

Breaking the Curse of Multiagency in Robust Multi-Agent Reinforcement Learning

Laixi Shi^{*†} Jingchu Gai^{*‡} Eric Mazumdar[†] Yuejie Chi[§] Adam Wierman[†]
Caltech PKU Caltech CMU Caltech

October 3, 2024

Abstract

Standard multi-agent reinforcement learning (MARL) algorithms are vulnerable to sim-to-real gaps. To address this, distributionally robust Markov games (RMGs) have been proposed to enhance robustness in MARL by optimizing the worst-case performance when game dynamics shift within a prescribed uncertainty set. Solving RMGs remains under-explored, from problem formulation to the development of sample-efficient algorithms. A notorious yet open challenge is if RMGs can escape the curse of multiagency, where the sample complexity scales exponentially with the number of agents. In this work, we propose a natural class of RMGs where the uncertainty set of each agent is shaped by both the environment and other agents’ strategies in a best-response manner. We first establish the well-posedness of these RMGs by proving the existence of game-theoretic solutions such as robust Nash equilibria and coarse correlated equilibria (CCE). Assuming access to a generative model, we then introduce a sample-efficient algorithm for learning the CCE whose sample complexity scales polynomially with all relevant parameters. To the best of our knowledge, this is the first algorithm to break the curse of multiagency for RMGs.

Keywords: multi-agent reinforcement learning, robust Markov games, game theory, distribution shift

Contents

1	Introduction	2
1.1	The curse of multiagency in robust MARL	2
1.2	Contributions	3
1.3	Related work	4
2	Preliminaries	5
2.1	Standard Markov games	5
2.2	Distributionally robust Markov games	6
3	Robust Markov Games with Fictitious Uncertainty Sets	7
3.1	A novel uncertainty set definition in RMGs	8
3.2	Properties of RMGs with fictitious rectangular uncertainty set	9
4	Sample-Efficient Learning: Algorithm and Theory	10
4.1	Problem setting and goal	10
4.2	Algorithm design	10
4.3	Theoretical guarantees	12

^{*}The first two authors contributed equally.

[†]Department of Computing Mathematical Sciences, California Institute of Technology, CA 91125, USA.

[‡]School of Mathematical Sciences, Peking University, Beijing, 100871, China.

[§]Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

5 Conclusion	14
A Preliminaries	19
A.1 Additional matrix and vector notation	19
A.2 Preliminary facts about FTRL	20
B Proof for Section 3	21
B.1 Proof of Theorem 1	21
B.2 Proof of auxiliary facts	24
C Proof of Theorem 2	28
C.1 Proof pipeline	28
C.2 Controlling B: adversarial online learning	29
C.3 Controlling terms A and C	34
C.4 Proof of auxiliary lemmas	44

1 Introduction

A flurry of problems naturally involve decision-making among multiple players with strategic objectives. Multi-agent reinforcement learning (MARL) serves as a powerful framework to address these challenges, demonstrating potential in various applications such as social dilemmas (Baker, 2020; Leibo et al., 2017; Zhang et al., 2024), autonomous driving (Lillicrap et al., 2015), robotics (Kober et al., 2013; Rusu et al., 2017), and games (Mnih et al., 2015; Vinyals et al., 2019). Despite the recent success of standard MARL, its transition from prototypes to reliable production is hindered by robustness concerns due to the complexity and variability of both the real-world environment and human behaviors. Specifically, environmental uncertainty can arise from sim-to-real gaps (Tobin et al., 2017), unexpected disturbance (Pinto et al., 2017), system noise, and adversarial attacks (Mahmood et al., 2018); agents’ behaviors are subject to unknown bounded rationality and variability (Tversky and Kahneman, 1974). The solution learned at training time can fail catastrophically when faced with a slightly shifted MARL problem during testing, resulting in a significant drop in overall outcomes and each agent’s individual payoff (Balaji et al., 2019; Shi et al., 2024; Slumbers et al., 2023; Yeh et al., 2021; Zeng et al., 2022; Zhang et al., 2020a).

To address robustness challenges, a promising framework is (distributionally) robust Markov games (RMGs) (Littman, 1994; Shapley, 1953). It is a robust counterpart to the common playground of standard MARL problems — Markov games (MGs) (Kardeş et al., 2011; Zhang et al., 2020c). In standard MGs, agents consider (competitive) personal objectives and simultaneously interact with each other within a shared unknown environment. The goal is to learn some solution concepts called equilibria, which are joint strategies/policies of agents that all of them stick with rationally with other agents fixed; for instance, Nash equilibria (NE) (Nash, 1951; Shapley, 1953), correlated equilibria (CE), and coarse correlated equilibria (CCE) (Aumann, 1987; Moulin and Vial, 1978). To promote robustness, RMGs differ from standard MGs by defining each agent’s payoff (objective) as its worst-case performance when the dynamics of the game shift within a prescribed uncertainty set centered around a nominal environment.

1.1 The curse of multiagency in robust MARL

Sample efficiency is a crucial metric for MARL due to the limited availability of data relative to the high dimensionality of the problem. In MARL, agents strive to learn a rationally optimal solution (equilibrium) through interactions with an unknown environment (Achiam et al., 2023; Silver et al., 2016; Vinyals et al., 2019). In contemporary applications, the environment is often extremely large-scale, while data acquisition can be prohibitively limited by high costs and stakes. As such, a notable challenge in terms of scalability for sample efficiency in MARL is known as *the curse of multiagency* — the sample complexity requirement scales exponentially with the number of agents (induced by the exponentially growing size of the joint action space). This issue has been recognized and studied in extensive MARL problems, but remains open for robust MARL. We concentrate on learning finite-horizon multi-player general-sum Markov games with a

generative model (Kearns and Singh, 1999), where the number of agents is n , the episode length is H , the size of the state space is S , and the size of the i -th agent’s action space is A_i , for $1 \leq i \leq n$.

- *Breaking the curse of multiagency in standard MARL.* A line of pioneering work (Bai and Jin, 2020; Jin et al., 2021; Li et al., 2023; Song et al., 2021) has recently introduced a new suite of algorithms using adaptive sampling that provably break the curse of multiagency in standard MGs. In particular, to find an ε -approximate CCE, Li et al. (2023) requires a sample complexity no more than

$$\tilde{O}\left(\frac{H^4 S \sum_{i=1}^n A_i}{\varepsilon^2}\right) \quad (1)$$

up to logarithmic factors, which depends only on the sum of individual actions, rather than the number of joint actions.

- *The persistent curse of multiagency in robust MARL.* The development of provable sample-efficient algorithms for RMGs is largely underexplored, with only a few recent studies (Blanchet et al., 2023; Kardeş et al., 2011; Ma et al., 2023; Shi et al., 2024; Zhang et al., 2020c). Focusing on a class of RMGs with uncertainty sets satisfying the (s, \mathbf{a}) -*rectangularity condition*, existing works all suffer from the curse of multiagency, significantly limiting their scalability. For example, using the total variation (TV) distance as the divergence function, the state-of-the-art (Shi et al., 2024), using non-adaptive sampling, finds an ε -approximate robust CCE with a sample complexity no more than

$$\tilde{O}\left(\frac{H^3 S \prod_{i=1}^n A_i}{\varepsilon^2} \min\left\{H, \frac{1}{\min_{1 \leq i \leq n} \sigma_i}\right\}\right) \quad (2)$$

up to logarithmic factors, where $\sigma_i \in [0, 1]$ is the uncertainty level for the i -th agent. As a result, the sample size requirement becomes prohibitive when the number of agents is large.

Consequently, there is a significant desire to explore paths that could break through the curse of multiagency in RMGs, which is much more involved than its standard counterpart due to complicated non-linearity introduced by planning for worst-case performances. Nevertheless, the family of RMGs is a much richer class of problems because of the flexibility in choosing the uncertainty sets to capture different robust design considerations. While convenient, the (s, \mathbf{a}) -*rectangularity condition* prevalent in current approaches can be overly restricted in practice, as each agent’s uncertainty set is assumed to be independent of other agents’ strategies and can be decoupled into independent subsets for each state-joint action pair (s, \mathbf{a}) , suggesting it might be challenging to break the curse of multiagency in the existing framework. Given these limitations, we are motivated to develop new classes of RMGs that can provide robust solutions applicable to more realistic MARL problems with sample-efficient algorithms. This raises an open question:

Can we design RMGs with practically-meaningful uncertainty sets that come with sample complexity guarantees breaking the curse of multiagency?

1.2 Contributions

We propose a new class of RMGs with a *fictitious* uncertainty set that explicitly captures uncertainties in the environment in view of other agents’ strategies, making it suitable for complex real-world scenarios. We begin by verifying the game-theoretic properties of the proposed class of RMGs to ensure the existence of robust variants of well-known standard equilibria notions, robust NE and robust CCE. Next, due to the general intractability of learning NE, we focus on designing algorithms that can provably overcome the curse of multiagency in learning an approximate robust CCE, referring to a joint policy where no agent can improve their benefit by more than ε through rational deviations.. Specifically, for sampling mechanisms to explore the unknown environment, we assume access to a generative model that can only draw samples from the nominal environment (Shi et al., 2024). The main contributions are summarized as follows.

- We introduce a new class of robust Markov games using fictitious uncertainty sets with *policy-induced* (s, a_i) -*rectangularity condition* (see Section 2.2 for details), which is a natural adaptation from robust

Algorithm	Uncertainty set	Equilibria	Sample complexity
P ² MPO (Blanchet et al., 2024)	(s, \mathbf{a}) -rectangularity	robust NE	$S^4 (\prod_{i=1}^n A_i)^3 H^4 / \varepsilon^2$
DR-NVI (Shi et al., 2024)	(s, \mathbf{a}) -rectangularity	robust NE/CE/CCE	$\frac{SH^3 \prod_{i=1}^n A_i}{\varepsilon^2} \min \left\{ H, \frac{1}{\min_{1 \leq i \leq n} \sigma_i} \right\}$
Robust-Q-FTRL (this work)	fictitious (s, a_i) -rectangularity	robust CCE	$\frac{SH^6 \sum_{1 \leq i \leq n} A_i}{\varepsilon^4} \min \left\{ H, \frac{1}{\min_{1 \leq i \leq n} \sigma_i} \right\}$

Table 1: Comparisons between our results and prior art for finding an ε -approximate equilibrium in finite-horizon multi-agent general-sum robust Markov games. We omit all logarithmic factors in the sample complexities here. Our result is the only computationally tractable algorithm that provably breaks the curse of multiagency.

single-agent RL to robust MARL. The uncertainty set for each agent i can be decomposed into independent subsets over each state and its own action tuple (s, a_i) , where each subset is a “ball” around the expected nominal transition determined by other agents’ policies and the nominal transition kernel, a distance function ρ , and the radius/uncertainty level σ_i . We verify several essential facts of this class of RMGs: the existence of the desired equilibrium — robust NE and robust CCE for this new class of RMGs using game-theoretical tools such as fixed-point theorem; the existence of best-response policies and robust Bellman equations.

- We consider the total variation (TV) distance as the distance metric ρ for uncertainty sets due to its popularity in both theory (Blanchet et al., 2023; Panaganti and Kalathil, 2022; Shi et al., 2023, 2024) and practice (Lee et al., 2021; Pan et al., 2023; Szita et al., 2003). We propose Robust-Q-FTRL that can provably find ε -approximate robust CCE with high probability, as long as the sample size exceeds

$$\tilde{O} \left(\frac{SH^6 \sum_{i=1}^n A_i}{\varepsilon^4} \min \left\{ H, \frac{1}{\min_{1 \leq i \leq n} \sigma_i} \right\} \right) \quad (3)$$

up to logarithmic factors, where $\sigma_i \in (0, 1]$ is the uncertainty level for the i -th agent. To the best of our knowledge, this is the first algorithm to break the curse of multiagency in sample complexity of RMGs. It provably finds an ε -approximate robust CCE using a sample size that is polynomial to all salient parameters. Table 1 provides a detailed comparison to prior works in robust MARL, where our results significantly improve upon prior art (2) (Shi et al., 2024) by reducing the exponential dependency on the size of each agent’s action space to a linear dependency. To achieve this, we utilize adaptive sampling and online adversarial learning tools, coupled by a tailored design and analysis for robust MARL due to the nonlinearity of the robust value function, which contrasts with the linear payoff functions in standard MARL with respect to the transition kernel.

Notation. In this paper, we denote $[T] := \{1, 2, \dots, T\}$ for any positive integer $T > 0$. We define $\Delta(\mathcal{S})$ as the simplex over a set \mathcal{S} . For any policy π and function $Q(\cdot)$ defined over a domain \mathcal{B} , the variance of Q under π is given by $\text{Var}_\pi(Q) := \sum_{a \in \mathcal{B}} \pi(a) [Q(a) - \mathbb{E}_\pi[Q]]^2$. We define $x = [x(s, \mathbf{a})]_{(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}} \in \mathbb{R}^{S^A}$ as any vector that represents values for each state-action pair, and $x = [x(s, a_i)]_{(s, a_i) \in \mathcal{S} \times \mathcal{A}_i} \in \mathbb{R}^{S^{A_i}}$ as any vector representing agent-wise state-action values. Similarly, we denote $x = [x(s)]_{s \in \mathcal{S}}$ as any vector representing values for each state. For $\mathcal{X} := (\mathcal{S}, \{A_i\}_{i \in [n]}, H, \{\sigma_i\}_{i \in [n]}, \frac{1}{\varepsilon}, \frac{1}{\delta})$, let $f(\mathcal{X}) = O(g(\mathcal{X}))$ denote that there exists a universal constant $C_1 > 0$ such that $f \leq C_1 g$. Furthermore, the notation $\tilde{O}(\cdot)$ is defined similarly to $O(\cdot)$ but hides logarithmic factors.

1.3 Related work

Breaking curse of multiagency for standard Markov games. Breaking the curse of multiagency is a major and prevalent challenge in sequential games. In standard multi-agent general-sum MGs, it has been

shown that learning a Nash equilibrium requires an exponential sample complexity (Bai and Jin, 2020; Rubinstein, 2017; Song et al., 2021). However, for other types of equilibria, such as CE and CCE, many works have successfully broken the curse of multiagency. Specifically, for finite-horizon general-sum MGs in the tabular setting with finite state and action spaces, Jin et al. (2021) developed the V-learning algorithm for learning CE and CCE with the sample complexity of $\tilde{O}(H^6 S (\max_{i \in [n]} A_i)^2 / \epsilon^2)$ and $\tilde{O}(H^6 S \max_{i \in [n]} A_i / \epsilon^2)$, respectively; Daskalakis et al. (2023) achieved a sample complexity of $\tilde{O}(H^{11} S^3 \max_{i \in [n]} A_i / \epsilon^3)$ for learning a CCE. Beyond tabular settings, Wang et al. (2023) and Cui et al. (2023) extended these results to linear function approximation, achieving sample complexities of $\tilde{O}(d^4 H^6 (\max_{i \in [n]} A_i^5) / \epsilon^2)$ and $\tilde{O}(H^{10} d^4 \log(\max_{i \in [n]} A_i) / \epsilon^4)$, respectively, where d is the dimension of the linear features. For Markov potential games, a subclass of MGs, Song et al. (2021) provided a centralized algorithm that learns a NE with a sample complexity of $\tilde{O}(H^4 S^2 \max_{i \in [n]} A_i / \epsilon^3)$.

Finite-sample analysis for distributionally robust Markov games. Robust Markov games under environmental uncertainty are largely underexplored, with only a few provable algorithms (Blanchet et al., 2023; Kardeş et al., 2011; Ma et al., 2023; Shi et al., 2024; Zhang et al., 2020a). Existing sample complexity analyses all suffer from the daunting curse of multiagency issues, or impose an extremely restricted uncertainty level that can fail to deliver the desired robustness (Blanchet et al., 2024; Ma et al., 2023; Shi et al., 2024). Specifically, they all consider a class of RMGs with the (s, \mathbf{a}) -rectangularity condition, where the uncertainty sets for each agent can be decomposed into independent sets over each (s, \mathbf{a}) pair. Shi et al. (2024) considered the generative model with an uncertainty set measured by the TV distance, Blanchet et al. (2023) treated a different sampling mechanism with offline data for both the TV distance and KL divergence. In addition, Ma et al. (2023) required the uncertainty level be much smaller than the accuracy-level and an instance-dependent parameter (i.e., $\sigma_i \leq \max\{\frac{\epsilon}{SH^2}, \frac{p_{\min}}{H}\}$ for all $i \in [n]$). This can thus fail to maintain the desired robustness, especially when the accuracy requirement is high (i.e., $\epsilon \rightarrow 0$) or the RMG has small minimal positive transition probabilities (i.e., $p_{\min} \rightarrow 0$).

Robust MARL. Standard MARL algorithms may overfit the training environment and could fail dramatically due to the perturbations and variability of both agents’ behaviors and the shared environment, leading to performance drop and large deviation from the equilibrium. To address this, this work considers a robust variant of MARL adopting the distributionally robust optimization (DRO) framework that has primarily been investigated in supervised learning (Bertsimas et al., 2018; Blanchet and Murthy, 2019; Duchi and Namkoong, 2018; Gao, 2020; Rahimian and Mehrotra, 2019) and has attracted a lot of attention in promoting robustness in single-agent RL (Badrinath and Kalathil, 2021; Iyengar, 2005; Nilim and El Ghaoui, 2005; Shi and Chi, 2024; Shi et al., 2023; Wang et al., 2024; Zhou et al., 2021). Beyond the RMG framework considered in this work, recent research has advanced the robustness of MARL algorithms from various perspectives, including resilience to uncertainties or attacks on states (Han et al., 2022; Zhou and Liu, 2023), the type of agents (Zhang et al., 2021), other agents’ policies (Kannan et al., 2023; Li et al., 2019), offline data poisoning (McMahan et al., 2024; Wu et al., 2024), and nonstationary environment (Szita et al., 2003). A recent review can be found in Vial et al. (2022).

2 Preliminaries

In this section, we begin with some background on multi-agent general-sum standard Markov games (MGs) in finite-horizon settings, followed by a general framework of a robust variant of standard MGs — distributionally robust Markov games.

2.1 Standard Markov games

A finite-horizon *multi-agent general-sum Markov game* (MG) is characterized by the tuple

$$\mathcal{MG} = \{\mathcal{S}, \{\mathcal{A}_i\}_{1 \leq i \leq n}, P, r, H\}.$$

This setup features n agents each striving to maximize their individual long-term cumulative rewards within a shared environment. At each time step, all agents observe the same state over the state space $\mathcal{S} = \{1, \dots, S\}$

within the shared environment. For each agent i ($i \in [n]$), $\mathcal{A}_i = \{1, \dots, A_i\}$ denotes its action space containing A_i possible actions. The joint action space for all agents (resp. the subset excluding the i -th agent) is defined as $\mathcal{A} := \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ (resp. $\mathcal{A}_{-i} := \prod_{j \neq i} \mathcal{A}_j$ for any $i \in [n]$). We use the notation $\mathbf{a} \in \mathcal{A}$ (resp. $\mathbf{a}_{-i} \in \mathcal{A}_{-i}$) to denote a joint action profile involving all agents (resp. all except the i -th agent). In addition, the probability transition kernel $P = \{P_h\}_{1 \leq h \leq H}$, with each $P_h : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$, describes the dynamics of the game: $P_h(s' | s, \mathbf{a})$ is the probability of transitioning from state $s \in \mathcal{S}$ to state $s' \in \mathcal{S}$ at time step h when agents choose the joint action profile $\mathbf{a} \in \mathcal{A}$. The reward function of the game is $r = \{r_{i,h}\}_{1 \leq i \leq n, 1 \leq h \leq H}$, with each $r_{i,h} : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$ normalized to the unit interval. For any $(i, h, s, \mathbf{a}) \in [n] \times [H] \times \mathcal{S} \times \mathcal{A}$, $r_{i,h}(s, \mathbf{a})$ represents the immediate reward received by the i -th agent in state s when the joint action profile \mathbf{a} is taken. Last but not least, $H > 0$ represents the horizon length.

Markov policies and value functions. In this work, we concentrate on Markov policies that the action selection rule depends only on the current state s , independent from previous trajectory. Namely, the i -th ($i \in [n]$) agent chooses actions according to $\pi_i = \{\pi_{i,h} : \mathcal{S} \mapsto \Delta(\mathcal{A}_i)\}_{1 \leq h \leq H}$. Here, $\pi_{i,h}(a | s)$ represents the probability of selecting action $a \in \mathcal{A}_i$ in state s at time step h . As such, the joint Markov policy of all agents can be denoted as $\pi = (\pi_1, \dots, \pi_n) : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A})$, i.e., given any $s \in \mathcal{S}$ and $h \in [H]$, the joint action profile $\mathbf{a} \in \mathcal{A}$ of all agents is chosen following the distribution $\pi_h(\cdot | s) = (\pi_{1,h}, \pi_{2,h}, \dots, \pi_{n,h})(\cdot | s) \in \Delta(\mathcal{A})$.

To continue, for any given joint policy π and transition kernel P of a \mathcal{MG} , the i -th agent's long-term cumulative reward can be characterized by the value function $V_{i,h}^{\pi,P} : \mathcal{S} \mapsto \mathbb{R}$ (resp. Q-function $Q_{i,h}^{\pi,P} : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$) as below: for all $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$,

$$V_{i,h}^{\pi,P}(s) := \mathbb{E}_{\pi,P} \left[\sum_{t=h}^H r_{i,t}(s_t, \mathbf{a}_t) \mid s_h = s \right], \quad Q_{i,h}^{\pi,P}(s, \mathbf{a}) := \mathbb{E}_{\pi,P} \left[\sum_{t=h}^H r_{i,t}(s_t, \mathbf{a}_t) \mid s_h = s, \mathbf{a}_h = \mathbf{a} \right]. \quad (4)$$

In this context, the expectation is calculated over the trajectory $\{(s_t, \mathbf{a}_t)\}_{h \leq t \leq H}$ produced by following the joint policy π under the transition kernel P .

2.2 Distributionally robust Markov games

A general distributionally robust Markov game (RMG) is represented by the tuple

$$\mathcal{RMG} = \{\mathcal{S}, \{\mathcal{A}_i\}_{1 \leq i \leq n}, \{\mathcal{U}_\rho^{\sigma_i}(P^0)\}_{1 \leq i \leq n}, r, H\}.$$

Here, $\mathcal{S}, \{\mathcal{A}_i\}_{1 \leq i \leq n}, r, H$ are defined in the same manner as those in standard MGs (see Section 2.1). RMGs differ from standard MGs: for each agent i ($1 \leq i \leq n$), the transition kernel is not fixed but can vary within its own prescribed uncertainty set $\mathcal{U}_\rho^{\sigma_i}(P^0)$ centered around some *nominal* kernel $P^0 : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$ that represents a reference (such as the training environment). The shape and the size of the uncertainty set $\{\mathcal{U}_\rho^{\sigma_i}(P^0)\}_{i \in [n]}$ are further specified by a divergence function ρ and the uncertainty levels $\{\sigma_i\}_{i \in [n]}$, serving as the ‘‘distance’’ metric and the radius respectively.

Various choices of the divergence function have been considered in the literature of robust RL, including but not limited to f -divergence (such as total variation, χ^2 divergence, and Kullback-Leibler (KL) divergence) (Lu et al., 2024; Shi and Chi, 2024; Wang et al., 2024; Yang et al., 2022; Zhou et al., 2021) and Wasserstein distance (Xu et al., 2023). Adopting uncertainty sets with different structures leads to distinct RMGs, as they address distinct types of uncertainty and game-theoretical solutions. This paper focuses on variability in environmental dynamics (transition kernels), though uncertainty in agents' reward functions could also be considered similarly but is omitted for brevity.

Robust value functions and best-response policies. For any RMG, each agent seeks to maximize its worst-case performance in the presence of other agents' behaviors despite perturbations in the environment dynamics, as long as the kernel transitions remain within its prescribed uncertainty set $\mathcal{U}_\rho^{\sigma_i}(P^0)$. Mathematically, given any joint policy $\pi : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A})$, the worst-case performance of any agent i is characterized by the *robust value function* $V_{i,h}^{\pi,\sigma_i}$ and the *robust Q-function* $Q_{i,h}^{\pi,\sigma_i}$: for all $(i, h, s, a_i) \in [n] \times [H] \times \mathcal{S} \times \mathcal{A}_i$,

$$V_{i,h}^{\pi,\sigma_i}(s) := \inf_{P \in \mathcal{U}_\rho^{\sigma_i}(P^0)} V_{i,h}^{\pi,P}(s) \quad \text{and} \quad Q_{i,h}^{\pi,\sigma_i}(s, a_i) := \inf_{P \in \mathcal{U}_\rho^{\sigma_i}(P^0)} Q_{i,h}^{\pi,P}(s, a_i). \quad (5)$$

Note that different from (4), here the Q-function for any i -th agent is defined only over its own action $a_i \in \mathcal{A}_i$ rather than the joint action $\mathbf{a} \in \mathcal{A}$.

To continue, we denote π_{-i} as the policy for all agents except for the i -th agent. By optimizing the i -th agent's policy $\pi'_i : \mathcal{S} \times [H] \rightarrow \Delta(\mathcal{A}_i)$ (independent from π_{-i}), we define the maximum of the robust value function as

$$V_{i,h}^{*,\pi_{-i},\sigma_i}(s) := \max_{\pi'_i: \mathcal{S} \times [H] \rightarrow \Delta(\mathcal{A}_i)} V_{i,h}^{\pi'_i \times \pi_{-i}, \sigma_i}(s) = \max_{\pi'_i: \mathcal{S} \times [H] \rightarrow \Delta(\mathcal{A}_i)} \inf_{P \in \mathcal{U}_\rho^{\sigma_i}(P^0)} V_{i,h}^{\pi'_i \times \pi_{-i}, P}(s) \quad (6)$$

for all $(i, h, s) \in [n] \times [H] \times \mathcal{S}$. The policy that achieves the maximum of the robust value function for all $(i, h, s) \in [n] \times [H] \times \mathcal{S}$ is called a *robust best-response policy*.

Solution concepts for robust Markov games. In view of the conflicting objectives between agents, establishing equilibrium becomes the goal of solving RMGs. As such, we introduce two kinds of solution concepts — robust NE and robust CCE — robust variants of standard NE and CCE (usually considered in standard MGs) specified to the form of RMGs.

- *Robust NE.* A product policy $\pi = \pi_1 \times \pi_2 \times \dots \times \pi_n : \mathcal{S} \times [H] \mapsto \prod_{i=1}^n \Delta(\mathcal{A}_i)$ is said to be a *robust NE* if

$$V_{i,1}^{\pi, \sigma_i}(s) = V_{i,1}^{*,\pi_{-i}, \sigma_i}(s), \quad \forall (s, i) \in \mathcal{S} \times [n]. \quad (7)$$

Given the strategies of the other agents π_{-i} , when each agent wants to optimize its worst-case performance when the environment and other agents' policy stay within its own uncertainty set $\mathcal{U}_\rho^{\sigma_i}(P^0)$, robust NE means that no player can benefit by unilaterally diverging from its present strategy.

- *Robust CCE.* A distribution over the joint product policy $\xi := \{\xi_h\}_{h \in [H]} : [H] \mapsto \Delta(\mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i))$ is said to be a *robust CCE* if it holds that

$$\mathbb{E}_{\pi \sim \xi} [V_{i,1}^{\pi, \sigma_i}(s)] \geq \mathbb{E}_{\pi \sim \xi} [V_{i,1}^{*,\pi_{-i}, \sigma_i}(s)], \quad \forall (i, s) \in [n] \times \mathcal{S}. \quad (8)$$

Considering all agents follow the policy drawn from the distribution ξ , i.e., $\pi_h \sim \xi_h$ for all $h \in [H]$, when the distribution of all agents but the i -th agent's policy is fixed as the marginal distribution of ξ , robust CCE indicates that no agent can benefit from deviating from its current policy.

Note that, for standard MGs, CCE is defined as a possibly correlated joint policy $\pi^{\text{CCE}} : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A})$ (Aumann, 1987; Moulin and Vial, 1978) if it holds that

$$V_{i,1}^{\pi^{\text{CCE}}, P}(s) \geq \max_{\pi'_i: \mathcal{S} \times [H] \rightarrow \Delta(\mathcal{A}_i)} V_{i,1}^{\pi'_i \times \pi_{-i}^{\text{CCE}}, P}(s), \quad \forall (s, i) \in \mathcal{S} \times [n]. \quad (9)$$

This correlated policy π^{CCE} can also be viewed as a distribution ξ over the product policy space since each joint action \mathbf{a} can be seen as a deterministic product policy. Careful readers may note that the definition (9) of CCE in standard MGs is in a different form from the one (8) in RMGs, as the latter does not include the expectation operator $\mathbb{E}_{\pi \sim \xi}[\cdot]$ with respect to the policy distribution (ξ) over the value function. We emphasize that the definition with the expectation operator outside of the value (or cost) function with respect to a distribution of product pure strategies in (8) is a natural formulation originating from game theory (Moulin et al., 2014; Moulin and Vial, 1978). In standard MARL and previous robust MARL studies, the definition in (9) is typically used because (9) and (8) are identical in those situations, as the expectation operator and the corresponding value functions are linear with respect to the joint policy, allowing them to be interchanged (Li et al., 2023; Shi et al., 2024).

3 Robust Markov Games with Fictitious Uncertainty Sets

Given the definition of general RMGs, a natural question arises: what kinds of uncertainty sets should we consider to achieve the desired robustness in our solutions? To address this, we focus on a specific class of RMGs characterized by a type of natural yet powerful uncertainty sets.

3.1 A novel uncertainty set definition in RMGs

We propose a new class of uncertainty sets, named *fictitious* uncertainty sets, which count in the uncertainty induced by both the environment and agents' behaviors in a correlated manner. Before introducing the uncertainty sets, we provide some auxiliary notations as below. We denote a vector of any transition kernel $P : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$ or $P^0 : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$ respectively as

$$\forall (s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A} : P_{h,s,\mathbf{a}} := P_h(\cdot | s, \mathbf{a}) \in \mathbb{R}^{1 \times S}, \quad P_{h,s,\mathbf{a}}^0 := P_h^0(\cdot | s, \mathbf{a}) \in \mathbb{R}^{1 \times S}. \quad (10)$$

For any (possibly correlated) joint Markov policy (defined in section 2.1) $\pi : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A})$, we define the expected nominal transition kernel conditioned on the situation that the i -th agent chooses some action $a_i \in \mathcal{A}_i$ and other agents play according to the conditional policy (i.e., $\mathbf{a}_{-i} \sim \pi_h(\cdot | s, a_i)$) given $s \in \mathcal{S}$ and a_i as below: for each time step $h \in [H]$,

$$\forall (h, s, a_i) \in [H] \times \mathcal{S} \times \mathcal{A}_i : P_{h,s,a_i}^{\pi_{-i}} = \mathbb{E}_{\mathbf{a}_{-i} \sim \pi_h(\cdot | s, a_i)} [P_{h,s,\mathbf{a}}^0] = \sum_{\mathbf{a}_{-i} \in \mathcal{A}_{-i}} \frac{\pi_h(a_i, \mathbf{a}_{-i} | s)}{\pi_{i,h}(a_i | s)} [P_{h,s,\mathbf{a}}^0]. \quad (11)$$

Armed with the above definitions, now we are in a position to define the *fictitious* uncertainty sets, denoted as $\{\mathcal{U}_\rho^{\sigma_i}(P^0, \cdot)\}_{i \in [n]}$, which satisfy a *policy-induced* (s, a_i) -*rectangularity condition*.

Definition 1. For any joint policy $\pi : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A})$, divergence function $\rho : \Delta(\mathcal{S}) \times \Delta(\mathcal{S}) \mapsto \mathbb{R}^+$ and accessible uncertainty levels $\sigma_i \geq 0$ for all $i \in [n]$, the *fictitious uncertainty sets* $\{\mathcal{U}_\rho^{\sigma_i}(P^0, \pi)\}_{i \in [n]}$ satisfy the *policy-induced* (s, a_i) -*rectangularity condition*:

$$\begin{aligned} \forall i \in [n] : \mathcal{U}_\rho^{\sigma_i}(P^0, \pi) &:= \otimes \mathcal{U}_\rho^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}}), \\ \text{s.t. } \forall (h, s, a_i) \in [H] \times \mathcal{S} \times \mathcal{A}_i : \mathcal{U}_\rho^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}}) &:= \left\{ P \in \Delta(\mathcal{S}) : \rho(P, P_{h,s,a_i}^{\pi_{-i}}) \leq \sigma_i \right\}, \end{aligned} \quad (12)$$

where \otimes represents the Cartesian product.

In words, conditioned on a fixed joint policy π , the uncertainty set $\mathcal{U}_\rho^{\sigma_i}(P^0, \pi)$ for each i -th agent can be decomposed into a Cartesian product of subsets over each state and agent-action pair (s, a_i) . Each uncertainty subset $\mathcal{U}_\rho^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})$ over (s, a_i) is defined as a "ball" around a reference — the expected nominal transition kernel $P_{h,s,a_i}^{\pi_{-i}}$ conditioned on both transition kernel and agents' behavior π .

Further discussions of fictitious uncertainty sets. It is in order to remark on the proposed type of uncertainty sets, in comparison with prior works.

- *A natural adaptation from single-agent robust RL.* When agents follow some joint policy $\pi : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A})$, fixing other agents' policy π_{-i} , from the perspective of each individual agent i , RMGs with our policy-induced (s, a_i) -rectangularity condition will degrade to a single-agent robust RL problem with the widely used (s, a_i) -rectangularity condition in the single-agent literature (Iyengar, 2005; Zhou et al., 2021). Namely, from any agent i 's viewpoint, in a RMG, it has an "overall environment" player that can not only manipulate the environmental dynamics but also other players' policy π_{-i} .
- *Allowing uncertainty from both the environment and agents' behaviors in a correlated manner.* One essential feature of our proposed uncertainty set is that it is shaped by both the environment and agents' strategies in a (possibly) correlated manner. Specifically, for any agent i and a given policy π , any uncertainty subset $\mathcal{U}_\rho^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})$ (over any (h, s, a_i)) is constructed as a neighborhood around a nominal center $P_{h,s,a_i}^{\pi_{-i}}$ (see (11)) that depends on both the nominal environment P^0 and other agents' conditional strategies $\pi_h(\cdot | s, a_i)$.
- *Comparisons to prior works.* Prior works on provable sample-efficient algorithms have focused on a different type of uncertainty sets with (s, \mathbf{a}) -rectangularity condition (Blanchet et al., 2023; Ma et al., 2023; Shi et al., 2024). This class of uncertainty sets decouples the uncertainty into independent subsets for each state-joint action pair (s, \mathbf{a}) , accounting for the uncertainty of the environment and agents' strategies independently. In comparison, the proposed uncertainty set lifts this independence assumption across subsets over different $(s, a_i, \mathbf{a}_{-i})$ for any $\mathbf{a}_{-i} \in \mathcal{A}_{-i}$, enabling the environment and agents' strategies to shape the uncertainty set in a correlated manner.

3.2 Properties of RMGs with fictitious rectangular uncertainty set

Throughout the paper, we focus on the class of RMGs with the above proposed fictitious uncertainty sets, represented as

$$\mathcal{RMG}_\pi = \{\mathcal{S}, \{\mathcal{A}_i\}_{1 \leq i \leq n}, \{\mathcal{U}_\rho^{\sigma_i}(P^0, \cdot)\}_{1 \leq i \leq n}, r, H\}$$

and abbreviated as fictitious RMGs in the remaining of the paper. In this section, we present key facts about fictitious RMGs related to best-response policies, equilibria, and the corresponding one-step lookahead robust Bellman equations. The proofs are postponed to Appendix B.

First, we introduce the following lemma, which verifies the existence of a robust best-response policy that achieves the maximum robust value function (cf. (6)).

Lemma 1. *For any $i \in [n]$, given $\pi_{-i} : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)$, there exists at least one policy $\tilde{\pi}_i : \mathcal{S} \times [H] \rightarrow \Delta(\mathcal{A}_i)$ for the i -th agent that can simultaneously attain $V_{i,h}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i}(s) = V_{i,h}^{*, \pi_{-i}, \sigma_i}(s)$ for all $s \in \mathcal{S}$ and $h \in [H]$. We refer this policy as the robust best-response policy.*

Existence of robust NE and robust CCE. fictitious RMGs can be viewed as hierarchical games with $n + nS \sum_{i=1}^n A_i$ agents. This includes the original n agents and n additional sets of $S \sum_{i=1}^n A_i$ independent adversaries, each determining the worst-case transitions for one agent over a state plus agent-wise-action pair. Considering the solution concepts — robust NE and robust CCE — introduced in Section 2.2, the following theorem verifies the existence of them for any fictitious RMGs using Kakutani’s fixed-point theorem (Kakutani, 1941), focusing on robust NE firstly.

Theorem 1 (Existence of robust NE). *For any $\mathcal{RMG}_\pi = \{\mathcal{S}, \{\mathcal{A}_i\}_{1 \leq i \leq n}, \{\mathcal{U}_\rho^{\sigma_i}(P^0, \cdot)\}_{1 \leq i \leq n}, r, H\}$ with an uncertainty set defined in Definition 1, there exists at least one robust NE.*

Analogous to standard Markov games, since $\{\text{robust NE}\} \subseteq \{\text{robust CCE}\}$, Theorem 1 indicates the existence of robust CCEs directly.

Robust Bellman equations. Fortunately, the class of fictitious RMGs feature a robust counterpart of the Bellman equation — *robust Bellman equation*. Specifically, for any joint policy $\pi : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A})$, the robust value function can be expressed as

$$V_{i,h}^{\pi, \sigma_i}(s) = \inf_{\mathcal{U}_\rho^{\sigma_i}(P^0, \pi)} \mathbb{E} \left[\sum_{t=h}^H r_i(s_t, a_t) \mid s_h = s \right] = \mathbb{E}_{\mathbf{a} \sim \pi_h(s)} [r_{i,h}(s, \mathbf{a})] + \mathbb{E}_{a_i \sim \pi_{i,h}(s)} \left[\inf_{\mathcal{U}_\rho^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} PV_{i,h+1}^{\pi, \sigma_i} \right]. \quad (13)$$

It can be verified directly by definition. The robust Bellman equation described above is intrinsically linked to the *policy-induced (s, a_i) -rectangularity* condition (cf. (12)) of the uncertainty set. This condition leads to a well-posed and computationally-tractable class of RMGs by allowing the decomposition from an overall uncertainty set to independent subsets across different agents, time steps, and each state-action pair (s, a_i) .

Note that the specified robust Bellman equation is different for a joint correlated policy and a joint product policy, induced by different expected nominal transition kernels. In particular, for any joint product policy $\pi : \mathcal{S} \times [H] \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i)$, the expected nominal transition kernel conditioned on the i -th agent’s action $a_i \in \mathcal{A}_i$, current state $s \in \mathcal{S}$, and the policy π can be expressed by

$$P_{h,s,a_i}^{\pi_{-i}} = \mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i}(\cdot \mid s, a_i)} [P_{h,s,\mathbf{a}}^0] = \mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,h}(\cdot \mid s)} [P_{h,s,(a_i, \mathbf{a}_{-i})}^0] \quad (14)$$

for any $(i, h, s, a_i) \in [n] \times [H] \times \mathcal{S} \times \mathcal{A}_i$, where the last equality holds since the policy π is a product policy, and the distribution of \mathbf{a}_{-i} is independent of a_i . It is observed that the expected nominal transition kernel $P_{h,s,a_i}^{\pi_{-i}}$ for a product policy π is independent of the i -th agent’s policy given (s, a_i) . This differs from (11) for a possibly correlated policy, where (11) can generally depend on the i -th agent’s policy.

The robust Bellman equation described above is intrinsically linked to the *policy-induced (s, a_i) -rectangularity* condition (cf. (12)) of the uncertainty set. This condition leads to a well-posed and computationally-tractable class of RMGs by allowing the decomposition from an overall uncertainty set to independent subsets across different agents, time steps, and each state-action pair (s, a_i) .

4 Sample-Efficient Learning: Algorithm and Theory

In this section, we focus on designing sample-efficient algorithms for solving fictitious RMGs when agents need to collect data by interacting with the unknown shared environment in order to learn the equilibria. To proceed, we shall first specify the data collection mechanism and the divergence function for the uncertainty set. Then we propose a sample-efficient algorithm Robust-Q-FTRL that leverages a carefully-designed adaptive sampling strategy to break the curse of multiagency.

4.1 Problem setting and goal

Recall that the uncertainty sets are constructed by specifying a divergence function ρ and the uncertainty level to control its shape and size. In this work, we focus on using the TV distance as the divergence function ρ for the uncertainty set, following Lee et al. (2021); Pan et al. (2023); Shi et al. (2023, 2024); Szita et al. (2003), defined by

$$\forall P, P' \in \Delta(\mathcal{S}) : \quad \rho_{\text{TV}}(P, P') := \frac{1}{2} \|P - P'\|_1. \quad (15)$$

For convenience, throughout the paper, we abbreviate $\mathcal{U}^{\sigma_i}(\cdot) := \mathcal{U}_{\rho_{\text{TV}}}^{\sigma_i}(\cdot)$ when there is no ambiguity.

Data collection mechanism: a generative model. We assume the agents interact with the environment through a generative model (simulator) (Kearns and Singh, 1999), which is a widely used sampling mechanism in both single-agent RL and MARL (Li et al., 2022; Zhang et al., 2020b). Specifically, at any time step h , we can collect an arbitrary number of independent samples from any state and joint action tuple $(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}$, generated based on the true *nominal* transition kernel P^0 :

$$s_{h,s,\mathbf{a}}^i \stackrel{i.i.d}{\sim} P_h^0(\cdot | s, \mathbf{a}), \quad i = 1, 2, \dots \quad (16)$$

Goal. Consider any fictitious RMGs $\mathcal{RMG}_\pi = \{\mathcal{S}, \{\mathcal{A}_i\}_{1 \leq i \leq n}, \{\mathcal{U}^{\sigma_i}(P^0)\}_{1 \leq i \leq n}, r, H\}$. While learning exact robust equilibria is computationally challenging and may not be necessary in practice, instead in this work, we focus on finding an approximate robust CCE (defined in (8)). Namely, a distribution $\xi := \{\xi_h\}_{h \in [H]} : [H] \mapsto \Delta(\mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i))$ is said to be an ε -robust CCE if

$$\text{gap}_{\text{CCE}}(\xi) := \max_{s \in \mathcal{S}, 1 \leq i \leq n} \left\{ \mathbb{E}_{\pi \sim \xi} [V_{i,1}^{*,\pi-i,\sigma_i}(s)] - \mathbb{E}_{\pi \sim \xi} [V_{i,1}^{\pi,\sigma_i}(s)] \right\} \leq \varepsilon. \quad (17)$$

Armed with a generative model of the nominal environment, the goal is to learn a robust CCE using as few samples as possible.

4.2 Algorithm design

With the sampling mechanism over a generative model in hand, we propose an algorithm called Robust-Q-FTRL to learn an ε -robust CCE in a sample-efficient manner, summarized in Algorithm 2 in the appendix. Robust-Q-FTRL draws inspiration from Q-FTRL developed in the standard MG literature (Li et al., 2022), but empowers tailored designs for learning in fictitious RMGs to achieve a robust equilibrium and to tackle statistical challenges arising from agents' nonlinear objectives. Overall, Robust-Q-FTRL takes a single pass to learn recursively from the final time step $h = H$ to $h = 1$. At each time step $h \in [H]$, an online learning process with K iterations will be executed. Before introducing the algorithm, we first concentrate on two essential steps customized for learning in fictitious RMGs.

Constructing the empirical model via N -sample estimation. For each time step h , we denote $\pi_{i,h}^k$ as the current learning policy of the i -th agent before the beginning of the k -th iteration for any $k \in [K]$. And we denote the joint product policy as $\pi_h^k = (\pi_{1,h}^k, \dots, \pi_{n,h}^k)$. During each iteration k , for each agent $i \in [n]$, we require to generate N independent samples from the generative model over each $(s, a_i) \in \mathcal{S} \times \mathcal{A}_i$ to obtain an empirical model, detailed in Algorithm 1. It includes an empirical reward function represented by $r_{i,h}^k \in \mathbb{R}^{S_{A_i}}$ and transition kernels denoted by $P_{i,h}^k \in \mathbb{R}^{S_{A_i} \times S}$. Note that different from standard MGs, we need to generate N samples instead of 1 sample per iteration to handle the additional statistical challenges induced by the non-linear objective of agents (N will be specified momentarily).

Algorithm 1: N -sample estimation ($\pi_h = \{\pi_{j,h}\}_{j \in [n]}, i, h$).

1 **Initialization:** the reward $\widehat{r} = 0 \in \mathbb{R}^{SA_i}$ and the transition model $\widehat{P} = 0 \in \mathbb{R}^{SA_i \times S}$.

2 **for** $(s, a_i) \in \mathcal{S} \times \mathcal{A}_i$ **do**

3 **for** $t = 1$ **to** N **do**

4 Sample $\mathbf{a}^t(s, a_i) = [a_j(s, a_i)]_{1 \leq j \leq n}$ constructed by independent actions drawn from policy:

$$a_j(s, a_i) \stackrel{\text{ind.}}{\sim} \pi_{j,h}(\cdot | s) \quad (j \neq i) \quad \text{and} \quad a_i(s, a_i) = a_i. \quad (18)$$

5 Sample from the generative model:

$$r_{i,h}^t(s, a_i) = r_{i,h}(s, \mathbf{a}^t(s, a_i)), \quad s_{s,a_i}^t \sim P_h(\cdot | s, \mathbf{a}^t(s, a_i)). \quad (19)$$

6 Set $\widehat{r}(s, a_i) = \frac{1}{N} \sum_{t \in [N]} r_{i,h}^t(s, a_i)$ and $\widehat{P}(s' | s, a_i) = \frac{1}{N} \sum_{t \in [N]} \mathbb{1}\{s_{s,a_i}^t = s'\}$.

7 **Return:** empirical model $(\widehat{r}, \widehat{P})$.

Estimating robust Q-function of the current policy π_h^k . We denote $\widehat{V}_{i,h} \in \mathbb{R}^S$ as the estimation of the i -th agent's robust value function at time step h . For any agent i , with the empirical reward function $r_{i,h}^k$, empirical kernel $P_{i,h}^k$, and the estimated robust value function $\widehat{V}_{i,h+1}$ at the next step in hand, the robust Q-function $\{q_{i,h}^k\}$ of current policy π_h^k can be estimated as:

$$\forall (i, h, s, a_i) \in [n] \times [H] \times \mathcal{S} \times \mathcal{A}_i: \quad q_{i,h}^k(s, a_i) = r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h}^k)} \mathcal{P} \widehat{V}_{i,h+1}. \quad (20)$$

Unlike the linear function w.r.t. $P_{i,h}^k$ in standard MGs, (20) lacks a closed form and introduces an additional inner optimization problem. Solving (20) directly is computationally challenging due to the need to optimize over an S -dimensional probability simplex, with complexity growing exponentially with the state space size S . Fortunately, by applying strong duality, we can solve (20) equivalently via its dual problem with tractable computation (Iyengar, 2005):

$$q_{i,h}^k(s, a_i) = r_{i,h}^k(s, a_i) + \max_{\alpha \in [\min_s \widehat{V}_{i,h+1}(s), \max_s \widehat{V}_{i,h+1}(s)]} \left\{ P_{i,h}^k \left[\widehat{V}_{i,h+1} \right]_{\alpha} - \sigma_i \left(\alpha - \min_{s'} \left[\widehat{V}_{i,h+1} \right]_{\alpha} (s') \right) \right\}, \quad (21)$$

where $[V]_{\alpha}$ denotes the clipped version of any vector $V \in \mathbb{R}^S$ determined by some level $\alpha \geq 0$, namely,

$$[V]_{\alpha}(s) := \begin{cases} \alpha, & \text{if } V(s) > \alpha, \\ V(s), & \text{otherwise.} \end{cases} \quad (22)$$

The above modules are key components of Robust-Q-FTRL, serving for constructing nonlinear robust objectives in the online learning process and ensuring the desired statistical accuracy.

Overall pipeline of Robust-Q-FTRL. With these technical modules in place, we introduce Robust-Q-FTRL, which follows a similar online learning procedure as Q-FTRL for standard MGs (Li et al., 2022). The complete procedure is summarized in Algorithm 2. We denote $Q_{i,h}^k \in \mathbb{R}^{SA_i}$ as the estimated robust Q-function of the equilibrium for the i -th agent at the k -th iteration of time step h . To begin with, Robust-Q-FTRL initialize the robust value function, robust Q-function $\widehat{V}_{i,H+1}(s) = Q_{i,h}^0(s, a_i) = 0$, and the policy $\pi_{i,h}^1(a_i | s) = 1/A_i$ for all $i \in [n]$. Then subsequently from the final time step $h = H$ to $h = 1$, for each step h , a K iterations online learning process will be executed. At each k -th iteration, given current policy π_h^k , as described above, an empirical model $(\{r_{i,h}^k\}_{i \in [n]})$ and $(\{P_{i,h}^k\}_{i \in [n]})$ is constructed by N -sample estimation (cf. algorithm 1). Then the robust Q-function $\{q_{i,h}^k\}_{i \in [n]}$ of the current policy π_h^k is estimated by (21).

Now we are ready to specify the loss objective and proceed the online learning procedure. With the current one-step update $\{q_{i,h}^k\}$, we update the Q-estimate as $Q_{i,h}^k = (1 - \alpha_k)Q_{i,h}^{k-1} + \alpha_k q_{i,h}^k$. Here, $\{\alpha_k\}_{k \in [K]}$

is a series of rescaled linear learning rates with some $c_\alpha \geq 24$,

$$\forall k \in [K] : \quad \alpha_k = \frac{c_\alpha \log K}{k-1 + c_\alpha \log K} \quad \text{and} \quad \alpha_k^n = \begin{cases} \alpha_k \prod_{i=k+1}^n (1 - \alpha_i), & \text{if } 0 < k < n \leq K \\ \alpha_n & \text{if } k = n \end{cases}. \quad (23)$$

Let the Q-estimate be the online learning loss objective at this moment, we apply the Follow-the-Regularized-Leader strategy (Li et al., 2022; Shalev-Shwartz, 2012) to update the corresponding policy as below:

$$\pi_{i,h}^{k+1}(a_i | s) = \frac{\exp(\eta_{k+1} Q_{i,h}^k(s, a_i))}{\sum_{a'} \exp(\eta_{k+1} Q_{i,h}^k(s, a'))} \quad \text{with} \quad \eta_{k+1} = \sqrt{\frac{\log K}{\alpha_k H}}, \quad k = 1, 2, \dots$$

This is a widely used adaptive sampling and learning procedure for MARL problems.

After completing K iterations for time step h , we finalize the robust value function estimation by setting it to its confidence upper bound, incorporating carefully designed optimistic bonus terms $\{\beta_{i,h}\}$ as: for all $(i, h, s) \in [n] \times [H] \times \mathcal{S}$,

$$\beta_{i,h}(s) = c_b \sqrt{\frac{\log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{KH}} \sum_{k=1}^K \alpha_k^K \left\{ \text{Var}_{\pi_{i,h}^k(\cdot|s)}(q_{i,h}^k(s, \cdot)) + H \right\}, \quad (24)$$

where c_b denotes some absolute constant, $\delta \in (0, 1)$ is the high probability threshold, Finally, after the recursive learning process ends for all time steps $h = H, H-1, \dots, 1$, we output a distribution of product policy $\hat{\xi} = \{\hat{\xi}_h\}_{h \in [H]}$ over all the policies $\{\pi_h^k = (\pi_{1,h}^k \times \dots \times \pi_{n,h}^k)\}_{h \in [H], k \in [K]}$ occurs during the process that defined as

$$\forall (h, k) \in [H] \times [K] : \quad \xi_h(\pi_h^k) := \alpha_k. \quad (25)$$

4.3 Theoretical guarantees

In this section, we provide the theoretical guarantees for the sample complexity of our proposed algorithm Robust-Q-FTRL, shown as below:

Theorem 2 (Upper bound). *Using the TV uncertainty set defined in (15). Consider any $\delta \in (0, 1)$ and any fictitious RMGs $\mathcal{RMG}_\pi = \{\mathcal{S}, \{A_i\}_{1 \leq i \leq n}, \{\mathcal{U}^{\sigma_i}(P^0, \cdot)\}_{1 \leq i \leq n}, r, H\}$ with $\sigma_i \in (0, 1]$ for all $i \in [n]$. For any $\varepsilon \leq \sqrt{\min\left\{H, \frac{1}{\min_{1 \leq i \leq n} \sigma_i}\right\}}$, Algorithm 2 can output an ε -robust CCE $\hat{\xi}$, i.e.,*

$$\text{gap}_{\text{CCE}}(\hat{\xi}) := \max_{s \in \mathcal{S}, 1 \leq i \leq n} \left\{ \mathbb{E}_{\pi \sim \hat{\xi}} [V_{i,1}^{*,\pi-i,\sigma_i}(s)] - \mathbb{E}_{\pi \sim \hat{\xi}} [V_{i,1}^{\pi,\sigma_i}(s)] \right\} \leq \varepsilon$$

with probability at least $1 - \delta$, as long as

$$N \geq \frac{C_1 H^2}{\varepsilon^2} \min \left\{ \frac{1}{\min_{1 \leq i \leq n} \sigma_i}, H \right\}, \quad K \geq \frac{C_1 H^3}{\varepsilon^2}. \quad (27)$$

Here C_1 is some universal large enough constant. Namely, it is sufficient if the total number of samples acquired in the learning process obeys

$$N_{\text{all}} := HKNS \sum_{1 \leq i \leq n} A_i \geq \frac{(C_1)^2 H^6 S \sum_{1 \leq i \leq n} A_i}{\varepsilon^4} \min \left\{ H, \frac{1}{\min_{1 \leq i \leq n} \sigma_i} \right\}.$$

Before we jump into more discussions of the above theorem, in addition, we introduce the information-theoretic minimax lower bound for this problem as well.

Algorithm 2: Robust-Q-FTRL

- 1 *Input:* learning rates $\{\alpha_k\}$ and $\{\eta_{k+1}\}$, number of iterations K per time step, and number of samples N per iteration.
 - 2 *Initialization:* $\widehat{V}_{i,H+1}(s) = Q_{i,h}^0(s, a_i) = 0$ and $\pi_{i,h}^1(a_i | s) = 1/A_i$ for all $i \in [n]$ and then all $(h, s, a_i) \in [H] \times \mathcal{S} \times \mathcal{A}_i$.
// start recursive learning process.
 - 3 **for** $h = H, H - 1, \dots, 1$ **do**
 - 4 **for** $k = 1, 2, \dots, K$ **do**
 - 5 **for** $i = 1, 2, \dots, n$ **do**
 - 6 // construct empirical models and estimate current robust Q-function
 - 7 $(r_{i,h}^k, P_{i,h}^k) \leftarrow N$ -sample estimation $(\pi_h^k = \{\pi_{j,h}^k\}_{j \in [n]}, i, h)$. (Algorithm 1)
 - 8 Estimate the robust Q-function $q_{i,h}^k$ of current π_h^k according to (21).
 - 9 // Online learning procedure
 - 10 Update the Q-estimate $Q_{i,h}^k = (1 - \alpha_k)Q_{i,h}^{k-1} + \alpha_k q_{i,h}^k$ and apply FTRL:
$$\forall (s, a_i) \in \mathcal{S} \times \mathcal{A}_i : \pi_{i,h}^{k+1}(a_i | s) = \frac{\exp(\eta_{k+1} Q_{i,h}^k(s, a_i))}{\sum_{a'} \exp(\eta_{k+1} Q_{i,h}^k(s, a'))}.$$
 - 11 // set the final robust value estimate at time step h .
 - 12 **for** $i = 1, 2, \dots, n$ **do**
 - 13 For all $s \in \mathcal{S}$: set $\beta_{i,h}(s)$ to be the optimistic bonus term in (24) and
$$\widehat{V}_{i,h}(s) = \min \left\{ \sum_{k=1}^K \alpha_k^K \langle \pi_{i,h}^k(\cdot | s), q_{i,h}^k(s, \cdot) \rangle + \beta_{i,h}(s), H - h + 1 \right\}, \quad (26)$$
- 14 *Output:* a set of policies $\{\pi_h^k = (\pi_{1,h}^k \times \dots \times \pi_{n,h}^k)\}_{k \in [K], h \in [H]}$ and a distribution $\widehat{\xi} = \{\widehat{\xi}_h\}_{h \in [H]}$ over them. For any time step h , $\widehat{\xi}_h$ is the distribution over $\{\pi_h^k\}_{k \in [K]}$ so that $\widehat{\xi}_h(\pi_h^k) = \alpha_k^K$.
-

Lower bound for learning in fictitious RMGs. Considering the instances of fictitious RMGs that the action space for all the agents except the i -th agent contains only a single action, i.e., $A_j = 1$ for all $j \neq i$. As such, all the agents $j \neq i$ will take a fixed action and the game reduces to a single-agent robust MDP with (s, a) -rectangularity condition (Zhou et al., 2021). So the goal of finding the robust equilibrium — robust NE/CCE also degrades to finding the optimal policy of the i -th agent. Invoking the results from Shi et al. (2024, Theorem 2), the lower bound for the class of fictitious RMGs is achieved directly: consider any tuple $\{S, \{A_i\}_{1 \leq i \leq n}, \{\sigma_i\}_{1 \leq i \leq n}, H\}$ obeying $\sigma_i \in (0, 1 - c_0]$ with $0 < c_0 \leq \frac{1}{4}$ being any small enough positive constant, and $H > 16 \log 2$. Let

$$\varepsilon \leq \begin{cases} \frac{c_1}{H}, & \text{if } \sigma_i \leq \frac{c_1}{2H}, \\ 1 & \text{otherwise} \end{cases} \quad (28)$$

for any $c_1 \leq \frac{1}{4}$. We can construct a set of fictitious RMGs $\mathcal{M} = \{\mathcal{RMG}_i\}_{i \in [I]}$, such that for any dataset generated from the nominal environment with in total N_{all} independent samples over all state-action pairs, we have

$$\inf_{\widehat{\xi} \in [H] \rightarrow \Delta(\mathcal{S} \rightarrow \prod_{i=1}^n \mathcal{A}_i)} \max_{\mathcal{MG}_i \in \mathcal{M}} \left\{ \mathbb{P}_{\mathcal{MG}_i}(\text{gap}_{\text{CCE}}(\widehat{\xi}) > \varepsilon) \right\} \geq \frac{1}{8}, \quad (29)$$

provided that

$$N_{\text{all}} \leq \frac{C_2 S H^3 \max_{1 \leq i \leq n} A_i}{\varepsilon^2} \min \left\{ H, \frac{1}{\min_{1 \leq i \leq n} \sigma_i} \right\}. \quad (30)$$

Here, the infimum is taken over all estimators $\widehat{\xi}$, $\mathbb{P}_{\mathcal{RMG}_i}$ denotes the probability when the game is \mathcal{MG}_i for all $\mathcal{MG}_i \in \mathcal{M}$, and C_2 is some small enough constant.

Armed with both the upper bound (Theorem 2) and lower bound in (30), we are now ready to discuss the implications of our sample complexity results.

Breaking the curse of multiagency in the sample complexity for RMGs. Theorem 2 demonstrates that for any fictitious RMGs, Robust-Q-FTRL algorithm finds an ϵ -robust CCE when the total number of samples exceeds

$$\tilde{O}\left(\frac{SH^6 \sum_{1 \leq i \leq n} A_i}{\epsilon^4} \min\left\{H, \frac{1}{\min_{1 \leq i \leq n} \sigma_i}\right\}\right). \quad (31)$$

To the best of our knowledge, Robust-Q-FTRL with the above sample complexity in (31) is the first algorithm for RMGs breaking the curse of multiagency, regardless of the types of uncertainty sets. Our sample complexity depends linearly on the sum of each agent’s actions $\sum_{i=1}^n A_i$ rather than their product $\prod_{i=1}^n A_i$ —making the algorithm highly scalable as the number of agents increases. Nonetheless, there still exist gaps between our upper bound and the lower bound—especially in terms of the dependency on the horizon length H and the accuracy level ϵ —an interesting direction to investigate in the future.

Comparisons with prior works. All prior works focus on learning equilibria for a different kind of robust MGs with (s, \mathbf{a}) -rectangular uncertainty sets (Blanchet et al., 2023; Ma et al., 2023; Shi et al., 2024). However, the state-of-the-art sample complexity $\tilde{O}\left(\frac{SH^3 \prod_{i=1}^n A_i}{\epsilon^2} \min\left\{H, \frac{1}{\min_{1 \leq i \leq n} \sigma_i}\right\}\right)$ (Shi et al., 2024) still suffers from the curse of multiagency with an exponential dependency on the number of agents when all agents have equal action spaces, which uses nonadaptive sampling. Our work circumvents the curse of multiagency by resorting to a tailored adaptive sampling and online learning procedure, together with the introduction of a new class of fictitious RMGs, providing a fresh perspective to learning RMGs.

Technical insights. For sample complexity analysis, while previous works have addressed the curse of multiagency in sequential games like standard Markov games (MGs) and Markov potential games, these methods are not directly applicable to RMGs. Prior approaches assume a linear relationship between the value function and the transition kernel, allowing statistical errors across K iterations to cancel out. However, in RMGs, the robust value function, due to its distributionally robust requirement, is highly nonlinear and often lacks a closed form, making it impossible to linearly aggregate statistical errors. To tackle the nonlinear challenges in RMGs, we design a variance-style bonus term through non-trivial decomposition and control of auxiliary statistical errors caused by nonlinearity, resulting in a tight upper bound on regret during the online learning process.

5 Conclusion

Robustness in MARL presents greater challenges than in single-agent RL due to the strategic interactions between agents in a game-theoretic setting. This work proposes a new class of RMGs with fictitious uncertainty sets that naturally extends from robust single-agent RL and addresses more realistic scenarios where each agent’s uncertainty is influenced by both the environment and the behavior of others. We then propose Robust-Q-FTRL, the first algorithm to break the curse of multiagency in robust Markov games regardless of the uncertainty set definitions, with sample complexity scaling polynomially with all key parameters. This opens up new research directions in MARL, such as uncertainty set selection, equilibrium refinement, and sample-efficient algorithm design.

Acknowledgements

The work of Y. Chi is supported in part by the grants NSF CCF-2106778 and CNS-2148212, and by funds from federal agency and industry partners as specified in the Resilient & Intelligent NextG Systems (RINGS) program. The work of L. Shi is supported in part by the Resnick Institute and Computing, Data, and Society Postdoctoral Fellowship at California Institute of Technology. The work of E. Mazumdar is supported in

part from NSF-2240110. The work of A. Wierman is supported in part from the NSF through CNS-2146814, CPS-2136197, CNS-2106403, NGSDF-2105648.

References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. (2023). Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica: Journal of the Econometric Society*, pages 1–18.
- Badrinath, K. P. and Kalathil, D. (2021). Robust reinforcement learning using least squares policy iteration with provable performance guarantees. In *International Conference on Machine Learning*, pages 511–520. PMLR.
- Bai, Y. and Jin, C. (2020). Provable self-play algorithms for competitive reinforcement learning. In *International Conference on Machine Learning*, pages 551–560. PMLR.
- Baker, B. (2020). Emergent reciprocity and team formation from randomized uncertain social preferences. *Advances in neural information processing systems*, 33:15786–15799.
- Balaji, B., Mallya, S., Genc, S., Gupta, S., Dirac, L., Khare, V., Roy, G., Sun, T., Tao, Y., Townsend, B., et al. (2019). DeepRacer: Educational autonomous racing platform for experimentation with sim2real reinforcement learning. *arXiv preprint arXiv:1911.01562*.
- Bertsimas, D., Gupta, V., and Kallus, N. (2018). Data-driven robust optimization. *Mathematical Programming*, 167(2):235–292.
- Blanchet, J., Lu, M., Zhang, T., and Zhong, H. (2023). Double pessimism is provably efficient for distributionally robust offline reinforcement learning: Generic algorithm and robust partial coverage. *arXiv preprint arXiv:2305.09659*.
- Blanchet, J., Lu, M., Zhang, T., and Zhong, H. (2024). Double pessimism is provably efficient for distributionally robust offline reinforcement learning: Generic algorithm and robust partial coverage. *Advances in Neural Information Processing Systems*, 36.
- Blanchet, J. and Murthy, K. (2019). Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600.
- Cui, Q., Zhang, K., and Du, S. (2023). Breaking the curse of multiagents in a large state space: RL in markov games with independent linear function approximation. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 2651–2652. PMLR.
- Daskalakis, C., Golowich, N., and Zhang, K. (2023). The complexity of markov equilibrium in stochastic games. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 4180–4234. PMLR.
- Duchi, J. and Namkoong, H. (2018). Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750*.
- Gao, R. (2020). Finite-sample guarantees for wasserstein distributionally robust optimization: Breaking the curse of dimensionality. *arXiv preprint arXiv:2009.04382*.
- Han, S., Su, S., He, S., Han, S., Yang, H., and Miao, F. (2022). What is the solution for state adversarial multi-agent reinforcement learning? *arXiv preprint arXiv:2212.02705*.
- Iyengar, G. N. (2005). Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280.
- Jin, C., Liu, Q., Wang, Y., and Yu, T. (2021). V-learning—a simple, efficient, decentralized algorithm for multiagent RL. *arXiv preprint arXiv:2110.14555*.

- Kakutani, S. (1941). A generalization of brouwer’s fixed point theorem.
- Kannan, S. S., Venkatesh, V. L., and Min, B.-C. (2023). Smart-LLM: Smart multi-agent robot task planning using large language models. *arXiv preprint arXiv:2309.10062*.
- Kardeş, E., Ordóñez, F., and Hall, R. W. (2011). Discounted robust stochastic games and an application to queueing control. *Operations research*, 59(2):365–382.
- Kearns, M. J. and Singh, S. P. (1999). Finite-sample convergence rates for Q-learning and indirect algorithms. In *Advances in neural information processing systems*, pages 996–1002.
- Kober, J., Bagnell, J. A., and Peters, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274.
- Lattimore, T. and Szepesvári, C. (2020). *Bandit algorithms*. Cambridge University Press.
- Lee, J., Jeon, W., Lee, B., Pineau, J., and Kim, K.-E. (2021). Optidice: Offline policy optimization via stationary distribution correction estimation. In *International Conference on Machine Learning*, pages 6120–6130. PMLR.
- Leibo, J. Z., Zambaldi, V., Lanctot, M., Marecki, J., and Graepel, T. (2017). Multi-agent reinforcement learning in sequential social dilemmas. *arXiv preprint arXiv:1702.03037*.
- Li, G., Cai, C., Chen, Y., Wei, Y., and Chi, Y. (2024). Is Q-learning minimax optimal? a tight sample complexity analysis. *Operations Research*, 72(1):222–236.
- Li, G., Chi, Y., Wei, Y., and Chen, Y. (2022). Minimax-optimal multi-agent RL in Markov games with a generative model. *Advances in Neural Information Processing Systems*, 35:15353–15367.
- Li, G., Yan, Y., Chen, Y., and Fan, J. (2023). Minimax-optimal reward-agnostic exploration in reinforcement learning. *arXiv preprint arXiv:2304.07278*.
- Li, S., Wu, Y., Cui, X., Dong, H., Fang, F., and Russell, S. (2019). Robust multi-agent reinforcement learning via minimax deep deterministic policy gradient. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 4213–4220.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. (2015). Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier.
- Lu, M., Zhong, H., Zhang, T., and Blanchet, J. (2024). Distributionally robust reinforcement learning with interactive data collection: Fundamental hardness and near-optimal algorithm. *arXiv preprint arXiv:2404.03578*.
- Ma, S., Chen, Z., Zou, S., and Zhou, Y. (2023). Decentralized robust v-learning for solving markov games with model uncertainty. *Journal of Machine Learning Research*, 24(371):1–40.
- Mahmood, A. R., Korenkevych, D., Vasan, G., Ma, W., and Bergstra, J. (2018). Benchmarking reinforcement learning algorithms on real-world robots. In *Conference on robot learning*, pages 561–591. PMLR.
- McMahan, J., Artiglio, G., and Xie, Q. (2024). Roping in uncertainty: Robustness and regularization in markov games. *arXiv preprint arXiv:2406.08847*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., and Ostrovski, G. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.
- Moulin, H., Ray, I., and Gupta, S. S. (2014). Coarse correlated equilibria in an abatement game. Technical report, Cardiff Economics Working Papers.

- Moulin, H. and Vial, J.-P. (1978). Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3):201–221.
- Nash, J. (1951). Non-cooperative games. *Annals of mathematics*, pages 286–295.
- Nilim, A. and El Ghaoui, L. (2005). Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798.
- Pan, Y., Chen, Y., and Lin, F. (2023). Adjustable robust reinforcement learning for online 3d bin packing. *arXiv preprint arXiv:2310.04323*.
- Panaganti, K. and Kalathil, D. (2022). Sample complexity of robust reinforcement learning with a generative model. In *International Conference on Artificial Intelligence and Statistics*, pages 9582–9602. PMLR.
- Pinto, L., Davidson, J., Sukthankar, R., and Gupta, A. (2017). Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, pages 2817–2826. PMLR.
- Rahimian, H. and Mehrotra, S. (2019). Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*.
- Rubinstein, A. (2017). Settling the complexity of computing approximate two-player nash equilibria. *ACM SIGecom Exchanges*, 15(2):45–49.
- Rusu, A. A., Večerík, M., Rothörl, T., Heess, N., Pascanu, R., and Hadsell, R. (2017). Sim-to-real robot learning from pixels with progressive nets. In *Conference on robot learning*, pages 262–270. PMLR.
- Shalev-Shwartz, S. (2007). Online learning: Theory, algorithms, and applications.
- Shalev-Shwartz, S. (2012). Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194.
- Shalev-Shwartz, S. and Singer, Y. (2007). A primal-dual perspective of online learning algorithms. *Machine Learning*, 69(2):115–142.
- Shapley, L. S. (1953). Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100.
- Shi, L. and Chi, Y. (2024). Distributionally robust model-based offline reinforcement learning with near-optimal sample complexity. *Journal of Machine Learning Research*, 25(200):1–91.
- Shi, L., Li, G., Wei, Y., Chen, Y., Geist, M., and Chi, Y. (2023). The curious price of distributional robustness in reinforcement learning with a generative model. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pages 79903–79917.
- Shi, L., Mazumdar, E., Chi, Y., and Wierman, A. (2024). Sample-efficient robust multi-agent reinforcement learning in the face of environmental uncertainty. In *Forty-first International Conference on Machine Learning*.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489.
- Slumbers, O., Mguni, D. H., Blumberg, S. B., Mcaleer, S. M., Yang, Y., and Wang, J. (2023). A game-theoretic framework for managing risk in multi-agent systems. In *International Conference on Machine Learning*, pages 32059–32087. PMLR.
- Song, Z., Mei, S., and Bai, Y. (2021). When can we learn general-sum Markov games with a large number of players sample-efficiently? *arXiv preprint arXiv:2110.04184*.
- Szita, I., Takács, B., and Lorincz, A. (2003). ϵ -mdps: Learning in varying environments. *Journal of Machine Learning Research*, 3(1).

- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., and Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30. IEEE.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *science*, 185(4157):1124–1131.
- Vial, D., Shakkottai, S., and Srikant, R. (2022). Robust multi-agent bandits over undirected graphs. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 6(3):1–57.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., et al. (2019). Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354.
- Wang, H., Shi, L., and Chi, Y. (2024). Sample complexity of offline distributionally robust linear Markov decision processes. *arXiv preprint arXiv:2403.12946*.
- Wang, Y., Liu, Q., Bai, Y., and Jin, C. (2023). Breaking the curse of multiagency: Provably efficient decentralized multi-agent RL with function approximation. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 2793–2848. PMLR.
- Wu, Y., McMahan, J., Zhu, X., and Xie, Q. (2024). Data poisoning to fake a nash equilibria for markov games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 15979–15987.
- Xu, Z., Panaganti, K., and Kalathil, D. (2023). Improved sample complexity bounds for distributionally robust reinforcement learning. *arXiv preprint arXiv:2303.02783*.
- Yang, W., Zhang, L., and Zhang, Z. (2022). Toward theoretical understandings of robust Markov decision processes: Sample complexity and asymptotics. *The Annals of Statistics*, 50(6):3223–3248.
- Yeh, C., Meng, C., Wang, S., Driscoll, A., Rozi, E., Liu, P., Lee, J., Burke, M., Lobell, D. B., and Ermon, S. (2021). Sustainbench: Benchmarks for monitoring the sustainable development goals with machine learning. *arXiv preprint arXiv:2111.04724*.
- Zeng, L., Qiu, D., and Sun, M. (2022). Resilience enhancement of multi-agent reinforcement learning-based demand response against adversarial attacks. *Applied Energy*, 324:119688.
- Zhang, H., Chen, H., Boning, D., and Hsieh, C.-J. (2021). Robust reinforcement learning on state observations with learned optimal adversary. *arXiv preprint arXiv:2101.08452*.
- Zhang, H., Chen, H., Xiao, C., Li, B., Liu, M., Boning, D., and Hsieh, C.-J. (2020a). Robust deep reinforcement learning against adversarial perturbations on state observations. *Advances in Neural Information Processing Systems*, 33:21024–21037.
- Zhang, K., Kakade, S., Basar, T., and Yang, L. (2020b). Model-based multi-agent RL in zero-sum Markov games with near-optimal sample complexity. *Advances in Neural Information Processing Systems*, 33:1166–1178.
- Zhang, K., Sun, T., Tao, Y., Genc, S., Mallya, S., and Basar, T. (2020c). Robust multi-agent reinforcement learning with model uncertainty. *Advances in neural information processing systems*, 33:10571–10583.
- Zhang, R., Shamma, J., and Li, N. (2024). Equilibrium selection for multi-agent reinforcement learning: A unified framework. *arXiv preprint arXiv:2406.08844*.
- Zhou, Z., Bai, Q., Zhou, Z., Qiu, L., Blanchet, J., and Glynn, P. (2021). Finite-sample regret bound for distributionally robust offline tabular reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3331–3339. PMLR.
- Zhou, Z. and Liu, G. (2023). Robustness testing for multi-agent reinforcement learning: State perturbations on critical agents. *arXiv preprint arXiv:2306.06136*.

A Preliminaries

Denoting the vectors $x = [x_i]_{1 \leq i \leq n}$ and $y = [y_i]_{1 \leq i \leq n}$, we use the notation $x \leq y$ (or $x \geq y$) to signify that $x_i \leq y_i$ (or $x_i \geq y_i$) for every $1 \leq i \leq n$. The Hadamard product of two vectors x and y in \mathbb{R}^S is denoted as $x \circ y = [x(s) \cdot y(s)]_{s \in \mathcal{S}}$. In addition, for any series of vectors $\{x_i\}_{i \in [S]}$, $\text{diag}(x_1, x_2, \dots, x_S)$ denote a block diagonal matrix by placing each given vector x_i along the diagonal, with zeros filling the off-diagonal blocks. $\mathbf{0}$ (or $\mathbf{1}$) represents the all-zero (or all-one) vector, while $e_i \in \mathbb{R}^S$ denotes a basis vector of dimension S with 1 in the i -th position and 0 elsewhere.

A.1 Additional matrix and vector notation

Before continuing, we introduce or recall some matrix and vector notation that will be used throughout the paper. In particular, for any joint policy $\pi : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A})$ and any $(i, h) \in [n] \times [H]$:

Matrices for policy. We introduce three matrices associated with π , i.e., $\Pi_h^\pi \in \mathbb{R}^{S \times S} \prod_{i=1}^n A_i$, $\Pi_h^{\pi-i} \in \mathbb{R}^{S \times S} \prod_{j \neq i} A_j$, and $\Pi_h^{\pi_i} \in \mathbb{R}^{S \times S A_i}$, which are defined as block diagonal matrices that adhere to the following properties:

- The matrix Π_h^π is given by $\text{diag}(\pi_h(1)^\top, \pi_h(2)^\top, \dots, \pi_h^\top(S))$, where $\pi_h(s) = [\pi_h(\mathbf{a} | s)]_{\mathbf{a} \in \mathcal{A}} \in \Delta(\mathcal{A})$ for each $s \in \mathcal{S}$ represents the joint policy vectors across all agents.
- The matrix $\Pi_h^{\pi-i}$ can be expressed as $\text{diag}(\pi_{-i,h}(1)^\top, \pi_{-i,h}(2)^\top, \dots, \pi_{-i,h}^\top(S))$, where $\pi_{-i,h}(s) = [\pi_h(\mathbf{a}_{-i} | s)]_{\mathbf{a}_{-i} \in \mathcal{A}_{-i}} \in \Delta(\mathcal{A}_{-i})$ for all $s \in \mathcal{S}$ denotes the joint policy vectors from all agents except agent i .
- The matrix $\Pi_h^{\pi_i}$ is defined as $\text{diag}(\pi_{i,h}(1)^\top, \pi_{i,h}(2)^\top, \dots, \pi_{i,h}^\top(S))$, where $\pi_{i,h}(s) = [\pi_{i,h}(a_i | s)]_{a_i \in \mathcal{A}_i} \in \Delta(\mathcal{A}_i)$ for each $s \in \mathcal{S}$ represents the policy of the i -th agent.

Reward vectors. We recall the definition of $r_{i,h}$ and introduce the reward vectors $r_{i,h}^\pi$ and $r_{i,h}^{\pi-i}$ as follows:

- Let $r_{i,h} = [r_{i,h}(s, \mathbf{a})]_{(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}} \in \mathbb{R}^S \prod_{i=1}^n A_i$ represent the reward function for the i -th player at time step h , where \mathcal{S} is the state space and \mathcal{A} is the action space.
- The reward vector $r_{i,h}^\pi \in \mathbb{R}^S$ corresponds to the joint policy $\pi = \{\pi_h\}_{h \in [H]}$ at time step h . Specifically, for each $s \in \mathcal{S}$, $r_{i,h}^\pi(s) = \mathbb{E}_{\mathbf{a} \sim \pi_h(s)}[r_{i,h}(s, \mathbf{a})]$, where the expectation is taken over the actions \mathbf{a} drawn from policy π_h in state s .
- The reward vector $r_{i,h}^{\pi-i} \in \mathbb{R}^{S A_i}$ corresponds to the joint policy $\pi_{-i} = \{\pi_{-i,h}\}_{h \in [H]}$ at time step h , excluding agent i . Specifically, for all $s \in \mathcal{S}$ and $a_i \in \mathcal{A}_i$, $r_{i,h}^{\pi-i}(s, a_i) = \mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,h}(s)}[r_{i,h}(s, \mathbf{a})]$, where the expectation is over the actions \mathbf{a}_{-i} drawn from the joint policy $\pi_{-i,h}$ for all agents except agent i .

Matrices for transition variants. We first introduce the following notations related to transitions associated with the nominal transition kernel and the policy π :

- Define $P_h^0 \in \mathbb{R}^S \prod_{i=1}^n A_i \times S$, the matrix representing the nominal transition kernel at time step h . Specifically, for any $(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}$, $P_{h,s,\mathbf{a}}^0 \in \mathbb{R}^{1 \times S}$ represents the row corresponding to the state-action pair (s, \mathbf{a}) .
- Define $P_h^{\pi-i} \in \mathbb{R}^{S A_i \times S}$, the matrix representing the nominal transition kernel at time step h , associated with the joint policy π_{-i} . Specifically, for all $s, s' \in \mathcal{S}$ and $a_i \in \mathcal{A}_i$, $P_{h,s,a_i}^{\pi-i}(s') = \mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,h}(s)}[P_{h,s,\mathbf{a}}^0(s')]$. Here, $P_{h,s,a_i}^{\pi-i} \in \mathbb{R}^{1 \times S}$ represents the row corresponding to the state-action pair (s, a_i) .
- Let $\widehat{P}_{i,h}^{\pi-i} \in \mathbb{R}^{S A_i \times S}$ denote the empirical transition kernel matrix at time step h , associated with the joint policy π_{-i} and agent i . Similarly, $\widehat{P}_{h,s,a_i}^{\pi-i} \in \mathbb{R}^{1 \times S}$ represents the row corresponding to the state-action pair (s, a_i) .
- Define $\underline{P}_h^\pi \in \mathbb{R}^{S \times S}$ as $\underline{P}_h^\pi := \Pi_h^\pi P_h^0$, where Π_h^π is the policy matrix at time step h under joint policy π .
- Define $\widehat{\underline{P}}_{i,h}^\pi \in \mathbb{R}^{S \times S}$ as $\widehat{\underline{P}}_{i,h}^\pi := \Pi_h^{\pi_i} \widehat{P}_{i,h}^{\pi-i}$, where $\Pi_h^{\pi_i}$ denotes the policy matrix at time step h under policy π_i .

We introduce matrix notations for transitions that are associated not only with the nominal transition and policy π , but also with value functions:

- For time step $h \in [H]$, joint policy π , and a value vector $V \in \mathbb{R}^S$, we define $P_{i,h}^{\pi-i,V} \in \mathbb{R}^{SA_i \times S}$ as the matrix representing the worst-case transition probability kernel within the uncertainty set for agent i , centered around the nominal kernel. The row corresponding to the state-action pair (s, a_i) in $P_{i,h}^{\pi-i,V}$, denoted as $P_{i,h,s,a_i}^{\pi-i,V} \in \mathbb{R}^S$, is given by:

$$P_{i,h,s,a_i}^{\pi-i,V} = \operatorname{argmin}_{\mathcal{P} \in \mathcal{U}_\rho^{\sigma_i}(P_{h,s,a_i}^{\pi-i})} \mathcal{P}V. \quad (32a)$$

We also define the transition matrices for specific value vectors as:

$$P_{i,h}^{\pi,V} := P_{i,h}^{\pi-i,V_{i,h+1}^{\pi,\sigma_i}} \quad \text{and} \quad P_{i,h,s,a_i}^{\pi,V} := P_{i,h,s,a_i}^{\pi-i,V_{i,h+1}^{\pi,\sigma_i}} = \operatorname{argmin}_{\mathcal{P} \in \mathcal{U}_\rho^{\sigma_i}(P_{h,s,a_i}^{\pi-i})} \mathcal{P}V_{i,h+1}^{\pi,\sigma_i}. \quad (32b)$$

Finally, we define square matrices $\underline{P}_{i,h}^{\pi,V} \in \mathbb{R}^{S \times S}$ as: $\underline{P}_{i,h}^{\pi,V} := \Pi_h^\pi P_{i,h}^{\pi-i,V}$.

- By replacing the nominal transition kernel with the empirical transition kernel, we similarly define $\widehat{P}_{i,h}^{\pi-i,V}$ as the worst-case probability transition kernel within the uncertainty set for agent i , centered around the empirical kernel $\widehat{P}_{i,h}^{\pi-i}$. The row corresponding to the state-action pair (s, a_i) in $\widehat{P}_{i,h}^{\pi-i,V}$ is denoted as $\widehat{P}_{i,h,s,a_i}^{\pi-i,V} \in \mathbb{R}^S$ and is defined as:

$$\widehat{P}_{i,h,s,a_i}^{\pi-i,V} = \operatorname{argmin}_{\mathcal{P} \in \mathcal{U}_\rho^{\sigma_i}(\widehat{P}_{i,h,s,a_i}^{\pi-i})} \mathcal{P}V. \quad (32c)$$

The transition matrices $\widehat{P}_{i,h}^{\pi,V}$ for specific value vectors are defined as:

$$\widehat{P}_{i,h}^{\pi,V} := \widehat{P}_{i,h}^{\pi-i,V_{i,h+1}^{\pi,\sigma_i}} \quad \text{and} \quad \widehat{P}_{i,h,s,a_i}^{\pi,V} := \widehat{P}_{i,h,s,a_i}^{\pi-i,V_{i,h+1}^{\pi,\sigma_i}} = \operatorname{argmin}_{\mathcal{P} \in \mathcal{U}_\rho^{\sigma_i}(\widehat{P}_{i,h,s,a_i}^{\pi-i})} \mathcal{P}V_{i,h+1}^{\pi,\sigma_i}, \quad (32d)$$

Additionally, we define square matrices $\widehat{\underline{P}}_{i,h}^{\pi,V} \in \mathbb{R}^{S \times S}$ as: $\widehat{\underline{P}}_{i,h}^{\pi,V} := \Pi_h^\pi \widehat{P}_{i,h}^{\pi-i,V}$.

Variance. We now introduce notations for variance corresponding to a specific probability distribution. For a probability vector $P \in \mathbb{R}^{1 \times S}$ and a vector $V \in \mathbb{R}^S$, we denote the variance of V with respect to P as $\operatorname{Var}_P(V)$, defined as:

$$\operatorname{Var}_P(V) := P(V \circ V) - (PV) \circ (PV), \quad (33)$$

Additionally, for a transition kernel $P^{\pi-i} \in \mathbb{R}^{SA_i \times S}$ and a vector $V \in \mathbb{R}^S$, we define $\operatorname{Var}_{P^{\pi-i}}(V) \in \mathbb{R}^{SA_i}$ as a vector of variances. The entry corresponding to (s, a_i) in $\operatorname{Var}_{P^{\pi-i}}(V)$ is given by:

$$\operatorname{Var}_{P^{\pi-i}}(s, a_i) := \operatorname{Var}_{P_{s,a_i}^{\pi-i}}(V), \quad (34)$$

where $P_{s,a_i}^{\pi-i}$ denotes the row of the transition matrix corresponding to state s and action a_i .

A.2 Preliminary facts about FTRL

Our proposed algorithm (see Algorithm 2) is inspired by online adversarial learning. In this section, we introduce the formulation of online learning and review key aspects of a widely-used algorithm, the Follow-the-Regularized-Leader (FTRL) algorithm.

Problem setting: online learning for weighted average loss. We consider an online learning problem over K steps, commonly found in adversarial learning settings (Lattimore and Szepesvári, 2020). The learner is presented with an action set \mathcal{A} , and loss functions $f_1, \dots, f_K : \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ are provided for each step. At each time step k , the learner selects a distribution over the action set, $\pi_k \in \Delta(\mathcal{A})$, and observes the loss function $f_k(\pi_k)$. The goal of the learner is to minimize the weighted average loss over the K steps, which is

defined as: $L^K = \sum_{k=1}^K \alpha_k^K f_k(\pi_k)$. To evaluate the learner's performance, the regret for the online learning process is defined as:

$$R^K = \sum_{k=1}^K \alpha_k^K f_k(\pi_k) - \left[\min_{\pi \in \Delta(\mathcal{A})} \sum_{k=1}^K \alpha_k^K f_k(\pi) \right]. \quad (35)$$

FTRL and its regret bound. A widely-used method for solving the online learning problem described above is the Follow-the-Regularized-Leader (FTRL) algorithm, introduced by [Shalev-Shwartz \(2007\)](#); [Shalev-Shwartz and Singer \(2007\)](#). At each step $k + 1$, the learner selects a soft-greedy action by solving:

$$\pi_{k+1} = \arg \min_{\pi \in \Delta(\mathcal{A})} \left[\sum_{i=1}^k \alpha_i^k f_i(\pi) + F_k(\pi) \right], \quad k = 1, 2, \dots, \quad (36)$$

where $F_k(\pi)$ represents a convex regularization function. The following theorem provides a refined regret bound for the FTRL algorithm when the loss function is linear with respect to the policy.

Theorem 3 (Theorem 3 in [Li et al. \(2022\)](#)). *For all $k \in [K]$ and policy π , the loss function is defined as $f_k(\pi) = \langle \pi_k, l_k \rangle$, where $l_k \in \mathbb{R}^{|\mathcal{A}|}$ represents a loss vector. The learner's choice π_{k+1} in episode $k + 1$ is updated according to the FTRL algorithm:*

$$\pi_{k+1}(a) = \arg \min_{\pi \in \Delta(\mathcal{A})} \{ \langle \pi, L_k \rangle + F_k(\pi) \} = \frac{\exp(-\eta_{k+1} L_k(a))}{\sum_{a' \in \mathcal{A}} \exp(-\eta_{k+1} L_k(a'))}, \quad \text{for all } a \in \mathcal{A}, \quad (37)$$

where the regularization function is given by $F_k(\pi) = \sum_{a \in \mathcal{A}} \frac{1}{\eta_{k+1}} \pi(a) \log(\pi(a))$. Suppose $0 < \alpha_1 \leq 1$ and $\eta_1 = \eta_2(1 - \alpha_1)$, and for all $k \geq 2$, assume $0 < \alpha_k < 1$ and $0 < \eta_{k+1}(1 - \alpha_k) \leq \eta_k$. Define:

$$\hat{\eta}_k := \begin{cases} \eta_2, & \text{if } k = 1, \\ \frac{\eta_k}{1 - \alpha_k}, & \text{if } k > 1. \end{cases} \quad (38)$$

Then, the regret of the FTRL algorithm is bounded by:

$$\begin{aligned} R_n &\leq \max_{a \in \mathcal{A}} \left[\sum_{k=1}^K \alpha_k^K \langle \pi_k, l_k \rangle - \sum_{k=1}^K \alpha_k^K l_k(a) \right] \\ &\leq \frac{5}{3} \sum_{k=1}^n \alpha_k^n \hat{\eta}_k \alpha_k \text{Var}_{\pi_k}(l_k) + \frac{\log A}{\eta_{n+1}} + 3 \sum_{k=1}^n \alpha_k^n \hat{\eta}_k^2 \alpha_k^2 \|l_k\|_\infty^3 \mathbb{I} \left(\hat{\eta}_k \alpha_k \|l_k\|_\infty > \frac{1}{3} \right). \end{aligned} \quad (39)$$

B Proof for Section 3

B.1 Proof of Theorem 1

Step 1: preliminaries. First, we introduce some useful definition and existing facts that are standard in real analysis and game theory literature.

Definition 2 (Upper semi-continuous). *A point-to-set mapping $x \in \mathcal{X} \mapsto \phi(x) \in \mathcal{Y}$ is upper semi-continuous if $\lim_{n \rightarrow \infty} x^n = x_0, y^n \in \phi(x^n), \lim_{n \rightarrow \infty} y^n = y_0$ imply that $y_0 \in \phi(x_0)$.*

Theorem 4 (Kakutani's fixed point Theorem ([Kakutani, 1941](#))). *If X is a closed, bounded, and convex set in a Euclidean space, and ϕ is a upper semi-continuous correspondence mapping X into the family of all closed convex subsets of X , then there exists $x \in X$ so that $x \in \phi(x)$.*

Step 2: constructing an auxiliary single-step game. Focusing on finite-horizon RMG $\mathcal{MG}_{\text{rob}} = \{\mathcal{S}, \{\mathcal{A}_i\}_{1 \leq i \leq n}, \{\mathcal{U}_\rho^{\sigma_i}(P^0)\}_{1 \leq i \leq n}, r, H\}$, we shall verify the theorem by firstly consider a one-step game and then apply the results recursively to the sequential Markov games.

Without loss of generality, we focus on any of the steps $h \in [H]$ and construct an auxiliary one-step game. Towards this, we first introduce a fixed value function $V_{i,h+1} \in \mathbb{R}^S$ with $0 \leq V_{i,h+1} \leq H$ for the i -th agent, representing the possible value function obtained at the next time step $h+1$. Focusing on time step h , for any joint product policy $\pi : \mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i)$, we abuse the notation defined in (14) to denote the expected nominal transition kernel over each (s, a_i) as:

$$P_{h,s,a_i}^{\pi_{-i}} = \mathbb{E}_{\pi(\mathbf{a}_{-i} | s, a_i)} \left[P_{h,s,(a_i, \mathbf{a}_{-i})}^0 \right] = \mathbb{E}_{\pi_{-i}(\mathbf{a}_{-i} | s)} \left[P_{h,s,(a_i, \mathbf{a}_{-i})}^0 \right]. \quad (40)$$

Armed with this, for any joint product policy $\pi : \mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i)$, we can define the payoffs to maximize for the players as below:

$$\forall s \in \mathcal{S} : \quad f_{i,s}(\pi_i(s), \pi_{-i}(s); V_{i,h+1}) = \mathbb{E}_{\mathbf{a} \sim \pi(s)} [r_{i,h}(s, \mathbf{a})] + \mathbb{E}_{a_i \sim \pi_i(s)} \left[\inf_{\mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} P V_{i,h+1} \right], \quad (41)$$

which is defined analogous to the robust Bellman equation (cf. (13)) by replacing a real robust value function vector (associated with some policy) to some fixed vector $V_{i,h+1}$.

Now we are ready to introduce the following useful mapping: for any $\pi : \mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i)$,

$$\phi(\pi) := \left\{ u \mid u_i(s) \in \operatorname{argmax}_{\pi'_i(s) \in \Delta(\mathcal{A}_i)} f_{i,s}(\pi'_i(s), \pi_{-i}(s); V_{i,h+1}), \forall (i, s) \in [n] \times \mathcal{S} \right\}. \quad (42)$$

Step 3: the existence of NE of the auxiliary game. To apply Theorem 4, there are three required conditions. First, we know that the space of product policy is $X = \{\pi : \mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i)\}$ is a closed, bounded and convex set in Euclidean space.

- **Verifying that $\phi(\pi)$ is an upper semi-continuous correspondence.** Before starting, we introduce the following two useful lemmas with the proof postponed to Appendix B.2.2 and B.2.3.

Lemma 2. *The set of function $\{f_{i,s}(\pi'_i(s), \pi_{-i}(s); V_{i,h+1}), 0 \leq V_{i,h+1} \leq H\}$ is equicontinuous with respect to $\pi'_i(s), \pi_{-i}(s)$ for all $(i, s) \in [n] \times \mathcal{S}$.*

Lemma 3. *For any $i \in [n]$ and then $x_{-i} : \mathcal{S} \mapsto \prod_{j \neq i, j \in [n]} \Delta(\mathcal{A}_j)$, the functions*

$$\forall s \in \mathcal{S} : \quad g_{i,s}(x_{-i}(s), V_{i,h+1}) := \max_{\pi'_i \in \Delta(\mathcal{S})} f_{i,s}(\pi'_i(s), x_{-i}(s); V_{i,h+1}) \quad (43)$$

are continuous with respect to $x_{-i}(s)$ and the set $\{g_{i,s}(\cdot, V) \mid V \in \mathbb{R}^S, 0 \leq V \leq H\}$ is equicontinuous.

Armed with above lemmas, we are in the position to prove this condition. We suppose there are two sequence $\lim_{n \rightarrow \infty} x^n = x^0, y^n \in \phi(x^n), \lim_{n \rightarrow \infty} y^n = y^0$. Recall the definition of a upper semi-continuous correspondence (cf. Definition 2), we are supposed to show that $y^0 \in \phi(x^0)$, i.e.,

$$\forall (i, s) \in [n] \times \mathcal{S} : \quad f_{i,s}(y_i^0(s), x_{-i}^0(s); V_{i,h+1}) = \max_{\pi'_i \in \Delta(\mathcal{S})} f_{i,s}(\pi'_i(s), x_{-i}^0(s); V_{i,h+1}). \quad (44)$$

Towards this, we have

$$\begin{aligned} & |f_{i,s}(y_i^0(s), x_{-i}^0(s); V_{i,h+1}) - g_{i,s}(x_{-i}^0(s), V_{i,h+1})| \\ & \leq |f_{i,s}(y_i^0(s), x_{-i}^0(s); V_{i,h+1}) - f_{i,s}(y_i^n(s), x_{-i}^n(s); V_{i,h+1})| \\ & \quad + |f_{i,s}(y_i^n(s), x_{-i}^n(s); V_{i,h+1}) - g_{i,s}(x_{-i}^n(s), V_{i,h+1})| \\ & \stackrel{(i)}{=} |f_{i,s}(y_i^0(s), x_{-i}^0(s); V_{i,h+1}) - f_{i,s}(y_i^n(s), x_{-i}^n(s); V_{i,h+1})| + |g_{i,s}(x_{-i}^n(s), V_{i,h+1}) - g_{i,s}(x_{-i}^0(s), V_{i,h+1})| \\ & \rightarrow 0 \quad \text{as } n \rightarrow \infty, \end{aligned} \quad (45)$$

where the first inequality follows from the triangle inequality, (i) holds by the assumption $y^n \in \phi(x^n)$ so that $f_{i,s}(y_i^n(s), x_{-i}^n(s); V_{i,h+1}) = \max_{\pi'_i \in \Delta(\mathcal{S})} f_{i,s}(\pi'_i(s), x_{-i}^n(s); V_{i,h+1})$, and the last line can be verified by the continuity implied by Lemma 2 and Lemma 3.

- **Verifying $\phi(\pi)$ is convex for any $\pi \in X$.** Finally, we gonna work on the convexity of $\phi(\pi)$ for any $\pi \in X$. To begin with, by the definition of $\phi(\pi)$ in (42), we know that $\phi(\pi) \subseteq X$ and the maximum of the continuous function $f_{i,s}(\pi_i(s), \pi_{-i}(s); V_{i,h+1})$ (cf. Lemma 2) on a compact set exists, i.e., $\phi(x) \neq \emptyset$. Suppose there exists two Nash equilibrium $z : \mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i)$, $v : \mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i)$ and $z, v \in \phi(\pi)$. Then we have that for any $(i, s) \in [n] \times \mathcal{S}$,

$$f_{i,s}(z_i(s), \pi_{-i}(s); V_{i,h+1}) = f_{i,s}(v_i(s), \pi_{-i}(s); V_{i,h+1}) = \max_{u_i(s) \in \Delta(\mathcal{A}_i)} f_{i,s}(u_i(s), \pi_{-i}(s); V_{i,h+1}). \quad (46)$$

To continue, for any $0 \leq \lambda \leq 1$, one has

$$\begin{aligned} & \max_{u_i(s) \in \Delta(\mathcal{A}_i)} f_{i,s}(u_i(s), \pi_{-i}(s); V_{i,h+1}) = \lambda f_{i,s}(z_i(s), \pi_{-i}(s); V_{i,h+1}) + (1-\lambda) f_{i,s}(v_i(s), \pi_{-i}(s); V_{i,h+1}) \\ & = \lambda \left(\mathbb{E}_{a_i \sim z_i(s)} \left[r_{i,h}^{\pi_{-i}}(s, a_i) \right] + \mathbb{E}_{a_i \sim z_i(s)} \left[\inf_{U^{\sigma_i} \left(P_{h,s,a_i}^{\pi_{-i}} \right)} PV_{i,h+1} \right] \right) \\ & \quad + (1-\lambda) \left(\mathbb{E}_{a_i \sim v_i(s)} \left[r_{i,h}^{\pi_{-i}}(s, a_i) \right] + \mathbb{E}_{a_i \sim v_i(s)} \left[\inf_{U^{\sigma_i} \left(P_{h,s,a_i}^{\pi_{-i}} \right)} PV_{i,h+1} \right] \right) \end{aligned} \quad (47)$$

$$\begin{aligned} & = \mathbb{E}_{a_i \sim [\lambda z_i(s) + (1-\lambda)v_i(s)]} \left[r_{i,h}^{\pi_{-i}}(s, a_i) \right] + \mathbb{E}_{a_i \sim [\lambda z_i(s) + (1-\lambda)v_i(s)]} \left[\inf_{U^{\sigma_i} \left(P_{h,s,a_i}^{\pi_{-i}} \right)} PV_{i,h+1} \right] \\ & = f_{i,s}(\lambda z_i(s) + (1-\lambda)v_i(s), \pi_{-i}(s); V_{i,h+1}). \end{aligned} \quad (48)$$

where we denote $r_{i,h}^{\pi_{-i}}(s, a_i) := \mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i}(s)} [r_{i,h}(s, (a_i, \mathbf{a}_{-i}))]$. Hence, we show that $\lambda z_i(s) + (1-\lambda)v_i(s) \in \phi(\pi)$ for all $(i, s) \in [n] \times \mathcal{S}$ and $0 \leq \lambda \leq 1$, thus verify that $\phi(\pi)$ is convex for any $\pi \in X$.

Step 4: the existence of robust NE in RMGs. Armed with above results, now we consider a general form to show that there exists a policy $\pi : [H] \times \mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i)$ that satisfies

$$\forall (i, h, s) \in [n] \times [H] \times \mathcal{S} : \quad V_{i,h}^{\pi, \sigma_i}(s) = V_{i,h}^{\star, \pi_{-i}, \sigma_i}(s). \quad (49)$$

We shall prove this by induction.

- **The base case.** Starting with the final step $h = H$, we recall that by definition,

$$\forall (i, s) \in [n] \times \mathcal{S} : \quad V_{i,H+1}^{\pi, \sigma_i}(s) = 0. \quad (50)$$

To apply the results in the one-step game constructed in Step 2, we consider the one-step game at $h = H$ and using the payoff function (cf. (41))

$$\forall s \in \mathcal{S} : \quad f_{i,s}(\pi_i(s), \pi_{-i}(s); V_{i,H+1}^{\pi, \sigma_i}) = \mathbb{E}_{\mathbf{a} \sim \pi(s)} [r_{i,h}(s, \mathbf{a})]. \quad (51)$$

We know that there exists a policy π so that

$$\forall (i, s) \in [n] \times \mathcal{S} : \quad V_{i,H}^{\pi, \sigma_i}(s) = V_{i,H}^{\star, \pi_{-i}, \sigma_i}(s) \quad (52)$$

by setting π_H as the NE of the one-step game.

- **Induction.** Assuming that there exists a policy π so that for subsequent steps $h+1, \dots, H$,

$$\forall (i, h, s) \in [n] \times \{h+1, \dots, H\} \times \mathcal{S} : \quad V_{i,h}^{\pi, \sigma_i}(s) = V_{i,h}^{\star, \pi_{-i}, \sigma_i}(s), \quad (53)$$

which are achieved by determining certain policies for $\{\pi_{h+1}, \pi_{h+2}, \dots, \pi_H\}$. We are supposed to prove that at time step h , we can ensure our policy π satisfying

$$\forall (i, s) \in [n] \times \mathcal{S} : \quad V_{i,h}^{\pi, \sigma_i}(s) = V_{i,h}^{\star, \pi_{-i}, \sigma_i}(s) \quad (54)$$

by choosing a proper policy π_h at the time step h .

Towards this, it is observed that

$$\begin{aligned}
V_{i,h}^{*,\pi_{-i},\sigma_i}(s) &= \max_{\pi'_i: \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)} V_{i,h}^{\pi'_i \times \pi_{-i}, \sigma_i}(s) \\
&= \max_{\pi'_i: \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)} \mathbb{E}_{\mathbf{a} \sim \pi'_{i,h}(s) \times \pi_{-i,h}(s)} [r_{i,h}(s, \mathbf{a})] + \mathbb{E}_{a_i \sim \pi'_{i,h}(s)} \left[\inf_{P \in \mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} PV_{i,h+1}^{\pi'_i \times \pi_{-i}, \sigma_i} \right] \\
&= \max_{\pi'_{i,h}(s) \in \Delta(\mathcal{A}_i)} \mathbb{E}_{\mathbf{a} \sim \pi'_{i,h}(s) \times \pi_{-i,h}(s)} [r_{i,h}(s, \mathbf{a})] \\
&\quad + \max_{\pi'_{i,h}(s) \in \Delta(\mathcal{A}_i)} \mathbb{E}_{a_i \sim \pi'_{i,h}(s)} \max_{\pi'_{i,h^+}: \mathcal{S} \times h^+ \mapsto \Delta(\mathcal{A}_i)} \left[\inf_{P \in \mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} PV_{i,h+1}^{\pi'_i \times \pi_{-i}, \sigma_i} \right] \\
&= \max_{\pi'_{i,h}(s) \in \Delta(\mathcal{A}_i)} \left[\mathbb{E}_{\mathbf{a} \sim \pi'_{i,h}(s) \times \pi_{-i,h}(s)} [r_{i,h}(s, \mathbf{a})] + \mathbb{E}_{a_i \sim \pi'_{i,h}(s)} \left[\inf_{P \in \mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} PV_{i,h+1}^{*,\pi_{-i},\sigma_i} \right] \right]. \tag{55}
\end{aligned}$$

where we denote $h^+ = \{h+1, h+2, \dots, H\}$ as the set that includes all the time steps after h until the end of the episode, and the last equality follows from the fact

$$\begin{aligned}
\max_{\pi'_{i,h^+}: \mathcal{S} \times h^+ \mapsto \Delta(\mathcal{A}_i)} \left[\inf_{\mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} PV_{i,h+1}^{\pi'_i \times \pi_{-i}, \sigma_i} \right] &= \inf_{\mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} P \max_{\pi'_{i,h^+}: \mathcal{S} \times h^+ \mapsto \Delta(\mathcal{A}_i)} V_{i,h+1}^{\pi'_i \times \pi_{-i}, \sigma_i} \\
&= \inf_{\mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} PV_{i,h+1}^{*,\pi_{-i},\sigma_i}, \tag{56}
\end{aligned}$$

which holds by the definition of $V_{i,h+1}^{*,\pi_{-i},\sigma_i}$. Now invoking the results in the auxiliary one-step game with $V_{i,h+1} = V_{i,h+1}^{*,\pi_{-i},\sigma_i}$, one has that there exists a policy with π_h that satisfies

$$\forall (i, s) \in [n] \times \mathcal{S}: \quad V_{i,h}^{\pi, \sigma_i}(s) = V_{i,h}^{*,\pi_{-i},\sigma_i}(s). \tag{57}$$

Combining the results in the base case and induction, we complete the proof by recursively choosing $\pi_h: \mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i)$ for $h = H, H-1, \dots, 1$ as the NE of the corresponding one-step game at time step h and arrive at

$$\forall (i, s) \in [n] \times \mathcal{S}: \quad V_{i,1}^{\pi, \sigma_i}(s) = V_{i,1}^{*,\pi_{-i},\sigma_i}(s). \tag{58}$$

B.2 Proof of auxiliary facts

B.2.1 Proof of Lemma 1

The proof is obtained by recursively showing that for each (h, s) , there exist a policy. Then the product policy of them will be that final policy

Without loss of generality, we consider any $i \in [n]$ with the other agents' policy $\pi_i: \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)$ fixed. We shall prove this lemma by induction.

- **The base case.** Consider the base case $h = H$. Conditioned on the other agents' policy $\pi_i: \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)$, the maximum of the robust value function of the i -th agent can be expressed by

$$\begin{aligned}
\forall s \in \mathcal{S}: \quad V_{i,H}^{*,\pi_{-i},\sigma_i}(s) &= \max_{\pi'_i: \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)} V_{i,H}^{\pi'_i \times \pi_{-i}, \sigma_i}(s) \\
&= \max_{\pi'_i: \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)} \mathbb{E}_{a_i \sim \pi'_{i,H}(s)} \left[\mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,H}(s)} [r_{i,H}(s, \mathbf{a})] \right] \\
&= \max_{\pi'_{i,H}(s) \sim \Delta(\mathcal{A}_i)} \mathbb{E}_{a_i \sim \pi'_{i,H}(s)} \left[\mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,H}(s)} [r_{i,H}(s, \mathbf{a})] \right]. \tag{59}
\end{aligned}$$

Since the maximum of the continuous function $\mathbb{E}_{\mathbf{a}_i \sim \pi'_{i,H}(s)} [\mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,H}(s)} [r_{i,H}(s, \mathbf{a})]]$ on a compact set $\Delta(\mathcal{A}_i)$ exists, by setting

$$\forall s \in \mathcal{S} : \quad \tilde{\pi}_{i,H}(s) = \operatorname{argmax}_{\pi'_{i,H}(s) \sim \Delta(\mathcal{A}_i)} \mathbb{E}_{\mathbf{a}_i \sim \pi'_{i,H}(s)} [\mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,H}(s)} [r_{i,H}(s, \mathbf{a})]], \quad (60)$$

we arrive at

$$\forall s \in \mathcal{S} : \quad V_{i,H}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i}(s) = V_{i,H}^{\star, \pi_{-i}, \sigma_i}(s). \quad (61)$$

This complete the proof for the base case.

- **Induction.** Assuming that for $t = h + 1, h + 2, \dots, H$, we have

$$\forall s \in \mathcal{S} : \quad V_{i,t}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i}(s) = V_{i,t}^{\star, \pi_{-i}, \sigma_i}(s). \quad (62)$$

Then, we want to prove for the step h , where the maximum of the robust value function of the i -th agent can be expressed as: for all $s \in \mathcal{S}$,

$$\begin{aligned} & V_{i,h}^{\star, \pi_{-i}, \sigma_i}(s) \\ &= \max_{\pi'_{i,h} : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)} V_{i,h}^{\pi'_{i,h} \times \pi_{-i}, \sigma_i}(s) \\ &= \max_{\pi'_{i,h} : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)} \mathbb{E}_{\mathbf{a}_i \sim \pi'_{i,h}(s)} [\mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,h}(s)} [r_{i,h}(s, \mathbf{a})]] + \mathbb{E}_{\mathbf{a}_i \sim \pi_{i,h}(s)} \left[\inf_{\mathcal{U}_\rho^{\sigma_i}(P_{h,s,\mathbf{a}_i}^{\pi_{-i}})} PV_{i,h+1}^{\star, \pi_{-i}, \sigma_i} \right] \\ &\stackrel{(i)}{=} \max_{\pi'_{i,h} : \mathcal{S} \times [H] \mapsto \Delta(\mathcal{A}_i)} \mathbb{E}_{\mathbf{a}_i \sim \pi'_{i,h}(s)} [\mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,h}(s)} [r_{i,h}(s, \mathbf{a})]] + \mathbb{E}_{\mathbf{a}_i \sim \pi_{i,h}(s)} \left[\inf_{\mathcal{U}_\rho^{\sigma_i}(P_{h,s,\mathbf{a}_i}^{\pi_{-i}})} PV_{i,h+1}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i} \right] \\ &= \max_{\pi'_{i,h}(s) \sim \Delta(\mathcal{A}_i)} \mathbb{E}_{\mathbf{a}_i \sim \pi'_{i,h}(s)} [\mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,h}(s)} [r_{i,h}(s, \mathbf{a})]] + \mathbb{E}_{\mathbf{a}_i \sim \pi_{i,h}(s)} \left[\inf_{\mathcal{U}_\rho^{\sigma_i}(P_{h,s,\mathbf{a}_i}^{\pi_{-i}})} PV_{i,h+1}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i} \right]. \quad (63) \end{aligned}$$

where (i) holds by the induction assumption in (62). Similarly to the base case, the maximum of the continuous function $\mathbb{E}_{\mathbf{a}_i \sim \pi'_{i,h}(s)} [\mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,h}(s)} [r_{i,h}(s, \mathbf{a})]] + \mathbb{E}_{\mathbf{a}_i \sim \pi_{i,h}(s)} \left[\inf_{\mathcal{U}_\rho^{\sigma_i}(P_{h,s,\mathbf{a}_i}^{\pi_{-i}})} PV_{i,h+1}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i} \right]$ on a compact set $\Delta(\mathcal{A}_i)$ exists. So without conflict, for all $s \in \mathcal{S}$, we can set

$$\begin{aligned} & \tilde{\pi}_{i,h}(s) \\ &= \operatorname{argmax}_{\pi'_{i,h}(s) \sim \Delta(\mathcal{A}_i)} \mathbb{E}_{\mathbf{a}_i \sim \pi'_{i,h}(s)} [\mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,h}(s)} [r_{i,h}(s, \mathbf{a})]] + \mathbb{E}_{\mathbf{a}_i \sim \pi_{i,h}(s)} \left[\inf_{\mathcal{U}_\rho^{\sigma_i}(P_{h,s,\mathbf{a}_i}^{\pi_{-i}})} PV_{i,h+1}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i} \right], \quad (64) \end{aligned}$$

since the function $\inf_{\mathcal{U}_\rho^{\sigma_i}(P_{h,s,\mathbf{a}_i}^{\pi_{-i}})} PV_{i,h+1}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i}$ and especially $V_{i,h+1}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i}$ are independent from the policy in the first h steps $(\{\tilde{\pi}_{i,t}(s)\}_{s \in \mathcal{S}, t \in [h]})$.

Consequently, (64) directly implies that

$$\forall s \in \mathcal{S} : \quad V_{i,h}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i}(s) = V_{i,h}^{\star, \pi_{-i}, \sigma_i}(s). \quad (65)$$

Combining the results in base case and the induction, we complete the proof by showing that

$$\forall (h, s) \in [H] \times \mathcal{S} : \quad V_{i,h}^{\tilde{\pi}_i \times \pi_{-i}, \sigma_i}(s) = V_{i,h}^{\star, \pi_{-i}, \sigma_i}(s). \quad (66)$$

B.2.2 Proof of Lemma 2

First, we define the distance between any two policy $\pi, \pi' \in X = \{\pi : \mathcal{S} \mapsto \prod_{i \in [n]} \Delta(A_i)\}$ as below:

$$d(\pi, \pi') := \max_{i \in [n]} \max_{(s, a_i) \in \mathcal{S} \times \mathcal{A}_i} |\pi_i(a_i | s) - \pi'_i(a_i | s)|. \quad (67)$$

To prove the continuity, given any $\epsilon > 0$, we want to show that there exists $\delta(\epsilon) > 0$ such that if

$$d(\pi, \pi') < \delta(\epsilon), \quad (68)$$

then

$$|f_{i,s}(\pi_i(s), \pi_{-i}(s); V_{i,h+1}) - f_{i,s}(\pi'_i(s), \pi'_{-i}(s); V_{i,h+1})| < \epsilon \quad (69)$$

for any fixed $\{V_{i,h+1}\}_{i \in [n]}$ with $0 \leq V_{i,h+1} \leq H$ for all $i \in [n]$. Towards this, we observe that

$$\begin{aligned} & |f_{i,s}(\pi_i(s), \pi_{-i}(s); V_{i,h+1}) - f_{i,s}(\pi'_i(s), \pi'_{-i}(s); V_{i,h+1})| \\ &= \left| \mathbb{E}_{\mathbf{a} \sim \pi(s)} [r_{i,h}(s, \mathbf{a})] + \mathbb{E}_{a_i \sim \pi_i(s)} \left[\inf_{\mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} PV_{i,h+1} \right] \right. \\ & \quad \left. - \mathbb{E}_{\mathbf{a} \sim \pi'(s)} [r_{i,h}(s, \mathbf{a})] + \mathbb{E}_{a_i \sim \pi'_i(s)} \left[\inf_{\mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi'_{-i}})} PV_{i,h+1} \right] \right| \\ &\leq \left| \mathbb{E}_{\mathbf{a} \sim \pi(s)} [r_{i,h}(s, \mathbf{a})] - \mathbb{E}_{\mathbf{a} \sim \pi'(s)} [r_{i,h}(s, \mathbf{a})] \right| \\ & \quad + \left| \mathbb{E}_{a_i \sim \pi_i(s)} \left[\inf_{\mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} PV_{i,h+1} \right] - \mathbb{E}_{a_i \sim \pi'_i(s)} \left[\inf_{\mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi'_{-i}})} PV_{i,h+1} \right] \right|. \end{aligned} \quad (70)$$

The first term can be bounded by

$$\begin{aligned} & \left| \mathbb{E}_{\mathbf{a} \sim \pi(s)} [r_{i,h}(s, \mathbf{a})] - \mathbb{E}_{\mathbf{a} \sim \pi'(s)} [r_{i,h}(s, \mathbf{a})] \right| \\ &\leq \sum_{\mathbf{a} \in \mathcal{A}} \left| \prod_{i \in [n]} \pi_i(a_i | s) - \prod_{i \in [n]} \pi'_i(a_i | s) \right| \max_{(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}} r_{i,h}(s, \mathbf{a}) \\ &\leq \sum_{\mathbf{a} \in \mathcal{A}} \left| \prod_{i \in [n]} \pi_i(a_i | s) - \prod_{i \in [n]} \pi'_i(a_i | s) \right|, \end{aligned} \quad (71)$$

where the last inequality holds by the definition of reward function $\max_{(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}} r_{i,h}(s, \mathbf{a}) \leq 1$ for all $(i, h) \in [n] \times [H]$. To continue, we first define the difference between $\delta_i(s, a_i) := \pi'_i(a_i | s) - \pi_i(a_i | s)$. Therefore, we have

$$\begin{aligned} & \left| \prod_{i \in [n]} \pi_i(a_i | s) - \prod_{i \in [n]} \pi'_i(a_i | s) \right| = \left| \prod_{i \in [n]} \pi_i(a_i | s) - \prod_{i \in [n]} (\pi_i(a_i | s) + \delta_i(s, a_i)) \right| \\ &= \left| \sum_{|\mathcal{Y}| \geq 1, \mathcal{Y} \subseteq [n]} \left(\prod_{i \in \mathcal{Y}} \delta_i(s, a_i) \right) \cdot \left(\prod_{i \in \mathcal{Y}^c} \pi_i(a_i | s) \right) \right| \\ &\leq \sum_{|\mathcal{Y}| \geq 1, \mathcal{Y} \subseteq [n]} \left| \left(\prod_{i \in \mathcal{Y}} \delta_i(s, a_i) \right) \cdot \left(\prod_{i \in \mathcal{Y}^c} \pi_i(a_i | s) \right) \right| \leq (2^n - 1) \delta(\epsilon), \end{aligned} \quad (72)$$

where the last inequality holds by (68). Plugging (72) back to (71) indicates that

$$\left| \mathbb{E}_{\mathbf{a} \sim \pi(s)} [r_{i,h}(s, \mathbf{a})] - \mathbb{E}_{\mathbf{a} \sim \pi'(s)} [r_{i,h}(s, \mathbf{a})] \right| \leq \prod_{i \in [n]} A_i (2^n - 1) \delta(\epsilon). \quad (73)$$

For the second term in (70), we observe that

$$\begin{aligned}
& \left| \mathbb{E}_{a_i \sim \pi_i(s)} \left[\inf_{P \in \mathcal{U}^{\sigma_i} \left(P_{h,s,a_i}^{\pi_{-i}} \right)} PV_{i,h+1} \right] - \mathbb{E}_{a_i \sim \pi'_i(s)} \left[\inf_{P \in \mathcal{U}^{\sigma_i} \left(P_{h,s,a_i}^{\pi'_{-i}} \right)} PV_{i,h+1} \right] \right| \\
& \leq \left| \mathbb{E}_{a_i \sim \pi_i(s)} \left[\inf_{P \in \mathcal{U}^{\sigma_i} \left(P_{h,s,a_i}^{\pi_{-i}} \right)} PV_{i,h+1} \right] - \mathbb{E}_{a_i \sim \pi_i(s)} \left[\inf_{P \in \mathcal{U}^{\sigma_i} \left(P_{h,s,a_i}^{\pi'_{-i}} \right)} PV_{i,h+1} \right] \right| \\
& \quad + \left| \mathbb{E}_{a_i \sim \pi_i(s)} \left[\inf_{P \in \mathcal{U}^{\sigma_i} \left(P_{h,s,a_i}^{\pi'_{-i}} \right)} PV_{i,h+1} \right] - \mathbb{E}_{a_i \sim \pi'_i(s)} \left[\inf_{P \in \mathcal{U}^{\sigma_i} \left(P_{h,s,a_i}^{\pi'_{-i}} \right)} PV_{i,h+1} \right] \right| \\
& \stackrel{(i)}{\leq} \mathbb{E}_{a_i \sim \pi_i(s)} \left[\max_{\alpha \in [\min_s V_{i,h+1}(s), \max_s V_{i,h+1}(s)]} \left| \mathbb{E}_{\pi_{-i}(a_{-i} | s)} \left[P_{h,s,(a_i, a_{-i})}^0 \right] [V_{i,h+1}]_{\alpha} \right. \right. \\
& \quad \left. \left. - \mathbb{E}_{\pi'_{-i}(a_{-i} | s)} \left[P_{h,s,(a_i, a_{-i})}^0 \right] [V_{i,h+1}]_{\alpha} \right| \right] + \sum_{a_i \in \mathcal{A}_i} |\pi'_i(a_i | s) - \pi_i(a_i | s)| \inf_{P \in \mathcal{U}^{\sigma_i} \left(P_{h,s,a_i}^{\pi'_{-i}} \right)} PV_{i,h+1} \\
& \stackrel{(ii)}{\leq} \sum_{a_{-i} \in \mathcal{A}_i} \left| \prod_{j \neq i} \pi_j(a_j | s) - \prod_{j \neq i} \pi'_j(a_j | s) \right| H + HA_i \Delta(\epsilon) \\
& \stackrel{(iii)}{\leq} H \prod_{j \neq i, j \in [n]} A_j (2^{n-1} - 1) \delta(\epsilon) + HA_i \Delta(\epsilon) \leq 2H \prod_{i \in [n]} A_i (2^n - 1) \cdot \delta(\epsilon), \tag{74}
\end{aligned}$$

where the first inequality holds by the triangle inequality, and (i) follows from applying the dual form of TV distance

$$\inf_{P \in \mathcal{U}^{\sigma_i}(P)} \mathcal{P}V = \max_{\alpha \in [\min_s V(s), \max_s V(s)]} \left\{ P[V]_{\alpha} - \sigma_i \left(\alpha - \min_{s'} [V]_{\alpha}(s') \right) \right\}, \tag{75}$$

and the maximum operator is 1-Lipschitz, (ii) arises from the fact that $\|V_{i,h+1}\|_{\infty} \leq H$, and (iii) can be verified by following the same pipeline of (72). Combining (73) and (74), one has

$$\left| f_{i,s}(\pi_i(s), \pi_{-i}(s); V_{i,h+1}) - f_{i,s}(\pi'_i(s), \pi'_{-i}(s); V_{i,h+1}) \right| \leq 3H \prod_{i \in [n]} A_i (2^n - 1) \cdot \delta(\epsilon). \tag{76}$$

Consequently, letting $\delta_1(\epsilon) = \frac{\min\{\epsilon, 1\}}{3H \prod_{i \in [n]} A_i (2^n - 1)}$, we have when $d(\pi, \pi') < \delta_1(\epsilon)$,

$$\left| f_{i,s}(\pi_i(s), \pi_{-i}(s); V_{i,h+1}) - f_{i,s}(\pi'_i(s), \pi'_{-i}(s); V_{i,h+1}) \right| < \epsilon.$$

B.2.3 Proof of Lemma 3

Without loss of generality, we consider any $i \in [n]$. Consider $x_{-i} : \mathcal{S} \mapsto \prod_{j \neq i, j \in [n]} \Delta(\mathcal{A}_j)$ and $y_{-i} : \mathcal{S} \mapsto \prod_{j \neq i, j \in [n]} \Delta(\mathcal{A}_j)$. Before continuing, for all $s \in \mathcal{S}$, we denote

$$\begin{aligned}
u_{i,s}^* &:= \operatorname{argmax}_{\pi'_i \in \Delta(\mathcal{S})} f_{i,s}(\pi'_i(s), x_{-i}(s); V_{i,h+1}), \\
v_{i,s}^* &:= \operatorname{argmax}_{\pi'_i \in \Delta(\mathcal{S})} f_{i,s}(\pi'_i(s), y_{-i}(s); V_{i,h+1}).
\end{aligned} \tag{77}$$

Then we have for any $s \in \mathcal{S}$,

$$\begin{aligned}
& g_{i,s}(x_{-i}(s), V_{i,h+1}) - g_{i,s}(y_{-i}(s), V_{i,h+1}) \\
& = \max_{\pi'_i \in \Delta(\mathcal{S})} f_{i,s}(\pi'_i(s), x_{-i}(s); V_{i,h+1}) - \max_{\pi'_i \in \Delta(\mathcal{S})} f_{i,s}(\pi'_i(s), y_{-i}(s); V_{i,h+1}) \\
& = f_{i,s}(u_{i,s}^*, x_{-i}(s); V_{i,h+1}) - f_{i,s}(v_{i,s}^*, y_{-i}(s); V_{i,h+1}) \\
& \leq f_{i,s}(u_{i,s}^*, x_{-i}(s); V_{i,h+1}) - f_{i,s}(u_{i,s}^*, y_{-i}(s); V_{i,h+1}) \rightarrow 0 \quad \text{as } y_{-i}(s) \rightarrow x_{-i}(s), \tag{78}
\end{aligned}$$

where the last line holds by Lemma (2) which shows that the function $f_{i,s}$ is continuous. Similarly, one has

$$\begin{aligned} & g_{i,s}(x_{-i}(s), V_{i,h+1}) - g_{i,s}(y_{-i}(s), V_{i,h+1}) \\ & \geq f_{i,s}(v_{i,s}^*, x_{-i}(s); V_{i,h+1}) - f_{i,s}(u_{i,s}^*, y_{-i}(s); V_{i,h+1}) \rightarrow 0 \quad \text{as } y_{-i}(s) \rightarrow x_{-i}(s). \end{aligned} \quad (79)$$

We complete the proof by showing that

$$|g_{i,s}(x_{-i}(s), V_{i,h+1}) - g_{i,s}(y_{-i}(s), V_{i,h+1})| \rightarrow 0 \quad \text{as } y_{-i}(s) \rightarrow x_{-i}(s). \quad (80)$$

C Proof of Theorem 2

We will present the proof of Theorem 2 by first outlining the proof structure, followed by a step-by-step explanation of the key components. Auxiliary proofs will be provided at the end of this section.

C.1 Proof pipeline

To proof Theorem 2, recall the goal is to show that

$$\forall (i, s) \in [n] \times \mathcal{S} : \quad \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,1}^{*,\pi_{-i},\sigma_i}(s)] - \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,1}^{\pi,\sigma_i}(s)] \leq \varepsilon, \quad (81)$$

where $\widehat{\xi} = \{\widehat{\xi}_h\}_{h \in [H]}$ is the output distribution over the set of policies $\{\pi_h^k = (\pi_{1,h}^k \times \cdots \times \pi_{n,h}^k)\}_{k \in [K], h \in [H]}$ from Algorithm 2. Namely, $\pi \sim \widehat{\xi}$ means

$$\forall h \in [H] : \quad \pi_h \sim \widehat{\xi}_h, \quad \text{where} \quad \widehat{\xi}_h(\pi_h^k) = \alpha_k^K. \quad (82)$$

We first introduce the best-response policy for player i :

$$\widehat{\pi}_i^* = [\widehat{\pi}_{i,h}^*]_{h \in [H]} := \arg \max_{\pi'_i : \mathcal{S} \times [H] \rightarrow \Delta(\mathcal{A}_i)} \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,1}^{\pi'_i, \pi_{-i}}].$$

Recall that value function $\mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,h}^{\pi,\sigma_i}]$ satisfies the following Bellman equation for all $(i, s, h) \in [n] \times \mathcal{S} \times [H]$:

$$\begin{aligned} \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,H+1}^{\pi,\sigma_i}(s)] &= 0, \\ \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,h}^{\pi,\sigma_i}(s)] &= \mathbb{E}_{\pi \sim \widehat{\xi}} \left\{ \sum_{\mathbf{a} \in \mathcal{A}} \pi_h(\mathbf{a} | s) r_{i,h}(s, \mathbf{a}) + \mathbb{E}_{a_i \sim \pi_{i,h}} \left[\inf_{\mathcal{P} \in \mathcal{U}_i^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}})} \mathcal{P} \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,h+1}^{\pi,\sigma_i}] \right] \right\}, \\ &= \sum_{k=1}^K \sum_{\mathbf{a} \in \mathcal{A}} \alpha_k^K \pi_h^k(\mathbf{a} | s) r_{i,h}(s, \mathbf{a}) + \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k} \left[\inf_{\mathcal{P} \in \mathcal{U}_i^{\sigma_i}(P_{h,s,a_i}^{\pi_{-i}^k})} \mathcal{P} \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,h+1}^{\pi,\sigma_i}] \right], \end{aligned}$$

where $P_{h,s,a_i}^{\pi_{-i}^k}$ is defined as:

$$P_{h,s,a_i}^{\pi_{-i}^k} = \mathbb{E}_{\mathbf{a}_{-i} \sim \pi_{-i,h}^k(\cdot | s)} [P_{h,s,(a_i, \mathbf{a}_{-i})}^0] = \sum_{\mathbf{a}_{-i} \in \mathcal{A}_{-i}} \pi_{-i,h}^k(\mathbf{a}_{-i} | s) [P_{h,s,(a_i, \mathbf{a}_{-i})}^0].$$

We decompose the error in the value functions as follows:

$$\begin{aligned} & \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,h}^{*,\pi_{-i}}] - \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,h}^{\pi}] \\ & \leq \underbrace{\mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,h}^{*,\pi_{-i}}] - \mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h}^{\widehat{\pi}_i^*, \pi_{-i}}]}_A + \underbrace{\mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h}^{\widehat{\pi}_i^*, \pi_{-i}}] - \mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h}^{\pi}]}_B + \underbrace{\mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h}^{\pi}] - \mathbb{E}_{\pi \sim \widehat{\xi}} [V_{i,h}^{\pi}]}_C. \end{aligned} \quad (83)$$

We define the following auxiliary value functions for all $s \in \mathcal{S}$:

$$\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi(s) \right] = \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[r_{i,h}^k(s, a_i) \right] + \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^\pi \right] \right], \quad (84a)$$

$$\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^{\tilde{\pi}_i^*, \pi-i}(s) \right] = \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \tilde{\pi}_{i,h}^*(s)} \left[r_{i,h}^k(s, a_i) \right] + \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \tilde{\pi}_{i,h}^*(s)} \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^{\tilde{\pi}_i^*, \pi-i} \right] \right], \quad (84b)$$

$$\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^{*, \pi-i}(s) \right] = \max_{a_i \in \mathcal{A}_i} \sum_{k=1}^K \alpha_k^K \left[r_{i,h}^k(s, a_i) + \left(\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^{*, \pi-i} \right] \right) \right], \quad (84c)$$

where for all $s \in \mathcal{S}$, we also have

$$\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,H+1}^\pi(s) \right] = \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,H+1}^{\tilde{\pi}_i^*, \pi-i}(s) \right] = \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,H+1}^{*, \pi-i}(s) \right] = 0$$

Here, we use the fact that $\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^{*, \pi-i} \right] \geq \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^{\tilde{\pi}_i^*, \pi-i} \right]$. Using the error decomposition in (83), we will now individually bound the three terms, A , B , and C , in the following sections.

C.2 Controlling B: adversarial online learning

C.2.1 Step 1: showing that $\hat{V}_{i,h}$ is an entry-wise upper bound on $\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^{*, \pi-i} \right]$

The following lemma demonstrates that the value estimate $\hat{V}_{i,h}$ for the i^{th} player serves as an optimistic estimate of the auxiliary value $\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^{*, \pi-i} \right]$, as defined in (84).

Lemma 4. *With probability at least $1 - \delta$, it holds that*

$$\hat{V}_{i,h} \geq \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^{*, \pi-i} \right], \quad \text{for all } (i, h) \in [n] \times [H].$$

Proof. See Appendix C.4.1 □

The following lemma demonstrates that the value estimate $\hat{V}_{i,h}$ for the i^{th} player serves as an optimistic estimate of the auxiliary value $\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^{*, \pi-i} \right]$, as defined in (84).

Lemma 5. *For value vector $\hat{V}_{i,h}$ and $\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi \right]$, it holds that*

$$\hat{V}_{i,h} \geq \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi \right], \quad \text{for all } (i, h) \in [n] \times [H].$$

Proof. See Appendix C.4.2 □

C.2.2 Step 2: constructing recursion

To begin with, according to the definition of $\hat{V}_{i,h}(s)$ and $\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi(s) \right]$, we have

$$\begin{aligned} & \hat{V}_{i,h}(s) - \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi(s) \right] \\ &= \min \left\{ \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \hat{V}_{i,h+1} \right] + \beta_{i,h}(s), H - h + 1 \right\} \end{aligned}$$

$$\begin{aligned}
& - \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^\pi \right] \right] \\
\leq & \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \hat{V}_{i,h+1} \right] + \beta_{i,h}(s) \\
& - \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^\pi \right] \right] \\
= & \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k} \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \hat{V}_{i,h+1} \right] + \beta_{i,h}(s) - \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k} \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^\pi \right] \right]
\end{aligned} \tag{85}$$

To simplify the notations, we define transition kernel associated estimated value function similarly as (32). For all $k \in [K]$, we define matrix notations $\hat{P}_{i,h}^{\pi^k, \hat{V}}$ and $\hat{P}_{i,h}^{\hat{\pi}^k, \bar{V}}$ as:

$$\begin{aligned}
\hat{P}_{i,h}^{\pi^k, \hat{V}} & := \hat{P}_{i,h}^{\pi^k, \hat{V}_{i,h+1}} \quad \text{and} \quad \hat{P}_{i,h,s,a_i}^{\pi^k, \hat{V}} := \hat{P}_{i,h,s,a_i}^{\pi^k, \hat{V}_{i,h+1}} = \operatorname{argmin}_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^{\pi^k})} \mathcal{P} \hat{V}_{i,h+1}, \\
\hat{P}_{i,h}^{\hat{\pi}^k, \bar{V}} & := \hat{P}_{i,h}^{\hat{\pi}^k, \mathbb{E}_{\pi \sim \hat{\xi}}[\bar{V}_{i,h+1}^\pi]} \quad \text{and} \quad \hat{P}_{i,h,s,a_i}^{\hat{\pi}^k, \bar{V}} := \hat{P}_{i,h,s,a_i}^{\hat{\pi}^k, \mathbb{E}_{\pi \sim \hat{\xi}}[\bar{V}_{i,h+1}^\pi]} = \operatorname{argmin}_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^{\hat{\pi}^k})} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}}[\bar{V}_{i,h+1}^\pi].
\end{aligned}$$

Additionally, we define square matrices $\underline{P}_{i,h}^{\pi^k, \bar{V}} \in \mathbb{R}^{S \times S}$ and $\underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \in \mathbb{R}^{S \times S}$ as: $\underline{P}_{i,h}^{\pi^k, \bar{V}} := \Pi_h^{\pi^k} \hat{P}_{i,h}^{\pi^k, \bar{V}}$ and $\underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} := \Pi_h^{\hat{\pi}^k} \hat{P}_{i,h}^{\hat{\pi}^k, \bar{V}}$. We rewrite the result of (85) in a vector form, we can obtain that

$$\begin{aligned}
& \hat{V}_{i,h} - \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi \right] \\
\leq & \sum_{k=1}^K \alpha_k^K \Pi_h^{\pi^k} \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^{\pi^k})} \mathcal{P} \hat{V}_{i,h+1} \right] + \beta_{i,h} - \sum_{k=1}^K \alpha_k^K \Pi_h^{\pi^k} \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^{\pi^k})} \mathcal{P} \bar{V}_{i,h+1}^\pi \right] \\
= & \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\pi^k, \hat{V}} \hat{V}_{i,h+1} + \beta_{i,h} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^\pi \right] \\
\leq & \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \left(\hat{V}_{i,h+1} - \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^\pi \right] \right) + \beta_{i,h}.
\end{aligned}$$

To continue, we first introduce an lemma of the upper bound for bonus vector $\beta_{i,h}$.

Lemma 6. *The bonus vector $\beta_{i,h}$ is bounded by the following inequality:*

$$\beta_{i,h} \leq 3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \left(H \cdot 1 + \sum_{k=1}^K \alpha_k^K \operatorname{Var}_{\underline{P}_{i,h}^{\pi^k, \hat{V}}} \hat{V}_{i,h+1} \right)$$

Proof. See Appendix C.4.3 □

To proceed, we introduce some notations for convenience. Let e_s denote the S -dimensional standard basis vector, with support on the s -th element. Additionally, we define:

$$b_h^h = e_s \quad \text{and} \quad b_h^j = e_s^\top \left[\prod_{r=h}^{j-1} \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,r}^{\hat{\pi}^k, \bar{V}} \right) \right], \quad \forall j = h+1, \dots, H. \tag{86}$$

Armed with above notations and fact, for any $s \in \mathcal{S}$, we have

$$\begin{aligned}
\widehat{V}_{i,h}(s) - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi(s) \right] &= \left\langle e_s, \widehat{V}_{i,h} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right] \right\rangle = \sum_{j=h}^H \left\langle b_h^j, \beta_{i,j} \right\rangle \\
&\leq \sum_{j=h}^H \left\langle b_h^j, 3c_b H \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \mathbf{1} \right\rangle + \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, 3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \widehat{V}_{i,j+1} \right\rangle \\
&= 3c_b \sqrt{\frac{H^3 \log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{K}} + 3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \widehat{V}_{i,j+1} \right\rangle. \tag{87}
\end{aligned}$$

With elementary inequality $\sqrt{\text{Var}_P(V + V')} \leq \sqrt{\text{Var}_P(V)} + \sqrt{\text{Var}_P(V')}$ for any transition kernel $P \in \mathbb{R}^{\mathcal{S}}$ and vector $V, V' \in \mathbb{R}^{\mathcal{S}}$, we further decompose (87) as

$$\begin{aligned}
&\widehat{V}_{i,h}(s) - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi(s) \right] \\
&\leq 3c_b \sqrt{\frac{H^3 \log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{K}} + 3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \widehat{V}_{i,j+1} \right\rangle \\
&\leq 3c_b \sqrt{\frac{H^3 \log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{K}} + 3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{k=1}^K \alpha_k^K \sum_{j=h}^H \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \left(\widehat{V}_{i,j+1} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle \\
&\quad + 3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{k=1}^K \alpha_k^K \sum_{j=h}^H \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle \\
&\leq 3c_b \sqrt{\frac{H^3 \log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{K}} + \underbrace{3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \left(\widehat{V}_{i,j+1} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle}_{\mathcal{D}_1} \\
&\quad + \underbrace{3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) - \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle}_{\mathcal{D}_2} \\
&\quad + \underbrace{3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle}_{\mathcal{D}_3}.
\end{aligned}$$

We now control the three terms $\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3$ separately.

Controlling \mathcal{D}_1 . We can directly obtain the following upper bound on \mathcal{D}_1 :

$$\begin{aligned}
\mathcal{D}_1 &= 3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,h}^{\pi^k, \widehat{V}}} \left(\widehat{V}_{i,j+1} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle \\
&\leq 3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \left\| \text{Var}_{\widehat{P}_{i,h}^{\pi^k, \widehat{V}}} \left(\widehat{V}_{i,j+1} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\|_\infty \cdot \mathbf{1} \right\rangle \\
&\leq 3c_b \sqrt{\frac{\log^3 \left(\frac{KS \sum_{i=1}^n A_i}{\delta} \right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \left\| \widehat{V}_{i,j+1} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right\|_\infty^2 \cdot \mathbf{1} \right\rangle
\end{aligned}$$

$$\begin{aligned}
&\stackrel{(i)}{\leq} 3c_b \sqrt{\frac{H \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \left\| \widehat{V}_{i,j+1} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right\|_\infty \cdot 1 \right\rangle \\
&\leq 3c_b \sqrt{\frac{H^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} \max_{h \leq j \leq H} \left\| \widehat{V}_{i,j+1} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right\|_\infty
\end{aligned} \tag{88}$$

where (i) follows from the elementary upper bound $\left\| \widehat{V}_{i,j+1} \right\|_\infty \leq H$, $\left\| \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right\|_\infty \leq H$ for all $h \leq j \leq H$.

Before deriving the upper bounds for the terms \mathcal{D}_2 and \mathcal{D}_3 , we first introduce the following auxiliary lemmas, which will be instrumental in the subsequent derivation.

Lemma 7. *For all $(i, h) \in [n] \times [H]$, the estimated robust value function $\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right]$ satisfies the following inequality:*

$$\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi(s) \right] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi(s) \right] \leq \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\}.$$

Proof. See Appendix C.4.4. □

With Lemma 7, we have the following lemma on variance base on different transition probability in the same uncertainty set, and we leave the proof to Appendix C.4.5.

Lemma 8. *For a transition kernel $P' \in \mathbb{R}^{\mathcal{S}}$ and any $\tilde{P} \in \mathbb{R}^{\mathcal{S}}$ such that $\tilde{P} \in \mathcal{U}^{\sigma_i}(P')$, the following bound holds for all $(i, h) \in [n] \times [H]$:*

$$\left| \text{Var}_{P'} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right] \right) - \text{Var}_{\tilde{P}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right] \right) \right| \leq \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\}. \tag{89a}$$

Controlling \mathcal{D}_2 . We can directly apply Lemma 8 and arrive at

$$\begin{aligned}
&\left| \text{Var}_{\widehat{P}_{i,h}^{\pi^k, \widehat{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^\pi \right] \right) - \text{Var}_{\widehat{P}_{i,h}^{\pi^k, \overline{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^\pi \right] \right) \right| \\
&\leq \left| \text{Var}_{\widehat{P}_{i,h}^{\pi^k, \widehat{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^\pi \right] \right) - \text{Var}_{\widehat{P}_{i,h}^{\pi^k}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^\pi \right] \right) \right| \\
&\quad + \left| \text{Var}_{\widehat{P}_{i,h}^{\pi^k}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^\pi \right] \right) - \text{Var}_{\widehat{P}_{i,h}^{\pi^k, \overline{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^\pi \right] \right) \right| \\
&\leq 2 \min \left\{ \frac{1}{\sigma_i}, H \right\}.
\end{aligned}$$

We insert (90) back to the expression of \mathcal{D}_2 , and we can obtain that

$$\begin{aligned}
\mathcal{D}_2 &= 3c_b \sqrt{\frac{\log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \widehat{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) - \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \overline{V}}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle \\
&\leq 3c_b \sqrt{\frac{\log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{KH}} \sum_{j=h}^H \left\langle b_h^j, 2 \min \left\{ \frac{1}{\sigma_i}, H \right\} 1 \right\rangle \\
&= 6c_b \sqrt{\frac{H \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} \min \left\{ \frac{1}{\sigma_i}, H \right\}.
\end{aligned} \tag{90}$$

Controlling \mathcal{D}_3 . We first apply Lemma 12, and we can directly deduce that

$$\begin{aligned} \mathcal{D}_3 &= 3c_b \sqrt{\frac{\log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{KH}} \sum_{j=h}^H \sum_{k=1}^K \alpha_k^K \left\langle b_h^j, \text{Var}_{\widehat{P}_{i,j}^{\pi^k, \nabla}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle \\ &\leq 3c_b \sqrt{\frac{\log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{KH}} \sum_{j=h}^H \left\langle b_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,j}^{\pi^k, \nabla}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle \end{aligned}$$

We now introduce the following lemma on $\sum_{j=h}^H \left\langle b_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,j}^{\pi^k, \nabla}} \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right\rangle$, which is an empirical-transition version of Lemma 16.

Lemma 9. *Let $\delta \in (0, 1)$. With probability at least $1 - \delta$, the following condition holds for all $(h, i) \in [H] \times [n]$:*

$$\sum_{j=h}^H \left\langle b_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,j}^{\pi^k, \nabla}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle \leq 3H \left(\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi(s) \right] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi(s) \right] \right). \quad (91)$$

Proof. See Appendix C.4.6. □

Therefore, we can further achieve the following upper bound of \mathcal{D}_3 by applying Lemma 9:

$$\begin{aligned} \mathcal{D}_3 &\leq 3c_b \sqrt{\frac{\log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{KH}} \sum_{j=h}^H \left\langle b_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,j}^{\pi^k, \nabla}} \left(\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right) \right\rangle \\ &\leq 9c_b \sqrt{\frac{H \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} \left(\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi(s) \right] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi(s) \right] \right) \\ &\stackrel{(i)}{\leq} 9c_b \sqrt{\frac{H \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} \min \left\{ \frac{1}{\sigma_i}, H \right\} \end{aligned} \quad (92)$$

where (i) holds due to Lemma 8.

C.2.3 Step 3: summing up the result

We combine the result of (88), (90), (92), yielding

$$\begin{aligned} \widehat{V}_{i,h} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right] &\leq 3c_b \sqrt{\frac{H^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} 1 + \mathcal{D}_1 + \mathcal{D}_2 + \mathcal{D}_3 \\ &\leq c_b \sqrt{\frac{H \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} \left(3H + 15 \min \left\{ \frac{1}{\sigma_i}, H \right\} \right) 1 \\ &\quad + 3c_b \sqrt{\frac{H^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} \max_{h \leq j \leq H} \left\| \widehat{V}_{i,j+1} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right\|_\infty 1. \end{aligned}$$

Moreover, Lemma 5 implies that $\widehat{V}_{i,h} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right] = \left| \widehat{V}_{i,h} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right] \right|$, which indicates that

$$\begin{aligned} &\max_{h \in [H]} \left\| \widehat{V}_{i,h} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right] \right\|_\infty \\ &\leq c_b \sqrt{\frac{H \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} \left(3H + 15 \min \left\{ \frac{1}{\sigma_i}, H \right\} \right) \\ &\quad + 3c_b \sqrt{\frac{H^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} \max_{h \leq j \leq H} \left\| \widehat{V}_{i,j+1} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,j+1}^\pi \right] \right\|_\infty \end{aligned}$$

$$\begin{aligned}
&\stackrel{(i)}{\leq} 18c_b \sqrt{\frac{H^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} + \frac{1}{2} \max_{h \in [H]} \left\| \widehat{V}_{i,h} - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right] \right\|_\infty \\
&\leq 36c_b \sqrt{\frac{H^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}}
\end{aligned}$$

where (i) holds by taking $K \geq 12c_b^2 H^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)$, and involving the basic facts that $\widehat{V}_{i,H+1} = \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,H+1}^\pi \right] = 0$. Eventually, we can achieve the following upper bound of term B :

$$\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^{\star, \pi-i} \right] - \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right] \leq 36c_b \sqrt{\frac{H^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} 1. \quad (93)$$

C.3 Controlling terms A and C

In this section, we derive an upper bound for the difference between the true value function and the estimated value function. We consider a more general case involving a given set of policies $\{\widehat{\pi}_h^k\}_{(h,k) \in [H] \times [K]}$, where either $\widehat{\pi}_h^k = \pi_h^k$ for all $(h,k) \in [H] \times [K]$, or $\widehat{\pi}_h^k = \widehat{\pi}_i^* \times \pi_{-i,h}^k$ for all $(h,k) \in [H] \times [K]$. Additionally, we define a distribution over the set of policies $\zeta := \{\zeta_h\}_{h \in [H]}$, with $\zeta_h : [H] \mapsto \Delta(\mathcal{S} \mapsto \prod_{i \in [n]} \Delta(\mathcal{A}_i))$, where $\zeta_h(\widehat{\pi}_h^k) = \alpha_k^K$ for all $(h,k) \in [H] \times [K]$. Our objective is to derive an upper bound for $\left| \mathbb{E}_{\pi \sim \zeta} \left[\overline{V}_{i,h}^\pi(s) \right] - \mathbb{E}_{\pi \sim \zeta} \left[\overline{V}_{i,h}^\pi(s) \right] \right|$, where for all $s \in \mathcal{S}$, $\mathbb{E}_{\pi \sim \zeta} \left[\overline{V}_{i,h}^\pi(s) \right]$ is defined as

$$\mathbb{E}_{\pi \sim \zeta} \left[\overline{V}_{i,h}^\pi(s) \right] = \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \widehat{\pi}_{i,h}^k}(s) [r_{i,h}^k(s, a_i)] + \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \widehat{\pi}_{i,h}^k}(s) \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(\widehat{P}_{i,h,s,a_i}^{\widehat{\pi}_i^k})} \mathcal{P} \mathbb{E}_{\pi \sim \zeta} \left[\overline{V}_{i,h+1}^\pi \right] \right],$$

with $\mathbb{E}_{\pi \sim \zeta} \left[\overline{V}_{i,H+1}^\pi(s) \right] = 0$. Here, $r_{i,h}^k(s, a_i)$ represents the empirical estimation of $r_{i,h}^k(s, a_i)$, and $\widehat{P}_{i,h,s,a_i}^{\widehat{\pi}_i^k}$ denotes the empirical estimation of $P_{h,s,a_i}^{\widehat{\pi}_i^k}$ for all $(h, s, a_i, k) \in [H] \times \mathcal{S} \times \mathcal{A}_i \times [K]$. For notational clarity, we define the empirical reward vector $\widehat{r}_{i,h}^k \in \mathbb{R}^{\mathcal{S}}$, such that $\widehat{r}_{i,h}^k(s) = \mathbb{E}_{a_i \sim \widehat{\pi}_{i,h}^k}(s) [r_{i,h}^k(s, a_i)]$ for all $s \in \mathcal{S}$.

We first introduce the following two lemmas in terms of estimation error of transition model and reward function:

Lemma 10. *Let $\delta \in (0, 1)$ and consider any $(h, i, k) \in [H] \times [n] \times [K]$. With a probability of at least $1 - \delta$, for any fixed value vector $V \in \mathbb{R}^{\mathcal{S}}$, where $0 \leq V(s) \leq H$ for all $s \in \mathcal{S}$, the following inequality holds:*

$$\begin{aligned}
\left| P_{i,h}^{\widehat{\pi}_i^k, V} - \widehat{P}_{i,h}^{\widehat{\pi}_i^k, V} \right| &\leq 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i N H K}{\delta}\right)}{N}} \sqrt{\text{Var}_{P_{i,h}^{\widehat{\pi}_i^k}}(V)} + \frac{\log\left(\frac{18S \sum_{i=1}^n A_i N H K}{\delta}\right)}{N} 1 \\
&\leq 3 \sqrt{\frac{H^2 \log\left(\frac{18S \sum_{i=1}^n A_i N K H}{\delta}\right)}{N}} 1,
\end{aligned}$$

where $\text{Var}_{P_{i,h}^{\widehat{\pi}_i^k}}(\cdot)$ is as defined in (34).

Proof. See Appendix C.4.7. □

Lemma 11. *There exists a constant c_r such that for any fixed pair $(h, i) \in [H] \times [n]$, with probability at least $1 - \delta$, the following inequality holds:*

$$\left| \sum_{k=1}^K \alpha_k^K \widehat{r}_{i,h}^k - \sum_{k=1}^K \alpha_k^K \overline{r}_{i,h}^k \right| \leq c_r \sqrt{\frac{\log\left(\frac{KS}{\delta}\right)}{K}} 1.$$

Proof. See Appendix C.4.8. \square

For clarity of presentation, we extend the definitions in (32) and introduce additional notations related to transitions associated with the estimated value function. With a slight abuse of notation, we define the matrix notations $\widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}}$ and $\widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}}$ as follows for all $(i, h, k) \in [n] \times [H] \times [K]$:

$$\begin{aligned} P_{i,h}^{\widehat{\pi}^k, \widehat{V}} &:= P_{i,h}^{\widehat{\pi}^k, \mathbb{E}_{\pi \sim \zeta}[\widehat{V}_{i,h+1}^\pi]} \quad \text{and} \quad P_{i,h,s,a_i}^{\widehat{\pi}^k, \widehat{V}} := P_{i,h,s,a_i}^{\widehat{\pi}^k, \mathbb{E}_{\pi \sim \zeta}[\widehat{V}_{i,h+1}^\pi]} = \operatorname{argmin}_{\mathcal{P} \in \mathcal{U}_\rho^{\sigma_i} \left(P_{h,s,a_i}^{\widehat{\pi}^k} \right)} \mathcal{P} \mathbb{E}_{\pi \sim \zeta} \left[\widehat{V}_{i,h+1}^\pi \right], \\ \widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} &:= \widehat{P}_{i,h}^{\widehat{\pi}^k, \mathbb{E}_{\pi \sim \zeta}[\widehat{V}_{i,h+1}^\pi]} \quad \text{and} \quad \widehat{P}_{i,h,s,a_i}^{\widehat{\pi}^k, \widehat{V}} := \widehat{P}_{i,h,s,a_i}^{\widehat{\pi}^k, \mathbb{E}_{\pi \sim \zeta}[\widehat{V}_{i,h+1}^\pi]} = \operatorname{argmin}_{\mathcal{P} \in \mathcal{U}_\rho^{\sigma_i} \left(\widehat{P}_{h,s,a_i}^{\widehat{\pi}^k} \right)} \mathcal{P} \mathbb{E}_{\pi \sim \zeta} \left[\widehat{V}_{i,h+1}^\pi \right]. \end{aligned}$$

Additionally, we define the square matrices $\widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} \in \mathbb{R}^{S \times S}$ and $P_{i,h}^{\widehat{\pi}^k, \widehat{V}} \in \mathbb{R}^{S \times S}$ as:

$$\widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} := \Pi_h^{\widehat{\pi}^k} \widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}}, \quad P_{i,h}^{\widehat{\pi}^k, \widehat{V}} := \Pi_h^{\widehat{\pi}^k} P_{i,h}^{\widehat{\pi}^k, \widehat{V}}.$$

At any time step $h \in [H]$, we have

$$\begin{aligned} & \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi] - \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h}^\pi] \\ & \stackrel{(i)}{=} \sum_{k=1}^K \alpha_k^K r_{i,h}^{\widehat{\pi}^k} + \sum_{k=1}^K \alpha_k^K \Pi_h^{\widehat{\pi}^k} \left[\inf_{\mathcal{P} \in \mathcal{U}_\rho^{\sigma_i} \left(P_{h,s,a_i}^{\widehat{\pi}^k} \right)} \mathcal{P} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] \right] \\ & \quad - \sum_{k=1}^K \alpha_k^K \widehat{r}_{i,h}^{\widehat{\pi}^k} - \sum_{k=1}^K \alpha_k^K \Pi_h^{\widehat{\pi}^k} \left[\inf_{\mathcal{P} \in \mathcal{U}_\rho^{\sigma_i} \left(\widehat{P}_{h,s,a_i}^{\widehat{\pi}^k} \right)} \mathcal{P} \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h+1}^\pi] \right] \\ & \stackrel{(ii)}{=} \sum_{k=1}^K \alpha_k^K r_{i,h}^{\widehat{\pi}^k} + \sum_{k=1}^K \alpha_k^K P_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] - \sum_{k=1}^K \alpha_k^K \widehat{r}_{i,h}^{\widehat{\pi}^k} - \sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h+1}^\pi] \end{aligned} \quad (94)$$

where (i) holds by the robust Bellman equation in (13) with matrix notation in (32), (ii) arises from the definition in (84). Moreover, through simple observation, we directly have $P_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] \leq \widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h+1}^\pi]$ for all $(h, k) \in [H] \times [K]$. Thus, we further control (94) as

$$\begin{aligned} & \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi] - \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h}^\pi] \\ & = \sum_{k=1}^K \alpha_k^K r_{i,h}^{\widehat{\pi}^k} + \sum_{k=1}^K \alpha_k^K P_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] - \sum_{k=1}^K \alpha_k^K \widehat{r}_{i,h}^{\widehat{\pi}^k} - \sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h+1}^\pi] \\ & = \sum_{k=1}^K \alpha_k^K \left[\left(r_{i,h}^{\widehat{\pi}^k} - \widehat{r}_{i,h}^{\widehat{\pi}^k} \right) + \left(P_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] - \widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h+1}^\pi] \right) \right. \\ & \quad \left. + \left(\widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h+1}^\pi] - \widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h+1}^\pi] \right) \right] \\ & \leq \sum_{k=1}^K \alpha_k^K \left(P_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] - \widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h+1}^\pi] \right) \\ & \quad + \underbrace{\sum_{k=1}^K \alpha_k^K \left[\left| r_{i,h}^{\widehat{\pi}^k} - \widehat{r}_{i,h}^{\widehat{\pi}^k} \right| + \left| P_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] - \widehat{P}_{i,h}^{\widehat{\pi}^k, \widehat{V}} \mathbb{E}_{\pi \sim \zeta} [\widehat{V}_{i,h+1}^\pi] \right]}_{:= a_{i,h}^\zeta}. \end{aligned} \quad (95)$$

Applying (95) recursively leads to

$$\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi] - \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] \leq \sum_{j=h}^H \left[\prod_{r=h}^{j-1} \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,r}^{\hat{\pi}^k, \bar{V}} \right) \right] a_{i,j}^\zeta, \quad (96)$$

where the inequality holds by adopting the following notations:

$$\begin{aligned} \left[\prod_{r=h}^{h-1} \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,r}^{\hat{\pi}^k, \bar{V}} \right) \right] &= I, \\ \left[\prod_{r=h}^{j-1} \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,r}^{\hat{\pi}^k, \bar{V}} \right) \right] &= \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \cdot \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h+1}^{\hat{\pi}^k, \bar{V}} \right) \cdots \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j-1}^{\hat{\pi}^k, \bar{V}} \right). \end{aligned}$$

Next, similar to (95), we can achieve that

$$\begin{aligned} &\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] - \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi] \\ &\stackrel{(i)}{=} \sum_{k=1}^K \alpha_k^K \bar{r}_{i,h}^{\hat{\pi}^k} + \sum_{k=1}^K \alpha_k^K \hat{\underline{P}}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] - \sum_{k=1}^K \alpha_k^K r_{i,h}^{\hat{\pi}^k} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] \\ &= \sum_{k=1}^K \alpha_k^K \left[\left(\bar{r}_{i,h}^{\hat{\pi}^k} - r_{i,h}^{\hat{\pi}^k} \right) + \left(\hat{\underline{P}}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] - \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \right) \right. \\ &\quad \left. + \left(\underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] - \underline{P}_{i,h}^{\hat{\pi}^k, V} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] \right) \right] \\ &\stackrel{(ii)}{\leq} \sum_{k=1}^K \alpha_k^K \left(\underline{P}_{i,h}^{\hat{\pi}^k, V} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] - \underline{P}_{i,h}^{\hat{\pi}^k, V} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^\pi] \right) \\ &\quad + \sum_{k=1}^K \alpha_k^K \left[\left| r_{i,h}^{\hat{\pi}^k} - \bar{r}_{i,h}^{\hat{\pi}^k} \right| + \left| \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] - \hat{\underline{P}}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \right| \right], \end{aligned}$$

where (i) holds due to robust Bellman equation, and (ii) holds due to the direct observation that $\underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \leq \underline{P}_{i,h}^{\hat{\pi}^k, V} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi]$. Then following the routine of achieving (96), we can obtain that

$$\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] - \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi] \leq \sum_{j=h}^H \left[\prod_{r=h}^{j-1} \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,r}^{\hat{\pi}^k, V} \right) \right] a_{i,j}^\zeta. \quad (97)$$

Summing up (97) and (96), one has

$$\begin{aligned} \left| \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] - \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi] \right| &\leq \max \left\{ \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi] - \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi], \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] - \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi] \right\} \\ &\leq \max \left\{ \sum_{j=h}^H \left[\prod_{r=h}^{j-1} \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,r}^{\hat{\pi}^k, \bar{V}} \right) \right] a_{i,j}^\zeta, \sum_{j=h}^H \left[\prod_{r=h}^{j-1} \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,r}^{\hat{\pi}^k, V} \right) \right] a_{i,j}^\zeta \right\}, \end{aligned} \quad (98)$$

where the max operator is taken entry-wise for vectors. To continue, we apply Lemma 10 and Lemma 11, and we can obtain the following upper bound on $a_{i,j}^\zeta$ for all $(i, j) \in [n] \times [H]$:

$$a_{i,h}^\zeta = \sum_{k=1}^K \alpha_k^K \left[\left| r_{i,h}^{\hat{\pi}^k} - \bar{r}_{i,h}^{\hat{\pi}^k} \right| + \left| \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] - \hat{\underline{P}}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \right| \right]$$

$$\begin{aligned}
&\leq 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} \sqrt{\text{Var}_{\underline{P}_h^{\hat{\pi}^k}}\left(\mathbb{E}_{\pi \sim \zeta}\left[\bar{V}_{i,h+1}^\pi\right]\right)} \\
&\quad + \frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N} \mathbf{1} + c_r \sqrt{\frac{\log\left(\frac{K S n H}{\delta}\right)}{K}} \mathbf{1},
\end{aligned}$$

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

C.3.1 Controlling the first term in (98)

To simplify notation, let us introduce some additional symbols. Recall that e_s represents the standard basis vector in S -dimensional space associated with the s -th component. We define

$$d_h^h = e_s \quad \text{and} \quad d_h^j = e_s^\top \left[\prod_{r=h}^{j-1} \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,r}^{\hat{\pi}^k, \bar{V}} \right) \right] \quad \text{for } j = h+1, \dots, H. \quad (99)$$

With these notations in place, for any $s \in \mathcal{S}$, we consider

$$\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] - \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi(s)] = \langle e_s, \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi] - \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] \rangle = \sum_{j=h}^H \langle d_h^j, a_{i,j}^\zeta \rangle.$$

Applying Lemma 10, we obtain

$$\begin{aligned}
&\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] - \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi(s)] \\
&\leq \sum_{j=h}^H \left\langle d_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N H}{\delta}\right)}{N}} \sqrt{\text{Var}_{\underline{P}_j^{\hat{\pi}^k}}\left(\mathbb{E}_{\pi \sim \zeta}\left[\bar{V}_{i,j+1}^\pi\right]\right)} \right\rangle \\
&\quad + \frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N H}{\delta}\right)}{N} + c_r \sqrt{\frac{\log\left(\frac{K S n H}{\delta}\right)}{K}} \\
&\leq \frac{H \log\left(\frac{18S \sum_{i=1}^n A_i K N H}{\delta}\right)}{N} + c_r \sqrt{\frac{H^2 \log\left(\frac{K S n H}{\delta}\right)}{K}} \\
&\quad + \sum_{j=h}^H \left\langle d_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N H}{\delta}\right)}{N}} \sqrt{\text{Var}_{\underline{P}_j^{\hat{\pi}^k}}\left(\mathbb{E}_{\pi \sim \zeta}\left[\bar{V}_{i,j+1}^\pi\right]\right)} \right\rangle. \quad (100)
\end{aligned}$$

By applying the triangle inequality, we can further decompose the term of interest as follows:

$$\begin{aligned}
&\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] - \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi(s)] \\
&\leq \frac{H \log\left(\frac{18S \sum_{i=1}^n A_i K N H}{\delta}\right)}{N} + c_r \sqrt{\frac{H^2 \log\left(\frac{K S n H}{\delta}\right)}{K}} \\
&\quad + \underbrace{\sum_{j=h}^H \left\langle d_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N H}{\delta}\right)}{N}} \sqrt{\text{Var}_{\underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}}\left(\mathbb{E}_{\pi \sim \zeta}\left[\bar{V}_{i,j+1}^\pi\right]\right)} \right\rangle}_{\mathcal{B}_1} \\
&\quad + \underbrace{\sum_{j=h}^H 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N H}{\delta}\right)}{N}} \left\langle d_h^j, \sqrt{\left| \text{Var}_{\underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}}\left(\mathbb{E}_{\pi \sim \zeta}\left[\bar{V}_{i,j+1}^\pi\right]\right) - \text{Var}_{\underline{P}_j^{\hat{\pi}^k}}\left(\mathbb{E}_{\pi \sim \zeta}\left[\bar{V}_{i,j+1}^\pi\right]\right) \right|} \right\rangle}_{\mathcal{B}_2}. \quad (101)
\end{aligned}$$

We then analyze the bounds for the terms \mathcal{B}_1 and \mathcal{B}_2 separately.

Controlling \mathcal{B}_1 . First, we introduce the following lemma and corresponding inequality to establish control over the term $\sum_{k=1}^K \alpha_k^K \sqrt{\text{Var}_{\underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right)}$:

Lemma 12. For any transition kernels $P_1, \dots, P_m \in \mathbb{R}^S$, and any weight $a_1, \dots, a_m \in [0, 1]$ with $a_1 + \dots + a_m = 1$, one has

$$\sum_{i=1}^m a_i \sqrt{\text{Var}_{P_i}(V)} \leq \sqrt{\text{Var}_{\sum_{i=1}^m a_i P_i}(V)},$$

where V denote any fixed value vector $V \in \mathbb{R}^S$ with $0 \leq V(s) \leq H$ for all $s \in S$.

Proof. Initially, since $f(x) = \sqrt{x}$ is a concave function, we have

$$\sum_{i=1}^m a_i \sqrt{\text{Var}_{P_i}(V)} = \sqrt{\sum_{i=1}^m a_i \text{Var}_{P_i}(V)}.$$

Moreover, according to the definition of variance in (33), we obtain that

$$\sum_{i=1}^m a_i \text{Var}_{P_i}(V) = \sum_{i=1}^m a_i (\mathbb{E}_{P_i}(V \circ V) - (\mathbb{E}_{P_i} V \circ \mathbb{E}_{P_i} V)) \leq \sum_{i=1}^m a_i \mathbb{E}_{P_i}(V \circ V) - \left(\sum_{i=1}^m a_i \mathbb{E}_{P_i} V \right) \circ \left(\sum_{i=1}^m a_i \mathbb{E}_{P_i} V \right),$$

where the last inequality holds due to the elementary fact that $f(x) = x^2$ is a convex function. Therefore, we have proven the result of the lemma. \square

With Lemma 12, we can further control \mathcal{B}_1 with

$$\begin{aligned} \mathcal{B}_1 &= \sum_{j=h}^H \left\langle d_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\text{Var}_{\underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right)} \right\rangle \\ &\leq 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sum_{j=h}^H \left\langle d_h^j, \sqrt{\text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right)} \right\rangle \\ &\leq 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{H \sum_{j=h}^H \left\langle d_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right) \right\rangle}. \end{aligned} \quad (102)$$

The last inequality holds due to Cauchy-Schwartz inequality. To further achieve the upper bound of \mathcal{B}_1 , we introduce the following lemma of $\sum_{j=h}^H \left\langle d_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right) \right\rangle$:

Lemma 13. Consider any $\delta \in (0, 1)$. With probability at least $1 - \delta$, one has

$$\begin{aligned} \forall (h, i) \in [H] \times [n]: \quad &\sum_{j=h}^H \left\langle d_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right) \right\rangle \\ &\leq 3H \left(\max_{s \in S} \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h+1}^\pi(s) \right] - \min_{s \in S} \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h+1}^\pi(s) \right] \right) \left(1 + 2H \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i n K N H}{\delta} \right)}{N}} \right). \end{aligned} \quad (103)$$

Proof. See Appendix C.4.9. \square

Lemma 14. For all $(i, h) \in [n] \times [H]$, the estimated robust value function $\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi]$ satisfies the following inequality:

$$\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi(s)] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi(s)] \leq \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\}.$$

Proof. The proof of Lemma 14 closely parallels that of Lemma 7. Therefore, we omit the details here for brevity and clarity. \square

Apply Lemma 13 to (102), we arrive at

$$\begin{aligned} \mathcal{B}_1 &\leq 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{H \sum_{j=h}^H \left\langle d_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K P_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] \right) \right\rangle} \\ &\leq \sqrt{3H^2 \left(\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi(s)] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi(s)] \right)} \\ &\quad \cdot 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \cdot \left(1 + 2H \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i n K N H}{N} \right)}{N}} \right) \\ &\stackrel{(i)}{\leq} 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i n K N n H}{\delta} \right)}{N}} \sqrt{3H^2 \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\} \left(1 + 2H \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i n K N H}{N} \right)}{N}} \right)} \\ &\leq 6 \sqrt{\frac{H^2 \min \{1/\sigma_i, H\} \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}}, \end{aligned} \tag{104}$$

where (i) holds by applying Lemma 7 and Lemma 14, and the final inequality follows by taking $N \geq 4H^2 \log \left(\frac{18S \sum_{i=1}^n A_i K n N H}{\delta} \right)$.

Controlling \mathcal{B}_2 . Initially, with similar analysis as Lemma 8, we have the following lemma:

Lemma 15. For transition kernel $P' \in \mathbb{R}^{\mathcal{S}}$ and any $\tilde{P} \in \mathbb{R}^{\mathcal{S}}$ such that $\tilde{P} \in \mathcal{U}^{\sigma_i}(P')$, the following bounds are established for all $(i, h) \in [n] \times [H]$:

$$\left| \text{Var}_{P'} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] \right) - \text{Var}_{\tilde{P}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] \right) \right| \leq \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\}.$$

With Lemma 15, we observe that

$$\begin{aligned} &\left| \text{Var}_{P_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] \right) - \text{Var}_{P_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] \right) \right| \\ &\stackrel{(i)}{=} \left| \Pi_j^{\pi_i} \left(\text{Var}_{P_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] \right) - \text{Var}_{P_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] \right) \right) \right| \\ &\stackrel{(ii)}{\leq} \left\| \text{Var}_{P_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] \right) - \text{Var}_{P_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] \right) \right\|_{\infty} 1 \\ &\leq \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\} 1, \end{aligned} \tag{105}$$

where (i) and (ii) follows from the definition of matrix notations Π_j^π (cf A.1) and $P_j^{\hat{\pi}^k}, P_{i,j}^{\hat{\pi}^k, \bar{V}}$ (cf A.1), and the last inequality holds by applying Lemma 15 with $P' = P_{j,s,a_i}^{\pi^k}, \tilde{P} = P_{i,j,s,a_i}^{\hat{\pi}^k, \bar{V}}$ for all $(s, a_i) \in \mathcal{S} \times \mathcal{A}_i$.

Plugging back (105) to (101), it can be verified that

$$\begin{aligned}
\mathcal{B}_2 &= \sum_{j=h}^H 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} \left\langle d_h^j, \sqrt{\left| \text{Var}_{\underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right) - \text{Var}_{\underline{P}_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right) \right|} \right\rangle \\
&\leq \sum_{j=h}^H 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} \left\langle d_h^j, \sqrt{\min\left\{ \frac{1}{\sigma_i}, H \right\} 1} \right\rangle \\
&\leq 2 \sqrt{\frac{H^2 \min\left\{ \frac{1}{\sigma_i}, H \right\} \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} \tag{106}
\end{aligned}$$

Consequently, combining (104) and (106), (101) can be bounded by

$$\begin{aligned}
&\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] - \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi(s)] \\
&\leq \frac{H \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N} + c_r \sqrt{\frac{H^2 \log\left(\frac{K S n H}{\delta}\right)}{K}} + \mathcal{B}_1 + \mathcal{B}_2 \\
&\leq \frac{H \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N} + c_r \sqrt{\frac{H^2 \log\left(\frac{K S n H}{\delta}\right)}{K}} + 8 \sqrt{\frac{H^2 \min\left\{ \frac{1}{\sigma_i}, H \right\} \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} \\
&\leq c_r \sqrt{\frac{H^2 \log\left(\frac{K S n H}{\delta}\right)}{K}} + 12 \sqrt{\frac{H^2 \min\left\{ \frac{1}{\sigma_i}, H \right\} \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}}, \tag{107}
\end{aligned}$$

where the last inequality holds by taking $N \geq 4H^2 \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)$.

C.3.2 Controlling the second term in (98)

To do so, similar to (99), we define

$$w_h^h = e_s \quad \text{and} \quad w_h^j = e_s^\top \left[\prod_{r=h}^{j-1} \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,r}^{\hat{\pi}^k, V} \right) \right] \quad \forall j = h+1, \dots, H. \tag{108}$$

With the above notations in mind, following the routine of (100) gives: for any $s \in \mathcal{S}$,

$$\begin{aligned}
&\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] - \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi(s)] \\
&\leq \sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} \sqrt{\text{Var}_{\underline{P}_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right)} \right. \\
&\quad \left. + \frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N} 1 + c_r \sqrt{\frac{\log\left(\frac{K S n H}{\delta}\right)}{K}} 1 \right\rangle \\
&\leq \frac{H \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N} + c_r \sqrt{\frac{H^2 \log\left(\frac{K S n H}{\delta}\right)}{K}} \\
&\quad + \sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} \sqrt{\text{Var}_{\underline{P}_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right)} \right\rangle. \tag{109}
\end{aligned}$$

Furthermore, following the routine established in (101), we can decompose the expression as follows:

$$\begin{aligned}
& \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] - \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi(s)] \\
& \leq \frac{H \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N} + c_r \sqrt{\frac{H^2 \log \left(\frac{K S n H}{\delta} \right)}{K}} \\
& \quad + \sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\text{Var}_{P_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] \right)} \right\rangle \\
& \stackrel{(i)}{\leq} \frac{H \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N} + c_r \sqrt{\frac{H^2 \log \left(\frac{K S n H}{\delta} \right)}{K}} \\
& \quad + \underbrace{\sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\text{Var}_{P_{i,j}^{\hat{\pi}^k, V}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right)} \right\rangle}_{\mathcal{B}_3} \\
& \quad + \underbrace{\sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\left| \text{Var}_{P_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right) - \text{Var}_{P_{i,j}^{\hat{\pi}^k, V}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right) \right|} \right\rangle}_{\mathcal{B}_4} \\
& \quad + \underbrace{\sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\text{Var}_{P_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] - \mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right)} \right\rangle}_{\mathcal{B}_5} \tag{110}
\end{aligned}$$

where (i) holds due to the triangle inequality and the fundamental inequality $\sqrt{\text{Var}_P(V + V')} \leq \sqrt{\text{Var}_P(V)} + \sqrt{\text{Var}_P(V')}$ for any transition kernel $P \in \mathbb{R}^S$ and vectors $V, V' \in \mathbb{R}^S$.

Next, we will control the three main terms $\mathcal{B}_3, \mathcal{B}_4, \mathcal{B}_5$ in (110) separately as outlined below:

Controlling \mathcal{B}_3 . Initially, we apply Lemma 12, and we can obtain the following upper bound of \mathcal{B}_3 :

$$\begin{aligned}
\mathcal{B}_3 &= \sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\text{Var}_{P_{i,j}^{\hat{\pi}^k, V}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right)} \right\rangle \\
&\leq 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sum_{j=h}^H \left\langle w_h^j, \sqrt{\text{Var}_{\sum_{k=1}^K \alpha_k^K P_{i,j}^{\hat{\pi}^k, V}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right)} \right\rangle.
\end{aligned}$$

We further apply Cauchy-Schwartz inequality, and we can obtain that

$$\begin{aligned}
\mathcal{B}_3 &\leq 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sum_{j=h}^H \left\langle w_h^j, \sqrt{\text{Var}_{\sum_{k=1}^K \alpha_k^K P_{i,j}^{\hat{\pi}^k, V}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right)} \right\rangle \\
&\leq 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{H \sum_{j=h}^H \left\langle w_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K P_{i,j}^{\hat{\pi}^k, V}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right) \right\rangle}
\end{aligned}$$

In addition, to further bound the term of interest, we introduce the following lemma and inequalities for $\sum_{j=h}^H \left\langle w_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K P_{i,j}^{\hat{\pi}^k, V}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right) \right\rangle$.

Lemma 16. For any joint policy π , we have for all $(h, i) \in [H] \times [n]$:

$$\sum_{j=h}^H \left\langle w_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, v}} (\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi]) \right\rangle \leq 3H \left(\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] \right). \quad (111)$$

Proof. See Appendix C.4.10. \square

Lemma 17. For all $(i, h) \in [n] \times [H]$, the estimated robust value function $\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi]$ satisfies the following inequality:

$$\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)] \leq \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\}.$$

Proof. The proof of Lemma 17 closely parallels that of Lemma 7. Therefore, we omit the details here for brevity and clarity. \square

Then applying Lemma 16 and Lemma 17 yields

$$\begin{aligned} \mathcal{B}_3 &\leq 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{H \sum_{j=h}^H \left\langle w_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, v}} (\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi]) \right\rangle} \\ &\leq 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{3H^2 \left(\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^{\pi, \sigma_i}(s)] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^{\pi, \sigma_i}(s)] \right)} \\ &\leq 4 \sqrt{\frac{H^2 \min \{1/\sigma_i, H\} \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}}, \end{aligned} \quad (112)$$

where the last inequality follows from Lemma 17.

Controlling \mathcal{B}_4 and \mathcal{B}_5 With similar analysis as Lemma 8, we have the following lemma:

Lemma 18. For any joint policy π , transition kernel $P' \in \mathbb{R}^S$, and any $\tilde{P} \in \mathbb{R}^S$ such that $\tilde{P} \in \mathcal{U}^{\sigma_i}(P')$, the following bounds are established for all $(i, h) \in [n] \times [H]$:

$$|\text{Var}_{P'} (\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi]) - \text{Var}_{\tilde{P}} (\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi])| \leq \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\}.$$

We apply Lemma 18, and we can directly obtain the following upper bound of \mathcal{B}_4 :

$$\begin{aligned} \mathcal{B}_4 &= \sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\left| \text{Var}_{\underline{P}_j^{\hat{\pi}^k}} (\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi]) - \text{Var}_{\underline{P}_{i,j}^{\hat{\pi}^k, v}} (\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi]) \right|} \right\rangle \\ &\leq \sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\min \left\{ \frac{1}{\sigma_i}, H \right\} 1} \right\rangle \\ &= 2 \sqrt{\frac{H^2 \min \{1/\sigma_i, H\} \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \end{aligned} \quad (113)$$

Then the remainder of the proof shall focus on \mathcal{B}_5 . Recalling the definition in (110), one has

$$\mathcal{B}_5 = \sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\text{Var}_{\underline{P}_j^{\hat{\pi}^k}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] - \mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right)} \right\rangle$$

$$\begin{aligned}
&\leq \sum_{j=h}^H \left\langle w_h^j, 2 \sum_{k=1}^K \alpha_k^K \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sqrt{\left\| \text{Var}_{\mathcal{P}_j^{\pi^k}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[V_{i,j+1}^\pi \right] \right) \right\|_\infty} \mathbf{1} \right\rangle \\
&\leq 2 \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \sum_{j=h}^H \left\langle w_h^j, \left\| \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[V_{i,j+1}^\pi \right] \right\|_\infty \mathbf{1} \right\rangle \\
&\leq 2 \sqrt{\frac{H^2 \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \max_{h \leq j \leq H} \left\| \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[V_{i,j+1}^\pi \right] \right\|_\infty. \tag{114}
\end{aligned}$$

Summing up (112), (113), and (114) and inserting back to (110), we conclude

$$\begin{aligned}
&\mathbb{E}_{\pi \sim \zeta} \left[V_{i,h}^\pi(s) \right] - \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h}^\pi(s) \right] \\
&\leq \frac{H \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N} + c_r \sqrt{\frac{H^2 \log \left(\frac{K S n H}{\delta} \right)}{K}} + \mathcal{B}_3 + \mathcal{B}_4 + \mathcal{B}_5 \\
&\leq \frac{H \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N} + c_r \sqrt{\frac{H^2 \log \left(\frac{K S n H}{\delta} \right)}{K}} + 6 \sqrt{\frac{H^2 \min\{1/\sigma_i, H\} \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \\
&\quad + 2 \sqrt{\frac{H^2 \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \max_{h \leq j \leq H} \left\| \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[V_{i,j+1}^\pi \right] \right\|_\infty \tag{115}
\end{aligned}$$

C.3.3 Summing up the results: upper bound for term A and C

Inserting (107) and (115) back into (98), we observe that

$$\begin{aligned}
&\left| \mathbb{E}_{\pi \sim \zeta} \left[V_{i,h}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h}^\pi \right] \right| \\
&\leq \max \left\{ \mathbb{E}_{\pi \sim \zeta} \left[V_{i,h}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h}^\pi \right], \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[V_{i,h}^\pi \right] \right\} \\
&\leq \max \left\{ c_r \sqrt{\frac{H^2 \log \left(\frac{K S n H}{\delta} \right)}{K}} \mathbf{1} + 12 \sqrt{\frac{H^2 \min \left\{ \frac{1}{\sigma_i}, H \right\} \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \mathbf{1}, \frac{H \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N} \mathbf{1} \right. \\
&\quad + 2 \sqrt{\frac{H^2 \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \max_{h \leq j \leq H} \left\| \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[V_{i,j+1}^\pi \right] \right\|_\infty \mathbf{1} + c_r \sqrt{\frac{H^2 \log \left(\frac{K S n H}{\delta} \right)}{K}} \mathbf{1} \\
&\quad \left. + 6 \sqrt{\frac{H^2 \min \left\{ \frac{1}{\sigma_i}, H \right\} \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \mathbf{1} \right\},
\end{aligned}$$

which indicates that

$$\begin{aligned}
&\max_{h \in [H]} \left\| \mathbb{E}_{\pi \sim \zeta} \left[V_{i,h}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h}^\pi \right] \right\|_\infty \tag{116} \\
&\leq c_r \sqrt{\frac{H^2 \log \left(\frac{K S n H}{\delta} \right)}{K}} + 12 \sqrt{\frac{H^2 \min \left\{ \frac{1}{\sigma_i}, H \right\} \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} + \frac{H \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N} \\
&\quad + 2 \sqrt{\frac{H^2 \log \left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta} \right)}{N}} \max_{h \in [H]} \left\| \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h+1}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[V_{i,h+1}^\pi \right] \right\|_\infty
\end{aligned}$$

$$\begin{aligned}
&\stackrel{(i)}{\leq} c_r \sqrt{\frac{H^2 \log\left(\frac{KS_n H}{\delta}\right)}{K}} + 12 \sqrt{\frac{H^2 \min\left\{\frac{1}{\sigma_i}, H\right\} \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} \\
&\quad + \frac{H \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N} + \frac{1}{2} \max_{h \in [H]} \left\| \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h}^\pi \right] - \mathbb{E}_{\pi \sim \zeta} \left[V_{i,h}^\pi \right] \right\|_\infty \\
&\leq 2c_r \sqrt{\frac{H^2 \log\left(\frac{KS_n H}{\delta}\right)}{K}} + 24 \sqrt{\frac{H^2 \min\left\{\frac{1}{\sigma_i}, H\right\} \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} + \frac{2H \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}
\end{aligned} \tag{117}$$

where (i) holds when $N \geq 4H^2 \log\left(\frac{18S \sum_{i=1}^n A_i N n H}{\delta}\right)$, and invoking the basic fact that $\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,H+1}^\pi \right] = \mathbb{E}_{\pi \sim \zeta} \left[V_{i,H+1}^\pi \right] = 0$. With (117), we can achieve the following upper bound on term A and term C :

$$\begin{aligned}
&\mathbb{E}_{\pi \sim \hat{\xi}} \left[V_{i,h}^{*,\pi-i} \right] - \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^{\tilde{\pi}^*,\pi-i} \right] \\
&\leq 2c_r \sqrt{\frac{H^2 \log\left(\frac{KS_n H}{\delta}\right)}{K}} + 24 \sqrt{\frac{H^2 \min\left\{\frac{1}{\sigma_i}, H\right\} \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} + \frac{2H \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}
\end{aligned} \tag{118}$$

$$\begin{aligned}
&\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi \right] - \mathbb{E}_{\pi \sim \hat{\xi}} \left[V_{i,h}^\pi \right] \\
&\leq 2c_r \sqrt{\frac{H^2 \log\left(\frac{KS_n H}{\delta}\right)}{K}} + 24 \sqrt{\frac{H^2 \min\left\{\frac{1}{\sigma_i}, H\right\} \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}} + \frac{2H \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}
\end{aligned} \tag{119}$$

C.3.4 Summing up the results

Summing up the results in (118), (93), (119), we can achieve the upper bound of our target:

$$\begin{aligned}
\mathbb{E}_{\pi \sim \hat{\xi}} \left[V_{i,1}^{*,\pi-i} \right] - \mathbb{E}_{\pi \sim \hat{\xi}} \left[V_{i,1}^\pi \right] &\leq 36c_b \sqrt{\frac{H^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i}{\delta}\right)}{K}} + 4c_r \sqrt{\frac{H^2 \log\left(\frac{KS_n H}{\delta}\right)}{K}} \\
&\quad + \frac{4H \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N} + 48 \sqrt{\frac{H^2 \min\left\{\frac{1}{\sigma_i}, H\right\} \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right)}{N}}.
\end{aligned}$$

Therefore, there exists a constant C , such that when N and K satisfies:

$$N \geq CH^2 \min\left\{\frac{1}{\min_{1 \leq i \leq n} \sigma_i}, H\right\} \log\left(\frac{18S \sum_{i=1}^n A_i K N n H}{\delta}\right) \frac{1}{\epsilon^2}, \quad K \geq CH^3 \log^3\left(\frac{KS \sum_{i=1}^n A_i H}{\delta}\right) \frac{1}{\epsilon^2}$$

we achieve $\max_{i \in [n]} \mathbb{E}_{\pi \sim \hat{\xi}} \left[V_{i,1}^{*,\pi-i} \right] - \mathbb{E}_{\pi \sim \hat{\xi}} \left[V_{i,1}^\pi \right] \leq \epsilon \cdot 1$ with probability at least $1 - \delta$. Therefore, the total number of samples we need is at least

$$N_{\text{all}} = HS \sum_{i=1}^n KN = \tilde{O}\left(\frac{S \max_{1 \leq i \leq n} A_i H^6}{\epsilon^4} \min\left\{\frac{1}{\min_{1 \leq i \leq n} \sigma_i}, H\right\}\right).$$

Thus, we finish the proof of Theorem 2.

C.4 Proof of auxiliary lemmas

C.4.1 Proof of Lemma 4

Before proving Lemma 4, we first introduce the following lemma regarding the properties of the learning rate.

Lemma 19 (Li et al. (2023, Lemma 1)). *For any $k \geq 1$, one has*

$$\alpha_1 = 1, \quad \sum_{i=1}^k \alpha_i^k = 1, \quad \max_{1 \leq i \leq k} \alpha_i^k \leq \frac{2c_\alpha \log K}{k}. \quad (120a)$$

In addition, if $k \geq c_\alpha \log K + 1$ and $c_\alpha \geq 24$, then one has

$$\max_{1 \leq i \leq k/2} \alpha_i^k \leq \frac{1}{K^6}. \quad (120b)$$

We will now prove the lemma with induction argument. Initially, the base step $H + 1$ trivially holds true, since we have

$$\widehat{V}_{i,H+1} = \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,H+1}^{*,\pi-i} \right] = 0.$$

Next, we assume that the lemma holds for step $h + 1$, namely

$$\widehat{V}_{i,h+1} \geq \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^{*,\pi-i} \right]$$

and attempt to justify the validity of Lemma 4 for step h . Let l_k denote $l_k = -q_{i,h}^k(s, \cdot), \forall k \geq 1$, then the update rule of Algorithm 2 can be viewed as the FTRL algorithm applied to the loss vectors $\{l_k\}_{k \in [K]}$. According to the definition of $\{\eta_k\}_{k \in [K]}$ and $\{\alpha_k\}_{k \in [K]}$, we have

$$\left(\frac{\eta_k}{\eta_{k+1}} \right)^2 = \frac{\alpha_k}{\alpha_{k-1}} = \frac{k-2+c_\alpha \log K}{k-1+c_\alpha \log K} \geq \frac{k-1}{k-1+c_\alpha \log K} = 1 - \alpha_k > (1 - \alpha_k)^2. \quad (121)$$

This property (121) permits us to invoke Theorem 3 to obtain

$$\begin{aligned} & \max_{a_i \in \mathcal{A}_i} \sum_{k=1}^K \alpha_k^K q_{i,h}^k(s, a_i) - \sum_{k=1}^K \alpha_k^K \langle \pi_{i,h}^k, q_{i,h}^k(s, \cdot) \rangle \\ &= \max_{a_i \in \mathcal{A}_i} \left\{ \sum_{k=1}^K \alpha_k^K \langle \pi_{i,h}^k(s), l_k \rangle - \sum_{k=1}^K \alpha_k^K l_k(a_i) \right\} \\ &\leq \frac{5}{3} \sum_{k=2}^K \alpha_k^K \frac{\eta_k \alpha_k}{1 - \alpha_k} \text{Var}_{\pi_{i,h}^k(s)} \left(q_{i,h}^k(s, \cdot) \right) + \frac{\log A_i}{\eta_{K+1}} + \tau_{i,h} \end{aligned}$$

where $\tau_{i,h}$ is defined as

$$\tau_{i,h} := \frac{5}{3} \alpha_1^K \eta_2 \|q_{i,h}^1\|_\infty^2 + \left\{ 3 \sum_{k=2}^K \alpha_k^K \frac{\eta_k^2 \alpha_k^2}{(1 - \alpha_k)^2} \|q_{i,h}^k\|_\infty^3 \mathbf{1} \left(\frac{\eta_k \alpha_k}{1 - \alpha_k} \|q_{i,h}^k\|_\infty > \frac{1}{3} \right) \right\} + 3 \alpha_1^K \eta_2^2 \|q_{i,h}^1\|_\infty^3. \quad (122)$$

According to the definition of $\{\alpha_k\}_{k=1}^K$ and $\{\eta_k\}_{k=1}^K$, we have the following fact:

$$1 - \alpha_k = 1 - \frac{c_\alpha \log K}{k-1+c_\alpha \log K} \geq \begin{cases} 1 - \frac{c_\alpha \log K}{1+c_\alpha \log K} = \frac{1}{1+c_\alpha \log K} \geq \frac{1}{2c_\alpha \log K}, & \text{if } k \geq 2, \\ 1 - \frac{c_\alpha \log K}{K/2+c_\alpha \log K} = \frac{K}{K+2c_\alpha \log K} \geq \frac{1}{2}, & \text{if } k \geq K/2 + 1, \end{cases} \quad (123a)$$

$$\eta_k \alpha_k = \sqrt{\frac{\log K}{\alpha_{k-1} H}} \cdot \alpha_k \leq \sqrt{\frac{\log K}{\alpha_k H}} \cdot \alpha_k = \sqrt{\frac{\alpha_k \log K}{H}} \leq \sqrt{\frac{2c_\alpha \log^2 K}{kH}}. \quad (123b)$$

Therefore, we can re-control $\max_{a_i \in \mathcal{A}_i} \sum_{k=1}^K \alpha_k^K q_{i,h}^k(s, a_i) - \sum_{k=1}^K \alpha_k^K \langle \pi_{i,h}^k, q_{i,h}^k(s, \cdot) \rangle$ with

$$\max_{a_i \in \mathcal{A}_i} \sum_{k=1}^K \alpha_k^K q_{i,h}^k(s, a_i) - \sum_{k=1}^K \alpha_k^K \langle \pi_{i,h}^k, q_{i,h}^k(s, \cdot) \rangle \quad (124)$$

$$\begin{aligned}
&\leq \frac{5}{3} \sum_{k=2}^K \alpha_k^K \frac{\eta_k \alpha_k}{1 - \alpha_k} \text{Var}_{\pi_{i,h}^k(s)} \left(q_{i,h}^k(s, \cdot) \right) + \frac{\log A_i}{\eta_{K+1}} + \tau_{i,h} \\
&\stackrel{(i)}{\leq} \underbrace{\frac{5}{3} \sum_{k=2}^{K/2} \frac{(2c_\alpha)^{1.5} \log^2 K}{\sqrt{kH}} \alpha_k^K \text{Var}_{\pi_{i,h}^k(s)} \left(q_{i,h}^k(s, \cdot) \right)}_{\mathcal{C}_1} \\
&\quad + \frac{20}{3} \sum_{k=K/2+1}^K \alpha_k^K \sqrt{\frac{c_\alpha \log^2 K}{KH}} \text{Var}_{\pi_{i,h}^k(s)} \left(q_{i,h}^k(s, \cdot) \right) + \underbrace{\frac{\log A_i}{\eta_{K+1}}}_{\mathcal{C}_2} + \underbrace{\tau_{i,h}}_{\mathcal{C}_3}, \tag{125}
\end{aligned}$$

Now we separately control the four terms $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$ in (125).

- For term \mathcal{C}_1 , we have

$$\begin{aligned}
\sum_{k=2}^{K/2} \frac{\alpha_k^K \log^2 K}{\sqrt{kH}} \text{Var}_{\pi_{i,h}^k(s)} \left(q_{i,h}^k(s, \cdot) \right) &\leq \sum_{k=2}^{K/2} \frac{\log^2 K}{K^6 \sqrt{kH}} \text{Var}_{\pi_{i,h}^k(s)} \left(q_{i,h}^k(s, \cdot) \right) \\
&\leq \sum_{k=2}^{K/2} \frac{\log^2 K}{K^6 \sqrt{kH}} \|q_{i,h}^k(s, \cdot)\|_\infty^2 \leq \frac{H^{3/2} \log^2 K}{K^6} \sum_{k=2}^{K/2} \frac{1}{\sqrt{k}} \\
&\leq \frac{2H^{3/2} \log^2 K}{K^6} \cdot \sqrt{K/2} \leq \frac{2H^{3/2} \log^2 K}{K^5}, \tag{126}
\end{aligned}$$

where the third inequality holds due to the elementary bound $\|q_{i,h}^k(s, \cdot)\|_\infty \leq H$.

- For term \mathcal{C}_2 , we have

$$\frac{\log A_i}{\eta_{K+1}} = \log A_i \sqrt{\frac{\alpha_K H}{\log K}} \leq \sqrt{\frac{2c_\alpha H \log^2 A_i}{K}}, \tag{127}$$

where the first equality holds due to the definition of η_{K+1} .

- For term \mathcal{C}_3 , we initially have

$$\frac{\eta_k \alpha_k}{1 - \alpha_k} \|q_{i,h}^k\|_\infty \leq \frac{\sqrt{\frac{2c_\alpha \log^2 K}{kH}}}{\frac{1}{2c_\alpha \log K}} \cdot H = \sqrt{\frac{8c_\alpha^3 H \log^4 K}{k}}. \tag{128}$$

Clearly, the right-hand side of (128) is upper bounded by $1/3$ for all k obeying $k \geq c_9 H^2 \log^4 \frac{K}{\delta}$ for some large enough constant $c_9 > 0$. Consequently, one can derive

$$\begin{aligned}
\tau_{i,h} &= \frac{5}{3} \alpha_1^K \eta_2 \|q_{i,h}^1\|_\infty^2 + \left\{ 3 \sum_{k=2}^K \alpha_k^K \frac{\eta_k^2 \alpha_k^2}{(1 - \alpha_k)^2} \|q_{i,h}^k\|_\infty^3 \mathbf{1} \left(\frac{\eta_k \alpha_k}{1 - \alpha_k} \|q_{i,h}^k\|_\infty > \frac{1}{3} \right) \right\} + 3 \alpha_1^K \eta_2 \|q_{i,h}^1\|_\infty^3 \\
&\leq \frac{5}{3K^6} \sqrt{\frac{\log K}{H}} \|q_{i,h}^1\|_\infty^2 + \frac{(2c_\alpha \log K)^2}{K^6} \left\{ 3 \sum_{k=2}^{c_9 H^2 \log^4 \frac{K}{\delta}} \eta_k^2 \alpha_k^2 \|q_{i,h}^k\|_\infty^3 \right\} + \frac{3 \log K}{K^6 H} \|q_{i,h}^1\|_\infty^3 \\
&\leq \frac{24c_\alpha^3 \log^4 K}{K^6 H} \left\{ \sum_{k=1}^K \frac{1}{k} H^3 \right\} \\
&\leq \frac{24c_\alpha^3 H^3 \log^5 K}{K^6} \leq \frac{1}{K^4}, \tag{129}
\end{aligned}$$

where the second line comes from (123) and the fact that $K/2 > c_9 H \log^4 \frac{K}{\delta}$.

Combining previous three items, we can obtain that

$$\begin{aligned}
& \max_{a_i \in \mathcal{A}_i} \sum_{k=1}^K \alpha_k^K q_{i,h}^k(s, a_i) - \sum_{k=1}^K \alpha_k^K \langle \pi_{i,h}^k, q_{i,h}^k(s, \cdot) \rangle \\
& \leq \frac{5(2c_\alpha)^{1.5}}{3} \cdot \frac{2H^{3/2} \log^2 K}{K^5} + \frac{20}{3} \sqrt{\frac{c_\alpha \log^2 K}{KH}} \sum_{k=K/2+1}^K \alpha_k^K \text{Var}_{\pi_{i,h}^k(s)}(q_{i,h}^k(s, \cdot)) + \sqrt{\frac{2c_\alpha H \log^2 A_i}{K}} + \frac{1}{K^4} \\
& \leq 10 \sqrt{\frac{c_\alpha \log^3(KA_i)}{KH}} \sum_{k=1}^K \alpha_k^K \text{Var}_{\pi_{i,h}^k(s)}(q_{i,h}^k(s, \cdot)) + 2 \sqrt{\frac{c_\alpha H \log^3(KA_i)}{K}}, \tag{130}
\end{aligned}$$

According to the definition of $q_{i,h}^k(s, a_i)$ in the update rule of Algorithm 2, we have

$$\begin{aligned}
& \max_{a_i \in \mathcal{A}_i} \sum_{k=1}^K \alpha_k^K \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \widehat{V}_{i,h+1} \right] - \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \widehat{V}_{i,h+1} \right] \\
& = \max_{a_i \in \mathcal{A}_i} \sum_{k=1}^K \alpha_k^K q_{i,h}^k(s, a_i) - \sum_{k=1}^K \alpha_k^K \langle \pi_{i,h}^k, q_{i,h}^k(s, \cdot) \rangle \\
& \leq 10 \sqrt{\frac{c_\alpha \log^3(KA_i)}{KH}} \sum_{k=1}^K \alpha_k^K \text{Var}_{\pi_{i,h}^k(s)}(q_{i,h}^k(s, \cdot)) + 2 \sqrt{\frac{c_\alpha H \log^3(KA_i)}{K}} = \beta_{i,h}(s)
\end{aligned}$$

Moreover, according to the induction hypothesis, we have for all $s \in \mathcal{S}$

$$\begin{aligned}
\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^{\star, \pi-i}(s) \right] & = \max_{a_i \in \mathcal{A}_i} \sum_{k=1}^K \alpha_k^K \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^{\star, \pi-i} \right] \right] \\
& \leq \max_{a_i \in \mathcal{A}_i} \sum_{k=1}^K \alpha_k^K \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \widehat{V}_{i,h+1} \right] \\
& \leq \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \widehat{V}_{i,h+1} \right] + \beta_{i,h}(s) \\
& \leq \widehat{V}_{i,h}(s).
\end{aligned}$$

Thus, we finished the proof of the lemma.

C.4.2 Proof of Lemma 5

We will prove the lemma with induction argument. Initially, the base step $H+1$ trivially holds true, since we have

$$\widehat{V}_{i,H+1} = \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,H+1}^\pi \right] = 0.$$

Next, we assume that the lemma holds for step $h+1$, namely

$$\widehat{V}_{i,h+1} \geq \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^\pi \right].$$

According to the definition of $\widehat{V}_{i,h}$ and $\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi \right]$, we have

$$\mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h}^\pi(s) \right] = \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \widehat{\xi}} \left[\overline{V}_{i,h+1}^\pi \right] \right]$$

$$\begin{aligned}
&\leq \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \widehat{V}_{i,h+1} \right] \\
&\leq \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \widehat{V}_{i,h+1} \right] + \beta_{i,h}(s).
\end{aligned}$$

Since, we also trivially have $\mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h}^\pi(s)] \leq H - h + 1$, we can deduce that for all $s \in \mathcal{S}$

$$\begin{aligned}
&\mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h}^\pi(s)] \\
&\leq \min \left\{ \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[r_{i,h}^k(s, a_i) + \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \widehat{V}_{i,h+1} \right] + \beta_{i,h}(s), H - h + 1 \right\} = \widehat{V}_{i,h}(s).
\end{aligned}$$

Thus, we finished the proof of the lemma.

C.4.3 Proof of Lemma 6

Recall that for all $s \in \mathcal{S}$, bonus term $\beta_{i,h}(s)$ is defined as

$$\beta_{i,h}(s) = c_b \sqrt{\frac{\log^3(\frac{KS \sum_{i=1}^n A_i}{\delta})}{KH}} \sum_{k=1}^K \alpha_k^K \left\{ \text{Var}_{\pi_{i,h}^k(\cdot|s)}(q_{i,h}^k(s, \cdot)) + H \right\}. \quad (131)$$

For any $k \in [K]$, we have the following inequality for $\text{Var}_{\pi_{i,h}^k(\cdot|s)}(q_{i,h}^k(s, \cdot))$:

$$\begin{aligned}
&\text{Var}_{\pi_{i,h}^k(\cdot|s)}(q_{i,h}^k(s, \cdot)) \\
&\leq 2\text{Var}_{\pi_{i,h}^k(\cdot|s)}(r_{i,h}^k(s, \cdot)) + 2\text{Var}_{\pi_{i,h}^k(\cdot|s)} \left(\sum_{s'} \widehat{P}_{i,h}^{\pi_{i,h}^k, \widehat{V}}(s' | s, \cdot) \widehat{V}_{i,h+1}(s') \right) \\
&\stackrel{(i)}{\leq} 2 + 2 \left[\sum_{a_i \in \mathcal{A}_i} \pi_{i,h}^k(a_i | s) \widehat{P}_{i,h}^{\pi_{i,h}^k, \widehat{V}}(\cdot | s, a_i) \left(\widehat{V}_{i,h+1} \circ \widehat{V}_{i,h+1} \right) - \left(\sum_{a_i \in \mathcal{A}_i} \pi_{i,h}^k(a_i | s) \widehat{P}_{i,h}^{\pi_{i,h}^k, \widehat{V}}(\cdot | s, a_i) \widehat{V}_{i,h+1} \right)^2 \right] \\
&= 2 + 2 \left\langle e_s, \text{Var}_{\widehat{P}_{i,h}^{\pi_{i,h}^k, \widehat{V}}} \widehat{V}_{i,h+1} \right\rangle. \quad (132)
\end{aligned}$$

where e_s denotes an S -dimensional standard basis supported on the s -th element, and (i) holds due to the elementary fact that $|r_{i,h}^k(s, a_i)| \leq 1$ and $|\widehat{P}_{i,h}^{\pi_{i,h}^k, \widehat{V}}(s' | s, a_i)| \leq 1$ for all $s, s' \in \mathcal{S}, a_i \in \mathcal{A}_i$. We insert the result of (132) back to (131), and rewrite the result in vector form, we can achieve that

$$\beta_{i,h} \leq 3c_b \sqrt{\frac{\log^3(\frac{KS \sum_{i=1}^n A_i}{\delta})}{KH}} \left(H \cdot 1 + \text{Var}_{\widehat{P}_{i,h}^{\pi_{i,h}^k, \widehat{V}}} \widehat{V}_{i,h+1} \right)$$

C.4.4 Proof of Lemma 7

To prove Lemma 7, we start by analyzing the value function of policy π under uncertainty set σ_i . We first establish bounds on $\min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h}^\pi(s)]$:

$$\begin{aligned}
&\min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h}^\pi(s)] \\
&= \min_{s \in \mathcal{S}} \sum_{k=1}^K \alpha_k^K \left[\mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} [r_{i,h}^k(s, a_i)] + \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h+1}^\pi(s)] \right] \right] \\
&\geq \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \widehat{\xi}} [\overline{V}_{i,h+1}^\pi(s)].
\end{aligned}$$

This follows from the robust Bellman equation (13).

Next, we examine $\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h}^\pi(s)]$:

$$\begin{aligned}
& \max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h}^\pi(s)] \\
&= \max_{s \in \mathcal{S}} \sum_{k=1}^K \alpha_k^K \left[\mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} [r_{i,h}^k(s, a_i)] + \mathbb{E}_{a_i \sim \pi_{i,h}^k(s)} \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi(s)] \right] \right] \\
&\leq 1 + \sum_{k=1}^K \alpha_k^K \max_{(s,a_i) \in \mathcal{S} \times \mathcal{A}_i} \left[\inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{h,s,a_i}^{\pi-i})} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi(s)] \right]. \tag{133}
\end{aligned}$$

We now construct an auxiliary distribution vector $P'_{h,s,a_i} \in \mathbb{R}^{\mathcal{S}}$ by strictly reducing some elements of $P_{h,s,a_i}^{\pi-i}$ such that:

$$0 \leq P'_{h,s,a_i} \leq P_{i,h,s,a_i}^k \quad \text{and} \quad \sum_{s' \in \mathcal{S}} P_{i,h,s,a_i}^k(s') - P'_{h,s,a_i}(s') = \|P'_{h,s,a_i} - P_{i,h,s,a_i}^k\|_1 = \sigma_i. \tag{134}$$

Let $e_{s_{i,h}^*}$ denote the standard basis vector supported on $s_{i,h}^*$. We can show:

$$\frac{1}{2} \left\| P'_{h,s,a_i} + \sigma_i [e_{s_{i,h}^*}]^\top - P_{i,h,s,a_i}^k \right\|_1 \leq \frac{1}{2} \|P'_{h,s,a_i} - P_{i,h,s,a_i}^k\|_1 + \frac{1}{2} \left\| \sigma_i [e_{s_{i,h}^*}]^\top \right\|_1 \leq \sigma_i, \tag{135}$$

where the first inequality follows from the triangle inequality of the total variation distance.

From (135), we conclude that:

$$\begin{aligned}
& \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i}(P_{i,h,s,a_i}^k)} \mathcal{P} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi] \\
&\leq \left(P'_{h,s,a_i} + \sigma_i [e_{s_{i,h}^*}]^\top \right) \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi] \\
&\leq \|P'_{h,s,a_i}\|_1 \left\| \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi] \right\|_\infty + \sigma_i \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi(s_{i,h}^*)] \\
&\leq (1 - \sigma_i) \max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi(s)] + \sigma_i \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi(s)].
\end{aligned} \tag{136}$$

Substituting (136) into (133) yields:

$$\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h}^\pi(s)] \leq 1 + (1 - \sigma_i) \max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi(s)] + \sigma_i \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi(s)]. \tag{137}$$

Combining (133) and (137) gives:

$$\begin{aligned}
& \max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h}^\pi(s)] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h}^\pi(s)] \\
&\leq 1 + (1 - \sigma_i) \left(\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi(s)] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+1}^\pi(s)] \right) \\
&\leq 1 + (1 - \sigma_i) \left[1 + (1 - \sigma_i) \left(\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+2}^\pi(s)] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h+2}^\pi(s)] \right) \right] \\
&\leq \dots \leq \frac{1 - (1 - \sigma_i)^{H-h}}{\sigma_i} \leq \frac{1}{\sigma_i}. \tag{138}
\end{aligned}$$

Combining this with the basic fact that $\max_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h}^\pi(s)] - \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} [\bar{V}_{i,h}^\pi(s)] \leq H - h + 1$, we complete the proof.

C.4.5 Proof for Lemma 8

We introduce the following notation for the value function at time h :

$$\forall h \in [H], \quad \bar{V}_{i,h}^{\text{span}} := \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi(s) \right] - \min_{s' \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi(s') \right], \quad (139)$$

which normalizes the value function $\bar{V}_{i,h}^\pi$. This definition leads to the following bound:

$$\left\| \bar{V}_{i,h}^{\text{span}} \right\|_\infty \leq \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\}, \quad (140)$$

a result derived using Lemma 7. With this notation established, we now consider any transition kernel $P' \in \mathbb{R}^{\mathcal{S}}$ and any $\tilde{P} \in \mathbb{R}^{\mathcal{S}}$ such that $\tilde{P} \in \mathcal{U}^{\sigma_i}(P')$. For all $(i, h) \in [n] \times [H]$, we analyze the variance difference between the value functions under these kernels:

$$\begin{aligned} \left| \text{Var}_{P'} \left(\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi \right] \right) - \text{Var}_{\tilde{P}} \left(\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi \right] \right) \right| &= \left| \text{Var}_{P'} \left(\bar{V}_{i,h}^{\text{span}} \right) - \text{Var}_{\tilde{P}} \left(\bar{V}_{i,h}^{\text{span}} \right) \right| \\ &\leq \left\| \tilde{P} - P' \right\|_1 \left\| \bar{V}_{i,h}^{\text{span}} \right\|_\infty \\ &\leq \sigma_i \left(\min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\} \right)^2 \leq \min \left\{ \frac{1}{\sigma_i}, H - h + 1 \right\}. \end{aligned} \quad (141)$$

C.4.6 Proof of Lemma 9

Analogous to Appendix C.4.9, we introduce some auxiliary values and reward functions to control

$$\sum_{j=h}^H \left\langle b_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k \hat{P}_{i,j}^{\pi^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,j+1}^\pi \right] \right) \right\rangle$$

as below for any time step h and agent i .

Definition 3. For any time step $h \in [H]$ and the i -th agent, we denote $\bar{V}_h^{\min} := \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi(s) \right]$ as the minimum value of all the entries in vector $\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi \right]$. We further define $\bar{V}'_h := \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi \right] - \bar{V}_h^{\min} \mathbf{1}$ as the truncated value function. Eventually for reward function, we define $\bar{r}_{i,h}^{\min} = \sum_{k=1}^K \alpha_k \mathbb{E}_{a_i \sim \pi_{i,h}^k} r_{i,h}^k(\cdot, a_i) + \left(\bar{V}_{h+1}^{\min} - \bar{V}_h^{\min} \right) \mathbf{1}$ as the truncated reward function..

Then applying the robust Bellman's consistency equation in (13) gives

$$\begin{aligned} \bar{V}'_h &= \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h}^\pi \right] - \bar{V}_h^{\min} \mathbf{1} = \sum_{k=1}^K \alpha_k \mathbb{E}_{a_i \sim \pi_{i,h}^k} r_{i,h}^k(\cdot, a_i) + \sum_{k=1}^K \alpha_k \hat{P}_{i,h}^{\pi^k, \bar{V}} \mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^\pi \right] - \bar{V}_h^{\min} \mathbf{1} \\ &= \sum_{k=1}^K \alpha_k \mathbb{E}_{a_i \sim \pi_{i,h}^k} r_{i,h}^k(\cdot, a_i) + \left(\bar{V}_{h+1}^{\min} - \bar{V}_h^{\min} \right) \mathbf{1} + \sum_{k=1}^K \alpha_k \hat{P}_{i,h}^{\pi^k, \bar{V}} \bar{V}'_{h+1} \end{aligned} \quad (142)$$

$$= \bar{r}_{i,h}^{\min} + \sum_{k=1}^K \alpha_k \hat{P}_{i,h}^{\pi^k, \bar{V}} \bar{V}'_{h+1}. \quad (143)$$

The above fact leads to

$$\begin{aligned} &\text{Var}_{\sum_{k=1}^K \alpha_k \hat{P}_{i,h}^{\pi^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}_{i,h+1}^\pi \right] \right) \\ &\stackrel{(i)}{=} \text{Var}_{\sum_{k=1}^K \alpha_k \hat{P}_{i,h}^{\pi^k, \bar{V}}} \left(\bar{V}'_{h+1} \right) \end{aligned}$$

$$= \sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \left(\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, \bar{V}} \bar{V}'_{h+1} \right) \circ \left(\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, \bar{V}} \bar{V}'_{h+1} \right)$$

where (i) follows from the fact that $\text{Var}_{\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, V}}(V - b1) = \text{Var}_{\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, V}}(V)$ for any value vector $V \in \mathbb{R}^S$ and scalar b . According to (143) and (13), we have

$$\begin{aligned} \text{Var}_{\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}'_{i,h+1} \right] \right) &= \sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \left(\bar{V}'_h - \bar{r}_{i,h}^{\min} \right)^{\circ 2} \\ &= \sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \bar{V}'_h \circ \bar{V}'_h + 2\bar{V}'_h \circ \bar{r}_{i,h}^{\min} - \bar{r}_{i,h}^{\min} \circ \bar{r}_{i,h}^{\min} \\ &\leq \sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \bar{V}'_h \circ \bar{V}'_h + 2\|\bar{V}'_h\|_{\infty}, \end{aligned}$$

where the last inequality arises from $\bar{r}_{i,h}^{\min} \leq \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi^k} r_{i,h}^k(\cdot, a_i) \leq 1$ since $\bar{V}'_{h+1} - \bar{V}'_h \leq 0$ by definition.

Consequently, combining (155) and the definition of b_h^j in (99), we arrive at

$$\begin{aligned} &\sum_{j=h}^H \left\langle b_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, V}} \left(\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}'_{i,j+1} \right] \right) \right\rangle \\ &= \sum_{j=h}^H \left(b_h^j \right)^{\top} \left(\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, \bar{V}} \left(\bar{V}'_{j+1} \circ \bar{V}'_{j+1} \right) - \bar{V}'_j \circ \bar{V}'_j + 2\|\bar{V}'_h\|_{\infty} \mathbf{1} \right) \\ &\stackrel{(i)}{\leq} \sum_{j=h}^H \left[\left(b_h^j \right)^{\top} \left(\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, \bar{V}} \left(\bar{V}'_{j+1} \circ \bar{V}'_{j+1} \right) - \bar{V}'_j \circ \bar{V}'_j \right) \right] + 2H \|\bar{V}'_h\|_{\infty} \end{aligned}$$

where (i) and the last inequality hold by the fact $\|\bar{V}'_h\|_{\infty} \geq \|\bar{V}'_{h+1}\|_{\infty} \geq \dots \geq \|\bar{V}'_H\|_{\infty}$. Further according to basic calculus, we have

$$\begin{aligned} &\sum_{j=h}^H \left\langle b_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \widehat{P}_{i,h}^{\pi^k, V}} \left(\mathbb{E}_{\pi \sim \hat{\xi}} \left[\bar{V}'_{i,j+1} \right] \right) \right\rangle \\ &= \sum_{j=h}^H \left[\left(b_h^{j+1} \right)^{\top} \left(\bar{V}'_{j+1} \circ \bar{V}'_{j+1} \right) - \left(b_h^j \right)^{\top} \left(\bar{V}'_j \circ \bar{V}'_j \right) \right] + 2H \|\bar{V}'_h\|_{\infty} \\ &\leq \|b_h^{H+1}\|_1 \|\bar{V}'_{H+1} \circ \bar{V}'_{H+1}\|_{\infty} + 2H \|\bar{V}'_h\|_{\infty} \\ &\leq 3H \|\bar{V}'_h\|_{\infty}. \end{aligned} \tag{144}$$

C.4.7 Proof of Lemma 10

To prove the inequality involving $P_{i,h}^{\widehat{\pi}^k, V} V$ and $\widehat{P}_{i,h}^{\widehat{\pi}^k, V} V$, we start by analyzing the absolute difference between these terms:

$$\begin{aligned} \left| P_{i,h,s,a_i}^{\widehat{\pi}^k, V} V - \widehat{P}_{i,h,s,a_i}^{\widehat{\pi}^k, V} V \right| &= \left| \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i} \left(P_{h,s,a_i}^{\widehat{\pi}^k} \right)} \mathcal{P}V - \inf_{\mathcal{P} \in \mathcal{U}^{\sigma_i} \left(\widehat{P}_{i,h,s,a_i}^{\widehat{\pi}^k} \right)} \mathcal{P}V \right| \\ &\stackrel{(i)}{=} \left| \mathbb{E}_{a_i \sim \widehat{\pi}_{i,h}^k} \max_{\alpha \in [\min_s V(s), \max_s V(s)]} \left[P_{h,s,a_i}^{\widehat{\pi}^k} [V]_{\alpha} - \sigma_i \left(\alpha - \min_{s'} [V]_{\alpha}(s') \right) \right] \right| \end{aligned}$$

$$\begin{aligned}
& - \mathbb{E}_{a_i \sim \hat{\pi}_{i,h}^k} \max_{\alpha \in [\min_s V(s), \max_s V(s)]} \left| \hat{P}_{i,h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha - \sigma_i \left(\alpha - \min_{s'} [V]_\alpha(s') \right) \right| \\
& \leq \mathbb{E}_{a_i \sim \hat{\pi}_{i,h}^k} \max_{\alpha \in [\min_s V(s), \max_s V(s)]} \left| P_{h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha - \hat{P}_{i,h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha \right|, \tag{145}
\end{aligned}$$

where (i) follows from applying the robust Bellman equation (13), and the last inequality uses the fact that the maximum operator is 1-Lipschitz.

Next, we apply Bernstein's inequality to bound the difference between $P_{h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha$ and $\hat{P}_{i,h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha$ for fixed α , k , and (s, a_i) . With probability at least $1 - \delta$, we have:

$$\left| P_{h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha - \hat{P}_{i,h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha \right| \leq \sqrt{\frac{2 \log \left(\frac{2}{\delta} \right)}{N}} \sqrt{\text{Var}_{P_{h,s,a_i}^{\hat{\pi}_{i,h}^k}} ([V]_\alpha)} + \frac{2H \log \left(\frac{2}{\delta} \right)}{3N}. \tag{146}$$

To extend this bound to all (s, a_i) , we use a uniform bound over an ε_1 -net for α . The net size $|N_{\varepsilon_1}| \leq \frac{3H}{\varepsilon_1}$ allows us to apply the union bound:

$$\begin{aligned}
\left| P_{h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha - \hat{P}_{i,h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha \right| & \leq \max_{\alpha \in N_{\varepsilon_1}} \left| P_{h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha - \hat{P}_{i,h,s,a_i}^{\hat{\pi}_{i,h}^k} [V]_\alpha \right| + \varepsilon_1 \\
& \leq \sqrt{\frac{2 \log \left(\frac{2S \sum_{i=1}^n A_i |N_{\varepsilon_1}| Kn}{\delta} \right)}{N}} \sqrt{\text{Var}_{P_{h,s,a_i}^{\hat{\pi}_{i,h}^k}} (V)} + \frac{2H \log \left(\frac{2S \sum_{i=1}^n A_i |N_{\varepsilon_1}| Kn}{\delta} \right)}{3N} + \varepsilon_1 \\
& \leq \sqrt{\frac{2 \log \left(\frac{2S \sum_{i=1}^n A_i NK n}{\delta} \right)}{N}} \sqrt{\text{Var}_{P_{h,s,a_i}^{\hat{\pi}_{i,h}^k}} (V)} + \frac{H \log \left(\frac{2S \sum_{i=1}^n A_i NK n}{\delta} \right)}{N}, \tag{147}
\end{aligned}$$

where the last steps use that $\varepsilon_1 = \frac{H \log \left(\frac{2S \sum_{i=1}^n A_i NK n}{\delta} \right)}{3N}$ and $|N_{\varepsilon_1}| \leq 9N$.

Inserting this back into (145) gives:

$$\begin{aligned}
\left| P_{i,h}^{\hat{\pi}_{i,h}^k, V} - \hat{P}_{i,h}^{\hat{\pi}_{i,h}^k, V} \right| & \leq \sqrt{\frac{2 \log \left(\frac{2S \sum_{i=1}^n A_i NK n}{\delta} \right)}{N}} \sqrt{\text{Var}_{P_h^{\hat{\pi}_{i,h}^k}} (V)} + \frac{H \log \left(\frac{2S \sum_{i=1}^n A_i NK n}{\delta} \right)}{N} 1 \\
& \leq 3 \sqrt{\frac{H^2 \log \left(\frac{2S \sum_{i=1}^n A_i NK n}{\delta} \right)}{N}} 1.
\end{aligned}$$

This completes the proof by showing that the bound holds uniformly over all $(s, a_i) \in \mathcal{S} \times \mathcal{A}_i$.

C.4.8 Proof of Lemma 11

Before proving Lemma 11, we first state a modified version of the Freedman inequality for martingales, which is crucial for our analysis.

Theorem 5. *Suppose $Y_n = \sum_{k=1}^n X_k \in \mathbb{R}$, where $\{X_k\}$ is a real-valued scalar sequence such that*

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

for some constant $R > 0$. Define

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

where \mathbb{E}_{k-1} denotes the conditional expectation given $\{X_j\}_{j < k}$. For any $\kappa > 0$, with probability at least $1 - \delta$, the following holds:

$$|Y_n| \leq \sqrt{8W_n \log \frac{3n}{\delta}} + 5R \log \frac{3n}{\delta} \leq \kappa W_n + \left(\frac{2}{\kappa} + 5R \right) \log \frac{3n}{\delta}. \tag{148}$$

Proof. Suppose deterministically that $W_n \leq \sigma^2$ for some σ^2 . According to Li et al. (2024), with probability at least $1 - \delta$, we have

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

for any positive integer $K \geq 1$. Utilizing the trivial bound $W_n \leq nR^2$, set $\sigma^2 = nR^2$ and $K = \log_2 n$. Then:

$$\begin{aligned} |Y_n| &\leq \sqrt{8 \max \{W_n, R^2\} \log \frac{4 \log_2 n}{\delta}} + \frac{4}{3} R \log \frac{4 \log_2 n}{\delta} \\ &\leq \sqrt{8W_n \log \frac{3n}{\delta}} + \sqrt{8R^2 \log \frac{3n}{\delta}} + \frac{4}{3} R \log \frac{3n}{\delta} \\ &\leq \sqrt{8W_n \log \frac{3n}{\delta}} + 5R \log \frac{3n}{\delta}, \end{aligned}$$

where we used $4 \log_2 n \leq 3n$ for any integer $n \geq 1$. This establishes the first inequality in (148). The second inequality follows from the elementary inequality $2ab \leq a^2 + b^2$. \square

To prove Lemma 11, we apply Lemma 5. Define

$$R := \max_{k \in [K]} |\alpha_k^K \langle \widehat{\pi}_{i,h}^k(s), r_{i,h}^k(s, \cdot) \rangle| \leq \left\{ \max_{k \in [K]} \alpha_k^K \right\} \left\{ \max_{k \in [K]} \|\widehat{\pi}_{i,h}^k(s)\|_1 \|r_{i,h}^k\|_\infty \right\} \leq \frac{2c_\alpha \log K}{K},$$

where the first line uses Lemma 19. We further define

$$\begin{aligned} W_K &= \sum_{k=1}^K (\alpha_k^K)^2 \text{Var}_{h,k-1} (\langle \widehat{\pi}_{i,h}^k(s), r_{i,h}^k(s, \cdot) \rangle) \\ &\leq \left\{ \max_{k \in [K]} \alpha_k^K \right\} \left\{ \sum_{k=1}^K \alpha_k^K \text{Var}_{h,k-1} (\langle \widehat{\pi}_{i,h}^k(s), r_{i,h}^k(s, \cdot) \rangle) \right\} \\ &\leq \frac{2c_\alpha \log K}{K} \sum_{k=1}^K \alpha_k^K \text{Var}_{h,k-1} (r_{i,h}^k(s)), \end{aligned}$$

where we use variance operator $\text{Var}_{h,k-1}[\cdot]$ to denote the variance conditional on what happens before the beginning of the k -th round of data collection for step h . Applying Freedman's inequality (Lemma 5) with $\kappa_1 = \sqrt{K \log(K/\delta)}$, we obtain

$$\begin{aligned} &\left| \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \widehat{\pi}_{i,h}^k} r_{i,h}^k(s, a_i) - \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \widehat{\pi}_{i,h}^k} r_{i,h}^k(s, a_i) \right| \\ &\leq \kappa_1 W_K + \left(\frac{2}{\kappa_1} + 5R_1 \right) \log \left(\frac{3K}{\delta} \right) \\ &\leq 2c_\alpha \sqrt{\frac{\log^3(K/\delta)}{K}} \sum_{k=1}^K \alpha_k^K \text{Var}_{h,k-1} (r_{i,h}^k(s)) + \left(2\sqrt{\frac{1}{K \log(K/\delta)}} + \frac{10c_\alpha \log K}{K} \right) \log \left(\frac{3K}{\delta} \right) \\ &\leq 2c_\alpha \sqrt{\frac{\log^3(K/\delta)}{K}} \sum_{k=1}^K \alpha_k^K \text{Var}_{h,k-1} (r_{i,h}^k(s)) + 4\sqrt{\frac{\log(3K/\delta)}{K}}, \end{aligned}$$

with probability at least $1 - \delta$. Taking a union bound over all $s \in \mathcal{S}$, there exists an absolute constant c_r such that

$$\sum_{k=1}^K \alpha_k^K \left| r_{i,h}^k - \bar{r}_{i,h}^k \right| \leq c_r \sqrt{\frac{\log(KS/\delta)}{K}}.$$

C.4.9 Proof of Lemma 13

In this section, we want to take the accessible range of the robust value function $\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi]$ into consideration when controlling $\sum_{j=h}^H \left\langle d_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,j+1}^\pi] \right) \right\rangle$. Towards this, we introduce some auxiliary values and reward functions as below.

Definition 4. For any time step $h \in [H]$ and the i -th agent, we denote $\bar{V}_h^{\min} := \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi(s)]$ as the minimum value of all the entries in vector $\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi]$. We further define $\bar{V}'_h := \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] - \bar{V}_h^{\min} \mathbf{1}$ as the truncated value function. Eventually for reward function, we define $\bar{r}_{i,h}^{\min} = \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \hat{\pi}_{i,h}^k} r_{i,h}^k(\cdot, a_i) + (\bar{V}_{h+1}^{\min} - \bar{V}_h^{\min}) \mathbf{1}$ as the truncated reward function..

With above notation, we introduce the following fact of \bar{V}'_h :

$$\begin{aligned}
\bar{V}'_h &= \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi] - \bar{V}_h^{\min} \mathbf{1} \\
&\stackrel{(i)}{=} \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \hat{\pi}_{i,h}^k} r_{i,h}^k(\cdot, a_i) + \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] - \bar{V}_h^{\min} \mathbf{1} \\
&= \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \hat{\pi}_{i,h}^k} r_{i,h}^k(\cdot, a_i) + \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \\
&\quad + \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] - \bar{V}_h^{\min} \mathbf{1} \\
&= \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \hat{\pi}_{i,h}^k} r_{i,h}^k(\cdot, a_i) + (\bar{V}_{h+1}^{\min} - \bar{V}_h^{\min}) \mathbf{1} \\
&\quad + \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \bar{V}'_{h+1} + \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \\
&= \bar{r}_{i,h}^{\min} + \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \bar{V}'_{h+1} + \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi], \tag{149}
\end{aligned}$$

where (i) holds by the robust Bellman's consistency equation of $\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h}^\pi]$. With the above fact in hand, we control $\text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \right)$ as follows:

$$\begin{aligned}
&\text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \right) \\
&\stackrel{(i)}{=} \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}}} \left(\bar{V}'_{h+1} \right) \\
&= \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \bar{V}'_{h+1} \right) \circ \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \bar{V}'_{h+1} \right), \tag{150}
\end{aligned}$$

where (i) follows from the fact that $\text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}}} (V - b \mathbf{1}) = \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}}} (V)$ for any value vector $V \in \mathbb{R}^S$ and scalar b . Additionally according to (149), we have

$$\begin{aligned}
&\text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \right) \\
&= \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \left(\bar{V}'_h - \bar{r}_{i,h}^{\min} - \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \mathbb{E}_{\pi \sim \zeta} [\bar{V}_{i,h+1}^\pi] \right)^{\circ 2}
\end{aligned}$$

$$\begin{aligned}
&= \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \bar{V}'_h \circ \bar{V}'_h + 2\bar{V}'_h \circ \left(\bar{r}_{i,h}^{\min} + \left(\sum_{k=1}^K \alpha_k^K \hat{\underline{P}}_{i,h}^{\hat{\pi}^k, \bar{V}} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h+1}^\pi \right] \right) \\
&- \left(\bar{r}_{i,h}^{\min} + \left(\sum_{k=1}^K \alpha_k^K \hat{\underline{P}}_{i,h}^{\hat{\pi}^k, \bar{V}} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h+1}^\pi \right] \right)^{\circ 2}
\end{aligned}$$

Furthermore, we have

$$\begin{aligned}
&\text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h+1}^\pi \right] \right) \\
&= \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \bar{V}'_h \circ \bar{V}'_h + 2\bar{V}'_h \circ \left(\bar{r}_{i,h}^{\min} + \left(\sum_{k=1}^K \alpha_k^K \hat{\underline{P}}_{i,h}^{\hat{\pi}^k, \bar{V}} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h+1}^\pi \right] \right) \\
&- \left(\bar{r}_{i,h}^{\min} + \left(\sum_{k=1}^K \alpha_k^K \hat{\underline{P}}_{i,h}^{\hat{\pi}^k, \bar{V}} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h+1}^\pi \right] \right)^{\circ 2} \\
&\stackrel{(i)}{\leq} \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \bar{V}'_h \circ \bar{V}'_h + 2\|\bar{V}'_h\|_\infty \left(1 + \left| \left(\sum_{k=1}^K \alpha_k^K \hat{\underline{P}}_{i,h}^{\hat{\pi}^k, \bar{V}} - \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \right) \mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,h+1}^\pi \right] \right| \right)
\end{aligned} \tag{151}$$

$$\leq \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, \bar{V}} \left(\bar{V}'_{h+1} \circ \bar{V}'_{h+1} \right) - \bar{V}'_h \circ \bar{V}'_h + 2\|\bar{V}'_h\|_\infty + 6\|V'_h\|_\infty \sqrt{\frac{H^2 \log \left(\frac{18S \sum_{i=1}^n A_i n H N K}{\delta} \right)}{N}} 1, \tag{152}$$

holds with probability at least $1 - \delta$, where (i) arises from $r_{i,h}^{\min} \leq \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \hat{\pi}_{i,h}^k} r_{i,h}^k(\cdot, a_i) \leq 1$ since $V_{h+1}^{\min} - V_h^{\min} \leq 0$ by definition, and the last inequality holds by Lemma 10. Finally, combining (152) and the definition of d_h^j in (99), the term of interest can be controlled as

$$\begin{aligned}
&\sum_{j=h}^H \left\langle d_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right) \right\rangle \\
&= \sum_{j=h}^H (d_h^j)^\top \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}} \left(\bar{V}'_{j+1} \circ \bar{V}'_{j+1} \right) - \bar{V}'_j \circ \bar{V}'_j + 2\|\bar{V}'_j\|_\infty + 6\|\bar{V}'_j\|_\infty \sqrt{\frac{H^2 \log \left(\frac{18S \sum_{i=1}^n A_i n H N K}{\delta} \right)}{N}} 1 \right) \\
&\stackrel{(i)}{\leq} \sum_{j=h}^H \left[(d_h^j)^\top \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}} \left(\bar{V}'_{j+1} \circ \bar{V}'_{j+1} \right) - \bar{V}'_j \circ \bar{V}'_j \right) \right] + 2H\|\bar{V}'_h\|_\infty \\
&\quad + 6H^2\|\bar{V}'_h\|_\infty \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i n H N K}{\delta} \right)}{N}}
\end{aligned}$$

where (i) holds by the fact $\|\bar{V}'_h\|_\infty \geq \|\bar{V}'_{h+1}\|_\infty \geq \dots \geq \|\bar{V}'_H\|_\infty$. With further basic calculus, we can finally obtain that

$$\begin{aligned}
&\sum_{j=h}^H \left\langle d_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, \bar{V}}} \left(\mathbb{E}_{\pi \sim \zeta} \left[\bar{V}_{i,j+1}^\pi \right] \right) \right\rangle \\
&= \sum_{j=h}^H \left[(d_h^{j+1})^\top \left(\bar{V}'_{j+1} \circ \bar{V}'_{j+1} \right) - (d_h^j)^\top \left(\bar{V}'_j \circ \bar{V}'_j \right) \right] + 2H\|\bar{V}'_h\|_\infty + 6H^2\|\bar{V}'_h\|_\infty \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i n H N K}{\delta} \right)}{N}} \\
&\leq \|d_h^{H+1}\|_1 \left\| \bar{V}'_{H+1} \circ \bar{V}'_{H+1} \right\|_\infty + 2H\|\bar{V}'_h\|_\infty + 6H^2\|\bar{V}'_h\|_\infty \sqrt{\frac{\log \left(\frac{18S \sum_{i=1}^n A_i n H N K}{\delta} \right)}{N}}
\end{aligned}$$

$$\begin{aligned}
&\leq 3H\|\bar{V}'_h\|_\infty + 6H^2\|\bar{V}'_h\|_\infty \sqrt{\frac{\log\left(\frac{18S\sum_{i=1}^n A_i n H N K}{\delta}\right)}{N}} \\
&= 3H\|\bar{V}'_h\|_\infty \left(1 + 2H\sqrt{\frac{\log\left(\frac{18S\sum_{i=1}^n A_i n H N K}{\delta}\right)}{N}}\right),
\end{aligned}$$

C.4.10 Proof of Lemma 16

Analogous to Appendix C.4.9, we introduce some auxiliary values and reward functions to control

$$\sum_{j=h}^H \left\langle w_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, V}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^\pi] \right) \right\rangle$$

as below: for any time step h and the i -th agent

Definition 5. For any time step $h \in [H]$ and the i -th agent, we denote $V_h^{\min} := \min_{s \in \mathcal{S}} \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi(s)]$ as the minimum value of all the entries in vector $\mathbb{E}_{\pi \sim \zeta} [V_{i,h}^\pi]$. We further define $V'_h := C_{i,h}^\pi - V_h^{\min} \mathbf{1}$ as the truncated value function. Eventually for reward function, we define $r_{i,h}^{\min} = \sum_{k=1}^K \alpha_k^K \mathbb{E}_{a_i \sim \pi_{i,h}^k} r_{i,h}^{\hat{\pi}^k}(\cdot, a_i) + (V_{h+1}^{\min} - V_h^{\min}) \mathbf{1}$ as the truncated reward function..

Then applying the robust Bellman's consistency equation in (13) gives

$$\begin{aligned}
V'_h = \mathbb{E}_{\pi \sim \zeta} [V_{i,h}^{\pi, \sigma_i}] - V_h^{\min} \mathbf{1} &= \sum_{k=1}^K \alpha_k^K r_{i,h}^{\hat{\pi}^k} + \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} \mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^{\pi, \sigma_i}] - V_h^{\min} \mathbf{1} \\
&= \sum_{k=1}^K \alpha_k^K r_{i,h}^{\hat{\pi}^k} + (V_{h+1}^{\min} - V_h^{\min}) \mathbf{1} + \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} V'_{h+1} \tag{153}
\end{aligned}$$

$$= r_{i,h}^{\min} + \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} V'_{h+1}. \tag{154}$$

The above fact leads to

$$\begin{aligned}
&\text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V}} \left(\mathbb{E}_{\pi \sim \zeta} [V_{i,h+1}^{\pi, \sigma_i}] \right) \\
&\stackrel{(i)}{=} \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V}} (V'_{h+1}) \\
&= \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} (V'_{h+1} \circ V'_{h+1}) - \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} V'_{h+1} \right) \circ \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} V'_{h+1} \right) \\
&\stackrel{(ii)}{=} \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} (V'_{h+1} \circ V'_{h+1}) - (V'_h - r_{i,h}^{\min})^{\circ 2} \\
&= \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} (V'_{h+1} \circ V'_{h+1}) - V'_h \circ V'_h + 2V'_h \circ r_{i,h}^{\min} - r_{i,h}^{\min} \circ r_{i,h}^{\min} \\
&\leq \sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V} (V'_{h+1} \circ V'_{h+1}) - V'_h \circ V'_h + 2\|V'_h\|_\infty \mathbf{1}, \tag{155}
\end{aligned}$$

where (i) follows from the fact that $\text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V}} (V - b \mathbf{1}) = \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V}} (V)$ for any value vector $V \in \mathbb{R}^{\mathcal{S}}$ and scalar b , (ii) holds by (154) and (13), and the last inequality arises from $r_{i,h}^{\min} \leq r_{i,h}^\pi \leq 1$ since $V_{h+1}^{\min} - V_h^{\min} \leq 0$ by definition.

Consequently, combining (155) and the definition of w_h^j in (108), we arrive at

$$\begin{aligned}
& \sum_{j=h}^H \left\langle w_h^j, \text{Var}_{\sum_{k=1}^K \alpha_k^K \underline{P}_{i,h}^{\hat{\pi}^k, V}} (\mathbb{E}_{\pi \sim \zeta} [V_{i,j+1}^{\pi, \sigma_i}]) \right\rangle \\
& \leq \sum_{j=h}^H (w_h^j)^\top \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, V} (V'_{j+1} \circ V'_{j+1}) - V'_j \circ V'_j + 2\|V'_h\|_\infty \mathbf{1} \right) \\
& \stackrel{(i)}{\leq} \sum_{j=h}^H \left[(w_h^j)^\top \left(\sum_{k=1}^K \alpha_k^K \underline{P}_{i,j}^{\hat{\pi}^k, V} (V'_{j+1} \circ V'_{j+1}) - V'_j \circ V'_j \right) \right] + 2H\|V'_h\|_\infty \\
& = \sum_{j=h}^H \left[(w_h^{j+1})^\top (V'_{j+1} \circ V'_{j+1}) - (w_h^j)^\top (V'_j \circ V'_j) \right] + 2H\|V'_h\|_\infty \\
& \leq \|w_h^{H+1}\|_1 \|V'_{H+1} \circ V'_{H+1}\|_\infty + 2H\|V'_h\|_\infty \\
& \leq 3H\|V'_h\|_\infty,
\end{aligned} \tag{156}$$

where (i) and the last inequality hold by the fact $\|V'_h\|_\infty \geq \|V'_{h+1}\|_\infty \geq \dots \geq \|V'_H\|_\infty$ and basic calculus.