Optimal variance-reduced stochastic approximation in Banach spaces

Wenlong Mou[*,◊]   Koulik Khamaru[*,†]   Martin J. Wainwright[◊,†,◦]
Peter L. Bartlett[◊,†]   Michael I. Jordan[◊,†]

Department of Electrical Engineering and Computer Sciences[◊]
Department of Statistics[†]
UC Berkeley

Department of Electrical Engineering and Computer Sciences[◦]
Department of Mathematics[◦]
Massachusetts Institute of Technology

November 30, 2022

Abstract

We study the problem of estimating the fixed point of a contractive operator defined on a separable Banach space. Focusing on a stochastic query model that provides noisy evaluations of the operator, we analyze a variance-reduced stochastic approximation scheme, and establish non-asymptotic bounds for both the operator defect and the estimation error, measured in an arbitrary semi-norm. In contrast to worst-case guarantees, our bounds are instance-dependent, and achieve the local asymptotic minimax risk non-asymptotically. For linear operators, contractivity can be relaxed to multi-step contractivity, so that the theory can be applied to problems like average reward policy evaluation problem in reinforcement learning. We illustrate the theory via applications to stochastic shortest path problems, two-player zero-sum Markov games, as well as average-reward policy evaluation.

1   Introduction

In this paper, we consider a class of stochastic fixed-point problems defined in Banach spaces. In particular, let $\mathbb{V}$ be a separable Banach space with its associated norm $\|\cdot\|$, and suppose that $\mathbf{h} : \mathbb{V} \to \mathbb{V}$ is an operator on the Banach space. Of interest to us are solutions θ^* to the fixed-point equation

$$\theta^* = \mathbf{h}(\theta^*). \tag{1}$$

When the operator $\mathbf{h}$ is contractive, the Banach fixed point theorem (e.g., [DG03]) ensures the existence and uniqueness of the fixed point. The bulk of our analysis focuses on this contractive case, but we also allow for weaker multi-stage contraction in certain settings.

---

[*] Wenlong Mou and Koulik Khamaru contributed equally to this work.

Fixed points of this type lie at the core of many mathematical areas, including differential and integral equations [Tes12; Kir11], game theory [Sto89], optimization and variational inequalities [Nes03; RW09], as well as dynamic programming and reinforcement learning [Ber19; Put05]. In these settings, the contraction property not only plays an instrumental role in existence and uniqueness proofs, but also leads to efficient methods for computing fixed points. Our focus will be on the extension of such methods to problems in which the operator $\mathbf{h}$ can be observed only via a stochastic oracle that, when given a query point θ, returns a noisy version of the operator evaluation $\mathbf{h}(\theta)$. Such random observation models necessitate the use of stochastic approximation schemes. A fundamental question associated with such schemes is their *statistical complexity*: how many noisy operator evaluations are required to estimate the fixed point θ^* to a pre-specified accuracy? In this paper, we undertake a fine-grained yet relatively general analysis of this question. Notably, our analysis captures the way in which statistical complexity depends on the geometry of the Banach space, as well as the structure of the fixed point θ^* itself.

An important sub-class of Banach spaces are Hilbert spaces, with the Euclidean case ($\mathbb{V} = \mathbb{R}^d$ with the usual inner product) being one special example. The behavior of stochastic approximation for many Hilbert spaces is relatively well understood. In this case, the space $\mathbb{V}$ is endowed with an inner product $\langle \cdot, \cdot \rangle_{\mathbb{V}}$ that induces the norm $\|x\| = \sqrt{\langle x, x \rangle_{\mathbb{V}}}$. For example, for the Euclidean space $(\mathbb{R}^d, \| \cdot \|_2)$, if we set $\mathbf{h}(x) := x - \beta^{-1} \nabla f(x)$ for a β-smooth and strongly convex function f, then solving the fixed-point equation (1) is equivalent to minimizing the function f. A rich theory has been developed around this stochastic optimization problem [BCN18; Nem+09], giving rise to the concepts of averaging [PJ92; Rup88], acceleration [GL12; GL13], and variance reduction [JZ13; NST21; Li+20], along with associated characterizations of optimality [MB11; DR16; Mou+20].

In contrast, relatively less is known in the general setting of Banach spaces. One of the simplest examples is $\mathbb{R}^d$ equipped with a non-Euclidean norm, such as the ℓ_∞-norm. To be clear, non-Euclidean set-ups of this type have been studied in the literature on stochastic optimization and stochastic variational inequalities, with the method of mirror descent being a representative example [NY83; JNT11; KLL20]. Our study, however, deviates from this line of research. The difference stems from the formulation of the problem itself: the operator $\mathbf{h}$ in equation (1) is a mapping from $\mathbb{V}$ to itself, whereas the operators studied in variational inequalities map a Banach space to its dual. This difference leads to a different path of analysis, as taken here.

At least initially, it might seem that non-Euclidean geometry should pose little difficulty for stochastic approximation: all norms are equivalent in the finite-dimensional case, and as is known from standard theory (e.g., [Bor09]), asymptotic convergence depends ultimately on the limiting ODE defined by the scheme. From a non-asymptotic point of view, however, the picture becomes more nuanced: a natural desideratum is that the bounds depend on the *geometric complexity* of $\mathbb{V}$, as opposed to its (possibly much larger) ambient dimension. The difference between the two can be significant. As one concrete example, when solving fixed-point equations that arise in tabular Markov decision processes, the ambient dimension is the size of state-action space, whereas one can obtain ℓ_∞-norm bounds that have only logarithmic dependency on the dimension (see, e.g., [Wai19b; Kha+20]). Our first goal, therefore, is to develop a unified and geometry-aware theory for a certain class of stochastic approximation procedures in Banach spaces.

Our second goal is to establish bounds that are instance-dependent, and so move us beyond a classical worst-case analysis. Any method for stochastic approximation corresponds to a particular type of recursive statistical estimator, so that the the classical statistical theory of

2

local asymptotic minimax can be brought into play [Häj72; Vaa00]. This theory provides a framework for deriving lower bounds on the error of any estimator that depend explicitly on (a local neighborhood of) the instance under consideration. As for the form that such bounds should take in our setting, recall that a sum of i.i.d. random variables in Banach spaces is known (under mild regularity conditions) to satisfy a central limit theorem (see Ledoux and Talagrand [LT13], Section 10). These two lines of asymptotic analysis, in conjunction, indicate that the "right" complexity for estimation in a Banach space $\mathbb{V}$ should involve the expected norm of a Gaussian random element with covariance structure specified by the noise in the stochastic oracle. Given this fundamental limit, it is natural to seek an estimator whose *non-asymptotic* risk matches this quantity, with possible higher-order terms which, again, depends only the geometric complexity of the norm $\|\cdot\|$ (and not the ambient dimension).

In order to address these goals, we analyze an extension of the ROOT-SGD algorithm, a stochastic approximation (SA) algorithm introduced in past work involving a subset of the current authors [Li+20]. We adapt the scheme to solve general fixed-point problems and establish instance-dependent non-asymptotic guarantees in general Banach spaces. More specifically:

- We establish sharp non-asymptotic bounds on the *operator defect* $\|\mathbf{h}(\theta_n) - \theta_n\|$ of the iterate θ_n after n rounds. The leading-order term, defined in terms of a Gaussian complexity induced by the noisy evaluations of the operator $\mathbf{h}$, matches the the optimal Gaussian limit. To the best of our knowledge, this is the first non-asymptotic bound for SA procedures with general non-Euclidean norm that depends directly on the geometric complexity of the underlying space.

- Under a local linearization assumption on the operator $\mathbf{h}$, we establish a sharp instance-dependent upper bound on the *estimation error* $\left\|\theta_n - \theta^*\right\|_C$, measured by any semi-norm $\left\|\cdot\right\|_C$ that is dominated by $\|\cdot\|$. The leading-order term of this bound is a Gaussian complexity involving the dual ball of the semi-norm $\left\|\cdot\right\|_C$, and its interaction with locally linear approximations of the operator around θ^*.

- When the operator $\mathbf{h}$ is affine, we establish an improved result that matches the leading-order term in the nonlinear case, and with an even lower sample complexity. We also generalize this result to settings in which $\mathbf{h}$ itself is not necessarily contractive, but its m-step composition is contractive.

- Finally, we illustrate some specific consequences of our theory for different examples, including stochastic shortest path problems, Markov games, and average-reward policy evaluation.

1.1 Related work

In this section, we survey existing literature on stochastic approximation and its variance-reduced analogues.

Stochastic approximation and asymptotic guarantees: The study of stochastic approximation methods dates back to the seminal work of Robbins and Monro [RM51], as well as Kiefer and Wolfowitz [KW52], who established asymptotic convergence for various classes of one-dimensional problems. Subsequent work by Ljung [Lju77a; Lju77b] and Kushner and Clark [KC78] provided general criteria for convergence to a stable limit, in particular by using an ordinary differential equation (ODE) to track the trajectory of SA procedures. The ODE method has been substantially refined in a long line of subsequent work [Ben96; KY03; Bor09; BMP12]. In addition to pointwise convergence, there is a rich body of work characterizing

the asymptotic distribution of SA trajectories [Kha66; Kus84; KS84]. We refer the reader to the monographs [KY03; Bor09; BMP12] for more background and details on these results.

The idea of improving SA schemes by averaging the iterates was proposed in independent work by Polyak and Juditsky [Pol90; PJ92] as well as Ruppert [Rup88]. Averaging the iterates allows for the use of more aggressive stepsize choices, and Gaussian limiting behavior is achieved over a broad range. The form of this limiting distribution is known to optimal in the sense of local asymptotic minimax [Häj72; Vaa00; DR16]. The idea of iterate averaging underlies many important aspects of large-scale statistical learning, leading to improved algorithms in different settings [BM13; DR16; Tri+18] and laying the foundations of online statistical inference [Che+20a]. The ROOT-SGD algorithm [Li+20] that inspired our approach is motivated by the averaging scheme, but combines variance reduction with averaging of the gradient sequence (as opposed to the sequence of iterates).

Non-asymptotic guarantees for stochastic approximation: Recent years have witnessed significant interest in obtaining non-asymptotic guarantees of the standard SA scheme (see equation (3) in the sequel). For instance, Qu and Wierman [QW20] directly analyzed the iterates of SA algorithms in the asynchronous setting, whereas Chen et al. [Che+20b] derived non-asymptotic bounds on stochastic approximation methods using Lyapunov functions. Using the generalized Moreau envelope, they constructed a smooth Lyapunov function, and show that the iterates of a standard SA scheme have a negative drift with respect to this Lyapunov function. Such Lyapunov techniques have been used to derive non-asymptotic guarantees for SA schemes in variety of settings (e.g., [Che+21a; Che+19; Che+21b; ZZM21]). Wainwright [Wai19b] proved non-asymptotic guarantees for stochastic approximation algorithms under a cone-contractive assumption. For general contractive fixed-point problems in Banach spaces, Gupta et al., [GJG18] developed general criteria for the asymptotic convergence of mini-batch fixed-point iterations; and recently, Borkar [Bor21] established non-asymptotic concentration inequalities for the iterates, albeit with potentially dimension-dependent pre-factors. It should be noted that the standard SA scheme (3), while guaranteed to converge to the fixed point, may do so at a sub-optimal rate when measured in a minimax sense; for example, the papers [Wai19b; Li+21] demonstrate the non-optimality of this approach for the Q-learning problem in reinforcement learning.

Non-asymptotic guarantees that are instance-dependent—meaning that they go beyond worst-case and are adaptive to the difficulty—have been established for several stochastic approximation procedures. For stochastic gradient (SG) methods in the Euclidean setting, such bounds have been established for Polyak-Ruppert-averaged SG [MB11; GP17] and variance-reduced SG algorithms [Fro+15; Li+20], with the sample complexity and high-order terms being improved over time. For reinforcement learning problems, such type of guarantees have been established in the $\| \cdot \|_\infty$ norm for temporal difference methods [Kha+20] and Q-learning [Kha+21] under a generative model, as well as Markovian trajectories [Mou+21; LLP21] under the ℓ_2-norm. In the context of stochastic optimization, the paper [Li+20] provides fine-grained bound for ROOT-SGD with a *unity* pre-factor on the leading-order instance-dependent term. The bounds in our paper, on the other hand, involve constants that need not be optimal in this sense. It is an interesting future direction of research to establish similar non-asymptotic bounds for ROOT-SA with the sharp unity pre-factor.

Variance-reduced stochastic approximation algorithms: In order to obtain optimal SA procedures, different forms of variance reduction have been analyzed. The idea of variance

reduction in stochastic approximation is classical; in the specific context of stochastic gradient methods, the papers [JZ13; DBLJ14; SLRB17] proposed versions of variance reduction that accelerate convergence by careful averaging and re-centering of the gradient sequence. In this special case of stochastic optimization, the fixed-point operator $\mathbf{h}$ is obtained from the gradient update operator (cf. the discussion in Section 1); under suitable convexity and smoothness conditions, it is contractive under the ℓ_2-norm. In more recent work, several fully online schemes for variance-reduced stochastic optimization have been developed and analyzed, including SARAH [Ngu+17; NST21], STORM [CO19] and ROOT-SGD [Li+20]. The ROOT-SGD scheme uses recursive $1/t$-averaging of gradients, and has been been shown to be optimal for various convex problems in both asymptotic and non-asymptotic settings; see the paper [Li+20] and references therein for more details.

In the context of reinforcement learning (RL) problems, the operator $\mathbf{h}$ often corresponds to some type of Bellman operator [Ber12b; Ber19], known to be contractive under the ℓ_∞-norm. Unfortunately, the key techniques used to design optimal methods for RL differ considerably from those used in the stochastic optimization literature. Concretely, in order to obtain optimal RL algorithms, it is often necessary to exploit monotonicity properties of the Bellman operator, combined with variance reduction schemes [Sid+18; Wai19c; Kha+20; Kha+21]. Consequently, the literature is currently lacking a more unified perspective on how to obtain optimal SA schemes in a general setting. The main contribution of our paper is to fill this gap by proposing and analyzing a single variance-reduced stochastic approximation algorithm for finding the fixed point of *any contractive* operator. In this way, our analysis does not depend on the exact form of the contraction norm $\|\cdot\|$.

Notation: We use $\mathbb{V}^*$ to denote the dual space of the Banach space $\mathbb{V}$, i.e., the space of all bounded linear functionals on $\mathbb{V}$. We define the dual norm $\|y\|_* := \sup_{x \in \mathbb{V} \setminus \{0\}} \langle x, y \rangle / \|x\|$. We define the unit norm ball $\mathbb{B} := \{x \in \mathbb{V}, \|x\| \leq 1\}$ in $\mathbb{V}$, as well the dual norm unit ball $\mathbb{B}^* := \{y \in \mathbb{V}^* \mid \|y\|_* \leq 1\}$.

Given a bounded linear operator $A : \mathbb{V} \to \mathbb{V}$, the adjoint operator $A^* : \mathbb{V}^* \to \mathbb{V}^*$ is characterized by the property

$$\langle Ax, y \rangle = \langle x, A^*y \rangle \qquad \text{for all } x \in \mathbb{V} \text{ and } y \in \mathbb{V}^*.$$

The operator norm of a bounded linear operator A on $\mathbb{V}$ is given by $\|A\|_{\mathbb{V}} := \sup_{x \in \mathbb{V} \setminus \{0\}} \frac{\|Ax\|}{\|x\|}$. Similarly, we can define the operator norm $\|\cdot\|_{\mathbb{V}^*}$ of a bounded linear operator mapping from $\mathbb{V}^*$ to itself. For any bounded linear operator A that maps $\mathbb{V}$ to itself, we have the equivalence $\|A^*\|_{\mathbb{V}^*} = \|A\|_{\mathbb{V}}$.

2 Problem set-up and the ROOT-SA Algorithm

In this section, we begin with a precise description of the class of problems that we study, along with the assumptions imposed. We then describe the ROOT-SA algorithm analyzed in this paper.

2.1 Problem formulation

Consider a separable Banach space $(\mathbb{V}, \|\cdot\|)$, and an operator $\mathbf{h}$ mapping from $\mathbb{V}$ to itself. Assuming sufficient regularity to guarantee the existence and uniqueness of the fixed-point

θ^* of the operator $\mathbf{h}$, we study stochastic approximation procedures for estimating the fixed point, i.e., for approximately solving the equation $\mathbf{h}(\theta) = \theta$.

In many practical applications, we may not have access to the operator $\mathbf{h}$ itself; instead, at each time t, we have access to a stochastic oracle $\mathbf{H}_t$ that, when queried at some $\theta \in \mathbb{V}$, returns a noisy version $\mathbf{H}_t(\theta)$ of the operator evaluation $\mathbf{h}(\theta)$. We impose the following conditions on the stochastic operators $\{\mathbf{H}_t\}_{t \geq 1}$ and the population operator $\mathbf{h}$:

Assumptions

(A1) There is a scalar $\gamma \in [0, 1)$ such that the operator $\mathbf{h} : \mathbb{V} \to \mathbb{V}$ is γ-contractive—viz.

$$\|\mathbf{h}(\theta_1) - \mathbf{h}(\theta_2)\| \leq \gamma \|\theta_1 - \theta_2\| \quad \text{for all } \theta_1, \theta_2 \in \mathbb{V}.$$

(A2) For each $t = 1, 2, \ldots$, the stochastic operator $\mathbf{H}_t : \mathbb{V} \mapsto \mathbb{V}$ is almost surely (a.s.) L-Lipschitz:

$$\|\mathbf{H}_t(\theta_1) - \mathbf{H}_t(\theta_2)\| \leq L \|\theta_1 - \theta_2\| \qquad \text{a.s. for all } \theta_1, \theta_2 \in \mathbb{V}.$$

(A3) For any fixed $\theta \in \mathbb{V}$, the noise variables $\{\varepsilon_t(\theta) := \mathbf{H}_t(\theta) - \mathbf{h}(\theta)\}_{t \geq 1}$ are zero-mean and i.i.d., and $\|\varepsilon_t(\theta^*)\| \leq b_*$ almost surely for all $t = 1, 2, \ldots$.

A few remarks are in order. By the Banach fixed point theorem (e.g., [DG03]), the contractivity condition in Assumption (A1) ensures that $\mathbf{h}$ has a unique fixed point θ^*. The bulk of our analysis imposes Assumption (A1), with the exception of Section 3.3, where it is relaxed to a multi-stage contraction assumption in the special case of linear operators. Throughout this paper, we assume that $\gamma \geq \frac{3}{4}$ for the ease of presentation. Note that this assumption can be made without loss of generality, since an operator that is γ-contractive for some $\gamma \in [0, 3/4)$ is also 3/4-contractive.

Assumption (A2) requires the stochastic operator $\mathbf{H}_t$ to be Lipschitz, with the associated constant L allowed to be much larger than one—that is, there is no requirement that $\mathbf{H}_t$ be contractive or non-expansive. This setup should be contrasted with past work on cone-contractive operators [Wai19b; Wai19c] or ℓ_∞-norm contractions [Kha+20; Kha+21], in which the stochastic operator $\mathbf{H}_t$ itself is required to be contractive. In the special case of stochastic optimization in $\mathbb{R}^d$, this type of sample-level Lipschitz condition is widely used, especially for variance-reduced procedures (cf. [JZ13; NST21; Li+20]).

As for Assumption (A3), it imposes bounds only on the noise function when evaluated at the fixed point θ^* of the operator $\mathbf{h}$. In conjunction with Assumption (A2), this bound implies that $\|\varepsilon_t(\theta)\| \leq b_* + (L + \gamma)\|\theta - \theta^*\|$, allowing the norm of the noise $\varepsilon_t(\theta)$ to grow linearly with $\|\theta - \theta^*\|$. It is worth remarking that by using slightly more involved concentration arguments, it is possible to relax the almost sure bounds in Assumptions (A2) and (A3). More precisely, it suffices to impose a p^{th}-moment condition on all projections:

$$\sup_{u \in \Gamma} \mathbb{E}\left[\langle u, \mathbf{H}_1(\theta_1) - \mathbf{H}_1(\theta_2)\rangle^p\right] \leq p! \cdot L^p \|\theta_1 - \theta_2\|^p \qquad \text{for all } \theta_1, \theta_2 \in \mathbb{V}, \text{ and} \tag{2a}$$

$$\sup_{u \in \Gamma} \mathbb{E}\left[\langle u, \varepsilon_1(\theta^*)\rangle^p\right] \leq p! \cdot b_*^p, \tag{2b}$$

for all $p \geq 2$. Here Γ is a skeleton set whose convex hull generates the dual norm ball.

2.2 The ROOT-SA algorithm

Stochastic approximation algorithms are methods for solving fixed-point equations based on noisy observations. In the simplest of such schemes, one starts with initial point θ_0, and then performs the recursive update

$$\theta_{t+1} = \theta_t + \alpha_t \left\{ \mathbf{H}_t(\theta_t) - \theta_t \right\}, \tag{3}$$

where $\{\alpha_t\}_{t \geq 0}$ is a sequence of positive stepsizes, typically in the interval $(0, 1)$. At any given step t, conditioned on θ_t, the quantity $\mathbf{H}_t(\theta_t)$ is an unbiased estimate of $\mathbf{h}(\theta_t)$, and the noise in the observation model is given by $\mathbf{H}_t(\theta_t) - \mathbf{h}(\theta_t)$. Under the contractivity assumptions (A1) on the operator $\mathbf{h}$ and moment bounds on the observation noise $\{\mathbf{H}_t(\theta_t) - \mathbf{h}(\theta_t)\}_{t \geq 1}$, the sequence $\{\theta_t\}$ converges almost surely to the unique fixed point θ^*; moreover, the rate of convergence of θ_t to θ^* is governed by the conditional variance of $\mathbf{H}_t(\theta_t)$ around its conditional mean $\mathbf{h}(\theta_t)$. See the standard texts [KY03; Bor09; BMP12] for results of this type.

Algorithm 1 ROOT-SA : A recursive SA algorithm

 1: Given (a) Initialization $\theta_0 \in \mathbb{V}$, (b) Burn-in $B_0 \geq 2$, and (c) stepsize $\alpha > 0$
 2: **for** $t = 1, \dots, T$ **do**
 3:     **if** $t \leq B_0$ **then**
 4:         $v_t = \frac{1}{B_0} \sum\limits_{t=1}^{B_0} \left\{ \mathbf{H}_t(\theta_0) - \theta_0 \right\}, \quad$ and $\quad \theta_t = \theta_0.$
 5:     **else**
 6:         $v_t = \left(\mathbf{H}_t(\theta_{-1}) - \theta_{-1} \right) + \frac{t-1}{t} \left\{ v_{t-1} - \left(\mathbf{H}_t(\theta_{-2}) - \theta_{-2} \right) \right\},$
 7:         $\theta_t = \theta_{t-1} + \alpha v_t.$
 8:     **end if**
 9: **end for**
10: **return** θ_T

The goal of variance reduction is to improve the basic stochastic approximation scheme (3) by replacing $\mathbf{H}_t(\theta_t) - \theta_t$ with an alternative quantity v_t that has lower variance. In this paper, we study a simple version of such a variance-reduction scheme, as described in Algorithm 1. Our algorithm is inspired by Recursive-one-over-t SGD (ROOT-SGD) algorithm proposed and analyzed in the past work [Li+20] involving a subset of the current authors. The ROOT-SGD algorithm was developed for stochastic optimization; it exploits a two-time scale framework that averages the gradient while performing variance reduction. Our ROOT-SA algorithm extends this same idea to the more general setting of stochastic approximation for fixed-point finding in Banach spaces. While the algorithms are similar in spirit, the analysis in this paper uses completely different techniques, since it applies to the Banach-space setting for general operators, as opposed to the Euclidean setting and gradient operators of convex functions. The key technical difficulties lie in the absence of inner product structure.

3 Main results

In this section, we state our main results and discuss some of their consequences. At a high level, our main results consist of various non-asymptotic bounds on the behavior of the ROOT-SA algorithm in a number of different (semi)-norms. In Section 3.1, we derive bounds on the *operator defect* $\|\mathbf{h}(\theta_t) - \theta_t\|$, which measures how far the t^{th}-iterate θ_t of Algorithm 1 is

from being a fixed point of the population operator $\mathbf{h}$. In other settings, we are interested in bounds on the *estimation error* $\|\theta_t - \theta^*\|$; accordingly, Section 3.2 is devoted to such results, along with bounds on various kinds of semi-norms. Finally, in Section 3.3, we discuss how to obtain refined results in the special case of linear operators, for which the contractivity assumption (A1) can be relaxed.

Central to our bounds are the second-order properties of the i.i.d. noise sequence $\{\varepsilon_t(\theta^*)\}_{t \geq 1}$. In particular, we let $W \in \mathbb{V}$ be a zero-mean Gaussian random element with covariance structure

$$\mathbb{E}\left[\langle W, y \rangle \cdot \langle W, z \rangle\right] = \mathbb{E}\left[\langle \varepsilon_1(\theta^*), y \rangle \cdot \langle \varepsilon_1(\theta^*), z \rangle\right] \quad \text{for all } y, z \in \mathbb{V}^*. \tag{4}$$

Various statistics of this Banach-space-valued random variable, including its mean $\mathbb{E}[\|W\|]$ and variance in certain directions, specify the leading instance-dependent terms of our results.

3.1 Upper bounds on operator defect

We begin by stating some non-asymptotic upper bounds on the so-called operator defect $\|\mathbf{h}(\theta_t) - \theta_t\|$, which measures the error in the iterate θ_t as a fixed point. As noted above, the Gaussian element W with covariance structure (4) plays a central role. In addition to the expected norm $\mathbb{E}[\|W\|]$, our result involves a certain type of maximal variance over a skeleton set—namely, a subset Γ of the dual ball $\mathbb{V}^\star$ such that $\mathbb{B}^* = \text{conv}(\Gamma)$, so that the norm $\| \cdot \|$ has the variational representation $\|x\| = \sup_{y \in \Gamma} \langle x, y \rangle$. Given a set of this type, we define the Γ-maximal variance

$$\sigma_\Gamma^2(W) := \sup_{u \in \Gamma} \mathbb{E}\left[\langle u, W \rangle^2\right]. \tag{5}$$

With this definition, we have the following:

Theorem 1. *Under Assumptions (A1)—(A3) and a given failure probability $\delta \in (0, 1)$, there is a range of stepsizes $\alpha > 0$ such that with burn-in period $B_0 := \frac{c}{(1-\gamma)^2 \alpha} \log\left(\frac{n}{\delta}\right)$, sample size $n \geq 2B_0$, the last iterate θ_n of Algorithm 1 satisfies*

$$\|\mathbf{h}(\theta_n) - \theta_n\| \leq \underbrace{\frac{c}{\sqrt{n}}\left\{\mathbb{E}[\|W\|] + \sqrt{\sigma_\Gamma^2(W) \log(\tfrac{1}{\delta})}\right\}}_{\text{Instance-dependent}} + \underbrace{\frac{cB_0}{n} \cdot \|\theta_0 - \mathbf{h}(\theta_0)\|}_{\text{Initial error}} + \underbrace{\mathcal{H}_n(\delta, \alpha)}_{\text{Higher-order}} \tag{6}$$

with probability at least $1 - \delta$. Here $\mathcal{H}_n(\delta, \alpha)$ is a higher-order term defined below (cf. equation (8b)).

See Section 5.1 for a proof of this theorem.

Instance-dependent term: As we discuss in the sequel, with a stepsize $\alpha \asymp 1/\sqrt{n}$, the dominant quantity in this upper bound is the instance-dependent term defined by the Gaussian process W. So as to appreciate its significance, we note that for any $\delta \in (0, 1)$, a near-optimal tail bound for the Gaussian process W is given by

$$\mathbb{P}\left[\|W\| \geq \mathbb{E}[\|W\|] + c\sqrt{\sigma_\Gamma^2(W) \log(\tfrac{1}{\delta})}\right] \leq \delta.$$

8

(For instance, see Section 3.1 in Ledoux and Talagrand [LT13]). Thus, the instance-dependent term in Theorem 1 matches the behavior of the limiting Gaussian random variable W, up to constant factors and high-order terms.

Of course, it is natural to wonder whether this instance-dependent term is actually optimal for stochastic approximation, or more generally, for any procedure used to estimate the fixed point based on stochastic observations. As we discuss in Section 3.2, for a finite-dimensional space $\mathbb{V}$, this term matches the fundamental lower bound provided by local asymptotic minimax theory, so that—at least in general—it cannot be improved.

3.1.1 Stepsizes and higher-order term

The range of permissible stepsizes and the higher-order term involve certain Dudley entropy integrals, which we now define. Let $\mathbb{B}$ denote the unit ball in the space $(\mathbb{V}, \|\cdot\|)$, and let $\mathbb{B}^*$ denote the dual norm ball. Recall that Γ is a skeleton set such that $\mathbb{B}^* = \text{conv}(\Gamma)$.

Given a metric ρ on the dual space, we define (for any $q \geq 1$) the Dudley entropy integral

$$\mathcal{J}_q(\Gamma, \rho_n) := \int_0^\infty \left[\log N(s; \Gamma, \rho_n) \right]^{1/q} ds,$$

where $N(s; \Gamma, \rho)$ denotes the cardinality of a minimal s-covering of the skeleton set $\Gamma \subseteq \mathbb{V}^\star$ under ρ. Of particular interest are the cases $q = 2$ and $q = 1$, which arise in the cases of sub-Gaussian and sub-exponential tails, respectively.

In the simplest case (when $\mathbb{B}^*$ is totally compact, such as the finite-dimensional case), we can let ρ be the dual norm $\|\cdot\|_*$. However, to handle the general infinite-dimensional case and also sharpen our results, we make use of the following pseudo-metric on the skeleton set

$$\rho_n(x, y) := \sup_{e \in n\Omega \cap \mathbb{B}} \langle x - y, e \rangle \quad \text{defined for all pairs} \quad x, y \in \Gamma, \tag{7}$$

where Ω is the range of the operator. Note that the additional restriction $e \in n\Omega$ makes ρ_n a weaker pseudo-metric than the dual norm, and in particular, we have $\rho_n(x, y) \leq \|x - y\|_*$ for any $x, y \in \mathbb{V}^*$. This weakening is especially important in the infinite-dimensional case, where the skeleton set Γ is not compact under the original norm $\|\cdot\|_*$.

With this notation, for a given tolerance probability $\delta \in (0, 1)$, the range of permissible stepsizes is given by

$$\alpha \in \left(0, \frac{(1-\gamma)^2}{cL^2 \mathcal{J}_2^2(\Gamma, \rho_n) \log\left(\frac{n}{\delta}\right)}\right], \tag{8a}$$

and the higher-order term is given by

$$\mathcal{H}_n(\delta, \alpha) := c\frac{b_*}{(1-\gamma)} \left[\frac{1}{n} + \frac{\alpha L}{\sqrt{n}} \mathcal{J}_2(\Gamma, \rho_n) \log(\frac{n}{\delta})\right] \cdot \left\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\frac{1}{\delta})\right\}. \tag{8b}$$

3.1.2 Stepsize choice and restarting

Note that Theorem 1 holds for a range of stepsizes, and the stepsize plays a role in the burn-in length $B_0 = \frac{c}{(1-\gamma)^2\alpha} \log\left(\frac{n}{\delta}\right)$. In conjunction, these requirements induce the following lower bound on the sample size

$$n \geq \frac{c}{(1-\gamma)^4} L^2 \mathcal{J}_2^2(\Gamma, \rho_n) \log^2(\frac{n}{\delta}). \tag{9a}$$

9

Given a sample size n satisfying this requirement, suppose that we choose the stepsize

$$\alpha = \left\{ L\mathcal{J}_2(\Gamma, \rho_n) \log\left(\tfrac{n}{\delta}\right) \sqrt{n} \right\}^{-1}. \tag{9b}$$

With this choice, when evaluated at iteration $t = n$, it can be shown that the bound (6) simplifies to

$$\|\mathbf{h}(\theta_n) - \theta_n\| \leq \frac{c}{\sqrt{n}} \left\{ \mathbb{E}[\|W\|] + \sqrt{\sigma_\Gamma^2(W) \log(\tfrac{1}{\delta})} \right\} + \frac{cB_0}{n} \cdot \|\theta_0 - \mathbf{h}(\theta_0)\|$$

$$+ c\frac{b_*}{(1-\gamma)n} \left\{ \mathcal{J}_1(\Gamma, \rho_n) + \log\left(\tfrac{1}{\delta}\right) \right\}. \tag{9c}$$

Thus, we see that the instance-dependent term (with its $1/\sqrt{n}$ decay) dominates the other two terms, which decay at the faster $1/n$ rate.

The final aspect that can be refined is the dependence of the bound (9c) on the initial error. As stated, this dependence is sub-optimal, but can be refined via a simple restarting procedure, leading to an improved bound

$$\|\mathbf{h}(\theta_n) - \theta_n\| \leq \frac{c}{\sqrt{n}} \left\{ \mathbb{E}[\|W\|] + \sqrt{\sigma_\Gamma^2(W) \log(\tfrac{1}{\delta})} \right\} + c\frac{b_*}{(1-\gamma)n} \left\{ \mathcal{J}_1(\Gamma, \rho_n) + \log\left(\tfrac{1}{\delta}\right) \right\}, \tag{10}$$

as long as the initial operator defect $\|\theta_0 - \mathbf{h}(\theta_0)\|$ is controlled by a finite-degree polynomial of the sample size n. See Appendix A for the details of this procedure. In Corollary 1 and Theorem 2 to follow, we assume that such re-starting scheme has been applied, so that the contribution from initial gap $\|\theta_0 - \mathbf{h}(\theta_0)\|$ is negligible.

3.1.3 Semi-norm bounds on the operator defect

There are various practical settings in which it is of interest to obtain a bound in some semi-norm $\|\cdot\|_C$, as opposed to the original Banach space norm $\|\cdot\|$. As a simple example, in the Euclidean setting, i.e. $\mathbb{V} = \mathbb{R}^d$, one might have an operator that is contractive in the ℓ_2-norm, but be interested in deriving bounds in the ℓ_∞-norm. As a second example, in various applications, one only cares about the error in some fixed direction $v \in \mathbb{R}^d$, so that the semi-norm $\|\theta\|_C := |v^\top \theta|$ is the relevant quantity.

In this section, we state a family of bounds applicable to any semi-norm $\|\cdot\|_C$ of the form

$$\|\theta\|_C := \sup_{v \in C} \langle v, \theta \rangle \qquad \text{where } C \subset \mathbb{V}^\star \text{ is a symmetric and convex subset.} \tag{11}$$

Note that a wide class of interesting semi-norms can be generated in this way.

A crude bound can be obtained by relating the semi-norm to the Banach space norm. In particular, when the *norm domination factor* $\mathcal{D} := \sup_{v \in C} \|v\|_*$ is finite, then any $\theta \in \mathbb{V}$ satisfies the upper bound

$$\|\theta\|_C \leq \left(\sup_{v \in C} \|v\|_*\right) \cdot \|\theta\| := \mathcal{D} \cdot \|\theta\|, \tag{12}$$

and a direct application of Theorem 1 yields

$$\|\mathbf{h}(\theta_n) - \theta_n\|_C \leq \mathcal{D} \cdot \|\mathbf{h}(\theta_n) - \theta_n\| \lesssim \frac{\mathcal{D}}{\sqrt{n}} \cdot \left\{ \mathbb{E}[\|W\|] + \sqrt{\sigma_\Gamma^2(W) \log(\tfrac{1}{\delta})} \right\}.$$

This bound is potentially weak for two reasons: (a) the leading term depends directly on $\mathcal{D}$, which can be large and possibly dependent on the ambient dimension of the problem; and (b) it depends in a global way on the Gaussian random element W, via the skeleton set Γ as opposed to C, which can be much smaller.

It is natural to expect that one could prove bounds with a leading term specified in terms of $\mathbb{E}\|W\|_C$ along with the refined variance functional $\sigma_C^2 := \sup_{u \in C} \mathbb{E}[\langle u, W \rangle^2]$. This refinement is the content of the following:

Corollary 1. *Under the conditions of Theorem 1, the iterate θ_n satisfies the bound*

$$\left\|\mathbf{h}(\theta_n) - \theta_n\right\|_C \leq \frac{c}{\sqrt{n}} \left\{ \mathbb{E}[\|W\|_C] + \sqrt{\sigma_C^2(W) \log(\tfrac{1}{\delta})} \right\} + \mathcal{H}_n^\dagger(\delta, \alpha) \tag{13}$$

with probability at least $1 - \delta$, where

$$\mathcal{H}_n^\dagger(\delta, \alpha) := \frac{\mathcal{D}}{(1-\gamma)} \left\{ L\mathcal{J}_2(\mathbb{B}^*, \rho_n) \log(\tfrac{n}{\delta})\sqrt{\tfrac{\alpha}{n}} + \tfrac{1}{n\sqrt{\alpha}} \right\} \left\{ \mathbb{E}[\|W\|] + \sqrt{\sigma_\Gamma^2(W) \log(\tfrac{n}{\delta})} \right\}$$
$$+ \frac{c\mathcal{D}Lb_*}{1-\gamma} \left\{ \tfrac{\sqrt{\alpha}}{n} + \tfrac{\alpha}{\sqrt{n}} \right\} \mathcal{J}_2(\Gamma, \rho_n)\mathcal{J}_1(\Gamma, \rho_n) \log^2(\tfrac{n}{\delta}).$$

See Section 5.2 for the proof.

With the stepsize choice (9b), the higher-order term scales as

$$\mathcal{H}_n^\dagger(\delta, \alpha) = \left\{ \mathbb{E}[\|W\|] + \sqrt{\sigma_\Gamma^2 \log(\tfrac{1}{\delta})} \right\} \cdot \widetilde{\mathcal{O}}\left(\frac{\mathcal{D}}{(1-\gamma)n^{3/4}} \right) + \widetilde{\mathcal{O}}\left(\frac{\mathcal{D}}{(1-\gamma)n} \right),$$

where $\widetilde{\mathcal{O}}$ subsumes various constants and logarithmic factors. Any dependence on global features of the Banach space appears only in this higher-order term, which goes to zero at a rate faster than $1/\sqrt{n}$.

3.2 Upper bounds on the estimation error

Thus far, our analysis has focused on bounding the operator defect $\theta_n - \mathbf{h}(\theta_n)$ in various (semi)-norms. In this section, we turn to problem of deriving upper bounds on the estimation error $\|\theta_n - \theta^*\|$, which is the primary goal in various applications of the SA methodology. Bounds on the operator defect imply bounds on this quantity: indeed, some simple calculation[1] yields the bound

$$\|\theta_n - \theta^*\| \leq \tfrac{1}{1-\gamma} \cdot \|\theta_n - \mathbf{h}(\theta_n)\|. \tag{14}$$

Although this bound is useful—and sharp in a worst-case sense— it can certainly be improved in general.

In this section, we develop a result (to be stated as Theorem 2) that gives a sharper bound on the estimation error $\|\theta_n - \theta^*\|$ when it is possible to construct linear approximations of the operator $\mathbf{h}$ in a neighborhood of θ^*. More precisely, we impose the following local linearity condition.

---

[1] By the triangle inequality, we have $\|\theta_n - \theta^*\| \leq \|\theta_n - \mathbf{h}(\theta_n)\| + \|\mathbf{h}(\theta_n) - \mathbf{h}(\theta^*)\|$. From the contractivity assumption (A1), we have $\|\mathbf{h}(\theta_n) - \mathbf{h}(\theta^*)\| \leq \gamma\|\theta_n - \theta^*\|$, and rearranging yields the claim.

Assumption: Local linearity

(A4) For any $s > 0$, there exists a set $\mathcal{A}_s$ of bounded linear operators on $\mathbb{V}$ such that

$$\|\theta - \theta^*\| \leq \sup_{A \in \mathcal{A}_s} \|(I - A)^{-1}(\mathbf{h}(\theta) - \theta)\| \qquad \text{for all } \theta \in \mathbb{B}(\theta^*, s). \tag{15}$$

As before, let W be a centered Gaussian random variable in $\mathbb{V}$ with the same covariance structure as $\varepsilon_1(\theta^*) := \mathbf{H}_1(\theta^*) - \mathbf{h}(\theta^*)$—that is

$$\mathbb{E}\left[\langle W, y \rangle \cdot \langle W, z \rangle\right] = \mathbb{E}\left[\langle \varepsilon_1(\theta^*), y \rangle \cdot \langle \varepsilon_1(\theta^*), z \rangle\right] \qquad \text{for all } y, z \in \mathbb{V}^\star.$$

Our bounds in this section are stated in terms of the solution to a fixed-point equation involving functionals of the Gaussian noise W. For any $s > 0$, define

$$\mathcal{G}(s) := \mathbb{E}\left[\sup_{\substack{y \in \Gamma \\ A \in \mathcal{A}_s}} \langle W, (I - A)^{-1}y \rangle\right], \quad \text{and} \quad \nu^2(s) := \sup_{\substack{y \in \Gamma \\ A \in \mathcal{A}_s}} \mathbb{E}\left[\langle y, (I - A)^{-1}W \rangle^2\right]. \tag{16}$$

Given a stepsize α satisfying (8a) and a tolerance probability $\delta \in (0, \frac{1}{1+\log(1/(1-\gamma))})$, we define the function

$$\mathcal{H}_n(\alpha, \delta) := \frac{\log(\frac{n}{\delta})}{(1-\gamma)^2}\left\{\left[\mathcal{J}_2(\mathbb{B}^*, \rho_n)L\sqrt{\frac{\alpha}{n}} + \frac{1}{n\sqrt{\alpha}}\right] \cdot \mathbb{E}[\|W\|] + \left[\frac{\mathcal{J}_2(\mathbb{B}^*, \rho_n)L\alpha}{\sqrt{n}} + \frac{1}{n}\right] \cdot b_*\left[\mathcal{J}_1(\Gamma, \rho_n) + \log(\frac{n}{\delta})\right]\right\}. \tag{17a}$$

This quantity serves as a higher-order term in our analysis. We consider the following fixed-point equation in the variable s:

$$s = \frac{\mathcal{G}(2s)}{\sqrt{n}} + \nu(2s)\sqrt{\frac{\log(1/\delta)}{n}} + \mathcal{H}_n(\alpha, \delta). \tag{17b}$$

As discussed below equation (19) to follow, equation (17b) has a non-empty and bounded set of non-negative solutions; let s_n^* be the largest such solution.

Theorem 2. *Suppose that Assumptions (A1)– (A4) are in force, and that for some $\delta \in (0, 1)$, we run Algorithm 1 using a stepsize α in the interval (8a) and burn-in period $B_0 = \frac{c}{(1-\gamma)^2\alpha}\log\left(\frac{n}{\delta}\right)$. Then the final iterate θ_n satisfies the bound*

$$\|\theta_n - \theta^*\| \leq c \cdot s_n^* \quad \text{with probability at least } 1 - \delta. \tag{18}$$

See Appendix B for the proof of this theorem.

Note that our contractivity assumption implies that functions $\mathcal{G}$ and ν defined in equation (16) are uniformly bounded—viz.

$$\mathcal{G}(s) = \mathbb{E}\left[\sup_{y \in \Gamma, A \in \mathcal{A}_s} \langle W, (I - A)^{-1}y \rangle\right] \leq \frac{\mathbb{E}[\|W\|]}{1 - \gamma}, \quad \text{and}$$

$$\nu^2(s) := \sup_{y \in \Gamma, A \in \mathcal{A}_s} \mathbb{E}\left[\langle y, (I - A)^{-1}W \rangle^2\right] \leq \frac{1}{(1-\gamma)^2}\sup_{y \in \Gamma}\mathbb{E}\left[\langle y, W \rangle^2\right]. \tag{19}$$

These inequalities (19), in conjunction with Theorem 1, guarantee that the fixed-point equation (17b) has a non-empty and bounded set of solutions; consequently, the maximum solution

s_n^* is well-defined. Moreover, this calculation also reveals that the bound from Theorem 2 is always superior to the naive bound (14).

Note that only the high-order term $\mathcal{H}_n(\alpha, \delta)$ depends on the stepsize. By taking the optimal stepsize $\alpha_n = \left\{ L\mathcal{J}_2(\Gamma, \rho_n) \log \left(\frac{n}{\delta} \right) \sqrt{n} \right\}^{-1}$, this term becomes

$$\mathcal{H}_n(\alpha_n, \delta) := \frac{\log(\frac{n}{\delta})}{(1-\gamma)^2} \left\{ \frac{\sqrt{L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}}{n^{3/4}} \cdot \mathbb{E}[\|W\|] + \frac{b_*}{n} \left[\mathcal{J}_1(\Gamma, \rho_n) + \log(\frac{n}{\delta}) \right] \right\}, \qquad (20)$$

which consists of two terms: an $\mathcal{O}(n^{-3/4})$ term depending on the expected norm $\mathbb{E}[\|W\|]$ that captures the second moment of the noise, and an $\mathcal{O}(n^{-1})$ term depending on the worst-case upper bound on the noise, as well as the Dudley integral. Under our stepsize choice, the high-order terms not only decay at a faster rate with sample size n, but also capture the underlying complexity of the norm $\|\cdot\|$, instead of the ambient dimension of the space $\mathbb{V}$.

3.2.1 Asymptotic optimality

Theorem 2 provides a non-asymptotic bound involving the Gaussian process $(I - A)^{-1}W$ for some $A \in \mathcal{A}_s$. It is natural to ask whether or not this bound is improvable. In certain cases it is straightforward to address this question using local asymptotic minimax theory (cf. [LeC53; Häj72; Vaa00]).

Let us suppose that $\mathbb{V}$ is finite-dimensional, and the operator $\mathbf{h}$ differentiable in an open neighborhood of the point θ^*. In this case, we can use known results to state a lower bound involving the random variable $(I - A_0)^{-1}W$, where $A_0 = \nabla \mathbf{h}(\theta^*)$. In order to state this lower bound precisely, we consider problems indexed by distributions $\mathbb{Q}$ in a local neighborhood of the target distribution $\mathbb{P}$. For any $\mathbb{Q}$, our goal is to solve the fixed point equation $\theta = \mathbb{E}_{\mathbf{H} \sim \mathbb{Q}}[\mathbf{H}(\theta)]$ using i.i.d. samples of the random operator $\mathbf{H}$. Under suitable tail assumptions on the distributions $\mathbb{P}$ and $\mathbb{Q}$, for any estimator $\widetilde{\theta}_n$ that maps a sequence of observed operators $\{\mathbf{H}_t\}_{t=1}^n$ to the vector space $\mathbb{V}$, an adaptation of Theorem 1 from the paper [DR16] (with loss function corresponding to the Banach norm) yields the lower bound

$$\liminf_{\Delta \to \infty} \liminf_{n \to \infty} \sup_{\mathbb{Q} \,|\, D_{\mathrm{KL}}(\mathbb{Q} \,\|\, \mathbb{P}) \leq \frac{\Delta}{n}} \mathbb{E}\left[\sqrt{n} \|(\widetilde{\theta}_n - \theta^*(\mathbb{Q}))\| \right] \geq \mathbb{E}\left[\|(I - A_0)^{-1}W\| \right], \qquad (21)$$

Thus, when estimating θ^* in the Banach norm $\|\cdot\|$, the asymptotic lower bound is given by $\mathbb{E}\left[\|(I - A_0)^{-1}W\| \right]$.

Let us compare this fundamental limit to the behavior of the `ROOT-SA` estimator. We take a sequence of stepsizes $\{\alpha_n\}_{n \geq 1}$ such that $\alpha_n \to 0^+$ and $n\alpha_n \to \infty$. With this choice, applying Theorem 2 yields that the `ROOT-SA` estimator θ_n satisfies the bound

$$\limsup_{n \to \infty} \mathbb{P}\left[\|\theta_n - \theta^*\| \geq c \cdot \mathbb{E}\left[\|(I - A_0)^{-1}W\| \right] \right] \leq \tfrac{1}{3}, \qquad (22)$$

for some universal constant $c > 0$, showing its behavior is controlled by the same functional that appears in the LAM lower bound.

3.2.2 Semi-norm bounds on the estimation error

Recall the setup of Section 3.1.3. We now refine these results by providing an upper bound on $\|\theta_n - \theta^*\|_C$, where $\|\cdot\|_C$ is a semi-norm of the form (11), assumed to satisfy the domination condition (12). Moreover, we assume the following modification of the local linearity condition holds.

Assumption: Local linearity in semi-norm

(A4)$'$ For any $s > 0$, there is a set $\mathcal{A}_s$ of bounded linear operators on $\mathbb{V}$ such that

$$\left\|\theta - \theta^*\right\|_C \leq \sup_{A \in \mathcal{A}_s} \left\|(I - A)^{-1}\big(\mathbf{h}(\theta) - \theta\big)\right\| \qquad \text{for all } \theta \in \mathbb{B}(\theta^*, s). \tag{23}$$

As a refinement of the definition (16), we introduce the complexity terms

$$\mathcal{G}_C(s) := \mathbb{E}\Big[\sup_{\substack{y \in C \\ A \in \mathcal{A}_s}} \langle W,\, (I - A)^{-1}y\rangle\Big], \quad \text{and} \quad \nu_C^2(s) := \sup_{\substack{y \in C \\ A \in \mathcal{A}_s}} \mathbb{E}\Big[\langle y,\, (I - A)^{-1}W\rangle^2\Big]. \tag{24}$$

Given a stepsize α satisfying the bound (8a) and a tolerance probability $\delta \in (0, \frac{1}{\log(1/(1-\gamma))})$, we define $s_{C,n}^* > 0$ to be the largest solution to the fixed-point equation

$$s = \frac{\mathcal{G}_C(2s)}{\sqrt{n}} + \nu_C(2s)\sqrt{\frac{\log(1/\delta)}{n}} + \mathcal{D} \cdot \mathcal{H}_n(\alpha, \delta), \tag{25}$$

where the higher-order term $\mathcal{H}_n(\alpha, \delta)$ was previously defined (17a).

Corollary 2. *Under Assumptions (A1)– (A3) and (A4)$'$, the estimate θ_n from Algorithm 1 satisfies*

$$\left\|\theta_n - \theta^*\right\|_C \leq c \cdot s_{C,n}^* \quad \text{with probability at least } 1 - \delta. \tag{26}$$

See Appendix B for the proof of this corollary.

3.3 Linear operators with multi-step contraction

In the special case where $\mathbf{h}$ is a bounded linear operator in $\mathbb{V}$, the contraction assumption (A1) can be significantly weakened. In particular, it suffices to require that a multi-step composition of the operator be contractive.

Assumption: Multi-step contraction

(A1)$'$ For some integer $m \geq 1$, the affine operator $\mathbf{h}(\theta) = A\theta + b$ is m-stage contractive, meaning that

$$\|A\|_{\mathbb{V}} \leq 1 \quad \text{and} \quad \|A^m\|_{\mathbb{V}} \leq \tfrac{1}{2}. \tag{27}$$

Note that assumption (A1)$'$ implies that the linear operator $(I - A)$ is invertible; in particular, we have the operator norm bound

$$\|(I - A)^{-1}\|_{\mathbb{V}} \leq \sum_{k=0}^{\infty} \sup_{v \in \mathbb{B}} \|A^k v\| = \sum_{k=0}^{\infty} \sum_{j=0}^{m-1} \|A^{mk+j}\|_{\mathbb{V}} \leq \sum_{k=0}^{\infty} \sum_{j=0}^{m-1} \|A^m\|_{\mathbb{V}}^k \cdot \|A^j\|_{\mathbb{V}} \leq 2m. \tag{28}$$

As before, let W be a centered Gaussian random variable in $\mathbb{V}$ with the same covariance structure as $\varepsilon(\theta^*) := \mathbf{H}(\theta^*) - \mathbf{h}(\theta^*)$; that is, $\mathbb{E}\left[\langle W, y\rangle \cdot \langle W, z\rangle\right] = \mathbb{E}\left[\langle\varepsilon(\theta^*), y\rangle \cdot \langle\varepsilon(\theta^*), z\rangle\right]$ for all $y, z \in \mathbb{V}^\star$.

Tuning parameters: Given a desired failure probability $\delta \in (0,1)$, and a total sample size n, we run Algorithm 1 with the following choices of parameters:

$$\text{Stepsize choice:} \quad \alpha \leq \frac{c}{mL^2 \mathcal{J}_2(\mathbb{B}^*, \rho_n)^2 \cdot \log^2 \frac{n}{\delta}} \tag{29a}$$

$$\text{Burn-in time:} \quad B_0 = \frac{cm}{\alpha} \log(\frac{n}{\delta}), \tag{29b}$$

where c is an universal constant.

Theorem 3. *Under Assumptions (A1)′, (A2) and (A3), and given a sample size $n \geq 2B_0$, consider Algorithm 1 run using tuning parameters from equation (29a) and (29b). Then for any given $t \in [B_0, n]$, the iterate θ_t satisfies*

$$\|\mathbf{h}(\theta_n) - \theta_n\| \leq \frac{c}{\sqrt{n}} \Big\{ \mathbb{E}[\|W\|] + \sqrt{\sigma_\Gamma^2 \log(\tfrac{1}{\delta})} \Big\} + \frac{cB_0}{n}\|\theta_0 - \mathbf{h}(\theta_0)\| + \mathcal{H}_n^\circ(\delta, \alpha) \tag{30}$$

with probability $1 - \delta$, where

$$\mathcal{H}_n^\circ(\delta, \alpha) := cb_* \Big\{ \tfrac{1}{n} + \tfrac{\alpha L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{n}} \log(\tfrac{n}{\delta}) \Big\} \Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \Big\},$$

See Appendix C.1 for the proof of this theorem.

Observe that for a linear operator $\mathbf{h}(\theta) = A\theta + b$ that satisfies the contractivity condition (cf. Assumption (A1)′), the inverse $(\mathcal{I} - A)^{-1}$ exists, and we have

$$\theta - \theta^* = (\mathcal{I} - A)^{-1}(\mathbf{h}(\theta) - \theta^*).$$

Consequently, given any semi-norm $\|\cdot\|_C$ of the form (11) satisfying condition (12), an argument similar to Corollary 2 yields the following guarantee. In stating it, we assume that the restarting scheme from Appendix A has been applied to remove dependence on the initial condition.

Corollary 3. *Under the conditions of Theorem 3, running the `ROOT-SA` algorithm with the restarting scheme yields an iterate θ_n such that*

$$\big\|\theta_n - \theta^*\big\|_C \leq \frac{c}{\sqrt{n}} \Big\{ \mathbb{E}\big[\big\|(I-A)^{-1}W\big\|_C\big] + \sqrt{\sup_{u \in C} \mathbb{E}\big[\langle u, (I-A)^{-1}W\rangle^2\big] \log(\tfrac{1}{\delta})} \Big\} + \mathcal{H}_n^\diamond(\delta, \alpha) \tag{31}$$

with probability at least $1 - \delta$.

See Appendix C.2 for the proof of this corollary, along with the definition of $\mathcal{H}_n^\diamond(\delta, \alpha)$.

Since the problem itself is linear, the class $\mathcal{A}_s$ of linear operators is singleton, and the estimation error upper bounds can be expressed directly through $\mathbb{E}\big[\big\|(I-A)^{-1}W\big\|_C\big]$, without resorting to fixed-point equations. Compared with the high-order terms defined by equation (17a) in the general case, the high order terms in equation (31) (the second and third line of the equation) save a factor of $\frac{1}{1-\gamma}$ in the contractive case, while generalizing to the multi-step contraction case. Furthermore, similar to the discussion in Section 3.1.2, the step α can be tuned based on the sample size n and knowledge about other problem parameters, so as to minimize the high-order terms $\mathcal{H}_n^\circ(\delta, \alpha)$ and $\mathcal{H}_n^\diamond(\delta, \alpha)$. The resulting error bounds contain high-order terms similar to equations (17a) and (20), the factor $(1-\gamma)^{-2}$ replaced by the integer m.

15

4 Consequences for specific use cases

Thus far, we have stated a number of general results. In this section, we discuss the consequences of these results for three classes of problems that fall within the framework of this paper. In the main text, we discuss in detail the problem of stochastic shortest paths in Section 4.1 and average-reward policy evaluation in Section 4.2. We defer discussion of methods for solving two-player zero-sum Markov games to Appendix D.

4.1 Computing stochastic shortest paths

We begin with the problem of computing stochastic shortest paths [YB13; BT91], or SSPs for short. It provides an illustration of the general theory using a Banach space defined by a certain weighted ℓ_∞-norm. On one hand, SSPs can be formulated in terms of Markov decision process (MDP) with a finite state space $\mathcal{X}$ and action space $\mathcal{U}$. Thus, although they might appear to be a special case of an MDP, in fact, they are sufficiently general to encompass both finite-horizon MDPs as well as discounted MDPs. Thus, the conclusions obtained in this section apply to a fairly broad class of problems.

An MDP is defined by a collection of probability transition kernels $\{\mathbf{P}_u(\cdot \mid x)\}_{(x,u) \in \mathcal{X} \times \mathcal{U}}$, where the transition kernel $\mathbf{P}_u(x' \mid x)$ denotes the probability of transition to the state x' when an action u is taken at the current state x. The MDP is equipped with a cost function $c : \mathcal{X} \times \mathcal{U} \mapsto \mathbb{R}$, and the value $c(x, u)$ corresponds to cost incurred upon performing the action u in state x. To formulate a *stochastic shortest path* (SSP) problem, we assume that state 1 is absorbing and cost-free, meaning that

$$c(x = 1, u) = 0 \quad \text{and} \quad \mathbb{P}_u(x' \mid x = 1) = \mathbf{1}_{x'=1} \quad \text{for all actions } u \in \mathcal{U}. \tag{32}$$

A stationary policy π is a mapping $\mathcal{X} \mapsto \mathcal{U}$ such that $\pi(x) \in \mathcal{U}$ denotes the action to be taken in the state x. We assume that the total infinite-horizon cost incurred by any stationary policy π is finite—viz. $\mathbb{E}_{x_0=x}\Big[\sum_{k=1}^\infty |c(x_k, \pi(x_k))| \Big] < \infty$ for all $x \in \mathcal{X}$. Such stationary policy π is called a *proper policy*, and our goal is to obtain proper policy $\pi^\star$ that minimizes the total cost.

Associated with any proper policy π is its Q-function

$$\theta^\pi(x, u) := \mathbb{E}\Big[\sum_{k=0}^\infty c(x_k, u_k) \mid x_0 = x, u_0 = u \Big], \quad \text{where } u_k = \pi(x_k) \quad \text{for all } k = 1, 2, \ldots.$$

An optimal policy can be obtained from the optimal Q-function, given by $\theta^\star(x, u) := \inf_{\pi \in \Pi} \theta^\pi(x, u)$.

4.1.1 Bellman operator and contractivity

Observe that for any policy π, the cost-free absorbing state property (32) ensures that $\theta^\pi(1, u) = 0$, and as a result $\theta^\star(1, u) = 0$ for all actions $u \in \mathcal{U}$. In terms of the shorthand $\mathcal{X}_{-1} := \mathcal{X} \setminus \{1\}$, classical theory [BT91; YB13] guarantees that the optimal Q-function restricted to the set $\mathcal{X}_{-1} \times \mathcal{U}$ is the unique fixed point of the Bellman operator

$$\mathbf{h}(\theta)(x, u) = c(x, u) + \sum_{x' \in \mathcal{X}_{-1}} \mathbf{P}_u(x' \mid x) \min_{u' \in \mathcal{U}} \theta(x', u') \qquad (x, u) \in \mathcal{X}_{-1} \times \mathcal{U}. \tag{33}$$

For SSP problems with finite state and action spaces, any Q-function can be viewed an element of $\mathbb{R}^{|\mathcal{X}_{-1} \times \mathcal{U}|}$, in which case the Bellman operator $\mathbf{h}$ can be viewed as acting on $\mathbb{R}^D$

where $D := |\mathcal{X}_{-1} \times \mathcal{U}|$. For a vector $\mathbf{w} := \{w_1, \ldots, w_D\}$ of strictly positive weights, we define a weighted ℓ_∞-norm on $\mathbb{R}^D$ via $\|\theta\|_w := \max\limits_{i=1,\ldots,D} \frac{|\theta_i|}{w_i}$. From known results on SSP problems [BT91; Tse90], one can use a hitting time analysis to define a weight vector $\mathbf{w}$ such that, for any $\theta_1, \theta_2 \in \mathbb{R}^D$, we have

$$\|\mathbf{h}(\theta_1) - \mathbf{h}(\theta_2)\|_w \leq \left(1 - \frac{1}{w_{\max}}\right) \cdot \|\theta_1 - \theta_2\|_w \tag{34}$$

where $w_{\max} = \max\limits_{i=1,\ldots,D} w_i \geq 1$. Thus, the Bellman operator $\mathbf{h}$ is $\left(1 - \frac{1}{w_{\max}}\right)$-contractive in the weighted ℓ_∞-norm, so that our general theory can be applied with this choice of Banach space.

4.1.2 Generative observation model

We analyze the `ROOT-SA` algorithm under a stochastic oracle known as the *generative observation model* for the SSP problem. For any state-action pair (x, u), the generative model allows us to draw next-state and cost samples from the MDP $(r, \mathbf{P})$. More precisely, we have access to a collection of n i.i.d. samples of the form $\{(\mathbf{Z}_k, C_k)\}_{k=1}^n$, where both $\mathbf{Z}_k$ and C_k are random matrices in $\mathbb{R}^{|\mathcal{X}_{-1}| \cdot |\mathcal{U}|}$. For each state-action pair (x, u), the entry $\mathbf{Z}_k(x, u)$ is drawn according to the transition kernel $\mathbf{P}_u(\cdot \mid x)$, whereas the entry $C_k(x, u)$ is a random variable with mean $c(x, u)$; this corresponds to a noisy observation of the cost function. We assume that the random cost $C_k(x, u)$ is upper bounded by $c_{\max}$ in absolute value. Here the cost samples $\{C_k(x, u)\}_{(x,u) \in \mathcal{X} \times \mathcal{U}}$ are independent across all state-action pairs, and the cost samples $\{C_k\}$ are independent of the transition samples $\{\mathbf{Z}_k\}$.

The empirical Bellman operator: Given a sample $(\mathbf{Z}, C)$ from our observation model, we define the single-sample empirical Bellman operator $\mathbf{H}(\cdot)$ on the space of Q-functions, whose action on a Q-function θ is given by

$$\mathbf{H}(\theta)(x, u) := C(x, u) + \sum_{x' \in \mathcal{X}_{-1}} \mathbf{Z}_u(x' \mid x) \min_{u' \in \mathcal{U}} \theta(x', u'). \tag{35}$$

Here we have introduced $\mathbf{Z}_u(x' \mid x) := \mathbf{1}_{\mathbf{Z}(x,u)=x'}$. We are ready to state our guarantees for the stochastic shortest path problem.

4.1.3 Guarantees for stochastic shortest path

It is easy to see that the operators $\mathbf{h}(\cdot)$ and $\mathbf{H}(\cdot)$, defined respectively in equations (33) and (35), satisfy Assumptions (A1)- (A3) with the weighted ℓ_∞-norm $\|\cdot\|_w$. In order to obtain an optimal policy from an estimate θ_n of the optimal Q function, it is natural to obtain performance bounds in the $\|\cdot\|_\infty$ norm, and we do so by invoking Corollaries 1 and 2 with $\|\cdot\|_C = \|\cdot\|_\infty$.

Accordingly, consider a Gaussian random vector W with $W \sim \mathcal{N}\left(0, \mathrm{cov}(\mathbf{H}(\theta^*) - \theta^*)\right)$, and define

$$\overline{W} = \mathbb{E}[\|W\|_\infty], \qquad \nu^2 := \sup_{x \in \mathcal{X}_{-1}, u \in \mathcal{U}} \mathbb{E}[W_{x,u}^2], \quad \text{and} \quad b_* := \frac{c_{\max}}{w_{\min}} + \|\theta^*\|_w. \tag{36}$$

For a given failure probability $\delta \in (0, 1)$, our result applies to the algorithm with parameters

$$\alpha = c_1 \left\{ \sqrt{n \log \left(|\mathcal{X}| \cdot |\mathcal{U}|\right)} \cdot \log(n/\delta) \right\}^{-1}, \quad \text{and} \quad B_0 = \frac{c_2 w_{\max}^2}{\alpha} \log(\tfrac{n}{\delta}), \qquad (37a)$$

We also choose the initialization θ_0 and the number of restarts R such that

$$\log \left(\frac{\|\theta_0 - \mathbf{h}(\theta_0)\| \sqrt{n}}{\mathcal{W}} \right) \le c_0 \log n \qquad \text{and} \qquad R \ge 2 c_0 \log n, \qquad (37b)$$

where c_0, c_1, c_2 are appropriate universal constants. We obtain the following guarantee:

Corollary 4. *Given a sample size n such that $\frac{n}{\log n} \ge c' \log(|\mathcal{X}| \cdot |\mathcal{U}|) \cdot w_{max}^4 \log(1/\delta)$, running Algorithm 1 with the tuning parameter choices (37) yields an estimate θ_n such that*

$$\|\mathbf{h}(\theta_n) - \theta_n\|_\infty \le \frac{c}{\sqrt{n}} \cdot \left\{ \overline{\mathcal{W}} + \nu \sqrt{\log(\tfrac{1}{\delta})} \right\} + c b_* w_{max}^2 \frac{\log(|\mathcal{X}| \cdot |\mathcal{U}|)}{n} \log^2(\tfrac{n}{\delta}),$$

with probability at least $1 - \delta$.

Note that when we invoke Corollary 1 to obtain this corollary, the second term is absorbed into the leading-order term under the sample size lower bound $\frac{n}{\log n} \ge c' \log(|\mathcal{X}| \cdot |\mathcal{U}|) \cdot w_{\max}^4 \log(1/\delta)$. In particular, the semi-norm domination factor is $\mathcal{D} = w_{\max}$ in this case, and we have the following inequalities:

$$\mathcal{D} \cdot \mathbb{E}\left[\|W\|_w\right] \le \frac{w_{\max}}{w_{\min}} \mathbb{E}\left[\|W\|_\infty\right] \le w_{\max} \overline{\mathcal{W}}, \quad \text{and}$$

$$\mathcal{D} \cdot \sup_{\|y\|_{1/w} \le 1} \sqrt{\mathbb{E}\left[\langle y, W \rangle^2\right]} \le \frac{w_{\max}}{w_{\min}} \sup_{\|y\|_\infty \le 1} \sqrt{\mathbb{E}\left[\langle y, W \rangle^2\right]} \le w_{\max} \nu,$$

which makes the second term of equation (13) dominated by the first term.

Next, in order to obtain an upper bound on the estimation error $\|\theta_n - \theta^*\|_\infty$ we need a few more definitions. For a given Q-function θ, we say π is a greedy policy of θ if and only if

$$\pi(x) = \arg\min_u \theta(x, u) \qquad \text{for all} \quad x \in \mathcal{X}_{-1},$$

and denote Π^θ as the set of all greedy policies of θ. Note that the greedy policies of a given Q-function may not be unique. Using this greedy policy, we can define the right-linear operator

$$\mathbf{P}^{\pi_\theta} \theta(x, u) = \sum_{x'} \mathbf{P}_u(x' \mid x) \theta(x', \pi_\theta(x')).$$

We also define a set $\mathcal{A}_s$ of linear operators as

$$\mathcal{A}_s = \{\mathbf{P}^{\pi_\theta} \mid \pi_\theta \text{ is a greedy policy of } \theta \text{ with } \theta \in \mathbb{B}(\theta^\star, s)\}. \qquad (38)$$

Let $\mathbb{B}(\theta^*, s) := \{\theta \mid \|\theta - \theta^*\|_\infty \le s\}$ denote the ℓ_∞-ball of radius s around θ^*. We use $\pi_\star$ to denote the greedy policy associated with the optimal Q-function $\theta^\star$. In Appendix G.2, we show that the local linearity assumption ((A4)') is satisfied for the Bellman operator (33) with the set of operators $\mathcal{A}_s$ from equation (38), and with $\|\cdot\|_C = \|\cdot\|_\infty$.

Given a tolerance probability $\delta \in (0, \frac{1}{\log(1/(1-\gamma))}))$, let s_n^* denotes the largest positive solution to the fixed-point equation

$$s_n = \frac{1}{\sqrt{n}}\Big\{ \mathbb{E}\Big[\sup_{\substack{\theta \in \mathbb{B}_\infty(\theta^*, s_n) \\ \pi \in \Pi^\theta}} \big\|(\mathcal{I} - \mathbf{P}^\pi)^{-1}W\big\|_\infty \Big] + \sup_{\substack{\theta \in \mathbb{B}_\infty(\theta^*, s_n), \pi \in \Pi^\theta \\ (x,u) \in \mathcal{X}_{-1} \times \mathcal{U}}} \Big(\mathbb{E}\big[\delta_{x,u}^\top(\mathcal{I} - \mathbf{P}^\pi)^{-1}W\big]\log(1/\delta) \Big)^{1/2} \Big\}$$

$$+ w_{\max}^2 \log(\tfrac{n}{\delta})\Big\{ \tfrac{w_{\max}}{w_{\min}}\big(\tfrac{\log(|\mathcal{X}_{-1}|\cdot|\mathcal{U}|)}{n}\big)^{3/4}\ \mathbb{E}[\|W\|_\infty] + b_* w_{\max}\tfrac{\log(|\mathcal{X}_{-1}|\cdot|\mathcal{U}|)}{n} \Big\}. \quad (39)$$

Here we have defined the indicator function $\delta_{x,u} = \mathbf{1}_{(x',u')=(x,u)}$. We obtain the following corollary:

Corollary 5. *Under the setup of Corollary 4, the estimate θ_n satisfies the bound*

$$\|\theta_n - \theta^*\|_\infty \leq c \cdot s_n^* \quad \text{with probability at least } 1 - \delta. \quad (40)$$

A few remarks are in order. First, the bound depends on the size of state-action space only poly-logarithmically, and depends on the quantity $w_{\max}$ through two sources: the contraction parameter and the norm domination factor between $\| \cdot \|_\infty$ and $\| \cdot \|_w$. Second, let Π^* be the set of all optimal policies for the SSP problem, for sample size n large enough,[2] the ball $\mathbb{B}_\infty(\theta^*, s_n)$ will eventually shrink to the singleton θ^*, and the supremum in the fixed-point equation (39) is taken over $\pi \in \Pi^*$. Therefore, using $\mathcal{H}_n$ to denote higher-order terms, the solution s_n takes the form

$$s_n = \frac{1}{\sqrt{n}}\Big\{ \mathbb{E}\Big[\sup_{\pi \in \Pi^*} \big\|(\mathcal{I} - \mathbf{P}^\pi)^{-1}W\big\|_\infty \Big] + \sup_{\substack{\pi \in \Pi^* \\ (x,u) \in \mathcal{X}_{-1} \times \mathcal{U}}} \Big(\mathbb{E}\big[\delta_{x,u}^\top(\mathcal{I} - \mathbf{P}^\pi)^{-1}W\big]\log(\tfrac{1}{\delta}) \Big)^{1/2} \Big\} + \mathcal{H}_n$$

$$\leq \frac{1}{\sqrt{n}} \max_{\substack{\pi \in \Pi^* \\ (x,u) \in \mathcal{X}_{-1} \times \mathcal{U}}} \sqrt{\mathbb{E}\Big[\big(\delta_{x,u}^\top(\mathcal{I} - \mathbf{P}^\pi)^{-1}(\mathbf{H}(\theta^*) - \theta^*)\big)^2 \Big]} \cdot \sqrt{\log \frac{|\mathcal{X}|\cdot|\mathcal{U}|\cdot|\Pi^*|}{\delta}} + \mathcal{H}_n.$$

Up to a factor of $\sqrt{\log \frac{|\mathcal{X}|\cdot|\mathcal{U}|\cdot|\Pi^*|}{\delta}}$, this matches the two-point lower bound in the paper [Kha+21] (in the discounted MDP case). When specializing to the cases where the optimal policy is unique, or satisfies the Lipschitz-type assumptions in the paper [Kha+21], the upper bound above also recovers the leading-order term in that paper. We conjecture that the leading-order term of the solution s_n to the fixed-point equation is actually optimal for large n. It would be interesting to verify this conjecture, and establish some kind of optimality over suitably defined problem classes.

When specialized to the γ-discounted MDPs, the sample size requirement in Corollary 5 scales as $\mathcal{O}\big((1-\gamma)^{-4}\big)$. This requirement is worse than corresponding requirements in the paper [Kha+21], at least in certain regimes. Intuitively, this is the price we pay when moving to the general case where only the contraction of the population-level operator is assumed, instead of the sample-level contraction.

4.2   Average cost policy evaluation

As a second illustration, we turn to a problem where the operator is *not* contractive, but does satisfy a form of multi-step contractivity needed to apply the theory from Section 3.3. This

---

[2]The sample size requirement may depend on the gap between the value of optimal and sub-optimal actions, as in the prior work [Kha+21].

example also involves an error measure that is only a semi-norm in the original space, but can converted to a norm in a Banach space by taking a suitable quotient.

More specifically, consider an undiscounted Markov reward process (MRP) with state space $\mathcal{X}$, probability transition kernel $\mathbf{P} \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ and cost function $c : \mathcal{X} \to \mathbb{R}$. When the Markov chain is irreducible and ergodic, there is a unique stationary distribution ξ. Letting $\mu_* := \mathbb{E}_{x \sim \xi}[c(x)]$ denote the average cost under this stationary distribution, our goal is to estimate the value function

$$\theta^*(x) := \sum_{t=0}^{\infty} \mathbf{P}^t \{ c(x) - \mu_* \}.$$

It is known that the value function θ^* and average cost μ_* jointly satisfy the Bellman equation

$$\mu_* + \theta^*(x) - \mathbf{P}\theta^*(x) - c(x) = 0 \quad \text{for all } x \in \mathcal{X}. \tag{41}$$

See the sources [Der66; TVR99] for more background.

In practical application of policy evaluations problems, of primary interest are the relative differences between the value function at different state-actions pairs. Thus, the primary goal is to estimate the function θ^*, with the average cost μ_* being a nuisance parameter. As shown in the sequel (see Section 4.2.1), by considering the span semi-norm in an appropriate vector space $\mathbb{V}$, it is possible to estimate θ^* without estimating μ_*.

Observation models and relevant operators: As before, we consider a generative observation model, where we observe a collection of n i.i.d. samples of the form $\{(\mathbf{Z}_k, C_k)\}_{k=1}^n$, where both $\mathbf{Z}_k \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{X}|}$ and $C_k \in \mathbb{R}^{|\mathcal{X}|}$. For each state $x \in \mathcal{X}$, the row x of the matrix $\mathbf{Z}_k$ is an indicator vector $\mathbf{1}_{s'}$, where the state s' is drawn according to the transition kernel $\mathbf{P}(\cdot \mid x)$; the entry $C_k(x)$ is a random variable with mean $c(x)$ and uniformly bounded by σ_r, corresponding to a noisy observation of the reward function.

The population and empirical Bellman operators for the average-cost policy evaluation can be written as follows:

$$\mathbf{h}(\theta) := \mathbf{P}\theta + c, \quad \text{and} \quad \mathbf{H}_k(\theta) := \mathbf{Z}_k\theta + C_k.$$

It can be seen that both $\mathbf{h}$ and $\mathbf{H}_k$ are linear operators, satisfying $\mathbb{E}[\mathbf{H}_k] = \mathbf{h}$.

In the rest of this subsection, we define a semi-norm and discuss the multi-step contraction properties of the operator $\mathbf{h}$, and then present the main consequences of Theorem 3 and Corollary 3 for such models.

4.2.1 The semi-norm and multi-step contraction

Consider the Banach space $\mathbb{V}$ given by

$$\mathbb{V} = \mathbb{R}^{|\mathcal{X}|} / \{ \theta + \alpha\mathbf{1} \mid \alpha \in \mathbb{R} \}, \tag{42}$$

where each element of $\mathbb{V}$ is an equivalence class of the form $\{ \theta + \alpha\mathbf{1} : \alpha \in \mathbb{R} \}$, equipped with the span norm

$$\|\theta\|_{\text{span}} := \max_{x \in \mathcal{X}} \theta(s) - \min_{x \in \mathcal{X}} \theta(s) \quad \text{for all } \theta \in \mathbb{V}.$$

20

Note that $\|\cdot\|_{\mathrm{span}}$ is a semi-norm on $\mathbb{R}^{\mathcal{X}}$, but a norm on the quotient space $\mathbb{V}$. For reinforcement learning problems, this choice is natural, since we often care only about the relative advantages of state-action pairs, in which case the average cost μ_* is irrelevant.

Under the norm $\|\cdot\|_{\mathrm{span}}$ on $\mathbb{V}$, the operator $\mathbf{h}$ is non-expansive, but not necessarily a contraction. However, under suitable conditions, it can be shown to contractive in a multi-step sense. In order to do so, we impose the following mixing time condition.

Assumption: Mixing time

(MT) There exists a positive integer t_{mix} such that

$$d_{\mathrm{TV}}(\delta_x^\top \mathbf{P}^{t_{\mathrm{mix}}}, \delta_y^\top \mathbf{P}^{t_{\mathrm{mix}}}) \le \tfrac{1}{2} \qquad \text{for any } x, y \in \mathcal{X}.$$

Here the vector $\delta_x \in \mathbb{R}^{\mathcal{X}}$ is the unit basis vector with a single one in entry $x \in \mathcal{X}$.

Under Assumption (MT), for any $\theta \in \mathbb{V}$, we have

$$\|\mathbf{P}^{2t_{\mathrm{mix}}}\theta\|_{\mathrm{span}} = \max_{x \in \mathcal{X}}\left\{\delta_x^\top \mathbf{P}^{2t_{\mathrm{mix}}}\theta\right\} - \min_{x \in \mathcal{X}}\left\{\delta_x^\top \mathbf{P}^{2t_{\mathrm{mix}}}\theta\right\} \overset{(i)}{\le} 2\max_{x \in \mathcal{X}}\left|\delta_x^\top \mathbf{P}^{2t_{\mathrm{mix}}}\theta - \xi^\top \mathbf{P}^{2t_{\mathrm{mix}}}\theta\right|$$

$$\le 2d_{\mathrm{TV}}(\delta_x \mathbf{P}^{2t_{\mathrm{mix}}}, \xi \mathbf{P}^{2t_{\mathrm{mix}}}) \cdot \|\theta\|_{\mathrm{span}} \overset{(ii)}{\le} \tfrac{1}{2}\|\theta\|_{\mathrm{span}}, \quad (43)$$

where step (i) is a direct consequence of triangle inequality, and in step (ii), we exploits the bound $d_{\mathrm{TV}}(\delta_x \mathbf{P}^{2t_{\mathrm{mix}}}, \xi \mathbf{P}^{2t_{\mathrm{mix}}}) \le \tfrac{1}{2} d_{\mathrm{TV}}(\delta_x \mathbf{P}^{t_{\mathrm{mix}}}, \xi \mathbf{P}^{t_{\mathrm{mix}}}) \le \tfrac{1}{4}$, obtained by applying the mixing time condition (MT) twice. Consequently, we see that the multi-step contraction assumption (A1)$'$ holds if the operator is composed $m = 2t_{\mathrm{mix}}$ times.

4.2.2   Estimation error upper bounds

Having defined the norm $\|\cdot\|_{\mathrm{span}}$ and the established the multi-step contraction property (43), we are ready to derive a guarantee for average-cost policy evaluation. This involves the Gaussian random variable

$$W \sim \mathcal{N}\left(0, \mathrm{cov}\left(\mathbf{H}(\theta^*) - \theta^*\right)\right),$$

as well as $\overline{\mathcal{W}} := \mathbb{E}\left[\|W\|_{\mathrm{span}}\right]$. For a given failure probability $\delta \in (0,1)$, our result applies to the algorithm with parameters

$$\alpha = c_1\left\{\sqrt{n\log|\mathcal{X}|}\cdot\log(\tfrac{n}{\delta})\right\}^{-1}, \quad \text{and} \quad B_0 = \tfrac{ct_{\mathrm{mix}}}{\alpha}\log(\tfrac{n}{\delta}). \tag{44a}$$

We also choose the initialization θ_0 and the number of restarts R such that

$$\log\left(\frac{\|\theta_0 - \mathbf{h}(\theta_0)\|\sqrt{n}}{\overline{\mathcal{W}}}\right) \le c_0\log n \qquad \text{and} \qquad R \ge 2c_0\log n, \tag{44b}$$

where c, c_0, c_1 are appropriate universal constants. We have the following guarantee:

Corollary 6. *Suppose Assumption (MT) holds, and the sample size n is lower bounded as $\frac{n}{\log^2 n} \ge c' t_{\mathrm{mix}}^2 \log(|\mathcal{X}|)\cdot\log(1/\delta)$. Then the estimate θ_n from Algorithm 1, obtained using tuning parameters satisfying conditions (44), satisfies the bound*

$$\|\theta_n - \theta^*\|_{\mathrm{span}} \le \frac{c}{\sqrt{n}}\left\{\mathbb{E}\left[\|(\mathcal{I}-\mathbf{P})^\dagger W\|_{\mathrm{span}}\right] + \sqrt{\sup_{x_1, x_2 \in \mathcal{X}}\mathbb{E}\left[((\delta_{x_1} - \delta_{x_2})(\mathcal{I}-\mathbf{P})^\dagger W)^2\right]\log(1/\delta)}\right\}$$

$$+ ct_{\mathrm{mix}}\left\{\left[\tfrac{\log|\mathcal{X}|}{n}\right]^{3/4}\overline{\mathcal{W}} + \tfrac{\log|\mathcal{X}|}{n}\left(\sigma_r + \|\theta^*\|_{\mathrm{span}}\right)\right\}\log^2(\tfrac{n}{\delta}), \quad (45)$$

with probability at least $1 - \delta$.

A few remarks are in order. First, the linear operator $(\mathcal{I} - \mathbf{P})$ is not invertible in $\mathbb{R}^{\mathcal{X}}$, with the all-one vector lying in its nullspace. However, it is invertible in the quotient space $\mathbb{V}$, with the pseudo-inverse $(\mathcal{I} - \mathbf{P})^\dagger$ being a representation of its inverse in the coordinate system of $\mathbb{R}^{\mathcal{X}}$, which appears in the bound. Second, as with the previous two cases, the bound depends on the size of state space only poly-logarithmically; it depends quadratically on the mixing time t_{mix}, as shown through the required lower bound on n. Taking the γ-discounted MRP as a special case of the average-cost framework,[3] Corollary 6 improves the results of previous work [Kha+20] in two aspects:

- Corollary 6 is valid whenever sample size satisfies $n \gtrsim (1 - \gamma)^{-2}$ up to log factors, which improves the previous $(1 - \gamma)^{-3}$ dependency from the past work;
- The instance-dependent quantity in the paper [Kha+20] is replaced with an optimal one matching the local asymptotic minimax limit.

These improvements are made possible by making use of the linear structure in policy evaluation problems. More importantly, Corollary 6 applies to a more general class of problems, where the mixing time t_{mix} replaces the role of effective horizon.

In terms of other related work, the quadratic mixing time dependence (i.e., sample size scaling as $\mathcal{O}(t_{\mathrm{mix}}^2)$) matches that of the paper [JS20]. On one hand, our results are more refined in that we give instance-dependent guarantees. On the other hand, their results apply to Markov decision processes with actions. Thus, an open and interesting direction of future work is to extend our instance-dependent bounds to the case of average-cost MDPs with policy optimization.

5 Proofs

This section is devoted to the proofs of our main results—namely, Theorems 1 and 2—along with the associated corollaries. So as to facilitate reading of the proofs, we reproduce here the two main recursions that define the algorithm:

$$v_t = \mathbf{H}_t(\theta_{t-1}) - \theta_{t-1} + \tfrac{t-1}{t} \left(v_{t-1} - \mathbf{H}_t(\theta_{t-2}) + \theta_{t-2} \right), \quad \text{and} \tag{46a}$$

$$\theta_t = \theta_{t-1} + \alpha v_t. \tag{46b}$$

Throughout the proofs, we make use of the shorthand $\overline{\mathcal{W}} = \mathbb{E}[\|W\|]$, and $\nu = \sigma_\Gamma(W)$.

5.1 Proof of Theorem 1

Our proof is based on a bootstrapping argument, and can be broken down into four steps:

- First, we establish recursions that relate $\|\mathbf{h}(\theta_t) - \theta_t\|$ and $\|v_t\|$.
- Second, we prove coarse upper bounds on $\|\mathbf{h}(\theta_t) - \theta_t\|$ and $\|v_t\|$.
- Third, starting with the sub-optimal bounds from Step 2, we iteratively refine them using a bootstrapping argument and the recursions from Step 1.
- In the fourth step, we improve higher-order terms in the bounds.

For the purposes of analysis, it is useful to define the auxiliary sequence

$$z_t := \left\{ \mathbf{h}(\theta_{t-1}) - \theta_{t-1} \right\} - v_t, \quad \text{for } t = B_0, B_0 + 1, \cdots. \tag{47}$$

---

[3]This can be done by adding an absorbing state $\perp$ to the state space. At a rate of $(1 - \gamma)$, the Markov process is killed and moved to the absorbing state. In such case, the unique stationary distribution is the singleton at $\perp$, and the mixing time assumption is satisfied with $t_{\mathrm{mix}} = \frac{c}{1-\gamma}$ for universal constant $c > 0$.

Our strategy is to control $\|\mathbf{h}(\theta_t) - \theta_t\|$ by proving upper bounds on $\|z_{t+1}\|$ and $\|v_{t+1}\|$.

Let $r_\theta(t)$ and $r_v(t)$, respectively, denote high probability bounds on the quantities $\|\mathbf{h}(\theta_t) - \theta_t\|$ and $\|v_t\|$. It is useful to introduce the notion of an *admissible sequence*: for some $\kappa \geq 0$, the sequence $\{r(t)\}_{t \geq B_0}$ is said to be *κ-admissible* if

(i) The sequence $\{r(t)\}_{t \geq B_0}$ is non-increasing.

(ii) The sequence $\{t^\kappa \cdot r(t)\}_{t \geq B_0}$ is non-decreasing.

We say that the sequence is admissible if it is κ-admissible for some $\kappa \geq 0$. For notational simplicity, we sometimes use the sequences with time index less than B_0, in such cases, we denote $r_v(t) := r_v(B_0)$ and $r_v(t) := r_\theta(B_0)$ for $t \in [1, B_0]$.

Observe that κ-admissible sequences are also β-admissible sequences for any $\beta > \kappa$. For the sake of notational convenience, we use the shorthands r_θ and r_v to denote the estimate sequences $\{r_\theta(t)\}_{t \geq B_0}$ and $\{r_v(t)\}_{t \geq B_0}$, respectively. Given an admissible pair (r_θ, r_v) and an integer $n > 0$, define the events

$$\mathcal{E}_n^{(\theta)}(r_\theta) := \left\{ \sup_{B_0 \leq t \leq n} \frac{\|\mathbf{h}(\theta_t) - \theta_t\|}{r_\theta(t)} \leq 1 \right\}, \quad \text{and} \quad \mathcal{E}_n^{(v)}(r_v) := \left\{ \sup_{B_0 \leq t \leq n} \frac{\|v_t\|}{r_v(t)} \leq 1 \right\}. \quad (48)$$

A key portion of our proof involves ensuring that the estimate sequences r_θ and r_v are κ-admissible for carefully chosen values of κ. With these concepts and notation in place, we are now ready to start the main argument.

5.1.1 Step 1: Relation between $\|\mathbf{h}(\theta_t) - \theta_t\|$ and $\|v_t\|$

From the definition (47), we have the relation $\mathbf{h}(\theta_t) - \theta_t = z_{t+1} + v_{t+1}$. As mentioned before, we prove an upper bound on $\|\mathbf{h}(\theta_t) - \theta_t\|$ by proving upper bounds on $\|z_{t+1}\|$ and $\|v_{t+1}\|$. We do so using two auxiliary lemmas, the first of which depends on a stepsize α satisfying the bound (8a)—namely:

$$\alpha \leq \frac{(1-\gamma)^2}{cL^2 \mathcal{J}_2^2(\Gamma, \rho_n) \log\left(\frac{n}{\delta}\right)}. \quad (49)$$

Lemma 1. *Suppose that Assumptions (A1), (A3) and (A2) are in force, and that (r_θ, r_v) are κ-admissible sequences for some $\kappa \in [0, 2]$. Then given a stepsize α satisfying the bound (49) and a burn-in period $B_0 \geq \frac{100}{(1-\gamma)\alpha}$ conditioned on the event $\mathcal{E}_n^{(v)}(r_v) \cap \mathcal{E}_n^{(\theta)}(r_\theta)$, for each $t \in [B_0, n]$ we have*

$$\|v_t\| \leq \frac{1+\gamma}{2} r_v(t) + \frac{8}{t\alpha} r_\theta(t) + \frac{c}{\sqrt{\alpha}} \left\{ \overline{W} + \nu \sqrt{\log(\frac{1}{\delta})} \right\} + \frac{cb_*}{t} \left\{ \log(\frac{1}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n) \right\} + 6(1-\gamma)\left(\frac{B_0}{t}\right)^2 \|v_{B_0}\|, \quad (50)$$

with probability at least $1 - \delta$.

See Section 5.3.1 for the proof of this lemma.

Lemma 2. *Under the same conditions as Lemma 1, for each $t \in [B_0, n]$, we have:*

$$\|z_t\| \leq \frac{c}{\sqrt{t}}\left\{\overline{W} + \nu\sqrt{\log(\frac{1}{\delta})}\right\} + \frac{b_*}{t}\left\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\frac{1}{\delta})\right\}$$

$$+ \frac{cL}{t}\left\{\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\frac{1}{\delta})}\right\}\left\{\alpha\Big(\sum_{s=B_0}^{t-1} s^2 r_v^2(s)\Big)^{1/2} + \frac{1}{1-\gamma}\Big(\sum_{s=1}^{t-1} r_\theta^2(s)\Big)^{1/2}\right\}, \quad (51)$$

with probability $1 - \delta$.

This lemma is a special case of Lemma 5, which is proved in Section 5.3.2.

Note that although the two lemmas are for a single time index $t \in [B_0, n]$, it is easy to transform them to guarantees that are uniform over $t \in [B_0, n]$. In particular, applying a union bound for $t = B_0, B_0 + 1, \cdots, n$, and by replacing δ with $\delta' = \delta/n$, the bounds (50) and (51) are valid uniformly over $t \in [B_0, n]$.

We use these two lemmas in our bootstrapping argument. In particular, beginning with the relation $\mathbf{h}(\theta_t) - \theta_t = z_{t+1} + v_{t+1}$, applying the triangle inequality yields the bound $\|\mathbf{h}(\theta_t) - \theta_t\| \leq \|z_{t+1}\| + r_v(t+1)$ on the event $\mathcal{E}_n^{(v)}(r_v)$. Our analysis shows that by starting with an initial estimate $(r_\theta(t), r_v(t))$, the bounds (50) and (51) allow us to obtain an improved estimate $(r_\theta^+(t), r_v^+(t))$ such that

$$\|\mathbf{h}(\theta_t) - \theta_t\| \leq r_\theta^+(t) < r_\theta(t), \quad \text{and}$$
$$\|v_t\| \leq r_v^+(t) < r_v(t)$$

with high probability. We quantify the improvement in $(r_\theta^+(t), r_v^+(t))$, and repeatedly apply this argument so as to "bootstrap" the bound and ultimately obtain sharp estimates for $r_\theta(t)$ and $r_v(t)$.

5.1.2 Step 2: Setup for the bootstrapping argument

Throughout this step, we require that the estimate sequences r_θ and r_v be $\frac{1}{2}$-admissible and 1-admissible, respectively. As shown in this section, these choices allow us to obtain upper bounds on $\|\mathbf{h}(\theta_t) - \theta_t\|$ and $\|v_t\|$ that decay at the rates $1/\sqrt{t}$ and $1/t$, respectively.

We assume that the pair (r_v^+, r_θ^+) satisfy the initialization condition

$$r_v^+(B_0) \geq \|v_{B_0}\|, \quad \text{and} \quad r_\theta^+(B_0) \geq \|\mathbf{h}(\theta_0) - \theta_0\|, \quad (52a)$$

and for each integer $t \in [B_0, n]$, the bounds

$$r_v^+(t) \geq \frac{1+\gamma}{2}r_v(t) + \frac{8}{\alpha t}r_\theta(t) + \frac{c}{t\sqrt{\alpha}}\left\{\overline{W} + \nu\sqrt{\log(\frac{n}{\delta})}\right\}$$

$$+ \frac{cb_*}{t}\left\{\log(\frac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\right\} + 6(1-\gamma)\left(\frac{B_0}{t}\right)^2\|v_{B_0}\|, \quad (52b)$$

and

$$r_\theta^+(t) \geq \left\{1 + c\alpha\sqrt{t}L\left[\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log\left(\frac{n}{\delta}\right)}\right]\right\}r_v(t) + \frac{2cL}{(1-\gamma)\sqrt{t}}\mathcal{J}_2(\mathbb{B}^*, \rho_n)\log(\frac{n}{\delta}) \cdot r_\theta(t)$$

$$+ \frac{c}{\sqrt{t}}\left\{\overline{W} + \nu\sqrt{\log\left(\frac{n}{\delta}\right)}\right\} + \frac{cb_*}{t}\left\{\log(\frac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\right\}. \quad (52c)$$

24

Under these conditions, by combining the bounds (50) and (51) and applying a union bound over $t \in [B_0, n]$, we find that

$$\mathbb{P}\Big[\mathcal{E}_n^{(\theta)}(r_\theta^+) \cap \mathcal{E}_n^{(v)}(r_v^+)\Big] \geq \mathbb{P}\Big[\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)\Big] - \delta,$$

valid for any pair (r_v, r_θ) that are $\frac{1}{2}$ and 1-admissible, respectively.

We consider sequences of a particular form $r_v^{(i)}(t) = \frac{\psi_v^{(i)}}{t\sqrt{\alpha}}$ and $r_\theta^{(i)}(t) = \frac{\psi_\theta^{(i)}}{\sqrt{t}}$, for pairs of positive reals $(\psi_v^{(i)}, \psi_\theta^{(i)})$ independent of t. Clearly, with such forms, the sequence $r_\theta^{(i)}$ is $\frac{1}{2}$-admissible, and the sequence $r_v^{(i)}$ is 1-admissible. However, if we directly substitute the sequences $(r_v^{(i)}(t), r_v^{(i)}(t))$ of such forms into the relations (52b)-(52c), the resulting sequences (r_θ^+, r_v^+) are no longer be of the desired form. So in order to unify the coefficients in equations (52b)-(52c) into the same time scale, given $\alpha > 0$, we define the burn-in time

$$B_0 = \frac{c}{(1-\gamma)^2\alpha}\log(\tfrac{n}{\delta}). \tag{53a}$$

For each $t = B_0, B_0 + 1\ldots$, the coefficients in (52b) and (52c) then satisfy the bounds

$$\frac{8}{\alpha t} \leq \frac{1-\gamma}{3}\cdot\frac{1}{\sqrt{\alpha t}}, \quad \frac{1}{\sqrt{\alpha t}} \leq \frac{1-\gamma}{6}, \quad \text{and} \quad \frac{1}{(1-\gamma)\sqrt{t}}\log(\tfrac{n}{\delta}) \leq \sqrt{\alpha}, \tag{53b}$$

Therefore, if we construct a two-dimensional vector sequence $\psi^{(i)} = \begin{bmatrix} \psi_v^{(i)} & \psi_\theta^{(i)} \end{bmatrix}^T$ satisfying the recursive relation $\psi^{(i+1)} = Q\psi^{(i)} + b$, where

$$Q := \begin{bmatrix} \frac{1+\gamma}{2} & \frac{1-\gamma}{3} \\ \frac{1-\gamma}{6} + cL\mathcal{J}_2(\mathbb{B}^*, \rho_n)\sqrt{\alpha}\log(\tfrac{n}{\delta}) & 2cL\mathcal{J}_2(\mathbb{B}^*, \rho_n)\log(\tfrac{n}{\delta})\sqrt{\alpha} \end{bmatrix}, \quad \text{and} \tag{54a}$$

$$b := \begin{bmatrix} c\big\{\overline{\mathcal{W}} + \nu\sqrt{\log(\tfrac{n}{\delta})}\big\} + cb_*\sqrt{\alpha}\big\{\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\big\} + (1-\gamma)B_0\sqrt{\alpha}\|v_{B_0}\| \\ \big\{\overline{\mathcal{W}} + \nu\sqrt{\log(\tfrac{n}{\delta})}\big\} + cb_*\sqrt{\alpha}\big\{\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\big\} + \sqrt{B_0}\|\mathbf{h}(\theta_0) - \theta_0\| \end{bmatrix} \tag{54b}$$

satisfy the requirement (78). Thus, we are led to the probability bound

$$\mathbb{P}\Big[\mathcal{E}_n^{(\theta)}(r_\theta^{(i+1)}) \cap \mathcal{E}_n^{(v)}(r_v^{(i+1)})\Big] \geq \mathbb{P}\Big[\mathcal{E}_n^{(\theta)}(r_\theta^{(i)}) \cap \mathcal{E}_n^{(v)}(r_v^{(i)})\Big] - \delta, \tag{55}$$

for the sequences $r_\theta^{(i)}(t) = \psi_\theta^{(i)}/\sqrt{t}$ and $r_v^{(i)}(t) = \psi_v^{(i)}/(\sqrt{\alpha}t)$. In order to initialize the argument, we need a coarse bound on the pair $(\|v_t\|, \|\mathbf{h}(\theta_t) - \theta_t\|)$; the following lemma provides the requisite bound:

Lemma 3. *Under Assumptions (A3) and (A2), we have*

$$\|\theta_t - \theta^*\| + \|v_t\| \leq e^{1+L\alpha t}\big(b_* + \|\theta_0 - \theta^*\|\big),$$

almost surely for each $t = 0, 1, 2, \ldots$.

See Appendix F.1 for the proof of this claim.

Based on Lemma 3, it follows that for each integer $t \in [1, n]$, we have (almost surely) the bound

$$\|v_t\| \leq r_v^{(0)}(t) := \frac{n}{t}e^{1+L\alpha t}\big\{b_* + \|\theta_0 - \theta^*\|\big\}, \quad \text{and}$$

$$\|\mathbf{h}(\theta_t) - \theta_t\| \stackrel{(i)}{\leq} r_\theta^{(0)}(t) := 2\cdot\sqrt{\frac{n}{t}}e^{1+L\alpha t}\big\{b_* + \|\theta_0 - \theta^*\|\big\},$$

where step (i) follows from the bound $\|\mathbf{h}(\theta_t) - \theta_t\| \leq \|\theta_t - \theta^*\| + \|\mathbf{h}(\theta_t) - \mathbf{h}(\theta^*)\| \leq 2\cdot\|\theta_t - \theta^*\|$.

By construction, the sequences $r_v^{(0)}$ and $r_\theta^{(0)}$ are 1-admissible and $\frac{1}{2}$-admissible, respectively, and by Lemma 3, the event $\mathcal{E}_n^{(\theta)}(r_\theta^{(0)}) \cap \mathcal{E}_n^{(v)}(r_v^{(0)})$ happens almost surely.

5.1.3 Step 3: Bootstrapping step

Recursing the bound (55) for i steps yields

$$\mathbb{P}\big[\mathcal{E}_n^{(v)}(r_v^{(i)}) \cap \mathcal{E}_n^{(\theta)}(r_\theta^{(i)})\big] \geq \mathbb{P}\big[\mathcal{E}_n^{(v)}(r_v^{(0)}) \cap \mathcal{E}_n^{(\theta)}(r_\theta^{(0)})\big] - i\delta = 1 - i\delta.$$

It remains to analyze the sequence $\psi^{(i)} = \begin{bmatrix} \psi_v^{(i)} & \psi_\theta^{(i)} \end{bmatrix}^T$ as the number of bootstrap steps i increases. We do so by analyzing the recursion relation $\psi^{(i+1)} = Q\psi^{(i)} + b$ with the matrix Q given in equation (54).

Observe that the stepsize condition (49) ensures that

$$cL\mathcal{J}_2(\mathbb{B}^*, \rho_n) \log(\tfrac{n}{\delta}) \cdot \sqrt{\alpha} \leq \tfrac{1-\gamma}{6}. \tag{56}$$

Consequently, the matrix Q from equation (49) is entrywise upper bounded by the matrix

$$\widetilde{Q} = \begin{bmatrix} \frac{1+\gamma}{2} & \frac{1-\gamma}{3} \\ \frac{1-\gamma}{3} & \frac{1}{2} \end{bmatrix}$$

This fact implies that for any vector $u \in \mathbb{R}^2$ with non-negative entries, we have the upper bound $Qu \preceq_{\mathrm{orth}} \widetilde{Q}u$, where $\preceq_{\mathrm{orth}}$ denotes the orthant ordering. Straightforward calculation yields the bound $\|\widetilde{Q}\|_{\mathrm{op}} \leq 1 - \frac{1-\gamma}{8}$. Putting together the pieces, we find that for each $N = 1, 2, \ldots$, conditioned on the event $\mathcal{E}_n^{(v)}(r_v^{(N)}) \cap \mathcal{E}_n^{(\theta)}(r_\theta^{(N)})$, we have

$$\psi^{(N)} = \Big(\sum_{i=0}^{N-1} Q^i \Big) b_\psi + Q^N \begin{bmatrix} \psi_v^{(0)} \\ \psi_\theta^{(0)} \end{bmatrix} \preceq_{\mathrm{orth}} \Big(\sum_{i=0}^{N-1} \widetilde{Q}^i \Big) b + \widetilde{Q}^N \begin{bmatrix} \psi_v^{(0)} \\ \psi_\theta^{(0)} \end{bmatrix}$$

$$\preceq_{\mathrm{orth}} (I - \widetilde{Q})^{-1} b + e^{\frac{(1-\gamma)}{8}N} \big(\psi_v^{(0)} + \psi_\theta^{(0)} \big) \mathbf{1}_2.$$

We take $N = \lceil \frac{cLn}{1-\gamma} \log n \rceil$. Replacing δ with δ/N and substituting into the above inequalities then yields

$$t\sqrt{\alpha} \cdot \|v_t\| \leq \psi_v^{(N)} \leq \tfrac{c}{1-\gamma} \Big\{ \overline{W} + \nu \sqrt{\log(\tfrac{n}{\delta})} \Big\} + \tfrac{cb_*\sqrt{\alpha}}{1-\gamma} \Big\{ \log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n) \Big\}$$
$$+ cB_0\sqrt{\alpha}\|v_{B_0}\| + \sqrt{B_0}\|\mathbf{h}(\theta_0) - \theta_0\|, \quad (57a)$$

and

$$\sqrt{t}\|\mathbf{h}(\theta_t) - \theta_t\| \leq \psi_\theta^{(N)} \leq c\Big\{ \overline{W} + \nu \sqrt{\log(\tfrac{n}{\delta})} \Big\} + cb_*\sqrt{\alpha} \Big\{ \log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n) \Big\}$$
$$+ cB_0(1-\gamma)\sqrt{\alpha}\|v_{B_0}\| + \sqrt{B_0}\|\mathbf{h}(\theta_0) - \theta_0\|, \quad (57b)$$

with probability at least $1 - \delta$, uniformly for each $t \in B_0, B_0 + 1, \cdots, n$.

It remains to provide upper bounds on $\|v_{B_0}\|$.

Lemma 4. *Under Assumptions (A1) and (A3), and a burn-in period given by equation* (53a), *we have*

$$\|v_{B_0}\| \leq 2\|\mathbf{h}(\theta_0) - \theta_0\| + \tfrac{c}{\sqrt{B_0}} \Big\{ \overline{W} + \nu \sqrt{\log(\tfrac{1}{\delta})} \Big\} + \tfrac{cb_*}{B_0} \Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \Big\}.$$

with probability at least $1 - \delta$.

26

See Appendix F.2 for the proof.

Combining Lemma 4 and bound (57a), we find that

$$\|\mathbf{h}(\theta_t) - \theta_t\| \leq \frac{c}{\sqrt{t}}\left(\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\right) + \frac{cb_*\sqrt{\alpha}}{\sqrt{t}}\left\{\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\right\}$$

$$+ \|\mathbf{h}(\theta_0) - \theta_0\|\frac{\sqrt{B_0}}{\sqrt{t}}\log^{\frac{3}{2}}(\tfrac{n}{\delta}) \quad (58)$$

with probability at least $1 - \delta$, uniformly for all integers $t \in [B_0, n]$.

Although this bound has optimal dependence on $\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}$, its dependence on the terms $\|\mathbf{h}(\theta_0) - \theta_0\|$ and $\mathcal{J}_1(\Gamma, \rho_n)$ and $\log(n/\delta)$ in the bound (58) can be sharpened. This motivates the second phase of the bootstrap argument.

5.1.4 Step 4: Improving higher-order terms

Given the pair $(\psi_v^{(N)}, \psi_\theta^{(N)})$ defined by[4] the right-hand side of (57), conditioned on the event $\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)$ with $r_v(t) = \frac{\psi_v^{(N)}}{t\sqrt{\alpha}}$ and the sequence $r_\theta(t) = \frac{\psi_\theta^{(N)}}{\sqrt{t}}$, invoking the bound (51) from Lemma 2 we have

$$\|\mathbf{h}(\theta_t) - \theta_t\| \leq \frac{c}{\sqrt{t}}\left(\overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})}\right) + \frac{cb_*}{t}\left(\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta})\right)$$

$$+ \left\{\frac{1}{t} + c\frac{\alpha L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log(\tfrac{n}{\delta})\right\}\frac{\psi_v^{(N)}}{\sqrt{\alpha}} + 2c\frac{L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{(1-\gamma)t}\log(\tfrac{n}{\delta}) \cdot \psi_\theta^{(N)}$$

$$\leq \frac{c}{\sqrt{t}}\left\{\overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})}\right\} + c'b_*\left\{\frac{1}{(1-\gamma)t} + \frac{\alpha L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{(1-\gamma)\sqrt{t}} \cdot \log(\tfrac{n}{\delta})\right\}\left\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta})\right\}$$

$$+ c'\left\{\frac{1}{(1-\gamma)t\sqrt{\alpha}} + \frac{L\mathcal{J}_2(\mathbb{B}^*, \rho_n)\sqrt{\alpha}}{(1-\gamma)\sqrt{t}} \cdot \log(\tfrac{n}{\delta}) + \frac{L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{(1-\gamma)t}\log(\tfrac{n}{\delta})\right\}\left\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\right\}$$

$$+ c'\left\{\frac{1}{t\sqrt{\alpha}} + \frac{\sqrt{\alpha}L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log(\tfrac{n}{\delta}) + \frac{L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{(1-\gamma)t}\log(\tfrac{n}{\delta})\right\}\sqrt{B_0}\|\mathbf{h}(\theta_0) - \theta_0\|$$

$$+ c'\left\{\frac{1}{t} + \frac{\alpha L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log(\tfrac{n}{\delta}) + \frac{\sqrt{\alpha}L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{t}\log(\tfrac{n}{\delta})\right\}B_0\|v_{B_0}\|,$$

which holds with probability at least $1 - \delta$. Given the burn-in period satisfying equation (53a) and stepsize satisfying equation (56), by combining with the bound on $\|v_{B_0}\|$ from Lemma 4, we find that $\|\mathbf{h}(\theta_t) - \theta_t\| \leq \widetilde{r}_\theta(t)$ with at least probability $1 - \delta$, uniformly for any integer $t \in [n]$, where

$$\widetilde{r}_\theta(t) := \frac{c_1}{\sqrt{t}}\left\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\right\} + \frac{c_2 b_*}{1-\gamma}\left\{\frac{1}{t} + \frac{\alpha L}{\sqrt{t}} \cdot \mathcal{J}_2(\mathbb{B}^*, \rho_n)\log(\tfrac{n}{\delta})\right\}\left\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta})\right\}$$

$$+ c_2\left\{\frac{\alpha B_0 L}{\sqrt{t}} \cdot \mathcal{J}_2(\mathbb{B}^*, \rho_n)\log(\tfrac{n}{\delta}) + \frac{B_0}{t}\right\}\|\mathbf{h}(\theta_0) - \theta_0\|. \quad (59)$$

By substituting our upper bound in terms of $\widetilde{r}_\theta$ into equation (52b), we obtain a recursive inequality that takes an admissible sequence r_v and generates a sequence r_v^+ such that

$$\mathbb{P}\left[\mathcal{E}_n^{(v)}(r_v^+)\right] \geq \mathbb{P}\left[\mathcal{E}_n^{(v)}(r_v)\right] - \delta.$$

For any positive integer N_1, we can apply the recursive inequality for N_1 times with $\delta' = \delta/N_1$; doing so yields a sharper bound for $\|v_t\|$. In particular, with probability at least $1 - \delta$, we

---

[4]We redefine $(\psi_v^{(N)}, \psi_\theta^{(N)})$ using the right-hand side of (57)

have

$$\|v_t\| \leq \frac{2}{1-\gamma}\left\{\frac{c}{t\sqrt{\alpha}}\left[\overline{\mathcal{W}} + \nu\sqrt{\log(\frac{nN_1}{\delta})}\right] + \frac{cb_*}{t}\left[\log(\frac{nN_1}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\right]\right\}$$
$$+ \frac{8}{(1-\gamma)\alpha t}\widetilde{r}_\theta(t) + \left(\frac{B_0}{t}\right)^2\|v_{B_0}\| + \left(\frac{1+\gamma}{2}\right)^{N_1} \cdot \frac{\psi_v^{(N)}}{t\sqrt{\alpha}}.$$

We take $N_1 := \lceil\frac{10\log n}{1-\gamma}\rceil$, and a stepsize and burn-in period satisfying the conditions (53a) and (56). With these choices, some algebra yields $\|v_t\| \leq \widetilde{r}_v(t)$ holds with probability at least $1-\delta$, uniformly for each integer $t \in [B_0, n]$, where

$$\widetilde{r}_v(t) := \frac{c'}{1-\gamma}\left\{\frac{1}{t\sqrt{\alpha}}\left[\overline{\mathcal{W}} + \nu\sqrt{\log(\frac{n}{\delta})}\right] + \frac{b_*}{t}\left[\log(\frac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\right]\right\} + 2\left(\frac{B_0}{t}\right)^2\|\theta_0 - \mathbf{h}(\theta_0)\|. \quad (60)$$

It can be seen that the sequences $\widetilde{r}_v$ and $\widetilde{r}_\theta$ are 2-admissible. Substituting their definitions into the bound (51) from Lemma 2 we find that the inequality

$$\|z_t\| \leq \frac{c}{\sqrt{t}}\left\{\overline{\mathcal{W}} + \nu\sqrt{\log(\frac{1}{\delta})}\right\} + \frac{cb_*}{t}\left\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\frac{1}{\delta})\right\}$$
$$+ \frac{cL}{t}\left\{\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\frac{1}{\delta})}\right\}\left\{\alpha\left(\sum_{s=B_0}^{t-1} s^2 r_v^2(s)\right)^{1/2} + \frac{1}{1-\gamma} \cdot \left(\sum_{s=1}^{t-1} r_\theta^2(s)\right)^{1/2}\right\}$$

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

Under the stepsize and burn-in period conditions (53a) and (56), some algebra yields:

$$\|z_t\| \leq \frac{c}{\sqrt{t}}\left(\overline{\mathcal{W}} + \nu\sqrt{\log(\frac{1}{\delta})}\right)$$
$$+ cb_*\left\{\frac{1}{t} + \frac{\alpha L}{\sqrt{t}} \cdot \mathcal{J}_2(\mathbb{B}^*, \rho_n)\log(\frac{n}{\delta})\right\}\left\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\frac{1}{\delta})\right\} + c\frac{B_0}{t}\|\theta_0 - \mathbf{h}(\theta_0)\|.$$

Combining with equation (60) yields the upper bound

$$\|\mathbf{h}(\theta_t) - \theta_t\| \leq \frac{c}{\sqrt{t}}\left(\overline{\mathcal{W}} + \nu\sqrt{\log(\frac{1}{\delta})}\right) + cb_*\left\{\frac{1}{t} + \frac{\alpha L}{\sqrt{t}}\mathcal{J}_2(\mathbb{B}^*, \rho_n)\log(\frac{n}{\delta})\right\}\left\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\frac{1}{\delta})\right\}$$
$$+ c\frac{B_0}{t}\|\theta_0 - \mathbf{h}(\theta_0)\|,$$

which completes the proof of the Theorem 1.

Besides, we also note that by taking a union bound over time steps $t \in \{B_0, B_0+1, \ldots, n\}$, we have the lower bound $\mathbb{P}\left[\mathcal{E}_n^{(\theta)}(r_\theta^*)\right] \geq 1-\delta$, where

$$r_\theta^*(t) := \frac{c}{\sqrt{t}}\left\{\overline{\mathcal{W}} + \nu\sqrt{\log(\frac{n}{\delta})}\right\}$$
$$+ cb_*\left\{\frac{1}{t} + \frac{\alpha L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}}\log(\frac{n}{\delta})\right\}\left\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\frac{n}{\delta})\right\} + c\frac{B_0}{t}\|\theta_0 - \mathbf{h}(\theta_0)\|.$$

5.2 Proof of Corollary 1

The proof of this corollary is based on a modification of Lemma 2. We introduce the shorthand

$$\overline{\mathcal{W}} := \mathbb{E}[\|W\|], \qquad\qquad \overline{\mathcal{W}}_C := \mathbb{E}[\|W\|_C],$$
$$\nu := \sqrt{\sup_{u \in \Gamma}\mathbb{E}\left[\langle u, W\rangle^2\right]} \quad \text{and} \quad \nu_C := \sqrt{\sup_{u \in C}\mathbb{E}\left[\langle u, W\rangle^2\right]}.$$

We begin by stating a lemma—a generalization of Lemma 2—that bounds the supremum of an averaged process. In the proof of Corollary 1, we only use a special case of Lemma 5, but the generality is useful later.

Recall the events $\mathcal{E}_n^{(\theta)}(r_\theta)$ and $\mathcal{E}_n^{(v)}(r_v)$ defined in equation (48). Given a bounded symmetric convex set $\mathcal{S} \subseteq \mathbb{V}^*$, we define the dimension factor $\mathcal{D}_\mathcal{S} := \sup_{u \in \mathcal{S}} \|u\|_*$. Moreover, we assume that there exists a constant $\mu > 0$ such that

$$\|\theta - \theta^*\| \leq \tfrac{1}{\mu}\|\mathbf{h}(\theta) - \theta\| \text{ for any } \theta \in \mathbb{V}. \tag{61}$$

We point out that under assumption (A1), the last condition is satisfied for $\mu = 1 - \gamma$. The condition (61) also allows us to analysis behavior of operators which satisfies a multi-step contraction assumption (A1)$'$ (cf. the proof of Theorem 3).

Lemma 5. *Suppose that the Assumptions (A2) and (A3) are in force, the sequences r_θ and r_v are κ-admissible for some $\kappa \in (0,2]$, and condition (61) holds. Then conditioned on the event $\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)$, we have*

$$\sup_{u \in \mathcal{S}} \langle u,\, z_t \rangle \leq \tfrac{c}{\sqrt{t}}\Big\{ \mathbb{E}\big[\sup_{u \in \mathcal{S}} \langle u,\, W \rangle\big] + \Big(\sup_{u \in \mathcal{S}} \mathbb{E}\big[\langle u,\, W \rangle^2\big] \log(\tfrac{1}{\delta}) \Big)^{1/2} \Big\} + \tfrac{\mathcal{D}_\mathcal{S} b_*}{t}\Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \Big\}$$

$$+ \tfrac{c \mathcal{D}_\mathcal{S} L}{t}\Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \Big\} \Big\{ \alpha \Big(\sum_{s=B_0}^{t-1} s^2 r_v^2(s) \Big)^{1/2} + \tfrac{1}{\mu} \cdot \Big(\sum_{s=1}^{t-1} r_\theta^2(s) \Big)^{1/2} \Big\}, \tag{62}$$

with probability at least $1 - \delta$, uniformly for all integers $t \in [B_0, n]$.

See Section 5.3.2 for the proof of this lemma.

Taking this lemma as given, we now proceed with proof of Corollary 1. As mentioned before, under assumption (A1)$'$, condition (61) is satisfied with $\mu = 1 - \gamma$. Applying Lemma 5 with $\mathcal{S} = C$ implies that

$$\big\| z_t \big\|_C \leq \tfrac{c}{\sqrt{t}}\Big\{ \overline{W}_C + \nu_C \sqrt{\log(\tfrac{1}{\delta})} \Big\} + c\tfrac{\mathcal{D} b_*}{t}\Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \Big\}$$

$$+ \tfrac{c \mathcal{D} L}{t}\Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \Big\} \Big\{ \alpha \Big(\sum_{s=B_0}^{t-1} s^2 r_v^2(s) \Big)^{1/2} + \tfrac{1}{1-\gamma} \Big(\sum_{s=1}^{t-1} r_\theta^2(s) \Big)^{1/2} \Big\}. \tag{63}$$

Now all we have to do is substitute an appropriate value of the sequences r_v and r_θ. Note that the estimate sequences $\widetilde{r}_v$ and $\widetilde{r}_\theta$ from equations (60) and (59), respectively, are 2-admissible; moreover, they provide upper bounds on the quantities $\|v_t\|$ and $\|z_t\|$ respectively. Next, using the stepsize and burn-in conditions (56) and (53a), we find that

$$\Big(\tfrac{1}{t} \sum_{s=B_0}^{t-1} s^2 \widetilde{r}_v^2(s) \Big)^{1/2} \leq \tfrac{c}{(1-\gamma)\sqrt{\alpha}}\Big\{ \overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})} \Big\}$$

$$+ \tfrac{c \cdot b_*}{(1-\gamma)}\Big\{ \log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n) \Big\} + \tfrac{2c\, B_0^{3/2}\, \|\theta_0 - \mathbf{h}(\theta_0)\|}{t},$$

and

$$\Big(\sum_{s=1}^{t-1} \widetilde{r}_\theta^2(s) \Big)^{1/2} \leq c \cdot \Big\{ \overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})} \Big\} \cdot \sqrt{\log t}$$

$$+ \tfrac{c b_*}{1-\gamma}\Big\{ \tfrac{1}{\sqrt{B_0}} + \alpha L \sqrt{\log t}\big[\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \big] \Big\} \Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta}) \Big\}$$

$$+ c\Big\{ \alpha B_0 L \sqrt{\log(t)}\big[\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \big] + \sqrt{B_0} \Big\} \cdot \|\mathbf{h}(\theta_0) - \theta_0\|;$$

29

both with probability at least $1-\delta$. Finally, substituting the last two bounds to the bound (63), and applying the conditions on stepsize (56) and burn-in period (53a), and using the fact $\left\|\mathbf{h}(\theta_t) - \theta_t\right\|_C \leq \left\|z_t\right\|_C + \mathcal{D} \cdot \|v_t\|$ we have

$$
\begin{aligned}
\left\|\mathbf{h}(\theta_t) - \theta_t\right\|_C &\leq \tfrac{c}{\sqrt{t}}\Big\{\overline{\mathcal{W}}_C + \nu_C\sqrt{\log(\tfrac{1}{\delta})}\Big\} + c\tfrac{\mathcal{D}L}{(1-\gamma)}\Big\{\mathcal{J}_2(\mathbb{B}^*,\rho_n)\log(\tfrac{n}{\delta})\sqrt{\tfrac{\alpha}{t}} + \tfrac{1}{t\sqrt{\alpha}}\Big\}\Big\{\overline{\mathcal{W}} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\} \\
&\quad + \tfrac{c\mathcal{D}Lb_*}{1-\gamma}\Big\{\tfrac{\sqrt{\alpha}}{t} + \tfrac{\alpha}{\sqrt{t}}\Big\}\mathcal{J}_2(\Gamma,\rho_n)\mathcal{J}_1(\Gamma,\rho_n)\log^2(\tfrac{n}{\delta}) + \tfrac{\mathcal{D}B_0}{t}\cdot\|\mathbf{h}(\theta_0) - \theta_0\|.
\end{aligned}
$$

This completes the proof of Corollary 1.

5.3 Proofs of key Lemmas for Theorem 1

In this section, we provide a detailed proofs of Lemmas 1 and 5 , which play a central role in the proofs of Theorem 1 and Corollary 1.

5.3.1 Proof of Lemma 1

We recursively expand the update rule for v_t from Algorithm 1, and obtain the identity:

$$
\begin{aligned}
tv_t &= (t-1)\left(v_{t-1} - \theta_{t-1} + \theta_{t-2} - \mathbf{H}_t(\theta_{t-2}) + \mathbf{H}_t(\theta_{t-1})\right) + \mathbf{H}_t(\theta_{t-1}) - \theta_{t-1} \\
&= (1-\alpha)(t-1)v_{t-1} + (t-1)\big(\mathbf{H}_t(\theta_{t-1}) - \mathbf{H}_t(\theta_{t-2})\big) + \big(\mathbf{H}_t(\theta_{t-1}) - \theta_{t-1}\big) \\
&= (1-\alpha)^\tau(t-\tau)v_{t-\tau} \\
&\quad + \sum_{j=1}^{\tau}(1-\alpha)^{j-1}\Big[(t-j)\big(\mathbf{H}_{t-j+1}(\theta_{t-j}) - \mathbf{H}_{t-j+1}(\theta_{t-j-1})\big) + \mathbf{H}_{t-j+1}(\theta_{t-j}) - \theta_{t-j}\Big],
\end{aligned}
$$

where the positive integer τ will be chosen later.

Consequently, we have the bound

$$
\begin{aligned}
t\|v_t\| &\leq (1-\alpha)^\tau(t-\tau)\|v_{t-\tau}\| + \sum_{j=1}^{\tau}(1-\alpha)^{j-1}(t-j)\|\mathbf{h}(\theta_{t-j}) - \mathbf{h}(\theta_{t-j-1})\| \\
&\quad + \Big\|\sum_{j=1}^{\tau}(1-\alpha)^{j-1}\Big((t-j)(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j})) + \mathbf{H}_{t-j+1}(\theta_{t-j}) - \theta_{t-j}\Big)\Big\| \\
&\leq (1-\alpha)^\tau(t-\tau)\|v_{t-\tau}\| + \sum_{j=1}^{\tau}(1-\alpha)^{j-1}\Big((t-j)\gamma\alpha\|v_{t-j}\| + \|\mathbf{h}(\theta_{t-j}) - \theta_{t-j}\|\Big) \\
&\quad + \Big\|\sum_{j=1}^{\tau}(1-\alpha)^{j-1}\Big((t-j)(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j-1})) + \varepsilon_{t-j+1}(\theta_{t-j})\Big)\Big\|.
\end{aligned}
$$

The estimate sequence is r_v κ-admissible for some $\kappa \in [0,2]$, so that the map $t \mapsto t^2 \cdot r_v(t)$ is non-decreasing. Thus, on the event $\mathcal{E}_n^{(v)}(r_v)$ for a burn-in $B_0 \geq \tfrac{12\tau}{1-\gamma}$, we have the upper bound

$$
(t-j)\|v_{t-j}\| \leq (t-j)r_v(t-j) \leq \tfrac{t^2}{t-j}r_v(t) \leq \tfrac{1}{1-\tau/B_0}\cdot tr_v(t) \leq \Big\{1 + \tfrac{1-\gamma}{6}\Big\}\cdot tr_v(t),
$$

valid for each integer $j \in [\tau]$.

Therefore, on the event $\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)$, we have the bound

$$t\|v_t\| \leq \left(1 + \tfrac{1-\gamma}{6}\right) \cdot \left\{(1-\alpha)^\tau + \gamma\alpha \sum_{j=1}^{\tau}(1-\alpha)^{j-1}\right\} \cdot tr_v(t) + \sum_{j=1}^{\tau} r_\theta(t-j) + T_1 + T_2, \quad (64)$$

where $T_2 := \left\|\sum_{j=1}^{\tau}(1-\alpha)^{j-1}\varepsilon_{t-j+1}(\theta_{t-j})\right\|$, and

$$T_1 := \|\sum_{j=1}^{\tau}(1-\alpha)^{j-1}(t-j)\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j-1})\big)\|.$$

We simplify the first two terms on the right-hand side of bound (64) by appropriately choosing the triple (τ, α, B_0). The later two terms T_1 and T_2 are norms of zero-mean random vectors in Banach spaces. First, we provide upper bound on these two noise terms.

Upper bound on T_1: First, we observe that the sum consists of the $(1-\alpha)$-weighted differences $(1-\alpha)^{j-1}(t-j)\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j-1})\big)$ that form a martingale difference sequence with respect to the natural filtration $(\mathcal{F}_t)_{t\geq 0}$. On the event $\mathcal{E}_n^{(v)}(r_v)$, we have that

$$\|(1-\alpha)^{j-1}(t-j)\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j-1})\big)\| \leq (t-j)\alpha L r_v(t-j)$$
$$\leq \tfrac{t^2}{t-j}\alpha L r_v(t) \leq 2t\alpha L r_v(t), \quad \text{a.s.}$$

The last inequality is due to the non-decreasing property of the function $t \mapsto t^2 r_v(t)$ and the fact that $t \geq B_0 > 2\tau$.

Since Ω is symmetric and convex by assumption, the difference $\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j-1})$ belongs to the set 2Ω. Conditioning on the event $\mathcal{E}_n^{(v)}(r_v)$ and invoking Lemma 8 yields

$$\left\|\tfrac{1}{\tau}\sum_{j=1}^{\tau}(1-\alpha)^{j-1}(t-j)\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j-1})\big)\right\| \leq \tfrac{ct\alpha L r_v(t)}{\sqrt{\tau}}\left\{\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\right\},$$

with probability at least $1 - \delta$.

Upper bound on T_2: In order to bound the last term in the decomposition (64), we decompose it into two parts:

$$\sum_{j=1}^{\tau}(1-\alpha)^{j-1}\varepsilon_{t-j+1}(\theta_{t-j}) = \sum_{j=1}^{\tau}(1-\alpha)^{j-1}\varepsilon_{t-j+1}(\theta^*) + \sum_{j=1}^{\tau}(1-\alpha)^{j-1}\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta^*)\big).$$

The former term is sum of independent random variables, while the latter is a martingale. Note that by Assumption (A1), we have $\|\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta^*)\| \leq \tfrac{L r_v(t-j+1)}{1-\gamma}$ on the event $\mathcal{E}_n^{(\theta)}(r_\theta)$. Invoking Lemma 7 yields

$$\|\tfrac{1}{\tau}\sum_{j=1}^{\tau}(1-\alpha)^{j-1}\varepsilon_{t-j+1}(\theta^*)\| \leq \tfrac{c}{\sqrt{\tau}}\left\{\overline{\mathcal{W}} + \nu\sqrt{\log(1/\delta)}\right\} + \tfrac{cb_*}{\tau}\left\{\log(1/\delta) + \mathcal{J}_1(\Gamma, \rho_n)\right\},$$

with probability at least $1 - \delta$.

Using the Lipschitz assumption (A2) and the contraction assumption (A1), we have

$$\|\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta^*)\| \leq L\|\theta_{t-j} - \theta^*\| \leq \tfrac{L}{1-\gamma}\|\mathbf{h}(\theta_{t-j}) - \theta_{t-j}\|.$$

31

Furthermore, since Ω is symmetric and convex, we have that $\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta^*) \in 2\Omega$. Conditioning on the event $\mathcal{E}_n^{(\theta)}(r_\theta)$ and invoking Lemma 8 yields

$$\left\|\tfrac{1}{\tau}\sum_{j=1}^{\tau}(1-\alpha)^{j-1}\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta^*)\big)\right\| \leq \tfrac{c}{\sqrt{\tau}} \cdot \tfrac{Lr_\theta(t-\tau+1)}{1-\gamma}\Big(\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\Big),$$

with probability at least $1 - \delta$.

Combining the pieces: Substituting the above concentration bounds into the decomposition (64) yields the upper bound

$$t \cdot \|v_t\| \leq \left\{(1-\alpha)^\tau + \gamma\alpha\sum_{j=1}^{\tau}(1-\alpha)^j\right\} \cdot \left\{1 + \tfrac{1-\gamma}{6}\right\} \cdot tr_v(t) + \tau \cdot r_\theta(t-\tau+1)$$
$$+ c\sqrt{\tau} \cdot \tfrac{Lr_\theta(t-\tau+1)}{1-\gamma}\left\{\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\right\} + c\sqrt{\tau}\left\{\overline{\mathcal{W}} + \nu\sqrt{\log(\tfrac{1}{\delta})}\right\}$$
$$+ cb_*\left\{\log(\tfrac{1}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\right\} + c\sqrt{\tau} \cdot t\alpha Lr_v(t) \cdot \left\{\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\right\}.$$

Re-arranging the terms in the last bound yields

$$t \cdot \|v_t\| \leq \left\{\gamma + (1-\gamma)\big((1-\alpha)^\tau + \tfrac{1}{3}\big) + cL\alpha\sqrt{\tau}\big(\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(1/\delta)}\big)\right\}tr_v(t)$$
$$+ \left\{\tau + c\tfrac{L\sqrt{\tau}}{1-\gamma}\big(\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(1/\delta)}\big)\right\}r_\theta(t-\tau+1)$$
$$+ c\sqrt{\tau}\Big(\overline{\mathcal{W}} + \nu\sqrt{\log(\tfrac{1}{\delta})}\Big) + cb_*\left\{\log(\tfrac{1}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\right\}.$$

Case I: $t \geq B_0 + \lceil 2\alpha^{-1}\rceil$

Taking $\tau = \lceil 2\alpha^{-1}\rceil \leq t - B_0$ and given a stepsize α satisfying the bound

$$6cL\sqrt{\alpha} \cdot \big(\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{n}{\delta})}\big) < 1 - \gamma, \tag{65}$$

we have the upper bounds

$$\tfrac{cL}{1-\gamma}\sqrt{\tau}\big(\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\big) \leq \tfrac{4}{\alpha}, \quad \text{and}$$
$$\gamma + (1-\gamma) \cdot \left\{(1-\alpha)^\tau + \tfrac{1}{3} + cL\alpha\sqrt{\tau}\big(\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\big)\right\} \leq \tfrac{1+\gamma}{2}.$$

Furthermore, since the function $t \mapsto t^2 \cdot r_\theta(t)$ is non-decreasing, for burn-in period $B_0 \geq 4\tau$, we have

$$r_\theta(t-\tau+1) \leq \tfrac{t^2}{(t-\tau+1)^2}r_\theta(t) \leq \tfrac{16}{9}r_\theta(t) \qquad \text{for all} \qquad t \geq B_0.$$

Substituting the bounds yields

$$t\|v_t\| \leq \tfrac{1+\gamma}{2}t \cdot r_v(t) + \tfrac{8}{\alpha}r_\theta(t) + \tfrac{c}{\sqrt{\alpha}}\Big(\overline{\mathcal{W}} + \nu\sqrt{\log(\tfrac{1}{\delta})}\Big) + c\left(\log(\tfrac{1}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\right), \tag{66}$$

which completes the proof of Lemma 1 in the case of $t \geq B_0 + 2/\alpha$.

Case II: $B_0 \leq t \leq B_0 + \lceil 2\alpha^{-1} \rceil$:

This case requires a special treatment, since the number τ of recursive expansion steps cannot be taken as large as $\lceil 2/\alpha \rceil$. Instead, we choose $\tau = t - B_0$, and expand the recursions backwards up to the beginning of the iterates. In this case, following the same arguments as above, on the event $\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)$, the error decomposition (64) takes the form

$$t \cdot \|v_t\| \leq (1-\alpha)^\tau B_0 \|v_{B_0}\| + \left(1 + \tfrac{1-\gamma}{6}\right) \cdot \left\{ \gamma\alpha \sum_{j=1}^\tau (1-\alpha)^{j-1} \right\} \cdot tr_v(t) + \sum_{j=1}^\tau r_\theta(t-j) + T_1 + T_2. \tag{67}$$

Substituting the upper bounds on the terms T_1 and T_2 yields

$$t \cdot \|v_t\| \leq (1-\alpha)^\tau B_0 \|v_{B_0}\| + \gamma\alpha \sum_{j=1}^\tau (1-\alpha)^j \cdot \left\{ 1 + \tfrac{1-\gamma}{6} \right\} \cdot tr_v(t) + \tau \cdot r_\theta(t - \tau + 1)$$

$$+ c\sqrt{\tau} \cdot \tfrac{Lr_\theta(t-\tau+1)}{1-\gamma} \left\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \right\} + c\sqrt{\tau} \left\{ \overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})} \right\}$$

$$+ cb_* \left\{ \log(\tfrac{1}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n) \right\} + c\sqrt{\tau} \cdot t\alpha L r_v(t) \cdot \left\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \right\}.$$

For a time index $t \in [B_0, B_0 + 2/\alpha]$, we have the decomposition

$$(1-\alpha)^\tau B_0 \cdot \|v_{B_0}\| \leq \left((1-\alpha)^\tau - 3(1-\gamma)\right) \cdot q \left\{ 1 + \frac{2}{\alpha B_0} \right\} t \cdot r_v(t) + 3(1-\gamma) B_0 \cdot \|v_{B_0}\|$$

$$\leq \left\{ (1-\alpha)^\tau - 2(1-\gamma) \right\} \cdot tr_v(t) + 6(1-\gamma) \frac{{B_0}^2}{t} \cdot \|v_{B_0}\|$$

Given a stepsize α satisfying the requirement (49), choosing the number of steps such that $\tau = t - B_0 \leq 2/\alpha$ leads to the inequalities

$$\left\{ (1-\alpha)^\tau - 2(1-\gamma) \right\} + \gamma\alpha \sum_{j=1}^\tau (1-\alpha)^j \cdot \left\{ 1 + \tfrac{1-\gamma}{6} \right\} + c\sqrt{\tau} \cdot \alpha L \cdot \left\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \right\} \leq \frac{1+\gamma}{2}, \quad \text{and}$$

$$\frac{cL}{1-\gamma} \sqrt{\tau} \left(\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(1/\delta)} \right) \leq \frac{4}{\alpha} + \tau \leq \frac{8}{\alpha}.$$

Putting together these bounds completes the proof in the second case.

5.3.2 Proof of Lemma 5

Expanding the update rule for z_t from Algorithm 1 we obtain the three-term decomposition $t \cdot z_t = B_0 \cdot z_{B_0} + M_t + \Psi_t$, where

$$M_t := \sum_{s=B_0}^{t-1} \varepsilon_s(\theta_{s-1}), \quad \text{and} \quad \Psi_t := \sum_{s=B_0}^{t-1} (s-1) \left\{ \varepsilon_s(\theta_{s-1}) - \varepsilon_s(\theta_{s-2}) \right\}.$$

It suffices to control each of these three terms in the semi-norm induced by the set $\mathcal{S}$.

Beginning with the martingale $\{M_t\}_{t \geq B_0}$, we further break it down into two parts:

$$M_t = \sum_{s=B_0}^{t-1} \varepsilon_s(\theta^*) + \sum_{s=B_0}^{t-1} \left(\varepsilon_s(\theta_{s-1}) - \varepsilon_s(\theta^*) \right) := M_t^* + \widetilde{M}_t.$$

The term M_t^* is sum of i.i.d. random variables. Invoking Lemma 7 and using the fact that the set $\mathcal{S}$ is contained within $\mathcal{D}_\mathcal{S}\Gamma$, we have the bound

$$\sup_{u \in \mathcal{S}} \langle u, M^*(t) \rangle \leq c\sqrt{t} \Big\{ \mathbb{E}\big[\sup_{u \in \mathcal{S}} \langle u, W \rangle \big] + \Big(\sup_{u \in \upsilon} \mathbb{E}\big[\langle u, W \rangle^2 \big] \cdot \log(\tfrac{1}{\delta}) \Big)^{1/2} \Big\}$$
$$+ c\mathcal{D}_\mathcal{S}b_* \Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \Big\}, \quad (68)$$

where W is the centered Gaussian process with covariance matching that of $\varepsilon(\theta^*)$.

Next we bound the terms $\langle u, \widetilde{M}(t) \rangle$ and $\langle u, \Psi(t) \rangle$. First, we claim that conditioned on the event $\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)$, we have

$$\|\widetilde{M}(t)\| \leq c\frac{L}{\mu} \Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{n}{\delta})} \Big\} \cdot \Big(\sum_{k=B_0}^{t-1} r_\theta^2(k) \Big)^{1/2}, \quad \text{and} \quad (69a)$$

$$\|\Psi(t)\| \leq c\alpha L \Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{n}{\delta})} \Big\} \cdot \Big(\sum_{s=B_0}^{t} s^2 r_v^2(s) \Big)^{1/2}, \quad (69b)$$

both bounds holding with probability at least $1 - \delta$.

The proof of these two inequalities can be found at the end of this subsection. Since the set $\mathcal{S}$ is contained within $\mathcal{D}_\mathcal{S}\Gamma$, it follows that

$$\sup_{u \in \mathcal{S}} \langle u, \widetilde{M}(t) \rangle \leq \mathcal{D}_\mathcal{S} \|\widetilde{M}(t)\|, \quad \text{and} \quad \sup_{u \in \mathcal{S}} \langle u, \Psi(t) \rangle \leq \mathcal{D}_\mathcal{S} \|\Psi(t)\|.$$

Finally, observe that $B_0 z_{B_0} = \sum_{t=1}^{B_0} \varepsilon_t(\theta^*) + \sum_{t=1}^{B_0} \{\varepsilon_t(\theta_0) - \varepsilon_t(\theta^*)\}$. By Lemma 7 we have

$$\sup_{u \in \mathcal{S}} \langle u, \sum_{t=1}^{B_0} \varepsilon_t(\theta^*) \rangle \leq c\sqrt{B_0} \Big\{ \mathbb{E}\big[\sup_{u \in \mathcal{S}} \langle u, W \rangle \big] + \Big(\sup_{u \in \upsilon} \mathbb{E}\big[\langle u, W \rangle^2 \big] \cdot \log(\tfrac{1}{\delta}) \Big)^{1/2} \Big\}$$
$$+ c\mathcal{D}_\mathcal{S}b_* \Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \Big\},$$

with probability at least $1 - \delta$. On the other hand, using Lemma 8, we have

$$\sup_{u \in \mathcal{S}} \langle u, \sum_{t=1}^{B_0} \big(\varepsilon_t(\theta_0) - \varepsilon_t(\theta^*) \big) \rangle \leq \mathcal{D}_\mathcal{S} \| \sum_{t=1}^{B_0} \{ \varepsilon_t(\theta_0) - \varepsilon_t(\theta^*) \} \|$$
$$\leq cL\mathcal{D}_\mathcal{S} \|\theta_0 - \theta^*\| \cdot \sqrt{B_0} \Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \Big\}$$
$$\leq c\mathcal{D}_\mathcal{S} \cdot \frac{L}{\mu} \sqrt{B_0} r_\theta(B_0) \Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \Big\},$$

with probability at least $1 - \delta$. Combining the two bounds, we conclude that

$$\sup_{u \in \mathcal{S}} \langle u, z_{B_0} \rangle \leq \frac{c}{\sqrt{B_0}} \Big\{ \mathbb{E}\big[\sup_{u \in \mathcal{S}} \langle u, W \rangle \big] + \Big(\sup_{u \in \upsilon} \mathbb{E}\big[\langle u, W \rangle^2 \big] \cdot \log(\tfrac{1}{\delta}) \Big)^{1/2} \Big\}$$
$$+ \frac{c\mathcal{D}_\mathcal{S}b_*}{B_0} \big(\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \big) + \frac{c\mathcal{D}_\mathcal{S} \cdot Lr_\theta(B_0)}{\mu\sqrt{B_0}} \big(\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \big(\log(\tfrac{1}{\delta}) \big)^{1/2} \big), \quad (70)$$

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

We now put together the bounds (68), (69a) (69b), and (70). By doing so, we are guaranteed that conditioned on the event $\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)$, for each integer $t \in [B_0, n]$, we have

$$
\sup_{u \in \mathcal{S}} \langle u, z_t \rangle \leq \frac{c}{\sqrt{t}} \Big\{ \mathbb{E}\big[\sup_{u \in \mathcal{S}} \langle u, W \rangle \big] + \Big(\sup_{u \in \mathcal{S}} \mathbb{E}\big[\langle u, W \rangle^2 \big] \log(\tfrac{1}{\delta}) \Big)^{1/2} \Big\} + \frac{\mathcal{D}_\mathcal{S} b_*}{t} \Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \Big\}
$$
$$
+ c \frac{\mathcal{D}_\mathcal{S} \cdot \alpha L}{\sqrt{t}} \Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \Big\} \Big(\frac{1}{t} \sum_{s=B_0}^{t-1} s^2 r_v^2(s) \Big)^{1/2}
$$
$$
+ 2c \frac{\mathcal{D}_\mathcal{S} \cdot L}{\mu t} \Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \Big\} \cdot \Big(\sum_{s=B_0}^{t-1} r_\theta^2(s) + B_0 r_\theta^2(B_0) \Big)^{1/2}.
$$

The claim of Lemma 5 now follows by noting $\Big(\sum_{s=B_0}^{t-1} r_\theta^2(s) + B_0 r_\theta^2(B_0) \Big)^{1/2} = \Big(\sum_{s=1 B_0}^{t-1} r_\theta^2(s) \Big)^{1/2}$.
It remains to prove inequalities (69a) and (69b).

Proof of the bound (69a): Conditioned on the event $\mathcal{E}_n^{(\theta)}(r_\theta)$, we have the upper bounds

$$
\|\varepsilon_s(\theta_{s-1}) - \varepsilon_s(\theta^*)\| = \|\mathbf{H}_s(\theta_{s-1}) - \mathbf{H}_s(\theta^*) - \mathbf{h}(\theta_{s-1}) + \mathbf{h}(\theta^*)\| \leq L\|\theta_{s-1} - \theta^*\|
$$
$$
\leq \frac{L}{\mu} r_\theta(s-1),
$$

where the last inequality follows from the assumption $\|\theta_{s-1} - \theta^*\| \leq \frac{1}{\mu}\|\mathbf{h}(\theta_{s-1}) - \mathbf{h}(\theta^*)\|$ (cf. assumption (61)).

On the event $\mathcal{E}_n^{(\theta)}(r_\theta)$, we apply Lemma 8 to the martingale differences $\{\varepsilon_s(\theta_{s-1}) - \varepsilon_s(\theta^*)\}_{s=B_0+1}^t$, and find that

$$
\Big\| \sum_{s=B_0+1}^t \big(\varepsilon_s(\theta_{s-1}) - \varepsilon_s(\theta^*) \big) \Big\| \leq c \Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \Big\} \frac{L}{\mu} \Big(\sum_{s=B_0}^{t-1} r_\theta^2(s) \Big)^{1/2},
$$

with probability at least $1 - \delta$, as claimed in inequality (69a).

Proof of bound (69b): We now control the martingale sequence $\{\Psi_t\}_{t \geq B_0}$. Conditioned on the event $\mathcal{E}_n^{(v)}(r_v)$, we have

$$
\|(s-1)\{\varepsilon_s(\theta_{s-1}) - \varepsilon_s(\theta_{s-2})\}\| \leq (s-1)L \cdot \|\theta_{s-1} - \theta_{s-2}\| = (s-1)\alpha L \|v_{s-1}\|
$$
$$
\leq (s-1)\alpha L r_v(s-1),
$$

valid for any integer $s \in [B_0, t]$. By Lemma 8, on the event $\mathcal{E}_n^{(v)}(r_v)$, we have

$$
\Big\| \sum_{s=B_0+1}^t (s-1)\{\varepsilon_s(\theta_{s-1}) - \varepsilon_s(\theta_{s-2})\} \Big\| \leq c\alpha \Big\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \Big\} L \Big(\sum_{s=B_0}^{t-1} s^2 r_v^2(s) \Big)^{1/2},
$$

with probability at least $1 - \delta$, which establishes the claim. This completes the proof of Lemma 2.

6 Discussion

In this paper, we have analyzed `ROOT-SA`, a variance-reduced stochastic approximation procedure designed for solving contractive fixed-point equations in Banach spaces. This procedure builds upon the `ROOT-SGD` algorithm [Li+20] for stochastic optimization, as studied in past work by a subset of the current authors. Our main contribution was to derive non-asymptotic upper bounds on the error of the `ROOT-SA` iterates in any semi-norm. We showed that these bounds are sharp in the sense that the leading order term matches the optimal risk characterized by local minimax theory. Furthermore, the sample complexity needed for such an instance-dependent optimal statistical behavior scale with the intrinsic complexity of the norm (measured in Dudley integral of the dual ball under certain metric), instead of the problem dimension. Our main results, while formulated for general Banach spaces and contractions, have interesting consequences for specific classes of problems. Here we illustrated with applications to dynamical programming and game theory, including stochastic shortest path problems, minimax Markov games, as well as average-cost policy evaluation. In terms of proof techniques, our analysis is rather different than much other theory on stochastic approximation that relies the inner product available in the Hilbert setting. To the best of our knowledge, our paper is the first to provide sharp non-asymptotic bounds for stochastic approximation *without* requiring such inner product structure.

Our work leaves open a number of open questions, among them:

- **Optimal sample complexity for SA schemes:** One open question in our analysis concerns how the minimal sample size scales with the contraction factor $\gamma \in (0,1)$, Our main results in this paper (Theorem 1 and 2) have a scaling condition of the form $n \gtrsim (1-\gamma)^{-4}$. These results are novel even under this quadratic scaling, and also optimal for large n. However, it is not yet clear whether this quadratic scaling is necessary, or rather an artifact of our proof technique. In certain special cases, the quadratic scaling can be avoided; for example, in the special case of $\mathbf{h}$ being an affine operator, our results (see Theorem 3 and Corollary 3) show that the quadratic scaling $\mathcal{O}\big((1-\gamma)^{-2}\big)$ is sufficient. Furthermore, in the classical Euclidean setting and in the special case gradient-update operator $\mathbf{h} : \theta \mapsto \theta - \beta^{-1}\nabla f(\theta)$ for a μ-strongly-convex and β-smooth function f, the paper [Li+20] establishes instance-optimal bounds that require only $\mathcal{O}\big((1-\gamma)^{-2}\big)$ samples. An interesting open problem, therefore, is to determine the minimum sample size for which non-asymptotic bounds of the form stated in this paper hold in the general Banach space setting.

- **Online statistical inference procedures:** In this paper, we focused exclusively on computing point estimates of the fixed point. However, a natural question is the construction of confidence sets for the solution θ^* to the fixed-point equation. Ideally, such confidence set should be efficiently computable, asymptotically exact, while capturing the desirable non-asymptotic properties satisfied by our estimator. Focusing stochastic optimization in the Euclidean setting and Polyak-Ruppert-averaged SGD, the paper [Che+20a] proposed an online estimator for the covariance that partly achieves these goals. In a concurrent piece of work involving a subset of the authors [Xia+22], confidence sets and early stopping rules are developed in the special case of policy evaluation and optimization for discounted MDPs. It is an interesting direction for future research to construct confidence sets with improved guarantees in the general setting, based purely on the algorithm's trajectory alone.

- **General operator equations beyond the contractive setting:** In the Euclidean setting, stochastic approximation procedures for nonlinear equations share geometric struc-

ture, giving rise to key concepts such as monotonicity and smoothness. This story becomes more complex for Banach spaces, with there being at least two distinct methods of analysis depending on the set-up. On the one hand, if the operator $\mathbf{h}$ is mapping from the space $\mathbb{V}$ to itself, then convergence is governed by contraction properties of the operator. On the other hand, if $\mathbf{h}$ maps from the Banach space $\mathbb{V}$ to its dual space $\mathbb{V}^*$, then a monotonicity condition with respect to the Bregman divergence plays a key role (see e.g. [JNT11; KLL20]). This paper focuses on the former case, in which $\mathbf{h}$ maps the Banach space to itself, but it is an interesting direction of future research to provide instance-dependent guarantees for various stochastic approximation procedures in the latter case, and examine their optimality properties. Even more broadly, it is interesting to consider stochastic approximation procedures for solving general non-linear equations defined on pairs of Banach spaces.

Acknowledgements

References

[BCN18]   L. Bottou, F. E. Curtis, and J. Nocedal. "Optimization methods for large-scale machine learning". In: *SIAM Review* 60.2 (2018), pp. 223–311 (cit. on p. 2).

[Ben96]   M. Benaim. "A dynamical system approach to stochastic approximations". In: *SIAM Journal on Control and Optimization* 34.2 (1996), pp. 437–472 (cit. on p. 3).

[Ber12a]   D. P. Bertsekas. *Approximate Dynamic Programming.* Athena Scientific Belmont, 2012 (cit. on p. 61).

[Ber12b]   D. P. Bertsekas. "Weighted sup-norm contractions in dynamic programming: A review and some new applications". In: *Dept. Elect. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, USA, Tech. Rep. LIDS-P-2884* (2012) (cit. on p. 5).

[Ber19]   D. P. Bertsekas. *Reinforcement Learning and Optimal Control.* Athena Scientific Belmont, MA, 2019 (cit. on pp. 2, 5, 56).

[BM13]   F. Bach and E. Moulines. "Non-strongly-convex smooth stochastic approximation with convergence rate $O(1/n)$". In: *arXiv preprint arXiv:1306.2119* (2013) (cit. on p. 4).

[BMP12]   A. Benveniste, M. Métivier, and P. Priouret. *Adaptive Algorithms and Stochastic Approximations.* Vol. 22. Springer Science & Business Media, 2012 (cit. on pp. 3, 4, 7).

[Bor09]   V. S. Borkar. *Stochastic Approximation: A Dynamical Systems Viewpoint.* Vol. 48. Springer, 2009 (cit. on pp. 2–4, 7).

[Bor21]   V. S. Borkar. "A concentration bound for contractive stochastic approximation". In: *Systems & Control Letters* 153 (2021), p. 104947 (cit. on p. 4).

[BT91]   D. P. Bertsekas and J. N. Tsitsiklis. "An analysis of stochastic shortest path problems". In: *Mathematics of Operations Research* 16.3 (1991), pp. 580–595 (cit. on pp. 16, 17).

[Che+19]   Z. Chen, S. Zhang, T. T. Doan, S. T. Maguluri, and J.-P. Clarke. "Performance of Q-learning with linear function approximation: Stability and finite-time analysis". In: *arXiv preprint arXiv:1905.11425* (2019) (cit. on p. 4).

[Che+20a]   X. Chen, J. D. Lee, X. T. Tong, and Y. Zhang. "Statistical inference for model parameters in stochastic gradient descent". In: *The Annals of Statistics* 48.1 (2020), pp. 251–273 (cit. on pp. 4, 36).

[Che+20b]   Z. Chen, S. Theja Maguluri, S. Shakkottai, and K. Shanmugam. "Finite-sample analysis of stochastic approximation using smooth convex envelopes". In: *arXiv e-prints* (2020), arXiv–2002 (cit. on p. 4).

[Che+21a]   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". In: *arXiv preprint arXiv:2102.01567* (2021) (cit. on p. 4).

[Che+21b]   Z. Chen, S. T. Maguluri, S. Shakkottai, and K. Shanmugam. "Finite-Sample Analysis of Off-Policy TD-Learning via Generalized Bellman Operators". In: *Advances in Neural Information Processing Systems* 34 (2021) (cit. on p. 4).

[CO19]   A. Cutkosky and F. Orabona. "Momentum-based variance reduction in non-convex SGD". In: *arXiv preprint arXiv:1905.10018* (2019) (cit. on p. 5).

[DBLJ14]   A. Defazio, F. Bach, and S. Lacoste-Julien. "SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives". In: *Advances in Neural Information Processing Systems*. 2014, pp. 1646–1654 (cit. on p. 5).

[Der66]   C. Derman. "Denumerable state Markovian decision processes-average cost criterion". In: *The Annals of Mathematical Statistics* 37.6 (1966), pp. 1545–1553 (cit. on p. 20).

[DG03]   J. Dugundji and A. Granas. *Fixed Point Theory*. Springer Verlag, 2003 (cit. on pp. 1, 6).

[DR16]   J. Duchi and F. Ruan. "Asymptotic optimality in stochastic optimization". In: *arXiv preprint arXiv:1612.05612* (2016) (cit. on pp. 2, 4, 13).

[Fro+15]   R. Frostig, R. Ge, S. M. Kakade, and A. Sidford. "Competing with the empirical risk minimizer in a single pass". In: *Conference on Learning Theory*. PMLR. 2015, pp. 728–763 (cit. on p. 4).

[GJG18]   A. Gupta, R. Jain, and P. Glynn. "Probabilistic contraction analysis of iterated random operators". In: *arXiv preprint arXiv:1804.01195* (2018) (cit. on p. 4).

[GL12]   S. Ghadimi and G. Lan. "Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization I: A generic algorithmic framework". In: *SIAM Journal on Optimization* 22.4 (2012), pp. 1469–1492 (cit. on p. 2).

[GL13]   S. Ghadimi and G. Lan. "Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization, II: shrinking procedures and optimal algorithms". In: *SIAM Journal on Optimization* 23.4 (2013), pp. 2061–2089 (cit. on p. 2).

[GP17]   S. Gadat and F. Panloup. "Optimal non-asymptotic bound of the Ruppert-Polyak averaging without strong convexity". In: *arXiv preprint arXiv:1709.03342* (2017) (cit. on p. 4).

[Häj72]   J. Hájek. "Local asymptotic minimax and admissibility in estimation". In: *Theory of Statistics*. University of California Press, 1972, pp. 175–194 (cit. on pp. 3, 4, 13).

[JNT11]   A. Juditsky, A. Nemirovski, and C. Tauvel. "Solving variational inequalities with stochastic mirror-prox algorithm". In: *Stochastic Systems* 1.1 (2011), pp. 17–58 (cit. on pp. 2, 37).

[JS20]   Y. Jin and A. Sidford. "Efficiently solving MDPs with stochastic mirror descent". In: *International Conference on Machine Learning*. PMLR. 2020, pp. 4890–4900 (cit. on p. 22).

[JZ13]   R. Johnson and T. Zhang. "Accelerating stochastic gradient descent using predictive variance reduction". In: *Advances in Neural Information Processing Systems* 26 (2013), pp. 315–323 (cit. on pp. 2, 5, 6).

[KC78]   H. J. Kushner and D. S. Clark. *Stochastic approximation methods for constrained and unconstrained systems*. Vol. 26. Applied Mathematical Sciences. Springer-Verlag, New York-Berlin, 1978, pp. x+261 (cit. on p. 3).

[Kha+20]   K. Khamaru, A. Pananjady, F. Ruan, M. J. Wainwright, and M. I. Jordan. "Is temporal difference learning optimal? An instance-dependent analysis". In: *arXiv preprint arXiv:2003.07337* (2020) (cit. on pp. 2, 4–6, 22).

[Kha+21]   K. Khamaru, E. Xia, M. J. Wainwright, and M. I. Jordan. "Instance-optimality in optimal value estimation: Adaptivity via variance-reduced Q-learning". In: *arXiv preprint arXiv:2106.14352* (2021) (cit. on pp. 4–6, 19, 56).

[Kha66]   R. Khas'minskii. "On stochastic processes defined by differential equations with a small parameter". In: *Theory of Probability and Its Applications* 11.2 (1966), pp. 211–228 (cit. on p. 4).

[Kir11]   A. Kirsch. *An Introduction to the Mathematical Theory of Inverse Problems*. Vol. 120. Springer, 2011 (cit. on p. 2).

[KLL20]   G. Kotsalis, G. Lan, and T. Li. "Simple and optimal methods for stochastic variational inequalities, II: Markovian noise and policy evaluation in reinforcement learning". In: *arXiv preprint arXiv:2011.08434* (2020) (cit. on pp. 2, 37).

[KS84]   H. J. Kushner and A. Shwartz. "An invariant measure approach to the convergence of stochastic approximations with state dependent noise". In: *SIAM Journal on Control and Optimization* 22.1 (1984), pp. 13–27 (cit. on p. 4).

[Kus84]   H. J. Kushner. *Approximation and weak convergence methods for random processes, with applications to stochastic systems theory*. Vol. 6. MIT press, 1984 (cit. on p. 4).

[KW52]   J. Kiefer and J. Wolfowitz. "Stochastic estimation of the maximum of a regression function". In: *The Annals of Mathematical Statistics* (1952), pp. 462–466 (cit. on p. 3).

[KY03]   H. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Vol. 35. Springer Science & Business Media, 2003 (cit. on pp. 3, 4, 7).

[LeC53]   L. LeCam. "On some asymptotic properties of maximum likelihood estimates and related Bayes estimates". In: *Univ. California Pub. Statist.* 1 (1953), pp. 277–330 (cit. on p. 13).

[Li+20]   C. J. Li, W. Mou, M. J. Wainwright, and M. I. Jordan. "Root-SGD: Sharp nonasymptotics and asymptotic efficiency in a single algorithm". In: *arXiv preprint arXiv:2008.12690* (2020) (cit. on pp. 2–7, 36).

[Li+21]   G. Li, C. Cai, Y. Chen, Y. Gu, Y. Wei, and Y. Chi. "Is Q-learning minimax optimal? a tight sample complexity analysis". In: *arXiv preprint arXiv:2102.06548* (2021) (cit. on p. 4).

[Lju77a]   L. Ljung. "Analysis of recursive stochastic algorithms". In: *IEEE Transactions on Automatic Control* 22.4 (1977), pp. 551–575 (cit. on p. 3).

[Lju77b]   L. Ljung. "On positive real transfer functions and the convergence of some recursive schemes". In: *IEEE Transactions on Automatic Control* 22.4 (1977), pp. 539–551 (cit. on p. 3).

[LLP21]   T. Li, G. Lan, and A. Pananjady. "Accelerated and instance-optimal policy evaluation with linear function approximation". In: *arXiv preprint arXiv:2112.13109* (2021) (cit. on p. 4).

[LT13]   M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer Science & Business Media, 2013 (cit. on pp. 3, 9).

[MB11]   E. Moulines and F. R. Bach. "Non-asymptotic analysis of stochastic approximation algorithms for machine learning". In: *Advances in Neural Information Processing Systems.* 2011, pp. 451–459 (cit. on pp. 2, 4).

[Mou+20] W. Mou, C. J. Li, M. J. Wainwright, P. L. Bartlett, and M. I. Jordan. "On linear stochastic approximation: Fine-grained Polyak-Ruppert and non-asymptotic concentration". In: *Conference on Learning Theory.* PMLR. 2020, pp. 2947–2997 (cit. on p. 2).

[Mou+21] W. Mou, A. Pananjady, M. J. Wainwright, and P. L. Bartlett. "Optimal and instance-dependent guarantees for Markovian linear stochastic approximation". In: *arXiv preprint arXiv:2112.12770* (2021) (cit. on p. 4).

[Nem+09] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. "Robust stochastic approximation approach to stochastic programming". In: *SIAM Journal on Optimization* 19.4 (2009), pp. 1574–1609 (cit. on p. 2).

[Nes03]   Y. Nesterov. *Introductory lectures on convex optimization: A basic course.* Vol. 87. Springer Science & Business Media, 2003 (cit. on p. 2).

[Ngu+17] L. M. Nguyen, J. Liu, K. Scheinberg, and M. Takáč. "SARAH: A novel method for machine learning problems using stochastic recursive gradient". In: *International Conference on Machine Learning.* PMLR. 2017, pp. 2613–2621 (cit. on p. 5).

[NST21]   L. M. Nguyen, K. Scheinberg, and M. Takáč. "Inexact SARAH algorithm for stochastic optimization". In: *Optimization Methods and Software* 36.1 (2021), pp. 237–258 (cit. on pp. 2, 5, 6).

[NY83]   A. S. Nemirovski and D. B. Yudin. *Problem Complexity and Method Efficiency in Optimization.* Wiley-Interscience, 1983 (cit. on p. 2).

[Pat97]   S. D. Patek. "Stochastic and shortest path games: theory and algorithms". PhD thesis. Massachusetts Institute of Technology, 1997 (cit. on pp. 54, 55).

[Per+15] J. Perolat, B. Scherrer, B. Piot, and O. Pietquin. "Approximate dynamic programming for two-player zero-sum Markov games". In: *International Conference on Machine Learning.* PMLR. 2015, pp. 1321–1329 (cit. on pp. 54, 55).

[PJ92]   B. T. Polyak and A. B. Juditsky. "Acceleration of stochastic approximation by averaging". In: *SIAM Journal on Control and Optimization* 30.4 (1992), pp. 838–855 (cit. on pp. 2, 4).

[Pol90]   B. T. Polyak. "A new method of stochastic approximation type". In: *Automat. i Telemekh* 7.98-107 (1990), p. 2 (cit. on p. 4).

[Put05]   M. L. Puterman. *Markov decision processes: Discrete stochastic dynamic programming.* Wiley, 2005 (cit. on p. 2).

[QW20]   G. Qu and A. Wierman. "Finite-Time Analysis of Asynchronous Stochastic Approximation and Q-Learning". In: *Conference on Learning Theory.* PMLR. 2020, pp. 3185–3205 (cit. on p. 4).

[RM51]   H. Robbins and S. Monro. "A stochastic approximation method". In: *The Annals of Mathematical Statistics* (1951), pp. 400–407 (cit. on p. 3).

[Rup88]   D. Ruppert. *Efficient estimations from a slowly convergent Robbins-Monro process.* Tech. rep. Cornell University Operations Research and Industrial Engineering, 1988 (cit. on pp. 2, 4).

[RW09]   R. T. Rockafellar and R. J.-B. Wets. *Variational analysis.* Vol. 317. Springer Science & Business Media, 2009 (cit. on p. 2).

[Sid+18] A. Sidford, M. Wang, X. Wu, L. F. Yang, and Y. Ye. "Near-optimal time and sample complexities for solving discounted Markov decision process with a generative model". In: *arXiv preprint arXiv:1806.01492* (2018) (cit. on p. 5).

[SLRB17]   M. Schmidt, N. Le Roux, and F. Bach. "Minimizing finite sums with the stochastic average gradient". In: *Mathematical Programming* 162.1-2 (2017), pp. 83–112 (cit. on p. 5).

[Sto89]   N. L. Stokey. *Recursive Methods in Economic Dynamics*. Harvard University Press, 1989 (cit. on p. 2).

[Sze98]   C. Szepesvári. "The asymptotic convergence-rate of Q-learning". In: *Advances in Neural Information Processing Systems* (1998), pp. 1064–1070 (cit. on p. 60).

[Tal06]   M. Talagrand. *The Generic Chaining: Upper and Lower Bounds of Stochastic Processes*. Springer Science and Business Media, 2006 (cit. on p. 57).

[Tal96]   M. Talagrand. "New concentration inequalities in product spaces". In: *Inventiones Mathematicae* 126.3 (1996), pp. 505–563 (cit. on pp. 57, 58).

[Tes12]   G. Teschl. *Ordinary differential equations and dynamical systems*. Vol. 140. American Mathematical Soc., 2012 (cit. on p. 2).

[Tri+18]   N. Tripuraneni, N. Flammarion, F. Bach, and M. I. Jordan. "Averaging stochastic gradient descent on Riemannian manifolds". In: *Conference On Learning Theory*. PMLR. 2018, pp. 650–687 (cit. on p. 4).

[Tse90]   P. Tseng. "Solving H-horizon, stationary Markov decision problems in time proportional to $\log(H)$". In: *Operations Research Letters* 9.5 (1990), pp. 287–297 (cit. on p. 17).

[TVR99]   J. N. Tsitsiklis and B. Van Roy. "Average cost temporal-difference learning". In: *Automatica* 35.11 (1999), pp. 1799–1808 (cit. on p. 20).

[Vaa00]   A. W. van der Vaart. *Asymptotic Statistics*. Vol. 3. Cambridge University Press, 2000 (cit. on pp. 3, 4, 13).

[VH14]   R. Van Handel. *Probability in high dimension*. Tech. rep. PRINCETON UNIV NJ, 2014 (cit. on p. 59).

[Wai19a]   M. J. Wainwright. *High-dimensional Statistics: A Non-asymptotic Viewpoint*. Cambridge University Press, 2019 (cit. on p. 57).

[Wai19b]   M. J. Wainwright. "Stochastic approximation with cone-contractive operators: Sharper ℓ_∞-bounds for Q-learning". In: *arXiv preprint arXiv:1905.06265* (2019) (cit. on pp. 2, 4, 6).

[Wai19c]   M. J. Wainwright. "Variance-reduced Q-learning is minimax optimal". In: *arXiv preprint arXiv:1906.04697* (2019) (cit. on pp. 5, 6, 56, 60).

[WD92]   C. J. Watkins and P. Dayan. "Q-learning". In: *Machine Learning* 8.3-4 (1992), pp. 279–292 (cit. on pp. 56, 60).

[Xia+22]   E. Xia, K. Khamaru, M. J. Wainwright, and M. I. Jordan. "Instance-dependent confidence and early stopping for reinforcement learning". In: *arXiv preprint* (2022) (cit. on p. 36).

[YB13]   H. Yu and D. P. Bertsekas. "Q-learning and policy iteration algorithms for stochastic shortest path problems". In: *Annals of Operations Research* 208.1 (2013), pp. 95–132 (cit. on p. 16).

[ZZM21]   S. Zhang, Z. Zhang, and S. T. Maguluri. "Finite Sample Analysis of Average-Reward TD Learning and Q-Learning". In: *Advances in Neural Information Processing Systems* 34 (2021) (cit. on p. 4).

A   Restarting procedure

In this section, we describe a simple restarting procedure that allows us to refine the dependency of all of our bounds on the initial condition. This restarting procedure requires $\mathcal{O}(B_0 \log n)$ additional samples.

Algorithm 2 `ROOT-SA` with re-starting

1: Given (a) Initialization $\theta_0 \in \mathbb{V}$, (b) stepsize $\alpha > 0$, (c) number of restarting epochs R
2: Set $\theta_0^{(0)} = \theta_0$
3: **for** $i = 1, \ldots, R$ **do**
4:    Run Algorithm 1 with initial point $\theta_0^{(i-1)}$, stepsize α, burn-in period $B_0 = \frac{c}{(1-\gamma^2)\alpha} \log(\frac{n}{\delta})$,
       and sample size $t_0 = 2cB_0$; generate the sequence $(\theta_t^{(i-1)})_{t=0,1,\cdots,t_0}$.
5:    Set $\theta_0^{(i)} = \theta_{t_0}^{(i-1)}$.
6: **end for**
7: Run Algorithm 1 with initial point $\theta_0^\circ = \theta_0^{(R)}$, stepsize α, burn-in period $B_0 = \frac{c}{(1-\gamma^2)\alpha} \log(\frac{n}{\delta})$, and sample size $T = n - 2cRB_0$; generate the sequence $(\theta_t^\circ)_{t=0,1,\cdots,T}$.
8: **return** θ_T°

Given some fixed number $R \geq 1$ of restarting epochs, we can run the `ROOT-SA` algorithm for R consecutive short epochs, each with length $2cB_0$, with the constant c being the one in equation (9c). The last iterate θ_{2cB_0} of each short epoch is used as the initial point of the subsequent epoch, and in the end, the output of last short epoch is used as the initial point $\widetilde{\theta}_0$ to run a final single-epoch instantiation of `ROOT-SA` on the rest of the data stream. The detail of the re-starting procedure is described in Algorithm 2. In total, this restarting procedure uses an additional $2cB_0R$ samples, and the initialization of the last epoch satisfies the bound

$$\|\widetilde{\theta}_0 - \mathbf{h}(\widetilde{\theta}_0)\| \leq \frac{c}{\sqrt{B_0}} \big(\overline{\mathcal{W}} + \nu\sqrt{\log(1/\delta)}\big) + \frac{cb_*}{B_0(1-\gamma)} \big(\mathcal{J}_1(\Gamma, \rho_n) + \log(1/\delta)\big) + \frac{\|\theta_0 - \mathbf{h}(\theta_0)\|}{2^R}. \quad (71)$$

By choosing $R \geq \log\big(\|\mathbf{h}(\theta_0) - \theta_0\|\sqrt{n}/\overline{\mathcal{W}}\big)$ with a restarting sample size $2cB_0R$, we can ensure that $\frac{\|\theta_0 - \mathbf{h}(\theta_0)\|}{2^R} \leq \overline{\mathcal{W}}/\sqrt{n}$.

Our standard restarting procedure is based on the following conditions. We assume that the initialization θ_0 is such that the number of restarts R satisfies

$$\texttt{Initialization:} \qquad \log\left(\frac{\|\theta_0 - \mathbf{h}(\theta_0)\|\sqrt{n}}{\overline{\mathcal{W}}}\right) \leq c_0 \log n, \qquad (72\text{a})$$

for a universal constant $c_0 > 0$. In words, the condition ensures that the operator defect $\|\mathbf{h}(\theta_0) - \theta_0\|$ for the initialization θ_0 is not exponentially large compared to $\overline{\mathcal{W}}$. We set the number of restarts R as

$$\texttt{Number of restarts:} \qquad R = 2c_0 \log n \qquad (72\text{b})$$

These conditions ensure that performing R many restarts requires at most

$$2cB_0 \log\left(\frac{\|\mathbf{h}(\theta_0) - \theta_0\|\sqrt{n}}{\overline{\mathcal{W}}}\right) \lesssim 4c_0 B_0 \log(n)$$

additional samples, assuming that the original sample size is lower bounded as $n \gtrsim \frac{L^2 \mathcal{J}_2(\mathbb{B}^*, \rho_n)^2}{(1-\gamma)^4}$. Substituting this bound back to the bounds from Theorem 1 with the optimal stepsize choice (9b), we find that

$$\|\mathbf{h}(\theta_n) - \theta_n\| \leq \frac{c}{\sqrt{n}} \Big\{ \mathbb{E}[\|W\|] + \sqrt{\sigma_\Gamma^2(W) \log(\tfrac{1}{\delta})} \Big\} + \frac{cb_*}{(1-\gamma)n} \cdot \Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \Big\}. \qquad (73)$$

B Proofs of Theorem 2 and Corollary 2

In this section, we prove Theorem 2 and Corollary 2. In fact, Corollary 2 is actually a generalization of Theorem 2; the theorem follows from the corollary by setting $\|\cdot\|_C = \|\cdot\|$. Accordingly, we devote our effort to proving the corollary.

B.1 Proof of Corollary 2

Define the pair

$$r_\theta^*(t) := \frac{c}{\sqrt{t}}\big(\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\big) + \frac{cb_*}{1-\gamma}\Big\{\frac{1}{t} + \frac{\alpha L}{\sqrt{t}} \cdot \mathcal{J}_2(\mathbb{B}^*,\rho_n)\log(\tfrac{n}{\delta})\Big\} \Big\{\mathcal{J}_1(\Gamma,\rho_n) + \log(\tfrac{1}{\delta})\Big\}$$
$$+ c\frac{B_0}{t}\|\theta_0 - \mathbf{h}(\theta_0)\|, \quad \text{and} \quad (74a)$$

$$r_v^*(t) := \frac{c}{1-\gamma}\Big\{\frac{1}{t\sqrt{\alpha}}\big(\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\big) + \frac{b_*}{t}\Big\{\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma,\rho_n)\Big\}\Big\} + 2\big(\tfrac{B_0}{t}\big)^2\|\theta_0 - \mathbf{h}(\theta_0)\|. \tag{74b}$$

Invoking Theorem 1 and applying a union bound over the iterates, we have the pair of bounds $\|\mathbf{h}(\theta_t) - \theta_t\| \leq r_\theta^*(t)$, and and $\|v_t\| \leq r_v^*(t)$, uniformly for $t = B_0, B_0 + 1, \cdots, n$ with probability at least $1 - \delta$. Using the restarting scheme with parameter choice (72b), we can guarantee that the initial operator defect $\|\mathbf{h}(\theta_0) - \theta_0\|$ satisfies the upper bound:

$$\|\theta_0 - \mathbf{h}(\theta_0)\| \leq \frac{c}{\sqrt{B_0}}\big(\overline{W} + \nu\sqrt{\log(1/\delta)}\big) + \frac{cb_*}{B_0(1-\gamma)}\big(\mathcal{J}_1(\Gamma,\rho_n) + \log(1/\delta)\big).$$

By the linearization condition (A4)′, for any $\theta \in \mathbb{B}(\theta^*, s_0)$, we have

$$s := \big\|\theta - \theta^*\big\|_C \leq \sup_{A \in \mathcal{A}_s} \big\|(I - A)^{-1}(\mathbf{h}(\theta) - \theta)\big\|_C.$$

In order to obtain an upper bound on $\big\|\theta_n - \theta^*\big\|_C$, it suffices to provide an bound for the quantity $\sup_{A\in\mathcal{A}_s}\big\|(I - A)^{-1}(\mathbf{h}(\theta_{n-1}) - \theta_{n-1})\big\|_C$ for any given $s > 0$. Recall that $h(\theta_{n-1}) - \theta_{n-1} = v_n - z_n$, by definition, and in the rest of this section we provide upper bounds on $\big\|v_n\big\|_C$ and $\big\|z_n\big\|_C$

Upper bound on $\big\|v_n\big\|_C$

Observe that $\big\|(I - A)^{-1}v_n\big\|_C \leq \frac{\mathcal{D}}{1-\gamma}\|v_n\|$. Thus, if we invoke the bound $\|v_n\| \leq r_v^*(t)$, where r_v^* is defined in (74b), we are guaranteed that

$$\big\|(I - A)^{-1}v_n\big\|_C \leq \frac{c'\mathcal{D}}{(1-\gamma)^2}\Big\{\frac{1}{n\sqrt{\alpha}}\Big[\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big] + \frac{b_*}{n}\Big[\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma,\rho_n)\Big]\Big\}$$
$$+ \frac{c\mathcal{D}}{1-\gamma}\Big\{\frac{B_0}{n}\Big\}^2\|\theta_0 - \mathbf{h}(\theta_0)\|.$$

with probability at least $1 - \delta$.

Upper bound on $\left\|z_n\right\|_C$

In order to establish a sharp upper bound on the term $\sup_{A\in\mathcal{A}_s}\left\|(I-A)^{-1}z_{n+1}\right\|$, we define the class of test functions $\mathcal{S} := \left\{(I-A^*)^{-1}u : A\in\mathcal{A}_s, u\in C\right\}$. Substituting the the bounds (74a) and (74b) in Lemma 5 with we find that for any given $s>0$, the quantity $\sup_{A\in\mathcal{A}_s}\left\|(I-A)^{-1}z_n\right\|_C$ is upper bounded as

$$\frac{c}{\sqrt{n}}\left\{\mathbb{E}\Big[\sup_{A\in\mathcal{A}_s}\left\|(I-A)^{-1}W\right\|_C\Big] + \nu(s)\sqrt{\log(\tfrac{1}{\delta})}\right\} + \frac{cb_*\mathcal{D}}{n(1-\gamma)}\left\{\mathcal{J}_1(\Gamma,\rho_n) + \log(\tfrac{1}{\delta})\right\}$$

$$+ \frac{cL\mathcal{D}}{(1-\gamma)^2 n}\left\{\mathcal{J}_2(\mathbb{B}^*,\rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\right\}\left\{\alpha\Big(\sum_{s=B_0}^{n-1}s^2 r_v^{*2}(s)\Big)^{1/2} + \frac{1}{1-\gamma}\Big(\sum_{s=1}^{n-1}r_\theta^{*2}(s)\Big)^{1/2}\right\}, \quad (75)$$

with probability at least $1-\delta$.

Putting together the pieces

The last two bounds are valid for a fixed value of s. In order to derive the fixed-point condition in Theorem 2 and Corollary 2, however, we need a bound that holds uniformly over s in a suitable range, which we now do. Define the quantity

$$\overline{R}_n := \frac{c}{(1-\gamma)\sqrt{n}}\left\{\mathbb{E}\big[\left\|W\right\|_C\big] + \Big(\sup_{u\in C}\mathbb{E}\big[\langle u,\,W\rangle^2\big]\log(\tfrac{1}{\delta})\Big)^{1/2}\right\} + \mathcal{D}\cdot\mathcal{H}_n(\alpha,\delta),$$

and let $\underline{R} := \frac{1-\gamma}{1+\gamma}\overline{R}$.

It can be seen that the solutions to equation (25) all belong to the interval $[\underline{R},\overline{R}]$. In particular, contraction assumption (A1), we find that

$$\frac{1}{1+\gamma} \le \left\|\!\left\|(I-A)^{-1}\right\|\!\right\|_{\mathrm{op}} \le \frac{1}{1-\gamma}, \qquad \text{valid for any } A\in\mathcal{A}_s,$$

which leads to the bounds $\frac{1}{1+\gamma}\overline{\mathcal{W}}_C \le \mathcal{G}_C(s) \le \frac{1}{1-\gamma}\overline{\mathcal{W}}_C$ and $\frac{\nu_C}{1+\gamma} \le \nu_C(s) \le \frac{\nu_C}{1-\gamma}$ for any $s>0$.

By Theorem 1 and the contractive assumption (A1), we have the upper bound

$$\mathbb{P}\big[\left\|\theta_n - \theta^*\right\|_C \ge \overline{R}_n\big] \le \delta.$$

Consider the sequence $s_\ell = 2^{\ell-1}\underline{R}_n$ for $\ell = 1, 2, \ldots, k$, where $k := \log_2(\lceil\overline{R}_n/\underline{R}_n\rceil)$. It forms a doubling grid $\mathcal{M}_n := \{s_1, s_2, \cdots, s_k\}$ on the interval $[\underline{R}_n, \overline{R}_n]$, and it can be seen that k satisfies the upper bound

$$k \le \log(\tfrac{1+\gamma}{1-\gamma}) \le 1 + \log(\tfrac{1}{1-\gamma}).$$

Taking a union bound over $s\in\mathcal{M}_n$, we find that the bound (75) holds with probability at least $1-k\delta$, uniformly over $s\in\mathcal{M}_n$. For any $s\in[\underline{R}_n,\overline{R}_n]$, define the index $\ell(s) := \max\{\ell \mid s_\ell \le s\}$. On the event above, we can conclude that

$$\sup_{A\in\mathcal{A}_s}\left\|(I-A)^{-1}z_n\right\|_C \le \sup_{A\in\mathcal{A}_{s_{\ell(s)+1}}}\left\|(I-A)^{-1}z_n\right\|_C$$

$$\le \frac{c}{\sqrt{n}}\left\{\mathcal{G}_C(2s) + \nu_C(2s)\sqrt{\log(\tfrac{1}{\delta}) + \log(\log n)}\right\} + \frac{c\mathcal{D}b_*}{(1-\gamma)n}\left\{\mathcal{J}_1(\Gamma,\rho_n) + \log(\tfrac{1}{\delta}) + \log(\log n)\right\}$$

$$+ \frac{c\mathcal{D}L}{(1-\gamma)^2 n}\left\{\mathcal{J}_2(\mathbb{B}^*,\rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\right\}\left\{\alpha\Big(\sum_{s=B_0}^{n-1}s^2 r_v^{*2}(s)\Big)^{1/2} + \frac{1}{1-\gamma}\Big(\sum_{s=1}^{n-1}r_\theta^{*2}(s)\Big)^{1/2}\right\},$$

44

for $s \in [\underline{R}_n, \overline{R}_n]$. Here we have used the facts that $\mathcal{G}_C(\cdot)$ and $\nu_C(\cdot)$ are non-decreasing functions. We now substitute our expressions for r_θ^* and r_v^*, and conclude that conditioned on the event $\mathcal{E}_n^{(\theta)} \cap \mathcal{E}_n^{(v)} \cap \{\|\theta_n - \theta^*\| \leq \overline{R}_n\}$, we have

$$
\begin{aligned}
s_n &\leq \sup_{A \in \mathcal{A}_{s_n}} \left\| (I - A)^{-1} z_n \right\|_C + \tfrac{\mathcal{D}}{1-\gamma} r_v^*(n) \\
&\leq \tfrac{c}{\sqrt{n}} \left(\mathcal{G}(2s_n) + \nu(2s_n) \sqrt{\log(\tfrac{1}{\delta})} \right) + \underline{R}_n + \tfrac{c\mathcal{D}b_*}{(1-\gamma)n} \left(\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \right) \\
&\quad + \tfrac{c\mathcal{D}L}{(1-\gamma)^2 \sqrt{n}} \cdot \mathcal{J}_2(\mathbb{B}^*, \rho_n) \log(\tfrac{n}{\delta}) \cdot \left(\sqrt{\alpha} \overline{\mathcal{W}} + \alpha b_* \left\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta}) \right\} + \sqrt{\tfrac{B_0}{n}} \|\theta_0 - \mathbf{h}(\theta_0)\| \right) \\
&\quad + \tfrac{c'\mathcal{D}}{(1-\gamma)^2} \left\{ \tfrac{1}{n\sqrt{\alpha}} \overline{\mathcal{W}} \sqrt{\log \tfrac{n}{\delta}} + \tfrac{b_*}{n} \left[\log \tfrac{n}{\delta} + \mathcal{J}_1(\Gamma, \rho_n) \right] \right\} + \tfrac{c\mathcal{D}}{1-\gamma} \left(\tfrac{B_0}{n} \right)^2 \|\theta_0 - \mathbf{h}(\theta_0)\| \\
&\leq \tfrac{c}{\sqrt{n}} \left\{ \mathcal{G}_C(2s_n) + \nu_C(2s_n) \sqrt{\log(\tfrac{1}{\delta})} \right\} + \tfrac{c\mathcal{D} \log(n/\delta)}{(1-\gamma)^2} \left\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) L \sqrt{\tfrac{\alpha}{n}} + \tfrac{1}{n\sqrt{\alpha}} \right\} \cdot \overline{\mathcal{W}} \\
&\quad + \tfrac{c\mathcal{D}b_* \log(\tfrac{n}{\delta})}{(1-\gamma)^2} \left\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) L \tfrac{\alpha}{\sqrt{n}} + \tfrac{1}{n} \right\} \cdot \left\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta}) \right\} + \tfrac{\mathcal{D}B_0}{(1-\gamma)n} \cdot \|\theta_0 - \mathbf{h}(\theta_0)\|,
\end{aligned}
$$

with probability at least $1 - \delta$, valid for any $\delta \in (0, 1/k)$, where $k = 1 + \log \tfrac{1}{1-\gamma}$.

Finally, noting that $\mathbb{P}\big[\mathcal{E}_n^{(\theta)} \cap \mathcal{E}_n^{(v)} \cap \{\|\theta_n - \theta^*\| \leq \overline{R}_n\}\big] \geq 1 - \delta$, and using the initialization conditions (72), we obtain the bound that was claimed in Corollary 2.

C Proofs for multi-step contractions

This section is devoted to the proofs of our results on multi-step contractions, with Theorem 3 proved in Appendix C.1 and Corollary 3 in Appendix C.2.

C.1 Proof of Theorem 3

The proof of this theorem is similar to that of Theorem 1, but is based on an improved version of Lemma 1, stated as Lemma 6. At a high level, there are three main steps:

1. First, we use Lemma 2 and Lemma 6 to establish a relation between $\|\mathbf{h}(\theta_t) - \theta_t\|$ and $\|v_t\|$.

2. Second, starting with the coarse bound on $\|\mathbf{h}(\theta_t) - \theta_t\|$ and $\|v_t\|$ from Lemma 3, we iteratively refine our bounds using the relation from Step 1.

3. Finally, we improve the higher-order terms in these bounds.

C.1.1 Step 1: Relating $\|\mathbf{h}(\theta_t) - \theta_t\|$ and $\|v_t\|$

We first state a sharpening of Lemma 1 that holds for a multi-step contractive linear operator (see Assumption (A1)′).

Lemma 6. *Under assumptions (A1)′, (A3), and (A2), there exists a universal constant $c > 0$ such that for stepsize α satisfying the bound*

$$
c\sqrt{m\alpha} \cdot L\mathcal{J}_2(\mathbb{B}^*, \rho_n) \cdot \log \tfrac{n}{\delta} \leq \tfrac{1}{3}, \tag{76a}
$$

45

and the burn-in period $B_0 \geq \frac{cm}{\alpha}$, given any κ-admissible sequences $r_\theta(t)$ and $r_v(t)$ with $0 < \kappa \leq 2$, on the event $\mathcal{E}_n^{(v)}(r_v) \cap \mathcal{E}_n^{(\theta)}(r_\theta)$, the following bound holds uniformly with respect to $t \in [B_0, n]$, with probability $1 - \delta$:

$$\|v_t\| \leq \frac{2r_v(t)}{3} + \frac{cmr_\theta(t)}{t\alpha} + \frac{c}{t}\sqrt{\frac{m}{\alpha}}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\} + \frac{cb_*}{t}\big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta})\big\} + 4\big(\tfrac{B_0}{t}\big)^2\|v_{B_0}\|. \tag{76b}$$

See Section C.1.4 for the proof of this lemma.

In addition, by Lemma 5 and the operator norm bound on $(I - A)^{-1}$, conditioned on the event $\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)$, we have the following bound uniformly for $t \in [B_0, n]$,

$$\|z_t\| \leq \frac{c}{\sqrt{t}}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\} + \frac{cb_*}{t}\big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta})\big\}$$
$$+ \frac{cL}{t}\Big\{\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{n}{\delta})}\Big\}\Big\{\alpha\Big(\sum_{s=B_0}^{t-1} s^2 r_v^2(s)\Big)^{1/2} + m\Big(\sum_{s=1}^{t-1} r_\theta^2(s)\Big)^{1/2}\Big\}, \tag{77}$$

with probability at least $1 - \delta$.

C.1.2 Step 2: Bounds using bootstrapping

Akin to the proof of Theorem 1, we impose the restrictions that the estimate sequences (r_θ, r_v) are $\frac{1}{2}$- and 1-admissible, respectively.

Consider a new pair (r_v^+, r_θ^+) satisfying the initial bounds $r_v^+(B_0) \geq \|v_{B_0}\|$ and $r_\theta^+(B_0) \geq \|\mathbf{h}(\theta_0) - \theta_0\|$, and such that

$$r_v^+(t) \geq \tfrac{2}{3}r_v(t) + \frac{1}{t\sqrt{\alpha}} \cdot \frac{cm}{\sqrt{t\alpha}} \cdot \sqrt{t}r_\theta(t)$$
$$+ \frac{c}{t}\sqrt{\frac{m}{\alpha}}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\} + c\frac{b_*}{t}\Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta})\Big\} + 4\big(\tfrac{B_0}{t}\big)^2\|v_{B_0}\|, \quad \text{and} \tag{78a}$$

$$r_\theta^+(t) \geq \frac{c}{\sqrt{t}}\Big\{\frac{1}{\sqrt{\alpha t}} + \sqrt{\alpha}L\mathcal{J}_2(\mathbb{B}^*, \rho_n)\log\tfrac{n}{\delta}\Big\} \cdot \Big\{t\sqrt{\alpha}r_v(t)\Big\} + 2c\frac{Lm}{\sqrt{t}}\mathcal{J}_2(\mathbb{B}^*, \rho_n)\log\tfrac{n}{\delta} \cdot r_\theta(t),$$
$$+ \frac{c}{\sqrt{t}}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\} + \frac{cb_*}{t}\Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta})\Big\}. \tag{78b}$$

for each integer $t \in [B_0, n]$.

By combining the bounds (76b) and (77), we are guaranteed that

$$\mathbb{P}\Big[\mathcal{E}_n^{(\theta)}(r_\theta^+) \cap \mathcal{E}_n^{(v)}(r_v^+)\Big] \geq \mathbb{P}\Big[\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)\Big] - \delta,$$

Our goal is to construct two series of admissible sequences $\big(r_v^{(i)}, r_v^{(i)}\big)$ with $i = 0, 1, \cdots$, such that the pair $(\|v_t\|, \|\mathbf{h}(\theta_t) - \theta_t\|)_{t \geq B_0}$ are dominated by $\big(r_v^{(i)}(t), r_v^{(i)}(t)\big)_{t \geq 0}$, with high probability. Concretely, we consider sequences of a particular form $r_v^{(i)}(t) = \frac{\psi_v^{(i)}}{t\sqrt{\alpha}}$ and $r_\theta^{(i)}(t) = \frac{\psi_\theta^{(i)}}{\sqrt{t}}$, for pairs of positive reals $\big(\psi_v^{(i)}, \psi_\theta^{(i)}\big)$ independent of t. Apparently, with such forms, the sequence $r_\theta^{(i)}$ is $\frac{1}{2}$-admissible, and the sequence $r_v^{(i)}$ is 1-admissible. However, if we directly substitute the sequences $\big(r_v^{(i)}(t), r_v^{(i)}(t)\big)$ of such forms into the iteration (78), the resulting

46

sequences (r_θ^+, r_v^+) will no longer be of the desired form. So in order to unify the coefficients in equation (78) into the same time scale, given a stepsize $\alpha > 0$, we define the burn-in time

$$B_0 = \frac{cm}{\alpha} \log(\frac{n}{\delta}). \tag{79a}$$

For each $t = B_0, B_0 + 1 \dots$, the coefficients in (78) then satisfy the bounds

$$\frac{cm}{\sqrt{\alpha t}} \leq \frac{1}{6}\sqrt{m}, \quad \frac{c}{\sqrt{\alpha t}} \leq \frac{1}{12\sqrt{m}}, \quad \text{and} \quad \frac{m}{\sqrt{t}} \log(\frac{n}{\delta}) \leq \sqrt{\alpha m}. \tag{79b}$$

Therefore, if we construct a two-dimensional vector sequence $\psi^{(i)} = \begin{bmatrix} \psi_v^{(i)} & \psi_\theta^{(i)} \end{bmatrix}^T$ satisfying the recursive relation $\psi^{(i+1)} = Q\psi^{(i)} + b$, where

$$Q := \begin{bmatrix} 2/3 & \frac{\sqrt{m}}{6} \\ \frac{1}{12\sqrt{m}} + cL\mathcal{J}_2(\mathbb{B}^*, \rho_n)\sqrt{\alpha} \cdot \log(\frac{n}{\delta}) & 2cL\mathcal{J}_2(\mathbb{B}^*, \rho_n)\sqrt{\alpha m} \log(\frac{n}{\delta}) \end{bmatrix}, \quad \text{and}$$

$$b := c \cdot \begin{bmatrix} \sqrt{m}\left(\overline{W} + \nu\sqrt{\log(n/\delta)}\right) + b_*\sqrt{\alpha}\left(\log(\frac{1}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\right) + B_0\sqrt{\alpha}\|v_{B_0}\| \\ \left(\overline{W} + \nu\sqrt{\log(n/\delta)}\right) + b_*\sqrt{\frac{\alpha}{m}}\left(\log(n/\delta) + \mathcal{J}_1(\Gamma, \rho_n)\right) + \sqrt{B_0}\|\mathbf{h}(\theta_0) - \theta_0\| \end{bmatrix}, \tag{80}$$

they will satisfy the requirement (78), leading to the probability bound:

$$\mathbb{P}\left[\mathcal{E}_n^{(\theta)}(r_\theta^{(i+1)}) \cap \mathcal{E}_n^{(v)}(r_v^{(i+1)})\right] \geq \mathbb{P}\left[\mathcal{E}_n^{(\theta)}(r_\theta^{(i)}) \cap \mathcal{E}_n^{(v)}(r_v^{(i)})\right] - \delta, \tag{81}$$

for the sequences $r_\theta^{(i)}(t) = \psi_\theta^{(i)}/\sqrt{t}$ and $r_v^{(i)}(t) = \psi_v^{(i)}/(\sqrt{\alpha}t)$.

It remains to specify an initial condition for the recursion above. Note that Lemma 3 implies that we have

$$\|\theta_t - \theta^*\| + \|v_t\| \leq e^{1+L\alpha t}\left(b_* + \|\theta_0 - \theta^*\|\right)$$

almost surely. So we can take the initialization:

$$\psi_v^{(0)} := n\sqrt{\alpha}e^{1+L\alpha n}(b_* + \|\theta_0 - \theta^*\|), \quad \text{and} \quad \psi_\theta^{(0)} := \sqrt{n}e^{1+L\alpha n}(b_* + \|\theta_0 - \theta^*\|),$$

for which the bounds $\|v_t\| \leq \frac{\psi_v^{(0)}}{t\sqrt{\alpha}}$ and $\|\theta_t - \mathbf{h}(\theta_t)\| \leq \frac{\psi_\theta^{(0)}}{\sqrt{t}}$ hold almost surely.

Given such an initial condition and the recursion (81), we find that

$$\mathbb{P}\left[\mathcal{E}_n^{(v)}(r_v^{(i)}) \cap \mathcal{E}_n^{(\theta)}(r_\theta^{(i)})\right] \geq \mathbb{P}\left[\mathcal{E}_n^{(v)}(r_v^{(0)}) \cap \mathcal{E}_n^{(\theta)}(r_\theta^{(0)})\right] - i\delta = 1 - i\delta.$$

It remains to understand the behavior of $\psi^{(i)}$ for large values of the index i, i.e. the after i iterations of the bootstrapping argument. We do so by solving the recursion $\psi^{(i+1)} = Q\psi^{(i)} + b$. Let us define a new matrix

$$\widetilde{Q} := \begin{bmatrix} 2/3 & \frac{\sqrt{m}}{6} \\ \frac{1}{6\sqrt{m}} & 2/3 \end{bmatrix} \overset{(i)}{=} \begin{bmatrix} \sqrt{m} & \sqrt{m} \\ 1 & -1 \end{bmatrix} \cdot \begin{bmatrix} \frac{5}{6} & 0 \\ 0 & \frac{1}{2} \end{bmatrix} \cdot \begin{bmatrix} \sqrt{m} & \sqrt{m} \\ 1 & -1 \end{bmatrix}^{-1},$$

where the equivalence (i) follows by a direct calculation. Note that the stepsize condition (29a) ensures that

$$cL\mathcal{J}_2(\mathbb{B}^*, \rho_n) \log\frac{n}{\delta} \cdot \sqrt{m\alpha} \leq \frac{1}{12}, \tag{82}$$

47

then the matrix $\widetilde{Q}$ is coordinate-wise larger than the matrix Q from equation (80), and consequently we are guaranteed that $Qu \preceq_{\text{orth}} \widetilde{Q}u$ for any 2-dimensional vector $u \succeq_{\text{orth}} 0$. Thus, for each integer $N = 1, 2, \ldots,$ we have the upper bounds

$$
\begin{bmatrix} \psi_v^{(N)} \\ \psi_\theta^{(N)} \end{bmatrix} = \left(\sum_{i=0}^{N-1} Q^i \right) b + Q^N \begin{bmatrix} \psi_v^{(0)} \\ \psi_\theta^{(0)} \end{bmatrix} \preceq_{\text{orth}} \left(\sum_{i=0}^{N-1} \widetilde{Q}^i \right) b_\psi + \widetilde{Q}^N \begin{bmatrix} \psi_v^{(0)} \\ \psi_\theta^{(0)} \end{bmatrix}
$$
$$
\preceq_{\text{orth}} (I - \widetilde{Q})^{-1} b + e^{-N/6} \sqrt{m} \left(\psi_v^{(0)} + \psi_\theta^{(0)} \right) \mathbf{1}_2.
$$

By taking $N = cLn \log n$, replacing δ with δ/N and substituting with the above bounds, we find that

$$
t\sqrt{\alpha} \cdot \|v_t\| \leq \psi_v^{(N)} \leq c\sqrt{m} \left\{ \overline{W} + \nu \sqrt{\log(\tfrac{n}{\delta})} \right\}
$$
$$
+ cb_* \sqrt{\alpha} \left(\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n) \right) + cB_0 \sqrt{\alpha} \|v_{B_0}\| + \sqrt{B_0 m} \|\mathbf{h}(\theta_0) - \theta_0\|, \quad (83a)
$$

along with

$$
\sqrt{t} \|\mathbf{h}(\theta_t) - \theta_t\| \leq \psi_\theta^{(N)} \leq c \left\{ \overline{W} + \nu \sqrt{\log(\tfrac{n}{\delta})} \right\} + cb_* \sqrt{\alpha} \left\{ \log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n) \right\}
$$
$$
+ cB_0 \sqrt{\alpha/m} \|v_{B_0}\| + \sqrt{B_0} \|\mathbf{h}(\theta_0) - \theta_0\|, \quad (83b)
$$

valid uniformly over $t \in \{B_0, B_0 + 1, \cdots, n\}$ with probability at least $1 - \delta$.

The latter bound, when combined with the Lemma 4 yields an upper bound on $\|\mathbf{h}(\theta_t) - \theta_t\|$ which has the correct leading-order term, i.e., the correct dependence on the term $\overline{W} + \nu \sqrt{\log \tfrac{n}{\delta}}$. In order to refine the dependence on the terms $\|\mathbf{h}(\theta_0) - \theta_0\|$ and $\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)$, we need do another round of bootstrapping.

C.1.3   Step 3: Improving the higher-order terms

With a slight abuse of notation, let the 2-vector $\psi^{(N)} := (\psi_v^{(N)}, \psi_\theta^{(N)})$ be defined by the right-hand side of equation (83), and consider the choices $r_v(t) := \frac{\psi_v^{(N)}}{t\sqrt{\alpha}}$ and $r_\theta(t) := \frac{\psi_\theta^{(N)}}{\sqrt{t}}$. Conditioned on the event $\mathcal{E}_n^{(\theta)}(r_\theta) \cap \mathcal{E}_n^{(v)}(r_v)$, we have

$$
\|\mathbf{h}(\theta_t) - \theta_t\| \leq \frac{c}{\sqrt{t}} \left\{ \overline{W} + \nu \sqrt{\log(\tfrac{1}{\delta})} \right\} + \frac{cb_*}{t} \left\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \right\}
$$
$$
+ \left\{ \frac{1}{t} + c \frac{\alpha L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log(\tfrac{n}{\delta}) \right\} \frac{\psi_v^{(N)}}{\sqrt{\alpha}} + 2c \frac{mL \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{t} \log(\tfrac{n}{\delta}) \cdot \psi_\theta^{(N)}
$$
$$
\leq \left\{ \frac{c}{\sqrt{t}} + \sqrt{\frac{m}{\alpha}} \left[\frac{1}{t} + c \frac{\alpha L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log(\tfrac{n}{\delta}) \right] \right\} \left\{ \overline{W} + \nu \sqrt{\log(\tfrac{n}{\delta})} \right\}
$$
$$
+ c'b_* \left\{ \frac{1}{t} + \frac{\alpha L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log(\tfrac{n}{\delta}) \right\} \left\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta}) \right\}
$$
$$
+ c' \left\{ \frac{1}{t} \sqrt{\frac{m}{\alpha}} + \frac{\sqrt{\alpha m} L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log(\tfrac{n}{\delta}) + \frac{mL \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{t} \log(\tfrac{n}{\delta}) \right\} \sqrt{B_0} \|\mathbf{h}(\theta_0) - \theta_0\|
$$
$$
+ c' \left\{ \frac{1}{t} + \frac{\alpha L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log(\tfrac{n}{\delta}) + \frac{\sqrt{\alpha/m} L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{t} \log(\tfrac{n}{\delta}) \right\} B_0 \|v_{B_0}\|,
$$

with probability at least $1 - \delta$.

Given a burn-in period B_0 satisfying (79a) and step size satisfying (82), using the bound on $\|v_{B_0}\|$ from Lemma 4, we have the upper bound $\|\mathbf{h}(\theta_t) - \theta_t\| \leq \widetilde{r}_\theta(t)$, with probability at

least $1 - \delta$, uniformly over all integers $t \in [n]$, where

$$\widetilde{r}_\theta(t) := \tfrac{c_1}{\sqrt{t}}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\} + c_2 b_*\Big\{\tfrac{1}{t} + \tfrac{\alpha L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log^3(\tfrac{n}{\delta})\Big\} \Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta})\Big\}$$
$$+ c_2\Big\{\tfrac{\alpha B_0 L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}} \cdot \log(\tfrac{n}{\delta}) + \tfrac{B_0}{t}\Big\} \|\mathbf{h}(\theta_0) - \theta_0\|.$$

By substituting the upper bound $\widetilde{r}_\theta$ into equation (78a), we obtain a recursive inequality that takes as input an admissible sequence $r_v(t)$, and generates as output a new sequence $r_v^+(t)$ such that

$$\mathbb{P}\Big[\mathcal{E}_n^{(v)}(r_v^+)\Big] \geq \mathbb{P}\Big[\mathcal{E}_n^{(v)}(r_v)\Big] - \delta.$$

Taking any integer $N_1 > 0$, by applying the recursive inequality for N_1 times with $\delta' = \delta/N_1$, we get a sharper bound for $\|v_t\|$ with probability $1 - \delta$:

$$\|v_t\| \leq 3c\left[\tfrac{\sqrt{m}}{t\sqrt{\alpha}}\Big(\overline{W} + \nu\sqrt{\log(\tfrac{nN_1}{\delta})}\Big) + \tfrac{b_*}{t}\big(\log(\tfrac{nN_1}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\big)\right]$$
$$+ \tfrac{cm}{\alpha t}\widetilde{r}_\theta(t) + c\big(\tfrac{B_0}{t}\big)^2 \|v_{B_0}\| + \big(\tfrac{1+\gamma}{2}\big)^{N_1} \cdot \tfrac{\psi_v^{(N)}}{t\sqrt{\alpha}}.$$

Taking $N_1 := 10 \log n$, for stepsize and burn-in period satisfying the conditions (79a) and (82), some algebra yields that $\|v_t\| \leq \widetilde{r}_v(t)$ with probability at least $1 - \delta$, uniformly for each integer $t \in [B_0, n]$, where

$$\widetilde{r}_v(t) := c'\Big\{\tfrac{1}{t}\sqrt{\tfrac{m}{\alpha}}\big[\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\big] + \tfrac{b_*}{t}\big[\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\big]\Big\} + 2c'\big(\tfrac{B_0}{t}\big)^2\|\theta_0 - \mathbf{h}(\theta_0)\| \quad (84)$$

for a universal constant $c' > 0$.

It can be seen that the sequences $\widetilde{r}_v$ and $\widetilde{r}_\theta$ are 2-admissible. Substituting their definitions into the bound (77). we find that the inequality

$$\|z_t\| \leq \tfrac{c}{\sqrt{t}}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})}\Big\} + \tfrac{cb_*}{t}\Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta})\Big\}$$
$$+ \tfrac{cL}{t}\Big\{\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\Big\} \Big\{\alpha\Big(\sum_{s=B_0}^{t-1} s^2 r_v^2(s)\Big)^{1/2} + m\Big(\sum_{s=1}^{t-1} r_\theta^2(s)\Big)^{1/2}\Big\}$$

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

Under the conditions (82) and (79a), some algebra yields:

$$\|z_t\| \leq \tfrac{c}{\sqrt{t}}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})}\Big\}$$
$$+ cb_*\Big\{\tfrac{1}{t} + \tfrac{\alpha L}{\sqrt{t}} \cdot \mathcal{J}_2(\mathbb{B}^*, \rho_n)\log(\tfrac{n}{\delta})\Big\} \Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta})\Big\} + c\tfrac{B_0}{t}\|\theta_0 - \mathbf{h}(\theta_0)\|.$$

Combining with equation (84) yields the upper bound

$$\|\mathbf{h}(\theta_t) - \theta_t\| \leq \tfrac{c}{\sqrt{t}}\big(\overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})}\big)$$
$$+ cb_*\Big\{\tfrac{1}{t} + \tfrac{\alpha L}{\sqrt{t}}\mathcal{J}_2(\mathbb{B}^*, \rho_n)\log(\tfrac{n}{\delta})\Big\} \Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta})\Big\} + c\tfrac{B_0}{t}\|\theta_0 - \mathbf{h}(\theta_0)\|, \quad (85)$$

which completes the proof of equation (30).

Besides, by taking a union bound over time steps $t \in \{B_0, B_0 + 1, \ldots, n\}$, we have the lower bound $\mathbb{P}\big[\mathcal{E}_n^{(\theta)}(r_\theta^*)\big] \geq 1 - \delta$, where

$$r_\theta^*(t) := \frac{c}{\sqrt{t}}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\}$$
$$+ cb_*\Big\{\tfrac{1}{t} + \tfrac{\alpha L \mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}}\log(\tfrac{n}{\delta})\Big\}\Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta})\Big\} + c\tfrac{B_0}{t}\|\theta_0 - \mathbf{h}(\theta_0)\|.$$

C.1.4 Proof of Lemma 6

Starting with recursion satisfied by v_t, we have

$$t \cdot v_t = (t-1)\big\{v_{t-1} + \theta_{t-2} - \mathbf{H}_t(\theta_{t-1}) - \theta_{t-1} + \mathbf{H}_t(\theta_{t-2})\big\} + \big\{\mathbf{H}(\theta_{t-1}) - \theta_{t-1}\big\}$$
$$= \big\{(1-\alpha)I + \alpha A\big\} \cdot (t-1)v_{t-1} - (t-1)\big\{\varepsilon_t(\theta_{t-1}) - \varepsilon_t(\theta_{t-2})\big\} + \varepsilon_t(\theta_{t-1}) + \big\{\mathbf{h}(\theta_{t-1}) - \theta_{t-1}\big\}.$$

For any positive integer τ, we can expand the above expression for τ steps so as to obtain

$$t \cdot v_t = \big((1-\alpha)I + \alpha A\big)^\tau (t-\tau)v_{t-\tau} - \sum_{j=1}^{\tau}(t-j)\big((1-\alpha)I + \alpha A\big)^{j-1}\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j-1})\big)$$

$$+ \sum_{j=1}^{\tau}\big((1-\alpha)I + \alpha A\big)^{j-1}\varepsilon_{t-j+1}(\theta_{t-j}) + \sum_{j=1}^{\tau}\big((1-\alpha)I + \alpha A\big)^{j-1}\big(\mathbf{h}(\theta_{t-j}) - \theta_{t-j}\big). \quad (86)$$

In addition, our analysis makes use of the following auxiliary bound

$$\big\|\big((1-\alpha)I + \alpha A\big)^t\big\|_{\mathbb{V}} \leq \min\Big\{1, 2\big(1 - \tfrac{\alpha}{2m}\big)^t\Big\}, \quad (87)$$

valid for all $t = 1, 2, \ldots$. See the end of this subsection for the proof of this claim.

Taking this bound as given, we proceed with the proof of this lemma. First, substituting the bound (87) into the decomposition (86) yields the bound

$$t \cdot \|v_t\| \leq 2\big(1 - \tfrac{\alpha}{2m}\big)^\tau (t-\tau)\|v_{t-\tau}\| + \|\Psi_{t-\tau,\tau}\| + \|M_{t-\tau,\tau}\| + \sum_{j=1}^{\tau}\|\mathbf{h}(\theta_{t-j}) - \theta_{t-j}\|. \quad (88)$$

where we define the terms

$$\Psi_{t-\tau,\tau} := \sum_{j=1}^{\tau}(t-j)\big((1-\alpha)I + \alpha A\big)^{j-1}\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j-1})\big), \quad \text{and} \quad (89a)$$

$$M_{t-\tau,\tau} := \sum_{j=1}^{\tau}\big((1-\alpha)I + \alpha A\big)^{j-1}\varepsilon_{t-j+1}(\theta_{t-j}). \quad (89b)$$

Now we bound the terms in the decomposition (88). On the event $\mathcal{E}_n^{(v)}(r_v)$, each term in the summation defining $\Psi_{t-\tau,\tau}$ satisfies an almost-sure upper bound:

$$\big((1-\alpha)I + \alpha A\big)^{j-1}\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta_{t-j-1})\big) \leq (t-j)L\alpha\|v_{t-j}\| \leq (t-j)L\alpha r_v(t-j).$$

Since the sequence r_v is admissible, for burn-in time $B_0 \geq 2\tau$, we have that $(t-j)r_\theta(t-j) \leq \frac{t^2}{(t-j)}r_v(t) \leq 2tr_v(t)$. Note that the terms in $\Psi_{t-\tau,\tau}$ form a martingale difference sequence,

adapted to the natural filtration $(\mathcal{F}_t)_{t\geq 0}$. Invoking the martingale concentration inequality from Lemma 8 yields the bound

$$\|\Psi_{t-\tau,\tau}\| \leq c\sqrt{\tau}\big\{\mathcal{J}_2(\mathbb{B}^*,\rho_n) + \sqrt{\log(1/\delta)}\big\} \cdot L\alpha tr_v(t), \qquad (90)$$

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

As for the term $M_{t-\tau,\tau}$, we use a decomposition similar to the one used in the proof of Lemma 1:

$$M_{t-\tau,\tau} = \sum_{j=1}^{\tau}\big((1-\alpha)I + \alpha A\big)^{j-1}\varepsilon_{t-j+1}(\theta^*) + \sum_{j=1}^{\tau}\big((1-\alpha)I + \alpha A\big)^{j-1}\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta^*)\big)$$

$$=: M^*_{t-\tau,\tau} + \widetilde{M}_{t-\tau,\tau}.$$

The term $M^*_{t-\tau,\tau}$ is sum of independent random variables in $\mathbb{V}$, with each term satisfying the conditions

$$\big\|\big((1-\alpha)I + \alpha A\big)^{j-1}\varepsilon_{t-j+1}(\theta^*)\big\| \leq \cdot\|\varepsilon_{t-j+1}(\theta^*)\|, \quad \text{and} \quad (1-\alpha)I + \alpha A\big)^{j-1}\varepsilon_{t-j+1}(\theta^*) \in \Omega.$$

Invoking the concentration inequality from Lemma 7 yields the bound

$$\|M^*_{t-\tau,\tau}\| \leq c\sqrt{\tau}\left(\overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})}\right) + cb_*\left(\mathcal{J}_1(\Gamma,\rho_n) + \log(\tfrac{1}{\delta})\right),$$

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

For the excess noise term $\widetilde{M}_{t-\tau,\tau}$, we note that conditioned on the event $\mathcal{E}_n^{(\theta)}(r_\theta)$, we have the upper bound

$$\left\|\big((1-\alpha)I + \alpha A\big)^{j-1}\big(\varepsilon_{t-j+1}(\theta_{t-j}) - \varepsilon_{t-j+1}(\theta^*)\big)\right\| \leq 2Lmr_\theta(t-j).$$

For an admissible sequence r_θ and burn-in period $B_0 \geq 2\tau$, we have that $r_\theta(t-j) \leq \frac{t^2}{(t-j)^2}r_\theta(t) \leq 4r_\theta(t)$ for any $j \in [\tau]$. Furthermore, the terms in $\widetilde{M}_{t-\tau,\tau}$ form a martingale difference sequence adapted to the natural filtration. By Lemma 8, on the event $\mathcal{E}_n^{(\theta)}(r_\theta)$, we have the martingale concentration inequality:

$$\|\widetilde{M}_{t-\tau,\tau}\| \leq c\sqrt{\tau}\big(\mathcal{J}_2(\mathbb{B}^*,\rho_n) + \sqrt{\log(1/\delta)}\big) \cdot Lmr_\theta(t).$$

Finally, for the last term in the decomposition (88), we note that on the event $\mathcal{E}_n^{(\theta)}(r_\theta)$, we have the bounds:

$$\|\mathbf{h}(\theta_{t-j}) - \theta_{t-j}\| \leq r_\theta(t-j) \leq \tfrac{t^2}{(t-j)^2}r_\theta(t) \leq 4r_\theta(t).$$

In order to prove the final results, as with the proof of Lemma 1, we consider the cases of $t \geq B_0 + 2m/\alpha$ and $t \leq B_0 + 2m/\alpha$ separately.

When $t \geq B_0 + 2m/\alpha$, collecting above bounds, by taking $\tau = 2m/\alpha$, we find that

$$t \cdot \|v_t\| \leq \left\{\tfrac{1}{3} + c\sqrt{\tau}\big[\mathcal{J}_2(\mathbb{B}^*,\rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\big] \cdot L\alpha\right\}tr_v(t)$$

$$+ \left\{c\sqrt{\tau}\big[\mathcal{J}_2(\mathbb{B}^*,\rho_n) + \sqrt{\log(\tfrac{1}{\delta})}\big] \cdot Lm + \tau\right\}r_\theta(t)$$

$$+ c\sqrt{\tau}\left\{\overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})}\right\} + cb_*\left\{\mathcal{J}_1(\Gamma,\rho_n) + \log(\tfrac{1}{\delta})\right\}.$$

Given a stepsize α such that

$$c\sqrt{m\alpha} \cdot \left\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \right\} \cdot L \leq \tfrac{1}{3}, \tag{91}$$

the above inequality implies that

$$t \cdot \|v_t\| \leq \tfrac{2}{3} tr_v(t) + \tfrac{cm}{\alpha} r_\theta(t) + c\sqrt{\tfrac{m}{\alpha}} \left\{ \overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})} \right\} + cb_* \right\} \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \right\},$$

which completes the proof of the first case.

On the other hand, when $t \leq B_0 + 2m/\alpha$, we let $\tau = t - B_0$, and find that:

$$t \cdot \|v_t\| \leq 2B_0 \cdot \|v_{B_0}\| + c\sqrt{\tau} \left[\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \right] \cdot L\alpha tr_v(t)$$
$$+ \left\{ c\sqrt{\tau} \left[\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \right] \cdot Lm + \tau \right\} r_\theta(t)$$
$$+ c\sqrt{\tau} \left\{ \overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})} \right\} + cb_* \left\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \right\}.$$

Note that for $t \in [B_0, B_0 + 2m/\alpha]$, we have that $2B_0 \cdot \|v_{B_0}\| \leq 4\frac{B_0{}^2}{t}\|v_{B_0}\|$. Assuming the stepsize condition (91), we conclude the inequality:

$$t \cdot \|v_t\| \leq \tfrac{2}{3} tr_v(t) + \tfrac{cm}{\alpha} r_\theta(t) + c\sqrt{\tfrac{m}{\alpha}} \left\{ \overline{W} + \nu\sqrt{\log(\tfrac{1}{\delta})} \right\} + cb_* \right\} \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \right\} + \frac{B_0{}^2}{t}\|v_{B_0}\|,$$

Proof of equation (87): Applying the triangle inequality yields

$$\left\| \left((1-\alpha)I + \alpha A \right)^t \right\|_{\mathbb{V}} \leq \sum_{k=0}^{t} \binom{t}{k} (1-\alpha)^k \alpha^{t-k} \|A^{t-k}\|_{\mathbb{V}}. \tag{92}$$

Since $\|A^t\|_{\mathbb{V}} \leq \|A\|_{\mathbb{V}}^t \leq 1$ for each $t = 0, 1, 2, \ldots$, we have

$$\left\| \left((1-\alpha)I + \alpha A \right)^t \right\|_{\mathbb{V}} \leq \sum_{k=0}^{t} \binom{t}{k} (1-\alpha)^k \alpha^{t-k} \leq 1.$$

On the other hand, we note that for any time index $i \in \mathbb{N}_+$, using the m-step contraction condition (A1)', we have that:

$$\|A^i\|_{\mathbb{V}} \leq \|A^m\|_{\mathbb{V}}^{\lfloor \frac{i}{m} \rfloor} \cdot \|A^{i-m\lfloor \frac{i}{m} \rfloor}\|_{\mathbb{V}} \leq 2^{-\lfloor \frac{i}{m} \rfloor} = 2^{1-i/m}.$$

Applying this inequality with $i = t - k$ and substituting into equation (92), we have that:

$$\left\| \left((1-\alpha)I + \alpha A \right)^t \right\|_{\mathbb{V}} \leq \sum_{k=0}^{t} \binom{t}{k} (1-\alpha)^k \alpha^{t-k} \cdot 2^{1-\frac{t-k}{m}} \leq 2 \left(1 - \alpha + \alpha \cdot \left(1 - \tfrac{1}{2m} \right) \right)^t = 2 \left(1 - \tfrac{\alpha}{2m} \right)^t.$$

C.2   Proof of Corollary 3

In this section, we prove the stated claim with the higher-order term defined as

$$\mathcal{H}_n^\diamond(\delta, \alpha) = cm\mathcal{D} \left\{ L\mathcal{J}_2(\mathbb{B}^*, \rho_n) \log(\tfrac{n}{\delta})\sqrt{\tfrac{\alpha m}{n}} + \tfrac{1}{n}\sqrt{\tfrac{m}{\alpha}} \right\} \left\{ \overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})} \right\}$$
$$+ cmb_*\mathcal{D} \left\{ \tfrac{1}{n} + \tfrac{\alpha L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{n}} \log(\tfrac{n}{\delta}) \right\} \left\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta}) \right\}$$

52

Recall that by Theorem 3 and a union bound, for the restarting procedure described in Appendix A, the event $\mathcal{E}_n^{(\theta)}(r_\theta^*) \cap \mathcal{E}_n^{(v)}(r_v^*)$ occurs with probability $1 - \delta$, for the function pair (r_θ^*, r_v^*) given by

$$r_v^*(t) := c\Big[\tfrac{1}{t}\sqrt{\tfrac{m}{\alpha}}\Big(\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big) + \tfrac{b_*}{t}\Big\{\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\Big\}\Big] \tag{93a}$$

$$r_\theta^*(t) := \tfrac{c}{\sqrt{t}}\Big\{(\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\} + cb_*\Big\{\tfrac{1}{t} + \tfrac{\alpha L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{t}}\log^3(\tfrac{n}{\delta})\Big\}\Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta})\Big\}. \tag{93b}$$

Since $\mathbf{h}$ is an affine operator, we have the decomposition

$$\big\|\theta_n - \theta^*\big\|_C \le \big\|(I - A)^{-1}v_{n+1}\big\|_C + \big\|(I - A)^{-1}z_{n+1}\big\|_C.$$

By the operator norm bound (28) and the bound (93a) on the norm $\|v_t\|$, we have

$$\big\|(I - A)^{-1}v_{t+1}\big\|_C \le c'\mathcal{D}m\Big[\tfrac{1}{t}\sqrt{\tfrac{m}{\alpha}}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\} + \tfrac{b_*}{t}\big\{\log(\tfrac{n}{\delta}) + \mathcal{J}_1(\Gamma, \rho_n)\big\}\Big].$$

For the term $\|(I - A)^{-1}z_{n+1}\|$, we consider the class of test functions $\mathcal{S} := \big\{(I - A^*)^{-1}u \mid u \in C\big\}$. Invoking Lemma 5 with $\mu = (1 - \gamma)$ yields that $\big\|(I - A^{-1})z_n\big\|_C$ is at most

$$\tfrac{c}{\sqrt{n}}\Big\{\mathbb{E}\big[\big\|(I - A)^{-1}W\big\|_C\big] + \Big(\sup_{u \in C} \mathbb{E}\big[\langle u,\, (I - A)W\rangle^2\big]\log(\tfrac{1}{\delta})\Big)^{1/2}\Big\} + c\tfrac{\mathcal{D}mb_*}{n}\Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta})\Big\}$$

$$+ \tfrac{c\mathcal{D}mL}{t}\Big\{\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \Big(\log(\tfrac{1}{\delta})\Big)^{1/2}\Big\}\Big\{\alpha\big(\sum_{s=B_0}^{n-1} s^2 r_v^*(s)^2\big)^{1/2} + m\big(\sum_{s=1}^{n-1} r_\theta^*(s)^2\big)^{1/2}\Big\},$$

with probability at least $1 - \delta$. Combining above results, some algebra yields that

$$\big\|\theta_n - \theta^*\big\|_C \le \tfrac{c}{\sqrt{n}}\Big\{\mathbb{E}\big[\big\|(I - A)^{-1}W\big\|_C\big] + \Big(\sup_{u \in C} \mathbb{E}\big[\langle u,\, (I - A)W\rangle^2\big]\log(\tfrac{1}{\delta})\Big)^{1/2}\Big\}$$

$$+ cm\mathcal{D}\Big\{L\mathcal{J}_2(\mathbb{B}^*, \rho_n)\log(\tfrac{n}{\delta})\sqrt{\tfrac{\alpha m}{n}} + \tfrac{1}{n}\sqrt{\tfrac{m}{\alpha}}\Big\}\Big\{\overline{W} + \nu\sqrt{\log(\tfrac{n}{\delta})}\Big\}$$

$$+ cmb_*\mathcal{D}\Big\{\tfrac{1}{n} + \tfrac{\alpha L\mathcal{J}_2(\mathbb{B}^*, \rho_n)}{\sqrt{n}}\log(\tfrac{n}{\delta})\Big\}\Big\{\mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{n}{\delta})\Big\},$$

with probability at least $1 - \delta$. This completes the proof of Corollary 3.

D   Two-player zero-sum Markov games

In this section, we explore the consequences of our general theory for two-player zero-sum Markov games. This class of problems results from a marriage between MDPs and two player zero-sum games: it is used the model two agents who play multiple rounds of a zero-sum game, and each has the goal to maximize their expected long-term reward. Markov games are characterized by a six-tuple $\{\mathcal{X}, \mathcal{U}_1, \mathcal{U}_2, \mathbf{P}, r, \gamma\}$. Let $\mathcal{X}$ denote the state space, and let $\mathcal{U}_1$ and $\mathcal{U}_2$ denote the action sets for players one and two, respectively. Here we focus on games with finite state and action space, i.e., $|\mathcal{X} \times \mathcal{U}_1 \times \mathcal{U}_2| < \infty$.

The probability transition kernel $\{\mathbf{P}_{u_1, u_2}(x' \mid x) \mid (x, u_1, u_2) \in \mathcal{X} \times \mathcal{U}_1 \times \mathcal{U}_2\}$, encodes the transition to the next state given the actions of the players. In particular, the scalar $\mathbf{P}_u(x' \mid x)$ denotes the probability of transition to the state x', when at state x player 1 takes the action u_1 and player 2 takes the action $\mathcal{U}_2$. The MDP is equipped with a reward function

$r : \mathcal{X} \times \mathcal{U}_1 \times \mathcal{U}_2 \mapsto \mathbb{R}$ such that the scalar $r(x, u_1, u_2)$ denotes the cost received at state x when player 1 takes the action u_1 and player 2 takes the action u_2. Finally, the scalar $\gamma \in (0, 1)$ is a parameter reflecting the discounting of future rewards.

For each player $i \in \{1, 2\}$, a stationary policy π_i is a mapping $\mathcal{X} \to \mathcal{P}(\mathcal{U}_i)$, where $\mathcal{P}(\mathcal{U}_i)$ denotes the set of probability distributions over the finite action set $\mathcal{U}_i$. In other words, the actions taken by the players can be random, and for any state $x \in \mathcal{X}$, the distribution $\pi_i(\cdot \mid x)$ is a probability distribution on the set of actions $\mathcal{U}_i$ to be taken by player i. We use Π_1 and Π_2 to denote the set of all policies for players 1 and 2, respectively.

Assuming player 1 is following policy π_1, and player 2 is following policy π_2, the value $V(\cdot \mid \pi_1, \pi_2) : \mathbb{R}^{|\mathcal{X}|} \mapsto \mathbb{R}$ of player 1 is defined as the expected sum of discounted rewards in an infinite sample path:

$$V(x \mid \pi_1, \pi_2) = \mathbb{E}\Big[\sum_{k=1}^{\infty} r(x_k, u_{1k}, u_{2k} \mid x_0 = x) \Big], \quad \text{where } u_{1k} \sim \pi_1(\cdot \mid x_k) \text{ and } u_{2k} \sim \pi_2(\cdot \mid x_k). \tag{94}$$

Given that the game is zero-sum, the reward for player 2 with initial state x is $-V(x \mid \pi_1, \pi_2)$. Players 1 and 2 want to choose their policies π_1 and π_2 that maximize their respective reward for all values of initial state x.

Nash equilibrium: A natural notion of equilibrium in two-player zero-sum Markov games is the Nash equilibrium. A policy pair $(\pi_1^\star, \pi_2^\star)$ is called a *Nash equilibrium* if for all initial states $x \in \mathcal{X}$

$$V(x \mid \pi_1^\star, \pi_2^\star) \geq V(x \mid \pi_1, \pi_2^\star) \quad \text{for all policies} \quad \pi_1 \in \Pi_1, \quad \text{and}$$
$$-V(x \mid \pi_1^\star, \pi_2^\star) \geq -V(x \mid \pi_1^\star, \pi_2) \quad \text{for all policies} \quad \pi_2 \in \Pi_2. \tag{95}$$

In words, the policy $\pi_1^\star$ is the best response for player 1 assuming player 2 is playing policy $\pi_2^\star$, and the policy $\pi_2^\star$ is the best response for player 2 assuming player 1 is playing policy $\pi_1^\star$. Thus, neither player has any incentive to deviate from the policy pair $(\pi_1^\star, \pi_2^\star)$. In two-player zero-sum Markov games, a Nash equilibrium always exists, and it is equivalent to the minimax solution [Per+15; Pat97]. Concretely, there exist policies $(\pi_1^\star, \pi_2^\star)$ such that

$$V^\star(x) = V(x \mid \pi_1^\star, \pi_2^\star) = \min_{\pi_1} \max_{\pi_2} V(x \mid \pi_1, \pi_2) = \max_{\pi_1} \min_{\pi_2} V(x \mid \pi_1, \pi_2) \text{ for all } x \in \mathcal{X}. \tag{96}$$

The function $V^\star$ is known as the *value* of the game.

D.1 Q-function and the Bellman fixed-point equation

One method for finding a pair of policies $(\pi_1^\star, \pi_2^\star)$ that achieves the equilibrium (96) is by computing the optimal state-action value functions or the optimal Q-function $\theta^\star$. It is known [Pat97; Per+15] to be the fixed point of the Bellman operator

$$\mathbf{h}(\theta)(x, u_1, u_2) = c(x, u_1, u_2)$$
$$+ \gamma \cdot \sum_{x' \in \mathcal{X}} \mathbf{P}_{u_1, u_2}(x' \mid x) \max_{\pi_1} \min_{\pi_2} \sum_{u_1', u_2'} \pi_1(u_1' \mid x') \cdot \pi_2(u_2' \mid x') \cdot \theta(x', u_1', u_2'). \tag{97}$$

Notably, when the number of states and actions are finite, the minimax problem on the right-hand side of equation (97) can be computed by solving the two-player zero-sum matrix game

with the payoff matrix $\{\theta(x', u_1, u_2) \mid u_1 \in \mathcal{U}_1, u_2 \in \mathcal{U}_2\}$. Finally, for Markov games with finite state and action spaces, the Q-function θ can be conveniently represented as an element of $\mathbf{R}^{|\mathcal{X}| \times |\mathcal{U}_1| \times |\mathcal{U}_2|}$, and the Bellman operator $\mathbf{h}$ is an operator on $\mathbb{R}^{|\mathcal{X}| \times |\mathcal{U}_1| \times |\mathcal{U}_2|}$.

A simple calculation yields that the Bellman operator is γ-contractive in the ℓ_∞-norm [Pat97; Per+15], and as a result, the optimal Q-function is the unique fixed point of the operator $\mathbf{h}$. We can thus apply our general Banach space theory to derive bounds on the `ROOT-SA` procedure.

D.2 The generative model and empirical Bellman operator

We analyze the behavior of the `ROOT-SA` algorithm under a stochastic oracle known as the generative model. A sample from this model consists of a pair of real-valued tensors $(\mathbf{Z}, R)$, each with dimensions $|\mathcal{X}| \times |\mathcal{U}_1| \times |\mathcal{U}_2|$. For each triple (x, u_1, u_2), the entry $\mathbf{Z}_{(}x, u_1, u_2)$ is drawn according to the transition kernel $\mathbf{P}_{u_1, u_2}(\cdot \mid x)$, whereas the entry $R(x, u_1, u_2)$ is a zero-mean random variable with mean $r(x, u_1, u_2)$, corresponding to a noisy observation of the reward function. The transition and reward samples across entries of the tensors are independently sampled, and we assume that the rewards are bounded in absolute value by $r_{\max}$.

Given a sample $(\mathbf{Z}, R)$ from our observation model, we can define the single-sample empirical Bellman operator

$$\mathbf{H}(\theta)(x, u_1, u_2) := R(x, u_1, u_2)$$
$$+ \sum_{x' \in \mathcal{X}} \mathbf{Z}_{u_1, u_2}\left(x' \mid x\right) \max_{\pi_1} \min_{\pi_2} \sum_{u_1', u_2'} \pi_1(u_1' \mid x') \cdot \pi_2(u_2' \mid x') \cdot \theta(x', u_1', u_2'), \quad (98)$$

where we have introduced the notation $\mathbf{Z}_{u_1, u_2}\left(x' \mid x\right) := \mathbf{1}_{\mathbf{Z}(x, u_1, u_2) = x'}$. With these definitions in hand, we are now ready to state our guarantees for two-player zero-sum Markov games.

D.3 Guarantees for two-player zero-sum Markov games

Let W be a zero-mean Gaussian random vector with covariance $\mathrm{cov}(\mathbf{H}(\theta^*) - \theta^*)$, and define

$$\overline{\mathcal{W}} = \mathbb{E}[\|W\|_\infty], \qquad \nu^2 := \sup_{x \in \mathcal{X}, u_1 \in \mathcal{U}_1, u_2 \in \mathcal{U}_2} \mathbb{E}[W_{x, u_1, u_2}^2], \quad \text{and} \quad b_* := r_{\max} + \|\theta^*\|_\infty. \quad (99)$$

For a given failure probability $\delta \in (0, 1)$, our result applies to the algorithm with parameters

$$\alpha = c_1 \left\{ \sqrt{n \log |\mathcal{X} \times \mathcal{U}_1 \times \mathcal{U}_2|} \cdot \log(\tfrac{n}{\delta}) \right\}^{-1}, \quad \text{and} \quad B_0 = \tfrac{c_2}{(1-\gamma)^2 \alpha} \log(\tfrac{n}{\delta}). \quad (100\text{a})$$

We also choose the initialization θ_0 and the number of restarts R such that

$$\log\left(\tfrac{\|\theta_0 - \mathbf{h}(\theta_0)\| \sqrt{n}}{\mathcal{W}}\right) \leq c_0 \log n \qquad \text{and} \qquad R \geq 2c_0 \log n \quad (100\text{b})$$

for appropriate universal constants c_0, c_1 and c_2. With this setup, a direct application of Theorem 1 yields the following:

Corollary 7. *Given a sample size n such that $\frac{n}{\log n} \geq \frac{c' \log(|\mathcal{X}| \cdot |\mathcal{U}_1| \cdot |\mathcal{U}_2|)}{(1-\gamma)^4} \log(\frac{1}{\delta})$, running Algorithm 1 with the tuning parameter choices (100) yields an estimate θ_n such that*

$$\|\mathbf{h}(\theta_n) - \theta_n\|_\infty \leq \tfrac{c}{\sqrt{n}} \cdot \left\{ \overline{\mathcal{W}} + \nu \sqrt{\log(\tfrac{1}{\delta})} \right\} + \tfrac{cb_*}{1-\gamma} \cdot \tfrac{\log(|\mathcal{X}| \cdot |\mathcal{U}_1| \cdot |\mathcal{U}_2|)}{n} \log^2(\tfrac{n}{\delta}).$$

with probability at least $1 - \delta$.

Note that the bound in Corollary (7) depends on the size of state-action space $|\mathcal{X}|\cdot|\mathcal{U}_1|\cdot|\mathcal{U}_2|$ only poly-logarithmically. Moreover, one can obtain an upper bound on the estimation error $\|\theta_n - \theta^*\|_\infty$ using the bound (14).

A special case of interest is when the set of actions for player two is a singleton, i.e., $|\mathcal{U}_2| = 1$. Observe that in this case the optimal state-action value estimation problem for the two-player zero-sum Markov game reduces to the optimal value estimation problem of an appropriate MDP in the discounted setting [Ber19; WD92; Wai19c]. In Appendix G.1, we show that the Bellman operator associated with the optimal value estimation problem of an MDP in the discounted setting satisfies the local linearity assumption (A4). Consequently, an argument similar to Corollary 4 yields an upper bound on the estimation error $\|\theta_n - \theta^*\|_\infty$ which matches the instance dependent lower bound (up to logarithmic terms) from the paper [Kha+21] for large n.[5] Finally, it is an important direction of future work to investigate whether the local linearity assumption (A4) holds when $|\mathcal{U}_2| > 1$.

E Some Concentration Inequalities in Banach Spaces

Our analysis makes use of some concentration inequalities for Banach-space-valued random variables, which we state and prove here.

E.1 Statement of the results

We begin with a bound for a sequence $\{X_i\}_{i=1}^n$ of i.i.d. zero-mean random elements. Our bound involves a zero-mean Gaussian random variable W in $\mathbb{V}$ such that

$$\mathbb{E}\left[\langle W,\, y\rangle \cdot \langle W,\, z\rangle\right] = \mathbb{E}\left[\langle X_1,\, y\rangle \cdot \langle X_1,\, z\rangle\right] \qquad \text{for all } y, z \in \mathbb{V}^*.$$

Lemma 7. *Let $\{X_i\}_{i=1}^n$ be independent zero-mean random elements taking values in $\Omega \subseteq \mathbb{V}$ with $\|X_i\| \leq 1$ almost surely for each $i = 1, 2, \cdots, n$. Then there exists a universal constant $c > 0$ such that for any $\delta \in (0,1)$ and any bounded symmetric convex set $\mathcal{S} \subseteq \Gamma$, we have*

$$\frac{1}{n}\sup_{u\in\mathcal{S}}\langle u,\, \sum_{i=1}^n X_i\rangle \leq \frac{c}{\sqrt{n}}\left\{\mathbb{E}\left[\sup_{u\in\mathcal{S}}\langle u,\, W\rangle\right] + \sqrt{\sup_{u\in\mathcal{S}}\mathbb{E}[\langle u,\, W\rangle^2]\cdot\log(\tfrac{1}{\delta})}\right\} + \frac{c}{n}\left\{\log(\tfrac{1}{\delta}) + \mathcal{J}_1(\mathcal{S}, \rho_n)\right\},$$

(101)

with probability at least $1 - \delta$.

See Appendix E.2 for the proof of this claim.

We next state a bound for the martingale case:

Lemma 8. *Let $\{X_t\}_{t=1}^n$ be a martingale in $\mathbb{V}$ adapted to the filtration $\{\mathcal{F}_t\}_{t=1}^n$. Assume that there exists a deterministic sequence $\{b_t\}_{t=1}^n$ such that $b_t \geq \frac{1}{n}$ and $\|X_t\| \leq b_t$ almost surely for each $t = 1, 2, \cdots, n$. Then there exists a universal constant $c > 0$ such that for any $\delta \in (0,1)$*

$$\|\sum_{i=1}^n X_i\| \leq c\left(\mathcal{J}_2(\Gamma, \rho_n) + \sqrt{\log(1/\delta)}\right) \cdot \sqrt{\sum_{i=1}^n b_i^2},$$

(102)

with probability at least $1 - \delta$.

---

[5]The sample size requirement for achieving the lower bound [Kha+21] may depend on the gap between the value of optimal and sub-optimal actions.

See Appendix E.3 for the proof of this claim.

E.2 Proof of Lemma 7

Our proof is based on a combination of Talagrand's concentration inequality [Tal96], the generic chaining [Tal06] and a functional Bernstein inequality [Wai19a]. The left-hand-side of the desired inequality is the supremum of an empirical process. Define the associated Rademacher complexity $\mathcal{R}_n(\mathcal{S}) := \frac{1}{n}\mathbb{E}[\sup_{y \in \mathcal{S}} R_n(y)]$, where $R_n(y) := \frac{1}{n}\sum_{i=1}^n \zeta_i \langle y, \, X_i \rangle$ with $\{\zeta_i\}_{i=1}^n$ an i.i.d. sequence of Rademacher random variables. The expectation is taken over the randomness of both the Rademacher sequence $\{\zeta_i\}_{i=1}^n$ and the random elements $(X_i)_{i=1}^n$.

Our first lemma is a type of functional Bernstein inequality; it bounds the supremum of the empirical process by the Rademacher complexity and some additional deviation terms:

Lemma 9. *Under the assumptions of Lemma 7, we have*

$$\frac{1}{n}\sup_{u \in \mathcal{S}}\langle u, \, \sum_{i=1}^n X_i \rangle \le 3 \cdot \mathcal{R}_n(\mathcal{S}) + 8\sqrt{\sup_{u \in \mathcal{S}}\langle u, \, W \rangle^2 \cdot \frac{\log(\frac{1}{\delta})}{n}} + c \cdot \frac{\log(\frac{1}{\delta})}{n},$$

with probability at least $1 - \delta$.

See Section E.2.1 for the proof of this claim.

We now use this auxiliary claim to complete the proof of Lemma 7. It suffices to upper bound the Rademacher complexity $\mathcal{R}_n(\mathcal{S})$. We define the pseudometrics

$$\rho_*(x, y) := \sqrt{\mathbb{E}\left[\langle x - y, \, X_1 \rangle^2\right]} \quad \text{and} \quad \rho_n(x, y) := \sup_{e \in \Omega \cap \mathbb{B}} \langle x - y, \, e \rangle, \qquad \text{for all } x, y \in \mathbb{V}^\star.$$

Recalling that $R_n(y) = \frac{1}{n}\sum_{i=1}^n \zeta_i \langle y, \, X_i \rangle$, applying Bernstein's inequality yields

$$\mathbb{P}\big[|R_n(y) - R_n(z)| > t\big] \le 2\exp\left\{-\min\left(\frac{n\alpha^2}{2\rho_*(y,z)^2}, \frac{n\alpha}{\rho_n(y,z)}\right)\right\} \qquad \text{for any } \alpha > 0.$$

For $q \ge 1$, we let γ_q denote the q^{th}-order generic chaining functional of Talagrand. With this notation, we have

$$\mathcal{R}_n(\mathcal{S}) = \mathbb{E}\left[\sup_{y \in \mathcal{S}} R_n(y)\right] \overset{(i)}{\le} \frac{c}{\sqrt{n}} \cdot \gamma_2(\mathcal{S}, \rho_\star) + \frac{1}{n}\gamma_1(\mathcal{S}, \rho_n)$$

$$\overset{(ii)}{\le} \frac{c}{\sqrt{n}} \cdot \mathbb{E}\left[\sup_{u \in \mathcal{S}}\langle u, \, W \rangle\right] + \frac{1}{n}\mathcal{J}_1(\mathcal{S}, \rho_n).$$

Here step (i) follows from the generic chaining theorem (see Theorem 1.2.7 [Tal06]). In step (ii), we bound the first term using the generic chaining lower bound (see Theorem 2.1.1 [Tal06]) and bound the second term using the fact that γ_1 functional is upper bounded by the Dudley entropy integral of order 1. This completes the proof of Lemma 7. It remains to prove Lemma 9.

E.2.1   Proof of Lemma 9

The proof of this lemma is based on Talagrand's concentration inequality for the suprema of empirical process [Tal96] and a symmetrization argument. Define the random variance $\widehat{\sigma}^2 := \frac{1}{n}\sup_{y\in\mathcal{S}}\sum_{i=1}^n \langle y, X_i\rangle^2$. Since the random variables X_i are bounded and $\mathcal{S}\subseteq\Gamma$, we have $|\sup_{y\in\mathcal{S}}\langle y, X_i\rangle| \leq 1$. Invoking Talagrand's concentration inequality [Tal96] yields the tail bound

$$\mathbb{P}\Big[\sup_{u\in\mathcal{S}}\langle u, \overline{X}_n\rangle \geq \mathbb{E}\big[\sup_{u\in\mathcal{S}}\langle u, \overline{X}_n\rangle\big] + \alpha\Big] \leq \exp\Big\{\tfrac{-n\alpha^2}{56\mathbb{E}[\widehat{\sigma}^2]+4\alpha}\Big\}, \quad \text{valid for all } \alpha > 0.$$

Consequently, for any $\delta \in (0,1)$, we have

$$\sup_{u\in\mathcal{S}}\langle u, \overline{X}_n\rangle \leq \mathbb{E}\big[\sup_{u\in\mathcal{S}}\langle u, \overline{X}_n\rangle\big] + 8\sqrt{\tfrac{\log(\frac{1}{\delta})}{n}\mathbb{E}[\widehat{\sigma}^2]} + 4\cdot\tfrac{\log(\frac{1}{\delta})}{n}$$

with probability at least $1 - \delta$.

It remains to upper bound the expected supremum $\mathbb{E}\big[\sup_{u\in\mathcal{S}}\langle u, \overline{X}_n\rangle\big]$ and the variance term $\widehat{\sigma}^2$. By a standard symmetrization argument, we have

$$\mathbb{E}\Big[\sup_{u\in\mathcal{S}}\langle u, \overline{X}_n\rangle\Big] \leq \tfrac{2}{n}\mathbb{E}\Big[\sup_{u\in\mathcal{S}}\sum_{i=1}^n \zeta_i\langle u, X_i\rangle\Big] = 2\mathcal{R}_n(\mathcal{S}),$$

Moving onto the bound on $\widehat{\sigma}^2$, we have

$$\widehat{\sigma}^2 \leq \tfrac{1}{n}\sup_{y\in\mathcal{S}}\sum_{i=1}^n \Big\{\langle y, X_i\rangle^2 - \mathbb{E}\big[\langle y, X_i\rangle^2\big]\Big\} + \tfrac{1}{n}\sup_{y\in\mathcal{S}}\sum_{i=1}^n \mathbb{E}\big[\langle y, X_i\rangle^2\big] = Z_n + \sup_{y\in\mathcal{S}}\mathbb{E}\big[\langle y, X_i\rangle^2\big],$$

where $Z_n := \frac{1}{n}\sup_{y\in\mathcal{S}}\sum_{i=1}^n \big(\langle y, X_i\rangle^2 - \mathbb{E}\big[\langle y, X_i\rangle^2\big]\big)$. Note that each term $|\langle y, X_i\rangle|$ is almost surely bounded by 1, and the map $a\mapsto a^2$ is 2-Lipschitz over the interval $[-1, 1]$. Consequently, letting $\{\zeta_i\}_{i=1}^n$ denote an i.i.d. sequence of Rademacher variables, we have

$$\mathbb{E}[Z_n] \overset{(i)}{\leq} \tfrac{2}{n}\mathbb{E}\Big[\sup_{y\in\mathcal{S}}\sum_{i=1}^n \zeta_i\langle y, X_i\rangle^2\Big] \overset{(ii)}{\leq} \tfrac{4}{n}\cdot\mathbb{E}\Big[\sup_{y\in\mathcal{S}}\sum_{i=1}^n \zeta_i\langle y, X_i\rangle\Big]$$

$$\overset{(iii)}{\leq} \tfrac{n}{64\log(\frac{1}{\delta})}\mathcal{R}_n^2(\mathcal{S}) + \tfrac{128}{n}\log(\tfrac{1}{\delta}),$$

where step (i) follows from a symmetrization argument; step (ii) follows from the Ledoux–Talagrand contraction; and step (iii) follows from the Cauchy–Schwarz inequality. Overall, we have

$$8\sqrt{\tfrac{\log(\frac{1}{\delta})}{n}\mathbb{E}[\widehat{\sigma}^2]} \leq \mathcal{R}_n + 8\sqrt{\sup_{u\in\mathcal{S}}\mathbb{E}\big[\langle u, X_1\rangle^2\big]\cdot\tfrac{\log(\frac{1}{\delta})}{n}} + \tfrac{c\log(\frac{1}{\delta})}{n}.$$

Putting together the pieces yields the bound of Lemma 9.

E.3   Proof of Lemma 8

For each vector $u \in \mathbb{V}_*$, we define the random variable $M_n(u) := \frac{1}{n} \sum_{i=1}^n \langle X_i, u \rangle$. Clearly, the sequence $\{M_t\}_{t \geq 1}$ is a scalar martingale adapted to the filtration $(\mathcal{F}_t)_{t \geq 0}$. Since $b_t \geq \frac{1}{n}$, we have

$$|\langle X_t, u \rangle| \leq b_t \cdot \sup_{x \in \mathbb{B} \cap b_t^{-1}\Omega} \langle x, u \rangle \leq b_t \rho_n(u, 0)$$

almost surely for each $t = 1, 2, \ldots$, where $\rho_n(\cdot, \cdot)$ is a pseudo-metric on the dual space $\mathbb{V}_*$ defined in (7). For any $u_1, u_2 \in \mathbb{V}^*$, the Azuma-Hoeffding inequality implies that

$$\mathbb{P}\Big[|M_n(u_1) - M_n(u_2)| \geq \alpha \Big] \leq \exp\Big\{ -\frac{n\alpha^2}{\rho_n(u_1, u_2)^2 \sum_{i=1}^n b_i^2} \Big\} \qquad \text{for each } \alpha > 0.$$

Applying the Dudley chaining tail bound (see e.g. [VH14], Theorem 5.29) to the sub-Gaussian process $\{M_n(u)\}_{u \in \Gamma}$, there exist universal constants $c, c_1 > 0$ such that

$$\mathbb{P}\left[\sup_{u \in \Gamma} M_n(u) \geq c \sqrt{\sum_{i=1}^n b_i^2} \cdot \Big(\int_0^1 \sqrt{\log N(s; \Gamma, \rho_n)} ds + t \Big) \right] \leq c e^{-c_1 t^2} \quad \text{for each } t > 0.$$

Setting $t = \sqrt{c_1^{-1} \log(1/\delta)}$ yields the claim.

F   Proofs of Auxiliary Lemmas

In this section we prove various auxiliary Lemmas that we use throughout the main proof Section 5.

F.1   Proof of Lemma 3

From the recursive relation, we have the upper bounds $\|\theta_t - \theta^*\| \leq \|\theta_{t-1} - \theta^*\| + \alpha\|v_t\|$, as well as

$$\|v_t\| \leq \frac{t-1}{t}\|v_{t-1}\| + \frac{1}{t}\Big\{ \|\theta_{t-1} - \theta_{t-2}\| + \|\mathbf{H}_t(\theta_{t-1}) - \mathbf{H}_t(\theta_{t-2})\| \Big\} + \frac{1}{t}\|\mathbf{H}_t(\theta_{t-1}) - \theta_{t-1}\|$$

$$\leq \Big\{ 1 + \frac{\alpha(L+1)}{t} \Big\}\|v_{t-1}\| + \frac{2L}{t}\|\theta_{t-1} - \theta^*\| + \frac{1}{t}b_*.$$

Putting these two inequalities together yields the vector-based recursion

$$\begin{bmatrix} \|\theta_t - \theta^*\| \\ \|v_t\| \end{bmatrix} \leq \begin{bmatrix} 1 & \alpha \\ 0 & 1 \end{bmatrix} \cdot \left(\begin{bmatrix} 1 & 0 \\ \frac{2L}{t} & 1 + \frac{\alpha(L+1)}{t} \end{bmatrix} \begin{bmatrix} \|\theta_{t-1} - \theta^*\| \\ \|v_{t-1}\| \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1}{t}b_* \end{bmatrix} \right),$$

where the inequality is taken elementwise. Solving this vector recursion yields

$$\|\theta_t - \theta^*\| + \|v_t\| \leq e^{1+\alpha Lt} \left(b_* + \|\theta_0 - \theta^*\| \right),$$

valid for any $t \geq B_0 \geq \frac{1}{\alpha}$.

F.2   Proof of Lemma 4

By definition, we have

$$v_{B_0} = \frac{1}{B_0} \sum_{t=1}^{B_0} \big(\mathbf{H}_t(\theta_0) - \theta_0\big) = \big(\mathbf{h}(\theta_0) - \theta_0\big) + \frac{1}{B_0} \sum_{t=1}^{B_0} \varepsilon_t(\theta^*) + \frac{1}{B_0} \sum_{t=1}^{B_0} \big(\varepsilon_t(\theta_0) - \varepsilon_t(\theta^*)\big).$$

Lemma 7 guarantees that

$$\| \sum_{t=1}^{B_0} \varepsilon_t(\theta^*) \| \leq c\sqrt{B_0}\Big\{ \overline{\mathcal{W}} + \nu\sqrt{\log(\tfrac{1}{\delta})} \Big\} + c\big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \big\},$$

with probability $1 - \delta$. Moreover, for each integer $t \in [B_0]$, we have:

$$\| \varepsilon_t(\theta_0) - \varepsilon_t(\theta^*) \| \leq L \| \theta_0 - \theta^* \| \leq \tfrac{L}{1-\gamma} \| \mathbf{h}(\theta_0) - \theta_0 \|.$$

Lemma 8 implies that

$$\| \sum_{t=1}^{B_0} \big(\varepsilon_t(\theta_0) - \varepsilon_t(\theta^*)\big) \| \leq \tfrac{cL}{1-\gamma} \| \mathbf{h}(\theta_0) - \theta_0 \| \sqrt{B_0} \left\{ \mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \right\}$$

with probability at least $1 - \delta$.

By combining these bounds, we find that

$$\| v_{B_0} \| \leq \| \mathbf{h}(\theta_0) - \theta_0 \| \Big\{ 1 + \tfrac{cL}{(1-\gamma)\sqrt{B_0}} \big[\mathcal{J}_2(\mathbb{B}^*, \rho_n) + \sqrt{\log(\tfrac{1}{\delta})} \big] \Big\}$$
$$+ \tfrac{c}{\sqrt{B_0}} \big(\overline{\mathcal{W}} + \nu\sqrt{\log(\tfrac{1}{\delta})} \big) + \tfrac{c}{B_0} \Big\{ \mathcal{J}_1(\Gamma, \rho_n) + \log(\tfrac{1}{\delta}) \Big\}$$

with probability at least $1 - \delta$. Substituting the burn-in time bound (53a) yields the final claim.

G   Comments on Theorem 2

In Section G.1, we prove that the Bellman optimality operator associated with the optimal Q-function estimation problem satisfies the local linearity condition (A4). Using a similar argument, in Section G.2 we show that the Bellman fixed-point operator for the stochastic shortest path problem satisfies the local linearity condition.

G.1   Verifying local linearity for Bellman optimality operator

In this section, we verify that the local linearity assumption (A4) holds for the Bellman optimality operator for Q-learning [WD92; Sze98; Wai19c]. Consider a tabular MDP $M = (r, \mathbf{P}, \gamma)$ with state space $\mathcal{S}$ and action space $\mathcal{A}$. For any state-action pair $(x, u) \in \mathcal{S} \times \mathcal{A}$, the scalar $r(x, u)$ denotes the reward when the action u is taken at state x, and the scalar $\mathbf{P}_u(x' \mid x)$ denotes the probability of transitioning to state x' when the action u is chosen at state x.

One way to estimate an optimal policy is to calculate the optimal Q-function. Associated with a (deterministic) policy π is its Q-*function*

$$\theta^\pi(x, u) := \mathbb{E}\Big[\sum_{k=0}^{\infty} r(x_k, u_k) \mid x_0 = x, u_0 = u \Big], \quad \text{where } u_k = \pi(x_k) \quad \text{for all } k = 1, 2, \ldots.$$

The optimal Q-function is given by $\theta^\star(x,u) := \sup_{\pi \in \Pi} \theta^\pi(x,u)$, and an optimal policy can be obtained as $\pi_\star(x) = \arg\max_u \theta^\star(x,u)$.

The Bellman optimality operator $\mathbf{h}$ acts on the space of Q-functions; more precisely, its action on a given Q-function θ is given by

$$\mathbf{h}(\theta)(x,u) = r(x,u) + \gamma \sum_{x'} \mathbf{P}_u(x' \mid x) \cdot \max_{u'} \theta(x',u') \quad \text{for all } (x,u) \in \mathcal{S} \times \mathcal{A}. \tag{103}$$

By standard results [Ber12a], the operator $\mathbf{h}$ is γ-contractive in the ℓ_∞-norm, and the optimal state-action value function $\theta^\star$ is its unique fixed point.

For a given Q-function θ, the associated greedy policy π_θ is given by

$$\pi_\theta(x) = \arg\max_u \theta(x,u), \tag{104}$$

where we break any ties by taking the smallest action (in the enumeration order) that achieves the maximum. Using this greedy policy, we can define the right linear operator

$$\mathbf{P}^{\pi_\theta}\theta(x,u) = \sum_{x'} \mathbf{P}_u(x' \mid x)\theta(x',\pi_\theta(x')).$$

Let $\mathbb{B}(\theta^*,s) := \{\theta \mid \|\theta - \theta^*\|_\infty \leq s\}$ denote the ℓ_∞-ball of radius s around θ^*, and define the set

$$\mathcal{A}_s = \{\gamma \cdot \mathbf{P}^{\pi_\theta} \mid \pi_\theta \text{ is a greedy policy of } \theta \text{ with } \theta \in \mathbb{B}(\theta^\star,s)\} \tag{105}$$

of linear operators. We use $\pi_\star$ to denote the greedy policy associated with the optimal Q-function $\theta^\star$. By definition, the Q-functions θ and θ^* satisfy the fixed-point relations

$$\mathbf{h}(\theta) = r + \gamma\mathbf{P}^{\pi_\theta}\theta \qquad \text{and} \qquad \theta^\star = r + \gamma\mathbf{P}^{\pi_\star}\theta^\star.$$

Rearranging the last two equations yields

$$\mathbf{h}(\theta) - \theta = r + \gamma\mathbf{P}^{\pi_\theta}\theta - \theta = (\mathcal{I} - \gamma\mathbf{P}^{\pi_\star})(\theta^\star - \theta) + (\gamma\mathbf{P}^{\pi_\theta} - \gamma\mathbf{P}^{\pi_\star})\theta \tag{106a}$$

$$\mathbf{h}(\theta) - \theta = r + \gamma\mathbf{P}^{\pi_\theta}\theta - \theta = (\mathcal{I} - \gamma\mathbf{P}^{\pi_\theta})(\theta^\star - \theta) + (\gamma\mathbf{P}^{\pi_\theta} - \gamma\mathbf{P}^{\pi_\star})\theta^\star. \tag{106b}$$

Next we claim that

$$(\mathcal{I} - \gamma\mathbf{P}^{\pi_\star})^{-1}(\gamma\mathbf{P}^{\pi_\theta} - \gamma\mathbf{P}^{\pi_\star})\theta \overset{(a)}{\succcurlyeq} 0 \qquad \text{and} \qquad (\mathcal{I} - \gamma\mathbf{P}^{\pi_\theta})^{-1}(\gamma\mathbf{P}^{\pi_\theta} - \gamma\mathbf{P}^{\pi_\star})\theta^\star \overset{(b)}{\preccurlyeq} 0. \tag{107}$$

Indeed, since the policy π_θ is greedy for θ, we have the element-wise inequality $(\gamma\mathbf{P}^{\pi_\theta} - \gamma\mathbf{P}^{\pi_\star})\theta \succcurlyeq 0$. The matrix $(\mathcal{I} - \gamma\mathbf{P}^{\pi_\star})^{-1}$ has non-negative entries, so that element-wise inequality (a) holds. A similar argument, using the fact that $\pi_\star$ is greedy for $\theta^\star$, yields the element-wise (b).

With the last observation in hand, combining the element-wise inequalities (107) with the two expressions of the Bellman defect (106) yields

$$|\theta^\star - \theta| \preccurlyeq \max\{|(\mathcal{I} - \gamma\mathbf{P}^{\pi_\theta})^{-1}(\mathbf{h}(\theta) - \theta)|, |(\mathcal{I} - \gamma\mathbf{P}^{\pi_\star})^{-1}(\mathbf{h}(\theta) - \theta)|\}.$$

Finally, note that the operator $\gamma\mathbf{P}^{\pi_\theta} \in \mathcal{A}_s$ for any $\theta \in \mathbb{B}(\theta^\star,s)$. Putting together the pieces we conclude that for all $\theta \in \mathbb{B}(\theta^*,s)$

$$\|\theta - \theta^*\|_\infty \leq \sup_{A \in \mathcal{A}_s} \|(\mathcal{I} - A)^{-1}(\mathbf{h}(\theta) - \theta)\|$$

Thus, we deduce that the local linearity condition (A4) is satisfied for the Bellman optimality operator $\mathbf{h}$ from equation (103) with $\|\cdot\|_C = \|\cdot\|_\infty$.

G.2 Verifying local linearity for the SSP operator

Recall from Section 4.1 the definition of a stochastic shortest path (SSP) problem $(r, \mathbf{P})$ with optimal-Q value $\theta^\star$. For a given Q-function θ, consider the greedy policy $\Pi_\theta(x) = \arg\min_u \theta(x, u)$. We can use it to define the right linear operator $\mathbf{P}^{\pi_\theta}\theta(x, u) = \sum_{x'} \mathbf{P}_u(x' \mid x)\theta(x', \pi_\theta(x'))$. Letting $\mathbb{B}(\theta^*, s) := \{\theta \mid \|\theta - \theta^*\| \leq s\}$ denote the ℓ_∞-ball of radius s around θ^*, we define the set

$$\mathcal{A}_s = \{\mathbf{P}^{\pi_\theta} \mid \pi_\theta \text{ is a greedy policy of } \theta \text{ with } \theta \in \mathbb{B}(\theta^\star, s)\} \tag{108}$$

of linear operators. We use $\pi_\star$ to denote the greedy policy associated with the optimal Q-function $\theta^\star$.

With this setup in hand, following the same argument as Section G.1, the local linearity assumption (A4) for the Bellman operator (33) can be verified with the set $\mathcal{A}_s$ of local linear operators defined in equation (108), and with $\|\cdot\|_C = \|\cdot\|_\infty$.