Variance-Aware Off-Policy Evaluation with Linear Function Approximation

Yifei Min*
Department of Statistics and Data Science
Yale University
CT 06511
yifei.min@yale.edu

Tianhao Wang*
Department of Statistics and Data Science
Yale University
CT 06511
tianhao.wang@yale.edu

Dongruo Zhou
Department of Computer Science
University of California, Los Angeles
CA 90095
drzhou@cs.ucla.edu

Quanquan Gu
Department of Computer Science
University of California, Los Angeles
CA 90095
qgu@cs.ucla.edu

Abstract

We study the off-policy evaluation (OPE) problem in reinforcement learning with linear function approximation, which aims to estimate the value function of a target policy based on the offline data collected by a behavior policy. We propose to incorporate the variance information of the value function to improve the sample efficiency of OPE. More specifically, for time-inhomogeneous episodic linear Markov decision processes (MDPs), we propose an algorithm, `VA-OPE`, which uses the estimated variance of the value function to reweight the Bellman residual in Fitted Q-Iteration. We show that our algorithm achieves a tighter error bound than the best-known result. We also provide a fine-grained characterization of the distribution shift between the behavior policy and the target policy. Extensive numerical experiments corroborate our theory.

1 Introduction

Reinforcement learning (RL) has been a hot spot in both theory and practice in the past decade. Many efficient algorithms have been proposed and theoretically analyzed for finding the optimal policy adopted by an agent to maximize the long-term cumulative rewards. In contrast to online RL where the agent actively interacts with the environment, offline RL (a.k.a., batch RL) [25, 24] aims to extract information from past data and use this information to learn the optimal policy. There has been much empirical success of offline RL in various application domains [4, 6, 36, 39, 35].

Among various tasks of offline RL, an important task is called *off-policy evaluation* (OPE), which evaluates the performance of a target policy π given offline data generated by a behavior policy $\bar{\pi}$. Most existing theoretical works on OPE are in the setting of tabular MDPs [33, 26, 11, 16, 43, 45, 46], where the state space $\mathcal{S}$ and the action space $\mathcal{A}$ are both finite. However, real-world applications often have high-dimensional or even infinite-dimensional state and action spaces, where function approximation is required for computational tractability and generalization. While provably efficient online RL with linear function approximation has been widely studied recently [44, 17, 47, 15, 3, 51], little work has been done for analyzing OPE with linear function approximation, with one notable exception by Duan et al. [10]. More specifically, Duan et al. [10] analyzed a regression-based

*Equal contribution.

35th Conference on Neural Information Processing Systems (NeurIPS 2021).

Fitted Q-Iteration method (FQI-OPE) that achieves an $\widetilde{\mathcal{O}}(H^2\sqrt{(1+d(\pi,\bar{\pi}))/N})$ error for linear MDPs [44, 17], where H is the planning horizon, N is the sample size, and $d(\pi,\bar{\pi})$ represents the distribution shift between the behavior policy and the target policy. They also proved a sample complexity lower bound for a subclass of linear MDPs, for which their algorithm is nearly minimax optimal. However, as we will show later, the H^2 dependence is not tight since they discard the useful variance information contained in the offline data. Consequently, their result is only optimal for a small class of MDPs of which the value functions have large variance. The H^2 dependence in the sample complexity also makes their algorithm less sample-efficient for long-horizon problems, which is one of the major challenges in RL.

Extracting useful information from the data is particularly important for offline RL since the agent cannot sample additional data by interacting with the environment, as compared to online RL. In this paper, we propose a new algorithm that incorporates the variance information of the value functions to improve the sample efficiency of OPE. This allows us to achieve a deeper understanding and tighter error bounds of OPE with linear function approximation. In detail, we consider time-inhomogeneous linear MDPs [44, 17] where the transition probability and reward function are assumed to be linear functions of a known feature mapping and may vary from stage to stage.

The main contributions of this paper are summarized as follows:

- We develop VA-OPE (Variance-Aware Off-Policy Evaluation), an algorithm for OPE that effectively utilizes the variance information from the offline data. The core idea behind the proposed algorithm is to calibrate the Bellman residual in the regression by an estimator of the conditional variance of the value functions, such that data points of higher quality can receive larger important weights.
- We show that our algorithm achieves $\widetilde{\mathcal{O}}(\sum_h(\mathbf{v}_h^\top\mathbf{\Lambda}_h^{-1}\mathbf{v}_h)^{1/2}/\sqrt{K})$ policy evaluation error, where $\mathbf{v}_h$ is the expectation of the feature vectors under target policy and $\mathbf{\Lambda}_h$ is the uncentered covariance matrix under behavior policy weighted by the conditional variance of the value function. Our algorithm achieves a tighter error bound and milder dependence on H than FQI-OPE [10], and provides a tighter characterization of the distribution shift between the behavior policy and the target policy, which is also verified by extensive numerical experiments.
- Our analysis is based on a novel two-step proof technique. In the first step, we use backward induction to establish worst-case uniform convergence[2] results for the estimators of the value functions. In the second step, the convergence of OPE estimator is proved by tightening the uniform convergence result based on an average-case analysis. Our proof strategy provides a generic way for analyzing (weighted) ridge regression methods that are carried out in a backward and iterated fashion. The analyses in both steps might be of independent interest.

Notation We use lower case letters to denote scalars and use lower and upper case boldface letters to denote vectors and matrices respectively. For any vector $\mathbf{x}\in\mathbb{R}^d$ and any positive semi-definite matrix $\mathbf{\Sigma}\in\mathbb{R}^{d\times d}$, we denote by $\|\mathbf{x}\|_2$ the Euclidean norm and $\|\mathbf{\Sigma}\|$ the operator norm, and define $\|\mathbf{x}\|_{\mathbf{\Sigma}}=\sqrt{\mathbf{x}^\top\mathbf{\Sigma}\mathbf{x}}$. For any positive integer n, we denote by $[n]$ the set $\{1,\ldots,n\}$. For any finite set A, we denote by $|A|$ the cardinality of A. For two sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n=\mathcal{O}(b_n)$ if there exists an absolute constant C such that $a_n\le Cb_n$, and we write $a_n=\Omega(b_n)$ if there exists an absolute constant C such that $a_n\ge Cb_n$. We use $\widetilde{\mathcal{O}}(\cdot)$ to further hide the logarithmic factors.

2 Preliminaries

2.1 Markov Decision Processes

We consider the time-inhomogeneous episodic Markov Decision Process (MDP), which is represented by a tuple $M(\mathcal{S},\mathcal{A},H,\{r_h\}_{h=1}^H,\{\mathbb{P}_h\}_{h=1}^H)$. In specific, we denote the state space by $\mathcal{S}$ and the action space by $\mathcal{A}$, and $H>0$ is the horizon length of each episode. At each stage $h\in[H]$, $r_h:\mathcal{S}\times\mathcal{A}\to[0,1]$ is the reward function, and $\mathbb{P}_h(s'|s,a)$ is the transition probability function which represents the probability for state s to transit to state s' given action a. A policy π consists of H mappings $\{\pi_h\}_{h=1}^H$ from $\mathcal{S}$ to the simplex on $\mathcal{A}$, such that for any $(h,s)\in[H]\times\mathcal{S}$, $\pi_h(\cdot|s)$ is a probability distribution over $\mathcal{A}$. Here a policy can be either deterministic (point mass) or stochastic.

[2] By uniform convergence we mean the convergence of the estimated value functions in ℓ_∞-norm to their true values, which is different from the uniform convergence over all policies in Yin et al. [46].

For any policy π, we define the associated action-value function $Q_h^\pi(s,a)$ and value function $V_h^\pi(s)$ at each stage $h \in [H]$ as follows:

$$Q_h^\pi(s,a) = \mathbb{E}_\pi\left[\sum_{i=h}^{H} r_i(s_i,a_i)\middle| s_h = s, a_h = a\right], \qquad V_h^\pi(s) = \int_{\mathcal{A}} Q_h^\pi(s,a)\mathrm{d}\pi_h(a|s), \tag{2.1}$$

where $a_i \sim \pi_i(\cdot|s_i)$ and $s_{i+1} \sim \mathbb{P}_i(\cdot|s_i,a_i)$. For any function $V : \mathcal{S} \to \mathbb{R}$, we introduce the following shorthand notation for the conditional expectation and variance of V:

$$[\mathbb{P}_h V](s,a) = \mathbb{E}_{s'\sim\mathbb{P}_h(\cdot|s,a)}[V(s')], \qquad [\mathbb{V}_h V](s,a) = [\mathbb{P}_h V^2](s,a) - ([\mathbb{P}_h V](s,a))^2. \tag{2.2}$$

Time-inhomogeneous linear MDPs. We consider a special class of MDPs called *linear MDPs* [44, 17]. Note that most of the existing works on RL with linear function approximation rely on this assumption.

Assumption 2.1. $M(\mathcal{S},\mathcal{A},H,\{r_h\}_{h=1}^H,\{\mathbb{P}_h\}_{h=1}^H)$ is called a linear MDP with a *known* feature mapping $\boldsymbol{\phi} : \mathcal{S}\times\mathcal{A} \to \mathbb{R}^d$, if for any $h \in [H]$, there exist $\boldsymbol{\gamma}_h$ and $\boldsymbol{\mu}_h \in \mathbb{R}^d$, such that for any state-action pair $(s,a) \in \mathcal{S}\times\mathcal{A}$, it holds that

$$\mathbb{P}_h(\cdot \mid s,a) = \langle \boldsymbol{\phi}(s,a), \boldsymbol{\mu}_h(\cdot)\rangle, \qquad r_h(s,a) = \langle \boldsymbol{\phi}(s,a), \boldsymbol{\gamma}_h\rangle. \tag{2.3}$$

We assume that at any stage h, for any state-action pair $(s,a) \in \mathcal{S}\times\mathcal{A}$, the reward received by the agent is given by $r = r_h(s,a) + \epsilon_h(s,a)$, where $r_h(s,a) \in [0,1]$ is the expected reward and $\epsilon_h(s,a)$ is the random noise. We assume that the noise is zero-mean and independent of anything else.

Without loss of generality, we assume that $\|\boldsymbol{\gamma}_h\|_2 \le 1$ and $\|\boldsymbol{\phi}(s,a)\|_2 \le 1$ for all $(s,a) \in \mathcal{S}\times\mathcal{A}$. We also assume that $r_h(s,a)+\epsilon_h(s,a) \le 1$, $|\epsilon_h(s,a)| \le 1$ almost surely and thus $\mathrm{Var}(\epsilon_h(s,a)) \le 1$ for all $h \in [H]$ and $(s,a) \in \mathcal{S}\times\mathcal{A}$. Moreover, we assume that $\max_{h\in[H]} \left\|\int_{\mathcal{S}} f(s)\mathrm{d}\boldsymbol{\mu}_h(s)\right\|_2 \le \sqrt{d}$ for all bounded function $f : \mathcal{S} \to \mathbb{R}$ such that $\sup_{s\in\mathcal{S}} |f(s)| \le 1$.

The above assumption on linear MDPs implies the following proposition for the action-value functions.

Proposition 2.2 (Proposition 2.3, [17])**.** For a linear MDP, for any policy π, there exist weights $\{\mathbf{w}_h^\pi, h \in [H]\}$ such that for any $(s,a,h) \in \mathcal{S}\times\mathcal{A}\times[H]$, we have $Q_h^\pi(s,a) = \langle \boldsymbol{\phi}(s,a), \mathbf{w}_h^\pi\rangle$. Moreover, we have $\|\mathbf{w}_h^\pi\|_2 \le 2H\sqrt{d}$ for all $h \in [H]$.

Following this proposition, we may further show that the value functions are also linear functions, but of different features. We define $\boldsymbol{\phi}_h^\pi(s) = \int_{\mathcal{A}} \boldsymbol{\phi}(s,a)\mathrm{d}\pi_h(a|s)$ for all $s \in [S]$ and $h \in [H]$. Then by (2.1) we have

$$V_h^\pi(s) = \int_{\mathcal{A}} \boldsymbol{\phi}(s,a)^\top \mathbf{w}_h^\pi \mathrm{d}\pi(a|s) = \langle \boldsymbol{\phi}_h^\pi(s), \mathbf{w}_h^\pi\rangle.$$

2.2 Off-policy Evaluation

The purpose of OPE is to evaluate a (known) target policy π given an offline dataset generated by a different (unknown) behavior policy $\bar{\pi}$. In this paper, our goal is to estimate the expectation of the value function induced by π over a fixed initial distribution ξ_1, i.e.,

$$v_1^\pi = \mathbb{E}_{s\sim\xi_1}[V_1^\pi(s)].$$

To faciliate the presentation, we further introduce some important notations. For all $h \in [H]$, let ν_h be the occupancy measure over $\mathcal{S}\times\mathcal{A}$ at stage h induced by the transition $\mathbb{P}$ and the behavior policy $\bar{\pi}$, that is, for any $E \subseteq \mathcal{S}\times\mathcal{A}$,

$$\nu_h(E) = \mathbb{E}\left[(s_h,a_h) \in E \mid s_1 \sim \xi_1,\ a_i \sim \bar{\pi}(\cdot|s_i),\ s_{i+1} \sim \mathbb{P}_i(\cdot|s_i,a_i),\ 1 \le i \le h\right]. \tag{2.4}$$

For simplicity, we write $\mathbb{E}_h[f(s,a)] = \mathbb{E}_{\bar{\pi},h}[f(s,a)] = \int_{\mathcal{S}\times\mathcal{A}} f(s,a)\mathrm{d}\nu_h(s,a)$ for any function f on $\mathcal{S}\times\mathcal{A}$. Similarly, we use $\mathbb{E}_{\pi,h}[f(s,a)]$ to denote the expectation of f with respect to the occupancy measure at stage h induced by the transition $\mathbb{P}$ and the target policy π.

We define the following uncentered covariance matrix under behavior policy for all $h \in [H]$:

$$\boldsymbol{\Sigma}_h = \mathbb{E}_{\bar{\pi},h}\left[\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top\right]. \tag{2.5}$$

Intuitively, these matrices measure the coverage of the offline data in the state-action space. It is known that the success of OPE necessitates a good coverage [10, 42]. Therefore here we make the same coverage assumption on the offline data.

Assumption 2.3 (Coverage)**.** For all $h \in [H]$, $\kappa_h := \lambda_{\min}(\mathbf{\Sigma}_h) > 0$. Denote $\kappa = \min_{h\in[H]} \kappa_h$.

A key difference in our result is that, instead of depending on $\mathbf{\Sigma}_h$ directly, the error bound depends on the following weighted version of the covariance matrices defined as

$$\mathbf{\Lambda}_h := \mathbb{E}_{\bar{\pi},h}\left[\sigma_h(s,a)^{-2}\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top\right], \tag{2.6}$$

for all $h \in [H]$, where each $\sigma_h : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ is defined as

$$\sigma_h(s,a) := \sqrt{\max\{1, \mathbb{V}_h V^\pi_{h+1}(s,a)\} + 1}. \tag{2.7}$$

Note that in the definition of $\sigma_h(\cdot,\cdot)$, taking the maximum and adding an extra 1 is purely for technical reason and is related to its estimator $\widehat{\sigma}_h(\cdot,\cdot)$, which we will introduce and explain later in Section 3.2. In general, one can think of $\sigma_h^2(s,a) \approx \mathbb{V}_h V^\pi_{h+1}(s,a)$. Therefore, compared with the raw covariance matrix $\mathbf{\Sigma}_h$, $\mathbf{\Lambda}_h$ further incorporates the variance of the value functions under the target policy. This is the key to obtaining a tighter instance-dependent error bound.

Definition 2.4 (Variance-aware coverage)**.** We define $\iota_h := \lambda_{\min}(\mathbf{\Lambda}_h)$ and $\iota = \min_{h\in[H]} \iota_h$.

Since $\sup_{(s,a)\in\mathcal{S}\times\mathcal{A}} \sigma_h(s,a)^2$ is bounded from above, by (2.6) and Assumption 2.3, we immediately have $\iota_h \geq \kappa_h/[\sup_{(s,a)\in\mathcal{S}\times\mathcal{A}} \sigma_h(s,a)^2] > 0$ for all $h \in [H]$, and thus $\iota > 0$. Even if Assumption 2.3 does not hold, we can always restrict to the subspace $\text{span}\{\boldsymbol{\phi}(s_h,a_h)\}$. For convenience of presentation, we make Assumption 2.3 in this paper.

Next, we introduce the assumption on the sampling process of the offline data.

Assumption 2.5 (Stage-sampling Data)**.** We have two offline datasets $\mathcal{D}$ and $\check{\mathcal{D}}$ where each dataset consists of data from H stages: $\mathcal{D} = \{\mathcal{D}_h\}_{h\in[H]}$ and $\check{\mathcal{D}} = \{\check{\mathcal{D}}_h\}_{h\in[H]}$. For the dataset $\mathcal{D}$, we assume $\mathcal{D}_{h_1}$ is independent of $\mathcal{D}_{h_2}$ for $h_1 \neq h_2$. For each stage h, we have $\mathcal{D}_h = \{(s_{k,h}, a_{k,h}, r_{k,h}, s'_{k,h})\}_{k\in[K]}$, where we assume for each $k \in [K]$, the data point $(s_{k,h}, a_{k,h}, r_{k,h}, s'_{k,h})$ is sampled identically and independently in the following way: $(s_{k,h}, a_{k,h}) \sim \nu_h(\cdot,\cdot)$ where $\nu_h(\cdot,\cdot)$ is the occupancy measure defined in (2.4), and $s'_{k,h} \sim \mathbb{P}_h(\cdot|s_{k,h}, a_{k,h})$. The same holds for $\check{\mathcal{D}}$, and we write $\check{\mathcal{D}}_h = \{(\check{s}_{k,h}, \check{a}_{k,h}, \check{r}_{k,h}, \check{s}'_{k,h})\}_{k\in[K]}$. Note that here $s'_{k,h} \neq s_{k,h+1}$.

Assumptions 2.5 is standard in the offline RL literature [46, 10]. Note that in the assumption, there is a data splitting, i.e., one can view it as the whole dataset $\mathcal{D} \cup \check{\mathcal{D}}$ being split into two halves. The datasets $\mathcal{D}$ and $\check{\mathcal{D}}$ will then be used for two different purposes in Algorithm 1 as will be made clear in the next section. We would like to remark that the only purpose of the splitting is to avoid a lengthy analysis. There is no need to perform the data splitting in practice. Also, in our implementation and experiments, we do not split the data.

3 Algorithm

To ease the notation, we denote $\boldsymbol{\phi}_{k,h} = \boldsymbol{\phi}(s_{k,h}, a_{k,h})$, $\check{\boldsymbol{\phi}}_{k,h} = \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})$, $\widehat{\sigma}_{k,h} = \widehat{\sigma}_h(s_{k,h}, a_{k,h})$ and $r_{k,h} = r_h(s_{k,h}, a_{k,h}) + \epsilon_{k,h}$ for all $(h,k) \in [H] \times [K]$.

3.1 Regression-Based Value Function Estimation

By Proposition 2.2, it suffices to estimate the vectors $\{\mathbf{w}^\pi_h, h \in [H]\}$. A popular approach is to apply the Least-Square Value Iteration (LSVI) [17] which relies on the Bellman equation, $Q^\pi_h(s,a) = r_h(s,a) + [\mathbb{P}_h V^\pi_{h+1}](s,a)$, that holds for all $h \in [H]$ and $(s,a) \in \mathcal{S} \times \mathcal{A}$. By viewing $V^\pi_{h+1}(s'_{k,h})$ as an unbiased estimate of $[\mathbb{P}_h V^\pi_{h+1}](s_{k,h}, a_{k,h})$, the idea of the LSVI-type method is to solve the following ridge regression problem:

$$\widehat{\mathbf{w}}^\pi_h := \operatorname*{argmin}_{\mathbf{w}\in\mathbb{R}^d} \lambda\|\mathbf{w}\|_2^2 + \sum_{k=1}^K \left[\langle \boldsymbol{\phi}_{k,h}, \mathbf{w}\rangle - r_{k,h} - V^\pi_{h+1}(s'_{k,h})\right]^2, \tag{3.1}$$

for some regularization parameter $\lambda > 0$. Since we do not know the exact values of V^π_{h+1} in (3.1), we replace it by an estimator $\widehat{V}^\pi_{h+1}$, and then recursively solve the lease-square problem in a backward

manner, which enjoys a closed-form solution as follows

$$\widehat{\mathbf{w}}_h^\pi = \left[\sum_{k=1}^K \boldsymbol{\phi}_{k,h}\boldsymbol{\phi}_{k,h}^\top + \lambda \mathbf{I}_d\right]^{-1} \sum_{k=1}^K \boldsymbol{\phi}_{k,h}\left[r_{k,h} + \widehat{V}_{h+1}^\pi(s'_{k,h})\right].$$

This has been used in the LSVI-UCB algorithm proposed by Jin et al. [17] and the FQI-OPE algorithm studied by Duan et al. [10], for online learning and OPE of linear MDPs respectively. For this kind of algorithms, the key difficulty in the analysis lies in bounding the Bellman error:

$$\left[\sum_{k=1}^K \boldsymbol{\phi}_{k,h}\boldsymbol{\phi}_{k,h}^\top + \lambda \mathbf{I}_d\right]^{-1} \sum_{k=1}^K \boldsymbol{\phi}_{k,h}\big([\mathbb{P}_h\widehat{V}_{h+1}^\pi](s_{k,h},a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h})\big).$$

Jin et al. [17] applied a Hoeffding-type inequality to bound the Bellman error. Although Duan et al. [10] applied Freedman's inequality in their analysis, their algorithm design overlooks the variance information in the data and consequently they can only adopt a crude upper bound on the conditional variance of the value function, i.e., $\mathbb{V}_h V_{h+1}^\pi \le (H-h)^2$, which simply comes from $\sup_{s\in\mathcal{S}} V_{h+1}^\pi(s) \le H-h$. Therefore, it prevents [10] from getting a tight instance-dependent error bound for OPE. This is further verified by our numerical experiments in Appendix A which show that the performance of FQI-OPE degrades for large H. This motivates us to utilize the variance information in the data for OPE.

3.2 The Proposed Algorithm

In particular, we present our main algorithm as displayed in Algorithm 1. Due to the greedy nature of the value functions, we adopt a backward estimation scheme.

Weighted ridge regression. For any $h \in [H]$, let $\widehat{\mathbf{w}}_{h+1}^\pi$ be the estimate of $\mathbf{w}_{h+1}^\pi$ computed at the previous step, and correspondingly $\widehat{V}_{h+1}^\pi(\cdot) = \langle \boldsymbol{\phi}_{h+1}^\pi(\cdot), \widehat{\mathbf{w}}_{h+1}^\pi\rangle$. Instead of the ordinary ridge regression (3.1), we consider the following weighted ridge regression:

$$\widehat{\mathbf{w}}_h^\pi := \operatorname*{argmin}_{\mathbf{w}\in\mathbb{R}^d} \lambda\|\mathbf{w}\|_2^2 + \sum_{k=1}^K \left[\langle \boldsymbol{\phi}_{k,h}, \mathbf{w}\rangle - r_{k,h} - \widehat{V}_{h+1}^\pi(s'_{k,h})\right]^2 / \widehat{\sigma}_{k,h}^2, \tag{3.2}$$

where $\widehat{\sigma}_{k,h} = \widehat{\sigma}_h(s_{k,h}, a_{k,h})$ for all $(h,k) \in [H]\times[K]$ with $\widehat{\sigma}_h(\cdot,\cdot)$ being a proper estimate of $\sigma_h(\cdot,\cdot)$ defined in (2.7). We then have the following closed-form solution (Line 9 and 7 of Alg. 1):

$$\widehat{\mathbf{w}}_h^\pi = \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \boldsymbol{\phi}_{k,h}\left(r_{k,h} + \widehat{V}_{h+1}^\pi(s'_{k,h})\right) / \widehat{\sigma}_{k,h}^2, \text{ with } \widehat{\boldsymbol{\Lambda}}_h = \sum_{k=1}^K \widehat{\sigma}_{k,h}^{-2}\boldsymbol{\phi}_{k,h}\boldsymbol{\phi}_{k,h}^\top + \lambda\mathbf{I}_d. \tag{3.3}$$

In the above estimator, we use the dataset $\mathcal{D}$ to estimate the value functions. Next, we apply an LSVI-type method to estimate σ_h using the dataset $\check{\mathcal{D}}$.

Variance estimator. By (2.2), we can write

$$[\mathbb{V}_h V_{h+1}^\pi](s,a) = [\mathbb{P}_h (V_{h+1}^\pi)^2](s,a) - \left([\mathbb{P}_h V_{h+1}^\pi](s,a)\right)^2. \tag{3.4}$$

For the first term in (3.4), by Assumption 2.1 we have

$$[\mathbb{P}_h (V_{h+1}^\pi)^2](s,a) = \int_{\mathcal{S}} V_{h+1}^\pi(s')^2 \mathrm{d}\mathbb{P}_h(s'|s,a) = \boldsymbol{\phi}(s,a)^\top \int_{\mathcal{S}} V_{h+1}^\pi(s')^2 \,\mathrm{d}\boldsymbol{\mu}_h(s'),$$

which suggests that $\mathbb{P}_h(V_{h+1}^\pi)^2$ also has a linear representation. Thus we adopt a linear estimator $\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\beta}}_h^\pi\rangle$ where $\widehat{\boldsymbol{\beta}}_h^\pi$ (Line 4) is the solution to the following ridge regression problem:

$$\widehat{\boldsymbol{\beta}}_h^\pi = \operatorname*{argmin}_{\boldsymbol{\beta}\in\mathbb{R}^d} \sum_{k=1}^K \left[\langle \check{\boldsymbol{\phi}}_{k,h}, \boldsymbol{\beta}\rangle - [\widehat{V}_{h+1}^\pi]^2(\check{s}'_{k,h})\right]^2 + \lambda\|\boldsymbol{\beta}\|_2^2 = \widehat{\boldsymbol{\Sigma}}_h^{-1} \sum_{k=1}^K \check{\boldsymbol{\phi}}_{k,h}\widehat{V}_{h+1}^\pi(\check{s}'_{k,h})^2. \tag{3.5}$$

Algorithm 1 Variance-Aware Off-Policy Evaluation (`VA-OPE`)

1: **Input:** target policy $\pi = \{\pi_h\}_{h\in[H]}$, datasets $\mathcal{D} = \{\{(s_{k,h}, a_{k,h}, r_{k,h}, s'_{k,h})\}_{h\in[H]}\}_{k\in[K]}$ and $\check{\mathcal{D}} = \{\{(\check{s}_{k,h}, \check{a}_{k,h}, \check{r}_{k,h}, \check{s}'_{k,h})\}_{h\in[H]}\}_{k\in[K]}$, initial distribution ξ_1, $\widehat{\mathbf{w}}^{\pi}_{H+1} = \mathbf{0}$
2: **for** $h = H, H-1, \ldots, 1$ **do**
3: $\quad \widehat{\boldsymbol{\Sigma}}_h \leftarrow \sum_{k=1}^{K} \check{\boldsymbol{\phi}}_{k,h}\check{\boldsymbol{\phi}}_{k,h}^{\top} + \lambda \mathbf{I}_d$
4: $\quad \widehat{\boldsymbol{\beta}}_h \leftarrow \widehat{\boldsymbol{\Sigma}}_h^{-1} \sum_{k=1}^{K} \check{\boldsymbol{\phi}}_{k,h}\widehat{V}^{\pi}_{h+1}(\check{s}'_{k,h})^2$
5: $\quad \widehat{\boldsymbol{\theta}}_h \leftarrow \widehat{\boldsymbol{\Sigma}}_h^{-1} \sum_{k=1}^{K} \check{\boldsymbol{\phi}}_{k,h}\widehat{V}^{\pi}_{h+1}(\check{s}'_{k,h})$
6: $\quad \widehat{\sigma}_h(\cdot,\cdot) \leftarrow \sqrt{\max\{1, \widehat{\mathbb{V}}_h \widehat{V}^{\pi}_{h+1}(\cdot,\cdot)\} + 1}$
7: $\quad \widehat{\boldsymbol{\Lambda}}_h \leftarrow \sum_{k=1}^{K} \boldsymbol{\phi}_{k,h}\boldsymbol{\phi}_{k,h}^{\top}/\widehat{\sigma}^2_{k,h} + \lambda \mathbf{I}_d$
8: $\quad Y_{k,h} \leftarrow r_{k,h} + \langle \boldsymbol{\phi}^{\pi}_h(s'_{k,h}), \widehat{\mathbf{w}}^{\pi}_{h+1}\rangle$
9: $\quad \widehat{\mathbf{w}}^{\pi}_h \leftarrow \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^{K} \boldsymbol{\phi}_{k,h} Y_{k,h}/\widehat{\sigma}^2_{k,h}$
10: $\quad \widehat{Q}^{\pi}_h(\cdot,\cdot) \leftarrow \langle \boldsymbol{\phi}(\cdot,\cdot), \widehat{\mathbf{w}}^{\pi}_h\rangle, \quad \widehat{V}^{\pi}_h(\cdot) \leftarrow \langle \boldsymbol{\phi}^{\pi}_h(\cdot), \widehat{\mathbf{w}}^{\pi}_h\rangle$
11: **end for**
12: **Output:** $\widehat{v}^{\pi}_1 \leftarrow \int_{\mathcal{S}} \widehat{V}^{\pi}_1(s)\, \mathrm{d}\xi_1(s)$

Similarly, we estimate the second term in (3.4) by $\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\theta}}^{\pi}_h\rangle$, where $\widehat{\boldsymbol{\theta}}^{\pi}_h$ (Line 5) is given by

$$\widehat{\boldsymbol{\theta}}_h = \operatorname*{argmin}_{\boldsymbol{\theta}\in\mathbb{R}^d} \sum_{k=1}^{K} \left[\langle \check{\boldsymbol{\phi}}_{k,h}, \boldsymbol{\theta}\rangle - \widehat{V}^{\pi}_{h+1}(\check{s}'_{k,h})\right]^2 + \lambda\|\boldsymbol{\theta}\|_2^2 = \widehat{\boldsymbol{\Sigma}}_h^{-1} \sum_{k=1}^{K} \check{\boldsymbol{\phi}}_{k,h}\widehat{V}^{\pi}_{h+1}(\check{s}'_{k,h}), \tag{3.6}$$

with $\widehat{\boldsymbol{\Sigma}}_h = \sum_{k=1}^{K} \check{\boldsymbol{\phi}}_{k,h}\check{\boldsymbol{\phi}}_{k,h}^{\top} + \lambda \mathbf{I}_d$. Combining (3.5) and (3.6), we estimate $\mathbb{V}_h V^{\pi}_{h+1}$ by

$$[\widehat{\mathbb{V}}_h \widehat{V}^{\pi}_{h+1}](\cdot,\cdot) = \langle \boldsymbol{\phi}(\cdot,\cdot), \widehat{\boldsymbol{\beta}}^{\pi}_h\rangle_{[0,(H-h+1)^2]} - \left[\langle \boldsymbol{\phi}(\cdot,\cdot), \widehat{\boldsymbol{\theta}}^{\pi}_h\rangle_{[0,H-h+1]}\right]^2, \tag{3.7}$$

where the subscript $[0,(H-h+1)^2]$ denotes the clipping into the given range, and similar for the subscript $[0, H-h+1]$. We do such clipping due to the fact that $V^{\pi}_{h+1} \in [0, H-h]$. We add 1 to deal with the approximation error in $\widehat{V}^{\pi}_{h+1}$.

Based on $\widehat{\mathbb{V}}_h \widehat{V}^{\pi}_{h+1}$, the final variance estimator $\widehat{\sigma}_h(\cdot,\cdot)$ (Line 6) is defined as

$$\widehat{\sigma}_h(\cdot,\cdot) = \sqrt{\max\{1, \widehat{\mathbb{V}}_h \widehat{V}^{\pi}_{h+1}(\cdot,\cdot)\} + 1}.$$

In order to deal with the situation where $\widehat{\mathbb{V}}_h \widehat{V}^{\pi}_{h+1} < 0$ or is very close to 0, we take maximum between $\widehat{\mathbb{V}}_h \widehat{V}^{\pi}_{h+1}$ and 1. Also, to account for the noise in the observed rewards, we add an extra 1 which is an upper bound of the noise variance by Assumption 2.1.

Final estimator. Recursively repeat the above procedure for $h = H, H-1, \ldots, 1$, and we obtain $\widehat{V}_1$. Then the final estimator for v^{π}_1 (Line 12) is defined as $\widehat{v}^{\pi}_1 = \int_{\mathcal{S}} \widehat{V}^{\pi}_1(s)\, \mathrm{d}\xi_1(s)$.

Intuition behind $\boldsymbol{\Lambda}_h$. To illustrate the intuition behind the weighted covariance matrix Λ_h, here we provide some brief heuristics. Let $\{(s_{k,h}, a_{k,h}, s'_{k,h})\}_{k\in[K]}$ be i.i.d. samples such that $(s_{k,h}, a_{k,h}) \sim \nu$ for some distribution ν over $\mathcal{S}\times\mathcal{A}$ and $s'_{k,h} \sim \mathbb{P}_h(\cdot|s_{k,h}, a_{k,h})$. Define

$$\mathbf{e}_k = \boldsymbol{\phi}(s_{k,h}, a_{k,h}) \left([\mathbb{P}_h V^{\pi}_{h+1}](s_{k,h}, a_{k,h}) - V^{\pi}_{h+1}(s'_{k,h})\right) / \left[\mathbb{V}_h V^{\pi}_{h+1}\right](s_{k,h}, a_{k,h})^2$$

for all $k \in [K]$. Note that $\mathbf{e}_k$'s are i.i.d zero-mean random vectors and a simple calculation yields

$$\mathrm{Cov}(\mathbf{e}_k) = \mathbb{E}\left[[\mathbb{V}_h V^{\pi}_{h+1}](s_{k,h}, a_{k,h})^{-2} \boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}(s_{k,h}, a_{k,h})^{\top}\right].$$

This coincides with (2.6). Suppose $\mathrm{Cov}(\mathbf{e}_k) \succ 0$, then by the central limit theorem, it holds that

$$\frac{1}{\sqrt{K}} \sum_{k=1}^{K} \mathbf{e}_k \xrightarrow{d} \mathcal{N}(0, \mathrm{Cov}(\mathbf{e}_k)).$$

Therefore, $\mathrm{Cov}(\mathbf{e}_k)^{-1}$, or equivalently $\boldsymbol{\Lambda}_h^{-1}$, can be seen as the Fisher information matrix associated with the weighted product of the Bellman error and the feature vectors. This is a tighter characterization of the convergence rate than bounding $\mathbb{V}_h V^{\pi}_{h+1}$ by its naive upper bound $(H-h)^2$.

4 Theoretical Results

In this section, we introduce our main theoretical results and give an overview of the proof technique.

4.1 OPE Error Bound

Our main result is a refined average-case OPE analysis that yields a tighter error bound in Theorem 4.1. The proof is in Appendix D. To simplify the notation, we define:

$$C_{h,2} = \sum_{i=h}^{H} \frac{H-h+1}{\sqrt{2\iota_h}}, \quad C_{h,3} = \frac{(H-h+1)^2}{2}, \quad C_{h,4} = \left(\|\mathbf{\Lambda}_h\| \cdot \|\mathbf{\Lambda}_h^{-1}\|\right)^{1/2}.$$

Theorem 4.1. Set $\lambda = 1$. Under Assumptions 2.1, 2.3 and 2.5, if K satisfies

$$K \geq C \cdot C_3 \cdot d^2 \left[\log\left(\frac{dH^2K}{\kappa\delta}\right)\right]^2, \tag{4.1}$$

where C is some problem-independent universal constant and

$$C_3 := \max\left\{\max_{h\in[H]} \frac{C_{h,3}\cdot C_{h,2}^2}{8\iota_h^2}, \frac{H^4}{\kappa^2}, \frac{H^2}{\kappa^2}\cdot \max_{h\in[H]}\frac{C_{h,3}}{2}\cdot \max_{h\in[H]}\frac{C_{h,3}}{\iota_h}\right\},$$

then with probability at least $1-\delta$, the output of Algorithm 1 satisfies

$$|v_1^\pi - \widehat{v}_1^\pi| \leq C\cdot\left[\sum_{h=1}^{H}\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}\right]\cdot\sqrt{\frac{\log(16H/\delta)}{K}} + C\cdot C_4\cdot\log\left(\frac{16H}{\delta}\right)\cdot\left(\frac{1}{K^{3/4}}+\frac{1}{K}\right),$$

where $\mathbf{v}_h^\pi := \mathbb{E}_{\pi,h}[\boldsymbol{\phi}(s_h,a_h)]$ and $C_4 := \sum_{h=1}^{H}\sqrt{C_{h,4}\cdot C_{h,2}\cdot\frac{(H-h+1)d}{4\iota_h}\cdot\log\left(\frac{dH^2K}{\kappa\delta}\right)}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}$.

Theorem 4.1 suggests that Algorithm 1 provably achieves a tighter instance-dependent error bound for OPE than that in [10]. In detail, the dominant term in our bound is $\widetilde{\mathcal{O}}(\sum_{h=1}^{H}\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}/\sqrt{K})$, as compared to the $\widetilde{\mathcal{O}}(\sum_{h=1}^{H}(H-h+1)\|\mathbf{v}_h^\pi\|_{\mathbf{\Sigma}_h^{-1}}/\sqrt{K})$ term in [10]. By (2.5) and (2.6), our bound is at least as good as the latter since $\mathbf{\Sigma}_h \preceq [(H-h+1)^2+1]\mathbf{\Lambda}_h$. More importantly, it is instance-dependent and tight for the general class of linear MDPs: for those where $\mathbb{V}_h V_{h+1}^\pi$ is close to its crude upper bound $(H-h+1)^2$, our bound recovers the prior result. When $\mathbb{V}_h V_{h+1}^\pi$ is small, `VA-OPE` benefits from incorporating the variance information and our bound gets tightened accordingly.

Remark 4.2. Note that we do not require $\mathbb{V}_h V_{h+1}^\pi(s,a)$ to be uniformly small for all s, a, and h. From the bound and (2.6), as long as the variances are smaller than $(H-h+1)^2$ on average of $(s,a)\in\mathcal{S}\times\mathcal{A}$ and in sum of h, the bound is improved. It is also worth noting that the lower bound proved in [10] only holds for a subclass of linear MDPs with $\mathbb{V}_h V_{h+1}^\pi = \Omega((H-h+1)^2)$, and thus their minimax-optimality does not hold for general linear MDPs. For more detailed comparison we refer the reader to Appendix B.

Remark 4.3. Conceptually, the term $\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}$ serves as a more precise characterization of the distribution shift between the behavior policy $\bar{\pi}$ and the target policy π in a variance-aware manner. This enables our algorithm to utilize the data more effectively. Compared with online RL where one can sample new data, OPE is more 'data-hungry': one cannot decide the overall quality of the data. Thus it is especially beneficial to put more focus on targeted values with less uncertainty. This is also the intuitive reason why our algorithm can achieve a tighter error bound.

4.2 Overview of the Proof Technique

Here we provide an overview of the proof for Theorem 4.1. Due to the parallel estimation of the the value functions and their variances, the analysis of `VA-OPE` is much more challenging compared with that of `FQI-OPE`. As a result, we need to develop a novel proof technique. First, we have the following error decomposition.

Lemma 4.4. For any $h \in [H]$, let $\widehat{V}_h^\pi$ be the output of Algorithm 1. Then it holds that

$$V_h^\pi(s) - \widehat{V}_h^\pi(s) = \int_{\mathcal{A}} [\mathbb{P}_h(V_{h+1}^\pi - \widehat{V}_{h+1}^\pi)](s,a)\mathrm{d}\pi_h(a|s) + \lambda \boldsymbol{\phi}_h^\pi(s)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1}\mathbf{w}_h^\pi \tag{4.2}$$
$$+ \boldsymbol{\phi}_h^\pi(s)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1}\left[-\lambda \int_{\mathcal{S}} \left(V_{h+1}^\pi(s') - \widehat{V}_{h+1}^\pi(s')\right)\mathrm{d}\boldsymbol{\mu}_h(s') + \sum_{k=1}^K \boldsymbol{\phi}_{k,h}\widehat{\sigma}_{k,h}^{-2}\Delta_{k,h}\right],$$

where $\Delta_{k,h} = [\mathbb{P}_h\widehat{V}_{h+1}^\pi](s_{k,h}, a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h}) - \epsilon_{k,h}$. In particular, recall that $\widehat{v}_1^\pi = \mathbb{E}[\widehat{V}_1^\pi(s_1) \mid s_1 \sim \xi_1]$ and the OPE error can be decomposed as

$$v_1^\pi - \widehat{v}_1^\pi = -\lambda \sum_{h=1}^H (\mathbf{v}_h^\pi)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \int_{\mathcal{S}} \left(V_{h+1}^\pi(s) - \widehat{V}_{h+1}^\pi(s)\right)\boldsymbol{\mu}_h(s)\mathrm{d}s$$
$$+ \sum_{h=1}^H (\mathbf{v}_h^\pi)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \boldsymbol{\phi}_{k,h}\widehat{\sigma}_{k,h}^{-2}\Delta_{k,h} + \lambda \sum_{h=1}^H (\mathbf{v}_h^\pi)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1}\mathbf{w}_h^\pi. \tag{4.3}$$

The OPE error bound (Theorem 4.1) is proved by bounding the three terms separately in (4.3). This decomposition is different from [10] in that $\widehat{\boldsymbol{\Sigma}}_h$ is replaced by $\widehat{\boldsymbol{\Lambda}}_h$. This prevents us from adopting a matrix embedding-type proof as used in the prior work.

The key is to show the convergence of $\widehat{\boldsymbol{\Lambda}}_h$ to its population counterpart. However, by definition of $\widehat{\boldsymbol{\Lambda}}_h$, to establish such a result, it first requires the convergence of $\widehat{V}_{h+1}^\pi$ to V_{h+1}^π in a uniform manner, i.e., a high probability bound for $\sup_{s\in\mathcal{S}} |\widehat{V}_h^\pi(s) - V_h^\pi(s)|$. To show this, we leverage the decomposition in (4.2) and a backward induction technique, and prove a uniform convergence result which states that with high probability, for all $h \in [H]$, Algorithm 1 can guarantee

$$\sup_{s\in\mathcal{S}} \left|\widehat{V}_h^\pi(s) - V_h^\pi(s)\right| \le \widetilde{\mathcal{O}}\left(\frac{1}{\sqrt{K}}\right).$$

This result is formalized as Theorem C.2 and proved in Appendix C. To the best of our knowledge, Theorem C.2 is the first to establish the uniform convergence of the estimation error for the value functions in offline RL with linear function approximation. We believe this result is of independent interest and may be broadly useful in OPE.

5 Numerical Experiments

In this section, we provide numerical experiments to evaluate our algorithm `VA-OPE`, and compare it with `FQI-OPE`.

We construct a linear MDP instance as follows. The MDP has $|\mathcal{S}| = 2$ states and $|\mathcal{A}| = 100$ actions, with the feature dimension $d = 10$. The behavior policy then chooses action $a = 0$ with probability p and $a \in \{1, \cdots, 99\}$ with probability $1 - p$ and uniformly over $\{1, \cdots, 99\}$. The target policy π always chooses $a = 0$ no matter which state it is, making state 0 and 1 absorbing. The parameter p can be used to control the distribution shift between the behavior and target policies. Here $p \to 0$ leads to small distribution shift, and $p \to 1$ leads to large distribution shift. The initial distribution ξ_1 is uniform over $|\mathcal{S}|$. For more details about the construction of the linear MDP and parameter configuration, please refer to Appendix A.

We compare the performance of the two algorithms on the synthetic MDP described above under different choices of horizon length H. We plot the log-scaled OPE error versus $\sqrt{K}$ in Figure 1. It is clear that `VA-OPE` is at least as good as `FQI-OPE` in all the cases. Specifically, for small H (Figure 1a), their performance is very comparable, which is as expected. As H increases, we can see from Figure 1a, 1b and 1c that `VA-OPE` starts to dominate `FQI-OPE`, and the advantage is more significant for larger H, as suggested by Theorem 4.1. Due to space limit, a comprehensive comparison under different parameter settings is deferred to Appendix A.

6 Related Work

Off-policy evaluation. There is a large body of literature on OPE for tabular MDPs. Since the seminal work by Precup [33], various importance sampling-based estimators have been studied in

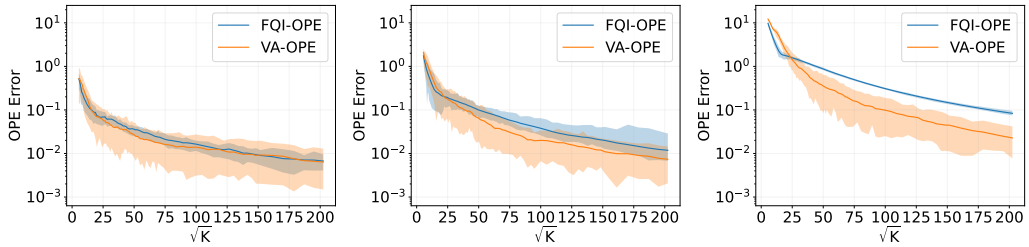

(a) $H = 5$. (b) $H = 10$. (c) $H = 30$.

Figure 1: Comparison of VA-OPE and FQI-OPE under different settings of horizon length H. VA-OPE's advantage becomes more significant as H increases, matching the theoretical prediction. The results are averaged over 50 trials and the error bars denote an empirical [10%,90%] confidence interval. The y-axis is log-scaled OPE error and x-axis is $\sqrt{K}$. For more details please see Appendix A.

the literature [26, 27, 38]. By using marginalized importance sampling methods [28, 43, 21, 45], one is able to further break the "curse-of-horizon". Moreover, various doubly robust estimators [11, 16, 12, 37] have been developed to achieve variance reduction. Most recently, it is shown by Yin et al. [46] that uniform convergence over all possible policy is also achievable. However, all the aforementioned works are limited to tabular MDPs. There is also a notable line of work on the estimation of the stationary distribution ratio between the target policy and the behavior policy using a primal-dual formulation [29, 48, 8]. However, a theoretical guarantee for the OPE error is not given in the work. More recently, Chen et al. [7] studied OPE in the *infinite-horizon* setting with linear function approximation.

There are many others topics related to OPE, for example, policy gradient [20, 30, 2], conservative policy iteration [19], off-policy temporal-difference learning [34], off-policy Q-learning [23], safe policy iteration [32] and pessimism in RL [22, 18], to mention a few. We refer the reader to the excellent survey by Levine et al. [25] for a more detailed introduction.

Online RL with linear function approximation. RL with function approximation has been actively studied as an extension of the tabular setting. Yang and Wang [44] studied discounted linear MDPs with a generative model, and Jin et al. [17] proposed an efficient LSVI-UCB algorithm for linear MDPs without a generative model. It has been shown by Du et al. [9] that MDP with misspecified linear function approximation could be exponentially hard to learn. Linear MDPs under various settings have also been studied by Zanette et al. [47], Neu and Pike-Burke [31] and He et al. [14].

A parallel line of work studies linear mixture MDPs [15, 3, 5, 50] (a.k.a., linear kernel MDPs [51]) where the transition kernel is a linear function of a ternary feature mapping $\psi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{R}^d$. In particular, Zhou et al. [50] achieved a nearly minimax regret bound by carefully utilizing the variance information of the value functions. Zhang et al. [49] constructed a variance-aware confidence set for time-homogeneous linear mixture MDPs. However, both works are focused on online RL rather than offline RL. It requires novel algorithm designs to exploit the variance information for offline tasks like OPE. What's more, the analysis in the offline setting deviates a lot from that for online RL where one can easily apply the law of total variance to obtain tighter bounds.

7 Conclusion and Future Work

In this paper, we incorporate the variance information into OPE and propose VA-OPE, an algorithm that provably achieves tighter error bound. Our $\widetilde{O}(\sum_h (\mathbf{v}_h^\top \mathbf{\Lambda}_h^{-1} \mathbf{v}_h)^{1/2}/\sqrt{K})$ error bound has a sharper dependence on the distribution shift between the behavior policy and the target policy.

Our work suggests several promising future directions. Theoretically, it remains open to provide an instance-dependent lower bound for the OPE error. Also, beyond the linear function approximation, it is interesting to establish similar results under more general function approximation schemes. Empirically, can we exploit the algorithmic insight of our algorithm to develop practically more

data-effective OPE algorithms for complex real-world RL tasks? We wish to explore these directions in the future.

Acknowledgments and Disclosure of Funding

We thank Mengdi Wang and Yaqi Duan for helpful discussions during the preparation of the paper. We also thank the anonymous reviewers for their helpful comments. DZ and QG are partially supported by the National Science Foundation CAREER Award 1906169, IIS-1904183 and AWS Machine Learning Research Award. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

References

[1] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.

[2] A. Agarwal, S. M. Kakade, J. D. Lee, and G. Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *arXiv preprint arXiv:1908.00261*, 2019.

[3] A. Ayoub, Z. Jia, C. Szepesvari, M. Wang, and L. Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pages 463–474. PMLR, 2020.

[4] F. Bertoluzzo and M. Corazza. Testing different reinforcement learning configurations for financial trading: Introduction and applications. *Procedia Economics and Finance*, 3:68–77, 2012.

[5] Q. Cai, Z. Yang, C. Jin, and Z. Wang. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pages 1283–1294. PMLR, 2020.

[6] D. Charles, M. Chickering, and P. Simard. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14, 2013.

[7] L. Chen, B. Scherrer, and P. L. Bartlett. Infinite-horizon offline reinforcement learning with linear function approximation: Curse of dimensionality and algorithm. *arXiv preprint arXiv:2103.09847*, 2021.

[8] B. Dai, O. Nachum, Y. Chow, L. Li, C. Szepesvári, and D. Schuurmans. Coindice: Off-policy confidence interval estimation. In *Advances in Neural Information Processing Systems*, 2020.

[9] S. S. Du, S. M. Kakade, R. Wang, and L. F. Yang. Is a good representation sufficient for sample efficient reinforcement learning? In *International Conference on Learning Representations*, 2020.

[10] Y. Duan, Z. Jia, and M. Wang. Minimax-optimal off-policy evaluation with linear function approximation. In *International Conference on Machine Learning*, pages 2701–2709. PMLR, 2020.

[11] M. Dudík, J. Langford, and L. Li. Doubly robust policy evaluation and learning. In *International Conference on Machine Learning*. PMLR, 2011.

[12] M. Farajtabar, Y. Chow, and M. Ghavamzadeh. More robust doubly robust off-policy evaluation. In *International Conference on Machine Learning*, pages 1447–1456. PMLR, 2018.

[13] D. A. Freedman. On tail probabilities for martingales. *the Annals of Probability*, pages 100–118, 1975.

[14] J. He, D. Zhou, and Q. Gu. Logarithmic regret for reinforcement learning with linear function approximation. In *International Conference on Machine Learning*. PMLR, 2021.

[15] Z. Jia, L. Yang, C. Szepesvari, and M. Wang. Model-based reinforcement learning with value-targeted regression. In *Learning for Dynamics and Control*, pages 666–686. PMLR, 2020.

[16] N. Jiang and L. Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661. PMLR, 2016.

[17] C. Jin, Z. Yang, Z. Wang, and M. I. Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020.

[18] Y. Jin, Z. Yang, and Z. Wang. Is pessimism provably efficient for offline rl? *arXiv preprint arXiv:2012.15085*, 2020.

[19] S. Kakade and J. Langford. Approximately optimal approximate reinforcement learning. In *In Proc. 19th International Conference on Machine Learning*. Citeseer, 2002.

[20] S. M. Kakade. A natural policy gradient. In *Advances in neural information processing systems*, volume 14, 2001.

[21] N. Kallus and M. Uehara. Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning. *arXiv preprint arXiv:1909.05850*, 2019.

[22] R. Kidambi, A. Rajeswaran, P. Netrapalli, and T. Joachims. Morel: Model-based offline reinforcement learning. *arXiv preprint arXiv:2005.05951*, 2020.

[23] A. Kumar, J. Fu, G. Tucker, and S. Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[24] S. Lange, T. Gabel, and M. Riedmiller. Batch reinforcement learning. In *Reinforcement learning*, pages 45–73. Springer, 2012.

[25] S. Levine, A. Kumar, G. Tucker, and J. Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.

[26] L. Li, W. Chu, J. Langford, and X. Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 297–306, 2011.

[27] L. Li, R. Munos, and C. Szepesvári. Toward minimax off-policy value estimation. In *Artificial Intelligence and Statistics*, pages 608–616. PMLR, 2015.

[28] Q. Liu, L. Li, Z. Tang, and D. Zhou. Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

[29] O. Nachum, Y. Chow, B. Dai, and L. Li. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Advances in Neural Information Processing Systems*, volume 32, pages 2318–2328. Curran Associates, Inc., 2019.

[30] O. Nachum, B. Dai, I. Kostrikov, Y. Chow, L. Li, and D. Schuurmans. Algaedice: Policy gradient from arbitrary experience. *arXiv preprint arXiv:1912.02074*, 2019.

[31] G. Neu and C. Pike-Burke. A unifying view of optimism in episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 33, 2020.

[32] M. Pirotta, M. Restelli, A. Pecorino, and D. Calandriello. Safe policy iteration. In *International Conference on Machine Learning*, pages 307–315. PMLR, 2013.

[33] D. Precup. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, page 80, 2000.

[34] D. Precup, R. S. Sutton, and S. Dasgupta. Off-policy temporal-difference learning with function approximation. In *International Conference on Machine Learning*, pages 417–424, 2001.

[35] D. Quillen, E. Jang, O. Nachum, C. Finn, J. Ibarz, and S. Levine. Deep reinforcement learning for vision-based robotic grasping: A simulated comparative evaluation of off-policy methods. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6284–6291. IEEE, 2018.

[36] L. Tang, R. Rosales, A. Singh, and D. Agarwal. Automatic ad format selection via contextual bandits. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 1587–1594, 2013.

[37] Z. Tang, Y. Feng, L. Li, D. Zhou, and Q. Liu. Doubly robust bias reduction in infinite horizon off-policy estimation. In *International Conference on Learning Representations*, 2020.

[38] P. Thomas and E. Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148. PMLR, 2016.

[39] P. S. Thomas, G. Theocharous, M. Ghavamzadeh, I. Durugkar, and E. Brunskill. Predictive off-policy policy evaluation for nonstationary decision problems, with applications to digital marketing. In *AAAI*, pages 4740–4745, 2017.

[40] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

[41] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.

[42] R. Wang, Y. Wu, R. Salakhutdinov, and S. M. Kakade. Instabilities of offline rl with pre-trained neural representation. In *International Conference on Machine Learning*. PMLR, 2021.

[43] T. Xie, Y. Ma, and Y.-X. Wang. Towards optimal off-policy evaluation for reinforcement learning with marginalized importance sampling. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[44] L. Yang and M. Wang. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, pages 6995–7004, 2019.

[45] M. Yin and Y.-X. Wang. Asymptotically efficient off-policy evaluation for tabular reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3948–3958. PMLR, 2020.

[46] M. Yin, Y. Bai, and Y.-X. Wang. Near-optimal provable uniform convergence in offline policy evaluation for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1567–1575. PMLR, 2021.

[47] A. Zanette, A. Lazaric, M. Kochenderfer, and E. Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pages 10978–10989. PMLR, 2020.

[48] R. Zhang, B. Dai, L. Li, and D. Schuurmans. Gendice: Generalized offline estimation of stationary values. In *International Conference on Learning Representations*, 2020.

[49] Z. Zhang, J. Yang, X. Ji, and S. S. Du. Variance-aware confidence set: Variance-dependent bound for linear bandits and horizon-free bound for linear mixture mdp. *arXiv preprint arXiv:2101.12745*, 2021.

[50] D. Zhou, Q. Gu, and C. Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*. PMLR, 2021.

[51] D. Zhou, J. He, and Q. Gu. Provably efficient reinforcement learning for discounted mdps with feature mapping. In *International Conference on Machine Learning*. PMLR, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]

(c) Did you discuss any potential negative societal impacts of your work? [N/A] Explain: our work is seeking to develop a mathematical understanding of the off-policy evaluation in reinforcement learning.

(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...

   (a) Did you state the full set of assumptions of all theoretical results? [Yes]

   (b) Did you include complete proofs of all theoretical results? [Yes]

3. If you ran experiments...

   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]

   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]

   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]

   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

   (a) If your work uses existing assets, did you cite the creators? [N/A]

   (b) Did you mention the license of the assets? [N/A]

   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Details of the Experiments

A.1 A Synthetic Linear MDP Example

We construct a synthetic linear MDP example based on a hard example proposed in Section 5 in [10], which was used to illustrate their lower bound. However, the feature dimension $d = 2$ in their illustrative example is too small to show discrepancy between our algorithm and theirs. Therefore, we construct an example sharing a similar structure but of much larger feature dimension and size of action space.

MDP instance. In specific, our MDP instance contains $|\mathcal{S}| = 2$ states and $|\mathcal{A}| = 100$ actions, and the feature dimension is $d = 10$. We denote $\mathcal{S} = \{0, 1\}$ and $\mathcal{A} = \{0, 1, \ldots, 99\}$ respectively. For each action $a \in [99]$, we represent it by a binary encoding vector $\mathbf{a} \in \mathbb{R}^8$ with each entry being either 1 or -1. With a slight abuse of notation, we interchangebly use a and and its vector representation $\mathbf{a}$.

We define

$$\delta(s, a) = \begin{cases} 1 & \text{if } \mathbb{1}\{s = 0\} = \mathbb{1}\{a = 0\}, \\ 0 & \text{otherwise.} \end{cases}$$

Then the feature mapping is given by

$$\boldsymbol{\phi}(s, a) = (\mathbf{a}^\top, \delta(s, a), 1 - \delta(s, a))^\top \in \mathbb{R}^{10}.$$

Let $\{\alpha_h\}_{h \in [H]}$ be a sequence of integers taking values in $\{0, 1\}$. For each $s \in \mathcal{S}$, the vector-valued measures are defined as

$$\boldsymbol{\mu}_h(s) = (0, \ldots, 0, (1 - s) \oplus \alpha_h, s \oplus \alpha_h)$$

for all $h \in [H]$, where $\oplus$ denotes the 'XOR' sign. Finally, we define $\boldsymbol{\gamma}_h \equiv \boldsymbol{\gamma} = (0, \ldots, 0, 1, 0) \in \mathbb{R}^{10}$. Thus the transition is $\mathbb{P}_h(s' \mid s, a) = \langle \boldsymbol{\phi}(s, a), \boldsymbol{\mu}_h(s') \rangle$ and the expected reward is $r_h(s, a) = \langle \boldsymbol{\phi}(s, a), \boldsymbol{\gamma} \rangle$. It is straightforward to verify that this is a valid time-inhomogeneous linear MDP.

Behavior and target policy. The target policy is given by $\pi(s) = 0$ for both $s = 0, 1$. The behavior policy is determined by a parameter $p \in (0, 1)$: with probability $1 - p$, the behavior policy chooses $a = 0$, and with probability $(1 - p)/99$ it chooses $a = i$ for each $i \in [99]$. This p can be used to control the distribution shift between the behavior and target policies. Note that p close to 0 induces small distribution shift, while larger p leads to large distribution shift. Moreover, we set the initial distribution ξ_1 to be uniform over $\mathcal{S}$.

We remark that in our implementation of `VA-OPE` we do *not* apply data splitting, i.e., $\mathcal{D} = \check{\mathcal{D}}$ and therefore no data is wasted. As is mentioned in the main text, the only purpose of the data splitting is to avoid an otherwise lengthy theoretical analysis. Therefore, for each fixed K, both algorithms use a dataset of size K sampled under the behavior policy.

A.2 Impact of the Planning Horizon

We first study the impact of the planning horizon H on the performance. We run our algorithm `VA-OPE` and the baseline method `FQI-OPE` with $\lambda = 1$ on the linear MDP instance constructed in the previous subsection under different values of H. We fix the initial distribution to be $\xi_1 = [1/2, 1/2]$ and p to be 0.6. The results are reported in Figure 2.

To explain the results, let us first recall the dominant term in our error bound and that in [10] (ignoring the logarithmic and constant factors):

$$D_{\texttt{VA}} = \frac{\sum_{h=1}^{H} \|\mathbf{v}_h\|_{\boldsymbol{\Lambda}_h^{-1}}}{\sqrt{K}} \qquad \text{vs} \qquad D_{\texttt{FQI}} = \frac{\sum_{h=1}^{H} (H - h + 1) \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Sigma}_h^{-1}}}{\sqrt{K}}. \tag{A.1}$$

As mentioned in the discussion following Theorem 4.1, it holds that $D_{\texttt{VA}} \leq D_{\texttt{FQI}}$. Indeed, this is reflected by the error plots where the error of `VA-OPE` is smaller than that of `FQI-OPE` except for very small K.

Moreover, as careful readers may have already observed, the discrepancy between $D_{\texttt{VA}}$ and $D_{\texttt{FQI}}$ would be amplified as the value of H increases. Again, our simulation results confirm this theoretical

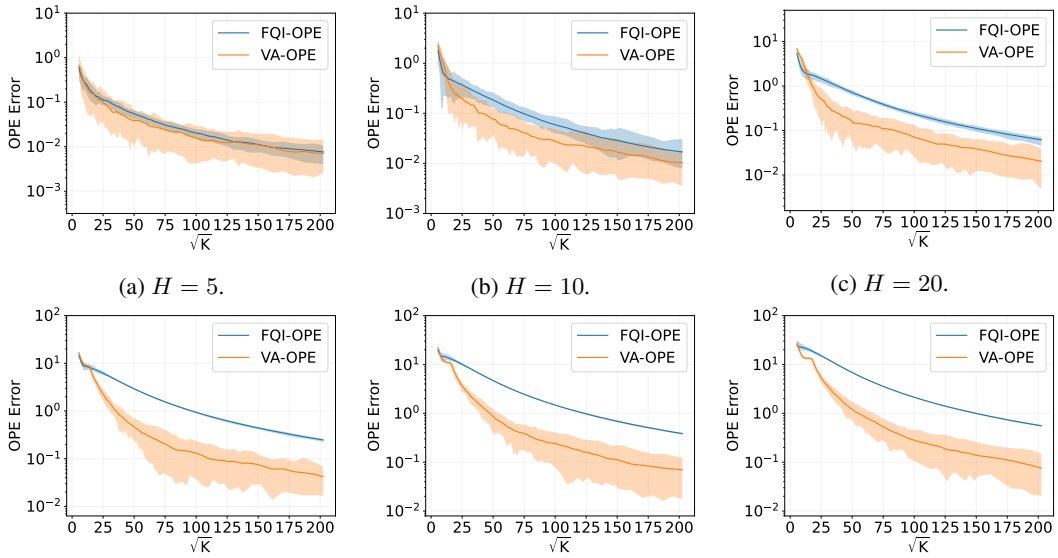

(a) $H = 5$. (b) $H = 10$. (c) $H = 20$.

(d) $H = 40$. (e) $H = 50$. (f) $H = 60$.

Figure 2: OPE error vs. $\sqrt{K}$. The results are averaged over 50 trials and the error bars are chosen to be the empirical [10%, 90%] confidence intervals. For a proper comparison, each sub-plot corresponds to a different setting of H by keeping everything else the same: $|\mathcal{S}| = 2$, $|\mathcal{A}| = 100$, $p = 0.6$.

observation as we can see by comparing the subplots of Figure 2. For larger values of H, `VA-OPE` tends to enjoy a much faster convergence rate. We would like to emphasize that this performance gain is especially beneficial for long-horizon tasks.

These findings also shed light on the minimax optimality of the OPE problem. The previous `FQI-OPE` algorithm is nearly minimax optimal only for a subclass of linear MDPs where $\mathbb{V}_h V^{\pi}_{h+1} = \Theta((H-h)^2)$. As suggested by our theory and confirmed by the numerical experiments, our algorithm `VA-OPE` achieves a tighter instance-dependent error for general linear MDPs. We would like to establish the universal minimax lower bound in the future work, and we believe that `VA-OPE` is a promising candidate for achieving minimax optimality.

We would also like to remark that the width of the error bars of `VA-OPE` is similar to that of `FQI-OPE`. It only appears wider on the plots since the y-axis is $\log_{10}$-scaled.

A.3 Impact of Distribution Shift

We also illustrate the impact of distribution shift between the behavior policy and the target policy on the performance, which can be controlled by the value of p. In Figure 3[3], we compare the performance of `VA-OPE` and `FQI-OPE` under different values of p.

The subplots in the same row share the same value of H. It is clear that for larger distribution shift, the performance of `VA-OPE` is superior. The reason behind this is that for fixed H, the ratio $D_{\texttt{FQI}}/D_{\texttt{VA}}$ increases as p increases. We further investigate this in the next subsection.

A.4 Comparison of the Dominant Terms

Finally we compare the dominant terms in the error upper bound of `VA-OPE` and `FQI-OPE` as defined in (A.1). Since both $D_{\texttt{VA}}$ and $D_{\texttt{FQI}}$ are theoretical values as the expectation over the occupancy measure induced by the transition kernel and the behavior/target policy, we simply estimate them by averaging over 1,000,000 independent trajectories. As presented in Figure 4, our characterization

[3]Note that the range of the y-axis differs among different rows.

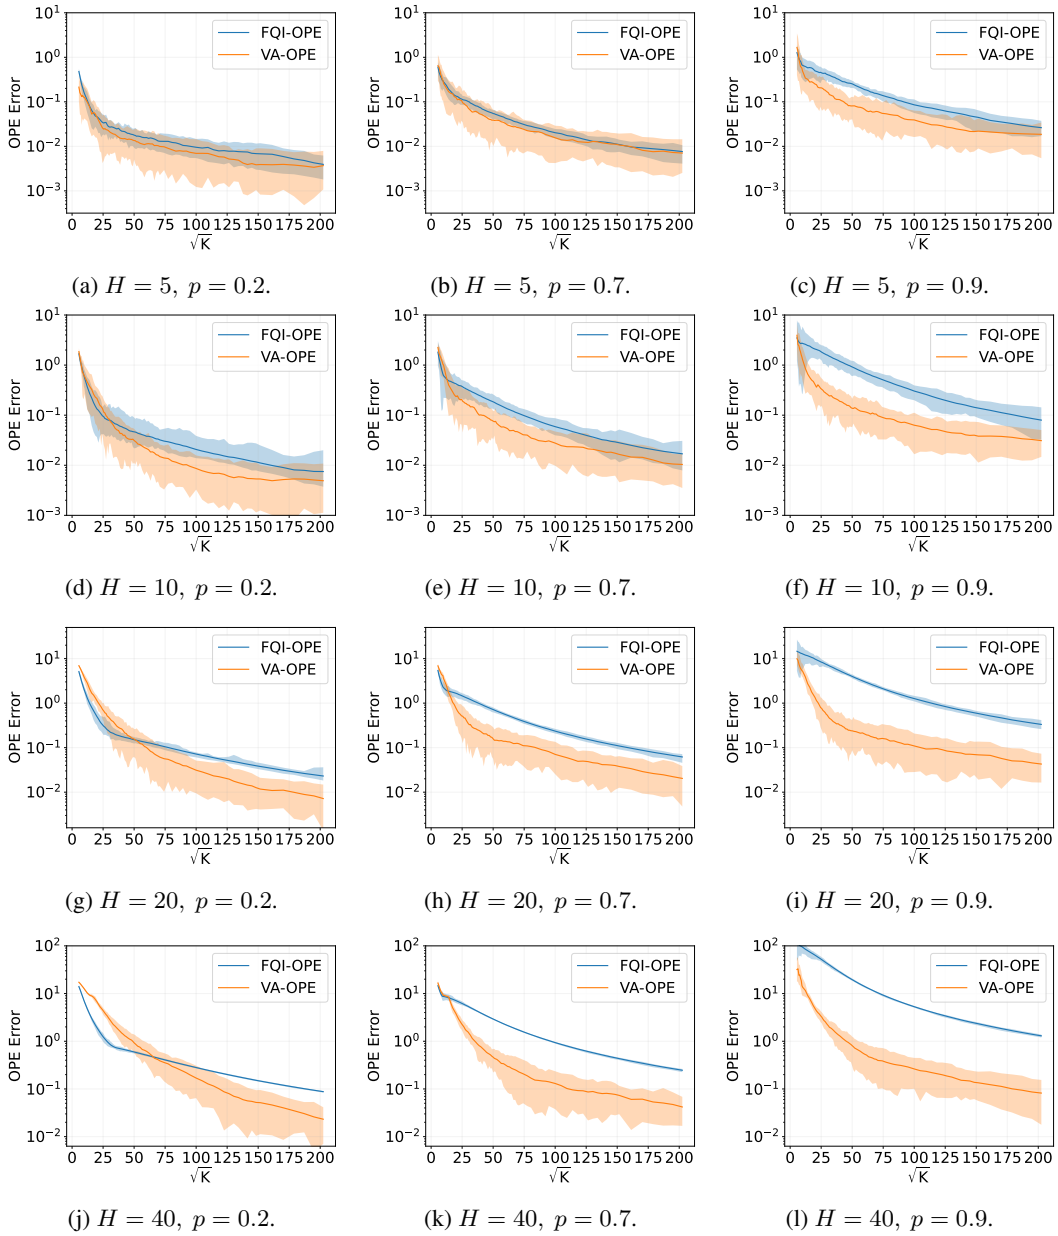

(a) $H = 5,\ p = 0.2.$ (b) $H = 5,\ p = 0.7.$ (c) $H = 5,\ p = 0.9.$

(d) $H = 10,\ p = 0.2.$ (e) $H = 10,\ p = 0.7.$ (f) $H = 10,\ p = 0.9.$

(g) $H = 20,\ p = 0.2.$ (h) $H = 20,\ p = 0.7.$ (i) $H = 20,\ p = 0.9.$

(j) $H = 40,\ p = 0.2.$ (k) $H = 40,\ p = 0.7.$ (l) $H = 40,\ p = 0.9.$

Figure 3: Log-scaled OPE error vs. $\sqrt{K}$ under different levels of distribution shift and horizon H. The level of distribution shift is controlled by the parameter p, where larger p corresponds to larger distribution shift.

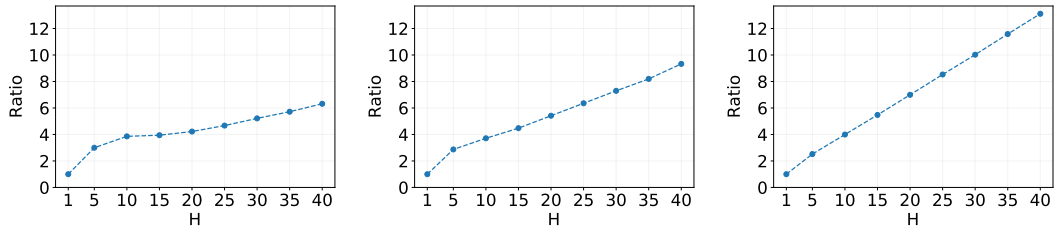

(a) $p = 0.3$. (b) $p = 0.6$. (c) $p = 0.9$.

Figure 4: Ratio between dominant terms vs. H. The results are generated by averaging over 1,000,000 trajectories.

of the distribution shift, $\sum_{h=1}^{H} \|\mathbf{v}_h\|_{\mathbf{\Lambda}_h^{-1}}$, is tighter. This is the main reason for the performance discrepancy that we have seen in the preceding subsections.

A.5 Hardware Details

All experiments are performed on an internal cluster with CPU and 30 GB of memory.

B Further Comparison with Duan et al. [10]

Consider the dominant term in the OPE error (omiting the logarithmic coefficients) given by Theorem 2 in Duan et al. [10], which was shown to be $\sum_{h=1}^{H}(H-h+1)\|\mathbf{v}_h^\pi\|_{\mathbf{\Sigma}_h^{-1}}/\sqrt{K}$ from their proof. As comparison, recall that our dominant term is about $\sum_{h=1}^{H} \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}/\sqrt{K}$. The definition of $\mathbf{\Sigma}_h$ in (2.5) and that of $\mathbf{\Lambda}_h$ in (2.6) immediately imply $\mathbf{\Sigma}_h \preceq [(H-h+1)^2+1] \cdot \mathbf{\Lambda}_h$ as σ_h^2 is bounded above by $(H-h+1)^2+1$. Therefore, it holds that

$$\sum_{h=1}^{H}(H-h+1)\|\mathbf{v}_h^\pi\|_{\mathbf{\Sigma}_h^{-1}} \geq \sum_{h=1}^{H} \frac{(H-h+1)\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}}{\sqrt{(H-h+1)^2+1}}. \tag{B.1}$$

The RHS of (B.1) is close to $\sum_{h=1}^{H} \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}$ if H is large. Moreover, when $\mathbb{V}_h V_{h+1}^\pi$ is small, the RHS of (B.1) can be much smaller than the LHS with appropriate choice of η_h. In other words, our bound is tighter than that of Duan et al. [10] in all scenarios, especially when $\mathbb{V}_h V_{h+1}^\pi$ is small.

Consider, for example, a scenario where the conditional variance of V_h^π, $h \in [H]$ is less than $H-h+1$, which is smaller than the crude upper bound of $(H-h+1)^2$ by a factor of $(H-h+1)$. Then by choosing $\eta_h = 1$ and $\sigma_r = 1$, we would have $\sigma_h^2 \equiv H-h+2$, and

$$\frac{(H-h+1)\|\mathbf{v}_h^\pi\|_{\mathbf{\Sigma}_h^{-1}}}{\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}} = \frac{H-h+1}{\sqrt{H-h+2}},$$

which suggests that $\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}$ is smaller than its counterpart by a factor of $(H-h+1)/\sqrt{H-h+2}$. Also, as mentioned in the main text, the conditional variance of V_h^π does **not** need to be uniformly smaller than $H-h+1$ for all $(s,a) \in \mathcal{S} \times \mathcal{A}$. It only needs to be small on average.

Regarding their lower bound (Theorem 3), it only holds for a subclass of all MDP instances where the conditional variance $\mathbb{V}_h V_{h+1}^\pi$ is on the order of $(H-h+1)^2$. Indeed, the theorem assumes there exists a high-value subset of states $\overline{\mathcal{S}}$ and a low-value subset of states $\underline{\mathcal{S}}$ under the target policy π such that $V_{h+1}^\pi(s) \geq \frac{3}{4}(H-h+1)$ if $s \in \overline{\mathcal{S}}$ and $V_{h+1}^\pi(s) \leq \frac{1}{4}(H-h+1)$ if $s \in \underline{\mathcal{S}}$. They also require there is non-zero probability $\overline{p} \geq c > 0$ and $\underline{p} \geq c > 0$ of transitting into $\overline{\mathcal{S}}$ and $\underline{\mathcal{S}}$ respectively. These assumptions immediately imply $\mathbb{V}_h V_{h+1}^\pi = \Omega((H-h+1)^2)$. Therefore, the prior result is only (nearly) minimax for a very small class of MDPs. This is confirmed by our numerical experiments in Appendix A where we compare the OPE error of `VA-OPE` and `FQI-OPE` under different settings of H. The results show that `VA-OPE`'s advantage over `FQI-OPE` increases as H becomes larger. It thus remains open to derive an instance-dependent lower bound that matches our upper bound.

C The Uniform Convergence Result

C.1 Important Remark

Throughout the appendix, we consider and analyze a slightly more general form of Algorithm 1. We now explain.

Recall that in (2.7), we define $\sigma_h(\cdot,\cdot)$ as

$$\sigma_h(s,a) = \sqrt{\max\{1, \mathbb{V}_h V_{h+1}^\pi(s,a)\} + 1},$$

and the corresponding estimator is given by

$$\widehat{\sigma}_h(\cdot,\cdot) \leftarrow \sqrt{\max\{1, \widehat{\mathbb{V}}_h \widehat{V}_{h+1}^\pi(\cdot,\cdot)\} + 1}.$$

Here taking maximum with 1 is to deal with the situation where $\widehat{\mathbb{V}}_h \widehat{V}_{h+1}^\pi(\cdot,\cdot)$ is close to zero or negative, and the second 1 is to account for the variance of the rewards. Now as a more general scheme, we replace both with adjustable parameters: η_h and σ_r^2 such that $\eta_h \geq 1$ and $0 \leq \sigma_r \leq 1$. Thereby, for each $h \in [H]$, we have

$$\widehat{\sigma}_h(\cdot,\cdot) \leftarrow \sqrt{\max\{\eta_h, \widehat{\mathbb{V}}_h \widehat{V}_{h+1}^\pi(\cdot,\cdot)\} + \sigma_r^2}\,.$$

We allow the flexibility of the choices of $\{\eta_h\}_{h\in[H]}$ and σ_r in part for generality and theoretical interests. These parameters will appear in the final results for the uniform convergence and the OPE error bound. The general algorithm is then presented as in Algorithm 2.

Algorithm 2 `VA-OPE` (general form)

Input: target policy $\pi = \{\pi_h\}_{h\in[H]}$, datasets $\mathcal{D} = \{\{(s_{k,h}, a_{k,h}, r_{k,h}, s'_{k,h})\}_{h\in[H]}\}_{k\in[K]}$ and $\check{\mathcal{D}} = \{\{(\check{s}_{k,h}, \check{a}_{k,h}, \check{r}_{k,h}, \check{s}'_{k,h})\}_{h\in[H]}\}_{k\in[K]}$, initial distribution ξ_1, $\widehat{\mathbf{w}}_{H+1}^\pi = 0$, λ, σ_r, $\{\eta_h\}_{h\in[H]}$
for $h = H, H-1, \ldots, 1$ **do**
  $\widehat{\boldsymbol{\Sigma}}_h \leftarrow \sum_{k=1}^K \check{\boldsymbol{\phi}}_{k,h}\check{\boldsymbol{\phi}}_{k,h}^\top + \lambda \mathbf{I}_d$
  $\widehat{\boldsymbol{\beta}}_h \leftarrow \widehat{\boldsymbol{\Sigma}}_h^{-1} \sum_{k=1}^K \check{\boldsymbol{\phi}}_{k,h} \widehat{V}_{h+1}^\pi(\check{s}'_{k,h})^2$
  $\widehat{\boldsymbol{\theta}}_h \leftarrow \widehat{\boldsymbol{\Sigma}}_h^{-1} \sum_{k=1}^K \check{\boldsymbol{\phi}}_{k,h} \widehat{V}_{h+1}^\pi(\check{s}'_{k,h})$
  $\widehat{\sigma}_h(\cdot,\cdot) \leftarrow \sqrt{\max\{\eta_h, \widehat{\mathbb{V}}_h \widehat{V}_{h+1}^\pi(\cdot,\cdot)\} + \sigma_r^2}$
  $\widehat{\boldsymbol{\Lambda}}_h \leftarrow \sum_{k=1}^K \boldsymbol{\phi}_{k,h}\boldsymbol{\phi}_{k,h}^\top / \widehat{\sigma}_{k,h}^2 + \lambda \mathbf{I}_d$
  $Y_{k,h} \leftarrow r_{k,h} + \langle \boldsymbol{\phi}_h^\pi(s'_{k,h}), \widehat{\mathbf{w}}_{h+1}^\pi \rangle$
  $\widehat{\mathbf{w}}_h^\pi \leftarrow \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \boldsymbol{\phi}_{k,h} Y_{k,h} / \widehat{\sigma}_{k,h}^2$
  $\widehat{Q}_h^\pi(\cdot,\cdot) \leftarrow \langle \boldsymbol{\phi}(\cdot,\cdot), \widehat{\mathbf{w}}_h^\pi \rangle, \quad \widehat{V}_h^\pi(\cdot) \leftarrow \langle \boldsymbol{\phi}_h^\pi(\cdot), \widehat{\mathbf{w}}_h^\pi \rangle$
end for
Output: $\widehat{v}_1^\pi \leftarrow \int_{\mathcal{S}} \widehat{V}_1^\pi(s)\, \mathrm{d}\xi_1(s)$

Correspondingly, throughout the appendix we redefine for each $h \in [H]$:

$$\sigma_h(s,a) = \sqrt{\max\{\eta_h, \mathbb{V}_h V_{h+1}^\pi(s,a)\} + \sigma_r^2}, \tag{C.1}$$

and thus $\boldsymbol{\Lambda}_h$ defined in (2.6) also becomes (η_h, σ_r^2)-related.

Besides generality, this is actually also meaningful, because let's consider, for example, a situation where the agent actually knows that the reward is deterministic (i.e. there is no noise in the observed reward). Then the agent can choose $\sigma_r = 0$ (though this will not give a huge boost to the OPE error bound since the determinant factor in $\|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}}$ is the variance $\mathbb{V}_h V_{h+1}^\pi$).

C.2 Recap of Notations

Before presenting the theorems and proof, let's walk through the algorithm and remind the readers of the notations.

Variance estimation Recall the dataset $\check{\mathcal{D}} = \{\check{\mathcal{D}}_h\}_{h\in[H]}$, where $\check{\mathcal{D}}_h = \{(\check{s}_{k,h}, \check{a}_{k,h}, \check{r}_{k,h}, \check{s}'_{k,h})\}_{k\in[K]}$. For each h, the dataset $\check{\mathcal{D}}_h$ is used to compute the function $\widehat{\sigma}_h(\cdot,\cdot)$, which is an estimator for the conditional variance of V^π_{h+1}. To be more clear, let go through the inner loop of Algorithm 2.

In the main text, due to the space limit, we use the abbreviation:

$$\check{\boldsymbol{\phi}}_{k,h} = \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h}).$$

For each h, the (biased and un-normalized) sample covariance matrix $\widehat{\boldsymbol{\Sigma}}_h$ (line 3) is given as

$$\widehat{\boldsymbol{\Sigma}}_h = \sum_{k=1}^{K} \check{\boldsymbol{\phi}}_{k,h}\check{\boldsymbol{\phi}}_{k,h}^\top + \lambda \mathbf{I}_d,$$

and its normalized population counterpart $\boldsymbol{\Sigma}_h$ is defined by (2.5) as

$$\boldsymbol{\Sigma}_h = \mathbb{E}_{\bar{\pi},h}\left[\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top\right].$$

Then the Algorithm computes $\widehat{\mathbb{V}}_h\widehat{V}^\pi_{h+1}$, which is an estimator of $\mathbb{V}_h V^\pi_{h+1}$, as the following:

$$[\widehat{\mathbb{V}}_h\widehat{V}^\pi_{h+1}](\cdot,\cdot) = \langle \boldsymbol{\phi}(\cdot,\cdot), \widehat{\boldsymbol{\beta}}^\pi_h\rangle_{[0,(H-h+1)^2]} - \left[\langle \boldsymbol{\phi}(\cdot,\cdot), \widehat{\boldsymbol{\theta}}^\pi_h\rangle_{[0,H-h+1]}\right]^2,$$

where $\widehat{\boldsymbol{\beta}}^\pi_h$ and $\widehat{\boldsymbol{\theta}}^\pi_h$ are computed in Algorithm 2 based on the estimated value function $\widehat{V}^\pi_{h+1}$ from last iteration, and the dataset $\check{D}_h$. Finally, the function $\widehat{\sigma}_h$ is computed.

Value function estimation Once we have the variance estimator $\widehat{\sigma}_h$, we can apply weighted regression to estimate the value function V^π_h using the dataset $\mathcal{D}_h = \{(s_{k,h}, a_{k,h}, r_{k,h}, s'_{k,h})\}_{k\in[K]}$. This is described by line 7 to 10 in Algorithm 2. Please note that we have adopt the abbreviation:

$$\boldsymbol{\phi}_{k,h} = \boldsymbol{\phi}(s_{k,h}, a_{k,h}),\ \widehat{\sigma}_{k,h} = \widehat{\sigma}_h(s_{k,h}, a_{k,h}).$$

Note that the weighted sample covariance matrix $\widehat{\boldsymbol{\Lambda}}_h$ in Algorithm 2 is given as

$$\widehat{\boldsymbol{\Lambda}}_h = \sum_{k=1}^{K} \boldsymbol{\phi}_{k,h}\boldsymbol{\phi}_{k,h}^\top/\widehat{\sigma}^2_{k,h} + \lambda \mathbf{I}_d,$$

with its normalized population counterpart $\boldsymbol{\Lambda}_h$ defined by (2.6) as

$$\boldsymbol{\Lambda}_h = \mathbb{E}_{\bar{\pi},h}\left[\sigma_h(s,a)^{-2}\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top\right].$$

Also note that in the offline dataset $\mathcal{D}$, for each $\mathcal{D}_h$ and the data point $(s_{k,h}, a_{k,h}, r_{k,h}, s'_{k,h})$ in $\mathcal{D}_h$, the reward $r_{k,h}$ is the random reward given by $r_{k,h} = r_h(s_{k,h}, a_{k,h}) + \epsilon_{k,h}$, where $r_h(\cdot,\cdot)$ is an unknown deterministic function representing the (conditional) mean and $\epsilon_{k,h}$ is some independent random noise. We only observe $r_{k,h}$ and not $\epsilon_{k,h}$.

Function classes Based on this characterization of the value functions, we define the following function class for each $h \in [H]$ and $L > 0$:

$$\mathcal{V}_h(L) := \left\{ V(s) = \langle \boldsymbol{\phi}^\pi_h(s), \mathbf{w}\rangle \middle| \mathbf{w} \in \mathbb{R}^d, \|\mathbf{w}\|_2 \le L,\ \sup_{s\in\mathcal{S}} |V(s)| \le H-h+2 \right\}. \tag{C.2}$$

One can see that functions in $\mathcal{V}_h(L)$ are parametrized by vectors $\mathbf{w} \in \mathbb{R}^d$. From Proposition 2.2, it is clear that $V^\pi_h \in \mathcal{V}_h(2H\sqrt{d})$ for all $h \in [H]$.

We define the following function class for each $h \in [H]$ and $L_1, L_2 > 0$:

$$\mathcal{T}_h(L_1, L_2) := \left\{ \sigma(\cdot,\cdot) = \sqrt{\max\left\{\eta_h, \langle \boldsymbol{\phi}(\cdot,\cdot), \boldsymbol{\beta}\rangle_{[0,(H-h+1)^2]} + \left[\langle \boldsymbol{\phi}(\cdot,\cdot), \boldsymbol{\theta}\rangle_{[0,H-h+1]}\right]^2\right\} + \sigma_r^2} \middle| \|\boldsymbol{\beta}\| \le L_1, \|\boldsymbol{\theta}\| \le L_2 \right\}, \tag{C.3}$$

which is parametrized by $\boldsymbol{\beta}, \boldsymbol{\theta} \in \mathbb{R}^d$. Later we will see that, with high probability, for all $h \in [H]$, we have $\widehat{\sigma}_h \in T_h(L_1, L_2)$ with above choice of $L_1 = H^2\sqrt{\frac{Kd}{\lambda}}$ and $L_2 = H\sqrt{\frac{Kd}{\lambda}}$, which is an immediate result by Theorem C.2 and Lemma H.15. Also note that $\sigma_h \in \mathcal{T}_h(L_1, L_2)$, which is clear from (3.4).

C.3 Formal Statement of Uniform Convergence Theorem

A weaker data sampling assumption. Recall Assumption 2.5 on the data sampling process introduced in the main text. It turns out that the uniform convergence result (Theorem C.2) holds under a weaker assumption which is the following.

Assumption C.1 (Trajectory-sampling Data)**.** We have two offline datasets $\mathcal{D}$ and $\check{\mathcal{D}}$ where each dataset consists of K trajectories with horizon length equal to H. Each trajectory is independently generated by the behavior policy $\bar{\pi}$. That is, $\mathcal{D} = \{\mathcal{D}_k\}_{k\in[K]}$, where each $\mathcal{D}_k$ is given by $\mathcal{D}_k = \{(s_{k,h}, a_{k,h}, r_{k,h})\}_{h\in[H]}$ such that $a_{k,h} \sim \pi_h(\cdot|s_{k,h})$ and $s_{k,h+1} \sim \mathbb{P}_h(\cdot|s_{k,h}, a_{k,h})$. For each $(k,h) \in [K] \times [H]$, the random reward $r_{k,h} = r_h(s_{k,h}, a_{k,h}) + \epsilon_{k,h}$, where $r_h(s_{k,h}, a_{k,h})$ is the (unknown) expected reward and $\epsilon_{k,h}$ is the noise. Similarly, we have $\check{\mathcal{D}} = \{\check{\mathcal{D}}_k\}_{k\in[K]}$, where $\check{\mathcal{D}}_k = \{(\check{s}_{k,h}, \check{a}_{k,h}, \check{r}_{k,h})\}_{h\in[H]}$. Here we denote $s'_{k,h} = s_{k,h+1}$ for simplicity.

Note that Assumption 2.5 is stronger than Assumption C.1 in the sense that Assumption 2.5 assumes an extra independence between the data points sampled at different stages. Therefore, as will be clear from the proof, since Theorem C.2 is established under Assumption C.1, it automatically holds under the stronger Assumption 2.5.

We now introduce the uniform convergence theorem. To simplify the notation, we define:

$$C_{h,1} = \sum_{i=h}^{H} \frac{1}{\iota_h}, \quad C_{h,2} = \sum_{i=h}^{H} \sqrt{\frac{C_{h,3}}{\iota_h}}, \quad C_{h,3} = \frac{(H-h+1)^2}{\eta_h + \sigma_r^2}, \quad C_{h,4} = \left(\|\mathbf{\Lambda}_h\| \cdot \|\mathbf{\Lambda}_h^{-1}\|\right)^{1/2}.$$

Note that by setting $\eta_h = \sigma_r = 1$ we recover the same $C_{h,2}$, $C_{h,3}$ as in the main text.

Theorem C.2 (Uniform Convergence)**.** Set $\lambda = 1$ and $\eta_h \in (0, (H-h+1)^2]$ for all $h \in [H]$ in Algorithm 1. Under Assumption 2.1, 2.3 and C.1, there exists some universal constant C such that if K satisfies

$$K \geq C \cdot \frac{H^2 d^2}{\kappa^2} \log\left(\frac{dHK}{\kappa\delta}\right) \cdot \max_{h\in[H]} \frac{(H-h+1)^2}{(\eta_h + \sigma_r^2)^2} \cdot \max_{h\in[H]} \frac{(H-h+1)^2}{\iota_h(\eta_h + \sigma_r^2)}, \tag{C.4}$$

then with probability at least $1-\delta$, it holds for all $h \in [H]$ that $\sup_{s\in\mathcal{S}} \left|\widehat{V}_h^\pi(s)\right| \leq H - h + 2$, and

$$\sup_{s\in\mathcal{S}} \left|\widehat{V}_h^\pi(s) - V_h^\pi(s)\right| \leq C \cdot \frac{C_{h,2} d}{\sqrt{K}} \log\left(\frac{dH^2K}{\kappa\delta}\right) + C \cdot \frac{C_{h,1} H\sqrt{d}}{K}.$$

We now present the proof of Theorem C.2. The proof relies on a backward induction argument, i.e., we will show $|\widehat{V}_h^\pi(s) - V_h^\pi(s)|$ is uniformly small for $h = H, H-1, \cdots, 1$. For this purpose, we need to use the first form of error decomposition (4.2) in Lemma 4.4.

C.4 Step 1: Base Case at Stage $h = H$

We first bound the approximation error at the last stage $h = H$. From the algorithm we have $\widehat{V}_{H+1}^\pi \equiv V_{H+1}^\pi \equiv 0$. Therefore, we have $\widehat{\boldsymbol{\theta}}_H = \widehat{\boldsymbol{\beta}}_H = 0$, $\widehat{\sigma}_H \equiv \sqrt{\eta_H + \sigma_r^2}$, and

$$\widehat{\mathbf{\Lambda}}_H = \frac{1}{\eta_H + \sigma_r^2} \sum_{k=1}^{K} \boldsymbol{\phi}(s_{k,H}, a_{k,H}) \boldsymbol{\phi}(s_{k,H}, a_{k,H})^\top + \lambda \mathbf{I}_d.$$

By the error decomposition in (4.2), we have

$$V_H^\pi(s) - \widehat{V}_H^\pi(s) = \underbrace{-\boldsymbol{\phi}_H^\pi(s)^\top \widehat{\mathbf{\Lambda}}_H^{-1} \sum_{k=1}^{K} \frac{\boldsymbol{\phi}(s_{k,H}, a_{k,H})}{\widehat{\sigma}_H(s_{k,H}, a_{k,H})^2} \epsilon_{k,H}}_{\Delta_1} + \underbrace{\lambda \boldsymbol{\phi}_H^\pi(s)^\top \widehat{\mathbf{\Lambda}}_H^{-1} \mathbf{w}_H^\pi}_{\Delta_2}. \tag{C.5}$$

We will bound the two terms separately.

To bound $|\Delta_1|$, we first apply Cauchy-Schwartz inequality to obtain that

$$|\Delta_1| \leq \|\boldsymbol{\phi}_H^\pi(s)\|_{\widehat{\mathbf{\Lambda}}_H^{-1}} \cdot \left\| \sum_{k=1}^{K} \frac{\boldsymbol{\phi}(s_{k,H}, a_{k,H})}{\widehat{\sigma}_H(s_{k,H}, a_{k,H})^2} \epsilon_{k,H} \right\|_{\widehat{\mathbf{\Lambda}}_H^{-1}} \tag{C.6}$$

By Lemma H.5, with probability at least $1-\delta$, we have

$$\|\boldsymbol{\phi}_H^\pi(s)\|_{\widehat{\boldsymbol{\Lambda}}_H^{-1}} \le \frac{2}{\sqrt{K}} \cdot \|\boldsymbol{\phi}_H^\pi(s)\|_{\boldsymbol{\Lambda}_H^{-1}} \tag{C.7}$$

for all $s \in \mathcal{S}$, as long as K satisfies that

$$K \ge \max\left\{\frac{512\|\boldsymbol{\Lambda}_H^{-1}\|^2}{(\eta_H+\sigma_r^2)^2}\log\left(\frac{2d}{\delta}\right), 2\lambda\|\boldsymbol{\Lambda}_H^{-1}\|\right\} = \frac{512\|\boldsymbol{\Lambda}_H^{-1}\|^2}{(\eta_H+\sigma_r^2)^2}\log\left(\frac{2d}{\delta}\right).$$

Note that $\mathrm{Var}(\epsilon_{k,h}) = \sigma_r^2 \le \widehat{\sigma}_H(s_{k,H}, a_{k,H})^2$ for all $k \in [K]$. Then by Theorem H.10 we have

$$\begin{aligned}
&\left\|\sum_{k=1}^K \widehat{\sigma}_H(s_{k,H}, a_{k,H})^{-2}\boldsymbol{\phi}(s_{k,H}, a_{k,H})\epsilon_{k,H}\right\|_{\widehat{\boldsymbol{\Lambda}}_H^{-1}} \\
&\le 8\sqrt{d\log\left(1+\frac{K}{\lambda d(\eta_H+\sigma_r^2)}\right)\cdot\log\left(\frac{4K^2}{\delta}\right)} + 4\sqrt{\frac{1}{\eta_H+\sigma_r^2}}\log\left(\frac{4K^2}{\delta}\right) \\
&\le 12\sqrt{d}\log\left(\frac{4K^2}{\delta}\right)
\end{aligned} \tag{C.8}$$

with probability at least $1-\delta$.

Then by (C.6), it suffices to take a union bound over (C.7) and (C.8) to conclude that if $K \ge 128\|\boldsymbol{\Lambda}_H^{-1}\|^2\log(2d/\delta)/(\eta_H+\sigma_r^2)$ then

$$|\Delta_1| \le \frac{12\sqrt{d}}{\sqrt{K}}\|\boldsymbol{\phi}_H^\pi(s)\|_{\boldsymbol{\Lambda}_H^{-1}} \cdot \log\left(\frac{4K^2}{\delta}\right) \tag{C.9}$$

with probability at least $1-\delta$.

At the same time, we can bound $|\Delta_2|$ using the same argument.

$$|\Delta_2| \le \lambda\|\boldsymbol{\phi}_H^\pi(s)\|_{\widehat{\boldsymbol{\Lambda}}_H^{-1}} \cdot \|\mathbf{w}_H^\pi\|_{\widehat{\boldsymbol{\Lambda}}_H^{-1}} \le \frac{4\lambda}{K} \cdot \|\boldsymbol{\phi}_H^\pi(s)\|_{\boldsymbol{\Lambda}_H^{-1}} \cdot \|\mathbf{w}_H^\pi\|_{\boldsymbol{\Lambda}_H^{-1}}, \tag{C.10}$$

where the second inequality holds on the same event as does (C.9).

Finally, we combine (C.5), (C.9) and (C.10), and obtain that if $K \ge 512\|\Lambda_H^{-1}\|^2/(\eta_H+\sigma_r^2)^2\log(4d/\delta)$ then

$$\begin{aligned}
\sup_{s\in\mathcal{S}}\left|V_H^\pi(s) - \widehat{V}_H^\pi(s)\right| &\le \frac{12\sqrt{d}}{\sqrt{K}}\log\left(\frac{4K^2}{\delta}\right)\cdot\sup_{s\in\mathcal{S}}\|\boldsymbol{\phi}_H^\pi(s)\|_{\boldsymbol{\Lambda}_H^{-1}} + \frac{4\lambda}{K}\cdot\sup_{s\in\mathcal{S}}\|\boldsymbol{\phi}_H^\pi(s)\|_{\boldsymbol{\Lambda}_H^{-1}}\cdot\|\mathbf{w}_H^\pi\|_{\boldsymbol{\Lambda}_H^{-1}} \\
&\le \frac{12\sqrt{d\|\boldsymbol{\Lambda}_H^{-1}\|}}{\sqrt{K}}\log\left(\frac{4K^2}{\delta}\right) + \frac{8\lambda H\sqrt{d}\|\boldsymbol{\Lambda}_H^{-1}\|}{K}
\end{aligned}$$

with probability at least $1-\delta$, where the last inequality follows from Assumption 2.1, Proposition 2.2 and the choice that $\lambda = 1$. Note that since $\widehat{\sigma}_H(\cdot,\cdot) \le 1+\sigma_r^2$, we have $\boldsymbol{\Lambda}_H \succeq \boldsymbol{\Sigma}_H/(1+\sigma_r^2)$, which implies that $\|\boldsymbol{\Lambda}_H^{-1}\| \le 2\|\boldsymbol{\Sigma}_H^{-1}\|$ as $\sigma_r^2 \le 1$. Then we further have

$$\sup_{s\in\mathcal{S}}\left|V_H^\pi(s) - \widehat{V}_H^\pi(s)\right| \le \frac{12\sqrt{2d}}{\sqrt{K\kappa_H}} + \frac{16\lambda H^3\sqrt{d}}{K\kappa_H}$$

Meanwhile, we can bound $\sup_{s\in\mathcal{S}}|\widehat{V}_H^\pi(s)|$ as follows

$$\sup_{s\in\mathcal{S}}|\widehat{V}_H^\pi(s)| \le \sup_{s\in\mathcal{S}} V_H^\pi(s) + \frac{12\sqrt{2d}}{\sqrt{K\kappa_H}} + \frac{16\lambda H\sqrt{d}}{K\kappa_H} \le 2,$$

when K satisfies that $K \ge 600(\lambda+1)(d+H\sqrt{d})/\kappa_H$.

In conclusion, we have

$$\sup_{s\in\mathcal{S}}|\widehat{V}_H^\pi(s)| \le 2,$$

and

$$\sup_{s\in\mathcal{S}}\left|V_H^\pi(s)-\widehat{V}_H^\pi(s)\right|\le\frac{12\sqrt{2d}}{\sqrt{K\kappa_H}}+\frac{16\lambda H\sqrt{d}}{K\kappa_H},$$

given that K satisfies

$$K\ge\max\left\{\frac{2048}{\kappa_H^2(\eta_H+\sigma_r^2)^2}\log\left(\frac{2d}{\delta}\right),600(\lambda+1)\frac{d+H\sqrt{d}}{\kappa_H}\right\}\tag{C.11}$$

C.5 Step 2: Induction Hypothesis

For the induction hypothesis, we assume that if for all sufficiently large K, with probability at least $1-(H-h)\delta$, the following event (denoted as $\mathcal{E}_{h+1}$) holds:

$$\sup_{s\in\mathcal{S}}|\widehat{V}_{h+2}^\pi(s)|\le H-h,\ \sup_{s\in\mathcal{S}}|\widehat{V}_{h+1}^\pi(s)|\le H-h+1,\ \sup_{s}|\widehat{V}_{h+1}^\pi(s)-V_{h+1}^\pi(s)|\le\alpha_{H-h},$$

where $\alpha_{H-h}\le(\eta_h+\sigma_r^2)/[8(H-h+1)]$.

We claim that if K satisfies

$$K\ge\frac{3600(H-h+1)^4d^2}{\kappa_h^2(\eta_h+\sigma_r^2)^2}\cdot\log\left(\frac{dHK}{\kappa_h\delta}\right)\tag{C.12}$$

then with probability at least $1-(H-h+1)\delta$, the following event (denoted by $\mathcal{E}_h$) holds:

$$\begin{aligned}
|\widehat{V}_{h+1}^\pi|&<H-h+1,\\
|\widehat{V}_h^\pi|&<H-h+2,\\
\sup_s|\widehat{V}_h^\pi(s)-V_h^\pi(s)|&\le\left(1+\frac{8\lambda}{\iota_hK}\right)\alpha_{H-h}+\frac{2\lambda H\sqrt{d}}{\iota_hK}\\
&\quad+\frac{20}{\sqrt{K}}\cdot\left(\frac{d}{\sqrt{\iota_h}}+\frac{d(H-h+1)}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right)\cdot\log\left(\frac{d(H-h+1)^2K}{\kappa_h(\eta_h+\sigma_r^2)\delta}\right)
\end{aligned}$$

We again bound the three terms in the error decomposition (4.2) simultaneously. Let $\widetilde{\mathcal{E}}_h$ be the event given by Lemma F.7 for h such that $\mathbb{P}(\widetilde{\mathcal{E}}_h)\ge1-\delta$, where we have $L=(1+1/H)d\sqrt{K/\lambda}$.

Let's consider the event $\widetilde{\mathcal{E}}_h\cap\mathcal{E}_{h+1}$, which satisfies $\mathbb{P}\{\widetilde{\mathcal{E}}_h\cap\mathcal{E}_{h+1}\}\ge1-(H-h+1)\delta$ by a union bound. Note that on $\mathcal{E}_{h+1}$, we have $|\widehat{V}_{h+1}^\pi|\le H-h+1$. Furthermore, since $|\widehat{V}_{h+2}^\pi|\le H-h$ on $\mathcal{E}_{h+1}$, again by Lemma H.15 with $B=H$, we see that $\widehat{V}_{h+1}^\pi\in\mathcal{V}_{h+1}(L)$. Therefore, by Lemma F.7 it holds on $\widetilde{\mathcal{E}}_h\cap\mathcal{E}_{h+1}$ that

$$\left\|\left(\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right)^{-1}\right\|\le\frac{8}{\iota_h},\tag{C.13}$$

and

$$\begin{aligned}
&\left|\boldsymbol{\phi}(s,a)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\sum_{k=1}^K\widehat{\sigma}_h(s_{k,h},a_{k,h})^{-2}\boldsymbol{\phi}(s_{k,h},a_{k,h})\left(\mathbb{P}_hV(s_{k,h},a_{k,h})-V(s'_{k,h})-\epsilon_{k,h}\right)\right|\\
&\le\frac{20}{\sqrt{K}}\cdot\left(\frac{d}{\sqrt{\iota_h}}+\frac{d(H-h+1)}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right)\cdot\log\left(\frac{d(H-h+1)^2K}{\kappa_h(\eta_h+\sigma_r^2)\delta}\right),
\end{aligned}\tag{C.14}$$

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

Since it holds on $\mathcal{E}_{h+1}$ that $\sup_{s\in\mathcal{S}}|\widehat{V}_{h+1}^\pi(s)-V_{h+1}^\pi(s)|\le\alpha_{H-h}$, we have

$$\sup_{s\in\mathcal{S}}\left|[\mathbb{J}_h\mathbb{P}_h(V_{h+1}^\pi-\widehat{V}_{h+1}^\pi)](s)\right|\le\alpha_{H-h}.\tag{C.15}$$

Also by (C.13) and $\sup_{s\in\mathcal{S}} |\widehat{V}^{\pi}_{h+1}(s) - V^{\pi}_{h+1}(s)| \le \alpha_{H-h}$, it follows from Cauchy-Schwartz inequality that

$$\sup_{s\in\mathcal{S}} \left| \lambda \boldsymbol{\phi}^{\pi}_h(s)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \int_{\mathcal{S}} \left(V^{\pi}_{h+1}(s) - \widehat{V}^{\pi}_{h+1}(s) \right) \boldsymbol{\mu}_h(s) \mathrm{d}s \right| \le \frac{8\lambda \alpha_{H-h}}{\iota_h K}. \tag{C.16}$$

Similarly, for the last term in (4.2) we have by Cauchy-Schwartz inequality that

$$\sup_{s\in\mathcal{S}} \left| \lambda \boldsymbol{\phi}^{\pi}_h(s)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \mathbf{w}^{\pi}_h \right| \le \lambda \sup_{s\in\mathcal{S}} \|\boldsymbol{\phi}^{\pi}_h(s)\|_2 \cdot \|\widehat{\boldsymbol{\Lambda}}_h^{-1}\| \cdot \|\mathbf{w}^{\pi}_h\|_2 \le \frac{2\lambda H\sqrt{d}}{\iota_h K}, \tag{C.17}$$

where the second inequality follows from Proposition 2.2.

Finally, combining (C.14), (C.15), (C.16) and (C.17), we obtain by the error decomposition (4.2) that

$$\begin{aligned} \sup_{s\in\mathcal{S}} \left| V^{\pi}_h(s) - \widehat{V}^{\pi}_h(s) \right| \le \frac{20}{\sqrt{K}} \cdot \left(\frac{d}{\sqrt{\iota_h}} + \frac{d(H-h+1)}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}} \right) \cdot \log\left(\frac{d(H-h+1)^2 K}{\kappa_h(\eta_h+\sigma_r^2)\delta} \right) \\ + \left(1 + \frac{8\lambda}{\iota_h K}\right) \alpha_{H-h} + \frac{2\lambda H\sqrt{d}}{\iota_h K} \end{aligned} \tag{C.18}$$

Note that when K satisfies (C.12), we would have $\sup_{s\in\mathcal{S}} |V^{\pi}_h(s) - \widehat{V}^{\pi}_h(s)| \le 1$, and thus

$$\sup_{s\in\mathcal{S}} |\widehat{V}^{\pi}_h(s)| \le \sup_{s\in\mathcal{S}} |V^{\pi}_h(s)| + \sup_{s\in\mathcal{S}} \left| V^{\pi}_h(s) - \widehat{V}^{\pi}_h(s) \right| \le H - h + 2. \tag{C.19}$$

Therefore, by (C.18) and (C.19), we conclude that $\widetilde{\mathcal{E}}_h \cap \mathcal{E}_{h+1} \subseteq \mathcal{E}_h$, which implies that

$$\mathbb{P}\{\mathcal{E}_h\} \ge \mathbb{P}\{\widetilde{\mathcal{E}}_h \cap \mathcal{E}_{h+1}\} \ge 1 - (H-h+1)\delta.$$

C.6 Step 3: Recursion

Let $\kappa = \min_{h\in[H]} \kappa_h$. Suppose K satisfies that

$$K \ge \frac{3600 H^2 d^2}{\kappa^2} \log\left(\frac{dHK}{\kappa\delta}\right) \cdot \max_{h\in[H]} \frac{(H-h+1)^2}{(\eta_h+\sigma_r^2)^2} \cdot \max_{h\in[H]} \frac{(H-h+1)^2}{\iota_h(\eta_h+\sigma_r^2)} \tag{C.20}$$

We also define the following quantity

$$\xi_{H-h} = \frac{2\lambda H\sqrt{d}}{\iota_h K} + \frac{20}{\sqrt{K}} \cdot \left(\frac{d}{\sqrt{\iota_h}} + \frac{d(H-h+1)}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}} \right) \cdot \log\left(\frac{d(H-h+1)^2 K}{\kappa_h(\eta_h+\sigma_r^2)\delta} \right) \tag{C.21}$$

for all $h \in [H]$. With the choice of K in (C.20), the following holds

$$\xi_{H-h} \le \min\left\{ \frac{1}{2He}, \frac{1}{16H} \cdot \min_{h\in[H]} \frac{\eta_h + \sigma_r^2}{H-h+1} \right\} \tag{C.22}$$

for all $h \in [H]$.

First by the base case at stage $h = H$ from subsection C.4, we have

$$\sup_{s\in\mathcal{S}} \left| V^{\pi}_H(s) - \widehat{V}^{\pi}_H(s) \right| \le \frac{12\sqrt{2d}}{\sqrt{K\kappa_H}} + \frac{16\lambda H\sqrt{d}}{K\kappa_H} := \alpha_0. \tag{C.23}$$

Also by the choice of K in (C.20), we have

$$\alpha_0 \le \min\left\{ \frac{1}{2e}, \frac{1}{16e} \cdot \min_{h\in[H]} \frac{\eta_h + \sigma_r^2}{H-h+1} \right\}. \tag{C.24}$$

Then by the induction step at stage $h = H-1$ from subsection C.5, we have

$$\sup_{s\in\mathcal{S}} \left| V_{H-1}(s) - \widehat{V}^{\pi}_{H-1}(s) \right| \le \alpha_1,$$

with

$$\alpha_1 \le \left(1 + \frac{\lambda C_{H-1} H}{K}\right)\alpha_0 + \xi_1 \le \left(1 + \frac{1}{H}\right)\alpha_0 + \xi_1 \le \min\left\{1, \frac{\eta_H + \sigma_r^2}{16}\right\}$$

We then define $\alpha_{H-h} = (1 + 1/H)\alpha_{H-h+1} + \xi_{H-h}$ recursively for all $h \in [H-1]$. Note that for all $i \in [H-h]$ we have

$$\alpha_i \le \left(1 + \frac{1}{H}\right)^i \alpha_0 + \sum_{j=0}^{i}\left(1 + \frac{1}{H}\right)^{i-j}\xi_j \le e \cdot \alpha_0 + e \cdot \sum_{h=0}^{i}\xi_j \le \min\left\{1, \frac{\eta_{H-i} + \sigma_r^2}{8(i+1)}\right\},$$

where the first inequality follows from the fact that $(1 + 1/n)^n \le e$ for all positive integer n, and the second inequality is due to (C.22) and (C.24).

Therefore, we may apply the induction step from the previous subsection to all $h \in [H-1]$ and obtain that

$$\sup_{s\in\mathcal{S}}\left|\widehat{V}_h^\pi\right| \le H - h + 2$$

and

$$\begin{aligned}\sup_{s\in\mathcal{S}}\left|\widehat{V}_h^\pi(s) - V_h^\pi(s)\right| &\le \left(1 + \frac{1}{H}\right)\alpha_{H-h} + \iota_{H-h} \\ &\le \left(1 + \frac{1}{H}\right)^{H-h}\alpha_0 + \sum_{i=0}^{H-h}\left(1 + \frac{1}{H}\right)^{H-h-i}\xi_i \\ &\le e \cdot \alpha_0 + e \cdot \sum_{i=0}^{H-h}\xi_i \end{aligned} \tag{C.25}$$

with probability at least $1 - H\delta$ simultaneously for all $h \in [H]$.

Therefore, replacing δ by δ/H and plugging (C.21) and (C.23) into (C.25), we obtain that

$$\begin{aligned}&\sup_{s\in\mathcal{S}}\left|\widehat{V}_h^\pi(s) - V_h^\pi(s)\right| \\ &\le \frac{12e\sqrt{2d}}{\sqrt{K\kappa_H}} + \frac{16e\lambda H\sqrt{d}}{K\kappa_H} + \frac{2e\lambda H\sqrt{d}}{K}\sum_{i=h}^{H-1}\frac{1}{\iota_h} + \frac{40ed}{\sqrt{K}}\log\left(\frac{dH^2K}{\kappa\delta}\right)\sum_{i=h}^{H-1}\frac{H-h+1}{\sqrt{\iota_h(\eta_h + \sigma_r^2)}}.\end{aligned} \tag{C.26}$$

We further define

$$C_{h,1} = \sum_{i=h}^{H}\frac{1}{\iota_h},\ C_{h,2} = \sum_{i=h}^{H}\frac{H-h+1}{\sqrt{\iota_h(\eta_h + \sigma_r^2)}},$$

then we can simplify and rearrange (C.26) to

$$\sup_{s\in\mathcal{S}}\left|\widehat{V}_h^\pi(s) - V_h^\pi(s)\right| \le C \cdot \frac{C_{h,2}d}{\sqrt{K}}\log\left(\frac{dH^2K}{\kappa\delta}\right) + C \cdot \frac{C_{h,1}H\sqrt{d}}{K}.$$

This completes the proof of Theorem C.2.

D Proof of OPE Convergence

As stated in Appendix C.1, we consider the general form of Theorem 4.1. Recall the following notation:

$$C_{h,1} = \sum_{i=h}^{H}\frac{1}{\iota_h},\ C_{h,2} = \sum_{i=h}^{H}\sqrt{\frac{C_{h,3}}{\iota_h}},\ C_{h,3} = \frac{(H-h+1)^2}{\eta_h + \sigma_r^2},\ C_{h,4} = \left(\|\mathbf{\Lambda}_h\| \cdot \|\mathbf{\Lambda}_h^{-1}\|\right)^{1/2}.$$

Theorem D.1 (General form of Theorem 4.1)**.** Set $\lambda = 1$, $\eta_h \in (0, (H-h+1)^2]$ for all $h \in [H]$ and $\sigma_r^2 \leq 1$. Under Assumptions 2.1, 2.3 and 2.5, if K satisfies

$$K \geq C \cdot C_3 \cdot d^2 \left[\log\left(\frac{dH^2K}{\kappa\delta}\right)\right]^2, \tag{D.1}$$

where C is some problem-independent universal constant and

$$C_3 := \max\left\{\max_{h\in[H]} \frac{C_{h,3}\cdot C_{h,2}^2}{\iota_h^2(\eta_h+\sigma_r^2)^3}\ ,\ \frac{H^4}{\sigma_r^4\kappa^2},\ \frac{H^2}{\sigma_r^4\kappa^2}\cdot\max_{h\in[H]}\frac{C_{h,3}}{\eta_h+\sigma_r^2}\cdot\max_{h\in[H]}\frac{C_{h,3}}{\iota_h}\right\}.$$

Then with probability at least $1-\delta$, it holds that

$$|v_1^\pi - \widehat{v}_1^\pi| \leq C\cdot\left[\sum_{h=1}^H \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}\right]\cdot\sqrt{\frac{\log(16H/\delta)}{K}} + C\cdot C_4\cdot\log\left(\frac{16H}{\delta}\right)\cdot\left(\frac{1}{K^{3/4}}+\frac{1}{K}\right),$$

where $C_4 := \sum_{h=1}^H \left\{\sqrt{C_{h,4}\cdot C_{h,2}\cdot\frac{(H-h+1)d}{\iota_h(\eta_h+\sigma_r^2)^2}\cdot\log\left(\frac{dH^2K}{\kappa\delta}\right)}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}\right\}$.

Note that by setting $\eta_h = \sigma_r = 1$ we recover Theorem 4.1.

The proof is based on the recursive error decomposition given by (4.3) and the prerequisite result on uniform convergence. We will show the OPE convergence conditioned on the high probability event of uniform convergence established by Theorem C.2.

Recall the error decomposition for the OPE problem given by (4.3) (proof in Section E):

$$\begin{aligned}
v_1^\pi - \widehat{v}_1^\pi &= -\lambda\sum_{h=1}^H(\mathbf{v}_h^\pi)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\int_{\mathcal{S}}\left(V_{h+1}^\pi(s)-\widehat{V}_{h+1}^\pi(s)\right)\boldsymbol{\mu}_h(s)\mathrm{d}s \\
&\quad+\sum_{h=1}^H(\mathbf{v}_h^\pi)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\sum_{k=1}^K\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})}{\widehat{\sigma}_h(s_{k,h},a_{k,h})^2}\left([\mathbb{P}_h\widehat{V}_{h+1}^\pi](s_{k,h},a_{k,h})-\widehat{V}_{h+1}^\pi(s'_{k,h})-\epsilon_{k,h}\right) \\
&\quad+\lambda\sum_{h=1}^H(\mathbf{v}_h^\pi)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\mathbf{w}_h^\pi \\
&:= E_1+E_2+E_3.
\end{aligned} \tag{D.2}$$

It suffices to prove that each term can be bounded with high probability and then we can take a union bound. By the result of Theorem C.2, we can condition on the event where both $\widehat{V}_{h+1}^\pi$ and $\widehat{\sigma}_h$ are good estimators of their population counterparts.

Remark D.2. All the lemmas in the remaining of this Section D will be proved under Assumptions 2.1 and 2.5. So we do not explicitly add these two assumptions into the description of the lemmas.

Also, recall the function classes $\mathcal{V}_h(L)$ and $\mathcal{T}_h(L_1, L_2)$ defined by (C.2) and (C.3). In the remaining of this section, we will assume L, L_1 and L_2 to be

$$L = \frac{H+1}{\sqrt{\eta+\sigma_r^2}}\sqrt{\frac{Kd}{\lambda}},\ L_1 = H^2\sqrt{\frac{Kd}{\lambda}},\ L_2 = H\sqrt{\frac{Kd}{\lambda}}.$$

The reason that we can make the above assumption is that, conditioning on the high probability event of uniform convergence (Theorem C.2), it follows immediately from Lemma H.15 that we have $\widehat{\sigma}_h \in \mathcal{T}_h(L_1, L_2)$, and $\widehat{V}_h^\pi \in \mathcal{V}_h(L)$ for all $h \in [H]$ with the above choice of L, L_1 and L_2.

D.1 Bounding the E_2 Term in the OPE Decomposition

We consider the term E_2 first. Decompose E_2 into $E_2 = \sum_{h=1}^H E_{2,h}$ where for each $h \in [H]$, $E_{2,h}$ is given as

$$E_{2,h} := (\mathbf{v}_h^\pi)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\sum_{k=1}^K\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})}{\widehat{\sigma}_h(s_{k,h},a_{k,h})^2}\left([\mathbb{P}_h\widehat{V}_{h+1}^\pi](s_{k,h},a_{k,h})-\widehat{V}_{h+1}^\pi(s'_{k,h})-\epsilon_{k,h}\right).$$

Further decompose $E_{2,h}$ into $E_{2,h} = \sum_{k=1}^{K} e_{h,k}$ where

$$e_{h,k} = (\mathbf{v}_h^\pi)^\top \widehat{\mathbf{\Lambda}}_h^{-1} \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2} \left([\mathbb{P}_h \widehat{V}_{h+1}^\pi](s_{k,h}, a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h}) - \epsilon_{k,h} \right).$$

In the following lemma, we consider the term $E_{2,h}$ for arbitrarily fixed $h \in [H]$. To simplify the notation, we omit the subscript h and take $e_k = e_{h,k}$ once h is fixed.

Lemma D.3. For any $h \in [H]$, condition on $\widehat{V}_{h+1}^\pi \in \mathcal{V}_{h+1}(L)$ and the induced $\widehat{\sigma}_h(\cdot,\cdot) \in \mathcal{T}_h(L_1, L_2)$ being fixed, such that $\widehat{\sigma}_h$ satisfies for all (s,a)

$$\left| \widehat{\sigma}_h^2(s,a) - \sigma_r^2 - \max\left\{ \eta_h,\ \mathbb{V}_h \widehat{V}_{h+1}^\pi(s,a) \right\} \right| \le \frac{C(H-h+1)^2\sqrt{d}}{\sqrt{K}}, \tag{D.3}$$

for some $C > 0$. If K satisfies (D.7), then with conditional probability at least $1-\delta$, we have

$$|E_{2,h}| \le 2\sqrt{2\log\left(\frac{4}{\delta}\right) B_h} \cdot \frac{1}{\sqrt{K}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}} + \frac{8}{3}\log\left(\frac{4}{\delta}\right) \cdot \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}} \cdot \left\|\mathbf{G}_h^{-1}\right\|^{1/2} \cdot \frac{1}{K},$$

where B_h is a $\widehat{V}_{h+1}^\pi$-dependent constant and $\mathbf{G}_h$ is a $\widehat{\sigma}_h$-dependent matrix given by

$$\begin{aligned} B_h &= \max_{(s,a)\sim\nu_h} \frac{\mathbb{V}_h \widehat{V}_{h+1}^\pi(s,a) + \sigma_r^2}{\max\left\{\eta_h, \mathbb{V}_h \widehat{V}_{h+1}^\pi(s,a)\right\} + \sigma_r^2 - \frac{C(H-h+1)^2\sqrt{d}}{\sqrt{K}}} \\ &\sim 1 + \widetilde{\mathcal{O}}(1/\sqrt{K}), \\ \mathbf{G}_h &\coloneqq \mathbb{E}_h\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\widehat{\sigma}_h^2(s,a)} \middle| \widehat{\sigma}_h \right], \end{aligned}$$

Remark D.4. Lemma D.3 will be combined with Lemma F.2, which gives an explicit formula for the constant C with high probability, as will be shown in Lemma D.5.

Proof of Lemma D.3. By definition,

$$e_k = (\mathbf{v}_h^\pi)^\top \widehat{\mathbf{\Lambda}}_h^{-1} \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2} \left([\mathbb{P}_h \widehat{V}_{h+1}^\pi](s_{k,h}, a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h}) - \epsilon_{k,h} \right)$$

for all $k \in [K]$. From Algorithm 1, it is clear that the function$\widehat{V}_{h+1}^\pi(\cdot)$ depends on the dataset $\check{\mathcal{D}}_i, \mathcal{D}_i$ for $i \ge h+1$, and the function $\widehat{\sigma}_h(\cdot,\cdot)$ depends on $\widehat{V}_{h+1}^\pi$ and the dataset $\check{\mathcal{D}}_h$, which are all independent of the dataset $\mathcal{D}_h$ under Assumption C.1. Therefore, conditioning on $\widehat{V}_{h+1}$ and $\widehat{\sigma}_h$ will not change the distribution of $\mathcal{D}_h$.

Define $F_h = \{(s_{k,h}, a_{k,h}), k \in [K]\}$, and for now we further condition on F_h being fixed. Then $\widehat{\mathbf{\Lambda}}_h$ and $\widehat{\sigma}_{k,h} \coloneqq \widehat{\sigma}_h(s_{k,h}, a_{k,h}),\ k \in [K]$ are both fixed. Define the filtration $\{\mathcal{F}_k\}_{k\in[K]}$ conditioned on F_h as $\mathcal{F}_k = \sigma\{s'_{1,h}, \epsilon_{1,h}, \cdots, s'_{k-1,h}, \epsilon_{k-1,h} | F_h\}$ for $1 < k \le K$, and $\mathcal{F}_1$ as the empty σ-field. Then $\mathbb{E}[e_k \mid \mathcal{F}_k] = 0$ implies that $\{e_k\}_{k\in[K]}$ is a martingale difference sequence. Since $\widehat{V}_{h+1}^\pi \in \mathcal{V}_{h+1}(L)$ and $\widehat{\sigma}_h \in \mathcal{T}_h(L_1, L_2)$, we have $\widehat{\sigma}_h(s,a)^2 \ge \eta_h + \sigma_r^2$ for all (s,a) and $|\widehat{V}_{h+1}^\pi(s)| \le H-h+1$ for all s. Also by Assumption 2.1 we have $|\epsilon_{k,h}| \le 1$ almost surely. This then implies

$$|e_k| \le \underbrace{\frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \|\mathbf{v}_h^\pi\|_{\widehat{\mathbf{\Lambda}}_h^{-1}} \cdot \|\boldsymbol{\phi}(s_{k,h}, a_{k,h})\|_{\widehat{\mathbf{\Lambda}}_h^{-1}}}_{c_{h,k}},$$

and

$$\begin{aligned} \mathrm{Var}(e_k | F_h, \mathcal{F}_k) &= \left[(\mathbf{v}_h^\pi)^\top \widehat{\mathbf{\Lambda}}_h^{-1} \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})} \right]^2 \cdot \mathbb{E}\left[\left(\frac{\mathbb{P}_h \widehat{V}_{h+1}^\pi(s_{k,h}, a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h}) - \epsilon_{k,h}}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})} \right)^2 \middle| F_h, \mathcal{F}_k \right] \\ &\le \left[(\mathbf{v}_h^\pi)^\top \widehat{\mathbf{\Lambda}}_h^{-1} \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}(s_{k,h}, a_{k,h})^\top}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2} \widehat{\mathbf{\Lambda}}_h^{-1} \mathbf{v}_h^\pi \right] \\ &\quad \cdot \left(\frac{\mathbb{V}_h \widehat{V}_{h+1}^\pi(s_{k,h}, a_{k,h}) + \sigma_r^2}{\max\left\{\eta_h, \mathbb{V}_h \widehat{V}_{h+1}^\pi(s_{k,h}, a_{k,h})\right\} + \sigma_r^2 - \frac{C(H-h+1)^2\sqrt{d}}{\sqrt{K}}} \right), \end{aligned}$$

where the last step is from Assumption 2.1 that $\epsilon_{k,h}$ is independent random noise satisfying $\mathrm{Var}[\epsilon_{k,h} \mid s_{k,h}, a_{k,k}] \leq \sigma_r^2$, and (D.3). Denote $c_h = \max_{k\in[K]}\{c_{h,k}\}$. We then have

$$c_h \leq \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \|\mathbf{v}_h^\pi\|_{\widehat{\mathbf{\Lambda}}_h^{-1}} \cdot \|\widehat{\mathbf{\Lambda}}_h^{-1}\|^{1/2}. \tag{D.4}$$

For simplicity, denote

$$\begin{aligned} b_{h,k} &:= \frac{\mathbb{V}_h \widehat{V}_{h+1}^\pi(s_{k,h}, a_{k,h}) + \sigma_r^2}{\max\left\{\eta_h, \mathbb{V}_h \widehat{V}_{h+1}^\pi(s_{k,h}, a_{k,h})\right\} + \sigma_r^2 - \frac{C(H-h+1)^2\sqrt{d}}{\sqrt{K}}} \\ &\leq 1 + \frac{\frac{C(H-h+1)^2\sqrt{d}}{\sqrt{K}}}{\eta_h + \sigma_r^2 - \frac{C(H-h+1)^2\sqrt{d}}{\sqrt{K}}} \\ &\sim 1 + \widetilde{\mathcal{O}}(1/\sqrt{K}), \end{aligned}$$

and $b_h := \max_k\{b_{h,k}\}$. Therefore, we further have

$$\begin{aligned} \sum_{k=1}^K \mathrm{Var}(e_k|F_h, \mathcal{F}_k) &\leq (\mathbf{v}_h^\pi)^\top \widehat{\mathbf{\Lambda}}_h^{-1} \left(\sum_{k=1}^K \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}(s_{k,h}, a_{k,h})^\top}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2}\right) \widehat{\mathbf{\Lambda}}_h^{-1}\mathbf{v}_h^\pi \cdot b_h \\ &= (\mathbf{v}_h^\pi)^\top \widehat{\mathbf{\Lambda}}_h^{-1} \left(\widehat{\mathbf{\Lambda}}_h - \lambda \mathbf{I}_d\right) \widehat{\mathbf{\Lambda}}_h^{-1}\mathbf{v}_h^\pi \cdot b_h \\ &\leq b_h \cdot \|\mathbf{v}_h^\pi\|^2_{\widehat{\mathbf{\Lambda}}_h^{-1}}, \end{aligned}$$

since $(\widehat{\mathbf{\Lambda}}_h)^{-1/2}(\widehat{\mathbf{\Lambda}}_h - \lambda\mathbf{I}_d)(\widehat{\mathbf{\Lambda}}_h)^{-1/2}$ is a contraction. Then by Freedman's inequality H.2, we have

$$\mathbb{P}\left(\left|\sum_{k=1}^K e_k\right| \geq \epsilon \middle| F_h\right) \leq 2\exp\left(-\frac{\epsilon^2/2}{b_h\|\mathbf{v}_h^\pi\|^2_{\widehat{\mathbf{\Lambda}}_h^{-1}} + c_h\epsilon/3}\right),$$

since b_h and c_h are fixed once we condition on $\widehat{V}_{h+1}^\pi$, $\widehat{\sigma}_h$ and F_h. It follows that with conditional (on $\widehat{V}_{h+1}^\pi, \widehat{\sigma}_h, F_h$) probability at least $1-\delta$,

$$\left|\sum_{k=1}^K e_k\right| \leq \sqrt{2\log\left(\frac{2}{\delta}\right) b_h} \cdot \|\mathbf{v}_h^\pi\|_{\widehat{\mathbf{\Lambda}}_h^{-1}} + \frac{2}{3}\log\frac{2}{\delta} \cdot c_h. \tag{D.5}$$

Define the matrix $\mathbf{G}_h$ as the conditional expectation given as

$$\mathbf{G}_h := \mathbb{E}_h\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\widehat{\sigma}_h(s,a)^2}\middle|\widehat{\sigma}_h\right], \tag{D.6}$$

by recalling the notation $\mathbb{E}_h[f(s,a)] = \int_{\mathcal{S}\times\mathcal{A}} f(s,a)\mathrm{d}\nu_h(s,a)$ for any function f on $\mathcal{S}\times\mathcal{A}$, with $\nu_h(\cdot,\cdot)$ being the occupancy measure of the MDP for stage h induced by the behavior policy $\bar{\pi}$. Now, since conditioning on $\widehat{V}_{h+1}^\pi$ and $\widehat{\sigma}_h$ does not change the distribution of F_h, by Lemma H.5, if K satisfies

$$K \geq \max\left\{512(\eta_h + \sigma_r^2)^{-2}\|\mathbf{G}_h^{-1}\|^2 \log\left(\frac{2d}{\delta}\right), 4\lambda\|\mathbf{G}_h^{-1}\|\right\}, \tag{D.7}$$

then over the space of F_h, there exists an event $\mathcal{E}_h$ such that $\mathbb{P}(\mathcal{E}_h) \geq 1-\delta$ and for all $F_h \in \mathcal{E}_h$ we have

$$\|\mathbf{u}\|_{\widehat{\mathbf{\Lambda}}_h^{-1}} \leq \frac{2}{\sqrt{K}} \cdot \|\mathbf{u}\|_{\mathbf{G}_h^{-1}} \tag{D.8}$$

for all $\mathbf{u} \in \mathbb{R}^d$. Combining (D.4), (D.5) and (D.8), we conclude that, with conditional probability (on $\widehat{V}_{h+1}^\pi, \widehat{\sigma}_h$ only) at least $1-2\delta$,

$$\left|\sum_{k=1}^K e_k\right| \leq \sqrt{2\log\left(\frac{2}{\delta}\right) B_h} \cdot \frac{2}{\sqrt{K}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}} + \frac{2}{3}\log\frac{2}{\delta} \cdot \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}} \cdot \left\|\mathbf{G}_h^{-1}\right\|^{1/2} \cdot \frac{4}{K},$$

where B_h is a $\widehat{V}^\pi_{h+1}$-dependent constants given by

$$B_h = \max_{(s,a)\sim\nu_h} \frac{\mathbb{V}_h \widehat{V}^\pi_{h+1}(s,a) + \sigma_r^2}{\max\left\{\eta_h, \mathbb{V}_h \widehat{V}^\pi_{h+1}(s,a)\right\} + \sigma_r^2 - \frac{C(H-h+1)^2\sqrt{d}}{\sqrt{K}}}$$
$$\sim 1 + \widetilde{\mathcal{O}}(1/\sqrt{K}),$$

and $\mathbf{G}_h$ is a $\widehat{\sigma}_h$-dependent matrix given by (D.6). Replacing δ with $\delta/2$ finishes the proof. □

In the next lemma, we relax the conditioning on $\widehat{\sigma}_h$ and condition on $\widehat{V}^\pi_{h+1}$ only.

Lemma D.5. For any $h \in [H]$, condition on $\widehat{V}^\pi_{h+1} \in \mathcal{V}_{h+1}(L)$ being fixed and satisfying $\sup_s |\widehat{V}^\pi_{h+1}(s) - V^\pi_{h+1}(s)| \le \rho$ for some $\rho \ge 0$, if K satisfies (D.14) and

$$K \ge \max\left\{\frac{911}{(\eta_h + \sigma_r^2)^2 \iota_h^2} \cdot \log\left(\frac{4d}{\delta}\right), \frac{6\lambda}{\iota_h}\right\}, \tag{D.9}$$

then with conditional probability at least $1-\delta$, we have

$$\begin{aligned}
|E_{2,h}| \le\ & 2\sqrt{2\log\left(\frac{8}{\delta}\right) B_h} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \frac{1}{\sqrt{K}} \\
&+ 2\sqrt{2\log\left(\frac{8}{\delta}\right) B_h} \cdot \sqrt{C_0 C_1 \cdot \frac{1}{\iota_h}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot (K^{1/4}\sqrt{\widetilde{\rho}}) \cdot \frac{1}{K^{3/4}} \\
&+ \frac{8}{3}\log\left(\frac{8}{\delta}\right) \cdot \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2}\sqrt{C_1} \cdot \frac{1}{\sqrt{\iota_h}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \frac{1}{K} \\
&+ \frac{8}{3}\log\left(\frac{8}{\delta}\right) \cdot \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \sqrt{C_0} \cdot C_1 \cdot \frac{1}{\iota_h} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \sqrt{\widetilde{\rho}} \cdot \frac{1}{K}
\end{aligned}$$

where

$$\begin{aligned}
C_0 &= \left(\frac{\|\mathbf{\Lambda}_h\|}{\iota_h}\right)^{1/2}, \\
C_1 &= \frac{1}{1 - \widetilde{\rho}/\iota_h}, \\
\widetilde{\rho} &= \frac{1}{(\eta_h + \sigma_r^2)^2} \cdot \left(\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}} + 4(H-h+1)\cdot\rho\right), \\
C_{K,h,\delta} &= 12\sqrt{2} \cdot \frac{1}{\sqrt{\kappa_h}} \cdot \left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d}\log\frac{8}{\delta}\right]^{1/2} + 12\lambda \cdot \frac{1}{\kappa_h}.
\end{aligned}$$

and B_h is a $\widehat{V}^\pi_{h+1}$-dependent constant given by

$$B_h = \max_{(s,a)\sim\nu_h} \frac{\mathbb{V}_h \widehat{V}^\pi_{h+1}(s,a) + \sigma_r^2}{\max\left\{\eta_h, \mathbb{V}_h \widehat{V}^\pi_{h+1}(s,a)\right\} + \sigma_r^2 - \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}}.$$

Remark D.6. Conditioning on the event in Theorem C.2, we have $\rho \sim \widetilde{\mathcal{O}}(1/\sqrt{K})$ and thus $\widetilde{\rho} \sim \mathcal{O}(1/\sqrt{K})$, which means the term $\sqrt{K}\widetilde{\rho}$ is a constant up to a logarithmic factor. This indicates that in the upper bound of $|E_{2,h}|$, only the first term is of order $\widetilde{\mathcal{O}}(1/\sqrt{K})$.

Proof of Lemma D.5. For simplicity, denote the function $\sigma_V(\cdot,\cdot)$ as

$$\sigma_V(\cdot,\cdot) := \sqrt{\max\left\{\eta_h,\ \mathbb{V}_h \widehat{V}^\pi_{h+1}(\cdot,\cdot)\right\} + \sigma_r^2},$$

and recall that $\widehat{\sigma}_h(\cdot,\cdot)$ is an estimator for $\sigma_V(\cdot,\cdot)$ generated using the dataset $\check{\mathcal{D}}_h$. Also recall the definition

$$\sigma_h(\cdot,\cdot) := \sqrt{\max\left\{\eta_h,\ \mathbb{V}_h V^\pi_{h+1}(\cdot,\cdot)\right\} + \sigma_r^2}.$$

First of all, by Lemma F.2, with probability at least $1-\delta$ over the space of $\check{\mathcal{D}}_h$, the following event happens:

$$\sup_{s,a} |\widehat{\sigma}_h^2(s,a) - \sigma_V^2(s,a)| \le \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}, \tag{D.10}$$

where $C_{K,h,\delta}$ is given by

$$C_{K,h,\delta} = 12\sqrt{2}\cdot\frac{1}{\sqrt{\kappa_h}}\cdot\left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d}\log\frac{4}{\delta}\right]^{1/2} + 12\lambda\cdot\frac{1}{\kappa_h}.$$

Denote the above event of $\widehat{\sigma}_h$ by $\mathcal{E}_{\widehat{\sigma}}$. For each fixed $\widehat{\sigma}_h \in \mathcal{E}_{\widehat{\sigma}}$, we can then apply Lemma D.3 with C replace by $C_{K,h,\delta}$. This gives that, for any h, condition on $\widehat{V}^\pi_{h+1}$ and $\widehat{\sigma}_h$, with probability at least $1-\delta$,

$$\begin{aligned}
&|E_{2,h}| \\
&\le 2\sqrt{2\log\left(\frac{4}{\delta}\right)B_h}\cdot\frac{1}{\sqrt{K}}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}} + \frac{8}{3}\log\left(\frac{4}{\delta}\right)\cdot\frac{2(H-h+1)+1}{\eta_h+\sigma_r^2}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}}\cdot\left\|\mathbf{G}_h^{-1}\right\|^{1/2}\cdot\frac{1}{K},
\end{aligned} \tag{D.11}$$

where

$$\begin{aligned}
B_h &:= \max_{(s,a)\sim\nu_h} \frac{\mathbb{V}_h\widehat{V}^\pi_{h+1}(s,a)+\sigma_r^2}{\max\left\{\eta_h, \mathbb{V}_h\widehat{V}^\pi_{h+1}(s,a)\right\}+\sigma_r^2 - \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}} \sim 1+\widetilde{\mathcal{O}}(1/\sqrt{K}), \\
\mathbf{G}_h &:= \mathbb{E}_h\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\widehat{\sigma}_h^2(s,a)}\middle|\widehat{\sigma}_h\right].
\end{aligned}$$

However, note that in the upper bound of $|E_{2,h}|$, the $\|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}}$ and $\left\|\mathbf{G}_h^{-1}\right\|^{1/2}$ term are $\widehat{\sigma}_h$-dependent, and so is the lower bound of the sample complexity given by (D.7). Therefore, it remains to derive a uniform upper bound of $E_{2,h}$ for all $\widehat{\sigma}_h \in \mathcal{E}_{\widehat{\sigma}}$, and a uniform lower bound of K.

To get this, first note that since $\widehat{V}^\pi_{h+1}, V^\pi_{h+1} \in \mathcal{V}_{h+1}(L)$ and $\sup_s|\widehat{V}^\pi_{h+1}(s) - V^\pi_{h+1}(s)| \le \rho$, we have

$$\sup_{s,a}|\sigma_V^2(s,a) - \sigma_h^2(s,a)| \le 4(H-h+1)\rho.$$

Using triangular inequality and (D.10) gives

$$\sup_{s,a}|\widehat{\sigma}_h^2(s,a) - \sigma_h^2(s,a)| \le \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}} + 4(H-h+1)\cdot\rho,$$

for all $\widehat{\sigma}_h \in \mathcal{E}_{\widehat{\sigma}}$.

Note that by definition,

$$\begin{aligned}
\|\mathbf{G}_h - \boldsymbol{\Lambda}_h\| &= \left\|\mathbb{E}_h\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\widehat{\sigma}_h^2(s,a)}\right] - \mathbb{E}_h\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma_h^2(s,a)}\right]\right\| \\
&= \left\|\mathbb{E}_h\left[\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top\frac{\widehat{\sigma}_h^2(s,a)-\sigma_h^2(s,a)}{\widehat{\sigma}_h^2(s,a)\cdot\sigma_h^2(s,a)}\right]\right\| \\
&\le \frac{1}{(\eta_h+\sigma_r^2)^2}\cdot\left(\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}} + 4(H-h+1)\cdot\rho\right) \\
&:= \widetilde{\rho} \\
&\sim \widetilde{\mathcal{O}}(1/\sqrt{K}),
\end{aligned}$$

where the inequality is from $\|\boldsymbol{\phi}(\cdot,\cdot)\| \le 1$ and $|\sigma(\cdot,\cdot)| \ge \sqrt{\eta_h + \sigma_r^2}$ for all $\sigma(\cdot,\cdot) \in \mathcal{T}_h$. Combine the above inequality with Lemma H.3, and we have

$$\left\|\mathbf{G}_h^{-1}\right\| \le \frac{\left\|\boldsymbol{\Lambda}_h^{-1}\right\|}{1 - \left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \|\boldsymbol{\Lambda}_h - \mathbf{G}_h\|} \le \frac{\left\|\boldsymbol{\Lambda}_h^{-1}\right\|}{1 - \left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \widetilde{\rho}}, \tag{D.12}$$

and by Lemma H.3 again,

$$\begin{aligned}
\|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}} &\le \left[1 + \sqrt{\left(\left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \|\boldsymbol{\Lambda}_h\|\right)^{1/2} \cdot \left\|\mathbf{G}_h^{-1}\right\| \cdot \widetilde{\rho}}\right] \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} \\
&\le \left[1 + \sqrt{\left(\left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \|\boldsymbol{\Lambda}_h\|\right)^{1/2} \cdot \frac{\left\|\boldsymbol{\Lambda}_h^{-1}\right\|}{1 - \left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \widetilde{\rho}} \cdot \widetilde{\rho}}\right] \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} \\
&= \left[1 + \sqrt{\left(\left\|\boldsymbol{\Lambda}_h^{-1}\right\|\right)^{3/2} \cdot \left(\|\boldsymbol{\Lambda}_h\|\right)^{1/2} \cdot \frac{1}{1 - \left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \widetilde{\rho}} \cdot \widetilde{\rho}}\right] \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} \\
&= \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} + \sqrt{\left(\left\|\boldsymbol{\Lambda}_h^{-1}\right\|\right)^{3/2} \cdot \left(\|\boldsymbol{\Lambda}_h\|\right)^{1/2} \cdot \frac{1}{1 - \left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \widetilde{\rho}}} \cdot \sqrt{\widetilde{\rho}} \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}}.
\end{aligned} \tag{D.13}$$

Note that the above holds when K is sufficiently large such that $\left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \widetilde{\rho}$ is less than, for example,

$$\left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \widetilde{\rho} \le 1/4. \tag{D.14}$$

We are now ready to derive an upper bound independent of $\widehat{\sigma}_h$. First define

$$C_0 = \left(\left\|\boldsymbol{\Lambda}_h^{-1}\right\|\right)^{1/2} \cdot \left(\|\boldsymbol{\Lambda}_h\|\right)^{1/2}, \quad C_1 = \frac{1}{1 - \left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \widetilde{\rho}}.$$

Then we have $\|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}} \le \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} + \sqrt{C_0 C_1 \left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \widetilde{\rho}} \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}}$, and $\left\|\mathbf{G}_h^{-1}\right\|^{1/2} \le \sqrt{C_1} \cdot \left\|\boldsymbol{\Lambda}_h^{-1}\right\|^{1/2}$. It follows that

$$\|\mathbf{v}_h^\pi\|_{\mathbf{G}_h^{-1}} \cdot \left\|\mathbf{G}_h^{-1}\right\|^{1/2} \le \sqrt{C_1} \cdot \left\|\boldsymbol{\Lambda}_h^{-1}\right\|^{1/2} \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} + C_1 \cdot \left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \sqrt{C_0} \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} \cdot \sqrt{\widetilde{\rho}}.$$

Plug into (D.11), and we have that, condition on $\widehat{V}_{h+1}^\pi$, with probability at least $1 - 2\delta$,

$$\begin{aligned}
|E_{2,h}| \le\ & 2\sqrt{2\log\left(\frac{4}{\delta}\right) B_h} \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} \cdot \frac{1}{\sqrt{K}} \\
& + 2\sqrt{2\log\left(\frac{4}{\delta}\right) B_h} \cdot \sqrt{C_0 C_1 \left\|\boldsymbol{\Lambda}_h^{-1}\right\|} \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} \cdot \sqrt{\widetilde{\rho}} \cdot \frac{1}{\sqrt{K}} \\
& + \frac{8}{3}\log\left(\frac{4}{\delta}\right) \cdot \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \sqrt{C_1} \cdot \left\|\boldsymbol{\Lambda}_h^{-1}\right\|^{1/2} \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} \cdot \frac{1}{K} \\
& + \frac{8}{3}\log\left(\frac{4}{\delta}\right) \cdot \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot C_1 \cdot \left\|\boldsymbol{\Lambda}_h^{-1}\right\| \cdot \sqrt{C_0} \cdot \|\mathbf{v}_h^\pi\|_{\boldsymbol{\Lambda}_h^{-1}} \cdot \sqrt{\widetilde{\rho}} \cdot \frac{1}{K}
\end{aligned}$$

Replacing δ by $\delta/2$ and using $1/\iota_h = \|\boldsymbol{\Lambda}_h^{-1}\|$ gives the desired upper bound. It remains to show the lower bound. By (D.7), (D.12), and (D.14), a uniform version of D.7 is given by

$$\begin{aligned}
K &\ge \max\left\{512(\eta_h + \sigma_r^2)^{-2} \cdot \frac{16}{9}\|\boldsymbol{\Lambda}_h^{-1}\|^2 \log\left(\frac{4d}{\delta}\right), 4\lambda \cdot \frac{4}{3}\|\boldsymbol{\Lambda}_h^{-1}\|\right\} \\
&> \max\left\{\frac{911}{(\eta_h + \sigma_r^2)^2 \iota_h^2} \cdot \log\left(\frac{4d}{\delta}\right), \frac{6\lambda}{\iota_h}\right\}.
\end{aligned} \tag{D.15}$$

□

Lemma D.7. If K satisfies (D.17), (D.20), (D.21) and

$$K \ge \max_{h \in [H]} \max\left\{\frac{911}{(\eta_h + \sigma_r^2)^2 \iota_h^2} \log\left(\frac{8Hd}{\delta}\right), \frac{6\lambda}{\iota_h}\right\}, \tag{D.16}$$

then with probability at least $1-\delta$, we have

$$
\begin{aligned}
|E_2| \le\ & 2\sqrt{2\log\left(\frac{16H}{\delta}\right)B}\cdot\left[\sum_{h=1}^{H}\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}\right]\cdot\frac{1}{\sqrt{K}} \\
&+\frac{16\sqrt{2}}{3}\log\left(\frac{16H}{\delta}\right)\cdot\sqrt{B}\cdot\left[\sum_{h=1}^{H}A_1(h)\right]\cdot\frac{1}{K^{3/4}} \\
&+\frac{16\sqrt{2}}{3}\log\left(\frac{16H}{\delta}\right)\cdot\sqrt{B}\cdot\left[\sum_{h=1}^{H}(A_2(h)+A_3(h))\right]\cdot\frac{1}{K},
\end{aligned}
$$

where

$$
\begin{aligned}
A_1(h) &= \sqrt{C_0(h)\cdot\frac{1}{\iota_h}}\cdot\left(K^{1/4}\sqrt{\widetilde{\rho}(h)}\right)\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}, \\
A_2(h) &= \frac{2(H-h+1)+1}{\eta_h+\sigma_r^2}\cdot\frac{1}{\sqrt{\iota_h}}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}, \\
A_3(h) &= \frac{2(H-h+1)+1}{\eta_h+\sigma_r^2}\cdot\sqrt{C_0(h)}\cdot\sqrt{\widetilde{\rho}}\cdot\frac{1}{\iota_h}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}, \\
C_0(h) &= \left(\frac{\|\mathbf{\Lambda}_h\|}{\iota_h}\right)^{1/2}, \\
\widetilde{\rho}(h) &= \frac{1}{(\eta_h+\sigma_r^2)^2}\cdot\left(\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}+4(H-h+1)\cdot\widetilde{C}(h)\cdot\frac{d}{\sqrt{K}}\right), \\
C_{K,h,\delta} &= 12\sqrt{2}\cdot\frac{1}{\sqrt{\kappa_h}}\cdot\left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right)+\frac{1}{d}\log\frac{16H}{\delta}\right]^{1/2}+12\lambda\cdot\frac{1}{\kappa_h}, \\
\widetilde{C}(h) &= C\cdot C_{h,2}\cdot\log\left(\frac{dH^2K}{\kappa\delta}\right),
\end{aligned}
$$

and C is some universal constant, B is a problem-dependent constant given by

$$
B=\max_{h\in[H]}\max_{V\in\mathcal{V}_{h+1}(L)}\max_{(s,a)\sim\nu_h}\frac{\mathbb{V}_hV(s,a)+\sigma_r^2}{\max\{\eta_h,\mathbb{V}_hV(s,a)\}+\sigma_r^2-\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}},
$$

and $C_{h,2}$ are the same constants as in Theorem C.2.

Proof of Lemma D.7. First, by Theorem C.2, if K satisfies

$$
K\ge C\cdot\frac{H^2d^2}{\kappa^2}\log\left(\frac{dHK}{\kappa\delta}\right)\cdot\max_{h\in[H]}\frac{(H-h+1)^2}{(\eta_h+\sigma_r^2)^2}\cdot\max_{h\in[H]}\frac{(H-h+1)^2}{\iota_h(\eta_h+\sigma_r^2)}, \tag{D.17}
$$

for some problem-independent constant C, then with probability at least $1-\delta/2$, for all $h\in[H]$, we have

$$
\sup_s\left|\widehat{V}_{h+1}^\pi(s)-V_{h+1}^\pi(s)\right|\le\widetilde{C}\cdot\frac{d}{\sqrt{K}},
$$

where

$$
\widetilde{C}\coloneqq C\cdot C_{h,2}\cdot\log\left(\frac{dH^2K}{\kappa\delta}\right)+C\cdot C_{h,1}\cdot\frac{H}{\sqrt{dK}}, \tag{D.18}
$$

and

$$
C_{h,1}=\sum_{i=h}^{H}\frac{1}{\iota_h},\qquad C_{h,2}=\sum_{i=h}^{H}\frac{H-h+1}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}.
$$

For simplicity, we define

$$\widetilde{C}(h) = C \cdot C_{h,2} \cdot \log\left(\frac{dH^2K}{\kappa\delta}\right),$$

for some different constant C, since the first term on the RHS of (D.18) is much larger than the second one by using $\eta_h \le (H-h+1)^2$.

Now we can combine Theorem C.2 and Lemma D.5 with the parameter ρ replaced by $\widetilde{C} \cdot \frac{d}{\sqrt{K}}$, take a union bound over all H terms, and conclude that, with probability at least $1-\delta$, the result of Lemma D.5 holds for all $h \in [H]$:

$$\begin{aligned}
|E_{2,h}| \le\ & 2\sqrt{2\log\left(\frac{16H}{\delta}\right) B_h} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \frac{1}{\sqrt{K}} \\
& + 2\sqrt{2\log\left(\frac{16H}{\delta}\right) B_h} \cdot \sqrt{C_0 C_1 \cdot \frac{1}{\iota_h}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot (K^{1/4}\sqrt{\widetilde{\rho}(h)}) \cdot \frac{1}{K^{3/4}} \\
& + \frac{8}{3}\log\left(\frac{16H}{\delta}\right) \cdot \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \sqrt{C_1} \cdot \frac{1}{\sqrt{\iota_h}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \frac{1}{K} \\
& + \frac{8}{3}\log\left(\frac{16H}{\delta}\right) \cdot \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \sqrt{C_0} \cdot C_1 \cdot \frac{1}{\iota_h} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \sqrt{\widetilde{\rho}(h)} \cdot \frac{1}{K}
\end{aligned} \tag{D.19}$$

where for each $h \in [H]$, $\widetilde{\rho}(h)$ is given as

$$\widetilde{\rho}(h) = \frac{1}{(\eta_h + \sigma_r^2)^2} \cdot \left(\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}} + 4(H-h+1) \cdot \widetilde{C}(h) \cdot \frac{d}{\sqrt{K}}\right),$$

$$C_{K,h,\delta} = 12\sqrt{2} \cdot \frac{1}{\sqrt{\kappa_h}} \cdot \left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d}\log\frac{16H}{\delta}\right]^{1/2} + 12\lambda \cdot \frac{1}{\kappa_h}.$$

Now, by the expression of B_h, if K satisfies

$$K \ge \frac{4C_{K,h,\delta}^2 H^4 d}{\sigma_r^4} \ge 1152 \cdot \max_{h\in[H]} \frac{1}{\kappa_h^2} \cdot \left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d}\log\frac{16H}{\delta}\right] \cdot \frac{H^4 d}{\sigma_r^4}, \tag{D.20}$$

then we have $B_h \le 2$ for all $h \in [H]$. Also, by (D.14), K also needs to be large enough so that

$$\max_{h\in[H]} \left\{\|\mathbf{\Lambda}_h^{-1}\| \cdot \widetilde{\rho}(h)\right\} = \max_{h\in[H]} \{\widetilde{\rho}(h)/\iota_h\} \le 1/4, \tag{D.21}$$

which implies $C_1 \le 4/3$ for all h. We can then simplify (D.19) into

$$\begin{aligned}
& |E_{2,h}| \\
& \le 2\sqrt{2\log\left(\frac{16H}{\delta}\right) B_h} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \frac{1}{\sqrt{K}} \\
& \quad + \frac{16\sqrt{2}}{3}\log\left(\frac{16H}{\delta}\right) \cdot \sqrt{B_h} \cdot A_1(h) \cdot \frac{1}{K^{3/4}} \\
& \quad + \frac{16\sqrt{2}}{3}\log\left(\frac{16H}{\delta}\right) \cdot \sqrt{B_h} \cdot [A_2(h) + A_3(h)] \cdot \frac{1}{K},
\end{aligned}$$

where

$$\begin{aligned}
A_1(h) &= \sqrt{C_0(h) \cdot \frac{1}{\iota_h}} \cdot \left(K^{1/4}\sqrt{\widetilde{\rho}(h)}\right) \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}, \\
A_2(h) &= \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \frac{1}{\sqrt{\iota_h}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}, \\
A_3(h) &= \frac{2(H-h+1)+1}{\eta_h + \sigma_r^2} \cdot \sqrt{C_0(h)} \cdot \sqrt{\widetilde{\rho}} \cdot \frac{1}{\iota_h} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}.
\end{aligned}$$

By denoting

$$B \coloneqq \max_{h\in[H]} \max_{V\in\mathcal{V}_{h+1}(L)} \max_{(s,a)\sim\nu_h} \frac{\mathbb{V}_h V(s,a)+\sigma_r^2}{\max\{\eta_h, \mathbb{V}_h V(s,a)\}+\sigma_r^2 - \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}},$$

which is less than 2 by (D.20),and using $|E_2| \le \sum_{h=1}^K |E_{2,h}|$, we prove the upper bound. The lower bound of K comes from (D.17), (D.20), (D.21) and (D.15) .

□

D.2 Bounding the E_1 Term in the OPE Decomposition

Consider the term E_1 in (D.2):

$$|E_1| \le \lambda \sum_{h=1}^H \left| (\mathbf{v}_h^\pi)^\top \widehat{\mathbf{\Lambda}}_h^{-1} \int_{\mathcal{S}} \left(V_{h+1}^\pi(s) - \widehat{V}_{h+1}^\pi(s) \right) \boldsymbol{\mu}_h(s) \mathrm{d}s \right| \coloneqq \sum_{h=1}^H |E_{1,h}|,$$

where for each $h \in [H]$,

$$E_{1,h} \coloneqq \lambda (\mathbf{v}_h^\pi)^\top \widehat{\mathbf{\Lambda}}_h^{-1} \int_{\mathcal{S}} \left(V_{h+1}^\pi(s) - \widehat{V}_{h+1}^\pi(s) \right) \boldsymbol{\mu}_h(s) \mathrm{d}s.$$

Lemma D.8. Under the same event where the result of Lemma D.7 holds, if K satisfies (D.16), (D.17), (D.20) and (D.21), we have

$$|E_1| \le 4\sqrt{2}\lambda \left[\sum_{h=1}^H A_4(h) \right] \cdot \frac{H\sqrt{d}}{K},$$

where for each h,

$$A_4(h) = \frac{H-h+1}{H} \cdot \frac{1}{\sqrt{\iota_h}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \left[1 + \sqrt{2C_0(h) \cdot \frac{1}{\iota_h} \cdot \widetilde{\rho}(h)} \right],$$

and $C_0(h)$ and $\widetilde{\rho}(h)$ are same constants as in Lemma D.7.

Proof of Lemma D.8. By (D.8) and (D.13) , we have that

$$\begin{aligned}
\|\mathbf{u}\|_{\widehat{\mathbf{\Lambda}}_h^{-1}} &\le \frac{2}{\sqrt{K}} \cdot \left\{ \|\mathbf{u}\|_{\mathbf{\Lambda}_h^{-1}} + \sqrt{\left(\|\mathbf{\Lambda}_h^{-1}\|\right)^{3/2} \cdot \left(\|\mathbf{\Lambda}_h\|\right)^{1/2} \cdot \frac{1}{1-\|\mathbf{\Lambda}_h^{-1}\| \cdot \widetilde{\rho}(h)}} \cdot \sqrt{\widetilde{\rho}(h)} \cdot \|\mathbf{u}\|_{\mathbf{\Lambda}_h^{-1}} \right\} \\
&= \frac{2}{\sqrt{K}} \cdot \left\{ \|\mathbf{u}\|_{\mathbf{\Lambda}_h^{-1}} + \sqrt{C_0(h) \cdot \frac{1}{\iota_h} \cdot \frac{1}{1-\|\mathbf{\Lambda}_h^{-1}\| \cdot \widetilde{\rho}(h)}} \cdot \sqrt{\widetilde{\rho}(h)} \cdot \|\mathbf{u}\|_{\mathbf{\Lambda}_h^{-1}} \right\}, \qquad \text{(D.22)}
\end{aligned}$$

for all $\mathbf{u} \in \mathbb{R}^d$, where the constants take the same values as given in Lemma D.7, i.e.,

$$\begin{aligned}
\widetilde{\rho}(h) &= \frac{1}{(\eta_h+\sigma_r^2)^2} \cdot \left(\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}} + 4(H-h+1) \cdot \widetilde{C}(h) \cdot \frac{d}{\sqrt{K}} \right), \\
C_{K,h,\delta} &= 12\sqrt{2} \cdot \frac{1}{\sqrt{\kappa_h}} \cdot \left[\frac{1}{2} \log\left(\frac{\lambda+K}{\lambda} \right) + \frac{1}{d} \log \frac{16H}{\delta} \right]^{1/2} + 12\lambda \cdot \frac{1}{\kappa_h}, \\
\widetilde{C}(h) &= C \cdot C_{h,2} \cdot \log\left(\frac{dH^2K}{\kappa\delta} \right),
\end{aligned}$$

with $C_{h,2}$ being the same constant as in Theorem C.2. Also, since the result of Lemma G.4 holds, we have

$$\left\| \frac{\widehat{\mathbf{\Lambda}}_h}{K} - \mathbf{\Lambda}_h \right\| \le \frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}} \cdot \left(\log \frac{16Hd}{\delta} \right)^{1/2} + \frac{\lambda}{K} + \widetilde{\rho}(h) \le 2\widetilde{\rho}(h),$$

where the first step is by replacing δ with $\delta/4H$and the second step is by the choice of K in Lemma D.7. Then It follows from Lemma H.3 that

$$\left\|\widehat{\mathbf{\Lambda}}_h^{-1}\right\| \leq \frac{\left\|(K\mathbf{\Lambda}_h)^{-1}\right\|}{1-\left\|(K\mathbf{\Lambda}_h)^{-1}\right\|\cdot\left\|\widehat{\mathbf{\Lambda}}_h - K\mathbf{\Lambda}_h\right\|} \leq \frac{1}{K}\cdot\frac{\left\|\mathbf{\Lambda}_h^{-1}\right\|}{1-2\widetilde{\rho}(h)\cdot\left\|\mathbf{\Lambda}_h^{-1}\right\|} \leq \frac{2\left\|\mathbf{\Lambda}_h^{-1}\right\|}{K}, \tag{D.23}$$

since $2\widetilde{\rho}(h)\cdot\left\|\mathbf{\Lambda}_h^{-1}\right\| \leq 1/2$ by (D.21). Also, since on the event of Lemma D.7, we have $|V_{h+1}^{\pi}(s) - \widehat{V}_{h+1}^{\pi}(s)| \leq 2(H-h+1)$, Assumption 2.1 then implies

$$\left\|\int_{\mathcal{S}}\left(V_{h+1}^{\pi}(s) - \widehat{V}_{h+1}^{\pi}(s)\right)\boldsymbol{\mu}_h(s)\mathrm{d}s\right\|_2 \leq 2(H-h+1)\sqrt{d}.$$

Together with (D.22) and (D.23) and Cauchy-Schwartz inequality, we conclude that

$$\begin{aligned}|E_{1,h}| &= \left|\lambda(\mathbf{v}_h^{\pi})^{\top}\widehat{\mathbf{\Lambda}}_h^{-1}\int_{\mathcal{S}}\left(V_{h+1}^{\pi}(s) - \widehat{V}_{h+1}^{\pi}(s)\right)\boldsymbol{\mu}_h(s)\mathrm{d}s\right| \\ &\leq \lambda\cdot\frac{4\sqrt{2}(H-h+1)\sqrt{d}\left\|\mathbf{\Lambda}_h^{-1}\right\|^{1/2}}{K}\cdot\|\mathbf{v}_h^{\pi}\|_{\mathbf{\Lambda}_h^{-1}}\cdot\left\{1+\sqrt{2C_0(h)\cdot\frac{1}{\iota_h}\cdot\widetilde{\rho}(h)}\right\},\end{aligned}$$

and thus

$$|E_1| \leq 4\sqrt{2}\lambda\frac{H\sqrt{d}}{K}\left[\sum_{h=1}^{H}A_4(h)\right],$$

where for each h,

$$A_4(h) := \frac{H-h+1}{H}\cdot\frac{1}{\sqrt{\iota_h}}\cdot\|\mathbf{v}_h^{\pi}\|_{\mathbf{\Lambda}_h^{-1}}\cdot\left[1+\sqrt{2C_0(h)\cdot\frac{1}{\iota_h}\cdot\widetilde{\rho}(h)}\right],$$

and $C_0(h)$ and $\widetilde{\rho}(h)$ are same constants as in Lemma D.7.

□

D.3 Bounding the E_3 Term in the OPE Decomposition

It remains to bound the term E_3 in (E.3) given by:

$$E_3 := \lambda\sum_{h=1}^{H}(\mathbf{v}_h^{\pi})^{\top}\widehat{\mathbf{\Lambda}}_h^{-1}\mathbf{w}_h^{\pi} = \sum_{h=1}^{H}E_{3,h},$$

where $E_{3,h} = \lambda(\mathbf{v}_h^{\pi})^{\top}\widehat{\mathbf{\Lambda}}_h^{-1}\mathbf{w}_h^{\pi}$. Similar to Lemma D.8, we have the following lemma.

Lemma D.9. Under the same event where the result of Lemma D.7 and Lemma D.8 holds, if K satisfies (D.16), (D.17), (D.20) and (D.21), we have

$$|E_3| \leq 4\sqrt{2}\lambda\left(\sum_{h=1}^{H}A_5(h)\right)\cdot\frac{H\sqrt{d}}{K},$$

where

$$A_5(h) = \frac{1}{\sqrt{\iota_h}}\cdot\|\mathbf{v}_h^{\pi}\|_{\mathbf{\Lambda}_h^{-1}}\cdot\left\{1+\sqrt{2C_0(h)\cdot\frac{1}{\iota_h}\cdot\widetilde{\rho}(h)}\right\},$$

and $C_0(h)$ and $\widetilde{\rho}(h)$ are same constants as in Lemma D.7.

Proof of Lemma D.9. First note that

$$\begin{aligned}|E_{3,h}| &\leq \lambda\cdot\|\mathbf{v}_h^{\pi}\|_{\widehat{\mathbf{\Lambda}}_h^{-1}}\cdot\|\mathbf{w}_h^{\pi}\|_{\widehat{\mathbf{\Lambda}}_h^{-1}} \\ &\leq \lambda\cdot\|\mathbf{v}_h^{\pi}\|_{\widehat{\mathbf{\Lambda}}_h^{-1}}\cdot\|\mathbf{w}_h^{\pi}\|_2\cdot\left\|\widehat{\mathbf{\Lambda}}_h^{-1}\right\|^{1/2} \\ &\leq \lambda\cdot\frac{2}{\sqrt{K}}\cdot\left\{\|\mathbf{v}_h^{\pi}\|_{\mathbf{\Lambda}_h^{-1}} + \sqrt{2C_0(h)\cdot\left\|\mathbf{\Lambda}_h^{-1}\right\|\cdot\widetilde{\rho}(h)}\cdot\|\mathbf{v}_h^{\pi}\|_{\mathbf{\Lambda}_h^{-1}}\right\}\cdot\frac{\sqrt{2}}{\sqrt{K}}\left\|\mathbf{\Lambda}_h^{-1}\right\|^{1/2}\cdot 2H\sqrt{d} \\ &= 4\sqrt{2}\lambda\cdot\frac{1}{\sqrt{\iota_h}}\cdot\|\mathbf{v}_h^{\pi}\|_{\mathbf{\Lambda}_h^{-1}}\cdot\left\{1+\sqrt{2C_0(h)\cdot\frac{1}{\iota_h}\cdot\widetilde{\rho}(h)}\right\}\cdot\frac{H\sqrt{d}}{K},\end{aligned}$$

where the third step is by (D.22), (D.23) and Proposition 2.2. We then conclude that

$$|E_3| \le 4\sqrt{2}\lambda \left(\sum_{h=1}^{H} A_5(h)\right) \cdot \frac{H\sqrt{d}}{K},$$

where for each $h \in [H]$,

$$A_5(h) = \frac{1}{\sqrt{\iota_h}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \left\{1 + \sqrt{2C_0(h) \cdot \frac{1}{\iota_h} \cdot \widetilde{\rho}(h)}\right\}.$$

□

D.4 Proof of Theorem D.1

Proof of Theorem D.1. By (4.3), and Lemmas D.7, D.8 and D.9, we have that with probability at least $1-\delta$,

$$\begin{aligned}
|v_1^\pi - \widehat{v}_1^\pi| \le & \sqrt{2\log\left(\frac{16H}{\delta}\right) B} \cdot \left[\sum_{h=1}^{H} \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}\right] \cdot \frac{1}{\sqrt{K}} \\
& + \frac{16\sqrt{2}}{3}\log\left(\frac{16H}{\delta}\right) \cdot \sqrt{B} \cdot \left[\sum_{h=1}^{H} A_1(h)\right] \cdot \frac{1}{K^{3/4}} \\
& + \frac{16\sqrt{2}}{3}\log\left(\frac{16H}{\delta}\right) \cdot \sqrt{B} \cdot \left[\sum_{h=1}^{H} (A_2(h) + A_3(h))\right] \cdot \frac{1}{K} \\
& + 4\sqrt{2}\lambda \left[\sum_{h=1}^{H} (A_4(h) + A_5(h))\right] \cdot \frac{H\sqrt{d}}{K}.
\end{aligned} \tag{D.24}$$

We now compute a lower bound for K. This comes from the lower bound of K required by Theorem C.2, Lemma D.7, Lemma D.8 and Lemma D.9. Recall (D.21), (D.16), (D.17) and (D.20):

$$\begin{aligned}
& K \ge \max_{h\in[H]} \frac{16\left\|\mathbf{\Lambda}_h^{-1}\right\|^2}{(\eta_h + \sigma_r^2)^4} \cdot \left(C_{K,h,\delta}(H-h+1)^2\sqrt{d} + 4(H-h+1)\cdot \widetilde{C}(h)\cdot d\right)^2, \\
& K \ge \max_{h\in[H]} \max\left\{\frac{911}{(\eta_h+\sigma_r^2)^2\iota_h^2}\log\left(\frac{8Hd}{\delta}\right), \frac{6\lambda}{\iota_h}\right\}, \\
& K \ge C \cdot \frac{H^2d^2}{\kappa^2}\log\left(\frac{dHK}{\kappa\delta}\right) \cdot \max_{h\in[H]} \frac{(H-h+1)^2}{(\eta_h+\sigma_r^2)^2} \cdot \max_{h\in[H]} \frac{(H-h+1)^2}{\iota_h(\eta_h+\sigma_r^2)}, \\
& K \ge 1152 \cdot \left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d}\log\frac{16H}{\delta}\right] \cdot \frac{H^4 d}{\kappa^2\sigma_r^4}.
\end{aligned} \tag{D.25}$$

It remains to simplify the expression. For the first lower bound in (D.25), note that $\widetilde{C}(h) \ge C_{K,h,\delta}\cdot H$, and thus

$$C_{K,h,\delta}(H-h+1)^2\sqrt{d} + 4(H-h+1)\cdot\widetilde{C}(h)\cdot d < 8(H-h+1)\cdot\widetilde{C}(h)\cdot d. \tag{D.26}$$

Therefore, it suffices to let K satisfy

$$K \ge \max_{h\in[H]} C^2 \cdot \frac{1024}{\iota_h^2(\eta_h+\sigma^2)^4} \cdot (H-h+1)^2 d^2 \cdot \left[\sum_{i=h}^{H} \frac{H-h+1}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right]^2 \cdot \left[\log\left(\frac{dH^2K}{\kappa\delta}\right)\right]^2, \tag{D.27}$$

where C is the problem-independent universal constant from the proof of Theorem C.2. The second lower bound in (D.25) is much smaller than (D.27) and thus can be omitted. We then consider the third and the fourth lower bound together. They can be combined into

$$K \ge C \cdot \frac{H^2d^2}{\sigma_r^4\kappa^2} \cdot \max\left\{\max_{h\in[H]} \frac{(H-h+1)^2}{(\eta_h+\sigma_r^2)^2} \cdot \max_{h\in[H]} \frac{(H-h+1)^2}{\iota_h(\eta_h+\sigma_r^2)}, H^2\right\} \cdot \log\left(\frac{dHK}{\kappa\delta}\right). \tag{D.28}$$

Denote

$$C_{h,3} \coloneqq \frac{(H-h+1)^2}{\eta_h + \sigma_r^2}.$$

Then (D.27) is simplified to

$$K \geq C \max_{h\in[H]} \frac{C_{h,3} d^2}{\iota_h^2 (\eta_h + \sigma_r^2)^3} \cdot \left[\sum_{i=h}^{H} \sqrt{\frac{C_{h,3}}{\iota_h}}\right]^2 \cdot \left[\log\left(\frac{dH^2K}{\kappa\delta}\right)\right]^2, \tag{D.29}$$

and (D.28) can be simplified

$$K \geq C \cdot \frac{H^2 d^2}{\sigma_r^4 \kappa^2} \cdot \max\left\{\max_{h\in[H]} \frac{C_{h,3}}{\eta_h + \sigma_r^2} \cdot \max_{h\in[H]} \frac{C_{h,3}}{\iota_h}, H^2\right\} \cdot \log\left(\frac{dHK}{\kappa\delta}\right). \tag{D.30}$$

We then combine (D.29) and (D.30) and get that

$$K \geq C \cdot C_3(h) \cdot d^2 \left[\log\left(\frac{dH^2K}{\kappa\delta}\right)\right]^2,$$

where C is some problem-independent universal constant and

$$C_3(h) \coloneqq \max\left\{\max_{h\in[H]} \frac{C_{h,3}}{\iota_h^2(\eta_h+\sigma_r^2)^3} \cdot \left[\sum_{i=h}^{H}\sqrt{\frac{C_{h,3}}{\iota_h}}\right]^2, \frac{H^2}{\sigma_r^4\kappa^2} \cdot \left(\max_{h\in[H]} \frac{C_{h,3}}{\eta_h+\sigma_r^2} \cdot \max_{h\in[H]} \frac{C_{h,3}}{\iota_h}\right), \frac{H^4}{\sigma_r^4\kappa^2}\right\}.$$

To simplify the upper bound given by (D.24), first note that by the choice of K, we have $B < 2$. By (D.26), we have that $\widetilde{\rho}(h)$ satisfies

$$\widetilde{\rho}(h) \leq 8C \frac{1}{\sqrt{K}} \cdot \frac{(H-h+1)d}{(\eta_h+\sigma_r^2)^2} \cdot \left[\sum_{i=h}^{H} \frac{H-h+1}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right] \cdot \log\left(\frac{dH^2K}{\kappa\delta}\right). \tag{D.31}$$

It follows that

$$A_1(h) \leq C \cdot \sqrt{C_0(h) \cdot \frac{(H-h+1)d}{(\eta_h+\sigma_r^2)^2} \cdot \left[\sum_{i=h}^{H} \frac{H-h+1}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right] \cdot \frac{1}{\iota_h} \cdot \log\left(\frac{dH^2K}{\kappa\delta}\right)} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}, \tag{D.32}$$

for some (different) universal constant C. Also, it is not hard to see $A_2(h)$ and $A_3(h)$ are less than the RHS of (D.32) up to a constant factor by our choice of K, which gives

$$\begin{aligned}
&\sum_{h=1}^{H} A_2(h) + A_3(h) \\
&\leq C \sum_{h=1}^{H} \sqrt{C_0(h) \cdot \frac{(H-h+1)d}{(\eta_h+\sigma_r^2)^2} \cdot \left[\sum_{i=h}^{H} \frac{H-h+1}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right] \cdot \frac{1}{\iota_h} \cdot \log\left(\frac{dH^2K}{\kappa\delta}\right)} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}},
\end{aligned} \tag{D.33}$$

for some universal constant C. To bound $A_4(h) + A_5(h)$, note that $A_4(h) \leq A_5(h)$ and thus

$$A_4(h) + A_5(h) \leq \frac{2}{\sqrt{\iota_h}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \left\{1 + \sqrt{2C_0(h) \cdot \frac{1}{\iota_h} \cdot \widetilde{\rho}(h)}\right\}, \tag{D.34}$$

where

$$\begin{aligned}
&\frac{2}{\sqrt{\iota_h}} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}} \cdot \sqrt{2C_0(h) \cdot \frac{1}{\iota_h} \cdot \widetilde{\rho}(h)} \\
&\leq \frac{C}{K^{1/4}} \sqrt{C_0(h) \frac{(H-h+1)d}{(\eta_h+\sigma_r^2)^2} \cdot \left[\sum_{i=h}^{H} \frac{H-h+1}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right] \cdot \frac{1}{\iota_h^2} \cdot \log\left(\frac{dH^2K}{\kappa\delta}\right)} \cdot \|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}.
\end{aligned}$$

Recall (D.24). By our choice of K, it is clear that

$$H\sqrt{d}\cdot\frac{2}{\sqrt{\iota_h}}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}\cdot\sqrt{2C_0(h)\cdot\frac{1}{\iota_h}\cdot\widetilde{\rho}(h)}$$
$$\leq C\cdot\sqrt{C_0(h)\cdot\frac{(H-h+1)d}{(\eta_h+\sigma_r^2)^2}\cdot\left[\sum_{i=h}^{H}\frac{H-h+1}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right]\cdot\frac{1}{\iota_h}\cdot\log\left(\frac{dH^2K}{\kappa\delta}\right)}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}},$$

where the RHS of the above is exactly the RHS of (D.32) up to a constant factor. Therefore, we can combine $[\sum_{h=1}^{H}A_4(h)+A_5(h)]H\sqrt{d}$ with (D.33), and together with (D.34), the last two terms on the RHS of (D.24) can be upper bounded by

$$\frac{16\sqrt{2}}{3}\log\left(\frac{16H}{\delta}\right)\cdot\sqrt{B}\cdot\left[\sum_{h=1}^{H}(A_2(h)+A_3(h))\right]\cdot\frac{1}{K}+4\sqrt{2}\lambda\left[\sum_{h=1}^{H}(A_4(h)+A_5(h))\right]\cdot\frac{H\sqrt{d}}{K}$$
$$\leq C\cdot C_4\cdot\log\left(\frac{16H}{\delta}\right)\cdot\frac{1}{K}, \tag{D.35}$$

where C is some universal constant and C_4 is given by

$$C_4\coloneqq\sum_{h=1}^{H}\left\{\sqrt{C_0(h)\cdot\frac{(H-h+1)d}{(\eta_h+\sigma_r^2)^2}\cdot\left[\sum_{i=h}^{H}\frac{H-h+1}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right]\cdot\frac{1}{\iota_h}\cdot\log\left(\frac{dH^2K}{\kappa\delta}\right)}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}\right\}$$
$$=\sum_{h=1}^{H}\left\{\sqrt{C_0(h)\cdot C_{h,2}\cdot\frac{(H-h+1)d}{\iota_h(\eta_h+\sigma_r^2)^2}\cdot\log\left(\frac{dH^2K}{\kappa\delta}\right)}\cdot\|\mathbf{v}_h^\pi\|_{\mathbf{\Lambda}_h^{-1}}\right\}.$$

Plugging in the formula for $C_0(h)$ given in Lemma D.7 finishes the proof. Note that in Theorem C.2, the notation $C_0(h)$ is changed to $C_{h,4}$.

□

E Proof of Error Decomposition

Proof. Since $Q_h^\pi(s,a)=\boldsymbol{\phi}(s,a)^\top\mathbf{w}_h^\pi=r_h(s,a)+[\mathbb{P}_hV_{h+1}^\pi](s,a)$ for some vector $\mathbf{w}_h^\pi\in\mathbb{R}^d$, we further have

$$Q_h^\pi(s,a)=\boldsymbol{\phi}(s,a)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\left(\sum_{k=1}^{K}\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})\boldsymbol{\phi}(s_{k,h},a_{k,h})^\top}{\widehat{\sigma}_h(s_{k,h},a_{k,h})^2}+\lambda\mathbf{I}_d\right)\mathbf{w}_h^\pi$$
$$=\boldsymbol{\phi}(s,a)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\sum_{k=1}^{K}\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})}{\widehat{\sigma}_h(s_{k,h},a_{k,h})^2}Q_h^\pi(s_{k,h},a_{k,h})+\lambda\boldsymbol{\phi}(s,a)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\mathbf{w}_h^\pi$$
$$=\boldsymbol{\phi}(s,a)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\sum_{k=1}^{K}\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})}{\widehat{\sigma}_h(s_{k,h},a_{k,h})^2}\left(r_h(s_{k,h},a_{k,h})+[\mathbb{P}_hV_{h+1}^\pi](s_{k,h},a_{k,h})\right)$$
$$+\lambda\boldsymbol{\phi}(s,a)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\mathbf{w}_h^\pi.$$

It follows that

$$Q_h^\pi(s,a)-\widehat{Q}_h^\pi(s,a)=\boldsymbol{\phi}(s_h,a_h)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\sum_{k=1}^{K}\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})}{\widehat{\sigma}_h(s_{k,h},a_{k,h})^2}\left([\mathbb{P}_hV_{h+1}^\pi](s_{k,h},a_{k,h})-\widehat{V}_{h+1}^\pi(s'_{k,h})-\epsilon_{k,h}\right)$$
$$+\lambda\boldsymbol{\phi}(s_h,a_h)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\mathbf{w}_h^\pi$$
$$=\boldsymbol{\phi}(s,a)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\sum_{k=1}^{K}\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})}{\widehat{\sigma}_h(s_{k,h},a_{k,h})^2}\left([\mathbb{P}_hV_{h+1}^\pi](s_{k,h},a_{k,h})-[\mathbb{P}_h\widehat{V}_{h+1}^\pi](s_{k,h},a_{k,h})\right)$$
$$+\boldsymbol{\phi}(s,a)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\sum_{k=1}^{K}\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})}{\widehat{\sigma}_h(s_{k,h},a_{k,h})^2}\left([\mathbb{P}_h\widehat{V}_{h+1}^\pi](s_{k,h},a_{k,h})-\widehat{V}_{h+1}^\pi(s'_{k,h})-\epsilon_{k,h}\right)$$
$$+\lambda\boldsymbol{\phi}(s,a)^\top\widehat{\mathbf{\Lambda}}_h^{-1}\mathbf{w}_h^\pi.$$

where $\epsilon_{k,h}$ is the noise in reward. Note that

$$[\mathbb{P}_h V_{h+1}^\pi](s_{k,h}, a_{k,h}) - [\mathbb{P}_h \widehat{V}_{h+1}^\pi](s_{k,h}, a_{k,h}) = \int_{\mathcal{S}} \left(V_{h+1}^\pi(s) - \widehat{V}_{h+1}^\pi(s)\right) \langle \boldsymbol{\phi}(s_{k,h}, a_{k,h}), \boldsymbol{\mu}_h(s)\rangle \mathrm{d}s$$
$$= \boldsymbol{\phi}(s_{k,h}, a_{k,h})^\top \int_{\mathcal{S}} \left(V_{h+1}^\pi(s) - \widehat{V}_{h+1}^\pi(s)\right) \boldsymbol{\mu}_h(s) \mathrm{d}s,$$

and thus

$$\begin{aligned}
&Q_h^\pi(s,a) - \widehat{Q}_h^\pi(s,a) \\
&= \boldsymbol{\phi}(s,a)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}(s_{k,h}, a_{k,h})^\top}{\widehat{\sigma}(s_{k,h}, a_{k,h})^2} \int_{\mathcal{S}} \left(V_{h+1}^\pi(s') - \widehat{V}_{h+1}^\pi(s')\right) \boldsymbol{\mu}_h(s') \mathrm{d}s' \\
&\quad + \boldsymbol{\phi}(s,a)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2} \left([\mathbb{P}_h \widehat{V}_{h+1}^\pi](s_{k,h}, a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h}) - \epsilon_{k,h}\right) \\
&\quad + \lambda \boldsymbol{\phi}(s_h, a_h)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \mathbf{w}_h^\pi \\
&= [\mathbb{P}_h(V_{h+1}^\pi - \widehat{V}_{h+1}^\pi)](s,a) - \lambda \boldsymbol{\phi}(s,a)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \int_{\mathcal{S}} \left(V_{h+1}^\pi(s) - \widehat{V}_{h+1}^\pi(s)\right) \boldsymbol{\mu}_h(s) \mathrm{d}s \\
&\quad + \boldsymbol{\phi}(s,a)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2} \left([\mathbb{P}_h \widehat{V}_{h+1}^\pi](s_{k,h}, a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h}) - \epsilon_{k,h}\right) \\
&\quad + \lambda \boldsymbol{\phi}(s,a)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \mathbf{w}_h^\pi. \qquad \text{(E.1)}
\end{aligned}$$

Then by the Bellman equation, we have

$$\begin{aligned}
V_h^\pi(s) - \widehat{V}_h^\pi(s) &= \mathbb{J}_h(Q_h^\pi - \widehat{Q}_h^\pi)(s) \\
&= \mathbb{J}_h \mathbb{P}_h(V_{h+1}^\pi - \widehat{V}_{h+1}^\pi)(s) - \lambda \mathbb{J}_h \boldsymbol{\phi}(s)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \int_{\mathcal{S}} \left(V_{h+1}^\pi(s') - \widehat{V}_{h+1}^\pi(s')\right) \boldsymbol{\mu}_h(s') \mathrm{d}s' \\
&\quad + \mathbb{J}_h \boldsymbol{\phi}(s)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2} \left([\mathbb{P}_h \widehat{V}_{h+1}^\pi](s_{k,h}, a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h}) - \epsilon_{k,h}\right) \\
&\quad + \lambda \mathbb{J}_h \boldsymbol{\phi}(s)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \mathbf{w}_h^\pi, \qquad \text{(E.2)}
\end{aligned}$$

where $\mathbb{J}_h f(\cdot) = \int_{\mathcal{A}} f(\cdot, a)\pi_h(a|\cdot)\mathrm{d}a$ for any function $f : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$. Recursively expanding the above equation, we obtain

$$\begin{aligned}
&V_1^\pi(s) - \widehat{V}_1^\pi(s) \\
&= -\lambda \sum_{h=1}^H \left(\prod_{i=1}^{h-1} \mathbb{J}_i \mathbb{P}_i\right) \mathbb{J}_h \boldsymbol{\phi}(s)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \int_{\mathcal{S}} \left(V_{h+1}^\pi(s') - \widehat{V}_{h+1}^\pi(s')\right) \boldsymbol{\mu}_h(s') \mathrm{d}s' \\
&\quad + \sum_{h=1}^H \left(\prod_{i=1}^{h-1} \mathbb{J}_i \mathbb{P}_i\right) \mathbb{J}_h \boldsymbol{\phi}(s)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2} \left([\mathbb{P}_h \widehat{V}_{h+1}^\pi](s_{k,h}, a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h}) - \epsilon_{k,h}\right) \\
&\quad + \lambda \sum_{h=1}^H \left(\prod_{i=1}^{h-1} \mathbb{J}_i \mathbb{P}_i\right) \mathbb{J}_h \boldsymbol{\phi}(s_1)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \mathbf{w}_h^\pi. \qquad \text{(E.3)}
\end{aligned}$$

Here with a slight abuse of notation we define $\prod_{i=1}^{h-1} \mathbb{J}_i \mathbb{P}_i = 1$ when $h = 1$. We then have

$$\begin{aligned}
v_1^\pi - \widehat{v}_1^\pi &= -\lambda \sum_{h=1}^H (\mathbf{v}_h^\pi)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \int_{\mathcal{S}} \left(V_{h+1}^\pi(s) - \widehat{V}_{h+1}^\pi(s)\right) \boldsymbol{\mu}_h(s) \mathrm{d}s \\
&\quad + \sum_{h=1}^H (\mathbf{v}_h^\pi)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2} \left([\mathbb{P}_h \widehat{V}_{h+1}^\pi](s_{k,h}, a_{k,h}) - \widehat{V}_{h+1}^\pi(s'_{k,h}) - \epsilon_{k,h}\right) \\
&\quad + \lambda \sum_{h=1}^H (\mathbf{v}_h^\pi)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \mathbf{w}_h^\pi \\
&:= E_1 + E_2 + E_3, \qquad \text{(E.4)}
\end{aligned}$$

where for simplicity we write $\mathbf{v}_h^\pi = \mathbb{E}_\pi\left[\left(\prod_{i=1}^{h-1}\mathbb{J}_i\mathbb{P}_i\right)\mathbb{J}_h\boldsymbol{\phi}(s_1)\middle| s_1\sim\xi_1\right] = \mathbb{E}_{\pi,h}[\boldsymbol{\phi}(s_h,a_h)]$ by recalling the definition of $\mathbb{E}_{\pi,h}[\cdot]$ given in the text following (2.4). □

F Lemmas for Uniform Convergence

All lemmas in this section are under the Assumption of Theorem C.2.

F.1 Convergence of $\widehat{\sigma}$

Lemma F.1. For any $h\in[H]$ and any $\widehat{V}_{h+1}^\pi\in\mathcal{V}_{h+1}(L)$, with probability at least $1-\delta$, it holds for all $(s,a)\in\mathcal{S}\times\mathcal{A}$ that

$$\left|\langle\boldsymbol{\phi}(s,a),\widehat{\boldsymbol{\beta}}_h\rangle_{[0,(H-h+1)^2]}-\mathbb{P}_h(\widehat{V}_{h+1})^2(s,a)\right|\le C'_{K,\delta}\cdot\frac{(H-h+1)^2\sqrt{d}}{\sqrt{K}}\left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right)+\frac{1}{d}\log\frac{4}{\delta}\right]^{1/2},$$

and

$$\left|\langle\boldsymbol{\phi}(s,a),\widehat{\boldsymbol{\theta}}_h\rangle_{[0,H-h+1]}-\mathbb{P}_h(\widehat{V}_{h+1})(s,a)\right|\le C'_{K,\delta}\cdot\frac{(H-h+1)\sqrt{d}}{\sqrt{K}}\left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right)+\frac{1}{d}\log\frac{4}{\delta}\right]^{1/2},$$

where

$$C'_{K,\delta}\coloneqq 4\sqrt{2}\frac{1}{\sqrt{\kappa_h}}+4\lambda\cdot\frac{1}{\kappa_h}\cdot\left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right)+\frac{1}{d}\log\frac{4}{\delta}\right]^{-1/2}.$$

Proof of Lemma F.1. First we consider $\langle\boldsymbol{\phi}(s,a),\widehat{\boldsymbol{\beta}}_h\rangle_{[0,(H-h+1)^2]}$. Note that, since $\mathbb{P}_h(\widehat{V}_{h+1}^\pi)^2(s,a)\in[0,(H-h+1)^2]$,

$$|\langle\boldsymbol{\phi}(s,a),\widehat{\boldsymbol{\beta}}_h\rangle_{[0,(H-h+1)^2]}-\mathbb{P}_h(\widehat{V}_{h+1}^\pi)^2(s,a)|\le|\langle\boldsymbol{\phi}(s,a),\widehat{\boldsymbol{\beta}}_h\rangle-\mathbb{P}_h(\widehat{V}_{h+1}^\pi)^2(s,a)|.$$

It then suffices to bound the RHS.

$$\begin{aligned}
&\langle\boldsymbol{\phi}(s,a),\widehat{\boldsymbol{\beta}}_h\rangle-\mathbb{P}_h(\widehat{V}_{h+1})^2(s,a)\\
&=\boldsymbol{\phi}(s,a)^\top(\widehat{\boldsymbol{\Sigma}}_h)^{-1}\sum_{k=1}^K\boldsymbol{\phi}(\check{s}_{k,h},\check{a}_{k,h})\widehat{V}_{h+1}^\pi(\check{s}'_{k,h})^2-\mathbb{P}_h(\widehat{V}_{h+1}^\pi)^2(s,a)\\
&=\boldsymbol{\phi}(s,a)^\top(\widehat{\boldsymbol{\Sigma}}_h)^{-1}\sum_{k=1}^K\boldsymbol{\phi}(\check{s}_{k,h},\check{a}_{k,h})\widehat{V}_{h+1}^\pi(\check{s}'_{k,h})^2-\boldsymbol{\phi}(s,a)^\top\int_{\mathcal{S}}(\widehat{V}_{h+1}^\pi)^2(s')\mathrm{d}\boldsymbol{\mu}_h(s').
\end{aligned}$$

Note that

$$\begin{aligned}
&\boldsymbol{\phi}(s,a)^\top\int_{\mathcal{S}}(\widehat{V}_{h+1}^\pi)^2(s')\mathrm{d}\boldsymbol{\mu}_h(s')\\
&=\boldsymbol{\phi}(s,a)^\top(\widehat{\boldsymbol{\Sigma}}_h)^{-1}\left(\sum_{k=1}^K\boldsymbol{\phi}(\check{s}_{k,h},\check{a}_{k,h})\boldsymbol{\phi}(\check{s}_{k,h},\check{a}_{k,h})^\top+\lambda\mathbf{I}_d\right)\int_{\mathcal{S}}(\widehat{V}_{h+1}^\pi)^2(s')\mathrm{d}\boldsymbol{\mu}_h(s')\\
&=\boldsymbol{\phi}(s,a)^\top(\widehat{\boldsymbol{\Sigma}}_h)^{-1}\sum_{k=1}^K\boldsymbol{\phi}(\check{s}_{k,h},\check{a}_{k,h})\mathbb{P}_h(\widehat{V}_{h+1}^\pi)^2(\check{s}_{k,h},\check{a}_{k,h})+\lambda\boldsymbol{\phi}(s,a)(\widehat{\boldsymbol{\Sigma}}_h)^{-1}\int_{\mathcal{S}}(\widehat{V}_{h+1}^\pi)^2(s')\mathrm{d}\boldsymbol{\mu}_h(s'),
\end{aligned}$$

and it follows that

$$\begin{aligned}
&\langle\boldsymbol{\phi}(s,a),\widehat{\boldsymbol{\beta}}_h\rangle-\mathbb{P}_h(\widehat{V}_{h+1})^2(s,a)\\
&=\underbrace{\boldsymbol{\phi}(s,a)^\top(\widehat{\boldsymbol{\Sigma}}_h)^{-1}\sum_{k=1}^K\boldsymbol{\phi}(\check{s}_{k,h},\check{a}_{k,h})\left[(\widehat{V}_{h+1}^\pi)^2(\check{s}'_{k,h})-\mathbb{P}_h(\widehat{V}_{h+1}^\pi)^2(\check{s}_{k,h},\check{a}_{k,h})\right]}_{A_1(s,a)}\\
&\quad\underbrace{-\lambda\boldsymbol{\phi}(s,a)^\top(\widehat{\boldsymbol{\Sigma}}_h)^{-1}\int_{\mathcal{S}}(\widehat{V}_{h+1}^\pi)^2(s')\mathrm{d}\boldsymbol{\mu}_h(s')}_{A_2(s,a)}.
\end{aligned}$$

To bound $|A_1|$, we first apply Cauchy-Schwartz inequality to obtain that

$$|E_1(s,a)| \le \|\boldsymbol{\phi}(s,a)\|_{\widehat{\boldsymbol{\Sigma}}_h^{-1}} \cdot \left\| \sum_{k=1}^{K} \boldsymbol{\phi}(\check{s}_{k,H}, \check{a}_{k,H}) \left[(\widehat{V}_{h+1}^{\pi})^2(\check{s}'_{k,h}) - \mathbb{P}_h(\widehat{V}_{h+1}^{\pi})^2(\check{s}_{k,h}, \check{a}_{k,h}) \right] \right\|_{\widehat{\boldsymbol{\Sigma}}_h^{-1}}.$$

By Lemma H.5, if K satisfies

$$K \ge \max\left\{ 512 \|\boldsymbol{\Sigma}_h^{-1}\|^2 \log\left(\frac{4d}{\delta}\right), 4\lambda \|\boldsymbol{\Sigma}_h^{-1}\| \right\}, \tag{F.1}$$

then with probability at least $1-\delta/2$, for all $(s,a) \in \mathcal{S} \times \mathcal{A}$,

$$\|\boldsymbol{\phi}(s,a)\|_{\widehat{\boldsymbol{\Sigma}}_h^{-1}} \le \frac{2}{\sqrt{K}} \cdot \|\boldsymbol{\phi}(s,a)\|_{\boldsymbol{\Sigma}_h^{-1}}.$$

By Lemma G.5, for fixed $\widehat{V}_{h+1}^{\pi}$, with probability at least $1-\delta/2$, we have

$$\begin{aligned} &\left\| \sum_{k=1}^{K} \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h}) \left[(\widehat{V}_{h+1}^{\pi})^2(\check{s}'_{k,h}) - \mathbb{P}_h(\widehat{V}_{h+1}^{\pi})^2(\check{s}_{k,h}, \check{a}_{k,h}) \right] \right\|_{\widehat{\boldsymbol{\Sigma}}_h^{-1}} \\ &\le 2\sqrt{2}(H-h+1)^2 \left[\frac{d}{2} \log\left(\frac{\lambda+K}{\lambda}\right) + \log\frac{4}{\delta} \right]^{1/2}. \end{aligned}$$

Combining the two inequalities above, we have that, with probability at least $1-\delta$,

$$\begin{aligned} |A_1(s,a)| &\le 2\sqrt{2}(H-h+1)^2 \left[\frac{d}{2} \log\left(\frac{\lambda+K}{\lambda}\right) + \log\frac{4}{\delta} \right]^{1/2} \cdot \frac{2}{\sqrt{K}} \cdot \|\boldsymbol{\phi}(s,a)\|_{\Sigma_h^{-1}} \\ &\le 4\sqrt{2}\|\Sigma_h^{-1}\|^{1/2} \left[\frac{1}{2} \log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d} \log\frac{4}{\delta} \right]^{1/2} \cdot \frac{(H-h+1)^2\sqrt{d}}{\sqrt{K}}, \end{aligned}$$

for all (s,a). At the same time, we can bound A_2 as

$$\begin{aligned} |A_2(s,a)| &\le \lambda \|\boldsymbol{\phi}(s,a)\|_{\widehat{\boldsymbol{\Sigma}}_h^{-1}} \cdot \left\| \int_{\mathcal{S}} (\widehat{V}_{h+1}^{\pi})^2(s') \mathrm{d}\boldsymbol{\mu}_h(s') \right\|_{\widehat{\boldsymbol{\Sigma}}_h^{-1}} \\ &\le \lambda \cdot \frac{2}{\sqrt{K}} \|\boldsymbol{\phi}(s,a)\|_{\boldsymbol{\Sigma}_h^{-1}} \cdot \frac{2}{\sqrt{K}} \left\| \int_{\mathcal{S}} (\widehat{V}_{h+1}^{\pi})^2(s') \mathrm{d}\boldsymbol{\mu}_h(s') \right\|_{\boldsymbol{\Sigma}_h^{-1}} \\ &\le 4\lambda \left\|\boldsymbol{\Sigma}_h^{-1}\right\| \cdot \frac{(H-h+1)^2\sqrt{d}}{K}, \end{aligned}$$

where the last step is by Assumption 2.1. We then conclude that, if K satisfies (F.1), then with probability at least $1-\delta$, for all (s,a),

$$\begin{aligned} &\left| \langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\beta}}_h \rangle - \mathbb{P}_h(\widehat{V}_{h+1})^2(s,a) \right| \\ &\le |A_1(s,a)| + |A_2(s,a)| \\ &\le 4\sqrt{2}\|\boldsymbol{\Sigma}_h^{-1}\|^{1/2} \left[\frac{1}{2} \log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d} \log\frac{4}{\delta} \right]^{1/2} \cdot \frac{(H-h+1)^2\sqrt{d}}{\sqrt{K}} + 4\lambda \left\|\boldsymbol{\Sigma}_h^{-1}\right\| \cdot \frac{(H-h+1)^2\sqrt{d}}{K} \\ &= 4\sqrt{2} \frac{1}{\sqrt{\kappa_h}} \left[\frac{1}{2} \log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d} \log\frac{4}{\delta} \right]^{1/2} \cdot \frac{(H-h+1)^2\sqrt{d}}{\sqrt{K}} + \frac{4\lambda}{\kappa_h} \cdot \frac{(H-h+1)^2\sqrt{d}}{K}, \end{aligned}$$

where in the last step we use the definition $\kappa_h := \lambda_{\min}(\boldsymbol{\Sigma}_h)$. Note that by Assumption 2.3, we have $\kappa_h > 0$ for all $h \in [H]$. At the same time, we can bound $\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\theta}}_h \rangle_{[0,H-h+1]} - \mathbb{P}_h(\widehat{V}_{h+1})(s,a)$ in a similar way as

$$\begin{aligned} &\left| \langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\theta}}_h \rangle_{[0,H-h+1]} - \mathbb{P}_h(\widehat{V}_{h+1})(s,a) \right| \\ &\le 4\sqrt{2} \frac{1}{\sqrt{\kappa_h}} \left[\frac{1}{2} \log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d} \log\frac{4}{\delta} \right]^{1/2} \cdot \frac{(H-h+1)\sqrt{d}}{\sqrt{K}} + \frac{4\lambda}{\kappa_h} \cdot \frac{(H-h+1)\sqrt{d}}{K}. \end{aligned}$$

□

Lemma F.2. For any $h \in [H]$ and any $\widehat{V}_{h+1}^{\pi} \in \mathcal{V}_{h+1}(L)$, with probability at least $1-\delta$, it holds for all $(s,a) \in \mathcal{S} \times \mathcal{A}$ that

$$\begin{aligned}\left|\widehat{\sigma}_h^2(s,a) - \sigma_r^2 - \max\left\{\eta_h,\ \mathbb{V}_h \widehat{V}_{h+1}^{\pi}(s,a)\right\}\right| &\leq \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}} \\ &\leq \frac{20(H-h+1)^2\sqrt{d}}{\kappa_h\sqrt{K}} \cdot \sqrt{\log\left(\frac{K}{\lambda\delta}\right)},\end{aligned}$$

where

$$C_{K,h,\delta} = 12\sqrt{2}\frac{1}{\sqrt{\kappa_h}} \cdot \left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d}\log\frac{4}{\delta}\right]^{1/2} + 12\lambda\frac{1}{\kappa_h}. \tag{F.2}$$

Proof of Lemma F.2. Recall that by definition,

$$\mathbb{V}_h \widehat{V}_{h+1}^{\pi}(s,a) = \mathbb{P}_h(\widehat{V}_{h+1}^{\pi})^2(s,a) - \left(\mathbb{P}_h \widehat{V}_{h+1}^{\pi}(s,a)\right)^2.$$

We then have

$$\begin{aligned}&\left|\left[\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\beta}}_h\rangle_{[0,(H-h+1)^2]} - \left(\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\theta}}_h\rangle_{[0,H-h+1]}\right)^2\right] - \mathbb{V}_h \widehat{V}_{h+1}^{\pi}(s,a)\right| \\ &\leq \left|\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\beta}}_h\rangle_{[0,(H-h+1)^2]} - \mathbb{P}_h(\widehat{V}_{h+1}^{\pi})^2(s,a)\right| + 2(H-h+1)\cdot\left|\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\theta}}_h\rangle_{[0,H-h+1]} - \mathbb{P}_h \widehat{V}_{h+1}^{\pi}(s,a)\right|,\end{aligned}$$

and the rest follows from Lemma F.1 and the fact that $\max\{\eta_h, \cdot\}$ is a contraction mapping. □

Lemma F.3. For any $h \in [H-1]$, let $V \in \mathcal{V}_{h+1}(L) \cap \{V : \sup_{s\in\mathcal{S}} |V(s) - V_{h+1}^{\pi}(s)| \leq \rho\}$ for some sufficiently small $\rho < (\eta_h + \sigma_r^2)/[12(H-h+1)]$. Suppose K satisfies that

$$K \geq \frac{3600(H-h+1)^4 d}{\kappa_h^2 \inf_{s,a}\sigma_h(s,a)^2} \cdot \log\left(\frac{Kd}{\lambda\delta}\right) \tag{F.3}$$

Then for any $\delta \in (0,1)$, it holds with probability at least $1-\delta$ that

$$\left\|\left(\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}\right)^{-1}\right\| \leq \frac{4}{\iota_h}.$$

Proof of Lemma F.3. By Lemma F.2, there exists an event $\check{\mathcal{E}}$ over $\{(\check{s}_{k,h}, \check{a}_{k,h}), k \in [K]\}$ such that $\mathbb{P}(\check{\mathcal{E}}) \geq 1-\delta$ and on this event it holds for all $(s,a) \in \mathcal{S}\times\mathcal{A}$ that

$$\left|\widehat{\sigma}_h^2(s,a) - \sigma_r^2 - \max\left\{\eta_h,\ \mathbb{V}_h V_{h+1}^{\pi}(s,a)\right\}\right| \leq \frac{20(H-h+1)^2\sqrt{d}}{\kappa_h\sqrt{K}}\sqrt{\log\left(\frac{K}{\lambda\delta}\right)} + 4(H-h+1)\cdot\rho.$$

Then by (F.3) and the assumption on ρ, we have

$$\frac{1}{3}\sigma_h(s,a) \leq \widehat{\sigma}_h(s,a) \leq \frac{5}{3}\sigma_h(s,a) \tag{F.4}$$

for all $(s,a) \in \mathcal{S}\times\mathcal{A}$. In the following argument we condition on $\check{\mathcal{E}}$, and this will not affect the distribution of $\{(s_{k,h}, a_{k,h}), k \in [K]\}$ by independence.

Recall that

$$\widehat{\boldsymbol{\Lambda}}_h = \sum_{k=1}^{K} \widehat{\sigma}_h(s_{k,h}, a_{k,h})^{-2}\boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}(s_{k,h}, a_{k,h})^{\top} + \lambda\mathbf{I}_d.$$

Since $\widehat{\sigma}_h \geq \inf_{s,a}\sigma_h(s,a)/3$, it then follows from Lemma H.1 that

$$\left\|\frac{\widehat{\boldsymbol{\Lambda}}_h}{K} - \mathbb{E}\left[\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}\right]\right\| \leq \frac{12\sqrt{2}}{\sqrt{K}\cdot\inf_{s,a}\sigma_h(s,a)^2} \cdot \sqrt{\log\left(\frac{2d}{\delta}\right)}. \tag{F.5}$$

To bound $\|(\widehat{\mathbf{\Lambda}}_h/K)^{-1}\|$, we use the fact that

$$\begin{aligned}\left\|\left(\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right)^{-1}\right\| &\le \left\|\mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]^{-1}\right\| + \left\|\left(\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right)^{-1} - \mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]^{-1}\right\| \\ &\le \left\|\mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]^{-1}\right\| + \left\|\left(\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right)^{-1}\right\| \cdot \left\|\mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]^{-1}\right\| \cdot \left\|\frac{\widehat{\mathbf{\Lambda}}_h}{K} - \mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]\right\|,\end{aligned}$$

which implies

$$\begin{aligned}\left\|\left(\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right)^{-1}\right\| &\le \left\|\mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]^{-1}\right\| \cdot \left(1 - \left\|\mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]^{-1}\right\| \cdot \left\|\frac{\widehat{\mathbf{\Lambda}}_h}{K} - \mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]\right\|\right)^{-1} \\ &\le \left(\left\|\mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]^{-1}\right\|^{-1} - \left\|\frac{\widehat{\mathbf{\Lambda}}_h}{K} - \mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right]\right\|\right)^{-1} \qquad \text{(F.6)}\end{aligned}$$

Note that by (F.4), we have

$$\begin{aligned}\mathbb{E}\left[\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right] &= \frac{1}{K}\sum_{k=1}^{K}\mathbb{E}\left[\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})\boldsymbol{\phi}(s_{k,h},a_{k,h})^\top}{\widehat{\sigma}_h(s_{k,h},a_{k,h})^2}\right] + \frac{\lambda}{K}\mathbf{I}_d \\ &\succeq \frac{1}{2K}\sum_{k=1}^{K}\mathbb{E}\left[\frac{\boldsymbol{\phi}(s_{k,h},a_{k,h})\boldsymbol{\phi}(s_{k,h},a_{k,h})^\top}{\sigma_h(s,a)^2}\right] + \frac{\lambda}{K}\mathbf{I}_d \\ &= \frac{1}{2}\mathbf{\Lambda}_h + \frac{\lambda}{K}\mathbf{I}_d. \qquad \text{(F.7)}\end{aligned}$$

Finally combining (F.5), (F.6) and (F.7) yields

$$\left\|\left(\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right)^{-1}\right\| \le \frac{1}{\left(\frac{\iota_h}{2} + \frac{\lambda}{K}\right) - \frac{12\sqrt{2}}{\sqrt{K}\cdot\inf_{s,a}\sigma_h(s,a)^2}\cdot\sqrt{\log\left(\frac{2d}{\delta}\right)}} \le \frac{4}{\iota_h},$$

where the second inequality follows from (F.3). □

Lemma F.4. For any $h \in [H-1]$, let ρ be some positive constant such that $\rho < (\eta_h + \sigma_r^2)/[12(H-h+1)]$. For any $\delta \in (0,1)$, suppose K satisfies that

$$K \ge \frac{3600(H-h+1)^4 d^2}{\kappa_h^2 \inf_{s,a}\sigma_h(s,a)^2}\cdot\log\left(\frac{d(H-h+1)KL}{\iota_h\kappa_h\lambda\delta}\right). \qquad \text{(F.8)}$$

Then it holds with probability at least $1-\delta$ that

$$\left\|\left(\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right)^{-1}\right\| \le \frac{8}{\iota_h}.$$

for all $V \in \mathcal{V}_{h+1}(L) \cap \{V : \sup_{s\in\mathcal{S}}|V(s) - V_{h+1}^{\pi}(s)| \le \rho\}$.

Proof of Lemma F.4. Let $\epsilon > 0$ be a constant to be determined later and $\mathcal{C}_V$ be a $\epsilon-$cover of $\mathcal{V}_{h+1}(L) \cap \{V : \sup_{s\in\mathcal{S}}|V(s) - V_{h+1}^{\pi}(s)| \le \rho\}$. By Lemma F.3, the choice of K in (F.8) and a union bound, we have

$$\left\|\left(\frac{\widehat{\mathbf{\Lambda}}_h}{K}\right)^{-1}\right\| \le \frac{4}{\iota_h} \qquad \text{(F.9)}$$

for all $V \in \mathcal{C}_V$, given K satisfies that

$$K \ge \frac{3600(H-h+1)^4 d}{\kappa_h^2 \inf_{s,a}\sigma_h(s,a)^2}\cdot\log\left(\frac{Kd\mathcal{N}_\epsilon}{\lambda\delta}\right) \qquad \text{(F.10)}$$

where $\mathcal{N}_\epsilon$ is the $\epsilon-$covering number of $\mathcal{V}_{h+1}(L) \cap \{V : \sup_{s\in\mathcal{S}} |V(s) - V^\pi_{h+1}(s)| \le \rho\}$.

For any $V_1 \in \mathcal{V}_{h+1}(L) \cap \left\{V : \sup_{s\in\mathcal{S}} |V(S) - V^\pi_{h+1}(s)| \le \rho\right\}$, there exists $V_2 \in \mathcal{C}_V$ such that $\sup_{s\in\mathcal{S}} |V_1(s) - V_2(s)| \le \epsilon$. Let σ_1 and $\widehat{\mathbf{\Lambda}}_{h,1}$ be the variance estimator and the weighted covariance induced by V_1, and σ_2 and $\widehat{\mathbf{\Lambda}}_{h,2}$ that of V_2. Then we have

$$
\begin{aligned}
&\left|\sigma_1^2(s,a) - \sigma_2^2(s,a)\right| \\
&\le \left|\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\beta}}_{h,1} - \widehat{\boldsymbol{\beta}}_{h,2}\rangle\right| + 2(H-h+1)\left|\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\theta}}_{h,1} - \widehat{\boldsymbol{\theta}}_{h,2}\rangle\right| \\
&\le \left|\widehat{\mathbf{\Sigma}}_h^{-1}\sum_{k=1}^K \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})(V_1^2(\check{s}'_{k,h}) - V_2^2(\check{s}'_{k,h}))\right| + \left|\widehat{\mathbf{\Sigma}}_h^{-1}\sum_{k=1}^K \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})(V_1(\check{s}'_{k,h}) - V_2(\check{s}'_{k,h}))\right| \\
&\le \frac{4(H-h+2)^2K}{\kappa_h}\cdot\epsilon,
\end{aligned}
\tag{F.11}
$$

where the second inequality is due to Assumption 2.1 and the third inequality follows from the fact that $V_1, V_2 \in \mathcal{V}_{h+1}(L)$.

Therefore, we can bound the difference between $\widehat{\mathbf{\Lambda}}_{h,1}$ and $\widehat{\mathbf{\Lambda}}_{h,2}$ as follows.

$$
\begin{aligned}
\left\|\frac{\widehat{\mathbf{\Lambda}}_{h,1}}{K} - \frac{\widehat{\mathbf{\Lambda}}_{h,2}}{K}\right\| &= \left\|\frac{1}{K}\sum_{k=1}^K \boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}(s_{k,h}, a_{k,h})^\top \cdot \frac{\sigma_1(s,a)^2 - \sigma_2(s,a)^2}{\sigma_1(s,a)^2\sigma_2(s,a)^2}\right\| \\
&\le \frac{1}{K}\sum_{k=1}^K \frac{|\sigma_1(s,a)^2 - \sigma_2(s,a)^2|}{\sigma_1(s,a)^2\sigma_2(s,a)^2} \\
&\le \frac{4(H-h+2)^2K}{\kappa_h(\eta_h+\sigma_r^2)^2}\cdot\epsilon,
\end{aligned}
\tag{F.12}
$$

where the first inequality follows from Assumption 2.1, and the second inequality is due to (F.11). When ϵ is small enough, by (F.9) we have

$$
\lambda_{\min}(\widehat{\mathbf{\Lambda}}_{h,1}/K) \ge \lambda_{\min}(\widehat{\mathbf{\Lambda}}_{h,2}/K) - \|\widehat{\mathbf{\Lambda}}_{h,1} - \widehat{\mathbf{\Lambda}}_{h,2}\|/K \ge \frac{\iota_h}{4} - \frac{4(H-h+2)^2K}{\kappa_h(\eta_h+\sigma_r^2)}\cdot\epsilon,
$$

which further implies that

$$
\left\|\left(\frac{\widehat{\mathbf{\Lambda}}_{h,1}}{K}\right)^{-1}\right\| \le \left(\frac{\iota_h}{4} - \frac{4(H-h+2)^2K}{\kappa_h(\eta_h+\sigma_r^2)}\cdot\epsilon\right)^{-1} \le \frac{8}{\iota_h}
$$

if we choose $\epsilon = \iota_h\kappa_h(\eta_h+\sigma_r^2)/[32(H-h+2)^2K]$. In this case, by Lemma H.13, we have

$$
\log\mathcal{N}_\epsilon \le d\cdot\left(1 + \frac{64L(H-h+2)^2K}{\iota_h\kappa_h(\eta_h+\sigma_r^2)}\right). \tag{F.13}
$$

Therefore, by (F.10), (F.12) and (F.13), it suffices to choose K such that

$$
K \ge \frac{3600(H-h+1)^4d^2}{\kappa_h^2\inf_{s,a}\sigma_h(s,a)^2}\cdot\log\left(\frac{d(H-h+1)KL}{\iota_h\kappa_h\lambda\delta}\right).
$$

□

F.2 Bernstein Inequality for the Self-Normalized Martingales

Lemma F.5. For any $h \in [H-1]$ and any fixed $\widehat{V}^\pi_{h+1} \in \mathcal{V}_{h+1}(L)$, let $\widehat{\sigma}_h$ be as defined in Line 6 of Algorithm 2 and $\widehat{\mathbf{\Lambda}}_h$ be as defined in (3.3). Suppose K satisfies that

$$
K \ge \frac{1600(H-h+1)^4d}{\kappa_h^2(\eta_h+\sigma_r^2)^2}\cdot\log\left(\frac{K}{\lambda\delta}\right) \tag{F.14}
$$

Then for any $\delta \in (0,1)$, it holds with probability at least $1-\delta$ that

$$\left\|\sum_{k=1}^{K} \widehat{\sigma}_h(s_{k,h},a_{k,h})^{-2}\boldsymbol{\phi}(s_{k,h},a_{k,h})\left(\mathbb{P}_h\widehat{V}^{\pi}_{h+1}(s_{k,h},a_{k,h})-\widehat{V}^{\pi}_{h+1}(s'_{k,h})-\epsilon_{k,h}\right)\right\|_{\widehat{\boldsymbol{\Lambda}}_h^{-1}}$$
$$\leq \sqrt{2d\log\left(1+\frac{K}{\lambda d(\eta_h+\sigma_r^2)}\right)\cdot\log\left(\frac{4K^2}{\delta}\right)}+\frac{4(2H-2h+3)}{\sqrt{\eta_h+\sigma_r^2}}\log\left(\frac{4K^2}{\delta}\right)$$

Proof of Lemma F.5. Let $\check{\mathcal{E}}$ be the event given by Lemma F.2, on which it holds for all $(s,a)\in\mathcal{S}\times\mathcal{A}$ that

$$\left|\widehat{\sigma}_h^2(s,a)-\sigma_r^2-\max\{\eta_h,\mathbb{V}_h\widehat{V}^{\pi}_{h+1}(s,a)\}\right|\leq\frac{20(H-h+1)^2\sqrt{d}}{\kappa_h\sqrt{K}}\cdot\sqrt{\log\left(\frac{K}{\lambda\delta}\right)}. \tag{F.15}$$

Now conditioning on $\check{\mathcal{E}}$, it will not affect the distribution of $\{(s_{k,h},a_{k,h}),k\in[K]\}$ by independence. In the following argument, we omit the explicit notation for conditioning on $\check{\mathcal{E}}$ for simplicity.

Define $\mathbf{x}_k=\phi(s_{k,h},a_{k,h})/\widehat{\sigma}_h(s_{k,h},a_{k,h})$, which is a deterministic function of $(s_{k,h},a_{k,h})$ since $\widehat{\boldsymbol{\theta}}_h$ and $\widehat{\boldsymbol{\beta}}_h$ are fixed. Define $\zeta_k=\left(\mathbb{P}_h\widehat{V}^{\pi}_{h+1}(s_{k,h},a_{k,h})-\widehat{V}^{\pi}_{h+1}(s'_{k,h})-\epsilon_{k,h}\right)/\widehat{\sigma}_{k,h}$, which is a function of $s_{k,h},a_{k,h},\epsilon_{k,h},s'_{k,h}$. Now we define the filtration $\{\mathcal{F}_k\}_{k=0}^K$ by $\mathcal{F}_0=\sigma(s_{1,h},a_{1,h})$, $\mathcal{F}_1=\sigma(s_{1,h},a_{1,h},\epsilon_{1,h},s'_{1,h},s_{2,h},a_{2,h})$, $\cdots$, $\mathcal{F}_k=\sigma(s_{1,h},a_{1,h},\epsilon_{1,h},s'_{1,h},\cdots,s_{k,h},a_{k,h},\epsilon_{k,h},s'_{k,h},s_{k+1,h},a_{k+1,h})$ for $k=1,\cdots,K-1$, and $\mathcal{F}_K=\sigma(\mathcal{F}_{K-1},\epsilon_{K,h},s'_{K,h})$. Then we see that $\mathbf{x}_k$ is $\mathcal{F}_{k-1}$-measurable and ζ_k is $\mathcal{F}_k$-measurable. Furthermore, since $\mathbb{E}[\widehat{V}^{\pi}_{h+1}(s'_{k,h})\mid\mathcal{F}_{k-1}]=\mathbb{P}_h\widehat{V}^{\pi}_{h+1}(s_{k,h},a_{k,h})$, $\mathbb{E}[\epsilon_{k,h}\mid\mathcal{F}_{k-1}]=0$ and $\widehat{\sigma}_{k,h}$ is $\mathcal{F}_{k-1}$-measurable, $\zeta_k\mid\mathcal{F}_{k-1}$ has zero-mean. Also, by construction we have $|\zeta_k|\leq(2H-2h+3)/\sqrt{\eta_h+\sigma_r^2}$, and it follows from (F.15) that

$$\mathrm{Var}(\zeta_k\mid\mathcal{F}_{k-1})\leq\frac{\mathbb{V}_h\widehat{V}^{\pi}_{h+1}(s_{k,h},a_{k,h})+\sigma_r^2}{\max\{\eta_h,\mathbb{V}_h\widehat{V}^{\pi}_{h+1}(s_{k,h},a_{k,h})\}+\sigma_r^2-\frac{20(H-h+1)^2\sqrt{d}}{\kappa_h\sqrt{K}}\cdot\sqrt{\log\left(\frac{K}{\lambda\delta}\right)}}\leq 2,$$

as long as K satisfies (F.14).

Then by Theorem H.10, with probability at least $1-\delta$, we have

$$\left\|\sum_{k=1}^{K}\mathbf{x}_k\zeta_k\right\|_{\widehat{\boldsymbol{\Lambda}}_h^{-1}}\leq\sqrt{2d\log\left(1+\frac{K}{\lambda d(\eta_h+\sigma_r^2)}\right)\cdot\log\left(\frac{4K^2}{\delta}\right)}+\frac{4(2H-2h+3)}{\sqrt{\eta_h+\sigma_r^2}}\log\left(\frac{4K^2}{\delta}\right).$$

Since $\mathbb{P}(\check{\mathcal{E}})\geq 1-\delta$, the overall probability is at least $(1-\delta)^2\geq 1-2\delta$ by independence. Finally replacing δ by $\delta/2$ completes the proof. □

Lemma F.6. Let $\epsilon>0$ be a constant. For any $h\in[H]$ and $\delta\in(0,1)$, suppose K satisfies that

$$K\geq\frac{1600(H-h+1)^4d^2}{\kappa_h^2(\eta_h+\sigma_r^2)^2}\cdot\log\left(\frac{(H-h+1)^2KL}{\lambda\kappa_h(\eta_h+\sigma_r^2)\delta}\right) \tag{F.16}$$

where $\mathcal{N}_\epsilon$ is the ϵ-covering number of $\mathcal{V}_{h+1}(L)$. Then with probability at least $1-\delta$, it holds for all function $V\in\mathcal{V}_{h+1}(L)$ that

$$\left\|\sum_{k=1}^{K}\widehat{\sigma}_h(s_{k,h},a_{k,h})^{-2}\boldsymbol{\phi}(s_{k,h},a_{k,h})\left(\mathbb{P}_hV(s_{k,h},a_{k,h})-V(s'_{k,h})-\epsilon_{k,h}\right)\right\|^2_{\widehat{\boldsymbol{\Lambda}}_h^{-1}}$$
$$\leq 50\left(d+\frac{\sqrt{d}(H-h+1)}{\sqrt{\eta_h+\sigma_r^2}}\right)^2\cdot\log^2\left(\frac{K(H-h+1)^2L}{\kappa_h(\eta_h+\sigma_r^2)\delta}\right). \tag{F.17}$$

Proof of Lemma F.6. For the simplicity of presentation, we first define some notations. We define the following vector,

$$\mathbf{v}_V := \sum_{k=1}^{K} \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h}) \left(\mathbb{P}_h V(s_{k,h}, a_{k,h}) - V(s'_{k,h}) - \epsilon_{k,h}\right)}{\widehat{\sigma}_h(s_{k,h}, a_{k,h})^2}.$$

and the following matrix

$$\boldsymbol{\Gamma}_V := \sum_{k=1}^{K} \boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}(s_{k,h}, a_{k,h})^\top / \widehat{\sigma}_h^2(s_{k,h}, a_{k,h}) + \lambda \mathbf{I}_d.$$

It remains to show that, with probability at least $1-\delta$, for any function $V \in \mathcal{V}_{h+1}(L)$, $\mathbf{v}_V^\top \boldsymbol{\Gamma}_V^{-1} \mathbf{v}_V$ is no greater than the R.H.S. of (F.17). In the following argument, for $V_1 \in \mathcal{V}_{h+1}(L)$, we denote $\mathbf{v}_1 = \mathbf{v}_{V_1}$, $\boldsymbol{\Gamma}_1 = \boldsymbol{\Gamma}_{V_1}$ and σ_1 the variance estimator induced by V_1, and similar for V_2.

Let $\mathcal{C}_V$ be the smallest ϵ-cover of $\mathcal{V}_{h+1}(L)$, and $\mathcal{N}_\epsilon = |\mathcal{C}_V|$ the ϵ-covering number of $\mathcal{V}_{h+1}(L)$. For any $V_1 \in \mathcal{V}_{h+1}(L)$, there exists $V_2 \in \mathcal{C}_V$ such that $\text{dist}(V_1, V_2) = \sup_s |V_1(s) - V_2(s)| \le \epsilon$. Note that we have the following decomposition:

$$\mathbf{v}_1^\top \mathbf{A}_1^{-1} \mathbf{v}_1 \le \mathbf{v}_2^\top \mathbf{A}_2^{-1} \mathbf{v}_2 + \left|\mathbf{v}_1^\top \mathbf{A}_1^{-1} \mathbf{v}_1 - \mathbf{v}_2^\top \mathbf{A}_2^{-1} \mathbf{v}_2\right|. \tag{F.18}$$

By Lemma H.13, when K satisfies (F.16), we have

$$K \ge \frac{1600(H-h+1)^4 d}{\kappa_h^2 (\eta_h + \sigma_r^2)^2} \cdot \log\left(\frac{K\mathcal{N}_\epsilon}{\lambda\delta}\right)$$

Then by Lemma F.5 and a union bound, we have that, with probability at least $1-\delta$,

$$\mathbf{v}_2^\top \mathbf{A}_2^{-1} \mathbf{v}_2 \le \left(\sqrt{2d \log\left(1 + \frac{K}{\lambda d(\eta_h + \sigma_r^2)}\right) \cdot \log\left(\frac{4K^2\mathcal{N}_\epsilon}{\delta}\right)} + \frac{4(2H-2h+3)}{\sqrt{\eta_h + \sigma_r^2}} \log\left(\frac{4K^2\mathcal{N}_\epsilon}{\delta}\right)\right)^2. \tag{F.19}$$

It remains to bound the second term in (F.18). We first bound $\|\mathbf{v}_1 - \mathbf{v}_2\|_2$.

$$\begin{aligned}
&\|\mathbf{v}_1 - \mathbf{v}_2\| \\
&= \left\| \sum_{k=1}^{K} \boldsymbol{\phi}(s_{k,h}, a_{k,h}) \left(\frac{\mathbb{P}_h V_1(s_{k,h}, a_{k,h}) - V_1(s'_{k,h}) - \epsilon_{k,h}}{\sigma_1^2(s_{k,h}, a_{k,h})} - \frac{\mathbb{P}_h V_2(s_{k,h}, a_{k,h}) - V_2(s'_{k,h}) - \epsilon_{k,h}}{\sigma_2^2(s_{k,h}, a_{k,h})} \right) \right\| \\
&\le \sum_{k=1}^{K} \|\boldsymbol{\phi}(s_{k,h}, a_{k,h})\|_2 \cdot \left| \frac{\mathbb{P}_h V_1(s_{k,h}, a_{k,h}) - V_1(s'_{k,h}) - \epsilon_{k,h}}{\sigma_1^2(s_{k,h}, a_{k,h})} - \frac{\mathbb{P}_h V_2(s_{k,h}, a_{k,h}) - V_2(s'_{k,h}) - \epsilon_{k,h}}{\sigma_2^2(s_{k,h}, a_{k,h})} \right| \\
&\le \sum_{k=1}^{K} \left| \frac{\left(\mathbb{P}_h V_1(s_{k,h}, a_{k,h}) - V_1(s'_{k,h}) - \epsilon_{k,h}\right)}{\sigma_1^2(s_{k,h}, a_{k,h})} - \frac{\left(\mathbb{P}_h V_2(s_{k,h}, a_{k,h}) - V_2(s'_{k,h}) - \epsilon_{k,h}\right)}{\sigma_2^2(s_{k,h}, a_{k,h})} \right|
\end{aligned} \tag{F.20}$$

where the first inequality follows from Cauchy-Schwartz inequality and the second inequality is due to Assumption 2.1.

Note that for any real-valued function $f_1(\cdot), f_2(\cdot)$ and positive function $g_1(\cdot), g_2(\cdot)$ bounded away from 0, we have

$$\begin{aligned}
\left|\frac{f_1}{g_1} - \frac{f_2}{g_2}\right| &= \left|\frac{f_1 g_2 - f_1 g_1 + f_1 g_1 - g_1 f_2}{g_1 g_2}\right| \\
&\le \left|\frac{f_1(g_2 - g_1)}{g_1 g_2}\right| + \left|\frac{g_1(f_1 - f_2)}{g_1 g_2}\right| \\
&\le \frac{1}{\inf g_1 \inf g_2} \left[(\sup |f_1|) \cdot |g_2 - g_1| + (\sup g_1) \cdot |f_1 - f_2|\right].
\end{aligned} \tag{F.21}$$

Now, by the construction we have $\sigma_1^2(\cdot,\cdot) \in [\eta_h + \sigma_r^2, (H-h+1)^2 + \sigma_r^2]$, and $\mathbb{P}_h V_1(\cdot,\cdot) - V_1(\cdot) - \epsilon_{k,h} \in [-2H+2h-3, 2H-2h+3]$, and the same for σ_2 and V_2. Also note that for all (s,a),

$$
\begin{aligned}
&\left|\sigma_1^2(s,a) - \sigma_2^2(s,a)\right| \\
&\le \left|\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\beta}}_{h,1} - \widehat{\boldsymbol{\beta}}_{h,2}\rangle\right| + 2(H-h+1)\left|\langle \boldsymbol{\phi}(s,a), \widehat{\boldsymbol{\theta}}_{h,1} - \widehat{\boldsymbol{\theta}}_{h,2}\rangle\right| \\
&\le \left|\widehat{\boldsymbol{\Sigma}}_h^{-1}\sum_{k=1}^K \boldsymbol{\phi}(\check{s}_{k,h},\check{a}_{k,h})(V_1^2(\check{s}'_{k,h}) - V_2^2(\check{s}'_{k,h}))\right| + \left|\widehat{\boldsymbol{\Sigma}}_h^{-1}\sum_{k=1}^K \boldsymbol{\phi}(\check{s}_{k,h},\check{a}_{k,h})(V_1(\check{s}'_{k,h}) - V_2(\check{s}'_{k,h}))\right| \\
&\le \frac{4(H-h+2)^2K}{\kappa_h}\cdot\epsilon,
\end{aligned}
\tag{F.22}
$$

where the second inequality is due to Assumption 2.1 and the third inequality follows from the fact that $V_1, V_2 \in \mathcal{V}_{h+1}(L)$.

Denote $\mathbf{u} = \mathbf{v}_2 - \mathbf{v}_1$. Combining (F.20), (F.21) and (F.22) yields that

$$
\begin{aligned}
\|\mathbf{u}\| &\le \sum_{k=1}^K \frac{1}{(\eta_h+\sigma_r^2)^2}\left[\frac{4(2H-2h+3)(H-h+2)^2K}{\kappa_h}\cdot\epsilon + ((H-h+1)^2+\sigma_r^2)\cdot 2\epsilon\right] \\
&\le \frac{10(H-h+1)^2K^2}{\kappa_h(\eta_h+\sigma_r^2)^2}\epsilon,
\end{aligned}
\tag{F.23}
$$

where the second inequality is due to the fact that $\sigma_r^2 \le 1$.

Next, by Lemma H.7 and (F.22), we have

$$
\|\mathbf{A}_1^{-1} - \mathbf{A}_2^{-1}\| \le \frac{8K^2((H-h+1)^2+\sigma_r^2)}{\lambda^2\kappa_h(\eta_h+\sigma_r^2)^2}\cdot\epsilon. \tag{F.24}
$$

Also note that

$$
\begin{aligned}
\left|\mathbf{v}_1^\top\mathbf{A}_1^{-1}\mathbf{v}_1 - \mathbf{v}_2^\top\mathbf{A}_2^{-1}\mathbf{v}_2\right| &= \left|\mathbf{v}_1^\top\mathbf{A}_1^{-1}\mathbf{v}_1 - (\mathbf{v}_1+\mathbf{u})^\top\mathbf{A}_2^{-1}(\mathbf{v}_1+\mathbf{u})\right| \\
&\le \left|\mathbf{v}_1^\top(\mathbf{A}_1^{-1}-\mathbf{A}_2^{-1})\mathbf{v}_1\right| + 2\left|\mathbf{v}_1^\top\mathbf{A}_2^{-1}\mathbf{u}\right| + \left|\mathbf{u}^\top\mathbf{A}_2^{-1}\mathbf{u}\right|.
\end{aligned}
\tag{F.25}
$$

By the definition, we have $\|\mathbf{v}_1\|_2, \|\mathbf{v}_2\|_2 \le (2H-2h+3)K/(\eta_h+\sigma_r^2)$, and $\|\mathbf{A}_1^{-1}\|, \|\mathbf{A}_2^{-1}\| \le 1/\lambda$. It then follows from (F.24) and (F.25) that

$$
\begin{aligned}
\left|\mathbf{v}_1^\top\mathbf{A}_1^{-1}\mathbf{v}_1 - \mathbf{v}_2^\top\mathbf{A}_2^{-1}\mathbf{v}_2\right| &\le \frac{(2H-2h+3)^2K^2}{(\eta_h+\sigma_r^2)^2}\cdot\frac{8K^2((H-h+1)^2+\sigma_r^2)}{\lambda^2\kappa_h(\eta_h+\sigma_r^2)^2}\cdot\epsilon \\
&\quad + \frac{2(2H-2h+3)K}{\lambda(\eta_h+\sigma_r^2)}\cdot\frac{10K(H-h+1)^2K^2}{\kappa_h(\eta_h+\sigma_r^2)^2}\cdot\epsilon \\
&\quad + \frac{100(H-h+1)^4K^4}{\lambda\kappa_h^2(\eta_h+\sigma_r^2)^4}\cdot\epsilon^2 \\
&\le \frac{200(H-h+1)^4K^4}{\kappa_h^2(\eta_h+\sigma_r^2)^2}\cdot\epsilon,
\end{aligned}
$$

by the choice of $\lambda = 1$. We then choose $\epsilon = \kappa_h^2(\eta_h+\sigma_r^2)^2/[200(H-h+1)^4K^5]$, and thus

$$
\left|\mathbf{v}_1^\top\mathbf{A}_1^{-1}\mathbf{v}_1 - \mathbf{v}_2^\top\mathbf{A}_2^{-1}\mathbf{v}_2\right| \le \frac{1}{K} \tag{F.26}
$$

Now by Lemma (H.13), we have

$$
\mathcal{N}_\epsilon \le \left(1 + \frac{400(H-h+1)^4K^5L}{\kappa_h^2(\eta_h+\sigma_r^2)^2}\right)^d. \tag{F.27}
$$

Then combining (F.18), (F.19), (F.26) and (F.27) yields

$$
\begin{aligned}
&\left\|\sum_{k=1}^K \widehat{\sigma}_h(s_{k,h},a_{k,h})^{-2}\boldsymbol{\phi}(s_{k,h},a_{k,h})\left(\mathbb{P}_hV(s_{k,h},a_{k,h}) - V(s'_{k,h}) - \epsilon_{k,h}\right)\right\|_{\widehat{\boldsymbol{\Lambda}}_h^{-1}}^2 \\
&\le 50\left(d + \frac{d(H-h+1)}{\sqrt{\eta_h+\sigma_r^2}}\right)^2\cdot\log^2\left(\frac{K(H-h+1)^2L}{\kappa_h(\eta_h+\sigma_r^2)\delta}\right).
\end{aligned}
$$

This completes the proof. $\square$

F.3 Bounding the error terms

Finally, we prove the following key lemma for completing the induction step in the proof of Theorem C.2.

Lemma F.7. Set $L = (1+1/H)d\sqrt{K/\lambda}$. For any $h \in [H-1]$, let ρ be some positive constant such that $\rho < (\eta_h + \sigma_r^2)/[12(H-h+1)]$. For any $\delta \in (0,1)$, suppose K satisfies that

$$K \geq \frac{3600(H-h+1)^4 d^2}{\kappa_h^2(\eta_h+\sigma_r^2)^2} \cdot \log\left(\frac{dHK}{\kappa_h \delta}\right) \tag{F.28}$$

Then the following two events hold simultaneously with probability at least $1-\delta$:

1. $\widetilde{\mathcal{E}}_1$: for all $V \in \mathcal{V}_{h+1}(L) \cap \{V : \sup_{s\in\mathcal{S}} |V(s) - V_{h+1}^{\pi}(s)| \leq \rho\}$,
$$\left\| \left(\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}\right)^{-1} \right\| \leq \frac{8}{\iota_h}; \tag{F.29}$$

2. $\widetilde{\mathcal{E}}_2$: for all function $V(\cdot) \in \mathcal{V}_{h+1}(L) \cap \{V : \sup_{s\in\mathcal{S}} |V(s) - V_{h+1}^{\pi}(s)| \leq \rho\}$ and all (s,a) pairs,
$$\left| \boldsymbol{\phi}(s,a)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \widehat{\sigma}_h(s_{k,h},a_{k,h})^{-2} \boldsymbol{\phi}(s_{k,h},a_{k,h}) \left(\mathbb{P}_h V(s_{k,h},a_{k,h}) - V(s'_{k,h}) - \epsilon_{k,h}\right) \right|$$
$$\leq \frac{20}{\sqrt{K}} \cdot \left(\frac{d}{\sqrt{\iota_h}} + \frac{d(H-h+1)}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right) \cdot \log\left(\frac{d(H-h+1)^2 K}{\kappa_h(\eta_h+\sigma_r^2)\delta}\right)$$

Proof of Lemma F.7. We want to show $\mathbb{P}\{\widetilde{\mathcal{E}}_1 \cap \widetilde{\mathcal{E}}_2\} \geq 1-\delta$. It follows from Lemma F.4 and (F.28) that $\mathbb{P}(\widetilde{\mathcal{E}}_1) \geq 1-\delta$.

To show that $\mathbb{P}(\widetilde{\mathcal{E}}_2) \geq 1-\delta$, first by Lemma F.6, we have

$$\begin{aligned}
&\left\| \sum_{k=1}^K \widehat{\sigma}_h(s_{k,h},a_{k,h})^{-2} \boldsymbol{\phi}(s_{k,h},a_{k,h}) \left(\mathbb{P}_h V(s_{k,h},a_{k,h}) - V(s'_{k,h}) - \epsilon_{k,h}\right) \right\|_{\widehat{\boldsymbol{\Lambda}}_h^{-1}}^2 \\
\leq & 50 \left(d + \frac{d(H-h+1)}{\sqrt{\eta_h+\sigma_r^2}}\right)^2 \cdot \log^2\left(\frac{K(H-h+1)^2 L}{\kappa_h(\eta_h+\sigma_r^2)\delta}\right),
\end{aligned} \tag{F.30}$$

for all $V \in \mathcal{V}_{h+1}(L)$.

It follows from Cauchy-Schwartz inequality that

$$\begin{aligned}
&\boldsymbol{\phi}(s,a)^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \widehat{\sigma}_h(s_{k,h},a_{k,h})^{-2} \boldsymbol{\phi}(s_{k,h},a_{k,h}) \left(\mathbb{P}_h V(s_{k,h},a_{k,h}) - V(s'_{k,h}) - \epsilon_{k,h}\right) \\
&\leq \|\boldsymbol{\phi}(s,a)\|_{\widehat{\boldsymbol{\Lambda}}_h^{-1}} \cdot \left\| \sum_{k=1}^K \widehat{\sigma}_h(s_{k,h},a_{k,h})^{-2} \boldsymbol{\phi}(s_{k,h},a_{k,h}) \left([\mathbb{P}_h V](s_{k,h},a_{k,h}) - V(s'_{k,h}) - \epsilon_{k,h}\right) \right\|_{\widehat{\boldsymbol{\Lambda}}_h^{-1}} \\
&\leq \|\widehat{\boldsymbol{\Lambda}}_h^{-1}\|^{1/2} \cdot \left\| \sum_{k=1}^K \widehat{\sigma}_h(s_{k,h},a_{k,h})^{-2} \boldsymbol{\phi}(s_{k,h},a_{k,h}) \left([\mathbb{P}_h V](s_{k,h},a_{k,h}) - V(s'_{k,h}) - \epsilon_{k,h}\right) \right\|_{\widehat{\boldsymbol{\Lambda}}_h^{-1}} \\
&\leq \frac{20}{\sqrt{K}} \cdot \left(\frac{d}{\sqrt{\iota_h}} + \frac{d(H-h+1)}{\sqrt{\iota_h(\eta_h+\sigma_r^2)}}\right) \cdot \log\left(\frac{d(H-h+1)^2 K}{\kappa_h(\eta_h+\sigma_r^2)\delta}\right),
\end{aligned}$$

where the second inequality follows from Assumption 2.1 and the third inequality follows from (F.29) and (F.30). Note that this holds for all $(s,a) \in \mathcal{S} \times \mathcal{A}$ as we directly bound the operator norm of $\widehat{\boldsymbol{\Lambda}}_h$. Replacing δ by $\delta/2$ completes the proof. $\square$

G Lemmas for OPE Convergence

G.1 Concentration of $\widehat{\sigma}$

Recall that in the algorithm, to estimate the variance, we use $\widehat{\boldsymbol{\beta}}_h$ and $\widehat{\boldsymbol{\theta}}_h$ which are estimated using the function $\widehat{V}_{h+1}^{\pi}$ and $\{\check{s}_{k,h}, \check{a}_{k,h}, \check{s}'_{k,h}\}_{k\in[K]}$.

For the next lemma we denote the function $\sigma(\cdot,\cdot)$ as computed from some function $V(\cdot)$ and data $\check{\mathcal{D}}_h$.

Lemma G.1. Let $\rho \geq 0$. For any $V \in \mathcal{V}_{h+1}(L) \cap \{V : \sup_s |V(s) - V_{h+1}^{\pi}(s)| \leq \rho\}$, with probability at least $1-\delta$, we have

$$\left|\sigma^2(s,a) - \sigma_r^2 - \max\left\{\eta_h,\ \mathbb{V}_h V_{h+1}^{\pi}(s,a)\right\}\right| \leq \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}} + 4(H-h+1)\cdot\rho,$$

for all $(s,a) \in \mathcal{S}\times\mathcal{A}$ where

$$C_{K,h,\delta} = 12\sqrt{2}\cdot\frac{1}{\sqrt{\kappa_h}}\cdot\left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d}\log\frac{4}{\delta}\right]^{1/2} + 12\lambda\cdot\frac{1}{\kappa_h}.$$

Proof of Lemma G.1. By Lemma F.2, with probability at least $1-\delta$, we have

$$\left|\sigma^2(s,a) - \sigma_r^2 - \max\left\{\eta_h,\ \mathbb{V}_h V(s,a)\right\}\right| \leq \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}.$$

Note that if two functions $f_1, f_2 : \mathcal{S}\to\mathbb{R}$ satisfies $\sup_s |f_1(s) - f_2(s)| \leq \rho$, $\sup_s |f_1(s)| \leq H-h+1$, and $\sup_s |f_2(s)| \leq H-h+1$, then for all (s,a),

$$|\mathbb{V}_h f_1(s,a) - \mathbb{V}_h f_2(s,a)| \leq 4(H-h+1)\cdot\rho.$$

Then using the triangular inequality and $|\mathbb{V}_h V(s,a) - \mathbb{V}_h V_{h+1}^{\pi}(s,a)| \leq 4(H-h+1)\rho$ finishes the proof. □

G.2 Concentration of Weighted Sample Covariance Matrices

In this subsection, we study the concentration of the matrices $\widehat{\boldsymbol{\Lambda}}_h$, $h\in[H]$ to their population counterparts. Recall from Algorithm 1 that for each $h\in[H]$, the matrix $\widehat{\boldsymbol{\Lambda}}_h$ is generated using the function $\widehat{\sigma}_h(\cdot,\cdot)$ and the dataset $\mathcal{D}_h = \{(s_{k,h}, a_{k,h}, r_{k,h}, s'_{k,h})\}_{k\in[K]}$. Since the function $\widehat{\sigma}_h(\cdot,\cdot)$ itself is generated by $\widehat{V}_{h+1}(\cdot)$ and the dataset $\check{\mathcal{D}}_h = \{(\check{s}_{k,h}, \check{a}_{k,h}, \check{r}_{k,h}, \check{s}'_{k,h})\}_{k\in[K]}$, we can equivalently view $\widehat{\boldsymbol{\Lambda}}_h$ as generated by $\widehat{V}_{h+1}(\cdot)$ and the datasets $\check{\mathcal{D}}_h$ and $\mathcal{D}_h$. In the remaining of the subsection, we will omit the subscript and superscript when it is clear and simply write

$$\widehat{\boldsymbol{\Lambda}}_h = \sum_{k=1}^{K}\boldsymbol{\phi}(s_{k,h},a_{k,h})\boldsymbol{\phi}^{\top}(s_{k,h},a_{k,h})/\sigma^2(s_{k,h},a_{k,h}) + \lambda\mathbf{I}_d,$$

where $\sigma(\cdot,\cdot)$ is generated using the function $V(\cdot)$ and the dataset $\check{\mathcal{D}}_h$ as described in Algorithm 1. We also denote

$$\sigma_V^2(\cdot,\cdot) := \max\left\{\eta_h,\ \mathbb{V}_h V(\cdot,\cdot)\right\} + \sigma_r^2.$$

By Lemma F.2, we know that with high probability, $\sigma^2(\cdot,\cdot)$ will be a good estimator for $\sigma_V^2(\cdot,\cdot)$. This will be used to show the concentration of the matrix $\widehat{\boldsymbol{\Lambda}}_h$. We start from the next lemma.

Lemma G.2. For any $h\in[H]$, conditioning on $\sigma(\cdot,\cdot)\in\mathcal{T}_h(L_1,L_2)$ being fixed, with conditional probability at least $1-\delta$,

$$\left\|\frac{\widehat{\boldsymbol{\Lambda}}_h}{K} - \mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^{\top}}{\sigma^2(s,a)}\right]\right\| \leq \frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}}\cdot\left(\log\frac{2d}{\delta}\right)^{1/2} + \frac{\lambda}{K}.$$

Proof of Lemma G.2. Since $\sigma(\cdot,\cdot)$ is a function of V and the dataset $\check{\mathcal{D}}_h$ which is independent of $\mathcal{D}_h$, conditioning on $\sigma(\cdot,\cdot)$ won't change the distribution of $\mathcal{D}_h$. In other words, $\boldsymbol{\phi}(s_{k,h},a_{k,h})/\sigma(s_{k,h},a_{k,h})$, $k\in[K]$ can be viewed as independent random vectors. Then by Lemma H.4, we have that, with conditional probability at least $1-\delta$,

$$\left\|\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}-\mathbb{E}_{\bar{\pi},h}\left[\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}\right]\right\|\leq\frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}}\cdot\sqrt{\log\left(\frac{2d}{\delta}\right)},$$

and thus

$$\begin{aligned}\left\|\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}-\mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma^2(s,a)}\right]\right\|&\leq\left\|\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}-\mathbb{E}_{\bar{\pi},h}\left[\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}\right]\right\|+\left\|\mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma^2(s,a)}\right]-\mathbb{E}_{\bar{\pi},h}\left[\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}\right]\right\|\\&\leq\frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}}\cdot\sqrt{\log\left(\frac{2d}{\delta}\right)}+\frac{\lambda}{K}.\end{aligned}$$

□

Next, combine Lemma G.2 and the event that $\sigma^2(\cdot,\cdot)$ is a good estimator for $\sigma_V^2(\cdot,\cdot)$, we get the following lemma.

Lemma G.3. For any $h\in[H]$, condition on $V\in\mathcal{V}_{h+1}(L)$ being fixed, with conditional probability at least $1-\delta$,

$$\begin{aligned}&\left\|\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}-\mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma_V^2(s,a)}\right]\right\|\\&\leq\frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}}\cdot\left(\log\frac{4d}{\delta}\right)^{1/2}+\frac{\lambda}{K}+\frac{1}{(\eta_h+\sigma_r^2)^2}\cdot\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}},\end{aligned}$$

where

$$C_{K,h,\delta}=12\sqrt{2}\cdot\frac{1}{\sqrt{\kappa_h}}\cdot\left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right)+\frac{1}{d}\log\frac{8}{\delta}\right]^{1/2}+12\lambda\cdot\frac{1}{\kappa_h}.$$

Proof of Lemma G.3. First note that condition on $\sigma(\cdot,\cdot)\in\mathcal{T}_h(L_1,L_2)$ such that $\sup_{s,a}|\sigma^2(s,a)-\sigma_V^2(s,a)|\leq\rho$ for some $\rho\geq 0$, we have

$$\begin{aligned}&\left\|\mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma^2(s,a)}\right]-\mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma_V^2(s,a)}\right]\right\|\\&\leq\mathbb{E}_{\bar{\pi},h}\left[\left\|\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top\right\|\sup_{s,a}\left(\frac{1}{\sigma^2(s,a)}-\frac{1}{\sigma_V^2(s,a)}\right)\right]\\&\leq\frac{1}{(\eta_h+\sigma_r^2)^2}\cdot\rho,\end{aligned}$$

since $\sigma^2(s,a)$ and $\sigma_V^2(s,a)$ are lower bounded by $\eta_h+\sigma_r^2$. Then by Lemma G.2, we have that, conditioning on fixed $\sigma(\cdot,\cdot)$ s.t. $\sup_{s,a}|\sigma^2(s,a)-\sigma_V^2(s,a)|\leq\rho$, with conditional probability at least $1-\delta$,

$$\left\|\frac{\widehat{\boldsymbol{\Lambda}}_h}{K}-\mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma_V^2(s,a)}\right]\right\|\leq\frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}}\cdot\left(\log\frac{2d}{\delta}\right)^{1/2}+\frac{\lambda}{K}+\frac{1}{(\eta_h+\sigma_r^2)^2}\cdot\rho. \tag{G.1}$$

Since conditioning on $V(\cdot)$ won't change the distribution of $\check{D}_h$ under Assumption C.1, by Lemma F.2, with probability at least $1-\delta$, it holds for all $(s,a)\in\mathcal{S}\times\mathcal{A}$ that

$$\left|\sigma^2(s,a)-\sigma_V^2(s,a)\right|\leq\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}, \tag{G.2}$$

where

$$C_{K,h,\delta} = 12\sqrt{2} \cdot \frac{1}{\sqrt{\kappa_h}} \cdot \left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d}\log\frac{4}{\delta}\right]^{1/2} + 12\lambda \cdot \frac{1}{\kappa_h}.$$

Combine (G.1) and (G.2), and we get that, condition on V, with probability at least $1-2\delta$,

$$\begin{aligned}
&\left\| \frac{\widehat{\boldsymbol{\Lambda}}_h}{K} - \mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma_V^2(s,a)}\right] \right\| \\
&\leq \frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}} \cdot \left(\log\frac{2d}{\delta}\right)^{1/2} + \frac{\lambda}{K} + \frac{1}{(\eta_h+\sigma_r^2)^2} \cdot \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}.
\end{aligned}$$

Replacing δ with $\delta/2$ finishes the proof. □

Finally, combining Lemma G.3 and the event of uniform convergence, we can bound the distance between $\widehat{\boldsymbol{\Lambda}}_h$ and its population counterpart $\boldsymbol{\Lambda}_h$.

Lemma G.4. For any $h \in [H]$, condition on $V \in \mathcal{V}_{h+1}(L) \cap \{V : \sup_s |V(s) - V_{h+1}^\pi(s)| \leq \rho\}$, with conditional probability at least $1-\delta$, we have

$$\begin{aligned}
&\left\| \frac{\widehat{\boldsymbol{\Lambda}}_h}{K} - \boldsymbol{\Lambda}_h \right\| \\
&\leq \frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}} \cdot \left(\log\frac{4d}{\delta}\right)^{1/2} + \frac{\lambda}{K} + \frac{1}{(\eta_h+\sigma_r^2)^2} \cdot \left(\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}} + 4(H-h+1)\cdot\rho\right),
\end{aligned}$$

where

$$C_{K,h,\delta} = 12\sqrt{2} \cdot \frac{1}{\sqrt{\kappa_h}} \cdot \left[\frac{1}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \frac{1}{d}\log\frac{8}{\delta}\right]^{1/2} + 12\lambda \cdot \frac{1}{\kappa_h}.$$

Proof of Lemma G.4. First note that by Lemma G.3, with probability at least $1-\delta$,

$$\begin{aligned}
&\left\| \frac{\widehat{\boldsymbol{\Lambda}}_h}{K} - \mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma_V^2(s,a)}\right] \right\| \\
&\leq \frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}} \cdot \left(\log\frac{4d}{\delta}\right)^{1/2} + \frac{\lambda}{K} + \frac{1}{(\eta_h+\sigma_r^2)^2} \cdot \frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}}.
\end{aligned}$$

On the other hand, by $\sup_s |V(s) - V_{h+1}^\pi(s)| \leq \rho$ and $|V(s)|, |V_{h+1}^\pi(s)| \leq H-h+1$, we have $\sup_s |\sigma_V^2(s) - \sigma_h^2(s)| \leq 4(H-h+1)\rho$. It implies that

$$\left\| \mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma_V^2(s,a)}\right] - \mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma_h^2(s,a)}\right] \right\| \leq \frac{1}{(\eta_h+\sigma_r^2)^2} \cdot 4(H-h+1)\cdot\rho.$$

Then by triangular inequality, we conclude that

$$\begin{aligned}
&\left\| \frac{\widehat{\boldsymbol{\Lambda}}_h}{K} - \mathbb{E}_{\bar{\pi},h}\left[\frac{\boldsymbol{\phi}(s,a)\boldsymbol{\phi}(s,a)^\top}{\sigma_h^2(s,a)}\right] \right\| \\
&\leq \frac{4\sqrt{2}}{(\eta_h+\sigma_r^2)\sqrt{K}} \cdot \left(\log\frac{4d}{\delta}\right)^{1/2} + \frac{\lambda}{K} + \frac{1}{(\eta_h+\sigma_r^2)^2} \cdot \left(\frac{C_{K,h,\delta}(H-h+1)^2\sqrt{d}}{\sqrt{K}} + 4(H-h+1)\cdot\rho\right).
\end{aligned}$$

Finally, recall the definition of $\boldsymbol{\Lambda}_h$ given by (2.6). □

G.3 Bound for the self-normalized martingales

Lemma G.5. For any $h \in [H]$, condition on $\widehat{V}^{\pi}_{h+1} \in \mathcal{V}_{h+1}(L)$ s.t. $\sup_s \left|\widehat{V}^{\pi}_{h+1}(s)\right| \leq B$, with conditional probability at least $1-\delta$,

$$\left\|\sum_{k=1}^{K} \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h}) \left[(\widehat{V}^{\pi}_{h+1})(\check{s}'_{k,h}) - \mathbb{P}_h(\widehat{V}^{\pi}_{h+1})(\check{s}_{k,h}, \check{a}_{k,h})\right]\right\|^2_{\widehat{\boldsymbol{\Sigma}}_h^{-1}} \leq 8B^2 \left[\frac{d}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \log\frac{2}{\delta}\right],$$

$$\left\|\sum_{k=1}^{K} \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h}) \left[(\widehat{V}^{\pi}_{h+1})^2(\check{s}'_{k,h}) - \mathbb{P}_h(\widehat{V}^{\pi}_{h+1})^2(\check{s}_{k,h}, \check{a}_{k,h})\right]\right\|^2_{\widehat{\boldsymbol{\Sigma}}_h^{-1}} \leq 8B^4 \left[\frac{d}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \log\frac{2}{\delta}\right].$$

Proof of Lemma G.5. Denote $\mathbf{x}_k = \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})$, and $\eta_k = (\widehat{V}^{\pi}_{h+1})(\check{s}'_{k,h}) - \mathbb{P}_h(\widehat{V}^{\pi}_{h+1})(\check{s}_{k,h}, \check{a}_{k,h})$.

Define the filtration $\{\mathcal{F}_k\}_{k=0}^K$ by $\mathcal{F}_0 = \sigma(\check{s}_{1,h}, \check{a}_{1,h})$, $\mathcal{F}_1 = \sigma(\check{s}_{1,h}, \check{a}_{1,h}, \check{s}'_{1,h}, \check{s}_{2,h}, \check{a}_{2,h})$, $\cdots$, $\mathcal{F}_k = \sigma(\check{s}_{1,h}, \check{a}_{1,h}, \check{s}'_{1,h}, \cdots, \check{s}_{k,h}, \check{a}_{k,h}, \check{s}'_{k,h}, \check{s}_{k+1,h}, \check{a}_{k+1,h})$ for $k = 1, \cdots, K-1$, and $\mathcal{F}_K = \sigma(\mathcal{F}_{K-1}, \check{s}'_{K,h})$. Then we see that $\mathbf{x}_k$ is $\mathcal{F}_{k-1}$-measurable, and η_k is $\mathcal{F}_k$-measurable. Furthermore, since $\mathbb{E}[(\widehat{V}^{\pi}_{h+1})^2(\check{s}'_{k,h}) \mid \mathcal{F}_{k-1}] = \mathbb{P}_h(\widehat{V}^{\pi}_{h+1})^2(\check{s}_{k,h}, \check{a}_{k,h})$, $\eta_k \mid \mathcal{F}_{k-1}$ is zero-mean. Also, $|\eta_k| \leq 2B$, which implies that $\eta_k \mid \mathcal{F}_{k-1}$ is $2B$-subgaussian. Then by H.9, with probability at least $1-\delta/2$,

$$\left\|\sum_{k=1}^{K} \mathbf{x}_k \eta_k\right\|^2_{\widehat{\boldsymbol{\Sigma}}_h^{-1}} \leq 8B^2 \log\left(\frac{\det(\widehat{\boldsymbol{\Sigma}}_h)^{1/2}\det(\lambda I)^{-1/2}}{\delta/2}\right).$$

Recall that $\widehat{\boldsymbol{\Sigma}}_h = \sum_{k=1}^K \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})\boldsymbol{\phi}^\top(\check{s}_{k,h}, \check{a}_{k,h}) + \lambda \mathbf{I}_d$ where $\|\boldsymbol{\phi}\| \leq 1$. It follows that

$$\det(\widehat{\boldsymbol{\Sigma}}_h) \leq (\lambda + K)^d.$$

We then conclude that

$$\left\|\sum_{k=1}^{K} \mathbf{x}_k \eta_k\right\|^2_{\widehat{\boldsymbol{\Sigma}}_h^{-1}} \leq 8B^2 \left[\frac{d}{2}\log\left(\frac{\lambda+K}{\lambda}\right) + \log\frac{2}{\delta}\right].$$

The second inequality is similar. Taking a union bound finishes the proof. □

H Auxiliary Lemmas

H.1 Concentration Inequalities

Lemma H.1 (Matrix McDiarmid inequality, Tropp 40)**.** Let $\mathbf{z}_k$, $k = 1, \cdots, K$ be independent random vectors in $\mathbb{R}^d$, and let $\mathbf{H}$ be a function that maps K vectors to a $d \times d$ symmetric matrix. Assume there exists a sequence of fixed symmetric matrices $\{\mathbf{A}_k\}_{k\in[K]}$ such that

$$\left(\mathbf{H}(\mathbf{z}_1, \cdots, \mathbf{z}_k, \cdots, \mathbf{z}_K) - \mathbf{H}(\mathbf{z}_1, \cdots, \mathbf{z}'_k, \cdots, \mathbf{z}_K)\right)^2 \preceq \mathbf{A}_k^2,$$

where $\mathbf{z}_k, \mathbf{z}'_k$ ranges over all possible values for each $k \in [K]$. Define σ^2 as

$$\sigma^2 := \left\|\sum_k \mathbf{A}_k^2\right\|.$$

Then, for any $t > 0$,

$$\mathbb{P}\left\{\lambda_{\max}\left(\mathbf{H}(\mathbf{z}) - \mathbb{E}\mathbf{H}(\mathbf{z})\right) \geq t\right\} \leq d \cdot \exp\left(\frac{-t^2}{8\sigma^2}\right),$$

where $\mathbf{z} = (\mathbf{z}_1, \cdots, \mathbf{z}_K)$.

Lemma H.2 (Freedman's inequality for martingales, Freedman 13)**.** Consider a martingale difference sequence $\{e_k,\ k=1,2,3,\cdots\}$ with filtration $\mathcal{F}_k := \sigma\{e_1,\cdots,e_{k-1}\}$, for $k=1,2,\cdots$. Assume e_k is uniformly bounded:

$$|e_k| \le R \quad \text{almost surely for} \quad k=1,2,3,\cdots$$

Then for all $\epsilon \ge 0$ and $\sigma^2 > 0$,

$$\mathbb{P}\left\{\exists K>0 : \left|\sum_{k=1}^{K} e_k\right| \ge \epsilon,\ \sum_{k=1}^{K} \mathrm{Var}[e_k \mid \mathcal{F}_k] \le \sigma^2\right\} \le 2\exp\left(\frac{-\epsilon^2/2}{\sigma^2 + R\epsilon/3}\right).$$

H.2 Basic Matrix Inequalities

Lemma H.3. Assume $\mathbf{G}_1$ and $\mathbf{G}_2 \in \mathbb{R}^{d\times d}$ are two positive semi-definite matrices. Then we have

$$\left\|\mathbf{G}_1^{-1}\right\| \le \left\|\mathbf{G}_2^{-1}\right\| + \left\|\mathbf{G}_1^{-1}\right\| \cdot \left\|\mathbf{G}_2^{-1}\right\| \cdot \|\mathbf{G}_1 - \mathbf{G}_2\|$$

and

$$\|\mathbf{u}\|_{\mathbf{G}_1^{-1}} \le \left[1 + \sqrt{\left(\left\|\mathbf{G}_2^{-1}\right\| \cdot \|\mathbf{G}_2\|\right)^{1/2} \cdot \left\|\mathbf{G}_1^{-1}\right\| \cdot \|\mathbf{G}_1 - \mathbf{G}_2\|}\right] \cdot \|\mathbf{u}\|_{\mathbf{G}_2^{-1}},$$

for all $\mathbf{u} \in \mathbb{R}^d$.

Proof of Lemma H.3. The first inequality is by

$$\left\|\mathbf{G}_1^{-1}\right\| \le \left\|\mathbf{G}_2^{-1}\right\| + \left\|\mathbf{G}_2^{-1} - \mathbf{G}_1^{-1}\right\| \le \left\|\mathbf{G}_2^{-1}\right\| + \left\|\mathbf{G}_2^{-1}\right\| \cdot \|\mathbf{G}_2 - \mathbf{G}_1\| \cdot \left\|\mathbf{G}_1^{-1}\right\|.$$

To prove the second one, note that

$$\begin{aligned}
\|\mathbf{u}\|_{\mathbf{G}_1^{-1}} &= \sqrt{\mathbf{u}^\top \mathbf{G}_1^{-1}\mathbf{u}} \\
&= \sqrt{\mathbf{u}^\top(\mathbf{G}_1^{-1} - \mathbf{G}_2^{-1})\mathbf{u} + \mathbf{u}^\top \mathbf{G}_2^{-1}\mathbf{u}} \\
&= \sqrt{\mathbf{u}^\top \mathbf{G}_2^{-1/2}\left[\mathbf{I} + (\mathbf{G}_2^{1/2}\mathbf{G}_1^{-1}\mathbf{G}_2^{1/2} - \mathbf{I})\right]\mathbf{G}_2^{-1/2}\mathbf{u}} \\
&\le \left(1 + \left\|\mathbf{G}_2^{1/2}\mathbf{G}_1^{-1}\mathbf{G}_2^{1/2} - \mathbf{I}\right\|^{1/2}\right) \cdot \|\mathbf{u}\|_{\mathbf{G}_2^{-1}},
\end{aligned}$$

and the rest follows from

$$\begin{aligned}
\left\|\mathbf{G}_2^{1/2}\mathbf{G}_1^{-1}\mathbf{G}_2^{1/2} - \mathbf{I}\right\| &= \left\|\mathbf{G}_2^{1/2}(\mathbf{G}_1^{-1} - \mathbf{G}_2^{-1})\mathbf{G}_2^{1/2}\right\| \\
&= \left\|\mathbf{G}_2^{1/2}\mathbf{G}_1^{-1}(\mathbf{G}_1 - \mathbf{G}_2)\mathbf{G}_2^{-1}\mathbf{G}_2^{1/2}\right\| \\
&\le \left(\left\|\mathbf{G}_2^{-1}\right\| \cdot \|\mathbf{G}_2\|\right)^{1/2} \cdot \left\|\mathbf{G}_1^{-1}\right\| \cdot \|\mathbf{G}_1 - \mathbf{G}_2\|.
\end{aligned}$$

□

Lemma H.4. Let $\boldsymbol{\varphi} : \mathcal{S}\times\mathcal{A} \to \mathbb{R}^d$ be a bounded function such that $|\boldsymbol{\varphi}(s,a)| \le C$ for all $(s,a) \in \mathcal{S}\times\mathcal{A}$. For any $K>0$ and $\lambda>0$, define $\bar{\mathbf{G}}_K = \sum_{k=1}^{K} \boldsymbol{\varphi}(s_k,a_k)\boldsymbol{\varphi}(s_k,a_k)^\top + \lambda\mathbf{I}_d$ where (s_k,a_k)'s are i.i.d samples from some distribution ν over $\mathcal{S}\times\mathcal{A}$. Then with probability at least $1-\delta$, it holds that

$$\left\|\frac{\bar{\mathbf{G}}_K}{K} - \mathbb{E}_\nu\left[\frac{\bar{\mathbf{G}}_K}{K}\right]\right\| \le \frac{4\sqrt{2}C^2}{\sqrt{K}}\left(\log\frac{2d}{\delta}\right)^{1/2}.$$

Proof of Lemma H.4. Denote $\mathbf{x}_k = \boldsymbol{\varphi}(s_k, a_k)$. Denote $\widetilde{\boldsymbol{\Sigma}}_h$ as the matrix obtained by replacing the k-th vector $\mathbf{x}_k$ in $\widehat{\boldsymbol{\Sigma}}_h$ by $\widetilde{\mathbf{x}}_k$ and leaving the rest $K-1$ vectors unchanged. Then we have

$$\begin{aligned}
\left(\frac{\widehat{\boldsymbol{\Sigma}}_h}{K} - \frac{\widetilde{\boldsymbol{\Sigma}}_h}{K}\right)^2 &= \left(\frac{\mathbf{x}_k\mathbf{x}_k^\top - \widetilde{\mathbf{x}}_k\widetilde{\mathbf{x}}_k^\top}{K}\right)^2 \\
&\preceq \frac{1}{K^2}\left(2\mathbf{x}_k\mathbf{x}_k^\top\mathbf{x}_k\mathbf{x}_k^\top + 2\widetilde{\mathbf{x}}_k\widetilde{\mathbf{x}}_k^\top\widetilde{\mathbf{x}}_k\widetilde{\mathbf{x}}_k^\top\right) \\
&\preceq \frac{1}{K^2}\left(2C^4\mathbf{I}_d + 2C^4\mathbf{I}_d\right) \\
&= \frac{4C^4}{K^2}\cdot\mathbf{I}_d \\
&:= \boldsymbol{A}_k^2,
\end{aligned}$$

where the first inequality uses the fact that $(\boldsymbol{A}-\boldsymbol{B})^2 \preceq 2\boldsymbol{A}^2 + 2\boldsymbol{B}^2$ for all p.s.d. matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, the second inequality is from $\|\boldsymbol{\varphi}\| \le C$. Note that we have

$$\left\|\sum_k \boldsymbol{A}_k^2\right\| = \frac{4C^4}{K}.$$

Then by Lemma H.1, we have: for all $t > 0$,

$$\mathbb{P}\left\{\left\|\frac{\widehat{\boldsymbol{\Sigma}}_h}{K} - \mathbb{E}\left[\frac{\widehat{\boldsymbol{\Sigma}}_h}{K}\right]\right\| \ge t\right\} \le 2d\cdot\exp\left(\frac{-t^2K}{32C^4}\right).$$

Equivalently, with probability at least $1-\delta$,

$$\left\|\frac{\widehat{\boldsymbol{\Sigma}}_h}{K} - \mathbb{E}\left[\frac{\widehat{\boldsymbol{\Sigma}}_h}{K}\right]\right\| \le \frac{4\sqrt{2}C^2}{\sqrt{K}}\left(\log\frac{2d}{\delta}\right)^{1/2}.$$

This completes the proof. □

Lemma H.5. Let $\boldsymbol{\varphi} : \mathcal{S}\times\mathcal{A} \to \mathbb{R}^d$ be a bounded function such that $\|\boldsymbol{\varphi}(s,a)\|_2 \le C$ for all $(s,a) \in \mathcal{S}\times\mathcal{A}$. For any $K > 0$ and $\lambda > 0$, define $\bar{\mathbf{G}}_K = \sum_{k=1}^K \boldsymbol{\varphi}(s_k,a_k)\boldsymbol{\varphi}(s_k,a_k)^\top + \lambda\mathbf{I}_d$ where (s_k,a_k)'s are i.i.d samples from some distribution ν over $\mathcal{S}\times\mathcal{A}$. Let $\mathbf{G} = \mathbb{E}_\nu[\boldsymbol{\varphi}(s,a)\boldsymbol{\varphi}(s,a)^\top]$. Then for any $\delta \in (0,1)$, if K satisfies that

$$K \ge \max\left\{512C^4\|\mathbf{G}^{-1}\|^2\log\left(\frac{2d}{\delta}\right), 4\lambda\|\mathbf{G}^{-1}\|\right\}. \tag{H.1}$$

Then with probability at least $1-\delta$, it holds simultaneously for all $\mathbf{u} \in \mathbb{R}^d$ that

$$\|\mathbf{u}\|_{\bar{\mathbf{G}}_K^{-1}} \le \frac{2}{\sqrt{K}}\|\mathbf{u}\|_{\mathbf{G}^{-1}}.$$

Proof of Lemma H.5. Note that

$$\begin{aligned}
\|\mathbf{u}\|_{\bar{\mathbf{G}}_K^{-1}} &= \frac{1}{\sqrt{K}}\sqrt{\mathbf{u}^\top\mathbf{G}^{-1}\mathbf{u} + \mathbf{u}^\top\left[\left(\frac{\bar{\mathbf{G}}_K}{K}\right)^{-1} - \mathbf{G}^{-1}\right]\mathbf{u}} \\
&= \frac{1}{\sqrt{K}}\sqrt{\mathbf{u}^\top\mathbf{G}^{-1}\mathbf{u} + \mathbf{u}^\top\mathbf{G}^{-1/2}\left[\mathbf{G}^{1/2}\left(\frac{\bar{\mathbf{G}}_K}{K}\right)^{-1}\mathbf{G}^{1/2} - \mathbf{I}_d\right]\mathbf{G}^{-1/2}\mathbf{u}} \\
&\le \frac{1}{\sqrt{K}}\left(1 + \left\|\mathbf{G}^{1/2}\left(\frac{\bar{\mathbf{G}}_K}{K}\right)^{-1}\mathbf{G}^{1/2} - \mathbf{I}_d\right\|^{1/2}\right)\|\mathbf{u}\|_{\mathbf{G}^{-1}},
\end{aligned} \tag{H.2}$$

where the last inequality follows from Cauchy-Schwartz inequality.

It then reduces to bound $\left\|\mathbf{G}^{1/2}\left(\bar{\mathbf{G}}_K/K\right)^{-1}\mathbf{G}^{1/2}-\mathbf{I}_d\right\|$, which can be further bounded by

$$\left\|\mathbf{G}^{1/2}\left(\frac{\bar{\mathbf{G}}_K}{K}\right)^{-1}\mathbf{G}^{1/2}-\mathbf{I}_d\right\| \leq \left\|\left[\mathbf{G}^{-1/2}\frac{\bar{\mathbf{G}}_K}{K}\mathbf{G}^{-1/2}\right]^{-1}\right\| \cdot \left\|\mathbf{I}_d-\mathbf{G}^{-1/2}\frac{\bar{\mathbf{G}}_K}{K}\mathbf{G}^{-1/2}\right\|. \tag{H.3}$$

By Lemma H.4, we have

$$\left\|\frac{\bar{\mathbf{G}}_K}{K}-\mathbb{E}\left[\frac{\bar{\mathbf{G}}_K}{K}\right]\right\| \leq \frac{4\sqrt{2}C^2}{\sqrt{K}}\left(\log\frac{2d}{\delta}\right)^{1/2}$$

with probability at least $1-\delta$, and thus

$$\begin{aligned}\left\|\mathbf{I}-\mathbf{G}^{-1/2}\frac{\bar{\mathbf{G}}_K}{K}\mathbf{G}^{-1/2}\right\| &\leq \left[\left\|\frac{\bar{\mathbf{G}}_K}{K}-\mathbb{E}\left[\frac{\bar{\mathbf{G}}_K}{K}\right]\right\|+\left\|\mathbb{E}\left[\frac{\bar{\mathbf{G}}_K}{K}\right]-\mathbf{G}\right\|\right]\cdot\|\mathbf{G}^{-1}\| \\ &\leq \frac{4\sqrt{2}C^2\|\mathbf{G}^{-1}\|}{\sqrt{K}}\sqrt{\log\frac{2d}{\delta}}+\frac{\lambda\|\mathbf{G}^{-1}\|}{K} \\ &\leq \frac{1}{2}\end{aligned} \tag{H.4}$$

where the last inequality follows from the assumption (H.1). Therefore,

$$\lambda_{\min}\left(\mathbf{G}^{-1/2}\frac{\bar{\mathbf{G}}_K}{K}\mathbf{G}^{-1/2}\right) \geq 1-\left\|\mathbf{I}-\mathbf{G}^{-1/2}\frac{\bar{\mathbf{G}}_K}{K}\mathbf{G}^{-1/2}\right\| \geq \frac{1}{2}$$

with probability at least $1-\delta$. This further implies that

$$\left\|\left[\mathbf{G}^{-1/2}\frac{\bar{\mathbf{G}}_K}{K}\mathbf{G}^{-1/2}\right]^{-1}\right\| = \lambda_{\min}\left(\mathbf{G}^{-1/2}\frac{\bar{\mathbf{G}}_K}{K}\mathbf{G}^{-1/2}\right)^{-1} \leq 2. \tag{H.5}$$

Combining (H.3), (H.4) and (H.5) yields that

$$\left\|\mathbf{G}^{1/2}\left(\frac{\bar{\mathbf{G}}_K}{K}\right)^{-1}\mathbf{G}^{1/2}-\mathbf{I}_d\right\| \leq 1 \tag{H.6}$$

with probability at least $1-\delta$. Then plug (H.6) back into (H.2), and we obtain that

$$\|\mathbf{u}\|_{\bar{\mathbf{G}}_K^{-1}} \leq \frac{2}{\sqrt{K}}\|\mathbf{u}\|_{\mathbf{G}^{-1}}$$

with probability at least $1-\delta$. Note that in the above argument we only need to bound $\left\|\mathbf{G}^{1/2}\left(\bar{\mathbf{G}}_K/K\right)^{-1}\mathbf{G}^{1/2}-\mathbf{I}_d\right\|$ which is independent of the choice of $\mathbf{u}$, thus it holds for all $\mathbf{u}\in\mathbb{R}^d$ simultaneously. This completes the proof. □

H.3 Inequalities for Sample Covariance Matrices

Here we introduce some useful lemmas about the inverse Gram matrix.

Lemma H.6 (Lemma D.1, Jin et al. 17)**.** Let $\mathbf{\Lambda}_t=\sum_{i=1}^t\mathbf{x}_i\mathbf{x}_i^\top+\lambda I$ where $\lambda>0$ and $\mathbf{x}_i\in\mathbb{R}^d$. Then

$$\sum_{i=1}^t\mathbf{x}_i^\top\mathbf{\Lambda}_t^{-1}\mathbf{x}_i \leq d.$$

Proof of Lemma H.6. Note that

$$\sum_{i=1}^t\mathbf{x}_i^\top\mathbf{\Lambda}_t^{-1}\mathbf{x}_i=\sum_{i=1}^t\mathrm{tr}(\mathbf{x}_i^\top\mathbf{\Lambda}_t^{-1}\mathbf{x}_i)=\mathrm{tr}\left(\mathbf{\Lambda}_t^{-1}\sum_{i=1}^t\mathbf{x}_i\mathbf{x}_i^\top\right).$$

Using the eigen-decomposition $\sum_{i=1}^t\mathbf{x}_i\mathbf{x}_i^\top=\mathbf{U}\mathrm{diag}(\lambda_1,\cdots,\lambda_d)\mathbf{U}^\top$, we have $\mathbf{\Lambda}_t=\mathbf{U}\mathrm{diag}(\lambda_1+1,\cdots,\lambda_d+1)\mathbf{U}^\top$, and it follows that

$$\mathrm{tr}\left(\mathbf{\Lambda}_t^{-1}\sum_{i=1}^t\mathbf{x}_i\mathbf{x}_i^\top\right)=\sum_{j=1}^d\frac{\lambda_j}{\lambda_j+\lambda} \leq d.$$

□

Lemma H.7. For any $h \in [H]$ and $L_1, L_2 > 0$, let $\sigma_1, \sigma_2 \in \mathcal{T}_h(L_1, L_2)$ such that $\sup_{s,a} |\sigma_1(s,a) - \sigma_2(s,a)| \le \epsilon$. Define

$$\mathbf{\Lambda}_1 := \sum_{k=1}^{K} \boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}(s_{k,h}, a_{k,h})^\top / \sigma_1(s_{k,h}, a_{k,h})^2 + \lambda \mathbf{I}_d \ ,$$

$$\mathbf{\Lambda}_2 := \sum_{k=1}^{K} \boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}(s_{k,h}, a_{k,h})^\top / \sigma_2(s_{k,h}, a_{k,h})^2 + \lambda \mathbf{I}_d \ .$$

Then under Assumption 2.1, it holds that

$$\|\mathbf{\Lambda}_1 - \mathbf{\Lambda}_2\| \le \frac{2K\sqrt{(H-h+1)^2 + \sigma_r^2} \cdot \epsilon}{(\eta_h + \sigma_r^2)^2},$$

and

$$\|\mathbf{\Lambda}_1^{-1} - \mathbf{\Lambda}_2^{-1}\| \le \frac{2K\sqrt{(H-h+1)^2 + \sigma_r^2} \cdot \epsilon}{\lambda^2(\eta_h + \sigma_r^2)^2}.$$

Proof of Lemma H.7. We have

$$\mathbf{\Lambda}_1 - \mathbf{\Lambda}_2 = \sum_{k=1}^{K} \boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}^\top(s_{k,h}, a_{k,h}) \left(\frac{1}{\sigma_1^2(s_{k,h}, a_{k,h})} - \frac{1}{\sigma_2^2(s_{k,h}, a_{k,h})} \right)$$

and thus

$$\begin{aligned}
\|\mathbf{\Lambda}_1 - \mathbf{\Lambda}_2\| &\le \sum_{k=1}^{K} \|\boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}^\top(s_{k,h}, a_{k,h})\| \cdot \left| \frac{1}{\sigma_1^2(s_{k,h}, a_{k,h})} - \frac{1}{\sigma_2^2(s_{k,h}, a_{k,h})} \right| \\
&\le \sum_{k=1}^{K} \left| \frac{1}{\sigma_1^2(s_{k,h}, a_{k,h})} - \frac{1}{\sigma_2^2(s_{k,h}, a_{k,h})} \right| \\
&= \sum_{k=1}^{K} \left| \frac{|\sigma_1(s_{k,h}, a_{k,h}) + \sigma_2(s_{k,h}, a_{k,h})| \cdot |\sigma_1(s_{k,h}, a_{k,h}) - \sigma_2(s_{k,h}, a_{k,h})|}{\sigma_1^2(s_{k,h}, a_{k,h})\sigma_2^2(s_{k,h}, a_{k,h})} \right| \\
&\le K \cdot \frac{2\sqrt{(H-h+1)^2 + \sigma_r^2}}{(\eta_h + \sigma_r^2)^2} \cdot \epsilon
\end{aligned}$$

where the first inequality is from the assumption that $\|\boldsymbol{\phi}(s,a)\| \le 1$ for all $(s,a) \in \mathcal{S} \times \mathcal{A}$ and the second inequality is by $\sigma^2(\cdot) \in [\eta_h + \sigma_r^2, (H-h+1)^2 + \sigma_r^2]$. It then follows that

$$\begin{aligned}
\|\mathbf{\Lambda}_1^{-1} - \mathbf{\Lambda}_2^{-1}\| &= \|\mathbf{\Lambda}_1^{-1} (\mathbf{\Lambda}_1 - \mathbf{\Lambda}_2) \mathbf{\Lambda}_2^{-1}\| \\
&\le \|\mathbf{\Lambda}_1^{-1}\| \cdot \|\mathbf{\Lambda}_1 - \mathbf{\Lambda}_2\| \cdot \|\mathbf{\Lambda}_2^{-1}\| \\
&\le \frac{2K\sqrt{(H-h+1)^2 + \sigma_r^2} \cdot \epsilon}{\lambda^2(\eta_h + \sigma_r^2)^2},
\end{aligned}$$

where in the last inequality we use $\|\mathbf{\Lambda}_1^{-1}\|, \ \|\mathbf{\Lambda}_1^{-2}\| \le 1/\lambda$. □

Lemma H.8. For any $h \in [H]$ and $L_1, L_2 > 0$, let $\sigma_1, \sigma_2 \in \mathcal{T}_h(L_1, L_2)$ such that $\sup_{s,a} |\sigma_1^2(s,a) - \sigma_2^2(s,a)| \le \epsilon$. Then it holds that

$$\|\mathbf{\Lambda}_1 - \mathbf{\Lambda}_2\| \le \frac{K}{(\eta + \sigma_r^2)^2} \cdot \epsilon, \ \|\mathbf{\Lambda}_1^{-1} - \mathbf{\Lambda}_2^{-1}\| \le \frac{K}{\lambda^2 (\eta + \sigma_r^2)^2} \cdot \epsilon.$$

Proof of Lemma H.8. Note that

$$\begin{aligned}
\|\mathbf{\Lambda}_1 - \mathbf{\Lambda}_2\| &\le \sum_{k=1}^{K} \|\boldsymbol{\phi}(s_{k,h}, a_{k,h})\boldsymbol{\phi}^\top(s_{k,h}, a_{k,h})\| \cdot \left| \frac{1}{\sigma_1^2(s_{k,h}, a_{k,h})} - \frac{1}{\sigma_2^2(s_{k,h}, a_{k,h})} \right| \\
&\le \sum_{k=1}^{K} \left| \frac{\sigma_2^2(s_{k,h}, a_{k,h}) - \sigma_1^2(s_{k,h}, a_{k,h})}{\sigma_1^2(s_{k,h}, a_{k,h}) \cdot \sigma_2^2(s_{k,h}, a_{k,h})} \right|,
\end{aligned}$$

and the rest follows from the proof of Lemma H.7. □

H.4 Bounds for self-normalized vector-valued martingales

Here we introduce some concentration inequalities that can be applied to bound the self-normalized martingales.

Theorem H.9 (Hoeffding inequality for self-normalized martingales, Abbasi-Yadkori et al. 1)**.** Let $\{\eta_t\}_{t=1}^{\infty}$ be a real-valued stochastic process. Let $\{\mathcal{F}_t\}_{t=0}^{\infty}$ be a filtration, such that η_t is $\mathcal{F}_t$-measurable. Assume $\eta_t \mid \mathcal{F}_{t-1}$ is zero-mean and R-subgaussian for some $R > 0$, i.e.,

$$\forall \lambda \in \mathbb{R}, \quad \mathbb{E}\left[e^{\lambda \eta_t | \mathcal{F}_{t-1}}\right] \leq e^{\lambda^2 R^2/2}.$$

Let $\{\mathbf{x}_t\}_{t=1}^{\infty}$ be an $\mathbb{R}^d$-valued stochastic process where $\mathbf{x}_t$ is $\mathcal{F}_{t-1}$-measurable. Assume $\mathbf{\Lambda}_0$ is a $d \times d$ positive definite matrix, and define $\mathbf{\Lambda}_t = \mathbf{\Lambda}_0 + \sum_{s=1}^{t} \mathbf{x}_s \mathbf{x}_s^\top$. Then, for any $\delta > 0$, with probability at least $1 - \delta$, for all $t > 0$,

$$\left\| \sum_{s=1}^{t} \mathbf{x}_s \eta_s \right\|_{\mathbf{\Lambda}_t^{-1}}^2 \leq 2R^2 \log\left(\frac{\det(\mathbf{\Lambda}_t)^{1/2} \det(\mathbf{\Lambda}_0)^{-1/2}}{\delta} \right).$$

Theorem H.10 (Bernstein inequality for self-normalized martingales, Zhou et al. 50)**.** Let $\{\eta_t\}_{t=1}^{\infty}$ be a real-valued stochastic process. Let $\{\mathcal{F}_t\}_{t=0}^{\infty}$ be a filtration, such that η_t is $\mathcal{F}_t$-measurable. Assume η_t also satisfies

$$|\eta_t| \leq R, \ \mathbb{E}[\eta_t \mid \mathcal{F}_{t-1}] = 0, \ \mathbb{E}[\eta_t^2 \mid \mathcal{F}_{t-1}] \leq \sigma^2.$$

Let $\{\mathbf{x}_t\}_{t=1}^{\infty}$ be an $\mathbb{R}^d$-valued stochastic process where $\mathbf{x}_t$ is $\mathcal{F}_{t-1}$-measurable and $\|\mathbf{x}_t\| \leq L$. Let $\mathbf{\Lambda}_t = \lambda \mathbf{I}_d + \sum_{s=1}^{t} \mathbf{x}_s \mathbf{x}_s^\top$. Then, for any $\delta > 0$, with probability at least $1 - \delta$, for all $t > 0$,

$$\left\| \sum_{s=1}^{t} \mathbf{x}_s \eta_s \right\|_{\mathbf{\Lambda}_t^{-1}} \leq 8\sigma \sqrt{d \log\left(1 + \frac{tL^2}{\lambda d}\right) \cdot \log\left(\frac{4t^2}{\delta}\right)} + 4R \log\left(\frac{4t^2}{\delta}\right).$$

H.5 Auxiliary Results for Self-normalized Martingales

Assume the function $\sigma_1(\cdot, \cdot)$ is computed using the function $V_1(\cdot)$ in the same way $\widehat{\sigma}_h$ is computed using $\widehat{V}_{h+1}^{\pi}$ as in Algorithm 1. In this way, we can view σ_1 as a function parameterized by V_1. And similar for σ_2 and V_2.

Lemma H.11. Assume V_1 and $V_2 \in \mathcal{V}_{h+1}(L)$ and satisfy $\sup_s |V_1(s) - V_2(s)| \leq \epsilon$. Then

$$\sup_{s,a} |\sigma_1(s,a) - \sigma_2(s,a)| \leq 2\sqrt{\frac{K(H-h+1)}{\lambda}} \cdot \sqrt{\epsilon},$$
$$\sup_{s,a} |\sigma_1^2(s,a) - \sigma_2^2(s,a)| \leq \frac{4K(H-h+1)}{\lambda} \epsilon.$$

Proof of Lemma H.11. By the proof of Lemma H.14, we have

$$\sup_{s,a} |\sigma_1(s,a) - \sigma_2(s,a)| \leq \sup_{s,a} \sqrt{|\sigma_1^2(s,a) - \sigma_2^2(s,a)|} \leq \sqrt{\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\| + 2(H-h+1) \cdot \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|}.$$

Note that

$$\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\| = \left\| (\widehat{\mathbf{\Sigma}}_h)^{-1} \sum_{k=1}^{K} \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})(V_1 - V_2)(\check{s}'_{k,h}) \right\| \leq \frac{K}{\lambda} \epsilon,$$

where we use $\left\| (\widehat{\mathbf{\Sigma}}_h)^{-1} \right\| \leq 1/\lambda$ and $\|\boldsymbol{\phi}(s,a)\| \leq 1$ for all (s,a). Similarly, we can show

$$\begin{aligned} \|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\| &\leq \left\| (\widehat{\mathbf{\Sigma}}_h)^{-1} \sum_{k=1}^{K} \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})(V_1(\check{s}'_{k,h})^2 - V_2(\check{s}'_{k,h})^2) \right\| \\ &\leq \left\| (\widehat{\mathbf{\Sigma}}_h)^{-1} \right\| \cdot K \cdot 2(H-h+1)\epsilon \leq \frac{2K(H-h+1)}{\lambda} \epsilon. \end{aligned}$$

Altogether, we have

$$\sup_{s,a} |\sigma_1(s,a) - \sigma_2(s,a)| \leq 2\sqrt{\frac{K(H-h+1)}{\lambda}} \cdot \sqrt{\epsilon},$$

and

$$\sup_{s,a} |\sigma_1^2(s,a) - \sigma_2^2(s,a)| \leq \frac{4K(H-h+1)}{\lambda}\epsilon.$$

□

H.6 Covering numbers of the function classes

Here we compute the covering numbers of the function classes $\mathcal{V}_h$ and $\mathcal{T}_h$.

Lemma H.12 (Covering number of the Euclidean Ball)**.** For any $\epsilon > 0$, the ϵ-covering number of the ball of radius r under the Euclidean norm satisfies $\mathcal{N}_\epsilon \leq (1 + 2r/\epsilon)^d$.

A proof of this classical result can be found, for example, in the work by Vershynin [41]. Now we give the covering number of the function class $\mathcal{V}_h(L)$ for all $h \in [H]$ and $L > 0$.

Lemma H.13. For any $h \in [H]$ and any $L > 0$, let $\mathcal{V}_h(L)$ be as defined in (C.2). Let $\mathcal{N}_\epsilon$ denote the ϵ-covering number of $\mathcal{V}_h(L)$ with respect to the distance $\text{dist}(V_1, V_2) = \sup_s |V_1(s) - V_2(s)|$. Then under Assumption 2.1, it holds that

$$\mathcal{N}_\epsilon \leq \left(1 + \frac{2L}{\epsilon}\right)^d.$$

Proof of Lemma H.13. For any $V_1, V_2 \in \mathcal{V}_h(L)$ parametrized by $\mathbf{w}_1$ and $\mathbf{w}_2$ respectively, we have

$$\text{dist}(V_1, V_2) = \sup_s |\langle \boldsymbol{\phi}_h^\pi(s), \mathbf{w}_1 - \mathbf{w}_2 \rangle| \leq \|\mathbf{w}_1 - \mathbf{w}_2\|_2 \cdot \sup_s \|\boldsymbol{\phi}_h^\pi(s)\|_2 \leq \|\mathbf{w}_1 - \mathbf{w}_2\|_2,$$

where the first inequality is by Cauchy-Schwarz inequality and the second inequality uses the assumption that $\|\boldsymbol{\phi}(s,a)\| \leq 1$.

Let $\mathcal{C}_\mathbf{w}(\epsilon)$ be an $\epsilon-$cover of the Euclidean ball $\{\mathbf{w} \in \mathbb{R}^d | \ \|\mathbf{w}\|_2 \leq L\}$. Then for any $V_1 \in \mathcal{V}_h(L)$, there exists a $V_2 \in \mathcal{V}_h(L)$ parametrized by $\mathbf{w}_2 \in \mathcal{C}_\mathbf{w}(\epsilon)$ such that $\text{dist}(V_1, V_2) \leq \epsilon$. Then we see that

$$\mathcal{N}_\epsilon \leq |\mathcal{C}_\mathbf{w}(\epsilon)| \leq \left(1 + \frac{2L}{\epsilon}\right)^d,$$

where the second inequality follows from Lemma H.12. □

Lemma H.14. For any $h \in [H]$ and $L_1, L_2 > 0$, let $\mathcal{T}_h(L_1, L_2)$ be as defined in (C.3). Let $\mathcal{N}_\epsilon$ denote the ϵ-covering number of $\mathcal{T}_h(L_1, L_2)$ with respect to the distance $\text{dist}(\sigma_1, \sigma_2) = \sup_{s,a} |\sigma_1(s,a) - \sigma_2(s,a)|$. Then under Assumption 2.1, it holds that

$$\mathcal{N}_\epsilon \leq \left(1 + \frac{4L_1}{\epsilon^2}\right)^d \cdot \left(1 + \frac{8(H-h+1)L_2}{\epsilon^2}\right)^d.$$

Proof of Lemma H.14. For any $\sigma_1, \sigma_2 \in \mathcal{T}$ which are parameterized by $(\boldsymbol{\beta}_1, \boldsymbol{\theta}_1)$ and $(\boldsymbol{\beta}_2, \boldsymbol{\theta}_2)$ respectively, we have

$$\begin{aligned}
&\text{dist}(\sigma_1, \sigma_2) \\
&= \sup_{s,a} |\sigma_1(s,a) - \sigma_2(s,a)| \\
&\leq \sup_{s,a} \sqrt{|\sigma_1^2(s,a) - \sigma_2^2(s,a)|} \\
&\leq \sup_{s,a} \sqrt{\left|\langle \boldsymbol{\phi}(s,a), \boldsymbol{\beta}_1 \rangle - \langle \boldsymbol{\phi}(s,a), \boldsymbol{\beta}_2 \rangle\right| + \left|[\langle \boldsymbol{\phi}(s,a), \boldsymbol{\theta}_1 \rangle_{[0,H-h+1]}]^2 - [\langle \boldsymbol{\phi}(s,a), \boldsymbol{\theta}_2 \rangle_{[0,H-h+1]}]^2\right|} \\
&\leq \sup_{s,a} \sqrt{\left|\langle \boldsymbol{\phi}(s,a), \boldsymbol{\beta}_1 \rangle - \langle \boldsymbol{\phi}(s,a), \boldsymbol{\beta}_2 \rangle\right| + 2(H-h+1) \cdot |\langle \boldsymbol{\phi}(s,a), \boldsymbol{\theta}_1 \rangle - \langle \boldsymbol{\phi}(s,a), \boldsymbol{\theta}_2 \rangle|} \\
&\leq \sqrt{\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\| + 2(H-h+1) \cdot \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|}.
\end{aligned}$$

where the first inequality uses the fact that $|a-b| \leq \sqrt{|a^2-b^2|}$ for any $a, b \geq 0$, the second and the third inequalities follows from the fact that $\max\{\eta_h, \cdot\}$ and the clipping $\{\cdot\}_{[0,(H-h+1)^2]}$, $\{\cdot\}_{[0,H-h+1]}$ are all contraction maps, and the last inequality is by Cauchy-Schwarz inequality and the assumption that $\|\boldsymbol{\phi}(s,a)\| \leq 1$.

In order to have $\mathrm{dist}(\sigma_1, \sigma_2) \leq \epsilon$, it suffices to have $\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}_2\| \leq \epsilon^2/2$ and $2(H-h+1)\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\| \leq \epsilon^2/2$. By Lemma H.12, in order to $\epsilon^2/2$-cover $\{\boldsymbol{\beta} : \|\boldsymbol{\beta}\| < L_1\}$ and $\epsilon^2/(4(H-h+1))$-cover $\{\boldsymbol{\theta} : \|\boldsymbol{\theta}\| \leq L_2\}$ we need

$$\mathcal{N}_\beta \leq \left(1 + \frac{4L_1}{\epsilon^2}\right)^d, \mathcal{N}_\theta \leq \left(1 + \frac{8(H-h+1)L_2}{\epsilon^2}\right)^d.$$

Altogether, to ϵ-cover $\mathcal{T}$, we have

$$\mathcal{N}_\epsilon \leq \mathcal{N}_\beta \cdot \mathcal{N}_\theta \leq \left(1 + \frac{4L_1}{\epsilon^2}\right)^d \cdot \left(1 + \frac{8(H-h+1)L_2}{\epsilon^2}\right)^d.$$

□

H.7 Bounds for the Regression Estimators

Lemma H.15. Assume $\sup_s \left|\widehat{V}^\pi_{h+1}(s)\right| \leq B$ for some $B \geq 0$. Then $\widehat{\boldsymbol{\theta}}_h$, $\widehat{\boldsymbol{\beta}}_h$ and $\widehat{\mathbf{w}}^\pi_h$ in Algorithm 1 satisfy the following:

$$\|\widehat{\boldsymbol{\theta}}_h\| \leq B\sqrt{\frac{Kd}{\lambda}}, \quad \|\widehat{\boldsymbol{\beta}}_h\| \leq B^2\sqrt{\frac{Kd}{\lambda}}, \quad \|\widehat{\mathbf{w}}^\pi_h\| \leq \frac{B+1}{\sqrt{\eta + \sigma_r^2}}\sqrt{\frac{Kd}{\lambda}}.$$

Proof of Lemma H.15. For any vector $\mathbf{v} \in \mathbb{R}^d$, we have

$$\begin{aligned}
|\mathbf{v}^\top \widehat{\boldsymbol{\theta}}_h| &= \left|\mathbf{v}^\top (\widehat{\boldsymbol{\Sigma}}_h)^{-1} \sum_{k=1}^K \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h}) \widehat{V}^\pi_{h+1}(\check{s}'_{k,h})\right| \\
&\leq \sum_{k=1}^K |\mathbf{v}^\top (\widehat{\boldsymbol{\Sigma}}_h)^{-1} \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})| \cdot \sup_s |\widehat{V}^\pi_{h+1}(s)| \\
&\leq B \cdot \sqrt{\left[\sum_{k=1}^K \mathbf{v}^\top (\widehat{\boldsymbol{\Sigma}}_h)^{-1} \mathbf{v}\right] \cdot \left[\sum_{k=1}^K \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})^\top (\widehat{\boldsymbol{\Sigma}}_h)^{-1} \boldsymbol{\phi}(\check{s}_{k,h}, \check{a}_{k,h})\right]} \\
&\leq B\|\mathbf{v}\|_2 \sqrt{\frac{K}{\lambda}} \cdot \sqrt{d},
\end{aligned}$$

where the second inequality is by Cauchy-Schwarz inequality, and the last inequality uses $\|(\widehat{\boldsymbol{\Sigma}}_h)^{-1}\| \leq 1/\lambda$ and Lemma H.6. It follows that $\|\widehat{\boldsymbol{\theta}}_h\| \leq B\sqrt{\frac{Kd}{\lambda}}$. Similarly, we have $\|\widehat{\boldsymbol{\beta}}_h\| \leq B^2\sqrt{\frac{Kd}{\lambda}}$ since $\sup_s |\widehat{V}^\pi_{h+1}(s)|^2 \leq B^2$. To bound $\|\widehat{\mathbf{w}}^\pi_h\|$, note that

$$\begin{aligned}
|\mathbf{v}^\top \widehat{\mathbf{w}}^\pi_h| &= \left|\mathbf{v}^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \sum_{k=1}^K \boldsymbol{\phi}(s_{k,h}, a_{k,h}) Y_{k,h}/\widehat{\sigma}^2_{k,h}\right| \\
&\leq \frac{B+1}{\sqrt{\eta_h + \sigma_r^2}} \cdot \sum_{k=1}^K \left|\mathbf{v}^\top \widehat{\boldsymbol{\Lambda}}_h^{-1} \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_{k,h}}\right| \\
&\leq \frac{B+1}{\sqrt{\eta_h + \sigma_r^2}} \cdot \sqrt{\left[\sum_{k=1}^K \mathbf{v}^\top (\widehat{\boldsymbol{\Lambda}}_h)^{-1} \mathbf{v}\right] \cdot \left[\sum_{k=1}^K \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_{k,h}}^\top (\widehat{\boldsymbol{\Lambda}}_h)^{-1} \frac{\boldsymbol{\phi}(s_{k,h}, a_{k,h})}{\widehat{\sigma}_{k,h}}\right]} \\
&\leq \frac{B+1}{\sqrt{\eta_h + \sigma_r^2}} \cdot \|\mathbf{v}\|_2 \sqrt{\frac{K}{\lambda}} \cdot \sqrt{d} \\
&= \left(\frac{B+1}{\sqrt{\eta_h + \sigma_r^2}} \sqrt{\frac{Kd}{\lambda}}\right) \|\mathbf{v}\|_2,
\end{aligned}$$

where the first inequality comes from

$$\left|\frac{Y_{k,h}}{\widehat{\sigma}_{k,h}}\right| = \left|\frac{r_{k,h} + \widehat{V}^{\pi}_{h+1}(s'_{k,h})}{\widehat{\sigma}_{k,h}}\right| \leq (B+1) \cdot \frac{1}{\sqrt{\eta_h + \sigma_r^2}},$$

and note that by assumption $|r_{k,h}| \leq 1 \; a.s.$, and by the clipping in the algorithm, $\widehat{\sigma}_{k,h} \geq \sqrt{\eta_h + \sigma_r^2}$.

□