}_k(x) \leq V^*(x) \leq \overline{V}_k(x)$. Concretely, the target value of the new Q -value is

$$\hat{Q}_k^{*b}(x, a) + \overline{V}_{k-1}(x_{h'}) + \text{bonus} \quad (4.4)$$

To make the update stable, following the standard framework proposed in Jin et al. (2018), we linearly interpolate the target value in (4.4) and the existing Q value with a learning rate α_k . We can derive a lower bound for the Q -values similarly and the resulting upper and lower bounds for the V -values. This part of the algorithm is described between Line 13 and Line 14 in Algorithm 1.

Our Bellman update is very reminiscent to the multi-step rollout Bellman updates that have been used successfully in practice (Sutton and Barto, 1998). However, in contrast to them, the Monte-Carlo rollout horizon in our algorithms adaptively depends on whether we have found the optimal actions in the following states.

As alluded before, given the upper and lower bounds, we can potentially eliminate more sub-optimal actions and build a small set of viable actions $A_{k+1}(x)$ (See Line 23 and Line 24 in Algorithm 1.)

The obtained new upper bound $\overline{Q}_k(x, a)$ and lower bound $\underline{Q}_k(x, a)$ will induce a new policy π_{k+1} . Instead of using UCB, we take actions that maximizes the length of the confidence interval

$$\pi_{k+1}(x) = \operatorname{argmax}_{a \in A_{k+1}(x)} \overline{Q}_k(x, a) - \underline{Q}_k(x, a) \quad (4.5)$$

We rollout with policy π_{k+1} in the next episode as in Line 5 of Alg. 1.

Comparison to previous algorithms. Compared with the existing model-free algorithms with regret guarantees, such as the one in Jin et al. (2018), there are two main differences between our algorithm and UCB-based algorithms:

- (1) In estimating \overline{Q} and \overline{V} , through a Q -function decomposition, we give tighter estimation for the first part $Q_k^{*b}(x, a)$, instead of directly summing up next level states' upper bounds.
- (2) Instead of choosing actions with the largest upper bounds as in UCB, we choose the actions with largest confidence interval lengths and eliminate an action when it can be excluded from the potential optimal actions.

Technical nuances. For notational convenience, we use \perp to denote a special termination state, and consider it to be on the level of $H + 1$. The value functions for this state is set to be zero in all cases, and we consider $\perp \notin G_k$ for all k . We also note that once $x \in G_k$ for some k , it will remain there for forever, and we will no longer update $Q_k(x, a)$ anymore—we will always take the unique optimal action as soon as $x \in G_k$ and the Q -values are no longer relevant anymore.

4.1. Proof Outline

In this section, we listed several key components of the proof for the case where each state has a unique optimal action (c.f. Corollary 1.2). Technical proofs are deferred to Appendix A. The proof for the general case (c.f. Theorem 1.1) is deferred to Appendix B. We first introduce some of the key notations and concepts in the analysis of the algorithm.

Key notations and concepts. We use $n_k(x, a)$ to denote the number of visits to the state-action pair (x, a) before and including episode k . For any t , let $k[t](x, a)$ be the episode number of the t -th visit to the state-action pair (x, a) . We will only use this notation when the algorithm indeed visits (x, a) for at least t times. When the pair (x, a) is clear in some context, we oftentimes omit (x, a) and just write $k[t]$ for simplicity. For any state x , let $x'_{k[t]}$ denotes the first undecided state (according to $G_{k[t]}$) in the episode $k[t]$ after the state x .

4.1.1. BACKGROUNDS ON LEARNING RATES AND CONCENTRATION PROPERTIES.

Our general framework follows the recent analysis of Q -learning algorithms (Jin et al., 2018), in terms of the choice of learning rates. We first define the quantity α_n^t that shows up in the analysis frequently when we expand the update rules for the Q -functions:

$$\alpha_n^0 \triangleq \prod_{j=1}^n (1 - \alpha_j) \quad \text{and} \quad \alpha_n^t \triangleq \alpha_t \prod_{j=t+1}^n (1 - \alpha_j) \quad (4.6)$$

Intuitively, α_n^t effectively measures how the update of $Q(x, a)$ at the n -visit to (x, a) depends on the past Q -value at the t -th visit of (x, a) , as characterized in the following statement:

$$\bar{Q}_{k[n]}(x, a) = \alpha_n^0 H + \sum_{t=1}^n \alpha_n^t \left(\hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) + b_t(x, a) \right) \quad (4.7)$$

The statement (4.7) can be obtained by a straightforward recursive expansion of the update rule in Line 13 of Alg. 1. (For a complete proof, see Section A.)

Similarly to the standard analysis of Q -learning, we will control $\bar{V}_{k[t]-1}(x'_{k[t]})$ on the RHS of equation (4.7) by recursion and $\hat{Q}_{k[t]}^{*b}(x, a)$ on the RHS by concentration inequality. The former part requires innovations but the latter part follows standard concentration inequality.

Lemma 4.1 (Concentration) *With probability at least $1 - \delta$ over the randomness of the environment, for all episodes $k \in [K]$, the following concentration inequalities hold:*

$$\forall x \in \mathcal{S} \setminus G_k, a \in A_k(x), \quad \left| \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) - Q_{k[t]}^{*b}(x, a) \right) \right| \leq \frac{1}{2} b_{n_k}(x, a) \quad (4.8)$$

$$\forall x \in \mathcal{S} \setminus G_k, a \in A_k(x), \left| \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(V^*(x'_{k[t]}) - Q_{k[t]}^{*r}(x, a) \right) \right| \leq \frac{1}{2} b_{n_k}(x, a) \quad (4.9)$$

For the sake of simplicity, in the following analysis, we use \mathcal{E}_j^B to denote the union of those inequalities' validity in (4.8) over episode $k = 1, \dots, j$, \mathcal{E}_j^R the union of those inequalities' validity in (4.9) over episode $k = 1, \dots, j$, and let $\mathcal{E}_j \triangleq \mathcal{E}_j^B \cap \mathcal{E}_j^R$. We also write \mathcal{E}^B for \mathcal{E}_K^B , \mathcal{E}^R for \mathcal{E}_K^R , and \mathcal{E} for $\mathcal{E}^B \cap \mathcal{E}^R$.

4.1.2. KEY STEPS IN THE PROOFS

Now we list several key lemmas in the proof. The following lemma shows our confidence intervals about the Q -function and the V -function are valid.

Lemma 4.2 (Valid Confidence Interval) *For all $(x, a) \in \mathcal{S} \times \mathcal{A}$ and at any episode k , when event \mathcal{E}_k happens, the upper and lower confidence bounds in Algorithm 1 are valid:*

$$\bar{V}_k(x) \geq V^*(x) \geq \underline{V}_k(x) \quad \text{and} \quad \bar{Q}_k(x, a) \geq Q^*(x, a) \geq \underline{Q}_k(x, a) \quad (4.10)$$

With this lemma, we can easily show that we never eliminate the optimal action.

Proposition 4.3 (Action Elimination) *When the event \mathcal{E}_{k-1} happens, for all $x \in \mathcal{S}$, all the optimal actions for x are in the set $A_k(x)$. As a direct consequence, $\forall x \in G_k$, the set $A_k(x)$ contains the unique optimal action for x .*

Now we turn to bounding the regret. The following lemma shows the length of the confidence interval is an upper bound of the regret, conditioned on the event \mathcal{E}_k .

Lemma 4.4 (Bounding Regret By the Confidence Interval Length) *For any episode k , conditioning on the event \mathcal{E}_{k-1} , the regret can be bounded by the confidence interval length of those undecided states that are not in G_k :*

$$(V_0^* - V_0^{\pi_k}) \Big|_{\mathcal{E}_{k-1}, \mathcal{F}_{k-1}} \leq 2\mathbb{E} \left[\sum_{h=1}^H \left(\bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \cdot \mathbb{I}[s_{k,h} \notin G_k] \Big|_{\mathcal{E}_{k-1}, \mathcal{F}_{k-1}} \right]$$

We note this lemma is different from the decomposition for the UCB-based algorithms, which admit the property that $\bar{Q}_k(s_{k,h}, a_{k,h}) \geq \bar{Q}_k(s_{k,h}, a^*(s_{k,h})) \geq V^*(s_{k,h})$ and then use the estimation error, $\bar{Q}_k(s_{k,h}, a_{k,h}) - Q^*(s_{k,h}, a_{k,h})$, as an upper bound for regret. However, since we do not always choose the action that maximizes the estimated Q -value, we need a new upper bound on the regret. Fortunately, for our analysis, with the action elimination mechanism, the regret can be simply bounded by the maximal confidence interval.

With Lemma 4.4 at hand, it suffices to bound the above the confidence intervals. Our analysis relies on the clip function

$$\text{clip}[x|y] \triangleq x \cdot \mathbb{I}[x \geq y].$$

We obtain the following recursion.

Proposition 4.5 (Confidence Interval Length Recursion) *Suppose \mathcal{E}_{k-1} happens. Suppose $(x, a) = (s_{k,h}, a_{k,h})$ is a state-action pair visited in the k -th episode where $x \notin G_k$ is an undecided state. Let ΔQ_{past} be a shorthand for*

$$\Delta Q_{\text{past}} \triangleq \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \left(\bar{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) - \underline{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) \right) \cdot \mathbb{I} \left[x'_{k[t]} \notin G_{k[t]} \right] \quad (4.11)$$

We have the following recursion bound for the confidence interval length of an undecided state:

$$\left(\bar{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a) \right) \cdot \mathbb{I} [x \notin G_k] \quad (4.12)$$

$$\leq \alpha_{n_{k-1}}^0 H + \left(1 + \frac{1}{H} \right) \Delta Q_{\text{past}} + \begin{cases} \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \frac{\Delta(x, a)}{4H} \right] & \text{if } a \neq a^*(x) \\ \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \frac{\Delta_{\min}(x)}{4H} \right] & \text{if } a = a^*(x) \end{cases} \quad (4.13)$$

The clipping operation was proposed in [Simchowitz and Jamieson \(2019\)](#) to derive gap-dependent logarithmic regret bounds. We use one particular property about this operation (c.f. Claim [A.13](#)). The main difference between their use and ours is that for the second case, $a = a^*(x)$, in Proposition 4.5, [Simchowitz and Jamieson \(2019\)](#) introduce half-clipping trick and results in $O(\frac{1}{\Delta_{\min}})$ regret for each state, while we use action elimination mechanism to give $O(\frac{1}{\Delta_{\min}(x)})$ regret. This is crucial to avoid the S/Δ_{\min} dependency.

To proceed, we can solve the recursion by induction, and obtain the following lemma.

Lemma 4.6 (Solving Recursion) *We define clipped reward function as*

$$\check{b}_{n_k}(s_{k,h}, a_{k,h}) \triangleq \text{clip} \left[4b_{n_k}(s_{k,h}, a_{k,h}) \mid \max \left(\frac{\Delta(s_{k,h}, a_{k,h})}{4H}, \frac{\Delta_{\min}(s_{k,h})}{4H} \right) \right] \quad (4.14)$$

When event \mathcal{E} happens, we can upper bound the regret by a linear combination of clipped reward:

$$\sum_{k=1}^K \sum_{h=1}^H \left(\bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \cdot \mathbb{I} [s_{k,h} \notin G_k] \quad (4.15)$$

$$\leq e^2 H^2 S A + e^2 H \sum_{k=1}^K \sum_{h=1}^H \check{b}_{n_{k-1}}(s_{k,h}, a_{k,h}) \cdot \mathbb{I} [s_{k,h} \notin G_k] \quad (4.16)$$

To finish the proof of Corollary 1.2, we use the upper bound of the failure probability and property of the clipping trick (Claim [A.13](#)). See details in Section [A](#).

5. Lower Bound

Here we present our formal lower bound.




Figure 2: A family of hard instances that make all consistent algorithms incur an $\Omega(SA \log K / \Delta_{\min})$ regret. M is the base MDP and M_{ij} is another MDP we wish to distinguish from M .

Theorem 5.1 *Given integers $S, A, H \geq \log_2(S)$, $0 < \Delta_{\min} < \min(\frac{1}{8}, 1/H)$, and $S \leq |Z_{\text{mul}}| \leq \frac{SA}{2}$, there exists an MDP which has S states, A actions, H levels, and satisfies $\sum_{\substack{(s,a,h) \in S \times A \times [H]: \\ \Delta_h(s,a) > 0}} \frac{1}{\Delta_h(s,a)} = c_1 \log(S) / \Delta_{\min}$ for some absolute constants $c_1 > 0$. On this MDP, there exists an absolute constant $c_2 > 0$ such that as $K \rightarrow \infty$, any consistent algorithm suffers a regret at least $c_2 |Z_{\text{mul}}| \log K / \Delta_{\min}$.*

In this section we describe our main ideas for proving Theorem 5.1. A graphical illustration of our hard-instance construction is shown in Figure 2. At a high level, we use an MDP to simulate a multi-armed bandit problem with SA arms, and then we choose the canonical hard instance in multi-armed bandit to prove the lower bound.

We construct an MDP M whose states form a complete binary tree, where there are $|S| = 2n - 1$ states and n leaves. We label their last horizon's states with $\{x_1, \dots, x_n\}$. All states previous to the last horizon have two actions $\{a_1, a_2\}$, while states on the last level have $|\mathcal{A}| = A$ actions $\{a_1, \dots, a_A\}$. All transitions are deterministic and follow the binary tree structure. The actions taken in states on the last level drawn on the rightest column is the only place where non-zero rewards are given to the agent. Following the standard proof of the lower bound for multi-armed bandits, we assume all rewards follow a Bernoulli distribution, whose mean is labeled on the top of the action. Only one state x_1 has one $\frac{1}{2} + \gamma$ reward action. All x_1 's other actions and all other states' actions have reward

mean $\frac{1}{2}$. Equivalently, this is a multi-armed bandit problem with SA arms and the only non-zero reward is on (x_1, a_1) .

Now we construct a set of MDPs $\{M_{ij}\}_{i \in \{2, \dots, |S|\}, j \in \{1, \dots, A\}}$. For MDP $M_{i,j}$, all transitions and rewards are the same with M , except that one action a_j of state x_i has reward $\frac{1}{2} + 2\gamma$. We note that these MDPs only have four possible rewards $\{0, \frac{1}{2}, \frac{1}{2} + \gamma, \frac{1}{2} + 2\gamma\}$.

To prove the lower bound, we follow the standard technique in multi-armed bandits (e.g., [Lattimore and Szepesvári \(2020\)](#)) to show $\Omega(SA \log K / \Delta_{\min})$ regret. See Section C for details.

6. Conclusion

In this paper, we design a new algorithm enjoying an improved gap-dependent regret bound for episodic finite-horizon MDPs. This new regret bound is significant tighter than previous bounds when all states have a unique optimal action. The two innovations involved are the use of adaptive multi-step bootstrap in Q -value estimation and choosing the action that has the largest confidence interval. We also prove a new regret lower bound showing that achieving the tighter regret bound for general MDPs is impossible.

Acknowledgment

TM acknowledges support of Google Faculty Award, NSF IIS 2045685, Lam Research, and JD.com

References

- Peter Auer and Ronald Ortner. Logarithmic online regret bounds for undiscounted reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 49–56, 2007.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, pages 263–272, 2017.
- Ronen I. Brafman and Moshe Tennenholtz. R-max - a general polynomial time algorithm for near-optimal reinforcement learning. *J. Mach. Learn. Res.*, 3(Oct):213–231, March 2003. ISSN 1532-4435.
- Sebastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1), 2012.
- Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2818–2826, 2015.
- Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying PAC and regret: Uniform PAC bounds for episodic reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 5717–5727, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Christoph Dann, Lihong Li, Wei Wei, and Emma Brunskill. Policy certificates: Towards accountable reinforcement learning. In *Proceedings of the 36th International Conference on Machine Learning*, pages 1507–1516, 2019.

- Jiafan He, Dongruo Zhou, and Quanquan Gu. Logarithmic regret for reinforcement learning with linear function approximation. *arXiv preprint arXiv:2011.11566*, 2020.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(Apr):1563–1600, 2010.
- Nan Jiang and Alekh Agarwal. Open problem: The dependence of sample complexity lower bounds on planning horizon. In *Conference On Learning Theory*, pages 3395–3398, 2018.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873, 2018.
- Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2-3):209–232, 2002.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Thodoris Lykouris, Max Simchowitz, Aleksandrs Slivkins, and Wen Sun. Corruption robust exploration in episodic reinforcement learning. *arXiv preprint arXiv:1911.08689*, 2019.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Jungseul Ok, Alexandre Proutiere, and Damianos Tranos. Exploration in structured reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 8874–8882, 2018.
- 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.
- Aleksandrs Slivkins. Introduction to multi-armed bandits. *arXiv preprint arXiv:1904.07272*, 2019.
- 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.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 1998.
- Csaba Szepesvári. Algorithms for reinforcement learning. *Synthesis lectures on artificial intelligence and machine learning*, 4(1):1–103, 2010.
- Ambuj Tewari and Peter L Bartlett. Optimistic linear programming gives logarithmic regret for irreducible MDPs. In *Advances in Neural Information Processing Systems*, pages 1505–1512, 2008.
- Ruosong Wang, Simon S Du, Lin F Yang, and Sham M Kakade. Is long horizon reinforcement learning more difficult than short horizon reinforcement learning? In *Advances in Neural Information Processing Systems*, 2020.
- Kunhe Yang, Lin F Yang, and Simon S Du. Q-learning with Logarithmic Regret. *arXiv preprint arXiv:2006.09118*, 2020.

Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pages 7304–7312, 2019.

Zihan Zhang, 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*, 2020a.

Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. In *Advances in Neural Information Processing Systems*, 2020b.

Appendix A. Missing Proofs in Section 4.1

In this section we prove the following result under Assumption 3.1. Note Theorem A.1 implies Corollary 1.2.

Theorem A.1 (Main Result Under Assumption 3.1) *Under Assumption 3.1, for fixed K , with probability at least $1 - \delta$, we have the following regret upper bound*

$$\text{Regret}_K \leq O \left(H^2 SA + \sum_{x \in \mathcal{S}} \left(\sum_{a \neq a^*(x)} \frac{H^5}{\Delta(x, a)} \right) \log \left(\frac{SAK}{\delta} \right) \right) \quad (\text{A.1})$$

Here we also briefly summarize why we have H^5 the bound. First, since we use α_i as the learning rate in Line 13 of Algorithm 1, we need to set bonus as large as $\Omega(\sqrt{\frac{H^3}{n}})$. Second, we use the clipping trick to clip the bonus at $\Omega(\frac{\Delta}{H})$ (Lemma 4.5), so the summation over bonus until it is clipped will yield $O(\frac{H^4}{\Delta})$. Finally, our use of confidence interval to decompose regret (Lemma 4.4) and solving the recursion (Lemma 4.6) incurs another factor of H . We leave it as a future direction is to improve the dependency on H .

Proposition A.2 (Lemma 4.1 of Jin et al. (2018)) *Recall that $\alpha_t = \frac{H+1}{H+t}$. Define $\alpha_t^0 \triangleq \prod_{j=1}^t (1 - \alpha_j)$, and $\alpha_t^i \triangleq \alpha_i \prod_{j=i+1}^t (1 - \alpha_j)$. Then, we have the following properties:*

- (1) $\begin{cases} \sum_{i=1}^t \alpha_t^i = 1 \text{ and } \alpha_t^0 = 0 & \forall t \geq 1 \\ \sum_{i=1}^t \alpha_t^i = 0 \text{ and } \alpha_t^0 = 1 & \text{for } t = 0 \end{cases}$
- (2) $\forall t \geq 1, \frac{1}{\sqrt{t}} \leq \sum_{i=1}^t \frac{\alpha_t^i}{\sqrt{i}} \leq \frac{2}{\sqrt{t}}$
- (3) $\forall i \geq 1, \sum_{t=i}^{\infty} \alpha_t^i = 1 + \frac{1}{H}$
- (4) $\forall i \geq 1, \sum_{i=1}^t (\alpha_t^i)^2 \leq \frac{2H}{t}$

A.1. Proof of lemma 4.2

Here we prove an extended version of Lemma 4.2, where we additionally prove an upper bound of $\bar{V}_k(x) - \underline{Q}_k(x, a)$. This inequality is useful for bounding regret in later analysis.

Lemma A.3 (Extended version of Lemma 4.2) *For all $(x, a) \in \mathcal{S} \setminus G_k \times A_k(x)$ and at any episode k , when event \mathcal{E}_k happens, the upper and lower confidence bounds are valid:*

$$\bar{V}_k(x) \geq V^*(x) \geq \underline{V}_k(x) \quad (\text{A.2})$$

$$\bar{Q}_k(x, a) \geq Q^*(x, a) \geq \underline{Q}_k(x, a) \quad (\text{A.3})$$

Moreover, for any k, x and $a \in A_{k+1}(x)$, we have that

$$\bar{V}_k(x) - \underline{Q}_k(x, a) \leq 2 \max_{a' \in A_{k+1}(x)} \left(\bar{Q}_k(x, a') - \underline{Q}_k(x, a') \right) \quad (\text{A.4})$$

Proof [Proof of lemma A.3]

We first use induction to prove this inequality (A.2). The induction proceeds in two dimensions, episode k and horizon h . We first check that the initialization in Algorithm 1 is valid. Then we assume that the induction is valid from episode 0 to $k - 1$. In the following, we prove that the argument is correct for a fixed episode k .

We consider the other induction dimension, horizon h , in a reversed order. For the base case $h = H + 1$, where there is only one state \perp . By the initialization in Algorithm 1, we have $\bar{V}_k(\perp) = \underline{V}_k(\perp) = 0$. Therefore the induction argument is valid for the base case. Now, we assume that for state $x \in \{\mathcal{S}_i\}_{i>h}$, the induction argument is valid. Then, we want to prove that the argument is valid for $x \in \mathcal{S}_h$.

For the sake of simplicity, we denote n_k to be an abbreviation for $n_k(x_k, a_k)$, the number of times we have visited the state action pair (x, a) , $k[t]$ to be the episode of the t -th arrival to the specified state x , $x'_{k[t]}$ to be the first not in $G_{k[t]}$ state arrived starting at x on episode $k[t]$.

According to the update rule in Algorithm 1, line (13), we have the following expression for $\bar{Q}_k(x, a)$:

$$\bar{Q}_k(x, a) = \min \left(H, \alpha_{n_k}^0 H + \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) + b_t(x, a) \right) \right) \quad (\text{A.5})$$

Because by definition $Q^*(x, a) \leq H$, the argument is true when the minimum of the two value equals H . Now we assume that the minimum takes the second term.

Then, we have the following:

$$\bar{Q}_k(x, a) - Q^*(x, a) \quad (\text{A.6})$$

$$= \alpha_{n_k}^0 H + \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) + b_t(x, a) \right) - Q^*(x, a) \quad (\text{A.7})$$

$$= \alpha_{n_k}^0 H + \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) + b_t(x, a) - Q^*(x, a) \right) - \mathbb{I}[n_k = 0] Q^*(x, a)$$

(By property (1) in Proposition A.2)

$$= \alpha_{n_k}^0 H + \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) - Q^*(x, a) \right) + \sum_{t=1}^{n_k} \alpha_{n_k}^t b_t(x, a) - \mathbb{I}[n_k = 0] Q^*(x, a) \quad (\text{A.8})$$

$$\geq \alpha_{n_k}^0 H + \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) - Q^*(x, a) \right) + \mathbb{I}[n_k > 0] b_{n_k}(x, a) - \mathbb{I}[n_k = 0] Q^*(x, a)$$

(By the definition of $b_k(x, a)$ and property (2) in Proposition A.2)

When $n_k = 0$, the first term above equals H , the last term equals $Q^*(x, a)$, and the second and third term becomes zero, so the RHS is greater than 0. Now, we consider the case when $n_k > 0$, then only the second and third term are non-zero. Next, we want to show that the second term is larger than $-b_k(x, a)$.

$$\begin{aligned}
 & \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) - Q^*(x, a) \right) \tag{A.9} \\
 = & \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) - V^*(x'_{k[t]}) + V^*(x'_{k[t]}) - Q_{k[t]}^{*b}(x, a) - Q_{k[t]}^{*r}(x, a) \right) \\
 & \text{(By the decomposition of } Q^*(x, a) \text{ defined in Line (4.1))} \\
 = & \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) - Q_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) - V^*(x'_{k[t]}) + V^*(x'_{k[t]}) - Q_{k[t]}^{*r}(x, a) \right) \tag{A.10} \\
 = & \underbrace{\sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) - Q_{k[t]}^{*b}(x, a) \right)}_{\text{Bounded by event } \mathcal{E}_k^B} + \underbrace{\sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\bar{V}_{k[t]-1}(x'_{k[t]}) - V^*(x'_{k[t]}) \right)}_{\text{Bounded by the induction argument}} \\
 & + \underbrace{\sum_{t=1}^{n_k} \alpha_{n_k}^t \left(V^*(x'_{k[t]}) - Q_{k[t]}^{*r}(x, a) \right)}_{\text{Bounded by event } \mathcal{E}_k^R} \tag{A.11} \\
 \geq & -\frac{1}{2}b_k(x, a) - 0 - \frac{1}{2}b_k(x, a) \tag{A.12} \\
 \geq & -b_k(x, a) \tag{A.13}
 \end{aligned}$$

A similar argument shows $\underline{Q}_k(x, a) \leq Q^*(x, a)$ as well.

Next we prove that for $x \in S_h$, $\bar{V}_k(x) \geq V^*(x)$ and $\underline{V}_k(x) \leq V^*(x)$. By the updating rule in (15), and (16) of Algorithm 1, we have

$$\bar{V}_k(x) = \max_{a \in A_k(x)} \bar{Q}_k(x, a) \geq \max_{a \in A_k(x)} Q^*(x, a) = V^*(x, a) \tag{A.14}$$

$$\underline{V}_k(x) = \max_{a \in A_k(x)} \underline{Q}_k(x, a) \leq \max_{a \in A_k(x)} Q^*(x, a) = V^*(x, a) \tag{A.15}$$

Finally we will show equation (A.4). Recall that by Line 16 of Alg. 1, we have that $\bar{V}_k(x) = \max_{a \in A_k(x)} \bar{Q}_k(x, a)$. Assume that the max is attained at a^* . Then, for all $a \in A_{k+1}(x)$:

$$\begin{aligned}
 \bar{V}_k(x) - \underline{Q}_k(x, a) &= \bar{Q}_k(x, a^*) - \underline{Q}_k(x, a) \\
 &= \left(\bar{Q}_k(x, a^*) - \underline{Q}_k(x, a^*) \right) + \left(\underline{Q}_k(x, a^*) - \bar{Q}_k(x, a) \right) + \left(\bar{Q}_k(x, a) - \underline{Q}_k(x, a) \right) \\
 &\leq 2 \max_{a' \in A_{k+1}(x)} \left(\bar{Q}_k(x, a') - \underline{Q}_k(x, a') \right) \tag{A.16}
 \end{aligned}$$

where in the last inequality we use the fact that $\underline{Q}_k(x, a^*) \leq \underline{V}_k(x) \leq \bar{Q}_k(x, a)$ for all $a \in A_{k+1}(x)$. \blacksquare

A.2. Proof of proposition 4.3

Restatement of Proposition 4.3 (Action Elimination)

When the event \mathcal{E}_{k-1} happens, for all $x \in \mathcal{S}$, all the optimal actions for x are in the set $A_k(x)$. As a direct consequence, $\forall x \in G_k$, the set $A_k(x)$ contains the unique optimal action for x .

Proof [Proof of proposition 4.3]

Suppose an action a was excluded from the set $A_k(x)$ at episode k by Alg. 1. It implies that $\bar{Q}_k(x, a) < \underline{V}_k(x)$. By Lemma 4.2, we have $\bar{Q}_k(x, a) < \underline{V}_k(x) \leq V^*(x)$, which indicates a is not an optimal action. Therefore, all actions that are eliminated are suboptimal and the set $A_k(x)$ always contain all the optimal actions.

Finally because any state x in G_k satisfies $|A_k(x)| = 1$, the set $A_k(x)$ must contain the unique optimal action. ■

A.3. Proof of lemma 4.4

Restatement of Lemma 4.4 (Bounding Regret By the Confidence Interval Length)

For any episode k , conditioning on the event \mathcal{E}_{k-1} , the regret can be bounded by the confidence interval length of those undecided states that are not in G_k :

$$(V_0^* - V_0^{\pi_k}) \Big|_{\mathcal{E}_{k-1}, \mathcal{F}_{k-1}} \leq 2\mathbb{E} \left[\sum_{h=1}^H \left(\bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \cdot \mathbb{I}[s_{k,h} \notin G_k] \Big|_{\mathcal{E}_{k-1}, \mathcal{F}_{k-1}} \right]$$

Proof [Proof of lemma 4.4] Conditioning on the event \mathcal{E}_{k-1} and the filtration \mathcal{F}_{k-1} of all the random variables generated until the beginning of epoch k , we can bound the V -value of policy π_k .

$$\begin{aligned} (V_0^* - V_0^{\pi_k}) \Big|_{\mathcal{E}_{k-1}, \mathcal{F}_{k-1}} &= \mathbb{E} \left[\sum_{h=1}^H V^*(s_{k,h}) - Q^*(s_{k,h}, a_{k,h}) \Big|_{\mathcal{E}_{k-1}, \mathcal{F}_{k-1}} \right] \\ &= \mathbb{E} \left[\sum_{h=1}^H (V^*(s_{k,h}) - Q^*(s_{k,h}, a_{k,h})) \cdot \mathbb{I}[s_{k,h} \notin G_k] \Big|_{\mathcal{E}_{k-1}, \mathcal{F}_{k-1}} \right] \\ &\quad \text{(By Lemma 4.3, for all } s_{k,h} \in G_k, \pi_k(s_{k,h}) = a_{s,h} \in Z_{\text{opt}}(x) \text{ and } V^*(s_{k,h}) = Q^*(s_{k,h}, a_{k,h})) \\ &\leq \mathbb{E} \left[\sum_{h=1}^H \left(\bar{V}_{k-1}(s_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \cdot \mathbb{I}[s_{k,h} \notin G_k] \Big|_{\mathcal{E}_{k-1}, \mathcal{F}_{k-1}} \right] \end{aligned} \tag{A.17}$$

where the last inequality follows from by Lemma 4.2

Invoking equation (A.4) of Lemma 4.2 with episode k , $x = s_{k,h}$, and $a = a_{k,h}$, noting that in Line 5 of Alg. 1, we chose $a_{k,h} = \operatorname{argmax}_{a' \in A_k(x)} \left(\bar{Q}_{k-1}(s_{k,h}, a') - \underline{Q}_{k-1}(s_{k,h}, a') \right)$, we have

$$\bar{V}_{k-1}(s_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \leq 2 \left(\bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \tag{A.18}$$

Plugging in (A.18) into equation (A.17) completes the proof. ■

A.4. Proofs of Proposition 4.5

Lemma A.6 *Suppose (x, a) is visited at the episode k . Conditioning on the event \mathcal{E}_k , letting $n_k = n_k(x, a)$ and $k[t] = k[t](x, a)$, we have,*

$$\bar{Q}_k(x, a) - \underline{Q}_k(x, a) \leq \alpha_{n_k}^0 H + 4\mathbb{I}[n_k > 0] \cdot b_k(x, a) + \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\bar{V}_{k[t]-1}(x'_{k[t]}) - \underline{V}_{k[t]-1}(x'_{k[t]}) \right) \quad (\text{A.19})$$

Lemma A.7 *Suppose sequences u_n and w_n satisfy*

$$u_n = (1 - \alpha_n)u_{n-1} + \alpha_n w_n \quad (\text{A.20})$$

for all $n \geq 1$ and $u_0 = H$. Then,

$$u_n = \alpha_n^0 H + \sum_{1 \leq t \leq n} \alpha_n^t w_t \quad (\text{A.21})$$

Proof [Proof of lemma A.7] We recursively expand u_n according to line (A.20) and get a linear combination of w_t . We can use induction to prove that the coefficient of w_t in the expansion of u_n is α_n^t . For the base case $t = n$, we have $\alpha_n = \alpha_n^n$. Supposing the coefficient of w_t in u_{n-1} is α_{n-1}^t , then we can deduce that the coefficient in u_n is $(1 - \alpha_n) \cdot \alpha_{n-1}^t = \alpha_n^t$, according to the definition of learning rate introduced in Line (4.6). ■

Proof [Proof of lemma A.6] Recall that $\bar{Q}_k(s_{k,h}, a_{k,h})$ is updated in Line (13) of Algorithm 1. Fixing $(x, a) = (s_{k,h}, a_{k,h})$ and let $u_t = \bar{Q}_{k[t]}(x, a)$ and $w_t = \hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) + b_t(x, a)$, then all the historical Q -value update for $(s_{k,h}, a_{k,h})$ can be abstracted as

$$u_t = \min\{H, (1 - \alpha_t)u_{t-1} + \alpha_t w_t\} \quad (\text{A.22})$$

Expanding the update recursively (using Lemma A.7), we obtain that

$$u_{n_k} = \alpha_{n_k}^0 H + \sum_{1 \leq t \leq n_k} \alpha_{n_k}^t w_t \quad (\text{A.23})$$

which can be rewritten as

$$\bar{Q}_k(x, a) \leq \alpha_{n_k}^0 H + \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) + \bar{V}_{k[t]-1}(x'_{k[t]}) \right) + \sum_{t=1}^{n_k} \alpha_{n_k}^t b_t(x, a) \quad (\text{A.24})$$

Because $b_t(x, a)$ is decreasing in t and by the second property of α in Proposition A.2, we have

$$\sum_{t=1}^{n_k} \alpha_{n_k}^t b_t(x, a) \leq 2\mathbb{I}[n_k > 0] b_{n_k}(x, a) \quad (\text{A.25})$$

Therefore,

$$\bar{Q}_k(x, a) \leq \alpha_{n_k}^0 H + 2\mathbb{I}[n_k > 0] \cdot b_{n_k}(x, a) + \sum_{t=1}^{n_k} \alpha_{n_k}^t \hat{Q}_{k[t]}^{*b}(x, a) + \sum_{t=1}^{n_k} \alpha_{n_k}^t \bar{V}_{k[t]-1}(x'_{k[t]}) \quad (\text{A.26})$$

At last, a similar argument can be applied to obtain a lower bound for $\underline{Q}_k(x, a)$. Subtracting the lower bound for $\underline{Q}_k(x, a)$ from Line (A.24) will complete the proof. \blacksquare

Claim A.8 *For any three positive numbers a, b , and c satisfying $a + b \geq c$, and for any $x \in (0, 1)$, the following holds:*

$$a + b \leq \text{clip}\left[a \mid \frac{xc}{2}\right] + (1+x)b \quad (\text{A.27})$$

We recall that $\text{clip}[x|y] \triangleq \mathbb{I}[x \geq y]x$ is defined at the beginning of section 4.1.

Proof If $a \geq \frac{xc}{2}$, then $\text{clip}\left[a \mid \frac{xc}{2}\right] \geq a$ and $(1+x)b \geq b$ and the claim follows. Otherwise, assume $a < \frac{xc}{2}$. Then, because $a + b \geq c$ and $x \leq 1$, we have $xb \geq x(c - a) \geq x(c - \frac{xc}{2}) = x(1 - \frac{x}{2})c \geq \frac{xc}{2} \geq a$. It follows that $a + b \leq (1+x)b \leq \text{clip}\left[a \mid \frac{xc}{2}\right] + (1+x)b$. \blacksquare

Restatement of Proposition 4.5 (Confidence Interval Length Recursion)

Suppose \mathcal{E}_{k-1} happens. Suppose $(x, a) = (s_{k,h}, a_{k,h})$ is a state-action pair visited in the k -th episode where $x \notin G_k$ is an undecided state. Let ΔQ_{past} be a shorthand for

$$\Delta Q_{\text{past}} \triangleq \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \left(\bar{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) - \underline{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) \right) \cdot \mathbb{I}\left[x'_{k[t]} \notin G_{k[t]}\right] \quad (\text{A.28})$$

We have the following recursion bound for the confidence interval length of an undecided state:

$$\left(\bar{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a) \right) \cdot \mathbb{I}[x \notin G_k] \quad (\text{A.29})$$

$$\leq \alpha_{n_{k-1}}^0 H + \left(1 + \frac{1}{H}\right) \Delta Q_{\text{past}} + \begin{cases} \text{clip}\left[4b_{n_{k-1}}(x, a) \mid \frac{\Delta(x, a)}{4H}\right] & \text{if } a \neq a^*(x) \\ \text{clip}\left[4b_{n_{k-1}}(x, a) \mid \frac{\Delta_{\min}(x)}{4H}\right] & \text{if } a = a^*(x) \end{cases} \quad (\text{A.30})$$

Proof [Proof of proposition 4.5]

For a fixed $(x, a) = (s_{k,h}, a_{k,h})$ and $x \notin G_k$, note that if $n_{k-1}(x, a) = 0$, then $\alpha_{n_{k-1}}^0 = 1$ and the inequality is true because the confidence interval has a trivial upperbound H . Otherwise, we can ignore the first term on (A.31), which will make the further analysis simpler. We first consider the case $a \neq a^*(x)$. We have the following upper bound for gap using confidence interval.

$$\begin{aligned} \Delta(x, a) &= Q^*(x, a^*(x)) - Q^*(x, a) && \text{(By the definition of gap)} \\ &= V^*(x) - Q^*(x, a) && \text{(By the definition of } V \text{ function)} \\ &\leq \bar{V}_{k-1}(x) - \underline{Q}_{k-1}(x, a) && \text{(By Lemma 4.2)} \\ &\leq 2 \max_{a' \in A_k(x)} \left(\bar{Q}_{k-1}(x, a') - \underline{Q}_{k-1}(x, a') \right) && \text{(By equation (A.4) of Lemma A.3)} \end{aligned}$$

$$\leq 2 \left(\bar{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a) \right) \quad (\text{Because } a = \pi_k(x) \text{ maximizes confidence interval})$$

Next, we decompose confidence interval using Lemma A.6.

$$\frac{\Delta(x, a)}{2} \leq \bar{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a) \quad (\text{A.31})$$

$$\leq \alpha_{n_{k-1}}^0 H + 4b_{n_{k-1}}(x, a) + \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \left(\bar{V}_{k[t]-1}(x'_{k[t]}) - \underline{V}_{k[t]-1}(x'_{k[t]}) \right) \quad (\text{A.32})$$

Then we can apply the clipping trick.

$$\begin{aligned} & \bar{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a) \\ & \leq \alpha_{n_{k-1}}^0 H + \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \frac{\Delta(x, a)}{4H} \right] + \left(1 + \frac{1}{H} \right) \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \left(\bar{V}_{k[t]-1}(x'_{k[t]}) - \underline{V}_{k[t]-1}(x'_{k[t]}) \right) \\ & \hspace{15em} (\text{By Claim A.8}) \\ & \leq \alpha_{n_{k-1}}^0 H + \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \frac{\Delta(x, a)}{4H} \right] \\ & \quad + \left(1 + \frac{1}{H} \right) \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \left(\bar{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) - \underline{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) \right) \\ & \hspace{15em} (\text{Because } a'_{k[t]} = \pi_{k[t]}(x'_{k[t]}) \text{ maximizes the confidence interval}) \\ & = \alpha_{n_{k-1}}^0 H + \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \frac{\Delta(x, a)}{4H} \right] \\ & \quad + \left(1 + \frac{1}{H} \right) \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \left(\bar{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) - \underline{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) \right) \cdot \mathbb{I} \left[x'_{k[t]} \notin G_{k[t]} \right] \\ & \hspace{15em} (\text{Because, by definition, } x'_{k[t]} \notin G_{k[t]}) \end{aligned}$$

At last, we can add the indicator to LHS too, because the proposition's statement only considers "undecided" state $x \notin G_k$.

To apply similar proof of the $a \neq a^*(x)$ case, we need to get a similar lower bound for the confidence interval of the selected state action pair like Line (A.31). We note that $x \notin G_k$ means that $|A_k(x)| > 1$, so according to our unique optimal action assumption 3.1, at least one sub-optimal action is still in $A_k(x)$. Similarly, we have the following:

$$\begin{aligned} \Delta_{\min}(x) & \leq 2 \max_{a' \in A_k(x)} \left(\bar{Q}_{k-1}(x, a') - \underline{Q}_{k-1}(x, a') \right) \quad (\text{A.33}) \\ & = 2 \left(\bar{Q}_{k-1}(x, a^*(x)) - \underline{Q}_{k-1}(x, a^*(x)) \right) \\ & \hspace{15em} (\text{Because } \pi_k(x) = a^*(x) \text{ maximizes the confidence interval}) \end{aligned}$$

The remaining deduction follows the $a \neq a^*(x)$ case. ■

A.5. Proof of lemma 4.6

Restatement of Lemma 4.6 (Solving Recursion)

We define clipped reward function as

$$\check{b}_{n_k}(s_{k,h}, a_{k,h}) \triangleq \text{clip} \left[4b_{n_k}(s_{k,h}, a_{k,h}) \mid \max \left(\frac{\Delta(s_{k,h}, a_{k,h})}{4H}, \frac{\Delta_{\min}(s_{k,h})}{4H} \right) \right] \quad (\text{A.34})$$

When event \mathcal{E} happens, we can upper bound the regret by a linear combination of clipped reward:

$$\sum_{k=1}^K \sum_{h=1}^H \left(\bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \cdot \mathbb{I}[s_{k,h} \notin G_k] \quad (\text{A.35})$$

$$\leq e^2 H^2 SA + e^2 H \sum_{k=1}^K \sum_{h=1}^H \check{b}_{n_{k-1}}(s_{k,h}, a_{k,h}) \cdot \mathbb{I}[s_{k,h} \notin G_k] \quad (\text{A.36})$$

Proof [Proof of lemma 4.6] We repeatedly use Proposition 4.5 to expand the confidence interval. For the summation of confidence intervals over episodes $1, \dots, K$ at a fixed horizon h , we want to express it as the sum of clipped reward at later horizons' states and prove the coefficient before each state has a desired upper bound. The expression below is the format of linear combination for the summation of confidence intervals at a fixed level h . Instead of calculating the exact coefficient, we will prove a coefficient upper bound only related with h, h' : $w(h, h')$.

$$\begin{aligned} & \sum_k \left(\bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \cdot \mathbb{I}[s_{k,h} \notin G_k] \\ & \leq \sum_{k', h'} w(h, h') \cdot \check{b}_{n_{k'-1}}(s_{k', h'}, a_{k', h'}) \cdot \mathbb{I}[s_{k', h'} \notin G_{k'}] \end{aligned} \quad (\text{A.37})$$

Considering a fixed bonus on the RHS, $\check{b}_{n_{k'-1}}(x', a')$ on level h' , according to the last term in Proposition 4.5, it will be contained in some previous state's confidence interval. We suppose the previous not in $G_{k'}$ state on episode k' lied on horizon h_1 , which has notation s_{k', h_1} , and we chose action a_{k', h_1} there. We can observe that only the expansion of $(x, a) = (s_{k', h_1}, a_{k', h_1})$'s confidence interval on episode $k (k \geq k')$, i.e. $\bar{Q}_k(x, a) - \underline{Q}_k(x, a)$, will contain $\check{b}_{n_{k'-1}}(x', a')$. From property (3) of proposition A.2, we know $\sum_{n=t}^{\infty} \alpha_n^t \leq 1 + \frac{1}{H}$, so we can have the following reduction $w(h, h') \leq w(h, h_1) \cdot (1 + \frac{1}{H})^2$ and $w(h, h) = 1$. The square here comes from the property of α and the leading coefficient in the last term of Proposition 4.5. By induction, we can prove that $w(h, h') \leq (1 + \frac{1}{H})^{2(h'-h)}$. Therefore the contribution of $\check{b}_{n_{k'-1}}(x', a')$ to the whole regret summation is upper bounded by $\sum_{h \leq h'} w(h, h') \leq e^2 H$. The calculation of $\alpha_n^0 H$ is similar. Combining these two parts will produce the desired result. \blacksquare

A.6. Proof of Theorem A.1

Restatement of Theorem A.1 (Main Result Under Assumption 3.1)

Under Assumption 3.1, for fixed K , with probability at least $1 - \delta$, we have the following regret upper bound

$$\text{Regret}_K \leq O \left(H^2 SA + \sum_{x \in \mathcal{S}} \left(\sum_{a \neq a^*(x)} \frac{H^5}{\Delta(x, a)} \right) \log \left(\frac{SAK}{\delta} \right) \right) \quad (\text{A.38})$$

Proof [Proof of theorem A.1] We use the notation $\text{clip}[x|y]$ as defined in Proposition 4.5 and $\check{b}_k(x, a)$ as defined in Lemma 4.6. Recall $\mathcal{E} \subseteq \mathcal{E}_{k-1}$ and $\Pr\{\mathcal{E}\} \geq 1 - \delta$. Therefore, with probability at least $1 - \delta$, we have the following relations on regret:

$$\begin{aligned} & \sum_{k=1}^K (V_0^* - V_0^{\pi_k}) \Big|_{\mathcal{E}_{k-1}} \quad (\text{A.39}) \\ & \leq \sum_{k=1}^K 2\mathbb{E} \left[\sum_{h=1}^H \left(\bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \cdot \mathbb{I}[s_{k,h} \notin G_k] \Big|_{\mathcal{E}_{k-1}} \right] \quad (\text{By Lemma 4.4}) \\ & = 2\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \left(\bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \cdot \mathbb{I}[s_{k,h} \notin G_k] \Big|_{\mathcal{E}} \right] \\ & \quad (\text{Transform } \mathcal{E}_{k-1} \text{ to } \mathcal{E}) \end{aligned}$$

$$\leq 2e^2 H^2 SA + 2e^2 H \mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \check{b}_{n_{k-1}}(s_{k,h}, a_{k,h}) \cdot \mathbb{I}[s_{k,h} \notin G_k] \Big|_{\mathcal{E}} \right] \quad (\text{A.40})$$

$$\leq 2e^2 H^2 SA + 2e^2 H \mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \check{b}_{n_{k-1}}(s_{k,h}, a_{k,h}) \Big|_{\mathcal{E}} \right] \quad (\text{A.41})$$

$$\begin{aligned} & \leq 2e^2 H^2 SA + 128e^2 H \sum_{x \in \mathcal{S}} \left(\left(\sum_{a \neq a^*} \frac{H^4}{\Delta(x, a)} \right) + \frac{H^4}{\Delta_{\min}(x)} \right) \log \left(\frac{SAK}{\delta} \right) \\ & \quad (\text{By Claim A.13 proving that } \sum_{n=1}^{\infty} \text{clip} \left[\frac{c}{\sqrt{n}} \Big|_{\epsilon} \right] \leq \frac{4c^2}{\epsilon} \text{ for any constant } c) \\ & \leq O \left(H^2 SA + \sum_{x \in \mathcal{S}} \left(\sum_{a \neq a^*} \frac{H^5}{\Delta(x, a)} \right) \log \left(\frac{SAK}{\delta} \right) \right) \quad (\text{A.42}) \end{aligned}$$

where in line A.40, we use Lemma 4.6. ■

A.7. Supporting claims

Restatement of Lemma 4.1 (Concentration)

With probability at least $1 - \delta$ over the randomness of the environment, for all episodes $k \in [K]$, the following concentration inequalities hold:

$$\forall x \in \mathcal{S} \setminus G_k, a \in A_k(x), \quad \left| \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) - Q_{k[t]}^{*b}(x, a) \right) \right| \leq \frac{1}{2} b_{n_k}(x, a) \quad (\text{A.43})$$

$$\forall x \in \mathcal{S} \setminus G_k, a \in A_k(x), \left| \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(V^*(x'_{k[t]}) - Q_{k[t]}^{*r}(x, a) \right) \right| \leq \frac{1}{2} b_{n_k}(x, a) \quad (\text{A.44})$$

Proof [Proof of lemma 4.1] We prove Line (A.43) first. We can observe that $\hat{Q}_{k[t]}^{*b}(x, a) - Q_{k[t]}^{*b}(x, a)$ is a martingale difference sequence w.r.t the filtration being the sigma field generated by all the random variables until episode $k[t]$. By the property of α_t^i (see Proposition A.2) and according to Azuma-Hoeffding inequality, we have for fixed x, a, k , w.p. $1 - \frac{\delta}{SAK}$,

$$\left| \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\hat{Q}_{k[t]}^{*b}(x, a) - Q_{k[t]}^{*b}(x, a) \right) \right| \leq \sqrt{2H^2 \sum_{t=1}^{n_k} (\alpha_{n_k}^t)^2 \log \left(\frac{SAK}{\delta} \right)} \leq 2\sqrt{\frac{H^3}{n_k} \log \left(\frac{SAK}{\delta} \right)} \quad (\text{A.45})$$

Next, we prove Line (A.44). By definition, $x'_{k[t]}$ represents the reward division between $Q_{k[t]}^{*b}(x, a)$ and $Q_{k[t]}^{*r}(x, a)$, so we know that the expected V^* function of $x'_{k[t]}$ equals the Q^{*r} function of x on episode $k[t]$, where the randomness comes from the uncertainty of $x'_{k[t]}$. Therefore $\mathbb{E}[V^*(x'_k)] = Q_k^{*r}(x, a)$. The remaining proof is similar to proving Line (A.43). ■

Claim A.13 (bounded summation for clipped function) *The summation of a clipped function which scales proportionally to the inverse of the square root of the variable n has the following bound:*

$$\sum_{n=1}^{\infty} \text{clip} \left[\frac{c}{\sqrt{n}} \mid \epsilon \right] \leq \frac{4c^2}{\epsilon} \quad (\text{A.46})$$

Proof When $n \geq \lceil c^2 \epsilon^{-2} \rceil$, $\text{clip} \left[\frac{c}{\sqrt{n}} \mid \epsilon \right] = 0$, so we only calculate the first $\lceil c^2 \epsilon^{-2} \rceil$ terms. Then we have:

$$\sum_{n=1}^{\lceil c^2 \epsilon^{-2} \rceil} \frac{c}{\sqrt{n}} \leq \frac{4c^2}{\epsilon} \quad (\text{A.47})$$

■

Appendix B. Regret Analysis for General MDPs

In this section we prove Theorem 1.1.

Theorem B.1 (Main Regret Bound) *For fixed K , with probability at least $1 - \delta$, we have the following regret upper bound for our algorithm:*

$$\text{Regret}_K \leq O \left(H^2 SA + \left(\sum_{x \in \mathcal{S}} \sum_{a \notin Z_{\text{opt}}(x)} \frac{H^5}{\Delta(x, a)} + \frac{H^5 |Z_{\text{mul}}|}{\Delta_{\text{min}}} \right) \log \left(\frac{SAK}{\delta} \right) \right) \quad (\text{B.1})$$

Note Theorem B.1 implies Theorem 1.1 in Section 1

Definition B.2 (Range Function) For each episode $k \in [K]$ and state action pair $(x, a) \in \mathcal{S} \setminus G_k \times A_k(x)$, n_k represents abbreviation for $n_k(x, a)$, we define the following quantities as range function.

$$\Delta Q_{k-1}(x, a) = \alpha_{n_{k-1}}^0 H + 4b_{n_{k-1}}(x, a) + \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \Delta V_{k[t]-1}(x'_{k[t]}) \quad (\text{B.2})$$

Similarly, for each episode $k \in [K]$ and state $x \in \mathcal{S} \setminus G_k$, we define

$$\Delta V_{k-1}(x) = \Delta Q_{k-1}(x, a_k) \text{ where } a_k = \operatorname{argmax}_{a \in A_k(x)} \bar{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a) \quad (\text{B.3})$$

We want to show that the range function ΔQ and ΔV defined above are valid upper bound for original confidence interval.

Lemma B.3 (Valid Upper Bound for Confidence Interval) For any episode k , state x , and action a , we have the following lower bound for the range function:

$$\Delta Q_k(x, a) \geq \bar{Q}_k(x, a) - \underline{Q}_k(x, a) \quad (\text{B.4})$$

$$\Delta V_k(x) \geq \bar{V}_k(x) - \underline{V}_k(x) \quad (\text{B.5})$$

Proof [Proof of Lemma B.3] We use induction to prove this lemma. For the base case where we denote any transition destination after horizon H to be \perp , we define $\Delta V_k(\perp) = 0$ for any $k \in [K]$ and the argument is valid. Now, we assume that for $x \in \{\mathcal{S}_i\}_{i>h} \cup \perp$, we have $\Delta Q_k(x, a) \geq \bar{Q}_k(x, a) - \underline{Q}_k(x, a)$ for any k, a and $\Delta V_k(x) \geq \bar{V}_k(x) - \underline{V}_k(x)$ for any k . We want to prove that the argument is also valid for $x \in \mathcal{S}_h$.

$$\begin{aligned} \bar{Q}_k(x, a) - \underline{Q}_k(x, a) &\leq \alpha_{n_k}^0 H + \mathbb{I}[n > 0] 4b_k(x, a) + \sum_{t=1}^{n_k} \alpha_{n_k}^t \left(\bar{V}_{k[t]-1}(x'_{k[t]}) - \underline{V}_{k[t]-1}(x'_{k[t]}) \right) \\ &\leq \alpha_{n_k}^0 H + 4b_k(x, a) + \sum_{t=1}^{n_k} \alpha_{n_k}^t \Delta V_{k[t]-1}(x'_{k[t]}) \\ &= \Delta Q_k(x, a) \end{aligned} \quad \begin{array}{l} \text{(By Lemma A.6)} \\ \text{(By the induction argument)} \end{array} \quad (\text{B.6})$$

According to the updating rule of upper and lower bound of V function in Alg.1, we have

$$\begin{aligned} \Delta V_k(x) &= \Delta Q_k(x, a'_{k+1}) \quad (\text{By Definition B.2, } a'_{k+1} = \operatorname{argmax}_{a \in A_{k+1}(x)} \bar{Q}_k(x, a) - \underline{Q}_k(x, a)) \\ &\geq \bar{Q}_k(x, a'_k) - \underline{Q}_k(x, a'_k) \end{aligned} \quad (\text{B.7})$$

$$= \max_{a \in A_k(x)} \bar{Q}_k(x, a) - \underline{Q}_k(x, a) \quad (\text{B.8})$$

$$\geq \max_{a \in A_k(x)} \bar{Q}_k(x, a) - \underline{V}_k(x) \quad (\text{B.9})$$

$$= \left(\max_{a \in A_k(x)} \bar{Q}_k(x, a) \right) - \underline{V}_k(x) \quad (\text{B.10})$$

$$= \bar{V}_k(x) - \underline{V}_k(x) \quad (\text{B.11})$$

Therefore the induction argument is also valid for $x \in \mathcal{S}_h$. ■

We next utilize the half-clipping trick to clip $b_{n_k}(x, a^*(x))$ at $\Omega(\frac{\Delta_{\min}}{H})$ and get a gap-dependent regret upper bound for our algorithm.

Definition B.4 (Half-Clipped Range Function) $\forall k \in [K], x \in \mathcal{S} \setminus G_k, a \in A_k(x)$, we define half-clipped range functions by directly clipping $b_{n_{k-1}}(x, a)$ at $\Omega(\frac{\Delta_{\min}}{H})$:

$$\Delta \ddot{Q}_{k-1}(x, a) = \alpha_{n_{k-1}}^0 H + \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \frac{\Delta_{\min}}{4H} \right] + \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \Delta \ddot{V}_{k-1}(x'_{k[t]}) \quad (\text{B.12})$$

$$\Delta \ddot{V}_{k-1}(x) = \Delta \ddot{Q}_{k-1}(x, a_k) \text{ where } a_k = \operatorname{argmax}_{a \in A_k(x)} \bar{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a) \quad (\text{B.13})$$

The half-clipped range function defined above only lose at most $O(\Delta_{\min})$ compared with their unclipped counterparts ΔQ and ΔV .

Proposition B.5 (lower bound for half-clipped range function) $\forall k \in [K], x \in \mathcal{S}, a \in \mathcal{A}$, we have the following lower bound for half-clipped range function defined in Definition B.4 :

$$\Delta \ddot{Q}_k(x, a) \geq \Delta Q_k(x, a) - \frac{\Delta_{\min}}{4} \quad (\text{B.14})$$

$$\Delta \ddot{V}_k(x, a) \geq \Delta V_k(x, a) - \frac{\Delta_{\min}}{4} \quad (\text{B.15})$$

Proof According to Definition B.4, one step expansion will lose at most $\frac{\Delta_{\min}}{4H}$ because of the clipping function. Our MDP has horizon H , so any half-clipped range function will lose at most $\frac{\Delta_{\min}}{4}$ compared with its corresponding range function. ■

Recall that in Lemma 4.4, we use the sum of Q-functions' confidence intervals to upper bound regret. Now, we will prove that this upper bound is still valid if we replace actual confidence interval with half-clipped range function.

Lemma B.6 (decompose regret into sum of half-clipped range function) $\forall k \in [K]$, conditioning on event \mathcal{E}_{k-1} , our algorithm's regret can be upper bounded by half-clipped range functions.

$$V_0^* - V_0^{\pi_k} \Big| \mathcal{E}_{k-1}, \mathcal{F}_{k-1} \leq 4\mathbb{E} \left[\sum_{h=1}^H \Delta \ddot{Q}_{k-1}(s_{k,h}, a_{k,h}) \cdot \mathbb{I}[s_{k,h} \notin G_k] \Big| \mathcal{E}_{k-1}, \mathcal{F}_{k-1} \right] \quad (\text{B.16})$$

Proof We use \mathcal{F}_{k-1} to denote the filtration generated by all the random variables before episode k . By similar expansion used in the proof of Lemma 4.4, we have

$$V_{k,0}^* - V_{k,0}^{\pi_k} \Big| \mathcal{E}_{k-1}, \mathcal{F}_{k-1} \quad (\text{B.17})$$

$$\leq 2\mathbb{E} \left[\sum_{h=1}^H \left(\bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \right) \cdot \mathbb{I}[a_{k,h} \notin Z_{\text{opt}}(s_{k,h})] \Big| \mathcal{E}_{k-1}, \mathcal{F}_{k-1} \right] \quad (\text{B.18})$$

$$\leq 2\mathbb{E} \left[\sum_{h=1}^H \Delta Q_{k-1}(s_{k,h}, a_{k,h}) \cdot \mathbb{I}[a_{k,h} \notin Z_{\text{opt}}(s_{k,h})] \Big| \mathcal{E}_{k-1}, \mathcal{F}_{k-1} \right] \quad (\text{B.19})$$

For the selected sub-optimal action on episode k , it satisfies

$$\Delta Q_{k-1}(s_{k,h}, a_{k,h}) \geq \bar{Q}_{k-1}(s_{k,h}, a_{k,h}) - \underline{Q}_{k-1}(s_{k,h}, a_{k,h}) \geq \frac{\Delta(s_{k,h}, a_{k,h})}{2} \quad (\text{By Line (A.31)})$$

According to the previous Proposition B.5 that the half-clipped range function decreases by at most $\frac{\Delta_{\min}}{4}$, when $a_{k,h} \notin Z_{\text{opt}}(s_{k,h})$, we have

$$\Delta \ddot{Q}_{k-1}(s_{k,h}, a_{k,h}) \geq \frac{1}{2} \Delta Q_{k-1}(s_{k,h}, a_{k,h}) \quad (\text{B.20})$$

Finally, $\forall k, h, x, a$, we have $\mathbb{I}[a_{k,h} \notin Z_{\text{opt}}(s_{k,h})] \leq \mathbb{I}[s_{k,h} \notin G_k]$ (i.e. our algorithm never recommends a sub-optimal action after it has found the best action). Replacing ΔQ_{k-1} with $\Delta \ddot{Q}_{k-1}$ and $\mathbb{I}[s_{k,h} \notin G_k]$ with $\mathbb{I}[a_{k,h} \notin Z_{\text{opt}}(s_{k,h})]$ will yield the desired result. \blacksquare

In the following proposition, we incorporate previous clipping trick for suboptimal actions and unique best action into current range function analysis.

Proposition B.7 (upper bound for range function) *Suppose \mathcal{E}_{k-1} happens. Suppose $(x, a) = (s_{k,h}, a_{k,h})$ is a state-action pair visited in the k -th episode where $x \notin G_k$ is an undecided state. Let $\Delta \ddot{Q}_{\text{past}}$ be a shorthand for*

$$\Delta \ddot{Q}_{\text{past}} \triangleq \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \Delta \ddot{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) \cdot \mathbb{I}[x'_{k[t]} \notin G_{k[t]}] \quad (\text{B.21})$$

Then, we have the recursion for the CI length of undecided state:

$$\begin{aligned} & \Delta \ddot{Q}_{k-1}(x, a_k) \cdot \mathbb{I}[x \notin G_k] \\ & \leq \alpha_{n_{k-1}}^0 H + \left(1 + \frac{1}{H}\right) \Delta \ddot{Q}_{\text{past}} + \begin{cases} \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \max \left(\frac{\Delta(x, a)}{8H}, \frac{\Delta_{\min}}{4H} \right) \right] & \text{if } a \notin Z_{\text{opt}}^h(x) \\ \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \max \left(\frac{\Delta_{\min}(x)}{8H}, \frac{\Delta_{\min}}{4H} \right) \right] & \text{if } a \in Z_{\text{opt}}(x) \end{cases} \end{aligned} \quad (\text{B.22})$$

Proof We first prove the $a \notin Z_{\text{opt}}(x)$ for (B.22). To use the clipping trick here, we need to prove a lower bound for $\Delta \ddot{Q}_{k-1}(x, a)$ like Line (A.31).

$$\Delta \ddot{Q}_{k-1}(x, a) \geq \frac{1}{2} \Delta Q_{k-1}(x, a) \quad (\text{By Line (B.20)})$$

$$\begin{aligned}
 &\geq \frac{1}{2} \left(\overline{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a) \right) && \text{(By Lemma B.3)} \\
 &\geq \frac{\Delta(x, a)}{4} && \text{(By Line (A.31))}
 \end{aligned}$$

Then, we can apply the clipping trick used in Proposition 4.5:

$$\Delta \ddot{Q}_{k-1}(x, a) = \alpha_{n_{k-1}}^0 H + \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \frac{\Delta_{\min}}{4H} \right] + \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \Delta \ddot{V}_{k-1}(x'_k[t]) \quad (\text{B.23})$$

$$\leq \alpha_{n_{k-1}}^0 H + \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \max \left(\frac{\Delta(x, a)}{8H}, \frac{\Delta_{\min}}{4H} \right) \right] \quad (\text{B.24})$$

$$+ \left(1 + \frac{1}{H} \right) \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \Delta \ddot{V}_{k[t]-1}(x'_{k[t]}) \quad (\text{B.25})$$

$$= \alpha_{n_{k-1}}^0 H + \text{clip} \left[4b_{n_{k-1}}(x, a) \mid \max \left(\frac{\Delta(x, a)}{8H}, \frac{\Delta_{\min}}{4H} \right) \right] \quad (\text{B.26})$$

$$+ \left(1 + \frac{1}{H} \right) \sum_{t=1}^{n_{k-1}} \alpha_{n_{k-1}}^t \Delta \ddot{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]}) \cdot \mathbb{I} \left[x'_{k[t]} \notin G_{k[t]} \right] \quad (\text{B.27})$$

where in the last line we can use $\Delta \ddot{Q}_{k[t]-1}(x'_{k[t]}, a'_{k[t]})$ to replace $\Delta \ddot{V}_{k[t]-1}(x'_{k[t]})$ because by Definition B.4 that $\Delta \ddot{V}_{k-1}(x) = \Delta \ddot{Q}_{k-1}(x, a_k)$ where $a_k = \arg \max_{a \in A_k(x)} \overline{Q}_{k-1}(x, a) - \underline{Q}_{k-1}(x, a)$.

Finally, adding an indicator to LHS will produce the desired result.

Proof for the second case $a \in Z_{\text{opt}}(x)$ of (B.22) is similar to Proposition 4.5. Note $\Delta_{\min}(x) = 0$ if $|Z_{\text{opt}}(x)| > 1$

■

Lemma B.8 (iterated clipping) *Conditioning on event \mathcal{E} , we can upper bound the regret by a linear combination of clipped reward defined as:*

$$\ddot{b}_k(s_{k,h}, a_{k,h}) = \text{clip} \left[4b_k(s_{k,h}, a_{k,h}) \mid \max \left(\frac{\Delta(s_{k,h}, a_{k,h})}{8H}, \frac{\Delta_{\min}(s_{k,h})}{8H} \cdot \mathbb{I} [|Z_{\text{opt}}(s_{k,h})| = 1], \frac{\Delta_{\min}}{4H} \right) \right] \quad (\text{B.28})$$

which is a generalized version of the clipped reward defined in Lemma 4.6 :

$$\sum_{k=1}^K \sum_{h=1}^H \Delta \ddot{Q}_{k-1}(s_{k,h}, a_{k,h}) \cdot \mathbb{I} [s_{k,h} \notin G_k] \leq e^2 H^2 S A + e^2 H \sum_{k=1}^K \sum_{h=1}^H \ddot{b}_{k-1}(s_{k,h}, a_{k,h}) \cdot \mathbb{I} [s_{k,h} \notin G_k] \quad (\text{B.29})$$

Proof The proof idea follows the same as Lemma 4.6.

■

Now we are ready to bound the regret.

Restatement of Theorem B.1 (Main Regret Bound)

For fixed K , with probability at least $1 - \delta$, we have the following regret upper bound for our algorithm:

$$\text{Regret}_K \leq O \left(H^2 SA + \left(\sum_{x \in \mathcal{S}} \sum_{a \notin Z_{\text{opt}}(x)} \frac{H^5}{\Delta(x, a)} + \frac{H^5 |Z_{\text{mul}}|}{\Delta_{\min}} \right) \log \left(\frac{SAK}{\delta} \right) \right) \quad (\text{B.30})$$

Proof First, we transform regret into the summation of clipped reward. Recall $\mathcal{E} \subseteq \mathcal{E}_{k-1}$ and $\Pr\{\mathcal{E}\} \geq 1 - \delta$. With probability at least $1 - \delta$, we have the following relations on regret:

$$\begin{aligned} & \sum_{k=1}^K (V_0^* - V_0^{\pi_k}) | \mathcal{E}_{k-1} \quad (\text{B.31}) \\ & \leq \sum_{k=1}^K 4\mathbb{E} \left[\sum_{h=1}^H \Delta \ddot{Q}_{k-1}(s_{k,h}, a_{k,h}) \cdot \mathbb{I}[s_{k,h} \notin G_k] \middle| \mathcal{E}_{k-1} \right] \quad (\text{By Lemma B.6}) \\ & = 4\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \Delta \ddot{Q}_{k-1}(s_{k,h}, a_{k,h}) \cdot \mathbb{I}[s_{k,h} \notin G_k] \middle| \mathcal{E} \right] \quad (\text{Transform } \mathcal{E}_{k-1} \text{ to } \mathcal{E}) \\ & \leq 4e^2 H^2 SA + 4e^2 H\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \ddot{b}_{k-1}(s_{k,h}, a_{k,h}) \cdot \mathbb{I}[s_{k,h} \notin G_k] \middle| \mathcal{E} \right] \quad (\text{By Lemma B.8}) \\ & \leq 4e^2 H^2 SA + 4e^2 H\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \ddot{b}_{k-1}(s_{k,h}, a_{k,h}) \middle| \mathcal{E} \right] \quad (\text{B.32}) \end{aligned}$$

Next, we use Claim A.13 to upper bound $\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \ddot{b}_{k-1}(s_{k,h}, a_{k,h}) \middle| \mathcal{E} \right]$.

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \ddot{b}_{k-1}(s_{k,h}, a_{k,h}) \middle| \mathcal{E} \right] \quad (\text{B.33})$$

$$\leq 128 \sum_{x \in \mathcal{S}} \left(\left(\sum_{a \notin Z_{\text{opt}}(x)} \frac{H^4}{\Delta(x, a)} \right) + \frac{H^4 \mathbb{I}[|Z_{\text{opt}}(x)| = 1]}{\Delta_{\min}(x)} + \frac{H^4 |Z_{\text{opt}}(x)| \mathbb{I}[|Z_{\text{opt}}(x)| > 1]}{\Delta_{\min}} \right) \log \left(\frac{SAK}{\delta} \right) \quad (\text{By Claim A.13})$$

$$\leq O \left(\sum_{x \in \mathcal{S}} \left(\left(\sum_{a \notin Z_{\text{opt}}(x)} \frac{H^4}{\Delta(x, a)} \right) + \frac{H^4 |Z_{\text{opt}}(x)| \mathbb{I}[|Z_{\text{opt}}(x)| > 1]}{\Delta_{\min}} \right) \log \left(\frac{SAK}{\delta} \right) \right) \quad (\text{B.34})$$

$$\leq O \left(\left(\sum_{x \in \mathcal{S}} \sum_{a \notin Z_{\text{opt}}(x)} \frac{H^4}{\Delta(x, a)} + \frac{H^4 |Z_{\text{mul}}|}{\Delta_{\min}} \right) \log \left(\frac{SAK}{\delta} \right) \right) \quad (\text{B.35})$$

Plugging Line (B.35) into Line (B.32) will produce the wanted result. ■

By discarding the clipping, we can also get a gap independent expected regret upper bound for our algorithm. We remark that the dependency on H in our bounds are not tight. We leave it as a future work to obtain a bound with a tight dependency on H .

Corollary B.10 (Gap-independent Bound) *For general MDPs and fixed K , with probability at least $1 - \delta$, our algorithm has the following gap independent regret upper bound:*

$$\text{Regret}_K \leq O\left(H^2 SA + \sqrt{H^5 SAT \log(SAK/\delta)}\right) \quad (\text{B.36})$$

Proof Starting at Line (B.32), we use another way to upper bound Line (B.33).

$$\mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H \ddot{b}_k(s_{k,h}, a_{k,h}) \middle| \mathcal{E} \right] \leq \mathbb{E} \left[\sum_{k=1}^K \sum_{h=1}^H b_k(s_{k,h}, a_{k,h}) \middle| \mathcal{E} \right] \quad (\text{B.37})$$

$$\leq \mathbb{E} \left[\sum_{(x,a) \in \mathcal{S} \times \mathcal{A}} \sum_{n=1}^{n_K(x,a)} \sqrt{\frac{H^3 \log\left(\frac{SAK}{\delta}\right)}{n}} \middle| \mathcal{E} \right] \quad (\text{B.38})$$

$$\leq S \cdot A \cdot \sqrt{H^3 \log\left(\frac{SAK}{\delta}\right) \frac{T}{S \cdot A}} \quad (\text{B.39})$$

$$\leq 2\sqrt{H^3 \log\left(\frac{SAK}{\delta}\right) SAT} \quad (\text{B.40})$$

Plugging Line (B.40) into Line (B.32) will produce the desired result. ■

Appendix C. New Instance Dependent Lower Bound Regarding Minimal Gap: Proof of Theorem 5.1

In this section, we prove our new lower bound. We first introduce some necessary definitions.

Definition C.1 (consistent algorithm) *We say an algorithm ALG is consistent if for $\forall 0 < \alpha < 1$ and any MDP M , when K approaches infinity, its incurred regret satisfies*

$$\lim_{K \rightarrow +\infty} \frac{\text{Regret}_K(M, \text{ALG})}{K^\alpha} = 0 \quad (\text{C.1})$$

Definition C.2 *Let \mathbb{P} and \mathbb{Q} be probability measures on the same measurable space (Ω, \mathcal{F}) . Relative entropy is defined as*

$$D(\mathbb{P}, \mathbb{Q}) = \mathbb{E} \left[\log \left(\frac{d\mathbb{P}}{d\mathbb{Q}} \right) \right] \quad (\text{C.2})$$

Lemma C.3 (Divergence Decomposition) *Let one MDP M has transition probability and reward distribution $\{\mathbb{P}_{s,a}, \mathbb{R}_{s,a}\}$ and another MDP M' has the same transition probability but different reward $\{\mathbb{P}_{s,a}, \mathbb{R}'_{s,a}\}$. We fix an algorithm ALG, and let $\mathbb{P}_{M, \text{ALG}}$ and $\mathbb{P}_{M', \text{ALG}}$ be the probability measure over state-action pairs of running algorithm ALG on model M and M' . Then, we have the following equality:*

$$D(\mathbb{P}_{M, \text{ALG}}, \mathbb{P}_{M', \text{ALG}}) = \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathbb{E}_{M, \text{ALG}}[n_K(s, a)] D(\mathbb{R}_{s,a}, \mathbb{R}'_{s,a}) \quad (\text{C.3})$$

Proof [Proof of Lemma C.3] The proof mostly follows that of Lemma 15.1 in [Lattimore and Szepesvári \(2020\)](#). We use $\mathcal{P}, \mathcal{R}, \mathcal{R}'$ to denote these distributions' respective probability density function and let π_k be the policy induced by ALG. According to our MDP's procedure, we write down the expression for distribution $\mathbb{P}_{M,\text{ALG}}$'s density function. To make the expression compact, we concatenate K episodes and use subscript $s_{k,h}$, $a_{k,h}$, and $r_{k,h}$ to represent the variables on the episode k , horizon h .

$$\mathcal{P}_{M,\text{ALG}}(s_{1,1}, a_{1,1}, r_{1,1}, \dots, s_{K,H}, a_{K,H}, r_{K,H}) \quad (\text{C.4})$$

$$= \prod_{k=1}^K \mathcal{P}_0(s_{k,1}) \prod_{h=1}^H \pi_k(s_{k,h}, a_{k,h}) \mathcal{R}_{s_{k,h}, a_{k,h}}(r_{k,h}) \mathcal{P}_{s_{k,h}, a_{k,h}}(s_{k,h+1}) \quad (\text{C.5})$$

We can similarly get the expression for $\mathcal{P}_{M',\text{ALG}}$, using its \mathbb{R}' . By canceling out the shared function, we have the following equality:

$$\log \left(\frac{d\mathbb{P}_{M,\text{ALG}}}{d\mathbb{P}_{M',\text{ALG}}} \right) = \sum_{k=1}^K \sum_{h=1}^H \log \left(\frac{\mathcal{R}_{s_{k,h}, a_{k,h}}(r_{k,h})}{\mathcal{R}'_{s_{k,h}, a_{k,h}}(r_{k,h})} \right) \quad (\text{C.6})$$

$$\text{D}(\mathbb{P}_{M,\text{ALG}}, \mathbb{P}_{M',\text{ALG}}) = \mathbb{E}_{M,\text{ALG}} \left[\log \left(\frac{d\mathbb{P}_{M,\text{ALG}}}{d\mathbb{P}_{M',\text{ALG}}} \right) \right] \quad (\text{By Definition C.2})$$

$$= \mathbb{E}_{M,\text{ALG}} \left[\sum_{k=1}^K \sum_{h=1}^H \log \left(\frac{\mathcal{R}_{s_{k,h}, a_{k,h}}(r_{k,h})}{\mathcal{R}'_{s_{k,h}, a_{k,h}}(r_{k,h})} \right) \right] \quad (\text{By Line C.6})$$

$$= \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{M,\text{ALG}} \left[\log \left(\frac{\mathcal{R}_{s_{k,h}, a_{k,h}}(r_{k,h})}{\mathcal{R}'_{s_{k,h}, a_{k,h}}(r_{k,h})} \right) \right] \quad (\text{C.7})$$

$$= \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{M,\text{ALG}} \left[\text{D}(\mathbb{R}_{s_{k,h}, a_{k,h}}, \mathbb{R}'_{s_{k,h}, a_{k,h}}) \right] \quad (\text{C.8})$$

$$= \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathbb{E}_{M,\text{ALG}} \left[\sum_{k=1}^K \sum_{h=1}^H \mathbb{I}[(s_{k,h}, a_{k,h}) = (s, a)] \text{D}(\mathbb{R}_{s_{k,h}, a_{k,h}}, \mathbb{R}'_{s_{k,h}, a_{k,h}}) \right] \quad (\text{C.9})$$

$$= \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathbb{E}_{M,\text{ALG}}[n_K(s, a)] \text{D}(\mathbb{R}_{s,a}, \mathbb{R}'_{s,a}) \quad (\text{C.10})$$

■

We also need the following inequality.

Lemma C.4 (Bretagnolle–Huber inequality) *Let \mathbb{P} and \mathbb{Q} be probability measures on the same measurable space (Ω, \mathcal{F}) , and let $A \in \mathcal{F}$ be an arbitrary event, $A^c = \Omega \setminus A$ be its complement. Then we have the following inequality:*

$$\mathbb{P}(A) + \mathbb{Q}(A^c) \geq \frac{1}{2} e^{-\text{D}(\mathbb{P}, \mathbb{Q})} \quad (\text{C.11})$$

Claim C.5 For two Bernoulli distribution $\mathbb{B}(\frac{1}{2})$ and $\mathbb{B}(\frac{1}{2} + x)$ with $x \leq 1/4$, their relative entropy satisfies $D(\mathbb{B}(\frac{1}{2}), \mathbb{B}(\frac{1}{2} + x)) \leq \frac{8x^2}{3}$

Proof

$$D(\mathbb{B}(1/2), \mathbb{B}(1/2 + x)) = \frac{1}{2} \left(\ln \frac{1/2}{(1/2 + x)} + \ln \frac{1/2}{(1/2 - x)} \right) \quad (\text{C.12})$$

$$= -\frac{1}{2} \ln 4(1/2 + x)(1/2 - x) \quad (\text{C.13})$$

$$= -\frac{1}{2} \ln 1 - 4x^2 \quad (\text{C.14})$$

$$\leq -\frac{1}{2} \cdot \frac{-4x^2}{1 - 4x^2} \quad \left(\frac{x}{1+x} < \ln 1 + x < x \text{ for } x > -1 \right)$$

$$= \frac{2x^2}{1 - 4x^2} \quad (\text{C.15})$$

$$\leq \frac{8x^2}{3} \quad (x \leq \frac{1}{4})$$

■

Now we are ready to show this hard instance gives us the desired lower bound.

Theorem C.6 (Regret Lower Bound for a Hard Instance) For the hard instance described above Figure 2, any consistent algorithm incurs expected regret at least $\frac{3(n-1)A \ln K}{32\gamma}$, larger than $\frac{SA \ln K}{32\Delta_{\min}}$ in terms of S, A, K, Δ_{\min} .

Proof [Proof of Theorem C.6] In Figure 2, we construct our family of hard instance for $|Z_{\text{mul}}| \approx \frac{SA}{2}$. For other $S \leq |Z_{\text{mul}}| \leq \frac{SA}{2}$, we can similarly construct their instance family by reducing the number of state action pairs on the last layer. For any consistent algorithm ALG, any fixed $i \in [2, n]$, we define event $A_{i,j} = \{n_K(x_i, a_j) \geq \frac{K}{2}\}$. We use $a^*(x)$ to denote the optimal action for state x . By Bretagnolle–Huber inequality in Lemma C.4, we have

$$\mathbb{P}_{M, \text{ALG}}(A_{i,j}) + \mathbb{P}_{M_{i,j}, \text{ALG}}(A_{i,j}^c) \geq \frac{1}{2} e^{-D(\mathbb{P}_{M, \text{ALG}}, \mathbb{P}_{M_{i,j}, \text{ALG}})} \quad (\text{C.16})$$

$$\geq \frac{1}{2} e^{-\mathbb{E}[n_K(x_i, a_j)] \cdot D(\mathbb{B}(1/2), \mathbb{B}(1/2 + 2\gamma))} \quad (\text{C.17})$$

$$\geq \frac{1}{2} e^{-\mathbb{E}[n_K(x_i, a_j)] \cdot \frac{32\gamma^2}{3}} \quad (\text{C.18})$$

By our assumption that ALG is consistent, we have the following inequality

$$\text{Regret}_K(M, \text{ALG}) + \text{Regret}_K(M_{i,j}, \text{ALG}) \geq \mathbb{P}_{M, \text{ALG}}(A_{i,j}) \cdot \frac{K}{2} \cdot \gamma + \mathbb{P}_{M_{i,j}, \text{ALG}}(A_{i,j}^c) \cdot \frac{K}{2} \cdot \gamma \quad (\text{C.19})$$

$$= \frac{K\gamma}{2} (\mathbb{P}_{M, \text{ALG}}(A_{i,j}) + \mathbb{P}_{M_{i,j}, \text{ALG}}(A_{i,j}^c)) \quad (\text{C.20})$$

$$\geq \frac{K\gamma}{4} e^{-\mathbb{E}_{M, \text{ALG}}[n_K(x_i, a_j)] \cdot \frac{32\gamma^2}{3}} \quad (\text{C.21})$$

In line C.19, visiting (x_i, a_j) in M incurs regret γ and not visiting (x_i, a_j) in $M_{i,j}$ incurs regret γ , so the two terms on the RHS lower bounds the cumulative regret in M and $M_{i,j}$. Now, let's lower bound the value $\mathbb{E}_{M, \text{ALG}} [n_K(x_i, a_j)]$ starting from an algebraic manipulations on line C.21 and then divide it by $\ln K$.

$$\liminf_{K \rightarrow +\infty} \frac{\mathbb{E}_{M, \text{ALG}} [n_K(x_i, a_j)]}{\ln K} \geq \liminf_{K \rightarrow +\infty} \frac{3}{32\gamma^2 \ln K} \ln \frac{K\gamma}{4(\text{Regret}_K(M, \text{ALG}) + \text{Regret}_k(M_{i,j}, \text{ALG}))} \quad (\text{C.22})$$

$$\begin{aligned} &\geq \frac{3}{32\gamma^2} \liminf_{K \rightarrow +\infty} \frac{\ln \frac{K\gamma}{4K^\alpha}}{\ln K} && (\text{Valid for any } 0 < \alpha < 1) \\ &\geq \frac{3}{32\gamma^2} \liminf_{K \rightarrow +\infty} \frac{(1-\alpha) \ln K + \ln \frac{\gamma}{4}}{\ln K} && (\text{C.23}) \end{aligned}$$

$$\geq \frac{3(1-\alpha)}{32\gamma^2} \quad (\text{C.24})$$

$$\geq \frac{3}{32\gamma^2} \quad (\text{C.25})$$

Line C.25 is valid because we can arbitrarily set $0 < \alpha < 1$. Moreover, Line C.25 works for any $i \in [2, n], j \in [1, A]$. Then we have the lower bound for algorithm ALG's regret on MDP M .

$$\mathbb{E} [\text{Regret}_K(M, \text{ALG})] \geq \sum_{i=2}^n \sum_{j=1}^A \mathbb{E}_{M, \text{ALG}} [n_K(x_i, a_j)] \cdot \gamma \quad (\text{C.26})$$

$$\geq \sum_{i=2}^n \sum_{j=1}^A \frac{3 \ln K}{32\gamma^2} \cdot \gamma \quad (\text{C.27})$$

$$= \frac{3(n-1)A \ln K}{32\gamma} \quad (\text{C.28})$$

$$\geq \frac{SA \ln K}{32\Delta_{\min}} \quad (\text{C.29})$$

■

Note Theorem 5.1 follows from Theorem C.6 directly because the instance in Theorem C.6 satisfies the requirements in Theorem 5.1.