



Settling the Sample Complexity of Online Reinforcement Learning

ZIHAN ZHANG, Department of Electrical and Computer Engineering, Princeton University, Princeton, United States

YUXIN CHEN, Department of Statistics and Data Science, University of Pennsylvania, Philadelphia, United States

JASON LEE, Department of Electrical and Computer Engineering, Princeton University, Princeton, United States

SIMON S. DU, Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, United States

A central issue lying at the heart of online reinforcement learning (RL) is data efficiency. While a number of recent works achieved asymptotically minimal regret in online RL, the optimality of these results is only guaranteed in a “large-sample” regime, imposing enormous burn-in cost in order for their algorithms to operate optimally. How to achieve minimax-optimal regret without incurring any burn-in cost has been an open problem in RL theory.

We settle this problem for finite-horizon inhomogeneous Markov decision processes. Specifically, we prove that a modified version of MVP (Monotonic Value Propagation), an optimistic model-based algorithm proposed by Zhang et al. [82], achieves a regret on the order of (modulo log factors)

$$\min\{\sqrt{SAH^3K}, HK\},$$

where S is the number of states, A is the number of actions, H is the horizon length, and K is the total number of episodes. This regret matches the minimax lower bound for the entire range of sample size $K \geq 1$, essentially eliminating any burn-in requirement. It also translates to a PAC sample complexity (i.e., the number of episodes needed to yield ϵ -accuracy) of $\frac{SAH^3}{\epsilon^2}$ up to log factor, which is minimax-optimal for the full ϵ -range. Further, we extend our theory to unveil the influences of problem-dependent quantities like the optimal value/cost and certain variances. The key technical innovation lies in a novel analysis paradigm (based on a new concept called “profiles”) to decouple complicated statistical dependency across the sample trajectories — a long-standing challenge facing the analysis of online RL in the sample-starved regime.

Y. Chen is supported in part by the Sloan Research Fellowship, the Google Research Scholar Award, the AFOSR grant FA9550-22-1-0198, the ONR grant N00014-22-1-2354, and the NSF grants CCF-2221009 and CCF-1907661. JDL acknowledges support of Open Philanthropy, NSF IIS-2107304, NSF CCF-2212262, ONR Young Investigator Award, NSF CAREER Award 2144994, and NSF CCF-2019844. SSD acknowledges the support of NSF IIS-2110170, NSF DMS-2134106, NSF CCF-2212261, NSF IIS-2143493, NSF CCF-2019844, NSF IIS-2229881, and the Sloan Research Fellowship.

Authors' Contact Information: Zihan Zhang, Department of Electrical and Computer Engineering, Princeton University, Princeton, New Jersey, United States; e-mail: zz5478@princeton.edu; Yuxin Chen, Department of Statistics and Data Science, University of Pennsylvania, Philadelphia, Pennsylvania, United States; e-mail: yuxinc@wharton.upenn.edu; Jason Lee, Department of Electrical and Computer Engineering, Princeton University, Princeton, New Jersey, United States; email: jasonlee@princeton.edu; Simon S. Du, Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, Washington, United States; e-mail: ssdu@cs.washington.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 0004-5411/2025/06-ART22

<https://doi.org/10.1145/3733592>

CCS Concepts: • **Theory of computation** → **Online learning theory; Reinforcement learning; Regret bounds;**

Additional Key Words and Phrases: Burn-in cost

ACM Reference Format:

Zihan Zhang, Yuxin Chen, Jason Lee, and Simon S. Du. 2025. Settling the Sample Complexity of Online Reinforcement Learning. *J. ACM* 72, 3, Article 22 (June 2025), 63 pages. <https://doi.org/10.1145/3733592>

1 Introduction

In **reinforcement learning (RL)**, an agent is often asked to learn optimal decisions (i.e., the ones that maximize cumulative reward) through real-time “trial-and-error” interactions with an unknown environment. This task is commonly dubbed as *online RL*, underscoring the critical role of adaptive online data collection and differentiating it from other RL settings that rely upon pre-collected data. A central challenge in achieving sample-efficient online RL boils down to how to optimally balance exploration and exploitation during data collection, namely, how to tradeoff the potential revenue of exploring unknown terrain/dynamics against the benefit of exploiting past experience. While decades-long effort has been invested toward unlocking the capability of online RL, how to *fully* characterize – and more importantly, attain – its fundamental performance limit remains largely unsettled.

In this article, we take an important step toward settling the sample complexity limit of online RL, focusing on tabular **Markov Decision Processes (MDPs)** with finite horizon and finite state-action space. More concretely, imagine that one seeks to learn a near-optimal policy of a time-inhomogeneous MDP with S states, A actions, and horizon length H , and is allowed to execute the MDP of interest K times to collect K sample episodes each of length H . This canonical problem is among the most extensively studied in the RL literature, with formal theoretical pursuit dating back to more than 25 years ago (e.g., Kearns and Singh [34]). Numerous works have since been devoted to improving the sample efficiency and/or refining the analysis framework [6, 7, 11, 17, 20, 26, 29, 32, 42, 49, 81, 82, 84]. As we shall elucidate momentarily, however, information-theoretic optimality has only been achieved in the “large-sample” regime. When it comes to the most challenging sample-hungry regime, there remains a substantial gap between the state-of-the-art regret upper bound and the best-known minimax lower bound, which motivates the research of this article.

1.1 Inadequacy of Prior Art: Enormous Burn-in Cost

While past research has obtained asymptotically optimal (i.e., optimal when K approaches infinity) regret bounds in the aforementioned setting, all of these results incur an enormous burn-in cost – that is, the minimum sample size needed for an algorithm to operate sample-optimally – which we explain in the sequel. For simplicity of presentation, we assume that each immediate reward lies within the normalized range $[0, 1]$ when discussing the prior art.

Minimax lower bound. To provide a theoretical benchmark, we first make note of the best-known minimax regret lower bound developed by Domingues et al. [20], Jin et al. [29]¹:

$$(\text{minimax lower bound}) \quad \Omega \left(\min \left\{ \sqrt{SAH^3K}, HK \right\} \right), \quad (1)$$

¹Let $\mathcal{X} = \{S, A, H, K, \frac{1}{\delta}\}$, where $1 - \delta$ is the target success rate (to be seen shortly). The notation $f(\mathcal{X}) = O(g(\mathcal{X}))$ (or $f(\mathcal{X}) \lesssim g(\mathcal{X})$) indicates the existence of some universal constant $c_1 > 0$ such that $f(\mathcal{X}) \leq c_1 g(\mathcal{X})$; $f(\mathcal{X}) = \Omega(g(\mathcal{X}))$ (or $f(\mathcal{X}) \gtrsim g(\mathcal{X})$) means that there exists some universal constant $c_2 > 0$ such that $f(\mathcal{X}) \geq c_2 g(\mathcal{X})$; and $f(\mathcal{X}) = \Theta(g(\mathcal{X}))$ (or $f(\mathcal{X}) \asymp g(\mathcal{X})$) means that $f(\mathcal{X}) \lesssim g(\mathcal{X})$ and $f(\mathcal{X}) \gtrsim g(\mathcal{X})$ hold simultaneously. Moreover, $\tilde{O}(\cdot)$, $\tilde{\Omega}(\cdot)$ and $\tilde{\Theta}(\cdot)$ are defined analogously, except that all logarithmic dependency on the quantities of \mathcal{X} are hidden.

assuming that the immediate reward at each step falls within $[0, 1]$ and imposing no restriction on K . Given that a regret of $O(HK)$ can be trivially achieved (as the sum of rewards in any K episodes cannot exceed HK), we shall sometimes drop the HK term and simply write

$$(\text{minimax lower bound}) \quad \Omega(\sqrt{SAH^3K}) \quad \text{if } K \geq SAH. \quad (2)$$

Prior upper bounds and burn-in cost. We now turn to the upper bounds developed in prior literature. For ease of presentation, we shall assume

$$K \geq SAH, \quad (3)$$

in the rest of this subsection unless otherwise noted. Log factors are also ignored in the discussion below.

The first article that achieves asymptotically optimal regret is Azar et al. [6], which came up with a model-based algorithm called UCBVI that enjoys a regret bound $\tilde{O}(\sqrt{SAH^3K} + H^3S^2A)$. A close inspection reveals that this regret matches the minimax lower bound (2) if and only if

$$(\text{burn-in cost of Azar et al. [6]}) \quad K \gtrsim S^3AH^3, \quad (4)$$

due to the presence of the lower-order term H^3S^2A in the regret bound. This burn-in cost is clearly undesirable, since the sample size available in many practical scenarios might be far below this requirement.

In light of its fundamental importance in contemporary RL applications (which often have very large dimensionality and relatively limited data collection capability), reducing the burn-in cost without compromising sample efficiency has emerged as a central problem in recent pursuit of RL theory [1, 18, 40, 42, 47, 49, 61, 81, 82, 87]. The state-of-the-art regret upper bounds for finite-horizon inhomogeneous MDPs can be summarized below (depending on the size of K):

$$[49] \quad \tilde{O}(\sqrt{SAH^3K} + SAH^4), \quad (5a)$$

$$[82, 87] \quad \tilde{O}(\sqrt{SAH^3K} + S^2AH^2), \quad (5b)$$

meaning that even the most advanced prior results fall short of sample optimality unless

$$(\text{best burn-in cost in past works}) \quad K \gtrsim \min\{SAH^5, S^3AH\}. \quad (6)$$

The interested reader is referred to Table 1 for more details about existing regret upper bounds and their associated sample complexities.

In summary, no prior theory was able to achieve optimal sample complexity in the data-hungry regime

$$SAH \leq K \lesssim \min\{SAH^5, S^3AH\}, \quad (7)$$

suffering from a significant barrier of either a long horizon (as in the term SAH^5) or a large state space (as in the term S^3AH). In fact, the information-theoretic limit is yet to be determined within this regime (i.e., neither the achievability results nor the lower bounds had been shown to be tight), although it has been conjectured by Ménard et al. [49] that the lower bound (1) reflects the correct scaling for any sample size K .²

Comparisons with other RL settings and key challenges. In truth, the incentives to minimize the burn-in cost and improve data efficiency arise in multiple other settings beyond online RL. For instance, in an idealistic setting that assumes access to a simulator (or a generative model) — a model that allows the learner to query arbitrary state-action pairs to draw samples — a recent

²Note that the original conjecture in Ménard et al. [49] was $\tilde{O}(\sqrt{SAH^3K} + SAH^2)$. Combining it with the trivial upper bound HK allows one to remove the term SAH^2 (with a little algebra).

Table 1. Comparisons between Our Result and Prior Works that Achieve Asymptotically Optimal Regret for Finite-horizon Inhomogeneous MDPs (with All Log Factors Omitted), where S (respectively A) is the Number of States (Respectively Actions), H is the Planning Horizon, and K is the Number of Episodes

Algorithm	Regret upper bound	Range of K that attains optimal regret	Sample complexity (or PAC bound)
MVP (this work, Theorem 1.1)	$\min \{\sqrt{SAH^3K}, HK\}$	$[1, \infty)$	$\frac{SAH^3}{\varepsilon^2}$
UCBVI [6]	$\min \{\sqrt{SAH^3K} + S^2AH^3, HK\}$	$[S^3AH^3, \infty)$	$\frac{SAH^3}{\varepsilon^2} + \frac{S^2AH^3}{\varepsilon}$
ORLC [18]	$\min \{\sqrt{SAH^3K} + S^2AH^4, HK\}$	$[S^3AH^5, \infty)$	$\frac{SAH^3}{\varepsilon^2} + \frac{S^2AH^4}{\varepsilon}$
EULER [81]	$\min \{\sqrt{SAH^3K} + S^{3/2}AH^3(\sqrt{S} + \sqrt{H}), HK\}$	$[S^2AH^3(\sqrt{S} + \sqrt{H}), \infty)$	$\frac{SAH^3}{\varepsilon^2} + \frac{S^2AH^3(\sqrt{S} + \sqrt{H})}{\varepsilon}$
UCB-Adv [84]	$\min \{\sqrt{SAH^3K} + S^2A^{3/2}H^{33/4}K^{1/4}, HK\}$	$[S^6A^4H^{27}, \infty)$	$\frac{SAH^3}{\varepsilon^2} + \frac{S^{8/3}A^2H^{11}}{\varepsilon^{4/3}}$
MVP [82]	$\min \{\sqrt{SAH^3K} + S^2AH^2, HK\}$	$[S^3AH, \infty)$	$\frac{SAH^3}{\varepsilon^2} + \frac{S^2AH^2}{\varepsilon}$
UCBMQ [49]	$\min \{\sqrt{SAH^3K} + SAH^4, HK\}$	$[SAH^5, \infty)$	$\frac{SAH^3}{\varepsilon^2} + \frac{SAH^4}{\varepsilon}$
Q-Earlysettled-Adv [42]	$\min \{\sqrt{SAH^3K} + SAH^6, HK\}$	$[SAH^9, \infty)$	$\frac{SAH^3}{\varepsilon^2} + \frac{SAH^6}{\varepsilon}$
Lower bound [20]	$\min \{\sqrt{SAH^3K}, HK\}$	n/a	$\frac{SAH^3}{\varepsilon^2}$

The third column reflects the burn-in cost, and the sample complexity (or PAC bound) refers to the number of episodes needed to yield ε accuracy. The results provided here account for all $K \geq 1$ or all $\varepsilon \in (0, H]$. Our article is the only one that gives regret (respectively PAC) bound matching the minimax lower bound for the entire range of K (respectively ε).

work Li et al. [43] developed a perturbed model-based approach that is provably optimal without incurring any burn-in cost. Analogous results have been obtained in Li et al. [47] for offline RL – a setting that requires policy learning to be performed based on historical data – unveiling the full-range optimality of a pessimistic model-based algorithm.

Unfortunately, the algorithmic and analysis frameworks developed in the above two works fail to accommodate the online counterpart. The main hurdle stems from the complicated statistical dependency intrinsic to episodic online RL; for instance, in online RL, the empirical transition probabilities and the running estimates of the value function are oftentimes statistically dependent in an intertwined manner (unless we waste data). How to decouple the intricate statistical dependency without compromising data efficiency constitutes the key innovation of this work. More precise, in-depth technical discussions will be provided in Section 4.

1.2 A Peek at Our Main Contributions

We are now positioned to summarize the main findings of this article. Focusing on time-inhomogeneous finite-horizon MDPs, our main contributions can be divided into two parts: the first part fully settles the minimax-optimal regret and sample complexity of online RL, whereas the second part further extends and augments our theory to make apparent the impacts of certain problem-dependent quantities. Throughout this subsection, the regret metric $\text{Regret}(K)$ captures the cumulative sub-optimality gap (i.e., the gap between the performance of the policy iterates and that of the optimal policy) over all K episodes, to be formally defined in (17).

1.2.1 Settling the Optimal Sample Complexity with No Burn-in Cost. Our first result *fully* determines the sample complexity limit of online RL in a minimax sense, allowing one to attain the optimal regret regardless of the number K of episodes that can be collected.

THEOREM 1.1. *For any $K \geq 1$ and any $0 < \delta < 1$, there exists an algorithm (see Algorithm 1) obeying*

$$\text{Regret}(K) \lesssim \min \left\{ \sqrt{SAH^3 K \log^5 \frac{SAHK}{\delta}}, HK \right\} \quad (8)$$

with probability at least $1 - \delta$.

The optimality of our regret bound (8) can be readily seen given that it matches the minimax lower bound (1) (modulo some logarithmic factor). One can also easily translate the above regret bound into sample complexity or **probably approximately correct (PAC)** bounds: the proposed algorithm is able to return an ε -suboptimal policy with high probability using at most

$$(\text{sample complexity}) \quad \tilde{O} \left(\frac{SAH^3}{\varepsilon^2} \right) \quad \text{episodes}, \quad (9)$$

(or equivalently, $\tilde{O} \left(\frac{SAH^4}{\varepsilon^2} \right)$ sample transitions as each episode has length H). Remarkably, this result holds true for the entire ε range (i.e., any $\varepsilon \in (0, H]$), essentially eliminating the need of any burn-in cost. It is noteworthy that even in the presence of an idealistic generative model, this order of sample size is un-improvable [5, 43].

The algorithm proposed herein is a modified version of *MVP: Monotonic Value Propagation*. Originally proposed by Zhang et al. [82], the *MVP* method falls under the category of model-based approaches, a family of algorithms that construct explicit estimates of the probability transition kernel before value estimation and policy learning. Notably, a technical obstacle that obstructs the progress in understanding model-based algorithms arises from the exceedingly large model dimensionality: given that the dimension of the transition kernel scales proportionally with S^2 , all existing analyses for model-based online RL fell short of effectiveness unless the sample size already far exceeds S^2 [6, 82]. To overcome this undesirable source of burn-in cost, a crucial step is to empower the analysis framework in order to accommodate the highly sub-sampled regime (i.e., a regime where the sample size scales linearly with S), which we shall elaborate on in Section 4. The full proof of Theorem 1.1 will be provided in Section 5.

1.2.2 Extension: Optimal Problem-dependent Regret Bounds. In practice, RL algorithms often perform far more appealingly than what their worst-case performance guarantees would suggest. This motivates a recent line of works that investigate optimal performance in a more problem-dependent fashion [19, 25, 30, 62, 66, 68, 69, 77, 79, 81, 86, 87]. Encouragingly, the proposed algorithm automatically achieves optimality on a more refined problem-dependent level, without requiring prior knowledge of additional problem-specific knowledge. This results in several extended theorems that take into account different problem-dependent quantities.

The first extension below investigates how the optimal value influences the regret bound.

THEOREM 1.2 (OPTIMAL VALUE-DEPENDENT REGRET). *For any $K \geq 1$, Algorithm 1 satisfies*

$$\mathbb{E}[\text{Regret}(K)] \lesssim \min \left\{ \sqrt{SAH^2 Kv^*}, Kv^* \right\} \log^5(SAHK), \quad (10)$$

where v^ is the value of the optimal policy averaged over the initial state distribution (to be formally defined in (43)).*

Moreover, there is also no shortage of applications where the use of a cost function is preferred over a value function [2, 4, 37, 72]. For this purpose, we provide another variation based upon the optimal cost.

THEOREM 1.3 (OPTIMAL COST-DEPENDENT REGRET). *For any $K \geq 1$ and any $0 < \delta < 1$, Algorithm 1 achieves*

$$\text{Regret}(K) \leq \tilde{O} \left(\min \left\{ \sqrt{SAH^2 K c^*} + SAH^2, K(H - c^*) \right\} \right), \quad (11)$$

with probability exceeding $1 - \delta$, where c^ denotes the cost of the optimal policy averaged over the initial state distribution (to be formally defined in (45)).*

It is worth noting that: despite the apparent similarity between the statements of Theorem 1.2 and Theorem 1.3, they do not imply each other, although their proofs are very similar to each other.

Finally, we establish another regret bound that reflects the effect of certain variance quantities of interest.

THEOREM 1.4 (OPTIMAL VARIANCE-DEPENDENT REGRET). *For any $K \geq 1$ and any $0 < \delta < 1$, Algorithm 1 obeys*

$$\text{Regret}(K) \leq \tilde{O} \left(\min \left\{ \sqrt{SAHK\text{var}} + SAH^2, KH \right\} \right), \quad (12)$$

with probability at least $1 - \delta$, where var is a quantity defined in (49).

As we shall see later in (49), var is a variance-type metric. Two remarks concerning the above extensions are in order:

- In the worst-case scenarios, the quantities v^* , c^* and var can all be as large as the order of H , in which case Theorems 1.2–1.4 readily recover Theorem 1.1. In contrast, the advantages of Theorems 1.2–1.4 over Theorem 1.1 become more evident in those favorable cases (e.g., the situation where $v^* \ll H$ or $c^* \ll H$, or the case when the environment is nearly deterministic (so that $\text{var} \approx 0$)).
- Interestingly, the regret bounds in Theorems 1.2–1.4 all contain a lower-order term SAH^2 , and one might naturally wonder whether this term is essential. To demonstrate the unavoidable nature of this term and hence the optimality of Theorems 1.2–1.4, we will provide matching lower bounds, to be detailed in Section 6.

1.3 Related Works

Let us take a moment to discuss several related theoretical works on tabular RL. Note that there has also been an active line of research that exploits low-dimensional function approximation to further reduce sample complexity, which is beyond the scope of this article.

Our discussion below focuses on two mainstream approaches that have received widespread adoption: the model-based approach and the model-free approach. In a nutshell, model-based algorithms decouple model estimation and policy learning, and often use the learned transition kernel to compute the value function and find a desired policy. In stark contrast, the model-free approach attempts to estimate the optimal value function and optimal policy directly without explicit estimation of the model. In general, model-free algorithms only require $O(SAH)$ memory — needed when storing the running estimates for Q-functions and value functions — while the model-based counterpart might require $\Omega(S^2AH)$ space in order to store the estimated transition kernel.

Sample complexity for RL with a simulator. As an idealistic setting that separates the consideration of exploration from that of estimation, RL with a simulator (or generative model) has been studied by numerous works, allowing the learner to draw independent samples for any state-action pairs [1, 5, 9, 14, 15, 23, 32, 33, 39, 40, 43, 53, 59–61, 70, 71]. While both model-based and model-free approaches are capable of achieving asymptotic sample optimality [1, 5, 61, 71], all model-free algorithms that enjoy asymptotically optimal sample complexity suffer from dramatic burn-in cost.

Thus far, only the model-based approach has been shown to fully eliminate the burn-in cost for both discounted infinite-horizon MDPs and inhomogeneous finite-horizon MDPs [43]. The full-range optimal sample complexity for time-homogeneous finite-horizon MDPs in the presence of a simulator remains open.

Sample complexity for offline RL. The subfield of offline RL is concerned with learning based purely on a pre-collected dataset [38]. A frequently used mathematical model assumes that historical data are collected (often independently) using some behavior policy, and the key challenges (compared with RL with a simulator) come from distribution shift and incomplete data coverage. The sample complexity of offline RL has been the focus of a large strand of recent works, with asymptotically optimal sample complexity achieved by multiple algorithms [31, 41, 44, 54–56, 58, 74, 75, 78, 80]. Akin to the simulator setting, the fully optimal sample complexity (without burn-in cost) has only been achieved via the model-based approach when it comes to discounted infinite-horizon and inhomogeneous finite-horizon settings [41]. All asymptotically optimal model-free methods incur substantial burn-in cost. The case with time-homogeneous finite-horizon MDPs also remains unsettled.

Sample complexity for online RL. Obtaining optimal sample complexity (or regret bound) in online RL without incurring any burn-in cost has been one of the most fundamental open problems in RL theory. In fact, the past decades have witnessed a flurry of activity toward improving the sample efficiency of online RL, partial examples including [3, 8, 11, 12, 16, 17, 20–22, 25–27, 29, 32, 34–36, 42, 45–47, 49–52, 57, 63–65, 67, 73, 76, 81–84]. Unfortunately, no work has been able to conquer this problem completely: the state-of-the-art result for model-based algorithms still incurs a burn-in that scales at least quadratically in S [82], while the burn-in cost of the best model-free algorithms (particularly with the aid of variance reduction introduced in Zhang et al. [84]) still suffers from highly sub-optimal horizon dependency [42].

1.4 Notation

Before proceeding, let us introduce a set of notation to be used throughout. Let 1 and 0 indicate respectively the all-one vector and the all-zero vector. Let e_s denote the s th standard basis vector (which has 1 at the s th coordinate and 0 otherwise). For any set \mathcal{X} , $\Delta(\mathcal{X})$ represents the set of probability distributions over the set \mathcal{X} . For any positive integer N , we denote $[N] = \{1, \dots, N\}$. For any two vectors x, y with the same dimension, we use $\langle x, y \rangle$ (or $x^\top y$) to denote the inner product of x and y . For any integer $S > 0$, any probability vector $p \in \Delta([S])$ and another vector $v = [v_i]_{1 \leq i \leq S}$, we denote by

$$\mathbb{V}(p, v) := \langle p, v^2 \rangle - (\langle p, v \rangle)^2 = \langle p, (v - \langle p, v \rangle 1)^2 \rangle, \quad (13)$$

the associated variance, where $v^2 = [v_i^2]_{1 \leq i \leq S}$ represents element-wise square of v . For any two vectors $a = [a_i]_{1 \leq i \leq n}$ and $b = [b_i]_{1 \leq i \leq n}$, the notation $a \geq b$ (respectively $a \leq b$) means $a_i \geq b_i$ (respectively $a_i \leq b_i$) holds simultaneously for all i . Without loss of generality, we assume throughout that K is a power of 2 to streamline presentation.

2 Problem Formulation

In this section, we introduce the basics of tabular online RL, as well as some basic assumptions to be imposed throughout.

Basics of finite-horizon MDPs. This article concentrates on time-inhomogeneous (or nonstationary) finite-horizon MDPs. Throughout the article, we employ $\mathcal{S} = \{1, \dots, S\}$ to denote the state space, $\mathcal{A} = \{1, \dots, A\}$ the action space, and H the planning horizon. The notation $P = \{P_h : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})\}_{1 \leq h \leq H}$ denotes the probability transition kernel of the MDP; for any current state s

at any step h , if action a is taken, then the state at the next step $h+1$ of the environment is randomly drawn from $P_{s,a,h} := P_h(\cdot | s, a) \in \Delta(\mathcal{S})$. Also, the notation $R = \{R_{s,a,h} \in \Delta([0, H])\}_{1 \leq h \leq H, s \in \mathcal{S}, a \in \mathcal{A}}$ indicates the reward distribution; that is, while executing action a in state s at step h , the agent receives an immediate reward — which is non-negative and possibly stochastic — drawn from the distribution $R_{s,a,h}$. We shall also denote by $r = \{r_h(s, a)\}_{1 \leq h \leq H, s \in \mathcal{S}, a \in \mathcal{A}}$ the mean reward function, so that $r_h(s, a) := \mathbb{E}_{r' \sim R_{s,a,h}}[r'] \in [0, H]$ for any (s, a, h) -tuple. Additionally, a deterministic policy $\pi = \{\pi_h : \mathcal{S} \rightarrow \mathcal{A}\}_{1 \leq h \leq H}$ stands for an action selection rule, so that the action selected in state s at step h is given by $\pi_h(s)$. The readers can consult standard textbooks (e.g., Bertsekas [10]) for more extensive descriptions.

In each episode, a trajectory $(s_1, a_1, r'_1, s_2, \dots, s_H, a_H, r'_H)$ is rolled out as follows: the learner starts from an initial state s_1 independently drawn from some fixed (but unknown) distribution $\mu \in \Delta(\mathcal{S})$; for each step $1 \leq h \leq H$, the learner takes action a_h , gains an immediate reward $r'_h \sim R_{s_h, a_h, h}$, and the environment transits to the state s_{h+1} at step $h+1$ according to $P_{s_h, a_h, h}$. Note that both the reward and the state transition are independently drawn from their respective distributions, depending solely on the current state-action-step triple but not any previous outcomes. All of our results in this article operate under the following assumption on the total reward.

ASSUMPTION 1. *For any possible trajectory $(s_1, a_1, r'_1, \dots, s_H, a_H, r'_H)$, one always has $0 \leq \sum_{h=1}^H r'_h \leq H$.*

As can be easily seen, Assumption 1 is less stringent than another common choice that assumes $r'_h \in [0, 1]$ for any h in any episode. In particular, Assumption 1 allows for sparse and spiky rewards along an episode; more discussions can be found in [28, 73].

Value function and Q-function. For any given policy π , one can define the value function $V^\pi = \{V_h^\pi : \mathcal{S} \rightarrow \mathbb{R}\}$ and the Q-function $Q^\pi = \{Q_h^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}\}$ such that

$$V_h^\pi(s) := \mathbb{E}_\pi \left[\sum_{j=h}^H r'_j \middle| s_h = s \right], \quad \forall (s, h) \in \mathcal{S} \times [H], \quad (14a)$$

$$Q_h^\pi(s, a) := \mathbb{E}_\pi \left[\sum_{j=h}^H r'_j \middle| (s_h, a_h) = (s, a) \right], \quad \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H], \quad (14b)$$

where the expectation $\mathbb{E}_\pi[\cdot]$ is taken over the randomness of an episode $\{(s_h, a_h, r'_h)\}_{1 \leq h \leq H}$ generated under policy π , that is, $a_j = \pi_j(s_j)$ for every $h \leq j \leq H$ (respectively $h < j \leq H$) is chosen in the definition of V_h^π (respectively Q_h^π). Accordingly, we define the optimal value function and the optimal Q-function respectively as

$$V_h^*(s) := \max_\pi V_h^\pi(s), \quad \forall (s, h) \in \mathcal{S} \times [H], \quad (15a)$$

$$Q_h^*(s, a) := \max_\pi Q_h^\pi(s, a) \quad \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]. \quad (15b)$$

Throughout this article, we shall often abuse the notation by letting both V_h^π and V_h^* (respectively Q_h^π and Q_h^*) represent S -dimensional (respectively SA -dimensional) vectors containing all elements of the corresponding value functions (respectively Q-functions). Two important properties are worth mentioning: (a) the optimal value and the optimal Q-function are linked by the Bellman equation:

$$Q_h^*(s, a) = r_h(s, a) + \langle P_{h,s,a}, V_{h+1}^* \rangle, \quad V_h^*(s) = \max_{a'} Q_h^*(s, a'), \quad \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]; \quad (16)$$

(b) there exists a deterministic policy, denoted by π^* , that achieves optimal value functions and Q-functions for all state-action-step tuples simultaneously, that is

$$V_h^{\pi^*}(s) = V_h^*(s) \quad \text{and} \quad Q_h^{\pi^*}(s, a) = Q_h^*(s, a), \quad \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H].$$

Data collection protocol and performance metrics. During the learning process, the learner is allowed to collect K episodes of samples (using arbitrary policies it selects). More precisely, in the k th episode, the learner is given an independently generated initial state $s_1^k \sim \mu$, and executes policy π^k (chosen based on data collected in previous episodes) to obtain a sample trajectory $\{(s_h^k, a_h^k, r_h^k)\}_{1 \leq h \leq H}$, with s_h^k , a_h^k and r_h^k denoting the state, action, and immediate reward at step h of this episode.

To evaluate the learning performance, a widely used metric is the (cumulative) regret over all K episodes:

$$\text{Regret}(K) := \sum_{k=1}^K \left(V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k) \right), \quad (17)$$

and our goal is to design an online RL algorithm that minimizes $\text{Regret}(K)$ regardless of the allowable sample size K . It is also well-known (see, e.g., Jin et al. [29]) that a regret bound can often be readily translated into a PAC sample complexity result, the latter of which counts the number of episodes needed to find an ε -optimal policy $\widehat{\pi}$ in the sense that $\mathbb{E}_{s_1 \sim \mu} [V_1^*(s_1) - V_1^{\widehat{\pi}}(s_1)] \leq \varepsilon$ (note that the expectation here is taken not only over the randomness of the initial state, but also over the randomness in selecting $\widehat{\pi}$). For instance, the reduction argument in Jin et al. [29] reveals that: if an algorithm achieves $\text{Regret}(K) \leq f(S, A, H)K^{1-\alpha}$ for some function f and some parameter $\alpha \in (0, 1)$, then by randomly selecting a policy from $\{\pi^k\}_{1 \leq k \leq K}$ as $\widehat{\pi}$ one achieves $\mathbb{E}_{s_1 \sim \mu} [V_1^*(s_1) - V_1^{\widehat{\pi}}(s_1)] \lesssim f(S, A, H)K^{-\alpha}$, thus resulting in a sample complexity bound of $\left(\frac{f(S, A, H)}{\varepsilon}\right)^{1/\alpha}$.

3 A Model-based Algorithm: Monotonic Value Propagation

In this section, we formally describe our algorithm: a simple variation of the model-based algorithm called *Monotonic Value Propagation* proposed by Zhang et al. [82]. We present the full procedure in Algorithm 1, and point out several key ingredients.

— *Optimistic updates using upper confidence bounds (UCB)*. The algorithm implements the optimism principle in the face of uncertainty by adopting the frequently used UCB-based framework (see, e.g., UCBVI by Azar et al. [6]). More specifically, the learner calculates the optimistic Bellman equation backward (from $h = H, \dots, 1$): it first computes an empirical estimate $\widehat{P} = \{\widehat{P}_h \in \mathbb{R}^{S \times A \times S}\}_{1 \leq h \leq H}$ of the transition probability kernel as well as an empirical estimate $\widehat{r} = \{\widehat{r}_h \in \mathbb{R}^{S \times A}\}_{1 \leq h \leq H}$ of the mean reward function, and then maintains upper estimates for the associated value function and Q-function using

$$Q_h(s, a) \leftarrow \min \{\widehat{r}_h(s, a) + \langle \widehat{P}_{s, a, h}, V_{h+1} \rangle + b_h(s, a), H\}, \quad (18a)$$

$$V_h(s) \leftarrow \max_a Q_h(s, a), \quad (18b)$$

for all state-action pairs. Here, Q_h (respectively V_h) indicates the running estimate for the Q-function (respectively value function), whereas $b_h(s, a) \geq 0$ is some suitably chosen bonus term that compensates for the uncertainty. The above opportunistic Q-estimate in turn allows one to obtain a policy estimate (via a simple greedy rule), which will then be executed to collect new data. The fact that we first estimate the model (i.e., the transition kernel and mean rewards) makes it a model-based approach. Noteworthily, the empirical model $(\widehat{P}, \widehat{r})$

ALGORITHM 1: Monotonic Value Propagation (MVP) [82]

1 **input:** state space \mathcal{S} , action space \mathcal{A} , horizon H , total number of episodes K , confidence parameter δ , $c_1 = \frac{460}{9}$, $c_2 = 2\sqrt{2}$, $c_3 = \frac{544}{9}$.

2 **initialization:** set $\delta' \leftarrow \frac{\delta}{200S\mathcal{A}H^2K^2}$, and for all $(s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$, set $\theta_h(s, a) \leftarrow 0$, $\kappa_h(s, a) \leftarrow 0$, $N_h^{\text{all}}(s, a, s') \leftarrow 0$, $N_h(s, a, s') \leftarrow 0$, $N_h(s, a) \leftarrow 0$, $Q_h(s, a) \leftarrow H$, $V_h(s) \leftarrow H$.

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

4 Set π^k such that $\pi_h^k(s) = \arg \max_a Q_h(s, a)$ for all $s \in \mathcal{S}$ and $h \in [H]$. /* policy update. */

5 **for** $h = 1, 2, \dots, H$ **do**

6 Observe s_h^k , take action $a_h^k = \arg \max_a Q_h(s_h^k, a)$, receive r_h^k , observe s_{h+1}^k . /* sampling. */

7 $(s, a, s') \leftarrow (s_h^k, a_h^k, s_{h+1}^k)$.

8 Update $N_h^{\text{all}}(s, a) \leftarrow N_h^{\text{all}}(s, a) + 1$, $N_h(s, a, s') \leftarrow N_h(s, a, s') + 1$, $\theta_h(s, a) \leftarrow \theta_h(s, a) + r_h^k$,

9 $\kappa_h(s, a) \leftarrow \kappa_h(s, a) + (r_h^k)^2$.

10 /* perform updates using data of this epoch.

11 **if** $N_h^{\text{all}}(s, a) \in \{1, 2, \dots, 2^{\log_2 K}\}$ **then**

12 $N_h(s, a) \leftarrow \sum_{\tilde{s}} N_h(s, a, \tilde{s})$. // number of visits to (s, a, h) in this epoch.

13 $\hat{r}_h(s, a) \leftarrow \frac{\theta_h(s, a)}{N_h(s, a)}$. // empirical rewards of this epoch.

14 $\hat{\sigma}_h(s, a) \leftarrow \frac{\kappa_h(s, a)}{N_h(s, a)}$. // empirical squared rewards of this epoch.

15 $\hat{P}_{s, a, h}(\tilde{s}) \leftarrow \frac{N_h(s, a, \tilde{s})}{N_h(s, a)}$ for all $\tilde{s} \in \mathcal{S}$. // empirical transition for this epoch.

16 Set TRIGGERED = TRUE, and $\theta_h(s, a) \leftarrow 0$, $\kappa_h(s, a) \leftarrow 0$, $N_h(s, a, \tilde{s}) \leftarrow 0$ for all $\tilde{s} \in \mathcal{S}$.

17 /* optimistic Q-estimation using empirical model of this epoch. */

18 **if** TRIGGERED = TRUE **then**

19 Set TRIGGERED = FALSE, and $V_{H+1}(s) \leftarrow 0$ for all $s \in \mathcal{S}$.

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

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

22 $b_h(s, a) \leftarrow c_1 \sqrt{\frac{\mathbb{V}(\hat{P}_{s, a, h}, V_{h+1}) \log \frac{1}{\delta'}}{\max\{N_h(s, a), 1\}}} + c_2 \sqrt{\frac{(\hat{\sigma}_h(s, a) - (\hat{r}_h(s, a))^2) \log \frac{1}{\delta'}}{\max\{N_h(s, a), 1\}}} + c_3 \frac{H \log \frac{1}{\delta'}}{\max\{N_h(s, a), 1\}}$, (19)

23 $Q_h(s, a) \leftarrow \min \{\hat{r}_h(s, a) + \langle \hat{P}_{s, a, h}, V_{h+1} \rangle + b_h(s, a), H\}$, $V_h(s) \leftarrow \max_a Q_h(s, a)$. (20)

shall be updated multiple times as new samples continue to arrive, and hence the updating rule (18) will be invoked multiple times as well.

— *An epoch-based procedure and a doubling trick.* Compared to the original UCBVI [6], one distinguishing feature of MVP is to update the empirical transition kernel and empirical rewards in an epoch-based fashion, as motivated by a doubling update framework adopted in Jaksch et al. [26]. More concretely, the whole learning process is divided into consecutive epochs via a simple doubling rule; namely, whenever there exists a (s, a, h) -tuple whose visitation count reaches a power of 2, we end the current epoch, reconstruct the empirical model (cf. lines 11 and 13 of Algorithm 1), compute the Q-function and value function using the newly updated transition kernel and rewards (cf. (20)), and then start a new epoch with an updated sampling policy. This stands in stark contrast with the original UCBVI, which computes new

estimates for the transition model, Q-function and value function in every episode. With this doubling rule in place, the estimated transition probability vector for each (s, a, h) -tuple will be updated by no more than $\log_2 K$ times, a feature that plays a pivotal role in significantly reducing some sort of covering number needed in our covering-based analysis (as we shall elaborate on shortly in Section 4). In each epoch, the learned policy is induced by the optimistic Q-function estimate – computed based on the empirical transition kernel of the *current* epoch – which will then be employed to collect samples in *all* episodes of the next epoch. More technical explanations of the doubling update rule will be provided in Section 4.2.

- *Monotonic bonus functions.* Another crucial step in order to ensure near-optimal regret lies in careful designs of the data-driven bonus terms $\{b_h(s, a)\}$ in (18a). Here, we adopt the monotonic Bernstein-style bonus function for MVP originally proposed in Zhang et al. [82], to be made precise in (19). Compared to the bonus function in Euler [81] and UCBVI [6], the monotonic bonus form has a cleaner structure that effectively avoids large lower-order terms. Note that in order to enable variance-aware regret, we also need to keep track of the empirical variance of the (stochastic) immediate rewards. For more details about “monotonicity” and bonus functions, we refer the readers to Appendix C.1.

Remark 1. We note that a doubling update rule has also been used in the original MVP [82]. A subtle difference between our modified version and the original one lies in that: when the visitation count for some (s, a, h) reaches 2^i for some integer $i \geq 1$, we only use the second half of the samples (i.e., the $\{2^{i-1} + l\}_{l=1}^{2^i-1}$ -th samples) to compute the empirical model, whereas the original MVP makes use of all the 2^i samples. This modified step turns out to be helpful in our analysis, while still preserving sample efficiency in an otherwise sense (since the latest batch always contains at least half of the samples).

4 Key Technical Innovations

In this section, we point out the key technical hurdles the previous approach encounters when mitigating the burn-in cost, and put forward a new strategy to overcome such hurdles. For ease of presentation, let us introduce a set of augmented notation to indicate several running iterates in Algorithm 1, which makes clear the dependency on the episode number k and will be used throughout all of our analysis.

- $\widehat{P}_{s, a, h}^k \in \mathbb{R}^S$: the latest update of the empirical transition probability vector $\widehat{P}_{s, a, h}$ before the k th episode.
- $\widehat{r}_h^k(s, a) \in [0, H]$: the latest update of the empirical reward $\widehat{r}_h(s, a)$ before the k th episode.
- $\widehat{\sigma}_h^k(s, a) \in [0, H^2]$: the latest update of the empirical squared reward $\widehat{\sigma}_h(s, a)$ before the k th episode.
- $b_h^k(s, a) \geq 0$: the latest update of the bonus term $b_h(s, a)$ before the k th episode.
- $N_h^{k, \text{all}}(s, a)$: the total visitation count of the (s, a, h) -tuple before the beginning of the k th episode.
- $N_h^k(s, a)$: the visitation count $N_h(s, a)$ of the (s, a, h) -tuple of the latest doubling batch used to compute $\widehat{P}_{s, a, h}$ before the k th episode. When $N_h^{k, \text{all}}(s, a) = 0$, we define $N_h^k(s, a) = 1$ for ease of presentation.
- $V_h^k \in \mathbb{R}^S$: the value function estimate V_h before the beginning of the k th episode.
- $Q_h^k \in \mathbb{R}^{SA}$: the Q-function estimate Q_h before the beginning of the k th episode.

Another notation for the empirical transition probability vector is also introduced below:

- For any $j \geq 2$ (respectively $j = 1$), let $\widehat{P}_{s,a,h}^{(j)}$ be the empirical transition probability vector for (s, a, h) computed using the j th batch of data, i.e., the $\{2^{j-2} + i\}_{i=1}^{2^{j-2}}$ -th samples (respectively the 1st sample) for (s, a, h) . For completeness, we take $\widehat{P}_{s,a,h}^{(0)} = \frac{1}{S} \mathbf{1}$ for the 0th batch.
- Similarly, let $\widehat{r}_h^{(j)}(s, a)$ (respectively $\widehat{\sigma}_h^{(j)}(s, a)$) denote the empirical reward (respectively empirical squared reward) w.r.t. (s, a, h) based on the j th batch of data.

4.1 Technical Barriers in Prior Theory for UCBVI

Let us take a close inspection on prior regret analysis for UCB-based model-based algorithms, in order to illuminate the part that calls for novel analysis. To simplify presentation, this subsection assumes deterministic rewards so that each empirical reward is replaced by its mean.

Let us look at the original UCBVI algorithm proposed by Azar et al. [6]. Standard decomposition arguments employed in the literature (e.g., Azar et al. [6], Jaksch et al. [26], Zhang et al. [82]) decompose the regret as follows:

$$\begin{aligned} \text{Regret}(K) \leq & \sum_{k,h} \left(\widehat{P}_{s_h^k, a_h^k, h}^{k,\text{all}} - P_{s_h^k, a_h^k, h} \right) V_{h+1}^k + \sum_{k,h} b_h^k(s_h^k, a_h^k) \\ & + \sum_{k,h} \left(P_{s_h^k, a_h^k, h} - e_{s_h^k, h} \right) (V_{h+1}^k - V_{h+1}^{\pi^k}); \end{aligned} \quad (21)$$

see also the derivation in Section 5. Here, we abuse the notation by letting V_{h+1}^k (respectively b_h^k) be the value function estimate (respectively bonus term) of UCBVI before the k th episode, and in the meantime, we let $\widehat{P}_{s,a,h}^{k,\text{all}}$ represent the empirical transition probability for the (s, a, h) -tuple computed using *all* samples before the k th episode (note that we add the superscript all to differentiate it from its counterpart in our algorithm). In order to achieve full-range optimal regret, one needs to bound the three terms on the right-hand side of (21) carefully, among which two are easy to handle.

- It is known that the second term (i.e., the aggregate bonus) on the right-hand side of (21) can be controlled in a rate-optimal manner if we adopt a suitably chosen Bernstein-style bonus; see, e.g., Zhang et al. [82], which will also be made clear shortly in Section 5.
- In the meantime, the third term on the right-hand side of (21) can be easily coped with by means of standard martingale concentration bounds (e.g., the Freedman inequality).

It then comes down to controlling the first term on the right-hand side of (21). This turns out to be the most challenging part, owing to the complicated statistical dependency between $\widehat{P}_{s_h^k, a_h^k, h}^{k,\text{all}}$ and V_{h+1}^k . To see this, note that $\widehat{P}_{s,a,h}^{k,\text{all}}$ is constructed based on *all* previous samples of (s, a, h) , which has non-negligible influences upon V_{h+1}^k as V_{h+1}^k is computed based on previous samples. At least two strategies have been proposed to circumvent this technical difficulty, which we take a moment to discuss.

- *Strategy 1: replacing V_{h+1}^k with V_{h+1}^{\star} for large k .* Most prior analysis for model-based algorithms [6, 17, 81, 82] decomposes

$$\begin{aligned} & \sum_{k,h} \left(\widehat{P}_{s_h^k, a_h^k, h}^{k,\text{all}} - P_{s_h^k, a_h^k, h} \right) V_{h+1}^k \\ & = \sum_{k,h} \left(\widehat{P}_{s_h^k, a_h^k, h}^{k,\text{all}} - P_{s_h^k, a_h^k, h} \right) V_{h+1}^{\star} + \sum_{k,h} \left(\widehat{P}_{s_h^k, a_h^k, h}^{k,\text{all}} - P_{s_h^k, a_h^k, h} \right) (V_{h+1}^k - V_{h+1}^{\star}). \end{aligned} \quad (22)$$

The rationale behind this decomposition is as follows:

- (i) given that V_{h+1}^* is fixed and independent from the data, the first term on the right-hand side of (22) can be bounded easily using Freedman's inequality;
- (ii) the second term on the right-hand side of (22) would vanish as V_{h+1}^k and V_{h+1}^* become exceedingly close (which would happen as k becomes large enough).

Such arguments, however, fall short of tightness when analyzing the initial stage of the learning process: given that $V_{h+1}^k - V_{h+1}^*$ cannot be sufficiently small at the beginning, this approach necessarily results in a huge burn-in cost.

— *Strategy 2: a covering-based argument.* Let us discuss informally another potential strategy that motivates our analysis. We first take a closer look at the relationship between $\widehat{P}_{s,a,h}^{k,\text{all}}$ and V_{h+1}^k . Abusing notation by letting $N_h^{k,\text{all}}(s, a)$ be the total number of visits to a (s, a, h) -tuple before the k th episode in UCBVI, we can easily observe that $\widehat{P}_{s,a,h}^{k,\text{all}}$ and V_{h+1}^k are statistically independent conditioned on the set $\{N_h^{k,\text{all}}(s, a)\}_{(s, a, k) \in \mathcal{S} \times \mathcal{A} \times [K]}$. Consequently, if we “pretend” that $\{N_h^{k,\text{all}}(s, a)\}$ are pre-fixed and independent of $\{\widehat{P}_{s,a,h}^{k,\text{all}}\}$, then one can invoke standard concentration inequalities to obtain a high-probability bound on $\sum_{k,h} (\widehat{P}_{s_h^k, a_h^k, h}^{k,\text{all}} - P_{s_h^k, a_h^k, h}) V_{h+1}^k$ in a desired manner. The next step would then be to invoke a union bound over all possible configurations of $\{N_h^{k,\text{all}}(s, a)\}$, so as to eliminate the above independence assumption. The main drawback of this approach, however, is that there are exponentially many (e.g., in K) possible choices of $\{N_h^{k,\text{all}}(s, a)\}$, inevitably loosening the regret bound.

4.2 Our Approach

In light of the covering-based argument in Section 4.1, one can only hope that this analysis strategy to work if substantial compression (i.e., a significantly reduced covering number) of the visitation counts is plausible. This motivates our introduction of the doubling batches as described in Section 3, so that for each (s, a, h) -tuple, the empirical model $\widehat{P}_{s,a,h}$ and its associated visitation count $N_h(s, a)$ (for the associated batch) are updated at most $\log_2 K$ times (see line 9 of Algorithm 1). Compared to the original UCBVI that recomputes the transition model in every episode, our algorithm allows for significant reduction of the covering number of the visitation counts, thanks to its much less frequent updates.

Similar to (21), we are in need of bounding the following term when analyzing Algorithm 1:

$$\sum_{k,h} \left(\widehat{P}_{s_h^k, a_h^k, h}^{k,\text{all}} - P_{s_h^k, a_h^k, h} \right) V_{h+1}^k. \quad (23)$$

In what follows, we present our key ideas that enable tight analysis of this quantity, which constitute our main technical innovations. The complete regret analysis for Algorithm 1 is postponed to Section 5.

4.2.1 *Key Concept: Profiles.* One of the most important concepts underlying our analysis for Algorithm 1 is the so-called “profile”, defined below.

Definition 4.1 (Profile). Consider any combination $\{N_h^{k,\text{all}}(s, a)\}_{(s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times [H] \times [K]}$. For any $k \in [K]$, define

$$\forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H] : \quad I_{s,a,h}^k := \begin{cases} \max \{j \in \mathbb{N} : 2^{j-1} \leq N_h^{k,\text{all}}(s, a)\}, & \text{if } N_h^{k,\text{all}}(s, a) > 0; \\ 0, & \text{if } N_h^{k,\text{all}}(s, a) = 0. \end{cases} \quad (24a)$$

The profile for the k th episode ($1 \leq k \leq K$) and the total profile are then defined respectively as

$$\mathcal{I}^k := \{I_{s,a,h}^k\}_{(s,a,h) \in \mathcal{S} \times \mathcal{A} \times [H]}, \quad (24b)$$

$$\text{and} \quad \mathcal{I} := \{\mathcal{I}^k\}_{k=1}^K. \quad (24c)$$

Clearly, once a total profile \mathcal{I} w.r.t. $\{N_h^{k,\text{all}}(s, a)\}$ is given, one can write

$$\widehat{P}_{s,a,h}^k = \widehat{P}_{s,a,h}^{(I_{s,a,h}^k)}, \quad \forall (s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times [H] \times [K]. \quad (25)$$

In other words, a total profile specifies all the time instances and locations when the empirical model is updated. Given that each $N_h^k(s, a)$ is recomputed only when the associated empirical model is updated (see line 10 of Algorithm 1), the total profile also provides a succinct representation of the set $\{N_h^k(s, a)\}$.

In order to quantify the degree of compression, Definition 4.1 offers when representing the update times and locations, we provide an upper bound on the number of possible total profiles in the lemma below.

LEMMA 4.2. *Suppose that $K \geq SAH \log_2 K$. Then the number of all possible total profiles w.r.t. Algorithm 1 is at most*

$$(4SAHK)^{SAH \log_2 K + 1}.$$

PROOF. Define the following set (which will be useful in subsequent analysis as well)

$$C := \left\{ \mathcal{I} = \{\mathcal{I}^1, \dots, \mathcal{I}^K\} \mid \mathcal{I}^1 \leq \mathcal{I}^2 \leq \dots \leq \mathcal{I}^K, \mathcal{I}^k \in \{0, 1, \dots, \log_2 K\}^{SAH} \text{ for all } 1 \leq k \leq K \right\}. \quad (26)$$

Due to the monotonicity constraints, it is easily seen that the total profile of any set $\{N_h^k(s, a)\}$ must lie within C . It then boils down to proving that $|C| \leq (4SAHK)^{SAH \log_2 K + 1}$, which can be accomplished via elementary combinatorial calculations. The complete proof is deferred to Appendix B.1. \square

In comparison to using $\{N_h^{k,\text{all}}(s, a)\}$ to encode all update times and locations — which might have exponentially many (in K) possibilities — the use of doubling batches in Algorithm 1 allows for remarkable compression (as the exponent of the number of possibilities only scales logarithmically in K).

4.2.2 Decoupling the Statistical Dependency. In this subsection, we discuss our approach to decoupling statistical dependency across the trajectories.

An expanded view of randomness w.r.t. state transitions. To facilitate analysis, we find it helpful to look at a different yet closely related way to generate independent samples from a generative model.

Definition 4.3 (An Expanded Sample Set from a Generative Model). Let $\mathcal{D}^{\text{expand}}$ be a set of $SAHK$ independent samples generated as follows: for each $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$, draw K independent samples $(s, a, h, s'^{(i)})$ obeying $s'^{(i)} \stackrel{\text{ind.}}{\sim} P_{s,a,h}$ ($1 \leq i \leq K$).

Crucially, $\mathcal{D}^{\text{expand}}$ can be viewed as an expansion of the original dataset — denoted by $\mathcal{D}^{\text{original}}$ — collected in online learning, as we can couple the data collection processes of $\mathcal{D}^{\text{original}}$ and $\mathcal{D}^{\text{expand}}$ as follows:

- (i) generate $\mathcal{D}^{\text{expand}}$ before the beginning of the online process;
- (ii) during the online learning process, whenever a sample needs to be drawn from (s, a, h) , one can take an unused sample of (s, a, h) from $\mathcal{D}^{\text{expand}}$ without replacement.

Remarkably, this coupling allows one to conduct analysis alternatively based on the expanded sample set $\mathcal{D}^{\text{expand}}$, which is sometimes more convenient (as we shall detail momentarily). Unless otherwise noted, all analyses in our proof assume that $\mathcal{D}^{\text{original}}$ and $\mathcal{D}^{\text{expand}}$ are *coupled* through the above simulation process.

In the sequel, we let $\widehat{P}_{s,a,h}^{(j)}$ (cf. the beginning of Section 4) denote the empirical probability vector based on the j th batch of data (i.e., those samples associated with (s, a, h) indexed with $2^{j-2} + 1, \dots, 2^{j-1}$) from $\mathcal{D}^{\text{original}}$ and $\mathcal{D}^{\text{expand}}$ interchangeably, as long as it is clear from the context.

A starting point: a basic decomposition. We now describe our approach to tackling the complicated statistical dependency between $\widehat{P}_{s,a,h}^k$ and V_{h+1}^k . To begin with, take $\mathcal{I}^{\text{true}} = \{\mathcal{I}^{1,\text{true}}, \dots, \mathcal{I}^{K,\text{true}}\}$ with $\mathcal{I}^{k,\text{true}} = \{I_{s,a,h}^{k,\text{true}}\}$, which denotes the total profile w.r.t. the true visitation counts in the online learning process; let $k_{l,j,s,a,h}$ denote the episode index of the sample that visits (s, a, h) for the $(2^{l-1} + j)$ -th time in the online learning process; and we take $V_{h+1}^k = 0$ for any $k > K$. Then from relation (25) we can write

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \left\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \right\rangle &= \sum_{k=1}^K \sum_{h=1}^H \left\langle \widehat{P}_{s_h^k, a_h^k, h}^{(I_{s,a,h}^{k,\text{true}})} - P_{s_h^k, a_h^k, h}, V_{h+1}^k \right\rangle \\ &= \sum_{l=0}^{\log_2 K} \sum_{s,a,h} \left\langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, \sum_{k=1}^K \mathbb{1} \left\{ (s_h^k, a_h^k) = (s, a), I_{s,a,h}^{k,\text{true}} = l \right\} V_{h+1}^k \right\rangle \\ &\leq \sum_{l=1}^{\log_2 K} \sum_{s,a,h} \left\langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, \sum_{k=1}^K \mathbb{1} \left\{ (s_h^k, a_h^k) = (s, a), I_{s,a,h}^{k,\text{true}} = l \right\} V_{h+1}^k \right\rangle + SAH^2 \\ &= \sum_{l=1}^{\log_2 K} \sum_{j=1}^{2^{l-1}} \left\{ \sum_{s,a,h} \left\langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, V_{h+1}^{k_{l,j,s,a,h}} \right\rangle \right\} + SAH^2. \end{aligned} \quad (27)$$

Here, the third line makes use of the fact that $0 \leq V_{h+1}^k(s) \leq H$ for all $s \in \mathcal{S}$. The decomposition (27) motivates us to first control the term $\sum_{s,a,h} \langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, V_{h+1}^{k_{l,j,s,a,h}} \rangle$, leading to the following 3-step analysis strategy.

- (1) For any given total profile $\mathcal{I} \in \mathcal{C}$ and any fixed $1 \leq l \leq \log_2 K$, develop a high-probability bound on a weighted sum taking the following form

$$\sum_{s,a,h} \left(\widehat{P}_{s,a,h}^{(l)} - P_{s,a,h} \right) X_{h+1,s,a}, \quad (28)$$

where each vector $X_{h+1,s,a}$ is any deterministic function of \mathcal{I} and the samples collected for steps $h' \geq h + 1$. Given the statistical independence between $\widehat{P}_{s,a,h}^{(l)}$ and those samples for steps $h' \geq h + 1$ (in the view of $\mathcal{D}^{\text{expand}}$), we can bound (28) using standard martingale concentration inequalities.

- (2) Take the union bound over all possible $\mathcal{I} \in \mathcal{C}$ — with the aid of Lemma 4.2 — to obtain a uniform control of the term (28), simultaneously accounting for all $\mathcal{I} \in \mathcal{C}$ and all associated sequences $\{X_{h+1,s,a}\}$.
- (3) We then demonstrate that the above uniform bounds can be applied to the decomposition (27) to obtain a desired bound.

Remark 2. Note that in Step (1) of the above description, the estimate $\widehat{P}_{s,a,h}^{(l)}$ and those samples for steps $h' > h$ are statistically independent for fixed profile \mathcal{I} .

Main steps. We now carry out the above three steps.

Steps (1) and (2). Let us first specify the types of vectors $\{X_{h,s,a}\}$ mentioned above in (28). For each total profile $\mathcal{I} \in C$ (cf. (26)), consider any set $\{\mathcal{X}_{h,\mathcal{I}}\}_{1 \leq h \leq H}$ obeying: for each $1 \leq h \leq H$,

– $\mathcal{X}_{h+1,\mathcal{I}}$ is given by a *deterministic* function of \mathcal{I} and

$$\left\{ \widehat{P}_{s,a,h'}^{(I_{s,a,h}^k)}, \widehat{r}_{h'}^{(I_{s,a,h}^k)}(s, a), \widehat{\sigma}_{h'}^{(I_{s,a,h}^k)}(s, a) \right\}_{h < h' \leq H, (s, a, k) \in \mathcal{S} \times \mathcal{A} \times [K]};$$

– $\|X\|_\infty \leq H$ for each vector $X \in \mathcal{X}_{h,\mathcal{I}}$;

– $\mathcal{X}_{h,\mathcal{I}}$ is a set of no more than $K + 1$ non-negative vectors in \mathbb{R}^S , and contains the all-zero vector 0.

Given such a construction of $\{\mathcal{X}_{h,\mathcal{I}}\}$, we can readily conduct Steps (1) and (2), with a uniform concentration bound stated below.

LEMMA 4.4. *Suppose that $K \geq SAH \log_2 K$, and construct a set $\{\mathcal{X}_{h,\mathcal{I}}\}_{1 \leq h \leq H}$ for each $\mathcal{I} \in C$ satisfying the above properties. Then with probability at least $1 - \delta'$,*

$$\begin{aligned} \sum_{s,a,h \in \mathcal{S} \times \mathcal{A} \times [H]} \langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, X_{h+1,s,a} \rangle &\leq \sum_{s,a,h \in \mathcal{S} \times \mathcal{A} \times [H]} \max \left\{ \langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, X_{h+1,s,a} \rangle, 0 \right\} \\ &\leq \sqrt{\frac{8}{2^{l-2}} \sum_{s,a,h} \mathbb{V}(P_{s,a,h}, X_{h+1,s,a}) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} \\ &\quad + \frac{4H}{2^{l-2}} \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right), \end{aligned} \quad (29)$$

holds simultaneously for all $\mathcal{I} \in C$, all $2 \leq l \leq \log_2 K + 1$, and all sequences $\{X_{h,s,a}\}_{(s,a,h) \in \mathcal{S} \times \mathcal{A} \times [H]}$ obeying $X_{h,s,a} \in \mathcal{X}_{h+1,\mathcal{I}}$, $\forall (s,a,h) \in \mathcal{S} \times \mathcal{A} \times [H]$. Here, we recall that $\delta' = \frac{\delta}{200SAH^2K^2}$.

PROOF. We first invoke the Freedman inequality to bound the target quantity for any fixed $\mathcal{I} \in C$, any fixed integer l , and any fixed feasible sequence $\{X_{h,s,a}\}$, before applying the union bound to establish uniform control. See Appendix B.2 for details. \square

Step (3). Next, we turn to Step (3), which is accomplished via the following lemma. Note that we also provide upper bounds for two additional quantities: $\sum_{k,h} \max\{\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle, 0\}$ and $\sum_{k,h} \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, (V_{h+1}^k)^2 \rangle$, which will be useful in subsequent analysis.

LEMMA 4.5. *Suppose that $K \geq SAH \log_2 K$. With probability exceeding $1 - \delta'$, we have*

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \left\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \right\rangle &\leq \sum_{k=1}^K \sum_{h=1}^H \max \left\{ \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle, 0 \right\} \\ &\leq \sqrt{16(\log_2 K) \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} \\ &\quad + 49SAH^2 \log_2^3 K + 8H(\log_2 K) \log \frac{1}{\delta'} \end{aligned}$$

and

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H \left\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, (V_{h+1}^k)^2 \right\rangle \\ & \leq 8H \sqrt{(\log_2 K) \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} \\ & \quad + 49SAH^3 \log_2^3 K + 8H^2 (\log_2 K) \log \frac{1}{\delta'}. \end{aligned}$$

PROOF. This result is proved by combining the uniform bound in Lemma 4.4 with the decomposition (27). See Appendix B.3. \square

Thus far, we have obtained high-probability bounds on the most challenging terms. The complete proof of Theorem 1.1 will be presented next in Section 5.

5 Proof of Theorem 1.1

This section is devoted to proving Theorem 1.1. For notational convenience, let B be a logarithmic term

$$B = 4,000(\log_2 K)^3 \log(3SAH) \log \frac{1}{\delta'}, \quad (30)$$

where we recall that δ is the confidence parameter in Algorithm 1 and $\delta' = \frac{\delta}{200SAH^2K^2}$. When $K \leq BSAH$, the claimed result in Theorem 1.1 holds trivially since

$$\text{Regret}(K) = \sum_{k=1}^K \left(V_1^\star(s_1^k) - V_1^{\pi^k}(s_1^k) \right) \leq HK = \min \left\{ \sqrt{BSAH^3K}, HK \right\}.$$

As a result, it suffices to focus on the scenario with

$$K \geq BSAH \quad \text{with } B = 4,000(\log_2 K)^3 \log(3SAH) \log \frac{1}{\delta'}. \quad (31)$$

Our regret analysis for Algorithm 1 consists of several steps described below.

Step 1: the optimism principle. To begin with, we justify that the running estimates of Q-function and value function in Algorithm 1 are always upper bounds on the optimal Q-function and the optimal value function, respectively, thereby guaranteeing optimism in the face of uncertainty.

LEMMA 5.1 (OPTIMISM). *With probability exceeding $1 - 4SAHK\delta'$, one has*

$$Q_h^k(s, a) \geq Q_h^\star(s, a) \quad \text{and} \quad V_h^k(s) \geq V_h^\star(s), \quad (32)$$

for all (s, a, h, k) .

PROOF. See Appendix C.1. \square

Step 2: regret decomposition. In view of the optimism shown in Lemma 5.1, the regret can be upper bounded by

$$\text{Regret}(K) = \sum_{k=1}^K (V_1^\star(s_1^k) - V_1^{\pi^k}(s_1^k)) \leq \sum_{k=1}^K (V_1^k(s_1^k) - V_1^{\pi^k}(s_1^k)), \quad (33)$$

with probability at least $1 - 4SAHK\delta'$. In order to control the right-hand side of (33), we first make note of the following upper bound on $V_1^k(s_1^k)$.

LEMMA 5.2. For every $1 \leq k \leq K$, one has

$$V_1^k(s_1^k) \leq \sum_{h=1}^H \left(\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle + b_h^k(s_h^k, a_h^k) + \widehat{r}_h^k(s_h^k, a_h^k) + \langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, V_{h+1}^k \rangle \right).$$

PROOF OF LEMMA 5.2. From the construction of V_h^k and Q_h^k , it is seen that, for each $1 \leq h \leq H$,

$$\begin{aligned} V_h^k(s_h^k) &= Q_h^k(s_h^k, a_h^k) \leq \widehat{r}_h^k(s_h^k, a_h^k) + \widehat{P}_{s_h^k, a_h^k, h}^k V_{h+1}^k + b_h^k(s_h^k, a_h^k) \\ &= \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle + b_h^k(s_h^k, a_h^k) + \widehat{r}_h^k(s_h^k, a_h^k) + \langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, V_{h+1}^k \rangle + V_{h+1}^k(s_{h+1}^k). \end{aligned}$$

Applying this relation recursively over $1 \leq h \leq H$ gives

$$\begin{aligned} V_1^k(s_1^k) & \\ &\leq \sum_{h=1}^H \left(\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle + b_h^k(s_h^k, a_h^k) + \widehat{r}_h^k(s_h^k, a_h^k) + \langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, V_{h+1}^k \rangle \right) + V_{H+1}^k(s_{H+1}^k), \end{aligned}$$

which combined with $V_{H+1}^k = 0$ concludes the proof. \square

Combine Lemma 5.2 with (33) to show that, with probability at least $1 - 4SAHK\delta'$,

$$\begin{aligned} \text{Regret}(K) &\leq \underbrace{\sum_{k=1}^K \sum_{h=1}^H \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle}_{=: T_1} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H b_h^k(s_h^k, a_h^k)}_{=: T_2} \\ &\quad + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, V_{h+1}^k \rangle}_{=: T_3} + \underbrace{\sum_{k=1}^K \left(\sum_{h=1}^H \widehat{r}_h^k(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k) \right)}_{=: T_4}, \end{aligned} \quad (34)$$

leaving us with four terms to control. In particular, T_1 has already been upper bounded in Section 4.2, and hence we shall describe how to bound T_2, \dots, T_4 in the sequel.

Step 3.1: bounding the terms T_2, T_3 , and T_4 . In this section, we seek to bound the terms T_2, T_3 , and T_4 defined in the regret decomposition (34). To do so, we find it helpful to first introduce the following quantities that capture some sort of aggregate variances:

$$T_5 := \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k), \quad (35a)$$

$$T_6 := \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k), \quad (35b)$$

with T_5 denoting certain empirical variance and T_6 the true variance. With these quantities in place, we claim that the following bounds hold true.

LEMMA 5.3. *With probability exceeding $1 - 15SAH^2K^2\delta'$, one has*

$$T_2 \leq 61\sqrt{2SAH(\log_2 K)\left(\log \frac{1}{\delta'}\right)T_5} + 8\sqrt{SAH^3K(\log_2 K)\log \frac{1}{\delta'}} + 151SAH^2(\log_2 K)\log \frac{1}{\delta'}, \quad (36a)$$

$$|T_3| \leq \sqrt{8T_6 \log \frac{1}{\delta'}} + 3H \log \frac{1}{\delta'}, \quad (36b)$$

$$|T_4| \leq 6\sqrt{2SAH^3K(\log_2 K)\log \frac{1}{\delta'}} + 55SAH^2(\log_2 K)\log \frac{1}{\delta'}. \quad (36c)$$

PROOF. See Appendix C.2. \square

Step 3.2: bounding the aggregate variances T_5 and T_6 . The previous bounds on T_2 and T_3 stated in Lemma 5.3 depend respectively on the aggregate variance T_5 and T_6 (cf. (35a) and (35b)), which we would like to control now. By introducing the following quantities:

$$T_7 := \sum_{k=1}^K \sum_{h=1}^H \left\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, (V_{h+1}^k)^2 \right\rangle, \quad (37a)$$

$$T_8 := \sum_{k=1}^K \sum_{h=1}^H \left\langle P_{s_h^k, a_h^k, h}^k - e_{s_{h+1}^k}, (V_{h+1}^k)^2 \right\rangle, \quad (37b)$$

$$T_9 := \sum_{k=1}^K \sum_{h=1}^H \max \left\{ \left\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \right\rangle, 0 \right\}, \quad (37c)$$

we can upper bound T_5 and T_6 through the following lemma. As a remark, the proof of this lemma makes use of the law of total variance, as in previous work like Azar et al. [6].

LEMMA 5.4. *With probability at least $1 - 4SAHK\delta'$,*

$$T_5 \leq T_7 + T_8 + 2HT_2 + 6KH^2, \quad (38a)$$

$$T_6 \leq 2HT_2 + 6KH^2 + \sqrt{32H^2T_6 \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'} + 2HT_9, \quad (38b)$$

$$|T_8| \leq \sqrt{32H^2T_6 \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'}. \quad (38c)$$

PROOF. See Appendix C.3. \square

Step 3.3: bounding the terms T_1 , T_7 , and T_9 . Taking a look at the above bounds on T_2, \dots, T_6 , we see that one still needs to deal with the terms T_1 , T_7 , and T_9 (see (34), (37a), and (37c), respectively). As it turns out, these quantities have already been bounded in Section 4. Specifically, Lemma 4.5 tells us that: with probability at least $1 - \delta'$,

$$T_1 \leq T_9 \leq \sqrt{BSAH \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k) + BSAH^2} = \sqrt{BSAHT_6} + BSAH^2, \quad (39a)$$

$$T_7 \leq H \sqrt{BSAH \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k) + BSAH^3} = H\sqrt{BSAHT_6} + BSAH^3, \quad (39b)$$

where we recall that $B = 4,000(\log_2 K)^3 \log(3SAH) \log \frac{1}{\delta'}$.

Step 4: putting all the pieces together. The previous bounds (36), (38), and (39) indicate that: with probability at least $1 - 100SAH^2K^2\delta'$, one has

$$T_2 \leq \sqrt{BSAHT_5} + \sqrt{BSAH^3K} + BSAH^2, \quad (40a)$$

$$T_3 \leq \sqrt{BT_6} + HB, \quad (40b)$$

$$T_4 \leq \sqrt{BSAH^3K} + BSAH^2, \quad (40c)$$

$$T_5 \leq T_7 + T_8 + 2HT_2 + 6KH^2, \quad (40d)$$

$$T_6 \leq \sqrt{BH^2T_6} + 2HT_2 + 2HT_9 + BH^2 + 6KH^2, \quad (40e)$$

$$T_8 \leq \sqrt{BH^2T_6} + BH^2, \quad (40f)$$

$$T_1 \leq \sqrt{BSAHT_6} + BSAH^2, \quad (40g)$$

$$T_7 \leq H\sqrt{BSAHT_6} + BSAH^3, \quad (40h)$$

$$T_9 \leq \sqrt{BSAHT_6} + BSAH^2, \quad (40i)$$

where we again use $B = 4,000(\log_2 K)^3 \log(3SAH) \log \frac{1}{\delta'}$.

To solve the inequalities (40), we resort to the elementary AM-GM inequality: if $a \leq \sqrt{bc} + d$ for some $b, c \geq 0$, then it follows that $a \leq \epsilon b + \frac{1}{2\epsilon}c + d$ for any $\epsilon > 0$. This basic inequality combined with (40) gives

$$\begin{aligned} HT_2 &\leq \epsilon T_5 + \left(\frac{1}{2\epsilon} + 1\right) BSAH^3 + \frac{3}{2} BSAH^3 + \frac{1}{2} KH^2, \\ T_6 &\leq \epsilon T_6 + 2HT_2 + 2HT_9 + \left(1 + \frac{1}{2\epsilon}\right) BH^2 + 6KH^2, \\ HT_9 &\leq \epsilon T_6 + \left(\frac{1}{2\epsilon} + 1\right) BSAH^3, \\ T_8 &\leq \epsilon T_6 + \left(\frac{1}{2\epsilon} + 1\right) BH^2, \\ T_7 &\leq \epsilon T_6 + \left(\frac{1}{2\epsilon} + 1\right) BSAH^3, \end{aligned}$$

which in turn result in

$$T_5 \leq T_7 + T_8 + 2HT_2 + 6KH^2 \leq 2\epsilon T_5 + 2\epsilon T_6 + \left(\frac{1}{\epsilon} + 2\right) BSAH^3 + 6KH^2;$$

$$T_6 \leq \epsilon T_6 + 2HT_2 + 2HT_9 + \left(1 + \frac{1}{2\epsilon}\right) BH^2 + 6KH^2 \leq 3\epsilon T_6 + 2\epsilon T_5 + \left(\frac{3}{\epsilon} + 8\right) BSAH^3 + 7KH^2.$$

By taking $\epsilon = 1/20$, we arrive at

$$T_5 + T_6 \lesssim BSAH^3 + KH^2 \asymp KH^2, \quad (41)$$

where the last relation holds due to our assumption $K \geq SAHB$ (cf. (31)). Substituting this into (40) yields

$$T_1 \lesssim \sqrt{BSAH^3K}, \quad T_2 \lesssim \sqrt{BSAH^3K}, \quad T_3 \lesssim \sqrt{BKH^2} \quad \text{and} \quad T_4 \lesssim \sqrt{BSAH^3K}, \quad (42)$$

provided that $K \geq SAHB$. These bounds taken collectively with (34) readily give

$$\text{Regret}(K) \lesssim \sqrt{BSAH^3K}.$$

Combining the two scenarios (i.e., $K \geq BSAH$ and $K \leq BSAH$) reveals that with probability at least $1 - 100SAH^2K^2\delta'$,

$$\text{Regret}(K) \lesssim \min \left\{ \sqrt{BSAH^3K}, HK \right\} \lesssim \min \left\{ \sqrt{BSAH^3K \log^5 \frac{SAHK}{\delta'}}, HK \right\}.$$

The proof of Theorem 1.1 is thus completed by recalling that $\delta' = \frac{\delta}{200SAH^2K^2}$.

6 Extensions

In this section, we develop more refined regret bounds for Algorithm 1 in order to reflect the role of several problem-dependent quantities. Detailed proofs are postponed to Appendix D and Appendix F.

6.1 Value-based Regret Bounds

Thus far, we have not yet introduced the crucial quantity v^* in Theorem 1.2, which we define now. When the initial states are drawn from μ , we define v^* to be the weighted optimal value:

$$v^* := \mathbb{E}_{s \sim \mu} [V_1^*(s)]. \quad (43)$$

Encouragingly, the value-dependent regret bound we develop in Theorem 1.2 is still minimax-optimal, as asserted by the following lower bound.

THEOREM 6.1. *Consider any $p \in [0, 1]$ and $K \geq 1$. For any learning algorithm, there exists an MDP with S states, A actions and horizon H obeying $v^* \leq Hp$ and*

$$\mathbb{E}[\text{Regret}(K)] \gtrsim \min \left\{ \sqrt{SAH^3Kp}, KHp \right\}. \quad (44)$$

In fact, the construction of the hard instance (as in the proof of Theorem 6.1) is quite simple. Design a new branch with 0 reward and set the probability of reaching this branch to be $1-p$. Also, with probability p , we direct the learner to a hard instance with regret $\Omega(\min\{\sqrt{SAH^3Kp}, KHp\})$ and optimal value H . This guarantees that the optimal value obeys $v^* \leq Hp$ and that the expected regret is at least

$$\Omega \left(\min \left\{ \sqrt{SAH^3Kp}, KHp \right\} \right) \gtrsim \min \left\{ \sqrt{SAH^2Kv^*}, Kv^* \right\}.$$

See Appendix G for more details.

6.2 Cost-based Regret Bounds

Next, we turn to the cost-aware regret bound as in Theorem 1.3. Note that all other results except for Theorem 1.3 (and a lower bound in this subsection) are about rewards as opposed to cost. In order to facilitate discussion, let us first formally formulate the cost-based scenarios.

Suppose that the reward distributions $\{R_{h,s,a}\}_{(s,a,h)}$ are replaced with the cost distributions $\{C_{h,s,a}\}_{(s,a,h)}$, where each distribution $C_{h,s,a} \in \Delta([0, H])$ has mean $c_h(s, a)$. In the h th step of an episode, the learner pays an immediate cost $c_h \sim C_{h,s_h,a_h}$ instead of receiving an immediate reward r_h , and the objective of the learner is instead to minimize the total cost $\sum_{h=1}^H c_h$ (in an expected sense). The optimal cost quantity c^* is then defined as

$$c^* := \min_{\pi} \mathbb{E}_{\pi, s_1 \sim \mu} \left[\sum_{h=1}^H c_h \right]. \quad (45)$$

In this cost-based setting, we find it convenient to re-define the Q -function and value function as follows:

$$Q_h^{\pi, \text{cost}}(s, a) := \mathbb{E}_\pi \left[\sum_{h'=h}^H c_{h'} \mid (s_h, a_h) = (s, a) \right], \quad \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H],$$

$$V_h^{\pi, \text{cost}}(s) := \mathbb{E}_\pi \left[\sum_{h'=h}^H c_{h'} \mid s_h = s \right], \quad \forall (s, h) \in \mathcal{S} \times [H],$$

where we adopt different fonts to differentiate them from the original Q -function and value function. The optimal cost function is then defined by

$$Q_h^{\star, \text{cost}}(s, a) = \min_{\pi} Q_h^{\pi, \text{cost}}(s, a) \quad \text{and} \quad V_h^{\star, \text{cost}}(s) = \min_{\pi} V_h^{\pi, \text{cost}}(s).$$

Given the definitions above, we overload the notation $\text{Regret}(K)$ to denote the regret for the cost-based scenario as

$$\text{Regret}(K) := \sum_{k=1}^K \left(V_1^{\pi^k, \text{cost}}(s_1^k) - V_1^{\star, \text{cost}}(s_1^k) \right).$$

One can also simply regard the cost minimization problem as reward maximization with negative rewards by choosing $r_h = -c_h$. This way allows us to apply Algorithm 1 directly, except that (20) is replaced by

$$Q_h(s, a) \leftarrow \max \left\{ \min \left\{ \widehat{r}_h(s, a) + \widehat{P}_{s, a, h} V_{h+1} + b_h(s, a), 0 \right\}, -H \right\}. \quad (46)$$

Note that the proof of Theorem 1.3 closely resembles that of Theorem 1.2, which can be found in Appendix E.

To confirm the tightness of Theorem 1.3, we develop the following matching lower bound, which resorts to a similar hard instance as in the proof of Theorem 6.1.

THEOREM 6.2. *Consider any $p \in [0, 1/4]$ and any $K \geq 1$. For any algorithm, one can construct an MDP with S states, A actions and horizon H obeying $c^\star \asymp H p$ and*

$$\mathbb{E}[\text{Regret}(K)] \gtrsim \min \left\{ \sqrt{SAH^3Kp} + SAH^2, KH(1-p) \right\} \asymp \min \left\{ \sqrt{SAH^2Kc^\star} + SAH^2, KH \right\}.$$

The proof of this lower bound can be found in Appendix G.2.

6.3 Variance-dependent Regret Bound

The final regret bound presented in Theorem 1.4 depends on some sort of variance metrics. Toward this end, let us first make precise the variance metrics of interest:

(i) The first variance metric is defined as

$$\text{var}_1 := \max_{\pi} \mathbb{E}_\pi \left[\sum_{h=1}^H \mathbb{V}(P_{s_h, a_h, h}, V_{h+1}^\star) + \sum_{h=1}^H \text{Var}(R_h(s_h, a_h)) \right], \quad (47)$$

where $\{(s_h, a_h)\}_{1 \leq h \leq H}$ represents a sample trajectory under policy π . This captures the maximal possible expected sum of variance with respect to the optimal value function $\{V_h^\star\}_{h=1}^H$.

(ii) Another useful variance metric is defined as

$$\text{var}_2 := \max_{\pi, s} \text{Var}_\pi \left[\sum_{h=1}^H r_h \mid s_1 = s \right], \quad (48)$$

where $\{r_h\}_{1 \leq h \leq H}$ denotes a sample sequence of immediate rewards under policy π . This indicates the maximal possible variance of the accumulative reward.

The interested reader is referred to Zhou et al. [87] for further discussion about these two metrics. Our final variance metric is then defined as

$$\text{var} := \min \{\text{var}_1, \text{var}_2\}. \quad (49)$$

With the above variance metrics in mind, we can then revisit Theorem 1.4. As a special case, when the transition model is fully deterministic, the regret bound in Theorem 1.4 simplifies to

$$\text{Regret}(K) \leq \tilde{O}(\min \{SAH^2, HK\}),$$

for any $K \geq 1$, which is roughly the cost of visiting each state-action pair. The full proof of Theorem 1.4 is postponed to Appendix F.

To finish up, let us develop a matching lower bound to corroborate the tightness and optimality of Theorem 1.4.

THEOREM 6.3. *Consider any $p \in [0, 1]$ and any $K \geq 1$. For any algorithm, one can find an MDP instance with S states, A actions, and horizon H satisfying $\max\{\text{var}_1, \text{var}_2\} \leq H^2p$ and*

$$\mathbb{E}[\text{Regret}(K)] \geq \min \{\sqrt{SAH^3Kp} + SAH^2, KH\}.$$

The proof of Theorem 6.3 resembles that of Theorem 6.1, except that we need to construct a hard instance when $K \leq SAH/p$. For this purpose, we construct a fully deterministic MDP (i.e., all of its transitions are deterministic and all rewards are fixed), and show that the learner has to visit about half of the state-action-layer tuples in order to learn a near-optimal policy. The proof details are deferred to Appendix G.

7 Discussion

Focusing on tabular online RL in time-inhomogeneous finite-horizon MDPs, this article has established the minimax-optimal regret (respectively sample complexity) – up to log factors – for the entire range of sample size $K \geq 1$ (respectively target accuracy level $\varepsilon \in (0, H]$), thereby fully settling an open problem at the core of recent RL theory. The MVP JAC studied herein is model-based in nature. Remarkably, the model-based approach remains the only family of algorithms that is capable of obtaining minimax optimality without burn-ins, regardless of the data collection mechanism in use (e.g., online RL, offline RL, and the simulator setting). We have further unlocked the optimality of this algorithm in a more refined manner, making apparent the effect of several problem-dependent quantities (e.g., optimal value/cost, variance statistics) upon the fundamental performance limits. The new analysis and algorithmic techniques put forward herein might shed important light on how to conquer other RL settings as well.

Moving forward, there are multiple directions that anticipate further theoretical pursuit. To begin with, is it possible to develop a model-free algorithm – which often exhibits more favorable memory complexity compared to the model-based counterpart – that achieves full-range minimax optimality? As alluded to previously, existing paradigms that rely on reference-advantage decomposition (or variance reduction) seem to incur a high burn-in cost [42, 84], thus calling for new ideas to overcome this barrier. Additionally, multiple other tabular settings (e.g., time-homogeneous finite-horizon MDPs, discounted infinite-horizon MDPs) have also suffered from similar issues regarding the burn-in requirements [27, 82]. Take time-homogeneous finite-horizon MDPs for example: in order to achieve optimal sample efficiency, one needs to carefully deal with the statistical dependency incurred by aggregating data from across different time steps to estimate the same transition matrix (due to the homogeneous nature of P), which results in more intricate issues than the time-homogeneous counterpart. We believe that resolving these two open problems will greatly enhance our theoretical understanding about online RL and beyond.

Acknowledgments

We thank for Qiwen Cui for helpful discussions.

References

- [1] Alekh Agarwal, Sham Kakade, and Lin F. Yang. 2020. Model-based reinforcement learning with a generative model is minimax optimal. In *The 33rd Annual Conference on Learning Theory*. 67–83.
- [2] Alekh Agarwal, Akshay Krishnamurthy, John Langford, Haipeng Luo, et al. 2017. Open problem: First-order regret bounds for contextual bandits. In *The 30th Annual Conference on Learning Theory*. 4–7.
- [3] Shipra Agrawal and Randy Jia. 2017. Optimistic posterior sampling for reinforcement learning: Worst-case regret bounds. In *Proceedings of the Advances in Neural Information Processing Systems*. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Curran Associates, Inc., 1184–1194.
- [4] Zeyuan Allen-Zhu, Sébastien Bubeck, and Yuanzhi Li. 2018. Make the minority great again: First-order regret bound for contextual bandits. In *Proceedings of the International Conference on Machine Learning*. 186–194.
- [5] Mohammad Gheshlaghi Azar, Rémi Munos, and Hilbert J. Kappen. 2013. Minimax PAC bounds on the sample complexity of reinforcement learning with a generative model. *Machine Learning* 91, 3 (2013), 325–349.
- [6] Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. 2017. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*. 263–272.
- [7] Yu Bai, Tengyang Xie, Nan Jiang, and Yu-Xiang Wang. 2019. Provably efficient Q-learning with low switching cost. In *Proceedings of the Advances in Neural Information Processing Systems*. 8004–8013.
- [8] Peter L. Bartlett and Ambuj Tewari. 2009. REGAL: A regularization based algorithm for reinforcement learning in weakly communicating MDPs. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*.
- [9] Carolyn L. Beck and Rayadurgam Srikant. 2012. Error bounds for constant step-size Q-learning. *Systems and Control Letters* 61, 12 (2012), 1203–1208.
- [10] Dimitri Bertsekas. 2019. *Reinforcement Learning and Optimal Control*. Athena Scientific.
- [11] Ronen I. Brafman and Moshe Tennenholtz. 2003. R-max - a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research* 3, Oct (2003), 213–231.
- [12] Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. 2019. Provably efficient exploration in policy optimization. In *Proceedings of the 37th International Conference on Machine Learning*. PMLR, 1283–1294.
- [13] Liyu Chen, Mehdi Jafarnia-Jahromi, Rahul Jain, and Haipeng Luo. 2021. Implicit finite-horizon approximation and efficient optimal algorithms for stochastic shortest path. In *Proceedings of the 35th International Conference on Neural Information Processing Systems* 34 (2021), 10849–10861.
- [14] Zaiwei Chen, Siva Theja Maguluri, Sanjay Shakkottai, and Karthikeyan Shanmugam. 2020. Finite-sample analysis of contractive stochastic approximation using smooth convex envelopes. *Advances in Neural Information Processing Systems* 33 (2020), 8223–8234.
- [15] Qiwen Cui and Lin F. Yang. 2021. Minimax sample complexity for turn-based stochastic game. In *Proceedings of the Uncertainty in Artificial Intelligence*. 1496–1504.
- [16] Christoph Dann and Emma Brunskill. 2015. Sample complexity of episodic fixed-horizon reinforcement learning. In *Proceedings of the Advances in Neural Information Processing Systems*. 2818–2826.
- [17] Christoph Dann, Tor Lattimore, and Emma Brunskill. 2017. Unifying PAC and regret: Uniform PAC bounds for episodic reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* 30 (2017), 5717–5727.
- [18] Christoph Dann, Lihong Li, Wei Wei, and Emma Brunskill. 2019. Policy certificates: Towards accountable reinforcement learning. In *Proceedings of the 36th International Conference on Machine Learning*. 1507–1516.
- [19] Christoph Dann, Teodor Vanislavov Marinov, Mehryar Mohri, and Julian Zimmert. 2021. Beyond value-function gaps: Improved instance-dependent regret bounds for episodic reinforcement learning. In *Proceedings of the 35th International Conference on Neural Information Processing Systems* 34 (2021), 1–12.
- [20] Omar Darwiche Domingues, Pierre Ménard, Emilie Kaufmann, and Michal Valko. 2021. Episodic reinforcement learning in finite mdps: Minimax lower bounds revisited. In *Algorithmic Learning Theory*. 578–598.
- [21] Kefan Dong, Yuanhao Wang, Xiaoyu Chen, and Liwei Wang. 2019. Q-learning with UCB exploration is sample efficient for infinite-horizon MDP. In *International Conference on Learning Representations*.
- [22] Yonathan Efroni, Nadav Merlis, Mohammad Ghavamzadeh, and Shie Mannor. 2019. Tight regret bounds for model-based reinforcement learning with greedy policies. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems* 32 (2019), 12224–12234.
- [23] Eyal Even-Dar and Yishay Mansour. 2003. Learning rates for Q-learning. *Journal of Machine Learning Research* 5 (2003), 1–25.
- [24] David A. Freedman. 1975. On tail probabilities for martingales. *The Annals of Probability* 3, 1 (1975), 100–118.

- [25] Ronan Fruin, Matteo Pirotta, Alessandro Lazaric, and Ronald Ortner. 2018. Efficient bias-span-constrained exploration-exploitation in reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*. 30 (2018), 1578–1586.
- [26] Thomas Jaksch, Ronald Ortner, and Peter Auer. 2010. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research* 11 (2010), 1563–1600.
- [27] Xiang Ji and Gen Li. 2023. Regret-optimal model-free reinforcement learning for discounted MDPs with short burn-in time. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*. 80674–80689.
- [28] Nan Jiang and Alekh Agarwal. 2018. Open problem: The dependence of sample complexity lower bounds on planning horizon. In *Proceedings of the Conference On Learning Theory*. 3395–3398.
- [29] Chi Jin, Zeyuan Allen-Zhu, Sébastien Bubeck, and Michael I. Jordan. 2018. Is Q-learning provably efficient?. In *Proceedings of the Advances in Neural Information Processing Systems*. 4863–4873.
- [30] Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. 2020. Reward-free exploration for reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning*. PMLR, 4870–4879.
- [31] Ying Jin, Zhioran Yang, and Zhaoran Wang. 2021. Is pessimism provably efficient for offline RL?. In *Proceedings of the International Conference on Machine Learning*. 5084–5096.
- [32] Sham M. Kakade. 2003. *On the Sample Complexity of Reinforcement Learning*. Ph. D. Dissertation. University of London London, England.
- [33] Michael Kearns and Satinder Singh. 1998. Finite-sample convergence rates for Q-learning and indirect algorithms. In *Proceedings of the 12th International Conference on Neural Information Processing Systems* 11 (1998), 996–1002.
- [34] Michael J. Kearns and Satinder P. Singh. 1998. Near-optimal reinforcement learning in polynomial time. In *Proceedings of the 15th International Conference on Machine Learning*. 260–268.
- [35] J. Zico Kolter and Andrew Y. Ng. 2009. Near-Bayesian exploration in polynomial time. In *Proceedings of the 26th Annual International Conference on Machine Learning*. 513–520.
- [36] Tor Lattimore and Marcus Hutter. 2012. PAC bounds for discounted MDPs. In *Proceedings of the International Conference on Algorithmic Learning Theory*. Springer, 320–334.
- [37] Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, and Mengxiao Zhang. 2020. Bias no more: High-probability data-dependent regret bounds for adversarial bandits and mdps. *Advances in Neural Information Processing Systems* 33 (2020), 15522–15533.
- [38] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. 2020. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643* (2020).
- [39] Gen Li, Changxiao Cai, Yuxin Chen, Yuting Wei, and Yuejie Chi. 2024. Is Q-learning minimax optimal? A tight sample complexity analysis. *Operations Research* 72, 1 (2024), 222–236.
- [40] Gen Li, Yuejie Chi, Yuting Wei, and Yuxin Chen. 2022. Minimax-optimal multi-agent RL in Markov games with a generative model. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*. 35 (2022), 15353–15367.
- [41] Gen Li, Laixi Shi, Yuxin Chen, Yuejie Chi, and Yuting Wei. 2024. Settling the sample complexity of model-based offline reinforcement learning. *The Annals of Statistics* 52, 1 (2024), 233–260.
- [42] Gen Li, Laixi Shi, Yuxin Chen, Yuantao Gu, and Yuejie Chi. 2021. Breaking the sample complexity barrier to regret-optimal model-free reinforcement learning. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* 34 (2021), 17762–17776.
- [43] Gen Li, Yuting Wei, Yuejie Chi, and Yuxin Chen. 2024. Breaking the sample size barrier in model-based reinforcement learning with a generative model. *Operations Research* 72, 1 (2024), 203–221.
- [44] Gen Li, Yuting Wei, Yuejie Chi, Yuantao Gu, and Yuxin Chen. 2021. Sample complexity of asynchronous Q-learning: Sharper analysis and variance reduction. *IEEE Transactions on Information Theory* 68, 1 (2021), 448–473.
- [45] Gen Li, Yuling Yan, Yuxin Chen, and Jianqing Fan. 2024. Minimax-optimal reward-agnostic exploration in reinforcement learning. In *The 37th Annual Conference on Learning Theory*. PMLR, 3431–3436.
- [46] Gen Li, Wenhao Zhan, Jason D. Lee, Yuejie Chi, and Yuxin Chen. 2023. Reward-agnostic fine-tuning: Provable statistical benefits of hybrid reinforcement learning. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* 36 (2024), 55582–55615.
- [47] Yuanzhi Li, Ruosong Wang, and Lin F. Yang. 2021. Settling the horizon-dependence of sample complexity in reinforcement learning. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*.
- [48] Andreas Maurer and Massimiliano Pontil. 2009. Empirical Bernstein bounds and sample variance penalization. In *Proceedings of the Conference on Learning Theory*.
- [49] Pierre Ménard, Omar Darwiche Domingues, Xuedong Shang, and Michal Valko. 2021. UCB momentum Q-learning: Correcting the bias without forgetting. In *Proceedings of the International Conference on Machine Learning*. 7609–7618.
- [50] Gergely Neu and Ciara Pike-Burke. 2020. A unifying view of optimism in episodic reinforcement learning. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*. 1392–1403.

[51] Ian Osband, Daniel Russo, and Benjamin Van Roy. 2013. (More) efficient reinforcement learning via posterior sampling. In *Proceedings of the Advances in Neural Information Processing Systems*. 3003–3011.

[52] Aldo Pacchiano, Philip Ball, Jack Parker-Holder, Krzysztof Choromanski, and Stephen Roberts. 2020. On optimism in model-based reinforcement learning. In *Uncertainty in Artificial Intelligence*. PMLR, 1413–1423.

[53] Ashwin Pananjady and Martin J. Wainwright. 2020. Instance-dependent ℓ_∞ -bounds for policy evaluation in tabular reinforcement learning. *IEEE Transactions on Information Theory* 67, 1 (2020), 566–585.

[54] Guannan Qu and Adam Wierman. 2020. Finite-time analysis of asynchronous stochastic approximation and Q-Learning. In *Proceedings of the 33rd Conference on Learning Theory*.

[55] Paria Rashidinejad, Banghua Zhu, Cong Ma, Jiantao Jiao, and Stuart Russell. 2021. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *Advances in Neural Information Processing Systems* 34 (2021), 11702–11716.

[56] Tongzheng Ren, Jialian Li, Bo Dai, Simon S. Du, and Sujay Sanghavi. 2021. Nearly horizon-free offline reinforcement learning. *Advances in Neural Information Processing Systems* 34 (2021), 15621–15634.

[57] Daniel Russo. 2019. Worst-case regret bounds for exploration via randomized value functions. In *Proceedings of the Advances in Neural Information Processing Systems*. 14433–14443.

[58] Laixi Shi, Gen Li, Yuting Wei, Yuxin Chen, and Yuejie Chi. 2022. Pessimistic Q-learning for offline reinforcement learning: Towards optimal sample complexity. In *Proceedings of the International Conference on Machine Learning*. 19967–20025.

[59] Laixi Shi, Gen Li, Yuting Wei, Yuxin Chen, Matthieu Geist, and Yuejie Chi. 2023. The curious price of distributional robustness in reinforcement learning with a generative model. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*. 79903–79917.

[60] Aaron Sidford, Mengdi Wang, Xian Wu, Lin Yang, and Yinyu Ye. 2018. Near-optimal time and sample complexities for solving Markov decision processes with a generative model. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. 5192–5202.

[61] Aaron Sidford, Mengdi Wang, Xian Wu, and Yinyu Ye. 2018. Variance reduced value iteration and faster algorithms for solving Markov decision processes. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics. 770–787.

[62] Max Simchowitz and Kevin G. Jamieson. 2019. Non-asymptotic gap-dependent regret bounds for tabular MDPs. In *Proceedings of the Advances in Neural Information Processing Systems*. 1153–1162.

[63] Alexander L. Strehl, Lihong Li, Eric Wiewiora, John Langford, and Michael L. Littman. 2006. PAC model-free reinforcement learning. In *Proceedings of the 23rd International Conference on Machine Learning*. ACM, 881–888.

[64] Alexander L. Strehl and Michael L. Littman. 2008. An analysis of model-based interval estimation for Markov decision processes. *Journal of Computer and System Sciences* 74, 8 (2008), 1309–1331.

[65] István Szita and Csaba Szepesvári. 2010. Model-based reinforcement learning with nearly tight exploration complexity bounds. In *Proceedings of the 27th International Conference on Machine Learning*.

[66] Mohammad Sadegh Talebi and Odalric-Ambrym Maillard. 2018. Variance-aware regret bounds for undiscounted reinforcement learning in mdps. In *Algorithmic Learning Theory*. PMLR, 770–805.

[67] Jean Tarbouriech, Runlong Zhou, Simon S. Du, Matteo Pirotta, Michal Valko, and Alessandro Lazaric. 2021. Stochastic shortest path: Minimax, parameter-free and towards horizon-free regret. In *Proceedings of the 35th International Conference on Neural Information Processing Systems* 34 (2021), 6843–6855.

[68] Andrea Tirinzoni, Matteo Pirotta, and Alessandro Lazaric. 2021. A fully problem-dependent regret lower bound for finite-horizon MDPs. *arXiv preprint arXiv:2106.13013* (2021).

[69] Andrew J Wagenmaker, Yifang Chen, Max Simchowitz, Simon Du, and Kevin Jamieson. 2022. First-order regret in reinforcement learning with linear function approximation: A robust estimation approach. In *Proceedings of the International Conference on Machine Learning*. 22384–22429.

[70] Martin J. Wainwright. 2019. Stochastic approximation with cone-contractive operators: Sharp ℓ_∞ -bounds for Q-learning. *arXiv preprint arXiv:1905.06265* (2019).

[71] Martin J. Wainwright. 2019. Variance-reduced Q-learning is minimax optimal. *arXiv preprint arXiv:1906.04697* (2019).

[72] Kaiwen Wang, Kevin Zhou, Runzhe Wu, Nathan Kallus, and Wen Sun. 2023. The benefits of being distributional: Small-loss bounds for reinforcement learning. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*. 22775–2312.

[73] Ruosong Wang, Simon S. Du, Lin F. Yang, and Sham M. Kakade. 2020. Is long horizon reinforcement learning more difficult than short horizon reinforcement learning?. In *Proceedings of the Advances in Neural Information Processing Systems*. 9075–9085.

[74] Xinqi Wang, Qiwen Cui, and Simon S. Du. 2022. On gap-dependent bounds for offline reinforcement learning. *Advances in Neural Information Processing Systems* 35 (2022), 14865–14877.

- [75] Tengyang Xie, Nan Jiang, Huan Wang, Caiming Xiong, and Yu Bai. 2021. Policy finetuning: Bridging sample-efficient offline and online reinforcement learning. In *Proceedings of the 35th International Conference on Neural Information Processing Systems NeurIPS 2021*. 27395–27407.
- [76] Zhihan Xiong, Ruqi Shen, Qiwen Cui, Maryam Fazel, and Simon S. Du. 2022. Near-optimal randomized exploration for tabular markov decision processes. *Advances in Neural Information Processing Systems* 35 (2022), 6358–6371.
- [77] Haike Xu, Tengyu Ma, and Simon Du. 2021. Fine-grained gap-dependent bounds for tabular MDPs via adaptive multi-step bootstrap. In *Proceedings of the Conference on Learning Theory*. 4438–4472.
- [78] Yuling Yan, Gen Li, Yuxin Chen, and Jianqing Fan. 2023. The efficacy of pessimism in asynchronous Q-learning. *IEEE Transactions on Information Theory* 69, 11 (2023), 7185–7219.
- [79] Kunhe Yang, Lin Yang, and Simon Du. 2021. Q-learning with logarithmic regret. In *International Conference on Artificial Intelligence and Statistics*. 1576–1584.
- [80] Ming Yin, Yaqi Duan, Mengdi Wang, and Yu-Xiang Wang. 2022. Near-optimal offline reinforcement learning with linear representation: Leveraging variance information with pessimism. In *International Conference on Learning Representation*.
- [81] Andrea Zanette and Emma Brunskill. 2019. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *Proceedings of the International Conference on Machine Learning*. 7304–7312.
- [82] Zihan Zhang, Xiangyang Ji, and Simon Du. 2021. Is reinforcement learning more difficult than bandits? A near-optimal algorithm escaping the curse of horizon. In *Proceedings of the Conference on Learning Theory*. 4528–4531.
- [83] Zihan Zhang, Xiangyang Ji, and Simon Du. 2022. Horizon-free reinforcement learning in polynomial time: The power of stationary policies. In *Proceedings of the Conference on Learning Theory*. 3858–3904.
- [84] Zihan Zhang, Yuan Zhou, and Xiangyang Ji. 2020. Almost optimal model-free reinforcement learning via reference-advantage decomposition. In *Proceedings of the Advances in Neural Information Processing Systems*. 15198–15207.
- [85] Zihan Zhang, Yuan Zhou, and Xiangyang Ji. 2021. Model-free reinforcement learning: From clipped pseudo-regret to sample complexity. In *Proceedings of the International Conference on Machine Learning*. 12653–12662.
- [86] Heyang Zhao, Jiafan He, Dongruo Zhou, Tong Zhang, and Quanquan Gu. 2023. Variance-dependent regret bounds for linear bandits and reinforcement learning: Adaptivity and computational efficiency. In *The 36th Annual Conference on Learning Theory*. PMLR, 4977–5020.
- [87] Runlong Zhou, Zhang Zihan, and Simon Shaolei Du. 2023. Sharp variance-dependent bounds in reinforcement learning: Best of both worlds in stochastic and deterministic environments. In *Proceedings of the International Conference on Machine Learning*. 42878–42914.

Appendices

A Preliminary Facts

In this section, we gather several useful results that prove useful in our analysis. We use $\mathbb{1}\{\mathcal{E}\}$ to denote the indicator of the event \mathcal{E} . The first result below is a user-friendly version of the celebrated Freedman inequality [24], a martingale counterpart to the Bernstein inequality. See Zhang et al. [85, Lemma 11] for the proof.

LEMMA A.1 (FREEDMAN'S INEQUALITY). *Let $(M_n)_{n \geq 0}$ be a martingale such that $M_0 = 0$ and $|M_n - M_{n-1}| \leq c$ ($\forall n \geq 1$) hold for some quantity $c > 0$. Define $\text{Var}_n := \sum_{k=1}^n \mathbb{E}[(M_k - M_{k-1})^2 | \mathcal{F}_{k-1}]$ for every $n \geq 0$, where \mathcal{F}_k is the σ -algebra generated by (M_1, \dots, M_k) . Then for any integer $n \geq 1$ and any $\epsilon, \delta > 0$, one has*

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

Next, letting $\text{Var}(X)$ represent the variance of X , we record a basic inequality connecting $\text{Var}(X^2)$ with $\text{Var}(X)$ for any bounded random variable X .

LEMMA A.2 (LEMMA 30 IN [13]). *Let X be a random variable, and denote by C_{\max} the largest possible value of X . Then we have $\text{Var}(X^2) \leq 4C_{\max}^2 \text{Var}(X)$.*

Now, we turn to an intimate connection between the sum of a sequence of bounded non-negative random variables and the sum of their associated conditional random variables (with

each random variable conditioned on the past), which is a consequence of basic properties about supermartingales.

LEMMA A.3 (LEMMA 10 IN [83]). *Let X_1, X_2, \dots be a sequence of random variables taking value in $[0, l]$. For any $k \geq 1$, let \mathcal{F}_k be the σ -algebra generated by (X_1, X_2, \dots, X_k) , and define $Y_k := \mathbb{E}[X_k | \mathcal{F}_{k-1}]$. Then for any $\delta > 0$, we have*

$$\begin{aligned} \mathbb{P} \left[\exists n, \sum_{k=1}^n X_k \geq 3 \sum_{k=1}^n Y_k + l \log \frac{1}{\delta} \right] &\leq \delta \\ \mathbb{P} \left[\exists n, \sum_{k=1}^n Y_k \geq 3 \sum_{k=1}^n X_k + l \log \frac{1}{\delta} \right] &\leq \delta. \end{aligned}$$

The next two lemmas are concerned with concentration inequalities for the sum of i.i.d. bounded random variables: the first one is a version of the Bennet inequality, and the second one is an empirical Bernstein inequality (which replaces the variance in the standard Bernstein inequality with the empirical variance).

LEMMA A.4 (BENNET'S INEQUALITY). *Let Z, Z_1, \dots, Z_n be i.i.d. random variables with values in $[0, 1]$ and let $\delta > 0$. Define $\mathbb{V}Z = \mathbb{E}[(Z - \mathbb{E}Z)^2]$. Then one has*

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

LEMMA A.5 (THEOREM 4 IN MAURER AND PONTIL [48]). *Consider any $\delta > 0$ and any integer $n \geq 2$. Let Z, Z_1, \dots, Z_n be a collection of i.i.d. random variables falling within $[0, 1]$. Define the empirical mean $\bar{Z} := \frac{1}{n} \sum_{i=1}^n Z_i$ and empirical variance $\widehat{V}_n := \frac{1}{n} \sum_{i=1}^n (Z_i - \bar{Z})^2$. Then we have*

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

Moreover, we record a simple fact concerning the visitation counts $\{N_h^k(s_h^k, a_h^k)\}$.

LEMMA A.6. *Recall the definition of $N_h^k(s_h^k, a_h^k)$ in Algorithm 1. It holds that*

$$\sum_{k=1}^K \sum_{h=1}^H \frac{1}{\max\{N_h^k(s_h^k, a_h^k), 1\}} \leq 2SAH \log_2 K. \quad (50)$$

PROOF. In view of the doubling batch update rule, it is easily seen that: for any given (s, a, h) ,

$$\sum_{k=1}^K \frac{1}{\max\{N_h^k(s_h^k, a_h^k), 1\}} \mathbb{1}\{(s, a) = (s_h^k, a_h^k)\} \leq 2 \log_2 K, \quad (51)$$

since each (s, a, h) is associated with at most $\log_2 K$ epochs. Summing over (s, a, h) completes the proof. \square

As it turns out, Lemma A.6 together with the Freedman inequality allows one to control the difference between the empirical rewards and the true mean rewards, as stated below.

LEMMA A.7. *With probability exceeding $1 - 2SAHK\delta'$, it holds that*

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \left| \widehat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k) \right| &\leq 4\sqrt{2SAH^2(\log_2 K) \log \frac{1}{\delta'}} \sqrt{\sum_{k=1}^K \sum_{h=1}^H r_h(s_h^k, a_h^k)} \\ &\quad + 52SAH^2(\log_2 K) \log \frac{1}{\delta'}; \\ \sum_{k=1}^K \sum_{h=1}^H \widehat{r}_h^k(s_h^k, a_h^k) &\leq 2 \sum_{k=1}^K \sum_{h=1}^H r_h(s_h^k, a_h^k) + 60SAH^2(\log_2 K) \log \frac{1}{\delta'}. \end{aligned}$$

As an immediate consequence of Lemma A.7 and the basic fact $\sum_{k,h} r_h(s_h^k, a_h^k) \leq KH$, we have

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \widehat{r}_h^k(s_h^k, a_h^k) &\leq 2 \sum_{k=1}^K \sum_{h=1}^H r_h(s_h^k, a_h^k) + 60SAH^2(\log_2 K) \log \frac{1}{\delta'} \\ &\leq 2KH + 60SAH^2(\log_2 K) \log \frac{1}{\delta'} \leq 3KH, \end{aligned} \quad (52)$$

with probability exceeding $1 - 2SAHK\delta'$, where the last inequality holds true under the assumption 31.

PROOF OF LEMMA A.7. In view of Lemma A.5 and the union bound, with probability $1 - 2SAHK\delta'$ we have

$$\begin{aligned} \left| \widehat{r}_h^k(s, a) - r_h(s, a) \right| &\leq 2\sqrt{2} \sqrt{\frac{\left(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \right) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + \frac{28H \log \frac{1}{\delta'}}{3N_h^k(s, a)} \\ &\leq 2\sqrt{2} \sqrt{\frac{H\widehat{r}_h^k(s, a) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + \frac{28H \log \frac{1}{\delta'}}{3N_h^k(s, a)}, \end{aligned}$$

simultaneously for all (s, a, h, k) obeying $N_h^k(s, a) > 2$, where we take advantage of the basic fact $\widehat{\sigma}_h^k(s_h^k, a_h^k) \leq H\widehat{r}_h^k(s, a)$ (since each immediate reward is upper bounded by H). Solve the inequality above to obtain

$$\left| \widehat{r}_h^k(s, a) - r_h(s, a) \right| \leq 4\sqrt{\frac{Hr_h(s, a) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + 24 \frac{H \log \frac{1}{\delta'}}{N_h^k(s, a)}. \quad (53)$$

It is then seen that

$$\begin{aligned} \sum_{k,h} \left| \widehat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k) \right| &\leq 4SAH^2 + \sum_{k,h} \left(4\sqrt{\frac{Hr_h(s_h^k, a_h^k) \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} + 24 \frac{H \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)} \right) \\ &\leq 4SAH^2 + 4\sqrt{\sum_{k,h} \frac{H \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} \cdot \sqrt{\sum_{k,h} r_h(s_h^k, a_h^k)} + 24 \sum_{k,h} \frac{H \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}. \end{aligned}$$

Here, the second inequality arises from Cauchy–Schwarz, whereas the term $4SAH^2$ accounts for those state-action pairs with $N_h^k(s, a) \leq 2$ (since there are at most $2SAH$ such occurrences and it

holds that $|\hat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k)| \leq 2H$. This together with Lemma A.6 then leads to

$$\begin{aligned} \sum_{k,h} |\hat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k)| &\leq 4SAH^2 + 4\sqrt{2SAH^2(\log_2 K) \log \frac{1}{\delta'}} \sqrt{\sum_{k,h} r_h(s_h^k, a_h^k)} \\ &\quad + 48SAH^2(\log_2 K) \log \frac{1}{\delta'} \\ &\leq 4\sqrt{2SAH^2(\log_2 K) \log \frac{1}{\delta'}} \sqrt{\sum_{k,h} r_h(s_h^k, a_h^k)} + 52SAH^2(\log_2 K) \log \frac{1}{\delta'}. \end{aligned}$$

Moreover, the AM-GM inequality implies that

$$\begin{aligned} \sum_{k,h} \hat{r}_h^k(s_h^k, a_h^k) - \sum_{k,h} r_h(s_h^k, a_h^k) &\leq \sum_{k=1}^K \sum_{h=1}^H r_h(s_h^k, a_h^k) + 8SAH^2(\log_2 K) \log \frac{1}{\delta'} + 52SAH^2(\log_2 K) \log \frac{1}{\delta'} \\ \implies \sum_{k,h} \hat{r}_h^k(s_h^k, a_h^k) &\leq 2 \sum_{k,h} r_h(s_h^k, a_h^k) + 60SAH^2(\log_2 K) \log \frac{1}{\delta'}, \end{aligned}$$

thus concluding the proof. \square

B Proofs of Key Lemmas in Section 4

B.1 Proof of Lemma 4.2

It suffices to develop an upper bound on the cardinality of C (cf. (26)). Setting

$$M = \log_2 K \quad \text{and} \quad N = SAH, \quad (54)$$

we find it helpful to introduce the following useful sets:

$$C^{\text{distinct}}(l) := \left\{ \mathcal{I} = \{\mathcal{I}^1, \dots, \mathcal{I}^l\} \mid \mathcal{I}^1 \leq \dots \leq \mathcal{I}^l, \mathcal{I}^\tau \in \{0, 1, \dots, M\}^N \text{ and } \mathcal{I}^\tau \neq \mathcal{I}^{\tau+1} (\forall \tau) \right\}; \quad (55a)$$

$$C^{\text{distinct}} := \bigcup_{l \geq 1} C^{\text{distinct}}(l). \quad (55b)$$

In words, $C^{\text{distinct}}(l)$ can be viewed as the set of non-decreasing length- l paths in $\{0, 1, \dots, M\}^N$, with all points on a path being distinct; C^{distinct} thus consists of all such paths regardless of the length.

We first establish a connection between $|C|$ and $|C^{\text{distinct}}|$. Define the operator $\text{Proj} : C \rightarrow C^{\text{distinct}}$ that maps each $\mathcal{I} \in C$ to $\mathcal{I}^{\text{distinct}} \in C^{\text{distinct}}$, where $\mathcal{I}^{\text{distinct}}$ is composed of all distinct elements in \mathcal{I} (in other words, this operator simply removes redundancy in \mathcal{I}). Let us looking at the following set

$$\mathcal{B}(\mathcal{I}^{\text{distinct}}) := \{ \mathcal{I} \in C \mid \text{Proj}(\mathcal{I}) = \mathcal{I}^{\text{distinct}} \}$$

for each $\mathcal{I}^{\text{distinct}} \in C^{\text{distinct}}$. Since $\mathcal{I}^{\text{distinct}}$ is a non-decreasing path with all its points being distinct, there are at most $MN + 1$ elements in each $\mathcal{I}^{\text{distinct}}$. Hence, the size of $\mathcal{B}(\mathcal{I}^{\text{distinct}})$ is at most the number of solutions to the following equations

$$\sum_{i=1}^{MN+1} x_i = K \quad \text{and} \quad x_i \in \mathbb{N} \text{ for all } 1 \leq i \leq MN + 1.$$

Elementary combinatorial arguments then reveal that

$$|\mathcal{B}(\mathcal{I}^{\text{distinct}})| \leq \binom{K + MN}{MN} \leq (K + MN)^{MN} \leq (2K)^{MN},$$

for each $\mathcal{I}^{\text{distinct}}$, provided that $K \geq MN = SAH \log_2 K$. We then arrive at

$$|C| \leq |C^{\text{distinct}}| \cdot (2K)^{MN}. \quad (56)$$

Everything then boils down to bounding $|C^{\text{distinct}}|$. To do so, let us first look at the set $C^{\text{distinct}}(MN + 1)$, as each path in C^{distinct} cannot have length more than $MN + 1$. For each $\mathcal{I}^{\text{distinct}} = \{\tilde{\mathcal{I}}^1, \tilde{\mathcal{I}}^2, \dots, \tilde{\mathcal{I}}^{MN+1}\} \in C^{\text{distinct}}(MN + 1)$, it is easily seen that

$$-\tilde{\mathcal{I}}^1 = [0, 0, \dots, 0]^{\top} \text{ and } \tilde{\mathcal{I}}^{MN+1} = [M, M, \dots, M]^{\top}.$$

– For each $1 \leq \tau \leq MN$, $\tilde{\mathcal{I}}^{\tau}$ and $\tilde{\mathcal{I}}^{\tau+1}$ differ only in one element (i.e., their Hamming distance is 1).

In other words, we can view $\mathcal{I}^{\text{distinct}}$ as an MN -step path from $[0, 0, \dots, 0]^{\top}$ to $[M, M, \dots, M]^{\top}$, with each step moving in one dimension. Clearly, each step has at most N directions to choose from, meaning that there are at most N^{MN} such paths. This implies that

$$|C^{\text{distinct}}(MN + 1)| \leq N^{MN}.$$

To finish up, we further observe that for each $\mathcal{I}^{\text{distinct}} \in C^{\text{distinct}}$, there exists some $\tilde{\mathcal{I}}^{\text{distinct}} \in C^{\text{distinct}}(MN + 1)$ such that $\mathcal{I}^{\text{distinct}} \subseteq \tilde{\mathcal{I}}^{\text{distinct}}$. This observation together with basic combinatorial arguments indicates that

$$|C^{\text{distinct}}| \leq 2^{MN+1} |C^{\text{distinct}}(MN + 1)| \leq (2N)^{MN+1},$$

which taken collectively with (56) leads to the advertised bound

$$|C| \leq (2K)^{MN} |C^{\text{distinct}}| \leq (4KN)^{MN+1} \leq (4KN)^{MN+1}.$$

B.2 Proof of Lemma 4.4

Let us begin by considering any fixed total profile $\mathcal{I} \in C$, any fixed integer l obeying $2 \leq l \leq \log_2 K + 1$, and any given feasible sequence $\{X_{h,s,a}\}_{(s,a,h) \in S \times \mathcal{A} \times [H]}$. Recall that (i) $\hat{P}_{s,a,h}^{(l)}$ is computed based on the l th batch of data comprising 2^{l-2} independent samples from $\mathcal{D}^{\text{expand}}$ (see Definition 4.3); and (ii) each $X_{h+1,s,a}$ is given by a deterministic function of \mathcal{I} and the empirical models for steps $h' \in [h+1, H]$. Consequently, Lemma A.1 together with Definition 4.3 tells us that: with probability at least $1 - \delta'$, one has

$$\begin{aligned} & \sum_{s,a,h} \langle \hat{P}_{s,a,h}^{(l)} - P_{s,a,h}, X_{h+1,s,a} \rangle \\ & \leq \sqrt{\frac{8}{2^{l-2}} \sum_{s,a,h} \mathbb{V}(P_{s,a,h}, X_{h+1,s,a}) \log \frac{3 \log_2(SAHK)}{\delta'}} + \frac{4H}{2^{l-2}} \log \frac{3 \log_2(SAHK)}{\delta'}, \end{aligned} \quad (57)$$

where we view the left-hand side of (57) as a martingale sequence from $h = H$ back to $h = 1$.

Moreover, given that each $X_{h,s,a}$ has at most $K + 1$ different choices (since we assume $|\mathcal{X}_{h,\mathcal{I}}| \leq K + 1$), there are no more than $(K + 1)^{SAH} \leq (2K)^{SAH}$ possible choices of the feasible sequence $\{X_{h,s,a}\}_{(s,a,h) \in S \times \mathcal{A} \times [H]}$. In addition, it has been shown in Lemma 4.2 that there are no more than $(4SAHK)^{2SAH \log_2 K}$ possibilities of the total profile \mathcal{I} . Taking the union bound over all these choices

and replacing δ' in (57) with $\delta' / ((4SAHK)^{2SAH \log_2 K} (2K)^{SAH} \log_2 K)$, we can demonstrate that with probability at least $1 - \delta'$,

$$\begin{aligned}
& \sum_{s,a,h} \langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, X_{h+1,s,a} \rangle \\
& \leq \sqrt{\frac{8}{2^{l-2}} \sum_{s,a,h} \mathbb{V}(P_{s,a,h}, X_{h+1,s,a}) \left(2SAH \log_2 K \log(4SAHK) + SAH \log(2K) + \log \frac{3 \log_2^2(SAHK)}{\delta'} \right)} \\
& \quad + \frac{4H}{2^{l-2}} \left(2SAH \log_2 K \log(4SAHK) + SAH \log(2K) + \log \frac{3 \log_2^2(SAHK)}{\delta'} \right) \\
& \leq \sqrt{\frac{8}{2^{l-2}} \sum_{s,a,h} \mathbb{V}(P_{s,a,h}, X_{h+1,s,a}) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} + \frac{4H}{2^{l-2}} \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right), \quad (58)
\end{aligned}$$

holds simultaneously for all $\mathcal{I} \in \mathcal{C}$, all $2 \leq l \leq \log_2 K + 1$, and all feasible sequences $\{X_{h,s,a}\}_{(s,a,h) \in \mathcal{S} \times \mathcal{A} \times [H]}$.

Finally, recalling our assumption $0 \in \mathcal{X}_{h+1,\mathcal{I}}$, we see that for every total profile \mathcal{I} and its associated feasible sequence $\{X_{h,s,a}\}$,

$$\sum_{s,a,h} \max \left\{ \langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, X_{h+1,s,a} \rangle, 0 \right\} \in \left\{ \sum_{s,a,h} \langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, \tilde{X}_{h+1,s,a} \rangle \mid \tilde{X}_{h+1,s,a} \in \mathcal{X}_{h+1,\mathcal{I}}, \forall (s,a,h) \right\},$$

holds true. Consequently, the uniform upper bound on the right-hand side of (58) continues to be a valid upper bound on $\sum_{s,a,h} \max \{ \langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, X_{h+1,s,a} \rangle, 0 \}$. This concludes the proof.

B.3 Proof of Lemma 4.5

We begin by making the following claim, which we shall establish toward the end of this subsection.

CLAIM 1. *With probability exceeding $1 - \delta'$,*

$$\begin{aligned}
& \sum_{s,a,h} \left\langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, V_{h+1}^{k_{l,j,s,a,h}} \right\rangle \\
& \leq \sqrt{\frac{8}{2^{l-2}} \sum_{s,a,h} \mathbb{V}(P_{s,a,h}, V_{h+1}^{k_{l,j,s,a,h}}) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} + \frac{4H}{2^{l-2}} \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right), \quad (59)
\end{aligned}$$

holds simultaneously for all $l = 1, \dots, \log_2 K$ and all $j = 1, \dots, 2^{l-1}$, where $k_{l,j,s,a,h}$ stands for the episode index of the sample that visits (s, a, h) for the $(2^{l-1} + j)$ -th time in the online learning process.

Assuming the validity of Claim 1 for the moment, we can combine this claim with the decomposition (27) and applying the Cauchy–Schwarz inequality to reach

$$\begin{aligned}
\sum_{k=1}^K \sum_{h=1}^H \left\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \right\rangle &\leq \sum_{l=1}^{\log_2 K} \sum_{j=1}^{2^{l-1}} \sum_{s, a, h} \left\langle \widehat{P}_{s, a, h}^{(l)} - P_{s, a, h}, V_{h+1}^{k_{l, j, s, a, h}} \right\rangle + SAH^2 \\
&\leq \sum_{l=1}^{\log_2 K} \sum_{j=1}^{2^{l-1}} \sqrt{\frac{8}{2^{l-2}} \sum_{s, a, h} \mathbb{V}(P_{s, a, h}, V_{h+1}^{k_{l, j, s, a, h}}) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} \\
&\quad + \sum_{l=1}^{\log_2 K} \sum_{j=1}^{2^{l-1}} \frac{4H}{2^{l-2}} \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right) + SAH^2 \\
&\leq \sum_{l=1}^{\log_2 K} \sqrt{16 \sum_{j=1}^{2^{l-1}} \sum_{s, a, h} \mathbb{V}(P_{s, a, h}, V_{h+1}^{k_{l, j, s, a, h}}) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} \\
&\quad + \sum_{l=1}^{\log_2 K} 8H \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right) + SAH^2 \\
&\leq \sqrt{16(\log_2 K) \sum_{l=1}^{\log_2 K} \sum_{j=1}^{2^{l-1}} \sum_{s, a, h} \mathbb{V}(P_{s, a, h}, V_{h+1}^{k_{l, j, s, a, h}}) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} \\
&\quad + \left(48SAH^2 \log_2^3 K + 8H(\log_2 K) \log \frac{1}{\delta'} \right) + SAH^2 \\
&\leq \sqrt{16(\log_2 K) \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} \\
&\quad + 49SAH^2 \log_2^3 K + 8H(\log_2 K) \log \frac{1}{\delta'}.
\end{aligned}$$

Here, the last inequality is valid due to our assumption $V_{h+1}^k = 0$ ($\forall k > K$) and the identity

$$\begin{aligned}
&\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k) \\
&= \sum_{l=1}^{\log_2 K} \sum_{s, a, h} \sum_{j=1}^{2^{l-1}} \mathbb{V}(P_{s, a, h}, V_{h+1}^{k_{l, j, s, a, h}}) + \sum_{k=1}^K \sum_{h=1}^H \mathbb{1}\{N_h^{k, \text{all}}(s_h^k, a_h^k) = 0\} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k).
\end{aligned}$$

This establishes our advertised bound on $\sum_{k, h} \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle$, provided that Claim 1 is valid.

Before proceeding to the proof of Claim 1, we note that the other two quantities $\sum_{k, h} \max \{ \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle, 0 \}$ and $\sum_{k, h} \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, (V_{h+1}^k)^2 \rangle$ can be upper bounded using exactly the same arguments, which we omit for the sake of brevity. In particular,

ALGORITHM 2: Monotonic Value Propagation for a given total profile $\mathcal{I} \in C$ (MVP(\mathcal{I}))

```

1 initialization: set  $V_{H+1}^{k, \mathcal{I}}(s) \leftarrow H$  for all  $s \in \mathcal{S}$  and  $1 \leq k \leq K$ .
2 for  $k = 1, 2, \dots, K$  do
3   for  $h = H, H-1, \dots, 1$  do
4     for  $(s, a) \in \mathcal{S} \times \mathcal{A}$  do
5        $j \leftarrow l_{s, a, h}^k, n \leftarrow 2^{j-2},$ 
        $b_h(s, a) \leftarrow c_1 \sqrt{\frac{\mathbb{V}(\hat{P}_{s, a, h}^{(j)}, V_{h+1}^{k, \mathcal{I}}) \log \frac{1}{\delta'}}{\max\{n, 1\}}} + c_2 \sqrt{\frac{(\hat{\sigma}_h^{(j)}(s, a) - \hat{r}_h^{(j)}(s, a))^2 \log \frac{1}{\delta'}}{\max\{n, 1\}}} + c_3 \frac{H \log \frac{1}{\delta'}}{\max\{n, 1\}},$ 
        $Q_h^{k, \mathcal{I}}(s, a) \leftarrow \min \left\{ \hat{r}_h^{(j)}(s, a) + \langle \hat{P}_{s, a, h}^{(j)}, V_{h+1}^{k, \mathcal{I}} \rangle + b_h(s, a), H \right\},$ 
        $V_h^{k, \mathcal{I}}(s) \leftarrow \max_a Q_h^{k, \mathcal{I}}(s, a).$ 

```

the latter quantity further satisfies

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h=1}^H \left\langle \hat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, (V_{h+1}^k)^2 \right\rangle \\
& \leq \sqrt{16(\log_2 K) \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, (V_{h+1}^k)^2) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} \\
& \quad + 49SAH^3 \log_2^3 K + 8H^2(\log_2 K) \log \frac{1}{\delta'} \\
& \leq 8H \sqrt{(\log_2 K) \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k) \left(6SAH \log_2^2 K + \log \frac{1}{\delta'} \right)} \\
& \quad + 49SAH^3 \log_2^3 K + 8H^2(\log_2 K) \log \frac{1}{\delta'},
\end{aligned}$$

where the last inequality follows from Lemma A.2 and the fact that $0 \leq V_{h+1}^k(s) \leq H$ for all $s \in \mathcal{S}$.

PROOF OF CLAIM 1. To invoke Lemma 4.5 to prove this claim, we need to choose the set $\{\mathcal{X}_{h, \mathcal{I}}\}$ properly to include the true value function estimates $\{V_h^k\}$. To do so, we find it helpful to first introduce an auxiliary algorithm tailored to each total profile. Specifically, for each $\mathcal{I} \in C$ (cf. (26)), consider the following updates operating upon the expanded sample set $\mathcal{D}^{\text{expand}}$.

If we construct

$$\mathcal{X}_{h, \mathcal{I}} := \left\{ V_h^{k, \mathcal{I}} \mid 1 \leq k \leq K \right\} \cup \{0\}, \quad \forall h \in [H] \text{ and } \mathcal{I} \in C, \quad (60)$$

then it can be easily seen that $\{\mathcal{X}_{h, \mathcal{I}}\}$ satisfies the properties stated right before Lemma 4.4. As a consequence, applying Lemma 4.4 yields

$$\begin{aligned}
& \sum_{s,a,h} \left\langle \widehat{P}_{s,a,h}^{(l)} - P_{s,a,h}, X_{h+1,s,a} \right\rangle \\
& \leq \sqrt{\frac{8}{2^{l-2}} \sum_{s,a,h} \mathbb{V}(P_{s,a,h}, X_{h+1,s,a}) \left(6SAH \log_2 K + \log \frac{1}{\delta'} \right)} + \frac{4H}{2^{l-2}} \left(6SAH \log_2 K + \log \frac{1}{\delta'} \right),
\end{aligned} \tag{61}$$

simultaneously for all $l = 1, \dots, \log_2 K$, all $\mathcal{I} \in C$, and all sequences $\{X_{h,s,a}\}$ obeying $X_{h,s,a} \in \mathcal{X}_{h,\mathcal{I}}, \forall (s,a,h)$.

To finish up, denote by $\mathcal{I}^{\text{true}}$ the true total profile resulting from the online learning process. Given the way we couple $\mathcal{D}^{\text{expand}}$ and $\mathcal{D}^{\text{original}}$ (see the beginning of Section 4.2.2), we can easily see that the true value function estimate $\{V_h^k\}$ obeys

$$V_h^k = V_h^{k, \mathcal{I}^{\text{true}}} \in \mathcal{X}_{h, \mathcal{I}^{\text{true}}}, \quad 1 \leq k \leq K. \tag{62}$$

The claimed result then follows immediately from (62) and the uniform bound (61). \square

C Proofs of Auxiliary Lemmas in Section 5

C.1 Proof of Lemma 5.1

To begin with, we find it helpful to define the following function

$$f(p, v, n) := \langle p, v \rangle + \max \left\{ \frac{20}{3} \sqrt{\frac{\mathbb{V}(p, v) \log \frac{1}{\delta'}}{n}}, \frac{400}{9} \frac{H \log \frac{1}{\delta'}}{n} \right\},$$

for any vector $p \in \Delta^S$, any non-negative vector $v \in \mathbb{R}^S$ obeying $\|v\|_\infty \leq H$, and any positive integer n . We claim that

$$f(p, v, n) \text{ is non-decreasing in each entry of } v. \tag{63}$$

To justify this claim, consider any $1 \leq s \leq S$, and let us freeze p, n and all but the s th entries of v . It then suffices to observe that (i) f is a continuous function, and (ii) except for at most two possible choices of $v(s)$ that obey $\frac{20}{3} \sqrt{\frac{\mathbb{V}(p, v) \log \frac{1}{\delta'}}{n}} = \frac{400}{9} \frac{H \log \frac{1}{\delta'}}{n}$, one can use the properties of p and v to calculate

$$\begin{aligned}
\frac{\partial f(p, v, n)}{\partial v(s)} &= p(s) + \frac{20}{3} \mathbb{1} \left\{ \frac{20}{3} \sqrt{\frac{\mathbb{V}(p, v) \log \frac{1}{\delta'}}{n}} \geq \frac{400}{9} \frac{H \log \frac{1}{\delta'}}{n} \right\} \frac{p(s)(v(s) - \langle p, v \rangle) \sqrt{\log \frac{1}{\delta'}}}{\sqrt{n} \mathbb{V}(p, v)} \\
&= p(s) + \mathbb{1} \left\{ \sqrt{n} \mathbb{V}(p, v) \log \frac{1}{\delta'} \geq \frac{20}{3} H \log \frac{1}{\delta'} \right\} \frac{\frac{20}{3} H \log \frac{1}{\delta'}}{\sqrt{n} \mathbb{V}(p, v) \log \frac{1}{\delta'}} \cdot \frac{p(s)(v(s) - \langle p, v \rangle)}{H} \\
&\geq \min \left\{ p(s) + p(s) \frac{(v(s) - \langle p, v \rangle)}{H}, p(s) \right\} \\
&\geq p(s) \min \left\{ \frac{H + v(s) - \langle p, v \rangle}{H}, 1 \right\} \geq 0,
\end{aligned}$$

thus establishing the claim (63).

We now proceed to the proof of Lemma 5.1. Consider any (h, k, s, a) , and we divide into two cases.

Case 1: $N_h^k(s, a) \leq 2$. In this case, the following trivial bounds arise directly from the update rule (19):

$$Q_h^k(s, a) = H \geq Q_h^\star(s, a) \quad \text{and} \quad V_h^k(s) = H \geq V_h^\star(s).$$

Case 2: $N_h^k(s, a) > 2$. Suppose now that $Q_{h+1}^k \geq Q_{h+1}^\star$, which also implies that $V_{h+1}^k \geq V_{h+1}^\star$. If $Q_h^k(s, a) = H$, then $Q_h^k(s, a) \geq Q_h^\star(s, a)$ holds trivially, and hence it suffices to look at the case with $Q_h^k(s, a) < H$. According to the update rule in (19), it holds that

$$\begin{aligned} & Q_h^k(s, a) \\ &= \widehat{r}_h^k(s, a) + \langle \widehat{P}_{s, a, h}^k, V_{h+1}^k \rangle \\ & \quad + c_1 \sqrt{\frac{\mathbb{V}(\widehat{P}_{s, a, h}^k, V_{h+1}^k) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + c_2 \sqrt{\frac{(\widehat{\sigma}_h^k(s, a) - (\widehat{r}_h^k(s, a))^2) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + c_3 \frac{H \log \frac{1}{\delta'}}{N_h^k(s, a)} \\ & \geq \widehat{r}_h^k(s, a) + 2\sqrt{2} \sqrt{\frac{(\widehat{\sigma}_h^k(s, a) - (\widehat{r}_h^k(s, a))^2) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + \frac{48H \log \frac{1}{\delta'}}{3N_h^k(s, a)} + f(\widehat{P}_{s, a, h}^k, V_{h+1}^k, N_h^k(s, a)) \\ & \geq \widehat{r}_h^k(s, a) + 2\sqrt{2} \sqrt{\frac{(\widehat{\sigma}_h^k(s, a) - (\widehat{r}_h^k(s, a))^2) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + \frac{48H \log \frac{1}{\delta'}}{3N_h^k(s, a)} + f(\widehat{P}_{s, a, h}^k, V_{h+1}^\star, N_h^k(s, a)), \quad (64) \end{aligned}$$

for any (s, a) , where the last inequality results from the claim (63) and the property $V_{h+1}^k \geq V_{h+1}^\star$. Moreover, applying Lemma A.5 and recalling the definition of $\widehat{\sigma}_h^k(s, a)$, we have

$$\begin{aligned} & \mathbb{P} \left\{ \left| \langle \widehat{P}_{s, a, h}^k - P_{s, a, h}, V_{h+1}^\star \rangle \right| > 2 \sqrt{\frac{\mathbb{V}(\widehat{P}_{s, a, h}^k, V_{h+1}^\star) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + \frac{14H \log \frac{1}{\delta'}}{3N_h^k(s, a)} \right\} \\ & \leq \mathbb{P} \left\{ \left| \langle \widehat{P}_{s, a, h}^k - P_{s, a, h}, V_{h+1}^\star \rangle \right| > \sqrt{\frac{2\mathbb{V}(\widehat{P}_{s, a, h}^k, V_{h+1}^\star) \log \frac{1}{\delta'}}{N_h^k(s, a) - 1}} + \frac{7H \log \frac{1}{\delta'}}{3N_h^k(s, a) - 1} \right\} \leq 2\delta', \quad (65a) \end{aligned}$$

and

$$\mathbb{P} \left\{ \left| \widehat{r}_h^k(s, a) - r_h(s, a) \right| > 2 \sqrt{\frac{(\widehat{\sigma}_h^k(s, a) - (\widehat{r}_h^k(s, a))^2) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + \frac{28H \log \frac{1}{\delta'}}{3N_h^k(s, a)} \right\} \leq 2\delta'. \quad (65b)$$

These two inequalities imply that with probability exceeding $1 - 4\delta'$,

$$\begin{aligned} r_h(s, a) & \leq \widehat{r}_h^k(s, a) + 2\sqrt{2} \sqrt{\frac{(\widehat{\sigma}_h^k(s, a) - (\widehat{r}_h^k(s, a))^2) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + \frac{28H \log \frac{1}{\delta'}}{3N_h^k(s, a)}; \\ f(\widehat{P}_{s, a, h}^k, V_{h+1}^\star, N_h^k(s, a)) & = \langle P_{s, a, h}, V_{h+1}^\star \rangle + \langle \widehat{P}_{s, a, h}^k - P_{s, a, h}, V_{h+1}^\star \rangle \\ & \quad + \max \left\{ \frac{20}{3} \sqrt{\frac{\mathbb{V}(\widehat{P}_{s, a, h}^k, V_{h+1}^\star) \log \frac{1}{\delta'}}{N_h^k(s, a)}}, \frac{400}{9} \frac{H \log \frac{1}{\delta'}}{N_h^k(s, a)} \right\} \\ & \geq \langle P_{s, a, h}, V_{h+1}^\star \rangle. \end{aligned}$$

Substitution into (64) gives: with probability at least $1 - 4\delta'$,

$$Q_h^k(s, a) \geq r_h(s, a) + \langle P_{s, a, h}, V_{h+1}^* \rangle = Q_h^*(s, a).$$

Putting all this together. With the above two cases in place, one can invoke standard induction arguments to deduce that: with probability at least $1 - 4SAHK\delta'$, one has $Q_h^k(s, a) \geq Q_h^*(s, a)$ and $V_h^k = \max_a Q_h^k(s, a) \geq \max_a Q_h^*(s, a) = V_h^*(s)$ for every (s, a, h, k) . The proof is thus completed.

C.2 Proof of Lemma 5.3

C.2.1 Bounding T_2 . We first establish the bound (36a) on T_2 . To begin with, T_2 can be decomposed using the definition (19) of the bonus term:

$$\begin{aligned} T_2 &= \sum_{k=1}^K \sum_{h=1}^H b_h^k(s_h^k, a_h^k) \\ &= \frac{460}{9} \sum_{k, h} \sqrt{\frac{\mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k) \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} + 2\sqrt{2} \sum_{k, h} \sqrt{\frac{(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2) \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} \\ &\quad + \frac{544}{9} \sum_{k, h} \frac{H \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}. \end{aligned} \quad (66)$$

Applying the Cauchy–Schwarz inequality and invoking Lemma A.6, we obtain

$$\begin{aligned} T_2 &\leq \frac{460}{9} \sqrt{\sum_{k, h} \frac{\log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} \sqrt{\sum_{k, h} \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k)} \\ &\quad + 2\sqrt{2} \sqrt{\sum_{k, h} \frac{\log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} \sqrt{\sum_{k, h} (\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2)} + \frac{544H \log \frac{1}{\delta'}}{9} \sum_{k, h} \frac{1}{N_h^k(s_h^k, a_h^k)} \\ &\leq \frac{460}{9} \sqrt{2SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right) \sum_{k, h} \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k)} \\ &\quad + 4\sqrt{SAH(\log_2 K) \log \frac{1}{\delta'}} \sqrt{\sum_{k, h} (\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2)} + \frac{1,088}{9} SAH^2(\log_2 K) \log \frac{1}{\delta'}. \end{aligned} \quad (67)$$

Using the basic fact $\widehat{\sigma}_h^k(s_h^k, a_h^k) \leq H\widehat{r}_h^k(s, a)$ (since each immediate reward is at most H) and the definition (35a) of T_5 , we can continue the bound in (67) to derive

$$\begin{aligned} T_2 &\leq \frac{460}{9} \sqrt{2SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right) T_5} \\ &\quad + 4\sqrt{SAH^2(\log_2 K) \log \frac{1}{\delta'}} \sqrt{\sum_{k, h} \widehat{r}_h^k(s_h^k, a_h^k)} + \frac{1,088}{9} SAH^2(\log_2 K) \log \frac{1}{\delta'}. \end{aligned} \quad (68)$$

Applying Lemma A.7 to bound $\sum_{k,h} \widehat{r}_h^k(s_h^k, a_h^k)$ and using the basic fact $\sum_{k,h} r_h(s_h^k, a_h^k) \leq KH$, we can employ a little algebra to deduce that

$$T_2 \leq 61 \sqrt{2SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right)} T_5 + 8 \sqrt{SAH^3 K (\log_2 K) \log \frac{1}{\delta'}} + 155SAH^2 (\log_2 K) \log \frac{1}{\delta'}$$

with probability exceeding $1 - 2SAHK\delta'$.

C.2.2 Bounding T_3 . Next, let us prove the bound (36b) on $|T_3|$. Recall that $V_{h+1}^k(s)$ denotes the value function estimate of state s before the k th episode, which corresponds to the value estimate computed at the end of the *previous epoch*. This important fact implies that conditional on (s_h^k, a_h^k) , the vector $e_{s_{h+1}^k}$ is statistically independent of V_{h+1}^k and has conditional mean $P_{s_h^k, a_h^k, h}$, allowing us to invoke the Freedman inequality for martingales (see Lemma A.1) to control the sum of $\langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, V_{h+1}^k \rangle$. Recalling the definition of T_6 in (35b), we can see from Lemma A.1 that

$$|T_3| \leq 2\sqrt{2} \cdot \sqrt{T_6 \log \frac{1}{\delta'}} + \log \frac{1}{\delta'} + 2H \log \frac{1}{\delta'} \leq 2\sqrt{2} \cdot \sqrt{T_6 \log \frac{1}{\delta'}} + 3H \log \frac{1}{\delta'} \quad (69)$$

with probability at least $1 - 10SAH^2K^2\delta'$.

C.2.3 Bounding T_4 . We now turn attention to the bound (36c) on $|T_4|$. Recall that

$$T_4 = \sum_{k=1}^K \sum_{h=1}^H \left(\widehat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k) \right) + \sum_{k=1}^K \left(\sum_{h=1}^H r_h(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k) \right), \quad (70)$$

and we shall bound the two terms above separately.

– Regarding the first term on the right-hand side of (70), we can apply Lemma A.7 and the fact $\sum_{k,h} r_h(s_h^k, a_h^k) \leq KH$ to show that

$$\left| \sum_{k=1}^K \sum_{h=1}^H \left(\widehat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k) \right) \right| \leq 4 \sqrt{2SAH^3 K (\log_2 K) \log \frac{1}{\delta'}} + 52SAH^2 (\log_2 K) \log \frac{1}{\delta'} \quad (71)$$

holds with probability at least $1 - 2SAHK\delta'$.

– With regards to the second term on the right-hand side of (70), we note that conditional on π^k , $E_k := \sum_{h=1}^H r_h(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k)$ is a zero-mean random variable bounded in magnitude by H . According to Lemma A.1,

$$\begin{aligned} \left| \sum_{k=1}^K E_k \right| &\leq 2\sqrt{2} \cdot \sqrt{\sum_{k=1}^K \text{Var}(E_k) \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'} \\ &\leq 2\sqrt{2KH^2 \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'} \end{aligned} \quad (72)$$

holds with probability exceeding $1 - 4\delta' \log_2(KH)$, where $\text{Var}(E_k)$ denotes the variance of E_k conditioned on what happens before the k th episode, and the last inequality follows since $|E_k| \leq H$ always holds.

Substituting (71) and (72) into (70) reveals that with probability at least $1 - 3SAHK\delta'$,

$$|T_4| \leq 6 \sqrt{2SAH^3 K (\log_2 K) \log \frac{1}{\delta'}} + 55SAH^2 (\log_2 K) \log \frac{1}{\delta'}. \quad (73)$$

C.3 Proof of Lemma 5.4

Regarding the term T_5 , direct calculation gives

$$\begin{aligned}
T_5 &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k) = \sum_{k=1}^K \sum_{h=1}^H \left(\left\langle \widehat{P}_{s_h^k, a_h^k, h}^k, (V_{h+1}^k)^2 \right\rangle - \left(\langle \widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle \right)^2 \right) \\
&= \underbrace{\sum_{k=1}^K \sum_{h=1}^H \left\langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, (V_{h+1}^k)^2 \right\rangle}_{= T_7} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H \left\langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (V_{h+1}^k)^2 \right\rangle}_{= T_8} \\
&\quad + \sum_{k=1}^K \sum_{h=1}^H (V_{h+1}^k(s_{h+1}^k))^2 - \sum_{k=1}^K \sum_{h=1}^H \left(\langle \widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle \right)^2 \\
&= T_7 + T_8 + \sum_{k=1}^K \sum_{h=2}^H (V_h^k(s_h^k))^2 - \sum_{k=1}^K \sum_{h=1}^H \left(\langle \widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle \right)^2 \\
&\leq T_7 + T_8 + 2H \sum_{k=1}^K \sum_{h=1}^H \max \left\{ V_h^k(s_h^k) - \langle \widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle, 0 \right\} \\
&\leq T_7 + T_8 + 2H \sum_{k=1}^K \sum_{h=1}^H b_h^k(s_h^k, a_h^k) + 2H \sum_{k=1}^K \sum_{h=1}^H \widehat{r}_h^k(s_h^k, a_h^k), \tag{74} \\
&\leq T_7 + T_8 + 2HT_2 + 6KH^2, \tag{75}
\end{aligned}$$

with probability at least $1 - 3\delta' \log(KH^3)$. Here, the third line utilizes the fact that $V_{H+1}^k = 0$, the first inequality holds since

$$\begin{aligned}
(V_h^k(s_h^k))^2 - \left(\langle \widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle \right)^2 &= \left(V_h^k(s_h^k) + \langle \widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle \right) \left(V_h^k(s_h^k) - \langle \widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle \right) \\
&\leq 2H \max \left\{ V_h^k(s_h^k) - \langle \widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle, 0 \right\},
\end{aligned}$$

the penultimate line makes use of the property $V_h^k(s_h^k) = Q_h^k(s_h^k, a_h^k)$ and the update rule (20), whereas the last line applies property (52) and the definition (36a) of T_2 .

Akin to the above bound on T_5 , we can show that with probability at least $1 - 3SAHK\delta'$,

$$\begin{aligned}
T_6 &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k) = \sum_{k=1}^K \sum_{h=1}^H \left\langle P_{s_h^k, a_h^k, h}, (V_{h+1}^k)^2 \right\rangle - \sum_{k=1}^K \sum_{h=1}^H \left(\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle \right)^2 \\
&= \sum_{k=1}^K \sum_{h=1}^H \left\langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (V_{h+1}^k)^2 \right\rangle + \sum_{k=1}^K \sum_{h=2}^H (V_h^k(s_h^k))^2 - \sum_{k=1}^K \sum_{h=1}^H \left(\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle \right)^2 \\
&\leq T_8 + 2H \sum_{k=1}^K \sum_{h=1}^H \max \left\{ V_h^k(s_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle, 0 \right\}
\end{aligned}$$

$$\begin{aligned}
&\leq T_8 + 2H \sum_{k=1}^K \sum_{h=1}^H \max \left\{ V_h^k(s_h^k) - \langle \hat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle, 0 \right\} \\
&\quad + 2H \sum_{k=1}^K \sum_{h=1}^H \max \left\{ \langle \hat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle, 0 \right\} \\
&\leq T_8 + 2H \sum_{k=1}^K \sum_{h=1}^H b_h^k(s_h^k, a_h^k) + 2H \sum_{k=1}^K \sum_{h=1}^H \hat{r}_h^k(s_h^k, a_h^k) + 2HT_9, \\
&\leq T_8 + 2HT_2 + 6KH^2 + 2HT_9. \tag{76}
\end{aligned}$$

Finally, note that the above bounds on T_5 and T_6 both depend on the term T_8 (cf. (37b)), which we would like to cope with now. Using Freedman's inequality (cf. Lemma A.1) and the fact that $\text{Var}(X^2) \leq 4H^2\text{Var}(X)$ for any random variable X with support on $[-H, H]$ (cf. Lemma A.2), we reach

$$|T_8| \leq 2\sqrt{2} \sqrt{\sum_{k,h} \mathbb{V}(\hat{P}_{s_h^k, a_h^k, h}^k, (V_{h+1}^k)^2) \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'} \leq \sqrt{32H^2 T_6 \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'} \tag{78}$$

with probability at least $1 - 3\delta' \log(KH^3)$. Substitution into (75) and (77) establishes (38).

D Proof of the Value-based Regret Bound (Proof of Theorem 1.2)

Recall that

$$B = 4,000(\log_2 K)^3 \log(3SAH) \log \frac{1}{\delta'} \quad \text{with } \delta' = \frac{\delta}{200SAH^2K^2}. \tag{79}$$

Consider first the scenario where $K \leq \frac{BSAH^2}{v^*}$: the regret bound can be upper bounded by

$$\begin{aligned}
\mathbb{E}[\text{Regret}(K)] &= \mathbb{E} \left[\sum_{k=1}^K \left(V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k) \right) \right] \leq \mathbb{E} \left[\sum_{k=1}^K V_1^*(s_1^k) \right] = K \mathbb{E}_{s_1 \sim \mu} [V_1^*(s_1)] \\
&= Kv^* = \min \left\{ \sqrt{BSAH^2Kv^*}, Kv^* \right\}. \tag{80}
\end{aligned}$$

As a result, the remainder of the proof is dedicated to the case with

$$K \geq \frac{BSAH^2}{v^*}. \tag{81}$$

To begin with, recall that the proof of Theorem 1.1 in Section 5 consists of bounding the quantities T_1, \dots, T_9 (see (34), (35), and (37)) and recall that $\delta' = \frac{\delta}{200SAH^2K^2}$. In order to establish Theorem 1.2, we need to develop tighter bounds on some of these quantities (i.e., T_2 , T_4 , T_5 , and T_6) to reflect their dependency on v^* (cf. (43)).

Bounding T_2 . Recall that we have shown in (68) that

$$\begin{aligned}
T_2 &\leq \frac{460}{9} \sqrt{2SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right) T_5} \\
&\quad + 4 \sqrt{SAH^2(\log_2 K) \log \frac{1}{\delta'}} \sqrt{\sum_{k,h} \hat{r}_h^k(s_h^k, a_h^k)} + \frac{1088}{9} SAH^2(\log_2 K) \log \frac{1}{\delta'}.
\end{aligned}$$

In view of the definition of T_4 (cf. (34)) as well as the fact that $\sum_{k=1}^K V_1^*(s_1^k) \leq 3Kv^* + H \log \frac{1}{\delta'}$ holds with probability at least $1 - \delta'$ (see Lemma A.3), we arrive at

$$\sum_{k,h} \hat{r}_h^k(s_h^k, a_h^k) \leq T_4 + \sum_k V_1^{\pi^k}(s_1^k) \leq T_4 + \sum_k V_1^*(s_1^k) \leq T_4 + 3Kv^* + H \log \frac{1}{\delta'}, \tag{82}$$

which in turn gives

$$T_2 \leq \frac{460}{9} \sqrt{2SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right) T_5} + 4 \sqrt{SAH^2(\log_2 K) \log \frac{1}{\delta'} \sqrt{T_4 + 3Kv^*} + 130SAH^2(\log_2 K) \log \frac{1}{\delta'}}. \quad (83)$$

Bounding T_4 . When it comes to the quantity T_4 (cf. (34)), we make the observation that

$$T_4 = \underbrace{\sum_{k=1}^K \left(\sum_{h=1}^H \hat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k) \right)}_{=: \check{T}_1} + \underbrace{\sum_{k=1}^K \left(\sum_{h=1}^H r_h(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k) \right)}_{=: \check{T}_2}. \quad (84)$$

Repeating the arguments for (82) yields

$$\sum_{k,h} r_h(s_h^k, a_h^k) \leq \check{T}_2 + \sum_k V_1^{\pi^k}(s_1^k) \leq \check{T}_2 + \sum_k V_1^*(s_1^k) \leq \check{T}_2 + 3Kv^* + H \log \frac{1}{\delta'} \quad (85)$$

with probability at least $1 - \delta'$. Combining this with Lemma A.7, we see that

$$\begin{aligned} \check{T}_1 &\leq 4 \sqrt{2SAH^2 \log_2 K \log \frac{1}{\delta'} \sqrt{\sum_{k=1}^K \sum_{h=1}^H r_h(s_h^k, a_h^k) + 52SAH^2(\log_2 K) \log \frac{1}{\delta'}}} \\ &\leq 4 \sqrt{2SAH^2 \log_2 K \log \frac{1}{\delta'} \sqrt{\check{T}_2 + 3Kv^*} + 60SAH^2(\log_2 K) \log \frac{1}{\delta'}} \end{aligned} \quad (86)$$

with probability exceeding $1 - 3SAHK\delta'$. In addition, Lemma A.1 tells us that

$$\begin{aligned} \check{T}_2 &\leq 2 \sqrt{2 \sum_{k=1}^K \mathbb{E}_{\pi^k, s_1 \sim \mu} \left[\left(\sum_{h=1}^H r_h(s_h, a_h) \right)^2 \right] \log \frac{1}{\delta'} + 3H^2 \log \frac{1}{\delta'}} \\ &\leq 2 \sqrt{2H \sum_{k=1}^K \mathbb{E}_{\pi^k, s_1 \sim \mu} \left[\sum_{h=1}^H r_h(s_h, a_h) \right] \log \frac{1}{\delta'} + 3H \log \frac{1}{\delta'}} \\ &\leq 2 \sqrt{2KHv^* \log \frac{1}{\delta'} + 3H \log \frac{1}{\delta'}}, \end{aligned} \quad (87)$$

$$\leq 2Kv^* + 5H \log \frac{1}{\delta'}, \quad (88)$$

with probability at least $1 - 2SAHK\delta'$, where the expectation operator $\mathbb{E}_{\pi^k, s_1 \sim \mu}[\cdot]$ is taken over the randomness of a trajectory $\{(s_h, a_h)\}$ generated under policy π^k and initial state $s_1 \sim \mu$, the last line arises from the AM-GM inequality, and the penultimate line makes use of Assumption 1 and the fact that

$$\mathbb{E}_{\pi^k, s_1 \sim \mu} \left[\sum_{h=1}^H r_h(s_h, a_h) \right] = \mathbb{E}_{s_1 \sim \mu} \left[V_1^{\pi^k}(s_1) \right] \leq \mathbb{E}_{s_1 \sim \mu} \left[V_1^*(s_1) \right] = v^*.$$

Taking (86), (87), and (88) together, we can demonstrate that with probability exceeding $1 - 5SAHK\delta'$,

$$\check{T}_1 \leq 13\sqrt{SAH^2Kv^*(\log_2 K) \log \frac{1}{\delta'}} + 80SAH^2(\log_2 K) \log \frac{1}{\delta'}, \quad (89a)$$

$$\check{T}_2 \leq 2\sqrt{2KHv^* \log \frac{1}{\delta'}} + 3H \log \frac{1}{\delta'}. \quad (89b)$$

Substitution into (84) reveals that: with probability exceeding $1 - 5SAHK\delta'$,

$$T_4 \leq 15\sqrt{SAH^2Kv^*(\log_2 K) \log \frac{1}{\delta'}} + 83SAH^2(\log_2 K) \log \frac{1}{\delta'}. \quad (90)$$

Bounding T_5 . Recall that we have proven in (74) that

$$T_5 \leq T_7 + T_8 + 2HT_2 + 2H \sum_{k=1}^K \sum_{h=1}^H \hat{r}_h^k(s_h^k, a_h^k). \quad (91)$$

With (85) and (88) in place, we can deduce that, with probability at least $1 - 3SAHK\delta'$,

$$\sum_{k,h} r_h(s_h^k, a_h^k) \leq \check{T}_2 + 3Kv^* + H \log \frac{1}{\delta'} \leq 5Kv^* + 6H \log \frac{1}{\delta'}. \quad (92)$$

Moreover, under the assumption (81), we can further bound (89a) as

$$\check{T}_1 \leq \sqrt{BSAH^2Kv^*} + BSAH^2 \leq 2Kv^*,$$

with probability exceeding $1 - 3SAHK\delta'$, which combined with (92) and the assumption (81) results in

$$\sum_{k,h} \hat{r}_h^k(s_h^k, a_h^k) = \sum_{k,h} r_h(s_h^k, a_h^k) + \check{T}_1 \leq 7Kv^* + 6H \log \frac{1}{\delta'} \leq 8Kv^*. \quad (93)$$

Substitution into (91) indicates that: with probability exceeding $1 - 6SAHK\delta'$,

$$T_5 \leq T_7 + T_8 + 2HT_2 + 16HKv^*. \quad (94)$$

Bounding T_6 . Making use of our bounds (76), (38c), and (93), we can readily derive

$$\begin{aligned} T_6 &\leq T_8 + 2HT_2 + 2HT_9 + 2H \sum_{k=1}^K \sum_{h=1}^H \hat{r}_h^k(s_h^k, a_h^k) \\ &\leq \sqrt{32T_6 \log \frac{1}{\delta'}} + 2HT_9 + 16HKv^* + 3H^2 \log \frac{1}{\delta'} + 2HT_2, \end{aligned} \quad (95)$$

with probability at least $1 - 16SAH^2K^2\delta'$.

Putting all the pieces together. Recalling our choice of B (cf. (79)), we can see from (83), (36b), (90), (94), (95), (38c), (39a), and (39b) that

$$T_2 \leq \sqrt{BSAHT_5} + \sqrt{BSAH^2(T_4 + 3Kv^*)} + BSAH^2, \quad (96a)$$

$$T_3 \leq \sqrt{BT_6} + BH, \quad (96b)$$

$$T_4 \leq \sqrt{BSAH^2Kv^*} + BSAH^2, \quad (96c)$$

$$T_5 \leq T_7 + T_8 + 2HT_2 + 16HKv^*, \quad (96d)$$

$$T_6 \leq \sqrt{BT_6} + 2HT_9 + 16HKv^* + BH^2 + 2HT_2, \quad (96e)$$

$$T_8 \leq \sqrt{BH^2T_6} + BH^2, \quad (96f)$$

$$T_1 \leq T_9 \leq \sqrt{BSAHT_6} + BSAH^2, \quad (96g)$$

$$T_7 \leq H\sqrt{BSAHT_6} + BSAH^3. \quad (96h)$$

Solving (96) under the assumption $K \geq \frac{BSAH^2}{v^*}$ allows us to demonstrate that

$$T_6 \lesssim BHKv^*, \quad (97a)$$

$$T_1 \leq T_9 \lesssim \sqrt{B^2SAH^2Kv^*}, \quad (97b)$$

$$T_7 + T_8 \lesssim \sqrt{B^2SAH^4Kv^*}, \quad (97c)$$

$$T_5 \lesssim BHKv^*, \quad (97d)$$

$$T_2 \lesssim \sqrt{B^2SAH^2Kv^*}, \quad (97e)$$

$$T_3 \lesssim \sqrt{B^2HKv^*}, \quad (97f)$$

$$T_4 \lesssim \sqrt{BSAH^2Kv^*}, \quad (97g)$$

with probability exceeding $1 - 200SAH^2K^2\delta'$. Putting these bounds together with (34), we arrive at

$$\text{Regret}(K) \leq T_1 + T_2 + T_3 + T_4 \lesssim B\sqrt{SAH^2Kv^*},$$

with probability exceeding $1 - 200SAH^2K^2\delta'$. Replacing δ' with $\frac{\delta}{200SAH^2K^2}$ and taking $\delta = \frac{1}{2KH}$ give

$$\begin{aligned} \mathbb{E}[\text{Regret}(K)] &\lesssim (1 - \delta)B\sqrt{SAH^2Kv^*} + \delta Kv^* \lesssim B\sqrt{SAH^2Kv^*} + 1 \asymp B\sqrt{SAH^2Kv^*} \\ &\asymp \min\{B\sqrt{SAH^2Kv^*}, BKv^*\} \asymp \min\{\sqrt{SAH^2Kv^*}, Kv^*\} \log^5(SAHK), \end{aligned}$$

provided that $K \geq \frac{BSAH^2}{v^*}$. Taking this collectively with (80) concludes the proof.

E Proof of the Cost-based Regret Bound (Proof of Theorem 1.3)

We now turn to the proof of Theorem 1.3. For notational convenience, we shall use r to denote the negative cost (namely, $r_h = -c_h$, $\hat{r}_h = -\hat{c}_h$, and so on) throughout this section. We shall also use the following notation (and similar quantities like Q_h^k , V_h^k , \dots)

$$\begin{aligned} Q_h(s, a) &\leftarrow \max \left\{ \min \left\{ \hat{r}_h(s, a) + \hat{P}_{s, a, h}V_{h+1} + b_h(s, a), 0 \right\}, -H \right\}, \\ V_h(s) &\leftarrow \max_a Q_h(s, a), \end{aligned}$$

in order to be consistent with the reward-based setting.

Akin to the proof of Theorem 1.2, we need to bound the quantities T_1, \dots, T_9 introduced previously (see (34), (35), and (37)). We note that the analysis for T_1, T_3, T_7, T_8 , and T_9 in Appendix D readily applies to the negative reward case herein. Thus, it suffices to develop bounds on T_2, T_4, T_5 , and T_6 to capture their dependency on c^* , which forms the main content of the remainder of this section.

Bounding T_2 . Recall from (66) that

$$T_2 = \frac{460}{9} \sum_{k,h} \sqrt{\frac{\mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k) \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} + 2\sqrt{2} \sum_{k,h} \sqrt{\frac{(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2) \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} + \frac{544}{9} \sum_{k,h} \frac{H \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}. \quad (98)$$

In what follows, let us bound the three terms on the right-hand side of (98) separately.

– For the first and the third terms on the right-hand side of (98), invoking the Cauchy–Schwarz inequality and Lemma A.6 gives

$$\begin{aligned} \sum_{k,h} \sqrt{\frac{\mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k) \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} &\leq \sqrt{2SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right) \sum_{k,h} \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k)} \\ &= \sqrt{2SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right) T_5}, \end{aligned} \quad (99)$$

with T_5 defined in (35a), and in addition,

$$\sum_{k,h} \frac{H \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)} \leq 2SAH^2(\log_2 K) \log \frac{1}{\delta'}. \quad (100)$$

– Let us turn to the second term on the right-hand side of (98). Observing the basic fact that

$$\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \leq -H\widehat{r}_h^k(s_h^k, a_h^k),$$

we can combine it with Lemma A.6 to derive

$$\begin{aligned} \sqrt{\frac{(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2) \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} &\leq \sqrt{2SAH(\log_2 K) \log \frac{1}{\delta'}} \sqrt{H \sum_{k,h} -\widehat{r}_h^k(s_h^k, a_h^k)} \\ &\leq \sqrt{2SAH^2(\log_2 K) \log \frac{1}{\delta'}} \sqrt{-T_4 + 3Kc^* + \sum_{k=1}^K \left(-V_1^{\pi^k}(s_1^k) + V_1^*(s_1^k) \right) + \sum_{k=1}^K \left(-V_1^*(s_1^k) - 3c^* \right)}, \end{aligned} \quad (101)$$

where the last inequality invokes the definition of T_4 (see (34)). By virtue of Lemma A.3 and the definition (45) of c^* , one can show that

$$\sum_{k=1}^K -V_1^*(s_1^k) \leq 3Kc^* + H \log \frac{1}{\delta'},$$

with probability exceeding $1 - \delta'$. In addition, we note that

$$\sum_{k=1}^K \left(-V_1^{\pi^k}(s_1^k) + V_1^*(s_1^k) \right) = \text{Regret}(K) = T_1 + T_2 + T_3 + T_4. \quad (102)$$

Taking these properties together with (101) yields

$$\begin{aligned} & \sqrt{\frac{\left(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \right) \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} \\ & \leq \sqrt{2SAH^2(\log_2 K) \log \frac{1}{\delta'}} \sqrt{T_1 + T_2 + T_3 + 2|T_4| + 3Kc^* + H \log \frac{1}{\delta'}}. \end{aligned}$$

Putting the above results together, we can deduce that, with probability exceeding $1 - \delta'$,

$$\begin{aligned} T_2 & \leq 90 \sqrt{SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right) T_5} \\ & + 4 \sqrt{SAH^2(\log_2 K) \log \frac{1}{\delta'}} \sqrt{T_1 + T_2 + T_3 + 2|T_4| + 3Kc^* + H \log \frac{1}{\delta'}} + 130SAH^2(\log_2 K) \log \frac{1}{\delta'}. \end{aligned} \quad (103)$$

Bounding T_4 . When it comes to the quantity T_4 , we recall that

$$T_4 = \underbrace{\sum_{k=1}^K \left(\sum_{h=1}^H \widehat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k) \right)}_{=: \check{T}_1} + \underbrace{\sum_{k=1}^K \left(\sum_{h=1}^H r_h(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k) \right)}_{=: \check{T}_2}. \quad (104)$$

To control T_4 , we first make note of the following result that bounds the empirical reward (for the case with negative rewards), which assists in bounding the term \check{T}_1 .

LEMMA E.1. *With probability at least $1 - 2SAHK\delta'$, it holds that*

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H \left| \widehat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k) \right| \\ & \leq 4 \sqrt{2SAH^2(\log_2 K) \log \frac{1}{\delta'}} \cdot \sqrt{\sum_{k=1}^K \sum_{h=1}^H (-r_h(s_h^k, a_h^k))} + 52SAH^2(\log_2 K) \log \frac{1}{\delta'}. \end{aligned}$$

PROOF. The proof basically follows the same arguments as in the proof of Lemma A.7, except that r is now replaced with $-r$. \square

Lemma E.1 tells us that with probability at least $1 - 3SAHK\delta'$,

$$\begin{aligned} |\check{T}_1| & \leq 4 \sqrt{2SAH^2(\log_2 K) \log \frac{1}{\delta'}} \cdot \sqrt{\sum_{k=1}^K \sum_{h=1}^H (-r_h(s_h^k, a_h^k))} + 52SAH^2(\log_2 K) \log \frac{1}{\delta'} \\ & \leq 4 \sqrt{2SAH^2(\log_2 K)} \cdot \sqrt{-\check{T}_2 + 3Kc^* + \sum_{k=1}^K (-V_1^*(s_1^k) - 3c^*)} + 52SAH^2(\log_2 K) \log \frac{1}{\delta'} \\ & \leq 4 \sqrt{2SAH^2(\log_2 K)} \cdot \sqrt{\check{T}_2 + 3Kc^* + 60SAH^2(\log_2 K) \log \frac{1}{\delta'}}. \end{aligned} \quad (105)$$

Here, the last line uses the fact (see Lemma A.3) that, with probability exceeding $1 - \delta'$,

$$\sum_{k=1}^K (-V_1^\star(s_1^k)) \leq 3Kc^\star + H \log \frac{1}{\delta'}. \quad (106)$$

In addition, the Freedman inequality in Lemma A.1 combined with (106) implies that, with probability at least $1 - 3SAHK\delta$,

$$\begin{aligned} |\check{T}_2| &\leq 2 \sqrt{2 \sum_{k=1}^K \mathbb{E}_{\pi^k} \left[\left(\sum_{h=1}^H r_h(s_h, a_h) \right)^2 \middle| s_1 = s_1^k \right] \log \frac{1}{\delta} + 3H \log \frac{1}{\delta}} \\ &\leq 2 \sqrt{2H \sum_{k=1}^K \mathbb{E}_{\pi^k} \left[\sum_{h=1}^H -r_h(s_h, a_h) \middle| s_1 = s_1^k \right] \log \frac{1}{\delta} + 3H \log \frac{1}{\delta}} \\ &= 2 \sqrt{2H \left(\sum_{k=1}^K \left(-V_1^{\pi^k}(s_1^k) + V_1^\star(s_1^k) \right) + \sum_{k=1}^K \left(-V_1^\star(s_1^k) - 3c^\star \right) + 3Kc^\star \right) \log \frac{1}{\delta} + 3H \log \frac{1}{\delta}}, \quad (107) \end{aligned}$$

$$\leq 3Kc^\star + T_1 + T_2 + T_3 + T_4 + 9H \log \frac{1}{\delta}. \quad (108)$$

Combining (105), (107), with (108) reveals that, with probability at least $1 - 4SAHK\delta$,

$$\begin{aligned} |\check{T}_1| &\leq 16 \sqrt{SAH^2(Kc^\star + T_1 + T_2 + T_3 + T_4)(\log_2 K) \log \frac{1}{\delta} + 200SAH^2(\log_2 K) \log \frac{1}{\delta}} \\ |\check{T}_2| &\leq 2 \sqrt{2H(3Kc^\star + T_1 + T_2 + T_3 + T_4) \log \frac{1}{\delta} + 9H \log \frac{1}{\delta}}. \end{aligned}$$

As a result, substitution into (104) leads to

$$|T_4| \leq 22 \sqrt{SAH^2(Kc^\star + T_1 + T_2 + T_3 + T_4)(\log_2 K) \log \frac{1}{\delta} + 209SAH^2(\log_2 K) \log \frac{1}{\delta}}. \quad (109)$$

Bounding T_5 . Invoking the arguments in (38a) and recalling the update rule (46), we obtain

$$\begin{aligned} T_5 &\leq \sum_{k=1}^K \sum_{h=1}^H \left\langle \hat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}, (V_{h+1}^k)^2 \right\rangle + \sum_{k=1}^K \sum_{h=1}^H \left\langle P_{s_h^k, a_h^k, h} - e_{s_h^k}, (V_{h+1}^k)^2 \right\rangle \\ &\quad + 2H \sum_{k=1}^K \sum_{h=1}^H \left[-r_h(s_h^k, a_h^k) \right]. \end{aligned} \quad (110)$$

Moreover, we recall that

$$\sum_{k=1}^K \sum_{h=1}^H \left[-r_h(s_h^k, a_h^k) \right] = -\check{T}_2 - \sum_{k=1}^K V_1^{\pi^k}(s_1) \leq -\check{T}_2 + \sum_{k=1}^K V_1^\star(s_1^k). \quad (111)$$

By virtue of (106), one sees that with probability at least $1 - 5SAHK\delta$,

$$\sum_{k=1}^K \sum_{h=1}^H \left[-r_h(s_h^k, a_h^k) \right] \leq 2 \sqrt{2H(3Kc^\star + T_1 + T_2 + T_3 + T_4) \log \frac{1}{\delta} + 3Kc^\star + 10H \log \frac{1}{\delta}}. \quad (112)$$

Consequently, we arrive at

$$T_5 \leq T_7 + T_8 + 2HT_2 + 4\sqrt{2H^3(3Kc^* + T_1 + T_2 + T_3 + T_4)\log\frac{1}{\delta}} + 6HKc^* + 20H^2\log\frac{1}{\delta}, \quad (113)$$

with probability exceeding $1 - 5SAHK\delta$.

Bounding T_6 . Invoking the arguments in (38a), (106) and (111), and recalling the update rule (46), we can demonstrate that

$$\begin{aligned} T_6 &\leq 2\sqrt{8T_6\log\frac{1}{\delta}} + 3H^2\log\frac{1}{\delta} + 2H\sum_{k=1}^K\sum_{h=1}^H \max\{\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle - V_h^k(s_h^k), 0\} \\ &\leq 2\sqrt{8T_6\log\frac{1}{\delta}} + 3H^2\log\frac{1}{\delta} + 2HT_9 + 2H\sum_{k=1}^K\sum_{h=1}^H [-r_h(s_h^k, a_h^k)] \\ &\leq 2\sqrt{8T_6\log\frac{1}{\delta}} + 3H^2\log\frac{1}{\delta} + 2HT_9 \\ &\quad + 2H\left(2\sqrt{2H(3Kc^* + T_1 + T_2 + T_3 + T_4)\log\frac{1}{\delta}} + 3Kc^* + 10H\log\frac{1}{\delta}\right) \end{aligned} \quad (114)$$

with probability at least $1 - 3SAHK\delta$.

Putting all this together. Armed with the preceding bounds, we are ready to establish the claimed regret bound. By solving (103), (36b), (109), (113), (114), (38c), (39a), and (39b), we can show that, with probability exceeding $1 - 100SAH^2K\delta$,

$$\begin{aligned} T_6 &\lesssim HKc^* + BSAH^3, \\ T_1 &\lesssim \sqrt{BSAH^2Kc^*} + BSAH^2, \\ T_7 + T_8 &\lesssim \sqrt{BSAH^4Kc^*} + BSAH^3, \\ T_5 &\lesssim HKc^* + BSAH^2, \\ T_2 &\lesssim \sqrt{BSAH^2Kc^*} + BSAH^2, \\ T_3 &\lesssim \sqrt{BHKc^*} + BSAH^2. \end{aligned}$$

We then readily conclude that the total regret is bounded by

$$O(\sqrt{BSAH^2Kc^*} + BSAH^2).$$

In addition, the regret bound is trivially upper bounded by $O(K(H - c^*))$. The proof is thus completed by combining these two regret bounds and replacing δ' with $\frac{\delta}{100SAH^2K}$.

F Proof of the Variance-dependent Regret Bounds (Proof of Theorem 1.4)

In this section, we turn to establishing Theorem 1.4. The proof primarily contains two parts, as summarized in the following lemmas.

LEMMA F.1. *With probability exceeding $1 - \delta/2$, Algorithm 1 obeys*

$$\text{Regret}(K) \leq \tilde{O}\left(\min\{\sqrt{SAHK\text{var}_1} + SAH^2, KH\}\right).$$

LEMMA F.2. *With probability at least $1 - \delta/2$, Algorithm 1 satisfies*

$$\text{Regret}(K) \leq \tilde{O}\left(\min\{\sqrt{SAHK\text{var}_2} + SAH^2, KH\}\right).$$

Putting these two regret bounds together and rescaling δ to $\delta/2$, we immediately conclude the proof of Theorem 1.4. The remainder of this section is thus devoted to establishing Lemma F.1 and Lemma F.2.

F.1 Proof of Lemma F.1

Before proceeding, we recall that

$$\begin{aligned} T_4 &= \sum_{k=1}^K \left(\sum_{h=1}^H \widehat{r}_h^k(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k) \right), \\ T_5 &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}, V_{h+1}^k), \\ T_6 &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k), \end{aligned}$$

and that

$$B = 4,000(\log_2 K)^3 \log(3SAH) \log \frac{1}{\delta'} \quad \text{and} \quad \delta' = \frac{\delta}{200SAH^2K^2}.$$

F.1.1 Bounding T_2 . Recall that when proving (36a), we have demonstrated that (see (67))

$$\begin{aligned} T_2 &\leq \frac{460}{9} \sqrt{2SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right) T_5} \\ &\quad + 4 \sqrt{SAH(\log_2 K) \log \frac{1}{\delta'}} \sqrt{\sum_{k,h} \left(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \right)} + \frac{1,088}{9} SAH^2(\log_2 K) \log \frac{1}{\delta'}. \end{aligned} \tag{115}$$

This motivates us to bound the sum $\sum_{k,h} (\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2)$, which we accomplish via the following lemma.

LEMMA F.3. *With probability at least $1 - 4SAHK\delta'$, one has*

$$\sum_{k,h} \left(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \right) \leq 6K\text{var}_1 + 242SAH^3(\log_2 K) \log \frac{1}{\delta'}. \tag{116}$$

Combining Lemma F.3 with (115), we can readily derive

$$\begin{aligned} T_2 &\leq \frac{460}{9} \sqrt{2SAH(\log_2 K) \left(\log \frac{1}{\delta'} \right) T_5} + 12 \sqrt{SAH(\log_2 K) \log \frac{1}{\delta'}} \sqrt{2K\text{var}_1} \\ &\quad + 157SAH^2(\log_2 K) \log \frac{1}{\delta'} \end{aligned} \tag{117}$$

with probability at least $1 - 4SAHK\delta'$.

PROOF OF LEMMA F.3. For notational convenience, let us define the variance of $R_h(s, a)$ as $v_h(s, a)$.

Firstly, we control each $\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2$ with $v_h(s, a)$. Fix (s, a, h, k) . Applying Lemma A.3 shows that, with probability at least $1 - 2\delta'$,

$$N_h^k(s, a) \left(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \right) \leq 3N_h^k v_h(s, a) + H^2 \log \frac{1}{\delta'}. \tag{118}$$

This allows us to deduce that, with probability exceeding $1 - 2SAHK\delta'$,

$$\begin{aligned} \sum_{k,h} \left(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \right) &\leq 3 \sum_{k,h} v_h(s_h^k, a_h^k) + \sum_{k,h} \frac{H^2 \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)} \\ &\leq 3 \sum_{k,h} v_h(s_h^k, a_h^k) + 2SAH^3(\log_2 K) \log \frac{1}{\delta'}. \end{aligned} \quad (119)$$

It then suffices to bound the sum $\sum_{k,h} v_h(s_h^k, a_h^k)$. Toward this end, let

$$\widetilde{V}_h^k(s) := \mathbb{E}_{\pi^k} \left[\sum_{h'=h}^H v_{h'}(s_{h'}, a_{h'}) \mid s_h = s \right],$$

be the value function with rewards taken to be $\{v_h(s, a)\}$ and the policy selected as π^k . It is clearly seen that

$$\widetilde{V}_h^k(s, a) \leq H^2.$$

In view of Lemma A.1, we can obtain

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H v_h(s_h^k, a_h^k) - \sum_{k=1}^K \widetilde{V}_1^k(s_1^k) &= \sum_{k=1}^K \left(\sum_{h=1}^H \langle e_{s_{h+1}^k} - P_{s_h^k, a_h^k, h}, \widetilde{V}_{h+1}^k \rangle \right) \\ &\leq 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \widetilde{V}_{h+1}^k) \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'}, \end{aligned} \quad (120)$$

with probability at least $1 - 2SAHK\delta'$. Moreover, invoking Lemma A.1 once again reveals that

$$\begin{aligned} &\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \widetilde{V}_{h+1}^k) \\ &= \sum_{k=1}^K \sum_{h=1}^H \langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (\widetilde{V}_{h+1}^k)^2 \rangle \\ &\quad + \sum_{k=1}^H \sum_{h=1}^H \left((\widetilde{V}_{h+1}^k(s_{h+1}^k))^2 - (\widetilde{V}_h^k(s_h^k))^2 \right) + \sum_{k=1}^K \sum_{h=1}^H \left((\widetilde{V}_h^k(s_h^k))^2 - (\langle P_{s_h^k, a_h^k, h}, \widetilde{V}_{h+1}^k \rangle)^2 \right) \\ &\leq 2 \sqrt{8H^4 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \widetilde{V}_{h+1}^k) \log \frac{1}{\delta'}} + 2H^2 \sum_{k=1}^K \sum_{h=1}^H v_h(s_h^k, a_h^k) + 3H^4 \log \frac{1}{\delta'} \\ &\leq 4H^2 \sum_{k=1}^K \sum_{h=1}^H v_h(s_h^k, a_h^k) + 42H^4 \log \frac{1}{\delta'}, \end{aligned} \quad (121)$$

with probability at least $1 - 2SAHK\delta'$. Combine (120) and (121) to yield

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H v_h(s_h^k, a_h^k) &\leq \sum_{k=1}^K \widetilde{V}_1^k(s_1^k) + 2 \sqrt{8H^2 \sum_{k=1}^K \sum_{h=1}^H v_h(s_h^k, a_h^k) \log \frac{1}{\delta'}} + 84H^4 \log^2 \frac{1}{\delta'} + 3H^2 \log \frac{1}{\delta'} \\ &\leq 2 \sum_{k=1}^K \widetilde{V}_1^k(s_1^k) + 80H^2 \log \frac{1}{\delta'} \\ &\leq 2K\text{var}_1 + 80H^2 \log \frac{1}{\delta'}, \end{aligned} \quad (122)$$

with probability exceeding $1 - 4SAHK\delta'$. \square

F.1.2 Bounding T_4 . We now move on to the term T_4 , which can be written as $T_4 = \check{T}_1 + \check{T}_2$ with

$$\begin{aligned}\check{T}_1 &= \sum_{k=1}^K \sum_{h=1}^H \left(\hat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k) \right) \\ \check{T}_2 &= \sum_{k=1}^K \left(\sum_{h=1}^H r_h(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k) \right).\end{aligned}$$

This leaves us with two quantities to control.

To begin with, let us look at \check{T}_1 . In view of Lemma A.4 and the union bound over (s, a, h, k) , we see that, with probability at least $1 - 2SAHK\delta'$,

$$\hat{r}_h^k(s, a) - r_h(s, a) \leq \sqrt{\frac{2v_h(s, a) \log \frac{1}{\delta'}}{N_h^k(s, a)}} + \frac{H \log \frac{1}{\delta'}}{N_h^k(s, a)}. \quad (123)$$

As a result, we obtain

$$\begin{aligned}|\check{T}_1| &\leq \sum_{k=1}^K \sum_{h=1}^H \left(\sqrt{\frac{2v_h(s_h^k, a_h^k) \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)}} + \frac{H \log \frac{1}{\delta'}}{N_h^k(s_h^k, a_h^k)} \right) \\ &\leq \sqrt{4SAH(\log_2 K) \log \frac{1}{\delta'}} \cdot \sqrt{\sum_{k=1}^K \sum_{h=1}^H v_h(s_h^k, a_h^k) + 2SAH^2(\log_2 K) \log \frac{1}{\delta'}}.\end{aligned} \quad (124)$$

In view of (122), with probability exceeding $1 - 4SAHK\delta'$ we have

$$\sum_{k=1}^K \sum_{h=1}^H v_h(s_h^k, a_h^k) \leq 2K\text{var}_1 + 80H^2 \log \frac{1}{\delta'}. \quad (125)$$

Consequently, we arrive at

$$|\check{T}_1| \leq \sqrt{8SAHK\text{var}_1(\log_2 K) \log \frac{1}{\delta'}} + 20SAH^2(\log_2 K) \log \frac{1}{\delta'}. \quad (126)$$

Next, we proceed to bound \check{T}_2 . Toward this, we make the observation that

$$\check{T}_2 = \sum_{k=1}^K \sum_{h=1}^H \langle e_{s_{h+1}^k} - P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k} \rangle. \quad (127)$$

Applying Lemma A.1 shows that, with probability at least $1 - 2SAHK\delta'$,

$$\begin{aligned}|\check{T}_2| &\leq 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) \log \frac{1}{\delta'} + 3H \log \frac{1}{\delta'}} \\ &\leq 2 \sqrt{4 \sum_{k=1}^K \sum_{h=1}^H (\mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\star}) + \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\star} - V_{h+1}^{\pi^k})) \log \frac{1}{\delta'} + 3H \log \frac{1}{\delta'}}.\end{aligned} \quad (128)$$

Continue the calculation to derive

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^* - V_{h+1}^{\pi^k}) \\
&= \sum_{k=1}^K \sum_{h=1}^H \left(\langle P_{s_h^k, a_h^k, h}, (V_{h+1}^* - V_{h+1}^{\pi^k})^2 \rangle - (\langle P_{s_h^k, a_h^k, h}, V_{h+1}^* - V_{h+1}^{\pi^k} \rangle)^2 \right) \\
&\leq \sum_{k=1}^K \sum_{h=1}^H \left\langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (V_{h+1}^* - V_{h+1}^{\pi^k})^2 \right\rangle \\
&\quad + 2H \sum_{k=1}^K \sum_{h=1}^H \max \left\{ \left(V_h^*(s_h^k) - r_h(s_h^k, a_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^* \rangle \right) \right. \\
&\quad \left. - \left(V_h^{\pi^k}(s_h^k) - r_h(s_h^k, a_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k} \rangle \right), 0 \right\} \\
&\leq 2 \sqrt{8H^2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^* - V_{h+1}^{\pi^k}) \log \frac{1}{\delta'}} \\
&\quad + 2H \sum_{k=1}^K \sum_{h=1}^H \left(V_h^*(s_h^k) - r_h(s_h^k, a_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^* \rangle \right) + 3H^2 \log \frac{1}{\delta'}. \quad (129)
\end{aligned}$$

Here, (129) holds with probability at least $1 - 2SAHK\delta'$, a consequence of Lemma A.1 and Lemma A.2.

To further bound the right-hand side of (129), we develop the following upper bound:

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h=1}^H \left(V_h^*(s_h^k) - r_h(s_h^k, a_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^* \rangle \right) \\
&= \sum_{k=1}^K \left(V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k) \right) + \sum_{k=1}^K \left(V_1^{\pi^k}(s_1^k) - \sum_{h=1}^H r_h(s_h^k, a_h^k) \right) + \sum_{k=1}^K \sum_{h=1}^H \langle e_{s_{h+1}^k} - P_{s_h^k, a_h^k, h}, V_{h+1}^* \rangle. \quad (130)
\end{aligned}$$

Note that the first term on the right-hand side (130) is exactly $\text{Regret}(K) = T_1 + T_2 + T_3 + T_4$, the second term on the right-hand side (130) corresponds to $-T_4$, whereas the third term on the right-hand side (130) can be bounded by

$$\sum_{k=1}^K \sum_{h=1}^H \langle e_{s_{h+1}^k} - P_{s_h^k, a_h^k, h}, V_{h+1}^* \rangle \leq 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^*) \log \frac{1}{\delta'}} + 3H \log \frac{1}{\delta'}, \quad (131)$$

with probability at least $1 - 2SAHK\delta'$. It then implies the validity of the following bound with probability exceeding $1 - 8SAHK\delta'$:

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h=1}^H \left(V_h^*(s_h^k) - r_h(s_h^k, a_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^* \rangle \right) \\
&\leq T_1 + T_2 + T_3 + 2|T_4| + 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^*) \log \frac{1}{\delta'}} + 55H \log \frac{1}{\delta'}. \quad (132)
\end{aligned}$$

Combining these bounds with (129), we can use a little algebra to further obtain

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\star} - V_{h+1}^{\pi^k}) \\ & \leq 4H \left(T_1 + T_2 + T_3 + 2|T_4| + 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\star}) \log \frac{1}{\delta'}} \right) + 262H^2 \log \frac{1}{\delta'}, \end{aligned} \quad (133)$$

with probability at least $1 - 8SAHK\delta'$. If we define $T_{10} = \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\star})$, then substituting (133) into (128) yields: with probability exceeding $1 - 10SAHK\delta'$,

$$\begin{aligned} |\check{T}_2| & \leq 2\sqrt{8K\text{var}_1 \log \frac{1}{\delta'}} + 8\sqrt{H \left(T_1 + T_2 + T_3 + 2|T_4| + 2\sqrt{2T_{10} \log \frac{1}{\delta'}} \right) \log \frac{1}{\delta'}} + 107H \log \frac{1}{\delta'} \\ & \leq 11\sqrt{T_{10} \log \frac{1}{\delta'}} + 16\sqrt{H(T_1 + T_2 + T_3 + 2|T_4|) \log \frac{1}{\delta'}} + 115H \log \frac{1}{\delta'}. \end{aligned} \quad (134)$$

Combining the above bound on $|\check{T}_2|$ with (126), with probability exceeding $1 - 10SAHK\delta'$

$$\begin{aligned} |T_4| & \leq |\check{T}_1| + |\check{T}_2| \\ & \leq 18\sqrt{SAHT_{10}(\log_2 K) \log \frac{1}{\delta'}} + 16\sqrt{H(T_1 + T_2 + T_3 + 2|T_4|) \log \frac{1}{\delta'}} + 135SAH^2(\log_2 K) \log \frac{1}{\delta'}, \end{aligned}$$

which together with a little algebra yields

$$|T_4| \leq 36\sqrt{SAHT_{10}(\log_2 K) \log \frac{1}{\delta'}} + 32\sqrt{H(T_1 + T_2 + T_3) \log \frac{1}{\delta'}} + 306SAH^2(\log_2 K) \log \frac{1}{\delta'}. \quad (135)$$

F.1.3 Bounding T_5 and T_6 . We now turn attention to the terms T_5 and T_6 . Toward this, we start with the following lemma.

LEMMA F.4. *With probability at least $1 - 2SAHK\delta'$, one has*

$$T_5 \leq 5T_6 + 8BSAH^3. \quad (136)$$

PROOF OF LEMMA F.4. Direct computation gives

$$\begin{aligned} & \sum_{k,h} \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k) \\ & = \sum_{k,h} \left(\langle \widehat{P}_{s_h^k, a_h^k, h}^k, (V_{h+1}^k)^2 \rangle - (\langle \widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle)^2 \right) \\ & \leq \sum_{k,h} \left(\langle P_{s_h^k, a_h^k, h}^k, (V_{h+1}^k)^2 \rangle - (\langle P_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle)^2 \right) + \sum_{k,h} \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}^k, (V_{h+1}^k)^2 \rangle \end{aligned} \quad (137)$$

$$\begin{aligned} & + 2H \sum_{k,h} \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle \\ & \leq \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k) + \sum_{k,h} \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}^k, (V_{h+1}^k)^2 \rangle + 2H \sum_{k,h} \langle \widehat{P}_{s_h^k, a_h^k, h}^k - P_{s_h^k, a_h^k, h}^k, V_{h+1}^k \rangle \\ & = T_5 + T_7 + 2HT_1. \end{aligned} \quad (138)$$

Invoking Lemma 4.5 to bound T_7 and T_1 , we obtain

$$\begin{aligned} \sum_{k,h} \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k) &\leq \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k) + 6 \sqrt{\sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k) BSAH^3} + 3BSAH^3 \\ &\leq 5 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k) + 8BSAH^3, \end{aligned} \quad (139)$$

with probability exceeding $1 - 2SAHK\delta'$. \square

In view of Lemma F.4, it suffices to bound $T_6 = \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k)$. Given that $\text{Var}(X + Y) \leq 2(\text{Var}(X) + \text{Var}(Y))$ holds for any two random variables X, Y , we have

$$\begin{aligned} T_6 &= \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k) \leq 2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^{\star}) + 2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k - V_{h+1}^{\star}) \\ &\leq 3K\text{var}_1 + \sum_{k=1}^K \left(\sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^{\star}) - 3\text{var}_1 \right) + 2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k - V_{h+1}^{\star}). \end{aligned} \quad (140)$$

To further upper bound the right-hand side of (140), we make note of the following lemmas.

LEMMA F.5. *With probability at least $1 - 4SAHK\delta'$, it holds that*

$$T_{10} - 2K\text{var}_1 = \sum_{k=1}^K \left(\sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^{\star}) - 2\text{var}_1 \right) \leq 80H^2 \log \frac{1}{\delta'}. \quad (141)$$

LEMMA F.6. *With probability at least $1 - 2\delta'$, it holds that*

$$\sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k - V_{h+1}^{\star}) \leq 4 \sqrt{BH^2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k) + 4H \sum_{k,h} b_h^k(s_h^k, a_h^k) + 3BSAH^3}.$$

Combining Lemma F.5 and Lemma F.6 with (140), we see that with probability at least $1 - 6SAHK\delta'$,

$$\begin{aligned} T_6 &= \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}^k, V_{h+1}^k) \\ &\leq 4K\text{var}_1 + 8\sqrt{BSAH^3 T_6} + 8HT_2 + 7BSAH^3, \end{aligned}$$

and as a result,

$$T_6 \leq 8K\text{var}_1 + 16HT_2 + 78BSAH^3. \quad (142)$$

This taken collectively with Lemma F.4 yields, with probability at least $1 - 8SAHK\delta'$,

$$T_5 = \sum_{k,h} \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}^k, V_{h+1}^k) \leq 40K\text{var}_1 + 80HT_2 + 398BSAH^3. \quad (143)$$

To finish our bounds on T_5 and T_6 , it remains to establish Lemma F.5 and Lemma F.6.

PROOF OF LEMMA F.5. Let $\bar{R}_h^{\star}(s, a) = \mathbb{V}(P_{s, a, h}, V_{h+1}^{\star})$, and define

$$\bar{V}_h^k(s) = \mathbb{E} \left[\sum_{h'=h}^H \bar{R}_{h'}(s_{h'}, a_{h'}) \mid s_h = s \right].$$

Then $\bar{V}_h^k(s) \leq \text{var}_1 \leq H^2$. It then follows that

$$\begin{aligned} \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\star}) - \text{var}_1 &= \sum_{h=1}^H \bar{R}_h^{\star}(s_h^k, a_h^k) - \text{var}_1 \\ &\leq \sum_{h=1}^H \bar{R}_h^{\star}(s_h^k, a_h^k) - \bar{V}_1^k(s_1^k) \\ &= \sum_{h=1}^H (e_{s_{h+1}^k} - P_{s_h^k, a_h^k, h}) \bar{V}_{h+1}^k. \end{aligned} \quad (144)$$

Note that \bar{V}^k depends only on π^k , which is determined before the beginning of the k th episode. Consequently, applying Lemma A.1 reveals that, with probability at least $1 - 2SAHK\delta'$,

$$\begin{aligned} &\sum_{k=1}^K \left(\sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\star}) - \bar{V}_1^k(s_1^k) \right) \\ &\leq 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \bar{V}_{h+1}^k) \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'}. \end{aligned} \quad (145)$$

Regarding the sum of variance terms on the right-hand side of (145), one can further bound

$$\begin{aligned} &\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \bar{V}_{h+1}^k) \\ &= \sum_{k=1}^K \sum_{h=1}^H \left(\langle P_{s_h^k, a_h^k, h}, (\bar{V}_{h+1}^k)^2 \rangle - (\langle P_{s_h^k, a_h^k, h}, \bar{V}_{h+1}^k \rangle)^2 \right) \\ &= \sum_{k=1}^K \sum_{h=1}^H \langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (\bar{V}_{h+1}^k)^2 \rangle \\ &\quad + \sum_{k=1}^K \sum_{h=1}^H \left((\bar{V}_{h+1}^k(s_{h+1}^k))^2 - (\bar{V}_h^k(s_h^k))^2 \right) + \sum_{k=1}^K \sum_{h=1}^H \left((\bar{V}_h^k(s_h^k))^2 - (\langle P_{s_h^k, a_h^k, h}, \bar{V}_{h+1}^k \rangle)^2 \right) \\ &\leq 2 \sqrt{8H^4 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \bar{V}_{h+1}^k) \log \frac{1}{\delta'}} + 2H^2 \sum_{k=1}^K \sum_{h=1}^H \bar{R}_h(s_h^k, a_h^k) + 3H^4 \log \frac{1}{\delta'}, \end{aligned} \quad (146)$$

with probability at least $1 - 2SAHK\delta'$. Here, the last inequality arises from Lemmas A.1 and A.2 as well as the fact that $\bar{V}_h^k(s_h^k) = \bar{R}_h(s_h^k, a_h^k) + P_{s_h^k, a_h^k, h} \bar{V}_{h+1}^k$. It then follows from elementary algebra that

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \bar{V}_{h+1}^k) \leq 4H^2 \sum_{k=1}^K \sum_{h=1}^H \bar{R}_h(s_h^k, a_h^k) + 42H^4 \log \frac{1}{\delta'}. \quad (147)$$

Substituting (147) into (145) gives

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\star}) \leq \sum_{k=1}^H \bar{V}_1^k(s_1^k) + 2 \sqrt{8H^2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\star}) \log \frac{1}{\delta'}} + 21H^2 \log \frac{1}{\delta'},$$

thus indicating that

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^*) \leq 2 \sum_{k=1}^K \bar{V}_1^k(s_1^k) + 84H^2 \log \frac{1}{\delta'} \leq 2K\text{var}_1 + 84H^2 \log \frac{1}{\delta'}.$$

The proof of Lemma F.5 is thus completed. \square

PROOF OF LEMMA F.6. We make the observation that

$$\begin{aligned} & \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^*) \\ &= \sum_{k,h} \left(\langle P_{s_h^k, a_h^k, h}, (V_{h+1}^k - V_{h+1}^*)^2 \rangle - (\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^* \rangle)^2 \right) \\ &= \sum_{k,h} \left(\langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (V_{h+1}^k - V_{h+1}^*)^2 \rangle \right) \\ & \quad + \sum_{k,h} \left((V_{h+1}^k(s_{h+1}^k) - V_{h+1}^*(s_{h+1}^k))^2 - (\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^* \rangle)^2 \right) \\ &= \sum_{k,h} \left(\langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (V_{h+1}^k - V_{h+1}^*)^2 \rangle \right) + \sum_{k,h} \left((V_h^k(s_h^k) - V_h^*(s_h^k))^2 - (\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^* \rangle)^2 \right). \end{aligned} \tag{148}$$

According to Lemmas A.1 and A.2, we see that with probability exceeding $1 - \delta'$,

$$\begin{aligned} & \sum_{k,h} \langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (V_{h+1}^k - V_{h+1}^*)^2 \rangle \\ & \leq 2\sqrt{2} \sqrt{4H^2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^*) \log \frac{1}{\delta'} + 3H^2 \log \frac{1}{\delta'}}. \end{aligned} \tag{149}$$

In addition, with probability at least $1 - \delta'$ one has

$$\begin{aligned} & \sum_{k,h} \left((V_h^k(s_h^k) - V_h^*(s_h^k))^2 - (\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^* \rangle)^2 \right) \\ & \leq 2H \sum_{k,h} \max \{V_h^k(s_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle - (V_h^*(s_h^k) - \langle P_h^k, V_{h+1}^* \rangle), 0\} \\ & \leq 2H \sum_{k,h} \max \{V_h^k(s_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle - r_h(s_h^k, a_h^k), 0\} \\ & \leq 2H \sum_{k,h} \max \{\langle \hat{P}_{s_h^k, a_h^k, h} - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle, 0\} + 2H \sum_{k,h} b_h^k \\ & \leq 2 \sqrt{BSAH^3 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k)} + 2H \sum_{k,h} b_h^k(s_h^k, a_h^k) + BSAH^3. \end{aligned} \tag{150}$$

It then follows that, with probability at least $1 - 2\delta'$,

$$\sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^*) \leq 4 \sqrt{BSAH^3 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k)} + 4H \sum_{k,h} b_h^k(s_h^k, a_h^k) + 3BSAH^3, \tag{151}$$

thereby concluding the proof. \square

F.1.4 Putting All This Together. To finish up, let us rewrite the inequalities (40g)–(40f) as follows, with (40a), (40c), (40d), and (40e) replaced by (117), (135) (143), and (142), respectively:

$$\begin{aligned}
T_1 &\leq \sqrt{128BSAH} + 24BSAH^2, \\
T_7 &\leq H\sqrt{512BSAH} + 24BSAH^3, \\
T_9 &\leq \sqrt{128BSAH} + 24BSAH^2, \\
T_2 &\leq 100\sqrt{BSAH} + 140BSAH^2, \\
T_3 &\leq \sqrt{8BT} + 3H \log \frac{1}{\delta'}, \\
T_4 &\leq \sqrt{BSAH} + 32\sqrt{BH(T_1 + T_2 + T_3)} + BSAH^2, \\
T_5 &\leq 40K\text{var}_1 + 80HT_2 + 398BSAH^3, \\
T_6 &\leq 8K\text{var}_1 + 16HT_2 + 78BSAH^3, \\
T_8 &\leq \sqrt{32BH^2T_6} + 3BH^2,
\end{aligned}$$

where we recall that $B = 4,000(\log_2 K)^3 \log(3SA) \log \frac{1}{\delta'}$. In addition, it follows from Lemma F.5 that

$$T_{10} \leq 2K\text{var}_1 + 80BH^2.$$

Solving the inequalities above reveals that, with probability exceeding $1 - 200SAH^2K^2\delta'$,

$$\text{Regret}(K) = T_1 + T_2 + T_3 + T_4 \leq O\left(\sqrt{BSAHK\text{var}_1} + BSAH^2\right). \quad (152)$$

One can thus conclude the proof by recalling that $\delta' = \frac{\delta}{200SAH^2K^2}$.

F.2 Proof of Lemma F.2

Following similar arguments as in the proof of Lemma F.1, we focus on bounding T_2 , T_4 , T_5 , and T_6 in terms of var_2 .

F.2.1 Bounding T_2 . Recall that δ' is defined as $\delta' = \frac{\delta}{200SAH^2K^2}$, and that we have demonstrated in (67) that

$$\begin{aligned}
T_2 &\leq \frac{460}{9}\sqrt{2SAHT_5(\log_2 K) \log \frac{1}{\delta'}} \\
&\quad + 4\sqrt{SAH(\log_2 K) \log \frac{1}{\delta'}} \sqrt{\sum_{k,h} \left(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \right)} + \frac{1088}{9}SAH^2(\log_2 K) \log \frac{1}{\delta'}.
\end{aligned} \quad (153)$$

To bound the right-hand side of (153), we resort to the following lemma.

LEMMA F.7. *With probability at least $1 - 4SAHK\delta'$, one has*

$$\sum_{k,h} \left(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \right) \leq 6K\text{var}_2 + 242H^2(\log_2 K) \log \frac{1}{\delta'}. \quad (154)$$

PROOF. Recall that in Lemma F.3, we have shown that with probability at least $1 - 4SAHK\delta'$,

$$\sum_{k=1}^K \sum_{h=1}^H \left(\widehat{\sigma}_h^k(s_h^k, a_h^k) - (\widehat{r}_h^k(s_h^k, a_h^k))^2 \right) \leq 3 \sum_{k=1}^K \widetilde{V}_1^k(s_1^k) + 2SAH^3(\log_2 K) \log \frac{1}{\delta'}. \quad (155)$$

We then complete the proof by observing that

$$\begin{aligned}\widetilde{V}_1^k(s_1^k) &\leq \widetilde{V}_1^k(s_1^k) + \mathbb{E}_{\pi^k} \left[\sum_{h=1}^H \mathbb{V}(P_{s_h, a_h, h}, V_{h+1}^{\pi^k}) \mid s_1 = s_1^k \right] \\ &= \text{Var}_{\pi^k} \left[\sum_{h=1}^H r_h(s_h, a_h) \mid s_1 = s_1^k \right] \leq \text{var}_2.\end{aligned}\quad (156)$$

□

Combining Lemma F.7 with (153) gives: with probability at least $1 - 4SAHK\delta'$,

$$\begin{aligned}T_2 &\leq \frac{460}{9} \sqrt{2SAHT_5(\log_2 K) \log \frac{1}{\delta'}} + 12 \sqrt{SAH(\log_2 K) \log \frac{1}{\delta'}} \sqrt{2K\text{var}_2} \\ &\quad + 157SAH^2(\log_2 K) \log \frac{1}{\delta'}.\end{aligned}\quad (157)$$

F.2.2 Bounding T_4 . Recall that $T_4 = \check{T}_1 + \check{T}_2$, where

$$\begin{aligned}\check{T}_1 &= \sum_{k=1}^K \sum_{h=1}^H \left(\widehat{r}_h^k(s_h^k, a_h^k) - r_h(s_h^k, a_h^k) \right), \\ \check{T}_2 &= \sum_{k=1}^K \left(\sum_{h=1}^H r_h(s_h^k, a_h^k) - V_1^{\pi^k}(s_1^k) \right).\end{aligned}$$

Repeating similar arguments employed in the proof of Lemma F.3 and using (124), we see that with probability exceeding $1 - 6SAHK\delta'$,

$$\begin{aligned}|\check{T}_1| &\leq \sqrt{4SAH(\log_2 K) \log \frac{1}{\delta'}} \cdot \sqrt{\sum_{k=1}^K \sum_{h=1}^H v_h(s_h^k, a_h^k) + 2SAH^2(\log_2 K) \log \frac{1}{\delta'}} \\ &\leq \sqrt{8SAHK\text{var}_2(\log_2 K) \log \frac{1}{\delta'}} + 20SAH^2(\log_2 K) \log \frac{1}{\delta'}.\end{aligned}$$

In addition, from Lemma A.1 and the definition of var_2 , we see that

$$|\check{T}_2| \leq 2\sqrt{2K\text{var}_2 \log \frac{1}{\delta'}} + 3H \log \frac{1}{\delta'}, \quad (158)$$

with probability at least $1 - 2SAHK\delta'$. Therefore, with probability at least $1 - 8SAHK\delta'$, it holds that

$$T_4 \leq 4\sqrt{2SAHK\text{var}_2(\log_2 K) \log \frac{1}{\delta'}} + 23SAH^2(\log_2 K) \log \frac{1}{\delta'}. \quad (159)$$

F.2.3 Bounding T_5 and T_6 . Recall that Lemma F.4 asserts that with probability exceeding $1 - 2\delta'$,

$$T_5 \leq 5T_6 + 8BSAH^3.$$

Hence, it suffices to bound T_6 .

From the elementary inequality $\text{Var}(X + Y) \leq 2\text{Var}(X) + 2\text{Var}(Y)$, we obtain

$$\begin{aligned} T_6 &= \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k) \leq 2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) + 2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k}) \\ &\leq 3K\text{var}_2 + \sum_{k=1}^K \left(\sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) - 3\text{var}_2 \right) + 2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k}). \end{aligned} \quad (160)$$

To bound the right-hand side of (160), we resort to the following two lemmas.

LEMMA F.8. *With probability at least $1 - 4SAHK\delta'$, it holds that*

$$\sum_{k=1}^K \left(\sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) - 2\text{var}_2 \right) \leq 80H^2 \log \frac{1}{\delta'}. \quad (161)$$

LEMMA F.9. *With probability exceeding $1 - 4SAKH\delta'$, it holds that*

$$\sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k}) \leq 4 \sqrt{BH^2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k)} + 4H \sum_{k,h} b_h^k(s_h^k, a_h^k) + 3BSAH^3.$$

With Lemmas F.8 and F.9 in place, we can demonstrate that with probability at least $1 - 6SAHK\delta'$,

$$\begin{aligned} T_6 &\leq 2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) + 2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k}) \\ &\leq 4K\text{var}_2 + 8\sqrt{BSAH^3 T_6} + 8HT_2 + 7BSAH^3, \\ \implies T_6 &\leq 8K\text{var}_2 + 16HT_2 + 78BSAH^3. \end{aligned} \quad (162)$$

Taking this result together with Lemma F.4 gives, with probability exceeding $1 - 8SAHK\delta'$,

$$T_5 = \sum_{k,h} \mathbb{V}(\widehat{P}_{s_h^k, a_h^k, h}, V_{h+1}^k) \leq 40K\text{var}_2 + 80HT_2 + 398BSAH^3. \quad (163)$$

To finish establishing the above bounds on T_5 and T_6 , it suffices to prove Lemmas F.8 and F.9, which we accomplish in the sequel.

PROOF OF LEMMA F.8. For notational convenience, define

$$\check{R}_h^k(s, a) = \mathbb{V}(P_{s, a, h}, V_{h+1}^{\pi^k}) \quad \text{and} \quad \check{V}_h^k(s) = \mathbb{E} \left[\sum_{h'=h}^H \check{R}_{h'}^k(s_{h'}, a_{h'}) \mid s_h = s \right].$$

It is easily seen that $\check{V}_h^k(s) \leq \text{var}_2 \leq H^2$.

We also make the observation that

$$\begin{aligned} \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) - \text{var}_2 &= \sum_{h=1}^H \check{R}_h^k(s_h^k, a_h^k) - \text{var}_2 \\ &\leq \sum_{h=1}^H \check{R}_h^k(s_h^k, a_h^k) - \check{V}_1^k(s_1^k) \\ &= \sum_{h=1}^H \langle e_{s_{h+1}^k} - P_{s_h^k, a_h^k, h}, \check{V}_{h+1}^k \rangle. \end{aligned} \quad (164)$$

Note that \check{V}^k only depends on π^k , which is determined before the k th episode starts. Lemma A.1 then tells us that, with probability at least $1 - 2SAHK\delta'$,

$$\begin{aligned} & \sum_{k=1}^K \left(\sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) - \check{V}_1^k(s_1^k) \right) \\ & \leq 2 \sqrt{2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \check{V}_{h+1}^k) \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'}. \end{aligned} \quad (165)$$

Further, it is observed that with probability at least $1 - 2SAHK\delta'$,

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \check{V}_{h+1}^k) \\ & = \sum_{k=1}^K \sum_{h=1}^H \left(\langle P_{s_h^k, a_h^k, h}, (\check{V}_{h+1}^k)^2 \rangle - (\langle P_{s_h^k, a_h^k, h}, \check{V}_{h+1}^k \rangle)^2 \right) \\ & = \sum_{k=1}^K \sum_{h=1}^H \langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (\check{V}_{h+1}^k)^2 \rangle \\ & \quad + \sum_{k=1}^K \sum_{h=1}^H \left((\check{V}_{h+1}^k(s_{h+1}^k))^2 - (\check{V}_h^k(s_h^k))^2 \right) + \sum_{k=1}^K \sum_{h=1}^H \left((\check{V}_h^k(s_h^k))^2 - (\langle P_{s_h^k, a_h^k, h}, \check{V}_{h+1}^k \rangle)^2 \right) \\ & \leq 2 \sqrt{8H^4 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \check{V}_{h+1}^k) \log \frac{1}{\delta'}} + 2H^2 \sum_{k=1}^K \sum_{h=1}^H \check{R}_h(s_h^k, a_h^k) + 3H^4 \log \frac{1}{\delta'}. \end{aligned} \quad (166)$$

Here, the last inequality results from Lemmas A.1, A.2 and the fact that $\check{V}_h^k(s_h^k) = \check{R}_h(s_h^k, a_h^k) + \langle P_{s_h^k, a_h^k, h}, \check{V}_{h+1}^k \rangle$. It then follows that

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, \check{V}_{h+1}^k) \leq 4H^2 \sum_{k=1}^K \sum_{h=1}^H \check{R}_h(s_h^k, a_h^k) + 42H^4 \log \frac{1}{\delta'}. \quad (167)$$

Taking (165) and (167) together leads to

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) \leq \sum_{k=1}^H \check{V}_1^k(s_1^k) + 2 \sqrt{8H^2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) \log \frac{1}{\delta'}} + 21H^2 \log \frac{1}{\delta'},$$

which further implies that

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^{\pi^k}) \leq 2 \sum_{k=1}^K \check{V}_1^k(s_1^k) + 84H^2 \log \frac{1}{\delta'} \leq 2K\text{var}_2 + 84H^2 \log \frac{1}{\delta'}.$$

This concludes the proof. \square

PROOF OF LEMMA F.6. A little algebra gives

$$\begin{aligned}
& \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h} V_{h+1}^k - V_{h+1}^{\pi^k}) \\
&= \sum_{k,h} \left(\langle P_{s_h^k, a_h^k, h}, (V_{h+1}^k - V_{h+1}^{\pi^k})^2 \rangle - (\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k} \rangle)^2 \right) \\
&= \sum_{k,h} \left(\langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (V_{h+1}^k - V_{h+1}^{\pi^k})^2 \rangle \right) \\
&\quad + \sum_{k,h} \left((V_{h+1}^k(s_{h+1}^k) - V_{h+1}^{\pi^k}(s_{h+1}^k))^2 - (\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k} \rangle)^2 \right) \\
&= \sum_{k,h} \left(\langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (V_{h+1}^k - V_{h+1}^{\pi^k})^2 \rangle \right) + \sum_{k,h} \left((V_h^k(s_h^k) - V_h^{\pi^k}(s_h^k))^2 - (\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k} \rangle)^2 \right). \tag{168}
\end{aligned}$$

From Lemmas A.1 and A.2, we can show that with probability $1 - 2SAKH\delta'$,

$$\sum_{k,h} \langle P_{s_h^k, a_h^k, h} - e_{s_{h+1}^k}, (V_{h+1}^k - V_{h+1}^{\pi^k})^2 \rangle \tag{169}$$

$$\leq 2\sqrt{2} \sqrt{4H^2 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k}) \log \frac{1}{\delta'}} + 3H^2 \log \frac{1}{\delta'}. \tag{170}$$

Additionally, with probability at least $1 - 2SAKH\delta'$,

$$\begin{aligned}
& \sum_{k,h} \left\{ (V_h^k(s_h^k) - V_h^{\pi^k}(s_h^k))^2 - (\langle P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k} \rangle)^2 \right\} \\
&\leq 2H \sum_{k,h} \max \left\{ V_h^k(s_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle - (V_h^{\pi^k}(s_h^k) - \langle P_h^k, V_{h+1}^{\pi^k} \rangle), 0 \right\} \\
&= 2H \sum_{k,h} \max \left\{ V_h^k(s_h^k) - \langle P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle - r_h(s_h^k, a_h^k), 0 \right\} \\
&\leq 2H \sum_{k,h} \max \left\{ \langle \widehat{P}_{s_h^k, a_h^k, h} - P_{s_h^k, a_h^k, h}, V_{h+1}^k \rangle, 0 \right\} + 2H \sum_{k,h} b_h^k(s_h^k, a_h^k) \\
&\leq 2 \sqrt{BSAH^3 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k)} + 2H \sum_{k,h} b_h^k(s_h^k, a_h^k) + BSAH^3. \tag{171}
\end{aligned}$$

It then follows that

$$\sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k - V_{h+1}^{\pi^k}) \leq 4 \sqrt{BSAH^3 \sum_{k,h} \mathbb{V}(P_{s_h^k, a_h^k, h}, V_{h+1}^k)} + 4H \sum_{k,h} b_h^k(s_h^k, a_h^k) + 3BSAH^3, \tag{172}$$

with probability at least $1 - 4SAKH\delta'$. The proof is thus complete. \square

F.2.4 Putting All Pieces Together. Recall that $B = 4,000(\log_2 K)^3 \log(3SA) \log \frac{1}{\delta'}$. The last step is to rewrite the inequalities (40g) – (40f) as follows with (40a), (40c), (40d), and (40e) replaced by

(157), (159), (163), and (162) respectively:

$$\begin{aligned}
T_1 &\leq \sqrt{128BSAH T_6} + 24BSAH^2, \\
T_7 &\leq H\sqrt{512BSAH T_6} + 24BSAH^3, \\
T_9 &\leq \sqrt{128BSAH T_6} + 24BSAH^2, \\
T_2 &\leq 100\sqrt{BSAH T_5} + 140BSAH^2, \\
T_3 &\leq \sqrt{8BT_6} + 3H \log \frac{1}{\delta'}, \\
T_4 &\leq \sqrt{BSAH K \text{var}_2} + BSAH^2; \\
T_5 &\leq 40K \text{var}_2 + 80HT_2 + 398BSAH^3, \\
T_6 &\leq 8K \text{var}_2 + 16HT_2 + 78BSAH^3, \\
T_8 &\leq \sqrt{32BH^2 T_6} + 3BH^2,
\end{aligned}$$

which are valid with probability at least $1 - 200SAH^2K^2\delta'$. Solving the inequalities listed above, we can readily conclude that

$$\text{Regret}(K) = T_1 + T_2 + T_3 + T_4 \leq O\left(\sqrt{BSAH K \text{var}_2} + BSAH^2\right). \quad (173)$$

This finishes the proof by recalling that $\delta' = \frac{\delta}{200SAH^2K^2}$.

G Minimax Lower Bounds

In this section, we establish the lower bounds advertised in this article.

G.1 Proof of Theorem 6.1

Consider any given (S, A, H) . We start by establishing the following lemma.

LEMMA G.1. *Consider any $K' \geq 1$. For any algorithm, there exists an MDP instance with S states, A actions, and horizon H , such that the regret in K' episodes is at least*

$$\text{Regret}(K') = \Omega(f(K')) = \Omega\left(\min\left\{\sqrt{SAH^3K'}, K'H\right\}\right).$$

PROOF OF LEMMA G.1. Our construction of the hard instance is based on the hard instance JAO-MDP constructed in Jaksch et al. [26], Jin et al. [29]. In Jin et al. [29, Appendix.D], the authors already showed that when $K' \geq C_0SAH$ for some constant $C_0 > 0$, the minimax regret lower bound is $\Omega(\sqrt{SAH^3K'})$. Hence, it suffices for us to focus on the regime where $K' \leq C_0SAH$. Without loss of generality, we assume $S = A = 2$, and the argument to generalize it to arbitrary (S, A) is standard and hence omitted for brevity.

Recall the construction of JAO-MDP in Jaksch et al. [26]. Let the two states be x and y , and the two actions be a and b . The reward is always equal to x in state 1 and $1/2$ in state y . The probability transition kernel is given by

$$P_{x,a} = P_{x,b} = [1 - \delta, \delta], \quad P_{y,a} = [1 - \delta, \delta], \quad P_{y,b} = [1 - \delta - \epsilon, \delta + \epsilon],$$

where we choose $\delta = C_1/H$ and $\epsilon = 1/H$. Then the mixing time of the MDP is roughly $O(H)$. By choosing C_1 large enough, we can ensure that the MDP is C_3 -mixing after the first half of the horizons for some proper constant $C_3 \in (0, 1/2)$.

It is then easy to show that action b is the optimal action for state y . Moreover, whenever action a is chosen in state y , the learner needs to pay regret $\Omega(\epsilon H) = \Omega(1)$. In addition, to differentiate

action a from action b in state y with probability at least $1 - \frac{1}{10}$, the learner needs at least $\Omega(\frac{\epsilon}{\delta^2}) = \Omega(H)$ rounds — let us call it C_4H rounds for some proper constant $C_4 > 0$. As a result, in the case where $K' \leq C_4H$, the minimax regret is at least $\Omega(K'H^2\epsilon) = \Omega(K'H)$. When $C_4H \leq K' \leq C_0SAH = 4C_0H$, the minimax regret is at least $\Omega(C_4H^2) = \Omega(K'H)$. This concludes the proof. \square

With Lemma G.1, we are ready to prove Theorem 6.1. Let \mathcal{M} be the hard instance for $K' = \max\{\frac{1}{10}Kp, 1\}$ constructed in the proof of Lemma G.1. We construct an MDP \mathcal{M}' as below.

- In the first step, for any state s , with probability p , the leaner transitions to a copy of \mathcal{M} , and with probability $1 - p$, the learner transitions to a dumb state with 0 reward.

It can be easily verified that $v^* \leq pH$. Let $X = X_1 + X_2 + \dots + X_k$, where $\{X_i\}_{i=1}^K$ are i.i.d. Bernoulli random variables with mean p . Let $g(X, K')$ denote the minimax regret on the hard instance \mathcal{M} in X episodes. Given that $g(X, K')$ is non-decreasing in X , one sees that

$$\text{Regret}(K) \geq \mathbb{E}[g(X, K')].$$

- In the case where $Kp \geq 10$, Lemma A.3 tells us that with probability at least $1/2$, $X \geq \frac{1}{10}Kp = K'$, and then it holds that

$$\mathbb{E}[g(X, K')] \geq \frac{1}{2}g(K', K') = \frac{1}{2}f(K') = \frac{1}{2}\Omega\left(\min\left\{\sqrt{SAH^3K'}, K'H\right\}\right) = \Omega(\sqrt{SAH^3Kp}, KHp).$$

- In the case where $Kp < 10$, with probability exceeding $1 - (1 - p)^K \geq (1 - e^{-Kp}) \geq \frac{Kp}{30}$, one has $X \geq 1$. Then one has

$$\mathbb{E}[g(X, K')] \geq \frac{Kp}{30} \cdot g(1, K') = \frac{Kp}{30} \cdot g(1, 1) = \Omega(KHp).$$

The preceding bounds taken together complete the proof.

G.2 Proof of Theorem 6.2

Without loss of generality, assume that $S = A = 2$ (as in the proof of Theorem 6.1), and recall the assumption that $p \leq 1/4$. In what follows, we construct a hard instance for which the learner needs to identify the correct action for each step.

Let $\mathcal{S} = \{s_1, s_2\}$, and take the initial state to be s_1 . The transition kernel and cost are chosen as follows.

- Select $a_h^* \in \{a_1, a_2\}$ for every $h \in [H]$.
- For each action a and each step h , set $P_{s_2, a, h} = e_{s_2}$ and $c_h(s_2, a) = 0$.
- For each step h and each action $a \neq a_h^*$, set $P_{s_1, a, h} = e_{s_2}$ and $c_h(s_1, a) = 1$.
- Set $P_{s_1, a_h^*, h} = e_{s_1}$ and $c_h(s_1, a_h^*) = p$.

It can be easily checked that $c^* = Hp$ (the cost obtained by choosing action a_h^* for each step h).

Note that in the above construction, the a_h^* 's are selected independently across different steps. Thus, to identify the optimal action a_h^* for at least half of the H steps, we need at least $\Omega(H)$ episodes. This implies that: there exists a constant $C_5 > 0$ such that in the first $K \leq C_5H$ episodes, the cost of the learner is at least $\Omega(H(1 - p))$. As a result, the minimax regret is at least

$$\Omega(K(H - c^*)) = \Omega(KH(1 - p)),$$

when $K \leq C_5H$. Similarly, in the case where $C_5H \leq K \leq \frac{100H}{p}$, the minimax regret is at least

$$\Omega(H(H - c^*)) = \Omega(H^2(1 - p)).$$

We then turn to the case where $K \geq \frac{100H}{p}$. Let \mathcal{M} be the hard instance having the same transition as the instance constructed in the proof of Lemma G.1, and set the cost as $1/2$ (respectively 1) for state x (respectively state y), with respect to $K' = Kp/10 \geq 10H$ (a quantity defined therein). Let

\mathcal{M}' be the MDP such that: in the first step, with probability p , the learner transitions to a copy of \mathcal{M} , and with probability $1 - p$, the learner transitions to a dumb state with 0 cost. Then we clearly have $c^* = \Theta(Hp)$. It follows from Lemma A.3 that, with probability exceeding $1/2$, one has $X \geq \frac{1}{3}Kp - \log 2 \geq \frac{1}{6}Kp$, where X is again defined in the proof of Lemma G.1. Then one has

$$\text{Regret}(K) \geq \frac{1}{2} \Omega \left(\min \{ \sqrt{H^3 K'}, K' H \} \right) = \Omega(\sqrt{H^3 K}).$$

The proof is thus completed by combining the above minimax regret lower bounds for the three regimes $K \in [1, C_5 H]$, $K \in (C_5 H, \frac{100H}{p}]$ and $K \in (\frac{100H}{p}, \infty]$.

G.3 Proof of Theorem 6.3

When $K \geq SAH/p$, the lower bound in Theorem 6.1 readily applies because the regret is at least $\Omega(\sqrt{SAH^3 K})$ and the variance var is at most pH^2 . When $SAH \leq K \leq SAH/p$, the regret is at least $\Omega(SAH^2) = \Omega(\min\{\sqrt{SAH^3 K} + SAH^2, KH\})$. As a result, it suffices to focus on the case where $1 \leq K \leq SAH$. Toward this end, we only need the following lemma, which suffices for us to complete the proof.

LEMMA G.2. *Consider any $1 \leq K \leq SAH$. There exists an MDP instance with S states, A actions, horizon H , and $\text{var}_1 = \text{var}_2 = 0$, such that the regret is at least $\Omega(KH)$.*

PROOF. Let us construct an MDP with deterministic transition; more precisely, for each (s, a, h) , there is some s' such that $P_{s, a, h, s'} = 1$ and $P_{s, a, h, s''} = 0$ for any $s'' \neq s'$. The reward function is also chosen to be deterministic. In this case, it is easy to verify that $\text{var}_1 = \text{var}_2 = 0$.

We first assume $S = 2$. For any action a and horizon h , we set $P_{s_2, a, h} = e_{s_2}$ and $r_h(s_2, a) = 0$. For any action $a \neq a^*$ and h , we also set $P_{s_1, a, h} = e_{s_2}$ and $r_h(s_2, a) = 0$. At last, we set $P_{s_1, a^*, h} = e_{s_1}$ and $r_h(s_1, a^*) = 1$. In other words, there are a dumb state and a normal state in each step. The learner would naturally hope to find the correct action to avoid the dumb state. Obviously, $V_1^*(s_1) = H$. To find an $\frac{H}{2}$ -optimal policy, the learner needs to identify a^* for the first $\frac{H}{2}$ steps, requiring at least $\Omega(HA)$ rounds in expectation. As a result, the minimax regret is at least $\Omega(KH)$ when $K \leq cHA$ for some proper constant $c > 0$.

Let us refer to the hard instance above as a *hard chain*. For general S , we can construct $d := \frac{S}{2}$ hard chains. Let the two states in the i th hard chain be $(s_1(i), s_2(i))$. We set the initial distribution to be the uniform distribution over $\{s_1(i)\}_{i=1}^d$. Then $V_1^*(s_1(i)) = H$ holds for any $1 \leq i \leq d$. Let $\text{Regret}_i(K)$ be the expected regret resulting from the i th hard chain. When $K \geq 100S$, Lemma A.3 tells us that with probability at least $\frac{1}{2}$, $s_1(i)$ is visited for at least $\frac{K}{10S} \geq 10$ times. As a result, we have

$$\text{Regret}_i(K) \geq \frac{1}{2} \cdot \Omega \left(\frac{KH}{S} \right).$$

Summing over i , we see that the total regret is at least $\sum_{i=1}^d \text{Regret}_i(K) = \Omega(KH)$. When $K < 100S$, with probability at least $1 - (1 - \frac{1}{S})^K \geq 0.0001 \frac{K}{S}$, we know that $s_1(i)$ is visited for at least one time. Therefore, it holds that $\text{Regret}_i(K) \geq \Omega(\frac{KH}{S})$. Summing over i , we obtain

$$\text{Regret}(K) = \sum_{i=1}^K \text{Regret}_i(K) = \Omega(KH)$$

as claimed. \square

Received 11 June 2024; revised 19 February 2025; accepted 23 April 2025