

---

# Sample Complexity of Distributionally Robust Off-Dynamics Reinforcement Learning with Online Interaction

---

**Yiting He \***

University of Science and Technology of China  
heyiting@mail.ustc.edu.cn

**Zhishuai Liu \***

Duke University  
zhishuai.liu@duke.edu

**Pan Xu**

Duke University  
pan.xu@duke.edu

## Abstract

Off-dynamics reinforcement learning (RL), where training and deployment transition dynamics are different, can be formulated as learning in a robust Markov decision process (RMDP) where uncertainties in transition dynamics are imposed. Existing literature mostly assumes access to generative models allowing arbitrary state-action queries or pre-collected datasets with a good state coverage of the deployment environment, bypassing the challenge of exploration. In this work, we study a more realistic and challenging setting where the agent is limited to online interaction with the training environment. To capture the intrinsic difficulty of exploration in online RMDPs, we introduce the supremal visitation ratio, a novel quantity that measures the mismatch between the training dynamics and the deployment dynamics. We show that if this ratio is unbounded, online learning becomes exponentially hard. We propose the first computationally efficient algorithm that achieves sublinear regret in online RMDPs with  $f$ -divergence based transition uncertainties. We also establish matching regret lower bounds, demonstrating that our algorithm achieves optimal dependence on both the supremal visitation ratio and the number of interaction episodes. Finally, we validate our theoretical results through comprehensive numerical experiments.

## 1 Introduction

Off-dynamics reinforcement learning (RL) [10, 20] has recently gained significant attention in scenarios where the transition dynamics of the deployment environment differs from that of the training environment. Such problems can be modeled as learning a robust Markov decision process (RMDP) [27, 12, 23], where the objective is to learn a policy that performs well when uncertainties are imposed into the transition dynamics. Two major frameworks have been proposed in the literature to incorporate the uncertainty of transition dynamics in RMDPs. The first, Constrained Robust Markov Decision Process (CRMDP), seeks the best policy under the worst-case transition dynamics within a predefined uncertainty set. The second, Regularized Robust Markov Decision Process (RRMDP), replaces the hard constraint on uncertainty sets with a regularization term that quantifies the divergence between the training and deployment dynamics.

---

\*Equal contribution

Existing research on learning CRMDPs can be categorized into three settings: (1) *Learning with a generative model (simulator)*. In this setting, the agent can query transitions at any state-action pair an arbitrary number of times [24, 43, 41, 30]. (2) *Learning with an offline dataset*. Here, the agent learns from a pre-collected dataset, typically assumed to be generated by a behavior policy from the nominal MDP [26, 4, 29, 17, 31]. Effective robust policy learning relies on sufficient coverage of the dataset on states in the deployment environment. (3) *Learning through online interaction*. More recent works have considered online learning of CRMDPs through direct interaction with the training environment to collect data [16, 19, 15], focusing on a specific case where the uncertainty set is defined via total variation (TV) distance<sup>2</sup>. RRMDP was introduced to get rid of the constrained optimization in the formulation of CRMDPs for the tractability of robust policy learning [42, 45]. In particular, [42] studied RRMDPs with general  $f$ -divergence-based regularization under a generative model setting, while [45] analyzed RRMDPs with Kullback-Leibler divergence-based regularization in the offline setting, relying on similar data coverage assumptions as in offline CRMDPs. Recently, [25] extended this work to general  $f$ -divergence-based regularization in the offline setting and explored RRMDPs with total variation regularization and fail-states in a hybrid online-offline setting.

Despite these advances, in more realistic applications where simulators or pre-collected datasets are not available, the problem of efficient online exploration in RMDPs remains understudied. Unlike standard MDPs, where exploration aims to reduce uncertainty within a fixed transition model, in RMDPs, the agent can only gather experience from a nominal environment, yet it must generalize to potentially shifted dynamics at deployment. This presents a fundamental challenge of *information deficit* in online RMDPs, requiring the development of exploration strategies that proactively account for distributional shifts. More specifically, the *information deficit* in online RMDPs arises when states that are rarely visited in the nominal environment become critical in the deployment environment. For instance, consider a state  $s$  in the nominal MDP that is extremely difficult to visit, e.g., with exponentially small visitation probability, resulting in limited data collection. If, in the deployment environment, the dynamics shift increases the visitation probability of  $s$ , the agent must make informed decisions at this state despite having little prior experience. In standard MDPs, such rare states typically have negligible effects on policy learning, but in RMDPs, they can critically impact performance, making online learning in RMDPs significantly more challenging.

To overcome this information deficit issue, existing research on online RMDPs adopt a fail-state type of assumption—there exist states with zero reward that only transit among themselves [16, 19, 15]. As we show in Proposition D.4 and the discussion following it, these assumptions essentially ensure that worst-case distribution shifts occur in a deterministic direction, which eliminating the information deficit issue and makes provably efficient online learning possible. However, such nice properties do not hold in RMDPs with general  $f$ -divergence based uncertainty sets or regularization, limiting all existing research on online RMDPs to CRMDPs with TV-distance based uncertainty sets. In this work, we answer the following fundamental question:

*Under what conditions can provably efficient online learning of RMDPs be achieved?*

We investigate tabular RMDPs with finite states and actions. We show that if the nominal environment is sufficiently exploratory—i.e., the agent can collect enough information through interaction—then sample-efficient online learning should be achievable for broader classes of RMDPs. We rigorously prove that the sample complexity of any online learning algorithm should be proportional to the difficulty of exploration. Our contributions are summarized as follows.

- We introduce the *supremal visitation ratio*  $C_{vr}$  (see Assumption 3.1) as a measure of exploration difficulty in RMDPs. We develop the first computationally efficient algorithm, Online Robust Bellman Iteration (ORBIT), for CRMDPs and RRMDPs based on the total variation (TV), Kullback-Leibler (KL), and  $\chi^2$  divergences, and prove regret bounds that explicitly depend on  $C_{vr}$ .
- We establish regret lower bounds, demonstrating that the supremal visitation ratio  $C_{vr}$  is an unavoidable term in the sample complexity of online RMDP learning. This result confirms that  $C_{vr}$  serves as a fundamental measure of exploration difficulty and a sufficient condition for provably

---

<sup>2</sup>We note that [9] also studied online CRMDPs, but we found essential flaws in proofs of their Lemmas A.2 and C.5, which invalidates their results.

efficient online learning in RMDPs. As a corollary, we construct hard instances to demonstrate that if  $C_{vr}$  is unbounded, general online learning in CRMDPs can become exponentially difficult.

- We conduct comprehensive numerical experiments to validate our theoretical findings. In a simulated MDP, we show that the performance of learned policies degrades as  $C_{vr}$  increases. We evaluate our algorithms in a simulated RMDP and the Frozen Lake environment, highlighting their effectiveness when distribution shifts are significant.

## 2 Preliminary

**Constrained Robust MDP (CRMDP)** We denote a finite horizon CRMDP as CRMDP  $(\mathcal{S}, \mathcal{A}, P^o, r, \mathcal{U}^\rho(P^o), H)$ , where  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space,  $P^o = \{P_h^o\}_{h=1}^H$  is the nominal transition kernel,  $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  is the reward function,  $\mathcal{U}^\rho(P^o)$  is the uncertainty set centered around the nominal kernel,  $\rho$  is the uncertainty level,  $H$  is the horizon length. In this work, we specifically focus on general  $f$ -divergence defined  $(s, a)$ -rectangular uncertainty sets [12],  $\mathcal{U}^\rho(P^o) = \otimes_{(s,a,h) \in \mathcal{S} \times \mathcal{A} \times [H]} \mathcal{U}_h^\rho(s, a)$ , where  $\mathcal{U}_h^\rho(s, a) = \{P \in \Delta(\mathcal{S}) | D_f(P || P_h^o(\cdot | s, a)) \leq \rho\}$ . The robust value function and  $Q$ -function are defined as  $V_h^{\pi, \rho}(s) = \inf_{P \in \mathcal{U}^\rho(P^o)} \mathbb{E}_{\pi, P} [\sum_{t=h}^H r_t(s_t, a_t) | s_h = s]$ ,  $Q_h^{\pi, \rho}(s, a) = \inf_{P \in \mathcal{U}^\rho(P^o)} \mathbb{E}_{\pi, P} [\sum_{t=h}^H r_t(s_t, a_t) | s_h = s, a_h = a]$ . The optimal robust value function and optimal robust  $Q$ -function are defined as:  $V_h^{*, \rho}(s) = \sup_{\pi \in \Pi} V_h^{\pi, \rho}(s)$ ,  $Q_h^{*, \rho}(s, a) = \sup_{\pi \in \Pi} Q_h^{\pi, \rho}(s, a)$ , where  $\Pi$  is the set of all policies. Correspondingly, the optimal robust policy is the policy that achieves the optimal robust value function  $\pi_h^* = \operatorname{arg\,sup}_{\pi \in \Pi} V_h^{\pi, \rho}(s)$ . For CRMDPs, [12] proved the robust Bellman optimality equations

$$Q_h^{*, \rho}(s, a) = r_h(s, a) + \inf_{P_h \in \mathcal{U}_h^\rho(P^o)} \mathbb{E}_{P_h} [V_{h+1}^{*, \rho}(s, a)], \quad V_h^{*, \rho}(s) = \max_{a \in \mathcal{A}} Q_h^{*, \rho}(s, a),$$

where  $\mathbb{E}_{P_h} [V_{h+1}^{*, \rho}](s, a) := \mathbb{E}_{s' \sim P_h(\cdot | s, a)} [V_{h+1}^{*, \rho}(s')]$ .

**Regularized Robust MDP (RRMDP)** A finite horizon RRMDP can be denoted as RRMDP  $(\mathcal{S}, \mathcal{A}, P^o, r, \beta, R, H)$ , where  $\beta$  is the regularizer parameter,  $R$  is a penalty on distribution shift, and we set  $R$  to be the probability divergence  $D$  throughout this paper. RRMDPs replace the uncertainty set constraint in CRMDPs with a regularization term. Specifically, the robust value function and  $Q$ -function under the regularized setting are defined as

$$V_h^{\pi, \beta}(s) = \inf_{P \in \Delta(\mathcal{S})} \mathbb{E}_{\pi, P} \left[ \sum_{t=h}^H r_t(s_t, a_t) + \beta \cdot D(P_t(\cdot | s_t, a_t), P_t^o(\cdot | s_t, a_t)) \middle| s_h = s \right],$$

$$Q_h^{\pi, \beta}(s, a) = \inf_{P \in \Delta(\mathcal{S})} \mathbb{E}_{\pi, P} \left[ \sum_{t=h}^H r_t(s_t, a_t) + \beta \cdot D(P_t(\cdot | s_t, a_t), P_t^o(\cdot | s_t, a_t)) \middle| s_h = s, a_h = a \right].$$

For RRMDPs, [42] showed the robust Bellman optimality equations:

$$Q_h^{*, \beta}(s, a) = r_h(s, a) + \inf_{P_h \in \Delta(\mathcal{S})} [\mathbb{E}_{P_h} [V_{h+1}^{*, \beta}](s, a) + \beta D(P_h(\cdot | s, a), P_h^o(\cdot | s, a))], \quad V_h^{*, \beta}(s) = \max_{a \in \mathcal{A}} Q_h^{*, \beta}(s, a).$$

**Learning Goal** We have an agent interacting with the nominal environment for  $K$  episodes. At the start of episode  $k$  with initial state  $s_1^k$ , the agent chooses a policy  $\pi^k$  based on the history information. Then it executes  $\pi^k$  and collects a new trajectory. The agent's goal is to minimize the cumulative regret after  $K$  episodes, defined as  $\operatorname{Regret}(K) = \sum_{k=1}^K [V_1^{*, \rho}(s_1^k) - V_1^{\pi^k, \rho}(s_1^k)]$  for CRMDPs, and  $\operatorname{Regret}(K) = \sum_{k=1}^K [V_1^{*, \beta}(s_1^k) - V_1^{\pi^k, \beta}(s_1^k)]$  for RRMDPs.

## 3 Algorithm design and theoretical guarantee

In this section, we present our Algorithm 1 and provide theoretical guarantees. We postpone more details on Algorithm 1 to Appendix C, and more results of theoretical guarantee to Appendix D.

To characterize the inherent difficulty in learning online RMDPs with general  $f$ -divergences, we propose an intrinsic metric that captures the essential of the problem, based on visitation measures in both the nominal and worst-case environments. Specifically, we have the following assumption.

---

**Algorithm 1** Online Robust Bellman Iteration (ORBIT)

---

**Require:** uncertainty level  $\rho > 0$  (for CRMDPs), or regularizer  $\beta > 0$  (for RRMDPs).

```
1: for  $k = 1, \dots, K$  do
2:    $\hat{V}_{H+1}^k(\cdot) \leftarrow 0$ .
3:   for  $h = H, \dots, 1$  do
4:     for  $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$  do
5:       Update  $\hat{Q}_h^k(s, a)$ : CRMDP–refer to Appendix C.2; RRMDP–refer to Appendix C.3
6:     end for
7:     for  $\forall s \in \mathcal{S}$  do
8:        $\pi_h^k(s) \leftarrow \operatorname{argmax}_{a \in \mathcal{A}} \hat{Q}_h^k(s, a)$ ,  $\hat{V}_h^k(s) \leftarrow \max_{a \in \mathcal{A}} \hat{Q}_h^k(s, a)$ .
9:     end for
10:  end for
11:  Collect trajectory  $\tau^k$  by executing  $\pi^k$ ;
12:  Update  $n_h^k, \hat{r}_h^{k+1}, \hat{P}_h^{k+1}$  according to (C.1).
13: end for
```

---

**Assumption 3.1** (Bounded visitation measure ratio). *Under the definition of Definition D.3, we define  $C_{vr} := \sup_{\pi, h, s} \frac{q_h^\pi(s)}{d_h^\pi(s)}$  as the supremal ratio between the nominal visitation measure and the worst-case visitation measure. We assume that  $C_{vr}$  is polynomial in  $H, S$  and  $A$ .*

We focus on CRMDPs with TV, KL and  $\chi^2$  divergences defined uncertainty sets in the main paper, results on RRMDPs are similar and postponed to Appendix D due to space limit.

**Theorem 3.2** (CRMDP regret upper bounds). *Assume Assumption 3.1 holds for CRMDPs with TV, KL and  $\chi^2$  divergence defined uncertainty sets. Assume Assumption D.7 holds for the CRMDP-KL setting. Then for any  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ , Algorithm 1 satisfies*

$$\text{Regret}(K) = \begin{cases} \tilde{O}(C_{vr} S^2 A H^2 + C_{vr}^{\frac{1}{2}} S^{\frac{3}{2}} A^{\frac{1}{2}} H^2 \sqrt{K}) & (\text{TV}) \\ \tilde{O}\left(\left(1 + \frac{H\sqrt{S}}{\rho C_{MP}}\right)(C_{vr} S A H + C_{vr}^{\frac{1}{2}} S^{\frac{1}{2}} A^{\frac{1}{2}} H \sqrt{K})\right) & (\text{KL}) \\ \tilde{O}((1 + \sqrt{\rho})(C_{vr} S^2 A H^2 + C_{vr}^{\frac{1}{2}} S^{\frac{3}{2}} A^{\frac{1}{2}} H^2 \sqrt{K})) & (\chi^2) \end{cases}$$

**Theorem 3.3** (CRMDP regret lower bound). *For CRMDPs with TV, KL and  $\chi^2$  divergence defined uncertainty sets, for any learning algorithm  $\xi$ , there exists a CRMDP  $\mathcal{M}$  satisfying Assumption 3.1, such that  $\mathbb{E}[\text{Regret}_{\mathcal{M}}(\xi, K)] = \Omega(C_{vr}^{\frac{1}{2}} \sqrt{K})$ .*

**Remark 3.4.** Comparing Theorem 3.2 and Theorem 3.3, we observe the order of  $C_{vr}$  in the dominant terms of the upper bounds matches that in the lower bounds. The upper bounds thus align with the lower bounds in the two most critical parameters governing sample complexity:  $C_{vr}$  and  $K$ . This indicates the fundamental presence of the information deficit issue in the online learning of robust policies, which stems from the discrepancy between the nominal and worst-case transitions and can be characterized by  $C_{vr}$ .

## 4 Conclusion

We investigated online robust reinforcement learning within the context of tabular CRMDPs and RRMDPs, demonstrating that when the nominal MDP is sufficiently exploratory, sample-efficient online learning becomes feasible. We quantified the exploration efficiency of RMDPs through a novel quantity called the supremal visitation ratio. We constructed hard instances to show that a moderate supremal visitation ratio is necessary for ensuring sample-efficient online learning. We developed computationally efficient algorithms and provided regret analyses with both upper and lower bounds, which indicates our algorithm has an optimal dependency on the supremal visitation ratio and the number of episodes. We also conducted numerical experiments on diverse environments to validate our theory and show the robustness of our proposed algorithm, which is postponed to Appendix E due to space limit.

## References

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 24, 2011.
- [2] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- [3] Kishan Panaganti Badrinath and Dileep Kalathil. Robust reinforcement learning using least squares policy iteration with provable performance guarantees. In *International Conference on Machine Learning*, pages 511–520. PMLR, 2021.
- [4] Jose Blanchet, Miao Lu, Tong Zhang, and Han Zhong. Double pessimism is provably efficient for distributionally robust offline reinforcement learning: Generic algorithm and robust partial coverage. *Advances in Neural Information Processing Systems*, 36:66845–66859, 2023.
- [5] Nicolo Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- [6] Longchao Da, Justin Turnau, Thirulogasankar Pranav Kutralingam, Alvaro Velasquez, Paulo Shakarian, and Hua Wei. A survey of sim-to-real methods in rl: Progress, prospects and challenges with foundation models. *arXiv preprint arXiv:2502.13187*, 2025.
- [7] Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying pac and regret: Uniform pac bounds for episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 30, 2017.
- [8] Siddharth Desai, Ishan Durugkar, Haresh Karnan, Garrett Warnell, Josiah Hanna, and Peter Stone. An imitation from observation approach to transfer learning with dynamics mismatch. *Advances in Neural Information Processing Systems*, 33:3917–3929, 2020.
- [9] Jing Dong, Jingwei Li, Baoxiang Wang, and Jingzhao Zhang. Online policy optimization for robust markov decision process. In *Conference on Uncertainty in Artificial Intelligence*, 2024.
- [10] Benjamin Eysenbach, Shreyas Chaudhari, Swapnil Asawa, Sergey Levine, and Ruslan Salakhutdinov. Off-dynamics reinforcement learning: Training for transfer with domain classifiers. In *International Conference on Learning Representations*, 2021.
- [11] Yihong Guo, Yixuan Wang, Yuanyuan Shi, Pan Xu, and Anqi Liu. Off-dynamics reinforcement learning via domain adaptation and reward augmented imitation. In *Advances in Neural Information Processing Systems*, volume 37, pages 136326–136360, 2024.
- [12] Garud N Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.
- [13] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [14] Gen Li, Laixi Shi, Yuxin Chen, Yuejie Chi, and Yuting Wei. Settling the sample complexity of model-based offline reinforcement learning. *The Annals of Statistics*, 52(1):233–260, 2024.
- [15] Zhishuai Liu, Weixin Wang, and Pan Xu. Upper and lower bounds for distributionally robust off-dynamics reinforcement learning. *arXiv preprint arXiv:2409.20521*, 2024.
- [16] Zhishuai Liu and Pan Xu. Distributionally robust off-dynamics reinforcement learning: Provable efficiency with linear function approximation. In *International Conference on Artificial Intelligence and Statistics*, pages 2719–2727. PMLR, 2024.
- [17] Zhishuai Liu and Pan Xu. Minimax optimal and computationally efficient algorithms for distributionally robust offline reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 37, pages 86602–86654, 2024.
- [18] Zhishuai Liu and Pan Xu. Linear mixture distributionally robust markov decision processes. *arXiv preprint arXiv:2505.18044*, 2025.

- [19] Miao Lu, Han Zhong, Tong Zhang, and Jose Blanchet. Distributionally robust reinforcement learning with interactive data collection: Fundamental hardness and near-optimal algorithms. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [20] Jiafei Lyu, Kang Xu, Jiacheng Xu, Jing-Wen Yang, Zongzhang Zhang, Chenjia Bai, Zongqing Lu, Xiu Li, et al. Odr1: A benchmark for off-dynamics reinforcement learning. *Advances in Neural Information Processing Systems*, 37:59859–59911, 2024.
- [21] Shie Mannor, Ofir Mebel, and Huan Xu. Robust mdps with k-rectangular uncertainty. *Mathematics of Operations Research*, 41(4):1484–1509, 2016.
- [22] Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- [23] Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.
- [24] Kishan Panaganti and Dileep Kalathil. Sample complexity of robust reinforcement learning with a generative model. In *International Conference on Artificial Intelligence and Statistics*, pages 9582–9602. PMLR, 2022.
- [25] Kishan Panaganti, Adam Wierman, and Eric Mazumdar. Model-free robust  $\phi$ -divergence reinforcement learning using both offline and online data. In *Proceedings of the 41st International Conference on Machine Learning*, pages 39324–39363, 2024.
- [26] Kishan Panaganti, Zaiyan Xu, Dileep Kalathil, and Mohammad Ghavamzadeh. Robust reinforcement learning using offline data. *Advances in Neural Information Processing Systems*, 35:32211–32224, 2022.
- [27] Jay K Satia and Roy E Lave Jr. Markovian decision processes with uncertain transition probabilities. *Operations Research*, 21(3):728–740, 1973.
- [28] Abdolreza Sayyareh. A new upper bound for kullback-leibler divergence. *Appl. Math. Sci*, 67:3303–3317, 2011.
- [29] Laixi Shi and Yuejie Chi. Distributionally robust model-based offline reinforcement learning with near-optimal sample complexity. *Journal of Machine Learning Research*, 25(200):1–91, 2024.
- [30] Laixi Shi, Gen Li, Yuting Wei, Yuxin Chen, Matthieu Geist, and Yuejie Chi. The curious price of distributional robustness in reinforcement learning with a generative model. *Advances in Neural Information Processing Systems*, 36, 2024.
- [31] Cheng Tang, Zhishuai Liu, and Pan Xu. Robust offline reinforcement learning with linearly structured  $\phi$ -divergence regularization. In *Forty-second International Conference on Machine Learning*, 2025.
- [32] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [33] He Wang, Laixi Shi, and Yuejie Chi. Sample complexity of offline distributionally robust linear markov decision processes. In *Reinforcement Learning Conference*, 2024.
- [34] Ruhan Wang, Yu Yang, Zhishuai Liu, Dongruo Zhou, and Pan Xu. Return augmented decision transformer for off-dynamics reinforcement learning. *arXiv preprint arXiv:2410.23450*, 2024.
- [35] Yue Wang and Shaofeng Zou. Online robust reinforcement learning with model uncertainty. *Advances in Neural Information Processing Systems*, 34:7193–7206, 2021.
- [36] Tsachy Weissman, Erik Ordentlich, Gadiel Seroussi, Sergio Verdu, and Marcelo J Weinberger. Inequalities for the l1 deviation of the empirical distribution. *Hewlett-Packard Labs, Tech. Rep*, page 125, 2003.

- [37] Xiaoyu Wen, Chenjia Bai, Kang Xu, Xudong Yu, Yang Zhang, Xuelong Li, and Zhen Wang. Contrastive representation for data filtering in cross-domain offline reinforcement learning. In *Proceedings of the 41st International Conference on Machine Learning*, pages 52720–52743, 2024.
- [38] Wolfram Wiesemann, Daniel Kuhn, and Berç Rustem. Robust markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013.
- [39] Huan Xu and Shie Mannor. The robustness-performance tradeoff in markov decision processes. *Advances in Neural Information Processing Systems*, 19, 2006.
- [40] Kang Xu, Chenjia Bai, Xiaoteng Ma, Dong Wang, Bin Zhao, Zhen Wang, Xuelong Li, and Wei Li. Cross-domain policy adaptation via value-guided data filtering. *Advances in Neural Information Processing Systems*, 36:73395–73421, 2023.
- [41] Zaiyan Xu, Kishan Panaganti, and Dileep Kalathil. Improved sample complexity bounds for distributionally robust reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 9728–9754. PMLR, 2023.
- [42] Wenhao Yang, Han Wang, Tadashi Kozuno, Scott M Jordan, and Zhihua Zhang. Robust markov decision processes without model estimation. *arXiv preprint arXiv:2302.01248*, 2023.
- [43] Wenhao Yang, Liangyu Zhang, and Zhihua Zhang. Toward theoretical understandings of robust markov decision processes: Sample complexity and asymptotics. *The Annals of Statistics*, 50(6):3223–3248, 2022.
- [44] Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pages 7304–7312. PMLR, 2019.
- [45] Runyu Zhang, Yang Hu, and Na Li. Soft robust mdps and risk-sensitive mdps: Equivalence, policy gradient, and sample complexity. In *The Twelfth International Conference on Learning Representations*, 2024.
- [46] Tan Zhang, Kefang Zhang, Jiatao Lin, Wing-Yue Geoffrey Louie, and Hui Huang. Sim2real learning of obstacle avoidance for robotic manipulators in uncertain environments. *IEEE Robotics and Automation Letters*, 7(1):65–72, 2021.
- [47] Zhengqing Zhou, Zhengyuan Zhou, Qinxun Bai, Linhai Qiu, Jose Blanchet, and Peter Glynn. Finite-sample regret bound for distributionally robust offline tabular reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3331–3339. PMLR, 2021.

## A Notations

For any positive integer  $H \in \mathbb{Z}_+$ , we denote  $[H] = \{1, 2, \dots, H\}$ . For any set  $\mathcal{S}$ , define  $\Delta(\mathcal{S})$  as the set of probability distributions over  $\mathcal{S}$ . Let  $P, Q \in \Delta(\mathcal{S})$  and  $P \ll Q$ . For a convex function  $f : [0, +\infty) \rightarrow (-\infty, +\infty]$  such that  $f(x)$  is finite for all  $x > 0$ ,  $f(1) = 0$  and  $f(0) = \lim_{t \rightarrow 0^+} f(t)$ . The  $f$ -divergence of  $P$  from  $Q$ , which measures their difference, is defined as  $D_f(P\|Q) = \int_{\Omega} f(\frac{dP}{dQ}) dQ$ . In our paper, we consider three common  $f$ -divergences including total variation (TV) distance with  $f(t) = \frac{1}{2}|t - 1|$ , Kullback-Leibler (KL) divergence with  $f(t) = t \ln t$ , and  $\chi^2$ -divergence with  $f(t) = (t - 1)^2$ . We use  $\mathcal{O}(\cdot)$  to hide absolute constant factors and  $\tilde{\mathcal{O}}(\cdot)$  to further hide logarithmic factors. For any two integers  $a$  and  $b$ , we denote  $a \vee b := \max\{a, b\}$ .

## B Related Work

**CRMDPs and RRMDPs** The framework of CRMDPs was first introduced in the context of optimal control [12, 23, 39, 38, 21], where the nominal MDP is assumed to be exactly known, and robust policies are obtained by solving a constrained max-min optimization problem. Subsequent works extended CRMDPs to the learning setting with access to a generative model [47, 43, 24, 30]. More recently, CRMDPs have been studied in the offline learning setting, where only a pre-collected dataset from the nominal MDP is available through a behavior policy [29, 26, 4, 33, 17, 18]. To ensure that a robust policy can be learned from a reasonably sized offline dataset, these works make assumptions about the behavior policy (and implicitly, the nominal MDP) to guarantee sufficient coverage. Such assumptions include the robust single-policy clipped concentrability [29], robust partial coverage [4], and uniformly well coverage assumptions [17, 33]. The framework of RRMDPs was more recently proposed by [42] and [45], who studied it under the generative model setting and the offline setting, respectively. This line of work was extended to function approximation settings by [25] and [31], considering both hybrid offline-online and purely offline scenarios.

It is worth noting that CRMDPs are sometimes referred to in the literature as Robust MDPs (RMDPs) or Distributionally Robust MDPs (DRMDPs). To distinguish them from the regularized robust framework, we adopt the term CRMDPs. Similarly, RRMDPs appear under various names, including penalized robust MDPs [42], soft robust MDPs [45], and robust  $\phi$ -regularized MDPs [25]. We use the term RRMDPs to clearly differentiate them from CRMDPs while remaining consistent with the literature.

**Online RMDPs** [35, 3] studied the online learning for infinite-horizon RMDPs with  $R$ -contamination and more general uncertainty sets, respectively. Their algorithmic design and theoretical analysis rely on assuming access to exploratory policies, which implicitly assumes that the nominal MDP is sufficiently exploratory. In contrast, we revisit this challenge from a different angle, focusing on the information deficit issue induced by distributional shift. Our work differs from theirs in two key aspects. First, we introduce a novel quantity to characterize the hardness of exploration in the nominal MDP and analyze it thoroughly via both upper and lower bounds on the sample complexity. Second, instead of assuming access to exploratory policies, we design algorithms that explicitly incorporate exploration strategies tailored for finite-horizon tabular CRMDPs and RRMDPs with  $(s, a)$ -rectangular uncertainty sets defined by general  $f$ -divergences.

[16, 19, 15] focused on online robust RL under the specific setting of CRMDPs with uncertainty sets defined by the TV-distance, coupled with assumptions such as the existence of fail-states or vanishing minimal values. From an information-theoretic perspective, we show that these assumptions effectively circumvent the information deficit by constraining the direction of the distribution shift. In contrast, our work seeks to identify a more general sufficient condition for provably efficient online learning in CRMDPs—one that applies to arbitrary divergence-based uncertainty sets and does not rely on the fail-state or vanishing minimal value assumption.

**Off-dynamics RL** A substantial body of empirical work addresses off-dynamics RL through the lens of domain adaptation and transfer learning [10, 8, 46, 40, 37, 11, 34, 20, 6], among others. In this



paper, we focus on the robust MDP (RMDP) formulation of off-dynamics RL. We refer readers to the above works for complementary approaches along this orthogonal line of research.

## C Online Robust Bellman Iteration (ORBIT)

In this section, we first present a meta-algorithm for online tabular RMDPs with general  $f$ -divergence defined uncertainty sets or regularization terms. We then instantiate the algorithm for CRMDPs with TV, KL and  $\chi^2$ -divergences defined uncertainty sets and RRMDPs with TV, KL and  $\chi^2$ -divergences defined regularization terms, respectively.

### C.1 Algorithm Interpretation

We present our meta-algorithm, Online Robust Bellman Iteration (ORBIT), in Algorithm 1. The algorithm follows a value iteration framework and integrates optimistic estimation and the robust Bellman optimality equation for estimating the robust  $Q$ -functions. In each episode  $k \in [K]$ , ORBIT consists of two stages. In the first stage (Lines 3 to 10), Algorithm 1 iteratively updates the value function and  $Q$ -function estimations in a backward manner. In the second stage (Lines 11 to 12), we collect trajectory  $\tau^k = (s_1^k, a_1^k, r_1^k, \dots, s_H^k, a_H^k, r_H^k)$  by executing  $\pi^k$ . After a new trajectory is collected, ORBIT updates the empirical reward function and transition kernel as follows

$$\begin{aligned} n_h^k(s, a) &= \sum_{i=1}^k \mathbb{1} \{s_h^i = s, a_h^i = a\}, \\ \hat{r}_h^{k+1}(s, a) &= \frac{\sum_{i=1}^k r_h^i(s, a) \cdot \mathbb{1} \{s_h^i = s, a_h^i = a\}}{n_h^k(s, a) \vee 1}, \\ \hat{P}_h^{k+1}(s'|s, a) &= \frac{\sum_{i=1}^k \mathbb{1} \{s_h^i = s, a_h^i = a, s_{h+1}^i = s'\}}{n_h^k(s, a) \vee 1}. \end{aligned} \quad (\text{C.1})$$

Algorithm 1 updates  $Q$ -functions at each  $(s, a)$  according to different RMDP settings and choices of  $f$ -divergences. We differentiate these cases using specific labels: CRMDP-TV, CRMDP-KL, CRMDP- $\chi^2$ , RRMDP-TV, RRMDP-KL, and RRMDP- $\chi^2$ . Finally, Algorithm 1 adopts the greedy policy of the estimated  $Q$ -function as the estimated optimal policy at episode  $k$ .

For the robust  $Q$ -function estimation, we leverage the robust optimality Bellman equation. Incorporating the optimism principle in the face of uncertainty [1] in the  $Q$ -function update, we have

$$\hat{Q}_h^k(s, a) = \min \{ \text{RB}_h^k(s, a) + b_h^k(s, a), H - h + 1 \}. \quad (\text{C.2})$$

There are two components in (C.2): a robust Bellman estimator  $\text{RB}_h^k(s, a)$  and a bonus term  $b_h^k(s, a)$ . Next, we will instantiate this meta-algorithm for CRMDPs and RRMDPs with various  $f$ -divergences, and provide explicit formulation for robust Bellman estimation and bonus design.

### C.2 ORBIT under Constrained Robust MDPs

We first focus on CRMDPs and detail the update of robust  $Q$ -functions (C.2) in various settings. To solve the optimization problem in the robust Bellman equation, we resort to strong duality results in the following.

**CRMDP-TV** In CRMDPs with TV-distance defined uncertainty sets, estimators in (C.2) are defined as follows

$$\text{RB}_h^k(s, a) = \hat{r}_h^k(s, a) - \inf_{\eta \in [0, H]} \left( \mathbb{E}_{\hat{P}_h^k} [(\eta - \hat{V}_{h+1}^{k, \rho})_+] (s, a) + \rho \left( \eta - \min_{s \in \mathcal{S}} \hat{V}_{h+1}^{k, \rho}(s) \right)_+ - \eta \right), \quad (\text{C.3})$$

$$b_h^k(s, a) = 2H \sqrt{\frac{2S^2 \ln(12SAH^2K^2/\delta)}{n_h^{k-1}(s, a) \vee 1}} + \frac{1}{K}, \quad (\text{C.4})$$

where (C.3) represents the empirical version of the robust Bellman operator and (C.4) is the bonus. The dual formulation for TV-distance and the optimism of the estimated  $Q$ -function are established in Appendix H.1.

**CRMDP-KL** In CRMDPs with KL-divergence defined uncertainty sets, estimators in (C.2) are defined as follows

$$\begin{aligned} \text{RB}_h^k(s, a) &= \hat{r}_h^k(s, a) - \inf_{\nu \in [0, \frac{H}{\rho}]} (\nu \ln \mathbb{E}_{\hat{P}_h^k} [\exp(-\nu^{-1} \hat{V}_{h+1}^{k, \rho})](s, a) + \nu \rho), \\ b_h^k(s, a) &= \left(1 + \frac{2H\sqrt{S}}{\rho C_{MP}}\right) \sqrt{\frac{2 \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}}, \end{aligned}$$

where  $C_{MP}$  is defined in Assumption D.7. The dual formulation for KL-divergence and the optimism of the estimated  $Q$ -function are proved in Appendix H.2.

**CRMDP- $\chi^2$**  For ORBIT in CRMDPs with  $\chi^2$ -divergence defined uncertainty sets, we have

$$\begin{aligned} \text{RB}_h^k(s, a) &= \hat{r}_h^k(s, a) + \sup_{\lambda \in [0, H]} \left( \mathbb{E}_{\hat{P}_h^k} [\hat{V}_{h+1}^{k, \rho} - \lambda](s, a) - \sqrt{\rho \text{Var}_{\hat{P}_h^k}(\hat{V}_{h+1}^{k, \rho} - \lambda)} \right), \\ b_h^k(s, a) &= (2 + \sqrt{\rho})H \sqrt{\frac{2S^2 \ln(192SAH^3K^3/\delta)}{n_h^{k-1}(s, a) \vee 1}} + \frac{1 + \sqrt{\rho}}{K}. \end{aligned}$$

The dual formulation for  $\chi^2$ -divergence and the optimism of the estimated  $Q$ -function are proved in Appendix H.3.

### C.3 ORBIT under Regularized Robust MDPs

We then focus on RRMDPs and detail the update of robust  $Q$ -functions (C.2) in various settings.

**RRMDP-TV** In RRMDPs with TV-distance regularization terms, estimators in (C.2) are defined as follows

$$\begin{aligned} \text{RB}_h^k(s, a) &= \hat{r}_h^k(s, a) - \mathbb{E}_{\hat{P}_h^k} \left[ \left( \min_{s \in \mathcal{S}} \hat{V}_{h+1}^{k, \beta}(s) + \beta - \hat{V}_{h+1}^{k, \beta}(s) \right)_+ \right](s, a) + \left( \min_{s \in \mathcal{S}} \hat{V}_{h+1}^{k, \beta}(s) + \beta \right), \\ b_h^k(s, a) &= 2H \sqrt{\frac{2S \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}}. \end{aligned}$$

The dual formulation for TV-distance and the optimism of the estimated  $Q$ -function are proved in Appendix I.1.

**RRMDP-KL** For ORBIT in RRMDPs with KL-divergence regularization terms, we have

$$\begin{aligned} \text{RB}_h^k(s, a) &= \hat{r}_h^k(s, a) - \beta \ln \mathbb{E}_{\hat{P}_h^k} [\exp(-\beta^{-1} \hat{V}_{h+1}^{k, \beta})](s, a), \\ b_h^k(s, a) &= (1 + \beta e^{\beta^{-1}H} \sqrt{S}) \sqrt{\frac{2 \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}}. \end{aligned}$$

The dual formulation for KL-divergence and the optimism of the estimated  $Q$ -function are provided in Appendix I.2.

**RRMDP- $\chi^2$**  For ORBIT in RRMDPs with  $\chi^2$ -divergence regularization terms, we have

$$\begin{aligned} \text{RB}_h^k(s, a) &= \hat{r}_h^k(s, a) + \sup_{\lambda \in [0, H]} \left( \mathbb{E}_{\hat{P}_h^k} [\hat{V}_{h+1}^{k, \beta} - \lambda](s, a) - \frac{1}{4\beta} \text{Var}_{\hat{P}_h^k} [\hat{V}_{h+1}^{k, \beta} - \lambda](s, a) \right), \\ b_h^k(s, a) &= \left( 2H + \frac{3H^2}{4\beta} \right) \sqrt{\frac{2S^2 \ln(48SAH^3K^2/\delta)}{n_h^{k-1}(s, a) \vee 1}} + \frac{1 + 4\beta}{4\beta K}. \end{aligned}$$

The dual formulation for  $\chi^2$ -divergence and the optimism of the estimated  $Q$ -function are provided in Appendix I.3.

## D Theoretical Results

In this section, we provide theoretical understandings on the online learning of RMDPs. We start with a new perspective—the information deficit issue—to understand existing conditions for provably efficient online learning. Motivated by existing solutions to address the information deficit issue, we propose a new metric, the supremal visitation ratio, to quantify the hardness in exploration under online RMDPs. Further, we provide upper and lower bounds, involving the supremal visitation ratio, on the regret of Algorithm 1 in all settings.

### D.1 Learnability of Online RMDPs

Focusing on CRMDPs with TV-distance defined uncertainty sets, [16] and [19] identified that the following assumptions can admit provably efficient online learning. In particular, [16] made the following fail-states assumption.

**Condition D.1** (Fail-states). [16, Condition 4.3] *There exists a subset  $\mathcal{S}_f \subset \mathcal{S}$  of fail-states such that  $r_h(s, a) = 0$ ,  $P_h^o(\mathcal{S}_f | s, a) = 1$ ,  $\forall (s, a, h) \in \mathcal{S}_f \times \mathcal{A} \times [H]$ .*

[19] extended Condition D.1 to Condition D.2, but both assumptions essentially serve the same purpose, eliminating  $\min_{s \in \mathcal{S}} V(s)$  in the dual formulation of the optimization problem in the CRMDP-TV setting.

**Condition D.2** (Vanishing minimal value). [19, Assumption 4.1] *The RMDP satisfies that  $\min_{s \in \mathcal{S}} V_1^{*,\rho}(s) = 0$ .*

To explain the rationale behind Condition D.1, we first define the visitation measure as follows.

**Definition D.3** (Visitation measure). *Under both CRMDPs and RRMDPs, for any policy  $\pi$ , we define the worst-case transition corresponding to  $\pi$  as*

$$\begin{aligned} P_h^{w,\pi}(\cdot | s, a) &= \operatorname{argmin}_{P_h \in \mathcal{U}_h^\rho(P_h^o)} \mathbb{E}_{P_h} [V_{h+1}^{\pi,\rho}](s, a) \text{ for CRMDPs,} \\ P_h^{w,\pi}(\cdot | s, a) &= \operatorname{argmin}_{P_h \in \Delta(\mathcal{S})} \mathbb{E}_{P_h} [V_{h+1}^{\pi,\beta}](s, a) + \beta \cdot D(P_h(\cdot | s, a), P_h^o(\cdot | s, a)) \text{ for RRMDPs.} \end{aligned}$$

At timestep  $h \in [H]$ , we denote  $d_h^\pi(\cdot)$  as the visitation measure on  $\mathcal{S}$  induced by the policy  $\pi$  under  $P^o$ , and  $q_h^\pi(\cdot)$  as the visitation measure on  $\mathcal{S}$  induced by the policy  $\pi$  under  $P^{w,\pi}$ . For convenience, we also write  $P^{w,k} := P^{w,\pi^k}$ ,  $d^k := d^{\pi^k}$  and  $q^k := q^{\pi^k}$ .

We show that Condition D.1 implies the following property on the CRMDP.

**Proposition D.4.** *For CRMDPs with TV-distance defined uncertainty set satisfying Condition D.1, for any  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}$ ,  $s' \in \mathcal{S} \setminus \mathcal{S}_f$  and policy  $\pi$ , we have  $P_h^{w,\pi}(s' | s, a) \leq P_h^o(s' | s, a)$ .*

Proposition D.4 shows that, for any state-action pair, the transition probability to a non-fail state is smaller in the worst-case environment than in the nominal environment. Consequently, non-fail states that are rarely visited in the nominal environment remain rarely visited in the worst-case environment, hence it would not incur large regret when making decisions at these states. Meanwhile, by definition, states in  $\mathcal{S}_f$  lead to precisely zero value no matter what action is taken and thus no regret could be incurred at these states. This implies that Condition D.1 or Condition D.2 ensures the information obtained from exploration in the nominal environment is sufficient for decision making in the worst-case environment, thus bypassing the information deficit issue.

Note that both Condition D.1 and Condition D.2 are specifically designed for CRMDPs with TV-distance defined uncertainty sets. In more general  $f$ -divergence contexts, such as RMDPs with KL-divergence or  $\chi^2$ -divergence defined uncertainty sets or regularization terms, the property described in Proposition D.4 does not hold. Consequently, learning RMDPs through online interaction is in general a challenging open problem [19] without additional assumptions. To characterize the inherent difficulty in learning online RMDPs with general  $f$ -divergences, we propose a more intrinsic metric

that captures the essential of the problem, based on visitation measures in both the nominal and worst-case environments. Specifically, we have the following assumption.

**Assumption D.5** (Bounded visitation measure ratio). *Under the definition of Definition D.3, we define  $C_{vr} := \sup_{\pi, h, s} \frac{q_h^\pi(s)}{d_h^\pi(s)}$  as the supremal ratio between the nominal visitation measure and the worst-case visitation measure. We assume that  $C_{vr}$  is polynomial in  $H, S$  and  $A$ .*

**Remark D.6** (Reduction to non-robust setting). *In non-robust settings, Assumption 3.1 is always satisfied with  $C_{vr} = 1$ , since  $q_h^\pi(s) = d_h^\pi(s)$ . This indicates our results also apply to the non-robust setting as a special case.*

Visitation measure is a common metric in the offline literature such as [14, Definition 3] and [29, Assumption 1]. Assumption 3.1 ensures that the information obtained in the nominal environment can be effectively used for estimation in the worst-case environment.

## D.2 Regret Bound for Constrained Robust MDP

We first focus on CRMDPs with TV, KL and  $\chi^2$  divergences defined uncertainty sets. Before we present our results, we introduce an extra assumption for the CRMDP-KL setting.

**Assumption D.7.** *We assume there exists a constant  $C_{MP} > 0$ , such that for any  $(h, s, a, s') \in [H] \times \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ , if  $P_h^o(s'|s, a) > 0$ , then  $P_h^o(s'|s, a) > C_{MP}$ .*

**Remark D.8.** *For CRMDPs with KL-divergence defined uncertainty sets, Assumption D.7 guarantees the regularity of dual formulation. We note that similar assumptions also appear in [43, Theorem 3.2] and [29, Theorem 3], both study CRMDPs with KL-divergence defined uncertainty sets.*

**Theorem D.9** (CRMDP regret upper bounds). *Assume Assumption 3.1 holds for CRMDPs with TV, KL and  $\chi^2$  divergence defined uncertainty sets. Assume Assumption D.7 holds for the CRMDP-KL setting. Then for any  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ , Algorithm 1 satisfies*

$$\text{Regret}(K) = \begin{cases} \tilde{O}(C_{vr}S^2AH^2 + C_{vr}^{\frac{1}{2}}S^{\frac{3}{2}}A^{\frac{1}{2}}H^2\sqrt{K}) & (\text{TV}) \\ \tilde{O}\left(\left(1 + \frac{H\sqrt{S}}{\rho C_{MP}}\right)(C_{vr}SAH + C_{vr}^{\frac{1}{2}}S^{\frac{1}{2}}A^{\frac{1}{2}}H\sqrt{K})\right) & (\text{KL}) \\ \tilde{O}((1 + \sqrt{\rho})(C_{vr}S^2AH^2 + C_{vr}^{\frac{1}{2}}S^{\frac{3}{2}}A^{\frac{1}{2}}H^2\sqrt{K})) & (\chi^2) \end{cases}$$

**Remark D.10.** *Theorem 3.2 presents the first provably sublinear result in the online RMDP literature for KL and  $\chi^2$  defined uncertainty sets. The differences in dominant terms stem from how value function errors are amplified in dual formulations through Bellman equations and induction. Although the radius  $\rho$  is not explicitly part of the results for TV-distance and  $\chi^2$ -divergence, it implicitly affects the bounds through  $C_{vr}$  in Assumption 3.1. A larger  $\rho$  loosens constraints, increases distribution shifts, and consequently, requires a higher  $C_{vr}$  and leads to increased regret bound.*

It is worth noting that under the policy selection scheme in Algorithm 1, our regret upper bounds still hold if we relax the definition of  $C_{vr}$  in Assumption 3.1 to be defined as the supremum visitation ratio over deterministic policies. By the standard online-to-batch conversion [5], the regret bounds in Theorem 3.2 immediately imply the following sample complexity results:

**Corollary D.11.** *Under the same setup in Theorem 3.2, when*

$$K = \begin{cases} \tilde{O}(C_{vr}S^3AH^4/\epsilon^2) & (\text{TV}) \\ \tilde{O}\left(\left(1 + \frac{H\sqrt{S}}{\rho C_{MP}}\right)^2 C_{vr}SAH^2/\epsilon^2\right) & (\text{KL}) \\ \tilde{O}((1 + \sqrt{\rho})^2 C_{vr}S^3AH^4/\epsilon^2) & (\chi^2) \end{cases}$$

*w.p. at least  $1 - \delta$ , the uniform mixture of the policies produced by Algorithm 1 is  $\epsilon$ -optimal.*

To see how tight the upper bounds in Theorem 3.2 are, we provide the following lower bounds.

**Theorem D.12** (CRMDP regret lower bound). *For CRMDPs with TV, KL and  $\chi^2$  divergence defined uncertainty sets, for any learning algorithm  $\xi$ , there exists a CRMDP  $\mathcal{M}$  satisfying Assumption 3.1, such that  $\mathbb{E}[\text{Regret}_{\mathcal{M}}(\xi, K)] = \Omega(C_{vr}^{\frac{1}{2}}\sqrt{K})$ .*

**Remark D.13.** Comparing Theorem 3.2 and Theorem 3.3, we observe the order of  $C_{vr}$  in the dominant terms of the upper bounds matches that in the lower bounds. The upper bounds thus align with the lower bounds in the two most critical parameters governing sample complexity:  $C_{vr}$  and  $K$ . This indicates the fundamental presence of the information deficit issue in the online learning of robust policies, which stems from the discrepancy between the nominal and worst-case transitions and can be characterized by  $C_{vr}$ .

Based on the proof of Theorem 3.3, we construct hard instances to illustrate the necessity of Assumption 3.1 in guaranteeing sample efficient online learning in CRMDPs.

**Lemma D.14** (CRMDP hard instances). *For CRMDPs with TV, KL and  $\chi^2$  divergence defined uncertainty sets, for any learning algorithm  $\xi$ , there exists a CRMDP  $\mathcal{M}$  with  $C_{vr} = 2^{2^A}$ , such that  $\mathbb{E}[\text{Regret}_{\mathcal{M}}(\xi, K)] = \Omega(2^A \sqrt{K})$ .*

**Remark D.15.** Lemma D.14 shows that, in the absence of additional assumptions, any online learning algorithm may perform poorly in CRMDPs. The hard instances are constructed by selecting a critical state that has an exponentially small visitation measure in the nominal environment, and make it has a visitation measure of constant order in the worst-case environment. As a result, an agent requires an exponential number of episodes to explore sufficient information about this state, while suffering a constant regret per episode when taking a suboptimal action.

### D.3 Regret Bound for Regularized Robust MDP

We then focus on RRMDPs with TV, KL and  $\chi^2$  divergences defined regularization.

**Theorem D.16** (RRMDP regret upper bound). *Assume Assumption 3.1 holds for RRMDPs with TV, KL and  $\chi^2$  divergence defined regularization terms. Then for any  $\delta \in (0, 1)$  with probability at least  $1 - \delta$ , Algorithm 1 satisfies*

$$\text{Regret}(K) = \begin{cases} \tilde{O}(C_{vr} S^{\frac{3}{2}} A H^2 + C_{vr}^{\frac{1}{2}} S A^{\frac{1}{2}} H^2 \sqrt{K}) & (\text{TV}) \\ \tilde{O}((1 + \beta e^{\beta^{-1} H} \sqrt{S})(C_{vr} S A H + C_{vr}^{\frac{1}{2}} S^{\frac{1}{2}} A^{\frac{1}{2}} H \sqrt{K})) & (\text{KL}) \\ \tilde{O}((1 + \frac{H}{\beta})(C_{vr} S^2 A H^2 + C_{vr}^{\frac{1}{2}} S^{\frac{3}{2}} A^{\frac{1}{2}} H^2 \sqrt{K})) & (\chi^2) \end{cases}$$




Figure 1: Figure 1(a) shows the comparison of the learned policy and the optimal policy in Appendix E.1 (Illustration of the Effect of  $C_{vr}$  on Robustness), where the optimal policy represents the ground truth optimal policy, the learned policy is obtained by Algorithm 1. Figures 1(b) to 1(d) present the comparison between our algorithm ORBIT and the non-robust algorithm in Appendix E.2 (Learning on Simulated RMDPs).

**Remark D.17.** Theorem D.16 presents the first provably sublinear results in the online RRMDP settings. The result in Theorem D.16 corresponding to the TV-distance is smaller by a factor of  $O(\sqrt{S})$  compared to Theorem 3.2. This efficiency gain arises because the dual formulation for RRMDPs eliminates the need for constructing an  $\epsilon$ -net, which also speeds up solving the inner optimization problem and demonstrates computational advantages. Notably, Assumption D.7 is not required in the RRMDP-KL setting, as the dual formulation for KL has a closed-form solution, which also reduces the computational cost. In the  $\chi^2$ -divergence setting, the upper bound is larger than that in the constrained setting by a factor of  $O(H)$ . This gap derives from the differences in dual formulations in two RMDPs, where Lemma I.9 in the RRMDP- $\chi^2$  setting does not admit a square root compared to Lemma H.15 in the CRMDP- $\chi^2$  setting.

**Theorem D.18** (RRMDP regret lower bound). *For RRMDPs with TV, KL and  $\chi^2$  divergence defined regularization terms, for any learning algorithm  $\xi$ , there exists a RRMDP  $\mathcal{M}$  satisfying Assumption 3.1, such that  $\mathbb{E}[\text{Regret}_{\mathcal{M}}(\xi, K)] = \Omega(C_{vr}^{\frac{1}{2}}\sqrt{K})$ .*

Comparing Theorem D.16 and Theorem D.18, we observe that the order of  $C_{vr}$  in the dominant terms of upper bounds matches that in the lower bound. Together with the observation in Remark 3.4, we can conclude that  $C_{vr}$  is a tight measure for evaluating exploration difficulty in RMDPs.

## E Experiments

In this section, we conduct numerical experiments to thoroughly verify the theoretical findings in previous sections. All numerical experiments were conducted on a server equipped with Intel(R) Xeon(R) Gold 5118 CPU @ 2.30GHz.

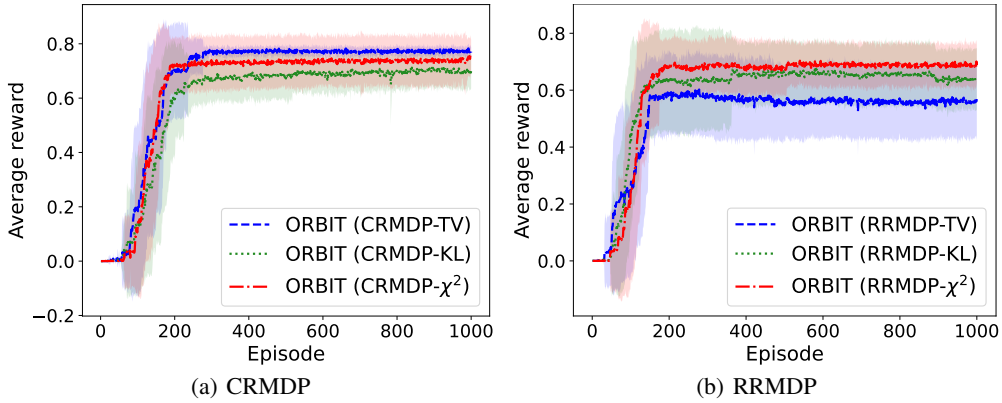


Figure 2: The convergence of ORBIT in Appendix E.3 (Learning the Frozen Lake Problem). We use the policies obtained after each training episode to evaluate the convergence.

### E.1 Illustration of the Effect of $C_{vr}$ on Robustness

The supremal visitation ratio  $C_{vr}$  measures the difficulty in exploration. With a fixed number of episodes, our results Theorem 3.2 and Theorem D.16 show that  $C_{vr}$  would increase the sub-optimality gap of the learned policies. In this section, we construct a toy example (see Figure 4) with  $H = 3$ ,  $\mathcal{S} = \{s_0, \dots, s_5\}$ , and  $\mathcal{A} = \{0, \dots, 9\}$ , focusing on the CRMDP-TV setting. More details about the environment can be found in Appendix F.1. The visitation measure of each states are influenced by a hyper-parameter  $\beta$ , and we can calculate that  $C_{vr} = 3 + \frac{1}{\beta}$  in this case.

As we can see from Figure 1(a), when we increase  $C_{vr}$  by decreasing  $\beta$  in the nominal environment, it becomes harder for the agent to explore the nominal environment sufficiently to learn the optimal policy in the perturbed environment, and thus deteriorate the performance of the learned policy and enlarge the sub-optimality gap. This aligns well with our theoretical results.

### E.2 Learning on Simulated RMDPs

Next, we design a simple MDP with learning horizon  $H = 3$ , the state space is  $\mathcal{S} = \{s_0, \dots, s_4\}$ , and the action space is  $\mathcal{A} = \{0, \dots, 4\}$ . The source environment and target environment are illustrated in Figures 5(a) and 5(b). More details about the environment can be found in Appendix F.2.

The experiment results are presented in Figures 1(b) to 1(d). We can see that the policies under CRMDPs perform similarly to their RRMDP counterpart. And compared to the non-robust algorithm,




Figure 3: Results in Appendix E.3 (Learning the Frozen Lake Problem). Figure 3(a) presents the average time taken for training in various settings. Figures 3(b) to 3(d) present the comparison between our algorithm ORBIT and the non-robust algorithm. We use the last episode policy  $\pi^K$  for the comparison.

all robust policies are less sensitive to the environment perturbation. In particular, the performance of the non-robust algorithm drops drastically with respect to the perturbation. When the perturbation exceeds 0.6, all robust policies outperform the non-robust algorithm. This confirms the robustness of our proposed algorithm.

### E.3 Learning the Frozen Lake Problem

Now we test our algorithm in a hard-to-explore setting, the Frozen Lake problem. In this scenario, the agent’s objective is to traverse a frozen lake from the Start (S) to the Goal (G) without falling into any Holes (H), navigating over the Frozen (F) surface. A hyper-parameter  $P_{\text{perturb}}$  is used to measure the perturbation in the test environment. More details about the environment can be found in Appendix F.3.

First, we evaluate the convergence of our algorithm by tracking the average reward throughout the training process in a single target environment with a fixed perturbation model,  $P_{\text{perturb}}$ . Specifically, we compute the average reward of the policy  $\pi^k$  obtained after each episode  $k$ . As shown in Figures 2(a) and 2(b), our algorithm consistently converges by the end of training. The corresponding average training time is reported in Figure 3(a).

For RRMDPs with TV and KL divergence defined regularization terms, the dual formulations of the  $Q$ -functions admit closed-form solutions, simplifying the training process and resulting in lower computation complexity compared to CRMDPs. Lemma I.9 in the RRMDP- $\chi^2$  setting, on the other hand, requires solving optimization problems to get the dual formulations, leading to higher computational complexity than Lemmas H.1 and I.1 in both two RMDPs with TV-distance and Lemmas H.11 and I.5 in both two RMDPs with KL-divergence, though still lower than Lemma H.15 in the CRMDP- $\chi^2$  setting. Notably, the training time for the CRMDP-TV setting is not significantly higher than that of RRMDP-TV, as we incorporate an additional optimization algorithm (detailed in Algorithm 2) to accelerate computation. Also, due to our additional optimizations, increased exploration may result in longer training times.

We also evaluate the robustness of the policy  $\pi^K$  after  $K$  iterations of updates in various target environments with different  $P_{\text{perturb}}$ , by calculating the average reward obtained in each target environment. As shown in Figures 3(b) to 3(d), our robust algorithm outperforms the corresponding non-robust version in most cases.

## F Additional Details on Experiments.

In this section, we provide more details about the experiments conducted in Appendix E.

## F.1 More Details on Appendix E.1 (Illustration of the Effect of $C_{vr}$ on Robustness)

**Construction** We consider a simulated RMDP Figure 4(a) with horizon length  $H = 3$ , state space  $\mathcal{S} = \{s_0, \dots, s_5\}$ , and action space  $\mathcal{A} = \{0, \dots, 9\}$ . At each episode, the initial state is always  $s_0$ . The nominal transition at the first stage is independent of actions taken,  $P^o(s_1|s_0) = 1 - P^o(s_2|s_0) = \beta$ , where  $\beta$  is a hyperparameter. At the second stage, if the current state is  $s_2$ , it will transit to  $s_5$  with probability 1; if the current state is  $s_1$ , then we have  $P^o(s_3|s_1, a = 0) = \frac{5}{6}$ ,  $P^o(s_3|s_1, a) = \frac{1}{2}, \forall a \in \{1, \dots, 9\}$  and  $P^o(s_4|s_1, a) = 1 - P^o(s_3|s_1, a)$ . Only  $s_3$  and  $s_5$  can generate reward, with  $r(s_3, a) = 1, r(s_5, a) = \frac{1}{2}, \forall a \in \mathcal{A}$ . By construction, the action taken at  $s_1$  determines the final reward, and there are actually two kinds of actions:  $\tilde{a} = 0$  if  $a = 0$  and  $\tilde{a} = 1$  if  $a \in \mathcal{A}/0$ . Though actions in  $\mathcal{A}/0$  are equivalent, they are set to increase the harness in exploration. We construct a TV-distance defined uncertainty set with radius  $\rho = \frac{1}{3}$ .

It is easy to observe that, regardless of the policy  $\pi$  chosen,  $V^{\pi, \rho}(s_4) \leq V^{\pi, \rho}(s_3)$  and  $V^{\pi, \rho}(s_1) \leq V^{\pi, \rho}(s_2)$ , therefore the worst-case transition probability for any policy  $\pi$  is  $P^{w, \pi}(s_3|s_1, a = 0) = \frac{1}{2}$ ,  $P^{w, \pi}(s_3|s_1, a) = \frac{1}{6}, \forall a \in \{1, \dots, 9\}$  and  $P^{w, \pi}(s_4|s_1, a) = 1 - P^{w, \pi}(s_3|s_1, a)$ . Thus,  $V^{\pi, \rho}(s_1) = \frac{1}{2} - \frac{\tilde{a}}{3} \leq \frac{1}{2} = V^{\pi, \rho}(s_2)$ . Furthermore, the transition probability  $P^{w, \pi}(s_1|s_0, a) = 1 - P^{w, \pi}(s_2|s_0, a) = \beta + \frac{1}{3}, \forall a$ , where  $\beta \in (0, \frac{2}{3})$  is a hyper-parameter. We can easily verify that all those transitions are within  $[0, 1]$  and therefore well defined. With this analysis, we can calculate the visitation measure for each policy in Table 1 and derive  $C_{vr} = 3 + \frac{1}{\beta}$ .

**Implementation** Under the optimal policy (taking  $a = 0$  at  $s = s_1$ ), the expected reward is  $\mathbb{E}_{\pi^*}[r] = \frac{1}{2}$ , regardless of  $\beta$ . We set the number of episodes  $K = 1000$  to simulate a scenario with limited exploration and run Algorithm 1 in the CRMDP-TV setting. We test learned robust policies in the worst-case target environment and calculate the average reward among 2000 runs. The experimental results are based on 50 replications and plotted in Figure 1(a).

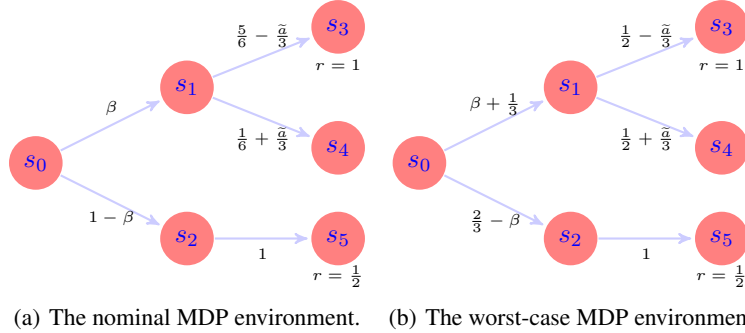


Figure 4: The constructions of the nominal MDP and the worst-case MDP environments in Appendix E.1.

Table 1: The visitation measure of each state in Appendix E.1, the maximum of  $\frac{q_h^\pi(s)}{d_h^\pi(s)}$  is achieved by taking  $\tilde{a} = 0$  at  $s_4$ .

|              | $s_0$ | $s_1$                 | $s_2$                 | $s_3$  | $s_4$  | $s_5$                 |
|--------------|-------|-----------------------|-----------------------|--|--|-----------------------|
| $d_h^\pi(s)$ | 1     | $\beta$               | $1 - \beta$           | $(\frac{5}{6} - \frac{\tilde{a}}{3})\beta$                 | $(\frac{1}{6} + \frac{\tilde{a}}{3})\beta$                 | $1 - \beta$           |
| $q_h^\pi(s)$ | 1     | $\beta + \frac{1}{3}$ | $\frac{2}{3} - \beta$ | $(\frac{1}{2} - \frac{\tilde{a}}{3})(\beta + \frac{1}{3})$ | $(\frac{1}{2} + \frac{\tilde{a}}{3})(\beta + \frac{1}{3})$ | $\frac{2}{3} - \beta$ |

## F.2 More Details on Appendix E.2 (Learning on Simulated RMDPs)

**Construction** We consider a simple MDP Figure 5(a) as the source environment. The learning horizon  $H = 3$ , the state space is  $\mathcal{S} = \{s_0, \dots, s_4\}$ , and the action space is  $\mathcal{A} = \{0, \dots, 4\}$ . The initial state is always  $s_0$ , where it can transit to  $s_1, s_3$  and  $s_4$  with probability  $P^o(s_1|s_0, a) = 0.4 + \frac{a}{10}$ ,



$P^o(s_3|s_0, a) = 0.1$  and  $P^o(s_4|s_0, a) = 0.5 - \frac{a}{10}$  correspondingly. From  $s_1$ , it can transit to  $s_2$  and  $s_3$  with probability  $P^o(s_2|s_1, a) = \frac{a}{10}$  and  $P^o(s_3|s_1, a) = 1 - \frac{a}{10}$  by taking action  $a$ . From  $s_2$ , it can transit to  $s_3$  and  $s_4$  with probability  $P^o(s_3|s_2, a) = 1 - \frac{a}{10}$  and  $P^o(s_4|s_2, a) = \frac{a}{10}$  by taking action  $a$ . The  $s_3$  and  $s_4$  are absorbing states. The agent is rewarded  $\frac{a}{20}$  by taking action  $a$  at  $s_0, s_1$  and  $s_2$ , rewarded 1 regardless of the action taken at  $s_4$ , and rewarded 0 regardless of the action taken at  $s_3$ .

The target environment Figure 5(b) is obtained by perturbing the first step in the source environment. To be specific, with perturbation rate  $q$ , the transition probability from  $s_0$  to  $s_3$  and  $s_4$  is  $P^w(s_3|s_0, a) = 0.1 + q \times (0.5 - \frac{a}{10})$  and  $P^w(s_4|s_0, a) = (1 - q) \times (0.5 - \frac{a}{10})$ , while  $P^w(s_1|s_0, a) = 0.4 + \frac{a}{10}$  stays the same.

**Implementation** We set  $K = 1,000$  in Algorithm 1 and evaluate the learned policy in target environments with  $q \in \{0, 0.05, 0.1, \dots, 1\}$ , respectively. In each target environment, the average reward among 500 runs is calculated for evaluation. All experimental results are based on 20 replications. The choice of uncertainty level  $\rho$ , regularizer  $\beta$ , and constant  $c_{\text{bonus}}$  are provided in Table 2. For the non-robust algorithm, we simply set  $\beta = 10000$  for Algorithm 1 instantiated with TV-distance defined regularization. It can be justified by Lemma I.1 that the extremely large regularization would not tolerate any perturbation, and thus the learned policy is basically the optimal policy under the source environment. The results are plotted in Figures 1(b) to 1(d).

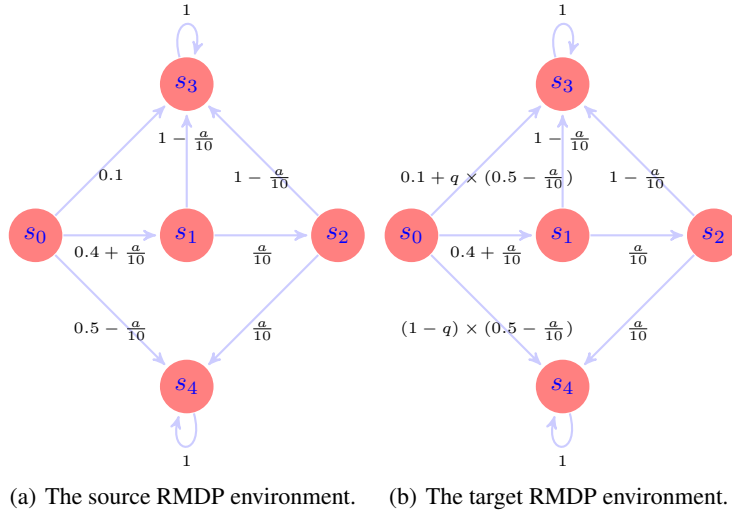


Figure 5: The source and target RMDP environments in Appendix E.2, where the target environment is constructed by perturbing the first step.

### F.3 More Details on Appendix E.3 (Learning the Frozen Lake Problem)

**Construction** In this scenario, the agent’s objective is to traverse a frozen lake from the Start (S) to the Goal (G) without falling into any Holes (H), navigating over the Frozen (F) surface. The agent’s movement is influenced by a hyper-parameter  $P_{\text{slip}} = 0.1$ , which determines the probability of the agent successfully moving in the intended direction. Specifically, the agent moves towards the intended direction with a probability of  $1 - P_{\text{slip}} = 0.9$ , and with probability  $P_{\text{slip}}/2 = 0.05$ , it will veer off in either perpendicular direction. The Goal state is an absorbing state, once reached the agent will stay there. The agent earns a reward of 1 if and only if it is at the goal (G) at step  $H - 1$ . For evaluation, after the agent selects an intended action, with a probability of  $P_{\text{perturb}}$ , the agent actually takes action towards the opposite direction instead. The difficulty of this environment arises from two sources: 1) it is a sparse reward MDP, and 2) the influence onto the movement in the source environment makes the exploration of the goal state very hard.

**Implementation** We use the default map in the OpenAI Gym library, which is illustrated in Example F.1, and set  $H = 25$  and  $K = 1,000$  in Algorithm 1. The hyperparameter  $\rho$  in the constrained setting,  $\beta$  in the regularized setting, and  $c_{\text{bonus}}$  are tuned from  $\{0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 1\}$ , with the final choice presented in Table 2. All experimental results are based on 20 replications.

We implement several optimizations to accelerate the algorithm. First, we reuse the counts  $n_h^k(s, a)$  from (C.1) across different steps, as the transition dynamics remain consistent throughout. Additionally, we observe that the inner optimization process, which is the primary bottleneck of the algorithm, does not need to be fully recalculated at each episode. For instance, if the agent falls into a hole after the first step, the information about the environment remains largely unchanged. To take advantage of this, we create a map where the keys are hashes of the optimization parameters and the values are the corresponding results, allowing us to reuse prior computations efficiently. We also applied Algorithm 2 to obtain the result in the CRMDP-TV setting, which can directly be derived from definition.

We assess the convergence of our algorithms by calculating the average reward among 500 runs in the single target environment with  $P_{\text{perturb}} = 0.1$ , of the policy  $\pi^k$  obtained after each episode  $k$  during the training. The convergence results are plotted in Figures 2(a) and 2(b), and the average training time is plotted in Figure 3(a). We also evaluate the robustness of the policy  $\pi^K$  after  $K$  episodes across various target environments with  $P_{\text{perturb}} \in \{0, 0.05, 0.1, \dots, 0.3\}$ . For each target environment, we compute the average reward among 500 runs, and the results are shown in Figures 3(b) to 3(d).

**Example F.1** (Illustration of Frozen Lake environment). *The environment of the Frozen Lake problem is illustrated as follows, where  $S$  denotes “Start”,  $G$  denotes “Goal”,  $H$  denotes “Hole” and  $F$  denotes “Frozen”.*

$$\begin{bmatrix} S & F & F & F & F & F & F & F \\ F & F & F & F & F & F & F & F \\ F & F & F & H & F & F & F & F \\ F & F & F & F & F & H & F & F \\ F & F & F & H & F & F & F & F \\ F & H & H & F & F & F & H & F \\ F & H & F & F & H & F & H & F \\ F & F & F & H & F & F & F & G \end{bmatrix}$$

#### F.4 A More Computationally Efficient Solver for the CRMDP-TV Setting

As shown in Lemma H.1, the update formulation of Algorithm 1 in the CRMDP-TV setting involves solving an optimization problem in its dual formulation. To reduce the computational complexity, we introduce Algorithm 2, which simplifies this procedure.

We explain the rationale behind Algorithm 2 as follows. The original formulation of the CRMDP-TV problem is given by  $Q_h(s, a) = r_h(s, a) + \inf_{\text{TV}(P \| P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}](s, a)$ . It is easy to see that the worst-case scenario is reached when the transition probabilities for states with the highest value functions are reduced by a total of  $\rho$ , and those for states with the lowest value functions are increased by  $\rho$ . This greedy approach avoids the need to solve the optimization problem for  $\eta$  as described in Lemma H.1.

#### F.5 Hyper-parameters for Experiments in Appendix E

Here, we provide the hyper-parameters used in training in the experiments section. Note that we reformulate the bonus term as  $c_{\text{bonus}}/\sqrt{K}$  in practical experiments.

---

**Algorithm 2** A more efficient solver for the CRMDP-TV setting

---

**Require:** robust set radius  $\rho > 0$ , transition array  $P[S]$ , value function array  $V[S]$  ( $S > 1$ )

```

1: args = argsort(V).
2:  $V = V[\text{args}]$ ,  $P = P[\text{args}]$ .
3:  $\text{pnt}_0 = 1$ ,  $\text{pnt}_1 = S$ .
4:  $\text{rho}_0 = 0$ ,  $\text{rho}_1 = 0$ .
5: while  $\text{rho}_0 < \rho$  do
6:    $\text{tmp} = \min(\rho - \text{rho}_0, 1 - P[\text{pnt}_0])$ .
7:    $P[\text{pnt}_0] = P[\text{pnt}_0] + \text{tmp}$ .
8:    $\text{rho}_0 = \text{rho}_0 + \text{tmp}$ .
9:    $\text{pnt}_0 = \text{pnt}_0 + 1$ .
10: end while
11: while  $\text{rho}_1 < \rho$  do
12:    $\text{tmp} = \min(\rho - \text{rho}_1, P[\text{pnt}_1])$ .
13:    $P[\text{pnt}_1] = P[\text{pnt}_1] - \text{tmp}$ .
14:    $\text{rho}_1 = \text{rho}_1 + \text{tmp}$ .
15:    $\text{pnt}_1 = \text{pnt}_1 - 1$ .
16: end while
17: Output:  $\sum_{i=1}^S P[i] * V[i]$ .
```

---

Table 2: hyper-parameters for Appendix E.2 (Learning on Simulated RMDPs)

| Setting              | $\rho$ or $\beta$ | $c_{\text{bonus}}$ |
|----------------------|-------------------|--------------------|
| non-robust           | —                 | 1                  |
| constrained TV       | 0.5               | 1                  |
| constrained KL       | 0.5               | 1                  |
| constrained $\chi^2$ | 1                 | 1                  |
| regularized TV       | 0.1               | 1                  |
| regularized KL       | 0.1               | 1                  |
| regularized $\chi^2$ | 0.1               | 1                  |

## G Proof of Proposition D.4

*Proof of Proposition D.4.* We prove it by contradiction. We assume that there exists  $s^*$ ,  $a^*$  and  $s'$  such that  $V_{h+1}^{\pi, \rho}(s') > 0$  and  $P_h^{w, \pi}(s'|s^*, a^*) > P_h^o(s'|s^*, a^*)$ . We pick  $\tilde{s} \in \mathcal{S}_f$  arbitrarily and consider the following transition measure  $P'_h$ :

$$P'_h(s|s^*, a^*) = \begin{cases} P_h^{w, \pi}(s|s^*, a^*) & s \notin \{s', \tilde{s}\}, \\ P_h^o(s'|s^*, a^*) & s = s', \\ P_h^{w, \pi}(\tilde{s}|s^*, a^*) + P_h^{w, \pi}(s'|s^*, a^*) - P_h^o(s'|s^*, a^*) & s = \tilde{s}. \end{cases}$$

It is easy to verify that  $P'_h \in \Delta(\mathcal{S})$ ,  $\text{TV}(P'_h \| P_h^o) \leq \text{TV}(P_h^{w, \pi} \| P_h^o)$ , therefore  $P'_h$  is a valid transition measure in the transition uncertainty set. Based on the definition of fail-states, we have  $V_{h+1}^{\pi, \rho}(\tilde{s}) = 0 < V_{h+1}^{\pi, \rho}(s')$  and thus  $\mathbb{E}_{P'_h}[V_{h+1}^{\pi, \rho}] < \mathbb{E}_{P_h^{w, \pi}}[V_{h+1}^{\pi, \rho}]$ , which contradicts the fact that  $P_h^{w, \pi}$  is the worst-case transition.  $\square$

## H Proofs of Results in Constrained RMDPs

### H.1 Proof of Theorem 3.2 (Constrained TV Setting)

Before proving this theorem, we first present several technical lemmas that will be useful in the proof. For convenience, we also write  $P^{w, k} := P^{w, \pi^k}$ ,  $d^k := d^{\pi^k}$  and  $q^k := q^{\pi^k}$ .

Table 3: hyper-parameters for Appendix E.3 (Learning the Frozen Lake Problem)

| Setting              | $\rho$ or $\beta$ | $c_{\text{bonus}}$ |
|----------------------|-------------------|--------------------|
| non-robust           | –                 | 0.001              |
| constrained TV       | 0.15              | 0.001              |
| constrained KL       | 0.15              | 0.01               |
| constrained $\chi^2$ | 0.5               | 0.01               |
| regularized TV       | 0.1               | 0.003              |
| regularized KL       | 0.1               | 0.001              |
| regularized $\chi^2$ | 0.05              | 0.01               |

We first give the closed form solution of constrained TV update formulation. This dual formulation has also been proved by [12, Lemma 4.3], but the formulation of our result is slightly different from theirs. Note that [19] used the same formulation, but their expression is incorrect by a factor of  $\frac{1}{2}$ . This error arises because they directly cited [43], which employed the  $L_1$  distance rather than the TV-distance. So we prove it again for the sake of completeness.

**Lemma H.1** (Dual formulation). *For the optimization problem  $Q_h(s, a) = r_h(s, a) + \inf_{\text{TV}(P\|P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}](s, a)$ , we have its dual formulation as follows*

$$Q_h = r_h - \inf_{\eta \in [0, H]} \left( \mathbb{E}_{P_h^o}[(\eta - V_{h+1}(s))_+] + \rho \left( \eta - \min_{s \in \mathcal{S}} V_{h+1}(s) \right)_+ - \eta \right). \quad (\text{H.1})$$

*Proof.* Consider the optimization problem

$$Q_h = r_h + \inf_{\text{TV}(P\|P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}] = r_h + \inf_{\text{TV}(P\|P_h^o) \leq \rho} \sum_{s \in \mathcal{S}} P(s) V_{h+1}(s).$$

We denote  $\varphi(t) = |t - 1|/2$ , then the Lagrangian can be written as

$$\mathcal{L}(P, \eta) = \sum_{s \in \mathcal{S}} P(s) V_{h+1}(s) + \nu \left( \sum_{s \in \mathcal{S}} P_h^o(s) \varphi \left( \frac{P(s)}{P_h^o(s)} \right) - \rho \right) - \sum_{s \in \mathcal{S}} \lambda(s) P(s) + \eta \left( 1 - \sum_{s \in \mathcal{S}} P(s) \right).$$

We denote  $g(s) = P(s)/P_h^o(s)$ , then we have

$$\begin{aligned} \inf_P \mathcal{L}(P, \eta) &= - \sup_g \sum_{s \in \mathcal{S}} P_h^o(s) [g(s) [(\lambda(s) + \eta) - V_{h+1}(s)] - \nu \varphi(g(s))] - \nu \rho + \eta \\ &= - \nu \mathbb{E}_{P_h^o} \left[ \varphi^* \left( \frac{(\lambda(s) + \eta) - V_{h+1}(s)}{\nu} \right) \right] - \nu \rho + \eta, \end{aligned}$$

where the second equation is from the definition of dual function  $\varphi^*(y) = \sup_x (y^\top x - \varphi(x))$ . From  $f$ -divergence literature [41], we know that for TV-distance,

$$\varphi^*(s) = \begin{cases} -\frac{1}{2} & \text{for } s \leq -\frac{1}{2}; \\ s & \text{for } -\frac{1}{2} < s \leq \frac{1}{2}; \\ +\infty & \text{for } s > \frac{1}{2}. \end{cases} \quad (\text{H.2})$$

So we have

$$\begin{aligned}
Q_h &= r_h + \sup_{\nu \geq 0, \lambda \geq 0, \eta} \left( \inf_P \mathcal{L}(P, \eta) \right) \\
&= r_h - \inf_{\nu \geq 0, \lambda \geq 0, \eta} \left( \nu \mathbb{E}_{P_h^\circ} \left[ \varphi^* \left( \frac{(\lambda(s) + \eta) - V_{h+1}(s)}{\nu} \right) \right] + \nu \rho - \eta \right) \\
&= r_h - \inf_{\nu \geq 0, \lambda \geq 0, \eta, \frac{\lambda(s) + \eta - V_{h+1}(s)}{\nu} \leq \frac{1}{2}} \left( \nu \mathbb{E}_{P_h^\circ} \left[ \max \left( \frac{(\lambda(s) + \eta) - V_{h+1}(s)}{\nu}, -\frac{1}{2} \right) \right] + \nu \rho - \eta \right) \tag{H.3}
\end{aligned}$$

$$\begin{aligned}
&= r_h - \inf_{\nu \geq 0, \lambda \geq 0, \eta, \frac{\lambda(s) + \eta - V_{h+1}(s)}{\nu} \leq \frac{1}{2}} \left( \mathbb{E}_{P_h^\circ} \left[ \left( (\lambda(s) + \eta) - V_{h+1}(s) + \frac{\nu}{2} \right)_+ \right] + \nu \rho - \eta - \frac{\nu}{2} \right) \\
&= r_h - \inf_{\nu \geq 0, \lambda \geq 0, \eta', \nu \geq \lambda(s) + \eta' - V_{h+1}(s)} \left( \mathbb{E}_{P_h^\circ} [(\lambda(s) + \eta' - V_{h+1}(s))_+] + \nu \rho - \eta' \right) \tag{H.4}
\end{aligned}$$

$$= r_h - \inf_{\lambda \geq 0, \eta'} \left( \mathbb{E}_{P_h^\circ} [(\lambda(s) + \eta' - V_{h+1}(s))_+] + \rho \max_{s \in \mathcal{S}} (\lambda(s) + \eta' - V_{h+1}(s))_+ - \eta' \right) \tag{H.5}$$

$$= r_h - \inf_{\eta'} \left( \mathbb{E}_{P_h^\circ} [(\eta' - V_{h+1}(s))_+] + \rho \left( \eta' - \min_{s \in \mathcal{S}} V_{h+1}(s) \right)_+ - \eta' \right) \tag{H.6}$$

$$= r_h - \inf_{\eta' \in [0, H]} \left( \mathbb{E}_{P_h^\circ} [(\eta' - V_{h+1}(s))_+] + \rho \left( \eta' - \min_{s \in \mathcal{S}} V_{h+1}(s) \right)_+ - \eta' \right), \tag{H.7}$$

where (H.3) follows from the definition of  $\varphi^*$  (H.2), we redefine  $\eta' = \eta + \frac{\nu}{2}$  in (H.4), (H.5) holds because the result increases monotonically with respect to  $\nu$  thus the minimum value is attained at  $\nu = \max_{s \in \mathcal{S}} (\lambda(s) + \eta' - V_{h+1}(s))_+$ , (H.6) holds because the result increases monotonically with respect to  $\lambda$ , (H.7) holds because the result increases monotonically with respect to  $\eta'$  when  $\eta' \leq 0$  and increases monotonically with respect to  $\eta'$  when  $\eta' \geq H$ .  $\square$

In the next lemma, we prove the optimism of estimation  $Q^k$ , which helps control the estimation error  $Q^* - Q^k$ .

**Lemma H.2** (Optimism). *If we set the bonus term as follows*

$$\text{bonus}_h^k(s, a) = 2H \sqrt{\frac{2S^2 \ln(12SAH^2K^2/\delta)}{n_h^{k-1}(s, a) \vee 1}} + \frac{1}{K}, \tag{H.8}$$

then for any policy  $\pi$  and any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , with probability at least  $1 - 2\delta$ , we have  $Q_h^{k, \rho}(s, a) \geq Q_h^{\pi, \rho}(s, a)$ . Specially, by setting  $\pi = \pi^*$ , we have  $Q_h^{k, \rho}(s, a) \geq Q_h^{\pi^*, \rho}(s, a)$ .

*Proof.* We prove this by induction. First, when  $h = H + 1$ ,  $Q_{H+1}^{k, \rho}(s, a) = 0 = Q_{H+1}^{\pi, \rho}(s, a)$  holds trivially.

Assume  $Q_{h+1}^{k, \rho}(s, a) \geq Q_{h+1}^{\pi, \rho}(s, a)$  holds, since  $\pi^k$  is the greedy policy, we have

$$V_{h+1}^{k, \rho}(s) = Q_{h+1}^{k, \rho}(s, \pi_{h+1}^k(s)) \geq Q_{h+1}^{k, \rho}(s, \pi_{h+1}(s)) \geq Q_{h+1}^{\pi, \rho}(s, \pi_{h+1}(s)) = V_{h+1}^{\pi, \rho}(s),$$

where the first inequality is because we choose  $\pi^k$  as the greedy policy.

Recall that we denote  $Q_h^{k, \rho}$  as the optimistic estimation in  $k$ -th episode, that is,

$$Q_h^{k, \rho}(s, a) = \min \left\{ \text{bonus}_h^k(s, a) + \hat{r}_h^k(s, a) + \inf_{\text{TV}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P [V_{h+1}^{k, \rho}](s, a), H - h + 1 \right\}.$$

If  $Q_h^{k, \rho}(s, a) = H - h + 1$ , then it follows immediately that

$$Q_h^{k, \rho}(s, a) = H - h + 1 \geq Q_h^{\pi, \rho}(s, a)$$

by the definition of  $Q_h^{\pi,\rho}(s, a)$ . Otherwise, we can infer that

$$\begin{aligned}
Q_h^{k,\rho} - Q_h^{\pi,\rho} &= \text{bonus}_h^k + \hat{r}_h^k + \inf_{\text{TV}(P\|\hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - r_h - \inf_{\text{TV}(P\|P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi,\rho}] \\
&= \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\text{TV}(P\|\hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - \inf_{\text{TV}(P\|P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] \\
&\quad + \inf_{\text{TV}(P\|P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - \inf_{\text{TV}(P\|P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi,\rho}] \\
&\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\text{TV}(P\|\hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - \inf_{\text{TV}(P\|P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] \quad (\text{H.9}) \\
&= \text{bonus}_h^k + \hat{r}_h^k - r_h - \inf_{\eta \in [0, H]} \left( \mathbb{E}_{\hat{P}_h^k}[(\eta - V_{h+1}^{k,\rho}(s))_+] + \rho \left( \eta - \min_{s \in \mathcal{S}} V_{h+1}^{k,\rho}(s) \right)_+ - \eta \right) \\
&\quad + \inf_{\eta \in [0, H]} \left( \mathbb{E}_{P_h^o}[(\eta - V_{h+1}^{k,\rho}(s))_+] + \rho \left( \eta - \min_{s \in \mathcal{S}} V_{h+1}^{k,\rho}(s) \right)_+ - \eta \right) \quad (\text{H.10}) \\
&\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\eta \in [0, H]} \left\{ \left( \mathbb{E}_{P_h^o}[(\eta - V_{h+1}^{k,\rho}(s))_+] + \rho \left( \eta - \min_{s \in \mathcal{S}} V_{h+1}^{k,\rho}(s) \right)_+ - \eta \right) \right. \\
&\quad \left. - \left( \mathbb{E}_{\hat{P}_h^k}[(\eta - V_{h+1}^{k,\rho}(s))_+] + \rho \left( \eta - \min_{s \in \mathcal{S}} V_{h+1}^{k,\rho}(s) \right)_+ - \eta \right) \right\} \\
&= \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\eta \in [0, H]} \left\{ \mathbb{E}_{P_h^o}[(\eta - V_{h+1}^{k,\rho}(s))_+] - \mathbb{E}_{\hat{P}_h^k}[(\eta - V_{h+1}^{k,\rho}(s))_+] \right\} \\
&\geq \text{bonus}_h^k - \underbrace{|\hat{r}_h^k - r_h|}_{(i)} - \underbrace{\sup_{\eta \in [0, H]} \left| \mathbb{E}_{P_h^o}[(\eta - V_{h+1}^{k,\rho}(s))_+] - \mathbb{E}_{\hat{P}_h^k}[(\eta - V_{h+1}^{k,\rho}(s))_+] \right|}_{(ii)}, \quad (\text{H.11})
\end{aligned}$$

where (H.9) is from the induction assumption, we plug in the dual formulation (H.1) in (H.10).

For term (i) in (H.11), from Lemma L.1 and a union bound, with probability at least  $1 - \delta$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| \leq \sqrt{\frac{\ln(2SAHK/\delta)}{2n_h^{k-1}(s, a) \vee 1}}, \quad (\text{H.12})$$

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

We denote  $V(\eta) = (\eta - V_{h+1}^{k,\rho}(s))_+ \in [0, H]$  and  $\mathcal{V} = \{V \in \mathbb{R}^S : \|V\|_\infty \leq H\}$ . To bound term (ii) in (H.11), we create a  $\epsilon$ -net  $\mathcal{N}_\mathcal{V}(\epsilon)$  for  $\mathcal{V}$ . From Lemma L.4, it holds that  $\ln |\mathcal{N}_\mathcal{V}(\epsilon)| \leq |S| \cdot \ln(3H/\epsilon)$ .

Therefore, by the definition of  $\mathcal{N}_\mathcal{V}(\epsilon)$ , for any fixed  $V$ , there exists a  $V' \in \mathcal{N}_\mathcal{V}(\epsilon)$  such that  $\|V - V'\|_\infty \leq \epsilon$ , that is

$$\begin{aligned}
|\mathbb{E}_{P_h^o}[V] - \mathbb{E}_{\hat{P}_h^k}[V]| &\leq |\mathbb{E}_{P_h^o}[V] - \mathbb{E}_{P_h^o}[V']| + |\mathbb{E}_{P_h^o}[V'] - \mathbb{E}_{\hat{P}_h^k}[V']| + |\mathbb{E}_{\hat{P}_h^k}[V'] - \mathbb{E}_{\hat{P}_h^k}[V]| \\
&\leq \|P_h^o\|_1 \|V - V'\|_\infty + |\mathbb{E}_{P_h^o}[V'] - \mathbb{E}_{\hat{P}_h^k}[V']| + \|\hat{P}_h^k\|_1 \|V - V'\|_\infty \\
&\leq \sup_{V' \in \mathcal{N}_\mathcal{V}(\epsilon)} |\mathbb{E}_{P_h^o}[V'] - \mathbb{E}_{\hat{P}_h^k}[V']| + 2\epsilon, \quad (\text{H.13})
\end{aligned}$$

where the second inequality follows from the Holder's inequality.

For any fixed  $V$ , we apply Lemma L.3 and have

$$|\mathbb{E}_{P_h^o}[V] - \mathbb{E}_{\hat{P}_h^k}[V]| \leq \|P_h^o - \hat{P}_h^k\|_1 \cdot \|V\|_\infty \leq H \sqrt{\frac{2S \ln(2/\delta)}{n_h^{k-1} \vee 1}}, \quad (\text{H.14})$$

with probability at least  $1 - \delta$ .

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

$$\begin{aligned} \sup_{\eta \in [0, H]} |\mathbb{E}_{P_h^o}[(\eta - V_{h+1}^{k, \rho}(s))_+] - \mathbb{E}_{\hat{P}_h^k}[(\eta - V_{h+1}^{k, \rho}(s))_+]| &\leq \sup_{\eta \in [0, H]} |\mathbb{E}_{P_h^o}[V(\eta)] - \mathbb{E}_{\hat{P}_h^k}[V(\eta)]| \\ &\leq \sup_{V \in \mathcal{N}_V(\epsilon)} |\mathbb{E}_{P_h^o}[V] - \mathbb{E}_{\hat{P}_h^k}[V]| + 2\epsilon \end{aligned} \quad (\text{H.15})$$

$$\leq H \sqrt{\frac{2S \ln(2SAHK|\mathcal{N}_V(\epsilon)|/\delta)}{n_h^{k-1} \vee 1}} + 2\epsilon \quad (\text{H.16})$$

$$\begin{aligned} &\leq H \sqrt{\frac{2S^2 \ln(6SAH^2K/\epsilon\delta)}{n_h^{k-1} \vee 1}} + 2\epsilon \\ &= H \sqrt{\frac{2S^2 \ln(12SAH^2K^2/\delta)}{n_h^{k-1} \vee 1}} + \frac{1}{K}. \end{aligned} \quad (\text{H.17})$$

for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , where (H.15) follows from (H.13), (H.16) is from (H.14) and a union bound, we set  $\epsilon = 1/2K$  in (H.17).

Apply the union bound again and combine (H.11) with (H.12), (H.17), the definition of bonus and induction assumption. With probability at least  $1 - 2\delta$ , we have  $Q_h^{k, \rho}(s, a) \geq Q_h^{\pi, \rho}(s, a)$  for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ . This completes the proof.  $\square$

**Lemma H.3.** *With probability at least  $1 - \delta$ , for any  $(k, s, a) \in [K] \times \mathcal{S} \times \mathcal{A}$ , the sum of estimation errors can be bounded as follows*

$$Q_1^{k, \rho} - Q_1^{\pi^k, \rho} \leq \sum_{h=1}^H \mathbb{E}_{\{P_h^{w, k}\}_{h=1}^H, \pi^k} [2\text{bonus}_h^k].$$

*Proof.* From the proof of Lemma H.2, we see that with probability at least  $1 - \delta$ , for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| + \left| \inf_{\text{TV}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}](s, a) - \inf_{\text{TV}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}](s, a) \right| \leq \text{bonus}_h^k(s, a). \quad (\text{H.18})$$

Recall that we define  $P_h^{w, k} = \underset{\text{TV}(P \parallel P_h^o) \leq \rho}{\text{argmin}} \mathbb{E}_P[V_{h+1}^{\pi^k, \rho}]$  as the worst-case transition in Definition D.3,

we have

$$\begin{aligned} Q_h^{k, \rho} - Q_h^{\pi^k, \rho} &\leq \text{bonus}_h^k + \hat{r}_h^k + \inf_{\text{TV}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] - r_h - \inf_{\text{TV}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi^k, \rho}] \\ &= \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\text{TV}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] - \inf_{\text{TV}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] \\ &\quad + \inf_{\text{TV}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] - \inf_{\text{TV}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi^k, \rho}] \\ &\leq 2\text{bonus}_h^k + \inf_{\text{TV}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] - \inf_{\text{TV}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi^k, \rho}] \end{aligned} \quad (\text{H.19})$$

$$\leq 2\text{bonus}_h^k + \mathbb{E}_{P_h^{w, k}}[V_{h+1}^{k, \rho} - V_{h+1}^{\pi^k, \rho}] \quad (\text{H.20})$$

$$= 2\text{bonus}_h^k + \mathbb{E}_{P_h^{w, k}, \pi^k}[Q_{h+1}^{k, \rho} - Q_{h+1}^{\pi^k, \rho}], \quad (\text{H.21})$$

where the (H.19) holds because of (H.18), (H.20) and (H.21) use the definition of  $P_h^{w, k}$  and  $\pi^k$  accordingly. Apply (H.21) recursively, we can obtain the result.  $\square$

Next, in Lemma H.9, we provide an upper bound on the sum of the expectations of  $\sqrt{\frac{1}{n_h^{k-1} \vee 1}}$  under the worst-case environment. In order to prove this lemma, we follow a similar procedure to that of [44], whose setting differs from ours as we consider the non-stationary dynamics.

**Lemma H.4.** (*Failure Events*) We define the following failure events:

$$F_k = \left\{ \exists s, a, h : n_h^{k-1}(s, a) \leq \frac{1}{2} \sum_{i < k} d_h^i(s, a) - \ln \left( \frac{SAHK}{\delta} \right) \right\}.$$

Then we have  $P\left(\bigcup_{k=1}^K F_k\right) \leq 1 - \delta$ .

*Proof.* Consider a fixed  $s \in \mathcal{S}, a \in \mathcal{A}, h \in [H]$ . We define  $\mathcal{F}_k$  to be the  $\sigma$ -field induced by the first  $k - 1$  episodes and  $X_k$  to be the indicator whether  $(s, a)$  was visited in episode  $k$  at step  $h$ . The probability  $\mathbb{P}(s = s_h^k, a = a_h^k | \pi_k)$  of whether  $X_k = 1$  is  $\mathcal{F}_k$ -measurable, therefore we can apply Lemma L.5 with  $W = \ln(\frac{SAHK}{\delta})$ . The proof is finished by applying a union bound over  $s, a, h, k$ .  $\square$

**Definition H.5.** (*The Good Set*) We define

$$L_k = \left\{ (s, a, h) : \frac{1}{4} \sum_{i < k} d_h^i(s, a) \geq \ln \left( \frac{SAHK}{\delta} \right) + 1 \right\}.$$

**Lemma H.6.** (*Visitation Ratio*) Outside the failure event, if  $(s, a, h) \in L_k$ , we have

$$n_h^{k-1}(s, a) \geq \frac{1}{4} \sum_{i \leq k} d_h^i(s, a).$$

*Proof.* Outside the failure events defined in Lemma H.4, we have

$$\begin{aligned} n_h^{k-1}(s, a) &> \frac{1}{2} \sum_{i < k} d_h^i(s, a) - \ln \left( \frac{SAHK}{\delta} \right) \\ &= \frac{1}{4} \sum_{i < k} d_h^i(s, a) + \frac{1}{4} \sum_{i < k} d_h^i(s, a) - \ln \left( \frac{SAHK}{\delta} \right) \\ &\geq \frac{1}{4} \sum_{i < k} d_h^i(s, a) + 1 \geq \frac{1}{4} \sum_{i < k} d_h^i(s, a) + d_h^k(s, a) \geq \frac{1}{4} \sum_{i \leq k} d_h^i(s, a). \end{aligned}$$

where the second inequality uses  $(s, a, h) \in L_k$  and the definition of  $L_k$  in Definition H.5.  $\square$

**Lemma H.7.** (*Minimal Contribution*) Outside the failure event, we have

$$\sum_{k=1}^K \sum_{(s, a, h) \notin L_k} d_h^k(s, a) = \tilde{\mathcal{O}}(SAH).$$

*Proof.* We have

$$\begin{aligned} \sum_{k=1}^K \sum_{(s, a, h) \notin L_k} d_h^k(s, a) &= \sum_{(s, a, h)} \sum_{k=1}^K d_h^k(s, a) \mathbb{1}\{(s, a, h) \notin L_k\} \\ &\leq \sum_{(s, a, h)} \left( \sum_{k < K} d_h^k(s, a) \mathbb{1}\{(s, a, h) \notin L_k\} + 1 \right) \\ &< \sum_{(s, a, h)} \left( 4 \ln \left( \frac{SAHK}{\delta} \right) + 5 \right) \\ &= \tilde{\mathcal{O}}(SAH). \end{aligned}$$

where the first inequality uses the definition of  $L_k$  in Definition H.5.  $\square$



**Lemma H.8.** (*Visitation Ratio*) Outside the failure event, it holds that

$$\sum_{k=1}^K \sum_{(s,a,h) \in L_k} \frac{d_h^k(s,a)}{n_h^{k-1}(s,a)} = \tilde{\mathcal{O}}(SAH).$$

*Proof.* Outside the failure events defined in Lemma H.4, we have

$$\begin{aligned} \sum_{k=1}^K \sum_{(s,a,h) \in L_k} \frac{d_h^k(s,a)}{n_h^{k-1}(s,a)} &= \sum_{k=1}^K \sum_{(s,a,h)} \frac{d_h^k(s,a)}{n_h^{k-1}(s,a)} \mathbb{1}\{(s,a,h) \in L_k\} \\ &\leq 4 \sum_{k=1}^K \sum_{(s,a,h)} \frac{d_h^k(s,a)}{\sum_{i \leq k} d_h^i(s,a)} \mathbb{1}\{(s,a,h) \in L_k\}, \end{aligned}$$

where the inequality follows from Lemma H.6.

Next, for any fixed  $(h, s, a) \in L_k$  for some  $k_0$ , since  $\sum_{i < k} d_h^i(s, a)$  is strictly increasing with  $k$ , there exists a critical episode  $\tilde{k} \leq k_0$  such that  $(h, s, a) \in L_k$  holds for all  $k \geq \tilde{k}$  and  $(h, s, a) \notin L_k$  holds for all  $k < \tilde{k}$ . From the definition of  $\tilde{k}$  and Definition H.5, we know  $\sum_{i < \tilde{k}} d_h^i(s, a) \geq 4 \ln(\frac{SAHK}{\delta}) + 4 > 4$ . Therefore,

$$\begin{aligned} \sum_{k=1}^K \frac{d_h^k(s,a)}{\sum_{i \leq k} d_h^i(s,a)} \mathbb{1}\{(s,a,h) \in L_k\} &= \sum_{k=1}^K \frac{d_h^k(s,a)}{\sum_{i < \tilde{k}} d_h^i(s,a) + \sum_{\tilde{k} \leq i \leq k} d_h^i(s,a)} \mathbb{1}\{(s,a,h) \in L_k\} \\ &\leq \sum_{k=1}^K \frac{d_h^k(s,a)}{4 + \sum_{\tilde{k} \leq i \leq k} d_h^i(s,a)} \mathbb{1}\{(s,a,h) \in L_k\} \\ &\leq \sum_{\tilde{k} \leq k \leq K} \frac{d_h^k(s,a)}{4 + \sum_{\tilde{k} \leq i \leq k} d_h^i(s,a)}, \end{aligned}$$

where the third inequality comes from the definition of  $\tilde{k}$ .

To simplify the notations, we define  $v_1 = d_h^{\tilde{k}}(s, a)$ ,  $v_2 = d_h^{\tilde{k}+1}(s, a)$ ,  $\dots$ ,  $v_{K-\tilde{k}+1} = d_h^K(s, a)$ . And in order to bound the above summation, we also define the functions  $F(x) = \sum_{i=1}^{\lfloor x \rfloor} v_i + v_{\lceil x \rceil}(x - \lfloor x \rfloor)$  and  $f(x) = v_{\lceil x \rceil}$ . It is easy to verify that the derivative of  $F(x)$  is  $f(x)$ . Then we write

$$\sum_{\tilde{k} \leq k \leq K} \frac{d_h^k(s,a)}{4 + \sum_{\tilde{k} \leq i \leq k} d_h^i(s,a)} = \sum_{k=1}^{K-\tilde{k}+1} \frac{v_k}{4 + \sum_{i=1}^k v_i} = \sum_{k=1}^{K-\tilde{k}+1} \frac{f(k)}{4 + F(k)}.$$

Additionally, we have that  $F(x) \leq \sum_{i=1}^{\lfloor x \rfloor} v_i + v_{\lceil x \rceil}(\lceil x \rceil - \lfloor x \rfloor) = \sum_{i=1}^{\lceil x \rceil} v_i = F(\lceil x \rceil)$  and  $f(x) = f(\lceil x \rceil)$ .

Then, we have

$$\begin{aligned} \sum_{k=1}^{K-\tilde{k}+1} \frac{f(k)}{4 + F(k)} &= \int_0^{K-\tilde{k}+1} \frac{f(\lceil x \rceil)}{4 + F(\lceil x \rceil)} dx \\ &\leq \int_0^{K-\tilde{k}+1} \frac{f(x)}{4 + F(x)} dx \\ &= \ln(4 + F(K - \tilde{k} + 1)) - \ln(4 + F(0)) \leq \tilde{\mathcal{O}}(\ln(K)). \end{aligned}$$

We obtain the result by summing over all the  $(s, a, h)$  pairs.  $\square$

**Lemma H.9.** *Outside the failure event, it holds that*

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} \left[ \sqrt{\frac{1}{n_h^{k-1}(s, a) \vee 1}} \right] = \tilde{\mathcal{O}}(\sqrt{C_{vr}SAH^2K} + C_{vr}SAH).$$

*Proof.* Outside the failure event, we can calculate as follows

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} \left[ \sqrt{\frac{1}{n_h^{k-1}(s, a) \vee 1}} \right] \\ &= \sum_{k=1}^K \sum_{(h,s,a)} q_h^k(s, a) \sqrt{\frac{1}{n_h^{k-1}(s, a) \vee 1}} \\ &\leq \sum_{k=1}^K \sum_{(h,s,a) \in L_k} q_h^k(s, a) \sqrt{\frac{1}{n_h^{k-1}(s, a) \vee 1}} + \sum_{k=1}^K \sum_{(h,s,a) \notin L_k} q_h^k(s, a) \end{aligned} \quad (\text{H.22})$$

$$= \sum_{k=1}^K \sum_{(h,s,a) \in L_k} q_h^k(s, a) \sqrt{\frac{1}{n_h^{k-1}(s, a)}} + \sum_{k=1}^K \sum_{(h,s,a) \notin L_k} q_h^k(s, a) \quad (\text{H.23})$$

$$\leq \sum_{k=1}^K \sum_{(h,s,a) \in L_k} q_h^k(s, a) \sqrt{\frac{1}{n_h^{k-1}(s, a)}} + C_{vr} \sum_{k=1}^K \sum_{(h,s,a) \notin L_k} d_h^k(s, a) \quad (\text{H.24})$$

$$\leq \sum_{k=1}^K \sum_{(h,s,a) \in L_k} q_h^k(s, a) \sqrt{\frac{1}{n_h^{k-1}(s, a)}} + \tilde{\mathcal{O}}(C_{vr}SAH) \quad (\text{H.25})$$

$$\leq \sqrt{\sum_{k=1}^K \sum_{(h,s,a)} q_h^k(s, a)} \sqrt{\sum_{k=1}^K \sum_{(h,s,a) \in L_k} \frac{q_h^k(s, a)}{n_h^{k-1}(s, a)}} + \tilde{\mathcal{O}}(C_{vr}SAH) \quad (\text{H.26})$$

$$\leq \sqrt{\sum_{k=1}^K \sum_{(h,s,a)} q_h^k(s, a)} \sqrt{C_{vr} \sum_{k=1}^K \sum_{(h,s,a) \in L_k} \frac{d_h^k(s, a)}{n_h^{k-1}(s, a)}} + \tilde{\mathcal{O}}(C_{vr}SAH) \quad (\text{H.27})$$

$$\begin{aligned} &\leq \sqrt{KH} \cdot \sqrt{\tilde{\mathcal{O}}(C_{vr}SAH)} + \tilde{\mathcal{O}}(C_{vr}SAH) \\ &= \tilde{\mathcal{O}}(\sqrt{C_{vr}SAH^2K} + C_{vr}SAH), \end{aligned} \quad (\text{H.28})$$

where (H.22) decomposes the summation into two parts and makes use of the fact that  $\sqrt{\frac{1}{n_h^{k-1}(s, a) \vee 1}} \leq 1$ , (H.23) holds because  $n_h^{k-1}(s, a) \geq \ln(\frac{SAHK}{\delta}) + 1 \geq 1$  by combining Definition H.5 and Lemma H.6, (H.24) and (H.27) are from our assumption Assumption 3.1, (H.25) and (H.28) are from Lemma H.7 and Lemma H.8 accordingly, and (H.26) is the Cauchy-Schwartz inequality.  $\square$

We are now ready to prove the main theorem that establishes the regret bound of ORBIT in the CRMDP-TV setting.

**Theorem H.10** (Restatement of Theorem 3.2 in TV-distance setting). *For CRMDP with  $(s, a)$ -rectangular TV-distance defined uncertainty set satisfying Assumption 3.1, with probability at least  $1 - \delta$ , the regret of Algorithm 1 satisfies*

$$\text{Regret} = \tilde{\mathcal{O}}(C_{vr}S^2AH^2 + C_{vr}^{\frac{1}{2}}S^{\frac{3}{2}}A^{\frac{1}{2}}H^2\sqrt{K}).$$

*Proof.* Setting  $\delta' = \delta/4$  in Lemmas H.2 and H.4, then with probability at least  $1 - \delta$ , we get

$$\begin{aligned} \text{Regret} &= \sum_{k=1}^K (V_1^{*,\rho} - V_1^{\pi^k,\rho}) \\ &= \sum_{k=1}^K (V_1^{*,\rho} - V_1^{k,\rho}) + \sum_{k=1}^K (V_1^{k,\rho} - V_1^{\pi^k,\rho}) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} [2\text{bonus}_h^k] \end{aligned} \quad (\text{H.29})$$

$$= 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} \left[ 2H \sqrt{\frac{2S^2 \ln(12SAH^2K^2/\delta)}{n_h^{k-1}(s,a) \vee 1}} + \frac{1}{K} \right] \quad (\text{H.30})$$

$$= \tilde{\mathcal{O}}(C_{vr} S^2 A H^2 + C_{vr}^{\frac{1}{2}} S^{\frac{3}{2}} A^{\frac{1}{2}} H^2 \sqrt{K}), \quad (\text{H.31})$$

where (H.29) is the combination of Lemma H.2 and Lemma H.3, we plug in the bonus (H.8) in (H.30), (H.31) is from Lemma H.9.  $\square$

## H.2 Proof of Theorem 3.2 (Constrained KL Setting)

We first give the closed form solution of constrained KL update formulation. This dual formulation has also been proved by [12, Lemma 4.1], but the range of  $\nu$  in our result is more precise compared to theirs.

**Lemma H.11** (Dual formulation). *For the optimization problem  $Q_h(s, a) = r_h(s, a) + \inf_{\text{KL}(P \| P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}](s, a)$ , we have its dual formulation as follows*

$$Q_h = r_h - \inf_{\nu \in [0, \frac{H}{\rho}]} (\nu \ln \mathbb{E}_{P_h^o} [e^{-\nu^{-1} V_{h+1}}] + \nu \rho). \quad (\text{H.32})$$

*Proof.* Consider the optimization problem

$$Q_h = r_h + \inf_{\text{KL}(P \| P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}] = r_h + \inf_{\text{KL}(P \| P_h^o) \leq \rho} \sum_{s \in \mathcal{S}} P(s) V_{h+1}(s).$$

The Lagrangian can be written as

$$\mathcal{L}(P, \eta) = \sum_{s \in \mathcal{S}} P(s) V_{h+1}(s) + \nu \left( \sum_{s \in \mathcal{S}} P(s) \ln \left( \frac{P(s)}{P_h^o(s)} \right) - \rho \right) - \sum_{s \in \mathcal{S}} \lambda(s) P(s) + \eta \left( 1 - \sum_{s \in \mathcal{S}} P(s) \right).$$

We set the derivative of  $\mathcal{L}$  w.r.t.  $P(s)$  to zero

$$\frac{\partial \mathcal{L}}{\partial P(s)} = V_{h+1}(s) + \nu \ln \left( \frac{P(s)}{P_h^o(s)} \right) + \nu - [\lambda(s) + \eta] = 0. \quad (\text{H.33})$$

We denote  $P'$  as the worst-case transition that satisfies (H.33), then we have

$$\begin{aligned} P'(s) &= P_h^o(s) e^{-\nu^{-1} V_{h+1}(s) + \nu^{-1} [\lambda(s) + \eta] - 1}, \\ \inf_P \mathcal{L}(P, \eta) &= -\nu \mathbb{E}_{P_h^o} [e^{-\nu^{-1} V_{h+1}(s) + \nu^{-1} [\lambda(s) + \eta] - 1}] - \nu \rho + \eta. \end{aligned}$$

Therefore,

$$\begin{aligned} Q_h &= r_h + \sup_{\nu \geq 0, \lambda \geq 0, \eta} (\inf_P \mathcal{L}(P, \eta)) \\ &= r_h - \inf_{\nu \geq 0, \lambda \geq 0, \eta} (\nu \mathbb{E}_{P_h^o} [e^{-\nu^{-1} V_{h+1}(s) + \nu^{-1} [\lambda(s) + \eta] - 1}] + \nu \rho - \eta) \\ &= r_h - \inf_{\nu \geq 0, \eta} (\nu \mathbb{E}_{P_h^o} [e^{-\nu^{-1} V_{h+1}(s) + \nu^{-1} \eta - 1}] + \nu \rho - \eta) \end{aligned} \quad (\text{H.34})$$

$$= r_h - \inf_{\nu \geq 0} (\nu \ln \mathbb{E}_{P_h^o} [e^{-\nu^{-1} V_{h+1}}] + \nu \rho). \quad (\text{H.35})$$

where (H.34) holds because the result increases monotonically with respect to  $\lambda$ , (H.35) holds by calculating the derivation with respect to  $\eta$  and thus setting  $\eta = -\nu(\ln \mathbb{E}_{P_h^o}[e^{-\nu^{-1}V_{h+1}}] - 1)$ .

We denote  $\tilde{\nu} = \operatorname{argmin}_{\nu \geq 0} (\nu \ln \mathbb{E}_{P_h^o}[e^{-\nu^{-1}V_{h+1}}] + \nu\rho)$ , from the strong duality, it is easy to infer that

$$\tilde{\nu} \ln \mathbb{E}_{P_h^o}[e^{-\tilde{\nu}^{-1}V_{h+1}}] + \tilde{\nu}\rho \leq 0. \quad (\text{H.36})$$

And since  $V_{h+1} \leq H$ , we have

$$\tilde{\nu} \ln \mathbb{E}_{P_h^o}[e^{-\tilde{\nu}^{-1}V_{h+1}}] \geq -H. \quad (\text{H.37})$$

Combine (H.36) and (H.37) into (H.35), we have  $\tilde{\nu} \in [0, \frac{H}{\rho}]$ . That is,

$$\begin{aligned} Q_h &= r_h - \inf_{\nu \geq 0} (\nu \ln \mathbb{E}_{P_h^o}[e^{-\nu^{-1}V_{h+1}}] + \nu\rho) \\ &= r_h - \inf_{\nu \in [0, \frac{H}{\rho}]} (\nu \ln \mathbb{E}_{P_h^o}[e^{-\nu^{-1}V_{h+1}}] + \nu\rho). \end{aligned}$$

This finishes the proof.  $\square$

Similar to Lemma H.2, we prove the optimism of estimation  $Q^k$  and control  $Q^* - Q^k$ .

**Lemma H.12 (Optimism).** *If we set the bonus term as follows*

$$\text{bonus}_h^k(s, a) = \left(1 + \frac{2H\sqrt{S}}{\rho C_{MP}}\right) \sqrt{\frac{2 \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}}, \quad (\text{H.38})$$

*then for any policy  $\pi$  and any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , with probability at least  $1 - 2\delta$ , we have  $Q_h^{k,\rho}(s, a) \geq Q_h^{\pi,\rho}(s, a)$ . Specially, by setting  $\pi = \pi^*$ , we have  $Q_h^{k,\rho}(s, a) \geq Q_h^{\pi^*,\rho}(s, a)$ .*

*Proof.* We prove this by induction. First, when  $h = H + 1$ ,  $Q_{H+1}^{k,\rho}(s, a) = 0 = Q_{H+1}^{\pi,\rho}(s, a)$  holds trivially.

Assume  $Q_{h+1}^{k,\rho}(s, a) \geq Q_{h+1}^{\pi,\rho}(s, a)$  holds, since  $\pi^k$  is the greedy policy, we have

$$V_{h+1}^{k,\rho}(s) = Q_{h+1}^{k,\rho}(s, \pi_{h+1}^k(s)) \geq Q_{h+1}^{k,\rho}(s, \pi_{h+1}(s)) \geq Q_{h+1}^{\pi,\rho}(s, \pi_{h+1}(s)) = V_{h+1}^{\pi,\rho}(s),$$

where the first inequality is because we choose  $\pi^k$  as the greedy policy.

Recall that we denote  $Q_h^{k,\rho}$  as the optimistic estimation in  $k$ -th episode, that is,

$$Q_h^{k,\rho}(s, a) = \min \left\{ \text{bonus}_h^k(s, a) + \hat{r}_h^k(s, a) + \inf_{\text{KL}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}](s, a), H - h + 1 \right\}.$$

If  $Q_h^{k,\rho}(s, a) = H - h + 1$ , then it follows immediately that

$$Q_h^{k,\rho}(s, a) = H - h + 1 \geq Q_h^{\pi,\rho}(s, a)$$

by the definition of  $Q_h^{\pi,\rho}(s, a)$ . Otherwise, we can infer that

$$\begin{aligned}
Q_h^{k,\rho} - Q_h^{\pi,\rho} &= \text{bonus}_h^k + \hat{r}_h^k + \inf_{\text{KL}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - r_h - \inf_{\text{KL}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi,\rho}] \\
&= \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\text{KL}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - \inf_{\text{KL}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] \\
&\quad + \inf_{\text{KL}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - \inf_{\text{KL}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi,\rho}] \\
&\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\text{KL}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - \inf_{\text{KL}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] \quad (\text{H.39})
\end{aligned}$$

$$\begin{aligned}
&= \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\nu \in [0, \frac{H}{\rho}]} \left( \nu \ln \mathbb{E}_{P_h^o} [e^{-\nu^{-1} V_{h+1}^{k,\rho}}] + \nu \rho \right) \\
&\quad - \inf_{\nu \in [0, \frac{H}{\rho}]} \left( \nu \ln \mathbb{E}_{\hat{P}_h^k} [e^{-\nu^{-1} V_{h+1}^{k,\rho}}] + \nu \rho \right) \quad (\text{H.40})
\end{aligned}$$

$$\begin{aligned}
&\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\nu \in [0, \frac{H}{\rho}]} \left( \nu \ln \mathbb{E}_{P_h^o} [e^{-\nu^{-1} V_{h+1}^{k,\rho}}] - \nu \ln \mathbb{E}_{\hat{P}_h^k} [e^{-\nu^{-1} V_{h+1}^{k,\rho}}] \right) \\
&\geq \text{bonus}_h^k - \underbrace{|\hat{r}_h^k - r_h|}_{(i)} - \underbrace{\sup_{\nu \in [0, \frac{H}{\rho}]} \left| \nu \ln \mathbb{E}_{\hat{P}_h^k} [e^{-\nu^{-1} V_{h+1}^{k,\rho}}] - \nu \ln \mathbb{E}_{P_h^o} [e^{-\nu^{-1} V_{h+1}^{k,\rho}}] \right|}_{(ii)}, \quad (\text{H.41})
\end{aligned}$$

where (H.39) is from the induction assumption, we plug in the dual formulation Lemma H.11 in (H.40).

For term (i) in (H.41), from Lemma L.1 and a union bound, with probability at least  $1 - \delta$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| \leq \sqrt{\frac{\ln(2SAHK/\delta)}{2n_h^{k-1}(s, a) \vee 1}}, \quad (\text{H.42})$$

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

To bound term (ii) in (H.11), we have

$$\begin{aligned}
&\sup_{\nu \in [0, \frac{H}{\rho}]} \left| \nu \ln \mathbb{E}_{\hat{P}_h^k} [e^{-\nu^{-1} V_{h+1}^{k,\rho}}] - \nu \ln \mathbb{E}_{P_h^o} [e^{-\nu^{-1} V_{h+1}^{k,\rho}}] \right| \\
&= \sup_{\nu \in [0, \frac{H}{\rho}]} \left| \nu \ln \left( 1 + \frac{\sum_{s \in \mathcal{S}} (\hat{P}_h^k(s) - P_h^o(s)) e^{-\nu^{-1} V_{h+1}^{k,\rho}(s)}}{\sum_{s \in \mathcal{S}} P_h^o(s) e^{-\nu^{-1} V_{h+1}^{k,\rho}(s)}} \right) \right| \\
&\leq \sup_{\nu \in [0, \frac{H}{\rho}]} 2\nu \left| \frac{\sum_{s \in \mathcal{S}} (\hat{P}_h^k(s) - P_h^o(s)) e^{-\nu^{-1} V_{h+1}^{k,\rho}(s)}}{\sum_{s \in \mathcal{S}} P_h^o(s) e^{-\nu^{-1} V_{h+1}^{k,\rho}(s)}} \right| \quad (\text{H.43})
\end{aligned}$$

$$\leq \sup_{\nu \in [0, \frac{H}{\rho}]} 2\nu \cdot \max_{s \in \mathcal{S}, P_h^o(s) \neq 0} \left| \frac{\hat{P}_h^k(s) - P_h^o(s)}{P_h^o(s)} \right| \quad (\text{H.44})$$

$$\leq \frac{2H}{\rho C_{MP}} \cdot \max_{s \in \mathcal{S}} |\hat{P}_h^k(s) - P_h^o(s)| \quad (\text{H.45})$$

$$\begin{aligned}
&\leq \frac{2H}{\rho C_{MP}} \cdot \|\hat{P}_h^k - P_h^o\|_1 \\
&\leq \frac{2H}{\rho C_{MP}} \sqrt{\frac{2S \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}}, \quad (\text{H.46})
\end{aligned}$$

for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , where (H.43) is because  $\ln(1+x) \leq 2|x|$ , (H.44) follows from the Holder's inequality, noting that we have  $n_h^{k-1}(s, a) = 0$  when  $P_h^o(s) = 0$  and thus  $\hat{P}_h^k(s) = 0$  from (C.1), (H.45) uses Assumption D.7, (H.46) is from (H.14) and a union bound.

Apply the union bound again and combine (H.41) with (H.42), (H.46), the definition of bonus and induction assumption. With probability at least  $1 - 2\delta$ , we have  $Q_h^{k,\rho}(s, a) \geq Q_h^{\pi^k,\rho}(s, a)$  for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ . This completes the proof.  $\square$

With similar proof to Lemma H.3, we can control the item  $Q^k - Q^{\pi^k}$ .

**Lemma H.13.** *With probability at least  $1 - \delta$ , for any  $(k, s, a) \in [K] \times \mathcal{S} \times \mathcal{A}$ , the sum of estimation errors can be bounded as follows*

$$Q_1^{k,\rho} - Q_1^{\pi^k,\rho} \leq \sum_{h=1}^H \mathbb{E}_{\{P_h^{w,k}\}_{h=1}^H, \pi^k} [2\text{bonus}_h^k].$$

*Proof.* From the proof of Lemma H.12, we see that with probability at least  $1 - \delta$ , for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| + \left| \inf_{\text{KL}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}](s, a) - \inf_{\text{KL}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}](s, a) \right| \leq \text{bonus}_h^k(s, a). \quad (\text{H.47})$$

Recall that we define  $P_h^{w,k} = \underset{\text{KL}(P \parallel P_h^o) \leq \rho}{\text{argmin}} \mathbb{E}_P[V_{h+1}^{\pi^k,\rho}]$  as the worst-case transition in Definition D.3, we have

$$\begin{aligned} Q_h^{k,\rho} - Q_h^{\pi^k,\rho} &\leq \text{bonus}_h^k + \hat{r}_h^k + \inf_{\text{KL}(P \parallel \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - r_h - \inf_{\text{KL}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi^k,\rho}] \\ &\leq 2\text{bonus}_h^k + \inf_{\text{KL}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - \inf_{\text{KL}(P \parallel P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi^k,\rho}] \end{aligned} \quad (\text{H.48})$$

$$= 2\text{bonus}_h^k + \mathbb{E}_{P_h^{w,k}, \pi^k} [Q_{h+1}^{k,\rho} - Q_{h+1}^{\pi^k,\rho}]. \quad (\text{H.49})$$

where (H.48) uses (H.47). Apply (H.49) recursively, we can obtain the result.  $\square$

Combine everything together the same way as the proof of Theorem H.10, we have

**Theorem H.14** (Restatement of Theorem 3.2 in KL-divergence setting). *For CRMDP with  $(s, a)$ -rectangular KL-divergence defined uncertainty set satisfying Assumption 3.1 and Assumption D.7, with probability at least  $1 - \delta$ , the regret of Algorithm 1 satisfies*

$$\text{Regret} = \tilde{O}\left(\left(1 + \frac{H\sqrt{S}}{\rho C_{MP}}\right)(C_{vr}SAH + C_{vr}^{\frac{1}{2}}S^{\frac{1}{2}}A^{\frac{1}{2}}H\sqrt{K})\right).$$

*Proof.* Setting  $\delta' = \delta/4$  in Lemmas H.4 and H.12, then with probability at least  $1 - \delta$ , we get

$$\begin{aligned} \text{Regret} &= \sum_{k=1}^K (V_1^{*,\rho} - V_1^{\pi^k,\rho}) \\ &= \sum_{k=1}^K (V_1^{*,\rho} - V_1^{k,\rho}) + \sum_{k=1}^K (V_1^{k,\rho} - V_1^{\pi^k,\rho}) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} [2\text{bonus}_h^k] \end{aligned} \quad (\text{H.50})$$

$$= 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} \left[ \left( 1 + \frac{2H\sqrt{S}}{\rho C_{MP}} \right) \sqrt{\frac{2 \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}} \right] \quad (\text{H.51})$$

$$= \tilde{\mathcal{O}} \left( \left( 1 + \frac{H\sqrt{S}}{\rho C_{MP}} \right) (C_{vr}SAH + C_{vr}^{\frac{1}{2}} S^{\frac{1}{2}} A^{\frac{1}{2}} H\sqrt{K}) \right), \quad (\text{H.52})$$

where (H.50) is the combination of Lemma H.12 and Lemma H.13, we plug in the bonus (H.38) in (H.51), (H.52) is from Lemma H.9.  $\square$

### H.3 Proof of Theorem 3.2 (Constrained $\chi^2$ Setting)

We first give the closed form solution of constrained  $\chi^2$  update formulation. This dual formulation has also been proved by [12, Lemma 4.2], but the range of  $\lambda$  in our result is more precise compared to theirs.

**Lemma H.15** (Dual formulation). *For the optimization problem  $Q_h(s, a) = r_h(s, a) + \inf_{\chi^2(P\|P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}](s, a)$ , we have its dual formulation as follows*

$$Q_h = r_h + \sup_{\lambda \in [0, H]} \left( \mathbb{E}_{P_h^o} [V_{h+1} - \lambda] - \sqrt{\rho \text{Var}_{P_h^o}(V_{h+1} - \lambda)} \right). \quad (\text{H.53})$$

*Proof.* Consider the optimization problem

$$Q_h = r_h + \inf_{\chi^2(P\|P_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}] = r_h + \inf_{\chi^2(P\|P_h^o) \leq \rho} \sum_{s \in \mathcal{S}} P(s) V_{h+1}(s).$$

The Lagrangian can be written as

$$\mathcal{L}(P, \eta) = \sum_{s \in \mathcal{S}} P(s) V_{h+1}(s) + \nu \left( \sum_{s \in \mathcal{S}} P_h^o(s) \left( \frac{P(s)}{P_h^o(s)} - 1 \right)^2 - \rho \right) - \sum_{s \in \mathcal{S}} \lambda(s) P(s) + \eta \left( 1 - \sum_{s \in \mathcal{S}} P(s) \right).$$

We set the derivative of  $\mathcal{L}$  w.r.t.  $P(s)$  to zero

$$\frac{\partial \mathcal{L}}{\partial P(s)} = V_{h+1}(s) + 2\nu \left( \frac{P(s)}{P_h^o(s)} - 1 \right) - [\lambda(s) + \eta] = 0. \quad (\text{H.54})$$

We denote  $P'$  as the worst-case transition that satisfies (H.54), then we have

$$\begin{aligned} P'(s) &= P_h^o(s) \left( 1 - \frac{V_{h+1}(s) - [\lambda(s) + \eta]}{2\nu} \right), \\ \inf_P \mathcal{L}(P, \eta) &= -\mathbb{E}_{P_h^o} \left[ \frac{1}{4\nu} (V_{h+1}(s) - [\lambda(s) + \eta])^2 - (V_{h+1}(s) - [\lambda(s) + \eta]) \right] - \nu\rho + \eta. \end{aligned}$$

Therefore,

$$\begin{aligned}
Q_h &= r_h + \sup_{\nu \geq 0, \lambda \geq 0, \eta} \left( \inf_P \mathcal{L}(P, \eta) \right) \\
&= r_h - \inf_{\nu \geq 0, \lambda \geq 0, \eta} \left( \mathbb{E}_{P_h^o} \left[ \frac{1}{4\nu} (V_{h+1}(s) - [\lambda(s) + \eta])^2 - (V_{h+1}(s) - [\lambda(s) + \eta]) \right] + \nu\rho - \eta \right) \\
&= r_h + \sup_{\nu \geq 0, \lambda \geq 0} \left( \mathbb{E}_{P_h^o} [V_{h+1} - \lambda] - \frac{1}{4\nu} \text{Var}_{P_h^o}(V_{h+1} - \lambda) - \nu\rho \right) \tag{H.55}
\end{aligned}$$

$$= r_h + \sup_{\lambda \geq 0} \left( \mathbb{E}_{P_h^o} [V_{h+1} - \lambda] - \sqrt{\rho \text{Var}_{P_h^o}(V_{h+1} - \lambda)} \right) \tag{H.56}$$

$$= r_h + \sup_{\lambda \in [0, H]} \left( \mathbb{E}_{P_h^o} [V_{h+1} - \lambda] - \sqrt{\rho \text{Var}_{P_h^o}(V_{h+1} - \lambda)} \right), \tag{H.57}$$

where (H.55) holds by calculating the derivation with respect to  $\eta$  and thus setting  $\eta = \mathbb{E}_{P_h^o}[V_{h+1} - \lambda]$ , (H.56) is from the basic inequality  $a + b \geq 2\sqrt{ab}$ , (H.57) holds because the result increases monotonically with respect to  $\lambda$  when  $\lambda \geq H$ .  $\square$

Similar to Lemma H.2, we prove the optimism of estimation  $Q^k$  and control  $Q^* - Q^k$ .

**Lemma H.16 (Optimism).** *If we set the bonus term as follows*

$$\text{bonus}_h^k(s, a) = (2 + \sqrt{\rho})H \sqrt{\frac{2S^2 \ln(192SAH^3K^3/\delta)}{n_h^{k-1}(s, a) \vee 1}} + (1 + \sqrt{\rho}) \frac{1}{K}, \tag{H.58}$$

then for any policy  $\pi$  and any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , with probability at least  $1 - 3\delta$ , we have  $Q_h^{k, \rho}(s, a) \geq Q_h^{\pi, \rho}(s, a)$ . Specially, by setting  $\pi = \pi^*$ , we have  $Q_h^{k, \rho}(s, a) \geq Q_h^{\pi^*, \rho}(s, a)$ .

*Proof.* We prove this by induction. First, when  $h = H + 1$ ,  $Q_{H+1}^{k, \rho}(s, a) = 0 = Q_{H+1}^{\pi, \rho}(s, a)$  holds trivially.

Assume  $Q_{h+1}^{k, \rho}(s, a) \geq Q_{h+1}^{\pi, \rho}(s, a)$  holds, since  $\pi^k$  is the greedy policy, we have

$$V_{h+1}^{k, \rho}(s) = Q_{h+1}^{k, \rho}(s, \pi_{h+1}^k(s)) \geq Q_{h+1}^{k, \rho}(s, \pi_{h+1}(s)) \geq Q_{h+1}^{\pi, \rho}(s, \pi_{h+1}(s)) = V_{h+1}^{\pi, \rho}(s),$$

where the first inequality is because we choose  $\pi^k$  as the greedy policy.

Recall that we denote  $Q_h^{k, \rho}$  as the optimistic estimation in  $k$ -th episode, that is,

$$Q_h^{k, \rho}(s, a) = \min \left\{ \text{bonus}_h^k(s, a) + \hat{r}_h^k(s, a) + \inf_{\chi^2(P \| \hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}](s, a), H - h + 1 \right\}.$$

If  $Q_h^{k, \rho}(s, a) = H - h + 1$ , then it follows immediately that

$$Q_h^{k, \rho}(s, a) = H - h + 1 \geq Q_h^{\pi, \rho}(s, a)$$



by the definition of  $Q_h^{\pi, \rho}(s, a)$ . Otherwise, we can infer that

$$\begin{aligned}
Q_h^{k, \rho} - Q_h^{\pi, \rho} &= \text{bonus}_h^k + \hat{r}_h^k + \inf_{\chi^2(P \|\hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] - r_h - \inf_{\chi^2(P \|\hat{P}_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi, \rho}] \\
&= \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\chi^2(P \|\hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] - \inf_{\chi^2(P \|\hat{P}_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] \\
&\quad + \inf_{\chi^2(P \|\hat{P}_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] - \inf_{\chi^2(P \|\hat{P}_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi, \rho}] \\
&\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\chi^2(P \|\hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] - \inf_{\chi^2(P \|\hat{P}_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k, \rho}] \quad (\text{H.59}) \\
&= \text{bonus}_h^k + \hat{r}_h^k - r_h + \sup_{\lambda \in [0, H]} \left( \mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k, \rho} - \lambda] - \sqrt{\rho \text{Var}_{\hat{P}_h^k}(V_{h+1}^{k, \rho} - \lambda)} \right) \\
&\quad - \sup_{\lambda \in [0, H]} \left( \mathbb{E}_{P_h^o}[V_{h+1}^{k, \rho} - \lambda] - \sqrt{\rho \text{Var}_{P_h^o}(V_{h+1}^{k, \rho} - \lambda)} \right) \quad (\text{H.60}) \\
&\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\lambda \in [0, H]} \left\{ \left( \mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k, \rho} - \lambda] - \sqrt{\rho \text{Var}_{\hat{P}_h^k}(V_{h+1}^{k, \rho} - \lambda)} \right) \right. \\
&\quad \left. - \left( \mathbb{E}_{P_h^o}[V_{h+1}^{k, \rho} - \lambda] - \sqrt{\rho \text{Var}_{P_h^o}(V_{h+1}^{k, \rho} - \lambda)} \right) \right\} \\
&\geq \text{bonus}_h^k - |\hat{r}_h^k - r_h| - \sup_{\lambda \in [0, H]} \left| \left( \mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k, \rho} - \lambda] - \sqrt{\rho \text{Var}_{\hat{P}_h^k}(V_{h+1}^{k, \rho} - \lambda)} \right) \right. \\
&\quad \left. - \left( \mathbb{E}_{P_h^o}[V_{h+1}^{k, \rho} - \lambda] - \sqrt{\rho \text{Var}_{P_h^o}(V_{h+1}^{k, \rho} - \lambda)} \right) \right| \\
&\geq \text{bonus}_h^k - \underbrace{|\hat{r}_h^k - r_h|}_{(i)} - \underbrace{\sup_{\lambda \in [0, H]} \left| \mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k, \rho} - \lambda] - \mathbb{E}_{P_h^o}[V_{h+1}^{k, \rho} - \lambda] \right|}_{(ii)} \\
&\quad - \underbrace{\sqrt{\rho} \sup_{\lambda \in [0, H]} \left| \sqrt{\text{Var}_{\hat{P}_h^k}(V_{h+1}^{k, \rho} - \lambda)} - \sqrt{\text{Var}_{P_h^o}(V_{h+1}^{k, \rho} - \lambda)} \right|}_{(iii)}, \quad (\text{H.61})
\end{aligned}$$

where (H.59) is from the induction assumption, we plug in the dual formulation Lemma H.15 in (H.60), (H.61) is because  $\sup f(x) + \sup g(x) \geq \sup(f - g)(x)$ .

For term (i) in (H.61), from Lemma L.1 and a union bound, with probability at least  $1 - \delta$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| \leq \sqrt{\frac{\ln(2SAHK/\delta)}{2n_h^{k-1}(s, a) \vee 1}}, \quad (\text{H.62})$$

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

We denote  $V(\lambda) = V_{h+1}^{k, \rho} - \lambda \in [-H, H]$  and  $\mathcal{V} = \{V \in \mathbb{R}^S : \|V\|_\infty \leq H\}$ . To bound term (ii) in (H.61), we create a  $\epsilon$ -net  $\mathcal{N}_\mathcal{V}(\epsilon)$  for  $\mathcal{V}$ . From Lemma L.4, it holds that  $\ln |\mathcal{N}_\mathcal{V}(\epsilon)| \leq |S| \cdot \ln(3H/\epsilon)$ . Then follow the same proof as (H.17), with probability at least  $1 - \delta$ , we have

$$\begin{aligned}
\sup_{\lambda \in [0, H]} \left| \mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k, \rho} - \lambda] - \mathbb{E}_{P_h^o}[V_{h+1}^{k, \rho} - \lambda] \right| &\leq \sup_{V \in \mathcal{N}_\mathcal{V}(\epsilon)} \left| \mathbb{E}_{P_h^o}[V] - \mathbb{E}_{\hat{P}_h^k}[V] \right| + 2\epsilon \\
&\leq H \sqrt{\frac{2S^2 \ln(12SAH^2K^2/\delta)}{n_h^{k-1} \vee 1}} + \frac{1}{K}, \quad (\text{H.63})
\end{aligned}$$

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

For term (iii) in (H.61), by the definition of  $\mathcal{N}_V(\epsilon)$ , for any fixed  $V$ , there exists a  $V' \in \mathcal{N}_V(\epsilon)$  such that  $\|V - V'\|_\infty \leq \epsilon$ , that is

$$\begin{aligned}
& \left| \sqrt{\text{Var}_{P_h^o}(V)} - \sqrt{\text{Var}_{\hat{P}_h^k}(V)} \right| \\
& \leq \left| \sqrt{\text{Var}_{P_h^o}(V)} - \sqrt{\text{Var}_{P_h^o}(V')} \right| + \left| \sqrt{\text{Var}_{P_h^o}(V')} - \sqrt{\text{Var}_{\hat{P}_h^k}(V')} \right| + \left| \sqrt{\text{Var}_{\hat{P}_h^k}(V')} - \sqrt{\text{Var}_{\hat{P}_h^k}(V)} \right| \\
& \leq \sqrt{|\text{Var}_{P_h^o}(V) - \text{Var}_{P_h^o}(V')|} + \sqrt{|\text{Var}_{P_h^o}(V') - \text{Var}_{\hat{P}_h^k}(V')|} + \sqrt{|\text{Var}_{\hat{P}_h^k}(V') - \text{Var}_{\hat{P}_h^k}(V)|} \\
& \leq \sqrt{|\mathbb{E}_{P_h^o}[V^2 - V'^2]|} + \sqrt{|\mathbb{E}_{P_h^o}^2[V] - \mathbb{E}_{P_h^o}^2[V']|} + \sqrt{|\mathbb{E}_{\hat{P}_h^k}[V'^2 - V^2]|} + \sqrt{|\mathbb{E}_{\hat{P}_h^k}^2[V'] - \mathbb{E}_{\hat{P}_h^k}^2[V]|} \\
& \quad + \left| \sqrt{\text{Var}_{P_h^o}(V')} - \sqrt{\text{Var}_{\hat{P}_h^k}(V')} \right| \\
& \leq \left| \sqrt{\text{Var}_{P_h^o}(V')} - \sqrt{\text{Var}_{\hat{P}_h^k}(V')} \right| + 4\sqrt{2H\epsilon} \\
& \leq \sup_{V' \in \mathcal{N}_V(\epsilon)} \left| \sqrt{\text{Var}_{P_h^o}(V')} - \sqrt{\text{Var}_{\hat{P}_h^k}(V')} \right| + 4\sqrt{2H\epsilon}, \tag{H.64}
\end{aligned}$$

where the second inequality follows from  $|\sqrt{x} - \sqrt{y}| \leq \sqrt{|x - y|}$  and the third inequality follows from  $\sqrt{x + y} \leq \sqrt{x} + \sqrt{y}$ .

For any fixed  $V$ , we apply Lemma L.2 and have

$$\left| \sqrt{\text{Var}_{P_h^o}(V)} - \sqrt{\text{Var}_{\hat{P}_h^k}(V)} \right| \leq H \sqrt{\frac{2 \ln(2/\delta)}{n_h^{k-1} \vee 1}}. \tag{H.65}$$

with probability at least  $1 - \delta$ .

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

$$\begin{aligned}
\sup_{\lambda \in [0, H]} \left| \sqrt{\text{Var}_{\hat{P}_h^k}(V_{h+1}^{k, \rho} - \lambda)} - \sqrt{\text{Var}_{P_h^o}(V_{h+1}^{k, \rho} - \lambda)} \right| & \leq \sup_{\lambda \in [0, H]} \left| \sqrt{\text{Var}_{P_h^o}(V(\lambda))} - \sqrt{\text{Var}_{\hat{P}_h^k}(V(\lambda))} \right| \\
& \leq \sup_{V \in \mathcal{N}_V(\epsilon)} \left| \sqrt{\text{Var}_{P_h^o}(V)} - \sqrt{\text{Var}_{\hat{P}_h^k}(V)} \right| + 4\sqrt{2H\epsilon} \tag{H.66} \\
& \leq H \sqrt{\frac{2 \ln(2SAHK|\mathcal{N}_V(\epsilon)|/\delta)}{n_h^{k-1} \vee 1}} + 4\sqrt{2H\epsilon} \tag{H.67} \\
& \leq H \sqrt{\frac{2S \ln(6SAH^2K/\epsilon\delta)}{n_h^{k-1} \vee 1}} + 4\sqrt{2H\epsilon} \\
& = H \sqrt{\frac{2S \ln(192SAH^3K^3/\delta)}{n_h^{k-1} \vee 1}} + \frac{1}{K}, \tag{H.68}
\end{aligned}$$

for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , where (H.66) follows from (H.64), (H.67) is from (H.65) and a union bound, we set  $\epsilon = 1/32HK^2$  in (H.68).

Apply the union bound again and combine (H.61) with (H.62), (H.63), (H.68), the definition of bonus and induction assumption. With probability at least  $1 - 3\delta$ , we have  $Q_h^{k, \rho}(s, a) \geq Q_h^{\pi, \rho}(s, a)$  for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ . This completes the proof.  $\square$

With similar proof to Lemma H.3, we can control the item  $Q^k - Q^{\pi^k}$ .

**Lemma H.17.** *With probability at least  $1 - \delta$ , for any  $(k, s, a) \in [K] \times \mathcal{S} \times \mathcal{A}$ , the sum of estimation errors can be bounded as follows*

$$Q_1^{k,\rho} - Q_1^{\pi^k,\rho} \leq \sum_{h=1}^H \mathbb{E}_{\{P_h^{w,k}\}_{h=1}^H, \pi^k} [2\text{bonus}_h^k].$$

*Proof.* From the proof of Lemma H.12, we see that with probability at least  $1 - \delta$ , for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| + \left| \inf_{\chi^2(P \|\hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}](s, a) - \inf_{\chi^2(P \|\hat{P}_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}](s, a) \right| \leq \text{bonus}_h^k(s, a). \quad (\text{H.69})$$

Recall that we define  $P_h^{w,k} = \underset{\chi^2(P \|\hat{P}_h^o) \leq \rho}{\operatorname{argmin}} \mathbb{E}_P[V_{h+1}^{\pi^k,\rho}]$  as the worst-case transition in Definition D.3, we have

$$\begin{aligned} Q_h^{k,\rho} - Q_h^{\pi^k,\rho} &\leq \text{bonus}_h^k + \hat{r}_h^k + \inf_{\chi^2(P \|\hat{P}_h^k) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - r_h - \inf_{\chi^2(P \|\hat{P}_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi^k,\rho}] \\ &\leq 2\text{bonus}_h^k + \inf_{\chi^2(P \|\hat{P}_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{k,\rho}] - \inf_{\chi^2(P \|\hat{P}_h^o) \leq \rho} \mathbb{E}_P[V_{h+1}^{\pi^k,\rho}] \end{aligned} \quad (\text{H.70})$$

$$= 2\text{bonus}_h^k + \mathbb{E}_{P_h^{w,k}, \pi^k} [Q_{h+1}^{k,\rho} - Q_{h+1}^{\pi^k,\rho}], \quad (\text{H.71})$$

where (H.70) uses (H.69). Apply (H.71) recursively, we can obtain the result.  $\square$

**Theorem H.18** (Restatement of Theorem 3.2 in  $\chi^2$ -divergence setting). *For CRMDP with  $(s, a)$ -rectangular  $\chi^2$ -divergence defined uncertainty set satisfying Assumption 3.1, with probability at least  $1 - \delta$ , the regret of Algorithm 1 satisfies*

$$\text{Regret} = \tilde{O}\left((1 + \sqrt{\rho})(C_{vr}S^2AH^2 + C_{vr}^{\frac{1}{2}}S^{\frac{3}{2}}A^{\frac{1}{2}}H^2\sqrt{K})\right).$$

*Proof.* Setting  $\delta' = \delta/5$  in Lemmas H.4 and H.16, then with probability at least  $1 - \delta$ , we get

$$\begin{aligned} \text{Regret} &= \sum_{k=1}^K (V_1^{*,\rho} - V_1^{\pi^k,\rho}) = \sum_{k=1}^K (V_1^{*,\rho} - V_1^{k,\rho}) + \sum_{k=1}^K (V_1^{k,\rho} - V_1^{\pi^k,\rho}) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} [2\text{bonus}_h^k] \end{aligned} \quad (\text{H.72})$$

$$= 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} \left[ (2 + \sqrt{\rho})H \sqrt{\frac{2S^2 \ln(192SAH^3K^3/\delta)}{n_h^{k-1}(s, a) \vee 1}} + (1 + \sqrt{\rho})\frac{1}{K} \right] \quad (\text{H.73})$$

$$= \tilde{O}\left((1 + \sqrt{\rho})(C_{vr}S^2AH^2 + C_{vr}^{\frac{1}{2}}S^{\frac{3}{2}}A^{\frac{1}{2}}H^2\sqrt{K})\right), \quad (\text{H.74})$$

where (H.72) is the combination of Lemma H.16 and Lemma H.17, we plug in the bonus (H.58) in (H.73), (H.74) is from Lemma H.9.  $\square$

## I Proofs of Results in Regularized RMDPs

### I.1 Proof of Theorem D.16 (Regularized TV Setting)

We first give the closed form solution of regularized TV update formulation. This dual formulation has also been proved by [25], but our formulation is more concise.

**Lemma I.1** (Dual formulation). *For the optimization problem  $Q_h(s, a) = r_h(s, a) + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}] + \beta \text{TV}(P \| P_h^o))(s, a)$ , we have its dual formulation as follows*

$$Q_h = r_h - \mathbb{E}_{P_h^o} \left[ \left( \min_{s \in S} V_{h+1}(s) + \beta - V_{h+1}(s) \right)_+ \right] + \left( \min_{s \in S} V_{h+1}(s) + \beta \right). \quad (\text{I.1})$$

*Proof.* Consider the optimization problem

$$\begin{aligned} Q_h &= r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}] + \beta \text{TV}(P \| P_h^o)) \\ &= r_h + \inf_{P \in \Delta(S)} \left( \sum_{s \in S} P(s) V_{h+1}(s) + \beta P_h^o(s) \varphi \left( \frac{P(s)}{P_h^o(s)} \right) \right). \end{aligned}$$

where  $\varphi(t) = |t - 1|/2$ , and then the Lagrangian can be written as

$$\mathcal{L}(P, \eta) = \sum_{s \in S} P(s) V_{h+1}(s) + \beta \sum_{s \in S} P_h^o(s) \varphi \left( \frac{P(s)}{P_h^o(s)} \right) - \sum_{s \in S} \lambda(s) P(s) + \eta \left( 1 - \sum_{s \in S} P(s) \right).$$

We denote  $g(s) = P(s)/P_h^o(s)$ , then we have

$$\begin{aligned} \inf_P \mathcal{L}(P, \eta) &= - \sup_g \sum_{s \in S} P_h^o(s) [g(s) [(\lambda(s) + \eta) - V_{h+1}(s)] - \beta \varphi(g(s))] + \eta \\ &= - \beta \mathbb{E}_{P_h^o} \left[ \varphi^* \left( \frac{(\lambda(s) + \eta) - V_{h+1}(s)}{\beta} \right) \right] + \eta, \end{aligned}$$

where the second equation is from the definition of dual function  $\varphi^*(y) = \sup_x (y^\top x - \varphi(x))$ .

So we have

$$\begin{aligned} Q_h &= r_h + \sup_{\lambda \geq 0, \eta} (\inf_P \mathcal{L}(P, \eta)) \\ &= r_h - \inf_{\lambda \geq 0, \eta} \left( \beta \mathbb{E}_{P_h^o} \left[ \varphi^* \left( \frac{(\lambda(s) + \eta) - V_{h+1}(s)}{\beta} \right) \right] - \eta \right) \\ &= r_h - \inf_{\lambda \geq 0, \eta, \frac{\lambda(s) + \eta - V_{h+1}(s)}{\beta} \leq \frac{1}{2}} \left( \beta \mathbb{E}_{P_h^o} \left[ \max \left( \frac{(\lambda(s) + \eta) - V_{h+1}(s)}{\beta}, -\frac{1}{2} \right) \right] - \eta \right) \quad (\text{I.2}) \end{aligned}$$

$$\begin{aligned} &= r_h - \inf_{\lambda \geq 0, \eta, \frac{\lambda(s) + \eta - V_{h+1}(s)}{\beta} \leq \frac{1}{2}} \left( \mathbb{E}_{P_h^o} \left[ \left( (\lambda(s) + \eta) - V_{h+1}(s) + \frac{\beta}{2} \right)_+ \right] - \eta - \frac{\beta}{2} \right) \\ &= r_h - \inf_{\lambda \geq 0, \eta', \lambda(s) + \eta' - V_{h+1}(s) \leq \beta} (\mathbb{E}_{P_h^o} [(\lambda(s) + \eta' - V_{h+1}(s))_+] - \eta') \quad (\text{I.3}) \end{aligned}$$

$$= r_h - \inf_{\eta' \leq V_{h+1}(s) + \beta} (\mathbb{E}_{P_h^o} [(\eta' - V_{h+1}(s))_+] - \eta') \quad (\text{I.4})$$

$$= r_h - \mathbb{E}_{P_h^o} \left[ \left( \min_{s \in S} V_{h+1}(s) + \beta - V_{h+1}(s) \right)_+ \right] + \left( \min_{s \in S} V_{h+1}(s) + \beta \right), \quad (\text{I.5})$$

where (I.2) follows from the definition of  $\varphi^*$  (H.2), we redefine  $\eta' = \eta + \frac{\beta}{2}$  in (I.3), (I.4) holds because the result increases monotonically with respect to  $\lambda$ , (I.5) holds because the result increases monotonically with respect to  $\eta'$ .  $\square$

Similar to Lemma H.2, we prove the optimism of estimation  $Q^k$  and control  $Q^* - Q^k$ .

**Lemma I.2** (Optimism). *If we set the bonus term as follows*

$$\text{bonus}_h^k(s, a) = 2H \sqrt{\frac{2S \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}}, \quad (\text{I.6})$$

then for any policy  $\pi$  and any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , with probability at least  $1 - 2\delta$ , we have  $Q_h^{k, \beta}(s, a) \geq Q_h^{\pi, \beta}(s, a)$ . Specially, by setting  $\pi = \pi^*$ , we have  $Q_h^{k, \beta}(s, a) \geq Q_h^{\pi^*, \beta}(s, a)$ .

*Proof.* We prove this by induction. First, when  $h = H + 1$ ,  $Q_{H+1}^{k,\beta}(s, a) = 0 = Q_{H+1}^{\pi,\beta}(s, a)$  holds trivially.

Assume  $Q_{h+1}^{k,\beta}(s, a) \geq Q_{h+1}^{\pi,\beta}(s, a)$  holds, since  $\pi^k$  is the greedy policy, we have

$$V_{h+1}^{k,\beta}(s) = Q_{h+1}^{k,\beta}(s, \pi_{h+1}^k(s)) \geq Q_{h+1}^{k,\beta}(s, \pi_{h+1}(s)) \geq Q_{h+1}^{\pi,\beta}(s, \pi_{h+1}(s)) = V_{h+1}^{\pi,\beta}(s),$$

where the first inequality is because we choose  $\pi^k$  as the greedy policy.

Recall that we denote  $Q_h^{k,\beta}$  as the optimistic estimation in  $k$ -th episode, that is,

$$Q_h^{k,\beta}(s, a) = \min \left\{ \text{bonus}_h^k(s, a) + \hat{r}_h^k(s, a) + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \| \hat{P}_h^k))(s, a), H - h + 1 \right\}.$$

If  $Q_h^{k,\beta}(s, a) = H - h + 1$ , then it follows immediately that

$$Q_h^{k,\beta}(s, a) = H - h + 1 \geq Q_h^{\pi,\beta}(s, a)$$

by the definition of  $Q_h^{\pi,\beta}(s, a)$ . Otherwise, we can infer that

$$\begin{aligned} Q_h^{k,\beta} - Q_h^{\pi,\beta} &= \text{bonus}_h^k + \hat{r}_h^k + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \| \hat{P}_h^k)) - r_h - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi,\beta}] + \beta \text{TV}(P \| P_h^o)) \\ &= \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \| \hat{P}_h^k)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \| P_h^o)) \\ &\quad + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \| P_h^o)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi,\beta}] + \beta \text{TV}(P \| P_h^o)) \\ &\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] - \mathbb{E}_P[V_{h+1}^{\pi,\beta}]) \\ &\quad + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \| \hat{P}_h^k)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \| P_h^o)) \end{aligned} \quad (\text{I.7})$$

$$\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \| \hat{P}_h^k)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \| P_h^o)) \quad (\text{I.8})$$

$$\begin{aligned} &= \text{bonus}_h^k + \hat{r}_h^k - r_h - \mathbb{E}_{P_h^o} \left[ \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta - V_{h+1}^{k,\beta}(s) \right)_+ \right] + \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta \right) \\ &\quad + \mathbb{E}_{\hat{P}_h^k} \left[ \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta - V_{h+1}^{k,\beta}(s) \right)_+ \right] - \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta \right) \end{aligned} \quad (\text{I.9})$$

$$\begin{aligned} &= \text{bonus}_h^k + \hat{r}_h^k - r_h - \mathbb{E}_{P_h^o} \left[ \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta - V_{h+1}^{k,\beta}(s) \right)_+ \right] + \mathbb{E}_{\hat{P}_h^k} \left[ \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta - V_{h+1}^{k,\beta}(s) \right)_+ \right] \\ &\geq \text{bonus}_h^k - \underbrace{|\hat{r}_h^k - r_h|}_{(i)} - \underbrace{\left| \mathbb{E}_{P_h^o} \left[ \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta - V_{h+1}^{k,\beta}(s) \right)_+ \right] - \mathbb{E}_{\hat{P}_h^k} \left[ \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta - V_{h+1}^{k,\beta}(s) \right)_+ \right] \right|}_{(ii)} \end{aligned} \quad (\text{I.10})$$

where (I.7) is from  $\inf f(x) - \inf g(x) \geq \inf(f - g)(x)$ , (I.8) is from the induction assumption, we plug in the dual formulation (I.1) in (I.9).

For term (i) in (I.10), from Lemma L.1 and a union bound, with probability at least  $1 - \delta$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| \leq \sqrt{\frac{\ln(2SAHK/\delta)}{2n_h^{k-1}(s, a) \vee 1}} \quad (\text{I.11})$$

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

For term (ii) in (I.10), follow the same proof as (H.14) and a union bound, with probability at least  $1 - \delta$ , we have

$$\left| \mathbb{E}_{P_h^o} \left[ \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta - V_{h+1}^{k,\beta}(s) \right)_+ \right] - \mathbb{E}_{\hat{P}_h^k} \left[ \left( \min_{s \in \mathcal{S}} V_{h+1}^{k,\beta}(s) + \beta - V_{h+1}^{k,\beta}(s) \right)_+ \right] \right| \leq H \sqrt{\frac{2S \ln(2SAHK/\delta)}{n_h^{k-1} \vee 1}} \quad (\text{I.12})$$

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

Apply the union bound again and combine (I.10) with (I.11), (I.12), the definition of bonus and induction assumption. With probability at least  $1 - 2\delta$ , we have  $Q_h^{k,\rho}(s, a) \geq Q_h^{\pi,\rho}(s, a)$  for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ . This completes the proof.  $\square$

With similar proof to Lemma H.3, we can control the item  $Q^k - Q^{\pi^k}$ .

**Lemma I.3.** *With probability at least  $1 - \delta$ , for any  $(k, s, a) \in [K] \times \mathcal{S} \times \mathcal{A}$ , the sum of estimation errors can be bounded as follows*

$$Q_1^{k,\beta} - Q_1^{\pi^k,\beta} \leq \sum_{h=1}^H \mathbb{E}_{\{P_h^{w,k}\}_{h=1}^H, \pi^k} [2\text{bonus}_h^k].$$

*Proof.* From the proof of Lemma I.2, we see that with probability at least  $1 - \delta$ , for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$\begin{aligned} & |\hat{r}_h^k(s, a) - r_h(s, a)| + \left| \inf_{P \in \Delta(\mathcal{S})} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \|\hat{P}_h^k))(s, a) \right. \\ & \quad \left. - \inf_{P \in \Delta(\mathcal{S})} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \|P_h^o))(s, a) \right| \leq \text{bonus}_h^k(s, a). \end{aligned} \quad (\text{I.13})$$

Recall that we define  $P_h^{w,k} = \operatorname{argmin}_{P \in \Delta(\mathcal{S})} (\mathbb{E}_P[V_{h+1}^{\pi^k,\beta}] + \beta \text{TV}(P \|P_h^o))$  as the worst-case transition in Definition D.3, we have

$$\begin{aligned} Q_h^{k,\beta} - Q_h^{\pi^k,\beta} & \leq \text{bonus}_h^k + \hat{r}_h^k + \inf_{P \in \Delta(\mathcal{S})} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \|\hat{P}_h^k)) - r_h - \inf_{P \in \Delta(\mathcal{S})} (\mathbb{E}_P[V_{h+1}^{\pi^k,\beta}] + \beta \text{TV}(P \|P_h^o)) \\ & \leq 2\text{bonus}_h^k + \inf_{P \in \Delta(\mathcal{S})} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{TV}(P \|P_h^o)) - \inf_{P \in \Delta(\mathcal{S})} (\mathbb{E}_P[V_{h+1}^{\pi^k,\beta}] + \beta \text{TV}(P \|P_h^o)) \\ & \quad (\text{I.14}) \end{aligned}$$

$$= 2\text{bonus}_h^k + \mathbb{E}_{P_h^{w,k}, \pi^k} [Q_{h+1}^{k,\beta} - Q_{h+1}^{\pi^k,\beta}]. \quad (\text{I.15})$$

where (I.14) uses (I.13). Apply (I.15) recursively, we can obtain the result.  $\square$

**Theorem I.4** (Restatement of Theorem D.16 in TV-distance setting). *For RRMDP with  $(s, a)$ -rectangular TV-distance defined regularization term satisfying Assumption 3.1, with probability at least  $1 - \delta$ , the regret of Algorithm 1 satisfies*

$$\text{Regret} = \tilde{\mathcal{O}}(C_{vr} S^{\frac{3}{2}} A H^2 + C_{vr}^{\frac{1}{2}} S A^{\frac{1}{2}} H^2 \sqrt{K}).$$

*Proof.* Setting  $\delta' = \delta/4$  in Lemmas H.4 and I.2, then with probability at least  $1 - \delta$ , we get

$$\begin{aligned} \text{Regret} & = \sum_{k=1}^K (V_1^{*,\beta} - V_1^{\pi^k,\beta}) = \sum_{k=1}^K (V_1^{*,\beta} - V_1^{k,\beta}) + \sum_{k=1}^K (V_1^{k,\beta} - V_1^{\pi^k,\beta}) \\ & \leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} [2\text{bonus}_h^k] \end{aligned} \quad (\text{I.16})$$

$$= 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} \left[ 2H \sqrt{\frac{2S \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}} \right] \quad (\text{I.17})$$

$$= \tilde{\mathcal{O}}(C_{vr} S^{\frac{3}{2}} A H^2 + C_{vr}^{\frac{1}{2}} S A^{\frac{1}{2}} H^2 \sqrt{K}). \quad (\text{I.18})$$

where (I.16) is the combination of Lemma I.2 and Lemma I.3, we plug in the bonus (I.6) in (I.17), (I.18) is from Lemma H.9.  $\square$

## I.2 Proof of Theorem D.16 (Regularized KL Setting)

We first give the closed form solution of regularized KL update formulation. This dual formulation has also been proved by [25], but our formulation is more concise. [45] applied the equivalence between regularized RMDPs and risk-sensitive MDPs to obtain the same result, which is different from our proof.

**Lemma I.5** (Dual formulation). *For the optimization problem  $Q_h(s, a) = r_h(s, a) + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}] + \beta \text{KL}(P \| P_h^o))(s, a)$ , we have its dual formulation as follows*

$$Q_h = r_h - \beta \ln \mathbb{E}_{P_h^o} [e^{-\beta^{-1} V_{h+1}}]. \quad (\text{I.19})$$

*Proof.* Consider the optimization problem

$$\begin{aligned} Q_h &= r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}] + \beta \text{KL}(P \| P_h^o)) \\ &= r_h + \inf_{P \in \Delta(S)} \left( \sum_{s \in S} P(s) V_{h+1}(s) + \beta P(s) \ln \frac{P(s)}{P_h^o(s)} \right). \end{aligned}$$

The Lagrangian can be written as

$$\mathcal{L}(P, \eta) = \sum_{s \in S} \left[ P(s) V_{h+1}(s) + \beta P(s) \ln \left( \frac{P(s)}{P_h^o(s)} \right) \right] - \sum_{s \in S} \lambda(s) P(s) + \eta \left( 1 - \sum_{s \in S} P(s) \right).$$

We set the derivative of  $\mathcal{L}$  w.r.t.  $P(s)$  to zero

$$\frac{\partial \mathcal{L}}{\partial P(s)} = V_{h+1}(s) + \beta \ln \left( \frac{P(s)}{P_h^o(s)} \right) + \beta - [\lambda(s) + \eta] = 0. \quad (\text{I.20})$$

We denote  $P'$  as the worst-case transition that satisfies (I.20), then we have

$$\begin{aligned} P'(s) &= P_h^o(s) e^{-\beta^{-1} V_{h+1}(s) + \beta^{-1} [\lambda(s) + \eta] - 1}, \\ \inf_P \mathcal{L}(P, \eta) &= -\beta \mathbb{E}_{P_h^o} [e^{-\beta^{-1} V_{h+1}(s) + \beta^{-1} [\lambda(s) + \eta] - 1}] + \eta. \end{aligned} \quad (\text{I.21})$$

We have

$$\begin{aligned} Q_h &= r_h + \sup_{\lambda \geq 0, \eta} \left( \inf_P \mathcal{L}(P, \eta) \right) \\ &= r_h - \inf_{\lambda \geq 0, \eta} \left( \beta \mathbb{E}_{P_h^o} [e^{-\beta^{-1} V_{h+1}(s) + \beta^{-1} [\lambda(s) + \eta] - 1}] - \eta \right) \\ &= r_h - \inf_{\eta} \left( \beta \mathbb{E}_{P_h^o} [e^{-\beta^{-1} V_{h+1}(s) + \beta^{-1} \eta - 1}] - \eta \right) \end{aligned} \quad (\text{I.22})$$

$$= r_h - \beta \ln \mathbb{E}_{P_h^o} [e^{-\beta^{-1} V_{h+1}}]. \quad (\text{I.23})$$

where (I.22) holds because the result increases monotonically with respect to  $\lambda$ , (I.23) holds by calculating the derivation with respect to  $\eta$  and thus setting  $\eta = -\beta (\ln \mathbb{E}_{P_h^o} [e^{-\beta^{-1} V_{h+1}}] - 1)$ .  $\square$

Similar to Lemma H.2, we prove the optimism of estimation  $Q^k$  and control  $Q^* - Q^k$ .

**Lemma I.6** (Optimism). *If we set the bonus term as follows*

$$\text{bonus}_h^k(s, a) = (1 + \beta e^{\beta^{-1} H} \sqrt{S}) \sqrt{\frac{2 \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}}, \quad (\text{I.24})$$

then for any policy  $\pi$  and any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , with probability at least  $1 - 2\delta$ , we have  $Q_h^{k, \beta}(s, a) \geq Q_h^{\pi, \beta}(s, a)$ . Specially, by setting  $\pi = \pi^*$ , we have  $Q_h^{k, \beta}(s, a) \geq Q_h^{\pi^*, \beta}(s, a)$ .

*Proof.* We prove this by induction. First, when  $h = H + 1$ ,  $Q_{H+1}^{k,\beta}(s, a) = 0 = Q_{H+1}^{\pi,\beta}(s, a)$  holds trivially.

Assume  $Q_{h+1}^{k,\beta}(s, a) \geq Q_{h+1}^{\pi,\beta}(s, a)$  holds, since  $\pi^k$  is the greedy policy, we have

$$V_{h+1}^{k,\beta}(s) = Q_{h+1}^{k,\beta}(s, \pi_{h+1}^k(s)) \geq Q_{h+1}^{k,\beta}(s, \pi_{h+1}(s)) \geq Q_{h+1}^{\pi,\beta}(s, \pi_{h+1}(s)) = V_{h+1}^{\pi,\beta}(s),$$

where the first inequality is because we choose  $\pi^k$  as the greedy policy.

Recall that we denote  $Q_h^{k,\beta}$  as the optimistic estimation in  $k$ -th episode, that is,

$$Q_h^{k,\beta}(s, a) = \min \left\{ \text{bonus}_h^k(s, a) + \hat{r}_h^k(s, a) + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| \hat{P}_h^k))(s, a), H - h + 1 \right\}.$$

If  $Q_h^{k,\beta}(s, a) = H - h + 1$ , then it follows immediately that

$$Q_h^{k,\beta}(s, a) = H - h + 1 \geq Q_h^{\pi,\beta}(s, a)$$

by the definition of  $Q_h^{\pi,\beta}(s, a)$ . Otherwise, we can infer that

$$\begin{aligned} Q_h^{k,\beta} - Q_h^{\pi,\beta} &= \text{bonus}_h^k + \hat{r}_h^k + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| \hat{P}_h^k)) - r_h - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi,\beta}] + \beta \text{KL}(P \| P_h^o)) \\ &= \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| \hat{P}_h^k)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| P_h^o)) \\ &\quad + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| P_h^o)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi,\beta}] + \beta \text{KL}(P \| P_h^o)) \\ &\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] - \mathbb{E}_P[V_{h+1}^{\pi,\beta}]) \\ &\quad + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| \hat{P}_h^k)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| P_h^o)) \end{aligned} \quad (\text{I.25})$$

$$\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| \hat{P}_h^k)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| P_h^o)) \quad (\text{I.26})$$

$$= \text{bonus}_h^k + \hat{r}_h^k - r_h + \beta \ln \mathbb{E}_{P_h^o}[e^{-\beta^{-1} V_{h+1}^{k,\beta}}] - \beta \ln \mathbb{E}_{\hat{P}_h^k}[e^{-\beta^{-1} V_{h+1}^{k,\beta}}] \quad (\text{I.27})$$

$$\geq \text{bonus}_h^k - \underbrace{|\hat{r}_h^k - r_h|}_{(i)} - \underbrace{\beta |\ln \mathbb{E}_{\hat{P}_h^k}[e^{-\beta^{-1} V_{h+1}^{k,\beta}}] - \ln \mathbb{E}_{P_h^o}[e^{-\beta^{-1} V_{h+1}^{k,\beta}}]}_{(ii)}, \quad (\text{I.28})$$

where (I.25) is from  $\inf f(x) - \inf g(x) \geq \inf(f - g)(x)$ , (I.26) is from the induction assumption, we plug in the dual formulation (I.19) in (I.27).

For term (i) in (I.28), from Lemma L.1 and a union bound, with probability at least  $1 - \delta$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| \leq \sqrt{\frac{\ln(2SAHK/\delta)}{2n_h^{k-1}(s, a) \vee 1}} \quad (\text{I.29})$$

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

For term (ii) in (I.28), following the same proof as (H.14) and a union bound, with probability at least  $1 - \delta$ , we have

$$|\ln \mathbb{E}_{\hat{P}_h^k}[e^{-\beta^{-1} V_{h+1}^{k,\beta}}] - \ln \mathbb{E}_{P_h^o}[e^{-\beta^{-1} V_{h+1}^{k,\beta}}]| \leq e^{\beta^{-1} H} \sqrt{\frac{2S \ln(2SAHK/\delta)}{n_h^{k-1} \vee 1}} \quad (\text{I.30})$$

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

Apply the union bound again and combine (I.28) with (I.29), (I.30), the definition of bonus and induction assumption. With probability at least  $1 - 2\delta$ , we have  $Q_h^{k,\rho}(s, a) \geq Q_h^{\pi,\rho}(s, a)$  for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ . This completes the proof.  $\square$



With similar proof to Lemma H.13, we can control the item  $Q^k - Q^{\pi^k}$ .

**Lemma I.7.** *With probability at least  $1 - \delta$ , for any  $(k, s, a) \in [K] \times \mathcal{S} \times \mathcal{A}$ , the sum of estimation errors can be bounded as follows*

$$Q_1^{k,\beta} - Q_1^{\pi^k,\beta} \leq \sum_{h=1}^H \mathbb{E}_{\{P_h^{w,k}\}_{h=1}^H, \pi^k} [2\text{bonus}_h^k].$$

*Proof.* From the proof of Lemma I.6, we see that with probability at least  $1 - \delta$ , for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$\begin{aligned} & |\hat{r}_h^k(s, a) - r_h(s, a)| + \left| \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| \hat{P}_h^k))(s, a) \right. \\ & \quad \left. - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| P_h^o))(s, a) \right| \leq \text{bonus}_h^k(s, a). \end{aligned} \quad (\text{I.31})$$

Recall that we define  $P_h^{w,k} = \underset{P \in \Delta(S)}{\text{argmin}} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| P_h^o))$  as the worst-case transition in Definition D.3, we have

$$\begin{aligned} Q_h^{k,\beta} - Q_h^{\pi^k,\beta} & \leq \text{bonus}_h^k + \hat{r}_h^k + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| \hat{P}_h^k)) - r_h - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi^k,\beta}] + \beta \text{KL}(P \| P_h^o)) \\ & \leq 2\text{bonus}_h^k + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \text{KL}(P \| P_h^o)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi^k,\beta}] + \beta \text{KL}(P \| P_h^o)) \end{aligned} \quad (\text{I.32})$$

$$= 2\text{bonus}_h^k + \mathbb{E}_{P_h^{w,k}, \pi^k} [Q_{h+1}^{k,\beta} - Q_{h+1}^{\pi^k,\beta}]. \quad (\text{I.33})$$

where (I.32) uses (I.31). Apply (I.33) recursively, we can obtain the result.  $\square$

**Theorem I.8** (Restatement of Theorem D.16 in KL-divergence setting). *For RRMDP with  $(s, a)$ -rectangular KL-divergence defined regularization term satisfying Assumption 3.1, with probability at least  $1 - \delta$ , the regret of Algorithm 1 satisfies*

$$\text{Regret} = \tilde{\mathcal{O}}((1 + \beta e^{\beta^{-1}H} \sqrt{S})(C_{vr}SAH + C_{vr}^{\frac{1}{2}}S^{\frac{1}{2}}A^{\frac{1}{2}}H\sqrt{K})).$$

*Proof.* Setting  $\delta' = \delta/4$  in Lemmas H.4 and I.6, then with probability at least  $1 - \delta$ , we get

$$\begin{aligned} \text{Regret} &= \sum_{k=1}^K (V_1^{*,\beta} - V_1^{\pi^k,\beta}) = \sum_{k=1}^K (V_1^{*,\beta} - V_1^{k,\beta}) + \sum_{k=1}^K (V_1^{k,\beta} - V_1^{\pi^k,\beta}) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} [2\text{bonus}_h^k] \end{aligned} \quad (\text{I.34})$$

$$= 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} \left[ (1 + \beta e^{\beta^{-1}H} \sqrt{S}) \sqrt{\frac{2 \ln(2SAHK/\delta)}{n_h^{k-1}(s, a) \vee 1}} \right] \quad (\text{I.35})$$

$$= \tilde{\mathcal{O}}((1 + \beta e^{\beta^{-1}H} \sqrt{S})(C_{vr}SAH + C_{vr}^{\frac{1}{2}}S^{\frac{1}{2}}A^{\frac{1}{2}}H\sqrt{K})). \quad (\text{I.36})$$

where (I.34) is the combination of Lemma I.6 and Lemma I.7, we plug in the bonus (I.24) in (I.35), (I.36) is from Lemma H.9.  $\square$

### I.3 Proof of Theorem D.16 (Regularized $\chi^2$ Setting)

We first give the closed form solution of regularized  $\chi^2$  update formulation. This dual formulation has also been proved by [25], but our formulation is more concise.

**Lemma I.9** (Dual formulation). *For the optimization problem  $Q_h(s, a) = r_h(s, a) + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}] + \beta \chi^2(P \| P_h^o))(s, a)$ , we have its dual formulation as follows*

$$Q_h = r_h + \sup_{\lambda \in [0, H]} (\mathbb{E}_{P_h^o}[V_{h+1} - \lambda] - \frac{1}{4\beta} \text{Var}_{P_h^o}(V_{h+1} - \lambda)). \quad (\text{I.37})$$

*Proof.* Consider the optimization problem

$$\begin{aligned} Q_h &= r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}] + \beta \chi^2(P \| P_h^o)) \\ &= r_h + \inf_{P \in \Delta(S)} \sum_{s \in S} \left( P(s) V_{h+1}(s) + \beta P_h^o(s) \left( \frac{P(s)}{P_h^o(s)} - 1 \right)^2 \right). \end{aligned}$$

The Lagrangian can be written as

$$\mathcal{L}(P, \eta) = \sum_{s \in S} \left[ P(s) V_{h+1}(s) + \beta P_h^o(s) \left( \frac{P(s)}{P_h^o(s)} - 1 \right)^2 \right] - \sum_{s \in S} \lambda(s) P(s) + \eta \left( 1 - \sum_{s \in S} P(s) \right).$$

We set the derivative of  $\mathcal{L}$  w.r.t.  $P(s)$  to zero

$$\frac{\partial \mathcal{L}}{\partial P(s)} = V_{h+1}(s) + 2\beta \left( \frac{P(s)}{P_h^o(s)} - 1 \right) - [\lambda(s) + \eta] = 0. \quad (\text{I.38})$$

We denote  $P'$  as the worst-case transition that satisfies (I.38), then we have

$$P'(s) = P_h^o(s) \left( 1 - \frac{V_{h+1}(s) - [\lambda(s) + \eta]}{2\beta} \right), \quad (\text{I.39})$$

$$\inf_P \mathcal{L}(P, \eta) = -\mathbb{E}_{P_h^o} \left[ \frac{1}{4\beta} (V_{h+1}(s) - [\lambda(s) + \eta])^2 - (V_{h+1}(s) - [\lambda(s) + \eta]) \right] + \eta.$$

We have

$$\begin{aligned} Q_h &= r_h + \sup_{\lambda \geq 0, \eta} (\inf_P \mathcal{L}(P, \eta)) \\ &= r_h - \inf_{\lambda \geq 0, \eta} \left( \mathbb{E}_{P_h^o} \left[ \frac{1}{4\beta} (V_{h+1}(s) - [\lambda(s) + \eta])^2 - (V_{h+1}(s) - [\lambda(s) + \eta]) \right] - \eta \right) \\ &= r_h + \sup_{\lambda \geq 0} \left( \mathbb{E}_{P_h^o}[V_{h+1} - \lambda] - \frac{1}{4\beta} \text{Var}_{P_h^o}(V_{h+1} - \lambda) \right) \end{aligned} \quad (\text{I.40})$$

$$= r_h + \sup_{\lambda \in [0, H]} \left( \mathbb{E}_{P_h^o}[V_{h+1} - \lambda] - \frac{1}{4\beta} \text{Var}_{P_h^o}(V_{h+1} - \lambda) \right), \quad (\text{I.41})$$

where (I.40) holds by calculating the derivation with respect to  $\eta$  and thus setting  $\eta = \mathbb{E}_{P_h^o}[V_{h+1} - \lambda]$ , (I.41) holds because the result increases monotonically with respect to  $\lambda$  when  $\lambda \geq H$ .  $\square$

Similar to Lemma H.2, we prove the optimism of estimation  $Q^k$  and control  $Q^* - Q^k$ .

**Lemma I.10** (Optimism). *If we set the bonus term as follows*

$$\text{bonus}_h^k(s, a) = \left( 2 + \frac{3H}{4\beta} \right) H \sqrt{\frac{2S^2 \ln(48SAH^3K^2/\delta)}{n_h^{k-1}(s, a) \vee 1}} + \left( 1 + \frac{1}{4\beta} \right) \frac{1}{K}, \quad (\text{I.42})$$

then for any policy  $\pi$  and any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , with probability at least  $1 - 3\delta$ , we have  $Q_h^{k, \beta}(s, a) \geq Q_h^{\pi, \beta}(s, a)$ . Specially, by setting  $\pi = \pi^*$ , we have  $Q_h^{k, \beta}(s, a) \geq Q_h^{\pi^*, \beta}(s, a)$ .

*Proof.* We prove this by induction. First, when  $h = H + 1$ ,  $Q_{H+1}^{k, \beta}(s, a) = 0 = Q_{H+1}^{\pi, \beta}(s, a)$  holds trivially.

Assume  $Q_{h+1}^{k,\beta}(s, a) \geq Q_{h+1}^{\pi^k, \beta}(s, a)$  holds, since  $\pi^k$  is the greedy policy, we have

$$V_{h+1}^{k,\beta}(s) = Q_{h+1}^{k,\beta}(s, \pi_{h+1}^k(s)) \geq Q_{h+1}^{k,\beta}(s, \pi_{h+1}(s)) \geq Q_{h+1}^{\pi^k, \beta}(s, \pi_{h+1}(s)) = V_{h+1}^{\pi^k, \beta}(s),$$

where the first inequality is because we choose  $\pi^k$  as the greedy policy.

Recall that we denote  $Q_h^{k,\beta}$  as the optimistic estimation in  $k$ -th episode, that is,

$$Q_h^{k,\beta}(s, a) = \min \left\{ \text{bonus}_h^k(s, a) + \hat{r}_h^k(s, a) + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \chi^2(P \| \hat{P}_h^k))(s, a), H - h + 1 \right\}.$$

If  $Q_h^{k,\beta}(s, a) = H - h + 1$ , then it follows immediately that

$$Q_h^{k,\beta}(s, a) = H - h + 1 \geq Q_h^{\pi^k, \beta}(s, a)$$

by the definition of  $Q_h^{\pi^k, \beta}(s, a)$ . Otherwise, we can infer that

$$\begin{aligned} Q_h^{k,\beta} - Q_h^{\pi^k, \beta} &= \text{bonus}_h^k + \hat{r}_h^k + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \chi^2(P \| \hat{P}_h^k)) - r_h - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi^k, \beta}] + \beta \chi^2(P \| P_h^o)) \\ &= \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \chi^2(P \| \hat{P}_h^k)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \chi^2(P \| P_h^o)) \\ &\quad + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \chi^2(P \| P_h^o)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi^k, \beta}] + \beta \chi^2(P \| P_h^o)) \\ &\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] - \mathbb{E}_P[V_{h+1}^{\pi^k, \beta}]) \\ &\quad + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \chi^2(P \| \hat{P}_h^k)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \chi^2(P \| P_h^o)) \end{aligned} \quad (\text{I.43})$$

$$\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \chi^2(P \| \hat{P}_h^k)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k,\beta}] + \beta \chi^2(P \| P_h^o)) \quad (\text{I.44})$$

$$\begin{aligned} &= \text{bonus}_h^k + \hat{r}_h^k - r_h + \sup_{\lambda \in [0, H]} (\mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k,\beta} - \lambda] - \frac{1}{4\beta} \text{Var}_{\hat{P}_h^k}(V_{h+1}^{k,\beta} - \lambda)) \\ &\quad - \sup_{\lambda \in [0, H]} (\mathbb{E}_{P_h^o}[V_{h+1}^{k,\beta} - \lambda] - \frac{1}{4\beta} \text{Var}_{P_h^o}(V_{h+1}^{k,\beta} - \lambda)) \end{aligned} \quad (\text{I.45})$$

$$\begin{aligned} &\geq \text{bonus}_h^k + \hat{r}_h^k - r_h + \inf_{\lambda \in [0, H]} \left\{ (\mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k,\beta} - \lambda] - \frac{1}{4\beta} \text{Var}_{\hat{P}_h^k}(V_{h+1}^{k,\beta} - \lambda)) \right. \\ &\quad \left. - (\mathbb{E}_{P_h^o}[V_{h+1}^{k,\beta} - \lambda] - \frac{1}{4\beta} \text{Var}_{P_h^o}(V_{h+1}^{k,\beta} - \lambda)) \right\} \\ &\geq \text{bonus}_h^k - |\hat{r}_h^k - r_h| - \sup_{\lambda \in [0, H]} |(\mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k,\beta} - \lambda] - \frac{1}{4\beta} \text{Var}_{\hat{P}_h^k}(V_{h+1}^{k,\beta} - \lambda)) \\ &\quad - (\mathbb{E}_{P_h^o}[V_{h+1}^{k,\beta} - \lambda] - \frac{1}{4\beta} \text{Var}_{P_h^o}(V_{h+1}^{k,\beta} - \lambda))| \\ &\geq \text{bonus}_h^k - \underbrace{|\hat{r}_h^k - r_h|}_{(i)} - \underbrace{\sup_{\lambda \in [0, H]} |\mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k,\beta} - \lambda] - \mathbb{E}_{P_h^o}[V_{h+1}^{k,\beta} - \lambda]|}_{(ii)} \\ &\quad - \underbrace{\frac{1}{4\beta} \sup_{\lambda \in [0, H]} |\text{Var}_{\hat{P}_h^k}(V_{h+1}^{k,\beta} - \lambda) - \text{Var}_{P_h^o}(V_{h+1}^{k,\beta} - \lambda)|}_{(iii)}, \end{aligned} \quad (\text{I.46})$$

where (I.43) is from  $\inf f(x) - \inf g(x) \geq \inf(f - g)(x)$ , (I.44) is from the induction assumption, we plug in the dual formulation (I.37) in (I.45).

For term (i) in (I.46), from Lemma L.1 and a union bound, with probability at least  $1 - \delta$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| \leq \sqrt{\frac{\ln(2SAHK/\delta)}{2n_h^{k-1}(s, a) \vee 1}}, \quad (\text{I.47})$$

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

To bound term (ii) in (I.46), following the same discussion as (H.63), with probability at least  $1 - \delta$ , we have

$$\begin{aligned} \sup_{\boldsymbol{\lambda} \in [0, H]} |\mathbb{E}_{\hat{P}_h^k}[V_{h+1}^{k, \rho} - \boldsymbol{\lambda}] - \mathbb{E}_{P_h^o}[V_{h+1}^{k, \rho} - \boldsymbol{\lambda}]| &\leq \sup_{V \in \mathcal{N}_{\mathcal{V}}(\epsilon)} |\mathbb{E}_{P_h^o}[V] - \mathbb{E}_{\hat{P}_h^k}[V]| + 2\epsilon \\ &\leq H \sqrt{\frac{2S^2 \ln(12SAH^2K^2/\delta)}{n_h^{k-1} \vee 1}} + \frac{1}{K}, \end{aligned} \quad (\text{I.48})$$

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

We denote  $V(\boldsymbol{\lambda}) = V_{h+1}^{k, \rho} - \boldsymbol{\lambda} \in [-H, H]$  and  $\mathcal{V} = \{V \in \mathbb{R}^S : \|V\|_\infty \leq H\}$ . To bound term (iii) in (I.46), we create a  $\epsilon$ -net  $\mathcal{N}_{\mathcal{V}}(\epsilon)$  for  $\mathcal{V}$ . From Lemma L.4, it holds that  $\ln |\mathcal{N}_{\mathcal{V}}(\epsilon)| \leq |S| \cdot \ln(3H/\epsilon)$ .

Therefore, by the definition of  $\mathcal{N}_{\mathcal{V}}(\epsilon)$ , for any fixed  $V$ , there exists a  $V' \in \mathcal{N}_{\mathcal{V}}(\epsilon)$  such that  $\|V - V'\|_\infty \leq \epsilon$ , that is

$$\begin{aligned} |\text{Var}_{P_h^o}(V) - \text{Var}_{\hat{P}_h^k}(V)| &\leq |\text{Var}_{P_h^o}(V) - \text{Var}_{P_h^o}(V')| + |\text{Var}_{P_h^o}(V') - \text{Var}_{\hat{P}_h^k}(V')| + |\text{Var}_{\hat{P}_h^k}(V') - \text{Var}_{\hat{P}_h^k}(V)| \\ &\leq |\mathbb{E}_{P_h^o}[V^2 - V'^2]| + |\mathbb{E}_{P_h^o}^2[V] - \mathbb{E}_{P_h^o}^2[V']| + |\mathbb{E}_{\hat{P}_h^k}[V'^2 - V^2]| + |\mathbb{E}_{\hat{P}_h^k}^2[V'] - \mathbb{E}_{\hat{P}_h^k}^2[V]| \\ &\quad + |\text{Var}_{P_h^o}(V') - \text{Var}_{\hat{P}_h^k}(V')| \\ &\leq |\text{Var}_{P_h^o}(V') - \text{Var}_{\hat{P}_h^k}(V')| + 8H\epsilon \\ &\leq \sup_{V' \in \mathcal{N}_{\mathcal{V}}} |\text{Var}_{P_h^o}(V') - \text{Var}_{\hat{P}_h^k}(V')| + 8H\epsilon. \end{aligned} \quad (\text{I.49})$$

For any fixed  $V$ , following the same analysis as (H.14), we have

$$\begin{aligned} |\text{Var}_{P_h^o}(V) - \text{Var}_{\hat{P}_h^k}(V)| &= |(\mathbb{E}_{P_h^o}[V^2] - \mathbb{E}_{P_h^o}^2[V]) - (\mathbb{E}_{\hat{P}_h^k}[V^2] - \mathbb{E}_{\hat{P}_h^k}^2[V])| \\ &\leq |\mathbb{E}_{P_h^o}[V^2] - \mathbb{E}_{\hat{P}_h^k}[V^2]| + |\mathbb{E}_{P_h^o}^2[V] - \mathbb{E}_{\hat{P}_h^k}^2[V]| \\ &\leq H^2 \sqrt{\frac{2S \ln(2/\delta)}{n_h^{k-1} \vee 1}} + (\mathbb{E}_{P_h^o}[V] + \mathbb{E}_{\hat{P}_h^k}[V]) \cdot |\mathbb{E}_{P_h^o}[V] - \mathbb{E}_{\hat{P}_h^k}[V]| \\ &\leq H^2 \sqrt{\frac{2S \ln(2/\delta)}{n_h^{k-1} \vee 1}} + 2H^2 \sqrt{\frac{2S \ln(2/\delta)}{n_h^{k-1} \vee 1}} \\ &\leq 3H^2 \sqrt{\frac{2S \ln(2/\delta)}{n_h^{k-1} \vee 1}} \end{aligned} \quad (\text{I.50})$$

with probability at least  $1 - \delta$ .

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

$$\begin{aligned} \sup_{\lambda \in [0, H]} |\text{Var}_{\hat{P}_h^k}(V_{h+1}^{k, \beta} - \lambda) - \text{Var}_{P_h^o}(V_{h+1}^{k, \beta} - \lambda)| &\leq \sup_{\lambda \in [0, H]} |\text{Var}_{P_h^o}(V(\lambda)) - \text{Var}_{\hat{P}_h^k}(V(\lambda))| \\ &\leq \sup_{V \in \mathcal{N}_V(\epsilon)} |\text{Var}_{P_h^o}(V) - \text{Var}_{\hat{P}_h^k}(V)| + 8H\epsilon \end{aligned} \quad (\text{I.51})$$

$$\leq 3H^2 \sqrt{\frac{2S \ln(2SAHK|\mathcal{N}_V|/\delta)}{n_h^{k-1} \vee 1}} + 8H\epsilon \quad (\text{I.52})$$

$$\begin{aligned} &\leq 3H^2 \sqrt{\frac{2S^2 \ln(6SAH^2K/\epsilon\delta)}{n_h^{k-1} \vee 1}} + 8H\epsilon \\ &= 3H^2 \sqrt{\frac{2S^2 \ln(48SAH^3K^2/\delta)}{n_h^{k-1} \vee 1}} + \frac{1}{K}, \end{aligned} \quad (\text{I.53})$$

for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , where (I.51) follows from (I.49), (I.52) is from (I.50) and a union bound, we set  $\epsilon = \frac{1}{8HK}$  in (I.53).

Apply the union bound again and combine (I.46) with (I.47), (I.48), (I.53), the definition of bonus and induction assumption. With probability at least  $1 - 3\delta$ , we have  $Q_h^{k, \rho}(s, a) \geq Q_h^{\pi, \rho}(s, a)$  for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ . This completes the proof.  $\square$

With similar proof to Lemma H.3, we can control the item  $Q^k - Q^{\pi^k}$ .

**Lemma I.11.** *With probability at least  $1 - \delta$ , for any  $(k, s, a) \in [K] \times \mathcal{S} \times \mathcal{A}$ , the sum of estimation errors can be bounded as follows*

$$Q_1^{k, \beta} - Q_1^{\pi^k, \beta} \leq \sum_{h=1}^H \mathbb{E}_{\{P_h^{w, k}\}_{h=1}^H, \pi^k} [2\text{bonus}_h^k].$$

*Proof.* From the proof of Lemma I.2, we see that with probability at least  $1 - \delta$ , for any  $(k, h, s, a) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A}$ , we have

$$|\hat{r}_h^k(s, a) - r_h(s, a)| + \left| \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k, \beta}] + \beta\chi^2(P\|\hat{P}_h^k))(s, a) \right| \quad (\text{I.54})$$

$$- \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k, \beta}] + \beta\chi^2(P\|P_h^o))(s, a) \leq \text{bonus}_h^k(s, a). \quad (\text{I.55})$$

Recall that we define  $P_h^{w, k} = \argmin_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k, \beta}] + \beta\chi^2(P\|P_h^o))$  as the worst-case transition in

Definition D.3, we have

$$\begin{aligned} Q_h^{k, \beta} - Q_h^{\pi^k, \beta} &\leq \text{bonus}_h^k + \hat{r}_h^k + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k, \beta}] + \beta\chi^2(P\|\hat{P}_h^k)) - r_h - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi^k, \beta}] + \beta\chi^2(P\|P_h^o)) \\ &\leq 2\text{bonus}_h^k + \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{k, \beta}] + \beta\chi^2(P\|P_h^o)) - \inf_{P \in \Delta(S)} (\mathbb{E}_P[V_{h+1}^{\pi^k, \beta}] + \beta\chi^2(P\|P_h^o)) \end{aligned} \quad (\text{I.56})$$

$$= 2\text{bonus}_h^k + \mathbb{E}_{P_h^{w, k}, \pi^k} [Q_{h+1}^{k, \beta} - Q_{h+1}^{\pi^k, \beta}]. \quad (\text{I.57})$$

where (I.56) uses (I.55). Apply (I.57) recursively, we can obtain the result.  $\square$

**Theorem I.12** (Restatement of Theorem D.16 in  $\chi^2$ -divergence setting). *For RRMDP with  $(s, a)$ -rectangular  $\chi^2$ -divergence defined regularization term satisfying Assumption 3.1, with probability at least  $1 - \delta$ , the regret of Algorithm 1 satisfies*

$$\text{Regret} = \tilde{O}\left(\left(1 + \frac{H}{\beta}\right)(C_{vr}S^2AH^2 + C_{vr}^{\frac{1}{2}}S^{\frac{3}{2}}A^{\frac{1}{2}}H^2\sqrt{K})\right).$$

*Proof.* Setting  $\delta' = \delta/5$  in Lemmas H.4 and I.10, then with probability at least  $1 - \delta$ , we get

$$\begin{aligned} \text{Regret} &= \sum_{k=1}^K (V_1^{*,\beta} - V_1^{\pi^k,\beta}) = \sum_{k=1}^K (V_1^{*,\beta} - V_1^{k,\beta}) + \sum_{k=1}^K (V_1^{k,\beta} - V_1^{\pi^k,\beta}) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} [2\text{bonus}_h^k] \end{aligned} \quad (\text{I.58})$$

$$= 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{P_h^{w,k}, \pi^k} \left[ \left( 2 + \frac{3H}{4\beta} \right) H \sqrt{\frac{2S^2 \ln(48SAH^3K^2/\delta)}{n_h^{k-1}(s, a) \vee 1}} + \left( 1 + \frac{1}{4\beta} \right) \frac{1}{K} \right] \quad (\text{I.59})$$

$$= \tilde{\mathcal{O}} \left( \left( 1 + \frac{H}{\beta} \right) (C_{vr} S^2 A H^2 + C_{vr}^{\frac{1}{2}} S^{\frac{3}{2}} A^{\frac{1}{2}} H^2 \sqrt{K}) \right), \quad (\text{I.60})$$

where (I.58) is the combination of Lemma I.10 and Lemma I.11, we plug in the bonus (I.42) in (I.59), (I.60) is from Lemma H.9.  $\square$

## J Proofs of Results in Lower Bounds

### J.1 Proof of Lower Bounds on the Regret

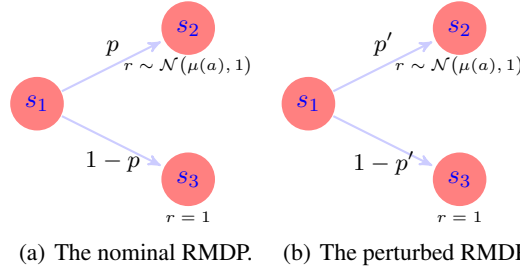


Figure 6: Constructions of the nominal environment and the worst-case environment in the proofs of regret lower bounds, the value on each arrow represents the transition probability.

**Theorem J.1** (Combination of Theorem 3.3 and Theorem D.18). *For CRMDPs with TV, KL and  $\chi^2$  divergence defined uncertainty sets and RRMDPs with TV, KL and  $\chi^2$  divergence defined regularization terms, for any learning algorithm  $\xi$ , there exists a RMDP  $\mathcal{M}$  satisfying Assumption 3.1, such that  $\text{Regret}_{\mathcal{M}}(\xi, K) = \Omega(C_{vr}^{\frac{1}{2}} \sqrt{K})$ .*

*Proof.* The proof here follows the high level idea in [13, Theorem 15.2]. We consider two RMDPs  $\mathcal{M}_1$  and  $\mathcal{M}_2$ , as illustrated in Figure 6, where  $H = 2$ ,  $\mathcal{S} = \{s_1, s_2, s_3\}$  and  $\mathcal{A} = \{a_1, a_2, \dots, a_{|\mathcal{A}|}\}$ . The only difference between these two RMDPs is the mean reward at  $s_2$ .  $s_1$  is always the initial state, and can transit to  $s_2$  with fixed probability  $p$  and  $s_3$  with fixed probability  $1 - p$  regardless of the action. We assume that  $K > |\mathcal{A}|$  and  $p \geq |\mathcal{A}|/K$  to facilitate the constructions, this bound becomes loose as  $K$  grows large.

The agent receives a reward drawn from a unit normal distribution,  $r \sim \mathcal{N}(\mu(a), 1)$ , at state  $s_2$ , and a fixed reward  $r = 1$  at state  $s_3$ , where  $\mu(a) \in [0, 1)$ . This choice of  $\mu(a)$  ensures that the robust value function satisfies  $V^{\pi, \rho}(s_2) < V^{\pi, \rho}(s_3)$ , implying that the transition probability from  $s_1$  to  $s_2$  will not decrease in the worst-case environment compared to the nominal environment. In  $\mathcal{M}_1$ , the mean reward vector at  $s_2$  is given by  $\mu_1 = (\Delta, 0, \dots, 0) \in \mathbb{R}^{|\mathcal{A}|}$  and  $\Delta < \frac{1}{2}$  is a constant to be specified.

We introduce some notations that will be useful in the following discussion. The agent follows a learning algorithm  $\xi$  and, in the  $k$ -th episode, selects a policy  $\pi^k$  to interact with the environment and collect a trajectory  $(o^k, a^k, r^k)$ , where  $o^k$  denotes the state to which the agent transitions from  $s_1$  after taking an arbitrary action;  $a^k$  denotes the action the agent takes at step  $h = 2$ , i.e., at state  $s_2$  or  $s_3$ ;

and  $r^k$  denotes the corresponding reward received at  $s_2$  or  $s_3$ , all in episode  $k$ . The joint distribution over the trajectories collected across all  $K$  episodes,  $\{(o^k, a^k, r^k)\}_{k=1}^K$ , induces a distribution denoted by  $\mathbb{P}_i^o$ .

The random variables corresponding to  $o^k, a^k, r^k$  are denoted by  $O^k, A^k, R^k$ , respectively. We use  $\mathbb{E}_i^o$  to denote the expectation under  $\mathbb{P}_i^o$ . As similarly noted by [2], it suffices to consider deterministic strategies without loss of generality. Let  $S_j(K)$  denote the number of episodes in which the agent selects action  $a_j$  if visiting state  $s_2$  according to the learned policy

$$S_j(K) = \sum_{k=1}^K \mathbb{1}\{\pi^k(s_2) = a_j\}, \quad (\text{J.1})$$

and let  $T_j(K)$  denote the number of episodes in which action  $a_j$  is actually taken at state  $s_2$ ,

$$T_j(K) = \sum_{k=1}^K \mathbb{1}\{O^k = s_2, A^k = a_j\}. \quad (\text{J.2})$$

We set  $c = \underset{j>1}{\operatorname{argmin}} \mathbb{E}_1^o[T_j(n)]$  as the least chosen action (excluding  $a_1$ ) under environment  $\mathcal{M}_1$ .

Since the expected number of times the agent transitions from  $s_1$  to  $s_2$  over  $K$  episodes is  $Kp$ , and given the selection of  $c$ , it follows that

$$\mathbb{E}_1^o[T_c(n)] \leq \frac{Kp}{|\mathcal{A}| - 1}. \quad (\text{J.3})$$

We now define the mean reward at  $s_2$  in environment  $\mathcal{M}_2$  as

$$\mu_2 = (\Delta, 0, \dots, 0, \underset{c\text{-th}}{2\Delta}, 0, \dots, 0).$$

That is,  $\mu_2(a_j) = \mu_1(a_j)$  for all  $j \neq c$  and  $\mu_2(a_c) = 2\Delta$ . Clearly, the optimal policy at state  $s_2$  selects action  $a_1$  in  $\mathcal{M}_1$  and action  $a_c$  in  $\mathcal{M}_2$ .

Recall the definition of  $C_{vr}$  in Assumption 3.1 and the visitation measure notations introduced in Definition D.3. Since the visitation measure to  $s_1$  is always 1, and the visitation measure to  $s_3$ , which is equivalently the transition probability from  $s_1$  to  $s_3$ , does not increase in the worst-case environment compared to the nominal one, the maximum visitation measure ratio can only be attained at state  $s_2$ . In the nominal environment, the visitation measure to  $s_2$  is  $d^\pi(s_2) = p$  for any policy  $\pi$ . In the worst-case environment of  $\mathcal{M}_1$ , let  $\tilde{p} = \max_{\pi} q^\pi(s_2)$  denote the maximum visitation measure to  $s_2$ , that is, the largest transition probability from  $s_1$  to  $s_2$  across all policies, under both the CRMDP and RRMDP settings. In what follows, we show that this value  $\tilde{p}$  remains unchanged in  $\mathcal{M}_2$ .

To proceed, we present the following lemma, which provides a lower bound on the sum of the total regret in environments  $\mathcal{M}_1$  and  $\mathcal{M}_2$ .

**Lemma J.2.** *For TV, KL and  $\chi^2$  divergence defined CRMDP and RRMDP settings, there exists a corresponding constant  $\zeta$  such that*

$$\mathbb{E}[\operatorname{Regret}_{\mathcal{M}_1}(\xi, K) + \operatorname{Regret}_{\mathcal{M}_2}(\xi, K)] \geq \zeta \cdot \tilde{p} \cdot \exp(-\operatorname{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o)) K \Delta, \quad (\text{J.4})$$

$$\text{where } \zeta = \begin{cases} \frac{1}{4} & \text{for CRMDP settings,} \\ \frac{1}{4} & \text{for RRMDP-TV setting,} \\ \frac{e^{-\beta^{-1}}}{4(1+\beta^{-1})} & \text{for RRMDP-KL setting,} \\ \frac{1}{16} & \text{for RRMDP-}\chi^2 \text{ setting.} \end{cases}$$

We present the following lemma to simplify the expression for  $\operatorname{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o)$ .

**Lemma J.3.** *For the distributions  $\mathbb{P}_1^o$  and  $\mathbb{P}_2^o$ , the following property holds:*

$$\operatorname{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o) = \sum_{j=1}^{|\mathcal{A}|} \mathbb{E}_1^o[T_j(K)] \cdot \operatorname{KL}(P_{r_{\mathcal{M}_1}(s_2, a_j)}, P_{r_{\mathcal{M}_2}(s_2, a_j)}), \quad (\text{J.5})$$

where  $T_j(K)$  is defined in (J.2).

From the construction of reward function, we can simplify (J.5) as

$$\begin{aligned}
\text{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o) &= \sum_{j=1}^{|\mathcal{A}|} \mathbb{E}_1^o[T_j(K)] \cdot \text{KL}(r_{\mathcal{M}_1}(s_2, a_j), r_{\mathcal{M}_2}(s_2, a_j)) \\
&= \mathbb{E}_1^o[T_c(K)] \cdot \text{KL}(\mathcal{N}(0, 1), \mathcal{N}(2\Delta, 1)) \\
&= 2\Delta^2 \cdot \mathbb{E}_1^o[T_c(K)] \\
&\leq \frac{2\Delta^2 Kp}{|\mathcal{A}| - 1},
\end{aligned} \tag{J.6}$$

where (J.6) follows from Lemma L.7, and (J.7) applies the result of (J.3).

We now return to bounding (J.4). By plugging in (J.7) and setting  $\Delta = \sqrt{(|\mathcal{A}| - 1)/(4Kp)}$ , which satisfies  $\Delta < 1/2$  due to the given range of  $K$  and  $p$ , we obtain

$$\begin{aligned}
2 \max \{ \mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K)], \mathbb{E}[\text{Regret}_{\mathcal{M}_2}(\xi, K)] \} &\geq \mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K) + \text{Regret}_{\mathcal{M}_2}(\xi, K)] \\
&\geq \zeta \cdot \tilde{p} \cdot \exp(-\text{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o)) K \Delta \\
&\geq \zeta \cdot \tilde{p} \cdot \exp\left(-\frac{2\Delta^2 Kp}{|\mathcal{A}| - 1}\right) K \Delta \\
&\geq \Omega\left(\frac{e^{-\frac{1}{2}} \sqrt{\tilde{p}}}{2} \sqrt{C_{vr}(|\mathcal{A}| - 1)K}\right) \\
&= \Omega(C_{vr}^{\frac{1}{2}} \sqrt{K}).
\end{aligned} \tag{J.8}$$

This finishes the proof.  $\square$

## J.2 Proof of Lemma D.14

To establish the results for hard instances, we instantiate the constructed example in Theorem J.1 by appropriately selecting the nominal transition  $p$  and the uncertainty set radius  $\rho$ . We assume that  $K \geq 2^{2|\mathcal{A}|+1}|\mathcal{A}|$  throughout the proof.

**For TV-distance.** We set  $p = \frac{1}{2^{2|\mathcal{A}|+1}-2} \geq |\mathcal{A}|/K$  and  $\rho = \frac{1}{2}$ . First, we solve the worst-case transition in the first step explicitly. Follow the definition of TV-distance, we have

$$\begin{aligned}
\tilde{p} &= \underset{p'}{\text{argmin}} p' V^\pi(s_2) + (1 - p') \\
\text{s.t. } & p' - p \leq \rho.
\end{aligned} \tag{J.9}$$

Since (J.9) decrease monotonically with  $\tilde{p}$ , the solution of this optimization problem is  $\tilde{p} = p + \rho = \frac{2^{2|\mathcal{A}|}}{2^{2|\mathcal{A}|+1}-2}$ . Therefore,  $C_{vr} = \tilde{p}/p = 2^{2|\mathcal{A}|}$ .

From (J.8), we see that the expected regret for constructed environments  $\mathcal{M}_1$  and  $\mathcal{M}_2$  satisfies

$$\max \{ \mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K)], \mathbb{E}[\text{Regret}_{\mathcal{M}_2}(\xi, K)] \} \geq \Omega\left(\frac{e^{-\frac{1}{2}} \sqrt{\tilde{p}}}{2} \sqrt{C_{vr}(|\mathcal{A}| - 1)K}\right) = \Omega(2^{|\mathcal{A}|} \sqrt{K}).$$

This finishes the proof.

**For  $\chi^2$ -divergence.** We set  $p = \frac{1}{2^{2|\mathcal{A}|+1}} \geq |\mathcal{A}|/K$  and  $\rho = \frac{(2^{2|\mathcal{A}|}-1)^2}{2^{2|\mathcal{A}|+1}-1}$ . First, we solve the worst-case transition in the first step explicitly. Follow the definition of  $\chi^2$ -divergence, we have

$$\begin{aligned}
\tilde{p} &= \underset{p'}{\text{argmin}} p' V^\pi(s_2) + (1 - p') \\
\text{s.t. } & p \left(\frac{p'}{p} - 1\right)^2 + (1 - p) \left(\frac{1 - p'}{1 - p} - 1\right)^2 \leq \rho.
\end{aligned} \tag{J.10}$$



Since (J.10) decrease monotonically with  $\tilde{p}$ , the solution of this optimization problem is  $\tilde{p} = p + \sqrt{\rho p(1-p)} = \frac{1}{2}$ . Therefore,  $C_{vr} = \tilde{p}/p = 2^{2|\mathcal{A}|}$ .

From (J.8), we see that the expected regret for constructed environments  $\mathcal{M}_1$  and  $\mathcal{M}_2$  satisfies

$$\max \{ \mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K)], \mathbb{E}[\text{Regret}_{\mathcal{M}_2}(\xi, K)] \} \geq \Omega \left( \frac{e^{-\frac{1}{2}} \sqrt{\tilde{p}}}{2} \sqrt{C_{vr}(|\mathcal{A}| - 1)K} \right) = \Omega(2^{|\mathcal{A}|} \sqrt{K}).$$

This finishes the proof.

**For KL-divergence.** We set  $p = \frac{1}{2^{2|\mathcal{A}|+1}} \geq |\mathcal{A}|/K$  and  $\rho = \frac{(2^{2|\mathcal{A}|}-1)^2}{2^{2|\mathcal{A}|+1}-1}$ . First, we solve the worst-case transition in the first step explicitly. Follow the definition of KL-divergence, we have

$$\begin{aligned} \tilde{p} &= \underset{p'}{\operatorname{argmin}} p' V^\pi(s_2) + (1 - p') \\ \text{s.t. } & p' \log \left( \frac{p'}{p} \right) + (1 - p) \log \left( \frac{1 - p'}{1 - p} \right) \leq \rho. \end{aligned}$$

Though it is difficult to obtain a closed-form solution for this optimization problem, Lemma L.8 implies that its optimal value is no less than that of the previous  $\chi^2$ -based optimization problem. That is,  $\tilde{p} \geq p + \sqrt{\rho p(1-p)} = \frac{1}{2}$ . Therefore,  $C_{vr} = \tilde{p}/p \geq 2^{2|\mathcal{A}|}$ .

From (J.8), we see that the expected regret for constructed environments  $\mathcal{M}_1$  and  $\mathcal{M}_2$  satisfies

$$\max \{ \mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K)], \mathbb{E}[\text{Regret}_{\mathcal{M}_2}(\xi, K)] \} \geq \Omega \left( \frac{e^{-\frac{1}{2}} \sqrt{\tilde{p}}}{2} \sqrt{C_{vr}(|\mathcal{A}| - 1)K} \right) = \Omega(2^{|\mathcal{A}|} \sqrt{K}).$$

This finishes the proof.

## K Proofs of Supporting Lemmas in the Proofs of Lower Bounds

### K.1 Proof of Lemma J.2

We provide the proofs for the three CRMDP settings in Appendix K.1.1, as they share a similar underlying structure. The proofs for the RRMDP-TV, RRMDP-KL, and RRMDP- $\chi^2$  settings are given in Appendix K.1.2, Appendix K.1.3, and Appendix K.1.4, respectively.

#### K.1.1 The CRMDP Settings

First, we consider the optimization problem for the worst-case transition from state  $s_1$  to states  $s_2$  and  $s_3$ , given by

$$\begin{aligned} P^w &= \underset{p': D(P^w \| P^o) \leq \rho}{\operatorname{argmin}} [p' \cdot V^{\pi, \rho}(s_2) + (1 - p') \cdot V^{\pi, \rho}(s_3)] \\ &= \underset{p': D(P^w \| P^o) \leq \rho}{\operatorname{argmin}} [p' \cdot V^{\pi, \rho}(s_2) + (1 - p')], \\ \text{where } & P^w = (p', 1 - p') \quad \text{and} \quad P^o = (p, 1 - p), \end{aligned}$$

where the second equality holds because  $V^{\pi, \rho}(s_2) < 1$  and  $V^{\pi, \rho}(s_3) = 1$  by construction. As a result, the objective function  $p' \cdot V^{\pi, \rho}(s_2) + (1 - p')$  decreases monotonically with respect to  $p'$ . Therefore, the worst-case transition probability is characterized by  $\tilde{p} = \sup\{p' : D(P^w \| P^o) \leq \rho\}$ , where  $D$  denotes an  $f$ -divergence (TV, KL or  $\chi^2$ ). This implies that  $\tilde{p}$  remains invariant across different policies  $\pi$  and environments  $\mathcal{M}_1$  and  $\mathcal{M}_2$ .

For environment  $\mathcal{M}_1$ , the total regret incurred by selecting a suboptimal policy  $\pi^k(s_2) \neq a_1$  in each episode can be bounded as

$$\begin{aligned}\mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K)] &= \sum_{k=1}^K \mathbb{E}_1^o[V^{\pi^*, \rho}(s_1) - V^{\pi^k, \rho}(s_1)] \\ &= \sum_{k=1}^K \mathbb{E}_1^o[(\tilde{p} \cdot V^{\pi^*, \rho}(s_2) + (1 - \tilde{p}) \cdot V^{\pi^*, \rho}(s_3)) - (\tilde{p} \cdot V^{\pi^k, \rho}(s_2) + (1 - \tilde{p}) \cdot V^{\pi^k, \rho}(s_3))] \\ &\quad \quad \quad (\text{K.1})\end{aligned}$$

$$\begin{aligned}&= \sum_{k=1}^K \mathbb{E}_1^o[\tilde{p} \cdot (V^{\pi^*, \rho}(s_2) - V^{\pi^k, \rho}(s_2))] \\ &= \sum_{k=1}^K \mathbb{E}_1^o[\tilde{p} \cdot \mathbb{1}\{\pi^k(s_2) \neq a_1\} \Delta] \\ &\quad \quad \quad (\text{K.2})\end{aligned}$$

$$\begin{aligned}&= \tilde{p} \cdot \Delta \cdot \mathbb{E}_1^o\left[\sum_{k=1}^K \mathbb{1}\{\pi^k(s_2) \neq a_1\}\right] \\ &= \tilde{p} \cdot \Delta \cdot \mathbb{E}_1^o[K - S_1(K)] \\ &\quad \quad \quad (\text{K.3})\end{aligned}$$

$$\geq \frac{\tilde{p}K\Delta}{2} \mathbb{P}_1^o\left(S_1(K) \leq \frac{K}{2}\right), \quad (\text{K.4})$$

where (K.1) follows from the definition of the robust value function under the CRMDP setting, using  $\tilde{p}$  as the worst-case transition probability, (K.2) relies on the construction of the reward function at state  $s = s_2$ , (K.3) uses the definition of  $S_j(K)$  as given in (J.1), and (K.4) follows from Markov's inequality.

For environment  $\mathcal{M}_2$ , following the same analysis, the total regret incurred by selecting a suboptimal policy  $\pi^k(s_2) \neq a_c$  in each episode can be bounded as

$$\begin{aligned}\mathbb{E}[\text{Regret}_{\mathcal{M}_2}(\xi, K)] &= \sum_{k=1}^K \mathbb{E}_2^o[V^{\pi^*, \rho}(s_1) - V^{\pi^k, \rho}(s_1)] \\ &= \sum_{k=1}^K \mathbb{E}_2^o[\tilde{p} \cdot (V^{\pi^*, \rho}(s_2) - V^{\pi^k, \rho}(s_2))] \\ &= \sum_{k=1}^K \mathbb{E}_2^o\left[\tilde{p} \cdot \left(\mathbb{1}\{\pi^k(s_2) = a_1\} \Delta + \sum_{j>1, j \neq c} \mathbb{1}\{\pi^k(s_2) = a_j\} \Delta\right)\right] \\ &\geq \sum_{k=1}^K \mathbb{E}_2^o[\tilde{p} \cdot (\mathbb{1}\{\pi^k(s_2) = a_1\} \Delta)] \\ &= \tilde{p} \cdot \Delta \cdot \mathbb{E}_2^o\left[\sum_{k=1}^K \mathbb{1}\{\pi^k(s_2) = a_1\}\right] \\ &= \tilde{p} \cdot \Delta \cdot \mathbb{E}_2^o[S_1(K)] \\ &\geq \frac{\tilde{p}K\Delta}{2} \mathbb{P}_2^o\left(S_1(K) > \frac{K}{2}\right). \\ &\quad \quad \quad (\text{K.5})\end{aligned}$$

By combining (K.4) and (K.5) with Lemma L.6, we obtain

$$\begin{aligned}\mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K) + \text{Regret}_{\mathcal{M}_2}(\xi, K)] &\geq \frac{\tilde{p}K\Delta}{2} \left( \mathbb{P}_1^o\left(S_1(K) \leq \frac{K}{2}\right) + \mathbb{P}_1^o\left(S_2(K) > \frac{K}{2}\right) \right) \\ &\geq \frac{\tilde{p}K\Delta}{4} \exp(-\text{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o)).\end{aligned}$$

We complete the proof by setting  $\zeta = \frac{1}{4}$ .

### K.1.2 The RRMDP-TV Setting

For the RRMDP-TV setting, the robust value function at state  $s_1$ , denoted by  $V^{\pi, \beta}(s_1)$ , is defined as

$$\begin{aligned} V^{\pi, \beta}(s_1) &= \inf_{p'} \{p' \cdot V^{\pi, \beta}(s_2) + (1 - p') \cdot V^{\pi, \beta}(s_3) + \beta|p' - p|\} \\ &= \inf_{p'} \{p' \cdot V^{\pi, \beta}(s_2) + (1 - p') + \beta(p' - p)\}, \end{aligned}$$

where the second equality holds because  $V^{\pi, \beta}(s_2) < 1$  and  $V^{\pi, \beta}(s_3) = 1$  by construction. Given that  $V_1^{\pi, \beta}(s_2)$  takes value in  $\{0, \Delta, 2\Delta\}$ , we set  $\beta \in [0, 1 - 2\Delta]$ , so that  $V_2^{\pi, \beta}(s_2) - 1 + \beta \leq 0$ . Therefore, the term  $[V_2^{\pi, \beta}(s_2) - 1 + \beta]p'$  decreases monotonically with  $p'$ , and the infimum is attained at  $\tilde{p} = 1$ . Consequently, we obtain

$$V^{\pi, \beta}(s_1) = V^{\pi, \beta}(s_2) + \beta(1 - p).$$

For environment  $\mathcal{M}_1$ , the total regret incurred by selecting a suboptimal policy  $\pi^k(s_2) \neq a_1$  in each episode can be bounded as

$$\begin{aligned} \mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K)] &= \sum_{k=1}^K \mathbb{E}_1^o[V^{\pi^*, \beta}(s_1) - V^{\pi^k, \beta}(s_1)] \\ &= \sum_{k=1}^K \mathbb{E}_1^o[\mathbb{1}\{\pi^k(s_2) \neq a_1\} \cdot ((\Delta + \beta(1 - p)) - \beta(1 - p))] \\ &= \tilde{p} \cdot \Delta \cdot \mathbb{E}_1^o\left[\sum_{k=1}^K \mathbb{1}\{\pi^k(s_2) \neq a_1\}\right] \tag{K.6} \end{aligned}$$

$$= \tilde{p} \cdot \Delta \cdot \mathbb{E}_1^o[K - S_1(K)] \tag{K.7}$$

$$\geq \frac{\tilde{p}K\Delta}{2} \mathbb{P}_1^o\left(S_1(K) \leq \frac{K}{2}\right), \tag{K.8}$$

where (K.6) holds because  $\tilde{p} = 1$  as previously derived, (K.7) uses the definition of  $S_j(K)$  as given in (J.1), and (K.8) follows from Markov's inequality.

For environment  $\mathcal{M}_2$ , following the same analysis, the total regret incurred by selecting a suboptimal policy  $\pi^k(s_2) \neq a_c$  in each episode can be bounded as

$$\begin{aligned} \mathbb{E}[\text{Regret}_{\mathcal{M}_2}(\xi, K)] &= \sum_{k=1}^K \mathbb{E}_2^o[V^{\pi^*, \beta}(s_1) - V^{\pi^k, \beta}(s_1)] \\ &= \sum_{k=1}^K \mathbb{E}_2^o\left[\mathbb{1}\{\pi^k(s_2) = a_1\} \cdot ((2\Delta + \beta(1 - p)) - (\Delta + \beta(1 - p)))\right. \\ &\quad \left.+ \sum_{j>1, j \neq c} \mathbb{1}\{\pi^k(s_2) = a_j\} \cdot ((2\Delta + \beta(1 - p)) - \beta(1 - p))\right] \\ &\geq \tilde{p} \cdot \Delta \cdot \mathbb{E}_2^o\left[\sum_{k=1}^K \mathbb{1}\{\pi^k(s_2) = a_1\}\right] \\ &= \tilde{p} \cdot \Delta \cdot \mathbb{E}_2^o[S_1(K)] \\ &\geq \frac{\tilde{p}K\Delta}{2} \mathbb{P}_2^o\left(S_1(K) > \frac{K}{2}\right). \tag{K.9} \end{aligned}$$

By combining (K.4) and (K.5) with Lemma L.6, we obtain

$$\begin{aligned} \mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K) + \text{Regret}_{\mathcal{M}_2}(\xi, K)] &\geq \frac{\tilde{p}K\Delta}{2} \left( \mathbb{P}_1^o\left(S_1(K) \leq \frac{K}{2}\right) + \mathbb{P}_2^o\left(S_1(K) > \frac{K}{2}\right) \right) \\ &\geq \frac{\tilde{p}K\Delta}{4} \exp(-\text{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o)). \end{aligned}$$

We complete the proof by setting  $\zeta = \frac{1}{4}$ .

### K.1.3 The RRMDP-KL Setting

For the RRMDP-KL setting, the robust value function at state  $s_1$ , denoted by  $V^{\pi, \beta}(s_1)$ , is defined as

$$V^{\pi, \beta}(s_1) = \inf_{p'} \left\{ p' \cdot V^{\pi, \beta}(s_2) + (1 - p') \cdot V^{\pi, \beta}(s_3) + \beta \left[ p' \log \left( \frac{p'}{p} \right) + (1 - p') \log \left( \frac{1 - p'}{1 - p} \right) \right] \right\}. \quad (\text{K.10})$$

Using the construction that  $V^{\pi, \beta}(s_3) = 1$  and setting the derivative of the objective function in (K.10) with respect to  $p'$  to zero, we obtain

$$p' = \frac{pe^{-\beta^{-1}V^{\pi, \beta}(s_2)}}{pe^{-\beta^{-1}V^{\pi, \beta}(s_2)} + (1 - p)e^{-\beta^{-1}}}. \quad (\text{K.11})$$

The worst-case transition probability  $\tilde{p}$  is defined as the maximum value of  $p'$  over all policies  $\pi$ . Since (K.11) is monotonically decreasing in  $V^{\pi, \beta}(s_2)$ , we set  $V^{\pi, \beta}(s_2) = 0$  in (K.11) to obtain

$$\tilde{p} = \frac{p}{p + (1 - p)e^{-\beta^{-1}}}. \quad (\text{K.12})$$

Finally, substituting (K.11) back into (K.10), we derive the expression for the robust value function  $V^{\pi, \beta}(s_1)$  as

$$V^{\pi, \beta}(s_1) = -\beta \log \left( pe^{-\beta^{-1}V^{\pi, \beta}(s_2)} + (1 - p)e^{-\beta^{-1}} \right). \quad (\text{K.13})$$

For the environment  $\mathcal{M}_1$ , if  $\pi^k(s_2) \neq a_1$ , then the episodic regret can be bounded by

$$\begin{aligned} V^{\pi^*, \beta}(s_1) - V^{\pi^k, \beta}(s_1) &= -\beta \log \left( pe^{-\beta^{-1}\Delta} + (1 - p)e^{-\beta^{-1}} \right) + \beta \log \left( p + (1 - p)e^{-\beta^{-1}} \right) \\ &= -\beta \log \left( \frac{pe^{-\beta^{-1}\Delta} + (1 - p)e^{-\beta^{-1}}}{p + (1 - p)e^{-\beta^{-1}}} \right) \\ &= -\beta \log \left( 1 + \frac{p(e^{-\beta^{-1}\Delta} - 1)}{p + (1 - p)e^{-\beta^{-1}}} \right) \\ &\geq \beta \left( \frac{p}{p + (1 - p)e^{-\beta^{-1}}} \right) (1 - e^{-\beta^{-1}\Delta}) \end{aligned} \quad (\text{K.14})$$

$$\geq \beta \left( \frac{p}{p + (1 - p)e^{-\beta^{-1}}} \right) \left( \frac{\beta^{-1}\Delta}{1 + \beta^{-1}\Delta} \right) \quad (\text{K.15})$$

$$\geq \tilde{p} \left( \frac{1}{1 + \beta^{-1}} \right) \Delta, \quad (\text{K.16})$$

where (K.14) is because  $\log(1 + x) \leq x$  for  $x \geq -1$ , (K.15) is because  $1 - e^{-x} \geq \frac{x}{1+x}$  for  $x \geq 0$ , (K.16) plugs in (K.12) and uses the selection that  $\Delta \leq \frac{1}{2}$ .

Summing (K.16) over all  $K$  episodes, we obtain

$$\begin{aligned} \mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K)] &= \sum_{k=1}^K \mathbb{E}_1^o[V^{\pi^*, \beta}(s_1) - V^{\pi^k, \beta}(s_1)] \\ &= \sum_{k=1}^K \mathbb{E}_1^o \left[ \mathbb{1}\{\pi^k(s_2) \neq a_1\} \cdot \tilde{p} \left( \frac{1}{1 + \beta^{-1}} \right) \Delta \right] \\ &= \frac{\tilde{p}\Delta}{1 + \beta^{-1}} \cdot \mathbb{E}_1^o \left[ \sum_{k=1}^K \mathbb{1}\{\pi^k(s_2) \neq a_1\} \right] \\ &= \frac{\tilde{p}\Delta}{1 + \beta^{-1}} \cdot \mathbb{E}_1^o[K - S_1(K)] \end{aligned} \quad (\text{K.17})$$

$$\geq \frac{\tilde{p}K\Delta}{2(1 + \beta^{-1})} \mathbb{P}_1^o \left( S_1(K) \leq \frac{K}{2} \right), \quad (\text{K.18})$$

where (K.17) uses the definition of  $S_j(K)$  as given in (J.1), and (K.18) follows from Markov's inequality.

For the environment  $\mathcal{M}_2$ , if  $\pi^k(s_2) = a_1$ , then the episodic regret can be bounded by

$$\begin{aligned}
V^{\pi^*, \beta}(s_1) - V^{\pi^k, \beta}(s_1) &= -\beta \log \left( pe^{-2\beta^{-1}\Delta} + (1-p)e^{-\beta^{-1}} \right) + \beta \log \left( pe^{-\beta^{-1}\Delta} + (1-p)e^{-\beta^{-1}} \right) \\
&= -\beta \log \left( \frac{pe^{-2\beta^{-1}\Delta} + (1-p)e^{-\beta^{-1}}}{pe^{-\beta^{-1}\Delta} + (1-p)e^{-\beta^{-1}}} \right) \\
&= -\beta \log \left( 1 + \frac{p(e^{-2\beta^{-1}\Delta} - e^{-\beta^{-1}})}{pe^{-\beta^{-1}\Delta} + (1-p)e^{-\beta^{-1}}} \right) \\
&\geq \beta \left( \frac{p}{pe^{-\beta^{-1}\Delta} + (1-p)e^{-\beta^{-1}}} \right) (1 - e^{-\beta^{-1}\Delta}) e^{-\beta^{-1}\Delta} \\
&\geq \beta e^{-\beta^{-1}} \left( \frac{p}{p + (1-p)e^{-\beta^{-1}}} \right) \left( \frac{\beta^{-1}\Delta}{1 + \beta^{-1}\Delta} \right) \tag{K.19} \\
&\geq \tilde{p} \left( \frac{e^{-\beta^{-1}}}{1 + \beta^{-1}} \right) \Delta, \tag{K.20}
\end{aligned}$$

where (K.19) uses  $e^{-\beta^{-1}} \leq e^{-\beta^{-1}\Delta} \leq 1$ , (K.20) plugs in (K.12) and uses the selection that  $\Delta \leq \frac{1}{2}$ .

Summing (K.20) over all  $K$  episodes, following the same analysis, we obtain

$$\begin{aligned}
\mathbb{E}[\text{Regret}_{\mathcal{M}_2}(\xi, K)] &= \sum_{k=1}^K \mathbb{E}_2^o[V^{\pi^*, \beta}(s_1) - V^{\pi^k, \beta}(s_1)] \\
&\geq \sum_{k=1}^K \mathbb{E}_2^o \left[ \mathbb{1}\{\pi^k(s_2) = a_1\} \cdot \tilde{p} \left( \frac{e^{-\beta^{-1}}}{1 + \beta^{-1}} \right) \Delta \right] \\
&= \frac{e^{-\beta^{-1}} \tilde{p} \Delta}{1 + \beta^{-1}} \cdot \mathbb{E}_2^o \left[ \sum_{k=1}^K \mathbb{1}\{\pi^k(s_2) = a_1\} \right] \\
&= \frac{e^{-\beta^{-1}} \tilde{p} \Delta}{1 + \beta^{-1}} \cdot \mathbb{E}_2^o[S_1(K)] \\
&\geq \frac{e^{-\beta^{-1}} \tilde{p} K \Delta}{2(1 + \beta^{-1})} \mathbb{P}_2^o \left( S_1(K) > \frac{K}{2} \right). \tag{K.21}
\end{aligned}$$

By combining (K.18) and (K.21) with Lemma L.6, we obtain

$$\begin{aligned}
\mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K) + \text{Regret}_{\mathcal{M}_2}(\xi, K)] &\geq \frac{e^{-\beta^{-1}} \tilde{p} K \Delta}{2(1 + \beta^{-1})} \left( \mathbb{P}_1^o \left( S_1(K) \leq \frac{K}{2} \right) + \mathbb{P}_2^o \left( S_1(K) > \frac{K}{2} \right) \right) \\
&\geq \frac{e^{-\beta^{-1}} \tilde{p} K \Delta}{4(1 + \beta^{-1})} \exp(-\text{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o)).
\end{aligned}$$

We complete the proof by setting  $\zeta = \frac{e^{-\beta^{-1}}}{4(1 + \beta^{-1})}$ .

#### K.1.4 The RRMDP- $\chi^2$ Setting

For the RRMDP- $\chi^2$  setting, the robust value function at state  $s_1$ , denoted by  $V^{\pi, \beta}(s_1)$ , is defined as

$$V^{\pi, \beta}(s_1) = \inf_{p'} \left\{ p' \cdot V^{\pi, \beta}(s_2) + (1 - p') \cdot V^{\pi, \beta}(s_3) + \beta \left[ p \left( \frac{p'}{p} - 1 \right)^2 + (1 - p) \left( \frac{1 - p'}{1 - p} - 1 \right)^2 \right] \right\}. \tag{K.22}$$

Using the construction that  $V^{\pi,\beta}(s_3) = 1$  and setting the derivative of the objective function in (K.22) with respect to  $p'$  to zero, we obtain

$$p' = p + \frac{p(1-p)(1-V^{\pi,\beta}(s_2))}{2\beta}. \quad (\text{K.23})$$

The worst-case transition probability  $\tilde{p}$  is defined as the maximum value of  $p'$  over all policies  $\pi$ . Since (K.23) is monotonically decreasing in  $V^{\pi,\beta}(s_2)$ , we set  $V^{\pi,\beta}(s_2) = 0$  in (K.23) to obtain

$$\tilde{p} = p + \frac{p(1-p)}{2\beta}. \quad (\text{K.24})$$

Finally, substituting (K.23) back into (K.22), we derive the expression for the robust value function  $V^{\pi,\beta}(s_1)$  as

$$V^{\pi,\beta}(s_1) = pV^{\pi,\beta}(s_2) + (1-p) - \frac{p(1-p)(1-V^{\pi,\beta}(s_2))^2}{4\beta}. \quad (\text{K.25})$$

For the environment  $\mathcal{M}_1$ , if  $\pi^k(s_2) \neq a_1$ , then the episodic regret can be bounded by

$$\begin{aligned} V^{\pi^*,\beta}(s_1) - V^{\pi,\beta}(s_1) &= \left( p\Delta + (1-p) - \frac{p(1-p)(1-\Delta)^2}{4\beta} \right) - \left( (1-p) - \frac{p(1-p)}{4\beta} \right) \\ &= p\Delta + \frac{p(1-p)}{4\beta}(2\Delta - \Delta^2) \\ &\geq \frac{1}{2}p\Delta + \frac{p(1-p)}{4\beta}\Delta \\ &\geq \frac{1}{2}\tilde{p}\Delta, \end{aligned} \quad (\text{K.26})$$

$$\geq \frac{1}{2}\tilde{p}\Delta, \quad (\text{K.27})$$

where (K.26) uses the selection that  $\Delta \leq \frac{1}{2}$  and (K.27) plugs in (K.24).

Summing (K.27) over all  $K$  episodes, we obtain

$$\begin{aligned} \mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K)] &= \sum_{k=1}^K \mathbb{E}_1^o[V^{\pi^*,\beta}(s_1) - V^{\pi^k,\beta}(s_1)] \\ &= \sum_{k=1}^K \mathbb{E}_1^o\left[\mathbb{1}\{\pi^k(s_2) \neq a_1\} \cdot \frac{\tilde{p}}{2}\Delta\right] \\ &= \frac{\tilde{p}\Delta}{2} \cdot \mathbb{E}_1^o\left[\sum_{k=1}^K \mathbb{1}\{\pi^k(s_2) \neq a_1\}\right] \\ &= \frac{\tilde{p}\Delta}{2} \cdot \mathbb{E}_1^o[K - S_1(K)] \\ &\geq \frac{\tilde{p}K\Delta}{4} \mathbb{P}_1^o\left(S_1(K) \leq \frac{K}{2}\right), \end{aligned} \quad (\text{K.28})$$

$$\geq \frac{\tilde{p}K\Delta}{4} \mathbb{P}_1^o\left(S_1(K) \leq \frac{K}{2}\right), \quad (\text{K.29})$$

where (K.28) uses the definition of  $S_j(K)$  as given in (J.1), and (K.29) follows from Markov's inequality.

For the environment  $\mathcal{M}_2$ , if  $\pi^k(s_2) = a_1$ , then the episodic regret can be bounded by

$$\begin{aligned} V^{\pi^*,\beta}(s_1) - V^{\pi,\beta}(s_1) &\geq \left( 2p\Delta + (1-p) - \frac{p(1-p)(1-2\Delta)^2}{4\beta} \right) - \left( p\Delta + (1-p) - \frac{p(1-p)(1-\Delta)^2}{4\beta} \right) \\ &= p\Delta + \frac{p(1-p)}{4\beta}(2\Delta - 3\Delta^2) \\ &= p\Delta + \frac{p(1-p)}{4\beta}\Delta(2 - 3\Delta) \\ &\geq \frac{1}{4}p\Delta + \frac{p(1-p)}{8\beta}\Delta \\ &\geq \frac{1}{4}\tilde{p}\Delta, \end{aligned} \quad (\text{K.30})$$

$$\geq \frac{1}{4}\tilde{p}\Delta, \quad (\text{K.31})$$

where (K.30) uses the selection that  $\Delta \leq \frac{1}{2}$  and (K.31) plugs in (K.24).

Summing (K.31) over all  $K$  episodes, following the same analysis, we obtain

$$\begin{aligned}
\mathbb{E}[\text{Regret}_{\mathcal{M}_2}(\xi, K)] &= \sum_{k=1}^K \mathbb{E}_2^o[V^{\pi^*, \beta}(s_1) - V^{\pi^k, \beta}(s_1)] \\
&\geq \sum_{k=1}^K \mathbb{E}_2^o\left[\mathbb{1}\{\pi^k(s_2) = a_1\} \cdot \frac{\tilde{p}}{4} \Delta\right] \\
&= \frac{\tilde{p} \Delta}{4} \cdot \mathbb{E}_2^o\left[\sum_{k=1}^K \mathbb{1}\{\pi^k(s_2) = a_1\}\right] \\
&= \frac{\tilde{p} \Delta}{4} \cdot \mathbb{E}_2^o[S_1(K)] \\
&\geq \frac{\tilde{p} K \Delta}{8} \mathbb{P}_2^o\left(S_1(K) > \frac{K}{2}\right). \tag{K.32}
\end{aligned}$$

By combining (K.29) and (K.32) with Lemma L.6, we obtain

$$\begin{aligned}
\mathbb{E}[\text{Regret}_{\mathcal{M}_1}(\xi, K) + \text{Regret}_{\mathcal{M}_2}(\xi, K)] &\geq \frac{\tilde{p} K \Delta}{8} \left( \mathbb{P}_1^o\left(S_1(K) \leq \frac{K}{2}\right) + \mathbb{P}_2^o\left(S_1(K) > \frac{K}{2}\right) \right) \\
&\geq \frac{\tilde{p} K \Delta}{16} \exp(-\text{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o)).
\end{aligned}$$

We complete the proof by setting  $\zeta = \frac{1}{16}$ .

## K.2 Proof of Lemma J.3

In order to prove this, we borrow the same technology as [13, Lemma 15.1]. First, by the definition of KL-divergence, we have that

$$\text{KL}(\mathbb{P}_1^o, \mathbb{P}_2^o) = \mathbb{E}_1^o\left[\log\left(\frac{d\mathbb{P}_1^o}{d\mathbb{P}_2^o}\right)\right]. \tag{K.33}$$

Following the notation defined in Theorem J.1, we calculate the Radon-Nikodym derivative of  $\mathbb{P}_1^o$  as follows

$$p_1^o(o^1, a^1, r^1, \dots, o^K, a^K, r^K) = \prod_{k=1}^K \Pr(o^k | s_1) \cdot \pi^k(a^k | o^1, a^1, r^1, \dots, o^{k-1}, a^{k-1}, r^{k-1}, o^k) \cdot \Pr(r_{\mathcal{M}_1}(o^k, a^k) = r^k),$$

The density of  $\mathbb{P}_2$  is exactly identical except that  $r_{\mathcal{M}_1}$  is replaced by  $r_{\mathcal{M}_2}$ , which gives rise to

$$\begin{aligned}
\log\left(\frac{d\mathbb{P}_1^o}{d\mathbb{P}_2^o}(o^1, a^1, r^1, \dots, o^K, a^K, r^K)\right) &= \sum_{k=1}^K \log \frac{\Pr(r_{\mathcal{M}_1}(o^k, a^k) = r^k)}{\Pr(r_{\mathcal{M}_2}(o^k, a^k) = r^k)} \\
&= \sum_{k=1}^K \mathbb{1}\{o^k = s_2\} \cdot \log \frac{\Pr(r_{\mathcal{M}_1}(s_2, a^k) = r^k)}{\Pr(r_{\mathcal{M}_2}(s_2, a^k) = r^k)}, \tag{K.34}
\end{aligned}$$

where (K.34) is because the agent receives a fixed reward  $r = 1$  when  $o^k = s_3$ .

Taking expectations on both sides of (K.34), we obtain

$$\mathbb{E}_1^o \left[ \log \left( \frac{d\mathbb{P}_1^o}{d\mathbb{P}_2^o}(O^1, A^1, R^1, \dots, O^K, A^K, R^K) \right) \right] \quad (\text{K.35})$$

$$\begin{aligned} &= \sum_{k=1}^K \mathbb{E}_1^o \left[ \mathbb{1}\{O^k = s_2\} \cdot \log \left( \frac{P_{r_{\mathcal{M}_1}(s_2, A^k)}(R^k)}{P_{r_{\mathcal{M}_2}(s_2, A^k)}(R^k)} \right) \right] \\ &= \sum_{k=1}^K \mathbb{E}_1^o \left[ \mathbb{E}_1^o \left[ \mathbb{1}\{O^k = s_2\} \cdot \log \left( \frac{P_{r_{\mathcal{M}_1}(s_2, A^k)}(R^k)}{P_{r_{\mathcal{M}_2}(s_2, A^k)}(R^k)} \right) \middle| O^k, A^k \right] \right] \\ &= \sum_{k=1}^K \mathbb{E}_1^o \left[ \mathbb{1}\{O^k = s_2\} \cdot \mathbb{E}_1^o \left[ \log \left( \frac{P_{r_{\mathcal{M}_1}(s_2, A^k)}(R^k)}{P_{r_{\mathcal{M}_2}(s_2, A^k)}(R^k)} \right) \middle| O^k, A^k \right] \right] \end{aligned} \quad (\text{K.36})$$

$$= \sum_{k=1}^K \mathbb{E}_1^o [\mathbb{1}\{O^k = s_2\} \cdot \text{KL}(P_{r_{\mathcal{M}_1}(s_2, A^k)}, P_{r_{\mathcal{M}_2}(s_2, A^k)})] \quad (\text{K.37})$$

$$\begin{aligned} &= \sum_{j=1}^{|\mathcal{A}|} \mathbb{E}_1^o \left[ \sum_{k=1}^K \mathbb{1}\{O^k = s_2, A^k = a_j\} \cdot \text{KL}(P_{r_{\mathcal{M}_1}(s_2, a_j)}, P_{r_{\mathcal{M}_2}(s_2, a_j)}) \right] \\ &= \sum_{j=1}^{|\mathcal{A}|} \mathbb{E}_1^o [T_j(K)] \cdot \text{KL}(P_{r_{\mathcal{M}_1}(s_2, a_j)}, P_{r_{\mathcal{M}_2}(s_2, a_j)}), \end{aligned} \quad (\text{K.38})$$

where (K.36) is because  $\mathbb{1}\{O^k = s_2\}$  is measurable with respect to the  $\sigma$ -field generated by  $O^k$  and  $A^k$ , (K.37) follows from the definition of KL divergence, (K.38) follows from the definition of  $T_j$  in (J.2). Combining (K.38) with (K.33), we conclude the proof.

## L Auxiliary Lemmas

Here, we present some auxiliary lemmas which are useful in the proof.

**Lemma L.1** (Hoeffding's inequality). [32, Theorem 2.2.6] *Let  $X_1, \dots, X_T$  be independent random variables. Assume that  $X_t \in [0, M]$  for every  $t$  with  $M > 0$ . Let  $S_T = \frac{1}{T} \sum_{t=1}^T X_t$ , then for any  $\epsilon > 0$ , we have*

$$\mathbb{P}(|S_T - \mathbb{E}[S_T]| \geq \epsilon) \leq 2 \exp \left( - \frac{2T\epsilon^2}{M^2} \right).$$

**Lemma L.2** (Self-bounding variance inequality). [22, Theorem 10] *Let  $X_1, \dots, X_T$  be independent and identically distributed random variables with finite variance. Assume that  $X_t \in [0, M]$  for every  $t$  with  $M > 0$ . Let  $S_T = \frac{1}{T} \sum_{t=1}^T X_t^2 - (\frac{1}{T} \sum_{t=1}^T X_t)^2$ , then for any  $\epsilon > 0$ , we have*

$$\mathbb{P}(|S_T - \sqrt{\text{Var}(X_1)}| \geq \epsilon) \leq 2 \exp \left( - \frac{T\epsilon^2}{2M^2} \right).$$

**Lemma L.3.** [36, Theorem 2.1] *Let  $P$  be a probability distribution over  $\mathcal{S} = \{s_1, \dots, s_S\}$ ,  $X_1, \dots, X_T$  be independent and identically distributed random variables distributed according to  $P$ . Let  $\hat{P}(s) = \frac{1}{T} \sum_{t=1}^T \mathbb{1}\{X_t = s\}$ , then for any  $\epsilon > 0$ , we have*

$$\mathbb{P}(\|P - \hat{P}\|_1 \geq \epsilon) \leq 2^S \exp \left( - \frac{T\epsilon^2}{2} \right).$$

**Lemma L.4.** [24, Lemma 7] *We define  $\mathcal{V} = \{V \in \mathbb{R}^S : \|V\|_\infty \leq V_{\max}\}$ . Let  $\mathcal{N}_{\mathcal{V}}(\epsilon)$  be a minimal  $\epsilon$ -cover of  $\mathcal{V}$  with respect to the distance metric  $d(V, V') = \|V - V'\|_\infty$  for some fixed  $\epsilon \in (0, 1)$ . Then we have*

$$\log |\mathcal{N}_{\mathcal{V}}(\epsilon)| \leq |\mathcal{S}| \cdot \log \left( \frac{3V_{\max}}{\epsilon} \right).$$



**Lemma L.5.** [7, Lemma F.4] Let  $\mathcal{F}_i$  for  $i = 1, 2, \dots$  be a filtration and  $X_1, \dots, X_n$  be a sequence of Bernoulli random variables with  $\mathbb{P}(X_i = 1 | \mathcal{F}_{i-1}) = P_i$  being  $\mathcal{F}_{i-1}$ -measurable and  $X_i$  being  $\mathcal{F}_i$ -measurable. It holds that

$$\mathbb{P}\left(\exists n : \sum_{t=1}^n X_t \leq \sum_{t=1}^n P_t/2 - W\right) \leq e^{-W}.$$

**Lemma L.6** (Bretagnolle-Huber inequality). [13, Theorem 14.2] Let  $P$  and  $Q$  be probability measures on the same measurable space  $(\Omega, \mathcal{F})$ , and let  $A \in \mathcal{F}$  be an arbitrary event. Then

$$P(A) + Q(A^c) \geq \frac{1}{2} \exp(-\text{KL}(P, Q)).$$

**Lemma L.7.** [13, Section 4.2] The KL-divergence between two Gaussian distributions with means  $\mu_1, \mu_2$  and common variance  $\sigma^2$  is

$$\text{KL}(\mathcal{N}(\mu_1, \sigma^2), \mathcal{N}(\mu_2, \sigma^2)) = \frac{(\mu_1 - \mu_2)^2}{2\sigma^2}.$$

**Lemma L.8.** [28, Theorem 3.1] Let  $P$  and  $Q$  be two probability distributions, then

$$\text{KL}(P \| Q) \leq \chi^2(P \| Q).$$