PERFORMATIVE POLICY GRADIENT: ASCENT TO OPTIMALITY IN PERFORMATIVE REINFORCEMENT LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Post-deployment machine learning algorithms often influence the environments they act in, and thus *shift* the underlying dynamics that the standard reinforcement learning (RL) methods ignore. While designing optimal algorithms in this *performative* setting has recently been studied in supervised learning, the RL counterpart remains under-explored. In this paper, we prove the performative counterparts of the performance difference lemma and the policy gradient theorem in RL, and further introduce the **Performative Policy Gradient** algorithm (PePG). PePG is the first policy gradient algorithm designed to account for performativity in RL. Under softmax parametrisation, and also with and without entropy regularisation, we prove that PePG converges to *performatively optimal policies*, i.e. policies that remain optimal under the distribution shifts induced by themselves. Thus, PePG significantly extends the prior works in Performative RL that achieves *performative stability* but not optimality. Furthermore, our empirical analysis on standard performative RL environments validate that PePG outperforms standard policy gradient algorithms and the existing performative RL algorithms aiming for stability.

1 Introduction

Reinforcement Learning (RL) studies the dynamic decision making problems under incomplete information (Sutton & Barto, 1998). Since an RL algorithm tries and optimises an utility function over a sequence of interactions with an unknown environment, RL has emerged as a powerful tool for algorithmic decision making. Specially, in the last decade, RL has underpinned some of the celebrated successes of AI, such as championing Go with AlphaGo (Silver et al., 2014), controlling particle accelerators (St. John et al., 2021), aligning Large Language Models (LLMs) (Bai et al., 2022), reasoning (Havrilla et al.), to name a few. But the existing paradigm of RL assumes that the underlying environment with which the algorithm interacts stays static over time and the goal of the algorithm is to find the utility-maximising, aka optimal policy for choosing actions over time for this specific environment. But this assumption does not hold universally.

In this digital age, algorithms are not passive. Their decisions also shape the environment they interact with, inducing distribution shifts. This phenomenon that predictive AI models often trigger actions that influences their own outcomes is termed as *performativity*. In the supervised learning setting, the study of *performative prediction* is pioneered by Perdomo et al. (2020), and then followed by an extensive literature encompassing stochastic optimisation, control, multi-agent RL, games (Izzo et al., 2021; 2022; Miller et al., 2021; Li & Wai, 2022; Narang et al., 2023; Piliouras & Yu, 2023; Góis et al., 2024; Barakat et al., 2025) etc.

There has been several attempts to achieve performative optimality or stability for real-life tasks—recommendation systems (Eilat & Rosenfeld, 2023), measuring the power of firms (Hardt et al., 2022; Mofakhami et al., 2023), healthcare (Zhang et al., 2022) etc. Performativity of algorithms is also omnipresent in practically deployed RL systems. For example, an RL algorithm deployed in a recommender system does not only aim to maximise the user satisfaction but also shifts the preferences of the users in the long-term (Chaney et al., 2018; Mansoury et al., 2020). To clarify the impact of performativity, let us consider an example.

Example 1 (Performative RL in loan approval). Let us consider a loan approval problem, where an applicant obtains a loan (or get rejected) according to their credit score x, and x depends on the capital of the applicant and that of the population. At each time t, a loan applicant arrives with a credit score x_t sampled from $\mathcal{N}(\mu_t, \sigma^2)$. The bank chooses whom to give a loan by applying a softmax binary classifier $\pi_{\theta} : \mathbb{R} \to \mathbb{R}$

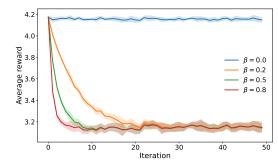


Figure 1: Average reward (over 10 runs) obtained by ERM and Performative Optimal policies across performative strength β .

 $\{0,1\}$ on x with threshold parameter θ . This decision has two effects. (a) The bank receives a positive payoff R, if the loan

applicant who was granted a loan repays, or else, loses by L. Thus, the bank's expected utility for policy π_{θ} is $U(\theta, \mu) = \mathbb{E}_{x \sim \mathcal{N}(\mu, \sigma^2)} \left[\pi_{\theta}(x) (\mathbb{P}(repayment|x)R - (1 - \mathbb{P}(repayment|x))L) \right]$. (b) Since the amount of capital both the applicant and the population influence the credit score, we model that the change in the population mean μ_{t+1} depends on the bank's policy, via a grant rate $\mathbb{E}_{x \sim \mathcal{N}(\mu_t, \sigma_t^2)} \left[\pi_{\theta}(x) \right]$. Specifically, $\mu_{t+1} = (1 - \beta)\mu_t + \beta f\left(\mathbb{E}_{x \sim \mathcal{N}(\mu_t, \sigma_t^2)} \left[\pi_{\theta}(x) \right] \right)$, where $\beta \in [0, 1]$ is the performative strength and $f: \mathbb{R} \to [-M, M]$. Now, if one ignores the performative nature of this decision making problem, and try to find out the optimal with respect to a static credit distribution, it obtains $\theta^{ERM} \triangleq \arg \max_{\theta} U(\theta, \mu_0)$. In contrast, if it considers performativity, it obtains $\theta^{Perf} \triangleq \arg \max_{\theta} U(\theta, \mu^*(\theta))$. In Figure 1, we show that the average reward obtained by both the solutions are significantly different. This demonstrates why performativity is a common phenomenon across algorithmic decision making problems, and how it changes the resulting optimal solution. Further details are in Appendix B.

These problem scenarios have motivated the study of performative RL. Though Bell et al. (2021) were the first to propose a setting where the transition and reward of an underlying MDP depend non-deterministically on the deployed policy, Mandal et al. (2023) formally introduced *Performative RL*, and its solution concepts, i.e., performatively stable and optimal policies. Performative stable policies do not get affected or changed due to distribution shifts after deployment. Performatively optimal policies yield the highest expected return once deployed in the performative RL environment. Mandal et al. (2023) proposed direct optimization and ascent based techniques that attains performative stability upon repeated retraining. Extending this work, Rank et al. (2024) and Mandal & Radanovic (2024) manage to solve the same problem with delayed retraining for gradually shifting and linear MDPs. However, there exists no algorithm yet in performative RL that provably converges to the performative optimal policy.

As we know from the RL literature, the Policy Gradient (PG) type of algorithms that treats policy as a parametric function and updates the parameters through gradient ascent algorithms are efficient and scalable (Williams, 1992; Sutton et al., 1999; Kakade, 2001). Some examples of successful and popular policy gradient methods include TRPO (Schulman et al., 2015), PPO (Schulman et al., 2017), DDPG (Silver et al., 2014), SAC (Mnih et al., 2016), which are widely used in modern RL applications. Recent theoretical advances also establish finite-sample convergence guarantees and complexity analyses (Agarwal et al., 2021; Yuan et al., 2022) of PG algorithms. Motivated by the simplicity and universality of the PG algorithms, we ask these two questions in the context of performative RL:

- 1. How to design PG-type algorithms for performative RL environments to achieve optimality?
 2. What are the minimal conditions under which PG-type algorithms converge to the performatively optimal policy?
- **Our contributions** address these questions affirmatively, and showcases the difference of optimality-seeking and stability-seeking algorithms in performative RL.
- **I. Algorithm Design:** We propose the first Performative Policy Gradient algorithm, PePG, for performative RL environments. Specifically, we extend the classical vanilla PG and entropy-regularised PG algorithms to Performative RL settings. Though the general algorithm design stays same, we derive a performative policy gradient theorem that shows, evaluation of the gradient involves two novel gradient terms in performative RL (a) the expected gradient of reward, and (b) the expected gradient of log-transition probabilities times its impact on the expected cumulative return. We leverage this theorem to propose an estimator of the performative policy gradient under any differentiable parametrisation.
- **II. Convergence to Performative Optimality.** We further analyse PePG (with and without entropy regularisation) for softmax policies, and softmax Performative Markov Decision Processes (PeMDPs), i.e. the MDPs with softmax transition probabilities and linear rewards with respect to the parameters of the softmax policy. We provide a minimal recipe to prove convergence of PePG using (a) smoothness of the performative value function, and (b) approximate gradient domination lemma for performative policy gradients. This allows us to show that PePG converges to an ϵ -ball around performative optimal policy in $\Omega\left(\frac{|\mathcal{S}||\mathcal{A}|^2}{\epsilon^2(1-\gamma)}\right)$ iterations, where $|\mathcal{S}|$ and $|\mathcal{A}|$ are the number of states and actions, respectively.
- Specifically, Mandal et al. (2023) frames the question of using policy gradient to find stable policies as an open problem. The authors further contemplate, as PG functions in the policy space, whether it is possible to converge towards a stable policy. In this paper, we affirmatively solve an extension to this open problem for tabular softmax PeMDPs with softmax policies.
- III. Stability- vs. Optimality-seeking Algorithms in Performative RL. We further theoretically and numerically contrast the performances of stability-seeking and optimality-seeking algorithms. Theoretically, we derive the performative performance difference lemma that distinguished the effect of policy update in these two types of algorithms. Numerically, we compare the performances of PePG with the state-of-the-art MDRR (Mixed Delayed Repeated Retraining (Rank et al., 2024)) algorithm for finding performatively stable policies in the multi-agent environment proposed by (Mandal et al., 2023). We show that PePG yields significantly higher values functions than MDRR, while MDRR achieves either similar or lower distance from stable state-action distribution than PePG.

PRELIMINARIES: FROM RL TO PERFORMATIVE RL

Now, we formalise the RL and performative RL problems, and provide the basics of policy gradient algorithms in RL.

RL: Infinite-horizon Discounted MDPs

In RL, we mostly study Markov Decision Processes (MDPs) defined via the tuple (S, A, P, r, γ) , where $S \subseteq \mathbb{R}^d$ is the state space and $A \subseteq \mathbb{R}^d$ is the action space. Both the spaces are assumed to be compact. At any time step $t \in \mathbb{N}$, an agent plays an action $a_t \in \mathcal{A}$ at a state $s_t \in \mathcal{S}$. It transits the MDP environment to a state s_{t+1} according to a transition kernel $\mathbf{P}(\cdot \mid s_t, a_t) \in \Delta(\mathcal{S})$. The agent further receives a reward $r(s_t, a_t) \in \mathbb{R}$ quantifying the goodness of taking action a_t at s_t . The strategy to take an action is represented by a stochastic map, called *policy*, i.e. $\pi: \mathcal{S} \to \Delta(\mathcal{A})$. Given an initial state distribution $\rho \in \Delta(S)$, the goal is to find the optimal policy π^* that maximises the expected discounted sum of rewards, i.e., the value function: $V_{\pi}(\rho) \triangleq \mathbb{E}_{s_0 \sim \rho, s_{t+1} \sim \mathbb{P}(\cdot | s_t, \pi(s_t))} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, \pi(s_t)) \right]$, where $\gamma \in (0, 1)$ is called the discount factor. γ indicates how much a previous reward matters in the next step, and bounds the effective horizon of a policy to $\frac{1}{1-\gamma}$.

Algorithm 1 Vanilla Policy Gradient

- 1: **Input:** Learning rate $\eta > 0$.
- 2: **Initialize:** Policy parameter $\theta_0(s, a) \forall s \in \mathcal{S}, a \in \mathcal{A}$.
- 3: **for** t = 1 to T **do**
- Estimate the gradient $\nabla_{\boldsymbol{\theta}}V^{\boldsymbol{\pi}}(\boldsymbol{\rho})\mid_{\boldsymbol{\theta}=\boldsymbol{\theta}_{t}}$ **Gradient ascent step:** $\boldsymbol{\theta}_{t+1}\leftarrow\boldsymbol{\theta}_{t}+\eta\nabla_{\boldsymbol{\theta}}V^{\boldsymbol{\pi}}(\boldsymbol{\rho})\mid_{\boldsymbol{\theta}=\boldsymbol{\theta}_{t}}$
- 6: end for

108

109 110

111 112

113 114

115

116

117

118

119

120 121 122

123

124

125

126

127

128

129

130

131

132

133

134

135

136 137

142

143

144 145

146 147

148

149

150

151

152 153 154

155 156

157

158

159 160

161

Policy Gradient (PG) Algorithms. PG-type algorithms maximise the value function by directly optimising the policy through a gradient over value function (Williams, 1992). To compute the gradient, we choose a parametric family of policies π_{θ} for some $\theta \in \mathbb{R}^d$ (e.g. direct (Agarwal et al., 2021; Wang & Zou, 2022), softmax (Agarwal et al., 2021; Mei et al., 2020), Gaussian (Ciosek & Whiteson, 2020; Ghavamzadeh & Engel, 2006)). Specifically, vanilla PG (Algorithm 1), performs a gradient ascent on the policy parameter at each step $t \in \mathbb{N}$. As the goal is to maximise

 $V^{\pi}(\rho)$, we update θ towards $\nabla_{\theta}V^{\pi}(\rho)$, which is the direction improving the value $V^{\pi}(\rho)$ with a fixed learning rate $\eta > 0$. For vanilla PG, the policy gradient takes the convenient form leading to estimators computable only with policy rollouts.

Theorem 1 (Policy Gradient Theorem (Sutton et al., 1999)). Fix a differentiable paramterisation $\theta \mapsto \pi_{\theta}(a \mid s)$ and an initial distribution ρ . Let us define the Q-value function $Q^{\pi_{\theta}}(s, a) \triangleq \mathbb{E}_{s_{t+1} \sim \mathbf{P}_{\pi}(\cdot | s_t, \pi(s_t))} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, \pi(s_t)) \mid s_0 = s, a_0 = a \right]$, and advantage function $A^{\pi_{\theta}}(s, a) \triangleq Q^{\pi_{\theta}}(s, a) - V^{\pi_{\theta}}(s)$. Then,

$$\nabla_{\boldsymbol{\theta}} V^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\boldsymbol{\rho}) = \frac{1}{1 - \gamma} \mathbb{E}_{\boldsymbol{\tau} \sim \mathbb{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \left[\sum_{t=0}^{\infty} \gamma^t Q^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s, a) \nabla_{\boldsymbol{\theta}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}}(a \mid s) \right] = \mathbb{E}_{\boldsymbol{\tau} \sim \mathbb{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \left[\sum_{t=0}^{\infty} \gamma^t A^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s, a) \nabla_{\boldsymbol{\theta}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}}(a \mid s) \right].$$

Since the value function is not concave in the policy parameters, achieving optimality with PG has been a challenge. But practical scalability and efficiency of these algorithms has motivated a long-line of work to understand the minimum conditions and parametric forms of policies leading to convergence to the optimal policy (Agarwal et al., 2021; Mei et al., 2020; Wang & Zou, 2022; Yuan et al., 2022). Our work extends these algorithmic techniques and theoretical insights to performative RL.

2.2 Performative RL: Infinite-horizon Discounted PeMDPs

Given a policy set $\pi \in \Pi$, we denote the Performative Markov Decision Process (PeMDP) is defined as the set of MDPs $\{\mathcal{M}(\boldsymbol{\pi}) \mid \boldsymbol{\pi} \in \Pi\}$, where each MDP is a tuple $\mathcal{M}(\boldsymbol{\pi}) \triangleq (\mathcal{S}, \mathcal{A}, \mathbf{P}_{\boldsymbol{\pi}}, r_{\boldsymbol{\pi}}, \gamma)$. Note, that the transition kernel and rewards distribution are no more invariant with respect to the policy. They shift with the deployed policy $\pi \in \Delta(A)$ (Mandal et al., 2023; Mandal & Radanovic, 2024). In this setting, the probability of generating a trajectory $\tau_{\pi} \triangleq (s_t, a_t)_{t=0}^{\infty}$ under policy π with underlying MDP $\mathcal{M}(\pi')$ is given by $\mathbb{P}^{\pi}_{\pi'}(\tau \mid \rho) \triangleq \rho(s_0) \prod_{t=0}^{\infty} \pi(a_t \mid s_t) \mathbb{P}_{\pi'}(s_{t+1} \mid s_t, a_t)$, where $\rho \in \Delta(\mathcal{S})$ is the initial state distribution. Furthermore, the state-action occupancy measure for deployed policy π and environment-inducing policy π' is defined as $d^{\pi}_{\pi',\rho} \triangleq \frac{1}{1-\gamma} \mathbb{E}_{\tau \sim \mathbb{P}^{\pi}_{\pi'}} [\sum_{t=0}^{\infty} \gamma^t \mathbb{1}(s_t = s, a_t = a) \mid s_0 \sim \rho]$. Now, we are ready to define the performative expected return, referred as the performative value function that we aim to maximise while solving PeMDP.

Definition 1 (Performative Value Function). Given a policy $\pi \in \Pi$ and an initial state distribution $\rho \in \Delta(S)$, the performative value function $V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\boldsymbol{\rho})$ is

$$V_{\pi}^{\pi}(\boldsymbol{\rho}) \triangleq \mathbb{E}_{\tau \sim \mathbb{P}_{\pi}^{\pi}} \left[\sum_{t=0}^{\infty} \gamma^{t} r_{\pi}(s_{t}, \pi(s_{t})) \mid s_{0} \sim \boldsymbol{\rho} \right]. \tag{1}$$

¹Hereafter, for relevant quantities, π in superscript denotes the deployed policy, and π' in the subscript denoted the environmentinducing, i.e. the policy inducing the transition kernel and reward function that the algorithm interacts with.

Equation (1) gives the total expected return that captures the performativity aspect in PeMDPs as the underlying dynamics changes with a deployed policy $\pi(\cdot \mid s)$. Note that, we can maximise performative value function in two ways: (i) considering π as both the environment-inducing policy and the policy the RL agent deploys, or (ii) deploying π to fix it as the environment-inducing policy and agent plays another policy π' . At this vantage point, let us introduce the notion of optimality and stability of policies in PeMDPs (Mandal et al., 2023).

Definition 2 (Performative Optimality). A policy π_o^* is performatively optimal if it maximizes the performative value function.

$$\boldsymbol{\pi}_{o}^{\star} \in \operatorname*{arg\,max}_{\boldsymbol{\pi}} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\boldsymbol{\rho}). \tag{2}$$

Thus, if we play the policy π in the environment induced by policy π to maximise the expected return, we land on the performatively optimal policy.

Definition 3 (Performative Stability). A policy π_s^* is performatively stable if there is no gain in performative value function due to deploying any other policy than π_s^* in the environment induced by π_s^* .

$$\pi_s^* \in \arg\max_{\pi \in \Delta(\mathcal{A})} V_{\pi_s^*}^{\pi}(\rho). \tag{3}$$

As noted by Mandal et al. (2023), a performatively optimal policy may not be performatively stable, i.e., π_o^* may not be optimal for a changed underlying environment $\mathcal{M}(\pi_o^*)$, when it is deployed. Also, in general, the performative value function of π_o^* might be equal to or higher than that of π_s^* . In this paper, we design PG algorithms computing the performative optimal policy for a given set of MDPs, and reinstate their differences with performatively stable policies.

The existing literature on PeMDPs (Mandal et al., 2023; Mandal & Radanovic, 2024; Rank et al., 2024; Pollatos et al., 2025; Chen et al., 2024) focused primarily on finding a performatively stable policy, i.e. a π_s^* according to Definition 3. In practice, while the notion of stable policies matters for very specific applications, a stable policy may not always suffice. But they might show large sub-optimality gaps, which are often not desired for real-life tasks. We fill up this gap in literature and propose the first provably converging and computationally efficient PG algorithm for PeMDPs. Later on, we also empirically show the deficiency of the existing stability finding algorithms if we aim for optimality (Section 5).

Entropy Regularised PeMDPs. Entropy regularisation has emerged as a simple but powerful technique in classical RL to design smooth and efficient algorithms with sufficient exploration. Thus, we study another variant of the performative value function that is regularised using discounted entropy (Mei et al., 2020; Neu et al., 2017; Liu et al., 2019; Zhao et al., 2019). In this setting, the original value function in Definition 1 is regularised using the discounted entropy $H_{\pi}(\rho) \triangleq \mathbb{E}_{\tau \sim \mathbb{P}_{\pi}^{\pi}} \left[-\sum_{t=0}^{\infty} \gamma^{t} \log \pi(a_{t} \mid s_{t}) \right]$. This is equivalent to maximising the expected reward with a shifted reward function $\tilde{r}_{\pi}(\pi(s_{t}), s_{t}) = r_{\pi}(\pi(s_{t}), s_{t}) - \lambda \log(\pi(a_{t} \mid s_{t}))$ for some $\lambda \geq 0$. \tilde{r}_{π} is referred as the "soft-reward" in MDP literature (Wang & Uchibe, 2024; Herman et al., 2016; Shi et al., 2019). This allows us to define the soft performative value function.

Definition 4 (Entropy Regularised (or Soft) Performative Value Function). Given a policy $\pi \in \Pi$, a starting state distribution $\rho \in \Delta(S)$, and a regularisation parameter $\lambda \geq 0$, the soft performative value function $V_{\pi}^{\pi}(\rho)$ is

$$\tilde{V}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\boldsymbol{\rho}) \triangleq \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \left(r_{\boldsymbol{\pi}}(s_{t}, \boldsymbol{\pi}(s_{t})) - \lambda \log \boldsymbol{\pi}(a_{t} \mid s_{t}) \right) \mid s_{0} \sim \boldsymbol{\rho} \right] = \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \tilde{r}_{\boldsymbol{\pi}}(s_{t}, \boldsymbol{\pi}(s_{t})) \mid s_{0} \sim \boldsymbol{\rho} \right] .$$
(4)

Since policies belong to the probability simplex, the entropy regularisation naturally lends to smoother and stable PG algorithms. Later, we show that the discounted entropy is a smooth function of the policy parameters for PeMDPs extending the optimization-wise benefits of entropy regularisation to PeMDPs. Additionally, using the notion of soft rewards, we can further define soft performatively optimal and stable policies for entropy regularised PeMDPs. Leveraging it, we unifiedly design PG algorithms for both the unregularised and the entropy regularised PeMDPs.

3 POLICY GRADIENT ALGORITHMS IN PERFORMATIVE RL

In this section, we first study the impact of policy updates in PeMDPs. Then, we leverage it to derive the performative policy gradient theorem and design Performative Policy Gradient (PePG) algorithm for any differentiable parametric policy class.

3.1 IMPACT OF POLICY UPDATES ON PEMDPS

Performance difference lemma has been central in RL to understand the impact of changing policies in terms of value functions (Kakade & Langford, 2002a). It has been also central to analysing and developing PG-type methods (Agarwal et al., 2021; Silver et al., 2014; Kallel et al., 2024). But the existing versions of performance difference cannot handle performativity. Here, we derive the performative version of the performance difference lemma that quantifies the shift in the performative value function due to change the deployed and environment-inducing policies.

Lemma 1 (Performative Performance Difference Lemma). The difference in performative value functions induced by π and $\pi' \in \Pi$ while starting from the initial state distribution ρ is

$$V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\boldsymbol{\rho}) - V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}'}(\boldsymbol{\rho}) = \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}',\rho}^{\boldsymbol{\pi}}} [A_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}'}(s,a)]$$

$$+ \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi},\rho}^{\boldsymbol{\pi}}} [(r_{\boldsymbol{\pi}}(s,a) - r_{\boldsymbol{\pi}'}(s,a)) + \gamma (\mathbf{P}_{\boldsymbol{\pi}}(\cdot|s,a) - \mathbf{P}_{\boldsymbol{\pi}'}(\cdot|s,a))^{\top} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot)].$$
 (5)

where $A_{\pi'}^{\pi'}(s,a) \triangleq Q_{\pi'}^{\pi'}(s,a) - V_{\pi'}^{\pi'}(s)$ is the performative advantage function for any state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$.

The crux of the proof is decomposing the performative value through environment-inducing and deployed policies

$$V_{\pi}^{\pi}(s_0) - V_{\pi'}^{\pi'}(s_0) = \underbrace{V_{\pi}^{\pi}(s_0) - V_{\pi'}^{\pi}(s_0)}_{\text{performative shift term}} + \underbrace{V_{\pi'}^{\pi}(s_0) - V_{\pi'}^{\pi'}(s_0)}_{\text{performance difference term}}.$$

(1) Connection to Classical RL. In classical RL, the performance difference lemma yields $V^{\pi}(\rho) - V^{\pi'}(\rho) = \frac{1}{1-\gamma}\mathbb{E}_{(s,a)\sim d^{\pi}_{\rho}}[A^{\pi'}(s,a)]$. The first term in Lemma 1 is equivalent to the classical result in the environment induced by π' . But due to environment shift, two more terms appear in the performative performance difference incorporating the impacts of reward shifts and transition shifts. (2) Connection to Performative Stability. If we ignore the reward and transition shift terms, the performance difference term $V^{\pi}_{\pi'}(s_0) - V^{\pi'}_{\pi'}(s_0)$ quantifies the impact of changing the deployed policy from π' to π in an environment induced by π' . Thus, a stability seeking algorithm would like to minimise this term, while an optimality seeking algorithm has to incorporate all of the terms.

Now, we ask: how much do the new environment shift terms change the performative performance difference?

For simplicity, we focus on the commonly studied PeMDPs with bounded rewards and gradually shifting environments, i.e. the ones with Lipschitz transitions and rewards with respect to the deployed policies (Rank et al., 2024).

Assumption 1 (Bounded reward). We assume that the rewards are bounded in $[-R_{\max}, R_{\max}]$.

This is the only assumption needed through the paper and is standard in MDP literature (Mei et al., 2020; Li & Yang, 2023). **Lemma 2** (Bounding Performance Difference for Gradually Shifting Environments). *Let us assume that both rewards and transitions are Lipschitz functions of policy, i.e.* $||r_{\pi} - r_{\pi'}|| \le L_r ||\pi - \pi'||$ and $||\mathbf{P}_{\pi} - \mathbf{P}_{\pi'}|| \le L_{\mathbf{P}} ||\pi - \pi'||$,

wards and transitions are Lipschitz functions of policy, i.e. $||r_{\pi} - r_{\pi'}|| \le L_r ||\pi - \pi'||$ and $||\mathbf{P}_{\pi} - \mathbf{P}_{\pi'}|| \le L_{\mathbf{P}} ||\pi - \pi'||$ for some $L_r, L_{\mathbf{P}} \ge 0$. Then, under Assumption 1, the performative shift in the sub-optimality gap of a policy π_{θ} satisfies

$$\left|V_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\boldsymbol{\rho}) - V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\boldsymbol{\rho}) - \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim d_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{\star},\boldsymbol{\rho}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{\star}}[A_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a)]\right| \leq \frac{2\sqrt{2}}{1 - \gamma} (L_{r} + \frac{\gamma}{1 - \gamma} L_{\mathbf{P}} R_{\max}) \mathbb{E}_{s_{0} \sim \boldsymbol{\rho}} D_{\mathbf{H}} \left(\boldsymbol{\pi}_{o}^{\star}(\cdot|s_{0}) \|\boldsymbol{\pi}_{\boldsymbol{\theta}}(\cdot|s_{0})\right) . \tag{6}$$

where $D_{\rm H}(\mathbf{x}||\mathbf{y})$ denotes the Hellinger distance between \mathbf{x} and \mathbf{y} .

Implication. Lemma 2 shows novel characterisation of the extra cost we have to pay to adapt to performativity of the environment in terms of Hellinger distance between the true performatively optimal policy π_o^* and any other parametrised policy π_θ . This implies that the order of difference between the optimal performative value function and that of any stability-seeking algorithm is $\Theta(\frac{1}{1-\gamma})$. This significantly improves the known order of sub-optimality achieved by existing algorithms. Specifically, Mandal et al. (2023) show that using repeated policy optimisation algorithms converges to a suboptimality gap $\mathcal{O}\left(\max\{\frac{S^{5/3}A^{1/3}\epsilon^{2/3}}{(1-\gamma)^{14/3}},\frac{\epsilon S}{(1-\gamma)^4}\}\right)$. Thus, we see an opportunity to improve on the existing works and design algorithms that can achieve suboptimality gap of order $\Theta(\frac{1}{1-\gamma})$.

3.2 ALGORITHM DESIGN: PERFORMATIVE POLICY GRADIENT (PePG)

To achieve performative optimality, the goal is to maximise value function at the end of learning process. Gradient ascent is a standard first-order optimisation method to find maxima of a function. Similar to Algorithm 1, the crux of performative policy gradient method lies in the ascent step:

$$\boldsymbol{\theta}_{t+1} \leftarrow \begin{cases} \boldsymbol{\theta}_t + \eta_t \nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\tau) \mid_{\boldsymbol{\theta} = \boldsymbol{\theta}_t} & \text{, for unregularised objective} \\ \boldsymbol{\theta}_t + \eta_t \nabla_{\boldsymbol{\theta}} \tilde{V}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\tau) \mid_{\boldsymbol{\theta} = \boldsymbol{\theta}_t} & \text{, for Entropy-regularised objective.} \end{cases}$$
(7)

Given this ascent step, we have to evaluate the gradient at each time step from the rollouts of the present policy. In classical PG, the policy gradient theorem serves this purpose (Williams, 1992; Sutton et al., 1999; Silver et al., 2014). Thus, we derive the performative counterpart of the classic policy gradient theorem.

Algorithm 2 PePG: Performative Policy Gradient

- 1: **Input:** Transition Feature Map $\psi(s) \forall s \in \mathcal{S}, \xi \in [-R_{\max}, R_{\max}]$ and discount factor γ .
- 2: **Initialize:** Initial policy parameters θ_0 , initial value function parameters ϕ_0
- 3: **for** $k = 1, 2, \dots$ **do**

271

272

273

274

276

277

278

279

280

281

283 284 285

286 287 288

289

290 291 292

294 295 296

297 298

299 300 301

302

304

305

306 307

308

315 316

317

318

319

321

323

- Collect trajectories: $\mathcal{D}_k = \{\tau_i\}_{i=1}^I$, where each $\tau_i \triangleq \{(s_{i,t}, a_{i,t}, s_{i,t+1}, r_{i,t})\}_{t=0}^{T-1}$ by playing $\boldsymbol{\pi}_{\boldsymbol{\theta}_k} = \boldsymbol{\pi}(\boldsymbol{\theta}_k)$ Compute returns $R_k \triangleq \{R_{k,i}\}_{i=1}^I$, where $R_{k,i} = \{R_{k,i,t}\}_{t=0}^{T-1}$ 4:
- 5:
- Compute advantage estimates $\hat{A}_k(\tau_i)$ using value function $\hat{V}_{\phi_k}(\tau_i)$ for each $\tau_i \in \mathcal{D}_k$ (estimate of $V_{\pi_{\theta_k}}^{\pi_{\theta_k}}(\tau_i)$ obtained from fitted value network with parameters ϕ_k)
- 7: **Gradient estimation:** Estimate policy gradient using (10)
- **Gradient ascent step:** Update policy parameters using (7)
- Fit value function $V_{\phi_{k+1}}$:

$$\phi_{k+1} \leftarrow \arg\min_{\phi} \frac{1}{I \cdot T} \sum_{i=1}^{I} \sum_{t=0}^{T-1} \left(\hat{V}_{\phi_k}(s_t \in \tau_i) - R_{k,i,t} \right)^2$$

10: **end for**

Theorem 2 (Performative Policy Gradient Theorem). The gradient of the performative value function w.r.t θ is as follows: (a) For unregularised objective,

$$\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}^{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(\tau) = \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}^{\boldsymbol{\pi_{\boldsymbol{\theta}}}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \left(A_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}^{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(s_{t}, a_{t}) \left(\nabla_{\boldsymbol{\theta}} \log \boldsymbol{\pi_{\boldsymbol{\theta}}}(a_{t} \mid s_{t}) + \nabla_{\boldsymbol{\theta}} \log P_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(s_{t+1} | s_{t}, a_{t}) \right) + \nabla_{\boldsymbol{\theta}} r_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(s_{t}, a_{t}) \right) \right], \quad (8)$$

(b) For entropy-regularized objective, we define $\tilde{A}^{\pi_{\theta}}_{\pi_{\theta}}(s,a) = Q^{\pi_{\theta}}_{\pi_{\theta}}(s,a) - V^{\pi_{\theta}}_{\pi_{\theta}}(s) - \lambda \log \pi_{\theta}(a|s)$, and get

$$\nabla_{\boldsymbol{\theta}} \tilde{V}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(\tau) = \mathbb{E}_{\boldsymbol{\tau} \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \left(\tilde{A}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \left(\nabla_{\boldsymbol{\theta}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}}(a_{t} \mid s_{t}) + \nabla_{\boldsymbol{\theta}} \log P_{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t+1} | s_{t}, a_{t}) \right) + \nabla_{\boldsymbol{\theta}} \tilde{r}_{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t} | \boldsymbol{\theta}) \right) \right]. \quad (9)$$

PePG: To elaborate on the design of PePG (Algorithm 2), we focus only on the REINFORCE update and softmax policy parametrisation. With the appropriate parameter choices, and initialisation of the policy parameter θ and value function parameter ϕ , for each episode $k=1,2,\ldots$, PePG collects I trajectories to calculate return R^i and estimates advantage function \hat{A}_k (Line 4-6). For a particular trajectory τ_i , the estimated advantage for a given state-action is $A_{\pi\theta_k}^{\pi\theta_k}(s_t^i, a_t^i) =$ $R_{t,k}^{i} - V_{\phi_k}(s_t^{i})$, where $R^{i} = \sum_{t=0}^{T-1} \gamma^{t} r_{\pi_{\theta_k}}(s_t^{i}, a_t^{i})$.

Gradient Estimation (Line 7). With the necessary estimates in hand for all the collected I trajectories, PePG computes average gradient estimate over all the trajectories using

$$\widehat{\nabla_{\boldsymbol{\theta}_{k}} V_{\boldsymbol{\pi}_{\boldsymbol{\theta}_{k}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}_{k}}}} = \frac{1}{I} \sum_{i=1}^{I} \sum_{t=0}^{T} \gamma^{t} (A_{\boldsymbol{\pi}_{\boldsymbol{\theta}_{k}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}_{k}}}(s_{t}^{i}, a_{t}^{i}) \left(\nabla_{\boldsymbol{\theta}_{k}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}_{k}}(a_{t}^{i} \mid s_{t}^{i}) + \nabla_{\boldsymbol{\theta}_{k}} \log P_{\boldsymbol{\pi}_{\boldsymbol{\theta}_{k}}}(s_{t+1}^{i} \mid s_{t}^{i}, a_{t}^{i}) \right) + \nabla_{\boldsymbol{\theta}_{k}} r_{\boldsymbol{\pi}_{\boldsymbol{\theta}_{k}}}(s_{t}^{i}, a_{t}^{i} \mid \boldsymbol{\theta}_{k}))$$

$$(10)$$

where all the individual gradients $\nabla_{\theta_k} \log P_{\pi_{\theta_k}}$, $\nabla_{\theta_k} r_{\pi_{\theta_k}}$ and $\nabla_{\theta_k} \log \pi_{\theta_k}$ have the closed form expressions for softmax parametrisation according to Equation (26). Further, in Line 8, PePG updates the policy parameter for the next episode using a gradient ascent step leveraging the estimated average gradient over all I trajectories. Specifically, we plug in $\nabla_{\theta_k} V_{\pi_{\theta_k}}^{\pi_{\theta_k}}$ to both the unregularised and entropy-regularised update rules are given in Equation (7). For the next episode, we again run a regression to update the value network plugging in the current estimates and resume the learning process further.

CONVERGENCE ANALYSIS OF PePG: SOFTMAX POLICIES AND SOFTMAX PEMDPS

For rigorous theoretical analysis of PePG, we restrict ourselves to softmax policy class, and softmax PeMDPs. We define the softmax PeMDPs as the ones having softmax transition kernesls with feature map $\psi(\cdot): \mathcal{S} \to \mathbb{R}$, and linear reward functions

with respect to the policy parameters, for all state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$. Specifically, the class of softmax PeMDPs is $\{\mathcal{M}(\boldsymbol{\theta}) = \mathcal{M}(\boldsymbol{\pi}_{\boldsymbol{\theta}}) \mid \boldsymbol{\theta} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}\}$ such that

$$\boldsymbol{\pi}_{\boldsymbol{\theta}}(a|s) = \frac{e^{\boldsymbol{\theta}_{s,a}}}{\sum_{a'} e^{\boldsymbol{\theta}_{s,a'}}}, \, \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s'|s,a) = \frac{e^{\boldsymbol{\theta}_{s,a}\psi(s')}}{\sum_{s''} e^{\boldsymbol{\theta}_{s,a}\psi(s'')}}, \, r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a) = \mathcal{P}_{[-R_{\max},R_{\max}]}[\boldsymbol{\xi}\boldsymbol{\theta}_{s,a}],$$
(11)

where ψ is non-negative and upper bounded by $\psi_{\max} > 0$, and $\xi \in [0, R_{\max}]$ to align with Assumption 1.

Thus, we derive the derivatives of policy, transitions, and rewards as

$$\frac{\partial}{\partial \boldsymbol{\theta}_{s',a'}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}}(a|s) = \mathbb{1}[s=s', a=a'] - \boldsymbol{\pi}_{\boldsymbol{\theta}}(a'|s)\mathbb{1}[s=s'],$$

$$\frac{\partial}{\partial \boldsymbol{\theta}_{s',a'}} \log \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s''|s, a) = \psi(s'')\mathbb{1}[s=s', a=a'] \left(1 - \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s''|s, a)\right), \frac{\partial}{\partial \boldsymbol{\theta}_{s',a'}} r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s, a) = \xi \mathbb{1}[s=s', a=a'].$$
(12)

Given the derivatives, we can now readily estimate the policy gradient and deploy PePG for softmax PeMDPs.

Convergence Analysis: Challenges and Three Step Analysis. The main challenge to prove convergence of PePG is that the performative value function is not concave in the paramterisation θ , in general, and also in softmax PeMDPs. The similar issue occurs while proving convergence of PG-type algorithms in classical RL, which has been overcome by leveraging smoothness properties of the value functions and by deriving the local Polyak-Lojasiewicz (PL)-type conditions, known as *gradient domination*, with respect to the policy paramterisation. Leveraging these insights, we devise a three step convergence analysis for PePG.

Step 1: Smoothness of Performative Value Functions. First, we prove that the unregularised performative value function is $\mathcal{O}(\frac{|\mathcal{A}|}{(1-\gamma)^2})$ smooth. As we show that the entropy is also a smooth function for softmax PeMDPs, then under proper choice of the regaularisation parameter, i.e., $\lambda = \frac{1-\gamma}{1+2\log|\mathcal{A}|}$, entropy regularised performative value function is also $\mathcal{O}(\frac{|\mathcal{A}|}{(1-\gamma)^2})$ smooth. Since gradient ascent/descent methods can work well in smooth functions, we proceed thoroughly.

Step 2: Gradient Domination for Softmax PeMDPs. Now, the next step is to relate the performative performance difference with the performative policy gradient. This allows us to connect the per iteration improvement in the performative value function with the performative gradient descent at that step. These are known as PL-type inequalities. For non-concave objectives, PL inequalities guarantee convergence to global maxima by showing that the gradient of the objective at any parameter dominates the sub-optimality w.r.t. that parameter.

Lemma 3 (Performative Gradient Domination for Softmax PeMDPs). For PeMDPs defined in (11) and $C_{\psi} > -1$, we get

(a) For unregularised value function:

$$V_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\rho) - V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\rho) \leq \frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{(1 + C_{\psi})} \left\| \frac{d_{\boldsymbol{\pi}_{o}, \rho}^{\boldsymbol{\pi}_{o}}}{d_{\boldsymbol{\pi}_{\theta}, \nu}^{\boldsymbol{\pi}_{\theta}}} \right\|_{\infty} \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\nu)\|_{2} + \frac{R_{\max}}{1 - \gamma} (|\mathcal{A}| + \frac{\gamma}{1 - \gamma} \psi_{\max}).$$
 (13)

(b) For entropy-regularised value function, $\tilde{V}_{\pi_{\alpha}^{*}}^{\pi_{\alpha}^{*}}(\rho) - \tilde{V}_{\pi_{\theta}}^{\pi_{\theta}}(\rho) \leq$

$$\frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{(1+C_{\psi})} \left\| \frac{d_{\boldsymbol{\pi}\boldsymbol{\theta},\boldsymbol{\rho}}^{\boldsymbol{\pi}^{\diamond}}}{d_{\boldsymbol{\pi}\boldsymbol{\theta},\boldsymbol{\nu}}^{\boldsymbol{\pi}^{\diamond}}} \right\|_{\infty} \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}^{\diamond}}(\boldsymbol{\nu})\|_{2} + \frac{1}{1-\gamma} \left(R_{\max}|\mathcal{A}| + \frac{\gamma}{1-\gamma} \psi_{\max}(R_{\max} + \lambda \log|\mathcal{A}|) \right) + \frac{\lambda}{1-\gamma} (1+2\log|\mathcal{A}|)$$
(14)

Step 3: Iterative Application of Gradient Domination for Smooth Functions. Now, we can apply gradient domination along with the classic iterative convergence proof of gradient ascent for smooth functions. The intuition is that since the perstep sub-optimality is dominated by the gradient and the smooth functions are bounded by quadratic envelopes of parameters, applying gradient ascent iteratively would bring the sub-optimality down to small error level after enough iterations. We formalise this in Theorem 3.

Theorem 3 (Convergence of PePG in softmax PeMDPs). Let $Cov \triangleq \max_{\theta, \nu} \left\| \frac{d_{\pi_{\theta}, \nu}^{\pi_{\theta}}}{d_{\pi_{\theta}, \nu}^{\pi_{\theta}}} \right\|_{\infty}$. The gradient ascent algorithm on

 $V_{\pi_{\theta}}^{\pi_{\theta}}(\rho)$ (Equation (7)) with step size $\eta = \Omega((1-\gamma)^2/|\mathcal{A}|)$ satisfies, for all distributions $\rho \in \Delta(\mathcal{S})$,

- (a) For unregularised case, $\min_{t < T} \left\{ V_{\boldsymbol{\pi}_{\boldsymbol{\theta}_{t}}^{\star}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}_{t}}^{\star}}(\boldsymbol{\rho}) V_{\boldsymbol{\pi}_{\boldsymbol{\theta}_{t}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}_{t}}}(\boldsymbol{\rho}) \right\} \leq \epsilon \text{ when } T = \Omega\left(\frac{R_{\max}|\mathcal{S}||\mathcal{A}|^{2}}{\epsilon^{2}(1-\gamma)^{3}}\mathrm{Cov}^{2}\right), \text{ and } \epsilon = \Omega\left(\frac{1}{1-\gamma}\right).$
- (b) For entropy regularised case, if we set $\lambda = \frac{(1-\gamma)}{1+2\log|\mathcal{A}|}$, we get $\min_{t < T} \left\{ \tilde{V}_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\rho) \tilde{V}_{\boldsymbol{\pi}_{\theta}^{(t)}}^{\boldsymbol{\pi}_{\theta}^{(t)}}(\rho) \right\} \leq \epsilon$ when $T = \Omega\left(\frac{R_{\max}|\mathcal{S}||\mathcal{A}|^{2}}{\epsilon^{2}(1-\gamma)^{3}}\mathrm{Cov}^{2}\right)$, and $\epsilon = \Omega\left(\frac{1}{1-\gamma}\right)$.

Algorithms	Regulariser λ	Min. #samples	Environment
RPO FS (Mandal et al., 2023)	$\mathcal{O}\left(rac{\left \mathcal{S} ight +\gamma\left \mathcal{S} ight ^{5/2}}{(1-\epsilon)(1-\gamma)^4} ight)$	$\frac{ A S ^3(B+\sqrt{ A })^2}{\delta^4(1-\gamma)^6\lambda^2}\ln\left(\frac{i+1}{p}\right)$	Direct PeMDPs + quadratic-regul. on occupancy
MDRR (Rank et al., 2024)	$\mathcal{O}\left(rac{ \mathcal{S} +\gamma \mathcal{S} ^{5/2}}{(1-\epsilon)(1-\gamma)^4} ight)$	$\frac{ A S ^3(B+\sqrt{ A })^2}{\delta^4(1-\gamma)^6\lambda^2}\ln\left(\frac{i+1}{p}\right)$	$\label{eq:definition} \mbox{Direct PeMDPs} + \mbox{quadratic-regul. on occupancy}$
PePG (This paper)	$\frac{1-\gamma}{1+2\log(\mathcal{A})}$	$rac{ \mathcal{S} \mathcal{A} ^2}{\epsilon^2(1-\gamma)^3}$	$softmax\ PeMDPs + entropy\ regul.\ on\ policy$
PePG (This paper)	0	$rac{ \dot{\mathcal{S}} \mathcal{A} ^2}{\epsilon^2(1-\gamma)^3}$	unregularised softmax PeMDPs

Table 1: Comparison of theoretical performance of SOTA stability-seeking algorithms against PePG.

Implications. (1) We observe that PePG converges to an ϵ -optimal policy in $\frac{|\mathcal{S}||\mathcal{A}|^2}{\epsilon^2(1-\gamma)^3}$ iterations. This reduces the sample complexity required for the existing stability-seeking algorithms by at least an order $\frac{|\mathcal{S}|^2}{(1-\gamma)^3}$, and shows efficiency of using PePG than the algorithms directly optimising the occupancy measures. (2) Additionally, the regularisation parameters needed for the existing algorithms are pretty big and bigger than $\frac{|\mathcal{S}|}{(1-\gamma)^4}$. This is counter-intuitive and does not match the experimental observations. Here, we prove that setting the regularisation parameter to $\frac{1-\gamma}{1+2\log|\mathcal{A}|}$ suffices for proving convergence to optimality. (3) The minimum number of samples required to achieve convergence is proportional to the square of coverage for the softmax PeMDP. This is a ubiquitous quantity dictating convergence of PG-methods in classical RL (Agarwal et al., 2021; Mei et al., 2020), and retraining methods in performative RL (Mandal et al., 2023; Rank et al., 2024).

5 EXPERIMENTAL ANALYSIS

In this section, we empirically compare the performance of PePG in the performative reinforcement learning setting and analyse its behaviour against the state-of-the-art stability-finding methods. ²

Performative RL Environment. We evaluate PePG in the Gridworld test-bed (Mandal et al., 2023), which has become a standard benchmark in performative RL. This environment consists of a grid where two agents A_1 (the principal) and A_2 (the follower), jointly control an actor navigating from start positions (S) to the goal (G) while avoiding hazards. The environment dynamics are as follows: Agent A_1 proposes a control policy for the actor by selecting one of four directional actions. Agent A_2 can either accept this action (not intervene) or override it with its own directional choice. This creates a performative environment for A_1 , as its effective policy outcomes depend on A_2 's responses to its deployed strategy.

The cost structure follows: visiting blank cells (S) incurs penalty of -0.01, goal cells (F) cost -0.02, hazard cells (H) impose a severe penalty of -0.5, and any intervention by A_2 results in an additional cost of -0.05 for the intervening agent. The response model also follows that of Mandal et al. (2023), i.e., the agent A_2 responds to A_1 's policy using a Boltzmann softmax operator. Given A_1 's current policy π_1 , we compute the optimal Q-function $Q^{*|\pi_1}$ for each follower agent A_j relative to a perturbed version of the grid world, where each cell types matches A_1 's environment with probability 0.7. We then define an average Q-function over the follower agents and determine the collective response policy via Boltzmann softmax $Q^{*|\pi_1}(s,a) = \frac{1}{n} \sum_{j=2}^{n+1} Q_j^{*|\pi_1}(s,a), \pi_2(a|s) = \frac{\exp(\beta \cdot Q^{*|\pi_1}(s,a))}{\sum_{a'} \exp(\beta \cdot Q^{*|\pi_1}(s,a'))}$.

It is important to note that our experimental setup deliberately uses the immediate response model from the original performative RL framework, rather than the gradually shifting environment introduced by Rank et al. (2024) that assumes slow shifts in the environment. Our choice to use the immediate response model presents a more challenging performative setting where the environment responds instantaneously to policy changes. This allows us to demonstrate that unlike MDRR (Rank et al., 2024), PePG can handle the fundamental performative challenge without requiring environmental assumptions that artificially slows down the feedback loop, thereby highlighting the robustness of the proposed PePG approach.

Experimental Setup. We evaluate PePG (with and without entropy regularisation) alongside Mixed Delayed Repeated Retraining (MDRR), which represents the current state-of-the-art in performative reinforcement learning under gradually shifting environments (Rank et al., 2024), and Repeated Policy Optimization with Finite Samples (RPO FS). MDRR has demonstrated significant improvements over traditional repeated retraining methods, by leveraging historical data from multiple deployments, while RPO FS is included as the baseline method from (Mandal et al., 2023) for direct comparison with the original performative RL approach.

²Anonymous code repository of PePG implementation is Link. Further ablation studies w.r.t. hyperparameters are in Appendix H.

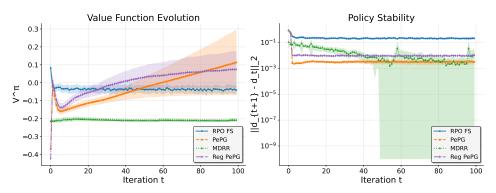


Figure 2: Comparison of evolution in expected average return (both regularised and unregularised) and stability of PePG with SOTA stability-achieving methods. Each algorithm is run for 20 random seeds and 100 iterations.

All experiments use a 8×8 grid with $\gamma=0.9$, exploration parameter $\epsilon=0.5$ for initial policy construction, one follower agent A_2 , and 100 trajectory samples per iteration. The algorithms share common parameters of T=100 iterations. For regularization, RPO FS and MDRR use $\lambda=0.1$ from their original experiments, while entropy-regularized PePG uses $\lambda=2.0$ (ablation studies for this choice are provided in the appendix). PePG uses learning rate $\eta=0.1$, MDRR employs memory weight v=1.1 for historical data utilization, delayed round parameter k=3, and FTRL parameters N=B=10, while RPO FS follows the finite-sample optimization from Mandal et al.

Results and Observations. Our experimental evaluation across 100 iterations reveals fundamental differences between PePG and MDRR and RPO in the immediate response performative setting. We used shorter training compared to (Rank et al., 2024), as this time-frame sufficiently demonstrates RPO and MDRR's stability convergence and PePG's progression toward optimality.

I. Results: Optimality: The left panel reveals a clear performance hierarchy among the four methods. PePG achieves the highest value function performance, with standard PePG reaching approximately 0.1 and regularized PePG (Reg PePG) reaching 0.05, both showing consistent improvement from initial values around -0.15 and still progressing upward at the end of the 100 iteration window. This steady upward progression highlights PePG's effectiveness in discovering better performative equilibria rather than settling for the first stable solution encountered. RPO FS remains relatively stable around -0.05 throughout training, while MDRR stabilizes at the lowest performance level of approximately -0.2 and remains flat throughout training.

II. Results: Comparison of Optimality- and Stability-seeking Algorithms. The results expose a critical limitation of algorithms designed primarily for stability rather than optimality. MDRR successfully achieves its design goal, with the right panel showing decreasing toward zero in the stability metric $||d_{t+1} - d_t||_2$ (the L_2 distance between occupancy measures of consecutive policy iterations), indicating policy stabilization. However, this stability comes at the cost of solution quality, as MDRR becomes trapped in a suboptimal point. The method prioritised finding any stable point over finding an optimal solution. In contrast, both PePG variants exhibit higher policy variability as they actively explore for better solutions. RPO FS maintains moderate stability around 10^{-1} but with limited performance improvement.

6 DISCUSSIONS, LIMITATIONS, AND FUTURE WORKS

We study the problem of Performative Reinforcement learning in tabular MDPs (PeMDPs) using softmax parametrised policies with entropy-regularised objective function, where any action taken by the agent cause potential shift in the MDP's underlying reward and transition dynamics. We are the first to develop PG-type algorithm, PePG, that attains performatively optimality against the existing performative stability-seeking algorithms, affirmatively solving an extended open problem in (Mandal et al., 2023). We also derive the novel performative counterpart of classic Performance Difference Lemma and Policy Gradient Theorem that affirmatively captures this performative nature of the environment we act. We provide a sufficient conditions to prove that PePG converges to an ϵ -ball around performative optimal policy in $\Omega\left(\frac{|\mathcal{S}||\mathcal{A}|^2}{\epsilon^2(1-\gamma)}\right)$ iterations.

As we develop a PG-type algorithm, it will be interesting to see how much can we reduce variance (Wu et al., 2018; Papini et al., 2018) while achieving optimality. We are still in the tabular setting with finite set of state-actions. A potential future direction would be to scale PePG to continuous state-space with large number of state-actions.

REFERENCES

- Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98):1–76, 2021. 2, 3, 4, 8, 16
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv* preprint arXiv:2204.05862, 2022. 1
- Anas Barakat, John Lazarsfeld, Georgios Piliouras, and Antonios Varvitsiotis. Multi-agent online control with adversarial disturbances. *arXiv preprint arXiv:2506.18814*, 2025. 1, 16
- James Bell, Linda Linsefors, Caspar Oesterheld, and Joar Skalse. Reinforcement learning in newcomblike environments. *Advances in Neural Information Processing Systems*, 34:22146–22157, 2021. 2, 16
- Gavin Brown, Shlomi Hod, and Iden Kalemaj. Performative prediction in a stateful world. In *International conference on artificial intelligence and statistics*, pp. 6045–6061. PMLR, 2022. 16
- Songfu Cai, Fei Han, and Xuanyu Cao. Performative control for linear dynamical systems. *Advances in Neural Information Processing Systems*, 37:70617–70658, 2024. 16
- Allison JB Chaney, Brandon M Stewart, and Barbara E Engelhardt. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM conference on recommender systems*, pp. 224–232, 2018. 1
- Qianyi Chen, Ying Chen, and Bo Li. Practical performative policy learning with strategic agents. *arXiv preprint* arXiv:2412.01344, 2024. 4
- Kamil Ciosek and Shimon Whiteson. Expected policy gradients for reinforcement learning. *Journal of Machine Learning Research*, 21(52):1–51, 2020. 3
- Itay Eilat and Nir Rosenfeld. Performative recommendation: diversifying content via strategic incentives. In *International Conference on Machine Learning*, pp. 9082–9103. PMLR, 2023. 1, 16
- Mohammad Ghavamzadeh and Yaakov Engel. Bayesian policy gradient algorithms. *Advances in neural information processing systems*, 19, 2006. 3
- António Góis, Mehrnaz Mofakhami, Fernando P Santos, Gauthier Gidel, and Simon Lacoste-Julien. Performative prediction on games and mechanism design. *arXiv preprint arXiv:2408.05146*, 2024. 1, 16
- Moritz Hardt, Meena Jagadeesan, and Celestine Mendler-Dünner. Performative power. *Advances in Neural Information Processing Systems*, 35:22969–22981, 2022. 1, 16
- Alexander Havrilla, Yuqing Du, Sharath Chandra Raparthy, Christoforos Nalmpantis, Jane Dwivedi-Yu, Eric Hambro, Sainbayar Sukhbaatar, and Roberta Raileanu. Teaching large language models to reason with reinforcement learning. In *AI for Math Workshop@ ICML 2024.* 1
- Michael Herman, Tobias Gindele, Jörg Wagner, Felix Schmitt, and Wolfram Burgard. Inverse reinforcement learning with simultaneous estimation of rewards and dynamics. In *Artificial intelligence and statistics*, pp. 102–110. PMLR, 2016. 4
- Zachary Izzo, Lexing Ying, and James Zou. How to learn when data reacts to your model: performative gradient descent. In *International Conference on Machine Learning*, pp. 4641–4650. PMLR, 2021. 1, 16
- Zachary Izzo, James Zou, and Lexing Ying. How to learn when data gradually reacts to your model. In *International Conference on Artificial Intelligence and Statistics*, pp. 3998–4035. PMLR, 2022. 1, 16
- Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *Proceedings of the nine-teenth international conference on machine learning*, pp. 267–274, 2002a. 4
- Sham M Kakade. A natural policy gradient. Advances in neural information processing systems, 14, 2001. 2
- Sham M. Kakade. A natural policy gradient. In *Advances in Neural Information Processing Systems 14 (NIPS 2001)*, pp. 1531–1538, 2002. 16
- Sham M. Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *International Conference on Machine Learning*, 2002b. URL https://api.semanticscholar.org/CorpusID:31442909. 17

Mahdi Kallel, Debabrota Basu, Riad Akrour, and Carlo D'Eramo. Augmented bayesian policy search. In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024. URL https://openreview.net/forum?id=OvlcyABNQT. 4

- Vijay R. Konda and John N. Tsitsiklis. Actor-critic algorithms. In *Advances in Neural Information Processing Systems 13* (NIPS 2000), pp. 1008–1014, 2000. 16
- Qiang Li and Hoi-To Wai. State dependent performative prediction with stochastic approximation. In *International Conference on Artificial Intelligence and Statistics*, pp. 3164–3186. PMLR, 2022. 1, 16
- Qiang Li, Chung-Yiu Yau, and Hoi-To Wai. Multi-agent performative prediction with greedy deployment and consensus seeking agents. *Advances in Neural Information Processing Systems*, 35:38449–38460, 2022. 16
- Shengshi Li and Lin Yang. Horizon-free learning for markov decision processes and games: stochastically bounded rewards and improved bounds. In *International Conference on Machine Learning*, pp. 20221–20252. PMLR, 2023. 5
- Jingbin Liu, Xinyang Gu, and Shuai Liu. Policy optimization reinforcement learning with entropy regularization. *arXiv* preprint arXiv:1912.01557, 2019. 4
- Debmalya Mandal and Goran Radanovic. Performative reinforcement learning with linear markov decision process. *arXiv* preprint arXiv:2411.05234, 2024. 2, 3, 4, 16
- Debmalya Mandal, Stelios Triantafyllou, and Goran Radanovic. Performative reinforcement learning, 2023. URL https://arxiv.org/abs/2207.00046. 2, 3, 4, 5, 8, 9, 16
- Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. Feedback loop and bias amplification in recommender systems. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pp. 2145–2148, 2020. 1
- Jincheng Mei, Chenjun Xiao, Csaba Szepesvari, and Dale Schuurmans. On the global convergence rates of softmax policy gradient methods. In *International conference on machine learning*, pp. 6820–6829. PMLR, 2020. 3, 4, 5, 8, 22, 23, 32
- Celestine Mendler-Dünner, Juan Perdomo, Tijana Zrnic, and Moritz Hardt. Stochastic optimization for performative prediction. *Advances in Neural Information Processing Systems*, 33:4929–4939, 2020. 16
- John P Miller, Juan C Perdomo, and Tijana Zrnic. Outside the echo chamber: Optimizing the performative risk. In *International Conference on Machine Learning*, pp. 7710–7720. PMLR, 2021. 1, 16
- Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pp. 1928–1937, 2016. 2, 16
- Mehrnaz Mofakhami, Ioannis Mitliagkas, and Gauthier Gidel. Performative prediction with neural networks. In *International Conference on Artificial Intelligence and Statistics*, pp. 11079–11093. PMLR, 2023. 1, 16
- Adhyyan Narang, Evan Faulkner, Dmitriy Drusvyatskiy, Maryam Fazel, and Lillian J Ratliff. Multiplayer performative prediction: Learning in decision-dependent games. *Journal of Machine Learning Research*, 24(202):1–56, 2023. 1, 16
- Gergely Neu, Anders Jonsson, and Vicenç Gómez. A unified view of entropy-regularized markov decision processes. *arXiv* preprint arXiv:1705.07798, 2017. 4
- Matteo Papini, Damiano Binaghi, Giuseppe Canonaco, Matteo Pirotta, and Marcello Restelli. Stochastic variance-reduced policy gradient. In *International conference on machine learning*, pp. 4026–4035. PMLR, 2018. 9
- Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *International Conference on Machine Learning*, pp. 7599–7609. PMLR, 2020. 1, 16
- Georgios Piliouras and Fang-Yi Yu. Multi-agent performative prediction: From global stability and optimality to chaos. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pp. 1047–1074, 2023. 1, 16
- Vasilis Pollatos, Debmalya Mandal, and Goran Radanovic. On corruption-robustness in performative reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 19939–19947, 2025. 4
- Ben Rank, Stelios Triantafyllou, Debmalya Mandal, and Goran Radanovic. Performative reinforcement learning in gradually shifting environments. *arXiv preprint arXiv:2402.09838*, 2024. 2, 4, 5, 8, 9, 16

Mitas Ray, Lillian J Ratliff, Dmitriy Drusvyatskiy, and Maryam Fazel. Decision-dependent risk minimization in geometrically decaying dynamic environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 8081–8088, 2022. 16

- John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 1889–1897, 2015. 2, 16
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. In *arXiv preprint arXiv:1707.06347*, 2017. 2, 16
- Wenjie Shi, Shiji Song, and Cheng Wu. Soft policy gradient method for maximum entropy deep reinforcement learning. *arXiv* preprint arXiv:1909.03198, 2019. 4
- David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pp. 387–395. Pmlr, 2014. 1, 2, 4, 5
- Jason St. John, Christian Herwig, Diana Kafkes, Jovan Mitrevski, William A Pellico, Gabriel N Perdue, Andres Quintero-Parra, Brian A Schupbach, Kiyomi Seiya, Nhan Tran, et al. Real-time artificial intelligence for accelerator control: A study at the fermilab booster. *Physical Review Accelerators and Beams*, 24(10):104601, 2021.
- Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction, volume 1. MIT press, 1998. 1
- Richard S. Sutton, David A. McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems 12 (NIPS 1999)*, pp. 1057–1063. The MIT Press, 1999. 2, 3, 5, 16
- Jiexin Wang and Eiji Uchibe. Reward-punishment reinforcement learning with maximum entropy. In 2024 International Joint Conference on Neural Networks (IJCNN), pp. 1–7. IEEE, 2024. 4
- Xiaolu Wang, Chung-Yiu Yau, and Hoi To Wai. Network effects in performative prediction games. In *International Conference on Machine Learning*, pp. 36514–36540. PMLR, 2023. 16
- Yue Wang and Shaofeng Zou. Policy gradient method for robust reinforcement learning. In *International conference on machine learning*, pp. 23484–23526. PMLR, 2022. 3
- Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992. 2, 3, 5, 16
- Cathy Wu, Aravind Rajeswaran, Yan Duan, Vikash Kumar, Alexandre M Bayen, Sham Kakade, Igor Mordatch, and Pieter Abbeel. Variance reduction for policy gradient with action-dependent factorized baselines. *arXiv* preprint arXiv:1803.07246, 2018. 9
- Tengyu Xu, Haoyang Wang, and Yingbin Liang. Improved sample complexity for policy gradient via Poincaré inequality. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pp. 15728–15739, 2020. 16
- Rui Yuan, Robert M Gower, and Alessandro Lazaric. A general sample complexity analysis of vanilla policy gradient. In *International Conference on Artificial Intelligence and Statistics*, pp. 3332–3380. PMLR, 2022. 2, 3, 16
- Angela Zhang, Lei Xing, James Zou, and Joseph C Wu. Shifting machine learning for healthcare from development to deployment and from models to data. *Nature biomedical engineering*, 6(12):1330–1345, 2022. 1, 16
- Kaiqing Zhang, Zhuoran Yang, and Tamer Basar. On the global convergence of policy gradient methods with variance reduction. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pp. 11106–11115. PMLR, 2020. 16
- Rui Zhao, Xudong Sun, and Volker Tresp. Maximum entropy-regularized multi-goal reinforcement learning. In *International Conference on Machine Learning*, pp. 7553–7562. PMLR, 2019. 4

Appendix

Table of Contents

A	Notations	14
В	Details of the Toy Example: Loan Approvement Problem	15
C	Detailed Related Works	16
D	Impact of Policy Updates on PeMDPs (Section 3.1)	17
E	Smoothness of Performative Value Function and Entropy Regulariser	19
F	Derivation of Performative Policy Gradients	25
G	Convergence of PePG: Proofs of Section 4 G.1 Proofs for Unregularised Value Function	27 27
	G.2 Proofs for Entropy-regularised or Soft Value Function	32
Н	Ablation study on entropy regularisation	36
I	Technical Lemmas	37

A NOTATIONS

Notation	Description
S	state space
\mathcal{A}	action space
γ	discount factor
$\pi_{ heta}$	policy parametrized by $ heta$
$\Pi(\Theta)$	policy space
P_{π}	transition under the environment induced by policy π
$r_{m{\pi}}$	reward under the environment induced by policy π
$oldsymbol{\pi}_s^{\star}$	performatively stable policy
$oldsymbol{\pi}_o^{\star}$	performatively optimal policy
$\mathbf{P}_{\boldsymbol{\pi}_{o}^{\star}}$	reward under the environment induced by performatively optimal policy
$r_{oldsymbol{\pi}_o^{\star}}$	reward under the environment induced by performatively optimal policy
$egin{array}{l} r_{oldsymbol{\pi}^{\star}_{o}} \ d^{oldsymbol{\pi}^{\star}_{o}}_{oldsymbol{\pi}^{\star}_{o}} \ V^{oldsymbol{\pi}^{\star}_{o}}_{oldsymbol{\sigma}^{\star}_{o}} \ d^{oldsymbol{\pi}_{1}}_{oldsymbol{\pi}_{2}} \end{array}$	state-action occupancy of optimal policy
$V^{oldsymbol{\pi}^{\star}_{o}}_{oldsymbol{\pi}^{\star}_{o}}$	value function of optimal policy
$\boldsymbol{d}_{\boldsymbol{\pi}_2}^{\boldsymbol{\pi}_1}$	state-action occupancy of playing policy $oldsymbol{\pi}_2$ in the environment induced by policy $oldsymbol{\pi}_1$
$V_{m{\pi}_1}^{m{\pi}_2}$	value function for playing policy $oldsymbol{\pi}_2$ in the environment induced by policy $oldsymbol{\pi}_1$
$Q_{m{\pi}_1}^{m{\pi}_2}$	Q-value function for playing policy π_2 in the environment induced by policy π_1
$A_{m{\pi}_1}^{m{\pi}_2}$	advantage function for playing policy $oldsymbol{\pi}_2$ in the environment induced by policy $oldsymbol{\pi}_1$
Δ_K	K-dimensional simplex

B DETAILS OF THE TOY EXAMPLE: LOAN APPROVEMENT PROBLEM

Environment. We consider a population of loan applicants represented by a scalar feature $x \in \mathbb{R}$, distributed as $x \sim \mathcal{N}(\mu, \sigma^2)$, where μ is the population mean and $\sigma > 0$ is fixed.

Bank's Policy. The bank chooses a threshold policy parameterized by $\theta \in \mathbb{R}$. A loan is granted to an applicant x if $x \geq \theta$. To smooth analysis, we use a differentiable approximation: $\pi_{\theta}(x) = \sigma(k(x-\theta))$, where $\sigma(z) = \frac{1}{1+e^{-z}}$ is the logistic sigmoid and k > 0 controls smoothness.

Rewards. If a loan is granted to applicant x, the bank receives a random payoff:

$$r(x) = \begin{cases} +R & \text{if applicant repays,} \\ -L & \text{if applicant defaults} \end{cases}$$

with repayment probability $\mathbb{P}(\text{repay} \mid x) = \sigma(\gamma x - c)$, where $\gamma > 0$ controls sensitivity and c is a calibration constant. The expected reward from granting to x is

$$u(x) = \sigma(\gamma x - c) \cdot R - (1 - \sigma(\gamma x - c)) \cdot L.$$

Expected Utility. Given distribution $x \sim \mathcal{N}(\mu, \sigma^2)$, the bank's expected utility for policy θ is

$$U(\theta, \mu) = \mathbb{E}_{x \sim \mathcal{N}(\mu, \sigma^2)} [\pi_{\theta}(x) \cdot u(x)].$$

Performative Feedback. The population mean μ depends on the bank's policy, via the grant rate: $g(\theta, \mu) = \mathbb{E}_{x \sim \mathcal{N}(\mu, \sigma^2)} [\pi_{\theta}(x)]$.

We assume a bounded performative update rule: $\mu_{t+1} = (1 - \beta)\mu_t + \beta \cdot f(g(\theta, \mu_t))$, where $\beta \in [0, 1]$ is the performative strength and $f(g) \in [-M, M]$ maps the grant rate to a feasible population mean.

At equilibrium, the induced feature distribution satisfies the fixed point condition:

$$\mu^*(\theta) = (1 - \beta)\mu^*(\theta) + \beta f(g(\theta, \mu^*(\theta))).$$

Optimization Problems. ERM Optimum. Ignoring performative effects (i.e. assuming $\mu=\mu_0$ is fixed), the ERM-optimal policy solves $\theta^{\text{ERM}} = \arg\max_{\mu} U(\theta,\mu_0).$

Performative Optimum. Accounting for performative feedback, the performative-optimal policy solves

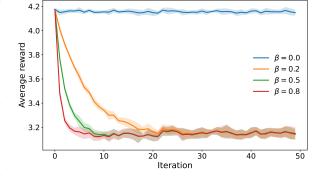
$$\theta^{\text{Perf}} = \arg \max_{\theta} U(\theta, \mu^*(\theta)).$$

Learning via Reinforcement Learning

An RL agent plays policies θ_t sequentially. At each round t:

- 1. Sample $x \sim \mathcal{N}(\mu_t, \sigma^2)$.
- 2. Grant loan with probability $\pi_{\theta_t}(x)$.
- 3. Observe reward r_t .
- 4. Update θ_{t+1} using policy gradient (REINFORCE).
- 5. Update population mean via performative dynamics:

$$\mu_{t+1} = (1 - \beta)\mu_t + \beta f(g(\theta_t, \mu_t)).$$



C DETAILED RELATED WORKS

 Performative Prediction. The study of performative prediction started with the pioneering work of (Perdomo et al., 2020), where they leveraged repeated retraining with the aim to converge towards a performatively stable point. We see extension of this work trying to achieve performative optimality (Izzo et al., 2021; 2022; Miller et al., 2021). This further opened a plethora of works in various other domains such as Multi-agent systems (Narang et al., 2023; Li et al., 2022; Piliouras & Yu, 2023), control systems (Cai et al., 2024; Barakat et al., 2025), stochastic optimisation (Li & Wai, 2022; Mendler-Dünner et al., 2020), games (Wang et al., 2023; Góis et al., 2024) etc. There has been several attempt of achieve performative optimality or stability for real-life tasks like recommendation (Eilat & Rosenfeld, 2023), to measure the power of firms (Hardt et al., 2022; Mofakhami et al., 2023), in healthcare (Zhang et al., 2022) etc. Another interesting setting is the *stateful* performative prediction i.e. prediction under gradual shifts in the distribution (Brown et al., 2022; Izzo et al., 2022; Ray et al., 2022), that paved the way for incorporating performative prediction in Reinforcement Learning.

Performative Reinforcement Learning. Bell et al. (2021) were the first to propose a setting where the transition and reward of an underlying MDP depend non-deterministically on the deployed policy, thus capturing the essence of performativity to some extent. However, Mandal et al. (2023) can be considered the pioneer in introducing the notion of "*Performative Reinforcement Learning*" and its solution concepts, performatively stable and optimal policy. They propose direct optimization and ascent based techniques which manage to attain performative stability upon repeated retraining. Extensions to this work, Rank et al. (2024) and Mandal & Radanovic (2024) manage to solve the same problem with delayed retraining for linear MDPs. However, there exists no literature that proposes a performative RL algorithm that converges to the performative optimal policy.

Specifically, Mandal et al. (2023) frames the question of using policy gradient to find stable policies as an open problem. The authors further contemplate, as PG functions in the policy space, whether it is possible to converge towards a stable policy. Thus, in this paper, we affirmatively solve an extension (rather a harder problem) of this open problem for tabular MDPs with softmax policies.

Policy Gradient Algorithms. Policy gradient algorithms build a central paradigm in reinforcement learning, directly optimizing parametrised policies by estimating the gradient of expected return. The foundational policy gradient theorem (Sutton et al., 1999) established an expression for this gradient in terms of the score and action-value function, while Williams (1992) introduced the REINFORCE algorithm, providing an unbiased likelihood-ratio estimator. Convergence properties of stochastic gradient ascent in policy space were analysed in these early works. Subsequently, Konda & Tsitsiklis (2000) formalized actor–critic methods via two-timescale stochastic approximation, and Kakade (2002) proposed the natural policy gradient, leveraging the Fisher information geometry to accelerate learning. Extensions to trust region methods (Schulman et al., 2015), proximal policy optimization (Schulman et al., 2017), and entropy-regularized objectives (Mnih et al., 2016) have made policy gradient methods widely practical in high-dimensional settings. Recent theoretical advances provide finite-sample convergence guarantees and complexity analyses (Agarwal et al., 2021; Yuan et al., 2022), as well as robustness to distributional shift and adversarial perturbations (Zhang et al., 2020; Xu et al., 2020). Collectively, this body of work establishes policy gradient methods as both practically effective and theoretically grounded method for solving MDP.

D IMPACT OF POLICY UPDATES ON PEMDPS (SECTION 3.1)

Proof of Lemma 1. We do this proof in majorly two steps. First step involves a decomposition of the difference in value function into two terms: (i) difference in value function after deploying the same policy while agent plays two different policies i.e. the difference that explains stability of the deployed policy, and (ii) difference in value function for deploying two different policies i.e. performance difference for changing the deployed policy. While the second term can be bounded using classic performance difference lemma, in the next and final step, we control the stability inducing term (i).

Step 1: Decomposition. We start by decomposing the performative performance difference to get a stability and a performance difference terms separately.

$$V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s_0) - V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}'}(s_0) = \underbrace{V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s_0) - V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(s_0)}_{\text{performative shift term}} + \underbrace{V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(s_0) - V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}'}(s_0)}_{\text{performance difference term}}$$

$$= V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s_0) - V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(s_0) + \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim d_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(\cdot|s_0)}[A_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}'}(s,a)] \tag{15}$$

The last equality is a consequence of the classical performance difference lemma (Kakade & Langford, 2002b).

Step 2: Controlling the performative stability term. First, let us define $\mathbf{P}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s',s) \triangleq \sum_{a \in \mathcal{A}} \mathbf{P}_{\boldsymbol{\pi}}(s'|s,a)\boldsymbol{\pi}(a|s)$, and $\langle \mathbf{P}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot,s_0), V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot) \rangle \triangleq \sum_{s \in \mathcal{S}} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s)\mathbf{P}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s,s_0)$.

Now, we observe that

$$V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s_0) - V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(s_0) = \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot|s_0)} \Big[r_{\boldsymbol{\pi}}(s,a) - r_{\boldsymbol{\pi}'}(s,a) \Big]$$

$$+ \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot|s_0)} \Big[\gamma \langle \mathbf{P}_{\boldsymbol{\pi}}(\cdot|s,a) - \mathbf{P}_{\boldsymbol{\pi}'}(\cdot|s,a), V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(\cdot) \rangle + \gamma \langle \mathbf{P}_{\boldsymbol{\pi}}(\cdot|s,a), V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(\cdot) - V_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(\cdot) \rangle \Big]$$

Combining steps 1 and 2 and taking expectation over $s_0 \sim \rho$, we get

$$V_{\pi}^{\pi}(\rho) - V_{\pi'}^{\pi'}(\rho) = \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim d_{\pi',\rho}^{\pi}} [A_{\pi'}^{\pi'}(s,a)] + \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim d_{\pi,\rho}^{\pi}} \Big[(r_{\pi}(s,a) - r_{\pi'}(s,a)) + \gamma (\mathbf{P}_{\pi}(\cdot|s,a) - \mathbf{P}_{\pi'}(\cdot|s,a))^{\top} V_{\pi}^{\pi}(\cdot) \Big].$$

Proof of Lemma 2. We do this proof in three steps. We start from the final expression in Lemma 1, then in step 2 we impose bounds on reward and transition differences leveraging the Lipschitz assumption. Lastly, we bound the policy difference in first order norm using relation between Total Variation (TV) and Hellinger distance.

Step 1: From Lemma 1, we have

$$V_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(s_{0}) - V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}^{\star}}(s_{0}) = \left| V_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(s_{0}) - V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}^{\star}}(s_{0}) \right|$$

$$= \frac{1}{1 - \gamma} \left| \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\cdot,\cdot|s_{0})} [A_{\boldsymbol{\pi}_{\theta}^{\star}}^{\boldsymbol{\pi}_{\theta}^{\star}}(s,a)] + \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}_{o}^{\star}}(\cdot,\cdot|s_{0})} \left[r_{\boldsymbol{\pi}_{o}^{\star}}(s,a) - r_{\boldsymbol{\pi}_{\theta}}(s,a) + \gamma (\mathbf{P}_{\boldsymbol{\pi}_{o}^{\star}} - \mathbf{P}_{\boldsymbol{\pi}_{\theta}})^{\top} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}} \right] \right|$$

$$\stackrel{(a)}{\leq} \frac{1}{1 - \gamma} \left| \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\cdot,\cdot|s_{0})} [A_{\boldsymbol{\pi}_{\theta}^{\star}}^{\boldsymbol{\pi}_{\theta}}(s,a)] \right|$$

$$\stackrel{(b)}{\leq} \frac{1}{1 - \gamma} \left| \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\cdot,\cdot|s_{0})} [A_{\boldsymbol{\pi}_{\theta}^{\star}}^{\boldsymbol{\pi}_{\theta}}(s,a)] \right|$$

$$+ \frac{1}{1 - \gamma} \left| \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}_{o}^{\star}}(\cdot,\cdot|s_{0})} [r_{\boldsymbol{\pi}_{o}^{\star}}(s,a) - r_{\boldsymbol{\pi}_{\theta}}(s,a)] \right|$$

$$+ \frac{\gamma}{1 - \gamma} \left| \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}_{o}^{\star}}(\cdot,\cdot|s_{0})} \left[(\mathbf{P}_{\boldsymbol{\pi}_{o}^{\star}} - \mathbf{P}_{\boldsymbol{\pi}_{\theta}})^{\top} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}} \right] \right|$$

$$(16)$$

where both (a) and (b) are due to applying triangle inequality

Step 2: Using Jensen's Inequality together with the fact that $d_{\pi}^{\pi}(s, a|s_0) \leq 1$ and leveraging the Lipschitzness assumption on reward and transitions, we get

$$\left| \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot,\cdot|s_0)} \left[r_{\boldsymbol{\pi}_o^{\star}}(s,a) - r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a) + \gamma (\mathbf{P}_{\boldsymbol{\pi}_o^{\star}} - \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}})^{\top} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}} \right] \right| \leq L_r \left\| \boldsymbol{\pi}_o^{\star} - \boldsymbol{\pi}_{\boldsymbol{\theta}} \right\|_1 + \gamma L_{\mathbf{P}} \left\| \mathbf{V}_{\boldsymbol{\pi}_o^{\star}}^{\boldsymbol{\pi}_o^{\star}} \right\|_{\infty} \left\| \boldsymbol{\pi}_o^{\star} - \boldsymbol{\pi}_{\boldsymbol{\theta}} \right\|_1$$

Finally, due to Assumption 1, we get $\|\mathbf{V}_{\boldsymbol{\pi}_{o}}^{\boldsymbol{\pi}_{o}^{*}}\|_{\infty} \leq \frac{R_{\max}}{1-\gamma}$, and thus,

$$\left| \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot,\cdot|s_0)} \left[r_{\boldsymbol{\pi}_o^{\star}}(s,a) - r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a) + \gamma (\mathbf{P}_{\boldsymbol{\pi}_o^{\star}} - \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}})^{\top} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}} \right] \right| \leq L_r \left\| \boldsymbol{\pi}_o^{\star} - \boldsymbol{\pi}_{\boldsymbol{\theta}} \right\|_1 + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} R_{\max} \left\| \boldsymbol{\pi}_o^{\star} - \boldsymbol{\pi}_{\boldsymbol{\theta}} \right\|_1$$

Step 3: We know $\|\boldsymbol{\pi}_o^{\star} - \boldsymbol{\pi}_{\boldsymbol{\theta}}\|_1 = 2\text{TV}\left(\boldsymbol{\pi}_o^{\star} \parallel \boldsymbol{\pi}_{\boldsymbol{\theta}}\right) \leq D_{\text{H}}\left(\boldsymbol{\pi}_o^{\star} \lVert \boldsymbol{\pi}_{\boldsymbol{\theta}}\right)$. Thus,

$$\left| \mathbb{E}_{(s,a) \sim d_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot,\cdot|s_0)} \left[r_{\boldsymbol{\pi}_o^{\star}}(s,a) - r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a) + \gamma (\mathbf{P}_{\boldsymbol{\pi}_o^{\star}} - \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}})^{\top} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}} \right] \right| \leq 2\sqrt{2} \left(L_r + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} R_{\max} \right) D_{\mathbf{H}} \left(\boldsymbol{\pi}_o^{\star}(\cdot \mid s_0) \| \boldsymbol{\pi}_{\boldsymbol{\theta}}(\cdot \mid s_0) \right)$$

$$(18)$$

We conclude this proof by putting the upper bound in Equation (18) in Equation (17) and taking expectation over $s_0 \sim \rho$ to get the desired expression.

E SMOOTHNESS OF PERFORMATIVE VALUE FUNCTION AND ENTROPY REGULARISER

Lemma 4 (Performative Smoothness Lemma). Let $\pi_{\alpha} \triangleq \pi_{\theta+\alpha u}$, and let $V_{\alpha}^{\alpha}(s_0)$ be the corresponding value at a fixed state s_0 , i.e., $V_{\alpha}^{\alpha}(s_0) \triangleq V_{\pi_{\alpha}}^{\pi_{\alpha}}(s_0)$. If the following conditions hold true,

$$\sum_{a \in \mathcal{A}} \left| \frac{d\mathbf{\pi}_{\alpha}(a \mid s_{0})}{d\alpha} \right|_{\alpha=0} \leq C_{1}, \quad \sum_{a \in \mathcal{A}} \left| \frac{d^{2}\mathbf{\pi}_{\alpha}(a \mid s_{0})}{d\alpha^{2}} \right|_{\alpha=0} \leq C_{2}, \sum_{s \in \mathcal{S}} \left| \frac{d\mathbf{P}_{\alpha}(s \mid s_{0}, a_{0})}{d\alpha} \right|_{\alpha=0} \leq T_{1}, \\
\sum_{s \in \mathcal{S}} \left| \frac{d^{2}\mathbf{P}_{\alpha}(s \mid s_{0}, a_{0})}{d\alpha^{2}} \right|_{\alpha=0} \leq T_{2}, \sum_{a \in \mathcal{A}} \left| \frac{dr_{\alpha}(s_{0}, a)}{d\alpha} \right|_{\alpha=0} \leq R_{1}, \quad \sum_{a \in \mathcal{A}} \left| \frac{d^{2}r_{\alpha}(s_{0}, a)}{d\alpha^{2}} \right|_{\alpha=0} \leq R_{2},$$

we get

$$\max_{\|u\|_2=1} \left\| \frac{\mathrm{d}^2 V_{\alpha}^{\alpha}(s_0)}{\mathrm{d}\alpha^2} \right|_{\alpha=0} \le \frac{C_2}{1-\gamma} + 2C_1\beta_1 + C_2\beta_2 \triangleq L,$$

where
$$\beta_1 = \frac{\gamma}{(1-\gamma)^2}(C_1+T_1) + \frac{R_1}{1-\gamma}$$
 and $\beta_2 = \frac{2\gamma^2}{(1-\gamma)^3}(C_1+T_1)^2 + \frac{\gamma}{(1-\gamma)^2}(C_2+2C_1T_1+T_2) + \frac{2\gamma R_1}{(1-\gamma)^2}(C_2+2C_1T_1+T_2) + \frac{R_2}{(1-\gamma)^2}$.

Proof. **Step 1:** To prove the second order smoothness of the value function we start by taking its second derivative. Consider the expected return under policy π_{α} :

$$V_{\alpha}^{\alpha}(s_0) = \sum_{a} \boldsymbol{\pi}_{\alpha}(a \mid s_0) Q_{\alpha}^{\alpha}(s_0, a)$$

Differentiating twice with respect to α , we obtain:

$$\frac{\mathrm{d}^2 V_\alpha^\alpha(s_0)}{\mathrm{d}\alpha^2} = \sum_a \frac{\mathrm{d}^2 \pmb{\pi}_\alpha(a \mid s_0)}{\mathrm{d}\alpha^2} Q_\alpha^\alpha(s_0, a) + 2 \sum_a \frac{\mathrm{d} \pmb{\pi}_\alpha(a \mid s_0)}{\mathrm{d}\alpha} \frac{\mathrm{d}Q_\alpha^\alpha(s_0, a)}{\mathrm{d}\alpha} + \sum_a \pmb{\pi}_\alpha(a \mid s_0) \frac{\mathrm{d}^2 Q_\alpha^\alpha(s_0, a)}{\mathrm{d}\alpha^2}$$

 $Q^{\alpha}_{\alpha}(s_0, a_0)$ is the Q-function corresponding to the policy π_{α} at state s_0 and action a_0 . Observe that $Q^{\alpha}_{\alpha}(s_0, a_0)$ can further be written as:

$$Q_{\alpha}^{\alpha}(s_0, a_0) = e_{(s_0, a_0)}^{\top} (I - \gamma \tilde{\mathbf{P}}(\alpha))^{-1} r_{\alpha} = e_{(s_0, a_0)}^{\top} M(\alpha) r_{\alpha}$$

where $M(\alpha) \triangleq (I - \gamma \mathbf{P}(\alpha))^{-1}$ and $\tilde{\mathbf{P}}(\alpha)$ is the state-action transition matrix under policy $\boldsymbol{\pi}_{\alpha}$, defined as:

$$[\tilde{\mathbf{P}}(\alpha)](s', a' \mid s, a) \triangleq \boldsymbol{\pi}_{\alpha}(a' \mid s') \mathbf{P}_{\alpha}(s' \mid s, a)$$

Differentiating $Q^{\alpha}_{\alpha}(s,a)$ with respect to α gives:

$$\frac{\mathrm{d}Q_{\alpha}^{\alpha}(s_0, a_0)}{\mathrm{d}\alpha} = \gamma e_{(s_0, a_0)}^{\top} M(\alpha) \frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha} M(\alpha) r_{\alpha} + e_{(s_0, a_0)}^{\top} M(\alpha) \frac{\mathrm{d}r_{\alpha}}{\mathrm{d}\alpha}$$

And correspondingly,

$$\frac{\mathrm{d}^{2}Q_{\alpha}^{\alpha}(s_{0}, a_{0})}{\mathrm{d}\alpha^{2}} = 2\gamma^{2}e_{(s_{0}, a_{0})}^{\mathsf{T}}M(\alpha)\frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha}M(\alpha)\frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha}M(\alpha)r_{\alpha} + \gamma e_{(s_{0}, a_{0})}^{\mathsf{T}}M(\alpha)\frac{\mathrm{d}^{2}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha^{2}}M(\alpha)r_{\alpha}
+ \gamma e_{(s_{0}, a_{0})}^{\mathsf{T}}M(\alpha)\frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha}M(\alpha)\frac{\mathrm{d}r_{\alpha}}{\mathrm{d}\alpha} + e_{(s_{0}, a_{0})}^{\mathsf{T}}M(\alpha)\frac{\mathrm{d}^{2}r_{\alpha}}{\mathrm{d}\alpha^{2}}
+ \gamma e_{(s_{0}, a_{0})}^{\mathsf{T}}M(\alpha)\frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha}M(\alpha)\frac{\mathrm{d}r_{\alpha}}{\mathrm{d}\alpha} \tag{19}$$

Step 2: Now we need to find the derivative of $\tilde{\mathbf{P}}(\alpha)$ w.r.t α in order to substitute in (19). Hence, we can differentiate $\tilde{\mathbf{P}}(\alpha)$ with respect to α to obtain:

$$\frac{d\tilde{\mathbf{P}}(\alpha)}{d\alpha}\bigg|_{\alpha=0} (s', a' \mid s, a) = \frac{d\boldsymbol{\pi}_{\alpha}(a' \mid s')}{d\alpha}\bigg|_{\alpha=0} \mathbf{P}_{\alpha}(s' \mid s, a) + \frac{d\mathbf{P}_{\alpha}(s' \mid s, a)}{d\alpha}\bigg|_{\alpha=0} \boldsymbol{\pi}_{\alpha}(a' \mid s')$$

Now, for an arbitrary vector x, we have:

$$\left[\frac{d\tilde{\mathbf{P}}(\alpha)}{d\alpha} \bigg|_{\alpha=0} \mathbf{x} \right]_{(s,a)} = \sum_{s',a'} \frac{d\boldsymbol{\pi}_{\alpha}(a'\mid s')}{d\alpha} \bigg|_{\alpha=0} \mathbf{P}_{\alpha}(s'\mid s,a) \mathbf{x}_{s',a'} + \sum_{s',a'} \frac{d\mathbf{P}_{\alpha}(s'\mid s,a)}{d\alpha} \bigg|_{\alpha=0} \boldsymbol{\pi}_{\alpha}(a'\mid s') \mathbf{x}_{s',a'}$$

Taking the maximum over unit vectors \mathbf{u} in ℓ_2 -norm:

$$\begin{aligned} \max_{\|\mathbf{u}\|_{2}=1} \left\| \frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha} \right|_{\alpha=0} \mathbf{x} \right\|_{\infty} &\leq \max_{\|\mathbf{u}\|_{2}=1} \left| \sum_{s',a'} \frac{\mathrm{d}\boldsymbol{\pi}_{\alpha}(a'\mid s')}{\mathrm{d}\alpha} \right|_{\alpha=0} \mathbf{P}_{\alpha}(s'\mid s,a) \mathbf{x}_{s',a'} \right| \\ &+ \max_{\|\mathbf{u}\|_{2}=1} \left| \sum_{s',a'} \frac{\mathrm{d}\mathbf{P}_{\alpha}(s'\mid s,a)}{\mathrm{d}\alpha} \right|_{\alpha=0} \boldsymbol{\pi}_{\alpha}(a'\mid s') \mathbf{x}_{s',a'} \right| \\ &\leq \max_{s,a} \sum_{s'} \mathbf{P}_{\alpha}(s'\mid s,a) \sum_{a'} \left| \frac{\mathrm{d}\boldsymbol{\pi}_{\alpha}(a'\mid s')}{\mathrm{d}\alpha} \right|_{\alpha=0} \cdot \|\mathbf{x}\|_{\infty} \\ &+ \max_{s,a} \sum_{a'} \boldsymbol{\pi}_{\alpha}(a'\mid s') \sum_{s'} \left| \frac{\mathrm{d}\mathbf{P}_{\alpha}(s'\mid s,a)}{\mathrm{d}\alpha} \right|_{\alpha=0} \cdot \|\mathbf{x}\|_{\infty} \\ &\leq \max_{s,a} \sum_{s'} \mathbf{P}_{\alpha}(s'\mid s,a) \|\mathbf{x}\|_{\infty} C_{1} + \max_{s,a} \sum_{a'} \boldsymbol{\pi}(a'\mid s') \|\mathbf{x}\|_{\infty} T_{1} \\ &\leq C_{1} \|\mathbf{x}\|_{\infty} + T_{1} \|\mathbf{x}\|_{\infty} = (C_{1} + T_{1}) \|\mathbf{x}\|_{\infty} \end{aligned}$$

By the definition of the ℓ_{∞} -norm, we conclude:

$$\max_{\|\mathbf{u}\|_{2}=1} \left\| \frac{\mathrm{d}\mathbf{P}_{\alpha}}{\mathrm{d}\alpha} \right|_{\alpha=0} \mathbf{x} \right\|_{\infty} \le (C_{1} + T_{1}) \|\mathbf{x}\|_{\infty} \tag{20}$$

Similarly, differentiating $\tilde{\mathbf{P}}(\alpha)$ twice w.r.t. α , we get

$$\left[\frac{\mathrm{d}^{2}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha^{2}}\Big|_{\alpha=0}\right]_{(s,a)\to(s',a')} = \frac{\mathrm{d}^{2}\boldsymbol{\pi}_{\alpha}(a'\mid s')}{(\mathrm{d}\alpha)^{2}}\Big|_{\alpha=0} \mathbf{P}_{\alpha}(s'\mid s,a) + \frac{\mathrm{d}^{2}\mathbf{P}_{\alpha}(s'\mid s,a)}{\mathrm{d}\alpha^{2}}\Big|_{\alpha=0} \boldsymbol{\pi}_{\alpha}(a'\mid s') + 2\left.\frac{\mathrm{d}\boldsymbol{\pi}_{\alpha}(a'\mid s')}{\mathrm{d}\alpha}\Big|_{\alpha=0} \frac{\mathrm{d}\mathbf{P}_{\alpha}(s'\mid s,a)}{\mathrm{d}\alpha}\Big|_{\alpha=0}$$

Hence, we can consider the following norm bound:

$$\max_{\|\mathbf{u}\|_{2}=1} \left\| \frac{\mathrm{d}^{2} \tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha^{2}} \right\|_{\alpha=0} \mathbf{x} \right\|_{\infty} \leq C_{2} \|\mathbf{x}\|_{\infty} + 2C_{1}T_{1} \|\mathbf{x}\|_{\infty} + T_{2} \|\mathbf{x}\|_{\infty} = (C_{2} + 2C_{1}T_{1} + T_{2}) \|\mathbf{x}\|_{\infty}$$
(21)

Step 3: Now we need to put the pieces back together in order to calculate the second derivative of V_{α}^{α} w.r.t α . Let us recall $M(\alpha)$. Using the power series expansion of the matrix inverse, we can write $M(\alpha)$ as:

$$M(\alpha) = (I - \gamma \tilde{\mathbf{P}}(\alpha))^{-1} = \sum_{n=0}^{\infty} \gamma^n \tilde{\mathbf{P}}(\alpha)^n$$

which implies that $M(\alpha) \ge 0$ (component-wise), and

$$M(\alpha)\mathbf{1} = \frac{1}{1-\gamma}\mathbf{1},$$

i.e., each row of $M(\alpha)$ is positive and sums to $\frac{1}{1-\gamma}$.

This implies:

$$\max_{\|u\|_2=1} \|M(\alpha)\mathbf{x}\|_{\infty} \le \frac{1}{1-\gamma} \|\mathbf{x}\|_{\infty}.$$

This gives, using the expressions for $\frac{\mathrm{d}^2 Q_{\alpha}^{\alpha}(s_0, a_0)}{\mathrm{d}\alpha^2}$ and $\frac{\mathrm{d} Q_{\alpha}^{\alpha}(s_0, a_0)}{\mathrm{d}\alpha}$, an upper bound on their magnitudes based on $\|\mathbf{x}\|_{\infty}$ and constants arising from bounds on the derivatives of $\tilde{\mathbf{P}}(\alpha)$ and r_{α} .

$$\begin{split} & \max_{\|\mathbf{u}\|_{2}=1} \left\| \frac{\mathrm{d}^{2} Q_{\alpha}^{\alpha}(s_{0}, a_{0})}{\mathrm{d}\alpha^{2}} \right\|_{\infty} \\ \leq & 2\gamma^{2} \left\| M(\alpha) \frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha} M(\alpha) \frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha} M(\alpha) r_{\alpha} \right\|_{\infty} + \gamma \left\| M(\alpha) \frac{\mathrm{d}^{2}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha^{2}} M(\alpha) r_{\alpha} \right\|_{\infty} \\ & + \gamma \left\| M(\alpha) \frac{\mathrm{d}^{2}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha^{2}} M(\alpha) \frac{\mathrm{d}r_{\alpha}}{\mathrm{d}\alpha} \right\|_{\infty} + \left\| M(\alpha) \frac{\mathrm{d}^{2}r_{\alpha}}{\mathrm{d}\alpha^{2}} \right\|_{\infty} + 2\gamma \left\| M(\alpha) \frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha} M(\alpha) \frac{\mathrm{d}r_{\alpha}}{\mathrm{d}\alpha} \right\|_{\infty} \end{split}$$

Bounding using known bounds on transitions and rewards:

$$\max_{\|\mathbf{u}\|_{2}=1} \left\| \frac{\mathrm{d}^{2} Q_{\alpha}^{\alpha}(s_{0}, a_{0})}{\mathrm{d}\alpha^{2}} \right\|_{\infty} \leq \frac{2\gamma^{2}}{(1-\gamma)^{3}} (C_{1} + T_{1})^{2} + \frac{\gamma}{(1-\gamma)^{2}} (C_{2} + 2C_{1}T_{1} + T_{2}) + \frac{2\gamma R_{1}}{(1-\gamma)^{2}} (C_{2} + 2C_{1}T_{1} + T_{2}) + \frac{R_{2}}{1-\gamma} + \frac{\gamma C_{1}R_{1}}{(1-\gamma)^{2}} = \beta_{2}$$

Corresponding bound on the first derivative is:

$$\max_{\|\mathbf{u}\|_{2}=1} \left\| \frac{\mathrm{d}Q_{\alpha}^{\alpha}(s_{0}, a_{0})}{\mathrm{d}\alpha} \right\|_{\infty} \leq \gamma \left\| M(\alpha) \frac{\mathrm{d}\tilde{\mathbf{P}}(\alpha)}{\mathrm{d}\alpha} M(\alpha) \frac{\mathrm{d}r_{\alpha}}{\mathrm{d}\alpha} \right\|_{\infty} + \left\| M(\alpha) \frac{\mathrm{d}r_{\alpha}}{\mathrm{d}\alpha} \right\|_{\infty} \\
\leq \frac{\gamma}{(1-\gamma)^{2}} (C_{1} + T_{1}) + \frac{R_{1}}{1-\gamma} = \beta_{1}$$

Step 4: Finally, putting all the bounds together to evaluate the upper bound of the desired quantity, we get,

$$\max_{\|\mathbf{u}\|_{2}=1} \left\| \frac{\mathrm{d}^{2} V_{\alpha}^{\alpha}(s_{0})}{\mathrm{d}\alpha^{2}} \right\|_{\infty} \leq \frac{C_{2}}{1-\gamma} + 2C_{1}\beta_{1} + \beta_{2}$$
(22)

Corollary 1. For softmax PeMDPs, we characterise

$$C_1 = 2$$
 , $C_2 = 6, T_1 = L_{\mathbf{P}} = \max_{s} |\psi(s)| \triangleq \psi_{\max}$, $T_2 = \max_{s} |\psi(s)|^2, R_1 = L_r = \xi |\mathcal{A}|$, $R_2 = 0$

Thus,

$$\max_{\|u\|_2=1} \left\| \left. \frac{\mathrm{d}^2 V_\alpha^\alpha(s_0)}{\mathrm{d}\alpha^2} \right|_{\alpha=0} \right\| \leq \mathcal{O}\left(\frac{|\mathcal{A}|}{(1-\gamma)^2} \right) ,$$

Proof. We use the expressions already found in (26) to state the following:

$$\sum_{a \in A} \left| \frac{\mathrm{d}}{\mathrm{d}\alpha} \boldsymbol{\pi}_{\boldsymbol{\theta} + \alpha \mathbf{u}}(a \mid s) \right|_{\alpha = 0} \leq \sum_{a \in A} \boldsymbol{\pi}_{\boldsymbol{\theta}}(a \mid s) \left| \mathbf{u}_{s}^{\top} \left(\mathbf{e}_{a} - \boldsymbol{\pi}(\cdot \mid s) \right) \right| \leq \max_{a \in A} \left(\mathbf{u}_{s}^{\top} \mathbf{e}_{a} + \mathbf{u}_{s}^{\top} \boldsymbol{\pi}(\cdot \mid s) \right) \leq 2.$$

Similarly, differentiating once again w.r.t. α , we get

$$\sum_{a \in \mathcal{A}} \left| \frac{\mathrm{d}^2}{\mathrm{d}\alpha^2} \boldsymbol{\pi}_{\boldsymbol{\theta} + \alpha \mathbf{u}}(a \mid s) \right|_{\alpha = 0} \leq \max_{a \in \mathcal{A}} \left(\mathbf{u}_s^\top \mathbf{e}_a \mathbf{e}_a^\top \mathbf{u}_s + \mathbf{u}_s^\top \mathbf{e}_a \boldsymbol{\pi}(\cdot \mid s)^\top \mathbf{u}_s + \mathbf{u}_s^\top \boldsymbol{\pi}(\cdot \mid s) \mathbf{e}_a^\top \mathbf{u}_s \right)$$

$$+2 \mathbf{u}_s^{\top} \boldsymbol{\pi}(\cdot \mid s) \boldsymbol{\pi}(\cdot \mid s)^{\top} \mathbf{u}_s + \mathbf{u}_s^{\top} \operatorname{diag}(\boldsymbol{\pi}(\cdot \mid s)) \mathbf{u}_s) \leq 6.$$

And hence for transition we get,

$$\sum_{s' \in \mathcal{S}} \left| \frac{\mathrm{d}}{\mathrm{d}\alpha} \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta} + \alpha \mathbf{u}}}(a \mid s) \right|_{\alpha = 0} \leq \sum_{s' \in \mathcal{S}} |\psi(s')| \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s' \mid s, a) |\mathbf{u}_{s,a}(1 - \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\cdot \mid s, a))| \leq |\mathbf{u}_{s,a}| \max_{s} |\psi(s)| \leq \max_{s} |\psi(s)|$$

And similarly, it can be shown that:

$$\sum_{a \in \mathcal{A}} \left| \frac{\mathrm{d}^2}{\mathrm{d}\alpha^2} \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta} + \alpha \mathbf{u}}}(a \mid s) \right|_{\alpha = 0} \leq |\mathbf{u}_{s,a}|^2 \max_{s} |\psi(s)|^2 \leq \max_{s} |\psi(s)|^2$$

Similarly for rewards we get:

$$\sum_{a \in \mathcal{A}} \left| \frac{\mathrm{d}}{\mathrm{d}\alpha} \, r_{\pi_{\theta + \alpha \mathbf{u}}}(a \mid s) \right|_{\alpha = 0} \right| \leq \xi |\mathcal{A}| \qquad , \qquad \sum_{a \in \mathcal{A}} \left| \frac{\mathrm{d}^2}{\mathrm{d}\alpha^2} \, r_{\pi_{\theta + \alpha \mathbf{u}}}(a \mid s) \right|_{\alpha = 0} \right| = 0$$

Hence, we can use the following choice of constants for softmax parametrization,

$$C_1 = 2$$
 , $C_2 = 6$
 $T_1 = L_{\mathbf{P}} = \max_{s} |\psi(s)|$, $T_2 = \max_{s} |\psi(s)|^2$
 $R_1 = L_r = \xi |\mathcal{A}|$, $R_2 = 0$

to get the desired order of $\max_{\|u\|_2=1} \left\| \frac{\mathrm{d}^2 V_\alpha^\alpha(s_0)}{\mathrm{d}\alpha^2} \right|_{\alpha=0} \right\|$.

Lemma 5 (Smoothness of Entropy Regularizer). *Define the discounted entropy regularizer as:*

$$\mathcal{H}_{\boldsymbol{\pi}_{\boldsymbol{\theta}_{\alpha}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}_{\alpha}}}(s) = \mathbb{E}_{\tau \sim \mathbf{P}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}} \left[\sum_{t=0}^{\infty} -\gamma^{t} \log \boldsymbol{\pi}_{\boldsymbol{\theta}_{\alpha}}(a_{t} \mid s_{t}) \right]$$

Under the same assumptions as **4**, *the following holds:*

$$\max_{\|u\|_2=1} \left\| \frac{\partial^2 \mathcal{H}^{\pi_{\theta_{\alpha}}}_{\pi_{\theta_{\alpha}}}(s)}{\partial \alpha^2} \right\|_{\alpha=0} \leq \beta_{\lambda}$$

where

$$\beta_{\lambda} = 2\gamma^{2} \cdot \frac{3 \cdot (1 + \log|\mathcal{A}|)}{1 - \gamma} + \gamma \cdot \frac{2 \cdot \log|\mathcal{A}|}{(1 - \gamma)^{2}} (C_{1} + T_{1})$$
$$+ 2\gamma \cdot \frac{\log|\mathcal{A}|}{(1 - \gamma)^{2}} (C_{2} + 2C_{1}T_{1} + T_{2}) + \frac{\log|\mathcal{A}|}{(1 - \gamma)^{3}} (C_{1} + T_{1})^{2}$$

Proof. **Step 1:** Define the state-wise entropy term:

$$h_{\boldsymbol{\theta}_{\alpha}}(s) = -\sum_{a} \boldsymbol{\pi}_{\boldsymbol{\theta}_{\alpha}}(a \mid s) \log \boldsymbol{\pi}_{\boldsymbol{\theta}_{\alpha}}(a \mid s).$$

From Mei et al. (2020) (Lemma 7) we report that,

$$\left\| \frac{\partial h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha} \right\|_{\infty} \le 2 \cdot \log |\mathcal{A}| \cdot \|u\|_{2}, \qquad \left\| \frac{\partial^{2} h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha^{2}} \right\|_{\infty} \le 3 \cdot (1 + \log |\mathcal{A}|) \cdot \|\mathbf{u}\|_{2}^{2}. \tag{23}$$

Additionally, Mei et al. (2020) also presents a second result expressing the second derivative of the entropy w.r.t α ,

$$\begin{split} \frac{\partial^{2} \mathcal{H}_{\boldsymbol{\pi}\boldsymbol{\theta}_{\alpha}}^{\boldsymbol{\pi}\boldsymbol{\theta}_{\alpha}}(s)}{\partial \alpha^{2}} = & 2\gamma^{2} \, \mathbf{e}_{s}^{\top} M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \\ & + \gamma \, \mathbf{e}_{s}^{\top} M(\alpha) \frac{\partial^{2} \mathbf{P}(\alpha)}{\partial \alpha^{2}} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} + 2\gamma \, \mathbf{e}_{s}^{\top} M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) \frac{\partial h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha} + \mathbf{e}_{s}^{\top} M(\alpha) \frac{\partial^{2} h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha^{2}}. \end{split}$$

Step 2: Now we proceed with bounding the absolute value of each term which will contribute towards bounding the overall second derivative of the regulariser.

For the last term,

$$\begin{split} \left| \mathbf{e}_{s}^{\top} M(\alpha) \frac{\partial^{2} h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha^{2}} \right|_{\alpha=0} & \leq \| \mathbf{e}_{s}^{\top} \|_{1} \cdot \left\| M(\alpha) \frac{\partial^{2} h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha^{2}} \right|_{\alpha=0} \right\|_{\infty} \\ & \leq \frac{1}{1-\gamma} \cdot \left\| \frac{\partial^{2} h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha^{2}} \right|_{\alpha=0} \right\|_{\infty} \\ & \leq \frac{3 \cdot (1 + \log |\mathcal{A}|)}{1-\gamma} \cdot \| \mathbf{u} \|_{2}^{2}. \end{split}$$

For the second last term,

$$\begin{split} \left| \mathbf{e}_{s}^{\top} M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) \frac{\partial h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha} \right|_{\alpha=0} & \leq \left\| M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) \frac{\partial h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha} \right|_{\alpha=0} \right\|_{\infty} \\ & \leq \frac{1}{1-\gamma} \cdot \left\| \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) \frac{\partial h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha} \right|_{\alpha=0} \right\|_{\infty} \\ & \leq \frac{(C_{1} + T_{1}) \cdot \|\mathbf{u}\|_{2}}{1-\gamma} \cdot \left\| M(\alpha) \frac{\partial h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha} \right|_{\alpha=0} \right\|_{\infty} \\ & \leq \frac{(C_{1} + T_{1}) \cdot \|\mathbf{u}\|_{2}}{(1-\gamma)^{2}} \cdot \left\| \frac{\partial h_{\boldsymbol{\theta}_{\alpha}}}{\partial \alpha} \right|_{\alpha=0} \right\|_{\infty} \\ & \leq \frac{2 \cdot \log |\mathcal{A}|}{(1-\gamma)^{2}} (C_{1} + T_{1}) \cdot \|\mathbf{u}\|_{2}^{2}. \end{split}$$

For the second term,

$$\begin{aligned} \left| \mathbf{e}_{s}^{\top} M(\alpha) \frac{\partial^{2} \mathbf{P}(\alpha)}{\partial \alpha^{2}} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right|_{\alpha=0} & \leq \left\| M(\alpha) \frac{\partial^{2} \mathbf{P}(\alpha)}{\partial \alpha^{2}} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right|_{\alpha=0} \right\|_{\infty} \\ & \leq \frac{1}{1-\gamma} \cdot \left\| \frac{\partial^{2} \mathbf{P}(\alpha)}{\partial \alpha^{2}} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right|_{\alpha=0} \right\|_{\infty} \\ & \leq \frac{\|\mathbf{u}\|_{2}^{2}}{1-\gamma} \cdot \left\| M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right|_{\alpha=0} \left\|_{\infty} (C_{2} + 2C_{1}T_{1} + T_{2}) \right. \\ & \leq \frac{\|\mathbf{u}\|_{2}^{2}}{(1-\gamma)^{2}} \cdot \left\| h_{\boldsymbol{\theta}_{\alpha}} \right|_{\alpha=0} \left\|_{\infty} (C_{2} + 2C_{1}T_{1} + T_{2}) \right. \\ & \leq \frac{\log |\mathcal{A}|}{(1-\gamma)^{2}} (C_{2} + 2C_{1}T_{1} + T_{2}) \cdot \|\mathbf{u}\|_{2}^{2}. \end{aligned}$$

For the first term,

$$\left|\mathbf{e}_s^\top M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right|_{\alpha = 0} \right| \leq \left\| M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right|_{\alpha = 0} \left\|_{\infty} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \right\|_{\infty} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} M(\alpha) h_{\boldsymbol{\theta}_{\alpha}} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha = 0} \left\| \mathbf{P}(\alpha) \frac{\partial \mathbf{P}(\alpha)}{\partial \alpha} \|_{\alpha = 0} \right\|_{\alpha =$$

$$\leq \frac{1}{1-\gamma} \cdot \|\mathbf{u}\|_2 \cdot \frac{1}{1-\gamma} \cdot \|\mathbf{u}\|_2 \cdot \frac{1}{1-\gamma} \cdot \log |\mathcal{A}| \cdot (C_1 + T_1)^2$$
$$= \frac{\log |\mathcal{A}|}{(1-\gamma)^3} (C_1 + T_1)^2 \cdot \|\mathbf{u}\|_2^2.$$

Step 3: Now combining all the above equations, we get the final expression,

$$\max_{\|\mathbf{u}\|_{2}=1} \left\| \frac{\partial^{2} \mathcal{H}_{\boldsymbol{\pi}_{\boldsymbol{\theta}_{\alpha}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}_{\alpha}}}(s)}{\partial \alpha^{2}} \right|_{\alpha=0} \leq \beta_{\lambda}$$

where

$$\beta_{\lambda} = 2\gamma^{2} \cdot \frac{3 \cdot (1 + \log|\mathcal{A}|)}{1 - \gamma} + \gamma \cdot \frac{2 \cdot \log|\mathcal{A}|}{(1 - \gamma)^{2}} (C_{1} + T_{1})$$
$$+ 2\gamma \cdot \frac{\log|\mathcal{A}|}{(1 - \gamma)^{2}} (C_{2} + 2C_{1}T_{1} + T_{2}) + \frac{\log|\mathcal{A}|}{(1 - \gamma)^{3}} (C_{1} + T_{1})^{2}$$

By definition of smoothness, the "soft performative value function" \tilde{V}_{π}^{π} is Lipschitz smooth with Lipschitz constant L_{λ} where $L_{\lambda} \triangleq L + \beta_{\lambda}$. Once again, we can choose C_1, C_2, T_1, T_2 according to Corollary 1 for simplification to get the order $\beta_{\lambda} = \Omega\left(\frac{\log |\mathcal{A}|}{(1-\gamma)^3}\right)$. Thus, the final bound for L_{λ} as

$$L_{\lambda} = \max \{ \Omega(L), \Omega(\beta_{\lambda}) \} = \Omega\left(\frac{|\mathcal{A}|}{(1-\gamma)^2}\right)$$

F DERIVATION OF PERFORMATIVE POLICY GRADIENTS

Proof of Theorem 2. We prove each part of this theorem separately i.e. first we derive explicit closed form gradient for unregularised performative value function.

Proof of part (a). Let us remind the definition of the regularised objective function,

$$f_{\theta}(\tau) = \sum_{t=0}^{\infty} \gamma^t r_{\pi_{\theta}}(s_t, a_t)$$

Thus, $\nabla_{\theta} V_{\pi_{\theta}}^{\pi_{\theta}}(\tau) = \nabla_{\theta} \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}}[f_{\theta}(\tau)]$. Further,

$$\begin{split} \nabla_{\theta} \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}}[f_{\theta}(\tau)] &= \nabla_{\theta} \sum_{\tau} \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau) f_{\theta}(\tau) \\ &= \sum_{\tau} \nabla_{\theta} (\mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau) f_{\theta}(\tau)) \\ &= \sum_{\tau} (\nabla_{\theta} \, \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau)) f_{\theta}(\tau) + \sum_{\tau} \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau) (\nabla_{\theta} f_{\theta}(\tau)) \\ &\stackrel{\text{(a)}}{=} \sum_{\tau} \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau) (\nabla_{\theta} \log \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau)) f_{\theta}(\tau) + \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}}[\nabla_{\theta} f_{\theta}(\tau)] \\ &= \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[(\nabla_{\theta} \log \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau)) f_{\theta}(\tau) \right] + \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}}[\nabla_{\theta} f_{\theta}(\tau)] \end{split}$$

where (a) holds since $\nabla_{\theta} \log \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau) = \frac{\nabla_{\theta} \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau)}{\mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau)}$

We build a sample-based estimator for $\nabla_{\theta} \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}}[f_{\theta}(\tau)]$. Let $\tau^{(1)}, \ldots, \tau^{(n)}$ be n empirical samples from $\mathbf{P}_{\pi_{\theta}}^{\pi_{\theta}}$ (which are obtained by playing the policy π_{θ} for n times, with T steps for each run). Thus, we estimate the gradient by,

$$\mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[(\nabla_{\theta} \log \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau)) f_{\theta}(\tau) \right] \approx \frac{1}{n} \sum_{i=1}^{n} (\nabla_{\theta} \log \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}(\tau^{(i)})) f_{\theta}(\tau^{(i)})$$
(24)

Here recall that ρ is used to denote the density of the distribution of s_0 . Hence, we further have,

$$\log \mathbb{P}_{\boldsymbol{\pi_{\theta}}}^{\boldsymbol{\pi_{\theta}}}(\tau) = \log \rho(s_0) + \log \boldsymbol{\pi_{\theta}}(a_0 \mid s_0) + \log \mathbf{P}_{\boldsymbol{\pi_{\theta}}}(s_1 \mid s_0, a_0) + \log \boldsymbol{\pi_{\theta}}(a_1 \mid s_1) + \log \mathbf{P}_{\boldsymbol{\pi_{\theta}}}(s_2 \mid s_1, a_1) + \dots + \log \mathbf{P}_{\boldsymbol{\pi_{\theta}}}(s_t \mid s_{T-1}, a_{T-1})$$

Taking the gradient with respect to θ , we obtain

$$\nabla_{\boldsymbol{\theta}} \log \mathbb{P}_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}^{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(\tau) = \nabla_{\boldsymbol{\theta}} \log \boldsymbol{\pi_{\boldsymbol{\theta}}}(a_0 \mid s_0) + \nabla_{\boldsymbol{\theta}} \log \boldsymbol{\pi_{\boldsymbol{\theta}}}(a_1 \mid s_1) + \dots + \nabla_{\boldsymbol{\theta}} \log \boldsymbol{\pi_{\boldsymbol{\theta}}}(a_{T-1} \mid s_{T-1}) + \nabla_{\boldsymbol{\theta}} \log \mathbf{P}_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(s_1 \mid s_0, a_0) + \nabla_{\boldsymbol{\theta}} \log \mathbf{P}_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(s_2 \mid s_1, a_1) \dots + \nabla_{\boldsymbol{\theta}} \log \mathbf{P}_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(s_T \mid s_{T-1}, a_{T-1})$$

Hence, substituting the value of $\nabla_{\theta} \log(\mathbf{P}_{\pi_{\theta}}^{\pi_{\theta}})$ in equation (24) we get,

$$\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\tau) = \nabla_{\boldsymbol{\theta}} \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \left[f_{\boldsymbol{\theta}}(\tau) \right] = \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \left[\left(\sum_{t=0}^{\infty} \nabla_{\boldsymbol{\theta}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}}(a_t \mid s_t) \right) \cdot \left(\sum_{t=0}^{\infty} \gamma^t r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s_t, a_t) \right) \right] \\ + \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \left[\left(\sum_{t=1}^{T} \nabla_{\boldsymbol{\theta}} \log \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s_t | s_{t-1}, a_{t-1}) \right) \cdot \left(\sum_{t=0}^{\infty} \gamma^t r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s_t, a_t) \right) \right] \\ + \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \left[\sum_{t=0}^{\infty} \gamma^t \nabla_{\boldsymbol{\theta}} r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s_t, a_t) \right]$$

$$= \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \left[\sum_{t=0}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s_{t}, a_{t}) \nabla_{\boldsymbol{\theta}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}}(a_{t} \mid s_{t}) \right]$$

$$+ \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \left[\sum_{t=1}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s_{t}, a_{t}) \nabla_{\boldsymbol{\theta}} \log \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s_{t} \mid s_{t-1}, a_{t-1}) \right]$$

$$+ \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \nabla_{\boldsymbol{\theta}} r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s_{t}, a_{t}) \right] .$$

Hence, we conclude the proof for part (a) of the theorem.

Now we move on to the proof of part (b), i.e., deriving explicit gradient form for entropy-regularised value function

Proof of part (b). Let $\tilde{r}_{\pi_{\theta}}(s_t, a_t) = r_{\pi_{\theta}}(s_t, a_t) - \lambda \log \pi_{\theta}(a_t|s_t)$. Again we start from the definition of regularised objective function.

$$f_{\theta}(\tau) = \sum_{t=0}^{\infty} \gamma^t \tilde{r}_{\pi_{\theta}}(s_t, a_t)$$

So,
$$\nabla_{\pmb{\theta}} \tilde{V}^{\pi_{\pmb{\theta}}}_{\pi_{\pmb{\theta}}}(\tau) = \nabla_{\pmb{\theta}} \mathbb{E}_{\tau \sim \mathbb{P}^{\pi_{\pmb{\theta}}}_{\pi_{\pmb{\theta}}}}[f_{\pmb{\theta}}(\tau)]$$

Therefore, Equation (24) yields,

$$\begin{split} \nabla_{\theta}V_{\pi_{\theta}}^{\pi_{\theta}}(\tau) &= \nabla_{\theta}\mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\left(\sum_{t=0}^{\infty} \nabla_{\theta} \log \pi_{\theta}(a_{t} \mid s_{t}) \right) \cdot \left(\sum_{t=0}^{\infty} \gamma^{t} \tilde{r}_{\pi_{\theta}}(s_{t}, a_{t}) \right) \right] \\ &+ \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\left(\sum_{t=1}^{\infty} \nabla_{\theta} \log \mathbf{P}_{\pi_{\theta}}(s_{t} | s_{t-1}, a_{t-1}) \right) \cdot \left(\sum_{t=0}^{\infty} \gamma^{t} \tilde{r}_{\pi_{\theta}}(s_{t}, a_{t}) \right) \right] \\ &+ \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s_{t}, a_{t}) \right] \\ &= \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\sum_{t=1}^{\infty} \gamma^{t} \tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s_{t}, a_{t}) \nabla_{\theta} \log \mathbf{P}_{\pi_{\theta}}(s_{t} | s_{t-1}, a_{t-1}) \right] + \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s_{t}, a_{t}) \nabla_{\theta} \log \mathbf{P}_{\pi_{\theta}}(s_{t} | s_{t-1}, a_{t-1}) \right] + \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s_{t}, a_{t}) \nabla_{\theta} \log \mathbf{P}_{\pi_{\theta}}(s_{t} | s_{t-1}, a_{t-1}) \right] \\ &+ \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s_{t}, a_{t}) \nabla_{\theta} \log \mathbf{P}_{\pi_{\theta}}(s_{t} | s_{t-1}, a_{t-1}) \right] \\ &+ \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s_{t}, a_{t}) \nabla_{\theta} \log \mathbf{P}_{\pi_{\theta}}(s_{t} | s_{t-1}, a_{t-1}) \right] \\ &+ \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \nabla_{\theta} r_{\pi_{\theta}}(s_{t}, a_{t}) \right] \\ &- \lambda \mathbb{E}_{\tau \sim \mathbb{P}_{\pi_{\theta}}^{\pi_{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \nabla_{\theta} \log \pi_{\theta}(a_{t} | s_{t}) \right] \end{aligned}$$

Hence, we conclude proof of part (b).

Combining the results found in part (a) and (b) we get the final expressions in Theorem 2.

G CONVERGENCE OF PePG: PROOFS OF SECTION 4

G.1 Proofs for Unregularised Value Function

Lemma 6 (Performative Policy Gradient for Softmax PeMDPs). *Given softmax PeMDPs defined by* (11), *for all* $(s, a, s') \in (S, A, S)$, *derivative of the performative value function w.r.t* $\theta_{s,a}$ *satisfies:*

(a) For unregularised objective:

$$\frac{\partial V_{\pi_{\theta}}^{\pi_{\theta}}(\rho)}{\partial \boldsymbol{\theta}_{s,a}} \ge \frac{1}{1 - \gamma} d_{\pi_{\theta}}^{\pi_{\theta}}(s, a|\rho) \left(A_{\pi_{\theta}}^{\pi_{\theta}}(s, a)(1 + C_{\psi}) + \xi \right) \tag{25}$$

Proof.] Notice that:

$$\frac{\partial}{\partial \boldsymbol{\theta}_{s',a'}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}}(a|s) = \mathbb{1}[s = s', a = a'] - \boldsymbol{\pi}_{\boldsymbol{\theta}}(a'|s)\mathbb{1}[s = s']$$

$$\frac{\partial}{\partial \boldsymbol{\theta}_{s',a'}} \log \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s''|s, a) = \psi(s'')\mathbb{1}[s = s', a = a'] \left(1 - \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s''|s, a)\right)$$

$$\frac{\partial}{\partial \boldsymbol{\theta}_{s',a'}} r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s, a) = \xi \mathbb{1}[s = s', a = a'].$$
(26)

Now, we get from Theorem 2,

$$\begin{split} \frac{\partial}{\partial \boldsymbol{\theta}_{s,a}} V_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\sigma}}(\tau) &= \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\sigma}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \left(A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\sigma}}(s_{t}, a_{t}) \frac{\partial}{\partial \boldsymbol{\theta}_{s,a}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}}(a_{t} \mid s_{t}) \right. \\ &\quad + A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\sigma}}(s_{t}, a_{t}) \frac{\partial}{\partial \boldsymbol{\theta}_{s,a}} \log \boldsymbol{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t+1} \mid s_{t}, a_{t}) \\ &\quad + \frac{\partial}{\partial \boldsymbol{\theta}_{s,a}} r_{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \right) \right] \\ &= \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \left(A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\sigma}}(s_{t}, a_{t}) \left(\mathbb{I}[s_{t} = s, a_{t} = a] - \boldsymbol{\pi}\boldsymbol{\theta}(a \mid s) \mathbb{I}[s_{t} = s] \right) \right. \\ &\quad + A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \psi(s_{t+1}) \mathbb{I}[s_{t} = s, a_{t} = a] \left. \left(1 - \mathbf{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t+1} \mid s, a) \right) \right. \\ &\quad + \xi \mathbb{I}[s_{t} = s, a_{t} = a] \right) \right] \\ &= \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \mathbb{I}[s_{t} = s, a_{t} = a] \right] - \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} \boldsymbol{\pi}\boldsymbol{\theta}(a \mid s) \mathbb{I}[s_{t} = s] A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \right. \\ &\quad + \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \psi(s_{t+1}) \mathbb{I}[s_{t} = s, a_{t} = a] \right] \\ &\quad - \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \psi(s_{t+1}) \mathbb{I}[s_{t} = s, a_{t} = a] \mathbf{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t+1} \mid s, a) \right] \\ &\quad + \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \psi(s_{t+1}) \mathbb{I}[s_{t} = s, a_{t} = a] \mathbf{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t+1} \mid s, a) \right] \\ &\quad + \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \psi(s_{t+1}) \mathbb{I}[s_{t} = s, a_{t} = a] \right] \\ &\quad - \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \psi(s_{t+1}) \mathbb{I}[s_{t} = s, a_{t} = a] \right] \\ &\quad - \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\pi}\boldsymbol{\theta}}(s_{t}, a_{t}) \psi(s_{t+1}) \mathbb{I}[s_{t} = s, a_{t} = a] \right] \\ &\quad - \mathbb{E}_{\tau \sim \mathbb{P}_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^{t} A_{\boldsymbol{\pi}\boldsymbol{\theta}}^{\boldsymbol{\theta}}(s_{t}, a_{t}) \psi(s_{t+1}) \mathbb{I}[s_{t} = s, a_{t} = a] \right]$$

Finally, for some initial state distribution ρ , the gradient can be expressed as:

$$\frac{\partial V_{\pi_{\theta}}^{\pi_{\theta}}(\rho)}{\partial \theta_{s,a}} \ge \frac{1}{1-\gamma} d_{\pi_{\theta}}^{\pi_{\theta}}(s,a|\rho) \left(A_{\pi_{\theta}}^{\pi_{\theta}}(s,a)(1+C_{\psi}) + \xi \right) \tag{27}$$

Proof of Lemma 3– Part (a). We know by Lemma 1 that

$$\begin{split} V_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\rho) - V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\rho) &= \frac{1}{1 - \gamma} \mathbb{E}_{(s, a) \sim d_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{o}^{\star}}(\cdot|\rho)} [A_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(s, a)] \\ &+ \frac{1}{1 - \gamma} \mathbb{E}_{(s, a) \sim d_{\boldsymbol{\pi}, \rho}^{\boldsymbol{\pi}}} [(r_{\boldsymbol{\pi}}(s, a) - r_{\boldsymbol{\pi}^{\prime}}(s, a)) + \gamma (\mathbf{P}_{\boldsymbol{\pi}}(\cdot|s, a) - \mathbf{P}_{\boldsymbol{\pi}^{\prime}}(\cdot|s, a))^{\top} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot)] \; . \end{split}$$

Step 1: Upper bounding Term 1.

$$\mathbb{E}_{(s,a)\sim d_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{\star}}(\cdot|\rho)}[A_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a)] = \sum_{s,a} d_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{\star}}(s,a|\rho) A_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a)$$

$$= \sum_{s,a} \frac{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{\star}}(s,a|\rho)}{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a|\nu)} d_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a|\nu) A_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a)$$

$$\leq \left\| \frac{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}},\rho}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{\star}}}{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}},\nu}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \right\| \sum_{s,a} d_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a|\nu) A_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a)$$

$$(28)$$

Now, we leverage the gradient of softmax performative MDPs to obtain

$$\begin{split} \sum_{s,a} d^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a|\nu) A^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a) &\leq \frac{1-\gamma}{1+C_{\psi}} \sum_{s,a} \frac{\partial V^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\nu)}{\partial \boldsymbol{\theta}_{s,a}} - \frac{1}{1+C_{\psi}} \xi \\ &= \frac{1-\gamma}{1+C_{\psi}} \mathbf{1}^{\top} \nabla_{\boldsymbol{\theta}} V^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\nu) - \frac{1}{1+C_{\psi}} \xi \\ &= \sqrt{|\mathcal{S}||\mathcal{A}|} \frac{1-\gamma}{1+C_{\psi}} \frac{1}{\sqrt{|\mathcal{S}||\mathcal{A}|}} \mathbf{1}^{\top} \nabla_{\boldsymbol{\theta}} V^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\nu) - \frac{1}{1+C_{\psi}} \xi \\ &\leq \sqrt{|\mathcal{S}||\mathcal{A}|} \frac{1-\gamma}{1+C_{\psi}} \|\nabla_{\boldsymbol{\theta}} V^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\nu)\|_{2} - \frac{1}{1+C_{\psi}} \xi \end{split}$$

Hence, substituting it back to (28) and using the fact that $\left\| \frac{d_{\pi\theta,\rho}^{\pi^*}}{d_{\pi\theta,\nu}^{\theta}} \right\|_{\infty} \ge 1$ (from lemma 10) we get,

$$\frac{1}{1-\gamma} \mathbb{E}_{(s,a) \sim d_{\pi}^{\pi}(\cdot|s_0)} [A_{\pi'}^{\pi'}(s,a)] \leq \frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{1+C_{\psi}} \left\| \frac{d_{\pi_{\theta},\rho}^{\pi_{\theta}^{*}}}{d_{\pi_{\theta},\nu}^{\pi_{\theta}}} \right\|_{\infty} \|\nabla_{\theta} V_{\pi_{\theta}}^{\pi_{\theta}}(\nu)\|_{2} - \frac{\xi}{(1+C_{\psi})(1-\gamma)}$$

Step 2: Upper bounding Term 2. For Lipschitz rewards and transitions, we further get that

Term 2 =
$$\frac{1}{1-\gamma} \mathbb{E}_{(s,a)\sim d_{\boldsymbol{\pi},\rho}^{\boldsymbol{\pi}}} \left[(r_{\boldsymbol{\pi}}(s,a) - r_{\boldsymbol{\pi}'}(s,a)) + \gamma (\mathbf{P}_{\boldsymbol{\pi}}(\cdot|s,a) - \mathbf{P}_{\boldsymbol{\pi}'}(\cdot|s,a))^{\top} V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot) \right].$$

$$\leq \frac{2\sqrt{2}}{1-\gamma} (L_r + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} R_{\max}) D_{\mathbf{H}} \left(\boldsymbol{\pi}_o^{\star}(s_0) \| \boldsymbol{\pi}_{\boldsymbol{\theta}}(s_0) \right)$$

Proof of Theorem 3– Part (a). Step 1: By the L-smoothness of $V_{\pi_{\theta}}^{\pi_{\theta}}$ we get,

$$|V_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}^{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(\rho) - V_{\boldsymbol{\pi_{\boldsymbol{\theta}}'}}^{\boldsymbol{\pi_{\boldsymbol{\theta}}'}}(\rho) - \langle \nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi_{\boldsymbol{\theta}}}}^{\boldsymbol{\pi_{\boldsymbol{\theta}}}}(\rho), \boldsymbol{\theta} - \boldsymbol{\theta}' \rangle| \leq \frac{L}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2$$

1513 Thus,

$$|V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t+1)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t+1)}}(\rho) - V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho) - \eta \|\nabla_{\boldsymbol{\theta}}V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho)\|^2| \leq \frac{L}{2}\|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_{t}\|^2$$

This implies that

$$\begin{split} V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t+1)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t+1)}}(\rho) - V_{\boldsymbol{\pi}_{\boldsymbol{\sigma}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho) &\geq \eta \|\nabla V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho)\|^{2} - \frac{L}{2} \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_{t}\|^{2} \\ \Longrightarrow V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t+1)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t+1)}}(\rho) - V_{\boldsymbol{\pi}_{\boldsymbol{\sigma}}^{\star}}^{\boldsymbol{\pi}_{\boldsymbol{\sigma}}^{\star}}(\rho) &\geq V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho) - V_{\boldsymbol{\pi}_{\boldsymbol{\sigma}}^{\star}}^{\boldsymbol{\pi}_{\boldsymbol{\sigma}}^{\star}}(\rho) + \eta \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho)\|^{2} - \frac{L}{2} \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_{t}\|^{2} \\ &= V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho) - V_{\boldsymbol{\pi}_{\boldsymbol{\sigma}}^{\star}}^{\boldsymbol{\pi}_{\boldsymbol{\sigma}}^{(t)}}(\rho) + \eta (1 - \frac{L\eta}{2}) \|\nabla V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho)\|^{2} \end{split}$$

Define
$$r_t \triangleq \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho)\|_2$$
 and $\delta_t \triangleq V_{\boldsymbol{\pi}_{\boldsymbol{\phi}}^{\star}}^{\boldsymbol{\pi}_{\boldsymbol{\sigma}}^{\star}}(\rho) - V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}(\rho)$

$$\eta (1 - \frac{L\eta}{2})r_t^2 \le \delta_t - \delta_{t+1}$$

$$(1 - \frac{L\eta}{2})w_t r_t^2 \le \frac{w_t}{\eta} (\delta_t - \delta_{t+1}) \le \frac{w_{t-1}}{\eta} \delta_t - \frac{w_t}{\eta} \delta_{t+1}$$

Now telescoping over $t = 0, 1, \dots, T - 1$ we get

$$(1 - \frac{L\eta}{2}) \sum_{t} w_{t} r_{t}^{2} \leq \frac{w_{-1}}{\eta} \delta_{0} - \frac{w_{T-1}}{\eta} \delta_{T}$$

$$(1 - \frac{L\eta}{2}) \min_{t} r_{t}^{2} \leq \frac{w_{-1}}{\eta(\sum_{t} w_{t})} \delta_{0} - \frac{w_{T-1}}{\eta(\sum_{t} w_{t})} \delta_{T} \leq \frac{1}{\eta T} \delta_{0}$$

$$\min_{t} r_{t}^{2} \leq \frac{1}{\eta T (1 - \frac{L\eta}{2})} \delta_{0}$$
(29)

Step 2: Now, we know that

$$(V_{\boldsymbol{\pi}_{o}^{\star}}^{\star}(\rho) - V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\rho))^{2} \leq \left(\frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{1 + C_{\psi}} \left\| \frac{d_{\boldsymbol{\pi}_{\theta}, \rho}^{\boldsymbol{\pi}_{o}^{\star}}}{d_{\boldsymbol{\pi}_{\theta}, \nu}^{\boldsymbol{\pi}_{\theta}}} \right\|_{\infty} \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\nu)\|_{2} - \frac{\xi}{(1 + C_{\psi})(1 - \gamma)} + \frac{2\sqrt{2}}{1 - \gamma} (L_{r} + \frac{\gamma}{1 - \gamma} L_{\mathbf{P}} R_{\max}) D_{\mathbf{H}} (\boldsymbol{\pi}_{o}^{\star}(\rho) \| \boldsymbol{\pi}_{\theta}(\rho)) \right)^{2} \\ \leq 2 \left(\frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{1 + C_{\psi}} \left\| \frac{d_{\boldsymbol{\pi}_{\theta}, \rho}^{\boldsymbol{\pi}_{o}^{\star}}}{d_{\boldsymbol{\pi}_{\theta}, \nu}^{\boldsymbol{\pi}_{\theta}}} \right\|_{\infty} \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\nu)\|_{2} \right)^{2} + 4 \left(\frac{\xi}{(1 + C_{\psi})(1 - \gamma)}\right)^{2} \\ + 32 \left(\frac{1}{1 - \gamma} (L_{r} + \frac{\gamma}{1 - \gamma} L_{\mathbf{P}} R_{\max}) D_{\mathbf{H}} (\boldsymbol{\pi}_{o}^{\star}(\rho) \| \boldsymbol{\pi}_{\theta}(\rho)) \right)^{2}$$

This implies that

$$2\left(\frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{1+C_{\psi}}\left\|\frac{d_{\pi_{\theta},\rho}^{\pi_{\sigma}^{\star}}}{d_{\pi_{\theta},\nu}^{\star}}\right\|_{\infty}\left\|\nabla_{\theta}V_{\pi_{\theta}}^{\pi_{\theta}}(\nu)\right\|_{2}\right)^{2} \geq (V_{\pi_{\sigma}^{\star}}^{\star}(\rho) - V_{\pi_{\theta}}^{\star}(\rho))^{2} - 4\left(\frac{\xi}{(1+C_{\psi})(1-\gamma)}\right)^{2}$$

$$-32\left(\frac{1}{1-\gamma}(L_{r} + \frac{\gamma}{1-\gamma}L_{\mathbf{P}}R_{\max})D_{\mathbf{H}}\left(\pi_{\sigma}^{\star}(\rho)\right)\|\pi_{\theta}(\rho)\right)\right)^{2}$$

$$\Rightarrow \frac{|\mathcal{S}||\mathcal{A}|}{(1+C_{\psi})^{2}}\left\|\frac{d_{\pi_{\theta},\rho}^{\star}}{d_{\pi_{\theta},\nu}^{\star}}\right\|_{\infty}^{2}\left\|\nabla_{\theta}V_{\pi_{\theta}}^{\pi_{\theta}}(\nu)\right\|_{2}^{2} \geq \frac{1}{2}(V_{\pi_{\sigma}^{\star}}^{\star}(\rho) - V_{\pi_{\theta}}^{\pi_{\theta}}(\rho))^{2} - 2\left(\frac{\xi}{(1+C_{\psi})(1-\gamma)}\right)^{2}$$

$$-16\frac{1}{(1-\gamma)^{2}}\left(L_{r} + \frac{\gamma}{1-\gamma}L_{\mathbf{P}}R_{\max}\right)^{2}$$

$$\implies \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\boldsymbol{\nu})\|_{2}^{2} \ge \frac{(1+C_{\psi})^{2}}{|\mathcal{S}||\mathcal{A}|} \frac{1}{\left\|\frac{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}},\boldsymbol{\rho}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}}{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}},\boldsymbol{\nu}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}}\right\|_{\infty}^{2}} \frac{1}{2} (V_{\boldsymbol{\pi}_{\boldsymbol{\delta}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}}^{\boldsymbol{\pi}_{\boldsymbol{\delta}}^{\boldsymbol{\star}}}(\boldsymbol{\rho}) - V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\boldsymbol{\rho}))^{2} - 2 \frac{1}{|\mathcal{S}||\mathcal{A}|} \frac{\xi^{2}}{\left\|\frac{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}},\boldsymbol{\rho}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}}{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}},\boldsymbol{\nu}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}}\right\|_{\infty}^{2}} - 16 \frac{(1+C_{\psi})^{2}}{|\mathcal{S}||\mathcal{A}|} \frac{1}{\left\|\frac{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}},\boldsymbol{\rho}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}}{d_{\boldsymbol{\pi}_{\boldsymbol{\theta}},\boldsymbol{\nu}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}}\right\|_{\infty}^{2}} \frac{1}{(1-\gamma)^{2}} \left(L_{r} + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} R_{\max}\right)^{2}$$

Step 3: Now, replacing this in the iterative upper bound, we get

$$\frac{(1+C_{\psi})^{2}}{|\mathcal{S}||\mathcal{A}|} \frac{1}{d_{\pi_{\theta},\nu}^{\star^{*}}} \Big\|_{\infty}^{2} \frac{1}{2} (V_{\pi_{\theta}}^{\star^{*}}(\rho) - V_{\pi_{\theta}}^{\pi_{\theta}}(\rho))^{2} \leq 2 \frac{1}{|\mathcal{S}||\mathcal{A}|} \frac{\xi^{2}}{d_{\pi_{\theta},\nu}^{\star^{*}}} \Big\|_{\infty}^{2} \frac{\xi^{2}}{(1-\gamma)^{2}} + 16 \frac{(1+C_{\psi})^{2}}{|\mathcal{S}||\mathcal{A}|} \frac{1}{d_{\pi_{\theta},\nu}^{\star^{*}}} \Big\|_{\infty}^{2} \frac{1}{(1-\gamma)^{2}} \Big(L_{r} + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} R_{\max}\Big)^{2} + \frac{1}{\eta T (1 - \frac{L\eta}{2})} \delta_{0}$$

$$(V_{\pi_{\theta}}^{\star^{*}}(\rho) - V_{\pi_{\theta}}^{\pi_{\theta}}(\rho))^{2} \leq \frac{4\xi^{2}}{(1+C_{\psi})^{2} (1-\gamma)^{2}} + \frac{32}{(1-\gamma)^{2}} \Big(L_{r} + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} R_{\max}\Big)^{2} + \frac{2\delta_{0}}{\eta T (1 - \frac{L\eta}{2})} \frac{|\mathcal{S}||\mathcal{A}|}{d_{\pi_{\theta},\nu}^{\pi_{\theta},\nu}} \Big\|_{\infty}^{2}$$

$$(30)$$

Step 4: Now, we solve for T

$$\begin{split} \min_{t \leq T} (V_{\pmb{\pi}_{o}^{\star}}^{\star}(\rho) - V_{\pmb{\pi}_{\theta}^{(t)}}^{\pmb{\pi}_{\theta}^{(t)}}(\rho))^{2} &\leq \frac{4\xi^{2}}{(1 + C_{\psi})^{2}(1 - \gamma)^{2}} \\ &+ \frac{32}{(1 - \gamma)^{2}} \Big(L_{r} + \frac{\gamma}{1 - \gamma} L_{\mathbf{P}} R_{\max}\Big)^{2} + \frac{2\delta_{0}}{\eta T (1 - \frac{L\eta}{2})} \frac{|\mathcal{S}||\mathcal{A}|}{(1 + C_{\psi})^{2}} \frac{d_{\pmb{\pi}_{\theta}, \rho}^{\pi_{\delta}}}{d_{\pmb{\pi}_{\theta}, \nu}^{\pi_{\theta}}} \Big\|_{\infty}^{2} \leq \epsilon^{2} \end{split}$$

Hence, solving for T we get,

$$T \geq \frac{2\delta_0|\mathcal{S}||\mathcal{A}| \left\| \frac{d_{\pi_{\theta}, \rho}^{\pi_{\theta}^{\star}}}{d_{\pi_{\theta}, \nu}^{\pi_{\theta}^{\star}}} \right\|_{\infty}^{2}}{\eta(1 - \frac{L\eta}{2})(1 + C_{\psi})^{2} \left[\epsilon^{2} - \frac{4\xi^{2}}{(1 + C_{\psi})^{2}(1 - \gamma)^{2}} - \frac{32}{(1 - \gamma)^{2}} \left(L_{r} + \frac{\gamma}{1 - \gamma} L_{\mathbf{P}} R_{\max} \right)^{2} \right]}$$

Choosing $\eta = \frac{1}{2L}$ and using the fact that $\delta_0 \leq \frac{2R_{\max}}{1-\gamma}$ we get the final expression,

$$T \ge \frac{32R_{\max}L|\mathcal{S}||\mathcal{A}|(1-\gamma)}{3\epsilon^{2}(1-\gamma)^{2}(1+C_{\psi})^{2}-12\xi^{2}-96(1+C_{\psi})^{2}(L_{p}+L_{r}R_{\max}\frac{\gamma}{1-\gamma})^{2}}{d_{\pi_{\theta},\nu}^{\pi_{\theta}}} \left\| \frac{d_{\pi_{\theta},\rho}^{\pi_{\phi}^{*}}}{d_{\pi_{\theta},\nu}^{\pi_{\theta}}} \right\|_{\infty}^{2}$$

This implies a non-trivial bound

$$\epsilon^2 > \frac{4\xi^2}{(1+C_{\psi})^2(1-\gamma)^2} + \frac{32}{(1-\gamma)^2} \left(L_r - \frac{\gamma}{1-\gamma}L_{\mathbf{P}}R_{\max}\right)^2$$

1620	$2\sqrt{2}\xi$ 8 (γ
1621	$\implies \epsilon > \frac{2\sqrt{2\xi}}{(1+C_{\psi})(1-\gamma)} + \frac{8}{(1-\gamma)} \left(L_r + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} R_{\max}\right)$
1622	
1623	$> rac{2\sqrt{2}}{(1+C_\psi)(1-\gamma)}\left(\xi+2\sqrt{2}(1+C_\psi)\Big(L_r+L_{\mathbf{P}}R_{\max}rac{\gamma}{1-\gamma}\Big) ight)$
1624	$(1+C_{\psi})(1-\gamma) \left(\frac{1}{2} + \frac{1}{2$
1625	

G.2 Proofs for Entropy-regularised or Soft Value Function

Lemma 7 (Regularized Performative Policy Difference: Generic Upper Bound). *Under Assumption 1*, the sub-optimality gap of a policy π_{θ} is

$$\tilde{V}_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(s_{0}) - \tilde{V}_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(s_{0}) \leq \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim d_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{o}^{\star}}(\cdot|s_{0})} [\tilde{A}_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(s,a)]
+ \frac{1}{1 - \gamma} \Big(L_{r} + \frac{\gamma}{1 - \gamma} L_{\mathbf{P}}(R_{\max} + \lambda \log|\mathcal{A}|) \Big) D_{\mathbf{H}} (\boldsymbol{\pi}_{o}^{\star}(\cdot|s_{0}) \|\boldsymbol{\pi}_{\theta}(\cdot|s_{0}))
- \frac{\lambda}{1 - \gamma} \mathbb{E}_{s \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot|s_{0})} D_{\mathbf{KL}} (\boldsymbol{\pi}_{o}^{\star}(\cdot|s_{0}) \|\boldsymbol{\pi}_{\theta}(\cdot|s_{0}))$$
(31)

Proof. Lemma 1 for regularized rewards reduces to,

$$\tilde{V}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s_0) - \tilde{V}_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}'}(s_0) = \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(\cdot|s_0)} [\tilde{A}_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}'}(s,a)]
+ \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot|s_0)} \Big([\tilde{r}_{\boldsymbol{\pi}}(s,a) - \tilde{r}_{\boldsymbol{\pi}'}(s,a)] + \gamma (\mathbf{P}_{\boldsymbol{\pi}} - \mathbf{P}_{\boldsymbol{\pi}'})^{\top} \tilde{V}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s_0) \Big).$$
(32)

Therefore,

$$\tilde{r}_{\boldsymbol{\pi}_{o}^{\star}}(s, a) - \tilde{r}_{\boldsymbol{\pi}}(s, a) = r_{\boldsymbol{\pi}_{o}^{\star}}(s, a) - r_{\boldsymbol{\pi}}(s, a) + \lambda \Big(\log \boldsymbol{\pi}(a|s) - \log \boldsymbol{\pi}_{o}^{\star}(a|s)\Big)$$

Therefore,

$$V_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s_{0}) - \tilde{V}_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}'}(s_{0}) = \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}}(\cdot|s_{0})} [\tilde{A}_{\boldsymbol{\pi}'}^{\boldsymbol{\pi}'}(s,a)]$$

$$+ \frac{1}{1 - \gamma} \mathbb{E}_{(s,a) \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot|s_{0})} \Big([\tilde{r}_{\boldsymbol{\pi}}(s,a) - \tilde{r}_{\boldsymbol{\pi}'}(s,a)] + \gamma (\mathbf{P}_{\boldsymbol{\pi}} - \mathbf{P}_{\boldsymbol{\pi}'})^{\top} \tilde{V}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s_{0}) \Big) .$$

$$+ \frac{\lambda}{1 - \gamma} \mathbb{E}_{s \sim \boldsymbol{d}_{\boldsymbol{\pi}^{\star}}^{\boldsymbol{\pi}^{\star}}(\cdot|s_{0})} \Big[\mathbb{E}_{a \sim \boldsymbol{\pi}^{\star}(\cdot|s_{0})} [\log \boldsymbol{\pi}(a|s) - \log \boldsymbol{\pi}_{o}^{\star}(a|s)] \Big]$$

$$(34)$$

Now,

$$\mathbb{E}_{a_0 \sim \boldsymbol{\pi}^{\star}(\cdot|s_0)}[\log \boldsymbol{\pi}(a|s) - \log \boldsymbol{\pi}^{\star}(a|s)] = -D_{\mathrm{KL}}(\boldsymbol{\pi}_o^{\star}(\cdot|s) \parallel \boldsymbol{\pi}(\cdot|s))$$

Following Mei et al. (2020),

$$\|\tilde{V}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}\|_{\infty} \leq \frac{R_{\max} + \lambda \log |\mathcal{A}|}{1 - \gamma}$$

A point to keep in mind is that the $\|\cdot\|_{\infty}$ -norm here is taken with respect to all starting states $s_0 \in \mathcal{S}$. Putting all these together we get our desired expression.

Lemma 8 (Regularized Performative Policy gradient for softmax policies and softmax MDPs). For a class of PeMDPs $\mathcal{M} \triangleq (\mathcal{S}, \mathcal{A}, \boldsymbol{\pi}, \mathbf{P}_{\boldsymbol{\pi}}, r_{\boldsymbol{\pi}}, \boldsymbol{\theta}, \rho)$ consider softmax parametrization for policy $\boldsymbol{\pi}_{\boldsymbol{\theta}} \in \Delta(\boldsymbol{\theta} \in \boldsymbol{\Theta})$ and transition dynamics $\mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}$ and linear parametrization for reward $r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}$. For all $(s, a, s') \in (\mathcal{S}, \mathcal{A}, \mathcal{S})$, derivative of the expected return w.r.t $\boldsymbol{\theta}_{s,a}$ satisfies:

$$\frac{\partial \tilde{V}_{\pi_{\theta}}^{\pi_{\theta}}(\rho)}{\partial \theta_{s,a}} \ge \frac{1}{1-\gamma} d_{\pi_{\theta}}^{\pi_{\theta}}(s,a|\rho) \left(\tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s,a)(1+C_{\psi}) + \xi \right) - \frac{\lambda}{1-\gamma} (1+\log|\mathcal{A}|)$$
(35)

where $\bar{\psi}(s') = \psi(s') - \mathbb{E}_{s'' \sim P_{\pi_{\theta}}} [\psi(s'')]$ and $C_{\psi} \triangleq \min_{s'} \bar{\psi}(s')$.

Proof. Notice that:

$$\frac{\partial}{\partial \boldsymbol{\theta}_{s',a'}} \log \boldsymbol{\pi}_{\boldsymbol{\theta}}(a|s) = \mathbb{1}[s = s', a = a'] - \boldsymbol{\pi}_{\boldsymbol{\theta}}(a'|s)\mathbb{1}[s = s']$$

$$\frac{\partial}{\partial \boldsymbol{\theta}_{s',a'}} \log \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s''|s, a) = \psi(s'')\mathbb{1}[s = s', a = a'] (1 - \mathbf{P}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s''|s, a))$$

$$\frac{\partial}{\partial \boldsymbol{\theta}_{s',a'}} r_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(s,a) = \xi \mathbb{1}[s=s',a=a']. \tag{36}$$

Now, we get from Theorem 2

Finally, for some initial state distribution s_0 , the gradient can be expressed as:

$$\frac{\partial \tilde{V}_{\pi_{\theta}}^{\pi_{\theta}}(\rho)}{\partial \theta_{s,a}} \ge \frac{1}{1-\gamma} d_{\pi_{\theta}}^{\pi_{\theta}}(s,a|\rho) \left(\tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s,a)(1+C_{\psi}) + \xi \right) - \frac{\lambda}{1-\gamma} \\
- \frac{\lambda}{1-\gamma} \sum_{a} d_{\pi_{\theta}}^{\pi_{\theta}}(s,a|s_{0}) \pi_{\theta}(a|s) \log \frac{1}{\pi_{\theta}(a|s)} \\
\ge \frac{1}{1-\gamma} d_{\pi_{\theta}}^{\pi_{\theta}}(s,a|\rho) \left(\tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s,a)(1+C_{\psi}) + \xi \right) - \frac{\lambda}{1-\gamma} (1+\log|\mathcal{A}|) \tag{37}$$

 inequality holds:

$$\tilde{V}_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\rho) - \tilde{V}_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\rho) \leq \frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{(1 + C_{\psi})} \left\| \frac{\boldsymbol{d}_{\boldsymbol{\pi}_{\theta}, \rho}^{\boldsymbol{\pi}_{o}}}{\boldsymbol{d}_{\boldsymbol{\pi}_{\theta}, \nu}^{\boldsymbol{\pi}_{\theta}}} \right\|_{\infty} \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\nu)\|_{2} - \frac{\xi}{(1 + C_{\psi})(1 - \gamma)} + \frac{1}{1 - \gamma} \left(L_{r} + \frac{\gamma}{1 - \gamma} L_{\mathbf{P}}(R_{\max} + \lambda \log |\mathcal{A}|) \right) + \frac{\lambda}{1 - \gamma} (1 + 2 \log |\mathcal{A}|) \tag{38}$$

Proof. We know,

$$-D_{\mathrm{KL}}\left(\boldsymbol{\pi}_{o}^{\star}(\cdot|s) \parallel \boldsymbol{\pi}_{\boldsymbol{\theta}}(\cdot|s)\right) \leq \log|\mathcal{A}|$$

Lemma 9 (Regularized Performative Gradient Domination: Part(b) of Lemma 3). For regularized PeMDPs the following

This follows from the fact that KL divergence is upper bounded by entropy which is further upper bounded by $\log |\mathcal{A}|$. Hence,

$$-\mathbb{E}_{s \sim \boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(\cdot|s_0)} D_{\mathrm{KL}} \left(\boldsymbol{\pi}_{o}^{\star}(\cdot|s) \parallel \boldsymbol{\pi}_{\boldsymbol{\theta}}(\cdot|s) \right) \leq \log |\mathcal{A}|$$

From (37), it is to see that following the exact steps as lemma 3, we can achieve

$$\sum_{s,a} d_{\pi_{\theta}}^{\pi_{\theta}}(s,a) \tilde{A}_{\pi_{\theta}}^{\pi_{\theta}}(s,a) \leq \sqrt{|\mathcal{S}||\mathcal{A}|} \frac{1-\gamma}{1+C_{\psi}} \|\nabla_{\theta} \tilde{V}_{\pi_{\theta}}^{\pi_{\theta}}(\nu)\|_{2} - \frac{\xi}{1+C_{\psi}} + \frac{\lambda}{1-\gamma} (2\log|\mathcal{A}|+1)$$

Therefore, using 31 we finally get,

$$\tilde{V}_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\rho) - \tilde{V}_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\rho) \leq \frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{(1+C_{\psi})} \left\| \frac{d_{\boldsymbol{\pi}_{\theta},\rho}^{\boldsymbol{\pi}_{\theta}}}{d_{\boldsymbol{\pi}_{\theta},\nu}^{\boldsymbol{\pi}_{\theta}}} \right\|_{\infty} \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\nu)\|_{2} - \frac{\xi}{(1+C_{\psi})(1-\gamma)} \frac{\lambda}{1-\gamma} (2\log|\mathcal{A}|+1) \\
+ \frac{1}{1-\gamma} \left(L_{r} + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} (R_{\max} + \lambda \log|\mathcal{A}|) \right) D_{\mathbf{H}} \left(\boldsymbol{\pi}_{o}^{\star}(\cdot|\rho) \| \boldsymbol{\pi}_{\theta}(\cdot|\rho) \right) \\
\leq \frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{(1+C_{\psi})} \left\| \frac{d_{\boldsymbol{\pi}_{\theta},\rho}^{\boldsymbol{\pi}_{o}}}{d_{\boldsymbol{\pi}_{\theta},\nu}^{\boldsymbol{\pi}_{\theta}}} \right\|_{\infty} \|\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\pi}_{\theta}}^{\boldsymbol{\pi}_{\theta}}(\nu)\|_{2} - \frac{\xi}{(1+C_{\psi})(1-\gamma)} + \frac{\lambda}{1-\gamma} (2\log|\mathcal{A}|+1) \\
+ \frac{1}{1-\gamma} \left(L_{r} + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} (R_{\max} + \lambda \log|\mathcal{A}|) \right) \tag{40}$$

Now substitute the values of L_r and L_P to get the desired inequality.

Theorem 4 (Convergence of Entropy-regularised PePG: Part(b) of Theorem 3).

Proof. Step 1: In Lemma 5 we proved that the soft-value function $\tilde{V}_{\pi\theta}^{\pi\theta}$ is L_{λ} -smooth. Combining this fact with (29) of Theorem 3 we get the following inequality:

$$\min_{t} r_t^2 \le \frac{1}{nT(1 - \frac{L_{\lambda}\eta}{\delta_0})} \tilde{\delta}_0$$

where
$$r_t \triangleq \|\nabla_{\boldsymbol{\theta}} \tilde{V}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(t)}}\|_2^2$$
 and $\tilde{\delta}_0 \triangleq \tilde{V}_{\boldsymbol{\pi}_o^{\star}}^{\boldsymbol{\pi}_o^{\star}}(\rho) - \tilde{V}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(0)}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{(0)}}(\rho) \leq \frac{2(R_{\max} + \log |\mathcal{A}|)}{1 - \gamma}$.

Again, we know,

$$(V_{\boldsymbol{\pi}_{o}^{\star}}^{\boldsymbol{\pi}_{o}^{\star}}(\rho) - V_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\rho))^{2} \leq \left(\frac{\sqrt{|\mathcal{S}||\mathcal{A}|}}{(1 + C_{\psi})} \left\| \frac{\boldsymbol{d}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}^{\star, \rho}}^{\boldsymbol{\pi}_{o}^{\star}}}{\boldsymbol{d}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}, \nu}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}} \right\|_{\infty} \|\nabla_{\boldsymbol{\theta}} \tilde{V}_{\boldsymbol{\pi}_{\boldsymbol{\theta}}}^{\boldsymbol{\pi}_{\boldsymbol{\theta}}}(\nu)\|_{2} - \frac{\xi}{(1 + C_{\psi})(1 - \gamma)} \right)$$

$$+ \frac{2\sqrt{2}}{1-\gamma} \left(L_r + \frac{\gamma}{1-\gamma} L_{\mathbf{P}}(R_{\max} + \lambda \log |\mathcal{A}|) \right) - \frac{\lambda}{1-\gamma} (2\log |\mathcal{A}| - 1) \right)^2$$

$$\leq \frac{2|\mathcal{S}||\mathcal{A}|}{(1+C_{\psi})^2} \left\| \frac{d_{\pi_{\theta},\rho}^{\pi_{\theta}}}{d_{\pi_{\theta},\nu}^{\pi_{\theta}}} \right\|_{\infty}^2 \|\nabla_{\theta} \tilde{V}_{\pi_{\theta}}^{\pi_{\theta}}(\nu)\|_2^2 + 4 \left(\frac{\xi^2}{(1+C_{\psi})^2 (1-\gamma)^2} \right)$$

$$+ \frac{32}{(1-\gamma)^2} \left(L_r + \frac{\gamma}{1-\gamma} L_{\mathbf{P}}(R_{\max} + \lambda \log |\mathcal{A}|) \right)^2 + \frac{4\lambda^2}{(1-\gamma)^2} (2\log |\mathcal{A}| + 1)^2$$

Finally, using the sub-steps followed in Step 3 of 3 we get,

$$\min_{t \leq T} (V_{\boldsymbol{\pi}_{o}^{\star}}^{\star}(\rho) - V_{\boldsymbol{\pi}_{\theta}^{(t)}}^{\boldsymbol{\pi}_{o}^{(t)}}(\rho))^{2} \leq \frac{4\xi^{2}}{(1 + C_{\psi})^{2}(1 - \gamma)^{2}} + \frac{32}{(1 - \gamma)^{2}} \left(L_{r} + \frac{\gamma}{1 - \gamma} L_{\mathbf{P}}(R_{\max} + \lambda \log |\mathcal{A}|)\right)^{2} \\
+ \frac{4\lambda^{2}}{(1 - \gamma)^{2}} (2\log |\mathcal{A}| + 1)^{2} + \frac{2\delta_{0}}{\eta T(1 - \frac{L_{\lambda}\eta}{2})} \frac{|\mathcal{S}||\mathcal{A}|}{\eta T(1 - \frac{L_{\lambda}\eta}{2})} \frac{|\mathcal{S}||\mathcal{A}|}{(1 + C_{\psi})^{2}} \leq \epsilon^{2}$$

Hence, we arrive at the final expression for T,

$$T \geq \frac{2\tilde{\delta}_{0}|\mathcal{S}||\mathcal{A}|\left\|\frac{d_{\pi_{\theta}, \rho}^{\pi_{\theta}^{*}}}{d_{\pi_{\theta}, \nu}^{\pi_{\theta}}}\right\|_{\infty}^{2}}{\eta(1 - \frac{L_{\lambda}\eta}{2})(1 + C_{\psi})^{2} \left[\epsilon^{2} - \frac{4\xi^{2}}{(1 + C_{\psi})^{2}(1 - \gamma)^{2}} - \frac{32}{(1 - \gamma)^{2}}\left(L_{r} + \frac{\gamma}{1 - \gamma}L_{\mathbf{P}}(R_{\max} + \lambda \log|\mathcal{A}|)\right)^{2} - \frac{4\lambda^{2}}{(1 - \gamma)^{2}}(2\log|\mathcal{A}| + 1)^{2}\right]}$$

Hence,

 Choosing $\eta = 1/2L_{\lambda}$, we get,

$$T \geq \frac{(32R_{\max} + 32\lambda \log |\mathcal{A}|)L_{\lambda}|\mathcal{S}||\mathcal{A}|}{3\epsilon^{2}(1 - \gamma)(1 + C_{\psi})^{2} - 12\xi^{2} - 96(1 + C_{\psi})^{2}\left(L_{r} + \frac{\gamma}{1 - \gamma}L_{\mathbf{P}}(R_{\max} + \lambda \log |\mathcal{A}|)\right)^{2} - 12\lambda^{2}(1 + C_{\psi})^{2}(2\log |\mathcal{A}| + 1)^{2}}$$

And consequently,

$$\epsilon^{2} \geq \frac{4\xi^{2}}{(1+C_{\psi})^{2}(1-\gamma)^{2}} + \frac{32}{(1-\gamma)^{2}} \left(L_{r} + \frac{\gamma}{1-\gamma} L_{\mathbf{P}}(R_{\max} + \lambda \log |\mathcal{A}|)\right)^{2} + \frac{4\lambda^{2}}{(1-\gamma)^{2}} (2\log |\mathcal{A}| - 1)^{2}$$

$$\Longrightarrow \epsilon \geq \frac{2\sqrt{2}\xi}{(1+C_{\psi})(1-\gamma)} + \frac{8}{(1-\gamma)} \left(L_{r} + \frac{\gamma}{1-\gamma} L_{\mathbf{P}}(R_{\max} + \lambda \log |\mathcal{A}|)\right) + \frac{2\sqrt{2}\lambda}{1-\gamma} (2\log |\mathcal{A}| + 1)$$

Choose $\lambda = \frac{(1-\gamma)}{(1+2\log|\mathcal{A}|)}$ to get

$$\epsilon \ge \frac{8}{(1-\gamma)} \Big(L_r + \frac{\gamma}{1-\gamma} L_{\mathbf{P}} R_{\max} + \frac{\gamma}{3} \Big) \Big) + 2\sqrt{2}$$

H ABLATION STUDY ON ENTROPY REGULARISATION

Effect of Entropy Regularization 0.2 0.1 Value Function V^{π} 0.0 -0.1-0.2Entropy Reg. -0.3 $\lambda = 0.01$ $\lambda = 0.5$ $\lambda = 1$ -0.4 $\lambda = 2$ Ò **Iteration**

Figure 3: Ablation study for PePG for different values of regularised λ with 20 random seeds, each for 100 iterations

We conducted an ablation study across four entropy regularization strengths ($\lambda \in \{0.01, 0.5, 1, 2\}$ to determine the optimale balance between exploration and convergence stability in RegPePG. The results demonstrate that $\lambda = 2$ achieves the highest final performance (0.05), while smaller values ($\lambda \le 1$) converge to similar suboptimal levels around -0.01 to 0, indicating that stronger entropy regularization enables more effective exploration of the policy space in performative settings.

I TECHNICAL LEMMAS

Lemma 10. For any $\pi, \pi' \in \Pi(\Theta)$, the following non-trivial lower bound holds,

$$\left\| \frac{oldsymbol{d_{\pi'}}}{oldsymbol{d_{\pi}}} \right\|_{20} \geq 1$$

Proof.

$$\left\| \frac{\boldsymbol{d}_{\boldsymbol{\pi}'}}{\boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}} \right\|_{\infty} = \max_{s,a} \frac{\boldsymbol{d}_{\boldsymbol{\pi}'}(s,a)}{\boldsymbol{d}_{\boldsymbol{\pi}}(s,a)} \ge \frac{1}{\sum_{s,a} w_{s,a}} \sum_{s,a} \frac{\boldsymbol{d}_{\boldsymbol{\pi}'}(s,a)}{\boldsymbol{d}_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s,a)} \cdot w_{s,a}$$

Choose $w_{s,a} = d_{\boldsymbol{\pi}}^{\boldsymbol{\pi}}(s,a)$ Hence, we get,

$$\max_{s,a} \frac{d_{\pi'}(s,a)}{d_{\pi}(s,a)} \ge \frac{\sum_{s,a} d_{\pi'}(s,a)}{\sum_{s,a} d_{\pi}^{\pi}(s,a)} = 1$$

The last equality holds from the fact that the state-action occupancy measure is a distribution over $\mathcal{S} \times \mathcal{A}$. Hence, $\sum_{s,a} d_{\pi'}(s,a) = \sum_{s,a} d_{\pi}(s,a) = 1$