

Cooperative Online Learning in Stochastic and Adversarial MDPs

Tal Lancewicki*

Tel Aviv University

lancewicki@mail.tau.ac.il

Aviv Rosenberg*

Tel Aviv University

avivros007@gmail.com

Yishay Mansour

Tel Aviv University and Google Research

mansour.yishay@gmail.com

Abstract

We study cooperative online learning in stochastic and adversarial Markov decision process (MDP). That is, in each episode, m agents interact with an MDP simultaneously and share information in order to minimize their individual regret. We consider environments with two types of randomness: *fresh* – where each agent’s trajectory is sampled i.i.d, and *non-fresh* – where the realization is shared by all agents (but each agent’s trajectory is also affected by its own actions). More precisely, with non-fresh randomness the realization of every cost and transition is fixed at the start of each episode, and agents that take the same action in the same state at the same time observe the same cost and next state. We thoroughly analyze all relevant settings, highlight the challenges and differences between the models, and prove nearly-matching regret lower and upper bounds. To our knowledge, we are the first to consider cooperative reinforcement learning (RL) with either non-fresh randomness or in adversarial MDPs.

Keywords: Regret minimization, Adversarial MDPs, Stochastic MDPs, Multi-agent RL, Cooperation, Reinforcement Learning, Online Learning.

1. Introduction

Cooperative multi-agent reinforcement learning (MARL; see [Zhang et al. \(2021a\)](#)) achieved impressive empirical success in many applications such as cyber-physical systems ([Adler and Blue, 2002](#); [Wang et al., 2016](#)), finance ([Lee et al., 2002, 2007](#)) and sensor/communication networks ([Cortes et al., 2004](#); [Choi et al., 2009](#)). Many of the theoretical work on MARL focus on Markov Games (MGs) ([Jin et al., 2021](#)), where agents have a shared state, the transition and cost is affected by all agents’ actions, and usually the goal is to converge to an equilibrium. On the other hand, in cooperative learning ([Lidard et al., 2021](#)), agents do not affect each other and the notion of equilibrium becomes irrelevant. In this model, agents share information with each other and the goal here is to utilize the shared information in order to obtain a significant improvement upon single-agent performance. This model is a generalization of the extensively studied Multi-agent multi-armed bandit (MAB) model ([Cesa-Bianchi et al., 2016](#)) and reveals many new challenges that are unique to MDPs.

In this paper we initiate the study of two topics not addressed before in the Cooperative MARL literature. First, we differentiate between two types of randomness: *fresh* – where each agent’s trajectory is sampled i.i.d, and *non-fresh* – where at any time the cost and transition kernel’s randomness is shared by all agents. More precisely, if at the same time two different agents perform the same action in the same state, they observe the same cost and the same next state. Second, we consider cooperation in the challenging adversarial MDP setting that generalizes stochastic MDPs and allows to model temporal changes in the environment through costs that change arbitrarily and are chosen by an adversary.

While previous works focus mostly on fresh randomness, non-fresh randomness models are just as well-motivated since different agents might experience the same dynamics and rewards when visiting the same state simultaneously. Conceptually, non-fresh randomness models cases where the randomness is more a function of the time than the agent. For example, drones that fly together experience similar weather conditions and autonomous vehicles that drive on the

*. Equal contribution.

Table 1: Summary of our regret upper and lower bounds for m agents facing K episode interaction with an MDP that has S states, A actions and horizon H . The bounds ignore poly-logarithmic factors and lower order terms. (*) The algorithm requires $m = \sqrt{K}$ agents.

Algorithm	Regret	Lower Bound	Randomness	Cost	Transition
coop-ULCVI	$\sqrt{\frac{H^3 SAK}{m}}$	$\sqrt{\frac{H^3 SAK}{m}}$	fresh	stochastic	unknown
coop-O-REPS	$\sqrt{H^2 K} + \sqrt{\frac{H^2 SAK}{m}}$	$\sqrt{H^2 K} + \sqrt{\frac{H^2 SAK}{m}}$	fresh	adversarial	known
coop-UOB-REPS	$\sqrt{H^2 K} + \sqrt{\frac{H^4 S^2 AK}{m}}$	$\sqrt{H^2 K} + \sqrt{\frac{H^3 SAK}{m}}$	fresh	adversarial	unknown
coop-ULCAE	$\sqrt{H^5 SK} + \sqrt{\frac{H^7 SAK}{\sqrt{m}}}$	$\sqrt{H^2 SK} + \sqrt{\frac{H^3 SAK}{m}}$	non-fresh	stochastic	unknown
coop-nf-O-REPS	$\sqrt{H^2 SK} + \sqrt{\frac{H^2 SAK}{m}}$	$\sqrt{H^2 SK} + \sqrt{\frac{H^2 SAK}{m}}$	non-fresh	adversarial	known
coop-nf-UOB-REPS	$\sqrt{H^4 S^2 K}$ (*)	$\sqrt{H^2 SK} + \sqrt{\frac{H^3 SAK}{m}}$	non-fresh	adversarial	unknown

same roads on the same time encounter the same traffic congestion. Moreover, the non-fresh randomness model is theoretically highly challenging, as we show in this paper. We indicate a gap in the lower bounds between fresh and non-fresh randomness and identify the weaknesses of current optimistic approaches in handling non-fresh randomness, thus requiring us to develop new algorithmic techniques.

Our main contributions can be summarized as follows. First, we derive multi-agent versions of known regret minimization algorithms in stochastic and adversarial MDPs, and thoroughly analyze their regret in the fresh randomness model. To complement our bounds, we formally prove matching lower bounds (the adversarial MDP with unknown transitions lower bound nearly-matches, as optimal regret has not been achieved even for a single agent). Second, we point to the failure of optimistic methods under non-fresh randomness and prove lower bounds that reveal a significant gap from the fresh randomness case. Our novel constructions for these lower bounds carefully take advantage of the agents’ shared random seed to make sure that they cannot explore different areas of the environment simultaneously. Third, we develop a novel multi-agent action-elimination based algorithm for stochastic MDP with non-fresh randomness that forces the agents to scatter at a carefully chosen time so that exploration is maximized. Through novel analysis of the relations between the policies of different agents and the error propagation, we prove near-optimal regret for the algorithm. Finally, for adversarial MDP (where action-elimination is not possible) with non-fresh randomness, we design a novel exploration mechanism to replace optimism and show that it can achieve near-optimal regret for a large number of agents. Section 1 summarizes all our regret lower and upper bounds.

1.1 Related Work

Multi-agent multi-armed bandit. Cooperation was previously studied in both stochastic MAB (Dubey et al., 2020; Wang et al., 2020; Madhushani et al., 2021; Landgren et al., 2021) and adversarial MAB (Cesa-Bianchi et al., 2016, 2019; Bar-On and Mansour, 2019; Ito et al., 2020). While we extend some ideas from the MAB literature to RL, many of the challenges that this paper faces do not arise in MAB. Notably, fresh vs. non-fresh randomness, which is the main focus of the paper, is unique to MDPs as it involves dynamics (i.e., state transitions) that do not exist in MAB.

A different line of work studies *collisions* in cooperative MAB (Bubeck et al. (2021)), where agents are penalized by choosing the same action. In these works, the goal of the cooperation is to minimize the number of collisions rather than improve the agents’ performance. Thus, this line of work which has different motivation is only marginally related to our work. See Bubeck et al. (2021) for a more thorough literature review on collisions in MAB.

MARL. There is a long line of research on the theoretical aspects of MARL, mainly focusing on MGs (Littman, 1994; Bai and Jin, 2020; Bai et al., 2020; Xie et al., 2020; Zhang et al., 2020; Liu et al., 2021; Jin et al., 2021). As mentioned before, this literature is only partially related to our model since it aims to converge to an equilibrium rather than minimize individual regret. More related is the literature on decentralized MARL that considers stochastic MDPs with

fresh randomness. However, the theoretical guarantees provided by these works are either asymptotic or less tight than our bounds (Zhang et al., 2018a,b, 2021b; Lidard et al., 2021).

Single-agent RL. There is a rich literature on regret minimization in both stochastic (Jaksch et al., 2010; Azar et al., 2017; Jin et al., 2018, 2020a,b; Yang and Wang, 2019; Zanette et al., 2020a,b) and adversarial (Zimin and Neu, 2013; Rosenberg and Mansour, 2019a,b, 2021; Jin and Luo, 2020; Cai et al., 2020; Luo et al., 2021) MDPs. Note that for a single agent, fresh and non-fresh randomness are identical.

2. Preliminaries

A finite-horizon episodic MDP \mathcal{M} is defined by the tuple $(\mathcal{S}, \mathcal{A}, H, p, \{c^k\}_{k=1}^K)$. \mathcal{S} (of size S) and \mathcal{A} (of size A) are finite state and action spaces, H is the horizon and K is the number of episodes. p is a transition function such that the probability to move to state s' when taking action a in state s at time h is $p_h(s'|s, a)$. $c^k \in [0, 1]^{HSA}$ is the cost function for episode k . In the *adversarial* setting the sequence of cost functions $\{c^k\}_{k=1}^K$ is chosen by an oblivious adversary before the interaction starts, while in the *stochastic* setting the costs are sampled i.i.d from a stationary distribution (that does not depend on k) with mean $c_h^k(s, a) = c_h(s, a)$. The adversarial MDP model generalizes stochastic MDPs.

A policy π is a function such that $\pi_h(a|s)$ gives the probability to take action a in state s at time h . If π is deterministic we often abuse notation and use $\pi_h(s)$ for the action chosen by the policy. Given a cost function c , the value $V_h^\pi(s)$ of π is the expected cost when starting from state s at time h , i.e., $V_h^\pi(s) = \mathbb{E}^{p, \pi}[\sum_{h'=h}^H c_{h'}(s_{h'}, a_{h'}) | s_h = s]$ where the notation $\mathbb{E}^{p, \pi}[\cdot]$ means that actions are chosen by π and transitions are determined by p . We also define the Q -function $Q_h^\pi(s, a) = \mathbb{E}^{p, \pi}[\sum_{h'=h}^H c_{h'}(s_{h'}, a_{h'}) | s_h = s, a_h = a]$ that satisfies the Bellman equations (Sutton and Barto, 2018):

$$Q_h^\pi(s, a) = c_h(s, a) + \mathbb{E}_{p_h(\cdot|s, a)}[V_{h+1}^\pi] \quad ; \quad V_h^\pi(s) = \langle \pi_h(\cdot | s), Q_h^\pi(s, \cdot) \rangle,$$

where $\mathbb{E}_{r(\cdot)}[f]$ denotes the expectation of $f(x)$ where x is sampled from the distribution r , and $\langle \cdot, \cdot \rangle$ is the dot product.

Multi-agent interaction. A team of m agents interacts with the MDP \mathcal{M} . At the beginning of episode k , every agent $v \in [m]$ picks a policy $\pi^{k, v}$ and starts in the initial state $s_1^{k, v} = s_{\text{init}}$. At time $h = 1, \dots, H$, each agent observes its current state $s_h^{k, v}$ and samples an action $a_h^{k, v} \sim \pi_h^{k, v}(\cdot | s_h^{k, v})$. In the *fresh randomness* model, the next state is sampled independently for each agent, i.e., $s_{h+1}^{k, v} \sim p_h(\cdot | s_h^{k, v}, a_h^{k, v})$. For *non-fresh randomness*, the next state is sampled once for each state-action pair $S_h^k(s, a) \sim p_h(\cdot | s, a)$ ahead of the episode, and then every agent v that takes action a in s at time h transitions to the same state $S_h^k(s, a)$, i.e., $s_{h+1}^{k, v} = S_h^k(s_h^{k, v}, a_h^{k, v})$. Similarly, the cost $C_h^{k, v}$ suffered by the agent is either sampled independently when randomness is fresh, or sampled once for each state-action pair (s, a) ahead of the episode when randomness is non-fresh. Note that for adversarial MDPs the costs are not stochastic so it is always the case that $C_h^{k, v} = c_h^k(s_h^{k, v}, a_h^{k, v})$. At the end of the episode, the team observes the trajectories and costs of all agents $\{s_h^{k, v}, a_h^{k, v}, C_h^{k, v}\}_{h=1, v=1}^{H, m}$ (i.e., bandit feedback).

Regret. Let $V^{k, \pi}$ be the value function of π with respect to c^k . The pseudo-regret of an agent v is the cumulative difference between the values of its policies and the values of the best fixed policy in hindsight. The performance of the team is measured by the maximal individual pseudo-regret (which is stronger than the average regret):

$$R_K = \max_{v \in [m]} \sum_{k=1}^K V_1^{k, \pi^{k, v}}(s_{\text{init}}) - \min_{\pi} \sum_{k=1}^K V_1^{k, \pi}(s_{\text{init}}).$$

For stochastic MDP, we use the more common definition of the regret in which, for $k \in [K]$, $V_h^{k, \pi}(s) = V_h^\pi(s)$ for the cost function c (the mean of the costs distribution).

Occupancy measures. For policy π , let q^π be its occupancy measure such that $q_h^\pi(s)$ is the probability to visit s at time h playing π and $q_h^\pi(s, a) = q_h^\pi(s)\pi_h(a | s)$. By definition, $V_1^{k, \pi}(s_{\text{init}}) = \langle q^\pi, c^k \rangle$, so we can write the regret in terms of occupancy measures as follows: $R_K = \max_{v \in [m]} \sum_{k=1}^K \langle q^{\pi^{k, v}}, c^k \rangle - \min_{q \in \Delta(\mathcal{M})} \sum_{k=1}^K \langle q, c^k \rangle$, where $\Delta(\mathcal{M})$ is the set of valid occupancy measures which corresponds to the set of stochastic policies and is a convex polytope in \mathbb{R}^{HSA} defined by $O(HSA)$ linear inequalities.

Additional notations. The notation $\tilde{O}(\cdot)$ hides constants, lower order terms and poly-logarithmic factors, including $\log(K/\delta)$ for some confidence parameter δ . For $n \in \mathbb{N}$ we denote $[n] = \{1, 2, \dots, n\}$, the indicator of event E is denoted by $\mathbb{I}\{E\}$, and $x \vee y = \max\{x, y\}$. π^* denotes the optimal policy (best in hindsight for the adversarial case).

3. Fresh Randomness

The main principle that guides us in the design of algorithms for fresh randomness is the following: even if all agents play the same policy, the team still gathers “ m times more data”. Thus, we take a single-agent regret minimization algorithm `ALG`, let all the agents play the policy that it outputs, and update it based on the observations of all agents.

For the stochastic setting we propose an optimistic algorithm we call `coop-ULCVI` based on the single-agent algorithms of Azar et al. (2017); Dann et al. (2019). The algorithm maintains empirical estimates of the transition probabilities and costs, based on samples from all agents. At the beginning of episode k it constructs an optimistic estimate \underline{Q}^k of the optimal Q -function Q^* so that with high probability (w.h.p) $\underline{Q}_h^k(s, a) \leq Q_h^*(s, a)$. The agents all play the same deterministic policy which is greedy with respect to \underline{Q}^k : $\pi_h^{k,v}(s) = \arg \max_a \underline{Q}_h^k(s, a)$. Even if agents arrive together at the same state and take the same action, we still get multiple i.i.d samples. Hence, the empirical estimates are based on m times more samples compared to the non-cooperative (single-agent) setting. This is the key property that allows us to prove the following improved regret bound. Detailed description of the `coop-ULCVI` algorithm and the proof of Theorem 1 appear in Appendix B.

Theorem 1 *For stochastic MDP with fresh randomness, `coop-ULCVI` ensures w.h.p $R_K = \tilde{O}(\sqrt{\frac{H^3 SAK}{m}})$.*

This bound improves upon Lidard et al. (2021) by a factor of \sqrt{H} , and is in fact optimal up to logarithmic factors as shown by our lower bound in Appendix A. The lower bound is built on a simple observation (Ito et al., 2020): minimizing the sum of regrets of m agents in K episodes is harder than minimizing the regret of a single agent in mK episodes. By single-agent lower bound (Domingues et al., 2021), the sum of regrets is $\Omega(\sqrt{H^3 SAKm})$, so that the lower bound on the average regret matches our regret bound in Theorem 1.

For the adversarial setting we propose `coop-O-REPS` which is based on the single-agent O-REPS algorithm Zimin and Neu (2013). Essentially, this is Online Mirror Descent (Shalev-Shwartz et al., 2011) on the set of occupancy measures with entropy regularization. More specifically, in episode k all agents play policy π^k computed as follows:

$$q^{\pi^k} = \arg \min_{q \in \Delta(\mathcal{M})} \eta \langle q, \hat{c}^{k-1} \rangle + \text{KL}(q \parallel q^{\pi^{k-1}}), \quad (1)$$

where $\text{KL}(\cdot \parallel \cdot)$ is the KL-divergence, η is a learning rate and \hat{c}^k is an importance sampling estimator. The main difference in our algorithm is the new estimator that incorporates the observations from all the agents as follows:

$$\hat{c}_h^k(s, a) = \frac{c_h^k(s, a) \mathbb{I}\{\exists v : s_h^{k,v} = s, a_h^{k,v} = a\}}{W_h^k(s, a) + \gamma}, \quad \text{where} \quad W_h^k(s, a) = \Pr(\exists v : s_h^{k,v} = s, a_h^{k,v} = a), \quad (2)$$

and γ is a bias added for high-probability regret (Neu, 2015). The quantity $W_h^k(s, a)$ will play a major role in the analysis of all our algorithms. Conveniently, with fresh-randomness $W_h^k(s, a) = 1 - (1 - q_h^{\pi^k}(s, a))^m$ as it is the complement of the event that all agents do not visit (s, a) at time h . Thus, the algorithm can be implemented efficiently similarly to the single-agent algorithm (Zimin and Neu, 2013).

For unknown transitions we propose `coop-UOB-REPS` based on single-agent UOB-REPS (Jin et al., 2020a), which is similar to `coop-O-REPS` but uses an estimate $\Delta(\mathcal{M}, k)$ of the set of occupancy measures which contains the true set $\Delta(\mathcal{M})$ w.h.p. Note that without knowing p , we cannot compute W_h^k . Instead, we use an optimistic estimate U_h^k which upper bounds W_h^k w.h.p. The full algorithms and analysis for the adversarial setting with fresh randomness appear in Appendices D and E.

Theorem 2 *For adversarial MDP with fresh randomness, `coop-O-REPS` ensures $R_K = \tilde{O}(H\sqrt{K} + \sqrt{\frac{H^2 SAK}{m}})$ for known dynamics, and `coop-UOB-REPS` ensures w.h.p $R_K = \tilde{O}(H\sqrt{K} + \sqrt{\frac{H^4 S^2 AK}{m}})$ for unknown dynamics.*

In Appendix A we show these bounds are optimal, except for an extra \sqrt{HS} factor in the second term of our unknown dynamics bound which we cannot hope to remove here since it is still an open problem even for single-agent. Notice the additional $H\sqrt{K}$ term that does not appear in the stochastic setting. It follows from the lower bound for single-agent adversarial MDP with full-information feedback (not bandit feedback), which is equivalent to our setting in the best case scenario where the agents manage to visit all state-actions.

Proof sketch for Theorem 2 By standard analysis, the regret scales with two terms: *penalty* of order H/η , and *stability* of order $\eta \sum_{k,h,s,a} q_h^{\pi^k}(s,a) \hat{c}_h^k(s,a)^2$. We show that the stability (which accounts for the estimator’s variance) decreases as the number of agents increases. In particular we prove that,

$$q_h^{\pi^k}(s,a) \leq \left(\frac{1}{m} + q_h^{\pi^k}(s,a)\right) W_h^k(s,a) \implies q_h^{\pi^k}(s,a) \hat{c}_h^k(s,a)^2 \leq \left(\frac{1}{m} + q_h^{\pi^k}(s,a)\right) \hat{c}_h^k(s,a) \quad (3)$$

Then, by concentration, the stability amounts to $\eta HK(1 + SA/m)$, and optimizing over η gives the desired bound.

With unknown dynamics, there is an additional error term that comes from the estimation of the occupancy measures set and the bias of the estimator (in particular U_h^k). This error is handled similarly to the stochastic case. ■

4. The Challenges of Non-fresh Randomness

Unlike fresh randomness, in the non-fresh randomness setting the total amount of feedback is not necessarily m times the feedback of a single agent. In fact, any algorithm that uses deterministic policies (e.g., optimistic algorithms) simply fails in this setting. The reason is that all agents follow the exact same trajectory since the policy does not introduce any randomness and the transitions are fixed ahead of the episode. Thus, the total amount of feedback is exactly the same as a single-agent would have gathered, which means $\Omega(\sqrt{H^3SAK})$ regret with no benefit from multiple agents.

The next theorem shows that the *non-fresh randomness* setting is significantly harder than fresh randomness in terms of the statistical lower bound. While the regret for stochastic MDP with *fresh randomness* scales only logarithmically with K for large enough m , $H\sqrt{SK}$ regret is unavoidable under *non-fresh randomness* even if $m \rightarrow \infty$.

Theorem 3 For any $S, A, H, m \in \mathbb{N}$ and $K \geq SAH$, and for any algorithm ALG , there exists a stochastic MDP with *non-fresh randomness* such that ALG suffers expected average regret of at least $\Omega(H\sqrt{SK} + \sqrt{\frac{H^3SAK}{m}})$.

Proof sketch We construct the following MDP illustrated in Figure 1. All agents start in s_0 . Taking action a_1 transitions to one of the MAB states s_1, \dots, s_S with probability $1/S$ to each. Taking any other action $a \neq a_1$ transitions to a bad state s_b which is a sink with maximal cost 1. Each MAB state encodes a hard MAB instance: one action gives Bernoulli cost with mean $1/2 + \epsilon$, while the rest of the actions give Bernoulli cost with mean $1/2$. From the MAB states all actions transition back to s_0 with probability 1.




Figure 1: Lower bound construction for non-fresh randomness.

Since the bad state has higher cost than every MAB state and does not contribute to exploration at all, we can assume that all agents choose action a_1 every time they arrive to s_0 . Recall that transitions are non-fresh, so all agents visit exactly the same states. This is the critical point in our construction, as it means that exploration is limited to the A actions in the states that all agents visit (and cannot remove S from the regret).

Choosing ϵ as in standard MAB lower bounds, we get that the regret from each of MAB state is $\Omega(\sqrt{(1 + A/m)X})$, where X is the total number of visits to that state. For last, $X \approx K/S$ w.h.p which implies that the regret from each MAB state is $\Omega(\sqrt{(1 + A/m)K/S})$. Summing over all states and time steps, we get the desired lower bound. We note that this is a simplified version, missing a factor of \sqrt{H} in the second term. For the full construction and other lower bounds, see Appendix A. ■

Algorithm 1 COOP-ULCAE

- 1: **initialize:** $\mathcal{A}_h^0(s) = \mathcal{A}$ for every $s \in \mathcal{S}$ and $h \in [H]$.
 - 2: **for** $k = 1, \dots, K$ **do**
 - 3: **Compute** $\underline{Q}^k, \overline{Q}^k$ based on empirical estimates, and set optimistic policy $\underline{\pi}_h^k(s) \in \arg \min_{a \in \mathcal{A}} \underline{Q}_h^k(s, a)$.
 - 4: **Eliminate** sub-optimal actions: remove a from $\mathcal{A}_h^k(s)$ if $\exists a' \in \mathcal{A}_h^k(s)$ s.t. $\underline{Q}_h^k(s, a) > \overline{Q}_h^k(s, a')$.
 - 5: **Set policies:** $\forall v \in [m]$ sample $h_v \in [H]$ uniformly at random and set: $\pi^{k,v} = \begin{cases} \underline{\pi}^k & \text{with probability } 1 - \epsilon \\ \pi^{k,h_v} & \text{with probability } \epsilon, \end{cases}$
 where $\pi_h^{k,h'} = \underline{\pi}_h^k$ for any $h \neq h'$ and uniform over $\mathcal{A}_h^k(s)$ at (h, s) .
 - 6: Play episode k , observe feedback and update empirical estimates.
 - 7: **end for**
-

In Section 5 we face non-fresh randomness in stochastic MDP and present the COOP-ULCAE algorithm based jointly on optimism and action-elimination. It is important to note that, much like optimistic algorithms, existing RL action elimination algorithms (e.g., Xu et al. (2021)) are deterministic and thus fail in the cooperative non-fresh randomness setting, even though they achieve optimal regret for single-agent. Moreover, naive ways to make these algorithms use stochastic policies that succeed in cooperative MAB, such as uniform exploration of non-eliminated arms, lead to sub-optimal regret in RL because exploration must be controlled more carefully to ensure important states are reached with large enough probability. Hence, we develop a novel exploration method for our algorithm which guarantees that agents can deviate from the optimistic policy and explore potentially optimal actions with minimal effect on the regret.

For adversarial MDP, non-fresh randomness introduces an additional challenge. Due to correlations between the trajectories of different agents, there is no clear and simple relation between $W_h^k(s, a)$ and $q_h^{\pi^k}(s, a)$ as in the fresh randomness setting. In fact, Eq. (3) does not hold anymore. In Section 6.1 we present a sophisticated technique for bounding the ratio $q_h^{\pi^k}(s, a)/W_h^k(s, a)$ through a Linear Programming formulation. This allows us to prove optimal regret bounds for adversarial cost and known dynamics.

Existing algorithms for adversarial cost and *unknown* dynamics are optimistic in essence, and as mentioned before, such algorithms fail to utilize cooperation under non-fresh randomness. In Section 6.2 we overcome this challenge with a novel exploration mechanism and prove regret that does not depend on A (up to logarithmic factors) if there are at least \sqrt{K} agents. However, finding the optimal regret for general m still remains an important open question.

5. Non-fresh Randomness - Stochastic MDP

As outlined in Section 4, under non-fresh randomness, we cannot let all agents play the optimistic policy as in the fresh randomness case. Hence, we want agents to occasionally deviate from the optimistic policy for the purpose of exploration. A naive approach would be to let agents explore a random action with probability ϵ and to follow the optimistic policy with probability $1 - \epsilon$. That way, we get $m\epsilon$ more feedback for $m \geq 1/\epsilon$ and the regret of playing the optimistic policy would scale as $\sqrt{SAK/(m\epsilon)}$ (ignoring dependency in H). On the other hand, deviating with an arbitrary action can lead to cost of order of H , which happens for approximately ϵK episodes, so one would have to set $\epsilon \leq \sqrt{SA/K}$ in order to obtain improvement over single agent regret. Thus, the number of agents must be at least $m \geq 1/\epsilon \geq \sqrt{K/(SA)}$ for an improvement. In this section we significantly reduce the number of agents required for a gain in the regret, and show that it can depend on A alone.

Another natural approach, which leads to optimal regret in cooperative MAB, is *action-elimination*, i.e., eliminate all actions that are clearly sub-optimal and explore uniformly at random over non-eliminated actions. However, this approach would also fail in RL because it does not explore efficiently enough. More precisely, agents deviate too much from the optimistic policy so we cannot guarantee that they visit important states. A closer look at action-elimination algorithms for RL (Xu et al., 2021) reveals that they use deterministic policies for this very reason.

Our algorithm, cooperative upper lower confidence action-elimination (COOP-ULCAE), is presented in Algorithm 1 and in its full version (together with the full analysis) in Appendix C. It takes inspiration from the two previous approaches but utilizes multi-agent exploration in a nearly optimal way. It explores only over non-eliminated actions,

but also makes sure that deviation from the optimistic policy is minimal, thus avoiding “non-important” states. This is achieved by playing a random non-eliminated action only at one step during the episode, selected uniformly at random.

Formally, the algorithm maintains a set of *active actions* in each state $\mathcal{A}_h^k(s)$, consisting of only potentially-optimal actions. In episode k , it computes optimistic and pessimistic estimates of Q^* , \underline{Q}^k and \overline{Q}^k , respectively, such that w.h.p $\underline{Q}_h^k(s, a) \leq Q_h^*(s, a) \leq \overline{Q}_h^k(s, a)$. Hence, if for actions a and a' , $\underline{Q}_h^k(s, a) > \overline{Q}_h^k(s, a')$, then a is clearly sub-optimal and we can eliminate it. The policies of the agents are determined as follows: agent v plays the optimistic policy (greedy with respect to \underline{Q}^k) with probability $1 - \epsilon$, and with probability ϵ she plays the optimistic policy except for one random time step h_v where she takes a uniformly random active action. The key idea is that deviating on a single time step with an active arm would have only minor affect on the regret, so we can set ϵ much larger compared to the naive ϵ -exploration approach described in the beginning of this section.

Theorem 4 *For stochastic MDP with non-fresh randomness, coop-ULCAE ensures with high probability*

$$R_K = \tilde{O}\left(\sqrt{H^5 SK} + \sqrt{\frac{H^7 SAK}{\sqrt{m}}}\right).$$

If $m \geq H^4 A^2$, the first term is dominant, in which case the regret is nearly optimal and matches our lower bound (Theorem 3) up to $H^{3/2}$. Otherwise, we have optimal dependence in S, A, K but there is still a gap of H^2 and more importantly $1/\sqrt[4]{m}$. Determining the optimal dependency in m for this setting is an important open question.

Proof sketch for Theorem 4 To simplify presentation, we ignore $\text{poly}(H)$ factors in the proof sketch and use the notation $V^\pi = V_1^\pi(s_{\text{init}})$. For agent v , we first break the regret into episodes in which she plays the optimistic policy $\mathcal{K}_{\text{OP}}^v$ and episodes in which she plays an exploration policy $\mathcal{K}_{\text{EXP}}^v$: $\sum_{k \in \mathcal{K}_{\text{OP}}^v} V^{\pi^k} - V^{\pi^*} + \sum_{k \in \mathcal{K}_{\text{EXP}}^v} V^{\pi^{k, h_v}} - V^{\pi^*}$.

Then, we show that the regret of playing π^k is bounded by the difference between the optimistic and pessimistic estimates of Q^* over the trajectory of π^k . This difference shrinks with the confidence radius and mainly scales as,

$$R_{\text{OP}}^v := \sum_{k \in \mathcal{K}_{\text{OP}}^v} V^{\pi^k} - V^{\pi^*} \lesssim \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \frac{q_h^{\pi^k}(s, a)}{\sqrt{n_h^k(s, a) \vee 1}}, \quad (4)$$

where $n_h^k(s, a) = \sum_{j=1}^{k-1} \mathbb{I}\{\exists v : s_h^{j, v} = s, a_h^{j, v} = a\}$ is the number of times some agent visited (s, a) at time h before episode k . Now, one can show that $n_h^k(s, a)$ is approximately the sum of probabilities that *some* agent visits (h, s, a) , i.e., $n_h^k(s, a) \approx \sum_{j=1}^{k-1} W_h^j(s, a)$. Trivially $W_h^j(s, a) \geq q_h^{\pi^j}(s, a)$, but we can further utilize the exploration of all the agents to bound $W_h^j(s, a)$ in terms of $q_h^{\pi^j}(s)$ and not $q_h^{\pi^j}(s, a)$, as follows: With probability $1 - (1 - \epsilon/H A)^m$ some agent plays the policy $\pi^{j, h}$ and takes action a at time h . In that case, she would arrive to s in time h with probability $q_h^{\pi^{j, h}}(s)$. Recall that $\pi^{j, h}$ and π^j are identical up to time h , and so $q_h^{\pi^{j, h}}(s) = q_h^{\pi^j}(s)$. Also, it is possible to show that $1 - (1 - \epsilon/A)^m \approx m\epsilon/A$ whenever $\epsilon \leq A/m$. Thus, we get the better bound:

$$n_h^k(s, a) \approx \sum_{j=1}^{k-1} W_h^j(s, a) \gtrsim \frac{m\epsilon}{A} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) \stackrel{\text{Eq. (4)}}{\implies} R_{\text{OP}}^v \lesssim \sqrt{\frac{A}{m\epsilon}} \sum_{k, h, s} \frac{\sum_a q_h^{\pi^k}(s, a)}{\sqrt{\sum_{j=1}^{k-1} q_h^{\pi^j}(s)}} \lesssim \sqrt{\frac{SAK}{m\epsilon}}, \quad (5)$$

where the last relation uses $\sum_{a \in \mathcal{A}} q_h^{\pi^k}(s, a) = q_h^{\pi^k}(s)$, the Cauchy–Schwarz inequality, and standard arguments (Rosenberg et al., 2020, Lemma B.18) to bound $\sum_k q_h^{\pi^k}(s) / \sum_{j=1}^{k-1} q_h^{\pi^j}(s) \lesssim \log K$.

For the regret in the rest of the rounds we utilize the fact that when the agent plays an exploration policy, she deviates from the optimistic policy using an active action. Particularly, we show that similar to the regret of the optimistic policy, the regret of the exploration episodes scales with the difference between the optimistic and pessimistic estimates of Q^* over the trajectory of π^{k, h_v} , but with additional penalty due to the deviation which is overall bounded by R_{OP}^v :

$$R_{\text{EXP}}^v := \sum_{k \in \mathcal{K}_{\text{EXP}}^v} V^{\pi^{k, h_v}} - V^{\pi^*} \lesssim R_{\text{OP}}^v + \sum_{k \in \mathcal{K}_{\text{EXP}}^v} \sum_{h, s, a} \frac{q_h^{\pi^{k, h_v}}(s, a)}{\sqrt{n_h^k(s, a) \vee 1}}. \quad (6)$$

While we can bound $n_h^k(s, a)$ as before in terms of $q_h^{\pi^j}(s)$, it can be very different than $q_h^{\pi^{k,h_v}}(s)$. Intuitively, once the agent deviated from the optimal policy, we have a much weaker guarantee on the quality of our confidence sets in the states that she reaches since the cooperative exploration is done over the trajectory of the optimistic policy. Thus, we cannot use similar arguments to the ones we used to bound R_{OP}^v , and in particular Eq. (5). Instead, we only utilize samples gathered by v in $\mathcal{K}_{\text{EXP}}^v$, and bound $n_h^k(s, a) \gtrsim \sum_{j \in \mathcal{K}_{\text{EXP}}^v, j < k} q_h^{\pi^{j,h_v}}(s, a)$. Using the fact that the number of exploration episodes for v is approximately $|\mathcal{K}_{\text{EXP}}^v| \approx \epsilon K$ and standard arguments, the second term in Eq. (6) is bounded by $\sqrt{SAK}\epsilon$. To finish, combine the bounds and set $\epsilon = \min\{\frac{A}{m}, \frac{1}{\sqrt{m}}\}$. ■

6. Non-fresh Randomness - Adversarial MDP

6.1 Known Transitions

Before tackling the most challenging model – adversarial MDP with non-fresh randomness and unknown transitions, we first study the case of known transitions. While some of the challenges that we tackled in Section 5 are alleviated when transitions are known, in this section we face additional challenges that stem from the fact that now an adversary is choosing the sequence of cost functions instead of them being sampled from a fixed distribution.

More precisely, under non-fresh randomness there are strong correlations between the trajectories of different agents. This is in stark contrast to the fresh randomness setting where, by playing the same policy for all agents, we obtained different i.i.d samples which enabled us to prove that our estimator has reduced variance and therefore get an improved regret bound (Theorem 2). The correlations between the agents’ trajectories introduce two main challenges: a statistical challenge and a computational challenge.

The statistical challenge resembles the ones we faced in the stochastic case (Section 5) – whenever the policy is close to being deterministic, the trajectories of the agents are almost identical. However, since costs are adversarial, here we need different techniques that are compatible with adversarial online learning. More formally, for a near-deterministic policy π^k , we have $W_h^k(s, a) \approx q_h^{\pi^k}(s, a)$ which means that there is almost no benefit from the cooperation between the agents. Even though algorithms for adversarial environments inherently choose stochastic policies, it is still unclear a-priori how to bound the ratio $q_h^{\pi^k}(s, a)/W_h^k(s, a)$, except for the trivial bound of 1 which leads to single-agent regret guarantees. Recall that for fresh randomness we bounded this ratio in Eq. (3) which relies on the independence of the agents’ trajectories given the policy π^k . However, this bound no longer holds and we develop a new bound that is suitable to the non-fresh randomness setting and builds on a novel LP formulation.

The computational challenge follows because we no longer have a closed-form expression to compute $W_h^k(s, a)$ like we had under fresh randomness. We propose to solve this challenge by Monte Carlo estimation of $W_h^k(s, a)$ which we show does not damage the final regret guarantees.

For this setting we propose `coop-nf-O-REPS`, presented (together with its analysis) in Appendix F. It follows the same update rule (Eq. (1)) of `coop-O-REPS`, but instead of $W_h^k(s, a)$ (which is now hard to compute) in the definition of the importance sampling estimator (Eq. (2)), it uses a Monte Carlo estimate $\widetilde{W}_h^k(s, a)$. The estimate is computed by simulating the run of multiple agents playing policy π^k over the MDP for $\widetilde{O}(K)$ times and taking the fraction of times where some agent visited the state-action pair (s, a) at time h . The approximation error adds a small bias of order $O(1/\sqrt{K})$ which affects the total regret by only a constant factor. We note that the computational complexity of this algorithm is similar to standard O-REPS-based algorithms which are known to have $\text{poly}(K, H, S, A)$ per-episode computational complexity (Dick et al., 2014).

Theorem 5 *For adversarial MDP with non-fresh randomness and known dynamics, `coop-nf-O-REPS` ensures w.h.p, $R_K = \widetilde{O}(H\sqrt{SK} + \sqrt{\frac{H^2SAK}{m}})$.*

The above regret bound is optimal up logarithmic factors. The lower bound can be found in Appendix A, and features a similar construction to the one in the proof of Theorem 3. Note that in Theorem 3 there is an extra \sqrt{H} factor in the second term, which only appears for unknown dynamics.

Proof sketch for Theorem 5 Similarly to the proof of Theorem 2, the regret scale with the penalty term H/η , and the stability term $\eta \sum_{k,h,s,a} q_h^{\pi^k}(s,a) \hat{c}_h^k(s,a)^2$. Bounding the approximation error of $\widetilde{W}_h^k(s,a)$ by $\gamma/2$ gives us:

$$(\text{stability}) \lesssim \eta \sum_{k,h,s,a} \frac{q_h^{\pi^k}(s,a)}{W_h^k(s,a)} \hat{c}_h^k(s,a) \quad (7)$$

To further bound the right-hand-side we use a standard concentration bound of $\hat{c}_h^k(s,a)$ around $c_h^k(s,a) \leq 1$. It remains to bound the ratio $q_h^{\pi^k}(s,a)/W_h^k(s,a)$ which is our main technical novelty in this proof. Let $M_h^k(s)$ be the random variable that represents the number of agents that arrive at state s in time h and denote $p_i = \Pr[M_h^k(s) = i]$. Let $\mathbb{E}^k[\cdot]$ denote an expectation conditioned on everything that occurred before the start of episode k . By definition,

$$W_h^k(s,a) = \mathbb{E}^k \left[1 - (1 - \pi_h^k(a|s))^{M_h^k(s)} \right] \geq \mathbb{E}^k \left[\frac{M_h^k(s) \pi_h^k(a|s)}{1 + M_h^k(s) \pi_h^k(a|s)} \right] = \sum_{i=0}^m \frac{p_i i \pi_h^k(a|s)}{1 + i \pi_h^k(a|s)}, \quad (8)$$

where the inequality holds deterministically for every realization of $M_h^k(s)$ (see Lemma 41). Note that the expected value of $M_h^k(s)$ is $m q_h^{\pi^k}(s)$, so the right-hand-side of Eq. (8) is bounded by the value of the following linear program:

$$\min_{p_0, \dots, p_m} \sum_{i=0}^m p_i \frac{i \pi_h^k(a|s)}{1 + i \pi_h^k(a|s)} \quad \text{s.t.} \quad \sum_{i=0}^m p_i i = m q_h^k(s) \quad ; \quad \sum_{i=0}^m p_i = 1.$$

Now, we can solve the LP by considering the dual problem (see Lemma 49), and get that $W_h^k(s,a) \geq \frac{m q_h^{\pi^k}(s) \pi_h^k(a|s)}{1 + m \pi_h^k(a|s)}$. Hence, Eq. (7) is bounded by $\eta HSK(1 + A/m)$, and we obtain the claim by optimizing over η . ■

6.2 Unknown Transitions

When facing adversarial MDPs with unknown transitions and non-fresh randomness, we encounter all the challenges presented in previous sections. The combination of these challenges makes this model especially hard from an algorithmic perspective. Specifically, the only way (currently) to obtain regret bounds in adversarial MDPs with unknown dynamics is via optimism. Unfortunately, as discussed in Sections 4 and 5, optimistic methods fail under non-fresh randomness. Moreover, our solution for the stochastic case is based on action-elimination so it cannot be extended to adversarial costs. Instead, in this section we present a novel exploration mechanism which guarantees near-optimal regret for large enough number of agents. Importantly, if we used optimism, regret would not improve even for $m \rightarrow \infty$.

We propose the `coop-nf-UOB-REPS` algorithm, presented in Algorithm 2 and in full version (together with analysis) in Appendix G. Similarly to `coop-O-REPS` and `coop-nf-O-REPS`, it maintains a policy π^k through the O-REPS update rule (Eq. (9)), however unlike the previous algorithms, some agents play a different policy than π^k for the purpose of exploration. We now present the two key features that allow our algorithm to perform efficient exploration in this challenging setting.

First, we equip the algorithm with a novel exploration mechanism: for every (h, a, k) we assign an agent $\sigma(h, a, k)$ to follow π^k up to time h , and then take action a . The rest of the agents follow the policy π^k . This exploration mechanism is motivated by `coop-ULCAE`, but since costs are adversarial, we cannot eliminate actions and thus have much weaker guarantees on the regret of the exploration policies. As a result, we require many agents so that each agent would explore less often. In particular, there are HAK targets to explore and whenever $m \geq \sqrt{K}$ we can choose σ so that each agent performs exploration for at most $HA\sqrt{K}$ episodes. Second, to avoid the complex dependencies between the agents' trajectories, we use a new importance sampling estimator that ignores all agents except for $\sigma(h, a, k)$, i.e.,

$$\hat{c}_h^k(s,a) = \frac{c_h^k(s,a) \mathbb{I}\{s_h^{k,\sigma(h,a,k)} = s\}}{u_h^k(s) + \gamma},$$

where $u_h^k(s) \approx q_h^{\pi^k}(s)$. Notice that this estimator is approximately unbiased (up to γ and approximation errors) since $\mathbb{I}\{a_h^{k,\sigma(h,a,k)} = a\} = 1$. Finally, since we do not know the set occupancy measures under unknown dynamics, we use an approximation $\Delta(\mathcal{M}, k)$ based on empirical estimates.

Algorithm 2 COOP-NF-UOB-REPS

1: **initialize:** define a mapping $\sigma : [H] \times \mathcal{A} \times [K] \rightarrow [m]$.

2: **for** $k = 1, \dots, K$ **do**

3: **Compute** $\pi_h^k(a | s) = q_h^k(s, a)/q_h^k(s)$ for:

$$q^k = \arg \min_{q \in \Delta(\mathcal{M}, k)} \eta \langle q, \hat{c}^{k-1} \rangle + \text{KL}(q \| q^{k-1}). \quad (9)$$

4: **Set policies for agents:** For every (h, a, k) set the policy of agent $v = \sigma(h, a, k)$ to be:

$$\pi_{h'}^{k,v}(a' | s) = \begin{cases} \mathbb{I}\{a' = a\} & h' = h \\ \pi_{h'}^k(a' | s) & h' \neq h, \end{cases}$$

and for the rest of the agents set $\pi^{k,v} = \pi^k$.

5: **Play** episode k , observe feedback, update transition empirical estimates, and compute cost estimator \hat{c}^k .

6: **end for**

Theorem 6 Assume $S \geq A$ and $m \geq \sqrt{K}$. For adversarial MDP with non-fresh randomness and unknown dynamics, coop-nf-o-reps ensures w.h.p, $R_K = \tilde{O}(H^2 S \sqrt{K})$.

The above result shows that optimal regret is attainable up to factor of \sqrt{HS} . Recall that the extra factor \sqrt{HS} compared to the lower bound of Theorem 3 also appear in the state-of-the-art upper bound for single-agent. There is still a significant gap on the number of agents required for optimal regret, and finding the minimal number of agents to ensure such regret still remains an important open problem.

Proof [Proof sketch for Theorem 6] The proof focuses on bounding the regret of the O-REPS policies $\{\pi^k\}_{k=1}^K$, since the number of episodes that each agent does not play these policies is at most $HA\sqrt{K}$, resulting in extra regret of at most $H^2 S \sqrt{K}$. Similarly to the proof of Theorem 5 we need to bound the stability term, but here the unknown dynamics also introduce additional approximation errors. The analysis of the stability term resembles the proof of Theorem 5 but utilizes the fact that $q_h^k(s, a)/u_h^k(s) \lesssim \pi_h^k(a | s)$. The analysis of the approximation errors takes inspiration from the proof of Theorem 4, and shows that it scales with the sum of confidence radius over the trajectory of π^k :

$$(*) = \sum_{h,s,a,k} \frac{H\sqrt{S}q_h^{\pi^k}(s,a)}{\sqrt{n_h^k(s,a)}} \approx \sum_{h,s,a,k} \frac{H\sqrt{S}q_h^{\pi^k}(s,a)}{\sqrt{\sum_{j=1}^{k-1} W_h^j(s,a)}}$$

We then utilize agents' exploration to lower bound $W_h^k(s, a)$. Recall that agent $\sigma(h, s, k)$ follows π^k until time h , so she arrives at s with probability $q_h^{\pi^k}(s)$ and then takes action a deterministically. Hence, $W_h^k(s, a) \geq q_h^{\pi^k}(s)$, yielding:

$$(*) \lesssim \sum_{h,s,k} \frac{H\sqrt{S} \sum_a q_h^{\pi^k}(s,a)}{\sqrt{\sum_{j=1}^{k-1} q_h^{\pi^j}(s)}} \lesssim H^2 S \sqrt{K}.$$

■

7. Conclusions and Future Work

In this paper we studied cooperation in multi-agent RL. We introduced the non-fresh randomness model and characterized its challenges compared to standard fresh randomness. We provided nearly-matching regret lower and upper bounds in all relevant settings, and developed novel techniques for handling different types of randomness in various models.

Our work leaves two important directions for future work. First, our regret bounds for non-fresh randomness with unknown transitions are not tight for both stochastic and adversarial MDPs. Second, we assume agents communicate through a fully-connected graph. Extending our results to general communication graphs (as in MAB) is an interesting future direction that would also require analyzing delayed feedback (Lancewicki et al., 2020; Howson et al., 2021).

Acknowledgments

Acknowledgements This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 882396), by the Israel Science Foundation (grant number 993/17), Tel Aviv University Center for AI and Data Science (TAD), and the Yandex Initiative for Machine Learning at Tel Aviv University.

References

- Jeffrey L Adler and Victor J Blue. A cooperative multi-agent transportation management and route guidance system. *Transportation Research Part C: Emerging Technologies*, 10(5-6):433–454, 2002.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 263–272. JMLR. org, 2017.
- Yu Bai and Chi Jin. Provable self-play algorithms for competitive reinforcement learning. In *International Conference on Machine Learning*, pages 551–560. PMLR, 2020.
- Yu Bai, Chi Jin, and Tiancheng Yu. Near-optimal reinforcement learning with self-play. *Advances in Neural Information Processing Systems*, 33, 2020.
- Yogev Bar-On and Yishay Mansour. Individual regret in cooperative nonstochastic multi-armed bandits. *Advances in Neural Information Processing Systems*, 32:3116–3126, 2019.
- Sébastien Bubeck, Thomas Budzinski, and Mark Sellke. Cooperative and stochastic multi-player multi-armed bandit: Optimal regret with neither communication nor collisions. In *Conference on Learning Theory*, pages 821–822. PMLR, 2021.
- Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pages 1283–1294. PMLR, 2020.
- Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Delay and cooperation in nonstochastic bandits. *Journal of Machine Learning Research*, 20(17):1–38, 2019.
- Nicolò Cesa-Bianchi, Claudio Gentile, Yishay Mansour, and Alberto Minora. Delay and cooperation in nonstochastic bandits. In *Conference on Learning Theory*, pages 605–622. PMLR, 2016.
- Jongeeun Choi, Songhwai Oh, and Roberto Horowitz. Distributed learning and cooperative control for multi-agent systems. *Automatica*, 45(12):2802–2814, 2009.
- Alon Cohen, Yonathan Efroni, Yishay Mansour, and Aviv Rosenberg. Minimax regret for stochastic shortest path. *Advances in Neural Information Processing Systems*, 34, 2021.
- Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on robotics and Automation*, 20(2):243–255, 2004.
- Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying pac and regret: uniform pac bounds for episodic reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5717–5727, 2017.
- Christoph Dann, Lihong Li, Wei Wei, and Emma Brunskill. Policy certificates: Towards accountable reinforcement learning. In *International Conference on Machine Learning*, pages 1507–1516. PMLR, 2019.
- Travis Dick, Andras Gyorgy, and Csaba Szepesvari. Online learning in markov decision processes with changing cost sequences. In *International Conference on Machine Learning*, pages 512–520. PMLR, 2014.

- Omar Darwiche Domingues, Pierre Ménard, Emilie Kaufmann, and Michal Valko. Episodic reinforcement learning in finite mdps: Minimax lower bounds revisited. In *Algorithmic Learning Theory*, pages 578–598. PMLR, 2021.
- Abhimanyu Dubey et al. Cooperative multi-agent bandits with heavy tails. In *International Conference on Machine Learning*, pages 2730–2739. PMLR, 2020.
- Yonathan Efroni, Nadav Merlis, Aadirupa Saha, and Shie Mannor. Confidence-budget matching for sequential budgeted learning. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 2937–2947. PMLR, 2021. URL <http://proceedings.mlr.press/v139/efroni21a.html>.
- Benjamin Howson, Ciara Pike-Burke, and Sarah Filippi. Delayed feedback in episodic reinforcement learning. *arXiv preprint arXiv:2111.07615*, 2021.
- Shinji Ito, Daisuke Hatano, Hanna Sumita, Kei Takemura, Takuro Fukunaga, Naonori Kakimura, and Ken-Ichi Kawarabayashi. Delay and cooperation in nonstochastic linear bandits. *Advances in Neural Information Processing Systems*, 33:4872–4883, 2020.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873, 2018.
- Chi Jin, Tiancheng Jin, Haipeng Luo, Suvrit Sra, and Tiancheng Yu. Learning adversarial markov decision processes with bandit feedback and unknown transition. In *International Conference on Machine Learning*, pages 4860–4869. PMLR, 2020a.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143, 2020b.
- Chi Jin, Qinghua Liu, Yuanhao Wang, and Tiancheng Yu. V-learning—a simple, efficient, decentralized algorithm for multiagent rl. *arXiv preprint arXiv:2110.14555*, 2021.
- Tiancheng Jin and Haipeng Luo. Simultaneously learning stochastic and adversarial episodic mdps with known transition. *Advances in Neural Information Processing Systems*, 33, 2020.
- Tal Lencewicz, Aviv Rosenberg, and Yishay Mansour. Learning adversarial markov decision processes with delayed feedback. *arXiv preprint arXiv:2012.14843*, 2020.
- Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. Distributed cooperative decision making in multi-agent multi-armed bandits. *Automatica*, 125:109445, 2021.
- Jae Won Lee, Byoung-Tak Zhang, et al. Stock trading system using reinforcement learning with cooperative agents. In *Proceedings of the Nineteenth International Conference on Machine Learning*, pages 451–458, 2002.
- Jae Won Lee, Jonghun Park, O Jangmin, Jongwoo Lee, and Euyseok Hong. A multiagent approach to q-learning for daily stock trading. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 37(6): 864–877, 2007.
- Justin Lidard, Udari Madhushani, and Naomi Ehrich Leonard. Provably efficient multi-agent reinforcement learning with fully decentralized communication. *arXiv preprint arXiv:2110.07392*, 2021.
- Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994.
- Qinghua Liu, Tiancheng Yu, Yu Bai, and Chi Jin. A sharp analysis of model-based reinforcement learning with self-play. In *International Conference on Machine Learning*, pages 7001–7010. PMLR, 2021.

- Haipeng Luo, Chen-Yu Wei, and Chung-Wei Lee. Policy optimization in adversarial mdps: Improved exploration via dilated bonuses. *Advances in Neural Information Processing Systems*, 34, 2021.
- Udai Madhushani, Abhimanyu Dubey, Naomi Leonard, and Alex Pentland. One more step towards reality: Cooperative bandits with imperfect communication. *Advances in Neural Information Processing Systems*, 34, 2021.
- Gergely Neu. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. *Advances in Neural Information Processing Systems*, 28:3168–3176, 2015.
- Aviv Rosenberg and Yishay Mansour. Online convex optimization in adversarial markov decision processes. In *International Conference on Machine Learning*, pages 5478–5486. PMLR, 2019a.
- Aviv Rosenberg and Yishay Mansour. Online stochastic shortest path with bandit feedback and unknown transition function. *Advances in Neural Information Processing Systems*, 32:2212–2221, 2019b.
- Aviv Rosenberg and Yishay Mansour. Stochastic shortest path with adversarially changing costs. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 2936–2942. ijcai.org, 2021. doi: 10.24963/ijcai.2021/404. URL <https://doi.org/10.24963/ijcai.2021/404>.
- Aviv Rosenberg, Alon Cohen, Yishay Mansour, and Haim Kaplan. Near-optimal regret bounds for stochastic shortest path. In *International Conference on Machine Learning*, pages 8210–8219. PMLR, 2020.
- Yevgeny Seldin, Peter Bartlett, Koby Crammer, and Yasin Abbasi-Yadkori. Prediction with limited advice and multi-armed bandits with paid observations. In *International Conference on Machine Learning*, pages 280–287. PMLR, 2014.
- Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and trends in Machine Learning*, 4(2):107–194, 2011.
- Lior Shani, Yonathan Efroni, Aviv Rosenberg, and Shie Mannor. Optimistic policy optimization with bandit feedback. In *International Conference on Machine Learning*, pages 8604–8613. PMLR, 2020.
- Matthew Streeter and H Brendan McMahan. Less regret via online conditioning. *arXiv preprint arXiv:1002.4862*, 2010.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Po-An Wang, Alexandre Proutiere, Kaito Ariu, Yassir Jedra, and Alessio Russo. Optimal algorithms for multiplayer multi-armed bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 4120–4129. PMLR, 2020.
- Shiyong Wang, Jiafu Wan, Daqiang Zhang, Di Li, and Chunhua Zhang. Towards smart factory for industry 4.0: a self-organized multi-agent system with big data based feedback and coordination. *Computer networks*, 101:158–168, 2016.
- Qiaomin Xie, Yudong Chen, Zhaoran Wang, and Zhuoran Yang. Learning zero-sum simultaneous-move markov games using function approximation and correlated equilibrium. In *Conference on Learning Theory*, pages 3674–3682. PMLR, 2020.
- Haike Xu, Tengyu Ma, and Simon S. Du. Fine-grained gap-dependent bounds for tabular mdps via adaptive multi-step bootstrap. In Mikhail Belkin and Samory Kpotufe, editors, *Conference on Learning Theory, COLT 2021, 15-19 August 2021, Boulder, Colorado, USA*, volume 134 of *Proceedings of Machine Learning Research*, pages 4438–4472. PMLR, 2021. URL <http://proceedings.mlr.press/v134/xu21a.html>.
- Lin Yang and Mengdi Wang. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, pages 6995–7004. PMLR, 2019.

- Andrea Zanette, David Brandfonbrener, Emma Brunskill, Matteo Pirodda, and Alessandro Lazaric. Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*, pages 1954–1964, 2020a.
- Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pages 10978–10989. PMLR, 2020b.
- Kaiqing Zhang, Zhuoran Yang, and Tamer Basar. Networked multi-agent reinforcement learning in continuous spaces. In *2018 IEEE conference on decision and control (CDC)*, pages 2771–2776. IEEE, 2018a.
- Kaiqing Zhang, Zhuoran Yang, Han Liu, Tong Zhang, and Tamer Basar. Fully decentralized multi-agent reinforcement learning with networked agents. In *International Conference on Machine Learning*, pages 5872–5881. PMLR, 2018b.
- Kaiqing Zhang, Sham Kakade, Tamer Basar, and Lin Yang. Model-based multi-agent rl in zero-sum markov games with near-optimal sample complexity. *Advances in Neural Information Processing Systems*, 33, 2020.
- Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of Reinforcement Learning and Control*, pages 321–384, 2021a.
- Kaiqing Zhang, Zhuoran Yang, Han Liu, Tong Zhang, and Tamer Basar. Finite-sample analysis for decentralized batch multi-agent reinforcement learning with networked agents. *IEEE Transactions on Automatic Control*, 2021b.
- Alexander Zimin and Gergely Neu. Online learning in episodic markovian decision processes by relative entropy policy search. In *Neural Information Processing Systems 26*, 2013.

Additional notations. While in the main paper the notation \lesssim hides lower order terms and logarithmic factors, in the appendix it only hides constant factors, i.e., $x \lesssim y$ if and only if $x = O(y)$. We use the notation $\mathbb{E}^k[\cdot]$ to denote an expectation conditioned on everything that occurred before the beginning of episode k . Furthermore, $\mathbb{E}^k[\cdot | \pi]$ denotes an expectation conditioned on everything that occurred before the beginning of episode k , and when playing episode k using the policy π . $n_h^k(s, a)$ denotes the number of samples we have from (s, a, h) in the beginning of episode k . More precisely, for fresh randomness $n_h^k(s, a) = \sum_{j=1}^{k-1} \sum_{v=1}^m \mathbb{I}\{s_h^{j,v} = s, a_h^{j,v} = a\}$, and for non-fresh randomness $n_h^k(s, a) = \sum_{j=1}^{k-1} \mathbb{I}\{\exists v : s_h^{j,v} = s, a_h^{j,v} = a\}$. Finally, $\text{Var}_{r(\cdot)}[f]$ denotes the variance of $f(x)$ where x is sampled from the distribution r . For transition function p' and policy π , $q_h^{p', \pi}(s, a, s') = \Pr[s_h = s, a_h = a, s_{h+1} = s' | p', \pi]$ denotes its occupancy measure, $q_h^{p', \pi}(s, a) = \sum_{s'} q_h^{p', \pi}(s, a, s')$ and $q_h^{p', \pi}(s) = \sum_a q_h^{p', \pi}(s, a)$. When the transition function is p , we often use the shorter notation $q^\pi = q^{p, \pi}$. π^* denotes the optimal policy (or best in hindsight), V^* denotes its values function and q^* its occupancy measure.

Appendix A. Lower bounds

In this section we provide proofs for the lower bounds that appear in Section 1.

A.1 Fresh randomness

Theorem 7 (Lower bound for stochastic MDP with fresh randomness) *Let $S, A, H, m \in \mathbb{N}$ and $K \geq SAH$. For any algorithm ALG there exists a stochastic MDP \mathcal{M} with fresh randomness such that: (i) \mathcal{M} has $\Theta(S)$ states, $\Theta(A)$ actions and horizon $\Theta(H)$; (ii) Running ALG with m agents for K episodes suffers expected average regret of at least $\Omega(\sqrt{\frac{H^3 SAK}{m}})$.*

Proof The proof is similar to the proof of Ito et al. (2020, Theorem 4). Notice that cooperative regret minimization with m agents for K episodes is harder than single agent regret minimization for Km episodes, because we can solve the second problem using an algorithm for the first problem. Simply let the agents play the first m episodes one by one, feed the feedback to the algorithm and then again let the agents play sequentially. This implies that the cumulative expected regret of the agents is at least $\Omega(\sqrt{H^3 SAKm})$ by standard lower bounds for MDPs (Domingues et al., 2021). Thus, the average expected regret is at least $\Omega(\sqrt{\frac{H^3 SAK}{m}})$. ■

Theorem 8 (Lower bound for adversarial MDP with fresh randomness and known transition) *Let $S, A, H, m \in \mathbb{N}$ and $K \geq SAH$. For any algorithm ALG there exists an adversarial MDP \mathcal{M} with fresh randomness such that: (i) \mathcal{M} has $\Theta(S)$ states, $\Theta(A)$ actions and horizon $\Theta(H)$; (ii) Running ALG with m agents for K episodes, when the transition function is known, suffers expected average regret of at least $\Omega(\sqrt{H^2 K} + \sqrt{\frac{H^2 SAK}{m}})$.*

Proof The lower bound is obtained by a combination of the following two constructions, i.e., with probability 1/2 the MDP has the structure of the first construction and with probability 1/2 of the other one:

1. Similarly to the proof of Theorem 7, cooperative regret minimization with m for K episodes is harder than single agent regret minimization for Km episodes. Invoking the $\Omega(\sqrt{H^2 SAKm})$ lower bound of Zimin and Neu (2013) for adversarial MDP with bandit feedback and known transition, this gives us the $\Omega(\sqrt{\frac{H^2 SAK}{m}})$ lower bound.
2. Cooperative regret minimization with bandit feedback is harder than single agent regret minimization with full-information feedback because the transition function is known so the agents share information only about the cost function (which is fully revealed under full-information feedback). Thus, for the second construction

we can simply use the construction of [Zimin and Neu \(2013\)](#) for the $\Omega(\sqrt{H^2K})$ lower bound of single agent adversarial MDP with known transition and full-information feedback. ■

Theorem 9 (Lower bound for adversarial MDP with fresh randomness and unknown transition) *Let $S, A, H, m \in \mathbb{N}$ and $K \geq SAH$. For any algorithm ALG there exists an adversarial MDP \mathcal{M} with fresh randomness such that: (i) \mathcal{M} has $\Theta(S)$ states, $\Theta(A)$ actions and horizon $\Theta(H)$; (ii) Running ALG with m agents for K episodes, when the transition function is unknown, suffers expected average regret of at least $\Omega(\sqrt{H^2K} + \sqrt{\frac{H^3SAK}{m}})$.*

Proof Similarly to the proof of [Theorem 8](#) we use a combination of two constructions. The first one is the construction from [Theorem 7](#) which gives the $\Omega(\sqrt{\frac{H^3SAK}{m}})$ lower bound, and the second one is the construction from [Theorem 8](#) which gives the $\Omega(\sqrt{H^2K})$ lower bound. ■

A.2 Non-fresh randomness

Theorem 10 (Lower bound for adversarial MDP with non-fresh randomness and known transition) *Let $S, A, H, m \in \mathbb{N}$ and $K \geq SAH$. For any algorithm ALG there exists an adversarial MDP \mathcal{M} with non-fresh randomness such that: (i) \mathcal{M} has $\Theta(S)$ states, $\Theta(A)$ actions and horizon $\Theta(H)$; (ii) Running ALG with m agents for K episodes, when the transition function is known, suffers expected average regret of at least $\Omega(\sqrt{H^2SK} + \sqrt{\frac{H^2SAK}{m}})$.*

Proof Consider the following MDP with horizon $2H$. There are A actions a_1, a_2, \dots, a_A and $S + 2$ states: the initial state s_0 , a bad state s_b and the MAB states s_1, s_2, \dots, s_S . The agent starts in the initial state s_0 where action a_1 transitions to each of the MAB states s_1, \dots, s_S with probability $1/S$, and all the other actions a_2, \dots, a_A transition to the bad state s_b . In the bad state the cost is always 1 and all the actions just stay in it with probability 1. Each MAB state s_i encodes a hard multi-arm bandit problem for each horizon step h . That is, all the actions transition back to the initial state s_0 , but one action (sampled uniformly at random) suffers cost 0 with probability $1/2 + \epsilon$ (and otherwise 1) while the other actions suffer cost 0 with probability $1/2$, where $\epsilon \approx \sqrt{SA/K}$ which is standard for MAB/RL lower bounds.

Without loss of generality we can assume that all of the agents always choose action a_1 in the initial state because otherwise they transition to the bad state and suffer maximal cost. Critically, this means that all of the agents visit the same state in every time step (because of non-fresh randomness).

Denote by $T_{i,h}$ the number of visits to MAB state s_i in step h . We utilize the lower bound for cooperation in multi-arm bandit ([Seldin et al., 2014](#); [Ito et al., 2020](#)) in order to bound the average expected regret from below by

$$\Omega \left(\mathbb{E} \left[\sum_{i=1}^S \sum_{h=1}^H \sqrt{\left(1 + \frac{A}{m}\right) T_{i,2h}} \right] \right) = \Omega \left(SH \sqrt{1 + \frac{A}{m}} \mathbb{E}[\sqrt{X}] \right),$$

for $X \sim \text{Bin}(n = K, p = 1/S)$ because in each even step size $2h$ one of the MAB states is sampled uniformly at random. By [Lemma 13](#), we have $\mathbb{E}[\sqrt{X}] \geq \Omega(\sqrt{np})$ for $n \geq 1/p^2$ which proves the lower bound $\Omega(\sqrt{H^2SK} + \sqrt{\frac{H^2SAK}{m}})$ for $K \geq S^2$. We note that a more involved proof of the lower bound in each state reveals that the standard assumption $K \geq HSA$ is sufficient. For more details see the proof of [Rosenberg et al. \(2020, Theorem 2.7\)](#). ■

Theorem 11 (Lower bound for stochastic MDP with non-fresh randomness) *Let $S, A, H, m \in \mathbb{N}$ and $K \geq SAH$. For any algorithm ALG there exists a stochastic MDP \mathcal{M} with non-fresh randomness such that: (i) \mathcal{M} has $\Theta(S)$ states,*

$\Theta(A)$ actions and horizon $\Theta(H)$; (ii) Running ALG with m agents for K episodes suffers expected average regret of at least $\Omega(\sqrt{H^2SK} + \sqrt{\frac{H^3SAK}{m}})$.

Proof Similarly to the proof of Theorem 8 we use a combination of two constructions. The first one is presented in the rest of the proof and gives the $\Omega(\sqrt{\frac{H^3SAK}{m}})$ lower bound, and the second one is the construction from Theorem 10 which gives the $\Omega(\sqrt{H^2SK})$ lower bound.

Consider the following MDP with horizon $2H + 1$. There are A actions a_1, a_2, \dots, a_A and $2S + 3$ states: the initial state s_0 , a bad state s_b , a good state s_g , the MAB states s_1, s_2, \dots, s_S , and the wait states s_1^w, \dots, s_S^w . The agent starts in the initial state s_0 where action a_1 transitions to each of the wait states s_1^w, \dots, s_S^w with probability $1/S$, and all the other actions a_2, \dots, a_A transition to the bad state s_b . In the bad state the cost is always 1 and all the actions just stay in it with probability 1, while in the good state the cost is always 0 and all the actions just stay in it with probability 1.

For each $i \in [S]$, the pair of states (s_i, s_i^w) encodes a hard multi-arm bandit problem with HA actions and costs either 0 or $\Omega(H)$. In the next paragraph we describe how the MAB problem is encoded, but first notice that this achieves the desired lower bound. In each episode all of the agents visit the same MAB problem and do not obtain any information about the other ones. Thus, similarly to the proof of Theorem 10, we can utilize the lower bound for cooperation in MAB (Seldin et al., 2014) in order to prove the lower bound:

$$\Omega\left(\sum_{i=1}^S H\sqrt{\frac{HA}{m} \cdot \frac{K}{S}}\right) = \Omega\left(\sqrt{\frac{H^3SAK}{m}}\right).$$

Finally, we describe how to encode a hard MAB instance through the pair of state (s_i, s_i^w) . In the wait state s_i^w the action a_1 transitions to s_i with probability 1, and the action a_2 stays in state s_i^w if the step h is at most $H + 2$, otherwise it transitions to the bad state s_b . All the other actions a_3, \dots, a_A always transition to the bad state. In state s_i all the actions transition to the good state s_g with probability $1/2$ and to the bad state s_b with probability $1/2$, except for one action in a specific time step (both sampled uniformly at random) that transition to the good state with probability $1/2 + \epsilon$ (and to the bad state with probability $1/2 - \epsilon$) for some $\epsilon \approx \sqrt{SA/K}$ which is standard for MAB/RL lower bounds.

Notice that this is in fact MAB with HA actions since the learner needs to pick both the right action and the right horizon step. Moreover, the cost is either 0 if the agents is successful in transitioning to the good state, or $\Theta(H)$ if the learner transitions to the bad state (in any case that the good state is not reached, the bad state will be reached before time step $H + 2$). ■

Theorem 12 (Lower bound for adversarial MDP with non-fresh randomness and unknown transition) *Let $S, A, H, m \in \mathbb{N}$ and $K \geq SAH$. For any algorithm ALG there exists an adversarial MDP \mathcal{M} with non-fresh randomness such that: (i) \mathcal{M} has $\Theta(S)$ states, $\Theta(A)$ actions and horizon $\Theta(H)$; (ii) Running ALG with m agents for K episodes, when the transition function is unknown, suffers expected average regret of at least $\Omega(\sqrt{H^2SK} + \sqrt{\frac{H^3SAK}{m}})$.*

Proof Follows immediately from Theorem 11 since adversarial MDPs generalize stochastic MDPs. ■

A.3 Auxiliary lemmas

Lemma 13 *Let $X \sim \text{Bin}(n, p)$ and assume that $n \geq 1/p^2$. Then, $\mathbb{E}[\sqrt{X}] \geq 0.01\sqrt{np}$.*

Proof By Markov inequality we have:

$$\mathbb{E}[\sqrt{X}] \geq \frac{\sqrt{np}}{10} \Pr\left[\sqrt{X} \geq \frac{\sqrt{np}}{10}\right] = \frac{\sqrt{np}}{10} \Pr\left[X \geq \frac{np}{100}\right] = \frac{\sqrt{np}}{10} \left(1 - \Pr\left[X < \frac{np}{100}\right]\right).$$

Thus, it suffices to show that $\Pr\left[X < \frac{np}{100}\right] \leq 9/10$ which follows immediately from Hoeffding inequality and the assumption that $n \geq 1/p^2$. ■

Algorithm 3 COOPERATIVE UPPER LOWER CONFIDENCE VALUE ITERATION (COOP-ULCVI)

- 1: **input:** state space \mathcal{S} , action space \mathcal{A} , horizon H , confidence parameter δ , number of episodes K , number of agents m .
 - 2: **initialize:** $n_h^1(s, a) = 0, n_h^1(s, a, s') = 0, C_h^1(s, a) = 0 \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 3: **for** $k = 1, \dots, K$ **do**
 - 4: set $\hat{p}_h^k(s' | s, a) \leftarrow \frac{n_h^k(s, a, s')}{n_h^k(s, a) \vee 1}, \hat{c}_h^k(s, a) \leftarrow \frac{C_h^k(s, a)}{n_h^k(s, a) \vee 1} \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 5: compute $\{\pi_h^k(s)\}_{s, h}$ via OPTIMISTIC-PESSIMISTIC VALUE ITERATION (Algorithm 4).
 - 6: set $n_h^{k+1}(s, a) \leftarrow n_h^k(s, a), n_h^{k+1}(s, a, s') \leftarrow n_h^k(s, a, s'), C_h^{k+1}(s, a) \leftarrow C_h^k(s, a) \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 7: **for** $v = 1, \dots, m$ **do**
 - 8: observe initial state $s_1^{k, v}$.
 - 9: **for** $h = 1, \dots, H$ **do**
 - 10: pick action $a_h^{k, v} = \pi_h^k(s_h^{k, v})$, suffer cost $C_h^{k, v} \sim c_h(s_h^{k, v}, a_h^{k, v})$ and observe next state $s_{h+1}^{k, v} \sim p_h(\cdot | s_h^{k, v}, a_h^{k, v})$.
 - 11: update $n_h^{k+1}(s_h^{k, v}, a_h^{k, v}) \leftarrow n_h^{k+1}(s_h^{k, v}, a_h^{k, v}) + 1, n_h^{k+1}(s_h^{k, v}, a_h^{k, v}, s_{h+1}^{k, v}) \leftarrow n_h^{k+1}(s_h^{k, v}, a_h^{k, v}, s_{h+1}^{k, v}) + 1$.
 - 12: update $C_h^{k+1}(s_h^{k, v}, a_h^{k, v}) \leftarrow C_h^{k+1}(s_h^{k, v}, a_h^{k, v}) + C_h^{k, v}$.
 - 13: **end for**
 - 14: **end for**
 - 15: **end for**
-

Appendix B. The coop-ULCVI algorithm for stochastic MDPs with fresh randomness

For the setting of stochastic MDPs with fresh randomness we propose the Cooperative Upper Lower Confidence Value Iteration algorithm (COOP-ULCVI; see Algorithm 3). The idea is simple: all the agents run the same optimistic policy, but the estimated costs and transition models are updated based on the trajectories of all of them. Since the randomness is fresh in this setting, we expect the agents to observe m times more information. Next, we prove the following optimal regret bound for COOP-ULCVI.

Theorem 14 *With probability $1 - \delta$, the individual regret of each agent of COOP-ULCVI is*

$$R_K = O\left(\sqrt{\frac{H^3 S A K}{m}} \log \frac{m K H S A}{\delta} + H^3 S^2 A \log^2 \frac{m K H S A}{\delta}\right).$$

B.1 The good event, optimism and pessimism

Define the following events (for $\tau = 3 \log \frac{6 S A H K m}{\delta}$):

$$\begin{aligned} E^c(k) &= \left\{ \forall (s, a, h) : |\hat{c}_h^k(s, a) - c_h(s, a)| \leq \sqrt{\frac{2\tau}{n_h^k(s, a) \vee 1}} \right\} \\ E^p(k) &= \left\{ \forall (s, a, s', h) : |p_h(s' | s, a) - \hat{p}_h^k(s' | s, a)| \leq \sqrt{\frac{2p_h(s' | s, a)\tau}{n_h^k(s, a) \vee 1}} + \frac{2\tau}{n_h^k(s, a) \vee 1} \right\} \\ E^{pv1}(k) &= \left\{ \forall (s, a, h) : |(\hat{p}_h^k(\cdot | s, a) - p_h(\cdot | s, a)) \cdot V_{h+1}^*| \leq \sqrt{\frac{2\text{Var}_{p_h(\cdot | s, a)}(V_{h+1}^*)\tau}{n_h^k(s, a) \vee 1}} + \frac{5H\tau}{n_h^k(s, a) \vee 1} \right\} \\ E^{pv2}(k) &= \left\{ \forall (s, a, h) : \left| \sqrt{\text{Var}_{p_h(\cdot | s, a)}(V_{h+1}^*)} - \sqrt{\text{Var}_{\hat{p}_h^k(\cdot | s, a)}(V_{h+1}^*)} \right| \leq \sqrt{\frac{12H^2\tau}{n_h^k(s, a) \vee 1}} \right\} \end{aligned}$$

Algorithm 4 OPTIMISTIC-PESSIMISTIC VALUE ITERATION

1: **input:** state space \mathcal{S} , action space \mathcal{A} , horizon H , confidence parameter δ , number of episodes K , number of agents m , visit counters n^k , empirical transition function \hat{p}^k , empirical cost function \hat{c}^k .

2: **initialize:** $V_{H+1}^k(s) = \bar{V}_{H+1}^k(s) = 0$ for all $s \in \mathcal{S}$.

3: **for** $h = H, H-1, \dots, 1$ **do**

4: **for** $s \in \mathcal{S}$ **do**

5: **for** $a \in \mathcal{A}$ **do**

6: set the bonus $b_h^k(s, a) = b_h^k(s, a; c) + b_h^k(s, a; p)$ defined as follows (for $\tau = 3 \log \frac{6SAHKm}{\delta}$),

$$b_h^k(s, a; c) = \sqrt{\frac{2\tau}{n_h^k(s, a) \vee 1}}$$

$$b_h^k(s, a; p) = \sqrt{\frac{2\text{Var}_{\hat{p}_h^k(\cdot|s,a)}(V_{h+1}^k)\tau}{n_h^k(s, a) \vee 1}} + \frac{44H^2S\tau}{n_h^k(s, a) \vee 1} + \frac{1}{16H} \mathbb{E}_{\hat{p}_h^k(\cdot|s,a)} \left[\bar{V}_{h+1}^k - V_{h+1}^k \right].$$

7: compute optimistic and pessimistic Q-functions:

$$\underline{Q}_h^k(s, a) = \hat{c}_h^k(s, a) - b_h^k(s, a) + \mathbb{E}_{\hat{p}_h^k(\cdot|s,a)}[V_{h+1}^k]$$

$$\bar{Q}_h^k(s, a) = \hat{c}_h^k(s, a) + b_h^k(s, a) + \mathbb{E}_{\hat{p}_h^k(\cdot|s,a)}[\bar{V}_{h+1}^k].$$

8: **end for**

9: set $\pi_h^k(s) \in \arg \min_{a \in \mathcal{A}} \underline{Q}_h^k(s, a)$.

10: set $\underline{V}_h^k(s) = \max\{\underline{Q}_h^k(s, \pi_h^k(s)), 0\}$, $\bar{V}_h^k(s) = \min\{\bar{Q}_h^k(s, \pi_h^k(s)), H\}$.

11: **end for**

12: **end for**

The basic good event, which is the intersection of the above events, is the one used in Efroni et al. (2021). The following lemma establishes that the good event holds with high probability. The proof is supplied in Efroni et al. (2021, Lemma 13) by applying standard concentration results.

Lemma 15 (The First Good Event) *Let $\mathbb{G}_1 = \cap_{k=1}^K E^c(k) \cap_{k=1}^K E^p(k) \cap_{k=1}^K E^{pv1}(k) \cap_{k=1}^K E^{pv2}(k)$ be the basic good event. It holds that $\Pr(\mathbb{G}_1) \geq 1 - \delta/2$.*

Under the first good event, we can prove that the value is optimistic using standard techniques (similar to Efroni et al. (2021, Lemma 14)).

Lemma 16 (Upper Value Function is Pessimistic, Lower Value Function is Optimistic) *Conditioned on the first good event \mathbb{G}_1 , it holds that $\underline{V}_h^k(s) \leq V_h^*(s) \leq V_h^{\pi^k}(s) \leq \bar{V}_h^k(s)$ for every $k = 1, \dots, K$, $s \in \mathcal{S}$ and $h = 1, \dots, H$.*

Finally, using similar techniques to Efroni et al. (2021, Lemma 21), we can prove an additional high probability event which hold alongside the basic good event \mathbb{G}_1 . To that end, we define the filtration $\{\mathcal{F}^k\}_{k \geq 1}$ as the σ -algebra that contains the information on all observed data until the beginning of episode k (including the initial state of episode k). In addition, we define the filtration $\{\mathcal{F}_h^k\}_{k \geq 1, h \geq 1}$ as the σ -algebra that contains the information on all observed data until step h of episode k (including the h -th state of episode k).

Lemma 17 (The Good Event) Let \mathbb{G}_1 be the event defined in Lemma 15. The second good event is the intersection of two events $\mathbb{G}_2 = E^{OP} \cap E^{\text{Var}}$ defined as follows:

$$E^{OP} = \left\{ \forall h \in [H], v \in [m] : \sum_{k=1}^K \mathbb{E}[\bar{V}_h^k(s_h^{k,v}) - \underline{V}_h^k(s_h^{k,v}) \mid \mathcal{F}_h^k] \leq 18H^2\tau + \left(1 + \frac{1}{2H}\right) \sum_{k=1}^K \bar{V}_h^k(s_h^{k,v}) - \underline{V}_h^k(s_h^{k,v}) \right\}$$

$$E^{\text{Var}} = \left\{ \forall v \in [m] : \sum_{k=1}^K \sum_{h=1}^H \text{Var}_{p_h(\cdot | s_h^{k,v}, a_h^{k,v})}(V_{h+1}^k) \leq 4H^3\tau + 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}[\text{Var}_{p_h(\cdot | s_h^{k,v}, a_h^{k,v})}(V_{h+1}^k) \mid \mathcal{F}^k] \right\}.$$

Then, the good event $\mathbb{G} = \mathbb{G}_1 \cap \mathbb{G}_2$ holds with probability at least $1 - \delta$.

B.2 Proof of Theorem 14

Lemma 18 (Key Recursion Bound) Conditioning on the good event \mathbb{G} , the following bound holds for all $h \in [H]$ and $v \in [m]$.

$$\begin{aligned} \sum_{k=1}^K \bar{V}_h^k(s_h^{k,v}) - \underline{V}_h^k(s_h^{k,v}) &\leq 18H^2\tau + \sum_{k=1}^K \frac{226H^2S\tau}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1} + \sum_{k=1}^K \frac{2\sqrt{2\tau}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} \\ &\quad + \sum_{k=1}^K 2\sqrt{2\tau} \frac{\sqrt{\text{Var}_{p_h(\cdot | s_h^{k,v}, a_h^{k,v})}(V_{h+1}^k)}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} + \left(1 + \frac{1}{2H}\right)^2 \sum_{k=1}^K \bar{V}_{h+1}^k(s_{h+1}^{k,v}) - \underline{V}_{h+1}^k(s_{h+1}^{k,v}). \end{aligned}$$

Proof We bound each of the terms in the sum as follows:

$$\begin{aligned} \bar{V}_h^k(s_h^{k,v}) - \underline{V}_h^k(s_h^{k,v}) &\leq 2b_h^k(s_h^{k,v}, a_h^{k,v}; c) + 2b_h^k(s_h^{k,v}, a_h^{k,v}; p) + \mathbb{E}_{\hat{p}_h^k(\cdot | s_h^{k,v}, a_h^{k,v})}[\bar{V}_{h+1}^k - \underline{V}_{h+1}^k] \\ &= 2b_h^k(s_h^{k,v}, a_h^{k,v}; c) + 2b_h^k(s_h^{k,v}, a_h^{k,v}; p) \\ &\quad + \mathbb{E}_{p_h(\cdot | s_h^{k,v}, a_h^{k,v})}[\bar{V}_{h+1}^k - \underline{V}_{h+1}^k] + (\hat{p}_h^k - p_h)(\cdot | s_h^{k,v}, a_h^{k,v}) \cdot (\bar{V}_{h+1}^k - \underline{V}_{h+1}^k) \\ &\leq 2b_h^k(s_h^{k,v}, a_h^{k,v}; c) + 2b_h^k(s_h^{k,v}, a_h^{k,v}; p) \\ &\quad + \frac{8H^2S\tau}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1} + \left(1 + \frac{1}{4H}\right) \mathbb{E}_{p_h(\cdot | s_h^{k,v}, a_h^{k,v})}[\bar{V}_{h+1}^k - \underline{V}_{h+1}^k], \end{aligned} \tag{10}$$

where the last relation holds by Cohen et al. (2021, Lemma B.13) which upper bounds

$$(\hat{p}_h^k - p_h)(\cdot | s_h^{k,v}, a_h^{k,v}) \cdot (\bar{V}_{h+1}^k - \underline{V}_{h+1}^k) \leq \frac{8H^2S\tau}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1} + \frac{1}{4H} \mathbb{E}_{p_h(\cdot | s_h^{k,v}, a_h^{k,v})}[\bar{V}_{h+1}^k - \underline{V}_{h+1}^k]$$

by setting $\alpha = 4H, C_1 = C_2 = 2$ and bounding $H\tau(2C_2 + \alpha SC_1/2) \leq 8H^2S\tau$ (the assumption of the lemma holds since the event $\cap_k E^p(k)$ holds). Taking the sum over $k \in [K]$ we get that

$$\begin{aligned} \sum_{k=1}^K \bar{V}_h^k(s_h^{k,v}) - \underline{V}_h^k(s_h^{k,v}) &\leq \sum_{k=1}^K 2b_h^k(s_h^{k,v}, a_h^{k,v}; c) + \sum_{k=1}^K 2b_h^k(s_h^{k,v}, a_h^{k,v}; p) + \sum_{k=1}^K \frac{8H^2S\tau}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1} \\ &\quad + \left(1 + \frac{1}{4H}\right) \sum_{k=1}^K \mathbb{E}_{p_h(\cdot | s_h^{k,v}, a_h^{k,v})}[\bar{V}_{h+1}^k - \underline{V}_{h+1}^k]. \end{aligned} \tag{11}$$

The first sum is bounded by definition by

$$\sum_{k=1}^K b_h^k(s_h^{k,v}, a_h^{k,v}; c) \leq \sum_{k=1}^K \sqrt{\frac{2\tau}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}},$$

and the second sum is bounded in [Efroni et al. \(2021, Lemma 24\)](#) by

$$\begin{aligned} \sum_{k=1}^K b_h^k(s_h^{k,v}, a_h^{k,v}) &\leq \sum_{k=1}^K \frac{109H^2 S\tau}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1} + \sqrt{2\tau} \sum_{k=1}^K \frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} \\ &\quad + \frac{1}{8H} \sum_{k=1}^K \mathbb{E}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}[\bar{V}_{h+1}^k - V_{h+1}^k]. \end{aligned}$$

Plugging this into (11) and rearranging the terms we get

$$\begin{aligned} \sum_{k=1}^K \bar{V}_h^k(s_h^{k,v}) - \underline{V}_h^k(s_h^{k,v}) &\leq \sum_{k=1}^K \frac{2\sqrt{2\tau}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} + 2\sqrt{2\tau} \sum_{k=1}^K \frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} \\ &\quad + \sum_{k=1}^k \frac{226H^2 S\tau}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1} + \left(1 + \frac{1}{2H}\right) \sum_{k=1}^K \mathbb{E}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}[\bar{V}_{h+1}^k - \underline{V}_{h+1}^k] \\ &\leq 18H^2\tau + \sum_{k=1}^K \frac{2\sqrt{2\tau}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} + \sum_{k=1}^k \frac{226H^2 S\tau}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1} \\ &\quad + \sum_{k=1}^K 2\sqrt{2\tau} \frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} + \left(1 + \frac{1}{2H}\right)^2 \sum_{k=1}^K \bar{V}_{h+1}^k(s_{h+1}^{k,v}) - \underline{V}_{h+1}^k(s_{h+1}^{k,v}), \end{aligned}$$

where the last inequality follows since the second good event holds. ■

Proof of Theorem 14 Start by conditioning on the good event which holds with probability greater than $1 - \delta$. Applying the optimism-pessimism of the upper and lower value function we get

$$\sum_{k=1}^K V_1^{\pi^k}(s_1^{k,v}) - V_1^*(s_1^{k,v}) \leq \frac{1}{m} \sum_{v=1}^m \sum_{k=1}^K \bar{V}_1^k(s_1^{k,v}) - \underline{V}_1^k(s_1^{k,v}). \quad (12)$$

Iteratively applying [Lemma 18](#) and bounding the exponential growth by $(1 + \frac{1}{2H})^{2H} \leq e \leq 3$, the following upper bound on the cumulative regret is obtained.

$$\begin{aligned} (12) &\leq 54H^2\tau + \frac{1}{m} \sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \frac{678H^2 S\tau}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1} \\ &\quad + \frac{1}{m} \sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \frac{6\sqrt{2\tau}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} + \frac{1}{m} \sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \frac{6\sqrt{2\tau} \text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k})}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v})}}. \quad (13) \end{aligned}$$

We now bound each of the three sums in Eq. (13). We bound the first sum in Eq. (13) via standard analysis as follows:

$$\begin{aligned}
\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \frac{1}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1} &= \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a) \vee 1} \\
&= \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{n_h^k(s, a) \geq m\} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a) \vee 1} \\
&\quad + \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{n_h^k(s, a) < m\} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a) \vee 1} \\
&\leq 2HSA\tau + \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{n_h^k(s, a) < m\} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a) \vee 1} \\
&\leq 2HSA\tau + 2mHSA \leq 4mHSA\tau, \tag{14}
\end{aligned}$$

where the first inequality is by Lemma 19.

The second sum in Eq. (13) is bounded as follows,

$$\begin{aligned}
\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \frac{1}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} &\leq \sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \frac{\mathbb{I}\{n_h^k(s_h^{k,v}, a_h^{k,v}) \geq m\}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} + 2mHSA \\
&\leq \sqrt{\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H 1} \sqrt{\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \frac{\mathbb{I}\{n_h^k(s_h^{k,v}, a_h^{k,v}) \geq m\}}{n_h^k(s_h^{k,v}, a_h^{k,v}) \vee 1}} + 2mHSA \\
&\leq \sqrt{KHm} \sqrt{2HSA\tau} + 24HSA\tau = \sqrt{2mH^2SAK\tau} + 2mHSA,
\end{aligned}$$

where the first inequality is similar to Eq. (14), the second is by Cauchy–Schwarz, and the third is by Lemma 19.

The third sum in Eq. (13) is bounded by applying the Cauchy–Schwarz inequality as follows,

$$\begin{aligned}
\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v})}} &\leq \sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \mathbb{I}\{n_h^k(s_h^{k,v}, a_h^{k,v}) \geq m\} \frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^{k,v}, a_h^{k,v})}} + 2mH^2SA \\
&\leq \sqrt{\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k})} \sqrt{\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \frac{\mathbb{I}\{n_h^k(s_h^{k,v}, a_h^{k,v}) \geq m\}}{n_h^k(s_h^{k,v}, a_h^{k,v})}} + 2mH^2SA \\
&\leq \sqrt{\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k})} \sqrt{2HSA\tau} + 2mH^2SA \\
&\leq \sqrt{2HSA\tau} \sqrt{\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \mathbb{E} \left[\text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k}) \mid \mathcal{F}^k \right]} + 4H^3\tau + 2mH^2SA \\
&\leq \sqrt{2HSA\tau} \sqrt{\sum_{v=1}^m \sum_{k=1}^K \sum_{h=1}^H \mathbb{E} \left[\text{Var}_{p_h(\cdot|s_h^{k,v}, a_h^{k,v})}(V_{h+1}^{\pi^k}) \mid \mathcal{F}^k \right]} + 5mH^2SA\tau \\
&\stackrel{(*)}{=} \sqrt{2HSA\tau} \sqrt{\sum_{v=1}^m \sum_{k=1}^K \mathbb{E} \left[\left(V_1^{\pi^k}(s_1^{k,v}) - \sum_{h=1}^H c_h(s_h^{k,v}, a_h^{k,v}) \right)^2 \mid \mathcal{F}^k \right]} + 5mH^2SA\tau \\
&\leq \sqrt{2HSA\tau} \sqrt{\sum_{v=1}^m \sum_{k=1}^K H^2} + 5mH^2SA\tau \leq \sqrt{2mH^3SAK\tau} + 5mH^2SA\tau.
\end{aligned}$$

where the first inequality is similar to Eq. (14), the third inequality is by Lemma 19, the fourth is by event E^{Var} , and (*) is by the law of total variance (Cohen et al., 2021, Lemma B.14). ■

B.3 Auxiliary lemmas

Lemma 19 *It holds that*

$$\sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{n_h^k(s, a) \geq m\} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a) \vee 1} \leq 2HSA \log(Km).$$

Proof By Rosenberg et al. (2020, Lemma B.18), we have that

$$\begin{aligned} \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{n_h^k(s, a) \geq m\} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a) \vee 1} &\leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} 2 \log \left(\sum_{k=1}^K \sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\} \right) \\ &\leq 2HSA \log(Km). \end{aligned}$$

■

Algorithm 5 COOPERATIVE UPPER LOWER CONFIDENCE ACTION ELIMINATION (COOP-ULCAE)

- 1: **input:** state space \mathcal{S} , action space \mathcal{A} , horizon H , confidence parameter δ , number of episodes K , number of agents m , exploration parameter $\epsilon > 0$.
 - 2: **initialize:** $n_h^1(s, a) = 0, n_h^1(s, a, s') = 0, C_h^1(s, a) = 0, \mathcal{A}_h^0(s) = \mathcal{A} \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 3: **for** $k = 1, \dots, K$ **do**
 - 4: set $\hat{p}_h^k(s' | s, a) \leftarrow \frac{n_h^k(s, a, s')}{n_h^k(s, a) \vee 1}, \hat{c}_h^k(s, a) \leftarrow \frac{C_h^k(s, a)}{n_h^k(s, a) \vee 1} \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 5: compute $\{\underline{\pi}_h^k(s)\}_{s, h}$ via OPTIMISTIC-PESSIMISTIC VALUE ITERATION (Algorithm 4).
 - 6: set $\mathcal{A}_h^k(s) \leftarrow \mathcal{A}_h^{k-1}(s)$ for every s, h .
 - 7: remove sub-optimal actions for every s, h : if $\exists a, a' \in \mathcal{A}_h^k(s)$ s.t. $\underline{Q}_h^k(s, a) > \overline{Q}_h^k(s, a')$, then $\mathcal{A}_h^k(s) \leftarrow \mathcal{A}_h^k(s) \setminus \{a\}$.
 - 8: set $I_h^k(s, a, s') = 0, I_h^k(s, a) = 0, IC_h^k(s, a) = 0 \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 9: **for** $v = 1, \dots, m$ **do**
 - 10: sample $h_v \in [H]$ uniformly at random.
 - 11: set $\pi^{k, v} = \begin{cases} \underline{\pi}^k & \text{with probability } 1 - \epsilon \\ \pi^{k, h_v} & \text{with probability } \epsilon \end{cases}$, where: $\pi_h^{k, h'}(a | s) = \begin{cases} \underline{\pi}_h^k(a | s) & h \neq h' \\ \frac{1}{|\mathcal{A}_h^k(s)|} & h = h' \end{cases}$.
 - 12: observe initial state $s_1^{k, v}$.
 - 13: **for** $h = 1, \dots, H$ **do**
 - 14: pick action $a_h^{k, v} \sim \pi_h^{k, v}(\cdot | s_h^{k, v})$, suffer cost $C_h^{k, v}$ and observe next state $s_{h+1}^{k, v}$.
 - 15: update $I_h^k(s_h^{k, v}, a_h^{k, v}) \leftarrow 1, I_h^k(s_h^{k, v}, a_h^{k, v}, s_{h+1}^{k, v}) \leftarrow 1, IC_h^k(s_h^{k, v}, a_h^{k, v}) \leftarrow C_h^{k, v}$.
 - 16: **end for**
 - 17: **end for**
 - 18: set $n_h^{k+1}(s, a) \leftarrow n_h^k(s, a) + I_h^k(s, a), n_h^{k+1}(s, a, s') \leftarrow n_h^k(s, a, s') + I_h^k(s, a, s') \forall (s, a, s', h)$.
 - 19: set $C_h^{k+1}(s, a) \leftarrow C_h^k(s, a) + IC_h^k(s, a) \forall (s, a, h)$.
 - 20: **end for**
-

Appendix C. The coop-ULCAE algorithm for stochastic MDPs with non-fresh randomness

For the setting of stochastic MDPs with non-fresh randomness we propose the Cooperative Upper Lower Confidence Action Elimination algorithm (coop-ULCAE; see Algorithm 5). Recall that if all the agents play the optimistic policy (like coop-ULCVI), the regret will not improve since the randomness is non-fresh. Thus, we want each agent to diverge from the trajectory of the optimistic policy at some point. To that end, at some step each agent takes a random action. At the other steps it follows the optimistic policy to make sure that its regret does not increase. Finally, since all actions have probability to be explored, we eliminate sub-optimal actions to avoid unnecessary over exploration.

Theorem 20 *With probability $1 - \delta$, setting $\epsilon = \min\{\frac{HA}{m}, \frac{1}{\sqrt{m}}\}$, the individual regret of each agent of coop-ULCAE is*

$$R_K = O\left(\sqrt{H^5 SK} \log \frac{mHSAK}{\delta} + \sqrt{\frac{H^7 SAK}{\sqrt{m}}} \log \frac{mHSAK}{\delta} + \sqrt{\frac{H^8 SAK}{m}} \log \frac{mHSAK}{\delta} + H^5 S^2 A \log^2 \frac{mHSAK}{\delta} + \frac{H^6 S^2 A^2}{\sqrt{m}} \log^2 \frac{mHSAK}{\delta}\right).$$

C.1 The good event, optimism and pessimism

Define the following events (for $\tau = 3 \log \frac{6SAHKm}{\delta}$):

$$\begin{aligned}
E^c(k) &= \left\{ \forall (s, a, h) : |\hat{c}_h^k(s, a) - c_h(s, a)| \leq \sqrt{\frac{2\tau}{n_h^k(s, a) \vee 1}} \right\} \\
E^p(k) &= \left\{ \forall (s, a, s', h) : |p_h(s'|s, a) - \hat{p}_h^k(s'|s, a)| \leq \sqrt{\frac{2p_h(s'|s, a)\tau}{n_h^k(s, a) \vee 1} + \frac{2\tau}{n_h^k(s, a) \vee 1}} \right\} \\
E^{pv1}(k) &= \left\{ \forall (s, a, h) : |((\hat{p}_h^k(\cdot|s, a) - p_h(\cdot|s, a)) \cdot V_{h+1}^*)| \leq \sqrt{\frac{2\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^*)\tau}{n_h^k(s, a) \vee 1} + \frac{5H\tau}{n_h^k(s, a) \vee 1}} \right\} \\
E^{pv2}(k) &= \left\{ \forall (s, a, h) : \left| \sqrt{\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^*)} - \sqrt{\text{Var}_{\hat{p}_h^k(\cdot|s, a)}(V_{h+1}^*)} \right| \leq \sqrt{\frac{12H^2\tau}{n_h^k(s, a) \vee 1}} \right\}
\end{aligned}$$

The basic good event, which is the intersection of the above events, is the one used in Efroni et al. (2021). The following lemma establishes that the good event holds with high probability. The proof is supplied in Efroni et al. (2021, Lemma 13) by applying standard concentration results.

Lemma 21 (The First Good Event) *Let $\mathbb{G}_1 = \cap_{k=1}^K E^c(k) \cap_{k=1}^K E^p(k) \cap_{k=1}^K E^{pv1}(k) \cap_{k=1}^K E^{pv2}(k)$ be the basic good event. It holds that $\Pr(\mathbb{G}_1) \geq 1 - \delta/2$.*

Under the first good event, we can prove that the value is optimistic using standard techniques (similar to Efroni et al. (2021, Lemma 14)).

Lemma 22 (Upper Value Function is Pessimistic, Lower Value Function is Optimistic) *Conditioned on the first good event \mathbb{G}_1 , it holds that $\underline{V}_h^k(s) \leq V_h^*(s) \leq \overline{V}_h^k(s) \leq \bar{V}_h^k(s)$ and that $\underline{Q}_h^k(s, a) \leq Q_h^*(s, a) \leq \overline{Q}_h^k(s, a) \leq \bar{Q}_h^k(s, a)$ for every $k = 1, \dots, K$, $s \in \mathcal{S}$, $a \in \mathcal{A}$ and $h = 1, \dots, H$. Moreover, $\pi_h^k(s) \in \mathcal{A}_h^k(s)$ for every $k = 1, \dots, K$, $s \in \mathcal{S}$ and $h = 1, \dots, H$.*

Finally, we define the following events that are more specific to our algorithmic action elimination framework:

$$\begin{aligned}
E^{n1} &= \left\{ \forall (k, h, s) \in [K] \times [H] \times \mathcal{S} \forall a \in \mathcal{A}_h^k(s) : n_h^k(s, a) \geq \frac{m\epsilon}{4HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \log \frac{6HSA}{\delta} \right\} \\
E^{n2} &= \left\{ \forall (k, h, s, a, v) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A} \times [m] : n_h^k(s, a) \geq \frac{1}{2} \sum_{j=1}^{k-1} q_h^{\pi^j, v}(s, a) - \log \frac{6mHSA}{\delta} \right\} \\
E^\epsilon &= \left\{ \forall (h', v) \in [H] \times [m] : \sum_{k=1}^K \mathbb{I}\{\pi^{k, v} = \pi^{k, h'}\} \leq \frac{\epsilon}{H}K + \sqrt{K \log \frac{6mH}{\delta}} \right\}
\end{aligned}$$

Lemma 23 (The Second Good Event) *Let $\mathbb{G}_2 = E^{n1} \cap E^{n2} \cap E^\epsilon$ be the second good event. It holds that $\Pr(\mathbb{G}_2) \geq 1 - \delta/2$.*

As a direct consequence, we get that the good event \mathbb{G} which is the intersection of \mathbb{G}_1 and \mathbb{G}_2 holds with probability $1 - \delta$.

Lemma 24 (The Good Event) *Let \mathbb{G}_1 be the first good event defined in Lemma 21, and \mathbb{G}_2 be the second good event defined in Lemma 23. Then, the good event $\mathbb{G} = \mathbb{G}_1 \cap \mathbb{G}_2$ holds with probability $1 - \delta$.*

Proof of Lemma 23 We show that each of the events $\neg E^{n1}$, $\neg E^{n2}$, $\neg E^\epsilon$ occur with probability at most $\delta/6$. Then, by a union bound we obtain the statement.

$\Pr[\neg E^{n1}] \leq \delta/6$: Without loss of generality, assume that in each episode, each agent uniformly randomizes a permutation over all actions, $\sigma^{k,v}$, and in case of exploration takes the first active arm in the permutation $\sigma^{k,v}$: $\arg \min_{a \in \mathcal{A}_h^k(s)} \sigma^{k,v}(a)$. For any $a \in \mathcal{A}_h^k(s)$,

$$\begin{aligned}
n_h^k(s, a) &= \sum_{j=1}^{k-1} \mathbb{I}\{\exists v : s_h^{j,v} = s, a_h^{j,v} = a\} \\
&\geq \sum_{j=1}^{k-1} \mathbb{I}\{\exists v : s_h^{j,v} = s, a_h^{j,v} = a, h_v = h, \pi^{j,v} = \pi^{j,h_v}, \sigma_s^{j,v}(a) = 1\} \\
&\stackrel{(*)}{\geq} \sum_{j=1}^{k-1} \mathbb{I}\{\exists v : s_h^{j,v} = s, h_v = h, \pi^{j,v} = \pi^{j,h_v}, \sigma_s^{j,v}(a) = 1\} \\
&\stackrel{(**)}{=} \sum_{j=1}^{k-1} \mathbb{I}\{s_h^{j,\pi^j} = s\} \mathbb{I}\{\exists v : h_v = h, \pi^{j,v} = \pi^{j,h_v}, \sigma_s^{j,v}(a) = 1\}, \tag{15}
\end{aligned}$$

For $(*)$, recall that $a \in \mathcal{A}_h^k(s)$, which implies that $a \in \mathcal{A}_h^j(s)$. Therefore, if $h_v = h, \pi^{j,v} = \pi^{j,h_v}, \sigma_s^{j,v}(a) = 1$ (that is, the agent explores h , and a is the first action in the permutation), then $a_h^k = a$. $(**)$ is because each agent that randomize $h_v = h$ follows the (deterministic) optimistic until horizon h . Since $h_v, \sigma_s^{j,v}$ and the event $\{\pi^{j,v} = \pi^{j,h_v}\}$ are randomized independently,

$$\begin{aligned}
\mathbb{E} \left[\mathbb{I}\{s_h^{j,\pi^j} = s\} \mathbb{I}\{\exists v : h_v = h, \pi^{j,v} = \pi^{j,h_v}, \sigma_s^{j,v}(a) = 1\} \mid \mathcal{F}^j \right] &= \\
&= q_h^{\pi^j}(s) \Pr[\exists v : h_v = h, \pi^{j,v} = \pi^{j,h_v}, \sigma_s^{j,v}(a) = 1] \\
&= q_h^{\pi^j}(s) [1 - \Pr[\forall v : h_v \neq h \vee \pi^{j,v} \neq \pi^{j,h_v} \vee \sigma_s^{j,v}(a) \neq 1]] \\
&= q_h^{\pi^j}(s) [1 - (\Pr[h_1 \neq h \vee \pi^{j,1} \neq \pi^{j,h_1} \vee \sigma_s^{j,1}(a) \neq 1])^m] \\
&= q_h^{\pi^j}(s) [1 - (1 - \Pr[h_1 = h, \pi^{j,1} = \pi^{j,h_1}, \sigma_s^{j,1}(a) = 1])^m] \\
&= q_h^{\pi^j}(s) \left[1 - \left(1 - \frac{\epsilon}{HA} \right)^m \right] \\
&= q_h^{\pi^j}(s) \left[1 - \left(\left(1 - \frac{\epsilon}{HA} \right)^{\frac{HA}{\epsilon}} \right)^{m \frac{\epsilon}{HA}} \right] \\
&\geq q_h^{\pi^j}(s) [1 - e^{-\frac{m\epsilon}{HA}}] \tag{((1-x^{-1})^x \leq e)} \\
&\geq q_h^{\pi^j}(s) \left[\frac{m\epsilon}{HA} - \frac{1}{2} \left(\frac{m\epsilon}{HA} \right)^2 \right] \tag{(e^{-x} \leq 1 - x + \frac{x^2}{2})} \\
&\geq q_h^{\pi^j}(s) \frac{m\epsilon}{2HA}. \tag{(\epsilon \leq \frac{HA}{m})}
\end{aligned}$$

By Dann et al. (2017)[Lemma F.4] and Eq. (15) we have, $\Pr \left[\exists k : n_h^k(s, a) < \frac{m\epsilon}{4HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \log \frac{6HSA}{\delta} \right] \leq \frac{\delta}{6HSA}$. By taking the union bound over all $h \in [H], s \in \mathcal{S}$ and $a \in \mathcal{A}$, we get $\Pr(\neg E^{n1}) \leq \delta/6$.

$\Pr[\neg E^{n2}] \leq \delta/6$: For any $v \in [m]$,

$$n_h^k(s, a) = \sum_{j=1}^{k-1} \mathbb{I}\{\exists v' \in [m] : s_h^{j,v'} = s, a_h^{j,v'} = a\} \geq \sum_{j=1}^{k-1} \mathbb{I}\{s_h^{j,v} = s, a_h^{j,v} = a\}.$$

Again, by [Dann et al. \(2017, Lemma F.4\)](#), we get $\Pr \left[\exists k \in [K] : n_h^k(s, a) < \frac{1}{2} \sum_{j=1}^{k-1} q_h^{\pi^{j,v}}(s, a) - \log \frac{6mHSA}{\delta} \right] \leq \frac{\delta}{6mHSA}$. Taking the union bound we get $\Pr[\neg E^{n^2}] \leq \delta/6$.

$\Pr[\neg E^\epsilon] \leq \delta/6$: Directly from Hoeffding's inequality and a union bound. ■

C.2 Proof of Theorem 20

Proof of Theorem 20 By [Lemma 24](#), the good event holds with probability $1 - \delta$. We now analyze the regret under the assumption that the good event holds. We start by decomposing the regret according to the policy played by agent v :

$$\begin{aligned}
R_K &= \sum_{k=1}^K V_1^{\pi^{k,v}}(s_1^{k,v}) - V_1^*(s_1^{k,v}) \\
&= \sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \underline{\pi}^k\} \left(V_1^{\underline{\pi}^k}(s_1^{k,v}) - V_1^*(s_1^{k,v}) \right) + \sum_{k=1}^K \sum_{h'=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \left(V_1^{\pi^{k,h'}}(s_1^{k,v}) - V_1^*(s_1^{k,v}) \right) \\
&\leq \sum_{k=1}^K V_1^{\underline{\pi}^k}(s_1^{k,v}) - V_1^*(s_1^{k,v}) + \sum_{k=1}^K \sum_{h'=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \left(V_1^{\pi^{k,h'}}(s_1^{k,v}) - V_1^*(s_1^{k,v}) \right).
\end{aligned}$$

For the first term we use [Lemma 31](#), then [Lemma 26](#) and then [Lemma 25](#):

$$\begin{aligned}
\sum_{k=1}^K V_1^{\underline{\pi}^k}(s_1^{k,v}) - V_1^*(s_1^{k,v}) &\lesssim H \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}^k \left[\frac{\sqrt{\tau \text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\underline{\pi}^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} + \frac{H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \underline{\pi}^k \right] \\
&\lesssim H \sqrt{\tau} \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}^k \left[\frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\underline{\pi}^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} \mid \underline{\pi}^k \right] + \frac{H^5 S^2 A^2 \tau^2}{m\epsilon} \\
&\lesssim \sqrt{\frac{H^6 S A K \tau^2}{m\epsilon}} + \frac{H^5 S^2 A^2 \tau^2}{m\epsilon}.
\end{aligned}$$

For the second term we use Lemma 30, then Lemmas 26 and 27 and then Lemmas 28 and 29:

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h'=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \left(V_1^{\pi^{k,h'}}(s_1^{k,v}) - V_1^*(s_1^{k,v}) \right) \lesssim \\
& \lesssim H \sum_{k=1}^K \sum_{h'=1}^H \sum_{h=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \mathbb{E}^k \left[\frac{\sqrt{\tau \text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} + \frac{H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^{k,h'} \right] \\
& \quad + H^2 \sum_{k=1}^K \sum_{h'=1}^H \sum_{h=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \mathbb{E}^k \left[\frac{\sqrt{\tau \text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} + \frac{H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \underline{\pi}^k \right] \\
& \lesssim H \sqrt{\tau} \sum_{k=1}^K \sum_{h'=1}^H \sum_{h=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \mathbb{E}^k \left[\frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} \mid \pi^{k,h'} \right] + H^5 S^2 A \tau^2 \\
& \quad + H^2 \sqrt{\tau} \sum_{k=1}^K \sum_{h'=1}^H \sum_{h=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \mathbb{E}^k \left[\frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} \mid \underline{\pi}^k \right] + \frac{H^6 S^2 A^2 \tau^2}{m \epsilon} \\
& \lesssim \sqrt{H^7 S A K \epsilon \tau^2} + \tau \sqrt{H^7 S A \tau^3} K^{1/4} + H^5 S^2 A \tau^2 + \sqrt{\frac{H^8 S A K \tau}{m}} + \sqrt{\frac{H^9 S A \tau^2}{m \epsilon}} K^{1/4} + \frac{H^6 S^2 A^2 \tau^2}{m \epsilon}.
\end{aligned}$$

Setting $\epsilon = \min \left\{ \frac{H A}{m}, \frac{1}{\sqrt{m}} \right\}$, we get:

$$\begin{aligned}
R_K & \lesssim \sqrt{\frac{H^6 S A K \tau^2}{m \epsilon}} + \frac{H^6 S^2 A^2 \tau^2}{m \epsilon} + \sqrt{H^7 S A K \epsilon \tau^2} + \tau \sqrt{H^7 S A \tau^3} K^{1/4} \\
& \quad + H^5 S^2 A \tau^2 + \sqrt{\frac{H^8 S A K \tau}{m}} + \sqrt{\frac{H^9 S A \tau^2}{m \epsilon}} K^{1/4} \\
& \lesssim \sqrt{H^5 S K \tau^2} + \sqrt{\frac{H^7 S A K \tau^2}{\sqrt{m}}} + \sqrt{\frac{H^8 S A K \tau}{m}} + H^5 S^2 A \tau^2 + \frac{H^6 S^2 A^2 \tau^2}{\sqrt{m}} \\
& \quad + \sqrt{H^7 S A \tau^3} K^{1/4} + \sqrt{H^8 S \tau^2} K^{1/4} + \sqrt{\frac{H^9 S A \tau^2}{\sqrt{m}}} K^{1/4} \\
& \lesssim \sqrt{H^5 S K \tau^2} + \sqrt{\frac{H^7 S A K \tau^2}{\sqrt{m}}} + \sqrt{\frac{H^8 S A K \tau}{m}} + H^5 S^2 A \tau^2 + \frac{H^6 S^2 A^2 \tau^2}{\sqrt{m}},
\end{aligned}$$

where the last inequality follows because the $K^{1/4}$ terms are dominant only when K is small, and in these cases the constant terms are larger. \blacksquare

C.3 Bounds on the cumulative bonuses

Lemma 25 *Under the good event, if $\frac{m \epsilon}{H A} \leq 1$,*

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{E}^k \left[\frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} \mid \underline{\pi}^k \right] \lesssim \sqrt{\frac{H^4 S A K \tau}{m \epsilon}} + \frac{H^3 S A^2 \tau}{m \epsilon}.$$

Proof By the event E^{n_1} , we have:

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}^k \left[\frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k)} \vee 1} \mid \pi^k \right] = \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \frac{q_h^{\pi^k}(s, a) \sqrt{\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s, a)} \vee 1} \\
& \leq \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{h=1}^H \sum_{k=1}^K \frac{q_h^{\pi^k}(s, a) \sqrt{\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})}}{\sqrt{(\frac{m\epsilon}{4HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \tau) \vee 1}} \\
& \leq \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{h=1}^H \sum_{k: \frac{m\epsilon}{8HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) \leq \tau} \underbrace{q_h^{\pi^k}(s, a) \sqrt{\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})}}_{\leq H} \\
& \quad + \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{h=1}^H \sum_{k: \frac{m\epsilon}{8HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) > \tau} \frac{q_h^{\pi^k}(s, a) \sqrt{\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})}}{\sqrt{(\frac{m\epsilon}{4HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \tau) \vee 1}} \\
& \lesssim \frac{H^3 A^2 S \tau}{m\epsilon} + \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{h=1}^H \sum_{k=1}^K \frac{q_h^{\pi^k}(s, a) \sqrt{\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})}}{\sqrt{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1}} \\
& \leq \frac{H^3 A^2 S \tau}{m\epsilon} + \sqrt{\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{h=1}^H \sum_{k=1}^K q_h^{\pi^k}(s, a) \text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})} \sqrt{\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{h=1}^H \sum_{k=1}^K \frac{q_h^{\pi^k}(s, a)}{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1}} \\
& = \frac{H^3 A^2 S \tau}{m\epsilon} + \sqrt{\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{h=1}^H \sum_{k=1}^K q_h^{\pi^k}(s, a) \text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})} \sqrt{\sum_{s \in \mathcal{S}} \sum_{h=1}^H \sum_{k=1}^K \frac{q_h^{\pi^k}(s)}{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1}},
\end{aligned}$$

where the last inequality is Cauchy–Schwarz inequality. Using the law of total variance (Cohen et al., 2021, Lemma B.14),

$$\begin{aligned}
& \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^{\pi^k}(s, a) \text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k}) = \mathbb{E} \left[\sum_{h=1}^H \text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k}) \mid \pi^k \right] \\
& = \mathbb{E} \left[\left(V_1^{\pi^k}(s_1^k) - \sum_h c_h(s_h^k, a_h^k) \right)^2 \mid \pi^k \right] \leq H^2.
\end{aligned} \tag{16}$$

Using Rosenberg et al. (2020, Lemma B.18), since $\frac{m\epsilon}{HA} \leq 1$,

$$\sum_{s \in \mathcal{S}} \sum_{h=1}^H \sum_{k=1}^K \frac{q_h^{\pi^k}(s)}{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1} = \frac{HA}{m\epsilon} \sum_{s \in \mathcal{S}} \sum_{h=1}^H \sum_{k=1}^K \frac{\frac{m\epsilon}{HA} q_h^{\pi^k}(s)}{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1} \lesssim \frac{H^2 S A \tau}{m\epsilon}. \tag{17}$$

Combining the last three inequalities completes the proof. ■

Lemma 26 Under the good event, $\sum_{k=1}^K \sum_{h=1}^H \mathbb{E}^k \left[\frac{1}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^k \right] \lesssim \frac{H^2 S A^2 \tau}{m\epsilon}$.

Proof By the event E^{n1} , we have:

$$\begin{aligned}
\sum_{k=1}^K \sum_{h=1}^H \mathbb{E}^k \left[\frac{1}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^k \right] &= \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{q_h^{\pi^k}(s, a)}{n_h^k(s, a) \vee 1} \leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{q_h^{\pi^k}(s, a)}{\left(\frac{m\epsilon}{4HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \tau \right) \vee 1} \\
&\leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k: \frac{m\epsilon}{8HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) \leq \tau} q_h^{\pi^k}(s, a) + \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k: \frac{m\epsilon}{8HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) > \tau} \frac{q_h^{\pi^k}(s, a)}{\left(\frac{m\epsilon}{4HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \tau \right) \vee 1} \\
&\lesssim \frac{H^2 S A^2 \tau}{m\epsilon} + \frac{H A}{m\epsilon} \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{q_h^{\pi^k}(s, a)}{\left(\sum_{j=1}^{k-1} q_h^{\pi^j}(s) \right) \vee 1} \lesssim \frac{H^2 S A^2 \tau}{m\epsilon},
\end{aligned}$$

where the last inequality follows from [Rosenberg et al. \(2020, Lemma B.18\)](#). ■

Lemma 27 Let $h' \in [H]$. Under the good event, $\sum_{k=1}^K \sum_{h=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \mathbb{E}^k \left[\frac{1}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^{k,h'} \right] \lesssim H S A \tau$.

Proof By the event E^{n2} , we have:

$$\begin{aligned}
\sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \sum_{h=1}^H \mathbb{E}^k \left[\frac{1}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^{k,h'} \right] &= \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \frac{q_h^{\pi^{k,h'}}(s, a)}{n_h^k(s, a) \vee 1} \\
&\leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \frac{q_h^{\pi^{k,h'}}(s, a)}{\left(\frac{1}{2} \sum_{j=1}^{k-1} q_h^{\pi^j, v}(s, a) - \tau \right) \vee 1} \\
&\leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a)}{\left(\frac{1}{2} \sum_{j=1}^{k-1} \mathbb{I}\{\pi^{j,v} = \pi^{j,h'}\} q_h^{\pi^{j,h'}}(s, a) - \tau \right) \vee 1} \\
&\leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k: \frac{1}{4} \sum_{j=1}^{k-1} \mathbb{I}\{\pi^{j,v} = \pi^{j,h'}\} q_h^{\pi^{j,h'}}(s, a) \leq \tau} \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a) \\
&\quad + \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k: \frac{1}{4} \sum_{j=1}^{k-1} \mathbb{I}\{\pi^{j,v} = \pi^{j,h'}\} q_h^{\pi^{j,h'}}(s, a) > \tau} \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a)}{\left(\frac{1}{2} \sum_{j=1}^{k-1} \mathbb{I}\{\pi^{j,v} = \pi^{j,h'}\} q_h^{\pi^{j,h'}}(s, a) - \tau \right) \vee 1} \\
&\lesssim H S A \tau + \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a)}{\left(\sum_{j=1}^{k-1} \mathbb{I}\{\pi^{j,v} = \pi^{j,h'}\} q_h^{\pi^{j,h'}}(s, a) \right) \vee 1} \lesssim H S A \tau,
\end{aligned}$$

where the last inequality follows [Rosenberg et al. \(2020, Lemma B.18\)](#). ■

Lemma 28 Let $h' \in [H]$. Under the good event,

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \mathbb{E}^k \left[\frac{\sqrt{\text{Var}_{p_h(\cdot | s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} \mid \pi^{k,h'} \right] \lesssim \sqrt{H^3 S A K \epsilon \tau} + \tau \sqrt{H^3 S A K}^{1/4} + H S A \tau.$$

Proof First we bound $\sqrt{\text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}$ by H . Now, under the good event,

$$\begin{aligned}
& \sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \sum_{h=1}^H \mathbb{E}^k \left[\frac{1}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} \mid \pi^{k,h'} \right] = \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \frac{q_h^{\pi^{k,h'}}(s, a)}{\sqrt{n_h^k(s, a) \vee 1}} \\
& \leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \frac{q_h^{\pi^{k,h'}}(s, a)}{\sqrt{\left(\frac{1}{2} \sum_{j=1}^{k-1} q_h^{\pi^{j,v}}(s, a) - \tau\right) \vee 1}} \\
& \leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a)}{\sqrt{\left(\frac{1}{2} \sum_{j=1}^{k-1} \mathbb{I}\{\pi^{j,v} = \pi^{j,h'}\} q_h^{\pi^{j,h'}}(s, a) - \tau\right) \vee 1}} \\
& \lesssim HSA\tau + \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a)}{\sqrt{\left(\sum_{j=1}^{k-1} \mathbb{I}\{\pi^{j,v} = \pi^{j,h'}\} q_h^{\pi^{j,h'}}(s, a)\right) \vee 1}} \\
& \leq \sqrt{\sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a)}{\left(\sum_{j=1}^{k-1} \mathbb{I}\{\pi^{j,v} = \pi^{j,h'}\} q_h^{\pi^{j,h'}}(s, a)\right) \vee 1}} \sqrt{\sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a)} \\
& \quad + HSA\tau,
\end{aligned}$$

where the first inequality is by E^{n^2} , the second inequality is done by breaking the sum to ks such that $\frac{1}{4} \sum_{j=1}^{k-1} q_h^{\pi^{j,v}}(s, a) \leq \tau$ and the rest of the ks , as done in the proof of Lemma 27 for example, and the last is Cauchy–Schwarz inequality. By Rosenberg et al. (2020, Lemma B.18),

$$\sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a)}{\left(\sum_{j=1}^{k-1} \mathbb{I}\{\pi^{j,v} = \pi^{j,h'}\} q_h^{\pi^{j,h'}}(s, a)\right) \vee 1} \lesssim HSA\tau.$$

By the good event E^ϵ ,

$$\sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^{k,h'}}(s, a) = H \sum_{k=1}^K \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \leq K\epsilon + \sqrt{K\tau}.$$

■

Lemma 29 Under the good event,

$$\sum_{k=1}^K \sum_{h'=1}^H \sum_{h=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \mathbb{E}^k \left[\frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} \mid \pi^k \right] \lesssim \sqrt{\frac{H^4 SAK\tau}{m}} + \sqrt{\frac{H^5 SA\tau^2}{m\epsilon}} K^{1/4} + \frac{H^3 A^2 S\tau}{m\epsilon}.$$

Proof By the event E^{n1} , we have:

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h'=1}^H \sum_{h=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \mathbb{E}^k \left[\frac{\sqrt{\text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} \mid \underline{\pi}^k \right] = \\
& = \sum_{k=1}^K \sum_{h'=1}^H \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \frac{q_h^{\pi^k}(s, a) \sqrt{\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} \\
& \leq \sum_{k=1}^K \sum_{h'=1}^H \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \frac{q_h^{\pi^k}(s, a) \sqrt{\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})}}{\sqrt{(\frac{m\epsilon}{4HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \tau) \vee 1}} \\
& \lesssim \frac{H^3 A^2 S \tau}{m\epsilon} + \sum_{k=1}^K \sum_{h'=1}^H \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \frac{q_h^{\pi^k}(s, a) \sqrt{\text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})}}{\sqrt{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1}} \\
& \leq \sqrt{\sum_{k, h', h, s, a} \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^k}(s, a) \text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})} \sqrt{\sum_{k, h', h, s, a} \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^k}(s, a)}{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1}} \\
& \quad + \frac{H^3 A^2 S \tau}{m\epsilon} \\
& = \sqrt{\sum_{k, h'} \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \sum_{h, s, a} q_h^{\pi^k}(s, a) \text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k})} \sqrt{\sum_{k, h', h, s} \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^k}(s)}{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1}} \\
& \quad + \frac{H^3 A^2 S \tau}{m\epsilon},
\end{aligned}$$

where the last inequality is by Cauchy-Schwarz. By Eq. (16) and the good event E^ϵ ,

$$\begin{aligned}
\sum_{k=1}^K \sum_{h'=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^{\pi^k}(s, a) \text{Var}_{p_h(\cdot|s, a)}(V_{h+1}^{\pi^k}) & \leq H^2 \sum_{k=1}^K \sum_{h'=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \\
& \leq H^2 K \epsilon + H^3 \sqrt{K \log \frac{mH}{\delta'}}.
\end{aligned}$$

For last,

$$\sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{h'=1}^H \frac{\mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} q_h^{\pi^k}(s)}{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1} \leq \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \frac{q_h^{\pi^k}(s)}{(\frac{m\epsilon}{HA} \sum_{j=1}^{k-1} q_h^{\pi^j}(s)) \vee 1} \lesssim \frac{H^2 S A \tau}{m\epsilon},$$

where the first inequality is since $\sum_{h'=1}^H \mathbb{I}\{\pi^{k,v} = \pi^{k,h'}\} \leq 1$ and the last is as in Eq. (17). Combining the last three inequalities completes the proof. \blacksquare

C.4 Bounding the regret of $\pi^{k,h'}$ and the optimistic policy

Lemma 30 *Under the good event, for every $k \in [K]$ it holds that:*

$$\begin{aligned} V_1^{\pi^{k,h'}}(s_1^k) - V_1^*(s_1^k) &\lesssim H \sum_{h=1}^H \mathbb{E}^k \left[\frac{\sqrt{\tau \text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} + \frac{H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^{k,h'} \right] \\ &\quad + H^2 \sum_{h=1}^H \mathbb{E}^k \left[\frac{\sqrt{\tau \text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} + \frac{H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^k \right]. \end{aligned}$$

Proof Apply Lemma 32 for every $k \in [K]$ and then apply Lemmas 35 to 38 iteratively. ■

Lemma 31 *Under the good event, for every $k \in [K]$ it holds that:*

$$V_1^{\pi^k}(s_1^k) - V_1^*(s_1^k) \lesssim H \sum_{h=1}^H \mathbb{E}^k \left[\frac{\sqrt{\tau \text{Var}_{p_h(\cdot|s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} + \frac{H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^k \right].$$

Proof Similar to the proof of Lemma 30. ■

Lemma 32 *Let $h' \in [H]$. Under the good event, for every $k \in [K]$ it holds that:*

$$\begin{aligned} V_1^{\pi^{k,h'}}(s_1^k) - V_1^*(s_1^k) &\leq \sum_{h=1}^H \mathbb{E}^k \left[\overline{Q}_h^k(s_h^k, a_h^k) - \underline{Q}_h^k(s_h^k, a_h^k) \mid \pi^{k,h'} \right] \\ &\quad + \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) - \underline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'} \right]. \end{aligned}$$

Proof It holds that:

$$\begin{aligned} V_1^{\pi^{k,h'}}(s_1^k) - V_1^*(s_1^k) &= V_1^{\pi^{k,h'}}(s_1^k) - V_1^*(s_1^k) = \sum_h \mathbb{E}^k \left[\langle Q_h^*(s_h^k, \cdot), \pi^{k,h'}(\cdot \mid s_h^k) - \pi_h^*(\cdot \mid s_h^k) \rangle \mid \pi^{k,h'} \right] \\ &= \sum_h \mathbb{E}^k \left[Q_h^*(s_h^k, a_h^k) - V_h^*(s_h^k) \mid \pi^{k,h'} \right] \leq \sum_h \mathbb{E}^k \left[\overline{Q}_h^k(s_h^k, a_h^k) - \underline{V}_h^k(s_h^k) \mid \pi^{k,h'} \right] \\ &= \sum_h \mathbb{E}^k \left[\overline{Q}_h^k(s_h^k, a_h^k) - \underline{Q}_h^k(s_h^k, \underline{\pi}_h^k(s_h^k)) \mid \pi^{k,h'} \right] \\ &= \sum_{h \neq h'} \mathbb{E}^k \left[\overline{Q}_h^k(s_h^k, a_h^k) - \underline{Q}_h^k(s_h^k, \underline{\pi}_h^k(s_h^k)) \mid \pi^{k,h'} \right] \\ &\quad + \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, a_{h'}^k) - \underline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'} \right] \\ &= \sum_{h \neq h'} \mathbb{E}^k \left[\overline{Q}_h^k(s_h^k, a_h^k) - \underline{Q}_h^k(s_h^k, a_h^k) \mid \pi^{k,h'} \right] \\ &\quad + \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, a_{h'}^k) - \underline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'} \right] \\ &\leq \sum_h \mathbb{E}^k \left[\overline{Q}_h^k(s_h^k, a_h^k) - \underline{Q}_h^k(s_h^k, a_h^k) \mid \pi^{k,h'} \right] \\ &\quad + \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) - \underline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'} \right]. \end{aligned}$$

The first inequality is by Lemma 22, the last equality is since $\pi^{k,h'} = \underline{\pi}^k$ for any $h \neq h'$ and the last inequality is by Lemma 34. ■

Lemma 33 *Under the good event, for every $k \in [K]$ it holds that:*

$$V_1^{\underline{\pi}^k}(s_1^k) - V_1^*(s_1^k) \leq \sum_{h=1}^H \mathbb{E}^k \left[\overline{Q}_h^k(s_h^k, a_h^k) - \underline{Q}_h^k(s_h^k, a_h^k) \mid \underline{\pi}^k \right].$$

Proof Similar to the proof of Lemma 32. ■

C.5 Auxiliary lemmas

Lemma 34 *If $a, a' \in \mathcal{A}_h^k(s)$ then,*

$$\overline{Q}_h^k(s, a) - \underline{Q}_h^k(s_h^k, a') \leq \overline{Q}_h^k(s, a) - \underline{Q}_h^k(s, a) + \overline{Q}_h^k(s, a') - \underline{Q}_h^k(s, a').$$

Proof Since $a, a' \in \mathcal{A}_h^k(s)$, we have that $\underline{Q}_h^k(s, a) \leq \overline{Q}_h^k(s, a')$. Thus:

$$\overline{Q}_h^k(s, a) - \underline{Q}_h^k(s, a') = \overline{Q}_h^k(s, a) - \underline{Q}_h^k(s, a) + \overline{Q}_h^k(s, a') - \underline{Q}_h^k(s, a') + \underbrace{\underline{Q}_h^k(s, a) - \overline{Q}_h^k(s, a')}_{\leq 0}.$$

Lemma 35 (Recursion with Optimistic Next-Action) *Let $h \neq h' - 1$. Under the good event, for every $k \in [K]$ it holds that:*

$$\begin{aligned} \mathbb{E}^k \left[\overline{Q}_h^k(s_h^k, a_h^k) - \underline{Q}_h^k(s_h^k, a_h^k) \mid \pi^{k,h'} \right] &\leq \mathbb{E}^k \left[\frac{8\sqrt{\tau \text{Var}_{p_h(\cdot | s_h^k, a_h^k)}(V_{h+1}^{\underline{\pi}^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} + \frac{118H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^{k,h'} \right] \\ &+ \left(1 + \frac{1}{4H} \right) \mathbb{E}^k \left[\overline{Q}_{h+1}^k(s_{h+1}^k, a_{h+1}^k) - \underline{Q}_{h+1}^k(s_{h+1}^k, a_{h+1}^k) \mid \pi^{k,h'} \right]. \end{aligned}$$

Proof By definition of the optimistic and pessimistic Q -functions, we have:

$$\begin{aligned}
& \mathbb{E}^k [\overline{Q}_h^k(s_h^k, a_h^k) - \underline{Q}_h^k(s_h^k, a_h^k) \mid \pi^{k, h'}] = \\
& = \mathbb{E}^k [2b_h^k(s_h^k, a_h^k; c) + 2b_h^k(s_h^k, a_h^k; p) \mid \pi^{k, h'}] \\
& \quad + \mathbb{E}^k \left[\mathbb{E}_{\tilde{p}_h^k(\cdot \mid s_h^k, a_h^k)} [\overline{V}_{h+1}^k - \underline{V}_{h+1}^k] \mid \pi^{k, h'} \right] \\
& \leq \mathbb{E}^k [2b_h^k(s_h^k, a_h^k; c) + 2b_h^k(s_h^k, a_h^k; p) \mid \pi^{k, h'}] \\
& \quad + \mathbb{E}^k \left[\frac{18H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} + \left(1 + \frac{1}{16H}\right) \mathbb{E}_{p_h(\cdot \mid s_h^k, a_h^k)} [\overline{V}_{h+1}^k - \underline{V}_{h+1}^k] \mid \pi^{k, h'} \right] \\
& = \mathbb{E}^k \left[2b_h^k(s_h^k, a_h^k; c) + 2b_h^k(s_h^k, a_h^k; p) + \frac{18H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^{k, h'} \right] \\
& \quad + \left(1 + \frac{1}{16H}\right) \mathbb{E}^k \left[\mathbb{E}_{p_h(\cdot \mid s_h^k, a_h^k)} [\overline{V}_{h+1}^k(s_{h+1}^k) - \underline{V}_{h+1}^k(s_{h+1}^k)] \mid \pi^{k, h'} \right] \\
& \leq \mathbb{E}^k \left[\frac{8\sqrt{\tau \text{Var}_{p_h(\cdot \mid s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} + \frac{118H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^{k, h'} \right] \\
& \quad + \left(1 + \frac{1}{4H}\right) \mathbb{E}^k \left[\mathbb{E}_{p_h(\cdot \mid s_h^k, a_h^k)} [\overline{V}_{h+1}^k - \underline{V}_{h+1}^k] \mid \pi^{k, h'} \right]
\end{aligned}$$

where the first inequality is by [Cohen et al. \(2021, Lemma B.13\)](#), and the second one is by [Cohen et al. \(2021, Lemma B.6\)](#). Let $\bar{\pi}_h^k(s) = \arg \min_a \overline{Q}_h^k(s, a)$. For the last term we have:

$$\begin{aligned}
& \mathbb{E}^k \left[\mathbb{E}_{p_h(\cdot \mid s_h^k, a_h^k)} [\overline{V}_{h+1}^k - \underline{V}_{h+1}^k] \mid \pi^{k, h'} \right] = \\
& = \sum_{h, s, a, s'} q_h^{\pi^{k, h'}}(s, a) p_h(s' \mid s, a) \left(\overline{V}_{h+1}^k(s') - \underline{V}_{h+1}^k(s') \right) \\
& = \sum_{h, s} q_{h+1}^{\pi^{k, h'}}(s) \left(\overline{V}_{h+1}^k(s) - \underline{V}_{h+1}^k(s) \right) \\
& = \mathbb{E}^k \left[\overline{V}_{h+1}^k(s_{h+1}^k) - \underline{V}_{h+1}^k(s_{h+1}^k) \mid \pi^{k, h'} \right] \\
& = \mathbb{E}^k \left[\overline{Q}_{h+1}^k(s_{h+1}^k, \bar{\pi}_{h+1}^k(s_{h+1}^k)) - \underline{Q}_{h+1}^k(s_{h+1}^k, \underline{\pi}_{h+1}^k(s_{h+1}^k)) \mid \pi^{k, h'} \right] \\
& \leq \mathbb{E}^k \left[\overline{Q}_{h+1}^k(s_{h+1}^k, a_{h+1}^k) - \underline{Q}_{h+1}^k(s_{h+1}^k, \underline{\pi}_{h+1}^k(s_{h+1}^k)) \mid \pi^{k, h'} \right] \\
& = \mathbb{E}^k \left[\overline{Q}_{h+1}^k(s_{h+1}^k, a_{h+1}^k) - \underline{Q}_{h+1}^k(s_{h+1}^k, a_{h+1}^k) \mid \pi^{k, h'} \right],
\end{aligned}$$

where the last equality follows because $h \neq h' - 1$. ■

Lemma 36 (Recursion with Non-Optimistic Next-Action) Let $h' \in [H]$. Under the good event, for every $k \in [K]$ it holds that:

$$\begin{aligned} \mathbb{E}^k \left[\overline{Q}_{h'-1}^k(s_{h'-1}^k, a_{h'-1}^k) - \underline{Q}_{h'-1}(s_{h'-1}^k, a_{h'-1}^k) \mid \pi^{k,h'} \right] &\leq \\ &\leq \mathbb{E}^k \left[\frac{8\sqrt{\tau \text{Var}_{p_{h'-1}(\cdot | s_{h'-1}^k, a_{h'-1}^k)}(V_{h'}^{\pi^k})}}{\sqrt{n_{h'-1}^k(s_{h'-1}^k, a_{h'-1}^k)} \vee 1} + \frac{118H^2 S\tau}{n_{h'-1}^k(s_{h'-1}^k, a_{h'-1}^k) \vee 1} \mid \pi^{k,h'} \right] \\ &+ \left(1 + \frac{1}{4H}\right) \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, a_{h'}^k) - \underline{Q}_{h'}^k(s_{h'}^k, a_{h'}^k) \mid \pi^{k,h'} \right] \\ &+ \left(1 + \frac{1}{4H}\right) \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) - \underline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'} \right]. \end{aligned}$$

Proof Similarly to Lemma 35, we have,

$$\begin{aligned} \mathbb{E}^k \left[\overline{Q}_{h'-1}^k(s_{h'-1}^k, a_{h'-1}^k) - \underline{Q}_{h'-1}(s_{h'-1}^k, a_{h'-1}^k) \mid \pi^{k,h'} \right] &\leq \\ &\leq \mathbb{E}^k \left[\frac{8\sqrt{\tau \text{Var}_{p_{h'-1}(\cdot | s_{h'-1}^k, a_{h'-1}^k)}(V_{h'}^{\pi^k})}}{\sqrt{n_{h'-1}^k(s_{h'-1}^k, a_{h'-1}^k)} \vee 1} + \frac{118H^2 S\tau}{n_{h'-1}^k(s_{h'-1}^k, a_{h'-1}^k) \vee 1} \mid \pi^{k,h'} \right] \\ &+ \left(1 + \frac{1}{4H}\right) \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, a_{h'}^k) - \underline{V}_{h'}^k(s_{h'}^k) \mid \pi^{k,h'} \right]. \end{aligned}$$

Now, by Lemma 34,

$$\begin{aligned} \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, a_{h'}^k) - \underline{V}_{h'}^k(s_{h'}^k) \mid \pi^{k,h'} \right] &= \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, a_{h'}^k) - \underline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'} \right] \\ &\leq \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, a_{h'}^k) - \underline{Q}_{h'}^k(s_{h'}^k, a_{h'}^k) \mid \pi^{k,h'} \right] \\ &\quad + \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) - \underline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'} \right]. \end{aligned}$$

■

Lemma 37 (Exploration Penalty Term Recursion) Let $h' \in [H]$. Under the good event, for every $k \in [K]$ it holds that:

$$\begin{aligned} \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) - \underline{Q}_{h'}(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'} \right] &\leq \\ &\leq \mathbb{E}^k \left[\frac{8\sqrt{\tau \text{Var}_{p_{h'}(\cdot | s_{h'}^k, a_{h'}^k)}(V_{h'+1}^{\pi^k})}}{\sqrt{n_{h'}^k(s_{h'}^k, a_{h'}^k)} \vee 1} + \frac{118H^2 S\tau}{n_{h'}^k(s_{h'}^k, a_{h'}^k) \vee 1} \mid \underline{\pi}^k \right] \\ &+ \left(1 + \frac{1}{4H}\right) \mathbb{E}^k \left[\overline{Q}_{h'+1}^k(s_{h'+1}^k, a_{h'+1}^k) - \underline{Q}_{h'+1}^k(s_{h'+1}^k, a_{h'+1}^k) \mid \underline{\pi}^k \right]. \end{aligned}$$

Proof Again, similar to Lemma 35,

$$\begin{aligned} \mathbb{E}^k \left[\overline{Q}_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) - \underline{Q}_{h'}(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'} \right] &\leq \\ &\leq \mathbb{E}^k \left[\frac{8\sqrt{\tau \text{Var}_{p_{h'}(\cdot | s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k))}(V_{h'+1}^{\pi^k})}}{\sqrt{n_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k))} \vee 1} + \frac{118H^2 S\tau}{n_{h'}^k(s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k)) \vee 1} \mid \pi^{k,h'} \right] \\ &+ \left(1 + \frac{1}{4H}\right) \mathbb{E}^k \left[\mathbb{E}_{p_{h'}(\cdot | s_{h'}^k, \underline{\pi}_{h'}^k(s_{h'}^k))}[\overline{V}_{h'+1}^k - \underline{V}_{h'+1}^k] \mid \pi^{k,h'} \right] \end{aligned}$$

Note that $q_{h'}^{\pi^{k,h'}}(s) = q_{h'}^{\pi^k}(s)$ since until step h' the policies are the same, i.e., $\pi_h^{h',k} = \pi_h^k$ for all $h < h'$. Hence, denoting the first term above by $\mathbb{E}^k[x(s_{h'}^k, \pi_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'}]$, we can write

$$\begin{aligned}\mathbb{E}^k[x(s_{h'}^k, \pi_{h'}^k(s_{h'}^k)) \mid \pi^{k,h'}] &= \sum_s q_{h'}^{\pi^{k,h'}}(s)x(s, \pi_{h'}^k(s)) = \sum_s q_{h'}^{\pi^k}(s)x(s, \pi_{h'}^k(s)) \\ &= \sum_{s,a} q_{h'}^{\pi^k}(s,a)x(s,a) = \mathbb{E}^k[x(s_{h'}^k, a_{h'}^k) \mid \pi^k],\end{aligned}$$

where the third equality is since π^k is deterministic. In a similar way, the second term can be bounded by,

$$\begin{aligned}\mathbb{E}^k\left[\mathbb{E}_{p_{h'}(\cdot \mid s_{h'}^k, \pi_{h'}^k(s_{h'}^k))}[\bar{V}_{h+1}^k - \underline{V}_{h+1}^k] \mid \pi^{k,h'}\right] &= \\ &= \sum_{s,s'} q_{h'}^{\pi^{k,h'}}(s)p_{h'}(\cdot \mid s_{h'}^k, \pi_{h'}^k(s_{h'}^k))\left(\bar{V}_{h'+1}^k(s') - \underline{V}_{h'+1}^k(s')\right) \\ &= \sum_{s,s'} q_{h'}^{\pi^k}(s)p_{h'}(\cdot \mid s_{h'}^k, \pi_{h'}^k(s_{h'}^k))\left(\bar{V}_{h'+1}^k(s') - \underline{V}_{h'+1}^k(s')\right) \\ &= \sum_s q_{h'+1}^{\pi^k}(s)\left(\bar{V}_{h'+1}^k(s') - \underline{V}_{h'+1}^k(s')\right) \\ &= \mathbb{E}^k\left[\bar{V}_{h'+1}^k(s_{h'+1}^k) - \underline{V}_{h'+1}^k(s_{h'+1}^k) \mid \pi^k\right] \\ &\leq \mathbb{E}^k\left[\bar{Q}_{h'+1}^k(s_{h'+1}^k, a_{h'+1}^k) - \underline{V}_{h'+1}^k(s_{h'+1}^k) \mid \pi^k\right] \\ &= \mathbb{E}^k\left[\bar{Q}_{h'+1}^k(s_{h'+1}^k, a_{h'+1}^k) - \underline{Q}_{h'+1}^k(s_{h'+1}^k, a_{h'+1}^k) \mid \pi^k\right].\end{aligned}$$

Lemma 38 (Recursion Optimistic Policy) *Let $h \in [H]$. Under the good event, for every $k \in [K]$ it holds that:*

$$\begin{aligned}\mathbb{E}^k\left[\bar{Q}_h^k(s_h^k, a_h^k) - \underline{Q}_h^k(s_h^k, a_h^k) \mid \pi^k\right] &\leq \mathbb{E}^k\left[\frac{8\sqrt{\tau \text{Var}_{p_h(\cdot \mid s_h^k, a_h^k)}(V_{h+1}^{\pi^k})}}{\sqrt{n_h^k(s_h^k, a_h^k) \vee 1}} + \frac{118H^2 S \tau}{n_h^k(s_h^k, a_h^k) \vee 1} \mid \pi^k\right] \\ &\quad + \left(1 + \frac{1}{4H}\right) \mathbb{E}^k\left[\bar{Q}_{h+1}^k(s_{h+1}^k, a_{h+1}^k) - \underline{Q}_{h+1}^k(s_{h+1}^k, a_{h+1}^k) \mid \pi^k\right].\end{aligned}$$

Proof Similar to the proof of Lemma 37. ■

Algorithm 6 COOPERATIVE O-REPS (COOP-O-REPS)

- 1: **input:** state space \mathcal{S} , action space \mathcal{A} , horizon H , transition function p , number of episodes K , number of agents m , exploration parameter γ , learning rate η .
 - 2: **initialize:** $\pi_h^1(a | s) = 1/A, q_h^1(s, a) = q_h^{\pi^1}(s, a) \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 3: **for** $k = 1, \dots, K$ **do**
 - 4: **for** $v = 1, \dots, m$ **do**
 - 5: observe initial state $s_1^{k,v}$.
 - 6: **for** $h = 1, \dots, H$ **do**
 - 7: pick action $a_h^{k,v} \sim \pi_h^k(\cdot | s_h^{k,v})$.
 - 8: suffer and observe cost $c_h^k(s_h^{k,v}, a_h^{k,v})$.
 - 9: observe next state $s_{h+1}^{k,v} \sim p_h(\cdot | s_h^{k,v}, a_h^{k,v})$.
 - 10: **end for**
 - 11: **end for**
 - 12: compute $W_h^k(s, a) = \Pr[\exists v : s_h^{k,v} = s, a_h^{k,v} = a | \pi^k] = 1 - (1 - q_h^k(s, a))^m \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 13: compute $\hat{c}_h^k(s, a) = \frac{c_h^k(s, a) \mathbb{I}\{\exists v: s_h^{k,v} = s, a_h^{k,v} = a\}}{W_h^k(s, a) + \gamma} \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 14: compute $q^{k+1} = \arg \min_{q \in \Delta(\mathcal{M})} \eta \langle q, \hat{c}^k \rangle + \text{KL}(q \| q^k)$.
 - 15: compute $\pi_h^{k+1}(a | s) = \frac{q_h^k(s, a)}{\sum_{a' \in \mathcal{A}} q_h^k(s, a')} \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 16: **end for**
-

Appendix D. The coop-O-REPS algorithm for adversarial MDPs with fresh randomness and known p

For the setting of adversarial MDPs with fresh randomness and known transitions we propose the Cooperative O-REPS algorithm (coop-O-REPS; see Algorithm 6). The idea is simple: all the agents run the same O-REPS algorithm, but the estimated costs are updated based on the trajectories of all of them. Since the randomness is fresh in this setting, we expect the agents to observe m times more information. Next, we prove the following optimal regret bound for coop-O-REPS.

Similarly to Zimin and Neu (2013), We use the notations $\Delta(\mathcal{M})$ and $\text{KL}(\cdot \| \cdot)$ for the set of occupancy measures of the MDP \mathcal{M} and the KL-divergence between occupancy measures, respectively.

Theorem 39 *With probability $1 - \delta$, setting $\eta = \gamma = \sqrt{\frac{\log \frac{HSA}{\delta}}{(1 + \frac{SA}{m})K}}$, the individual regret of each agent of coop-O-REPS is*

$$R_K = O \left(H \sqrt{K \log \frac{HSA}{\delta}} + \sqrt{\frac{H^2 SAK}{m} \log \frac{HSA}{\delta}} + \frac{HSA}{m} \log \frac{HSA}{\delta} \right).$$

D.1 The good event

Define the following events:

$$\begin{aligned} E^c &= \left\{ \sum_{k=1}^K \langle \mathbb{E}^k[\hat{c}^k] - \hat{c}^k, q^k \rangle \leq 4H \sqrt{K \log \frac{3}{\delta}} \right\} \\ E^{\hat{c}} &= \left\{ \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left(\frac{1}{m} + q_h^k(s, a) \right) (\hat{c}_h^k(s, a) - 2c_h^k(s, a)) \leq \frac{10HSA \log \frac{3HSA}{\delta}}{m\gamma} + \frac{10H \log \frac{3HSA}{\delta}}{\gamma} \right\} \\ E^* &= \left\{ \sum_{k=1}^K \langle \hat{c}^k - c^k, q^{\pi^*} \rangle \leq \frac{H \log \frac{3HSA}{\delta}}{\gamma} \right\} \end{aligned}$$

The good event is the intersection of the above events. The following lemma establishes that the good event holds with high probability.

Lemma 40 (The Good Event) *Let $\mathbb{G} = E^c \cap E^{\hat{c}} \cap E^*$ be the good event. It holds that $\Pr[\mathbb{G}] \geq 1 - \delta$.*

Proof We show that each of the events $\neg E^c$, $\neg E^{\hat{c}}$, $\neg E^*$ occur with probability at most $\delta/3$. Then, by a union bound we obtain the statement. Notice that:

$$\begin{aligned}
\langle \hat{c}^k, q^k \rangle &= \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) \hat{c}_h^k(s, a) \leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) \frac{\mathbb{I}\{\exists v : s_h^{k,v} = s, a_h^{k,v} = a\}}{W_h^k(s, a) + \gamma} \\
&\leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \frac{q_h^k(s, a)}{1 - (1 - q_h^k(s, a))^m} \mathbb{I}\{\exists v : s_h^{k,v} = s, a_h^{k,v} = a\} \\
&\leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left(\frac{1}{m} + q_h^k(s, a) \right) \mathbb{I}\{\exists v : s_h^{k,v} = s, a_h^{k,v} = a\} \\
&\leq \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) + \frac{1}{m} \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mathbb{I}\{\exists v : s_h^{k,v} = s, a_h^{k,v} = a\} \leq H + \frac{1}{m} \cdot Hm = 2H,
\end{aligned}$$

where the third inequality is by Lemma 41, and the last inequality follows because for each step h the agents visit at most m state-action pairs. Thus, event E^c holds by Azuma inequality.

Event $E^{\hat{c}}$ holds by Cohen et al. (2021, Lemma E.2) since $\sum_{h,s,a} \left(\frac{1}{m} + q_h^k(s, a) \right) \hat{c}_h^k(s, a) \leq \frac{1}{\gamma} \left(\frac{HSA}{m} + H \right)$ and $\mathbb{E}^k[\hat{c}_h^k(s, a)] \leq c_h^k(s, a)$. Event E^* holds by Jin et al. (2020a, Lemma 14). ■

D.2 Proof of Theorem 39

Proof of Theorem 39 By Lemma 40, the good event holds with probability $1 - \delta$. We now analyze the regret under the assumption that the good event holds. We start by decomposing the regret as follows:

$$\begin{aligned}
R_K &= \sum_{k=1}^K V_1^{k, \pi^k}(s_1^{k,v}) - V_1^{k, \pi^*}(s_1^{k,v}) = \sum_{k=1}^K \langle c^k, q^k - q^{\pi^*} \rangle \\
&= \underbrace{\sum_{k=1}^K \langle c^k - \hat{c}^k, q^k \rangle}_{(A)} + \underbrace{\sum_{k=1}^K \langle \hat{c}^k, q^k - q^{\pi^*} \rangle}_{(B)} + \underbrace{\sum_{k=1}^K \langle \hat{c}^k - c^k, q^{\pi^*} \rangle}_{(C)}.
\end{aligned}$$

Term (A) can be further decomposed as:

$$(A) = \sum_{k=1}^K \langle c^k - \hat{c}^k, q^k \rangle = \sum_{k=1}^K \langle c^k - \mathbb{E}^k[\hat{c}^k], q^k \rangle + \sum_{k=1}^K \langle \mathbb{E}^k[\hat{c}^k] - \hat{c}^k, q^k \rangle.$$

The second term is bounded by $4H\sqrt{K\log\frac{3}{\delta}}$ by the good event E^c , and for the first term:

$$\begin{aligned} \sum_{k=1}^K \langle c^k - \mathbb{E}^k[\hat{c}^k], q^k \rangle &= \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) c_h^k(s, a) \left(1 - \frac{\mathbb{E}^k[\mathbb{I}\{\exists v : s_h^{k,v} = s, a_h^{k,v} = a\}]}{W_h^k(s, a) + \gamma} \right) \\ &= \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) c_h^k(s, a) \left(1 - \frac{W_h^k(s, a)}{W_h^k(s, a) + \gamma} \right) \leq \gamma \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \frac{q_h^k(s, a)}{W_h^k(s, a) + \gamma} \\ &\leq \gamma \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \frac{q_h^k(s, a)}{1 - (1 - q_h^k(s, a))^m} \leq \gamma \sum_{k, h, s, a} \left(\frac{1}{m} + q_h^k(s, a) \right) = \gamma HK \left(1 + \frac{SA}{m} \right), \end{aligned}$$

where the last inequality is by Lemma 41.

Term (B) is bounded by OMD (see, e.g., Zimin and Neu (2013)) as follows:

$$\begin{aligned} (B) &= \sum_{k=1}^K \langle \hat{c}^k, q^k - q^{\pi^*} \rangle \leq \frac{H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) \hat{c}_h^k(s, a)^2 \\ &\leq \frac{H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) \frac{\hat{c}_h^k(s, a)}{W_h^k(s, a) + \gamma} \\ &\leq \frac{H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \hat{c}_h^k(s, a) \frac{q_h^k(s, a)}{1 - (1 - q_h^k(s, a))^m} \\ &\leq \frac{H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left(\frac{1}{m} + q_h^k(s, a) \right) \hat{c}_h^k(s, a) \\ &\leq \frac{H \log(HSA)}{\eta} + 2\eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left(\frac{1}{m} + q_h^k(s, a) \right) c_h^k(s, a) + \frac{10\eta HSA \log \frac{3HSA}{\delta}}{m\gamma} + \frac{10\eta H \log \frac{3HSA}{\delta}}{\gamma} \\ &\lesssim \frac{H \log \frac{HSA}{\delta}}{\eta} + \frac{\eta HSAK}{m} + \eta HK + \frac{\eta HSA \log \frac{3HSA}{\delta}}{m\gamma} + \frac{\eta H \log \frac{3HSA}{\delta}}{\gamma}, \end{aligned}$$

where the forth inequality is by Lemma 41, and the fifth inequality is by the good event $E^{\hat{c}}$.

Term (C) is bounded by $\frac{H \log \frac{3H}{\delta}}{\gamma}$ by the good event E^* . Putting the three terms together gives the final regret bound when setting $\eta = \gamma = \sqrt{\frac{\log \frac{HSA}{\delta}}{(1 + \frac{SA}{m})K}}$. ■

D.3 Auxiliary lemmas

Lemma 41 Let $x \in (0, 1)$. Then, $\frac{x}{1 - (1-x)^m} \leq \frac{1}{m} + x$.

Proof Using AM-GM inequality,

$$\begin{aligned} ((1-x)^m(1+xm))^{\frac{1}{m+1}} &\leq \frac{m(1-x) + 1 + xm}{m+1} = 1 \\ \implies (1-x)^m &\leq \frac{1}{1+xm} \\ \implies 1 - (1-x)^m &\geq \frac{xm}{1+xm} \\ \implies \frac{x}{1 - (1-x)^m} &\leq \frac{1}{m} + x. \end{aligned}$$
■

Algorithm 7 COOPERATIVE UOB-REPS (COOP-UOB-REPS)

- 1: **input:** state space \mathcal{S} , action space \mathcal{A} , horizon H , confidence parameter δ , number of episodes K , number of agents m , exploration parameter γ , learning rate η .
 - 2: **initialize:** $n_h^1(s, a) = 0, n_h^1(s, a, s') = 0, \pi_h^1(a | s) = 1/A, q_h^1(s, a, s') = 1/S^2 A \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 3: **for** $k = 1, \dots, K$ **do**
 - 4: set $n_h^{k+1}(s, a) \leftarrow n_h^k(s, a), n_h^{k+1}(s, a, s') \leftarrow n_h^k(s, a, s') \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 5: **for** $v = 1, \dots, m$ **do**
 - 6: observe initial state $s_1^{k,v}$.
 - 7: **for** $h = 1, \dots, H$ **do**
 - 8: pick action $a_h^{k,v} \sim \pi_h^k(\cdot | s_h^{k,v})$.
 - 9: suffer and observe cost $c_h^k(s_h^{k,v}, a_h^{k,v})$.
 - 10: observe next state $s_{h+1}^{k,v} \sim p_h(\cdot | s_h^{k,v}, a_h^{k,v})$.
 - 11: update $n_h^{k+1}(s_h^{k,v}, a_h^{k,v}) \leftarrow n_h^{k+1}(s_h^{k,v}, a_h^{k,v}) + 1, n_h^{k+1}(s_h^{k,v}, a_h^{k,v}, s_{h+1}^{k,v}) \leftarrow n_h^{k+1}(s_h^{k,v}, a_h^{k,v}, s_{h+1}^{k,v}) + 1$.
 - 12: **end for**
 - 13: **end for**
 - 14: set $\hat{p}_h^{k+1}(s' | s, a) \leftarrow \frac{n_h^{k+1}(s, a, s')}{n_h^{k+1}(s, a) \vee 1} \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 15: compute confidence set for $\epsilon_h^{k+1}(s' | s, a) = 4\sqrt{\frac{\hat{p}_h^{k+1}(s' | s, a) \ln \frac{HSAK}{4\delta}}{n_h^{k+1}(s, a) \vee 1}} + 10\frac{\ln \frac{HSAK}{4\delta}}{n_h^{k+1}(s, a) \vee 1}$:
$$\mathcal{P}^{k+1} = \{p' | \forall (s, a, s', h) : |\hat{p}_h^{k+1}(s' | s, a) - p'_h(s' | s, a)| \leq \epsilon_h^{k+1}(s' | s, a)\}.$$
 - 16: compute $u_h^k(s) = \max_{p' \in \mathcal{P}^k} q_h^{p', \pi^k}(s) = \max_{p' \in \mathcal{P}^k} \Pr[s_h = s | \pi^k, p'] \forall s \in \mathcal{S}$.
 - 17: compute $U_h^k(s, a) = 1 - (1 - u_h^k(s, a))^m \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 18: compute $\hat{c}_h^k(s, a) = \frac{c_h^k(s, a) \mathbb{I}\{\exists v: s_h^{k,v} = s, a_h^{k,v} = a\}}{U_h^k(s, a) + \gamma} \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 19: compute $q^{k+1} = \arg \min_{q \in \Delta(\mathcal{M}, k+1)} \eta \langle q, \hat{c}^k \rangle + \text{KL}(q \| q^k)$.
 - 20: compute $\pi_h^{k+1}(a | s) = \frac{q_h^{k+1}(s, a)}{\sum_{a' \in \mathcal{A}} q_h^{k+1}(s, a')} \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$, where $q_h^{k+1}(s, a) = \sum_{s' \in \mathcal{S}} q_h^{k+1}(s, a, s')$.
 - 21: **end for**
-

Appendix E. The coop-UOB-REPS algorithm for adversarial MDPs with fresh randomness and unknown p

For the setting of adversarial MDPs with fresh randomness and unknown transitions we propose the Cooperative UOB-REPS algorithm (coop-UOB-REPS; see Algorithm 7). The idea is simple: all the agents run the same UOB-REPS algorithm, but the estimated costs and transitions are updated based on the trajectories of all of them. Since the randomness is fresh in this setting, we expect the agents to observe m times more information. Next, we prove the following regret bound for coop-UOB-REPS. Note that this bound is optimal up to a \sqrt{HS} factor. Removing this extra \sqrt{HS} is an open-problem even for adversarial MDPs with a single agent.

Similarly to Rosenberg and Mansour (2019a), We use the notation $\Delta(\mathcal{M}, k)$ for the set of occupancy measures whose induced transition function is within the confidence set \mathcal{P}^k .

Theorem 42 *With probability $1 - \delta$, setting $\eta = \gamma = \sqrt{\frac{\log \frac{mKHSA}{\delta}}{(1 + \frac{SA}{m})K}}$, the individual regret of each agent of coop-UOB-REPS is*

$$R_K = O \left(H \sqrt{K \log \frac{mKHSA}{\delta}} + \sqrt{\frac{H^4 S^2 AK}{m} \log^3 \frac{mKHSA}{\delta}} + H^3 S^3 A \log^3 \frac{mKHSA}{\delta} \right).$$

E.1 The good event

Denote $\epsilon_h^k(s' | s, a) = \sqrt{\frac{2\hat{p}_h^k(s'|s,a) \log \frac{30KHSA}{\delta}}{n_h^k(s,a)\sqrt{1}}} + \frac{2 \log \frac{30KHSA}{\delta}}{n_h^k(s,a)\sqrt{1}}$, $\epsilon_h^k(s, a) = \sum_{s' \in \mathcal{S}} \epsilon_h^k(s' | s, a)$, $\tilde{\epsilon}_h^k(s' | s, a) = 8\sqrt{\frac{p_h(s'|s,a) \log \frac{30KHSA}{\delta}}{n_h^k(s,a)\sqrt{1}}} + \frac{100 \log \frac{30KHSA}{\delta}}{n_h^k(s,a)\sqrt{1}}$, and $\tilde{\epsilon}_h^k(s, a) = \sum_{s'} \tilde{\epsilon}_h^k(s' | s, a)$. Define the following events:

$$\begin{aligned}
E^P &= \left\{ \forall (k, s, a, s', h) : |p_h(s'|s, a) - \hat{p}_h^k(s'|s, a)| \leq \epsilon_h^k(s' | s, a) \right\} \\
E^{on1} &= \left\{ \forall v \in [m] : \sum_{k, h, s, a} \left(q_h^{\pi^k}(s, a) - \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\} \right) \min\{2, \epsilon_h^k(s, a)\} \leq 10\sqrt{K \log \frac{30KHSA m}{\delta}} \right\} \\
E^{on2} &= \left\{ \forall v \in [m] : \sum_{k, h, s, a} q_h^{\pi^k}(s, a) \tilde{\epsilon}_h^k(s, a) \leq 2 \sum_{k, h, s, a} \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\} \tilde{\epsilon}_h^k(s, a) + 100HS \log^2 \frac{30KHSA m}{\delta} \right\} \\
E^{on3} &= \left\{ \forall v \in [m] : \sum_{k, s, a, h} \frac{q_h^{\pi^k}(s, a)}{n_h^k(s, a)} \leq 2 \sum_{k, s, a, h} \frac{\mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a)} + H \log \frac{m}{\delta} \right\} \\
E^c &= \left\{ \sum_{k=1}^K \langle \mathbb{E}^k[\hat{c}^k] - \hat{c}^k, q^k \rangle \leq 4H\sqrt{K \log \frac{6}{\delta}} \right\} \\
E^{\hat{c}} &= \left\{ \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left(\frac{1}{m} + q_h^k(s, a) \right) (\hat{c}_h^k(s, a) - c_h^k(s, a)) \leq \frac{HSA \log \frac{3HSA}{\delta}}{\gamma} \right\} \\
E^* &= \left\{ \sum_{k=1}^K \langle \hat{c}^k - c^k, q^{\pi^*} \rangle \leq \frac{H \log \frac{3HSA}{\delta}}{\gamma} \right\}
\end{aligned}$$

The good event is the intersection of the above events. The following lemma establishes that the good event holds with high probability.

Lemma 43 (The Good Event) *Let $\mathbb{G} = E^P \cap E^{on1} \cap E^{on2} \cap E^{on3} \cap E^c \cap E^{\hat{c}} \cap E^*$ be the good event. It holds that $\Pr[\mathbb{G}] \geq 1 - \delta$.*

Proof E^{on2} and E^{on3} follows [Cohen et al. \(2021, Lemma E.2\)](#). The rest are similar to the proofs of [Lemmas 17 and 40](#) and to proofs in [Jin et al. \(2020a\)](#). ■

E.2 Proof of Theorem 42

Proof of Theorem 42 By Lemma 43, the good event holds with probability $1 - \delta$. We now analyze the regret under the assumption that the good event holds. We start by decomposing the regret as follows:

$$\begin{aligned} R_K &= \sum_{k=1}^K V_1^{k, \pi^k}(s_1^{k,v}) - V_1^{k, \pi^*}(s_1^{k,v}) = \sum_{k=1}^K \langle c^k, q^{\pi^k} - q^{\pi^*} \rangle \\ &= \underbrace{\sum_{k=1}^K \langle c^k, q^{\pi^k} - q^k \rangle}_{(A)} + \underbrace{\sum_{k=1}^K \langle c^k - \hat{c}^k, q^k \rangle}_{(B)} + \underbrace{\sum_{k=1}^K \langle \hat{c}^k, q^k - q^{\pi^*} \rangle}_{(C)} + \underbrace{\sum_{k=1}^K \langle \hat{c}^k - c^k, q^{\pi^*} \rangle}_{(D)}. \end{aligned}$$

Let $\tau = \log \frac{KHSAm}{\delta}$ be a logarithmic term. Term (A) can be decomposed using the value difference lemma (see, e.g., Shani et al. (2020)):

$$\begin{aligned} (A) &= \sum_{k=1}^K \langle c^k, q^{\pi^k} - q^k \rangle \leq 2H \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^{\pi^k}(s, a) \|p_h(\cdot | s, a) - \hat{p}_h^k(\cdot | s, a)\|_1 \\ &\leq 2H \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^{\pi^k}(s, a) \min\{2, \epsilon_h^k(s, a)\} \\ &\lesssim \frac{H}{m} \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\} \epsilon_h^k(s, a) + H\sqrt{K\tau} \\ &\lesssim \frac{H\sqrt{S\tau}}{m} \sum_{k,h,s,a} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{\sqrt{n_h^k(s, a) \vee 1}} + \frac{HS\tau}{m} \sum_{k,h,s,a} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a) \vee 1} + H\sqrt{K\tau}, \end{aligned}$$

where the second inequality is by event E^p , and the third inequality uses event E^{on1} and Cauchy–Schwarz inequality. We now bound each of the two sums separately. For the second sum recall that $n_h^k(s, a) = \sum_{j=1}^{k-1} \sum_{v=1}^m \mathbb{I}\{s_h^{j,v} = s, a_h^{j,v} = a\}$, thus we have:

$$\begin{aligned} \sum_{k,h,s,a} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a) \vee 1} &\leq 2HSAm + \sum_{h,s,a} \sum_{k: n_h^k(s,a) \geq m} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s, a)} \\ &= 2HSAm + \sum_{h,s,a} \sum_{k: n_h^k(s,a) \geq m} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{\sum_{j=1}^{k-1} \sum_{v=1}^m \mathbb{I}\{s_h^{j,v} = s, a_h^{j,v} = a\}} \\ &\leq 2HSAm + HSA \log(Km), \end{aligned} \tag{18}$$

where the last inequality is by [Rosenberg et al. \(2020, Lemma B.18\)](#). For the first term:

$$\begin{aligned}
\sum_{k,h,s,a} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{\sqrt{n_h^k(s,a) \vee 1}} &\leq 2HSA m + \sum_{h,s,a} \sum_{k: n_h^k(s,a) \geq m} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{\sqrt{n_h^k(s,a)}} \\
&= 2HSA m + \sum_{h,s,a} \sum_{k: n_h^k(s,a) \geq m} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{\sqrt{n_h^{k+1}(s,a)}} \sqrt{\frac{n_h^{k+1}(s,a)}{n_h^k(s,a)}} \\
&\leq 2HSA m + \sum_{h,s,a} \sum_{k: n_h^k(s,a) \geq m} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{\sqrt{n_h^{k+1}(s,a)}} \sqrt{\frac{n_h^k(s,a) + m}{n_h^k(s,a)}} \\
&\leq 2HSA m + 2 \sum_{h,s,a} \sum_{k: n_h^k(s,a) \geq m} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{\sqrt{n_h^{k+1}(s,a)}} \\
&= 2HSA m + 2 \sum_{h,s,a} \sum_{k: n_h^k(s,a) \geq m} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{\sqrt{\sum_{j=1}^k \sum_{v=1}^m \mathbb{I}\{s_h^{j,v} = s, a_h^{j,v} = a\}}} \\
&\leq 2HSA m + 4 \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sqrt{\sum_{k=1}^K \sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}} \\
&\leq 2HSA m + 4 \sum_{h=1}^H \sqrt{SA \sum_{k=1}^K \sum_{v=1}^m \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}} \\
&= 2HSA m + 4H\sqrt{SAKm}, \tag{19}
\end{aligned}$$

where the forth inequality is by [Streeter and McMahan \(2010, Lemma 1\)](#), and the last inequality is by Jensen's inequality. Putting these together we get that: $(A) \lesssim \sqrt{\frac{H^4 S^2 AK\tau}{m}} + H^2 S^2 A\tau^2 + H\sqrt{K\tau}$.

Term (B) can be further decomposed as:

$$(B) = \sum_{k=1}^K \langle c^k - \hat{c}^k, q^k \rangle = \sum_{k=1}^K \langle c^k - \mathbb{E}^k[\hat{c}^k], q^k \rangle + \sum_{k=1}^K \langle \mathbb{E}^k[\hat{c}^k] - \hat{c}^k, q^k \rangle.$$

The second term is bounded by $4H\sqrt{K\log\frac{6}{\delta}}$ by the good event E^c . For the first term, let

$$W_h^{\pi^k}(s, a) = \Pr[\exists v : s_h^{k,v} = s, a_h^{k,v} = a \mid \pi^k] = 1 - (1 - q_h^{\pi^k}(s, a))^m:$$

$$\begin{aligned} \sum_{k=1}^K \langle c^k - \mathbb{E}^k[\hat{c}^k], q^k \rangle &= \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) c_h^k(s, a) \left(1 - \frac{\mathbb{E}^k[\mathbb{I}\{\exists v : s_h^{k,v} = s, a_h^{k,v} = a\}]}{U_h^k(s, a) + \gamma} \right) \\ &= \sum_{k,h,s,a} q_h^k(s, a) c_h^k(s, a) \left(1 - \frac{W_h^{\pi^k}(s, a)}{U_h^k(s, a) + \gamma} \right) \leq \sum_{k,h,s,a} \frac{q_h^k(s, a)}{W_h^{\pi^k}(s, a)} (U_h^k(s, a) - W_h^{\pi^k}(s, a) + \gamma) \\ &\leq \sum_{k,h,s,a} \left(\frac{1}{m} + q_h^k(s, a) \right) (U_h^k(s, a) - W_h^{\pi^k}(s, a) + \gamma) \\ &= \sum_{k,h,s,a} \left(\frac{1}{m} + q_h^k(s, a) \right) \left((1 - q_h^{\pi^k}(s, a))^m - (1 - u_h^k(s, a))^m \right) + \frac{\gamma HSAK}{m} + \gamma HK \\ &\leq \sum_{k,h,s,a} \left(\frac{1}{m} + q_h^k(s, a) \right) m(1 - q_h^{\pi^k}(s, a))^m (u_h^k(s, a) - q_h^{\pi^k}(s, a)) + \frac{\gamma HSAK}{m} + \gamma HK \\ &\leq \sum_{k,h,s,a} u_h^k(s, a) - q_h^{\pi^k}(s, a) + \frac{\gamma HSAK}{m} + \gamma HK \\ &\quad + \sum_{k,h,s,a} m q_h^k(s, a) (1 - q_h^{\pi^k}(s, a))^m (u_h^k(s, a) - q_h^{\pi^k}(s, a)) \\ &\leq 3 \sum_{k,h,s,a} (u_h^k(s, a) - q_h^{\pi^k}(s, a)) \log m + \frac{\gamma HSAK}{m} + \gamma HK, \end{aligned}$$

where the second inequality is by Lemma 41, the third inequality is by convexity of the function $f(x) = (1-x)^m$ for $x \in [0, 1]$, and the last inequality follows because $m x(1-x)^m \leq \log m$ for every $x \in [0, 1]$ since if $1-x \leq 1 - \frac{2\log m}{m}$ then $(1-x)^m \leq \frac{1}{m^2}$; otherwise $x \leq \frac{2\log m}{m}$. Finally, $\sum_{k,h,s,a} (u_h^k(s, a) - q_h^{\pi^k}(s, a))$ is bounded by Lemma 46.

Term (C) is bounded by OMD (see, e.g., Rosenberg and Mansour (2019a)) as follows:

$$\begin{aligned} (C) &= \sum_{k=1}^K \langle \hat{c}^k, q^k - q^{\pi^*} \rangle \leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) \hat{c}_h^k(s, a)^2 \\ &\leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k,h,s,a} q_h^k(s, a) \frac{\hat{c}_h^k(s, a)}{U_h^k(s, a) + \gamma} \leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k,h,s,a} \hat{c}_h^k(s, a) \frac{q_h^k(s, a)}{1 - (1 - q_h^k(s, a))^m} \\ &\leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k,h,s,a} \left(\frac{1}{m} + q_h^k(s, a) \right) \hat{c}_h^k(s, a) \\ &\leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k,h,s,a} \left(\frac{1}{m} + q_h^k(s, a) \right) c_h^k(s, a) + \frac{\eta HSA \log \frac{3HSA}{\delta}}{\gamma} \\ &\leq \frac{2H \log \frac{HSA}{\delta}}{\eta} + \frac{\eta HSAK}{m} + \eta HK + \frac{\eta HSA \log \frac{3HSA}{\delta}}{\gamma}, \end{aligned}$$

where the fourth inequality is by Lemma 41, and the fifth inequality is by the good event $E^{\hat{c}}$.

Term (D) is bounded by $\frac{H \log \frac{3H}{\delta}}{\gamma}$ by the good event E^* . Putting the three terms together gives the final regret bound when setting $\eta = \gamma = 1/\sqrt{(1 + \frac{SA}{m})K}$. ■

E.3 Auxiliary Lemmas

The following Lemma is by [Jin et al. \(2020a, Lemma 8\)](#), [Cohen et al. \(2021, Lemma B.13\)](#).

Lemma 44 *Under the good event we have,*

$$\forall(k, s, a, s', h) : \quad |p_h(s'|s, a) - \hat{p}_h^k(s'|s, a)| \leq \tilde{\epsilon}_h^k(s' | s, a).$$

The following Lemma is part of the proof of [Jin et al. \(2020a, Lemma 4\)](#). We provide the proof here for completeness.

Lemma 45 *Let $q_h^{\pi^k}(\tilde{s}, \tilde{a} | \tilde{s}'; h)$ be the probability to visit (\tilde{s}, \tilde{a}) in time \tilde{h} given that we visited \tilde{s}' in time h . Under the good event,*

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} |u_h^k(s, a) - q_h^{\pi^k}(s, a)| &\lesssim H \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \tilde{\epsilon}_h^k(s, a) q_h^{\pi^k}(s, a) \\ &+ HS \sum_{k=1}^K \sum_{1 \leq h < \tilde{h} \leq H} \sum_{s \in \mathcal{S}, a \in \mathcal{A}, s' \in \mathcal{S}} \sum_{\tilde{s} \in \mathcal{S}, \tilde{a} \in \mathcal{A}} \tilde{\epsilon}_h^k(s' | s, a) q_h^{\pi^k}(s, a) \min \left\{ 2, \sum_{\tilde{s}' \in \mathcal{S}} \tilde{\epsilon}_h^k(\tilde{s}' | \tilde{s}, \tilde{a}) \right\} q_h^{\pi^k}(\tilde{s}, \tilde{a} | s'; h + 1). \end{aligned} \quad (20)$$

Proof Let $q^{k,s,h}$ be the occupancy measure such that $q_h^{k,s,h}(s) = u_h^k(s)$, and let $p^{k,s,h}$ be the transition that corresponds to $q^{k,s,h}$. Let $\sigma_h(s)$ be the set of all trajectories that end in s in time h , i.e., $\sigma_h(s) = \{s_1, a_1, \dots, s_{h-1}, a_{h-1}, s_h\}$ where $s_h = s$. We have:

$$\begin{aligned} u_h^k(s, a) &= q_h^{k,s,h}(s, a) = \pi_h^k(a | s) \sum_{\sigma_h(s)} \prod_{h'=1}^{h-1} \pi_{h'}^k(a_{h'} | s_{h'}) p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \\ q_h^{\pi^k}(s, a) &= \pi_h^k(a | s) \sum_{\sigma_h(s)} \prod_{h'=1}^{h-1} \pi_{h'}^k(a_{h'} | s_{h'}) p_{h'}(s_{h'+1} | s_{h'}, a_{h'}). \end{aligned}$$

Then,

$$|u_h^k(s, a) - q_h^{\pi^k}(s, a)| = \pi_h^k(a | s) \sum_{\sigma_h(s)} \prod_{h'=1}^{h-1} \pi_{h'}^k(a_{h'} | s_{h'}) \left| \prod_{h'=1}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) - \prod_{h'=1}^{h-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \right|.$$

We can rewrite the following term as,

$$\begin{aligned}
& \left| \prod_{h'=1}^{h-1} p^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) - \prod_{h'=1}^{h-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \right| \\
&= \left| \sum_{l=2}^{h-1} \prod_{h'=1}^{l-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \prod_{h'=l}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) + \prod_{h'=1}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \right. \\
&\quad \left. - \prod_{h'=1}^{h-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) - \sum_{l=2}^{h-1} \prod_{h'=1}^{l-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \prod_{h'=l}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \right| \\
&= \left| \sum_{l=1}^{h-1} \prod_{h'=1}^{l-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \prod_{h'=l}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \right. \\
&\quad \left. - \sum_{l=2}^h \prod_{h'=1}^{l-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \prod_{h'=l}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \right| \\
&= \left| \sum_{l=1}^{h-1} \prod_{h'=1}^{l-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \prod_{h'=l}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \right. \\
&\quad \left. - \sum_{l=1}^{h-1} \prod_{h'=1}^l p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \prod_{h'=l+1}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \right| \\
&= \sum_{l=1}^{h-1} \left| p_l^{k,s,h}(s_{l+1} | s_l, a_l) - p_l(s_{l+1} | s_l, a_l) \right| \prod_{h'=1}^{l-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \prod_{h'=l+1}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}).
\end{aligned}$$

Hence,

$$\begin{aligned}
& |u_h^k(s, a) - q_h^{\pi^k}(s, a)| \\
&\leq \pi_h^k(a | s) \sum_{\sigma_h(s)} \prod_{h'=1}^{h-1} \pi_{h'}^k(a_{h'} | s_{h'}) \sum_{l=1}^{h-1} \left| p_l^{k,s,h}(s_{l+1} | s_l, a_l) - p_l(s_{l+1} | s_l, a_l) \right| \\
&\quad \cdot \prod_{h'=1}^{l-1} p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \prod_{h'=l+1}^{h-1} p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \\
&\leq \sum_{l=1}^{h-1} \sum_{\sigma_h(s)} \left| p_l^{k,s,h}(s_{l+1} | s_l, a_l) - p_l(s_{l+1} | s_l, a_l) \right| \left(\pi_l^k(a_l | s_l) \prod_{h'=1}^{l-1} \pi_{h'}^k(a_{h'} | s_{h'}) p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \right) \\
&\quad \cdot \left(\pi_h^k(a | s) \prod_{h'=l+1}^{h-1} \pi_{h'}^k(a_{h'} | s_{h'}) p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \right) \\
&= \sum_{l=1}^{h-1} \sum_{s_l \in \mathcal{S}, a_l \in \mathcal{A}, s_{l+1} \in \mathcal{S}} \left| p_l^{k,s,h}(s_{l+1} | s_l, a_l) - p_l(s_{l+1} | s_l, a_l) \right| \left(\sum_{\sigma_l(s_l)} \pi_l^k(a_l | s_l) \prod_{h'=1}^{l-1} \pi_{h'}^k(a_{h'} | s_{h'}) p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \right) \\
&\quad \cdot \left(\sum_{a_{l+1} \in \mathcal{A}} \sum_{\{s_{h''} \in \mathcal{S}, a_{h''} \in \mathcal{A}\}_{h''=l+2}^{h-1}} \pi_h^k(a | s) \prod_{h'=l+1}^{h-1} \pi_{h'}^k(a_{h'} | s_{h'}) p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) \right) \\
&= \sum_{l=1}^{h-1} \sum_{s_l \in \mathcal{S}, a_l \in \mathcal{A}, s_{l+1} \in \mathcal{S}} \left| p_l^{k,s,h}(s_{l+1} | s_l, a_l) - p_l(s_{l+1} | s_l, a_l) \right| q_l^{\pi^k}(s_l, a_l) \cdot q_h^{k,s,h}(s, a | s_{l+1}),
\end{aligned}$$

where we ease notation and denote $q_h^{k,s,h}(s, a | s_{l+1}) = q_h^{k,s,h}(s, a | s_{l+1}; l+1)$. Similarly, we can show that,

$$\begin{aligned}
& |q_h^{k,s,h}(s, a | s_{l+1}) - q_h^{\pi^k}(s, a | s_{l+1})| \\
& \lesssim \sum_{h'=l+1}^{h-1} \sum_{s_{h'} \in \mathcal{S}, a_{h'} \in \mathcal{A}, s_{h'+1} \in \mathcal{S}} \left| p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) - p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \right| q_{h'}^{\pi^k}(s_{h'}, a_{h'} | s_{l+1}) q_{h'}^{k,s,h}(s, a | s_{h'+1}) \\
& \leq \pi_h^k(a | s) \sum_{h'=l+1}^{h-1} \sum_{s_{h'} \in \mathcal{S}, a_{h'} \in \mathcal{A}, s_{h'+1} \in \mathcal{S}} \left| p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) - p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \right| q_{h'}^{\pi^k}(s_{h'}, a_{h'} | s_{l+1}),
\end{aligned}$$

where the last is since $q_{h'}^{k,s,h}(s, a | s_{h'+1}) \leq \pi_h^k(a | s)$. Combining the last two,

$$\begin{aligned}
& \sum_{h,s,a,k} |u_h^k(s,a) - q_h^{\pi^k}(s,a)| \\
& \lesssim \sum_{h,s,a,k} \sum_{l=1}^{h-1} \sum_{s_l \in \mathcal{S}, a_l \in \mathcal{A}, s_{l+1} \in \mathcal{S}} \left| p_l^{k,s,h}(s_{l+1} | s_l, a_l) - p_l(s_{l+1} | s_l, a_l) \right| q_l^{\pi^k}(s_l, a_l) \cdot q_h^{k,s,h}(s, a | s_{l+1}) \\
& \leq \sum_{h,s,a,k} \sum_{l=1}^{h-1} \sum_{s_l \in \mathcal{S}, a_l \in \mathcal{A}, s_{l+1} \in \mathcal{S}} \tilde{\epsilon}_l^k(s_{l+1} | s_l, a_l) q_l^{\pi^k}(s_l, a_l) \cdot q_h^{\pi^k}(s, a | s_{l+1}) \\
& \quad + \sum_{h,s,a,k} \sum_{l=1}^{h-1} \sum_{s_l \in \mathcal{S}, a_l \in \mathcal{A}, s_{l+1} \in \mathcal{S}} \tilde{\epsilon}_l^k(s_{l+1} | s_l, a_l) q_l^{\pi^k}(s_l, a_l) \pi_h^k(a | s) \\
& \quad \cdot \left(\sum_{h'=l+1}^{h-1} \sum_{s_{h'} \in \mathcal{S}, a_{h'} \in \mathcal{A}, s_{h'+1} \in \mathcal{S}} \left| p_{h'}^{k,s,h}(s_{h'+1} | s_{h'}, a_{h'}) - p_{h'}(s_{h'+1} | s_{h'}, a_{h'}) \right| q_{h'}^{\pi^k}(s_{h'}, a_{h'} | s_{l+1}) \right) \\
& \leq \sum_{k,h} \sum_{l=1}^{h-1} \sum_{s_l \in \mathcal{S}, a_l \in \mathcal{A}, s_{l+1} \in \mathcal{S}} \tilde{\epsilon}_l^k(s_{l+1} | s_l, a_l) q_l^{\pi^k}(s_l, a_l) \cdot \left(\sum_{s,a} q_h^{\pi^k}(s, a | s_{l+1}) \right) \\
& \quad + \sum_{h,s,k} \sum_{l=1}^{h-1} \sum_{s_l \in \mathcal{S}, a_l \in \mathcal{A}, s_{l+1} \in \mathcal{S}} \tilde{\epsilon}_l^k(s_{l+1} | s_l, a_l) q_l^{\pi^k}(s_l, a_l) \sum_a \pi_h^k(a | s) \\
& \quad \cdot \left(\sum_{h'=l+1}^{h-1} \sum_{s_{h'} \in \mathcal{S}, a_{h'} \in \mathcal{A}} \min \left\{ 2, \sum_{s_{h'+1} \in \mathcal{S}} \tilde{\epsilon}_{h'}^k(s_{h'+1} | s_{h'}, a_{h'}) \right\} q_{h'}^{\pi^k}(s_{h'}, a_{h'} | s_{l+1}) \right) \\
& \leq H \sum_{k=1}^K \sum_{1 \leq l \leq H} \sum_{s_l \in \mathcal{S}, a_l \in \mathcal{A}, s_{l+1} \in \mathcal{S}} \tilde{\epsilon}_l^k(s_{l+1} | s_l, a_l) q_l^{\pi^k}(s_l, a_l) \\
& \quad + HS \sum_{k=1}^K \sum_{1 \leq l < h' \leq H} \sum_{s_l \in \mathcal{S}, a_l \in \mathcal{A}, s_{l+1} \in \mathcal{S}} \sum_{s_{h'} \in \mathcal{S}, a_{h'} \in \mathcal{A}} \tilde{\epsilon}_l^k(s_{l+1} | s_l, a_l) q_l^{\pi^k}(s_l, a_l) \\
& \quad \cdot \min \left\{ 2, \sum_{s_{h'+1} \in \mathcal{S}} \tilde{\epsilon}_{h'}^k(s_{h'+1} | s_{h'}, a_{h'}) \right\} q_{h'}^{\pi^k}(s_{h'}, a_{h'} | s_{l+1}) \\
& = H \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}, a \in \mathcal{A}, s' \in \mathcal{S}} \tilde{\epsilon}_h^k(s' | s, a) q_h^{\pi^k}(s, a) \\
& \quad + HS \sum_{k=1}^K \sum_{1 \leq h < \tilde{h} \leq H} \sum_{s \in \mathcal{S}, a \in \mathcal{A}, s' \in \mathcal{S}} \sum_{\tilde{s} \in \mathcal{S}, \tilde{a} \in \mathcal{A}} \tilde{\epsilon}_h^k(s' | s, a) q_h^{\pi^k}(s, a) \min \left\{ 2, \sum_{\tilde{s}' \in \mathcal{S}} \tilde{\epsilon}_{\tilde{h}}^k(\tilde{s}' | \tilde{s}, \tilde{a}) \right\} q_{\tilde{h}}^{\pi^k}(\tilde{s}, \tilde{a} | s'; h+1),
\end{aligned}$$

where the last inequality is by Lemma 44 and since $p_l^{k,s,h}(\cdot | s, a), p_l(\cdot | s, a)$ are probability distributions $\forall (s, a, l)$. \blacksquare

Lemma 46 *Under the good event,*

$$\sum_{h,s,a,k} |u_h^k(s,a) - q_h^{\pi^k}(s,a)| \lesssim \sqrt{\frac{H^4 S^2 A K \tau}{m}} + H^3 S^3 A \tau^2.$$

Proof We first bound $\sum_{h,s,a,k} |u_h^k(s,a) - q_h^{\pi^k}(s,a)|$ using Lemma 45. Now, for the first term in Eq. (20):

$$\begin{aligned}
\sum_{k,h,s,a} \tilde{\epsilon}_h^k(s,a) q_h^{\pi^k}(s,a) &= \frac{1}{m} \sum_{v=1}^m \sum_{k,h,s,a} \tilde{\epsilon}_h^k(s,a) q_h^{\pi^k}(s,a) \leq \frac{2}{m} \sum_{v=1}^m \sum_{k,h,s,a} \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\} \tilde{\epsilon}_h^k(s,a) + 18HS\tau^2 \\
&\lesssim \frac{\sqrt{S\tau}}{m} \sum_{k,h,s,a} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{\sqrt{n_h^k(s,a) \vee 1}} \\
&\quad + \frac{S\tau}{m} \sum_{k,h,s,a} \frac{\sum_{v=1}^m \mathbb{I}\{s_h^{k,v} = s, a_h^{k,v} = a\}}{n_h^k(s,a) \vee 1} + HS\tau^2 \\
&\lesssim \sqrt{\frac{H^2 S^2 AK\tau}{m}} + HS^2 A\tau^2,
\end{aligned}$$

where the first inequality is by event E^{on2} and the last inequality is by Eqs. (18) and (19). Plugging the definition of $\tilde{\epsilon}$, we break the second sum in Eq. (20) as follows:

$$\begin{aligned}
&\sum_{k=1}^K \sum_{1 \leq h < \bar{h} \leq H} \sum_{s,a,s'} \sum_{\tilde{s}, \tilde{a}} \tilde{\epsilon}_h^k(s' | s, a) q_h^{\pi^k}(s, a) \cdot \min \left\{ 2, \sum_{\tilde{s}'} \tilde{\epsilon}_h^k(\tilde{s}' | \tilde{s}, \tilde{a}) \right\} q_h^{\pi^k}(\tilde{s}, \tilde{a} | s'; h+1) \\
&\lesssim \sum_{k=1}^K \sum_{1 \leq h < \bar{h} \leq H} \sum_{s,a,s'} \sum_{\tilde{s}, \tilde{a}, \tilde{s}'} \sqrt{\frac{p_h^k(s'|s,a)\tau}{n_h^k(s,a) \vee 1}} q_h^{\pi^k}(s, a) \cdot \sqrt{\frac{p_h^k(\tilde{s}'|\tilde{s},\tilde{a})\tau}{n_h^k(\tilde{s},\tilde{a}) \vee 1}} q_h^{\pi^k}(\tilde{s}, \tilde{a} | s'; h+1) \\
&\quad + \sum_{k=1}^K \sum_{1 \leq h < \bar{h} \leq H} \sum_{s,a,s'} \sum_{\tilde{s}, \tilde{a}} \sqrt{\frac{p_h^k(s'|s,a)\tau}{n_h^k(s,a) \vee 1}} q_h^{\pi^k}(s, a) \cdot \min \left\{ 2, \frac{S\tau}{n_h^k(\tilde{s}, \tilde{a}) \vee 1} \right\} q_h^{\pi^k}(\tilde{s}, \tilde{a} | s'; h+1) \\
&\quad + \sum_{k=1}^K \sum_{1 \leq h < \bar{h} \leq H} \sum_{s,a,s'} \sum_{\tilde{s}, \tilde{a}} \frac{\tau}{n_h^k(s,a) \vee 1} q_h^{\pi^k}(s, a) \cdot \min \left\{ 2, \sum_{\tilde{s}'} \tilde{\epsilon}_h^k(\tilde{s}' | \tilde{s}, \tilde{a}) \right\} q_h^{\pi^k}(\tilde{s}, \tilde{a} | s'; h+1) \\
&\lesssim \underbrace{\sum_{k=1}^K \sum_{1 \leq h < \bar{h} \leq H} \sum_{s,a,s'} \sum_{\tilde{s}, \tilde{a}, \tilde{s}'} \sqrt{\frac{p_h^k(s'|s,a)\tau}{n_h^k(s,a) \vee 1}} q_h^{\pi^k}(s, a) \cdot \sqrt{\frac{p_h^k(\tilde{s}'|\tilde{s},\tilde{a})\tau}{n_h^k(\tilde{s},\tilde{a}) \vee 1}} q_h^{\pi^k}(\tilde{s}, \tilde{a} | s'; h+1)}_{(i)} \\
&\quad + \underbrace{\sum_{k=1}^K \sum_{1 \leq h < \bar{h} \leq H} \sum_{s,a,s'} \sum_{\tilde{s}, \tilde{a}, \tilde{s}'} q_h^{\pi^k}(s, a) p_h^k(s'|s, a) \cdot \frac{\tau}{n_h^k(\tilde{s}, \tilde{a}) \vee 1} q_h^{\pi^k}(\tilde{s}, \tilde{a} | s'; h+1)}_{(ii)} \\
&\quad + \underbrace{\sum_{k=1}^K \sum_{1 \leq h < \bar{h} \leq H} \sum_{s,a,s'} \sum_{\tilde{s}, \tilde{a}} \frac{\tau}{n_h^k(s,a) \vee 1} q_h^{\pi^k}(s, a) \cdot q_h^{\pi^k}(\tilde{s}, \tilde{a} | s'; h+1)}_{(iii)},
\end{aligned}$$

where the last inequality follows because $\sqrt{xy} \leq x + y$ for every $x, y \geq 0$. Term (i) is bounded as follows:

$$\begin{aligned}
(i) &= \tau \sum_{1 \leq h < \tilde{h} \leq H} \sum_{k=1}^K \sum_{s,a,s'} \sum_{\tilde{s},\tilde{a},\tilde{s}'} \sqrt{\frac{q_h^{\pi^k}(s,a)q_h^{\pi^k}(\tilde{s},\tilde{a} | s'; h+1)p_h^k(\tilde{s}'|\tilde{s},\tilde{a})}{n_h^k(s,a) \vee 1}} \cdot \sqrt{\frac{q_h^{\pi^k}(s,a)q_h^{\pi^k}(\tilde{s},\tilde{a} | s'; h+1)p_h^k(s'|s,a)}{n_h^k(\tilde{s},\tilde{a}) \vee 1}} \\
&\leq \tau \sum_{1 \leq h < \tilde{h} \leq H} \sqrt{\sum_{k,s,a,s',\tilde{s},\tilde{a},\tilde{s}'} \frac{q_h^{\pi^k}(s,a)q_h^{\pi^k}(\tilde{s},\tilde{a} | s'; h+1)p_h^k(\tilde{s}'|\tilde{s},\tilde{a})}{n_h^k(s,a) \vee 1}} \cdot \sqrt{\sum_{k,s,a,s',\tilde{s},\tilde{a},\tilde{s}'} \frac{q_h^{\pi^k}(s,a)q_h^{\pi^k}(\tilde{s},\tilde{a} | s'; h+1)p_h^k(s'|s,a)}{n_h^k(\tilde{s},\tilde{a}) \vee 1}} \\
&= \tau \sum_{1 \leq h < \tilde{h} \leq H} \sqrt{S \sum_{k=1}^K \sum_{s,a} \frac{q_h^{\pi^k}(s,a)}{n_h^k(s,a) \vee 1}} \cdot \sqrt{\sum_{k=1}^K \sum_{s',\tilde{s},\tilde{a},\tilde{s}'} \frac{q_{h+1}^{\pi^k}(s')q_h^{\pi^k}(\tilde{s},\tilde{a} | s'; h+1)}{n_h^k(\tilde{s},\tilde{a}) \vee 1}} \\
&\leq HS\tau \sqrt{\sum_{k=1}^K \sum_{s,a,h} \frac{q_h^{\pi^k}(s,a)}{n_h^k(s,a) \vee 1}} \cdot \sqrt{\sum_{k=1}^K \sum_{\tilde{s},\tilde{a},\tilde{h}} \frac{q_h^{\pi^k}(\tilde{s},\tilde{a})}{n_h^k(\tilde{s},\tilde{a}) \vee 1}} \leq HS\tau \sum_{k=1}^K \sum_{s,a,h} \frac{q_h^{\pi^k}(s,a)}{n_h^k(s,a) \vee 1} \lesssim H^2 S^2 A \tau^2,
\end{aligned}$$

where the last inequality is by event E^{on3} and Eq. (18). Term (ii) is bounded as follows:

$$\begin{aligned}
(ii) &= S \sum_{k=1}^K \sum_{1 \leq h < \tilde{h} \leq H} \sum_{\tilde{s},\tilde{a},\tilde{s}'} \frac{\tau}{n_h^k(\tilde{s},\tilde{a}) \vee 1} q_h^{\pi^k}(\tilde{s},\tilde{a} | s'; h+1) \sum_{s,a} q_h^{\pi^k}(s,a) p_h^k(s'|s,a) \\
&= S \sum_{k=1}^K \sum_{1 \leq h < \tilde{h} \leq H} \sum_{\tilde{s},\tilde{a}} \frac{\tau}{n_h^k(\tilde{s},\tilde{a}) \vee 1} \sum_{s'} q_h^{\pi^k}(\tilde{s},\tilde{a} | s'; h+1) q_{h+1}^{\pi^k}(s') \\
&= HS\tau \sum_{k=1}^K \sum_{\tilde{h}} \sum_{\tilde{s},\tilde{a}} \frac{q_h^{\pi^k}(\tilde{s},\tilde{a})}{n_h^k(\tilde{s},\tilde{a}) \vee 1} \lesssim H^2 S^2 A \tau^2.
\end{aligned}$$

Term (iii) is bounded as follows:

$$\begin{aligned}
(iii) &\leq \sum_{k=1}^K \sum_{1 \leq h < \tilde{h} \leq H} \sum_{s,a,s'} \frac{\tau}{n_h^k(s,a) \vee 1} q_h^{\pi^k}(s,a) \sum_{\tilde{s},\tilde{a}} q_h^{\pi^k}(\tilde{s},\tilde{a} | s'; h+1) \\
&\leq HS\tau^2 \sum_{k=1}^K \sum_{s,a,h} \frac{q_h^{\pi^k}(s,a)}{n_h^k(s,a) \vee 1} \lesssim H^2 S^2 A \tau^2.
\end{aligned}$$

■

Algorithm 8 COOPERATIVE O-REPS WITH NON-FRESH RANDOMNESS (COOP-NF-O-REPS)

- 1: **input:** state space \mathcal{S} , action space \mathcal{A} , horizon H , transition function p , number of episodes K , number of agents m , exploration parameter γ , learning rate η , confidence parameter δ .
 - 2: **initialize:** $\pi_h^1(a | s) = 1/A, q_h^1(s, a) = q_h^{\pi^1}(s, a) \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 3: **for** $k = 1, \dots, K$ **do**
 - 4: **for** $v = 1, \dots, m$ **do**
 - 5: observe initial state $s_1^{k,v}$.
 - 6: **for** $h = 1, \dots, H$ **do**
 - 7: pick action $a_h^{k,v} \sim \pi_h^k(\cdot | s_h^{k,v})$, suffer cost $c_h^k(s_h^{k,v}, a_h^{k,v})$ and observe next state $s_{h+1}^{k,v}$.
 - 8: **end for**
 - 9: **end for**
 - 10: For every (s, a, h) compute $\widetilde{W}_h^k(s, a)$ – the estimate of $W_h^k(s, a) = \Pr[\exists v : s_h^{k,v} = s, a_h^{k,v} = a | \pi^k]$ using $N = 10\gamma^{-2} \log \frac{KHSAm}{\delta}$ samples (Algorithm 9).
 - 11: compute $\hat{c}_h^k(s, a) = \frac{c_h^k(s, a) \mathbb{I}\{\exists v: s_h^{k,v} = s, a_h^{k,v} = a\}}{\widetilde{W}_h^k(s, a) + \gamma} \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 12: compute $q^{k+1} = \arg \min_{q \in \Delta(\mathcal{M})} \eta \langle q, \hat{c}^k \rangle + \text{KL}(q \| q^k)$.
 - 13: compute $\pi_h^{k+1}(a | s) = \frac{q_h^k(s, a)}{\sum_{a' \in \mathcal{A}} q_h^k(s, a')} \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 14: **end for**
-

Appendix F. The `coop-nf-o-reps` algorithm for adversarial MDPs with non-fresh randomness and known p

For the setting of adversarial MDPs with non-fresh randomness and known transitions we propose the Cooperative O-REPS with non-fresh randomness algorithm (`coop-nf-o-reps`; see Algorithm 8). The idea is similar to the `coop-o-reps` algorithm for fresh randomness, but the key difference is that the probability to reach some state-action pair that the algorithm uses (i.e., $W_h^k(s, a)$) must be computed differently in order to suit non-fresh randomness. In fact, computing $W_h^k(s, a)$ becomes a difficult challenge once the randomness is non-fresh, and a naive computation takes exponential time. Instead we propose to estimate $W_h^k(s, a)$ from samples. That is, we simulate $N = 10\gamma^{-2} \log \frac{KHSAm}{\delta}$ i.i.d episodes in which all agents use policy π^k and then estimate $W_h^k(s, a)$ by the fraction of episodes in which (s, a) was reached in step h by at least one of the m agents. This way our algorithm keeps polynomial running time.

Theorem 47 *With probability $1 - \delta$, setting $\eta = \gamma = 1/\sqrt{(1 + \frac{SA}{m})K}$, the individual regret of each agent of `coop-nf-o-reps` is*

$$R_K = O \left(H \sqrt{SK \log \frac{HSA}{\delta}} + \sqrt{\frac{H^2SAK}{m} \log \frac{HSA}{\delta}} + \frac{HSA}{m} \log \frac{HSA}{\delta} + HS \log \frac{HSA}{\delta} \right).$$

Algorithm 9 ESTIMATE REACHABILITY PROBABILITY FOR NON-FRESH RANDOMNESS

```
1: input: state space  $\mathcal{S}$ , action space  $\mathcal{A}$ , transition function  $p$ , number of agents  $m$ , policy  $\pi$ , number of samples  $N$ ,  
   state-action-step triplet to estimate  $(\bar{s}, \bar{a}, \bar{h})$ .  
2: initialize indicator for reaching  $I(n) \leftarrow 0$  for  $n \in [N]$ .  
3: for  $n = 1, \dots, N$  do  
4:   initialize realized transitions  $p_h^r(s' | s, a) = 0 \forall (s, a, s', h)$ .  
5:   for  $h = 1, \dots, \bar{h}$  do  
6:     for  $(s, a) \in \mathcal{S} \times \mathcal{A}$  do  
7:       sample  $s' \sim p_h(\cdot | s, a)$  and set  $p_h^r(s' | s, a) = 1$ .  
8:     end for  
9:   end for  
10:  for  $v = 1, \dots, m$  do  
11:    observe initial state  $s_1^v = s_{\text{init}}$ .  
12:    for  $h = 1, \dots, \bar{h}$  do  
13:      pick action  $a_h^v \sim \pi_h(\cdot | s_h^v)$  and observe next state  $s_{h+1}^v \sim p_h^r(\cdot | s, a)$ .  
14:    end for  
15:    if  $s_h^v = \bar{s}, a_h^v = \bar{a}$  then  
16:      set  $I(n) \leftarrow 1$ .  
17:    break  
18:    end if  
19:  end for  
20: end for  
21: return  $\frac{1}{N} \sum_{n=1}^N I(n)$ .
```

F.1 The good event

Define the following events:

$$\begin{aligned} E^{app} &= \left\{ \forall (s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times [H] \times [K] : |\widetilde{W}_h^k(s, a) - W_h^k(s, a)| \leq \gamma/2 \right\} \\ E^c &= \left\{ \sum_{k=1}^K \langle \mathbb{E}^k[\hat{c}^k] - \hat{c}^k, q^k \rangle \leq 4H \sqrt{K \log \frac{6}{\delta}} \right\} \\ E^{\hat{c}} &= \left\{ \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left(\frac{1}{m} + \pi_h^k(a | s) \right) (\hat{c}_h^k(s, a) - 2c_h^k(s, a)) \leq \frac{10HSA \log \frac{6HSA}{\delta}}{m\gamma} + \frac{10HS \log \frac{6HSA}{\delta}}{\gamma} \right\} \\ E^* &= \left\{ \sum_{k=1}^K \langle \hat{c}^k - c^k, q^{\pi^*} \rangle \leq \frac{2H \log \frac{6HSA}{\delta}}{\gamma} \right\} \end{aligned}$$

The good event is the intersection of the above events. The following lemma establishes that the good event holds with high probability.

Lemma 48 (The Good Event) *Let $\mathbb{G} = E^{app} \cap E^c \cap E^{\hat{c}} \cap E^*$ be the good event. It holds that $\Pr[\mathbb{G}] \geq 1 - \delta$.*

Proof By Hoeffding inequality we have that $\Pr[\neg E^{app}] \leq \delta/6$, and the other events are similar to Lemma 40. ■

F.2 Proof of Theorem 47

Proof of Theorem 47 By Lemma 48, the good event holds with probability $1 - \delta$. We now analyze the regret under the assumption that the good event holds. We start by decomposing the regret as follows:

$$\begin{aligned} R_K &= \sum_{k=1}^K V_1^{k, \pi^k}(s_1^{k,v}) - V_1^{k, \pi^*}(s_1^{k,v}) = \sum_{k=1}^K \langle c^k, q^k - q^{\pi^*} \rangle \\ &= \underbrace{\sum_{k=1}^K \langle c^k - \hat{c}^k, q^k \rangle}_{(A)} + \underbrace{\sum_{k=1}^K \langle \hat{c}^k, q^k - q^{\pi^*} \rangle}_{(B)} + \underbrace{\sum_{k=1}^K \langle \hat{c}^k - c^k, q^{\pi^*} \rangle}_{(C)}. \end{aligned}$$

Term (A) can be further decomposed as:

$$(A) = \sum_{k=1}^K \langle c^k - \hat{c}^k, q^k \rangle = \sum_{k=1}^K \langle c^k - \mathbb{E}^k[\hat{c}^k], q^k \rangle + \sum_{k=1}^K \langle \mathbb{E}^k[\hat{c}^k] - \hat{c}^k, q^k \rangle.$$

The second term is bounded by $4H\sqrt{K \log \frac{6}{\delta}}$ by the good event E^c , and for the first term:

$$\begin{aligned} \sum_{k=1}^K \langle c^k - \mathbb{E}^k[\hat{c}^k], q^k \rangle &= \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) c_h^k(s, a) \left(1 - \frac{\mathbb{E}^k[\mathbb{I}\{\exists v : s_h^{k,v} = s, a_h^{k,v} = a\}]}{\widetilde{W}_h^k(s, a) + \gamma} \right) \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) c_h^k(s, a) \left(1 - \frac{\mathbb{E}^k[\mathbb{I}\{\exists v : s_h^{k,v} = s, a_h^{k,v} = a\}]}{W_h^k(s, a) + \gamma/2} \right) \\ &= \sum_{k,h,s,a} q_h^k(s, a) c_h^k(s, a) \left(1 - \frac{W_h^k(s, a)}{W_h^k(s, a) + \gamma/2} \right) \leq \gamma \sum_{k,h,s,a} \frac{q_h^k(s, a)}{W_h^k(s, a) + \gamma/2} \\ &\leq \gamma \sum_{k,h,s,a} \frac{q_h^k(s) \pi_h^k(a | s)}{W_h^k(s, a)} \leq \gamma \sum_{k,h,s,a} \sum_{a \in \mathcal{A}} \left(\frac{1}{m} + \pi_h^k(a | s) \right) = \frac{\gamma HSAK}{m} + \gamma HSK, \end{aligned}$$

where the first inequality is by the event E^{app} , and the last inequality is by Lemma 49.

Term (B) is bounded by OMD (see, e.g., [Zimin and Neu \(2013\)](#)) as follows:

$$\begin{aligned}
(B) &= \sum_{k=1}^K \langle \hat{c}^k, q^k - q^{\pi^*} \rangle \leq \frac{H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) \hat{c}_h^k(s, a)^2 \\
&\leq \frac{H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) \frac{\hat{c}_h^k(s, a)}{\widetilde{W}_h^k(s, a) + \gamma} \\
&\leq \frac{H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) \frac{\hat{c}_h^k(s, a)}{W_h^k(s, a) + \gamma/2} \\
&\leq \frac{H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \frac{q_h^k(s) \pi_h^k(a | s)}{W_h^k(s, a)} \hat{c}_h^k(s, a) \\
&\leq \frac{H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left(\frac{1}{m} + \pi_h^k(a | s) \right) \hat{c}_h^k(s, a) \\
&\leq \frac{H \log(HSA)}{\eta} + 2\eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left(\frac{1}{m} + \pi_h^k(a | s) \right) c_h^k(s, a) + \frac{10\eta HSA \log \frac{6HSA}{\delta}}{m\gamma} + \frac{10\eta HS \log \frac{6HSA}{\delta}}{\gamma} \\
&\lesssim \frac{H \log \frac{HSA}{\delta}}{\eta} + \frac{\eta HSAK}{m} + \eta HSK + \frac{\eta HSA \log \frac{6HSA}{\delta}}{m\gamma} + \frac{\eta HS \log \frac{6HSA}{\delta}}{\gamma},
\end{aligned}$$

where the forth inequality is by Lemma 49, and the fifth inequality is by the good event $E^{\hat{c}}$.

Term (C) is bounded by $\frac{2H \log \frac{6H}{\delta}}{\gamma}$ by the good event E^* . Putting the three terms together gives the final regret bound when setting $\eta = \gamma = \sqrt{\frac{\log \frac{HSA}{\delta}}{(1+\frac{A}{m})SK}}$. ■

F.3 Auxiliary lemmas

Lemma 49 *Let π be a policy and denote by $q_h^\pi(s)$ the probability to reach state s in time h when playing policy π . Assume that m agents use the same policy π in an MDP \mathcal{M} with non-fresh randomness, and denote by $W_h(s, a)$ the probability that at least one agent reaches (s, a) in time h . Then, for every $(s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$, it holds that:*

$$\frac{q_h^\pi(s) \pi_h(a | s)}{W_h(s, a)} \leq \frac{1}{m} + \pi_h(a | s).$$

Proof Let $M_h(s)$ be the number of agents that arrive at state s in time h . We have that,

$$\begin{aligned}
W_h(s, a) &= \Pr[\exists v \in [m] : s_h^v = s, a_h^v = a | \pi] = \mathbb{E} \left[1 - (1 - \pi_h(a | s))^{M_h(s)} | \pi \right] \\
&\geq \mathbb{E} \left[\frac{\pi_h(a | s)}{\frac{1}{M_h(s)} + \pi_h(a | s)} | \pi \right] = \mathbb{E} \left[\frac{M_h(s) \pi_h(a | s)}{1 + M_h(s) \pi_h(a | s)} | \pi \right], \tag{21}
\end{aligned}$$

where the inequality is by Lemma 41.

Notice that $\mathbb{E}[M_h(s) \mid \pi] = mq_h^\pi(s)$ by linearity of expectation. Therefore, Eq. (21) is bounded from below by the value of the following optimization problem:

$$\begin{aligned} \min_{p_0, \dots, p_m} \quad & \sum_{i=0}^m p_i \frac{i\pi_h(a \mid s)}{1 + i\pi_h(a \mid s)}, \\ \text{s.t.} \quad & \sum_{i=0}^m p_i i = mq_h^\pi(s), \\ & \sum_{i=0}^m p_i = 1, \\ & p_i \geq 0 \quad \forall i \in [m], \end{aligned}$$

where p_i represents $\Pr[M_h(s) = i]$. Since the coefficient of p_i in the constrains and the objective are non-negative, we can substitute the equality constrains with “ \geq ” constrains. We get the following standard form Linear Programming:

$$\begin{aligned} \min_{p \in \mathbb{R}^{m+1}} \quad & b^T p, \\ \text{s.t.} \quad & A^T p \geq c, \\ & p \geq 0, \end{aligned}$$

where,

$$b = \begin{pmatrix} 0 \\ \frac{\pi_h(a \mid s)}{1 + \pi_h(a \mid s)} \\ \vdots \\ \frac{m\pi_h(a \mid s)}{1 + m\pi_h(a \mid s)} \end{pmatrix}, \quad A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \\ \vdots & \vdots \\ 1 & m \end{pmatrix}, \quad c = (1, mq_h^\pi(s)).$$

The dual problem is,

$$\begin{aligned} \max_{x_1, x_2} \quad & (x_1 + x_2)mq_h^\pi(s) \\ \text{s.t.} \quad & x_1 \leq 0, \\ & x_1 + x_2 \leq \frac{\pi_h(a \mid s)}{1 + \pi_h(a \mid s)}, \\ & x_1 + 2x_2 \leq \frac{2\pi_h(a \mid s)}{1 + 2\pi_h(a \mid s)}, \\ & \vdots \\ & x_1 + mx_2 \leq \frac{m\pi_h(a \mid s)}{1 + m\pi_h(a \mid s)}, \\ & x_1, x_2 \geq 0. \end{aligned}$$

From the first and the last constrains we have $x_1 = 0$ and the rest of the constrains are equivalent to $x_2 \leq \frac{\pi_h(a \mid s)}{1 + m\pi_h(a \mid s)}$. Hence the maximum value is $\frac{m\pi_h(a \mid s)}{1 + m\pi_h(a \mid s)} q_h^\pi(s)$, which completes the proof. \blacksquare

Algorithm 10 COOPERATIVE UOB-REPS WITH NON-FRESH RANDOMNESS (COOP-NF-UOB-REPS)

- 1: **input:** state space \mathcal{S} , action space \mathcal{A} , horizon H , number of episodes K , number of agents $m = \sqrt{K}$, exploration parameter γ , learning rate η , confidence parameter δ .
 - 2: **initialize:** $n_h^1(s, a) = 0, n_h^1(s, a, s') = 0, \pi_h^1(a | s) = 1/A, q_h^1(s, a, s') = 1/S^2A \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 3: **initialize:** define a mapping $\sigma : [H] \times \mathcal{A} \times [K] \rightarrow [m]$ such that $\sigma(h, a, k) \neq \sigma(h', a', k)$ whenever $h \neq h'$ or $a \neq a'$, and such that each agent is assigned by σ exactly $HA\sqrt{K}$ times (i.e., $|\sigma^{-1}(v)| = HA\sqrt{K}$ for every $v \in [m]$).
 - 4: **for** $k = 1, \dots, K$ **do**
 - 5: set $I_h^k(s, a, s') = 0, I_h^k(s, a) = 0 \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 6: **for** $v = 1, \dots, m$ **do**
 - 7: observe initial state $s_1^{k,v}$.
 - 8: **for** $h = 1, \dots, H$ **do**
 - 9: **if** $\exists \tilde{a} \in \mathcal{A} : \sigma(h, \tilde{a}, k) = v$ **then**
 - 10: pick action $a_h^{k,v} = \tilde{a}$.
 - 11: **else**
 - 12: pick action $a_h^{k,v} \sim \pi_h^k(\cdot | s_h^{k,v})$.
 - 13: **end if**
 - 14: suffer cost $c_h^k(s_h^{k,v}, a_h^{k,v})$ and observe next state $s_{h+1}^{k,v}$.
 - 15: update $I_h^k(s_h^{k,v}, a_h^{k,v}) \leftarrow 1, I_h^k(s_h^{k,v}, a_h^{k,v}, s_{h+1}^{k,v}) \leftarrow 1$.
 - 16: **end for**
 - 17: set $n_h^{k+1}(s, a) \leftarrow n_h^k(s, a) + I_h^k(s, a), n_h^{k+1}(s, a, s') \leftarrow n_h^k(s, a, s') + I_h^k(s, a, s') \forall (s, a, s', h)$.
 - 18: set $\hat{p}_h^{k+1}(s' | s, a) \leftarrow \frac{n_h^{k+1}(s, a, s')}{n_h^{k+1}(s, a) \vee 1} \forall (s, a, s', h) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times [H]$.
 - 19: compute confidence set for $\epsilon_h^{k+1}(s' | s, a) = 4\sqrt{\frac{\hat{p}_h^{k+1}(s' | s, a) \ln \frac{HSAK}{4\delta}}{n_h^{k+1}(s, a) \vee 1}} + 10\frac{\ln \frac{HSAK}{4\delta}}{n_h^{k+1}(s, a) \vee 1}$:
$$\mathcal{P}^{k+1} = \{p' | \forall (s, a, s', h) : |\hat{p}_h^{k+1}(s' | s, a) - p'(s' | s, a)| \leq \epsilon_h^{k+1}(s' | s, a)\}.$$
 - 20: compute $u_h^k(s) = \max_{p' \in \mathcal{P}^k} q_h^{p', \pi^k}(s) = \max_{p' \in \mathcal{P}^k} \Pr[s_h = s | \pi^k, p'] \forall s \in \mathcal{S}$.
 - 21: compute $\hat{c}_h^k(s, a) = \frac{c_h^k(s, a) \mathbb{1}_{\{s_h^{k, \sigma(h, a, k)} = s\}}}{u_h^k(s) + \gamma} \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$.
 - 22: compute $q^{k+1} = \arg \min_{q \in \Delta(\mathcal{M}, k+1)} \eta(q, \hat{c}^k) + \text{KL}(q \| q^k)$.
 - 23: compute $\pi_h^{k+1}(a | s) = \frac{q_h^{k+1}(s, a)}{\sum_{a' \in \mathcal{A}} q_h^{k+1}(s, a')} \forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$, where $q_h^{k+1}(s, a) = \sum_{s' \in \mathcal{S}} q_h^{k+1}(s, a, s')$.
 - 24: **end for**
 - 25: **end for**
-

Appendix G. The coop-nf-UOB-REPS algorithm for adversarial MDPs with non-fresh randomness and unknown p

For the setting of adversarial MDPs with non-fresh randomness and unknown transitions we propose the Cooperative UOB-REPS with non-fresh randomness algorithm (coop-nf-UOB-REPS; see Algorithm 10). The idea is to combine the coop-nf-O-REPS algorithm for known transitions with ideas from the coop-ULCAE algorithm in order to handle unknown transitions under non-fresh randomness. The main challenge is that, unlike the stochastic case, we cannot eliminate sub-optimal actions. Thus, our method requires \sqrt{K} agents to attain near-optimal regret as opposed to the stochastic case where only H^2A^2 agents are required.

Theorem 50 Assume that coop-nf-UOB-REPS is run with $m = \sqrt{K}$ agents. With probability $1 - \delta$, setting $\eta = \gamma = \sqrt{\frac{\log \frac{KHSA}{\delta}}{SK}}$, the individual regret of each agent of coop-nf-O-REPS is

$$R_K = O \left(H^2 S \sqrt{K \log \frac{KHSA}{\delta}} + H^3 S^3 \log^2 \frac{KHSA}{\delta} \right).$$

G.1 The good event

Denote $\epsilon_h^k(s' | s, a) = \sqrt{\frac{2\hat{p}_h^k(s'|s,a) \log \frac{30KHSA}{\delta}}{n_h^k(s,a) \vee 1}} + \frac{2 \log \frac{30KHSA}{\delta}}{n_h^k(s,a) \vee 1}$ and $\epsilon_h^k(s, a) = \sum_{s' \in \mathcal{S}} \epsilon_h^k(s' | s, a)$. Define the following events:

$$\begin{aligned} E^p &= \left\{ \forall (k, s, a, s', h) : |p_h(s'|s, a) - \hat{p}_h^k(s'|s, a)| \leq \epsilon_h^k(s' | s, a) \right\} \\ E^{on} &= \left\{ \forall (k, h, s, a, v) \in [K] \times [H] \times \mathcal{S} \times \mathcal{A} : n_h^k(s, a) \geq \frac{1}{2} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \log \frac{6mHSA}{\delta} \right\} \\ E^c &= \left\{ \sum_{k=1}^K \langle \mathbb{E}^k[\hat{c}^k | \pi^k] - \hat{c}^k, q^k \rangle \leq 4HS \sqrt{K \log \frac{6}{\delta}} \right\} \\ E^{\hat{c}} &= \left\{ \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \pi_h^k(a | s) (\hat{c}_h^k(s, a) - 2c_h^k(s, a)) \leq \frac{10HS \log \frac{3HSA}{\delta}}{\gamma} \right\} \\ E^* &= \left\{ \sum_{k=1}^K \langle \hat{c}^k - c^k, q^{\pi^*} \rangle \leq \frac{H \log \frac{3HSA}{\delta}}{\gamma} \right\} \end{aligned}$$

The good event is the intersection of the above events. The following lemma establishes that the good event holds with high probability.

Lemma 51 (The Good Event) *Let $\mathbb{G} = E^p \cap E^{on} \cap E^c \cap E^{\hat{c}} \cap E^*$ be the good event. It holds that $\Pr[\mathbb{G}] \geq 1 - \delta$.*

Proof Similar to the proofs of Lemmas 24 and 48 and to proofs in Jin et al. (2020a). ■

G.2 Proof of Theorem 50

Proof of Theorem 50 By Lemma 51, the good event holds with probability $1 - \delta$. We now analyze the regret under the assumption that the good event holds. Note that each agent plays the OMD policy π^k in all except for $HA\sqrt{K}$ episodes. Thus, the regret is bounded by the regret of the policies $\{\pi^k\}_{k=1}^K$ plus a $H^2A\sqrt{K}$ term which is at most $H^2S\sqrt{K}$. Next, we focus on bounding the regret of $\{\pi^k\}_{k=1}^K$, starting with the following decomposition:

$$\begin{aligned} \sum_{k=1}^K V_1^{k, \pi^k}(s_1^{k,v}) - V_1^{k, \pi^*}(s_1^{k,v}) &= \sum_{k=1}^K \langle c^k, q^{\pi^k} - q^{\pi^*} \rangle \\ &= \underbrace{\sum_{k=1}^K \langle c^k, q^{\pi^k} - q^k \rangle}_{(A)} + \underbrace{\sum_{k=1}^K \langle c^k - \hat{c}^k, q^k \rangle}_{(B)} + \underbrace{\sum_{k=1}^K \langle \hat{c}^k, q^k - q^{\pi^*} \rangle}_{(C)} + \underbrace{\sum_{k=1}^K \langle \hat{c}^k - c^k, q^{\pi^*} \rangle}_{(D)}. \end{aligned}$$

Let $\tau = \log \frac{KHSAm}{\delta}$ be a logarithmic term. Term (A) can be decomposed using the value difference lemma (see, e.g., Shani et al. (2020)):

$$\begin{aligned} (A) &= \sum_{k=1}^K \langle c^k, q^{\pi^k} - q^k \rangle \leq 2H \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^{\pi^k}(s, a) \|p_h(\cdot | s, a) - \hat{p}_h^k(\cdot | s, a)\|_1 \\ &\leq H\sqrt{S\tau} \sum_{k,h,s,a} \frac{q_h^{\pi^k}(s, a)}{\sqrt{n_h^k(s, a) \vee 1}} + HS\tau \sum_{k,h,s,a} \frac{q_h^{\pi^k}(s, a)}{n_h^k(s, a) \vee 1}, \end{aligned}$$

where the second inequality is by event E^p . We now bound each of the two sums separately using the event E^{on} . For the second sum we have:

$$\begin{aligned} \sum_{k,h,s,a} \frac{q_h^{\pi^k}(s,a)}{n_h^k(s,a) \vee 1} &\leq \sum_{k,h,s,a} \frac{q_h^{\pi^k}(s,a)}{(\frac{1}{2} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \log \frac{6mHSA}{\delta}) \vee 1} = \sum_{k,h,s} \frac{q_h^{\pi^k}(s) \sum_a \pi_h^k(a|s)}{(\frac{1}{2} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \log \frac{6mHSA}{\delta}) \vee 1} \\ &\leq 2HS\tau + 2 \sum_{h,s} \sum_{k: \sum_{j=1}^{k-1} q_h^{\pi^j}(s) \geq 2 \log \frac{6mHSA}{\delta}} \frac{q_h^{\pi^k}(s)}{\sum_{j=1}^{k-1} q_h^{\pi^j}(s)} \lesssim HS\tau, \end{aligned}$$

where the last inequality is by [Rosenberg et al. \(2020, Lemma B.18\)](#). For the first term:

$$\begin{aligned} \sum_{k,h,s,a} \frac{q_h^{\pi^k}(s,a)}{\sqrt{n_h^k(s,a) \vee 1}} &\leq \sum_{k,h,s,a} \frac{q_h^{\pi^k}(s,a)}{\sqrt{(\frac{1}{2} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \log \frac{6mHSA}{\delta}) \vee 1}} = \sum_{k,h,s} \frac{q_h^{\pi^k}(s) \sum_a \pi_h^k(a|s)}{\sqrt{(\frac{1}{2} \sum_{j=1}^{k-1} q_h^{\pi^j}(s) - \log \frac{6mHSA}{\delta}) \vee 1}} \\ &\leq 2HS\tau + 2 \sum_{h,s} \sum_{k: \sum_{j=1}^{k-1} q_h^{\pi^j}(s) \geq 2 \log \frac{6mHSA}{\delta}} \frac{q_h^{\pi^k}(s)}{\sqrt{\sum_{j=1}^{k-1} q_h^{\pi^j}(s)}} \\ &= 2HS\tau + 2 \sum_{h,s} \sum_{k: \sum_{j=1}^{k-1} q_h^{\pi^j}(s) \geq 2 \log \frac{6mHSA}{\delta}} \frac{q_h^{\pi^k}(s)}{\sqrt{\sum_{j=1}^k q_h^{\pi^j}(s)}} \sqrt{\frac{\sum_{j=1}^k q_h^{\pi^j}(s)}{\sum_{j=1}^{k-1} q_h^{\pi^j}(s)}} \\ &= 2HS\tau + 4 \sum_{h,s} \sum_{k: \sum_{j=1}^{k-1} q_h^{\pi^j}(s) \geq 2 \log \frac{6mHSA}{\delta}} \frac{q_h^{\pi^k}(s)}{\sqrt{\sum_{j=1}^k q_h^{\pi^j}(s)}} \\ &\leq 2HS\tau + 8 \sum_{h,s} \sqrt{\sum_{k=1}^K q_h^{\pi^k}(s)} \leq 2HS\tau + 8 \sqrt{HS \sum_{k=1}^K \sum_{h,s} q_h^{\pi^k}(s)} = 2HS\tau + 8H\sqrt{SK}, \end{aligned}$$

where the third inequality is by [Streeter and McMahan \(2010, Lemma 1\)](#), and the last inequality is by Jensen's inequality. Putting these together we get that: (A) $\lesssim H^2 S \sqrt{K} \tau + H^2 S^2 \tau^2$.

Term (B) can be further decomposed as:

$$(B) = \sum_{k=1}^K \langle c^k - \hat{c}^k, q^k \rangle = \sum_{k=1}^K \langle c^k - \mathbb{E}^k[\hat{c}^k | \pi^k], q^k \rangle + \sum_{k=1}^K \langle \mathbb{E}^k[\hat{c}^k | \pi^k] - \hat{c}^k, q^k \rangle.$$

The second term is bounded by $4HS\sqrt{K \log \frac{6}{\delta}}$ by the good event E^c , and for the first term:

$$\begin{aligned} \sum_{k=1}^K \langle c^k - \mathbb{E}^k[\hat{c}^k | \pi^k], q^k \rangle &= \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s,a) c_h^k(s,a) \left(1 - \frac{\mathbb{E}^k[\mathbb{I}\{s_h^{k,\sigma(h,a,k)} = s\} | \pi^k]}{u_h^k(s) + \gamma} \right) \\ &= \sum_{k,h,s,a} q_h^k(s,a) c_h^k(s,a) \left(1 - \frac{q_h^{\pi^k}(s)}{u_h^k(s) + \gamma} \right) \leq \sum_{k,h,s} q_h^k(s) \left(1 - \frac{q_h^{\pi^k}(s)}{u_h^k(s) + \gamma} \right) \\ &\leq 2 \sum_{k,h,s} \frac{q_h^k(s)}{u_h^k(s)} (u_h^k(s) - q_h^{\pi^k}(s) + \gamma) \leq 2 \sum_{k,h,s} (u_h^k(s) - q_h^{\pi^k}(s)) + \gamma HSK, \end{aligned}$$

where the second equality is because agent $\sigma(h, a, k)$ plays policy π^k until step h . Finally, $\sum_{k,h,s} (u_h^k(s) - q_h^k(s))$ is bounded by similarly to Lemma 46.

Term (C) is bounded by OMD (see, e.g., [Rosenberg and Mansour \(2019a\)](#)) as follows:

$$\begin{aligned}
(B) &= \sum_{k=1}^K \langle \hat{c}^k, q^k - q^{\pi^*} \rangle \leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s, a) \hat{c}_h^k(s, a)^2 \\
&\leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_h^k(s) \pi_h^k(a | s) \frac{\hat{c}_h^k(s, a)}{u_h^k(s) + \gamma} \\
&\leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \pi_h^k(a | s) \hat{c}_h^k(s, a) \\
&\leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \pi_h^k(a | s) c_h^k(s, a) + \frac{\eta HS \log \frac{3HSA}{\delta}}{\gamma} \\
&\leq \frac{2H \log(HSA)}{\eta} + \eta \sum_{k=1}^K \sum_{h=1}^H \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \pi_h^k(a | s) + \frac{\eta HS \log \frac{3HSA}{\delta}}{\gamma} \\
&= \frac{2H \log(HSA)}{\eta} + \eta HSK + \frac{\eta HS \log \frac{3HSA}{\delta}}{\gamma},
\end{aligned}$$

where the forth inequality is by the good event $E^{\hat{c}}$.

Term (D) is bounded by $\frac{H \log \frac{3H}{\delta}}{\gamma}$ by the good event E^* . Putting the three terms together gives the final regret bound when setting $\eta = \gamma = \sqrt{\frac{\log \frac{KHS A}{\delta}}{SK}}$. ■