

# Pessimistic Q-Learning for Offline Reinforcement Learning: Towards Optimal Sample Complexity

Laixi Shi\*  
CMU

Gen Li†  
UPenn

Yuting Wei†  
UPenn

Yuxin Chen†‡  
UPenn

Yuejie Chi\*  
CMU

February 28, 2022

## Abstract

Offline or batch reinforcement learning seeks to learn a near-optimal policy using history data without active exploration of the environment. To counter the insufficient coverage and sample scarcity of many offline datasets, the principle of pessimism has been recently introduced to mitigate high bias of the estimated values. While pessimistic variants of model-based algorithms (e.g., value iteration with lower confidence bounds) have been theoretically investigated, their model-free counterparts — which do not require explicit model estimation — have not been adequately studied, especially in terms of sample efficiency. To address this inadequacy, we study a pessimistic variant of Q-learning in the context of finite-horizon Markov decision processes, and characterize its sample complexity under the single-policy concentrability assumption which does not require the full coverage of the state-action space. In addition, a variance-reduced pessimistic Q-learning algorithm is proposed to achieve near-optimal sample complexity. Altogether, this work highlights the efficiency of model-free algorithms in offline RL when used in conjunction with pessimism and variance reduction.

**Keywords:** offline/batch reinforcement learning, Q-learning, pessimism, variance reduction

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
1.1	Main contributions . . . . .	2
1.2	Related works . . . . .	3
1.3	Notation and paper organization . . . . .	4
<b>2</b>	<b>Background and problem formulation</b>	<b>5</b>
2.1	Tabular finite-horizon MDPs . . . . .	5
2.2	Offline RL under single-policy concentrability . . . . .	6
<b>3</b>	<b>Pessimistic Q-learning: algorithms and theory</b>	<b>7</b>
3.1	LCB-Q: a natural pessimistic variant of Q-learning . . . . .	7
3.2	Theoretical guarantees for LCB-Q . . . . .	8
3.3	LCB-Q-Advantage for near-optimal offline RL . . . . .	9
3.4	Theoretical guarantees for LCB-Q-Advantage . . . . .	11
<b>4</b>	<b>Analysis</b>	<b>12</b>
4.1	Analysis of LCB-Q . . . . .	12
4.2	Analysis of LCB-Q-Advantage . . . . .	16
<b>5</b>	<b>Discussions</b>	<b>20</b>

---

\*Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

†Department of Statistics and Data Science, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA.

‡Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA.

<b>A</b>	<b>Technical lemmas</b>	<b>20</b>
A.1	Preliminary facts . . . . .	20
A.2	Freedman’s inequality and its consequences . . . . .	21
<b>B</b>	<b>Proof of main lemmas for LCB-Q (Theorem 1)</b>	<b>25</b>
B.1	Proof of Lemma 1 . . . . .	25
B.2	Proof of Lemma 2 . . . . .	28
B.3	Proof of Lemma 3 . . . . .	30
<b>C</b>	<b>Proof of lemmas for LCB-Q-Advantage (Theorem 2)</b>	<b>32</b>
C.1	Proof of Lemma 4 . . . . .	33
C.2	Proof of Lemma 5 . . . . .	37
C.3	Proof of Lemma 6 . . . . .	42
C.4	Proof of Lemma 12 . . . . .	54

# 1 Introduction

Reinforcement Learning (RL) has achieved remarkable success in recent years, including matching or surpassing human performance in robotics control and strategy games (Mnih et al., 2015; Silver et al., 2017). Nevertheless, these success stories often come with nearly prohibitive cost, where an astronomical number of samples are required to train the learning algorithm to a satisfactory level. Scaling up and replicating the RL success in many real-world problems face considerable challenges, due to limited access to large-scale simulation data. In applications such as online advertising and clinical trials, real-time data collection could be expensive, time-consuming, or constrained in sample sizes as a result of experimental limitations.

On the other hand, it is worth noting that tons of samples might have already been accumulated and stored — albeit not necessarily with the desired quality — during previous data acquisition attempts. It is therefore natural to wonder whether such history data can be leveraged to improve performance in future deployments. In reality, the history data was often obtained by executing some (possibly unknown) behavior policy, which is typically not the desired policy. This gives rise to the problem of offline RL or batch RL (Lange et al., 2012; Levine et al., 2020),<sup>1</sup> namely, how to make the best use of history data to learn an improved or even optimal policy, without further exploring the environment. In stark contrast to online RL that relies on active interaction with the environment, the performance of offline RL depends critically not only on the quantity, but also the quality of history data (e.g., coverage over the space-action space), given that the agent is no longer collecting new samples for the purpose of exploring the unknown environment.

Recently, the principle of pessimism (or conservatism) — namely, being conservative in Q-function estimation when there are not enough samples — has been put forward as an effective way to solve offline RL (Buckman et al., 2020; Kumar et al., 2020). This principle has been implemented in, for instance, a model-based offline value iteration algorithm, which modifies classical value iteration (Azar et al., 2017) by subtracting a penalty term in the estimated Q-values and has been shown to achieve appealing sample efficiency (Jin et al., 2021; Rashidinejad et al., 2021; Xie et al., 2021b). It is noteworthy that the model-based approach is built upon the construction of an empirical transition kernel, and therefore, requires specific representation of the environment (see, e.g. Agarwal et al. (2020); Li et al. (2020)). It remains unknown whether the pessimism principle can be incorporated into model-free algorithms — another class of popular algorithms that performs learning without model estimation — in a provably effective fashion for offline RL.

## 1.1 Main contributions

In this paper, we consider finite-horizon non-stationary Markov decision processes (MDPs) with  $S$  states,  $A$  actions, and horizon length  $H$ . The focal point is to pin down the sample efficiency for pessimistic variants of model-free algorithms, under the mild single-policy concentrability assumption (cf. Assumption 1) of the batch dataset introduced in Rashidinejad et al. (2021); Xie et al. (2021b) (in short, this assumption captures how close the batch dataset is to an expert dataset, and will be formally introduced in Section 2.2). Given

<sup>1</sup>Throughout this paper, we will be using the term offline RL (resp. dataset) or batch RL (resp. dataset) interchangeably.

Algorithm	Type	Sample complexity
VI-LCB (Xie et al., 2021b)	model-based	$\frac{H^6 SC^*}{\varepsilon^2}$
PEVI-Adv (Xie et al., 2021b)	model-based	$\frac{H^4 SC^*}{\varepsilon^2}$
Q-LCB (this work)	model-free	$\frac{H^6 SC^*}{\varepsilon^2}$
Q-LCB-Adv (this work)	model-free	$\frac{H^4 SC^*}{\varepsilon^2}$
lower bound (Xie et al., 2021b)	n/a	$\frac{H^4 SC^*}{\varepsilon^2}$

Table 1: Comparisons between our results and prior art for finding an  $\varepsilon$ -optimal policy in finite-horizon non-stationary MDPs. The sample complexities included in the table are valid for sufficiently small  $\varepsilon$ , with all logarithmic factors omitted.

$K$  episodes of history data each of length  $H$  (which amounts to a total number of  $T = KH$  samples), our main contributions are summarized as follows.

- We first study a natural pessimistic variant of the Q-learning algorithm, which simply modifies the classical Q-learning update rule by subtracting a penalty term (via certain lower confidence bounds). We prove that pessimistic Q-learning finds an  $\varepsilon$ -optimal policy as soon as the sample size  $T$  exceeds the order of (up to log factor)

$$\frac{H^6 SC^*}{\varepsilon^2},$$

where  $C^*$  denotes the single-policy concentrability coefficient of the batch dataset. In comparison to the minimax lower bound  $\Omega(\frac{H^4 SC^*}{\varepsilon^2})$  developed in Xie et al. (2021b), the sample complexity of pessimistic Q-learning is at most a factor of  $H^2$  from optimal (modulo some log factor).

- To further improve the sample efficiency of pessimistic model-free algorithms, we introduce a variance-reduced variant of pessimistic Q-learning. This algorithm is guaranteed to find an  $\varepsilon$ -optimal policy as long as the sample size  $T$  is above the order of

$$\frac{H^4 SC^*}{\varepsilon^2} + \frac{H^5 SC^*}{\varepsilon}$$

up to some log factor. In particular, this sample complexity is minimax-optimal (namely, as low as  $\frac{H^4 SC^*}{\varepsilon^2}$  up to log factor) for small enough  $\varepsilon$  (namely,  $\varepsilon \leq (0, 1/H]$ ). The  $\varepsilon$ -range that enjoys near-optimality is much larger compared to  $\varepsilon \leq (0, 1/H^{2.5}]$  established in Xie et al. (2021b) for model-based algorithms.

Both of the proposed algorithms achieve low computation cost (i.e.,  $O(T)$ ) and low memory complexities (i.e.,  $O(\min\{T, SAH\})$ ). Additionally, more complete comparisons with prior sample complexities of pessimistic model-based algorithms (Xie et al., 2021b) are provided in Table 1. In comparison with model-based algorithms, model-free algorithms require drastically different technical tools to handle the complicated statistical dependency between the estimated Q-values at different time steps.

## 1.2 Related works

In this section, we discuss several lines of works which are related to ours, with an emphasis on value-based algorithms for tabular settings with finite state and action spaces.

**Offline RL.** One of the key challenges in offline RL lies in the insufficient coverage of the batch dataset, due to lack of interaction with the environment (Levine et al., 2020; Liu et al., 2020). To address this challenge, most of the recent works can be divided into two lines: 1) regularizing the policy to avoid visiting under-covered state and action pairs (Dadashi et al., 2021; Fujimoto et al., 2019); 2) penalizing the estimated values of the under-covered state-action pairs (Buckman et al., 2020; Kumar et al., 2020). Our work follows the latter line (also known as the principle of pessimism), which has garnered significant attention recently. In fact, pessimism has been incorporated into recent development of various offline RL approaches, such as policy-based approaches (Rezaeifar et al., 2021; Xie et al., 2021a; Zanette et al., 2021), model-based approaches (Jin et al., 2021; Kidambi et al., 2020; Rashidinejad et al., 2021; Uehara and Sun, 2021; Uehara et al., 2021; Xie et al., 2021b; Yin and Wang, 2021; Yin et al., 2022; Yu et al., 2021b, 2020), and model-free approaches (Kumar et al., 2020; Yan et al., 2022; Yu et al., 2021a).

**Finite-sample guarantees for pessimistic approaches.** While model-free approaches with pessimism (Kumar et al., 2020; Yu et al., 2021a) have achieved considerable empirical successes in offline RL, prior theoretical guarantees of pessimistic schemes have been confined almost exclusively to model-based approaches. Under the same single-policy concentrability assumption used in prior analyses of model-based approaches (Rashidinejad et al., 2021; Xie et al., 2021b; Yin et al., 2021a), the current paper provides the first finite-sample guarantees for model-free approaches with pessimism in the tabular case without explicit model construction. In addition, Yin and Wang (2021) directly employed the occupancy distributions of the behavior policy and the optimal policy in bounding the performance of a model-based approach, rather than the worst-case upper bound of their ratios as done under the single-policy concentrability assumption.

**Non-asymptotic guarantees for variants of Q-learning.** Q-learning, which is among the most famous model-free RL algorithms (Jaakkola et al., 1994; Watkins and Dayan, 1992; Watkins, 1989), has been adapted in a multitude of ways to deal with different RL settings. Theoretical analyses for Q-learning and its variants have been established in, for example, the online setting via regret analysis (Bai et al., 2019; Dong et al., 2019; Jafarnia-Jahromi et al., 2020; Jin et al., 2018; Li et al., 2021b; Yang et al., 2021; Zhang et al., 2020a,b,c), and the simulator setting via probably approximately correct (PAC) bounds (Chen et al., 2020; Li et al., 2021a; Wainwright, 2019). The variant that is most closely related to ours is asynchronous Q-learning, which aims to find the optimal Q-function from Markovian trajectories following some behavior policy (Beck and Srikant, 2012; Even-Dar and Mansour, 2003; Li et al., 2021c; Qu and Wierman, 2020; Yin et al., 2021a,a,b). Different from ours, these works typically require full coverage of the state-action space by the behavior policy, a much stronger assumption than the single-policy concentrability assumed in our offline RL setting.

**Variance reduction in RL.** Variance reduction, originally proposed to accelerate stochastic optimization (e.g., the SVRG algorithm proposed by Johnson and Zhang (2013)), has been successfully leveraged to improve the sample efficiency of various RL algorithms, including but not limited to policy evaluation (Du et al., 2017; Khamaru et al., 2020; Wai et al., 2019; Xu et al., 2019), planning (Sidford et al., 2018a,b), Q-learning and its variants (Li et al., 2021b,c; Wainwright, 2019; Zhang et al., 2020b), and offline RL (Xie et al., 2021b; Yin et al., 2021a).

### 1.3 Notation and paper organization

Let us introduce a set of notation that will be used throughout. We denote by  $\Delta(\mathcal{S})$  the probability simplex over a set  $\mathcal{S}$ , and introduce the notation  $[N] := \{1, \dots, N\}$  for any integer  $N > 0$ . For any vector  $x \in \mathbb{R}^{\mathcal{S}A}$  (resp.  $x \in \mathbb{R}^{\mathcal{S}}$ ) that constitutes certain values for each of the state-action pairs (resp. state), we shall often use  $x(s, a)$  (resp.  $x(s)$ ) to denote the entry associated with the  $(s, a)$  pair (resp. state  $s$ ). Similarly, we shall denote by  $x := \{x_h\}_{h \in [H]}$  the set composed of certain vectors for each of the time step  $h \in [H]$ . We let  $e_i$  represent the  $i$ -th standard basis vector, with the only non-zero element being in the  $i$ -th entry.

Let  $\mathcal{X} := (\mathcal{S}, A, H, T)$ . The notation  $f(\mathcal{X}) \lesssim g(\mathcal{X})$  (resp.  $f(\mathcal{X}) \gtrsim g(\mathcal{X})$ ) means that there exists a universal constant  $C_0 > 0$  such that  $|f(\mathcal{X})| \leq C_0|g(\mathcal{X})|$  (resp.  $|f(\mathcal{X})| \geq C_0|g(\mathcal{X})|$ ). In addition, we often overload scalar functions and expressions to take vector-valued arguments, with the interpretation that they are applied in an entrywise manner. For example, for a vector  $x = [x_i]_{1 \leq i \leq n}$ , we have  $x^2 = [x_i^2]_{1 \leq i \leq n}$ .

For any two vectors  $x = [x_i]_{1 \leq i \leq n}$  and  $y = [y_i]_{1 \leq i \leq n}$ , the notation  $x \leq y$  (resp.  $x \geq y$ ) means  $x_i \leq y_i$  (resp.  $x_i \geq y_i$ ) for all  $1 \leq i \leq n$ .

**Paper organization.** The rest of this paper is organized as follows. Section 2 introduces the backgrounds on finite-horizon MDPs and formulates the offline RL problem. Section 3 starts by introducing a natural pessimistic variant of Q-learning along with its sample complexity bound, and further enhances the sample efficiency via variance reduction in Section 3.3. Section 4 presents the proof outline and key lemmas. Finally, we conclude in Section 5 with a discussion and defer the proof details to the supplementary material.

## 2 Background and problem formulation

### 2.1 Tabular finite-horizon MDPs

**Basics.** This work focuses on an episodic finite-horizon MDP as represented by

$$\mathcal{M} = (\mathcal{S}, \mathcal{A}, H, \{P_h\}_{h=1}^H, \{r_h\}_{h=1}^H),$$

where  $H$  is the horizon length,  $\mathcal{S}$  is a finite state space of cardinality  $S$ ,  $\mathcal{A}$  is a finite action space of cardinality  $A$ , and  $P_h : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$  (resp.  $r_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ ) represents the probability transition kernel (resp. reward function) at the  $h$ -th time step ( $1 \leq h \leq H$ ). Throughout this paper, we shall adopt the following convenient notation

$$P_{h,s,a} := P_h(\cdot | s, a) \in [0, 1]^{1 \times S}, \quad (1)$$

which stands for the transition probability vector given the current state-action pair  $(s, a)$  at time step  $h$ . The parameters  $S$ ,  $A$  and  $H$  can all be quite large, allowing one to capture the challenges arising in MDPs with large state/action space and long horizon.

A policy (or action selection rule) of an agent is represented by  $\pi = \{\pi_h\}_{h=1}^H$ , where  $\pi_h : \mathcal{S} \rightarrow \Delta(\mathcal{A})$  specifies the associated selection probability over the action space at time step  $h$  (or more precisely, we let  $\pi_h(a | s)$  represent the probability of selecting action  $a$  in state  $s$  at step  $h$ ). When  $\pi$  is a deterministic policy, we abuse the notation and let  $\pi_h(s)$  denote the action selected by policy  $\pi$  in state  $s$  at step  $h$ . In each episode, the agent generates an initial state  $s_1 \in \mathcal{S}$  drawn from an initial state distribution  $\rho \in \Delta(\mathcal{S})$ , and rolls out a trajectory over the MDP by executing a policy  $\pi$  as follows:

$$\{s_h, a_h, r_h\}_{h=1}^H = \{s_1, a_1, r_1, \dots, s_H, a_H, r_H\}, \quad (2)$$

where at time step  $h$ ,  $a_h \sim \pi_h(\cdot | s_h)$  indicates the action selected in state  $s_h$ ,  $r_h = r_h(s_h, a_h)$  denotes the deterministic immediate reward, and  $s_{h+1}$  denotes the next state drawn from the transition probability vector  $P_{h,s_h,a_h} := P_h(\cdot | s_h, a_h)$ . In addition, let  $d_h^\pi(s)$  and  $d_h^\pi(s, a)$  denote respectively the occupancy distribution induced by  $\pi$  at time step  $h \in [H]$ , namely,

$$d_h^\pi(s) := \mathbb{P}(s_h = s | s_1 \sim \rho, \pi), \quad d_h^\pi(s, a) := \mathbb{P}(s_h = s | s_1 \sim \rho, \pi) \pi_h(a | s); \quad (3)$$

here and throughout, we denote  $[H] := \{1, \dots, H\}$ . Given that the initial state  $s_1$  is drawn from  $\rho$ , the above definition gives

$$d_1^\pi(s) = \rho(s) \quad \text{for any policy } \pi. \quad (4)$$

**Value function, Q-function, and optimal policy.** The value function  $V_h^\pi(s)$  of policy  $\pi$  in state  $s$  at step  $h$  is defined as the expected cumulative rewards when this policy is executed starting from state  $s$  at step  $h$ , i.e.,

$$V_h^\pi(s) := \mathbb{E} \left[ \sum_{t=h}^H r_t(s_t, a_t) \mid s_h = s \right], \quad (5)$$

where the expectation is taken over the randomness of the trajectory (2) induced by the policy  $\pi$  as well as the MDP transitions. Similarly, the Q-function  $Q_h^\pi(\cdot, \cdot)$  of a policy  $\pi$  at step  $h$  is defined as

$$Q_h^\pi(s, a) := r_h(s, a) + \mathbb{E} \left[ \sum_{t=h+1}^H r_t(s_t, a_t) \mid s_h = s, a_h = a \right], \quad (6)$$

where the expectation is again over the randomness induced by  $\pi$  and the MDP except that the state-action pair at step  $h$  is now conditioned to be  $(s, a)$ . By convention, we shall also set

$$V_{H+1}^\pi(s) = Q_{H+1}^\pi(s, a) = 0 \quad \text{for any } \pi \text{ and } (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (7)$$

A policy  $\pi^* = \{\pi_h^*\}_{h=1}^H$  is said to be an optimal policy if it maximizes the value function (resp. Q-function) *simultaneously* for all states (resp. state-action pairs) among all policies, whose existence is always guaranteed (Puterman, 2014). The resulting optimal value function  $V^* = \{V_h^*\}_{h=1}^H$  and optimal Q-functions  $Q^* = \{Q_h^*\}_{h=1}^H$  are denoted respectively by

$$V_h^*(s) := V_h^{\pi^*}(s) = \max_{\pi} V_h^\pi(s), \quad Q_h^*(s, a) := Q_h^{\pi^*}(s, a) = \max_{\pi} Q_h^\pi(s, a)$$

for any  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ . Throughout this paper, we assume that  $\pi^*$  is a *deterministic optimal policy*, which always exists (Puterman, 2014).

Additionally, when the initial state is drawn from a given distribution  $\rho$ , the expected value of a given policy  $\pi$  and that of the optimal policy at the initial step are defined respectively by

$$V_1^\pi(\rho) := \mathbb{E}_{s_1 \sim \rho} [V_1^\pi(s_1)] \quad \text{and} \quad V_1^*(\rho) := \mathbb{E}_{s_1 \sim \rho} [V_1^*(s_1)]. \quad (8)$$

**Bellman equations.** The Bellman equations play a fundamental role in dynamic programming (Bertsekas, 2017). Specifically, the value function and the Q-function of any policy  $\pi$  satisfy the following Bellman consistency equation:

$$Q_h^\pi(s, a) = r_h(s, a) + \mathbb{E}_{s' \sim P_{h,s,a}} [V_{h+1}^\pi(s')] \quad (9)$$

for all  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ . Moreover, the optimal value function and the optimal Q-function satisfy the Bellman optimality equation:

$$Q_h^*(s, a) = r_h(s, a) + \mathbb{E}_{s' \sim P_{h,s,a}} [V_{h+1}^*(s')] \quad (10)$$

for all  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ .

## 2.2 Offline RL under single-policy concentrability

Offline RL assumes the availability of a history dataset  $\mathcal{D}_\mu$  containing  $K$  episodes each of length  $H$ . These episodes are independently generated based on a certain policy  $\mu = \{\mu_h\}_{h=1}^H$  — called the *behavior policy*, resulting in a dataset

$$\mathcal{D}_\mu := \left\{ (s_1^k, a_1^k, r_1^k, \dots, s_H^k, a_H^k, r_H^k) \right\}_{k=0}^{K-1}.$$

Here, the initial states  $\{s_1^k\}_{k=1}^K$  are independently drawn from  $\rho \in \Delta(\mathcal{S})$  such that  $s_1^k \stackrel{\text{i.i.d.}}{\sim} \rho$ , while the remaining states and actions are generated by the MDP induced by the behavior policy  $\mu$ . The total number of samples is thus given by

$$T = KH.$$

With the notation (8) in place, the goal of offline RL amounts to finding an  $\varepsilon$ -optimal policy  $\hat{\pi} = \{\hat{\pi}_h\}_{h=1}^H$  satisfying

$$V_1^*(\rho) - V_1^{\hat{\pi}}(\rho) \leq \varepsilon$$

with as few samples as possible, and ideally, in a computationally fast and memory-efficient manner.

Obviously, efficient offline RL cannot be accomplished without imposing proper assumptions on the behavior policy, which also provide means to gauge the difficulty of the offline RL task through the quality of the history dataset. Following the recent works [Rashidinejad et al. \(2021\)](#); [Xie et al. \(2021b\)](#), we assume that the behavior policy  $\mu$  satisfies the following property called *single-policy concentrability*.

**Assumption 1** (single-policy concentrability). *The single-policy concentrability coefficient  $C^* \in [1, \infty)$  of a behavior policy  $\mu$  is defined to be the smallest quantity that satisfies*

$$\max_{(h,s,a) \in [H] \times \mathcal{S} \times \mathcal{A}} \frac{d_h^{\pi^*}(s,a)}{d_h^\mu(s,a)} \leq C^*, \quad (11)$$

where we adopt the convention  $0/0 = 0$ .

Intuitively, the single-policy concentrability coefficient measures the discrepancy between the optimal policy  $\pi^*$  and the behavior policy  $\mu$  in terms of the resulting density ratio of the respective occupancy distributions. It is noteworthy that a finite  $C^*$  does not necessarily require  $\mu$  to cover the entire state-action space; instead, it can be attainable when its coverage subsumes that of the optimal policy  $\pi^*$ . This is in stark contrast to, and in fact much weaker than, other assumptions that require either full coverage of the behavior policy (i.e.,  $\min_{(h,s,a) \in [H] \times \mathcal{S} \times \mathcal{A}} d_h^\mu(s,a) > 0$  ([Li et al., 2021c](#); [Yin et al., 2021a,b](#))), or uniform concentrability over all possible policies ([Chen and Jiang, 2019](#)). Additionally, the single-policy concentrability coefficient is minimized (i.e.,  $C^* = 1$ ) when the behavior policy  $\mu$  coincides with the optimal policy  $\pi^*$ , a scenario closely related to imitation learning or behavior cloning ([Rajaraman et al., 2020](#)).

### 3 Pessimistic Q-learning: algorithms and theory

In the current paper, we present two model-free algorithms — namely, LCB-Q and LCB-Q-Advantage — for offline RL, along with their respective theoretical guarantees. The first algorithm can be viewed as a pessimistic variant of the classical Q-learning algorithm, while the second one further leverages the idea of variance reduction to boost the sample efficiency. In this section, we begin by introducing LCB-Q.

#### 3.1 LCB-Q: a natural pessimistic variant of Q-learning

Before proceeding, we find it convenient to first review the classical Q-learning algorithm ([Watkins and Dayan, 1992](#); [Watkins, 1989](#)), which can be regarded as a stochastic approximation scheme to solve the Bellman optimality equation (10). Upon receiving a sample transition  $(s_h, a_h, r_h, s_{h+1})$  at time step  $h$ , Q-learning updates the corresponding entry in the Q-estimate as follows

$$Q_h(s_h, a_h) \leftarrow (1 - \eta)Q_h(s_h, a_h) + \eta \left\{ r_h(s_h, a_h) + V_{h+1}(s_{h+1}) \right\}, \quad (12)$$

where  $Q_h$  (resp.  $V_h$ ) indicates the running estimate of  $Q_h^*$  (resp.  $V_h^*$ ), and  $0 < \eta < 1$  is the learning rate. In comparison to model-based algorithms that require estimating the probability transition kernel based on all the samples, Q-learning, as a popular kind of model-free algorithms, is simpler and enjoys more flexibility without explicitly constructing the model of the environment. The wide applicability of Q-learning motivates one to adapt it to accommodate offline RL.

Inspired by recent advances in incorporating the pessimism principle for offline RL ([Jin et al., 2021](#); [Rashidinejad et al., 2021](#)), we study a pessimistic variant of Q-learning called LCB-Q, which modifies the Q-learning update rule as follows

$$Q_h(s_h, a_h) \leftarrow (1 - \eta_n)Q_h(s_h, a_h) + \eta_n \left\{ r_h(s_h, a_h) + V_{h+1}(s_{h+1}) - b_n \right\}, \quad (13)$$

where  $\eta_n$  is the learning rate depending on the number of times  $n$  that the state-action pair  $(s_h, a_h)$  has been visited at step  $h$ , and the penalty term  $b_n > 0$  (cf. line 10 of Algorithm 1) reflects the uncertainty of the corresponding Q-estimate and implements pessimism in the face of uncertainty. The entire algorithm, which is a *single-pass* algorithm that only requires reading the offline dataset once, is summarized in Algorithm 1.



---

**Algorithm 1:** LCB-Q for offline RL

---

```
1 Parameters: some constant  $c_b > 0$ , target success probability  $1 - \delta \in (0, 1)$ , and  $\iota = \log\left(\frac{SAT}{\delta}\right)$ .
2 Initialize  $Q_h(s, a) \leftarrow 0$ ,  $N_h(s, a) \leftarrow 0$ , and  $V_h(s) \leftarrow 0$  for all  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H + 1]$ ;  $\hat{\pi}$  s.t.
    $\hat{\pi}_h(s) = 1$  for all  $(h, s) \in [H] \times \mathcal{S}$ .
3 for Episode  $k = 1$  to  $K$  do
4   Sample a new trajectory  $\{s_h, a_h, r_h\}_{h=1}^H$  from  $\mathcal{D}_\mu$ . // sampling from batch dataset
   // update the policy
5   for Step  $h = 1$  to  $H$  do
6      $N_h(s_h, a_h) \leftarrow N_h(s_h, a_h) + 1$ . // update the counter
7      $n \leftarrow N_h(s_h, a_h)$ ;  $\eta_n \leftarrow \frac{H+1}{H+n}$ . // update the learning rate
8      $b_n \leftarrow c_b \sqrt{\frac{H^3 \iota^2}{n}}$ . // update the bonus term
     // run the Q-learning update with LCB
9      $Q_h(s_h, a_h) \leftarrow Q_h(s_h, a_h) + \eta_n \{r_h(s_h, a_h) + V_{h+1}(s_{h+1}) - Q_h(s_h, a_h) - b_n\}$ .
     // update the value estimates
10     $V_h(s_h) \leftarrow \max \{V_h(s_h), \max_a Q_h(s_h, a)\}$ .
11    If  $V_h(s_h) = \max_a Q_h(s_h, a)$ : update  $\hat{\pi}_h(s) \leftarrow \arg \max_a Q_h(s, a)$ .
12 Output: the policy  $\hat{\pi}$ .
```

---

### 3.2 Theoretical guarantees for LCB-Q

The proposed LCB-Q algorithm manages to achieve an appealing sample complexity as formalized by the following theorem.

**Theorem 1.** Consider any  $\delta \in (0, 1)$ . Suppose that the behavior policy  $\mu$  satisfies Assumption 1 with single-policy concentrability coefficient  $C^* \geq 1$ . Let  $c_b > 0$  be some sufficiently large constant, and take  $\iota := \log\left(\frac{SAT}{\delta}\right)$ . Assume that  $T > SC^* \iota$ , then the policy  $\hat{\pi}$  returned by Algorithm 1 satisfies

$$V_1^*(\rho) - V_1^{\hat{\pi}}(\rho) \leq c_a \sqrt{\frac{H^6 SC^* \iota^3}{T}} \quad (14)$$

with probability at least  $1 - \delta$ , where  $c_a > 0$  is some universal constant.

As asserted by Theorem 1, the LCB-Q algorithm is guaranteed to find an  $\varepsilon$ -optimal policy with high probability, as long as the total sample size  $T = KH$  exceeds

$$\tilde{O}\left(\frac{H^6 SC^*}{\varepsilon^2}\right), \quad (15)$$

where  $\tilde{O}(\cdot)$  hides logarithmic dependencies. When the behavior policy is close to the optimal policy, the single-policy concentrability coefficient  $C^*$  is closer to 1; if this is the case, then our bound indicates that the sample complexity does not depend on the size  $A$  of the action space, which can be a huge saving when the action space is enormous.

**Comparison with model-based pessimistic approaches.** A model-based approach — called Value Iteration with Lower Confidence Bounds (VI-LCB) — has been recently proposed for offline RL (Rashidinejad et al., 2021; Xie et al., 2021b). In the finite-horizon case, VI-LCB incorporates an additional LCB penalty into the classical value iteration algorithm, and updates *all* the entries in the Q-estimate simultaneously as follows

$$Q_h(s, a) \leftarrow r_h(s, a) + \hat{P}_{h,s,a} V_{h+1} - b_h(s, a), \quad (16)$$

with the aim of tuning down the confidence on those state-action pairs that have only been visited infrequently. Here,  $\hat{P}_{h,s,a}$  represents the empirical estimation of the transition kernel  $P_{h,s,a}$ , and  $b_h(s, a) > 0$  is



---

**Algorithm 2:** Offline LCB-Q-Advantage RL

---

```
1 Parameters: number of epochs  $M$ , universal constant  $c_b > 0$ , probability of failure  $\delta \in (0, 1)$ , and  
    $\iota = \log\left(\frac{SAT}{\delta}\right)$ ;  
2 Initialize:  
3  $Q_h(s, a), Q_h^{\text{LCB}}(s, a), \bar{Q}_h(s, a), \bar{\mu}_h(s, a), \bar{\mu}_h^{\text{next}}(s, a), N_h(s, a) \leftarrow 0$  for all  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ ;  
4  $V_h(s), \bar{V}_h(s), \bar{V}_h^{\text{next}}(s) \leftarrow 0$  for all  $(s, h) \in \mathcal{S} \times [H + 1]$ ;  
5  $\mu_h^{\text{ref}}(s, a), \sigma_h^{\text{ref}}(s, a), \mu_h^{\text{adv}}(s, a), \sigma_h^{\text{adv}}(s, a), \bar{\delta}_h(s, a), \bar{B}_h(s, a) \leftarrow 0$  for all  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ .  
6 for Epoch  $m = 1$  to  $M$  do  
7    $L_m = 2^m$ ; // specify the number of episodes in the current epoch  
8    $\hat{N}_h(s, a) = 0$  for all  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . // reset the epoch-wise counter  
   /* Inner-loop: update value-estimates  $V_h(s, a)$  and Q-estimates  $Q_h(s, a)$  */  
9   for In-epoch Episode  $t = 1$  to  $L_m$  do  
10    Sample a new trajectory  $\{s_h, a_h, r_h\}_{h=1}^H$ . // sampling from batch dataset  
11    for Step  $h = 1$  to  $H$  do  
12       $N_h(s_h, a_h) \leftarrow N_h(s_h, a_h) + 1$ ;  $n \leftarrow N_h(s_h, a_h)$ . // update the overall counter  
13       $\eta_n \leftarrow \frac{H+1}{H+n}$ ; // update the learning rate  
      // run the Q-learning update rule with LCB  
14       $Q_h^{\text{LCB}}(s_h, a_h) \leftarrow \text{update-lcb-q}()$ .  
      // update the Q-estimate with LCB and reference-advantage  
15       $\bar{Q}_h(s_h, a_h) \leftarrow \text{update-lcb-q-ra}()$ .  
      // update the Q-estimate  $Q_h$  and value estimate  $V_h$   
16       $Q_h(s_h, a_h) \leftarrow \max\{Q_h^{\text{LCB}}(s_h, a_h), \bar{Q}_h(s_h, a_h), Q_h(s_h, a_h)\}$ .  
17       $V_h(s_h) \leftarrow \max_a Q_h(s_h, a)$ .  
      // update the epoch-wise counter and  $\bar{\mu}_h^{\text{next}}$  for the next epoch  
18       $\hat{N}_h(s_h, a_h) \leftarrow \hat{N}_h(s_h, a_h) + 1$ ;  
19       $\bar{\mu}_h^{\text{next}}(s_h, a_h) \leftarrow \left(1 - \frac{1}{\hat{N}_h(s_h, a_h)}\right) \bar{\mu}_h^{\text{next}}(s_h, a_h) + \frac{1}{\hat{N}_h(s_h, a_h)} \bar{V}_{h+1}^{\text{next}}(s_{h+1})$ ;  
   /* Update the reference  $(\bar{V}_h, \bar{V}_h^{\text{next}})$  and  $(\bar{\mu}_h, \bar{\mu}_h^{\text{next}})$  */  
20   for  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H + 1]$  do  
21      $\bar{V}_h(s) \leftarrow \bar{V}_h^{\text{next}}(s)$ ;  $\bar{\mu}_h(s, a) \leftarrow \bar{\mu}_h^{\text{next}}(s, a)$ . // set  $\bar{V}_h$  and  $\bar{\mu}_h$  for the next epoch  
22      $\bar{V}_h^{\text{next}}(s) \leftarrow V_h(s)$ ;  $\bar{\mu}_h^{\text{next}}(s, a) \leftarrow 0$ . // restart  $\bar{\mu}_h^{\text{next}}$  and set  $\bar{V}_h^{\text{next}}$  for the next epoch  
Output: the policy  $\hat{\pi}$  s.t.  $\hat{\pi}_h(s) = \arg \max_a Q_h(s, a)$  for any  $(s, h) \in \mathcal{S} \times [H]$ .
```

---

chosen to capture the uncertainty level of  $(\hat{P}_{h,s,a} - P_{h,s,a})V_{h+1}$ . Working backward, the algorithm estimates the Q-value  $Q_h$  recursively over the time steps  $h = H, H-1, \dots, 1$ . In comparison with VI-LCB, our sample complexity bound for LCB-Q matches the bound developed for VI-LCB by Xie et al. (2021b), while enjoying enhanced flexibility without the need of specifying the transition kernel of the environment (as model estimation might potentially incur a higher memory burden).

### 3.3 LCB-Q-Advantage for near-optimal offline RL

The careful reader might notice that the sample complexity (15) derived for LCB-Q remains a factor of  $H^2$  away from the minimax lower bound (see Table 1). To further close the gap and improve the sample complexity, we propose a new variant called LCB-Q-Advantage, which leverages the idea of variance reduction to accelerate convergence (Johnson and Zhang, 2013; Li et al., 2021b,c; Sidford et al., 2018b; Wainwright, 2019; Xie et al., 2021b; Zhang et al., 2020b).

Inspired by the reference-advantage decomposition adopted in (Li et al., 2021b; Zhang et al., 2020b) for online Q-learning, LCB-Q-Advantage maintains a collection of reference values  $\{\bar{V}_h\}_{h=1}^H$ , which serve as running proxy for the optimal values  $\{V_h^*\}_{h=1}^H$  and allow for reduced variability in each iteration. To be




Figure 1: An illustration of the epoch-based LCB-Q-Advantage algorithm.

more specific, the LCB-Q-Advantage algorithm (cf. Algorithm 2 as well as the subroutines in Algorithm 3 that closely resemble Li et al. (2021b)) proceeds in an epoch-based style (the  $m$ -th epoch consists of  $L_m = 2^m$  episodes of samples), where the reference values are updated at the end of each epoch to be used in the next epoch, and the Q-estimates are iteratively updated during the remaining time of each epoch. By maintaining two auxiliary sequences of *pessimistic* Q-estimates — that is,  $Q^{\text{LCB}}$  constructed by the pessimistic Q-learning update, and  $\bar{Q}$  constructed by the pessimistic Q-learning update based on the reference-advantage decomposition — the Q-estimate is updated by taking the maximum over the three candidates (cf. line 16 of Algorithm 2)

$$Q_h(s, a) \leftarrow \max\{Q_h^{\text{LCB}}(s, a), \bar{Q}_h(s, a), Q_h(s, a)\} \quad (17)$$

when the state-action pair  $(s, a)$  is visited at the step  $h$ . We now take a moment to discuss the key ingredients of the proposed algorithm in further detail.

**Updating the references  $\bar{V}_h$  and  $\bar{\mu}_h$ .** At the end of each epoch, the reference values  $\{\bar{V}_h\}_{h=1}^H$ , as well as the associated running average  $\{\bar{\mu}_h\}_{h=1}^H$ , are determined using what happens during the current epoch. More specifically, the following update rules for  $\bar{V}_h$  and  $\bar{\mu}_h$  are carried out at the end of the  $m$ -th epoch:

$$\bar{V}_h(s) \leftarrow \bar{V}_h^{\text{next}}(s), \quad (18a)$$

$$\bar{\mu}_h(s, a) \leftarrow \frac{\sum_{t=1}^{L_m} \mathbb{1}(s_h^t = s, a_h^t = a) \bar{V}_{h+1}(s_{h+1}^t)}{\max\left\{\left\{\sum_{t=1}^{L_m} \mathbb{1}(s_h^t = s, a_h^t = a)\right\}, 1\right\}} \quad (18b)$$

for all  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . Here,  $\bar{V}_h(s)$  is assigned by  $\bar{V}_h^{\text{next}}(s)$ , which is maintained as the value estimate  $V_h(s)$  at the end of the  $(m-1)$ -th epoch, and the update of  $\bar{\mu}_h(s, a)$  is implemented in a recursive manner in the current  $m$ -th epoch. See also line 21 and line 19 of Algorithm 2.

**Learning Q-estimate  $\bar{Q}_h$  based on the reference-advantage decomposition.** Armed with the references  $\bar{V}_h$  and  $\bar{\mu}_h$  updated at the end of the previous  $(m-1)$ -th epoch, LCB-Q-Advantage iteratively updates the Q-estimate  $\bar{Q}_h$  in all episodes during the  $m$ -th epoch. At each time step  $h$  in any episode, whenever  $(s, a)$  is visited, LCB-Q-Advantage updates the reference Q-value as follows:

$$\bar{Q}_h(s, a) \leftarrow (1 - \eta) \bar{Q}_h(s, a) + \eta \left\{ r_h(s, a) + \underbrace{\hat{P}_{h,s,a}(V_{h+1} - \bar{V}_{h+1})}_{\text{estimate of } P_{h,s,a}(V_{h+1} - \bar{V}_{h+1})} + \underbrace{\bar{\mu}_h}_{\text{estimate of } P_{h,s,a} \bar{V}_{h+1}} - \bar{b}_h(s, a) \right\}. \quad (19)$$

Intuitively, we decompose the target  $P_{h,s,a} V_{h+1}$  into a reference part  $P_{h,s,a} \bar{V}_{h+1}$  and an advantage part  $P_{h,s,a} (V_{h+1} - \bar{V}_{h+1})$ , and cope with the two parts separately. In the sequel, let us take a moment to discuss three essential ingredients of the update rule (19), which shed light on the design rationale of our algorithm.

- Akin to LCB-Q, the term  $\hat{P}_{h,s,a}(V_{h+1} - \bar{V}_{h+1})$  serves as an unbiased stochastic estimate of  $P_{h,s,a}(V_{h+1} - \bar{V}_{h+1})$  if a sample transition  $(s, a, s_{h+1})$  at time step  $h$  is observed. If  $V_{h+1}$  stays close to the reference  $\bar{V}_{h+1}$  as the algorithm proceeds, the variance of this stochastic term can be lower than that of the stochastic term  $\hat{P}_{h,s,a} V_{h+1}$  in (13).

- The auxiliary estimate  $\bar{\mu}_h$  introduced in (18b) serves as a running estimate of the reference part  $P_{h,s,a}\bar{V}_{h+1}$ . Based on the update rule (18b), we design  $\bar{\mu}_h(s, a)$  to estimate the running mean of the reference part  $[P_{h,s,a}\bar{V}_{h+1}]$  using a number of previous samples. As a result, we expect the variability of this term to be well-controlled, particularly as the number of samples in each epoch grows exponentially (recall that  $L_m = 2^m$ ).
- In each episode, the term  $\bar{b}_h(s, a)$  serves as the additional confidence bound on the error between the estimates of the reference/advantage and the ground truth. More specifically,  $\mu_h^{\text{ref}}(s, a)$  and  $\sigma_h^{\text{ref}}(s, a)$  are respectively the running mean and 2nd moment of the reference part  $[P_{h,s,a}\bar{V}_{h+1}]$  (cf. lines 9-10 of Algorithm 3);  $\mu_h^{\text{adv}}(s, a)$  and  $\sigma_h^{\text{adv}}(s, a)$  represent respectively the running mean and 2nd moment of the advantage part  $[P_{h,s,a}(V_{h+1} - \bar{V}_{h+1})]$  (cf. lines 11-12 of Algorithm 3);  $\bar{B}_h(s, a)$  aggregates the empirical standard deviations of the reference and the advantage parts. The LCB penalty term  $\bar{b}_h(s, a)$  is updated using  $\bar{B}_h(s, a)$  and  $\bar{\delta}_h(s_h, a_h)$  (cf. lines 5-6 of Algorithm 3), taking into account the confidence bounds for both the reference and the advantage.

In a nutshell, the auxiliary sequences of the reference values are designed to help reduce the variance of the stochastic Q-learning updates, which taken together with the principle of pessimism play a crucial role in the improvement of sample complexity for offline RL.

### 3.4 Theoretical guarantees for LCB-Q-Advantage

Encouragingly, the proposed LCB-Q-Advantage algorithm provably achieves near-optimal sample complexity for sufficiently small  $\varepsilon$ , as demonstrated by the following theorem.

**Theorem 2.** *Consider any  $\delta \in (0, 1)$ , and recall that  $\iota = \log(\frac{SAT}{\delta})$  and  $T = KH$ . Suppose that  $c_b > 0$  is chosen to be a sufficiently large constant, and that the behavior policy  $\mu$  satisfies Assumption 1. Then there exists some universal constant  $c_g > 0$  such that with probability at least  $1 - \delta$ , the policy  $\hat{\pi}$  output by Algorithm 2 satisfies*

$$V_1^*(\rho) - V_1^{\hat{\pi}}(\rho) \leq c_g \left( \sqrt{\frac{H^4 SC^* \iota^5}{T}} + \frac{H^5 SC^* \iota^4}{T} \right). \quad (20)$$

As a consequence, Theorem 2 reveals that the LCB-Q-Advantage algorithm is guaranteed to find an  $\varepsilon$ -optimal policy (i.e.,  $V_1^*(\rho) - V_1^{\hat{\pi}}(\rho) \leq \varepsilon$ ) as long as the total sample size  $T$  exceeds

$$\tilde{O} \left( \frac{H^4 SC^*}{\varepsilon^2} + \frac{H^5 SC^*}{\varepsilon} \right). \quad (21)$$

For sufficiently small accuracy level  $\varepsilon$  (i.e.,  $\varepsilon \leq 1/H$ ), this results in a sample complexity of

$$\tilde{O} \left( \frac{H^4 SC^*}{\varepsilon^2} \right), \quad (22)$$

thereby matching the minimax lower bound developed in Xie et al. (2021b) up to logarithmic factor. Compared with the minimax lower bound  $\Omega(\frac{H^4 SA}{\varepsilon^2})$  in the online RL setting (Domingues et al., 2021), this suggests that offline RL can be fairly sample-efficient when the behavior policy closely mimics the optimal policy in terms of the resulting state-action occupancy distribution (a scenario where  $C^*$  is potentially much smaller than the size of the action space).

**Comparison with offline model-based approaches.** In the same offline finite-horizon setting, the state-of-art model-based approach called PEVI-Adv has been proposed by Xie et al. (2021b), which also leverage the idea of reference-advantage decomposition. In comparison with PEVI-Adv, LCB-Q-Advantage not only enjoys the flexibility of model-free approaches, but also achieves optimal sample complexity for a broader range of target accuracy level  $\varepsilon$ . More precisely, the  $\varepsilon$ -range for which the algorithm achieves sample optimality can be compared as follows:

$$\underbrace{\varepsilon \leq (0, H^{-1}]}_{\text{(Our LCB-Q-Advantage)}} \quad \text{vs.} \quad \underbrace{\varepsilon \leq (0, H^{-2.5}]}_{\text{(PEVI-Adv)}}, \quad (23)$$

offering an improvement by a factor of  $H^{1.5}$ .

---

**Algorithm 3:** Auxiliary functions

---

```

1 Function update-lcb-q():
2    $Q_h^{\text{LCB}}(s_h, a_h) \leftarrow (1 - \eta_n)Q_h^{\text{LCB}}(s_h, a_h) + \eta_n(r(s_h, a_h) + V_{h+1}(s_{h+1}) - c_b \sqrt{\frac{H^3 \iota^2}{n}}).$ 
3 Function update-lcb-q-ra():
4    $[\mu_h^{\text{ref}}, \sigma_h^{\text{ref}}, \mu_h^{\text{adv}}, \sigma_h^{\text{adv}}](s_h, a_h) \leftarrow \text{update-moments}();$ 
5    $[\bar{\delta}_h, \bar{B}_h](s_h, a_h) \leftarrow \text{update-bonus}();$ 
6    $\bar{b}_h(s_h, a_h) \leftarrow \bar{B}_h(s_h, a_h) + (1 - \eta_n) \frac{\bar{\delta}_h(s_h, a_h)}{\eta_n} + c_b \frac{H^{7/4} \iota}{n^{3/4}} + c_b \frac{H^2 \iota}{n};$ 
7    $\bar{Q}_h(s_h, a_h) \leftarrow (1 - \eta_n)\bar{Q}_h(s_h, a_h) + \eta_n(r_h(s_h, a_h) + V_{h+1}(s_{h+1}) - \bar{V}_{h+1}(s_{h+1}) + \bar{\mu}_h(s_h, a_h) - \bar{b}_h);$ 
8 Function update-moments():
9    $\mu_h^{\text{ref}}(s_h, a_h) \leftarrow (1 - \frac{1}{n})\mu_h^{\text{ref}}(s_h, a_h) + \frac{1}{n}\bar{V}_{h+1}^{\text{next}}(s_{h+1});$  // mean of the reference
10   $\sigma_h^{\text{ref}}(s_h, a_h) \leftarrow (1 - \frac{1}{n})\sigma_h^{\text{ref}}(s_h, a_h) + \frac{1}{n}(\bar{V}_{h+1}^{\text{next}}(s_{h+1}))^2;$  // 2nd moment of the reference
11   $\mu_h^{\text{adv}}(s_h, a_h) \leftarrow (1 - \eta_n)\mu_h^{\text{adv}}(s_h, a_h) + \eta_n(V_{h+1}(s_{h+1}) - \bar{V}_{h+1}(s_{h+1}));$  // mean of the advantage
12   $\sigma_h^{\text{adv}}(s_h, a_h) \leftarrow (1 - \eta_n)\sigma_h^{\text{adv}}(s_h, a_h) + \eta_n(V_{h+1}(s_{h+1}) - \bar{V}_{h+1}(s_{h+1}))^2.$  // 2nd moment of the
    advantage
13 Function update-bonus():
14   $B_h^{\text{next}}(s_h, a_h) \leftarrow c_b \sqrt{\frac{\iota}{n}} \left( \sqrt{\sigma_h^{\text{ref}}(s_h, a_h) - (\mu_h^{\text{ref}}(s_h, a_h))^2} + \sqrt{H} \sqrt{\sigma_h^{\text{adv}}(s_h, a_h) - (\mu_h^{\text{adv}}(s_h, a_h))^2} \right);$ 
15   $\bar{\delta}_h(s_h, a_h) \leftarrow B_h^{\text{next}}(s_h, a_h) - \bar{B}_h(s_h, a_h);$ 
16   $\bar{B}_h(s_h, a_h) \leftarrow B_h^{\text{next}}(s_h, a_h).$ 

```

---

## 4 Analysis

In this section, we outline the main steps needed to establish the main results in Theorem 1 and Theorem 2. Before proceeding, let us first recall the following rescaled learning rates

$$\eta_n = \frac{H+1}{H+n} \quad (24)$$

for the  $n$ -th visit of a given state-action pair at a given time step  $h$ , which are adopted in both LCB-Q and LCB-Q-Advantage. For notational convenience, we further introduce two sequences of related quantities defined for any integers  $N \geq 0$  and  $n \geq 1$ :

$$\eta_0^N := \begin{cases} \prod_{i=1}^N (1 - \eta_i) = 0, & \text{if } N > 0, \\ 1, & \text{if } N = 0, \end{cases} \quad \text{and} \quad \eta_n^N := \begin{cases} \eta_n \prod_{i=n+1}^N (1 - \eta_i), & \text{if } N > n, \\ \eta_n, & \text{if } N = n, \\ 0, & \text{if } N < n. \end{cases} \quad (25)$$

The following identity can be easily verified:

$$\sum_{n=0}^N \eta_n^N = 1. \quad (26)$$

### 4.1 Analysis of LCB-Q

To begin with, we intend to derive a recursive formula concerning the update rule of  $Q_h^k$  — the estimate of the Q-function at step  $h$  at the beginning of the  $k$ -th episode. Note that we have omitted the dependency of all quantities on the episode index  $k$  in Algorithm 1. For notational convenience and clearness, we rewrite Algorithm 1 as Algorithm 4 by specifying the dependency on the episode index  $k$  and shall often use the following set of short-hand notation when it is clear from context.

---

**Algorithm 4:** LCB-Q for offline RL (a rewrite of Algorithm 1 to specify dependency on  $k$ )

---

```

1 Parameters: some constant  $c_b > 0$ , target success probability  $1 - \delta \in (0, 1)$ , and  $\iota = \log(\frac{SAT}{\delta})$ .
2 Initialize  $Q_h^1(s, a) \leftarrow 0$ ;  $N_h^1(s, a) \leftarrow 0$  for all  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ ;  $V_h^1(s) \leftarrow 0$  for all
    $(s, h) \in \mathcal{S} \times [H + 1]$ ;  $\pi^1$  s.t.  $\pi_h^1(s) = 1$  for all  $(s, h) \in \mathcal{S} \times [H]$ .
3 for Episode  $k = 1$  to  $K$  do
4   Sample the  $k$ -th trajectory  $\{s_h^k, a_h^k, r_h^k\}_{h=1}^H$  from  $\mathcal{D}_\mu$ . // sampling from batch dataset
5   for Step  $h = 1$  to  $H$  do
6     for  $(s, a) \in \mathcal{S} \times \mathcal{A}$  do
7       // carry over the estimates and policy
7        $N_h^{k+1}(s, a) \leftarrow N_h^k(s, a)$ ;  $Q_h^{k+1}(s, a) \leftarrow Q_h^k(s, a)$ ;  $V_h^{k+1}(s) \leftarrow V_h^k(s)$ ;  $\pi_h^{k+1}(s) \leftarrow \pi_h^k(s)$ .
8        $N_h^{k+1}(s_h^k, a_h^k) \leftarrow N_h^k(s_h^k, a_h^k) + 1$ . // update the counter
9        $n \leftarrow N_h^{k+1}(s_h^k, a_h^k)$ ;  $\eta_n \leftarrow \frac{H+1}{H+n}$ . // update the learning rate
10       $b_n \leftarrow c_b \sqrt{\frac{H^3 \iota^2}{n}}$ . // update the bonus term
10      // update the Q-estimates with LCB
11       $Q_h^{k+1}(s_h^k, a_h^k) \leftarrow Q_h^k(s_h^k, a_h^k) + \eta_n \{r_h(s_h^k, a_h^k) + V_{h+1}^k(s_{h+1}^k) - Q_h^k(s_h^k, a_h^k) - b_n\}$ .
11      // update the value estimates
12       $V_h^{k+1}(s_h^k) \leftarrow \max \{V_h^k(s_h^k), \max_a Q_h^{k+1}(s_h^k, a)\}$ .
12      // update the policy
13      If  $V_h^{k+1}(s_h^k) = \max_a Q_h^{k+1}(s_h^k, a)$ : update  $\pi_h^{k+1}(s_h^k) = \arg \max_a Q_h^{k+1}(s_h^k, a)$ .

```

---

- $N_h^k(s, a)$ , or the shorthand  $N_h^k$ : the number of episodes that has visited  $(s, a)$  at step  $h$  before the beginning of the  $k$ -th episode.
- $k_h^n(s, a)$ , or the shorthand  $k_h^n$ : the index of the episode in which the state-action pair  $(s, a)$  is visited at step  $h$  for the  $n$ -th times. We also adopt the convention that  $k^0 = 0$ .
- $P_h^k \in \{0, 1\}^{1 \times S}$ : a row vector corresponding to the empirical transition at step  $h$  of the  $k$ -th episode, namely,

$$P_h^k(s) = \mathbb{1}(s = s_{h+1}^k) \quad \text{for all } s \in \mathcal{S}. \quad (27)$$

- $\pi^k = \{\pi_h^k\}_{h=1}^H$  with  $\pi_h^k(s) := \arg \max_a Q_h^k(s, a), \forall (h, s) \in [H] \times \mathcal{S}$ : the deterministic greedy policy at the beginning of the  $k$ -th episode.
- $\hat{\pi}$ : the final output  $\hat{\pi}$  of Algorithms 1 corresponds to  $\pi^{K+1}$  defined above; for notational simplicity, we shall treat  $\hat{\pi}$  as  $\pi^K$  in our analysis, which does not affect our result at all.

Consider any state-action pair  $(s, a)$ . According to the update rule in line 11 of Algorithm 4, we can express (with the assistance of the above notation)

$$Q_h^k(s, a) = Q_h^{k^{N_h^k}+1}(s, a) = (1 - \eta_{N_h^k}) Q_h^{k^{N_h^k}}(s, a) + \eta_{N_h^k} \{r_h(s, a) + V_{h+1}^{k^{N_h^k}}(s_{h+1}^{k^{N_h^k}}) - b_{N_h^k}\}, \quad (28)$$

where the first identity holds since  $k^{N_h^k}$  denotes the latest episode prior to  $k$  that visits  $(s, a)$  at step  $h$ , and the learning rate is defined in (24). Note that it always holds that  $k > k^{N_h^k}$ . Applying the above relation (28) recursively and using the notation (25) lead to

$$Q_h^k(s, a) = \eta_0^{N_h^k} Q_h^1(s, a) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( r_h(s, a) + V_{h+1}^{k^n}(s_{h+1}^{k^n}) - b_n \right). \quad (29)$$

As another important fact, the value estimate  $V_h^k$  is monotonically non-decreasing in  $k$ , i.e.,

$$V_h^{k+1}(s) \geq V_h^k(s) \quad \text{for all } (s, k, h) \in \mathcal{S} \times [K] \times [H], \quad (30)$$

which is an immediate consequence of the update rule in line 12 of Algorithm 4. Crucially, we observe that the iterate  $V_h^k$  forms a “pessimistic view” of  $V_h^{\pi^k}$  — and in turn  $V_h^*$  — resulting from suitable design of the penalty term. This observation is formally stated in the following lemma, with the proof postponed to Section B.1.

**Lemma 1.** *Consider any  $\delta \in (0, 1)$ , and suppose that  $c_b > 0$  is some sufficiently large constant. Then with probability at least  $1 - \delta$ ,*

$$\left| \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} \left( P_{h,s,a} - P_h^{k^n(s,a)} \right) V_{h+1}^{k^n(s,a)} \right| \leq \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n \quad (31)$$

holds simultaneously for all  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , and

$$V_h^k(s) \leq V_h^{\pi^k}(s) \leq V_h^*(s) \quad (32)$$

holds simultaneously for all  $(k, h, s) \in [K] \times [H] \times \mathcal{S}$ .

In a nutshell, the result (32) in Lemma 1 reveals that  $V_h^k$  is a pointwise lower bound on  $V_h^{\pi^k}$  and  $V_h^*$ , thereby forming a pessimistic estimate of the optimal value function. In addition, the property (31) in Lemma 1 essentially tells us that the weighted sum of the penalty terms dominates the weighted sum of the uncertainty terms, which plays a crucial role in ensuring the aforementioned pessimism property. As we shall see momentarily, Lemma 1 forms the basis of the subsequent proof.

We are now ready to embark on the analysis for LCB-Q, which is divided into multiple steps as follows.

**Step 1: decomposing estimation errors.** With the aid of Lemma 1, we can develop an upper bound on the performance difference of interest in (20) as follows

$$\begin{aligned} V_1^*(\rho) - V_1^{\hat{\pi}}(\rho) &= \mathbb{E}_{s_1 \sim \rho} [V_1^*(s_1)] - \mathbb{E}_{s_1 \sim \rho} [V_1^{\pi^K}(s_1)] \\ &\stackrel{(i)}{\leq} \mathbb{E}_{s_1 \sim \rho} [V_1^*(s_1)] - \mathbb{E}_{s_1 \sim \rho} [V_1^K(s_1)] \\ &\stackrel{(ii)}{\leq} \frac{1}{K} \sum_{k=1}^K \left( \mathbb{E}_{s_1 \sim \rho} [V_1^*(s_1)] - \mathbb{E}_{s_1 \sim \rho} [V_1^k(s_1)] \right) \\ &= \frac{1}{K} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_1^{\pi^*}(s) (V_1^*(s) - V_1^k(s)), \end{aligned} \quad (33)$$

where (i) results from Lemma 1 (i.e.,  $V_1^{\pi^K}(s) \geq V_1^K(s)$  for all  $s \in \mathcal{S}$ ), (ii) follows from the monotonicity property in (30), and the last equality holds since  $d_1^{\pi^*}(s) = \rho(s)$  (cf. (4)).

We then attempt to bound the quantity on the right-hand side of (33). Given that  $\pi^*$  is assumed to be a deterministic policy, we have  $d_h^{\pi^*}(s) = d_h^{\pi^*}(s, \pi^*(s))$ . Taking this together with the relations  $V_h^k(s) \geq \max_a Q_h^k(s, a) \geq Q_h^k(s, \pi_h^*(s))$  (see line 12 of Algorithm 4) and  $V_h^*(s) = Q_h^*(s, \pi_h^*(s))$ , we obtain

$$\begin{aligned} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) &= \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s, \pi_h^*(s)) (V_h^*(s) - V_h^k(s)) \\ &\leq \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s, \pi_h^*(s)) (Q_h^*(s, \pi_h^*(s)) - Q_h^k(s, \pi_h^*(s))) \\ &= \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) (Q_h^*(s, a) - Q_h^k(s, a)) \end{aligned} \quad (34)$$

for any  $h \in [H]$ , where the last identity holds since  $\pi^*$  is deterministic and hence

$$d_h^{\pi^*}(s, a) = 0 \quad \text{for any } a \neq \pi_h^*(s). \quad (35)$$

In view of (34), we need to properly control  $Q_h^*(s, a) - Q_h^k(s, a)$ . By virtue of (26), we can rewrite  $Q_h^*(s, a)$  as follows

$$\begin{aligned} Q_h^*(s, a) &= \sum_{n=0}^{N_h^k} \eta_n^{N_h^k} Q_h^*(s, a) = \eta_0^{N_h^k} Q_h^*(s, a) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} Q_h^*(s, a) \\ &= \eta_0^{N_h^k} Q_h^*(s, a) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (r_h(s, a) + P_{h,s,a} V_{h+1}^*), \end{aligned} \quad (36)$$

where the second line follows from Bellman's optimality equation (10). Combining (29) and (36) leads to

$$\begin{aligned} &Q_h^*(s, a) - Q_h^k(s, a) \\ &= \eta_0^{N_h^k} (Q_h^*(s, a) - Q_h^1(s, a)) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_{h,s,a} V_{h+1}^* - V_{h+1}^{k^n}(s_{h+1}^{k^n}) + b_n) \\ &= \eta_0^{N_h^k} (Q_h^*(s, a) - Q_h^1(s, a)) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} b_n + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n}) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_{h,s,a} - P_h^{k^n}) V_{h+1}^{k^n} \end{aligned} \quad (37)$$

$$\leq \eta_0^{N_h^k} H + 2 \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} b_n + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n}), \quad (38)$$

where we have made use of the definition in (27) by recognizing  $P_h^{k^n} V_{h+1}^{k^n} = V_{h+1}^{k^n}(s_{h+1}^{k^n})$  in (37), and the last inequality follows from the fact  $Q_h^*(s, a) - Q_h^1(s, a) = Q_h^*(s, a) - 0 \leq H$  and the bound (31) in Lemma 1. Substituting the above bound into (34), we arrive at

$$\begin{aligned} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) &\leq \underbrace{\sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \eta_0^{N_h^k(s,a)} H + 2 \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n}_{=: I_h} \\ &\quad + \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (V_{h+1}^* - V_{h+1}^{k^n}(s_{h+1}^{k^n})). \end{aligned} \quad (39)$$

**Step 2: establishing a crucial recursion.** As it turns out, the last term on the right-hand side of (39) can be used to derive a recursive relation that connects step  $h$  with step  $h+1$ , as summarized in the next lemma.

**Lemma 2.** *With probability at least  $1 - \delta$ , the following recursion holds:*

$$\begin{aligned} &\sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (V_{h+1}^* - V_{h+1}^{k^n}(s_{h+1}^{k^n})) \\ &\leq \left(1 + \frac{1}{H}\right) \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)) + 24 \sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 12 H C^* \log \frac{2H}{\delta}. \end{aligned} \quad (40)$$

Lemma 2 taken together with (39) implies that

$$\sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) \leq \left(1 + \frac{1}{H}\right) \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s))$$



$$+ I_h + 24\sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 12HC^* \log \frac{2H}{\delta}. \quad (41)$$

Invoking (41) recursively over the time steps  $h = H, H-1, \dots, 1$  with the terminal condition  $V_{H+1}^k = V_{H+1}^* = 0$ , we reach

$$\begin{aligned} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_1^{\pi^*}(s) (V_1^*(s) - V_1^k(s)) &\leq \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) \\ &\leq \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \left( I_h + 24\sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 12HC^* \log \frac{2H}{\delta} \right), \end{aligned} \quad (42)$$

which captures the estimation error resulting from the use of pessimism principle.

**Step 3: controlling the right-hand side of (42).** The right-hand side of (42) can be bounded through the following lemma, which will be proved in Appendix B.3.

**Lemma 3.** *Consider any  $\delta \in (0, 1)$ . With probability at least  $1 - \delta$ , we have*

$$\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \left( I_h + 24\sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 12HC^* \log \frac{2H}{\delta} \right) \lesssim H^2 SC^* \iota + \sqrt{H^5 SC^* K \iota^3}, \quad (43)$$

where we recall that  $\iota := \log \left( \frac{SAT}{\delta} \right)$ .

Combining Lemma 3 with (42) and (33) yields

$$\begin{aligned} V_1^*(\rho) - V_1^{\widehat{\pi}}(\rho) &\leq \frac{1}{K} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_1^{\pi^*}(s) (V_1^*(s) - V_1^k(s)) \\ &\leq \frac{1}{K} \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) \\ &\leq \frac{c_a}{2} \sqrt{\frac{H^5 SC^* \iota^3}{K}} + \frac{c_a}{2} \frac{H^2 SC^* \iota}{K} = \frac{c_a}{2} \sqrt{\frac{H^6 SC^* \iota^3}{T}} + \frac{c_a}{2} \frac{H^3 SC^* \iota}{T} \\ &\leq c_a \sqrt{\frac{H^6 SC^* \iota^3}{T}} \end{aligned} \quad (44)$$

for some sufficiently large constant  $c_a > 0$ , where the last inequality is valid as long as  $T > SC^* \iota$ . This concludes the proof of Theorem 1.

## 4.2 Analysis of LCB-Q-Advantage

We now turn to the analysis of LCB-Q-Advantage. Thus far, we have omitted the dependency of all quantities on the epoch number  $m$  and the in-epoch episode number  $t$  in Algorithms 2 and 3. While it allows for a more concise description of our algorithm, it might hamper the clarity of our proofs. In the following, we introduce the notation  $k$  to denote the current episode as follows:

$$k := \sum_{i=1}^{m-1} L_i + t, \quad (45)$$

which corresponds to the  $t$ -th in-epoch episode in the  $m$ -th epoch; here,  $L_m = 2^m$  stands for the total number of in-epoch episodes in the  $m$ -th epoch. With this notation in place, we can rewrite Algorithm 2 as Algorithm 5 in order to make clear the dependency on the episode index  $k$ , epoch number  $m$ , and in-epoch episode index  $t$ .

---

**Algorithm 5:** LCB-Q-Advantage (a rewrite of Algorithm 2 that specifies dependency on  $k$  or  $(m, t)$ .)

---

```

1 Parameters: number of epochs  $M$ , universal constant  $c_b > 0$ , target success probability
    $1 - \delta \in (0, 1)$ , and  $\iota = \log(\frac{SAT}{\delta})$ .
2 Initialize:
3  $Q_h^1(s, a), Q_h^{\text{LCB},1}(s, a), \bar{Q}_h^1(s, a), \bar{\mu}_h^1(s, a), \bar{\mu}_h^{\text{next},1}(s, a), N_h^1(s, a) \leftarrow 0$  for all  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ ;
4  $V_h^1(s), \bar{V}_h^1(s), \bar{V}_h^{\text{next},1}(s) \leftarrow 0$  for all  $(s, h) \in \mathcal{S} \times [H + 1]$ ;
5  $\mu_h^{\text{ref},1}(s, a), \sigma_h^{\text{ref},1}(s, a), \mu_h^{\text{adv},1}(s, a), \sigma_h^{\text{adv},1}(s, a), \bar{\delta}_h^1(s, a), \bar{B}_h^1(s, a) \leftarrow 0$  for all  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ .
6 for Epoch  $m = 1$  to  $M$  do
7    $L_m = 2^m$ . // specify the number of episodes in the current epoch
8    $\hat{N}_h^{(m,1)}(s, a) = 0$  for all  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . // reset the epoch-wise counter
   /* Inner-loop: update value-estimates  $V_h(s, a)$  and Q-estimates  $Q_h(s, a)$  */
9   for In-epoch Episode  $t = 1$  to  $L_m$  do
10    Set  $k \leftarrow \sum_{i=1}^{m-1} L_i + t$ . // set the episode index
11    Sample the  $k$ -th trajectory  $\{s_h^k, a_h^k, r_h^k\}_{h=1}^H$ . // sampling from batch dataset
12    Compute  $\pi^k$  s.t.  $\pi_h^k(s) = \arg \max_a Q_h^k(s, a)$  for all  $(s, h) \in \mathcal{S} \times [H]$ . // update the policy
13    for Step  $h = 1$  to  $H$  do
14      for  $(s, a) \in \mathcal{S} \times \mathcal{A}$  do
15        // carry over the estimates
16         $N_h^{k+1}(s, a) \leftarrow N_h^k(s, a); \hat{N}_h^{k+1}(s, a) \leftarrow \hat{N}_h^k(s, a); V_h^{k+1}(s) \leftarrow V_h^k(s);$ 
17         $Q_h^{\text{LCB},k+1}(s, a) \leftarrow Q_h^{\text{LCB},k}(s, a); \bar{Q}_h^{k+1}(s, a) \leftarrow \bar{Q}_h^k(s, a); Q_h^{k+1}(s, a) \leftarrow Q_h^k(s, a);$ 
18         $\bar{V}_h^{k+1}(s) \leftarrow \bar{V}_h^k(s); \bar{V}_h^{\text{next},k+1}(s) \leftarrow \bar{V}_h^{\text{next},k}(s); \bar{\mu}_h^{k+1}(s, a) \leftarrow \bar{\mu}_h^k(s, a).$ 
19         $N_h^{k+1}(s_h^k, a_h^k) \leftarrow N_h^k(s_h^k, a_h^k) + 1; n \leftarrow N_h^{k+1}(s_h^k, a_h^k)$ . // update the overall counter
20         $\eta_n \leftarrow \frac{H+1}{H+n}$ . // update the learning rate
21        // update the Q-estimate with LCB
22         $Q_h^{\text{LCB},k+1}(s_h^k, a_h^k) \leftarrow \text{update-lcb-q}()$ .
23        // update the Q-estimate with LCB and reference-advantage
24         $\bar{Q}_h^{k+1}(s_h^k, a_h^k) \leftarrow \text{update-lcb-q-ra}()$ .
25        // update the Q-estimate  $Q_h$  and value estimate  $V_h$ 
26         $Q_h^{k+1}(s_h^k, a_h^k) \leftarrow \max \{Q_h^{\text{LCB},k+1}(s_h^k, a_h^k), \bar{Q}_h^{k+1}(s_h^k, a_h^k), Q_h^k(s_h^k, a_h^k)\}.$ 
27         $V_h^{k+1}(s_h^k) \leftarrow \max_a Q_h^{k+1}(s_h^k, a).$ 
28        // update epoch-wise counter and  $\bar{\mu}_h^{\text{next}}(s, a)$  for the next epoch
29         $\hat{N}_h^{(m,t+1)}(s_h^k, a_h^k) \leftarrow \hat{N}_h^{(m,t)}(s_h^k, a_h^k) + 1.$ 
30         $\bar{\mu}_h^{\text{next},k+1}(s_h^k, a_h^k) \leftarrow \left(1 - \frac{1}{\hat{N}_h^{(m,t+1)}(s_h^k, a_h^k)}\right) \bar{\mu}_h^{\text{next},k}(s_h, a_h) + \frac{1}{\hat{N}_h^{(m,t+1)}(s_h^k, a_h^k)} \bar{V}_h^{\text{next},k}(s_{h+1}).$ 
31      /* Update the reference  $(\bar{V}_h, \bar{V}_h^{\text{next}})$  and  $(\bar{\mu}_h, \bar{\mu}_h^{\text{next}})$  */
32    for  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H + 1]$  do
33       $\bar{V}_h^{k+1}(s) \leftarrow \bar{V}_h^{\text{next},k+1}(s); \bar{\mu}_h^{k+1}(s, a) \leftarrow \bar{\mu}_h^{\text{next},k+1}(s, a)$ . // set  $\bar{V}_h$  and  $\bar{\mu}_h$  for the next epoch
34       $\bar{V}_h^{\text{next},k+1}(s) \leftarrow V_h^{k+1}(s); \bar{\mu}_h^{\text{next},k+1}(s, a) \leftarrow 0$ . // set  $\bar{\mu}_h^{\text{next}}$  and  $\bar{V}_h^{\text{next}}$  for the next epoch

```

---

**Output:** the policy  $\hat{\pi} = \pi^K$  with  $K = \sum_{m=1}^M L_m$ .

---

Before embarking on our main proof, we make two crucial observations which play important roles in our subsequent analysis. First, similar to the property (30) for LCB-Q, the update rule (cf. lines 16-17 of Algorithm 5) ensures the monotonic non-decreasing property of  $V_h(s)$  such that for all  $k \in [K]$ ,

$$V_h^{k+1}(s) \geq V_h^k(s), \quad \text{for all } (k, s, h) \in [K] \times \mathcal{S} \times [H]. \quad (46)$$

Secondly,  $V_h^k$  forms a “pessimistic view” of  $V_h^*$ , which is formalized in the lemma below; the proof is deferred

to Appendix C.1.

**Lemma 4.** *Let  $\delta \in (0, 1)$ . Suppose that  $c_b > 0$  is some sufficiently large constant. Then with probability at least  $1 - \delta$ , the value estimates produced by Algorithm 2 satisfy*

$$V_h^k(s) \leq V_h^{\pi^k}(s) \leq V^*(s) \quad (47)$$

for all  $(k, h, s) \in [K] \times [H + 1] \times \mathcal{S}$ .

With these two observations in place, we can proceed to present the analysis for LCB-Q-Advantage. To begin with, the performance difference of interest can be controlled similar to (33) as follows:

$$\begin{aligned} V_1^*(\rho) - V_1^{\hat{\pi}}(\rho) &= \mathbb{E}_{s_1 \sim \rho} [V_1^*(s_1)] - \mathbb{E}_{s_1 \sim \rho} [V_1^{\pi^K}(s_1)] \\ &\stackrel{(i)}{\leq} \mathbb{E}_{s_1 \sim \rho} [V_1^*(s_1)] - \mathbb{E}_{s_1 \sim \rho} [V_1^K(s_1)] \\ &\stackrel{(ii)}{\leq} \frac{1}{K} \sum_{k=1}^K \left( \mathbb{E}_{s_1 \sim \rho} [V_1^*(s_1)] - \mathbb{E}_{s_1 \sim \rho} [V_1^k(s_1)] \right) \\ &= \frac{1}{K} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_1^{\pi^*}(s) (V_1^*(s) - V_1^k(s)), \end{aligned} \quad (48)$$

where (i) follows from Lemma 4 (i.e.,  $V_1^{\pi^K}(s) \geq V_1^K(s)$  for all  $s \in \mathcal{S}$ ), (ii) holds due to the monotonicity in (46) and the last equality holds since  $d_1^{\pi^*}(s) = \rho(s)$  (cf. (4)). It then boils down to controlling the right-hand side of (48). Towards this end, it turns out that one can control a more general counterpart, i.e.,

$$\sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) \quad (49)$$

for any  $h \in [H]$ . This is accomplished via the following lemma, whose proof is postponed to Appendix C.2.

**Lemma 5.** *Let  $\delta \in (0, 1)$ , and recall that  $\iota := \log\left(\frac{SAT}{\delta}\right)$ . Suppose that  $c_a, c_b > 0$  are some sufficiently large constants. Then with probability at least  $1 - \delta$ , one has*

$$\sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) \leq J_h^1 + J_h^2 + J_h^3, \quad (50)$$

where

$$\begin{aligned} J_h^1 &:= \sum_{k=1}^K \sum_{s, a \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \left[ \eta_0^{N_h^k(s, a)} H + \frac{4c_b H^{7/4} \iota}{(N_h^k(s, a) \vee 1)^{3/4}} + \frac{4c_b H^2 \iota}{N_h^k(s, a) \vee 1} \right], \\ J_h^2 &:= 2 \sum_{k=1}^K \sum_{s, a \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \bar{B}_h^k(s, a), \\ J_h^3 &:= \left(1 + \frac{1}{H}\right) \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)) + 48 \sqrt{HC^* K \log \frac{2H}{\delta}} + 28c_a H^3 C^* \sqrt{S} \iota^2. \end{aligned} \quad (51)$$

As a direct consequence of Lemma 5, one arrives at a recursive relationship between time steps  $h$  and  $h + 1$  as follows:

$$\sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s))$$

$$\leq \left(1 + \frac{1}{H}\right) \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)) + 48\sqrt{HC^*K \log \frac{2H}{\delta}} + 28c_a H^3 C^* \sqrt{S} \iota^2 + J_h^1 + J_h^2. \quad (52)$$

Recurring over time steps  $h = H, H-1, \dots, 1$  with the terminal condition  $V_{H+1}^k = V_{H+1}^* = 0$ , we can upper bound the performance difference at  $h = 1$  as follows

$$\begin{aligned} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_1^{\pi^*}(s) (V_1^*(s) - V_1^k(s)) &\leq \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) \\ &\leq \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \left(48\sqrt{HC^*K \log \frac{2H}{\delta}} + 28c_a H^3 C^* \sqrt{S} \iota^2 + J_h^1 + J_h^2\right). \end{aligned} \quad (53)$$

To finish up, it suffices to upper bound each term in (53) separately. We summarize their respective upper bounds as follows; the proof is provided in Appendix C.3.

**Lemma 6.** Fix  $\delta \in (0, 1)$ , and recall that  $\iota := \log\left(\frac{SAT}{\delta}\right)$ . With probability at least  $1 - \delta$ , we have

$$\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} J_h^1 \lesssim H^{2.75} (SC^*)^{\frac{3}{4}} K^{\frac{1}{4}} \iota^2 + H^3 SC^* \iota^3, \quad (54a)$$

$$\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} J_h^2 \lesssim \sqrt{H^4 SC^* \iota^3 \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s))} + \sqrt{H^3 SC^* K \iota^5} + H^4 SC^* \iota^4, \quad (54b)$$

$$\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \left(48\sqrt{HC^*K \log \frac{2H}{\delta}} + 28c_a H^3 C^* \sqrt{S} \iota^2\right) \lesssim \sqrt{H^3 C^* K \log \frac{2H}{\delta}} + H^4 C^* \sqrt{S} \iota^2. \quad (54c)$$

Substituting the above upper bounds into (48) and (53) and recalling that  $T = HK$ , we arrive at

$$\begin{aligned} V_1^*(\rho) - V_1^{\hat{\pi}}(\rho) &\lesssim \frac{1}{K} \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) \\ &\lesssim \frac{1}{K} \left( \sqrt{H^4 SC^* \iota^3 \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s))} + \left( \sqrt{H^3 SC^* K \iota^5} + H^4 SC^* \iota^4 + H^{2.75} (SC^*)^{\frac{3}{4}} K^{\frac{1}{4}} \iota^2 \right) \right) \\ &\stackrel{(i)}{\asymp} \frac{1}{K} \left( \sqrt{H^4 SC^* \iota^3 \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s))} + \sqrt{H^3 SC^* K \iota^5} + H^4 SC^* \iota^4 \right) \\ &\stackrel{(ii)}{\lesssim} \frac{1}{K} \left( \sqrt{H^3 SC^* K \iota^5} + H^4 SC^* \iota^4 \right) \\ &\asymp \sqrt{\frac{H^4 SC^* \iota^5}{T}} + \frac{H^5 SC^* \iota^4}{T}, \end{aligned}$$

where (i) has made use of the AM-GM inequality:

$$2H^{2.75} (SC^*)^{\frac{3}{4}} K^{\frac{1}{4}} \leq \left( H^{0.75} (SC^*)^{\frac{1}{4}} K^{\frac{1}{4}} \right)^2 + \left( H^2 (SC^*)^{\frac{1}{2}} \right)^2 = \sqrt{H^3 SC^* K} + H^4 SC^*,$$

and (ii) holds by letting  $x := \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s))$  and solving the inequality  $x \lesssim \sqrt{H^4 SC^* \iota^3 x} + \sqrt{H^3 SC^* K \iota^5} + H^4 SC^* \iota^4$ . This concludes the proof.

## 5 Discussions

Focusing on model-free paradigms, this paper has developed near-optimal sample complexities for some variants of pessimistic Q-learning algorithms — armed with lower confidence bounds and variance reduction — for offline RL. These sample complexity results, taken together with the analysis framework developed herein, open up a few exciting directions for future research. For example, the pessimistic Q-learning algorithms can be deployed in conjunction with their optimistic counterparts (e.g., Jin et al. (2018); Li et al. (2021b); Zhang et al. (2020b)), when additional online data can be acquired to fine-tune the policy (Xie et al., 2021b). In addition, the  $\varepsilon$ -range for LCB-Q-Advantage to attain sample optimality remains somewhat limited (i.e.,  $\varepsilon \in (0, 1/H]$ ). Our ongoing work Li et al. (2022) suggests that a new variant of pessimistic model-based algorithm is sample-optimal for a broader range of  $\varepsilon$ , which in turn motivates further investigation into whether model-free algorithms can accommodate a broader  $\varepsilon$ -range too without compromising sample efficiency. Moving beyond the tabular setting, it would be of great importance to extend the algorithmic and theoretical framework to accommodate low-complexity function approximation (Nguyen-Tang et al., 2021).

## Acknowledgements

L. Shi and Y. Chi are supported in part by the grants ONR N00014-19-1-2404, NSF CCF-2106778, CCF-2007911, and CAREER award ECCS-1818571. Y. Wei is supported in part by the grant NSF DMS-2147546/2015447 and the NSF CAREER award DMS-2143215. Y. Chen is supported in part by the Alfred P. Sloan Research Fellowship, the grants AFOSR YIP award FA9550-19-1-0030, ONR N00014-19-1-2120, ARO YIP award W911NF-20-1-0097, NSF CCF-2106739/2221009, CCF-1907661, IIS-1900140/2218713 and IIS-2100158/2218773. L. Shi is also gratefully supported by the Leo Finzi Memorial Fellowship, Wei Shen and Xuehong Zhang Presidential Fellowship, and Liang Ji-Dian Graduate Fellowship at Carnegie Mellon University. Part of this work was done while G. Li, Y. Wei and Y. Chen were visiting the Simons Institute for the Theory of Computing.

## A Technical lemmas

### A.1 Preliminary facts

Our results rely heavily on proper choices of the learning rates. In what follows, we make note of several useful properties concerning the learning rates, which have been established in Jin et al. (2018); Li et al. (2021b).

**Lemma 7** (Lemma 1 in (Li et al., 2021b)). *For any integer  $N > 0$ , the following properties hold:*

$$\frac{1}{N^a} \leq \sum_{n=1}^N \frac{\eta_n^N}{n^a} \leq \frac{2}{N^a} \quad \text{for all } \frac{1}{2} \leq a \leq 1, \quad (55a)$$

$$\max_{1 \leq n \leq N} \eta_n^N \leq \frac{2H}{N}, \quad \sum_{n=1}^N (\eta_n^N)^2 \leq \frac{2H}{N}, \quad \sum_{N=n}^{\infty} \eta_n^N \leq 1 + \frac{1}{H}. \quad (55b)$$

In addition, we gather a few elementary properties about the Binomial distribution, which will be useful throughout the proof. The lemma below is adapted from Xie et al. (2021b, Lemma A.1).

**Lemma 8.** *Suppose  $N \sim \text{Binomial}(n, p)$ , where  $n \geq 1$  and  $p \in [0, 1]$ . For any  $\delta \in (0, 1)$ , we have*

$$\frac{p}{N \vee 1} \leq \frac{8 \log(\frac{1}{\delta})}{n}, \quad (56)$$

and

$$N \geq \frac{np}{8 \log(\frac{1}{\delta})} \quad \text{if } np \geq 8 \log\left(\frac{1}{\delta}\right), \quad (57a)$$

$$N \leq \begin{cases} e^2 np & \text{if } np \geq \log\left(\frac{1}{\delta}\right), \\ 2e^2 \log\left(\frac{1}{\delta}\right) & \text{if } np \leq 2 \log\left(\frac{1}{\delta}\right). \end{cases} \quad (57b)$$

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

*Proof.* To begin with, we directly invoke [Xie et al. \(2021b\)](#), Lemma A.1) which yields the results in (56) and (57a). Regarding (57b), invoking the Chernoff bound ([Vershynin, 2018](#), Theorem 2.3.1) with  $\mathbb{E}[N] = np$ , when  $np \geq \log\left(\frac{1}{\delta}\right)$ , it satisfies

$$\mathbb{P}(N \geq e^2 np) \leq e^{-np} \left( \frac{enp}{e^2 np} \right)^{e^2 np} \leq e^{-np} \leq \delta.$$

Similarly, when  $np \leq 2 \log\left(\frac{1}{\delta}\right)$ , we have

$$\begin{aligned} \mathbb{P}\left(N \geq 2e^2 \log\left(\frac{1}{\delta}\right)\right) &\stackrel{(i)}{\leq} e^{-np} \left( \frac{enp}{2e^2 \log\left(\frac{1}{\delta}\right)} \right)^{2e^2 \log\left(\frac{1}{\delta}\right)} \\ &\stackrel{(ii)}{\leq} e^{-np} \left( \frac{enp}{e^2 np} \right)^{2e^2 \log\left(\frac{1}{\delta}\right)} \leq e^{-2e^2 \log\left(\frac{1}{\delta}\right)} \leq \delta, \end{aligned}$$

where (i) results from [Vershynin \(2018, Theorem 2.3.1\)](#), and (ii) follows from the basic fact  $e^2 \log\left(\frac{1}{\delta}\right) \geq 2 \log\left(\frac{1}{\delta}\right) \geq np$ . Taking the union bound thus completes the proof.  $\square$

## A.2 Freedman's inequality and its consequences

Both the samples collected within each episode and the algorithms analyzed herein exhibit certain Markovian structure. As a result, concentration inequalities tailored to martingales become particularly effective for our analysis. In this subsection, we collect a few useful concentration results that will be applied multiple times in the current paper. These results might be of independent interest.

To begin with, the following theorem provides a user-friendly version of Freedman's inequality ([Freedman, 1975](#)); see [Li et al. \(2021a, Section C\)](#) for more details.

**Theorem 3** (Freedman's inequality). *Consider a filtration  $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots$ , and let  $\mathbb{E}_k$  stand for the expectation conditioned on  $\mathcal{F}_k$ . Suppose that  $Y_n = \sum_{k=1}^n X_k \in \mathbb{R}$ , where  $\{X_k\}$  is a real-valued scalar sequence obeying*

$$|X_k| \leq R \quad \text{and} \quad \mathbb{E}_{k-1}[X_k] = 0 \quad \text{for all } k \geq 1$$

for some quantity  $R < \infty$ . We also define

$$W_n := \sum_{k=1}^n \mathbb{E}_{k-1}[X_k^2].$$

In addition, suppose that  $W_n \leq \sigma^2$  holds deterministically for some given quantity  $\sigma^2 < \infty$ . Then for any positive integer  $m \geq 1$ , with probability at least  $1 - \delta$  one has

$$|Y_n| \leq \sqrt{8 \max\left\{W_n, \frac{\sigma^2}{2m}\right\} \log \frac{2m}{\delta}} + \frac{4}{3} R \log \frac{2m}{\delta}. \quad (58)$$

We shall also record some immediate consequence of Freedman's inequality tailored to our problem. Recall that  $N_h^i(s, a)$  denotes the number of times that  $(s, a)$  has been visited at step  $h$  before the beginning of the  $i$ -th episode, and  $k^n(s, a)$  stands for the index of the episode in which  $(s, a)$  is visited for the  $n$ -th time. The following concentration bound has been established in [Li et al. \(2021b, Lemma 7\)](#).

**Lemma 9.** *Let  $\{W_h^i \in \mathbb{R}^S \mid 1 \leq i \leq K, 1 \leq h \leq H+1\}$  and  $\{u_h^i(s, a, N) \in \mathbb{R} \mid 1 \leq i \leq K, 1 \leq h \leq H+1\}$  be a collections of vectors and scalars, respectively, and suppose that they obey the following properties:*

- $W_h^i$  is fully determined by the samples collected up to the end of the  $(h-1)$ -th step of the  $i$ -th episode;
- $\|W_h^i\|_\infty \leq C_w$ ;
- $u_h^i(s, a, N)$  is fully determined by the samples collected up to the end of the  $(h-1)$ -th step of the  $i$ -th episode, and a given positive integer  $N \in [K]$ ;
- $0 \leq u_h^i(s, a, N) \leq C_u$ ;
- $0 \leq \sum_{n=1}^{N_h^k(s,a)} u_h^{k_n^n(s,a)}(s, a, N) \leq 2$ .

In addition, consider the following sequence

$$X_i(s, a, h, N) := u_h^i(s, a, N)(P_h^i - P_{h,s,a})W_{h+1}^i \mathbb{1}\{(s_h^i, a_h^i) = (s, a)\}, \quad 1 \leq i \leq K, \quad (59)$$

with  $P_h^i$  defined in (27). Consider any  $\delta \in (0, 1)$ . Then with probability at least  $1 - \delta$ ,

$$\begin{aligned} & \left| \sum_{i=1}^k X_i(s, a, h, N) \right| \\ & \lesssim \sqrt{C_u \log^2 \frac{SAT}{\delta}} \sqrt{\sum_{n=1}^{N_h^k(s,a)} u_h^{k_n^n(s,a)}(s, a, N) \text{Var}_{h,s,a}(W_{h+1}^{k_n^n(s,a)})} + \left( C_u C_w + \sqrt{\frac{C_u}{N}} C_w \right) \log^2 \frac{SAT}{\delta} \end{aligned} \quad (60)$$

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

Next, we make note of an immediate consequence of Lemma 9 as follows.

**Lemma 10.** Let  $\{W_h^i \in \mathbb{R}^S \mid 1 \leq i \leq K, 1 \leq h \leq H+1\}$  be a collection of vectors satisfying the following properties:

- $W_h^i$  is fully determined by the samples collected up to the end of the  $(h-1)$ -th step of the  $i$ -th episode;
- $\|W_h^i\|_\infty \leq C_w$ .

For any positive  $N \geq H$ , we consider the following sequence

$$X_i(s, a, h, N) := \eta_{N_h^i(s,a)}^N (P_h^i - P_{h,s,a}) W_{h+1}^i \mathbb{1}\{(s_h^i, a_h^i) = (s, a)\}, \quad 1 \leq i \leq K, \quad (61)$$

with  $P_h^i$  defined in (27). Consider any  $\delta \in (0, 1)$ . With probability at least  $1 - \delta$ ,

$$\left| \sum_{i=1}^k X_i(s, a, h, N) \right| \lesssim \sqrt{\frac{H}{N} C_w^2 \log^2 \frac{SAT}{\delta}} \quad (62)$$

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

*Proof.* Taking  $u_h^i(s, a, N) = \eta_{N_h^i(s,a)}^N$ , one can see from (55b) in Lemma 7 that

$$|u_h^i(s, a, N)| \leq \frac{2H}{N} =: C_u.$$

Recognizing the trivial bound  $\text{Var}_{h,s,a}(W_{h+1}^{k_n^n(s,a)}) \leq C_w^2$ , we can invoke Lemma 9 to obtain that, with probability at least  $1 - \delta$ ,

$$\begin{aligned} \left| \sum_{i=1}^k X_i(s, a, h, N) \right| & \lesssim \sqrt{C_u \log^2 \frac{SAT}{\delta}} \sqrt{\sum_{n=1}^{N_h^k(s,a)} \eta_n^N C_w^2} + \left( C_u C_w + \sqrt{\frac{C_u}{N}} C_w \right) \log^2 \frac{SAT}{\delta} \\ & \lesssim \sqrt{\frac{H}{N} \log^2 \frac{SAT}{\delta}} \cdot C_w + \frac{H C_w}{N} \log^2 \frac{SAT}{\delta} \lesssim \sqrt{\frac{H C_w^2}{N} \log^2 \frac{SAT}{\delta}} \end{aligned}$$

holds simultaneously for all  $(k, h, s, a, N) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A} \times [K]$ , where the last line applies (55b) in Lemma 7 once again.  $\square$



Finally, we introduce another lemma by invoking Freedman's inequality in Theorem 3.

**Lemma 11.** *Let  $\{W_h^k(s, a) \in \mathbb{R}^S \mid (s, a) \in \mathcal{S} \times \mathcal{A}, 1 \leq k \leq K, 1 \leq h \leq H+1\}$  be a collection of vectors satisfying the following properties:*

- $W_h^k(s, a)$  is fully determined by the given state-action pair  $(s, a)$  and the samples collected up to the end of the  $(k-1)$ -th episode;
- $\|W_h^k(s, a)\|_\infty \leq C_w$ .

For any positive  $C_d \geq 0$ , we consider the following sequences

$$X_{h,k} := C_d \left[ \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k,a_h^k} W_{h+1}^k(s_h^k, a_h^k) - \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} W_{h+1}^k(s, a) \right], \quad 1 \leq k \leq K, \quad (63)$$

$$\bar{X}_{h,k} := C_d \left[ \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_h^k W_{h+1}^k(s_h^k, a_h^k) - \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} W_{h+1}^k(s, a) \right], \quad 1 \leq k \leq K. \quad (64)$$

Consider any  $\delta \in (0, 1)$ . Then with probability at least  $1 - \delta$ ,

$$\left| \sum_{k=1}^K X_{h,k} \right| \leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) [P_{h,s,a} W_{h+1}^k(s, a)]^2 \log \frac{2H}{\delta}} + 2C_d C^* C_w \log \frac{2H}{\delta} \quad (65)$$

$$\left| \sum_{k=1}^K \bar{X}_{h,k} \right| \leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} [W_{h+1}^k(s, a)]^2 \log \frac{2H}{\delta}} + 2C_d C^* C_w \log \frac{2H}{\delta} \quad (66)$$

hold simultaneously for all  $h \in [H]$ .

*Proof.* We intend to apply Freedman's inequality (cf. Theorem 3) to control  $\sum_{k=1}^K X_{h,k}$ . Considering any given time step  $h$ , it is easily verified that

$$\mathbb{E}_{k-1}[X_{h,k}] = 0, \quad \mathbb{E}_{k-1}[\bar{X}_{h,k}] = 0,$$

where  $\mathbb{E}_{k-1}$  denotes the expectation conditioned on everything happening up to the end of the  $(k-1)$ -th episode. To continue, we observe that

$$|X_{h,k}| \leq C_d \left( \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} + 1 \right) \|W_{h+1}^k(s, a)\|_\infty \leq 2C_d C^* C_w, \quad (67)$$

$$|\bar{X}_{h,k}| \leq C_d \left( \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} + 1 \right) \|W_{h+1}^k(s, a)\|_\infty \leq 2C_d C^* C_w, \quad (68)$$

where we use the assumptions  $\frac{d_h^{\pi^*}(s, a)}{d_h^\mu(s, a)} \leq C^*$  for all  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$  (cf. Assumption 1) and  $\|W_{h+1}^k(s_h^k, a_h^k)\|_\infty \leq C_w$ .

Recall that  $\Delta(\mathcal{S} \times \mathcal{A})$  is the probability simplex over the set  $\mathcal{S} \times \mathcal{A}$  of all state-action pairs, and we denote by  $d_h^\mu \in \Delta(\mathcal{S} \times \mathcal{A})$  the state-action visitation distribution induced by the behavior policy  $\mu$  at time step  $h \in [H]$ . With this in hand, we obtain

$$\begin{aligned} \sum_{k=1}^K \mathbb{E}_{k-1}[|X_{h,k}|^2] &\leq \sum_{k=1}^K C_d^2 \mathbb{E}_{k-1} \left[ \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k,a_h^k} W_{h+1}^k(s_h^k, a_h^k) - \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} W_{h+1}^k(s, a) \right]^2 \\ &\leq \sum_{k=1}^K C_d^2 \mathbb{E}_{(s_h^k, a_h^k) \sim d_h^\mu} \left[ \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k,a_h^k} W_{h+1}^k(s_h^k, a_h^k) \right]^2 \end{aligned}$$

$$\begin{aligned}
&= \sum_{k=1}^K C_d^2 \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{d_h^{\pi^*}(s,a)}{d_h^\mu(s,a)} d_h^{\pi^*}(s,a) [P_{h,s,a} W_{h+1}^k(s,a)]^2 \\
&\stackrel{(i)}{\leq} \sum_{k=1}^K C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) [P_{h,s,a} W_{h+1}^k(s,a)]^2 \tag{69}
\end{aligned}$$

$$\leq \sum_{k=1}^K C_d^2 \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} C^* d_h^{\pi^*}(s,a) \|W_{h+1}^k(s_h^k, a_h^k)\|_\infty^2 \leq C_d^2 C^* C_w^2 K, \tag{70}$$

where (i) follows from  $\frac{d_h^{\pi^*}(s,a)}{d_h^\mu(s,a)} \leq C^*$  (see Assumption 1) and the assumption  $\|W_{h+1}^k(s_h^k, a_h^k)\|_\infty \leq C_w$ . Similarly, we can derive

$$\begin{aligned}
\sum_{k=1}^K \mathbb{E}_{k-1}[|\bar{X}_{h,k}|^2] &\leq \sum_{k=1}^K C_d^2 \mathbb{E}_{k-1} \left[ \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_h^k W_{h+1}^k(s_h^k, a_h^k) - \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} W_{h+1}^k(s,a) \right]^2 \\
&\leq \sum_{k=1}^K C_d^2 \mathbb{E}_{(s_h^k, a_h^k) \sim d_h^\mu} \left[ \mathbb{E}_{P_h^k \sim P_{h,s_h^k, a_h^k}} \left[ \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_h^k W_{h+1}^k(s_h^k, a_h^k) \right]^2 \right] \\
&= \sum_{k=1}^K C_d^2 \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{d_h^{\pi^*}(s,a)}{d_h^\mu(s,a)} d_h^{\pi^*}(s,a) \mathbb{E}_{P_h^k \sim P_{h,s,a}} [P_h^k W_{h+1}^k(s,a)]^2 \\
&\stackrel{(i)}{\leq} \sum_{k=1}^K C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \mathbb{E}_{P_h^k \sim P_{h,s,a}} [P_h^k W_{h+1}^k(s,a)]^2 \tag{71}
\end{aligned}$$

$$= \sum_{k=1}^K C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} [W_{h+1}^k(s,a)]^2 \tag{72}$$

$$\leq \sum_{k=1}^K C_d^2 \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} C^* d_h^{\pi^*}(s,a) \|W_{h+1}^k(s,a)\|_\infty^2 \leq C_d^2 C^* C_w^2 K, \tag{73}$$

where (i) follows from  $\frac{d_h^{\pi^*}(s,a)}{d_h^\mu(s,a)} \leq C^*$  (see Assumption 1) and the assumption  $\|W_{h+1}^k(s_h^k, a_h^k)\|_\infty \leq C_w$ .

Plugging in the results in (67) and (69) (resps. (68) and (72)) to control  $\sum_{k=1}^K |X_{h,k}|$  (resps.  $\sum_{k=1}^K |\bar{X}_{h,k}|$ ), we invoke Theorem 3 with  $m = \lceil \log_2 K \rceil$  and take the union bound over  $h \in [H]$  to show that with probability at least  $1 - \delta$ ,

$$\begin{aligned}
\left| \sum_{k=1}^K X_{h,k} \right| &\leq \sqrt{8 \max \left\{ \sum_{k=1}^K C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) [P_{h,s,a} W_{h+1}^k(s,a)]^2, \frac{C_d^2 C^* C_w^2 K}{2^m} \right\} \log \frac{2H}{\delta}} \\
&\quad + \frac{8}{3} C_d C^* C_w \log \frac{2H}{\delta} \\
&\leq \sqrt{\sum_{k=1}^K 8 C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) [P_{h,s,a} W_{h+1}^k(s,a)]^2 \log \frac{2H}{\delta} + 6 C_d C^* C_w \log \frac{2H}{\delta}}
\end{aligned}$$

and

$$\begin{aligned}
\left| \sum_{k=1}^K \bar{X}_{h,k} \right| &\leq \sqrt{8 \max \left\{ \sum_{k=1}^K C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} [W_{h+1}^k(s,a)]^2, \frac{C_d^2 C^* C_w^2 K}{2^m} \right\} \log \frac{2H}{\delta}} \\
&\quad + \frac{8}{3} C_d C^* C_w \log \frac{2H}{\delta}
\end{aligned}$$

$$\leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} [W_{h+1}^k(s,a)]^2 \log \frac{2H}{\delta}} + 6C_d C^* C_w \log \frac{2H}{\delta}$$

holds simultaneously for all  $h \in [H]$ .  $\square$

## B Proof of main lemmas for LCB-Q (Theorem 1)

### B.1 Proof of Lemma 1

#### B.1.1 Proof of inequality (31)

To begin with, we shall control  $\sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (P_{h,s,a} - P_h^{k^n(s,a)}) V_{h+1}^{k^n(s,a)}$  by invoking Lemma 10. Let

$$W_{h+1}^i := V_{h+1}^i,$$

which satisfies

$$\|W_{h+1}^i\|_\infty \leq H =: C_w.$$

Applying Lemma 10 with  $N = N_h^k(s,a)$  reveals that, with probability at least  $1 - \delta$ ,

$$\left| \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (P_{h,s,a} - P_h^{k^n(s,a)}) V_{h+1}^{k^n(s,a)} \right| = \left| \sum_{i=1}^k X_i(s,a,h, N_h^k(s,a)) \right| \leq c_b \sqrt{\frac{H^3 \iota^2}{N_h^k(s,a)}} \quad (74a)$$

holds simultaneously for all  $(s,a,k,h) \in \mathcal{S} \times \mathcal{A} \times [K] \times [H]$ , provided that the constant  $c_b > 0$  is large enough and that  $N_h^k(s,a) > 0$ . If  $N_h^k(s,a) = 0$ , then we have the trivial bound

$$\left| \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (P_{h,s,a} - P_h^{k^n(s,a)}) V_{h+1}^{k^n(s,a)} \right| = 0. \quad (74b)$$

Additionally, from the definition  $b_n = c_b \sqrt{\frac{H^3 \iota^2}{n}}$ , we observe that

$$\begin{cases} \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n \in \left[ c_b \sqrt{\frac{H^3 \iota^2}{N_h^k(s,a)}}, 2c_b \sqrt{\frac{H^3 \iota^2}{N_h^k(s,a)}} \right], & \text{if } N_h^k(s,a) > 0 \\ \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n = 0, & \text{if } N_h^k(s,a) = 0 \end{cases} \quad (75)$$

holds simultaneously for all  $s,a,h,k \in \mathcal{S} \times \mathcal{A} \times [H] \times [K]$ , which follows directly from the property (55a) in Lemma 7.

Combining the above bounds (74) and (75), we arrive at the advertised result

$$\left| \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (P_{h,s,a} - P_h^{k^n(s,a)}) V_{h+1}^{k^n(s,a)} \right| \leq \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n.$$

#### B.1.2 Proof of inequality (32)

Note that the second inequality of (32) holds straightforwardly as

$$V_h^\pi(s) \leq V^*(s)$$

holds for any policy  $\pi$ . As a consequence, it suffices to establish the first inequality of (32), namely,

$$V_h^k(s) \leq V_h^{\pi^k}(s) \quad \text{for all } (s,h,k) \in \mathcal{S} \times [H] \times [K]. \quad (76)$$

Before proceeding, let us introduce the following auxiliary index

$$k_o(h, k, s) := \max \left\{ l : l < k \text{ and } V_h^l(s) = \max_a Q_h^l(s, a) \right\} \quad (77)$$

for any  $(h, k, s) \in [H] \times [K] \times \mathcal{S}$ , which denotes the index of the latest episode — before the end of the  $(k-1)$ -th episode — in which  $V_h(s)$  has been updated. In what follows, we shall often abbreviate  $k_o(h, k, s)$  as  $k_o(h)$  whenever it is clear from the context.

Towards establishing the relation (76), we proceed by means of an inductive argument. In what follows, we shall first justify the desired inequality for the base case when  $h+1 = H+1$  for all episodes  $k \in [K]$ , and then use induction to complete the argument for other cases. More specifically, consider any step  $h \in [H]$  in any episode  $k \in [K]$ , and suppose that the first inequality of (32) is satisfied for all previous episodes as well as all steps  $h' \geq h+1$  in the current episode, namely,

$$V_{h'}^{k'}(s) \leq V_{h'}^{\pi^{k'}}(s) \quad \text{for all } (k', h', s) \in [k-1] \times [H+1] \times \mathcal{S}, \quad (78a)$$

$$V_{h'}^k(s) \leq V_{h'}^{\pi^k}(s) \quad \text{for all } h' \geq h+1 \text{ and } s \in \mathcal{S}. \quad (78b)$$

We intend to justify that the following is valid

$$V_h^k(s) \leq V_h^{\pi^k}(s) \quad \text{for all } s \in \mathcal{S}, \quad (79)$$

assuming that the induction hypothesis (78) holds.

**Step 1: base case.** Let us begin with the base case when  $h+1 = H+1$  for all episodes  $k \in [K]$ . Recognizing the fact that  $V_{H+1}^\pi = V_{H+1}^k = 0$  for any  $\pi$  and any  $k \in [K]$ , we directly arrive at

$$V_{H+1}^k(s) \leq V_{H+1}^{\pi^k}(s) \quad \text{for all } (k, s) \in [K] \times \mathcal{S}. \quad (80)$$

**Step 2: induction.** To justify (79) under the induction hypothesis (78), we decompose the difference term to obtain

$$\begin{aligned} V_h^{\pi^k}(s) - V_h^k(s) &= V_h^{\pi^k}(s) - \max \left\{ \max_a Q_h^k(s, a), V_h^{k-1}(s) \right\} \\ &= Q_h^{\pi^k}(s, \pi_h^k(s)) - \max \left\{ \max_a Q_h^k(s, a), V_h^{k_o(h)}(s) \right\}, \end{aligned} \quad (81)$$

where the last line holds since  $V_h(s)$  has not been updated during episodes  $k_o(h), k_o(h) + 1, \dots, k-1$  (in view of the definition of  $k_o(h)$  in (77)). We shall prove that the right-hand side of (81) is non-negative by discussing the following two cases separately.

- Consider the case where  $V_h^k(s) = \max_a Q_h^k(s, a)$ . Before continuing, it is easily observed from the update rule in line 13 and line 12 of Algorithm 1 that:  $V_h(s)$  and  $\pi_h(s)$  are updated hand-in-hand for every  $h$ . Thus, it implies that

$$\pi_h^k(s) = \arg \max_a Q_h^k(s, a), \quad \text{when } V_h^k(s) = \max_a Q_h^k(s, a) \quad (82)$$

holds for all  $(k, h) \in [K] \times [H]$ . As a result, we express the term of interest as follows:

$$V_h^{\pi^k}(s) - V_h^k(s) = Q_h^{\pi^k}(s, \pi_h^k(s)) - \max_a Q_h^k(s, a) = Q_h^{\pi^k}(s, \pi_h^k(s)) - Q_h^k(s, \pi_h^k(s)). \quad (83)$$

To continue, we turn to controlling a more general term  $Q_h^{\pi^k}(s, a) - Q_h^k(s, a)$  for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ . Invoking the fact  $\eta_0^{N_h^k} + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} = 1$  (see (25) and (26)) leads to

$$Q_h^{\pi^k}(s, a) = \eta_0^{N_h^k} Q_h^{\pi^k}(s, a) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} Q_h^{\pi^k}(s, a).$$

This relation combined with (29) allows us to express the difference between  $Q_h^{\pi^k}$  and  $Q_h^k$  as follows

$$\begin{aligned}
Q_h^{\pi^k}(s, a) - Q_h^k(s, a) &= \eta_0^{N_h^k} \left( Q_h^{\pi^k}(s, a) - Q_h^1(s, a) \right) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left[ Q_h^{\pi^k}(s, a) - r_h(s, a) - V_{h+1}^{k^n}(s_{h+1}^{k^n}) + b_n \right] \\
&\stackrel{(i)}{=} \eta_0^{N_h^k} \left( Q_h^{\pi^k}(s, a) - Q_h^1(s, a) \right) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left[ P_{h,s,a} V_{h+1}^{\pi^k} - V_{h+1}^{k^n}(s_{h+1}^{k^n}) + b_n \right] \\
&\stackrel{(ii)}{\geq} \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left[ P_{h,s,a} V_{h+1}^{\pi^k} - V_{h+1}^{k^n}(s_{h+1}^{k^n}) + b_n \right] \\
&\stackrel{(iii)}{=} \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} \left( V_{h+1}^{\pi^k} - V_{h+1}^{k^n} \right) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left[ \left( P_{h,s,a} - P_h^{k^n} \right) V_{h+1}^{k^n} + b_n \right] \\
&\stackrel{(iv)}{\geq} \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left[ \left( P_{h,s,a} - P_h^{k^n} \right) V_{h+1}^{k^n} + b_n \right]. \tag{84}
\end{aligned}$$

Here, (i) invokes the Bellman equation  $Q_h^{\pi^k}(s, a) = r_h(s, a) + P_{h,s,a} V_{h+1}^{\pi^k}$ ; (ii) holds since  $Q_h^{\pi^k}(s, a) \geq 0 = Q_h^1(s, a)$ ; (iii) relies on the notation (27); and (iv) comes from the fact

$$V_{h+1}^{\pi^k} \geq V_{h+1}^k \geq V_{h+1}^{k^n},$$

owing to the induction hypothesis in (78) as well as the monotonicity of  $V_{h+1}$  in (30). Consequently, it follows from (84) that

$$\begin{aligned}
Q_h^{\pi^k}(s, a) - Q_h^k(s, a) &\geq \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} \left( P_{h,s,a} - P_h^{k^n(s,a)} \right) V_{h+1}^{k^n(s,a)} + \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n \\
&\geq \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n - \left| \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} \left( P_{h,s,a} - P_h^{k^n(s,a)} \right) V_{h+1}^{k^n(s,a)} \right| \geq 0 \tag{85}
\end{aligned}$$

for all state-action pair  $(s, a)$ , where the last inequality holds due to the bound (31) in Lemma 1. Plugging the above result into (83) directly establishes that

$$V_h^{\pi^k}(s) - V_h^k(s) = Q_h^{\pi^k}(s, \pi^k(s)) - Q_h^k(s, \pi^k(s)) \geq 0. \tag{86}$$

- When  $V_h^k(s) = V_h^{k_o(h)}(s)$ , it indicates that

$$V_h^{k_o(h)}(s) = \max_a Q_h^{k_o(h)}(s, a), \quad \pi_h^{k_o(h)}(s) = \arg \max_a Q_h^{k_o(h)}(s, a), \tag{87}$$

which follows from the definition of  $k_o(h)$  in (77) and the corresponding fact in (82). We also make note of the fact that

$$\pi_h^k(s) = \pi_h^{k_o(h)}(s), \tag{88}$$

which holds since  $V_h(s)$  (and hence  $\pi_h(s)$ ) has not been updated during episodes  $k_o(h), k_o(h)+1, \dots, k-1$  (in view of the definition (77)). Combining the above two results, we can show that

$$\begin{aligned}
V_h^{\pi^k}(s) - V_h^k(s) &= Q_h^{\pi^k}(s, \pi_h^k(s)) - V_h^{k_o(h)}(s) = Q_h^{\pi^k}(s, \pi_h^k(s)) - \max_a Q_h^{k_o(h)}(s, a) \\
&= Q_h^{\pi^k}(s, \pi_h^{k_o(h)}(s)) - Q_h^{k_o(h)}(s, \pi_h^{k_o(h)}(s)) \\
&\geq 0, \tag{89}
\end{aligned}$$

where the final line can be verified using exactly the same argument as in the previous case to show (84) and then (86). Here, we omit the proof of this step for brevity.

To conclude, substituting the relations (86) and (89) in the above two cases back into (81), we arrive at

$$V_h^{\pi^k}(s) - V_h^k(s) \geq 0$$

as desired in (79). This immediately completes the induction argument.

## B.2 Proof of Lemma 2

We make the observation that Lemma 2 would follow immediately if we could establish the following relation:

$$\begin{aligned} A_h &:= \sum_{k=1}^K \underbrace{\sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (V_{h+1}^* - V_{h+1}^{k^n(s,a)})}_{=: A_{h,k}} \\ &\leq \underbrace{\sum_{k=1}^K \left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s))}_{=: B_{h,k}} + 24 \sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 12 H C^* \log \frac{2H}{\delta}. \end{aligned} \quad (90)$$

The remainder of the proof is thus dedicated to proving (90).

To continue, let us first consider two auxiliary sequences  $\{Y_{h,k}\}_{k=1}^K$  and  $\{Z_{h,k}\}_{k=1}^K$  which are the empirical estimates of  $A_{h,k}$  and  $B_{h,k}$ , respectively. For any time step  $h$  in episode  $k$ ,  $Y_{h,k}$  and  $Z_{h,k}$  are defined as follows

$$\begin{aligned} Y_{h,k} &:= \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k, a_h^k} \sum_{n=1}^{N_h^k(s_h^k, a_h^k)} \eta_n^{N_h^k(s_h^k, a_h^k)} (V_{h+1}^* - V_{h+1}^{k^n(s_h^k, a_h^k)}), \\ Z_{h,k} &:= \left(1 + \frac{1}{H}\right) \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k, a_h^k} (V_{h+1}^* - V_{h+1}^k). \end{aligned}$$

To begin with, let us establish the relationship between  $\{Y_{h,k}\}_{k=1}^K$  and  $\{Z_{h,k}\}_{k=1}^K$ :

$$\begin{aligned} \sum_{k=1}^K Y_{h,k} &= \sum_{k=1}^K \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k, a_h^k} \sum_{n=1}^{N_h^k(s_h^k, a_h^k)} \eta_n^{N_h^k(s_h^k, a_h^k)} (V_{h+1}^* - V_{h+1}^{k^n(s_h^k, a_h^k)}) \\ &\stackrel{(i)}{=} \sum_{l=1}^K \frac{d_h^{\pi^*}(s_h^l, a_h^l)}{d_h^\mu(s_h^l, a_h^l)} P_{h,s_h^l, a_h^l} \left\{ \sum_{N=N_h^l(s_h^l, a_h^l)}^{N_h^K(s_h^l, a_h^l)} \eta_N^{N_h^K(s_h^l, a_h^l)} \right\} (V_{h+1}^* - V_{h+1}^l) \end{aligned} \quad (91)$$

$$\leq \left(1 + \frac{1}{H}\right) \sum_{k=1}^K \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k, a_h^k} (V_{h+1}^* - V_{h+1}^k) = \sum_{k=1}^K Z_{h,k}. \quad (92)$$

Here, (i) holds by replacing  $k^n(s_h^k, a_h^k)$  with  $l$  and gathering all terms that involve  $V_{h+1}^* - V_{h+1}^l$ ; in the last line, we have invoked the property  $\sum_{N=n}^{N_h^K(s,a)} \eta_n^N \leq \sum_{N=n}^\infty \eta_n^N = 1 + 1/H$  (see (55b)) together with the fact  $V_{h+1}^* - V_{h+1}^l \geq 0$  (see Lemma 1), and have further replaced  $l$  with  $k$ .

With the above relation in hand, in order to verify (90), we further decompose  $A_h$  into several terms

$$\begin{aligned} A_h &= \sum_{k=1}^K A_{h,k} = \sum_{k=1}^K Y_{h,k} + \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \stackrel{(i)}{\leq} \sum_{k=1}^K Z_{h,k} + \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \\ &= \sum_{k=1}^K B_{h,k} + \sum_{k=1}^K (Z_{h,k} - B_{h,k}) + \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \end{aligned} \quad (93)$$

where (i) follows from (92).

As a result, it remains to control  $\sum_{k=1}^K (Z_{h,k} - B_{h,k})$  and  $\sum_{k=1}^K (A_{h,k} - Y_{h,k})$  separately in the following.

**Step 1: controlling  $\sum_{k=1}^K (A_{h,k} - Y_{h,k})$ .** We shall first control this term by means of Lemma 11. Specifically, consider

$$W_{h+1}^k(s, a) := \sum_{n=1}^{N_h^k(s, a)} \eta_n^{N_h^k(s, a)} \left( V_{h+1}^* - V_{h+1}^{k^n(s, a)} \right), \quad C_d := 1 \quad (94)$$

which satisfies

$$\|W_{h+1}^k(s, a)\|_\infty \leq \sum_{n=1}^{N_h^k(s, a)} \eta_n^{N_h^k(s, a)} \left( \|V_{h+1}^*\|_\infty + \|V_{h+1}^{k^n(s, a)}\|_\infty \right) \leq 2H =: C_w. \quad (95)$$

Here we use the fact that  $\eta_0^{N_h^k} + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} = 1$  (see (25) and (26)). Then, applying Lemma 11 with (94), we have with probability at least  $1 - \delta$ , the following inequality holds true

$$\begin{aligned} \left| \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \right| &= \left| \sum_{k=1}^K X_{h,k} \right| \\ &\leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) [P_{h,s,a} W_{h+1}^k(s, a)]^2 \log \frac{2H}{\delta} + 2C_d C^* C_w \log \frac{2H}{\delta}} \\ &\stackrel{(i)}{\leq} \sqrt{8C^* \sum_{k=1}^K \|W_{h+1}^k(s, a)\|_\infty^2 \log \frac{2H}{\delta} + 4HC^* \log \frac{2H}{\delta}} \\ &\leq 8\sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 4HC^* \log \frac{2H}{\delta}, \end{aligned} \quad (96)$$

where (i) holds since  $|P_{h,s,a} W_{h+1}^k(s, a)| \leq \|P_{h,s,a}\|_1 \|W_{h+1}^k(s, a)\|_\infty = \|W_{h+1}^k(s, a)\|_\infty$ .

**Step 2: controlling  $\sum_{k=1}^K (Z_{h,k} - B_{h,k})$ .** Similarly, we shall control  $\sum_{k=1}^K (Z_{h,k} - B_{h,k})$  by invoking Lemma 11. Recall that

$$Z_{h,k} - B_{h,k} = \left(1 + \frac{1}{H}\right) \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k, a_h^k} (V_{h+1}^* - V_{h+1}^k) - \left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)), \quad (97)$$

and let us consider

$$W_{h+1}^k(s, a) := V_{h+1}^* - V_{h+1}^k, \quad C_d := \left(1 + \frac{1}{H}\right) \leq 2 \quad (98)$$

which satisfies

$$\|W_{h+1}^k(s, a)\|_\infty \leq \|V_{h+1}^*\|_\infty + \|V_{h+1}^k\|_\infty \leq 2H =: C_w. \quad (99)$$

Again, in view of Lemma 11, we can show that with probability at least  $1 - \delta$ ,

$$\begin{aligned} \left| \sum_{k=1}^K (B_{h,k} - Z_{h,k}) \right| &= \left| \sum_{k=1}^K X_{h,k} \right| \\ &\leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) [P_{h,s,a} W_{h+1}^k(s, a)]^2 \log \frac{2H}{\delta} + 2C_d C^* C_w \log \frac{2H}{\delta}} \\ &\stackrel{(i)}{\leq} \sqrt{32C^* \sum_{k=1}^K \|W_{h+1}^k(s, a)\|_\infty^2 \log \frac{2H}{\delta} + 8HC^* \log \frac{2H}{\delta}} \\ &\leq 16\sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 8HC^* \log \frac{2H}{\delta}, \end{aligned} \quad (100)$$

where (i) holds due to the fact  $\|P_{h,s,a}\|_1 = 1$ .



**Step 3: putting all this together.** Substitution results in (96) and (100) back into (93) completes the proof of (90) as follows

$$\begin{aligned} A_h &\leq \sum_{k=1}^K B_{h,k} + \left| \sum_{k=1}^K (Z_{h,k} - B_{h,k}) \right| + \left| \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \right| \\ &\leq \sum_{k=1}^K B_{h,k} + 24\sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 12HC^* \log \frac{2H}{\delta}. \end{aligned}$$

This in turn concludes the proof of Lemma 2.

### B.3 Proof of Lemma 3

Recall that the term of interest in (42) is given by

$$\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \left( 24\sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 12HC^* \log \frac{2H}{\delta} \right) + \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} I_h. \quad (101)$$

First, it is easily seen that

$$\left(1 + \frac{1}{H}\right)^{h-1} \leq \left(1 + \frac{1}{H}\right)^H \leq e \quad \text{for every } h = 1, \dots, H, \quad (102)$$

which taken collectively with the expression of the first term in (101) yields

$$\begin{aligned} \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \left( 24\sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 12HC^* \log \frac{2H}{\delta} \right) &\leq 24e \sum_{h=1}^H \left( \sqrt{H^2 C^* K \log \frac{2H}{\delta}} + HC^* \log \frac{2H}{\delta} \right) \\ &\lesssim \sqrt{H^4 C^* K \log \frac{H}{\delta}} + H^2 C^* \log \frac{H}{\delta}. \end{aligned} \quad (103)$$

As a result, it remains to control the second term in (101). Plugging the expression of  $I_h$  (cf. (39)) and invoking the fact (102) give

$$\begin{aligned} \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} I_h &= \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \eta_0^{N_h^k(s,a)} H \\ &\quad + 2 \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n \\ &\leq \underbrace{e \sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \eta_0^{N_h^k(s,a)} H}_{=:A} + \underbrace{2e \sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n}_{=:B}. \end{aligned} \quad (104)$$

**Step 1: controlling the quantities  $A$  and  $B$  in (104).** We first develop an upper bound on the quantity  $A$  in (104). Recognizing the fact that  $\eta_0^N = 0$  for any  $N > 0$  (see (25)), we have

$$\begin{aligned} A &= e \sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \eta_0^{N_h^k(s,a)} H \\ &\leq eH \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{k=1}^K \mathbb{1}(N_h^k(s, a) < 1) \end{aligned}$$

$$\begin{aligned}
&\leq eH \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \frac{8\iota}{d_h^\mu(s,a)} + eH \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=\lceil \frac{8\iota}{d_h^\mu(s,a)} \rceil}^K \mathbb{1}(N_h^k(s,a) < 1) \\
&= eH \sum_{h=1}^H \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s, \pi^*(s)) \frac{8\iota}{d_h^\mu(s, \pi^*(s))} + eH \sum_{h=1}^H \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s, \pi^*(s)) \sum_{k=\lceil \frac{8\iota}{d_h^\mu(s, \pi^*(s))} \rceil}^K \mathbb{1}(N_h^k(s, \pi^*(s)) < 1),
\end{aligned}$$

where the last equality holds since  $\pi^*$  is a deterministic policy (so that  $d_h^{\pi^*}(s,a) \neq 0$  only when  $a = \pi^*(s)$ ). Recalling  $\frac{d_h^{\pi^*}(s,a)}{d_h^\mu(s,a)} \leq C^*$  under Assumption 1, we can further bound  $A$  by

$$\begin{aligned}
A &\leq 8eH^2SC^*\iota + eH \sum_{h=1}^H \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s, \pi^*(s)) \sum_{k=\lceil \frac{8\iota}{d_h^\mu(s, \pi^*(s))} \rceil}^K \mathbb{1}(N_h^k(s, \pi^*(s)) < 1) \\
&= 8eH^2SC^*\iota,
\end{aligned} \tag{105}$$

where the last inequality follows since when  $k \geq \frac{8\iota}{d_h^\mu(s,a)}$ , one has — with probability at least  $1 - \delta$  — that

$$N_h^k(s,a) \geq \frac{kd_h^\mu(s,a)}{8\iota} \geq 1,$$

holds simultaneously for all  $(s,a,h,k) \in \mathcal{S} \times \mathcal{A} \times [K] \times [H]$  (as implied by (57a)).

Turning to the quantity  $B$  in (104), one can deduce that

$$\begin{aligned}
B &= 2e \sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} b_n \\
&\lesssim \sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sqrt{\frac{H^3\iota^2}{N_h^k(s,a) \vee 1}} = \sum_{h=1}^H \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s, \pi^*(s)) \sqrt{\frac{H^3\iota^2}{N_h^k(s, \pi^*(s)) \vee 1}},
\end{aligned} \tag{106}$$

where the inequality follows from inequality (75), and the last equality is valid since  $\pi^*$  is a deterministic policy. To further control the right hand side above, Lemma 8 provides an upper bound for  $\sqrt{1/(N_h^k(s, \pi^*(s)) \vee 1)}$  which in turn leads to

$$\begin{aligned}
B &\lesssim \sqrt{H^3\iota^3} \sum_{h=1}^H \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s, \pi^*(s)) \sqrt{\frac{1}{kd_h^\mu(s, \pi^*(s))}} \\
&\lesssim \sqrt{H^3C^*\iota^3} \sum_{h=1}^H \sum_{k=1}^K \sum_{s \in \mathcal{S}} \sqrt{d_h^{\pi^*}(s, \pi^*(s))} \sqrt{\frac{1}{k}} \\
&\lesssim \sqrt{H^5C^*K\iota^3} \max_h \sum_{s \in \mathcal{S}} \sqrt{d_h^{\pi^*}(s, \pi^*(s))} \\
&\lesssim \sqrt{H^5C^*K\iota^3} \cdot \left( \sqrt{S} \cdot \sqrt{\sum_{s \in \mathcal{S}} d_h^{\pi^*}(s, \pi^*(s))} \right) \asymp \sqrt{H^5SC^*K\iota^3},
\end{aligned} \tag{107}$$

where the second inequality follows from the fact  $\frac{d_h^{\pi^*}(s,a)}{d_h^\mu(s,a)} \leq C^*$  under Assumption 1, and the last line invokes the Cauchy-Schwarz inequality.

Taking the upper bounds on both  $A$  and  $B$  collectively establishes

$$\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} I_h \leq A + B \lesssim H^2SC^*\iota + \sqrt{H^5SC^*K\iota^3}. \tag{108}$$

**Step 2: putting everything together.** Combining (103) and (108) allows us to establish that

$$\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \left( I_h + 16\sqrt{H^2 C^* K \log \frac{2H}{\delta}} + 8HC^* \log \frac{2H}{\delta} \right) \lesssim H^2 SC^* \iota + \sqrt{H^5 SC^* K \iota^3},$$

as advertised.

## C Proof of lemmas for LCB-Q-Advantage (Theorem 2)

**Additional notation for LCB-Q-Advantage.** Let us also introduce, and remind the reader of, several notation of interest in Algorithm 5 as follows.

- $N_h^k(s, a)$  (resp.  $N_h^{(m,t)}(s, a)$ ) denotes the value of  $N_h(s, a)$  — the number of episodes that has visited  $(s, a)$  at step  $h$  at the *beginning* of the  $k$ -th episode (resp. the *beginning* of  $t$ -th episode of the  $m$ -th epoch); for the sake of conciseness, we shall often abbreviate  $N_h^k = N_h^k(s, a)$  (resp.  $N_h^{(m,t)} = N_h^{(m,t)}(s, a)$ ) when it is clear from context.
- $L_m = 2^m$ : the total number of in-epoch episodes in the  $m$ -th epoch.
- $k_h^n(s, a)$ : the index of the episode in which  $(s, a)$  is visited for the  $n$ -th time at time step  $h$ ;  $(m_h^n(s, a), t_h^n(s, a))$  denote respectively the index of the epoch and that of the in-epoch episode in which  $(s, a)$  is visited for the  $n$ -th time at step  $h$ ; for the sake of conciseness, we shall often use the shorthand  $k^n = k_h^n(s, a)$ ,  $(m^n, k^n) = (m_h^n(s, a), k_h^n(s, a))$  whenever it is clear from context.
- $Q_h^k(s, a)$ ,  $Q_h^{\text{LCB},k}(s, a)$ ,  $\bar{Q}_h^k(s, a)$  and  $V_h^k(s)$  are used to denote  $Q_h(s, a)$ ,  $Q_h^{\text{LCB}}(s, a)$ ,  $\bar{Q}_h(s, a)$ , and  $V_h(s)$  at the *beginning* of the  $k$ -th episode, respectively.
- $\bar{V}_h^k(s)$ ,  $\bar{V}_h^{\text{next},k}(s)$ ,  $\bar{\mu}_h^k(s, a)$ ,  $\bar{\mu}_h^{\text{next},k}(s, a)$  denote the values of  $\bar{V}_h(s)$ ,  $\bar{V}_h^{\text{next}}(s)$ ,  $\bar{\mu}_h(s, a)$  and  $\bar{\mu}_h^{\text{next}}(s, a)$  at the *beginning* of the  $k$ -th episode, respectively.
- $\hat{N}_h^{(m,t)}(s, a)$  represents  $\hat{N}_h(s, a)$  at the *beginning* of the  $t$ -th in-epoch episode in the  $m$ -th epoch.
- $\hat{N}_h^{\text{epo},m}(s, a)$  denotes  $\hat{N}_h^{(m,L_m+1)}(s, a)$ , representing the number of visits to  $(s, a)$  in the entire duration of the  $m$ -th epoch.
- $[\mu_h^{\text{ref},k}, \sigma_h^{\text{ref},k}, \mu_h^{\text{adv},k}, \sigma_h^{\text{adv},k}, \bar{\delta}_h^k, \bar{B}_h^k, \bar{b}_h^k]$ : the values of  $[\mu_h^{\text{ref}}, \sigma_h^{\text{ref}}, \mu_h^{\text{adv}}, \sigma_h^{\text{adv}}, \bar{\delta}_h, \bar{B}_h, \bar{b}_h]$  at the *beginning* of the  $k$ -th episode, respectively.

In addition, for a fixed vector  $V \in \mathbb{R}^{|S|}$ , let us define a variance parameter with respect to  $P_{h,s,a}$  as follows

$$\text{Var}_{h,s,a}(V) := \mathbb{E}_{s' \sim P_{h,s,a}} \left[ (V(s') - P_{h,s,a}V)^2 \right] = P_{h,s,a}(V^2) - (P_{h,s,a}V)^2. \quad (109)$$

This notation will be useful in the subsequent proof. We remind the reader that there exists a one-to-one mapping between the index of the episode  $k$  and the index pair  $(m, t)$  (i.e., the epoch  $m$  and in-epoch episode  $t$ ), as specified in (45). In the following, for any episode  $k$ , we recall the expressions of  $\bar{V}_{h+1}$  and  $\bar{\mu}_h$  (which is the running mean of  $\bar{V}_{h+1}$ ).

- Recalling the update rule of  $\bar{V}_h$  and  $\bar{V}_h^{\text{next}}$  in line 26 and line 27 of Algorithm 5, we observe that both the reference values for the current epoch  $\bar{V}_h$  and for the next epoch  $\bar{V}_h^{\text{next}}$  remain unchanged within each epoch. Additionally, for any epoch  $m$ ,  $\bar{V}_h$  takes the value of  $\bar{V}_h^{\text{next}}$  in the previous  $(m-1)$ -th epoch; namely, for any episode  $k$  happening in the  $m$ -th epoch, we have

$$\bar{V}_h^k = \bar{V}_h^{\text{next},k'} \quad (110)$$

for all episode  $k'$  within the  $(m-1)$ -th epoch.

- $\bar{\mu}_h^k$  serves as the estimate of  $P_{h,s,a}\bar{V}_{h+1}^k$  constructed by the samples in the previous  $(m-1)$ -th epoch (collected by updating  $\bar{\mu}_h^{\text{next}}$ ). Recall the update rule of  $\bar{\mu}_h$  in line 26 and line 24 of Algorithm 5: for any  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ , we can write  $\bar{\mu}_h^k$  as

$$\begin{aligned}\bar{\mu}_h^k(s, a) &= \bar{\mu}_h^{(m,1)}(s, a) = \bar{\mu}_h^{\text{next},(m,1)}(s, a) = \bar{\mu}_h^{\text{next},(m-1,L_{m-1})}(s, a) \\ &= \frac{\sum_{i=N_h^{(m-1,1)}+1}^{N_h^{(m,1)}} \bar{V}_{h+1}^{\text{next},k^i}(s_{h+1}^{k^i})}{\widehat{N}_h^{\text{epo},m-1}(s, a) \vee 1} = \frac{\sum_{i=N_h^{(m-1,1)}+1}^{N_h^{(m,1)}} \bar{V}_{h+1}^k(s_{h+1}^{k^i})}{\widehat{N}_h^{\text{epo},m-1}(s, a) \vee 1},\end{aligned}\quad (111)$$

where the last equality follows from (110) using the fact that the indices of episodes in which  $(s, a)$  is visited within the  $(m-1)$ -th epoch are  $\{i : i = N_h^{(m-1,1)} + 1, N_h^{(m-1,1)} + 2, \dots, N_h^{(m,1)}\}$ .

Finally, according to the update rules of  $\mu_h^{\text{adv},k^{n+1}}(s_h^k, a_h^k)$  and  $\sigma_h^{\text{adv},k^{n+1}}(s_h^k, a_h^k)$  in lines 11-12 of Algorithm 3, we have

$$\begin{aligned}\mu_h^{\text{adv},k^{n+1}}(s_h^k, a_h^k) &= \mu_h^{\text{adv},k^n+1}(s_h^k, a_h^k) = (1 - \eta_n)\mu_h^{\text{adv},k^n}(s_h^k, a_h^k) + \eta_n(V_{h+1}^{k^n}(s_{h+1}^{k^n}) - \bar{V}_{h+1}^{k^n}(s_{h+1}^{k^n})), \\ \sigma_h^{\text{adv},k^{n+1}}(s_h^k, a_h^k) &= \sigma_h^{\text{adv},k^n+1}(s_h^k, a_h^k) = (1 - \eta_n)\sigma_h^{\text{adv},k^n}(s_h^k, a_h^k) + \eta_n(V_{h+1}^{k^n}(s_{h+1}^{k^n}) - \bar{V}_{h+1}^{k^n}(s_{h+1}^{k^n}))^2.\end{aligned}$$

Applying this relation recursively and invoking the definitions of  $\eta_n^{N_h^k}$  in (25) give

$$\mu_h^{\text{adv},k^{N_h^k+1}}(s, a) = \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_h^{k^n} (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}), \quad \sigma_h^{\text{adv},k^{N_h^k+1}}(s, a) = \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_h^{k^n} (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})^2. \quad (112)$$

Similarly, according to the update rules of  $\mu_h^{\text{ref},k^{n+1}}(s, a)$  and  $\sigma_h^{\text{ref},k^{n+1}}(s, a)$  in lines 9-10 of Algorithm 3, we obtain

$$\begin{aligned}\mu_h^{\text{ref},k^{n+1}}(s, a) &= \mu_h^{\text{ref},k^n+1}(s, a) = \left(1 - \frac{1}{n}\right) \mu_h^{\text{ref},k^n}(s, a) + \frac{1}{n} \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n}), \\ \sigma_h^{\text{ref},k^{n+1}}(s, a) &= \sigma_h^{\text{ref},k^n+1}(s, a) = \left(1 - \frac{1}{n}\right) \sigma_h^{\text{ref},k^n}(s, a) + \frac{1}{n} \left(\bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n})\right)^2.\end{aligned}$$

Simple recursion leads to

$$\mu_h^{\text{ref},k^{N_h^k+1}}(s, a) = \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} P_h^{k^n} \bar{V}_{h+1}^{\text{next},k^n}, \quad \sigma_h^{\text{ref},k^{N_h^k+1}}(s, a) = \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} P_h^{k^n} (\bar{V}_{h+1}^{\text{next},k^n})^2. \quad (113)$$

## C.1 Proof of Lemma 4

Akin to the proof of Lemma 1, the second inequality of (47) holds trivially since

$$V_h^\pi(s) \leq V_h^*(s)$$

holds for any policy  $\pi$ . Thus, it suffices to focus on justifying the first inequality of (47), namely,

$$V_h^k(s) \leq V_h^{\pi^k}(s) \quad \forall (k, h, s) \in [K] \times [H] \times \mathcal{S}, \quad (114)$$

which we shall prove by induction.

**Step 1: introducing the induction hypothesis.** For notational simplicity, let us define

$$k_o(h, k, s) := \max \left\{ l : l < k \text{ and } V_h^l(s) = \max_a \max \left\{ Q_h^{\text{LCB},l}(s, a), \bar{Q}_h^l(s, a) \right\} \right\} \quad (115)$$

for any  $(h, k, s) \in [H] \times [K] \times \mathcal{S}$ . Here,  $k_o(h, k, s)$  denotes the index of the latest episode — right at the end of the  $(k-1)$ -th episode — in which  $V_h(s)$  has been updated, which shall be abbreviated as  $k_o(h)$  whenever it is clear from context.

In what follows, we shall first justify the advertised inequality for the base case where  $h = H + 1$  for all episodes  $k \in [K]$ , followed by an induction argument. Regarding the induction part, let us consider any  $k \in [K]$  and any  $h \in [H]$ , and suppose that

$$V_{h'}^{k'}(s) \leq V_{h'}^{\pi^{k'}}(s) \quad \text{for all } (k', h', s) \in [k-1] \times [H+1] \times \mathcal{S}, \quad (116a)$$

$$V_{h'}^k(s) \leq V_{h'}^{\pi^k}(s) \quad \text{for all } h' \geq h+1 \text{ and } s \in \mathcal{S}. \quad (116b)$$

We intend to justify

$$V_h^k(s) \leq V_h^{\pi^k}(s) \quad \forall s \in \mathcal{S}, \quad (117)$$

assuming that the induction hypotheses (116) hold.

**Step 2: controlling the confident bound  $\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1}$ .** Before proceeding, we first introduce an auxiliary result on bounding  $\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1}$ , which plays a crucial role. For any  $(s, a)$ , it is easily seen that

$$N_h^k(s, a) = 0 \quad \implies \quad \sum_{n=1}^{N_h^k(s, a)} \eta_n^{N_h^k(s, a)} \bar{b}_h^{k^n(s, a)+1} = 0. \quad (118)$$

When  $N_h^k(s, a) > 0$ , expanding the definitions of  $\bar{b}_h^{k^n+1}$  (cf. line 6 of Algorithm 3) and  $\bar{\delta}_h^{k+1}$  (cf. line 15 of Algorithm 3) leads to

$$\begin{aligned} & \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1} \\ &= \sum_{n=1}^{N_h^k} \eta_n \prod_{i=n+1}^{N_h^k} (1 - \eta_i) \cdot \left( \left(1 - \frac{1}{\eta_n}\right) \bar{B}_h^{k^n}(s, a) + \frac{1}{\eta_n} \bar{B}_h^{k^n+1}(s, a) \right) + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n^{3/4}} H^{7/4} \iota + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n} H^2 \iota \\ &= \sum_{n=1}^{N_h^k} \left( \prod_{i=n+1}^{N_h^k} (1 - \eta_i) \bar{B}_h^{k^n+1}(s, a) - \prod_{i=n}^{N_h^k} (1 - \eta_i) \bar{B}_h^{k^n}(s, a) \right) + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n^{3/4}} H^{7/4} \iota + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n} H^2 \iota \\ &\stackrel{(i)}{=} \sum_{n=1}^{N_h^k} \prod_{i=n+1}^{N_h^k} (1 - \eta_i) \bar{B}_h^{k^n+1}(s, a) - \sum_{n=2}^{N_h^k} \prod_{i=n}^{N_h^k} (1 - \eta_i) \bar{B}_h^{k^n}(s, a) + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n^{3/4}} H^{7/4} \iota + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n} H^2 \iota \\ &\stackrel{(ii)}{=} \sum_{n=1}^{N_h^k} \prod_{i=n+1}^{N_h^k} (1 - \eta_i) \bar{B}_h^{k^n+1}(s, a) - \sum_{n=1}^{N_h^k-1} \prod_{i=n+1}^{N_h^k} (1 - \eta_i) \bar{B}_h^{k^n+1}(s, a) + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n^{3/4}} H^{7/4} \iota + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n} H^2 \iota \\ &= \bar{B}_h^{k^{N_h^k}+1}(s, a) + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n^{3/4}} H^{7/4} \iota + c_b \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n} H^2 \iota, \end{aligned} \quad (119)$$

where we abuse the notation to let  $\prod_{i=j+1}^j (1 - \eta_i) = 1$ . Here, (i) holds since  $\bar{B}^{k^1}(s, a) = 0$ , (ii) follows from the fact that  $\bar{B}^{k^n+1}(s, a) = \bar{B}^{k^{n+1}}(s, a)$ , since  $(s, a)$  has not been visited at step  $h$  during the episodes between the indices  $k^n + 1$  and  $k^{n+1} - 1$ . Combining the above result in (119) with the properties  $\frac{1}{(N_h^k)^{3/4}} \leq$

$\sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n^{3/4}} \leq \frac{2}{(N_h^k)^{3/4}}$  and  $\frac{1}{N_h^k} \leq \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{n} \leq \frac{2}{N_h^k}$  (see Lemma 7), we arrive at

$$\bar{B}_h^{k^{N_h^k}+1}(s, a) + c_b \frac{H^{7/4} \iota}{(N_h^k)^{3/4}} + c_b \frac{H^2 \iota}{N_h^k} \leq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1} \leq \bar{B}_h^{k^{N_h^k}+1}(s, a) + 2c_b \frac{H^{7/4} \iota}{(N_h^k)^{3/4}} + 2c_b \frac{H^2 \iota}{N_h^k} \quad (120)$$

as long as  $N_h^k(s, a) > 0$ .

**Step 3: base case.** Let us look at the base case with  $h = H + 1$  for any  $k \in [K]$ . Recalling the facts that  $V_{H+1}^\pi = V_{H+1}^k = 0$  for any  $\pi$  and any  $k \in [K]$ , we reach

$$V_{H+1}^k(s) \leq V_{H+1}^{\pi^k}(s) \quad \text{for all } (k, s) \in [K] \times \mathcal{S}. \quad (121)$$

**Step 4: induction arguments.** We now turn to the induction arguments. Suppose that (116) holds for a pair  $(k, h) \in [K] \times [H]$ . Everything comes down to justifying (117) for time step  $h$  in the episode  $k$ .

First, we recall the update rule of  $V_h(s)$  in lines 21-22 of Algorithm 5:

$$V_h^k(s) = \max_a Q_h^k(s, a) = Q_h^k(s, \pi_h^k(s)) = \max \left\{ Q_h^{\text{LCB},k}(s, \pi_h^k(s)), \bar{Q}_h^k(s, \pi_h^k(s)), Q_h^{k-1}(s, \pi_h^k(s)) \right\}.$$

Then we shall verify (117) in three different cases.

- When  $V_h^k(s) = Q_h^{\text{LCB},k}(s, \pi_h^k(s))$ , the term of interest can be controlled by

$$V_h^{\pi^k}(s) - V_h^k(s) \stackrel{(i)}{=} Q_h^{\pi^k}(s, \pi_h^k(s)) - Q_h^{\text{LCB},k}(s, \pi_h^k(s)) \geq 0,$$

where (i) holds since  $\pi^k$  is set to be the greedy policy such that  $V_h^{\pi^k}(s) = Q_h^{\pi^k}(s, \pi_h^k(s))$ , and the last inequality follows directly from the analysis for LCB-Q (see (85)).

- When  $V_h^k(s) = \bar{Q}_h^k(s, \pi_h^k(s))$ , we obtain

$$V_h^{\pi^k}(s) - V_h^k(s) = Q_h^{\pi^k}(s, \pi_h^k(s)) - \bar{Q}_h^k(s, \pi_h^k(s)). \quad (122)$$

To prove the term on the right-hand side of (122) is non-negative, we proceed by developing a more general lower bound on  $Q_h^{\pi^k}(s, a) - \bar{Q}_h^k(s, a)$  for every  $(s, a) \in \mathcal{S} \times \mathcal{A}$ . Towards this, recalling the definition of  $N_h^k$  and  $k^n$ , we can express

$$\bar{Q}_h^k(s, a) = \bar{Q}_h^{k^{N_h^k+1}}(s, a).$$

Thus, according to the update rule (cf. line 7 in Algorithm 3), we arrive at

$$\begin{aligned} \bar{Q}_h^k(s, a) &= \bar{Q}_h^{k^{N_h^k+1}}(s, a) \\ &= (1 - \eta_{N_h^k}) \bar{Q}_h^{k^{N_h^k}}(s, a) + \eta_{N_h^k} \left\{ r_h(s, a) + V_{h+1}^{k^{N_h^k}}(s_{h+1}^{k^{N_h^k}}) - \bar{V}_{h+1}^{k^{N_h^k}}(s_{h+1}^{k^{N_h^k}}) + \bar{\mu}_h^{k^{N_h^k}}(s, a) - \bar{b}_h^{k^{N_h^k+1}} \right\}. \end{aligned}$$

Applying this relation recursively and invoking the definitions of  $\eta_0^{N_h^k}$  and  $\eta_n^{N_h^k}$  in (25) give

$$\bar{Q}_h^k(s, a) = \eta_0^{N_h^k} \bar{Q}_h^1(s, a) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left\{ r_h(s, a) + V_{h+1}^{k^n}(s_{h+1}^{k^n}) - \bar{V}_{h+1}^{k^n}(s_{h+1}^{k^n}) + \bar{\mu}_h^{k^n}(s, a) - \bar{b}_h^{k^n+1} \right\}. \quad (123)$$

Additionally, for any policy  $\pi^k$ , the basic relation  $\eta_0^{N_h^k} + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} = 1$  (see (26) and (25)) gives

$$Q_h^{\pi^k}(s, a) = \eta_0^{N_h^k} Q_h^{\pi^k}(s, a) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} Q_h^{\pi^k}(s, a). \quad (124)$$

Combing (123) and (124) leads to

$$\begin{aligned} Q_h^{\pi^k}(s, a) - \bar{Q}_h^k(s, a) &= \eta_0^{N_h^k} (Q_h^{\pi^k}(s, a) - \bar{Q}_h^1(s, a)) \\ &\quad + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left\{ Q_h^{\pi^k}(s, a) - r_h(s, a) - V_{h+1}^{k^n}(s_{h+1}^{k^n}) + \bar{V}_{h+1}^{k^n}(s_{h+1}^{k^n}) - \bar{\mu}_h^{k^n}(s, a) + \bar{b}_h^{k^n+1} \right\}. \end{aligned} \quad (125)$$

Plugging in the construction of  $\bar{\mu}_h$  in (111) and invoking the Bellman equation

$$Q_h^{\pi^k}(s, a) = r_h(s, a) + P_{h,s,a} V_{h+1}^{\pi^k}, \quad (126)$$

we arrive at

$$\begin{aligned} Q_h^{\pi^k}(s, a) - r_h(s, a) - V_{h+1}^{k^n}(s_{h+1}^{k^n}) + \bar{V}_{h+1}^{k^n}(s_{h+1}^{k^n}) - \bar{\mu}_h^{k^n}(s, a) + \bar{b}_h^{k^n+1} \\ = P_{h,s,a} V_{h+1}^{\pi^k} + \bar{V}_{h+1}^{k^n}(s_{h+1}^{k^n}) - V_{h+1}^{k^n}(s_{h+1}^{k^n}) - \frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} \bar{V}_{h+1}^{k^n}(s_{h+1}^{k^i})}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} + \bar{b}_h^{k^n+1} \\ = P_{h,s,a} V_{h+1}^{\pi^k} - V_{h+1}^{k^n}(s_{h+1}^{k^n}) + (P_h^{k^n} - P_{h,s,a}) \bar{V}_{h+1}^{k^n} + \left( P_{h,s,a} - \frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} P_h^{k^i}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} \right) \bar{V}_{h+1}^{k^n} + \bar{b}_h^{k^n+1} \\ = P_{h,s,a} (V_{h+1}^{\pi^k} - V_{h+1}^{k^n}) + \bar{b}_h^{k^n+1} + \xi_h^{k^n}, \end{aligned}$$

where

$$\xi_h^{k^n} := (P_h^{k^n} - P_{h,s,a}) (\bar{V}_{h+1}^{k^n} - V_{h+1}^{k^n}) + \left( P_{h,s,a} - \frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} P_h^{k^i}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} \right) \bar{V}_{h+1}^{k^n}. \quad (127)$$

Inserting the above result into (125) leads to the following decomposition

$$Q_h^{\pi^k}(s, a) - \bar{Q}_h^k(s, a) = \eta_0^{N_h^k} (Q_h^{\pi^k}(s, a) - \bar{Q}_h^1(s, a)) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left\{ P_{h,s,a} (V_{h+1}^{\pi^k} - V_{h+1}^{k^n}) + \bar{b}_h^{k^n+1} + \xi_h^{k^n} \right\} \quad (128)$$

$$\geq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (\bar{b}_h^{k^n+1} + \xi_h^{k^n}), \quad (129)$$

which holds by virtue of the following facts:

- (i) The initialization  $\bar{Q}_h^1(s, a) = 0$  and the non-negativity of  $Q_h^{\pi^k}(s, a)$  for any policy  $\pi$  and  $(s, a) \in \mathcal{S} \times \mathcal{A}$  lead to  $Q_h^{\pi^k}(s, a) - \bar{Q}_h^1(s, a) = Q_h^{\pi^k}(s, a) \geq 0$ .
- (ii) For any episode  $k^n$  appearing before  $k$ , making use of the induction hypothesis  $V_{h+1}^{\pi^k}(s) \geq V_{h+1}^k(s)$  in (116b) and the monotonicity of  $V_h(s)$  in (46), we obtain

$$V_{h+1}^{\pi^k}(s) - V_{h+1}^{k^n}(s) \geq V_{h+1}^k(s) - V_{h+1}^{k^n}(s) \geq 0. \quad (130)$$

The following lemma ensures that the right-hand side of (129) is non-negative. We postpone the proof of Lemma 12 to Appendix C.4 to streamline our discussion.

**Lemma 12.** *For any  $\delta \in (0, 1)$ , there exists some sufficiently large constant  $c_b > 0$ , such that with probability at least  $1 - \delta$ ,*

$$\left| \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \xi_h^{k^n} \right| \leq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1}, \quad \forall k \in [K]. \quad (131)$$

Taking this lemma together with the inequalities (122) and (129) yields

$$V_h^{\pi^k}(s) - V_h^k(s) = Q_h^{\pi^k}(s, a) - \bar{Q}_h^k(s, a) \geq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1} - \left| \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \xi_h^{k^n} \right| \geq 0.$$



- Next, consider the case where  $V_h^k(s) = Q_h^{k-1}(s, \pi_h^k(s))$ . In view of the definition of  $k_o(h)$  in (115), one has

$$V_h^k(s) = Q_h^{k-1}(s, \pi_h^k(s)) = Q_h^{k_o(h)}(s, \pi_h^k(s)) = \max \left\{ Q_h^{\text{LCB}, k_o(h)}(s, \pi_h^k(s)), \bar{Q}_h^{k_o(h)}(s, \pi_h^k(s)) \right\},$$

since  $Q_h(s, \pi_h^k(s))$  has not been updated during the episode  $k_o(h)$  and remains unchanged in the episodes  $k_o(h) + 1, k_o(h) + 2, \dots, k - 1$ . With this equality in hand, the term of interest in (117) can be controlled by

$$V_h^{\pi^k}(s) - V_h^k(s) = Q_h^{\pi^k}(s, \pi_h^k(s)) - \max \left\{ Q_h^{\text{LCB}, k_o(h)}(s, \pi_h^k(s)), \bar{Q}_h^{k_o(h)}(s, \pi_h^k(s)) \right\} \geq 0,$$

where the last inequality follows from the facts

$$\begin{aligned} Q_h^{\pi^k}(s, \pi_h^k(s)) - Q_h^{\text{LCB}, k_o(h)}(s, \pi_h^k(s)) &\stackrel{(i)}{\geq} 0, \\ Q_h^{\pi^k}(s, \pi_h^k(s)) - \bar{Q}_h^{k_o(h)}(s, \pi_h^k(s)) &\stackrel{(ii)}{\geq} 0. \end{aligned}$$

Here, (i) follows from the same analysis framework for showing (84) and (86); (ii) holds due to the following fact

$$Q_h^{\pi^k}(s, a) - \bar{Q}_h^{k_o(h)}(s, a) \geq \sum_{n=1}^{N_h^{k_o(h)}} \eta_n^{N_h^{k_o(h)}} (\bar{b}_h^{k^n+1} + \xi_h^{k^n}) \geq 0,$$

which is obtained directly by adapting (129) and then invoking (131) for  $k = k_o(h)$ ; since the analysis follows verbatim, we omit their proofs here.

Combining the above three cases verifies the induction hypothesis in (117), provided that (116) is satisfied.

**Step 5: putting everything together.** Combining the base case in Step 3 and induction arguments in Step 4, we can readily verify the induction hypothesis in Step 1, which in turn establishes Lemma 4.

## C.2 Proof of Lemma 5

For every  $h \in [H]$ , we can decompose

$$\begin{aligned} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) &\stackrel{(i)}{\leq} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s, \pi_h^*(s)) \left( Q_h^*(s, \pi_h^*(s)) - \bar{Q}_h^k(s, \pi_h^*(s)) \right) \\ &= \sum_{k=1}^K \sum_{s, a \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \left( Q_h^*(s, a) - \bar{Q}_h^k(s, a) \right), \end{aligned} \quad (132)$$

where (i) follows from the fact  $V_h^k(s) = \max_a Q_h^k(s, a) \geq \max_a \bar{Q}_h^k(s, a) \geq \bar{Q}_h^k(s, \pi_h^*(s))$  (see lines 21-22 in Algorithm 5). Here, the last equality is due to (35).

**Step 1: bounding  $Q_h^*(s, a) - \bar{Q}_h^k(s, a)$ .** The basic relation  $\eta_0^{N_h^k} + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} = 1$  (see (26) and (25)) gives

$$Q_h^*(s, a) = \eta_0^{N_h^k} Q_h^*(s, a) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} Q_h^*(s, a), \quad (133)$$

which combined with (123) leads to

$$Q_h^*(s, a) - \bar{Q}_h^k(s, a) = \eta_0^{N_h^k} (Q_h^*(s, a) - \bar{Q}_h^1(s, a))$$

$$+ \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left\{ Q_h^*(s, a) - r_h(s, a) - V_{h+1}^{k^n}(s_{h+1}^{k^n}) + \bar{V}_{h+1}^{k^n}(s_{h+1}^{k^n}) - \bar{\mu}_h^{k^n}(s, a) + \bar{b}_h^{k^n+1} \right\}. \quad (134)$$

Invoking the Bellman optimality equation

$$Q_h^*(s, a) = r_h(s, a) + P_{h,s,a} V_{h+1}^*, \quad (135)$$

we can decompose  $Q_h^*(s, a) - \bar{Q}_h^k(s, a)$  similar to (128) by inserting (127) as follows:

$$\begin{aligned} Q_h^*(s, a) - \bar{Q}_h^k(s, a) &= \eta_0^{N_h^k} (Q_h^*(s, a) - \bar{Q}_h^1(s, a)) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left\{ P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n}) + \bar{b}_h^{k^n+1} + \xi_h^{k^n} \right\} \\ &\stackrel{(i)}{\leq} \eta_0^{N_h^k} H + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (\bar{b}_h^{k^n+1} + \xi_h^{k^n}) + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n}) \\ &\stackrel{(ii)}{\leq} \eta_0^{N_h^k} H + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n}) + 2 \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1} \\ &\stackrel{(iii)}{\leq} \eta_0^{N_h^k} H + \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n}) + 2 \left( \bar{B}_h^k(s, a) + 2c_b \frac{H^{7/4}\iota}{(N_h^k \vee 1)^{3/4}} + 2c_b \frac{H^2\iota}{N_h^k \vee 1} \right), \end{aligned} \quad (136)$$

where (i) follows from the initialization  $\bar{Q}_h^1(s, a) = 0$  and the trivial upper bound  $Q_h^*(s, a) \leq H$  for any policy  $\pi$ , (ii) holds owing to the fact (see (131))

$$\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (\bar{b}_h^{k^n+1} + \xi_h^{k^n}) \leq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1} + \left| \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \xi_h^{k^n} \right| \leq 2 \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1}, \quad (137)$$

and (iii) comes from (120) with the fact  $\bar{B}_h^{k^{N_h^k+1}}(s, a) = \bar{B}_h^k(s, a)$ .

**Step 2: decomposing the error in (132).** Plugging (136) into (132) and rearranging terms yield

$$\begin{aligned} &\sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) \\ &\leq \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \left[ \eta_0^{N_h^k(s,a)} H + 2\bar{B}_h^k(s, a) + \frac{4c_b H^{7/4}\iota}{(N_h^k(s, a) \vee 1)^{3/4}} + \frac{4c_b H^2\iota}{N_h^k(s, a) \vee 1} \right] \\ &\quad + \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (V_{h+1}^* - V_{h+1}^{k^n(s,a)}) \\ &\leq \underbrace{\sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \left[ \eta_0^{N_h^k(s,a)} H + \frac{4c_b H^{7/4}\iota}{(N_h^k(s, a) \vee 1)^{3/4}} + \frac{4c_b H^2\iota}{N_h^k(s, a) \vee 1} \right]}_{=: J_h^1} + 2 \underbrace{\sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \bar{B}_h^k(s, a)}_{=: J_h^2} \\ &\quad + \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (V_{h+1}^* - V_{h+1}^{k^n(s,a)}). \end{aligned} \quad (139)$$

**Step 3: controlling the last term in (139).** If we could verify the following result

$$\sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (V_{h+1}^* - V_{h+1}^{k^n(s,a)})$$

$$\leq \underbrace{\left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)) + 48\sqrt{HC^*K \log \frac{2H}{\delta}} + 28c_a H^3 C^* \sqrt{S} \iota^2}_{=: J_h^3}, \quad (140)$$

then combining this result with inequality (139) would immediately establish Lemma 5. As a result, it suffices to verify the inequality (140), which shall be accomplished as follows.

**Proof of inequality (140).** We first make the observation that the left-hand side of inequality (140) is the same as what Lemma 2 shows. Therefore, we shall establish this inequality following the same framework as in Appendix B.2. To begin with, let us recall several definitions in Appendix B.2:

$$\begin{aligned} A_h &:= \underbrace{\sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (V_{h+1}^* - V_{h+1}^{k^n(s,a)})}_{=: A_{h,k}}, \\ B_{h,k} &:= \left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)), \\ Y_{h,k} &= \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k, a_h^k} \sum_{n=1}^{N_h^k(s_h^k, a_h^k)} \eta_n^{N_h^k(s_h^k, a_h^k)} (V_{h+1}^* - V_{h+1}^{k^n(s_h^k, a_h^k)}), \\ Z_{h,k} &= \left(1 + \frac{1}{H}\right) \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_{h,s_h^k, a_h^k} (V_{h+1}^* - V_{h+1}^k), \end{aligned} \quad (141)$$

and we also remind the reader of the relation in (93) as follows

$$A_h \leq \sum_{k=1}^K B_{h,k} + \sum_{k=1}^K (Z_{h,k} - B_{h,k}) + \sum_{k=1}^K (A_{h,k} - Y_{h,k}). \quad (142)$$

Equipped with these relations, we aim to control  $\sum_{k=1}^K (Z_{h,k} - B_{h,k})$  and  $\sum_{k=1}^K (A_{h,k} - Y_{h,k})$  respectively as follows.

- We first bound  $\sum_{k=1}^K (A_{h,k} - Y_{h,k})$ , which is similar to (96) (as controlled by Lemma 11). Repeating the argument and tightening the bound from the second line of (96), we have for all  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ , with probability at least  $1 - \delta$ ,

$$\begin{aligned} \left| \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \right| &\leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) [P_{h,s,a} W_{h+1}^k(s,a)]^2 \log \frac{2H}{\delta}} + 2C_d C^* C_w \log \frac{2H}{\delta} \\ &\leq \sqrt{8C^* \log \frac{2H}{\delta} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \left[ \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n(s,a)}) \right]^2} + 4HC^* \log \frac{2H}{\delta} \\ &\stackrel{(i)}{\leq} \sqrt{8C^* \log \frac{2H}{\delta} (36HK + 3c_a^2 H^6 SC^* \iota)} + 4HC^* \log \frac{2H}{\delta} \\ &\leq 32\sqrt{HC^*K \log \frac{2H}{\delta}} + 12c_a H^3 C^* \sqrt{S} \iota^2. \end{aligned} \quad (143)$$

Here, (i) holds by virtue of the following fact

$$\sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \left[ \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n(s,a)}) \right]^2 \leq 36HK + 3c_a^2 H^6 SC^* \iota, \quad (144)$$

whose proof is postponed to Appendix C.2.1.

- Next, we turn to  $\sum_{k=1}^K (Z_{h,k} - B_{h,k})$ , which can be bounded similar to (100) (as controlled via Lemma 11). Repeating the argument and tightening the bound from the second line of (100) yield

$$\begin{aligned} \left| \sum_{k=1}^K (B_{h,k} - Z_{h,k}) \right| &\leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) [P_{h,s,a} W_{h+1}^k(s,a)]^2 \log \frac{2H}{\delta}} + 2C_d C^* C_w \log \frac{2H}{\delta} \\ &\leq 8 \sqrt{C^* \log \frac{2H}{\delta} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) [P_{h,s,a} (V_{h+1}^* - V_{h+1}^k)]^2} + 8HC^* \log \frac{2H}{\delta}. \end{aligned} \quad (145)$$

To further control (145), we have

$$\begin{aligned} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) [P_{h,s,a} (V_{h+1}^* - V_{h+1}^k)]^2 &\stackrel{(i)}{\leq} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} (V_{h+1}^* - V_{h+1}^k)^2 \\ &\stackrel{(ii)}{\leq} H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} (V_{h+1}^* - V_{h+1}^k) \\ &\stackrel{(iii)}{\leq} 2HK + c_a^2 H^6 SC^* \iota. \end{aligned} \quad (146)$$

Here, (i) holds due to the non-negativity of the variance

$$\text{Var}_{h,s,a}(V_{h+1}^* - \bar{V}_{h+1}^k) = P_{h,s,a}(V_{h+1}^* - V_{h+1}^k)^2 - (P_{h,s,a}(V_{h+1}^* - V_{h+1}^k))^2 \geq 0; \quad (147)$$

(ii) follows from the basic property  $\|V_{h+1}^* - V_{h+1}^k\|_\infty \leq H$ ; to see why (iii) holds, we refer the reader to (154), which will be proven in Appendix C.2.1 as well. Inserting (146) back into (145) yields

$$\begin{aligned} \left| \sum_{k=1}^K (B_{h,k} - Z_{h,k}) \right| &\leq 8 \sqrt{C^* \log \frac{2H}{\delta} (2KH + c_a^2 H^6 SC^* \iota)} + 8HC^* \log \frac{2H}{\delta} \\ &\leq 16 \sqrt{HC^* K \log \frac{2H}{\delta}} + 16c_a H^3 C^* \sqrt{S} \iota. \end{aligned} \quad (148)$$

Substituting the inequalities (143) and (148) into (142), and using the definitions in (141), we arrive at

$$\begin{aligned} A_h &= \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} (V_{h+1}^* - V_{h+1}^{k^n(s,a)}) \\ &\leq \left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)) + \sum_{k=1}^K (Z_{h,k} - B_{h,k}) + \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \\ &\leq \left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)) + 32 \sqrt{HC^* K \log \frac{2H}{\delta}} + 12c_a H^3 C^* \sqrt{S} \iota^2 \\ &\quad + 16 \sqrt{HC^* K \log \frac{2H}{\delta}} + 16c_a H^3 C^* \sqrt{S} \iota \\ &\leq \left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)) + 48 \sqrt{HC^* K \log \frac{2H}{\delta}} + 28c_a H^3 C^* \sqrt{S} \iota^2, \end{aligned} \quad (149)$$

which directly verifies (140) and completes the proof.

### C.2.1 Proof of inequality (144)

**Step 1: rewriting the term of interest.** We first invoke Jensen's inequality to obtain

$$\left( \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n}) \right)^2 \leq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n}) \right)^2 \leq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} (V_{h+1}^* - V_{h+1}^{k^n})^2,$$

where the first inequality follows from  $\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} = 1$  (see (26) and (25)), and the last inequality holds by the non-negativity of the variance  $\text{Var}_{h,s,a}[V_{h+1}^* - V_{h+1}^{k^n}]$ . This allows one to derive

$$\begin{aligned} & \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \left[ \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_{h,s,a} \left( V_{h+1}^* - V_{h+1}^{k^n} \right) \right]^2 \\ & \leq \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( V_{h+1}^* - V_{h+1}^{k^n} \right)^2 \\ & \stackrel{(i)}{\leq} \left( 1 + \frac{1}{H} \right) \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right)^2 + 32 \sqrt{H^4 C^* K \log \frac{2H}{\delta}} + 32 H^2 C^* \log \frac{2H}{\delta}, \end{aligned} \quad (150)$$

where (i) can be verified in a way similar to the proof of Lemma 2 in Appendix B.2. We omit the details for conciseness.

**Step 2: controlling the first term in (150).** Let us introduce the following short-hand notation

$$k_{\text{stop}} := c_a^2 H^5 S C^* \iota,$$

and decompose the term in (150) as follows

$$\begin{aligned} & \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \sum_{k=1}^K \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right)^2 \stackrel{(i)}{\leq} H \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right) \\ & = H \sum_{k=1}^{k_{\text{stop}}} \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right) + H \sum_{k=k_{\text{stop}}+1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right). \end{aligned} \quad (151)$$

Here, (i) holds since  $0 \leq V_{h+1}^*(s) - V_{h+1}^k(s) \leq H$ . The first term in (151) satisfies

$$H \sum_{k=1}^{k_{\text{stop}}} \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right) \leq H \left( c_a \sqrt{H^5 S C^* \iota k_{\text{stop}}} + c_a H^2 S C^* \iota \right) \leq c_a^2 H^6 S C^* \iota, \quad (152)$$

where the first inequality holds by applying the results of LCB-Q in (44) with  $K = k_{\text{stop}}$ . The second term in (151) can be controlled as follows:

$$\begin{aligned} & H \sum_{k=k_{\text{stop}}+1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right) \leq H K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^*(s) - V_{h+1}^{k_{\text{stop}}}(s) \right) \\ & \leq H K \frac{1}{k_{\text{stop}}} \sum_{k=1}^{k_{\text{stop}}} \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right) \\ & \leq H K \left( c_a \sqrt{\frac{H^5 S C^* \iota}{k_{\text{stop}}}} + \frac{c_a H^2 S C^* \iota}{k_{\text{stop}}} \right) \leq 2 H K, \end{aligned} \quad (153)$$

where the first and the second inequalities hold by the monotonicity property  $V_{h+1}^{k+1} \geq V_{h+1}^k$  introduced in (46), and the final inequality follows from applying (44).

Inserting the results in (152) and (153) into (151) yields

$$\sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \sum_{k=1}^K \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right)^2 \leq H \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right) \leq 2 H K + c_a^2 H^6 S C^* \iota. \quad (154)$$

**Step 3: combining the above results.** Inserting the above result (154) back into (150), we reach:

$$\begin{aligned}
& \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \left[ \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_{h,s,a} \left( V_{h+1}^* - V_{h+1}^{k^n} \right) \right]^2 \\
& \leq \left( 1 + \frac{1}{H} \right) \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left( V_{h+1}^* - V_{h+1}^k \right)^2 + 32 \sqrt{H^4 C^* K \log \frac{2H}{\delta}} + 32 H^2 C^* \log \frac{2H}{\delta} \\
& \stackrel{(i)}{\leq} 4HK + 2c_a^2 H^6 S C^* \iota + 32 \sqrt{H^4 C^* K \log \frac{2H}{\delta}} + 32 H^2 C^* \log \frac{2H}{\delta} \\
& \stackrel{(ii)}{\leq} 36HK + 3c_a^2 H^6 S C^* \iota,
\end{aligned} \tag{155}$$

where (i) holds due to (154) and  $1 + \frac{1}{H} \leq 2$ , and (ii) results from the Cauchy-Schwarz inequality.

### C.3 Proof of Lemma 6

We shall verify the three inequalities in (54) separately.

#### C.3.1 Proof of inequality (54a)

We start by rewriting the term of interest using the expression of  $J_h^1$  in (51) as

$$\begin{aligned}
& \sum_{h=1}^H \left( 1 + \frac{1}{H} \right)^{h-1} J_h^1 \\
& = \sum_{h=1}^H \left( 1 + \frac{1}{H} \right)^{h-1} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \left[ \eta_0^{N_h^k(s,a)} H + \frac{4c_b H^{7/4} \iota}{(N_h^k(s,a) \vee 1)^{3/4}} + \frac{4c_b H^2 \iota}{N_h^k(s,a) \vee 1} \right] \\
& = \underbrace{\sum_{h=1}^H \left( 1 + \frac{1}{H} \right)^{h-1} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \eta_0^{N_h^k(s,a)} H}_{=:\mathcal{J}_1^1} + \underbrace{\sum_{h=1}^H \left( 1 + \frac{1}{H} \right)^{h-1} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \frac{4c_b H^{7/4} \iota}{(N_h^k(s,a) \vee 1)^{3/4}}}_{=:\mathcal{J}_1^2} \\
& \quad + \underbrace{\sum_{h=1}^H \left( 1 + \frac{1}{H} \right)^{h-1} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \frac{4c_b H^2 \iota}{N_h^k(s,a) \vee 1}}_{=:\mathcal{J}_1^3}.
\end{aligned} \tag{156}$$

Invoking (105) and (102) yields

$$\mathcal{J}_1^1 \lesssim H^2 S C^* \iota. \tag{157}$$

In terms of  $\mathcal{J}_1^2$ , one has

$$\begin{aligned}
\mathcal{J}_1^2 & = \sum_{h=1}^H \left( 1 + \frac{1}{H} \right)^{h-1} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \frac{4c_b H^{7/4} \iota}{(N_h^k(s,a) \vee 1)^{3/4}} \\
& \stackrel{(i)}{\lesssim} H^{7/4} \iota^2 \sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \frac{1}{(k d_h^\mu(s,a))^{\frac{3}{4}}} \\
& \stackrel{(ii)}{\lesssim} H^{7/4} \iota^2 (C^*)^{\frac{3}{4}} \sum_{h=1}^H \sum_{k=1}^K \frac{1}{k^{\frac{3}{4}}} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left( d_h^{\pi^*}(s,a) \right)^{\frac{1}{4}} \\
& = H^{7/4} \iota^2 (C^*)^{\frac{3}{4}} \sum_{h=1}^H \sum_{k=1}^K \frac{1}{k^{\frac{3}{4}}} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathbb{1}(a = \pi_h^*(s)) \left( d_h^{\pi^*}(s,a) \right)^{\frac{1}{4}},
\end{aligned}$$

where (i) holds due to (102) and  $\frac{1}{N_h^k(s,a) \vee 1} \leq \frac{8\iota}{kd_h^\mu(s,a)}$  from Lemma 8, and (ii) follows from the definition of  $C^*$  in Assumption 1. A direct application of Hölder's inequality leads to

$$\begin{aligned} \mathcal{J}_1^2 &\leq H^{7/4} \iota^2 (C^*)^{\frac{3}{4}} \sum_{h=1}^H \sum_{k=1}^K \frac{1}{k^{\frac{3}{4}}} \left( \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathbb{1}(a = \pi_h^*(s)) \right)^{3/4} \left( \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \right)^{1/4} \\ &\stackrel{(iii)}{\leq} H^{7/4} \iota^2 (SC^*)^{\frac{3}{4}} \sum_{h=1}^H \sum_{k=1}^K \frac{1}{k^{\frac{3}{4}}} \lesssim H^{2.75} (SC^*)^{\frac{3}{4}} K^{\frac{1}{4}} \iota^2, \end{aligned} \quad (158)$$

where (iii) follows since  $\pi^*$  is assumed to be a deterministic policy.

Similarly, we can derive an upper bound on  $\mathcal{J}_1^3$  as follows:

$$\begin{aligned} \mathcal{J}_1^3 &= \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \frac{4c_b H^2 \iota}{N_h^k(s,a) \vee 1} \\ &\stackrel{(i)}{\lesssim} H^2 \iota^2 \sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{d_h^{\pi^*}(s,a)}{kd_h^\mu(s,a)} \lesssim H^3 SC^* \iota^3, \end{aligned} \quad (159)$$

where (i) follows from the result in (102) and the fact  $\frac{1}{N_h^k(s,a) \vee 1} \leq \frac{8\iota}{kd_h^\mu(s,a)}$  (cf. Lemma 8), and the last relation results from the definition of  $C^*$  (cf. Assumption 1) and the assumption that  $\pi^*$  is a deterministic policy.

Putting the preceding results (157), (158) and (159) together, we conclude that

$$\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} J_h^1 \lesssim H^{2.75} (SC^*)^{\frac{3}{4}} K^{\frac{1}{4}} \iota^2 + H^3 SC^* \iota^3. \quad (160)$$

### C.3.2 Proof of inequality (54b)

Making use of the definition of  $\bar{B}_h^k(s,a)$  (cf. (14)) in the expression of  $J_h^2$  (cf. (51)), we obtain

$$\begin{aligned} \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} J_h^2 &= 2 \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \bar{B}_h^k(s,a) \\ &= 2 \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} c_b \sqrt{H\iota} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{\sigma_h^{\text{adv},k}(s,a) - (\mu_h^{\text{adv},k}(s,a))^2}{N_h^k(s,a) \vee 1}} \\ &\quad + 2 \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} c_b \sqrt{\iota} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{\sigma_h^{\text{ref},k}(s,a) - (\mu_h^{\text{ref},k}(s,a))^2}{N_h^k(s,a) \vee 1}} \\ &\lesssim \underbrace{\sqrt{H\iota} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{\sigma_h^{\text{adv},k}(s,a) - (\mu_h^{\text{adv},k}(s,a))^2}{N_h^k(s,a) \vee 1}}}_{=: \mathcal{J}_2^1} \\ &\quad + \underbrace{\sqrt{\iota} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{\sigma_h^{\text{ref},k}(s,a) - (\mu_h^{\text{ref},k}(s,a))^2}{N_h^k(s,a) \vee 1}}}_{=: \mathcal{J}_2^2}, \end{aligned} \quad (161)$$

where the last inequality follows from (102). In the following, we shall look at the two terms in (161) separately.

**Step 1: controlling  $\mathcal{J}_2^1$ .** Recalling the expressions of  $\sigma_h^{\text{adv},k}(s,a) = \sigma_h^{\text{adv},k^{N_h^k+1}}(s,a)$  in (112), we observe that the main part of  $\mathcal{J}_2^1$  in (161) satisfies

$$\begin{aligned}
& \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{\sigma_h^{\text{adv},k}(s,a) - \left(\mu_h^{\text{adv},k}(s,a)\right)^2}{N_h^k(s,a) \vee 1}} \leq \sqrt{\iota} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{k=1}^K \sqrt{d_h^{\pi^*}(s,a) \frac{d_h^{\pi^*}(s,a) \cdot \sigma_h^{\text{adv},k}(s,a)}{k d_h^\mu(s,a)}} \\
& = \sqrt{\iota} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{k=1}^K \sqrt{d_h^{\pi^*}(s,a) \frac{d_h^{\pi^*}(s,a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_h^{k^n} \left(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\right)^2}{k d_h^\mu(s,a)}} \\
& \stackrel{(i)}{\leq} \sqrt{C^* \iota} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{k=1}^K \sqrt{\frac{1}{k} \mathbf{1}(a = \pi_h^*(s)) \cdot d_h^{\pi^*}(s,a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_h^{k^n} \left(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\right)^2} \\
& \stackrel{(ii)}{\leq} \sqrt{C^* \iota} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_h^{k^n} \left(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\right)^2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mathbf{1}(a = \pi_h^*(s)) \frac{1}{k}} \\
& \lesssim \sqrt{H S C^* \iota^2} \sqrt{\sum_{k=1}^K \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_h^{k^n} \left(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\right)^2}, \tag{162}
\end{aligned}$$

where the first inequality is due to the fact  $\frac{1}{N_h^k(s,a) \vee 1} \leq \frac{8\iota}{k d_h^\mu(s,a)}$  from Lemma 8, (i) follows from the definition of  $C^*$  in Assumption 1 and (35), and (ii) follows from the Cauchy-Schwarz inequality. To continue, we claim the following bound holds, which will be proven in Appendix C.3.4:

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_h^{k^n} \left(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\right)^2 \\
& \lesssim H^2 \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) + K + H^5 \sqrt{S C^*} \iota^2. \tag{163}
\end{aligned}$$

Combining the above inequality with (162), we arrive at

$$\begin{aligned}
\mathcal{J}_2^1 & \lesssim \sqrt{H^2 S C^* \iota^3} \sqrt{H^2 \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) + K + H^5 \sqrt{S C^*} \iota^2} \\
& \lesssim \sqrt{H^4 S C^* \iota^3 \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) + \sqrt{H^2 S C^* K \iota^3} + H^{3.5} S C^* \iota^{2.5}}. \tag{164}
\end{aligned}$$

**Step 2: controlling  $\mathcal{J}_2^2$ .** Recalling the expressions of  $\mu_h^{\text{ref},k+1}(s,a) = \mu_h^{\text{ref},k^{N_h^k+1}}(s,a)$  and  $\sigma_h^{\text{ref},k+1}(s,a) = \sigma_h^{\text{ref},k^{N_h^k+1}}(s,a)$  in (113) to  $\mathcal{J}_2^2$  in (161), we can deduce that

$$\begin{aligned}
\mathcal{J}_2^2 & = \sqrt{\iota} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{\sigma_h^{\text{ref},k}(s,a) - \left(\mu_h^{\text{ref},k}(s,a)\right)^2}{N_h^k(s,a) \vee 1}} \\
& \leq \sqrt{\iota} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s,a) \vee 1} \underbrace{\left( \frac{\sum_{n=1}^{N_h^k(s,a)} \left(\bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n})\right)^2}{N_h^k(s,a) \vee 1} - \left( \frac{\sum_{n=1}^{N_h^k(s,a)} \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n})}{N_h^k(s,a) \vee 1} \right)^2 \right)}_{=: F_{h,k}}}. \tag{165}
\end{aligned}$$



We further decompose and bound  $F_{h,k}$  as follows:

$$\begin{aligned}
F_{h,k} &\stackrel{(i)}{\leq} \sqrt{\frac{\sum_{n=1}^{N_h^k(s,a)} (V_{h+1}^\star(s_{h+1}^{k^n}))^2}{N_h^k(s,a) \vee 1} - \left(\frac{\sum_{n=1}^{N_h^k(s,a)} \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n})}{N_h^k(s,a) \vee 1}\right)^2} \\
&= \sqrt{\frac{\sum_{n=1}^{N_h^k(s,a)} (V_{h+1}^\star(s_{h+1}^{k^n}))^2}{N_h^k(s,a) \vee 1} - \left(\frac{\sum_{n=1}^{N_h^k(s,a)} V_{h+1}^\star(s_{h+1}^{k^n})}{N_h^k(s,a) \vee 1}\right)^2} + \left(\frac{\sum_{n=1}^{N_h^k(s,a)} V_{h+1}^\star(s_{h+1}^{k^n})}{N_h^k(s,a) \vee 1}\right)^2 - \left(\frac{\sum_{n=1}^{N_h^k(s,a)} \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n})}{N_h^k(s,a) \vee 1}\right)^2 \\
&\stackrel{(ii)}{\leq} \underbrace{\sqrt{\frac{\sum_{n=1}^{N_h^k(s,a)} (V_{h+1}^\star(s_{h+1}^{k^n}))^2}{N_h^k(s,a) \vee 1} - \left(\frac{\sum_{n=1}^{N_h^k(s,a)} V_{h+1}^\star(s_{h+1}^{k^n})}{N_h^k(s,a) \vee 1}\right)^2}}_{G_{h,k}} + \underbrace{\sqrt{\frac{\sum_{n=1}^{N_h^k(s,a)} 2H (V_{h+1}^\star(s_{h+1}^{k^n}) - \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n}))}{N_h^k(s,a) \vee 1}}}_{=:L_{h,k}}, \tag{166}
\end{aligned}$$

where (i) follows from the fact that for some  $k' \in [K]$ ,  $\bar{V}_{h+1}^{\text{next},k^n} = V_{h+1}^{k'} \leq V_{h+1}^\star$  (see the update rule of  $\bar{V}^{\text{next}}$  in line 27 and the fact in (47)), and (ii) holds due to the fact that

$$\left(\frac{\sum_{n=1}^{N_h^k(s,a)} V_{h+1}^\star(s_{h+1}^{k^n})}{N_h^k(s,a) \vee 1}\right)^2 - \left(\frac{\sum_{n=1}^{N_h^k(s,a)} \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n})}{N_h^k(s,a) \vee 1}\right)^2 \leq 2H \frac{\sum_{n=1}^{N_h^k(s,a)} (V_{h+1}^\star(s_{h+1}^{k^n}) - \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n}))}{N_h^k(s,a) \vee 1}.$$

Inserting (166) back into (165), we arrive at

$$\begin{aligned}
\mathcal{J}_2^2 &\leq \sqrt{\iota} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^\star}(s,a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s,a) \vee 1}} (G_{h,k} + L_{h,k}) \\
&\stackrel{(i)}{\lesssim} \sqrt{\iota} \left( \sqrt{H^3 SC^\star K \iota^4} + H^4 SC^\star \iota^3 + \sqrt{H^3 SC^\star K \iota^2} + H^{2.5} SC^\star \iota^3 \right) \lesssim \sqrt{H^3 SC^\star K \iota^5} + H^4 SC^\star \iota^4, \tag{167}
\end{aligned}$$

where (i) follows from the following facts

$$\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^\star}(s,a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s,a) \vee 1}} L_{h,k} \lesssim \sqrt{H^3 SC^\star K \iota^4} + H^4 SC^\star \iota^3, \tag{168}$$

$$\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^\star}(s,a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s,a) \vee 1}} G_{h,k} \lesssim \sqrt{H^3 SC^\star K \iota^2} + H^{2.5} SC^\star \iota^3. \tag{169}$$

We postpone the proofs of (168) and (169) to Appendix C.3.5 and Appendix C.3.6, respectively.

**Putting the bounds together.** Substitute (164) and (167) back into (161) to yield

$$\begin{aligned}
\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} J_h^2 &\lesssim \sqrt{H^4 SC^\star \iota^3 \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^\star}(s) (V_h^\star(s) - V_h^k(s)) + \sqrt{H^2 SC^\star K \iota^3} + H^{3.5} SC^\star \iota^{2.5}} \\
&\quad + \sqrt{H^3 SC^\star K \iota^5} + H^4 SC^\star \iota^4 \\
&\lesssim \sqrt{H^4 SC^\star \iota^3 \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^\star}(s) (V_h^\star(s) - V_h^k(s)) + \sqrt{H^3 SC^\star K \iota^5} + H^4 SC^\star \iota^4}.
\end{aligned}$$

### C.3.3 Proof of inequality (54c)

Invoking inequality (102) directly leads to

$$\sum_{h=1}^H \left(1 + \frac{1}{H}\right)^{h-1} \left(48\sqrt{HC^*K \log \frac{2H}{\delta}} + 28c_a H^3 C^* \sqrt{S} \iota^2\right) \lesssim \sqrt{H^3 C^* K \log \frac{2H}{\delta}} + H^4 C^* \sqrt{S} \iota^2$$

as claimed.

### C.3.4 Proof of inequality (163)

We shall control the term in (163) in a way similar to the proof of Lemma 2 in Appendix B.2.

**Step 1: decomposing the terms of interest.** Akin to Appendix B.2, let us introduce the terms of interest and definitions as follows:

$$\begin{aligned} A_h &:= \underbrace{\sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_h^{k^n} \left(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\right)^2}_{=: A_{h,k}}, \\ B_{h,k} &:= \left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left(V_{h+1}^k(s) - \bar{V}_{h+1}^k(s)\right)^2, \\ Y_{h,k} &= \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} \sum_{n=1}^{N_h^k(s_h^k, a_h^k)} \eta_n^{N_h^k(s_h^k, a_h^k)} P_h^{k^n} \left(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\right)^2, \\ Z_{h,k} &= \left(1 + \frac{1}{H}\right) \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_h^k \left(V_{h+1}^k - \bar{V}_{h+1}^k\right)^2. \end{aligned} \quad (170)$$

With these definitions in place, we directly adapt the argument in (93) to arrive at

$$A_h \leq \sum_{k=1}^K B_{h,k} + \sum_{k=1}^K (Z_{h,k} - B_{h,k}) + \sum_{k=1}^K (A_{h,k} - Y_{h,k}). \quad (171)$$

As a consequence, it remains to control  $\sum_{k=1}^K (Z_{h,k} - B_{h,k})$  and  $\sum_{k=1}^K (A_{h,k} - Y_{h,k})$  separately.

**Step 2: controlling  $\sum_{k=1}^K (A_{h,k} - Y_{h,k})$ .** To control  $\sum_{k=1}^K (A_{h,k} - Y_{h,k})$ , we resort to Lemma 11 by setting

$$W_{h+1}^k(s, a) := \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} \left(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\right)^2, \quad C_d := 1, \quad (172)$$

which satisfies

$$\|W_{h+1}^k(s, a)\|_\infty \leq 4H^2 =: C_w.$$

Applying Lemma 11 with (172) yields that: with probability at least  $1 - \delta$ ,

$$\begin{aligned} \left| \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \right| &= \left| \sum_{k=1}^K \bar{X}_{h,k} \right| \\ &\leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} [W_{h+1}^k(s,a)]^2 \log \frac{2H}{\delta} + 2C_d C^* C_w \log \frac{2H}{\delta}} \end{aligned}$$

$$\lesssim \sqrt{C^\star \log \frac{2H}{\delta} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^\star}(s,a) P_{h,s,a} \left[ \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} \left( V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n} \right)^2 \right]^2} + C^\star H^2 \log \frac{2H}{\delta}. \quad (173)$$

To further control the first term in (173), it follows from Jensen's inequality that

$$P_{h,s,a} \left[ \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n} \right)^2 \right]^2 \leq P_{h,s,a} \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n} \right)^4, \quad (174)$$

which yields

$$\begin{aligned} & \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^\star}(s,a) P_{h,s,a} \left[ \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} \left( V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n} \right)^2 \right]^2 \\ & \leq \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^\star}(s,a) P_{h,s,a} \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n} \right)^4 \\ & \leq \left( 1 + \frac{1}{H} \right) \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^\star}(s) \left( V_{h+1}^k(s) - \bar{V}_{h+1}^k(s) \right)^4 + 32 \sqrt{H^8 C^\star K \log \frac{2H}{\delta}} + 32 H^4 C^\star \log \frac{2H}{\delta}. \end{aligned} \quad (175)$$

This can be verified similar to the proof for Lemma 2 in Appendix B.2. We omit the details for conciseness. To continue, it follows that

$$\begin{aligned} & \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^\star}(s) \left( V_{h+1}^k(s) - \bar{V}_{h+1}^k(s) \right)^4 \\ & \stackrel{(i)}{\leq} \sum_{m=1}^M \sum_{t=1}^{L_m} \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^\star}(s) \left( V_{h+1}^\star(s) - \bar{V}_{h+1}^{(m,t)}(s) \right)^4 \\ & \stackrel{(ii)}{=} \sum_{m=1}^M \sum_{t=1}^{L_m} \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^\star}(s) \left( V_{h+1}^\star(s) - V_{h+1}^{((m-1) \vee 1, 1)}(s) \right)^4 \\ & \stackrel{(iii)}{=} \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^\star}(s) \sum_{m=1}^M 2^m \left( V_{h+1}^\star(s) - V_{h+1}^{((m-1) \vee 1, 1)}(s) \right)^4 \\ & = 4 \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^\star}(s) \sum_{m=2}^{M-2} 2^{m-2} \left( V_{h+1}^\star(s) - V_{h+1}^{((m-1) \vee 1, 1)}(s) \right)^4 \\ & = 4 \sum_{m=2}^0 2^{m-2} \left( V_{h+1}^\star(s) - V_{h+1}^{(1,1)}(s) \right)^4 + 4 \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^\star}(s) \sum_{m=2}^{M-2} 2^{m-2} \left( V_{h+1}^\star(s) - V_{h+1}^{(m-1,1)}(s) \right)^4. \end{aligned}$$

Here, (i) holds by using the pessimistic property  $V^\star \geq V^k \geq \bar{V}^k$  for all  $k \in [K]$  (see (47)) and by regrouping the summands; (ii) follows from the fact (see updating rules in line 26 and line 27) that for any  $(m, s, h) \in [M] \times \mathcal{S} \times [H+1]$ ,

$$\bar{V}_h^{(m,t)}(s) = V_h^{((m-1) \vee 1, 1)}(s), \quad t = 1, 2, \dots, L_m; \quad (176)$$

and (iii) results from the choice of the parameter  $L_m = 2^m$ . In addition, we can further control

$$\sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^\star}(s) \left( V_{h+1}^k(s) - \bar{V}_{h+1}^k(s) \right)^4 \stackrel{(iv)}{\leq} 8H^4 + 4 \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^\star}(s) \sum_{m=1}^{M-2} \sum_{t=1}^{L_m} \left( V_{h+1}^\star(s) - V_{h+1}^{(m+1,1)}(s) \right)^4$$

$$\begin{aligned}
&\stackrel{(v)}{\leq} 8H^4 + 4 \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \sum_{m=1}^{M-2} \sum_{t=1}^{L_m} \left( V_{h+1}^*(s) - V_{h+1}^{(m,t)}(s) \right)^4 \\
&\leq 8H^4 + 4 \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \sum_{k=1}^K \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right)^4 \tag{177}
\end{aligned}$$

$$\begin{aligned}
&\leq 8H^4 + 4H^3 \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \sum_{k=1}^K \left( V_{h+1}^*(s) - V_{h+1}^k(s) \right) \\
&\stackrel{(vi)}{\lesssim} H^3 K + H^8 S C^* \iota. \tag{178}
\end{aligned}$$

Here, (iv) follows from the fact  $0 \leq V_{h+1}^*(s) - V_{h+1}^{(1,1)}(s) \leq H - 0 = H$ ; (v) holds since  $V_{h+1}^* \geq V_{h+1}^{(m+1,1)} = V_{h+1}^{(m,L_m)} \geq V_{h+1}^{(m,t)}$  for all  $t \in [L_m]$  (using the monotonic increasing property of  $V_{h+1}$  introduced in (46)); and (vi) follows from (154). Putting (178) and (175) together with (173), we arrive at

$$\begin{aligned}
\left| \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \right| &\lesssim \sqrt{C^* \log \frac{2H}{\delta} \left( H^3 K + H^8 S C^* \iota + \sqrt{H^8 C^* K \log \frac{2H}{\delta}} + H^4 C^* \log \frac{2H}{\delta} \right)} + C^* H^2 \log \frac{2H}{\delta} \\
&\lesssim \sqrt{H^3 C^* K \iota} + H^4 \sqrt{S C^* \iota^2}. \tag{179}
\end{aligned}$$

**Step 3: controlling  $\sum_{k=1}^K (Z_{h,k} - B_{h,k})$ .** Similarly, we also invoke Lemma 11 to control  $\sum_{k=1}^K (Z_{h,k} - B_{h,k})$ . Let's set

$$W_{h+1}^k(s, a) := \left( V_{h+1}^k(s, a) - \bar{V}_{h+1}^k \right)^2, \quad C_d := \left( 1 + \frac{1}{H} \right) \leq 2, \tag{180}$$

which satisfies

$$\|W_{h+1}^k(s, a)\|_{\infty} \leq 4H^2 =: C_w.$$

Applying Lemma 11 with (180) yields that: with probability at least  $1 - \delta$ ,

$$\begin{aligned}
\left| \sum_{k=1}^K (B_{h,k} - Z_{h,k}) \right| &= \left| \sum_{k=1}^K \bar{X}_{h,k} \right| \\
&\leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} [W_{h+1}^k(s, a)]^2 \log \frac{2H}{\delta}} + 2C_d C^* C_w \log \frac{2H}{\delta} \\
&\lesssim \sqrt{C^* \log \frac{2H}{\delta} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) P_{h,s,a} [V_{h+1}^k - \bar{V}_{h+1}^k]^4} + C^* H^2 \log \frac{2H}{\delta} \\
&\stackrel{(i)}{\lesssim} \sqrt{C^* \log \frac{2H}{\delta} (H^3 K + H^8 S C^* \iota)} + C^* H^2 \log \frac{2H}{\delta} \lesssim \sqrt{H^3 C^* K \iota} + H^4 \sqrt{S C^* \iota^2}, \tag{181}
\end{aligned}$$

where (i) follows from (177) and (178).

**Step 4: combining the results.** Inserting (181) and (179) back into (171), we can conclude that

$$\begin{aligned}
&\sum_{k=1}^K \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} P_h^{k^n} \left( V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n} \right)^2 = \sum_{h=1}^H A_h \\
&\leq \sum_{h=1}^H \sum_{k=1}^K B_{h,k} + \sum_{h=1}^H \sum_{k=1}^K (Z_{h,k} - B_{h,k}) + \sum_{h=1}^H \sum_{k=1}^K (A_{h,k} - Y_{h,k})
\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{h=1}^H \sum_{k=1}^K \left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left(V_{h+1}^k(s) - \bar{V}_{h+1}^k(s)\right)^2 + \sum_{h=1}^H \left| \sum_{k=1}^K (Z_{h,k} - B_{h,k}) \right| + \sum_{h=1}^H \left| \sum_{k=1}^K (A_{h,k} - Y_{h,k}) \right| \\
&\leq H \sum_{h=1}^H \sum_{k=1}^K \left(1 + \frac{1}{H}\right) \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) \left(V_{h+1}^k(s) - \bar{V}_{h+1}^k(s)\right) + \sqrt{H^5 C^* K \ell} + H^5 \sqrt{SC^*} \ell^2 \\
&\stackrel{(i)}{\lesssim} H \sum_{h=1}^H \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)) + K + H^5 \sqrt{SC^*} \ell^2 \\
&\lesssim H^2 \max_{h \in [H]} \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_h^{\pi^*}(s) (V_h^*(s) - V_h^k(s)) + K + H^5 \sqrt{SC^*} \ell^2,
\end{aligned} \tag{182}$$

where (i) follows from the same routine to obtain (177) and the Cauchy-Schwarz inequality.

### C.3.5 Proof of inequality (168)

**Step 1: decomposing the error in (168).** The term in (168) obeys

$$\begin{aligned}
&\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s,a) \vee 1}} L_{h,k} \\
&= \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s,a) \vee 1}} \sqrt{\frac{\sum_{n=1}^{N_h^k(s,a)} 2H \left( V_{h+1}^*(s_{h+1}^{k^n}) - \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n}) \right)}{N_h^k(s,a) \vee 1}} \\
&\stackrel{(i)}{\lesssim} \sqrt{H\ell} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{k=1}^K \sqrt{\frac{d_h^{\pi^*}(s,a)}{kd_h^\mu(s,a)}} \sqrt{\frac{d_h^{\pi^*}(s,a) \ell \sum_{n=1}^{N_h^k(s,a)} \left( V_{h+1}^*(s_{h+1}^{k^n}) - \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n}) \right)}{kd_h^\mu(s,a)}} \\
&\stackrel{(ii)}{\lesssim} \sqrt{HC^* \ell^2} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{k=1}^K \sqrt{\frac{\mathbb{1}(a = \pi^*(s))}{k}} \sqrt{\frac{d_h^{\pi^*}(s,a) \sum_{n=1}^{N_h^k(s,a)} \left( V_{h+1}^*(s_{h+1}^{k^n}) - \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n}) \right)}{kd_h^\mu(s,a)}} \\
&\stackrel{(iii)}{\lesssim} \sqrt{HC^* \ell^2} \sqrt{\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{k=1}^K \frac{d_h^{\pi^*}(s,a) \sum_{n=1}^{N_h^k(s,a)} \left( V_{h+1}^*(s_{h+1}^{k^n}) - \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n}) \right)}{kd_h^\mu(s,a)}} \sqrt{\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{k=1}^K \frac{\mathbb{1}(a = \pi^*(s))}{k}} \\
&\lesssim \sqrt{H^2 SC^* \ell^3} \sqrt{\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{d_h^{\pi^*}(s,a)}{d_h^\mu(s,a)} \sum_{k=1}^K \frac{1}{k} \sum_{n=1}^{N_h^k(s,a)} \left( V_{h+1}^*(s_{h+1}^{k^n(s,a)}) - \bar{V}_{h+1}^{\text{next},k^n}(s_{h+1}^{k^n(s,a)}) \right)} \\
&\stackrel{(iv)}{=} \sqrt{H^2 SC^* \ell^3} \sqrt{\sum_{h=1}^H \sum_{k=1}^K \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_h^k \sum_{k'=k}^K \frac{1}{k'} (V_{h+1}^* - \bar{V}_{h+1}^{\text{next},k})} \\
&\lesssim \sqrt{H^2 SC^* \ell^4} \sqrt{\sum_{h=1}^H \sum_{k=1}^K \frac{d_h^{\pi^*}(s_h^k, a_h^k) P_h^k}{d_h^\mu(s_h^k, a_h^k)} (V_{h+1}^* - \bar{V}_{h+1}^{\text{next},k})}.
\end{aligned} \tag{183}$$

Here, (i) follows from the fact  $\frac{1}{N_h^k(s,a) \vee 1} \leq \frac{8\ell}{kd_h^\mu(s,a)}$  (cf. Lemma 8); (ii) follows from the definition of  $C^*$  in Assumption 1; (iii) invokes the Cauchy-Schwarz inequality; (iv) can be obtained by regrouping the terms (the terms involving  $(V_{h+1}^* - \bar{V}_{h+1}^{\text{next},k})$  associated with index  $k$  will only be added during episodes  $k' = k, k+1, \dots, K$ ).

With this upper bound in hand, we further decompose

$$\begin{aligned}
& \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s,a) \vee 1}} L_{h,k} \lesssim \sqrt{H^2 SC^* \iota^4} \sqrt{\sum_{h=1}^H \sum_{k=1}^K \frac{d_h^{\pi^*}(s_h^k, a_h^k) P_h^k}{d_h^\mu(s_h^k, a_h^k)} (V_{h+1}^* - \bar{V}_{h+1}^{\text{next},k})} \\
& \stackrel{(i)}{\lesssim} \sqrt{H^2 SC^* \iota^4} \sqrt{\sum_{h=1}^H \sum_{k=1}^K \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_h^k (V_{h+1}^* - \bar{V}_{h+1}^k)} \\
& \stackrel{(ii)}{\lesssim} \sqrt{H^2 SC^* \iota^4} \sqrt{\sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} (V_{h+1}^* - \bar{V}_{h+1}^k)} \\
& \quad + \sqrt{H^2 SC^* \iota^4} \sqrt{\sum_{h=1}^H \sum_{k=1}^K \left( \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} - \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_h^k \right) (V_{h+1}^* - \bar{V}_{h+1}^k)}. \tag{184}
\end{aligned}$$

Here (i) holds due to the following observation: denoting by  $m$  the index of the epoch in which episode  $k$  occurs, we have

$$\bar{V}_{h+1}^{\text{next},k} = V_{h+1}^{(m,1)} \geq V_{h+1}^{((m-1 \vee 1),1)} = \bar{V}_{h+1}^k, \tag{185}$$

which invokes the monotonicity of  $V_{h+1}^k$  in (46). In addition, (ii) arises from the Cauchy-Schwarz inequality.

**Step 2: controlling the first term in (184).** The first term in (184) satisfies

$$\begin{aligned}
& \sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} (V_{h+1}^* - \bar{V}_{h+1}^k) = \sum_{h=1}^H \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \langle P_h(\cdot | s, a), V_{h+1}^* - \bar{V}_{h+1}^k \rangle \\
& \stackrel{(i)}{=} \sum_{h=1}^H \sum_{k=1}^K \sum_{s' \in \mathcal{S}} d_{h+1}^{\pi^*}(s') (V_{h+1}^*(s') - \bar{V}_{h+1}^k(s')) \\
& \stackrel{(ii)}{\lesssim} H^2 + \sum_{h=1}^H \sum_{k=1}^K \sum_{s \in \mathcal{S}} d_{h+1}^{\pi^*}(s) (V_{h+1}^*(s) - V_{h+1}^k(s)) \\
& \stackrel{(iii)}{\lesssim} HK + H^6 SC^* \iota, \tag{186}
\end{aligned}$$

where (i) holds due to the fact  $\sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_h(\cdot | s, a) = d_{h+1}^{\pi^*}(\cdot)$ , (ii) comes from the same argument employed to establish (177), and (iii) follows from (154).

**Step 3: controlling the second term in (184).** We shall invoke Lemma 11 for this purpose. To proceed, let

$$W_{h+1}^k(s, a) := V_{h+1}^* - \bar{V}_{h+1}^k, \quad C_d =: 1, \tag{187}$$

which satisfies

$$\|W_{h+1}^k(s, a)\|_\infty \leq H =: C_w.$$

Applying Lemma 11 with (187) yields, for all  $h \in [H]$ , with probability at least  $1 - \delta$

$$\left| \sum_{k=1}^K \left( \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} - \frac{d_h^{\pi^*}(s_h^k, a_h^k)}{d_h^\mu(s_h^k, a_h^k)} P_h^k \right) (V_{h+1}^* - \bar{V}_{h+1}^k) \right| = \left| \sum_{k=1}^K \bar{X}_{h,k} \right|$$

$$\begin{aligned}
&\leq \sqrt{\sum_{k=1}^K 8C_d^2 C^* \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} [W_{h+1}^k(s,a)]^2 \log \frac{2H}{\delta}} + 2C_d C^* C_w \log \frac{2H}{\delta} \\
&\lesssim \sqrt{C^* \log \frac{2H}{\delta} \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} (V_{h+1}^* - \bar{V}_{h+1}^k)^2} + HC^* \log \frac{2H}{\delta} \\
&\stackrel{(i)}{\lesssim} \sqrt{C^* \log \frac{2H}{\delta} \left( H^2 + \sum_{k=1}^K \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} (V_{h+1}^* - V_{h+1}^k)^2 \right)} + HC^* \log \frac{2H}{\delta} \\
&\stackrel{(ii)}{\lesssim} \sqrt{C^* \log \frac{2H}{\delta} (HK + H^6 SC^* \iota)} + HC^* \log \frac{2H}{\delta} \\
&\lesssim \sqrt{HC^* K \iota} + H^3 \sqrt{SC^* \iota}.
\end{aligned} \tag{188}$$

Here (i) follows from the same routine to arrive at (177), and (ii) comes from (154). As a result, the second term in (184) satisfies, with probability at least  $1 - \delta$ ,

$$\begin{aligned}
&\left| \sum_{h=1}^H \sum_{k=1}^K \left( \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} - \frac{d_h^{\pi^*}(s_h^k, a_h^k) P_h^k}{d_h^\mu(s_h^k, a_h^k)} \right) (V_{h+1}^* - \bar{V}_{h+1}^k) \right| \\
&\leq \sum_{h=1}^H \left| \sum_{k=1}^K \left( \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) P_{h,s,a} - \frac{d_h^{\pi^*}(s_h^k, a_h^k) P_h^k}{d_h^\mu(s_h^k, a_h^k)} \right) (V_{h+1}^* - \bar{V}_{h+1}^k) \right| \lesssim \sqrt{H^3 C^* K \iota} + H^4 \sqrt{SC^* \iota}.
\end{aligned} \tag{189}$$

**Step 4: combining the results.** Finally, inserting (186) and (189) into (184), we arrive at

$$\begin{aligned}
&\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s,a)}} L_{h,k} \\
&\lesssim \sqrt{H^2 SC^* \iota^4} \sqrt{HK + H^6 SC^* \iota} + \sqrt{H^2 SC^* \iota^4} \sqrt{\sqrt{H^3 C^* K \iota} + H^4 \sqrt{SC^* \iota}} \\
&\lesssim \sqrt{H^3 SC^* K \iota^4} + H^4 SC^* \iota^3 + \sqrt{H^2 SC^* \iota^4} \sqrt{HK + H^4 \sqrt{SC^* \iota}} \lesssim \sqrt{H^3 SC^* K \iota^4} + H^4 SC^* \iota^3,
\end{aligned} \tag{190}$$

where the last two inequalities follow from the Cauchy-Schwarz inequality.

### C.3.6 Proof of inequality (169)

Recall the expression of  $G_{h,k}$  in (166) as

$$\begin{aligned}
G_{h,k}^2 &= \frac{\sum_{n=1}^{N_h^k(s,a)} (V_{h+1}^*(s_{h+1}^{k,n}))^2}{N_h^k(s,a) \vee 1} - \left( \frac{\sum_{n=1}^{N_h^k(s,a)} V_{h+1}^*(s_{h+1}^{k,n})}{N_h^k(s,a) \vee 1} \right)^2 \\
&= \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k,n} (V_{h+1}^*)^2}{N_h^k(s,a) \vee 1} - \left( \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k,n} V_{h+1}^*}{N_h^k(s,a) \vee 1} \right)^2.
\end{aligned} \tag{191}$$

To continue, we make the following observation

$$\begin{aligned}
G_{h,k} &\leq \left\{ |G_{h,k}^2 - \text{Var}_{h,s,a}(V_{h+1}^*)| + \text{Var}_{h,s,a}(V_{h+1}^*) \right\}^{1/2} \\
&\leq |G_{h,k}^2 - \text{Var}_{h,s,a}(V_{h+1}^*)|^{1/2} + \sqrt{\text{Var}_{h,s,a}(V_{h+1}^*)}
\end{aligned} \tag{192}$$

due to the elementary inequality  $\sqrt{a^2 + b^2} \leq a + b$  for any  $a, b \geq 0$ . Here, we remind the reader that  $\text{Var}_{h,s,a}(V_{h+1}^*) = P_{h,s,a}(V_{h+1}^*)^2 - (P_{h,s,a} V_{h+1}^*)^2$  (cf. (109)). This allows us to rewrite

$$\begin{aligned} & \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s,a) \vee 1}} G_{h,k} \\ & \leq \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{|G_{h,k}^2 - \text{Var}_{h,s,a}(V_{h+1}^*)|}{N_h^k(s,a) \vee 1}} + \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s,a) \sum_{k=1}^K \sqrt{\frac{\text{Var}_{h,s,a}(V_{h+1}^*)}{N_h^k(s,a) \vee 1}}, \end{aligned} \quad (193)$$

leaving us with two terms to cope with.

**Step 1: controlling the first term of (193).** By definition, we have

$$\begin{aligned} |G_{h,k}^2 - \text{Var}_{h,s,a}(V_{h+1}^*)| &= \left| \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} (V_{h+1}^*)^2}{N_h^k(s,a) \vee 1} - \left( \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} V_{h+1}^*}{N_h^k(s,a) \vee 1} \right)^2 - P_{h,s,a}(V_{h+1}^*)^2 + (P_{h,s,a} V_{h+1}^*)^2 \right| \\ &\leq \left| \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} (V_{h+1}^*)^2}{N_h^k(s,a) \vee 1} - P_{h,s,a}(V_{h+1}^*)^2 \right| + \left| \left( \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} V_{h+1}^*}{N_h^k(s,a) \vee 1} \right)^2 - (P_{h,s,a} V_{h+1}^*)^2 \right| \\ &\leq \left| \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} (V_{h+1}^*)^2}{N_h^k(s,a) \vee 1} - P_{h,s,a}(V_{h+1}^*)^2 \right| + 2H \left| \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} V_{h+1}^*}{N_h^k(s,a) \vee 1} - P_{h,s,a} V_{h+1}^* \right|, \end{aligned} \quad (194)$$

where the last inequality holds due to

$$\begin{aligned} \left| \left( \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} V_{h+1}^*}{N_h^k(s,a) \vee 1} \right)^2 - (P_{h,s,a} V_{h+1}^*)^2 \right| &= \left| \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} V_{h+1}^*}{N_h^k(s,a) \vee 1} - P_{h,s,a} V_{h+1}^* \right| \cdot \left| \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} V_{h+1}^*}{N_h^k(s,a) \vee 1} + P_{h,s,a} V_{h+1}^* \right| \\ &\leq 2H \left| \frac{\sum_{n=1}^{N_h^k(s,a)} P_h^{k^n} V_{h+1}^*}{N_h^k(s,a) \vee 1} - P_{h,s,a} V_{h+1}^* \right|. \end{aligned}$$

We now control the two terms in (194) separately by invoking Lemma 9. For the first term in (194), let us set

$$W_{h+1}^i := (V_{h+1}^*)^2, \quad \text{and} \quad u_h^i(s,a,N) := \frac{1}{N \vee 1} := C_u, \quad (195)$$

which indicates that

$$\|W_{h+1}^i\|_\infty \leq H^2 =: C_w, \quad (196)$$

Applying Lemma 9 with (195) and  $N = N_h^k = N_h^k(s,a)$ , with probability at least  $1 - \frac{\delta}{2}$ , we arrive at

$$\begin{aligned} & \left| \frac{1}{N_h^k(s,a) \vee 1} \sum_{n=1}^{N_h^k} (P_h^{k^n} - P_{h,s,a})(V_{h+1}^*)^2 \right| = \left| \sum_{i=1}^k X_i(s,a,h,N_h^k) \right| \\ & \lesssim \sqrt{C_u \log^2 \frac{SAT}{\delta}} \sqrt{\sum_{n=1}^{N_h^k} u_h^{k^n}(s,a,N_h^k) \text{Var}_{h,s,a}(W_{h+1}^{k^n})} + \left( C_u C_w + \sqrt{\frac{C_u}{N_h^k \vee 1}} C_w \right) \log^2 \frac{SAT}{\delta} \\ & \asymp \sqrt{\frac{\iota^2}{N_h^k \vee 1}} \sqrt{\sum_{n=1}^{N_h^k} \frac{1}{N_h^k \vee 1} \|W_{h+1}^{k^n}\|_\infty^2} + \frac{H^2 \iota^2}{N_h^k \vee 1} \lesssim H^2 \iota^2 \sqrt{\frac{1}{N_h^k \vee 1}}. \end{aligned} \quad (197)$$



Similarly, for the second term in (194), with  $W_{h+1}^i := V_{h+1}^*$ , we have with probability at least  $1 - \frac{\delta}{2}$ ,

$$\frac{1}{N_h^k(s, a) \vee 1} \sum_{n=1}^{N_h^k} \left( P_h^{k^n} - P_{h,s,a} \right) V_{h+1}^* \lesssim H \iota^2 \sqrt{\frac{1}{N_h^k(s, a) \vee 1}}. \quad (198)$$

Inserting (197) and (198) back into (194) yields

$$|G_{h,k}^2 - \text{Var}_{h,s,a}(V_{h+1}^*)| \lesssim H^2 \iota^2 \sqrt{\frac{1}{N_h^k(s, a) \vee 1}}. \quad (199)$$

Consequently, the first term in (193) can be controlled as

$$\begin{aligned} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{k=1}^K \sqrt{\frac{|G_{h,k}^2 - \text{Var}_{h,s,a}(V_{h+1}^*)|}{N_h^k(s, a) \vee 1}} &\lesssim H \iota \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{k=1}^K \frac{1}{(N_h^k(s, a))^{\frac{3}{4}} \vee 1} \\ &\lesssim H^2 (SC^*)^{\frac{3}{4}} K^{\frac{1}{4}} \iota^2, \end{aligned} \quad (200)$$

where the last inequality holds due to (158).

**Step 2: controlling the second term of (193).** The second term can be decomposed as

$$\begin{aligned} &\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{k=1}^K \sqrt{\frac{\text{Var}_{h,s,a}(V_{h+1}^*)}{N_h^k(s, a) \vee 1}} \\ &\stackrel{(i)}{\lesssim} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{k=1}^K \sqrt{\frac{C^* \iota d_h^{\pi^*}(s, a) \text{Var}_{h,s,a}(V_{h+1}^*)}{k}} \mathbb{1}(a = \pi_h^*(s)) \\ &\stackrel{(ii)}{\lesssim} \sqrt{C^* \iota} \sqrt{\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{k=1}^K \text{Var}_{h,s,a}(V_{h+1}^*)} \sqrt{\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{k=1}^K \frac{1}{k} \mathbb{1}(a = \pi_h^*(s))} \\ &\lesssim \sqrt{HSC^* K \iota^2} \sqrt{\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \text{Var}_{h,s,a}(V_{h+1}^*)}, \end{aligned} \quad (201)$$

where (i) follows from the facts  $\frac{1}{N_h^k(s, a) \vee 1} \leq \frac{8\iota}{k d_h^{\pi^*}(s, a)}$  by Lemma 8 and the definition of  $C^*$  in Assumption 1, (ii) holds by the Cauchy-Schwarz inequality, and the final inequality comes from the fact that  $\pi^*$  is deterministic.

We are then left with bounding  $\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \text{Var}_{h,s,a}(V_{h+1}^*)$ . Note that

$$\begin{aligned} &\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \text{Var}_{h,s,a}(V_{h+1}^*) = \mathbb{E}_{s_1 \sim \rho, s_{h+1} \sim P_{h,s_h, \pi_h^*(s_h)}} \left[ \sum_{h=1}^H \text{Var}_{h,s_h, \pi_h^*(s_h)}(V_{h+1}^*) \right] \\ &\stackrel{(i)}{=} \mathbb{E}_{s_1 \sim \rho, s_{h+1} \sim P_{h,s_h, \pi_h^*(s_h)}} \left[ \sum_{h=1}^H \left( r_h(s_h, \pi_h^*(s_h)) + V_{h+1}^*(s_{h+1}) - V_h^*(s_h) \right)^2 \right] \\ &\stackrel{(ii)}{=} \mathbb{E}_{s_1 \sim \rho, s_{h+1} \sim P_{h,s_h, \pi_h^*(s_h)}} \left[ \sum_{h=1}^H \left( r_h(s_h, \pi_h^*(s_h)) + V_{h+1}^*(s_{h+1}) - V_h^*(s_h) \right) \right]^2 \\ &\stackrel{(iii)}{=} \mathbb{E}_{s_1 \sim \rho, s_{h+1} \sim P_{h,s_h, \pi_h^*(s_h)}} \left[ \left( \sum_{h=1}^H r_h(s_h, \pi_h^*(s_h)) \right) - V_1^*(s_1) \right]^2 \stackrel{(iv)}{\leq} H^2, \end{aligned} \quad (202)$$

where (i) follows from Bellman's optimality equation, (ii) follows from the Markov property, (iii) holds due to the fact that  $V_{H+1}^*(s) = 0$  for all  $s \in \mathcal{S}$ , and (iv) arises from the fact  $r_h(s, a) \leq 1$  for all  $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ .

Substituting (202) back into (201), we get

$$\sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{k=1}^K \sqrt{\frac{\text{Var}_{h,s,a}(V_{h+1}^*)}{N_h^k(s, a) \vee 1}} \lesssim \sqrt{H^3 SC^* K \iota^2}. \quad (203)$$

**Step 4: combining the results.** Combining (200) and (203) with (193) yields

$$\begin{aligned} \sum_{h=1}^H \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} d_h^{\pi^*}(s, a) \sum_{k=1}^K \sqrt{\frac{1}{N_h^k(s, a) \vee 1}} G_{h,k} &\lesssim H^2 (SC^*)^{\frac{3}{4}} K^{\frac{1}{4}} \iota^2 + \sqrt{H^3 SC^* K \iota^2} \\ &\lesssim \sqrt{H^3 SC^* K \iota^2} + H^{2.5} SC^* \iota^3. \end{aligned} \quad (204)$$

## C.4 Proof of Lemma 12

In view of (127), we can decompose the term of interest into

$$\left| \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} \xi_h^{k^n} \right| \leq |U_1| + |U_2|,$$

where

$$U_1 := \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_h^{k^n} - P_{h,s,a}) (\bar{V}_{h+1}^{k^n} - V_{h+1}^{k^n}), \quad (205a)$$

$$U_2 := \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( P_{h,s,a} - \frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} P_h^{k^i}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} \right) \bar{V}_{h+1}^{k^n}. \quad (205b)$$

Next, we turn to controlling these two terms separately with the assistance of Lemma 9.

**Step 1: controlling  $U_1$ .** In the following, we invoke Lemma 9 to control  $U_1$  in (205a). Let us set

$$W_{h+1}^i := \bar{V}_{h+1}^i - V_{h+1}^i, \quad \text{and} \quad u_h^i(s, a, N) := \eta_{N_h^i(s,a)}^N \geq 0,$$

which indicates that

$$\|W_{h+1}^i\|_\infty \leq \|\bar{V}_{h+1}^i\|_\infty + \|V_{h+1}^i\|_\infty \leq 2H =: C_w,$$

and

$$\max_{N,h,s,a \in (\{0\} \cup [K]) \times [H] \times \mathcal{S} \times \mathcal{A}} \eta_{N_h^i(s,a)}^N \leq \frac{2H}{N \vee 1} =: C_u. \quad (206)$$

Here, the last inequality follows since (according to Lemma 7 and the definition in (25))

$$\begin{aligned} \eta_{N_h^i(s,a)}^N &\leq \frac{2H}{N \vee 1}, & \text{if } 0 \leq N_h^i(s, a) \leq N; \\ \eta_{N_h^i(s,a)}^N &= 0, & \text{if } N_h^i(s, a) > N. \end{aligned}$$

To continue, it can be seen from (26) that

$$0 \leq \sum_{n=1}^N u_h^{k^n(s,a)}(s, a, N) = \sum_{n=1}^N \eta_n^N \leq 1 \quad (207)$$

holds for all  $(N, s, a) \in [K] \times \mathcal{S} \times \mathcal{A}$ . Therefore, choosing  $N = N_h^k(s, a) = N_h^k$  for any  $(s, a)$  and applying Lemma 9 with the above quantities, we arrive at

$$\begin{aligned}
|U_1| &= \left| \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_h^{k^n} - P_{h,s,a}) (\bar{V}_{h+1}^{k^n} - V_{h+1}^{k^n}) \right| = \left| \sum_{i=1}^k X_i(s, a, h, N_h^k) \right| \\
&\lesssim \sqrt{C_u \log^2 \frac{SAT}{\delta}} \sqrt{\sum_{n=1}^{N_h^k} u_h^{k^n}(s, a, N_h^k) \text{Var}_{h,s,a}(W_{h+1}^{k^n})} + \left( C_u C_w + \sqrt{\frac{C_u}{N \vee 1}} C_w \right) \log^2 \frac{SAT}{\delta} \\
&\asymp \sqrt{\frac{H \ell^2}{N_h^k \vee 1}} \sqrt{\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \text{Var}_{h,s,a}(\bar{V}_{h+1}^{k^n} - V_{h+1}^{k^n})} + \frac{H^2 \ell^2}{N_h^k \vee 1} \tag{208}
\end{aligned}$$

$$\lesssim \sqrt{\frac{H \ell^2}{N_h^k \vee 1}} \sqrt{\sigma_h^{\text{adv}, k^{N_h^k+1}}(s, a) - (\mu_h^{\text{adv}, k^{N_h^k+1}}(s, a))^2} + \frac{H^{7/4} \ell^2}{(N_h^k \vee 1)^{3/4}} + \frac{H^2 \ell^2}{N_h^k \vee 1}. \tag{209}$$

with probability at least  $1 - \delta$ . Here, the proof of the inequality (209) is postponed to Appendix C.4.1 in order to streamline the presentation of the analysis.

**Step 2: bounding  $U_2$ .** Making use of the result in (111), we arrive at

$$\frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} P_h^{k^i} \bar{V}_{h+1}^{k^i}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} = \frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} P_h^{k^i} \bar{V}_{h+1}^{\text{next}, k^i}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1}.$$

To continue, for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , we rewrite and rearrange  $U_2$  (cf. (205b)) as follows:

$$\begin{aligned}
U_2 &= \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( P_{h,s,a} - \frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} P_h^{k^i}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} \right) \bar{V}_{h+1}^{k^n} \\
&= \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( P_{h,s,a} \bar{V}_{h+1}^{k^n} - \frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} P_h^{k^i} \bar{V}_{h+1}^{\text{next}, k^i}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} \right) \\
&\stackrel{(i)}{=} \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( \frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} P_{h,s,a}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} \bar{V}_{h+1}^{k^n} - \frac{\sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} P_h^{k^i} \bar{V}_{h+1}^{\text{next}, k^i}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} \right) \\
&= \sum_{n=1}^{N_h^k} \frac{\eta_n^{N_h^k}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} \sum_{i=N_h^{(m^n-1,1)}+1}^{N_h^{(m^n,1)}} (P_{h,s,a} - P_h^{k^i}) \bar{V}_{h+1}^{\text{next}, k^i} \\
&\stackrel{(ii)}{=} \sum_{i=1}^{N_h^k} \left( \sum_{n=N_h^{(m^i+1,1)}+1}^{N_h^{(m^i+2,1)} \wedge N_h^k} \frac{\eta_n^{N_h^k}}{\hat{N}_h^{\text{epo}, m^n-1}(s, a) \vee 1} \right) (P_{h,s,a} - P_h^{k^i}) \bar{V}_{h+1}^{\text{next}, k^i} \\
&= \sum_{i=1}^{N_h^k} \left( \sum_{n=N_h^{(m^i+1,1)}+1}^{N_h^{(m^i+2,1)} \wedge N_h^k} \frac{\eta_n^{N_h^k}}{\hat{N}_h^{\text{epo}, m^i} \vee 1} \right) (P_{h,s,a} - P_h^{k^i}) \bar{V}_{h+1}^{\text{next}, k^i},
\end{aligned}$$

where (i) follows from the fact that  $N_h^{(m^n,1)} - N_h^{(m^n-1,1)} = \hat{N}_h^{\text{epo}, m^n-1}(s, a)$ , and (ii) is obtained by rearranging terms with respect to  $i$  (the terms with respect to  $\bar{V}_{h+1}^{\text{next}, k^i}$  will only be added during the epoch  $m^i + 1$ ), and the last equality holds since  $m^n - 1 = m^i$  for all  $n = N_h^{(m^i+1,1)} + 1, N_h^{(m^i+1,1)} + 2, \dots, N_h^{(m^i+2,1)}$ .

With the above relation in mind, we are ready to invoke Lemma 9 to control  $U_2$ . To continue, for any episode  $j \leq k$ , let us denote by  $m(j)$  the index of the epoch in which episode  $j$  happens (with slight abuse of notation). Let us set

$$W_{h+1}^j := \bar{V}_{h+1}^{\text{next},j}, \quad \text{and} \quad u_h^j(s, a, N) := \sum_{n=N_h^{(m(j)+1,1)}+1}^{N_h^{(m(j)+2,1)} \wedge N} \frac{\eta_n^N}{\widehat{N}_h^{\text{epo},m(j)}(s, a) \vee 1}.$$

As a result, we see that

$$\|W_{h+1}^j\|_\infty \leq \|\bar{V}_{h+1}^{\text{next},j}\|_\infty \leq H =: C_w$$

and the following fact (which will be established in Appendix C.4.2)

$$0 \leq u_h^j(s, a, N) = \sum_{n=N_h^{(m(j)+1,1)}+1}^{N_h^{(m(j)+2,1)} \wedge N} \frac{\eta_n^N}{\widehat{N}_h^{\text{epo},m(j)}(s, a) \vee 1} \leq \frac{64e^2\iota}{N \vee 1} =: C_u \quad (210)$$

holds for all  $(j, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$  with probability at least  $1 - \delta$ .

Given that  $N = N_h^k(s, a) = N_h^k$ , applying Lemma 9 with the above quantities, we can show that for any state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ ,

$$\begin{aligned} |U_2| &= \left| \sum_{i=1}^{N_h^k} \left( \sum_{n=N_h^{(m(i)+1,1)}+1}^{N_h^{(m(i)+2,1)} \wedge N_h^k} \frac{\eta_n^{N_h^k}}{\widehat{N}_h^{\text{epo},m(i)} \vee 1} \right) (P_{h,s,a} - P_h^{k,i}) \bar{V}_{h+1}^{\text{next},k,i} \right| = \left| \sum_{j=1}^k X_j(s, a, h, N_h^k) \right| \\ &\lesssim \sqrt{C_u \log^2 \frac{SAT}{\delta}} \sqrt{\sum_{i=1}^{N_h^k(s,a)} u_h^{k,i}(s, a, N) \text{Var}_{h,s,a}(W_{h+1}^{k,i}(s, a))} + \left( C_u C_w + \sqrt{\frac{C_u}{N \vee 1}} C_w \right) \log^2 \frac{SAT}{\delta} \\ &\lesssim \sqrt{\frac{\iota^3}{N_h^k \vee 1}} \sqrt{\frac{1}{N_h^k \vee 1} \sum_{i=1}^{N_h^k} \text{Var}_{h,s,a}(\bar{V}_{h+1}^{\text{next},k,i})} + \frac{H\iota^3}{N_h^k \vee 1} \\ &\lesssim \sqrt{\frac{\iota^3}{N_h^k \vee 1}} \sqrt{\sigma_h^{\text{ref},k,N_h^k+1}(s, a) - (\mu_h^{\text{ref},k,N_h^k+1}(s, a))^2} + \frac{H\iota^3}{(N_h^k \vee 1)^{3/4}}. \end{aligned} \quad (211)$$

To streamline the presentation of the analysis, we shall postpone the proof of (211) to Appendix C.4.3.

**Step 3: summing up.** Combining the bounds in (209) and (211) yields that: for any state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ ,

$$\begin{aligned} \left| \sum_{n=1}^{N_h^k(s,a)} \eta_n^{N_h^k(s,a)} \xi_h^{k^n} \right| &\leq |U_1| + |U_2| \\ &\lesssim \sqrt{\frac{H\iota^2}{N_h^k \vee 1}} \sqrt{\sigma_h^{\text{adv},k,N_h^k+1}(s, a) - (\mu_h^{\text{adv},k,N_h^k+1}(s, a))^2} \\ &\quad + \sqrt{\frac{\iota^3}{N_h^k \vee 1}} \sqrt{\sigma_h^{\text{ref},k,N_h^k+1}(s, a) - (\mu_h^{\text{ref},k,N_h^k+1}(s, a))^2} + c_b \frac{H^{7/4}\iota^2}{(N_h^k \vee 1)^{3/4}} + c_b \frac{H^2\iota^2}{N_h^k \vee 1} \\ &\leq \bar{B}_h^{k,N_h^k+1}(s, a) + c_b \frac{H^{7/4}\iota^2}{(N_h^k \vee 1)^{3/4}} + c_b \frac{H^2\iota^2}{N_h^k \vee 1} \end{aligned} \quad (212)$$

holds for some sufficiently large constant  $c_b > 0$ , where the last line follows from the definition of  $\bar{B}_h^{k,N_h^k+1}(s, a)$  in line 14 of Algorithm 3. As a consequence of the inequality (212), for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , one has

$$\left| \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \xi_h^{k^n} \right| \leq \bar{B}_h^{k,N_h^k+1}(s, a) + c_b \frac{H^{7/4}\iota^2}{(N_h^k \vee 1)^{3/4}} + c_b \frac{H^2\iota^2}{N_h^k \vee 1} \leq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \bar{b}_h^{k^n+1},$$

where the last inequality holds due to (120). We have thus concluded the proof of Lemma 12.

#### C.4.1 Proof of inequality (209)

To establish the inequality (209), it is sufficient to consider the difference

$$W_1 := \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \text{Var}_{h,s,a}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) - \sigma_h^{\text{adv},k^{N_h^k+1}}(s,a) + (\mu_h^{\text{adv},k^{N_h^k+1}}(s,a))^2.$$

Before continuing, it is easily verified that if  $N_h^k = N_h^k(s,a) = 0$ , the basic fact  $\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} = 0$  leads to  $W_1 = 0$ , and therefore, (209) holds directly. The remainder of the proof is thus dedicated to controlling  $W_1$  when  $N_h^k = N_h^k(s,a) \geq 1$ . Recalling the definition in (109)

$$\text{Var}_{h,s,a}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) = P_{h,s,a}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})^2 - \left(P_{h,s,a}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})\right)^2, \quad (213)$$

we can take this result together with (112) to yield

$$\begin{aligned} W_1 &= \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})^2 - \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left(P_{h,s,a}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})\right)^2 \\ &\quad + \left(\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_h^{k^n}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})\right)^2 - \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_h^{k^n}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})^2 \\ &\leq \underbrace{\left[\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_h^{k^n} - P_{h,s,a})(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})^2\right]}_{=: W_1^1} \\ &\quad + \underbrace{\left(\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_h^{k^n}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})\right)^2 - \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left(P_{h,s,a}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})\right)^2}_{=: W_1^2}. \end{aligned} \quad (214)$$

It then boils down to control the above two terms in (214) separately when  $N_h^k = N_h^k(s,a) \geq 1$ .

**Step 1: controlling  $W_1^1$ .** To control  $W_1^1$ , we shall invoke Lemma 9 by setting

$$W_{h+1}^i := (V_{h+1}^i - \bar{V}_{h+1}^i)^2, \quad \text{and} \quad u_h^i(s,a,N) := \eta_{N_h^i(s,a)}^N \geq 0,$$

which obey

$$\|W_{h+1}^i\|_\infty \leq \|\bar{V}_{h+1}^i\|_\infty^2 + \|V_{h+1}^i\|_\infty^2 \leq 2H^2 =: C_w.$$

Invoking the facts in (206) and (207), we arrive at

$$\frac{2H}{N \vee 1} =: C_u$$

and

$$0 \leq \sum_{n=1}^N u_h^{k^n(s,a)}(s,a,N) \leq 1, \quad \forall (N,s,a) \in [K] \times \mathcal{S} \times \mathcal{A}.$$

Therefore, choosing  $N = N_h^k(s, a) = N_h^k$  for any  $(s, a)$  and applying Lemma 9 with the above quantities, we arrive at, with probability at least  $1 - \delta$ ,

$$\begin{aligned}
|W_1^1| &= \left| \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_h^{k^n} - P_{h,s,a})(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})^2 \right| = \left| \sum_{i=1}^k X_i(s, a, h, N_h^k) \right| \\
&\lesssim \sqrt{C_u \log^2 \frac{SAT}{\delta}} \sqrt{\sum_{n=1}^{N_h^k} u_h^{k^n}(s, a, N_h^k) \text{Var}_{h,s,a}(W_{h+1}^{k^n})} + \left( C_u C_w + \sqrt{\frac{C_u}{N \vee 1}} C_w \right) \log^2 \frac{SAT}{\delta} \\
&\lesssim \sqrt{\frac{H}{N_h^k \vee 1}} \iota^2 \sqrt{\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \|W_{h+1}^{k^n}\|_\infty^2} + \frac{H^3 \iota^2}{N_h^k \vee 1} \lesssim \sqrt{\frac{H^5}{N_h^k \vee 1}} \iota^2 + \frac{H^3 \iota^2}{N_h^k \vee 1}.
\end{aligned} \tag{215}$$

**Step 2: controlling  $W_1^2$ .** Observe that Jensen's inequality gives

$$\left( \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) \right)^2 \leq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( P_{h,s,a} (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) \right)^2, \tag{216}$$

due to the fact  $\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} = 1$  (see (26) and (25)). Plugging the above relation into (214) gives

$$\begin{aligned}
W_1^2 &\leq \left( \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_h^{k^n} (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) \right)^2 - \left( \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} P_{h,s,a} (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) \right)^2 \\
&= \left( \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_h^{k^n} - P_{h,s,a}) (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) \right) \cdot \left( \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_h^{k^n} + P_{h,s,a}) (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) \right).
\end{aligned} \tag{217}$$

Note that the first term in (217) is exactly  $|U_1|$  defined in (205a), which can be controlled by invoking (208) to achieve that, with probability at least  $1 - \delta$ ,

$$\begin{aligned}
&\left| \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_h^{k^n} - P_{h,s,a}) (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) \right| \\
&\lesssim \sqrt{\frac{H}{N_h^k \vee 1}} \iota^2 \sqrt{\sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \text{Var}_{h,s,a}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n})} + \frac{H^2 \iota^2}{N_h^k \vee 1} \lesssim \sqrt{\frac{H^3 \iota^2}{N_h^k \vee 1}} + \frac{H^2}{N_h^k \vee 1} \iota^2,
\end{aligned} \tag{218}$$

where the final inequality holds since  $\text{Var}_{h,s,a}(V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) \lesssim H^2$  and the fact in (26). In addition, the second term in (217) can be controlled straightforwardly by

$$\left| \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} (P_h^{k^n} + P_{h,s,a}) (V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}) \right| \leq \sum_{n=1}^{N_h^k} \eta_n^{N_h^k} \left( \|P_h^{k^n}\|_1 + \|P_{h,s,a}\|_1 \right) \|V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\|_\infty \leq 2H,$$

where we have used the fact in (26),  $\|V_{h+1}^{k^n} - \bar{V}_{h+1}^{k^n}\|_\infty \leq H$  and  $\|P_h^{k^n}\|_1 = \|P_{h,s,a}\|_1 = 1$ .

Taking the above two facts collectively with (217) yields

$$W_1^2 \lesssim \sqrt{\frac{H^5 \iota^2}{N_h^k \vee 1}} + \frac{H^3 \iota^2}{N_h^k \vee 1}. \tag{219}$$

**Step 3: summing up.** Plugging the results in (215) and (219) back into (214), we have

$$W_1 \leq W_1^1 + W_1^2 \lesssim \sqrt{\frac{H^5 \iota^2}{N_h^k \vee 1}} + \frac{H^3 \iota^2}{N_h^k \vee 1},$$

which leads to the desired result (209) directly.

### C.4.2 Proof of inequality (210)

To begin with, let us recall two pieces of notation that shall be used throughout this proof:

1.  $m(j)$ : the index of the epoch in which the  $j$ -th episode occurs.
2.  $\widehat{N}_h^{\text{epo},m}(s, a)$ : the value of  $\widehat{N}_h^{(m, L_m+1)}(s, a)$ , representing the number of visits to  $(s, a)$  in the entire  $m$ -th epoch with length  $L_m = 2^m$ .

Applying (56) and taking the union bound over  $(m(j), h, s, a) \in [M] \times [H] \times \mathcal{S} \times \mathcal{A}$  yield

$$\widehat{N}_h^{\text{epo},m(j)}(s, a) \vee 1 \geq \frac{2^{m(j)} d_h^\mu(s, a)}{8 \log\left(\frac{SAT}{\delta}\right)} \quad (220)$$

with probability at least  $1 - \delta/2$ .

For any epoch  $m$ , if we denote by  $k_{\text{last}}(m)$  the index of the last episode in the  $m$ -th epoch, we can immediately see that

$$k_{\text{last}}(m) = \sum_{i=1}^m L_i = \sum_{i=1}^m 2^i = 2^{m+1} - 2 \leq 2^{m+1}. \quad (221)$$

Applying (56) again and taking the union bound over  $(m(j), h, s, a) \in [M] \times [H] \times \mathcal{S} \times \mathcal{A}$ , one can guarantee that for every  $n \in [N_h^{(m(j)+1,1)}, N_h^{(m(j)+2,1)}]$ , with probability at least  $1 - \delta/2$ ,

$$\begin{aligned} N_h^{(m(j)+1,1)} \leq n \leq N_h^{(m(j)+2,1)} &= N_h^{k_{\text{last}}(m(j)+1)} \\ &\leq N_h^{2^{m(j)+2}} \leq \begin{cases} e^2 2^{m(j)+2} d_h^\mu(s, a) & \text{if } 2^{m(j)+2} d_h^\mu(s, a) \geq \log\left(\frac{SAT}{\delta}\right) \\ 2e^2 \log\left(\frac{SAT}{\delta}\right) & \text{if } 2^{m(j)+2} d_h^\mu(s, a) \leq 2 \log\left(\frac{SAT}{\delta}\right) \end{cases}. \end{aligned} \quad (222)$$

Combine the above results to yield

$$\begin{cases} \widehat{N}_h^{\text{epo},m(j)}(s, a) \vee 1 \geq \frac{2^{m(j)} d_h^\mu(s, a)}{8 \log\left(\frac{SAT}{\delta}\right)} \stackrel{(i)}{\geq} \frac{1}{32e^2 \log\left(\frac{SAT}{\delta}\right)} n, & \text{if } 2^{m(j)+2} d_h^\mu(s, a) \geq \log\left(\frac{SAT}{\delta}\right), \\ \widehat{N}_h^{\text{epo},m(j)}(s, a) \vee 1 \geq 1 \stackrel{(iii)}{\geq} \frac{1}{2e^2 \log\left(\frac{SAT}{\delta}\right)} n & \text{if } 2^{m(j)+2} d_h^\mu(s, a) \leq 2 \log\left(\frac{SAT}{\delta}\right), \end{cases} \quad (223)$$

where (i) follows from (220), (ii) and (iii) hold due to (222). As a result, we arrive at

$$\begin{aligned} \sum_{n=N_h^{(m(j)+1,1)}+1}^{N_h^{(m(j)+2,1)} \wedge N} \frac{\eta_n^N}{\widehat{N}_h^{\text{epo},m(j)}(s, a) \vee 1} &\leq \sum_{n=N_h^{(m(j)+1,1)}+1}^{N_h^{(m(j)+2,1)} \wedge N} \frac{32e^2 \log\left(\frac{SAT}{\delta}\right) \eta_n^N}{n} \\ &\leq \sum_{n=N_h^{(m(j)+1,1)}+1}^N \frac{32e^2 \log\left(\frac{SAT}{\delta}\right) \eta_n^N}{n} \leq \frac{64e^2 \log\left(\frac{SAT}{\delta}\right)}{N \vee 1}, \end{aligned}$$

where the last inequality holds since  $\sum_{i=1}^N \frac{\eta_i^N}{i} \leq \frac{2}{N \vee 1}$  (see Lemma 7).

### C.4.3 Proof of inequality (211)

In this subsection, we intend to control the following term

$$W_2 := \frac{1}{N_h^k \vee 1} \sum_{n=1}^{N_h^k} \text{Var}_{h,s,a} \left( \overline{V}_{h+1}^{\text{next},k^n} \right) - \left( \sigma_h^{\text{ref},k^{N_h^k+1}}(s, a) - \left( \mu_h^{\text{ref},k^{N_h^k+1}}(s, a) \right)^2 \right)$$

for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ . First, it is easily seen that if  $N_h^k = 0$ , then we have  $W_2 = 0$  and thus (211) is satisfied. Therefore, the remainder of the proof is devoted to verifying (211) when  $N_h^k = N_h^k(s, a) \geq 1$ . Combining the expression (113) with the following definition

$$\text{Var}_{h,s,a} \left( \overline{V}_{h+1}^{\text{next},k^n} \right) = P_{h,s,a} \left( \overline{V}_{h+1}^{\text{next},k^n} \right)^2 - \left( P_{h,s,a} \overline{V}_{h+1}^{\text{next},k^n} \right)^2,$$

we arrive at

$$\begin{aligned}
W_2 &= \frac{1}{N_h^k \vee 1} \sum_{n=1}^{N_h^k} \left( P_{h,s,a} \left( \bar{V}_{h+1}^{\text{next},k^n} \right)^2 - \left( P_{h,s,a} \bar{V}_{h+1}^{\text{next},k^n} \right)^2 \right) \\
&\quad - \frac{1}{N_h^k \vee 1} \sum_{n=1}^{N_h^k} P_h^{k^n} \left( \bar{V}_{h+1}^{\text{next},k^n} \right)^2 + \left( \frac{1}{N_h^k \vee 1} \sum_{n=1}^{N_h^k} P_h^{k^n} \bar{V}_{h+1}^{\text{next},k^n} \right)^2 \\
&= \underbrace{\frac{1}{N_h^k \vee 1} \sum_{n=1}^{N_h^k} \left( P_{h,s,a} - P_h^{k^n} \right) \left( \bar{V}_{h+1}^{\text{next},k^n} \right)^2}_{=: W_2^1} + \underbrace{\left( \frac{1}{N_h^k \vee 1} \sum_{n=1}^{N_h^k} P_h^{k^n} \bar{V}_{h+1}^{\text{next},k^n} \right)^2 - \frac{1}{N_h^k \vee 1} \sum_{n=1}^{N_h^k} \left( P_{h,s,a} \bar{V}_{h+1}^{\text{next},k^n} \right)^2}_{=: W_2^2}.
\end{aligned} \tag{224}$$

In the sequel, we intend to control the terms in (224) separately.

**Step 1: controlling  $W_2^1$ .** The first term  $W_2^1$  can be controlled by invoking Lemma 9 and set

$$W_{h+1}^i := \left( \bar{V}_{h+1}^{\text{next},i} \right)^2, \quad \text{and} \quad u_h^i(s, a, N) := \frac{1}{N} =: C_u.$$

To proceeding, with the fact

$$\|W_{h+1}^i\|_\infty \leq \left\| \bar{V}_{h+1}^{\text{next},i} \right\|_\infty^2 \leq H^2 =: C_w$$

and  $N = N_h^k(s, a) = N_h^k$ , applying Lemma 9 with the above quantities, we have for all state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ ,

$$\begin{aligned}
|W_2^1| &= \left| \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} \left( P_{h,s,a} - P_h^{k^n} \right) \left( \bar{V}_{h+1}^{\text{next},k^n} \right)^2 \right| = \left| \sum_{i=1}^k X_i(s, a, h, N_h^k) \right| \\
&\lesssim \sqrt{C_u \log^2 \frac{SAT}{\delta}} \sqrt{\sum_{n=1}^{N_h^k(s,a)} u_h^{k^n(s,a)}(s, a, N) \text{Var}_{h,s,a} \left( W_{h+1}^{k^n(s,a)} \right)} + \left( C_u C_w + \sqrt{\frac{C_u}{N}} C_w \right) \log^2 \frac{SAT}{\delta} \\
&\lesssim \sqrt{\frac{\iota^2}{N_h^k}} \sqrt{\|W_{h+1}^i\|_\infty^2} + \frac{H^2 \iota^2}{N_h^k} \lesssim \sqrt{\frac{H^4 \iota^2}{N_h^k}} + \frac{H^2 \iota^2}{N_h^k}.
\end{aligned} \tag{225}$$

**Step 2: controlling  $W_2^2$ .** Towards controlling  $W_2^2$  in (224), we observe that by Jensen's inequality,

$$\left( \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} P_{h,s,a} \bar{V}_{h+1}^{\text{next},k^n} \right)^2 \leq \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} \left( P_{h,s,a} \bar{V}_{h+1}^{\text{next},k^n} \right)^2.$$

Equipped with this relation,  $W_2^2$  satisfies

$$\begin{aligned}
W_2^2 &\leq \left( \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} P_h^{k^n} \bar{V}_{h+1}^{\text{next},k^n} \right)^2 - \left( \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} P_{h,s,a} \bar{V}_{h+1}^{\text{next},k^n} \right)^2 \\
&= \left( \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} \left( P_h^{k^n} - P_{h,s,a} \right) \bar{V}_{h+1}^{\text{next},k^n} \right) \cdot \left( \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} \left( P_h^{k^n} + P_{h,s,a} \right) \bar{V}_{h+1}^{\text{next},k^n} \right).
\end{aligned} \tag{226}$$



As for the first term in (226), let us set

$$W_{h+1}^i := \bar{V}_{h+1}^{\text{next},i}, \quad \text{and} \quad u_h^i(s, a, N) := \frac{1}{N} =: C_u,$$

which satisfy

$$\|W_{h+1}^i\|_\infty \leq \|\bar{V}_{h+1}^{\text{next},i}\|_\infty \leq H =: C_w.$$

For any  $(s, a)$ , Lemma 9 together with the above quantities and  $N = N_h^k = N_h^k(s, a)$  gives

$$\begin{aligned} & \left| \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} (P_h^{k^n} - P_{h,s,a}) \bar{V}_{h+1}^{\text{next},k^n} \right| \\ & \lesssim \sqrt{C_u \log^2 \frac{SAT}{\delta}} \sqrt{\sum_{n=1}^{N_h^k(s,a)} u_h^{k^n(s,a)}(s, a, N) \text{Var}_{h,s,a}(W_{h+1}^{k^n(s,a)})} + \left( C_u C_w + \sqrt{\frac{C_u}{N}} C_w \right) \log^2 \frac{SAT}{\delta} \\ & \lesssim \sqrt{\frac{\iota^2}{N_h^k}} \sqrt{\|W_{h+1}^{k^n(s,a)}\|_\infty^2} + \frac{H \iota^2}{N_h^k} \lesssim \sqrt{\frac{H^2 \iota^2}{N_h^k}} + \frac{H \iota^2}{N_h^k} \end{aligned}$$

with probability at least  $1 - \delta$ . In addition, the second term can be bounded straightforwardly by

$$\left| \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} (P_h^{k^n} + P_{h,s,a}) \bar{V}_{h+1}^{\text{next},k^n} \right| \leq \frac{1}{N_h^k} \sum_{n=1}^{N_h^k} (\|P_h^{k^n}\|_1 + \|P_{h,s,a}\|_1) \|\bar{V}_{h+1}^{\text{next},k^n}\|_\infty \leq 2H,$$

where the last inequality is valid since  $\|\bar{V}_{h+1}^{\text{next},k^n}\|_\infty \leq H$  and  $\|P_h^{k^n}\|_1 = \|P_{h,s,a}\|_1 = 1$ . Substitution of the above two observations back into (226) yields

$$W_2^2 \lesssim \sqrt{\frac{H^4}{N_h^k \vee 1}} \iota^2 + \frac{H^2}{N_h^k \vee 1} \iota^2. \quad (227)$$

**Step 3: combining the above results.** Plugging the results in (225) and (227) into (224), we reach

$$W_2 \leq W_2^1 + W_2^2 \lesssim \sqrt{\frac{H^4}{N_h^k \vee 1}} \iota^2 + \frac{H^2}{N_h^k \vee 1} \iota^2,$$

thus establishing the desired inequality (211).

## References

- Agarwal, A., Kakade, S., and Yang, L. F. (2020). Model-based reinforcement learning with a generative model is minimax optimal. *Conference on Learning Theory*, pages 67–83.
- Azar, M. G., Osband, I., and Munos, R. (2017). Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, pages 263–272. JMLR. org.
- Bai, Y., Xie, T., Jiang, N., and Wang, Y.-X. (2019). Provably efficient Q-learning with low switching cost. *arXiv preprint arXiv:1905.12849*.
- Beck, C. L. and Srikant, R. (2012). Error bounds for constant step-size Q-learning. *Systems & control letters*, 61(12):1203–1208.
- Bertsekas, D. P. (2017). *Dynamic programming and optimal control (4th edition)*. Athena Scientific.

- Buckman, J., Gelada, C., and Bellemare, M. G. (2020). The importance of pessimism in fixed-dataset policy optimization. In *International Conference on Learning Representations*.
- Chen, J. and Jiang, N. (2019). Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, pages 1042–1051. PMLR.
- Chen, Z., Maguluri, S. T., Shakkottai, S., and Shanmugam, K. (2020). Finite-sample analysis of stochastic approximation using smooth convex envelopes. *arXiv preprint arXiv:2002.00874*.
- Dadashi, R., Rezaeifar, S., Vieillard, N., Hussenot, L., Pietquin, O., and Geist, M. (2021). Offline reinforcement learning with pseudometric learning. *arXiv preprint arXiv:2103.01948*.
- Domingues, O. D., Ménard, P., Kaufmann, E., and Valko, M. (2021). Episodic reinforcement learning in finite MDPs: Minimax lower bounds revisited. In *Algorithmic Learning Theory*, pages 578–598. PMLR.
- Dong, K., Wang, Y., Chen, X., and Wang, L. (2019). Q-learning with UCB exploration is sample efficient for infinite-horizon MDP. *arXiv preprint arXiv:1901.09311*.
- Du, S. S., Chen, J., Li, L., Xiao, L., and Zhou, D. (2017). Stochastic variance reduction methods for policy evaluation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1049–1058. JMLR. org.
- Even-Dar, E. and Mansour, Y. (2003). Learning rates for Q-learning. *Journal of machine learning Research*, 5(Dec):1–25.
- Freedman, D. A. (1975). On tail probabilities for martingales. *the Annals of Probability*, pages 100–118.
- Fujimoto, S., Meger, D., and Precup, D. (2019). Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pages 2052–2062. PMLR.
- Jaakkola, T., Jordan, M. I., and Singh, S. P. (1994). Convergence of stochastic iterative dynamic programming algorithms. In *Advances in neural information processing systems*, pages 703–710.
- Jafarnia-Jahromi, M., Wei, C.-Y., Jain, R., and Luo, H. (2020). A model-free learning algorithm for infinite-horizon average-reward MDPs with near-optimal regret. *arXiv preprint arXiv:2006.04354*.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. (2018). Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873.
- Jin, Y., Yang, Z., and Wang, Z. (2021). Is pessimism provably efficient for offline RL? In *International Conference on Machine Learning*, pages 5084–5096.
- Johnson, R. and Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, pages 315–323.
- Khamaru, K., Pananjady, A., Ruan, F., Wainwright, M. J., and Jordan, M. I. (2020). Is temporal difference learning optimal? an instance-dependent analysis. *arXiv preprint arXiv:2003.07337*.
- Kidambi, R., Rajeswaran, A., Netrapalli, P., and Joachims, T. (2020). Morel: Model-based offline reinforcement learning. *arXiv preprint arXiv:2005.05951*.
- Kumar, A., Zhou, A., Tucker, G., and Levine, S. (2020). Conservative Q-learning for offline reinforcement learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1179–1191. Curran Associates, Inc.
- Lange, S., Gabel, T., and Riedmiller, M. (2012). Batch reinforcement learning. In *Reinforcement learning*, pages 45–73. Springer.
- Levine, S., Kumar, A., Tucker, G., and Fu, J. (2020). Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*.

- Li, G., Cai, C., Chen, Y., Gu, Y., Wei, Y., and Chi, Y. (2021a). Is Q-learning minimax optimal? a tight sample complexity analysis. *arXiv preprint arXiv:2102.06548*.
- Li, G., Shi, L., Chen, Y., Chi, Y., and Wei, Y. (2022). Settling the sample complexity of model-based offline reinforcement learning. *in preparation*.
- Li, G., Shi, L., Chen, Y., Gu, Y., and Chi, Y. (2021b). Breaking the sample complexity barrier to regret-optimal model-free reinforcement learning. *Advances in Neural Information Processing Systems*, 34.
- Li, G., Wei, Y., Chi, Y., Gu, Y., and Chen, Y. (2020). Breaking the sample size barrier in model-based reinforcement learning with a generative model. In *Advances in Neural Information Processing Systems*, volume 33.
- Li, G., Wei, Y., Chi, Y., Gu, Y., and Chen, Y. (2021c). Sample complexity of asynchronous Q-learning: Sharper analysis and variance reduction. *IEEE Transactions on Information Theory*, 68(1):448–473.
- Liu, Y., Swaminathan, A., Agarwal, A., and Brunskill, E. (2020). Provably good batch reinforcement learning without great exploration. *arXiv preprint arXiv:2007.08202*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., and Ostrovski, G. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.
- Nguyen-Tang, T., Gupta, S., and Venkatesh, S. (2021). Sample complexity of offline reinforcement learning with deep ReLU networks. *arXiv preprint arXiv:2103.06671*.
- Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- Qu, G. and Wierman, A. (2020). Finite-time analysis of asynchronous stochastic approximation and Q-learning. *Conference on Learning Theory*, pages 3185–3205.
- Rajaraman, N., Yang, L., Jiao, J., and Ramchandran, K. (2020). Toward the fundamental limits of imitation learning. *Advances in Neural Information Processing Systems*, 33.
- Rashidinejad, P., Zhu, B., Ma, C., Jiao, J., and Russell, S. (2021). Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *Neural Information Processing Systems (NeurIPS)*.
- Rezaeifar, S., Dadashi, R., Vieillard, N., Hussenot, L., Bachem, O., Pietquin, O., and Geist, M. (2021). Offline reinforcement learning as anti-exploration. *arXiv preprint arXiv:2106.06431*.
- Sidford, A., Wang, M., Wu, X., Yang, L., and Ye, Y. (2018a). Near-optimal time and sample complexities for solving Markov decision processes with a generative model. In *Advances in Neural Information Processing Systems*, pages 5186–5196.
- Sidford, A., Wang, M., Wu, X., and Ye, Y. (2018b). Variance reduced value iteration and faster algorithms for solving Markov decision processes. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 770–787. SIAM.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359.
- Uehara, M. and Sun, W. (2021). Pessimistic model-based offline reinforcement learning under partial coverage. *arXiv preprint arXiv:2107.06226*.
- Uehara, M., Zhang, X., and Sun, W. (2021). Representation learning for online and offline RL in Low-rank MDPs. *arXiv preprint arXiv:2110.04652*.
- Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press.

- Wai, H.-T., Hong, M., Yang, Z., Wang, Z., and Tang, K. (2019). Variance reduced policy evaluation with smooth function approximation. *Advances in Neural Information Processing Systems*, 32:5784–5795.
- Wainwright, M. J. (2019). Variance-reduced Q-learning is minimax optimal. *arXiv preprint arXiv:1906.04697*.
- Watkins, C. J. and Dayan, P. (1992). Q-learning. *Machine learning*, 8(3-4):279–292.
- Watkins, C. J. C. H. (1989). Learning from delayed rewards. *PhD thesis, King’s College, University of Cambridge*.
- Xie, T., Cheng, C.-A., Jiang, N., Mineiro, P., and Agarwal, A. (2021a). Bellman-consistent pessimism for offline reinforcement learning. *arXiv preprint arXiv:2106.06926*.
- Xie, T., Jiang, N., Wang, H., Xiong, C., and Bai, Y. (2021b). Policy finetuning: Bridging sample-efficient offline and online reinforcement learning. *arXiv preprint arXiv:2106.04895*.
- Xu, T., Wang, Z., Zhou, Y., and Liang, Y. (2019). Reanalysis of variance reduced temporal difference learning. In *International Conference on Learning Representations*.
- Yan, Y., Li, G., Chen, Y., and Fan, J. (2022). The efficacy of pessimism in asynchronous Q-learning. *in preparation*.
- Yang, K., Yang, L., and Du, S. (2021). Q-learning with logarithmic regret. In *International Conference on Artificial Intelligence and Statistics*, pages 1576–1584. PMLR.
- Yin, M., Bai, Y., and Wang, Y.-X. (2021a). Near-optimal offline reinforcement learning via double variance reduction. *arXiv preprint arXiv:2102.01748*.
- Yin, M., Bai, Y., and Wang, Y.-X. (2021b). Near-optimal provable uniform convergence in offline policy evaluation for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1567–1575. PMLR.
- Yin, M. and Wang, Y.-X. (2021). Towards instance-optimal offline reinforcement learning with pessimism. *Advances in neural information processing systems*, 34.
- Yin, M., Wang, Y.-X., Duan, Y., and Wang, M. (2022). Near-optimal offline reinforcement learning with linear representation: Leveraging variance information with pessimism.
- Yu, T., Kumar, A., Chebotar, Y., Hausman, K., Levine, S., and Finn, C. (2021a). Conservative data sharing for multi-task offline reinforcement learning. *arXiv preprint arXiv:2109.08128*.
- Yu, T., Kumar, A., Rafailov, R., Rajeswaran, A., Levine, S., and Finn, C. (2021b). Combo: Conservative offline model-based policy optimization. *arXiv preprint arXiv:2102.08363*.
- Yu, T., Thomas, G., Yu, L., Ermon, S., Zou, J., Levine, S., Finn, C., and Ma, T. (2020). MOPO: Model-based offline policy optimization. *arXiv preprint arXiv:2005.13239*.
- Zanette, A., Wainwright, M. J., and Brunskill, E. (2021). Provable benefits of actor-critic methods for offline reinforcement learning. *arXiv preprint arXiv:2108.08812*.
- Zhang, Z., Ji, X., and Du, S. S. (2020a). Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. *arXiv preprint arXiv:2009.13503*.
- Zhang, Z., Zhou, Y., and Ji, X. (2020b). Almost optimal model-free reinforcement learning via reference-advantage decomposition. *Advances in Neural Information Processing Systems*, 33.
- Zhang, Z., Zhou, Y., and Ji, X. (2020c). Model-free reinforcement learning: from clipped pseudo-regret to sample complexity. *arXiv preprint arXiv:2006.03864*.