

# The Blessing of Heterogeneity in Federated Q-Learning: Linear Speedup and Beyond

**Jiin Woo**

JIINW@ANDREW.CMU.EDU

*Department of Electrical and Computer Engineering  
Carnegie Mellon University  
Pittsburgh, PA 15213, USA*

**Gauri Joshi**

GAURIJ@ANDREW.CMU.EDU

*Department of Electrical and Computer Engineering  
Carnegie Mellon University  
Pittsburgh, PA 15213, USA*

**Yuejie Chi**

YUEJIEC@ANDREW.CMU.EDU

*Department of Electrical and Computer Engineering  
Carnegie Mellon University  
Pittsburgh, PA 15213, USA*

**Editor:** Gergely Neu

## Abstract

In this paper, we consider federated Q-learning, which aims to learn an optimal Q-function by *periodically* aggregating local Q-estimates trained on local data alone. Focusing on infinite-horizon tabular Markov decision processes, we provide sample complexity guarantees for both the synchronous and asynchronous variants of federated Q-learning, which exhibit a linear speedup with respect to the number of agents and near-optimal dependencies on other salient problem parameters.

In the asynchronous setting, existing analyses of federated Q-learning, which adopt an equally weighted averaging of local Q-estimates, require that every agent covers the entire state-action space. In contrast, our improved sample complexity scales inverse proportionally to the minimum entry of the *average* stationary state-action occupancy distribution of all agents, thus only requiring the agents to *collectively cover* the entire state-action space, unveiling the *blessing of heterogeneity*. However, its sample complexity still suffers when the local trajectories are highly heterogeneous. In response, we propose a novel federated Q-learning algorithm with importance averaging, giving larger weights to more frequently visited state-action pairs, which achieves a robust linear speedup as if all trajectories are centrally processed, regardless of the heterogeneity of local behavior policies.

**Keywords:** federated RL, Q-learning, sample complexity, linear speedup, heterogeneity

## 1. Introduction

Reinforcement Learning (RL) (Sutton and Barto, 2018) is an area of machine learning for sequential decision making, aiming to learn an optimal policy that maximizes the total rewards via interactions with an unknown environment. RL is widely used in many real-world applications, such as autonomous driving, games, clinical trials, and recommendation systems. However, due to the high dimensionality of the state-action space, training of RL

agents typically requires a significant amount of computation and data to achieve desirable performance. Moreover, data collection can be extremely time-consuming with limited access in the wild, especially when performed by a single agent. On the other hand, it is possible to leverage multiple agents to collect data simultaneously, under the premise that they can learn a global policy collaboratively with the aid of a central server without the need of sharing local data. As a result, there is a growing need to conduct RL in a distributed or federated fashion.

Although there have been many studies analyzing federated learning (Kairouz et al., 2021) in other areas such as supervised machine learning (McMahan et al., 2017; Bonawitz et al., 2019; Wang et al., 2020b), there are only a few recent works focused on federated RL. They consider issues such as robustness to adversarial attacks (Wu et al., 2021; Fan et al., 2021), environment heterogeneity (Jin et al., 2022), as well as sample and communication complexities (Doan et al., 2021; Khodadadian et al., 2022; Shen et al., 2022). Encouragingly, some of these prior works offer non-asymptotic sample complexity analyses of federated RL algorithms that highlight a linear speedup of the required sample size in terms of the number of agents. However, the performance characterization of these federated algorithms is still far from complete.

### 1.1 Federated Q-learning: prior art and limitations

This paper focuses on Q-learning (Watkins and Dayan, 1992), one of the most celebrated model-free RL algorithms, which aims to learn the optimal Q-function directly without forming an estimate of the model. Two sampling protocols are typically studied: synchronous sampling and asynchronous sampling. With synchronous sampling, all state-action pairs are updated uniformly assuming access to a generative model or a simulator (Kearns and Singh, 1999). With asynchronous sampling, only the state-action pair that is visited by the behavior policy is updated at each time (Tsitsiklis, 1994). Despite its long history of theoretical investigation, the tight sample complexity of Q-learning in the single-agent setting has only recently been pinned down in Li et al. (2023). As we shall elucidate, there remains a large gap in terms of the sample complexity requirement between the federated setting and the single-agent setting in terms of dependencies on salient problem parameters.

To harness the power of multiple agents, Khodadadian et al. (2022) proposed and analyzed a federated variant of Q-learning with *asynchronous* sampling that periodically aggregates the local Q-estimates trained on local Markovian trajectories collected over  $K$  agents. To set the stage, consider an infinite-horizon tabular Markov decision process (MDP) with state space  $\mathcal{S}$ , action space  $\mathcal{A}$ , and a discount factor  $\gamma \in [0, 1)$ . To learn an  $\varepsilon$ -optimal Q-function estimate (in the  $\ell_\infty$  sense), Khodadadian et al. (2022) requires a per-agent sample size on the order of

$$\tilde{O}\left(\frac{|\mathcal{S}|^2}{K\mu_{\min}^5(1-\gamma)^9\varepsilon^2}\right) \quad (1)$$

for sufficiently small  $\varepsilon$ , where  $\mu_{\min} := \min_{1 \leq k \leq K} \min_{(s,a) \in \mathcal{S} \times \mathcal{A}} \mu_{\mathbf{b}}^k(s, a)$  is the minimum entry of the stationary state-action occupancy distributions  $\mu_{\mathbf{b}}^k$  of the sample trajectories over all agents, and  $\tilde{O}$  hides logarithmic terms. On the other hand, the sample requirement of

single-agent Q-learning (Li et al., 2023) for learning an  $\varepsilon$ -optimal Q-function is

$$\tilde{O}\left(\frac{1}{\mu_{\min}(1-\gamma)^4\varepsilon^2}\right) \quad (2)$$

for sufficiently small  $\varepsilon$ . Comparing the two sample complexity bounds reveals several drawbacks of existing analyses and raises the following natural questions.

- *Near-optimal sample size.* Despite the appealing linear speedup in terms of the number of agents  $K$  shown in Khodadadian et al. (2022), it has unfavorable dependencies on other salient problem parameters. In particular, since  $1/\mu_{\min} \geq |\mathcal{S}||\mathcal{A}|$ , the sample complexity in (1) will be better than that of the single-agent case in (2) only if  $K$  is at least above the order of  $\frac{|\mathcal{S}|^6|\mathcal{A}|^4}{(1-\gamma)^5}$ , which may not be practically feasible with large state-action space and long effective horizon. *Can we improve the dependency on the salient problem parameters for federated Q-learning while maintaining linear speedup?*
- *Benefits of heterogeneity.* Existing analyses in Khodadadian et al. (2022) require that each agent has a full coverage of the state-action space (i.e.,  $\mu_{\min} > 0$ ), which is as stringent as the single-agent setting. However, given that the insufficient coverage of individual agents can be complemented by each other when agents have heterogeneous local trajectories, it may not be necessary to require full coverage of the state-action space from every agent. *Can we exploit the heterogeneity in the agents' local trajectories and relax the coverage requirement on individual agents?*

## 1.2 Summary of our contributions

In this paper, we answer these questions in the affirmative, by providing a sample complexity analysis of federated Q-learning under both the synchronous and asynchronous settings. The main contributions are summarized as follows, with Table 1 providing a comparison with the prior art.

- *Sample complexity of federated synchronous Q-learning with equal averaging.* We show that with high probability, the sample complexity of federated synchronous Q-learning (FedSynQ) to learn an  $\varepsilon$ -optimal Q-function in the  $\ell_\infty$  sense is (see Theorem 1)

$$\tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|}{K(1-\gamma)^5\varepsilon^2}\right), \quad (3)$$

which exhibits a linear speedup with respect to the number of agents  $K$  and nearly matches the tight sample complexity bound of single-agent synchronous Q-learning up to a factor of  $1/(1-\gamma)$  in Li et al. (2023) for  $K = 1$ .

- *Sample complexity of federated asynchronous Q-learning with equal averaging.* We provide a sharpened sample complexity analysis of the algorithm developed in Khodadadian et al. (2022) for federated asynchronous Q-learning with equal averaging (FedAsynQ-EqAvg) that leads to new insights. To learn an  $\varepsilon$ -optimal Q-function in the  $\ell_\infty$  sense, FedAsynQ-EqAvg requires at most (see Theorem 2)

$$\tilde{O}\left(\frac{C_{\text{het}}}{K\mu_{\text{avg}}(1-\gamma)^5\varepsilon^2}\right) \quad (4)$$

sampling	reference	number of agents	coverage	sample complexity
synchronous	Wainwright (2019a); Chen et al. (2020)	1	full	$\frac{ \mathcal{S}  \mathcal{A} }{(1-\gamma)^5 \varepsilon^2}$
	Li et al. (2023)	1	full	$\frac{ \mathcal{S}  \mathcal{A} }{(1-\gamma)^4 \varepsilon^2}$
	FedSynQ (Theorem 1)	$K$	full	$\frac{ \mathcal{S}  \mathcal{A} }{K(1-\gamma)^5 \varepsilon^2}$
asynchronous	Qu and Wierman (2020)	1	full	$\frac{t_{\text{mix}}}{\mu_{\min}^2 (1-\gamma)^5 \varepsilon^2}$
	Li et al. (2021b)	1	full	$\frac{1}{\mu_{\min} (1-\gamma)^5 \varepsilon^2}$
	Li et al. (2023)	1	full	$\frac{1}{\mu_{\min} (1-\gamma)^4 \varepsilon^2}$
	FedAsynQ-EqAvg (Khodadadian et al., 2022)	$K$	full	$\frac{ \mathcal{S} ^2}{K \mu_{\min}^5 (1-\gamma)^9 \varepsilon^2}$
	FedAsynQ-EqAvg (Theorem 2)	$K$	partial	$\frac{C_{\text{het}}}{K \mu_{\text{avg}} (1-\gamma)^5 \varepsilon^2}$
	FedAsynQ-ImAvg (Theorem 3)	$K$	partial	$\frac{1}{K \mu_{\text{avg}} (1-\gamma)^5 \varepsilon^2}$

Table 1: Comparison of sample complexity upper bounds of single-agent and federated Q-learning algorithms under synchronous and asynchronous sampling protocols to learn an  $\varepsilon$ -optimal Q-function in the  $\ell_\infty$  sense, where logarithmic factors and burn-in costs are hidden. Here,  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space,  $\gamma$  is the discount factor,  $K$  is the total number of agents, and  $t_{\text{mix}}$  is the mixing time of the behavior policy. In addition,  $\mu_{\min} = \min_{k,s,a} \mu_{\text{b}}^k(s,a)$  denotes the minimum entry of the stationary state-action occupancy distributions  $\mu_{\text{b}}^k$  of all agents,  $\mu_{\text{avg}} := \min_{s,a} \frac{1}{K} \sum_{k=1}^K \mu_{\text{b}}^k(s,a)$  denotes the minimum entry of the average stationary state-action occupancy distribution of all agents, and  $C_{\text{het}} := \max_{k,s,a} K \mu_{\text{b}}^k(s,a) / (\sum_{k=1}^K \mu_{\text{b}}^k(s,a))$  captures the heterogeneity across the agents.

samples per agent for sufficiently small  $\varepsilon$  (ignoring the burn-in cost that depends on the mixing times of the Markovian trajectories over all agents), where  $\mu_{\text{avg}}$  is the minimum entry of the *average* stationary state-action occupancy distribution of all agents, i.e.,  $\mu_{\text{avg}} = \min_{s,a} \frac{1}{K} \sum_{k=1}^K \mu_{\text{b}}^k(s,a) \geq \mu_{\min}$ , and  $C_{\text{het}} = \max_{k,s,a} \frac{K \mu_{\text{b}}^k(s,a)}{\sum_{k=1}^K \mu_{\text{b}}^k(s,a)} \in [1, 1/\mu_{\text{avg}}]$  captures the heterogeneity of the behavior policies across agents. This sample complexity not only proves a linear speedup with respect to the number of agents, but also greatly sharpens the dependency on all the salient problem parameters — including  $1/(1-\gamma)$ ,  $|\mathcal{S}|$ , and  $1/\mu_{\min}$  — by orders of magnitudes compared to the bound obtained in Khodadadian et al. (2022). More importantly, it uncovers that as long as the agents collectively cover the entire state-action space (i.e.,  $\mu_{\text{avg}} > 0$ ), FedAsynQ-EqAvg still enables learning even when individual agents fail to cover the entire state-action space (i.e.,  $\mu_{\min} = 0$ ), unveiling the blessing of heterogeneity that was not elucidated in prior work (Khodadadian et al., 2022).

- *Sample complexity of federated asynchronous Q-learning with importance averaging.* Although heterogeneous behavior policies at agents may induce local trajectories cover-

ing different parts of the state-action space and relax the coverage requirement, equally weighting the local Q-estimates may hinder the convergence which is bottlenecked by the slowest converging agent. This is evident by the dependency on  $C_{\text{het}}$  in the sample complexity of FedAsynQ-EqAvg, which becomes larger when the local behavior policies are highly disparate. To address this issue, we propose a novel importance averaging scheme in federated Q-learning (FedAsynQ-ImAvg) that averages the local Q-estimates by assigning larger weights to more frequently updated local estimates. To learn an  $\varepsilon$ -optimal Q-function in the  $\ell_\infty$  sense, FedAsynQ-ImAvg requires at most (see Theorem 3)

$$\tilde{O}\left(\frac{1}{K\mu_{\text{avg}}(1-\gamma)^5\varepsilon^2}\right) \quad (5)$$

samples per agent for sufficiently small  $\varepsilon$  (ignoring the burn-in cost that depends on the mixing times of the Markovian trajectories over all agents). This improves over that of FedAsynQ-EqAvg by removing the dependency on  $C_{\text{het}}$ , which can be as large as  $1/\mu_{\text{avg}}$ . More importantly, this suggests that FedAsynQ-ImAvg achieves stable linear speedup with respect to the profile of the local behavior policies while maintaining the blessing of heterogeneity that eases the burden of individual agents' coverage.

### 1.3 Related work

**Analysis of single-agent Q-learning.** There has been extensive research on the convergence guarantees of Q-learning, focusing on the single-agent case. Many initial studies have analyzed the asymptotic convergence of Q-learning (Tsitsiklis, 1994; Szepesvári, 1998; Jaakkola et al., 1994; Borkar and Meyn, 2000). Later, Even-Dar and Mansour (2003); Beck and Srikant (2012); Wainwright (2019a); Chen et al. (2020); Li et al. (2023) have studied the sample complexity of Q-learning under synchronous sampling, and Even-Dar and Mansour (2003); Beck and Srikant (2012); Qu and Wierman (2020); Li et al. (2023, 2021b); Chen et al. (2021b) have investigated the finite-time convergence of Q-learning under asynchronous sampling (also referred to as Markovian sampling). In addition, Jin et al. (2018); Bai et al. (2019); Zhang et al. (2020); Li et al. (2021a); Yang et al. (2021) studied Q-learning with optimism for online RL, and Shi et al. (2022); Yan et al. (2022) dealt with Q-learning with pessimism for offline RL.

**Distributed and federated RL.** Several recent works have developed distributed versions of RL algorithms to accelerate training (Mnih et al., 2016; Espeholt et al., 2018; Assran et al., 2019). Theoretical analysis of convergence and communication efficiency of these distributed RL algorithms have also been considered in recent works. For example, a collection of works (Doan et al., 2019; Sun et al., 2020; Wang et al., 2020a; Wai, 2020; Chen et al., 2022a; Zeng et al., 2021) have analyzed the convergence of decentralized temporal difference (TD) learning. Furthermore, Chen et al. (2022b); Shen et al. (2022) have analyzed the finite-time convergence of distributed actor-critic algorithms and Chen et al. (2021a) proposed a communication-efficient policy gradient algorithm with provable convergence guarantees.

**Notation.** Throughout this paper, we denote by  $\Delta(\mathcal{S})$  the probability simplex over a set  $\mathcal{S}$ , and  $[K] := \{1, \dots, K\}$  for any positive integer  $K > 0$ . In addition,  $f(\cdot) = \tilde{O}(g(\cdot))$  or

$f \lesssim g$  (resp.  $f(\cdot) = \tilde{\Omega}(g(\cdot))$  or  $f \gtrsim g$ ) means that  $f(\cdot)$  is orderwise no larger than (resp. no smaller than)  $g(\cdot)$  modulo some logarithmic factors. The notation  $f \asymp g$  means  $f \lesssim g$  and  $f \gtrsim g$  hold simultaneously.

## 2. Model and background

In this section, we introduce the mathematical model and background of Markov decision processes.

**Infinite-horizon Markov decision process.** We consider an infinite-horizon Markov decision process (MDP), which is represented by  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$ . Here,  $\mathcal{S}$  and  $\mathcal{A}$  denote the state space and the action space, respectively,  $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  indicates the transition kernel such that  $P(s' | s, a)$  denotes the probability that action  $a$  in state  $s$  leads to state  $s'$ ,  $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  denotes a deterministic reward function, where  $r(s, a)$  is the immediate reward for action  $a$  in state  $s$ , and  $\gamma \in [0, 1)$  is the discount factor.

**Policy, value function, and Q-function.** A *policy* is an action-selection rule denoted by the mapping  $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ , such that  $\pi(a|s)$  is the probability of taking action  $a$  in state  $s$ . For a given policy  $\pi$ , the *value function*  $V^\pi : \mathcal{S} \rightarrow \mathbb{R}$ , which measures the expected discounted cumulative reward from an initial state  $s$ , is defined as

$$\forall s \in \mathcal{S} : \quad V^\pi(s) := \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s \right]. \quad (6)$$

Here, the expectation is taken with respect to the randomness of the trajectory  $\{s_t, a_t, r_t\}_{t=0}^{\infty}$ , sampled based on the transition kernel (i.e.,  $s_{t+1} \sim P(\cdot | s_t, a_t)$ ) and the policy  $\pi$  (i.e.,  $a_t \sim \pi(\cdot | s_t)$ ) for any  $t \geq 0$ . Similarly, the state-action value function (i.e., *Q-function*)  $Q^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ , which measures the expected discounted cumulative reward from an initial state-action pair  $(s, a)$ , is defined as

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A} : \quad Q^\pi(s, a) := r(s, a) + \mathbb{E} \left[ \sum_{t=1}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a \right].$$

Again here, the expectation is taken with respect to the randomness of the trajectory  $\{s_t, a_t, r_t\}_{t=1}^{\infty}$  generated similarly as above. Since the rewards lie within  $[0, 1]$ , it follows that for any policy  $\pi$ ,

$$0 \leq V^\pi \leq \frac{1}{1-\gamma}, \quad 0 \leq Q^\pi \leq \frac{1}{1-\gamma}. \quad (7)$$

**Optimal policy and Bellman's principle of optimality.** A policy that maximizes the value function uniformly over all states is called an *optimal policy* and denoted by  $\pi^*$ . Note that the existence of such an optimal policy is always guaranteed (Puterman, 2014), which also maximizes the Q-function simultaneously. The corresponding optimal value function and Q-function are denoted by  $V^* := V^{\pi^*}$  and  $Q^* := Q^{\pi^*}$ , respectively. It is well-known that the optimal Q-function  $Q^*$  can be determined as the unique fixed point of the Bellman operator  $\mathcal{T}$ , given by

$$\mathcal{T}(Q)(s, a) := r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} \left[ \max_{a' \in \mathcal{A}} Q(s', a') \right]. \quad (8)$$

Q-learning (Watkins and Dayan, 1992), perhaps the most widely used model-free RL algorithm, seeks to learn the optimal Q-function based on samples collected from the underlying MDP without estimating the model.

### 3. Federated synchronous Q-learning: algorithm and theory

In this section, we begin with understanding federated synchronous Q-learning, where all the state-action pairs are updated simultaneously assuming access to a generative model or simulator at all the agents.

#### 3.1 Problem setting

In the synchronous setting, each agent  $k \in [K]$  has access to a generative model, and generates a new sample

$$s_t^k(s, a) \sim P(\cdot | s, a) \quad (9)$$

for every state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$  *independently* at every iteration  $t$ . Our goal is to learn the optimal Q-function  $Q^*$  collaboratively by aggregating the local Q-learning estimates *periodically*.

**Review: synchronous Q-learning with a single agent.** To facilitate algorithmic development, let us recall the synchronous Q-learning update rule with a single agent. Starting with certain initialization  $Q_0$ , at every iteration  $t \geq 1$ , the Q-function is updated according to

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A} : \quad Q_t(s, a) = (1 - \eta)Q_{t-1}(s, a) + \eta \left( r(s, a) + \gamma \max_{a' \in \mathcal{A}} Q_{t-1}(s_t(s, a), a') \right), \quad (10)$$

where  $s_t(s, a) \sim P(\cdot | s, a)$  is drawn independently for every state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , and  $\eta$  denotes the constant learning rate. The sample complexity of synchronous Q-learning has been recently investigated and sharpened in a number of works, e.g. Li et al. (2023); Wainwright (2019a); Chen et al. (2020).

#### 3.2 Algorithm description

We propose a natural federated synchronous Q-learning algorithm called FedSynQ that alternates between local updates at agents and periodic averaging at a central server. The complete description is summarized in Algorithm 1. FedSynQ initializes a local Q-function as  $Q_0^k = Q_0$  at each agent  $k \in [K]$ . Suppose at the beginning of each iteration  $t \geq 1$ , each agent maintains a local Q-function estimate  $Q_{t-1}^k$  and a local value function estimate  $V_{t-1}^k$ , which are related via

$$\forall s \in \mathcal{S} : \quad V_t^k(s) := \max_{a \in \mathcal{A}} Q_t^k(s, a). \quad (11)$$

FedSynQ proceeds according to the following steps in the rest of the  $t$ -th iteration.

1. *Local updates:* Each agent first independently updates *all* entries of its Q-estimate  $Q_{t-1}^k$  to reach some *intermediate* estimate following the update rule:

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}: \quad Q_{t-\frac{1}{2}}^k(s, a) = (1 - \eta)Q_{t-1}^k(s, a) + \eta \left( r(s, a) + \gamma V_{t-1}^k(s_t^k(s, a)) \right), \quad (12)$$

where  $s_t^k(s, a)$  is drawn according to (9), and  $\eta \geq 0$  is the learning rate.

2. *Periodic averaging:* These intermediate estimates will be periodically averaged by the server to form the updated estimate  $Q_t^k$  at the end of the  $t$ -th iteration. Formally, denoting  $\tau \geq 1$  as the synchronization period, it follows

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}: \quad Q_t^k(s, a) = \begin{cases} \frac{1}{K} \sum_{k=1}^K Q_{t-\frac{1}{2}}^k(s, a) & \text{if } t \equiv 0 \pmod{\tau} \\ Q_{t-\frac{1}{2}}^k(s, a) & \text{otherwise} \end{cases}. \quad (13)$$

Denoting the number of total iterations by  $T$ , the algorithm outputs the final Q-estimate as the average of all local estimates, i.e.  $Q_T = \frac{1}{K} \sum_k Q_T^k$ . Without loss of generality, we assume the total number of iterations  $T$  is divisible by  $\tau$ , where  $C_{\text{round}} = T/\tau$  is the rounds of communication.

---

**Algorithm 1:** Federated Synchronous Q-learning (FedSynQ)
 

---

- 1: **inputs:** learning rate  $\eta$ , discount factor  $\gamma$ , number of agents  $K$ , synchronization period  $\tau$ , number of iterations  $T$ .
  - 2: **initialization:**  $Q_0^k = Q_0$  for all  $k$ .
  - 3: **for**  $t = 1, \dots, T$  **do**
  - 4:   **for**  $k \in [K]$  **do**
  - 5:     Draw  $s_t^k(s, a) \sim P(\cdot | s, a)$  for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ .
  - 6:     Compute  $Q_{t-\frac{1}{2}}^k$  according to (12).
  - 7:     Compute  $Q_t^k$  according to (13).
  - 8:   **end for**
  - 9: **end for**
  - 10: **return:**  $Q_T = \frac{1}{K} \sum_k Q_T^k$ .
- 

### 3.3 Performance guarantee

We are ready to provide the finite-time convergence analysis of Algorithm 1.

**Theorem 1 (Finite-time convergence of FedSynQ)** *Consider any given  $\delta \in (0, 1)$  and  $\varepsilon \in (0, \frac{1}{1-\gamma}]$ . Suppose that the initialization of Algorithm 1 satisfies  $0 \leq Q_0 \leq \frac{1}{1-\gamma}$ , and the synchronization period  $\tau$  obeys*

$$\tau \leq 1 + \frac{1}{\eta} \min \left\{ \frac{1-\gamma}{8\gamma}, \frac{1}{K} \right\}. \quad (14a)$$



There exist some sufficiently large constant  $c_T > 0$  and sufficiently small constant  $c_\eta > 0$ , both independent of problem parameters, such that with probability at least  $1 - \delta$ , the output of Algorithm 1 satisfies  $\|Q_T - Q^*\|_\infty \leq \varepsilon$ , provided that the sample size per agent  $T$  and the learning rate  $\eta$  satisfy

$$T \geq \frac{c_T}{K(1-\gamma)^5 \varepsilon^2} (\log((1-\gamma)^2 \varepsilon))^2 \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}, \quad (14b)$$

$$\eta = c_\eta K(1-\gamma)^4 \varepsilon^2 \frac{1}{\log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}}. \quad (14c)$$

Theorem 1 suggests that to achieve an  $\varepsilon$ -accurate Q-function estimate in an  $\ell_\infty$  sense, the number of samples required at each agent is no more than

$$\tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|}{K(1-\gamma)^5 \varepsilon^2}\right),$$

given that the agent collects  $|\mathcal{S}||\mathcal{A}|$  samples at each iteration. A few implications are in order.

**Linear speedup.** The sample complexity exhibits an appealing linear speedup with respect to the number of agents  $K$ . In comparison, the sharpest upper bound known for single-agent Q-learning (Li et al., 2023) is  $\tilde{O}\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^4 \min\{\varepsilon, \varepsilon^2\}}\right)$ , which matches with its algorithmic-dependent lower bound when  $\varepsilon \in (0, 1)$ . Therefore, our federated setting enables faster learning as soon as the number of agents satisfies

$$K \gtrsim \frac{1}{(1-\gamma) \max\{1, \varepsilon\}}$$

up to logarithmic factors. When  $K = 1$ , our bound nearly matches with the lower bound of single-agent Q-learning up to a factor of  $1/(1-\gamma)$ , indicating its near-optimality.

**Communication efficiency.** One key feature of our federated setting is the use of periodic averaging with the hope to improve communication efficiency. According to (14a), our theory requires that the synchronization period  $\tau$  be inversely proportional to the learning rate  $\eta$ , which suggests that more frequent communication is needed to compensate the discrepancy of local updates when the learning rate is large. To provide insights, consider the parameter regime when  $K \gtrsim \frac{1}{1-\gamma}$  and  $\varepsilon \lesssim \frac{1}{K(1-\gamma)^2}$ . Plugging the choice of the learning rate (14c) into the upper bound of  $\tau$  in (14a), we can choose the synchronization period as  $\tau \asymp \frac{1}{K^2(1-\gamma)^4 \varepsilon^2}$  up to logarithmic factors, leading to a communication complexity no larger than  $C_{\text{round}} = \frac{T}{\tau} \lesssim \frac{K}{1-\gamma}$ , which is almost independent of the final accuracy  $\varepsilon$ .

#### 4. Federated asynchronous Q-learning: algorithm and theory

In this section, we study the sample complexity of federated asynchronous Q-learning, where  $K$  agents sample local trajectories using different behavior policies. In particular, we propose a novel aggregation algorithm **FedAsynQ-lmAvg** that leverages the heterogeneity of these policies and dramatically improves the sample complexity.

#### 4.1 Problem setting

In the asynchronous setting, each agent  $k \in [K]$  independently collects a sample trajectory  $\{s_t^k, a_t^k, r_t^k\}_{t=0}^\infty$  from the same underlying MDP  $\mathcal{M}$  following some stationary *local* behavior policy  $\pi_b^k$  such that

$$a_t^k \sim \pi_b^k(\cdot | s_t^k), \quad r_t^k = r(s_t^k, a_t^k), \quad s_{t+1}^k \sim P(\cdot | s_t^k, a_t^k) \quad (15)$$

for all  $t \geq 0$ , where the initial state is initialized as  $s_0^k$  for each agent  $k$ . Note that the behavior policies  $\{\pi_b^k\}_{k \in [K]}$  are heterogeneous across agents and can be different from the optimal policy  $\pi^*$ . Contrary to the generative model considered in the synchronous setting, the samples collected under the asynchronous setting are no longer independent across time but are Markovian, making the analysis significantly more challenging. The sample trajectory at each agent can be viewed as sampling a time-homogeneous Markov chain over the set of state-action pairs. Throughout this paper, we make the following standard uniform ergodicity assumption (Paulin, 2015; Li et al., 2021b).

**Assumption 1 (Uniform ergodicity)** *For every agent  $k \in [K]$ , the Markov chain induced by the stationary behavior policy  $\pi_b^k$  is uniformly ergodic over the entire state-action space  $\mathcal{S} \times \mathcal{A}$ .*

Uniform ergodicity guarantees that the distribution of the state-action pair  $(s_t, a_t)$  of a trajectory converges to the stationary distribution of the Markov chain geometrically fast regardless of the initial state-action pair, and eventually, each state-action pair is visited in proportion to the stationary distribution.

**Key parameters.** Two important quantities concerning the resulting Markov chains will govern the performance guarantees. The first one is the stationary state-action distribution  $\mu_b^k$ , which is the stationary distribution of the Markov chain induced by  $\pi_b^k$  over all state-action pairs; the second one is  $t_{\text{mix}}^k$ , which is the mixing time of the same Markov chain given by

$$t_{\text{mix}}^k := \min \left\{ t \mid \max_{(s_0, a_0) \in \mathcal{S} \times \mathcal{A}} d_{\text{TV}}(P_t^k(\cdot | s_0, a_0), \mu_b^k) \leq \frac{1}{4} \right\}, \quad (16)$$

where  $P_t^k(\cdot | s_0, a_0)$  denote the distribution of  $(s_t, a_t)$  conditioned on  $(s_0, a_0)$  for agent  $k$ , and  $d_{\text{TV}}(\cdot, \cdot)$  is the total variation distance. Further, let the largest mixing time of all the Markov chains induced by local behavior policies be

$$t_{\text{mix}}^{\max} := \max_{k \in [K]} t_{\text{mix}}^k. \quad (17)$$

In words,  $t_{\text{mix}}^{\max}$  approximately indicates the time that the transition of every agent starts to follow its stationary distribution regardless of its initial state.

Let us further define a few key parameters that measure the coverage and heterogeneity of the stationary state-action distribution  $\mu_b^k$  across agents. First, define

$$\mu_{\min} := \min_{k \in [K]} \mu_{\min}^k, \quad \text{where} \quad \mu_{\min}^k := \min_{(s, a) \in \mathcal{S} \times \mathcal{A}} \mu_b^k(s, a). \quad (18)$$

State-action pairs with small stationary probabilities are visited less frequently, and therefore can become bottlenecks in improving the quality of Q-function estimates. Clearly,  $\mu_{\min} \leq \frac{1}{|\mathcal{S}||\mathcal{A}|}$ . In addition, denote

$$\mu_{\text{avg}} := \min_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{1}{K} \sum_{k=1}^K \mu_{\text{b}}^k(s, a). \quad (19)$$

In words,  $\mu_{\text{avg}}$  is the minimum entry of the *average* stationary state-action distribution of all agents. The difference between  $\mu_{\text{avg}}$  and  $\mu_{\min}$  stands out when an individual agent fails to cover the entire state-action space. While  $\mu_{\min} = 0$  in such a case,  $\mu_{\text{avg}}$  can still be positive as long as each state-action pair is explored by at least one of the agents, i.e.,  $\sum_{k=1}^K \mu_{\text{b}}^k(s, a) > 0$ . Note that  $\mu_{\text{avg}}$  is always greater than or equal to  $\mu_{\min}$  since

$$\mu_{\text{avg}} = \min_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{1}{K} \sum_{k=1}^K \mu_{\text{b}}^k(s, a) \geq \min_{(s,a) \in \mathcal{S} \times \mathcal{A}, k \in [K]} \mu_{\text{b}}^k(s, a) = \mu_{\min}. \quad (20)$$

Last but not least, we measure the heterogeneity of the stationary state-action distributions across agents by

$$C_{\text{het}} := \max_{k \in [K]} \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{\mu_{\text{b}}^k(s, a)}{\frac{1}{K} \sum_{k=1}^K \mu_{\text{b}}^k(s, a)}, \quad (21)$$

which satisfies  $1 \leq C_{\text{het}} \leq \min\{K, 1/\mu_{\text{avg}}\}$ , and in particular,  $C_{\text{het}} = 1$  when  $\mu_{\text{b}}^k = \mu_{\text{b}}$  are all equal.

**Review: asynchronous Q-learning with a single agent.** Recall the update rule of asynchronous Q-learning with a single agent, where at each iteration  $t \geq 1$ , upon receiving a transition  $(s_{t-1}, a_{t-1}, s_t)$ , the Q-estimate is updated via

$$Q_t(s, a) = \begin{cases} (1 - \eta)Q_{t-1}(s, a) + \eta(r(s, a) + \gamma \max_{a' \in \mathcal{A}} Q_{t-1}(s, a')), & \text{if } (s, a) = (s_{t-1}, a_{t-1}), \\ Q_{t-1}(s, a) & \text{otherwise,} \end{cases} \quad (22)$$

where  $\eta$  denotes the learning rate and  $V_t$  is defined in (11). The sample complexity of asynchronous Q-learning has been recently investigated in Li et al. (2021b, 2023); Qu and Wierman (2020).

## 4.2 Algorithm description

Similar to the synchronous setting, we describe a federated asynchronous Q-learning algorithm, called FedAsynQ (see Algorithm 2), that learns the optimal Q-function by periodically averaging the local Q-estimates with the aid of a central server. See Figure 1 for an illustration. Inheriting the notation of  $Q_t^k$  and  $V_t^k$  from the synchronous setting (cf. (11)), FedAsynQ proceeds as follows in the rest of the  $t$ -th iteration.

1. *Local updates:* Each agent  $k$  samples a transition  $(s_{t-1}^k, a_{t-1}^k, r_{t-1}^k, s_t^k)$  from its Markovian trajectory generated by the behavior policy  $\pi_{\text{b}}^k$  according to (15) and updates a *single*

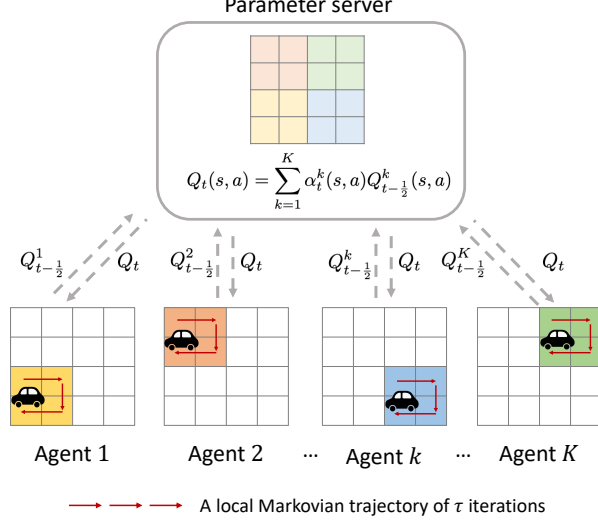


Figure 1: Federated asynchronous Q-learning with  $K$  agents and a parameter server. Each agent  $k$  performs  $\tau$  local updates on its local Q-table along a Markovian trajectory induced by behavior policy  $\pi_b^k$  and sends the Q-table to the server. The server averages and synchronizes the local Q-tables every  $\tau$  iterations. For importance averaging, the agents additionally send the number of visits over all the state-action pairs within each synchronization period, which is not pictured.

entry of its local Q-estimate  $Q_{t-1}^k$ :

$$Q_{t-\frac{1}{2}}^k(s, a) = \begin{cases} (1 - \eta)Q_{t-1}^k(s, a) + \eta(r_{t-1}^k + \gamma V_{t-1}^k(s_t^k)) & \text{if } (s, a) = (s_{t-1}^k, a_{t-1}^k) \\ Q_{t-1}^k(s, a), & \text{otherwise} \end{cases}, \quad (23)$$

where  $\eta$  denotes the learning rate.

2. *Periodic averaging:* The intermediate local estimates will be averaged every  $\tau$  iterations, where  $\tau \geq 1$  is the synchronization period. Here, we consider a more general weighted averaging scheme, where the updated estimate  $Q_t^k$  is:

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}: \quad Q_t^k(s, a) = \begin{cases} \sum_{k=1}^K \alpha_t^k(s, a) Q_{t-\frac{1}{2}}^k(s, a) & \text{if } t \equiv 0 \pmod{\tau} \\ Q_{t-\frac{1}{2}}^k(s, a) & \text{otherwise} \end{cases}, \quad (24)$$

where  $\alpha_t^k = [\alpha_t^k(s, a)]_{s \in \mathcal{S}, a \in \mathcal{A}} \in [0, 1]^{|S||A|}$  is an entry-wise weight assigned to agent  $k$  such that

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}: \quad \sum_{k=1}^K \alpha_t^k(s, a) = 1.$$

After a total of  $T$  iterations, for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , FedAsynQ outputs a global Q-estimate  $Q_T(s, a) = \sum_{k=1}^K \alpha_T^k(s, a) Q_T^k(s, a)$ . In the subsections below, we provide two possible ways (equal and importance weighting) to choose  $\alpha_t^k$  and their corresponding sample complexity analyses.

---

**Algorithm 2:** Federated Asynchronous Q-learning (FedAsynQ)
 

---

- 1: **inputs:** learning rate  $\{\eta\}$ , discount factor  $\gamma$ , number of agents  $K$ , synchronization period  $\tau$ , total number of iterations  $T$ .
  - 2: **initialization:**  $Q_0^k = Q_0$  for all  $k \in [K]$ .
  - 3: **for**  $t = 1, \dots, T$  **do**
  - 4:   **for**  $k \in [K]$  **do**
  - 5:     Draw action  $a_{t-1}^k \sim \pi_b^k(s_{t-1}^k)$ , observe reward  $r_{t-1}^k = r(s_{t-1}^k, a_{t-1}^k)$ , and draw next state  $s_t^k \sim P(\cdot | s_{t-1}^k, a_{t-1}^k)$ .
  - 6:     Compute  $Q_{t-\frac{1}{2}}^k$  according to (23).
  - 7:     Compute  $Q_t^k$  according to (24).
  - 8:   **end for**
  - 9: **end for**
  - 10: **return:**  $Q_T(s, a) = \sum_{k=1}^K \alpha_T^k(s, a) Q_T^k(s, a)$ , for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ .
- 

### 4.3 Performance guarantees with equal averaging

We begin with the most natural choice, which equally weights the local Q-estimates, that is,

$$\alpha_t^k(s, a) = \frac{1}{K}. \quad (25)$$

We call the resulting scheme FedAsynQ-EqAvg, which is also analyzed in Khodadadian et al. (2022). We have the following improved performance guarantee in the next theorem.

**Theorem 2 (Finite-time convergence of FedAsynQ-EqAvg)** *Consider any given  $\delta \in (0, 1)$  and  $\varepsilon \in (0, \frac{1}{1-\gamma}]$ . Suppose that the initialization of FedAsynQ-EqAvg satisfies  $0 \leq Q_0 \leq \frac{1}{1-\gamma}$ . There exist some sufficiently large constant  $c_T > 0$  and sufficiently small constant  $c_\eta > 0$ , both independent of problem parameters, such that with probability at least  $1 - \delta$ , the output of FedAsynQ-EqAvg satisfies  $\|Q_T - Q^*\|_\infty \leq \varepsilon$ , provided that the synchronization period  $\tau$ , the sample size per agent  $T$ , and the learning rate  $\eta$  satisfy*

$$\tau_0 \leq \tau \leq \frac{1}{4\eta} \min \left\{ \frac{1-\gamma}{4}, \frac{1}{K} \right\}, \quad (26a)$$

$$T \geq c_T \left( \frac{C_{\text{het}}}{K \mu_{\text{avg}} (1-\gamma)^5 \varepsilon^2} + T_0 \right) (\log((1-\gamma)^2 \varepsilon))^2 \log(TK) \log \frac{|\mathcal{S}| |\mathcal{A}| T^2 K}{\delta}, \quad (26b)$$

$$\eta = c_\eta \min \left\{ \frac{K(1-\gamma)^4 \varepsilon^2}{C_{\text{het}}}, \eta_0 \right\} \frac{1}{\log(TK) \log \frac{|\mathcal{S}| |\mathcal{A}| T^2 K}{\delta}}, \quad (26c)$$

where  $\tau_0 = \frac{2176 t_{\text{mix}}^{\max}}{\mu_{\text{avg}}} \log 8K \log \frac{4|\mathcal{S}| |\mathcal{A}| T^2}{\delta}$ ,  $T_0 = \frac{1}{\mu_{\text{avg}} (1-\gamma) \eta_0}$ , and  $\eta_0 = \frac{\mu_{\text{avg}} \min\{1-\gamma, K^{-1}\}}{t_{\text{mix}}^{\max}}$ , independent of  $\varepsilon$ .

Theorem 2 implies that to achieve an  $\varepsilon$ -accurate estimate (in the  $\ell_\infty$  sense), the sample complexity per agent of FedAsynQ-EqAvg is no more than

$$\tilde{O} \left( \frac{C_{\text{het}}}{K \mu_{\text{avg}} (1-\gamma)^5 \varepsilon^2} \right)$$

for sufficiently small  $\varepsilon$ , when the burn-in cost  $T_0$  — representing the impact of the mixing times — is amortized over time. A few implications are in order.

**Linear speedup without full coverage.** The sample complexity of FedAsynQ-EqAvg shows linear speedup with respect to the number of agents, which is especially pronounced when the local behavior policies are similar, i.e.,  $C_{\text{het}} \approx 1$ . Notably, the guarantee holds as long as all agents collectively cover the entire state-action space (i.e.,  $\mu_{\text{avg}} > 0$ ), unveiling the benefit of heterogeneity in local behavior policies. This is surprising in view of the convergence guarantee provided in Khodadadian et al. (2022), which requires each agent visits the entire state-action space (i.e.,  $\mu_{\min} = 0$ ). Moreover, our sample complexity has sharpened dependency on nearly all problem-dependent parameters compared to the bound  $\tilde{O}\left(\frac{|\mathcal{S}|^2}{K\mu_{\min}^5(1-\gamma)^9\varepsilon^2}\right)$  obtained in Khodadadian et al. (2022) by at least a factor of

$$\frac{\mu_{\text{avg}}|\mathcal{S}|^2}{C_{\text{het}}\mu_{\min}^5(1-\gamma)^4} \geq \frac{|\mathcal{S}|^5|\mathcal{A}|^3}{(1-\gamma)^4}.$$

For  $K = 1$ , the bound nearly matches with the sharpest upper bound  $\tilde{O}\left(\frac{1}{\mu_{\min}(1-\gamma)^4\varepsilon^2}\right)$  for the single-agent case (Li et al., 2023) up to a factor of  $1/(1-\gamma)$ , when ignoring the burn-in cost.

**Communication efficiency.** To provide further insights on the communication complexity of FedAsynQ-EqAvg, consider the regime when  $\varepsilon$  is sufficiently small and the number of agents is sufficiently large such that  $K \gtrsim \frac{1}{1-\gamma}$ . By plugging the choice of the learning rate (26c) into the upper bound of  $\tau$  in (26a), we can select the synchronization period as large as  $\tau \asymp \frac{C_{\text{het}}}{K^2(1-\gamma)^4\varepsilon^2}$  up to logarithmic factors, which ensures the communication complexity  $C_{\text{round}} = T/\tau$  is no more than  $\tilde{O}\left(\frac{K}{\mu_{\text{avg}}(1-\gamma)}\right)$ .

**Burn-in cost and mixing times.** In (26b), the second term  $T_0$  can be viewed as the burn-in cost needed for asynchronous Q-learning to eliminate the impact of Markovian noise and approximate synchronous Q-learning. The burn-in cost linearly scales with  $t_{\text{mix}}^{\max}$ , the largest mixing time of all the Markov chains of agents induced by their local behavior policies. Similarly, the convergence analysis in Khodadadian et al. (2022) also requires the sample size  $T$  to exceed a certain mixing time to guarantee the convergence of *all* Markov chains, similar to  $t_{\text{mix}}^{\max}$ . However, in Khodadadian et al. (2022), the effect of the mixing time on the convergence is not amortized over time, remaining non-negligible even for sufficiently small target accuracy  $\varepsilon$ . In contrast, in our sample complexity analysis, the burn-in time  $T_0$  is independent of the target accuracy  $\varepsilon$ , implying that the cost associated with the mixing times becomes negligible for sufficiently small  $\varepsilon$  and is amortized over time.

#### 4.4 Performance guarantees with importance averaging

In the asynchronous setting, heterogeneous behavior policies induce local trajectories that cover the state-action space in a non-uniform manner. As a result, agents may update the Q-estimate for a state-action pair at different frequencies, resulting in noisier Q-estimates of state-action pairs that an agent rarely visits. Equally-weighted averaging of such local Q-estimates is not efficient, because the convergence speed to the optimal Q-function for

each state-action pair is bottlenecked with the slowest converging agent that visits it least frequently. This is highlighted by the impact of the heterogeneity factor  $C_{\text{het}}$  in the sample complexity of FedAsynQ-EqAvg, which scales linearly with  $C_{\text{het}}$ , implying that increased heterogeneity among agents' trajectories may impede the convergence. For example, if only one agent exclusively visits a certain state-action pair  $(s, a)$  with probability one, while other agents never visit that particular state-action pair, the heterogeneity factor becomes  $C_{\text{het}} = K$  when  $K \leq 1/\mu_{\text{avg}}$ , canceling out the linear speedup.

Our key idea to prevent such inefficiency is to increase the contribution of frequently updated local Q-estimates, which are likely to have smaller errors. By assigning a weight inversely proportional to the error of the corresponding local estimate, we can balance the heterogeneous training progress of the local estimates and obtain an average estimate with much lower error. Combining this idea with the property that the local error decreases exponentially with the number of local visits, we propose an importance averaging scheme FedAsynQ-ImAvg with weights given by

$$\alpha_t^k(s, a) = \frac{(1 - \eta)^{-N_{t-\tau, t}^k(s, a)}}{\sum_{k'=1}^K (1 - \eta)^{-N_{t-\tau, t}^{k'}(s, a)}} \quad (27)$$

for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $k \in [K]$ , where  $N_{t-\tau, t}^k(s, a)$  represents the number of iterations between  $[t - \tau, t)$  when the agent  $k$  visits  $(s, a)$ . The weights in (27) can be calculated at the server based on the number of visits to each state-action pair by the agents in one synchronization period. Therefore, each agent needs to send its  $N_{t-\tau, t}^k(s, a)$  for each  $(s, a)$  along with its local Q-estimate, and FedAsynQ-ImAvg incurs twice the communication cost of FedAsynQ-EqAvg per iteration.

We have the following theorem on the finite-time convergence of FedAsynQ-ImAvg.

**Theorem 3 (Finite-time convergence of FedAsynQ-ImAvg)** *Consider any given  $\delta \in (0, 1)$  and  $\varepsilon \in (0, \frac{1}{1-\gamma}]$ . Suppose that the initialization of FedAsynQ-ImAvg satisfies  $0 \leq Q_0 \leq \frac{1}{1-\gamma}$ , and the synchronization period  $\tau$  obeys*

$$\tau \leq \frac{1}{4\eta} \min \left\{ \frac{1-\gamma}{4}, \frac{1}{K} \right\}. \quad (28a)$$

*There exist some sufficiently large constant  $c_T > 0$  and sufficiently small constant  $c_\eta > 0$ , both independent of problem parameters, such that with probability at least  $1 - \delta$ , the output of FedAsynQ-ImAvg satisfies  $\|Q_T - Q^*\|_\infty \leq \varepsilon$ , provided that the sample size per agent  $T$  and the learning rate  $\eta$  satisfy*

$$T \geq c_T \left( \frac{1}{K\mu_{\text{avg}}(1-\gamma)^5\varepsilon^2} + \tilde{T}_0 \right) (\log((1-\gamma)^2\varepsilon))^2 \log(TK) \log \frac{|S||\mathcal{A}|T^2K}{\delta}, \quad (28b)$$

$$\eta = c_\eta \min \{ K(1-\gamma)^4\varepsilon^2, \tilde{\eta}_0 \} \frac{1}{\log(TK) \log \frac{|S||\mathcal{A}|T^2K}{\delta}}, \quad (28c)$$

where  $\tilde{T}_0 = \frac{1}{\mu_{\text{avg}}(1-\gamma)\eta_0}$  and  $\tilde{\eta}_0 = \min \left\{ \frac{1}{t_{\text{mix}}^{\max}}, 1 - \gamma, K^{-1} \right\}$ , independent of  $\varepsilon$ .

Theorem 3 implies that to achieve an  $\varepsilon$ -accurate estimate (in the  $\ell_\infty$  sense), the sample complexity per agent of FedAsynQ-ImAvg is no more than

$$\tilde{O}\left(\frac{1}{K\mu_{\text{avg}}(1-\gamma)^5\varepsilon^2}\right)$$

for sufficiently small  $\varepsilon$ , when the burn-in cost  $\tilde{T}_0$  — representing the impact of the mixing times — is amortized over time. A few implications are in order.

**Linear speedup without the curse of heterogeneity.** The sample complexity of FedAsynQ-ImAvg is better than that of FedAsynQ-EqAvg, since it no longer depends on  $C_{\text{het}}$  which can be as large as  $1/\mu_{\text{avg}}$ . FedAsynQ-ImAvg not only overcomes potential insufficient local coverage by exploiting the complementary coverage of agents' behavior policies, but also achieves linear speedup with respect to the number of agents without suffering from the potential performance degradation due to the associated statistical heterogeneity as in FedAsynQ-EqAvg. In fact, the performance of FedAsynQ-ImAvg matches with centralized Q-learning as if we collect and process all data trajectories at the central server, up to the burn-in cost and logarithmic factors.

**Communication efficiency.** To provide further insights on the communication complexity of FedAsynQ-ImAvg, consider again the regime when  $\varepsilon$  is sufficiently small and  $K \gtrsim \frac{1}{1-\gamma}$ . To minimize the communication frequency while preserving the sample efficiency, we again plug the choice of the learning rate (28c) into (28a) and select the synchronization period as large as  $\tau \asymp \frac{1}{K^2(1-\gamma)^4\varepsilon^2}$  up to logarithmic factors. Then, this ensures the communication complexity  $C_{\text{round}} = T/\tau$  is no more than  $\tilde{O}\left(\frac{K}{\mu_{\text{avg}}(1-\gamma)}\right)$ .

## 5. Numerical experiments

In this section, we conduct numerical experiments to demonstrate the performance of the asynchronous Q-learning algorithms (FedAsynQ-EqAvg and FedAsynQ-ImAvg).

**Experimental setup.** Consider an MDP  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma)$  described in Figure 2, where  $\mathcal{S} = \{0, 1\}$  and  $\mathcal{A} = \{1, 2, \dots, m\}$ . The reward function  $r$  is set as  $r(s = 1, a) = 1$  and  $r(s = 0, a) = 0$  for any action  $a \in \mathcal{A}$ , and the discount factor is set as  $\gamma = 0.9$ . We now describe the transition kernel  $P$ . Here, we set the self-transitioning probabilities  $p_a := P(0|0, a)$  and  $q_a := P(1|1, a)$  uniformly at random from  $[0.4, 0.6]$  for each  $a \in \mathcal{A}$ , and set the probability of transitioning to the other state as  $P(1 - s|s, a) = 1 - P(s|s, a)$  for each  $s \in \mathcal{S}$ .

We evaluate the proposed federated asynchronous Q-learning algorithms on the above MDP with  $K$  agents selecting their behavior policies from  $\Pi = \{\pi_1, \pi_2, \dots, \pi_m\}$ , where the  $i$ -th policy always chooses action  $i$  for any state, i.e.,  $\pi_i(i|s) = 1$  for all  $s \in \mathcal{S}$ . Here, we assign  $\pi_i$  to agent  $k \in [K]$  if  $i \equiv k \pmod{m}$ . Note that if an agent has a behavior policy  $\pi_i$ , it can visit only two state-action pairs,  $(s = 0, a = i)$  and  $(s = 1, a = i)$ , as described in Figure 2. Thus, each agent covers a subset of the state-action space, and at least  $K = m$  agents are required to obtain local trajectories collectively covering the entire state-action space. Under this setting with  $m = 20$ , we run the algorithms for 100 simulations using



samples randomly generated from the MDP and policies assigned to the agents. The Q-function is initialized with entries uniformly at random from  $(0, \frac{1}{1-\gamma}]$  for each state-action pair.




Figure 2: An illustration of the constructed synthetic MDP  $\mathcal{M}$ . The red arrows represent transitioning paths when action  $a = i$  is taken in  $s = 0$  and  $s = 1$ . A trajectory induced by  $\pi_i$ , which executes only action  $i$  for any state, can cover only two state-action pairs,  $(s = 0, a = i)$  and  $(s = 1, a = i)$ .

**Faster convergence of FedAsynQ-ImAvg.** Figure 3 shows the normalized Q-estimate error  $(1 - \gamma)\|Q_T - Q^*\|_\infty$  with respect to the sample size  $T$ , with  $K = 20$  and  $\tau = 50$ . Given the trajectories of agents collectively cover the entire state-action space, the global Q-estimates of both FedAsynQ-EqAvg and FedAsynQ-ImAvg converge to the optimal Q-function, yet at different speeds. Although FedAsynQ-EqAvg converges in the end, we can see that it converges much slower compared to FedAsynQ-ImAvg, because each entry of the Q-function is trained by only one agent while the other  $m - 1$  agents never contribute useful information. However, the vacuous values of the  $m - 1$  agents significantly slow down the global convergence under equal averaging.

**Convergence speedup.** Figure 4 demonstrates the impact of the number of agents on the convergence speed of FedAsynQ-EqAvg and FedAsynQ-ImAvg. It can be observed that there is indeed a speedup in terms of the number of agents  $K$  with respect to the squared  $\ell_\infty$  error  $\|Q_T - Q^*\|_\infty^{-2}$ , which is poised to scale linearly with respect to the number of agents. In particular, the speedup is more rapid with FedAsynQ-ImAvg as  $K$  increases, while it increases much slower with FedAsynQ-EqAvg. This shows that FedAsynQ-ImAvg achieves much better convergence speedup in terms of the number of agents.

**Communication efficiency.** Figure 5 demonstrates the impact of the synchronization period  $\tau$  on the convergence of FedAsynQ-ImAvg and FedAsynQ-EqAvg. With frequent av-




Figure 3: The normalized  $\ell_\infty$  error of the Q-estimates  $(1 - \gamma)\|Q_T - Q^*\|_\infty$  with respect to the number of samples  $T$  for both FedAsynQ-EqAvg and FedAsynQ-ImAvg, with  $K = 20$  and  $\tau = 50$ . Here, the learning rates of FedAsynQ-ImAvg and FedAsynQ-EqAvg are set as  $\eta = 0.05$  and  $\eta = 0.2$ , where each algorithm converges to the same error floor at the fastest speed, respectively.

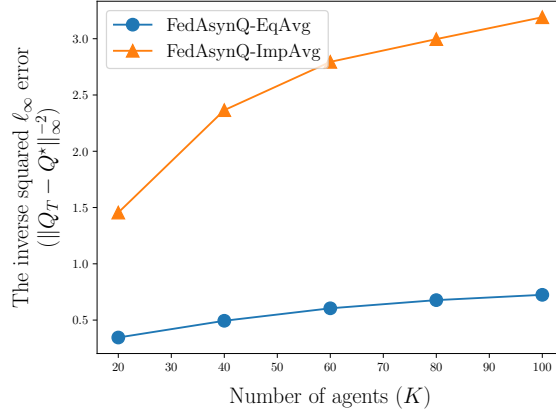


Figure 4: The inverse squared  $\ell_\infty$  error  $\|Q_T - Q^*\|_\infty^{-2}$  with respect to the number of agents  $K = 20, 40, 60, 80, 100$  for both FedAsynQ-EqAvg and FedAsynQ-ImAvg, with  $T = 300$  and  $\tau = 50$ .

eraging ( $\tau = 1$ ), FedAsynQ-ImAvg slightly outperforms FedAsynQ-EqAvg, but there is no significant difference because the heterogeneity between local Q-functions after just one local update is very small. The performance of FedAsynQ-EqAvg degrades as we increase  $\tau$  since FedAsynQ-EqAvg cannot cope with the increased heterogeneity between local Q-estimates as we increase the number of local steps. On the other end, the performance of FedAsynQ-ImAvg improves first (i.e.,  $\tau = 10, 25, 50$ ) as it balances the heterogeneity much better than FedAsynQ-EqAvg, but drops later if  $\tau$  is too large (i.e.,  $\tau = 75, 100$ ) due to the high variance of the averaged Q-estimates.




Figure 5: The normalized  $\ell_\infty$  error of the Q-estimates  $(1-\gamma)\|Q_T - Q^*\|_\infty$  with respect to the synchronization period  $\tau = 1, 10, 25, 50, 75, 100$  for both FedAsynQ-EqAvg and FedAsynQ-ImpAvg, with  $K = 20$  and  $T = 300$ .

## 6. Analysis outline

Let the matrix  $P \in \mathbb{R}^{|S| \times |\mathcal{A}| \times |\mathcal{A}|}$  represent the transition kernel of the underlying MDP, where  $P(s, a) = P(\cdot | s, a)$  is the probability vector corresponding to the state transition at the state-action pair  $(s, a)$ . For any vector  $V \in \mathbb{R}^{|S|}$ , we define the variance parameter  $\text{Var}_{s,a}(V)$  with respect to the probability vector  $P(s, a)$  as

$$\text{Var}_{s,a}(V) := \mathbb{E}_{s' \sim P(\cdot | s, a)} [V(s') - P(s, a)V]^2 = P(s, a)(V \circ V) - [P(s, a)V] \circ [P(s, a)V]. \quad (29)$$

Here,  $\circ$  denotes the Hadamard product such that  $a \circ b = [a_i b_i]_{i=1}^n$  for any vector  $a = [a_i]_{i=1}^n, b = [b_i]_{i=1}^n \in \mathbb{R}^n$ . With slight abuse of notation, we shall also assume  $V^* \in \mathbb{R}^{|S|}$ ,  $V_t^k \in \mathbb{R}^{|S|}$ ,  $Q^* \in \mathbb{R}^{|S| \times |\mathcal{A}|}$ ,  $Q_t^k \in \mathbb{R}^{|S| \times |\mathcal{A}|}$ ,  $Q_{t+\frac{1}{2}}^k \in \mathbb{R}^{|S| \times |\mathcal{A}|}$  and  $r \in \mathbb{R}^{|S| \times |\mathcal{A}|}$  represent the corresponding functions in the matrix/vector form.

### 6.1 Basic facts

We first state a few basic facts that hold both for the synchronous and the asynchronous settings. It is easy to establish, by induction, that all iterates satisfy for all  $1 \leq k \leq K$  and  $t \geq 0$  that

$$0 \leq Q_t^k \leq \frac{1}{1-\gamma}, \quad 0 \leq V_t^k \leq \frac{1}{1-\gamma}, \quad (30)$$

as long as  $0 \leq Q_0 = Q_0^k \leq \frac{1}{1-\gamma}$ ; see a similar argument, e.g., in Li et al. (2023, Lemma 4). In addition, observe that

$$\|V_t^k - V^*\|_\infty \leq \|Q_t^k - Q^*\|_\infty \quad (31)$$

since

$$\|V_t^k - V^*\|_\infty = \max_{s \in S} \left| \max_{a \in \mathcal{A}} Q_t^k(s, a) - \max_{a \in \mathcal{A}} Q^*(s, a) \right| \leq \max_{s \in S, a \in \mathcal{A}} |Q_t^k(s, a) - Q^*(s, a)| \leq \|Q_t^k - Q^*\|_\infty.$$

Letting  $Q_t$  be the average of the local Q-estimates at the end of the  $t$ -th iteration, i.e.,  $Q_t = \frac{1}{K} \sum_{k=1}^K Q_t^k$ , it follows from (13) and (24) that for all  $t \geq 0$  that

$$Q_t = \frac{1}{K} \sum_{k=1}^K Q_t^k = \frac{1}{K} \sum_{k=1}^K Q_{t-\frac{1}{2}}^k. \quad (32)$$

Denote the error between  $Q_t$  and  $Q^*$  by

$$\Delta_t = Q^* - Q_t,$$

which is the quantity we aim to control. From (30), it holds immediately that for all  $t \geq 0$ ,

$$\|\Delta_t\|_\infty \leq \frac{1}{1-\gamma}. \quad (33)$$

Next, we also introduce the following functions pertaining to periodic averaging. For any  $t$ ,

- define  $\iota(t) := \tau \lfloor \frac{t}{\tau} \rfloor$  as the most recent synchronization step until  $t$ ;
- define  $\phi(t) := \lfloor \frac{t}{\tau} \rfloor$  as the number of synchronization steps until  $t$ .

## 6.2 Proof outline of Theorem 1

Define the local empirical transition matrix at the  $t$ -th iteration  $P_t^k \in \{0, 1\}^{|S| \times |A| \times |S|}$  as

$$P_t^k((s, a), s') := \begin{cases} 1, & \text{if } s' = s_t^k(s, a) \\ 0, & \text{otherwise} \end{cases}, \quad (34)$$

then the local update rule (12) can be rewritten as

$$Q_{t-\frac{1}{2}}^k = (1-\eta)Q_{t-1}^k + \eta \left( r + \gamma P_t^k V_{t-1}^k \right). \quad (35)$$

The proof of Theorem 1 consists of the following steps.

**Step 1: error decomposition.** To analyze the error  $\Delta_t$ , we first decompose the error into three terms, each of which can be bounded in a simple form. From (32), it follows that

$$\begin{aligned} \Delta_t &= \frac{1}{K} \sum_{k=1}^K (Q^* - Q_{t-\frac{1}{2}}^k) \stackrel{(i)}{=} \frac{1}{K} \sum_{k=1}^K \left( (1-\eta)(Q^* - Q_{t-1}^k) + \eta(Q^* - r - \gamma P_t^k V_{t-1}^k) \right) \\ &\stackrel{(ii)}{=} (1-\eta)\Delta_{t-1} + \eta \frac{\gamma}{K} \sum_{k=1}^K (P V^* - P_t^k V_{t-1}^k) \\ &= (1-\eta)\Delta_{t-1} + \eta \frac{\gamma}{K} \sum_{k=1}^K (P - P_t^k) V_{t-1}^k + \eta \frac{\gamma}{K} \sum_{k=1}^K P(V^* - V_{t-1}^k), \end{aligned}$$

where (i) follows from (35), and (ii) follows from Bellman's optimality equation  $Q^* = r + \gamma PV^*$ . By recursion over the above relation, we obtain

$$\Delta_t = \underbrace{(1-\eta)^t \Delta_0}_{=:E_t^1} + \underbrace{\eta \frac{\gamma}{K} \sum_{i=1}^t (1-\eta)^{t-i} \sum_{k=1}^K (P - P_i^k) V_{i-1}^k}_{=:E_t^2} + \underbrace{\eta \frac{\gamma}{K} \sum_{i=1}^t (1-\eta)^{t-i} \sum_{k=1}^K P(V^* - V_{i-1}^k)}_{=:E_t^3}. \quad (36)$$

Here, the first term  $E_t^1$  denotes the initialization error stemming from the disparity between the initial Q-values and the optimal Q-values, which diminishes exponentially throughout iterations. The second term,  $E_t^2$ , comprises a weighted sum accounting for the difference between the true transition probability and the realized transition in each iteration, where the difference arises from the randomness of transitions. Lastly, the final term,  $E_t^3$ , represents a weighted sum of value estimation errors from preceding iterations, which introduces a recursive relation.

**Step 2: bounding the error terms.** Now, we obtain a bound of each of the error terms in (36) separately.

- **Bounding  $\|E_t^1\|_\infty$ .** Using the fact that all agents start with the same initial Q-values, i.e.,  $Q_0^k = Q_0$ , the first error term is bounded as follows:

$$\|E_t^1\|_\infty = (1-\eta)^t \|\Delta_0\|_\infty \leq \frac{(1-\eta)^t}{1-\gamma}, \quad (37)$$

where the last inequality follows from (33).

- **Bounding  $\|E_t^2\|_\infty$ .** Exploiting conditional independence across transitions in different iterations and applying Freedman's inequality (Freedman, 1975), the second error term is bounded using Lemma 4 below, whose proof is provided in Appendix B.1.

**Lemma 4** *For any given  $\delta \in (0, 1)$ , the following holds*

$$\|E_t^2\|_\infty \leq \frac{8\gamma}{1-\gamma} \sqrt{\frac{\eta}{K} \log \frac{|\mathcal{S}||\mathcal{A}|T}{\delta}} \quad (38)$$

for all  $0 \leq t \leq T$  with probability at least  $1-\delta$ , as long as  $\eta$  satisfies  $\eta \leq \frac{K}{2} (\log \frac{|\mathcal{S}||\mathcal{A}|T}{\delta})^{-1}$ .

- **Bounding  $\|E_t^3\|_\infty$ .** For  $E_t^3$ , we obtain the following recursive relation using Lemma 5 below, whose proof is provided in Appendix B.2.

**Lemma 5** *Let  $\beta$  be any integer that satisfies  $0 \leq \beta \leq \phi(T)$ . For any given  $\delta \in (0, 1)$ , the following holds*

$$\begin{aligned} \|E_t^3\|_\infty &\leq \frac{2\gamma}{1-\gamma} (1-\eta)^{\beta\tau} + \frac{16\gamma\eta\sqrt{\tau-1}}{(1-\gamma)} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|KT}{\delta}} \\ &\quad + \gamma(1+4\eta(\tau-1)) \max_{i(t)-\beta\tau \leq i < t} \|\Delta_i\|_\infty \end{aligned} \quad (39)$$

for all  $\beta\tau \leq t \leq T$  with probability at least  $1-\delta$ , as long as  $\eta$  satisfies  $\tau\eta < 1/2$ .

**Step 3: solving a recursive relation.** By putting all the bounds derived in the previous step together, for any  $\beta\tau \leq t \leq T$ , the total error bound can be written in a simple recursive form as follows:

$$\|\Delta_t\|_\infty \leq \zeta + \gamma(1 + 4\eta(\tau - 1)) \max_{\iota(t) - \beta\tau \leq i < t} \|\Delta_i\|_\infty \leq \zeta + \left(\frac{1 + \gamma}{2}\right) \max_{\iota(t) - \beta\tau \leq i < t} \|\Delta_i\|_\infty, \quad (40)$$

where in the first inequality we introduce the short-hand notation

$$\zeta := \frac{4(1 - \eta)^{\beta\tau}}{1 - \gamma} + \frac{8\gamma}{1 - \gamma} \sqrt{\frac{\eta}{K} \log \frac{|\mathcal{S}||\mathcal{A}|T}{\delta}} + \frac{16\gamma\eta\sqrt{\tau - 1}}{(1 - \gamma)} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|KT}{\delta}}, \quad (41)$$

and the second inequality follows from the assumption  $\tau - 1 \leq \frac{1 - \gamma}{8\gamma\eta}$ .

By invoking the recursive relation in (40)  $L$  times, where the choices of  $\beta$  and  $L$  will be made momentarily, it follows that for any  $L\beta\tau \leq t \leq T$ ,

$$\begin{aligned} \|\Delta_t\|_\infty &\leq \sum_{i=0}^{L-1} \left(\frac{1 + \gamma}{2}\right)^i \zeta + \left(\frac{1 + \gamma}{2}\right)^L \max_{\iota(t) - L\beta\tau \leq i < t} \|\Delta_i\|_\infty \\ &\leq \frac{2}{1 - \gamma} \zeta + \left(\frac{1 + \gamma}{2}\right)^L \left(\frac{1}{1 - \gamma}\right), \end{aligned} \quad (42)$$

where the second line uses the crude bound in (33).

Setting  $\beta = \left\lfloor \frac{1}{\tau} \sqrt{\frac{(1 - \gamma)T}{2\eta}} \right\rfloor$  and  $L = \left\lceil \sqrt{\frac{\eta T}{1 - \gamma}} \right\rceil$ , which ensures  $L\beta\tau \leq T$ , and plugging their choices into (41) and (42) at  $t = T$ , we obtain that

$$\begin{aligned} &\|\Delta_T\|_\infty \\ &\leq \frac{8(1 - \eta)^{\beta\tau}}{(1 - \gamma)^2} + \frac{16\gamma}{(1 - \gamma)^2} \sqrt{\frac{\eta}{K} \log \frac{|\mathcal{S}||\mathcal{A}|T}{\delta}} + \frac{32\gamma\eta\sqrt{\tau - 1}}{(1 - \gamma)^2} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|KT}{\delta}} + \left(\frac{1 + \gamma}{2}\right)^L \left(\frac{1}{1 - \gamma}\right) \\ &\leq \frac{32}{(1 - \gamma)^2} \left( \exp\left(-\frac{\sqrt{(1 - \gamma)\eta T}}{2}\right) + \gamma \sqrt{\frac{\eta}{K} \log \frac{|\mathcal{S}||\mathcal{A}|T}{\delta}} + \gamma\eta\sqrt{\tau - 1} \sqrt{\log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}} \right) \\ &\leq \frac{64}{(1 - \gamma)^2} \left( \exp\left(-\frac{\sqrt{(1 - \gamma)\eta T}}{2}\right) + \gamma \sqrt{\frac{\eta}{K} \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}} \right), \end{aligned} \quad (43)$$

where the second line follows from

$$\begin{aligned} (1 - \eta)^{\beta\tau} &\leq \exp(-\eta\beta\tau) \leq \exp\left(-\frac{\sqrt{(1 - \gamma)\eta T}}{2}\right), \\ \left(\frac{1 + \gamma}{2}\right)^L &= \left(1 - \frac{1 - \gamma}{2}\right)^L \leq \exp\left(-\frac{(1 - \gamma)L}{2}\right) \leq \exp\left(-\frac{\sqrt{(1 - \gamma)\eta T}}{2}\right), \end{aligned}$$

and the third line follows from the choice of the synchronization period such that

$$\tau - 1 \leq \frac{1}{\eta} \min \left\{ \frac{1 - \gamma}{8\gamma}, \frac{1}{K} \right\}. \quad (44)$$

Thus, for any given  $\varepsilon \in (0, \frac{1}{1-\gamma})$ , we can guarantee that  $\|\Delta_T\|_\infty \leq \varepsilon$  if

$$\begin{aligned} T &\geq c_T \frac{1}{K(1-\gamma)^5 \varepsilon^2} (\log((1-\gamma)^2 \varepsilon))^2 \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}, \\ \eta &= c_\eta K(1-\gamma)^4 \varepsilon^2 \frac{1}{\log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}} \end{aligned} \quad (45)$$

for some sufficiently large  $c_T$  and sufficiently small  $c_\eta$ .

### 6.3 Proof outline of Theorem 2

For simplicity, we introduce the following notation. Let  $\mathcal{U}_{v_1, v_2}^k(s, a)$  represent a set of iteration indices between  $[v_1, v_2)$  for some  $0 \leq v_1 \leq v_2 \leq T$  where agent  $k$  visits  $(s, a)$ , i.e.,

$$\mathcal{U}_{v_1, v_2}^k(s, a) := \left\{ u \in [v_1, v_2) : (s_u^k, a_u^k) = (s, a) \right\},$$

and  $N_{v_1, v_2}^k(s, a)$  denotes the number of visits of agent  $k$  on  $(s, a)$  during iterations between  $[v_1, v_2)$ , i.e.,

$$N_{v_1, v_2}^k(s, a) = |\mathcal{U}_{v_1, v_2}^k(s, a)|.$$

Define the local empirical transition matrix at the  $t$ -th iteration  $P_t^k \in \{0, 1\}^{|\mathcal{S}||\mathcal{A}| \times |\mathcal{S}|}$  as

$$P_t^k((s, a), s') := \begin{cases} 1 & \text{if } (s, a, s') = (s_{t-1}^k, a_{t-1}^k, s_t^k) \\ 0 & \text{otherwise} \end{cases}. \quad (46)$$

Then the local update rule (23) can be rewritten as

$$Q_{t-\frac{1}{2}}^k(s, a) = \begin{cases} (1-\eta)Q_{t-1}^k(s, a) + \eta(r_{t-1}^k + \gamma P_t^k(s, a)V_{t-1}^k) & \text{if } (s, a) = (s_{t-1}^k, a_{t-1}^k) \\ Q_{t-1}^k(s, a), & \text{otherwise} \end{cases}. \quad (47)$$

The proof of Theorem 2 consists of the following steps.

**Step 1: error decomposition.** Consider any  $0 \leq t \leq T$  such that  $t \equiv 0 \pmod{\tau}$ , i.e.,  $t$  is a synchronization step. To analyze  $\Delta_t$ , we first decompose the error for each  $(s, a) \in \mathcal{S} \times \mathcal{A}$  as follows:

$$\begin{aligned} \Delta_t(s, a) &= \frac{1}{K} \sum_{k=1}^K (Q^*(s, a) - Q_{t-\frac{1}{2}}^k(s, a)) \\ &= \left( \frac{1}{K} \sum_{k=1}^K (1-\eta)^{N_{t-\tau, t}^k(s, a)} \right) \Delta_{t-\tau}(s, a) \\ &\quad + \frac{\gamma}{K} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{t-\tau, t}^k(s, a)} \eta(1-\eta)^{N_{u+1, t}^k(s, a)} (P(s, a) - P_{u+1}^k(s, a)) V_u^k \\ &\quad + \frac{\gamma}{K} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{t-\tau, t}^k(s, a)} \eta(1-\eta)^{N_{u+1, t}^k(s, a)} P(s, a) (V^* - V_u^k), \end{aligned} \quad (48)$$

where we invoke the following recursive relation of the local error at iteration  $u$  such that  $(s_{u-1}, a_{u-1}) = (s, a)$ :

$$\begin{aligned}
 Q^\star(s, a) - Q_{u-\frac{1}{2}}^k(s, a) &= (1 - \eta)(Q^\star(s, a) - Q_{u-1}^k(s, a)) + \eta(Q^\star(s, a) - r_{u-1}^k - \gamma P_u^k(s, a) V_{u-1}^k) \\
 &= (1 - \eta)(Q^\star(s, a) - Q_{u-1}^k(s, a)) + \eta(\gamma P(s, a) V^\star - \gamma P_u^k(s, a) V_{u-1}^k) \\
 &= (1 - \eta)(Q^\star(s, a) - Q_{u-1}^k(s, a)) + \gamma \eta(P(s, a) - P_u^k(s, a)) V_{u-1}^k + \gamma P(s, a) (V^\star - V_{u-1}^k)
 \end{aligned} \tag{49}$$

Here, the second equality follows from Bellman's optimality equation. Denoting

$$\lambda_{v_1, v_2}(s, a) := \frac{1}{K} \sum_{k=1}^K (1 - \eta)^{N_{v_1, v_2}^k(s, a)} \tag{50}$$

for any integer  $0 \leq v_1 \leq v_2 \leq T$ , we apply recursion to the relation (48) over the synchronization periods, and obtain

$$\begin{aligned}
 \Delta_t(s, a) &= \left( \prod_{h=0}^{\phi(t)-1} \lambda_{h\tau, (h+1)\tau}(s, a) \right) \Delta_0(s, a) \\
 &\quad + \frac{\gamma}{K} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a) \right) \\
 &\quad \quad \times \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \eta(1 - \eta)^{N_{u+1, (h+1)\tau}^k(s, a)} (P(s, a) - P_{u+1}^k(s, a)) V_u^k \\
 &\quad + \frac{\gamma}{K} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a) \right) \\
 &\quad \quad \times \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \eta(1 - \eta)^{N_{u+1, (h+1)\tau}^k(s, a)} P(s, a) (V^\star - V_u^k) \\
 &= \underbrace{\omega_{0,t}(s, a) \Delta_0(s, a)}_{=: E_t^1(s, a)} + \underbrace{\gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} \omega_{u,t}^k(s, a) (P(s, a) - P_{u+1}^k(s, a)) V_u^k}_{=: E_t^2(s, a)} \\
 &\quad + \underbrace{\gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} \omega_{u,t}^k(s, a) P(s, a) (V^\star - V_u^k)}_{=: E_t^3(s, a)}, \tag{51}
 \end{aligned}$$

which is decomposed in a similar manner as (36). Here, we define

$$\omega_{0,t}(s, a) := \prod_{h=0}^{\phi(t)-1} \lambda_{h\tau, (h+1)\tau}(s, a), \tag{52a}$$



$$\omega_{u,t}^k(s, a) := \frac{1}{K} \eta (1 - \eta)^{N_{u+1, (\phi(u)+1)\tau}^k(s, a)} \prod_{l=\phi(u)+1}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a). \quad (52b)$$

We record the following useful lemma whose proof is provided in Appendix C.2.

**Lemma 6** *Consider integers  $v_1$  and  $v_2$  such that  $0 \leq v_1 \leq v_2 \leq t \leq T$ , where  $t \equiv 0 \pmod{\tau}$ , and a state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ . Suppose that  $\eta\tau \leq 1$ . The parameters defined in (52) satisfy*

$$\lambda_{v_1, v_2}(s, a) \leq \exp \left( -\frac{\eta}{2K} \sum_{k=1}^K N_{v_1, v_2}^k(s, a) \right), \quad (53a)$$

$$\omega_{0,t}(s, a) + \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} \omega_{u,t}^k(s, a) = 1, \quad (53b)$$

$$\sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, h'\tau}^k(s, a)} \omega_{u,t}^k(s, a) \leq \exp \left( -\frac{\eta}{2K} \sum_{k=1}^K N_{h'\tau, t}^k(s, a) \right), \quad \forall 0 \leq h' \leq \phi(t), \quad (53c)$$

$$\sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} (\omega_{u,t}^k(s, a))^2 \leq \frac{2\eta}{K}. \quad (53d)$$

**Step 2: bounding the error terms.** Here, we derive the bound of the error terms in (51) separately for all the state-action pairs  $(s, a) \in \mathcal{S} \times \mathcal{A}$ .

- **Bounding  $|E_t^1(s, a)|$ .** Using the initialization condition that  $Q_0(s, a) = Q_0^k(s, a)$  for every agent  $k \in [K]$ , we bound the first term for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  as follows:

$$|E_t^1(s, a)| \leq \omega_{0,t}(s, a) (\|Q_0\|_\infty + \|Q^*\|_\infty) \stackrel{(i)}{\leq} \frac{2\omega_{0,t}(s, a)}{1 - \gamma} \stackrel{(ii)}{\leq} \frac{2}{1 - \gamma} \exp \left( -\frac{\eta\mu_{\text{avg}}t}{8} \right), \quad (54)$$

where (i) holds because  $\|Q_0\|_\infty, \|Q^*\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)) and (ii) follows from the fact that

$$\omega_{0,t}(s, a) \leq \exp \left( -\frac{\eta}{2K} \sum_{k=1}^K N_{0,t}^k(s, a) \right) \leq \exp \left( -\frac{\eta\mu_{\text{avg}}t}{8} \right), \quad (55)$$

where the first inequality holds according to (53a) of Lemma 6, and the last inequality follows from the fact that  $\sum_{k=1}^K N_{0,t}^k(s, a) \geq \frac{K\mu_{\text{avg}}t}{4}$  for all  $(s, a, k, h) \in \mathcal{S} \times \mathcal{A} \times [K] \times [T]$  at least with probability  $1 - \delta$  according to Lemma 15 and the union bound, as long as  $t \geq t_{\text{th}}$ .

- **Bounding  $|E_t^2(s, a)|$ .** By carefully treating the statistical dependency via a decoupling argument and applying Freedman's inequality, we can obtain the following bound, whose proof is provided in Appendix C.3.

**Lemma 7** For any given  $\delta \in (0, 1)$ , the following holds for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $1 \leq t \leq T$ :

$$|E_t^2(s, a)| \leq \frac{7241\gamma}{(1-\gamma)} \sqrt{\frac{C_{\text{het}}\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \quad (56)$$

with probability at least  $1 - 4\delta$ , as long as  $\tau \geq t_{\text{th}}$  and

$$\frac{3}{T} \leq \eta \leq \min \left\{ \frac{1}{16\tau}, \frac{1}{4\tau K}, \frac{1}{128KC_{\text{het}} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \right\}.$$

- **Bounding  $|E_t^3(s, a)|$ .** For  $E_t^3$ , we can obtain the following recursive relation, whose proof is provided in Appendix C.4.

**Lemma 8** Let  $\beta$  be any integer that satisfies  $0 < \beta \leq \phi(T)$ . For any given  $\delta \in (0, 1)$ , the following holds

$$\begin{aligned} |E_t^3(s, a)| &\leq \frac{2\gamma}{1-\gamma} \exp\left(-\frac{\eta\mu_{\text{avg}}\beta\tau}{8}\right) + \frac{8\gamma\eta\sqrt{\tau-1}}{1-\gamma} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}} \\ &\quad + \frac{1+\gamma}{2} \max_{\phi(t)-\beta \leq h \leq \phi(t)-1} \|\Delta_{h\tau}\|_{\infty}, \end{aligned} \quad (57)$$

for all  $\beta\tau \leq t \leq T$  with probability at least  $1 - \delta$ , as long as  $\beta\tau \geq t_{\text{th}}$  and  $\eta \leq \min\{\frac{1-\gamma}{4\gamma\tau}, \frac{1}{2\tau}\}$ .

**Step 3: solving a recursive relation.** By putting all the bounds derived in the previous step together, for any  $\beta\tau \leq t \leq T$ , the total error bound can be written in a simple recursive form as follows:

$$\|\Delta_t\|_{\infty} \leq \theta + \frac{1+\gamma}{2} \max_{\phi(t)-\beta \leq h \leq \phi(t)-1} \|\Delta_{h\tau}\|_{\infty}, \quad (58)$$

where we define

$$\begin{aligned} \theta &:= \frac{4}{1-\gamma} \exp\left(-\frac{\eta\mu_{\text{avg}}\beta\tau}{8}\right) + \frac{7241\gamma}{(1-\gamma)} \sqrt{\frac{C_{\text{het}}\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \\ &\quad + \frac{8\gamma\eta\sqrt{\tau-1}}{1-\gamma} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}}. \end{aligned} \quad (59)$$

Then, by invoking the recursive relation for  $L_1$  times, where the choices of  $\beta$  and  $L_1$  will be made momentarily, it follows that for any  $L_1\beta\tau \leq t \leq T$ ,

$$\|\Delta_t\|_{\infty} \leq \sum_{l=0}^{L_1-1} \left(\frac{1+\gamma}{2}\right)^l \theta + \left(\frac{1+\gamma}{2}\right)^{L_1} \max_{\phi(t)-\beta L_1 \leq i \leq \phi(t)-1} \|\Delta_{i\tau}\|_{\infty} \leq \frac{2}{1-\gamma} \left(\theta + \left(\frac{1+\gamma}{2}\right)^{L_1}\right), \quad (60)$$

where the last inequality follows from (33).

Setting  $\beta = \left\lfloor \frac{1}{\tau} \sqrt{\frac{2(1-\gamma)T}{\mu_{\text{avg}}\eta}} \right\rfloor$  and  $L_1 = \left\lceil \frac{1}{2} \sqrt{\frac{\mu_{\text{avg}}\eta T}{(1-\gamma)}} \right\rceil$ , which ensures  $L_1\beta\tau \leq T$ , and plugging the choices into (59) and (60) at  $t = T$ , we obtain

$$\begin{aligned}
 \|\Delta_T\|_\infty &\leq \frac{8 \exp\left(-\frac{\eta\mu_{\text{avg}}\beta\tau}{8}\right)}{(1-\gamma)^2} + \frac{14481\gamma}{(1-\gamma)^2} \sqrt{\frac{C_{\text{het}}\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \\
 &\quad + \frac{16\gamma\eta\sqrt{\tau-1}}{(1-\gamma)^2} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}} + \frac{2}{1-\gamma} \left(\frac{1+\gamma}{2}\right)^L \\
 &\leq \frac{16}{(1-\gamma)^2} \exp\left(-\frac{\sqrt{(1-\gamma)\mu_{\text{avg}}\eta T}}{8}\right) + \frac{14481\gamma}{(1-\gamma)^2} \sqrt{\frac{C_{\text{het}}\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \\
 &\quad + \frac{16\gamma\eta\sqrt{\tau-1}}{(1-\gamma)^2} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}} \\
 &\leq \frac{14497}{(1-\gamma)^2} \left( \exp\left(-\frac{\sqrt{(1-\gamma)\mu_{\text{avg}}\eta T}}{8}\right) + \gamma \sqrt{\frac{C_{\text{het}}\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \right), \tag{61}
 \end{aligned}$$

where the second line follows from

$$\begin{aligned}
 \exp\left(-\frac{\eta\mu_{\text{avg}}\beta\tau}{8}\right) &\leq \exp\left(-\frac{\sqrt{(1-\gamma)\mu_{\text{avg}}\eta T}}{8}\right), \\
 \left(\frac{1+\gamma}{2}\right)^{L_1} &= \left(1 - \frac{1-\gamma}{2}\right)^{L_1} \leq \exp\left(-\frac{1-\gamma}{2}L_1\right) \leq \exp\left(-\frac{\sqrt{(1-\gamma)\mu_{\text{avg}}\eta T}}{4}\right),
 \end{aligned}$$

and the third line follows from the choice of the synchronization period such that

$$t_{\text{th}} \leq \tau \leq \frac{1}{4\eta} \min\left\{\frac{1-\gamma}{4}, \frac{1}{K}\right\}. \tag{62}$$

Thus, for any given  $\varepsilon \in (0, \frac{1}{1-\gamma}]$ , we can guarantee that  $\|\Delta_T\|_\infty \leq \varepsilon$  if

$$\begin{aligned}
 T &\geq c_T (\log((1-\gamma)^2\varepsilon))^2 \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta} \\
 &\quad \cdot \frac{1}{\mu_{\text{avg}}} \max\left\{\frac{C_{\text{het}}}{K(1-\gamma)^5\varepsilon^2}, \frac{t_{\text{mix}}^{\max}}{\mu_{\text{avg}}(1-\gamma) \min\{1-\gamma, K^{-1}\}}\right\}, \\
 \eta &= c_\eta \left(\log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}\right)^{-1} \min\left\{\frac{K(1-\gamma)^4\varepsilon^2}{C_{\text{het}}}, \frac{\mu_{\text{avg}} \min\{1-\gamma, K^{-1}\}}{t_{\text{mix}}^{\max}}\right\}
 \end{aligned}$$

for some sufficiently large  $c_T$  and sufficiently small  $c_\eta$ .

#### 6.4 Proof outline of Theorem 3

The proof of Theorem 3 consists of the following steps.

**Step 1: error decomposition.** Consider any  $0 \leq t \leq T$  such that  $t \equiv 0 \pmod{\tau}$ , i.e.,  $t$  is a synchronization step. To analyze  $\Delta_t$ , invoking the recursive relation of the local error (cf. (49)), we first decompose the error for each  $(s, a) \in \mathcal{S} \times \mathcal{A}$  as follows:

$$\begin{aligned}
 \Delta_t(s, a) &= \sum_{k=1}^K \alpha_t^k(s, a) (Q^*(s, a) - Q_{t-\frac{1}{2}}^k(s, a)) \\
 &= \left( \sum_{k=1}^K \alpha_t^k(s, a) (1 - \eta)^{N_{t-\tau, t}^k(s, a)} \right) \Delta_{t-\tau}(s, a) \\
 &\quad + \gamma \sum_{k=1}^K \alpha_t^k(s, a) \sum_{u \in \mathcal{U}_{t-\tau, t}^k(s, a)} \eta (1 - \eta)^{N_{u+1, t}^k(s, a)} (P(s, a) - P_{u+1}^k(s, a)) V_u^k \\
 &\quad + \gamma \sum_{k=1}^K \alpha_t^k(s, a) \sum_{u \in \mathcal{U}_{t-\tau, t}^k(s, a)} \eta (1 - \eta)^{N_{u+1, t}^k(s, a)} P(s, a) (V^* - V_u^k) \\
 &= \left( \frac{K}{\sum_{k'=1}^K (1 - \eta)^{-N_{t-\tau, t}^{k'}(s, a)}} \right) \Delta_{t-\tau}(s, a) \\
 &\quad + \gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{t-\tau, t}^k(s, a)} \frac{\eta (1 - \eta)^{-N_{t-\tau, u+1}^k(s, a)}}{\sum_{k'=1}^K (1 - \eta)^{-N_{t-\tau, t}^{k'}(s, a)}} (P(s, a) - P_{u+1}^k(s, a)) V_u^k \\
 &\quad + \gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{t-\tau, t}^k(s, a)} \frac{\eta (1 - \eta)^{-N_{t-\tau, u+1}^k(s, a)}}{\sum_{k'=1}^K (1 - \eta)^{-N_{t-\tau, t}^{k'}(s, a)}} P(s, a) (V^* - V_u^k), \tag{63}
 \end{aligned}$$

where the last line uses the definition of  $\alpha_t^k(s, a)$  in (27). Denoting

$$\tilde{\lambda}_{v_1, v_2}(s, a) := \frac{K}{\sum_{k=1}^K (1 - \eta)^{N_{v_1, v_2}^k(s, a)}} \tag{64}$$

for any integer  $0 \leq v_1 \leq v_2 \leq T$ , we apply recursion to the relation (63) over the synchronization period, and obtain

$$\begin{aligned}
 \Delta_t(s, a) &= \left( \prod_{h=0}^{\phi(t)-1} \tilde{\lambda}_{h\tau, (h+1)\tau}(s, a) \right) \Delta_0(s, a) \\
 &\quad + \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right) \\
 &\quad \quad \times \gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \frac{\eta (1 - \eta)^{-N_{h\tau, u+1}^k(s, a)}}{\sum_{k'=1}^K (1 - \eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}} (P(s, a) - P_{u+1}^k(s, a)) V_u^k \\
 &\quad + \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right)
 \end{aligned}$$

$$\begin{aligned}
 & \gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \frac{\eta(1-\eta)^{-N_{h\tau, u+1}^k(s, a)}}{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}} P(s, a) (V^* - V_u^k) \\
 &= \underbrace{\tilde{\omega}_{0,t}(s, a) \Delta_0(s, a)}_{=: E_t^1(s, a)} + \underbrace{\gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) (P(s, a) - P_{u+1}^k(s, a)) V_u^k}_{=: E_t^2(s, a)} \\
 &+ \underbrace{\gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) P(s, a) (V^* - V_u^k)}_{=: E_t^3(s, a)}, \tag{65}
 \end{aligned}$$

which is again decomposed similarly as (36). Here, we define

$$\tilde{\omega}_{0,t}(s, a) := \prod_{h=0}^{\phi(t)-1} \tilde{\lambda}_{h\tau, (h+1)\tau}(s, a), \tag{66a}$$

$$\tilde{\omega}_{u,t}^k(s, a) := \frac{\eta(1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)}}{\sum_{k'=1}^K (1-\eta)^{-N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}} \left( \prod_{l=\phi(u)+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right). \tag{66b}$$

We record the following useful lemma whose proof is provided in Appendix C.5.

**Lemma 9** *Consider any integers  $0 \leq v_1 \leq v_2 \leq t \leq T$  where  $t \equiv 0 \pmod{\tau}$  and any state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ . Suppose that  $\eta\tau \leq 1$ , then the parameters defined in (66) satisfy*

$$\frac{1}{3K} \leq \alpha_t^k(s, a) \leq \frac{3}{K}, \tag{67a}$$

$$\tilde{\omega}_{0,t}(s, a) \leq (1-\eta)^{\frac{1}{K} \sum_{k=1}^K N_{0,t}^k(s, a)}, \tag{67b}$$

$$\tilde{\omega}_{0,t}(s, a) + \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) = 1, \tag{67c}$$

$$\sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, h'\tau}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) \leq (1-\eta)^{\frac{1}{K} \sum_{k=1}^K N_{h'\tau, t}^k(s, a)}, \quad \forall 0 \leq h' \leq \phi(t), \tag{67d}$$

$$\sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} (\tilde{\omega}_{u,t}^k(s, a))^2 \leq \frac{6\eta}{K}. \tag{67e}$$

**Step 2: bounding the error terms.** Here, we derive the bound of each error term in (65) separately for all the state-action pairs  $(s, a) \in \mathcal{S} \times \mathcal{A}$ .

- **Bounding  $|E_t^1(s, a)|$ .** Using the initialization condition that  $Q_0(s, a) = Q_0^k(s, a)$  for every client  $k \in [K]$ , we bound the first term for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  as follows:

$$|E_t^1(s, a)| \leq \tilde{\omega}_{0,t}(\|Q_0\|_\infty + \|Q^*\|_\infty) \stackrel{(i)}{\leq} \frac{2\tilde{\omega}_{0,t}}{1-\gamma} \stackrel{(ii)}{\leq} \frac{2}{1-\gamma} (1-\eta)^{\frac{1}{K} \sum_{k=1}^K N_{0,t}^k(s, a)}$$

$$\stackrel{\text{(iii)}}{\leq} \frac{2}{1-\gamma} (1-\eta)^{\frac{1}{4}\mu_{\text{avg}}t}, \quad (68)$$

where (i) holds because  $\|Q_0\|_\infty, \|Q^*\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)), (ii) follows from (67b) of Lemma 9, and (iii) holds for all  $(s, a, t) \in \mathcal{S} \times \mathcal{A} \times [T]$  with probability at least  $1 - \delta$  according to Lemma 15, as long as  $t \geq t_{\text{th}}$ .

- **Bounding  $|E_t^2(s, a)|$ .** By carefully treating the statistical dependency via a decoupling argument and applying Freedman's inequality, we can obtain the following bound, whose proof is provided in Appendix C.6.

**Lemma 10** *For any given  $\delta \in (0, 1)$ , the following holds for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $1 \leq t \leq T$ :*

$$|E_t^2(s, a)| \leq \frac{2064\gamma}{(1-\gamma)} \sqrt{\frac{\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \quad (69)$$

with probability at least  $1 - 2\delta$ , as long as

$$\frac{3}{T} < \eta \leq \min \left\{ \frac{1}{16\tau}, \frac{K}{256 \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}}, \frac{1}{34816t_{\text{mix}}^{\max} \log(8K) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}} \right\}.$$

- **Bounding  $|E_t^3(s, a)|$ .** For  $E_t^3$ , similarly to Lemma 8, we can obtain the following recursive relation, whose proof is provided in Appendix C.7.

**Lemma 11** *Let  $\beta$  be any integer that satisfies  $\frac{t_{\text{th}}}{\tau} \leq \beta \leq \phi(T)$ . For any given  $\delta \in (0, 1)$ , the following holds*

$$\begin{aligned} |E_t^3(s, a)| &\leq \frac{2(1-\eta)^{\frac{\mu_{\text{avg}}\beta\tau}{4}}}{1-\gamma} + \frac{8\gamma\eta\sqrt{\tau-1}}{1-\gamma} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}} \\ &\quad + \frac{1+\gamma}{2} \max_{\phi(t)-\beta \leq h \leq \phi(t)-1} \|\Delta_{h\tau}\|_\infty, \end{aligned} \quad (70)$$

for all  $\beta\tau \leq t \leq T$  with probability at least  $1 - \delta$ , as long as  $\eta \leq \min\{\frac{1-\gamma}{4\gamma\tau}, \frac{1}{2\tau}\}$ .

**Step 3: solving a recursive relation.** By putting all the bounds derived in the previous step together, for any  $\beta\tau \leq t \leq T$ , the total error bound can be written in a simple recursive form as follows:

$$\|\Delta_t\|_\infty \leq \theta + \frac{1+\gamma}{2} \max_{\phi(t)-\beta \leq h \leq \phi(t)-1} \|\Delta_{h\tau}\|_\infty, \quad (71)$$

where we define

$$\begin{aligned} \tilde{\theta} &:= \frac{4}{1-\gamma} (1-\eta)^{\frac{\mu_{\text{avg}}\beta\tau}{4}} + \frac{2064\gamma}{(1-\gamma)} \sqrt{\frac{\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \\ &\quad + \frac{8\gamma\eta\sqrt{\tau-1}}{1-\gamma} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}}. \end{aligned} \quad (72)$$

Then, by invoking the recursive relation for  $L_2$  times, where the choices of  $\beta$  and  $L_2$  will be made momentarily, it follows that for any  $L_2\beta\tau \leq t \leq T$ ,

$$\|\Delta_t\|_\infty \leq \sum_{l=0}^{L_2-1} \left(\frac{1+\gamma}{2}\right)^l \tilde{\theta} + \left(\frac{1+\gamma}{2}\right)^{L_2} \max_{\phi(t)-\beta L \leq i \leq \phi(t)-1} \|\Delta_{i\tau}\|_\infty \leq \frac{2}{1-\gamma} \left(\theta + \left(\frac{1+\gamma}{2}\right)^{L_2}\right), \quad (73)$$

where the last inequality follows from (33).

Setting  $L_2 = \left\lceil \frac{1}{2} \sqrt{\frac{\mu_{\text{avg}} \eta T}{(1-\gamma)}} \right\rceil$  and  $\beta = \left\lfloor \frac{1}{\tau} \sqrt{\frac{2(1-\gamma)T}{\mu_{\text{avg}} \eta}} \right\rfloor$ , which ensures  $L_2\beta\tau \leq T$ , and plugging the choices into (72) and (73) at  $t = T$ , we obtain

$$\begin{aligned} \|\Delta_T\|_\infty &\leq \frac{8(1-\eta)^{\frac{\mu_{\text{avg}}\beta\tau}{4}}}{(1-\gamma)^2} + \frac{4128\gamma}{(1-\gamma)^2} \sqrt{\frac{\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \\ &\quad + \frac{16\gamma\eta\sqrt{\tau-1}}{(1-\gamma)^2} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}} + \frac{2}{1-\gamma} \left(\frac{1+\gamma}{2}\right)^{L_2} \\ &\leq \frac{16}{(1-\gamma)^2} \exp\left(-\frac{\sqrt{(1-\gamma)\mu_{\text{avg}}\eta T}}{4}\right) + \frac{4128\gamma}{(1-\gamma)^2} \sqrt{\frac{\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \\ &\quad + \frac{16\gamma\eta\sqrt{\tau-1}}{(1-\gamma)^2} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}} \\ &\leq \frac{4144}{(1-\gamma)^2} \left( \exp\left(-\frac{\sqrt{(1-\gamma)\mu_{\text{avg}}\eta T}}{4}\right) + \gamma \sqrt{\frac{\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \right), \end{aligned} \quad (74)$$

where the second line follows from

$$\begin{aligned} (1-\eta)^{\frac{\mu_{\text{avg}}\beta\tau}{4}} &\leq \exp\left(-\frac{\eta\mu_{\text{avg}}\beta\tau}{4}\right) \leq \exp\left(-\frac{\sqrt{(1-\gamma)\mu_{\text{avg}}\eta T}}{4}\right), \\ \left(\frac{1+\gamma}{2}\right)^{L_2} &= \left(1 - \frac{1-\gamma}{2}\right)^{L_2} \leq \exp\left(-\frac{1-\gamma}{2} L_2\right) \leq \exp\left(-\frac{\sqrt{(1-\gamma)\mu_{\text{avg}}\eta T}}{4}\right), \end{aligned}$$

and the third line follows from the choice of the synchronization period such that

$$\tau \leq \frac{1}{4\eta} \min\left\{\frac{1-\gamma}{4}, \frac{1}{K}\right\}. \quad (75)$$

Thus, for any given  $\varepsilon \in (0, \frac{1}{1-\gamma})$ , optimizing  $\eta$  and  $T$  to make (74) bounded by  $\varepsilon$  and recalling  $\beta\tau \geq t_{\text{th}}$ , we can guarantee that  $\|\Delta_T\|_\infty \leq \varepsilon$  if

$$\begin{aligned} T &\geq c_T (\log((1-\gamma)^2 \varepsilon))^2 \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta} \\ &\quad \cdot \frac{1}{\mu_{\text{avg}}} \max\left\{\frac{1}{K(1-\gamma)^5 \varepsilon^2}, \frac{t_{\text{mix}}^{\max}}{(1-\gamma)}, \frac{1}{(1-\gamma) \min\{1-\gamma, K^{-1}\}}\right\}, \end{aligned}$$

$$\begin{aligned}\eta &= c_\eta \min \left\{ K(1-\gamma)^4 \varepsilon^2 \frac{1}{\log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}}, \frac{1}{\mu_{\text{avg}} t_{\text{th}}}, \frac{1}{t_{\text{mix}}^{\max} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \right\} \\ &= c_\eta \left( \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta} \right)^{-1} \min \left\{ K(1-\gamma)^4 \varepsilon^2, \frac{1}{t_{\text{mix}}^{\max}}, \min \{1-\gamma, K^{-1}\} \right\}\end{aligned}$$

for some sufficiently large  $c_T$  and sufficiently small  $c_\eta$ .

## 7. Discussions

We presented a sample complexity analysis of federated Q-learning in both synchronous and asynchronous settings. Our sample complexity not only leads to linear speedup with respect to the number of agents, but also significantly improves the dependencies on other salient problem parameters over the prior art. For federated asynchronous Q-learning, we proposed a novel importance averaging scheme that weighs the agents' local Q-estimates according to the number of visits to each state-action pair. This allows agents to leverage the blessing of heterogeneity of their local behavior policies and collaboratively learn the optimal Q-function that otherwise would not be possible, without requiring each individual agent to cover the entire state-action space. Looking ahead, this work opens up many exciting future directions, some outlined below.

- *Improved sample complexity.* While our sample complexity bounds are near-optimal with respect to the size of the state-action space, it is still sub-optimal with respect to the effective horizon length as well as the mixing time when benchmarking with the sample complexity in the single-agent setting (Li et al., 2023). It will be interesting to close this gap, and further improve the sample complexity with variance reduction techniques (Wainwright, 2019b; Li et al., 2021b) in the federated setting.
- *Understanding communication asynchrony across agents.* As a starting point, our work assumes that all agents communicate with the server in a synchronous manner to perform periodic averaging. However, in practical federated networks, some agents might be stragglers due to communication slowdowns, which warrants further investigation (Kairouz et al., 2021).
- *Other RL settings and function approximation.* Besides the infinite-horizon tabular MDPs, it will be of great interest to extend our analysis framework to other RL settings including but not limited to the finite-horizon setting, the average reward setting, heterogeneous environments across the agents (Yang et al., 2024), as well as incorporating function approximation.
- *Federated offline RL.* In many applications, offline RL is attracting a growing amount of interest, which aims to explore history datasets to improve the learned policy without exploration, e.g. via pessimistic variants of Q-learning (Shi et al., 2022). It will be appealing to develop federated offline Q-learning algorithms to enable learning from geographically distributed history datasets.
- *Adaptive communication periods and communication efficiency.* In our work, we assume a constant communication period  $\tau$ , where the number of local steps between



synchronizations is fixed. This results in the number of communication rounds growing linearly with the number of agents  $K$  and the horizon factor  $\frac{1}{1-\gamma}$ , which is shy from the statistical-communication complexity established in Salgia and Chi (2024). A recent study Woo et al. (2024) shows that a federated offline RL algorithm can significantly improve communication efficiency without sacrificing sample efficiency by adaptively increasing the communication periods exponentially in a finite-horizon MDP. In this case, the required number of communication rounds scales only with the horizon length, independent of the number of agents. It would be interesting to explore whether adaptive communication periods could lead to similar improvements in communication efficiency in the infinite-horizon MDP setting.

## Acknowledgments

This work is supported in part by the grants NSF CCF-2007911, CCF-2106778, CNS-2148212, and ONR N00014-19-1-2404 to Y. Chi, NSF-CCF 2007834, CCF-2045694, CNS-2112471, and ONR N00014-23-1-2149 to G. Joshi, and the Hsu Chang Memorial Fellowship at Carnegie Mellon University to J. Woo.

## Appendix A. Preliminaries

We record a few useful inequalities that will be used throughout our analysis. To start with, our analysis leverages Freedman's inequality (Freedman, 1975), which we record a user-friendly version as follows.

**Theorem 12 (Theorem 6 in Li et al. (2023))** *Suppose that  $Y_n = \sum_{k=1}^n X_k \in \mathbb{R}$ , where  $\{X_k\}$  is a real-valued scalar sequence obeying*

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

Define

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

where we write  $\mathbb{E}_{k-1}$  for the expectation conditional on  $\{X_j\}_{j:j < k}$ . Then for any given  $\sigma^2 \geq 0$ , one has

$$\mathbb{P} \{ |Y_n| \geq \tau \text{ and } W_n \leq \sigma^2 \} \leq 2 \exp \left( -\frac{\tau^2/2}{\sigma^2 + R\tau/3} \right). \quad (76)$$

In addition, suppose that  $W_n \leq \sigma^2$  holds deterministically. For any positive integer  $m \geq 1$ , with probability at least  $1 - \delta$  one has

$$|Y_n| \leq \sqrt{8 \max \left\{ W_n, \frac{\sigma^2}{2m} \right\} \log \frac{2m}{\delta}} + \frac{4}{3} R \log \frac{2m}{\delta}. \quad (77)$$

Another useful relation concerns the concentration of empirical distributions of uniformly ergodic Markov chains, which is rephrased from Li et al. (2021b).

**Lemma 13 (Li et al. (2021b, Lemma 8))** *Consider any time homogeneous and uniformly ergodic Markov chain  $(X_0, X_1, X_2, \dots)$  with transition kernel  $P$ , finite state space  $\mathcal{X}$ , and stationary distribution  $\mu$ . Let  $t_{\text{mix}}$  be the mixing time of the Markov chain and  $\mu_{\min}$  be the minimum entry of the stationary distribution  $\mu$ . Consider any  $0 < \delta < 1$ . For any  $x \in \mathcal{X}$ , if  $t \geq \frac{443t_{\text{mix}}}{\nu} \log \frac{4|\mathcal{X}|}{\delta}$  for  $\nu \geq \mu(x)$ , then*

$$\forall y \in \mathcal{X} : \quad \mathbb{P}_{X_1=y} \left\{ \left| \sum_{i=1}^t \mathbb{1}\{X_i = x\} - t\mu(x) \right| \geq \frac{1}{2} t\nu \right\} \leq \frac{\delta}{|\mathcal{X}|}.$$

**Remark 14** *Lemma 13 is a slightly generalized version of in Li et al. (2021b, Lemma 8), where the concentration bound is characterized in terms of any given threshold  $\nu \geq \mu(x)$ , not scaling with the stationary distribution  $\mu(x)$ . It can be shown using the Bernstein's inequality for Markov chains (Paulin, 2015, Theorem 3.11) in the same manner as Li et al. (2021b, Lemma 8), except that the threshold is set to  $\frac{\nu t}{2}$  instead of  $\frac{\mu(x)t}{2}$ . We omit further details for conciseness and refer interested readers to the proof in Li et al. (2021b).*

In addition, we provide the concentration bound of the total number of visits of multiple agents with independent uniformly ergodic Markov chains, whose proof is provided in Appendix C.1. Denote

$$t_{\text{th}}(s, a) := \frac{2176t_{\text{mix}}^{\max} \log 8K \log \frac{4|S||\mathcal{A}|T^2}{\delta}}{\mu_{\text{avg}}(s, a)} \quad \text{and} \quad t_{\text{th}} := \frac{2176t_{\text{mix}}^{\max} \log 8K \log \frac{4|S||\mathcal{A}|T^2}{\delta}}{\mu_{\text{avg}}}. \quad (78)$$

Here,  $\mu_{\text{avg}}(s, a) := \frac{1}{K} \sum_{k=1}^K \mu_{\text{b}}^k(s, a)$  is the average behavior policy over all agents.

**Lemma 15** *Consider any  $\delta \in (0, 1)$ . Under the asynchronous sampling, for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $0 \leq u < v \leq T$  such that  $v - u \geq t_{\text{th}}(s, a)$ , the following holds :*

$$\frac{1}{4}(v - u)K\mu_{\text{avg}}(s, a) \leq \sum_{k=1}^K N_{u,v}^k(s, a) \leq 2(v - u)K\mu_{\text{avg}}(s, a) \quad (79)$$

with probability at least  $1 - \frac{\delta}{|S||\mathcal{A}|T^2}$ .

## Appendix B. Proofs for federated synchronous Q-learning (Section 3)

Define the following actions

$$a^*(s) = \arg \max_{a \in \mathcal{A}} Q^*(s, a), \quad a_i^k(s) = \arg \max_{a \in \mathcal{A}} Q_i^k(s, a), \quad a_i(s) = \arg \max_{a \in \mathcal{A}} \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a) \quad (80)$$

for any state  $s \in \mathcal{S}$ , which will be useful throughout the proof.

### B.1 Proof of Lemma 4

For notation simplicity, let  $z_i^k(s, a) := \eta(1 - \eta)^{t-i}(P(s, a) - P_i^k(s, a))V_{i-1}^k$ , then the entries of  $E_t^2 = [E_t^2(s, a)]$  can be written as

$$E_t^2(s, a) = \eta \frac{\gamma}{K} \sum_{i=1}^t (1 - \eta)^{t-i} \sum_{k=1}^K (P(s, a) - P_i^k(s, a))V_{i-1}^k = \frac{\gamma}{K} \sum_{i=1}^t \sum_{k=1}^K z_i^k(s, a), \quad (81)$$

which we plan to bound by invoking Freedman's inequality (cf. Theorem 12) using the fact  $z_i^k(s, a)$  is independent of the transition events of other agents  $k' \neq k$  at  $i$  and has zero mean conditioned on the events before iteration  $i$ , i.e.,

$$\mathbb{E}[z_i^k(s, a) | V_{i-1}^K, \dots, V_{i-1}^1, \dots, V_0^K, \dots, V_0^1] = 0, \quad \forall k \in [K], 1 \leq i \leq t. \quad (82)$$

Before applying Freedman's inequality, we first derive the following properties of the variable  $z_i^k(s, a)$ .

- First, we can bound

$$B_t(s, a) := \max_{k \in [K], 1 \leq i \leq t} |z_i^k(s, a)|$$

$$\leq \max_{k \in [K], 1 \leq i \leq t} \eta (\|P(s, a)\|_1 + \|P_i^k(s, a)\|_1) \|V_{i-1}^k\|_\infty \leq \frac{2\eta}{1-\gamma}, \quad (83)$$

where the first inequality uses  $(1-\eta)^{t-i} \leq 1$ , and the last inequality follows from  $\|P(s, a)\|_1 \leq 1$ ,  $\|P_i^k(s, a)\|_1 \leq 1$ , and  $\|V_{i-1}^k\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)).

• Next, we have

$$\begin{aligned} W_t(s, a) &:= \sum_{i=1}^t \sum_{k=1}^K \mathbb{E}[(z_i^k(s, a))^2 | V_{i-1}^K, \dots, V_{i-1}^1, \dots, V_0^K, \dots, V_0^1] \\ &= \sum_{i=1}^t \sum_{k=1}^K \text{Var}(z_i^k(s, a) | V_{i-1}^K, \dots, V_{i-1}^1, \dots, V_0^K, \dots, V_0^1) \\ &= \sum_{i=1}^t \sum_{k=1}^K \eta^2 (1-\eta)^{2(t-i)} \text{Var}_{s,a}(V_{i-1}^k) \\ &\leq \frac{2K}{(1-\gamma)^2} \sum_{i=1}^t \eta^2 (1-\eta)^{2(t-i)} \leq \frac{2\eta K}{(1-\gamma)^2} := \sigma^2, \end{aligned} \quad (84)$$

where we recall the definition of  $\text{Var}_{s,a}$  in (29). Here, the first inequality holds since

$$\text{Var}_{s,a}(V_{i-1}^k) \leq \|P(s, a)\|_1 (\|V_{i-1}^k\|_\infty)^2 + (\|P(s, a)\|_1 \|V_{i-1}^k\|_\infty)^2 \leq \frac{2}{(1-\gamma)^2}$$

and the last inequality follows from

$$\sum_{i=1}^t \eta^2 (1-\eta)^{2(t-i)} \leq \frac{\eta^2 (1 - (1-\eta)^{2t})}{1 - (1-\eta)^2} \leq \eta. \quad (85)$$

By substituting the above bounds (cf. (83) and (84)) and  $m = 1$  into Freedman's inequality (see Theorem 12), it follows that for any  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}$  and  $t \in [T]$ ,

$$\begin{aligned} \left| \sum_{i=1}^t \sum_{k=1}^K z_i^k(s, a) \right| &\leq \sqrt{8 \max\{W_t(s, a), \frac{\sigma^2}{2m}\} \log \frac{2m|\mathcal{S}||\mathcal{A}|T}{\delta}} + \frac{4}{3} B_t(s, a) \log \frac{2m|\mathcal{S}||\mathcal{A}|T}{\delta} \\ &\leq \sqrt{\frac{32\eta K}{(1-\gamma)^2} \log \frac{|\mathcal{S}||\mathcal{A}|T}{\delta}} + \frac{6\eta}{1-\gamma} \log \frac{|\mathcal{S}||\mathcal{A}|T}{\delta} \\ &\leq \frac{8}{1-\gamma} \sqrt{\eta K \log \frac{|\mathcal{S}||\mathcal{A}|T}{\delta}} \end{aligned} \quad (86)$$

with probability at least  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T}$ , where the last inequality holds under the assumption  $\eta \leq \frac{K}{2} (\log \frac{|\mathcal{S}||\mathcal{A}|T}{\delta})^{-1}$ . Applying the union bound over all  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}$  and  $t \in [T]$  then completes the proof.

## B.2 Proof of Lemma 5

For any  $\beta\tau \leq t \leq T$  and  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , we can decompose the entries of  $E_t^3 = [E_t^3(s, a)]$  as

$$\begin{aligned}
 |E_t^3(s, a)| &= \left| \frac{\eta\gamma}{K} \sum_{i=0}^{t-1} \sum_{k=1}^K (1-\eta)^{t-i-1} P(s, a)(V^\star - V_i^k) \right| \\
 &\leq \underbrace{\left| \frac{\eta\gamma}{K} \sum_{i=0}^{\iota(t)-\beta\tau-1} \sum_{k=1}^K (1-\eta)^{t-i-1} P(s, a)(V^\star - V_i^k) \right|}_{=: E_t^{3a}(s, a)} + \underbrace{\left| \frac{\eta\gamma}{K} \sum_{i=\iota(t)-\beta\tau}^{t-1} \sum_{k=1}^K (1-\eta)^{t-i-1} P(s, a)(V^\star - V_i^k) \right|}_{=: E_t^{3b}(s, a)}.
 \end{aligned} \tag{87}$$

We shall bound these two terms separately.

**Step 1: bounding  $E_t^{3a}(s, a)$ .** First, the bound of  $E_t^{3a}$  is obtained as follows:

$$\begin{aligned}
 E_t^{3a}(s, a) &\leq \eta \frac{\gamma}{K} \sum_{k=1}^K \sum_{i=0}^{\iota(t)-\beta\tau-1} (1-\eta)^{t-i} \|P(s, a)\|_1 (\|V^\star\|_\infty + \|V_i^k\|_\infty) \\
 &\leq \frac{2\eta\gamma}{1-\gamma} \sum_{i=0}^{\iota(t)-\beta\tau-1} (1-\eta)^{t-i-1} \leq \frac{2\gamma}{1-\gamma} (1-\eta)^{\beta\tau},
 \end{aligned} \tag{88}$$

where the second inequality holds due to the fact that  $\|P(s, a)\|_1 \leq 1$  and  $\|V^\star\|_\infty \leq \frac{1}{1-\gamma}$ ,  $\|V_i^k\|_\infty \leq \frac{1}{1-\gamma}$ , and the last inequality follows from

$$\sum_{i=0}^{\iota(t)-\beta\tau-1} (1-\eta)^{t-i-1} \leq (1-\eta)^{\beta\tau} + (1-\eta)^{\beta\tau+1} + \dots + (1-\eta)^{t-1} \leq \frac{(1-\eta)^{\beta\tau}}{1-(1-\eta)} \leq \frac{(1-\eta)^{\beta\tau}}{\eta}.$$

**Step 2: decomposing the bound on  $E_t^{3b}(s, a)$ .** Next,  $E_t^{3b}(s, a)$  can be bounded as follows

$$\begin{aligned}
 E_t^{3b}(s, a) &= \left| \frac{\eta\gamma}{K} \sum_{i=\iota(t)-\beta\tau}^{t-1} \sum_{k=1}^K (1-\eta)^{t-i-1} P(s, a)(V^\star - V_i^k) \right| \\
 &\leq \gamma \sum_{i=\iota(t)-\beta\tau}^{t-1} \eta (1-\eta)^{t-i-1} \left| \frac{1}{K} \sum_{k=1}^K P(s, a)(V^\star - V_i^k) \right| \\
 &\leq \gamma \sum_{i=\iota(t)-\beta\tau}^{t-1} \eta (1-\eta)^{t-i-1} \left\| \frac{1}{K} \sum_{k=1}^K (V^\star - V_i^k) \right\|_\infty,
 \end{aligned} \tag{89}$$

where the second inequality holds since  $\|P(s, a)\|_1 \leq 1$ . To continue, denoting

$$d_{v,w}^k(s, a) := Q_w^k(s, a) - Q_v^k(s, a), \tag{90}$$

we claim the following bound for any  $0 \leq i < T$ , which will be shown in Appendix B.2.1:

$$\left\| \frac{1}{K} \sum_{k=1}^K (V^\star - V_i^k) \right\|_\infty \leq \|\Delta_i\|_\infty + 2 \max_k \|d_{\iota(i),i}^k\|_\infty. \quad (91)$$

In view of (91), it boils down to control  $\max_k \|d_{\iota(i),i}^k\|_\infty$ . For any  $(s, a) \in \mathcal{S} \times \mathcal{A}$ ,  $k \in [K]$ , and  $0 \leq i < T$ , by the definition (90), it follows that

$$|d_{\iota(i),i}^k(s, a)| = \left| \sum_{j=\iota(i)}^{i-1} d_{j,j+1}^k(s, a) \right| \leq \underbrace{2\eta \sum_{j=\iota(i)}^{i-1} \|\Delta_j^k\|_\infty}_{:=B_1} + \underbrace{\gamma\eta \left| \sum_{j=\iota(i)}^{i-1} (P_{j+1}^k(s, a) - P(s, a))V^\star \right|}_{:=B_2}, \quad (92)$$

where

$$\Delta_j^k = Q^\star - Q_j^k. \quad (93)$$

The inequality (92) holds by the local update rule:

$$\begin{aligned} d_{j,j+1}^k(s, a) &= Q_{j+1}^k(s, a) - Q_j^k(s, a) \\ &= \eta(r(s, a) + \gamma P_{j+1}^k(s, a)V_j^k - Q_j^k(s, a)) \\ &\stackrel{(i)}{=} \eta(r(s, a) + \gamma P_{j+1}^k(s, a)V_j^k - r(s, a) - \gamma P(s, a)V^\star + Q^\star(s, a) - Q_j^k(s, a)) \\ &= \eta(\gamma P_{j+1}^k(s, a)V_j^k - \gamma P(s, a)V^\star + Q^\star(s, a) - Q_j^k(s, a)) \\ &= \gamma\eta P_{j+1}^k(s, a)(V_j^k - V^\star) + \gamma\eta(P_{j+1}^k(s, a) - P(s, a))V^\star + \eta\Delta_j^k(s, a) \\ &\leq 2\eta\|\Delta_j^k\|_\infty + \gamma\eta(P_{j+1}^k(s, a) - P(s, a))V^\star, \end{aligned} \quad (94)$$

where (i) follows from Bellman's optimality equation, and the last inequality follows from  $\|P_{j+1}^k(s, a)\|_1 \leq 1$  and  $\|V_j^k - V^\star\|_\infty \leq \|\Delta_j^k\|_\infty$  (cf. (31)).

Next, we bound each term in (92) separately.

- **Bounding  $B_1$ .** The local error  $\|\Delta_j^k\|_\infty$  is bounded as stated in the following lemma, whose proof is provided in Appendix B.2.2.

**Lemma 16** *Assume  $\tau\eta \leq \frac{1}{2}$ . For any given  $\delta \in (0, 1)$ , the following bound holds for any  $1 \leq i \leq T$  and  $k \in [K]$ :*

$$\|\Delta_i^k\|_\infty \leq \|\Delta_{\iota(i)}\|_\infty + \frac{2}{1-\gamma} \sqrt{\eta \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}} \quad (95)$$

with at least probability  $1 - \delta$ , where  $\iota(i)$  is the most recent synchronization step until  $i$ .

Using the fact that  $i - \iota(i) \leq \tau - 1$ , we can claim that

$$2\eta \sum_{j=\iota(i)}^{i-1} \|\Delta_j^k\|_\infty \leq 2\eta(\tau - 1)\|\Delta_{\iota(i)}\|_\infty + \frac{4\eta(\tau - 1)}{1-\gamma} \sqrt{\eta \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}}. \quad (96)$$

- **Bounding  $B_2$ .** Using the fact that the empirical transitions are independent and centered on the true transition probability, by invoking Hoeffding's inequality and the union bound, we can claim that the following holds for all  $(s, a, k, t) \in \mathcal{S} \times \mathcal{A} \times [K] \times [T]$ ,

$$\begin{aligned} \gamma\eta \left| \sum_{j=\iota(i)}^{i-1} (P_{j+1}^k(s, a) - P(s, a))V^\star \right| &\leq \frac{\gamma\eta}{1-\gamma} \sqrt{\frac{1}{2} \sum_{j=\iota(i)}^{i-1} \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}} \\ &\leq \frac{\gamma\eta}{1-\gamma} \sqrt{(\tau-1) \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}} \end{aligned} \quad (97)$$

with probability at least  $1 - \delta$  for any given  $\delta \in (0, 1)$ , where  $\tau$  is the synchronization period.

By substituting the bound of  $B_1$  and  $B_2$  into (92), and applying the union bound, we obtain that: for any given  $\delta \in (0, 1)$ , the following holds for any  $0 \leq i \leq T$  and  $k \in [K]$ :

$$\begin{aligned} \|d_{\iota(i), i}^k\|_\infty &\leq 2\eta(\tau-1)\|\Delta_{\iota(i)}\|_\infty + \frac{4\eta((\tau-1)\sqrt{\eta} + \sqrt{\tau-1})}{(1-\gamma)} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|KT}{\delta}} \\ &\leq 2\eta(\tau-1)\|\Delta_{\iota(i)}\|_\infty + \frac{8\eta\sqrt{\tau-1}}{(1-\gamma)} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|KT}{\delta}} \end{aligned} \quad (98)$$

with at least probability  $1 - \delta$ , where  $\iota(i)$  is the most recent synchronization step until  $i$ . Here, the second line uses the fact  $\eta\tau < 1$ .

By combining (98) and (91) and substituting it into (89) and using the fact that  $\sum_{i=\iota(t)-\beta\tau}^{t-1} \eta(1-\eta)^{t-i-1} \leq 1$ , we can obtain the bound  $E_t^{3b}(s, a)$  as follows:

$$\begin{aligned} |E_t^{3b}(s, a)| &\leq \frac{16\gamma\eta\sqrt{\tau-1}}{(1-\gamma)} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|KT}{\delta}} \\ &\quad + \gamma \sum_{i=\iota(t)-\beta\tau}^{t-1} \eta(1-\eta)^{t-i-1} (\|\Delta_i\|_\infty + 4\eta(\tau-1)\|\Delta_{\iota(i)}\|_\infty) \\ &\leq \frac{16\gamma\eta\sqrt{\tau-1}}{(1-\gamma)} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|KT}{\delta}} + \gamma(1 + 4\eta(\tau-1)) \max_{\iota(t)-\beta\tau \leq i < t} \|\Delta_i\|_\infty. \end{aligned} \quad (99)$$

**Step 3: putting all together.** Now, we have the bounds of  $E_t^{3a}$  and  $E_t^{3b}$  separately derived above. By combining the bounds in (87), we can finally claim the advertised bound and this completes the proof.

### B.2.1 PROOF OF (91)

On one end, it follows that for any  $s \in \mathcal{S}$ ,

$$\begin{aligned} \frac{1}{K} \sum_{k=1}^K (V^\star(s) - V_i^k(s)) &= Q^\star(s, a^\star(s)) - \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_i^k(s)) \\ &\leq Q^\star(s, a^\star(s)) - \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a^\star(s)) = \Delta_i(s, a^\star(s)), \end{aligned} \quad (100)$$

where we use the definitions in (80). On the other end, it follows that

$$\begin{aligned}
 & \frac{1}{K} \sum_{k=1}^K \left( V^\star(s) - V_i^k(s) \right) \\
 &= Q^\star(s, a^\star(s)) - \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_{\iota(i)}(s)) + \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_{\iota(i)}(s)) - \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_i^k(s)) \\
 &\geq Q^\star(s, a_{\iota(i)}(s)) - \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_{\iota(i)}(s)) + \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_{\iota(i)}(s)) - \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_i^k(s)) \\
 &= \Delta_i(s, a_{\iota(i)}(s)) + \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_{\iota(i)}(s)) - \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_i^k(s)), \tag{101}
 \end{aligned}$$

where the inequality follows from the fact that  $a^\star(s)$  is the optimal action for state  $s$ . Notice that the latter terms can be further lower bounded as

$$\begin{aligned}
 & \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_{\iota(i)}(s)) - \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_i^k(s)) \\
 &= \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_{\iota(i)}(s)) - \frac{1}{K} \sum_{k=1}^K Q_{\iota(i)}^k(s, a_{\iota(i)}(s)) + \frac{1}{K} \sum_{k=1}^K Q_{\iota(i)}^k(s, a_{\iota(i)}(s)) \\
 &\quad - \frac{1}{K} \sum_{k=1}^K Q_{\iota(i)}^k(s, a_i^k(s)) + \frac{1}{K} \sum_{k=1}^K Q_{\iota(i)}^k(s, a_i^k(s)) - \frac{1}{K} \sum_{k=1}^K Q_i^k(s, a_i^k(s)) \\
 &\geq \frac{1}{K} \sum_{k=1}^K (d_{\iota(i),i}^k(s, a_{\iota(i)}(s)) - d_{\iota(i),i}^k(s, a_i^k(s))), \tag{102}
 \end{aligned}$$

where the inequality follows from the definition (90) and the fact that

$$Q_{\iota(i)}^k(s, a_{\iota(i)}(s)) - Q_{\iota(i)}^k(s, a_i^k(s)) \geq 0.$$

The above holds, since  $Q_{\iota(i)}^k = Q_{\iota(i)}$  for all  $k \in [K]$  agents after periodic averaging at  $\iota(i)$ , and  $a_{\iota(i)}(s)$  is the optimal action at state  $s$  at time  $\iota(i)$  for every agent.

Combining (100), (101) and (102), we obtain

$$\begin{aligned}
 \Delta_i(s, a_{\iota(i)}(s)) + \frac{1}{K} \sum_{k=1}^K (d_{\iota(i),i}^k(s, a_{\iota(i)}(s)) - d_{\iota(i),i}^k(s, a_i^k(s))) &\leq \frac{1}{K} \sum_{k=1}^K \left( V^\star(s) - V_i^k(s) \right) \\
 &\leq \Delta_i(s, a^\star(s)),
 \end{aligned}$$

which immediately implies (91).



## B.2.2 PROOF OF LEMMA 16

By applying the decomposition in (36) to the local error for agent  $k$ , we decompose  $\Delta_i^k$  as follows:

$$\begin{aligned} \Delta_i^k(s, a) = & \underbrace{(1 - \eta)^{i-\iota(i)} \Delta_{\iota(i)}^k(s, a)}_{:=D_1} + \underbrace{\gamma \sum_{j=\iota(i)+1}^i \eta(1 - \eta)^{i-j} (P(s, a) - P_j^k(s, a)) V^*}_{:=D_2} \\ & + \underbrace{\gamma \sum_{j=\iota(i)+1}^i \eta(1 - \eta)^{i-j} P_j^k(s, a) (V^* - V_{j-1}^k)}_{:=D_3}. \end{aligned} \quad (103)$$

We shall bound each term separately.

- **Bounding  $D_1$ .** Since  $\Delta_{\iota(i)}^k = \Delta_{\iota(i)}$  for every agent  $k$  at the synchronization step  $\iota(i)$ ,

$$|D_1| \leq (1 - \eta)^{i-\iota(i)} \|\Delta_{\iota(i)}\|_\infty. \quad (104)$$

- **Bounding  $D_2$ .** In a similar manner to (97), by invoking Hoeffding inequality and using the fact that  $\sum_{j=\iota(i)+1}^i (\eta(1 - \eta)^{i-j})^2 \leq \eta$  (cf. (85)), we can claim that the following holds for all  $(s, a, k, t) \in \mathcal{S} \times \mathcal{A} \times [K] \times [T]$ ,

$$|D_2| \leq \gamma \sqrt{\sum_{j=\iota(i)+1}^i (\eta(1 - \eta)^{i-j})^2 \|V^*\|_\infty^2 \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}} \leq \frac{\gamma}{1 - \gamma} \sqrt{\eta \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}} \quad (105)$$

with probability at least  $1 - \delta$  for any given  $\delta \in (0, 1)$ .

- **Bounding  $D_3$ .** By bounding  $\|V^* - V_{j-1}^k\|_\infty$  with the local error  $\|\Delta_{j-1}^k\|_\infty$  (cf. (31)) and using  $\|P_j^k(s, a)\|_1 \leq 1$ , we have

$$|D_3| \leq \gamma \sum_{j=\iota(i)+1}^i \eta(1 - \eta)^{i-j} \|P_j^k(s, a)\|_1 \|V^* - V_{j-1}^k\|_\infty \leq \gamma \sum_{j=\iota(i)+1}^i \eta(1 - \eta)^{i-j} \|\Delta_{j-1}^k\|_\infty. \quad (106)$$

By combining the bounds obtained above in (103), we obtain the following recursive relation

$$\|\Delta_i^k\|_\infty \leq (1 - \eta)^{i-\iota(i)} \|\Delta_{\iota(i)}\|_\infty + \underbrace{\frac{\gamma}{1 - \gamma} \sqrt{\eta \log \frac{|\mathcal{S}||\mathcal{A}|KT}{\delta}}}_{:=\rho} + \gamma \sum_{j=\iota(i)+1}^i \eta(1 - \eta)^{i-j} \|\Delta_{j-1}^k\|_\infty. \quad (107)$$

By invoking the recursive relation with some algebraic calculations, we obtain the following bound

$$\begin{aligned}
 & \|\Delta_i^k\|_\infty \\
 & \leq (1-\eta)^{i-\iota(i)} \|\Delta_{\iota(i)}\|_\infty + \rho \\
 & \quad + \gamma \sum_{j_1=\iota(i)+1}^i \eta(1-\eta)^{i-j_1} \left( (1-\eta)^{j_1-1-\iota(i)} \|\Delta_{\iota(i)}\|_\infty + \rho + \gamma \sum_{j_2=\iota(i)+1}^{j_1-1} \eta(1-\eta)^{j_1-1-j_2} \|\Delta_{j_2-1}^k\|_\infty \right) \\
 & = \left( (1-\eta)^{i-\iota(i)} + \gamma \sum_{j_1=\iota(i)+1}^i \eta(1-\eta)^{i-1-\iota(i)} \right) \|\Delta_{\iota(i)}\|_\infty + \left( 1 + \gamma \sum_{j_1=\iota(i)+1}^i \eta(1-\eta)^{i-j_1} \right) \rho \\
 & \quad + \gamma^2 \sum_{j_1=\iota(i)+1}^i \sum_{j_2=\iota(i)+1}^{j_1-1} \eta^2(1-\eta)^{i-1-j_2} \|\Delta_{j_2-1}^k\|_\infty \\
 & \leq \left( (1-\eta)^{i-\iota(i)} + \gamma \sum_{j_1=\iota(i)+1}^i \eta(1-\eta)^{i-1-\iota(i)} \right) \|\Delta_{\iota(i)}\|_\infty + \left( 1 + \gamma \sum_{j_1=\iota(i)+1}^i \eta(1-\eta)^{i-j_1} \right) \rho \\
 & \quad + \gamma^2 \sum_{j_1=\iota(i)+1}^i \sum_{j_2=\iota(i)+1}^{j_1-1} \eta^2(1-\eta)^{i-1-j_2} \left( (1-\eta)^{j_2-1-\iota(i)} \|\Delta_{\iota(i)}\|_\infty + \rho + \dots \right) \\
 & \leq \|\Delta_{\iota(i)}\|_\infty \\
 & \quad \times \left( (1-\eta)^{i-\iota(i)} + \gamma \sum_{j_1=\iota(i)+1}^i \eta(1-\eta)^{i-1-\iota(i)} + \dots + \gamma^l \sum_{j_1=\iota(i)+1}^i \dots \sum_{j_l=\iota(i)+1}^{j_{l-1}-1} \eta^l(1-\eta)^{i-l-\iota(i)} \right) \\
 & \quad + \left( 1 + \gamma \sum_{j_1=\iota(i)+1}^i \eta(1-\eta)^{i-j_1} + \dots + \gamma^l \sum_{j_1=\iota(i)+1}^i \dots \sum_{j_l=\iota(i)+1}^{j_{l-1}-1} \eta^l(1-\eta)^{i-l+1-j_l} \right) \rho \\
 & \quad + \gamma^{l+1} \sum_{j_1=\iota(i)+1}^i \dots \sum_{j_{l+1}=\iota(i)+1}^{j_l-1} \eta^{l+1}(1-\eta)^{i-l-j_{l+1}} \left( \|\Delta_{j_{l+1}-1}^k\| \right) \\
 & \stackrel{(i)}{\leq} \sum_{l=0}^{i-\iota(i)} \gamma^l \binom{i-\iota(i)}{l} \eta^l (1-\eta)^{i-\iota(i)-l} \|\Delta_{\iota(i)}^k\|_\infty + \sum_{l=0}^{i-\iota(i)-1} \gamma^l \binom{i-\iota(i)}{l} \eta^l \rho \\
 & \leq ((1-\eta) + \gamma\eta)^{i-\iota(i)} \|\Delta_{\iota(i)}^k\|_\infty + (1+\gamma\eta)^{i-\iota(i)} \rho \\
 & \stackrel{(ii)}{\leq} \|\Delta_{\iota(i)}^k\|_\infty + 2\rho, \tag{108}
 \end{aligned}$$

where (i) follows from  $\Delta_{j_{i-\iota(i)}-1}^k = \Delta_{\iota(i)}^k$  since  $j_l \leq i-l+1$ ,

$$\sum_{j_1=\iota(i)+1}^i \sum_{j_2=\iota(i)+1}^{j_1-1} \dots \sum_{j_l=\iota(i)+1}^{j_{l-1}-1} \eta^l (1-\eta)^{i-l-\iota(i)} = \binom{i-\iota(i)}{l} \eta^l (1-\eta)^{i-l-\iota(i)},$$

$$\sum_{j_1=\iota(i)+1}^i \cdots \sum_{j_l=\iota(i)+1}^{j_{l-1}-1} \eta^l (1-\eta)^{i-l+1-j_l} \leq \sum_{j_1=\iota(i)+1}^i \cdots \sum_{j_l=\iota(i)+1}^{j_{l-1}-1} \eta^l \leq \binom{i-\iota(i)}{l} \eta^l,$$

and (ii) follows from  $(1 + \gamma\eta)^{i-\iota(i)} \leq (1 + \gamma\eta)^\tau \leq e^{\tau\eta} \leq 2$  since  $i - \iota(i) \leq \tau$  and  $\tau\eta \leq \frac{1}{2}$ . This completes the proof.

## Appendix C. Proofs for federated asynchronous Q-learning (Section 4)

### C.1 Proof of Lemma 15

To describe the joint probabilistic transitions of  $K$  agents formally, we first introduce the following Markov chain  $X_t = (X_t^1, \dots, X_t^K)$ ,  $t = 0, 1, \dots$ , where  $X_t^k \in \mathcal{S} \times \mathcal{A}$  is the state-action pair visited by agent  $k$  at time  $t$ . The joint transition kernel  $P$  of  $K$  agents is given by

$$P := \begin{pmatrix} P^1 & & & \\ & P^2 & & \\ & & \ddots & \\ & & & P^K \end{pmatrix}, \quad (109)$$

where  $P^k$  is the transition kernel of agent  $k$ ,  $k = 1, \dots, K$ . Since the agents are independent, the stationary distribution of the joint Markov chain is  $\mu$ , given by

$$\mu(x) := \prod_{k=1}^K \mu_b^k(x^k), \quad \forall x = (x^1, x^2, \dots, x^K) \in (\mathcal{S} \times \mathcal{A})^K, \quad (110)$$

where  $\mu_b^k$  denotes the stationary distribution of agent  $k$ , which are induced by its behavior policy  $\pi_b^k$ . Next, we define the mixing time of the joint Markov chain as follows:

$$t_{\text{mix}}(\epsilon) := \min \left\{ t \mid \sup_{x_0 \in (\mathcal{S} \times \mathcal{A})^K} d_{\text{TV}}(P_t(\cdot | x_0), \mu) \leq \epsilon \right\} \quad \text{and} \quad t_{\text{mix}} := t_{\text{mix}}\left(\frac{1}{4}\right), \quad (111)$$

where

$$P_t(\cdot | x_0) = \prod_{k=1}^K P_t^k(\cdot | x_0^k) \quad (112)$$

denotes the distribution of the joint state-action pairs of all agents after  $t$  transitions starting from  $x_0 = (x_0^1, \dots, x_0^K)$ . The mixing time of the joint Markov chain can be connected to those of the individual chains via the following relation

$$t_{\text{mix}}(\epsilon) \leq \max_k t_{\text{mix}}^k(\epsilon/K), \quad t_{\text{mix}} \leq 4 \log 8K \max_{k \in [K]} t_{\text{mix}}^k, \quad (113)$$

which will be proven at the end of the proof.

We now turn to the proof of Lemma 15. Define the event

$$\mathcal{B}_{u,v}(s, a) := \left\{ \left| \sum_{k=1}^K N_{u,v}^k(s, a) - (v - u) \sum_{k=1}^K \mu_b^k(s, a) \right| \geq \frac{1}{2} (v - u) \sum_{k=1}^K \mu_b^k(s, a) \right\}. \quad (114)$$

We first establish that

$$\max_{x_0 \in (\mathcal{S} \times \mathcal{A})^K} \mathbb{P} \left\{ \mathcal{B}_{u,v}(s, a) \mid \{(s_0^k, a_0^k)\}_{k=1}^K = x_0 \right\} \leq \frac{\delta}{|\mathcal{S}||\mathcal{A}|T^2} \quad (115)$$

for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $1 \leq u < v \leq T$  provided that  $u \geq t_{\text{th}}(s, a)/2$  and  $v - u \geq t_{\text{th}}(s, a)/2$ . To this end, we decompose the probability into two terms as follows:

$$\begin{aligned} & \mathbb{P} \left\{ \mathcal{B}_{u,v}(s, a) \mid \{(s_0^k, a_0^k)\}_{k=1}^K = x_0 \right\} \\ &= \underbrace{\mathbb{P} \left\{ \mathcal{B}_{u,v}(s, a) \mid \{(s_0^k, a_0^k)\}_{k=1}^K \sim \mu \right\}}_{=: G_1} \\ &+ \underbrace{\mathbb{P} \left\{ \mathcal{B}_{u,v}(s, a) \mid \{(s_0^k, a_0^k)\}_{k=1}^K = x_0 \right\} - \mathbb{P} \left\{ \mathcal{B}_{u,v}(s, a) \mid \{(s_0^k, a_0^k)\}_{k=1}^K \sim \mu \right\}}_{=: G_2}, \end{aligned}$$

and show each of the terms is bounded by  $\frac{\delta}{2|\mathcal{S}||\mathcal{A}|T^2}$  for any  $x_0 \in (\mathcal{S} \times \mathcal{A})^K$ . We shall derive the bounds of these two terms separately.

**Step 1: bounding  $G_1$ .** This is for the case that the distribution of the initial state follows the joint stationary distribution. Since the total number of visits can be written as

$$\sum_{k=1}^K N_{u,v}^k(s, a) = \sum_{k=1}^K \sum_{i=u+1}^v Z_i^k(s, a) = \sum_{i=u+1}^v \bar{Z}_i(s, a),$$

where

$$Z_i^k(s, a) = \begin{cases} 1, & \text{if } (s, a) \in (s_{i-1}^k, a_{i-1}^k) \\ 0, & \text{otherwise} \end{cases} \quad \text{and} \quad \bar{Z}_i(s, a) = \sum_{k=1}^K Z_i^k(s, a),$$

and

$$\nu_{u,v}(s, a) := \mathbb{E}_{(s_0^k, a_0^k) \sim \mu^k \forall k \in [K]} \left[ \sum_{i=u+1}^v \bar{Z}_i(s, a) \right] = (v - u) \sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a),$$

we can invoke Bernstein's inequality for Markov chains (Paulin, 2015, Theorem 3.11) and obtain

$$\begin{aligned} G_1 &= \mathbb{P}_{\{(s_0^k, a_0^k)\}_{k=1}^K \sim \mu} \left[ \left| \sum_{i=u+1}^v \bar{Z}_i(s, a) - \nu_{u,v}(s, a) \right| \geq \frac{1}{2} \nu_{u,v}(s, a) \right] \\ &\leq 2 \exp \left( - \frac{(\nu_{u,v}(s, a)/2)^2 \gamma_{\text{ps}}}{8((v - u) + 1/\gamma_{\text{ps}})V_f + 20C(\nu_{u,v}(s, a)/2)} \right). \end{aligned} \quad (116)$$

Here,  $\gamma_{\text{ps}}$  is the pseudo spectral gap satisfying

$$\gamma_{\text{ps}} \geq \frac{1}{2t_{\text{mix}}} \quad (117a)$$

for uniformly ergodic Markov chains according to Paulin (2015, Proposition 3.4). The parameters  $C$  and  $V_f$  are defined and bounded as follows

$$C := \max_{u < i \leq v} |\bar{Z}_i(s, a) - \mathbb{E}[\bar{Z}_i(s, a)]| \leq K, \quad (117b)$$

$$V_f := \text{Var}(\bar{Z}_i(s, a)) = \sum_{k=1}^K (1 - \mu_{\mathbf{b}}^k(s, a)) \mu_{\mathbf{b}}^k(s, a) \leq \sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a). \quad (117c)$$

Plugging (117) into (116), we have

$$\begin{aligned} G_1 &\leq 2 \exp \left( - \frac{(\nu_{u,v}(s, a))^2}{8t_{\text{mix}}(24(v-u)(\sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a)) + 10K\nu_{u,v}(s, a))} \right) \\ &\leq 2 \exp \left( - \frac{(v-u)(\sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a))}{8t_{\text{mix}}(24 + 10K)} \right) \leq \frac{\delta}{2|\mathcal{S}||\mathcal{A}|T^2}, \end{aligned} \quad (118)$$

where the last inequality holds since  $(v-u)$  is large enough to satisfy the following condition:

$$v - u \geq \frac{t_{\text{th}}(s, a)}{2} \geq \frac{1088(\max_{k \in [K]} t_{\text{mix}}^k) \log 8K \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}}{\frac{1}{K} \sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a)} \geq \frac{272t_{\text{mix}} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}}{\frac{1}{K} \sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a)}.$$

**Step 2: bounding  $G_2$ .** By the same argument of Li et al. (2021b, Section A.1), using the fact that the difference caused by the initial state becomes very small after sufficiently long time, we have

$$\begin{aligned} G_2 &:= \mathbb{P} \left\{ \mathcal{B}_{u,v}(s, a) \mid \{(s_0^k, a_0^k)\}_{k=1}^K = x_0 \right\} - \mathbb{P} \left\{ \mathcal{B}_{u,v}(s, a) \mid \{(s_0^k, a_0^k)\}_{k=1}^K \sim \mu \right\} \\ &\leq d_{\text{TV}}(P_u(\cdot | x_0), \mu) \leq \frac{\delta}{2|\mathcal{S}||\mathcal{A}|T^2}, \end{aligned} \quad (119)$$

where the last inequality holds due to

$$u \geq \frac{t_{\text{th}}(s, a)}{2} \geq 4 \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta} \max_{k \in [K]} t_{\text{mix}}^k \geq \max_{k \in [K]} t_{\text{mix}}^k \left( \frac{\delta}{2|\mathcal{S}||\mathcal{A}|T^2K} \right) \geq t_{\text{mix}} \left( \frac{\delta}{2|\mathcal{S}||\mathcal{A}|T^2} \right). \quad (120)$$

Here, the second inequality follows from the fact that  $t_{\text{mix}}^k(\epsilon) \leq 2t_{\text{mix}}^k \log_2 \frac{2}{\epsilon}$  (Paulin, 2015), and the last inequality follows from (113).

**Step 3: summing things up.** By combining the above bound, we complete the proof of (115), provided that  $u \geq t_{\text{th}}(s, a)/2$  and  $v-u \geq t_{\text{th}}(s, a)$ . Then, we can obtain the following bound for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $0 \leq u < v \leq T$ :

$$\begin{aligned} &\mathbb{P} \left\{ \frac{1}{4}(v-u) \sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a) \leq \sum_{k=1}^K N_{u,v}^k(s, a) \leq 2(v-u) \sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a) \right\} \\ &\leq \mathbb{P} \left\{ \left| \sum_{k=1}^K N_{u+\frac{t_{\text{th}}(s,a)}{2}, v}^k(s, a) - \left( v-u - \frac{t_{\text{th}}(s, a)}{2} \right) \sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a) \right| \right\} \end{aligned}$$

$$\begin{aligned}
 & \geq \frac{1}{2} \left( v - u - \frac{t_{\text{th}}(s, a)}{2} \right) \sum_{k=1}^K \mu_{\mathbf{b}}^k(s, a) \Big\} \\
 & = \max_{x_0 \in (\mathcal{S} \times \mathcal{A})^K} \mathbb{P} \left\{ \mathcal{B}_{u + \frac{t_{\text{th}}(s, a)}{2}, v}(s, a) \mid \{(s_0^k, a_0^k)\}_{k=1}^K = x_0 \right\} \leq \frac{\delta}{|\mathcal{S}| |\mathcal{A}| T^2}. \tag{121}
 \end{aligned}$$

**Proof of (113).** Notice that by the definition of  $d_{\text{TV}}$  and (112), we have

$$d_{\text{TV}}(P_t(\cdot | x_0), \mu) \leq \sum_{k=1}^K d_{\text{TV}}(P_t^k(\cdot | x_0^k), \mu_{\mathbf{b}}^k)$$

for any  $x_0 \in (\mathcal{S} \times \mathcal{A})^K$ . Hence, setting  $t = \max_{k \in [K]} t_{\text{mix}}^k(\frac{\epsilon}{K})$ , we have

$$\max_{x_0 \in (\mathcal{S} \times \mathcal{A})^K} d_{\text{TV}}(P_t(\cdot | x_0), \mu) \leq \sum_{k=1}^K \frac{\epsilon}{K} = \epsilon,$$

which immediately implies

$$t_{\text{mix}}(\epsilon) \leq \max_k t_{\text{mix}}^k(\epsilon/K).$$

The proof is complete by using the fact that  $t_{\text{mix}}(\epsilon) \leq 2t_{\text{mix}} \log_2 \frac{2}{\epsilon}$  (Paulin, 2015), which leads to

$$t_{\text{mix}} \leq \max_{k \in [K]} t_{\text{mix}}^k \left( \frac{1}{4K} \right) \leq 4 \log 8K \max_{k \in [K]} t_{\text{mix}}^k.$$

## C.2 Proof of Lemma 6

First, (53a) is derived as follows:

$$\begin{aligned}
 \lambda_{v_1, v_2}(s, a) &= \frac{1}{K} \sum_{k=1}^K (1 - \eta)^{N_{v_1, v_2}^k(s, a)} \leq \frac{1}{K} \sum_{k=1}^K \exp(-\eta N_{v_1, v_2}^k(s, a)) \\
 &\leq 1 - \frac{1}{2} \frac{1}{K} \sum_{k=1}^K \eta N_{v_1, v_2}^k(s, a) \\
 &\leq \exp \left( -\frac{\eta}{2K} \sum_{k=1}^K N_{v_1, v_2}^k(s, a) \right) \tag{122}
 \end{aligned}$$

using the fact that  $1 - x \leq \exp(-x) \leq 1 - \frac{x}{2}$  holds for any  $0 \leq x < 1$ , and  $\eta N_{h\tau, (h+1)\tau}^{k'}(s, a) \leq \eta\tau \leq 1$ .

Next, we obtain (53b) through the following derivation:

$$\begin{aligned}
 & \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} \omega_{u,t}^k(s, a) \\
 &= \sum_{k=1}^K \sum_{h=0}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \omega_{u,t}^k(s, a)
 \end{aligned}$$

$$\begin{aligned}
 &= \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a) \right) \sum_{k=1}^K \frac{1}{K} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \left( \eta(1-\eta)^{N_{u+1, (h+1)\tau}^k(s, a)} \right) \\
 &\stackrel{(i)}{=} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a) \right) \sum_{k=1}^K \frac{1}{K} (1 - (1-\eta)^{N_{h\tau, (h+1)\tau}^k(s, a)}) \\
 &\stackrel{(ii)}{=} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a) \right) (1 - \lambda_{h\tau, (h+1)\tau}(s, a)) \\
 &\stackrel{(iii)}{=} 1 - \lambda_{0,\tau} \lambda_{\tau, 2\tau} \cdots \lambda_{(\phi(t)-1)\tau, t} = 1 - \omega_{0,t}(s, a),
 \end{aligned} \tag{123}$$

where (i) follows from the geometric sum

$$\begin{aligned}
 \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \eta(1-\eta)^{N_{u+1, (h+1)\tau}^k(s, a)} &= \eta + \eta(1-\eta) + \cdots + \eta(1-\eta)^{N_{h\tau, (h+1)\tau}^k(s, a)-1} \\
 &= 1 - (1-\eta)^{N_{h\tau, (h+1)\tau}^k(s, a)},
 \end{aligned} \tag{124}$$

(ii) follows from the definition (50), and (iii) follows by cancellation.

Similarly, (53c) can be obtained with some algebraic calculations as follows:

$$\begin{aligned}
 \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, h'\tau}^k(s, a)} \omega_{u,t}^k(s, a) &= \sum_{k=1}^K \sum_{h=0}^{h'-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \omega_{u,t}^k(s, a) \\
 &\stackrel{(i)}{=} \sum_{h=0}^{h'-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a) \right) (1 - \lambda_{h\tau, (h+1)\tau}(s, a)) \\
 &\stackrel{(ii)}{\leq} \lambda_{h'\tau, (h'+1)\tau} \cdots \lambda_{(\phi(t)-1)\tau, t} - \lambda_{0,\tau} \lambda_{\tau, 2\tau} \cdots \lambda_{(\phi(t)-1)\tau, t} \\
 &\leq \lambda_{h'\tau, (h'+1)\tau} \cdots \lambda_{(\phi(t)-1)\tau, t} \\
 &\stackrel{(iii)}{\leq} \prod_{h=h'}^{\phi(t)-1} \exp \left( -\frac{\eta}{2K} \sum_{k=1}^K N_{h\tau, (h+1)\tau}^k(s, a) \right),
 \end{aligned} \tag{125}$$

where (i) follows from similar derivations as above, (ii) follows by cancellation, and (iii) follows from (53a).

Finally, (53d) is derived as follows:

$$\begin{aligned}
 &\sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} (\omega_{u,t}^k(s, a))^2 \\
 &= \sum_{k=1}^K \sum_{h=0}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} (\omega_{u,t}^k(s, a))^2
 \end{aligned}$$

$$\begin{aligned}
 &= \frac{1}{K} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a) \right)^2 \sum_{k=1}^K \frac{1}{K} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \left( \eta(1-\eta)^{N_{u+1, (h+1)\tau}^k(s, a)} \right)^2 \\
 &\stackrel{(i)}{\leq} \frac{2\eta}{K} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a) \right) \sum_{k=1}^K \frac{1}{K} \left( 1 - (1-\eta)^{N_{h\tau, (h+1)\tau}^k(s, a)} \right) \\
 &= \frac{2\eta}{K} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=(h+1)}^{\phi(t)-1} \lambda_{l\tau, (l+1)\tau}(s, a) \right) (1 - \lambda_{h\tau, (h+1)\tau}(s, a)) \\
 &\stackrel{(ii)}{\leq} \frac{2\eta}{K},
 \end{aligned}$$

where (i) holds since

$$\begin{aligned}
 \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \left( \eta(1-\eta)^{N_{u+1, (h+1)\tau}^k(s, a)} \right)^2 &= \eta^2 + \eta^2(1-\eta)^2 \cdots + \eta^2(1-\eta)^{2(N_{u+1, (h+1)\tau}^k(s, a)-1)} \\
 &\leq \eta \left( 1 - (1-\eta)^{2N_{u+1, (h+1)\tau}^k(s, a)} \right) \\
 &\leq 2\eta \left( 1 - (1-\eta)^{N_{u+1, (h+1)\tau}^k(s, a)} \right)
 \end{aligned} \tag{126}$$

and (ii) can be similarly derived to the proof of (53c) (cf. (125)).

### C.3 Proof of Lemma 7

Without loss of generality, we prove the claim for some fixed  $1 \leq t \leq T$  and  $(s, a) \in \mathcal{S} \times \mathcal{A}$ . For notation simplicity, let

$$y_{u,t}^k(s, a) = \begin{cases} \omega_{u,t}^k(s, a)(P(s, a) - P_{u+1}^k(s, a))V_u^k & \text{if } (s_u^k, a_u^k) = (s, a) \\ 0 & \text{otherwise} \end{cases}, \tag{127}$$

where

$$\omega_{u,t}^k(s, a) = \frac{\eta}{K} (1-\eta)^{N_{u+1, (\phi(u)+1)\tau}^k(s, a)} \prod_{h=\phi(u)+1}^{\phi(t)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right), \tag{128}$$

then  $E_t^2(s, a) = \gamma \sum_{k=1}^K \sum_{u=0}^{t-1} y_{u,t}^k(s, a)$ . However, due to the dependency between  $P_{u+1}^k(s, a)$  and  $\omega_{u,t}^k(s, a)$  arising from the Markovian sampling, it is difficult to track the sum of  $y := \{y_{u,t}^k(s, a)\}$  directly. To address this issue, we will first analyze the sum using a collection of approximate random variables  $\hat{y} = \{\hat{y}_{u,t}^k(s, a)\}$  drawn from a carefully constructed set  $\hat{\mathcal{Y}}$ , which is closely coupled with the target  $\{y_{u,t}^k(s, a)\}_{0 \leq u < t}$ , i.e.,

$$D(y, \hat{y}) := \left| \sum_{k=1}^K \sum_{u=0}^{t-1} (y_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a)) \right| \tag{129}$$



is sufficiently small. In addition,  $\hat{y}$  shall exhibit some useful statistical independence and thus easier to control its sum; we shall control this over the entire set  $\hat{\mathcal{Y}}$ . Finally, leveraging the proximity above, we can obtain the desired bound on  $y$  via triangle inequality. We now provide details on executing this proof outline, where the crust is in designing the set  $\hat{\mathcal{Y}}$  with a controlled size.

Before describing our construction, let's introduce the following useful event:

$$\mathcal{B}_M(s, a) := \bigcap_{u=0}^{t-M\tau} \left\{ \frac{1}{4} \mu_{\text{avg}}(s, a) K M \tau \leq \sum_{k=1}^K N_{u, u+M\tau}^k(s, a) \leq 2 \mu_{\text{avg}}(s, a) K M \tau \right\}, \quad (130)$$

where  $M = M(s, a) := \lfloor \frac{1}{8\eta\mu_{\text{avg}}(s, a)\tau} \rfloor$ . Note that  $M\tau \geq \tau \geq t_{\text{th}}$  (see (78) for the definition of  $t_{\text{th}}(s, a)$ ), and  $1 \leq 1/(16\eta\mu_{\text{avg}}(s, a)\tau) \leq M(s, a) \leq 1/(8\eta\mu_{\text{avg}}(s, a)\tau)$  if  $\eta\tau \leq 1/16$ . Then,  $\mathcal{B}_M(s, a)$  holds with probability at least  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T}$  according to Lemma 15. The rest of the proof shall be carried out under the event  $\mathcal{B}_M(s, a)$ .

**Step 1: constructing  $\hat{\mathcal{Y}}$ .** To decouple dependency between  $P_{u+1}^k(s, a)$  and  $\omega_{u,t}^k(s, a)$ , we will introduce approximates of  $\omega_{u,t}^k(s, a)$  that only depend on history until  $u$  by replacing a factor dependent on future with some constant. To gain insight, we first decompose  $\omega_{u,t}^k(s, a)$  as follows:

$$\begin{aligned} \omega_{u,t}^k(s, a) &= \frac{\eta}{K} (1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)} \frac{(1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)}}{\sum_{k'=1}^K (1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}} \prod_{h=\phi(u)}^{\phi(t)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right) \\ &= \underbrace{\frac{\eta}{K} (1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)} \prod_{h=\phi(u)}^{\phi(t)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right)}_{:= \bar{\omega}_{u,t}^k(s, a)} \\ &\quad + \underbrace{\frac{\eta}{K} (1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)} \left( \frac{(1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)}}{\sum_{k'=1}^K (1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}} - 1 \right) \prod_{h=\phi(u)}^{\phi(t)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right)}_{:= \chi_{u,t}^k(s, a)}. \end{aligned}$$

Considering that  $\chi_{u,t}^k(s, a)$  can be made small enough, which will be shown in the following step, we analyze the dominant factor  $\bar{\omega}_{u,t}^k(s, a)$  in detail as follows:

$$\begin{aligned} \bar{\omega}_{u,t}^k(s, a) &= \prod_{h=h_0(u, t)}^{\phi(u)-1} \left( \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right) \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right)^{-1} \right) \\ &\quad \times \frac{\eta}{K} (1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)} \prod_{h=\phi(u)}^{\phi(t)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right) \end{aligned}$$

$$\begin{aligned}
 &= \underbrace{\frac{\eta}{K} (1 - \eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)} \prod_{h=h_0(u, t)}^{\phi(u)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right)^{-1}}_{\text{dependent on history until } u} \\
 &\quad \times \underbrace{\prod_{h=h_0(u, t)}^{\phi(t)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right)}_{\text{dependent on history and future until } t} \\
 &= \underbrace{\frac{\eta}{K} (1 - \eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)} \prod_{h=h_0(u, t)}^{\phi(u)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right)^{-1}}_{:= x_u^k(s, a)} \\
 &\quad \times \underbrace{\prod_{l=1}^{l(u, t)} \prod_{h=\max\{0, \phi(t)-lM\}}^{\phi(t)-(l-1)M-1} \left( \frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right)}_{:= z_l(s, a)},
 \end{aligned} \tag{131}$$

where we denote  $h_0(u, t) = \max\{0, \phi(t) - l(u, t)M\}$ , with  $l(u, t) := \lceil \frac{t-u}{M\tau} \rceil$ .

Motivated by the above decomposition, we will construct  $\hat{\mathcal{Y}}$  by approximating the future-dependent parameter  $z_l(s, a)$  for  $1 \leq l \leq L$ , where we define

$$L := \min \left\{ \left\lceil \frac{t}{M\tau} \right\rceil, \lceil 128 \log(K/\eta) \rceil \right\}. \tag{132}$$

We note that  $L \leq 128 \log(TK)$  for  $\eta \geq 3/T$ . Using the fact that  $1 - x \leq \exp(-x) \leq 1 - \frac{x}{2}$  holds for any  $0 \leq x < 1$ , and  $\eta N_{h\tau, (h+1)\tau}^{k'}(s, a) \leq \eta\tau \leq \frac{1}{2}$ ,

$$\begin{aligned}
 \exp \left( -\frac{2\eta}{K} \sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a) \right) &\leq 1 - \frac{\eta}{K} \sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a) \\
 &\leq \frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \\
 &\leq \frac{1}{K} \sum_{k'=1}^K \exp(-\eta N_{h\tau, (h+1)\tau}^{k'}(s, a)) \\
 &\leq 1 - \frac{1}{2} \frac{1}{K} \sum_{k'=1}^K \eta N_{h\tau, (h+1)\tau}^{k'}(s, a) \\
 &\leq \exp \left( -\frac{\eta}{2K} \sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a) \right).
 \end{aligned} \tag{133}$$

Therefore, for  $1 \leq l < L$ , under  $\mathcal{B}_M(s, a)$ , the range of  $z_l(s, a)$  is bounded as follows:

$$z_l(s, a) \in \left[ \exp(-4\eta\mu_{\text{avg}}(s, a)M\tau), \exp(-\frac{1}{8}\eta\mu_{\text{avg}}(s, a)M\tau) \right].$$

Using this property, we construct a set of values that can cover possible realizations of  $z_l(s, a)$  in a fine-grained manner as follows:

$$\mathcal{Z} := \left\{ \exp \left( -\frac{1}{8} \eta \mu_{\text{avg}}(s, a) M \tau - \frac{i\eta}{K} \right) \mid i \in \mathbb{Z} : 0 \leq i < 4K \mu_{\text{avg}}(s, a) M \tau \right\}. \quad (134)$$

Note that the distance of adjacent elements of  $\mathcal{Z}$  is bounded by  $\eta/K e^{-1/8 \eta \mu_{\text{avg}}(s, a) M \tau}$ , and the size of the set is bounded by  $4K \mu_{\text{avg}}(s, a) M \tau$ . For  $l = L$ , because the number of iterations involved in  $z_L(s, a)$  can be less than  $M \tau$ , it follows that  $z_L(s, a) \in [\exp(-4 \eta \mu_{\text{avg}}(s, a) M \tau), 1]$ . Hence, we construct the set

$$\mathcal{Z}_0 := \left\{ \exp \left( -\frac{i\eta}{K} \right) \mid i \in \mathbb{Z} : 0 \leq i < 4K \mu_{\text{avg}}(s, a) M \tau \right\}. \quad (135)$$

In sum, we can always find  $(\hat{z}_1, \dots, \hat{z}_l, \dots, \hat{z}_L) \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$  where its entry-wise distance to  $(z_l(s, a))_{l \in [L-1]}$  (resp.  $z_L(s, a)$ ) is at most  $\eta/K e^{-1/8 \eta \mu_{\text{avg}}(s, a) M \tau}$  (resp.  $\eta/K$ ).

Moreover, we approximate  $x_u^k(s, a)$  by clipping it when the accumulated number of visits of all agents is not too large as follows:

$$\hat{x}_u^k(s, a) = \begin{cases} x_u^k(s, a) & \text{if } \sum_{k=1}^K N_{h_0(u, t)\tau, \phi(u)\tau}^k(s, a) \leq 2K \mu_{\text{avg}}(s, a) M \tau \\ 0 & \text{otherwise} \end{cases}. \quad (136)$$

Note that the clipping never occurs and  $\hat{x}_u^k(s, a) = x_u^k(s, a)$  for all  $u$  as long as  $\mathcal{B}_M(s, a)$  holds. To provide useful properties of  $\hat{x}_u^k(s, a)$  that will be useful later, we record the following lemma whose proof is provided in Appendix C.3.1.

**Lemma 17** *For any state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , consider any integers  $1 \leq t \leq T$  and  $1 \leq l \leq \lceil \frac{t}{M\tau} \rceil$ , where  $M = \lfloor \frac{1}{8\eta \mu_{\text{avg}}(s, a)\tau} \rfloor$ . Suppose that  $4\eta\tau \leq 1$ , then  $\hat{x}_u^k(s, a)$  defined in (136) satisfy*

$$\forall u \in [h_0, \phi(t) - (l-1)M] : \hat{x}_u^k(s, a) \leq \frac{9\eta}{K}, \quad (137a)$$

$$\sum_{h=h_0}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) \leq 16\eta \mu_{\text{avg}}(s, a) M \tau, \quad (137b)$$

$$\sum_{h=h_0}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K (\hat{x}_u^k(s, a))^2 \leq \frac{64\eta^2 \mu_{\text{avg}}(s, a) M \tau}{K}, \quad (137c)$$

where  $h_0 = \max\{0, \phi(t) - lM\}$ .

Finally, for each  $\mathbf{z} = (\hat{z}_1, \dots, \hat{z}_L) \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$ , setting

$$\hat{\omega}_{u,t}^k(s, a; \mathbf{z}) = \hat{x}_u^k(s, a) \prod_{l=1}^{l(u,t)} \hat{z}_l, \quad (138)$$

an approximate random sequence  $\hat{y}_z = \{\hat{y}_{u,t}^k(s, a; z)\}_{0 \leq u < t}$  can be constructed as follows:

$$\hat{y}_{u,t}^k(s, a; z) = \begin{cases} \hat{\omega}_{u,t}^k(s, a; z)(P(s, a) - P_{u+1}^k(s, a))V_u^k & \text{if } (s_u^k, a_u^k) = (s, a) \text{ and } l(u, t) \leq L \\ 0 & \text{otherwise} \end{cases}. \quad (139)$$

If  $t > LM\tau$ , for any  $u < t - LM\tau$ , i.e.,  $l(u, t) > L$ , we set  $\hat{y}_{u,t}^k(s, a; z) = 0$  since the magnitude of  $\omega_{u,t}^k(s, a)$  becomes negligible when the time difference between  $u$  and  $t$  is large enough, and the fine-grained approximation using  $\mathcal{Z}$  is no longer needed, as shall be seen momentarily. Finally, denote a collection of the approximates induced by  $\mathcal{Z}^{L-1} \times \mathcal{Z}_0$  as

$$\hat{\mathcal{Y}} = \{\hat{y}_z : z \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0\}.$$

**Step 2: bounding the approximation error  $D(y, \hat{y}_z)$ .** We now show that under  $\mathcal{B}_M(s, a)$ , there exists  $\hat{y}_z := \hat{y}_{z(y)} \in \hat{\mathcal{Y}}$  such that

$$D(y, \hat{y}_z) < \frac{525}{1 - \gamma} \sqrt{\frac{C_{\text{het}} \eta L}{K} \log \frac{4|S||\mathcal{A}|T^2}{\delta}} \quad (140)$$

with at least probability  $1 - 2\delta$ . To this end, we first decompose the approximation error as follows:

$$\begin{aligned} & \min_{\hat{y}_z \in \hat{\mathcal{Y}}} D(y, \hat{y}_z) \\ &= \min_{z \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=0}^{t-1} \left( y_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; z) \right) \right| \\ &\leq \max_{z \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=0}^{t-LM\tau-1} y_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; z) \right| \\ &\quad + \min_{z \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} y_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; z) \right| \\ &\leq \underbrace{\max_{z \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=0}^{t-LM\tau-1} y_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; z) \right|}_{=: D_1} \\ &\quad + \underbrace{\min_{z \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} (\bar{\omega}_{u,t}^k(s, a) - \hat{\omega}_{u,t}^k(s, a; z))(P(s, a) - P_{u+1}^k(s, a))V_u^k \right|}_{=: D_2} \\ &\quad + \underbrace{\left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} \chi_{u,t}^k(s, a)(P(s, a) - P_{u+1}^k(s, a))V_u^k \right|}_{=: D_3}, \end{aligned}$$

and will bound each term separately.

- **Bounding  $D_1$ .** This term appears only when  $t > LM\tau$ . Since  $\hat{y}_{u,t}^k(s, a; \mathbf{z}) = 0$  for all  $u < t - LM\tau$  regardless of  $\mathbf{z}$  by construction,

$$\begin{aligned}
 & \left| \sum_{k=1}^K \sum_{u=0}^{t-LM\tau-1} y_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right| \\
 & \leq \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t-LM\tau}^k(s, a)} \omega_{u,t}^k(s, a) \|P(s, a) - P_{u+1}^k(s, a)\|_1 \|V_u^k\|_\infty \\
 & \stackrel{(i)}{\leq} \frac{2}{1-\gamma} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t-LM\tau}^k(s, a)} \omega_{u,t}^k(s, a) \\
 & \leq \frac{2}{1-\gamma} \prod_{h=\phi(t)-LM}^{\phi(t)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right) \\
 & \stackrel{(ii)}{\leq} \frac{2}{1-\gamma} \exp \left( -\frac{\eta}{2K} \sum_{k'=1}^K N_{t-LM\tau, t}^{k'}(s, a) \right) \\
 & \stackrel{(iii)}{\leq} \frac{2}{1-\gamma} \exp \left( -\frac{1}{8} \eta \mu_{\text{avg}}(s, a) LM\tau \right) \\
 & \stackrel{(iv)}{\leq} \frac{2\eta}{(1-\gamma)K},
 \end{aligned}$$

where (i) holds since  $\|P(s, a)\|_1$ ,  $\|P_u^k(s, a)\|_1 \leq 1$  and  $\|V_{u-1}^k\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)), (ii) follows from (133), (iii) holds due to  $\mathcal{B}_M(s, a)$ , and (iv) holds because  $L \geq 128 \log \frac{K}{\eta} \geq \frac{8}{\eta \mu_{\text{avg}}(s, a) M\tau} \log \frac{K}{\eta}$  given that  $\eta \mu_{\text{avg}}(s, a) M\tau \geq 1/16$ .

- **Bounding  $D_2$ .** Since  $\hat{x}_u^k(s, a) = x_u^k(s, a)$  when  $\mathcal{B}_M(s, a)$  holds, in view of (139), we have

$$\begin{aligned}
 & \min_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} (\bar{\omega}_{u,t}^k(s, a) - \hat{\omega}_{u,t}^k(s, a; \mathbf{z})) (P(s, a) - P_{u+1}^k(s, a)) V_u^k \right| \\
 & \leq \min_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{t-LM\tau, t}^k(s, a)} |\bar{\omega}_{u,t}^k(s, a) - \hat{\omega}_{u,t}^k(s, a; \mathbf{z})| \|P(s, a) - P_{u+1}^k(s, a)\|_1 \|V_u^k\|_\infty \\
 & \leq \frac{2}{1-\gamma} \min_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left( \sum_{l=1}^L \sum_{h=\phi(t)-lM}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) \left| \prod_{l'=1}^l z_{l'}(s, a) - \prod_{l'=1}^l \hat{z}_{l'}^* \right| \right),
 \end{aligned}$$

where the last inequality holds since  $\|P(s, a)\|_1$ ,  $\|P_u^k(s, a)\|_1 \leq 1$  and  $\|V_{u-1}^k\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)) and the definition of  $\hat{\omega}_{u,t}^k(s, a; \mathbf{z})$  defined in (138).

Note that for any given  $\{z_l(s, a)\}_{l \in [L]}$ , under  $\mathcal{B}_M(s, a)$ , there exists  $\hat{\mathbf{z}}^* = (\hat{z}_1^*, \dots, \hat{z}_L^*, \dots, \hat{z}_L^*) \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$  such that  $|\hat{z}_l^* - z_l(s, a)| \leq \frac{\eta}{K} \exp(-1/8 \eta \mu_{\text{avg}}(s, a) M\tau)$  for  $l < L$  and

$|\hat{z}_L^* - z_L(s, a)| \leq \frac{\eta}{K}$ . Also, recall that  $z_l(s, a)$ ,  $\hat{z}_l^* \leq \exp(-1/8\eta\mu_{\text{avg}}(s, a)M\tau)$  for  $l < L$  and  $z_L(s, a)$ ,  $\hat{z}_L^* \leq 1$ . Then, for any  $l \leq L$  it follows that:

$$\begin{aligned} \left| \prod_{l'=1}^l z_{l'}(s, a) - \prod_{l'=1}^l \hat{z}_{l'}^* \right| &\leq \left( \left| \prod_{l'=1}^l z_{l'}(s, a) - \hat{z}_1^* \prod_{l'=2}^l z_{l'}(s, a) \right| + \cdots + \left| z_l \prod_{l'=1}^{l-1} \hat{z}_{l'}^* - \prod_{l'=1}^l \hat{z}_{l'}^* \right| \right) \\ &\leq \exp\left(-\frac{1}{8}(l-1)\eta\mu_{\text{avg}}(s, a)M\tau\right) \sum_{l'=1}^l \frac{\eta}{K} \\ &\leq \exp\left(-\frac{1}{8}(l-1)\eta\mu_{\text{avg}}(s, a)M\tau\right) \frac{L\eta}{K}. \end{aligned}$$

Then, applying the above bound and (137b) in Lemma 17,

$$\begin{aligned} D_2 &\leq \frac{2}{1-\gamma} \sum_{l=1}^L \sum_{h=\phi(t)-lM}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) \left| \prod_{l'=1}^l z_{l'}(s, a) - \prod_{l'=1}^l \hat{z}_{l'}^* \right| \\ &\leq \frac{2}{1-\gamma} \frac{L\eta}{K} \sum_{l=1}^L \exp\left(-\frac{1}{8}(l-1)\eta\mu_{\text{avg}}(s, a)M\tau\right) \\ &\quad \times \sum_{h=\phi(t)-lM}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) \\ &\leq \frac{2}{1-\gamma} \frac{L\eta}{K} \frac{1}{1 - \exp(-1/8\eta\mu_{\text{avg}}(s, a)M\tau)} (16\eta\mu_{\text{avg}}(s, a)M\tau) \\ &\stackrel{(i)}{\leq} \frac{2}{1-\gamma} \frac{L\eta}{K} \frac{16}{\eta\mu_{\text{avg}}(s, a)M\tau} 16\eta\mu_{\text{avg}}(s, a)M\tau \leq \frac{512\eta L}{(1-\gamma)K}, \end{aligned}$$

where (i) holds since  $\eta\mu_{\text{avg}}(s, a)M\tau/8 \leq 1$  and  $e^{-x} \leq 1 - \frac{1}{2}x$  for any  $0 \leq x \leq 1$ .

- **Bounding  $D_3$ .** Applying Freedman's inequality, we can obtain the following bound, whose proof is provided in Appendix C.3.2.

**Lemma 18** *Consider any  $\delta \in (0, 1)$  and  $L$  defined in (132). For any  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $1 \leq t \leq T$ , the following holds:*

$$D_3 \leq \frac{9}{1-\gamma} \sqrt{\frac{C_{\text{het}}\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}} \quad (141)$$

with probability at least  $1-2\delta$ , as long as  $\tau \geq t_{\text{th}}$ , and  $\eta \leq \min\{\frac{1}{4\tau K}, \frac{1}{KC_{\text{het}}L \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}}\}$ .

By combining the bounds obtained above,

$$\begin{aligned} \min_{\hat{y}_z \in \hat{\mathcal{Y}}} D(y, \hat{y}_z) &\leq \frac{2\eta}{(1-\gamma)K} + \frac{512\eta L}{(1-\gamma)K} + \frac{9}{1-\gamma} \sqrt{\frac{C_{\text{het}}\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}} \\ &\leq \frac{525}{1-\gamma} \sqrt{\frac{C_{\text{het}}\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}} \end{aligned}$$

since  $\eta \leq \frac{K}{128 \log(TK)} \leq K/L$  due to  $L \leq 128 \log(TK)$ .

**Step 3: concentration bound over  $\mathcal{Y}$ .** We now show that for all elements in  $\widehat{\mathcal{Y}} = \{\widehat{y}_{\mathbf{z}} : \mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0\}$  satisfy

$$\left| \sum_{k=1}^K \sum_{u=0}^{t-1} \widehat{y}_{u,t}^k(s, a; \mathbf{z}) \right| < \frac{115}{(1-\gamma)} \sqrt{\frac{\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta}} \quad (142)$$

with probability at least  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T}$ . It suffices to establish (142) for a fixed  $\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$  with probability at least  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T|\mathcal{Y}|}$ , where

$$|\widehat{\mathcal{Y}}| = |\mathcal{Z}^{L-1} \times \mathcal{Z}_0| \leq (4K\mu_{\text{avg}}(s, a)M\tau)^L \leq (K/\eta)^L \leq (TK)^L \quad (143)$$

because  $\eta\mu_{\text{avg}}(s, a)M\tau \leq 1/4$  and  $\eta \geq 1/T$ .

For any fixed  $\mathbf{z} = (\widehat{z}_1, \dots, \widehat{z}_L) \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$ , since  $\widehat{\omega}_{u,t}^k(s, a; \mathbf{z}) = \widehat{x}_u^k(s, a) \prod_{l=1}^{l(u,t)} \widehat{z}_l$  only depends on the events happened until  $u$ , which is independent to a transition at  $u+1$ . Thus, we can apply Freedman's inequality to bound the sum of  $\widehat{y}_{u,t}^k(s, a; \mathbf{z})$  since

$$\mathbb{E}[\widehat{y}_{u,t}^k(s, a; \mathbf{z}) | \mathcal{Y}_u] = 0, \quad (144)$$

where  $\mathcal{Y}_u$  denotes the history of visited state-action pairs and updated values of all agents until  $u$ , i.e.,  $\mathcal{Y}_u = \{(s_v^k, a_v^k), V_v^k\}_{k \in [K], v \leq u}$ . Before applying Freedman's inequality, we need to calculate the following quantities. First,

$$\begin{aligned} B_t(s, a) &:= \max_{k \in [K], 0 \leq u < t} |\widehat{y}_{u,t}^k(s, a; \mathbf{z})| \leq \widehat{x}_u^k(s, a) \prod_{l=1}^{l(u,t)} \widehat{z}_l \|P(s, a) - P_{u+1}^k(s, a)\|_1 \|V_u^k\|_{\infty} \\ &\leq \frac{18\eta}{(1-\gamma)K}, \end{aligned} \quad (145)$$

where the last inequality follows from  $\|P(s, a)\|_1, \|P_u^k(s, a)\|_1 \leq 1, \|V_{u-1}^k\|_{\infty} \leq \frac{1}{1-\gamma}$  (cf. (30)),  $\widehat{z}_l \leq 1$ , and (137a) in Lemma 17. Next, we can bound the variance as

$$\begin{aligned} W_t(s, a) &:= \sum_{u=0}^t \sum_{k=1}^K \mathbb{E}[(\widehat{y}_{u,t}^k(s, a; \mathbf{z}))^2 | \mathcal{Y}_u] \\ &= \sum_{l=1}^L \sum_{h=\max\{0, \phi(t)-lM\}}^{\phi(t)-(l-1)M-1} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} (\widehat{x}_u^k(s, a) \prod_{l'=1}^l \widehat{z}_{l'})^2 \text{Var}_{P(s, a)}(V_u^k) \\ &\stackrel{(i)}{\leq} \frac{2}{(1-\gamma)^2} \sum_{l=1}^L \left( \prod_{l'=1}^l \widehat{z}_{l'}^2 \right) \sum_{h=\max\{0, \phi(t)-lM\}}^{\phi(t)-(l-1)M-1} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} (\widehat{x}_u^k(s, a))^2 \\ &\stackrel{(ii)}{\leq} \frac{2}{(1-\gamma)^2} \sum_{l=1}^L \left( \prod_{l'=1}^l \widehat{z}_{l'}^2 \right) \frac{64\eta^2 \mu_{\text{avg}}(s, a)M\tau}{K} \\ &\stackrel{(iii)}{\leq} \frac{128\eta^2 \mu_{\text{avg}}(s, a)M\tau}{K(1-\gamma)^2} \sum_{l=1}^L \exp(-1/4(l-1)\eta\mu_{\text{avg}}(s, a)M\tau) \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{128\eta^2\mu_{\text{avg}}(s,a)M\tau}{K(1-\gamma)^2} \frac{1}{1 - \exp(-1/4\eta\mu_{\text{avg}}(s,a)M\tau)} \\
 &\stackrel{\text{(iv)}}{\leq} \frac{128\eta^2\mu_{\text{avg}}(s,a)M\tau}{K(1-\gamma)^2} \frac{8}{\eta\mu_{\text{avg}}(s,a)M\tau} = \frac{1024\eta}{K(1-\gamma)^2} =: \sigma^2,
 \end{aligned} \tag{146}$$

where (i) holds due to the fact that  $\|\text{Var}_P(V)\|_\infty \leq \|P\|_1(\|V\|_\infty)^2 + (\|P\|_1\|V\|_\infty)^2 \leq \frac{2}{(1-\gamma)^2}$  because  $\|V\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)) and  $\|P\|_1 \leq 1$ , (ii) follows from (137c) in Lemma 17, (iii) holds due to the range of  $\mathcal{Z}$  and  $\mathcal{Z}_0$  is bounded by  $\exp(-1/8\eta\mu_{\text{avg}}(s,a)M\tau)$  and 1, respectively, and (iv) holds since  $e^{-x} \leq 1 - \frac{1}{2}x$  for any  $0 \leq x \leq 1$  and  $\eta\mu_{\text{avg}}(s,a)M\tau/4 \leq 1$ .

Now, by substituting the above bounds of  $W_t$  and  $B_t$  into Freedman's inequality (see Theorem 12) and setting  $m = 1$ , it follows that for any  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}$ ,  $t \in [T]$  and  $\hat{y}_{\mathbf{z}} \in \hat{\mathcal{Y}}$ ,

$$\begin{aligned}
 &\left| \sum_{k=1}^K \sum_{u=0}^{t-1} \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right| \\
 &\leq \sqrt{8 \max\{W_t(s, a), \frac{\sigma^2}{2m}\} \log \frac{4m|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}{\delta}} + \frac{4}{3} B_t(s, a) \log \frac{4m|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}{\delta} \\
 &\leq \sqrt{8192 \frac{\eta}{K(1-\gamma)^2} \log \frac{4|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}{\delta}} + \frac{24\eta}{K(1-\gamma)} \log \frac{4|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}{\delta} \\
 &\stackrel{\text{(i)}}{\leq} \frac{115}{(1-\gamma)} \sqrt{\frac{\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta}},
 \end{aligned} \tag{147}$$

with at least probability  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}$ , where (i) holds because  $|\hat{\mathcal{Y}}| \leq (TK)^L$  (cf. (143)), and  $\frac{\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta} \leq 1$  when  $L \leq 128 \log(TK)$  and  $\eta \leq \frac{K}{128 \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta}}$ . Therefore, it follows that (142) holds.

**Step 4: putting things together.** We now putting all the results obtained in the previous steps together to achieve the claimed bound. Under  $\mathcal{B}_M(s, a)$ , there exists  $\hat{y}_{\mathbf{z}} := \hat{y}_{\mathbf{z}(y)} \in \hat{\mathcal{Y}}$  such that (140) holds. Hence,

$$\begin{aligned}
 \sum_{k=1}^K \sum_{u=0}^{t-1} y_{u,t}^k(s, a) &\leq \left| \sum_{k=1}^K \sum_{u=0}^{t-1} \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right| + D(y, \hat{y}_{\mathbf{z}}) \\
 &\leq \frac{115}{(1-\gamma)} \sqrt{\frac{\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta}} + \frac{525}{1-\gamma} \sqrt{\frac{C_{\text{het}}\eta L}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}} \\
 &\leq \frac{7241}{(1-\gamma)} \sqrt{\frac{C_{\text{het}}\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta}},
 \end{aligned}$$

where the second line holds due to (142) and (140), and the last line holds because  $L \leq 128 \log(TK)$ . By taking a union bound over all  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $t \in [T]$ , we complete the proof.



### C.3.1 PROOF OF LEMMA 17

For notational simplicity, let  $\bar{h}$  be the largest integer among  $h \in (h_0, \phi(t) - (l-1)M)$  such that

$$\sum_{k=1}^K N_{h_0\tau, (h-1)\tau}^k(s, a) \leq 2K\mu_{\text{avg}}(s, a)M\tau. \quad (148)$$

Then, the following holds:

$$\begin{aligned} \sum_{k=1}^K N_{h_0\tau, \bar{h}\tau}^k(s, a) &= \sum_{k=1}^K N_{(\bar{h}-1)\tau, \bar{h}\tau}^k(s, a) + \sum_{k=1}^K N_{h_0\tau, (\bar{h}-1)\tau}^k(s, a) \\ &\leq K\tau + 2K\mu_{\text{avg}}(s, a)M\tau. \end{aligned} \quad (149)$$

Also, for the following proofs, we provide an useful bound as follows:

$$\begin{aligned} \sum_{k'=1}^K \frac{(1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}}{K} &\leq \frac{\sum_{k'=1}^K e^{\eta N_{h\tau, (h+1)\tau}^{k'}(s, a)}}{K} \leq 1 + 2\eta \frac{\sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a)}{K} \\ &\leq \exp\left(2\eta \frac{\sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a)}{K}\right), \end{aligned} \quad (150)$$

which holds since  $1+x \leq e^x \leq 1+2x$  for any  $x \in [0, 1]$  and  $\eta N_{h\tau, (h+1)\tau}^{k'}(s, a) \leq \eta\tau \leq 1$ .

According to (136), for any integer  $u \in [\bar{h}\tau, t - (l-1)M\tau)$ ,  $\hat{x}_u^k(s, a)$  is clipped to zero. Now, we prove the bounds in Lemma 17 respectively.

**Proof of (137a).** For  $u \in [h_0\tau, \bar{h}\tau)$ ,

$$\begin{aligned} \hat{x}_u^k(s, a) &= \frac{\eta}{K} (1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)} \prod_{h=h_0(u, t)}^{\phi(u)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right)^{-1} \\ &\stackrel{(i)}{\leq} \frac{3\eta}{K} \prod_{h=h_0(u, t)}^{\phi(u)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right)^{-1} \\ &\stackrel{(ii)}{\leq} \frac{3\eta}{K} \exp\left(\frac{2\eta}{K} \sum_{k'=1}^K N_{h_0\tau, (\bar{h}-1)\tau}^{k'}(s, a)\right) \\ &\stackrel{(iii)}{\leq} \frac{3\eta}{K} \exp(4\eta\mu_{\text{avg}}(s, a)M\tau) \stackrel{(iv)}{\leq} \frac{9\eta}{K}, \end{aligned}$$

where (i) holds since  $(1+\eta)^x \leq e^{\eta x}$  and  $\eta N_{\phi(u)\tau, u+1}^k(s, a) \leq \eta\tau \leq 1$ , (ii) holds due to (133) and the fact that  $\phi(u) \leq \bar{h} - 1$ , (iii) follows from the condition of  $\bar{h}$  in (148), and (iv) holds because  $4\eta\mu_{\text{avg}}(s, a)M\tau \leq 1$ .

**Proof of (137b).** By the definition of  $\bar{h}$ , it follows that

$$\sum_{h=h_0}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) = \sum_{h=h_0}^{\bar{h}-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K x_u^k(s, a).$$

Using the following relation for each  $h$ :

$$\begin{aligned} & \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K x_u^k(s, a) \\ &= \frac{1}{K} \left( \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \eta(1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)} \right) \prod_{h'=h_0}^{h-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h'\tau, (h'+1)\tau}^{k'}(s, a)} \right)^{-1} \\ &= \left( \frac{1}{K} \sum_{k=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^k(s, a)} - 1 \right) \prod_{h'=h_0}^{h-1} \left( \frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{h'\tau, (h'+1)\tau}^{k'}(s, a)} \right)^{-1} \\ &\leq \left( \frac{1}{K} \sum_{k=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^k(s, a)} - 1 \right) \prod_{h'=h_0}^{h-1} \left( \frac{1}{K} \sum_{k=1}^K (1-\eta)^{-N_{h'\tau, (h'+1)\tau}^k(s, a)} \right), \end{aligned}$$

where the last inequality follows from Jensen's inequality, and applying (150), we can complete the proof as follows:

$$\begin{aligned} \sum_{h=h_0}^{\bar{h}-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K x_u^k(s, a) &\leq \prod_{h'=h_0}^{\bar{h}-1} \left( \frac{1}{K} \sum_{k=1}^K (1-\eta)^{-N_{h'\tau, (h'+1)\tau}^k(s, a)} \right) - 1 \\ &\leq \exp \left( \frac{2\eta \sum_{k'=1}^K N_{h_0\tau, \bar{h}\tau}^{k'}(s, a)}{K} \right) - 1 \\ &\stackrel{(i)}{\leq} \exp(4\eta\mu_{\text{avg}}(s, a)M\tau + 2\eta\tau) - 1 \\ &\stackrel{(ii)}{\leq} 16\eta\mu_{\text{avg}}(s, a)M\tau, \end{aligned}$$

where (i) follows from (149), and (ii) holds because  $e^x \leq 1+2x$  for any  $x \in [0, 1]$ ,  $2\eta\tau \leq 1/2$ , and  $4\eta\mu_{\text{avg}}(s, a)M\tau \leq 1/2$ .

**Proof of (137c).** Similarly,

$$\sum_{h=h_0}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K (\hat{x}_u^k(s, a))^2 = \sum_{h=h_0}^{\bar{h}-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K (x_u^k(s, a))^2.$$

Using the following relation for each  $h$ :

$$\sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K (x_u^k(s, a))^2$$

$$\begin{aligned}
 &= \frac{1}{K^2} \left( \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \eta^2 (1 - \eta)^{-2N_{\phi(u)\tau, u+1}^k(s, a)} \right) \prod_{h'=h_0}^{h-1} \left( \frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{h'\tau, (h'+1)\tau}^{k'}(s, a)} \right)^{-2} \\
 &\leq \frac{\eta}{K} \left( \frac{1}{K} \sum_{k=1}^K (1 - \eta)^{-2N_{h\tau, (h+1)\tau}^k(s, a)} - 1 \right) \prod_{h'=h_0}^{h-1} \left( \frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{h'\tau, (h'+1)\tau}^{k'}(s, a)} \right)^{-2} \\
 &\leq \frac{\eta}{K} \left( \frac{1}{K} \sum_{k=1}^K (1 - \eta)^{-2N_{h\tau, (h+1)\tau}^k(s, a)} - 1 \right) \prod_{h'=h_0}^{h-1} \left( \frac{1}{K} \sum_{k=1}^K (1 - \eta)^{-2N_{h'\tau, (h'+1)\tau}^k(s, a)} \right),
 \end{aligned}$$

where the last inequality follows from Jensen's inequality, and applying (150) under the condition  $2\eta\tau \leq 1$ , we can complete the proof as follows:

$$\begin{aligned}
 \sum_{h=h_0}^{\bar{h}-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K (x_u^k(s, a))^2 &\leq \frac{\eta}{K} \prod_{h'=h_0}^{\bar{h}-1} \left( \frac{1}{K} \sum_{k=1}^K (1 - \eta)^{-2N_{h'\tau, (h'+1)\tau}^k(s, a)} \right) - 1 \\
 &\leq \frac{\eta}{K} \left( \exp \left( 4\eta \frac{\sum_{k'=1}^K N_{h_0\tau, \bar{h}\tau}^{k'}(s, a)}{K} \right) - 1 \right) \\
 &\stackrel{(i)}{\leq} \frac{\eta}{K} (\exp(8\eta\mu_{\text{avg}}(s, a)M\tau + 4\eta\tau) - 1) \\
 &\stackrel{(ii)}{\leq} \frac{64\eta^2\mu_{\text{avg}}(s, a)M\tau}{K},
 \end{aligned}$$

where (i) follows from (149), and (ii) holds because  $e^x \leq 1 + 4x$  for any  $x \in [0, 2]$ ,  $4\eta\tau \leq 1$ , and  $8\eta\mu_{\text{avg}}(s, a)M\tau \leq 1$ .

### C.3.2 PROOF OF LEMMA 18

Recall that

$$\begin{aligned}
 \chi_{u,t}^k(s, a) &= \frac{\eta}{K} (1 - \eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)} \left( \frac{(1 - \eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)}}{\sum_{k'=1}^K (1 - \eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}} - 1 \right) \\
 &\quad \times \prod_{h=\phi(u)}^{\phi(t)-1} \left( \frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{h\tau, (h+1)\tau}^{k'}(s, a)} \right) \\
 &= \left( \frac{(1 - \eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)}}{\sum_{k'=1}^K (1 - \eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}} - 1 \right) \omega_{u,t}^k(s, a).
 \end{aligned}$$

We can observe that  $\chi_{u,t}^k(s, a)$  and  $\omega_{u,t}^k(s, a)$  are solely determined by the number of visits of agents during local steps, i.e.,  $(N_{h\tau, (h+1)\tau}^k(s, a))_{k \in [K], h \in [\phi(t) - LM, \phi(t) - 1]}$ . It thus suffice to consider  $\{\chi_{u,t}^k(s, a; \mathbf{N})\}_{0 \leq u < t, k \in [K]}$  and  $\{\omega_{u,t}^k(s, a; \mathbf{N})\}_{0 \leq u < t, k \in [K]}$  constructed with each of the possible combinations of number of visits for all  $k \in [K]$  and  $h \in [\phi(t) - LM, \phi(t) - 1]$

, i.e.,  $\mathbf{N} \in [0, \tau]^{KLM}$ . Then, by setting  $X = 9\sqrt{\frac{C_{\text{het}}\eta L}{K(1-\gamma)^2} \log \frac{4|S||\mathcal{A}|T^2}{\delta}}$  and taking an union bound,

$$\begin{aligned} & \mathbb{P} \left[ \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} \chi_{u,t}^k(s, a)(P(s, a) - P_{u+1}^k(s, a))V_u^k \right| \geq X \right] \\ &= \sum_{\mathbf{N} \in [0, \tau]^{KLM}} \mathbb{P} \left[ \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} \chi_{u,t}^k(s, a)(P(s, a) - P_{u+1}^k(s, a))V_u^k \right| \geq X, \chi_{u,t}^k(s, a) = \chi_{u,t}^k(s, a; \mathbf{N}) \right] \\ &\leq \sum_{\mathbf{N} \in [0, \tau]^{KLM}} \mathbb{P} \left[ \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} \chi_{u,t}^k(s, a; \mathbf{N})(P(s, a) - P_{u+1}^k(s, a))V_u^k \right| \geq X \right], \end{aligned}$$

and it suffices to show that

$$\mathbb{P} \left[ \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} \chi_{u,t}^k(s, a; \mathbf{N})(P(s, a) - P_{u+1}^k(s, a))V_u^k \right| \geq X \right] \leq \frac{\delta}{|S||\mathcal{A}|T(1+\tau)^{KLM}}.$$

Since  $\chi_{u,t}^k(s, a; \mathbf{N})$  is a constant, which does not depend on  $P_{u+1}^k(s, a)$ ,

$$\mathbb{E}[\chi_{u,t}^k(s, a; \mathbf{N})(P(s, a) - P_{u+1}^k(s, a))V_u^k | \mathcal{Y}_u] = 0, \quad (151)$$

where  $\mathcal{Y}_u$  denotes the history of visited state-action pairs and updated values of all agents until  $u$ , i.e.,  $\mathcal{Y}_u = \{(s_v^k, a_v^k), V_v^k\}_{k \in [K], v \leq u}$ , and thus, we can apply Freedman's inequality to bound the sum.

Before applying Freedman's inequality, we need to calculate the following quantities. First,

$$\begin{aligned} B_t(s, a) &:= \max_{k \in [K], t-LM\tau \leq u < t} |\chi_{u,t}^k(s, a; \mathbf{N})(P(s, a) - P_{u+1}^k(s, a))V_u^k| \\ &\leq \max_{k \in [K], t-LM\tau \leq u < t} \left( \left| 1 - \frac{\frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}}{(1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)}} \right| \right. \\ &\quad \left. \times \omega_{u,t}^k(s, a; \mathbf{N}) \|P(s, a) - P_{u+1}^k(s, a)\|_1 \|V_u^k\|_\infty \right) \\ &\stackrel{(i)}{\leq} \frac{2}{1-\gamma} \max_{k \in [K], t-LM\tau \leq u < t} \left| 1 - \frac{\frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}}{(1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)}} \right| \omega_{u,t}^k(s, a; \mathbf{N}) \\ &\stackrel{(ii)}{\leq} \frac{8\eta\mu_{\max}(s, a)\tau}{1-\gamma} \max_{k \in [K], t-LM\tau \leq u < t} \omega_{u,t}^k(s, a; \mathbf{N}) \stackrel{(iii)}{\leq} \frac{8\eta^2\mu_{\max}(s, a)\tau}{(1-\gamma)K}, \end{aligned}$$

where (i) holds because  $\|P(s, a)\|_1, \|P_u^k(s, a)\|_1 \leq 1, \|V_{u-1}^k\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)), (ii) follows from the fact that (which will be shown at the end of the proof)

$$\left| 1 - \frac{\frac{1}{K} \sum_{k'=1}^K (1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}}{(1-\eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)}} \right| \leq 4\eta\mu_{\max}(s, a)\tau, \quad (152)$$

with  $\mu_{\max}(s, a) := \max_k \mu_b^k(s, a)$ , and (iii) holds due to the fact that  $\omega_{u,t}^k(s, a; \mathbf{N}) \leq \frac{\eta}{K}$ .

Next, we can bound the variance as

$$\begin{aligned}
 W_t(s, a) &:= \sum_{u=\max\{0, t-LM\tau\}}^{t-1} \sum_{k=1}^K \mathbb{E} \left[ \left( \chi_{u,t}^k(s, a; \mathbf{N}) (P(s, a) - P_{u+1}^k(s, a)) V_u^k \right)^2 | \mathcal{Y}_u \right] \\
 &\stackrel{(i)}{\leq} (4\eta \mu_{\max}(s, a) \tau)^2 \sum_{h=\max\{0, \phi(t)-LM\}}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \left( \omega_{u,t}^k(s, a; \mathbf{N}) \right)^2 \text{Var}_{P(s, a)}(V_u^k) \\
 &\stackrel{(ii)}{\leq} \frac{2(4\eta \mu_{\max}(s, a) \tau)^2}{(1-\gamma)^2} \sum_{h=\max\{0, \phi(t)-LM\}}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \left( \omega_{u,t}^k(s, a; \mathbf{N}) \right)^2 \\
 &\stackrel{(iii)}{\leq} \frac{2(4\eta \mu_{\max}(s, a) \tau)^2}{(1-\gamma)^2} \frac{6\eta}{K} =: \sigma^2,
 \end{aligned}$$

where (i) follows from (152), (ii) holds due to the fact that  $\|\text{Var}_P(V)\|_\infty \leq \|P\|_1 (\|V\|_\infty)^2 + (\|P\|_1 \|V\|_\infty)^2 \leq \frac{2}{(1-\gamma)^2}$  because  $\|V\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)) and  $\|P\|_1 \leq 1$ , (iii) follows from (53d) in Lemma 6.

Now, by substituting the above bounds of  $W_t$  and  $B_t$  into Freedman's inequality (see Theorem 12) and setting  $m = 1$ , it follows that for any  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}$ ,  $t \in [T]$  and  $\mathbf{N} = (N_{h\tau, (h+1)\tau}^k(s, a))_{k \in [K], h \in [\phi(t)-LM, \phi(t)-1]} \in [0, \tau]^{KLM}$ ,

$$\begin{aligned}
 &\left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} \chi_{u,t}^k(s, a; \mathbf{N}) (P(s, a) - P_{u+1}^k(s, a)) V_u^k \right| \\
 &\leq \sqrt{8 \max\{W_t(s, a), \frac{\sigma^2}{2m}\} \log \frac{4m|\mathcal{S}||\mathcal{A}|T(1+\tau)^{KLM}}{\delta}} + \frac{4}{3} B_t(s, a) \log \frac{4m|\mathcal{S}||\mathcal{A}|T(1+\tau)^{KLM}}{\delta} \\
 &\leq \sqrt{96 \frac{(4\eta \mu_{\max}(s, a) \tau)^2 \eta}{K(1-\gamma)^2} \log \frac{4|\mathcal{S}||\mathcal{A}|T(1+\tau)^{KLM}}{\delta}} + \frac{12\eta^2 \mu_{\max}(s, a) \tau}{K(1-\gamma)} \log \frac{4|\mathcal{S}||\mathcal{A}|T(1+\tau)^{KLM}}{\delta} \\
 &\leq \sqrt{384 \frac{(4\eta \tau K)(\mu_{\max}(s, a)^2 \eta M \tau) L \eta}{K(1-\gamma)^2} \log \frac{4|\mathcal{S}||\mathcal{A}|T(1+\tau)}{\delta}} \\
 &\quad + \frac{12L\eta(\mu_{\max}(s, a) \eta M \tau)}{(1-\gamma)} \log \frac{4|\mathcal{S}||\mathcal{A}|T(1+\tau)}{\delta} \\
 &\stackrel{(i)}{\leq} \sqrt{48 \frac{C_{\text{het}} L \eta}{K(1-\gamma)^2} \log \frac{4|\mathcal{S}||\mathcal{A}|T(1+\tau)}{\delta}} + \frac{2C_{\text{het}} L \eta}{(1-\gamma)} \log \frac{4|\mathcal{S}||\mathcal{A}|T(1+\tau)}{\delta} \\
 &\stackrel{(ii)}{\leq} 9 \sqrt{\frac{C_{\text{het}} \eta L}{K(1-\gamma)^2} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}} \tag{153}
 \end{aligned}$$

with at least probability  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T(1+\tau)^{KLM}}$ , where we invoke the definition of  $C_{\text{het}}$  (cf. (21)). Here, (i) holds because  $\eta \tau K \leq 1/4$  and  $\mu_{\max}(s, a) \eta M \tau \leq C_{\text{het}} \mu_{\text{avg}}(s, a) \eta M \tau \leq \frac{C_{\text{het}}}{8}$ , and (ii) follows from the fact that  $\eta \leq \frac{1}{128K C_{\text{het}} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}} \leq \frac{1}{K C_{\text{het}} L \log \frac{4|\mathcal{S}||\mathcal{A}|T^2}{\delta}}$ .

**Proof of (152).** Using the fact that for  $0 < \eta < 1$ ,

$(1 - \eta)^{-n} \leq e^{\eta n} \leq 1 + 2\eta n$  if  $n \geq 0$  and  $\eta n \leq 1$ , and  $(1 - \eta)^n \geq 1 - \eta n$  if  $n \leq 0$  or  $n \geq 1$ ,

we can obtain the bounds as follows:

$$\begin{aligned} 1 - \frac{\eta}{K} \sum_{k'=1}^K N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a) &\leq \frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)} \\ &\leq \frac{\frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}}{(1 - \eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)}} \\ &\leq (1 - \eta)^{-N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)} \\ &\leq 1 + 2\eta N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a). \end{aligned}$$

Thus, recalling  $\mu_{\max}(s, a) := \max_k \mu_b^k(s, a)$ , and using the fact that for any  $(s, a, k, u) \in \mathcal{S} \times \mathcal{A} \times [K] \times [T]$ :

$$N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a) \leq 2\mu_{\max}(s, a)\tau$$

at least with probability  $1 - \delta$ , as long as  $\tau \geq 443 \left( \frac{t_{\max}^k}{\mu_{\max}(s, a)} \right) \log \frac{4|\mathcal{S}||\mathcal{A}|TK}{\delta}$ , which naturally holds if  $\tau \geq t_{\text{th}}$  (see (78) for the definition of  $t_{\text{th}}$ ), according to Lemma 13,

$$\begin{aligned} &\left| 1 - \frac{\frac{1}{K} \sum_{k'=1}^K (1 - \eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a)}}{(1 - \eta)^{N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a)}} \right| \\ &\leq 2\eta \max \left\{ N_{\phi(u)\tau, (\phi(u)+1)\tau}^k(s, a), \frac{1}{K} \sum_{k'=1}^K N_{\phi(u)\tau, (\phi(u)+1)\tau}^{k'}(s, a) \right\} \\ &\leq 4\eta \mu_{\max}(s, a)\tau. \end{aligned}$$

#### C.4 Proof of Lemma 8

For any  $t \geq \beta\tau$ , the error term can be decomposed as follows:

$$\begin{aligned} E_t^3(s, a) &= \gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} \omega_{u,t}^k(s, a) P(s, a) (V^* - V_u^k) \\ &= \underbrace{\gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, (\phi(t)-\beta)\tau}^k(s, a)} \omega_{u,t}^k(s, a) P(s, a) (V^* - V_u^k)}_{=: E_t^{3a}(s, a)} \\ &\quad + \underbrace{\gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{(\phi(t)-\beta)\tau, t}^k(s, a)} \omega_{u,t}^k(s, a) P(s, a) (V^* - V_u^k)}_{=: E_t^{3b}(s, a)}. \end{aligned} \tag{154}$$

We shall these two terms separately.

- **Bounding  $E_t^{3a}(s, a)$ .** First, the bound on  $E_t^{3a}(s, a)$  is derived as follows:

$$\begin{aligned}
 |E_t^{3a}(s, a)| &\leq \gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, (\phi(t)-\beta)\tau}^k(s, a)} \omega_{u,t}^k(s, a) \|P(s, a)\|_1 \|V^* - V_u^k\|_\infty \\
 &\stackrel{(i)}{\leq} \frac{2\gamma}{1-\gamma} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, (\phi(t)-\beta)\tau}^k(s, a)} \omega_{u,t}^k(s, a) \\
 &\stackrel{(ii)}{\leq} \frac{2\gamma}{1-\gamma} \exp\left(-\frac{\eta}{2K} \sum_{k=1}^K N_{(\phi(t)-\beta)\tau, t}^k(s, a)\right) \\
 &\stackrel{(iii)}{\leq} \frac{2\gamma}{1-\gamma} \exp\left(-\frac{\eta\mu_{\text{avg}}\beta\tau}{8}\right), \tag{155}
 \end{aligned}$$

where (i) holds because  $\|V_u^k\|_\infty, \|V^*\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)) and  $\|P(s, a)\|_1 \leq 1$ , (ii) holds due to (53c) in Lemma 6, and (iii) follows from the fact that  $\sum_{k=1}^K N_{(\phi(t)-\beta)\tau, t}^k(s, a) \geq \frac{K\mu_{\text{avg}}\beta\tau}{4}$  according to Lemma 15 as long as  $\beta\tau \geq t_{\text{th}}$ .

- **Bounding  $E_t^{3b}(s, a)$ .** Next, we bound  $E_t^{3b}(s, a)$  as follows:

$$\begin{aligned}
 |E_t^{3b}(s, a)| &\leq \gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{(\phi(t)-\beta)\tau, t}^k(s, a)} \omega_{u,t}^k(s, a) \|V^* - V_u^k\|_\infty \\
 &\stackrel{(i)}{\leq} \gamma \sum_{k=1}^K \sum_{h=\phi(t)-\beta}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \omega_{u,t}^k(s, a) (\|\Delta_{h\tau}\|_\infty + \|Q_u^k - Q_{h\tau}^k\|_\infty) \\
 &\stackrel{(ii)}{\leq} \gamma \sum_{k=1}^K \sum_{h=\phi(t)-\beta}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \omega_{u,t}^k(s, a) ((1 + 2\eta\tau)\|\Delta_{h\tau}\|_\infty + \sigma_{\text{local}}) \tag{156}
 \end{aligned}$$

where (i) follows from the following bound, which will be shown in Appendix C.4.1,

$$\|V^* - V_u^k\|_\infty \leq \|\Delta_{\iota(u)}^k\|_\infty + \|Q_u^k - Q_{\iota(u)}^k\|_\infty, \tag{157}$$

and (ii) holds due to the following lemma.

**Lemma 19** Assume  $\eta\tau \leq \frac{1}{2}$ . For any given  $\delta \in (0, 1)$ , the following holds for any  $k \in [K]$  and  $0 \leq u < T$ :

$$\|Q_u^k - Q_{\iota(u)}^k\|_\infty \leq 2\eta\tau \|\Delta_{\iota(u)}^k\|_\infty + \frac{8\gamma\eta\sqrt{\tau-1}}{1-\gamma} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}} \tag{158}$$

with probability at least  $1 - \delta$ .

Here, for notation simplicity, we denote  $\sigma_{\text{local}} := \frac{8\gamma\eta\sqrt{\tau-1}}{1-\gamma} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}}$ .

Then, with some algebraic calculations, we can obtain the bound on  $E_t^{3b}(s, a)$  as follows:

$$\begin{aligned}
 |E_t^{3b}(s, a)| &\stackrel{(i)}{\leq} \sigma_{\text{local}} + \gamma \sum_{h=\phi(t)-\beta}^{\phi(t)-1} (1 + 2\eta\tau) \|\Delta_{h\tau}\|_{\infty} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \omega_{u,t}^k(s, a) \\
 &\stackrel{(ii)}{\leq} \sigma_{\text{local}} + \frac{1+\gamma}{2} \max_{\phi(t)-\beta \leq h < \phi(t)} \|\Delta_{h\tau}\|_{\infty} \sum_{k=1}^K \sum_{h=\phi(t)-\beta}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \omega_{u,t}^k(s, a) \\
 &\stackrel{(iii)}{\leq} \sigma_{\text{local}} + \frac{1+\gamma}{2} \max_{\phi(t)-\beta \leq h < \phi(t)} \|\Delta_{h\tau}\|_{\infty}, \tag{159}
 \end{aligned}$$

where (i) holds according to (53b) of Lemma 6, (ii) holds when  $\eta$  is small enough that  $\eta \leq \frac{1-\gamma}{4\gamma\tau}$ , and (iii) follows from (53b) of Lemma 6.

Now we have the bounds of  $E_t^{3a}(s, a)$  and  $E_t^{3b}(s, a)$  separately obtained above. By combining the bounds in (154), we can claim the advertised bound, which completes the proof.

#### C.4.1 PROOF OF (157)

We prove the claim by showing

$$\Delta_{\iota(u)}^k(s, a_{\iota(u)}^k(s)) - d_{\iota(u), u}^k(s, a^*(s)) \leq V^*(s) - V_u^k(s) \leq \Delta_{\iota(u)}^k(s, a^*(s)) - d_{\iota(u), u}^k(s, a^*(s)) \tag{160}$$

for any  $s \in \mathcal{S}$ . The upper bound is derived as follows:

$$\begin{aligned}
 V^*(s) - V_u^k(s) &= Q^*(s, a^*(s)) - Q_u^k(s, a_u^k(s)) \\
 &\leq Q^*(s, a^*(s)) - Q_u^k(s, a^*(s)) \\
 &= Q^*(s, a^*(s)) - Q_{\iota(u)}^k(s, a^*(s)) - \underbrace{(Q_u^k(s, a^*(s)) - Q_{\iota(u)}^k(s, a^*(s)))}_{d_{\iota(u), u}^k(s, a^*(s))} \tag{161}
 \end{aligned}$$

using the fact that  $Q_u^k(s, a_u^k(s)) \geq Q_u^k(s, a^*(s))$ . Similarly, the lower bound is obtained as follows:

$$\begin{aligned}
 V^*(s) - V_u^k(s) &= Q^*(s, a^*(s)) - Q_u^k(s, a_u^k(s)) \\
 &= Q^*(s, a^*(s)) - Q_{\iota(u)}^k(s, a_{\iota(u)}^k(s)) + Q_{\iota(u)}^k(s, a_{\iota(u)}^k(s)) - Q_u^k(s, a_u^k(s)) \\
 &= Q^*(s, a^*(s)) - Q_{\iota(u)}^k(s, a_{\iota(u)}^k(s)) + Q_{\iota(u)}^k(s, a_{\iota(u)}^k(s)) - Q_{\iota(u)}^k(s, a_u^k(s)) - d_{\iota(u), u}^k(s, a_u^k(s)) \\
 &\geq Q^*(s, a_{\iota(u)}^k(s)) - Q_{\iota(u)}^k(s, a_{\iota(u)}^k(s)) + Q_{\iota(u)}^k(s, a_{\iota(u)}^k(s)) - Q_{\iota(u)}^k(s, a_u^k(s)) - d_{\iota(u), u}^k(s, a_u^k(s)) \\
 &\geq Q^*(s, a_{\iota(u)}^k(s)) - Q_{\iota(u)}^k(s, a_{\iota(u)}^k(s)) - d_{\iota(u), u}^k(s, a_u^k(s)) \tag{162}
 \end{aligned}$$

using the fact that  $Q^*(s, a_{\iota(u)}^k(s)) \leq Q^*(s, a^*(s))$  and  $Q_{\iota(u)}^k(s, a_{\iota(u)}^k(s)) \geq Q_{\iota(u)}^k(s, a_u^k(s))$ .



## C.4.2 PROOF OF LEMMA 19

For any  $0 \leq u < T$ ,  $k \in [K]$ , and  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , we can write the bound as

$$|Q_u^k(s, a) - Q_{\iota(u)}^k(s, a)| \leq \underbrace{2\eta \sum_{v \in \mathcal{U}_{\iota(u), u}^k(s, a)} \|\Delta_v^k\|_\infty}_{:=B_1} + \underbrace{\left| \gamma\eta \sum_{v \in \mathcal{U}_{\iota(u), u}^k(s, a)} (P_{v+1}^k(s, a) - P(s, a))V^\star \right|}_{:=B_2}. \quad (163)$$

The inequality holds by the local update rule:

$$\begin{aligned} Q_{v+1}^k(s, a) - Q_v^k(s, a) &= (1 - \eta)Q_v^k(s, a) + \eta(r(s, a) + \gamma P_{v+1}^k(s, a)V_v^k) - Q_v^k(s, a) \\ &= \eta(r(s, a) + \gamma P_{v+1}^k(s, a)V_v^k - Q_v^k(s, a)) \\ &= \eta(\gamma P_{v+1}^k(s, a)V_v^k - \gamma P(s, a)V^\star + Q^\star(s, a) - Q_v^k(s, a)) \\ &= \gamma\eta P_{v+1}^k(s, a)(V_v^k - V^\star) + \gamma\eta(P_{v+1}^k(s, a) - P(s, a))V^\star + \eta\Delta_v^k(s, a), \end{aligned} \quad (164)$$

and

$$\begin{aligned} |Q_u^k(s, a) - Q_{\iota(u)}^k(s, a)| &\leq \sum_{v \in \mathcal{U}_{\iota(u), u}^k(s, a)} |Q_{v+1}^k(s, a) - Q_v^k(s, a)| \\ &\leq \sum_{v \in \mathcal{U}_{\iota(u), u}^k(s, a)} \left( \eta|\Delta_v^k(s, a)| + \gamma\eta|P_{v+1}^k(s, a)(V_v^k - V^\star)| \right) \\ &\quad + \left| \gamma\eta \sum_{v \in \mathcal{U}_{\iota(u), u}^k(s, a)} (P_{v+1}^k(s, a) - P(s, a))V^\star \right| \\ &\leq \sum_{v \in \mathcal{U}_{\iota(u), u}^k(s, a)} 2\eta\|\Delta_v^k\|_\infty + \left| \gamma\eta \sum_{v \in \mathcal{U}_{\iota(u), u}^k(s, a)} (P_{v+1}^k(s, a) - P(s, a))V^\star \right|, \end{aligned} \quad (165)$$

where the last inequality holds since  $\|P_{v+1}^k(s, a)\|_1 \leq 1$  and  $\|V_v^k - V^\star\|_\infty \leq \|Q_v^k - Q^\star\|_\infty$  (cf. (31)).

Now, we shall bound each term separately.

- **Bounding  $B_1$ .** The local error  $\|\Delta_v^k\|_\infty$  is bounded as follows.

**Lemma 20** Assume  $\eta\tau \leq \frac{1}{2}$ . For any given  $\delta \in (0, 1)$ , the following holds for any  $k \in [K]$  and  $0 \leq u < T$ :

$$\|\Delta_u^k\|_\infty \leq \|\Delta_{\iota(u)}^k\|_\infty + \frac{2\gamma}{1 - \gamma} \sqrt{\eta \log \frac{|\mathcal{S}||\mathcal{A}|TK}{\delta}} \quad (166)$$

with probability at least  $1 - \delta$ .

Then, combining the fact that the number of local updates before the periodic averaging is at most  $\tau - 1$ , we can conclude that

$$\begin{aligned} 2\eta \sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \|\Delta_v^k\|_\infty &\leq 2\eta |\mathcal{U}_{\iota(u),u}^k(s,a)| \max_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \|\Delta_v^k\|_\infty \\ &\leq 2\eta(\tau - 1) \left( \|\Delta_{\iota(u)}^k\|_\infty + \frac{2}{1-\gamma} \sqrt{\eta \log \frac{|\mathcal{S}||\mathcal{A}|TK}{\delta}} \right). \end{aligned} \quad (167)$$

- **Bounding  $B_2$ .** Exploiting the independence of the transitions and applying the Hoeffding inequality and using the fact that  $|\mathcal{U}_{\iota(u),u}^k(s,a)| \leq \tau - 1$ ,  $B_2$  is bounded as follows:

$$\begin{aligned} B_2 &\leq \gamma\eta \sqrt{\sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} |(P_{v+1}^k(s,a) - P(s,a))V^*| \log \frac{|\mathcal{S}||\mathcal{A}|TK}{\delta}} \\ &\leq \frac{2\gamma\eta}{1-\gamma} \sqrt{(\tau - 1) \log \frac{|\mathcal{S}||\mathcal{A}|TK}{\delta}} \end{aligned} \quad (168)$$

for any  $k \in [K]$ ,  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , and  $0 \leq u < T$  with probability at least  $1 - \delta$ , where the last inequality follows from  $\|V^*\|_\infty \leq \frac{1}{1-\gamma}$ ,  $\|P_{v+1}^k(s, a)\|_1$ , and  $\|P(s, a)\|_1 \leq 1$ .

By substituting the bound on  $B_1$  and  $B_2$  into (163) and using the condition that  $\eta\tau < 1$ , we can claim the stated bound holds and this completes the proof.

#### C.4.3 PROOF OF LEMMA 20

For each state-action  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and agent  $k$ , by invoking the recursive relation (49) derived from the local Q-learning update in (23),  $\Delta_u^k$  is decomposed as follows:

$$\begin{aligned} \Delta_u^k(s, a) &= \underbrace{(1 - \eta)^{N_{\iota(u),u}^k(s,a)} \Delta_{\iota(u)}^k(s, a)}_{=:D_1} + \underbrace{\gamma \sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \eta(1 - \eta)^{N_{v+1,u}^k(s,a)} (P(s, a) - P_{v+1}^k(s, a)) V^*}_{=:D_2} \\ &\quad + \underbrace{\gamma \sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \eta(1 - \eta)^{N_{v+1,u}^k(s,a)} P_{v+1}^k(s, a) (V^* - V_v^k)}_{=:D_3}. \end{aligned} \quad (169)$$

Now, we obtain the bound on the three decomposed terms separately.

- **Bounding  $D_1$ .** The term  $D_1$  can be bounded by

$$|D_1| \leq (1 - \eta)^{N_{\iota(u),u}^k(s,a)} \|\Delta_{\iota(u)}^k\|_\infty. \quad (170)$$

- **Bounding  $D_2$ .** By applying the Hoeffding bound using the independence of transitions, the second term is bounded as follows:

$$|D_2| \leq \gamma \sqrt{\sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} (\eta(1 - \eta)^{N_{v+1,u}^k(s,a)})^2 (\|V^*\|_\infty)^2 \log \frac{|\mathcal{S}||\mathcal{A}|TK}{\delta}}$$

$$\leq \frac{\gamma}{1-\gamma} \sqrt{\eta \log \frac{|S||\mathcal{A}|TK}{\delta}} := \rho \quad (171)$$

with probability at least  $1 - \delta$ , where the last inequality holds due to the fact that  $\|V^*\|_\infty \leq \frac{1}{1-\gamma}$  and

$$\sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} (\eta(1-\eta)^{N_{v+1,u}^k(s,a)})^2 \leq \eta^2(1 + (1-\eta)^2 + (1-\eta)^4 + \dots) \leq \eta.$$

See (Li et al., 2021b, Lemma 1) for the detailed explanation of the bound.

• **Bounding  $D_3$ .** Lastly, we bound the third term as follows:

$$\begin{aligned} |D_3| &\leq \gamma \sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \eta(1-\eta)^{N_{v+1,u}^k(s,a)} \|P_{v+1}^k(s,a)\|_1 \|V^* - V_v^k\|_\infty \\ &\leq \gamma \sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \eta(1-\eta)^{N_{v+1,u}^k(s,a)} \|\Delta_v^k\|_\infty, \end{aligned} \quad (172)$$

where the last inequality follows from the fact that  $\|P_{v+1}^k(s,a)\|_1 = 1$  and

$$Q_v^k(s, a^*(s)) - Q^*(s, a^*(s)) \leq V_v^k(s) - V^*(s) \leq Q_v^k(s, a_v^k(s)) - Q^*(s, a_v^k(s))$$

for any  $s \in \mathcal{S}$ , where we denote  $a^*(s) = \arg \max_a Q^*(s, a)$ ,  $a_v^k(s) = \arg \max_a Q_v^k(s, a)$ .

By combining the bounds of the above three terms, we obtain the following recursive relation:

$$|\Delta_u^k(s, a)| \leq (1-\eta)^{N_{\iota(u),u}^k(s,a)} \|\Delta_{\iota(u)}^k\|_\infty + \rho + \gamma \sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \eta(1-\eta)^{N_{v+1,u}^k(s,a)} \|\Delta_v^k\|_\infty. \quad (173)$$

Using the recursive relation, we will prove that the following claim holds for any  $0 \leq m < \tau$  by induction:

$$\|\Delta_{\iota(u)+m}^k\|_\infty \leq \|\Delta_{\iota(u)}^k\|_\infty + 2\rho, \quad (174)$$

which completes the proof of Lemma 20. First, if  $m = 0$ , the claim is obviously true. Suppose the claim holds for  $\iota(u), \iota(u) + 1, \dots, \iota(u) + m - 1$ . Then, for  $u = \iota(u) + m$ , by invoking the recursive relation (173), we can show that the claim (174) holds for  $m$  as follows:

$$\begin{aligned} &|\Delta_{\iota(u)+m}^k(s, a)| \\ &\leq (1-\eta)^{N_{\iota(u),u}^k(s,a)} \|\Delta_{\iota(u)}^k\|_\infty + \rho + \gamma \sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \eta(1-\eta)^{N_{v+1,u}^k(s,a)} (\|\Delta_{\iota(u)}^k\|_\infty + 2\rho) \\ &= ((1-\eta)^{N_{\iota(u),u}^k(s,a)} + \gamma \sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \eta(1-\eta)^{N_{v+1,u}^k(s,a)}) \|\Delta_{\iota(u)}^k\|_\infty \end{aligned}$$

$$\begin{aligned}
 & + (1 + 2\gamma \sum_{v \in \mathcal{U}_{\iota(u),u}^k(s,a)} \eta(1 - \eta)^{N_{v+1,u}^k(s,a)})\rho \\
 & = ((1 - \eta)^{N_{\iota(u),u}^k(s,a)} + \gamma(1 - (1 - \eta)^{N_{\iota(u),u}^k(s,a)}))\|\Delta_{\iota(u)}^k\|_\infty + (1 + 2\gamma(1 - (1 - \eta)^{N_{\iota(u),u}^k(s,a)}))\rho \\
 & \leq \|\Delta_{\iota(u)}^k\|_\infty + 2\rho,
 \end{aligned} \tag{175}$$

where the last inequality holds since

$$(1 - \eta)^{N_{\iota(u),u}^k(s,a)} \geq (1 - \eta)^\tau \geq \left(\frac{1}{4}\right)^{\eta\tau} \geq \frac{1}{2}$$

provided that  $\eta\tau \leq \frac{1}{2}$ .

### C.5 Proof of Lemma 9

First, using the fact that

$$1 \leq (1 - \eta)^{-N_{t-\tau,t}^k(s,a)} \leq e^{\eta\tau} \leq 3$$

given that  $\eta\tau \leq 1$ , by the definition of  $\alpha_t^k$  (cf. (27)), we derive (67a) as follows:

$$\begin{aligned}
 \frac{1}{3K} & \leq \frac{1}{K \max_{k' \in [K]} (1 - \eta)^{-N_{t-\tau,t}^{k'}(s,a)}} \leq \alpha_t^k(s, a) = \frac{(1 - \eta)^{-N_{t-\tau,t}^k(s,a)}}{\sum_{k'=1}^K (1 - \eta)^{-N_{t-\tau,t}^{k'}(s,a)}} \\
 & \leq \frac{(1 - \eta)^{-N_{t-\tau,t}^k(s,a)}}{K} \leq \frac{3}{K}.
 \end{aligned}$$

Moving onto (67b), it follows that

$$\begin{aligned}
 \tilde{\omega}_{0,t}(s, a) & = \prod_{h=0}^{\phi(t)-1} \tilde{\lambda}_{h\tau, (h+1)\tau}(s, a) \\
 & = \prod_{h=0}^{\phi(t)-1} \sum_{k=1}^K \alpha_{(h+1)\tau}^k(s, a) (1 - \eta)^{N_{h\tau, (h+1)\tau}^k(s,a)} \\
 & \stackrel{(i)}{=} \prod_{h=0}^{\phi(t)-1} \frac{K}{\sum_{k=1}^K (1 - \eta)^{-N_{h\tau, (h+1)\tau}^k(s,a)}} \\
 & \stackrel{(ii)}{\leq} \prod_{h=0}^{\phi(t)-1} \frac{1}{(1 - \eta)^{-\frac{1}{K} \sum_{k=1}^K N_{h\tau, (h+1)\tau}^k(s,a)}} \\
 & = (1 - \eta)^{\sum_{h=0}^{\phi(t)-1} \frac{1}{K} \sum_{k=1}^K N_{h\tau, (h+1)\tau}^k(s,a)} = (1 - \eta)^{\frac{1}{K} \sum_{k=1}^K N_{0,t}^k(s,a)},
 \end{aligned}$$

where (i) follows from the definition of  $\alpha_t^k$  (cf. (27)), (ii) follows from Jensen's inequality.

Next, we obtain (67c) through the following derivation:

$$\sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s,a)} \tilde{\omega}_{u,t}^k(s, a) = \sum_{k=1}^K \sum_{h=0}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s,a)} \tilde{\omega}_{u,t}^k(s, a)$$

$$\begin{aligned}
 &= \sum_{k=1}^K \sum_{h=0}^{\phi(t)-1} \alpha_{(h+1)\tau}^k(s, a) \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \eta(1-\eta)^{N_{u+1, (h+1)\tau}^k(s, a)} \left( \prod_{l=h+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right) \\
 &= \sum_{k=1}^K \sum_{h=0}^{\phi(t)-1} \alpha_{(h+1)\tau}^k(s, a) \left( 1 - (1-\eta)^{N_{h\tau, (h+1)\tau}^k(s, a)} \right) \left( \prod_{l=h+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right) \\
 &\stackrel{(i)}{=} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=h+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right) \sum_{k=1}^K \alpha_{(h+1)\tau}^k(s, a) \left( 1 - (1-\eta)^{N_{h\tau, (h+1)\tau}^k(s, a)} \right) \\
 &\stackrel{(ii)}{=} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=h+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right) \left( 1 - \sum_{k=1}^K \alpha_{(h+1)\tau}^k(s, a) (1-\eta)^{N_{h\tau, (h+1)\tau}^k(s, a)} \right) \\
 &= \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=h+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right) \left( 1 - \tilde{\lambda}_{h\tau, (h+1)\tau}(s, a) \right) \\
 &\stackrel{(iii)}{=} 1 - \tilde{\lambda}_{0,\tau}(s, a) \tilde{\lambda}_{\tau, 2\tau}(s, a) \cdots \tilde{\lambda}_{(\phi(t)-1)\tau, t}(s, a) = 1 - \tilde{\omega}_{0,t}(s, a), \tag{176}
 \end{aligned}$$

where (i) follows by reordering the summation, (ii) follows by  $\sum_{k=1}^K \alpha_t^k(s, a) = 1$ , and (iii) holds by cancellation.

In a similar manner, (67d) is derived as follows:

$$\begin{aligned}
 \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, h'\tau}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) &= \sum_{k=1}^K \sum_{h=0}^{h'-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) \\
 &= \sum_{h=0}^{h'-1} \left( \prod_{l=h+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right) \left( 1 - \tilde{\lambda}_{h\tau, (h+1)\tau}(s, a) \right) \\
 &\leq \prod_{l=h'}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \\
 &\leq (1-\eta)^{\frac{1}{K} \sum_{k=1}^K N_{h'\tau, t}^k(s, a)},
 \end{aligned}$$

where the last inequality follows from

$$\prod_{l=h'}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) = \prod_{h=h'}^{\phi(t)-1} \frac{K}{\sum_{k=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^k(s, a)}} \leq \prod_{h=h'}^{\phi(t)-1} \frac{1}{(1-\eta)^{-\frac{1}{K} \sum_{k=1}^K N_{h\tau, (h+1)\tau}^k(s, a)}}$$

due to Jensen's inequality.

Finally, with basic algebraic calculations, (67e) is derived as follows:

$$\sum_{k=1}^K \sum_{u \in \mathcal{U}_{0,t}^k(s, a)} (\tilde{\omega}_{u,t}^k(s, a))^2 = \sum_{k=1}^K \sum_{h=0}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} (\tilde{\omega}_{u,t}^k(s, a))^2$$

$$\begin{aligned}
 &= \sum_{k=1}^K \sum_{h=0}^{\phi(t)-1} (\alpha_{(h+1)\tau}^k(s, a))^2 \left( \prod_{l=h+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right)^2 \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \left( \eta(1-\eta)^{N_{u+1, (h+1)\tau}^k(s, a)} \right)^2 \\
 &\stackrel{(i)}{\leq} 2 \sum_{k=1}^K \sum_{h=0}^{\phi(t)-1} (\alpha_{(h+1)\tau}^k(s, a))^2 \left( \prod_{l=h+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right)^2 \eta \left( 1 - (1-\eta)^{N_{h\tau, (h+1)\tau}^k(s, a)} \right) \\
 &\stackrel{(ii)}{\leq} \frac{6\eta}{K} \sum_{h=0}^{\phi(t)-1} \left( \prod_{l=h+1}^{\phi(t)-1} \tilde{\lambda}_{l\tau, (l+1)\tau}(s, a) \right)^2 \sum_{k=1}^K \alpha_{(h+1)\tau}^k(s, a) \left( 1 - (1-\eta)^{N_{h\tau, (h+1)\tau}^k(s, a)} \right) \\
 &\stackrel{(iii)}{\leq} \frac{6\eta}{K},
 \end{aligned}$$

where (i) holds because

$$\begin{aligned}
 \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} (\eta(1-\eta)^{N_{u+1, (h+1)\tau}^k(s, a)})^2 &= \eta^2 \frac{1 - (1-\eta)^{2(N_{h\tau, (h+1)\tau}^k(s, a))}}{1 - (1-\eta)^2} \\
 &\leq \eta(1 - (1-\eta)^{2(N_{h\tau, (h+1)\tau}^k(s, a))}) \\
 &\leq 2\eta(1 - (1-\eta)^{(N_{h\tau, (h+1)\tau}^k(s, a))}) \quad (177)
 \end{aligned}$$

given that  $2x - x^2 \geq x$  for  $x \leq 1$  and  $(1 - x^2) \leq 2(1 - x)$ , (ii) follows from (67a), and (iii) follows from the same reasoning of (176).

## C.6 Proof of Lemma 10

Without loss of generality, we prove the claim for some fixed  $1 \leq t \leq T$  and  $(s, a) \in \mathcal{S} \times \mathcal{A}$ . For notation simplicity, let

$$\tilde{y}_{u,t}^k(s, a) = \begin{cases} \tilde{\omega}_{u,t}^k(s, a)(P(s, a) - P_{u+1}^k(s, a))V_u^k & \text{if } (s_u^k, a_u^k) = (s, a) \\ 0 & \text{otherwise} \end{cases}, \quad (178)$$

where

$$\tilde{\omega}_{u,t}^k(s, a) = \frac{\eta(1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)}}{K} \prod_{h=\phi(u)}^{\phi(t)-1} \frac{K}{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}}, \quad (179)$$

then  $E_t^2(s, a) = \gamma \sum_{k=1}^K \sum_{u=0}^{t-1} \tilde{y}_{u,t}^k(s, a)$ . However, due to the dependency between  $P_{u+1}^k(s, a)$  and  $\tilde{\omega}_{u,t}^k(s, a)$  arising from the Markovian sampling, it is difficult to track the sum of  $\tilde{y} := \{\tilde{y}_{u,t}^k(s, a)\}$  directly. To address this issue, we will first analyze the sum using a collection of approximate random variables  $\hat{y} = \{\hat{y}_{u,t}^k(s, a)\}$  drawn from a carefully constructed set  $\hat{\mathcal{Y}}$ , which is closely coupled with the target  $\{\tilde{y}_{u,t}^k(s, a)\}_{0 \leq u < t}$ , i.e.,

$$D(\tilde{y}, \hat{y}) := \left| \sum_{k=1}^K \sum_{u=0}^{t-1} (\tilde{y}_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a)) \right| \quad (180)$$

is sufficiently small. In addition,  $\hat{y}$  shall exhibit some useful statistical independence and thus easier to control its sum; we shall control this over the entire set  $\hat{\mathcal{Y}}$ . Finally, leveraging the proximity above, we can obtain the desired bound on  $\tilde{y}$  via triangle inequality. We now provide details on executing this proof outline, where the crust is in designing the set  $\hat{\mathcal{Y}}$  with a controlled size.

Before describing our construction, let's introduce the following useful event:

$$\mathcal{B}_M := \bigcap_{u=0}^{t-M\tau} \left\{ \frac{1}{4} \mu_{\text{avg}}(s, a) K M \tau \leq \sum_{k=1}^K N_{u, u+M\tau}^k(s, a) \leq 2 \mu_{\text{avg}}(s, a) K M \tau \right\}, \quad (181)$$

where  $M = M(s, a) := \lfloor \frac{1}{8\eta\mu_{\text{avg}}(s, a)\tau} \rfloor$ . Note that  $M \geq \frac{1}{16\eta\mu_{\text{avg}}(s, a)\tau}$  since  $\eta\tau \leq 1/16$ . Combining this with the assumption  $\eta \leq \frac{1}{16t_{\text{th}}(s, a)\mu_{\text{avg}}(s, a)}$  (see (78) for the definition of  $t_{\text{th}}(s, a)$ ), it follows that  $M\tau \geq t_{\text{th}}(s, a)$  always holds. Then,  $\mathcal{B}_M$  holds with probability at least  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T}$  according to Lemma 15. The rest of the proof shall be carried out under the event  $\mathcal{B}_M$ .

**Step 1: constructing  $\hat{\mathcal{Y}}$ .** To decouple dependency between  $P_{u+1}^k(s, a)$  and  $\tilde{\omega}_{u,t}^k(s, a)$ , we will introduce approximates of  $\tilde{\omega}_{u,t}^k(s, a)$  that only depend on history until  $u$  by replacing a factor dependent on future with some constant. To gain insight, we factorize  $\tilde{\omega}_{u,t}^k(s, a)$  into two components as follows:

$$\begin{aligned} \tilde{\omega}_{u,t}^k(s, a) &= \prod_{h=h_0(u,t)}^{\phi(u)-1} \left( \frac{K}{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}} \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}}{K} \right) \\ &\quad \times \frac{\eta(1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)}}{K} \prod_{h=\phi(u)}^{\phi(t)-1} \frac{K}{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}} \\ &= \underbrace{\left( \prod_{h=h_0(u,t)}^{\phi(u)-1} \left( \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}}{K} \right) \frac{\eta(1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)}}{K} \right)}_{\text{dependent on history until } u} \\ &\quad \times \underbrace{\left( \prod_{h=h_0(u,t)}^{\phi(t)-1} \frac{K}{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}} \right)}_{\text{dependent on history and future until } t} \\ &= \underbrace{\left( \prod_{h=h_0(u,t)}^{\phi(u)-1} \left( \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}}{K} \right) \frac{\eta(1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)}}{K} \right)}_{:= x_u^k(s, a)} \\ &\quad \times \underbrace{\prod_{l=1}^{l(u,t)} \left( \prod_{h=\max\{0, \phi(t)-lM\}}^{\phi(t)-(l-1)M-1} \frac{K}{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}} \right)}_{:= z_l(s, a)}. \end{aligned} \quad (182)$$

where we denote  $l(u, t) := \lceil \frac{t-u}{M\tau} \rceil$  and  $h_0(u, t) = \max\{0, \phi(t) - l(u, t)M\}$ .

Motivated by the above decomposition, we will construct  $\hat{\mathcal{Y}}$  by approximating the future-dependent parameter  $z_l(s, a)$  for  $1 \leq l \leq L$ , where  $L := \min\{\lceil \frac{t}{M\tau} \rceil, \lceil 64 \log(K/\eta) \rceil\}$ . Using the fact that  $1 + x \leq \exp(x) \leq 1 + 2x$  holds for any  $0 \leq x < 1$ , and  $\eta \frac{\sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a)}{K} \leq \eta\tau \leq 1$ , and applying Jensen's inequality,

$$\begin{aligned} \exp\left(-\eta \frac{\sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a)}{K}\right) &\geq \frac{K}{\sum_{k'=1}^K (1 - \eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}} \\ &\geq \frac{K}{\sum_{k'=1}^K e^{\eta N_{h\tau, (h+1)\tau}^{k'}(s, a)}} \\ &\geq \frac{1}{1 + 2\eta \sum_{k'=1}^K \frac{N_{h\tau, (h+1)\tau}^{k'}(s, a)}{K}} \\ &\geq \exp\left(-2\eta \frac{\sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a)}{K}\right). \end{aligned}$$

Therefore, for  $1 \leq l < L$ , under  $\mathcal{B}_M$ , the range of  $z_l(s, a)$  is bounded as follows:

$$z_l(s, a) \in \left[ \exp(-4\eta\mu_{\text{avg}}(s, a)M\tau), \exp(-\frac{1}{4}\eta\mu_{\text{avg}}(s, a)M\tau) \right].$$

Using this property, we construct a set of values that can cover possible realizations of  $z_l(s, a)$  in a fine-grained manner as follows:

$$\mathcal{Z} := \left\{ \exp\left(-\frac{1}{4}\eta\mu_{\text{avg}}(s, a)M\tau - \frac{i\eta}{K}\right) \mid i \in \mathbb{Z} : 0 \leq i < 4K\mu_{\text{avg}}(s, a)M\tau \right\}. \quad (183)$$

Note that the distance of adjacent elements of  $\mathcal{Z}$  is bounded by  $\eta/K e^{-1/4\eta\mu_{\text{avg}}(s, a)M\tau}$ , and the size of the set is bounded by  $4K\mu_{\text{avg}}(s, a)M\tau$ . For  $l = L$ , because the number of iterations involved in  $z_L(s, a)$  can be less than  $M\tau$ , it follows that  $z_L(s, a) \in [\exp(-4\eta\mu_{\text{avg}}(s, a)M\tau), 1]$ . Hence, we construct the set

$$\mathcal{Z}_0 := \left\{ \exp\left(-\frac{i\eta}{K}\right) \mid i \in \mathbb{Z} : 0 \leq i < 4K\mu_{\text{avg}}(s, a)M\tau \right\}. \quad (184)$$

In sum, we can always find  $(\hat{z}_1, \dots, \hat{z}_l, \dots, \hat{z}_L) \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$  where its entry-wise distance to  $(z_l(s, a))_{l \in [L-1]}$  (resp.  $z_L(s, a)$ ) is at most  $\eta/K e^{-1/4\eta\mu_{\text{avg}}(s, a)M\tau}$  (resp.  $\eta/K$ ).

Moreover, we approximate  $x_u^k(s, a)$  by clipping it when the accumulated number of visits of all agents is not too large as follows:

$$\hat{x}_u^k(s, a) = \begin{cases} x_u^k(s, a) & \text{if } \sum_{k=1}^K N_{h_0(u, t)\tau, \phi(u)\tau}^k(s, a) \leq 2K\mu_{\text{avg}}(s, a)M\tau \\ 0 & \text{otherwise} \end{cases}. \quad (185)$$

Note that the clipping never occurs and  $\hat{x}_u^k(s, a) = x_u^k(s, a)$  for all  $u$  as long as  $\mathcal{B}_M$  holds. To provide useful properties of  $\hat{x}_u^k(s, a)$  that will be useful later, we record the following lemma whose proof is provided in Appendix C.6.1.



**Lemma 21** *For any state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , consider any integers  $1 \leq t \leq T$  and  $1 \leq l \leq \lceil \frac{t}{M\tau} \rceil$ , where  $M = \lfloor \frac{1}{8\eta\mu_{\text{avg}}(s,a)\tau} \rfloor$ . Suppose that  $4\eta\tau \leq 1$ , then  $\hat{x}_u^k(s, a)$  defined in (185) satisfy*

$$\forall u \in [h_0, \phi(t) - (l-1)M] : \hat{x}_u^k(s, a) \leq \frac{9\eta}{K}, \quad (186a)$$

$$\sum_{h=h_0}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) \leq 16\eta\mu_{\text{avg}}(s, a)M\tau, \quad (186b)$$

$$\sum_{h=h_0}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K (\hat{x}_u^k(s, a))^2 \leq \frac{64\eta^2\mu_{\text{avg}}(s, a)M\tau}{K}, \quad (186c)$$

where  $h_0 = \max\{0, \phi(t) - lM\}$ .

Finally, for each  $\mathbf{z} = (\hat{z}_1, \dots, \hat{z}_L) \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$ , setting  $\hat{\omega}_{u,t}^k(s, a; \mathbf{z}) = \hat{x}_u^k(s, a) \prod_{l=1}^{l(u,t)} \hat{z}_l$ , an approximate random sequence  $\hat{\mathbf{y}}_{\mathbf{z}} = \{\hat{y}_{u,t}^k(s, a; \mathbf{z})\}_{0 \leq u < t}$  can be constructed as follows:

$$\hat{y}_{u,t}^k(s, a; \mathbf{z}) = \begin{cases} \hat{\omega}_{u,t}^k(s, a; \mathbf{z})(P(s, a) - P_{u+1}^k(s, a))V_u^k & \text{if } (s_u^k, a_u^k) = (s, a) \text{ and } l(u, t) \leq L \\ 0 & \text{otherwise} \end{cases}. \quad (187)$$

If  $t > LM\tau$ , for any  $u < t - LM\tau$ , i.e.,  $l(u, t) > L$ , we set  $\hat{y}_{u,t}^k(s, a; \mathbf{z}) = 0$  since the magnitude of  $\hat{\omega}_{u,t}^k(s, a)$  becomes negligible when the time difference between  $u$  and  $t$  is large enough, and the fine-grained approximation using  $\mathcal{Z}$  is no longer needed, as shall be seen momentarily. Finally, denote a collection of the approximates induced by  $\mathcal{Z}^{L-1} \times \mathcal{Z}_0$  as

$$\hat{\mathcal{Y}} = \{\hat{\mathbf{y}}_{\mathbf{z}} : \mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0\}.$$

**Step 2: bounding the approximation error  $D(\tilde{\mathbf{y}}, \hat{\mathbf{y}}_{\mathbf{z}})$ .** We now show that under  $\mathcal{B}_M$ , there always exists  $\hat{\mathbf{y}}_{\mathbf{z}} := \hat{\mathbf{y}}_{\mathbf{z}(\tilde{\mathbf{y}})} \in \hat{\mathcal{Y}}$  such that

$$D(\tilde{\mathbf{y}}, \hat{\mathbf{y}}_{\mathbf{z}}) < \frac{129}{1-\gamma} \sqrt{\frac{L\eta}{K}}. \quad (188)$$

To this end, we first decompose the approximation error as follows:

$$\begin{aligned} & \min_{\hat{\mathbf{y}}_{\mathbf{z}} \in \hat{\mathcal{Y}}} D(\tilde{\mathbf{y}}, \hat{\mathbf{y}}_{\mathbf{z}}) \\ &= \min_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=0}^{t-1} \left( \hat{y}_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right) \right| \\ &\leq \underbrace{\max_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=0}^{t-LM\tau-1} \hat{y}_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right|}_{=: D_1} \end{aligned}$$

$$+ \underbrace{\min_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} \tilde{y}_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right|}_{=: D_2}$$

- **Bounding  $D_1$ .** This term appears only when  $t > LM\tau$ . Since  $\hat{y}_{u,t}^k(s, a; \mathbf{z}) = 0$  for all  $u < t - LM\tau$  regardless of  $\mathbf{z}$  by construction,

$$\begin{aligned} & \left| \sum_{k=1}^K \sum_{u=0}^{t-LM\tau-1} \tilde{y}_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right| \\ & \leq \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, t-LM\tau}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) \|P(s, a) - P_{u+1}^k(s, a)\|_1 \|V_u^k\|_\infty \\ & \stackrel{(i)}{\leq} \frac{2}{1-\gamma} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, t-LM\tau}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) \\ & \stackrel{(ii)}{\leq} \frac{2}{1-\gamma} (1-\eta)^{\frac{1}{K} \sum_{k=1}^K N_{t-LM\tau, t}^k(s, a)} \\ & \stackrel{(iii)}{\leq} \frac{2}{1-\gamma} e^{-\eta \frac{1}{4} \mu_{\text{avg}}(s, a) LM\tau} \\ & \stackrel{(iv)}{\leq} \frac{2\eta}{(1-\gamma)K}, \end{aligned}$$

where (i) holds since  $\|P(s, a)\|_1, \|P_u^k(s, a)\|_1 \leq 1$  and  $\|V_{u-1}^k\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)), (ii) follows from (67d) in Lemma 9, (iii) holds due to  $\mathcal{B}_M$ , and (iv) holds because  $L \geq 64 \log \frac{K}{\eta} \geq \frac{4}{\eta \mu_{\text{avg}}(s, a) M\tau} \log \frac{K}{\eta}$  given that  $\eta \mu_{\text{avg}}(s, a) M\tau \geq 1/16$ .

- **Bounding  $D_2$ .** Since  $\hat{x}_u^k(s, a) = x_u^k(s, a)$  when  $\mathcal{B}_M$  holds, in view of (187), we have

$$\begin{aligned} & \min_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} \tilde{y}_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right| \\ & \leq \min_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{t-LM\tau, t}^k(s, a)} |\tilde{\omega}_{u,t}^k(s, a) - \hat{\omega}_{u,t}^k(s, a; \mathbf{z})| \|P(s, a) - P_{u+1}^k(s, a)\|_1 \|V_u^k\|_\infty \\ & \leq \frac{2}{1-\gamma} \min_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left( \sum_{l=1}^L \sum_{h=\phi(t)-lM}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) \left| \prod_{l'=1}^l z_{l'}(s, a) - \prod_{l'=1}^l \hat{z}_{l'} \right| \right), \end{aligned}$$

where the last inequality holds since  $\|P(s, a)\|_1, \|P_u^k(s, a)\|_1 \leq 1$  and  $\|V_{u-1}^k\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)).

Note that for any given  $\{z_l(s, a)\}_{l \in [L]}$ , under  $\mathcal{B}_M$ , there exists  $\hat{\mathbf{z}}^* = (\hat{z}_1^*, \dots, \hat{z}_L^*) \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$  such that  $|\hat{z}_l^* - z_l(s, a)| \leq \frac{\eta}{K} \exp(-1/4 \eta \mu_{\text{avg}}(s, a) M\tau)$  for  $l < L$  and

$|\hat{z}_L^* - z_L(s, a)| \leq \frac{\eta}{K}$ . Also, recall that  $z_l(s, a)$ ,  $\hat{z}_l^* \leq \exp(-1/4\eta\mu_{\text{avg}}(s, a)M\tau)$  for  $l < L$  and  $z_L(s, a)$ ,  $\hat{z}_L^* \leq 1$ . Then, for any  $l \leq L$  it follows that:

$$\begin{aligned} \left| \prod_{l'=1}^l z_{l'}(s, a) - \prod_{l'=1}^l \hat{z}_{l'}^* \right| &\leq \left( \left| \prod_{l'=1}^l z_{l'}(s, a) - \hat{z}_1^* \prod_{l'=2}^l z_{l'}(s, a) \right| + \cdots + \left| z_l \prod_{l'=1}^{l-1} \hat{z}_{l'}^* - \prod_{l'=1}^l \hat{z}_{l'}^* \right| \right) \\ &\leq \exp\left(-\frac{1}{4}(l-1)\eta\mu_{\text{avg}}(s, a)M\tau\right) \sum_{l'=1}^l \frac{\eta}{K} \\ &\leq \exp\left(-\frac{1}{4}(l-1)\eta\mu_{\text{avg}}(s, a)M\tau\right) \frac{L\eta}{K}. \end{aligned}$$

Then, applying the above bound and (186b) in Lemma 21,

$$\begin{aligned} &\min_{\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0} \left| \sum_{k=1}^K \sum_{u=t-LM\tau}^{t-1} \tilde{y}_{u,t}^k(s, a) - \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right| \\ &\leq \frac{2}{1-\gamma} \sum_{l=1}^L \sum_{h=\phi(t)-(l-1)M-1}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) \left| \prod_{l'=1}^l z_{l'}(s, a) - \prod_{l'=1}^l \hat{z}_{l'}^* \right| \\ &\leq \frac{2}{1-\gamma} \frac{L\eta}{K} \sum_{l=1}^L \exp\left(-\frac{1}{4}(l-1)\eta\mu_{\text{avg}}(s, a)M\tau\right) \sum_{h=\phi(t)-(l-1)M-1}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) \\ &\leq \frac{2}{1-\gamma} \frac{L\eta}{K} \frac{1}{1 - \exp(-1/4\eta\mu_{\text{avg}}(s, a)M\tau)} (16\eta\mu_{\text{avg}}(s, a)M\tau) \\ &\stackrel{(i)}{\leq} \frac{2}{1-\gamma} \frac{L\eta}{K} \frac{8}{\eta\mu_{\text{avg}}(s, a)M\tau} 16\eta\mu_{\text{avg}}(s, a)M\tau \leq \frac{256L\eta}{(1-\gamma)K}, \end{aligned}$$

where (i) holds since  $1/4\eta\mu_{\text{avg}}(s, a)M\tau \leq 1$  and  $e^{-x} \leq 1 - \frac{1}{2}x$  for any  $0 \leq x \leq 1$ .

By combining the bounds obtained above and using the fact that  $\frac{4\eta L}{K} \leq 1$  and  $L \leq 64 \log(TK)$ , we can conclude that

$$\min_{\hat{\mathbf{y}}_{\mathbf{z}} \in \hat{\mathcal{Y}}} D(\tilde{\mathbf{y}}, \hat{\mathbf{y}}_{\mathbf{z}}) \leq \frac{2\eta}{(1-\gamma)K} + \frac{256L\eta}{(1-\gamma)K} \leq \frac{129}{1-\gamma} \sqrt{\frac{L\eta}{K}}.$$

**Step 3: concentration bound over  $\mathcal{Y}$ .** We now show that for all elements in  $\hat{\mathcal{Y}} = \{\hat{\mathbf{y}}_{\mathbf{z}} : \mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0\}$  satisfy

$$\left| \sum_{k=1}^K \sum_{u=0}^{t-1} \hat{y}_{u,t}^k(s, a; \mathbf{z}) \right| < \frac{624}{(1-\gamma)} \sqrt{\frac{\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2K}{\delta}} \quad (189)$$

with probability at least  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T}$ . It suffices to establish (189) for a fixed  $\mathbf{z} \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$  with probability at least  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T|\mathcal{Y}|}$ , where

$$|\hat{\mathcal{Y}}| = |\mathcal{Z}^{L-1} \times \mathcal{Z}_0| \leq (4K\mu_{\text{avg}}(s, a)M\tau)^L \leq (K/\eta)^L \leq (TK)^L. \quad (190)$$

For any fixed  $\mathbf{z} = (\hat{z}_1, \dots, \hat{z}_L) \in \mathcal{Z}^{L-1} \times \mathcal{Z}_0$ , since  $\hat{\omega}_{u,t}^k(s, a; \mathbf{z}) = \hat{x}_u^k(s, a) \prod_{l=1}^{l(u,t)} \hat{z}_l$  only depends on the events happened until  $u$ , which is independent to a transition at  $u+1$ . Thus, we can apply Freedman's inequality to bound the sum of  $\hat{y}_{u,t}^k(s, a; \mathbf{z})$  since

$$\mathbb{E}[\hat{y}_{u,t}^k(s, a; \mathbf{z}) | \mathcal{Y}_u] = 0, \quad (191)$$

where  $\mathcal{Y}_u$  denotes the history of visited state-action pairs and updated values of all agents until  $u$ , i.e.,  $\mathcal{Y}_u = \{(s_v^k, a_v^k), V_v^k\}_{k \in [K], v \leq u}$ . Before applying Freedman's inequality, we need to calculate the following quantities. First,

$$\begin{aligned} B_t(s, a) &:= \max_{k \in [K], 0 \leq u < t} |\hat{y}_{u,t}^k(s, a; \mathbf{z})| \leq \hat{x}_u^k(s, a) \prod_{l=1}^{l(u,t)} \hat{z}_l \|P(s, a) - P_{u+1}^k(s, a)\|_1 \|V_u^k\|_\infty \\ &\leq \frac{18\eta}{(1-\gamma)K}, \end{aligned} \quad (192)$$

where the last inequality follows from  $\|P(s, a)\|_1, \|P_u^k(s, a)\|_1 \leq 1, \|V_{u-1}^k\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)),  $\hat{z}_l \leq 1$ , and (186a) in Lemma 21. Next, we can bound the variance as

$$\begin{aligned} W_t(s, a) &:= \sum_{u=t-LM\tau}^{t-1} \sum_{k=1}^K \mathbb{E}[(\hat{y}_{u,t}^k(s, a; \mathbf{z}))^2 | \mathcal{Y}_u] \\ &= \sum_{l=1}^L \sum_{h=\max\{0, \phi(t)-lM\}}^{\phi(t)-(l-1)M-1} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} (\hat{x}_u^k(s, a) \prod_{l'=1}^l \hat{z}_{l'})^2 \text{Var}_{P(s, a)}(V_u^k) \\ &\stackrel{(i)}{\leq} \frac{2}{(1-\gamma)^2} \sum_{l=1}^L \left( \prod_{l'=1}^l \hat{z}_{l'}^2 \right) \sum_{h=\max\{0, \phi(t)-lM\}}^{\phi(t)-(l-1)M-1} \sum_{k=1}^K \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} (\hat{x}_u^k(s, a))^2 \\ &\stackrel{(ii)}{\leq} \frac{2}{(1-\gamma)^2} \sum_{l=1}^L \left( \prod_{l'=1}^l \hat{z}_{l'}^2 \right) \frac{64\eta^2 \mu_{\text{avg}}(s, a) M\tau}{K} \\ &\stackrel{(iii)}{\leq} \frac{128\eta^2 \mu_{\text{avg}}(s, a) M\tau}{K(1-\gamma)^2} \sum_{l=1}^L \exp(-1/2(l-1)\eta \mu_{\text{avg}}(s, a) M\tau) \\ &\leq \frac{128\eta^2 \mu_{\text{avg}}(s, a) M\tau}{K(1-\gamma)^2} \frac{1}{1 - \exp(-1/2\eta \mu_{\text{avg}}(s, a) M\tau)} \\ &\stackrel{(iv)}{\leq} \frac{128\eta^2 \mu_{\text{avg}}(s, a) M\tau}{K(1-\gamma)^2} \frac{4}{\eta \mu_{\text{avg}}(s, a) M\tau} = \frac{512\eta}{K(1-\gamma)^2} := \sigma^2, \end{aligned} \quad (193)$$

where (i) holds due to the fact that  $\|\text{Var}_P(V)\|_\infty \leq \|P\|_1 (\|V\|_\infty)^2 + (\|P\|_1 \|V\|_\infty)^2 \leq \frac{2}{(1-\gamma)^2}$  because  $\|V\|_\infty \leq \frac{1}{1-\gamma}$  (cf. (30)) and  $\|P\|_1 \leq 1$ , (ii) follows from (186c) in Lemma 21, (iii) holds due to the range of  $\mathcal{Z}$  and  $\mathcal{Z}_0$  is bounded by  $\exp(-1/4\eta \mu_{\text{avg}}(s, a) M\tau)$  and 1, respectively, and (iv) holds since  $e^{-x} \leq 1 - \frac{1}{2}x$  for any  $0 \leq x \leq 1$  and  $1/2\eta \mu_{\text{avg}}(s, a) M\tau \leq 1$ .

Now, by substituting the above bounds of  $W_t$  and  $B_t$  into Freedman's inequality (see Theorem 12) and setting  $m = 1$ , it follows that for any  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}$ ,  $t \in [T]$  and  $\hat{y}_z \in \hat{\mathcal{Y}}$ ,

$$\begin{aligned}
 & \left| \sum_{k=1}^K \sum_{u=0}^{t-1} \hat{y}_{u,t}^k(s, a; z) \right| \\
 & \leq \sqrt{8 \max \{W_t(s, a), \frac{\sigma^2}{2^m}\} \log \frac{4m|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}{\delta}} + \frac{4}{3} B_t(s, a) \log \frac{4m|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}{\delta} \\
 & \leq \sqrt{4096 \frac{\eta}{K(1-\gamma)^2} \log \frac{4|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}{\delta}} + \frac{24\eta}{K(1-\gamma)} \log \frac{4|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}{\delta} \\
 & \stackrel{(i)}{\leq} \frac{78}{(1-\gamma)} \sqrt{\frac{\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta}}, \tag{194}
 \end{aligned}$$

with at least probability  $1 - \frac{\delta}{|\mathcal{S}||\mathcal{A}|T|\hat{\mathcal{Y}}|}$ , where (i) holds because  $|\hat{\mathcal{Y}}| \leq (TK)^L$  given that  $\eta \mu_{\text{avg}}(s, a) M \tau \leq 1/4$ , and  $\frac{4\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta} \leq 1$ . Therefore, it follows that (189) holds.

**Step 4: putting things together.** We now putting all the results obtained in the previous steps together to achieve the claimed bound. Under  $\mathcal{B}_M$ , there always exists  $\hat{y}_z := \hat{y}_{z(\tilde{y})} \in \hat{\mathcal{Y}}$  such that (188) holds. Hence, setting  $q = \frac{2064}{(1-\gamma)} \sqrt{\frac{\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta}}$ ,

$$\begin{aligned}
 \sum_{k=1}^K \sum_{u=0}^{t-1} \tilde{y}_{u,t}^k(s, a) & \leq \left| \sum_{k=1}^K \sum_{u=0}^{t-1} \hat{y}_{u,t}^k(s, a; z) \right| + D(\tilde{y}, \hat{y}_z) \\
 & \leq \frac{78}{(1-\gamma)} \sqrt{\frac{\eta L}{K} \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta}} + \frac{129}{1-\gamma} \sqrt{\frac{L\eta}{K}} \\
 & \leq \frac{2064}{(1-\gamma)} \sqrt{\frac{\eta}{K} \log(TK) \log \frac{4|\mathcal{S}||\mathcal{A}|T^2 K}{\delta}}, \tag{195}
 \end{aligned}$$

where the second line holds due to (189) and (188), and the last line holds due to  $L \leq 64 \log(TK)$ . By taking a union bound over all  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $t \in [T]$ , we complete the proof.

#### C.6.1 PROOF OF LEMMA 21

For notational simplicity, let  $\bar{h}$  be the largest integer among  $h \in (h_0, \phi(t) - (l-1)M)$  such that

$$\sum_{k=1}^K N_{h_0\tau, (h-1)\tau}^k(s, a) \leq 2K \mu_{\text{avg}}(s, a) M \tau. \tag{196}$$

Then, the following holds:

$$\sum_{k=1}^K N_{h_0\tau, \bar{h}\tau}^k(s, a) = \sum_{k=1}^K N_{(\bar{h}-1)\tau, \bar{h}\tau}^k(s, a) + \sum_{k=1}^K N_{h_0\tau, (\bar{h}-1)\tau}^k(s, a)$$

$$\leq K\tau + 2K\mu_{\text{avg}}(s, a)M\tau. \quad (197)$$

Also, for the following proofs, we provide an useful bound as follows:

$$\begin{aligned} \sum_{k'=1}^K \frac{(1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}}{K} &\leq \frac{\sum_{k'=1}^K e^{\eta N_{h\tau, (h+1)\tau}^{k'}(s, a)}}{K} \leq 1 + 2\eta \frac{\sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a)}{K} \quad (198) \\ &\leq \exp\left(2\eta \frac{\sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s, a)}{K}\right), \end{aligned}$$

which holds since  $1 + x \leq e^x \leq 1 + 2x$  for any  $x \in [0, 1]$  and  $\eta N_{h\tau, (h+1)\tau}^{k'}(s, a) \leq \eta\tau \leq 1$ .

According to (185), for any integer  $u \in [\bar{h}\tau, t - (l-1)M\tau]$ ,  $\hat{x}_u^k(s, a)$  is clipped to zero. Now, we prove the bounds in Lemma 21 respectively.

**Proof of (186a).** For  $u \in [h_0\tau, \bar{h}\tau]$ ,

$$\begin{aligned} \hat{x}_u^k(s, a) &= \prod_{h=h_0}^{\phi(u)-1} \left( \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}}{K} \right) \frac{\eta(1-\eta)^{-N_{\phi(u)\tau, u+1}^k(s, a)}}{K} \\ &\stackrel{(i)}{\leq} \prod_{h=h_0}^{\phi(u)-1} \left( \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h\tau, (h+1)\tau}^{k'}(s, a)}}{K} \right) \frac{3\eta}{K} \\ &\stackrel{(ii)}{\leq} \exp\left(\frac{2\eta}{K} \sum_{k'=1}^K N_{h_0\tau, (\bar{h}-1)\tau}^{k'}(s, a)\right) \frac{3\eta}{K} \\ &\stackrel{(iii)}{\leq} \exp(4\eta\mu_{\text{avg}}(s, a)M\tau) \frac{3\eta}{K} \stackrel{(iv)}{\leq} \frac{9\eta}{K}, \quad (199) \end{aligned}$$

where (i) holds since  $(1+\eta)^x \leq e^{\eta x}$  and  $\eta N_{\phi(u)\tau, u+1}^k(s, a) \leq \eta\tau \leq 1$ , (ii) holds due to (198) and the fact that  $\phi(u) \leq \bar{h} - 1$ , (iii) follows from the definition of  $\bar{h}$  in (196), and (iv) holds because  $4\eta\mu_{\text{avg}}(s, a)M\tau \leq 1$ .

**Proof of (186b).** By the definition of  $\bar{h}$ , it follows that

$$\sum_{h=h_0}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K \hat{x}_u^k(s, a) = \sum_{h=h_0}^{\bar{h}-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K x_u^k(s, a).$$

Using the following relation for each  $h$ :

$$\begin{aligned} &\sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K x_u^k(s, a) \\ &= \left( \prod_{h'=h_0}^{h-1} \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h'\tau, (h'+1)\tau}^{k'}(s, a)}}{K} \right) \sum_{k=1}^K \frac{\sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \eta(1-\eta)^{-N_{h\tau, u+1}^k(s, a)}}{K} \end{aligned}$$

$$\begin{aligned}
 &= \left( \prod_{h'=h_0}^{h-1} \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h'\tau, (h'+1)\tau}^{k'}(s,a)}}{K} \right) \sum_{k=1}^K \frac{(1-\eta)^{-N_{h\tau, (h+1)\tau}^k(s,a)} - 1}{K} \\
 &= \left( \prod_{h'=h_0}^h \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h'\tau, (h'+1)\tau}^{k'}(s,a)}}{K} \right) - \left( \prod_{h'=h_0}^{h-1} \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h'\tau, (h'+1)\tau}^{k'}(s,a)}}{K} \right),
 \end{aligned}$$

and applying (198), we can complete the proof as follows:

$$\begin{aligned}
 \sum_{h=h_0}^{\bar{h}-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s,a)} \sum_{k=1}^K x_u^k(s,a) &\leq \prod_{h'=h_0}^{\bar{h}-1} \exp \left( \frac{2\eta \sum_{k'=1}^K N_{h'\tau, (h'+1)\tau}^{k'}(s,a)}{K} \right) - 1 \\
 &\leq \exp \left( \frac{2\eta \sum_{k'=1}^K N_{h_0\tau, \bar{h}\tau}^{k'}(s,a)}{K} \right) - 1 \\
 &\stackrel{(i)}{\leq} \exp(4\eta\mu_{\text{avg}}(s,a)M\tau + 2\eta\tau) - 1 \\
 &\stackrel{(ii)}{\leq} 16\eta\mu_{\text{avg}}(s,a)M\tau,
 \end{aligned}$$

where (i) follows from (197), and (ii) holds because  $e^x \leq 1 + 2x$  for any  $x \in [0, 1]$  and  $2\eta\tau \leq 4\eta\mu_{\text{avg}}(s,a)M\tau \leq 1/2$ .

**Proof of (186c).** Similarly,

$$\sum_{h=h_0}^{\phi(t)-(l-1)M-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s,a)} \sum_{k=1}^K (\hat{x}_u^k(s,a))^2 = \sum_{h=h_0}^{\bar{h}-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s,a)} \sum_{k=1}^K (x_u^k(s,a))^2.$$

Using the following relation for each  $h$ :

$$\begin{aligned}
 &\sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s,a)} \sum_{k=1}^K (x_u^k(s,a))^2 \\
 &= \left( \prod_{h'=h_0}^{h-1} \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h'\tau, (h'+1)\tau}^{k'}(s,a)}}{K} \right)^2 \sum_{k=1}^K \frac{\sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s,a)} \eta^2 (1-\eta)^{-2N_{h\tau, (h+1)\tau}^k(s,a)}}{K^2} \\
 &\leq \left( \prod_{h'=h_0}^{h-1} \frac{\sum_{k'=1}^K (1-\eta)^{-N_{h'\tau, (h'+1)\tau}^{k'}(s,a)}}{K} \right)^2 \sum_{k=1}^K \frac{\eta((1-\eta)^{-2N_{h\tau, (h+1)\tau}^k(s,a)} - 1)}{K^2} \\
 &\leq \frac{\eta}{K} \left( \prod_{h'=h_0}^{h-1} \exp \left( 2\eta \frac{\sum_{k'=1}^K N_{h'\tau, (h'+1)\tau}^{k'}(s,a)}{K} \right) \right)^2 \left( \exp \left( 4\eta \frac{\sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s,a)}{K} \right) - 1 \right) \\
 &= \frac{\eta}{K} \exp \left( 4\eta \frac{\sum_{k'=1}^K N_{h_0\tau, h\tau}^{k'}(s,a)}{K} \right) \left( \exp \left( 4\eta \frac{\sum_{k'=1}^K N_{h\tau, (h+1)\tau}^{k'}(s,a)}{K} \right) - 1 \right)
 \end{aligned}$$

$$= \frac{\eta}{K} \left( \exp \left( 4\eta \frac{\sum_{k'=1}^K N_{h_0\tau, (h+1)\tau}^{k'}(s, a)}{K} \right) - \exp \left( 4\eta \frac{\sum_{k'=1}^K N_{h_0\tau, h\tau}^{k'}(s, a)}{K} \right) \right), \quad (200)$$

where the inequality is derived similarly to (198) under the condition  $2\eta\tau \leq 1$ , we can complete the proof as follows:

$$\begin{aligned} \sum_{h=h_0}^{\bar{h}-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \sum_{k=1}^K (x_u^k(s, a))^2 &\leq \frac{\eta}{K} \left( \exp \left( 4\eta \frac{\sum_{k'=1}^K N_{h_0\tau, \bar{h}\tau}^{k'}(s, a)}{K} \right) - 1 \right) \\ &\stackrel{(i)}{\leq} \frac{\eta}{K} (\exp(8\eta\mu_{\text{avg}}(s, a)M\tau + 4\eta\tau) - 1) \\ &\stackrel{(ii)}{\leq} \frac{64\eta^2\mu_{\text{avg}}(s, a)M\tau}{K}, \end{aligned} \quad (201)$$

where (i) follows from (197), and (ii) holds because  $e^x \leq 1 + 4x$  for any  $x \in [0, 2]$  and  $4\eta\tau \leq 8\eta\mu_{\text{avg}}(s, a)M\tau \leq 1$ .

### C.7 Proof of Lemma 11

The proof follows a similar structure to that of Lemma 8. We omit common parts of the proofs and refer to Appendix C.4 to check the detailed derivations. First, we decompose the error term as follows:

$$\begin{aligned} E_t^3(s, a) &= \underbrace{\gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, (\phi(t)-\beta)\tau}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) P(s, a) (V^* - V_u^k)}_{=: E_t^{3a}(s, a)} \\ &\quad + \underbrace{\gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{(\phi(t)-\beta)\tau, t}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) P(s, a) (V^* - V_u^k)}_{=: E_t^{3b}(s, a)}. \end{aligned} \quad (202)$$

We shall bound these two terms separately.

- **Bounding  $E_t^{3a}(s, a)$ .** First, the bound of  $E_t^{3a}(s, a)$  is derived as follows:

$$\begin{aligned} |E_t^{3a}(s, a)| &\leq \gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{0, (\phi(t)-\beta)\tau}^k(s, a)} \tilde{\omega}_{u,t}^k(s, a) \|P(s, a)\|_1 \|V^* - V_u^k\|_\infty \\ &\stackrel{(i)}{\leq} \frac{2}{1-\gamma} (1-\eta)^{\frac{1}{K} \sum_{k=1}^K N_{(\phi(t)-\beta)\tau, t}^k(s, a)} \\ &\stackrel{(ii)}{\leq} \frac{2}{1-\gamma} (1-\eta)^{\frac{\mu_{\text{avg}}\beta\tau}{4}}, \end{aligned} \quad (203)$$



where (i) holds due to Lemma 9 (cf. (67d)), and (ii) follows from applying Lemma 15 that with probability at least  $1 - \delta$ ,

$$\sum_{k=1}^K N_{(\phi(t)-\beta)\tau, t}^k(s, a) \geq \frac{K\beta\tau\mu_{\text{avg}}}{4}$$

holds for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$  and  $0 \leq u < v \leq T$  as long as  $\beta\tau \geq t_{\text{th}}$ .

- **Bounding  $E_t^{3b}(s, a)$ .** Combining (157) and Lemma 19 to bound  $\|V^* - V_u^k\|_\infty$ , we bound  $E_t^{3b}(s, a)$  as follows:

$$\begin{aligned} |E_t^{3b}(s, a)| &\leq \gamma \sum_{k=1}^K \sum_{u \in \mathcal{U}_{(\phi(t)-\beta)\tau, t}^k(s, a)} \tilde{\omega}_{u, t}^k(s, a) \|V^* - V_u^k\|_\infty \\ &\leq \gamma \sum_{k=1}^K \sum_{h=\phi(t)-\beta}^{\phi(t)-1} \sum_{u \in \mathcal{U}_{h\tau, (h+1)\tau}^k(s, a)} \tilde{\omega}_{u, t}^k(s, a) ((1 + 2\eta\tau)\|\Delta_{h\tau}\|_\infty + \sigma_{\text{local}}) \\ &\leq \sigma_{\text{local}} + \frac{1 + \gamma}{2} \max_{\phi(t)-\beta \leq h < \phi(t)} \|\Delta_{h\tau}\|_\infty \end{aligned} \quad (204)$$

where we denote  $\sigma_{\text{local}} := \frac{8\gamma\eta\sqrt{\tau-1}}{1-\gamma} \sqrt{\log \frac{2|\mathcal{S}||\mathcal{A}|TK}{\delta}}$  for notational simplicity, and the last inequality follows from Lemma 9 (cf. (67c)) and the assumption that  $\eta \leq \frac{1-\gamma}{4\gamma\tau}$ .

Now we have the bounds of  $E_t^{3a}(s, a)$  and  $E_t^{3b}(s, a)$  separately obtained above. By combining the bounds in (202), we can claim the advertised bound, which completes the proof.

## References

- M. Assran, J. Romoff, N. Ballas, J. Pineau, and M. Rabbat. Gossip-based actor-learner architectures for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Y. Bai, T. Xie, N. Jiang, and Y.-X. Wang. Provably efficient Q-learning with low switching cost. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- C. L. Beck and R. Srikant. Error bounds for constant step-size Q-learning. *Systems & control letters*, 61(12):1203–1208, 2012.
- K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konečný, S. Mazzocchi, B. McMahan, et al. Towards federated learning at scale: System design. In *Proceedings of Machine Learning and Systems*, pages 374–388, 2019.
- V. S. Borkar and S. P. Meyn. The ODE method for convergence of stochastic approximation and reinforcement learning. *SIAM Journal on Control and Optimization*, 38(2):447–469, 2000.

- T. Chen, K. Zhang, G. B. Giannakis, and T. Başar. Communication-efficient policy gradient methods for distributed reinforcement learning. *IEEE Transactions on Control of Network Systems*, 9(2):917–929, 2021a.
- Z. Chen, S. T. Maguluri, S. Shakkottai, and K. Shanmugam. Finite-sample analysis of contractive stochastic approximation using smooth convex envelopes. In *Advances in Neural Information Processing Systems*, volume 33, pages 8223–8234, 2020.
- Z. Chen, S. T. Maguluri, S. Shakkottai, and K. Shanmugam. A Lyapunov theory for finite-sample guarantees of asynchronous Q-learning and TD-learning variants. *arXiv preprint arXiv:2102.01567*, 2021b.
- Z. Chen, Y. Zhou, and R. Chen. Multi-agent off-policy TDC with near-optimal sample and communication complexities. *Transactions on Machine Learning Research*, 2022a.
- Z. Chen, Y. Zhou, R.-R. Chen, and S. Zou. Sample and communication-efficient decentralized actor-critic algorithms with finite-time analysis. In *International Conference on Machine Learning*, volume 162, pages 3794–3834. PMLR, 2022b.
- T. Doan, S. Maguluri, and J. Romberg. Finite-time analysis of distributed TD(0) with linear function approximation on multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 1626–1635, 2019.
- T. T. Doan, S. T. Maguluri, and J. Romberg. Finite-time performance of distributed temporal-difference learning with linear function approximation. *SIAM Journal on Mathematics of Data Science*, 3(1):298–320, 2021.
- L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley, I. Dunning, S. Legg, and K. Kavukcuoglu. IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures. In *International Conference on Machine Learning*, pages 1406–1415, 2018.
- E. Even-Dar and Y. Mansour. Learning rates for Q-learning. *Journal of machine learning Research*, 5(Dec):1–25, 2003.
- X. Fan, Y. Ma, Z. Dai, W. Jing, C. Tan, and B. K. H. Low. Fault-tolerant federated reinforcement learning with theoretical guarantee. In *Advances in Neural Information Processing Systems*, volume 34, pages 1007–1021, 2021.
- D. A. Freedman. On tail probabilities for martingales. *The Annals of Probability*, 3(1):100–118, 1975.
- T. Jaakkola, M. I. Jordan, and S. P. Singh. Convergence of stochastic iterative dynamic programming algorithms. In *Advances in Neural Information Processing Systems*, pages 703–710, 1994.
- C. Jin, Z. Allen-Zhu, S. Bubeck, and M. I. Jordan. Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4863–4873, 2018.

- H. Jin, Y. Peng, W. Yang, S. Wang, and Z. Zhang. Federated reinforcement learning with environment heterogeneity. In *International Conference on Artificial Intelligence and Statistics*, pages 18–37, 2022.
- P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.
- M. J. Kearns and S. P. Singh. Finite-sample convergence rates for Q-learning and indirect algorithms. In *Advances in Neural Information Processing Systems*, pages 996–1002, 1999.
- S. Khodadadian, P. Sharma, G. Joshi, and S. T. Maguluri. Federated reinforcement learning: Linear speedup under Markovian sampling. In *International Conference on Machine Learning*, pages 10997–11057, 2022.
- G. Li, L. Shi, Y. Chen, and Y. Chi. Breaking the sample complexity barrier to regret-optimal model-free reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 34, 2021a.
- G. Li, Y. Wei, Y. Chi, Y. Gu, and Y. Chen. Sample complexity of asynchronous Q-learning: Sharper analysis and variance reduction. *IEEE Transactions on Information Theory*, 68(1):448–473, 2021b.
- G. Li, C. Cai, Y. Chen, Y. Wei, and Y. Chi. Is Q-learning minimax optimal? a tight sample complexity analysis. *Operations Research*, 2023.
- H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *International Conference on Artificial Intelligence and Statistics*, 2017.
- V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, volume 48, pages 1928–1937. PMLR, 2016.
- D. Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20, 2015.
- M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- G. Qu and A. Wierman. Finite-time analysis of asynchronous stochastic approximation and Q-learning. In *Conference on Learning Theory*, pages 3185–3205. PMLR, 2020.
- S. Salgia and Y. Chi. The sample-communication complexity trade-off in federated Q-learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- H. Shen, K. Zhang, M. Hong, and T. Chen. Towards understanding asynchronous advantage actor-critic: convergence and linear speedup. *arXiv preprint arXiv:2012.15511*, 2022.

- L. Shi, G. Li, Y. Wei, Y. Chen, and Y. Chi. Pessimistic Q-learning for offline reinforcement learning: Towards optimal sample complexity. In *International Conference on Machine Learning*, volume 162, pages 19967–20025. PMLR, 2022.
- J. Sun, G. Wang, G. B. Giannakis, Q. Yang, and Z. Yang. Finite-time analysis of decentralized temporal-difference learning with linear function approximation. In *International Conference on Artificial Intelligence and Statistics*, pages 4485–4495. PMLR, 2020.
- R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- C. Szepesvári. The asymptotic convergence-rate of Q-learning. In *Advances in Neural Information Processing Systems*, pages 1064–1070, 1998.
- J. N. Tsitsiklis. Asynchronous stochastic approximation and Q-learning. *Machine learning*, 16(3):185–202, 1994.
- H.-T. Wai. On the convergence of consensus algorithms with Markovian noise and gradient bias. In *Conference on Decision and Control*, pages 4897–4902. IEEE, 2020.
- M. J. Wainwright. Stochastic approximation with cone-contractive operators: Sharp  $\ell_\infty$ -bounds for Q-learning. *arXiv preprint arXiv:1905.06265*, 2019a.
- M. J. Wainwright. Variance-reduced Q-learning is minimax optimal. *arXiv preprint arXiv:1906.04697*, 2019b.
- G. Wang, S. Lu, G. Giannakis, G. Tesauro, and J. Sun. Decentralized TD tracking with linear function approximation and its finite-time analysis. In *Advances in Neural Information Processing Systems*, volume 33, 2020a.
- J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. In *Advances in Neural Information Processing Systems*, volume 33, Dec. 2020b. URL <https://arxiv.org/abs/2007.07481>.
- C. J. Watkins and P. Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- J. Woo, L. Shi, G. Joshi, and Y. Chi. Federated offline reinforcement learning: Collaborative single-policy coverage suffices. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pages 53165–53201. PMLR, 21–27 Jul 2024.
- Z. Wu, H. Shen, T. Chen, and Q. Ling. Byzantine-resilient decentralized policy evaluation with linear function approximation. *IEEE Transactions on Signal Processing*, 69:3839–3853, 2021.
- Y. Yan, G. Li, Y. Chen, and J. Fan. The efficacy of pessimism in asynchronous Q-learning. *arXiv preprint arXiv:2203.07368*, 2022.
- K. Yang, L. Yang, and S. Du. Q-learning with logarithmic regret. In *International Conference on Artificial Intelligence and Statistics*, pages 1576–1584. PMLR, 2021.

- T. Yang, S. Cen, Y. Wei, Y. Chen, and Y. Chi. Federated natural policy gradient and actor critic methods for multi-task reinforcement learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- S. Zeng, T. T. Doan, and J. Romberg. Finite-time analysis of decentralized stochastic approximation with applications in multi-agent and multi-task learning. In *Conference on Decision and Control*, pages 2641–2646. IEEE, 2021.
- Z. Zhang, Y. Zhou, and X. Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. In *Advances in Neural Information Processing Systems*, volume 33, 2020.