

GLOBAL CONVERGENCE OF POLICY GRADIENT IN AVERAGE REWARD MDPs

Anonymous authors

Paper under double-blind review

ABSTRACT

We present the first comprehensive finite-time global convergence analysis of policy gradient for infinite horizon average reward Markov decision processes (MDPs). Specifically, we focus on ergodic tabular MDPs with finite state and action spaces. Our analysis shows that the policy gradient iterates converge to the optimal policy at a sublinear rate of $O\left(\frac{1}{T}\right)$, where T represents the number of iterations. Performance bounds for discounted reward MDPs cannot be easily extended to average reward MDPs as the bounds grow proportional to the fifth power of the effective horizon. Recent work on such extensions makes a smoothness assumption that has not been verified. Thus, our primary contribution is in providing the first complete proof that the policy gradient algorithm converges globally for average-reward MDPs, without such an assumption. We also obtain the corresponding finite-time performance guarantees. In contrast to the existing discounted reward performance bounds, our performance bounds have an explicit dependence on constants that capture the complexity of the underlying MDP. Motivated by this observation, we reexamine and improve the existing performance bounds for discounted reward MDPs. We also present simulations that empirically validate the result.

1 INTRODUCTION

Average reward Markov Decision Processes (MDPs) find applications in domains where decisions are made over time to optimize long-term performance. Some of these applications include resource allocation, portfolio management in finance, healthcare, and robotics (Ghalme et al., 2021; Bielecki et al., 1999; Patrick & Begen, 2011; Mahadevan, 1996; Tadepalli & Ok, 1998). Approaches for determining the optimal policy can be broadly categorized into dynamic programming algorithms (such as value and policy iteration (Murthy et al., 2024; Abbasi-Yadkori et al., 2019; Gosavi, 2004)) and gradient-based algorithms. Although gradient-based algorithms are heavily used in practice (Schulman et al., 2015; Baxter & Bartlett, 2000), the theoretical analysis of their global convergence is a relatively recent undertaking.

While extensive research has been conducted on the global convergence of policy gradient methods in the context of discounted reward MDPs (Agarwal et al., 2020; Khodadadian et al., 2021), comparatively less attention has been given to its average reward counterpart. Contrary to average reward MDPs, the presence of a discount factor ($\gamma < 1$) serves as a source of contraction that alleviates the technical challenges involved in analyzing the performance of various algorithms in the context of discounted reward MDPs. Consequently, many algorithms designed for average reward MDPs are evaluated using the framework of discounted MDPs, where the discount factor approaches one (Grand-Clément & Petrik, 2024).

In the context of discounted reward MDPs, the projected policy gradient (PPG) algorithm converges as

$$\rho_\gamma^* - \rho_\gamma^{\pi_k} \leq O\left(\frac{1}{(1-\gamma)^5}\right) \quad (1)$$

where γ is the discount factor, π_k is the policy obtained at the k -th iteration of the PPG algorithm, ρ_γ^* represents the optimal value function, and $\rho_\gamma^{\pi_k}$ represents the value function iterates obtained through projected gradient ascent (Xiao, 2022b). Let ρ^π denote the average reward associated with

054 some policy π . It is well known that $\rho^\pi = \lim_{\gamma \rightarrow 1} (1 - \gamma) \rho_\gamma^\pi$ under some mild conditions (Puterman,
 055 1994; Bertsekas, 2007). Utilizing this relationship in equation 1, we observe the upper bound tends
 056 to infinity in the limit $\gamma \rightarrow 1$. Hence, it is necessary to devise an alternate approach to study the
 057 convergence of policy gradient in the context of average reward MDPs.

059 1.1 RELATED WORK

061 There is a wealth of literature on discounted reward MDPs. Fazel et al. (2018) were among the first to
 062 establish the global convergence of policy gradients, specifically within the domain of linear quadratic
 063 regulators. Bhandari & Russo (2024) established a connection between the policy gradient and policy
 064 iteration objectives, determining conditions under which policy gradient algorithms converge to
 065 the globally optimal solution. Agarwal et al. (2020) offer convergence bounds of $O(\frac{1}{\epsilon^2(1-\gamma)^6})$ for
 066 policy gradient and $O(\frac{1}{\epsilon(1-\gamma)^2})$ for natural policy gradient, where ϵ represents the suboptimality. It is
 067 noteworthy that while their convergence bounds for policy gradient rely on the cardinality of the state
 068 and action space, the convergence bounds for natural policy gradient are independent of them. Xiao
 069 (2022b) enhance the $O(\frac{1}{\epsilon^2(1-\gamma)^6})$ policy gradient bounds by refining the dependency on the discount
 070 factor, yielding improved bounds of $O(\frac{1}{\epsilon(1-\gamma)^5})$. Zhang et al. (2021) prove that variance-reduced
 071 versions of stochastic policy gradient also converge to the global optimal solution. They achieve
 072 this through a gradient truncation mechanism. Mei et al. (2020) analyze global convergence of
 073 softmax-based gradient methods and prove exponential rejection of suboptimal policies.

074 The global convergence of policy gradient methods has been extensively studied in the context of
 075 planning for average reward Markov Decision Processes (MDPs). Even-Dar et al. (2009) and Murthy
 076 & Srikant (2023) provide foundational results for natural policy gradient methods, proving their global
 077 convergence for finite state and action spaces. Extending this work, Grosz et al. (2024) analyze the
 078 more challenging case of infinite state spaces, establishing theoretical guarantees for convergence
 079 in this general setting. A critical contribution of our work lies in proving that the smoothness of
 080 the average reward holds for a large class of MDPs in the tabular setting. This smoothness property
 081 allows us to establish the first global convergence bounds for policy gradient methods in average
 082 reward MDPs, eliminating the need for previously unverified and restrictive smoothness assumptions.

083 In the learning context, actor-critic and gradient-based methods have been a central focus for analyzing
 084 average reward MDPs. Konda & Tsitsiklis (1999) investigate two-time-scale actor-critic algorithms,
 085 employing linear function approximation for value functions and demonstrating asymptotic local
 086 convergence to a stationary point. Similarly, Bhatnagar et al. (2009) analyze four gradient-based
 087 methods, including natural policy gradients, and establish asymptotic local convergence using
 088 an ODE-based framework. Addressing global convergence, Bai et al. (2023) study policy gradient
 089 methods under the assumption that the average reward is smooth with respect to the policy parameters,
 090 achieving a regret bound of $O(T^{1/4})$. Ganesh et al. (2024) consider the average reward natural actor-
 091 critic algorithm which does not rely on knowledge of mixing time, and provide a regret bound of
 092 $O(\sqrt{T})$. However, the results in both of these works rely on unverified smoothness assumptions, leaving
 093 open questions about their applicability to general MDPs. Our contributions further strengthen the
 094 understanding of when these assumptions might indeed be true.

096 1.2 CONTRIBUTIONS

098 In this subsection, we outline the key contributions of this paper.

- 099 • **Elimination of Smoothness Assumption:** Unlike previous work which assumes the un-
 100 derlying average cost function is smooth, we prove its smoothness by introducing a new
 101 analysis technique. This technique addresses the key difficulty of the lack of uniqueness
 102 of the value function in average-reward problems. We overcome this challenge by using a
 103 projection technique to ensure uniqueness and leveraging the properties of the projection to
 104 prove smoothness. This removes a significant assumption and strengthens the theoretical
 105 foundations of policy gradient methods in average reward MDPs.
- 106 • **Expression for Smooth Average Cost:** We derive an explicit expression for the average
 107 cost that is shown to be smooth in the policy π . This contribution is critical as it provides

a deeper understanding and new insights into the behaviour of the average reward MDPs under policy gradient methods.

- **Sublinear Convergence Bounds:** Using the above smoothness property, we present finite time bounds on the optimality gap over time, showing that the iterates approach the optimal policy with an overall regret of $O(\log(T))$. In contrast, the regret bounds in Bai et al. (2023) are atleast $O(T^{\frac{1}{4}})$ without learning error and with tabular parametrization. In place of the discount factor and the cardinality of the state and action spaces in the discounted setting, our finite-time performance bounds involves a different parameter which characterizes the complexity of the underlying MDP.
- **Extension to Discounted Reward MDPs:** Our analysis can also be applied to the discounted reward MDP problem to provide stronger results than the state of the art. In particular, we show that our performance bounds for discounted MDPs can be expressed in terms of a problem complexity parameter, which can be independent of the size of the state and action spaces in some problems.
- **Experimental Validation:** We simulate the performance of policy gradient across a simple class of MDPs to empirically evaluate its performance. The simulations illustrate the impact of MDP complexity on convergence rates. Unlike previous results, where the bounds depend solely on the size of the state and action spaces, these simulations demonstrate how the underlying structure of the MDP can result in significantly different convergence rates, even with fixed state and action spaces. These observations further validate the theoretical bounds derived for the convergence of projected policy gradient in average reward MDPs.

2 PRELIMINARIES

In this section, we introduce our model, address the limitations of applying the optimality gap bounds from discounted reward MDPs to the average reward scenario, present the gradient ascent update, and discuss the assumptions underlying our analysis.

2.1 AVERAGE REWARD MDP FORMULATION

We consider the class of infinite horizon average reward MDPs with finite state space \mathcal{S} and finite action space \mathcal{A} . The environment is modeled as a probability transition kernel denoted by \mathbb{P} . We consider a class of randomized policies $\Pi = \{\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})\}$, where a policy π maps each state to a probability vector over the action space. The transition kernel corresponding to a policy π is represented by $\mathbb{P}^\pi : \mathcal{S} \rightarrow \mathcal{S}$, where $\mathbb{P}^\pi(s'|s) = \sum_{a \in \mathcal{A}} \pi(a|s) \mathbb{P}(s'|s, a)$ denotes the single step probability of moving from state s to s' under policy π . Let $r(s, a)$ denote the single step reward obtained by taking action $a \in \mathcal{A}$ in state $s \in \mathcal{S}$. The single-step reward associated with a policy π at state $s \in \mathcal{S}$ is defined as $r^\pi(s) = \sum_{a \in \mathcal{A}} \pi(a|s) r(s, a)$.

The infinite horizon average reward objective ρ^π associated with a policy π is defined as:

$$\rho^\pi = \lim_{N \rightarrow \infty} \frac{\mathbb{E}_\pi \left[\sum_{n=0}^{N-1} r^\pi(s_n) \right]}{N}, \quad (2)$$

where the expectation is taken with respect to \mathbb{P}^π . The average reward is independent of the initial state distribution under some mild conditions (Ross, 1983; Bertsekas, 2007) and can be alternatively expressed as $\rho^\pi = \sum_{s \in \mathcal{S}} d^\pi(s) r^\pi(s)$, where $d^\pi(s)$ is the stationary measure corresponding to state s under the transition kernel \mathbb{P}^π , ensuring that d^π satisfies the equation $d^\pi \mathbb{P}^\pi = d^\pi$. Associated with a policy is a relative state value function $v^\pi \in \mathbb{R}^{|\mathcal{S}|}$ that satisfies the following average reward Bellman equation

$$\rho^\pi \mathbb{1} + v^\pi = r^\pi + \mathbb{P}^\pi v^\pi, \quad (3)$$

where $\mathbb{1}$ is the all ones vector (Puterman, 1994; Bertsekas, 2007). Note that v^π is unique up to an additive constant. Setting $\sum_{s \in \mathcal{S}} d^\pi(s) v^\pi(s) = 0$ imposes an additional constraint over v^π , providing a unique value function vector denoted by v_0^π , known as the basic differential reward function (Tsitsiklis & Van Roy, 1999). It can be shown that v_0^π can alternatively expressed as

162 $v_0^\pi(s) = \mathbb{E}_\pi [\sum_{n=0}^{\infty} (r^\pi(s_n) - \rho^\pi) | s_0 = s]$. Hence any element in the set $\{v_0^\pi + c\mathbb{1} : c \in \mathbb{R}\}$ is a
 163 solution v^π to the Bellman equation equation 3.

164 The relative state action value function $Q^\pi \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ associated with a policy π is defined as:

$$165 \quad Q^\pi(s, a) = r(s, a) + \sum_{\substack{s' \in \mathcal{S} \\ a' \in \mathcal{A}}} \mathbb{P}(s'|s, a) \pi(a'|s') Q^\pi(s', a') - \rho^\pi \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (4)$$

166 Similar to v^π , Q^π is also unique up to an additive constant. Analogously, every solution Q^π of
 167 equation 4 can be expressed as an element in the set $\{Q_0^\pi(s, a) + c\mathbb{1} : c \in \mathbb{R}\}$ where $Q_0^\pi(s, a) =$
 168 $\mathbb{E}_\pi [\sum_{n=0}^{\infty} (r^\pi(s_n) - \rho^\pi) | s_0 = s, a_0 = a]$. Upon averaging equation 4 with policy π , it follows that
 169 $v^\pi(s) = \sum_{a \in \mathcal{A}} \pi(a|s) Q^\pi(s, a)$. The average reward policy gradient theorem (Sutton & Barto, 2018)
 170 for policies parameterized by θ is given by:

$$171 \quad \frac{\partial \rho}{\partial \theta} = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \frac{\partial \pi(s, a)}{\partial \theta} Q^\pi(s, a) \quad (5)$$

172 As we focus on tabular policies in this paper, our parameterization aligns with the tabular policy,
 173 where θ is equivalent to π . The policy gradient update considered is defined below.

$$174 \quad \pi_{k+1} := \mathbf{Proj}_\Pi \left[\pi_k + \eta \frac{\partial \rho^\pi}{\partial \pi} \Bigg|_{\pi = \pi_k} \right] \quad \forall k \geq 0, \quad (6)$$

175 where \mathbf{Proj}_Π denotes the orthogonal projection in the Euclidean norm onto the space of randomized
 176 policies Π and η denotes the step size of the update. In the following subsection, we recall the
 177 policy gradient result within the framework of discounted reward MDPs and address why it cannot
 178 be directly applied to the average reward scenario.

179 2.2 RELATIONSHIP TO DISCOUNTED REWARD MDPs

180 Let $\rho_{\mu, \gamma}^\pi := \mu^T (1 - \gamma \mathbb{P}^\pi)^{-1} r^\pi$ represent the discounted reward value function associated with policy
 181 π , under the initial distribution $\mu \in \Delta \mathcal{S}$ and where γ represents the discount factor (Bertsekas, 2007).
 182 Consider the projected policy gradient update given below.

$$183 \quad \pi_{k+1} := \mathbf{Proj}_\Pi \left[\pi_k + \eta \frac{\partial \rho_{\mu, \gamma}^\pi}{\partial \pi} \Bigg|_{\pi = \pi_k} \right] \quad \forall k \geq 0. \quad (7)$$

184 When the step size $\eta = \frac{(1-\gamma)^3}{2\gamma|\mathcal{A}|}$, the iterates π_k generated from projected gradient ascent equation 7
 185 satisfy the following equation:

$$186 \quad \rho_{\mu, \gamma}^* - \rho_{\mu, \gamma}^{\pi_k} \leq \frac{256|\mathcal{S}||\mathcal{A}|}{k(1-\gamma)^5} \left\| \frac{d_{\mu, \gamma}^{\pi_k}}{\mu} \right\|_\infty^2, \quad (8)$$

187 where $\rho_{\mu, \gamma}^*$ represents the optimal value function under initial distribution μ , and $d_{\mu, \gamma}^{\pi^*} := (1 -$
 188 $\gamma)\mu^T (1 - \gamma \mathbb{P}^{\pi^*})^{-1}$ represents the state occupancy measure under optimal policy π^* (Xiao, 2022b).
 189 Under some mild conditions the average reward ρ^π associated with a policy π and the value function
 190 $\rho_{\mu, \gamma}^\pi(s)$ are related as below:

$$191 \quad \rho^\pi = \lim_{\gamma \rightarrow 1} (1 - \gamma) \rho_{\mu, \gamma}^\pi(s). \quad (9)$$

192 Note that the above relation (Bertsekas, 2007; Ross, 1983) holds for all $s \in \mathcal{S}$ and all $\mu \in \Delta \mathcal{S}$ since
 193 the average reward is independent of the initial state distribution. Upon leveraging the relation in
 194 equation 9 and multiplying equation 8 with $(1 - \gamma)$, it is apparent that the upper bound of equation 8
 195 in the limit of $\gamma \rightarrow 1$ tends to infinity. This is due to $(1 - \gamma)^4$ that remains in the denominator
 196 of equation 8 upon multiplying with $(1 - \gamma)$. Therefore it is necessary to devise an alternative proof
 197 technique in order to analyze the global convergence of policy gradient in the context of average
 198 reward MDPs. Prior to presenting the main result and its proof, we state the assumption used in our
 199 analysis.

200 **Assumption 1.** For every policy $\pi \in \Pi$, the transition matrix \mathbb{P}^π associated with the induced Markov
 201 chain is irreducible and aperiodic. This assumption also means that there exist constants $C_e < \infty$
 202 and $\lambda \in [0, 1)$ such that for any $k \in \mathbb{N}$ and any $\pi \in \Pi$, the Markov chain corresponding to \mathbb{P}^π is
 203 geometrically ergodic i.e., $\|(\mathbb{P}^\pi)^k - \mathbb{1}(d^\pi)^\top\|_\infty \leq C_e \lambda^k$.

3 MAIN RESULTS

Theorem 1. Let ρ^{π_k} be the average reward corresponding to the policy iterates π_k , obtained through the policy gradient update equation 6. Let ρ^* represent the optimal average reward, that is, $\rho^* = \max_{\pi \in \Pi} \rho^\pi$. There exist constants L_2^Π and C_{PL} , determined by the underlying MDP, such that when the step size $\eta < \frac{1}{L_2^\Pi}$, the following holds:

- For all MDPs it is true that,

$$\rho^* - \rho^{\pi_k} \leq \frac{1}{\frac{1}{\rho^* - \rho^{\pi_0}} + \nu k}, \quad \forall k \geq 0. \quad (10)$$

$$\text{where } \nu := \left(\frac{1}{32C_{PL}^2|\mathcal{S}|L_2^\Pi} \right) \left(1 + 4 \left(\frac{1}{32C_{PL}^2|\mathcal{S}|L_2^\Pi} \right) \right)^{-\frac{3}{2}}$$

- For simple MDPs (i.e. $L_2^\Pi \ll 1$) we obtain exponential convergence, that is

$$\rho^* - \rho^{\pi_k} \leq c^{-\frac{k}{2}} \left(\rho^* - \rho^{\pi_0} \right)^{\frac{1}{2k}}, \quad \forall k \geq 0 \quad (11)$$

$$\text{where } \frac{1}{c} = 32|\mathcal{S}|L_2^\Pi C_{PL}^2 < 1.$$

Remark: It is worth noting that the above bounds correspond to a regret of $O(\log(T))$. While regret is typically a concept associated with online learning settings, we present it here to facilitate the development of learning algorithms inspired by the findings of this work. All previous results on convergence for discounted MDPs are of the form $\frac{\sigma}{k^p}$ (Agarwal et al., 2020; Mei et al., 2022; Xiao, 2022a), where σ is a large constant. However, since the worst sub-optimality is 1, the bound $\frac{\sigma}{k^p}$ becomes less meaningful for initial k . In contrast, our bound $\frac{1}{\frac{1}{\rho^* - \rho^{\pi_0}} + \nu k}$ is meaningful from the very first iteration. The maximum sub-optimality is $\rho^* - \rho^{\pi_0}$ at $k = 0$, and it decreases monotonically thereafter.

Our second result, an observation that is novel to this work, shows that simple MDPs exhibit much faster (linear) convergence rates. This explains why techniques like reward shaping, which simplify the MDP, can be highly effective.

3.1 KEY IDEAS AND PROOF OUTLINE

A similar result was proved for discounted reward MDPs in Agarwal et al. (2020); Xiao (2022a). An important property pivotal to the global convergence analysis of the projected policy gradient is the smoothness of the discounted reward value function. Demonstrating the smoothness of the discounted reward value function is relatively straightforward due to the contractive properties of the discount factor. However, this poses a significant challenge in the context of average reward MDPs. Here, the absence of a discount factor as a source of contraction, coupled with the lack of uniqueness in the average reward value function, complicates the task of proving the smoothness of the average reward. Therefore, the first important property we prove is the smoothness of average reward.

3.1.1 SMOOTHNESS OF AVERAGE REWARD

A differentiable function $f : \mathcal{C} \rightarrow \mathbb{R}$ is called L -smooth if it satisfies

$$\|\nabla f(y) - \nabla f(x)\|_2 \leq L\|y - x\|_2 \quad \forall y, x \in \mathcal{C}. \quad (12)$$

where \mathcal{C} is some subset of \mathbb{R}^n . Further if the function is L -smooth, it satisfies the following property.

$$\left| f(y) - f(x) - \langle \nabla f(x), y - x \rangle \right| \leq \frac{L}{2} \|y - x\|^2 \quad \forall y, x \in \mathcal{C}, \quad (13)$$

From the above definition, it is apparent that if f is L -smooth then cf is $|c|L$ -smooth for any $c \in \mathbb{R}$. It can be shown that the infinite horizon discounted reward $V_{\mu, \gamma}^\pi$ is $\frac{2\gamma|\mathcal{A}|}{(1-\gamma)^3}$ -smooth. Leveraging the result in equation 9, one can see that the smoothness constant of ρ^π is $\lim_{\gamma \rightarrow 1} \frac{2\gamma|\mathcal{A}|}{(1-\gamma)^2} \rightarrow \infty$. Hence,

the smoothness of the discounted reward cannot be leveraged to show the smoothness of the average reward.

In this paper, we establish the smoothness of the infinite horizon average reward by first establishing the smoothness of the associated relative value function. We then leverage the average reward Bellman Equation 3 to establish the smoothness of the average reward in terms of the smoothness of the relative value function. However, since the relative value function is unique up to an additive constant, we consider the projection of the value function onto the subspace orthogonal to the $\mathbb{1}$ vector. This provides us with a unique representation of the value function whose smoothness can be evaluated.

Let $\Phi \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ be the projection matrix that maps any vector to its orthogonal projection in the Euclidean norm onto the subspace perpendicular to the $\mathbb{1}$ vector. Then the following lemma holds.

Lemma 1. *Let I be the identity matrix and $\mathbb{1}$ be the all ones vector, both of dimension $|\mathcal{S}|$. Then the orthogonal projection matrix is given by $\Phi = \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{|\mathcal{S}|} \right)$. The unique value function v_ϕ^π is obtained as a solution to the following fixed point equation,*

$$v_\phi^\pi = \Phi \left(r^\pi + \mathbb{P}^\pi v_\phi^\pi - \rho^\pi \mathbb{1} \right) \quad (14)$$

and can be alternatively represented as

$$v_\phi^\pi = (I - \Phi \mathbb{P}^\pi)^{-1} \Phi r^\pi. \quad (15)$$

Since \mathbb{P}^π has 1 as its Perron Frobenius eigenvalue, $(I - \mathbb{P}^\pi)$ is a singular matrix. It can be verified that $\Phi \mathbb{1} = 0$, hence $\mathbb{1}$ is an eigenvector of $\Phi \mathbb{P}^\pi$ for all π with a corresponding eigenvalue of 0. It can subsequently be proven that the rest of the eigenvalues of $\Phi \mathbb{P}^\pi$ are all less than one in terms of their absolute value and hence equation 15 is well defined.

With a unique closed form for the average reward value function established, the subsequent task is to determine its smoothness constant. Given that the smoothness constant of a function f corresponds to the largest eigenvalue of its Hessian, we adopt an analytical approach similar to that presented in Agarwal et al. (2020). This involves utilizing directional derivatives and evaluating the maximum rate of change of derivatives across all directions within the policy space. It's important to note that since we are maximizing over directions expressible as differences between any two policies within the policy space, the resulting Lipschitz and smoothness constants are referred to as the restricted Lipschitz and smoothness constants, respectively. The restricted smoothness of the average reward value function is stated below.

Lemma 2. *For any policy $\pi \in \Pi$, there exist constants $C_m, C_p, C_r, \kappa_r \in \mathbb{R}^+$ which are determined by the underlying MDP, such that the value function v_ϕ^π is $4(2C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r)$ -smooth.*

Since the average reward value function is Lipschitz and smooth with respect to its policy, one can directly utilize this property to establish the Lipschitzness and smoothness of the average reward. These results are characterized in the following lemmas.

Lemma 3. *For any policy $\pi \in \Pi$, there exist constants $C_m, C_p, C_r, \kappa_r \in \mathbb{R}^+$ which are determined by the underlying MDP, such that the average reward ρ^π is L_1^Π -Lipschitz.*

$$\left| \left\langle \frac{\partial \rho^\pi}{\partial \pi}, \pi' - \pi \right\rangle \right| \leq L_1^\Pi \|\pi' - \pi\|_2, \quad \forall \pi, \pi' \in \Pi, \quad (16)$$

where $L_1^\Pi = 2(C_r + C_p C_m \kappa_r + 2(C_m^2 C_p \kappa_r + C_m C_r))$

The restricted Lipschitzness of the average reward is utilized to prove its restricted smoothness.

Lemma 4. *For any policy $\pi \in \Pi$, there exist constants $C_m, C_p, C_r, \kappa_r \in \mathbb{R}^+$ which are determined by the underlying MDP, such that the average reward ρ^π is L_2^Π -smooth.*

$$\left| \left\langle \pi' - \pi, \frac{\partial^2 \rho^\pi}{\partial \pi^2} (\pi' - \pi) \right\rangle \right| \leq \frac{L_2^\Pi}{2} \|\pi' - \pi\|_2^2 \quad \forall \pi, \pi' \in \Pi, \quad (17)$$

where $L_2^\Pi = 4(C_p^2 C_m^2 \kappa_r + C_p C_m C_r + (C_p + 1)(C_m^2 C_p \kappa_r + C_m C_r) + 4(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r))$.

Table 1: Constants capturing the MDP Complexity

	Definition	Range	Remark
C_m	$\max_{\pi \in \Pi} \ (I - \Phi P^\pi)^{-1}\ $	$\frac{2C_e \mathcal{S} }{1-\lambda}$ See Assumption 1 for definition of C_e and λ	Lowest rate of mixing
C_p	$\max_{\pi, \pi' \in \Pi} \frac{\ \mathbb{P}^{\pi'} - \mathbb{P}^\pi\ }{\ \pi' - \pi\ _2}$	$[0, \sqrt{ \mathcal{A} }]$	Diameter of transition kernel
C_r	$\max_{\pi, \pi'} \frac{\ r^{\pi'} - r^\pi\ _\infty}{\ \pi' - \pi\ _2}$	$[0, \sqrt{ \mathcal{A} }]$	Diameter of reward function
κ_r	$\max_{\pi} \ \Phi r^\pi\ _\infty$	$[0, 2)$	Variance of reward function
$\frac{L_1^\Pi}{2}$	$C_r + C_p C_m \kappa_r + 2(C_m^2 C_p \kappa_r + C_m C_r)$	$[0, k_1 \sqrt{ \mathcal{A} } C_m^2]$	Restricted Lipschitz constant
$\frac{L_2^\Pi}{4}$	$C_p^2 C_m^2 \kappa_r + C_p C_m C_r + (C_p + 1)(C_m^2 C_p \kappa_r + C_m C_r) + 4(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r)$	$[0, k_2 \mathcal{A} C_m^3]$	Restricted smoothness constant

C_p, C_m are defined using operator norm w.r.t. L_∞ norm. Precisely,

$$C_m = \max_{\pi} \max_{\|v\|_\infty \leq 1} \|(I - \Phi P^\pi)^{-1} v\|_\infty, \text{ and } C_p = \max_{\pi, \pi' \in \Pi} \max_{\|v\|_\infty \leq 1} \frac{\|(P^{\pi'} - P^\pi)v\|_\infty}{\|\pi' - \pi\|_2}$$

Note that the restricted Lipschitz constant of the average reward is upper bounded by its general Lipschitz constant:

$$\max_{\pi' \in \Pi: \|\pi' - \pi\|_2 \leq 1} \left| \left\langle \frac{\partial \rho^\pi}{\partial \pi}, \pi' - \pi \right\rangle \right| \leq \max_{u \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}: \|u\|_2 \leq 1} \left| \left\langle \frac{\partial \rho^\pi}{\partial \pi}, u \right\rangle \right| \quad (18)$$

By confining our analysis of the smoothness constants to the policy class, we introduce a dependency of our convergence bounds on MDP-specific constants, including C_r, C_p, C_m, C_e and κ_r . These constants capture the complexity of the underlying MDP and are exclusive to the analysis presented in this paper, as there appears to be no such dependency observed in the global convergence bounds of Agarwal et al. (2020). A more detailed description of these constants can be found in Table 1, where k_1 and k_2 represent MDP-independent numeric constants. These constants, which rely on the characteristics of the MDP, suggest that the projected policy gradient may achieve faster convergence in MDPs with lower complexity as opposed to those with higher complexity. The range of these constants can be found in Appendix A. We now proceed to analyze the convergence of projected policy gradient utilizing the smoothness of the average reward.

3.1.2 CONVERGENCE OF POLICY GRADIENT

Using the smoothness property of the average reward, it is possible to show that the improvement in the successive average reward iterates is bounded from below by the product of the smoothness constant and the difference in the policy iterates, as described in the lemma below.

Lemma 5. *Let ρ^{π_k} be the average reward corresponding to the policy iterate π_k obtained from equation 6. Let L_2^Π be as in Lemma 4. Then,*

$$\rho^{\pi_{k+1}} - \rho^{\pi_k} \geq \frac{L_2^\Pi}{2} \|\pi_{k+1} - \pi_k\|^2, \quad \forall k \in \mathbb{N}. \quad (19)$$

Successively increasing iterates are not sufficient to guarantee finite time global convergence bounds. It is therefore necessary to bound the suboptimality associated with each iterate. We do so by leveraging the performance difference lemma stated below.

Lemma 6. Let ρ^* be the globally optimal average reward. Then for any $\pi \in \Pi$, the suboptimality of ρ^π can be expressed as:

$$\rho^* - \rho^\pi = \sum_s d^{\pi^*}(s) \sum_a Q^\pi(s, a) [\pi^*(a|s) - \pi(a|s)]. \quad (20)$$

Proof. The proof can be found in Cao (1999). \square

In the next lemma, we upper bound the right-hand side of equation 20 in terms of the gradient of ρ^π .

Lemma 7. The suboptimality of any $\pi \in \Pi$ satisfies:

$$\rho^* - \rho^\pi \leq C_{PL} \max_{\pi' \in \Pi} \left\langle \pi' - \pi, \frac{\partial \rho^\pi}{\partial \pi} \right\rangle, \quad \forall \pi \in \Pi, \quad (21)$$

where $C_{PL} = \max_{\pi \in \Pi} \max_{s \in \mathcal{S}} \frac{d^{\pi^*}(s)}{d^\pi(s)}$.

Note that C_{PL} is a constant that is proportional to the size of the state space. We do not know if the appearance of such a constant is inevitable or not; however, it should be noted such a constant appears in prior works on discounted reward problems as well (Agarwal et al., 2020; Xiao, 2022a).

It is possible to further upper bound the expression in Lemma 7 using the smoothness property of the average reward.

Lemma 8. Let π_k be the policy iterates generated by equation 6. Then for all $\pi' \in \Pi$ it is true that,

$$\left\langle \frac{\partial \rho_{\pi_{k+1}}}{\partial \pi_{k+1}}, \pi' - \pi_{k+1} \right\rangle \leq 4\sqrt{|\mathcal{S}|} L_2^\Pi \|\pi_{k+1} - \pi_k\|, \quad (22)$$

Lemmas 5,7 and 8 are combined to prove the result in Theorem 1.

3.2 EXTENSION TO DISCOUNTED REWARD MDPs

Existing performance bounds in the context of discounted reward MDPs require an iteration complexity of $O\left(\frac{|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^5\epsilon}\right)$ to achieve policies with suboptimality of ϵ (Xiao, 2022a). These bounds are independent of the hardness of the underlying MDP. Our approach improves on this bound, yielding an $O\left(\frac{|\mathcal{S}|L_2^\Pi}{\epsilon}\right)$ iteration complexity, where $L_2^\Pi = C_p^2 \widehat{C}_m^2 \kappa_r + C_p \widehat{C}_m C_r + (C_p + 1)(\widehat{C}_m^2 C_p \kappa_r + \widehat{C}_m C_r) + 4(\widehat{C}_m^3 C_p^2 \kappa_r + \widehat{C}_m^2 C_p C_r)$, where $\widehat{C}_m := \|(I - \gamma \mathbb{P}^\pi)^{-1}\|$. It is straightforward to see that $\widehat{C}_m \leq \frac{1}{1-\gamma}$. Hence, the iteration complexity improves to $O\left(\frac{L_2^\Pi |\mathcal{S}|}{(1-\gamma)^5\epsilon}\right)$, as the constants satisfy $\kappa_r \leq 2$ and $C_p, C_r \leq \sqrt{|\mathcal{A}|}$. Further, the approach considered in this paper provides faster convergence rates for MDPs with low complexity, i.e., MDPs that have low values of C_p or C_r . The exact performance bounds can be obtained from an approach similar to the one outlined in Kumar et al. (2023), where L_2^Π represents the restricted smoothness constant of the discounted return ρ_γ^π . This constant can be derived through a process analogous to the one described in this paper.

For instance, consider a trivial MDP for which $C_p = 0$ or $\kappa_r = 0$ (implies $C_r = 0$), i.e., an MDP where the transition kernel is independent of the action enacted. For this trivial MDP every policy is an optimal policy. The state of the art convergence guarantees (Xiao, 2022a), still requires $O(|\mathcal{S}||\mathcal{A}|\epsilon^{-1})$ iterations for ϵ close optimal policy. Whereas, the performance bounds presented in this paper predict $O(|\mathcal{S}|\epsilon^{-1})$ iterations for convergence. Thus the constant L_2^Π captures the hardness of the MDP. Therefore, MDPs with lower complexity, i.e., lower values of L_2^Π , converge faster than MDPs with higher complexity, thus improving on current complexity-independent bounds.

4 SIMULATIONS

Here, we present simulations corresponding to two MDP complexity measures. In order to study the convergence of projected policy gradient in the context of average reward MDPs and its dependence on the underlying MDP complexity, we present simulation results corresponding to two complexity measures: the cardinality of state and action spaces, and the diameter of the reward function.

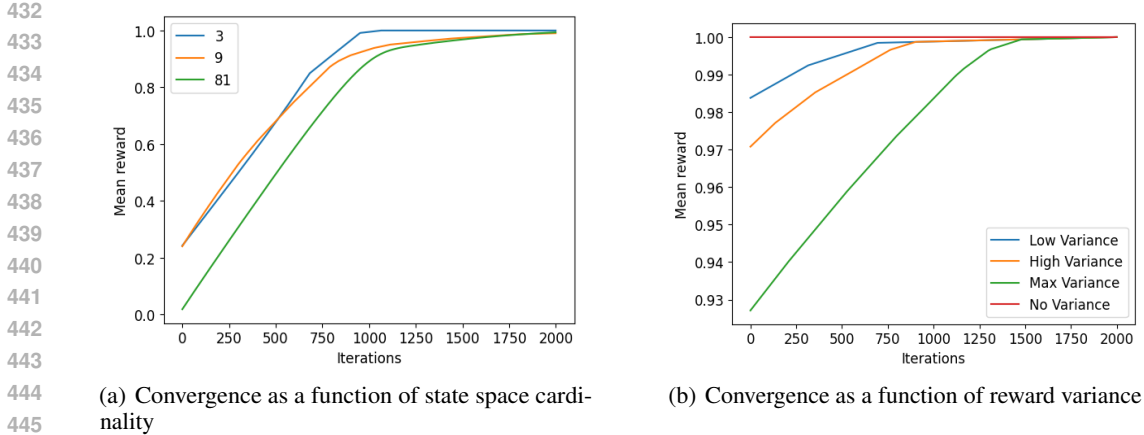



Figure 1: Improvement in average reward as a function of MDP complexity

Figure 1(a) considers MDPs with $(|\mathcal{S}|, |\mathcal{A}|) = \{(3, 3), (9, 9), (81, 81)\}$. We construct the transition kernel and the reward function in the same manner for all MDPs, which we discuss in the appendix. Projected policy gradient was implemented for 2000 iterations and the overall average reward is plotted as a function of iteration number. As expected, the convergence rate is slower when $(|\mathcal{S}|, |\mathcal{A}|)$ are larger due to the fact that the reward smoothness constant is larger. This reduction stems from small values of C_M, C_r, C_p , which are characteristic of MDPs with smaller state and action space cardinalities when the transition kernel and reward structures are similar. A less obvious result is that, even for MDPs with a fixed cardinality of state and action spaces, the rate of convergence can be considerably different as shown in Figure 1(b). For this simulation, we fix the state and action space cardinality at $(|\mathcal{S}|, |\mathcal{A}|) = (16, 16)$. We randomly generate a transition kernel, which remains constant across different single-step reward functions corresponding to varying reward variances. In particular, we consider four different reward variances - no variance, low, high and maximal variance. We recall the definition of C_r found in Table 1. We see that C_r scales with reward variance. Specifically, C_r is large when small changes to the policy result in significant modifications to the mean reward. Therefore, we anticipate that higher reward variance will lead to slower convergence. Additional details are found in the appendix. The observed convergence trend aligns with the theoretical bounds obtained, indicating that MDPs with small values of C_r tend to converge relatively faster.

Next, we discuss the impact of C_p on convergence of policy gradient. We consider MDPs of size 16, i.e., $(|\mathcal{S}|, |\mathcal{A}|) = (16, 16)$. We generate three different transition kernels. The first is uniform, so the actions do not change the transition probabilities, the second is deterministic (i.e., there exists some $s' \in \mathcal{S}$ such that $\mathbb{P}(s'|s, a) = 1$ for all $(s, a) \in (\mathcal{S}, \mathcal{A})$), and the last is non-uniform but stochastic transition kernel. We recall the definition of C_p from Table 1. We see that C_p is larger when the transition probabilities change by a greater amount with small changes to the policy. Thus, deterministic MDPs should have higher C_p values than ones that are more stochastic. Additional details are provided in the appendix. We run the policy gradient algorithm considered in this paper for each MDP setting and plot the overall change in average reward as a function of iterations, for 3000 iterations. Figure 2 indicates that policy gradient in MDPs corresponding to small values of C_p converges relatively faster than for MDPs corresponding to large values of C_p . Hence, the performance bounds obtained in Theorem 1 are in some sense, more representative of the empirical convergence trend of the policy gradient algorithm.

Note on Limitations and Future Work: Although we study tabular policies, this approach can be generalized to parametric class of policies. This study uses the exact value of the gradient and does not account for learning errors in the analysis. Nonetheless, we highlight that this is the first comprehensive proof of global convergence for policy gradient methods in average reward MDPs. Future work will focus on incorporating learning errors into this framework.

Figure 2: Convergence as a function of C_p

5 CONCLUSION

In this paper, we presented the first comprehensive finite-time global convergence analysis of policy gradient for infinite horizon average reward MDPs. Key contributions include eliminating the smoothness assumption from previous work, deriving an explicit expression for the smooth average cost, and proving sublinear convergence with a regret of $O(\log(T))$. Our findings offer a more general and robust understanding of policy gradient methods in average reward MDPs, addressing long-standing challenges such as the lack of uniqueness in value functions. We also extended our analysis to discounted reward MDPs, providing stronger performance bounds by incorporating a complexity parameter beyond state and action space sizes. The theoretical results were further supported by simulations that highlighted how the structure of the underlying MDP influences convergence rates. These insights open avenues for refining performance bounds and exploring real-world applications in both average and discounted reward MDPs.

REFERENCES

- Yasin Abbasi-Yadkori, Peter Bartlett, Kush Bhatia, Nevena Lazic, Csaba Szepesvari, and Gellért Weisz. Politex: Regret bounds for policy iteration using expert prediction. In *International Conference on Machine Learning*, pp. 3692–3702. PMLR, 2019.
- Alekh Agarwal, Sham M. Kakade, Jason D. Lee, and Gaurav Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift, 2020.
- Qinbo Bai, Washim Uddin Mondal, and Vaneet Aggarwal. Regret analysis of policy gradient algorithm for infinite horizon average reward markov decision processes. *arXiv preprint arXiv:2309.01922*, 2023.
- Jonathan Baxter and Peter L Bartlett. Direct gradient-based reinforcement learning. In *2000 IEEE International Symposium on Circuits and Systems (ISCAS)*, volume 3, pp. 271–274. IEEE, 2000.
- Amir Beck. *Introduction to nonlinear optimization: Theory, algorithms, and applications with MATLAB*. SIAM, 2014.
- Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control, Vol. II*. Athena Scientific, 3rd edition, 2007. ISBN 1886529302.
- Jalaj Bhandari and Daniel Russo. Global optimality guarantees for policy gradient methods. *Operations Research*, 2024.
- Shalabh Bhatnagar, Richard S Sutton, Mohammad Ghavamzadeh, and Mark Lee. Natural actor-critic algorithms. *Automatica*, 45(11):2471–2482, 2009.
- Tomasz Bielecki, Daniel Hernandez-Hernandez, and Stanley R Pliska. Value iteration for controlled markov chains with risk sensitive cost criterion. In *Proceedings of the 38th IEEE Conference on Decision and Control (Cat. No. 99CH36304)*, volume 1, pp. 126–130. IEEE, 1999.

- 540 Xi-Ren Cao. Single sample path-based optimization of markov chains. *Journal of Optimization*
541 *Theory and Applications*, 100:527–548, 1999. doi: 10.1023/A:1022634422482.
- 542
- 543 Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Online markov decision processes. *Mathe-*
544 *matics of Operations Research*, 34(3):726–736, 2009.
- 545 Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient
546 methods for the linear quadratic regulator. In *International conference on machine learning*, pp.
547 1467–1476. PMLR, 2018.
- 548
- 549 Swetha Ganesh, Washim Uddin Mondal, and Vaneet Aggarwal. Order-optimal global convergence for
550 average reward reinforcement learning via actor-critic approach. *arXiv preprint arXiv:2407.18878*,
551 2024.
- 552 Ganesh Ghalme, Vineet Nair, Vishakha Patil, and Yilun Zhou. Long-term resource allocation fairness
553 in average markov decision process (amd) environment. *arXiv preprint arXiv:2102.07120*, 2021.
- 554
- 555 Abhijit Gosavi. A reinforcement learning algorithm based on policy iteration for average reward:
556 Empirical results with yield management and convergence analysis. *Machine Learning*, 55:5–29,
557 2004.
- 558
- 559 Julien Grand-Clément and Marek Petrik. Reducing blackwell and average optimality to discounted
560 mdps via the blackwell discount factor. *Advances in Neural Information Processing Systems*, 36,
561 2024.
- 562 Isaac Grosf, Siva Theja Maguluri, and R Srikant. Convergence for natural policy gradient on
563 infinite-state average-reward markov decision processes. *arXiv preprint arXiv:2402.05274*, 2024.
- 564
- 565 Emmeran Johnson, Ciara Pike-Burke, and Patrick Rebeschini. Optimal convergence rate for exact
566 policy mirror descent in discounted markov decision processes, 2023. URL <https://arxiv.org/abs/2302.11381>.
- 567
- 568 Sajad Khodadadian, Prakirt Raj Jhunjunwala, Sushil Mahavir Varma, and Siva Theja Maguluri. On
569 the linear convergence of natural policy gradient algorithm. In *2021 60th IEEE Conference on*
570 *Decision and Control (CDC)*, pp. 3794–3799. IEEE, 2021.
- 571
- 572 Vijay R. Konda and John N. Tsitsiklis. Actor-critic algorithms. In *Neural Information Processing*
573 *Systems*, 1999. URL <https://api.semanticscholar.org/CorpusID:207779694>.
- 574
- 575 Navdeep Kumar, Inura Usmanova, Kfir Yehuda Levy, and Shie Mannor. Towards faster global
576 convergence of robust policy gradient methods. In *Sixteenth European Workshop on Reinforcement*
Learning, 2023. URL <https://openreview.net/forum?id=cWrwdbEBx5>.
- 577
- 578 Navdeep Kumar, Priyank Agrawal, Giorgia Ramponi, Kfir Yehuda Levy, and Shie Mannor. Improved
579 sample complexity for global convergence of actor-critic algorithms, 2024. URL <https://arxiv.org/abs/2410.08868>.
- 580
- 581 Jiakai Liu, Wenye Li, and Ke Wei. Elementary analysis of policy gradient methods, 2024. URL
582 <https://arxiv.org/abs/2404.03372>.
- 583
- 584 Sridhar Mahadevan. Average reward reinforcement learning: Foundations, algorithms, and empirical
585 results. *Machine learning*, 22(1):159–195, 1996.
- 586
- 587 Jincheng Mei, Chenjun Xiao, Csaba Szepesvari, and Dale Schuurmans. On the global convergence
588 rates of softmax policy gradient methods. In *International conference on machine learning*, pp.
6820–6829. PMLR, 2020.
- 589
- 590 Jincheng Mei, Chenjun Xiao, Csaba Szepesvari, and Dale Schuurmans. On the global convergence
591 rates of softmax policy gradient methods, 2022.
- 592
- 593 Yashaswini Murthy and R Srikant. On the convergence of natural policy gradient and mirror descent-
like policy methods for average-reward mdps. In *2023 62nd IEEE Conference on Decision and*
Control (CDC), pp. 1979–1984. IEEE, 2023.

- 594 Yashaswini Murthy, Mehrdad Moharrami, and R Srikant. Performance bounds for policy-based
595 average reward reinforcement learning algorithms. *Advances in Neural Information Processing*
596 *Systems*, 36, 2024.
- 597
- 598 Jonathan Patrick and Mehmet A Begen. Markov decision processes and its applications in healthcare.
599 *Handbook of healthcare delivery systems*. CRC, Boca Raton, 2011.
- 600
- 601 Martin L. Puterman. Markov decision processes: Discrete stochastic dynamic programming. In *Wiley*
602 *Series in Probability and Statistics*, 1994.
- 603
- 604 Sheldon M. Ross. *Introduction to Stochastic Dynamic Programming: Probability and Mathematical*.
605 Academic Press, Inc., USA, 1983. ISBN 0125984200.
- 606
- 607 John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region
608 policy optimization. In *International conference on machine learning*, pp. 1889–1897. PMLR,
609 2015.
- 610
- 611 Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press,
612 second edition, 2018. URL <http://incompleteideas.net/book/the-book-2nd.html>.
- 613
- 614 Prasad Tadepalli and DoKyeong Ok. Model-based average reward reinforcement learning. *Artificial*
615 *intelligence*, 100(1-2):177–224, 1998.
- 616
- 617 John N Tsitsiklis and Benjamin Van Roy. Average cost temporal-difference learning. *Automatica*, 35
618 (11):1799–1808, 1999.
- 619
- 620 Lin Xiao. On the convergence rates of policy gradient methods, 2022a. URL <https://arxiv.org/abs/2201.07443>.
- 621
- 622
- 623 Lin Xiao. On the convergence rates of policy gradient methods. *Journal of Machine Learning*
624 *Research*, 23(282):1–36, 2022b.
- 625
- 626 Junyu Zhang, Chengzhuo Ni, Csaba Szepesvari, Mengdi Wang, et al. On the convergence and sample
627 efficiency of variance-reduced policy gradient method. *Advances in Neural Information Processing*
628 *Systems*, 34:2228–2240, 2021.
- 629

631 A SMOOTHNESS OF AVERAGE REWARD

633 A.1 PROOF OF LEMMA 1

634 Consider the subspace orthogonal E to the all ones vector $\mathbb{1} \in \mathbb{R}^{|\mathcal{S}|}$ defined below:

$$636 E = \text{span} \left\{ \theta \in \mathbb{R}^{|\mathcal{S}|} : \theta^\top \mathbb{1} = 0. \right\} \quad (23)$$

637 The orthogonal projection v_ϕ of a vector v in the Euclidean norm onto the subspace E is defined as:

$$639 v_\phi = \arg \min_{u \in E} \|v - u\|_2 \quad (24)$$

640 It can be checked that the closed form expression for v_ϕ is given by:

$$641 v_\phi = \left(I - \frac{\mathbb{1} \mathbb{1}^\top}{|\mathcal{S}|} \right) v \quad (25)$$

642 where $I \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ is the identity matrix.

648 Consider the projection of the vector $r^\pi + \mathbb{P}^\pi v^\pi - \rho^\pi \mathbb{1}$ onto E for any policy $\pi \in \Pi$. The above
 649 projection is identical to the projection of $r^\pi + \mathbb{P}^\pi v^\pi$ onto E , since $\rho^\pi \mathbb{1}$ lies in the nullspace of Φ .
 650

$$651 \Phi(r^\pi - \rho^\pi \mathbb{1} + \mathbb{P}^\pi v) = (r^\pi - \rho^\pi \mathbb{1} + \mathbb{P}^\pi v) - \langle \mathbb{1}, r^\pi - \rho^\pi \mathbb{1} + \mathbb{P}^\pi v \rangle \frac{\mathbb{1}}{|\mathcal{S}|} \quad (26)$$

$$652 = (r^\pi + \mathbb{P}^\pi v) - \langle \mathbb{1}, r^\pi + \mathbb{P}^\pi v \rangle \frac{\mathbb{1}}{|\mathcal{S}|} \quad (27)$$

$$653 = r^\pi - \langle r^\pi, \mathbb{1} \rangle \frac{\mathbb{1}}{|\mathcal{S}|} + \mathbb{P}^\pi v - \langle \mathbb{1}, \mathbb{P}^\pi v \rangle \frac{\mathbb{1}}{|\mathcal{S}|} \quad (28)$$

$$654 = r^\pi - \langle r^\pi, \mathbb{1} \rangle \frac{\mathbb{1}}{|\mathcal{S}|} + \mathbb{P}^\pi v - \frac{\mathbb{1}}{|\mathcal{S}|} (\mathbb{1}^\top \mathbb{P}^\pi v) \quad (29)$$

$$655 = (I - \frac{\mathbb{1} \mathbb{1}^\top}{|\mathcal{S}|}) r^\pi + (I - \frac{\mathbb{1} \mathbb{1}^\top}{|\mathcal{S}|}) \mathbb{P}^\pi v \quad (30)$$

$$656 = \Phi [r^\pi + \mathbb{P}^\pi v]. \quad (31)$$

657 Consider the average reward Bellman equation corresponding to policy $\pi \in \Pi$:

$$658 \rho^\pi \mathbb{1} + v^\pi = r^\pi + \mathbb{P}^\pi v^\pi \quad (32)$$

659 Imposing an additional constraint $v^{\pi^\top} \mathbb{1} = 0$ yields a unique average reward value function denoted
 660 by v_ϕ^π . Moreover, it is true that,
 661

$$662 \Phi v_\phi^\pi + \Phi \rho^\pi \mathbb{1} = \Phi r^\pi + \Phi \mathbb{P}^\pi v_\phi^\pi \quad (33)$$

$$663 \implies \Phi v_\phi^\pi = \Phi r^\pi + \Phi \mathbb{P}^\pi v_\phi^\pi, \quad (34)$$

$$664 \stackrel{(a)}{\implies} v_\phi^\pi = \Phi [r^\pi + \mathbb{P}^\pi v_\phi^\pi], \quad (35)$$

665 where (a) is true because $v^{\pi^\top} \mathbb{1} = 0 \implies \Phi v^\pi = v^\pi$. Thus the projected value function with an
 666 unique representation is given by:
 667

$$668 v_\phi^\pi = [I - \Phi \mathbb{P}^\pi]^{-1} \Phi r^\pi, \quad (36)$$

669 and the existence of the inverse is proven in Subsection A.2, Lemma 12. An alternate expression for
 670 the projected value function is given by: $v_\phi^\pi = \left(I + \mathbb{1} \mathbb{1}^\top D - \frac{\mathbb{1} \mathbb{1}^\top}{|\mathcal{S}|} \right) v_0^\pi$, where $D \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ is a
 671 diagonal matrix whose entries correspond to the stationary measure over the states associated with
 672 policy π . See Tsitsiklis & Van Roy (1999) for more details.
 673

674 A.2 PROOF THAT EIGENVALUES OF $(I - \Phi \mathbb{P}^\pi)$ ARE NON-ZERO

675 In this subsection, we introduce the lemmas required to establish the proof of the eigenvalues of
 676 $\left(I - \frac{\mathbb{1} \mathbb{1}^\top}{|\mathcal{S}|} \right) \mathbb{P}^\pi$ being nonzero. We use the following notation: $\mathbb{1} \in \mathbb{R}^n$ represents the all ones vector
 677 and $I \in \mathbb{R}^{n \times n}$ is the identity matrix.
 678

679 **Lemma 9.** *Let $A \in \mathbb{R}^{n \times n}$ be a stochastic matrix. It is true that*

$$680 \left(\left(I - \frac{\mathbb{1} \mathbb{1}^\top}{n} \right) A \right)^k = \left(I - \frac{\mathbb{1} \mathbb{1}^\top}{n} \right) A^k. \quad (37)$$

702 *Proof.* For any $k \in \mathbb{N}$, consider,

$$704 \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A^k \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A = \left(A^k - \frac{\mathbb{1}\mathbb{1}^\top}{n} A^k\right) \left(A - \frac{\mathbb{1}\mathbb{1}^\top}{n} A\right) \quad (38)$$

$$706 = A^{k+1} - \frac{\mathbb{1}\mathbb{1}^\top}{n} A^{k+1} - A^k \frac{\mathbb{1}\mathbb{1}^\top}{n} A + \frac{\mathbb{1}\mathbb{1}^\top}{n} A^k \frac{\mathbb{1}\mathbb{1}^\top}{n} A \quad (39)$$

$$708 \stackrel{(a)}{=} A^{k+1} - \frac{\mathbb{1}\mathbb{1}^\top}{n} A^{k+1} - \frac{\mathbb{1}\mathbb{1}^\top}{n} A + \frac{\mathbb{1}\mathbb{1}^\top}{n} \frac{\mathbb{1}\mathbb{1}^\top}{n} A, \quad (40)$$

$$710 \stackrel{(b)}{=} A^{k+1} - \frac{\mathbb{1}\mathbb{1}^\top}{n} A^{k+1} - \frac{\mathbb{1}\mathbb{1}^\top}{n} A + \frac{\mathbb{1}\mathbb{1}^\top}{n} A, \quad (41)$$

$$712 = A^{k+1} - \frac{\mathbb{1}\mathbb{1}^\top}{n} A^{k+1}. \quad (42)$$

$$714 = \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A^{k+1} \quad (43)$$

718 where (a) is true because $A^k \mathbb{1} = \mathbb{1}$ and (b) follows from the fact that $\frac{\mathbb{1}^\top \mathbb{1}}{n} = \mathbb{1}$. From mathematical
719 induction it thus follows that,

$$721 \left(\left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A\right)^k = \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A^k \quad \forall k \in \mathbb{N}. \quad (44)$$

724 \square

725 **Lemma 10.** For any irreducible and aperiodic stochastic matrix $A \in \mathbb{R}^{n \times n}$, it is true that

$$727 \lim_{k \rightarrow \infty} \left(\left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A\right)^k = 0 \quad (45)$$

730 *Proof.* From Lemma 9 we have,

$$732 \left(\left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A\right)^k = \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A^k \quad \forall k \in \mathbb{N} \quad (46)$$

735 Since A is irreducible and aperiodic, the following limit converges to the stationary distribution
736 $d \in \mathbb{R}_+^n$ associated with A .

$$738 \lim_{k \rightarrow \infty} A^k = \mathbb{1}d^\top \quad (47)$$

739 Consider the following,

$$741 \lim_{k \rightarrow \infty} \left(\left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A\right)^k = \lim_{k \rightarrow \infty} \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) A^k \quad (48)$$

$$743 = \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) \lim_{k \rightarrow \infty} A^k \quad (49)$$

$$744 = \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n}\right) \mathbb{1}d^\top \quad (\text{from Equation equation 47}), \quad (50)$$

$$746 = \mathbb{1}d^\top - \frac{\mathbb{1}\mathbb{1}^\top}{n} \mathbb{1}d^\top \quad (51)$$

$$748 \stackrel{(a)}{=} \mathbb{1}d^\top - \mathbb{1}d^\top \quad (52)$$

$$749 = 0. \quad (53)$$

753 where (a) is true because $\frac{\mathbb{1}^\top \mathbb{1}}{n} = 1$. \square

754 **Lemma 11.** Let $A \in \mathbb{R}^{n \times n}$ be a matrix such that $\lim_{k \rightarrow \infty} A^k = 0$. Then $(I - A)^{-1} = \sum_{k=0}^{\infty} A^k$.

756 *Proof.* For any $K \in \mathbb{N}$, consider the following,

$$757 \quad (I - A) \left(\sum_{k=0}^K A^k \right) = I - A^{K+1}, \quad (54)$$

$$759 \quad \implies (I - A) \left(\lim_{K \rightarrow \infty} \sum_{k=0}^K A^k \right) = \lim_{K \rightarrow \infty} (I - A^{K+1}) \stackrel{(a)}{=} I, \quad (55)$$

761 where (a) follows from the fact that $\lim_{k \rightarrow \infty} A^k = 0$. Hence the inverse of $(I - A)$ can be expressed
762 as $(I - A)^{-1} = \sum_{k=0}^{\infty} A^k$. \square

766 **Lemma 12.** Let $A \in \mathbb{R}^{n \times n}$ be an irreducible and aperiodic stochastic matrix. Then the matrix
767 $\left(I - \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A \right)$ is invertible and its inverse is given by:

$$769 \quad \left(I - \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A \right)^{-1} = \sum_{k=0}^{\infty} \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A^k \quad (56)$$

772 *Proof.* Let λ_i be eigenvalues of $\left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A$. Then λ_i^k represents the eigenvalues of
773 $\left(\left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A \right)^k$. But from Lemma 10, we know that

$$774 \quad \lim_{k \rightarrow \infty} \left(\left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A \right)^k = 0 \quad (57)$$

776 Since eigenvalues are continuous functions of their corresponding matrices and all eigenvalues of a
777 zero matrix are zero, we thus have,

$$779 \quad \lim_{k \rightarrow \infty} \lambda_i^k = 0 \quad \forall i \in \{1, \dots, n\} \quad (58)$$

781 Equation 58 thus implies that $|\lambda_i| < 1, \forall i \in \{1, \dots, n\}$. Hence the matrix $\left(I - \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A \right)$
782 has all non zero eigenvalues and is thus invertible. From Lemma 11, we know that

$$783 \quad (I - A)^{-1} = \sum_{k=0}^{\infty} A^k \quad (59)$$

784 when $\lim_{k \rightarrow \infty} A^k = 0$. Since, $\lim_{k \rightarrow \infty} \left(\left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A \right)^k = 0$ from Lemma 10, we have the
785 following result,

$$786 \quad \left(I - \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A \right)^{-1} = \sum_{k=0}^{\infty} \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right) A^k \quad (60)$$

794 \square

795 From definition we have $\Phi = \left(I - \frac{\mathbb{1}\mathbb{1}^\top}{n} \right)$. Hence the inverse $(1 - \Phi \mathbb{P}^\pi)^{-1} = \sum_{k=0}^{\infty} \Phi (\mathbb{P}^\pi)^k$ exists
796 and is well defined for all $\pi \in \Pi$.

798 A.3 SMOOTHNESS OF THE AVERAGE REWARD VALUE FUNCTION v_ϕ^π

800 In order to prove the smoothness of the average reward value function and the infinite horizon average
801 reward, we consider an analysis inspired by Agarwal et al. (2020), where instead of computing the
802 maximum eigenvalue of the associated Hessian matrices, we consider the maximum value of the
803 directional derivative across all directions within the policy class.

804 Let $\pi, \pi' \in \Pi$ be any policies within the policy class. Then define π_α as a convex combination of
805 policies π and π' . That is

$$807 \quad \pi_\alpha := (1 - \alpha)\pi + \alpha\pi' \quad (61)$$

$$808 \quad = \pi + \alpha(\pi' - \pi) \quad (62)$$

$$809 \quad = \pi + \alpha u \quad (63)$$

810 where $u = \pi' - \pi$.

811 Since π_α is linear in α , it is true that

$$812 \quad \nabla_\alpha \pi_\alpha = \frac{d(\pi + \alpha u)}{d\alpha} = u, \quad \text{and} \quad \nabla_\alpha^2 \pi_\alpha = 0. \quad (64)$$

815 This thus implies,

$$816 \quad \|\nabla_\alpha \pi_\alpha\|_2 = \|u\|_2 \leq \|\pi' - \pi\|_1 \leq 2S, \quad \text{and} \quad \|\nabla_\alpha^2 \pi_\alpha\|_2 = 0, \quad (65)$$

818 Thus, π_α is both $\|u\|_2$ -Lipschitz and 0-smooth with respect to α , for all u that can be represented as
819 the difference of any two policies.

820 From the definition of \mathbb{P}^π , we have

$$821 \quad \mathbb{P}^{\pi_\alpha}(s'|s) = \sum_{a \in \mathcal{A}} \pi_\alpha(a|s) \mathbb{P}(s'|s, a) \quad (66)$$

$$822 \quad = \sum_{a \in \mathcal{A}} [\pi(a|s) + \alpha u(a|s)] \mathbb{P}(s'|s, a) \quad (67)$$

$$823 \quad \implies \frac{\partial \mathbb{P}^{\pi_\alpha}(s'|s)}{\partial \alpha} = \sum_{a \in \mathcal{A}} u(a|s) \mathbb{P}(s'|s, a). \quad (68)$$

829 That is,

$$830 \quad \nabla_\alpha \mathbb{P}^{\pi_\alpha} = \mathbb{P}^u, \quad \text{consequently} \quad \nabla_\alpha^2 \mathbb{P}^{\pi_\alpha} = 0. \quad (69)$$

832 From the definition of r^π , we have

$$833 \quad r^{\pi_\alpha}(s) = \sum_{a \in \mathcal{A}} \pi_\alpha(a|s) r(s, a) \quad (70)$$

$$834 \quad = \sum_{a \in \mathcal{A}} [\pi(a|s) + \alpha u(a|s)] r(s, a) \quad (71)$$

$$835 \quad \implies \frac{\partial r^{\pi_\alpha}(s)}{\partial \alpha} = \sum_{a \in \mathcal{A}} u(a|s) r(s, a). \quad (72)$$

841 That is,

$$842 \quad \nabla_\alpha r^{\pi_\alpha} = r^u, \quad \text{consequently} \quad \nabla_\alpha^2 r^{\pi_\alpha} = 0. \quad (73)$$

844 Hence the policy π_α , the associated reward r^{π_α} and the transition kernel \mathbb{P}^{π_α} are all Lipschitz and
845 smooth with respect to α .

846 **Lemma 13.** Let $A(\alpha) \in \mathbb{R}^{n \times n}$ be a matrix such that $(I - A(\alpha))$ is invertible for all $\alpha \in [0, 1]$.
847 Define $M(\alpha) := (I - A(\alpha))^{-1}$. Then it is true that,

$$848 \quad \frac{\partial^2 M(\alpha)}{\partial \alpha^2} = \frac{\partial M(\alpha)}{\partial \alpha} \frac{\partial A(\alpha)}{\partial \alpha} M(\alpha) + M(\alpha) \frac{\partial^2 A(\alpha)}{\partial \alpha^2} M(\alpha) + M(\alpha) \frac{\partial A(\alpha)}{\partial \alpha} \frac{\partial M(\alpha)}{\partial \alpha}. \quad (74)$$

851 *Proof.*

$$852 \quad M(\alpha) (I - A(\alpha)) = I \quad (75)$$

$$853 \quad \frac{\partial M(\alpha)}{\partial \alpha} (I - A(\alpha)) - M(\alpha) \frac{\partial A(\alpha)}{\partial \alpha} = 0 \quad (76)$$

$$854 \quad \frac{\partial M(\alpha)}{\partial \alpha} = M(\alpha) \frac{\partial A(\alpha)}{\partial \alpha} M(\alpha) \quad (77)$$

$$855 \quad \frac{\partial^2 M(\alpha)}{\partial \alpha^2} = \frac{\partial}{\partial \alpha} \left(M(\alpha) \frac{\partial A(\alpha)}{\partial \alpha} M(\alpha) \right), \quad (78)$$

$$856 \quad = \frac{\partial M(\alpha)}{\partial \alpha} \frac{\partial A(\alpha)}{\partial \alpha} M(\alpha) + M(\alpha) \frac{\partial^2 A(\alpha)}{\partial \alpha^2} M(\alpha) + M(\alpha) \frac{\partial A(\alpha)}{\partial \alpha} \frac{\partial M(\alpha)}{\partial \alpha} \quad (79)$$

863 \square

864 Consider the following definition utilized in the proofs of the upcoming lemmas.
865

$$866 \quad M^{\pi_\alpha} = [I - \Phi \mathbb{P}^{\pi_\alpha}]^{-1} \quad (80)$$

867 **Lemma 14.** Recall the definition of the projected average reward value function v_ϕ^π in Equation equa-
868 tion 36. Value function v_ϕ^π is $2C_m^2 C_p \kappa_r + 2C_m C_r$ -Lipschitz in Π , that is
869

$$870 \quad \left\langle \frac{\partial v_\phi^\pi}{\partial \pi}, \pi' - \pi \right\rangle \leq 2(C_m^2 C_p \kappa_r + C_m C_r) \|\pi' - \pi\|_2, \quad \forall \pi, \pi' \in \Pi. \quad (81)$$

873 *Proof.*

$$874 \quad v_\phi^{\pi_\alpha} = M^{\pi_\alpha} \Phi r^{\pi_\alpha} \quad (82)$$

$$875 \quad \Rightarrow \frac{\partial v_\phi^{\pi_\alpha}}{\partial \alpha} = \frac{\partial M^{\pi_\alpha}}{\partial \alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \frac{\partial r^{\pi_\alpha}}{\partial \alpha} \quad (83)$$

$$876 \quad = M^{\pi_\alpha} \frac{\partial \Phi \mathbb{P}^{\pi_\alpha}}{\partial \alpha} M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \frac{\partial r^{\pi_\alpha}}{\partial \alpha}, \quad (\text{from Lemma 13}), \quad (84)$$

$$877 \quad = M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u, \quad (\text{from equation 69 and equation 73}). \quad (85)$$

$$878 \quad \Rightarrow \left\| \frac{\partial v_\phi^{\pi_\alpha}}{\partial \alpha} \right\|_\infty = \|M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u\|_\infty \quad (86)$$

$$879 \quad \leq \|M^{\pi_\alpha}\|_\infty \|\Phi\|_\infty \|\mathbb{P}^u\|_\infty \|M^{\pi_\alpha}\|_\infty \|\Phi r^{\pi_\alpha}\|_\infty + \|M^{\pi_\alpha}\|_\infty \|\Phi r^u\|_\infty \quad (87)$$

$$880 \quad \leq 2C_m^2 C_p \kappa_r + 2C_m C_r. \quad (88)$$

881 The constants C_m, C_p, C_r and κ_r are characterized in Table 1 with their respective bounds in Lemma
882 18. \square

883 We can now build on the previous lemma to prove the smoothness of the average reward value
884 function.

885 **Lemma 15.** The value function v_ϕ^π is $8(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r)$ -smooth in Π . That is,
886

$$887 \quad \left\langle \pi' - \pi, \frac{\partial^2 v_\phi^\pi(s)}{\partial \pi^2} (\pi' - \pi) \right\rangle \leq 8(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r) \|\pi' - \pi\|_2^2 \quad \forall \pi', \pi \in \Pi, s \in \mathcal{S} \quad (89)$$

888 *Proof.* From Lemma 14, it is true that

$$889 \quad \frac{\partial v_\phi^{\pi_\alpha}}{\partial \alpha} = M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u$$

$$890 \quad \Rightarrow \frac{\partial^2 v_\phi^{\pi_\alpha}}{\partial \alpha^2} = \frac{\partial}{\partial \alpha} \left[M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u \right]$$

$$891 \quad = \frac{\partial M^{\pi_\alpha}}{\partial \alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \mathbb{P}^u \frac{\partial M^{\pi_\alpha}}{\partial \alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \frac{\partial r^{\pi_\alpha}}{\partial \alpha}$$

$$892 \quad + \frac{\partial M^{\pi_\alpha}}{\partial \alpha} \Phi r^u$$

$$893 \quad = M^{\pi_\alpha} \frac{\partial \Phi \mathbb{P}^{\pi_\alpha}}{\partial \alpha} M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \frac{\partial \Phi \mathbb{P}^{\pi_\alpha}}{\partial \alpha} M^{\pi_\alpha} \Phi r^{\pi_\alpha}$$

$$894 \quad + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \frac{\partial r^{\pi_\alpha}}{\partial \alpha} + M^{\pi_\alpha} \frac{\partial \Phi \mathbb{P}^{\pi_\alpha}}{\partial \alpha} M^{\pi_\alpha} \Phi r^u, \quad (\text{from Lemma 13}),$$

$$895 \quad = M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha}$$

$$896 \quad + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u, \quad (\text{from equation 69 and equation 73}).$$

918 Considering the L_∞ norm,

$$\begin{aligned}
919 \left\| \frac{\partial^2 v_\phi^{\pi_\alpha}}{\partial \alpha^2} \right\|_\infty &= 2 \|M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u\|_\infty \\
920 &\leq 2 \|M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha}\|_\infty + \|M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u\|_\infty \\
921 &\leq 8(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r).
\end{aligned}$$

925 Hence, we obtain,

$$\left\langle \pi' - \pi, \frac{\partial^2 v_\phi^{\pi_\alpha}(s)}{\partial \pi} (\pi' - \pi) \right\rangle \leq 8 (C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r) \|\pi' - \pi\|_2^2 \quad \forall \pi', \pi \in \Pi, s \in \mathcal{S} \quad (90)$$

931 \square

933 A.4 LIPSCHITZNESS OF THE INFINITE HORIZON AVERAGE REWARD ρ^π

934 The Lipschitzness and smoothness of the projected value function v_ϕ^π is leveraged through the average reward Bellman equation to prove the Lipschitzness and smoothness of the infinite horizon average reward.

935 **Lemma 16.** *Recall the average reward Bellman Equation corresponding to a policy π and projected value function v_ϕ^π in Equation equation 32. The average reward ρ^π is L_1^Π -Lipschitz.*

$$\left| \left\langle \frac{\partial \rho^\pi}{\partial \pi}, \pi' - \pi \right\rangle \right| \leq L_1^\Pi \|\pi' - \pi\|_2, \quad \forall \pi, \pi' \in \Pi, \quad (91)$$

944 where $L_1^\Pi = 2(C_r + C_p C_m \kappa_r + 2(C_m^2 C_p \kappa_r + C_m C_r))$

946 *Proof.* From Equation equation 32,

$$\rho^\pi \mathbb{1} = r^\pi + \mathbb{P}^\pi v_\phi^\pi - v_\phi^\pi. \quad (92)$$

949 Taking derivative with respect to α ,

$$\frac{\partial \rho^{\pi_\alpha}}{\partial \alpha} \mathbb{1} = \frac{\partial r^{\pi_\alpha}}{\partial \alpha} + \frac{\partial \mathbb{P}^{\pi_\alpha}}{\partial \alpha} v_\phi^{\pi_\alpha} + \mathbb{P}^{\pi_\alpha} \frac{\partial v_\phi^{\pi_\alpha}}{\partial \alpha} - \frac{\partial v_\phi^{\pi_\alpha}}{\partial \alpha} \quad (93)$$

$$= \Phi r^u + \Phi \mathbb{P}^u v^{\pi_\alpha} + (\mathbb{P}^{\pi_\alpha} - I)(M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u), \quad (\text{from Lemma 14}) \quad (94)$$

$$= \Phi r^u + \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + (\mathbb{P}^{\pi_\alpha} - I)(M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u), \quad (95)$$

$$\text{(from equation 73 and equation 69).} \quad (96)$$

$$(97)$$

962 Considering the L_∞ norm of the above expression,

$$\left| \frac{\partial \rho^{\pi_\alpha}}{\partial \alpha} \right| = \left\| \Phi r^u + \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + (\mathbb{P}^{\pi_\alpha} - I)(M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u) \right\|_\infty, \quad (98)$$

$$\leq \|\Phi r^u\|_\infty + \|\Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha}\|_\infty + \|\mathbb{P}^{\pi_\alpha} - I\|_\infty (\|M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha}\|_\infty + \|M^{\pi_\alpha} \Phi r^u\|_\infty), \quad (99)$$

$$\leq 2\|r^u\|_\infty + 2\|\mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha}\|_\infty + \|\mathbb{P}^{\pi_\alpha} - I\| (\|M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha}\|_\infty + \|M^{\pi_\alpha} \Phi r^u\|_\infty), \quad (100)$$

$$\leq 2C_r + 2C_p C_m \kappa_r + 2(2C_m^2 C_p \kappa_r + 2C_m C_r). \quad (101)$$

971 \square

972 A.5 SMOOTHNESS OF THE INFINITE HORIZON AVERAGE REWARD ρ^π

973
974 **Lemma 17.** *The average reward ρ^π is L_2^Π -smooth.*

$$975 \left\langle \pi' - \pi, \frac{\partial^2 \rho^\pi}{\partial \pi^2} (\pi' - \pi) \right\rangle \leq \frac{L_2^\Pi}{2} \|\pi' - \pi\|_2^2 \quad \forall \pi, \pi' \in \Pi, \quad (102)$$

976
977 where $L_2^\Pi = 4(C_p^2 C_m^2 \kappa_r + C_p C_m C_r + (C_p + 1)(C_m^2 C_p \kappa_r + C_m C_r) + 4(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r))$.

978
979 *Proof.* From Lemma 16, we have

$$980 \frac{\partial \rho^{\pi_\alpha}}{\partial \alpha} \mathbb{1} = \Phi r^u + \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + (\mathbb{P}^{\pi_\alpha} - I)(M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u). \quad (103)$$

981
982 Taking the derivative again, and repeatedly invoking Equations equation 69, equation 73 and Lemma 13, it follows that,

$$983 \frac{\partial^2 \rho^{\pi_\alpha}}{\partial \alpha^2} \mathbb{1} = 0 + \frac{\partial}{\partial \alpha} (\Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha}) + \frac{\partial}{\partial \alpha} ((\mathbb{P}^{\pi_\alpha} - I)(M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u))$$

$$984 = \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u + (\mathbb{P}^u)(M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u) \quad (104)$$

$$985 + (\mathbb{P}^{\pi_\alpha} - I) (M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} \quad (105)$$

$$986 + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u), \quad (106)$$

$$987 = \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u + (\mathbb{P}^u)(M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi r^u) \quad (107)$$

$$988 + 2(\mathbb{P}^{\pi_\alpha} - I) (M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u). \quad (108)$$

989
990 Considering the L_∞ norm of the above expression,

$$991 \left| \frac{\partial^2 \rho^{\pi_\alpha}}{\partial \alpha^2} \right| \leq \|\Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha}\|_\infty + \|\Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u\|_\infty + \|(\mathbb{P}^u)(M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} \quad (109)$$

$$992 + M^{\pi_\alpha} \Phi r^u)\|_\infty + 2\|(\mathbb{P}^{\pi_\alpha} - I) (M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^{\pi_\alpha} + M^{\pi_\alpha} \Phi \mathbb{P}^u M^{\pi_\alpha} \Phi r^u)\|_\infty$$

$$993 \leq 4(C_p^2 C_m^2 \kappa_r + C_p C_m C_r + (C_p + 1)(C_m^2 C_p \kappa_r + C_m C_r) + 4(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r)).$$

994 \square

995
996 **Remark:** The smoothness and Lipschitz constant analysis of both the average reward value functions and the infinite horizon average reward are constrained to all directions u , such that every $u = \pi - \pi'$ can be expressed as a difference of any two policies $\pi, \pi' \in \Pi$. Hence the smoothness and Lipschitz constants derived are restricted to the directions that can be expressed as this difference and hence are referred to as restricted smoothness/Lipschitzness.

1000 A.6 TABLE OF CONSTANTS CAPTURING MDP COMPLEXITY

1001 We restate the table of constants and their description here for the sake of convenience.

Table 2: Constants capturing the MDP Complexity

	Definition	Range	Remark
C_m	$\max_{\pi \in \Pi} \ (I - \Phi P^\pi)^{-1}\ $	$\frac{2C_e S }{1-\lambda}$	Lowest rate of mixing
C_p	$\max_{\pi, \pi' \in \Pi} \frac{\ P^{\pi'} - P^\pi\ }{\ \pi' - \pi\ _2}$	$[0, \sqrt{A}]$	Diameter of transition kernel
C_r	$\max_{\pi, \pi'} \frac{\ r^{\pi'} - r^\pi\ _\infty}{\ \pi' - \pi\ _2}$	$[0, \sqrt{A}]$	Diameter of reward function
κ_r	$\max_{\pi} \ \Phi r^\pi\ _\infty$	$[0, 2]$	Variance of reward function
$\frac{L_1^\Pi}{2}$	$C_r + C_p C_m \kappa_r + 2(C_m^2 C_p \kappa_r + C_m C_r)$	$[0, k_1 \sqrt{A} C_m^2]$	Restricted Lipschitz constant
$\frac{L_2^\Pi}{4}$	$C_p^2 C_m^2 \kappa_r + C_p C_m C_r + (C_p + 1)(C_m^2 C_p \kappa_r + C_m C_r) + 4(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r)$	$[0, k_2 A C_m^3]$	Restricted smoothness constant

Lemma 18. *The constants C_p, C_r, C_m, κ_r in Table 2 and other operator norms are bounded as below:*

1. $\|\Phi\| := \max_{\|v\|_\infty \leq 1} \|\Phi v\|_\infty \leq 2.$
2. $\|P^\pi\| = \max_{\|v\|_\infty \leq 1} \|P^\pi v\|_\infty \leq \max_{\|v\|_\infty \leq 1} \|v\|_\infty = 1.$
3. $\kappa_r = \max_{\pi} \|\Phi r^\pi\|_\infty \leq 2$
4. $C_m \leq \frac{2C_e S}{1-\lambda}$
5. $C_p = \max_{u = \frac{\pi' - \pi}{\|\pi' - \pi\|_2}, \pi', \pi \in \Pi} \max_{\|v\|_\infty \leq 1} \|P^u v\|_\infty \leq \sqrt{A}.$
6. $C_r = \max_{u = \frac{\pi' - \pi}{\|\pi' - \pi\|_2}, \pi', \pi \in \Pi} \|R^u\|_\infty \leq \sqrt{A}.$

Proof. 1. Consider the projection matrix Φ ,

$$\|\Phi\|_\infty = \max_{\|v\|_\infty \leq 1} \|\Phi v\|_\infty \leq \max_{\|v\|_\infty} \quad (110)$$

$$= \max_{s \in \mathcal{S}} \left| v(s) - \frac{\sum_{s \in \mathcal{S}} v(s)}{S} \right| \quad (111)$$

$$\leq \max_{s \in \mathcal{S}} |v(s)| + \left| \frac{\sum_{s \in \mathcal{S}} v(s)}{S} \right| \quad (112)$$

$$\leq 2\|v\|_\infty = 2 \quad (113)$$

2. The operator norm of \mathbb{P}^π is bounded as below:

$$\|P^\pi\| = \max_{\|v\|_\infty \leq 1} \|P^\pi v\|_\infty \quad (114)$$

$$\leq \max_{\|v\|_\infty \leq 1} \|v\|_\infty \quad (115)$$

$$\leq 1. \quad (116)$$

Equality is attained by the vector $v = \mathbb{1}$.

3. κ_r is bounded as below:

$$\kappa_r = \max_{\pi \in \Pi} \|\Phi r^\pi\|_\infty \quad (117)$$

$$= \max_{\pi \in \Pi} \left\| r^\pi - \frac{\sum_{s \in \mathcal{S}} r^\pi(s)}{|\mathcal{S}|} \mathbb{1} \right\|_\infty \quad (118)$$

$$\leq \max_{\pi \in \Pi} \|r^\pi\|_\infty + \left\| \frac{\sum_{s \in \mathcal{S}} r^\pi(s)}{|\mathcal{S}|} \mathbb{1} \right\|_\infty \quad (119)$$

$$\leq 2 \quad (120)$$

κ_r , in some sense, captures the variance of the single step reward function across the class of policies. Greater the variation of the r across different actions, greater the value of κ_r .

4. C_m is the maximum of the operator norm of the matrix $(I - \Phi \mathbb{P}^\pi)^{-1}$ across all policies $\pi \in \Pi$. It is determined as follows:

$$(I - \Phi \mathbb{P}^\pi)^{-1} = \sum_{k=0}^{\infty} (\Phi \mathbb{P}^\pi)^k = \sum_{k=0}^{\infty} \Phi (\mathbb{P}^\pi)^k, \quad (\text{from Lemma 12}), \quad (121)$$

$$\stackrel{(a)}{=} \sum_{k=0}^{\infty} \Phi ((\mathbb{P}^\pi)^k - \mathbb{1} (d^\pi)^\top), \quad (\text{as } \Phi \mathbb{1} (d^\pi)^\top = 0) \quad (122)$$

Let $v \in \mathbb{R}^{|\mathcal{S}|}$ such that $\|v\|_\infty \leq 1$. Then,

$$\implies \|(I - \Phi \mathbb{P}^\pi)^{-1} v\|_\infty \leq \sum_{k=0}^{\infty} \|\Phi ((\mathbb{P}^\pi)^k - \mathbb{1} (d^\pi)^\top) v\|_\infty \quad (123)$$

$$\leq \sum_{k=0}^{\infty} \|\Phi\|_\infty \|((\mathbb{P}^\pi)^k - \mathbb{1} (d^\pi)^\top) v\|_\infty \quad (124)$$

$$\leq \sum_{k=0}^{\infty} \|\Phi\|_\infty |\mathcal{S}| \|((\mathbb{P}^\pi)^k - \mathbb{1} (d^\pi)^\top)\|_\infty \|v\|_\infty \quad (125)$$

$$\stackrel{(b)}{\leq} \sum_{k=0}^{\infty} 2|\mathcal{S}| C_e \lambda^k \|v\|_\infty, \quad (126)$$

$$= \frac{2C_e |\mathcal{S}|}{1 - \lambda} \quad (127)$$

where d^π represents the stationary measure associated with the transition kernel \mathbb{P}^π , (a) follows from the fact that the projection matrix Φ projects vectors onto a subspace orthogonal to the subspace spanned by the all ones vector $\mathbb{1}$ and (b) is a consequence of the irreducibility and aperiodicity assumption of the Markov chain induced under all policies. More precisely, for any irreducible and aperiodic stochastic matrix A , it is true that:

$$\|A^n - \mathbb{1} d^\top\|_\infty \leq C_e \lambda^n, \quad (128)$$

for some constants $\lambda \in [0, 1)$, $C_e < \infty$, where d is stationary distribution of A . λ is the coefficient of mixing and captures the rate of geometric mixing of the Markov Chain. Hence, higher the value of λ , lower the rate of mixing.

5. C_p represents the diameter of the transition kernel as a function of the policy class and can be bound as below.

$$C_p := \max_{u = \frac{\pi' - \pi}{\|\pi' - \pi\|_2}, \pi', \pi \in \Pi} \max_{\|v\|_\infty \leq 1} \|\mathbb{P}^u v\|_\infty, \quad (129)$$

$$= \max_{\pi', \pi \in \Pi, \|v\|_\infty \leq 1} \frac{\|(\mathbb{P}^{\pi'} - \mathbb{P}^\pi)v\|_\infty}{\|\pi' - \pi\|_2}, \quad (130)$$

$$= \max_{\pi', \pi \in \Pi, \|v\|_\infty \leq 1} \max_{s \in \mathcal{S}} \frac{|(\mathbb{P}^{\pi'}v)(s) - (\mathbb{P}^\pi v)(s)|}{\|\pi' - \pi\|_2} \quad (131)$$

$$\leq \max_{\pi', \pi \in \Pi, \|v\|_\infty \leq 1} \max_{s \in \mathcal{S}} \frac{|(\mathbb{P}^{\pi'}v)(s) - (\mathbb{P}^\pi v)(s)|}{\|\pi'_s - \pi_s\|_2}, \quad (\text{since } \|\pi'_s - \pi_s\|_2 \leq \|\pi' - \pi\|_2), \quad (132)$$

$$= \max_{\pi', \pi \in \Pi, \|v\|_\infty \leq 1} \max_{s \in \mathcal{S}} \frac{|(\pi'_s)^\top \mathbb{P}_s v - (\pi_s)^\top \mathbb{P}_s v|}{\|\pi'_s - \pi_s\|_2}, \quad (\text{where } \mathbb{P}_s(a, s') = P(s'|s, a), \pi_s(a) := \pi(s, a)), \quad (133)$$

$$= \max_{\pi', \pi \in \Pi, \|v\|_\infty \leq 1} \max_{s \in \mathcal{S}} \frac{|(\pi'_s - \pi_s)^\top \mathbb{P}_s v|}{\|\pi'_s - \pi_s\|_2}, \quad (134)$$

$$\leq \max_{\pi', \pi \in \Pi, \|v\|_\infty \leq 1} \max_{s \in \mathcal{S}} \frac{\|\pi'_s - \pi_s\|_1 \|\mathbb{P}_s v\|_\infty}{\|\pi'_s - \pi_s\|_2}, \quad (\text{from Holder's inequality}) \quad (135)$$

$$= \max_{\pi', \pi \in \Pi} \max_{s \in \mathcal{S}} \frac{\|\pi'_s - \pi_s\|_1}{\|\pi'_s - \pi_s\|_2}, \quad (136)$$

$$\leq \sqrt{|\mathcal{A}|}. \quad (137)$$

6. C_r represents the diameter of the single step reward function as a function of the policy class and can be bound as below.

$$C_r := \max_{u = \frac{\pi' - \pi}{\|\pi' - \pi\|_2}, \pi', \pi \in \Pi} \|r^u\|_\infty \quad (138)$$

$$= \max_{\pi', \pi \in \Pi} \frac{\|r^{\pi'} - r^\pi\|_\infty}{\|\pi' - \pi\|_2} \quad (139)$$

$$= \max_{\pi', \pi \in \Pi} \max_{s \in \mathcal{S}} \frac{|r^{\pi'}(s) - r^\pi(s)|}{\|\pi' - \pi\|_2} \quad (140)$$

$$\leq \max_{\pi', \pi \in \Pi} \max_{s \in \mathcal{S}} \frac{|r^{\pi'}(s) - r^\pi(s)|}{\|\pi'_s - \pi_s\|_2}, \quad (\text{as } \|\pi'_s - \pi_s\|_2 \leq \|\pi' - \pi\|_2), \quad (141)$$

$$= \max_{\pi', \pi \in \Pi} \max_{s \in \mathcal{S}} \frac{|(\pi'_s)^\top r_s - (\pi_s)^\top r_s|}{\|\pi'_s - \pi_s\|_2}, \quad (\text{where } r_s(a) = r(s, a), \pi_s(a) = \pi(s, a)), \quad (142)$$

$$C_r = \max_{\pi', \pi \in \Pi} \max_{s \in \mathcal{S}} \frac{|(\pi'_s - \pi_s)^\top r_s|}{\|\pi'_s - \pi_s\|_2}, \quad (143)$$

$$\leq \max_{\pi', \pi \in \Pi} \max_{s \in \mathcal{S}} \frac{\|\pi'_s - \pi_s\|_1 \|r_s\|_\infty}{\|\pi'_s - \pi_s\|_2}, \quad (\text{from Holder's inequality}) \quad (144)$$

$$= \max_{\pi', \pi \in \Pi} \max_{s \in \mathcal{S}} \frac{\|\pi'_s - \pi_s\|_1}{\|\pi'_s - \pi_s\|_2}, \quad (145)$$

$$\leq \sqrt{|\mathcal{A}|}. \quad (146)$$

Since the directional derivatives considered are all within the policy class, the analysis gives rise to constants such as C_p and C_r , which are functions of the underlying policy class. These constants capture the MDP complexity by the virtue of their definition and are an artifact of this proof technique. \square

B CONVERGENCE OF AVERAGE REWARD PROJECTED POLICY GRADIENT

Lemma 19. For any convex set $\mathcal{X} \subseteq \mathbb{R}^d$, any point $a \in \mathcal{X}$, and any update direction $u \in \mathbb{R}^d$, let $b = \text{Proj}_{\mathcal{X}}(a + u)$ be the projection of $a + u$ onto \mathcal{X} . It is true that

1. $\langle u, b - a \rangle \geq \|b - a\|_2^2$.
2. $\langle c - b, u - (b - a) \rangle \leq 0, \quad \forall c \in \mathcal{X}$.

Proof. The formal proof can be found in Beck (2014).

However, the proof follows trivially from the geometrical representation of projection (see Figure 3, Kumar et al. (2023)), and the fact that the hyperplane separates a convex set from a point not in the set.




Figure 3: Convex Projection

Intuitively, the proof of the lemma can be interpreted as below.

1. Since the angle between vectors $(a - b)$ and $((a + u) - b)$ is greater than 90 degrees, it is true that $\langle a - b, (a + u) - b \rangle \leq 0$, which then directly implies $\|b - a\|_2^2 \leq \langle u, b - a \rangle$.
2. The angle between vectors $(c - b)$ and $((a + u) - b)$ is greater than 90 degrees $\forall c \in \mathcal{X}$, therefore $\langle c - b, u - (b - a) \rangle \leq 0$.

□

B.1 PROOF OF LEMMA 5

Lemma 20. The average reward iterates ρ^{π_k} generated from projected policy gradient satisfy the following,

$$\rho^{\pi_{k+1}} - \rho^{\pi_k} \geq \frac{L_2^\Pi}{2} \|\pi_{k+1} - \pi_k\|^2, \quad \forall k \geq 0.$$

where L_2^Π is the restricted smoothness constant associated with average reward ρ^π .

Proof. From the restricted smoothness of the average cost, we have

$$\begin{aligned} \rho^{\pi_{k+1}} &\geq \rho^{\pi_k} + \left\langle \frac{d\rho^\pi}{d\pi} \Bigg|_{\pi=\pi_k}, \pi_{k+1} - \pi_k \right\rangle - \frac{L_2^\Pi}{2} \|\pi_{k+1} - \pi_k\|^2, \\ &= \rho^{\pi_k} + L_2^\Pi \left\langle \frac{1}{L_2^\Pi} \frac{d\rho^\pi}{d\pi} \Bigg|_{\pi=\pi_k}, \pi_{k+1} - \pi_k \right\rangle - \frac{L_2^\Pi}{2} \|\pi_{k+1} - \pi_k\|^2, \\ &\geq \rho^{\pi_k} + L_2^\Pi \|\pi_{k+1} - \pi_k\|^2 - \frac{L_2^\Pi}{2} \|\pi_{k+1} - \pi_k\|^2. \end{aligned}$$

The last inequality follows from the projected gradient ascent policy update rule and item 1 of Lemma 19. Note that the proof only relies on the convexity of the projection set Π and the smoothness of the objective function. \square

B.2 PROOF OF LEMMA 7

Lemma 21. *The suboptimality of a policy π can be bounded from above as:*

$$\rho^* - \rho^\pi \leq C_{PL} \max_{\pi' \in \Pi} \left\langle \pi' - \pi, \frac{\partial \rho^\pi}{\partial \pi} \right\rangle, \quad \forall \pi \in \Pi, \quad (147)$$

where $C_{PL} = \max_{\pi, s} \frac{\partial \pi^*(s)}{\partial \pi(s)}$ and ρ^* is the optimal average reward.

Proof. Average Reward Performance Difference Lemma states that

$$\rho^* - \rho^\pi = \sum_{s \in \mathcal{S}} d^{\pi^*}(s) Q^\pi(s, a) [\pi^*(a|s) - \pi(a|s)] \quad (148)$$

$$\leq \max_{\pi'} \sum_{s \in \mathcal{S}} d^{\pi^*}(s) Q^\pi(s, a) [\pi'(a|s) - \pi(a|s)] \quad (149)$$

$$= \sum_{s \in \mathcal{S}} \frac{d^{\pi^*}(s)}{d^\pi(s)} d^\pi(s) \max_{\pi'_s} Q^\pi(s, a) [\pi'(a|s) - \pi(a|s)] \quad (150)$$

$$= \sum_{s \in \mathcal{S}} \frac{d^{\pi^*}(s)}{d^\pi(s)} \underbrace{d^\pi(s) \max_{\pi'_s} Q^\pi(s, a) [\pi'(a|s) - \pi(a|s)]}_{\geq 0, \quad (= 0 \text{ when } \pi'_s = \pi_s)} \quad (151)$$

$$\leq \sum_{s \in \mathcal{S}} \left(\max_{\pi, s} \frac{d^{\pi^*}(s)}{d^\pi(s)} \right) d^\pi(s) \max_{\pi'_s} Q^\pi(s, a) [\pi'(a|s) - \pi(a|s)] \quad (152)$$

$$= C_{PL} \max_{\pi'} \sum_{s \in \mathcal{S}} d^\pi(s) Q^\pi(s, a) [\pi'(a|s) - \pi(a|s)] \quad (153)$$

$$\stackrel{(a)}{=} C_{PL} \max_{\pi'} \left\langle \frac{\partial \rho^\pi}{\partial \pi}, \pi' - \pi \right\rangle, \quad (154)$$

where (a) follows from the average reward policy gradient theorem. \square

B.3 PROOF OF LEMMA 8

Lemma 22. *Let π_{k+1} represent the policy iterates obtained through projected policy gradient. For any policy $\pi' \in \Pi$, it is true that,*

$$\left\langle \frac{\partial \rho^{\pi_{k+1}}}{\partial \pi_{k+1}}, \pi' - \pi_{k+1} \right\rangle \leq 4\sqrt{|\mathcal{S}|} L_2^\Pi \|\pi_{k+1} - \pi_k\|_2, \quad (155)$$

Proof. For all $x, y \in C$, we have:

$$\left\langle \frac{\partial \rho^{\pi_{k+1}}}{\partial \pi_{k+1}}, \pi' - \pi_{k+1} \right\rangle = \left\langle \frac{\partial \rho^{\pi_{k+1}}}{\partial \pi_{k+1}} - \frac{\partial \rho^{\pi_k}}{\partial \pi_k} + \frac{\partial \rho^{\pi_k}}{\partial \pi_k}, \pi' - \pi_{k+1} \right\rangle \quad (156)$$

$$= \left\langle \frac{\partial \rho^{\pi_{k+1}}}{\partial \pi_{k+1}} - \frac{\partial \rho^{\pi_k}}{\partial \pi_k}, \pi' - \pi_{k+1} \right\rangle + \left\langle \frac{\partial \rho^{\pi_k}}{\partial \pi_k}, \pi' - \pi_{k+1} \right\rangle \quad (157)$$

$$\leq \left\| \frac{\partial \rho^{\pi_{k+1}}}{\partial \pi_{k+1}} - \frac{\partial \rho^{\pi_k}}{\partial \pi_k} \right\| \|\pi' - \pi_{k+1}\| + \left\langle \frac{\partial \rho^{\pi_k}}{\partial \pi_k}, \pi' - \pi_{k+1} \right\rangle \quad (158)$$

$$\stackrel{(a)}{\leq} L_2^\Pi \|\pi_{k+1} - \pi_k\| \|\pi' - \pi_{k+1}\| + \left\langle \frac{\partial \rho^{\pi_k}}{\partial \pi_k}, \pi' - \pi_{k+1} \right\rangle, \quad (159)$$

where (a) uses smoothness of average reward. Thus, we may continue the chain of inequalities as

$$\begin{aligned}
\text{equation 159} &= L_2^\Pi \|\pi_{k+1} - \pi_k\| \|\pi' - \pi_{k+1}\| + \left\langle \frac{\partial \rho^{\pi_k}}{\partial \pi_k} - L_2^\Pi(\pi_{k+1} - \pi_k), \pi' - \pi_{k+1} \right\rangle + L_2^\Pi \langle \pi_{k+1} - \pi_k, \pi' - \pi_{k+1} \rangle \\
&\leq 2L_2^\Pi \|\pi_{k+1} - \pi_k\| \|\pi' - \pi_{k+1}\| + \left\langle \frac{\partial \rho^{\pi_k}}{\partial \pi_k} - L_2^\Pi(\pi_{k+1} - \pi_k), \pi' - \pi_{k+1} \right\rangle \\
&\leq 2L_2^\Pi \|\pi_{k+1} - \pi_k\| \|\pi' - \pi_{k+1}\| + L_2^\Pi \underbrace{\left\langle \frac{1}{L_2^\Pi} \frac{\partial \rho^{\pi_k}}{\partial \pi_k} - (\pi_{k+1} - \pi_k), \pi' - \pi_{k+1} \right\rangle}_{\leq 0, \quad (\text{From item 2 of Lemma 19})} \\
&\leq 2L_2^\Pi \|\pi_{k+1} - \pi_k\| \|\pi' - \pi_{k+1}\| \\
&\leq 2L_2^\Pi \|\pi_{k+1} - \pi_k\| \mathbf{diam}(\Pi).
\end{aligned}$$

The diameter of the policy class Π , can be upper bounded as

$$\mathbf{diam}(\Pi)^2 = \max_{\pi, \pi'} \sum_s \|\pi'_s - \pi_s\|_2^2 \leq \max_{\pi', \pi} \sum_s \|\pi'_s - \pi_s\|_1^2 \leq 4S. \quad (160)$$

This yields the result. \square

Lemma 23. *The scaled sub-optimality $a_k := \rho^* - \rho_k$ follows the recursion*

$$ca_{k+1}^2 + a_{k+1} - a_k \leq 0, \quad (161)$$

where $c = \frac{1}{32L_2^\Pi |S| C_{PL}^2}$.

Proof. From Lemma 21, we know that,

$$\rho^* - \rho^{\pi_{k+1}} \leq C_{PL} \left\langle \pi' - \pi_{k+1}, \frac{\partial \rho^{\pi_{k+1}}}{\partial \pi_{k+1}} \right\rangle, \quad \forall \pi' \in \Pi, \quad (162)$$

From Lemma 22, we know that,

$$\left\langle \frac{\partial \rho^{\pi_{k+1}}}{\partial \pi_{k+1}}, \pi' - \pi_{k+1} \right\rangle \leq 4\sqrt{|S|} L_2^\Pi \|\pi_{k+1} - \pi_k\|_2. \quad (163)$$

From Lemma 20, we know that,

$$\|\pi_{k+1} - \pi_k\|_2 \leq \sqrt{\frac{2(\rho^{\pi_{k+1}} - \rho^{\pi_k})}{L_2^\Pi}}, \quad \forall k \geq 0. \quad (164)$$

Combining the above equations yields,

$$\rho^* - \rho^{\pi_{k+1}} \leq \sqrt{32C_{PL}^2 L_2^\Pi |S| (\rho^{\pi_{k+1}} - \rho^{\pi_k})} \quad (165)$$

This thus yields,

$$\frac{(\rho^* - \rho^{\pi_{k+1}})^2}{32C_{PL}^2 L_2^\Pi |S|} + (\rho^* - \rho^{\pi_{k+1}}) - (\rho^* - \rho^{\pi_k}) \leq 0. \quad (166)$$

\square

A more detailed interpretation of this Lemma can be found in Kumar et al. (2023).

B.4 RECURSION BOUND

In this subsection, we consider the sequence defined as

$$a_k - a_{k+1} \geq a_k^2,$$

where $p \geq 0$ and $0 \leq a_0 \leq 1$. Let f be linear interpolation of the sequence $\{a_k\}_{k \geq 0}$, formally defined as

$$f(x) := (1 - \alpha)a_k + \alpha a_{k+1}, \quad \text{where } k = \lfloor x \rfloor \text{ and } \alpha = x - \lfloor x \rfloor.$$

Let $g(0) := a_0$ and

$$\frac{dg(x)}{dx} = -g(x)^2, \quad \text{and} \quad \tau_k := g^{-1}(a_k).$$

Observe that g is a strictly decreasing function.

1350 **Proposition 1.** *If $\tau_k \geq k$ then*

$$1351 \quad g(x) \geq f(x), \quad \forall x \in [k, k+1].$$

1352 *Proof.* We have $f(k) = g(\tau_k) = a_k$ and for $\alpha \in (0, 1)$

$$1353 \quad \frac{dg(\tau_k + \alpha)}{dx} = -g(\tau_k + \alpha)^2, \quad (\text{by definition}) \quad (167)$$

$$1354 \quad \geq -g(\tau_k)^2, \quad (g \text{ is a decreasing function}) \quad (168)$$

$$1355 \quad = -f(k)^2, \quad (\text{as } f(k) = g(\tau_k)) \quad (169)$$

$$1356 \quad = -a_k^2 \quad (\text{by definition of } f) \quad (170)$$

$$1357 \quad \geq a_{k+1} - a_k, \quad (\text{by definition of } a_{k+1}) \quad (171)$$

$$1358 \quad = \frac{df(k + \alpha)}{dx}, \quad (\text{by definition of } f). \quad (172)$$

1359 Above together with continuity of f and g , for all $\alpha \in [0, 1]$, we have

$$1360 \quad f(k + \alpha) \leq g(\tau_k + \alpha), \quad (173)$$

$$1361 \quad \leq g(k + \alpha), \quad (\text{as } \tau_k \geq k \text{ and } g \text{ is a decreasing function}). \quad (174)$$

1362 Hence claim is proved. \square

1363 **Proposition 2.** *For all $k \geq 0$, we have*

$$1364 \quad \tau_k \geq k.$$

1365 *Proof.* Note that $\tau_0 = 0$ by definition $g(0) = a_0$. Now let $\tau_k \geq k$, then from Proposition 1, we have

$$1366 \quad g(x) \geq f(x), \quad \forall x \in [k, k+1] \quad (175)$$

$$1367 \quad \implies g(k+1) \geq f(k+1) \quad (176)$$

$$1368 \quad = a_{k+1} \quad (177)$$

$$1369 \quad \implies \tau_{k+1} \geq k+1, \quad (\text{as } g \text{ is a decreasing function}). \quad (178)$$

1370 Hence, by induction the claim is established. \square

1371 **Lemma 24. [Recursion Upper Bound]** *For $p \geq 2$, and $0 \leq a_0 \leq 1$, sequence $\{a_k\}_{k \geq 0}$ satisfying the recursion $a_k - a_{k+1} \geq a_k^2$, follows*

$$1372 \quad a_k \leq \frac{1}{\frac{1}{a_0} + k}, \quad \forall k \geq 1.$$

1373 *Proof.* From Proposition 2, we get $\tau_k \geq k$. Combining it with Proposition 1, we get

$$1374 \quad g(x) \geq f(x), \quad \forall x \geq 0.$$

1375 Now, we solve the o.d.e. to get

$$1376 \quad \frac{dg(x)}{dx} = -g(x)^2 \quad (179)$$

$$1377 \quad \implies dx \Big|_{x=0}^{x=k} = \int_{x=0}^k -\frac{dg(x)}{g^2(x)} \quad (180)$$

$$1378 \quad \implies k = \frac{1}{g(k)} - \frac{1}{g(0)} \quad (181)$$

$$1379 \quad \implies g(k) = \frac{g(0)}{1 + g(0)k} \quad (182)$$

$$1380 \quad \implies f(k) \leq \frac{a_0}{1 + a_0k}, \quad (\text{as } f(x) \leq g(x)) \quad (183)$$

$$1381 \quad \implies a_k \leq \frac{1}{\frac{1}{a_0} + k}. \quad (184)$$

1382 This proves the claim. \square

1404 **Lemma 25.** If $a_k - a_{k+1} \geq ca_k^2$ then

$$1405$$

$$1406 \quad a_k \leq \frac{1}{\frac{1}{a_0} + ck}.$$

$$1407$$

$$1408$$

1409
1410 *Proof.* We have

$$1411 \quad a_k - a_{k+1} \geq ca_k^2 \tag{185}$$

$$1412 \quad \implies ca_k - ca_{k+1} \geq (ca_k)^2 \tag{186}$$

$$1413 \quad \implies ca_k \leq \frac{1}{\frac{1}{ca_0} + k}, \quad (\text{from Lemma 24}) \tag{187}$$

$$1414 \quad \implies a_k \leq \frac{1}{\frac{1}{a_0} + ck}. \tag{188}$$

1415
1416
1417
1418
1419 □

1420
1421
1422 This subsection (proving the recursion upper bound) is inspired by a technique from Kumar et al.
1423 (2024). However, while their result is similar, it is not applicable to our case. Their result assumes
1424 that c is upper bounded by a constant, which does not hold in our setting. In our case, the smoothness
1425 constant (or hardness coefficient) can approach zero, causing the constant c to diverge to infinity. Our
1426 result is more general, and the proof technique we use is distinct.

1427 **Lemma 26.** Given $a_k - a_{k+1} \geq ca_{k+1}^2$, we have

$$1428$$

$$1429 \quad a_k \leq \frac{1}{\frac{1}{a_0} + ck},$$

$$1430$$

1431 where $\nu = c(1 + \frac{8c}{1-\gamma})^{-\frac{3}{2}}$ and $c = \frac{1}{32C_{PL}^2|S|L_2^\Pi}$.

1432
1433
1434
1435 *Proof.* We have

$$1436 \quad ca_{k+1}^2 + a_{k+1} - a_k \leq 0$$

$$1437 \quad \implies a_{k+1} \leq \frac{-1 + \sqrt{1 + 4ca_k}}{2c},$$

$$1438 \quad = \frac{-1 + f(0) + f'(0)4ca_k + f''(b)\frac{(4ca_k)^2}{2}}{2c}, \quad (\text{where } f(x) = \sqrt{1+x}, b \in [0, 4ca_k])$$

$$1439 \quad = \frac{2ca_k - 2c^2a_k^2(1+b)^{-\frac{3}{2}}}{2c}, \quad (\text{putting } f(0) = 1, f'(0) = \frac{1}{2}, f''(a) = \frac{(1+b)^{-\frac{3}{2}}}{4})$$

$$1440 \quad = \frac{2ca_k - 2c^2a_k^2(1+4ca_k)^{-\frac{3}{2}}}{2c}, \quad (\text{as } b \leq 4ca_k \text{ and } -(1+y)^{-\frac{3}{2}} \text{ is an increasing function})$$

$$1441 \quad \leq a_k - \frac{ca_k^2}{(1+4ca_k)^{\frac{3}{2}}}, \quad (\text{basic algebra})$$

$$1442 \quad \leq a_k - \nu a_k^2, \quad (\text{as } a_k \leq 1 \text{ and } \nu := c(1+4c)^{-\frac{3}{2}}).$$

1443
1444 We get the desired result from Lemma 25. □

1445
1446 **Lemma 27.** If $\frac{1}{c} = 32C_{PL}^2|S|L_2^\Pi < 1$ that is MDP is very easy (i.e. $L_2^\Pi \ll 1$) then the policy
1447 gradient converges exponentially fast, that is

$$1448$$

$$1449 \quad a_k \leq \left(\frac{1}{c}\right)^{\frac{k}{2}} a_0^{-2^k}.$$

$$1450$$

1458 *Proof.* From the above discussion, we have

$$1459 \quad ca_{k+1}^2 \leq a_k - a_{k+1} \quad (189)$$

$$1461 \quad \leq a_k, \quad (\text{as } a_{k+1} \geq 0 \text{ by definition}) \quad (190)$$

$$1463 \quad \Rightarrow a_{k+1} \leq \sqrt{\frac{a_k}{c}} \quad (191)$$

$$1465 \quad \leq \left(\frac{1}{c}\right)^{\frac{k+1}{2}} a_0^{\frac{1}{2^{k+1}}} \quad (192)$$

1467 □

1468

1469 B.5 PROOF OF THEOREM 1

1470

1471 We restate the theorem for the sake of convenience.

1472 **Theorem 2.** Let ρ^{π_k} be the average reward corresponding to the policy iterates π_k , obtained
 1473 through the policy gradient update equation 6. Let ρ^* represent the optimal average reward, that is,
 1474 $\rho^* = \max_{\pi \in \Pi} \rho^\pi$. There exist constants L_2^Π and C_{PL} which are determined by the underlying MDP
 1475 such that:

1476

1477 • For all MDPs it is true that,

$$1478 \quad \rho^* - \rho^{\pi_k} \leq \frac{1}{\frac{1}{\rho^* - \rho^{\pi_0}} + \nu k}, \quad \forall k \geq 0. \quad (193)$$

1480

1481 where $\nu := \left(\frac{1}{32C_{PL}^2|\mathcal{S}|L_2^\Pi}\right) \left(1 + 4\left(\frac{1}{32C_{PL}^2|\mathcal{S}|L_2^\Pi}\right)\right)^{-\frac{3}{2}}$

1482

1483 • For simple MDPs (i.e. $L_2^\Pi \ll 1$) we obtain exponential convergence, that is

$$1484 \quad \rho^* - \rho^{\pi_k} \leq c^{-\frac{k}{2}} (\rho^* - \rho^{\pi_0})^{\frac{1}{2^k}}, \quad \forall k \geq 0 \quad (194)$$

1485

1486 where $\frac{1}{c} = 32|\mathcal{S}|L_2^\Pi C_{PL}^2 < 1$.

1487

1488 *Proof.* Using Gradient Domination Lemma and Sufficient Increase Lemma as shown in Lemma 23,
 1489 we get the following recursion

$$1490 \quad a_k - a_{k+1} \geq ca_{k+1}^2, \quad (195)$$

1491

1492 where $a_k = J^* - J^{\pi_k}$ and $c = \frac{1}{32C_{PL}^2|\mathcal{S}|L_2^\Pi}$ is a small constant. Then we get the desired result by
 1493 solving the above recursion in Lemma 26 and Lemma 27 for complex MDPs and simple MDPs
 1494 respectively. □

1495

1496

1500 C SIMULATION DETAILS

1501

1502 C.1 CONVERGENCE WITH DIFFERENT ACTION AND STATE SPACE SIZE

1503

1504 In the first experiment, we compare the convergence of PG for a tabular MDP with $(S, A) \in$
 1505 $\{(3, 3), (9, 9), (81, 81)\}$. We set the reward kernel to be with maximal variance as described above.
 1506 For the transition kernel, we use the following matrix:

$$1507 \quad P(\cdot | s, \cdot) = \frac{1}{2} \left(\mathbf{1}_{S \times A} + \frac{1}{S} \right)$$

1508

1509 so $P(i|s, i) = \frac{1+\frac{1}{S}}{2}$ and $P(i|s, j) = \frac{1}{2S}$ for $i \neq j$. For the reward kernel we set the rewards of half
 1510 the actions to 1 and the rest to -1 , for every state.

1511

C.2 CONVERGENCE WITH DIFFERENT REWARD FUNCTIONS

In the second experiment, we compare the convergence of PG for a tabular MDP with $S = 16$ and $A = 16$. We set $r(s, a) = 0$ for any s and a except for one state which we denote by s_0 . We use the same randomly generated transition kernel and use the following procedure to generate the reward function:

- No variance: We assign each (s_0, a) pair a reward of 1.
- Low variance: We assign $\frac{1}{8}$ of the actions for s_0 a reward of -1 , and 1 otherwise.
- High variance: We assign $\frac{1}{4}$ of the actions for s_0 a reward of -1 , and 1 otherwise.
- Max variance: We assign $\frac{1}{2}$ of the actions for s_0 a reward of -1 , and 1 otherwise.

C.3 CONVERGENCE WITH DIFFERENT TRANSITION KERNELS

In the third experiment, we compare the convergence of PG for a tabular MDP with $S = 16$ and $A = 16$. We create three different MDPs with the same (S, A) values and the same reward function that is generated according to the process described above for high variance reward function. We then generate three different transition kernels:

- Uniform: We assign for all values of $s, a, s', P(s' | s, a) = \frac{1}{S}$.
- Non-uniform: We assign $P(i | s, i) = \frac{1}{2S} + \frac{1}{2}$ and $P(i | s, j) = \frac{1}{2S}$ for $i \neq j$.
- Deterministic: We look at $P(\cdot | s, \cdot)$ as an $S \times A$ matrix, and assign it a random permutation of the identity matrix. In this way the result MDP is deterministic but not trivial (so every state leads to a different one).

D ADDITIONAL DISCUSSION AND FUTURE WORK

Extension to Discounted Reward Setting. In the discounted reward setting, the return ρ^π , and the value function v^π is defined as

$$\rho^\pi = \mu^T (I - \gamma P^\pi) R^\pi, \quad v^\pi = (I - \gamma P^\pi) R^\pi, \quad (196)$$

where $\gamma \in [0, 1)$ is the discount factor Sutton & Barto (2018). The return ρ^π is proven to be $\frac{8}{(1-\gamma)^3}$ -smooth (Agarwal et al., 2020). Note that it is an MDP-agnostic bound. We can achieve an MDP instance-dependent bound with a very minor change in the smoothness analysis of the average-reward case.

Let us define,

$$M^{\pi_\alpha} := (I - \gamma P^{\pi_\alpha})^{-1}. \quad (197)$$

Observe that Lemma 13 holds for $A(\alpha) = \gamma P^{\pi_\alpha}$, which yields us

$$\frac{\partial^2 M(\alpha)}{\partial \alpha^2} = \frac{\partial M(\alpha)}{\partial \alpha} \frac{\partial A(\alpha)}{\partial \alpha} M(\alpha) + M(\alpha) \frac{\partial^2 A(\alpha)}{\partial \alpha^2} M(\alpha) + M(\alpha) \frac{\partial A(\alpha)}{\partial \alpha} \frac{\partial M(\alpha)}{\partial \alpha}, \quad (198)$$

where $M(\alpha)$ is shorthand for M^{π_α} .

Table 3: Constants capturing the MDP Complexity for Discounted reward

	Definition	Range	Remark
C_m	$\max_{\pi \in \Pi} \ (I - \gamma P^\pi)^{-1}\ $	$\frac{1}{1-\gamma}$	
C_p	$\gamma \max_{\pi, \pi' \in \Pi} \frac{\ P^{\pi'} - P^\pi\ }{\ \pi' - \pi\ _2}$	$[0, \gamma\sqrt{A}]$	Diameter of transition kernel
C_r	$\max_{\pi, \pi'} \frac{\ r^{\pi'} - r^\pi\ _\infty}{\ \pi' - \pi\ _2}$	$[0, \sqrt{A}]$	Diameter of reward function
κ_r	$\max_{\pi} \ r^\pi\ _\infty$	$[0, 1]$	Variance of reward function

Lemma 28. *The value function v_ϕ^π is $8(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r)$ -smooth in Π . That is,*

$$\left\langle \pi' - \pi, \frac{\partial^2 v_\phi^\pi(s)}{\partial \pi} (\pi' - \pi) \right\rangle \leq 8 (C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r) \|\pi' - \pi\|_2^2 \quad \forall \pi', \pi \in \Pi, s \in \mathcal{S} \quad (199)$$

Proof.

$$\begin{aligned} \frac{\partial^2 v^{\pi\alpha}}{\partial \alpha^2} &= M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} r^{\pi\alpha} + M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} r^{\pi\alpha} \\ &\quad + M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} r^u + M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} r^u, \quad (\text{from equation 69 and equation 73}). \end{aligned}$$

Considering the L_∞ norm,

$$\begin{aligned} \left\| \frac{\partial^2 v^{\pi\alpha}}{\partial \alpha^2} \right\|_\infty &= 2 \|M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} r^{\pi\alpha} + M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} r^u\|_\infty \\ &\leq 2 \|M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} r^{\pi\alpha}\|_\infty + \|M^{\pi\alpha} \gamma \mathbb{P}^u M^{\pi\alpha} r^u\|_\infty \\ &\leq 8(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r). \end{aligned}$$

Hence, we obtain,

$$\left\langle \pi' - \pi, \frac{\partial^2 v^\pi(s)}{\partial \pi} (\pi' - \pi) \right\rangle \leq 8 (C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r) \|\pi' - \pi\|_2^2 \quad \forall \pi', \pi \in \Pi, s \in \mathcal{S} \quad (200)$$

□

The above result implies the return ρ^π is $L_2^\Pi = 8(C_m^3 C_p^2 \kappa_r + C_m^2 C_p C_r)$ -smooth. This establishes the convergence of the projected policy gradient algorithm with an iteration complexity similar to that stated in Theorem 1 for the average reward case.

Follow up/concurrent work in discounted reward case: The work Liu et al. (2024) improved the iteration complexity of the policy gradient (discounted reward case) method to $O(\frac{A}{\epsilon})$ from the previous state-of-the-art iteration complexity of $O(\frac{SA}{\epsilon})$ Xiao (2022a); Mei et al. (2022). Liu et al. (2024) does not use smoothness of the return to establish sufficient increase lemma, instead leverages the performance difference lemma in a novel way.

- Our instance-dependent bound for the MDP stems from the smoothness of the return. As a result, it is unclear how the two approaches can be effectively combined to achieve tighter bounds, leaving this as an avenue for future research.
- However, the bound in Liu et al. (2024) can be asymptotically improved by a factor of $\frac{1}{1-\gamma}$, and more meaningful bounds for initial iterates can be obtained using the enhanced recursion-solving techniques presented in Kumar et al. (2024).

Our work can provide improved, alternative, or suboptimal results (depending on the parameters) for the discounted reward case. However, it remains the first to establish the global convergence of policy gradient methods for the average reward setting.

All existing works (Agarwal et al., 2020; Bhandari & Russo, 2024; Mei et al., 2022; Xiao, 2022a) on the discounted reward case have an iteration complexity that scales with the cardinalities of the state space S and the action space A , with the exception of the recent work by Liu et al. (2024), which depends only on A . This dependence on the size of the action space poses significant challenges when attempting to generalize to infinite state-action spaces.

Infinite Action Space. Our work (for both average and discounted reward case) has the MDP instance bounded bound of $\frac{SL_2^\Pi}{\epsilon}$, where L_2^Π is the smoothness constant that encodes the hardness of the MDP. This hardness coefficients that makes up L_2^Π , may be small/finite for even large/infinite action-space MDPs. However, it requires more careful study to determine the conditions for this to happen, which we leave for the future work.

Infinite State Space For the discounted reward setting, Liu et al. (2024) provides a state-independent bound of $O(\frac{A}{\epsilon})$. The approach taken in their work is fundamentally different from ours, and extending this technique to the average reward case presents an intriguing direction for future research.

In our work, which continues the line of research from Agarwal et al. (2020); Bhandari & Russo (2024); Mei et al. (2022); Xiao (2022a), the bound exhibits state dependence. This dependence arises from the diameter of the policy class, defined as $\text{diam}(\Pi)^2 = \sum_{\pi, \pi'} \|\pi - \pi'\|_2^2 \leq S$. While this quantity is inherently tied to the state space, it can potentially be bounded for infinite state spaces under certain structures, such as low-rank policy classes. Exploring this direction in greater detail is an intriguing avenue for future research. Another challenge is to circumvent the dependence on C_{PL} constant which captures the suboptimality of a policy. C_{PL} can be ∞ when the state space is infinite, necessitating a different approach to characterizing the suboptimality of a policy.

D.1 PARAMETRIZED POLICY CLASS LOWER BOUNDS

Parametrized policy class. Our work can also be extended to parameterized policy classes, such as softmax policies. For parameterized policy classes, the smoothness coefficient can be derived by augmenting our analysis with the chain and product rules, which is a straightforward extension. However, this may result in different hardness coefficients, making it an interesting direction for further exploration. We leave this investigation for future work.

Lower Bounds for policy gradient for average reward case. For the discounted reward setting, Mei et al. (2022) establishes a lower bound of $O(\epsilon^{-1})$ for policy gradient methods. Specifically, Theorem 9 of Mei et al. (2022) derives this lower bound using a bandit problem as a counterexample. Since a bandit is a special case of an MDP with a single state, it serves as an example for both discounted reward and average reward MDPs. Consequently, the same lower bound of $O(\epsilon^{-1})$ also applies to average reward MDPs, as implied by Theorem 9 of Mei et al. (2022).

Linear Rates of Policy Gradient with aggressively increasing step sizes. Policy gradient can be interpreted as a form of soft policy iteration, assuming all states are visited or updated by the policy. Specifically, as the learning rate increases, policy gradient behavior increasingly resembles policy iteration. Since policy iteration is known to converge linearly, which is significantly faster than the typical convergence rates of policy gradient methods, it is natural to expect linear convergence bounds for policy gradient with aggressively increasing learning rates. This has been established in several works, including Xiao (2022a); Johnson et al. (2023); Liu et al. (2024).

In most cases, the model is not known, and the exact gradient cannot be computed, requiring the use of stochastic gradient descent. In such noisy settings, using aggressive step sizes can lead to instability in the algorithm.

However, in Theorem 1, we demonstrated linear convergence rates for simple MDPs with constant step sizes. It may also be possible to achieve similar rates for the average reward case by employing aggressively increasing step sizes, which we leave as an interesting direction for future work.