

---

# Improved Sample Complexity Bounds for Distributionally Robust Reinforcement Learning

---

Zaiyan Xu\*  
Texas A&M University

Kishan Panaganti\*  
Texas A&M University

Dileep Kalathil  
Texas A&M University

## Abstract

We consider the problem of learning a control policy that is robust against the parameter mismatches between the training environment and testing environment. We formulate this as a distributionally robust reinforcement learning (DR-RL) problem where the objective is to learn the policy which maximizes the value function against the worst possible stochastic model of the environment in an uncertainty set. We focus on the tabular episodic learning setting where the algorithm has access to a generative model of the nominal (training) environment around which the uncertainty set is defined. We propose the Robust Phased Value Learning (RPVL) algorithm to solve this problem for the uncertainty sets specified by four different divergences: total variation, chi-square, Kullback-Leibler, and Wasserstein. We show that our algorithm achieves  $\tilde{O}(|\mathcal{S}||\mathcal{A}|H^5)$  sample complexity, which is uniformly better than the existing results by a factor of  $|\mathcal{S}|$ , where  $|\mathcal{S}|$  is number of states,  $|\mathcal{A}|$  is the number of actions, and  $H$  is the horizon length. We also provide the first-ever sample complexity result for the Wasserstein uncertainty set. Finally, we demonstrate the performance of our algorithm using simulation experiments.

## 1 Introduction

Training a Reinforcement Learning (RL) algorithm directly on the real-world system is expensive and possibly dangerous since the standard RL algorithms need a lot of data samples to learn even a reasonably performing policy. The traditional solution to this issue is to train the RL algorithm on a simulator before applying the trained policy

to the real-world system. However, there will inevitably be mismatches between the simulator model and the real-world system owing to a variety of factors, including approximation errors made during modeling, changes in the real-world parameters over time, and potential adversarial disruptions in the real-world. For instance, the mass, friction, sensor noise, floor terrain, and obstacles parameters of a mobile robot simulator environment can be different from that of the real-world environment. Standard RL algorithms often fail to perform well when faced with even small changes between the training and testing environments (sim-to-real gap) (Sünderhauf et al., 2018; Tobin et al., 2017; Peng et al., 2018).

Learning a policy that is robust against the model parameter mismatches between the training and testing environments is the goal of distributionally robust reinforcement learning (DR-RL). The framework of the robust Markov decision process (RMDP) (Iyengar, 2005; Nilim and El Ghaoui, 2005) is used to address the distributionally robust *planning* problem. The RMDP formulation considers a collection of models known as the *uncertainty set*, as opposed to the standard non-robust MDP which considers only one model. In DR-RL, the objective is to learn the optimal robust policy that performs well even under the worst model in the uncertainty set. The minimization over the uncertainty set makes the RMDP and DR-RL problems significantly more challenging than their non-robust counterparts.

The RMDP problem has been well studied in the literature (Xu and Mannor, 2010; Wiesemann et al., 2013; Yu and Xu, 2015; Mannor et al., 2016; Russel and Petrik, 2019), taking into account various types of uncertainty sets and computationally effective techniques. These studies, however, are only applicable to the planning problem. Learning algorithms have also been proposed to solve the DR-RL problem (Roy et al., 2017; Panaganti and Kalathil, 2021a; Wang and Zou, 2021), but they only provide asymptotic convergence guarantees. There are empirical works that address DR-RL problem using deep RL methods (Pinto et al., 2017; Derman et al., 2018, 2020; Mankowitz et al., 2020; Zhang et al., 2020a). However, these works do not provide any provable guarantees on the performance of the learned policy. Recently, there has also been works on using func-

---

\*Equal contribution. Proceedings of the 26<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2023, Valencia, Spain. PMLR: Volume 206. Copyright 2023 by the author(s).

Algorithm	Sample Complexity				
	TV	chi-square	Kullback-Leibler		Wasserstein
(Yang et al., 2021)	$\frac{ \mathcal{S} ^2  \mathcal{A}  H^5}{\rho^2 \epsilon^2}$	$\frac{(1+\rho)^2  \mathcal{S} ^2  \mathcal{A}  H^5}{(\sqrt{1+\rho}-1)^2 \epsilon^2}$	-	$\frac{ \mathcal{S} ^2  \mathcal{A}  H^5}{\rho^2 \underline{p}^2 \epsilon^2}$	-
(Zhou et al., 2021)	-	-	$\frac{\exp(\mathcal{O}(H))  \mathcal{S} ^2  \mathcal{A}  H^5}{\rho^2 \epsilon^2}$	-	-
(Panaganti and Kalathil, 2022)	$\frac{ \mathcal{S} ^2  \mathcal{A}  H^5}{\epsilon^2}$	$\frac{\rho  \mathcal{S} ^2  \mathcal{A}  H^5}{\epsilon^2}$	$\frac{\exp(\mathcal{O}(H))  \mathcal{S} ^2  \mathcal{A}  H^5}{\rho^2 \epsilon^2}$	-	-
<b>This work</b>	$\frac{ \mathcal{S}   \mathcal{A}  H^5}{\epsilon^2}$	$\frac{(1+\rho)^2  \mathcal{S}   \mathcal{A}  H^5}{(\sqrt{1+\rho}-1)^2 \epsilon^2}$	$\frac{\exp(\mathcal{O}(H))  \mathcal{S}   \mathcal{A}  H^5}{\rho^2 \epsilon^2}$	$\frac{ \mathcal{S}   \mathcal{A}  H^5}{\rho^2 \underline{p}^2 \epsilon^2}$	$\frac{(B_p + \rho^p)^2  \mathcal{S}   \mathcal{A}  H^5}{\rho^{2p} \epsilon^2}$
(Non-robust) Lower bound (Li et al., 2020)	$ \mathcal{S}   \mathcal{A}  H^4 / \epsilon^2$				

Table 1: Sample complexity result comparison between our work and the existing works for different uncertainty sets. Here,  $|\mathcal{S}|$  is number of states,  $|\mathcal{A}|$  is the number of actions,  $H$  is the horizon length,  $\rho$  is the uncertainty parameter,  $\epsilon \in (0, H)$  is the sub-optimality level,  $\underline{p}$  is the smallest positive state transition probability of the nominal model,  $p \in [1, \infty)$  is the Wasserstein distance parameter, and  $B_p$  is the maximal value of the Wasserstein distance. Results in the infinite horizon setting are adapted to finite horizon setting by using  $H$  instead of  $(1 - \gamma)^{-1}$ , where  $\gamma$  is the discount factor, and by adding an additional  $H$  to accommodate the non-stationary nature of the nominal model. *The sample complexity result of our work is uniformly better than all the existing works by a factor of  $|\mathcal{S}|$ . Our bound also matches with the non-robust RL lower bound in all factors, except in  $H$ .*

tion approximation approaches with offline data for DR-RL (Panaganti et al., 2022).

In this work, we develop a new model-based DR-RL algorithm, Robust Phased Value Learning (RPVL) algorithm, with *provable finite-sample performance guarantees in the tabular, finite-horizon RMDP setting*. We make the standard generative model assumption used in the model-based non-robust RL literature (Azar et al., 2013; Haskell et al., 2016; Agarwal et al., 2020; Li et al., 2020; Kalathil et al., 2021). More precisely, we assume that the algorithm has access to a generative model that can generate next-state samples for all state-action pairs according to the nominal model. We address the following important (sample complexity) question: *How many samples from the nominal model are required to learn an  $\epsilon$ -optimal robust policy with a high probability?*

The closest to our work in the DR-RL literature are Yang et al. (2021); Zhou et al. (2021); Panaganti and Kalathil (2022). All of these works consider the infinite horizon setting of the DR-RL problem where the nominal model is stationary and obtains a sample complexity  $\tilde{\mathcal{O}}(|\mathcal{S}|^2)$ . However, many real-world application have non-stationary dynamics and are episodic in nature (Choi et al., 2009; Schulman et al., 2013; Zhang et al., 2020a). So, in this work, we consider a non-stationary nominal model in a finite horizon (episodic) DR-RL setting. The sample complexity of our algorithm is  $\tilde{\mathcal{O}}(|\mathcal{S}| |\mathcal{A}| H^5)$ , which is superior by a factor  $|\mathcal{S}|$  compared to all the above mentioned existing works. Table 1 provides a detailed comparison of our result with that of the existing works. We note that, a very recent work (Shi and Chi, 2022) also considers the finite horizon DR-RL problem but in an offline RL setting. Their sample complexity also scales with  $\tilde{\mathcal{O}}(|\mathcal{S}|^2)$ , similar to the existing works when considering the worse-case concentrability

coefficient (Chen and Jiang, 2019; Xie and Jiang, 2020).

**Our Contributions:** (1) We propose a new model-based DR-RL algorithm called RPVL algorithm inspired by the *phased value iteration* (Kakade, 2003, Algorithm 4) that takes advantage of the non-stationary dynamics in each *phase*. In addition, we develop an uncertainty-set-specific covering number argument instead of the uniform covering number argument used in the prior works. Combining these, we are able to establish the sample complexity  $\tilde{\mathcal{O}}(|\mathcal{S}| |\mathcal{A}| H^5)$  for our RPVL algorithm. This not only improves the existing results but also matches with that of the non-robust RL lower bound in  $|\mathcal{S}|$  and  $|\mathcal{A}|$  (see Table 1).

(2) To the best of our knowledge, we provide the first-ever sample complexity result for the DR-RL problem with the Wasserstein uncertainty set (see Table 1).

(3) We demonstrate the performance of our RPVL algorithm on the Gambler’s Problem for four different uncertainty sets. We demonstrate that the learned RPVL algorithm policy is robust to changes in the model parameters for all uncertainty sets. We also show our algorithm converges with regard to the sample size. Finally, with respect to the optimality gap, we demonstrate its dependence on the uncertainty set radius parameter.

*Remark 1.* (i) The term robust RL is now broadly used in a wide variety of formulations, including in the data corruption setting (Lykouris et al., 2021) and in the adversarial learning setting (Pinto et al., 2017; Zhang et al., 2020b; Vinitzky et al., 2020; Zhang et al., 2020a). The problem addressed in our work is fundamentally different from all these, where we use the classical RMDP framework (Iyengar, 2005; Nilim and El Ghaoui, 2005), and compare with the existing works only using this framework. (ii) Distributionally robust optimization is now a well established area

(Duchi and Namkoong, 2018; Chen et al., 2020; Namkoong and Duchi, 2016), whose formulation is identical to the classical RMDP formulation. A growing number works (Zhou et al., 2021; Si et al., 2020; Shi and Chi, 2022) in the RL literature hence use the terminology *distributionally* robust to clearly establish this connection and also to avoid possible misinterpretation with the other broad use of the term robust RL. For the same reasons, we use the terminology *DR-RL* instead of robust RL.

## 2 Preliminaries and Problem Formulation

**Notations:**  $\Delta(\mathcal{S})$  denotes the probability simplex over any finite set  $\mathcal{S}$ . For any probability vector  $P \in \Delta(\mathcal{S})$  and real vector  $V \in \mathbb{R}^{|\mathcal{S}|}$ , let  $PV$  denote the expectation  $\mathbb{E}_{s \sim P}[V(s)]$ . For any positive integer  $M$ ,  $[M]$  denotes the set  $\{1, 2, \dots, M\}$ .  $\mathbf{1}$  denotes the vector of all ones whose dimension is determined from context. Let  $m(\mu, \nu) \subseteq \Delta(\mathcal{S} \times \mathcal{S})$  denote the set of all probability measures on  $\mathcal{S} \times \mathcal{S}$  such that the marginals along first and second dimensions are the distribution  $\nu$  and  $\mu$  respectively. That is, for any  $\omega \in m(\mu, \nu)$  we have  $\nu(y) = \int_x \omega(x, y) dx$  and  $\mu(x) = \int_y \omega(x, y) dy$ .

A **finite-horizon Markov Decision Process (MDP)** can be defined as a tuple  $(\mathcal{S}, \mathcal{A}, (P_h)_{h=1}^H, (r_h)_{h=1}^H, H)$ , where  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space,  $H$  is the horizon length, for any  $h \in [H]$ ,  $r_h: \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  is a known deterministic reward function, and  $P_h: \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$  is the transition probability function at time  $h$ . For any  $h \in [H]$ , and  $P_h(s'|s, a)$  represents the probability of transitioning to state  $s'$  when action  $a$  is taken at state  $s$ . We use  $P_{h,s,a}$  to denote the  $|\mathcal{S}|$ -dimensional probability vector taking value  $P_{h,s,a}(s') := P_h(s'|s, a)$  for any  $s' \in \mathcal{S}$ . We assume that  $|\mathcal{S}|$  and  $|\mathcal{A}|$  are finite. When it is clear from the context, we use the shorthand  $P$  for  $(P_h)_{h=1}^H$ .

A non-stationary Markov policy  $\pi = (\pi_h)_{h=1}^H$  is a sequence of decision rules such that  $\pi_h: \mathcal{S} \rightarrow \Delta(\mathcal{A})$  where  $\pi_h(a|s)$  specifies the probability of choosing action  $a$  in state  $s$  at time  $h$ . We consider the deterministic Markov policy class  $\Pi$ . That is, for any policy  $\pi \in \Pi$ ,  $\pi_h(s)$  is a deterministic decision rule. For any state  $s \in \mathcal{S}$  and time  $h \in [H]$ , we define the value function for policy  $\pi$  as  $V_h^{\pi, P}(s) := \mathbb{E}_{\pi, P} \left[ \sum_{t=h}^H r_t(s_t, a_t) \mid s_h = s, \pi \right]$ , where  $a_h \sim \pi_h(s_h)$  and  $s_{h+1} \sim P_h(\cdot | s_h, a_h)$ . Since we are particularly interested in the value of a policy starting from  $h = 1$ , we denote  $V^{\pi, P} := V_1^{\pi, P}$ . The optimal value function  $V_h^{*, P}$  and the optimal policy  $\pi_h^{*, P}$  of an MDP with the transition dynamics  $P$  are defined as

$$V_h^{*, P} = \max_{\pi \in \Pi} V_h^{\pi, P}, \quad \pi_h^{*, P} = \operatorname{argmax}_{\pi \in \Pi} V_h^{\pi, P}, \quad \forall h \in [H].$$

Since we assume that  $r_h(s, a) \in [0, 1]$  for all  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ , it is straightforward that any value function

is bounded by  $V_{\max} := H$ . We denote the set of all value functions  $\mathcal{V} = \{V \in \mathbb{R}^{|\mathcal{S}|} : 0 \leq \|V\|_{\infty} \leq H\}$ .

### (Distributionally) Robust Markov Decision Process:

A finite-horizon distributionally robust MDP (RMDP) can be defined as a tuple  $M = (\mathcal{S}, \mathcal{A}, \mathcal{P}, (r_h)_{h=1}^H, H)$ . At each time  $h \in [H]$ , instead of a single model, we consider the set of models within a ball centered at the nominal model  $P^o \equiv (P_h^o)_{h=1}^H$ . Formally, we define the uncertainty set as  $\mathcal{P} = \bigotimes_{h,s,a \in [H] \times \mathcal{S} \times \mathcal{A}} \mathcal{P}_{h,s,a}$  such that

$$\mathcal{P}_{h,s,a} = \{P \in \Delta(\mathcal{S}) : D(P, P_{h,s,a}^o) \leq \rho\}, \quad (1)$$

where  $D(\cdot, \cdot)$  is some distance metric between two probability measures, and  $\rho$  defines the radius of the uncertainty set. We note that, by construction, the uncertainty set  $\mathcal{P}$  satisfies  $(s, a)$ -rectangularity condition (Iyengar, 2005).

**Choices of Uncertainty Sets:** We focus on four different uncertainty sets - three corresponding to different  $f$ -divergences and one corresponding to Wasserstein metric.

(i) *Total variation uncertainty set:* Let  $\mathcal{P}^{\text{TV}} = \bigotimes_{h,s,a} \mathcal{P}_{h,s,a}^{\text{TV}}$ , where  $\mathcal{P}_{h,s,a}^{\text{TV}}$  is defined as in Eq. (1) with total variation distance

$$D_{\text{TV}}(P, P_{h,s,a}^o) = (1/2) \|P - P_{h,s,a}^o\|_1. \quad (2)$$

(ii) *Chi-square uncertainty set:* Let  $\mathcal{P}^{\chi} = \bigotimes_{h,s,a} \mathcal{P}_{h,s,a}^{\chi}$ , where  $\mathcal{P}_{h,s,a}^{\chi}$  is defined as in Eq. (1) with Chi-square distance

$$D_{\chi}(P, P_{h,s,a}^o) = \sum_{s' \in \mathcal{S}} \frac{(P(s') - P_{h,s,a}^o(s'))^2}{P_{h,s,a}^o(s')}. \quad (3)$$

(iii) *Kullback-Leibler uncertainty set:* Let  $\mathcal{P}^{\text{KL}} = \bigotimes_{h,s,a} \mathcal{P}_{h,s,a}^{\text{KL}}$ , where  $\mathcal{P}_{h,s,a}^{\text{KL}}$  is defined as in Eq. (1) with Kullback-Leibler distance

$$D_{\text{KL}}(P, P_{h,s,a}^o) = \sum_{s' \in \mathcal{S}} P(s') \log \left( \frac{P(s')}{P_{h,s,a}^o(s')} \right). \quad (4)$$

(iv) *Wasserstein uncertainty set:* Let  $\mathcal{P}^{\text{W}} = \bigotimes_{h,s,a} \mathcal{P}_{h,s,a}^{\text{W}}$ , where  $\mathcal{P}_{h,s,a}^{\text{W}}$  is defined as in Eq. (1) with Wasserstein distance

$$D_{\text{W}}(P, P_{h,s,a}^o) = \inf_{\nu \in m(P, P_{h,s,a}^o)} \int d^p(x, y) d\nu(x, y), \quad (5)$$

where the integration is over  $(x, y) \in \mathcal{S} \times \mathcal{S}$ ,  $p \in [1, \infty)$ , and  $m(P, P_{h,s,a}^o)$  denotes all probability measures on  $\mathcal{S} \times \mathcal{S}$  with marginals  $P$  and  $P_{h,s,a}^o$ . In addition, we set  $B_p := \max_{s,s'} d^p(s, s')$ .

**Robust Dynamic Programming:** Under the RMDP setting, for a policy  $\pi \in \Pi$  the robust value function  $V^{\pi}$  is defined as (Iyengar, 2005; Nilim and El Ghaoui, 2005)

$$V_h^{\pi}(s) = \inf_{P \in \mathcal{P}} V_h^{\pi, P}(s), \quad \forall (h, s) \in [H] \times \mathcal{S}. \quad (6)$$

The interpretation is that we evaluate a policy using the worst possible model in the uncertainty set. For consistency, we set  $V_{H+1}^\pi = 0$  for any  $\pi \in \Pi$ .

The optimal robust value function and the optimal policy are defined as,  $\forall (h, s) \in [H] \times \mathcal{S}$ ,

$$V_h^*(s) = \sup_{\pi \in \Pi} V_h^\pi(s), \quad \pi_h^* = \arg \sup_{\pi \in \Pi} V_h^\pi(s), \quad (7)$$

It is known that there exists at least one optimal deterministic Markov policy (Iyengar, 2005, Theorem 2.2).

To simplify the notations, denote

$$L_{\mathcal{P}_{h,s,a}} V = \inf\{PV : P \in \mathcal{P}_{h,s,a}\}. \quad (8)$$

The optimal robust value function satisfies the following *robust Bellman equation* (Iyengar, 2005, Theorem 2.1),

$$V_h^*(s) = \max_{a \in \mathcal{A}} \{r_h(s, a) + L_{\mathcal{P}_{h,s,a}} V_{h+1}^*\}, \quad (9)$$

for  $\forall (h, s) \in [H] \times \mathcal{S}$ . From the optimal robust value function, the optimal robust policy can be calculated as

$$\pi_h^*(s) = \operatorname{argmax}_{a \in \mathcal{A}} \{r_h(s, a) + L_{\mathcal{P}_{h,s,a}} V_{h+1}^*\}, \quad (10)$$

for  $\forall (h, s) \in [H] \times \mathcal{S}$ . The *robust dynamic programming* method involves performing the backward iteration using Eq. (9)-Eq. (10) to find  $V^* = (V^*)_{h=1}^H$  and  $\pi^* = (\pi^*)_{h=1}^H$ .

We also have the *robust Bellman consistency equation* (Iyengar, 2005, Theorem 2.1) which is the extension of its counterpart in standard MDP. For any  $\pi \in \Pi$  and  $\forall h \in [H]$ ,

$$V_h^\pi(s) = r_h(s, \pi_h(s)) + L_{\mathcal{P}_{h,s,\pi_h(s)}} V_{h+1}^\pi, \quad \forall s \in \mathcal{S}. \quad (11)$$

### 3 Algorithm and Sample Complexity

In order to compute  $V^*$  and  $\pi^*$ , the robust dynamic programming method requires the knowledge of the nominal model  $P^o$  and the radius of the uncertainty set  $\rho$ . Although  $\rho$  could be a design parameter, in most practical applications, we do not have access to the analytical form of the nominal model  $P^o$ . So, we assume that the nominal model  $P^o$  is unknown. Instead, similar to the standard non-robust RL setting, we assume that we only have access to samples from a generative model, which, given the state-action  $(s, a)$  and time  $h$  as inputs, will produce samples of the next state  $s'$  according to  $P_{h,s,a}^o$ . We propose a model-based algorithm called Robust Phased Value Learning (RPVL) algorithm (Algorithm 1) to learn the optimal robust policy.

RPVL algorithm is inspired by the phased value iteration in non-robust RL (Kakade, 2003, Algorithm 4). We first get a maximum likelihood estimate  $\hat{P}_h^o$  of the nominal non-stationary model  $P_h^o$  for each  $h \in [H]$  by following the standard approach in the literature (Kakade, 2003, Algorithm 4) (Azar et al., 2013, Algorithm 3). The term *phased*

---

#### Algorithm 1 Robust Phased Value Learning (RPVL)

---

- 1: **Input:** Uncertainty radius  $\rho$
  - 2: **Initialize:**  $\hat{V}_{H+1} = 0$
  - 3: **for**  $h = H, \dots, 1$  **do**
  - 4:   Compute the empirical uncertainty set  $\hat{\mathcal{P}}_{h,s,a}$  for all  $(s, a) \in (\mathcal{S}, \mathcal{A})$  according to Eq. (12)
  - 5:    $\hat{V}_h(s) = \max_a (r(s, a) + L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1})$ ,  $\forall s \in \mathcal{S}$
  - 6:    $\hat{\pi}_h(s) = \operatorname{argmax}_a (r(s, a) + L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1})$ ,  $\forall s \in \mathcal{S}$
  - 7: **end for**
  - 8: **Output:**  $\hat{\pi} = (\hat{\pi}_h)_{h=1}^H$
- 

indicates that we use different independent samples to estimate the nominal non-stationary model  $P_h^o$  for each step (phase)  $h \in [H]$ . Formally, for each  $h \in [H]$ , we generate  $N$  next-state samples corresponding to each state-action pairs. Let  $N_h(s, a, s')$  be the count of the state  $s'$  in the  $N$  total transitions from the state-action pair  $(s, a)$  in step  $h \in [H]$ . Now, the maximum likelihood estimate is given by  $\hat{P}_{h,s,a}^o(s') = N_h(s, a, s')/N$ . Given  $\hat{P}^o$ , we can get an empirical estimate of the uncertainty set  $\mathcal{P}$  as,

$$\hat{\mathcal{P}} = \bigotimes_{h,s,a \in [H] \times \mathcal{S} \times \mathcal{A}} \hat{\mathcal{P}}_{h,s,a}, \quad \text{where,} \quad (12)$$

$$\hat{\mathcal{P}}_{h,s,a} = \left\{ P \in \Delta(\mathcal{S}) : D(P, \hat{P}_{h,s,a}^o) \leq \rho \right\},$$

where  $D$  is one of the metrics specified in Eq. (2) - Eq. (5).

For finding an approximately optimal robust policy, we now consider the empirical RMDP  $\hat{M} = (\mathcal{S}, \mathcal{A}, \hat{\mathcal{P}}, (r_h)_{h=1}^H, H)$  and perform robust dynamic programming as given in Eq. (9)-Eq. (10). We formally give our proposed approach in Algorithm 1.

#### 3.1 Sample Complexity Results

Now we give the sample complexity results for all uncertainty sets based on the metrics specified in Eq. (2) - Eq. (5). We note that the total number of samples needed in the RPVL algorithm is  $N_{\text{total}} = N |\mathcal{S}| |\mathcal{A}| H$ .

**Theorem 1** (TV uncertainty set.). *Consider a finite-horizon RMDP with a total variation uncertainty set  $\mathcal{P}^{\text{TV}}$ . Fix  $\delta \in (0, 1)$ ,  $\rho > 0$ , and  $\epsilon \in (0, 8H)$ . Consider the RPVL algorithm, with the total number of samples  $N_{\text{total}} \geq N_{\text{TV}}$ , where*

$$N_{\text{TV}} = \frac{8H^5 |\mathcal{S}| |\mathcal{A}|}{\epsilon^2} \cdot \log \left( \frac{32H^3 |\mathcal{S}| |\mathcal{A}|}{\epsilon \delta} \right).$$

*Then,  $\|V_1^* - V_1^{\hat{\pi}}\|_\infty \leq \epsilon$ , with probability at least  $1 - \delta$ .*

**Remark 2** (Comparison with the sample complexity of the non-robust RL). It is known from the non-robust non-stationary RL literature (Azar et al., 2013) that, for  $\epsilon \in (0, 1)$ , no algorithm can learn an  $\epsilon$ -optimal policy with fewer than  $\tilde{\Omega}(|\mathcal{S}| |\mathcal{A}| H^4 / \epsilon^2)$  samples. Recently, Li et al. (2020) settled this sample complexity question with a

matching lower and upper bound  $\tilde{\Theta}(|\mathcal{S}||\mathcal{A}|H^4/\epsilon^2)$  for the full range of accuracy level  $\epsilon \in (0, H)$ . Our sample complexity bound given in Theorem 1 for the distributionally robust RL is  $\tilde{\mathcal{O}}(|\mathcal{S}||\mathcal{A}|H^5/\epsilon^2)$ , which matches with that of the non-robust RL bound in  $|\mathcal{S}|, |\mathcal{A}|$  and also accommodates the full range of accuracy level  $(0, H)$ , albeit worse by a factor  $H$ .

**Theorem 2** (Chi-square uncertainty set). *Consider a finite-horizon RMDP with a chi-square uncertainty set  $\mathcal{P}^\chi$ . Fix  $\delta \in (0, 1)$ ,  $\rho > 0$ , and  $\epsilon \in (0, 8H)$ . Consider the RPVL algorithm, with the total number of samples  $N_{\text{total}} \geq N_\chi$ , where*

$$N_\chi = \frac{128C_\rho^4 H^5 |\mathcal{S}||\mathcal{A}|}{(C_\rho - 1)^2 \epsilon^2} \cdot \log \left( \frac{2H |\mathcal{S}||\mathcal{A}| (1 + \frac{8C_\rho H^2}{\epsilon(C_\rho - 1)})}{\delta} \right),$$

and where  $C_\rho = \sqrt{1 + \rho}$ . Then,  $\|V^* - V^{\hat{\pi}}\|_\infty \leq \epsilon$  with probability at least  $1 - \delta$ .

We note that Remark 2 holds for the above theorem also.

For KL uncertainty set, we present two results.

**Theorem 3** (KL uncertainty set). *Consider a finite-horizon RMDP with a KL uncertainty set  $\mathcal{P}^{\text{KL}}$ . Fix  $\delta \in (0, 1)$ ,  $\rho > 0$ , and  $\epsilon \in (0, H)$ . Consider the RPVL algorithm, with the total number of samples  $N_{\text{total}} \geq N_{\text{KL}}$ , where*

$$N_{\text{KL}} = \frac{2 \exp(3H/\underline{\lambda}) H^5 |\mathcal{S}||\mathcal{A}|}{\rho^2 \epsilon^2} \cdot \log \left( \frac{8H |\mathcal{S}||\mathcal{A}|}{\underline{\lambda} \delta} \right),$$

and  $\underline{\lambda}$  is a problem dependent parameter independent of  $N_{\text{KL}}$ , then  $\|V^* - V^{\hat{\pi}}\|_\infty \leq \epsilon$  with probability at least  $1 - \delta$ .

We note that the sample complexity of the above result for the KL uncertainty set is worse by a term  $\exp(3H/\underline{\lambda})$  when compared to Theorem 1 and Theorem 2. We note that this exponential dependence on  $H$  is observed in the prior works also (Panaganti and Kalathil, 2022; Zhou et al., 2021). We can, however, replace the exponential dependence on horizon with a *problem dependent constant* as stated in the the results below.

**Theorem 4** (KL uncertainty set). *Consider a finite-horizon RMDP with a KL uncertainty set  $\mathcal{P}^{\text{KL}}$ . Fix  $\delta \in (0, 1)$ ,  $\rho > 0$ , and  $\epsilon \in (0, H)$ . Consider the RPVL algorithm, with the total number of samples  $N_{\text{total}} \geq N'_{\text{KL}}$ , where*

$$N'_{\text{KL}} = \frac{2H^5 |\mathcal{S}||\mathcal{A}|}{\rho^2 \underline{p}^2 \epsilon^2} \cdot \log \left( \frac{2H |\mathcal{S}|^2 |\mathcal{A}|}{\delta} \right),$$

and where  $\underline{p} = \min_{s', s, a, h: P_h^\circ(s'|s, a) > 0} P_h^\circ(s'|s, a)$ . Then,  $\|V^* - V^{\hat{\pi}}\|_\infty \leq \epsilon$  with probability at least  $1 - \delta$ .

The sample complexity of the RPVL algorithm with the KL uncertainty set here is  $\tilde{\mathcal{O}}(|\mathcal{S}||\mathcal{A}|H^5/\rho^2 \underline{p}^2 \epsilon^2)$ . We note that  $\underline{p}$  is a problem dependent constant and can be arbitrarily low. So,  $1/\underline{p}^2$  can even be larger than  $\exp(3H/\underline{\lambda})$ . So, this also does not offer a definitive answer to the question of finding a problem-independent sample complexity

bound for the KL uncertainty set, which is still an open question. We, however, emphasize that our sample complexity bounds given in Theorem 3 and Theorem 4 are superior compared to the existing results by a factor of  $|\mathcal{S}|$  (see Table 1 for details). In addition, Remark 2 holds for both Theorem 3 and Theorem 4.

**Theorem 5** (Wasserstein uncertainty set). *Consider a finite-horizon RMDP with a Wasserstein uncertainty set  $\mathcal{P}^{\text{W}}$ . Fix  $\delta \in (0, 1)$ ,  $\rho > 0$ , and  $\epsilon \in (0, H)$ . Consider the RPVL algorithm, with the total number of samples  $N_{\text{total}} \geq N_{\text{W}}$ , where*

$$N_{\text{W}} = \frac{8H^5 |\mathcal{S}||\mathcal{A}| (B_p + \rho^p)^2}{\rho^{2p} \epsilon^2} \cdot \log \left( \frac{16H^2 |\mathcal{S}||\mathcal{A}| (HB_p + (H \vee \rho^p))}{\rho^p \delta \epsilon} \right).$$

Then,  $\|V^* - V^{\hat{\pi}}\|_\infty \leq \epsilon$  with probability at least  $1 - \delta$ .

We note that Remark 2 also holds for the Wasserstein uncertainty set case as the sample complexity here is  $\tilde{\mathcal{O}}((B_p + \rho^p)^2 |\mathcal{S}||\mathcal{A}|H^5/\rho^{2p} \epsilon^2)$ .

## 4 Sample Complexity Analysis

Here we briefly explain the key analysis ideas of RPVL algorithm for obtaining the sample complexity in Theorem 1 - Theorem 5. The complete proofs are provided in Appendix B.1 - Appendix B.5.

*Step 1:* To bound  $\|V_1^* - V_1^{\hat{\pi}}\|_\infty$ , we split it up using the RPVL algorithm value function  $\hat{V}_1$  as  $\|V_1^* - V_1^{\hat{\pi}}\|_\infty \leq \|V_1^* - \hat{V}_1\|_\infty + \|\hat{V}_1 - V_1^{\hat{\pi}}\|_\infty$ . This split is useful since it is easier to bound these individually. We now state a result (Yang et al., 2021; Panaganti and Kalathil, 2022; Shi and Chi, 2022) which shows Lipschitz property of operators  $L_{\mathcal{P}_{h,s,a}}$  and  $L_{\hat{\mathcal{P}}_{h,s,a}}$  for all  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ .

**Lemma 1** (Panaganti and Kalathil, 2022, Lemma 1). *For any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$  and for any  $V_1, V_2 \in \mathbb{R}^{|\mathcal{S}|}$ , we have  $|L_{\mathcal{P}_{h,s,a}} V_1 - L_{\mathcal{P}_{h,s,a}} V_2| \leq \|V_1 - V_2\|_\infty$  and  $|L_{\hat{\mathcal{P}}_{h,s,a}} V_1 - L_{\hat{\mathcal{P}}_{h,s,a}} V_2| \leq \|V_1 - V_2\|_\infty$ .*

Using this and the *robust Bellman equations* Eq. (9) and Eq. (11) in tandem, we can establish the recursion

$$\|V_h^* - \hat{V}_h\|_\infty \leq \|V_{h+1}^* - \hat{V}_{h+1}^*\|_\infty + \max_{h,s,a} \left| L_{\mathcal{P}_{h,s,a}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1} \right|,$$

$$\|V_h^{\hat{\pi}} - \hat{V}_h\|_\infty \leq \|V_{h+1}^{\hat{\pi}} - \hat{V}_{h+1}^{\hat{\pi}}\|_\infty + \max_{h,s,a} \left| L_{\mathcal{P}_{h,s,a}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1} \right|.$$

Using the consistency of robust value functions at  $H + 1$ , i.e.,  $V_{H+1}^\pi = 0$  for any  $\pi \in \Pi$  and  $\hat{V}_{H+1} = 0$ , from the

recursion on  $h$ , we finally obtain

$$\begin{aligned} & \|V_1^* - V_1^{\hat{\pi}}\|_{\infty} \\ & \leq 2H \max_{h,s,a} |L_{\mathcal{P}_{h,s,a}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1}|. \end{aligned} \quad (13)$$

*Step 2:* This is the crucial step where we develop concentration inequality bounds for  $\max_{h,s,a} |L_{\mathcal{P}_{h,s,a}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1}|$  in Eq. (13). We emphasize that this step is different and depends on the structure of  $L_{\mathcal{P}_{h,s,a}}$  and  $L_{\hat{\mathcal{P}}_{h,s,a}}$  for all uncertainty sets considered in this work.

Obtaining a bound for  $\max_{h,s,a} |L_{\mathcal{P}_{h,s,a}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1}|$  is the most challenging part of our analysis. In the non-robust setting, this will be equivalent to the error term  $P_{h,s,a}^o V - \hat{P}_{h,s,a} V$ , which is unbiased and can be easily bounded using concentration inequalities. In the distributionally robust setting, establishing such a concentration is not immediate because the nonlinear nature of the function  $L(\cdot)$  leads to  $\mathbb{E}[L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1}] \neq L_{\mathcal{P}_{h,s,a}} \hat{V}_{h+1}$ . We note that the previous works (Panaganti and Kalathil, 2021b; Yang et al., 2021; Zhou et al., 2021; Panaganti and Kalathil, 2022; Shi and Chi, 2022) also develop such concentration inequalities for the distributionally robust setting. But all of them rely on developing the uniform bounds over the value function space yielding worse sample complexity with an extra  $|\mathcal{S}|$  factor. Instead, our strategy is to get amenable concentration inequalities with novel covering number arguments Propositions 2, 4, 6 and 9. In our analyses, we also make use of an important RPVL algorithm property that  $\hat{V}_{h+1}$  is independent of  $\hat{P}_{h,s,a}^o$  by construction. We emphasize that this is the first work to use this property in the DR-RL literature. This technique is inspired by (Kakade, 2003, Theorem 2.5.1) in robust RL literature. Below, we outline the analysis of the Step 2 separately for all the uncertainty sets.

**Total Variation Uncertainty Set:** We first state a result which gives an equivalent dual form characterization for the optimization problem  $L_{\mathcal{P}_{h,s,a}} V$ , where  $\mathcal{P}_{h,s,a}$  is the total variation uncertainty set  $\mathcal{P}_{h,s,a}^{\text{TV}}$ .

**Proposition 1** (Panaganti et al., 2022, Lemma 5). *Consider an RMDP  $M$  with the total variation uncertainty set  $\mathcal{P}^{\text{TV}}$ . Fix any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any value function  $V \in \mathcal{V}$ , the inner optimization problem  $L_{\mathcal{P}_{h,s,a}^{\text{TV}}} V$  has the following equivalent form*

$$\begin{aligned} L_{\mathcal{P}_{h,s,a}^{\text{TV}}} V = & - \inf_{\eta \in [0, 2H/\rho]} \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))_+] \\ & + \left( \eta - \inf_{s'' \in \mathcal{S}} V(s'') \right)_ + \cdot \rho - \eta. \end{aligned}$$

*Remark 3.* We remark that the dual form of the optimization problem  $L_{\mathcal{P}_{h,s,a}^{\text{TV}}} V$  above is a single variable convex optimization problem. This result hence is useful for both the analysis of sample complexity bound and for the practical implementation of our RPVL algorithm.

Now we develop our main concentration inequality result based on Proposition 1 and the RPVL algorithm property that  $\hat{V}_{h+1}$  is independent of  $\hat{P}_{h,s,a}^o$  by construction.

**Proposition 2.** *Fix any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\theta, \delta \in (0, 1)$  and  $\rho \in (0, 1]$ , we have, with probability at least  $1 - \delta$ ,  $|L_{\mathcal{P}_{h,s,a}^{\text{TV}}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}^{\text{TV}}} \hat{V}_{h+1}| \leq 2\theta + \sqrt{H^2 \log(4H/(\theta\delta))} / (2N)$ .*

*Remark 4.* We note that the  $\theta$  here is a covering number argument parameter (see Lemma 6) which is a novel development made in this work which also enabled us to improve the sample complexity.

Finally, the proof of Theorem 1 is now complete by the application of uniform bound over the  $[H] \times \mathcal{S} \times \mathcal{A}$  space, combined with Eq. (13).

**Chi-square Uncertainty Set:** We first develop the dual reformulation version of  $L_{\mathcal{P}_{h,s,a}^{\chi}} V$  by adapting the techniques from distributionally robust optimization (Duchi and Namkoong, 2018).

**Proposition 3.** *Consider an RMDP  $M$  with the chi-square uncertainty set  $\mathcal{P}^{\chi}$ . Fix any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any value function  $V \in \mathcal{V}$ , the inner optimization problem  $L_{\mathcal{P}_{h,s,a}^{\chi}} V$  has the following equivalent form*

$$\begin{aligned} & L_{\mathcal{P}_{h,s,a}^{\chi}} V \\ & = - \inf_{\eta \in \Xi} \left\{ C_{\rho} \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - \eta \right\}, \end{aligned}$$

where  $\Xi = [0, C_{\rho}H / (C_{\rho} - 1)]$  and  $C_{\rho} = \sqrt{\rho + 1}$ .

Remark 3 hold here as well. Now we develop our main concentration inequality result based on Proposition 3 and the RPVL algorithm property that  $\hat{V}_{h+1}$  is independent of  $\hat{P}_{h,s,a}^o$  by construction.

**Proposition 4.** *Fix any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\theta, \delta \in (0, 1)$  and  $\rho > 0$ , we have, with probability at least  $1 - \delta$ ,*

$$\begin{aligned} & |L_{\mathcal{P}_{h,s,a}^{\chi}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}^{\chi}} \hat{V}_{h+1}| \leq 2\theta + \\ & \frac{\sqrt{2}C_{\rho}^2 H}{(C_{\rho} - 1)\sqrt{N}} \left( \sqrt{\log \left( \frac{2(1 + C_{\rho}H/(\theta(C_{\rho} - 1)))}{\delta} \right)} + 1 \right). \end{aligned}$$

Remark 4 holds here as well (see Lemma 10). Finally, the proof of Theorem 2 is now complete by the application of uniform bound over the  $[H] \times \mathcal{S} \times \mathcal{A}$  space, combined with Eq. (13).

**KL Uncertainty Set:** Here also, the first step is to formulate the dual version of  $L_{\mathcal{P}_{h,s,a}^{\text{KL}}} V$  as given below.

**Proposition 5.** *Consider an RMDP  $M$  with the chi-square uncertainty set  $\mathcal{P}^{\text{KL}}$ . Fix any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ . For*

any value function  $V \in \mathcal{V}$ , the inner optimization problem  $L_{\mathcal{P}_{h,s,a}^{\text{KL}}} V$  has the following equivalent form

$$L_{\mathcal{P}_{h,s,a}^{\text{KL}}} V = - \inf_{\lambda \in [0, H/\rho]} (\lambda \rho + \lambda \log \left( \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right)).$$

Now we develop our main concentration inequality result based on Proposition 5 and the RPVL algorithm property that  $\widehat{V}_{h+1}$  is independent of  $\widehat{P}_{h,s,a}^o$  by construction.

**Proposition 6.** Fix any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\theta, \delta \in (0, 1)$  and  $\rho > 0$ , we have, with probability at least  $1 - \delta$ ,  $|L_{\mathcal{P}_{h,s,a}^{\text{KL}}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}^{\text{KL}}} \widehat{V}_{h+1}| \leq \frac{H}{\rho} \exp(H/\underline{\lambda}) \exp(\theta H) \sqrt{\log(4/(\theta \underline{\lambda} \delta))} / (2N)$ .

Remark 4 holds here as well (see Lemma 14). Finally, the proof of Theorem 3 is now complete by the application of uniform bound over the  $[H] \times \mathcal{S} \times \mathcal{A}$  space, combined with Eq. (13).

To prove Theorem 4, we follow similar steps above. We develop our main concentration inequality result based on Proposition 5, which now depends on  $\underline{p}$ , and the RPVL algorithm property that  $\widehat{V}_{h+1}$  is independent of  $\widehat{P}_{h,s,a}^o$  by construction.

**Proposition 7.** Fix any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\delta \in (0, 1)$  and  $\rho > 0$ , we have, with probability at least  $1 - \delta$ ,  $|L_{\mathcal{P}_{h,s,a}^{\text{KL}}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}^{\text{KL}}} \widehat{V}_{h+1}| \leq (H/\rho) \sqrt{\log(2|\mathcal{S}|/\delta)} / (2N\underline{p}^2)$ .

We note that the covering number argument is avoided for this result. Finally, the proof of Theorem 4 is now complete by the application of uniform bound over the  $[H] \times \mathcal{S} \times \mathcal{A}$  space, combined with Eq. (13).

**Wasserstein Uncertainty Set:** We make use of the recent results (Gao and Kleywegt, 2022) on Wasserstein distributionally robust optimization in order to develop the dual reformulation of  $L_{\mathcal{P}_{h,s,a}^{\text{W}}} V$  as given below.

**Proposition 8.** Consider an RMDP  $M$  with the Wasserstein uncertainty set  $\mathcal{P}^{\text{W}}$ . Fix any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any value function  $V \in \mathcal{V}$ , the inner optimization problem  $L_{\mathcal{P}_{h,s,a}^{\text{W}}} V$  has the following equivalent form

$$L_{\mathcal{P}_{h,s,a}^{\text{W}}} V = - \inf_{\lambda \in [0, H/\rho^p]} (\lambda \rho^p - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [\inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\}]).$$

Now we develop our main concentration inequality result based on Proposition 8 and the RPVL algorithm property that  $\widehat{V}_{h+1}$  is independent of  $\widehat{P}_{h,s,a}^o$  by construction.

**Proposition 9.** Fix any  $h, s, a \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\theta, \delta \in (0, 1)$  and  $\rho > 0$ , we have, with probability at least  $1 - \delta$ ,

$$|L_{\mathcal{P}_{h,s,a}^{\text{W}}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}^{\text{W}}} \widehat{V}_{h+1}| \leq 2\theta + \frac{H(B_p + \rho^p)}{\rho^p} \sqrt{\frac{\log \left( \frac{2HB_p + 2(H \vee \rho^p)}{\rho^p \theta \delta} \right)}{2N}}.$$

Remark 4 holds here as well (see Lemma 16). The proof of Theorem 5 is now complete by the application of uniform bound over the  $[H] \times \mathcal{S} \times \mathcal{A}$  space, combined with Eq. (13).

## 5 Experiments

We evaluate the performance of our algorithm on the *Gambler's Problem* environment (Sutton and Barto, 2018, Example 4.3). This environment is also used in the prior DR-RL works (Zhou et al., 2021; Panaganti and Kalathil, 2021b; Shi and Chi, 2022).

**Gambler's Problem:** In the gambler's problem, a gambler starts with a random balance and makes bets on a sequence of coin flips, winning the stake with heads and losing with tails. The game allows  $H$  total number of bets. The game also ends when the gambler's balance is either 0 or 50. This problem can be formulated as an episodic finite-horizon MDP, with a state space  $\mathcal{S} = \{0, 1, \dots, 50\}$  and an action space  $\mathcal{A} = \{0, 1, \dots, \min\{s, 50 - s\}\}$  at each state  $s$ . We set the horizon as  $H = 50$ . We denote the heads-up probability as  $p_h$ . Note that this particular MDP has a stationary transition kernel:  $p_h$  remains the same for all  $h \in \mathcal{H}$ . The gambler receives a reward of 1 if  $s = 50$ . The reward is 0 for all other cases. We use  $p_h^o = 0.6$  as the nominal model for training the algorithm unless we mention otherwise.

We study the following characteristics of the RPVL:

(1) *Convergence rate of the RPVL algorithm with respect to number of samples  $N$ .* For each uncertainty set, we train the RPVL policy  $\widehat{\pi}(N)$  with sample sizes  $N = 100, 300, 1000, 3000, 5000$ . We demonstrate the convergence by plotting the sub-optimality gap  $\|V^* - V^{\widehat{\pi}(N)}\|$  against  $N$ . We also plot  $1/\sqrt{N}$  for reference as our results show that the sub-optimality gap is upper bounded by some function of order  $1/\sqrt{N}$  with different coefficients for different uncertainty sets. Hence, we normalize the sub-optimality gap values of each uncertainty set by dividing maximum gap value respectively. As shown in the left plot in Fig. 1, the robust policies of all four uncertainty sets converge to their respective robust optimal policies when  $N$  increases.

(2) *Effect of robustness parameter  $\rho$ .* To show this, we consider the Wasserstein uncertainty set. We fix  $N = 1000$  and train policies with  $p_h^o = 0.6$  and  $\rho = 0.05, 0.1, 0.2, 0.3$ . In

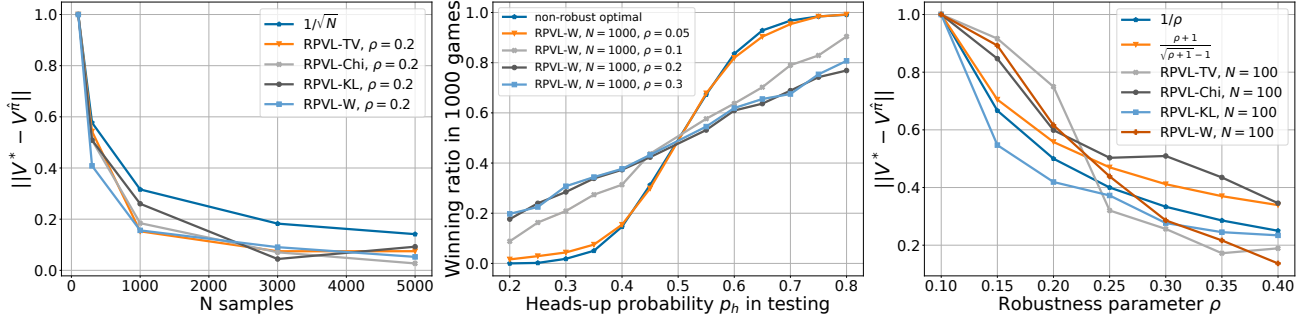


Figure 1: *Convergence of RPVL algorithm.* The left plot shows the rate of convergence with respect to the number of sample  $N$ . The middle plot shows the level of robustness of Wasserstein robust policies in testing environments with perturbed model parameter  $p_h$ . The right plot shows how sub-optimality gap changes with respect to the robustness parameter  $\rho$ .

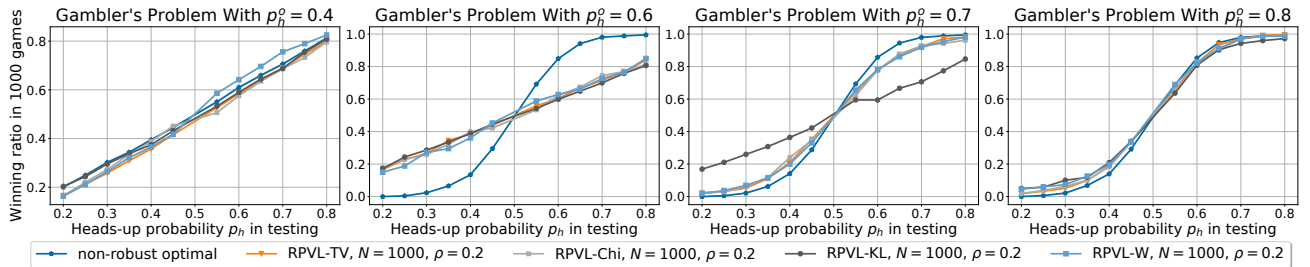


Figure 2: *Robustness fairness.* In each of the four plots, policies are trained with the nominal heads-up probability  $p_h^o$ . Each plot shows the level of robustness of the policies in perturbed testing environments with different heads-up probability  $p_h$ .

the middle plot in Fig. 1, it can be seen that the policies trained with lower  $\rho$  perform similarly to how non-robust optimal policy performs. As  $\rho$  increases, robust policies achieve better winning rate across the perturbed testing environments with  $p_h \leq 0.5$ . When  $p_h = 0.6$ , the testing environment coincides with the training environment, and the optimality of non-robust optimal policy guarantees its optimal performance.

(3) *Dependence on  $\rho$ .* When proving the main theorems, we find that the upper bound of the sub-optimality gap is some function of  $\rho$ , i.e., of order  $1/\rho$  for TV, KL, and Wasserstein and of order  $\frac{\rho+1}{\sqrt{\rho+1}-1}$  for Chi-square. We fix  $N = 100$  and perform policy evaluation on the RPVL policies trained with robustness parameter  $\rho$  in  $\{0.1, 0.15, \dots, 0.4\}$ . Note that these upper bounds have different coefficients for different uncertainty sets. Hence, we normalize the sub-optimality gap values in the same way as we do in (1). As shown in the right plot of Fig. 1, the upper bound for sub-optimality gap is indeed of the order we theorize in the proofs.

(4) *Why is robust policy robust?* Our answer follows from the analysis of *Gambler's Problem* in (Zhou et al., 2021). When the training  $p_h^o$  is less than 0.5, there is a family of optimal policies that achieve a winning ratio that equals to the testing  $p_h$ , and both the non-robust dynamic program-

ming and RPVL algorithm find such strategy (as shown in the 1st plot in Fig. 2). However, when  $p_h^o \geq 0.5$ , there is a unique non-robust optimal policy: simply bet one dollar at each coin flip. This explains that when the testing  $p_h$  increases, the winning ratio increases fast for the non-robust optimal policy, and hence it has an S-shaped curve. On the other hand, for all the robust policies with  $\rho = 0.2$ , the uncertainty sets contain adversarial models with  $p_h < 0.5$ . Thus, all four robust policies still learn the conservative betting strategy (as shown in 2nd plot in Fig. 2). Next, when the training  $p_h^o = 0.7$ , an uncertainty set with radius  $\rho = 0.2$  may or may not include the adversarial model with  $p_h < 0.5$  depending on how it measures the distance between probability measures and hence produces policies behaving differently (as shown in 3rd plot in Fig. 2). Lastly, in 4th plot in Fig. 2, since the uncertainty sets are centered way beyond  $p_h = 0.5$ , all four robust policies converge to the strategy used by non-robust optimal policy. This tells us an important characteristic of robust policies: if uncertainty set is well designed so that it includes meaningful adversarial models, then the corresponding robust policy will manifest robustness.

*Remark 5.* We would like to emphasize that our work is primarily a theoretical contribution, addressing the fundamental sample complexity question of *tabular* (finite state



and finite action) distributionally robust RL. So, the goal of this experiments section is mainly to validate the performance our RPVL algorithm in *tabular environment*. High dimensional continuous control tasks, such as the MuJoCo environment used in deep RL works, are not tabular and cannot be solved by the tabular methods, including ours. So, we do not consider such continuous control benchmark tasks in this works. We note that the prior works in DR-RL have also considered only tabular environments for experiments (Zhou et al., 2021; Panaganti and Kalathil, 2021b; Shi and Chi, 2022).

## 6 Conclusion

In this paper, we propose the Robust Phased Value Learning algorithm. Our RPVL algorithm approximates the robust Bellman updates in the standard robust dynamic programming, which is a model-based robust reinforcement learning algorithm. For four distinct uncertainty sets - total variation, chi-square, Kullback-Leibler, and Wasserstein - we present sample complexity results of the learned policy with regard to the optimal robust policy. In order to highlight the theoretical features of the RPVL algorithm, we showcase its performance on the *Gambler's Problem* environment.

This work aims to provide tighter sample complexity results for the finite-horizon robust reinforcement learning problem in the finite state space and action space regime. In comparison to other known prior works, the analyses of our RPVL algorithm provide an improvement in the sample complexity by a factor of  $|\mathcal{S}|$  uniformly for all uncertainty sets. We leave the investigation for factor  $H$  improvement to our future work.

## 7 Acknowledgements

The authors would like to thank Ruida Zhou for the valuable feedback at an earlier version. This work was supported in part by the National Science Foundation (NSF) grants NSF-CAREER-EPCN-2045783 and NSF ECCS 2038963. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsoring agencies.

## References

- Agarwal, A., Kakade, S., and Yang, L. F. (2020). Model-based reinforcement learning with a generative model is minimax optimal. In *Conference on Learning Theory*, pages 67–83. 2
- Azar, M. G., Munos, R., and Kappen, H. J. (2013). Minimax PAC bounds on the sample complexity of reinforcement learning with a generative model. *Mach. Learn.*, 91(3):325–349. 2, 4
- Boucheron, S., Lugosi, G., and Massart, P. (2013). *Concentration Inequalities: A Nonasymptotic Theory of Independence*. OUP Oxford. 12
- Chen, J. and Jiang, N. (2019). Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, pages 1042–1051. 2
- Chen, R., Paschalidis, I. C., et al. (2020). Distributionally robust learning. *Foundations and Trends® in Optimization*, 4(1-2):1–243. 3
- Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2009). Reinforcement learning and savings behavior. *The Journal of finance*, 64(6):2515–2534. 2
- Csiszár, I. (1963). Eine informationstheoretische Ungleichung und ihre Anwendung auf den Beweis der Ergodizität von Markoffschen Ketten. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 8:85–108. 12
- Derman, E., Mankowitz, D., Mann, T., and Mannor, S. (2020). A bayesian approach to robust reinforcement learning. In *Uncertainty in Artificial Intelligence*, pages 648–658. 1
- Derman, E., Mankowitz, D. J., Mann, T. A., and Mannor, S. (2018). Soft-robust actor-critic policy-gradient. In *AUAI press for Association for Uncertainty in Artificial Intelligence*, pages 208–218. 1
- Duchi, J. and Namkoong, H. (2018). Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750*. 3, 6, 12
- Gao, R. and Kleywegt, A. (2022). Distributionally robust stochastic optimization with wasserstein distance. *Mathematics of Operations Research*. 7, 23
- Haskell, W. B., Jain, R., and Kalathil, D. (2016). Empirical dynamic programming. *Mathematics of Operations Research*, 41(2):402–429. 2
- Iyengar, G. N. (2005). Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280. 1, 3, 4
- Kakade, S. M. (2003). *On the Sample Complexity of Reinforcement Learning*. Phd thesis, University of College London. 2, 4, 6, 15
- Kalathil, D., Borkar, V. S., and Jain, R. (2021). Empirical Q-Value Iteration. *Stochastic Systems*, 11(1):1–18. 2
- Li, G., Wei, Y., Chi, Y., Gu, Y., and Chen, Y. (2020). Breaking the sample size barrier in model-based reinforcement learning with a generative model. In *Advances in Neural Information Processing Systems*, volume 33, pages 12861–12872. 2, 4

- Lykouris, T., Simchowitz, M., Slivkins, A., and Sun, W. (2021). Corruption-robust exploration in episodic reinforcement learning. In *Conference on Learning Theory*, pages 3242–3245. 2
- Mankowitz, D. J., Levine, N., Jeong, R., Abdolmaleki, A., Springenberg, J. T., Shi, Y., Kay, J., Hester, T., Mann, T., and Riedmiller, M. (2020). Robust reinforcement learning for continuous control with model misspecification. In *International Conference on Learning Representations*. 1
- Mannor, S., Mebel, O., and Xu, H. (2016). Robust mdps with k-rectangular uncertainty. *Mathematics of Operations Research*, 41(4):1484–1509. 1
- Moses, A. K. and Sundaresan, R. (2011). Further results on geometric properties of a family of relative entropies. In *2011 IEEE International Symposium on Information Theory Proceedings*, pages 1940–1944. 12
- Namkoong, H. and Duchi, J. C. (2016). Stochastic gradient methods for distributionally robust optimization with f-divergences. *Advances in neural information processing systems*, 29. 3
- Nilim, A. and El Ghaoui, L. (2005). Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798. 1, 3, 20
- Panaganti, K. and Kalathil, D. (2021a). Robust reinforcement learning using least squares policy iteration with provable performance guarantees. In *Proceedings of the 38th International Conference on Machine Learning*, pages 511–520. 1
- Panaganti, K. and Kalathil, D. (2021b). Sample complexity of model-based robust reinforcement learning. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 2240–2245. IEEE. 6, 7, 9
- Panaganti, K. and Kalathil, D. (2022). Sample complexity of robust reinforcement learning with a generative model. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 9582–9602. PMLR. 2, 5, 6, 20, 21
- Panaganti, K., Xu, Z., Kalathil, D., and Ghavamzadeh, M. (2022). Robust reinforcement learning using offline data. *Advances in Neural Information Processing Systems*. 2, 6
- Peng, X. B., Andrychowicz, M., Zaremba, W., and Abbeel, P. (2018). Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 3803–3810. IEEE. 1
- Pinto, L., Davidson, J., Sukthankar, R., and Gupta, A. (2017). Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, pages 2817–2826. 1, 2
- Roy, A., Xu, H., and Pokutta, S. (2017). Reinforcement learning under model mismatch. In *Advances in Neural Information Processing Systems*, pages 3043–3052. 1
- Russel, R. H. and Petrik, M. (2019). Beyond confidence regions: Tight bayesian ambiguity sets for robust mdps. *Advances in Neural Information Processing Systems*. 1
- Schulman, J., Ho, J., Lee, A. X., Awwal, I., Bradlow, H., and Abbeel, P. (2013). Finding locally optimal, collision-free trajectories with sequential convex optimization. In *Robotics: science and systems*, volume 9, pages 1–10. Citeseer. 2
- Shapiro, A. (2017). Distributionally robust stochastic programming. *SIAM Journal on Optimization*, 27(4):2258–2275. 12
- Shi, L. and Chi, Y. (2022). Distributionally robust model-based offline reinforcement learning with near-optimal sample complexity. *arXiv preprint arXiv:2208.05767*. 2, 3, 5, 6, 7, 9
- Si, N., Zhang, F., Zhou, Z., and Blanchet, J. (2020). Distributionally robust policy evaluation and learning in offline contextual bandits. In *International Conference on Machine Learning*, pages 8884–8894. 3
- Sünderhauf, N., Brock, O., Scheirer, W., Hadsell, R., Fox, D., Leitner, J., Upcroft, B., Abbeel, P., Burgard, W., Milford, M., et al. (2018). The limits and potentials of deep learning for robotics. *The International journal of robotics research*, 37(4-5):405–420. 1
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press. 7
- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., and Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30. 1
- Vinitzky, E., Du, Y., Parvate, K., Jang, K., Abbeel, P., and Bayen, A. (2020). Robust reinforcement learning using adversarial populations. *arXiv preprint arXiv:2008.01825*. 2
- Wang, Y. and Zou, S. (2021). Online robust reinforcement learning with model uncertainty. *Advances in Neural Information Processing Systems*, 34:7193–7206. 1
- Wiesemann, W., Kuhn, D., and Rustem, B. (2013). Robust Markov decision processes. *Mathematics of Operations Research*, 38(1):153–183. 1
- Xie, T. and Jiang, N. (2020). Q\* approximation schemes for batch reinforcement learning: A theoretical comparison. In *Conference on Uncertainty in Artificial Intelligence*, pages 550–559. 2
- Xu, H. and Mannor, S. (2010). Distributionally robust Markov decision processes. In *Advances in Neural Information Processing Systems*, pages 2505–2513. 1

- Yang, W., Zhang, L., and Zhang, Z. (2021). Towards theoretical understandings of robust markov decision processes: Sample complexity and asymptotics. *arXiv preprint arXiv:2105.03863*. [2](#), [5](#), [6](#)
- Yu, P. and Xu, H. (2015). Distributionally robust counterpart in Markov decision processes. *IEEE Transactions on Automatic Control*, 61(9):2538–2543. [1](#)
- Zhang, H., Chen, H., Boning, D. S., and Hsieh, C.-J. (2020a). Robust reinforcement learning on state observations with learned optimal adversary. In *International Conference on Learning Representations*. [1](#), [2](#)
- Zhang, K., Hu, B., and Basar, T. (2020b). On the stability and convergence of robust adversarial reinforcement learning: A case study on linear quadratic systems. *Advances in Neural Information Processing Systems*, 33. [2](#)
- Zhou, Z., Bai, Q., Zhou, Z., Qiu, L., Blanchet, J., and Glynn, P. (2021). Finite-sample regret bound for distributionally robust offline tabular reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3331–3339. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#), [9](#), [21](#)

## ☕ Supplementary Materials ☕

### A Technical Results

**Lemma 2** (Hoeffding’s inequality (see Boucheron et al., 2013, Theorem 2.8)). *Let  $X_1, \dots, X_n$  be independent random variables such that  $X_i$  takes its values in  $[a_i, b_i]$  almost surely for all  $i \leq n$ . Let*

$$S = \sum_{i=1}^n (X_i - \mathbb{E}[X_i]).$$

Then for every  $t > 0$ ,

$$\mathbb{P}(S \geq t) \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$

Furthermore, if  $X_1, \dots, X_n$  are a sequence of independent, identically distributed random variables with mean  $\mu$ . Let  $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$ . Suppose that  $X_i \in [a, b], \forall i$ . Then for all  $t > 0$

$$\mathbb{P}(|\bar{X}_n - \mu| \geq t) \leq 2\exp\left(-\frac{2nt^2}{(b-a)^2}\right).$$

**Lemma 3** (Boucheron et al., 2013, Theorem 6.10). *Let  $h: \mathbb{R}^n \rightarrow \mathbb{R}$  be convex or concave and  $L$ -Lipschitz with respect to  $\|\cdot\|_2$ . Let  $Z_i$  be independent random variables with  $Z_i \in [a, b]$ .  $Z_{1:n}$  denotes the data vector  $(Z_1, \dots, Z_n)^T$ . For  $t \geq 0$ ,*

$$\mathbb{P}(|h(Z_{1:n}) - \mathbb{E}[h(Z_{1:n})]| \geq t) \leq 2\exp\left(-\frac{t^2}{2L^2(b-a)^2}\right).$$

**Lemma 4** (Duchi and Namkoong, 2018, Lemma 7). *The map  $\mathbb{R}^n \ni y \mapsto \left(\frac{1}{n} \sum_{i=1}^n |y_i|^{k_*}\right)^{\frac{1}{k_*}}$  is  $n^{-\frac{1}{k_*\sqrt{2}}}$ -Lipschitz with respect to  $\|\cdot\|_2$ .*

**Lemma 5** (Duchi and Namkoong, 2018, Lemma 8). *Let  $k_* \in [1, \infty)$  and let  $Y_i$  be an i.i.d. sequence of random variables satisfying  $\mathbb{E}[|Y|^{2k_*}] \leq C^{k_*} \mathbb{E}[|Y|^{k_*}]$  for some  $C \in \mathbb{R}_+$ . For any  $k_* \in [1, \infty)$ , we have*

$$\mathbb{E}\left[\left(\frac{1}{n} \sum_{i=1}^n |Y_i|^{k_*}\right)^{\frac{1}{k_*}}\right] \geq \mathbb{E}[|Y|^{k_*}]^{\frac{1}{k_*}} - \frac{2}{k_*} \sqrt{C} n^{-\frac{1}{k_*\sqrt{2}}},$$

where  $k = k_*/(k_* - 1)$  is the conjugate exponent of  $k_*$ .

#### A.1 Distributionally Robust Optimization (DRO) Results

The Total Variation, Chi-square, and Kullback-Liebler uncertainty sets are constructed with the  $f$ -divergence. The  $f$ -divergence between the distributions  $P$  and  $P^\circ$  is defined as

$$D_f(P \parallel P^\circ) = \int f\left(\frac{dP}{dP^\circ}\right) dP^\circ, \quad (14)$$

where  $f$  is a convex function (Csiszár, 1963; Moses and Sundaresan, 2011). We obtain different divergences for different forms of the function  $f$ , including some well-known divergences. For example,  $f(t) = |t - 1|/2$  gives Total Variation,  $f(t) = (t - 1)^2$  gives chi-square, and  $f(t) = t \log(t)$  gives Kullback-Liebler.

Let  $P^\circ$  be a distribution on the space  $\mathcal{X}$  and let  $l: \mathcal{X} \rightarrow \mathbb{R}$  be a loss function. We have the following result from the *distributionally robust optimization* literature, see e.g., (Duchi and Namkoong, 2018, Proposition 1) and (Shapiro, 2017, Section 3.2).

**Proposition 10.** *Let  $D_f$  be the  $f$ -divergence as defined in (14). Then,*

$$\sup_{D_f(P \parallel P^\circ) \leq \rho} \mathbb{E}_P[l(X)] = \inf_{\lambda > 0, \eta \in \mathbb{R}} \mathbb{E}_{P^\circ} \left[ \lambda f^* \left( \frac{l(X) - \eta}{\lambda} \right) \right] + \lambda \rho + \eta, \quad (15)$$

where  $f^*(s) = \sup_{t \geq 0} \{st - f(t)\}$  is the Fenchel conjugate.

Note that on the right hand side of (15), the expectation is taken only with respect to  $P^\circ$ . We will use the above result to derive the dual reformulation of the robust Bellman operator.

## B Proof of the Theorems

### B.1 Proof of Theorem 1

Here we provide the proofs of supporting lemmas and that of Theorem 1.

**Lemma 6** (Covering number (TV)). *Given a value function  $V \in \mathcal{V}$ , let  $\mathcal{U}_V = \{(\eta \cdot \mathbf{1} - V)_+ : \eta \in [0, H]\}$ . Fix any  $\theta \in (0, 1)$ . Denote*

$$\mathcal{N}_V(\theta) = \{(\eta \cdot \mathbf{1} - V)_+ : \eta \in \{\theta, 2\theta, \dots, N_\theta \cdot \theta\}\},$$

where  $N_\theta = \lceil H/\theta \rceil$ . Then  $\mathcal{N}_V(\theta)$  is a  $\theta$ -cover for  $\mathcal{U}_V$  with respect to  $\|\cdot\|_\infty$ , and its cardinality is bounded as  $|\mathcal{N}_V(\theta)| \leq 2H/\theta$ . Furthermore, for any  $\nu \in \mathcal{N}_V(\theta)$ , we have  $\|\nu\|_\infty \leq H$ .

*Proof.* First,  $N_\theta = \lceil H/\theta \rceil$  is the minimal number of subintervals of length  $\theta$  needed to cover  $[0, H]$ . Denote  $J_i = [(i-1)\theta, i\theta]$  to be the  $i$ -th subinterval,  $1 \leq i \leq N_\theta$ . Fix some  $\mu \in \mathcal{U}_V$ . Then  $\mu = (\eta \cdot \mathbf{1} - V)_+$ . Without loss of generality, assume this particular  $\eta \in J_i$ . Let  $\nu = ((i\theta) \cdot \mathbf{1} - V)_+$ . Now, for any  $s \in \mathcal{S}$ ,

$$\begin{aligned} |\nu(s) - \mu(s)| &= |(i\theta - V(s))_+ - (\eta - V(s))_+| \\ &\stackrel{(a)}{\leq} |i\theta - V(s) - \eta + V(s)| \\ &\leq |i\theta - (i-1)\theta| = \theta, \end{aligned}$$

where (a) follows from  $i\theta > \eta$  and the fact that  $\max\{x, 0\} - \max\{y, 0\} \leq x - y$ , if  $x > y$ . Taking maximum with respect to  $s$  on both sides, we get  $\|\nu - \mu\|_\infty \leq \theta$ . Since  $\nu \in \mathcal{N}_V(\theta)$ , this suggests  $\mathcal{N}_V(\theta)$  is a  $\theta$ -cover for  $\mathcal{U}_V$ . The cardinality bound directly follows from

$$|\mathcal{N}_V(\theta)| = N_\theta = \lceil H/\theta \rceil \leq H/\theta + 1 \leq 2H/\theta,$$

where the last inequality is due to  $0 < \theta < 1 \leq H$ . Now, for any  $\nu \in \mathcal{N}_V(\theta)$ , we can establish the following

$$\nu = (\eta \cdot \mathbf{1} - V)_+ \leq (H \mathbf{1} - V)_+ \leq H,$$

where the inequality is element-wise. □

**Lemma 7.** *Fix any  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . Fix any value function  $V \in \mathcal{V}$ . Let  $\mathcal{N}_V(\theta)$  be the  $\theta$ -cover of  $\mathcal{U}_V = \{(\eta \cdot \mathbf{1} - V)_+ : \eta \in [0, H]\}$  as described in Lemma 6. We then have*

$$\sup_{\eta \in [0, H]} \left| \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s, a)} [(\eta - V(s'))_+] - \mathbb{E}_{s' \sim P_h^o(\cdot|s, a)} [(\eta - V(s'))_+] \right| \leq \max_{\nu \in \mathcal{N}_V(\theta)} |\widehat{P}_{h, s, a}^o \nu - P_{h, s, a}^o \nu| + 2\theta.$$

*Proof.* For any  $\mu \in \mathcal{U}_V$ , there exists  $\nu \in \mathcal{N}_V(\theta)$  such that  $\|\mu - \nu\|_\infty \leq \theta$ . Now for such particular  $\mu$  and  $\nu$ , we have

$$\begin{aligned} |\widehat{P}_{h, s, a}^o \mu - P_{h, s, a}^o \mu| &\leq |\widehat{P}_{h, s, a}^o \mu - \widehat{P}_{h, s, a}^o \nu| + |\widehat{P}_{h, s, a}^o \nu - P_{h, s, a}^o \nu| + |P_{h, s, a}^o \nu - P_{h, s, a}^o \mu| \\ &\leq \|\widehat{P}_{h, s, a}^o\|_1 \|\mu - \nu\|_\infty + |\widehat{P}_{h, s, a}^o \nu - P_{h, s, a}^o \nu| + \|P_{h, s, a}^o\|_1 \|\nu - \mu\|_\infty \\ &\leq \max_{\nu \in \mathcal{N}_V(\theta)} |\widehat{P}_{h, s, a}^o \nu - P_{h, s, a}^o \nu| + 2\theta. \end{aligned}$$

Taking maximum over  $\mathcal{U}_V$  on both sides, we get

$$\sup_{\mu \in \mathcal{U}_V} |\widehat{P}_{h, s, a}^o \mu - P_{h, s, a}^o \mu| \leq \max_{\nu \in \mathcal{N}_V(\theta)} |\widehat{P}_{h, s, a}^o \nu - P_{h, s, a}^o \nu| + 2\theta.$$

Now note that by the definition of  $\mathcal{U}_V$ , we have

$$\sup_{\eta \in [0, H]} \left| \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s, a)} [(\eta - V(s'))_+] - \mathbb{E}_{s' \sim P_h^o(\cdot|s, a)} [(\eta - V(s'))_+] \right| = \sup_{\mu \in \mathcal{U}_V} |\widehat{P}_{h, s, a}^o \mu - P_{h, s, a}^o \mu|.$$

The desired result directly follows. □

**Lemma 8.** Fix any value function  $V \in \mathcal{V}$  and fix  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\theta, \delta \in (0, 1)$  and  $\rho > 0$ , we have, with probability at least  $1 - \delta$ ,

$$\left| L_{\mathcal{P}_{h,s,a}} V - L_{\widehat{\mathcal{P}}_{h,s,a}} V \right| \leq \sqrt{\frac{H^2 \log(4H/\theta\delta)}{2N}} + 2\theta,$$

where  $N$  is the number of samples used to approximate  $P_{h,s,a}^o$ .

*Proof.* Fix any value function  $V \in \mathcal{V}$  and  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . From Proposition 1, we have

$$\begin{aligned} L_{\mathcal{P}_{h,s,a}} V &= - \inf_{\eta \in [0, 2H/\rho]} \left\{ \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))_+] + \left( \eta - \inf_{s'' \in \mathcal{S}} V(s'') \right)_+ \cdot \rho - \eta \right\}, \\ L_{\widehat{\mathcal{P}}_{h,s,a}} V &= - \inf_{\eta \in [0, 2H/\rho]} \left\{ \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))_+] + \left( \eta - \inf_{s'' \in \mathcal{S}} V(s'') \right)_+ \cdot \rho - \eta \right\}. \end{aligned}$$

Fix any  $\rho > 0$ . Now it follows that

$$\begin{aligned} \left| L_{\mathcal{P}_{h,s,a}} V - L_{\widehat{\mathcal{P}}_{h,s,a}} V \right| &= \left| \inf_{\eta \in [0, 2H/\rho]} \left\{ \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))_+] + \left( \eta - \inf_{s'' \in \mathcal{S}} V(s'') \right)_+ \cdot \rho - \eta \right\} \right. \\ &\quad \left. - \inf_{\eta \in [0, 2H/\rho]} \left\{ \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))_+] + \left( \eta - \inf_{s'' \in \mathcal{S}} V(s'') \right)_+ \cdot \rho - \eta \right\} \right| \\ &\stackrel{(a)}{\leq} \sup_{\eta \in [0, 2H/\rho]} \left| \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))_+] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))_+] \right| \\ &\leq \max \left\{ \sup_{\eta \in [0, H]} \left| \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))_+] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))_+] \right|, \right. \\ &\quad \left. \sup_{\eta \in [H, 2H/\rho]} \left| \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))_+] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))_+] \right| \right\} \\ &\stackrel{(b)}{=} \max \left\{ \sup_{\eta \in [0, H]} \left| \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))_+] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))_+] \right|, \right. \\ &\quad \left. \left| \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} [V(s')] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [V(s')] \right| \right\} \\ &\stackrel{(c)}{\leq} \max \left\{ \max_{\nu \in \mathcal{N}_V(\theta)} |\widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu| + 2\theta, \left| \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} [V(s')] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [V(s')] \right| \right\}, \end{aligned} \tag{16}$$

where (a) follows from the fact that  $|\inf_x f(x) - \inf_x g(x)| \leq \sup_x |f(x) - g(x)|$ . For (b), recall that  $\|V\|_\infty \leq H$  for any  $V \in \mathcal{V}$ . Hence, the term  $\eta - V(s')$  is always non-negative for  $\eta \in [H, 2H/\rho]$ , which cancels out by linearity of the expectation. (c) follows from applying Lemma 7 to the first term. Recall that all  $\nu \in \mathcal{N}_V(\theta)$  is upper bounded by  $\nu_{\max} := H$ . Now we can apply Hoeffding's inequality (Lemma 2) to the first term in Eq. (16):

$$\mathbb{P} \left( |\widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu| \geq \epsilon \right) \leq 2 \exp \left( -\frac{2N\epsilon^2}{\nu_{\max}^2} \right) = 2 \exp \left( -\frac{2N\epsilon^2}{H^2} \right), \quad \forall \epsilon > 0.$$

Now choose  $\epsilon = \sqrt{\frac{H^2 \log(2|\mathcal{N}_V(\theta)|/\delta)}{2N}}$  and recall that  $|\mathcal{N}_V(\theta)| \leq 2H/\theta$  from Lemma 6. We have

$$\begin{aligned} \mathbb{P} \left( \left| \widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu \right| \geq \sqrt{\frac{H^2 \log(4H/\theta\delta)}{2N}} \right) &\leq \mathbb{P} \left( \left| \widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu \right| \geq \sqrt{\frac{H^2 \log(2|\mathcal{N}_V(\theta)|/\delta)}{2N}} \right) \\ &\leq \frac{\delta}{|\mathcal{N}_V(\theta)|}. \end{aligned}$$

Applying a union bound over  $\mathcal{N}_V(\theta)$ , we get

$$\max_{\nu \in \mathcal{N}_V(\theta)} |\widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu| \leq \sqrt{\frac{H^2 \log(4H/\theta\delta)}{2N}}, \tag{17}$$

with probability at least  $1 - \delta$ . Now we can also apply Hoeffding's inequality to the second term in Eq. (16). Recall that any value function is bounded by  $V_{\max} := H$ . We have

$$|\widehat{P}_{h,s,a}^o V - P_{h,s,a}^o V| \leq \sqrt{\frac{H^2 \log(2/\delta)}{2N}}, \quad (18)$$

with probability at least  $1 - \delta$ . Combining Eq. (16) - Eq. (18) completes the proof.  $\square$

**Corollary 1.** *Lemma 8 holds true for any random vector  $\widehat{V}$  independent of  $\widehat{P}_{h,s,a}^o$ .*

*Proof.* Note that Lemma 8 holds true for any fixed value function  $V$ . It is then true for any realization of the random vector  $\widehat{V}$ . Now since  $\widehat{V}$  is independent of  $\widehat{P}_{h,s,a}^o$ , the result directly follows from the law of total probability.  $\square$

**Proof of Proposition 2.** Recall that  $\widehat{V}_{h+1}$  is independent of  $\widehat{P}_{h,s,a}^o$  by construction. The result directly follows from Lemma 8 and Corollary 1.  $\square$

We now have all the ingredients to prove our main result.

**Proof of Theorem 1.** The proof is inspired by the proof of (Kakade, 2003, Theorem 2.5.1). Recall that the optimal robust value function of the RMDP  $M = (\mathcal{S}, \mathcal{A}, \mathcal{P}, (r_h)_{h=1}^H, [H])$  is characterized by a set of functions  $\{V_h^*\}_{h=1}^H$ . Also, at each time  $h$ , the RPVL algorithm outputs  $\widehat{\pi}_h$  and  $\widehat{V}_h$ . Now,

$$\|V_1^* - V_1^{\widehat{\pi}}\|_{\infty} \leq \|V_1^* - \widehat{V}_1\|_{\infty} + \|\widehat{V}_1 - V_1^{\widehat{\pi}}\|_{\infty}. \quad (19)$$

*Step 1: bounding  $V_1^* - \widehat{V}_1$  in Eq. (19).* For any state  $s \in \mathcal{S}$  and time  $h \in [H]$ , we have

$$\begin{aligned} |V_h^*(s) - \widehat{V}_h(s)| &= \left| \max_a \{r_h(s, a) + L_{\mathcal{P}_{h,s,a}} V_{h+1}^*\} - \max_a \{r_h(s, a) + L_{\widehat{\mathcal{P}}_{h,s,a}} \widehat{V}_{h+1}\} \right| \\ &\stackrel{(a)}{\leq} \max_a \left| L_{\mathcal{P}_{h,s,a}} V_{h+1}^* - L_{\widehat{\mathcal{P}}_{h,s,a}} \widehat{V}_{h+1} \right| \\ &\leq \max_a \left| L_{\mathcal{P}_{h,s,a}} V_{h+1}^* - L_{\mathcal{P}_{h,s,a}} \widehat{V}_{h+1} \right| + \max_a \left| L_{\mathcal{P}_{h,s,a}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}} \widehat{V}_{h+1} \right| \\ &\stackrel{(b)}{\leq} \|V_{h+1}^* - \widehat{V}_{h+1}\|_{\infty} + \max_a \left| L_{\mathcal{P}_{h,s,a}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}} \widehat{V}_{h+1} \right| \\ &\leq \|V_{h+1}^* - \widehat{V}_{h+1}\|_{\infty} + \max_{h,s,a} \left| L_{\mathcal{P}_{h,s,a}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}} \widehat{V}_{h+1} \right|, \end{aligned}$$

where (a) follows from  $|\max_x f(x) - \max_x g(x)| \leq \max_x |f(x) - g(x)|$ . (b) follows from Lemma 1. Taking maximum over  $\mathcal{S}$ , we get

$$\|V_h^* - \widehat{V}_h\|_{\infty} \leq \|V_{h+1}^* - \widehat{V}_{h+1}\|_{\infty} + \max_{h,s,a} \left| L_{\mathcal{P}_{h,s,a}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}} \widehat{V}_{h+1} \right|. \quad (20)$$

Now we use the trivial fact that  $V_{H+1}^* = 0 = \widehat{V}_{H+1}$  and recursively apply Eq. (20) to get

$$\|V_1^* - \widehat{V}_1\|_{\infty} \leq H \max_{h,s,a} \left| L_{\mathcal{P}_{h,s,a}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}} \widehat{V}_{h+1} \right|. \quad (21)$$

*Step 2: bounding  $\widehat{V}_1 - V_1^{\widehat{\pi}}$  in Eq. (19).* Again, fix any  $s \in \mathcal{S}$  and time  $h \in [H]$ . Note that  $\widehat{\pi}_h$  is the greedy policy with respect to  $\widehat{V}_{h+1}$  as defined in the REVI algorithm. Recall the robust Bellman (consistency) equation (Eq. (11)). For any deterministic policy  $\pi = (\pi_h)_{h \in [H]}$ , we have  $V_h^{\pi}(s) = r_h(s, \pi_h(s)) + L_{\mathcal{P}_{h,s,\pi_h(s)}} V_{h+1}^{\pi}$ . Now using these, we have

$$\begin{aligned} |V_h^{\widehat{\pi}}(s) - \widehat{V}_h(s)| &= \left| r_h(s, \widehat{\pi}_h(s)) + L_{\mathcal{P}_{h,s,\widehat{\pi}_h(s)}} V_{h+1}^{\widehat{\pi}} - r_h(s, \widehat{\pi}_h(s)) - L_{\widehat{\mathcal{P}}_{h,s,\widehat{\pi}_h(s)}} \widehat{V}_{h+1} \right| \\ &= \left| L_{\mathcal{P}_{h,s,\widehat{\pi}_h(s)}} V_{h+1}^{\widehat{\pi}} - L_{\mathcal{P}_{h,s,\widehat{\pi}_h(s)}} \widehat{V}_{h+1} + L_{\mathcal{P}_{h,s,\widehat{\pi}_h(s)}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,\widehat{\pi}_h(s)}} \widehat{V}_{h+1} \right| \\ &\leq \left| L_{\mathcal{P}_{h,s,\widehat{\pi}_h(s)}} V_{h+1}^{\widehat{\pi}} - L_{\mathcal{P}_{h,s,\widehat{\pi}_h(s)}} \widehat{V}_{h+1} \right| + \max_{h,s,a} \left| L_{\mathcal{P}_{h,s,a}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}} \widehat{V}_{h+1} \right| \\ &\leq \|V_{h+1}^{\widehat{\pi}} - \widehat{V}_{h+1}\|_{\infty} + \max_{h,s,a} \left| L_{\mathcal{P}_{h,s,a}} \widehat{V}_{h+1} - L_{\widehat{\mathcal{P}}_{h,s,a}} \widehat{V}_{h+1} \right|. \end{aligned}$$

Similar to the *step 1*, after taking maximum over  $\mathcal{S}$  and unrolling the recursion, we get

$$\left\| V_1^{\hat{\pi}} - \hat{V}_1 \right\|_{\infty} \leq H \max_{h,s,a} \left| L_{\mathcal{P}_{h,s,a}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1} \right|. \quad (22)$$

*Step 3: concentration on  $|L_{\mathcal{P}_{h,s,a}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1}|$ .* Applying Proposition 2 and taking a union bound over  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ , we get

$$\max_{h,s,a} \left| L_{\mathcal{P}_{h,s,a}} \hat{V}_{h+1} - L_{\hat{\mathcal{P}}_{h,s,a}} \hat{V}_{h+1} \right| \leq \sqrt{\frac{H^2 \log(4H^2 |\mathcal{S}| |\mathcal{A}| / \theta \delta)}{2N}} + 2\theta, \quad (23)$$

with probability at least  $1 - \delta$ . Now, plugging Eq. (21)-Eq. (23) into Eq. (19), we get

$$\left\| V_1^* - V_1^{\hat{\pi}} \right\|_{\infty} \leq 2H \sqrt{\frac{H^2 \log(4H^2 |\mathcal{S}| |\mathcal{A}| / \theta \delta)}{2N}} + 4H\theta, \quad (24)$$

with probability at least  $1 - \delta$ . We can choose  $\theta = \epsilon / (8H)$ . Note that since  $\epsilon \in (0, 8H)$ , this particular  $\theta$  is in  $(0, 1)$ . Now, if we choose

$$N \geq \frac{8H^4 \log(32H^3 |\mathcal{S}| |\mathcal{A}| / \epsilon \delta)}{\epsilon^2},$$

we get  $\|V_1^* - V_1^{\hat{\pi}}\|_{\infty} \leq \epsilon$  with probability at least  $1 - \delta$ .  $\square$

## B.2 Proof of Theorem 2

Here we provide the proofs of supporting lemmas and that of Theorem 2.

**Lemma 9.** *Let  $D_f$  be defined as in Eq. (14) with the convex function  $f(t) = (t - 1)^2$  corresponding to the Chi-square uncertainty set. Then*

$$D_f(P \| P^o) \leq \rho \quad \mathbb{E}_P[l(X)] = - \inf_{\eta \in \mathbb{R}} \left\{ \sqrt{\rho + 1} \sqrt{\mathbb{E}_{P^o}[(\eta - l(X))^2]} - \eta \right\}.$$

*Proof.* We first get the Fenchel conjugate of Chi-square divergence. We have

$$f^*(s) = \sup_{t \geq 0} (st - (t - 1)^2) = \sup_{t \geq 0} (-t^2 + (s + 2)t - 1).$$

It is easy to see that for  $s \leq -2$ , we have  $f^*(s) = -1$ . Whenever  $s > -2$ , from Calculus, the optimal  $t$  for this optimization problem is  $(s + 2)/2$ . Hence we get  $f^*(s) = \max\{s^2/4 + s, -1\}$ . Now, from Proposition 10 we get,

$$\begin{aligned} \sup_{D_f(P \| P^o) \leq \rho} \mathbb{E}_P[l(X)] &= \inf_{\lambda > 0, \eta \in \mathbb{R}} \left\{ \mathbb{E}_{P^o} \left[ \lambda f^* \left( \frac{l(X) - \eta}{\lambda} \right) \right] + \lambda \rho + \eta \right\} \\ &= \inf_{\lambda, \eta: \lambda > 0, \eta \in \mathbb{R}} \left\{ \lambda \mathbb{E}_{P^o} \left[ \max \left\{ \frac{(l(X) - \eta)^2}{4\lambda^2} + \frac{l(X) - \eta}{\lambda}, -1 \right\} \right] + \lambda \rho + \eta \right\} \\ &= \inf_{\lambda, \eta: \lambda > 0, \eta \in \mathbb{R}} \left\{ \mathbb{E}_{P^o} \left[ \max \left\{ \frac{(l(X) - \eta)^2}{4\lambda} + l(X) - \eta, -\lambda \right\} \right] + \lambda \rho + \eta \right\} \\ &= \inf_{\lambda, \eta: \lambda > 0, \eta \in \mathbb{R}} \left\{ \mathbb{E}_{P^o} \left[ \left( \frac{(l(X) - \eta)^2}{4\lambda} + l(X) - \eta + \lambda \right)_+ - \lambda \right] + \lambda \rho + \eta \right\} \\ &= \inf_{\lambda, \eta: \lambda > 0, \eta \in \mathbb{R}} \left\{ \frac{1}{4\lambda} \mathbb{E}_{P^o} [(l(X) - \eta + 2\lambda)^2] + \lambda(\rho - 1) + \eta \right\} \\ &= \inf_{\lambda, \eta': \lambda > 0, \eta' \in \mathbb{R}} \left\{ \frac{1}{4\lambda} \mathbb{E}_{P^o} [(l(X) - \eta')^2] + \lambda(\rho + 1) + \eta' \right\}, \end{aligned}$$

where the fourth equality follows from the fact that  $\max\{x, y\} = (x - y)_+ + y$  for any  $x, y \in \mathbb{R}$ , and the last equality follows by making the substitution  $\eta' = \eta - 2\lambda$ . Taking the optimal value of  $\lambda$ , i.e.,  $\lambda = \sqrt{\mathbb{E}_{P^o}[(l(X) - \eta')^2]} / (2\sqrt{\rho + 1})$ , we get

$$\sup_{D_f(P \| P^o) \leq \rho} \mathbb{E}_P[l(X)] = \inf_{\eta' \in \mathbb{R}} \left\{ \sqrt{\rho + 1} \sqrt{\mathbb{E}_{P^o}[(l(X) - \eta')^2]} + \eta' \right\}.$$



Now,

$$\begin{aligned}
 \inf_{D_f(P\|P^o) \leq \rho} \mathbb{E}_P[l(X)] &= - \sup_{D_f(P\|P^o) \leq \rho} \mathbb{E}_P[-l(X)] \\
 &= - \inf_{\eta' \in \mathbb{R}} \left\{ \sqrt{\rho + 1} \sqrt{\mathbb{E}_{P^o} [(-l(X) - \eta')^2]} + \eta' \right\} \\
 &= - \inf_{\eta \in \mathbb{R}} \left\{ \sqrt{\rho + 1} \sqrt{\mathbb{E}_{P^o} [(l(X) - \eta)^2]} - \eta \right\},
 \end{aligned}$$

which completes the proof.  $\square$

**Proof of Proposition 3.** Applying Lemma 9 to value function, we get

$$L_{\mathcal{P}_{h,s,a}} V = - \inf_{\eta \in \mathbb{R}} \left\{ \sqrt{\rho + 1} \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - \eta \right\}. \quad (25)$$

Now let  $h(\eta) = \sqrt{\rho + 1} \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - \eta$ .  $h$  is convex in dual variable  $\eta$ . Now observe that  $h(\eta) \geq 0$  when  $\eta \leq 0$  and  $h(\frac{\sqrt{\rho+1}}{\sqrt{\rho+1}-1}H) \geq 0$ . Hence it is sufficient to consider  $\eta \in [0, \frac{\sqrt{\rho+1}}{\sqrt{\rho+1}-1}H]$ . Setting  $C_\rho = \sqrt{\rho + 1}$ , we get the desired result.  $\square$

**Lemma 10.** Fix any  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . Let  $\mathcal{N}_\rho(\theta)$  be the  $\theta$ -cover of the interval  $[0, C_\rho H / (C_\rho - 1)]$ , where  $C_\rho = \sqrt{1 + \rho}$ . For any value function  $V \in \mathcal{V}$  and  $\eta \in [0, C_\rho H / (C_\rho - 1)]$ , we have

$$\begin{aligned}
 \sup_{\eta \in [0, C_\rho H / (C_\rho - 1)]} & \left| \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} \right| \\
 & \leq \max_{\nu \in \mathcal{N}_\rho(\theta)} \left| \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} - \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} \right| + 2\theta.
 \end{aligned}$$

Furthermore, we have  $|\mathcal{N}_\rho(\theta)| \leq \frac{C_\rho H}{(C_\rho - 1)\theta} + 1$ .

*Proof.* Fix any  $\eta \in [0, C_\rho H / (C_\rho - 1)]$ . Then there exists a  $\nu \in \mathcal{N}_\rho(\theta)$  such that  $|\eta - \nu| \leq \theta$ . Let  $X$  be a random variable that takes values in  $\{V(1), \dots, V(|\mathcal{S}|)\}$  with probability  $\mathbb{P}(X = V(s')) = P_h^o(s' | s, a)$ , for all  $s' \in \mathcal{S}$ . We use  $\|\cdot\|_{p,P}$  to denote the  $L^p$  norm in the measure space on  $\mathcal{S}$  possessing a measure determined by the probability mass function  $P$ . It leads to the  $L^p$  norm of a random variable:  $\|X\|_{p,P} = (\mathbb{E}_P |X|^p)^{1/p}$ . Now using these definitions, for the particular  $\eta$  and  $\nu$  we picked, we have

$$\begin{aligned}
 & \left| \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} \right| = \left| \|\eta - X\|_{2, \hat{P}_{h,s,a}^o} - \|\eta - X\|_{2, P_{h,s,a}^o} \right| \\
 & \leq \left| \|\eta - X\|_{2, \hat{P}_{h,s,a}^o} - \|\nu - X\|_{2, \hat{P}_{h,s,a}^o} \right| + \left| \|\nu - X\|_{2, \hat{P}_{h,s,a}^o} - \|\nu - X\|_{2, P_{h,s,a}^o} \right| + \left| \|\nu - X\|_{2, P_{h,s,a}^o} - \|\eta - X\|_{2, P_{h,s,a}^o} \right| \\
 & \stackrel{(a)}{\leq} \|\eta - X - \nu + X\|_{2, \hat{P}_{h,s,a}^o} + \left| \|\nu - X\|_{2, \hat{P}_{h,s,a}^o} - \|\nu - X\|_{2, P_{h,s,a}^o} \right| + \|\nu - X - \eta + X\|_{2, P_{h,s,a}^o} \\
 & \leq \max_{\nu \in \mathcal{N}_\rho(\theta)} \left| \|\nu - X\|_{2, \hat{P}_{h,s,a}^o} - \|\nu - X\|_{2, P_{h,s,a}^o} \right| + 2|\nu - \eta| \\
 & \leq \max_{\nu \in \mathcal{N}_\rho(\theta)} \left| \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} - \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} \right| + 2\theta,
 \end{aligned}$$

where (a) is due to reverse triangle inequality. Taking supremum over  $\eta \in [0, C_\rho H / (C_\rho - 1)]$  on the both sides of the above, we have the desired result. For the cardinality, it is trivial that

$$|\mathcal{N}_\rho(\theta)| \leq \left\lceil \frac{C_\rho H}{(C_\rho - 1)\theta} \right\rceil \leq \frac{C_\rho H}{(C_\rho - 1)\theta} + 1. \quad (26)$$

$\square$

**Lemma 11.** Fix any value function  $V \in \mathcal{V}$  and fix  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . We have the following inequality with probability at least  $1 - \delta$ :

$$\left| C_\rho \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} - C_\rho \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} \right| \leq \frac{\sqrt{2}C_\rho^2 H}{(C_\rho - 1)\sqrt{N}} \left( \sqrt{\log(2/\delta)} + 1 \right),$$

where  $C_\rho = \sqrt{1 + \rho}$ .

*Proof.* Fix any  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ ,  $\eta \in [0, C_\rho H / (C_\rho - 1)]$ , and  $V \in \mathcal{V}$  independent of  $P_{h,s,a}^o$ . Consider a sequence of i.i.d. samples  $\{X_i\}_{i=1}^N$  generated from  $P_{h,s,a}^o$ . Recall that we use the outcomes of this sequence of random variables to construct  $\hat{P}_{h,s,a}^o$ . Besides, we denote  $Y_i = (\eta - V(X_i))$ . Note that  $Y$  is upper bounded by  $Y_{\max} := C_\rho H / (C_\rho - 1)$ . To ease notation, denote  $g(\eta; P) = C_\rho (\mathbb{E}_P [Y^2])^{1/2}$ .

Lemma 4 implies that  $g(\eta; \hat{P}_{h,s,a}^o) = C_\rho (\frac{1}{N} \sum_{i=1}^N Y_i^2)^{1/2}$  is a  $\frac{C_\rho}{\sqrt{N}}$ -Lipschitz function of the vector  $(Y_1, \dots, Y_N)^T$  with respect to  $\|\cdot\|_2$ . Applying Lemma 3, we get

$$\left| g(\eta; \hat{P}_{h,s,a}^o) - \mathbb{E}_{P_{h,s,a}^o} [g(\eta; \hat{P}_{h,s,a}^o)] \right| \leq \frac{\sqrt{2}C_\rho^2 H}{(C_\rho - 1)\sqrt{N}} \sqrt{\log(2/\delta)}, \quad (27)$$

with probability at least  $1 - \delta$ . It remains to see that  $\mathbb{E}_{P_{h,s,a}^o} [g(\eta; \hat{P}_{h,s,a}^o)]$  and  $g(\eta; P_{h,s,a}^o)$  are close. Note that the random variable  $Y$  satisfies  $\mathbb{E}[|Y|^4] \leq Y_{\max}^2 \mathbb{E}[|Y|^2]$ . Since  $\mathbb{E}_{P_{h,s,a}^o} [(\frac{1}{N} \sum_{i=1}^N Y_i^2)^{1/2}] \leq (\mathbb{E}[Y^2])^{1/2}$  by Jensen's inequality, Lemma 5 implies that

$$\left| \mathbb{E}_{P_{h,s,a}^o} [g(\eta; \hat{P}_{h,s,a}^o)] - g(\eta; P_{h,s,a}^o) \right| \leq \frac{\sqrt{Y_{\max}}}{\sqrt{N}} = \sqrt{\frac{C_\rho H}{C_\rho - 1}} \frac{1}{\sqrt{N}} \leq \frac{C_\rho^2 \sqrt{H}}{(C_\rho - 1)\sqrt{N}}. \quad (28)$$

Combining Eq. (27) and Eq. (28), we get

$$\left| C_\rho \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} - C_\rho \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} \right| \leq \frac{\sqrt{2}C_\rho^2 H}{(C_\rho - 1)\sqrt{N}} \left( \sqrt{\log(2/\delta)} + 1 \right),$$

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

**Lemma 12.** Fix any value function  $V \in \mathcal{V}$  and  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\theta, \delta \in (0, 1)$  and  $\rho > 0$ , we have, with probability at least  $1 - \delta$ ,

$$\left| L_{\mathcal{P}_{h,s,a}} V - L_{\hat{\mathcal{P}}_{h,s,a}} V \right| \leq \frac{\sqrt{2}C_\rho^2 H}{(C_\rho - 1)\sqrt{N}} \left( \sqrt{\log\left(\frac{2(1 + C_\rho H / (\theta(C_\rho - 1)))}{\delta}\right)} + 1 \right) + 2\theta,$$

where  $C_\rho = \sqrt{1 + \rho}$  and  $N$  is the number of samples used to approximate  $P_{h,s,a}^o$ .

*Proof.* From Proposition 3, we get

$$\begin{aligned} L_{\mathcal{P}_{h,s,a}} V &= - \inf_{\eta \in [0, C_\rho H / (C_\rho - 1)]} \left\{ C_\rho \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - \eta \right\}, \\ L_{\hat{\mathcal{P}}_{h,s,a}} V &= - \inf_{\eta \in [0, C_\rho H / (C_\rho - 1)]} \left\{ C_\rho \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - \eta \right\}. \end{aligned}$$

Now for any  $\rho > 0$ , we have

$$\begin{aligned} \left| L_{\mathcal{P}_{h,s,a}} V - L_{\hat{\mathcal{P}}_{h,s,a}} V \right| &= \left| \inf_{\eta \in [0, C_\rho H / (C_\rho - 1)]} \left\{ C_\rho \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - \eta \right\} \right. \\ &\quad \left. - \inf_{\eta \in [0, C_\rho H / (C_\rho - 1)]} \left\{ C_\rho \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - \eta \right\} \right| \\ &\stackrel{(a)}{\leq} \sup_{\eta \in [0, C_\rho H / (C_\rho - 1)]} \left| C_\rho \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} - C_\rho \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\eta - V(s'))^2]} \right| \\ &\stackrel{(b)}{\leq} \max_{\nu \in \mathcal{N}_\rho(\theta)} \left| C_\rho \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} - C_\rho \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} \right| + 2\theta, \quad (29) \end{aligned}$$

where (a) is from  $|\inf_x f(x) - \inf_x g(x)| \leq \sup_x |f(x) - g(x)|$ . (b) is due to Lemma 10.

Now using Lemma 11 and applying a union bound over  $\nu \in \mathcal{N}_\rho(\theta)$ , we get the following inequalities with probability at least  $1 - \delta$

$$\begin{aligned} \max_{\nu \in \mathcal{N}_\rho(\theta)} & \left| C_\rho \sqrt{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} - C_\rho \sqrt{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [(\nu - V(s'))^2]} \right| \\ & \leq \frac{\sqrt{2}C_\rho^2 H}{(C_\rho - 1)\sqrt{N}} \left( \sqrt{\log \left( \frac{2|\mathcal{N}_\rho(\theta)|}{\delta} \right)} + 1 \right) \\ & \stackrel{(a)}{\leq} \frac{\sqrt{2}C_\rho^2 H}{(C_\rho - 1)\sqrt{N}} \left( \sqrt{\log \left( \frac{2(1 + C_\rho H/(\theta(C_\rho - 1)))}{\delta} \right)} + 1 \right), \end{aligned} \quad (30)$$

where (a) follows from Eq. (26). Combining Eq. (29) and Eq. (30), we get the desired inequality.  $\square$

**Proof of Proposition 4.** Recall that  $\hat{V}_{h+1}$  is independent of  $\hat{P}_{h,s,a}^o$  by construction. Similar to Corollary 1 and Proposition 2, the result directly follows from Lemma 12 and the law of total probability.  $\square$

We now have all the ingredients to prove our main result.

**Proof of Theorem 2.** The proof is almost identical to that of Theorem 1. By applying Proposition 4 and taking a union bound over  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$\|V_1^* - V_1^{\hat{\pi}}\|_\infty \leq \frac{2\sqrt{2}C_\rho^2 H^2}{(C_\rho - 1)\sqrt{N}} \left( \sqrt{\log \left( \frac{2H|\mathcal{S}||\mathcal{A}|(1 + \frac{C_\rho H}{\theta(C_\rho - 1)})}{\delta} \right)} + 1 \right) + 4H\theta,$$

with probability at least  $1 - \delta$ . We can choose  $\theta = \epsilon/(8H)$ . Note that since  $\epsilon \in (0, 8H)$ , this particular  $\theta$  is in  $(0, 1)$ . Now, if we choose

$$N \geq \frac{32C_\rho^4 H^4}{(C_\rho - 1)^2} \frac{1}{\epsilon^2} \left( \sqrt{\log \left( \frac{2H|\mathcal{S}||\mathcal{A}|(1 + \frac{8C_\rho H^2}{\epsilon(C_\rho - 1)})}{\delta} \right)} + 1 \right)^2$$

we get  $\|V_1^* - V_1^{\hat{\pi}}\|_\infty \leq \epsilon$  with probability at least  $1 - \delta$ . To simplify the above, consider the elementary inequality  $(a + b)^2 \leq 2(a^2 + b^2)$ , for any non-negative real number  $a$  and  $b$ . With this, we get our final sample complexity result:

$$H|\mathcal{S}||\mathcal{A}|N \geq \frac{64C_\rho^4}{(C_\rho - 1)^2} \frac{H^5|\mathcal{S}||\mathcal{A}|}{\epsilon^2} \left( \log \left( \frac{2H|\mathcal{S}||\mathcal{A}|(1 + \frac{8C_\rho H^2}{\epsilon(C_\rho - 1)})}{\delta} \right) + 1 \right).$$

$\square$

### B.3 Proof of Theorem 3

Here we provide the proofs of supporting lemmas and that of Theorem 3.

**Lemma 13.** Let  $D_f$  be defined as in Eq. (14) with the convex function  $f(t) = t \log(t)$  corresponding to the KL uncertainty set  $\mathcal{P}^{KL}$ . Then

$$D_f(P \| P^o)_{\leq \rho} \inf_{\lambda \geq 0} \left\{ \lambda \rho + \lambda \log \left( \mathbb{E}_{P^o} \left[ \exp \left( \frac{-l(X)}{\lambda} \right) \right] \right) \right\}.$$

*Proof.* We first derive the Fenchel conjugate of  $f$ . We have  $f^*(s) = \sup_{t \geq 0} \{st - t \log(t)\}$ . From calculus, the optimal  $t$  for this optimization problem is  $\exp(s - 1)$ . Plugging this back into the conjugate, we get  $f^*(s) = \exp(s - 1)$ . From

Proposition 10, we get

$$\begin{aligned} \sup_{D_f(P\|P^o) \leq \rho} \mathbb{E}_P[l(X)] &= \inf_{\lambda > 0, \eta \in \mathbb{R}} \mathbb{E}_{P^o} \left[ \lambda f^* \left( \frac{l(X) - \eta}{\lambda} \right) \right] + \lambda \rho + \eta \\ &= \inf_{\lambda \geq 0, \eta \in \mathbb{R}} \left\{ \mathbb{E}_{P^o} \left[ \lambda \exp \left( \frac{l(X) - \eta}{\lambda} - 1 \right) \right] + \lambda \rho + \eta \right\} \\ &= \inf_{\lambda \geq 0} \left\{ \lambda \rho + \lambda \log \left( \mathbb{E}_{P^o} \left[ \exp \left( \frac{l(X)}{\lambda} \right) \right] \right) \right\}, \end{aligned}$$

where the last equality is from infimizing over  $\eta$ . That is, the optimal  $\eta = \lambda \log \left( \mathbb{E}_{P^o} \left[ \exp \left( \frac{l(X)}{\lambda} - 1 \right) \right] \right)$ . Now this implies that

$$\begin{aligned} \inf_{D_f(P\|P^o) \leq \rho} \mathbb{E}_P[l(X)] &= - \sup_{D_f(P\|P^o) \leq \rho} \mathbb{E}_P[-l(X)] \\ &= - \inf_{\lambda \geq 0} \left\{ \lambda \rho + \lambda \log \left( \mathbb{E}_{P^o} \left[ \exp \left( \frac{-l(X)}{\lambda} \right) \right] \right) \right\}. \end{aligned}$$

This completes the proof.  $\square$

**Proof of Proposition 5.** Applying Lemma 13 to value function, we get

$$L_{\mathcal{P}_{h,s,a}} V = - \inf_{\lambda \geq 0} \left\{ \lambda \rho + \lambda \log \left( \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right) \right\}.$$

We denote  $h(\lambda) = \lambda \rho + \lambda \log \left( \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right)$ . Note that  $h$  is convex in the dual variable  $\lambda$ . Now fix any  $\lambda \geq H/\rho$ . We have

$$\lambda \rho + \lambda \log \left( \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right) \geq \lambda \rho + \lambda \log \left( \exp \left( \frac{-H}{\lambda} \right) \right) \geq 0.$$

In addition, observe that  $h$  is monotonically increasing in  $\lambda$  when  $\lambda \geq H/\rho$ . Thus, it is sufficient to optimize  $\lambda$  in  $[0, H/\rho]$ . This implies that

$$L_{\mathcal{P}_{h,s,a}} V = - \inf_{\lambda \in [0, H/\rho]} \left\{ \lambda \rho + \lambda \log \left( \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right) \right\}.$$

$\square$

**Lemma 14.** Fix any value function  $V \in \mathcal{V}$  and  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\theta, \delta \in (0, 1)$  and  $\rho > 0$ , we have, with probability at least  $1 - \delta$ ,

$$\left| L_{\mathcal{P}_{h,s,a}} V - L_{\hat{\mathcal{P}}_{h,s,a}} V \right| \leq \frac{H}{\rho} \exp(H/\underline{\lambda}) \exp(\theta H) \sqrt{\frac{\log(4/(\theta \underline{\lambda} \delta))}{2N}},$$

where  $N$  is the number of samples used to approximate  $P_{h,s,a}^o$ , and  $\underline{\lambda}$  is a problem dependent parameter but independent of  $N$ .

*Proof.* Fix any value function  $V \in \mathcal{V}$ ,  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ , and any  $\rho > 0$ . From Proposition 5, we have

$$\begin{aligned} L_{\mathcal{P}_{h,s,a}} V &= - \inf_{\lambda \in [0, H/\rho]} \left\{ \lambda \rho + \lambda \log \left( \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right) \right\}, \\ L_{\hat{\mathcal{P}}_{h,s,a}} V &= - \inf_{\lambda \in [0, H/\rho]} \left\{ \lambda \rho + \lambda \log \left( \mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right) \right\}. \end{aligned}$$

Suppose the optimization problem  $L_{\mathcal{P}_{h,s,a}} V$  above is achieved at  $\lambda^* = 0$ . It can be shown that when  $\lambda^* = 0$ ,  $L_{\mathcal{P}_{h,s,a}} V = L_{\hat{\mathcal{P}}_{h,s,a}} V = V_{\min}$  with high probability (and hence  $|L_{\mathcal{P}_{h,s,a}} V - L_{\hat{\mathcal{P}}_{h,s,a}} V| = 0$ ), where  $V_{\min} = \min_{s \in \mathcal{S}} V(s)$ , whenever  $N$  is greater than some problem dependent constant but independent of the optimality gap  $\epsilon$  defined in Theorem 3. (Nilim and El Ghaoui, 2005, Appendix C) provided a proof of this trivial case in terms of robust optimal control. (Panaganti and

Kalathil, 2022, Lemma 13) adapted the proof to RMDP and provided a detailed analysis. We omit the proof here and only discuss the case where  $\lambda^* \in (0, H/\rho]$ .

Now assume the optimal  $\lambda^*$  is achieved in  $(0, H/\rho]$ . Define  $\underline{\lambda} = \lambda^*/2$  if  $\lambda^* \in (0, 1)$  and  $\underline{\lambda} = 1/2$  if  $\lambda^* \geq 1$ . We define  $\underline{\lambda}$  this way because restricting  $\underline{\lambda} \in (0, 1)$  can later give us cleaner expression without compromising its function. Let the optimization problem  $L_{\hat{P}_{h,s,a}} V$  be achieved at  $\hat{\lambda}^*$ . Again from (Zhou et al., 2021, Lemma 4) (Panaganti and Kalathil, 2022, Lemma 13), it holds that  $\hat{\lambda}^* \in (\underline{\lambda}, H/\rho]$  whenever  $N$  is greater than some problem dependent constant but independent of the optimality gap  $\epsilon$  defined in Theorem 3, and hence omit this constant for  $N$  further in our analysis. Now it follows that

$$\begin{aligned}
 \left| L_{\mathcal{P}_{h,s,a}} V - L_{\hat{P}_{h,s,a}} V \right| &= \left| \inf_{\lambda \in (\underline{\lambda}, H/\rho]} \left\{ \lambda \rho + \lambda \log \left( \mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right) \right\} \right. \\
 &\quad \left. - \inf_{\lambda \in (\underline{\lambda}, H/\rho]} \left\{ \lambda \rho + \lambda \log \left( \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right) \right\} \right| \\
 &\stackrel{(a)}{\leq} \sup_{\lambda \in (\underline{\lambda}, H/\rho]} \left| \lambda \log \left( \frac{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right]}{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right]} \right) \right| \\
 &\leq \frac{H}{\rho} \sup_{\lambda \in (\underline{\lambda}, H/\rho]} \left| \log \left( \frac{\mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right]}{\mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right]} + 1 \right) \right| \\
 &\stackrel{(b)}{\leq} \frac{H}{\rho} \sup_{\lambda \in (\underline{\lambda}, H/\rho]} \frac{\left| \mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right|}{\left| \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp \left( \frac{-V(s')}{\lambda} \right) \right] \right|} \\
 &\stackrel{(c)}{\leq} \frac{H}{\rho} \sup_{\lambda' \in [\rho/H, 1/\underline{\lambda}]} \frac{\left| \mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} \left[ \exp(-\lambda' V(s')) \right] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp(-\lambda' V(s')) \right] \right|}{\left| \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp(-\lambda' V(s')) \right] \right|} \tag{31} \\
 &\stackrel{(d)}{\leq} \frac{H}{\rho} \exp(H/\underline{\lambda}) \sup_{\lambda' \in [\rho/H, 1/\underline{\lambda}]} \left| \mathbb{E}_{s' \sim \hat{P}_h^o(\cdot|s,a)} \left[ \exp(-\lambda' V(s')) \right] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \exp(-\lambda' V(s')) \right] \right| \\
 &= \frac{H}{\rho} \exp(H/\underline{\lambda}) \sup_{\lambda' \in [\rho/H, 1/\underline{\lambda}]} \left| \hat{P}_{h,s,a}^o \exp(-\lambda' V) - P_{h,s,a}^o \exp(-\lambda' V) \right|, \tag{32}
 \end{aligned}$$

where (a) follows from the fact that  $|\inf_x f(x) - \inf_x g(x)| \leq \sup_x |f(x) - g(x)|$ . (b) is due to  $|\log(1+x)| \leq |x|$ . (c) is from a substitution of  $\lambda' = 1/\lambda$ . (d) is because  $\exp(-\lambda' V(s'))$  is lower bounded by  $\exp(-H/\underline{\lambda})$ , for all  $s' \in \mathcal{S}$ . In the last equality,  $\exp(\cdot)$  denotes the element-wise exponential function.

Now let  $\mathcal{N}_\rho(\theta)$  be a  $\theta$ -cover of the interval  $[\rho/H, 1/\underline{\lambda}]$  with  $\theta \in (0, 1)$ . Recall that we defined  $\underline{\lambda} \in (0, 1)$  above. We have  $|\mathcal{N}_\rho(\theta)| \leq (1/\underline{\lambda} - \rho/H)/\theta + 1 \leq 2/(\theta\underline{\lambda})$ . Now fix any  $\lambda \in [\rho/H, 1/\underline{\lambda}]$ , there exists a  $\nu \in \mathcal{N}_\rho(\theta)$  such that  $|\lambda - \nu| \leq \theta$ . Now for these particular  $\lambda$  and  $\nu$ , we have

$$\begin{aligned}
 \left| \hat{P}_{h,s,a}^o \exp(-\lambda V) - P_{h,s,a}^o \exp(-\lambda V) \right| &= \left| \sum_{s' \in \mathcal{S}} \left( \hat{P}_h^o(s' | s, a) - P_h^o(s' | s, a) \right) \exp(-\nu V(s')) \exp((\nu - \lambda)V(s')) \right| \\
 &\stackrel{(a)}{\leq} \left| \sum_{s' \in \mathcal{S}} \left( \hat{P}_h^o(s' | s, a) - P_h^o(s' | s, a) \right) \exp(-\nu V(s')) \exp(\theta V_{\max}) \right| \\
 &= \exp(\theta H) \left| \hat{P}_{h,s,a}^o \exp(-\nu V) - P_{h,s,a}^o \exp(-\nu V) \right| \\
 &\leq \exp(\theta H) \max_{\nu \in \mathcal{N}_\rho(\theta)} \left| \hat{P}_{h,s,a}^o \exp(-\nu V) - P_{h,s,a}^o \exp(-\nu V) \right|,
 \end{aligned}$$

where (a) is due to  $V_{\max} = H$  and  $|\nu - \lambda| \leq \theta$ . Taking maximum over  $\lambda \in [\rho/H, 1/\underline{\lambda}]$  on both sides, we get

$$\sup_{\lambda \in [\rho/H, 1/\underline{\lambda}]} \left| \hat{P}_{h,s,a}^o \exp(-\lambda V) - P_{h,s,a}^o \exp(-\lambda V) \right| \leq \max_{\nu \in \mathcal{N}_\rho(\theta)} \left| \hat{P}_{h,s,a}^o \exp(-\nu V) - P_{h,s,a}^o \exp(-\nu V) \right| \exp(\theta H). \tag{33}$$

Observe that for all  $\nu \in \mathcal{N}_\rho(\theta)$  and  $s \in \mathcal{S}$ ,  $\exp(-\nu V(s)) \leq \exp(-\nu 0) \leq 1$ . That is,  $\|\exp(-\nu V)\|_\infty \leq 1$ . Since  $V$  is independent of  $\widehat{P}_{h,s,a}^o$ , we can apply Hoeffding's inequality (Lemma 2):

$$\mathbb{P}\left(\left|\widehat{P}_{h,s,a}^o \exp(-\nu V) - P_{h,s,a}^o \exp(-\nu V)\right| \geq \epsilon\right) \leq 2\exp\left(-\frac{2N\epsilon^2}{\|\exp(-\nu V)\|_\infty^2}\right) \leq 2\exp(-2N\epsilon^2), \quad \forall \epsilon > 0.$$

Now choose  $\epsilon = \sqrt{\frac{\log(2|\mathcal{N}_\rho(\theta)|/\delta)}{2N}}$ . We have

$$\begin{aligned} \mathbb{P}\left(\left|\widehat{P}_{h,s,a}^o \exp(-\nu V) - P_{h,s,a}^o \exp(-\nu V)\right| \geq \sqrt{\frac{\log(4/(\theta\lambda\delta))}{2N}}\right) \\ \leq \mathbb{P}\left(\left|\widehat{P}_{h,s,a}^o \exp(-\nu V) - P_{h,s,a}^o \exp(-\nu V)\right| \geq \sqrt{\frac{\log(2|\mathcal{N}_\rho(\theta)|/\delta)}{2N}}\right) \leq \frac{\delta}{|\mathcal{N}_\rho(\theta)|}. \end{aligned}$$

Applying a union bound over  $\nu \in \mathcal{N}_{\rho,V}(\theta)$ , we get

$$\max_{\nu \in \mathcal{N}_{\rho,V}(\theta)} \left| \widehat{P}_{h,s,a}^o \exp(-\nu V) - P_{h,s,a}^o \exp(-\nu V) \right| \leq \sqrt{\frac{\log(4/(\theta\lambda\delta))}{2N}}, \quad (34)$$

with probability at least  $1 - \delta$ . Combining Eq. (32)-Eq. (34) completes the proof.  $\square$

**Proof of Proposition 6.** Recall that  $\widehat{V}_{h+1}$  is independent of  $\widehat{P}_{h,s,a}^o$  by construction. Similar to Corollary 1 and Proposition 2, the result directly follows from Lemma 14 and the law of total probability.  $\square$

We now have all the ingredients to prove our main result.

**Proof of Theorem 3.** The proof is almost identical to that of Theorem 1. By applying Proposition 6 and taking a union bound over  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$\begin{aligned} \left\|V_1^* - V_1^{\widehat{\pi}}\right\|_\infty &\leq \frac{2H^2}{\rho} \exp(H/\underline{\lambda}) \exp(\theta H) \sqrt{\frac{\log(4H|\mathcal{S}||\mathcal{A}|/(\theta\lambda\delta))}{2N}} \\ &\leq \frac{2H^2}{\rho} \exp(H/\underline{\lambda} + H/2) \sqrt{\frac{\log(8H|\mathcal{S}||\mathcal{A}|/(\lambda\delta))}{2N}}, \end{aligned}$$

with probability at least  $1 - \delta$ , and the last inequality is because we can simply choose  $\theta = 1/2$ . Now if we choose

$$N \geq \frac{2\exp(3H/\underline{\lambda})}{\rho^2} \frac{H^4 \log(8H|\mathcal{S}||\mathcal{A}|/(\lambda\delta))}{\epsilon^2},$$

we get  $\|V_1^* - V_1^{\widehat{\pi}}\|_\infty \leq \epsilon$  with probability at least  $1 - \delta$ .  $\square$

#### B.4 Proof of Theorem 4

Here we provide the proofs of supporting lemmas and that of Theorem 4.

**Lemma 15.** Fix any value function  $V \in \mathcal{V}$  and  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\theta, \delta \in (0, 1)$  and  $\rho > 0$ , we have, with probability at least  $1 - \delta$ ,

$$\left|L_{\mathcal{P}_{h,s,a}} V - L_{\widehat{\mathcal{P}}_{h,s,a}} V\right| \leq \frac{H}{\rho} \sqrt{\frac{\log(2|\mathcal{S}|/\delta)}{2Np^2}},$$

where  $\underline{p} = \min_{s',s,a,h: P_h^o(s'|s,a) > 0} P_h^o(s'|s,a) > 0$  and  $N$  is the number of samples used to approximate  $P_{h,s,a}^o$ .

*Proof.* From Eq. (31), we get

$$\begin{aligned}
 \left| L_{\mathcal{P}_{h,s,a}} V - L_{\widehat{\mathcal{P}}_{h,s,a}} V \right| &\leq \frac{H}{\rho} \sup_{\lambda \in [\rho/H, 1/\Delta]} \frac{\left| \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} [\exp(-\lambda V(s'))] - \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [\exp(-\lambda V(s'))] \right|}{\left| \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [\exp(-\lambda V(s'))] \right|} \\
 &\leq \frac{H}{\rho} \sup_{\lambda \in [\rho/H, 1/\Delta]} \frac{\sum_{s' \in \mathcal{S}} \left| \widehat{P}_h^o(s' | s, a) - P_h^o(s' | s, a) \right| \exp(-\lambda V(s'))}{\sum_{s' \in \mathcal{S}} P_h^o(s' | s, a) \exp(-\lambda V(s'))} \\
 &\leq \frac{H}{\rho} \sup_{\lambda \in [\rho/H, 1/\Delta]} \frac{\sum_{s' \in \mathcal{S}} \left| \widehat{P}_h^o(s' | s, a) - P_h^o(s' | s, a) \right| \exp(-\lambda V(s'))}{\sum_{s' \in \mathcal{S}} P_h^o(s' | s, a) \exp(-\lambda V(s'))} \\
 &\stackrel{(a)}{\leq} \frac{H}{\rho} \max_{s': P_h^o(s'|s,a) > 0} \left| \frac{\widehat{P}_h^o(s' | s, a)}{P_h^o(s' | s, a)} - 1 \right|, \tag{35}
 \end{aligned}$$

where (a) follows from  $\sum_i a_i / \sum_i b_i \leq \max_i \{a_i / b_i\}$ , when  $b_i > 0$ . We denote  $\underline{p} = \min_{s', s, a, h: P_h^o(s'|s,a) > 0} P_h^o(s' | s, a)$ . Note that  $\underline{p}$  is a problem dependent constant and not dependent on  $(h, s, a)$ . Now fix any  $s' \in \mathcal{S}$  such that  $P_h^o(s' | s, a) > 0$ . Hoeffding's inequality tells us

$$\mathbb{P} \left( \left| \frac{\widehat{P}_h^o(s' | s, a)}{P_h^o(s' | s, a)} - 1 \right| \geq \epsilon \right) \leq 2 \exp \left( -\frac{2N\epsilon^2}{(1/\underline{p})^2} \right) = 2 \exp(-2N\underline{p}^2\epsilon^2).$$

Choosing  $\epsilon = \sqrt{\frac{\log(2|\mathcal{S}|/\delta)}{2N\underline{p}^2}}$  and applying a union bound over  $s' \in \mathcal{S}$ , we get

$$\max_{s': P_h^o(s'|s,a) > 0} \left| \frac{\widehat{P}_h^o(s' | s, a)}{P_h^o(s' | s, a)} - 1 \right| \leq \sqrt{\frac{\log(2|\mathcal{S}|/\delta)}{2N\underline{p}^2}},$$

with probability at least  $1 - \delta$ . Combining the above and Eq. (35), we get the desired result.  $\square$

**Proof of Proposition 7.** Recall that  $\widehat{V}_{h+1}$  is independent of  $\widehat{P}_{h,s,a}^o$  by construction. Similar to Corollary 1 and Proposition 2, the result directly follows from Lemma 15 and the law of total probability.  $\square$

We now have all the ingredients to prove our main result.

**Proof of Theorem 4.** The proof is almost identical to that of Theorem 1. By applying Proposition 7 and taking a union bound over  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$\left\| V_1^* - V_1^{\widehat{\pi}} \right\|_{\infty} \leq \frac{2H^2}{\rho} \sqrt{\frac{\log(2H|\mathcal{S}|^2|\mathcal{A}|/\delta)}{2N\underline{p}^2}},$$

with probability at least  $1 - \delta$ . Now if we choose

$$N \geq \frac{2H^4 \log(2H|\mathcal{S}|^2|\mathcal{A}|/\delta)}{\rho^2 \underline{p}^2 \epsilon},$$

then we get  $\|V^* - V^{\widehat{\pi}}\|_{\infty} \leq \epsilon$  with probability at least  $1 - \delta$ .  $\square$

## B.5 Proof of Theorem 5

Here we provide the proofs of supporting lemmas and that of Theorem 5.

**Proposition 11** (Gao and Kleywegt, 2022, Theorem 1). *Let  $P^o$  be a distribution on the bounded space  $\mathcal{X}$  and let  $l: \mathcal{X} \rightarrow \mathbb{R}$  be a bounded loss function. Then,*

$$\sup_{D_W(P \| P^o) \leq \rho} \mathbb{E}_P[l(X)] = \inf_{\lambda \geq 0} \mathbb{E}_{Y \sim P^o} \left[ \sup_{x \in \mathcal{X}} \{l(x) - \lambda d^p(x, Y)\} \right] + \lambda \rho^p.$$

**Proof of Proposition 8.** Fix any  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . We have

$$\begin{aligned}
 L_{\mathcal{P}_{h,s,a}} V &= \inf_{D_W(P \parallel P^o) \leq \rho} \mathbb{E}_{s \sim P} [V(s)] = - \sup_{D_W(P \parallel P^o) \leq \rho} \mathbb{E}_{s \sim P} [-V(s)] \\
 &\stackrel{(a)}{=} - \inf_{\lambda \geq 0} \left\{ \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \sup_{s'' \in \mathcal{S}} \{-V(s'') - \lambda d^p(s'', s')\} \right] + \lambda \rho^p \right\} \\
 &= \sup_{\lambda \geq 0} \left\{ \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] - \lambda \rho^p \right\} \\
 &\stackrel{(b)}{=} \sup_{\lambda \in [0, H/\rho^p]} \left\{ \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] - \lambda \rho^p \right\},
 \end{aligned}$$

where (a) follows from Proposition 11. For (b), let us first denote any optimizer in (a) to be  $\lambda^*$ . Observe that since  $V$  is non-negative, it follows that  $L_{\mathcal{P}_{h,s,a}} V$  is also non-negative. Now due to  $V_{\max} = H$ , we have that

$$0 \leq -\lambda^* \rho^p + \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \inf_{s''} \{V(s'') + \lambda d^p(s'', s')\} \right] \leq -\lambda^* \rho^p + \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} [V(s') + \lambda d^p(s', s')] \leq -\lambda^* \rho^p + H,$$

where in the last inequality we use that the distance metric satisfies  $d(s, s) = 0$ , for any  $s \in \mathcal{S}$ .  $\square$

**Lemma 16** (Covering number (Wasserstein)). *Fix any value function  $V \in \mathcal{V}$ , consider the following set of  $\mathbb{R}^{|\mathcal{S}|}$  vectors:*

$$\mathcal{U}_{\rho,V} = \left\{ \left( \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', 1)\}, \dots, \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', |\mathcal{S}|)\} \right)^T : \lambda \in [0, H/\rho^p] \right\}.$$

Let

$$\mathcal{N}_{\rho,V}(\theta) = \left\{ \left( \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', 1)\}, \dots, \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', |\mathcal{S}|)\} \right)^T : \lambda \in \left\{ \frac{\theta}{B_p}, \frac{2\theta}{B_p}, \dots, N_{\rho,\theta} \frac{\theta}{B_p} \right\} \right\},$$

where  $N_{\rho,\theta} = \left\lceil \frac{HB_p}{\rho^p \theta} \right\rceil$  and  $B_p = \max_{s', s''} d^p(s'', s')$ . Then  $\mathcal{N}_{\rho,V}(\theta)$  is a  $\theta$ -cover of  $\mathcal{U}_{\rho,V}$  with respect to  $\|\cdot\|_\infty$ , and its cardinality is bounded as  $|\mathcal{N}_{\rho,V}(\theta)| \leq \frac{HB_p + (H \vee \rho^p)}{\rho^p \theta}$ . Furthermore, for any  $\nu \in \mathcal{N}_{\rho,V}(\theta)$ , we have  $\|\nu\|_\infty \leq \frac{H(B_p + \rho^p)}{\rho^p}$ .

*Proof.* Fix any  $\theta \in (0, 1)$ . First note that  $N_{\rho,\theta}$  is the minimal number of subintervals of length  $\frac{\theta}{B_p}$  needed to cover  $[0, H/\rho^p]$ . Denote  $J_i = [(i-1)\frac{\theta}{B_p}, i\frac{\theta}{B_p}]$ ,  $1 \leq i \leq N_{\rho,\theta}$ . Fix some  $\mu \in \mathcal{U}_{\rho,V}$ . Then  $\mu$  must takes the form

$$\mu = \left( \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', 1)\}, \dots, \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', |\mathcal{S}|)\} \right)^T,$$

for some  $\lambda \in [0, H/\rho^p]$ . Without loss of generality, assume  $\lambda \in J_i$ . Now we pick

$$\nu = \left( \inf_{s'' \in \mathcal{S}} \left\{ V(s'') + i \frac{\theta}{B_p} d^p(s'', 1) \right\}, \dots, \inf_{s'' \in \mathcal{S}} \left\{ V(s'') + i \frac{\theta}{B_p} d^p(s'', |\mathcal{S}|) \right\} \right)^T.$$

Fix any  $s' \in \mathcal{S}$ , we have

$$\begin{aligned}
 |\mu(s') - \nu(s')| &= \left| \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} - \inf_{s'' \in \mathcal{S}} \left\{ V(s'') + i \frac{\theta}{B_p} d^p(s'', s) \right\} \right| \\
 &\stackrel{(a)}{\leq} \sup_{s'' \in \mathcal{S}} \left| \left( \lambda - i \frac{\theta}{B_p} \right) d^p(s'', s') \right| \\
 &\leq \left| \lambda - i \frac{\theta}{B_p} \right| \max_{s', s''} d^p(s'', s') = \left| \lambda - i \frac{\theta}{B_p} \right| B_p \\
 &\leq \left| (i-1) \frac{\theta}{B_p} - i \frac{\theta}{B_p} \right| B_p = \theta,
 \end{aligned}$$



where (a) is due to  $|\inf_x f(x) - \inf_x g(x)| \leq \sup_x |f(x) - g(x)|$ . Taking maximum over  $s' \in \mathcal{S}$  on both sides, we get  $\|\mu - \nu\|_\infty \leq \theta$ . Since  $\nu \in \mathcal{N}_{\rho, V}(\theta)$ , this suggests that  $\mathcal{N}_{\rho, V}(\theta)$  is a  $\theta$ -cover for  $\mathcal{U}_{\rho, V}$ .

To bound the cardinality of  $\mathcal{N}_{\rho, V}(\theta)$ , we consider two cases. If  $0 < \rho < 1$ , then  $\rho^p \theta < 1$  and

$$\left\lceil \frac{HB_p}{\rho^p \theta} \right\rceil \leq \frac{HB_p}{\rho^p \theta} + 1 \leq \frac{HB_p}{\rho^p \theta} + \frac{H}{\rho^p \theta} = \frac{HB_p + H}{\rho^p \theta}.$$

On the other hand, if  $\rho > 1$ , then since  $\theta \in (0, 1)$ , we have

$$\left\lceil \frac{HB_p}{\rho^p \theta} \right\rceil \leq \frac{HB_p}{\rho^p \theta} + 1 = \frac{HB_p}{\rho^p \theta} + \frac{\rho^p \theta}{\rho^p \theta} \leq \frac{HB_p}{\rho^p \theta} + \frac{\rho^p}{\rho^p \theta} = \frac{HB_p + \rho^p}{\rho^p \theta}.$$

Hence, we have  $|\mathcal{N}_{\rho, V}(\theta)| = N_{\rho, \theta} \leq \frac{HB_p + (H \vee \rho^p)}{\rho^p \theta}$ . Now we prove the last claim. Fix any  $\nu \in \mathcal{N}_{\rho, V}$ . Note that for any  $s' \in \mathcal{S}$ ,

$$\nu(s') = \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \leq H + \lambda B_p \leq H + \frac{H}{\rho^p} B_p = \frac{H(B_p + \rho^p)}{\rho^p}.$$

The result then follows from taking maximum over  $s' \in \mathcal{S}$  on both sides.  $\square$

**Lemma 17.** Fix any  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . Fix any value function  $V \in \mathcal{V}$ . Let  $\mathcal{N}_{\rho, V}(\theta)$  be the  $\theta$ -cover of the set

$$\mathcal{U}_{\rho, V} = \left\{ \left( \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', 1)\}, \dots, \{V(s'') + \lambda d^p(s'', |\mathcal{S}|)\} \right)^T : \lambda \in [0, H/\rho^p] \right\},$$

as described in Lemma 16. We then have

$$\begin{aligned} \sup_{\lambda \in [0, H/\rho^p]} \left| \mathbb{E}_{s' \sim P_h^\circ(\cdot | s, a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] - \mathbb{E}_{s' \sim \widehat{P}_h^\circ(\cdot | s, a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] \right| \\ \leq \max_{\nu \in \mathcal{N}_{\rho, V}(\theta)} \left| \widehat{P}_{h, s, a}^\circ \nu - P_{h, s, a}^\circ \nu \right| + 2\theta. \end{aligned}$$

*Proof.* The proof is identical to the proof of Lemma 7.  $\square$

**Lemma 18.** Fix any value function  $V \in \mathcal{V}$  and  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . For any  $\theta, \delta \in (0, 1)$  and  $\rho > 0$ , we have the following inequality with probability at least  $1 - \delta$

$$\left| L_{\mathcal{P}_{h, s, a}} V - L_{\widehat{\mathcal{P}}_{h, s, a}} V \right| \leq \frac{H(B_p + \rho^p)}{\rho^p} \sqrt{\frac{\log \left( \frac{2HB_p + 2(H \vee \rho^p)}{\rho^p \theta \delta} \right)}{2N}} + 2\theta,$$

where  $B_p = \max_{s', s''} d^p(s'', s')$ .

*Proof.* Fix any value function  $V \in \mathcal{V}$  independent of  $\widehat{P}_{h, s, a}^\circ$  and  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . From Proposition 8, we have

$$\begin{aligned} L_{\mathcal{P}_{h, s, a}} V &= \sup_{\lambda \in [0, H/\rho^p]} \left\{ \mathbb{E}_{s' \sim P_h^\circ(\cdot | s, a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] - \lambda \rho^p \right\}, \\ L_{\widehat{\mathcal{P}}_{h, s, a}} V &= \sup_{\lambda \in [0, H/\rho^p]} \left\{ \mathbb{E}_{s' \sim \widehat{P}_h^\circ(\cdot | s, a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] - \lambda \rho^p \right\}. \end{aligned}$$

Now it follows that

$$\begin{aligned}
 |L_{\mathcal{P}_{h,s,a}} V - L_{\widehat{\mathcal{P}}_{h,s,a}} V| &= \left| \sup_{\lambda \in [0, H/\rho^p]} \left\{ \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] - \lambda \rho^p \right\} \right. \\
 &\quad \left. - \sup_{\lambda \in [0, H/\rho^p]} \left\{ \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] - \lambda \rho^p \right\} \right| \\
 &\stackrel{(a)}{\leq} \sup_{\lambda \in [0, H/\rho^p]} \left| \mathbb{E}_{s' \sim P_h^o(\cdot|s,a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] \right. \\
 &\quad \left. - \mathbb{E}_{s' \sim \widehat{P}_h^o(\cdot|s,a)} \left[ \inf_{s'' \in \mathcal{S}} \{V(s'') + \lambda d^p(s'', s')\} \right] \right| \\
 &\stackrel{(b)}{\leq} \max_{\nu \in \mathcal{N}_{\rho, V}(\theta)} \left| \widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu \right| + 2\theta, \tag{36}
 \end{aligned}$$

where (a) follows from  $|\sup_x f(x) - \sup_x g(x)| \leq \sup_x |f(x) - g(x)|$ . (b) follows from Lemma 17.

Recall that all  $\nu \in \mathcal{N}_{\rho, V}(\theta)$  is bounded by  $\nu_{\max} := \frac{H(B_p + \rho^p)}{\rho^p}$ . Now we can apply Hoeffding's inequality (Lemma 2):

$$\mathbb{P} \left( \left| \widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu \right| \geq \epsilon \right) \leq 2 \exp \left( -\frac{2N\epsilon^2}{\nu_{\max}^2} \right) = 2 \exp \left( -\frac{2N\epsilon^2}{\left( \frac{H(B_p + \rho^p)}{\rho^p} \right)^2} \right), \quad \forall \epsilon > 0.$$

Now recall that  $|\mathcal{N}_{\rho, V}(\theta)| \leq \frac{HB_p + (H \vee \rho^p)}{\rho^p \theta}$  and choose

$$\epsilon = \frac{H(B_p + \rho^p)}{\rho^p} \sqrt{\frac{\log(2|\mathcal{N}_{\rho, V}(\theta)|/\delta)}{2N}}.$$

We then have

$$\begin{aligned}
 \mathbb{P} \left( \left| \widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu \right| \geq \frac{H(B_p + \rho^p)}{\rho^p} \sqrt{\frac{\log \left( \frac{2HB_p + 2(H \vee \rho^p)}{\rho^p \theta \delta} \right)}{2N}} \right) \\
 \leq \mathbb{P} \left( \left| \widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu \right| \geq \frac{H(B_p + \rho^p)}{\rho^p} \sqrt{\frac{\log(2|\mathcal{N}_{\rho, V}(\theta)|/\delta)}{2N}} \right) \\
 \leq \frac{\delta}{|\mathcal{N}_{\rho, V}(\theta)|}.
 \end{aligned}$$

Finally, applying a union bound over  $\mathcal{N}_{\rho, V}(\theta)$ , we get

$$\max_{\nu \in \mathcal{N}_{\rho, V}(\theta)} \left| \widehat{P}_{h,s,a}^o \nu - P_{h,s,a}^o \nu \right| \leq \frac{H(B_p + \rho^p)}{\rho^p} \sqrt{\frac{\log \left( \frac{2HB_p + 2(H \vee \rho^p)}{\rho^p \theta \delta} \right)}{2N}},$$

with probability at least  $1 - \delta$ . Combining the above and Eq. (36) completes the proof.  $\square$

**Proof of Proposition 9.** Recall that  $\widehat{V}_{h+1}$  is independent of  $\widehat{P}_{h,s,a}^o$  by construction. Similar to Corollary 1 and Proposition 2, the result directly follows from Lemma 18 and the law of total probability.  $\square$

We now have all the ingredients to prove our main result.

**Proof of Theorem 5.** The proof is almost identical to that of Theorem 1. By applying Proposition 9 and taking a union bound over  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$\left\| V_1^* - V_1^{\widehat{\pi}} \right\|_{\infty} \leq \frac{2H^2(B_p + \rho^p)}{\rho^p} \sqrt{\frac{\log \left( \frac{2H|\mathcal{S}||\mathcal{A}|(HB_p + (H \vee \rho^p))}{\rho^p \theta \delta} \right)}{2N}} + 4H\theta$$

with probability at least  $1 - \delta$ . We can choose  $\theta = \epsilon/(8H)$ . Note that since  $\epsilon \in (0, 8H)$ , this particular  $\theta$  is in  $(0, 1)$ . Now, if we choose

$$N \geq \frac{8H^4(B_p + \rho^p)^2}{\rho^{2p}\epsilon^2} \log \left( \frac{16H^2 |\mathcal{S}| |\mathcal{A}| (HB_p + (H \vee \rho^p))}{\rho^p \delta \epsilon} \right),$$

we get  $\|V_1^* - V_1^{\hat{\pi}}\|_\infty \leq \epsilon$  with probability at least  $1 - \delta$ . □