Finite-Time Theory for Momentum Q-learning

Bowen Weng[*1]   **Huaqing Xiong**[*1]   **Lin Zhao**[*2]   **Yingbin Liang**[1]   **Wei Zhang**[3]

[1]Department of Electrical and Computer Engineering, The Ohio State University, Columbus, Ohio, USA
[2]Department of Electrical and Computer Engineering , National University of Singapore , Singapore, Republic of Singapore
[3]Department of Mechanical and Energy Engineering, Southern University of Science and Technology, Shenzhen, China

Abstract

Existing studies indicate that momentum ideas in conventional optimization can be used to improve the performance of Q-learning algorithms. However, the finite-time analysis for momentum-based Q-learning algorithms is only available for the tabular case without function approximation. This paper analyzes a class of momentum-based Q-learning algorithms with finite-time convergence guarantee. Specifically, we propose the MomentumQ algorithm, which integrates the Nesterov's and Polyak's momentum schemes, and generalizes the existing momentum-based Q-learning algorithms. For the infinite state-action space case, we establish the convergence guarantee for MomentumQ with linear function approximation under Markovian sampling. In particular, we characterize a finite-time convergence rate which is provably faster than the vanilla Q-learning. This is the first finite-time analysis for momentum-based Q-learning algorithms with function approximation. For the tabular case under synchronous sampling, we also obtain a finite-time convergence rate that is slightly better than the SpeedyQ [Azar et al., 2011]. Finally, we demonstrate through various experiments that the proposed MomentumQ outperforms other momentum-based Q-learning algorithms.

1 INTRODUCTION

Reinforcement learning (RL) aims to design strategies for an agent to find a desirable policy through interacting with an environment in order to maximize an accumulative reward for a task. RL has received drastically growing attention in recent years and accomplished tremendous success in various application domains such as playing video games [Mnih

---

[*]equal contribution

et al., 2013], bipedal walking robot [Castillo et al., 2019], board game [Silver et al., 2017], to name a few. This paper focuses on Q-learning, which is a widely used model-free RL algorithm for finding the action-value function (known as the Q-function) of the optimal policy.

Q-learning was first proposed in Watkins and Dayan [1992] and has been studied extensively since then. For scenarios with a finite state-action space, the Q-function can be conveniently represented as a tabular function. The convergence of Q-learning in the tabular case was proved in Jaakkola et al. [1994]. In the case with a continuous state-action space, one typically approximates the Q-function with a parameterized function class of a relatively small parameter dimension. Among the rich approximation classes, linear function approximation [Bertsekas and Tsitsiklis, 1996, Sutton and Barto, 2018] and neural network function approximation [Mnih et al., 2013] are often adopted in the literature. We will review these studies in more details in Section 1.2.

The central idea of Q-learning algorithms is to solve an optimal Bellman equation [Bertsekas and Tsitsiklis, 1996] iteratively as a fixed point problem. Since the Bellman operator is expressed as the expected value over the underlying Markov decision processes (MDP) which is unknown, Q-learning (as a model-free algorithm) approximates it via its sampled version, and such an update can be viewed analogously to the first-order (stochastic) gradient descent algorithm [Baird, 1995]. This connection thus motivated several studies on accelerating Q-learning by incorporating various momentum schemes, such as Heavy-ball (HB) [Polyak, 1964] and Nesterov's accelerated gradient (NAG) [Nesterov, 2013] which were shown to accelerate gradient descent in conventional optimization algorithms. For example, speedy Q-learning (SpeedyQ) proposed in Azar et al. [2011] can be viewed as incorporating NAG to Q-learning and has been proved to accelerate the convergence with particularly designed learning rate in the tabular case. Devraj et al. [2019] applied HB and NAG to Q-learning separately with a matrix learning rate in the tabular case and analyzed their asymptotic properties. Vieillard et al. [2019] proposed momentum-based value iter-

ation by viewing the greedy policy as an analog of gradient ascent. However, the convergence is not guaranteed. To the best of our knowledge, the *finite-time* convergence rate has not been established for momentum-based Q-learning algorithms with function approximation. The focus of this paper is to address the above important question.

1.1 MAIN CONTRIBUTIONS

This paper investigates a general momentum-based Q-learning scheme (referred to as MomentumQ hereafter), which involves both NAG-type and HB-type of history information for accelerating Q-learning. The main contribution of this paper is three-fold.

First, we establish the finite-time convergence rate for MomentumQ with linear function approximation, and we show that this algorithm provably accelerates vanilla Q-learning. To the best of our knowledge, this is the first finite-time convergence guarantee for momentum-based Q-learning with linear function approximation.

Second, the only existing finite-time baseline bounds for momentum-based Q-learning is given by SpeedyQ [Azar et al., 2011] and Generalized SpeedQ [John et al., 2020] for the tabular case. Hence, to be able to compare with such a baseline, we also provide a finite-time analysis of MomentumQ in the tabular case. We show that it achieves a better (but order-wisely the same) convergence rate than SpeedyQ. Technically, due to the additional momentum terms in MomemtumQ, its analysis is more challenging than SpeedyQ and requires substantial new technical developments.

Finally, our numerical results show that the proposed MomentumQ outperforms the vanilla Q-learning as well as the other existing momentum-based Q-learning algorithms for both tabular and function approximation cases.

1.2 RELATED WORK

We review the most relevant studies on Q-learning here with a focus on the theoretical convergence analysis.

Q-learning with function approximation: When the state-action space is considerably large or even continuous, it is practical to properly discretize the space [Shah and Xie, 2018], or parameterize the Q-function with a certain function class. For linear MDP, Melo and Ribeiro [2007], Yang and Wang [2019] proposed provably sample-efficient Q-learning algorithms with linear function approximation. For more general MDPs with linear function approximation of the Q-function, finite-time convergence analysis was established in Zou et al. [2019], Chen et al. [2019] under Markovian sampling, in Du et al. [2019] on exploration samples and in Weng et al. [2020a] by incorporating Adam-type updates. Recently, Cai et al. [2019], Fan et al. [2019] and Xu and Gu [2019] established the convergence rate of

Q-learning with neural network approximation in the over-parameterized regime under i.i.d. and non-i.i.d sampling, respectively.

Tabular Q-learning: Q-learning was first proposed in Watkins and Dayan [1992] under finite state-action space. Regarding the theoretical studies, research of tabular Q-learning has focused on the asymptotic convergence which was usually studied via its connection to the corresponding stochastic approximation algorithm (see, for example, Tsitsiklis [1994], Jaakkola et al. [1994], Borkar and Meyn [2000], Melo [2001]). More recently, Lee and He [2019] provided asymptotic results for asynchronous Q-learning by formulating it as a switching affine system. Another research line has focused on the finite-time (i.e., non-asymptotic) analysis. Finite-time performance for Q-learning was first established in Szepesvári [1998]. Considering both synchronous and asynchronous Q-learning, Even-Dar and Mansour [2003] investigated the convergence rates under different choices of the learning rates. Sharper bounds on the finite-time convergence rate have been established in more recent work [Wainwright, 2019a, Qu and Wierman, 2020, Li et al., 2020].

Momentum-based Q-learning: For tabular Q-learning, several studies incorporated the momentum idea in conventional optimization to accelerate the convergence. Azar et al. [2011] proposed the SpeedyQ algorithm and characterized the finite-time performance. John et al. [2020] generalized SpeedyQ by introducing a relaxation parameter which was used to modify sample distribution and to redistribute the momentum. Devraj et al. [2019] extended HB and NAG with a matrix learning rate on the momentum. The asymptotic performance was analyzed under simplified assumptions. Vieillard et al. [2019] proposed a momentum-based value iteration and generalized the scheme to DQN. While some theoretical properties of the algorithms were explored in the tabular case, the convergence of the algorithm was not established. Among these studies, only Azar et al. [2011], John et al. [2020] characterized the finite-time rate for SpeedyQ in the tabular case, and such finite-time analysis for momentum-based Q-learning algorithms has not been provided for the function approximation case, which is the focus of this paper.

Other variants of Q-learning: Other than the above momentum-based Q-learning algorithms, which mainly exploit the acceleration ideas in conventional optimization, Q-learning also inspires a number of other variants, including residual Q-learning [Baird, 1995], phased Q-learning [Kearns and Singh, 1999], Zap Q-learning [Devraj and Meyn, 2017], periodic Q-learning [Lee and He, 2020], variance reduced Q-learning [Wainwright, 2019b, Li et al., 2020], and double Q-learning [Hasselt, 2010, Weng et al., 2020b, Xiong et al., 2020b] to name a few. These algorithms are proposed to speed up convergence rates or improve the performance by mitigating various issues in the implemen-

tation of Q-learning. In this paper, we mainly focus on the momentum-based Q-learning algorithm motivated by the optimization idea.

2 PRELIMINARIES

In this section, we provide the necessary background.

2.1 MARKOV DECISION PROCESS

We consider the standard reinforcement learning setting, where a learning agent interacts with a (possibly stochastic) environment modeled as a discrete-time discounted Markov decision process (MDP). Such an MDP is characterized by a quintuple $(\mathcal{X}, \mathcal{U}, P, R, \gamma)$, where $\mathcal{X}$ is the state space, $\mathcal{U}$ is the action space, $P : \mathcal{X} \times \mathcal{U} \times \mathcal{X} \mapsto [0, 1]$ is the probability transition kernel, namely, $P(\cdot|x, u)$ denotes the probability that the system takes the next state given the current state x and action u. In addition, $R : \mathcal{X} \times \mathcal{U} \mapsto [0, R_{\max}]$ denotes the reward function (or negative of the cost function) mapping the state-action pairs to a bounded subset of $\mathbb{R}$, and $\gamma \in (0, 1)$ is the discount factor. A policy $\pi : \mathcal{X} \mapsto \mathcal{U}$ represents a strategy to take actions, i.e., it captures the probability of taking each action at any given state. By following a policy π, we perform an action u_k with probability $\pi(u_k|x_k)$ at time k, observe a reward $r_k = R(x_k, u_k)$, and evolve to the next state x_{k+1} with the probability $P(x_{k+1}|x_k, u_k)$.

We define the value function as the expected return of following policy π and starting from state x, given by $V^\pi(x) = \mathbb{E}_P \sum_{k=0}^\infty \gamma^k r_k$, where $\mathbb{E}_P$ denotes the expectation with respect to the transition probability P. The Q-function is defined as the state-action value function $Q^\pi(x, u) = R(x, u) + \gamma \sum_{y \in \mathcal{X}} P(y|x, u) V^\pi(y)$, which is the return of performing action u at state s at the first step and following policy π thereafter.

2.2 TABULAR Q-LEARNING

Q-learning seeks to maximize the expected discounted return over policy π as formulated below.

$$\underset{\pi}{\text{maximize}} \quad V^\pi(x_0) = \mathbb{E}_P \left[\sum_{k=0}^\infty \gamma^k R(x_k, \pi(x_k)) \right],$$
$$\text{subject to} \quad x_{k+1} \sim P(\cdot|x_k, \pi(x_k)), \quad (1)$$

We let $\pi^\star$ denote the optimal stationary policy $\pi^\star : \mathcal{X} \mapsto \mathcal{U}$ which is the solution of the above optimization problem.

Define the Bellman operator $\mathcal{T}$ pointwisely as

$$\mathcal{T}Q(x, u) = R(x, u) + \gamma \mathbb{E}_P \max_{u' \in U(x')} Q(x', u'), \quad (2)$$

where $x' \sim P(\cdot|x, u)$ and $U(x)$ denotes the admissible set of actions at state x. It can be shown that the Bellman operator $\mathcal{T}$ is γ-contractive in the supremum norm $\|Q\| := \sup_{x,u} |Q(x, u)|$, i.e., it satisfies

$$\|\mathcal{T}Q(x, u) - \mathcal{T}Q'(x, u)\| \le \gamma \|Q(x, u) - Q'(x, u)\|. \quad (3)$$

Thus, $\mathcal{T}$ has a unique fixed point $Q^\star$ given by

$$Q^\star(x, u) = R(x, u) + \gamma \mathbb{E}_P \max_{u' \in U(x')} Q^\star(x', u'). \quad (4)$$

The above property suggests that starting with an arbitrary initial Q-function, we can apply the Bellman operator $\mathcal{T}$ iteratively to learn $Q^\star$. Hence, the optimal policy can be obtained from the optimal Q-function as:

$$\pi^\star(x) = \underset{u \in U(x)}{\text{argmax}} \, Q^\star(x, u), \forall x \in \mathcal{X}. \quad (5)$$

Note that the knowledge of the transition probability P is not needed in (5), which is one advantage of Q-learning.

In practice, exact evaluation of the Bellman operator (2) is usually not feasible due to the lack of the knowledge of the transition probability kernel. Instead, the sample-based *empirical Bellman operator* is used as an estimator. Specifically, for the kth round of iteration at the state-action pair (x, u), we sample the next state $y_k \sim P(\cdot|x, u)$, and then evaluate the empirical Bellman operator $\hat{\mathcal{T}}_k$ as

$$\hat{\mathcal{T}}_k Q_k(x, u) = R(x, u) + \gamma \max_{u' \in U(y_k)} Q_k(y_k, u'). \quad (6)$$

Then the update of tabular Q-learning is implemented as

$$Q_{k+1} = Q_k - \alpha_k(Q_k - \hat{\mathcal{T}}_k Q_k), \quad (7)$$

where α_k is the stepsize and we omit the dependence on (x, u) for simplicity when there is no confusion.

2.3 Q-LEARNING WITH LINEAR FUNCTION APPROXIMATION

For relatively large or even infinite state-action space $\mathcal{X} \times \mathcal{U}$, it is impractical to express the Q-function in an explicit tabular form with respect to each state-action pair. In such a case, the update rule of (7) is no longer directly applicable.

To handle such cases, a parametric function $\hat{Q}(x, u; \theta)$ is adopted as an approximation of the Q-function, where the parameter vector θ is of small dimension. Our focus here is the linear function class, which is often considered in the literature for establishing the finite-time analysis [Zou et al., 2019, Chen et al., 2019, Du et al., 2019]. Then the Q-function $\hat{Q}(x, u; \theta)$ can be written as

$$\hat{Q}(x, u; \theta) = \Phi(x, u)^T \theta, \quad (8)$$

where $\theta \in \mathbb{R}^d$, and $\Phi : \mathcal{S} \times \mathcal{A} \to \mathbb{R}^d$ is a vector function of size d, and the elements of Φ represent the nonlinear kernel

(feature) functions. Correspondingly, the updating rule of Q-learning with linear function approximation is given by

$$\theta_{k+1} = \theta_k - \alpha_k \Phi(x_k, u_k) \left[\Phi(x_k, u_k)^T \theta_k - R(x_k, u_k) \right.$$
$$\left. -\gamma \max_{u' \in U(x_{k+1})} \Phi(x_{k+1}, u')^T \theta_k \right], \qquad (9)$$

where α_k is the stepsize.

3 MOMENTUMQ ALGORITHM

In this section, we introduce the MomentumQ algorithm.

3.1 TABULAR MOMENTUMQ

Overall, MomentumQ integrates the Nesterov's momentum [Nesterov, 2013] and Polyak's Momentum [Polyak, 1964] together, with the learning rates flexibly interpolating between the two to optimize the momentum performance. Specifically, MomentumQ takes the form given by

$$S_k = (1 - a_k)Q_{k-1} + a_k \hat{\mathcal{T}}_k Q_{k-1},$$
$$P_k = (1 - a_k)Q_k + a_k \hat{\mathcal{T}}_k Q_k, \qquad (10)$$
$$Q_{k+1} = P_k + \underbrace{b_k(P_k - S_k)}_{\text{Nesterov's momentum}} + \underbrace{c_k(Q_k - Q_{k-1})}_{\text{Polyak's momentum}},$$

where a_k, b_k, c_k determine the learning rates. Algorithm 1 implements MomentumQ with a particular family of learning rates under synchronous sampling [Even-Dar and Mansour, 2003]. One special feature of the algorithm is the additional freedom introduced by the hyperparameter m. This hyperparameter actually stems from our finite analysis of MomentumQ in the tabular case. It helps to better bound the propagation of the learning error (i.e., $Q* - Q_k$, see also the proof of Lemma 7) via the particularly designed learning rate. Unlike SpeedyQ, which only admits a fixed set of learning rate (namely, $1/k$), MomentumQ's learning rate is more flexible due to the introduction of this hyperparameter. In practice, we found that in general setting $m = 1/\gamma + c$ for some arbitrary small number $c > 0$ (e.g., $c = 1$), will generally yields a very good performance in suppressing the overshoot of the learning error which was otherwise very significant in SpeedyQ learning. We will also see later in the simulation that the proposed algorithm accelerates the convergence for a number of arbitrarily chosen m that satisfies $m \geq 1/\gamma$.

Note that the proposed MomentumQ algorithm in (10) contains not only the momentum term $\hat{\mathcal{T}}_k Q_{k-1}$ in the update, but also the historical information Q_{k-1} explicitly. We compare MomentumQ with SpeedyQ [Azar et al., 2011] given by

$$Q_{k+1} = Q_k + a_k(\hat{\mathcal{T}}_k Q_k - Q_k) + (1 - a_k)(\hat{\mathcal{T}}_k Q_k - \hat{\mathcal{T}}_k Q_{k-1}). \qquad (11)$$

Algorithm 1 Synchronous Tabular MomentumQ

Input: Initial action-value function Q_0 and $Q_{-1} = Q_0$, discount factor γ, parameter $m \geq \frac{1}{\gamma}$, and iteration number T
for $k = 0, 1, 2, \cdots, T - 1$ **do**
    $a_k = \frac{1}{k+1}$, $b_k = k - m - 1$, $c_k = \frac{-k^2 + (m+1)k+1}{k+1}$;
    for each $(x, u) \in \mathcal{X} \times U(x)$ **do**
        Generate the next state sample $y_k \sim P(\cdot|x, u)$;
        $\hat{\mathcal{T}}_k Q_{k-1}(x, u) = R(x, u) + \gamma \max_{u \in U(y_k)} Q_{k-1}(y_k, u)$;
        $\hat{\mathcal{T}}_k Q_k(x, u) = R(x, u) + \gamma \max_{u \in U(y_k)} Q_k(y_k, u)$;
        $S_k(x, u) = (1 - a_k)Q_{k-1}(x, u) + a_k \hat{\mathcal{T}}_k Q_{k-1}(x, u)$;
        $P_k(x, u) = (1 - a_k)Q_k(x, u) + a_k \hat{\mathcal{T}}_k Q_k(x, u)$;
        $Q_{k+1}(x, u) = P_k(x, u) + b_k (P_k(x, u) - S_k(x, u))$
               $+ c_k(Q_k(x, u) - Q_{k-1}(x, u))$;
    end for
end for
Output: Q_T

We see from (11) that SpeedyQ contains only the momentum term $\hat{\mathcal{T}}_k Q_{k-1}$ in the update. In contrast, MomentumQ additionally incorporates the historical information Q_{k-1} explicitly. Indeed, the simulation in Section 5 indicates that MomentumQ effectively smooths out the large overshoots that are present in SpeedyQ and converges faster by incorporating additional historical terms. The finite-time analysis of MomentumQ is more challenging than SpeedyQ due to this difference, since the additional Q_{k-1} term increases the order of the recursion. We will discuss in more details later. Furthermore, (11) simply involves $\hat{\mathcal{T}}_k Q_k - \hat{\mathcal{T}}_k Q_{k-1}$ as the only momentum term, while our algorithm designs this part more systematically. We directly use two consecutive outputs of the empirical Bellman operators to update the Q-function and obtain S_k and P_k. Compared to (11), the resulted Nesterov's momentum is potentially a better estimation of the "gradient" for updating the Q function. This intuition is also verified in our numerical results.

3.2 MOMENTUMQ WITH LINEAR FUNCTION APPROXIMATION

For the case where the state-action space is considerably large, we use the linear function approximation to estimate the Q-function to overcome the curse of dimensionality.

Consider the case where the Q-function is approximated by a linear parameterized function. We propose MomentumQ for this case as

$$\theta_{k+1} = \theta_k + (b_k + c_k)(\theta_k - \theta_{k-1}) - a_k(1 + b_k)g_k + a_k b_k g_{k-1}, \qquad (12)$$

where

$$g_k := g(\theta_k; x_k, u_k, x_{k+1})$$

Algorithm 2 MomentumQ with linear approximation

Input: Initial parameters $\theta_{-1} = \theta_0$; discount factor γ; iteration number T.
for $k = 0, 1, 2, \cdots, T-1$ **do**
    Assign a_k, b_k, c_k;
    Sample $u_k \sim \pi$, $x_{k+1} \sim P(\cdot|x_k, u_k)$;
    Compute g_k as (13);
    Update $\theta_{k+1} = \theta_k + (b_k + c_k)(\theta_k - \theta_{k-1})$
                  $-a_k(1 + b_k)g_k + a_k b_k g_{k-1}$;
end for
Output: θ_{out}.

$$
= \Big(\Phi(x_k, u_k)^T \theta_k - \gamma \max_{u' \in U(x_{k+1})} \Phi(x_{k+1}, u')^T \theta_k
$$
$$
- R(x_k, u_k) \Big) \Phi(x_k, u_k). \tag{13}
$$

We focus on the more practical Markovian sampling model, in which the data tuples are sequentially drawn from a single trajectory under an unknown stationary distribution. More implementation details are referred to Algorithm 2.

4 FINITE-TIME CONVERGENCE RESULTS

In this section, we present our main results on the finite-time convergence rate guarantee for MomentumQ. We focus on linear function class and provide the first finite-time analysis for momentum-based Q-learning with function approximation. We also present our study of tabular MomentumQ in order to make a comparison with the only existing theory baseline for momentum-based Q-learning, which was established for tabular SpeedyQ.

4.1 MOMENTUMQ WITH LINEAR FUNCTION APPROXIMATION UNDER MARKOVIAN SAMPLING

In this section, we characterize the finite-time convergence guarantee for the proposed MomentumQ algorithm with linear function approximation under Markovian sampling. To proceed the convergence analysis, we first define

$$
\bar{g}(\theta) := \mathbb{E}_{\mu}[g(\theta; x, u, x')]
$$
$$
= \mathbb{E}_{\mu}\left[\Phi(x, u)^T \theta - R(x, u) - \gamma \max_{u' \in U(x')} \Phi(x', u')^T \theta \right] \Phi(x, u), \tag{14}
$$

where the expectation is taken over the stationary distribution of the sampling tuple (x, u, x').

We take the following standard assumptions in our analysis.

Assumption 1. *The columns of Φ are linearly independent and $\|\Phi\|_2 \le 1$.*

Assumption 2. *The term $\bar{g}(\cdot)$ has a unique root denoted as $\theta^\star$, i.e., $\bar{g}(\theta^\star) = 0$. There exists a constant $\delta > 0$, such that for any $\theta \in \mathbb{R}^d$ we have*

$$
(\theta - \theta^\star)^T \bar{g}(\theta) \ge \delta \|\theta - \theta^\star\|_2^2. \tag{15}
$$

Assumption 3. *The domain of the approximation parameters θ is contained in a ball $\mathcal{B}$ centered at $\theta = 0$ with a bounded diameter $D_{\max}$ and the optimal parameter $\theta^\star \in \mathcal{B}$. That is, there exists $D_{\max}$, such that $\|\theta - \theta'\|_2 \le D_{\max}, \forall \theta, \theta' \in \mathcal{B}$, and $\theta^\star \in \mathcal{B}$.*

Assumption 4. *There exist constants $\sigma > 0$ and $\rho \in (0, 1)$ such that*

$$
\sup_{x \in \mathcal{X}} d_{TV}(\mathbb{P}(x_k \in \cdot | x_0 = x), \mu) \le \sigma \rho^k \quad \forall k,
$$

where $d_{TV}(\mu, \nu)$ denotes the total-variation distance between the probability measures μ and ν.

Assumptions 1 and 2 are standard in the literature on theoretical analysis of Q-learning algorithms with linear function approximation [Bhandari et al., 2018, Chen et al., 2019, Zou et al., 2019]. The boundedness condition in Assumption 1 can be justified by normalization and hence does not lose generalization. Assumption 4 can easily hold for irreducible and aperiodic Markov chains, and is widely adopted in the literature on theoretical analysis of RL algorithms under Markovian sampling [Bhandari et al., 2018, Chen et al., 2019, Zou et al., 2019, Xu and Gu, 2019, Xiong et al., 2020a]. For Assumption 4, we further define the quantity of the mixing time $\tau^{mix}(\cdot)$ as follows, which denotes the duration of the time for the Markov chain to approach sufficiently close to its steady-state

$$
\tau^* := \tau^{mix}(\kappa) := \min \left\{ k = 1, 2, \dots \,|\, \sigma \rho^k \le \kappa \right\}. \tag{16}
$$

To understand the challenges of analyzing Markovian sampling in MomemtumQ, we first illustrate how a non-zero bias is introduced if the Markovian sampling is considered. For simplicity, we denote $O_k := (x_k, u_k, x_{k+1})$ as the data at time step k sampled from a Markov chain. Recall $g_k(\theta; O_k)$ in (13), and $\bar{g}(\theta) = \mathbb{E}[g(\theta; O_k)]$ in (14) where the expectation is taken over the marginal distribution of O_k since θ is fixed. However, if θ is random and dependent on O_k, the equality no longer holds. In particular, since θ_k is dependent on the historical tuples $\{O_1, O_2, \dots, O_k\}$, we have

$$
\bar{g}(\theta_k) \ne \mathbb{E}[g(\theta_k; O_k)|\theta_k].
$$

Thus, we have a non-zero bias due to Markovian sampling to approximate the expectation of $g_k^T(\theta_k - \theta^\star)$. Namely,

$$
\mathbb{E}[g_k^T(\theta_k - \theta^\star)] = \mathbb{E}[\bar{g}(\theta_k)^T(\theta_k - \theta^\star)]
$$
$$
+ \mathbb{E}[(g_k - \bar{g}(\theta_k))^T(\theta_k - \theta^\star)],
$$

where the second term on the right hand side captures the bias, which is the key challenge of the analysis under this setting. The following lemma develops an important upper bound on the bias term, which is a key step in the convergence analysis.

Lemma 1. *Suppose that Assumptions 1-4 hold and fix $\kappa > 0$ in (16). Let MomentumQ update as (12) by choosing non-increasing a_k, b_k, c_k and denote $\beta_k = b_k + c_k$ with $\beta_k \in (0, 1)$. Then we have*

$$
\mathbb{E}[(g_k - \bar{g}(\theta_k))^T (\theta_k - \theta^\star)]
$$
$$
\leq \begin{cases} \eta_1 \sum_{i=1}^{k-1} \beta_i + \eta_2 \sum_{i=1}^{k-1} a_i, & k \leq \tau^*; \\ 4D_{\max} G_{\max} \kappa + \eta_1 \tau^* \beta_{k-\tau^*} + \eta_2 \tau^* a_{k-\tau^*}, & k > \tau^*, \end{cases}
$$

where $\eta_1 = 2D_{\max}((1 + \gamma)D_{\max} + G_{\max}), \eta_2 = 6G_{\max}((1 + \gamma)D_{\max} + G_{\max})$ with $G_{\max} = 2D_{\max} + R_{\max}$.

With the bias term bounded, we are ready to provide the convergence result for MomentumQ with linear function approximation under Markovian sampling.

Theorem 1. *(MomentumQ with constant learning rate) Suppose that Assumptions 1-4 hold and fix $\kappa > 0$ in (16). Let $a_k = \alpha$, $b_k + c_k = \beta \lambda^k$ where $\beta, \lambda \in (0, 1)$ and $\alpha \in (0, \frac{1-\lambda}{2\delta})$. After running T steps of Algorithm 2 under Markovian sampling, we take the output $\theta_{out} = \theta_T$ and have*

$$
\mathbb{E} \|\theta_{out} - \theta^\star\|_2^2 \tag{17}
$$
$$
\leq \prod_{i=0}^{T-1} (1 - 2\alpha\delta(1 + b_i)) \|\theta_0 - \theta^\star\|_2^2
$$
$$
+ \beta \left(\frac{2\eta_1 \tau^*}{\delta} + \frac{C}{1 - 2\alpha\delta - \lambda} \right) (1 - 2\alpha\delta)^{T-1-\tau^*}
$$
$$
+ \frac{15G_{\max}^2 \alpha}{2\delta} + \frac{2\eta_2 \tau^* \alpha}{\delta} + \frac{8D_{\max} G_{\max} \kappa}{\delta}, \tag{18}
$$

where $C = 5D_{\max}^2 + 2\alpha D_{\max} G_{\max} + 4\alpha\eta_1 \tau^* \lambda$ with $G_{\max} = 2D_{\max} + R_{\max}$, and η_1, η_2 are defined in Lemma 1.

Theorem 1 indicates that the convergence behavior is determined by five terms. The first two terms capture the convergence rate as T changes, indicating that with a constant learning rate, MomentumQ enjoys an exponential convergence rate to a neighborhood of the global optimum. Since $\prod_{i=0}^{T-1}(1 - 2\alpha\delta(1+b_i)) < (1-2\alpha\delta)^T$, the dominant term of the convergence rate is the second term. The last three terms capture the convergence error. Since one usually chooses $\kappa = \alpha_k = \alpha$, the convergence error can be made as small as possible by choosing a sufficiently small learning rate.

As a comparison, the convergence of the vanilla Q-learning under similar assumptions and Markovian sampling is obtained in Chen et al. [2019] as $\mathbb{E} \|\theta_{out} - \theta^\star\|_2^2 \leq (1 -$

$2\delta\alpha)^T \|\theta_0 - \theta^\star\|_2^2 + \alpha C_1 + \kappa C_2$ for some constants C_1, C_2. Clearly, the dominant order in (18) can have a smaller coefficient than that of the vanilla Q-learning by setting a small β, so that MomentumQ can enjoy a better convergence rate.

In addition, one can also observe that α, β control a set of tradeoffs. First, while smaller α yields a smaller convergence error, it also slows down the convergence rate. As for β, although smaller β yields a smaller coefficient in the dominant term, it can also slow down the convergence rate because b_i in the first term needs to be small.

Next, we seek to remove the convergence error and balance the tradeoff caused by the choice of a_k. To this end, we can choose a diminishing learning rate and obtain the following theorem.

Theorem 2. *(MomentumQ with diminishing learning rate) Suppose that Assumptions 1-4 hold and fix $\kappa > 0$. Let $a_k = \frac{\alpha}{\sqrt{k}}, b_k + c_k = \beta \lambda^k$ with $\alpha > 0, \beta, \lambda \in (0, 1)$. After running T steps of Algorithm 2 under Markovian sampling, we take the output $\theta_{out} = \frac{1}{T} \sum_{k=1}^{T} \theta_k$ and have*

$$
\mathbb{E} \|\theta_{out} - \theta^\star\|_2^2
$$
$$
\leq \frac{D_{\max}^2/\alpha + 30\alpha G_{\max}^2 + 16\tau^* \alpha \eta_2}{2\delta\sqrt{T}} + \frac{8D_{\max} G_{\max} \kappa}{\delta}
$$
$$
+ \frac{1}{T} \left[\frac{5\beta D_{\max}^2}{2\alpha\delta(1 - \lambda)^2} + \frac{D_{\max} G_{\max} \beta \lambda + 4\tau^* \eta_1 \beta \lambda}{\delta(1 - \lambda)} \right],
$$

where η_1, η_2 are defined in Lemma 1.

In Theorem 2, if we choose $\kappa = \alpha_k = \alpha/\sqrt{T}$, then the mixing time $\tau^* = \mathcal{O}(\log T)$. Thus, MomentumQ converges to the global optimum at a rate of $\mathcal{O}(\log T/\sqrt{T})$ under a diminishing learning rate.

4.2 TABULAR MOMENTUMQ

In this subsection, we provide the finite-time analysis for tabular MomentumQ as listed in Algorithm 1. As we mention before, MomentumQ combines different types of momentum terms dynamically. This requires substantial new technical developments here in the convergence analysis.

We assume that the state space $\mathcal{X}$ and the action space $\mathcal{U}$ are finite with cardinalities $|\mathcal{X}|$ and $|\mathcal{U}|$, respectively. We denote $n = |\mathcal{X}| \cdot |\mathcal{U}|$. We also need the following assumption in our analysis.

Assumption 5. *The Q-function is uniformly bounded throughout the learning process. That is, $\exists V_{\max}$, such that $\|Q_k\| \leq V_{\max}, \forall k \geq 0$.*

Note that it is nontrivial to show the boundedness of the proposed iteration scheme. In fact, it is usually assumed for proving convergence of many such complicated stochastic

approximation algorithms [Kushner and Yin, 2003]. Alternatively, one can extend the ODE method [Borkar and Meyn, 2000] considerably to show the boundedness, which we leave for our future work.

To facilitate the analysis, we rewrite (10) in a more compact form as

$$Q_{k+1} = (1-a_k)Q_k + [b_k(1-a_k)+c_k](Q_k - Q_{k-1})$$
$$+ a_k \left[(1+b_k)\hat{\mathcal{T}}_k Q_k - b_k \hat{\mathcal{T}}_k Q_{k-1} \right]. \quad (19)$$

Our analysis first bounds the errors of approximating the exact Bellman operator $\mathcal{T}$ with empirical Bellman operators $\hat{\mathcal{T}}_k$. For convenience, we denote the last term in (19) by

$$\mathcal{D}_k [Q_k, Q_{k-1}] := (1+b_k)\hat{\mathcal{T}}_k Q_k - b_k \hat{\mathcal{T}}_k Q_{k-1}, \quad (20)$$

for all $k \geq 0$. Note that $\mathcal{D}_k$ is a function of all samples $\{y_1, y_2, \cdots, y_k\}$ for all state-action pairs (x, u) up to round k. Let $\mathcal{F}_k$ denote the filtration generated by the sequence of these random variables $\{y_1, y_2, \cdots, y_k\}$. We see that $\mathcal{D}_k \in \mathcal{F}_k$ and $Q_{k+1} \in \mathcal{F}_k$. Then if we define $\mathcal{D}[Q_k, Q_{k-1}]$ as the conditional expectation of $\mathcal{D}_k[Q_k, Q_{k-1}]$ given $\mathcal{F}_{k-1}$, we obtain by the definition of $\mathcal{T}$ that

$$\mathcal{D}[Q_k, Q_{k-1}] := \mathbb{E}_P (\mathcal{D}_k[Q_k, Q_{k-1}] | \mathcal{F}_{k-1})$$
$$= (1+b_k)\mathcal{T}Q_k - b_k \mathcal{T}Q_{k-1}. \quad (21)$$

Now define the error between $\mathcal{D}_k$ and $\mathcal{D}$ as follows:

$$\epsilon_k := \mathcal{D}[Q_k, Q_{k-1}] - \mathcal{D}_k[Q_k, Q_{k-1}]. \quad (22)$$

Clearly $\mathbb{E}_P (\epsilon_k | \mathcal{F}_{k-1}) = 0$. This shows that $\forall (x, u) \in \mathcal{X} \times U(x)$, the sequence of the estimation errors $\{\epsilon_k(x, u)\}_{k=0}^{T}$ is a martingale difference sequence with respect to the filtration $\mathcal{F}_k$. In other words, if we denote

$$E_k(x, u) := \sum_{j=0}^{k} \epsilon_j(x, u), \quad (23)$$

then E_k is a martingale sequence with respect to $\mathcal{F}_k$, $\forall (x, u) \in \mathcal{X} \times U(x)$ and $\forall k \geq 0$.

The following proposition provides the uniform bounds of $\mathcal{D}_k$ and ϵ_k.

Proposition 1. *Suppose Assumption 5 holds. Consider MomentumQ as in Algorithm 1. Then the terms $\mathcal{D}_k[Q_k, Q_{k-1}]$ defined in (20) and ϵ_k in (22) are uniformly bounded for all $k \geq 0$. Specifically, $\exists \bar{D} > 0$, s.t. $\|\mathcal{D}_k[Q_k, Q_{k-1}]\| \leq \bar{D}, \|\epsilon_k\| \leq 2\bar{D}, \forall k \geq 0$.*

The uniform bounds proved in Proposition 1 are critical in the derivation of the main theorem below.

Theorem 3. *Suppose Assumption 5 holds. Consider Algorithm 1 where $m \geq 1/\gamma$. Then, with probability at least*

$1 - \delta$, the output of MomentumQ satisfies for $T > m$:

$$\|Q^\star - Q_T\| \leq \frac{\tilde{h}V_{\max} + \bar{D}\sqrt{8(T - \lfloor m \rfloor - 1)\log \frac{2n}{\delta}}}{T(1-\gamma)},$$
$$(24)$$

where $\tilde{h} = 2\gamma(m + \lfloor m \rfloor + 2) + 2$, $\bar{D}$ is specified in Proposition 1, and $\lfloor m \rfloor$ denotes the largest integer that does not exceed m.

Using the Borel–Cantelli lemma, we immediately have the following corollary.

Corollary 1. *Q_k converges to $Q^\star$ almost surely at a rate of at least $\tilde{\mathcal{O}}(\frac{\sqrt{(T - \lfloor m \rfloor - 1)\log n}}{(1-\gamma)^2 T})$.*

This rate is slightly better than $\tilde{\mathcal{O}}(\frac{\sqrt{\log n}}{(1-\gamma)^2 \sqrt{T}})$ of SpeedyQ due to the presence of $m > 1$. We note that the Generalized SpeedyQ [John et al., 2020] only slightly improves the constant multipliers over SpeedyQ. In addition, implementation of Generalized SpeedyQ requires a complicated sampling procedure, which added extra computational complexity.

5 EXPERIMENTS

We evaluate the performance of the proposed MomentumQ scheme and compare it with other related Q-learning algorithms over a series of FrozenLake games (see the appendix for further specifications of the FrozonLake problem). We present the empirical results for both the tabular MomentumQ and MomentumQ with linear function approximation which are discussed above.

5.1 EXPERIMENTS ON TABULAR MOMENTUMQ

We compare MomentumQ with two other existing momentum-based Q-learning algorithms: SpeedyQ and the Nesterov stochastic approximation (NeSA) algorithm [Devraj et al., 2019]. In addition, we include comparisons with vanilla Q-learning to demonstrate the advantages of the proposed momentum techniques.

The experimental settings in this section are consistent with those of MomentumQ in Algorithm 1 and SpeedyQ in Azar et al. [2011, Algorithm 1]. Thus the numerical results should be able to give a convincing comparison between two algorithms. It is worth mentioning that the tabular MomentumQ has an additional hyperparameter m that can take a wide range of values (recall $m \geq 1/\gamma$). We experiment with several different m's. For relatively large m values (e.g., when $m > 10$), the learning rates are shifted to step from $1/(m+1)$, that is, $\alpha_k = 1/(m+k+1)$, for $k = 0, 1, 2, \ldots$. This is to avoid the large errors accumulated from initial iterations when $b_k < 0$, which are reflected in the constants

(a) FrozenLake-4×4 (b) FrozenLake-8×8

Figure 1: Comparing MomentumQ with NeSA, SpeedQ, and VanillaQ with tabular Q-function.

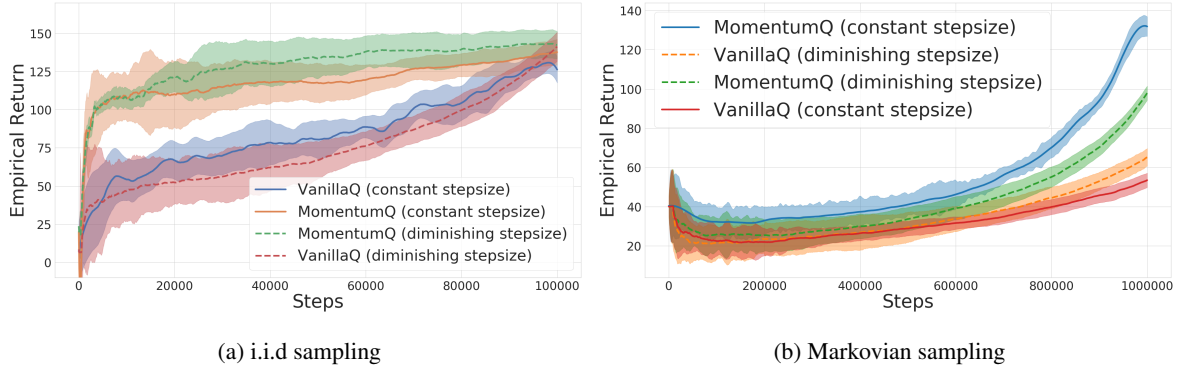
(a) i.i.d sampling   (b) Markovian sampling

Figure 2: Comparison of MomentumQ with VanillaQ in the FrozenLake-128×128 task with various learning rate schemes and sampling strategies.

in (24). Note that this shift does not change the obtained theoretical order of the convergence rate. We observe stable and consistent improvement in convergence and optimality across various tests when choosing large m values, which also aligns with our theoretical analysis.

Considering the randomness embedded in MDP of both FrozenLake games, we evaluate the performance of each algorithm with 20 different random seeds and then illustrate the average loss and standard deviation in Fig. 1a and Fig. 1b. For evaluation purpose, we have access to the true transition probability, and can find the ground truth optimal Q-function $Q^\star$ using dynamic programming. In both games, the loss at step k is then defined as $\|Q_k - Q^\star\|$. It can be seen from the plots that MomentumQ with a rather wide choices of m can all converge faster than vanilla Q-learning and SpeedyQ. It shows competitive performance against NeSA with smaller variance presented. Note that the high variance observed in the NeSA training aligns with the previous reported results from Devraj et al. [2019] under different tasks.

5.2 EXPERIMENTS WITH FUNCTION APPROXIMATION

We adopt the FrozenLake-128×128 as the benchmark task to evaluate the performance of MomentumQ with linear function approximation and compare it with the vanilla Q-learning (referred to as VanillaQ). Both algorithms are evaluated with different learning rate schemes (constant &

diminishing stepsize), as well as different sampling strategies (i.i.d. and Markovian). We note that SpeedyQ and NeSA have been proposed in the literature only for the tabular setting and are thus not included here for comparison.

Note that the i.i.d. sampling is an ideal assumption and cannot be satisfied perfectly in practice. For our implementation, we perform i.i.d. sampling strategy in a similar fashion to the experience replay [Mnih et al., 2013] typically used for DQN training. A data buffer, referred to as the experience, is accumulated with data points collected across multiple training steps in the past. At each training step, the training data is then randomly uniformly sampled from the data buffer. In contrast, the Markovian sampling takes the training samples in an "on-policy" manner where the collected data points are fed in to the Q-learning process right after. At step k, the performance of the algorithm is evaluated through the total return of 150 rounds of trials. Similarly to the tabular setup, we execute each algorithm 20 times with different random seeds and illustrate the average return and standard deviation in Fig. 2a with i.i.d. sampling and Fig. 2b with Markovian sampling.

Overall, the MomentumQ algorithm has exhibited superior performance than the vanilla Q-learning. In particular, training with i.i.d. sampling is significantly faster than the Markovian sampling, which can be also expected from our theoretical results. Within the same sampling strategy, MomentumQ is also faster in convergence than the vanilla Q-learning with the same learning rate scheme.

6 CONCLUSION

We proposed new momentum-based Q-learning algorithms for both the tabular and linear function approximation cases, which are respectively applicable to finite and continuous state-action spaces. We further characterized the convergence rate with empirical evaluation. The proposed algorithms accelerate the convergence in comparison to vanilla Q-learning on various challenging tasks under both tabular and parametric Q-learning settings.

References

Noga Alon and Joel H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, 3rd edition, 2008.

Mohammad Gheshlaghi Azar, Remi Munos, M Ghavamzadaeh, and Hilbert J Kappen. Speedy Q-learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2411–2419, 2011.

Leemon Baird. Residual algorithms: Reinforcement learning with function approximation. In *Machine Learning Proceedings 1995*, pages 30–37. Elsevier, 1995.

Dimitri P. Bertsekas and John N Tsitsiklis. *Neuro-Dynamic Programming*, volume 5. Athena Scientific, 1996.

Jalaj Bhandari, Daniel Russo, and Raghav Singal. A finite time analysis of temporal difference learning with linear function approximation. In *Conference on Learning Theory (COLT)*, 2018.

Vivek S Borkar and Sean P Meyn. The ode method for convergence of stochastic approximation and reinforcement learning. *SIAM Journal on Control and Optimization*, 38 (2):447–469, 2000.

Qi Cai, Zhuoran Yang, Jason D Lee, and Zhaoran Wang. Neural temporal-difference learning converges to global optima. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 11312–11322, 2019.

Guillermo A Castillo, Bowen Weng, Ayonga Hereid, Zheng Wang, and Wei Zhang. Reinforcement learning meets hybrid zero dynamics: A case study for rabbit. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 284–290, 2019.

Zaiwei Chen, Sheng Zhang, Thinh T. Doan, Siva Theja Maguluri, and John-Paul Clarke. Finite-time analysis of Q-learning with linear function approximation. *arXiv preprint arXiv:1905.11425*, 2019.

Adithya M Devraj and Sean Meyn. Zap Q-learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2235–2244, 2017.

Adithya M. Devraj, Ana Bušić, and Sean Meyn. On matrix momentum stochastic approximation and applications to Q-learning. In *57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 749–756, 2019.

Simon S Du, Yuping Luo, Ruosong Wang, and Hanrui Zhang. Provably efficient Q-learning with function approximation via distribution shift error checking oracle. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 8058–8068, 2019.

Rick Durrett. *Probability: Theory and Examples*. Cambridge University Press, Cambridge, 5 edition, 2019. doi: 10.1017/9781108591034.

Eyal Even-Dar and Yishay Mansour. Learning rates for Q-learning. *Journal of Machine Learning Research*, 5 (Dec):1–25, 2003.

Jianqing Fan, Zhaoran Wang, Yuchen Xie, and Zhuoran Yang. A theoretical analysis of deep Q-learning. *arXiv preprint arXiv:1901.00137*, 2019.

Hado V Hasselt. Double Q-learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2613–2621, 2010.

Tommi Jaakkola, Michael I Jordan, and Satinder P Singh. Convergence of stochastic iterative dynamic programming algorithms. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 703–710, 1994.

I. John, C. Kamanchi, and S. Bhatnagar. Generalized speedy q-learning. *IEEE Control Systems Letters*, 4(3):524–529, 2020.

Michael J Kearns and Satinder P Singh. Finite-sample convergence rates for Q-learning and indirect algorithms. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 996–1002, 1999.

Harold Kushner and G George Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35. Springer Science & Business Media, 2003.

Donghwan Lee and Niao He. A unified switching system perspective and ODE analysis of Q-learning algorithms. *arXiv preprint arXiv:1912.02270*, 2019.

Donghwan Lee and Niao He. Periodic Q-learning. *arXiv preprint arXiv:2002.09795*, 2020.

Gen Li, Yuting Wei, Yuejie Chi, Yuantao Gu, and Yuxin Chen. Sample complexity of asynchronous q-learning: Sharper analysis and variance reduction. *arXiv preprint arXiv:2006.03041*, 2020.

Francisco S Melo. Convergence of Q-learning: A simple proof. *Institute Of Systems and Robotics, Tech. Rep*, pages 1–4, 2001.

Francisco S Melo and M Isabel Ribeiro. Q-learning with linear function approximation. In *International Conference on Computational Learning Theory*, pages 308–322, 2007.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.

Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.

Guannan Qu and Adam Wierman. Finite-time analysis of asynchronous stochastic approximation and Q-learning. *arXiv preprint arXiv:2002.00260*, 2020.

Devavrat Shah and Qiaomin Xie. Q-learning with nearest neighbors. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3111–3121, 2018.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017. ISSN 1476-4687.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction.* MIT press, 2018.

Csaba Szepesvári. The asymptotic convergence-rate of Q-learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1064–1070, 1998.

John N Tsitsiklis. Asynchronous stochastic approximation and Q-learning. *Machine learning*, 16(3):185–202, 1994.

Nino Vieillard, Bruno Scherrer, Olivier Pietquin, and Matthieu Geist. Momentum in reinforcement learning. *arXiv preprint arXiv:1910.09322*, 2019.

Martin J Wainwright. Stochastic approximation with cone-contractive operators: Sharp ℓ_∞-bounds for Q-learning. *arXiv preprint arXiv:1905.06265*, 2019a.

Martin J Wainwright. Variance-reduced Q-learning is minimax optimal. *arXiv preprint arXiv:1906.04697*, 2019b.

Christopher J.C.H. Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8(3-4):279–292, 1992.

Bowen Weng, Huaqing Xiong, Yingbin Liang, and Wei Zhang. Analysis of Q-learning with adaptation and momentum restart for gradient descent. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20)*, pages 3051–3057, 2020a.

Wentao Weng, Harsh Gupta, Niao He, Lei Ying, and Srikant R. Provably-efficient double Q-learning. *arXiv preprint arXiv:arXiv:2007.05034*, 2020b.

Huaqing Xiong, Tengyu Xu, Yingbin Liang, and Wei Zhang. Non-asymptotic convergence of Adam-type reinforcement learning algorithms under markovian sampling. *arXiv preprint arXiv:2002.06286*, 2020a.

Huaqing Xiong, Lin Zhao, Yingbin Liang, and Wei Zhang. Finite-time analysis for double Q-learning. *arXiv preprint arXiv:2009.14257*, 2020b.

Pan Xu and Quanquan Gu. A finite-time analysis of Q-learning with neural network function approximation. *arXiv preprint arXiv:1912.04511*, 2019.

Lin Yang and Mengdi Wang. Sample-optimal parametric Q-learning using linearly additive features. In *International Conference on Machine Learning (ICML)*, pages 6995–7004, 2019.

Shaofeng Zou, Tengyu Xu, and Yingbin Liang. Finite-sample analysis for SARSA with linear function approximation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 8665–8675, 2019.

Supplementary Materials

A   SPECIFICATIONS OF FROZENLAKE PROBLEM

FrozenLake is a classic benchmark problem for Q-learning, in which an agent controls the movement of a character in an $n \times n$ grid world. Some tiles of the grid are walkable, and others lead to the agent falling into the water. Additionally, the movement direction of the agent is uncertain and only partially depends on the chosen direction. The agent is rewarded for finding a feasible path to a goal tile. As shown in Figure 3 with a Frozenlake-8×8 task, "S" is the safe starting point, "F" is the safe frozen surface, "H" stands for the hole that terminates the game, and "G" is the target state that comes with an immediate reward of 1. This forms a problem with the state-space size n^2, the action-space size 4 and the reward space $R = \{0, 1\}$. For tabular Q-learning algorithms with finite state-action problems of relatively small dimensions, FrozenLake-4×4 and FrozenLake-8×8 are two typical benchmark tasks. As the grid world becomes large, e.g., FrozenLake-128×128, Q-learning with linear function approximation is then adopted to solve the problem.


Figure 3: The FrozenLake-8×8 task environment.

B   PROOF OF LEMMA 1

We bound the expectation of bias via constructing a new Markov chain and applying some techniques from information theory. Before deriving the bound, we first introduce some technical lemmas.

Lemma 2. *Suppose Assumptions 1 and 3 hold. Then for g_k defined in (12), we have $\|g_k\|_2 \leq G_{\max}$ for all k, where $G_{\max} = 2D_{\max} + R_{\max}$.*

Proof. Following from the definition of g_k and the assumptions that $\|\Phi(x, u)\|_2 \leq 1, \|\theta\|_2 \leq D_{\max}$, and $\|R(x, u)\|_2 \leq R_{\max}$, we have

$$\|g_k\|_2 = \left\| (\Phi(x_k, u_k)^T \theta_k - R(x_k, u_k) - \gamma \max_{u' \in U(x_{k+1})} \Phi(x_{k+1}, u')^T \theta_k) \Phi(x_k, u_k) \right\|_2$$

$$\leq \left\| \Phi(x_k, u_k)^T \theta_k \right\|_2 + \|R(x_k, u_k)\|_2 + \max_{u' \in U(x_{k+1})} \left\| \Phi(x_{k+1}, u')^T \theta_k \right\|_2$$

$$\leq 2D_{\max} + R_{\max},$$

where we use Cauchy-Schwartz inequality and the triangle inequality. $\square$

For notational simplicity, throughout this section we use $O = (x, u, x')$ to denote the sample tuple and $O_k = (x_k, u_k, x_{k+1})$ to denote the sample tuple at time k.

Lemma 3. *Let $\xi(\theta; O) := (g(\theta; O) - \bar{g}(\theta))^T (\theta - \theta^\star)$. Then $\xi(\theta; O)$ is uniformly bounded by*

$$|\xi(\theta; O)| \leq 2D_{\max} G_{\max}, \quad \forall \theta \in \mathcal{B},$$

and it is Lipschitz continuous with

$$|\xi(\theta; O) - \xi(\theta'; O)| \leq 2((1 + \gamma)D_{\max} + G_{\max}) \|\theta - \theta'\|_2, \quad \forall \theta, \theta' \in \mathcal{B}.$$

Proof. The first statement is straightforward based on Assumption 3 and Lemma 2. That is,

$$|\xi(\theta; O)| \le \|g(\theta; O) - \bar{g}(\theta)\|_2 \|\theta - \theta^\star\|_2 \le 2D_{\max}G_{\max}.$$

Next to prove the Lipschitz condition, we first prove the Lipschitz condition of $g(\theta; O_k)$ with respect to θ.

$$\|g(\theta; O) - g(\theta'; O)\|_2 \overset{(i)}{\le} |\Phi(x, u)^T(\theta - \theta') + \gamma \max_{u' \in U(x')} \Phi(x', u')^T\theta' - \gamma \max_{u' \in U(x')} \Phi(x', u')^T\theta|$$

$$\overset{(ii)}{\le} |\Phi(x, u)^T(\theta - \theta')| + |\gamma \max_{u' \in U(x')} \Phi(x', u')^T\theta' - \gamma \max_{u' \in U(x')} \Phi(x', u')^T\theta|,$$

where (i) follows from Cauchy-Schwartz inequality and the assumption $\|\Phi\|_2 \le 1$, and (ii) follows from the triangle inequality.

Now we consider two cases. If the item in the second norm of (ii) is non-negative, we let $u^\star = \arg\max_{u' \in U(x')} \Phi(x', u')^T\theta'$. Then $\max_{u' \in U(x')} \Phi(x', u')^T\theta \ge \Phi(x', u^\star)^T\theta$. Thus, we continue to bound the above inequality as

$$\|g(\theta; O) - g(\theta'; O)\|_2 \le |\Phi(x, u)^T(\theta - \theta')| + \gamma\Phi(x', u^\star)^T(\theta' - \theta)$$
$$|\Phi(x, u)^T(\theta - \theta')| + \gamma|\Phi(x', u^\star)^T(\theta - \theta')| \tag{25}$$

Similarly, if this item is negative, we let $u^\star = \arg\max_{u' \in U(x')} \Phi(x', u')^T\theta$. Then $\max_{u' \in U(x')} \Phi(x', u')^T\theta' \ge \Phi(x', u^\star)^T\theta'$. Thus, we have

$$\|g(\theta; O) - g(\theta'; O)\|_2 \le |\Phi(x, u)^T(\theta - \theta')| + \gamma\Phi(x', u^\star)^T(\theta - \theta')$$
$$|\Phi(x, u)^T(\theta - \theta')| + \gamma|\Phi(x', u^\star)^T(\theta - \theta')| \tag{26}$$

Then it follows from (25) and (26) that

$$\|g(\theta; O) - g(\theta'; O)\|_2 \le (1 + \gamma)\|\theta - \theta'\|.$$

Similarly, we obtain the same result for $\bar{g}(\theta)$ as follows.

$$\|\bar{g}(\theta) - \bar{g}(\theta')\|_2 \le \underset{\mu}{\mathbb{E}} \|g_k(\theta) - g_k(\theta')\|_2 \le (1 + \gamma)\|\theta - \theta'\|_2.$$

Then we focus on obtaining the second statement,

$$|\xi(\theta; O) - \xi(\theta'; O)|$$
$$= |(g(\theta; O) - \bar{g}(\theta))^T(\theta - \theta^\star) - (g(\theta'; O) - \bar{g}(\theta'))^T(\theta' - \theta^\star)|$$
$$\le \|g(\theta; O) - \bar{g}(\theta)\|_2 \|\theta - \theta'\|_2 + \|\theta' - \theta^\star\|_2 \|(g(\theta; O) - \bar{g}(\theta)) - (g(\theta'; O) - \bar{g}(\theta'))\|_2$$
$$\overset{(i)}{\le} 2G_{\max}\|\theta - \theta'\|_2 + D_{\max}\|(g(\theta; O) - g(\theta'; O)) - (\bar{g}(\theta) - \bar{g}(\theta'))\|_2$$
$$\overset{(ii)}{\le} 2G_{\max}\|\theta - \theta'\|_2 + 2D_{\max}(1 + \gamma)\|\theta - \theta'\|_2$$
$$= 2((1 + \gamma)D_{\max} + G_{\max})\|\theta - \theta'\|_2,$$

where (i) follows from Assumption 3 and Lemma 2, and (ii) follows from triangle inequality and (25). $\square$

We use $X \to Z \to Y$ to indicate that the random variable X and Y are independent conditioned on Z.

Lemma 4. *[Bhandari et al., 2018, Lemma 9] Consider two random variables X and Y such that*

$$X \to x_k \to x_{k+\tau} \to Y, \tag{27}$$

for fixed k and $\tau > 0$. Suppose Assumption 4 holds. Let X', Y' are independent copies drawn from the marginal distributions of X and Y, that is $\mathbb{P}(X' = \cdot, Y' = \cdot) = \mathbb{P}(X = \cdot)\mathbb{P}(Y = \cdot)$. Then, for any bounded v, we have

$$|\mathbb{E}[v(X, Y)] - \mathbb{E}[v(X', Y')]| \le 2\|v\|_\infty (\sigma\rho^\tau).$$

We continue the proof of Lemma 1. We first develop the connection between $\xi(\theta_k; O_k)$ and $\xi(\theta_{k-\tau}; O_k)$ via Lemma 3. To do so, we first observe that

$$
\begin{aligned}
\|\theta_{i+1} - \theta_i\|_2 &= \|\beta_i(\theta_i - \theta_{i-1}) + a_i(1 + b_i)g_i + a_i b_i g_{i-1}\|_2 \\
&\overset{(i)}{\leq} \|\beta_i(\theta_i - \theta_{i-1})\|_2 + \|a_i(1 + b_i)g_i\|_2 + \|a_i b_i g_{i-1}\|_2 \\
&\overset{(ii)}{\leq} D_{\max}\beta_i + 3G_{\max}a_i,
\end{aligned}
$$

where (i) follows from the triangle inequality and (ii) from the Assumptions 3 and 2 and the fact $b_i < 1$. Then we have

$$
\|\theta_k - \theta_{k-\tau}\|_2 \leq \sum_{i=k-\tau}^{k-1} \|\theta_{i+1} - \theta_i\|_2 \leq D_{\max} \sum_{i=k-\tau}^{k-1} \beta_i + 3G_{\max} \sum_{i=k-\tau}^{k-1} a_i.
$$

Thus, we can relate $\xi(\theta_k; O_k)$ and $\xi(\theta_{k-\tau}; O_k)$ by using the Lipschitz property established in Lemma 3 as follows:

$$
\begin{aligned}
\xi(\theta_k; O_k) - \xi(\theta_{k-\tau}; O_k) &\leq |\xi(\theta_k; O_k) - \xi(\theta_{k-\tau}; O_k)| \\
&\leq 2((1 + \gamma)D_{\max} + G_{\max}) \|\theta_k - \theta_{k-\tau}\|_2 \\
&\leq 2((1 + \gamma)D_{\max} + G_{\max}) \left(D_{\max} \sum_{i=k-\tau}^{k-1} \beta_i + 3G_{\max} \sum_{i=k-\tau}^{k-1} a_i \right). \quad (28)
\end{aligned}
$$

Next, we bound $\mathbb{E}[\xi(\theta_{k-\tau}; O_k)]$ using Lemma 4. Observe that given any deterministic $\theta \in \mathcal{B}$, we have

$$
\mathbb{E}[\xi(\theta; O_k)] = (\mathbb{E}[g(\theta; O_k)] - \bar{g}(\theta))^T(\theta - \theta^\star) = 0.
$$

Since θ_0 is a fixed constant, we have $\mathbb{E}[\xi(\theta_0, O_k)] = 0$. Now we are ready to bound $\mathbb{E}[\xi(\theta_{k-\tau}, O_k)]$ via Lemma 4 by constructing a random process satisfying (27). To do so, consider random variables $\theta'_{k-\tau}$ and O'_k drawn independently from the marginal distribution of $\theta_{k-\tau}$ and O_k, so that $\mathbb{P}(\theta'_{k-\tau} = \cdot, O'_k = \cdot) = \mathbb{P}(\theta_{k-\tau} = \cdot)\mathbb{P}(O_k = \cdot)$. We further obtain $\mathbb{E}[\xi(\theta'_{k-\tau}, O'_k)] = \mathbb{E}[\mathbb{E}[\xi(\theta'_{k-\tau}, O'_k)|\theta'_{k-\tau}]] = 0$ since $\theta'_{k-\tau}$ and O'_k are independent. Combining Lemmas 3 and 4, we have

$$
\mathbb{E}[\xi(\theta_{k-\tau}, O_k)] \leq 2(2D_{\max}G_{\max})(\sigma\rho^\tau). \quad (29)
$$

Finally, we are ready to bound the bias. We first take expectation for both sides of (28) and obtain

$$
\mathbb{E}[\xi(\theta_k; O_k)] \leq \mathbb{E}[\xi(\theta_{k-\tau}; O_k)] + 2((1 + \gamma)D_{\max} + G_{\max}) \left(D_{\max} \sum_{i=k-\tau}^{k-1} \beta_i + 3G_{\max} \sum_{i=k-\tau}^{k-1} a_i \right).
$$

When $k \leq \tau^{mix}(\kappa)$, we choose $\tau = k$ and have

$$
\begin{aligned}
\mathbb{E}[\xi(\theta_k; O_k)] &\leq \mathbb{E}[\xi(\theta_0; O_k)] + 2((1 + \gamma)D_{\max} + G_{\max}) \left(D_{\max} \sum_{i=0}^{k-1} \beta_i + 3G_{\max} \sum_{i=0}^{k-1} a_i \right) \\
&= 2((1 + \gamma)D_{\max} + G_{\max}) \left(D_{\max} \sum_{i=0}^{k-1} \beta_i + 3G_{\max} \sum_{i=0}^{k-1} a_i \right).
\end{aligned}
$$

When $k > \tau^{mix}(\kappa)$, we choose $\tau = \tau^* := \tau^{mix}(\kappa)$ and have

$$
\begin{aligned}
&\mathbb{E}[\xi(\theta_k; O_k)] \\
&\quad \leq \mathbb{E}[\xi(\theta_{k-\tau^*}; O_k)] + 2((1 + \gamma)D_{\max} + G_{\max}) \left(D_{\max} \sum_{i=k-\tau^*}^{k-1} \beta_i + 3G_{\max} \sum_{i=k-\tau^*}^{k-1} a_i \right) \\
&\quad \overset{(i)}{\leq} 4D_{\max}G_{\max}(\sigma\rho^{\tau^*}) + 2((1 + \gamma)D_{\max} + G_{\max}) \left(D_{\max} \sum_{i=k-\tau^*}^{k-1} \beta_i + 3G_{\max} \sum_{i=k-\tau^*}^{k-1} a_i \right) \\
&\quad \overset{(ii)}{\leq} 4D_{\max}G_{\max}\kappa + 2((1 + \gamma)D_{\max} + G_{\max}) \left(D_{\max} \sum_{i=k-\tau^*}^{k-1} \beta_i + 3G_{\max} \sum_{i=k-\tau^*}^{k-1} a_i \right) \\
&\quad \overset{(iii)}{\leq} 4D_{\max}G_{\max}\kappa + 2((1 + \gamma)D_{\max} + G_{\max}) \left(D_{\max}\tau^*\beta_{k-\tau^*} + 3G_{\max}\tau^* a_{k-\tau^*} \right),
\end{aligned}
$$

where (i) follows from (29), (ii) follows due to the definition of the mixing time, and (iii) follows because a_k, β_k are non-increasing.

C PROOF OF THEOREM 1

Recall that MomentumQ with linear function approximation updates as (12). Given the unique fixed point $\theta^\star$ and denoting $b_k + c_k = \beta_k$, we have

$$
\begin{aligned}
\|\theta_{k+1} - \theta^\star\|_2^2 &= \|\theta_k - \theta^\star + \beta_k(\theta_k - \theta_{k-1}) - a_k(1 + b_k)g_k + a_k b_k g_{k-1}\|_2^2 \\
&= \|\theta_k - \theta^\star\|_2^2 + \|\beta_k(\theta_k - \theta_{k-1}) - a_k(1 + b_k)g_k + a_k b_k g_{k-1}\|_2^2 \\
&\quad + 2\langle \theta_k - \theta^\star, \beta_k(\theta_k - \theta_{k-1}) - a_k(1 + b_k)g_k + a_k b_k g_{k-1}\rangle \\
&= \|\theta_k - \theta^\star\|_2^2 + \|\beta_k(\theta_k - \theta_{k-1}) - a_k(1 + b_k)g_k + a_k b_k g_{k-1}\|_2^2 \\
&\quad + 2\langle \theta_k - \theta^\star, \beta_k(\theta_k - \theta_{k-1}) + a_k b_k g_{k-1}\rangle - 2a_k(1 + b_k)\langle \theta_k - \theta^\star, g_k\rangle.
\end{aligned}
$$

Next, taking the expectation over all the randomness up to time step k on both sides, we have

$$
\begin{aligned}
&\mathbb{E}\,\|\theta_{k+1} - \theta^\star\|_2^2 \\
&= \mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 + \mathbb{E}\,\|\beta_k(\theta_k - \theta_{k-1}) - a_k(1 + b_k)g_k + a_k b_k g_{k-1}\|_2^2 \\
&\quad + 2\mathbb{E}\langle \theta_k - \theta^\star, \beta_k(\theta_k - \theta_{k-1}) + a_k b_k g_{k-1}\rangle - 2a_k(1 + b_k)\mathbb{E}\langle \theta_k - \theta^\star, g_k\rangle \\
&\overset{(i)}{\leq} \mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 + \mathbb{E}\,\|\beta_k(\theta_k - \theta_{k-1}) - a_k(1 + b_k)g_k + a_k b_k g_{k-1}\|_2^2 \\
&\quad + 2\beta_k \mathbb{E}\,\|\theta_k - \theta^\star\|_2\,\|\theta_k - \theta_{k-1}\|_2 + 2a_k b_k \mathbb{E}\,\|\theta_k - \theta^\star\|_2\,\|g_{k-1}\|_2 - 2a_k(1 + b_k)\mathbb{E}\langle \theta_k - \theta^\star, g_k\rangle \\
&\overset{(ii)}{\leq} \mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 + 3\beta_k^2 \mathbb{E}\,\|\theta_k - \theta_{k-1}\|_2^2 + 3a_k^2(1 + b_k)^2 \mathbb{E}\,\|g_k\|_2^2 + 3a_k^2 b_k^2 \mathbb{E}\,\|g_{k-1}\|_2^2 \\
&\quad + 2\beta_k \mathbb{E}\,\|\theta_k - \theta^\star\|_2\,\|\theta_k - \theta_{k-1}\|_2 + 2a_k b_k \mathbb{E}\,\|\theta_k - \theta^\star\|_2\,\|g_{k-1}\|_2 - 2a_k(1 + b_k)\mathbb{E}\langle \theta_k - \theta^\star, g_k\rangle \\
&\overset{(iii)}{\leq} \mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 + 3\beta_k^2 D_{\max}^2 + 3a_k^2(1 + b_k)^2 G_{\max}^2 + 3a_k^2 b_k^2 G_{\max}^2 \\
&\quad + 2\beta_k D_{\max}^2 + 2a_k b_k D_{\max} G_{\max} - 2a_k(1 + b_k)\mathbb{E}\langle \theta_k - \theta^\star, g_k\rangle \\
&\overset{(iv)}{\leq} \mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 + 5\beta_k D_{\max}^2 + 15a_k^2 G_{\max}^2 + 2a_k b_k D_{\max} G_{\max} - 2a_k(1 + b_k)\mathbb{E}\langle \theta_k - \theta^\star, g_k\rangle,
\end{aligned}
\tag{30}
$$

where (i) follows from Cauchy-Schwartz inequality, (ii) holds due to the fact $(x + y + z)^2 \leq 3x^2 + 3y^2 + 3z^2$, (iii) holds because of the boundedness of the parameter domain in Assumption 3 and because of Lemma 2, and (iv) follows since $b_k \leq \beta_k < 1$.

Since the samples are generated in a non-i.i.d. manner, we have

$$
\begin{aligned}
\mathbb{E}\left[(\theta_k - \theta^\star)^T g_k\right] &= \mathbb{E}\left[(\theta_k - \theta^\star)^T \bar{g}(\theta_k)\right] + \mathbb{E}\left[(\theta_k - \theta^\star)^T(g_k - \bar{g}(\theta_k))\right] \\
&= \mathbb{E}\left[(\theta_k - \theta^\star)^T \bar{g}(\theta_k)\right] + \mathbb{E}[\xi(\theta_k; O_k)].
\end{aligned}
\tag{31}
$$

Then, we continue to bound (30) and obtain

$$
\begin{aligned}
&\mathbb{E}\,\|\theta_{k+1} - \theta^\star\|_2^2 \\
&\leq \mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 + 5\beta_k D_{\max}^2 + 15a_k^2 G_{\max}^2 + 2a_k b_k D_{\max} G_{\max} - 2a_k(1 + b_k)\mathbb{E}\langle \theta_k - \theta^\star, g_k\rangle \\
&= \mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 + 5\beta_k D_{\max}^2 + 15a_k^2 G_{\max}^2 + 2a_k b_k D_{\max} G_{\max} \\
&\quad - 2a_k(1 + b_k)\mathbb{E}\langle \theta_k - \theta^\star, \bar{g}(\theta_k)\rangle - 2a_k(1 + b_k)\mathbb{E}[\xi(\theta_k; O_k)] \\
&\leq \mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 + 5\beta_k D_{\max}^2 + 15a_k^2 G_{\max}^2 + 2a_k b_k D_{\max} G_{\max} - 2a_k(1 + b_k)\delta \mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 \\
&\quad - 2a_k(1 + b_k)\mathbb{E}[\xi(\theta_k; O_k)] \\
&= (1 - 2a_k \delta(1 + b_k))\mathbb{E}\,\|\theta_k - \theta^\star\|_2^2 + 5\beta_k D_{\max}^2 + 15a_k^2 G_{\max}^2 + 2a_k b_k D_{\max} G_{\max} \\
&\quad - 2a_k(1 + b_k)\mathbb{E}[\xi(\theta_k; O_k)],
\end{aligned}
\tag{32}
$$

where the last inequality follows from Assumption 2.

We consider a constant stepsize $\alpha_k = \alpha$. For notational simplicity, we denote $f_k = 5\beta_k D_{\max}^2 + 15a_k^2 G_{\max}^2 + 2a_k b_k D_{\max} G_{\max}$, and $\zeta_k = -2a_k(1 + b_k)\mathbb{E}[\xi(\theta_k; O_k)]$. Then for $k > \tau^*$ we have

$$
\begin{aligned}
&\mathbb{E}\|\theta_{k+1} - \theta^\star\|_2^2 \\
&\leq (1 - 2\alpha\delta(1 + b_k))\mathbb{E}\|\theta_k - \theta^\star\|_2^2 + f_k + \zeta_k \\
&\leq \cdots \\
&\leq \prod_{i=0}^{k}(1 - 2\alpha\delta(1 + b_i))\|\theta_0 - \theta^\star\|_2^2 + \sum_{i=0}^{k} f_i \prod_{j=i+1}^{k}(1 - 2\alpha\delta(1 + b_j)) \\
&\quad + \sum_{i=\tau^*+1}^{k} \zeta_i \prod_{j=i+1}^{k}(1 - 2\alpha\delta(1 + b_j)) + \sum_{i=0}^{\tau^*} \zeta_i \prod_{j=i+1}^{k}(1 - 2\alpha\delta(1 + b_j)) \\
&\leq \prod_{i=0}^{k}(1 - 2\alpha\delta(1 + b_i))\|\theta_0 - \theta^\star\|_2^2 + \sum_{i=0}^{k} f_i(1 - 2\alpha\delta)^{k-i} \\
&\quad + \sum_{i=\tau^*+1}^{k} \zeta_i(1 - 2\alpha\delta)^{k-i} + \sum_{i=0}^{\tau^*} \zeta_i(1 - 2\alpha\delta)^{k-i},
\end{aligned}
$$

where the last inequality follows because $b_k > 0$, $\forall k$. Further, we bound the term $\sum_{i=0}^{k}(1 - 2\alpha\delta)^{k-i} f_i$ as

$$
\begin{aligned}
&\sum_{i=0}^{k}(1 - 2\delta\alpha)^{k-i} f_i \\
&= 5D_{\max}^2 \sum_{i=0}^{k}(1 - 2\delta\alpha)^{k-i}\beta_i + 15\alpha^2 G_{\max}^2 \sum_{i=0}^{k}(1 - 2\delta\alpha)^{k-i} + 2\alpha D_{\max} G_{\max} \sum_{i=0}^{k}(1 - 2\delta\alpha)^{k-i} b_i \\
&\leq 15\alpha^2 G_{\max}^2 \sum_{i=0}^{k}(1 - 2\delta\alpha)^{k-i} + (5D_{\max}^2 + 2\alpha D_{\max} G_{\max})\sum_{i=0}^{k}(1 - 2\delta\alpha)^{k-i}\beta_i \\
&\leq \frac{15\alpha G_{\max}^2}{2\delta} + (5D_{\max}^2 + 2\alpha D_{\max} G_{\max})\beta(1 - 2\delta\alpha)^k \sum_{i=0}^{k}\left(\frac{\lambda}{1 - 2\delta\alpha}\right)^i \\
&\overset{\text{(i)}}{\leq} \frac{15\alpha G_{\max}^2}{2\delta} + (5D_{\max}^2 + 2\alpha D_{\max} G_{\max})\beta(1 - 2\delta\alpha)^k \frac{1}{1 - 2\delta\alpha - \lambda},
\end{aligned}
\tag{33}
$$

where (i) follows from $\alpha < \frac{1-\lambda}{2\delta}$. It remains to bound the last two tail terms. From Lemma 1, we obtain

$$
\zeta_i = \begin{cases}
2\alpha(1 + b_i)\left(\eta_1 \sum_{i=1}^{k-1}\beta_i + \eta_2 \sum_{i=1}^{k-1} a_i\right) \leq 4\alpha\left(\eta_1\tau^*\beta + \eta_2\tau^*\alpha\right), & i \leq \tau^*; \\
4\alpha\left(4D_{\max}G_{\max}\kappa + \eta_1\tau^*\beta_{i-\tau^*} + \eta_2\tau^*\alpha\right), & i > \tau^*,
\end{cases}
$$

where $\eta_1 = 2D_{\max}((1 + \gamma)D_{\max} + G_{\max}), \eta_2 = 6G_{\max}((1 + \gamma)D_{\max} + G_{\max})$. Then we obtain

$$
\begin{aligned}
&\sum_{i=\tau^*+1}^{k}\zeta_i(1 - 2\alpha\delta)^{k-i} + \sum_{i=0}^{\tau^*}\zeta_i(1 - 2\alpha\delta)^{k-i} \\
&\leq 4\eta_2\tau^*\alpha^2 \sum_{i=0}^{k}(1 - 2\alpha\delta)^{k-i} + 4\alpha\eta_1\tau^*\beta \sum_{i=0}^{\tau^*}(1 - 2\alpha\delta)^{k-i} \\
&\quad + 16D_{\max}G_{\max}\kappa\alpha \sum_{i=\tau^*+1}^{k}(1 - 2\alpha\delta)^{k-i} + 4\alpha\eta_1\tau^* \sum_{i=\tau^*+1}^{k}\beta_{i-\tau^*}(1 - 2\alpha\delta)^{k-i} \\
&\leq \frac{2\eta_2\tau^*\alpha}{\delta} + \frac{2\eta_1\tau^*\beta}{\delta}(1 - 2\alpha\delta)^{k-\tau^*} + \frac{8D_{\max}G_{\max}\kappa}{\delta} + 4\alpha\beta\eta_1\tau^* \sum_{i=\tau^*+1}^{k}\lambda^{i-\tau^*}(1 - 2\alpha\delta)^{k-i}
\end{aligned}
$$

$$= \frac{2\eta_2\tau^\star\alpha}{\delta} + \frac{2\eta_1\tau^\star\beta}{\delta}(1-2\alpha\delta)^{k-\tau^\star} + \frac{8D_{\max}G_{\max}\kappa}{\delta}$$

$$+ 4\alpha\beta\eta_1\tau^\star(1-2\alpha\delta)^{k-\tau^\star}\sum_{i=\tau^\star+1}^{k}\left(\frac{\lambda}{1-2\alpha\delta}\right)^{i-\tau^\star}$$

$$\le \frac{2\eta_2\tau^\star\alpha}{\delta} + \frac{2\eta_1\tau^\star\beta}{\delta}(1-2\alpha\delta)^{k-\tau^\star} + \frac{8D_{\max}G_{\max}\kappa}{\delta} + \frac{4\alpha\beta\eta_1\tau^\star\lambda}{1-2\alpha\delta-\lambda}(1-2\alpha\delta)^{k-\tau^\star},$$

where the last inequality follows due to the fact that $\alpha < \frac{1-\lambda}{2\delta}$. Thus, we can conclude that

$$\mathbb{E}\|\theta_{k+1}-\theta^\star\|_2^2$$

$$\le \prod_{i=0}^{k}(1-2\alpha\delta(1+b_i))\|\theta_0-\theta^\star\|_2^2 + \sum_{i=0}^{k}f_i(1-2\alpha\delta)^{k-i}$$

$$+ \sum_{i=\tau^\star+1}^{k}\zeta_i(1-2\alpha\delta)^{k-i} + \sum_{i=0}^{\tau^\star}\zeta_i(1-2\alpha\delta)^{k-i}$$

$$\le \prod_{i=0}^{k}(1-2\alpha\delta(1+b_i))\|\theta_0-\theta^\star\|_2^2 + \frac{15\alpha G_{\max}^2}{2\delta} + \frac{\beta(5D_{\max}^2+2\alpha D_{\max}G_{\max})(1-2\delta\alpha)^k}{1-2\delta\alpha-\lambda}$$

$$+ \frac{2\eta_2\tau^\star\alpha}{\delta} + \frac{2\eta_1\tau^\star\beta}{\delta}(1-2\alpha\delta)^{k-\tau^\star} + \frac{8D_{\max}G_{\max}\kappa}{\delta} + \frac{4\alpha\beta\eta_1\tau^\star\lambda}{1-2\alpha\delta-\lambda}(1-2\alpha\delta)^{k-\tau^\star}$$

$$\le \prod_{i=0}^{k}(1-2\alpha\delta(1+b_i))\|\theta_0-\theta^\star\|_2^2 + \frac{15\alpha G_{\max}^2}{2\delta} + \frac{2\eta_2\tau^\star\alpha}{\delta} + \frac{8D_{\max}G_{\max}\kappa}{\delta}$$

$$+ \beta\left(\frac{2\eta_1\tau^\star}{\delta} + \frac{5D_{\max}^2+2\alpha D_{\max}G_{\max}+4\alpha\eta_1\tau^\star\lambda}{1-2\delta\alpha-\lambda}\right)(1-2\delta\alpha)^{k-\tau^\star}.$$

D   PROOF OF THEOREM 2

Before proving this theorem, we introduce two lemmas of series sum that will help to streamline the presentation.

Lemma 5. *Let $a_k = \frac{\alpha}{\sqrt{k}}$ and $\beta_k = \beta\lambda^k$ with $\alpha > 0, \beta, \lambda \in (0,1)$ for $k = 1, 2, \ldots$. Then*

$$\sum_{k=1}^{T}\frac{\beta_k}{a_k} \le \frac{\beta}{\alpha(1-\lambda)^2}. \tag{34}$$

Proof. The proof is based on taking the standard sum of geometric sequences as follows:

$$\sum_{k=1}^{T}\frac{\beta_k}{a_k} = \sum_{k=1}^{T}\frac{\beta\lambda^k\sqrt{k}}{\alpha} \le \sum_{k=1}^{T}\frac{\beta\lambda^k k}{\alpha} = \frac{\beta}{\alpha(1-\lambda)}\left(\sum_{k=1}^{T}\lambda^k - T\lambda^T\right) \le \frac{\beta}{\alpha(1-\lambda)^2}.$$

$\square$

Lemma 6. *Let $a_k = \frac{\alpha}{\sqrt{k}}$. Then*

$$\sum_{k=1}^{T}a_k \le 2\alpha\sqrt{T}. \tag{35}$$

Proof. We use the comparison principle to bound the series sum as follows:

$$\sum_{k=1}^{T}a_k = \sum_{k=1}^{T}\frac{\alpha}{\sqrt{k}} \le \int_{1}^{T+1}\frac{\alpha}{\sqrt{t-1}}dt = 2\alpha\sqrt{t-1}\big|_1^{T+1} = 2\alpha\sqrt{T}.$$

$\square$

The proof of Theorem 2 is partially similar to that of Theorem 1. The steps are the same until (32), where we have

$$
\begin{aligned}
\mathbb{E}\left\|\theta_{k+1}-\theta^{\star}\right\|_{2}^{2} \\
\leq \mathbb{E}\left\|\theta_{k}-\theta^{\star}\right\|_{2}^{2}+5\beta_{k}D_{\max}^{2}+15a_{k}^{2}G_{\max}^{2}+2a_{k}b_{k}D_{\max}G_{\max}-2a_{k}(1+b_{k})\delta\mathbb{E}\left\|\theta_{k}-\theta^{\star}\right\|_{2}^{2} \\
-2a_{k}(1+b_{k})\mathbb{E}[\xi(\theta_{k};O_{k})].
\end{aligned}
$$

Then we continue the proof with rearranging the previous inequality:

$$
\begin{aligned}
2\delta\mathbb{E}\left\|\theta_{k}-\theta^{\star}\right\|_{2}^{2} \\
\leq 2(1+b_{k})\delta\mathbb{E}\left\|\theta_{k}-\theta^{\star}\right\|_{2}^{2} \\
\leq \frac{\mathbb{E}\left\|\theta_{k}-\theta^{\star}\right\|_{2}^{2}-\mathbb{E}\left\|\theta_{k+1}-\theta^{\star}\right\|_{2}^{2}}{a_{k}}+\frac{5\beta_{k}}{a_{k}}D_{\max}^{2}+15a_{k}G_{\max}^{2}+2b_{k}D_{\max}G_{\max}+4|\mathbb{E}[\xi(\theta_{k};O_{k})]|.
\end{aligned}
$$

Then we sum over time step k from 1 to $T(T>\tau^{*})$ and obtain

$$
\begin{aligned}
2\delta\sum_{k=1}^{T}\mathbb{E}\left\|\theta_{k}-\theta^{\star}\right\|_{2}^{2} \\
\leq\sum_{k=1}^{T}\frac{\mathbb{E}\left\|\theta_{k}-\theta^{\star}\right\|_{2}^{2}-\mathbb{E}\left\|\theta_{k+1}-\theta^{\star}\right\|_{2}^{2}}{a_{k}}+4\sum_{k=1}^{\tau^{*}}|\mathbb{E}[\xi(\theta_{k};O_{k})]|+4\sum_{k=\tau^{*}+1}^{T}|\mathbb{E}[\xi(\theta_{k};O_{k})]| \\
+5D_{\max}^{2}\sum_{k=1}^{T}\frac{\beta_{k}}{a_{k}}+15G_{\max}^{2}\sum_{k=1}^{T}a_{k}+2D_{\max}G_{\max}\sum_{k=1}^{T}b_{k} \\
=\frac{\left\|\theta_{1}-\theta^{\star}\right\|_{2}^{2}}{a_{1}}+\sum_{k=2}^{T}\mathbb{E}\left\|\theta_{k}-\theta^{\star}\right\|_{2}^{2}\left(\frac{1}{a_{k}}-\frac{1}{a_{k-1}}\right)-\frac{\mathbb{E}\left\|\theta_{T+1}-\theta^{\star}\right\|_{2}^{2}}{a_{T+1}} \\
+5D_{\max}^{2}\sum_{k=1}^{T}\frac{\beta_{k}}{a_{k}}+15G_{\max}^{2}\sum_{k=1}^{T}a_{k}+2D_{\max}G_{\max}\sum_{k=1}^{T}b_{k} \\
+4\sum_{k=1}^{\tau^{*}}|\mathbb{E}[\xi(\theta_{k};O_{k})]|+4\sum_{k=\tau^{*}+1}^{T}|\mathbb{E}[\xi(\theta_{k};O_{k})]| \\
\overset{(i)}{\leq}\frac{\left\|\theta_{1}-\theta^{\star}\right\|_{2}^{2}}{a_{1}}+D_{\max}^{2}\sum_{k=2}^{T}\left(\frac{1}{a_{k}}-\frac{1}{a_{k-1}}\right) \\
+5D_{\max}^{2}\sum_{k=1}^{T}\frac{\beta_{k}}{a_{k}}+15G_{\max}^{2}\sum_{k=1}^{T}a_{k}+2D_{\max}G_{\max}\sum_{k=1}^{T}b_{k} \\
+4\sum_{k=1}^{\tau^{*}}|\mathbb{E}[\xi(\theta_{k};O_{k})]|+4\sum_{k=\tau^{*}+1}^{T}|\mathbb{E}[\xi(\theta_{k};O_{k})]| \\
\overset{(ii)}{\leq}\frac{D_{\max}^{2}}{\alpha_{T}}+5D_{\max}^{2}\sum_{k=1}^{T}\frac{\beta_{k}}{a_{k}}+15G_{\max}^{2}\sum_{k=1}^{T}a_{k}+2D_{\max}G_{\max}\sum_{k=1}^{T}\beta_{k} \\
+4\sum_{k=1}^{\tau^{*}}|\mathbb{E}[\xi(\theta_{k};O_{k})]|+4\sum_{k=\tau^{*}+1}^{T}|\mathbb{E}[\xi(\theta_{k};O_{k})]| \\
\overset{(iii)}{\leq}\frac{D_{\max}^{2}\sqrt{T}}{\alpha}+\frac{5\beta D_{\max}^{2}}{\alpha(1-\lambda)^{2}}+30\alpha G_{\max}^{2}\sqrt{T}+\frac{2D_{\max}G_{\max}\beta\lambda}{1-\lambda} \\
+4\sum_{k=1}^{\tau^{*}}|\mathbb{E}[\xi(\theta_{k};O_{k})]|+4\sum_{k=\tau^{*}+1}^{T}|\mathbb{E}[\xi(\theta_{k};O_{k})]|,
\end{aligned}
$$

where (i) follows from Assumption 3 and the fact that $\alpha_{k}<\alpha_{k-1}$, and $\mathbb{E}\left\|\theta_{T+1}-\theta^{\star}\right\|^{2}/a_{T+1}>0$, (ii) holds due to Assumption 3, and (iii) follows from Lemmas 2, 5, and 6.

It remains to bound $4 \sum_{k=1}^{\tau^*} |\mathbb{E}[\xi(\theta_k; O_k)]| + 4 \sum_{k=\tau^*+1}^{T} |\mathbb{E}[\xi(\theta_k; O_k)]|$. We bound the tail term by using Lemma 1. For simplicity, in the following we denote

$$\eta_1 = 2D_{\max}((1+\gamma)D_{\max} + G_{\max}), \quad \eta_2 = 6G_{\max}((1+\gamma)D_{\max} + G_{\max}).$$

Following from Lemma 1, we have

$$\sum_{k=1}^{\tau^*} |\mathbb{E}[\xi(\theta_k; O_k)]| \leq \sum_{k=1}^{\tau^*} \eta_1 \sum_{i=1}^{k-1} \beta_i + \sum_{k=1}^{\tau^*} \eta_2 \sum_{i=1}^{k-1} a_i$$

$$\leq \tau^* \eta_1 \sum_{k=1}^{T} \beta_k + \tau^* \eta_2 \sum_{k=1}^{T} a_k$$

$$\leq \frac{\tau^* \eta_1 \beta \lambda}{1 - \lambda} + 2\tau^* \eta_2 \alpha \sqrt{T}.$$

Similarly, we obtain

$$\sum_{k=\tau^*+1}^{T} |\mathbb{E}[\xi(\theta_k; O_k)]| \leq \sum_{k=\tau^*+1}^{T} (4D_{\max}G_{\max}\kappa + \eta_1 \tau^* \beta_{k-\tau^*} + \eta_2 \tau^* a_{k-\tau^*})$$

$$\leq 4D_{\max}G_{\max}\kappa T + \tau^* \eta_1 \sum_{k=1}^{T-\tau^*} \beta_k + \tau^* \eta_2 \sum_{k=1}^{T-\tau^*} a_k$$

$$\leq 4D_{\max}G_{\max}\kappa T + \frac{\tau^* \eta_1 \beta \lambda}{1 - \lambda} + 2\tau^* \eta_2 \alpha \sqrt{T}.$$

Thus, we have

$$2\delta \sum_{k=1}^{T} \mathbb{E} \|\theta_k - \theta^\star\|_2^2$$

$$\leq \frac{D_{\max}^2 \sqrt{T}}{\alpha} + \frac{5\beta D_{\max}^2}{\alpha(1-\lambda)^2} + 30\alpha G_{\max}^2 \sqrt{T} + \frac{2D_{\max}G_{\max}\beta\lambda}{1-\lambda}$$

$$+ 4\sum_{k=1}^{\tau^*} |\mathbb{E}[\xi(\theta_k; O_k)]| + 4\sum_{k=\tau^*+1}^{T} |\mathbb{E}[\xi(\theta_k; O_k)]|$$

$$\leq \frac{D_{\max}^2 \sqrt{T}}{\alpha} + \frac{5\beta D_{\max}^2}{\alpha(1-\lambda)^2} + 30\alpha G_{\max}^2 \sqrt{T} + \frac{2D_{\max}G_{\max}\beta\lambda}{1-\lambda}$$

$$+ 16D_{\max}G_{\max}\kappa T + \frac{8\tau^* \eta_1 \beta \lambda}{1-\lambda} + 16\tau^* \eta_2 \alpha \sqrt{T}.$$

Finally, we apply Jensen's inequality and complete the proof as

$$\mathbb{E} \|\theta_{\text{out}} - \theta^\star\|_2^2 = \mathbb{E} \left\| \frac{1}{T} \sum_{k=1}^{T} \theta_k - \theta^\star \right\|_2^2 \leq \frac{1}{T} \sum_{k=1}^{T} \mathbb{E} \|\theta_k - \theta^\star\|_2^2$$

$$\leq \frac{D_{\max}^2/\alpha + 30\alpha G_{\max}^2 + 16\tau^* \alpha \eta_2}{2\delta\sqrt{T}} + \frac{8D_{\max}G_{\max}\kappa}{\delta}$$

$$+ \frac{1}{T} \left[\frac{5\beta D_{\max}^2}{2\alpha\delta(1-\lambda)^2} + \frac{D_{\max}G_{\max}\beta\lambda + 4\tau^* \eta_1 \beta \lambda}{\delta(1-\lambda)} \right].$$

E PROOF OF PROPOSITION 1

Proof. For convenience, we denote $\mathcal{M}Q_k(y_k) := \max_{u \in U(y_k)} Q_k(y_k, u)$, then $\hat{\mathcal{T}}_k Q_k = R + \mathcal{M}Q_k(y_k)$ and $\hat{\mathcal{T}}_k Q_{k-1} = R + \mathcal{M}Q_{k-1}(y_k)$. If $k = 0$, we have from (20) that

$$\|\mathcal{D}_0[Q_0, Q_{-1}]\| = \left\| \hat{\mathcal{T}}_0 Q_0 \right\| \leq \|R\| + \gamma \|\mathcal{M}Q_0(y_0)\|$$

$$\leq R_{\max} + \gamma V_{\max}.$$

Now, considering $k \geq 1$ we have

$$\|\mathcal{D}_k [Q_k, Q_{k-1}]\|$$

$$\overset{(i)}{\leq} \|R\| + \gamma \|(1+b_k)\mathcal{M}Q_k - b_k \mathcal{M}Q_{k-1}\|$$

$$\leq R_{\max} + \gamma \|(1+b_k)\mathcal{M}(Q_{k-1} - \alpha_{k-1}Q_{k-2} + \alpha_{k-1}\mathcal{D}_{k-1}[Q_{k-1}, Q_{k-2}]) - b_k \mathcal{M}Q_{k-1}\|$$

$$\overset{(ii)}{\leq} R_{\max} + \gamma \|Q_{k-1}\| + \gamma |1+b_k|a_{k-1}\|Q_{k-2}\| + \gamma |1+b_k|\alpha_{k-1}\|\mathcal{D}_{k-1}[Q_{k-1}, Q_{k-2}]\|, \tag{36}$$

where (i) follows from the triangle inequality and (ii) follows from the definition of the infinity norm.

To proceed to bound (36), we consider two cases. If $k < \frac{m}{2}$, there are at most a finite number of $\mathcal{D}_k$'s, which are obviously bounded. If $k \geq \frac{m}{2}$, we have $|1+b_k|a_{k-1} = \frac{|k-m|}{k} \leq 1$. It follows from (36) that

$$\|\mathcal{D}_k [Q_k, Q_{k-1}]\|$$

$$\leq R_{\max} + \gamma \|Q_{k-1}\| + \gamma \|Q_{k-2}\| + \gamma \|\mathcal{D}_{k-1}[Q_{k-1}, Q_{k-2}]\|$$

$$\overset{(i)}{\leq} R_{\max} + 2\gamma V_{\max} + \gamma \|\mathcal{D}_{k-1}[Q_{k-1}, Q_{k-2}]\|$$

$$\overset{(ii)}{\leq} R_{\max} \sum_{i=0}^{k-\lfloor m/2 \rfloor} \gamma^i + 2V_{\max} \sum_{i=1}^{k-\lfloor m/2 \rfloor} \gamma^i + \gamma^{k-\lfloor m/2 \rfloor} \|\mathcal{D}_{\lfloor m/2 \rfloor}[Q_{\lfloor m/2 \rfloor}, Q_{\lfloor m/2 \rfloor - 1}]\| \tag{37}$$

where $\lfloor x \rfloor$ denotes the largest integer that is no larger than x. Note that (i) follows from the boundedness of Q_k (Assumption 5), and (ii) follows from applying (i) to $\mathcal{D}_t$ for $t = k-1, k-2, \ldots, \lfloor m/2 \rfloor + 1$ iteratively. Since $\gamma < 1$, the first two items in (ii) are bounded. Obviously, the third item is also bounded. Therefore, there exists some constant $\bar{D}$, such that $\|\mathcal{D}_k\| \leq \bar{D}$, $\forall k \geq 0$.

The bound on ϵ_k follows directly from its definition as

$$\|\epsilon_k\| = \|\mathbb{E}_P (\mathcal{D}_k [Q_k, Q_{k-1}] (x,u)|\mathcal{F}_{k-1}) - \mathcal{D}_k [Q_k, Q_{k-1}]\|$$

$$\leq 2 \|\mathcal{D}_k [Q_k, Q_{k-1}]\| \leq 2\bar{D}.$$

Thus we conclude our proof. $\qquad\square$

F PROOF OF THEOREM 3

We first prove two lemmas that will be useful for establishing the main results. The first lemma derives the dynamics of Q_k in terms of E_k.

Lemma 7. *Consider MomentumQ as in Algorithm 1. For any $k \geq 1$, we have*

$$Q_k = \frac{1}{k}(Q_{k-1} - Q_0 + (k-m-1)\mathcal{T}Q_{k-1}) + \frac{1}{k}((m+1)\mathcal{T}Q_0 - E_{k-1}). \tag{38}$$

Proof. We prove the lemma by substituting the learning rates a_k, b_k, c_k in Algorithm 1 and using induction. From (19), we see that $Q_1 = \hat{\mathcal{T}}_1 Q_0 = \mathcal{T}Q_0 - E_0$, Thus (38) holds when $k = 1$. Now assume (38) holds for a certain integer $k > 1$ we prove it also holds for $k+1$. To see this, we rewrite (19) as

$$Q_{k+1} = \frac{1}{k+1}Q_k - \frac{1}{k+1}Q_{k-1} + \frac{k}{k+1}Q_k + \frac{1}{k+1}\left[(k-m)\hat{\mathcal{T}}_k Q_k - (k-m-1)\hat{\mathcal{T}}_k Q_{k-1}\right]$$

$$= \frac{1}{k+1}Q_k - \frac{1}{k+1}Q_{k-1} + \frac{1}{k+1}(Q_{k-1} - Q_0 + (k-m-1)\mathcal{T}Q_{k-1}$$

$$+ (m+1)\mathcal{T}Q_0 - E_{k-1}) + \frac{1}{k+1}\left[(k-m)\hat{\mathcal{T}}_k Q_k - (k-m-1)\hat{\mathcal{T}}_k Q_{k-1}\right]$$

$$= \frac{1}{k+1}Q_k - \frac{1}{k+1}Q_{k-1} + \frac{1}{k+1}(Q_{k-1} - Q_0 + (k-m-1)\mathcal{T}Q_{k-1}$$

$$+ (m+1)\mathcal{T}Q_0 - E_{k-1}) + \frac{1}{k+1}\left[(k-m)\mathcal{T}Q_k - (k-m-1)\mathcal{T}Q_{k-1} - \epsilon_k\right]$$

$$= \frac{1}{k+1}(Q_k - Q_0 + (k-m)\mathcal{T}Q_k + (m+1)\mathcal{T}Q_0 - E_k),$$

which shows that (38) holds for $k + 1$. Therefore, by induction (38) holds for all $k \geq 1$. $\qquad\square$

The second lemma derives the propagation of the approximation errors ϵ_k in the process of Q-function iteration, which can be proved conveniently using Lemma 7.

Lemma 8. *Suppose Assumption 5 holds and $m \geq \frac{1}{\gamma}$ as in Algorithm 1. Then for all $k \geq m + 1$, we have*

$$\|Q^\star - Q_k\| \leq \frac{\tilde{h} V_{\max}}{k(1 - \gamma)} + \frac{1}{k} \sum_{i=0}^{k - \lfloor m \rfloor - 2} \gamma^i \|E_{k-i}\|, \tag{39}$$

where $\tilde{h} = 2\gamma(m + \lfloor m \rfloor + 2) + 2$.

Proof. For $k \geq m + 1$, expand Q_k using (38) in Lemma 7 iteratively, yielding

$$
\begin{aligned}
\|Q^\star - Q_k\| &= \frac{1}{k} \|Q_0 - Q_{k-1} + (k - m - 1)(\mathcal{T}Q^\star - \mathcal{T}Q_{k-1}) + (m + 1)(\mathcal{T}Q^\star - \mathcal{T}Q_0) + E_k\| \\
&\overset{\text{(i)}}{\leq} \frac{\gamma(k - m - 1) + 1}{k} \|Q^\star - Q_{k-1}\| + \frac{\gamma(m + 1) + 1}{k} \|Q^\star - Q_0\| + \frac{\|E_k\|}{k} \\
&\overset{\text{(ii)}}{\leq} \frac{\gamma(k - 1)}{k} \|Q^\star - Q_{k-1}\| + \frac{2h}{k} V_{\max} + \frac{\|E_k\|}{k} \\
&\overset{\text{(iii)}}{\leq} \frac{\gamma^{k - \lfloor m \rfloor - 1}}{k} (\lfloor m \rfloor + 1) \|Q^\star - Q_{\lfloor m \rfloor + 1}\| + \frac{2h V_{\max}}{k} \sum_{i=0}^{k - \lfloor m \rfloor - 2} \gamma^i + \sum_{i=0}^{k - \lfloor m \rfloor - 2} \frac{\gamma^i}{k} \|E_{k-i}\| \\
&\leq 2 \frac{\gamma(\lfloor m \rfloor + 1) + h}{k(1 - \gamma)} V_{\max} + \frac{1}{k} \sum_{i=0}^{k - \lfloor m \rfloor - 2} \gamma^i \|E_{k-i}\|,
\end{aligned}
$$

where (i) follows from the triangle inequality and the contraction property (3), (ii) follows from Assumption 5 and because $m \geq \frac{1}{\gamma}$, $h = \gamma(m + 1) + 1$, and (iii) follows from applying (ii) to $\|Q^\star - Q_t\|$ for $t = k - 1, k - 2, \ldots, \lfloor m \rfloor + 2$ iteratively. Then (39) follows from the definition of $\tilde{h}$. $\qquad\square$

Lemma 9. *(Maximal Hoeffding-Azuma Inequality) [Alon and Spencer, 2008, Chapter 7]*
Let $\{M_1, M_2, \ldots, M_T\}$ be a martingale difference sequence with respect to a sequence of random variables $\{X_1, X_2, \ldots, X_T\}$ (i.e. $\mathbb{E}(M_{k+1} | X_1, X_2, \ldots, X_k) = 0, \forall 1 \leq k \leq T$) and uniformly bounded by $\bar{M} > 0$ almost surely. If we define $S_k = \sum_{i=1}^k M_i$, then for any $\varepsilon > 0$, we have

$$\mathbb{P}\left(\max_{1 \leq k \leq T} S_k > \varepsilon\right) \leq \exp\left(\frac{-\varepsilon^2}{2T\bar{M}^2}\right).$$

Now we are ready to prove the main results of Theorem 3.

Proof of Theorem 3. The proof applies Lemma 8 and the Maximal Hoeffding-Azuma Inequality (Lemma 9).

Applying Lemma 8 with $k = T$, we obtain

$$\|Q^\star - Q_T\| \leq \frac{\tilde{h} V_{\max}}{T(1 - \gamma)} + \frac{1}{T} \sum_{i=0}^{T - \lfloor m \rfloor - 2} \gamma^i \|E_{T-i}\|.$$

It suffices to bound the second term. For convenience, we denote $K = T - \lfloor m \rfloor - 2$. Observe that

$$\frac{1}{T} \sum_{i=0}^{K} \gamma^i \|E_{T-i}\| \leq \frac{1}{T} \sum_{i=0}^{K} \gamma^i \max_{0 \leq i \leq K} \|E_{T-i}\| \leq \frac{\max_{0 \leq i \leq K} \|E_{T-i}\|}{(1 - \gamma)T}. \tag{40}$$

In remains to bound $\max_{0 \leq i \leq K} \|E_{T-i}\|$. Notice that $\max_{0 \leq i \leq K} \|E_{T-i}\| = \max_{(x,u)} \max_{0 \leq i \leq K} |E_{T-i}(x,u)|$. For a given (x,u) and $\varepsilon > 0$, we have

$$
\begin{aligned}
&\mathbb{P}\left(\max_{0 \leq i \leq K} |E_{T-i}(x,u)| > \varepsilon\right) \\
&= \mathbb{P}\left(\left\{\max_{0 \leq i \leq K}(E_{T-i}(x,u)) > \varepsilon\right\} \bigcup \left\{\max_{0 \leq i \leq K}(-E_{T-i}(x,u)) > \varepsilon\right\}\right) \\
&= \mathbb{P}\left(\max_{0 \leq i \leq K}(E_{T-i}(x,u)) > \varepsilon\right) + \mathbb{P}\left(\max_{0 \leq i \leq K}(-E_{T-i}(x,u)) > \varepsilon\right),
\end{aligned}
\tag{41}
$$

where $\bar{D}$ is specified in Proposition 1. Since $\{\epsilon_k(x,u)\}_{k \geq 0}$ is a martingale difference sequence with respect to the filtration $\mathcal{F}_k$ as defined previously, we apply the Maximal Hoeffding-Azuma inequality (see Lemma 9) and obtain

$$
\mathbb{P}\left(\max_{0 \leq i \leq K}(E_{T-i}(x,u)) > \varepsilon\right) \leq \exp\left(\frac{-\varepsilon^2}{8(K+1)\bar{D}^2}\right),
$$

and

$$
\mathbb{P}\left(\max_{0 \leq i \leq K}(-E_{T-i}(x,u)) > \varepsilon\right) \leq \exp\left(\frac{-\varepsilon^2}{8(K+1)\bar{D}^2}\right).
$$

Then we further bound (41) as

$$
\mathbb{P}\left(\max_{0 \leq i \leq K} |E_{T-i}(x,u)| > \varepsilon\right) \leq 2\exp\left(\frac{-\varepsilon^2}{8(K+1)\bar{D}^2}\right).
$$

Since we consider a finite state-action space where the number of state-action pairs is defined by n, we use the union bound to obtain

$$
\mathbb{P}\left(\max_{0 \leq i \leq K} \|E_{T-i}\| > \varepsilon\right) \leq 2n\exp\left(\frac{-\varepsilon^2}{8(K+1)\bar{D}^2}\right).
$$

Letting $\delta = 2n\exp\left(\frac{-\varepsilon^2}{8(K+1)\bar{D}^2}\right)$, and we have

$$
\mathbb{P}\left(\max_{0 \leq i \leq K} \|E_{T-i}\| \leq \bar{D}\sqrt{8(K+1)\log\frac{2n}{\delta}}\right) \geq 1 - \delta,
$$

where $K = T - \lfloor m \rfloor - 2$. By substituting the above bound into (40) yields the desired result. $\qquad\square$

G PROOF OF COROLLARY 1

In Theorem 3, take $\delta = \frac{1}{T^2}$, and denote by A_T the event "inequality (24) holds". Then the conclusion of Theorem 3 becomes $\mathbb{P}[A_T] \geq 1 - \frac{1}{T^2}$, or equivalently, $\mathbb{P}[A_T^c] \leq \frac{1}{T^2}$, for all $T > m$, where the superscript c meaning taking the set complement. It follows that $\sum_{T=m+1}^{\infty} \mathbb{P}[A_T^c] \leq \sum_{T=m+1}^{\infty} \frac{1}{T^2} < \infty$. By the Borel–Cantelli lemma (see, for example, Chapter 2.3, Theorem 2.3.1 of [Durrett, 2019]), this implies $\mathbb{P}[A_T^c \ i.o.] = 0$, where i.o. stands for infinitely often. This is equivalent to the statement that Q_T converges to Q^* almost surely at a rate of at least $\mathcal{O}(\frac{\sqrt{(T-\lfloor m \rfloor -1)\log nT}}{(1-\gamma)^2 T})$, where note that in (24) the constant $\bar{D}$ is proportional to $\frac{1}{1-\gamma}$. Using the $\tilde{\mathcal{O}}$ notation which ignores the $\log T$ factor, the order of the convergence rate can be written as $\tilde{\mathcal{O}}(\frac{\sqrt{(T-\lfloor m \rfloor -1)\log n}}{(1-\gamma)^2 T})$. Thus it completes the proof. $\qquad\square$