Provably Efficient Multi-Task Reinforcement Learning with Model Transfer

Chicheng Zhang University of Arizona chichengz@cs.arizona.edu Zhi Wang University of California San Diego zhiwang@eng.ucsd.edu

Abstract

We study multi-task reinforcement learning (RL) in tabular episodic Markov decision processes (MDPs). We formulate a heterogeneous multi-player RL problem, in which a group of players concurrently face similar but not necessarily identical MDPs, with a goal of improving their collective performance through inter-player information sharing. We design and analyze an algorithm based on the idea of model transfer, and provide gap-dependent and gap-independent upper and lower bounds that characterize the intrinsic complexity of the problem.

1 Introduction

In many real-world applications, reinforcement learning (RL) agents can be deployed as a group to complete similar tasks at the same time. For example, in healthcare robotics, robots are paired with people with dementia to perform personalized cognitive training activities by learning their preferences [42, 21]; in autonomous driving, a set of autonomous vehicles learn how to navigate and avoid obstacles in various environments [27]. In these settings, each learning agent alone may only be able to acquire a limited amount of data, while the agents as a group have the potential to collectively learn faster through sharing knowledge among themselves. Multi-task learning [7] is a practical framework that can be used to model such settings, where a set of learning agents share/transfer knowledge to improve their collective performance.

Despite many empirical successes of multi-task RL (see, e.g., [51, 28, 27]) and transfer learning for RL (see, e.g., [26, 39]), a theoretical understanding of when and how information sharing or knowledge transfer can provide benefits remains limited. Exceptions include [16, 6, [11, [17, 32, 25]], which study multi-task learning from parameter- or representation-transfer perspectives. However, these works still do not provide a completely satisfying answer: for example, in many application scenarios, the reward structures and the environment dynamics are only slightly different for each task—this is, however, not captured by representation transfer [11, [17]] or existing works on clustering-based parameter transfer [16, [6]]. In such settings, is it possible to design provably efficient multi-task RL algorithms that have guarantees never worse than agents learning individually, while outperforming the individual agents in favorable situations?

In this work, we formulate an online multi-task RL problem that is applicable to the aforementioned settings. Specifically, inspired by a recent study on multi-task multi-armed bandits [43], we formulate the ϵ -Multi-Player Episodic Reinforcement Learning (abbreviated as ϵ -MPERL) problem, in which all tasks share the same state and action spaces, and the tasks are assumed to be similar—i.e., the dissimilarities between the environments of different tasks (specifically, the reward distributions and transition dynamics associated with the players/tasks) are bounded in terms of a dissimilarity parameter $\epsilon \ge 0$. This problem not only models concurrent RL [34, [16] as a special case by taking $\epsilon = 0$, but also captures richer multi-task RL settings when ϵ is nonzero. We study regret minimization for the ϵ -MPERL problem, specifically:

- 1. We identify a problem complexity notion named *subpar* state-action pairs, which captures the amenability to information sharing among tasks in ϵ -MPERL problem instances. As shown in the multi-task bandits literature (e.g., [43]), inter-task information sharing is *not* always helpful in reducing the players' collective regret. Subpar state-action pairs, intuitively speaking, are clearly suboptimal for all tasks, for which we can robustly take advantage of (possibly biased) data collected for other tasks to achieve a lower regret in a certain task.
- 2. In the setting where the dissimilarity parameter ϵ is known, we design a model-based algorithm MULTI-TASK-EULER (Algorithm), which is built upon state-of-the-art algorithms for learning single-task Markov decision processes (MDPs) [3, 46, 36], as well as algorithmic ideas of model transfer in RL [39]. MULTI-TASK-EULER crucially utilizes the dissimilarity assumption to robustly take advantage of information sharing among tasks, and achieves regret upper bounds in terms of subpar state-action pairs, in both (value function suboptimality) gap-dependent and gap-independent fashions. Specifically, compared with a baseline algorithm that does not utilize information sharing, MULTI-TASK-EULER has a regret guarantee that: (1) is never worse, i.e., it avoids negative transfer [33]; (2) can be much superior when there are a large number of subpar state-action pairs.
- 3. We also present gap-dependent and gap-independent regret lower bounds for the ϵ -MPERL problem in terms of subpar state-action pairs. These lower bounds nearly match the upper bounds when the episode length of the MDP is a constant. Together, the upper and lower bounds can be used to characterize the intrinsic complexity of the ϵ -MPERL problem.

2 Preliminaries

Throughout this paper, we denote by $[n] := \{1, \ldots, n\}$. For a set A in a universe U, we use $A^C = U \setminus A$ to denote its complement. Denote by $\Delta(\mathcal{X})$ the set of probability distributions over \mathcal{X} . For functions f, g, we use $f \leq g$ or f = O(g) (resp. $f \geq g$ or $f = \Omega(g)$) to denote that there exists some constant c > 0, such that $f \leq cg$ (resp. $f \geq cg$), and use f = g to denote $f \leq g$ and $f \geq g$ simultaneously. Define $a \lor b := \max(a, b)$, and $a \land b := \min(a, b)$. We use \mathbb{E} to denote the expectation operator, and use var to denote the variance operator. Throughout, we use $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ notation to hide polylogarithmic factors.

Multi-task RL in episodic MDPs. We have a set of M MDPs $\{\mathcal{M}_p = (H, S, \mathcal{A}, p_0, \mathbb{P}_p, r_p)\}_{p=1}^M$, each associated with a player $p \in [M]$. Each MDP \mathcal{M}_p is regarded as a task. The MDPs share the same episode length $H \in \mathbb{N}_+$, finite state space S, finite action space \mathcal{A} , and initial state distribution $p_0 \in \Delta(S)$. Let \bot be a default terminal state that is not contained in S. The transition probabilities $\mathbb{P}_p : S \times \mathcal{A} \to \Delta(S \cup \{\bot\})$ and reward distributions $r_p : S \times \mathcal{A} \to \Delta([0,1])$ of the players are not necessarily identical. We assume that the MDPs are layered, in that the state space S can be partitioned into disjoint subsets $(S_h)_{h=1}^H$, where p_0 is supported on S_1 , and for every $p \in [M]$, $h \in [H]$, and every $s \in S_h, a \in \mathcal{A}, \mathbb{P}_p(\cdot | s, a)$ is supported on S_{h+1} ; here, we define $S_{H+1} = \{\bot\}$. We denote by S := |S| the size of the state space, and $A := |\mathcal{A}|$ the size of the action space.

Interaction process. The interaction process between the players and the environment is as follows: at the beginning, both $(r_p)_{p=1}^M$ and $(\mathbb{P}_p)_{p=1}^M$ are unknown to the players. For each episode $k \in [K]$, conditioned on the interaction history up to episode k - 1, each player $p \in [M]$ independently interacts with its respective MDP \mathcal{M}_p ; specifically, player p starts at state $s_{1,p}^k \sim p_0$, and at every step (layer) $h \in [H]$, it chooses action $a_{h,p}^k$, transitions to next state $s_{h+1,p}^k \sim \mathbb{P}_p(\cdot \mid s_{h,p}^k, a_{h,p}^k)$ and receives a stochastic immediate reward $r_{h,p}^k \sim r_p(\cdot \mid s_{h,p}^k, a_{h,p}^k)$; after all players have finished their k-th episode, they can communicate and share their interaction history. The goal of the players is to maximize their expected collective reward $\mathbb{E}\left[\sum_{k=1}^K \sum_{p=1}^M \sum_{h=1}^H r_{h,p}^k\right]$.

¹This is a standard assumption (see, e.g., [44]). It is worth noting that any episodic MDP (with possibly nonstationary transition and reward) can be converted to a layered MDP with stationary transition and reward, with the state space size being *H* times the size of the original state space.

Policy and value functions. A deterministic, history-independent policy π is a mapping from S to A, which can be used by a player to make decisions in its respective MDP. For player p and step h, we use $V_{h,p}^{\pi} : S_h \to [0, H]$ and $Q_{h,p}^{\pi} : S_h \times A \to [0, H]$ to denote its respective value and action-value functions, respectively. They satisfy the following recurrence known as the Bellman equation:

$$\forall h \in [H]: \quad V_{h,p}^{\pi}(s) = Q_{h,p}^{\pi}(s,\pi(s)), \quad Q_{h,p}^{\pi}(s,a) = R_p(s,a) + (\mathbb{P}_p V_{h+1,p}^{\pi})(s,a) + (\mathbb{P}_p V_{$$

where we use the convention that $V_{H+1,p}^{\pi}(\perp) = 0$, and for $f : S_{h+1} \to \mathbb{R}$, $(\mathbb{P}_p f)(s, a) := \sum_{s' \in S_{h+1}} \mathbb{P}_p(s' \mid s, a) f(s')$, and $R_p(s, a) := \mathbb{E}_{\hat{r} \sim r_p(\cdot \mid s, a)} [\hat{r}]$ is the expected immediate reward of player p. For player p and policy π , denote by $V_{0,p}^{\pi} = \mathbb{E}_{s_1 \sim p_0} \left[V_{1,p}^{\pi}(s_1) \right]$ its expected reward.

For player p, we also define its optimal value function $V_{h,p}^{\star} : S_h \to [0, H]$ and the optimal actionvalue function $Q_{h,p}^{\star} : S_h \times \mathcal{A} \to [0, H]$ using the Bellman optimality equation:

$$\forall h \in [H]: \quad V_{h,p}^{\star}(s) = \max_{a \in \mathcal{A}} Q_{h,p}^{\star}(s,a), \quad Q_{h,p}^{\star}(s,a) = R_p(s,a) + (\mathbb{P}_p V_{h+1,p}^{\star})(s,a), \quad (1)$$

where we again use the convention that $V_{H+1,p}^{\star}(\perp) = 0$. For player p, denote by $V_{0,p}^{\star} = \mathbb{E}_{s_1 \sim p_0} \left[V_{1,p}^{\star}(s_1) \right]$ its optimal expected reward.

Given a policy π , as $V_{h,p}^{\pi}$ for different h's are only defined in the respective layer S_h , we "collate" the value functions $(V_{h,p}^{\pi})_{h=1}^{H}$ and obtain a single value function $V_p^{\pi} : S \cup \{\bot\} \to \mathbb{R}$. Formally, for every $h \in [H+1]$ and $s \in S_h$,

$$V_p^{\pi}(s) := V_{h,p}^{\pi}(s).$$

We define $Q_p^{\pi}, V_p^{\star}, Q_p^{\star}$ similarly. For player p, given its optimal action value function Q_p^{\star} , any of its greedy policies $\pi_p^{\star}(s) \in \operatorname{argmax}_{a \in \mathcal{A}} Q_p^{\star}(s, a)$ is optimal with respect to \mathcal{M}_p .

Suboptimality gap. For player p, we define the suboptimality gap of state-action pair (s, a) as $\operatorname{gap}_p(s, a) = V_p^*(s) - Q_p^*(s, a)$. We define the minimum suboptimality gap of player p as $\operatorname{gap}_{p,\min} = \min_{(s,a):\operatorname{gap}_p(s,a)>0} \operatorname{gap}_p(s, a)$, and the minimum suboptimality gap over all players as $\operatorname{gap}_{\min} = \min_{p \in [M]} \operatorname{gap}_{p,\min}$. For player $p \in [M]$, define $Z_{p,\operatorname{opt}} := \{(s,a): \operatorname{gap}_p(s,a) = 0\}$ as the set of optimal state-action pairs with respect to p.

Performance metric. We measure the performance of the players using their collective regret, i.e., over a total of K episodes, how much extra reward they would have collected in expectation if they were executing their respective optimal policies from the beginning. Formally, suppose for each episode k, player p executes policy $\pi^k(p)$, then the collective regret of the players is defined as:

$$\operatorname{Reg}(K) = \sum_{p=1}^{M} \sum_{k=1}^{K} \left(V_{0,p}^{\star} - V_{0,p}^{\pi^{k}(p)} \right).$$

Baseline: individual STRONG-EULER. A naive baseline for multi-task RL is to let each player run a separate RL algorithm without communication. For concreteness, we choose to let each player run the state-of-the-art STRONG-EULER algorithm [36] (see also its precursor EULER [46]), which enjoys minimax gap-independent [3] [8] and gap-dependent regret guarantees, and we refer to this strategy as individual STRONG-EULER. Specifically, as it is known that STRONG-EULER has a regret of $\tilde{O}(\sqrt{H^2SAK} + H^4S^2A)$, individual STRONG-EULER has a collective regret of $\tilde{O}(M\sqrt{H^2SAK} + MH^4S^2A)$. In addition, by a union bound and summing up the gap-dependent regret guarantees of STRONG-EULER for the M MDPs altogether, it can be checked that with probability $1 - \delta$, individual STRONG-EULER has a collective regret of order

$$\ln\left(\frac{MSAK}{\delta}\right)\left(\sum_{p\in[M]}\left(\sum_{(s,a)\in Z_{p,\text{opt}}}\frac{H^3}{\text{gap}_{p,\min}} + \sum_{(s,a)\in Z_{p,\text{opt}}^C}\frac{H^3}{\text{gap}_p(s,a)}\right) + MH^4S^2A\ln\frac{SA}{\text{gap}_{\min}}\right)$$
(2)

Our goal is to design multi-task RL algorithms that can achieve collective regret strictly lower than this baseline in both gap-dependent and gap-independent fashions when the tasks are similar.

Notion of similarity. Throughout this paper, we will consider the following notion of similarity between MDPs in the multi-task episodic RL setting.

Definition 1. A collection of MDPs $(\mathcal{M}_p)_{p=1}^M$ is said to be ϵ -dissimilar for $\epsilon \ge 0$, if for all $p, q \in [M]$, and $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\left|R_p(s,a) - R_q(s,a)\right| \le \epsilon, \ \|\mathbb{P}_p(\cdot \mid s,a) - \mathbb{P}_q(\cdot \mid s,a)\|_1 \le \frac{\epsilon}{H}$$

If this happens, we call $(\mathcal{M}_p)_{p=1}^M$ an ϵ -Multi-Player Episodic Reinforcement Learning (abbrev. ϵ -MPERL) problem instance.

If the MDPs in $(\mathcal{M}_p)_{p=1}^M$ are 0-dissimilar, then they are identical by definition, and our interaction protocol degenerates to the concurrent RL protocol [34]. Our dissimilarity notion is complementary to those of [6, 16]: they require the MDPs to be either identical, or have well-separated parameters for at least one state-action pair; in contrast, our dissimilarity notion allows the MDPs to be nonidentical and arbitrarily close.

We have the following intuitive lemma that shows the closeness of optimal value functions of different MDPs, in terms of the dissimilarity parameter ϵ :

Lemma 2. If
$$(\mathcal{M}_p)_{p=1}^M$$
 are ϵ -dissimilar, then for every $p, q \in [M]$, and $(s, a) \in \mathcal{S} \times \mathcal{A}$,
 $\left|Q_p^{\star}(s, a) - Q_q^{\star}(s, a)\right| \leq 2H\epsilon$; consequently, $\left|\operatorname{gap}_p(s, a) - \operatorname{gap}_q(s, a)\right| \leq 4H\epsilon$.

3 Algorithm

We now describe our main algorithm, MULTI-TASK-EULER (Algorithm []). Our model-based algorithm is built upon recent works on episodic RL that provide algorithms with sharp instance-dependent guarantees in the single task setting [46, 36]. In a nutshell, for each episode k and each player p, the algorithm performs optimistic value iteration to construct high-probability upper and lower bounds for the optimal value and action value functions V_p^* and Q_p^* , and uses them to guide its exploration and decision making process.

Empirical estimates of model parameters. For each player p, the construction of its value function bound estimates relies on empirical estimates on its transition probability and expected reward function. For both estimands, we use two estimators with complementary roles, which are at two different points of the bias-variance tradeoff spectrum: one estimator uses only the player's own data (termed *individual estimate*), which has large variance; the other estimator uses the data collected by all players (termed *aggregate estimate*), which has lower variance but can easily be biased, as transition probabilities and reward distributions are heterogeneous. Such an algorithmic idea of "model transfer", where one estimates model in one task using data collected from other tasks has appeared in prior works (e.g., [39]). Specifically, at the beginning of episode k, for every $h \in [H]$ and $(s, a) \in S_h \times A$, the algorithm has its empirical count of encountering (s, a) for each player p, along with its total empirical count across all players, respectively:

$$n_p(s,a) := \sum_{l=1}^{k-1} \mathbf{1}\left((s_{h,p}^l, a_{h,p}^l) = (s,a) \right), \ n(s,a) := \sum_{l=1}^{k-1} \sum_{p=1}^M \mathbf{1}\left((s_{h,p}^l, a_{h,p}^l) = (s,a) \right).$$
(3)

²The originally-stated gap-dependent regret bound of STRONG-EULER ([36], Corollary 2.1) uses a slightly different notion of suboptimality gap, which takes an extra minimum over all steps. A close examination of their proof shows that STRONG-EULER has regret bound (2) in layered MDPs. See also Remark 21 in Appendix .

Algorithm 1: MULTI-TASK-EULER

Input: Failure probability $\delta \in (0, 1)$, dissimilarity parameter $\epsilon \ge 0$. **Initialize:** Set $V_p(\perp) = 0$ for all p in [M], where \perp is the only state in S_{H+1} ; 1 for $k = 1, 2, \dots, K$ do for p = 1, 2, ..., M do 2 // Construct optimal value estimates for player pfor h = H, H - 1, ..., 1 do 3 for $(s, a) \in \mathcal{S}_h \times \mathcal{A}$ do 4 Compute: 5
$$\begin{split} \overline{\mathrm{ind}\text{-}Q}_p(s,a) &= \hat{R}_p(s,a) + (\hat{\mathbb{P}}_p \overline{V}_p)(s,a) + \mathrm{ind}\text{-}b_p(s,a);\\ \underline{\mathrm{ind}\text{-}Q}_p(s,a) &= \hat{R}_p(s,a) + (\hat{\mathbb{P}}_p \underline{V}_p)(s,a) - \mathrm{ind}\text{-}b_p(s,a); \end{split}$$
6 7
$$\begin{split} \overline{\underline{\mathsf{agg-}Q}}_p^p(s,a) &= \hat{R}(s,a) + (\hat{\mathbb{P}}\overline{V}_p)(s,a) + \mathtt{agg-}b_p(s,a);\\ \underline{\mathsf{agg-}Q}_p(s,a) &= \hat{R}(s,a) + (\hat{\mathbb{P}}\underline{V}_p)(s,a) - \mathtt{agg-}b_p(s,a); \end{split}$$
8 9 Update optimal action value function upper and lower bound estimates: 10 $\overline{Q}_p(s,a) = \min\left\{H - h + 1, \overline{\operatorname{ind-}Q}_p(s,a), \overline{\operatorname{agg-}Q}_p(s,a)\right\};$ 11 $\underline{Q}_p(s,a) = \max \left\{ 0, \underline{\mathrm{ind-}}Q_p(s,a), \underline{\mathrm{agg-}}Q_p(s,a) \right\};$ 12 for $s \in \mathcal{S}_h$ do 13 $\begin{array}{l} \text{Define } \pi^k(p)(s) = \operatorname{argmax}_{a \in \mathcal{A}} \overline{Q}_p(s,a) \ ; \\ \text{Update } \overline{V}_p(s) = \overline{Q}_p\left(s,\pi^k(p)(s)\right), \underline{V}_p(s) = \underline{Q}_p\left(s,\pi^k(p)(s)\right). \end{array}$ 14 15 // All players p interact with their respective environments, and update reward and transition estimates for p = 1, 2, ..., M do 16 Player p executes policy $\pi^k(p)$ on \mathcal{M}_p and obtains trajectory $(s_{h,p}^k, a_{h,p}^k, r_{h,p}^k)_{h=1}^H$. 17 Update individual estimates of transition probability $\hat{\mathbb{P}}_p$, reward \hat{R}_p and count $n_p(\cdot, \cdot)$ 18 using the first parts of Equations (3), (4) and (5). Update aggregate estimates of transition probability $\hat{\mathbb{P}}$, reward \hat{R} and count $n(\cdot, \cdot)$ using the 19 second parts of Equations (3), (4) and (5).

The individual and aggregate estimates of immediate reward R(s, a) are defined as:

$$\hat{R}_{p}(s,a) := \frac{\sum_{l=1}^{k-1} \mathbf{1} \left((s_{h,p}^{l}, a_{h,p}^{l}) = (s,a) \right) r_{h,p}^{l}}{n_{p}(s,a)}, \ \hat{R}(s,a) := \frac{\sum_{l=1}^{k-1} \sum_{p=1}^{M} \mathbf{1} \left((s_{h,p}^{l}, a_{h,p}^{l}) = (s,a) \right) r_{h,p}^{l}}{n(s,a)}$$
(4)

Similarly, for every $h \in [H]$ and $(s, a, s') \in S_h \times A \times S_{h+1}$, we also define the individual and aggregate estimates of transition probability as:

$$\hat{\mathbb{P}}_{p}(s' \mid s, a) := \frac{\sum_{l=1}^{k} \mathbf{1} \left((s_{h,p}^{l}, a_{h,p}^{l}, s_{h+1,p}^{l}) = (s, a, s') \right)}{n_{p}(s, a)},$$

$$\hat{\mathbb{P}}(s' \mid s, a) := \frac{\sum_{l=1}^{k} \sum_{p=1}^{M} \mathbf{1} \left((s_{h,p}^{l}, a_{h,p}^{l}, s_{h+1,p}^{l}) = (s, a, s') \right)}{n(s, a)}.$$
(5)

If n(s, a) = 0, we define $\hat{R}(s, a) := 0$ and $\hat{\mathbb{P}}(s' \mid s, a) := \frac{1}{|S_{h+1}|}$; and if $n_p(s, a) = 0$, we define $\hat{R}_p(s, a) := 0$ and $\hat{\mathbb{P}}_p(s' \mid s, a) := \frac{1}{|S_{h+1}|}$. The counts and reward estimates can be maintained by MULTI-TASK-EULER efficiently in an incremental manner.

Constructing value function estimates via optimistic value iteration. For each player p, based on these model parameter estimates, MULTI-TASK-EULER performs optimistic value iteration to

compute the value function estimates for states at all layers (lines [] to []). For the terminal layer H + 1, $V_{H+1}^{\star}(\perp) = 0$ trivially, so nothing needs to be done. For earlier layers $h \in [H]$, MULTI-TASK-EULER iteratively builds its value function estimates in a backward fashion. At the time of estimating values for layer h, the algorithm has already obtained optimal value estimates for layer h + 1. Based on the Bellman optimality equation (]], MULTI-TASK-EULER estimates $(Q_p^{\star}(s,a))_{s\in S_h,a\in \mathcal{A}}$ using model parameter estimates and its estimates of $(V_p^{\star}(s))_{s\in S_{h+1}}$, i.e., $(\overline{V}_p(s))_{s\in S_{h+1}}$ and $(\underline{V}_p(s))_{s\in S_{h+1}}$ (lines []).

Specifically, MULTI-TASK-EULER constructs estimates of $Q_p^*(s, a)$ for all $s \in S_h, a \in A$ in two different ways. First, it uses the individual estimates of model of player p to construct $\underline{ind} - Q_p$ and $\overline{ind} - Q_p$, upper and lower bound estimates of Q_p^* (lines and \mathfrak{D}); this construction is reminiscent of EULER and STRONG-EULER [46, 36], in that if we were only to use $\underline{ind} - Q_p$ and $\overline{ind} - Q_p$ as our optimal action value function estimates \overline{Q}_p and \underline{Q}_p , our algorithm becomes individual STRONG-EULER. The individual value function estimates are key to establishing MULTI-TASK-EULER's fall-back guarantees, ensuring that it never performs worse than the individual STRONG-EULER baseline. Second, it uses the aggregate estimate of model to construct $\underline{agg} - Q_p$ and $\overline{agg} - \overline{Q}_p$, also upper and lower bound estimates of Q_p^* (lines **6** and **7**); this construction is unique to the multitask learning setting, and is our new algorithmic contribution.

To ensure that $\overline{\text{agg-}Q_p}$ and $\overline{\text{ind-}Q_p}$ (resp. $\underline{\text{agg-}Q_p}$ and $\underline{\text{ind-}Q_p}$) are valid upper bounds (resp. lower bounds) of Q_p^{\star} , MULTI-TASK-EULER adds bonus terms ind- $b_p(s, a)$ and $\operatorname{agg-}b_p(s, a)$, respectively, in the optimistic value iteration process, to account for estimation error of the model estimates against the true models. Specifically, both bonus terms comprise three parts:

$$\begin{split} \operatorname{ind} b_p(s,a) &\coloneqq b_{\operatorname{rw}}\left(n_p(s,a),0\right) + b_{\operatorname{prob}}\left(\hat{\mathbb{P}}_p(\cdot \mid s,a), n_p(s,a), \overline{V}_p, \underline{V}_p, 0\right) + \\ b_{\operatorname{str}}\left(\hat{\mathbb{P}}_p(\cdot \mid s,a), n_p(s,a), \overline{V}_p, \underline{V}_p, 0\right), \\ \operatorname{agg-}b_p(s,a) &\coloneqq b_{\operatorname{rw}}\left(n(s,a),\epsilon\right) + b_{\operatorname{prob}}\left(\hat{\mathbb{P}}(\cdot \mid s,a), n(s,a), \overline{V}_p, \underline{V}_p, \epsilon\right) + \\ b_{\operatorname{str}}\left(\hat{\mathbb{P}}(\cdot \mid s,a), n(s,a), \overline{V}_p, \underline{V}_p, \epsilon\right), \end{split}$$

where

l

$$b_{\mathrm{rw}}(n,\kappa) \coloneqq 1 \wedge \kappa + \Theta\left(\sqrt{\frac{L(n)}{n}}\right),$$

$$b_{\mathrm{prob}}\left(q,n,\overline{V},\underline{V},\kappa\right) \coloneqq H \wedge 2\kappa + \Theta\left(\sqrt{\frac{\operatorname{var}_{s'\sim q}\left[\overline{V}(s')\right]L(n)}{n}} + \frac{1}{n}\right),$$

$$\sqrt{\frac{\mathbb{E}_{s'\sim q}\left[(\overline{V}(s')-\underline{V}(s'))^2\right]L(n)}{n}} + \frac{HL(n)}{n}\right),$$

$$b_{\mathrm{str}}\left(q,n,\overline{V},\underline{V},\kappa\right) \coloneqq \kappa + \Theta\left(\sqrt{\frac{S \mathbb{E}_{s'\sim q}\left[(\overline{V}(s')-\underline{V}(s'))^2\right]L(n)}{n}} + \frac{HSL(n)}{n}\right),$$

and $L(n) \eqsim \ln(\frac{MSAn}{\delta})$.

The bonus terms altogether ensures strong optimism [36], i.e.,

for any
$$p$$
 and (s, a) , $\overline{Q}_p(s, a) \ge R_p(s, a) + (\mathbb{P}_p \overline{V}_p)(s, a).$ (6)

In short, strong optimism is a stronger form of optimism (the weaker requirement that for any p and (s, a), $\overline{Q}_p(s, a) \ge Q_p^{\star}(s, a)$ and $\overline{V}_p(s) \ge V_p^{\star}(s)$), which allows us to use the clipping lemma

(Lemma B.6 of [36], see also Lemma 20 in Appendix (C.4) to obtain sharp gap-dependent regret guarantees. The three parts in the bonus term serve for different purposes towards establishing (6):

- 1. The first component accounts for the uncertainty in reward estimation: with probability $1 - O(\delta)$, $|\hat{R}_p(s, a) - R_p(s, a)| \leq b_{rw} (n_p(s, a), 0)$, and $|\hat{R}(s, a) - R_p(s, a)| \leq b_{rw} (n(s, a), \epsilon)$.
- 2. The second component accounts for the uncertainty in estimating $(\mathbb{P}_p V_p^{\star})(s, a)$: with probability $1 O(\delta)$, $\left| (\hat{\mathbb{P}}_p V_p^{\star})(s, a) (\mathbb{P}_p V_p^{\star})(s, a) \right| \le b_{\text{prob}} \left(\hat{\mathbb{P}}_p(\cdot \mid s, a), n_p(s, a), \overline{V}_p, \underline{V}_p, 0 \right)$ and $\left| (\hat{\mathbb{P}} V_p^{\star})(s, a) - (\mathbb{P}_p V_p^{\star})(s, a) \right| \le b_{\text{prob}} \left(\hat{\mathbb{P}}(\cdot \mid s, a), n(s, a), \overline{V}_p, \underline{V}_p, \epsilon \right).$
- 3. The third component accounts for the lower order terms for strong optimism: with probability $1 - O(\delta)$, $\left| (\hat{\mathbb{P}}_p - \mathbb{P}_p)(\overline{V}_p - V_p^{\star})(s, a) \right| \leq b_{\text{str}} \left(\hat{\mathbb{P}}_p(\cdot \mid s, a), n_p(s, a), \overline{V}_p, \underline{V}_p, 0 \right)$, and $\left| (\hat{\mathbb{P}} - \mathbb{P}_p)(\overline{V}_p - V_p^{\star})(s, a) \right| \leq b_{\text{str}} \left(\hat{\mathbb{P}}(\cdot \mid s, a), n(s, a), \overline{V}_p, \underline{V}_p, \epsilon \right)$.

Based on the above concentration inequalities and the definitions of bonus terms, it can be shown inductively that, with probability $1 - O(\delta)$, both $\overline{\text{agg-}Q_p}$ and $\overline{\text{ind-}Q_p}$ (resp. $\underline{\text{agg-}Q_p}$ and $\underline{\text{ind-}Q_p}$) are valid upper bounds (resp. lower bounds) of Q_p^* .

Finally, observe that for any $(s, a) \in S_h \times A$, $Q_p^*(s, a)$ has range [0, H - h + 1]. By taking intersections of all confidence bounds of Q_p^* it has obtained, MULTI-TASK-EULER constructs its final upper and lower bound estimates for $Q_p^*(s, a)$, $\overline{Q}_p(s, a)$ and $\underline{Q}_p(s, a)$ respectively, for $(s, a) \in S_h \times A$ (line \square to \square). Similar ideas on using data from multiple sources to construct confidence intervals and guide explorations have been used by [37], [43] for multi-task noncontextual and contextual bandits. Using the relationship between the optimal value $V_p^*(s)$ and and optimal action values $\left\{Q_p^*(s, a) : a \in A\right\}$, MULTI-TASK-EULER also constructs upper and lower bound estimates for $V_n^*(s)$, $\overline{V}_p(s)$ and $\underline{V}_p(s)$, respectively for $s \in S_h$ (line \square).

Executing optimistic policies. At each episode k, for each player p, its optimal action-value function upper bound estimate \overline{Q}_p induces a greedy policy $\pi^k(p) : s \mapsto \operatorname{argmax}_{a \in \mathcal{A}} \overline{Q}_p(s, a)$ (line 14); the player then executes this policy in this episode to collect a new trajectory and use this to update its individual model parameter estimates. After all players finish their episode k, the algorithm also updates its aggregate model parameter estimates (lines 16 to 19) using Equations (3), (4) and (5), and continues to the next episode.

4 Performance guarantees

Before stating the guarantees of Algorithm [] we define an instance-dependent complexity measure that characterizes the amenability to information sharing.

Definition 3. The set of subpar state-action pairs is defined as:

$$\mathcal{I}_{\epsilon} := \left\{ (s, a) \in \mathcal{S} \times \mathcal{A} : \exists p \in [M], \operatorname{gap}_{p}(s, a) > 96H\epsilon \right\},\$$

where we recall that $gap_n(s, a) = V_n^{\star}(s) - Q_n^{\star}(s, a)$.

Definition 3 generalizes the notion of subpar arms defined for multi-task multi-armed bandit learning [43] in two ways: first, it is with regards to state-action pairs as opposed to actions only; second, in RL, suboptimality gaps depend on optimal value function, which in turn depends on both immediate reward and subsequent long-term return.

To ease our later presentation, we also present the following lemma.

Lemma 4. For any $(s, a) \in \mathcal{I}_{\epsilon}$, we have that: (1) for all $p \in [M]$, $(s, a) \notin Z_{p,\text{opt}}$, where we recall that $Z_{p,\text{opt}} = \{(s, a) : \text{gap}_p(s, a) = 0\}$ is the set of optimal state-action pairs with respect to p; (2) for all $p, q \in [M]$, $\text{gap}_p(s, a) \ge \frac{1}{2}\text{gap}_q(s, a)$.

The lemma follows directly from Lemma 2; its proof can be found in the Appendix along with proofs of the following theorems. Item 1 implies that any subpar state-action pair is suboptimal for all players. In other words, for every player p, the state-action space $S \times A$ can be partitioned to three disjoint sets: $\mathcal{I}_{\epsilon}, Z_{p,\text{opt}}, (\mathcal{I}_{\epsilon} \cup Z_{p,\text{opt}})^C$. Item 2 implies that for any subpar (s, a), its suboptimal gaps with respect to all players are within a constant of each other.

4.1 Upper bounds

With the above definitions, we are now ready to present the performance guarantees of Algorithm We first present a gap-independent collective regret bound of MULTI-TASK-EULER.

Theorem 5 (Gap-independent bound). If $\{\mathcal{M}_p\}_{p=1}^M$ are ϵ -dissimilar, then MULTI-TASK-EULER satisfies that with probability $1 - \delta$,

$$\operatorname{Reg}(K) \leq \tilde{O}\left(M\sqrt{H^2|\mathcal{I}_{\epsilon}^C|K} + \sqrt{MH^2|\mathcal{I}_{\epsilon}|K} + MH^4S^2A\right).$$

We again compare this regret upper bound with individual STRONG-EULER's gap independent regret bound. Recall that individual STRONG-EULER guarantees that with probability $1 - \delta$,

$$\operatorname{Reg}(K) \leq \tilde{O}\left(M\sqrt{H^2SAK} + MH^4S^2A\right).$$

We focus on the comparison on the leading terms, i.e., the \sqrt{K} terms. As $M\sqrt{H^2SAK} \approx M\sqrt{H^2|\mathcal{I}_{\epsilon}|K} + M\sqrt{H^2|\mathcal{I}_{\epsilon}^C|K}$, we see that an improvement in the collective regret bound comes from the contributions from subpar state-action pairs: the $M\sqrt{H^2|\mathcal{I}_{\epsilon}|K}$ term is reduced to $\sqrt{MH^2|\mathcal{I}_{\epsilon}|K}$, a factor of $\tilde{O}(\sqrt{\frac{1}{M}})$ improvement. Moreover, if $|\mathcal{I}_{\epsilon}^C| \ll SA$ and $M \gg 1$, MULTI-TASK-EULER provides a regret bound of lower order than individual STRONG-EULER.

We next present a gap-dependent upper bound on its collective regret.

Theorem 6 (Gap-dependent upper bound). If $\{\mathcal{M}_p\}_{p=1}^M$ are ϵ -dissimilar, then MULTI-TASK-EULER satisfies with probability $1 - \delta$,

$$\operatorname{Reg}(K) \lesssim \ln(\frac{MSAK}{\delta}) \left(\sum_{p \in [M]} \left(\sum_{(s,a) \in \mathbb{Z}_{p,\text{opt}}} \frac{H^3}{\operatorname{gap}_{p,\min}} + \sum_{(s,a) \in (\mathcal{I}_{\epsilon} \cup \mathbb{Z}_{p,\text{opt}})^C} \frac{H^3}{\operatorname{gap}_p(s,a)} \right) + \sum_{(s,a) \in \mathcal{I}_{\epsilon}} \frac{H^3}{\min_p \operatorname{gap}_p(s,a)} \right) + \ln(\frac{MSAK}{\delta}) \cdot MH^4 S^2 A \ln \frac{MSA}{\operatorname{gap}_{\min}},$$

where we recall that $gap_{p,\min} = \min_{(s,a): gap_p(s,a)>0} gap_p(s,a)$, and $gap_{\min} = \min_p gap_{p,\min}$.

Comparing this regret bound with the regret bound obtained by the individual STRONG-EULER baseline, recall that by summing over the regret guarantees of STRONG-EULER for all players $p \in [M]$, and taking a union bound over all p, individual STRONG-EULER guarantees a collective regret bound of

$$\operatorname{Reg}(K) \lesssim \ln(\frac{MSAK}{\delta}) \left(\sum_{p \in [M]} \left(\sum_{(s,a) \in \mathbb{Z}_{p,\text{opt}}} \frac{H^3}{\operatorname{gap}_{p,\min}} + \sum_{(s,a) \in (\mathbb{Z}_{\epsilon} \cup \mathbb{Z}_{p,\text{opt}})^C} \frac{H^3}{\operatorname{gap}_p(s,a)} \right) + \sum_{(s,a) \in \mathbb{I}_{\epsilon}} \sum_{p \in [M]} \frac{H^3}{\operatorname{gap}_p(s,a)} \right) + \ln(\frac{MSAK}{\delta}) \cdot MH^4 S^2 A \ln \frac{SA}{\operatorname{gap}_{\min}},$$

that holds with probability $1 - \delta$. We again focus on comparing the leading terms, i.e., the terms that have polynomial dependences on the suboptimality gaps in the above two bounds. It can be seen that an improvement in the regret bound by MULTI-TASK-EULER comes from the contributions from the subpar state-action pairs: for each $(s, a) \in \mathcal{I}_{\epsilon}$, the regret bound is reduced from $\sum_{p \in [M]} \frac{H^3}{\text{gap}_p(s, a)}$ to

 $\frac{H^3}{\min_p \operatorname{gap}_p(s,a)}$, a factor of $O(\frac{1}{M})$ improvement. Recent work of [44] has shown that in the single-task setting, it is possible to replace $\sum_{(s,a)\in Z_{p,\operatorname{opt}}} \frac{H^3}{\operatorname{gap}_{p,\min}}$ with a sharper problem-dependent complexity term that depends on the multiplicity of optimal state-action pairs. We leave improving the guarantee of Theorem 6 in a similar manner as an interesting open problem.

Key to the proofs of Theorems 5 and 6 is a new bound on the *surplus* 36 of the value function estimates. Our new surplus bound is a minimum of two terms: one depends on the usual stateaction visitation counts of player p, the other depends on the task dissimilarity parameter ϵ and the state-action visitation counts of all players. Detailed proofs can be found at Appendix C

4.2 Lower bounds

To complement the above upper bounds, we now present gap-dependent and gap-independent regret lower bounds that also depend on our subpar state-action pair notion. Our lower bounds are inspired by regret lower bounds for episodic RL [36, 8] and multi-task bandits [43].

Theorem 7 (Gap-independent lower bound). For any $A \ge 2$, $H \ge 2$, $S \ge 4H$, $K \ge SA$, $M \in \mathbb{N}$, and $l, l^C \in \mathbb{N}$ with $l + l^C = SA$ and $l \le SA - 4(S + HA)$, there exists some ϵ that satisfies: for any algorithm Alg, there exists an ϵ -MPERL problem instance with S states, A actions, M players and an episode length of H such that $\left|\mathcal{I}_{\frac{\epsilon}{192H}}\right| \ge l$, and

$$\mathbb{E}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right] \geq \Omega\left(M\sqrt{H^2l^CK} + \sqrt{MH^2lK}\right).$$

We also present a gap-dependent lower bound. Before that, we first formally define the notion of sublinear regret algorithms: for any fixed ϵ , we say that an algorithm Alg is a sublinear regret algorithm for the ϵ -MPERL problem if there exists some C > 0 (that possibly depends on the state-action space, the number of players, and ϵ) and $\alpha < 1$ such that for all K and all ϵ -MPERL environments, $\mathbb{E}\left[\operatorname{Reg}_{Alg}(K)\right] \leq CK^{\alpha}$.

Theorem 8 (Gap-dependent lower bound). Fix $\epsilon \ge 0$. For any $S \in \mathbb{N}$, $A \ge 2$, $H \ge 2$, $M \in \mathbb{N}$, with $S \ge 2(H-1)$, let $S_1 = S - 2(H-1)$; and let $\{\Delta_{s,a,p}\}_{(s,a,p)\in[S_1]\times[A]\times[M]}$ be any set of values that satisfies: (1) each $\Delta_{s,a,p} \in [0, H/(48\sqrt{M})]$, (2) for every $(s,p) \in [S_1] \times [M]$, there exists at least one action $a \in [A]$ such that $\Delta_{s,a,p} = 0$, and (3) for every $(s,a) \in [S_1] \times [A]$ and $p, q \in [M]$, $|\Delta_{s,a,p} - \Delta_{s,a,q}| \le \epsilon/4$. There exists an ϵ -MPERL problem instance with S states, A actions, M players and an episode length of H, such that $S_1 = [S_1]$, $|S_h| = 2$ for all $h \ge 2$, and

$$\operatorname{gap}_p(s, a) = \Delta_{s, a, p}, \quad \forall (s, a, p) \in [S_1] \times [A] \times [M];$$

for this problem instance, any sublinear regret algorithm Alg for the ϵ -MPERL problem must satisfy:

$$\mathbb{E}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right] \ge \Omega \left(\ln K \left(\sum_{\substack{p \in [M] \ (s,a) \in \mathcal{I}_{(\epsilon/768H)}^C \\ \operatorname{gap}_p(s,a) > 0}} \frac{H^2}{\operatorname{gap}_p(s,a)} + \sum_{\substack{(s,a) \in \mathcal{I}_{(\epsilon/768H)} \\ \operatorname{gap}_p(s,a) > 0}} \frac{H^2}{\operatorname{min}_p \operatorname{gap}_p(s,a)} \right) \right)$$

Comparing the lower bounds with MULTI-TASK-EULER's regret upper bounds in Theorems and we see that the upper and lower bounds nearly match for any constant H. When H is large, a key difference between the upper and lower bounds is that the former are in terms of \mathcal{I}_{ϵ} , while the latter are in terms of $\mathcal{I}_{\Theta(\frac{\epsilon}{H})}$. We conjecture that our upper bounds can be improved by replacing \mathcal{I}_{ϵ} with $\mathcal{I}_{\Theta(\frac{\epsilon}{H})}$ —our analysis uses a clipping trick similar to [36], which may be the reason for a suboptimal dependence on H. We leave closing this gap as an open question.

5 Related Work

Regret minimization for MDPs. Our work belongs to the literature of regret minimization for MDPs, e.g., [5] [18] [8] [3] [9] [19] [10] [46] [36] [49] [45] [44]. In the episodic setting, [3] [10] [46] [36] [49] achieve minimax $\sqrt{H^2SAK}$ regret bounds for general stationary MDPs. Furthermore, the EULER algorithm [46] achieves adaptive problem-dependent regret guarantees when the total reward within an episode is small or when the environmental norm of the MDP is small. [36] refines EULER, proposing STRONG-EULER that provides more fine-grained gap-dependent $O(\log K)$ regret guarantees. [45] [44] show that the optimistic Q-learning algorithm [19] and its variants can also achieve gap-dependent logarithmic regret guarantees. Remarkably, [44] achieves a regret bound that improves over that of [36], in that it replaces the dependence on the number of optimal state-action pairs with the number of non-unique state-action pairs.

Transfer and lifelong learning for RL. A considerable portion of related works concerns transfer learning for RL tasks (see [40, 24, 50] for surveys from different angles), and many studies investigate a batch setting: given some source tasks and target tasks, transfer learning agents have access to batch data collected for the source tasks (and sometimes for the target tasks as well). In this setting, model-based approaches have been explored in e.g., [39]; theoretical guarantees for transfer of samples across tasks have been established in e.g., [25, 41]. Similarly, sequential transfer has been studied under the framework of lifelong RL in e.g., [38, 11, 15, 22]—in this setting, an agent faces a sequence of RL tasks and aims to take advantage of knowledge gained from previous tasks for better performance in future tasks; in particular, analyses on the sample complexity of transfer learning algorithms are presented in [6, 29] under the assumption that an upper bound on the total number of unique (and well-separated) RL tasks is known. We note that, in contrast, we study an online setting in which no prior data are available and multiple RL tasks are learned concurrently by RL agents.

Concurrent RL. Data sharing between multiple RL agents that learn concurrently has also been investigated in the literature. For example, in [20] 35, 16, 12, a group of agents interact in parallel with *identical* environments. Another setting is studied in [16], in which agents solve different RL tasks (MDPs); however, similar to 6 29, it is assumed that there is a finite number of unique tasks, and different tasks are well-separated, i.e., there is a minimum gap. In this work, we assume that players face similar but not necessarily identical MDPs, and we do not assume a minimum gap. $\boxed{17}$ study multi-task RL with linear function approximation with representation transfer, where it is assumed that the optimal value functions of all tasks are from a low dimensional linear subspace. Our setting and results are most similar to [32] and [13]. [32] study concurrent exploration in similar MDPs with continuous states in the PAC setting; however, their PAC guarantee does not hold for target error rate arbitrarily close to zero; in contrast, our algorithm has a fall-back guarantee, in that it always has a sublinear regret. Concurrent RL from similar linear MDPs has also been recently studied in 13: under the assumption of small heterogeneity between different MDPs (a setting very similar to ours), the provided regret guarantee involves a term that is linear in the number of episodes, whereas our algorithm in this paper always has a sublinear regret; concurrent RL under the assumption of large heterogeneity is also studied in that work, but additional contextual information is assumed to be available for the players to ensure a sublinear regret.

Other related topics and models. In many multi-agent RL models [47, 31], a set of learning agents interact with a common environment and have shared global states; in particular, [48] study the setting with heterogeneous reward distributions, and provide convergence guarantees for two policy gradient-based algorithms. In contrast, in our setting, our learning agents interact with separate environments. Multi-agent bandits with similar, heterogeneous reward distributions are investigated in [37, 43]; herein, we generalize their multi-task bandit problem setting to the episodic MDP setting.

6 Conclusion and Future Directions

In this paper, we generalize the multi-task bandit learning framework in [43] and formulate a multitask concurrent RL problem, in which tasks are similar but not necessarily identical. We provide a provably efficient model-based algorithm that takes advantage of knowledge transfer between different tasks. Our instance-dependent regret upper and lower bounds formalize the intuition that subpar state-action pairs are amenable to information sharing among tasks. There still remain gaps between our upper and lower bounds which can be closed by either a finer analysis or a better algorithm: first, the dependence on \mathcal{I}_{ϵ} in the upper bound does not match the dependence of $\mathcal{I}_{\Theta(\epsilon/H)}$ in the lower bound when H is large; second, the gap-dependent upper bound has $O(H^3)$ dependence, whereas the gap-dependent lower bound only has $\Omega(H^2)$ dependence; third, the additive dependence on the number of optimal state-action pairs can potentially be removed by new algorithmic ideas [44].

Furthermore, one major obstacle in deploying our algorithm in practice is its requirement for knowledge of ϵ ; an interesting avenue is to apply model selection strategies in bandits and RL to achieve adaptivity to unknown ϵ . Another interesting future direction is to consider more general parameter transfer for online RL, for example, in the context of function approximation.

7 Acknowledgements

We thank Kamalika Chaudhuri for helpful initial discussions, and thank Akshay Krishnamurthy and Tongyi Cao for discussing the applicability of adaptive RL in metric spaces to the multitask RL problem studied in this paper. CZ acknowledges startup funding support from the University of Arizona. ZW thanks the National Science Foundation under IIS 1915734 and CCF 1719133 for research support.

References

- [1] David Abel, Yuu Jinnai, Sophie Yue Guo, George Konidaris, and Michael Littman. Policy and value transfer in lifelong reinforcement learning. In *International Conference on Machine Learning*, pages 20–29. PMLR, 2018.
- [2] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Tuning bandit algorithms in stochastic environments. In *International conference on algorithmic learning theory*, pages 150–165. Springer, 2007.
- [3] Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pages 263–272. PMLR, 2017.
- [4] Peter Bartlett, Varsha Dani, Thomas Hayes, Sham Kakade, Alexander Rakhlin, and Ambuj Tewari. High-probability regret bounds for bandit online linear optimization. In *Proceedings* of the 21st Annual Conference on Learning Theory-COLT 2008, pages 335–342. Omnipress, 2008.
- [5] Peter L Bartlett and Ambuj Tewari. Regal: a regularization based algorithm for reinforcement learning in weakly communicating mdps. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 35–42, 2009.
- [6] Emma Brunskill and Lihong Li. Sample complexity of multi-task reinforcement learning. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, pages 122–131, 2013.
- [7] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- [8] Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. In Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 2, pages 2818–2826, 2015.
- [9] 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.
- [10] 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.
- [11] Carlo D'Eramo, Davide Tateo, Andrea Bonarini, Marcello Restelli, and Jan Peters. Sharing knowledge in multi-task deep reinforcement learning. In *International Conference on Learning Representations*, 2020.

- [12] Maria Dimakopoulou and Benjamin Van Roy. Coordinated exploration in concurrent reinforcement learning. In *International Conference on Machine Learning*, pages 1271–1279. PMLR, 2018.
- [13] Abhimanyu Dubey and Alex Pentland. Provably efficient cooperative multi-agent reinforcement learning with function approximation. arXiv preprint arXiv:2103.04972, 2021.
- [14] David A Freedman. On tail probabilities for martingales. *the Annals of Probability*, pages 100–118, 1975.
- [15] Francisco M Garcia and Philip S Thomas. A meta-mdp approach to exploration for lifelong reinforcement learning. arXiv preprint arXiv:1902.00843, 2019.
- [16] Zhaohan Guo and Emma Brunskill. Concurrent pac rl. In *Proceedings of the AAAI Conference* on Artificial Intelligence, volume 29, 2015.
- [17] Jiachen Hu, Xiaoyu Chen, Chi Jin, Lihong Li, and Liwei Wang. Near-optimal representation learning for linear bandits and linear rl. *arXiv preprint arXiv:2102.04132*, 2021.
- [18] Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- [19] Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? In Proceedings of the 32nd International Conference on Neural Information Processing Systems, pages 4868–4878, 2018.
- [20] R Matthew Kretchmar. Parallel reinforcement learning. In *The 6th World Conference on Systemics, Cybernetics, and Informatics.* Citeseer, 2002.
- [21] Alyssa Kubota, Emma IC Peterson, Vaishali Rajendren, Hadas Kress-Gazit, and Laurel D Riek. Jessie: Synthesizing social robot behaviors for personalized neurorehabilitation and beyond. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pages 121–130, 2020.
- [22] Nicholas C Landolfi, Garrett Thomas, and Tengyu Ma. A model-based approach for sampleefficient multi-task reinforcement learning. *arXiv preprint arXiv:1907.04964*, 2019.
- [23] Tor Lattimore and Csaba Szepesvári. Bandit algorithms. Cambridge University Press, 2020.
- [24] Alessandro Lazaric. Transfer in reinforcement learning: a framework and a survey. In *Reinforcement Learning*, pages 143–173. Springer, 2012.
- [25] Alessandro Lazaric and Marcello Restelli. Transfer from multiple mdps. In Advances in Neural Information Processing Systems, volume 24. Curran Associates, Inc., 2011.
- [26] Alessandro Lazaric, Marcello Restelli, and Andrea Bonarini. Transfer of samples in batch reinforcement learning. In *Proceedings of the 25th international conference on Machine learning*, pages 544–551, 2008.
- [27] Xinle Liang, Yang Liu, Tianjian Chen, Ming Liu, and Qiang Yang. Federated transfer reinforcement learning for autonomous driving. arXiv preprint arXiv:1910.06001, 2019.
- [28] Boyi Liu, Lujia Wang, and Ming Liu. Lifelong federated reinforcement learning: a learning architecture for navigation in cloud robotic systems. *IEEE Robotics and Automation Letters*, 4(4):4555–4562, 2019.
- [29] Yao Liu, Zhaohan Guo, and Emma Brunskill. Pac continuous state online multitask reinforcement learning with identification. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '16, page 438–446, 2016.
- [30] Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample variance penalization. *COLT*, 2009.
- [31] Afshin OroojlooyJadid and Davood Hajinezhad. A review of cooperative multi-agent deep reinforcement learning. *arXiv preprint arXiv:1908.03963*, 2019.
- [32] Jason Pazis and Ronald Parr. Efficient pac-optimal exploration in concurrent, continuous state mdps with delayed updates. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- [33] Michael T. Rosenstein, Zvika Marx, Leslie Pack Kaelbling, and Thomas G. Dietterich. To transfer or not to transfer. In *NIPS 2005 workshop on transfer learning*, 2005.

- [34] David Silver, Leonard Newnham, David Barker, Suzanne Weller, and Jason McFall. Concurrent reinforcement learning from customer interactions. In *International conference on machine learning*, pages 924–932. PMLR, 2013.
- [35] David Silver, Leonard Newnham, David Barker, Suzanne Weller, and Jason McFall. Concurrent reinforcement learning from customer interactions. In *International conference on machine learning*, pages 924–932. PMLR, 2013.
- [36] Max Simchowitz and Kevin Jamieson. Non-asymptotic gap-dependent regret bounds for tabular mdps. arXiv preprint arXiv:1905.03814, 2019.
- [37] Marta Soare, Ouais Alsharif, Alessandro Lazaric, and Joelle Pineau. Multi-task linear bandits. NIPS2014 Workshop on Transfer and Multi-task Learning : Theory meets Practice, 2014.
- [38] Fumihide Tanaka and Masayuki Yamamura. Multitask reinforcement learning on the distribution of mdps. In *Proceedings 2003 IEEE International Symposium on Computational Intelli*gence in Robotics and Automation, volume 3, pages 1108–1113. IEEE, 2003.
- [39] Matthew E Taylor, Nicholas K Jong, and Peter Stone. Transferring instances for model-based reinforcement learning. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 488–505. Springer, 2008.
- [40] Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(7), 2009.
- [41] Andrea Tirinzoni, Mattia Salvini, and Marcello Restelli. Transfer of samples in policy search via multiple importance sampling. In *International Conference on Machine Learning*, pages 6264–6274. PMLR, 2019.
- [42] Konstantinos Tsiakas, Cheryl Abellanoza, and Fillia Makedon. Interactive learning and adaptation for robot assisted therapy for people with dementia. In *Proceedings of the 9th ACM International Conference on PErvasive Technologies Related to Assistive Environments*, pages 1–4, 2016.
- [43] Zhi Wang, Chicheng Zhang, Manish Kumar Singh, Laurel Riek, and Kamalika Chaudhuri. Multitask bandit learning through heterogeneous feedback aggregation. In *International Conference on Artificial Intelligence and Statistics*, pages 1531–1539. PMLR, 2021.
- [44] Haike Xu, Tengyu Ma, and Simon S Du. Fine-grained gap-dependent bounds for tabular mdps via adaptive multi-step bootstrap. *arXiv preprint arXiv:2102.04692*, 2021.
- [45] Kunhe Yang, Lin Yang, and Simon Du. Q-learning with logarithmic regret. In International Conference on Artificial Intelligence and Statistics, pages 1576–1584. PMLR, 2021.
- [46] Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pages 7304–7312. PMLR, 2019.
- [47] Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. arXiv preprint arXiv:1911.10635, 2019.
- [48] 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, 2018.
- [49] Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learningvia reference-advantage decomposition. Advances in Neural Information Processing Systems, 33, 2020.
- [50] Zhuangdi Zhu, Kaixiang Lin, and Jiayu Zhou. Transfer learning in deep reinforcement learning: A survey. *arXiv preprint arXiv:2009.07888*, 2020.
- [51] Hankz Hankui Zhuo, Wenfeng Feng, Qian Xu, Qiang Yang, and Yufeng Lin. Federated reinforcement learning. *arXiv preprint arXiv:1901.08277*, 2019.

A Proofs of Lemmas 2 and 4

A.1 Proof of Lemma 2

Lemma 2. If $(\mathcal{M}_p)_{p=1}^M$ is ϵ -dissimilar, then for every $p, q \in [M]$, and $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\left|Q_{p}^{\star}(s,a) - Q_{q}^{\star}(s,a)\right| \leq 2H\epsilon$$

consequently, $\left| gap_p(s, a) - gap_q(s, a) \right| \le 4H\epsilon$.

Proof. For the first claim, we prove a stronger statement by backward induction on h, namely, for every $p, q \in [M]$, every $h \in [1, H + 1]$, and $(s, a) \in S_h \times A$,

$$\left|Q_p^{\star}(s,a) - Q_q^{\star}(s,a)\right| \le 2(H-h+1)\epsilon.$$

Base case: For h = H + 1, we have $Q_p^{\star}(s, a) = 0$ for every $(s, a) \in S_h \times A$, and $p \in [M]$. It follows trivially that $|Q_p^{\star}(s, a) - Q_q^{\star}(s, a)| = 0 \le 2(H - h + 1)\epsilon$.

Inductive case: Suppose by inductive hypothesis that for some $h \in [1, H]$ and, for every $(s, a) \in S_{h+1} \times A$ and $p, q \in [M]$, $\left|Q_p^{\star}(s, a) - Q_q^{\star}(s, a)\right| \leq 2(H - h)\epsilon$.

We first prove the following auxiliary statement: for every $s \in S_{h+1}$ and $p, q \in [M]$,

$$\left|V_{p}^{\star}(s) - V_{q}^{\star}(s)\right| \le 2(H - h)\epsilon.$$
(7)

Let $a_p = \operatorname{argmax}_{a \in \mathcal{A}} Q_p^{\star}(s, a)$ and $a_q = \operatorname{argmax}_{a \in \mathcal{A}} Q_q^{\star}(s, a)$. The above auxiliary statement can be easily proven by contradiction: without loss of generality, suppose that $V_p^{\star}(s) - V_q^{\star}(s) = Q_p^{\star}(s, a_p) - Q_q^{\star}(s, a_q) > 2(H - h)\epsilon$. Since $Q_q^{\star}(s, a_p) \geq Q_p^{\star}(s, a_p) - 2(H - h)\epsilon$, it follows that $Q_q^{\star}(s, a_p) > Q_q^{\star}(s, a_q)$, which contradicts the fact that $a_q = \operatorname{argmax}_{a \in \mathcal{A}} Q_q^{\star}(s, a)$.

We now return to the inductive proof, and we show that given the inductive hypothesis, for every $(s, a) \in S_h \times A$ and $p, q \in [M]$,

$$\begin{aligned} \left| Q_{p}^{\star}(s,a) - Q_{q}^{\star}(s,a) \right| \\ &\leq \left| R_{p}(s,a) - R_{q}(s,a) \right| + \left| \sum_{s' \in \mathcal{S}_{h+1}} \left[\mathbb{P}_{p}(s' \mid s,a) V_{p}^{\star}(s') - \mathbb{P}_{q}(s' \mid s,a) V_{q}^{\star}(s') \right] \right| \\ &\leq \epsilon + \left| \sum_{s' \in \mathcal{S}_{h+1}} \left[\mathbb{P}_{p}(s' \mid s,a) V_{p}^{\star}(s') - \mathbb{P}_{q}(s' \mid s,a) V_{p}^{\star}(s') \right] \right| + \left| \sum_{s' \in \mathcal{S}_{h+1}} \mathbb{P}_{q}(s' \mid s,a) \left(V_{p}^{\star}(s') - V_{q}^{\star}(s') \right) \right| \\ &\leq \epsilon + \left\| \mathbb{P}_{p}(\cdot \mid s,a) - \mathbb{P}_{q}(\cdot \mid s,a) \right\|_{1} \left(\max_{s' \in \mathcal{S}_{h+1}} \left| V_{p}^{\star}(s') \right| \right) + \left\| \mathbb{P}_{q}(\cdot \mid s,a) \right\|_{1} \left(\max_{s' \in \mathcal{S}_{h+1}} \left| V_{p}^{\star}(s') - V_{q}^{\star}(s') \right| \right) \\ &\leq \epsilon + \frac{\epsilon}{H} \cdot H + 2(H - h)\epsilon \\ &= 2(H - h + 1)\epsilon, \end{aligned}$$

where the first inequality follows from Eq. (1) and the triangle inequality; the second inequality follows from Definition 1 and the triangle inequality; the third inequality follows from Hölder's inequality; and the fourth inequality uses Definition 1 and Eq. (7).

For the second claim, we note that from the first claim, we have for any p, q, s,

$$\left|V_{p}^{\star}(s) - V_{q}^{\star}(s)\right| = \left|\max_{a \in \mathcal{A}} Q_{p}^{\star}(s, a) - \max_{a \in \mathcal{A}} Q_{p}^{\star}(s, a)\right| \le 2H\epsilon$$

therefore, for any p, q, s, a,

$$\left|\operatorname{gap}_p(s,a) - \operatorname{gap}_q(s,a)\right| \le \left|V_p^{\star}(s) - V_q^{\star}(s)\right| + \left|Q_p^{\star}(s,a) - Q_p^{\star}(s,a)\right| \le 4H\epsilon.$$

A.2 Proof of Lemma 4

Lemma 4 For any $(s, a) \in \mathcal{I}_{\epsilon}$, we have that: (1) for all $p \in [M]$, $(s, a) \notin Z_{p, \text{opt}}$, where we recall that $Z_{p, \text{opt}} = \left\{ (s, a) : \text{gap}_p(s, a) = 0 \right\}$ is the set of optimal state-action pairs with respect to p; (2) for all $p, q \in [M]$, $\text{gap}_p(s, a) \ge \frac{1}{2} \text{gap}_q(s, a)$.

Proof. For any $(s, a) \in \mathcal{I}_{\epsilon}$, there exists some p_0 such that $gap_{p_0}(s, a) > 96H\epsilon$. From Lemma we know that $|gap_p(s, a) - gap_{p_0}(s, a)| \le 4H\epsilon$. Therefore, for all p,

$$\operatorname{gap}_p(s, a) \ge \operatorname{gap}_{p_0}(s, a) - 4H\epsilon > 92H\epsilon \ge 0.$$

This proves the first item.

For the second item, for all $p, q \in [M]$,

$$\frac{\operatorname{gap}_p(s,a)}{\operatorname{gap}_q(s,a)} = \frac{\operatorname{gap}_q(s,a) - 4H\epsilon}{\operatorname{gap}_q(s,a)} \ge 1 - \frac{4H\epsilon}{\operatorname{gap}_q(s,a)} \ge 1 - \frac{4}{92} \ge \frac{1}{2}.$$

B Additional Definitions Used in the Proofs

In this section, we define a few useful notations that will be used in our proofs. For state-action pair $(s, a) \in S \times A$, player $p \in [M]$, episode $k \in [K]$:

- 1. Define $n^k(s, a)$ (resp. $n_p^k(s, a)$, $\hat{\mathbb{P}}^k$, $\hat{\mathbb{P}}_p^k$, \hat{R}^k , \hat{R}_p^k) to be the value of n(s, a) (resp. $n_p(s, a)$, $\hat{\mathbb{P}}, \hat{\mathbb{P}}_p, \hat{R}, \hat{R}_p$) at the *beginning* of episode k of MULTI-TASK-EULER.
- 2. Denote by \overline{Q}_p^k (resp. $\underline{Q}_p^k, \overline{V}_p^k, \underline{V}_p^k$, ind- $b_p^k(s, a)$, agg- $b_p^k(s, a)$) the values of \overline{Q}_p (resp. $\underline{Q}_p, \overline{V}_p, \underline{V}_p$, ind- $b_p(s, a)$, agg- $b_p(s, a)$) right after MULTI-TASK-EULER finishes its optimistic value iteration (line [5]) at episode k.
- 3. Define the *surplus* [36] (also known as the Bellman error) of (s, a) at episode k and player p as:

$$E_p^k(s,a) := \overline{Q}_p^k(s,a) - R_p(s,a) - (\mathbb{P}_p \overline{V}_p^k)(s,a)$$

4. Define $w_p^k(s,a) := \frac{n_p^k(s,a)}{n^k(s,a)}$ be the proportion of player p on (s,a) at the beginning of episode k; this induces (s,a)'s *mixture expected reward*:

$$\bar{R}^k(s,a) := \sum_{q=1}^M w_q^k(s,a) R_q(s,a),$$

and mixture transition probability:

$$\bar{\mathbb{P}}^k(\cdot \mid s, a) \coloneqq \sum_{q=1}^M w_q^k(s, a) \mathbb{P}_q(\cdot \mid s, a).$$

5. Define $\rho_p^k(s, a) := \mathbb{P}((s_h, a_h) = (s, a) \mid \pi^k(p), \mathcal{M}_p)$ to be the occupancy measure of $\pi^k(p)$ over \mathcal{M}_p on (s, a), where $h \in [H]$ is the layer s is in (so that $s \in S_h$). It can be seen that ρ_p^k , when restricted to $S_h \times \mathcal{A}$, is a probability distribution on this set.

Define
$$\rho^k(s, a) := \sum_{p=1}^{M} \rho_p^k(s, a)$$
; it can be seen that $\rho^k(s, a) \in [0, M]$. Define $\bar{n}_p^k(s, a) := \sum_{j=1}^{k} \rho_p^j(s, a)$, and $\bar{n}^k(s, a) := \sum_{j=1}^{k} \rho^j(s, a)$.

6. Define $N^k(s) := \sum_{a \in \mathcal{A}} n^k(s, a)$ and $N_p^k(s) := \sum_{a \in \mathcal{A}} n_p^k(s, a)$ to be the total number of encounters of state s by all players, and by player p only, respectively, at the beginning of episode k.

³These are the cumulative occupancy measures up to episode k, inclusively; this is in contrast with the definition of $n^k(s, a)$ and $n^k_p(s, a)$, which do not count the trajectories observed at episode k.

- 7. Define $N_1 \equiv M \ln(\frac{SAK}{\delta})$, and $N_2 \equiv \ln(\frac{MSAK}{\delta})$; define $\tau(s, a) := \min \left\{k : \bar{n}^k(s, a) \ge N_1\right\}$, and $\tau_p(s, a) := \min \left\{k : \bar{n}^k_p(s, a) \ge N_2\right\}$. With high probability, so long as $k \ge \tau(s, a)$ (resp. $k \ge \tau_p(s, a)$), $n^k(s, a)$ and $\bar{n}^k(s, a)$ (resp. $n^k_p(s, a)$) and $\bar{n}^k_p(s, a)$) are within a constant factor of each other; see Lemma [1]]
- 8. Define $g \check{a} p_p(s, a) := \frac{g a p_p(s, a)}{4H} \vee \frac{g a p_{p,\min}}{4H}$; recall the definitions of $g a p_p(s, a)$ and $g a p_{p,\min}$ in Section 2.

Define $\operatorname{Reg}(K, p) := \sum_{k=1}^{K} \left(V_{0,p}^{\star} - V_{0,p}^{\pi^{k}(p)} \right)$ as player *p*'s contribution to the collective regret; in this notation, $\operatorname{Reg}(K) = \sum_{p=1}^{M} \operatorname{Reg}(K, p)$.

Define the clipping function $\operatorname{clip}(\alpha, \Delta) := \alpha \mathbf{1}(\alpha \ge \Delta).$

We also adopt the following conventions in our proofs:

- 1. As ϵ -dissimilarity with $\epsilon > 2H$ does not impose any constraints on $\{\mathcal{M}_p\}_{p=1}^M$ (recall Definition 1), throughout the proof, we only focus on the regime that $\epsilon \leq 2H$.
- 2. We will use $\pi^k(p)$ and π^k_p interchangeably. To avoid notational clutter, we will also sometimes slightly abuse notation and use $V_{p,h}^{\pi^k}$, $V_p^{\pi^k}$ to denote $V_{p,h}^{\pi^k(p)}$, $V_p^{\pi^k(p)}$, respectively.

C Proof of the Upper Bounds

Proof outline. This section establishes the regret guarantees of MULTI-TASK-EULER (Theorems and **6**). The proof follows a similar outline as STRONG-EULER's analysis **[36**], with important modifications tailored to the multitask setting. The proof has the following structure:

- 1. Subsection C.1 defines a "clean" event E that we show happens with probability 1δ . When E happens, the observed samples are representative enough so that standard concentration inequalities apply. This will serve as the basis of our subsequent arguments.
- 2. Subsection C.2 shows that when *E* happens, the value function upper and lower bounds are valid; furthermore, MULTI-TASK-EULER satisfies strong optimism [36], in that all players' surpluses are always nonnegative for all state-action pairs at all time steps.
- 3. Subsection C.3 establishes a distribution-dependent upper bound on MULTI-TASK-EULER's surpluses when *E* happens, which is key to our regret theorems. In comparison with STRONG-EULER [36] in the single task setting, MULTI-TASK-EULER exploits inter-task similarity, so that its surpluses on state-action pair (s, a) for player *p* are further controlled by a new term that depends on the dissimilarity parameter ϵ , along with $n^k(s, a)$, the total visitation counts of (s, a) by all players.
- 4. Subsection C.4 uses the strong optimism property and the surplus bounds established in the previous two subsections to conclude our final gap-independent and gap-dependent regret guarantees, via the clipping lemma of [36] (see also Lemma [20]).
- 5. Finally, Subsection C.5 collects miscellaneous technical lemmas used in the proofs.

C.1 A clean event

Below we define a "clean" event E in which all concentration bounds used in the analysis hold, which we will show happens with high probability. Specifically, we will define $E = E_{ind} \cap E_{agg} \cap E_{sample}$, where $E_{ind}, E_{agg}, E_{sample}$ are defined respectively below.

In subsequent definitions of events, we will abbreviate $\forall k \in [K], h \in [H], p \in [M], s \in S_h, a \in A, s' \in S_{h+1}$ as $\forall k, h, p, s, a, s'$. Also, recall that $L(n) \approx \ln(\frac{MSAn}{\delta})$.

Define event E_{ind} as:

$$E_{\rm ind} = E_{\rm ind, rw} \cap E_{\rm ind, val} \cap E_{\rm ind, prob} \cap E_{\rm ind, var}, \tag{8}$$

$$E_{\text{ind,rw}} = \left\{ \forall k, h, p, s, a \cdot \left| \hat{R}_p^k(s, a) - R_p(s, a) \right| \le \sqrt{\frac{L(n^k(s, a))}{2n^k(s, a)}} \right\},\tag{9}$$

$$E_{\text{ind,val}} = \left\{ \forall k, h, p, s, a \cdot \left| (\hat{\mathbb{P}}_p^k V_p^\star - \mathbb{P}_p V_p^\star)(s, a) \right| \le 4\sqrt{\frac{\operatorname{var}_{\mathbb{P}_p}(\cdot | s, a) [V_p^\star] L(n_p^k(s, a))}{n_p^k(s, a)}} + \frac{2HL(n_p^k(s, a))}{n_p^k(s, a)} \right\},$$
(10)

$$E_{\text{ind,prob}} = \left\{ \forall k, h, p, s, a, s' \cdot \left| (\hat{\mathbb{P}}_p^k - \mathbb{P}_p)(s' \mid s, a) \right| \le 4\sqrt{\frac{L(n_p^k(s, a)) \cdot \mathbb{P}_p(s' \mid s, a)}{n_p^k(s, a)}} + \frac{2L(n_p^k(s, a))}{n_p^k(s, a)} \right\}$$
(11)

,

$$E_{\text{ind,var}} = \left\{ \forall k, h, p, s, a \cdot \left| \frac{1}{n_p^k(s, a)} \sum_{i=1}^{n_p^k(s, a)} (V_p^\star((s_i^p)') - (\mathbb{P}_p V_p^\star)(s, a))^2 - \operatorname{var}_{\mathbb{P}_p(\cdot|s, a)} [V_p^\star] \right|,$$

$$\leq 4\sqrt{\frac{H^2 \operatorname{var}_{\mathbb{P}_p(\cdot|s, a)} [V_p^\star] L(n_p^k(s, a))}{n_p^k(s, a)}} + \frac{2H^2 L(n_p^k(s, a))}{n_p^k(s, a)} \right\},$$
(12)

where in Equation (12), $(s_i^p)'$ denotes the next state player p transitions to, for the *i*-th episode it experiences (s, a). E_{ind} captures the concentration behavior of each player's individual model estimates.

Lemma 9. $\mathbb{P}(E_{\text{ind}}) \geq 1 - \frac{\delta}{3}$.

Proof. The proof follows a similar reasoning as the proof of e.g., [36] Proposition F.9] using Freedman's Inequality. We would like to show that each of $E_{\text{ind,rw}}$, $E_{\text{ind,val}}$, $E_{\text{ind,var}}$, $E_{\text{ind,var}}$ happens with probability $1 - \frac{\delta}{12}$, which would give the lemma statement by a union bound. For brevity, we only show that $\mathbb{P}(E_{\text{ind,var}}) \ge 1 - \frac{\delta}{12}$, and the other probability statements follow from a similar reasoning.

Fix $h \in [H]$, $(s, a) \in \mathcal{S}_h \times \mathcal{A}$, and $p \in [M]$. We will show

$$\mathbb{P}\left(\exists k \in [K] \cdot \left| \frac{1}{n_{p}^{k}(s,a)} \sum_{i=1}^{n_{p}^{k}(s,a)} (V_{p}^{\star}((s_{i}^{p})') - (\mathbb{P}_{p}V_{p}^{\star})(s,a))^{2} - \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\star}] \right|$$

$$\geq 4\sqrt{\frac{H^{2}\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\star}]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{2H^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \right) \leq \frac{\delta}{12MSA}.$$
(13)

For every $j \in \mathbb{N}_+$, define stopping time k_j as the *j*-th episode when (s, a) is experienced by player p, if such episode exists; otherwise, k_j is defined as ∞ . it suffices to show that

$$\mathbb{P}\left(\exists j \in \mathbb{N}_{+}, k_{j} < \infty \land \left|\frac{1}{j}\sum_{i=1}^{j}(V_{p}^{\star}((s_{i}^{p})') - (\mathbb{P}_{p}V_{p}^{\star})(s,a))^{2} - \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\star}]\right| \\ \geq 4\sqrt{\frac{H^{2}\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\star}]L(j)}{j}} + \frac{2H^{2}L(j)}{j}\right) \leq \frac{\delta}{12MSA}.$$
(14)

For every $k \in \mathbb{N}_+$, Define \mathcal{F}_{k-1} as the σ -field generated by all players' observations up to episode k-1, along with all players' observations at episode k up to them taking action at step h. Define

$$X_k := I((s_{h,p}^k, a_{h,p}^k) = (s,a)) \cdot \left((V_p^{\star}(s_{h+1,p}^k) - (\mathbb{P}_p V_p^{\star})(s,a))^2 - \operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\star}] \right),$$

it can be seen that X_k is \mathcal{F}_k -measurable, and $\mathbb{E}[X_k | \mathcal{F}_{k-1}] = 0$, i.e. $\{X_k\}_{k=1}^{\infty}$ is a martingale difference sequence adapted to $\{\mathcal{F}_k\}_{k=1}^{\infty}$. In addition,

$$\mathbb{E}\left[X_{k}^{2} \mid \mathcal{F}_{k-1}\right] \leq \mathbb{E}\left[I((s_{h,p}^{k}, a_{h,p}^{k}) = (s, a))(V_{p}^{\star}(s_{h+1,p}^{k}) - (\mathbb{P}_{p}V_{p}^{\star})(s, a))^{4} \mid \mathcal{F}_{k-1}\right] \\ \leq H^{2} \cdot I((s_{h,p}^{k}, a_{h,p}^{k}) = (s, a)) \cdot \operatorname{var}_{\mathbb{P}_{p}(\cdot \mid s, a)}[V_{p}^{\star}] =: U_{k};$$

Note that $|X_k/H^2| \le 1$; by [14, Corollary 1.4] applied to $\{X_k/H^2\}_{k=1}^{\infty}$, for any $\lambda \ge 0$,

$$\left\{Y_k(\lambda) = \exp\left(\lambda(\sum_{i=1}^k \frac{X_i}{H^2}) - \left((e^\lambda - \lambda - 1)\sum_{i=1}^k \frac{U_i}{H^4}\right)\right)\right\}_{k=0}^{\infty}$$

is a nonnegative supermartingale. Applying optional sampling theorem on $Y_k(\lambda)$ and stopping time k_j , we get $\mathbb{E}\left[Y_{k_j}(\lambda)I(k_j < \infty)\right] \leq \mathbb{E}\left[Y_0(\lambda)\right] = 1$. As a result, for any fixed thresholds $b, v \geq 0$ [?, see]Theorem 1.6]freedman1975tail,

$$\mathbb{P}\left(\sum_{i=1}^{k_j} X_i \ge b \land \sum_{i=1}^{k_j} U_i \le v \land k_j < \infty\right)$$
$$=\mathbb{P}\left(\sum_{i=1}^{k_j} \frac{X_i}{H^2} \ge \frac{b}{H^2} \land \sum_{i=1}^{k_j} \frac{U_i}{H^4} \le \frac{v}{H^4} \land k_j < \infty\right)$$
$$\le \exp\left(-\frac{(b/H^2)^2}{2(v/H^4) + 2(b/H^2)}\right) = \exp\left(-\frac{b^2}{2v + 2H^2b}\right)$$

Note that if $k_j < \infty$, $\sum_{i=1}^{k_j} X_i = \sum_{i=1}^j (V_p^{\star}((s_i^p)') - (\mathbb{P}_p V_p^{\star})(s, a))^2 - \operatorname{var}_{\mathbb{P}_p(\cdot|s, a)}[V_p^{\star}]$, and $\sum_{i=1}^{k_j} U_i = j \cdot H^2 \cdot \operatorname{var}_{\mathbb{P}_p(\cdot|s, a)}[V_p^{\star}]$, and the above inequality can be rewritten as: for any $b, v \ge 0$,

$$\mathbb{P}\left(\sum_{i=1}^{j} (V_p^{\star}((s_i^p)') - (\mathbb{P}_p V_p^{\star})(s, a))^2 - \operatorname{var}_{\mathbb{P}_p(\cdot|s, a)}[V_p^{\star}] \ge b \wedge j \cdot H^2 \cdot \operatorname{var}_{\mathbb{P}_p(\cdot|s, a)}[V_p^{\star}] \le v \wedge k_j < \infty\right)$$
$$\le \exp\left(-\frac{b^2}{2v + 2H^2b}\right).$$

Now, by the doubling argument of [4, Lemma 2] (observe that $j \cdot H^2 \cdot \operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\star}] \in [0, H^4j]$), we have that for all $j \in \mathbb{N}_+$:

$$\mathbb{P}\left(k_j < \infty \land \left|\frac{1}{j}\sum_{i=1}^{j} (V_p^{\star}((s_i^p)') - (\mathbb{P}_p V_p^{\star})(s,a))^2 - \operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\star}]\right|$$
$$\geq 4\sqrt{\frac{H^2 \operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\star}]L(j)}{j}} + \frac{2H^2L(j)}{j}\right) \leq \ln(4j) \cdot \frac{\delta}{48j^2 MSA}.$$

A union bound over all $j \in \mathbb{N}_+$ yields Equation (14).

Define event E_{agg} as:

$$E_{\text{agg}} = E_{\text{agg,rw}} \cap E_{\text{agg,val}} \cap E_{\text{agg,prob}} \cap E_{\text{agg,var}},\tag{15}$$

$$E_{\text{agg,rw}} = \left\{ \forall k, h, s, a \cdot \left| \hat{R}^k(s, a) - \bar{R}^k(s, a) \right| \le \sqrt{\frac{L(n^k(s, a))}{2n^k(s, a)}} \right\},\tag{16}$$

$$E_{\text{agg,val}} = \left\{ \forall k, h, p, s, a \cdot \left| (\hat{\mathbb{P}}^k V_p^\star - \bar{\mathbb{P}}^k V_p^\star)(s, a) \right|,$$
(17)

$$\leq 4\sqrt{\frac{\left(\sum_{q=1}^{M} w_{q}^{k}(s,a) \operatorname{var}_{\mathbb{P}_{q}(\cdot|s,a)}[V_{p}^{\star}]\right) L(n^{k}(s,a))}{n^{k}(s,a)} + \frac{2HL(n^{k}(s,a))}{n^{k}(s,a)}}\right\},$$
(18)

$$E_{\text{agg,prob}} = \left\{ \forall k, h, s, a, s' \cdot \left| (\hat{\mathbb{P}}^k - \bar{\mathbb{P}}^k)(s' \mid s, a) \right| \le 4\sqrt{\frac{\bar{\mathbb{P}}^k(s' \mid s, a) \cdot L(n^k(s, a))}{n^k(s, a)}} + \frac{2L(n^k(s, a))}{n^k(s, a)} \right\}$$
(19)

$$E_{\text{agg,var}} = \left\{ \forall k, h, p, s, a \cdot \left| \frac{1}{n^k(s, a)} \sum_{i=1}^{n^k(s, a)} (V_p^{\star}(s_i') - (\mathbb{P}_{p_i} V_p^{\star})(s, a))^2 - \sum_{q=1}^M w_q^k(s, a) \operatorname{var}_{\mathbb{P}_q(\cdot | s, a)} [V_p^{\star}] \right|$$
(20)

$$\leq 4\sqrt{\frac{H^2\left(\sum_{q=1}^{M} w_q^k(s, a) \mathrm{var}_{\mathbb{P}_q(\cdot | s, a)}[V_p^{\star}]\right) L(n^k(s, a))}{n^k(s, a)}} + \frac{2H^2 L(n^k(s, a))}{n^k(s, a)}}\right\},$$

where in Equation (20), s'_i and p_i denote the next state and the player index for the *i*-th time some player experiences (s, a), respectively, where within an episode, we order the experiences of the players by their indices from 1 to M. E_{agg} captures the concentration behavior of the aggregate model estimates.

Lemma 10. $\mathbb{P}(E_{\text{agg}}) \geq 1 - \frac{\delta}{3}$.

Proof. The proof follows a similar reasoning as the proof of e.g., [36] Proposition F.9] using Freedman's Inequality. We would like to show that each of $E_{agg,rw}$, $E_{agg,val}$, $E_{agg,prob}$, $E_{agg,var}$ happen with probability $1 - \frac{\delta}{12}$, which would give the lemma statement by a union bound. For brevity, we show that $\mathbb{P}(E_{agg,var}) \geq 1 - \frac{\delta}{12}$, and the other probability statements follow from a similar reasoning.

Fix $h \in [H]$, $(s, a) \in S_h \times A$ and $p \in [M]$. It suffices to show that

$$\mathbb{P}\left(\exists k \in [K] \cdot \left| \frac{1}{n^{k}(s,a)} \sum_{i=1}^{n^{k}(s,a)} \left((V_{p}^{\star}(s_{i}^{\prime}) - (\mathbb{P}_{p_{i}}V_{p}^{\star})(s,a))^{2} - \operatorname{var}_{\mathbb{P}_{p_{i}}(\cdot|s,a)}[V_{p}^{\star}] \right) \right| \\
\geq 4\sqrt{\frac{H^{2}\left(\sum_{i=1}^{n^{k}(s,a)} \operatorname{var}_{\mathbb{P}_{p_{i}}(\cdot|s,a)}[V_{p}^{\star}]\right) L(n^{k}(s,a))}{(n^{k}(s,a))^{2}}} + \frac{2H^{2}L(n^{k}(s,a))}{n^{k}(s,a)}}{\int ds} \leq \frac{\delta}{12MSA},$$
(21)

because $\frac{1}{n^k(s,a)} \sum_{i=1}^{n^k(s,a)} \operatorname{var}_{\mathbb{P}_{p_i}(\cdot|s,a)}[V_p^\star] = \sum_{q=1}^M w_q^k(s,a) \operatorname{var}_{\mathbb{P}_q(\cdot|s,a)}[V_p^\star].$

Define a *micro-episode* as an (episode, player) pair; we order them lexicographically, i.e. $(1,1) \leq \ldots \leq (1,M) \leq \ldots \leq (K,1) \leq \ldots \leq (K,M)$. For micro-episode (k,p), denote its *index* as l = (k-1)M + p; it can be easily seen that the ordering of micro-episodes' indices is consistent with their lexical ordering. For every $j \in \mathbb{N}_+$, define stopping time $l_j \in \mathbb{N}_+$ as follows: it is the

index of the *j*-th micro-episode when (s, a) is experienced by some player, if such micro-episode exists; and l_j is defined to be ∞ otherwise. With this notation, it suffices to show:

$$\mathbb{P}\left(\exists j \in \mathbb{N}_{+}, l_{j} < \infty \land \left| \frac{1}{j} \sum_{i=1}^{j} \left((V_{p}^{\star}(s_{i}^{\prime}) - (\mathbb{P}_{p_{i}}V_{p}^{\star})(s, a))^{2} - \operatorname{var}_{\mathbb{P}_{p_{i}}(\cdot|s, a)}[V_{p}^{\star}] \right) \right|$$

$$\geq 4\sqrt{\frac{H^{2}\left(\sum_{i=1}^{j} \operatorname{var}_{\mathbb{P}_{p_{i}}(\cdot|s, a)}[V_{p}^{\star}]\right) L(j)}{j^{2}}} + \frac{2H^{2}L(j)}{j} \right) \leq \frac{\delta}{12MSA},$$

$$(22)$$

For every $l \in \mathbb{N}_+$, Define \mathcal{F}_{l-1} as the σ -field generated by all players' observations up to microepisode l-1, along with micro-episode l's corresponding player (player index $((l-1) \mod M) + 1$)'s observations up to them taking action at step h. Define

$$X_{l} := I((s_{h,p}^{k}, a_{h,p}^{k}) = (s, a)) \cdot \left((V_{p}^{\star}(s_{h+1,p}^{k}) - (\mathbb{P}_{p}V_{p}^{\star})(s, a))^{2} - \operatorname{var}_{\mathbb{P}_{p}(\cdot|s, a)}[V_{p}^{\star}] \right),$$

where in the above expression, to avoid notation clutter, we use k and p to denote microepisode l's episode number and corresponding player number $k(l) = \lceil l/M \rceil$ and $p(l) = ((l-1) \mod M) + 1$, respectively.

It can be seen that X_l is \mathcal{F}_l -measurable, and $\mathbb{E}[X_l | \mathcal{F}_{l-1}] = 0$, i.e. $\{X_l\}_{l=1}^{\infty}$ is a martingale difference sequence adapted to $\{\mathcal{F}_l\}_{l=1}^{\infty}$. In addition,

$$\mathbb{E}\left[X_{l}^{2} \mid \mathcal{F}_{k-1}\right] \leq \mathbb{E}\left[I((s_{h,p}^{k}, a_{h,p}^{k}) = (s, a))(V_{p}^{\star}(s_{h+1,p}^{k}) - (\mathbb{P}_{p}V_{p}^{\star})(s, a))^{4} \mid \mathcal{F}_{k-1}\right]$$

$$\leq H^{2} \cdot I((s_{h,p}^{k}, a_{h,p}^{k}) = (s, a)) \cdot \operatorname{var}_{\mathbb{P}_{p}(\cdot \mid s, a)}[V_{p}^{\star}] =: U_{l};$$

Note that $|X_l/H^2| \le 1$; by [14, Corollary 1.4] applied to $\{X_l/H^2\}_{l=1}^{\infty}$, for any $\lambda \ge 0$,

$$\left\{Y_k(\lambda) = \exp\left(\lambda(\sum_{i=1}^l \frac{X_i}{H^2}) - \left((e^\lambda - \lambda - 1)\sum_{i=1}^l \frac{U_i}{H^4}\right)\right)\right\}_{l=0}^{\infty}$$

is a nonnegative supermartingale. Also, note that if $l_j < \infty$, $\sum_{i=1}^{l_j} X_i = \sum_{i=1}^j (V_p^{\star}(s'_i) - (\mathbb{P}_{p_i}V_p^{\star})(s,a))^2 - \operatorname{var}_{\mathbb{P}_{p_i}(\cdot|s,a)}[V_p^{\star}]$, and $\sum_{i=1}^{l_j} U_i = \sum_{i=1}^j H^2 \cdot \operatorname{var}_{\mathbb{P}_{p_i}(\cdot|s,a)}[V_p^{\star}]$.

Using the same reasoning as in the proof of Lemma (and observing that $\sum_{l=1}^{l_j} U_l \in [0, H^4 j]$), we have that for all $j \in \mathbb{N}_+$:

$$\mathbb{P}\left(k_{j} < \infty \land \left|\frac{1}{j}\sum_{i=1}^{j}\left((V_{p}^{\star}(s_{i}') - (\mathbb{P}_{p_{i}}V_{p}^{\star})(s,a))^{2} - \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\star}]\right)\right| \\
\geq 4\sqrt{\frac{H\sum_{i=1}^{j}\operatorname{var}_{\mathbb{P}_{p_{i}}(\cdot|s,a)}[V_{p}^{\star}]L(j)}{j^{2}}} + \frac{2H^{2}L(j)}{j}\right) \leq \ln(4j) \cdot \frac{\delta}{48j^{2}MSA}.$$

A union bound over all $j \in \mathbb{N}_+$ implies that Equation (22) holds.

Define

1

$$\begin{split} E_{\text{sample}} &= E_{\text{ind,sample}} \cap E_{\text{agg,sample}}, \\ E_{\text{agg,sample}} &= \left\{ \forall s, a, k \cdot \bar{n}^k(s, a) \ge N_1 \implies n^k(s, a) \ge \frac{1}{2} \bar{n}^k(s, a) \right\}, \\ E_{\text{ind,sample}} &= \left\{ \forall s, a, k, p \cdot \bar{n}_p^k(s, a) \ge N_2 \implies n_p^k(s, a) \ge \frac{1}{2} \bar{n}_p^k(s, a) \right\}, \\ \text{call from Section Pleter } N = M \ln \binom{SAK}{2} \text{ and } N = \ln \binom{MSAK}{2} \end{split}$$

where we recall from Section **B** that $N_1 \approx M \ln(\frac{SAK}{\delta})$, and $N_2 \approx \ln(\frac{MSAK}{\delta})$.

Lemma 11. $\mathbb{P}(E_{\text{sample}}) \geq 1 - \frac{\delta}{3}$.

Proof. We first show $\mathbb{P}(E_{\text{agg,sample}}) \geq 1 - \frac{\delta}{6}$. Specifically, fix $h \in [H]$ and $(s, a) \in \mathcal{S}_h \times \mathcal{A}$, define random variable $X_k = \sum_{p=1}^M \left(\mathbf{1} \left((s_{h,p}^k, a_{h,p}^k) = (s, a) \right) - \rho_p^k(s, a) \right)$. Also, define \mathcal{G}_k as the σ -algebra generated by all observations up to episode k. It can be readily seen that $\{X_k\}_{k=1}^K$ is a martingale difference sequence adapted to filtration $\{\mathcal{G}_k\}_{k=0}^K$. Freedman's inequality (specifically, Lemma 2 of [4]) implies that for every fixed k, with probability $1 - \frac{\delta}{6K}$,

$$\left| n^{k}(s,a) - \bar{n}^{k-1}(s,a) \right| \leq 4\sqrt{\bar{n}^{k-1}(s,a) \cdot M \ln\left(\frac{6SAK^{2}}{\delta}\right)} + 4M \ln\left(\frac{6SAK^{2}}{\delta}\right), \quad (23)$$

If Equation (23) happens, then by AM-GM inequality that $\sqrt{\bar{n}^{k-1}(s,a)} \cdot M \ln\left(\frac{6SAK^2}{\delta}\right) \leq \frac{1}{4}\bar{n}^{k-1}(s,a) + 16M \ln\left(\frac{6SAK^2}{\delta}\right)$, we have

$$\bar{n}^{k-1}(s,a) - n^k(s,a) \le \frac{1}{4}\bar{n}^{k-1}(s,a) + 20M\ln\left(\frac{6SAK^2}{\delta}\right),$$

implying that

$$n^{k}(s,a) \ge \frac{3}{4}\bar{n}^{k-1}(s,a) - 20M\ln\left(\frac{6SAK^{2}}{\delta}\right).$$

Additionally, as $\bar{n}^{k-1}(s,a) \geq \bar{n}^k(s,a) - M$ always holds, we have

$$n^{k}(s,a) \geq \frac{3}{4}\bar{n}^{k}(s,a) - 21M\ln\left(\frac{6SAK^{2}}{\delta}\right).$$

In summary, for any fixed k, with probability $1 - \frac{\delta}{6K}$, if $\bar{n}^k(s, a) \ge N_1 := 84M \ln\left(\frac{6SAK^2}{\delta}\right)$,

$$n^k(s,a) \ge \frac{1}{2}\bar{n}^k(s,a).$$

Taking a union bound over all $k \in [K]$, we have $\mathbb{P}(E_{\text{agg,sample}}) \geq 1 - \frac{\delta}{6}$.

It follows similarly that $\mathbb{P}(E_{\text{ind,sample}}) \geq 1 - \frac{\delta}{6}$; the only difference in the proof is that, we need to take an extra union bound over all $p \in [M]$ - hence an additional factor of M within $\ln(\cdot)$ in the definition of N_2 . The lemma statement follows from a union bound over these two statements. \Box

Lemma 12. $\mathbb{P}(E) \geq 1 - \delta$.

Proof. Follows from Lemmas 9, 10, and 11, along with a union bound.

C.2 Validity of value function bounds

In this section, we show that if the clean event E happens, then for all k and p, the value function estimates $\overline{Q}_p^k, \overline{Q}_p^k, \overline{V}_p^k, \underline{V}_p^k$ are valid upper and lower bounds of the optimal value functions Q_p^\star, V_p^\star (Lemma 13). As a by-product, we also give a general bound on the surplus (Lemma 14) which will be refined and used in the subsequent regret bound calculations. Before going into the proof of the above two lemmas, we need a technical lemma below (Lemma 13) that gives necessary concentration results which motivate the bonus constructions; its proof can be found at Section C.2.1.

Lemma 13. Fix $p \in [M]$. Suppose E happens, and suppose that for episode k and step h, we have that for all $s' \in S_{h+1}$, $\underline{V}_p^k(s') \leq V_p^*(s') \leq \overline{V}_p^k(s')$. Then, for all $(s, a) \in S_h \times A$:

1.

$$\left|\hat{R}_p^k(s,a) - R_p(s,a)\right| \le b_{\rm rw}\left(n_p^k(s,a),0\right),\tag{24}$$

$$\left|\hat{R}^{k}(s,a) - R_{p}(s,a)\right| \le b_{\rm rw}\left(n^{k}(s,a),\epsilon\right).$$
(25)

2.

$$\left| (\hat{\mathbb{P}}_{p}^{k} - \mathbb{P}_{p})(V_{p}^{\star})(s, a) \right| \leq b_{\text{prob}} \left(\hat{\mathbb{P}}_{p}^{k}(\cdot \mid s, a), n_{p}^{k}(s, a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, 0 \right),$$

$$\left| (\hat{\mathbb{P}}^{k} - \mathbb{P}_{p})(V_{p}^{\star})(s, a) \right| \leq b_{\text{prob}} \left(\hat{\mathbb{P}}^{k}(\cdot \mid s, a), n^{k}(s, a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, \epsilon \right).$$

$$(26)$$

$$\left| \hat{\mathbb{P}}^{k} - \mathbb{P}_{p} \right| (V_{p}^{\star})(s, a) \right| \leq b_{\text{prob}} \left(\hat{\mathbb{P}}^{k}(\cdot \mid s, a), n^{k}(s, a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, \epsilon \right).$$

$$(27)$$

3. For any $V_1, V_2 : S_{h+1} \to \mathbb{R}$ such that $\underline{V}_p^k \le V_1 \le V_2 \le \overline{V}_p^k$,

$$\left| \left(\hat{\mathbb{P}}_p^k - \mathbb{P}_p \right) (V_2 - V_1)(s, a) \right| \le b_{\text{str}} \left(\hat{\mathbb{P}}_p^k (\cdot \mid s, a), n_p^k(s, a), \overline{V}_p^k, \underline{V}_p^k, 0 \right),$$
(28)

$$\left| \left(\hat{\mathbb{P}}^k - \mathbb{P}_p \right) (V_2 - V_1)(s, a) \right| \le b_{\text{str}} \left(\hat{\mathbb{P}}^k (\cdot \mid s, a), n^k(s, a), \overline{V}_p^k, \underline{V}_p^k, \epsilon \right).$$
(29)

Lemma 14. If event E happens, and suppose that for episode k and step h, we have that for all $s' \in \mathcal{S}_{h+1}, \ \underline{V}_p^{\star}(s') \leq V_p^{\star}(s') \leq \overline{V}_p^{\star}(s').$ Then, for $(s, a) \in \mathcal{S}_h \times \mathcal{A}$,

$$\overline{Q}_p^k(s,a) - \left(R_p(s,a) + (\mathbb{P}_p \overline{V}_p^k)(s,a)\right) \in \left[0, (H-h+1) \wedge 2ind \cdot b_p^k(s,a) \wedge 2agg \cdot b_p^k(s,a)\right],$$
(30)

and

$$\left(R_p(s,a) + (\mathbb{P}_p\underline{V}_p^k)(s,a)\right) - \underline{Q}_p^k(s,a) \in \left[0, (H-h+1) \wedge 2ind \cdot b_p^k(s,a) \wedge 2agg \cdot b_p^k(s,a)\right],\tag{31}$$

where we recall that

$$ind-b_{p}^{k}(s,a) = b_{\mathrm{rw}}\left(n_{p}^{k}(s,a),0\right) + b_{\mathrm{prob}}\left(\hat{\mathbb{P}}_{p}^{k}(\cdot \mid s,a), n_{p}^{k}(s,a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, 0\right) + b_{\mathrm{str}}\left(\hat{\mathbb{P}}_{p}^{k}(\cdot \mid s,a), n_{p}^{k}(s,a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, 0\right),$$

$$agg-b_{p}^{k}(s,a) = b_{\mathrm{rw}}\left(n^{k}(s,a),\epsilon\right) + b_{\mathrm{prob}}\left(\hat{\mathbb{P}}^{k}(\cdot \mid s,a), n^{k}(s,a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, \epsilon\right) + b_{\mathrm{str}}\left(\hat{\mathbb{P}}^{k}(\cdot \mid s,a), n^{k}(s,a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, \epsilon\right).$$

Proof. We only show Equation (30) for brevity; Equation (31) follows from an exact symmetrical reasoning.

 $\text{Recall that } \overline{Q}_p^k(s,a) = \min \Big(\overline{\text{ind-}Q}_p^k(s,a), \overline{\text{agg-}Q}_p^k(s,a), H \Big). \text{ We compare each term in the } \min(\cdot)$ operator with $(R_p(s, a) + (\mathbb{P}_p \overline{V}_p^k)(s, a))$:

• For $\overline{\operatorname{ind}-Q_p^k}(s,a)$, using Lemma \square and our assumption on \overline{V}_p^k and \underline{V}_p^k on \mathcal{S}_{h+1} , we have:

$$\begin{aligned} \overline{\mathrm{ind}} &-\overline{Q}_p^k(s,a) - \left(R_p(s,a) + (\mathbb{P}_p \overline{V}_p^k)(s,a)\right) \\ &= \left(\hat{R}_p^k - R_p\right)(s,a) + b_{\mathrm{rw}}\left(n_p^k(s,a),0\right) \\ &+ \left(\left(\hat{\mathbb{P}}_p^k - \mathbb{P}_p\right)V_p^\star\right)(s,a) + b_{\mathrm{prob}}\left(\hat{\mathbb{P}}_p^k(\cdot \mid s,a), n_p^k(s,a), \overline{V}_p^k, \underline{V}_p^k, 0\right) \\ &+ \left(\hat{\mathbb{P}}_p^k - \mathbb{P}_p\right)(\overline{V}_p^k - V_p^\star)(s,a) + b_{\mathrm{str}}\left(\hat{\mathbb{P}}_p^k(\cdot \mid s,a), n_p^k(s,a), \overline{V}_p^k, \underline{V}_p^k, 0\right) \\ &\in [0, 2\mathrm{ind} \cdot b_p^k(s,a)]. \end{aligned}$$

• For $\overline{\text{agg-}Q}_p^k(s, a)$, using Lemma [3] and our assumptions on \overline{V}_p^k and \underline{V}_p^k over S_{h+1} , we have:

$$\begin{aligned} \overline{\operatorname{agg-}Q_p^k}(s,a) &- \left(R_p(s,a) + (\mathbb{P}_p \overline{V}_p^k)(s,a)\right) \\ &= (\hat{R}_p^k - R_p)(s,a) + b_{\operatorname{rw}}\left(n^k(s,a),\epsilon\right) \\ &+ ((\hat{\mathbb{P}}^k - \mathbb{P}_p)V_p^\star)(s,a) + b_{\operatorname{prob}}\left(\hat{\mathbb{P}}^k(\cdot \mid s,a), n^k(s,a), \overline{V}_p^k, \underline{V}_p^k, \epsilon\right) \\ &+ ((\hat{\mathbb{P}}^k - \mathbb{P}_p)(\overline{V}_p^k - V_p^\star))(s,a) + b_{\operatorname{str}}\left(\hat{\mathbb{P}}^k(\cdot \mid s,a), n^k(s,a), \overline{V}_p^k, \underline{V}_p^k, \epsilon\right) \\ &\in [0, 2\operatorname{agg-}b_p^k(s,a)], \end{aligned}$$

• For H - h + 1, we have:

$$(H - h + 1) - (R_p(s, a) + (\mathbb{P}_p \overline{V}_p^k)(s, a)) \in [0, H - h + 1],$$

where we use the observation that $R(s,a) \in [0,1]$, and $(\mathbb{P}_p \overline{V}_p^k)(s,a) \in [0, H-h]$, and their sum is in [0, H].

Combining the above three establishes that

$$\overline{Q}_p^k(s,a) - \left(R(s,a) + (\mathbb{P}_p \overline{V}_p^k)(s,a)\right) \in \left[0, (H-h+1) \land 2\mathrm{ind} \cdot b_p^k(s,a) \land 2\mathrm{agg} \cdot b_p^k(s,a)\right]. \quad \Box$$

Lemma 15. Under event E, for every $k \in [K]$, and every $p \in [M]$, and for every $h \in [H]$, For all $(s, a) \in S_h \times A$,

$$\underline{Q}_{p}^{k}(s,a) \leq Q_{p}^{\pi^{k}}(s,a) \leq Q_{p}^{\star}(s,a) \leq \overline{Q}_{p}^{k}(s,a),$$
(32)

and

$$\underline{V}_{p}^{k}(s) \leq V_{p}^{\pi^{k}}(s) \leq V_{p}^{\star}(s) \leq \overline{V}_{p}^{k}(s),$$
(33)

here, recall that $V_p^{\pi_k}$ is the value function of policy $\pi^k(p)$ with respect to \mathcal{M}_p defined in Section 2

Proof. The proof of this lemma extends [36]. Proposition F.1] to our multitask setting.

For every k and p, we show the above holds for all layers $h \in [H]$ and every $(s, a) \in S_h \times A$; to this end, we do backward induction on layer h.

Base case: For layer
$$h = H + 1$$
, we have $\underline{V}_p^k(\bot) = V_p^{\pi^k}(\bot) = V_p^{\star}(\bot) = \overline{V}_p^k(\bot) = 0$.

Inductive case: By our inductive hypothesis, for layer h + 1 and every $s \in S_{h+1}$,

$$\underline{V}_p^k(s) \le V_p^{\pi^k}(s) \le V_p^{\star}(s) \le \overline{V}_p^k(s)$$

We will show that Equations (32) and (33) holds holds for all $(s, a) \in S_h \times A$.

We first show Equation (32). First, $Q_p^{\pi^k}(s, a) \leq Q_p^{\star}(s, a)$ for all $(s, a) \in \mathcal{S}_h \times \mathcal{A}$ is trivial.

To show $Q_p^{\star}(s, a) \leq \overline{Q}_p^k(s, a)$ for all $(s, a) \in S_h \times A$, by Lemma 14 and inductive hypothesis, we have:

$$Q_p^{\star}(s,a) = R_p(s,a) + (\mathbb{P}_p V_p^{\star})(s,a) \le R_p(s,a) + (\mathbb{P}_p \overline{V}_p^{\star})(s,a) \le \overline{Q}_p^{\star}(s,a)$$

Likewise, we show $Q_p^{\pi^k}(s, a) \ge \underline{Q}_p^k(s, a)$ for all $(s, a) \in S_h \times A$, using Lemma 14 and inductive hypothesis:

$$Q_p^{\pi^k}(s,a) = R_p(s,a) + (\mathbb{P}_p V_p^{\pi^k})(s,a) \ge R_p(s,a) + (\mathbb{P}_p \overline{V}_p^k)(s,a) \ge \underline{Q}_p^k(s,a).$$

This completes the proof of Equation (32) for layer h.

We now show Equation (33) for layer h. Again $V_p^{\pi^k}(s) \leq V_p^{\star}(s)$ for all $s \in \mathcal{S}_h$ is trivial.

To show $V_p^{\star}(s) \leq \overline{V}_p^k(s)$ for all $s \in \mathcal{S}_h$, observe that

$$V_p^{\star}(s) = \max_{a \in \mathcal{A}} Q_p^{\star}(s, a) \le \max_{a \in \mathcal{A}} \overline{Q}_p^k(s, a) = \overline{V}_p^k(s).$$

To show $V_p^{\pi^k}(s) \geq \underline{V}_p^k(s)$ for all $s \in \mathcal{S}_h$, observe that

$$V_p^{\pi^k}(s) = Q_p^{\pi^k}(s, \pi^k(p)(s)) \ge \underline{Q}_p^k(s, \pi^k(p)(s)) = \underline{V}_p^k(s).$$

This completes the induction.

C.2.1 Proof of Lemma 13

Proof of Lemma [73] Equations (24), (26), and (28) essentially follow the same reasoning as in [36]; we still include their proofs for completeness. Equations (25), (27), and (29) are new, and require a more involved analysis. Our proof also relies on a technical lemma, namely Lemma [16]; we defer its statement and proof to the end of this subsection.

- 1. Equation (24) follows directly from the definition of $E_{\text{ind,rw}}$. Equation (25) follows from the definition of $E_{\text{agg,rw}}$, and the fact that $|\bar{R}^k(s,a) R_p(s,a)| \leq \epsilon$.
- 2. We prove Equation (26) as follows:

$$\begin{split} & \left| (\hat{\mathbb{P}}_{p}^{k} V^{\star} - \mathbb{P}_{p} V_{p}^{\star})(s, a) \right| \\ \leq O\left(\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V^{\star}]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{HL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \right) \\ \leq O\left(\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}^{k}(\cdot|s,a)}[V^{\star}]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{HL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \right) \\ \leq O\left(\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}^{k}(\cdot|s,a)}[\overline{V}_{p}^{k}]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{\|V_{p}^{\star} - \overline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{HL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \right) \\ \leq O\left(\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}^{k}(\cdot|s,a)}[\overline{V}_{p}^{k}]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{HL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \right) \\ \leq O\left(\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}^{k}(\cdot|s,a)}[\overline{V}_{p}^{k}]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}} \right) \\ \leq O_{\text{prob}}\left(\hat{\mathbb{P}}_{p}^{k}(\cdot|s,a), n_{p}^{k}(s,a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, 0 \right), \end{split}$$

where the first inequality is from the definition of $E_{\text{ind,val}}$; the second inequality is from Equation (34) of Lemma 16; the third inequality is from Lemma 24; the fourth inequality is from our assumption that for all $s' \in S_{h+1}$, $\underline{V}_p^k(s') \leq V^*(s') \leq \overline{V}_p^k(s')$, and thus $\left| (V_p^* - \underline{V}_p^k)(s') \right| \leq \left| (\overline{V}_p^k - \underline{V}_p^k)(s') \right|$ for all s' in the support of $\hat{\mathbb{P}}_p^k(\cdot \mid s, a)$.

We prove Equation (27) as follows:

$$\begin{split} & \left| (\hat{\mathbb{P}}^k - \mathbb{P}_p)(V_p^*)(s, a) \right| \\ \leq \epsilon + \left| (\hat{\mathbb{P}}^k - \bar{\mathbb{P}}^k)(V_p^*)(s, a) \right| \\ \leq \epsilon + O\left(\sqrt{\frac{\left(\sum_{q=1}^M w_q^k(s, a) \operatorname{var}_{\mathbb{P}_q(\cdot | s, a)}[V_p^*] \right) L(n^k(s, a))}{n^k(s, a)}} + \frac{HL(n^k(s, a))}{n^k(s, a)} \right) \\ \leq \epsilon + O\left(\sqrt{\frac{\operatorname{var}_{\hat{\mathbb{P}}^k(\cdot | s, a)}[V_p^*] L(n^k(s, a))}{n^k(s, a)}} + \sqrt{\frac{L(n^k(s, a))}{n^k(s, a)}} \cdot \epsilon H} + \frac{HL(n^k(s, a))}{n^k(s, a)} + \frac{HL(n^k(s, a))}{n^k(s, a)} \right) \\ \leq 2\epsilon + O\left(\sqrt{\frac{\operatorname{var}_{\hat{\mathbb{P}}^k(\cdot | s, a)}[\overline{V}_p^k] L(n^k(s, a))}{n^k(s, a)}} + \sqrt{\frac{||\overline{V}_p^k - V_p^*||_{\hat{\mathbb{P}}^k(\cdot | s, a)}^2 L(n^k(s, a))}{n^k(s, a)}} + \frac{HL(n^k(s, a))}{n^k(s, a)} \right) \\ \leq 2\epsilon + O\left(\sqrt{\frac{\operatorname{var}_{\hat{\mathbb{P}}^k(\cdot | s, a)}[\overline{V}_p^k] L(n^k(s, a))}{n^k(s, a)}} + \sqrt{\frac{||\overline{V}_p^k - V_p^k||_{\hat{\mathbb{P}}^k(\cdot | s, a)}^2 L(n^k(s, a))}{n^k(s, a)}} + \frac{HL(n^k(s, a))}{n^k(s, a)} \right) \\ \leq 2\epsilon + O\left(\sqrt{\frac{\operatorname{var}_{\hat{\mathbb{P}}^k(\cdot | s, a)}[\overline{V}_p^k] L(n^k(s, a))}{n^k(s, a)}} + \sqrt{\frac{||\overline{V}_p^k - V_p^k||_{\hat{\mathbb{P}}^k(\cdot | s, a)}^2 L(n^k(s, a))}{n^k(s, a)}} + \frac{HL(n^k(s, a))}{n^k(s, a)} \right) \\ \leq 2\epsilon + O\left(\sqrt{\frac{\operatorname{var}_{\hat{\mathbb{P}}^k(\cdot | s, a)}, n^k(s, a), \overline{V}_p^k, \underline{V}_p^k, \epsilon}{n^k(s, a)}} + \sqrt{\frac{||\overline{V}_p^k - V_p^k||_{\hat{\mathbb{P}}^k(\cdot | s, a)}^2 L(n^k(s, a))}{n^k(s, a)}} + \frac{HL(n^k(s, a))}{n^k(s, a)} \right) \\ \leq b_{\text{prob}}\left(\hat{\mathbb{P}}^k(\cdot | s, a), n^k(s, a), \overline{V}_p^k, \underline{V}_p^k, \epsilon \right), \end{split}$$

where the first inequality is from the observation that $\|\bar{\mathbb{P}}^{k}(\cdot | s, a) - \mathbb{P}_{p}(\cdot | s, a)\|_{1} \leq \frac{\epsilon}{H}$ and Lemma 25 the second inequality is from the definition of $E_{\text{agg,val}}$; the third inequality is from Equation (35) of Lemma 16 the fourth inequality is from Lemma 24 and the observation that for constant c > 0, $c\sqrt{\frac{L(n^{k}(s,a))}{n^{k}(s,a)}} \cdot \epsilon H \leq \epsilon + \frac{c^{2}}{4} \frac{L(n^{k}(s,a))}{n^{k}(s,a)}$ by AM-GM inequality; the fifth inequality is from our assumption that for all $s' \in S_{h+1}$, $\underline{V}_{p}^{k}(s') \leq V^{*}(s') \leq \overline{V}_{p}^{k}(s')$, and thus $\left| (V_{p}^{*} - \underline{V}_{p}^{k})(s') \right| \leq \left| (\overline{V}_{p}^{k} - \underline{V}_{p}^{k})(s') \right|$ for all s' in the support of $\hat{\mathbb{P}}^{k}(\cdot | s, a)$.

3. We prove Equation (28) as follows:

$$\begin{split} & \left| (\hat{\mathbb{P}}_{p}^{k} - \mathbb{P}_{p})(V_{2} - V_{1})(s, a) \right| \\ \leq & \sum_{s' \in \mathcal{S}_{h+1}} \left| (\hat{\mathbb{P}}_{p}^{k} - \mathbb{P}_{p})(s' \mid s, a) \right| \cdot (V_{2} - V_{1})(s') \\ \leq & O\left(\sum_{s' \in \mathcal{S}_{h+1}} \left(\sqrt{\frac{L(n_{p}^{k}(s, a)) \cdot \mathbb{P}_{p}(s' \mid s, a)}{n_{p}^{k}(s, a)}} + \frac{L(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)} \right) \cdot (V_{2} - V_{1})(s') \right) \\ \leq & O\left(\sum_{s' \in \mathcal{S}_{h+1}} \left(\sqrt{\frac{L(n_{p}^{k}(s, a)) \cdot \hat{\mathbb{P}}_{p}^{k}(s' \mid s, a)}{n_{p}^{k}(s, a)}} + \frac{L(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)} \right) \cdot (V_{2} - V_{1})(s') \right) \\ \leq & O\left(\sum_{s' \in \mathcal{S}_{h+1}} \sqrt{\hat{\mathbb{P}}_{p}^{k}(s' \mid s, a)} (\overline{V}_{p}^{k} - \underline{V}_{p}^{k})(s') \cdot \sqrt{\frac{L(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)}} + \sum_{s' \in \mathcal{S}_{h+1}} \frac{HL(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)} \right) \\ \leq & O\left(\sqrt{\frac{S \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\hat{\mathbb{P}}_{p}^{k}(\cdot \mid s, a)}^{2} L(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)}} + \frac{SHL(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)} \right) \\ \leq & O\left(\sqrt{\frac{S \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\hat{\mathbb{P}}_{p}^{k}(\cdot \mid s, a)}^{2} L(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)}} + \frac{SHL(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)}} \right) \\ \leq & O\left(\sqrt{\frac{S \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\hat{\mathbb{P}}_{p}^{k}(\cdot \mid s, a)}^{2} L(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)}} + \frac{SHL(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)}} \right) \\ \leq & O_{\text{str}} \left(\hat{\mathbb{P}}_{p}^{k}(\cdot \mid s, a), n(s, a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, 0 \right), \end{split}$$

where the first inequality is from the elementary fact that $\left|\sum_{i=1}^{n} a_i\right| \leq \sum_{i=1}^{n} |a_i|$; the second inequality is from the definition of $E_{\text{ind,prob}}$; the third inequality is from the definition of $E_{\text{ind,prob}}$ and Lemma 26; the fourth inequality is by algebra and $0 \leq (V_2 - V_1)(s') \leq c_1$ $\min(H, (\overline{V}_p^k - \underline{V}_p^k)(s'))$ for all $s' \in \mathcal{S}_{h+1}$; the fifth inequality is by Cauchy-Schwarz.

We now prove Equation (29):

$$\begin{split} & \left| (\hat{\mathbb{P}}^{k} - \mathbb{P}_{p})(V_{2} - V_{1})(s, a) \right| \\ \leq & \left| (\bar{\mathbb{P}}^{k} - \mathbb{P}_{p})(V_{2} - V_{1})(s, a) \right| + \left| (\hat{\mathbb{P}}^{k} - \bar{\mathbb{P}}^{k})(V_{2} - V_{1})(s, a) \right| \\ \leq & \epsilon + \sum_{s' \in \mathcal{S}_{h+1}} \left| (\hat{\mathbb{P}}^{k} - \bar{\mathbb{P}}^{k})(s' \mid s, a) \right| \cdot (V_{2} - V_{1})(s') \\ \leq & \epsilon + O\left(\sum_{s' \in \mathcal{S}_{h+1}} \left(\sqrt{\frac{L(n^{k}(s, a)) \cdot \bar{\mathbb{P}}^{k}(s' \mid s, a)}{n^{k}(s, a)}} + \frac{L(n^{k}(s, a))}{n^{k}(s, a)} \right) \cdot (V_{2} - V_{1})(s') \right) \\ \leq & \epsilon + O\left(\sum_{s' \in \mathcal{S}_{h+1}} \left(\sqrt{\frac{L(n^{k}(s, a)) \cdot \bar{\mathbb{P}}^{k}(s' \mid s, a)}{n^{k}(s, a)}} + \frac{L(n^{k}(s, a))}{n^{k}(s, a)} \right) \cdot (V_{2} - V_{1})(s') \right) \\ \leq & \epsilon + O\left(\sum_{s' \in \mathcal{S}_{h+1}} \sqrt{\hat{\mathbb{P}}^{k}(s' \mid s, a)} (\bar{V}_{p}^{k} - \underline{V}_{p}^{k})(s') \cdot \sqrt{\frac{L(n^{k}(s, a))}{n^{k}(s, a)}} + \sum_{s' \in \mathcal{S}_{h+1}} \frac{HL(n^{k}(s, a))}{n^{k}(s, a)} \right) \\ \leq & \epsilon + O\left(\sqrt{\frac{S \| \overline{V}_{p}^{k} - \underline{V}_{p}^{k} \|_{\mathbb{P}^{k}(\cdot \mid s, a)}^{2} L(n^{k}(s, a))}{n^{k}(s, a)}} + \frac{SHL(n^{k}(s, a))}{n^{k}(s, a)} \right) \\ \leq & \delta_{\text{str}} \left(\hat{\mathbb{P}}^{k}(\cdot \mid s, a), n(s, a), \overline{V}_{p}^{k}, \underline{V}_{p}^{k}, \epsilon \right), \end{split}$$

where the first inequality is triangle inequality; the second inequality is from the elementary fact that $\left|\sum_{i=1}^{n} a_{i}\right| \leq \sum_{i=1}^{n} |a_{i}|$, along with $\|\bar{\mathbb{P}}_{k}(\cdot \mid s, a) - \mathbb{P}_{p}(\cdot \mid s, a)\|_{1} \leq \frac{\epsilon}{H}$ and Lemma 25 the third inequality is from the definition of $E_{\text{agg,prob}}$; the fourth inequality is from the definition of $E_{\text{agg,prob}}$; the fourth inequality is from the definition of $E_{\text{agg,prob}}$; the fourth inequality is by algebra and $0 \leq (V_{2} - V_{1})(s') \leq \min(H, (\overline{V}_{p}^{k} - \underline{V}_{p}^{k})(s'))$ for all $s' \in S_{h+1}$; the last inequality is by Cauchy-Schwarz.

Lemma 2 relies on the following technical lemma on the concentrations of the conditional variances. Specifically, Equation (34) is well-known (see, e.g., [2, 30]); Equations (35) and (36) are new, and allow for heterogeneous data aggregation in the multi-task RL setting. We still include the proof of Equation (34) here, as it helps illustrate our ideas for proving the two new inequalities.

Lemma 16. If event E happens, then for any s, a, k, p, we have:

$$\left| \sqrt{\operatorname{var}_{\hat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}\left[V_{p}^{\star} \right]} - \sqrt{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}\left[V_{p}^{\star} \right]} \right| \lesssim H \sqrt{\frac{L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}},$$
(34)

2.

1.

$$\left| \sqrt{\operatorname{var}_{\hat{\mathbb{P}}^{k}(\cdot|s,a)}\left[V_{p}^{\star}\right]} - \sqrt{\sum_{q=1}^{M} w_{q}^{k}(s,a) \operatorname{var}_{\mathbb{P}_{q}(\cdot|s,a)}\left[V_{p}^{\star}\right]} \right| \lesssim \sqrt{H\epsilon} + H\sqrt{\frac{L(n^{k}(s,a))}{n^{k}(s,a)}},$$
(35)

and

$$\left|\sqrt{\operatorname{var}_{\widehat{\mathbb{P}}^{k}(\cdot|s,a)}\left[V_{p}^{\star}\right]} - \sqrt{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}\left[V_{p}^{\star}\right]}\right| \lesssim \sqrt{H\epsilon} + H\sqrt{\frac{L(n^{k}(s,a))}{n^{k}(s,a)}}, \quad (36)$$

Proof. 1. By the definition of E, we have

$$\left|\frac{1}{n_p^k(s,a)} \sum_{i=1}^{n_p^k(s,a)} (V_p^{\star}((s_i^p)') - (\mathbb{P}_p V_p^{\star})(s,a))^2 - \operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\star}]\right| \lesssim \sqrt{\frac{H^2 \operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\star}]L(n_p^k(s,a))}{n_p^k(s,a)}} + \frac{H^2 L(n_p^k(s,a))}{n_p^k(s,a)};$$

this, when combined with Lemma 26, implies that

$$\left| \sqrt{\frac{1}{n_p^k(s,a)} \sum_{i=1}^{n_p^k(s,a)} (V_p^\star((s_i^p)') - (\mathbb{P}_p V_p^\star)(s,a))^2} - \sqrt{\operatorname{var}_{\mathbb{P}_p}(\cdot|s,a)} [V_p^\star] \right| \le H \sqrt{\frac{L(n_p^k(s,a))}{n_p^k(s,a)}}.$$
 (37)

Now, observe that

$$\operatorname{var}_{\hat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}\left[V_{p}^{\star}\right] = \frac{1}{n_{p}^{k}(s,a)} \sum_{i=1}^{n_{p}^{k}(s,a)} (V_{p}^{\star}((s_{i}^{p})') - (\mathbb{P}_{p}V_{p}^{\star})(s,a))^{2} - ((\hat{\mathbb{P}}_{p}^{k}V_{p}^{\star})(s,a) - (\mathbb{P}_{p}V_{p}^{\star})(s,a))^{2},$$

which can be seen by applying Lemma 23 with X being the random variable that is drawn uniformly from $\left\{V_p^{\star}((s_i^p)')\right\}_{i=1}^{n_p^k(s,a)}$, which has expectation $\mu = (\hat{\mathbb{P}}_p^k V_p^{\star})(s,a)$, and setting $m = (\mathbb{P}_p V_p^{\star})(s,a)$. Recall that by the definition of event E, we have

$$\left| (\hat{\mathbb{P}}_{p}^{k} V_{p}^{\star})(s, a) - (\mathbb{P}_{p} V_{p}^{\star})(s, a) \right| \leq H \wedge \left(\sqrt{\frac{H^{2}L(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)}} + \frac{HL(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)} \right) \leq 2H \sqrt{\frac{L(n_{p}^{k}(s, a))}{n_{p}^{k}(s, a)}},$$

where the second inequality uses Lemma 27 Using the elementary fact that $|A - B| \le C \Rightarrow \sqrt{A} \le \sqrt{B} + \sqrt{C}$, we get that

$$\left| \sqrt{\operatorname{var}_{\mathbb{P}_{p}^{k}(\cdot|s,a)}\left[V_{p}^{\star}\right]} - \sqrt{\frac{1}{n^{k}(s,a)}\sum_{i=1}^{n^{k}(s,a)} (V_{p}^{\star}((s_{i}^{p})') - (\mathbb{P}_{p}V_{p}^{\star})(s,a))^{2}} \right|$$

$$\leq \left| (\hat{\mathbb{P}}_{p}^{k}V_{p}^{\star})(s,a) - (\mathbb{P}_{p}V_{p}^{\star})(s,a) \right| \lesssim H\sqrt{\frac{L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}.$$

$$(38)$$

Combining Equations (37) and (38), using algebra, we get

$$\left|\sqrt{\operatorname{var}_{\hat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}\left[V_{p}^{\star}\right]} - \sqrt{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}\left[V_{p}^{\star}\right]}\right| \lesssim H\sqrt{\frac{L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}},$$

establishing Equation (34).

2. We first show Equation (35). By the definition of E, we have

$$\begin{split} & \left| \frac{1}{n^{k}(s,a)} \sum_{i=1}^{n^{k}(s,a)} (V_{p}^{\star}(s_{i}') - (\mathbb{P}_{p_{i}}V_{p}^{\star})(s,a))^{2} - \sum_{p=1}^{M} w_{p}^{k}(s,a) \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\star}] \right| \\ & \lesssim \sqrt{\frac{H^{2}\left(\sum_{p=1}^{M} w_{p}^{k}(s,a) \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\star}]\right) L(n^{k}(s,a))}{n^{k}(s,a)}} + \frac{H^{2}L(n^{k}(s,a))}{n^{k}(s,a)}, \end{split}$$

this, combined with Lemma 26, implies that

$$\left| \sqrt{\frac{1}{n^{k}(s,a)} \sum_{i=1}^{n^{k}(s,a)} (V_{p}^{\star}(s_{i}^{\prime}) - (\mathbb{P}_{p_{i}}V_{p}^{\star})(s,a))^{2}} - \sqrt{\sum_{p=1}^{M} w_{p}^{k}(s,a) \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\star}]} \right| \lesssim H \sqrt{\frac{L(n^{k}(s,a))}{n^{k}(s,a)}}$$
(39)

For the first term on the left hand side, observe that for each i, $|(\mathbb{P}_{p_i}V_p^{\star})(s,a) - (\mathbb{P}_pV_p^{\star})(s,a)| \leq 1$ $H_{\overline{H}}^{\epsilon} = \epsilon, \text{ we therefore have } \left| (V_p^{\star}(s_i') - (\mathbb{P}_{p_i}V_p^{\star})(s,a))^2 - (V_p^{\star}(s_i') - (\mathbb{P}_pV_p^{\star})(s,a))^2 \right| \le 2H\epsilon \text{ by}$ 2*H*-Lipschitzness of function $f(x) = x^2$ on [-H, H]. By averaging over all *i*'s and taking square root, we have

$$\left| \sqrt{\frac{1}{n^{k}(s,a)} \sum_{i=1}^{n^{k}(s,a)} (V_{p}^{\star}(s_{i}^{\prime}) - (\mathbb{P}_{p_{i}}V_{p}^{\star})(s,a))^{2}} - \sqrt{\frac{1}{n^{k}(s,a)} \sum_{i=1}^{n^{k}(s,a)} (V_{p}^{\star}(s_{i}^{\prime}) - (\mathbb{P}_{p}V_{p}^{\star})(s,a))^{2}} \right| \lesssim \sqrt{H\epsilon}.$$
(40)

Furthermore,

$$\operatorname{var}_{\hat{\mathbb{P}}^{k}(\cdot|s,a)}\left[V_{p}^{\star}\right] = \frac{1}{n^{k}(s,a)} \sum_{i=1}^{n^{k}(s,a)} (V_{p}^{\star}(s_{i}') - (\mathbb{P}_{p}V_{p}^{\star})(s,a))^{2} - ((\hat{\mathbb{P}}^{k}V_{p}^{\star})(s,a) - (\mathbb{P}_{p}V_{p}^{\star})(s,a))^{2},$$

which can be seen by applying Lemma 23 with X being the random variable that is drawn uniformly from $\left\{V_p^{\star}(s_i)\right\}_{i=1}^{n^k(s,a)}$, which has expectation $\mu = (\hat{\mathbb{P}}^k V_p^{\star})(s,a)$, and setting $m = (\mathbb{P}_p V_p^{\star})(s,a)$. In addition,

$$\left| (\hat{\mathbb{P}}^k V_p^\star)(s, a) - (\mathbb{P}_p V_p^\star)(s, a) \right| \lesssim \epsilon + H \sqrt{\frac{L(n^k(s, a))}{n^k(s, a)}}$$

Together with our assumption that $\epsilon \leq 2H$ (which implies that $\epsilon \lesssim \sqrt{H\epsilon}$), this gives

$$\left|\sqrt{\operatorname{var}_{\hat{\mathbb{P}}^{k}(\cdot|s,a)}\left[V_{p}^{\star}\right]} - \sqrt{\frac{1}{n^{k}(s,a)}}\sum_{i=1}^{n^{k}(s,a)} \left(V_{p}^{\star}(s_{i}^{\prime}) - (\mathbb{P}_{p}V_{p}^{\star})(s,a)\right)^{2}}\right| \lesssim \sqrt{H\epsilon} + H\sqrt{\frac{L(n^{k}(s,a))}{n^{k}(s,a)}}.$$

$$(41)$$

Equation (35) is a direct consequence of Equations (39), (40) and (41) along with algebra.

We now show Equation (36) using Equation (35). By Lemma 25, for every q, $\left|\operatorname{var}_{\mathbb{P}_q(\cdot|s,a)}\left[V_p^{\star}\right] - \operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}\left[V_p^{\star}\right]\right| \leq 3H^2 \cdot \frac{\epsilon}{H} = 3H\epsilon$. Therefore, Jensen's inequality and the convexity of $|\cdot|$ implies $\left|\sum_{q=1}^{M} w_q^k(s, a) \operatorname{var}_{\mathbb{P}_p(\cdot|s, a)}[V_p^{\star}] - \operatorname{var}_{\mathbb{P}_p(\cdot|s, a)}\left[V_p^{\star}\right]\right| \le 3H\epsilon$, and $\left|\sqrt{\sum_{q=1}^{M} w_q^k(s, a) \operatorname{var}_{\mathbb{P}_p(\cdot|s, a)}[V_{-}^{\star}]} - \sqrt{\operatorname{var}_{\mathbb{P}_p(\cdot|s, a)}[V_{+}^{\star}]}\right| < \sqrt{\mu_{\epsilon}}$

$$\sqrt{\sum_{q=1}^{M} w_q^k(s, a) \operatorname{var}_{\mathbb{P}_q(\cdot|s, a)}[V_p^{\star}]} - \sqrt{\operatorname{var}_{\mathbb{P}_p(\cdot|s, a)}\left[V_p^{\star}\right]} \lesssim \sqrt{H\epsilon}$$

This, together with Equation (35), implies

$$\left| \sqrt{\operatorname{var}_{\hat{\mathbb{P}}^{k}(\cdot|s,a)} \left[V_{p}^{\star} \right]} - \sqrt{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)} \left[V_{p}^{\star} \right]} \right| \lesssim \sqrt{H\epsilon} + H\sqrt{\frac{L(n^{k}(s,a))}{n^{k}(s,a)}},$$

establishing Equation (36).

Simplifying the surplus bounds C.3

In this section, we show a distribution-dependent bound on the surplus terms, namely Lemma 19 which is key to establishing our regret bound. It can be seen as an extension of Proposition B.4 of 36 to our multitask setting using the MULTI-TASK-EULER algorithm, under the ϵ -dissimilarity assumption. Before we present Lemma 19 (Section C.3.1), we first show and prove two auxiliary lemmas, Lemma 17 and Lemma 18

Lemma 17 (Bounds on $\overline{V}_p^k - \underline{V}_p^k$, generalization of [36], Lemma F.8). If *E* happens, then for all $p \in [M], k \in [K], h \in [H+1]$ and $s \in S_h$,

$$(\overline{V}_p^k - \underline{V}_p^k)(s) \le 4\mathbb{E}\left[\sum_{t=h}^{H} \left(H \wedge ind \cdot b_p^k(s_t, a_t) \wedge agg \cdot b_p^k(s_t, a_t)\right) \mid s_h = s, \pi^k(p), \mathcal{M}_p\right]; \quad (42)$$

consequently,

$$(\overline{V}_p^k - \underline{V}_p^k)(s) \lesssim H \sum_{t=h}^H \mathbb{E}\left[\left(1 \wedge \sqrt{\frac{SL(n_p^k(s_t, a_t))}{n_p^k(s_t, a_t)}}\right) \mid s_h = s, \pi^k(p), \mathcal{M}_p\right].$$
(43)

Proof. First, Lemmas 15 and 14 together imply that if E holds, Equations (30) and (31) holds for all p, k, s, a. Under this premise, we show Equation (42) by backward induction.

Base case: for h = H + 1, we have that LHS is $(\overline{V}_p^k - \underline{V}_p^k)(\perp) = 0$ which is equal to the RHS.

inductive case: Suppose Equation (42) holds for all $s \in S_{h+1}$. Now consider $s \in S_h$. By the definitions of \overline{V}_p^k and \underline{V}_p^k ,

$$\begin{split} &(\overline{V}_p^k - \underline{V}_p^k)(s) \\ = \overline{Q}_p^k(s, \pi_p^k(s)) - \underline{Q}_p^k(s, \pi_p^k(s)) \\ \leq &(\mathbb{P}_p(\overline{V}_p^k - \underline{V}_p^k))(s, \pi_p^k(s)) + 4(H \wedge \operatorname{ind} b_p^k(s, \pi_p^k(s)) \wedge \operatorname{agg-} b_p^k(s, \pi_p^k(s))) \\ = &\mathbb{E}\left[4\min(H, \operatorname{ind-} b_p^k(s, a), \operatorname{agg-} b_p^k(s, a)) + (\overline{V}_p^k - \underline{V}_p^k)(s_{h+1}) \mid s_h = s, \pi_p^k, \mathcal{M}_p\right] \\ \leq &\mathbb{E}\left[4(H \wedge \operatorname{ind-} b_p^k(s, a) \wedge \operatorname{agg-} b_p^k(s, a)) + \mathbb{E}\left[\sum_{t=h+1}^H \left(H \wedge 2\operatorname{ind-} b_p^k(s_t, a_t) \wedge 2\operatorname{agg-} b_p^k(s_t, a_t)\right) \mid s_{h+1}\right] \mid s_h = s, \pi_p^k, \mathcal{M}_p\right] \\ \leq &4\mathbb{E}\left[\sum_{t=h}^H \left(H \wedge \operatorname{ind-} b_p^k(s_t, a_t) \wedge \operatorname{agg-} b_p^k(s_t, a_t)\right) \mid s_h = s, \pi_p^k, \mathcal{M}_p\right], \end{split}$$

where the first inequality is from Equations (30) and (31) for (s, a) and player p at episode k, and the second inequality is from the inductive hypothesis; the third inequality is by algebra. This completes the induction.

We now show Equation (43). By the definition of $\operatorname{ind} b_p^k(s, a)$ and algebra,

$$\begin{split} & \inf -b_{p}^{k}(s,a) \\ \lesssim \sqrt{\frac{\operatorname{var}_{\widehat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}\left[\overline{V}_{p}^{k}\right] L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{S \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\widehat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}^{2} L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{HSL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \\ \lesssim H \sqrt{\frac{SL(n_{p}^{k}(s_{t},a_{t}))}{n_{p}^{k}(s_{t},a_{t})}} + \frac{HSL(n_{p}^{k}(s_{t},a_{t}))}{n_{p}^{k}(s_{t},a_{t})}, \end{split}$$

where the second inequality uses $\operatorname{var}_{\widehat{\mathbb{P}}^k_p(\cdot|s,a)}\left[\overline{V}^k_p\right] \leq H^2$ and $\|\overline{V}^k_p - \underline{V}^k_p\|^2_{\widehat{\mathbb{P}}^k_p} \leq H^2$. As a consequence, using Lemma 27.

$$\begin{split} H \wedge \operatorname{ind-} b_p^k(s_t, a_t) \wedge \operatorname{agg-} b_p^k(s_t, a_t) \lesssim & H \wedge \left(H \sqrt{\frac{SL(n_p^k(s_t, a_t))}{n_p^k(s_t, a_t)}} + \frac{HSL(n_p^k(s_t, a_t))}{n_p^k(s_t, a_t)} \right) \\ \lesssim & H \left(1 \wedge \sqrt{\frac{SL(n_p^k(s_t, a_t))}{n_p^k(s_t, a_t)}} \right). \end{split}$$

Lemma 18. If E happens, we have the following statements holding for all p, k, s, a:

1. For two terms that appear in $ind-b_p^k(s, a)$, they are bounded respectively as:

$$\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2} \lesssim \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{H^{2}SL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}$$
(44)

$$\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}^{k}(\cdot|s,a)}\left[\overline{V}_{p}^{k}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} \lesssim \sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}\left[V_{p}^{\pi^{k}}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{H\sqrt{S}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \tag{45}}$$

2. For two terms that appear in $agg-b_p^k(s, a)$, they are bounded respectively as:

$$\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}^{k}(\cdot|s,a)}^{2} \lesssim 2\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{H^{2}SL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} + H\epsilon$$

$$(46)$$

$$\sqrt{\frac{\operatorname{var}_{\mathbb{P}^{k}(\cdot|s,a)}\left[\overline{V}_{p}^{k}\right]L(n^{k}(s,a))}{n^{k}(s,a)}} \lesssim \sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}\left[V_{p}^{\pi^{k}}\right]L(n^{k}(s,a))}{n^{k}(s,a)}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2}L(n^{k}(s,a))}{n^{k}(s,a)}} + \frac{H\sqrt{S}L(n^{k}(s,a))}{n^{k}(s,a)} + \sqrt{\frac{H\epsilon L(n^{k}(s,a))}{n^{k}(s,a)}} \tag{47}$$

Proof. First, Lemmas 15 and 14 together imply that if E happens, the value function upper and lower bounds are valid. Conditioned on E happening, we prove the two items respectively.

1. For Equation (44), using the definition of $E_{ind,prob}$ and AM-GM inequality, when E happens, we have for all p, k, s, a, s',

$$\hat{\mathbb{P}}_p^k(s' \mid s, a) \lesssim \mathbb{P}_p(s' \mid s, a) + \frac{L(n_p^k(s, a))}{n_p^k(s, a)}.$$
(48)

This implies that

$$\begin{split} \|\overline{V}_p^k - \underline{V}_p^k\|_{\mathbb{P}_p^k(\cdot|s,a)}^2 \\ &= \sum_{s' \in \mathcal{S}_{h+1}} \hat{\mathbb{P}}_p^k(s' \mid s,a) (\overline{V}_p^k(s') - \underline{V}_p^k(s'))^2 \\ &\lesssim \sum_{s' \in \mathcal{S}_{h+1}} \mathbb{P}_p^k(s' \mid s,a) (\overline{V}_p^k(s') - \underline{V}_p^k(s'))^2 + \sum_{s' \in \mathcal{S}_{h+1}} \frac{L(n_p^k(s,a))}{n_p^k(s,a)} \cdot H^2 \\ &\lesssim \|\overline{V}_p^k - \underline{V}_p^k\|_{\mathbb{P}_p(\cdot|s,a)}^2 + \frac{SH^2L(n_p^k(s,a))}{n_p^k(s,a)}, \end{split}$$

where the first inequality is from Equation (48), and the fact that $\overline{V}_p^k(s') - \underline{V}_p^k(s') \in [0, H]$ for any $s' \in S_{h+1}$; the second inequality is by algebra.

For Equation (45), we have:

$$\begin{split} &\sqrt{\frac{\operatorname{var}_{\widehat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}\left[\overline{V}_{p}^{k}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}} \\ \lesssim &\sqrt{\frac{\operatorname{var}_{\widehat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}\left[V_{p}^{\star}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\widehat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} \\ \lesssim &\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}}(\cdot|s,a)}\left[V_{p}^{\star}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{\sqrt{SHL(n_{p}^{k}(s,a))}}{n_{p}^{k}(s,a)} \\ \lesssim &\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}}(\cdot|s,a)}\left[V_{p}^{\pi^{k}}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{\sqrt{SHL(n_{p}^{k}(s,a))}}{n_{p}^{k}(s,a)} \\ \end{cases} \end{split}$$

where the first inequality is from Lemma 24 and the observation that when E happens, $\left|(\overline{V}_p^k - V_p^\star)(s')\right| \leq \left|(\overline{V}_p^k - \underline{V}_p^k)(s')\right|$ for all $s' \in \mathcal{S}_{h+1}$; the second inequality is from Equation (34) of Lemma 16 and Equation (44); the third inequality again uses Lemma 24 and the observation that when E happens, $\left|(V_p^\star - V_p^{\pi^k})(s')\right| \leq \left|(\overline{V}_p^k - \underline{V}_p^k)(s')\right|$ for all $s' \in \mathcal{S}_{h+1}$.

2. For Equation (46), using the definition of $E_{agg,prob}$ and AM-GM inequality, when E happens, we have for all p, k, s, a, s',

$$\hat{\mathbb{P}}^k(s' \mid s, a) \lesssim \bar{\mathbb{P}}^k(s' \mid s, a) + \frac{L(n^k(s, a))}{n^k(s, a)}.$$
(49)

This implies that

$$\begin{split} \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\hat{\mathbb{P}}^{k}(\cdot|s,a)}^{2} \\ &= \sum_{s' \in \mathcal{S}_{h+1}} \hat{\mathbb{P}}^{k}(s' \mid s,a)(\overline{V}_{p}^{k}(s') - \underline{V}_{p}^{k}(s'))^{2} \\ &\lesssim 2 \sum_{s' \in \mathcal{S}_{h+1}} \bar{\mathbb{P}}^{k}(s' \mid s,a)(\overline{V}_{p}^{k}(s') - \underline{V}_{p}^{k}(s'))^{2} + \sum_{s' \in \mathcal{S}_{h+1}} \frac{L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \cdot H^{2} \\ &\lesssim 2 \sum_{s' \in \mathcal{S}_{h+1}} \mathbb{P}_{p}(s' \mid s,a)(\overline{V}_{p}^{k}(s') - \underline{V}_{p}^{k}(s'))^{2} + \epsilon H + \frac{SH^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \\ &\lesssim \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{SH^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} + \epsilon H, \end{split}$$

where the first inequality is from Equation (49) and the fact that $\overline{V}_p^k(s') - \underline{V}_p^k(s') \in [0, H]$ for any $s' \in S_{h+1}$; the second inequality is from the observation that $\|\mathbb{P}_p(\cdot \mid s, a) - \overline{\mathbb{P}}^k(\cdot \mid s, a)\|_1 \leq \frac{\epsilon}{H}$; the third inequality is by algebra.

For Equation (47), we have:

$$\begin{split} &\sqrt{\frac{\operatorname{var}_{\widehat{\mathbb{P}}^{k}(\cdot|s,a)}\left[\overline{V}_{p}^{k}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} \\ \lesssim &\sqrt{\frac{\operatorname{var}_{\widehat{\mathbb{P}}^{k}(\cdot|s,a)}\left[V_{p}^{\star}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\widehat{\mathbb{P}}^{k}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} \\ \lesssim &\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}\left[V_{p}^{\star}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{\sqrt{SHL(n_{p}^{k}(s,a))}}{n_{p}^{k}(s,a)} + \sqrt{\frac{H\epsilon L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} \\ \lesssim &\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}\left[V_{p}^{\star}\right]L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2}L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{\sqrt{SHL(n_{p}^{k}(s,a))}}{n_{p}^{k}(s,a)} + \sqrt{\frac{H\epsilon L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}, \end{split}$$

where the first inequality is from Lemma 24 and the observation that when E happens, $\left| (\overline{V}_p^k - V_p^\star)(s') \right| \leq \left| (\overline{V}_p^k - \underline{V}_p^k)(s') \right|$ for $s' \in S_{h+1}$; the second inequality uses Equation (36) of Lemma 6 and Equation (46); the third inequality is from Lemma 24 and the observation that when E happens, $\left| (\overline{V}_p^\star - V_p^{\pi^k})(s') \right| \leq \left| (\overline{V}_p^k - \underline{V}_p^k)(s') \right|$ for $s' \in S_{h+1}$.

C.3.1 Distribution-dependent bound on the surplus terms

Lemma 19 (Surplus bound). If E happens, then for all p, k, s, a:

$$E_p^k(s,a) \lesssim B_p^{k,\text{lead}}(s,a) + \mathbb{E}\left[\sum_{t=h}^H B_p^{k,\text{fut}}(s_t,a_t) \mid (s_h,a_h) = (s,a), \pi^k(p), \mathcal{M}_p\right],$$

where

$$\begin{split} B_p^{k,\mathrm{lead}}(s,a) &= H \wedge \left(5\epsilon + O\left(\sqrt{\frac{\left(1 + \mathrm{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\pi^k}]\right)L(n^k(s,a))}{n^k(s,a)}}\right) \right) \\ & \wedge O\left(\sqrt{\frac{\left(1 + \mathrm{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\pi^k}]\right)L(n_p^k(s,a))}{n_p^k(s,a)}}\right), \\ B_p^{k,\mathrm{fut}}(s,a) &= H^3 \wedge O\left(\frac{H^3SL(n_p^k(s,a))}{n_p^k(s,a)}\right). \end{split}$$

Proof of Lemma [79] First, Lemmas [15] and [14] together imply that if E holds, for all p, k, s, a, $E_p^k(s, a) \leq 2 \left(H \wedge \operatorname{ind} b_p^k(s, a) \wedge \operatorname{agg-} b_p^k(s, a) \right)$. We now bound $\operatorname{ind-} b_p^k(s, a)$ and $\operatorname{agg-} b_p^k(s, a)$ respectively.

Bounding ind- $b_p^k(s, a)$: We have

$$\begin{split} & \operatorname{ind} b_{p}^{k}(s, a) \\ = O\left(\sqrt{\frac{\operatorname{var}_{\widehat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}[\overline{V}_{p}^{k}] \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{S \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\widehat{\mathbb{P}}_{p}^{k}(\cdot|s,a)}^{2} \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}\right) \\ \leq O\left(\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[\overline{V_{p}^{k}}] \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{S \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2} \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}\right) \\ \leq O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[\overline{V_{p}^{k}}]\right) \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \sqrt{\frac{S \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2} \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}\right) \\ \leq O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[\overline{V_{p}^{k}}]\right) \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2} + \frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}\right)}{n_{p}^{k}(s,a)}\right) \\ \leq O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[\overline{V_{p}^{k}}]\right) \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}} + \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2} + \frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}\right)}\right) \\ \leq O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[\overline{V_{p}^{k}}]\right) \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2} + \frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}\right)}\right) \\ \leq O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[\overline{V_{p}^{k}}]\right) \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2} + \frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}\right)}\right) \\ \leq O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[\overline{V}_{p}^{k}]\right) \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} + \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}^{k}(\cdot|s,a)}^{2} + \frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}\right)}\right) \\ \leq O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}^{k}(\cdot|s,a)}[\overline{V}_{p}^{k}]\right) \ L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} + \frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}\right)}\right) \\ \leq O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}^{k}(\cdot|s,a)}[\overline{V}_{p$$

where the first inequality is by expanding the definition of ind- $b_p^k(s, a)$ and algebra; the second inequality is from Equations Equation (44) and (45) of Lemma 18 along with algebra; the third inequality is by the basic fact that $\sqrt{A} + \sqrt{B} \lesssim \sqrt{A+B}$; the fourth inequality is by AM-GM inequality.

Bounding agg- $b_p^k(s, a)$: We have:

$$\begin{split} & \arg b_{p}^{k}(s,a) \\ \lesssim & 4\epsilon + O\left(\sqrt{\frac{\operatorname{var}_{\mathbb{P}^{k}(\cdot|s,a)}[\overline{V}_{p}^{k}] L(n^{k}(s,a))}{n^{k}(s,a)}} + \sqrt{\frac{L(n^{k}(s,a))}{n^{k}(s,a)}} + \sqrt{\frac{S||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}^{k}(\cdot|s,a)}^{2} L(n^{k}(s,a))}{n^{k}(s,a)}} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}\right) \\ \lesssim & 5\epsilon + O\left(\sqrt{\frac{\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}] L(n^{k}(s,a))}{n^{k}(s,a)}} + \sqrt{\frac{L(n^{k}(s,a))}{n^{k}(s,a)}} + \sqrt{\frac{S||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}_{p}(\cdot|s,a)}^{2} L(n^{k}(s,a))}{n^{k}(s,a)}} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}\right) \\ \lesssim & 5\epsilon + O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right) L(n^{k}(s,a))}{n^{k}(s,a)}} + \sqrt{\frac{S||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}_{p}(\cdot|s,a)}^{2} L(n^{k}(s,a))}{n^{k}(s,a)}} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}\right) \\ \leq & 5\epsilon + O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right) L(n^{k}(s,a))}{n^{k}(s,a)}} + ||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}}\right) \\ \leq & 5\epsilon + O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right) L(n^{k}(s,a))}{n^{k}(s,a)}} + ||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}}\right) \\ \leq & 5\epsilon + O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right) L(n^{k}(s,a))}{n^{k}(s,a)}} + ||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}}\right) \\ \leq & 5\epsilon + O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right) L(n^{k}(s,a))}{n^{k}(s,a)}} + ||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}}\right) \\ \leq & 5\epsilon + O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right) L(n^{k}(s,a))}{n^{k}(s,a)}} + ||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}}\right) \\ \leq & 5\epsilon + O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right) L(n^{k}(s,a))}{n^{k}(s,a)}} + ||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}}\right) \\ \leq & 5\epsilon + O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right) L(n^{k}(s,a))}{n^{k}(s,a)}} + ||\overline{V}_{p}^{k} - \underline{V}_{p}^{k}||_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \frac{SHL(n^{k}(s,a))}{n^{k}(s,a)}}\right) \\ \leq & 5\epsilon + O\left(\sqrt{$$

where the first inequality is by expanding the definition of $agg-b_p^k(s,a)$ and algebra; the second inequality is from Equations (47) and Equation (46) of Lemma 18 along with the observation that $\sqrt{\frac{S \epsilon H L(n^k(s,a))}{n^k(s,a)}} \leq \frac{S H L(n^k(s,a))}{n^k(s,a)} + \epsilon$ by AM-GM inequality; the third inequality is by the basic fact that $\sqrt{A} + \sqrt{B} \lesssim \sqrt{A + B}$; the fourth inequality is from AM-GM inequality.

Combining the above upper bounds, and using the observation that $\frac{L(n^k(s,a))}{n^k(s,a)} \leq \frac{L(n^k_p(s,a))}{n^k_p(s,a)}$, we get

$$\begin{split} & \operatorname{ind} b_p^k(s, a) \wedge \operatorname{agg-} b_p^k(s, a) \wedge H \\ \leq & O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_p}(\cdot | s, a) [V_p^{\pi^k}]\right) \ L(n_p^k(s, a))}{n_p^k(s, a)}}\right) \wedge \left(5\epsilon + O\left(\sqrt{\frac{\left(1 + \operatorname{var}_{\mathbb{P}_p}(\cdot | s, a) [V_p^{\pi^k}]\right) \ L(n^k(s, a))}{n^k(s, a)}}\right)\right) \wedge H \\ & + O\left(\|\overline{V}_p^k - \underline{V}_p^k\|_{\mathbb{P}_p}^2(\cdot | s, a)} + \left(\frac{SHL(n_p^k(s, a))}{n_p^k(s, a)} \wedge H\right)\right) \\ \leq & B^{k, \operatorname{lead}}(s, a) + O\left(\|\overline{V}_p^k - \underline{V}_p^k\|_{\mathbb{P}_p}^2(\cdot | s, a)} + \left(\frac{SHL(n_p^k(s, a))}{n_p^k(s, a)} \wedge H\right)\right). \end{split}$$

We now show that

$$\|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \left(\frac{SHL(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)} \wedge H\right) \lesssim \mathbb{E}\left[\sum_{t=h}^{H} B^{k,\operatorname{fut}}(s_{t},a_{t}) \mid (s_{h},a_{h}) = (s,a), \pi^{k}(p), \mathcal{M}_{p}\right],$$
(50)

which will conclude the proof. To this end, we simplify the left hand side of Equation (50) using Lemma 17:

$$\begin{split} \|\overline{V}_{p}^{k} - \underline{V}_{p}^{k}\|_{\mathbb{P}_{p}(\cdot|s,a)}^{2} + \left(\frac{SHL(n^{k}(s,a))}{n^{k}(s,a)} \wedge H\right) \\ \lesssim \mathbb{E}\left[\left(H\sum_{t=h+1}^{H} \mathbb{E}\left[\left(1 \wedge \sqrt{\frac{SL(n_{p}^{k}(s_{t},a_{t}))}{n_{p}^{k}(s_{t},a_{t})}}\right) \mid s_{h+1}\right]\right)^{2} \mid (s_{h},a_{h}) = (s,a), \pi^{k}(p), \mathcal{M}_{p}\right] + \left(\frac{SHL(n^{k}(s,a))}{n^{k}(s,a)} \wedge H\right) \\ \lesssim H^{3}\mathbb{E}\left[\sum_{t=h+1}^{H} \mathbb{E}\left[\left(1 \wedge \sqrt{\frac{SL(n_{p}^{k}(s_{t},a_{t}))}{n_{p}^{k}(s_{t},a_{t})}}\right)^{2} \mid s_{h+1}\right] \mid (s_{h},a_{h}) = (s,a), \pi^{k}(p), \mathcal{M}_{p}\right] + \left(\frac{SHL(n^{k}(s,a))}{n^{k}(s,a)} \wedge H\right) \\ \lesssim \mathbb{E}\left[\sum_{t=h}^{H} H^{3} \wedge \frac{H^{3}SL(n_{p}^{k}(s_{t},a_{t}))}{n_{p}^{k}(s_{t},a_{t})} \mid (s_{h},a_{h}) = (s,a), \pi^{k}(p), \mathcal{M}_{p}\right] \\ \lesssim \mathbb{E}\left[\sum_{t=h}^{H} B^{k,\text{fut}}(s_{t},a_{t}) \mid (s_{h},a_{h}) = (s,a), \pi^{k}(p), \mathcal{M}_{p}\right], \end{split}$$

where the first inequality is from Equation (43) of Lemma 17: the second inequality is by Cauchy-Schwarz and $\mathbb{E}[X]^2 \leq \mathbb{E}[X^2]$ for any random variable X; and the third inequality is by the law of total expectation and algebra.

C.4 Concluding the regret bounds

In this section, we present the proofs of Theorems 5 and 6.

To bound the collective regret of MULTI-TASK-EULER, we first recall the following general result from [36], which is useful to establish instance-dependent regret guarantees for episodic RL.

Lemma 20 (Clipping lemma, 36), Lemma B.6). Fix player $p \in [M]$; suppose for each episode k, it follows $\pi^k(p)$, the greedy policy with respect to \overline{Q}_p^k . In addition, there exists some event E and a collection of functions $\left\{B_p^{k,\text{lead}}, B_p^{k,\text{fut}}\right\}_{k \in [K]} \subset (\mathcal{S} \times \mathcal{A} \to \mathbb{R})$, such that if E happens, then for

all $k \in [K]$, $h \in [H]$ and $(s, a) \in S_h \times A$, the surplus of \overline{Q}_p^k satisfies that

$$0 \le E_p^k(s,a) \lesssim B_p^{k,\text{lead}}(s,a) + \mathbb{E}\left[\sum_{t=h}^H B_p^{k,\text{fut}}(s_t,a_t) \mid (s_h,a_h) = (s,a), \pi^k(p), \mathcal{M}_p\right],$$

then, on E:

$$\operatorname{Reg}(K,p) \lesssim \sum_{s,a} \sum_{k} \rho_p^k(s,a) \operatorname{clip}\left(B_p^{k,\operatorname{lead}}(s,a), \operatorname{g\check{a}p}_p(s,a)\right) + H \sum_{s,a} \sum_{k} \rho_p^k(s,a) \operatorname{clip}\left(B_p^{k,\operatorname{fut}}(s,a), \frac{\operatorname{gap}_{p,\min}}{8SAH^2}\right),$$

here, recall that $\operatorname{clip}(\alpha, \Delta) = \alpha \mathbf{1}(\alpha \ge \Delta)$ *, and* $\operatorname{g extstyle gap}_p(s, a) = \frac{\operatorname{gap}_p(s, a)}{4H} \lor \frac{\operatorname{gap}_{p,\min}}{4H}$.

Remark 21. *Our presentation of the clipping lemma is slightly different than the original one* [36] *Lemma B.6], in that:*

- 1. We consider layered MDPs, while [36] consider general stationary MDPs where one state may be experienced at multiple different steps in [H]. Specifically, in a layered MDP, the occupancy distributions $\omega_{k,h}$ defined in [36] is only supported over $S_h \times A$. As a result, in the presentation here, we no longer need to sum over h – this is already captured in the sum over all s across all layers.
- 2. Our presentation here is in the context of multitask RL, which is with respect to a player $p \in [M]$, its corresponding MDP \mathcal{M}_p , and its policies used throughout the process $\{\pi^k(p)\}_{k=1}^K$. As a result, all quantities have p as subscripts.
- 3. Since every (state, action, step) in every MDP is trivially 1-transition optimal (see Definition B.3 of [36]), when applying Lemma B.6 to \mathcal{M}_p , we set $\alpha_{s,a,h} = 1$ and $g\check{a}p_p(s,a) = \frac{gap_{p,\min}}{2H} \vee \frac{gap_p(s,a)}{4H}$.

We are now ready to prove Theorems 5 and 6. MULTI-TASK-EULER's main regret theorems.

C.4.1 Proof of Theorem 5

Proof of Theorem 5 From Lemma 20 and Lemma 19, we have that when E happens,

$$\operatorname{Reg}(K) = \sum_{p=1}^{M} \operatorname{Reg}(K, p)$$

$$\leq \underbrace{\sum_{s,a} \sum_{k,p} \rho_{p}^{k}(s, a) \operatorname{clip}\left(B^{k, \operatorname{lead}}(s, a), \operatorname{g\check{a}p}_{p}(s, a)\right)}_{(A)} + \underbrace{H \sum_{s,a} \sum_{k,p} \rho_{p}^{k}(s, a) \operatorname{clip}\left(B^{k, \operatorname{fut}}(s, a), \frac{\operatorname{gap}_{p, \min}}{8SAH^{2}}\right)}_{(B)},$$

$$(51)$$

We bound each term separately. We can directly use Lemma 22 to bound term (B) as:

$$H\sum_{s,a}\sum_{k,p}\rho_p^k(s,a)\operatorname{clip}\left(B^{k,\operatorname{fut}}(s,a),\frac{\operatorname{gap}_{p,\min}}{8SAH^2}\right) \lesssim MH^4S^2A\left(\ln\left(\frac{MSAK}{\delta}\right)\right)^2.$$
 (52)

For term (A), we will group the sum by $(s, a) \in \mathcal{I}_{\epsilon}$ and $(s, a) \notin \mathcal{I}_{\epsilon}$ separately.

Case 1: $(s,a) \in \mathcal{I}_{\epsilon}$. In this case, we have that for all p, $g \check{a} p_p(s,a) = \frac{g a p_p(s,a)}{4H} \ge 24\epsilon$. We simplify the corresponding term as follows:

$$\begin{split} &\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k,p}\rho_{p}^{k}(s,a)\operatorname{clip}\left(B^{k,\operatorname{lead}}(s,a),\operatorname{g\check{a}p}_{p}(s,a)\right) \\ &\leq \sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k,p}\rho_{p}^{k}(s,a)\operatorname{clip}\left(H\wedge\left(5\epsilon+O\left(\sqrt{\frac{(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}])L(n^{k}(s,a))}{n^{k}(s,a)}}\right)\right), \frac{\operatorname{min}_{p}\operatorname{gap}_{p}(s,a)}{4H}\right) \\ &\leq \sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k,p}\rho_{p}^{k}(s,a)\left(H\wedge\operatorname{clip}\left(5\epsilon+O\left(\sqrt{\frac{(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}])L(n^{k}(s,a))}{n^{k}(s,a)}}\right), \frac{\operatorname{min}_{p}\operatorname{gap}_{p}(s,a)}{4H}\right)\right) \\ &\lesssim \sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k,p}\rho_{p}^{k}(s,a)\left(H\wedge\sqrt{\frac{(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}])L(n^{k}(s,a))}{n^{k}(s,a)}}\right) \end{split}$$

where the first inequality is from the definition of $B^{k,\text{lead}}$; the second inequality is from the basic fact that $\operatorname{clip}(A \land B, C) \leq A \land \operatorname{clip}(B, C)$; the third inequality uses Lemma 28 with $a_1 = 5\epsilon$, $a_2 = \sqrt{\frac{(1+\operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\pi^k}])L(n^k(s,a))}{n^k(s,a)}}$, and $\Delta = \frac{\min_p \operatorname{gap}_p(s,a)}{4H}$, along with the observation that $\operatorname{clip}(5\epsilon, \frac{\min_p \operatorname{gap}_p(s,a)}{16H}) = 0$, since for all $(s,a) \in \mathcal{I}_{\epsilon}$ and all $p \in [M]$, $\operatorname{gap}_p(s,a) \geq 96\epsilon H$. We now decompose the inner sum over $k \sum_{i=1}^{K} t_i$ to $\sum_{i=1}^{\tau(s,a)-1}$ and $\sum_{i=1}^{K} t_i$. The first part is

We now decompose the inner sum over k, $\sum_{k=1}^{K}$, to $\sum_{k=1}^{\tau(s,a)-1}$ and $\sum_{k=\tau(s,a)}^{K}$. The first part is bounded by:

$$\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=1}^{\tau_{p}(s,a)-1}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)\left(H\wedge\sqrt{\frac{(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}])L(n^{k}(s,a))}{n^{k}(s,a)}}\right) \leq \sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=1}^{\tau_{p}(s,a)-1}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)H\leq SAHN_{1}$$

which is $\lesssim MHSA \ln \left(\frac{SAK}{\delta}\right)$. For the second part,

$$\begin{split} &\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=\tau(s,a)}^{K}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)\left(H\wedge\sqrt{\frac{(1+\operatorname{var}_{\mathbb{P}_{p}}(\cdot|s,a)}[V_{p}^{\pi k}])L(n^{k}(s,a))}{n^{k}(s,a)}}\right) \\ &\lesssim &\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=\tau(s,a)}^{K}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)\sqrt{\frac{(1+\operatorname{var}_{\mathbb{P}_{p}}(\cdot|s,a)}[V_{p}^{\pi k}])L(\bar{n}^{k}(s,a))}{\bar{n}^{k}(s,a)}} \\ &\lesssim &\sqrt{\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=\tau(s,a)}^{K}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)\cdot\frac{L(\bar{n}^{k}(s,a))}{\bar{n}^{k}(s,a)}}{\bar{n}^{k}(s,a)}}\cdot\sqrt{\frac{\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=\tau(p)}^{K}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)\left(1+\operatorname{var}_{\mathbb{P}_{p}}(\cdot|s,a)[V_{p}^{\pi k}]\right)}{\bar{n}^{k}(s,a)}}, \end{split}$$

where the first inequality is by dropping the " $H \wedge$ " operator; the second inequality is by Cauchy-Schwarz.

We bound each factor as follows: for the first factor,

$$\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=\tau(s,a)}^{K}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)\cdot\frac{L(\bar{n}^{k}(s,a))}{\bar{n}^{k}(s,a)} = \sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=\tau(s,a)}^{K}\rho^{k}(s,a)\cdot\frac{L(\bar{n}^{k}(s,a))}{\bar{n}^{k}(s,a)}$$
$$\leq L(MK)\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=\tau(s,a)}^{K}\frac{\rho^{k}(s,a)}{\bar{n}^{k}(s,a)}$$
$$\leq \sum_{(s,a)\in\mathcal{I}_{\epsilon}}L(MK)\cdot\int_{1}^{\bar{n}^{K}(s,a)}\frac{1}{u}du$$
$$\leq |\mathcal{I}_{\epsilon}|L(MK)^{2} \lesssim |\mathcal{I}_{\epsilon}|\left(\ln\left(\frac{MSAK}{\delta}\right)\right)^{2},$$

where the first inequality is because L is monotonically increasing, and $\bar{n}^k(s,a) \leq MK$; the second inequality is from the observation that $\rho^k(s,a) \in [0,M]$, $\bar{n}^k(s,a) \geq 2M$, and $u \mapsto \frac{1}{u}$ is monotonically decreasing; the last two inequalities are by algebra.

For the second factor,

$$\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=1}^{K}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)\left(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right)\lesssim MKH + \sum_{p=1}^{M}\sum_{k=1}^{K}\sum_{(s,a)\in\mathcal{S}\times\mathcal{A}}\rho_{p}^{k}(s,a)\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]$$
$$\lesssim MKH + \sum_{p=1}^{M}\sum_{k=1}^{K}\operatorname{Var}\left[\sum_{h=1}^{H}r_{h,p}^{k} \mid \pi^{k}(p)\right]$$
$$\lesssim MKH^{2}.$$
(53)

where the first inequality is by the fact that ρ_p^k are probability distributions over every layer $h \in [H]$; the last two inequalities are by a law of total variance identity (see, e.g., [3] Equation (26)]). To summarize, the second part is at most

$$\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k=\tau(s,a)}^{K}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)\left(H\wedge\sqrt{\frac{(1+\operatorname{var}_{\mathbb{P}_{p}}(\cdot|s,a)[V_{p}^{\pi^{k}}])L(n^{k}(s,a))}{n^{k}(s,a)}}\right)\lesssim\sqrt{MKH^{2}|\mathcal{I}_{\epsilon}|}\ln\left(\frac{MSAK}{\delta}\right).$$

Combining the bounds for the first and the second parts, we have:

$$\sum_{(s,a)\in\mathcal{I}_{\epsilon}}\sum_{k,p}\rho_{p}^{k}(s,a)\operatorname{clip}\left(B^{k,\operatorname{lead}}(s,a),\operatorname{g\check{a}p}_{p}(s,a)\right) \lesssim \left(\sqrt{MKH^{2}|\mathcal{I}_{\epsilon}|} + MHSA\right)\ln\left(\frac{MSAK}{\delta}\right)$$

Case 2: $(s, a) \notin \mathcal{I}_{\epsilon}$. We simplify the corresponding term as follows:

$$\begin{split} &\sum_{(s,a)\notin\mathcal{I}_{\epsilon}}\sum_{k,p}\rho_{p}^{k}(s,a)\operatorname{clip}\left(B^{k,\operatorname{lead}}(s,a),\operatorname{g\check{a}p}_{p}(s,a)\right) \\ &\lesssim \sum_{(s,a)\notin\mathcal{I}_{\epsilon}}\sum_{k,p}\rho_{p}^{k}(s,a)\operatorname{clip}\left(H\wedge\left(\sqrt{\frac{\left(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right)L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}\right), \frac{\operatorname{g\check{a}p}_{p}(s,a)}{4H}\right) \\ &\lesssim \sum_{(s,a)\notin\mathcal{I}_{\epsilon}}\sum_{k,p}\left(H\wedge\sqrt{\frac{\left(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right)L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}\right) \end{split}$$

For each p and (s, a), we now decompose the inner sum over k, $\sum_{k=1}^{K}$, to $\sum_{k=1}^{\tau_p(s,a)-1}$ and $\sum_{k=\tau_p(s,a)}^{K}$. The first part is bounded by:

$$\sum_{(s,a)\notin\mathcal{I}_{\epsilon}} \sum_{p=1}^{M} \sum_{k=1}^{\tau_{p}(s,a)-1} \rho_{p}^{k}(s,a) \left(H \wedge \sqrt{\frac{(1 + \operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}])L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}} \right) \leq \sum_{(s,a)\notin\mathcal{I}_{\epsilon}} \sum_{p=1}^{M} \sum_{k=1}^{\tau_{p}(s,a)-1} \rho_{p}^{k}(s,a) H \leq MHSAN_{2},$$

which is $\lesssim MHSA \ln \left(\frac{MSAK}{\delta}\right)$.

For the second part,

$$\begin{split} &\sum_{(s,a)\notin\mathcal{I}_{\epsilon}}\sum_{p=1}^{M}\sum_{k=\tau_{p}(s,a)}^{K}\rho_{p}^{k}(s,a)\left(H\wedge\sqrt{\frac{\left(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right)L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}\right)\\ &\lesssim\sum_{(s,a)\notin\mathcal{I}_{\epsilon}}\sum_{p=1}^{M}\sum_{k=\tau_{p}(s,a)}^{K}\rho_{p}^{k}(s,a)\sqrt{\frac{\left(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right)L(\bar{n}_{p}^{k}(s,a))}{\bar{n}_{p}^{k}(s,a)}}}\\ &\leq\sqrt{\sum_{(s,a)\notin\mathcal{I}_{\epsilon}}\sum_{p=1}^{M}\sum_{k=\tau_{p}(s,a)}^{K}\rho_{p}^{k}(s,a)\cdot\frac{L(\bar{n}_{p}^{k}(s,a))}{\bar{n}_{p}^{k}(s,a)}}{\bar{n}_{p}^{k}(s,a)}\cdot\sqrt{\sum_{(s,a)\notin\mathcal{I}_{\epsilon}}\sum_{k=1}^{K}\sum_{p=1}^{M}\rho_{p}^{k}(s,a)\left(1+\operatorname{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right)}} \end{split}$$

We bound each factor as follows: for the first factor,

$$\sum_{(s,a)\notin\mathcal{I}_{\epsilon}} \sum_{p=1}^{M} \sum_{k=\tau_{p}(s,a)}^{K} \rho_{p}^{k}(s,a) \cdot \frac{L(\bar{n}^{k}(s,a))}{\bar{n}^{k}(s,a)} \leq L(K) \cdot \sum_{(s,a)\notin\mathcal{I}_{\epsilon}} \sum_{p=1}^{M} \sum_{k=\tau_{p}(s,a)}^{K} \frac{\rho_{p}^{k}(s,a)}{\bar{n}^{k}(s,a)}$$
$$\leq L(K) \cdot \sum_{(s,a)\notin\mathcal{I}_{\epsilon}} \sum_{p=1}^{M} \int_{1}^{\bar{n}_{p}^{K}(s,a)} \frac{1}{u} du$$
$$\leq \left|\mathcal{I}_{\epsilon}^{C}\right| ML(K)^{2} \leq \left|\mathcal{I}_{\epsilon}^{C}\right| M\left(\ln\left(\frac{MSAK}{\delta}\right)\right)^{2}.$$

where the first inequality is because L is monotonically increasing, and $\bar{n}_p^k(s,a) \leq K$; the second inequality is from the observation that $\rho^k(s,a) \in [0,1]$, $\bar{n}^k(s,a) \geq 2$, and $u \mapsto \frac{1}{u}$ is monotonically decreasing; the last two inequalities are by algebra.

The second factor is again bounded by (53). Therefore, the second part of the sum is at most

$$\begin{split} &\sum_{(s,a)\notin\mathcal{I}_{\epsilon}}\sum_{p=1}^{M}\sum_{k=\tau_{p}(s,a)}^{K}\rho_{p}^{k}(s,a)\left(H\wedge\sqrt{\frac{\left(1+\mathrm{var}_{\mathbb{P}_{p}(\cdot|s,a)}[V_{p}^{\pi^{k}}]\right)L(n_{p}^{k}(s,a))}{n_{p}^{k}(s,a)}}\right) \\ &\leq \left(M\sqrt{KH^{2}|\mathcal{I}_{\epsilon}^{C}|}+MHSA\right)\ln\left(\frac{MSAK}{\delta}\right). \end{split}$$

Combining the bounds for the first and the second parts, we have:

$$\sum_{(s,a)\notin\mathcal{I}_{\epsilon}}\sum_{k,p}\rho_{p}^{k}(s,a)\operatorname{clip}\left(B^{k,\operatorname{lead}}(s,a),\operatorname{g\check{a}p}_{p}(s,a)\right) \lesssim \left(M\sqrt{KH^{2}\left|\mathcal{I}_{\epsilon}^{C}\right|} + MHSA\right)\ln\left(\frac{MSAK}{\delta}\right)$$

Now, combining the bounds for cases 1 and 2, we have that

$$(A) \le \left(\sqrt{MKH^2|\mathcal{I}_{\epsilon}|} + M\sqrt{KH^2|\mathcal{I}_{\epsilon}^C|} + MHSA\right) \cdot \ln\left(\frac{MSAK}{\delta}\right).$$
(54)

In conclusion, by the regret decomposition Equation (51), and Equations (54) and (52), we have:

$$\operatorname{Reg}(K) \leq \left(\sqrt{MH^2 |\mathcal{I}_{\epsilon}| K} + M\sqrt{H^2 |\mathcal{I}_{\epsilon}^C| K} + MH^4 S^2 A \ln\left(\frac{MSAK}{\delta}\right)\right) \ln\left(\frac{MSAK}{\delta}\right).$$

C.4.2 Proof of Theorem 6

Proof of Theorem [6] From Lemma [20], we have that when E happens,

$$\operatorname{Reg}(K) = \sum_{p=1}^{M} \operatorname{Reg}(K, p)$$

$$\leq \underbrace{\sum_{s,a} \sum_{k,p} \rho_p^k(s, a) \operatorname{clip}\left(B^{k, \operatorname{lead}}(s, a), \operatorname{g\check{a}p}_p(s, a)\right)}_{(A)} + \underbrace{H \sum_{s,a} \sum_{k,p} \rho_p^k(s, a) \operatorname{clip}\left(B^{k, \operatorname{fut}}(s, a), \frac{\operatorname{gap}_{p, \min}}{8SAH^2}\right)}_{(B)}$$

We focus on each term separately. We directly use Lemma 22 to bound term (B) as:

$$H\sum_{s,a}\sum_{k,p}\rho_p^k(s,a)\operatorname{clip}\left(B^{k,\operatorname{fut}}(s,a),\frac{\operatorname{gap}_{p,\min}}{8SAH^2}\right) \lesssim MH^4S^2A\ln\left(\frac{MSAK}{\delta}\right) \cdot \ln\frac{MSA}{\operatorname{gap}_{\min}}.$$
 (55)

For the (s, a)-th term in term (A), we will consider the cases of $(s, a) \in \mathcal{I}_{\epsilon}$ and $(s, a) \notin \mathcal{I}_{\epsilon}$ separately.

Case 1: $(s, a) \in \mathcal{I}_{\epsilon}$. In this case, we have that for all p, $g\check{a}p_p(s, a) = \frac{gap_p(s, a)}{4H} \ge 24\epsilon$. We simplify the corresponding term as follows:

$$\begin{split} &\sum_{k,p} \rho_p^k(s,a) \operatorname{clip} \left(B^{k,\operatorname{lead}}(s,a), \operatorname{g\check{a}p}_p(s,a) \right) \\ &\leq \sum_{k=1}^K \sum_{p=1}^M \rho_p^k(s,a) \operatorname{clip} \left(H \wedge \left(5\epsilon + O\left(\sqrt{\frac{(1 + \operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\pi^k}])L(n^k(s,a))}{n^k(s,a)}}\right) \right), \frac{\min_p \operatorname{gap}_p(s,a)}{4H} \right) \\ &\leq \sum_{k=1}^K \rho^k(s,a) \operatorname{clip} \left(H \wedge \left(5\epsilon + O\left(\sqrt{\frac{H^2L(n^k(s,a))}{n^k(s,a)}}\right) \right), \frac{\min_p \operatorname{gap}_p(s,a)}{4H} \right) \\ &\leq \sum_{k=1}^k \rho^k(s,a) \left(H \wedge \operatorname{clip} \left(5\epsilon + O\left(\sqrt{\frac{H^2L(n^k(s,a))}{n^k(s,a)}}\right), \frac{\min_p \operatorname{gap}_p(s,a)}{4H} \right) \right) \right) \\ &\lesssim \sum_{k=1}^K \rho^k(s,a) \left(H \wedge \operatorname{clip} \left(\sqrt{\frac{H^2L(n^k(s,a))}{n^k(s,a)}}, \frac{\min_p \operatorname{gap}_p(s,a)}{16H} \right) \right) \right) \end{split}$$

where the first inequality is by the definition of $B^{k,\text{lead}}$; the second inequality is from that $\operatorname{var}_{\mathbb{P}_p(\cdot|s,a)}[V_p^{\pi^k}] \leq H^2$; the third inequality is from that $\operatorname{clip}(A \wedge B, C) \leq A \wedge \operatorname{clip}(B, C)$; the third inequality uses Lemma 28 with $a_1 = 5\epsilon$, $a_2 = \sqrt{\frac{H^2L(n^k(s,a))}{n^k(s,a)}}$, and $\Delta = \frac{\min_p \operatorname{gap}_p(s,a)}{4H}$, along with the observation that $\operatorname{clip}(5\epsilon, \frac{\min_p \operatorname{gap}_p(s,a)}{16H}) = 0$, since for all $(s,a) \in \mathcal{I}_\epsilon$ and all $p \in [M]$, $\operatorname{gap}_p(s,a) \geq 96\epsilon H$.

We now decompose the inner sum over k, $\sum_{k=1}^{K}$, to $\sum_{k=1}^{\tau(s,a)-1}$ and $\sum_{k=\tau(s,a)}^{K}$. The first part's contribution is at most $N_1 \cdot H \lesssim MH \ln \left(\frac{SAK}{\delta}\right)$. For the second part, its contribution is at most:

$$\begin{split} &\sum_{k=\tau(s,a)}^{K} \rho^{k}(s,a) \left(H \wedge \operatorname{clip}\left(\sqrt{\frac{H^{2}L(n^{k}(s,a))}{n^{k}(s,a)}}, \frac{\min_{p} \operatorname{gap}_{p}(s,a)}{16H}\right) \right) \\ &\lesssim MH + \int_{1}^{\bar{n}^{K}(s,a)} \left(H \wedge \operatorname{clip}\left(\sqrt{\frac{H^{2}L(u)}{u}}, \frac{\min_{p} \operatorname{gap}_{p}(s,a)}{16H}\right) \right) du \\ &\lesssim MH + \frac{H^{3}}{\min_{p} \operatorname{gap}_{p}(s,a)} \ln\left(\frac{MSAK}{\delta}\right) \end{split}$$

where the second inequality is from Lemma 29 with $f_{\max} = H$, $C = H^2$, $\Delta = \frac{\min_p \operatorname{gap}_p(s,a)}{16H}$, $N = MSA, \xi = \delta, \Gamma = 1, n = \overline{n}^K(s, a) \leq K$. In summary, for all $(s, a) \in \mathcal{I}_{\epsilon}$,

$$\sum_{k,p} \rho_p^k(s,a) \operatorname{clip}\left(B^{k,\operatorname{lead}}(s,a), \operatorname{g\check{a}p}_p(s,a)\right) \leq \left(MH + \frac{H^3}{\min_p \operatorname{gap}_p(s,a)}\right) \ln\left(\frac{MSAK}{\delta}\right).$$

Case 2: $(s, a) \notin \mathcal{I}_{\epsilon}$. In this case, for each $p \in [M]$, we simplify the corresponding term as follows:

$$\sum_{k} \rho_p^k(s, a) \operatorname{clip} \left(B^{k, \operatorname{lead}}(s, a), \operatorname{g\check{a}p}_p(s, a) \right)$$
$$\lesssim \sum_{k=1}^{K} \rho_p^k(s, a) \left(H \wedge \operatorname{clip} \left(\sqrt{\frac{H^2 L(n_p^k(s, a))}{n_p^k(s, a)}}, \frac{\operatorname{g\check{a}p}_p(s, a)}{16H} \right) \right)$$

We now decompose the inner sum over k, $\sum_{k=1}^{K}$, to $\sum_{k=1}^{\tau_p(s,a)-1}$ and $\sum_{k=\tau_p(s,a)}^{K}$. The first part's contribution is at most $N_2 \cdot H \lesssim H \ln\left(\frac{MSAK}{\delta}\right)$.

For the second part, its contribution is at most:

$$\begin{split} &\sum_{k=\tau_p(s,a)}^{K} \rho_p^k(s,a) \left(H \wedge \operatorname{clip}\left(\sqrt{\frac{H^2 L(n^k(s,a))}{n^k(s,a)}}, \frac{\operatorname{g\check{a}p}_p(s,a)}{16H}\right) \right) \\ &\lesssim H + \int_1^{\bar{n}_p^K(s,a)} \left(H \wedge \operatorname{clip}\left(\sqrt{\frac{H^2 L(u)}{u}}, \frac{\operatorname{g\check{a}p}_p(s,a)}{16H}\right) \right) du \\ &\lesssim H + \frac{H^3}{\operatorname{g\check{a}p}_p(s,a)} \ln\left(\frac{MSAK}{\delta}\right) \end{split}$$

where the second inequality is from Lemma 29 with $f_{\max} = H$, $C = H^2$, $\Delta = \frac{g\check{a}_{P_p}(s,a)}{16H}$, N = MSA, $\xi = \delta$, $\Gamma = 1$, $n = \bar{n}_p^K(s,a) \leq K$. In summary, for any $(s,a) \in \mathcal{I}_{\epsilon}^C$ and $p \in [M]$,

$$\sum_{k} \rho_p^k(s, a) \operatorname{clip}\left(B^{k, \operatorname{lead}}(s, a), \operatorname{g\check{a}p}_p(s, a)\right) \lesssim \left(H + \frac{H^3}{\operatorname{g\check{a}p}_p(s, a)}\right) \ln\left(\frac{MSAK}{\delta}\right),$$

summing over p, we get:

$$\sum_{k,p} \rho_p^k(s,a) \operatorname{clip}\left(B^{k,\operatorname{lead}}(s,a), \operatorname{g\check{a}p}_p(s,a)\right) \lesssim \left(MH + \sum_{p=1}^M \frac{H^3}{\operatorname{g\check{a}p}_p(s,a)}\right) \ln\left(\frac{MSAK}{\delta}\right),$$

In summary, combining the regret bounds of cases 1 and 2 for term (A), along with Equation (55) for term (B), and observe that $g \check{a} p_p(s, a) = g a p_{p,\min}$ if $(s, a) \in Z_{p,\text{opt}}$, and $g \check{a} p_p(s, a) = g a p_p(s, a)$ otherwise, we have that on event E, MULTI-TASK-EULER satisfies:

$$\operatorname{Reg}(K) \lesssim \ln\left(\frac{MSAK}{\delta}\right) \left(\sum_{p \in [M]} \left(\sum_{(s,a) \in \mathbb{Z}_{p, \text{opt}}} \frac{H^3}{\operatorname{gap}_{p, \min}} + \sum_{(s,a) \in (\mathcal{I}_{\epsilon} \cup \mathbb{Z}_{p, \text{opt}})^C} \frac{H^3}{\operatorname{gap}_p(s,a)}\right) + \sum_{(s,a) \in \mathcal{I}_{\epsilon}} \frac{H^3}{\min_p \operatorname{gap}_p(s,a)}\right) + \ln\left(\frac{MSAK}{\delta}\right) \cdot MS^2 A H^4 \ln \frac{MSA}{\operatorname{gap}_{\min}}. \Box$$

Lemma 22 (Bounding the lower order terms). If E happens, then

$$\sum_{s,a} \sum_{k,p} \rho_p^k(s,a) \operatorname{clip}\left(B^{k,\operatorname{fut}}(s,a), \frac{\operatorname{gap}_{p,\min}}{8SAH^2}\right) \lesssim MH^3 S^2 A \ln\left(\frac{MSAK}{\delta}\right) \left(\ln\left(\frac{MSAK}{\delta}\right) \wedge \ln\left(\frac{MSA}{\operatorname{gap}_{\min}}\right)\right)$$

Proof. We expand the left hand side using the definition of $B^{k,\text{fut}}$, and the fact that $\text{gap}_{p,\min} \geq \text{gap}_{\min}$:

$$\sum_{k=1}^{K} \rho_p^k(s,a) \operatorname{clip}\left(B^{k,\operatorname{fut}}(s,a), \frac{\operatorname{gap}_{p,\min}}{8SAH^2}\right)$$
(56)

$$\lesssim \sum_{k=1}^{K} \rho_p^k(s, a) \left(H^3 \wedge \operatorname{clip}\left(\frac{H^3 SL(n_p^k(s, a))}{n_p^k(s, a)}, \frac{\operatorname{gap}_{\min}}{8SAH^2}\right) \right)$$
(57)

We now decompose the sum $\sum_{k=1}^{K}$ to $\sum_{k=1}^{\tau_p(s,a)-1}$ and $\sum_{k=\tau_p(s,a)}^{K}$. The first part can be bounded by

$$\sum_{k=1}^{\tau_p(s,a)-1} \rho_p^k(s,a) \left(H^3 \wedge \operatorname{clip}\left(\frac{H^3 SL(n_p^k(s,a))}{n_p^k(s,a)}, \frac{\operatorname{gap}_{\min}}{8SAH^2}\right) \right) \le \sum_{k=1}^{\tau_p(s,a)-1} H^3 \rho_p^k(s,a) \le H^3 N_2,$$

which is at most $O\left(H^3 \cdot \ln\left(\frac{MSAK}{\delta}\right)\right)$. For the second part, it can be bounded by:

$$\begin{split} &\sum_{k=\tau_p(s,a)}^{K} \rho_p^k(s,a) \left(H^3 \wedge \operatorname{clip}\left(\frac{H^3 SL(n_p^k(s,a))}{n_p^k(s,a)}, \frac{\operatorname{gap}_{\min}}{8SAH^2}\right) \right) \\ &\leq H^3 \cdot 1 + \int_1^{\bar{n}_p^K(s,a)} \left(H^3 \wedge \operatorname{clip}\left(\frac{H^3 SL(u)}{u}, \frac{\operatorname{gap}_{\min}}{8SAH^2}\right) \right) du \\ &\lesssim H^3 + H^3 \ln\left(\frac{MSA}{\delta}\right) + H^3 S \ln\left(\frac{MSAK}{\delta}\right) \left(\ln\left(\frac{MSAK}{\delta}\right) \wedge \ln\left(\frac{MHSA}{\operatorname{gap}_{\min}}\right) \right) \end{split}$$

where the second inequality is from Lemma 29 with $f_{\text{max}} = H^3$, $C = H^3S$, $\Delta = \frac{\text{gap}_{\text{min}}}{8SAH^2}$, N = MSA, $\xi = \delta$, $\Gamma = 1$, $n = \bar{n}_p^K(s, a) \leq K$. In addition, observe that $H \leq S$ by our layered MDP assumption, we have

$$\sum_{k} \rho_p^k(s, a) \operatorname{clip}\left(B^{k, \operatorname{lead}}(s, a), \frac{\operatorname{gap}_{\min}}{8SAH^2}\right) \lesssim H^3 S \ln\left(\frac{MSAK}{\delta}\right) \left(\ln\left(\frac{MSAK}{\delta}\right) \wedge \ln\left(\frac{MSA}{\operatorname{gap}_{\min}}\right)\right)$$

Summing over $s \in S$, $a \in A$, and $p \in [M]$, we get

$$\sum_{s,a} \sum_{k,p} \rho_p^k(s,a) \operatorname{clip}\left(B^{k,\operatorname{lead}}(s,a), \frac{\operatorname{gap}_{\min}}{8SAH^2}\right) \lesssim MH^3 S^2 A \ln\left(\frac{MSAK}{\delta}\right) \left(\ln\left(\frac{MSAK}{\delta}\right) \wedge \ln\left(\frac{MSA}{\operatorname{gap}_{\min}}\right)\right)$$

C.5 Miscellaneous lemmas

This subsection collects a few miscellaneous lemmas used throughout the upper bound proofs.

Lemma 23 (Bias-variance decomposition). For any random variable X with $\mathbb{E}[X] = \mu \in \mathbb{R}$, and any $m \in \mathbb{R}$, $\mathbb{E}[(X - m)^2] = \mathbb{E}[(X - \mu)^2] + (\mu - m)^2$.

Lemma 24 ([36], Lemma F.5). For random variables X and Y, $\left|\sqrt{\operatorname{var}[X]} - \sqrt{\operatorname{var}[Y]}\right| \leq \sqrt{\mathbb{E}\left[(X-Y)^2\right]}$.

Lemma 25. Suppose distributions P and Q are supported over [0, B], and $||P - Q||_1 \le \epsilon \le 2$. Then: $|\mathbb{F}_{X \to P}[X] = \mathbb{F}_{X \to Q}[X]| \le B\epsilon$

$$\left| \operatorname{var}_{X \sim P}[X] - \operatorname{var}_{X \sim Q}[X] \right| \le B\epsilon,$$

$$\left| \operatorname{var}_{X \sim P}[X] - \operatorname{var}_{X \sim Q}[X] \right| \le 3B^2\epsilon.$$

Proof. First,

$$\left|\mathbb{E}_{X \sim P}[X] - \mathbb{E}_{X \sim Q}[X]\right| = \left|\int_{0}^{B} x(p_{X}(x) - q_{X}(x))dx\right| \le \int_{0}^{B} |x| \left|p_{X}(x) - q_{X}(x)\right| dx \le B \|P - Q\|_{1} \le B\epsilon.$$

Second, observe that

$$\left|\mathbb{E}_{X\sim P}[X^2] - \mathbb{E}_{X\sim Q}[X^2]\right| \le B^2 \epsilon.$$

Meanwhile,

$$\left| \left(\mathbb{E}_{X \sim P}[X] \right)^2 - \left(\mathbb{E}_{X \sim Q}[X] \right)^2 \right| \le \left| \mathbb{E}_{X \sim P}[X] - \mathbb{E}_{X \sim Q}[X] \right| \left| \left| \mathbb{E}_{X \sim P}[X] + \mathbb{E}_{X \sim Q}[X] \right| \le 2B \cdot B\epsilon = 2B^2 \epsilon.$$

Combining the above, we have

$$\operatorname{var}_{X \sim P}[X] - \operatorname{var}_{X \sim Q}[X] \le 3B^2 \epsilon.$$

Lemma 26. For $A, B, C, D, E, F \ge 0$:

1. If
$$A \leq B + C\sqrt{A}$$
, then $\sqrt{A} \leq \sqrt{B} + C$.
2. If $|D - E| \leq \sqrt{EF} + F$, then we have $\left|\sqrt{D} - \sqrt{E}\right| \leq 2\sqrt{F}$.

- *Proof.* 1. The roots of $x^2 Cx B = 0$ are $\frac{C \pm \sqrt{C^2 + 4B}}{2}$, and therefore A must satisfy $\sqrt{A} \le \frac{C + \sqrt{C^2 + 4B}}{2} \le \frac{C + C + 2\sqrt{B}}{2} = C + \sqrt{B}$.
 - 2. First, $D E \le |D E| \le \sqrt{EF} + F$; this implies that $D \le E + 2\sqrt{EF} + F$, and therefore $\sqrt{D} \le \sqrt{E} + \sqrt{F}$.

On the other hand, $E \leq D+F+\sqrt{EF}$; therefore, applying item 1 with A = E, B = D+F, and $F = \sqrt{E}$, we have $\sqrt{E} \leq \sqrt{D+F} + \sqrt{F} \leq \sqrt{D} + 2\sqrt{F}$.

Lemma 27. For $a \ge 0$, $1 \land (a + \sqrt{a}) \le 1 \land 2\sqrt{a}$.

Proof. We consider the cases of $a \ge 1$ and a < 1 respectively. If $a \ge 1$, LHS = 1 = RHS. Otherwise, $a \le 1$; in this case, LHS = $1 \land (a + \sqrt{a}) \le 1 \land (\sqrt{a} + \sqrt{a}) = RHS.$

Lemma 28 (Special case of [36], Lemma B.5). For $a_1, a_2, \Delta \ge 0$, $\operatorname{clip}(a_1 + a_2, \Delta) \le 2 \operatorname{clip}(a_1, \Delta/4) + 2 \operatorname{clip}(a_2, \Delta/4)$.

Lemma 29 (Integral calculation, extracted from [36], Lemma B.9). Let $f(u) \leq \min(f_{\max}, \operatorname{clip}(g(u), \Delta))$, where $\Delta \in [0, \Gamma]$ and $\Gamma \geq 1$, and g(u) is nonincreasing. Let $N \geq 1$ and $\xi \in (0, \frac{1}{2})$. Then:

_		

1. If
$$g(u) \lesssim \sqrt{\frac{C \log \frac{Nu}{\xi}}{u}}$$
 for some $C > 0$ such that $\ln C \lesssim \ln N$, then

$$\int_{\Gamma}^{n} f(u/4) du \lesssim \sqrt{Cn \ln \frac{Nn}{\xi}} \wedge \frac{C}{\Delta} \ln \left(\frac{Nn}{\xi}\right).$$
2. If $g(u) \lesssim \frac{C \ln \frac{Nu}{\xi}}{\xi}$ for some $C > 0$ such that $\ln C \lesssim \ln N$, then

If
$$g(u) \lesssim \frac{\Box m - \varepsilon}{u}$$
 for some $C > 0$ such that $\ln C \lesssim \ln N$, then
$$\int_{\Gamma}^{n} f(u/4) du \lesssim f_{\max} \ln \frac{N}{\xi} + C \ln \frac{Nn}{\xi} \cdot \left(\ln \frac{Nn}{\xi} \wedge \ln \frac{N\Gamma}{\Delta} \right).$$

D Proof of the Lower Bounds

D.1 Auxiliary Lemmas

Lemma 30 (Regret decomposition, [36], Section H.2). *For any MPERL problem instance and any algorithm, we have*

$$\mathbb{E}\left[\operatorname{Reg}(K)\right] \ge \sum_{p=1}^{M} \sum_{(s,a)\in\mathcal{S}_1\times\mathcal{A}} \mathbb{E}\left[n_p^{K+1}(s,a)\right] \operatorname{gap}_p(s,a),\tag{58}$$

where we recall that S_1 is the subset of state space where the initial state distribution p_0 is supported on, $n_p^{K+1}(s, a)$ is the number of visits of (s, a) by player p at the beginning of the (K+1)-th episode (after the first K episodes). Furthermore, for any $(s, a) \in S_1 \times A$, we have

$$\sum_{p=1}^{M} \mathbb{E}\left[n_p^{K+1}(s,a)\right] \operatorname{gap}_p(s,a) \ge \mathbb{E}\left[n^{K+1}(s,a)\right] \left(\min_{p \in [M]} \operatorname{gap}_p(s,a)\right),\tag{59}$$

where we recall that $n^{K+1}(s, a) = \sum_{p=1}^{M} n_p^{K+1}(s, a)$.

Proof. Eq. (59) follows straightforwardly from the fact that for every $(s, a, p) \in S_1 \times A \times [M]$, $\min_{p' \in [M]} \operatorname{gap}_{p'}(s, a) \leq \operatorname{gap}_p(s, a)$.

We now prove Eq. (58). Let π_p^k denote $\pi^k(p)$. We have

$$\mathbb{E}\left[\operatorname{Reg}(K)\right] = \mathbb{E}\left[\sum_{p=1}^{M} \sum_{k=1}^{K} \sum_{s \in S_{1}} p_{0}(s_{1,p}^{k} = s) \left(V_{p}^{\star}(s) - V_{p}^{\pi_{p}^{k}}(s)\right)\right]$$

$$\geq \mathbb{E}\left[\sum_{p=1}^{M} \sum_{k=1}^{K} \sum_{s \in S_{1}} p_{0}(s_{1,p}^{k} = s) \left(V_{p}^{\star}(s) - Q_{p}^{\star}(s, \pi_{p}^{k}(s))\right)\right]$$

$$= \mathbb{E}\left[\sum_{p=1}^{M} \sum_{k=1}^{K} \sum_{s \in S_{1}} p_{0}(s) \operatorname{gap}_{p}(s, \pi_{p}^{k}(s))\right]$$

$$= \mathbb{E}\left[\sum_{p=1}^{M} \sum_{k=1}^{K} \sum_{s \in S_{1}} \mathbf{1}\left(s_{1,p}^{k} = s\right) \operatorname{gap}_{p}(s, \pi_{p}^{k}(s))\right]$$

$$= \mathbb{E}\left[\sum_{p=1}^{M} \sum_{k=1}^{K} \sum_{s \in S_{1} \times \mathcal{A}} \mathbf{1}\left(s_{1,p}^{k}, \pi_{p}^{k}(s) = (s, a)\right) \operatorname{gap}_{p}(s, a)\right]$$

$$= \sum_{p=1}^{M} \sum_{(s,a) \in S_{1} \times \mathcal{A}} \mathbb{E}\left[n_{p}^{K}(s, a)\right] \operatorname{gap}_{p}(s, a)$$

(60)

where the first equality is from the definition of collective regret; the first inequality is from the simple fact that $V_p^{\pi}(s) = Q_p^{\pi}(s, \pi(s)) \leq Q_p^{\star}(s, \pi(s))$ for any policy π ; the second equality is from the definition of suboptimality gaps; and the third equality is from the basic observation that $s_{1,p}^k \sim p_0$.

Lemma 31 (Divergence decomposition [23] [44]). For two MPERL problem instances, \mathfrak{M} and \mathfrak{M}' , which only differ in the transition probabilities $\{\mathbb{P}_p(\cdot | s, a)\}_{p \in [M], (s,a) \in S \times A}$, and for a fixed algorithm, let $\mathbb{P}_{\mathfrak{M}}$ and $\mathbb{P}_{\mathfrak{M}'}$ be the probability measures on the outcomes of running the algorithm on \mathfrak{M} and \mathfrak{M}' , respectively. Then,

$$\mathrm{KL}(\mathbb{P}_{\mathfrak{M}},\mathbb{P}_{\mathfrak{M}'}) = \sum_{p=1}^{M} \sum_{(s,a)\in\mathcal{S}\times\mathcal{A}} \mathbb{E}_{\mathfrak{M}}\left[n_p^{K+1}(s,a)\right] \mathrm{KL}\left(\mathbb{P}_p^{\mathfrak{M}}(\cdot \mid s,a), \mathbb{P}_p^{\mathfrak{M}'}(\cdot \mid s,a)\right),$$

where $\mathbb{P}_p^{\mathfrak{M}}(\cdot \mid s, a)$ and $\mathbb{P}_p^{\mathfrak{M}'}(\cdot \mid s, a)$ are the transition probabilities of the problem instance \mathfrak{M} and \mathfrak{M}' , respectively.

Lemma 32 (Bretagnolle-Huber inequality, [23], Theorem 14.2). Let \mathbb{P} and \mathbb{Q} be two distributions on the same measurable space, and A be an event. Then,

$$\mathbb{P}(A) + \mathbb{Q}(A^C) \ge \frac{1}{2} \exp\left(-\operatorname{KL}(\mathbb{P}, \mathbb{Q})\right)$$

Lemma 33 (see, e.g., [43], Lemma 25). For any $x, y \in [\frac{1}{4}, \frac{3}{4}]$, KL $(\text{Ber}(x), \text{Ber}(y)) \leq 3(x-y)^2$. Lemma 34. Let X be a Binomial random variable and $X \sim \text{Bin}(n, p)$, where $n \geq \frac{1}{p}$. Then,

$$\mathbb{E}\left[X^{\frac{3}{2}}\right] \le 2(np)^{\frac{3}{2}}.$$

Proof. Let $Y = X^2$, and $f(y) = y^{\frac{3}{4}}$. We have $\mathbb{E}[Y] = \mathbb{E}[X^2] = \operatorname{var}[X] + \mathbb{E}[X]^2 = (np)^2 + np(1-p) \le (np)^2 + np \le 2(np)^2$, where the last inequality follows from the assumption that $n \ge \frac{1}{p}$. By Jensen's inequality, we have $\mathbb{E}\left[X^{\frac{3}{2}}\right] = \mathbb{E}\left[f(Y)\right] \le f\left(\mathbb{E}[Y]\right) \le \left(2n^2p^2\right)^{\frac{3}{4}} \le 2(np)^{\frac{3}{2}}$. \Box

D.2 Gap independent lower bounds

Theorem 35 (Restatement of Theorem 7). For any $A \ge 2$, $H \ge 2$, $S \ge 4H$, $K \ge SA$, $M \in \mathbb{N}$, and $l, l^C \in \mathbb{N}$ such that $l + l^C = SA$ and $l \le SA - 4(S + HA)$, there exists some ϵ such that for any algorithm Alg, there exists an ϵ -MPERL problem instance with S states, A actions, M players and an episode length of H such that $\left|\mathcal{I}_{\frac{\epsilon}{192H}}\right| \ge l$, and

$$\mathbb{E}\left[\mathrm{Reg}_{\mathrm{Alg}}(K)\right] \geq \Omega\left(M\sqrt{H^2l^CK} + \sqrt{MH^2lK}\right).$$

Proof. The construction and techniques in this proof are inspired by [43, Section E.1] and [36]. Fix any algorithm Alg; we consider two cases:

- 1. $l > M l^C;$
- 2. $Ml^C \ge l$.

Case 1: $l > Ml^C$. Let $S_1 = S - 2(H - 1)$, and $b = \lceil \frac{l}{S_1} \rceil \ge 1$. Let $\Delta = \sqrt{\frac{l+1}{384MK}}$, and let $\epsilon = \frac{1}{2}H\Delta$. We note that under the assumption that $K \ge SA$, and the observation that $l \le SA$, we have $\Delta \le \frac{1}{4}$. We define $(b + 1)^{S_1} \epsilon$ -MPERL problem instances, each indexed by an element in $[b + 1]^{S_1}$. It suffices to show that, on at least one of these problem instances, $\mathbb{E}\left[\operatorname{Reg}_{Alg}(K)\right] \ge \Omega\left(\sqrt{MH^2lK}\right)$.

Construction. For $\mathbf{a} = (a_1, \dots, a_{S_1}) \in [b+1]^{S_1}$, we define the following ϵ -MPERL problem instance, $\mathfrak{M}(\mathbf{a}) = \{\mathcal{M}_p\}_{p=1}^M$, with S states, A actions, and an episode length of H, such that for each $p \in [M]$, \mathcal{M}_p is constructed as follows:

- $S_1 = [S_1]$, and p_0 is a uniform distribution over the states in S_1 .
- For $h \in [2, H]$, $S_h = \{S_1 + 2h 3, S_1 + 2h 2\}$.
- $\mathcal{A} = [A].$
- For each $(s,a) \in \mathcal{S} \times \mathcal{A}$, the reward distribution $r_p(s,a)$ is a Bernoulli distribution, $Ber(R_p(s,a))$, and we will specify $R_p(s,a)$ subsequently.
- For each state $s \in [S_1]$,

$$\mathbb{P}_p(S_1 + 1 \mid s, a) = \begin{cases} \frac{1}{2} + \Delta, & \text{if } a = a_s; \\ \frac{1}{2}, & \text{if } a \in [b+1] \setminus \{a_s\}; \\ 0, & \text{if } a \notin [b+1]; \end{cases}$$

and for each $a \in \mathcal{A}$, $\mathbb{P}_p(S_1 + 2 \mid s, a) = 1 - \mathbb{P}_p(S_1 + 1 \mid s, a)$, and $R_p(s, a) = 0$.

• For $h \in [2, H]$, and $a \in \mathcal{A}$, let

$$- \mathbb{P}_p \left(S_1 + 2h - 1 \mid S_1 + 2h - 3, a \right) = 1, \mathbb{P}_p \left(S_1 + 2h \mid S_1 + 2h - 3, a \right) = 0, \text{ and } R_p (S_1 + 2h - 3, a) = 1.$$

$$- \mathbb{P}_p \left(S_1 + 2h \mid S_1 + 2h - 2, a \right) = 0, \mathbb{P}_p \left(S_1 + 2h - 1 \mid S_1 + 2h - 2, a \right) = 1, \text{ and } R_p (S_1 + 2h - 2, a) = 0.$$

It can be easily verified that $\mathfrak{M}(\mathbf{a}) = \{\mathcal{M}_p\}_{p=1}^M$ is a 0-MPERL problem instance, and hence an ϵ -MPERL problem instance—the reward distributions and the transition probabilities are the same for all players, i.e., for every $p, q \in [M]$, and every $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\left|R_p(s,a) - R_q(s,a)\right| = 0 \le \epsilon, \quad \left|\mathbb{P}_p(\cdot \mid s,a) - \mathbb{P}_q(\cdot \mid s,a)\right| = 0 \le \frac{\epsilon}{H}.$$

Suboptimality gaps. We now calculate the suboptimality gaps of the state-action pairs in the above MDPs. For each $p \in [M]$ and each $(s, a) \in S \times A$,

$$gap_p(s, a) = V_p^{\star}(s) - Q_p^{\star}(s, a) = \max_{a'} Q_p^{\star}(s, a') - Q_p^{\star}(s, a)$$

In $\mathfrak{M}(\mathbf{a})$, it can be easily observed that for every $p \in [M]$, and every $(s, a) \in (\mathcal{S} \setminus \mathcal{S}_1) \times \mathcal{A}$, $\operatorname{gap}_p(s, a) = 0$. Now, for every $p \in [M]$, $(s, a) \in \mathcal{S}_1 \times \mathcal{A}$, we have

$$gap_{p}(s,a) = \max_{a'} Q_{p}^{\star}(s,a') - Q_{p}^{\star}(s,a) = (H-1) \left(\max_{a'} \mathbb{P}_{p}(S_{1}+1 \mid s,a') - \mathbb{P}_{p}(S_{1}+1 \mid s,a) \right).$$

It follows that, for every $p \in [M]$ and every state $s \in [S_1]$,

$$\operatorname{gap}_p(s,a) = \begin{cases} 0, & \text{if } a = a_s;\\ (H-1)\Delta, & \text{if } a \in [b+1] \setminus a_s;\\ (H-1)\left(\frac{1}{2} + \Delta\right), & \text{if } a \notin [b+1]. \end{cases}$$

Subpar state-action pairs. It can be verified that in $\mathfrak{M}(\mathbf{a})$, $\left|\mathcal{I}_{\frac{\epsilon}{192H}}\right| \geq l$. Indeed, since $(H-1)\Delta = (H-1)\frac{2\epsilon}{H} \geq \epsilon \geq \frac{\epsilon}{2} = 96H\frac{\epsilon}{192H}$, we have that $\mathcal{I}_{\frac{\epsilon}{192H}}$ is a superset of $\{(s,a): s \in [S_1], a \in [b+1] \setminus \{a_s\}\}$, whose size is at least $S_1b = S_1\lceil \frac{l}{S_1}\rceil \geq l$.

It suffices to prove that

$$\mathbb{E}_{\mathbf{a}\sim \text{Unif}([b+1]^{S_1})}\mathbb{E}_{\mathfrak{M}(\mathbf{a})}\left[\text{Reg}_{\text{Alg}}(K)\right] \geq \frac{1}{640}\sqrt{MH^2lK},$$

where we recall that $\mathbf{a} = (a_1, \dots, a_{S_1})$; furthermore, it suffices to show that, for any $s' \in [S_1]$,

$$\mathbb{E}_{\mathbf{a}\sim \mathrm{Unif}\left([b+1]^{S_1}\right)}\mathbb{E}_{\mathfrak{M}(\mathbf{a})}\left[N^{K+1}(s') - n^{K+1}(s', a_{s'})\right] \ge \frac{MK}{4S_1},\tag{61}$$

where $N^{K+1}(s') = \sum_{a \in \mathcal{A}} n^{K+1}(s', a)$; this is because it follows from Eq. (61) that

$$\mathbb{E}_{\mathbf{a}\sim \mathrm{Unif}\left([b+1]^{S_1}\right)} \mathbb{E}_{\mathfrak{M}(\mathbf{a})} \left[\mathrm{Reg}_{\mathrm{Alg}}(K) \right] \geq \sum_{s' \in \mathcal{S}_1} (H-1) \frac{\Delta}{4} \cdot \mathbb{E}_{\mathbf{a}\sim \mathrm{Unif}\left([b+1]^{S_1}\right)} \mathbb{E}_{\mathfrak{M}(\mathbf{a})} \left[N^{K+1}(s') - n^{K+1}(s', a_{s'}) \right]$$
$$\geq \sum_{s' \in \mathcal{S}_1} \frac{H}{2} \cdot \frac{\Delta}{4} \cdot \frac{MK}{4S_1}$$
$$\geq \frac{1}{640} \sqrt{MH^2 lK},$$

where the first inequality uses Lemma 30 (the regret decomposition lemma).

Without loss of generality, we prove Eq. (61) with s' = 1; the inequality for other s' values are shown symmetrically. To this end, we use a standard technique and define a set of helper problem instances. Specifically, for any $(a_2, a_3, \ldots, a_{S_1}) \in [b+1]^{S_1-1}$, we define a problem instance $\mathfrak{M}(0, a_2, \ldots, a_{s_1})$ such that it agrees with $\mathfrak{M}(a_1, a_2, \ldots, a_{s_1})$ on everything but $\mathbb{P}_p(\cdot | 1, a_1)$'s, i.e., in $\mathfrak{M}(0, a_2, \ldots, a_{s_1})$, for every $p \in [M]$,

$$\mathbb{P}_p(S_1 + 1 \mid 1, a_1) = \frac{1}{2}.$$

Now, for each $(j, a_2, \ldots, a_{s_1}) \in ([0] \cup [b+1]) \times [b+1]^{S_1-1}$, let $\mathbb{P}_{j,a_2,\ldots,a_{S_1}}$ denote the probability measure on the outcomes of running Alg on the problem instance $\mathfrak{M}(j, a_2, \ldots, a_{s_1})$. Further, for each $j \in \{0\} \cup [b+1]$, we define

$$\mathbb{P}_{j} = \frac{1}{(b+1)^{S_{1}-1}} \sum_{a_{2},\dots,a_{S_{1}} \in [b+1]^{S_{1}-1}} \mathbb{P}_{j,a_{2},\dots,a_{S_{1}}};$$

and we use \mathbb{E}_j to denote the expectation with respect to \mathbb{P}_j .

In subsequent calculations, for any index $m \in ([0] \cup [b+1]) \times [b+1]^{S_1-1}$, we also denote by $\mathbb{P}_m(\cdot \mid N^{K+1}(1))$ and $\mathbb{E}_m[\cdot \mid N^{K+1}(1)]$ the probability and expectation, respectively, conditional on a realization of $N^{K+1}(1)$ under \mathbb{P}_m . Observe that, for any $j \in \{0\} \cup [b+1]$,

$$\mathbb{P}_{j}(\cdot \mid N^{K+1}(1)) = \frac{\mathbb{P}_{j}(\cdot, N^{K+1}(1))}{\mathbb{P}_{j}(N^{K+1}(1))} \\
= \frac{\frac{1}{(b+1)^{S_{1}-1}} \sum_{a_{2},...,a_{S_{1}} \in [b+1]^{S_{1}-1}} \mathbb{P}_{j,a_{2},...,a_{S_{1}}}(\cdot, N^{K+1}(1))}{\mathbb{P}_{j}(N^{K+1}(1))} \\
= \frac{1}{(b+1)^{S_{1}-1}} \sum_{a_{2},...,a_{S_{1}} \in [b+1]^{S_{1}-1}} \frac{\mathbb{P}_{j,a_{2},...,a_{S_{1}}}(\cdot, N^{K+1}(1))}{\mathbb{P}_{j,a_{2},...,a_{S_{1}}}(N^{K+1}(1))} \\
= \frac{1}{(b+1)^{S_{1}-1}} \sum_{a_{2},...,a_{S_{1}} \in [b+1]^{S_{1}-1}} \mathbb{P}_{j,a_{2},...,a_{S_{1}}}(\cdot \mid N^{K+1}(1)), \quad (62)$$

where the first equality is from the definition of conditional probability; the second equality is from the definition of \mathbb{P}_j ; the third equality uses the fact that $\mathbb{P}_j(N^{K+1}(1)) = \mathbb{P}_{j,a_2,\ldots,a_{S_1}}(N^{K+1}(1))$ for any a_2, \ldots, a_{S_1} , which is true because $N^{K+1}(1)$ is independent of a_2, \ldots, a_{S_1} conditional on j; and the last equality, again, is from the definition of conditional probability.

We have, for each $j \in [b+1]$,

$$\mathbb{E}_{j}\left[n^{K+1}(1,j) \mid N^{K+1}(1)\right] - \mathbb{E}_{0}\left[n^{K+1}(1,j) \mid N^{K+1}(1)\right] \\
\leq N^{K+1}(1) \left\| \mathbb{P}_{j}\left(\cdot \mid N^{K+1}(1)\right) - \mathbb{P}_{0}\left(\cdot \mid N^{K+1}(1)\right) \right\|_{1} \\
\leq N^{K+1}(1) \cdot \frac{1}{(b+1)^{S_{1}-1}} \sum_{a_{2},...,a_{S_{1}} \in [b+1]^{S_{1}-1}} \left\| \mathbb{P}_{j,a_{2},...,a_{S_{1}}}\left(\cdot \mid N^{K+1}(1)\right) - \mathbb{P}_{0,a_{2},...,a_{S_{1}}}\left(\cdot \mid N^{K+1}(1)\right) \right\|_{1} \\
\leq N^{K+1}(1) \cdot \frac{1}{(b+1)^{S_{1}-1}} \sum_{a_{2},...,a_{S_{1}} \in [b+1]^{S_{1}-1}} \sqrt{2 \operatorname{KL}} \left(\operatorname{Ber}(\frac{1}{2} + \Delta), \operatorname{Ber}(\frac{1}{2})\right) \mathbb{E}_{0,a_{2},...,a_{S_{1}}}\left[n^{K+1}(1,j) \mid N^{K+1}(1)\right] \\
\leq N^{K+1}(1) \cdot \frac{1}{(b+1)^{S_{1}-1}} \sum_{a_{2},...,a_{S_{1}} \in [b+1]^{S_{1}-1}} \sqrt{6\Delta^{2} \mathbb{E}_{0,a_{2},...,a_{S_{1}}}\left[n^{K+1}(1,j) \mid N^{K+1}(1)\right]} \\
\leq N^{K+1}(1) \sqrt{6 \frac{l+1}{384MK} \cdot \mathbb{E}_{0}\left[n^{K+1}(1,j) \mid N^{K+1}(1)\right]} \\
= \frac{1}{8} N^{K+1}(1) \sqrt{\frac{l+1}{MK} \cdot \mathbb{E}_{0}\left[n^{K+1}(1,j) \mid N^{K+1}(1)\right]}.$$
(63)

where the first inequality is based on Lemma 25 and the fact that, conditional on $N^{K+1}(1)$, $n^{K+1}(1,j)$ has distribution supported on $[0, N^{K+1}(1)]$; the second inequality follows from Equation (62) and the triangle inequality; the third inequality uses Pinsker's inequality and Lemma 31 (the divergence decomposition lemma); the fourth inequality uses Lemma 33 and the fact that $\Delta \leq \frac{1}{4}$; and the last inequality follows from Jensen's inequality.

Since $N^{K+1}(1)$ has the same distribution under both \mathbb{P}_0 and any \mathbb{P}_j (which is $Bin(K, \frac{1}{S_1})$), taking expectation with respect to $N^{K+1}(1)$, we have that, for any $j \in [b+1]$,

$$\mathbb{E}_{j}\left[n^{K+1}(1,j)\right] - \mathbb{E}_{0}\left[n^{K+1}(1,j)\right] \leq \mathbb{E}_{0}\left[\frac{1}{8}N^{K+1}(1)\sqrt{\frac{l+1}{MK}} \cdot \mathbb{E}_{0}\left[n^{K+1}(1,j) \mid N^{K+1}(1)\right]\right]$$

In subsequent derivations, we can now avoid bounding the conditional expectation. Specifically, we have

$$\frac{1}{b+1} \sum_{j \in [b+1]} \mathbb{E}_{j} \left[n^{K+1}(1,j) \right] \\
\leq \frac{1}{b+1} \sum_{j \in [b+1]} \mathbb{E}_{0} \left[n^{K+1}(1,j) \right] + \frac{1}{b+1} \sum_{j \in [b+1]} \mathbb{E}_{0} \left[\frac{1}{8} N^{K+1}(1) \sqrt{\frac{l+1}{MK}} \cdot \mathbb{E}_{0} \left[n^{K+1}(1,j) \mid N^{K+1}(1) \right] \right] \\
\leq \frac{1}{b+1} \mathbb{E}_{0} \left[\sum_{j \in [b+1]} n^{K+1}(1,j) \right] + \mathbb{E}_{0} \left[\frac{1}{8} N^{K+1}(1) \sqrt{\frac{l+1}{MK}} \cdot \frac{1}{b+1} \sum_{j \in [b+1]} \mathbb{E}_{0} \left[n^{K+1}(1,j) \mid N^{K+1}(1) \right] \\
\leq \frac{1}{b+1} \mathbb{E}_{0} \left[N^{K+1}(1) \right] + \mathbb{E}_{0} \left[\frac{1}{8} \sqrt{\frac{l+1}{MK}} \cdot \frac{1}{b+1} \left(N^{K+1}(1) \right)^{\frac{3}{2}} \right] \\
\leq \frac{1}{b+1} \mathbb{E}_{0} \left[N^{K+1}(1) \right] + \frac{1}{8} \sqrt{\frac{S_{1}}{MK}} \cdot \mathbb{E}_{0} \left[\left(N^{K+1}(1) \right)^{\frac{3}{2}} \right],$$
(64)

where the first inequality follows from Eq. (63) and algebra; the second inequality uses linearity of expectation and Jensen's inequality; the third inequality uses the facts that $\sum_{j \in [b+1]} n^{K+1}(1,j) \leq N^{K+1}(1)$ and, for every $z \in [0] \cup [b+1]$,

$$\sum_{j \in [b+1]} \mathbb{E}_z \left[n^{K+1}(1,j) \mid N^{K+1}(1) \right] \le \sum_{j \in \mathcal{A}} \mathbb{E}_z \left[n^{K+1}(1,j) \mid N^{K+1}(1) \right] = N^{K+1}(1)$$

and the last inequality uses the linearity of expectation and the construction that $b = \lceil \frac{l}{S_1} \rceil$, which implies that $l \leq bS_1$ and therefore $l + 1 \leq bS_1 + 1 \leq bS_1 + S_1 = (b+1)S_1$.

It follows from Equation (64) that

$$\begin{aligned} \frac{1}{b+1} \sum_{j \in [b+1]} \mathbb{E}_j \left[n^{K+1}(1,j) \right] &\leq \frac{1}{b+1} \cdot \frac{MK}{S_1} + \frac{1}{8} \sqrt{\frac{S_1}{MK}} \cdot \mathbb{E}_0 \left[\left(N^{K+1}(1) \right)^{\frac{3}{2}} \right] \\ &\leq \frac{MK}{2S_1} + \frac{1}{4} \sqrt{\frac{S_1}{MK} \left(\frac{MK}{S_1} \right)^3} \\ &\leq \frac{3MK}{4S_1}, \end{aligned}$$

where the second inequality uses the fact that $\frac{1}{b+1} \leq \frac{1}{2}$ and Lemma 34 under the assumption that $K \geq S_1$.

It then follows that

$$\frac{1}{b+1} \sum_{j \in [b+1]} \mathbb{E}_j \left[N^{K+1}(1) - n^{K+1}(1,j) \right] \ge \frac{1}{b+1} \sum_{j \in [b+1]} \mathbb{E}_j \left[N^{K+1}(1) \right] - \frac{3MK}{4S_1} = \frac{MK}{4S_1}$$

and we have

$$\mathbb{E}_{\mathbf{a}\sim \text{Unif}([b+1]^{S_1})}\mathbb{E}_{\mathfrak{M}(\mathbf{a})}\left[N^{K+1}(1) - n^{K+1}(1,a_1)\right] \ge \frac{MK}{4S_1}$$

Case 2: $Ml^C \ge l$. Again, let $S_1 = S - 2(H - 1)$. Let $u = \lceil \frac{l}{S_1} \rceil$ and $v = A - u = A - \lceil \frac{l}{S_1} \rceil$. Furthermore, let $\Delta = \sqrt{\frac{vS_1}{384K}}$, and $\epsilon = 2H\Delta$. We note that under the assumption that $K \ge SA$ and the fact that $vS_1 \le SA$, we have $\Delta \le \frac{1}{4}$. We will define $v^{S_1 \times M} \epsilon$ -MPERL problem instances, each indexed by an element in $[v]^{S_1 \times M}$. It suffices to show that, on at least one of the instances, $\mathbb{E}\left[\operatorname{Reg}_{Alg}(K)\right] \ge \Omega\left(M\sqrt{H^2l^CK}\right)$.

Facts about v. There are two helpful facts about v that can be easily verified:

- $vS_1 \ge \frac{1}{2}l^C$. This is true because, by definition, $vS_1 \ge S_1A l S_1 = S_1A (SA l^C) S_1 = l^C (SA S_1A) S_1 = l^C (2(H-1)A + S_1)$; since, by assumption, $l \le SA 4(S + HA)$, we have $l^C \ge 4(HA + S) \ge 2(2(H-1)A + S_1)$; it then follows that $vS_1 \ge l^C (2(H-1)A + S_1) \ge \frac{1}{2}l^C$.
- $v \ge 2$. This is true because, as shown above, $vS_1 \ge \frac{1}{2}l^C$ and $l^C \ge 4(HA + S)$, which imply that $v \ge \frac{2(HA+S)}{S_1} \ge \frac{2S_1}{S_1} = 2$.

Construction. For $\mathbf{a} = (a_{1,1}, \ldots, a_{1,M}, a_{2,1}, \ldots, a_{S_1,M}) \in [v]^{S_1 \times M}$, we define the following ϵ -MPERL problem instance, $\mathfrak{M}(\mathbf{a}) = \{\mathcal{M}_p\}_{p=1}^M$, with S states, A actions, and an episode length of H, such that for each $p \in [M]$, \mathcal{M}_p is constructed in the same way as it is for case 1, except for the transition probabilities of $(s, a) \in S_1 \times A$:

• For each state $s \in [S_1]$,

$$\mathbb{P}_p(S_1 + 1 \mid s, a) = \begin{cases} \frac{1}{2} + \Delta, & \text{if } a = a_{s,p}; \\ \frac{1}{2}, & \text{if } a \in [v] \setminus \{a_{s,p}\}; \\ 0, & \text{if } a \notin [v]; \end{cases}$$

and for each $a \in A$, $\mathbb{P}_p(S_1 + 2 \mid s, a) = 1 - \mathbb{P}_p(S_1 + 1 \mid s, a)$, and $R_p(s, a) = 0$.

We now verify that $\mathfrak{M}(\mathbf{a})$ is an ϵ -MPMAB problem instance. It can be easily observed that the reward distributions are the same for all players, i.e., for every $p, q \in [M]$ and every $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\left|R_p(s,a) - R_q(s,a)\right| = 0 \le \epsilon.$$

Regarding the transition probabilities, for every $(s, a) \in ((S_1 \times (A \setminus [v]))) \cup ((S \setminus S_1) \times A)$, we observe that the transition probabilities are the same for all players. Furthermore, for every $p, q \in [M]$ and every $(s, a) \in S_1 \times [v]$,

$$\left\| \mathbb{P}_p\left(\cdot \mid s, a \right) - \mathbb{P}_q\left(\cdot \mid s, a \right) \right\|_1 \le 2\Delta = \frac{\epsilon}{H}$$

Therefore, $\mathfrak{M}(\mathbf{a})$ is an ϵ -MPMAB problem instance.

Suboptimality gaps. Similar to the arguments in Case 1, it can be shown that for every $p \in [M]$, and every $(s, a) \in (S \setminus S_1) \times A$, gap_p(s, a) = 0. And, for every $p \in [M]$, and every $s \in S_1$,

$$\operatorname{gap}_{p}(s,a) = \begin{cases} 0, & \text{if } a = a_{s,p};\\ (H-1)\Delta, & \text{if } a \in [v] \setminus a_{s,p};\\ (H-1)\left(\frac{1}{2} + \Delta\right), & \text{if } a \notin [v]. \end{cases}$$

Subpar state-action pairs. Based on the above construction, for every $(s, a) \in S_1 \times (\mathcal{A} \setminus [v])$ and every $p \in [M]$, $\operatorname{gap}_p(s, a) = (H - 1) \left(\frac{1}{2} + \Delta\right) \ge 3(H - 1)\Delta = \frac{3(H-1)}{2H}\epsilon \ge \frac{3}{4}\epsilon \ge 96H\left(\frac{\epsilon}{192H}\right)$, where the first inequality uses the fact that $\Delta \le \frac{1}{4}$. Therefore, $\mathcal{I}_{\frac{\epsilon}{192H}} = i$ is a superset of $[S_1] \times ([A] \setminus [v])$, whose cardinality is at least $(A - v)S_1 = uS_1 \ge l$, i.e., $\left|\mathcal{I}_{\frac{\epsilon}{192H}}\right| \ge l$.

Now, it suffices to prove that

$$\mathbb{E}_{\mathbf{a}\sim \text{Unif}([v]^{S_1\times M})}\mathbb{E}_{\mathfrak{M}(\mathbf{a})}\left[\text{Reg}_{\text{Alg}}(K)\right] \geq \frac{1}{240}M\sqrt{H^2l^CK},$$

where we recall that $\mathbf{a} = (a_{1,1}, \ldots, a_{1,M}, a_{2,1}, \ldots, a_{S_1,M})$. It suffices to show, for any $s' \in [S_1]$ and any $p' \in [M]$,

$$\mathbb{E}_{\mathbf{a}\sim\mathrm{Unif}\left([v]^{S_1\times M}\right)}\mathbb{E}_{\mathfrak{M}(\mathbf{a})}\left[N_{p'}^{K+1}(s')-n_{p'}^{K+1}(s',a_{s'})\right]\geq\frac{K}{4S_1},\tag{65}$$

where $N_{p'}^{K+1}(s') = \sum_{a \in \mathcal{A}} n_{p'}^{K+1}(s', a)$. To see this, by Lemma 30, we have

$$\mathbb{E}_{\mathbf{a}\sim\mathrm{Unif}\left([v]^{S_{1}\times M}\right)}\mathbb{E}_{\mathfrak{M}(\mathbf{a})}\left[\mathrm{Reg}_{\mathrm{Alg}}(K)\right] \geq \sum_{p=1}^{M} \sum_{s'\in\mathcal{S}_{1}} (H-1)\Delta \cdot \mathbb{E}_{\mathbf{a}\sim\mathrm{Unif}\left([v]^{S_{1}\times M}\right)}\mathbb{E}_{\mathfrak{M}(\mathbf{a})}\left[N_{p}^{K+1}(s') - n_{p}^{K+1}(s',a_{s'})\right]$$
$$\geq \frac{H-1}{4}MK\sqrt{\frac{vS_{1}}{384K}}$$
$$\geq \frac{1}{160}M\sqrt{H^{2}(vS_{1})K}$$
$$\geq \frac{1}{240}M\sqrt{H^{2}l^{C}K},$$

where the last inequality uses the fact that $vS_1 \geq \frac{1}{2}l^C$.

Without loss of generality, it suffices to prove Eq. (65) for s' = 1 and p' = 1; the other settings of (s', p') can be handled symmetrically. Similar to case 1, we define a set of helper problem instances: for any $(a_{1,2}, \ldots, a_{S_1,M}) \in [v]^{S_1 \times M - 1}$, we define a problem instance $\mathfrak{M}(0, a_{1,2}, \ldots, a_{S_1,M})$ such that it agrees with $\mathfrak{M}(a_{1,1}, a_{1,2}, \ldots, a_{S_1,M})$ on everything but $\mathbb{P}_1(\cdot \mid 1, a_1)$, namely, in $\mathfrak{M}(0, a_{1,2}, \ldots, a_{S_1,M})$, $\mathbb{P}_1(S_1 + 1 \mid 1, a_1) = \frac{1}{2}$.

For each $(j, a_{1,2}, \ldots, a_{S_1,M}) \in ([0] \cup [v]) \times [v]^{S_1 \times M - 1}$, let $\mathbb{P}_{j,a_{1,2},\ldots,a_{S_1,M}}$ denote the probability measure on the outcomes of running Alg on the problem instance $\mathfrak{M}(j, a_{1,2}, \ldots, a_{S_1,M})$. Further, for each $j \in \{0\} \cup [v]$, we define

$$\mathbb{P}_{j} = \frac{1}{v^{S_{1} \times M - 1}} \sum_{a_{1,2}, \dots, a_{S_{1},M} \in [v]^{S_{1} \times M - 1}} \mathbb{P}_{j, a_{1,2}, \dots, a_{S_{1},M}};$$

and we use \mathbb{E}_j to denote the expectation with respect to \mathbb{P}_j . In subsequent calculations, for any $m \in ([0] \cup [v]) \times [v]^{S_1 \times M-1}$, we also denote by $\mathbb{P}_m \left(\cdot \mid N_1^{K+1}(1) \right)$ and $\mathbb{E}_m \left[\cdot \mid N_1^{K+1}(1) \right]$ the probability and expectation conditional on a realization of $N_1^{K+1}(1)$ under \mathbb{P}_m . Similar to case 1, it can be shown that, for any $j \in \{0\} \cup [v]$,

$$\mathbb{P}_{j}(\cdot \mid N^{K+1}(1)) = \frac{1}{v^{S_{1} \times M - 1}} \sum_{a_{1,2},\dots,a_{S_{1},M} \in [v]^{S_{1} \times M - 1}} \mathbb{P}_{j,a_{1,2},\dots,a_{S_{1},M}}\left(\cdot \mid N^{K+1}(1)\right).$$
(66)

Now, for each $j \in [v]$, we have

$$\mathbb{E}_{j} \left[n_{1}^{K+1}(1,j) \mid N_{1}^{K+1}(1) \right] - \mathbb{E}_{0} \left[n_{1}^{K+1}(1,j) \mid N_{1}^{K+1}(1) \right] \\
\leq N_{1}^{K+1}(1) \left\| \mathbb{P}_{j} \left(\cdot \mid N_{1}^{K+1}(1) \right) - \mathbb{P}_{0} \left(\cdot \mid N_{1}^{K+1}(1) \right) \right\|_{1} \\
\leq N_{1}^{K+1}(1) \cdot \frac{1}{v^{S_{1} \times M-1}} \sum_{a_{1,2},\dots,a_{S_{1},M} \in [v]^{S_{1} \times M-1}} \left\| \mathbb{P}_{j,a_{1,2},\dots,a_{S_{1},M}} \left(\cdot \mid N_{1}^{K+1}(1) \right) - \mathbb{P}_{0,a_{1,2},\dots,a_{S_{1},M}} \left(\cdot \mid N_{1}^{K+1}(1) \right) \right\|_{1} \\
\leq N_{1}^{K+1}(1) \cdot \frac{1}{v^{S_{1} \times M-1}} \sum_{a_{1,2},\dots,a_{S_{1},M} \in [v]^{S_{1} \times M-1}} \sqrt{2 \operatorname{KL} \left(\operatorname{Ber}(\frac{1}{2} + \Delta), \operatorname{Ber}(\frac{1}{2}) \right) \mathbb{E}_{0,a_{2},\dots,a_{S_{1}}} \left[n_{1}^{K+1}(1,j) \mid N_{1}^{K+1}(1) \right] \\
\leq N_{1}^{K+1}(1) \cdot \frac{1}{v^{S_{1} \times M-1}} \sum_{a_{1,2},\dots,a_{S_{1},M} \in [v]^{S_{1} \times M-1}} \sqrt{6\Delta^{2} \mathbb{E}_{0,a_{2},\dots,a_{S_{1}}} \left[n_{1}^{K+1}(1,j) \mid N_{1}^{K+1}(1) \right] \\
\leq N_{1}^{K+1}(1) \cdot \sqrt{\frac{6vS_{1}}{384K} \cdot \mathbb{E}_{0} \left[n_{1}^{K+1}(1,j) \mid N_{1}^{K+1}(1) \right] }$$

$$(67)$$

where the first inequality is based on Lemma [25] and the fact that, conditional on $N_1^{K+1}(1)$, $n_1^{K+1}(1,j)$ has distribution supported on $[0, N_1^{K+1}(1)]$; the second inequality follows from Equation (66) and the triangle inequality; the third inequality uses Pinsker's inequality and Lemma [31] (the divergence decomposition lemma); the fourth inequality uses Lemma [33] and the fact that $\Delta \leq \frac{1}{4}$; and the last inequality follows from Jensen's inequality.

Using arguments similar to the ones shown for case 1, we have that

$$\begin{split} &\frac{1}{v}\sum_{j\in[v]}\mathbb{E}_{j}\left[n_{1}^{K+1}(1,j)\right] \\ \leq &\frac{1}{v}\mathbb{E}_{0}\left[n_{1}^{K+1}(1,j)\right] + \mathbb{E}_{0}\left[\frac{1}{8}N_{1}^{K+1}(1)\sqrt{\frac{vS_{1}}{K}\cdot\frac{1}{v}\sum_{j\in[v]}\mathbb{E}_{0}\left[n_{1}^{K+1}(1,j)\mid N_{1}^{K+1}(1)\right]}\right] \\ \leq &\frac{1}{v}\mathbb{E}_{0}\left[N^{K+1}(1)\right] + \frac{1}{8}\sqrt{\frac{S_{1}}{K}}\cdot\mathbb{E}_{0}\left[\left(N_{1}^{K+1}(1)\right)^{\frac{3}{2}}\right] \\ \leq &\frac{1}{v}\cdot\frac{K}{S_{1}} + \frac{1}{4}\sqrt{\frac{S_{1}}{K}\left(\frac{K}{S_{1}}\right)^{3}} \\ \leq &\frac{3K}{4S_{1}}, \end{split}$$

where the second to last inequality is from Lemma 34 under the assumption that $K \ge S_1$, and the last inequality uses the fact that $v \ge 2$.

It then follows that

$$\frac{1}{v}\sum_{j\in[v]}\mathbb{E}_j\left[N_1^{K+1}(1) - n_1^{K+1}(1,j)\right] \ge \frac{1}{v}\sum_{j\in[v]}\mathbb{E}_j\left[N_1^{K+1}(1)\right] - \frac{3K}{4S_1} = \frac{K}{4S_1},$$

and we thereby have shown that

$$\mathbb{E}_{\mathbf{a} \sim \mathrm{Unif}\left([v]^{S_1 \times M}\right)} \mathbb{E}_{\mathfrak{M}(\mathbf{a})}\left[N_1^{K+1}(1) - n_1^{K+1}(1, a_1)\right] \ge \frac{K}{4S_1}.$$

D.3 Gap dependent lower bound

Theorem 36 (Restatement of Theorem 8). Fix $\epsilon \ge 0$. For any $S \in \mathbb{N}$, $A \ge 2$, $H \ge 2$, $M \in \mathbb{N}$, such that $S \ge 2(H-1)$, let $S_1 = S - 2(H-1)$; and let $\{\Delta_{s,a,p}\}_{(s,a,p)\in[S_1]\times[A]\times[M]}$ be any set of values such that

- for every $(s, a, p) \in [S_1] \times [A] \times [M], \Delta_{s,a,p} \in [0, \frac{H}{48\sqrt{M}}];$
- for every $(s, p) \in [S_1] \times [M]$, there exists at least one action $a \in [A]$ such that $\Delta_{s,a,p} = 0$;
- and, for every $(s, a) \in [S_1] \times [A]$ and $p, q \in [M], |\Delta_{s,a,p} \Delta_{s,a,q}| \le \epsilon/4$.

There exists an ϵ -MPERL problem instance with S states, A actions, M players and an episode length of H, such that $S_1 = [S_1]$, $|S_h| = 2$ for all $h \ge 2$, and

$$\operatorname{gap}_p(s,a) = \Delta_{s,a,p}, \quad \forall (s,a,p) \in [S_1] \times [A] \times [M].$$

For this problem instance, any sublinear regret algorithm Alg for the ϵ -MPERL problem must have regret at least

$$\mathbb{E}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right] \ge \Omega\left(\ln K\left(\sum_{\substack{p \in [M] \ (s,a) \in \mathcal{I}_{(\epsilon/768H)}^C: \\ \operatorname{gap}_p(s,a) > 0}} \frac{H^2}{\operatorname{gap}_p(s,a)} + \sum_{\substack{(s,a) \in \mathcal{I}_{(\epsilon/768H)}}} \frac{H^2}{\operatorname{min}_p \operatorname{gap}_p(s,a)}\right)\right)$$

Proof. The construction and techniques in this proof are inspired by [36] and [43].

Proof outline. We will construct an ϵ -MPERL problem instance, \mathfrak{M} , and show that, for any sublinear regret algorithm and sufficiently large K, the following two claims are true:

1. for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ such that for all p, $gap_p(s, a) > 0$,

$$\mathbb{E}_{\mathfrak{M}}\left[n^{K}(s,a)\right] \ge \Omega\left(\frac{H^{2}}{\left(\min_{p}\operatorname{gap}_{p}(s,a)\right)^{2}}\ln K\right);$$
(68)

2. for any $(s,a)\in \mathcal{I}_{\frac{c}{768H}}^C$ and any $p\in [M]$ such that $\mathrm{gap}_p(s,a)>0,$

$$\mathbb{E}_{\mathfrak{M}}\left[n_p^K(s,a)\right] \ge \Omega\left(\frac{H^2}{\left(\operatorname{gap}_p(s,a)\right)^2}\ln K\right).$$
(69)

The rest then follows from Lemma 30 (the regret decomposition lemma).

Construction of \mathfrak{M} . Given any set of values $\{\Delta_{s,a,p}\}_{(s,a,p)\in[S_1]\times[A]\times[M]}$ that satisfies the assumptions in the theorem statement, we can construct a collection of MDPs $\{\mathcal{M}_p\}_{p=1}^M$, such that for each $p \in [M]$, \mathcal{M}_p is as follows, and $\mathfrak{M} = \{\mathcal{M}_p\}_{p=1}^M$ is an ϵ -MPERL problem instance:

• $S_1 = [S_1]$, and p_0 is a uniform distribution over the states in S_1 .

- For $h \in [2, H]$, $S_h = \{S_1 + 2h 3, S_1 + 2h 2\}$.
- $\mathcal{A} = [A].$
- For all $(s,a) \in \mathcal{S} \times \mathcal{A}$, the reward distribution $r_p(s,a)$ is a Bernoulli distribution, $Ber(R_p(s,a))$, and we specify $R_p(s,a)$ subsequently.
- For every $(s, a) \in S_1 \times [A]$, set $\bar{\Delta}_{s,a}^p = \frac{\Delta_{s,a,p}}{H-1}$. Then, let

$$\mathbb{P}_{p}\left(S_{1}+1 \mid s, a\right) = \frac{1}{2} - \bar{\Delta}_{s,a}^{p}, \quad \mathbb{P}_{p}\left(S_{1}+2 \mid s, a\right) = \frac{1}{2} + \bar{\Delta}_{s,a}^{p},$$

and $R_p(s, a) = 0$. Since $\Delta_{s,a,p} \in [0, H/48]$, $\overline{\Delta}_{s,a}^p \leq \frac{H}{48(H-1)} \leq \frac{1}{24}$, where the last inequality follows from the assumption that $H \geq 2$. Therefore, $\mathbb{P}_p(S_1 + 1 \mid s, a) \in [0, 1]$, and $\mathbb{P}_p(S_1 + 2 \mid s, a) \in [0, 1]$.

• For $h \in [2, H]$, and $a \in [A]$, let

 $- \mathbb{P}_p \left(S_1 + 2h - 1 \mid S_1 + 2h - 3, a \right) = 1, \mathbb{P}_p \left(S_1 + 2h \mid S_1 + 2h - 3, a \right) = 0, \text{ and } R_p (S_1 + 2h - 3, a) = 1.$ $- \mathbb{P}_p \left(S_1 + 2h \mid S_1 + 2h - 2, a \right) = 0, \mathbb{P}_p \left(S_1 + 2h - 1 \mid S_1 + 2h - 2, a \right) = 1, \text{ and } R_p (S_1 + 2h - 2, a) = 0.$

By the assumption that for every $(s, p) \in [S_1] \times [M]$, there exists at least one action $a \in [A]$ such that $\Delta_{s,a,p} = 0$, we have that there is at least one action a such that $\bar{\Delta}_{s,a}^p = 0$. We verify that for every $(s, a, p) \in [S_1] \times [A] \times [M]$,

$$gap_p(s,a) = V_p^*(s) - Q_p^*(s,a)$$
$$= \max_{a'} Q_p^*(s,a') - Q_p^*(s,a)$$
$$= (H-1)\bar{\Delta}_{s,a}^p$$
$$= \Delta_{s,a,p}.$$

We now verify that the above MPERL problem instance $\mathfrak{M} = \{\mathcal{M}_p\}_{p=1}^M$ is an ϵ -MPERL problem instance:

1. The reward distributions are the same for all players, namely, for all p, q,

$$|R_p(s,a) - R_q(s,a)| = 0 \le \epsilon, \forall (s,a) \in \mathcal{S} \times \mathcal{A}$$

2. Further, by the assumption that for every $(s,a) \in [S_1] \times [A]$ and $p,q \in [M]$, $|\Delta_{s,a,p} - \Delta_{s,a,q}| \le \epsilon/4$, we have that

$$\left|\bar{\Delta}_{s,a}^p - \bar{\Delta}_{s,a}^q\right| = \frac{\left|\Delta_{s,a,p} - \Delta_{s,a,q}\right|}{H - 1} \le \frac{\epsilon}{4(H - 1)} \le \frac{\epsilon}{2H}$$

It then follows that

$$\left\|\mathbb{P}_{p}\left(\cdot \mid s,a\right) - \mathbb{P}_{q}\left(\cdot \mid s,a\right)\right\|_{1} = 2\left|\bar{\Delta}_{s,a}^{p} - \bar{\Delta}_{s,a}^{q}\right| \leq \frac{\epsilon}{H}.$$

Meanwhile, for every $(s, a) \in (S \setminus S_1) \times A$

$$\|\mathbb{P}_p(\cdot \mid s, a) - \mathbb{P}_q(\cdot \mid s, a)\|_1 = 0 \le \frac{\epsilon}{H}$$

In summary, for every $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\|\mathbb{P}_p(\cdot \mid s, a) - \mathbb{P}_q(\cdot \mid s, a)\|_1 \le \frac{\epsilon}{H}$$

We are now ready to prove the two claims:

1. Proving claim 1 (Equation (68)):

Fix any $(s_0, a_0) \in [S_1] \times [A]$ such that $\overline{\Delta}_{s_0, a_0}^{\min} = \min_p \overline{\Delta}_{s_0, a_0}^p > 0$. It can be easily observed that $\operatorname{gap}_p(s_0, a_0) > 0$ for all p. Define $p_0 = \operatorname{argmin}_p \overline{\Delta}_{s_0, a_0}^p$. We can construct a new problem instance, \mathfrak{M}' , which agrees with \mathfrak{M} , except that

$$\forall p \in [M], \mathbb{P}_p\left(S_1 + 1 \mid s_0, a_0\right) = \frac{1}{2} - \bar{\Delta}_{s_0, a_0}^p + 2\bar{\Delta}_{s_0, a_0}^{\min}, \mathbb{P}_p\left(S_1 + 2 \mid s_0, a_0\right) = \frac{1}{2} + \bar{\Delta}_{s_0, a_0}^p - 2\bar{\Delta}_{s_0, a_0}^{\min}$$

 \mathfrak{M}' is an ϵ -MPERL problem instance. To see this, we note that the only change is in $\mathbb{P}_p(\cdot | s_0, a_0)$ for all $p \in [M]$. In this new instance, it is still true that for every $p, q \in [M]$,

$$\left\|\mathbb{P}_{p}\left(\cdot\mid s_{0},a_{0}\right)-\mathbb{P}_{q}\left(\cdot\mid s_{0},a_{0}\right)\right\|_{1}=2\left|\bar{\Delta}_{s_{0},a_{0}}^{p}-\bar{\Delta}_{s_{0},a_{0}}^{q}\right|\leq\frac{\epsilon}{H}$$

Fix any sublinear regret algorithm Alg for the ϵ -MPERL problem. By Lemma 31 (the divergence decomposition lemma), we have

$$\mathrm{KL}(\mathbb{P}_{\mathfrak{M}},\mathbb{P}_{\mathfrak{M}'}) = \sum_{p=1}^{M} \mathbb{E}_{\mathfrak{M}}\left[n_{p}^{K}(s_{0},a_{0})\right] \mathrm{KL}\left(\mathbb{P}_{p}^{\mathfrak{M}}(\cdot \mid s_{0},a_{0}),\mathbb{P}_{p}^{\mathfrak{M}'}(\cdot \mid s_{0},a_{0})\right),$$

where $\mathbb{P}_{\mathfrak{M}}$ and $\mathbb{P}_{\mathfrak{M}'}$ are the probability measures on the outcomes of running Alg on \mathfrak{M} and \mathfrak{M}' , respectively; $\mathbb{P}_p^{\mathfrak{M}}(\cdot \mid s_0, a_0)$, $\mathbb{P}_p^{\mathfrak{M}'}(\cdot \mid s_0, a_0)$ are the transition probabilities for (s_0, a_0) and player p in \mathfrak{M} and \mathfrak{M}' , respectively.

We observe that, for any $p \in [M]$,

$$\begin{aligned} \operatorname{KL}\left(\mathbb{P}_{p}^{\mathfrak{M}}(\cdot \mid s_{0}, a_{0}), \mathbb{P}_{p}^{\mathfrak{M}'}(\cdot \mid s_{0}, a_{0})\right) \\ = \operatorname{KL}\left(\operatorname{Ber}\left(\frac{1}{2} - \bar{\Delta}_{s_{0}, a_{0}}^{p}\right), \operatorname{Ber}\left(\frac{1}{2} - \bar{\Delta}_{s_{0}, a_{0}}^{p} + 2\bar{\Delta}_{s_{0}, a_{0}}^{\min}\right)\right) \\ \leq 12(\bar{\Delta}_{s_{0}, a_{0}}^{\min})^{2}, \end{aligned}$$

where the last inequality follows from Lemma 33 and the assumption that $\Delta_{s,a,p} \leq \frac{H}{48}$.

In addition, $\sum_{p=1}^{M} \mathbb{E}_{\mathfrak{M}} \left[n_{p}^{K}(s_{0}, a_{0}) \right] = \mathbb{E}_{\mathfrak{M}} \left[n^{K}(s_{0}, a_{0}) \right]$. It then follows that $\operatorname{KL}(\mathbb{P}_{\mathfrak{M}}, \mathbb{P}_{\mathfrak{M}'}) \leq 12 \mathbb{E}_{\mathfrak{M}} \left[n^{K}(s_{0}, a_{0}) \right] (\bar{\Delta}_{s_{0}, a_{0}}^{\min})^{2}.$ (70)

Now, in the original ϵ -MPERL problem instance, \mathfrak{M} , by Equation (58) and Markov's Inequality, we have

$$\mathbb{E}_{\mathfrak{M}}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right] \geq \frac{K}{4S_1}\left((H-1)\bar{\Delta}_{s_0,a_0}^{\min}\right)\mathbb{P}_{\mathfrak{M}}\left(n_{p_0}^K(s_0,a_0)\geq \frac{K}{4S_1}\right);$$

where we note that $\bar{\Delta}_{s_0,a_0}^{p_0} = \bar{\Delta}_{s_0,a_0}^{\min}$. In \mathfrak{M}' , the new ϵ -MPERL problem instance, we have

$$\begin{split} \mathbb{E}_{\mathfrak{M}'} \left[\operatorname{Reg}_{\operatorname{Alg}}(K) \right] &\geq \left((H-1) \bar{\Delta}_{s_0, a_0}^{\min} \right) \mathbb{E}_{\mathfrak{M}'} \left[\sum_{a \neq a_0} n_{p_0}(s_0, a) \right] \\ &= \left((H-1) \bar{\Delta}_{s_0, a_0}^{\min} \right) \mathbb{E}_{\mathfrak{M}'} \left[N_{p_0}^K(s_0) - n_{p_0}(s_0, a_0) \right] \\ &\geq \frac{K}{4S_1} \left((H-1) \bar{\Delta}_{s_0, a_0}^{\min} \right) \mathbb{P}_{\mathfrak{M}'} \left(N_{p_0}^K(s_0) - n_{p_0}(s_0, a_0) \geq \frac{K}{4S_1} \right) \\ &\geq \frac{K}{4S_1} \left((H-1) \bar{\Delta}_{s_0, a_0}^{\min} \right) \mathbb{P}_{\mathfrak{M}'} \left(N_{p_0}^K(s_0) \geq \frac{K}{2S_1}, n_{p_0}(s_0, a_0) \leq \frac{K}{4S_1} \right) \\ &\geq \frac{K}{4S_1} \left((H-1) \bar{\Delta}_{s_0, a_0}^{\min} \right) \left(\mathbb{P}_{\mathfrak{M}'} \left(n_{p_0}(s_0, a_0) \leq \frac{K}{4S_1} \right) - \exp(-\frac{K}{8S_1}) \right) \end{split}$$

where the first inequality is by Equation (58); the second inequality is by Markov's Inequality; the third inequality is by simple algebra; and the last inequality is by Chernoff bound that $\mathbb{P}_{\mathfrak{M}'}\left(N_{p_0}^K(s_0) < \frac{K}{2S_1}\right) \leq \exp(-\frac{K}{8S_1})$, and $\mathbb{P}(A \cap B) \geq \mathbb{P}(B) - \mathbb{P}(A^C)$ for events A, B.

It then follows that

$$\mathbb{E}_{\mathfrak{M}}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right] + \mathbb{E}_{\mathfrak{M}'}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right]$$

$$= \frac{K}{2}\left((H-1)\bar{\Delta}_{s_{0},a_{0}}^{\min}\right)\left(\mathbb{P}_{\mathfrak{M}}\left(n_{p_{0}}^{K}(s_{0},a_{0}) \geq \frac{K}{4S_{1}}\right) + \mathbb{P}_{\mathfrak{M}'}\left(n_{p_{0}}^{K}(s_{0},a_{0}) < \frac{K}{4S_{1}}\right) - \exp(-\frac{K}{8S_{1}})\right)$$

$$\geq \frac{K}{2}\left((H-1)\bar{\Delta}_{s_{0},a_{0}}^{\min}\right)\left(\frac{1}{2}\exp\left(-\operatorname{KL}(\mathbb{P}_{\mathfrak{M}},\mathbb{P}_{\mathfrak{M}'})\right) - \exp(-\frac{K}{8S_{1}})\right)$$

$$\geq \frac{K}{2}\left((H-1)\bar{\Delta}_{s_{0},a_{0}}^{\min}\right)\left(\frac{1}{2}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^{K}(s_{0},a_{0})\right](\bar{\Delta}_{s_{0},a_{0}}^{\min})^{2}\right) - \exp(-\frac{K}{8S_{1}})\right),$$

where the first inequality follows from Lemma 2 (the Bretagnolle-Huber inequality), and the second inequality follows from Eq. (70). Observe that $\mathbb{E}_{\mathfrak{M}}\left[n^{K}(s_{0}, a_{0})\right] \leq \frac{MK}{S_{1}}$; in addition, by our assumption that $\Delta_{s,a,p} \leq \frac{H}{48\sqrt{M}}$ for every (s, a, p), we have $\bar{\Delta}_{s_{0},a_{0}}^{\min} \leq \frac{1}{24\sqrt{M}}$. These together implies that $\frac{1}{4} \exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^{K}(s_{0}, a_{0})\right](\bar{\Delta}_{s_{0},a_{0}}^{\min})^{2}\right) \geq \frac{1}{4} \exp\left(-\frac{K}{48S_{1}}\right) \geq \exp\left(-\frac{K}{8S_{1}}\right)$, as long as $K \geq 20S_{1}$. Therefore, we have

$$\mathbb{E}_{\mathfrak{M}}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right] + \mathbb{E}_{\mathfrak{M}'}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right] \geq \frac{K}{2}\left((H-1)\bar{\Delta}_{s_0,a_0}^{\min}\right) \cdot \frac{1}{4}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0)\right](\bar{\Delta}_{s_0,a_0}^{\min})^2\right) \cdot \frac{1}{4}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0)\right](\bar{\Delta}_{s_0,a_0}^{\min})^2\right] \cdot \frac{1}{4}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0)\right](\bar{\Delta}_{s_0,a_0}^{\min})^2\right] \cdot \frac{1}{4}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0)\right](\bar{\Delta}_{s_0,a_0}^{\min})^2\right] \cdot \frac{1}{4}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0)\right](\bar{\Delta}_{s_0,a_0}^{\min})^2\right] \cdot \frac{1}{4}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0)\right](\bar{\Delta}_{s_0,a_0}^{\min})^2\right] \cdot \frac{1}{4}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0)\right](\bar{\Delta}_{s_0,a_0}^{\min})^2\right] \cdot \frac{1}{4}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0)\right](\bar{\Delta}_{s_0,a_0}^{\max})^2\right] \cdot \frac{1}{4}\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0)\right](\bar{\Delta}_{s_0,a_0}^{\max$$

Now, under the assumption that Alg is a sublinear regret algorithm, we have

$$\frac{K}{8} \left((H-1)\bar{\Delta}_{s_0,a_0}^{\min} \right) \exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^K(s_0,a_0) \right] (\bar{\Delta}_{s_0,a_0}^{\min})^2 \right) \le 2CK^{\alpha}$$

It follows that

$$\mathbb{E}_{\mathfrak{M}}\left[n^{K}(s_{0}, a_{0})\right] \geq \frac{1}{12\left(\bar{\Delta}_{s_{0}, a_{0}}^{\min}\right)^{2}} \ln\left(\frac{(H-1)\bar{\Delta}_{s_{0}, a_{0}}^{\min}K^{1-\alpha}}{16C}\right)$$
$$= \frac{(H-1)^{2}}{12\left(\min_{p} \operatorname{gap}_{p}(s_{0}, a_{0})\right)^{2}} \ln\left(\frac{\min_{p} \operatorname{gap}_{p}(s_{0}, a_{0})K^{1-\alpha}}{16C}\right)$$
$$\geq \frac{H^{2}}{48\left(\min_{p} \operatorname{gap}_{p}(s_{0}, a_{0})\right)^{2}} \ln\left(\frac{\min_{p} \operatorname{gap}_{p}(s_{0}, a_{0})K^{1-\alpha}}{16C}\right)$$

We then have

$$\mathbb{E}_{\mathfrak{M}}\left[n^{K}(s_{0}, a_{0})\right] \geq \Omega\left(\frac{H^{2}}{\left(\min_{p} \operatorname{gap}_{p}(s_{0}, a_{0})\right)^{2}} \ln K\right).$$

2. Proving Claim 2 (Equation (69)):

Fix any $(s_0, a_0) \in \mathcal{I}_{\frac{\epsilon}{768H}}^C$ and $p_0 \in [M]$ such that $\bar{\Delta}_{(s_0, a_0)}^{p_0} > 0$, which means that $\operatorname{gap}_{p_0}(s_0, a_0) > 0$. We have that for all $p \in [M]$,

$$\bar{\Delta}_{s_0,a_0}^p = \frac{\Delta_{s_0,a_0}^p}{H-1} = \frac{\operatorname{gap}_p(s_0,a_0)}{H-1} \le \frac{96H(\epsilon/(768H))}{(H-1)} \le \frac{\epsilon}{8(H-1)} \le \frac{\epsilon}{4H}.$$
 (71)

We construct a new problem instance, \mathfrak{M}' , which agrees with \mathfrak{M} except that

$$\mathbb{P}_{p_0}\left(S_1+1 \mid s_0, a_0\right) = \frac{1}{2} - \bar{\Delta}_{s_0, a_0}^{p_0} + 2\bar{\Delta}_{s_0, a_0}^{p_0} = \frac{1}{2} + \bar{\Delta}_{s_0, a_0}^{p_0}$$
$$\mathbb{P}_{p_0}\left(S_1+2 \mid s_0, a_0\right) = \frac{1}{2} + \bar{\Delta}_{s_0, a_0}^{p_0} - 2\bar{\Delta}_{s_0, a_0}^{p_0} = \frac{1}{2} - \bar{\Delta}_{s_0, a_0}^{p_0}$$

 \mathfrak{M}' is an ϵ -MPERL problem instance. To see this, we note that the only change is in $\mathbb{P}_{p_0}(\cdot \mid s_0, a_0)$. In this new instance, it is still true that for any $q \neq p_0$,

$$\left\|\mathbb{P}_{p_0}(\cdot \mid s_0, a_0) - \mathbb{P}_q(\cdot \mid s_0, a_0)\right\|_1 \le 2\left|\bar{\Delta}_{s_0, a_0}^{p_0} + \bar{\Delta}_{s_0, a_0}^q\right| \le \frac{\epsilon}{H}.$$

where the last inequality uses Equation (71) that $\bar{\Delta}_{s_0,a_0}^p \leq \frac{\epsilon}{4H}$ for every $p \in [M]$.

Fix any sublinear regret algorithm Alg. By Lemma \fbox (the divergence decomposition lemma), we have

$$\mathrm{KL}(\mathbb{P}_{\mathfrak{M}},\mathbb{P}_{\mathfrak{M}'}) = \mathbb{E}_{\mathfrak{M}}\left[n_{p_0}^K(s_0,a_0)\right] \mathrm{KL}\left(\mathbb{P}_{p_0}^{\mathfrak{M}}(\cdot \mid s_0,a_0),\mathbb{P}_{p_0}^{\mathfrak{M}'}(\cdot \mid s_0,a_0)\right)$$

Using a similar reasoning as before, and recall that $\bar{\Delta}_{s_0,a_0}^p \leq \frac{1}{24}$, we can show that

$$\mathrm{KL}(\mathbb{P}_{\mathfrak{M}}, \mathbb{P}_{\mathfrak{M}'}) \le 12\mathbb{E}_{\mathfrak{M}}\left[n_{p_0}^K(s_0, a_0)\right] (\bar{\Delta}_{s_0, a_0}^{p_0})^2 \le \frac{K}{48S_1},\tag{72}$$

and consequently, as long as $K \ge 20S_1$,

$$\frac{1}{4} \exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n^{K}(s_{0}, a_{0})\right] (\bar{\Delta}_{s_{0}, a_{0}}^{\min})^{2}\right) \geq \frac{1}{4} \exp(-\frac{K}{48S_{1}}) \geq \exp(-\frac{K}{8S_{1}}).$$

It then follows that

$$\mathbb{E}_{\mathfrak{M}}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right] + \mathbb{E}_{\mathfrak{M}'}\left[\operatorname{Reg}_{\operatorname{Alg}}(K)\right]$$

$$\geq \frac{K}{2}\left((H-1)\bar{\Delta}_{s_{0},a_{0}}^{p_{0}}\right)\left(\frac{1}{2}\exp\left(-\operatorname{KL}(\mathbb{P}_{\mathfrak{M}},\mathbb{P}_{\mathfrak{M}'})\right) - \exp\left(-\frac{K}{8S_{1}}\right)\right)$$

$$\geq \frac{K}{8}\left((H-1)\bar{\Delta}_{s_{0},a_{0}}^{p_{0}}\right)\exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n_{p_{0}}^{K}(s_{0},a_{0})\right](\bar{\Delta}_{s_{0},a_{0}}^{p_{0}})^{2}\right).$$

Now, under the assumption that Alg is a sublinear regret algorithm, we have

$$\frac{K}{8} \left((H-1)\bar{\Delta}_{s_0,a_0}^{p_0} \right) \exp\left(-12\mathbb{E}_{\mathfrak{M}}\left[n_{p_0}^K(s_0,a_0) \right] (\bar{\Delta}_{s_0,a_0}^{p_0})^2 \right) \le 2CK^{\alpha}.$$

It follows that

I

$$\mathbb{E}_{\mathfrak{M}}\left[n_{p_{0}}^{K}(s_{0}, a_{0})\right] \geq \frac{1}{12\left(\bar{\Delta}_{s_{0}, a_{0}}^{p_{0}}\right)^{2}} \ln\left(\frac{(H-1)\bar{\Delta}_{s_{0}, a_{0}}^{p_{0}}K^{1-\alpha}}{16C}\right)$$
$$\geq \frac{H^{2}}{24\left(\operatorname{gap}_{p_{0}}(s_{0}, a_{0})\right)^{2}} \ln\left(\frac{\operatorname{gap}_{p_{0}}(s_{0}, a_{0})K^{1-\alpha}}{16C}\right).$$

We then have that

$$\mathbb{E}_{\mathfrak{M}}\left[n_{p_0}^{K}(s_0, a_0)\right] \ge \Omega\left(\frac{H^2}{\left(\operatorname{gap}_{p_0}(s_0, a_0)\right)^2}\ln K\right).$$

1

`

Combing the two claims: We note that in \mathfrak{M} , for any $(s, a, p) \in (\mathcal{S} \setminus \mathcal{S}_1) \times \mathcal{A} \times [M]$, $\operatorname{gap}_p(s, a) = 0$. It then follows from Lemma 30 (the regret decomposition lemma) and the fact that for any $(s, a, p) \in \mathcal{I}_{\epsilon/768H} \times [M]$, $\operatorname{gap}_p(s, a) > 0$, that

$$\mathbb{E}\left[\operatorname{Reg}_{A\lg}(K)\right] \ge \sum_{p=1}^{M} \sum_{\substack{(s,a) \in \mathcal{S}_1 \times \mathcal{A}}} \mathbb{E}\left[n_p^K(s,a)\right] \operatorname{gap}_p(s,a)$$
$$\ge \Omega \left(\ln K \left(\sum_{\substack{p \in [M] \\ gap_p(s,a) > 0}} \sum_{\substack{(s,a) \in \mathcal{I}_{\epsilon/768H}^C : \\ gap_p(s,a) > 0}} \frac{H^2}{\operatorname{gap}_p(s,a)} + \sum_{\substack{(s,a) \in \mathcal{I}_{\epsilon/768H}}} \frac{H^2}{\operatorname{min}_p \operatorname{gap}_p(s,a)}\right)\right).$$