



HAL
open science

Score-Aware Policy-Gradient Methods and Performance Guarantees using Local Lyapunov Conditions

Céline Comte, Matthieu Jonckheere, Jaron Sanders, Albert Senen-Cerda

► **To cite this version:**

Céline Comte, Matthieu Jonckheere, Jaron Sanders, Albert Senen-Cerda. Score-Aware Policy-Gradient Methods and Performance Guarantees using Local Lyapunov Conditions. 2024. hal-04329790v2

HAL Id: hal-04329790

<https://hal.science/hal-04329790v2>

Preprint submitted on 14 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Score-Aware Policy-Gradient Methods and Performance Guarantees using Local Lyapunov Conditions

Applications to Product-Form Stochastic Networks and Queueing Systems

Céline Comte¹, Matthieu Jonckheere¹, Jaron Sanders², and Albert Senen-Cerda^{1,3}

¹CNRS and LAAS, France

²Eindhoven University of Technology, The Netherlands

³IRIT and Université Toulouse III Paul Sabatier, France

`celine.comte@cnrs.fr`, `matthieu.jonckheere@laas.fr`, `jaron.sanders@tue.nl`, `albert.senen-cerda@irit.fr`

Abstract

In this paper, we introduce a policy-gradient method for model-based reinforcement learning (RL) that exploits a type of stationary distributions commonly obtained from Markov decision processes (MDPs) in stochastic networks, queueing systems, and statistical mechanics.

Specifically, when the stationary distribution of the MDP belongs to an exponential family that is parametrized by policy parameters, we can improve existing policy gradient methods for average-reward RL. Our key identification is a family of gradient estimators, called score-aware gradient estimators (SAGEs), that enable policy gradient estimation without relying on value-function approximation in the aforementioned setting. This contrasts with other common policy-gradient algorithms such as actor-critic methods.

We first show that policy-gradient with SAGE locally converges, including in cases when the objective function is nonconvex, presents multiple maximizers, and the state space of the MDP is not finite. Under appropriate assumptions such as starting sufficiently close to a maximizer, the policy under stochastic gradient ascent (SGA) with SAGE has an overwhelming probability of converging to the associated optimal policy. Other key assumptions are that a local Lyapunov function exists, and a nondegeneracy property of the Hessian of the objective function holds locally around a maximizer.

Furthermore, we conduct a numerical comparison between a SAGE-based policy-gradient method and an actor-critic method. We specifically focus on several examples inspired from stochastic networks, queueing systems, and models derived from statistical physics, where parametrizable exponential families are commonplace. Our results demonstrate that a SAGE-based method finds close-to-optimal policies faster than an actor-critic method.

Keywords: reinforcement learning, policy-gradient method, exponential families, product-form stationary distribution, stochastic approximation

1 Introduction

Reinforcement learning (RL) has become the primary tool for optimizing controls in uncertain environments. Model-free RL, in particular, can be used to solve generic Markov decision processes (MDPs) with unknown dynamics with an agent that learns to maximize a reward incurred upon acting on the environment. In stochastic systems, examples of possible applications of RL can

be found in stochastic networks, queueing systems, and particle systems, where an optimal policy is desirable. For example, a policy yielding a good routing policy, an efficient scheduling, or an annealing schedule to reach a desired state.

As stochastic systems expand in size and complexity, however, the RL agent must deal with large state and action spaces. This leads to several computational concerns, namely, the combinatorial explosion of action choices, the computationally intensive exploration and evaluation of policies [30], and a larger complexity of the optimization landscape.

One way to circumvent issues pertaining to large state spaces and/or nonconvex objective functions is to include features of the underlying MDP in the RL algorithm. If the model class of the environment is known, a model-based RL approach estimates first an approximate model of the environment in the class that can later be used to solve an MDP describing its approximate dynamics. This approach is common in queueing networks [25, 3]. Nevertheless, solving an approximate MDP adds a computational burden if the number of states is large.

Policy-gradient methods are learning algorithms that instead directly optimize policy parameters through stochastic gradient ascent (SGA) [35]. These methods have gained attention and popularity due to their perceived ability to handle large state and action spaces in model-free settings [9, 20]. Policy-gradient methods rely on the estimation of value functions, which encode reward-weighted representations of the underlying model dynamics. Computing such functions, however, is challenging in high-dimensional settings and, differently to a model-based approach, key model features are initially unknown.

In this paper, we improve policy-gradient methods for some stochastic systems by incorporating model-specific information of the MDP to the gradient estimator. Specifically, we exploit the fact that long-term average behavior of such systems are described using exponential families of distributions. In the context of stochastic networks and queueing systems, this typically means that the Markov chains associated to fixed policies have a product-form stationary distribution. This structural assumption holds in various relevant scenarios, including Jackson and Whittle networks [32, Chapter 1], BCMP networks [5], and more recent models arising in datacenter scheduling and online matching [16]. By encoding this key model feature into policy-gradient methods, we aim to expand the current model-based RL techniques for control policies of stochastic systems.

Our primary contributions are the following:

- We present a new gradient estimator for policy-gradient methods that incorporates information from the stationary measure of the MDP. Under an average-reward and infinite-horizon learning setting, we namely consider policy parametrizations such that there is a known relationship between the policy on the one hand, and the MDP’s stationary distribution on the other hand. In practice, this translates to assuming that the stationary distribution forms an exponential family explicitly depending on the policy parameters. Using this structure, we define score-aware gradient estimators (SAGEs), a class of estimators that exploit the aforementioned assumption to estimate the policy gradient *without* relying on value or action–value functions.
- We show the local convergence of SAGE-based policy-gradient under broad assumptions, such as infinite state space, nonconvex objective functions, and unbounded rewards. To do so, we use an approach that employs a local Lyapunov stability assumption and does not rely on the structure of the stationary measure. Specifically, we show convergence with large probability, whenever the initial policy is close enough to an optimum and nondegeneracy condition of

the Hessian at the optimum holds. Remarkably, the local assumptions also allow for unstable policies to exist. The proof technique can also be adapted to other gradient-based methods and is of independent interest.

- We numerically evaluate the performance of SAGE-based policy-gradient on several models from stochastic networks, queueing systems, and statistical physics. Compared to an actor-critic algorithm, we observe that SAGE-based policy-gradient methods exhibit faster convergence and lower variance.

Our results suggest that exploiting model-specific information is a promising approach to improve RL algorithms, especially for stochastic networks and queueing systems. Sections 1.1 and 1.2 below describe our contributions in more details.

1.1 Score-aware gradient estimators (SAGEs)

We introduce SAGEs for MDPs following the exponential-family assumption in Section 4. These estimators leverage the structure of the stationary distribution, with the goal of reducing variance and favoring stable learning. Notably, their usage requires neither knowledge nor explicit estimation of model parameters, ensuring practical applicability. The key step of the derivation exploits information on the form of the *score* of exponential families—that is, the gradient of the logarithm of the probability mass function.

We can illustrate the working principle using a toy example on a countable state space \mathcal{S} : given a *sufficient statistic* $x : \mathcal{S} \rightarrow \mathbb{R}^n$, the associated exponential family in canonical form is the family of distributions with probability mass functions $p(\cdot|\theta) \propto \exp(\theta^\top x(\cdot))$ parametrized by $\theta \in \mathbb{R}^n$. Observe now that these distributions satisfy the relation

$$\frac{d \log(p(s|\theta))}{d\theta} = x(s) - \mathbb{E}_{S \sim p(\cdot|\theta)}[x(S)], \quad s \in \mathcal{S} \tag{1}$$

and that (1) gives an exact expression for the gradient of the score.

Now, a more general version of (1) that is also applicable beyond this toy example—see Theorem 1 below—allows us to bypass the commonly used policy-gradient theorem [35, §13.2], which ties the estimation of the gradient with that of first estimating value or action-value functions. A key aspect that SAGE practically exploits is that, in the models from queueing and statistical physics that we will study, we fully or partially know the sufficient statistic x . Furthermore, such models commonly possess an ‘effective dimension’ that is much lower than the size of state space and is reflected by the sufficient statistic. For example, in a load-balancing model we consider in the numerical section, an agnostic model-free RL algorithm would assume the number of states in \mathcal{S} to grow exponentially in the number n of servers. However, the state space latent representation is actually the job count at each server, which is efficiently encoded by the sufficient statistic with an n -dimensional vector.

1.2 Convergence of policy-gradient methods

We examine the convergence properties of the SAGE-based policy-gradient method theoretically in Section 5. Specifically, we consider the setting of policy-gradient RL with average rewards, which

consists of finding a parameter θ such that the parametric policy $\pi(\theta) = \pi(\cdot|\cdot, \theta)$ maximizes

$$J(\theta) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T R_t \right]. \quad (2)$$

Here, R_{t+1} denotes the reward that is given after choosing action A_t while being in state S_t , which happens with probability $\pi(A_t|S_t, \theta)$. As is common in episodic RL, we consider epochs, that is, time intervals where the parameter θ is fixed and a trajectory of the Markov chain is observed. For each epoch m , and under the exponential-family assumption for the stationary distribution, SAGE yields a gradient estimator H_m from a trajectory of state-action-reward tuples (S_t, A_t, R_{t+1}) sampled from a policy with Θ_m as an epoch-dependent parameter. Convergence analysis of the SAGE-based policy-gradient method aligns with ascent algorithms like SGA by considering updates at the end of epoch m with step-size $\alpha_m > 0$,

$$\Theta_{m+1} = \Theta_m + \alpha_m H_m. \quad (3)$$

Convergence analyzes for policy-gradient RL and SGA are quite standard; see Section 2. Our work specifically aligns with the framework of [14] that studies local convergence of unbiased stochastic gradient descent (SGD), that is, when the conditional estimator H_m of $\nabla J(\Theta_m)$ on the past \mathcal{F} is unbiased, which is typical in a supervised learning setting. An important part of our work consists in expanding the results of [14] to the case of Markovian data, leading to biased estimators (i.e., $\mathbb{E}[H_m|\mathcal{F}] \neq \nabla J(\Theta_m)$). In our RL setting, we handle potentially unbounded rewards and unbounded state spaces as well as the existence of unstable policies. We also assume an online application of the policy-gradient method, where restarts are impractical or costly: the last state of the prior epoch is used as the initial state for the next, distinguishing our work from typical episodic RL setups where an initial state S_0 is sampled from a predetermined distribution.

Our main result in Section 5 shows convergence of iterates in (3) to the set \mathcal{M} of maxima, assuming nondegeneracy of J on \mathcal{M} and existence of a local Lyapunov function. If SGA starts within a sufficiently small neighborhood V of a maximizer $\theta^* \in \mathcal{M}$, with appropriate epoch length and step-sizes, convergence to \mathcal{M} occurs with large probability: for any epoch $m > 0$ and $\epsilon > 0$,

$$\mathbb{P}[J(\theta^*) - J(\Theta_m) > \epsilon | \Theta_0 \in V] \leq O\left(\epsilon^{-2} m^{-\sigma-\kappa} + m^{1-\sigma/2-\kappa/2} + m^{-\kappa/2} + \frac{\alpha^2}{\ell}\right), \quad (4)$$

where the parameters $\sigma \in (2/3, 1)$, $\kappa > 0$, $\alpha \in (0, \alpha_0]$, and $\ell \in [\ell_0, \infty)$ depend on the step and batch sizes and can be tuned to make the bound in (4) arbitrarily small.

Our key assumption relies on the existence of a local Lyapunov function in the neighborhood V . Hence, we need only to assume stability of policies that are close to the optimum. This sets our work further apart from others in the RL literature, which typically require existence of a global Lyapunov function and/or finite state space. In fact, our numerical results in Section 6 show an instance where local stability suffices, highlighting the benefits of SAGE. The set \mathcal{M} of global maxima is also not required to be finite or convex, thanks to the local nondegeneracy assumption.

For large m , the bound in (4) can be made arbitrarily small by setting the initial step size α and batch size ℓ small and large, respectively. In (4), the chance that the policy escapes the set V , outside which stability cannot be guaranteed, does not vanish when $m \rightarrow \infty$; it remains as α^2/ℓ . We show that this term is inherent to the local assumptions. Specifically, for any $\beta > 0$, there are functions f such that $\mathbb{P}[f(\theta^*) - f(\Theta_m) > \epsilon | \Theta_0 \in V] > c\alpha^{2+\beta}/\ell$ for some $c > 0$. Hence, a lower

bound shows that the proof method cannot be improved without further using the global structure of H_m or J .

Denoting the total number of samples drawn at the end of epoch m by T and defining $\Theta_m = \tilde{\Theta}_T$, from (4) we obtain a sample complexity bound of our algorithm when the rewards are bounded. We namely show that for any $\epsilon > 0$, we can find appropriate step and batch sizes such that

$$\mathbb{E}[J^* - J(\tilde{\Theta}_T) \mid \tilde{\Theta}_0 \in V] = O(T^{-1/(3+\epsilon)} + \alpha^2/\ell). \quad (5)$$

For cases where the optimum is reached only as $\Theta_m \rightarrow \infty$, as with deterministic policies, we additionally show that adding a small entropy regularization term to $J(\theta)$ allows us to ensure not just that maxima are bounded but also that \mathcal{M} satisfies the nondegeneracy assumption required to show local convergence.

1.3 Numerical experiments

After the introduction and theoretical analysis of the SAGE-based policy-gradient algorithm, we finally assess its applicability in Section 6 by comparing its performance with the actor-critic algorithm on three models from queueing systems, stochastic networks, and statistical physics. Specifically, we consider an admission control problem on the M/M/1 queue, a load balancing system, and the Ising model with Glauber dynamics.

The numerical results on these examples suggest that, when applicable, SAGEs can expedite convergence towards an optimal policy (compared to actor-critic) by leveraging the structure of the stationary distribution. Furthermore, the lower variance of SAGE becomes decisive when stability is not guaranteed for all policies. Namely, we observe in an example that the SAGE-based policy-gradient method converges to a close-to-optimal policy even if some policies are unstable, provided that a stable policy is used as initialization. This behavior contrasts with actor-critic, whose output policies are not always stable. SAGE also reproduces a well-known phenomenon in annealing schedules for Ising models. Specifically, the agent momentarily increases the temperature in order to escape stable states that do not correspond to the global optimum.

2 Related works

The work in the present manuscript resides at the intersection of distinct lines of research. We therefore broadly review, relate, and position our work to other research in this section.

2.1 Gradient estimation, exponential families, and product forms

Operations on high-dimensional probability distributions, such as marginalization and inference, are numerically intractable in general. Exponential families—see Section 4.1 for a definition—are parametric sets of distributions that lead to more tractable operations and approximations while also capturing well-known probability distributions, such as probabilistic graphical models [37], popular in machine learning. In the context of stochastic networks and queueing systems, the stationary distribution of many product-form systems can be seen as forming an exponential family.

Our first contribution is related to several works on exponential families, product-form distributions, and probabilistic graphical models. Key parameters in these distributions are numerically intractable *a priori*, but can be expressed as expectations of random vectors that can be sampled

by simulation. The most basic and well-known result, which appears in Section 1.1 and will be exploited in Section 4.2, rewrites the gradient of the logarithm of the normalizing constant (a.k.a. the log-partition function) as the expectation of the model’s *sufficient statistics*. In probabilistic graphical models, this relation has been mainly used to learn a distribution that best describes a dataset via SGD [37, 21]. In stochastic networks, this relation has been applied to analyze systems with *known* parameters, for instance to predict their performance [11, 41, 6, 33, 34], to characterize their asymptotic behavior in scaling regimes [33, 34], for sensitivity analysis [11, 26], and occasionally to optimize control parameters via gradient ascent [26, 10, 33].

To the best of our knowledge, an approach similar to ours is found only in [31]. This work derives a gradient estimator and performs SGA in a class of product-form reversible networks. However, the procedure requires first estimating the stationary distribution, convergence is proven only for convex objective functions, and the focus is more on developing a distributed algorithm than on canonical RL. The algorithm in [18] is similarly noteworthy, although the focus there is on developing a distributed control algorithm specifically for wireless networks and not general product-form networks.

2.2 Stochastic gradient ascent (SGA) and policy-gradient methods

When a gradient is estimated using samples from a Markov chain, methods from Markov Chain Monte Carlo (MCMC) are commonly used [27]. In our case, we have moreover bias from being unable to restart the chain at each epoch. Convergence of biased SGD to approximate stationary points of smooth nonconvex functions—points θ such that $|\nabla J(\theta)| < \epsilon$ for some $\epsilon > 0$ —has been addressed in the literature [36, 4, 19, 12]. The asymptotic conditions for local convergence to a stationary point were first investigated in [36], where conditions for the asymptotic stochastic variance of the gradient estimator and bias were assumed (see Assumptions 2.1–2.3 in [36]). In [19], a nonasymptotic analysis of biased SGD is shown. Under Lipschitz assumptions on the transition probabilities and bounded variance of the gradient estimator H_m , in [19] it is shown that under appropriate step-sizes, for some $m^* \leq T$, $\mathbb{E}[|\nabla J(\Theta_{m^*})|^2] = O(\log(T)/\sqrt{T})$, where T is a time horizon. In [36, 19], these results are applied in a RL context. While these works demonstrate convergence to stationary points, our contribution lies in proving convergence to a maximum, albeit locally. This approach is essential for addressing scenarios with only local assumptions and potentially unstable (not positive recurrent) policies.

Finally, several recent works build on gradient domination for policy-gradient methods, addressing convexity limitations and ensuring global convergence [13, 2, 39, 22]. Notable differences to our work include our assumption of a single trajectory versus initiating the Markov chain from a predetermined distribution, as well as distinct structural assumptions on policy parametrization like natural gradients. We tackle challenges involving infinite state space and multiple maxima, aspects often overlooked in prior studies. Another unique aspect of our contribution lies in specialized gradient estimation schemes based on the exponential family assumption on the stationary distribution.

3 Problem formulation

3.1 Basic notation

The sets of nonnegative integers, positive integers, reals, and nonnegative reals are denoted by \mathbb{N} , \mathbb{N}_+ , \mathbb{R} , and $\mathbb{R}_{\geq 0}$, respectively. For a differentiable function $f : \theta \in \mathbb{R}^n \mapsto f(\theta) \in \mathbb{R}$, $\nabla f(\theta)$ denotes the gradient of f taken at $\theta \in \mathbb{R}^n$, that is, the n -dimensional column vector whose j -th component is the partial derivative of f with respect to θ_j , for $j \in \{1, 2, \dots, n\}$. If f is twice differentiable, $\text{Hess}_\theta f$ denotes the Hessian of f at θ , that is, the $n \times n$ matrix of second partial derivatives. For a differentiable vector function $f : \theta \in \mathbb{R}^n \mapsto f(\theta) = (f_1(\theta), \dots, f_d(\theta)) \in \mathbb{R}^d$, $Df(\theta)$ is the Jacobian matrix of f taken at θ , that is, the $d \times n$ matrix whose i -th row is $\nabla f_i(\theta)^\top$, for $i \in \{1, 2, \dots, d\}$. We define the operator norm of a matrix $A \in \mathbb{R}^{a \times b}$ as $|A|_{\text{op}} = \sup_{x \in \mathbb{R}^b: |x|=1} |Ax|$. We use uppercase to denote random variables and vectors, and a calligraphic font for their sets of outcomes.

3.2 Markov decision process (MDP)

We consider a Markov decision process (MDP) with countable state, action, and reward spaces \mathcal{S} , \mathcal{A} , and \mathcal{R} , respectively, and transition probability kernel $P : (s, a, r, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{R} \times \mathcal{S} \mapsto P(r, s'|s, a) \in [0, 1]$, where $P(r, s'|s, a)$ gives the conditional probability that the next reward–state pair is (r, s') given that the current state–action pair is (s, a) . With a slight abuse of notation, we introduce, for $s, s' \in \mathcal{S}$, $a \in \mathcal{A}$, and $r \in \mathcal{R}$,

$$P(r|s, a) = \sum_{s' \in \mathcal{S}} P(r, s'|s, a) \quad \text{and} \quad P(s'|s, a) = \sum_{r \in \mathcal{R}} P(r, s'|s, a).$$

All our results also generalize to absolutely continuous rewards; an example will appear in Section 6.1.

Following the framework of policy-gradient algorithms [35, Chapter 13], we assume that the agent is given a random policy parametrization $\pi : (s, \theta, a) \in \mathcal{S} \times \mathbb{R}^n \times \mathcal{A} \rightarrow \pi(a|s, \theta) \in (0, 1)$, such that $\pi(a|s, \theta)$ is the conditional probability that the next action is $a \in \mathcal{A}$ given that the current state is $s \in \mathcal{S}$ and the parameter vector $\theta \in \mathbb{R}^n$. We assume that the function $\theta \mapsto \pi(a|s, \theta)$ is differentiable for each $(s, a) \in \mathcal{S} \times \mathcal{A}$. The goal of the learning algorithm will be to find a parameter (vector) that maximizes the long-run average reward, as will be defined formally in Section 3.3.

As a concrete example, we will often consider a class of softmax policies that depend on a feature extraction map $\xi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^n$ as follows:

$$\pi(a|s, \theta) = \frac{e^{\theta^\top \xi(s, a)}}{\sum_{a' \in \mathcal{A}} e^{\theta^\top \xi(s, a')}} \quad s \in \mathcal{S}, \quad a \in \mathcal{A}, \quad (6)$$

The feature extraction map ξ may leverage prior information on the system dynamics. In queueing systems for instance, we may decide to make similar decisions in large states, as these states are typically visited rarely, and it may be beneficial to aggregate the information collected about them.

3.3 Stationary analysis and optimality criterion

Given $\theta \in \mathbb{R}^n$, if the agent applies the policy $\pi(\theta) : (s, a) \in \mathcal{S} \times \mathcal{A} \mapsto \pi(a|s, \theta)$ at every time step, the random state–action–reward sequence $((S_t, A_t, R_{t+1}), t \in \mathbb{N})$ obtained by running this policy is a Markov chain such that, for each $s, s' \in \mathcal{S}$, $a \in \mathcal{A}$, and $r \in \mathcal{R}$, we have $\mathbb{P}[A_t = a | S_t = s] = \pi(a|s, \theta)$

and $\mathbb{P}[R_{t+1} = r, S_{t+1} = s' | S_t = s, A_t = a] = P(r, s' | s, a)$. The dependency of the random variables on the parameter is left implicit to avoid cluttering notation. Leaving aside actions and rewards, the state sequence $(S_t, t \in \mathbb{N})$ also defines a Markov chain, with transition probability kernel $P(\theta) : (s, s') \in \mathcal{S} \times \mathcal{S} \mapsto P(s' | s, \theta)$ given by

$$P(s' | s, \theta) = \sum_{a \in \mathcal{A}} \pi(a | s, \theta) P(s' | s, a), \quad s, s' \in \mathcal{S}.$$

In the remainder, we will assume that Assumptions 1 and 2 below are satisfied.

Assumption 1. *There exists an open set $\Omega \subseteq \mathbb{R}^n$ such that, for each $\theta \in \Omega$, the Markov chain $(S_t, t \in \mathbb{N})$ with transition probability kernel $P(\theta)$ is irreducible and positive recurrent.*

Thanks to Assumption 1, for each $\theta \in \Omega$, the corresponding Markov chain $(S_t, t \in \mathbb{N})$ has a unique stationary distribution $p(\cdot | \theta)$. We say that a triplet (S, A, R) of random variables is a *stationary state-action-reward triplet*, and we write $(S, A, R) \sim \text{STAT}(\theta)$, if (S, A, R) follows the stationary distribution of the Markov chain $((S_t, A_t, R_{t+1}), t \in \mathbb{N})$, given by

$$\mathbb{P}[S = s, A = a, R = r] = p(s | \theta) \pi(a | s, \theta) P(r | s, a), \quad s \in \mathcal{S}, \quad a \in \mathcal{A}, \quad r \in \mathcal{R}. \quad (\text{STAT}(\theta))$$

Assumption 2. *For each $\theta \in \Omega$, the stationary state-action-reward triplet $(S, A, R) \sim \text{STAT}(\theta)$ is such that the random variables $|R|$, $|R \nabla \log p(S | \theta)|$, and $|R \nabla \log \pi(A | S, \theta)|$ have a finite expectation.*

By ergodicity [7, Theorem 4.1], the running average reward $\frac{1}{T} \sum_{t=1}^T R_t$ tends to $J(\theta)$ in (2) almost surely as T tends to infinity. $J(\theta)$ is called the *long-run average reward* and is also given by

$$J(\theta) = \mathbb{E}[R] = \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{r \in \mathcal{R}} p(s | \theta) \pi(a | s, \theta) P(r | s, a) r, \quad \theta \in \Omega. \quad (7)$$

Our end goal, further developed in Section 3.4, is to find a learning algorithm that maximizes the objective function J . For now, we only observe that the objective function $J : \theta \in \Omega \mapsto J(\theta)$ is differentiable thanks to Assumption 2, and that its gradient is given by

$$\nabla J(\theta) = \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{r \in \mathcal{R}} p(s | \theta) \pi(a | s, \theta) P(r | s, a) r (\nabla \log p(s | \theta) + \nabla \log \pi(a | s, \theta)), \quad \theta \in \Omega. \quad (8)$$

In general, computing $\nabla J(\theta)$ using (8) is challenging: (i) computing $\nabla \log p(s | \theta)$ is in itself challenging because $p(s | \theta)$ depends in a complex way on the unknown transition kernel $P(r, s' | s, a)$ and the parameter θ via the policy $\pi(\theta)$, and (ii) enumerating and thus summing over the state space \mathcal{S} is often practically infeasible (for instance, when the state space \mathcal{S} is infinite and/or high-dimensional). Our first contribution, in Section 4, is precisely a new family of estimators for the gradient (8).

3.4 Learning algorithm

In Section 3.3, we defined the objective function J by considering trajectories where the agent applied a policy $\pi(\theta)$ parametrized by a constant vector θ . Going back to a learning setting, we consider a state-action-reward sequence $((S_t, A_t, R_{t+1}), t \in \mathbb{N})$ and a parameter sequence $(\Theta_m, m \in \mathbb{N})$.

Algorithm 1 Generic policy-gradient algorithm. Examples of GRADIENT procedure, based on different estimators for the gradient ∇J , are given in Algorithms 2 and 3. All variables of Algorithm 1 are accessible within the GRADIENT procedure.

```

1: Input: • Observation times  $0 \triangleq t_0 < t_1 < t_2 < \dots$ 
           • Step size sequence  $\alpha_0, \alpha_1, \alpha_2, \dots > 0$ 
           • Positive and differentiable policy parametrization  $(s, \theta, a) \mapsto \pi(a|s, \theta)$ 
2: Initialization: Policy parameter  $\Theta_0 \in \Omega$  and initial state  $S_0 \in \mathcal{S}$ 
3: Main loop:
4: for  $m = 0, 1, 2, \dots$  do
5:   for  $t = t_m, \dots, t_{m+1} - 1$  do
6:     Sample  $A_t \sim \pi(\cdot|S_t, \Theta_m)$ 
7:     Take action  $A_t$  and observe  $R_{t+1}, S_{t+1}$ 
8:   end for
9:   Update  $\Theta_{m+1} \leftarrow \Theta_m + \alpha_m \text{GRADIENT}(m)$ 
10: end for

```

\mathbb{N}) obtained by updating the parameter periodically according to the gradient-ascent step $\Theta_{m+1} = \Theta_m + \alpha_m H_m$ introduced in (3), where H_m is provided by a family of learning algorithms, called *policy gradient*. The pseudocode of a generic policy-gradient algorithm, shown in Algorithm 1, is parametrized by a sequence $0 \triangleq t_0 < t_1 < t_2 < \dots$ of *observation times* and a sequence $\alpha_0, \alpha_1, \alpha_2, \dots > 0$ of *step sizes*. For each $m \in \mathbb{N}$, \mathcal{D}_m denotes batch m , obtained by applying policy $\pi(\Theta_m)$ at epoch m , given by

$$\mathcal{D}_m = ((S_t, A_t, R_{t+1}), t \in \{t_m, \dots, t_{m+1} - 1\}). \quad (9)$$

Given some initialization Θ_0 , Algorithm 1 calls a procedure GRADIENT that computes an estimate H_m of $\nabla J(\Theta_m)$ from \mathcal{D}_m , and it updates the parameter according to (3).

As discussed at the end of Section 3.3, finding an estimator H_m for $\nabla J(\Theta_m)$ directly from (7) is difficult in general. A common way to obtain H_m follows from the policy-gradient theorem [35, Chapter 13], which instead writes the gradient $\nabla J(\theta)$ using the action-value function q :

$$\nabla J(\theta) = \mathbb{E}[q(S, A) \nabla \log \pi(A|S, \theta)],$$

where $(S, A, R) \sim \text{STAT}(\theta)$, for each $\theta \in \Omega$. Consistently, in a model-free setting, policy-gradient methods like the actor-critic algorithm recalled in Appendix A.1 estimate $\nabla J(\Theta_m)$ by first estimating a value function. However, this approach can suffer from high-variance of the estimator, which slows down convergence, as described in Section 1. Some of these problems can be circumvented by exploiting the problem structure, as we will see now.

4 Score-aware gradient estimator (SAGE)

We now define the key structural assumption in our paper. Namely, that we have information on the impact of the policy parameter θ on the stationary distribution p . In Section 4.2, we will use this assumption to build a new family of estimators for the gradient ∇J that do not involve the state-value function, contrary to actor-critic. In Section 4.3, we will further explain how to use this insight to design a SAGE-based policy-gradient method.

4.1 Product-form and exponential family

As announced in the introduction, our end goal is to design a gradient estimator capable of exploiting information on the stationary distribution $p(\cdot|\theta)$ of the MDP when such information is available. Assumption 3 below formalizes this idea by assuming that the stationary distribution forms an exponential family parametrized by the policy parameter θ .

Assumption 3 (Stationary distribution). *There exist a scalar function $\Phi : \mathcal{S} \rightarrow \mathbb{R}_{>0}$, an integer $d \in \mathbb{N}_+$, a differentiable vector function $\rho : \Omega \rightarrow \mathbb{R}_{>0}^d$, and a vector function $x : \mathcal{S} \rightarrow \mathbb{R}^d$ such that the following two equivalent equations are satisfied:*

$$p(s|\theta) = \frac{1}{Z(\theta)} \Phi(s) \prod_{i=1}^d \rho_i(\theta)^{x_i(s)}, \quad s \in \mathcal{S}, \quad \theta \in \Omega, \quad (10\text{-PF})$$

$$\log p(s|\theta) = \log \Phi(s) + \log \rho(\theta)^\top x(s) - \log Z(\theta), \quad s \in \mathcal{S}, \quad \theta \in \Omega, \quad (10\text{-EF})$$

where the partition function $Z : \Omega \rightarrow \mathbb{R}_{>0}$ follows by normalization:

$$Z(\theta) = \sum_{s \in \mathcal{S}} \Phi(s) \prod_{i=1}^d \rho_i(\theta)^{x_i(s)} = \sum_{s \in \mathcal{S}} e^{\log \Phi(s) + \log \rho(\theta)^\top x(s)}, \quad \theta \in \Omega. \quad (11)$$

We will call Φ the balance function, ρ the load function, and x the sufficient statistics.

(10–PF) is the product-form variant of the stationary distribution, classical in queueing theory. (10–EF) is the exponential-family description of the distribution. This latter representation is more classical in machine learning [37] and will simplify our derivations. Let us briefly discuss the implications of this assumption as well as examples where this assumption is satisfied.

Assumption 3 implies that the stationary distribution p depends on the policy parameter θ only via the load function ρ . Yet, this assumption may not seem very restrictive *a priori*. Assuming for instance that the state space \mathcal{S} is finite, with $\mathcal{S} = \{s_1, s_2, \dots, s_N\}$, we can write the stationary distribution in the form (10) with $d = N$, $\rho_i(\theta) = p(s_i|\theta)$, $x_i(s) = \mathbf{1}[s = s_i]$, and $\Phi(s) = Z(\theta) = 1$, for each $\theta \in \mathbb{R}^n$, $s \in \mathcal{S}$, and $i \in \{1, 2, \dots, N\}$. However, writing the stationary distribution in this form is not helpful, in the sense that in general the function ρ will be prohibitively intricate. As we will see in Section 4.2, what will prove important in Assumption 3 is that the load function ρ is simple enough so that we can evaluate its Jacobian matrix function $D \log \rho$ numerically.

There is much literature on stochastic networks and queueing systems with a stationary distribution of the form (10–PF). Most works focus on performance evaluation, that is, evaluating $J(\theta)$ for some parameter $\theta \in \Omega$, assuming that the MDP’s transition probability kernel is known. In this context, the product-form (10–PF) arises in Jackson and Whittle networks [32, Chapter 1], BCMP networks [5], as well as more recent models arising in datacenter scheduling and online matching [16]¹. Building on this literature, in Section 6, we will consider policy parametrizations for control problems that also lead to a stationary distribution of the form (10).

In the next section, we exploit Assumption 3 to construct a gradient estimator that requires knowing the functions $D \log \rho$ and x but not the functions ρ , Φ , and Z .

¹Although the distributions recalled in [16, Theorems 3.9, 3.10, 3.13] do not seem to fit the framework of (10) *a priori* because the number of factors in the product can be arbitrarily large, some of these distributions can be rewritten in the form (10) by using an expanded state descriptor, as in [1, Equation (4), Corollary 2, and Theorem 6] and [28, Equation (7) and Proposition 3.1].

4.2 Score-aware gradient estimator (SAGE)

As our first contribution, Theorem 1 below gives simple expressions for $\nabla \log p(s|\theta)$ and $\nabla J(\theta)$ under Assumptions 1 to 3. Gradient estimators that will be formed using (13) will be called score-aware gradient estimators (SAGEs), to emphasize that the estimators rely on the simple expression (12) for the score $\nabla \log p(s|\theta)$. Particular cases of this result have been obtained in [11, 10, 26] for specific stochastic networks; our proof is shorter and more general thanks to the exponential form (10–EF).

Theorem 1. *Suppose that Assumptions 1 to 3 hold. For each $\theta \in \Omega$, we have*

$$\nabla \log p(s|\theta) = \text{D} \log \rho(\theta)^\top (x(s) - \mathbb{E}[x(S)]), \quad (12)$$

$$\nabla J(\theta) = \text{D} \log \rho(\theta)^\top \text{Cov}[R, x(S)] + \mathbb{E}[R \nabla \log \pi(A|S, \theta)], \quad (13)$$

where $(S, A, R) \sim \text{STAT}(\theta)$, $\text{Cov}[R, x(S)] = (\text{Cov}[R, x_1(S)], \dots, \text{Cov}[R, x_d(S)])^\top$, and the gradient and Jacobian operators, ∇ and D respectively, are taken with respect to θ .

Proof. Applying the gradient operator to the logarithm of (11) and simplifying yields

$$\nabla \log Z(\theta) = \text{D} \log \rho(\theta)^\top \mathbb{E}[x(S)]. \quad (14)$$

This equation is well-known and was already discussed in Section 2.1. Equation (12) follows by applying the gradient operator to (10–EF) and injecting (14). Equation (13) follows by injecting (12) into (8) and simplifying. \square

Assuming that the functions $\text{D} \log \rho$ and x are known in closed-form, Theorem 1 allows us to construct an estimator of $\nabla J(\theta)$ from a state-action-reward sequence $((S_t, A_t, R_{t+1}), t \in \{0, 1, \dots, T\})$ obtained by applying policy $\pi(\theta)$ at every time step as follows:

$$H = \text{D} \log \rho(\theta)^\top \bar{C} + \bar{E}, \quad (15)$$

where \bar{C} and \bar{E} are estimators of $\text{Cov}[R, x(S)]$ and $\mathbb{E}[R \nabla \log \pi(A|S, \theta)]$, respectively, obtained for instance by taking the sample mean and sample covariance. An estimator of the form (15) will be called a score-aware gradient estimator (SAGE). This idea will form the basis of the SAGE-based policy-gradient method that will be introduced in Section 4.3. Observe that such an estimator will typically be biased since the initial state S_0 is not stationary. Nonetheless, we will show in the proof of the convergence result in Section 5 that this bias does not prevent convergence.

The advantage of using a SAGE is twofold. First, the challenging task of estimating $\nabla J(\theta)$ is reduced to the simpler task of estimating the d -dimensional covariance $\text{Cov}[R, x(S)]$ and the n -dimensional expectation $\mathbb{E}[R \nabla \log \pi(A|S, \theta)]$, for which leveraging estimation techniques in the literature is possible. Also recall that the gradient estimator used in the actor–critic algorithm (Appendix A.1) relies on the state-value function, so that it requires estimating $|\mathcal{S}|$ values; we therefore anticipate SAGEs to yield better performance when $\max(n, d) \ll |\mathcal{S}|$; see examples from Sections 6.2 and 6.3. Second, as we will also observe in Section 6, SAGEs can “by-design” exploit information on the structure of the policy and stationary distribution. Actor–critic exploits this information only indirectly due to its dependency on the state-value function.

4.3 SAGE-based policy-gradient algorithm

Algorithm 2 introduces a SAGE-based policy-gradient method based on Theorem 1. For each $m \in \mathbb{N}$, the `GRADIENT`(m) procedure is called in the gradient-update step (Line 9) of Algorithm 1, at the end of epoch m , and returns an estimate of $\nabla J(\Theta_m)$ based on batch \mathcal{D}_m , defined in (9). Algorithm 2 can be understood as follows. According to Theorem 1, we have $\nabla J(\Theta_m) = \text{D log } \rho(\Theta_m)^\top \text{Cov}[R, x(S)] + \mathbb{E}[R \nabla \log \pi(A|S, \Theta_m)]$ with $(S, A, R) \sim \text{STAT}(\Theta_m)$. Lines 4 to 6 estimate $\text{Cov}[R, x(S)]$ using the usual sample covariance estimator. Line 7 estimates $\mathbb{E}[R \nabla \log \pi(A|S, \theta)]$ using the usual sample mean estimator. To simplify the signature of `GRADIENT`(m), we assume all variables from Algorithm 1, in particular batch \mathcal{D}_m , are accessible within Algorithm 2. The variable N_m computed on Line 3 is the batch size, i.e., the number of samples used to estimate the gradient $\nabla J(\Theta_m)$, and we assume it is greater than or equal to 2. An alternate implementation of the SAGE-based policy-gradient method that allows for batch sizes equal to 1 is given in Appendix A.2.

Algorithm 2 SAGE-based policy-gradient method, to be called on Line 9 of Algorithm 1.

```

1: Input: • Positive and differentiable policy parametrization  $(s, \theta, a) \mapsto \pi(a|s, \theta)$ 
           • Jacobian matrix function  $\theta \mapsto \text{D log } \rho(\theta)$ 
           • Feature function  $s \mapsto x(s)$ 
2: procedure GRADIENT( $m$ )
3:    $N_m \leftarrow t_{m+1} - t_m$ 
4:    $\bar{X}_m \leftarrow \frac{1}{N_m} \sum_{t=t_m}^{t_{m+1}-1} x(S_t)$ 
5:    $\bar{R}_m \leftarrow \frac{1}{N_m} \sum_{t=t_m}^{t_{m+1}-1} R_{t+1}$ 
6:    $\bar{C}_m \leftarrow \frac{1}{N_m-1} \sum_{t=t_m}^{t_{m+1}-1} (x(S_t) - \bar{X}_m)(R_{t+1} - \bar{R}_m)$ 
7:    $\bar{E}_m \leftarrow \frac{1}{N_m} \sum_{t=t_m}^{t_{m+1}-1} R_{t+1} \nabla \log \pi(A_t|S_t, \Theta_m)$ 
8:   return  $\text{D log } \rho(\Theta_m) \bar{C}_m + \bar{E}_m$ 
9: end procedure

```

Recall that our initial goal was to exploit information on the stationary distribution, when such information is available. Consistently, compared to actor-critic (Appendix A.1), the SAGE-based method of Algorithm 2 requires as input the Jacobian matrix function $\text{D log } \rho$ and the sufficient statistics x . In return, as we will see in Sections 5 and 6, the SAGEs-based method relies on a lower-dimensional estimator whenever $\max(n, d) \ll |\mathcal{S}|$, which can lead to an improved convergence.

5 A local convergence result

Our goal in this section is to study the limiting behavior of Algorithm 2. To do so, we will consider this algorithm as an SGA algorithm that uses biased gradient estimates. The gradient estimates are biased because they arise from the MCMC estimations from Lines 4–7 in Algorithm 2. Throughout the proof, we will assume for simplicity that the reward is a deterministic function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$.

Under this assumption, for each $m \in \mathbb{N}$, Algorithm 2 follows the gradient ascent step (3), with

$$H_m = D \log \rho(\Theta_m)^\top \bar{C}_m + \bar{E}_m, \text{ where } \begin{cases} \bar{X}_m = \frac{\sum_{t=t_m}^{t_{m+1}-1} x(S_t)}{t_{m+1} - t_m}, & \bar{R}_m = \frac{\sum_{t=t_m}^{t_{m+1}-1} r(S_t, A_t)}{t_{m+1} - t_m}, \\ \bar{C}_m = \frac{\sum_{t=t_m}^{t_{m+1}-1} (x(S_t) - \bar{X}_m)(r(S_t, A_t) - \bar{R}_m)}{t_{m+1} - t_m - 1}, \\ \bar{E}_m = \frac{\sum_{t=t_m}^{t_{m+1}-1} r(S_t, A_t) \nabla \log \pi(A_t | S_t, \Theta_m)}{t_{m+1} - t_m}, \end{cases} \quad (16)$$

The estimates \bar{X}_m , \bar{R}_m , and \bar{C}_m are functions of \mathcal{D}_m , while H_m and \bar{E}_m are functions of \mathcal{D}_m and Θ_m . We will additionally apply decreasing step sizes and increasing batch sizes of the form

$$\alpha_m = \frac{\alpha}{(m+1)^\sigma} \quad \text{and} \quad t_{m+1} = t_m + \ell m^{\frac{\sigma}{2} + \kappa} \quad \text{for each } m \in \mathbb{N}, \quad (17)$$

for some parameters $\alpha \in (0, \infty)$, $\ell \in (1, \infty)$, $\sigma \in (2/3, 1)$, and $\kappa \in [0, \infty)$.

Our goal—to study the limiting algorithmic behavior of Algorithm 2—is equivalent to studying the limiting algorithmic behavior of the stochastic recursion (3). In particular, we will show local convergence of the iterates of (3) and (16) to the following set of global maximizers:

$$\mathcal{M} = \{\theta \in \Omega : J(\theta) = J^*\}, \quad \text{where } J^* = \sup_{\theta \in \Omega} J(\theta). \quad (18)$$

We will assume \mathcal{M} to be nonempty, that is, $\mathcal{M} \neq \emptyset$. The assumptions that we consider (Assumption 7 below) allow us to assume that \mathcal{M} is only locally a manifold. Consequently, J can be nonconvex with noncompact level-subsets, and J is even allowed not to exist outside the local neighborhood; namely, the policies may be unstable. In this latter case $\theta \notin \Omega$, and we will use the convention that if θ yields an unstable policy, then $J(\theta) = \inf_{s \in \mathcal{S}, a \in \mathcal{A}} r(s, a) \geq -\infty$. While the previous assumptions allow for general objective functions, the convergence will be guaranteed only close to the set of maxima \mathcal{M} .

5.1 Assumptions pertaining to algorithmic convergence

We use the Markov chain of state-action pairs. Specifically, consider the pairs $\{(S_t, A_t)\}_{t \geq 0} \subset \mathcal{S} \times \mathcal{A}$, where A_t is generated according to policy $\pi(\cdot | S_t, \theta)$. For a given $\theta \in \Omega$, the one-step transition probability and the stationary distribution of this Markov chain are

$$P((s', a') | (s, a), \theta) = \pi(a' | s', \theta) P(s' | s, a), \quad (19)$$

$$\tilde{p}((s, a) | \theta) = p(s | \theta) \pi(a | s, \theta) \quad \text{for } (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (20)$$

The following are assumed:

Assumption 4. *There exists a function $\mathcal{L} : \mathcal{S} \times \mathcal{A} \rightarrow [1, \infty)$ such that, for any $\theta^* \in \mathcal{M}$, there exist a neighborhood U of θ^* in Ω and four constants $\lambda \in (0, 1)$, $C > 0$, $b \in \mathbb{R}^+$, and $v \geq 16$ such that, for each $\theta \in U$, the policy $\pi(\cdot | \cdot, \theta)$ is such that*

$$\sum_{(s', a') \in \mathcal{S} \times \mathcal{A}} P((s', a') | (s, a), \theta) (\mathcal{L}(s', a'))^v \leq \lambda (\mathcal{L}(s, a))^v + b, \quad \text{for each } (s, a) \in \mathcal{S} \times \mathcal{A},$$

and, for each $\ell \in \mathbb{N}_+$ and $(s, a), (s', a') \in \mathcal{S} \times \mathcal{A}$,

$$\left| P^\ell((s', a')|(s, a), \theta) - \tilde{p}((s', a')|\theta) \right| \leq C\lambda^\ell \mathcal{L}(s, a),$$

where $P^\ell(\theta)$ is the ℓ -step transition probability kernel of the Markov chain with transition probability kernel (19).

Assumption 5. *There exists a constant $C > 0$ such that $|\mathrm{D} \log \rho(\theta)|_{op} < C$ for each $\theta \in \Omega$.*

Assumption 6. *Let \mathcal{L} be the Lyapunov function from Assumption 4. For any $\theta^* \in \mathcal{M}$, if U is a local neighborhood satisfying the conditions of Assumption 4, then there exists a constant $C > 0$ such that for any $\theta \in U$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$,*

$$|x(s)| < C\mathcal{L}(s, a), \quad |r(s, a)| < C\mathcal{L}(s, a), \quad |r(s, a)\nabla \log \pi(a|s, \theta)| < C\mathcal{L}(s, a). \quad (21)$$

Assumption 7. *There exist an integer $\mathbf{n} \in \{0, 1, \dots, n-1\}$ and an open subset $U \subseteq \Omega$ such that (i) $\mathcal{M} \cap U$ is a nonempty \mathbf{n} -dimensional C^2 -submanifold of \mathbb{R}^n , and (ii) the Hessian of J at θ^* has rank $n - \mathbf{n}$, for each $\theta^* \in \mathcal{M} \cap U$.*

These assumptions have the following interpretation. Assumption 4 formalizes that the Markov chain is stable for policies close to the maximum. Remarkably, it does not assume that the chain is geometrically ergodic for all policies, only for those close to an optimal policy when $\theta \in \Omega$. This stability is guaranteed by a local Lyapunov function \mathcal{L} *uniformly* over some neighborhood close to a maximizer.

Assumptions 5 and 6 together guarantee that the estimator H_m concentrates around $\nabla J(\Theta_m)$ at an appropriate rate. Assumption 5 is easy to verify in our examples since ρ is always positive and bounded. Assumption 6 guarantees that the reward $r(s, a)$ and sufficient statistics $x(s)$ cannot grow fast enough in s to perturb the stability of the MDP. In many applications from queueing Assumption 6 holds. Namely, S is usually a normed space and the order of the Lyapunov function $\mathcal{L}(s, a)$ is exponential in the norm of the state $s \in S$, compared to the sufficient statistic x which has an order linear in the norm of s . We remark that, in a setting with a bounded reward function r and a bounded map x or with a finite state space, Assumption 6 becomes trivial.

Assumption 7 is a geometric condition: it guarantees that, locally around the set of maxima \mathcal{M} , in directions perpendicular to \mathcal{M} , J behaves approximately in a convex manner. Concretely, this means that $\mathrm{Hess}_\theta J$ has strictly negative eigenvalues in the directions normal to \mathcal{M} —also referred to as the Hessian being *nondegenerate*. Thus, there is one-to-one correspondence between local directions around $\theta \in \mathcal{M}$ that decrease J and directions that do not belong to the tangent space of \mathcal{M} . Strictly concave functions satisfy that $\mathbf{n} = 0$ and Assumption 7 is thus automatically satisfied in such cases. If $\mathcal{M} \cap U = \{\theta^*\}$ is a singleton, Assumption 7 reduces to assuming that $\mathrm{Hess}_{\theta^*} J$ is negative definite. Assumption 7 in a general setting can be difficult to verify, but by adding a regularization term, it can be guaranteed to hold in a broad sense (see Section 5.5).

5.2 Local convergence results

This is our main convergence result for the case that the set of maxima is not necessarily bounded.

Theorem 2 (Noncompact case). *Suppose that Assumptions 1 to 7 hold. For every maximizer $\theta^* \in \mathcal{M} \cap U$, there exist constants $c > 0$ and $\alpha_0 > 0$ such that, for each $\alpha \in (0, \alpha_0]$, there exists a*

nonempty neighborhood V of θ^* and $\ell_0 \geq 1$ such that, for each $\ell \in [\ell_0, \infty)$, $\sigma \in (2/3, 1)$, $\kappa \in [0, \infty)$ with $\sigma + \kappa > 1$, we have, for each $m \in \mathbb{N}_+$,

$$\mathbb{P}[J(\Theta_m) < J^* - \epsilon | \Theta_0 \in V] \leq c \left(\epsilon^{-2} m^{-\sigma-\kappa} + \frac{m^{1-\sigma-\kappa}}{\ell} + \frac{\alpha^2}{\ell} + \alpha m^{-\kappa/2} + \frac{\alpha m^{1-(\sigma+\kappa)/2}}{\sqrt{\ell}} \right), \quad (22)$$

where $(\Theta_m, m \in \mathbb{N})$ is a random sequence initialized with $\mathbb{P}[\Theta_0 \in V] > 0$ and built by recursively applying the gradient ascent step (3) with the gradient update (16) and the step and batch sizes (17) parameterized by these values of α , ℓ , σ , and κ .

In Theorem 2, by setting the parameters α , ℓ , σ , and κ in (17) appropriately, we can make the probability of Θ_m being ϵ -suboptimal arbitrarily small. Specifically, the step and batch sizes for each epoch allow us to control the variance of the estimators in (16). This shows that the SAGE-based policy-gradient method converges with large probability. The bound can be understood as follows. The term in (22) on the bound depending on ϵ characterizes the convergence rate assuming that all iterates up to time m remain in V . The remaining terms in (22) estimate the probability that the iterates escape the set V , which can be made small by tuning parameters that diminish the variance of the estimator H_m , such as setting κ or ℓ large—the batch size becomes larger.

Theorem 2 extends the result of [14, Thm. 25] to a Markovian setting with inability to restart. In our case, the bias can be controlled by using a longer batch size with exponent at least $\sigma/2$. Furthermore, we also use the Lyapunov function to keep track of the state of the MDP as we update the parameter in V and ensure stability. The proof sketch of Theorem 2 can be found in Section 5.6. In Appendix D, we also consider the case that $\mathcal{M} \cap U$ is compact, which can be used to improve Theorem 2. Note that the sequence $(\Theta_m, m \in \mathbb{N})$ from Theorem 2 is well defined even if unstable policies occur, since the update H_m from (16) is finite. In this case, recall that we have the convention that if θ yields an unstable policy then $J(\theta) = \inf_{s \in S, a \in A} r(s, a) \geq -\infty$. As a side remark, observe that Theorem 2 holds for any estimator \tilde{H}_m of the gradient $J(\Theta_m)$ provided that this estimator satisfies Lemma 1 stated in Section 5.6 below, and is finite.

5.3 Lower bound

As noted in Theorem 2, the rate in (22) includes the probability that the iterates escape V , outside which convergence cannot be guaranteed. Indeed, there is a term $O(\alpha^2/\ell)$ that characterizes the probability that the iterates escape the basin of attraction. For general settings, this term cannot be avoided, even in the unbiased case. In fact, the proposition below shows that for any $\beta > 0$ there are cases where there is a positive lower bound depending on $\alpha^{2+\beta}/\ell$. In Proposition 1 below, we consider an SGA setting with i.i.d. data, where the target is to maximize a function f using estimators H_m for the gradient $\nabla f(\Theta_m)$ at epoch m . In a non-RL setting, we usually have $H_m = H_m(\Theta_m, Z_m)$, where Z_m is a collection of i.i.d. random variables and \mathcal{F}_m denotes the sigma algebra of the random variables $\Theta_0, \dots, \Theta_m$ as well as Z_0, \dots, Z_{m-1} . For our result, we assume the iterates Θ_m satisfy (3), and $\eta_m = H_m - \nabla f(\Theta_m)$ satisfies the following unbiased conditional concentration bounds for some $C > 0$:

$$\mathbb{E}[\eta_m | \mathcal{F}_m] = 0 \quad \text{and} \quad \mathbb{E}[|\eta_m|^2 | \mathcal{F}_m] \leq \frac{C}{t_{m+1} - t_m}. \quad (23)$$

Proposition 1 below shows that Theorem 2 is almost sharp and characterizes the limitations of making local assumptions only. As we will see in Section 6.1, however, there are examples where only local convergence can be expected. The proof of Proposition 1 can be found in Appendix E.

Proposition 1. For any $\beta > 0$, there are functions $f \in C^\infty(\mathbb{R}^n)$ with a maximum $f^* = f(\theta^*)$ satisfying Assumption 7, such that if the iterates Θ_m satisfy (3) and the gradient estimator $H_m = \nabla f(\Theta_m) + \eta_m$ satisfies (23), there exists a constant $c > 0$ depending on f and independent of m such that for any $\epsilon \in (0, 1)$, $1 > \alpha > 0$, $\delta > 0$, $\ell \geq 1$ and any $\sigma \geq 0, \kappa \geq 0$, in (17) we have that

$$\mathbb{P}[f(\Theta_m) < f^* - \epsilon | \Theta_0 \in V] \geq c \frac{\alpha^{2+\beta}}{\ell}. \quad (24)$$

5.4 Sample Complexity

With the additional assumption that the reward r is bounded, we immediately obtain a typical sample complexity bound for the number of epochs.

Corollary 1. Under the same assumptions and setting as in Theorem 2, assume moreover that there exists some $b > 0$ such that $|r(s, a)| < b$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Then we have

$$\begin{aligned} \mathbb{E}[J^* - J(\Theta_m) | \Theta_0 \in V] &\leq 3 \left(\frac{cb}{2} \right)^{\frac{1}{3}} m^{-\frac{(\sigma+\kappa)}{3}} + 2 \frac{bc}{\ell} m^{1-(\sigma+\kappa)} + 2bc \frac{\alpha^2}{\ell} \\ &\quad + 2bc\alpha m^{-\kappa/2} + 2 \frac{bc\alpha}{\ell} m^{1-(\sigma+\kappa)/2}. \end{aligned}$$

Proof. Optimizing the following bound over $\epsilon > 0$,

$$\mathbb{E}[J^* - J(\Theta_m) | \Theta_0 \in V] \leq \mathbb{P}[J(\Theta_m) < J^* - \epsilon | \Theta_0 \in V] 2b + \epsilon, \quad (25)$$

immediately yields the result by using the bound from (22). \square

Since the algorithm uses an increasing batch size, Corollary 2 yields only the first-order sample complexity, that is, the complexity in the number of gradient evaluations—estimations or epochs in our case. We can obtain a sample complexity for the number of samples as well. To do so, we define the total number of samples up to epoch m as

$$T = t_{m+1} = \sum_{k=1}^m \ell k^{\kappa+\sigma/2} = \Theta \left(\frac{\ell}{\kappa + \frac{\sigma}{2} + 1} m^{\kappa + \frac{\sigma}{2} + 1} \right). \quad (26)$$

For an epoch l , we consider that Θ_l is fixed for any sample of that epoch and define for $t \in [t_l, t_{l+1}]$ that $\tilde{\Theta}_t = \Theta_l$, then we can use the bound (26) in Proposition 2 to obtain that

$$\mathbb{E}[J^* - J(\tilde{\Theta}_T) | \tilde{\Theta}_0 \in V] = O \left(T^{-\frac{(\sigma+\kappa)}{3(\kappa+\frac{\sigma}{2}+1)}} + T^{\frac{1-(\sigma+\kappa)/2}{(\kappa+\frac{\sigma}{2}+1)}} + T^{\frac{-\kappa}{2(\kappa+\frac{\sigma}{2}+1)}} + \frac{\alpha^2}{\ell} \right). \quad (27)$$

Note that by looking at the orders in (27), for any $\epsilon > 0$, there exists $\kappa_0(\epsilon) > 0$ such that if $\kappa \geq \kappa_0(\epsilon)$, then

$$\mathbb{E}[J^* - J(\tilde{\Theta}_T) | \tilde{\Theta}_0 \in V] = O \left(T^{-\frac{1}{3+\epsilon}} + \frac{\alpha^2}{\ell} \right). \quad (28)$$

Besides the term α^2/ℓ inherent to the local analysis, the term $T^{-1/(3+\epsilon)}$ in (28) cannot be easily compared with other common sample complexity results that assume global features for $J(\theta)$ or the gradient estimator H_m . However, we remark that one example of a similar order is obtained for policy-gradient methods under a weak (global) Polyak–Łojasiewicz condition for $J(\theta)$; see [40].

5.5 Local convergence with entropy regularization

A well-known phenomenon that can occur when using the softmax policy (6) is that, if the optimal policy is deterministic, the iterates converge to this optimal policy only when $\Theta_m \rightarrow \infty$. Problems where this occurs will thus not satisfy Assumption 7: the set of maxima will be empty. This phenomenon is illustrated in the example of Section 6.1. One prevalent method to mitigate the occurrence of maxima at the boundary involves incorporating a regularization term, often linked to relative entropy $\text{KL}[\tilde{\pi} \parallel \pi]$ of the policy π compared to a given $\tilde{\pi}$, defined below in (29).

Let $\tilde{\pi}$ be a policy of the same type as those defined in (6) and let ζ be a distribution on \mathcal{S} such that $\zeta(h^{-1}(i)) > 0$ for any $i \in \mathcal{I}$, where h is the index map defined for the class of policies that we use in (6). We define the regularization term as

$$\mathcal{R}_{\tilde{\pi}}(\theta) = \mathbb{E}_{S \sim \zeta}[\text{KL}[\tilde{\pi}(\cdot | S) \parallel \pi(\cdot | S, \theta)]] = \sum_{s \in \mathcal{S}} \zeta(s) \mathbb{E}_{A \sim \tilde{\pi}(\cdot | s)} \left[\log \left(\frac{\tilde{\pi}(A | s)}{\pi(A | s, \theta)} \right) \right]. \quad (29)$$

For some $b > 0$ we define

$$J_{\tilde{\pi}}(\theta) = J(\theta) - b\mathcal{R}_{\tilde{\pi}}(\theta). \quad (30)$$

We can show that adding (29) to $J(\theta)$ defined in (7) not only prevents maxima from being at the boundary, but also allows us to avoid using Assumption 7 altogether. The next proposition is proved in Appendix F.

Proposition 2. *Assume that we use the softmax policy from (6) and let $J(\theta)$ be defined as in (7). Then for almost every policy $\tilde{\pi}$ in the class of (6) with respect to its Lebesgue measure,*

1. *the function $J_{\tilde{\pi}}(\theta)$ in (30) satisfies Assumption 7 and the set of maximizers is bounded, and*
2. *Theorem 2 for $J_{\tilde{\pi}}(\theta)$ holds without Assumption 7.*

5.6 Proof outline for Theorem 2

We extend the local approach presented in [14, §5], that deals with convergence of SGD where the samples used to estimate the gradient are i.i.d. We consider instead an RL setting where data is Markovian and thus presents a bias. Fortunately, we can overcome its presence by adding an increasing batch size while tracking the states of the Markov chain via the local Lyapunov function from Assumption 4, which guarantees a stable MPD trajectory as long as the parameter is in a neighborhood close to the maximum. Below we give an outline of the technique employed. For the full proof we refer to Appendix C.

Structure of the proof

The proof of Theorems 2 consists of several parts. To show a bound on the probability that Θ_m is ϵ -suboptimal, we consider the event \mathcal{B}_m that all previous iterates $\Theta_0, \dots, \Theta_m$ belong to a local neighborhood V , and the complementary event $\bar{\mathcal{B}}_m$. We bound these separately. Firstly, on the event \mathcal{B}_m , we show in Lemma 2 that the iterates converge to \mathcal{M} , and we obtain a bound on the ϵ -suboptimal probability for this case. Secondly, the probability of the complement $\bar{\mathcal{B}}_m$ is written as the sum of the probabilities of two disjoint events, namely, (i) $\Theta_{m+1} \notin V$ and the distance of Θ_{m+1} to \mathcal{M} is larger than δ , and (ii) $\Theta_{m+1} \notin V$ and the distance of Θ_{m+1} to \mathcal{M} is less than δ . Intuitively, these events group the cases when Θ_{m+1} escapes V in ‘normal directions’ to \mathcal{M} and in ‘tangent

directions' to \mathcal{M} , respectively. We can bound the former by using concentration inequalities, but for the latter we need a maximal excursion bound (Lemma 3 below). Combining all bounds results in an upper bound on $\mathbb{P}[\mathcal{B}_m]$ (Lemma 4). The local properties of J are then be used to finish the proof. Crucially, we use throughout the proof that the local Lyapunov function guarantees stability of the Markov chain and the gradient estimator within V , as well as keeps track of the initial state for each epoch.

Preliminary step: Definition of the local neighborhood and bound strategy

For $\theta^* \in \mathcal{M} \cap U$, we now define a neighborhood $V_{\tau, \delta}(\theta^*)$ of θ^* where the algorithm will operate. Let $\bar{B}_\tau(\theta) := \{\theta \in \Omega : |\theta - \theta^*| \leq \tau\}$ denote a closed ball around θ^* with radius τ and $\text{dist}(\theta, L) = \sup_{\theta' \in L} |\theta - \theta'|$ for an open set L . Let U be the neighborhood of θ^* described in Assumptions 4 and 7. We define a tubular neighborhood of θ^* as follows:

$$V_{\tau, \delta}(\theta^*) := \{\theta \in \Omega \cap U : \text{dist}(\theta, M \cap U) = \text{dist}(\theta, \bar{B}_\tau(\theta^*) \cap M \cap U) < \delta\}. \quad (31)$$

Crucially, Assumption 7 implies that there exists $\delta_0, \tau_0 > 0$ such that for any $\delta \in (0, \delta_0]$ and $\tau \in (0, \tau_0]$ an equivalent definition of the set is then

$$V_{\tau, \delta}(\theta^*) = \{y + v : y \in (\bar{B}_\tau(\theta^*) \cap \mathcal{M} \cap U) \text{ and } v \in (T_y(\mathcal{M} \cap U))^\perp \text{ with } |v| < \delta, \mathbf{p}(y + v) = y\}. \quad (32)$$

Here, \mathbf{p} is the unique local projection onto $\mathcal{M} \cap U$, and $T_y(\mathcal{M} \cap U)^\perp$ denotes the cotangent space of $\mathcal{M} \cap U$ at y . For further details on this geometric statement, we refer to [14, Prop. 13] or [23, Thm. 6.24].

In the following, we let U denote the intersection of the neighborhoods from Assumptions 4 and 7, and \mathcal{L} the Lyapunov function from Assumption 4. For any $m \in \mathbb{N}_+$ define the event and filtration

$$\mathcal{B}_m := \bigcap_{l=1}^m \{\Theta_l \in V_{\tau, \delta}(\theta^*)\}, \quad (33)$$

$$\mathcal{F}_m := \sigma(\mathcal{D}_1 \cup \dots \cup \mathcal{D}_{m-1} \cup \{\Theta_0, \dots, \Theta_m\}). \quad (34)$$

Due to the local properties of J , Theorem 2 can be shown by bounding $\mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) \leq \epsilon | \mathcal{B}_0]$. By separating into the event \mathcal{B}_m and its complement, we can show that

$$\mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) \leq \epsilon | \mathcal{B}_0] \leq \mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) \mathbb{1}[\mathcal{B}_{m-1}] \geq \epsilon] + \mathbb{P}[\bar{\mathcal{B}}_m]. \quad (35)$$

The remaining steps of the proof consist of bounding both terms in the right-hand side of (35).

Step 1: The variance of the gradient estimator decreases, in spite of the bias

For each $m \in \mathbb{N}_+$, let

$$\eta_m := H_m - \nabla J(\Theta_m), \quad (36)$$

denote the difference between the gradient estimator H_m in (16) and the true gradient $\nabla J(\Theta_m)$. Lemma 1 below implies that the difference in (36) is, ultimately, small. From Assumption 4, since the state-action chain $\{(S_t, A_t)\}_{t \geq 0}$ has a Lyapunov function \mathcal{L} , so does the chain $\{S_t\}_{t > 0}$ with

$$\mathcal{L}_v(s) = \sum_{a \in \mathcal{A}} \mathcal{L}(s, a)^v \pi(a | s, \theta), \quad (37)$$

where $v \geq 16$ is the exponent from Assumption 4. We can define $\mathcal{L}_4(s)$ similarly. The following lemma bounds the variance of η_m on the event \mathcal{B}_m , which can be controlled with the local Lyapunov function. The proof of Lemma 1 is deferred to Appendix C.3.

Lemma 1. *Suppose that Assumptions 1–7 hold. There exists a constant $C > 0$ that depends on θ^* , U , and J such that for every $m \in \mathbb{N}_+$,*

$$|\mathbb{E}[\eta_m \mathbb{1}[\mathcal{B}_m] | \mathcal{F}_m]| \leq \frac{C}{t_{m+1} - t_m} \mathcal{L}_4(S_{t_m})^{1/2}, \quad (38)$$

$$\mathbb{E}[|\eta_m|^l \mathbb{1}[\mathcal{B}_m] | \mathcal{F}_m] \leq \frac{C}{(t_{m+1} - t_m)^{p/2}} \mathcal{L}_4(S_{t_m})^{l/2}, \quad \text{for every } l \in \{1, 2\}. \quad (39)$$

Lemma 1 helps to determine the bias incurred when starting at a different state than that of stationarity, and is used to bound the term $\text{dist}(\Theta_m, \mathcal{M} \cap U)$ from (35) in Lemma 2 below. Note that the definition of SAGE and Assumptions 5 and 6 are used. We remark, however, that any other estimator \tilde{H}_m of ∇J satisfying (38) and (39) from Lemma 1 would yield similar guarantees.

Step 2: Convergence on the event \mathcal{B}_{m-1} .

We turn to the first term on the right-hand side of (35) and examine, on the event \mathcal{B}_{m-1} , if the iterates converge. Using a similar proof strategy as that of [14, Proposition 20] for the unbiased non-Markovian case, we prove in Lemma 2 that the variance of the distance to the set of minima decreases under the appropriate step and batch sizes. The proof of Lemma 2 is in Appendix C.4.

Lemma 2. *Suppose that Assumptions 1–7 hold. There then exist $\tau_0, \alpha_0, \ell_0 > 0$ and $c > 0$ such that for any $\tau \in (0, \tau_0]$, $\alpha \in (0, \alpha_0]$ and $\ell \in [\ell_0, \infty)$ there also exists $\delta_0 > 0$ such that for any $\delta \in (0, \delta_0]$ and $m \in \mathbb{N}_+$,*

$$\mathbb{E}\left[\left(\text{dist}(\Theta_m, \mathcal{M} \cap U) \wedge \delta\right)^2 \mathbb{1}[\mathcal{B}_{m-1}]\right] \leq cm^{-\sigma-\kappa}. \quad (40)$$

Compared to the unbiased case in [14], Lemma 2 needs to use a larger batch size to deal with the bias of Lemma 1. A key result required is that on the event \mathcal{B}_{m-1} , the Lyapunov function is bounded in expectation. With Lemma 2 together with Markov’s inequality, a bound of order $\epsilon^{-2}m^{-\sigma-\kappa}$ for the first term in (35) follows.

Step 3: Excursion and the probability of staying in $V_{\tau, \delta}(\theta^*)$

We next focus on $\mathbb{P}[\bar{\mathcal{B}}_m]$. Since

$$\mathbb{P}[\mathcal{B}_m] \geq \mathbb{P}[\mathcal{B}_{m-1}] - \mathbb{P}[\Theta_m \notin V_{\tau, \delta}(\theta^*), \mathcal{B}_{m-1}], \quad (41)$$

we can use a recursive argument to obtain a lower bound, *if* we can bound first the probability

$$\begin{aligned} \mathbb{P}[\Theta_m \notin V_{\tau, \delta}(\theta^*), \mathcal{B}_{m-1}] &= \mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) > \delta, \mathcal{B}_{m-1}] \\ &\quad + \mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) \leq \delta, \Theta_m \notin V_{\tau, \delta}(\theta^*), \mathcal{B}_{m-1}]. \end{aligned} \quad (42)$$

The first term in (42) represents the event that the iterand Θ_m escapes the set $V_{\tau, \delta}(\theta^*)$ in directions ‘normal’ to \mathcal{M} , while the second term represents the escape in directions ‘tangent’ to \mathcal{M} —intuition derived from the fact that, in that latter event, we still have $\text{dist}(\Theta_m, \mathcal{M} \cap U) \leq \delta$.

The first term in (42) can be bounded by using the local geometric properties around minima in the set U and associating the escape probability with the probability that on the event \mathcal{B}_{m-1} escape can only occur if $|\eta_m|$ is large enough. The probability of this last event happening can then be controlled with the variance estimates from Lemma 1.

After a recursive argument, we have to consider the second term in (42) for all $l \leq m$. Fortunately, this term can be bounded by first looking at the maximal excursion event for the iterates $\{\Theta_l\}_{l=1}^m$. The proof can be found in Appendix C.5. Here, the Lyapunov function again plays a crucial role to control the variance of the gradient estimator on the events \mathcal{B}_l for $l \leq m$, compared to an unbiased and non-Markovian case.

Lemma 3. *Suppose that Assumptions 1–7 hold. Then there exist $\mathfrak{r}_0, \alpha_0, \ell_0 > 0$, and $c > 0$ such that for any $\mathfrak{r} \in (0, \mathfrak{r}_0]$, $\alpha \in (0, \alpha_0]$ and $\ell \in [\ell_0, \infty)$, there exist $\delta_0 > 0$ such that for any $\delta \in (0, \delta_0]$ and $m \geq 1$,*

$$\mathbb{E} \left[\max_{1 \leq l \leq m} |\Theta_l - \Theta_0| \mathbf{1}[\mathcal{B}_{l-1}] \right] < c\alpha \left(m^{1-3\sigma/2-\kappa/2} + \sqrt{\frac{1}{\ell}} m^{1-5\sigma/8-\kappa/2} \right). \quad (43)$$

Finally, with the previous steps we obtain a bound on $\mathbb{P}[\mathcal{B}_m]$ in Lemma 4 below². The proof of Lemma 4 can be found in Appendix C.6.

Lemma 4. *Suppose that Assumptions 1–7 hold and $\sigma + \kappa > 1$. There exist \mathfrak{r}_0, α_0 , such that for any $\mathfrak{r} \in (0, \mathfrak{r}_0]$, $\alpha \in (0, \alpha_0]$, there also exists a constant $c > 0$, $\delta_0 > 0$ such that for any $\delta \in (0, \delta_0]$, if $\Theta_0 \in V_{\mathfrak{r}/2, \delta}(\theta^*)$, there exists $\ell_0 > 0$ such that for any $\ell \in [\ell_0, \infty)$ and $m \in \mathbb{N}_+$,*

$$\mathbb{P}[\mathcal{B}_m] \geq \exp\left(-\frac{c\alpha^2}{\delta^2\ell}\right) - \frac{c}{\delta^4\ell} m^{1-\sigma-\kappa} - c\alpha \frac{(m^{1-3\sigma/2-\kappa/2} + \ell^{-1/2} m^{1-5\sigma/8-\kappa/2})}{(\mathfrak{r}/2 - 2\delta)_+}. \quad (44)$$

Step 4: Combining the bounds in (35).

The proof of Theorem 2 follows the same steps as are used to prove [14, Theorem 25] by substituting the modified bounds that we have obtained from Lemmas 2 and 4 in (35). The details can be found in Appendix C.2.

6 Examples and numerical results

In Section 5, we have shown convergence of a SAGE-based policy-gradient method under the assumptions in Section 5.1. We now numerically assess its performance in examples from stochastic networks and statistical physics that go beyond these assumptions. Specifically, we examine a single-server queue with admission control in Section 6.1, a load-balancing system in Section 6.2, and an example of the Ising model with Glauber dynamics in Section 6.3. These examples satisfy the assumptions of Section 4, required to implement the SAGE-based policy-gradient method, but not necessarily those of Section 5.1 used to prove convergence. More discussion about these assumptions is provided in Appendix B.

²In (43), if the Lyapunov function has only smaller moments than order ν , then condition on $\kappa \geq 0$ will become stricter. In particular, κ tunes the batch size required to sample from the tails of the stationary distribution and may be required to be positive depending the moments of the Lyapunov function. The terms σ and κ can be tuned to control the bias coming from variance and nonstationarity, and finite batch size, respectively.

Simulation setup. Plots are obtained by averaging 10 independent simulations, each lasting $T_{\max} = 10^6$ time steps. The initial parameter vector Θ_0 is taken to be the zero vector, yielding in each example a uniform policy over the action space. The SAGE-based algorithm (Algorithm 2) is run with batch size 100 and step size $\alpha_m = 10^{-1}$. The actor-critic algorithm (Appendix A.1) is run with batch size 1 and step sizes $\alpha_m = 10^{-3}$ and $\alpha_v = \alpha_{\bar{r}} = 10^{-2}$. It uses a tabular value function which we initially treat as containing all-zeros, and which in practice is expanded as states are visited.

6.1 Admission control in a single-server queue

Consider a queueing system where jobs arrive according to a Poisson process with rate $\lambda > 0$, service times are independent and exponentially distributed with rate $\mu > 0$, and the server applies an arbitrary nonidling nonanticipating scheduling policy such as first-come-first-served or processor-sharing. This model is also commonly known as an M/M/1 queue in the literature. When a job arrives, the agent decides to either admit or reject it: in the former case, the job is added to the queue, otherwise it is lost permanently. The agent receives a one-time reward $\gamma > 0$ for each admitted job (incentive to accept jobs) and incurs a holding cost $\eta > 0$ per job per time unit (incentive to reject jobs). The goal is to find an admission-control policy that achieves a trade-off between these two conflicting objectives.

The problem can be related to the framework of Section 3 as follows. For $t \in \mathbb{N}$, let S_t denote the number of jobs in the system right before the arrival of the $(t + 1)$ th job, and let A_t denote the decision of either admitting or rejecting this job. We have $\mathcal{S} = \mathbb{N}$ and $\mathcal{A} = \{\text{admit}, \text{reject}\}$. Also, let $(\Sigma_\tau, \tau \in \mathbb{R}_{\geq 0})$ denote the continuous-time process that describes the evolution of the number of jobs over time and $(T_t, t \in \mathbb{N})$ the sequence of job arrival times, so that $S_0 = \Sigma_0$ and $S_t = \lim_{\sigma \uparrow T_t} \Sigma_\sigma$ for $t \in \mathbb{N}_+$. Rewards are given by

$$R_{t+1} = r_{\text{disc}}(S_t, A_t) + \int_{T_t}^{T_{t+1}} r_{\text{cont}}(\Sigma_\tau) d\tau,$$

where $r_{\text{disc}}(s, a) = \gamma \mathbb{1}[a = \text{admit}]$ represents the one-time admission reward and $r_{\text{cont}}(s) = -\eta s$ the holding cost incurred continuously over time. We use this common reward structure in this example, but we remark that arbitrary reward functions r_{disc} and r_{cont} are possible.

For each $k \in \mathbb{N}$, we define a random policy parametrization π_k with threshold k and parameter vector³ $\theta = (\theta_0, \theta_1, \dots, \theta_k) \in \mathbb{R}^{k+1}$ as follows. Under policy π_k , an incoming job finding s jobs in the system is accepted with probability

$$\pi_k(\text{admit}|s, \theta) = \frac{1}{1 + e^{-\theta_{\min(s, k)}}}, \quad s \in \mathbb{N}. \quad (45)$$

Taking $k = 0$ yields a static (i.e., state-independent) random policy, while letting k tend to infinity yields a fully state-dependent random policy. We believe this parametrization makes intuitive sense because, in a stable queueing system, small states tend to be visited more frequently than large states.

Under policy parametrization π_k , Assumptions 1 to 3 are satisfied with $n = d = k + 1$, $\Omega = \{\theta \in \mathbb{R}^{k+1} : \pi_k(\text{admit}|s, \theta) < \frac{\mu}{\lambda}\}$, $\Phi(s) = (\frac{\lambda}{\mu})^s$ for each $s \in \mathcal{S}$, $x_i(s) = \mathbb{1}[s \geq i + 1]$ for each $i \in \{0, 1, \dots, k - 1\}$ and $x_k(s) = \max(s - k, 0)$, and $\rho_i(\theta) = \pi_k(\text{admit}|i, \theta)$ for each $i \in \{0, 1, \dots, k\}$.

³In this example, vectors and matrices are indexed starting at 0 (instead of 1) for notational convenience.

It follows that $\nabla \log \rho_i = \nabla \log \pi_k(\text{admit}|i, \cdot)$ for each $i \in \{0, 1, \dots, k\}$. Assumption 7 can be satisfied by adding a small relative entropy regularization term as shown in Proposition 2. We refer to Appendix B.1 for further details.

Numerical results in a stable queue. We study the impact of the policy threshold $k \in \mathbb{N}$ on the performance of SAGE and actor-critic. The parameters are $\lambda = 0.7$, $\mu = 1$, $\gamma = 5$, and $\eta = 1$, and we consider random policies π_k with various thresholds. We have $\Omega = \mathbb{R}^{k+1}$ because $\lambda < \mu$, i.e., the queue is always stable. As we can verify using Appendix B.1, if $k \leq 2$ the best policy is random, while if $k \geq 3$, the best policy (deterministically) admits incoming jobs if and only if there are at most 2 jobs in the system. Thus, if $k \geq 3$, the best policy is approximated when $\theta_i \rightarrow +\infty$ if $i \in \{0, 1, 2\}$ and $\theta_i \rightarrow -\infty$ if $i \in \{3, 4, \dots, k\}$. This deterministic policy is optimal among all Markovian policies. The initial policy is $\pi_k(\text{admit}|s, \Theta_0) = \frac{1}{2}$ for each $s \in \mathbb{N}$, and the system is initially empty, i.e., $S_0 = 0$ with probability 1.




Figure 1: Long-run average reward $J(\Theta_t)$ in the admission-control problem with $\lambda = 0.7$, $\mu = 1$, $\gamma = 5$, and $\eta = 1$. Using Appendix B.1, we can verify that the long-run average reward under the best policy is approximately 2.183 if $k = 0$, 2.566 if $k = 1$, and 2.795 if $k \geq 3$.

Figure 1 depicts the impact of the threshold k on the evolution of the long-run average reward $J(\Theta_t)$ (defined in (7) and computed using the formulas of Appendix B.1) under SAGE and actor-critic. Figure 2 shows the admission probabilities $\pi_3(\text{admit}|i, \Theta_t)$ for each $i \in \{0, 1, 2, 3\}$ (i.e., the admission probabilities under the policy with threshold $k = 3$). In both plots, the x-axis has a logarithmic scale starting at time $t = 10^2$, lines are obtained by averaging the results over




Figure 2: Admission probabilities under policy parametrization π_3 .

10 independent simulations, and transparent areas show the standard deviation. Both SAGE and actor-critic eventually converge to the maximal attainable long-run average reward, and under both algorithms the convergence is initially faster under policy π_0 than under π_1 , π_3 , and π_{100} . For a particular threshold k , the convergence is initially faster under actor-critic than under SAGE. However, the long-run average reward under SAGE increases monotonically from its initial value to its maximal value while, under actor-critic, there is a time period (comprised between 10^3 and 10^5 time steps) where the long-run average reward stagnates or even decreases. Similar qualitative remarks can be made when looking at the running average reward $\frac{1}{t} \sum_{t'=1}^t R_{t'}$ instead of the long-run average reward $J(\Theta_t)$. Figure 2b suggests that, under π_3 , this is because actor-critic first “overshoots” by increasing $\pi_3(\text{admit}|3, \Theta_t)$ too much and then decreasing $\pi_3(\text{admit}|2, \Theta_t)$ too much before eventually converging to the best admission probabilities. This overshooting is more pronounced with a small threshold k , but it is still visible with $k = 100$.

Figures 1 and 2 suggest actor-critic has more difficulty to correctly estimate the policy update compared to SAGE, especially under parametrizations π_k with small thresholds k . We conjecture this is due to the combination of two phenomena which reaches a peak when k is small. First, a close examination of the evolution of the value function under π_3 and π_{10} (not shown here) reveals that there is a transitory bias in the estimate of the value function. For instance, right after increasing the admission probability in state 0, the estimate of the value function at states 2 and 3 becomes negative, even if the optimal value function at these states is positive. Second, due to the policy parametrization, parameter θ_k is updated whenever a state $s \in \{k, k+1, k+2, \dots\}$ is visited (while, for each $i \in \{0, 1, \dots, k-1\}$, parameter θ_i is updated only when state i is visited). As a result, the correlated biases in the estimates of the value function at states $k, k+1, k+2, \dots$ add up and lead actor-critic to overshoot the update of θ_k , which has a knock-on effect on other states.

Numerical results in a possibly-unstable queue. Figure 3 is the counterpart of Figure 1 when the arrival rate is $\lambda = 1.4 > 1 = \mu$. Now the set of policy parameters for which the system is stable is $\Omega = \{\theta \in \mathbb{R}^{k+1} : \pi_k(\text{admit}|k, \theta) < \frac{\mu}{\lambda}\} \subsetneq \mathbb{R}^{k+1}$, with $\frac{\mu}{\lambda} \simeq 0.714$. For simplicity, we will say that a policy is *stable* if the Markov chain defined by the system state under this policy is positive recurrent (i.e., if $\pi_k(\text{admit}|k, \theta) < \frac{\mu}{\lambda}$), and *unstable* otherwise. This is an example where convergence can only be guaranteed locally, as not all policies are stable. Again using Appendix B.1, we can verify that if $k \leq 1$, the best policy is random, while if $k \geq 2$, the best policy (deterministically)

admits incoming jobs if and only if there are fewer than 2 jobs in the system. This deterministic policy is optimal among all Markovian policies. The initial policy is again the (stable) uniform policy, and the system is initially empty, i.e., $S_0 = 0$ with probability 1.




Figure 3: Long-run average reward in the admission-control problem with parameters $\lambda = 1.4$, $\mu = 1$, $\gamma = 5$, and $\eta = 1$. Using Appendix B.1, we can verify that the maximal value of the long-run average reward is approximately 1.091 if $k = 0$ and 1.880 if $k \geq 2$.

The first take-away of Figure 3 is that SAGE converges to a close-to-optimal policy despite the fact that some policies are unstable. The convergence of SAGE is actually faster under $\lambda = 1.4$ compared to $\lambda = 0.7$ (Figure 1). By looking at the evolution of the admission probability (not shown here), we conjecture this is due to the fact that the admission probability in states larger than or equal to 2 decreases much faster when $\lambda = 1.4$ compared to $\lambda = 0.7$, and that this probability has a significant impact on the long-run average reward. In none of the simulations does SAGE reach an unstable policy. This suggests that the updates of SAGE have lower chance of reaching unstable regions of the policy space per observed sample.

The second take-away of Figure 3 is that, on the contrary, actor-critic has difficulties coping with instability in this example. In all simulation runs used to plot this figure, the long-run average reward $J(\Theta_t)$ first decreases before possibly increasing again and converging to the best achievable long-run average reward. Under parametrizations π_0 , π_2 , and π_4 , unstable policies are visited for thousands of steps in all simulation runs, and a stable policy is eventually reached in only 7 out of 10 runs. Under parametrization π_0 , the long-run average reward under the last policy is close to the best only in 2 out of 10 runs. Under π_{100} , the policy remains stable throughout all runs, but

the long-run average reward temporarily decreases before increasing again.

6.2 Load-balancing system

Consider a cluster of n servers. Jobs arrive according to a Poisson process with rate $\lambda > 0$, and a new job is admitted if and only if there are fewer than $c \in \mathbb{N}_+$ jobs in the system. Each server $i \in \{1, 2, \dots, n\}$ processes jobs in its queue according to a nonidling, nonanticipating policy. The service time of each job at server i is exponentially distributed with rate $\mu_i > 0$, independently of all other random variables. The agent aims to maximize the admission probability by adequately distributing load across servers.

For each $t \in \mathbb{N}$, let $S_t = (S_{t,1}, S_{t,2}, \dots, S_{t,n})$ denote the vector containing the number of jobs at each server right before the arrival of the $(t+1)$ th job, and let $A_t \in \{1, 2, \dots, n\}$ denote the server to which this $(t+1)$ th job is assigned. (This decision is void if $S_{t,1} + \dots + S_{t,n} = c$ because the job is rejected anyway.) We have $\mathcal{S} = \{s \in \mathbb{N}^n : s_1 + s_2 + \dots + s_n \leq c\}$ and $\mathcal{A} = \{1, 2, \dots, n\}$. The agent obtains a reward of 1 if the job is accepted and 0 otherwise, that is, $R_{t+1} = \mathbb{1}[S_{t,1} + \dots + S_{t,n} \leq c-1]$ for each $t \in \mathbb{N}$.

We consider the following static policy parametrization, with parameter vector $\theta \in \mathbb{R}^n$: irrespective of the system state $s \in \mathcal{S}$, an incoming job is assigned to server i with probability

$$\pi(i|s, \theta) = \pi(i|\theta) = \frac{e^{\theta_i}}{\sum_{j=1}^n e^{\theta_j}}, \quad i \in \{1, 2, \dots, n\}. \quad (46)$$

Assumptions 1 to 3 are satisfied with $n = d$, $\Omega = \mathbb{R}^n$, $\Phi(s) = \prod_{i=1}^n (\frac{\lambda}{\mu_i})^{s_i}$ for each $s \in \mathcal{S}$, $x_i(s) = s_i$ for $i \in \{1, 2, \dots, n\}$ and $s \in \mathcal{S}$, and $\rho_i(\theta) = \pi(i|\theta)$ for $i \in \{1, 2, \dots, n\}$ and $\theta \in \mathbb{R}^n$. Also note that $\nabla \log \rho_i(\theta) = \nabla \log \pi(i|\theta)$ for $i \in \{1, 2, \dots, d\}$ and $\theta \in \mathbb{R}^n$. Except for Assumption 7, the remaining assumptions outlined in Section 5 are also satisfied. We refer to Appendix B.2 for more details. Lastly observe that, in spite of the policy being static and the state space being finite, the function J is still nonconvex for typical system parameters. In fact, our numerical experiments are done in nonconvex scenarios. Furthermore, note that this system can become challenging to optimize if c and n are large.

Numerical results We study the performance of SAGE and actor-critic under varying numbers of servers and service speed imbalance. Given an integer $n \in \mathbb{N}_{>0}$ multiple of 4 and $\delta > 1$, we consider the following cluster of n servers divided into 4 pools. For each $k \in \{1, 2, 3, 4\}$, pool k consists of the $\frac{n}{4}$ servers indexed from $(k-1)\frac{n}{4} + 1$ to $k\frac{n}{4}$, and each server i in this pool has service rate $\mu_i = \delta^{k-1}$. The total arrival rate is $\lambda = 0.7(\sum_{i=1}^n \mu_i)$ and the upper bound on the number of jobs in the system is $c = 10\frac{n}{4}$. Letting $\delta = 1$ gives a system where all servers have the same service speed, while increasing δ makes the server speeds more and more imbalanced. The initial policy is uniform, i.e., $\pi(i|\Theta_0) = \frac{1}{n}$ for each $i \in \{1, 2, \dots, n\}$, and the initial state is empty, i.e., $S_0 = 0$ with probability 1.

Figure 4 shows performance of SAGE and actor-critic in clusters of $n \in \{4, 20, 100\}$ servers. Solid lines show the evolution of the long-run average reward $J(\Theta_t)$, and dashed lines show the running average reward $\frac{1}{t} \sum_{\nu=1}^t R_\nu$. (Recall $J(\Theta_t)$ is the limit of the running average we would see if we ran the system under policy $\pi(\Theta_t)$. It is defined in (7) and can be computed as shown in Appendix B.2.) As before, transparent areas show the standard deviation around the average. The results under actor-critic are reported only for $n = 4$ servers, as this method already suffers

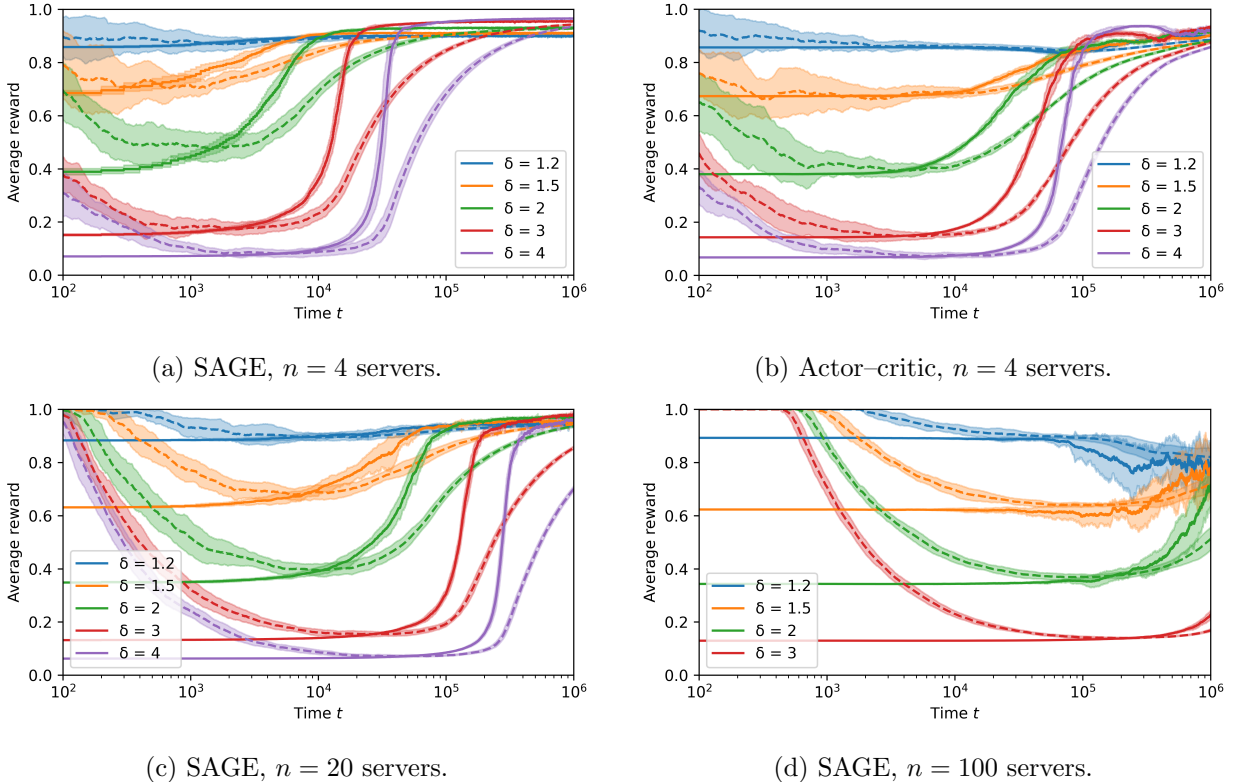


Figure 4: Impact of the number of servers and service-rate imbalance on the performance of SAGE and actor-critic in a load-balancing system. Solid lines show the long-run average reward $J(\Theta_t)$, while dashed lines show the running average reward, $\frac{1}{t} \sum_{t'=1}^t R_{t'}$. Simulations for $n = 100$ and $\delta = 4$ are omitted because numerical instability of Buzen’s algorithm (see Appendix B.2) prevents us from computing $J(\Theta_t)$ in this case.

from a combinatorial explosion in the state-action space for $n \in \{20, 100\}$. Indeed, while the memory complexity increases linearly with the number n of servers under SAGE, it increases with the cardinality $\binom{n+c}{c}$ of the state space under actor-critic⁴, which is already prohibitively large for $n \in \{20, 100\}$.

All four subfigures in Figure 4 show a consistent 2-phase pattern: first the running average reward $\frac{1}{t} \sum_{t'=1}^t R_{t'}$ converges to the initial long-run average reward $J(\Theta_0)$, and then the long-run average reward increases to reach the best value, with the running average reward catching up at a slower pace. This suggests that the gradient estimates under both algorithms remain close to zero until the system reaches approximate stationarity. A similar reasoning explains why the algorithms converge at a slower pace when we increase the imbalance factor δ (as the stationary distribution under the initial uniform policy $\pi(\Theta_0)$ puts mass on states that are further away from the initial empty state) or the number n of servers (as the mixing time increases).

Focusing on the system with $n = 4$ servers, Figures 4a and 4b show convergence occurs on the order of 10^5 time steps sooner under SAGE than under actor-critic. We conjecture that this is again due to the fact that actor-critic relies on estimating the state-value function, so that it

⁴As shown by applying the *stars and bars* method in combinatorics.

needs to estimate $\binom{n+c}{c}$ values. SAGE, on the other hand, exploits the structure of the stationary distribution and only needs to estimate a number of values that grows linearly with n (and is independent of c). We also note that actor-critic shows nonmonotonic convergence (i.e., $J(\Theta_t)$ decreases before increasing again between 10^5 and 10^6 time steps). We conjecture this is due to a similar phenomenon as described in Section 6.1.

6.3 Ising model and Glauber dynamics

Consider a system of spin particles spread over a two-dimensional lattice of shape $d_1 \times d_2$, for some $d_1, d_2 \in \{2, 3, 4, \dots\}$. Let $\mathcal{V} = \{1, 2, \dots, d_1\} \times \{1, 2, \dots, d_2\}$ denote the set of lattice coordinates. For any two coordinates $v = (v_1, v_2) \in \mathcal{V}$ and $w = (w_1, w_2) \in \mathcal{V}$, we write $v \sim w$ if and only if v and w are neighbors in the lattice, that is, if and only if $|v_1 - w_1| + |v_2 - w_2| = 1$.

A map $\sigma : \mathcal{V} \rightarrow \{-1, +1\}$ is called a *spin configuration*, and the set of all $2^{d_1 d_2}$ configurations is denoted by Σ . Given a configuration $\sigma \in \Sigma$, we refer to $\sigma(v) \in \{-1, +1\}$ as the *spin* (of the particle located) at v . If the system is in some configuration $\sigma \in \Sigma$, we say that the spin at $v \in \mathcal{V}$ is *flipped* if the system jumps to the configuration $\sigma_{-v} \in \Sigma$ such that $\sigma_{-v}(v) = -\sigma(v)$ and $\sigma_{-v}(w) = \sigma(w)$ for each $w \in \mathcal{V} \setminus \{v\}$. As we will formalize below, the agent's goal is to reach a configuration $\sigma \in \Sigma$ so that the *magnetization* on the left (resp. right) half of the lattice is close to $\xi_{\text{left}} \in (-1, +1)$ (resp. $\xi_{\text{right}} \in (-1, +1)$), i.e.,

$$\frac{2}{d_1 d_2} M_{\text{left}}(\sigma) \simeq \xi_{\text{left}}, \quad \frac{2}{d_1 d_2} M_{\text{right}}(\sigma) \simeq \xi_{\text{right}},$$

where

$$M_{\text{left}}(\sigma) = \sum_{v \in \mathcal{V}: v_2 \leq d_2/2} \sigma(v), \quad M_{\text{right}}(\sigma) = \sum_{v \in \mathcal{V}: v_2 > d_2/2} \sigma(v).$$

To each configuration $\sigma \in \Sigma$ is associated an *energy* $E(\sigma) \triangleq -JI(\sigma) - \mu F(\sigma)$, where I and F are called the *interaction* and *external field* terms, respectively, given by

$$I(\sigma) = \sum_{v, w \in \mathcal{V}: v \sim w} \sigma(v)\sigma(w), \quad F(\sigma) = \sum_{v \in \mathcal{V}} h(v)\sigma(v),$$

where the first sum runs over all pairs of neighboring coordinates (so that each pair appears once). Here, $J \in \mathbb{R}$ is the *coupling constant*, $\mu \in \mathbb{R}_{\geq 0}$ the *magnetic moment*, and $h : \mathcal{V} \rightarrow \mathbb{R}$ the *external magnetic field*. Under the dynamics defined below, the probability of a configuration $\sigma \in \Sigma$ will be proportional to $e^{-\beta E(\sigma)}$, where $\beta \in \mathbb{R}_{> 0}$ is the *inverse temperature*. If $J > 0$ (resp. $J < 0$), the interaction term I contributes to increasing the probability of configurations where neighboring spins have the same (resp. opposite) sign. Concurrently, due to the external-field term F , the spin at each $v \in \mathcal{V}$ is attracted in the direction pointed by the sign of $h(v)$. The coupling constant J and magnetic moment μ are fixed and known by the agent (as they depend on the particles), and the agent will fine-tune the inverse temperature β and coarse-tune the external magnetic field h .

Glauber dynamics Given a starting configuration, at every time step, the spin at a coordinate chosen uniformly at random is flipped (or not) with some probability that depends on the current configuration and the parameters set by the agent. This is cast as a Markov decision process as follows. The state and action spaces are given by $\mathcal{S} = \Sigma \times \mathcal{V}$ and $\mathcal{A} = \{\text{flip}, \text{not flip}\}$, respectively.

For each $s = (\sigma, v) \in \mathcal{S}$ and $a \in \mathcal{A}$, the state reached by taking action a in state s is given by $S' = (\sigma', V')$, where $\sigma' = \sigma_{-v}$ if $a = \text{flip}$ and $\sigma' = \sigma$ if $a = \text{not flip}$, and V' is chosen uniformly at random in \mathcal{V} , independently of the past states, actions, and rewards. The next reward r is the opposite of the sum of the absolute difference between the next magnetizations and the desired magnetizations, that is,

$$r = -\left| \xi_{\text{left}} - \frac{2}{d_1 d_2} M_{\text{left}}(\sigma') \right| - \left| \xi_{\text{right}} - \frac{2}{d_2 d_2} M_{\text{right}}(\sigma') \right|.$$

The agent controls a vector $\theta \in \mathbb{R}^3$ that determines the inverse temperature and the left and right external magnetic fields as follows:

$$\beta(\theta) = 1 + \tanh(\theta_1), \quad h_{\text{left}}(\theta) = \tanh(\theta_2), \quad h_{\text{right}}(\theta) = \tanh(\theta_3),$$

so that in particular $\beta(\theta) \in (0, 2)$, $h_{\text{left}}(\theta) \in (-1, 1)$, and $h_{\text{right}}(\theta) \in (-1, 1)$. The corresponding external magnetic field and external field term are

$$\begin{aligned} h(v|\theta) &= h_{\text{left}}(\theta) \mathbf{1}[v_2 \leq d_2/2] + h_{\text{right}}(\theta) \mathbf{1}[v_2 > d_2/2], \quad v \in \mathcal{V}, \\ F(\sigma|\theta) &= \sum_{v \in \mathcal{V}} h(v|\theta) \sigma(v) = h_{\text{left}}(\theta) M_{\text{left}}(\sigma) + h_{\text{right}}(\theta) M_{\text{right}}(\sigma), \quad \sigma \in \Sigma. \end{aligned}$$

Given $\theta \in \mathbb{R}^3$, for each $s = (\sigma, v) \in \Sigma$, the probability that the spin at the randomly-chosen coordinate v is flipped when the current configuration is σ is given by

$$\pi(\text{flip}|s, \theta) = \frac{1}{1 + e^{\delta(s|\theta)}}, \quad \text{with} \quad \delta(s|\theta) = 2\beta(\theta)\sigma(v) \left(J \sum_{w \in \mathcal{V}: w \sim v} \sigma(w) + \mu h(v|\theta) \right). \quad (47)$$

When $\theta \in \mathbb{R}^3$ is fixed, the dynamics defined by this system are called the Glauber dynamics [24, Section 3.3]. Note that, although we use the word *action* to match the terminology of MDPs, here an action should be seen as a random event in the environment, of which only the distribution π can be controlled by the agent via the parameter vector θ .

Product-form distribution We verify in Appendix B.3 that the stationary distribution of the system state under a particular choice of $\theta \in \mathbb{R}^3$ satisfies

$$p(s|\theta) \propto e^{\beta(\theta)(JI(\sigma) + \mu F(\sigma|\theta))}, \quad s = (\sigma, v) \in \mathcal{S}, \quad \theta \in \mathbb{R}^3. \quad (48)$$

Assumptions 1 to 3 are satisfied with $n = d = 3$, $\Omega = \mathbb{R}^3$, $\Phi(s) = 1$ for each $s \in \mathcal{S}$, $\log \rho_1(\theta) = \beta(\theta)J$, $\log \rho_2(\theta) = \beta(\theta)\mu h_{\text{left}}(\theta)$, and $\log \rho_3(\theta) = \beta(\theta)\mu h_{\text{right}}(\theta)$ for each $\theta \in \mathbb{R}^3$, and $x_1(s) = I(\sigma)$, $x_2(s) = M_{\text{left}}(\sigma)$, and $x_3(s) = M_{\text{right}}(\sigma)$ for each $s = (\sigma, v) \in \mathcal{S}$. All derivations are given in Appendix B.3.

Numerical results Figure 5 shows the performance of SAGE in a system with parameters $d_1 = 10$, $d_2 = 20$, $J = \mu = 1$, $\xi_{\text{left}} = -1$, and $\xi_{\text{right}} = 1$. We do not run simulations under actor-critic, as again the state space has size $2^{d_1 d_2} = 2^{200}$, which is out of reach for this method. The initial parameter vector is $\Theta_0 = 0$, yielding inverse temperature $\beta(\Theta_0) = 1$ and external fields $h_{\text{left}}(\Theta_0) = h_{\text{right}}(\Theta_0) = 0$. The initial configuration has spins 1 on the left-hand side and -1 on the

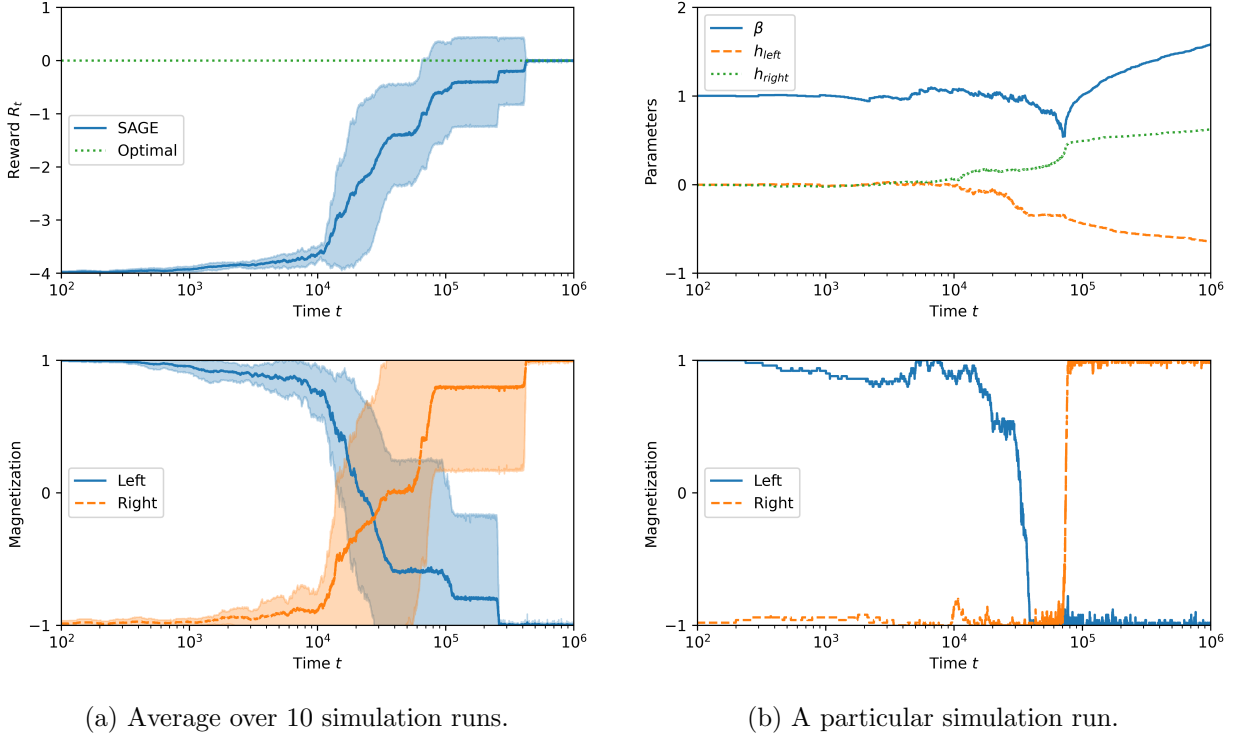


Figure 5: Performance of SAGE in the Ising model.

right-hand side, so that reaching the target configuration requires flipping every spin. In Figure 5a, the reward R_t seems to increase on average monotonically from -4 to 0 , which is consistent with the observation that the left (resp. right) magnetization decreases from 1 to -1 (resp. increases from -1 to 1). The increase of the reward is stepwise, with stages where it remains roughly constant for several thousand time steps. Lastly, the standard deviation increases significantly from about 10^4 to $3 \cdot 10^5$ time steps, and it becomes negligible afterwards.

To help us understand these observations, Figure 5b shows the evolution of the system parameters and of the magnetizations over a particular simulation run. The left magnetization M_{left} starts decreasing around 10^4 time steps (bottom plot), approximately when $h_{\text{left}}(\Theta_t)$ and $h_{\text{right}}(\Theta_t)$ become nonzero (top plot), to become roughly -1 around $3 \cdot 10^4$ time steps. At that moment, the system configuration is close to the all-spin-down configuration σ_{-1} such that $\sigma_{-1}(v) = -1$ for each $v \in \mathcal{V}$. The right magnetization starts increasing only when the inverse temperature $\beta(\Theta_t)$ has a sudden decrease (top plot). To make sense of this observation, consider $\pi(\text{flip}|s, \theta)$ as given by (47), where $s = (\sigma_{-1}, v)$ for some $v \in \{2, 3, \dots, d_1 - 1\} \times \{2, 3, \dots, d_2 - 1\}$. In $\delta(s|\theta)$, the first term $J \sum_{w \in \mathcal{V}: w \sim v} \sigma_{-1}(v) \sigma_{-1}(w)$ is equal to 4 , while the absolute value of the second term $\mu \sigma_{-1}(v) h(v|\theta)$ is at most 1 ; hence, if $\beta(\theta) \simeq 1$ as initially, $\pi(\text{flip}|s, \theta)$ is between $\frac{1}{1+e^{2(4+1)}} \simeq 4.5 \cdot 10^{-5}$ and $\frac{1}{1+e^{2(4-1)}} \simeq 2.5 \cdot 10^{-3}$. The brief decrease of $\beta(\theta)$ is an efficient way of increasing the flipping probability in all states, which allows the system to escape from σ_{-1} . Other simulation runs are qualitatively similar, but the times at which the qualitative changes occur and the side (left or right) that flips magnetization first vary, which explains the large standard deviation observed earlier.

7 Conclusion

In this paper, we incorporated model-specific information about MDPs into the gradient estimator in policy-gradient methods. Specifically, assuming that the stationary distribution is an exponential family, we derived score-aware gradient estimators (SAGEs) that do not require the computation of value functions (Theorem 1). As showcased in Section 6, this assumption is satisfied by models from stochastic networks, where the stationary distribution possesses a product-form structure, and by models from statistical mechanics, such as the Ising models with Glauber dynamics.

The numerical results in Section 6 show that in these systems, policy-gradient algorithms equipped with a SAGE outperform actor-critic. In these examples, the Jacobian of the load function $D \log \rho(\theta)$ can be computed explicitly in terms of the policy parameter θ . However, SAGE estimators can be harder to compute in more complex cases, for example when $D \log \rho(\theta)$ depends on some model parameters. Nevertheless, our examples showcase how it is possible to improve the current policy gradient methods by leveraging information on the MDP, and we expect extensions of SAGEs to cover more challenging cases, for example by combining SAGE with model selection by first estimating the model parameters appearing in $D \log \rho(\theta)$. We leave such extensions of SAGE for future work.

We have also shown with Theorem 2 that policy gradient with SAGE converges to the optimal policy under light assumptions, namely, the existence of a local Lyapunov function close to the optimum, which allows for unstable policies to exist, and a nondegeneracy property of the Hessian at maxima. The convergence occurs with a probability arbitrarily close to one provided that the iterates start close enough. In Corollary 1, the sample complexity of the algorithm is shown to be of order $O(T^{-1/(3+\epsilon)} + \alpha^2/\ell)$, where T is the number of samples drawn. Unlike most common convergence results, our assumptions are weak enough to allow for unstable policies to exist and thus, there is an unavoidable nonzero probability that the algorithm will be unstable. This fact is captured by the term α^2/ℓ . Remarkably, such instabilities are observed in one of the examples of Section 6. If we had made stronger assumptions such as the existence of a global Lyapunov function, then such phenomena would not have been captured by the analysis.

References

- [1] I. Adan, A. Bušić, J. Mairesse, and G. Weiss. Reversibility and further properties of FCFS infinite bipartite matching. *Mathematics of Operations Research*, 43(2):598–621, dec 2017. Publisher: INFORMS.
- [2] A. Agarwal, S. M. Kakade, J. D. Lee, and G. Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *The Journal of Machine Learning Research*, 22(1):4431–4506, 2021.
- [3] J. Anselmi, B. Gaujal, and L.-S. Rebuffi. Learning optimal admission control in partially observable queueing networks. *arXiv preprint arXiv:2308.02391*, 2023.
- [4] Y. F. Atchadé, G. Fort, and E. Moulines. On perturbed proximal gradient algorithms. *The Journal of Machine Learning Research*, 18(1):310–342, 2017.
- [5] F. Baskett, K. M. Chandy, R. R. Muntz, and F. G. Palacios. Open, closed, and mixed networks of queues with different classes of customers. *Journal of the ACM*, 22(2):248–260, apr 1975.

- [6] T. Bonald and J. Virtamo. Calculating the flow level performance of balanced fairness in tree networks. *Performance Evaluation*, 58(1):1–14, oct 2004.
- [7] P. Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Texts in Applied Mathematics. Springer-Verlag, 1999.
- [8] J. P. Buzen. Computational algorithms for closed queueing networks with exponential servers. *Communications of the ACM*, 16(9):527–531, sep 1973.
- [9] H. Daneshmand, J. Kohler, A. Lucchi, and T. Hofmann. Escaping saddles with stochastic gradients. In *International Conference on Machine Learning*, pages 1155–1164. PMLR, 2018.
- [10] E. de Souza e Silva and M. Gerla. Queueing network models for load balancing in distributed systems. *Journal of Parallel and Distributed Computing*, 12(1):24–38, may 1991.
- [11] E. de Souza e Silva and R. R. Muntz. Simple relationships among moments of queue lengths in product form queueing networks. *IEEE Transactions on Computers*, 37(9):1125–1129, sep 1988.
- [12] T. T. Doan, L. M. Nguyen, N. H. Pham, and J. Romberg. Finite-time analysis of stochastic gradient descent under Markov randomness. *arXiv preprint arXiv:2003.10973*, 2020.
- [13] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pages 1467–1476. PMLR, 2018.
- [14] B. Fehrman, B. Gess, and A. Jentzen. Convergence rates for the stochastic gradient descent method for non-convex objective functions. *Journal of Machine Learning Research*, 21:136, 2020.
- [15] G. Fort and E. Moulines. Convergence of the Monte Carlo expectation maximization for curved exponential families. *The Annals of Statistics*, 31(4):1220–1259, 2003.
- [16] K. Gardner and R. Righter. Product forms for FCFS queueing models with arbitrary server-job compatibilities: an overview. *Queueing Systems*, 96(1):3–51, oct 2020.
- [17] V. Guillemin and A. Pollack. *Differential topology*, volume 370. American Mathematical Soc., 2010.
- [18] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Transactions on Networking*, 18(3):960–972, 2009.
- [19] B. Karimi, B. Miasojedow, E. Moulines, and H.-T. Wai. Non-asymptotic analysis of biased stochastic approximation scheme. In *Conference on Learning Theory*, pages 1944–1974. PMLR, 2019.
- [20] S. Khadka and K. Tumer. Evolution-guided policy gradient in reinforcement learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

- [21] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, Cambridge, MA, 2009.
- [22] N. Kumar, Y. Murthy, I. Shufaro, K. Y. Levy, R. Srikant, and S. Mannor. On the global convergence of policy gradient in average reward Markov decision processes. *arXiv preprint arXiv:2403.06806*, 2024.
- [23] J. M. Lee. *Introduction to Smooth Manifolds*, volume 218. Springer Science & Business Media, 2013.
- [24] D. A. Levin and Y. Peres. *Markov Chains and Mixing Times*. American Mathematical Society, 2e édition edition.
- [25] B. Liu, Q. Xie, and E. Modiano. RL-QN: A reinforcement learning framework for optimal control of queueing systems. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems*, 7(1):1–35, 2022.
- [26] Z. Liu and P. Nain. Sensitivity results in open, closed and mixed product form queueing networks. *Performance Evaluation*, 13(4):237–251, 1991.
- [27] S. Mohamed, M. Rosca, M. Figurnov, and A. Mnih. Monte Carlo gradient estimation in machine learning. *The Journal of Machine Learning Research*, 21(1):5183–5244, 2020.
- [28] P. Moyal, A. Bušić, and J. Mairesse. A product form for the general stochastic matching model. *Journal of Applied Probability*, 58(2):449–468, jun 2021. Publisher: Cambridge University Press.
- [29] L. I. Nicolaescu et al. *An invitation to Morse theory*. Springer, 2011.
- [30] Y. Qian, J. Wu, R. Wang, F. Zhu, and W. Zhang. Survey on reinforcement learning applications in communication networks. *Journal of Communications and Information Networks*, 4(2):30–39, 2019.
- [31] J. Sanders, S. C. Borst, and J. S. H. van Leeuwen. Online network optimization using product-form Markov processes. *IEEE Transactions on Automatic Control*, 61(7):1838–1853, 2016.
- [32] R. Serfozo. *Introduction to Stochastic Networks*. Stochastic Modelling and Applied Probability. Springer-Verlag, 1999.
- [33] D. Shah. Message-passing in stochastic processing networks. *Surveys in Operations Research and Management Science*, 16(2):83–104, jul 2011.
- [34] V. Shah and G. de Veciana. High-performance centralized content delivery infrastructure: Models and asymptotics. *IEEE/ACM Transactions on Networking*, 23(5):1674–1687, oct 2015.
- [35] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT press, Cambridge, MA, USA, 2 edition, 2018.
- [36] V. B. Tadic and A. Doucet. Asymptotic bias of stochastic gradient search. *Annals of Applied Probability*, 27(6):3255–3304, 2017.

- [37] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.*, 1(1):1–305, 2018.
- [38] R. W. Wolff. Poisson arrivals see time averages. *Operations Research*, 30(2):223–231, apr 1982. Publisher: INFORMS.
- [39] L. Xiao. On the convergence rates of policy gradient methods. *The Journal of Machine Learning Research*, 23(1):12887–12922, 2022.
- [40] R. Yuan, R. M. Gower, and A. Lazaric. A general sample complexity analysis of vanilla policy gradient. In *International Conference on Artificial Intelligence and Statistics*, pages 3332–3380. PMLR, 2022.
- [41] S. Zachary and I. Ziedins. Loss networks and Markov random fields. *Journal of Applied Probability*, 36(2):403–414, jun 1999. Publisher: Cambridge University Press.

A Policy-gradient algorithms

A.1 Actor–critic algorithm

The actor–critic algorithm is first mentioned in Section 3.4 and compared to our SAGE-based policy-gradient algorithm in Section 6. We focus on the version of actor–critic described in [35, Section 13.6] for the average-reward criterion in infinite horizon. The algorithm relies on the following expression for $\nabla J(\theta)$, which is a variant of the *policy-gradient theorem* [35, Chapter 13]:

$$\nabla J(\theta) \propto \mathbb{E}[(R - J(\theta) + v(S') - v(S))\nabla \log \pi(A|S, \theta)],$$

where (S, A, R, S') is a quadruplet of random variables such that $S \sim p(\cdot | \theta)$, $A|S \sim \pi(\cdot | S, \theta)$, and $(R, S')|(S, A) \sim P(\cdot, \cdot | S, A)$ (so that in particular $(S, A, R) \sim \text{STAT}(\theta)$), and v is the state-value function.

Algorithm 3 Actor–critic algorithm [35, Section 13.6] to be called on Line 9 of Algorithm 1, with batch sizes equal to one.

- 1: **Input:** Positive and differentiable policy parametrization $(s, \theta, a) \mapsto \pi(a|s, \theta)$
 - 2: **Parameters:** Step sizes $\alpha_{\bar{R}} > 0$ and $\alpha_v > 0$
 - 3: **Initialization:**
 - $\bar{R} \leftarrow 0$
 - $V[s] \leftarrow 0$ for each $s \in \mathcal{S}$
 - 4: **procedure** GRADIENT(t)
 - 5: $\delta \leftarrow R_{t+1} - \bar{R} + V[S_{t+1}] - V[S_t]$
 - 6: Update $\bar{R} \leftarrow \bar{R} + \alpha_{\bar{R}}\delta$
 - 7: Update $V[S_t] \leftarrow V[S_t] + \alpha_v\delta$
 - 8: **return** $\delta \nabla \log \pi(A_t|S_t, \Theta_t)$
 - 9: **end procedure**
-

The pseudocode of the procedure GRADIENT used in the actor–critic algorithm is given in Algorithm 3. This procedure is to be implemented within Algorithm 1 with batch sizes equal to one, meaning that $t_{m+1} = t_m + 1$ for each $m \in \mathbb{N}$. We assume for simplicity that all variables

from Algorithm 1 are accessible inside Algorithm 3. The variable \bar{R} updated on Line 6 is a biased estimate of $J(\Theta_m)$, while the table V updated on Line 7 is a biased estimate of the state-value function under policy $\pi(\Theta_m)$. Compared to [35, Section 13.6], the value function is encoded by a table V and there are no eligibility traces. If the state space \mathcal{S} is infinite, the table V is initialized at zero over a finite subset of \mathcal{S} containing the initial state S_0 and expanded with zero padding whenever necessary.

A.2 SAGE-based policy-gradient method

Algorithm 4 SAGE-updated policy-gradient method, to be called on Line 9 of Algorithm 1.

- 1: **Input:**
 - Positive and differentiable policy parametrization $(s, \theta, a) \mapsto \pi(a|s, \theta)$
 - Jacobian matrix function $\theta \mapsto D \log \rho(\theta)$
 - Feature function $s \mapsto x(s)$
 - 2: **Parameters:** Memory factor $\nu \in [0, 1]$
 - 3: **Initialization:** Global variables $N_{-1}, M_{-1}, \bar{X}_{-1}, \bar{R}_{-1}, \bar{C}_{-1}, \bar{E}_{-1} \leftarrow 0$
 - 4: **procedure** GRADIENT(m)
 - 5: $N_m \leftarrow \nu N_{m-1} + (t_{m+1} - t_m)$
 - 6: $M_m \leftarrow \nu^2 M_{m-1} + (t_{m+1} - t_m)$
 - 7: **return** $D \log \rho(\Theta_m) \text{COVARIANCE}(m) + \text{EXPECTATION}(m)$
 - 8: **end procedure**
 - 9: **procedure** COVARIANCE(m)
 - 10: Update $\bar{X}_m \leftarrow \nu \bar{X}_{m-1} + \sum_{t=t_m}^{t_{m+1}-1} x(S_t)$
 - 11: Update $\bar{R}_m \leftarrow \nu \bar{R}_{m-1} + \sum_{t=t_m}^{t_{m+1}-1} R_{t+1}$
 - 12: Update $\bar{C}_m \leftarrow \nu \bar{C}_{m-1} + \sum_{t=t_m}^{t_{m+1}-1} (x(S_t) - \frac{1}{N_m} \bar{X}_m)(R_{t+1} - \frac{1}{N_m} \bar{R}_m)$
 - 13: **return** $\frac{N_m}{N_m^2 - M_m} \bar{C}_m$ if $N_m^2 > M_m$ else $\frac{1}{N_m} \bar{C}_m$
 - 14: **end procedure**
 - 15: **procedure** EXPECTATION(m)
 - 16: Update $\bar{E}_m \leftarrow \nu \bar{E}_{m-1} + \sum_{t=t_m}^{t_{m+1}-1} R_{t+1} \nabla \log \pi(A_t|S_t, \Theta_m)$
 - 17: **return** $\frac{1}{N_m} \bar{E}_m$
 - 18: **end procedure**
-

Algorithm 4 is an extension of Algorithm 2 that allows for batches of size 1. The main advantage of Algorithm 4 over Algorithm 2 is that it estimates $\nabla J(\Theta_m)$ based not only on batch \mathcal{D}_m , but also on previous batches, depending on the memory factor ν initialized on Line 2. To simplify the signature of procedures in Algorithm 4, we assume variables $N_m, M_m, \bar{X}_m, \bar{R}_m, \bar{C}_m$, and \bar{E}_m are global, and that all variables from Algorithm 1 are accessible within Algorithm 4, in particular batch \mathcal{D}_m . Line 5 in Algorithm 4 is the counterpart of Line 3 in Algorithm 2. Line 6 in Algorithm 4 is used on Line 13 to compute the counterpart of the quotient $1/(N_m - 1)$ in Line 6 in Algorithm 2. The COVARIANCE(m) procedure in Algorithm 4 is the counterpart of Lines 4–6 in Algorithm 2. The EXPECTATION(m) procedure in Algorithm 4 is the counterpart of Line 7 in Algorithm 2. Algorithm 2 can be seen as a special case of Algorithm 4 with memory factor $\nu = 0$. Note that terminology in Algorithm 4 differs slightly compared to Algorithm 2: bar notation refers to cumulative sums instead of averages.

The subroutines `COVARIANCE` and `EXPECTATION` compute biased covariance and mean estimates for $\text{Cov}[R, x(S)]$ and $\mathbb{E}[R \nabla \log \pi(A|S, \theta)]$, where $(S, A, R) \sim \text{STAT}(\Theta_m)$, consistently with Theorem 1. If the memory factor ν is zero, these procedures return the usual sample mean and covariance estimates taken over the last batch \mathcal{D}_m (as in Algorithm 2), and bias only comes from the fact that the system is not stationary. If ν is positive, estimates from previous batches are also taken into account, so that the bias is increased in exchange for a (hopefully) lower variance. In this case, the updates on Lines 10–12 and 16 calculate iteratively the weighted sample mean and covariance over the whole history, where observations from epoch $m - \underline{m}$ have weight $\nu^{\underline{m}}$, for each $\underline{m} \in \{0, 1, \dots, m\}$. When m is large, the mean returned by `EXPECTATION` is approximately equal to the sample mean over batches \mathcal{D}_{m-M} through \mathcal{D}_m , where M is a truncated geometric random variable, independent of all other random variables, such that $\mathbb{P}[M = \underline{m}] \propto \nu^{\underline{m}}$ for each $\underline{m} \in \{0, 1, \dots, m\}$; if batches have constant size c , then we take into account approximately the last $c(\mathbb{E}[M] + 1) = \frac{c}{1-\nu}$ steps.

B Examples

This appendix provides detailed derivations for the examples of Section 6. We consider the single-server queue with admission control of Section 6.1 in Appendix B.1, the load-balancing example of Section 6.2 in Appendix B.2, and the Ising model of Section 6.3 in Appendix B.3.

B.1 Single-server queue with admission control

Consider the example of Section 6.1, where jobs arrive according to a Poisson process with rate $\lambda > 0$, and service times are exponentially distributed with rate $\mu > 0$. Recall the long-run average reward is the difference between an admission reward proportional to the admission probability and a holding cost proportional to the mean queue size. We first verify that Assumptions 1 to 3 are satisfied, then we give a closed-form expression for the objective function, and lastly we discuss the assumptions of Section 5. We consider a random threshold-based policy π_k of the form (45) for some $k \in \mathbb{N}$ and some parameter $\theta \in \Omega$, where $\Omega = \{\theta \in \mathbb{R}^{k+1} : \pi_k(\text{admit}|k, \theta) < \frac{\mu}{\lambda}\}$.

Product-form stationary distribution. The evolution of the number of jobs in the system (either waiting or in service) defines a birth-and-death process with birth rate $\lambda \pi_k(\text{admit}|s, \theta)$ and death rate $\mu \mathbb{1}[s \geq 1]$ in state s , for each $s \in \{0, 1, 2, \dots\}$. This process is irreducible because its birth and death rates are positive, and it is positive recurrent because $\lambda \pi_k(\text{admit}|s, \theta) < \mu$ for each $s \in \mathcal{S}$ by definition of Ω . This verifies Assumption 1. The stationary distribution is given by

$$\begin{aligned} p(s|\theta) &= \frac{1}{Z(\theta)} \prod_{q=0}^{s-1} \left(\frac{\lambda}{\mu} \pi(\text{admit}|q, \theta) \right), \\ &= \frac{1}{Z(\theta)} \left(\frac{\lambda}{\mu} \right)^s \left[\prod_{i=0}^{k-1} \pi_k(\text{admit}|i, \theta) \mathbb{1}^{[s \geq i+1]} \right] \pi_k(\text{admit}|k, \theta)^{\max(s-k, 0)}, \quad s \in \mathbb{N}, \end{aligned} \quad (49)$$

where the second equality follows by injecting (45), and the value of $Z(\theta)$ follows by normalization. We recognize (10–PF) from Assumption 3, with $n = d = k + 1$, $\Phi(s) = \left(\frac{\lambda}{\mu}\right)^s$ for each $s \in \mathcal{S}$, $x_i(s) = \mathbb{1}[s \geq i + 1]$ for each $i \in \{0, 1, \dots, k - 1\}$ and $x_k(s) = \max(s - k, 0)$ for each $s \in \mathcal{S}$, and $\rho_i(\theta) = \pi_k(\text{admit}|i, \theta)$ for each $i \in \{0, 1, \dots, k\}$. The function ρ defined in this way is differentiable.

Assumption 3 is therefore satisfied, as the distribution of the system seen at arrival times is also (49) according to the PASTA property [38]. For each $s \in \mathbb{N}$ and $a \in \{\text{admit}, \text{reject}\}$, $\nabla \log \pi_k(a|s, \theta)$ is the $(k+1)$ -dimensional column vector with value $\mathbf{1}[a = \text{admit}] - \pi_k(\text{admit}|i, \theta)$ in component $i = \min(s, k)$ and zero elsewhere, and $D \log \rho(\theta)$ is the $(k+1)$ -dimensional diagonal matrix with diagonal coefficient $1 - \pi_k(\text{admit}|i, \theta)$ in position i , for each $i \in \{0, 1, \dots, k\}$. This can be used to verify that Assumption 5 is satisfied.

Objective function. The objective function is $J(\theta) = \gamma \mathbb{P}[A = \text{admit}] - \frac{\eta}{\lambda} \mathbb{E}[S]$, where

$$\begin{aligned} \mathbb{P}[A = \text{admit}] &= \sum_{i=0}^{k-1} p(i|\theta) \pi_k(\text{admit}|i, \theta) + \left(1 - \sum_{i=0}^{k-1} p(i|\theta)\right) \pi_k(\text{admit}|k, \theta), \\ \mathbb{E}[S] &= \sum_{i=0}^{k-1} i p(i|\theta) + \frac{p(k|\theta)}{1 - \rho_k(\theta)} \left(k + \frac{\frac{\lambda}{\mu} \rho_k(\theta)}{1 - \frac{\lambda}{\mu} \rho_k(\theta)}\right), \\ Z(\theta) &= \sum_{s=0}^{k-1} \left(\prod_{i=0}^{s-1} \frac{\lambda}{\mu} \rho_i(\theta)\right) + \left(\prod_{i=0}^{k-1} \frac{\lambda}{\mu} \rho_i(\theta)\right) \frac{1}{1 - \frac{\lambda}{\mu} \rho_k(\theta)}, \end{aligned}$$

with the convention that empty sums are equal to zero and empty products are equal to one. All calculations remain valid in the limit as $\pi_k(\text{admit}|i, \theta) \rightarrow 1$ for some $i \in \{0, 1, \dots, k\}$ (corresponding to $\theta_i \rightarrow +\infty$). In the limit as $\pi_k(\text{admit}|i, \theta) \rightarrow 0$ for some $i \in \{0, 1, \dots, k\}$, we can study the restriction of the birth-and-death process to the state space $\{0, 1, \dots, \underline{c}\}$, where $\underline{c} = \min\{i \in \{0, 1, \dots, k\} : \pi_i(\theta) = 0\}$.

Assumptions of Section 5. For any closed set $U \subset \Omega$, it can be shown that there exists a Lyapunov function \mathcal{L} uniformly over $\theta \in U$ such that $\mathcal{L}(s, a) = \exp(cs)$ for some $c > 0$, depending on U and the model parameters. Hence, Assumptions 4, 5 and 6 are satisfied. In general, Assumption 7 does not hold for this example because maxima occur only as $|\theta| \rightarrow \infty$. As suggested by Proposition 2, by adding a small regularization term, we can guarantee Assumption 7 while simultaneously ensuring that the maximizer is bounded. In practice, using a regularization term can additionally present some benefits such as avoiding vanishing gradients and saddle points.

B.2 Load-balancing system

We now consider the load-balancing example of Section 6.2. Recall that jobs arrive according to a Poisson process with rate $\lambda > 0$, there are n servers at which service times are distributed exponentially with rates $\mu_1, \mu_2, \dots, \mu_n$, respectively, and the system can contain at most c jobs, for some $c \in \mathbb{N}_+$. The goal is to choose a static random policy that maximizes the admission probability. We first verify that the system satisfies Assumptions 1 to 3, then we provide an algorithm to evaluate the objective function when the parameters are known; this is used in particular for performance comparison with the optimal policy in the numerical results. Lastly, we discuss the assumptions of Section 5. Throughout this section, we assume that we apply the policy $\pi(\theta)$ defined by (46) for some parameter $\theta \in \mathbb{R}^n$.

Product-form stationary distribution. That Assumption 1 is satisfied follows from the facts that the rates and probabilities $\lambda, \mu_1, \mu_2, \dots, \mu_n, \pi_1(\theta), \pi_2(\theta), \dots, \pi_n(\theta)$ are positive and that the

state space \mathcal{S} is finite. Assumption 2 is satisfied because the state space is finite. This system can be modeled either as a loss Jackson network with n queues (one queue for each server in the load-balancing system) or as a closed Jackson network with $n+1$ queues (one queue for each server in the system, plus another queue signaling available positions in the system, with service rate λ). Either way, we can verify (for instance by writing the balance equations) that the stationary distribution of the continuous-time Markov chain that describes the evolution of the system state is given by:

$$p(s|\theta) = \frac{1}{Z(\theta)} \prod_{i=1}^n \left(\frac{\lambda}{\mu_i} \pi_i(\theta) \right)^{s_i}, \quad s = (s_1, s_2, \dots, s_n) \in \mathcal{S}, \quad (50)$$

where $Z(\theta)$ follows by normalization. This is exactly (10–PF) from Assumption 3, with $n = d$, $\Omega = \mathbb{R}^n$, $\Phi(s) = \prod_{i=1}^n \left(\frac{\lambda}{\mu_i} \right)^{s_i}$ for each $s \in \mathcal{S}$, $x_i(s) = s_i$ for each $i \in \{1, 2, \dots, n\}$ and $s \in \mathcal{S}$, and $\rho_i(\theta) = \pi_i(\theta)$ for each $i \in \{1, 2, \dots, n\}$. The function ρ defined in this way is differentiable. Assumption 3 is therefore satisfied, as the distribution of the system seen at arrival times is also (50) according to the PASTA property. Besides the sufficient statistics x , the inputs of Algorithm 2 are $\nabla \log \pi(a|s, \theta) = \mathbf{1}_a - \pi(\theta)$, where $\mathbf{1}_a$ is the n -dimensional vector with one in component a and zero elsewhere and $\pi(\theta)$ is the policy seen as a (column) vector, and $D \log \rho(\theta) = \text{Id} - \mathbf{1} \pi(\theta)^\top$, where Id is the n -dimensional identity matrix, $\mathbf{1}$ is the n -dimensional vector with all-one components, and $\pi(\theta)^\top$ is the (row) vector obtained by transposing $\pi(\theta)$. This latter equation can be used to verify Assumption 5.

Objective function. When all parameters are known and the number of servers is not too large, the normalizing constant $Z(\theta)$ and admission probability $J(\theta)$ can be computed efficiently using a variant of Buzen’s algorithm [8] for loss networks. Define the array $G = (G_{\underline{c}, \underline{n}})_{\underline{c} \in \{0, 1, \dots, c\}, \underline{n} \in \{1, 2, \dots, n\}}$ by

$$G_{\underline{c}, \underline{n}} = \sum_{\substack{s \in \mathbb{N}^{\underline{n}}: \\ |s| \leq \underline{c}}} \prod_{i=1}^{\underline{n}} \left(\frac{\lambda}{\mu_i} \rho_i(\theta) \right)^{s_i}, \quad \underline{c} \in \{0, 1, \dots, c\}, \quad \underline{n} \in \{1, 2, \dots, n\}.$$

The dependency of G on θ is left implicit to alleviate notation. The normalizing constant and admission probability are given by $Z(\theta) = G_{c, n}$ and $J(\theta) = G_{c-1, n} / G_{c, n}$, respectively. Defining the array G allows us to calculate these metrics more efficiently than by direct calculation, as we have $G_{0, \underline{n}} = 1$ for each $\underline{n} \in \{1, 2, \dots, n\}$, and

$$\begin{aligned} G_{\underline{c}, 1} &= 1 + \frac{\lambda}{\mu_1} \rho_1(\theta) G_{\underline{c}-1, 1}, & \underline{c} &\in \{1, 2, \dots, c\}, \\ G_{\underline{c}, \underline{n}} &= G_{\underline{c}, \underline{n}-1} + \frac{\lambda}{\mu_{\underline{n}}} \rho_{\underline{n}}(\theta) G_{\underline{c}-1, \underline{n}}, & \underline{c} &\in \{1, 2, \dots, c\}, \quad \underline{n} \in \{2, 3, \dots, n\}. \end{aligned}$$

Assumptions of Section 5. Assumptions 4, 5, and 6 are automatically satisfied because the state space is finite (with $|\mathcal{S}| = \binom{n+c}{c}$). Verifying Assumption 7 is challenging since it requires computing $\text{Hess}_{\theta^*} J$ at the maximizer θ^* , which depends in an implicit manner on the parameters of the system such as the arrival rate λ , service rates $\mu_1, \mu_2, \dots, \mu_n$, and policy $\pi(\theta^*)$. However, the nondegeneracy property of the Hessian for smooth functions is a property that is commonly stable in the following sense: if a function satisfies this property, then it will still be satisfied after any small-enough smooth perturbation. In particular, smooth functions with isolated nondegenerate critical points—also known as Morse functions—are dense and form an open subset in the space of smooth functions; see [29, Section 1.2]. Thus, unless the example is adversarial or presents symmetries, we can expect Assumption 7 to hold.

B.3 Ising model and Glauber dynamics

Lastly, we focus on the example of Section 6.3. We consider the Markov chain defined by applying the policy (47) parameterized by some vector $\theta \in \mathbb{R}^3$: starting from an arbitrary initial configuration, at every time step, a coordinate is chosen uniformly at random, and the agent flips or not the spin at this coordinate according to the policy.

Product-form stationary distribution The Markov chain is irreducible because it has a positive probability of transitioning from any configuration to any other as follows: at every step, choose a coordinate at which the two configurations differ and flip the spin at this coordinate. The Markov chain is positive recurrent because its state space is finite. Hence, Assumptions 1 and 2 are satisfied. We now focus on proving Assumption 3.

Our goal is to verify that the Markov chain that describes the random evolution of the state admits the stationary distribution (48), which we recall here:

$$p(\sigma, v|\theta) = \frac{1}{Z(\theta)} e^{\beta(\theta)JI(\theta) + \beta(\theta)\mu F(\sigma|\theta)}, \quad s = (\sigma, v) \in \mathcal{S}, \quad \theta \in \mathbb{R}^3,$$

Observe that $p(\sigma, v|\theta)$ is independent of v , hence we can let $q(\sigma|\theta) \triangleq p(\sigma, \cdot|\theta)$ for each $\sigma \in \Sigma$. The key argument to prove that this is indeed the stationary distribution consists of observing that the policy (47) satisfies, for each $s = (\sigma, v) \in \mathcal{S}$,

$$\pi(\text{flip}|\sigma, v, \theta) = \frac{q(\sigma_{-v}|\theta)}{q(\sigma|\theta) + q(\sigma_{-v}|\theta)}, \quad \pi(\text{not flip}|\sigma, v, \theta) = \frac{q(\sigma|\theta)}{q(\sigma|\theta) + q(\sigma_{-v}|\theta)}, \quad (51)$$

where $\sigma_{-v} \in \Sigma$ is the configuration obtained by flipping the spin at v compared to σ , that is, $\sigma_{-v}(w) = \sigma(w)$ for each $w \in \mathcal{V} \setminus \{v\}$ and $\sigma_{-v}(v) = -\sigma(v)$.

The balance equation for a particular state $s = (\sigma, v) \in \mathcal{S}$ writes

$$p(\sigma, v|\theta) = \sum_{w \in \mathcal{V}} p(\sigma, w|\theta) \pi(\text{not flip}|\sigma, w, \theta) \frac{1}{d_1 d_2} + \sum_{w \in \mathcal{V}} p(\sigma_{-w}, w|\theta) \pi(\text{flip}|\sigma_{-w}, w, \theta) \frac{1}{d_1 d_2}.$$

Dropping the dependency on θ to simplify notation, and injecting (51) into the right-hand side of this balance equation, we obtain successively

$$\begin{aligned} & \sum_{w \in \mathcal{V}} p(\sigma, w) \frac{q(\sigma)}{q(\sigma) + q(\sigma_{-w})} \frac{1}{d_1 d_2} + \sum_{w \in \mathcal{V}} p(\sigma_{-w}, w) \frac{q((\sigma_{-w})_{-w})}{q(\sigma_{-w}) + q((\sigma_{-w})_{-w})} \frac{1}{d_1 d_2} \\ & \stackrel{(1)}{=} \sum_{w \in \mathcal{V}} (p(\sigma, w) + p(\sigma_{-w}, w)) \frac{q(\sigma)}{q(\sigma) + q(\sigma_{-w})} \frac{1}{d_1 d_2} \stackrel{(2)}{=} q(\sigma) \sum_{w \in \mathcal{V}} \frac{1}{d_1 d_2} = q(\sigma) \stackrel{(2)}{=} p(\sigma, v), \end{aligned}$$

where (1) follows by observing that $(\sigma_{-w})_{-w} = \sigma$ and (2) by recalling that $q(\sigma) = p(\sigma, w)$ for each $(\sigma, v) \in \mathcal{S}$. This proves that the distribution (48) is indeed the stationary distribution of the Markov chain that describes the evolution of the state under the policy (47).

Besides the sufficient statistics x , the inputs of Algorithm 2 are given, for each $\theta \in \mathbb{R}^3$ and $s = (\sigma, v) \in \mathcal{S}$, by

$$D \log \rho(\theta) = \begin{bmatrix} \beta'(\theta)J & 0 & 0 \\ \beta'(\theta)\mu\theta_{\text{left}}(\theta) & \beta(\theta)\mu h_{\text{left}}'(\theta) & 0 \\ \beta'(\theta)\mu\theta_{\text{right}}(\theta) & 0 & \beta(\theta)\mu h_{\text{right}}'(\theta) \end{bmatrix},$$

$$\nabla \log \pi(a|\sigma, v, \theta) = (\mathbf{1}[a = \text{not flip}] - \pi(\text{not flip}|\sigma, v, \theta)) \nabla \delta(\sigma, v|\theta),$$

$$\nabla \delta(\sigma, v|\theta) = 2 \begin{bmatrix} \beta'(\theta) \sigma(v) (J \sum_{w \in \mathcal{V}: w \sim v} \sigma(w) + \mu h(v|\theta)) \\ \beta(\theta) \sigma(v) \mu h_{\text{left}}'(\theta) \mathbf{1}[v_2 \leq d_2/2] \\ \beta(\theta) \sigma(v) \mu h_{\text{right}}'(\theta) \mathbf{1}[v_2 > d_2/2] \end{bmatrix},$$

where $\beta'(\theta)$ (resp. $h_{\text{left}}'(\theta)$, $h_{\text{right}}'(\theta)$) is to be understood as the partial derivative of β (resp. h_{left} , h_{right}) with respect to θ_1 (resp. θ_2 , θ_3).

C Proof of Theorem 2

C.1 Preliminaries

We are going to use concentration inequalities for Markov chains. Such results are common in the literature (for example, see [19]), and will be required to get a concentration bound of the plug-in estimators from (16).

Denote by $B_\epsilon(\theta)$ the open ball of radius ϵ centered at $\theta \in \Omega \subseteq \mathbb{R}^n$ and $\mathcal{Y} = \mathcal{S} \times \mathcal{A}$. Given a function $q : \mathcal{Y} \rightarrow \mathbb{R}$ and the Lyapunov function $\mathcal{L} : \mathcal{Y} \rightarrow [1, \infty)$ from Assumption 4, define

$$|q|_{\mathcal{L}} = \sup_{y \in \mathcal{Y}} \frac{|q(y)|}{\mathcal{L}(y)}. \quad (52)$$

Given a signed measure ν , we also define the seminorm

$$|\nu|_{\mathcal{L}} = \sup_{|q|_{\mathcal{L}} \leq 1} |\nu[q]| = \sup_{|q|_{\mathcal{L}} \leq 1} \left| \int q(y) \nu(dy) \right|. \quad (53)$$

Equations (52) to (53) imply that

$$|\nu[q]| \leq |\nu|_{\mathcal{L}} |q|_{\mathcal{L}}. \quad (54)$$

Note that we defined $|\cdot|_{\mathcal{L}}$ for a unidimensional function. Given instead m functions $q_i : \mathcal{Y} \rightarrow \mathbb{R}$, for the higher-dimensional function $q : \mathcal{Y} \rightarrow \mathbb{R}^m$ that satisfies for all $y \in \mathcal{Y}$, $q(y) = (q_1(y), \dots, q_l(y))$, we define $|q|_{\mathcal{L}} = \sqrt{\sum_{i=1}^l |q_i|_{\mathcal{L}}^2}$.

The following lemma yields the concentration inequalities required:

Lemma 5. *Let $\{Y_n\}_{n \geq 1}$ be a geometrically ergodic Markov chain with invariant distribution p and transition matrix $P(\cdot, \cdot)$. Let the Lyapunov function be $\mathcal{L} : \mathcal{Y} \rightarrow \mathbb{R}$. From geometric ergodicity, there exists $C > 0$ and $\lambda \in (0, 1)$ such that for any $y \in \mathcal{Y}$,*

$$|P^m(\cdot | y) - p(\cdot)|_{\mathcal{L}} \leq C \lambda^m. \quad (55)$$

Let $\mathcal{F} = \sigma(Y_1)$ be the σ -algebra of Y_1 . Let $q : \mathcal{Y} \rightarrow \mathbb{R}^m$ be a measurable function such that $|q|_{\mathcal{L}} < \infty$. For a finite trajectory Y_1, \dots, Y_M of the Markov chain, we define the empirical estimator for $p[q]$ as

$$\hat{p}_M[q] = \frac{1}{M} \sum_{i=1}^M q(Y_i). \quad (56)$$

With these assumptions, there exists C' depending on C and λ such that

$$\left| \mathbb{E} \left[p[q] - \hat{p}_M[q] | \mathcal{F} \right] \right| \leq \frac{C' |q|_{\mathcal{L}}}{M} \mathcal{L}(Y_1), \quad (57)$$

and for $l \in \{1, 2, 4\}$,

$$\mathbb{E}\left[|p[q] - \hat{p}_M[q]|^l \mid \mathcal{F}\right] \leq \frac{C'|q|_{\mathcal{L}}^l}{M^{l/2}} \mathcal{L}^l(Y_1). \quad (58)$$

Proof. We refer to [15, Prop. 12] for a proof of (58). What remains is to prove (57).

Observe that for $y \in \mathcal{Y}$, $P(y) = P(\cdot | y)$ is a distribution over \mathcal{Y} . Conditional on \mathcal{F} , there exists $C > 0$ such that

$$\begin{aligned} \left| \mathbb{E}\left[\frac{1}{M} \sum_{i=1}^M q(Y_i) - p[q] \mid \mathcal{F}\right] \right| &\leq \frac{1}{M} \sum_{i=1}^M |P^i(Y_1)[q] - p[q]| = \frac{1}{M} \sum_{i=1}^M |(P^i(Y_1) - p)[q]| \\ &\leq \frac{1}{M} \sum_{i=1}^M |P^i(\cdot | Y_1) - p(\cdot)|_{\mathcal{L}} |q|_{\mathcal{L}} \mathcal{L}(Y_1) \leq \frac{|q|_{\mathcal{L}}}{M} \sum_{i=1}^M C \lambda^i \mathcal{L}(Y_1) \\ &\leq \frac{C|q|_{\mathcal{L}}}{M(1-\lambda)} \mathcal{L}(Y_1). \end{aligned} \quad (59)$$

This concludes the proof. \square

In epoch m , the Markov chain $\{S_t\}_{t \in [t_m, t_{m+1}]}$ with control parameter Θ_m has a Lyapunov function \mathcal{L}_v . Intuitively, as a consequence of Assumption 4, we can show that the process does not drift to infinity on the event \mathcal{B}_m (despite the changing control parameter Θ_m).

Specifically, for $m > 0$, let $\{S_t\}_{i \in [t_m, t_{m+1}]}$ be the Markov chain trajectory with transition probabilities $P(\Theta_m)$, where Θ_m is given by the updates in (3) and (16) and initial state $S_0 \in \mathcal{S}$. Recall that \mathcal{B}_m is defined in (33). We can then prove the following:

Lemma 6. *Suppose Assumption 4 holds. There exists $D < \infty$ such that for $m > 0$, $\mathbb{E}[\mathcal{L}_v(S_{t_{m+1}}) \mathbf{1}[\mathcal{B}_m]] < D$.*

Proof. We will give an inductive argument. A similar argument can be found in [4].

First, observe that for $m = 0$, S_0 is fixed. There thus exists a D such that $\mathcal{L}_v(S_0) < D$.

Next, assume that $\mathbb{E}[\mathcal{L}_v(S_{t_m}) \mathbf{1}[\mathcal{B}_{m-1}]] < D$. On the event \mathcal{B}_m , Assumption 4 holds since $\Theta_1, \dots, \Theta_{m-1}, \Theta_m \in V_{\tau, \delta}(\theta^*) \subset U$. Thus, on the event \mathcal{B}_m , and when additionally conditioning on $S_{t_{m+1}-1}$ and Θ_m , the following holds true:

$$\begin{aligned} \mathbb{E}[\mathcal{L}_v(S_{t_{m+1}}) \mathbf{1}[\mathcal{B}_m]] &\leq \mathbb{E}[\mathbb{E}[\mathcal{L}_v(S_{t_{m+1}}) \mathbf{1}[\mathcal{B}_m] \mid S_{t_{m+1}-1}]] \\ &= \mathbb{E}[\mathbf{1}[\mathcal{B}_m] P_{\Theta_m} \mathcal{L}_v(S_{t_{m+1}-1})] \\ &\leq \mathbb{E}[\mathbf{1}[\mathcal{B}_m] [\lambda \mathcal{L}_v(S_{t_{m+1}-1}) + b]]. \end{aligned} \quad (60)$$

The last step followed from Assumption 4.

Observe finally that the bound in (60) can be iterated by conditioning on $S_{t_{m+1}-2}$; so on and so forth. After $t_{m+1} - t_m$ iterations, one obtains

$$\mathbb{E}[\mathcal{L}_v(S_{t_{m+1}}) \mathbf{1}[\mathcal{B}_m]] \leq \lambda \mathbb{E}[\mathcal{L}_v(S_{t_m}) \mathbf{1}[\mathcal{B}_m]] + \frac{b}{1-\lambda}. \quad (61)$$

Noting that $\mathbf{1}[\mathcal{B}_m] \leq \mathbf{1}[\mathcal{B}_{m-1}]$, the claim follows by induction if we choose D large enough such that $\lambda D + b/(1-\lambda) < D$. \square

C.2 Proof of Theorem 2

To prove Theorem 2, we more-or-less follow the arguments of [14, Thm. 25]. Modifications are however required because we consider a Markovian setting instead. Specifically, we rely on the bounds in Lemmas 2 to 4 instead of the bounds in [14, Prop. 20, Prop. 21, Prop. 24], respectively.

Let us begin by bounding

$$\mathbb{P}[J^* - J(\Theta_m) > \epsilon | \mathcal{B}_0]. \quad (62)$$

Here, $\mathcal{B}_0 = \{\Theta_0 \in V_{\tau, \delta}(\theta^*)\}$ —recall (33). Theorem 2 assumes that we initialize in a set V which we will specify later but satisfies $V \subset V_{\tau, \delta}(\theta^*)$. Since we can initialize Θ_0 with positive probability in V , we have that $\mathbb{P}[\mathcal{B}_0] \geq \mathbb{P}[\Theta_0 \in V] > 1/c > 0$ for some $c > 0$. Thus, we will focus on finding an upper bound of

$$\mathbb{P}[J^* - J(\Theta_m) > \epsilon | \mathcal{B}_0] \leq c \mathbb{P}[\{J^* - J(\Theta_m) > \epsilon\} \cap \mathcal{B}_0]. \quad (63)$$

Denote the orthogonal projection of Θ_m onto $\mathcal{M} \cap U$ by $\tilde{\Theta}_m = \mathfrak{p}(\Theta_m)$. We can relate the objective gap $J^* - J(\Theta_m)$ to the distance $D_m := \text{dist}(\Theta_m, \mathcal{M} \cap U)$ as follows. Since J is twice continuously differentiable with maximum J^* attained at $\mathcal{M} \cap U$, the function $J(\theta)$ with $\theta \in V_{\tau, \delta}(\theta^*)$ is locally Lipschitz with constant $\iota_{\tau, \delta}(\theta^*) > 0$. On the event \mathcal{B}_m , we have $\Theta_m \in V_{\tau, \delta}(\theta^*)$ and therefore we have the inequality

$$J^* - J(\Theta_m) = J(\tilde{\Theta}_m) - J(\Theta_m) \leq \iota_{\tau, \delta}(\theta^*) |\tilde{\Theta}_m - \Theta_m| = \iota_{\tau, \delta}(\theta^*) D_m. \quad (64)$$

Consequently, we have the bound

$$\mathbb{P}[\{J^* - J(\Theta_m) > \epsilon\} \cap \mathcal{B}_m] \leq \mathbb{P}\left[\left\{D_m \geq \frac{\epsilon}{\iota_{\tau, \delta}(\theta^*)}\right\} \cap \mathcal{B}_m\right] \quad (65)$$

If we define $\epsilon' = \epsilon / \iota_{\tau, \delta}(\theta^*)$, the right-hand side of (65) can also be written as

$$\begin{aligned} \mathbb{P}[\{D_m \geq \epsilon'\} \cap \mathcal{B}_m] &= \mathbb{E}[\mathbb{1}[D_m \geq \epsilon'] \mathbb{1}[\mathcal{B}_m]] \\ &= \mathbb{E}[\mathbb{1}[D_m \mathbb{1}[\mathcal{B}_m] \geq \epsilon']] = \mathbb{P}[D_m \mathbb{1}[\mathcal{B}_m] \geq \epsilon'] \end{aligned} \quad (66)$$

by the positivity of D_m .

Next, we use (i) the law of total probability noting that $\mathcal{B}_m \subset \mathcal{B}_0$, (ii) the bound (65) and the inequality $\mathbb{P}[A \cap B] \leq \mathbb{P}[A]$ for any two events A, B , and finally, (iii) the equality (66). We obtain

$$\begin{aligned} \mathbb{P}[\{J^* - J(\Theta_m) > \epsilon\} \cap \mathcal{B}_0] &\stackrel{(i)}{\leq} \mathbb{P}[\{J^* - J(\Theta_m) > \epsilon\} \cap \mathcal{B}_m] + \mathbb{P}[\{J^* - J(\Theta_m) > \epsilon\} \cap \overline{\mathcal{B}_m}] \\ &\stackrel{(ii)}{\leq} \mathbb{P}[\{D_m \geq \epsilon'\} \cap \mathcal{B}_m] + \mathbb{P}[\overline{\mathcal{B}_m}] \\ &\stackrel{(iii)}{\leq} \mathbb{P}[D_m \mathbb{1}[\mathcal{B}_m] \geq \epsilon'] + \mathbb{P}[\overline{\mathcal{B}_m}] \\ &\leq \mathbb{P}[D_m \mathbb{1}[\mathcal{B}_{m-1}] \geq \epsilon'] + \mathbb{P}[\overline{\mathcal{B}_m}] = \text{Term I} + \text{Term II}. \end{aligned} \quad (67)$$

Term I can be bounded by using Markov's inequality and Lemma 2. This shows that

$$\text{Term I} \leq c \epsilon'^{-2} m^{-\sigma - \kappa}. \quad (68)$$

Term II can be bounded by Lemma 4. Specifically, one finds that there exists a constant $c > 0$ such that, if $\Theta_0 \in V_{\tau/2, \delta}(\theta^*)$,

$$\text{Term II} \leq 1 - \exp\left(-\frac{c\alpha^2}{\delta^2\ell}\right) + c\delta^{-2}\ell^{-1}m^{1-\sigma-\kappa} + c\alpha \frac{(m^{1-3/2\sigma-\kappa/2} + \ell^{-1/2}m^{1-5\sigma/8-\kappa/2})}{(\tau/2 - 2\delta)_+}. \quad (69)$$

Note next that for any $\alpha \in (0, \alpha_0]$ and $c > 0$ there exists δ_0 such that for any $\delta \in (0, \delta_0]$ there exists ℓ_0 such that if $\ell \in [\ell_0, \infty)$ there exists a constant $c' > 0$ such that we have the inequality $1 - \exp(-c\alpha^2/\delta^2\ell) \leq c'\alpha^2/\delta^2\ell$. We can substitute this bound in (69) to yield

$$\text{Term II} \leq c' \frac{\alpha^2}{\delta^2\ell} + \text{idem.} \quad (70)$$

Bounding (67) by the sum of (68) and (70), and substituting the bound in (63) reveals that there exists a constant $c'' > 0$ such that if $\Theta_0 \in V_{\tau/2, \delta}(\theta^*)$ then

$$\begin{aligned} \mathbb{P}[J^* - J(\Theta_m) > \epsilon | \mathcal{B}_0] &\leq c''(\epsilon')^{-2} m^{-\sigma-\kappa} + c''\alpha^2\delta^{-2}\ell^{-1} + c''\delta^{-2}\ell^{-1}m^{1-\sigma-\kappa} \\ &\quad + c''\alpha \frac{(m^{1-3/2\sigma-\kappa/2} + \ell^{-1/2}m^{1-5\sigma/8-\kappa/2})}{(\tau/2 - 2\delta)_+}. \end{aligned} \quad (71)$$

Note that the exponents of m in (71) satisfy that since $\sigma \in (2/3, 1)$, $1 - 3/2\sigma - \kappa/2 \leq -\kappa/2$ as well as $1 - 5\sigma/8 - \kappa/2 < 1 - \sigma/2 - \kappa/2$. Finally, let the initialization set be $V = V_{\tau/2, \delta}(\theta^*)$. Note that since $\{\Theta_0 \in V\} \subset \mathcal{B}_0$ there exists a constant $c''' > 0$ such that

$$\mathbb{P}[J^* - J(\Theta_m) > \epsilon | \Theta_0 \in V] \leq c'''\mathbb{P}[J^* - J(\Theta_m) > \epsilon | \mathcal{B}_0]. \quad (72)$$

Substituting the upper bound (71) in (72) concludes the proof. \square

C.3 Proof of Lemma 1

For simplicity, we will denote $t_{m+1} - t_m = T_m$, $X_t = x(S_t)$ throughout this proof. We also temporarily omit the summation indices for the epoch. We note that the policies defined in (6) satisfy that for $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$(\nabla \log \pi(a|s, \theta))_{i, a'} = \begin{cases} \mathbf{1}[a = a'] - \pi(a'|s, \theta) & \text{if } i = h(s), \\ 0 & \text{otherwise.} \end{cases}$$

In particular, there exists $c_1 > 0$ such that for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, $|\nabla \log \pi(a|s, \theta)| < c_1$. The proof below, however, can also be extended to other policy classes.

C.3.1 Proof of (38)

Observe that if the event \mathcal{B}_m holds, that then the definitions in (16) also imply that

$$\begin{aligned} \eta_m &= \nabla J(\Theta_m) - H_m = \nabla J(\Theta_m) - (\text{D log } \rho(\Theta_m)^\top \bar{C}_m + \bar{E}_m) \\ &= \nabla J(\Theta_m) - (\text{D log } \rho(\Theta_m)^\top \frac{1}{T_{m+1}} \sum_{t=t_m}^{t_{m+1}-1} (X_t - \bar{X}_m)r(S_t, A_t) \\ &\quad + \frac{1}{T_m} \sum_{t=t_m}^{t_{m+1}-1} r(S_t, A_t)\nabla \log \pi(A_t|S_t, \Theta_m) \\ &= \text{D log } \rho(\Theta_m)^\top \left(\text{Cov}[R, S] - \frac{1}{T_m} \sum_{t=t_m}^{t_{m+1}-1} (X_t - \bar{X}_m)r(S_t, A_t) \right) \end{aligned}$$

$$\begin{aligned}
& + \left(\mathbb{E}[R \nabla \log \pi(A|S, \Theta_m)] - \frac{1}{T_m} \sum_{t=t_m}^{t_{m+1}-1} r(S_t, A_t) \nabla \log \pi(A_t|S_t, \Theta_m) \right) \\
& = D \log \rho(\Theta_m)^\top \tilde{\eta}_m + \tilde{\zeta}_m.
\end{aligned} \tag{73}$$

We will deal with the terms $\tilde{\eta}_m$ in and $\tilde{\zeta}_m$ in (73) one-by-one.

Dealing with the 1st term, $\tilde{\eta}_m$. Define

$$\begin{aligned}
A & = \mathbb{E}[(X - \mathbb{E}[X])R] - \frac{1}{T_m} \sum_t (X_t - \mathbb{E}[X])r(S_t, A_t), \\
B & = \frac{1}{T_m} \left(\sum_t r(S_t, A_t) \right) (\mathbb{E}[X] - \bar{X}_m),
\end{aligned} \tag{74}$$

and observe that

$$\tilde{\eta}_m = A + B. \tag{75}$$

We look first at A in (74). Recall that $\{Y_t\}_{t>0} = \{(S_t, A_t)\}_{t>0}$ is the chain of state-action pairs (see Section 5.1). Define the function $g : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^n$ as

$$g(y) = g((s, a)) = (x(s) - \mathbb{E}[x(s)])r(y). \tag{76}$$

Then, we can rewrite

$$A = \mathbb{E}[g(Y)] - \frac{1}{T_m} \sum_t g(Y_t). \tag{77}$$

We are now almost in position to apply Lemma 5 to A . Observe next that the law of total expectation implies that

$$\mathbb{E}[\eta_m \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m] = \sum_{a \in \mathcal{A}} \mathbb{E}[\eta_m \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] \pi(a | S_m, \Theta_m), \tag{78}$$

Without loss of generality, it therefore suffices to consider the case that we have one action $A_{t_m} = a \in \mathcal{A}$. For the first term we have that there exists a constant $c_2 > 0$ such that

$$\begin{aligned}
\left| \mathbb{E}[A \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] \right| & = \left| \mathbb{E} \left[\mathbb{E}[g(Y)] - \frac{1}{T_m} \sum_t g(Y_t) \mathbf{1}[\mathcal{B}_m] \middle| Y_0 = (S_{t_m}, A_{t_m}) \right] \right| \\
& \stackrel{\text{(Lemma 5)}}{=} \frac{c_2 |g|_{\mathcal{L}}}{T_m} \mathcal{L}((S_{t_m}, a)),
\end{aligned} \tag{79}$$

where we can use that $|g|_{\mathcal{L}} < \infty$ due to Assumption 6.

For the term B in (74). We can add and subtract again the following terms and obtain

$$\begin{aligned}
B & = \frac{1}{T_m} \left(\sum_t r(S_t, A_t) \right) (\mathbb{E}[X] - \bar{X}_m) - \mathbb{E}[R] (\mathbb{E}[X] - \bar{X}_m) \\
& \quad + \mathbb{E}[R] (\mathbb{E}[X] - \bar{X}_m) \\
& = C + D,
\end{aligned} \tag{80}$$

where

$$\begin{aligned} C &= (\mathbb{E}[X] - \bar{X}_m) \left(\frac{1}{T_m} \sum_t r(S_t, A_t) - \mathbb{E}[R] \right), \\ D &= \mathbb{E}[R] (\mathbb{E}[X] - \bar{X}_m). \end{aligned} \quad (81)$$

For the term D in (81) we can readily use the concentration of Lemma 5 to obtain

$$\mathbb{E}[\mathbb{E}[R] (\mathbb{E}[X] - \bar{X}_m) \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] \leq \mathbb{E}[R] \frac{|x(S)|_{\mathcal{L}}}{T_m} \mathcal{L}(S_{t_m}, a), \quad (82)$$

where we have $|x(S)|_{\mathcal{L}} < \infty$ from Assumption 6 and $\mathbb{E}[R] < J^*$.

For the term C , we use Cauchy–Schwartz together with Lemma 1. In particular, we have

$$\begin{aligned} & \left| \mathbb{E}[(\mathbb{E}[X] - \bar{X}_m) \left(\frac{1}{T_m} \sum_t r(S_t, A_t) - \mathbb{E}[R] \right) \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] \right| \leq \\ & \quad \left| \mathbb{E}[|\mathbb{E}[X] - \bar{X}_m|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] \right|^{1/2} \times \\ & \quad \left| \mathbb{E} \left[\left| \frac{1}{T_m} \sum_t r(S_t, A_t) - \mathbb{E}[R] \right|^2 \mathbf{1}[\mathcal{B}_m] \middle| \mathcal{F}_m, A_{t_m} = a \right] \right|^{1/2} \end{aligned} \quad (83)$$

For both terms we can repeat the same argument to that in (78) together with Lemma 5 to show that

$$\begin{aligned} & \left| \mathbb{E}[|\mathbb{E}[X] - \bar{X}_m|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] \right|^{1/2} \leq c_3 \frac{|X|_{\mathcal{L}}^{1/2}}{T_m^{1/2}} \mathcal{L}(S_{t_m}, a) \\ & \left| \mathbb{E} \left[\left| \frac{1}{T_m} \sum_t r(S_t, A_t) - \mathbb{E}[R] \right|^2 \mathbf{1}[\mathcal{B}_m] \middle| \mathcal{F}_m, A_{t_m} = a \right] \right|^{1/2} \leq c_4 \frac{|R|_{\mathcal{L}}^{1/2}}{T_m^{1/2}} \mathcal{L}(S_{t_m}, a) \end{aligned} \quad (84)$$

Therefore multiplying both bounds in (84) and using Assumption 6 to bound the \mathcal{L} -norms, we obtain that there exists $c_5 > 0$ such that

$$|\mathbb{E}[C | \mathcal{F}_m, A_{t_m} = a]| \leq \frac{c_5}{T_m} \mathcal{L}(S_{t_m}, a)^2. \quad (85)$$

Adding the bounds (79), (85), and (82) together we have now

$$|\mathbb{E}[\tilde{\eta}_m \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a]| \leq \frac{c_6}{T_m} \mathcal{L}^2(S_{t_m}, a). \quad (86)$$

Finally, averaging this bound over all actions in (78), we obtain

$$|\mathbb{E}[\tilde{\eta}_m \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m]| \leq \frac{c_7}{T_m} \left(\sum_a \mathcal{L}(S_{t_m}, a)^2 \pi(a | S_{t_m}, \Theta_m) \right) \leq \frac{c_7}{T_m} \mathcal{L}_4(S_{t_m})^{1/2}. \quad (87)$$

Now we use Assumption 5. We can write

$$\begin{aligned} |\mathbb{E}[\nabla \log(\Theta_m) \tilde{\eta}_m \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m]| &= |\nabla \log(\Theta_m) \mathbb{E}[\tilde{\eta}_m \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m]| \\ &\leq C |\mathbb{E}[\tilde{\eta}_m \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m]| \\ &\leq \frac{c_8}{T_m} \mathcal{L}(S_{t_m}) \end{aligned} \quad (88)$$

Dealing with the 2nd term, $\tilde{\zeta}_m$. Define a function of $Y = (S, A)$ as

$$g(Y) = r(Y)\nabla \log \pi(A|S, \theta), \quad (89)$$

so that

$$\zeta_m = \mathbb{E}[g(Y)] - \frac{1}{T_m} \sum_t g(Y_t). \quad (90)$$

By combining the argument of (78) with the fact that $|g(Y)|_{\mathcal{L}} < \infty$ by Assumption 6, we find that

$$|\mathbb{E}[\tilde{\zeta}_m \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m]| \leq \frac{c_9}{T_m} \mathcal{L}(S_{t_m}) \quad (91)$$

Adding (87) and (91) together with their largest exponents yields

$$\begin{aligned} |\mathbb{E}[\eta_m \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m]| &\leq \frac{c_{10}}{T_m} \sum_a \mathcal{L}(S_{t_m}, a)^2 \pi(a|S_{t_m}) \\ &\leq \frac{c_{10}}{T_m} \left(\sum_a \mathcal{L}(S_{t_m}, a)^4 \pi(a|S_{t_m}) \right)^{1/2} \leq \frac{c_{10}}{T_m} \mathcal{L}_4(S_{t_m})^{1/2}. \end{aligned} \quad (92)$$

This concludes the proof of (38). \square

C.3.2 Proof of (39)

Note that by using the fact that for a vector-valued random variable Z we have that $\mathbb{E}[|Z|^2] \geq \mathbb{E}[|Z|]^2$, the case for $p = 1$ follows from the case $p = 2$.

We focus on the case $p = 2$. By using the identity $(a + b) \leq 2a^2 + b^2$, we estimate

$$\begin{aligned} &\mathbb{E}[|D \log \rho(\Theta_m)^\top \tilde{\eta}_m + \tilde{\zeta}_m|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m] \\ &\leq 2(\mathbb{E}[|D \log \rho(\Theta_m)^\top \tilde{\eta}_m|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m] + \mathbb{E}[|\tilde{\zeta}_m|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m]) \\ &\stackrel{(5)}{\leq} 2c_1^2 \mathbb{E}[|\tilde{\eta}_m|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m] + 2\mathbb{E}[|\tilde{\zeta}_m|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m] \end{aligned} \quad (93)$$

say. We again use the law of total expectation with the action set in (78) and condition on the action $A_m = a$.

For the term involving $\tilde{\zeta}_m$ in (93) we can again use the definition of g in (89). We bound

$$\begin{aligned} \mathbb{E}[|\tilde{\zeta}_m|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] &= \mathbb{E}[|\mathbb{E}[g(Y) - \frac{1}{T_m} \sum_t g(Y)|^2 | Y_0 = (S_{t_m}, a)]| \\ &\stackrel{(\text{Lemma 5})}{\leq} \frac{c_2}{T_m} \mathcal{L}(S_{t_m}, a)^2. \end{aligned} \quad (94)$$

For the term involving $\tilde{\eta}_m$ in (93), we use the same definition for the terms A, C and D from (79), (85) and (82) as in the proof of (38). We have the bound

$$\begin{aligned} \mathbb{E}[|\tilde{\eta}_m|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] &\leq 3(\mathbb{E}[|A|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] + \mathbb{E}[|C|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] \\ &\quad + \mathbb{E}[|D|^2 \mathbf{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a]) \end{aligned} \quad (95)$$

For the terms pertaining to A and D in (95) the same argument as those used for $\tilde{\zeta}_m$ in (89) and (94) can be used to show that

$$\begin{aligned}\mathbb{E}[|A|^2 \mathbb{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] &\leq \frac{c_3}{T_m} \mathcal{L}(S_{t_m}, a)^2 \\ \mathbb{E}[|D|^2 \mathbb{1}[\mathcal{B}_m] | \mathcal{F}_m, A_{t_m} = a] &\leq \frac{c_4}{T_m} \mathcal{L}(S_{t_m}, a)^2.\end{aligned}\quad (96)$$

The only remaining term to bound in (95) is C . We use again Cauchy–Schwartz’s inequality

$$\begin{aligned}\mathbb{E}\left[\left|(\mathbb{E}[X] - \bar{X}_m) \left(\frac{1}{T_m} \sum_t r(S_t, A_t) - \mathbb{E}[R]\right) \mathbb{1}[\mathcal{B}_m]\right|^4 \middle| \mathcal{F}_m, A_{t_m} = a\right] &\leq \\ &\mathbb{E}\left[|\mathbb{E}[X] - \bar{X}_m|^2 \mathbb{1}[\mathcal{B}_m] \middle| \mathcal{F}_m, A_{t_m} = a\right]^{1/2} \times \\ &\left|\mathbb{E}\left[\left|\frac{1}{T_m} \sum_t r(S_t, A_t) - \mathbb{E}[R]\right|^4 \mathbb{1}[\mathcal{B}_m] \middle| \mathcal{F}_m, A_{t_m} = a\right]\right|^{1/2} m\end{aligned}\quad (97)$$

and by Lemma 5 the following hold

$$\begin{aligned}\left|\mathbb{E}\left[|\mathbb{E}[X] - \bar{X}_m|^4 \mathbb{1}[\mathcal{B}_m] \middle| \mathcal{F}_m, A_{t_m} = a\right]\right|^{1/2} &\leq c_5 \frac{|X|_{\mathcal{L}}^{1/2}}{T_m} \mathcal{L}(S_{t_m}, a)^2 \\ \left|\mathbb{E}\left[\left|\frac{1}{T_m} \sum_t r(S_t, A_t) - \mathbb{E}[R]\right|^4 \mathbb{1}[\mathcal{B}_m] \middle| \mathcal{F}_m, A_{t_m} = a\right]\right|^{1/2} &\leq c_6 \frac{|R|_{\mathcal{L}}^{1/2}}{T_m} \mathcal{L}(S_{t_m}, a)^2.\end{aligned}\quad (98)$$

The bound for C thus becomes

$$\mathbb{E}[|C|^2 | \mathcal{F}_m, A_{t_m} = a] \leq \frac{c_7}{T_m^2} \mathcal{L}(S_{t_m}, a)^4. \quad (99)$$

Upper bounding all terms by the largest exponents and adding over the different actions, we finally obtain

$$\mathbb{E}[|\eta_m|^2 \mathbb{1}[\mathcal{B}_m] | \mathcal{F}_m] \leq \frac{c_8}{T_m} \sum_a \mathcal{L}(S_{t_m}, a)^4 \pi(a | S_{t_m}), \Theta_m \leq \frac{c_9}{T_m} \mathcal{L}_4(S_{t_m}). \quad (100)$$

That is it. \square

C.4 Proof of Lemma 2

We will again use the notation $t_{m+1} - t_m = T_m$ and without loss of generality we will assume that $T_m = \ell m^{\sigma/2+\kappa}$ instead of $\lfloor \ell m^{\sigma/2+\kappa} \rfloor$. This can be assumed since for $m \geq 1$ there exist constants $c_l, c_u > 0$ such that $c_l \ell m^{\sigma/2+\kappa} \leq t_{m+1} - t_m \leq c_u \ell m^{\sigma/2+\kappa}$. The proof of Lemma 2 follows the same steps as in [14, Proposition 20]. However, we have to quickly diverge and adapt the estimates to the case that there the variance of H_m depends on the states of a Markov chain. From the assumptions, it can be shown that there is a unique differentiable orthogonal projection map $\mathbf{p} : V_{\tau, \delta}(\theta^*) \rightarrow \mathcal{M} \cap U$ from $V_{\tau, \delta}(\theta^*) \cap U$ onto $V_{\tau, \delta}(\theta^*) \cap \mathcal{M} \cap U$. The distance of Θ_m to the set of minima can then be upper bounded by the distance to the projection $\mathbf{p} : V_{\tau, \delta}(\theta^*) \rightarrow \mathcal{M} \cap U$ of Θ_{m-1} by

$$\begin{aligned}\text{dist}(\Theta_m, \mathcal{M} \cap U)^2 &\leq |\Theta_m - \mathbf{p}(\Theta_{m-1})|^2 \\ &\leq |\Theta_{m-1} - \mathbf{p}(\Theta_{m-1}) - \alpha_{m-1} \nabla J(\Theta_{m-1})|\end{aligned}$$

$$+ (\alpha_{m-1} \nabla J(\Theta_{m-1}) - \alpha_{m-1} H_{m-1})^2 \quad (101)$$

After expanding (101) and taking expectations, however, the effect of bias already appears, and we must diverge from the analysis from [14, (44)] thereafter. In particular, the effect of the bias of H_{m-1} needs to be handled in the terms

$$\mathbb{E} \left[2 \left\langle \Theta_{m-1} - \mathfrak{p}(\Theta_{m-1}) - \alpha_{m-1} \nabla J(\Theta_{m-1}), \alpha_{m-1} \nabla J(\Theta_{m-1}) - \alpha_{m-1} H_{m-1} \right\rangle \mathbb{1}[\mathcal{B}_{m-1}] \right], \quad (102)$$

and

$$\mathbb{E} \left[\left| \alpha_{m-1} \nabla J(\Theta_{m-1}) - \alpha_{m-1} H_{m-1} \right|^2 \mathbb{1}[\mathcal{B}_{m-1}] \right] = (\alpha_{m-1})^2 \mathbb{E} \left[|\eta_{m-1}|^2 \mathbb{1}[\mathcal{B}_{m-1}] \right]. \quad (103)$$

We specifically require bounds of these terms without relying on independence of the iterands.

We focus on (103) first. Recall for $m > 0$, that \mathcal{F}_m is the sigma algebra defined in (34). By using the tower property of the conditional expectation and conditioning on \mathcal{F}_{m-1} , from Lemma 1 together with the fact that $T_m < cT_{m-1}$ for some $c > 0$, we obtain directly

$$(103) = (\alpha_{m-1})^2 \mathbb{E} \left[\mathbb{E} \left[|\eta_{m-1}|^2 \mathbb{1}[\mathcal{B}_{m-1}] \mid \mathcal{F}_{m-1} \right] \right] \stackrel{(\text{Lemma 1})}{\leq} (\alpha_{m-1})^2 \frac{c_1}{T_m} \mathbb{E} [\mathcal{L}_4(S_{t_{m-1}})^2 \mathbb{1}[\mathcal{B}_{m-1}]]. \quad (104)$$

Let us next bound (102). Note that this term does not vanish due to dependence of the samples conditional on \mathcal{F}_{m-1} . In our case, however, we have a Markov chain trajectory whose kernel will depend on Θ_{m-1} . Let

$$Z_{m-1} = \Theta_{m-1} - \mathfrak{p}(\Theta_{m-1}) - \alpha_{m-1} \nabla J(\Theta_{m-1}). \quad (105)$$

We use the law of total expectation again on (102). Note that Z_{m-1} and \mathcal{B}_{m-1} are \mathcal{F}_{m-1} -measurable.

$$\begin{aligned} (102) &\leq 2\alpha_{m-1} \mathbb{E} \left[\left\langle \mathbb{1}[\mathcal{B}_{m-1}] Z_{m-1}, \mathbb{E}[\eta_{m-1} \mid \mathcal{F}_{m-1}] \right\rangle \right] \\ &\stackrel{(i)}{\leq} 2\alpha_{m-1} \mathbb{E} \left[|Z_{m-1}|^2 \mathbb{1}[\mathcal{B}_{m-1}] \right]^{1/2} \mathbb{E} \left[|\mathbb{E}[\eta_{m-1} \mid \mathcal{F}_{m-1}]|^2 \right]^{1/2} \\ &\stackrel{(ii)}{\leq} 2\alpha_{m-1} \mathbb{E} \left[|Z_{m-1}|^2 \mathbb{1}[\mathcal{B}_{m-1}] \right]^{1/2} \mathbb{E} \left[\mathbb{1}[\mathcal{B}_{m-1}] \mathcal{L}_4(S_{t_{m-1}})^2 \right]^{1/2} \frac{c_2}{T_m} \end{aligned} \quad (106)$$

where (i) have used Cauchy–Schwartz and (ii) Lemma 1 and the fact that for some $c > 0$, $T_m < cT_{m-1}$.

The terms in (104) and (106) containing $\mathcal{L}_4(S_{t_m})$ can be upper bounded as follows. From the definition of (37) and since $v \geq 16$, by a generalized mean inequality and the fact that $\mathcal{L}(s, a) \geq 1$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$ we have

$$\mathcal{L}_4(s) \leq \mathcal{L}_v(s)^{4/v} \leq \mathcal{L}_v(s)^{1/4}. \quad (107)$$

Now, by Lemma 6, there exists $D > 0$ such that for all $m \in \mathbb{N}$

$$\mathbb{E} \left[\mathbb{1}[\mathcal{B}_{m-1}] \mathcal{L}_4(S_{t_{m-1}})^2 \right] \leq \mathbb{E} \left[\mathbb{1}[\mathcal{B}_{m-2}] \mathcal{L}_4(S_{t_{m-1}})^2 \right] \stackrel{(107)}{\leq} \mathbb{E} \left[\mathbb{1}[\mathcal{B}_{m-2}] \mathcal{L}_v(S_{t_{m-1}}) \right] \leq D. \quad (108)$$

For the other term in (106), we can use the same bound used in [14, (41)]: There exists constants $y, c > 0$ depending on J, θ^* and \mathfrak{r}_0 such that on the event \mathcal{B}_{m-1} we have

$$\begin{aligned} |Z_{m-1}|^2 &\leq (1 - \alpha_{m-1}y)^2 \text{dist}(\Theta_{m-1}, \mathcal{M} \cap U)^2 + c(1 - \alpha_{m-1}y)\alpha_{m-1} \text{dist}(\Theta_{m-1}, \mathcal{M} \cap U)^3 \\ &\quad + c(\alpha_{m-1})^2 \text{dist}(\Theta_{m-1}, \mathcal{M} \cap U)^4. \end{aligned} \quad (109)$$

The bound in (109) characterizes the fact that, close to the manifold of maximizers, the projection is differentiable and can be approximated by an orthogonal expansion of J around the manifold of maximizers. The error terms of this expansion can be bounded depending on the Hessian at $\mathfrak{p}(\Theta_{m-1}) \in \mathcal{M} \cap U$, $\text{Hess}_{\mathfrak{p}(\Theta_{m-1})} J$. We refer to [14, Proposition 17] for a proof of this fact.

We will now use an induction argument to show the claim of the lemma. Namely, we will assume for the time being that for $m-1$ we have

$$\mathbb{E} \left[(\text{dist}(\Theta_{m-1}, \mathcal{M} \cap U) \wedge \delta)^2 \mathbb{1}[\mathcal{B}_{m-1}] \right] \leq \delta^2 c(\alpha) (m-1)^{-\sigma-\kappa}, \quad (110)$$

where $c(\alpha) > 0$ is a function of a to be determined. We want to show (110) for m . To do so we will use (109) to bound Z_{m-1} . Suppose that there exists a sequence $\{b_l\}_{l>0} \subset \mathbb{R}_+$ such that we have

$$\mathbb{E} \left[|Z_{m-1}|^2 \mathbb{1}[\mathcal{B}_{m-1}] \right] \leq b_{m-1}. \quad (111)$$

Using (111) in (106) yields that for some $c_3 > 0$ we have:

$$(102) \leq 2(b_{m-1})^{1/2} \alpha_{m-1} D^{1/2} \frac{c_3}{T_m}. \quad (112)$$

From the expansion of (101) and combining the bounds of (109) and (112) together we obtain

$$\begin{aligned} \mathbb{E} \left[\text{dist}(\Theta_m, \mathcal{M} \cap U)^2 \mathbb{1}[\mathcal{B}_{m-1}] \right] &\leq b_{m-1} + 2(b_{m-1})^{1/2} \alpha_{m-1} D^{1/2} \frac{c_3}{T_m} \\ &\quad + (\alpha_{m-1})^2 \frac{c_4}{T_m} D. \end{aligned} \quad (113)$$

We show now that from the induction hypothesis, if (110) holds, then we also have the bound

$$b_{m-1} \leq c(\alpha) \delta^2 m^{-\sigma-\kappa} - \delta^2 \frac{\alpha y}{2} c(\alpha) (m-1)^{-\sigma-\kappa} m^{-\sigma}. \quad (114)$$

Indeed, taking expectations in (109) and using the bound (110) yields

$$\begin{aligned} b_{m-1} &\leq (1 - \alpha_{m-1}y)^2 c(\alpha) (m-1)^{-\sigma-\kappa} + c(\alpha) (1 - \alpha_{m-1}y) \alpha_{m-1} \delta c(\alpha) (m-1)^{-\sigma-\kappa} \\ &\quad + c(\alpha) (\alpha_{m-1})^2 \delta^2 c(\alpha) (m-1)^{-\sigma-\kappa}. \end{aligned} \quad (115)$$

Recall that $\alpha_{m-1} = \alpha m^{-\sigma/2-\kappa}$. Adding and subtracting $c(\alpha) m^{-\sigma-\kappa}$ in (115), we obtain that

$$\begin{aligned} b_{m-1} &\leq c(\alpha) m^{-\sigma-\kappa} \\ &\quad + c(\alpha) m^{-\sigma} (m-1)^{-\sigma-\kappa} \left(m^\sigma - (m-1)^{\sigma+\kappa} m^{-\kappa} - 2\alpha y + \frac{\alpha^2 y}{m^\sigma} + \left(1 - \frac{\alpha y}{m^\sigma}\right) \alpha \delta + \delta^2 \frac{\alpha^2 y^2}{m^\sigma} \right) \end{aligned}$$

Note now that there exists $m_0(a) > 0$ such that if $m \geq m_0(a)$, we have

$$m^\sigma - (m-1)^{\sigma+\kappa}m^{-\kappa} - \alpha y + \frac{\alpha^2 y}{m^\sigma} < -\frac{\alpha y}{2}. \quad (116)$$

Indeed, note that the latter equation can be satisfied for $m \geq m_0(a)$ since there exists a constant $c > 0$ depending on σ and κ such that

$$\begin{aligned} m^\sigma - (m-1)^{\sigma+\kappa}m^{-\kappa} &\leq m^{-\kappa}(m^{\sigma+\kappa} - (m-1)^{\sigma+\kappa}) \\ &\leq m^{-\kappa}(\sigma + \kappa) \max[(m-1)^{\sigma+\kappa-1}, m^{\sigma+\kappa-1}] \\ &\leq c_5(\sigma + \kappa)m^{\sigma-1}. \end{aligned} \quad (117)$$

In this case we have that

$$m_0(\alpha) = \left(\frac{2c_5(\sigma + \kappa)}{y\alpha} \right)^{1-\sigma} > \frac{c'}{\alpha^{1-\sigma}}. \quad (118)$$

Then for $m > m_0(\alpha)$, we will have

$$b_m \leq c(\alpha)m^{-\sigma-\kappa} + c(\alpha)m^{-\sigma}(m-1)^{-\sigma-\kappa} \left(-\frac{\alpha y}{2} + \left(1 - \frac{\alpha y}{m^\sigma}\right)\alpha\delta + \delta^2 \frac{\alpha^2 y^2}{m^\sigma} \right).$$

Choose $\delta \in (0, \delta_1(\alpha)]$, where $\delta_1(\alpha)$ is a bound that we will choose appropriately, such that for any $m \geq m_0(\alpha)$ we have

$$\left(1 - \frac{\alpha y}{m^\sigma}\right)\alpha\delta + \delta^2 \frac{\alpha^2 y^2}{m^\sigma} \leq \alpha y. \quad (119)$$

Thus, from (110) we obtain (114). With (114) with an appropriate choice of $c(\alpha)$, we can now show (110) for m . We will namely choose $c(\alpha)$ as follows

$$c(\alpha) = \max\left(\frac{c'}{\alpha^{(1-\sigma)(\sigma+\kappa)}}, \frac{4C^2D + 4yCD\alpha\ell}{\delta^2\ell^2y^2} \right), \quad (120)$$

where recall that $\delta \in (0, \delta_1(\alpha)]$ and $\delta_1(\alpha)$ were chosen so that (119) holds. Let $L = \ell^{-1}$. Substituting the bound of (114) into (113) and recalling that $T_m = m^{\kappa+\sigma/2}\ell$ yields

$$\begin{aligned} \mathbb{E}\left[(\text{dist}(\Theta_m, \mathcal{M} \cap U))^2 \mathbb{1}[\mathcal{B}_{m-1}] \right] &\leq c(\alpha)\delta^2 m^{-\sigma-\kappa} - \frac{\alpha y}{2} c(\alpha)\delta^2 (m-1)^{-\sigma-\kappa} m^{-\sigma} \\ &\quad + 2(c(\alpha)\delta^2 m^{-\sigma-\kappa} - \frac{\alpha\lambda}{2} c(\alpha)\delta^2 (m-1)^{-\sigma-\kappa} m^{-\sigma})^{1/2} \alpha m^{-\sigma} D^{1/2} \frac{c_3}{T_m} + D m^{-2\sigma} \frac{\alpha^2 c_3}{T_m} \\ &\leq c(\alpha)\delta^2 m^{-\sigma-\kappa} + m^{-\sigma} (2\sqrt{c(\alpha)}\delta c_3 a D^{1/2} L m^{-\sigma-3\kappa/2} \\ &\quad + c_3 D \alpha^2 L m^{-3\sigma/2-\kappa} - c(\alpha)\delta^2 \alpha y (m-1)^{-\sigma-\kappa}) \\ &\leq c(\alpha)\delta^2 m^{-\sigma-\kappa} + m^{-\sigma} (m-1)^{-\sigma-\kappa} (2\sqrt{c(\alpha)}\delta c_3 a D^{1/2} L + c_3 D \alpha^2 L - c(\alpha)\delta^2 \alpha y). \end{aligned} \quad (121)$$

By the choice of $c(\alpha)$ in (120), for any $\kappa \geq 0$ we have the following inequality

$$2\sqrt{c(\alpha)}\delta c_5 D^{1/2} L + c_5 D \alpha L - c(\alpha)\delta^2 y < 0. \quad (122)$$

Hence, with this choice of $c(\alpha)$, in (121) the latter term in the right-hand side is negative for any $m \geq 2$ and the induction step follows if $m > m_0(\alpha)$. That is, we have for some $c > 0$ that and when $m > m_0(\alpha)$ that

$$\mathbb{E}\left[\text{dist}(4\Theta_m, \mathcal{M} \cap U)^2 \mathbb{1}[\mathcal{B}_{m-1}] \right] \leq c \max\left(\frac{\delta^2}{\alpha^{(1-\sigma)(\sigma+\kappa)}}, \frac{(1 + \alpha\ell)}{\ell^2} \right) m^{-\sigma-\kappa}. \quad (123)$$

We have left to show that the induction hypothesis holds in (110) for some m . Recall that $m > m_0(\alpha)$ is the only restriction we needed on the starting point for the induction argument to work— δ was already chosen depending on α in (119). From the choice

$$m_0(\alpha) \geq \frac{c'}{\alpha^{1-\sigma}}, \quad (124)$$

if $m \leq m_0(\alpha)$, the following slightly changed version of (110) will hold; namely

$$\mathbb{E}\left[(\text{dist}(\Theta_m, \mathcal{M} \cap U)^2 \wedge \delta^2) \mathbf{1}[\mathcal{B}_{m-1}]\right] \leq \delta^2 c(\alpha) m^{-\sigma-\kappa}. \quad (125)$$

Hence, by same arguments conducted with (125) instead of (110), we have shown by induction that (125) holds for $m > 0$.

For convenience, we will further show that there exists a constant $c_6 > 0$ such that for all $m > 0$ we have

$$\mathbb{E}\left[(\text{dist}(\Theta_m, \mathcal{M} \cap U)^2 \wedge \delta^2) \mathbf{1}[\mathcal{B}_{m-1}]\right] \leq c_6 m^{-\sigma-\kappa}. \quad (126)$$

Fix $c_6 > 0$. Choose $\delta_0 \leq \delta_1(\alpha)$ depending on α small enough and $\ell_0 > 0$ large enough such that for $\delta \in (0, \delta_0]$ and $\ell \in [\ell_0, \infty)$ we have that

$$\begin{aligned} \frac{c'\delta^2}{\alpha^{(1-\sigma)(\sigma+\kappa)}} &< c_6 \\ \frac{c(1+\alpha\ell)}{\ell^2} &< c_6, \end{aligned} \quad (127)$$

With the conditions in (127), the proof of the lemma follows noting that $\delta^2 c(\alpha) = \delta^2 c(\alpha, \ell) < c_6$.

C.5 Proof of Lemma 3

We will again use the notation that $t_{m+1} - t_m = T_m$ and without loss of generality assume that $T_m = \ell m^{\sigma/2+\kappa}$ as in Appendix C.4. The proof of Lemma 3 also mainly follows the steps of [14]. However, we again need to take care of the terms that the bias and lack of independence generate in the analysis.

The bounding starts noting the inequality

$$\mathbb{E}\left[\max_{1 \leq l \leq m} |\Theta_l - \Theta_0| \mathbf{1}[\mathcal{B}_{l-1}]\right] \leq \sum_{l=1}^m \mathbb{E}[|\Theta_l - \Theta_{l-1}|^2 \mathbf{1}[\mathcal{B}_{l-1}]]^{1/2}. \quad (128)$$

We will show that there exists a constant $c > 0$ such that for $l \in [m]$ we have

$$\mathbb{E}[|\Theta_{l+1} - \Theta_l|^2 \mathbf{1}[\mathcal{B}_l]]^{1/2} \leq c\alpha \left(l^{-3/2\sigma-\kappa/2} + \sqrt{\frac{1}{\ell}} l^{-5\sigma/8-\kappa/2} \right), \quad (129)$$

where the exponents of σ and κ already differ from the result in [14], and are required to account for the lack of independence and bias. Following the steps from [14], in the neighborhood $V_{\tau, \delta}(\theta^*)$, for each $l \leq m$ there is a random variable $\epsilon_l : \mathcal{B}_l \rightarrow \mathbb{R}^n$ and there exists a constant $c > 0$ such that

$$|\epsilon_l| < c \text{dist}(\Theta_l, \mathcal{M} \cap U)^2 \quad (130)$$

and such that on the event \mathcal{B}_l we have

$$\nabla J(\Theta_l) = \text{Hess}_{\mathbf{p}(\Theta_l)}(\Theta_l - \mathbf{p}(\Theta_l)) + \epsilon_l. \quad (131)$$

Recalling the definition of η_l in (36), we have then the equality

$$\Theta_{l+1} = \Theta_l - \alpha_l \text{Hess}_{\mathbf{p}(\Theta_l)}(\Theta_l - \mathbf{p}(\Theta_l)) - \alpha_l \epsilon_l + \alpha_l \eta_l \quad (132)$$

Define

$$\tilde{\Theta}_l = \Theta_l - \alpha_l \text{Hess}_{\mathbf{p}(\Theta_l)}(\Theta_l - \mathbf{p}(\Theta_l)). \quad (133)$$

We use the triangle inequality with in (132) separating $\Theta_{l+1} - \Theta_l$ as the summands of $\Theta_{l+1} - \tilde{\Theta}_l$ and $\tilde{\Theta}_l - \Theta_l$.

We estimate first $|\Theta_{l+1} - \tilde{\Theta}_l|^2$. In our case, after expanding $\mathbb{E}[|\Theta_{l+1} - \tilde{\Theta}_l|^2 \mathbf{1}[\mathcal{B}_l]]$, we diverge from [14, (58)] and we need to bound

$$\alpha_l^2 \mathbb{E} \left[\mathbf{1}[\mathcal{B}_l] \langle \epsilon_l, \eta_l \rangle \right]. \quad (134)$$

Similar to the proof of Lemma 2, we can condition on \mathcal{F}_l and using that ϵ_l and \mathcal{B}_l are \mathcal{F}_l -measurable together with the Cauchy–Schwartz inequality, we have

$$\begin{aligned} \alpha_l^2 \mathbb{E} \left[\mathbf{1}[\mathcal{B}_l] \langle \epsilon_l, \eta_l \rangle \right] &\leq \alpha_l^2 \mathbb{E} \left[\langle \mathbf{1}[\mathcal{B}_l] \epsilon_l, \mathbb{E}[\eta_l \mathbf{1}[\mathcal{B}_l] | \mathcal{F}_l] \rangle \right] \\ &\leq \alpha_l^2 \mathbb{E} \left[\mathbf{1}[\mathcal{B}_l] |\epsilon_l|^2 \right]^{1/2} \mathbb{E} \left[|\mathbb{E}[\eta_l \mathbf{1}[\mathcal{B}_l] | \mathcal{F}_l]|^2 \right]^{1/2} \end{aligned} \quad (135)$$

Since $\mathbf{1}[\mathcal{B}_m] \leq \mathbf{1}[\mathcal{B}_{m-1}]$, we can bound

$$\mathbb{E} \left[|\mathbb{E}[\eta_l \mathbf{1}[\mathcal{B}_l] | \mathcal{F}_l]|^2 \right]^{1/2} \stackrel{\text{(Lemma 1)}}{\leq} \mathbb{E} \left[\mathbf{1}[\mathcal{B}_l] \frac{c_2^2}{T_l^2} \mathcal{L}_v(S_{t_l}) \right]^{1/2} \stackrel{\text{(Lemma 6)}}{\leq} \frac{c_2}{T_l}. \quad (136)$$

For the remaining term in (135), recall that on the event \mathcal{B}_l , since $\Theta_l \in V_{\tau, \delta}(\theta^*)$, we have that $\text{dist}(\Theta_l, \mathcal{M} \cap U) \leq \delta$. Hence, we can bound for any $l > 0$ that

$$\begin{aligned} \mathbb{E} \left[\mathbf{1}[\mathcal{B}_l] |\epsilon_l|^2 \right]^{1/2} &\stackrel{\text{(130)}}{\leq} (\alpha_l)^2 \mathbb{E}[\text{dist}(\Theta_l, \mathcal{M} \cap U)^4 \mathbf{1}[\mathcal{B}_l]]^{1/2} \frac{c_3}{T_{l+1}} \\ &\leq (\alpha_l)^2 \delta^2 \mathbb{E}[\text{dist}(\Theta_l, \mathcal{M} \cap U)^2 \mathbf{1}[\mathcal{B}_l]]^{1/2} \frac{c_3}{T_{l+1}} \\ &\leq (\alpha_l)^2 \delta^2 \mathbb{E}[\text{dist}(\Theta_l, \mathcal{M} \cap U)^2 \mathbf{1}[\mathcal{B}_{l-1}]]^{1/2} \frac{c_3}{T_{l+1}} \\ &\stackrel{\text{(Lemma 2)}}{\leq} (\alpha_l)^2 \delta^2 l^{-\sigma/2 - \kappa/2} \frac{c_4}{T_l}. \end{aligned} \quad (137)$$

The estimation of the remaining terms in the expansion of $\mathbb{E}[|\Theta_l - \tilde{\Theta}_{l-1}|^2 \mathbf{1}[\mathcal{B}_{l-1}]]$ can be conducted in the same way as that in [14], to which we refer for the details to the interested reader. Together with the estimate of (137) that accounts for the biases we have that

$$\mathbb{E}[|\Theta_l - \tilde{\Theta}_{l-1}|^2 \mathbf{1}[\mathcal{B}_l]] \leq c_5 (\alpha_l)^2 \delta^2 \mathbb{E} \left[\text{dist}(\Theta_l, \mathcal{M} \cap U)^2 \mathbf{1}[\mathcal{B}_l] \right]$$

$$\begin{aligned}
& + 2\delta \mathbb{E} \left[\text{dist}(\Theta_l, \mathcal{M} \cap U)^2 \mathbb{1}[\mathcal{B}_l] \right]^{1/2} \frac{c_6}{T_l} + (\alpha_l)^2 \frac{c_7}{T_l} \\
& \leq c_8 (\alpha_l)^2 \left[\delta^2 l^{-\sigma-\kappa} + 2\delta l^{-\sigma/2-\kappa/2} \frac{1}{T_l} + \frac{1}{T_l} \right]. \tag{138}
\end{aligned}$$

Substituting $T_l = t_{l+1} - t_l = l^{\kappa+\sigma/2}\ell$ and using $\alpha_l < \alpha_{l-1} = \alpha l^{-\sigma}$ into (138) yields the bound

$$\begin{aligned}
\mathbb{E}[|\Theta_l - \tilde{\Theta}_{l-1}|^2 \mathbb{1}[\mathcal{B}_{l-1}]] & \leq c_9 \frac{\alpha^2}{l^{2\sigma}} \left(\delta^2 \frac{1}{l^{\sigma+\kappa}} + 2\delta \frac{1}{l^{\sigma+3\kappa/2\ell}} + \frac{1}{l^{\kappa+\sigma/2}} \right) \\
& \leq c_{10} \frac{\alpha^2}{l^{5\sigma/4+\kappa\ell}}, \tag{139}
\end{aligned}$$

where in the last inequality we have taken the term with the highest order. Using the previous bounds from Lemma 2 we can show that

$$\mathbb{E}[|\Theta_l - \tilde{\Theta}_l|^2 \mathbb{1}[\mathcal{B}_l]] \leq \alpha_l^2 \mathbb{E}[\text{dist}(\Theta_l, \mathcal{M} \cap U) \mathbb{1}[\mathcal{B}_l]] \leq c_{11} \frac{\alpha^2}{l^{3\sigma+\kappa}}, \tag{140}$$

so that using the triangle inequality and combining the bounds of (139) and (140) we obtain

$$\mathbb{E}[|\Theta_{l+1} - \Theta_l|^2 \mathbb{1}[\mathcal{B}_l]]^{1/2} \leq c_{12} \alpha \left(l^{-3/2\sigma-\kappa/2} + \sqrt{\ell}^{-1} l^{-5\sigma/8-\kappa/2} \right). \tag{141}$$

Hence, since $\sigma \in (2/3, 1)$ adding the bound (141) in (128) yields

$$\begin{aligned}
\mathbb{E} \left[\max_{1 \leq l \leq m} |\Theta_l - \Theta_0| \mathbb{1}[\mathcal{B}_{l-1}] \right] & \leq \sum_{l=1}^m c_{12} \alpha \left(l^{-3/2\sigma-\kappa/2} + \sqrt{\ell}^{-1} l^{-\sigma-\kappa/2} \right) \\
& \leq c_{13} \alpha \left(m^{1-3/2\sigma-\kappa/2} + \sqrt{\ell}^{-1} m^{1-5\sigma/8-\kappa/2} \right).
\end{aligned}$$

C.6 Proof of Lemma 4

The proof mimicks the proof strategy of [14, Prop. 24], but modifications are required due to our Markovian assumptions and appearances of biases. Specifically, we must carefully consider the adverse effects that these biases could have on the probability that the iterates exit the basin of attraction. Concretely, our effort will go into firstly proving the following sufficiently strong analogue of [14, (75)] that is applicable to our problem:

Lemma 7. *There exist constants $c_1, c_2 > 0$ such that*

$$\mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) > \delta, \mathcal{B}_{m-1}] \leq \frac{c_1 \alpha^2}{\delta^2 \ell m^{2\sigma}} \mathbb{P}[\mathcal{B}_{m-1}] + \frac{c_2}{\delta^4 \ell m^{\sigma+\kappa}}. \tag{142}$$

The proof of Lemma 7 can be found in Appendix C.6.1.

Once Lemma 7 has been established, we secondly estimate the combined probability that any of the iterates escape in directions tangential to the manifold. The proof of this fact, which is analogous to [14, (78)–(79)], can be found in Appendix C.6.2.

Lemma 8. *If $\Theta_0 \in V_{\tau/2, \delta}(\theta^*)$, then*

$$\sum_{l=1}^m \mathbb{P}[\text{dist}(\Theta_l, \mathcal{M} \cap U) < \delta, \Theta_l \notin V_{\tau, \delta}(\theta^*), \mathcal{B}_{l-1}] \leq \mathbb{P} \left[\max_{1 \leq l \leq m} |\Theta_l - \Theta_0| \mathbb{1}[\mathcal{B}_{l-1}] > R/2 - 2\delta \right]. \tag{143}$$

Proof that Lemmas 7 and 8 imply Lemma 4. First, note that the recursion

$$\mathbb{P}[\mathcal{B}_m] = \mathbb{P}[\Theta_m \in V_{\tau, \delta}(\theta^*), \mathcal{B}_{m-1}] = \mathbb{P}[\mathcal{B}_{m-1}] - \mathbb{P}[\Theta_m \notin V_{\tau, \delta}(\theta^*), \mathcal{B}_{m-1}] \quad (144)$$

can be iterated whenever we can control and bound the following probabilities

$$\begin{aligned} \mathbb{P}[\Theta_m \notin V_{\tau, \delta}(\theta^*), \mathcal{B}_{m-1}] &= \mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) > \delta, \mathcal{B}_{m-1}] \\ &\quad + \mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) \leq \delta, \Theta_m \notin V_{\tau, \delta}(\theta^*), \mathcal{B}_{m-1}]. \end{aligned} \quad (145)$$

Using Lemma 7 and induction on (144) and (145), it follows that for some $c > 0$,

$$\mathbb{P}[\mathcal{B}_m] \geq \prod_{l=1}^m \left(1 - \frac{c\alpha^2}{\delta^2 \ell l^{2\sigma}}\right)_+ - \sum_{l=1}^m \frac{c}{\ell \delta^4 l^{\sigma+\kappa}} - \sum_{l=1}^m \mathbb{P}[\text{dist}(\Theta_l, \mathcal{M} \cap U) < \delta, \Theta_l \notin V_{\tau, \delta}(\theta^*), \mathcal{B}_{l-1}]. \quad (146)$$

We use Lemma 8 together with Lemma 3 and Markov's inequality to obtain the bound

$$\sum_{l=1}^m \mathbb{P}[\text{dist}(\Theta_l, \mathcal{M} \cap U) < \delta, \Theta_l \notin V_{\tau, \delta}(\theta^*), \mathcal{B}_{l-1}] \leq c\alpha \frac{(m^{1-3/2\sigma-\kappa/2} + \ell^{-1/2} m^{1-5\sigma/8-\kappa/2})}{(\tau/2 - 2\delta)_+} \quad (147)$$

Thus, substituting (147) in (146), for some $c > 0$ we have

$$\mathbb{P}[\mathcal{B}_m] \geq \prod_{l=1}^m \left(1 - \frac{c\alpha^2}{\delta^2 \ell l^{2\sigma}}\right)_+ - \sum_{l=1}^m \frac{c}{\ell \delta^4 l^{\sigma+\kappa}} - c\alpha \frac{(m^{1-3/2\sigma-\kappa/2} + \ell^{-1/2} m^{1-5\sigma/8-\kappa/2})}{(\tau/2 - 2\delta)_+}. \quad (148)$$

Note first that since $\sigma \in (2/3, 1)$ and $\kappa \geq 0$, if $\sigma + \kappa \neq 1$, then there exists a constant $c_1 > 0$ such that

$$\sum_{l=1}^m \frac{c}{\ell \delta^4 l^{\sigma+\kappa}} \leq c_1 m^{1-\sigma-\kappa} \quad (149)$$

Lastly, there also exists a constant $c > 0$, $\alpha_0 > 0$, δ_0 such that if $\alpha \in (0, \alpha_0]$ and $\delta \in (0, \delta_0]$ then there exists $\ell_0 > 0$ such that if $\ell \in [\ell_0, \infty)$ then

$$\prod_{l=1}^m \left(1 - \frac{c\alpha^2}{\delta^2 \ell l^{2\sigma}}\right)_+ \geq \exp\left(-\frac{c\alpha^2}{\delta^2 \ell}\right) \quad (150)$$

Lower bounding (148) using (149) and (150) yields Lemma 4. \square

C.6.1 Proof of Lemma 7

We follow first [14, (69)], by fixing δ_1 small enough such that $\delta \in (0, \delta_1]$, on the event \mathcal{B}_{m-1} it is shown in [14] that we have the inequality

$$\text{dist}(\Theta_m, \mathcal{M} \cap U) \leq \left(1 - \frac{\lambda\alpha_{m-1}}{2}\right) \text{dist}(\Theta_{m-1}, \mathcal{M} \cap U) + \alpha_{m-1} |\eta_{m-1}|. \quad (151)$$

We consider now the event $\{\text{dist}(\Theta_m, \mathcal{M} \cap U) > \delta\} \cap \mathcal{B}_{m-1}$. This event occurs when in (151), either $\Theta_{m-1} \in V_{\tau, \delta/2}(\theta^*)$ and $|\eta_{m-1}| \geq \alpha_{m-1} \delta/2$, or $\Theta_{m-1} \in V_{\tau, \delta}(\theta^*) \setminus V_{\tau, \delta/2}(\theta^*)$ and the gradient term can have smaller size. Mathematically, this translates into the inequality

$$\mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) > \delta, \mathcal{B}_{m-1}] \leq \mathbb{P}\left[|\eta_{m-1}| \geq \frac{\delta}{2\alpha_{m-1}}, \Theta_{m-1} \in V_{\tau, \delta/2}(\theta^*), \mathcal{B}_{m-2}\right] \quad (152)$$

$$+ \mathbb{P}\left[|\eta_{m-1}| \geq \frac{\delta\lambda}{2}, \Theta_{m-1} \in V_{\tau,\delta}(\theta^*) \setminus V_{\tau,\delta/2}(\theta^*), \mathcal{B}_{m-2}\right] =: P_1 + P_2.$$

Contrary to what is done in the proof of [14, Prop. 24], we cannot use an independence property to estimate the probabilities P_1 and P_2 in (152). After all, the Markov chain's behavior at epoch $m-1$ depends on Θ_{m-1} .

In order to overcome this issue we will use the characterization of η_{m-1} in Lemma 1. Recall Lemma 1, and note that it implies

$$\begin{aligned} \mathbb{E}\left[\mathbb{1}[\mathcal{B}_{m-1}]\mathbb{1}\left[|\eta_{m-1}| \geq \frac{\delta}{2\alpha_{m-1}}\right] \middle| \mathcal{F}_{m-1}\right] &= \mathbb{P}\left[|\eta_{m-1}| \geq \frac{\delta}{2\alpha_{m-1}}, \mathcal{B}_{m-1} \middle| \mathcal{F}_{m-1}\right] \\ &\leq \frac{\mathbb{E}[|\eta_{m-1}|^2 \mathbb{1}[\mathcal{B}_{m-1}] \middle| \mathcal{F}_{m-1}]}{\frac{\delta^2}{4(\alpha_{m-1})^2}} \leq \frac{4c_2(\alpha_{m-1})^2 \mathcal{L}_4(S_{t_{m-1}})}{\delta^2 T_m} \end{aligned} \quad (153)$$

since there exist a constant $c > 0$ such that $T_m < cT_{m-1}$.

Bounding P_1 in (152). We can write

$$\begin{aligned} P_1 &\stackrel{(i)}{=} \mathbb{E}\left[\mathbb{1}\left[|\eta_{m-1}| \geq \frac{\delta}{2\alpha_{m-1}}\right] \mathbb{1}[\Theta_{m-1} \in V_{\tau,\delta/2}(\theta^*)] \mathbb{1}[\mathcal{B}_{m-2}] \mathbb{1}[\mathcal{B}_{m-1}]\right] \\ &= \mathbb{E}\left[\mathbb{1}[\Theta_{m-1} \in V_{\tau,\delta/2}(\theta^*)] \mathbb{1}[\mathcal{B}_{m-2}] \mathbb{E}\left[\mathbb{1}[\mathcal{B}_{m-1}] \mathbb{1}\left[|\eta_{m-1}| \geq \frac{\delta}{2\alpha_{m-1}}\right] \middle| \mathcal{F}_{m-1}\right]\right] \\ &\stackrel{(153)}{\leq} \frac{4c_2(\alpha_{m-1})^2}{T_m \delta^2} \mathbb{E}\left[\mathbb{1}[\Theta_{m-1} \in V_{\tau,\delta/2}(\theta^*)] \mathbb{1}[\mathcal{B}_{m-2}] \mathcal{L}_4(S_{t_{m-1}})\right] \end{aligned} \quad (154)$$

where for (i) we have used the fact that $\{\Theta_{m-1} \in V_{\tau,\delta/2}(\theta^*)\} \cap \mathcal{B}_{m-2} \subset \mathcal{B}_{m-1}$.

We deal now with the remaining term in (154). Differently to the independent and unbiased case we need to control the bias and use the tail probability that the Lyapunov function is larger than a certain bound in order to estimate the deviation probability. This step is the crucial different step compared to [14], where we have to explicitly use Assumption 4 and 6. Note that a Cauchy–Schwartz inequality in (154) will not yields an inequality strong enough. See the remark after the proof for further details.

Before bounding the remaining term in (154), we obtain the necessary inequalities. Recall from Lemma 6 that since $\mathbb{E}[\mathcal{L}_4(S_{t_{m-1}})^4 \mathbb{1}[\mathcal{B}_{m-2}]] < \mathbb{E}[\mathcal{L}_v(S_{t_{m-1}}) \mathbb{1}[\mathcal{B}_{m-2}]] < D < \infty$, then by Markov's inequality we have that there exists $D > 0$ such that for any $m > 0$,

$$\mathbb{P}[\mathcal{L}(S_{t_{m-1}}) > m^s, \mathcal{B}_{m-2}] \leq D^4 m^{-4s}. \quad (155)$$

Note also that under the moment assumptions the following holds

$$\begin{aligned} \mathbb{E}\left[\mathcal{L}(S_{t_{m-1}}) \mathbb{1}[\mathcal{B}_{m-2}] \mathbb{1}[\mathcal{L}(S_{t_{m-1}}) > m^s]\right] &= \int_{m^s}^{\infty} \mathbb{P}[\mathcal{L}(S_{t_{m-1}}) > t, \mathcal{B}_{m-2}] dt \\ &= \int_{m^s}^{\infty} \frac{D^4}{t^4} dt \leq D^4 m^{-3s+1}. \end{aligned} \quad (156)$$

We use the (156) to bound (154) as follows

$$\mathbb{E}\left[\mathbb{1}[\Theta_{m-1} \in V_{\tau,\delta/2}(\theta^*)] \mathcal{L}_4(S_{t_{m-1}}) \mathbb{1}[\mathcal{B}_{m-2}]\right]$$

$$\begin{aligned}
&\leq \mathbb{E} \left[\mathbb{1}[\Theta_{m-1} \in V_{\tau, \delta/2}(\theta^*)] \mathcal{L}_4(S_{t_{m-1}}) \mathbb{1}[\mathcal{B}_{m-2}] \left(\mathbb{1}[\mathcal{L}_4(S_{t_{m-1}}) > m^s] + \mathbb{1}[\mathcal{L}_4(S_{t_{m-1}}) \leq m^s] \right) \right] \\
&\leq \mathbb{E} \left[\mathbb{1}[\Theta_{m-1} \in V_{\tau, \delta/2}(\theta^*)] m^s \mathbb{1}[\mathcal{B}_{m-2}] \right] + \mathbb{E} \left[\mathcal{L}(S_{t_{m-1}}) \mathbb{1}[\mathcal{B}_{m-2}] \mathbb{1}[\mathcal{L}(S_{t_{m-1}}) > m^s] \right] \\
&\stackrel{(156)}{\leq} m^s \mathbb{P}[\Theta_{m-1} \in V_{\tau, \delta/2}(\theta^*), \mathcal{B}_{m-2}] + c_3 D m^{-3s+1} \leq m^s \mathbb{P}[\mathcal{B}_{m-1}] + c_3 D m^{-3s+1}. \tag{157}
\end{aligned}$$

Thus, using (157), we can bound P_1 in (152). Specifically,

$$P_1 \leq \frac{4c_4(\alpha_{m-1})^2}{T_m \delta^2} (m^s \mathbb{P}[\mathcal{B}_{m-1}] + m^{-3s+1}). \tag{158}$$

This completes our bound for P_1 .

Bounding P_2 in (152). Repeating the argumentation behind (158), we can show that

$$P_2 \leq \frac{4c_5}{T_m \lambda^2 \delta^2} \left(m^s \mathbb{P}[\Theta_{m-1} \in V_{\tau, \delta}(\theta^*) \setminus V_{\tau, \delta/2}(\theta^*), \mathcal{B}_{m-2}] + m^{-3s+1} \right). \tag{159}$$

Using the facts (i) $\{\Theta_{m-1} \in V_{\tau, \delta}(\theta^*) \setminus V_{\tau, \delta/2}(\theta^*)\} \subseteq \{\text{dist}(\Theta_{m-1}, \mathcal{M} \cap U) \geq \delta/2\}$, with (ii) an application of Lemma 2 and Markov's inequality, reveals that

$$\mathbb{P}[\Theta_{m-1} \in V_{\tau, \delta}(\theta^*) \setminus V_{\tau, \delta/2}(\theta^*), \mathcal{B}_{m-2}] \stackrel{(i)}{\leq} \mathbb{P} \left[\text{dist}(\Theta_{m-1}, \mathcal{M} \cap U) \geq \frac{\delta}{2}, \mathcal{B}_{m-2} \right] \stackrel{(ii)}{\leq} \frac{4}{\delta^2} c_6 m^{-\sigma-\kappa}. \tag{160}$$

Applying the bound in (159) to (160) yields

$$P_2 \leq \frac{4c_7}{T_m \lambda^2 \delta^4} \left(m^s m^{-\sigma-\kappa} + m^{-3s+1} \right). \tag{161}$$

This completes the bound for P_2 in (152).

A return to (152), and parameter selection. Let us now combine (157) and (161) and return to bounding the left-hand side of (152). Specifically, observe that we proved that

$$\begin{aligned}
\mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) > \delta, \mathcal{B}_{m-1}] &\leq \frac{4c_8(\alpha_{m-1})^2}{T_m \delta^2} (m^s \mathbb{P}[\mathcal{B}_{m-1}] + m^{-3s+1}) \\
&\quad + \frac{4c_9}{T_m \delta^4} (m^{s-\sigma-\kappa} + m^{-3s+1}). \tag{162}
\end{aligned}$$

We now specify $s = \kappa + \sigma/2$ in (162). Without loss of generality we will again assume that $T_m = \ell m^{\sigma/2+\kappa}$ instead of $\lfloor \ell m^{\sigma/2+\kappa} \rfloor$ —there is namely only a constant changed. By choosing the smallest exponents in m in (162) for all $m > 0$ we have

$$\mathbb{P}[\text{dist}(\Theta_m, \mathcal{M} \cap U) > \delta, \mathcal{B}_{m-1}] \leq c_{10} \frac{a^2}{\delta^2 \ell m^{2\sigma}} \mathbb{P}[\mathcal{B}_{m-1}] + \frac{c_{10}}{\delta^4 \ell} (m^{-3\sigma-4\kappa+1} + m^{-\sigma-\kappa}). \tag{163}$$

Since $\sigma \in (2/3, 1)$, then $-3\sigma - 4\kappa + 1 < -\sigma - \kappa$ for any $\kappa \geq 0$. Upper bounding the leading orders in m completes the proof of Lemma 7. \square

Remark. A Cauchy–Schwartz inequality in (154) would only yield a factor $\mathbb{P}[\mathcal{B}_{m-1}]^{1/2} > \mathbb{P}[\mathcal{B}_{m-1}]$, which would not be sufficient. Similarly, we could have used Lemma 6 directly and obtain a bound on $\mathbb{E}[\mathbb{1}[\mathcal{B}_{m-2}] \mathcal{L}_4(S_{t_{m-1}})]$. However, this would not give an inequality that can be iterated inductively and is sharp enough. We can directly simplify this term to obtain $\mathbb{P}(\mathcal{B}_{m-1})$ in the inequality only when $\mathcal{L}_4(S_{t_{m-1}})$ is bounded.

C.6.2 Proof of Lemma 8

In [14], it is [14, Lem. 23] that establishes [14, (78)–(79)] directly. Since [14, Lem. 23] is solely a geometric argument, and does not concern the stochastic process, it also applies in our Markovian setting. \square

D The compact case

In the case that the set of maxima \mathcal{M} is compact, we can improve the convergence rate of Theorem 2. We will namely assume the following

Assumption 8 (Compactness, Optional). *The open subset U defined in Assumption 7 is such that $\mathcal{M} \cap U$ is compact.*

Under the additional Assumption 8, in Appendix D.2 we show the following

Theorem 3 (Compact case). *Suppose that Assumptions 1 to 8 hold, except that (17) is now relaxed to allow for $\sigma \in (0, 1)$ and $\kappa \in [0, \infty)$. For every maximizer $\theta^* \in \mathcal{M}$, there exist constants $c > 0$ and $\alpha_0 > 0$ such that, for every $\alpha \in (0, \alpha_0]$, there exists a neighborhood V of θ^* such that there exists $\ell_0 > 0$ such that for any $\ell \in [\ell_0, \infty)$, $m \in \mathbb{N}_+$, and $\epsilon \in (0, 1)$,*

$$\mathbb{P}[J(\Theta_m) < J^* - \epsilon | \Theta_0 \in V] \leq c \left(\epsilon^{-2} m^{-\sigma - \kappa} + \frac{m^{1 - \sigma - \kappa}}{\ell} + \frac{\alpha^2}{\ell} \right). \quad (164)$$

The term proportional to $\alpha m^{-\kappa/2} + \alpha m^{1 - \sigma/2 - \kappa/2} \ell^{-\frac{1}{2}}$ is not in Theorem 3 compared to Theorem 2. This term estimates the probability that the iterates escape V along directions almost parallel to those of \mathcal{M} . As it turns out, in the compact case such event cannot occur. The bound in (164) thus holds when the set of maxima is, for example, a singleton $\mathcal{M} \cap U = \{x_0\}$.

D.1 Sample Complexity

In the compact case, we can similarly obtain a complexity bound like that of Corollary 1.

Corollary 2 (Sample Complexity, Compact case). *Under the same assumptions and setting as in Theorem 3, assume moreover that there exists some $b > 0$ such that $|r(s, a)| < b$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Then we have*

$$\mathbb{E}[J^* - J(\Theta_m) | \Theta_0 \in V] \leq 3 \left(\frac{cb}{2} \right)^{\frac{1}{3}} m^{-\frac{(\sigma + \kappa)}{3}} + 2bc \frac{m^{1 - (\sigma + \kappa)} + \alpha^2}{\ell}.$$

Proof. The proof is the same as that of Corollary 1, where we use (164) instead. \square

D.2 Proof of Theorem 3

The proof is the same as with Theorem 2, but we can omit the last term in (70) by showing that we can choose \mathfrak{r} arbitrarily large. The argument is as follows. If the manifold $\mathcal{M} \cap U$ is compact, it can be covered by a finite number k of local tubular neighborhoods $V_i = V_{\mathfrak{r}_i, \delta_i}(\theta_i)$ where $\theta_i \in \mathcal{M} \cap U$ and $\mathcal{M} \cap U \subset \cup_{i \in [k]} V_i$. Choose $\delta = \min_{i \in [k]} \delta_i$. Then, any $\theta \in U$ such that $\text{dist}(\theta, \mathcal{M} \cap U) < \delta$ will satisfy that $\mathfrak{p}(\theta) \in \mathcal{M} \cap U$, where \mathfrak{p} is the unique local orthogonal projection on $\mathcal{M} \cap U$ from (32).

Now, from compactness, for any $\theta^* \in \mathcal{M} \cap U$ there exists $\tilde{\tau} > 0$ such that $\mathcal{M} \cap U \subset B_{\tilde{\tau}}(\theta^*)$. For any $\tau \geq \tilde{\tau}$ we thus have that $V_{\tau, \delta}(\theta^*) = V_{\tilde{\tau}, \delta}(\theta^*)$ is a tubular neighborhood containing $\mathcal{M} \cap U$. Then, we can choose τ arbitrarily large and conclude that the last term in the bound for the probability in Theorem 2 vanishes if $\mathcal{M} \cap U$ is a compact manifold. More details on tubular neighborhoods and their existence for embedded manifolds can be found in [23].

E Proof of Proposition 1

We consider the following setting. Let $D < 1$. We consider $\theta \in \mathbb{R}$ and a function f such that in $\mathbb{R} \setminus [-D, D]$ satisfies $f(\theta) = 0$ and in $[-D/2, D/2]$ satisfies

$$f(\theta) = 1 - \theta^2. \quad (165)$$

In $[-D, -D/2] \cup [D/2, D]$, we define f such that it is smoothly and monotonically interpolated between $[-D/2, D/2]$ and $\mathbb{R} \setminus [-D, D]$.

We let H_m be such that $H_m = 0$ in $\mathbb{R} \setminus [-D, D]$. Hence, the set $\mathbb{R} \setminus [-D, D]$ is an absorbing set that is 1-suboptimal. In $[-D/2, D/2]$, we will consider $\eta_m = \nabla f(\Theta_m) - H_m$ to be a random variable that, conditional on \mathcal{F}_m , is unbiased and has a second moment for all m but approximates a heavy tailed random variable. In particular, for $\beta > 0$, we define η_m such that there exists $c > 0$ such that for any m , we have

$$\mathbb{P}[|\eta_m| > s | \mathcal{F}_m] \geq \frac{c}{s^{2+\beta} T_m} \quad \text{for } s > D. \quad (166)$$

Note that this constraint on η_m is compatible with the finite second moment condition from (23). If moreover $\alpha \leq 1$ and $\sqrt{\epsilon} < 2D$, then we can bound under the previous conditions

$$\begin{aligned} \mathbb{P}[f(\Theta_m) < f^* - \epsilon | \Theta_0 \in V] &\stackrel{(i)}{\geq} \mathbb{P}[f(\Theta_m) < f^* - \epsilon | \Theta_0 = \theta^{\min}] \\ &= \mathbb{P}[|\Theta_m| > \sqrt{\epsilon} | \Theta_0 = \theta^{\min}] \\ &\stackrel{(ii)}{\geq} \mathbb{P}\left[\sup_{l \leq m} |\Theta_l| > 2D | \Theta_0 = \theta^{\min}\right] \\ &\geq \mathbb{P}[|\Theta_1| > 2D | \Theta_0 = \theta^{\min}] \\ &= \mathbb{P}[|\theta^{\min} + \alpha_1 \eta_1| > 2D | \Theta_0 = \theta_0] \\ &\stackrel{(iii)}{\geq} \mathbb{P}[\alpha_1 |\eta_1| > D | \Theta_0] \\ &\stackrel{(166)}{\geq} c \frac{\alpha_1^{2+\beta}}{D^{2+\beta} T_1} \\ &\geq c \frac{\alpha^{2+\beta}}{D^{2+\beta} \ell}, \end{aligned} \quad (167)$$

where (i) we have used that for any $V = [-\delta, \delta]$ with $\delta < D$,

$$\begin{aligned} \mathbb{P}[f(\Theta_m) < f^* - \epsilon | \Theta_0 \in V] &= \int_{\theta \in V} \mathbb{P}[f(\Theta_m) < f^* - \epsilon | \Theta_0 = \theta] d\mathbb{P}[\Theta_0 = \theta | \Theta_0 \in V] \\ &\geq \min_{\theta \in V} \mathbb{P}[f(\Theta_m) < f^* - \epsilon | \Theta_0 = \theta] \end{aligned}$$

$$\geq \mathbb{P}[f(\Theta_m) < f^* - \epsilon | \Theta_0 = \theta^{\min}] \quad (168)$$

for some $\theta^{\min} \in V$. In (ii), we have used the fact that from the definition of f , we have the inclusion of events $\{\sup_{l \leq m} |\Theta_l| > 2D\} \in \{|\Theta_m| > 2D\}$, since the set $\mathbb{R} \setminus [-D, D]$ is absorbent for the process $\{\Theta_t\}_{t \geq 0}$. In (iii), we have used that θ^{\min} belongs at least to $[-D, D]$, since otherwise it cannot be the minimum as defined in (168). To guarantee that $\epsilon \in (0, 1)$ we may choose $D = 1/2$, for example.

F Proof of Proposition 2

To prove the proposition we will show that for almost all $\tilde{\pi}$ in the Lebesgue measure of the class of policies defined in (6), the function $J_{\tilde{\pi}}(\theta)$ is Morse. Morse functions are smooth functions f such that every critical point of f is nondegenerate, that is, for any x such that $\nabla_x f = 0$ we have that $\text{Hess}_x f$ is nonsingular. Hence, all critical points are isolated. If the function $J_{\tilde{\pi}}(\theta)$ is Morse and furthermore satisfies that $J_{\tilde{\pi}}(\theta) \rightarrow -\infty$ as $|\theta| \rightarrow \infty$, it will then have bounded isolated maxima.

We show first that for almost all $\tilde{\pi}$, the function $J_{\tilde{\pi}}(\theta)$ is a Morse function. To do so, we will implicitly use the fact that Morse functions are dense and form an open subset in the space of smooth functions (see [29]).

We introduce first notation. For a finite dimensional smooth manifold M , we denote by $T_x M$ and $T_x^* M$ the tangent and cotangent spaces at $x \in M$, respectively. When $M = \mathbb{R}^u$, for $f : \mathbb{R}^u \rightarrow \mathbb{R}$ we will denote the (covariant) derivative and gradient of f at x by $d_x f \in T_x^* M$ and $\nabla_x f \in T_x M$, respectively. In local coordinates (w_1, \dots, w_u) , we have namely

$$\begin{aligned} d_x f &= \sum_{i=1}^u \frac{\partial f(x)}{\partial w_i} dw_i \\ \nabla_x f &= \sum_{i=1}^u \frac{\partial f(x)}{\partial w_i} \frac{d}{dw_i}, \end{aligned} \quad (169)$$

where $dw_i(\frac{d}{dw_i}) = \mathbb{1}[i = j]$. In this notation and since $M = \mathbb{R}^u$, we have then

$$d_x(df) = \sum_{i=1}^u d_x \left(\frac{\partial f(x)}{\partial w_i} dw_i \right) = \sum_{i=1}^u \sum_{j=1}^u \frac{\partial^2 f(x)}{\partial w_j \partial w_i} dw_j \otimes dw_i = \text{Hess}_x f \in T_x^* M \otimes T_x^* M. \quad (170)$$

We require the following lemmas and definitions.

Definition 1. Let M and N be two manifolds and let B be a submanifold of N . We say a smooth map $f : M \rightarrow N$ is transversal to B if for every point $x \in M$ such that $f(x) \in B$ we have

$$d_x f(T_x M) + T_{f(x)} B = T_{f(x)} N. \quad (171)$$

We will use the following result that has its core an application of Sard's theorem that states that in a map between smooth manifolds, the set of critical points has measure zero in the image.

Lemma 9 (Parametric transversality theorem [17]). *Let Z, M and N be smooth manifolds and let B be a smooth submanifold of N . Let $F : Z \times M \rightarrow N$ be a smooth submersion, that is, the*

differential map is surjective everywhere. If F is transversal to B , then for almost every $z \in Z$, the map

$$F_z(m) = F(z, m) \quad (172)$$

is transversal to B .

When appropriate, we will make explicit the dependence of $v \in T_x^*M$ on x by writing $(x, v) \in T_x^*M$. We can now show the following,

Lemma 10. *Let $M = \mathbb{R}^u$ and let $f : M \rightarrow \mathbb{R}$ be a smooth map. Consider the map $\tilde{f} : M \rightarrow T^*M$ given for $x \in M$ by*

$$\tilde{f}(x) = (x, d_x f) \in T_x^*M. \quad (173)$$

*Let $B \subset T^*M$ be the zero section submanifold, that is, $B(x) = (x, 0) \in T_x^*M$ for every x . Then x is a nondegenerate critical point of f if and only if \tilde{f} is transversal to B at x and $\nabla_x f = 0$.*

Proof. x is a critical nondegenerate point if and only if $\nabla_x f = 0$ and $\text{Hess}_x f \in T_x^*M \otimes T_x^*M$ is nonsingular. For any $\nu \in T_x M$, we have then that

$$d_x \tilde{f}(\nu) = (\nu, \text{Hess}_x f(\nu)) \quad (174)$$

By definition, \tilde{f} is transversal to B if and only if for every $x \in M$,

$$\begin{aligned} d_x \tilde{f}(T_x M) + T_x M \oplus 0 &= (Id \oplus \text{Hess}_x(f))(T_x M) + T_x M \oplus 0 \\ &= T_x M \oplus \text{Hess}_x f(T_x M) \\ &= T_x M \oplus T_x^* M, \end{aligned} \quad (175)$$

which is true if and only if $\text{Hess}_x f$ is nonsingular. \square

From the last two lemmas it follows that by adding an appropriate perturbation to a function, the perturbed function is nondegenerate. This result is well-known in the literature in the context of genericity of Morse functions and can be generalized to general smooth manifolds; see [17].

Lemma 11. *Let $M = \mathbb{R}^u$. Let $f : M \rightarrow \mathbb{R}$ and $g_i : M \rightarrow \mathbb{R}$ for $i \in [l]$ be smooth functions such that for every $x \in M$, $\text{span}(\{d_x g_i\}_{i=1}^l) = T_x^*M$. Then for almost every $z = (z_1, \dots, z_l) \in \mathbb{R}^l$ we have that*

$$f_z(\cdot) = f(\cdot) + \sum_{i=1}^l z_i g_i(\cdot) \quad (176)$$

is a Morse function.

Proof. Define the smooth function $F : \mathbb{R}^l \times M \rightarrow T^*M$ given by

$$F(z, x) = (x, d_x f + \sum_{i=1}^l z_i d_x g_i) = (x, d_x f_z). \quad (177)$$

The derivative of this map at (z, x) evaluated at $(\eta, \chi) \in T_z \mathbb{R}^l \times T_x M$ is then

$$d_{(z,x)} F(\eta, \chi) = (\chi, \text{Hess}_x f_z(\chi) + \sum_{i=1}^l \eta_i d_x g_i) \in T_{F(z,x)}(T^*M) \simeq T_x M \oplus T_x^* M. \quad (178)$$

For every x , we have $\text{span}(\{d_x g_i\}_{i=1}^l) = T_x^* M$, then $d_{(z,x)} F(T_z \mathbb{R}^l, T_x M) = T_{F(z,x)}(T^* M)$ and $d_{(z,x)} F$ is surjective. Thus, F is a submersion and is therefore transversal to the zero section of $T^* M$ and by Lemma 9 for almost every $z \in Z$ the map $F_z(x) = F(z, x)$ is transversal to the zero section of $T^* M$. Finally, by Lemma 10 we can conclude that for almost every $z \in Z$, the critical points of f_z are nondegenerate, that is, f_z is a Morse function. \square

We are now in position to show the proposition. Recall from the definition of the policy in (6) that there is an index set \mathcal{I} and a function $h : \mathcal{S} \rightarrow \mathcal{I}$ that determines the parameter dependence of $\{\theta_{i,a} : (i, a) \in \mathcal{I} \times \mathcal{A}\}$. For $s \in \mathcal{I}$, let $z_{(s,i)} = \tilde{\pi}(a|i)$ and denote $\tilde{\zeta}(i) = \sum_{s \in \mathcal{S}: h(s)=i} \zeta(s)$. We can write

$$\begin{aligned}
d_\theta \mathcal{R}_{\tilde{\pi}}(\theta) &= b \sum_{s \in \mathcal{S}} \zeta(s) \sum_{a \in \mathcal{A}} \tilde{\pi}(a|s) d_\theta \log(\pi(a|s, \theta)) \\
&= b \sum_{s \in \mathcal{S}} \zeta(s) \sum_{a \in \mathcal{A}} \tilde{\pi}(a|s) \left(\sum_{a' \in \mathcal{A}} (\mathbb{1}[a = a'] - \pi(a'|s, \theta)) d\theta_{h(s), a'} \right) \\
&= b \sum_{s \in \mathcal{S}} \zeta(s) \sum_{a' \in \mathcal{A}} (\tilde{\pi}(a|s) - \pi(a'|s, \theta)) d\theta_{h(s), a'} \\
&= b \sum_{i \in \mathcal{I}} \sum_{a \in \mathcal{A}} \tilde{\zeta}(i) (\tilde{\pi}(a|i) - \pi(a|i, \theta)) d\theta_{i,a} \\
&= b \sum_{(i,a) \in \mathcal{I} \times \mathcal{A}} \tilde{\zeta}(i) (z_{(i,a)} - \pi(a|i, \theta)) d\theta_{i,a}
\end{aligned} \tag{179}$$

If $\tilde{\zeta}(i) > 0$ for all $i \in \mathcal{I}$, it is clear from (179) that the terms $\{d\theta_{i,a}\}_{(i,a) \in \mathcal{I} \times \mathcal{A}}$ span $T_\theta^* \mathbb{R}^{|\mathcal{A}| \times |\mathcal{I}|}$ for each θ , since $\pi(a|s, \theta) \neq 0$ for any finite θ . By Lemma 11 and the assumption on ζ , we immediately obtain that for almost all policies $\tilde{\pi}$, the function

$$J_{\tilde{\pi}}(\theta) = J(\theta) - b \mathcal{R}_{\tilde{\pi}}(\theta). \tag{180}$$

is Morse and has nondegenerate critical points—including the maximum. Finally, the set of maxima of (180) will be nonempty. Indeed, the function $-b \mathcal{R}_{\tilde{\pi}}(\theta) \rightarrow -\infty$ whenever for any $s \in \mathcal{S}$, $\pi(\cdot|s) \rightarrow \partial \Delta(\mathcal{S})$. Thus, by continuity, the set of maxima belongs to a compact set.