

BEYOND SOFTMAX AND ENTROPY: CONVERGENCE RATES OF POLICY GRADIENTS WITH f -SOFTARGMAX PARAMETERIZATION & COUPLED REGULARIZATION

Safwan Labbi¹, Daniil Tiapkin^{1,2}, Paul Mangold¹, Eric Moulines^{3,4}

¹ CMAP, CNRS, École Polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

² Université Paris-Saclay, CNRS, LMO, 91405, Orsay, France

³ Mohamed bin Zayed University of Artificial Intelligence, UAE

⁴ LRE EPITA, 94270 Le Kremlin-Bicêtre, France

{safwan.labbi, daniil.tiapkin, paul.mangold}@polytechnique.edu
eric.moulines@mbzuai.ac.ae

ABSTRACT

Policy gradient methods are known to be highly sensitive to the choice of policy parameterization. In particular, the widely used softmax parameterization can induce ill-conditioned optimization landscapes and lead to exponentially slow convergence. Although this can be mitigated by preconditioning, this solution is often computationally expensive. Instead, we propose replacing the softmax with an alternative family of policy parameterizations based on the generalized f -softargmax. We further advocate coupling this parameterization with a regularizer induced by the same f -divergence, which improves the optimization landscape and ensures that the resulting regularized objective satisfies a Polyak–Łojasiewicz inequality. Leveraging this structure, we establish the *first explicit non-asymptotic last-iterate convergence guarantees* for stochastic policy gradient methods for finite MDPs *without any form of preconditioning*. We also derive sample-complexity bounds for the unregularized problem and show that f -PG with Tsallis divergences achieves *polynomial sample complexity* in contrast to the exponential complexity incurred by the standard softmax parameterization.

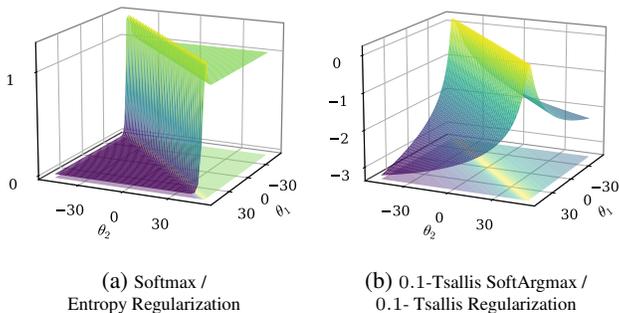


Figure 1: Regularized value landscapes (with temperature $\lambda = 1$) for a one-state, two-action MDP: softmax with entropy (left) versus α -Tsallis SoftArgmax with α -Tsallis regularization (right, $\alpha = 0.1$). The value of the classical coupling Entropy–Softmax is much flatter than for our proposed coupling Tsallis–Tsallis. Using the latter removes flat areas that are far from the solution, allowing *policy gradient methods to escape the gravitational pull*.

1 INTRODUCTION

Policy gradient methods are a cornerstone of modern reinforcement learning (RL) and underpin many of its most notable successes. Algorithms such as Trust-Region Policy Optimization (TRPO; Schulman et al., 2015) and Proximal Policy Optimization (PPO; Schulman et al., 2017) have demonstrated strong empirical performance across a wide range of domains (Berner et al., 2019; Akkaya et al., 2019). Despite these successes, it has become increasingly clear that the performance and convergence behavior of policy gradient methods are highly sensitive to seemingly low-level design choices, among which the choice of policy parameterization plays a central role (Hsu et al., 2020).

Table 1: Comparison of the performance of *scalable*⁽¹⁾ policy gradient methods with global convergence guarantees on the unregularized objective.

	Configuration		Stochastic	Last iterate	Explicit Rates ⁽³⁾
	Parameterization.	Regularization.			
Mei et al. (2020b)	softmax	✗	✗	✓	✗
Mei et al. (2020a)	EscortTransform	✗	✗	✓	✗
Zhang et al. (2021)	softmax	Log-Barrier	✓	✗	✓ ⁽⁴⁾ _{poly}
Liu et al. (2025)	Hadamard	✗	✗	✓	✗
Ours (Corollary 4.6)	softmax	Entropy	✓	✓	✓ ⁽⁵⁾ _{exp}
Ours (Corollary 4.7)	f_α -softargmax ⁽²⁾	α -Tsallis	✓	✓	✓ ⁽⁴⁾ _{poly}

(1) We refer by *scalability* to policy gradient methods that do not use any form of preconditioning; (2) refers to the parameterization induced by using the α -Tsallis divergence generator (see Table 2); (3) Explicit rates means an explicit dependency on all problem parameters and not on intractable quantities; (4) ✓_{poly} indicates an explicit convergence rate with explicit polynomial dependency on all problem parameters; (5) ✓_{exp} indicates an explicit convergence rate with exponential dependence on at least one parameter.

In discrete control scenarios, the default choice is the softmax parameterization, typically coupled with entropy regularization. While ubiquitous, several recent results have revealed fundamental limitations of softmax-based policy gradient methods (Mei et al., 2020a; Li et al., 2023). In particular, in the absence of regularization, the softmax parameterization can induce extremely flat regions in the optimization landscape, leading to an unavoidable exponential lower bound on the rate of convergence (Li et al., 2023). Although entropy regularization is sometimes introduced in an attempt to mitigate this issue, no polynomial convergence guarantees are currently known in this setting. Even with entropy regularization, the landscape remains flat (see Figure 1a for an illustration). These observations motivate treating the policy parameterization itself as a *design choice*: rather than varying the regularizer within softmax, we ask whether moving beyond softmax can fundamentally improve the conditioning of policy gradient methods.

In this paper, we follow the line of work of (Mei et al., 2020a; Liu et al., 2025) and propose a new flexible family of alternative parameterizations induced by divergence generators (denoted f in the following), which we refer to as *f -softargmax parameterizations*. We regularize the objective with the corresponding f -divergence, and we refer to this as a *coupled* parameterization–regularization pair (i.e., the same generator f induces both the parameterization and the regularizer). This viewpoint generalizes the classical softmax–entropy pairing, in which the policy is both induced and regularized by Shannon entropy. Similar constructions have recently shown theoretical and practical benefits for supervised learning (Blondel et al., 2020; Roulet et al., 2025), but remain largely unexplored in reinforcement learning. This takes policy gradient methods *beyond the softmax and its coupled entropy regularization*, yielding a better conditioned optimization landscape (see Figure 1b). In particular, we show that when these parameterizations are *coupled* with the regularizer induced by the corresponding f -divergence, policy gradient methods enjoy improved convergence rates. Remarkably, for the Tsallis divergence, this leads to convergence rates exponentially faster compared to the softmax–entropy pairing.

Formally, we study the f -regularized value function under the f -softargmax parameterization. We show that it satisfies a *non-uniform Łojasiewicz inequality* and a monotonicity property. This monotonicity allows us to restrict the optimization to regions that are easy to project onto and in which the Łojasiewicz coefficient is uniformly lower bounded, resulting in the *uniform Polyak–Łojasiewicz inequality* over the region of interest. Building on these observations, we establish *global last-iterate convergence guarantees* for stochastic policy gradient methods in the tabular setting, with *fully explicit constants*. To the best of our knowledge, these are the first guarantees of this type for policy gradient methods, even with entropy regularization and softmax parameterization, that do not rely on preconditioning or exponentially large batch sizes. For the KL-induced parameterization–regularization pair (softmax–entropy), the resulting uniform Polyak–Łojasiewicz constant is exponentially small in the problem parameters, recovering known exponential convergence rates (Ding et al., 2025; Labbi et al., 2026). In contrast, for Tsallis divergence generators, this constant scales only polynomially, reflecting a substantially better-conditioned optimization landscape. Addi-

tionally, our analysis shows that moving beyond the entropy-softmax pairing yields a better trade-off between regularization bias and sample complexity. In particular, Tsallis-type couplings yield *polynomial* last-iterate convergence guarantees even for the *unregularized* objective, improving upon the worst-case guarantees known for the standard softmax (see Table 1).

Overall, our contributions are threefold:

- We introduce f -softargmax policy parameterizations and study the regularity of the associated f -regularized value function as a function of the policy parameters. We show that it is smooth, satisfies a non-uniform Łojasiewicz inequality, and, by exploiting a monotonicity property, admits a uniform bound on a Polyak-Łojasiewicz constant on a region of interest that is easy to project onto.
- We prove global last-iterate convergence guarantees for stochastic tabular policy gradient, with fully explicit sample complexity bounds that apply to both regularized and unregularized objectives.
- We demonstrate that alternative couplings beyond entropy-softmax lead to improved sample complexity for an unregularized problem both theoretically and empirically. In particular, the Tsallis coupling yields polynomial dependencies on problem parameters, resulting in an exponential improvement over softmax, and provides additional flexibility for practical adaptation.

The paper is organized as follows. Section 2 introduces the necessary background. Section 3 presents the f -softargmax parameterization and the properties of the f -regularized value under this parameterization. Convergence rates for policy gradient are established in Section 4, and numerical experiments are reported in Section 5.

Related work. (*Policy gradient methods.*) Global convergence guarantees are known to hold for unregularized policy gradient methods with deterministic gradients, achieving sublinear rates with constant step-sizes (Mei et al., 2020b; Liu et al., 2024). However, the convergence rates of softmax-based policy gradient methods is exponential in the problem parameters (Mei et al., 2020a; Li et al., 2023). Two strategies were proposed to mitigate this issue: first, preconditioning, most notably through natural policy gradient methods (Kakade, 2001), which can alleviate ill-conditioning, but scales poorly to larger problems due to the nature of the updates. Second, log-barrier regularization (Zhang et al., 2021), which yields polynomial rates but has no last iterate convergence guarantees and is unstable in practice. In contrast, our approach avoids preconditioning altogether and therefore retains the scalability of standard policy gradient methods, while providing explicit polynomial convergence guarantees for the *last iterate*.

(*Alternative Parameterizations.*) Alternatives to softmax have been proposed and studied in optimization (Martins & Astudillo, 2016; Peters et al., 2019; Roulet et al., 2025). In RL, the study of alternative parameterizations is still in its early stages. The escort transform of Mei et al. (2020a) avoids exponential slowdowns in deterministic settings, but its guarantees rely on increasing step-sizes and do not extend to stochastic gradients. The Hadamard parameterization (Liu et al., 2025) yields local linear convergence in deterministic regimes, but without explicit constants. In this work, we propose a more flexible family of parameterizations that can adapt to various problems and provide explicit guarantees in the stochastic setting.

2 BACKGROUND

Reinforcement learning. Consider a discounted Markov decision process $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \gamma, \mathsf{P}, r, \rho)$ with finite state and action spaces \mathcal{S} and \mathcal{A} , discount factor $\gamma \in (0, 1)$, transition kernel $\mathsf{P}: \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{S})$, bounded reward function $r: \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, and initial distribution ρ . The value of a policy $\pi: \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$ is defined by

$$v_\pi(s) := \mathbb{E}_s^\pi \left[\sum_{t=0}^{\infty} \gamma^t r(S_t, A_t) \right], \quad (1)$$

where $S_0 = s$, $A_t \sim \pi(\cdot | S_t)$, and $S_{t+1} \sim \mathsf{P}(\cdot | S_t, A_t)$. For $\rho \in \mathcal{P}(\mathcal{S})$, we define $v_\pi(\rho) := \sum_s \rho(s) v_\pi(s)$. For any π , we define a corresponding discounted occupancy measure

$$d_\rho^\pi(s) := \frac{1}{1-\gamma} \mathbb{E}_s^\pi \left[\sum_{t=0}^{\infty} \gamma^t \mathbf{1}_s(S_t) \right].$$

Parameterizations on the simplex. Following Roulet et al. (2025), we study a family of parameterizations of the simplex based on divergence generators. For a given *generator* $f: (0, \infty) \rightarrow \mathbb{R}$

Table 2: Example of classical divergence generators f included in our framework and their associated f-softargmax operators.

Name	$f(\mathbf{u})$	f-softargmax(x, ν_{ref})[a] ⁽¹⁾
KL	$u \log u - (u - 1)$	$\nu_{\text{ref}}(a) \exp(x(a)) / (\sum_{b \in \mathcal{A}} \nu_{\text{ref}}(b) \exp(x(b)))$
Tsallis ($0 < \alpha < 1$)	$(u^\alpha - \alpha u + \alpha - 1) / \alpha(\alpha - 1)$	$\nu_{\text{ref}}(a) (1 + (\alpha - 1)(x(a) - \mu_x^\alpha))^{1/(\alpha - 1)}$
Jensen-Shannon	$u \log(u) - (u + 1) \log(\frac{u+1}{2})$	$\nu_{\text{ref}}(a) \exp(2(x(a) - \mu_x)) / (2 - \exp(2(x(a) - \mu_x)))$

(1) Here μ_x^α and μ_x are normalization factors that ensures that the weights sum to 1.

strictly convex with $f(1) = 0$, and reference distribution q with *full support*, we define

$$\begin{aligned} \text{f-softmax}(x, q) &:= \max_{\nu \in \mathcal{P}(\mathcal{A})} \{ \langle \nu, x \rangle - D^f(\nu \| q) \} , \\ \text{f-softargmax}(x, q) &:= \arg \max_{\nu \in \mathcal{P}(\mathcal{A})} \{ \langle \nu, x \rangle - D^f(\nu \| q) \} , \end{aligned}$$

where $D^f(p \| q)$ is the f -divergence between p and q (see [Csiszár, 1967](#), or [Appendix A](#)). Since $D^f(p \| q)$ is strictly convex in its first argument on the simplex, the output of the f-softargmax operator is well defined and unique, as it corresponds to the $\arg \max$ of a strictly concave function over a compact set. This construction recovers the classical softmax as a special case and yields a rich family of alternative parameterizations (see [Table 2](#)). Computing f-softargmax reduces to solving a one-dimensional root-finding problem, which can be done efficiently by dichotomy; see [Roulet et al. \(2025\)](#) and [Lemma B.1](#) for details.

f -regularized value functions. Given a reference policy π_{ref} , temperature $\lambda \geq 0$, and a divergence generator f , the f -regularized value function of π is defined by

$$v_\pi^f(s) := \mathbb{E}_s^\pi \left[\sum_{t=0}^{\infty} \gamma^t (r(S_t, A_t) - \lambda D^f(\pi(\cdot | S_t) \| \pi_{\text{ref}}(\cdot | S_t))) \right] ,$$

A key result ([Geist et al., 2019](#)) is that the optimal regularized value $v_\star^f(s) := \max_\pi v_\pi^f(s)$, together with optimal policy π_\star^f admits a closed-form Bellman characterization:

$$v_\star^f(s) = \max_{\nu \in \mathcal{P}(\mathcal{A})} \{ \langle \nu, q_\star^f(s, \cdot) \rangle - \lambda D^f(\nu \| \pi_{\text{ref}}(\cdot | s)) \} , \quad (2)$$

$$\pi_\star^f(\cdot | s) = \arg \max_{\nu \in \mathcal{P}(\mathcal{A})} \{ \langle \nu, q_\star^f(s, \cdot) \rangle - \lambda D^f(\nu \| \pi_{\text{ref}}(\cdot | s)) \} , \quad (3)$$

where $q_\star^f(s, a) := r(s, a) + \gamma P v_\star^f(s, a)$.

3 COUPLING PARAMETERIZATION & REGULARIZATION

We introduce a new class of policy parameterizations for reinforcement learning, which we refer to as f-softargmax policies. Let π_{ref} denote a full-support reference policy. For $\theta \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$, we define the f-softargmax policy by

$$\pi_\theta^f(\cdot | s) := \text{f-softargmax}(\theta(s, \cdot), \pi_{\text{ref}}(\cdot | s)), \quad \forall s \in \mathcal{S}. \quad (4)$$

This parameterization can be directly used within unregularized policy gradient methods (see [Appendix G](#)). However, in practice, unregularized methods tend to over-exploit and converge prematurely to suboptimal policies. This suggests that the choice of parameterization should be guided by the geometry of a suitably regularized objective, rather than considered in isolation.

To understand which regularization is naturally associated with the f-softargmax family, we examine the structure of the f -regularized problem (3). The optimal policy of this problem admits the following representation:

$$\pi_\star^f(\cdot | s) = \text{f-softargmax}(q_\star^f(s, \cdot) / \lambda, \pi_{\text{ref}}(\cdot | s)) .$$

In particular, if we choose logits $\theta_\star^f(s, \cdot) = q_\star^f(s, \cdot) / \lambda + b(s)$, where $b: \mathcal{S} \rightarrow \mathbb{R}$ is an arbitrary state-dependent baseline, then the f-softargmax mapping exactly recovers the optimal f -regularized policy, i.e., $\pi_{\theta_\star}^f = \pi_\star^f$.

This shows that the f -softargmax parameterization is not arbitrary: it is precisely matched to the geometry of the f -regularized problem. Under this parameterization, learning the policy is equivalent to learning the regularized optimal Q -function, and the associated f -divergence regularizer arises naturally from the variational characterization of the optimal policy.

To further formalize the benefits of such coupling, we now establish the smoothness, as well as a Polyak-Łojasiewicz inequality, of the f -regularized value with coupled parameterization $v_\theta^f := v_{\pi_\theta}^f$.

We derive these properties under the following two assumptions on f and π_{ref} .

Assumption $\mathbf{P}(\underline{\pi}_{\text{ref}})$. *There exists a number $\underline{\pi}_{\text{ref}} > 0$ such that $\min_{(s,a)} \pi_{\text{ref}}(a|s) > \underline{\pi}_{\text{ref}}$.*

Assumption $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$. *The generator function f satisfies:*

- (i) f is bounded, strictly convex on $[0; 1/\underline{\pi}_{\text{ref}}]$, $f(1) = 0$, and f is thrice differentiable on $(0, 1/\underline{\pi}_{\text{ref}})$;
- (ii) $\lim_{u \downarrow 0} f'(u) = -\infty$, and $\lim_{u \rightarrow 0} |f'(u)/f''(u)| < \infty$;
- (iii) there exists $\omega_f \in [1, \infty)$, and $\kappa_f \in (0, \infty)$, such that for any $u \in [0; 1/\underline{\pi}_{\text{ref}}]$, we have $1/(uf''(u)) \leq \omega_f$, and $|f'''(u)/f''(u)^2| \leq \kappa_f$;
- (iv) there exists $\iota_f \in (0, 1]$ such that f'' decreases on $[0; \iota_f]$ and for $u \in [\iota_f; 1/\underline{\pi}_{\text{ref}}]$, $f''(\iota_f) \geq f''(u)$.

These conditions are met by a broad class of commonly used divergence generators, like the KL, Tsallis with $\alpha \leq 1$, and Jensen-Shannon (see Appendix F).

Remark 3.1 *Tsallis divergences with $\alpha > 1$ violate condition (ii): since $f'(0)$ is finite, the induced policies are sparse, leading to non-smooth parameterizations.*

Under $\mathbf{P}(\underline{\pi}_{\text{ref}})$ and $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$, we can define the weights $w_\theta^f(a|s)$ and the sum $W_\theta^f(s)$, defined as

$$w_\theta^f(a|s) := \frac{1}{W_\theta^f(s)} \frac{\pi_{\text{ref}}(a|s)}{f''(\pi_\theta^f(a|s)/\pi_{\text{ref}}(a|s))}, \quad W_\theta^f(s) := \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''(\pi_\theta^f(a|s)/\pi_{\text{ref}}(a|s))}, \quad (5)$$

which will play a central role in our analysis. In the KL divergence case, we recover simple expressions $w_\theta^f(a|s) \equiv \pi_\theta^f(a|s)$ and $W_\theta^f(s) = 1$. Using the notations in (5), we can express the gradient of the regularized value.

Lemma 3.2. *Assume that, for some $\underline{\pi}_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$ and $\mathbf{P}(\underline{\pi}_{\text{ref}})$. For any $s \in \mathcal{S}$, we have*

$$\frac{\partial v_\theta^f(\rho)}{\partial \theta(s, \cdot)} = (1 - \gamma)^{-1} W_\theta^f(s) d_\rho^\theta(s) H(w_\theta^f(\cdot|s)) \left[q_\theta^f(s, \cdot) - \lambda \theta(s, \cdot) \right],$$

where for any vector $u \in \mathbb{R}^{|\mathcal{A}|}$, $H(u) := \text{diag}(u) - uu^\top$, $q^f := r + \gamma \mathbf{P} v_\theta^f$, and $d_\rho^\theta(s) := d_\rho^{\pi_\theta^f}(s)$.

Next, we introduce three quantities that arise naturally in the expression of the Hessian of $v_\theta^f(\rho)$.

$$y_f := \max_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s) \frac{|f'(\nu(a)/\pi_{\text{ref}}(a|s))|}{f''(\nu(a)/\pi_{\text{ref}}(a|s))}, \quad d_f := \max_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} D^f(\nu \| \pi_{\text{ref}}(\cdot|s)), \quad (6)$$

$$\zeta_f := \min_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''(\nu(a)/\pi_{\text{ref}}(a|s))} > 0. \quad (7)$$

Since they are bounded under $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$ and $\mathbf{P}(\underline{\pi}_{\text{ref}})$, we can establish the smoothness of $v_\theta^f(\rho)$.

Theorem 3.3. *Assume $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$ and $\mathbf{P}(\underline{\pi}_{\text{ref}})$ for some $\underline{\pi}_{\text{ref}} > 0$. For any $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, $\|\nabla^2 v_\theta^f(\rho)\|_2 \leq L_f$ with*

$$L_f := \mathcal{O} \left(\frac{\omega_f^2 + \omega_f \kappa_f + \lambda \cdot (\omega_f^2 d_f + \omega_f (\kappa_f d_f + y_f) + \omega_f + 2\kappa_f y_f)}{(1 - \gamma)^3} \right).$$

We refer to Appendix B.4 for a proof and a complete expression of L_f . We now introduce the classical exploration assumption (Mei et al., 2020a;b; Agarwal et al., 2021).

Assumption \mathbf{A}_ρ . *The coefficient $\rho_{\min} := \min_{s \in \mathcal{S}} \rho(s)$ of the initial distribution ρ satisfies $\rho_{\min} > 0$.*

Next, we derive a Non-Uniform Łojasiewicz inequality.

Theorem 3.4. Assume that, for some $\pi_{\text{ref}} > 0$, $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ hold. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_ρ . Then, it holds that

$$\|\nabla_\theta v_\theta^f(\rho)\|_2^2 \geq \mu_f(\theta)(v_\star^f(\rho) - v_\theta^f(\rho)) ,$$

with $\mu_f(\theta) := \lambda(1 - \gamma)\rho_{\min}^2(\zeta_f/\omega_f)^2 \min_{s,a} w_\theta^f(a|s)^2$.

We prove this theorem in Appendix C. To highlight the main steps of the proof, we give a sketch of the proof in the bandits setting, where the state space is a singleton.

Sketch of the proof in the bandits case. The proof consists of two steps: (1) we bound the suboptimality gap by the distance between the logit and the rescaled reward; (2) then link it to the gradient of the function.

Step 1: By (2), the optimal regularized value is equal to $v_\star^f = \lambda \text{f-softmax}(r/\lambda, \pi_{\text{ref}})$. Next, since $\text{f-softmax}(\theta, \pi_{\text{ref}}) = \langle \pi_\theta^f, \theta \rangle - \text{D}^f(\pi_\theta^f \| \pi_{\text{ref}})$, and $v_\theta^f = \langle \pi_\theta^f, r \rangle - \lambda \text{D}^f(\pi_\theta^f \| \pi_{\text{ref}})$, the suboptimality gap rewrite as

$$\begin{aligned} v_\star^f - v_\theta^f &= \lambda \cdot \text{f-softmax}\left(\frac{r}{\lambda}, \pi_{\text{ref}}\right) - \langle \pi_\theta^f, r \rangle + \lambda \text{D}^f(\pi_\theta^f \| \pi_{\text{ref}}) \\ &= \lambda \left[\text{f-softmax}\left(\frac{r}{\lambda}, \pi_{\text{ref}}\right) - \text{f-softmax}(\theta, \pi_{\text{ref}}) - \langle \pi_\theta^f, \frac{r}{\lambda} - \theta \rangle \right]. \end{aligned}$$

Combining $\pi_\theta^f = \nabla \text{f-softmax}(\theta, \pi_{\text{ref}})$, and that for $a \in \mathbb{R}$, $\text{f-softmax}(\theta + \alpha \mathbf{1}_{|\mathcal{A}|}, \pi_{\text{ref}}) = \text{f-softmax}(\theta, \pi_{\text{ref}}) + \alpha$ yields

$$\begin{aligned} v_\star^f - v_\theta^f &= \lambda \left[\text{f-softmax}\left(\frac{r}{\lambda}, \pi_{\text{ref}}\right) - \text{f-softmax}(\theta + K_\theta^f \mathbf{1}_{|\mathcal{A}|}, \pi_{\text{ref}}) \right. \\ &\quad \left. - \langle \nabla \text{f-softmax}(\theta + K_\theta^f \mathbf{1}_{|\mathcal{A}|}, \pi_{\text{ref}}), r/\lambda - \theta - K_\theta^f \mathbf{1}_{|\mathcal{A}|} \rangle \right], \end{aligned}$$

where we have defined $K_\theta^f = \langle r/\lambda - \theta, \mathbf{1}_{|\mathcal{A}|} \rangle / |\mathcal{A}|$. Defining $\zeta_\theta^f = r/\lambda - \theta - K_\theta^f \mathbf{1}_{|\mathcal{A}|}$ and using a second-order Taylor expansion of the function $\text{f-softmax}(\cdot)$ between, we obtain

$$v_\star^f - v_\theta^f = \frac{\lambda}{2} (\zeta_\theta^f)^\top \nabla^2 \text{f-softmax}_f(\xi) \zeta_\theta^f, \quad (8)$$

for some ξ on the segment joining $\theta + K_\theta^f \mathbf{1}_{|\mathcal{A}|}$ and r/λ . Next, by Lemma B.5, it holds that $\|\nabla^2 \text{f-softmax}(\xi)\|_2 \leq 2\omega_f$, which implies $v_\star^f - v_\theta^f \leq \lambda\omega_f \|\zeta_\theta^f\|_2^2$.

Step 2: Using Lemma 3.2, we have $\nabla v_\theta^f = \text{W}_\theta^f H(\text{w}_\theta^f) [r - \lambda\theta]$. Next, applying Lemma 23 of Mei et al. (2020b) (see Lemma I.4) gives

$$\|\nabla v_\theta^f\|_2^2 \geq \text{W}_\theta^f \min_{a \in \mathcal{A}} w_\theta^f \|\zeta_\theta^f\|_2^2 \geq \zeta_f \min_{a \in \mathcal{A}} w_\theta^f \|\zeta_\theta^f\|_2^2 .$$

where by Lemma B.6, we have $\text{W}_\theta^f \geq \zeta_f$. Combining the two bounds concludes the proof \square

For softmax parameterization coupled with entropy regularization, we retrieve a property outlined by Mei et al. (2020b). However, their proof is highly specific to the entropy-softmax pairing case and cannot be extended to general f -divergences, because it relies in an essential way on the logarithm's special properties. Indeed, their proof require rewriting the soft sub-optimality gap $v_\star^{\text{KL}}(\rho) - v_\theta^{\text{KL}}(\rho)$ as $\frac{\lambda}{1-\gamma} \sum_{s' \in \mathcal{S}} d_\rho^\theta(s') \text{D}^{\text{KL}}(\pi_\theta^{\text{KL}}(\cdot|s') \| \pi_\star^{\text{KL}}(\cdot|s))$ which is not possible for a general f . Our proof is more natural, as it simply relies on Taylor expansion to obtain the inequality (8).

From Non-Uniform Łojasiewicz to Polyak-Łojasiewicz. To obtain a Polyak-Łojasiewicz bound, we need to bound the coefficient $\mu_f(\theta)$ uniformly. To this end, we restrict the optimization to a smaller subspace, eliminating policies for which the regularization is too large. Given a policy π and $0 < \tau < \underline{\pi}_{\text{ref}}/2$, we define the projection-like operator \mathcal{U}_τ , for every $(s, a) \in \mathcal{S} \times \mathcal{A}$, as

$$\mathcal{U}_\tau(\pi)(a|s) = \begin{cases} \pi_{\text{ref}}(a|s)\tau, & \text{if } \pi(a|s) \leq \pi_{\text{ref}}(a|s)\tau/2, \\ \pi(a|s) - b_\pi(s), & \text{if } a = a_\pi^{\text{max}}(s), \\ \pi(a|s), & \text{otherwise,} \end{cases}$$

where we define $b_\pi(s) = \sum_{b \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(b|s)\tau - \pi(b|s)$, $a_\pi^{\text{max}}(s) = \arg \max_{a \in \mathcal{A}} \{\pi(a|s)/\pi_{\text{ref}}(a|s)\}$, where ties in the $\arg \max$ are resolved arbitrary, and

$$\mathcal{A}_\tau^\pi(s) := \{a \in \mathcal{A}, \pi(a|s)/\pi_{\text{ref}}(a|s) \leq \tau/2\} .$$

Algorithm 1 f -SoftArgmax Policy Gradient

- 1: **Initialization:** Learning rate $\eta > 0$, initial parameter θ_0 , divergence generator f , batch size B .
- 2: **for** $t = 0$ to $T - 1$ **do**
- 3: Collect $Z_t := (S_{t,0:H-1}^b, A_{t,0:H-1}^b)_{b=0}^{B-1}$ using $\pi_{\theta_t}^f$
- 4: Compute the gradient $g_{Z_t}^f(\theta_t)$ using (11)
- 5: Update $\theta_{t+1} = \mathcal{T}_{\tau_\lambda}(\theta_t + \eta g_{Z_t}^f(\theta_t))$
- 6: **Return** θ_T

This operator prevents policies from becoming "too deterministic": if the probability of any action gets too close to zero, it is increased above a threshold that depends on τ and π_{ref} . For a proper choice of τ , applying this operator on a policy returns a policy with a higher regularized value.

Theorem 3.5. Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$, and that ρ satisfies \mathbf{A}_ρ . Let

$$\tau_\lambda = \min([f']^{-1}(-\frac{16+8\gamma\lambda d_f}{\lambda(1-\gamma)^2\rho_{\min}}), [f']^{-1}(-4|f'(\frac{1}{2})|), \frac{\pi_{\text{ref}}}{2}) .$$

Then, for any policy π and for $\tilde{\pi} = \mathcal{U}_{\tau_\lambda}(\pi)$, it holds that $v_{\tilde{\pi}}^f(\rho) \geq v_\pi^f(\rho)$ and that $\tilde{\pi}(a|s) \geq \pi_{\text{ref}}\tau_\lambda$.

Since \mathcal{U}_τ operates in the space of policies, we lift it to the parameter space by defining an operator \mathcal{T}_τ such that for $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, we have $\pi_{\mathcal{T}_\tau\theta}^f = \mathcal{U}_\tau\pi_\theta^f$ (see Appendix D for an explicit construction). Finally, we show that with the choice of threshold from Theorem 3.5, we can give a uniform lower bound of μ_f on the set of restricted logits. This shows how the non-uniform Łojasiewicz inequality is upgraded to a Polyak–Łojasiewicz condition: it suffices to restrict the search to parameters that encode non-degenerate policies, since such policies are provably suboptimal.

Corollary 3.6. Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ and that ρ satisfies \mathbf{A}_ρ . If $\lambda \lesssim \frac{1}{(1-\gamma)^2\rho_{\min}} \min(\frac{1}{|f'(\frac{1}{2})|}, \frac{1}{|f'(\frac{1}{2})|}, \frac{1}{|f'(\frac{1}{2}\pi_{\text{ref}})|})$ where τ_λ is defined in Theorem 3.5.

Under this condition, it holds that $\inf_{\theta \in \mathbb{R}^d} \mu_f(\mathcal{T}_{\tau_\lambda}\theta) \geq \underline{\mu}_f$, where

$$\underline{\mu}_f := \lambda(1-\gamma)\rho_{\min}^2 \zeta_f^2 \pi_{\text{ref}}^2 (f^*)'' \left(\frac{-16-8\gamma\lambda d_f}{\lambda(1-\gamma)^2\rho_{\min}} \right)^2 / \omega_f^2 ,$$

where f^* is a convex conjugate of f .

4 CONVERGENCE ANALYSIS OF f -PG

In this section, we aim to optimize the f -regularized value under f -softargmax parameterization.

$$\max_{\theta \in \Theta} \{ J^f(\theta) := v_{\pi_\theta^f}^f(\rho) \} . \quad (9)$$

The f -PG algorithm. We introduce f -PG (Algorithm 1), an f -SoftArgmax policy gradient method with coupled regularization. At each iteration of f -PG, the agent samples a batch of independent truncated trajectories of length H from $\nu(\pi; \cdot)$ defined for a single truncated trajectory $z = (s_h, a_h)_{h=0}^{H-1} \in (\mathcal{S} \times \mathcal{A})^H$ by $\nu(\pi; z) := \rho(s_0)\pi(a_0|s_0)\prod_{h=0}^{H-1} \mathbf{P}(s_h|s_{h-1}, a_{h-1})\pi(a_h|s_h)$. Then, the agent performs the update

$$\theta_{t+1} = \mathcal{T}_{\tau_\lambda}(\theta_t + \eta \cdot g_{Z_t}^f(\theta_t)) , \quad \text{for } t \geq 0 , \quad (10)$$

where $\eta > 0$ is a learning rate, $\mathcal{T}_{\tau_\lambda} : \mathbb{R}^{|\mathcal{S}||\mathcal{A}|} \rightarrow \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$ is the projection-like operator defined in Section 3, and $g_{Z_t}^f(\theta_t)$ is a REINFORCE-like estimator (Williams, 1992) of $\nabla v_{\pi_{\theta_t}^f}^f(\rho)$ that uses a batch of B independent trajectories $Z_t \sim [\nu(\theta_t)]^{\otimes B}$. For a batch of trajectories $z = (s_{0:H-1}^b, a_{0:H-1}^b)_{b=0}^{B-1}$, this estimator is defined

$$g_z^f(\theta) := \frac{1}{B} \sum_{b=0}^{B-1} \sum_{h=0}^{H-1} \left\{ \sum_{\ell=0}^h \frac{\partial \log \pi_\theta^f(a_\ell^b|s_\ell^b)}{\partial \theta} \gamma^h r_\theta^f(s_h^b, a_h^b) - \lambda \gamma^h F_\theta^f(s_h^b) \right\} , \quad (11)$$

where $r_\theta^f(s_h^b, a_h^b) = r(s_h^b, a_h^b) - \lambda D^f(\pi_\theta^f(\cdot|s_h^b) \|\pi_{\text{ref}}(\cdot|s_h^b))$, and $F_\theta^f(s) \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$ is defined by

$$[F_\theta^f(s)]_{(s',b)} := \mathbf{1}_{s'}(s) W_\theta^f(s) w_\theta^f(b|s) \left(f' \left(\frac{\pi_\theta^f(b|s)}{\pi_{\text{ref}}(b|s)} \right) - \sum_{a \in \mathcal{A}} w_\theta^f(a|s) f' \left(\frac{\pi_\theta^f(a|s)}{\pi_{\text{ref}}(a|s)} \right) \right) , \quad \forall (s', b) \in \mathcal{S} \times \mathcal{A} .$$

Remark 4.1 (Connection with (Lazy) Mirror Descent.) *We stress that f -PG is fundamentally different from mirror descent. Let $\Phi(\pi) = \sum_{s \in \mathcal{S}} D^f(\pi(\cdot|s) \|\pi_{\text{ref}}(\cdot|s))$, mirror descent iterations are $\nabla \Phi(\tilde{\pi}_{t+1}) = \nabla \Phi(\tilde{\pi}_t) + \eta \nabla_{\pi} v_{\pi}^f(\rho)|_{\pi=\tilde{\pi}_t}$, $\pi_{t+1} = \arg \min_{\pi} \{\Phi(\pi) - \langle \nabla \Phi(\tilde{\pi}_{t+1}), \pi - \tilde{\pi}_{t+1} \rangle\}$.* (12)

where $\pi \in \mathcal{P}(\mathcal{A})^{|\mathcal{S}|}$ is a policy. Denoting $\theta_t = \nabla \Phi(\tilde{\pi}_t)$, one obtains updates that resemble (10) (without $\mathcal{T}_{\tau_\lambda}$), with one key difference: the gradient in (12) is taken with respect to the policy π whereas in (10) it is computed w.r.t the "dual" parameter θ (in the mirror descent terminology). Moreover, the update (12) can be expressed as, by the chain rule, $\theta_{t+1} = \theta_t + \eta [\frac{\partial \pi_{\theta}}{\partial \theta}]_{\theta=\theta_t}^{-1} \nabla_{\theta} J^f(\theta_t)$, which have an additional preconditioning term given by the inverse of the policy Jacobian. See Appendix H for more details.

Next, we bound the bias and variance of the gradient estimator (a proof is provided in Appendix E.1).

Lemma 4.2. *Assume that, for some $\underline{\pi}_{\text{ref}} > 0$, f satisfy $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$. For any $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, we have*

$$\|g^f(\theta) - \frac{\partial v_{\theta}^f(\rho)}{\partial \theta}\|_2 \leq \beta_f, \quad \mathbb{E}_Z[\|g^f(\theta) - g_Z^f(\theta)\|_2^2] \leq \frac{\sigma_f^2}{B},$$

where $Z \sim [\nu(\theta)]^{\otimes B}$, and where

$$\sigma_f^2 := \frac{12}{(1-\gamma)^4} \left[\omega_f^3 + \lambda^2 \gamma^2 \omega_f^3 d_f^2 + \lambda^2 (1-\gamma)^2 \omega_f^2 y_f^2 \right], \quad \beta_f := \frac{2\gamma^H (H+1)}{(1-\gamma)^2} \omega_f [2 + 2\lambda d_f + \lambda(1-\gamma)y_f].$$

Convergence analysis. We now present our main result for this section, which gives a convergence rate for f -PG with explicit constants for the regularized problem. This result is based on the regularity properties of the regularized value function, which we developed in Section 3.

Theorem 4.3. *Assume that, for some $\underline{\pi}_{\text{ref}} > 0$, f and $\underline{\pi}_{\text{ref}}$ satisfy $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$ and $\mathbf{P}(\underline{\pi}_{\text{ref}})$, and that ρ satisfies \mathbf{A}_ρ . Fix $\eta \leq 1/(2L_f)$, and λ and τ_λ as in Corollary 3.6. Then, for any $t \geq 0$, the iterates of f -PG satisfy*

$$\mathbb{E}[\Delta_t] \leq \left(1 - \frac{\eta \underline{\mu}_f}{4}\right)^t \Delta_0 + \frac{6\eta \sigma_f^2}{B \underline{\mu}_f} + \frac{6\beta_f^2}{\underline{\mu}_f},$$

with $\Delta_t = v_{\star}^f(\rho) - v_{\theta_t}^f(\rho)$, and σ_f^2 and β_f^2 from Corollary 3.6 and Lemma 4.2 respectively.

We provide a proof of this result in Appendix E.2. A crucial feature of this theorem is that it is explicit, as all the terms that appear can be expressed using problem-dependent constants. This allows us to derive the following sample complexity result for optimizing the regularized value.

Corollary 4.4. *Let $\epsilon > 0$ and set $H \lesssim (1-\gamma)^2 \cdot \log(1/\underline{\mu}_f)$, and $\eta \lesssim \min(L_f^{-1}, \epsilon B \underline{\mu}_f \cdot \sigma_f^{-2})$. Under assumptions of Theorem 4.3, the final iterate of f -PG satisfies $\mathbb{E}[v_{\star}^f(\rho) - v_{\theta_T}^f(\rho)] \leq \epsilon$ with*

$$T \lesssim \max\left(\frac{L_f}{\underline{\mu}_f}, \frac{\sigma_f^2}{\epsilon B \underline{\mu}_f^2}\right) \log\left(\frac{v_{\star}^f(\rho) - v_{\theta_0}^f(\rho)}{\epsilon}\right).$$

Importantly, Corollary 4.4 shows that f -PG achieves convergence rates comparable to stochastic gradient ascent in the strongly convex regime: $O(\kappa \log(1/\epsilon))$ in the low-variance setting, with condition number $\kappa > 0$, and $O((1/\epsilon) \log(1/\epsilon))$ in general.

Convergence for unregularized objective. A natural question is then how to compare different choices of regularizers, since each method optimizes a distinct regularized objective. By appropriately tuning the temperature λ , we recover the final sample complexity bound for the unregularized problem. We give a precise statement and proof of this result in Corollary E.10.

Corollary 4.5. *Let $0 < \epsilon < (1-\gamma)^{-3} \rho_{\min}^{-1} \min(|f'(\iota_f)|^{-1}, |f'(\frac{1}{2})|^{-1}, |f'(\frac{\underline{\pi}_{\text{ref}}}{2})|^{-1})$. Define $c_f = \min(1/d_f, 1/y_f, 1)$, and set $\lambda = (1-\gamma)\epsilon/4 \cdot c_f$. Under assumptions of Theorem 4.3, the final iterate of f -PG satisfies $\mathbb{E}[v_{\star}(\rho) - v_{\theta_T}(\rho)] \leq \epsilon$ with*

$$T \lesssim \frac{(f^*)''\left(\frac{-1}{\epsilon c_f (1-\gamma)^3 \rho_{\min}}\right)^{-4}}{\epsilon^3 B (1-\gamma)^8 \rho_{\min}^4 \underline{\pi}_{\text{ref}}^4} \log\left(\frac{6(v_{\star}^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon}\right),$$

$H \lesssim (1-\gamma)^{-2} \cdot \log(1/\underline{\mu}_f)$, and $\eta \lesssim a \wedge b$ with $a = (1-\gamma)^3$ and $b = \epsilon^2 (f^*)''\left(\frac{-1}{\epsilon c_f (1-\gamma)^6 \rho_{\min}}\right)^2 (1-\gamma)^3 B \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2$, where f^* is a convex conjugate of f .

This corollary shows that the convergence rate to the unregularized optimum is primarily controlled by the asymptotic behaviour of $(f^*)''$. As the target precision $\epsilon \rightarrow 0$, divergences for which $(f^*)''$ grows faster yield better conditioning, which in turn results in faster convergence.

Sample complexity for specific choices of f . We now provide a more complete interpretation of these results by stating sample complexity bounds for specific choices of f .

Corollary 4.6 (Complexity for Softmax-Entropy). *Let f be the Kullback-Leibler divergence generator. Let $\epsilon > 0$. Under the choice of $\eta, \lambda, H, \tau_\lambda$, and T of Corollary 4.5, the final iterate of f -PG achieves $\mathbb{E}[v_*(\rho) - v_{\theta_T}(\rho)] \leq \epsilon$ in $TBH \lesssim \frac{|\log(\pi_{\text{ref}})|^3}{\epsilon^4(1-\gamma)^{12}\rho_{\min}^5\pi_{\text{ref}}^4} \exp(\frac{|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^3\rho_{\min}})$ samples.*

This corollary shows that the number of samples required by the softmax policy gradient method is exponential in $1/(1-\gamma)$. This is in line with recent work on vanilla softmax policy gradient, which demonstrated that the number of steps is at least exponential in $1/(1-\gamma)$ (Li et al., 2023).

Corollary 4.7 (Complexity for α -Tsallis SoftArgmax with α -Tsallis regularization). *Let f be the α -Csiszár-Cressie-Read divergence generator for $\alpha \in (0, 1)$ (see Table 2 for its expression). Let $\epsilon > 0$. Under the choice of $\eta, \lambda, H, \tau_\lambda$, and T of Corollary 4.5, the last iterate of f -PG achieves $\mathbb{E}[v_*(\rho) - v_{\theta_T}(\rho)] \leq \epsilon$ in a number of samples*

$$TBH \lesssim \frac{|\log(\pi_{\text{ref}})|^3}{\epsilon^4\alpha^6(1-\gamma)^{12}\rho_{\min}^5\pi_{\text{ref}}^{4+7(1-\alpha)}} \left(1 + \frac{(1-\alpha)|\log(\pi_{\text{ref}})|}{\epsilon\alpha^2(1-\gamma)^3\rho_{\min}}\right)^{4+\frac{4}{1-\alpha}}.$$

We give detailed versions and prove these corollaries in Appendix F. These corollaries show that Tsallis SoftArgmax parameterization with coupled regularisation allows for faster learning, reducing the dependency on $(1-\gamma)^{-1}$ from exponential in Corollary 4.6 to polynomial in Corollary 4.7. Next, we approximate the choice of α that achieves the fastest convergence (according to our bounds).

Corollary 4.8. *Assume the conditions of Corollary 4.7 hold. The value of α that minimizes the sample complexity in Corollary 4.7 is given by $\alpha^*(\epsilon) = 11/(2\log(1/\epsilon)) + o(1/\log(1/\epsilon))$. Moreover, for ϵ sufficiently small, choosing $\alpha = \alpha^*(\epsilon)$ yields a sample complexity $\epsilon^{-1.2}\text{poly}(1/(1-\gamma), 1/\rho_{\min}, 1/\pi_{\text{ref}})$ up to logarithmic factors.*

These results show that the best choice of α is not $\alpha = 0$ nor $\alpha = 1$, but depends on the desired precision level. This corroborates results from the bandit literature (Zimmert & Seldin, 2021), and gives strong evidence that Tsallis-SoftArgmax with coupled regularization has the potential to accelerate RL algorithms. It also highlights the strength of our framework: one can choose, among multiple parameterizations, the one that is best suited for the problem at hand.

5 EXPERIMENTS

In this section, we demonstrate the generalizability of our framework by showing that our class of parameterizations, with its coupled regularization, can be readily integrated into modern on-policy reinforcement learning algorithms. For this purpose, we introduce and evaluate α -Tsallis PPO, a simple yet principled extension of Proximal Policy Optimization (PPO; Schulman et al. 2017)¹. Our approach is obtained by replacing both the policy parameterization and the entropy regularization in PPO with their Tsallis counterparts (see Appendix J for full experimental details). We compare the performance of α -Tsallis PPO against the standard PPO baseline (Schulman et al., 2017) on two families of environments that we describe below. Our code is available on GitHub: <https://github.com/Labbi-Safwan/f-regularised-policy-gradient>.

Noisy CartPole (Osband et al., 2020). This environment is a variant of the classic CartPole control task in which additive noise is injected into the reward signal. The underlying dynamics is unchanged: at each time step, the agent applies a left or right force to a cart in order to keep an inverted pendulum balanced, receiving a base reward of +1 for each step the pole remains upright, and the episode terminates when stability is lost or after a time limit. However, the reward returned by the environment is perturbed as $\tilde{r}_t = r_t + \sigma \xi_t$, $\xi_t \sim \mathcal{N}(0, 1)$, where $\sigma > 0$ controls the noise level. This preserves the dynamics and optimal policy, but increases the variance of observed returns.

¹Additional experiments on the exact f -PG algorithm are provided in Appendix J.

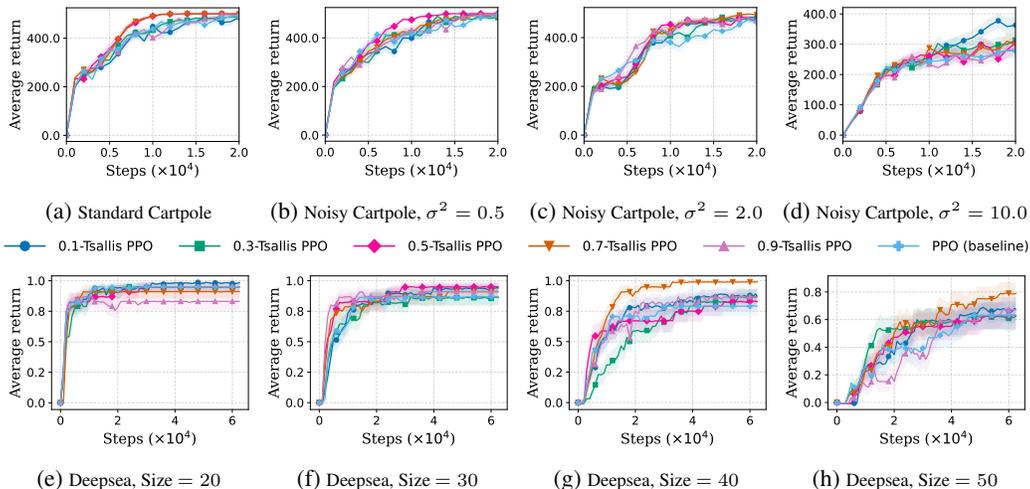


Figure 2: Learning curves for Noisy CartPole (top row) and DeepSea (bottom row) under different choices of the Tsallis parameter α . For Noisy CartPole, we report the standard *un-noised* CartPole environment (a) and reward-noisy variants with increasing noise levels (b–d). For DeepSea, we consider grid sizes $L \in \{20, 30, 40, 50\}$ (e–h). Each curve corresponds to the best temperature and step-size for a given α , and shaded regions indicate ± 1 standard error over 25 seeds. On Noisy CartPole, values $\alpha < 1$ consistently improve performance over the PPO baseline in the standard and low-noise settings, with the gap increasing as the reward noise grows. On DeepSea, the improvement over the PPO baseline becomes more pronounced with increasing L , where $\alpha = 0.7$ achieves the highest returns and the fastest learning.

DeepSea (Osband et al., 2019). DeepSea is an RL environment designed to study deep exploration under sparse rewards. The environment is a directed grid of size $L \times L$. The agent starts in the top-left corner $(0, 0)$ and, at each step, moves downward while choosing between two actions that shift agent’s position either left or right. Thus, each episode lasts exactly L steps and corresponds to selecting a binary action sequence of length L , which defines a unique path through the grid. Only a single trajectory, the one that selects the hidden *correct* (right) action at every depth reaches the rewarding terminal state at $(L-1, L-1)$. However, selecting the right action is not free: every time the agent moves right, it incurs a small movement cost $0.01/L$.

Problem-adaptive couplings yield better performance. Figure 2 illustrates that no single choice of α is uniformly optimal, and that different tasks favor different couplings of parameterization and regularization. On Noisy CartPole, the standard setting (Figure 2a) and the mildly noisy variant (Figure 2b) show a small but systematic advantage for $\alpha < 1$ over the PPO baseline, which becomes more pronounced as the reward noise increases (Figures 2c and 2d). By contrast, on DeepSea, where performance depends on discovering a single sparse-reward trajectory, the hardest instances (Figures 2g and 2h) favor an intermediate value $\alpha = 0.7$. These observations suggest that highly noisy environments and deep exploration problems may benefit from different regions of the Tsallis family, supporting the need for a tunable parameterization-regularization pair.

6 CONCLUSION

We proposed a new class of policy parameterizations based on operators induced by f -divergences. Equipped with a matching f -divergence regularizer, this framework generalizes the classical softmax-entropy pairing and allows flexible alternative parameterizations. Using Tsallis divergence instead of Shannon entropy, we showed that the resulting algorithm yields polynomial, rather than exponential, convergence guarantees for the unregularized RL problem. Empirically, this choice leads to improved performance in exploration-heavy and noisy environments. An important direction for future work is to extend these guarantees for adversarial MDPs, where Tsallis regularization has already proven effective in the bandit setting (Zimmert & Seldin, 2021).

ACKNOWLEDGEMENTS

We would like to thank Vincent Roulet for a fruitful discussion on Mirror Descent and for providing the implementation of f-softargmax operators in JAX. The work of S. Labbi, and P. Mangold has been supported by Technology Innovation Institute (TII), project Fed2Learn. The work of D. Tiapkin has been supported by the Paris Île-de-France Région in the framework of DIM AI4IDF. The work of E. Moulines has been partly funded by the European Union (ERC-2022-SYG-OCEAN-101071601). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

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. URL <http://jmlr.org/papers/v22/19-736.html>.
- Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.
- Mathieu Blondel, André’ F.T. Martins, and Vlad Niculae. Learning with fenchel-young losses. *Journal of Machine Learning Research*, 21(35):1–69, 2020. URL <http://jmlr.org/papers/v21/19-021.html>.
- Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Imre Csiszár. On information-type measure of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2:299–318, 1967.
- John M. Danskin. The theory of max-min, with applications. *SIAM Journal on Applied Mathematics*, 14(4):641–664, 1966. doi: 10.1137/0114053. URL <https://doi.org/10.1137/0114053>.
- Yuhao Ding, Junzi Zhang, Hyunin Lee, and Javad Lavaei. Beyond exact gradients: Convergence of stochastic soft-max policy gradient methods with entropy regularization. *IEEE Transactions on Automatic Control*, pp. 1–16, 2025.
- Mathieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized Markov decision processes. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2160–2169. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/geist19a.html>.
- Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Fundamentals of convex analysis*. Springer Science & Business Media, 2004.
- Chloe Ching-Yun Hsu, Celestine Mendler-Dünger, and Moritz Hardt. Revisiting design choices in proximal policy optimization. *arXiv preprint arXiv:2009.10897*, 2020.
- Anatoli Juditsky, Joon Kwon, and Éric Moulines. Unifying mirror descent and dual averaging. *Mathematical Programming*, 199(1):793–830, May 2023. ISSN 1436-4646. doi: 10.1007/s10107-022-01850-3. URL <https://doi.org/10.1007/s10107-022-01850-3>.
- Sham M Kakade. A natural policy gradient. In T. Dietterich, S. Becker, and Z. Ghahramani (eds.), *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001. URL https://proceedings.neurips.cc/paper_files/paper/2001/file/4b86abe48d358ecf194c56c69108433e-Paper.pdf.

- Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Safwan Labbi, Paul Mangold, Daniil Tiapkin, and Eric Moulines. On global convergence rates for federated softmax policy gradient under heterogeneous environments. In *The 29th International Conference on Artificial Intelligence and Statistics*, 2026. URL <https://openreview.net/forum?id=aPAq9U7MFm>.
- Gen Li, Yuting Wei, Yuejie Chi, and Yuxin Chen. Softmax policy gradient methods can take exponential time to converge. *Mathematical Programming*, 201(1):707–802, Sep 2023. ISSN 1436-4646. doi: 10.1007/s10107-022-01920-6. URL <https://doi.org/10.1007/s10107-022-01920-6>.
- Friedrich Liese and Igor Vajda. On divergences and informations in statistics and information theory. *IEEE Transactions on Information Theory*, 52(10):4394–4412, 2006.
- Jiacai Liu, Wenye Li, and Ke Wei. Elementary analysis of policy gradient methods, 2024. URL <https://arxiv.org/abs/2404.03372>.
- Jiacai Liu, Jinchi Chen, and Ke Wei. On the linear convergence of policy gradient under hadamard parameterization. *Information and Inference: A Journal of the IMA*, 14(1):iaaf003, 2025.
- Andre Martins and Ramon Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. In Maria Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1614–1623, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/martins16.html>.
- Jincheng Mei, Chenjun Xiao, Bo Dai, Lihong Li, Csaba Szepesvari, and Dale Schuurmans. Escaping the gravitational pull of softmax. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 21130–21140. Curran Associates, Inc., 2020a. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/f1cf2a082126bf02de0b307778ce73a7-Paper.pdf.
- Jincheng Mei, Chenjun Xiao, Csaba Szepesvari, and Dale Schuurmans. On the global convergence rates of softmax policy gradient methods. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6820–6829. PMLR, 13–18 Jul 2020b. URL <https://proceedings.mlr.press/v119/mei20b.html>.
- Arthur Mensch and Mathieu Blondel. Differentiable dynamic programming for structured prediction and attention. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3462–3471. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/mensch18a.html>.
- Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87 of *Applied Optimization*. Springer Science & Business Media, 2004. doi: 10.1007/978-1-4419-8853-9.
- Yurii Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1):221–259, Aug 2009. ISSN 1436-4646. doi: 10.1007/s10107-007-0149-x. URL <https://doi.org/10.1007/s10107-007-0149-x>.
- Ian Osband, Benjamin Van Roy, Daniel J Russo, and Zheng Wen. Deep exploration via randomized value functions. *Journal of Machine Learning Research*, 20(124):1–62, 2019.
- Ian Osband, Yotam Doron, Matteo Hessel, John Aslanides, Eren Sezener, Andre Saraiva, Katrina McKinney, Tor Lattimore, Csaba Szepesvári, Satinder Singh, Benjamin Van Roy, Richard Sutton, David Silver, and Hado van Hasselt. Behaviour suite for reinforcement learning. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rygf-kSYwH>.

- Ben Peters, Vlad Niculae, and André F. T. Martins. Sparse sequence-to-sequence models. In Anna Korhonen, David Traum, and Lluís Màrquez (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1504–1519, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1146. URL <https://aclanthology.org/P19-1146/>.
- Martin L. Puterman. *Discounted Markov Decision Problems*, chapter 6, pp. 142–276. John Wiley and Sons, Ltd, 1994. ISBN 9780470316887. doi: <https://doi.org/10.1002/9780470316887.ch6>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470316887.ch6>.
- Alfréd Rényi. On measures of entropy and information. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability, volume 1: contributions to the theory of statistics*, volume 4, pp. 547–562. University of California Press, 1961.
- Vincent Roulet, Tianlin Liu, Nino Vieillard, Michael Eli Sander, and Mathieu Blondel. Loss functions and operators generated by f-divergences. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=V1YfPJdlw>.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1889–1897, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/schulman15.html>.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv:1707.06347*, 2017.
- Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, May 1992. ISSN 1573-0565. doi: 10.1007/BF00992696. URL <https://doi.org/10.1007/BF00992696>.
- Lin Xiao. Dual averaging method for regularized stochastic learning and online optimization. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta (eds.), *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc., 2009. URL https://proceedings.neurips.cc/paper_files/paper/2009/file/7cce53cf90577442771720a370c3c723-Paper.pdf.
- Junzi Zhang, Jongho Kim, Brendan O’Donoghue, and Stephen Boyd. Sample efficient reinforcement learning with reinforce. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12):10887–10895, May 2021. doi: 10.1609/aaai.v35i12.17300. URL <https://ojs.aaai.org/index.php/AAAI/article/view/17300>.
- Julian Zimmert and Yevgeny Seldin. Tsallis-inf: An optimal algorithm for stochastic and adversarial bandits. *Journal of Machine Learning Research*, 22(28):1–49, 2021. URL <http://jmlr.org/papers/v22/19-753.html>.

Appendix

Table of Contents

A	Notations	14
B	Smoothness of the objective	16
B.1	Properties of the soft- f -argmax	16
B.2	Specification to the f -divergence generators of Table 2	20
B.3	Derivatives of the policy	22
B.4	Smoothness of Objective	23
C	Non-Uniform Łojasiewicz inequality	28
C.1	Lower Bounding the norm of the gradient	29
C.2	Bounding the Suboptimality Gap	30
D	Monotone Improvement Operators	31
E	Convergence analysis of Stochastic Policy Gradient	35
E.1	Bounding the bias and variance of the stochastic estimator	35
E.2	Sample complexity of Stochastic f -PG	40
E.3	Guarantees on the non-regularized problem	43
F	Application to common f-divergences	45
F.1	Kullback-Leibler	45
F.2	α -Tsallis	50
G	Discussion on unregularized Policy Gradient with f-SoftArgmax Parameterization	57
H	Links with Mirror Descent	58
I	Technical Lemmas	59
J	Experiments	61
J.1	Uncoupling the parameterization and the regularization	61
J.2	Tabular experiments	61
J.3	Full description of the deepRL experiments	62

A NOTATIONS

f -Divergence. Let $f: (0, \infty) \rightarrow \mathbb{R}$ be a strictly convex generator with $f(1) = 0$. Its adjoint (or reverse generator) is $f_{\dagger}(u) := u f(1/u)$, $u > 0$, which is convex, strictly convex if f is, and satisfies $f_{\dagger}(1) = 0$. Boundary conventions are $f(0) := \lim_{u \downarrow 0} f(u) \in (-\infty, \infty]$ and $f_{\dagger}(0) := \lim_{u \downarrow 0} f_{\dagger}(u) = \lim_{t \uparrow \infty} f(t)/t \in (-\infty, \infty]$. For $p, q \in \mathcal{P}(\mathcal{A})$ over finite \mathcal{A} , the f -divergence is

$$D^f(p||q) := \sum_{a \in \mathcal{A}: q(a) > 0} q(a) f(p(a)/q(a)) + f_{\dagger}(0) \cdot \sum_{a \in \mathcal{A}: q(a) = 0} p(a), \quad (13)$$

with conventions: $q(a)f(0)$ if $q(a) > 0$, $p(a) = 0$, and 0 if $p(a) = q(a) = 0$ (Rényi, 1961; Csiszár, 1967; Liese & Vajda, 2006). f -divergences satisfy $D^f(p||q) \in [0, \infty]$, are jointly convex (Csiszár, 1967), vanish iff $p = q$, and are not symmetric ($D^{f_{\dagger}}(p||q) = D^f(p||q)$).

Symbols	Meaning	Definition
\mathcal{S}	State space	Section 2
\mathcal{A}	Action space	Section 2
γ	Discount factor	Section 2
P	Transition kernel	Section 2
r	Reward function	Section 2
π	Policy	Section 2
π_{ref}	Reference policy used in the regularization problem	Section 2
f	Divergence generator	Section 2
λ	Temperature of the regularization	Section 2
ρ	Initial state distribution	Section 2
κ_f	Upper bound on $ f'''(x)/f''(x)^2 $	$\mathbf{A}_f(\pi_{\text{ref}})$
ω_f	Upper bound on $1/(xf''(x)^2)$	$\mathbf{A}_f(\pi_{\text{ref}})$
\mathbf{d}_f	Upper bound on a set of divergences	(6)
\mathbf{y}_f	Upper bound on a quantity that depends on f'' and f'	(6)
ζ_f	Lower bound on a quantity that depends on f''	(7)
v_π	Value function of a policy π	(1)
v_π^f	Regularized Value function of a policy π	(42)
P_π	Transition kernel induced by policy π	Section 2
q_π^f	Regularized Q-function of a policy π	(47)
d_ρ^π	discounted state visitation of a policy π	(48)
θ	Parameter of the policy (element of $\mathbb{R}^{ \mathcal{S} \mathcal{A} }$)	Section 3
π_θ^f	The soft- f -argmax policy associated with θ	(4)
w_θ^f	A matrix of size $\mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ such that for any $s \in \mathcal{S}$, $w_\theta^f(\cdot s) \in \mathcal{P}(\mathcal{A})$	(5)
$f'_\theta(a s)$	shorthand notation for $f'(\pi_\theta^f(a s)/\pi_{\text{ref}}(a s))$	(35)
$f''_\theta(a s)$	shorthand notation for $f''(\pi_\theta^f(a s)/\pi_{\text{ref}}(a s))$	(36)
$f'''_\theta(a s)$	short hand notation for $f'''(\pi_\theta^f(a s)/\pi_{\text{ref}}(a s))$	(37)
$W_\theta^f(s)$	A function of $f'_\theta(\cdot s)$	(38)
$Y_\theta^f(s)$	A function of $f'_\theta(\cdot s)$ and $f''_\theta(\cdot s)$	(38)
T	Number of iterations performed by f -PG	Algorithm 1
$g_Z^f(\theta)$	Stochastic estimator of the gradient at θ	(11)
H	Truncation horizon in f -PG	Algorithm 1
$g_Z^f(\theta)$	Stochastic estimator of the gradient at θ	(11)
β_f	Bias of the stochastic estimator at θ	Lemma 4.2
σ_f	Variance of the stochastic estimator at θ	Lemma 4.2
L_f	Local smoothness of the objective at θ	Theorem 3.3
$\mathcal{P}(\mathcal{A})$	Set of probability measures over \mathcal{A}	Section 2
$D^f(p q)$	f -divergence between two probability measures p and q	(13)

Distribution of the state-action sequence. The state–action sequence $(S_t, A_t)_{t \geq 0}$ defines a stochastic process on the canonical space $(\mathcal{S} \times \mathcal{A})^{\mathbb{N}}$. For any initial state $s_0 \in \mathcal{S}$, we denote by $\mathbb{P}_{s_0}^\pi$ the law of this process. That is, for any $n \in \mathbb{N}$ and any subset $B \subset (\mathcal{S} \times \mathcal{A})^n$,

$$\mathbb{P}_{s_0}^\pi(B) = \sum_{(a_0, \dots, a_{n-1}) \in \mathcal{A}^n} \sum_{(s_1, \dots, s_{n-1}) \in \mathcal{S}^{n-1}} \mathbb{1}_B((s_0, a_0), \dots, (s_{n-1}, a_{n-1})) \prod_{i=0}^{n-1} \pi(a_i|s_i) P(s_{i+1}|s_i, a_i),$$

with the convention s_0 is the given initial state. We denote by $\mathbb{E}_{s_0}^\pi$ the corresponding expectation operator. In particular, the state sequence $(s_t)_{t \geq 0}$ defines a Markov reward process (Section 2.1.6 in Puterman (1994)) with transition kernel

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

Norms. For $x \in \mathbb{R}^d$, we define the norms

$$\|x\|_\infty = \max_{i \in \{1, \dots, d\}} |x_i| , \quad \|x\|_1 = \sum_{i=1}^d |x_i| , \quad \|x\|_2 = \left(\sum_{i=1}^d |x_i|^2 \right)^{1/2} .$$

For a $d \times d$ matrix M , we denote by $\|M\|_\infty$, and $\|M\|_2$ respectively the max *row* sum, and the spectral norm:

$$\|M\|_\infty = \sup_{x \neq 0} \{\|Mx\|_\infty / \|x\|_\infty\} = \sup_{i \in \{1, \dots, d\}} \sum_{j=1}^d |M_{i,j}|, \quad \|M\|_2 = \sup_{x \neq 0} \{\|Mx\|_2 / \|x\|_2\}. \quad (14)$$

Recall that, for any $x \in \mathbb{R}^d$, $\|Mx\|_\infty \leq \|M\|_\infty \|x\|_\infty$ and $\|Mx\|_2 \leq \|M\|_2 \|x\|_2$.

Functional and matrix forms. For notational convenience, we also view P as a $(|\mathcal{S}| \cdot |\mathcal{A}|) \times |\mathcal{S}|$ matrix with entries $P_{(s,a),s'} = P(s' | s, a)$. Similarly, v_π^f is a vector of size $|\mathcal{S}|$ and q_π^f a vector of size $|\mathcal{S}| \times |\mathcal{A}|$. Finally, we identify the parameter $\theta \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ with its matrix representation $\theta \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$, indexed by $(s, a) \in \mathcal{S} \times \mathcal{A}$. This slight abuse of notation allows us to conveniently switch between functional and matrix views.

B SMOOTHNESS OF THE OBJECTIVE

In this section, we establish the smoothness of the regularized value function $v_\theta^f := v_{\pi_\theta}^f(\rho)$ with respect to the parameter θ . As a first step, we show that the policy π_θ^f is smooth under suitable assumptions on the divergence generator f and compute its first and second derivatives. To do so, we start by studying the properties of the soft- f -argmax operator and then apply the obtained results to derive properties of the policy π_θ^f .

B.1 PROPERTIES OF THE SOFT- f -ARGMAX

In this Section, we compute the derivative of $\nu_x^f(\cdot) := \text{f-softargmax}(x, \nu_{\text{ref}})$ and $\text{x-softmax}(\cdot) := \text{f-softmax}(x, \nu_{\text{ref}})$, where

$$\text{f-softmax}(x, \nu_{\text{ref}}) := \max_{\nu \in \mathcal{P}(\mathcal{A})} \left\{ \langle \nu, x \rangle - D^f(\nu \| \nu_{\text{ref}}) \right\}. \quad (15)$$

The function $\text{x-softmax}(\cdot)$ is the Fenchel–Legendre transform of $D^f(\cdot \| \nu_{\text{ref}})$, and the results in this section are therefore standard results from convex analysis, statements of which can be found in various forms in [Hiriart-Urruty & Lemaréchal \(2004\)](#); [Mensch & Blondel \(2018\)](#); [Geist et al. \(2019\)](#); [Roulet et al. \(2025\)](#).

In this section, we fix a reference probability distribution $\nu_{\text{ref}} \in \mathcal{P}(\mathcal{A})$ such that for any $a \in \mathcal{A}$, we have $\nu_{\text{ref}}(a) > \underline{\nu}_{\text{ref}}$ for some $\min_{a \in \mathcal{A}} \nu_{\text{ref}}(a) > \underline{\nu}_{\text{ref}} > 0$. For a given $x \in \mathbb{R}^{|\mathcal{A}|}$, we define

$$\nu_x^f(\cdot) := \text{f-softargmax}(x, \nu_{\text{ref}}) = \arg \max_{\nu \in \mathcal{P}(\mathcal{A})} \left\{ \langle \nu, x \rangle - D^f(\nu \| \nu_{\text{ref}}) \right\}. \quad (16)$$

The following result is a simplified version of ([Roulet et al., 2025](#), Proposition 1). For the sake of completeness, we provide a full proof.

Lemma B.1. *Assume that f is strictly convex on $[0, 1/\underline{\nu}_{\text{ref}}]$, differentiable on $(0, 1/\underline{\nu}_{\text{ref}})$, with $\lim_{x \downarrow 0^+} f'(u) = -\infty$, for some $\underline{\nu}_{\text{ref}} > 0$. Let ν_{ref} be a policy such that $\min_{a \in \mathcal{A}} \nu_{\text{ref}}(a) > \underline{\nu}_{\text{ref}}$. For any $x \in \mathbb{R}^{|\mathcal{A}|}$ and $a \in \mathcal{A}$, we have $0 < \nu_x^f(a)$. Moreover, for all $x \in \mathbb{R}^{|\mathcal{A}|}$, there exists a unique $\mu_x \in I_{\underline{\nu}_{\text{ref}}}(x)$, where*

$$I_{\underline{\nu}_{\text{ref}}}(x) := \left(\max_{a \in \mathcal{A}} x(a) - f'(1/\underline{\nu}_{\text{ref}}), \max_{a \in \mathcal{A}} x(a) - f'(1) \right), \quad (17)$$

such that

$$\nu_x^f(a) = \nu_{\text{ref}}(a) [f']^{-1}(x(a) - \mu_x). \quad (18)$$

Moreover, $\mu_x \in \mathbb{R}$ is the unique root of the equation

$$F(x, \mu) := \sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) [f']^{-1}(x(a) - \mu) - 1 = 0 \quad \text{for } \mu \in I_{\underline{\nu}_{\text{ref}}}(x). \quad (19)$$

Proof. Under the stated assumption, the map f' is strictly increasing on $(0, 1/\underline{\nu}_{\text{ref}}]$, hence injective; therefore it is invertible *onto its image*. Moreover,

$$\text{Dom}([f']^{-1}) = f'((0, 1/\underline{\nu}_{\text{ref}}]) = (-\infty, f'(1/\underline{\nu}_{\text{ref}})], \quad (20)$$

which is an interval and the function $[f']^{-1}$ is strictly increasing.

Fix $x \in \mathcal{A}$, and $(a, b) \in \mathcal{A} \times \mathcal{A}$. Recall from (16) the definition of the soft- f -argmax ,

$$\nu_x^f = \arg \max_{\nu \in \mathcal{P}(\mathcal{A})} \left\{ \langle \nu, x \rangle - D^f(\nu \| \nu_{\text{ref}}) \right\}$$

This is a strictly concave optimization problem over the probability simplex $\mathcal{P}(\mathcal{A})$ so it admits a unique maximizer. We now characterize the maximizer via the KKT conditions. Introduce multipliers $\mu \in \mathbb{R}$ for the equality constraint $\sum_{c \in \mathcal{A}} \nu(c) = 1$, and for every $c \in \mathcal{A}$, $\lambda(c) \in \mathbb{R}^+$ for the non-negativity constraints $\nu(c) \geq 0$. The Lagrangian reads

$$L(\nu, \mu, \{\lambda(c)\}_{c \in \mathcal{A}}) = \sum_{c \in \mathcal{A}} \nu(c) x(c) - D^f(\nu \| \nu_{\text{ref}}) + \mu \left(1 - \sum_{c \in \mathcal{A}} \nu(c) \right) - \sum_{c \in \mathcal{A}} \lambda(c) \nu(c) .$$

By the differentiability of f , differentiating the Lagrangian with respect to $\nu(a)$ gives

$$\frac{\partial L}{\partial \nu(a)} = x(a) - f' \left(\frac{\nu(a)}{\nu_{\text{ref}}(a)} \right) - \mu - \lambda(a) . \quad (21)$$

At the optimum $(\nu_x^f, \mu_x, \{\lambda_x(c)\}_{c \in \mathcal{A}})$, the KKT conditions yield:

$$\nu_x^f(c) \lambda_x(c) = 0, \quad \forall c \in \mathcal{A}, \quad (22)$$

$$x(c) - f' \left(\frac{\nu_x^f(c)}{\nu_{\text{ref}}(c)} \right) - \mu_x - \lambda_x(c) = 0, \quad \forall c \in \mathcal{A}. \quad (23)$$

Under the stated assumptions, $\lim_{x \rightarrow 0^+} f'(x) = -\infty$. Hence, if $\nu_x^f(c) = 0$ for some c , the stationarity condition (23) cannot hold with finite multipliers. Therefore $\nu_x^f(c) > 0$ for all $c \in \mathcal{A}$, which by (22) implies $\lambda_x(c) = 0$. Thus, for each $c \in \mathcal{A}$ the stationarity condition reduces to

$$x(c) - f' \left(\frac{\nu_x^f(c)}{\nu_{\text{ref}}(c)} \right) = \mu_x . \quad (24)$$

Note also that (24) also implies that for all $a \in \mathcal{A}$, $x(a) - \mu_x \in \text{Dom}(f')$, which implies, using (20) that $\max_{a \in \mathcal{A}} x(a) - f'(1/\underline{\nu}_{\text{ref}}) \leq \mu_x$. Together with (24), this shows (18). Note that μ_x is a root of (19), since using (18),

$$\sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) [f']^{-1}(x(a) - \mu_x) = \sum_{a \in \mathcal{A}} \nu_x^f(a) = 1.$$

Because $[f']^{-1}$ is strictly increasing on $(-\infty, f'(1/\underline{\nu}_{\text{ref}})]$, for each $x \in \mathbb{R}^{|\mathcal{A}|}$, the function $\mu \mapsto F(x, \mu)$ (see (19)) is strictly decreasing on $(-\infty, f'(1/\underline{\nu}_{\text{ref}})]$. Strict monotonicity gives the uniqueness of μ_x . Note finally that if $\mu > \max_{a \in \mathcal{A}} x(a) - f'(1)$, then $x(a) - \mu < f'(1)$, and since $[f']^{-1}$ is strictly increasing, $[f']^{-1}(x(a) - \mu) < 1$, showing that $F(x, \mu) < 0$, which concludes the proof. \square

Remark B.2 Let $f : (0, \tau) \rightarrow \mathbb{R}$ be a strictly convex and differentiable on $(0, \tau)$, $\text{Dom}(f') = (0, \tau)$, and f' is strictly increasing and continuous on $(0, \tau)$. Let $\alpha := \lim_{x \rightarrow 0} f'(x) \in [-\infty, +\infty)$ and $\beta := \lim_{x \rightarrow \tau} f'(x) \in (-\infty, +\infty]$. Then $f'((0, \tau)) = (\alpha, \beta)$, i.e., $f' : (0, \tau) \rightarrow (\alpha, \beta)$ is a strictly increasing bijection (hence admits a continuous inverse $[f']^{-1} : (\alpha, \beta) \rightarrow (0, \tau)$). Define the convex conjugate of f ,

$$f^*(y) := \sup_{x \in (0, \tau)} \{xy - f(x)\} . \quad (25)$$

The two following properties hold,

- (i) For every $y \in (\alpha, \beta)$, the supremum is attained at a unique point $x = (f')^{-1}(y)$.

- (ii) $\text{Dom}(f^*) \subseteq [\alpha, \beta]$ (with the convention that an infinite endpoint is excluded). Specifically, $\text{Dom}(f^*) \cap (\alpha, \beta) = (\alpha, \beta)$, and $f^*(y) = +\infty$ for $y \notin [\alpha, \beta]$. At an endpoint $y = \alpha$ (resp. $y = \beta$), if finite, $f^*(y) = \lim_{x \downarrow 0} (yx - f(x))$ (resp. $\lim_{x \uparrow \tau} (yx - f(x))$), so $f^*(\alpha)$ (resp. $f^*(\beta)$) is finite iff the corresponding one-sided limit is finite.

On the open interval (α, β) , the function f^* is differentiable and

$$(f^*)'(y) = [f']^{-1}(y), \quad y \in (\alpha, \beta),$$

We retrieve the statement in Proposition 1 of Roulet et al. (2025) by replacing $[f']^{-1}$ by $(f^*)'$. In most examples, there is no need to resort to the convex conjugate to compute the inverse.

Lemma B.3. Assume, in addition to Lemma B.1, that f is two-times continuously differentiable on $(0, \nu_{\text{ref}})$. Then, the function $x \mapsto \mu_x$ is continuously differentiable on $\mathbb{R}^{|\mathcal{A}|}$, and for all $a \in \mathcal{A}$,

$$\frac{\partial}{\partial x(a)} \mu_x = -\frac{\partial}{\partial x(a)} F(x, \mu_x) / \frac{\partial}{\partial \mu} F(x, \mu_x) \quad (26)$$

Proof. The function $(x, \mu) \mapsto F(x, \mu)$ is two-times continuously differentiable on the open set $U := \{(x, \mu) : x \in \mathbb{R}^{|\mathcal{A}|}, \mu \in I_{\nu_{\text{ref}}}(x)\} \subset \mathbb{R}^{|\mathcal{A}|} \times \mathbb{R}$. Let $x_0 \in \mathbb{R}^{|\mathcal{A}|}$. Since $[f']^{-1}$ is strictly increasing,

$$\frac{\partial}{\partial \mu} F(x, \mu_x) = \sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) \frac{1}{f''([f']^{-1}(x_0(a) - \mu_{x_0}))} > 0.$$

Hence, we may apply the implicit function theorem, which shows that there exists an open neighborhood V_{x_0} and a unique function $x \mapsto \mu_x$ on V_{x_0} , such that for all $x \in V_{x_0}$, $F(x, \mu_x) = 0$ and (26) holds. \square

We now introduce some compact notations that will be used throughout the sequel. For any $a \in \mathcal{A}$, define the first three derivatives of f evaluated at the probability ratio $\nu_x^f(a)/\nu_{\text{ref}}(a)$:

$$f'_x(a) := f' \left(\frac{\nu_x^f(a)}{\nu_{\text{ref}}(a)} \right), \quad f''_x(a) := f'' \left(\frac{\nu_x^f(a)}{\nu_{\text{ref}}(a)} \right), \quad f'''_x(a) := f''' \left(\frac{\nu_x^f(a)}{\nu_{\text{ref}}(a)} \right). \quad (27)$$

In addition, we introduce the quantities

$$W_x^f := \sum_{a \in \mathcal{A}} \frac{\nu_{\text{ref}}(a)}{f''_x(a)}. \quad (28)$$

Importantly, as f is strictly convex, its second derivative is strictly positive, making the preceding quantity well-defined. For any $a \in \mathcal{A}$, we also define the following normalized weights

$$w_x^f(a) = \frac{1}{W_x^f} \frac{\nu_{\text{ref}}(a)}{f''_x(a)}. \quad (29)$$

Lemma B.4. Assume $\mathbf{A}_f(\nu_{\text{ref}})$. Then the soft- f -argmax ν_x^f is twice continuously differentiable with respect to x . Moreover, for any $x \in \mathbb{R}^{|\mathcal{A}|}$, and $(a, b) \in \mathcal{A}^2$, we have

$$\frac{1}{W_x^f} \cdot \frac{\partial \nu_x^f(a)}{\partial x(b)} = \mathbf{1}_b(a) w_x^f(a) - w_x^f(a) w_x^f(b),$$

In addition, for any $(a, b, c) \in \mathcal{A}^3$, the second derivative satisfies

$$\begin{aligned} \frac{1}{W_x^f} \cdot \frac{\partial \nu_x^f(a)}{\partial x(b) \partial x(c)} &= -\mathbf{1}_b(a) \mathbf{1}_c(a) \frac{f'''_x(a)}{f''_x(a)^2} \cdot w_x^f(a) + \mathbf{1}_c(b) w_x^f(a) w_x^f(b) \frac{f'''_x(b)}{f''_x(b)^2} \\ &\quad + \mathbf{1}_b(a) w_x^f(a) w_x^f(c) \frac{f'''_x(a)}{f''_x(a)^2} + \mathbf{1}_c(a) w_x^f(a) w_x^f(b) \frac{f'''_x(a)}{f''_x(a)^2} \\ &\quad - w_x^f(a) \cdot w_x^f(b) \cdot w_x^f(c) \cdot \left[\frac{f'''_x(a)}{f''_x(a)^2} + \frac{f'''_x(b)}{f''_x(b)^2} + \frac{f'''_x(c)}{f''_x(c)^2} \right] \\ &\quad + w_x^f(a) \cdot w_x^f(b) \cdot w_x^f(c) \cdot \sum_{d \in \mathcal{A}} w_x^f(d) \cdot \frac{f'''_x(d)}{f''_x(d)^2}. \end{aligned}$$

Proof. Fix $x \in \mathbb{R}^{|\mathcal{A}|}$ and $(a, b, c) \in \mathcal{A}^3$. Define

$$F(x; \mu) = \sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) [f']^{-1}(x(a) - \mu_x) - 1 .$$

First derivative. Importantly, using Lemma B.1 we have that

$$F(x; \mu_x) = 0 .$$

Differentiating the previous identity with respect to $x(b)$, yields

$$\frac{\partial F(x; \mu_x)}{\partial x(b)} = \sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) \frac{1}{f''([f']^{-1}(x(a) - \mu_x))} \left(\mathbf{1}_b(a) - \frac{\partial \mu_x}{\partial x(b)} \right) ,$$

where we used that the derivative of $[f']^{-1}$ is $1/f''([f']^{-1})$. Next, using from Lemma B.1 that $\nu_x^f(a) = \nu_{\text{ref}}(a)[f']^{-1}(x(a) - \mu_x)$ yields

$$\frac{\partial F(x; \mu_x)}{\partial x(b)} = \sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) \frac{1}{f''_x(a)} \left(\mathbf{1}_b(a) - \frac{\partial \mu_x}{\partial x(b)} \right) = 0 ,$$

where $f''_x(a)$ is defined in (27). This implies

$$\frac{\partial \mu_x}{\partial x(b)} = w_x^f(b) , \quad (30)$$

where $w_x^f(b)$ is defined in (29). Now that we have computed the derivative of the normalization factor μ_x , we can compute the derivative of the policy. Starting from Lemma B.1, we have that

$$\nu_x^f(a) = \nu_{\text{ref}}(a)[f']^{-1}(x(a) - \mu_x) .$$

Differentiating the previous identity with respect to $x(b)$, yields

$$\frac{\partial \nu_x^f(a)}{\partial x(b)} = \frac{\nu_{\text{ref}}(a)}{f''_x(a)} \left[\mathbf{1}_b(a) - \frac{\partial \mu_x}{\partial x(b)} \right] = \frac{\nu_{\text{ref}}(a)}{f''_x(a)} [\mathbf{1}_b(a) - w_x^f(b)] ,$$

where in the last identity, we used the expression of the derivative of μ_x given in (30). Finally, using the definition of W_x^f given in (28) gives

$$\boxed{\frac{\partial \nu_x^f(a)}{\partial x(b)} = \mathbf{1}_b(a) \frac{\nu_{\text{ref}}(a)}{f''_x(a)} - \frac{\nu_{\text{ref}}(a)}{f''_x(a)} \frac{\nu_{\text{ref}}(b)}{f''_x(b)} \cdot \frac{1}{W_x^f}} , \quad (31)$$

Second derivative From (31), we aim to differentiate once more with respect to $x(c)$. First, note that it holds that

$$\frac{\partial f''_x(a)}{\partial x(c)} = \frac{f'''_x(a)}{\nu_{\text{ref}}(a)} \cdot \frac{\partial \nu_x^f(a)}{\partial x(c)} = \mathbf{1}_c(a) \frac{f'''_x(a)}{f''_x(a)} - \frac{\nu_{\text{ref}}(c) f'''_x(a)}{f''_x(a) f''_x(c)} \cdot \frac{1}{W_x^f} , \quad (32)$$

$$\frac{\partial W_x^f}{\partial x(c)} = - \sum_{d \in \mathcal{A}} \frac{\nu_{\text{ref}}(d)}{f''_x(d)^2} \frac{\partial f''_x(d)}{\partial x(c)} = - \frac{\nu_{\text{ref}}(c) f'''_x(c)}{f''_x(c)^3} + \frac{1}{W_x^f} \sum_{d \in \mathcal{A}} \frac{\nu_{\text{ref}}(c) \nu_{\text{ref}}(d) f'''_x(d)}{f''_x(d)^3 f''_x(c)} , \quad (33)$$

Now computing the second derivative of ν_x^f gives

$$\begin{aligned} \frac{\partial \nu_x^f(a)}{\partial x(b) \partial x(c)} &= -\mathbf{1}_b(a) \frac{\nu_{\text{ref}}(a)}{f''_x(a)^2} \frac{\partial f''_x(a)}{\partial x(c)} + \frac{\nu_{\text{ref}}(a) \nu_{\text{ref}}(b)}{f''_x(a)^2 f''_x(b)} \cdot \frac{1}{W_x^f} \frac{\partial f''_x(a)}{\partial x(c)} \\ &\quad + \frac{\nu_{\text{ref}}(a) \nu_{\text{ref}}(b)}{f''_x(a) f''_x(b)^2} \cdot \frac{1}{W_x^f} \frac{\partial f''_x(b)}{\partial x(c)} + \frac{\nu_{\text{ref}}(a) \nu_{\text{ref}}(b)}{f''_x(a) f''_x(b)} \cdot \frac{1}{(W_x^f)^2} \frac{\partial W_x^f}{\partial x(c)} \end{aligned}$$

Plugging in (32), and (33) in the preceding inequality yields

$$\frac{\partial \nu_x^f(a)}{\partial x(b) \partial x(c)} = -\mathbf{1}_b(a) \frac{\nu_{\text{ref}}(a)}{f''_x(a)^2} \left[\mathbf{1}_c(a) \frac{f'''_x(a)}{f''_x(a)} - \frac{\nu_{\text{ref}}(c) f'''_x(a)}{f''_x(a) f''_x(c)} \cdot \frac{1}{W_x^f} \right]$$

$$\begin{aligned}
& + \frac{\nu_{\text{ref}}(a)\pi_{\text{ref}}(b)}{f_x''(a)^2 f_x''(b)} \cdot \frac{1}{W_x^f} \left[\mathbf{1}_c(a) \frac{f_x'''(a)}{f_x''(a)} - \frac{\nu_{\text{ref}}(c) f_x'''(a)}{f_x''(a) f_x''(c)} \cdot \frac{1}{W_x^f} \right] \\
& + \frac{\nu_{\text{ref}}(a)\pi_{\text{ref}}(b)}{f_x''(a) f_x''(b)^2} \cdot \frac{1}{W_x^f} \left[\mathbf{1}_c(b) \frac{f_x'''(b)}{f_x''(b)} - \frac{\nu_{\text{ref}}(c) f_x'''(b)}{f_x''(b) f_x''(c)} \cdot \frac{1}{W_x^f} \right] \\
& + \frac{\nu_{\text{ref}}(a)\nu_{\text{ref}}(b)}{f_x''(a) f_x''(b)} \cdot \frac{1}{(W_x^f)^2} \left[-\frac{\nu_{\text{ref}}(c) f_x'''(c)}{f_x''(c)^3} + \frac{1}{W_x^f} \sum_{d \in \mathcal{A}} \frac{\nu_{\text{ref}}(c) \nu_{\text{ref}}(d) f_x'''(d)}{f_x''(d)^3 f_x''(c)} \right],
\end{aligned}$$

which concludes the proof. \square

The following lemma links the gradients f-softmax and f-softargmax operators.

Lemma B.5. Assume $\mathbf{A}_f(\nu_{\text{ref}})$. For any $x \in \mathbb{R}^{|\mathcal{A}|}$, it holds that

$$\frac{\partial \text{f-softmax}(x, \nu_{\text{ref}})}{\partial x} = \text{f-softargmax}(x, \nu_{\text{ref}}) \quad \text{and} \quad \left\| \frac{\partial^2 \text{f-softmax}(x, \nu_{\text{ref}})}{\partial x^2} \right\|_2 \leq 2 W_x^f .$$

Proof. For any $x \in \mathcal{A}$ and $\nu \in \mathcal{P}(\mathcal{A})$, define

$$h^f(x, \nu) = \langle \nu, x \rangle - D^f(\nu \| \nu_{\text{ref}}) . \quad (34)$$

Fix $b \in \mathcal{A}$ and note that

$$\frac{\partial h^f(x, \nu)}{\partial x(b)} = \nu(b) .$$

It holds that $\text{f-softmax}(x) = \max_{\nu \in \mathcal{P}(\mathcal{A})} h^f(x, \nu)$. As h^f is continuous in its two variables, $\mathcal{P}(\mathcal{A})$ is a compact set, and for every $x \in \mathbb{R}^{|\mathcal{A}|}$, the function $h^f(x, \cdot)$ admits a unique optimizer in $\mathcal{P}(\mathcal{A})$, then by Danskin's theorem (Lemma I.5)

$$\frac{\partial \text{f-softmax}(x, \nu_{\text{ref}})}{\partial x(b)} = \frac{\partial h^f(x, \nu^*(x))}{\partial x} = \nu^*(x) , \quad \text{where } \nu^*(x) = \arg \max_{\nu \in \mathcal{P}(\mathcal{A})} h^f(x, \nu) .$$

Finally, using that

$$\nu^*(x) = \text{f-softargmax}(x, \nu_{\text{ref}}) ,$$

establishes the first identity of the lemma. Using the fact that for any matrix $A \in \mathbb{R}^{d \times d}$, we have $\|A\|_2 \leq \sum_{i=1}^d \sum_{j=1}^d |a_{i,j}|$, implies

$$\left\| \frac{\partial^2 \text{f-softmax}(x, \nu_{\text{ref}})}{\partial x^2} \right\|_2 \leq \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{A}} \left| \frac{\partial \nu_x^f(a)}{\partial x(b)} \right| \leq \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{A}} W_x^f |1_b(a) w_x^f(a) - w_x^f(a) w_x^f(b)|$$

where in the last equality, we used Lemma B.4. Finally, applying the triangle inequality establishes the second claim of the lemma. \square

B.2 SPECIFICATION TO THE \mathbf{f} -DIVERGENCE GENERATORS OF TABLE 2

We now specify the particular forms that these quantities take under three choices of the function f : the KL divergence generator ($f(u) = u \log u$), the α -Csiszár–Cressie–Read divergence generator for $0 < \alpha < 1$, and the Hellinger divergence generator

KL case ($f(u) = u \log u$). Since $f'(u) = \log u + 1$, $f''(u) = 1/u$, and $[f']^{-1}(y) = \exp(y - 1)$, the soft- f operators specialize as follows (with base measure $\nu_{\text{ref}} \in \mathcal{P}(\mathcal{A})$):

$$\nu_x^{\text{KL}}(a) = \nu_{\text{ref}}(a) [f']^{-1}(x(a) - \mu_x) = \frac{\nu_{\text{ref}}(a) \exp(x(a))}{\sum_{b \in \mathcal{A}} \nu_{\text{ref}}(b) \exp(x(b))}, \quad a \in \mathcal{A},$$

where the normalizer is

$$\mu_x = -1 + \log \left(\sum_{b \in \mathcal{A}} \nu_{\text{ref}}(b) \exp(x(b)) \right).$$

The associated softmax function is the log-partition:

$$\text{KL-softmax}(x, \nu_{\text{ref}}) = \log \left(\sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) \exp(x(a)) \right).$$

For the curvature quantities in (28)–(29), since

$$f'' \left(\frac{\nu_x^{\text{KL}}(a)}{\nu_{\text{ref}}(a)} \right) = \left(\frac{\nu_x^{\text{KL}}(a)}{\nu_{\text{ref}}(a)} \right)^{-1} = \frac{\nu_{\text{ref}}(a)}{\nu_x^{\text{KL}}(a)},$$

we obtain

$$W_x^{\text{KL}} = \sum_{a \in \mathcal{A}} \frac{\nu_{\text{ref}}(a)}{f''(a)} = \sum_{a \in \mathcal{A}} \nu_x^{\text{KL}}(a) = 1$$

and the corresponding normalized weights are

$$w_x^{\text{KL}}(a) = \frac{1}{W_x^{\text{KL}}} \frac{\nu_{\text{ref}}(a)}{f''(a)} = \nu_x^{\text{KL}}(a), \quad a \in \mathcal{A}.$$

Tsallis- α case ($0 < \alpha < 1$). Let $f(u) = \frac{u^\alpha - \alpha u + \alpha - 1}{\alpha(\alpha - 1)}$, so that

$$f'(u) = \frac{u^{\alpha-1} - 1}{\alpha - 1}, \quad f''(u) = u^{\alpha-2}, \quad [f']^{-1}(y) = [1 + (\alpha - 1)y]^{\frac{1}{\alpha-1}}.$$

The soft- f operators specialize (with base measure $\nu_{\text{ref}} \in \mathcal{P}(\mathcal{A})$) to

$$\nu_x^{\text{TS}}(a) = \nu_{\text{ref}}(a) [f']^{-1}(x(a) - \mu_x) = \nu_{\text{ref}}(a) [1 + (\alpha - 1)(x(a) - \mu_x)]^{1/(\alpha-1)}, \quad a \in \mathcal{A},$$

where μ_x is the unique normalizer satisfying the constraint

$$\sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) [1 + (\alpha - 1)(x(a) - \mu_x)]^{\frac{1}{\alpha-1}} = 1,$$

with the domain condition $1 + (\alpha - 1)(x(a) - \mu_x) > 0$ for all $a \in \mathcal{A}$.

The associated softmax is

$$\text{TS-softmax}(x, \nu_{\text{ref}}) = \mu_x - \frac{1}{\alpha} + \frac{1}{\alpha} \sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) [1 + (\alpha - 1)(x(a) - \mu_x)]^{\frac{\alpha}{\alpha-1}}.$$

For the curvature quantities in (28)–(29), set

$$u_x(a) := \frac{\nu_x^{\text{TS}}(a)}{\nu_{\text{ref}}(a)} = [1 + (\alpha - 1)(x(a) - \mu_x)]^{\frac{1}{\alpha-1}}.$$

Since $f''(u_x(a)) = u_x(a)^{\alpha-2}$, we have

$$W_x^{\text{TS}} = \sum_{a \in \mathcal{A}} \frac{\nu_{\text{ref}}(a)}{f''(a)} = \sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) [1 + (\alpha - 1)(x(a) - \mu_x)]^{(2-\alpha)/(\alpha-1)},$$

and the corresponding normalized weights are, for $a \in \mathcal{A}$,

$$w_x^{\text{TS}}(a) = \frac{1}{W_x^{\text{TS}}} \frac{\nu_{\text{ref}}(a)}{f''(a)} = \frac{\nu_{\text{ref}}(a) [1 + (\alpha - 1)(x(a) - \mu_x)]^{(2-\alpha)/(\alpha-1)}}{\sum_{c \in \mathcal{A}} \nu_{\text{ref}}(c) [1 + (\alpha - 1)(x(c) - \mu_x)]^{(2-\alpha)/(\alpha-1)}}.$$

Jensen-Shannon. Let $f(u) = \frac{1}{2} (u \log(u) - (u + 1) \log(\frac{u+1}{2}))$. In this case, we have

$$f'(u) = \frac{1}{2} \log \left(\frac{2u}{u+1} \right), \quad f''(u) = \frac{1}{2u(u+1)}, \quad [f']^{-1}(u) = \frac{\exp(2u)}{2 - \exp(2u)}.$$

The JS-softargmax operator specialize (with base measure $\nu_{\text{ref}} \in \mathcal{P}(\mathcal{A})$) to

$$\nu_x^{\text{JS}}(a) = \nu_{\text{ref}}(a) [f']^{-1}(x(a) - \mu_x) = \nu_{\text{ref}}(a) \frac{\exp(2(x(a) - \mu_x))}{2 - \exp(2(x(a) - \mu_x))}, \quad a \in \mathcal{A},$$

where μ_x is the unique normalizer satisfying the constraint

$$\sum_{a \in \mathcal{A}} \nu_{\text{ref}}(a) \frac{\exp(2(x(a) - \mu_x))}{2 - \exp(2(x(a) - \mu_x))} = 1,$$

B.3 DERIVATIVES OF THE POLICY

Next, we exploit the expression of the derivatives of the soft- f -argmax derived in Lemma B.4 to compute the first and second derivatives of the policy. We begin by extending the notations defined in (27), (28), and (29) to encompass a dependence on the state. For any pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, define the first three derivatives of f evaluated at the likelihood ratio $\pi_\theta^f(a|s)/\pi_{\text{ref}}(a|s)$:

$$f'_\theta(a|s) := f'_{\theta(s,\cdot)}(a) = f' \left(\frac{\pi_\theta^f(a|s)}{\pi_{\text{ref}}(a|s)} \right), \quad (35)$$

$$f''_\theta(a|s) := f''_{\theta(s,\cdot)}(a) = f'' \left(\frac{\pi_\theta^f(a|s)}{\pi_{\text{ref}}(a|s)} \right), \quad (36)$$

$$f'''_\theta(a|s) := f'''_{\theta(s,\cdot)}(a) = f''' \left(\frac{\pi_\theta^f(a|s)}{\pi_{\text{ref}}(a|s)} \right). \quad (37)$$

In addition, for every $s \in \mathcal{S}$ we introduce the quantities

$$W_\theta^f(s) := W_{\theta(s,\cdot)}^f = \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''_\theta(a|s)}, \quad \text{and} \quad Y_\theta^f(s) := \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''_\theta(a|s)} |f'_\theta(a|s)|. \quad (38)$$

For every $(s, a) \in \mathcal{S} \times \mathcal{A}$, we define the normalized weights

$$w_\theta^f(a|s) = \frac{1}{W_\theta^f(s)} \frac{\pi_{\text{ref}}(a|s)}{f''_\theta(a|s)}. \quad (39)$$

The following lemma provides bounds on several key quantities that will appear in the appendix.

Lemma B.6. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$, respectively. For any parameter $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, it holds that*

$$\left\| W_\theta^f \right\|_\infty \leq \omega_f, \quad \left\| Y_\theta^f \right\|_\infty \leq y_f, \quad \max_{s \in \mathcal{S}} D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) \leq d_f, \quad \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{w_\theta^f(a|s)}{\pi_\theta^f(a|s)} \leq \omega_f.$$

Proof. The proof follows from the definition of the different quantities and $\mathbf{A}_f(\pi_{\text{ref}})$ \square

Using Lemma B.4, we get the following expression for the derivatives of the policy.

Corollary B.7. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$, respectively. Then the policy π_θ^f is twice continuously differentiable with respect to θ . Additionally, for all $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$ and $s \in \mathcal{S}$, there exists a unique $\mu_\theta(s) \in \mathbb{R}$ such that for any $a \in \mathcal{A}$, we have*

$$\pi_\theta^f(a|s) = \pi_{\text{ref}}(a|s) [f']^{-1}(\theta(s, a) - \mu_\theta(s)). \quad (40)$$

For any $s \in \mathcal{S}$, the $\theta \mapsto \mu_\theta(s)$ is continuously differentiable. For any $s' \neq s$, $\partial/\partial\theta(s', \cdot)\mu_\theta(s) = 0$ and

$$\frac{\partial \mu_\theta(s)}{\partial \theta(s, \cdot)} = w_\theta^f(\cdot|s).$$

Moreover, for any $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, $s \in \mathcal{S}$, and $(a, b) \in \mathcal{A} \times \mathcal{A}$, we have

$$\frac{1}{W_\theta^f(s)} \cdot \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, b)} = \mathbf{1}_b(a) w_\theta^f(a|s) - w_\theta^f(a|s) w_\theta^f(b|s),$$

In addition, for any $(a, b, c) \in \mathcal{A}^3$, the second derivative satisfies

$$\begin{aligned} \frac{1}{W_\theta^f(s)} \cdot \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, b) \partial \theta(s, c)} &= -\mathbf{1}_b(a) \mathbf{1}_c(a) \frac{f'''_\theta(a|s)}{f''_\theta(a|s)^2} \cdot w_\theta^f(a|s) + \mathbf{1}_c(b) w_\theta^f(a|s) w_\theta^f(b|s) \frac{f'''_\theta(b|s)}{f''_\theta(b|s)^2} \\ &\quad + \mathbf{1}_b(a) w_\theta^f(a|s) w_\theta^f(c|s) \frac{f'''_\theta(a|s)}{f''_\theta(a|s)^2} + \mathbf{1}_c(a) w_\theta^f(a|s) w_\theta^f(b|s) \frac{f'''_\theta(a|s)}{f''_\theta(a|s)^2} \end{aligned}$$

$$\begin{aligned}
& - w_\theta^f(a|s) \cdot w_\theta^f(b|s) \cdot w_\theta^f(c|s) \cdot \left[\frac{f_\theta'''(a|s)}{f_\theta''(a|s)^2} + \frac{f_\theta'''(b|s)}{f_\theta''(b|s)^2} + \frac{f_\theta'''(c|s)}{f_\theta''(c|s)^2} \right] \\
& + w_\theta^f(a|s) \cdot w_\theta^f(b|s) \cdot w_\theta^f(c|s) \cdot \sum_{d \in \mathcal{A}} w_\theta^f(d|s) \cdot \frac{f_\theta'''(d|s)}{f_\theta''(d|s)^2} .
\end{aligned}$$

Lemma B.8. Assume that, for some $\underline{\pi}_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$ and $\mathbf{P}(\underline{\pi}_{\text{ref}})$, respectively. Then, it holds that

$$\sum_{(a,b) \in \mathcal{A}^2} \left| \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s,b)} \right| \leq 2 W_\theta^f(s) , \quad \sum_{(a,b,c) \in \mathcal{A}^3} \left| \frac{\partial^2 \pi_\theta^f(a|s)}{\partial \theta(s,b) \partial \theta(s,c)} \right| \leq 8 \kappa_f W_\theta^f(s) .$$

Proof. Using the expression of the derivative of the policy provided in Corollary B.7, we have by the triangle inequality

$$\sum_{b \in \mathcal{A}} \left| \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s,b)} \right| = \sum_{b \in \mathcal{A}} W_\theta^f(s) w_\theta^f(a|s) \left| 1_b(a) - w_\theta^f(b|s) \right| = 2 W_\theta^f(s) w_\theta^f(a|s) (1 - w_\theta^f(a|s)) .$$

where we used that

$$\sum_{b \in \mathcal{A}} \left| 1_b(a) - w_\theta^f(b|s) \right| = 2(1 - w_\theta^f(a|s)) . \quad (41)$$

Hence, we have

$$\sum_{(a,b) \in \mathcal{A}^2} \left| \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s,b)} \right| \leq 2 W_\theta^f(s) .$$

Fix $a \in \mathcal{A}$. By using the expression of the second derivative of the policy provided in Corollary B.7 combined with the triangle inequality and (41), we get

$$\sum_{(b,c) \in \mathcal{A}^2} \left| \frac{\partial^2 \pi_\theta^f(a|s)}{\partial \theta(s,b) \partial \theta(s,c)} \right| \leq 4 W_\theta^f(s) \frac{f_\theta'''(a|s)}{f_\theta''(a|s)^2} w_\theta^f(a|s) + 4 W_\theta^f(s) w_\theta^f(a|s) \sum_{b \in \mathcal{A}} \frac{f_\theta'''(b|s)}{f_\theta''(b|s)^2} w_\theta^f(b|s) .$$

Next combining $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$ and (41), we obtain

$$\sum_{(b,c) \in \mathcal{A}^2} \left| \frac{\partial^2 \pi_\theta^f(a|s)}{\partial \theta(s,b) \partial \theta(s,c)} \right| \leq 8 W_\theta^f(s) w_\theta^f(a|s) \kappa_f .$$

Finally, summing over the actions concludes the proof. \square

B.4 SMOOTHNESS OF OBJECTIVE

Firstly, we recall that the regularized value satisfies the following fixed-point equation Geist et al. (2019):

$$v_\pi^f(s) = \sum_{a \in \mathcal{A}} \pi(a|s) r(s,a) - \lambda D^f(\pi(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) + \gamma \sum_{(a,s') \in \mathcal{A} \times \mathcal{S}} \pi(a|s) P(s'|s,a) v_\pi^f(s') . \quad (42)$$

In order to prove the smoothness of the objective, we will prove that all the second-order directional derivatives are bounded. Denote $\theta_\alpha = \theta + \alpha u$ where $\alpha \in \mathbb{R}$ and $u \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$. By Equation (42), for any $s \in \mathcal{S}$ it holds that

$$v_{\theta_\alpha}^f(s) = \mathbf{e}_s^\top M(\alpha) r_{\theta_\alpha}^f , \quad (43)$$

where $M(\alpha)$ is a matrix of $\mathbb{R}^{\mathcal{S} \times \mathcal{S}}$ defined by

$$M(\alpha) = (\text{Id} - \gamma P_{\theta_\alpha})^{-1} ,$$

and where $r_{\theta_\alpha}^f$, and P_{θ_α} are defined in Section 2. Taking the derivative of (43) with respect to α yields

$$\frac{\partial v_{\theta_\alpha}^f(s)}{\partial \alpha} = \gamma \mathbf{e}_s^\top M(\alpha) \frac{\partial P_{\theta_\alpha}}{\partial \alpha} M(\alpha) r_{\theta_\alpha}^f + \mathbf{e}_s^\top M(\alpha) \frac{\partial r_{\theta_\alpha}^f}{\partial \alpha} .$$

Taking the derivative of the preceding equation with respect to α gives

$$\begin{aligned} \frac{\partial^2 v_{\theta_\alpha}^f(s)}{\partial \alpha^2} &= 2\gamma^2 \mathbf{e}_s^\top M(\alpha) \frac{\partial P_{\theta_\alpha}}{\partial \alpha} M(\alpha) \frac{\partial P_{\theta_\alpha}}{\partial \alpha} M(\alpha) r_{\theta_\alpha}^f + \gamma \mathbf{e}_s^\top M(\alpha) \frac{\partial^2 P_{\theta_\alpha}}{\partial \alpha^2} M(\alpha) r_{\theta_\alpha}^f \\ &\quad + 2\gamma \mathbf{e}_s^\top M(\alpha) \frac{\partial P_{\theta_\alpha}}{\partial \alpha} M(\alpha) \frac{\partial r_{\theta_\alpha}^f}{\partial \alpha} + \mathbf{e}_s^\top M(\alpha) \frac{\partial^2 r_{\theta_\alpha}^f}{\partial \alpha^2} . \end{aligned} \quad (44)$$

In order to control the second-order directional derivative of the regularised value function, we establish first several properties of the quantities that appear in the preceding equality.

Lemma B.9. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$, respectively. We have*

$$\left\| \frac{\partial P_{\theta_\alpha}}{\partial \alpha} \Big|_{\alpha=0} \right\|_\infty \leq 2 \max_s \{W_\theta^f(s)\} \|u\|_2 ,$$

Similarly, we have

$$\left\| \frac{\partial^2 P_{\theta_\alpha}}{\partial \alpha^2} \Big|_{\alpha=0} \right\|_\infty \leq 8\kappa_f \max_s \{W_\theta^f(s)\} \|u\|_2^2 .$$

Proof. **Bounding the first-order directional derivative.** The derivative with respect to α is

$$\left[\frac{\partial P_{\theta_\alpha}}{\partial \alpha} \Big|_{\alpha=0} \right]_{s,s'} = \sum_{a \in \mathcal{A}} \left[\frac{\partial \pi_{\theta_\alpha}^f(a|s)}{\partial \alpha} \Big|_{\alpha=0} \right] \mathbf{P}(s'|s, a) .$$

Fix $s \in \mathcal{S}$. Because $\pi_{\theta_\alpha}^f(a|s)$ depends only on $\theta(s, \cdot)$, by the chain rule

$$\sum_{a \in \mathcal{A}} \left| \frac{\partial \pi_{\theta_\alpha}^f(a|s)}{\partial \alpha} \Big|_{\alpha=0} \right| = \sum_{a \in \mathcal{A}} \left| \left\langle \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, \cdot)}, u(s, \cdot) \right\rangle \right| \leq \sum_{a \in \mathcal{A}} \left\| \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, \cdot)} \right\|_1 \|u(s, \cdot)\|_2 ,$$

where in the last inequality we used Cauchy-Schwarz inequality and the fact that the L_1 norm dominates the L_2 norm. Now using Lemma B.8, we get that

$$\sum_{a \in \mathcal{A}} \left\| \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, \cdot)} \right\|_1 \|u(s, \cdot)\|_2 \leq 2 \|u\|_2 W_\theta^f(s) .$$

Bounding the second-order directional derivative. Similarly, taking the second derivative with respect to α yields

$$\left[\frac{\partial^2 P_{\theta_\alpha}}{\partial \alpha^2} \Big|_{\alpha=0} \right]_{s,s'} = \sum_{a \in \mathcal{A}} \left[\frac{\partial^2 \pi_{\theta_\alpha}^f(a|s)}{\partial \alpha^2} \Big|_{\alpha=0} \right] \mathbf{P}(s'|s, a) .$$

Fix $s \in \mathcal{S}$. It holds that

$$\sum_{a \in \mathcal{A}} \left| \frac{\partial^2 \pi_{\theta_\alpha}^f(a|s)}{\partial \alpha^2} \Big|_{\alpha=0} \right| = \sum_{a \in \mathcal{A}} \left| \left\langle \frac{\partial^2 \pi_\theta^f(a|s)}{\partial^2 \theta(s, \cdot)} u(s, \cdot), u(s, \cdot) \right\rangle \right| \leq \sum_{a \in \mathcal{A}} \left\| \frac{\partial^2 \pi_\theta^f(a|s)}{\partial^2 \theta(s, \cdot)} \right\|_2 \|u(s, \cdot)\|_2^2 ,$$

where in the last inequality, we used the Cauchy-Schwarz inequality and the definition of the matrix operator norm. Additionally, using that for any matrix $A \in \mathbb{R}^{d \times d}$, we have

$$\|A\|_2 \leq \sum_{i=1}^d \sum_{j=1}^d |a_{i,j}|$$

combined with Lemma B.8, we get that

$$\sum_{a \in \mathcal{A}} \left| \frac{\partial^2 \pi_{\theta_\alpha}^f(a|s)}{\partial \alpha^2} \Big|_{\alpha=0} \right| \leq 8 \|u(s, \cdot)\|_2^2 \kappa_f \max_{s \in \mathcal{S}} W_\theta^f(s) ,$$

which concludes the proof. \square

Lemma B.10. Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$, respectively. Then, the regularized reward satisfies

$$\|r_\theta^f\|_\infty \leq 1 + \lambda \max_{s \in \mathcal{S}} D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) .$$

Additionally, we have that

$$\left\| \frac{\partial r_{\theta_\alpha}^f}{\partial \alpha} \Big|_{\alpha=0} \right\|_\infty \leq 2 \max_{s \in \mathcal{S}} \left\{ W_\theta^f(s) + \lambda Y_\theta^f(s) \right\} \|u\|_2 ,$$

and that

$$\left\| \frac{\partial^2 r_{\theta_\alpha}^f}{\partial \alpha^2} \Big|_{\alpha=0} \right\|_\infty \leq \max_{s \in \mathcal{S}} \left\{ 4(2\kappa_f + \lambda) W_\theta^f(s) + 8\lambda\kappa_f Y_\theta^f(s) \right\} \|u\|_2^2 .$$

Proof. The bound on $\|r_\theta^f\|_\infty$ is immediate.

Bounding the first derivative. It holds that

$$\left\| \frac{\partial r_{\theta_\alpha}^f}{\partial \alpha} \Big|_{\alpha=0} \right\|_\infty = \max_{s \in \mathcal{S}} \left| \left\langle \frac{\partial r_\theta^f(s)}{\partial \theta}, u \right\rangle \right| = \max_{s \in \mathcal{S}} \left| \left\langle \frac{\partial r_\theta^f(s)}{\partial \theta(s, \cdot)}, u(s, \cdot) \right\rangle \right| \leq \max_{s \in \mathcal{S}} \left\| \frac{\partial r_\theta^f(s)}{\partial \theta(s, \cdot)} \right\|_1 \|u\|_\infty .$$

Computing the derivative of $r_\theta^f(s)$ with respect to $\theta(s, b)$ yields

$$\frac{\partial r_\theta^f(s)}{\partial \theta(s, b)} = \sum_{a \in \mathcal{A}} \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, b)} r(s, a) - \lambda \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, b)} f' \left(\frac{\pi_\theta^f(a|s)}{\pi_{\text{ref}}(a|s)} \right) .$$

Plugging in the expression of the derivative of the policy of Corollary B.7 in the preceding identity yields

$$\begin{aligned} \frac{\partial \tilde{r}_\theta(s)}{\partial \theta(s, b)} &= \frac{\pi_{\text{ref}}(b|s)}{f_\theta''(b|s)} r(s, b) - \lambda \frac{\pi_{\text{ref}}(b|s)}{f_\theta''(b|s)} f_\theta'(b|s) - \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s) \pi_{\text{ref}}(b|s)}{f_\theta''(a|s) f_\theta''(b|s)} \cdot \frac{r(s, a)}{W_\theta^f(s)} \\ &\quad + \lambda \frac{1}{W_\theta^f(s)} \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s) \pi_{\text{ref}}(b|s)}{f_\theta''(a|s) f_\theta''(b|s)} f_\theta'(a|s) . \end{aligned} \quad (45)$$

Taking the absolute value, applying the triangle inequality, and using that the rewards are bounded by 1 gives

$$\sum_{b \in \mathcal{A}} \left| \frac{\partial r_\theta^f(s)}{\partial \theta(s, b)} \right| \leq 2 W_\theta^f(s) + 2\lambda Y_\theta^f(s) .$$

Bounding the second derivative. It holds that

$$\left\| \frac{\partial^2 r_{\theta_\alpha}^f}{\partial \alpha^2} \Big|_{\alpha=0} \right\|_\infty = \max_{s \in \mathcal{S}} \left\| \frac{\partial^2 r_{\theta_\alpha}^f(s)}{\partial \alpha^2} \Big|_{\alpha=0} \right\| \leq \max_{s \in \mathcal{S}} \left\| \frac{\partial^2 r_\theta^f(s)}{\partial \theta(s, \cdot)^2} \right\|_2 \|u\|_2^2 ,$$

where in the last inequality, we used the Cauchy-Schwarz inequality and the definition of the matrix operator norm.. We now compute the second derivative of $r_\theta^f(s)$. Starting from (45), we get

$$\begin{aligned} \frac{\partial^2 r_\theta^f(s)}{\partial \theta(s, b) \partial \theta(s, c)} &= -\frac{\pi_{\text{ref}}(b|s)}{f_\theta''(b|s)^2} \left[\mathbf{1}_c(b) \frac{f_\theta'''(b|s)}{f_\theta''(b|s)} - \frac{\pi_{\text{ref}}(c|s) f_\theta'''(b|s)}{f_\theta''(b|s) f_\theta''(c|s)} \cdot \frac{1}{W_\theta^f(s)} \right] r(s, b) \\ &\quad + \lambda \frac{\pi_{\text{ref}}(b|s) f_\theta'(b|s)}{f_\theta''(b|s)^2} \left[\mathbf{1}_c(b) \frac{f_\theta'''(b|s)}{f_\theta''(b|s)} - \frac{\pi_{\text{ref}}(c|s) f_\theta'''(b|s)}{f_\theta''(b|s) f_\theta''(c|s)} \cdot \frac{1}{W_\theta^f(s)} \right] \\ &\quad - \lambda \left[\mathbf{1}_b(c) \frac{\pi_{\text{ref}}(b|s)}{f_\theta''(b|s)} - \frac{\pi_{\text{ref}}(b|s) \pi_{\text{ref}}(c|s)}{f_\theta''(b|s) f_\theta''(c|s)} \cdot \frac{1}{W_\theta^f(s)} \right] \end{aligned}$$

$$\begin{aligned}
& + \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)\pi_{\text{ref}}(b|s)}{f''_{\theta}(a|s)^2 f''_{\theta}(b|s)} \cdot \frac{r(s, a)}{W_{\theta}^f(s)} \left[\mathbf{1}_c(a) \frac{f'''_{\theta}(a|s)}{f''_{\theta}(a|s)} - \frac{\pi_{\text{ref}}(c|s)f'''_{\theta}(a|s)}{f''_{\theta}(a|s)f''_{\theta}(c|s)} \cdot \frac{1}{W_{\theta}^f(s)} \right] \\
& + \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)\pi_{\text{ref}}(b|s)}{f''_{\theta}(a|s)f''_{\theta}(b|s)^2} \cdot \frac{r(s, a)}{W_{\theta}^f(s)} \left[\mathbf{1}_c(b) \frac{f'''_{\theta}(b|s)}{f''_{\theta}(b|s)} - \frac{\pi_{\text{ref}}(c|s)f'''_{\theta}(b|s)}{f''_{\theta}(b|s)f''_{\theta}(c|s)} \cdot \frac{1}{W_{\theta}^f(s)} \right] \\
& + \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)\pi_{\text{ref}}(b|s)}{f''_{\theta}(a|s)f''_{\theta}(b|s)} \frac{r(s, a)}{W_{\theta}^f(s)^2} \left[\frac{-\pi_{\text{ref}}(c|s)f'''_{\theta}(c|s)}{f''_{\theta}(c|s)^3} + \sum_{d \in \mathcal{A}} \frac{\pi_{\text{ref}}(c|s)\pi_{\text{ref}}(d|s)f'''_{\theta}(d|s)}{W_{\theta}^f(s)f''_{\theta}(d|s)^3 f''_{\theta}(c|s)} \right] \\
& - \lambda \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)\pi_{\text{ref}}(b|s)}{f''_{\theta}(a|s)f''_{\theta}(b|s)} \frac{f'_{\theta}(b|s)}{W_{\theta}^f(s)^2} \left[\frac{-\pi_{\text{ref}}(c|s)f'''_{\theta}(c|s)}{f''_{\theta}(c|s)^3} + \sum_{d \in \mathcal{A}} \frac{\pi_{\text{ref}}(c|s)\pi_{\text{ref}}(d|s)f'''_{\theta}(d|s)}{W_{\theta}^f(s)f''_{\theta}(d|s)^3 f''_{\theta}(c|s)} \right] \\
& - \lambda \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)\pi_{\text{ref}}(b|s)}{f''_{\theta}(a|s)^2 f''_{\theta}(b|s)} \frac{f'_{\theta}(b|s)}{W_{\theta}^f(s)} \left[\mathbf{1}_c(a) \frac{f'''_{\theta}(a|s)}{f''_{\theta}(a|s)} - \frac{\pi_{\text{ref}}(c|s)f'''_{\theta}(a|s)}{f''_{\theta}(a|s)f''_{\theta}(c|s)} \cdot \frac{1}{W_{\theta}^f(s)} \right] \cdot \\
& - \lambda \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)\pi_{\text{ref}}(b|s)}{f''_{\theta}(a|s)f''_{\theta}(b|s)^2} \frac{f'_{\theta}(b|s)}{W_{\theta}^f(s)} \left[\mathbf{1}_c(b) \frac{f'''_{\theta}(b|s)}{f''_{\theta}(b|s)} - \frac{\pi_{\text{ref}}(c|s)f'''_{\theta}(b|s)}{f''_{\theta}(b|s)f''_{\theta}(c|s)} \cdot \frac{1}{W_{\theta}^f(s)} \right] \cdot \\
& + \lambda \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''_{\theta}(a|s)} \frac{1}{W_{\theta}^f(s)} \left[\mathbf{1}_b(c) \frac{\pi_{\text{ref}}(b|s)}{f''_{\theta}(b|s)} - \frac{\pi_{\text{ref}}(b|s)\pi_{\text{ref}}(c|s)}{f''_{\theta}(b|s)f''_{\theta}(c|s)} \cdot \frac{1}{W_{\theta}^f(s)} \right] \cdot
\end{aligned}$$

Taking the absolute value, applying the triangle inequality, using that under $\mathbf{A}_f(\pi_{\text{ref}})$, for all $x \in \mathbb{R}^+$, we have $|f'''(x)/f''(x)^2| \leq \kappa_f$, and using that the rewards are bounded by 1 gives

$$\sum_{b \in \mathcal{A}} \sum_{c \in \mathcal{A}} \left| \frac{\partial^2 r_{\theta_0}^f(s)}{\partial \theta(s, b) \partial \theta(s, c)} \right| \leq 8\kappa_f W_{\theta}^f(s) + 8\lambda\kappa_f Y_{\theta}^f(s) + 4\lambda W_{\theta}^f(s) ,$$

which concludes the proof. \square

Lemma B.11. Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. Then, for any $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$ and $u \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$,

$$\left| u^{\top} \frac{\partial^2 v_{\theta}^f(s)}{\partial \theta^2} u \right| \leq \left(\sum_{i=1}^3 \frac{L_{\lambda, f}^{(i)}(\theta)}{(1-\gamma)^i} \right) \|u\|_2^2 ,$$

where for $i \in \{1, 2, 3\}$, $L_{\lambda, f}^{(i)}(\theta)$, are defined as

$$\begin{aligned}
L_{\lambda, f}^{(1)}(\theta) & := 4(2\kappa_f + \lambda) \left\| W_{\theta}^f \right\|_{\infty} + 8\lambda\kappa_f \left\| Y_{\theta}^f \right\|_{\infty} , \\
L_{\lambda, f}^{(2)}(\theta) & := 8\gamma \left\| W_{\theta}^f \right\|_{\infty} \left(\kappa_f \{1 + \lambda \max_{s \in \mathcal{S}} D^f(\pi_{\theta}^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s))\} + \left\| W_{\theta}^f \right\|_{\infty} + \lambda \left\| Y_{\theta}^f \right\|_{\infty} \right) , \\
L_{\lambda, f}^{(3)}(\theta) & := 8\gamma^2 \left\| W_{\theta}^f \right\|_{\infty}^2 \left(1 + \lambda \max_{s \in \mathcal{S}} D^f(\pi_{\theta}^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) \right) .
\end{aligned}$$

Proof. By construction, we get

$$\left| u^{\top} \frac{\partial^2 v_{\theta}^f(s)}{\partial \theta^2} u \right| = \left| \frac{\partial^2 v_{\theta_{\alpha}}^f(s)}{\partial \alpha^2} \Big|_{\alpha=0} \right| .$$

Using (44), we get that

$$\left| \frac{\partial^2 v_{\theta_{\alpha}}^f(s)}{\partial \alpha^2} \Big|_{\alpha=0} \right| \leq \underbrace{\left| 2\gamma^2 \mathbf{e}_s^{\top} M(0) \frac{\partial P_{\theta_{\alpha}}}{\partial \alpha} \Big|_{\alpha=0} M(0) \frac{\partial P_{\theta_{\alpha}}}{\partial \alpha} \Big|_{\alpha=0} M(0) r_{\theta}^f \right|}_{(\text{A})}$$

$$\begin{aligned}
& + \underbrace{\left| \gamma \mathbf{e}_s^\top M(0) \frac{\partial^2 \mathbf{P}_{\theta_\alpha}}{\partial^2 \alpha} \Big|_{\alpha=0} M(0) r_\theta^f \right|}_{\text{(B)}} + \underbrace{\left| 2\gamma \mathbf{e}_s^\top M(0) \frac{\partial \mathbf{P}_{\theta_\alpha}}{\partial \alpha} \Big|_{\alpha=0} M(0) \frac{\partial r_\theta^f}{\partial \alpha} \Big|_{\alpha=0} \right|}_{\text{(C)}} \\
& + \underbrace{\left| \mathbf{e}_s^\top M(0) \frac{\partial^2 r_\theta^f}{\partial^2 \alpha} \Big|_{\alpha=0} \right|}_{\text{(D)}} .
\end{aligned}$$

We now bound each of these terms separately

Bounding (A). First note that, for any vector $x \in \mathbb{R}^S$ and $\alpha \in \mathbb{R}$, we have

$$\|M(\alpha)x\|_\infty \leq \frac{1}{1-\gamma} \|x\|_\infty , \quad (46)$$

This yields

$$\text{(A)} \leq 2\gamma^2 \left\| M(0) \frac{\partial \mathbf{P}_{\theta_\alpha}}{\partial \alpha} \Big|_{\alpha=0} M(0) \frac{\partial \mathbf{P}_{\theta_\alpha}}{\partial \alpha} \Big|_{\alpha=0} M(0) r_\theta^f \right\|_\infty \leq \frac{2\gamma^2}{(1-\gamma)^3} \left\| \frac{\partial \mathbf{P}_{\theta_\alpha}}{\partial \alpha} \Big|_{\alpha=0} \right\|_\infty^2 \|r_\theta^f\|_\infty$$

By using again (46), Lemma B.9, and Lemma B.10 we get

$$\text{(A)} \leq \frac{8\gamma^2 \max_s \{W_\theta^f(s)\}^2 \|u\|_2^2}{(1-\gamma)^3} (1 + \lambda \max_{s \in \mathcal{S}} D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s))) .$$

Bounding (B). Using (46), Lemma B.9, and Lemma B.10 we get

$$\begin{aligned}
\text{(B)} & \leq \frac{\gamma}{(1-\gamma)^2} \left\| \frac{\partial^2 \mathbf{P}_{\theta_\alpha}}{\partial^2 \alpha} \Big|_{\alpha=0} \right\|_\infty \|r_\theta^f\|_\infty \\
& \leq \frac{8\gamma}{(1-\gamma)^2} \kappa_f \max_s \{W_\theta^f(s)\} \|u\|_2^2 \{1 + \lambda \max_{s \in \mathcal{S}} D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s))\} .
\end{aligned}$$

Bounding (C). Similarly, using (46), Lemma B.9, and Lemma B.10 we get

$$\text{(C)} \leq \frac{8\gamma}{(1-\gamma)^2} \max_{s \in \mathcal{S}} \{W_\theta^f(s)\} \max_{s \in \mathcal{S}} \{W_\theta^f(s) + \lambda Y_\theta^f(s)\} \|u\|_2^2 .$$

Bounding (D). Using (46), Lemma B.9, and Lemma B.10 we get

$$\text{(D)} \leq \frac{1}{1-\gamma} \left\| \frac{\partial^2 r_\theta^f}{\partial^2 \alpha} \Big|_{\alpha=0} \right\|_\infty \leq \frac{1}{1-\gamma} \max_{s \in \mathcal{S}} \left\{ 4(2\kappa_f + \lambda) W_\theta^f(s) + 8\lambda \kappa_f Y_\theta^f(s) \right\} \|u\|_2^2 .$$

The proof is completed by collecting these upper bounds. \square

Theorem B.12. Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. Then for any $\theta, \theta' \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$, it holds that

$$\left| v_{\theta'}^f(\rho) - v_\theta^f(\rho) - \left\langle \frac{\partial v_\theta^f(\rho)}{\partial \theta}, \theta' - \theta \right\rangle \right| \leq \frac{L_f}{2} \|\theta' - \theta\|_2^2 .$$

where

$$L_f := \frac{8\omega_f (\gamma\omega_f + (1-\gamma)\kappa_f)}{(1-\gamma)^3} + 4\lambda \frac{2\gamma^2 \omega_f^2 d_f + 2\gamma(1-\gamma)\omega_f [\kappa_f d_f + y_f] + (1-\gamma)^2 [\omega_f + 2\kappa_f y_f]}{(1-\gamma)^3} .$$

Proof. Fix any vector $u \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$ and $\theta \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$. Using Lemma B.11, it holds that

$$\left| u^\top \frac{\partial^2 v_\theta^f(s)}{\partial \theta^2} u \right| \leq \left(\sum_{i=1}^3 \frac{L_{\lambda, f}^{(i)}(\theta)}{(1-\gamma)^i} \right) \|u\|_2^2 ,$$

where for $i \in \{1, 2, 3\}$, $L_{\lambda,f}^{(i)}(\theta)$, are defined as

$$\begin{aligned} L_{\lambda,f}^{(1)}(\theta) &:= 4(2\kappa_f + \lambda) \left\| W_\theta^f \right\|_\infty + 8\lambda\kappa_f \left\| Y_\theta^f \right\|_\infty, \\ L_{\lambda,f}^{(2)}(\theta) &:= 8\gamma \left\| W_\theta^f \right\|_\infty \left(\kappa_f \{1 + \lambda \max_{s \in \mathcal{S}} D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s))\} + \left\| W_\theta^f \right\|_\infty + \lambda \left\| Y_\theta^f \right\|_\infty \right), \\ L_{\lambda,f}^{(3)}(\theta) &:= 8\gamma^2 \left\| W_\theta^f \right\|_\infty^2 (1 + \lambda \max_{s \in \mathcal{S}} D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s))). \end{aligned}$$

Using Lemma B.6 combined with Lemma I.1 concludes the proof. \square

C NON-UNIFORM ŁOJASIEWICZ INEQUALITY

Firstly, define respectively q_θ^f and d_ρ^θ as the regularized Q-function and discounted state visitation associated with the policy π_θ^f , i.e.

$$q_\theta^f(s, a) = r(s, a) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) v_\theta^f(s'), \quad (47)$$

$$d_\rho^\theta(s) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \rho P_{\pi_\theta^f}^t(s). \quad (48)$$

The goal of this section is to prove that the global objective satisfies a non-uniform Łojasiewicz inequality, i.e we aim to show the following theorem

Theorem C.1. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. Assume in addition that the initial distribution ρ satisfy \mathbf{A}_ρ . Then, it holds that*

$$\left\| \frac{\partial v_\theta^f(\rho)}{\partial \theta} \right\|_2^2 \geq \mu_f(\theta) \left(v_\star^f(\rho) - v_\theta^f(\rho) \right),$$

where

$$\mu_f(\theta) := \frac{\lambda(1 - \gamma)\rho_{\min}^2 \zeta_f^2}{\omega_f^2} \min_{(s,a) \in \mathcal{S} \times \mathcal{A}} w_\theta^f(a|s)^2.$$

One of the main challenges in establishing such an inequality lies in connecting global information (the suboptimality gap) to local information (the gradient norm). Recall from Section 2 that if

$$\theta(s, a) = q_\star^f(s, a)/\lambda, \quad \forall a \in \mathcal{A},$$

then $\pi_\theta^f = \pi_\star^f$. This observation highlights that, under this parameterization and regularization, the key quantity is the closeness between θ and q_θ^f/λ . Formally, we will show that both the suboptimality gap and the gradient norm can be upper and lower bounded, respectively, by a quantity proportional to $\|\zeta_\theta^f(s)\|_2$, where we define

$$\zeta_\theta^f(s) := q_\theta^f(s, \cdot)/\lambda - \theta(s, \cdot) - K_\theta^f(s) \mathbf{1}_{|\mathcal{A}|}, \quad (49)$$

$$K_\theta^f(s) := \frac{\langle q_\theta^f(s, \cdot)/\lambda - \theta(s, \cdot), \mathbf{1}_{|\mathcal{A}|} \rangle}{|\mathcal{A}|}. \quad (50)$$

Note that ζ_θ^f is the projection of $q_\theta^f(s, \cdot)/\lambda - \theta(s, \cdot)$ onto the subspace orthogonal to $\mathbf{1}_{|\mathcal{A}|}$.

The proof proceeds in three steps:

1. Derive an explicit expression for the gradient of the objective and establish a lower bound in terms of $\|\zeta_\theta^f\|$.
2. Upper bound the suboptimality gap by a quantity directly related to $\|\zeta_\theta^f\|$.
3. Combine these two bounds to identify the corresponding non-uniform PL coefficient.

We now detail each step in turn.

C.1 LOWER BOUNDING THE NORM OF THE GRADIENT

Before deriving a lower bound on the norm of the gradient, we start by deriving an expression for the latter.

Lemma C.2. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. For any $s \in \mathcal{S}$ and $b \in \mathcal{A}$, we have*

$$\frac{1}{W_\theta^f(s)} \frac{\partial v_\theta^f(\rho)}{\partial \theta(s, b)} = \frac{d_\rho^\theta(s)}{1-\gamma} w_\theta^f(b|s) \left[q_\theta^f(s, b) - \lambda \theta(s, b) - \sum_{a \in \mathcal{A}} w_\theta^f(a|s) \left[q_\theta^f(s, a) - \lambda \theta(s, a) \right] \right] .$$

Proof. Fix $s \in \mathcal{S}$ and $b \in \mathcal{A}$. Additionally fix any $\tilde{s} \in \mathcal{S}$. Using (42), we have

$$v_\theta^f(\tilde{s}) = \sum_{a \in \mathcal{A}} \pi_\theta^f(a|\tilde{s}) r(\tilde{s}, a) - \lambda D^f(\pi_\theta^f(\cdot|\tilde{s}) || \pi_{\text{ref}}(\cdot|\tilde{s})) + \gamma \sum_{a \in \mathcal{A}} \sum_{\tilde{s}' \in \mathcal{S}} \pi_\theta^f(a|\tilde{s}) \mathbf{P}(\tilde{s}'|\tilde{s}, a) v_\theta^f(\tilde{s}')$$

Deriving the preceding recursion with respect to $\theta(s, b)$ yields

$$\begin{aligned} \frac{\partial v_\theta^f(\tilde{s})}{\partial \theta(s, b)} &= \sum_{a \in \mathcal{A}} \frac{\partial \pi_\theta^f(a|\tilde{s})}{\partial \theta(s, b)} \left[\underbrace{r(\tilde{s}, a) - \lambda f' \left(\frac{\pi_\theta^f(a|\tilde{s})}{\pi_{\text{ref}}(a|\tilde{s})} \right) + \gamma \sum_{\tilde{s}' \in \mathcal{S}} \mathbf{P}(\tilde{s}'|\tilde{s}, a) v_\theta^f(\tilde{s}')}_{Z(\tilde{s})} \right] \\ &\quad + \gamma \sum_{a \in \mathcal{A}} \sum_{\tilde{s}' \in \mathcal{S}} \pi_\theta^f(a|\tilde{s}) \mathbf{P}(\tilde{s}'|\tilde{s}, a) \frac{\partial v_\theta^f(\tilde{s}')}{\partial \theta(s, b)} . \end{aligned}$$

Using the definition of the regularized Q-function and writing the preceding recursion in a vector form yields

$$\frac{\partial v_\theta^f(\cdot)}{\partial \theta(s, b)} = Z(\cdot) + \gamma \mathbf{P}_\theta \frac{\partial v_\theta^f(\cdot)}{\partial \theta(s, b)} .$$

which implies

$$\rho^\top \frac{\partial v_\theta^f(\cdot)}{\partial \theta(s, b)} = \frac{\partial v_\theta^f(\rho)}{\partial \theta(s, b)} = \rho^\top (\text{Id} - \gamma \mathbf{P}_\theta)^{-1} Z(\cdot) .$$

Next, using the definition of the discounted state visitation (48) and the regularized Q-function (47) implies

$$\frac{\partial v_\theta^f(\rho)}{\partial \theta(s, b)} = \frac{1}{1-\gamma} \sum_{s' \in \mathcal{S}} d_\rho^\theta(s') \sum_{a \in \mathcal{A}} \frac{\partial \pi_\theta^f(a|s')}{\partial \theta(s, b)} \left[q_\theta^f(s', a) - \lambda f' \left(\frac{\pi_\theta^f(a|s')}{\pi_{\text{ref}}(a|s')} \right) \right] . \quad (51)$$

Using that $\sum_{a \in \mathcal{A}} \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, b)} = 0$ and that for $s \neq s'$, we have $\frac{\partial \pi_\theta^f(a|s')}{\partial \theta(s, b)} = 0$ yields

$$\frac{\partial v_\theta^f(\rho)}{\partial \theta(s, b)} = \frac{1}{1-\gamma} d_\rho^\theta(s) \sum_{a \in \mathcal{A}} \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, b)} \left[q_\theta^f(s, a) - \lambda f' \left(\frac{\pi_\theta^f(a|s)}{\pi_{\text{ref}}(a|s)} \right) - \lambda \mu_\theta(s) \right] .$$

where $\mu_\theta(s)$ is defined in Corollary B.7 and satisfies for any $a \in \mathcal{A}$

$$\theta(s, a) - f' \left(\frac{\pi_\theta^f(a|s)}{\pi_{\text{ref}}(a|s)} \right) = \mu_\theta(s) .$$

Thus, we obtain

$$\frac{\partial v_\theta^f(\rho)}{\partial \theta(s, b)} = \frac{1}{1-\gamma} d_\rho^\theta(s) \sum_{a \in \mathcal{A}} \frac{\partial \pi_\theta^f(a|s)}{\partial \theta(s, b)} \left[q_\theta^f(s, a) - \lambda \theta(s, a) \right] ,$$

Finally, plugging in the expression of the derivative of the policy derived in Corollary B.7 in the previous equality concludes the proof. \square

Using the previous lemma, we prove the following lower bound for the norm of the gradient.

Lemma C.3. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. Assume in addition that the initial distribution ρ satisfy \mathbf{A}_ρ . We have*

$$\left\| \frac{\partial v_\theta^f(\rho)}{\partial \theta} \right\|_2^2 \geq \lambda^2 \rho_{\min}^2 \min_{(s,a) \in \mathcal{S} \times \mathcal{A}} \{w_\theta^f(a|s)^2\} \min_{s \in \mathcal{S}} \{W_\theta^f(s)^2\} \sum_{s \in \mathcal{S}} \|\zeta_\theta(s)\|_2^2 .$$

Proof. It holds that

$$\left\| \frac{\partial v_\theta^f(\rho)}{\partial \theta} \right\|_2^2 = \sum_{s \in \mathcal{S}} \left\| \frac{\partial v_\theta^f(\rho)}{\partial \theta(s, \cdot)} \right\|_2^2 .$$

Fix $s \in \mathcal{S}$. Using Lemma C.2, we observe that

$$\frac{1}{W_\theta^f(s)} \frac{\partial v_\theta^f(\rho)}{\partial \theta(s, \cdot)} = \frac{d_\rho^\theta(s)}{1 - \gamma} H(w_\theta^f(\cdot|s)) \left[q_\theta^f(s, \cdot) - \lambda \theta(s, \cdot) \right] .$$

where for any vector $u \in \mathbb{R}^{|\mathcal{A}|}$, we define $H(u) := \text{diag}(u) - uu^\top$. Thus, we get that

$$\begin{aligned} \frac{1 - \gamma}{W_\theta^f(s)} \left\| \frac{\partial v_\theta^f(\rho)}{\partial \theta(s, \cdot)} \right\|_2 &= d_\rho^\theta(s) \left\| H(w_\theta^f(\cdot|s)) \left[q_\theta^f(s, \cdot) - \lambda \theta(s, \cdot) \right] \right\|_2 \\ &= \lambda d_\rho^\theta(s) \left\| H(w_\theta^f(\cdot|s)) \left[q_\theta^f(s, \cdot)/\lambda - \theta(s, \cdot) - K_\theta^f(s) \mathbf{1}_{|\mathcal{A}|} \right] \right\|_2 \quad (\text{using } H(u) \mathbf{1}_{|\mathcal{A}|} = 0 \text{ and (50)}) \\ &\geq \lambda d_\rho^\theta(s) \min_{a \in \mathcal{A}} w_\theta^f(a|s) \|\zeta_\theta(s)\|_2 \quad (\text{where } \zeta_\theta(s) \text{ is defined in (49) and using Lemma I.4}) . \end{aligned}$$

Finally, using $d_\rho^\theta(s) \geq (1 - \gamma)\rho(s)$ and \mathbf{A}_ρ concludes the proof. \square

C.2 BOUNDING THE SUBOPTIMALITY GAP

The first step is to connect the suboptimality gap to information localized at θ . This is achieved via the performance difference lemma for the regularized value function yields (see Lemma I.3)

$$v_\star^f(\rho) - v_\theta^f(\rho) = \sum_{s \in \mathcal{S}} \frac{\lambda d_\rho^{\pi_\star^f}(s)}{1 - \gamma} \underbrace{\left[\sum_{a \in \mathcal{A}} \pi_\star^f(a|s) \frac{q_\theta^f(s, a)}{\lambda} - D^f(\pi_\star^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) - \frac{v_\theta^f(s)}{\lambda} \right]}_{(\mathbf{A}(s))} . \quad (52)$$

Fix $s \in \mathcal{S}$. Using the definition of the regularized value functions and Q-functions combined with Equation (42), we have

$$v_\theta^f(s) = \langle \pi_\theta^f(\cdot|s), q_\theta^f(s, \cdot) \rangle - \lambda D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) .$$

This implies $\mathbf{A}(s) = \mathbf{A}_1(s) - \mathbf{A}_2(s) - \mathbf{A}_3(s)$ where

$$\begin{aligned} \mathbf{A}_1(s) &= \langle \pi_\star^f(\cdot|s), q_\theta^f(s, \cdot)/\lambda \rangle - D^f(\pi_\star^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) \\ \mathbf{A}_2(s) &= \langle \pi_\theta^f(\cdot|s), \theta(s, \cdot) \rangle - D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) \\ \mathbf{A}_3(s) &= \langle \pi_\theta^f(\cdot|s), q_\theta^f(s, \cdot)/\lambda - \theta(s, \cdot) \rangle . \end{aligned} \quad (53)$$

Using (4) and (15), combined with $\mathbf{A}_1(s) \leq \text{f-softmax}(q_\theta^f(s, \cdot)/\lambda, \pi_{\text{ref}}(\cdot|s))$ and the fact that $\mathbf{A}_2(s) = \text{f-softmax}(\theta(s, \cdot), \pi_{\text{ref}}(\cdot|s))$, gives

$$\begin{aligned} \mathbf{A}(s) &\leq \text{f-softmax}(q_\theta^f(s, \cdot)/\lambda, \pi_{\text{ref}}(\cdot|s)) - \text{f-softmax}(\theta(s, \cdot), \pi_{\text{ref}}(\cdot|s)) - \mathbf{A}_3(s) \\ &= \text{f-softmax}(q_\theta^f(s, \cdot)/\lambda, \pi_{\text{ref}}(\cdot|s)) - \text{f-softmax}(\theta(s, \cdot) + K_\theta^f(s) \mathbf{1}_{|\mathcal{A}|}) \\ &\quad - \langle \pi_\theta^f(\cdot|s), q_\theta^f(s, \cdot)/\lambda - \theta(s, \cdot) - K_\theta^f(s) \mathbf{1}_{|\mathcal{A}|} \rangle , \end{aligned} \quad (54)$$

where in the last equality we used that, for any $x \in \mathbb{R}^{|\mathcal{A}|}$ and $\alpha \in \mathbb{R}$,

$$\text{f-softmax}(x + \alpha \mathbf{1}_{|\mathcal{A}|}, \pi_{\text{ref}}(\cdot|s)) = \text{f-softmax}(x, \pi_{\text{ref}}(\cdot|s)) + \alpha .$$

The structure of $(\mathbf{A}(s))$ closely resembles that of a first-order Taylor expansion as by Lemma B.5, we have that

$$\frac{\partial \text{f-softmax}(\theta(s, \cdot), \pi_{\text{ref}}(\cdot|s))}{\partial \theta(s, \cdot)} = \text{f-softargmax}(\theta(s, \cdot), \pi_{\text{ref}}(\cdot|s)) = \pi_{\theta}^f(\cdot|s) . \quad (55)$$

Lemma C.4. Assume $\mathbf{A}_f(\pi_{\text{ref}})$. It holds that

$$v_{\star}^f(\rho) - v_{\theta}^f(\rho) \leq \frac{\lambda}{1-\gamma} \sup_{\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}} \{\|W_{\theta}\|_{\infty}\} \sum_{s \in \mathcal{S}} \|\zeta_{\theta}^f(s)\|_2^2$$

Proof. In this proof, we denote by $g_s(x) = \text{f-softmax}(x, \pi_{\text{ref}}(\cdot|s))$. Combining (52) and (54) yields

$$v_{\star}^f(\rho) - v_{\theta}^f(\rho) \leq \frac{\lambda}{1-\gamma} \sum_{s \in \mathcal{S}} d_{\rho_{\star}^f}^{\pi_{\star}^f}(s) B(s) , \quad (56)$$

where we have defined

$$B(s) = g_s(q_{\theta}^f(s, \cdot)/\lambda) - g_s(\theta(s, \cdot) + K_{\theta}^f(s)\mathbf{1}_{|\mathcal{A}|}) - \langle \pi_{\theta}^f(\cdot|s), q_{\theta}^f(s, \cdot)/\lambda - \theta(s, \cdot) - K_{\theta}^f(s)\mathbf{1}_{|\mathcal{A}|} \rangle .$$

Next, by Lemma B.5, it holds that

$$\left. \frac{\partial g_s(x)}{\partial x} \right|_{x=\theta(s, \cdot) + K_{\theta}^f(s)\mathbf{1}_{|\mathcal{A}|}} = \pi_{\theta}^f(\cdot|s) .$$

Standard one-dimensional Taylor theorem with Lagrange remainder shows that there exists $y \in \mathbb{R}^{|\mathcal{A}|}$ which belongs to the segment between $\theta(s, \cdot) + K_{\theta}^f(s)\mathbf{1}_{\mathcal{A}}$ and $q_{\theta}^f(s, \cdot)/\lambda$ such that

$$B(s) = \frac{1}{2} \left\langle \left. \frac{\partial^2 g_s(x)}{\partial x^2} \right|_{x=y}, \zeta_{\theta}^f(s), \zeta_{\theta}^f(s) \right\rangle$$

Using the bound on the spectral norm of the Hessian of g_s derived in Lemma B.5, we obtain

$$B(s) = \frac{1}{2} \left\langle \left. \frac{\partial^2 g_s(x)}{\partial x^2} \right|_{x=y}, \zeta_{\theta}^f, \zeta_{\theta}^f \right\rangle \leq \frac{1}{2} \left\| \left. \frac{\partial^2 g_s(x)}{\partial x^2} \right|_{x=y} \right\|_2 \left\| \zeta_{\theta}^f(s) \right\|_2^2 \leq W_y^f(s) \left\| \zeta_{\theta}^f(s) \right\|_2^2 .$$

Finally, bounding the discounted state visitation measure in (56) by 1 and plugging in the preceding bound on $B(s)$ concludes the proof. \square

The proof of Theorem C.1 follows immediately from Lemma C.3, Lemma C.4, Lemma B.6, and (7).

D MONOTONE IMPROVEMENT OPERATORS

A key challenge in analyzing stochastic policy gradient methods is that the Łojasiewicz inequality depends on θ and degenerates whenever the probability of an action becomes small. The goal of this section is therefore to show the existence of an operator IMP^f with two crucial properties: (i) for any policy, applying this operator produces a new policy with higher objective value, and (ii) every policy generated by this operator assigns at least a fixed minimum probability to every action. The main idea is to build the improvement operator such that it slightly augments the smallest probability weights, such that for any state action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ the probability ratio $\pi(a|s)/\pi_{\text{ref}}(a|s)$ stays above a certain threshold. We will show below that this procedure improves the global objective while keeping the probabilities uniformly bounded away from 0 when the threshold is properly chosen. Let $\underline{\pi}_{\text{ref}} > 0$ be such that $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$ and $\mathbf{P}(\underline{\pi}_{\text{ref}})$ hold. For any policy π , state $s \in \mathcal{S}$, $\tau < \underline{\pi}_{\text{ref}}/2$, we respectively define $\mathcal{A}_{\tau}^{\pi}(s)$, and $a_{\text{max}}^{\pi}(s)$ as

$$\mathcal{A}_{\tau}^{\pi}(s) := \{a \in \mathcal{A}, \pi(a|s)/\pi_{\text{ref}}(a|s) \leq \tau/2\} , \quad a_{\text{max}}^{\pi}(s) = \arg \max_{a \in \mathcal{A}} \{\pi(a|s)/\pi_{\text{ref}}(a|s)\} ,$$

where the $\arg \max$ is chosen at random in the case of ties. Note that the definition of τ ensures that $a_{\max}^\pi(s)$ does not belong to the set $\mathcal{A}_\tau^\pi(s)$ as

$$\max_{a \in \mathcal{A}} \frac{\pi(a|s)}{\pi_{\text{ref}}(a|s)} \geq 1 .$$

Finally, we define the improvement operator as follows:

$$\begin{aligned} \mathcal{U}_\tau : \mathcal{P}(\mathcal{A})^{\mathcal{S}} &\longrightarrow \mathcal{P}(\mathcal{A})^{\mathcal{S}}, \\ \pi &\longmapsto \mathcal{U}_\tau(\pi), \end{aligned} \quad (57)$$

where for every $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\mathcal{U}_\tau(\pi)(a|s) = \begin{cases} \pi_{\text{ref}}(a|s)\tau, & \text{if } \pi(a|s) \leq \pi_{\text{ref}}(a|s)\tau/2, \\ \pi(a|s) - \sum_{b \in \mathcal{A}_\tau^\pi(s)} (\pi_{\text{ref}}(b|s)\tau - \pi(b|s)), & \text{if } a = a_{\max}^\pi(s), \\ \pi(a|s), & \text{otherwise.} \end{cases}$$

The operator \mathcal{U}_τ builds $\mathcal{U}_\tau(\pi)(a|s)$ by (statewise) raising each $a \in \mathcal{A}_\tau^\pi(s)$ to $\pi_{\text{ref}}(a|s)\tau$, subtracting the total added mass from the single action $a_{\max}^\pi(s)$, and leaving other actions unchanged. If $\mathcal{A}_\tau^\pi(s) = \emptyset$, for all $s \in \mathcal{S}$, then $\text{IMP}^f(\pi) = \pi$. Note that mass conservation is immediate from the definition and the fact that $\tau < \underline{\pi_{\text{ref}}}/2$. Non-negativity of $\mathcal{U}_\tau(\pi)(a_{\max}^\pi(s)|s)$ follows because the removed mass is

$$\sum_{a \in \mathcal{A}_\tau^\pi(s)} \{\pi_{\text{ref}}(a|s)\tau - \pi(a|s)\} \leq \tau \sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) \leq \tau$$

Since $\pi(a_{\max}^\pi(s)|s) \geq \pi_{\text{ref}}(a_{\max}^\pi(s)|s) \geq \underline{\pi_{\text{ref}}}$, and $\tau \leq \underline{\pi_{\text{ref}}}/2$, we get that $\mathcal{U}_\tau(\pi)(a_{\max}^\pi(s)|s) \geq \tau/2$. This in particular shows that $\mathcal{U}_\tau(\pi)$ is a policy. As by **A_f($\underline{\pi_{\text{ref}}}$)** we have $\lim_{x \rightarrow 0^+} f'(x) = -\infty$, we consider $[f']^{-1} : (-\infty, f'(1/\underline{\pi_{\text{ref}}})] \mapsto \mathbb{R}_+$ the inverse of f' and define

$$\tau_\lambda := \min \left([f']^{-1} \left(-\frac{16 + 8\gamma\lambda d_f}{\lambda(1-\gamma)^2 \rho_{\min}} \right), [f']^{-1} \left(-4 \left| f' \left(\frac{1}{2} \right) \right| \right), \frac{1}{2} \underline{\pi_{\text{ref}}} \right) . \quad (58)$$

The following lemma establishes the crucial improvement property when $\tau = \tau_\lambda$.

Lemma D.1. *Assume that, for some $\underline{\pi_{\text{ref}}} > 0$, f and $\underline{\pi_{\text{ref}}}$ satisfy **A_f($\underline{\pi_{\text{ref}}}$)** and **P($\underline{\pi_{\text{ref}}}$)** respectively. Assume in addition that the initial distribution ρ satisfies **A _{ρ}** . For any policy π , it holds that*

$$v_{\mathcal{U}_{\tau_\lambda}(\pi)}^f(\rho) \geq v_\pi^f(\rho) .$$

Additionally, for any policy π , we have that

$$\mathcal{U}_{\tau_\lambda}(\pi)(a|s) \geq \underline{\pi_{\text{ref}}}\tau_\lambda .$$

Proof. Set an arbitrary policy π . For avoiding heavy notations, we will, through this proof, denote by $A_\tau^\pi = A_{\tau_\lambda}^\pi$. We consider the case where there is $s \in \mathcal{S}$ such that $\mathcal{A}_\tau^\pi(s) \neq \emptyset$ (alternatively $\mathcal{U}_{\tau_\lambda}(\pi) = \pi$, which makes the previous inequality immediately valid). Define $\tilde{\pi} = \mathcal{U}_{\tau_\lambda}(\pi)$. The following applies

$$\begin{aligned} v_{\tilde{\pi}}^f(\rho) - v_\pi^f(\rho) &= \sum_{s \in \mathcal{S}} d_\rho^{\tilde{\pi}}(s) \sum_{a \in \mathcal{A}} \left[\tilde{\pi}(a|s)r(s, a) - \lambda \pi_{\text{ref}}(a|s) f \left(\frac{\tilde{\pi}(a|s)}{\pi_{\text{ref}}(a|s)} \right) \right] \\ &\quad - \sum_{s \in \mathcal{S}} d_\rho^\pi(s) \sum_{a \in \mathcal{A}} \left[\pi(a|s)r(s, a) - \lambda \pi_{\text{ref}}(a|s) f \left(\frac{\pi(a|s)}{\pi_{\text{ref}}(a|s)} \right) \right] \\ &= \underbrace{\sum_{s \in \mathcal{S}} (d_\rho^{\tilde{\pi}}(s) - d_\rho^\pi(s)) \sum_{a \in \mathcal{A}} \left[\tilde{\pi}(a|s)r(s, a) - \lambda \pi_{\text{ref}}(a|s) f \left(\frac{\tilde{\pi}(a|s)}{\pi_{\text{ref}}(a|s)} \right) \right]}_{\text{(I)}} \\ &\quad + \underbrace{\sum_{s \in \mathcal{S}} d_\rho^\pi(s) \sum_{a \in \mathcal{A}} (\tilde{\pi}(a|s) - \pi(a|s))r(s, a)}_{\text{(II)}} \end{aligned}$$

$$+ \lambda \underbrace{\sum_{s \in \mathcal{S}} d_\rho^\pi(s) \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s) \left[f\left(\frac{\pi(a|s)}{\pi_{\text{ref}}(a|s)}\right) - f\left(\frac{\tilde{\pi}(a|s)}{\pi_{\text{ref}}(a|s)}\right) \right]}_{\text{(III)}} .$$

We now lower-bound each of the three terms separately.

Bounding (I). Using Lemma I.2, we have

$$\begin{aligned} \text{(I)} &\geq - \|d_\rho^\pi - d_\rho^\pi\|_1 \max_{s \in \mathcal{S}} \left| \sum_{a \in \mathcal{A}} \left[\tilde{\pi}(a|s)r(s, a) - \lambda \pi_{\text{ref}}(a|s) f\left(\frac{\tilde{\pi}(a|s)}{\pi_{\text{ref}}(a|s)}\right) \right] \right| \\ &\geq - \frac{\gamma}{1-\gamma} \sup_{s \in \mathcal{S}} \|\tilde{\pi}(\cdot|s) - \pi(\cdot|s)\|_1 \sup_{s \in \mathcal{S}} \left[1 + \lambda D^f(\tilde{\pi}(\cdot|s) \|\pi_{\text{ref}}(\cdot|s)) \right] \\ &\geq - \frac{2\gamma}{1-\gamma} \tau_\lambda \max_{s \in \mathcal{S}} \left\{ \sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) \right\} \sup_{\nu \in \mathcal{P}(\mathcal{A})} \sup_{s \in \mathcal{S}} \left[1 + \lambda D^f(\nu \|\pi_{\text{ref}}(\cdot|s)) \right] , \end{aligned}$$

where in the last inequality we used that (because we increase the probability of the actions in $\mathcal{A}_\tau^\pi(s)$ by τ_λ and remove the total added mass from the probability of $\pi(a_{\text{max}}^\pi(s))$)

$$\sup_{s \in \mathcal{S}} \|\tilde{\pi}(\cdot|s) - \pi(\cdot|s)\|_1 \leq 2 \max_{s \in \mathcal{S}} \left\{ \sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) \right\} \tau_\lambda .$$

Bounding (II). Using the triangle inequality yields

$$\text{(II)} \geq - \sup_{s \in \mathcal{S}} \|\tilde{\pi}(\cdot|s) - \pi(\cdot|s)\|_1 \geq -2 \max_{s \in \mathcal{S}} \left\{ \sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) \right\} \tau_\lambda .$$

Bounding (III). All the state-action pairs on which the original π allocates the same probability then the policy $\tilde{\pi}$ are equal to 0 in (III) allowing us to simplify this term

$$\begin{aligned} \text{(III)} &= \lambda \sum_{s \in \mathcal{S}} d_\rho^\pi(s) \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s) \left[f\left(\frac{\pi(a|s)}{\pi_{\text{ref}}(a|s)}\right) - f\left(\frac{\tilde{\pi}(a|s)}{\pi_{\text{ref}}(a|s)}\right) \right] \\ &= \lambda \sum_{s \in \mathcal{S}} d_\rho^\pi(s) \sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) \left[f\left(\frac{\pi(a|s)}{\pi_{\text{ref}}(a|s)}\right) - f\left(\frac{\tilde{\pi}(a|s)}{\pi_{\text{ref}}(a|s)}\right) \right] \\ &\quad + \lambda \sum_{s \in \mathcal{S}} \mathbf{1}(\mathcal{A}_\tau^\pi(s) \neq \emptyset) d_\rho^\pi(s) \pi_{\text{ref}}(a_{\text{max}}^\pi(s)|s) \left[f\left(\frac{\pi(a_{\text{max}}^\pi(s)|s)}{\pi_{\text{ref}}(a_{\text{max}}^\pi(s)|s)}\right) - f\left(\frac{\tilde{\pi}(a_{\text{max}}^\pi(s)|s)}{\pi_{\text{ref}}(a_{\text{max}}^\pi(s)|s)}\right) \right] . \end{aligned}$$

Since f is convex, for all $u, v \in [0; 1/\pi_{\text{ref}}]$, $f(u) - f(v) \geq f'(v)(u - v)$, we have

$$\begin{aligned} \text{(III)} &\geq \lambda \sum_{s \in \mathcal{S}} d_\rho^\pi(s) \sum_{a \in \mathcal{A}_\tau^\pi(s)} (\pi(a|s) - \tilde{\pi}(a|s)) f'(\tau_\lambda) \quad (\text{since } \tilde{\pi}(a|s)/\pi_{\text{ref}}(a|s) = \tau_\lambda) \\ &\quad + \lambda \sum_{s \in \mathcal{S}} \mathbf{1}(\mathcal{A}_\tau^\pi(s) \neq \emptyset) d_\rho^\pi(s) [\pi(a_{\text{max}}^\pi(s)|s) - \tilde{\pi}(a_{\text{max}}^\pi(s)|s)] f'\left(\frac{\tilde{\pi}(a_{\text{max}}^\pi(s)|s)}{\pi_{\text{ref}}(a_{\text{max}}^\pi(s)|s)}\right) , \end{aligned}$$

Next, using that

$$\frac{\tilde{\pi}(a_{\text{max}}^\pi(s)|s)}{\pi_{\text{ref}}(a_{\text{max}}^\pi(s)|s)} \geq \frac{\pi(a_{\text{max}}^\pi(s)|s) - \tau_\lambda}{\pi_{\text{ref}}(a_{\text{max}}^\pi(s)|s)} \geq \frac{\pi(a_{\text{max}}^\pi(s)|s) - \pi_{\text{ref}}/2}{\pi_{\text{ref}}(a_{\text{max}}^\pi(s)|s)} \geq 1 - \frac{1}{2} = \frac{1}{2} ,$$

combined with the monotonicity of f' and the fact that $\pi(a_{\text{max}}^\pi(s)|s) - \tilde{\pi}(a_{\text{max}}^\pi(s)|s) \geq 0$ yields

$$\text{(III)} \geq \lambda \sum_{s \in \mathcal{S}} d_\rho^\pi(s) \sum_{a \in \mathcal{A}_\tau^\pi(s)} (\pi(a|s) - \tilde{\pi}(a|s)) f'(\tau_\lambda) \quad (\text{since } \tilde{\pi}(a|s)/\pi_{\text{ref}}(a|s) = \tau_\lambda)$$

$$+ \lambda \sum_{s \in \mathcal{S}} \mathbf{1}(\mathcal{A}_\tau^\pi(s) \neq \emptyset) d_\rho^\pi(s) [\pi(a_{\max}^\pi(s)|s) - \tilde{\pi}(a_{\max}^\pi(s)|s)] f' \left(\frac{\tilde{\pi}(a_{\max}^\pi(s)|s)}{\pi_{\text{ref}}(a_{\max}^\pi(s)|s)} \right) ,$$

Additionally, since

$$0 \leq \pi(a_{\max}^\pi(s)|s) - \tilde{\pi}(a_{\max}^\pi(s)|s) \leq \sum_{a \in \mathcal{A}_\tau^\pi(s)} (\pi(a|s) - \tilde{\pi}(a|s)) \leq \tau_\lambda \sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) ,$$

implies

$$\begin{aligned} \text{(III)} &\geq -\frac{\lambda}{2} \sum_{s \in \mathcal{S}} d_\rho^\pi(s) \mathbf{1}(\mathcal{A}_\tau^\pi(s) \neq \emptyset) \left(\sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) \right) \tau_\lambda f'(\tau_\lambda) \\ &\quad + \lambda \sum_{s \in \mathcal{S}} d_\rho^\pi(s) \mathbf{1}(\mathcal{A}_\tau^\pi(s) \neq \emptyset) \left(\sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) \right) \tau_\lambda f'(1/2) , \\ &\geq -\frac{\lambda}{4} \sum_{s \in \mathcal{S}} d_\rho^\pi(s) \mathbf{1}(\mathcal{A}_\tau^\pi(s) \neq \emptyset) \left(\sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) \right) \tau_\lambda f'(\tau_\lambda) , \end{aligned}$$

where in the last inequality, we used that $f'(\tau_\lambda) \leq -4|f'(1/2)|$. Hence, by using [A_ρ](#), we can lower bound this term as follows

$$\text{(III)} \geq -\frac{\lambda}{4} (1 - \gamma) \min_{s \in \mathcal{S}} \{\rho(s)\} \max_{s \in \mathcal{S}} \left\{ \sum_{a \in \mathcal{A}_\tau^\pi(s)} \pi_{\text{ref}}(a|s) \right\} \tau_\lambda f'(\tau_\lambda) .$$

Collecting these lower bounds and using that

$$f'(\tau_\lambda) \leq -\frac{16 + 8\gamma\lambda d_f}{\lambda(1 - \gamma)^2 \rho_{\min}}$$

concludes the proof. \square

Finally, we define the operator that maps each policy to one corresponding parameter

$$M^f : \Pi \rightarrow \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$$

by

$$M^f(\pi)(s, a) := f' \left(\frac{\pi(a|s)}{\pi_{\text{ref}}(a|s)} \right) - f' \left(\frac{\pi(a_{|\mathcal{A}|}|s)}{\pi_{\text{ref}}(a_{|\mathcal{A}|}|s)} \right), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (59)$$

Finally, we define the improvement operator on the logit space as

$$\mathcal{T}_\tau := M^f \circ \mathcal{U}_\tau .$$

The following lemma shows that M^f successfully recovers a parameter that gives the policy and that \mathcal{T}_τ improves the value of the objective when $\lambda = \tau_\lambda$.

Lemma D.2. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy [A_f\(\$\pi_{\text{ref}}\$ \)](#) and [P\(\$\pi_{\text{ref}}\$ \)](#) respectively. Assume in addition that the initial distribution ρ satisfies [A_ρ](#). For any policy π , it holds that*

$$\pi_{M^f(\pi)}^f = \pi ,$$

Additionally, for any $\theta \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have that

$$v_{\mathcal{T}_{\tau_\lambda}^f}^f(\theta) \geq v_\theta^f , \quad \pi_{\mathcal{T}_{\tau_\lambda}^f}^f(\theta) \geq \pi_{\text{ref}} \tau_\lambda .$$

Proof. The proof follows immediately from a combination of equality (40) in Corollary [B.7](#), (59), and Lemma [D.1](#). \square

E CONVERGENCE ANALYSIS OF STOCHASTIC POLICY GRADIENT

In this section, we aim to derive under $\mathbf{A}_f(\pi_{\text{ref}})$ and \mathbf{A}_ρ non-asymptotic convergence rates for f -PG. First, we establish a bound on the bias and variance of the REINFORCE estimator defined in (11).

E.1 BOUNDING THE BIAS AND VARIANCE OF THE STOCHASTIC ESTIMATOR

First, recall the expression of the stochastic estimator of the gradient

$$\begin{aligned} g_z^f(\theta) &= \frac{1}{B} \sum_{b=0}^{B-1} \sum_{h=0}^{H-1} \sum_{\ell=0}^h \frac{\partial \log \pi_\theta^f(a_\ell | s_\ell)}{\partial \theta} \gamma^h r(s_h, a_h) \\ &\quad - \frac{\lambda}{B} \sum_{b=0}^{B-1} \sum_{h=0}^{H-1} \sum_{\ell=0}^{h-1} \frac{\partial \log \pi_\theta^f(a_\ell | s_\ell)}{\partial \theta} \gamma^h D^f(\pi_\theta^f(\cdot | s_h) \| \pi_{\text{ref}}(\cdot | s_h)) - \lambda \frac{1}{B} \sum_{b=0}^{B-1} \sum_{h=0}^{H-1} \gamma^h F_\theta^f(s_h), \end{aligned} \quad (60)$$

where $z = (s_{0:H-1}^b, a_{0:H-1}^b)_{b=0}^{B-1} \in (\mathcal{S} \cdot \mathcal{A})^{H \cdot B}$, and we recall that for any $s \in \mathcal{S}$, $F_\theta^f(s)$ is a vector of size $|\mathcal{S}| \times |\mathcal{A}|$ defined as

$$[F_\theta^f(s)]_{(s', b)} = \mathbf{1}_s(s') W_\theta^f(s) w_\theta^f(b|s) \left[f' \left(\frac{\pi_\theta^f(b|s)}{\pi_{\text{ref}}(b|s)} \right) - \sum_{a \in \mathcal{A}} w_\theta^f(a|s) f' \left(\frac{\pi_\theta^f(a|s)}{\pi_{\text{ref}}(a|s)} \right) \right],$$

and where W_θ^f , Y_θ^f , and w_θ^f are defined in (38) and (39). Finally, define the expected gradient estimator as

$$g^f(\theta) := \mathbb{E}_{Z \sim [\nu(\theta)]^{\otimes B}} [g_Z^f(\theta)], \quad (61)$$

Before bounding the bias and the variance, we give an explicit expression of the derivative of the log probability that appears in the expression of our stochastic gradient estimator. We also provide a bound on the derivative of the log probabilities and on the matrix $F_\theta^f(s)$ for any state $s \in \mathcal{S}$.

Lemma E.1. *Assume that, for some $\pi_{\text{ref}} > 0$, f satisfy $\mathbf{A}_f(\pi_{\text{ref}})$. For any $\theta \in \mathbb{R}^{|\mathcal{S}| \cdot |\mathcal{A}|}$, $(s, s', a, b) \in \mathcal{S}^2 \times \mathcal{A}^2$, we have*

$$\frac{\partial \log \pi_\theta^f(a|s)}{\partial \theta(s', b)} = \mathbf{1}_{s'}(s) \frac{W_\theta^f(s)}{\pi_\theta^f(a|s)} \left[\mathbf{1}_b(a) w_\theta^f(a|s) - w_\theta^f(a|s) w_\theta^f(b|s) \right].$$

Additionally, we have that

$$\left\| \frac{\partial \log \pi_\theta^f(a|s)}{\partial \theta} \right\|_2 \leq \frac{2 W_\theta^f(s) w_\theta^f(a|s)}{\pi_\theta^f(a|s)}, \quad \left\| F_\theta^f(s) \right\|_2 \leq 2 W_\theta^f(s) Y_\theta^f(s).$$

Proof. The proof follows from the log-derivative trick and the expression of the derivative of the policy provided in Corollary B.7. \square

Next, we establish a REINFORCE-type formula for the gradient of the objective.

Lemma E.2. *Assume that, for some $\pi_{\text{ref}} > 0$, f satisfy $\mathbf{A}_f(\pi_{\text{ref}})$. It holds that*

$$\begin{aligned} \frac{\partial v_\theta^f(\rho)}{\partial \theta(s, b)} &= \mathbb{E} \left[\sum_{t=0}^{\infty} \sum_{\ell=0}^t \frac{\partial \log \pi_\theta^f(A_\ell | S_\ell)}{\partial \theta(s, b)} \gamma^t r(S_t, A_t) \right] \\ &\quad - \lambda \mathbb{E} \left[\sum_{t=0}^{\infty} \sum_{\ell=0}^{t-1} \frac{\partial \log \pi_\theta^f(A_\ell | S_\ell)}{\partial \theta(s, b)} \gamma^t D^f(\pi_\theta^f(\cdot | S_t) \| \pi_{\text{ref}}(\cdot | S_t)) \right] \\ &\quad - \lambda \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \mathbf{1}_s(S_t) W_\theta^f(s) w_\theta^f(b|s) \left[f' \left(\frac{\pi_\theta^f(b|s)}{\pi_{\text{ref}}(b|s)} \right) - \sum_{a \in \mathcal{A}} w_\theta^f(a|s) f' \left(\frac{\pi_\theta^f(a|s)}{\pi_{\text{ref}}(a|s)} \right) \right] \right]. \end{aligned}$$

Proof. Fix a parameter $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, a horizon T , and a divergence generator f . For any truncated trajectory $z = (s_t, a_t)_{t=0}^{T-1} \in (\mathcal{S} \times \mathcal{A})^T$, we define its probability as

$$\nu_T^f(\theta; z) = \rho(s_0) \pi_\theta^f(a_0|s_0) \prod_{t=1}^{T-1} \mathbb{P}(s_t|s_{t-1}, a_{t-1}) \pi_\theta^f(a_t|s_t),$$

and the regularized return

$$R_{\theta,T}^f(z) = \sum_{t=0}^{T-1} \gamma^t \left(r(s_t, a_t) - \lambda \frac{\pi_{\text{ref}}^f(a_t|s_t)}{\pi_\theta^f(a_t|s_t)} f\left(\frac{\pi_\theta^f(a_t|s_t)}{\pi_{\text{ref}}^f(a_t|s_t)}\right) \right). \quad (62)$$

The finite-horizon objective is

$$J_T^f(\theta) = \sum_{z \in (\mathcal{S} \times \mathcal{A})^T} \nu_T^f(\theta; z) R_{\theta,T}^f(z).$$

Fix $(s, b) \in \mathcal{S} \times \mathcal{A}$. Differentiating this finite-horizon objective gives

$$\frac{\partial J_T^f(\theta)}{\partial \theta(s, b)} = \underbrace{\sum_{z \in (\mathcal{S} \times \mathcal{A})^T} \frac{\partial \nu_T^f(\theta; z)}{\partial \theta(s, b)} R_{\theta,T}^f(z)}_{\text{(A)}} + \underbrace{\sum_{z \in (\mathcal{S} \times \mathcal{A})^T} \nu_T^f(\theta; z) \frac{\partial R_{\theta,T}^f(z)}{\partial \theta(s, b)}}_{\text{(B)}}. \quad (63)$$

We now treat these two terms separately.

Term (A). Using the log-derivative trick and the fact that the only terms that depend on θ in $\nu_T^f(\theta; z)$ are the ones that depend on the policy itself, we obtain

$$\frac{\partial \nu_T^f(\theta; z)}{\partial \theta(s, b)} = \nu_T^f(\theta; z) \sum_{t=0}^{T-1} \frac{\partial \log \pi_\theta^f(a_t|s_t)}{\partial \theta(s, b)}. \quad (64)$$

Plugging expressions (62) and (64) in (A) then gives

$$\text{(A)} = \sum_{z \in (\mathcal{S} \times \mathcal{A})^T} \nu_T^f(\theta; z) \sum_{t=0}^{T-1} \sum_{\ell=0}^{T-1} \frac{\partial \log \pi_\theta^f(a_\ell|s_\ell)}{\partial \theta(s, b)} \gamma^t \left(r(s_t, a_t) - \lambda \frac{\pi_{\text{ref}}^f(a_t|s_t)}{\pi_\theta^f(a_t|s_t)} f\left(\frac{\pi_\theta^f(a_t|s_t)}{\pi_{\text{ref}}^f(a_t|s_t)}\right) \right).$$

Now observe that for $\ell > t$, the sum of the log-gradient derivatives over $z \in (\mathcal{S} \times \mathcal{A})^T$ is 0. Therefore (A) reduces to

$$\text{(A)} = \sum_{z \in (\mathcal{S} \times \mathcal{A})^T} \nu_T^f(\theta; z) \sum_{t=0}^{T-1} \sum_{\ell=0}^t \frac{\partial \log \pi_\theta^f(a_\ell|s_\ell)}{\partial \theta(s, b)} \gamma^t \left(r(s_t, a_t) - \lambda \frac{\pi_{\text{ref}}^f(a_t|s_t)}{\pi_\theta^f(a_t|s_t)} f\left(\frac{\pi_\theta^f(a_t|s_t)}{\pi_{\text{ref}}^f(a_t|s_t)}\right) \right). \quad (65)$$

Term (B). Taking the derivative of (62), we have

$$\frac{\partial R_{\theta,T}^f(z)}{\partial \theta(s, b)} = \lambda \sum_{t=0}^{T-1} \gamma^t \left(\frac{\partial \pi_\theta^f(a_t|s_t)}{\partial \theta(s, b)} \frac{\pi_{\text{ref}}^f(a_t|s_t)}{\pi_\theta^f(a_t|s_t)^2} f\left(\frac{\pi_\theta^f(a_t|s_t)}{\pi_{\text{ref}}^f(a_t|s_t)}\right) - \frac{\partial \pi_\theta^f(a_t|s_t)}{\partial \theta(s, b)} \frac{f'\left(\frac{\pi_\theta^f(a_t|s_t)}{\pi_{\text{ref}}^f(a_t|s_t)}\right)}{\pi_\theta^f(a_t|s_t)} \right).$$

Summing over $z \in (\mathcal{S} \times \mathcal{A})^T$ and using the log-derivative trick gives

$$\text{(B)} = \lambda \sum_{z \in (\mathcal{S} \times \mathcal{A})^T} \nu_T^f(\theta; z) \sum_{t=0}^{T-1} \gamma^t \frac{\partial \log \pi_\theta^f(a_t|s_t)}{\partial \theta(s, b)} \left[\frac{\pi_{\text{ref}}^f(a_t|s_t)}{\pi_\theta^f(a_t|s_t)} f\left(\frac{\pi_\theta^f(a_t|s_t)}{\pi_{\text{ref}}^f(a_t|s_t)}\right) - f'\left(\frac{\pi_\theta^f(a_t|s_t)}{\pi_{\text{ref}}^f(a_t|s_t)}\right) \right]. \quad (66)$$

Plugging the expressions (65) and (66) in (63) gives

$$\frac{\partial J_T^f(\theta)}{\partial \theta(s, b)} = \mathbb{E} \left[\sum_{t=0}^{T-1} \sum_{\ell=0}^t \frac{\partial \log \pi_\theta^f(A_\ell|S_\ell)}{\partial \theta(s, b)} \gamma^t \left(r(S_t, A_t) - \lambda \frac{\pi_{\text{ref}}^f(A_t|S_t)}{\pi_\theta^f(A_t|S_t)} f\left(\frac{\pi_\theta^f(A_t|S_t)}{\pi_{\text{ref}}^f(A_t|S_t)}\right) \right) \right]$$

$$+ \lambda \mathbb{E} \left[\sum_{t=0}^{T-1} \gamma^t \frac{\partial \log \pi_\theta^f(A_t|S_t)}{\partial \theta(s, b)} \left(\frac{\pi_{\text{ref}}(A_t|S_t)}{\pi_\theta^f(A_t|S_t)} f \left(\frac{\pi_\theta^f(A_t|S_t)}{\pi_{\text{ref}}(A_t|S_t)} \right) - f' \left(\frac{\pi_\theta^f(A_t|S_t)}{\pi_{\text{ref}}(A_t|S_t)} \right) \right) \right].$$

The previous term can be rewritten as

$$\begin{aligned} \frac{\partial J_T^f(\theta)}{\partial \theta(s, b)} &= \mathbb{E}_{Z \sim \nu_T^f(\theta)} \left[\sum_{t=0}^{T-1} \sum_{\ell=0}^t \frac{\partial \log \pi_\theta^f(A_\ell|S_\ell)}{\partial \theta(s, b)} \gamma^t r(S_t, A_t) \right] \\ &\quad - \lambda \mathbb{E}_{Z \sim \nu_T^f(\theta)} \left[\sum_{t=0}^{T-1} \sum_{\ell=0}^{t-1} \frac{\partial \log \pi_\theta^f(A_\ell|S_\ell)}{\partial \theta(s, b)} \gamma^t \frac{\pi_{\text{ref}}(A_t|S_t)}{\pi_\theta^f(A_t|S_t)} f \left(\frac{\pi_\theta^f(A_t|S_t)}{\pi_{\text{ref}}(A_t|S_t)} \right) \right] \\ &\quad - \lambda \mathbb{E}_{Z \sim \nu_T^f(\theta)} \left[\sum_{t=0}^{T-1} \gamma^t \frac{\partial \log \pi_\theta^f(A_t|S_t)}{\partial \theta(s, b)} f' \left(\frac{\pi_\theta^f(A_t|S_t)}{\pi_{\text{ref}}(A_t|S_t)} \right) \right]. \end{aligned}$$

Next, we apply the tower property by taking the conditional expectation with respect to $\mathcal{G}_t := \sigma(S_0, A_0, \dots, S_t)$ on the second expectation and using Lemma E.1 in the third expectation, gives

$$\begin{aligned} \frac{\partial J_T^f(\theta)}{\partial \theta(s, b)} &= \mathbb{E} \left[\sum_{t=0}^{T-1} \sum_{\ell=0}^t \frac{\partial \log \pi_\theta^f(A_\ell|S_\ell)}{\partial \theta(s, b)} \gamma^t r(S_t, A_t) \right] \\ &\quad - \lambda \mathbb{E} \left[\sum_{t=0}^{T-1} \sum_{\ell=0}^{t-1} \frac{\partial \log \pi_\theta^f(A_\ell|S_\ell)}{\partial \theta(s, b)} \gamma^t D^f(\pi_\theta^f(\cdot|S_t) \| \pi_{\text{ref}}(\cdot|S_t)) \right] \\ &\quad - \lambda \mathbb{E} \left[\sum_{t=0}^{T-1} \gamma^t \mathbf{1}_s(S_t) \frac{W_\theta^f(S_t) w_\theta^f(b|S_t)}{\pi_\theta^f(A_t|S_t)} [\mathbf{1}_b(A_t) - w_\theta^f(A_t|S_t)] f' \left(\frac{\pi_\theta^f(A_t|S_t)}{\pi_{\text{ref}}(A_t|S_t)} \right) \right]. \end{aligned}$$

Applying the tower property again by taking the conditional expectation with respect to \mathcal{G}_t in the third expectation gives

$$\begin{aligned} \frac{\partial J_T^f(\theta)}{\partial \theta(s, b)} &= \mathbb{E} \left[\sum_{t=0}^{T-1} \sum_{\ell=0}^t \frac{\partial \log \pi_\theta^f(A_\ell|S_\ell)}{\partial \theta(s, b)} \gamma^t r(S_t, A_t) \right] \\ &\quad - \lambda \mathbb{E} \left[\sum_{t=0}^{T-1} \sum_{\ell=0}^{t-1} \frac{\partial \log \pi_\theta^f(A_\ell|S_\ell)}{\partial \theta(s, b)} \gamma^t D^f(\pi_\theta^f(\cdot|S_t) \| \pi_{\text{ref}}(\cdot|S_t)) \right] \\ &\quad - \lambda \mathbb{E} \left[\sum_{t=0}^{T-1} \gamma^t \mathbf{1}_s(S_t) W_\theta^f(S_t) w_\theta^f(b|S_t) \left[f' \left(\frac{\pi_\theta^f(b|S_t)}{\pi_{\text{ref}}(b|S_t)} \right) - \sum_{a \in \mathcal{A}} w_\theta^f(a|S_t) f' \left(\frac{\pi_\theta^f(a|S_t)}{\pi_{\text{ref}}(a|S_t)} \right) \right] \right]. \end{aligned}$$

Taking $T \rightarrow +\infty$ and applying the dominated convergence theorem concludes the proof. \square

The following lemma establishes a bound on the variance and bias of the stochastic estimator.

Lemma E.3. Assume that, for some $\pi_{\text{ref}} > 0$, f satisfy **A_f**(π_{ref}). There exists a constant $\beta_f \geq 0$ such that, for any parameter $\theta \in \mathbb{R}^{|S|+|\mathcal{A}|}$, we have

$$\left\| g^f(\theta) - \frac{\partial v_\theta^f(\rho)}{\partial \theta} \right\|_2 \leq \beta_f,$$

where β_f is an upper bound on the bias defined as

$$\beta_f := \frac{2\gamma^H(H+1)}{(1-\gamma)^2} \omega_f [2 + 2\lambda d_f + \lambda(1-\gamma)y_f],$$

where ω_f is defined in **A_f**(π_{ref}) and d_f , and y_f are defined in (6).

Proof. Using the expression of the gradient truncated at H from (60) and (61), of the true gradient from Lemma E.2, and the triangle inequality, we have

$$\left\| g^f(\theta) - \frac{\partial v_\theta^f(\rho)}{\partial \theta} \right\|_2 \leq \sum_{t=H}^{\infty} \sum_{\ell=0}^t \gamma^t \left\| \mathbb{E}_\rho^{\pi_\theta^f} \left[\frac{\partial \log \pi_\theta^f(A_\ell|S_\ell)}{\partial \theta} r(S_t, A_t) \right] \right\|_2$$

$$+ \sum_{t=H}^{\infty} \sum_{\ell=0}^{t-1} \lambda \gamma^t \left\| \mathbb{E}_{\rho}^{\pi_{\theta}^f} \left[\frac{\partial \log \pi_{\theta}^f(A_{\ell}|S_{\ell})}{\partial \theta} D^f(\pi_{\theta}^f(\cdot|S_t) \|\pi_{\text{ref}}(\cdot|S_t)) \right] \right\|_2 + \lambda \sum_{t=H}^{\infty} \gamma^t \left\| \mathbb{E}_{\rho}^{\pi_{\theta}^f} [F_{\theta}^f(S_t)] \right\|_2.$$

Next, applying Lemma E.1 combined with the triangle inequality yields

$$\begin{aligned} \left\| g^f(\theta) - \frac{\partial v_{\theta}^f(\rho)}{\partial \theta} \right\|_2 &\leq \sum_{t=H}^{\infty} \sum_{\ell=0}^t \gamma^t \mathbb{E}_{\rho}^{\pi_{\theta}^f} \left[\frac{2 W_{\theta}^f(S_{\ell}) w_{\theta}^f(A_{\ell}|S_{\ell})}{\pi_{\theta}^f(A_{\ell}|S_{\ell})} |r(S_t, A_t)| \right] \\ &\quad + \sum_{t=H}^{\infty} \sum_{\ell=0}^{t-1} \lambda \gamma^t \mathbb{E}_{\rho}^{\pi_{\theta}^f} \left[\frac{2 W_{\theta}^f(S_{\ell}) w_{\theta}^f(A_{\ell}|S_{\ell})}{\pi_{\theta}^f(A_{\ell}|S_{\ell})} D^f(\pi_{\theta}^f(\cdot|S_t) \|\pi_{\text{ref}}(\cdot|S_t)) \right] \\ &\quad + \lambda \sum_{t=H}^{\infty} \gamma^t \mathbb{E}_{\rho}^{\pi_{\theta}^f} [2 W_{\theta}^f(S_t) Y_{\theta}^f(S_t)]. \end{aligned}$$

We define the following filtration, for $t \geq 0$,

$$\mathcal{G}_t = \sigma(S_0, A_0, \dots, S_t),$$

Next, applying the tower property of the conditional expectation by conditioning on \mathcal{G}_t , bounding the reward and the divergence respectively by 1, and $\max_{s \in \mathcal{S}} \sup_{\nu \in \mathcal{P}(\mathcal{A})} D^f(\nu \|\pi_{\text{ref}}(\cdot|s))$, and using that $w_{\theta}^f(\cdot|s) \in \mathcal{P}(\mathcal{A})$ for any $s \in \mathcal{S}$ yields

$$\begin{aligned} \left\| g^f(\theta) - \frac{\partial v_{\theta}^f(\rho)}{\partial \theta} \right\|_2 &\leq 2 \sum_{t=H}^{\infty} \sum_{\ell=0}^t \gamma^t \left\| W_{\theta}^f \right\|_{\infty} \\ &\quad + 2\lambda \sum_{t=H}^{\infty} \sum_{\ell=0}^{t-1} \gamma^t \left\| W_{\theta}^f \right\|_{\infty} \sup_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} D^f(\nu \|\pi_{\text{ref}}(\cdot|s)) + 2\lambda \sum_{t=H}^{\infty} \gamma^t \left\| W_{\theta}^f \right\|_{\infty} \left\| Y_{\theta}^f \right\|_{\infty}. \end{aligned}$$

Finally, using that

$$\sum_{t=H}^{\infty} \gamma^t \leq \frac{\gamma^H}{1-\gamma}, \quad \sum_{t=H}^{\infty} \gamma^t (t-1) \leq 2\gamma^H \frac{H}{(1-\gamma)^2}, \quad \sum_{t=H}^{\infty} \gamma^t t \leq 2\gamma^H \frac{H+1}{(1-\gamma)^2},$$

combined with Lemma B.6 completes the proof. \square

Lemma E.4. Assume that, for some $\pi_{\text{ref}} > 0$, f satisfy $\mathbf{A}_f(\pi_{\text{ref}})$. For any $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, it holds that

$$\mathbb{E}_{Z \sim [\nu(\theta)]^{\otimes B}} \left[\left\| g^f(\theta) - g_Z^f(\theta) \right\|_2^2 \right] \leq \frac{\sigma_f^2}{B},$$

where we have defined

$$\sigma_f^2 := \frac{12}{(1-\gamma)^4} [\omega_f^3 + \lambda^2 \gamma^2 \omega_f^3 d_f^2 + \lambda^2 (1-\gamma)^2 \omega_f^2 y_f^2],$$

and where ω_f , d_f , and y_f are defined in $\mathbf{A}_f(\pi_{\text{ref}})$ and (6).

Proof. Firstly, define for $\xi = (s_h, a_h)_{h=0}^{H-1} \in (\mathcal{S} \times \mathcal{A})^H$

$$\begin{aligned} u_{\xi}(\theta) &:= \sum_{h=0}^{H-1} \sum_{\ell=0}^h \frac{\partial \log \pi_{\theta}^f(a_{\ell}|s_{\ell})}{\partial \theta} \gamma^h r(s_h, a_h) \\ &\quad - \lambda \sum_{h=0}^{H-1} \sum_{\ell=0}^{h-1} \frac{\partial \log \pi_{\theta}^f(a_{\ell}|s_{\ell})}{\partial \theta} \gamma^h D^f(\pi_{\theta}^f(\cdot|s_h) \|\pi_{\text{ref}}(\cdot|s_h)) - \lambda \sum_{h=0}^{H-1} \gamma^h F_{\theta}^f(s_h), \end{aligned}$$

Importantly, for a given $Z \sim [\nu(\theta)]^{\otimes B}$, denoting by $Z = (Z_0, \dots, Z_{B-1})$, it holds that

$$g_Z^f(\theta) = \frac{1}{B} \sum_{b=0}^{B-1} u_{Z_b}(\theta), \quad \text{and } g^f(\theta) = \mathbb{E}_{\mathfrak{Z} \sim \nu(\theta)} [u_{\mathfrak{Z}}(\theta)].$$

Using that the variables (Z_0, \dots, Z_{B-1}) are independent and identically distributed, we get

$$\begin{aligned} \mathbb{E}_{Z \sim [\nu(\theta)]^{\otimes B}} \left[\left\| g^f(\theta) - g_Z^f(\theta) \right\|_2^2 \right] &= \mathbb{E}_{Z \sim [\nu(\theta)]^{\otimes B}} \left[\left\| \frac{1}{B} \sum_{b=0}^{B-1} u_{Z_b}(\theta) - g^f(\theta) \right\|_2^2 \right] \\ &= \frac{1}{B} \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left\| u_{\mathfrak{I}}(\theta) - g^f(\theta) \right\|_2^2 \right] \\ &\leq \frac{1}{B} \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left\| u_{\mathfrak{I}}(\theta) \right\|_2^2 \right], \end{aligned} \quad (67)$$

where in the last inequality, we used that the second moment of a random variable dominates its variance. Next, using Jensen's inequality combined with the convexity of the square function, we have

$$\begin{aligned} \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left\| u_{\mathfrak{I}}(\theta) \right\|_2^2 \right] &\leq 3 \mathbb{E}_{\mathfrak{I}} \left[\left\| \sum_{h=0}^{H-1} \sum_{\ell=0}^h \frac{\partial \log \pi_{\theta}^f(A_{\ell} | S_{\ell})}{\partial \theta} \gamma^h r(S_h, A_h) \right\|_2^2 \right] \\ &+ 3\lambda^2 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left\| \sum_{h=0}^{H-1} \sum_{\ell=0}^{h-1} \frac{\partial \log \pi_{\theta}^f(A_{\ell} | S_{\ell})}{\partial \theta} \gamma^h D^f(\pi_{\theta}^f(\cdot | S_h) \| \pi_{\text{ref}}(\cdot | S_h)) \right\|_2^2 \right] \\ &+ 3\lambda^2 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left\| \sum_{h=0}^{H-1} \gamma^h F_{\theta}^f(S_h) \right\|_2^2 \right], \end{aligned}$$

Applying the triangle inequality and the fact that the reward and the divergence are positive yields

$$\begin{aligned} \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left\| u_{\mathfrak{I}}(\theta) \right\|_2^2 \right] &\leq 3 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left(\sum_{h=0}^{H-1} \sum_{\ell=0}^h \gamma^h \left\| \frac{\partial \log \pi_{\theta}^f(A_{\ell} | S_{\ell})}{\partial \theta} \right\|_2 r(S_h, A_h) \right)^2 \right] \\ &+ 3\lambda^2 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left(\sum_{h=0}^{H-1} \sum_{\ell=0}^{h-1} \gamma^h \left\| \frac{\partial \log \pi_{\theta}^f(A_{\ell} | S_{\ell})}{\partial \theta} \right\|_2 D^f(\pi_{\theta}^f(\cdot | S_h) \| \pi_{\text{ref}}(\cdot | S_h)) \right)^2 \right] \\ &+ 3\lambda^2 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left(\sum_{h=0}^{H-1} \gamma^h \left\| F_{\theta}^f(S_h) \right\|_2 \right)^2 \right], \end{aligned}$$

Combining Lemma E.1 and the fact that the reward is bounded between 0 and 1 gives

$$\begin{aligned} \mathbb{E}_{\mathfrak{I}} \left[\left\| u_{\mathfrak{I}}(\theta) \right\|_2^2 \right] &\leq 3 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left(\sum_{h=0}^{H-1} \sum_{\ell=0}^h 2\gamma^{h/2} \cdot \gamma^{h/2} \frac{W_{\theta}^f(S_{\ell}) w_{\theta}^f(A_{\ell} | S_{\ell})}{\pi_{\theta}^f(A_{\ell} | S_{\ell})} \right)^2 \right] \\ &+ 3\lambda^2 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left(\sum_{h=0}^{H-1} \sum_{\ell=0}^{h-1} 2\gamma^{h/2} \cdot \gamma^{h/2} \frac{W_{\theta}^f(S_{\ell}) w_{\theta}^f(A_{\ell} | S_{\ell})}{\pi_{\theta}^f(A_{\ell} | S_{\ell})} \sup_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} D^f(\nu \| \pi_{\text{ref}}(\cdot | s)) \right)^2 \right] \\ &+ 3\lambda^2 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left(\sum_{h=0}^{H-1} 2\gamma^{h/2} \cdot \gamma^{h/2} \left\| W_{\theta}^f \right\|_{\infty} \left\| Y_{\theta}^f \right\|_{\infty} \right)^2 \right], \end{aligned}$$

Next, applying the Cauchy-Schwarz inequality gives

$$\begin{aligned} \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left\| u_{\mathfrak{I}}(\theta) \right\|_2^2 \right] &\leq 3 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\left(\sum_{h=0}^{H-1} \sum_{\ell=0}^h 4\gamma^h \right) \left(\sum_{h=0}^{H-1} \sum_{\ell=0}^h \gamma^h \frac{W_{\theta}^f(S_{\ell})^2 w_{\theta}^f(A_{\ell} | S_{\ell})^2}{\pi_{\theta}^f(A_{\ell} | S_{\ell})^2} \right) \right] \\ &+ 3\lambda^2 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\sum_{h=0}^{H-1} \sum_{\ell=0}^{h-1} 4\gamma^h \left(\sum_{h=0}^{H-1} \sum_{\ell=0}^{h-1} \gamma^h \frac{W_{\theta}^f(S_{\ell})^2 w_{\theta}^f(A_{\ell} | S_{\ell})^2}{\pi_{\theta}^f(A_{\ell} | S_{\ell})^2} \sup_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} D^f(\nu \| \pi_{\text{ref}}(\cdot | s))^2 \right) \right] \\ &+ 3\lambda^2 \mathbb{E}_{\mathfrak{I} \sim \nu(\theta)} \left[\sum_{h=0}^{H-1} 4\gamma^h \left(\sum_{h=0}^{H-1} \gamma^h \left\| W_{\theta}^f \right\|_{\infty}^2 \left\| Y_{\theta}^f \right\|_{\infty}^2 \right) \right], \end{aligned}$$

We define the following filtration, for $t \geq 0$,

$$\mathcal{G}_t = \sigma(S_0, A_0, \dots, S_t) ,$$

Next, applying the tower property of the conditional expectation by conditioning on \mathcal{G}_t , and using that $w_\theta^f(\cdot|s) \in \mathcal{P}(\mathcal{A})$ for any $s \in \mathcal{S}$ yields

$$\begin{aligned} \mathbb{E}_{\mathfrak{T} \sim \nu(\theta)} \left[\|u_{\mathfrak{T}}(\theta)\|_2^2 \right] &\leq 12 \left\| W_\theta^f \right\|_\infty^2 \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left\{ \frac{w_\theta^f(a|s)}{\pi_\theta^f(a|s)} \right\} \left(\sum_{h=0}^{H-1} \gamma^h (h+1) \right)^2 \\ &+ 12\lambda^2 \left\| W_\theta^f \right\|_\infty^2 \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \left\{ \frac{w_\theta^f(a|s)}{\pi_\theta^f(a|s)} \right\} \sup_{(s,\nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} \left\{ D^f(\nu \| \pi_{\text{ref}}(\cdot|s)) \right\} \left(\sum_{h=0}^{H-1} \gamma^h h \right)^2 \\ &+ 12\lambda^2 \left\| W_\theta^f \right\|_\infty^2 \left\| Y_\theta^f \right\|_\infty^2 \left(\sum_{h=0}^{H-1} \gamma^h \right)^2 . \end{aligned}$$

Next using

$$\sum_{t=0}^{H-1} \gamma^t \leq \frac{1}{1-\gamma} , \quad \sum_{t=0}^{H-1} \gamma^t t \leq \frac{\gamma}{(1-\gamma)^2} , \quad \sum_{t=0}^{H-1} \gamma^t (t+1) \leq \frac{1}{(1-\gamma)^2} ,$$

and plugging in the obtained bound in (67), combined with Lemma B.6 concludes the proof. \square

E.2 SAMPLE COMPLEXITY OF STOCHASTIC f -PG

We now derive convergence rates for f -PG. First, we define the following quantity, which will be the Polyak-Łojasiewicz constant of our function over the optimization space, where policies are guaranteed not to be too ill-conditioned, i.e., all their entries are larger than the τ_λ defined in (58),

$$\underline{\mu}_f := \frac{\lambda(1-\gamma)\rho_{\min}^2 \zeta_f^2}{\omega_f^2} \pi_{\text{ref}}^2 \min_{x \in [\tau_\lambda, \frac{1}{\pi_{\text{ref}}}] } f''(x)^{-2} . \quad (68)$$

As we will prove in this subsection, this quantity represents a lower bound of the non-uniform Łojasiewicz coefficient along the trajectory. In the following lemma, we give a simpler lower bound of $\underline{\mu}_f$ provided that λ is not too large.

Lemma E.5. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_ρ and that λ satisfies*

$$\lambda \leq \frac{4}{(1-\gamma)^2 \rho_{\min}} \min \left(\frac{4}{|f'(\iota_f)|} , \frac{1}{|f'(\frac{1}{2})|} , \frac{4}{|f'(\frac{1}{2}\pi_{\text{ref}})|} \right) .$$

In this case, it holds that

$$\underline{\mu}_f = \frac{\lambda(1-\gamma)\rho_{\min}^2 \zeta_f^2}{\omega_f^2} \pi_{\text{ref}}^2 (f^*)'' \left(-\frac{16 + 8\gamma\lambda d_f}{\lambda(1-\gamma)^2 \rho_{\min}} \right)^2 .$$

Additionally if $\lambda \leq 1/d_f$, then it holds that

$$\underline{\mu}_f \geq \frac{\lambda(1-\gamma)\rho_{\min}^2 \zeta_f^2}{\omega_f^2} \pi_{\text{ref}}^2 (f^*)'' \left(-\frac{24}{\lambda(1-\gamma)^2 \rho_{\min}} \right)^2 .$$

Proof. First note that the first condition on λ implies that $\tau_\lambda < \iota_f$, with ι_f defined in $\mathbf{A}_f(\pi_{\text{ref}})$, and thus

$$\min_{x \in [\tau_\lambda, \frac{1}{\pi_{\text{ref}}}] } f''(x)^{-2} = f''(\tau_\lambda)^{-2} .$$

Additionally, the second and third conditions on λ guarantee that the minimum of τ_λ in (58) is attained in the first term, that is

$$\tau_\lambda = [f']^{-1} \left(-\frac{16 + 8\gamma\lambda d_f}{\lambda(1-\gamma)^2 \rho_{\min}} \right) .$$

Finally, we recall that the convex conjugate of the divergence generator f defined in (25) satisfies, for any $y \in (-\infty, f'(\frac{1}{\pi_{\text{ref}}}))$,

$$(f^*)''(y) = \frac{1}{f''([f']^{-1}(y))} .$$

Thus, we obtain that $f''(\tau_\lambda)^{-2} = (f^*)''(\frac{-16+8\gamma\lambda d_f}{\lambda(1-\gamma)^2 \rho_{\min}})^2$ which concludes the proof. \square

In the following, we define the filtration adapted to the iterates of f -PG as

$$\mathcal{F}_t := \sigma \left(Z_t : t \in \{0, \dots, T-1\} \right) .$$

The following theorem gives convergence rates of f -PG.

Theorem E.6. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_ρ . Fix $\eta \leq 1/2L_f$, a given temperature λ , and consider the iterates $(\theta_t)_{t=0}^\infty$ of the algorithm f -PG. It holds almost surely that*

$$\inf_{t \geq 0} \mu_f(\theta_t) \geq \underline{\mu}_f . \quad (69)$$

Additionally, for any $t \geq 0$ we have that

$$v_\star^f(\rho) - \mathbb{E} \left[v_{\theta_t}^f(\rho) \right] \leq (1 - \underline{\mu}_f \eta / 4)^t (v_\star^f(\rho) - v_{\theta_0}^f(\rho)) + \frac{6\eta\sigma_f^2}{B\underline{\mu}_f} + \frac{6\beta_f^2}{\underline{\mu}_f} .$$

Proof. Recall that for any $t \in [T]$, we have that

$$\theta_t = \mathcal{T}_{\tau_\lambda}(\bar{\theta}_t) .$$

Hence, by Lemma D.2, for any $t \in [T]$ it holds that

$$\pi_{\theta_t}^f \geq \tau_\lambda \pi_{\text{ref}} .$$

Combining the previous inequality with the expression of the coefficient μ_f provided in Theorem C.1 proves the first statement of the lemma. Next, using Theorem B.12 gives

$$v_{\theta_{t+1}}^f(\rho) \geq v_{\theta_t}^f(\rho) + 2\eta \langle \nabla v_{\theta_t}^f(\rho), g_{Z_t}^f(\theta_t) \rangle - \frac{\eta^2 L_f}{2} \left\| g_{Z_t}^f(\theta_t) \right\|_2^2 .$$

Next, taking the conditional expectation with respect to \mathcal{F}_t and adding and subtracting $\nabla v_{\theta_t}^f(\rho)$ in the dot product gives

$$\begin{aligned} \mathbb{E} \left[v_{\theta_{t+1}}^f(\rho) \middle| \mathcal{F}_t \right] &\geq v_{\theta_t}^f(\rho) + 2\eta \left\| \nabla v_{\theta_t}^f(\rho) \right\|^2 + \underbrace{2\eta \langle \nabla v_{\theta_t}^f(\rho), g^f(\theta_t) - \nabla v_{\theta_t}^f(\rho) \rangle}_{(\mathbf{K}_1)} \\ &\quad - \underbrace{\frac{\eta^2 L_f}{2} \mathbb{E} \left[\left\| g_{Z_t}^f(\theta_t) \right\|_2^2 \middle| \mathcal{F}_t \right]}_{(\mathbf{K}_2)} . \end{aligned} \quad (70)$$

We now bound each of these terms separately.

Bounding \mathbf{K}_1 . Using the Cauchy-Schwarz inequality, yields

$$\begin{aligned} 2\eta \langle \nabla v_{\theta_t}^f(\rho), g^f(\theta_t) - \nabla v_{\theta_t}^f(\rho) \rangle &\geq -2\eta \left\| \nabla v_{\theta_t}^f(\rho) \right\|_2 \left\| g^f(\theta_t) - \nabla v_{\theta_t}^f(\rho) \right\|_2 \\ &= -2 \cdot \eta^{1/2} \left\| \nabla v_{\theta_t}^f(\rho) \right\|_2 \cdot \eta^{1/2} \beta_f , \end{aligned}$$

where in the last inequality we used Lemma E.3. Next, using Young's inequality gives

$$2\eta \langle \nabla v_{\theta_t}^f(\rho), g^f(\theta_t) - \nabla v_{\theta_t}^f(\rho) \rangle \geq -\eta \left\| \nabla v_{\theta_t}^f(\rho) \right\|_2^2 - \eta \beta_f^2 . \quad (71)$$

Bounding \mathbf{K}_2 . Using the convexity of the square function with Jensen's inequality gives

$$\begin{aligned} \mathbb{E} \left[\left\| \mathbf{g}_{Z_t}^f(\theta_t) \right\|_2^2 \middle| \mathcal{F}_t \right] &= \mathbb{E} \left[\left\| \mathbf{g}_{Z_t}^f(\theta_t) - \mathbf{g}^f Z_t(\theta_t) + \mathbf{g}^f Z_t(\theta_t) - \nabla v_{\theta_t}^f(\rho) + \nabla v_{\theta_t}^f(\rho) \right\|_2^2 \middle| \mathcal{F}_t \right] \\ &\leq 3\beta_f^2 + 3 \left\| \nabla v_{\theta_t}^f(\rho) \right\|_2^2 + \frac{3\sigma_f^2}{B}, \end{aligned} \quad (72)$$

where we used Lemma E.3 and Lemma E.4. Plugging in the bounds (71) on \mathbf{K}_1 and (72) on \mathbf{K}_2 in (70) gives

$$\mathbb{E} \left[v_{\theta_{t+1}}^f(\rho) \middle| \mathcal{F}_t \right] \geq v_{\theta_t}^f(\rho) + \eta \left\| \nabla v_{\theta_t}^f(\rho) \right\|_2^2 - \left(\eta + \frac{3\eta^2 L_f}{2} \right) \beta_f^2 - \frac{3\eta^2 L_f}{2} \left\| \nabla v_{\theta_t}^f(\rho) \right\|_2^2 - \frac{3\eta^2 L_f \sigma_f^2}{2B}.$$

Taking the expectation with respect to all the stochasticity, multiplying both sides by -1 , and adding $v_\star^f(\rho)$ gives

$$\begin{aligned} v_\star^f(\rho) - \mathbb{E} \left[v_{\theta_{t+1}}^f(\rho) \right] &\leq v_\star^f(\rho) - \mathbb{E} \left[v_{\theta_t}^f(\rho) \right] - \eta \left(1 - \frac{3\eta L_f}{2} \right) \left\| \nabla v_{\theta_t}^f(\rho) \right\|_2^2 \\ &\quad + \left(\eta + \frac{3\eta^2 L_f}{2} \right) \beta_f^2 + \frac{3\eta^2 L_f \sigma_f^2}{2B}. \end{aligned}$$

Next using Theorem C.1 combined with (69) yields

$$\begin{aligned} v_\star^f(\rho) - \mathbb{E} \left[v_{\theta_{t+1}}^f(\rho) \right] &\leq \left(1 - \eta \underline{\mu}_f \left(1 - \frac{3\eta L_f}{2} \right) \right) \left(v_\star^f(\rho) - \mathbb{E} \left[v_{\theta_t}^f(\rho) \right] \right) \\ &\quad + 2\eta \beta_f^2 + \frac{3\eta^2 L_f \sigma_f^2}{2B}, \end{aligned}$$

where we used $3\eta L_f/2 \leq 1$. Finally, using that $\eta \leq 1/2L_f$ to bound $1 - 3\eta L_f/2$ and unrolling the recursion concludes the proof. \square

Next, we provide the sample complexity of f -PG for solving the f -regularized objective.

Corollary E.7. Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_ρ . Fix any $\epsilon > 0$ and $\lambda > 0$. Setting

$$H \geq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log \left(\frac{216\omega_f^2}{\epsilon \underline{\mu}_f (1-\gamma)^4} [4 + 4\lambda^2 d_f^2 + \lambda^2 (1-\gamma)^2 y_f^2] \right), \quad (73)$$

and

$$\eta \leq \min \left(\frac{1}{2L_f}, \frac{\epsilon B \underline{\mu}_f}{18\sigma_f^2} \right), \quad (74)$$

and

$$T \geq \frac{4}{\underline{\mu}_f} \max \left(2L_f, \frac{18\sigma_f^2}{\epsilon B \underline{\mu}_f} \right) \cdot \log \left(\frac{3(v_\star^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon} \right), \quad (75)$$

guarantees that

$$v_\star^f(\rho) - \mathbb{E} \left[v_{\theta_t}^f(\rho) \right] \leq \epsilon.$$

Proof. As $\eta \leq 1/2L_f$, then by using Theorem E.6, it holds that

$$v_\star^f(\rho) - \mathbb{E} \left[v_{\theta_T}^f(\rho) \right] \leq \underbrace{(1 - \underline{\mu}_f \eta / 4)^T (v_\star^f(\rho) - v_{\theta_0}^f(\rho))}_{(\mathbf{U})} + \underbrace{\frac{6\eta\sigma_f^2}{B\underline{\mu}_f}}_{(\mathbf{V})} + \underbrace{\frac{6\beta_f^2}{\underline{\mu}_f}}_{(\mathbf{W})}.$$

Next, we aim to show that under our conditions on T , H , and η , each of these terms is smaller than $\epsilon/3$.

Bounding V. We start with the term \mathbf{V} , which gives a condition on the step-size. In particular, setting

$$\eta \leq \frac{\epsilon B \underline{\mu}_f}{18\sigma_f^2},$$

guarantees that $\mathbf{V} \leq \epsilon/3$, which, together with $\eta \leq 1/(2L_f)$, gives the condition (74).

Bounding U. In order to ensure that \mathbf{U} is smaller than $\epsilon/3$, we need T to satisfy

$$T \geq \frac{4}{\eta \underline{\mu}_f} \log \left(\frac{\epsilon}{3 \left(v_{\star}^f(\rho) - v_{\theta_0}^f(\rho) \right)} \right),$$

which, combined with the inequality $\log(1+x) \leq x$ for $x > -1$, ensures that \mathbf{U} is smaller than $\epsilon/3$, and the condition (75) follows from (74).

Bounding W. Using Lemma E.3 and Jensen's inequality, it holds that

$$\mathbf{W} \leq \frac{6}{\underline{\mu}_f} \cdot \frac{4\gamma^{2H}(H+1)^2}{(1-\gamma)^4} \omega_f^2 [12 + 12\lambda^2 d_f^2 + 3\lambda^2(1-\gamma)^2 y_f^2],$$

Next, we remark that for any $a > 0$, we have $a\gamma^{2H}(H+1)^2 \leq \epsilon/3$ for

$$H \geq \frac{1}{1-\gamma} \max \left(\frac{4}{1-\gamma}, \log \left(\frac{3a}{\epsilon} \right) \right).$$

Taking $a = \frac{6}{\underline{\mu}_f} \cdot \frac{4}{(1-\gamma)^4} \omega_f^2 [12 + 12\lambda^2 d_f^2 + 3\lambda^2(1-\gamma)^2 y_f^2]$ gives $\mathbf{W} < \epsilon/3$, provided that (73) holds. \square

E.3 GUARANTEES ON THE NON-REGULARIZED PROBLEM

A key criterion for evaluating the quality of a reinforcement learning algorithm is its sample efficiency in solving the original, unregularized objective. To this end, we recall a result from Geist et al. (2019), which establishes a connection between regularized and unregularized value functions, and further characterizes the performance of the optimal f -regularized policy when evaluated in the original unregularized MDP.

Lemma E.8 (Proposition 3 and Theorem 2 of Geist et al. (2019)). *For any policy π , and state $s \in \mathcal{S}$ it holds that*

$$|v_{\pi}^f(s) - v_{\pi}(s)| \leq \frac{\lambda d_f}{1-\gamma}.$$

Additionally, denote by π_{\star}^f the optimal regularized policy. For any state $s \in \mathcal{S}$, It holds that

$$v_{\star}^f(s) - v_{\pi_{\star}^f}(s) \leq \frac{\lambda d_f}{1-\gamma}.$$

The following theorem gives the convergence rate for the non-regularised problem.

Theorem E.9. *Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_{ρ} . Fix $\eta \leq 1/2L_f$, a given temperature λ , and consider the iterates $(\theta)_{t=0}^{\infty}$ of the algorithm f -PG. For any $t \geq 0$ we have that*

$$\mathbb{E}[v_{\star}(\rho) - v_{\theta_t}(\rho)] \leq (1 - \underline{\mu}_f \eta / 4)^t (v_{\star}^f(\rho) - v_{\theta_0}^f(\rho)) + \frac{6\eta\sigma_f^2}{B\underline{\mu}_f} + \frac{6\beta_f^2}{\underline{\mu}_f} + \frac{2\lambda d_f}{1-\gamma}.$$

Proof. The proof holds from Theorem E.6 and Lemma E.8. \square

Finally, we give the sample complexity of f -PG to solve the unregularised problem.

Corollary E.10. Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_ρ . Consider any constant $c_f > 0$ such that

$$c_f \leq \min\left(\frac{1}{d_f}, \frac{1}{y_f}, 1\right).$$

Fix any $(1 - \gamma)^{-1} > \epsilon > 0$ such that

$$\epsilon < \frac{16}{(1 - \gamma)^3 \rho_{\min}} \min\left(\frac{4}{|f'(\iota_f)|}, \frac{1}{|f'(\frac{1}{2})|}, \frac{4}{|f'(\frac{1}{2}\pi_{\text{ref}})|}\right), \text{ and set } \lambda = \frac{(1 - \gamma)\epsilon}{4} c_f. \quad (76)$$

Define

$$d(\epsilon) = (f^*)'' \left(\frac{-96}{\epsilon c_f (1 - \gamma)^3 \rho_{\min}}\right)^2, \quad \ell(\epsilon) = \log\left(\frac{6(v_\star^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon}\right). \quad (77)$$

Additionally, define the three following constants which depend only on f as

$$C_f^{(1)} = \frac{16\omega_f^2}{c_f \zeta_f^2}, \quad C_f^{(2)} = 48\omega_f(\gamma\omega_f + (1 - \gamma)\kappa_f), \quad C_f^{(3)} = \frac{3456\omega_f^5}{\zeta_f^2 c_f}. \quad (78)$$

Setting

$$H \geq \frac{4}{(1 - \gamma)^2} + \frac{1}{1 - \gamma} \log\left(\frac{297\omega_f^2 C_f^{(1)}}{\epsilon d(\epsilon)(1 - \gamma)^6 \rho_{\min}^2 \pi_{\text{ref}}^2}\right),$$

and

$$\eta \leq \min\left(\frac{(1 - \gamma)^3}{C_f^{(2)}}, \frac{\epsilon^2 d(\epsilon)(1 - \gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}{C_f^{(3)}}\right), \quad (79)$$

and

$$T \geq \frac{16C_f^{(1)} \ell(\epsilon)}{\epsilon d(\epsilon)(1 - \gamma)^2 \rho_{\min}^2 \pi_{\text{ref}}^2} \max\left(\frac{C_f^{(2)}}{(1 - \gamma)^3}, \frac{C_f^{(3)}}{\epsilon^2 d(\epsilon)(1 - \gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}\right), \quad (80)$$

guarantees that

$$v_\star(\rho) - \mathbb{E}[v_{\theta_t}(\rho)] \leq \epsilon.$$

Proof. First, note that the conditions (76) on ϵ and λ guarantee that

$$\lambda \leq \frac{4}{(1 - \gamma)^2 \rho_{\min}} \min\left(\frac{4}{|f'(\iota_f)|}, \frac{1}{|f'(\frac{1}{2})|}, \frac{4}{|f'(\frac{1}{2}\pi_{\text{ref}})|}\right),$$

Thus, using Lemma E.5, and the fact that $\epsilon < (1 - \gamma)^{-1}$, we have that

$$\underline{\mu}_f \geq \frac{\lambda(1 - \gamma)\rho_{\min}^2 \zeta_f^2}{\omega_f^2} \pi_{\text{ref}}^2 (f^*)'' \left(\frac{-24}{\lambda(1 - \gamma)^2 \rho_{\min}}\right)^2.$$

Plugging in the expression of λ from (76) yields the following simplified expression

$$\underline{\mu}_f \geq \frac{\epsilon c_f (1 - \gamma)^2 \rho_{\min}^2 \zeta_f^2}{4\omega_f^2} \pi_{\text{ref}}^2 (f^*)'' \left(\frac{-96}{\epsilon c_f (1 - \gamma)^3 \rho_{\min}}\right)^2 = \frac{4\epsilon(1 - \gamma)^2 \rho_{\min}^2 \pi_{\text{ref}}^2}{C_f^{(1)}} d(\epsilon), \quad (81)$$

where $d(\epsilon)$ and $C_f^{(1)}$ are defined respectively in (77) and (78). Additionally, combining Lemma E.3, Lemma E.4 and Theorem B.12 and the expression of λ yields

$$\sigma_f^2 \leq \frac{24}{(1 - \gamma)^4} \omega_f^3, \quad \beta_f \leq \frac{6\gamma^H H}{(1 - \gamma)^2} \omega_f, \quad L_f = \frac{13\omega_f(\omega_f + (1 - \gamma)\kappa_f)}{(1 - \gamma)^3}, \quad (82)$$

where we used that under $\mathbf{A}_f(\pi_{\text{ref}})$, we have $\omega_f \geq 1$. Using Theorem E.9 and the fact that $\lambda \leq (1 - \gamma)\epsilon/4d_f$, we have that

$$v_\star(\rho) - \mathbb{E}[v_{\theta_t}(\rho)] \leq \underbrace{(1 - \frac{\underline{\mu}_f \eta}{4})^t (v_\star^f(\rho) - v_{\theta_0}^f(\rho))}_{(\mathbf{U}')} + \underbrace{\frac{6\eta\sigma_f^2}{B\underline{\mu}_f}}_{(\mathbf{V}')} + \underbrace{\frac{6\beta_f^2}{\underline{\mu}_f}}_{(\mathbf{W}')} + \frac{\epsilon}{2}.$$

Next, we aim to show that under our conditions on T , H , and η , each of these terms is smaller than $\epsilon/6$.

Bounding \mathbf{V}' . We start with the term \mathbf{V}' which gives a condition on the step-size. Using (81) and (82), it holds that

$$\begin{aligned} \mathbf{V}' &\leq \frac{144\eta\omega_f^3}{(1-\gamma)^4 B \underline{\mu}_f} \leq \frac{576\eta\omega_f^5}{(1-\gamma)^6 B} \frac{1}{\epsilon c_f \rho_{\min}^2 \zeta_f^2 \pi_{\text{ref}}^2 (f^*)'' \left(\frac{-96}{\epsilon c_f (1-\gamma)^3 \rho_{\min}} \right)^2} \\ &= \frac{C_f^{(3)} \eta}{6\epsilon d(\epsilon) (1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}, \end{aligned}$$

where $C_f^{(3)}$ is defined in (78). In particular, setting

$$\eta \leq \frac{\epsilon^2 d(\epsilon) (1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}{C_f^{(3)}},$$

guarantees that $\mathbf{V}' \leq \epsilon/6$, which together with the condition $\eta \leq 1/(2L_f)$, gives the condition (79).

Bounding \mathbf{U}' . In order to ensure that \mathbf{U}' is smaller than $\epsilon/6$, we need T to satisfy

$$T \geq \frac{1}{\log(1 - \eta \underline{\mu}_f / 4)} \log \left(\frac{\epsilon}{6 (v_*^f(\rho) - v_{\theta_0}^f(\rho))} \right),$$

which combined with the inequality $\log(1+x) \leq x$ for $x > -1$, ensures that $\mathbf{V}' < \epsilon/6$ and the condition (80) follows from (79).

Bounding \mathbf{W}' . Using (82), it holds that

$$\mathbf{W}' \leq \frac{6}{\underline{\mu}_f} \cdot \frac{36\gamma^{2H}(H+1)^2}{(1-\gamma)^4} \omega_f^2.$$

Next, using that for any $a > 0$, we have $a\gamma^{2H}(H+1)^2 \leq \epsilon/6$ for

$$H \geq \frac{1}{1-\gamma} \max \left(\frac{4}{1-\gamma}, \log \left(\frac{6a}{\epsilon} \right) \right),$$

shows that under our condition on H , we have $\mathbf{W}' < \epsilon/6$. \square

F APPLICATION TO COMMON f -DIVERGENCES

In this section, we apply the results of the preceding section to two commonly used f -divergences, which are Kullback-Leibler and α -Tsallis.

F.1 KULLBACK-LEIBLER

Lemma F.1. Assume that, for some $\pi_{\text{ref}} > 0$, π_{ref} satisfy $\mathbf{P}(\pi_{\text{ref}})$. The function f defined by $f(u) = u \log(u)$ satisfies $\mathbf{A}_f(\pi_{\text{ref}})$, with

$$\omega_f = 1, \quad \kappa_f = 1, \quad \iota_f = 1.$$

Additionally, under the condition that

$$\lambda \leq \frac{4}{(1-\gamma)^2 \rho_{\min} (\log(2/\pi_{\text{ref}}) + 1)}, \quad (83)$$

we have

$$\begin{aligned} \zeta_f &= 1, \quad d_f \leq |\log(\pi_{\text{ref}})|, \quad y_f \leq 1 + 2|\log(\pi_{\text{ref}})|, \quad L_f = \frac{8}{(1-\gamma)^3} + 4\lambda \frac{3+4|\log(\pi_{\text{ref}})|}{(1-\gamma)^3}, \\ \beta_f &= \frac{2\gamma^H(H+1)}{(1-\gamma)^2} [2 + \lambda + 4\lambda|\log(\pi_{\text{ref}})|], \quad \sigma_f^2 \leq \frac{12}{(1-\gamma)^4} [1 + \lambda^2 (2 + 5|\log(\pi_{\text{ref}})|^2)], \\ \underline{\mu}_f &\geq \lambda(1-\gamma) \rho_{\min}^2 \pi_{\text{ref}}^2 \exp \left(-\frac{32 + 16\gamma\lambda|\log(\pi_{\text{ref}})|}{\lambda(1-\gamma)^2 \rho_{\min}} \right) / 9. \end{aligned}$$

Proof. Firstly, note that we have

$$f(u) = u \log(u) , \quad f'(u) = \log(u) + 1 , \quad f''(u) = u^{-1} , \quad f'''(u) = -u^{-2} .$$

Satisfying $\mathbf{A}_f(\pi_{\text{ref}})$. Observe that (i) and (ii) are immediately valid from the expression above of the derivatives of f . Moreover, we have

$$1/uf''(u) = 1 , \quad \frac{|f'''(u)|}{f''(u)^2} = 1$$

showing that (iii) of $\mathbf{A}_f(\pi_{\text{ref}})$, is satisfied with $\omega_f = \kappa_f = 1$. Finally, as f'' is a strictly decreasing function on \mathbb{R}_+ then (iv) is valid with $\iota_f = 1$.

Bounding the constants. Next, we bound sequentially each of the constants that appear in the statement of the lemma. For any $s \in \mathcal{S}$ and $\nu \in \mathcal{P}(\mathcal{A})$, we have that

$$\sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right)} = \sum_{a \in \mathcal{A}} \nu(a) = 1 ,$$

Thus using (7), we have that $\zeta_f = 1$. It holds that

$$\begin{aligned} d_f &= \max_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s) f\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) \\ &= \max_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} \sum_{a \in \mathcal{A}} \nu(a) \log\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) \\ &\leq \max_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} \sum_{a \in \mathcal{A}} \nu(a) \log(\nu(a)) - \nu(a) \log(\pi_{\text{ref}}) , \end{aligned}$$

which gives $d_f \leq -\log(\pi_{\text{ref}})$. Next, we have

$$\begin{aligned} y_f &= \max_{(s, \nu) \in \mathcal{S} \times \mathcal{P}(\mathcal{A})} \sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right)} \left| f'\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) \right| \\ &= \max_{(s, \nu) \in \mathcal{P}(\mathcal{A})} \sum_{a \in \mathcal{A}} \nu(a) \left| \log\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) + 1 \right| \\ &= 1 - \log(\pi_{\text{ref}}) + \max_{\nu \in \mathcal{P}(\mathcal{A})} - \sum_{a \in \mathcal{A}} \nu(a) \log(\nu(a)) \\ &\leq 1 + 2|\log(\pi_{\text{ref}})| , \end{aligned}$$

where in the last inequality, we used that the entropy of a distribution on \mathcal{A} is bounded by $\log(|\mathcal{A}|)$ and the fact that $\pi_{\text{ref}} \leq 1/|\mathcal{A}|$. Next, using Theorem B.12 and the bounds above, we have

$$\begin{aligned} L_f &= \frac{8\omega_f(\gamma\omega_f + (1-\gamma)\kappa_f)}{(1-\gamma)^3} + 4\lambda \frac{2\gamma^2\omega_f^2 d_f + 2\gamma(1-\gamma)\omega_f[\kappa_f d_f + y_f] + (1-\gamma)^2[\omega_f + 2\kappa_f y_f]}{(1-\gamma)^3} \\ &\leq \frac{8}{(1-\gamma)^3} + 4\lambda \frac{3 + 4|\log(\pi_{\text{ref}})|}{(1-\gamma)^3} . \end{aligned}$$

Using Lemma E.3 gives

$$\beta_f = \frac{2\gamma^H(H+1)}{(1-\gamma)^2} \omega_f [2 + 2\lambda d_f + \lambda(1-\gamma)y_f] \leq \frac{2\gamma^H(H+1)}{(1-\gamma)^2} [2 + \lambda + 4\lambda|\log(\pi_{\text{ref}})|] .$$

Using Lemma E.4 gives

$$\sigma_f^2 = \frac{12}{(1-\gamma)^4} [\omega_f^3 + \lambda^2 \gamma^2 \omega_f^3 d_f^2 + \lambda^2 (1-\gamma)^2 \omega_f^2 y_f^2] \leq \frac{12}{(1-\gamma)^4} [1 + \lambda^2 (2 + 5|\log(\pi_{\text{ref}})|^2)] .$$

Next, note that (83), guarantees that we have

$$\lambda \leq \frac{4}{(1-\gamma)^2 \rho_{\min}} \min \left(\frac{4}{|f'(\iota_f)|}, \frac{1}{|f'(\frac{1}{2})|}, \frac{4}{|f'(\frac{1}{2}\pi_{\text{ref}})|} \right)$$

Thus, using Lemma E.5, we have

$$\begin{aligned}\underline{\mu}_f &= \frac{\lambda(1-\gamma)\rho_{\min}^2 \zeta_f^2}{\omega_f^2} \underline{\pi}_{\text{ref}}^2 (f^*)'' \left(-\frac{16+8\gamma\lambda d_f}{\lambda(1-\gamma)^2 \rho_{\min}} \right)^2 \\ &\geq \lambda(1-\gamma)\rho_{\min}^2 \underline{\pi}_{\text{ref}}^2 \exp\left(-\frac{32+16\gamma\lambda|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2 \rho_{\min}}\right) / 9 ,\end{aligned}$$

where in the last equality, we used that the convex conjugate of $f(u) = u \log(u)$ is $f^*(y) = \exp(y-1)$ and that $\exp(-2) \geq 1/9$. \square

In the next two corollaries, we apply Corollary E.7 and Corollary E.10 to get more explicitly the sample complexity of f -PG with entropy regularization.

Corollary F.2. *Assume that, for some $1/4 \geq \underline{\pi}_{\text{ref}} > 0$, $\underline{\pi}_{\text{ref}}$ satisfy $\mathbf{P}(\underline{\pi}_{\text{ref}})$. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_ρ . Fix any $(1-\gamma)^{-1} \geq \epsilon > 0$, $\lambda > 0$ and B such that*

$$\lambda \leq \min\left(\frac{4}{(1-\gamma)^2 \rho_{\min} (\log(2/\underline{\pi}_{\text{ref}}) + 1)}, 1\right), \quad B \leq \frac{216|\log(\underline{\pi}_{\text{ref}})| \exp\left(\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2 \rho_{\min}}\right)}{\epsilon \lambda (1-\gamma)^2 \underline{\pi}_{\text{ref}}^2 \rho_{\min}^2}.$$

Setting

$$H \geq \frac{1}{1-\gamma} \log\left(\frac{29160 \log(\underline{\pi}_{\text{ref}})^2}{\epsilon \lambda (1-\gamma)^5 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}\right) + \frac{52|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^3 \rho_{\min}}, \quad (84)$$

and

$$\eta \leq \frac{\epsilon B \lambda (1-\gamma)^5 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}{15552|\log(\underline{\pi}_{\text{ref}})|^2} \exp\left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2 \rho_{\min}}\right),$$

and

$$T \geq \frac{559872|\log(\underline{\pi}_{\text{ref}})|^2}{\lambda^2 \epsilon B (1-\gamma)^6 \rho_{\min}^4 \underline{\pi}_{\text{ref}}^4} \exp\left(\frac{96|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2 \rho_{\min}}\right),$$

guarantees that

$$v_\star^f(\rho) - \mathbb{E}\left[v_{\theta_t}^f(\rho)\right] \leq \epsilon ,$$

where f is the Kullback-Leibler divergence generator. Thus, the sample complexity of f -PG to learn an ϵ -solution of the entropy regularized problem is

$$TBH \approx \frac{1}{\lambda^3 \epsilon B} \frac{|\log(\underline{\pi}_{\text{ref}})|^3}{(1-\gamma)^9 \rho_{\min}^5 \underline{\pi}_{\text{ref}}^4} \exp\left(\frac{|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2 \rho_{\min}}\right)$$

Proof. To prove this corollary, we show that the assumptions of Corollary E.7 hold. First, $\mathbf{A}_f(\underline{\pi}_{\text{ref}})$ holds as a consequence of Lemma F.1. Then, we show that the condition (73) in Corollary E.7 holds, that is $H \leq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log\left(\frac{216\omega_f^2}{\epsilon \underline{\mu}_f (1-\gamma)^4} \left[4 + 4\lambda^2 d_f^2 + \lambda^2(1-\gamma)^2 y_f^2\right]\right)$. To this end, we remark that

$$\begin{aligned}&\frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log\left(\frac{216\omega_f^2}{\epsilon \underline{\mu}_f (1-\gamma)^4} \left[4 + 4\lambda^2 d_f^2 + \lambda^2(1-\gamma)^2 y_f^2\right]\right) \\ &\leq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log\left(\frac{216}{\epsilon \underline{\mu}_f (1-\gamma)^4} \left[4 + \lambda^2(5 + 6 \log(\underline{\pi}_{\text{ref}})^2)\right]\right) \\ &\leq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log\left(\frac{1944 \left[4 + \lambda^2(5 + 6 \log(\underline{\pi}_{\text{ref}})^2)\right]}{\epsilon \lambda (1-\gamma)^5 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}\right) + \frac{32 + 16\gamma\lambda |\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^3 \rho_{\min}} \\ &\leq \frac{1}{1-\gamma} \log\left(\frac{29160 \log(\underline{\pi}_{\text{ref}})^2}{\epsilon \lambda (1-\gamma)^5 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}\right) + \frac{52|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^3 \rho_{\min}} \leq H ,\end{aligned}$$

where we used the lower bound on μ_f provided in Lemma F.1 in the second inequality, as well as $\lambda < 1$ and $\underline{\pi}_{\text{ref}} \leq 1/4$ in the last two inequalities. Furthermore, Lemma F.1 with $\lambda \leq 1$ gives

$$L_f \leq \frac{36|\log(\underline{\pi}_{\text{ref}})|}{(1-\gamma)^3}, \quad (85)$$

$$\sigma_f^2 \leq \frac{96|\log(\underline{\pi}_{\text{ref}})|^2}{(1-\gamma)^4}, \quad (86)$$

$$\underline{\mu}_f \geq \frac{\lambda(1-\gamma)\rho_{\min}^2\underline{\pi}_{\text{ref}}^2}{9} \exp\left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2\rho_{\min}}\right). \quad (87)$$

Using these three bounds on smoothness, variance and Polyak-Łojasiewicz coefficients, we obtain

$$\begin{aligned} \min\left(\frac{1}{2L_f}, \frac{\epsilon B \underline{\mu}_f}{18\sigma_f^2}\right) &\geq \min\left(\frac{(1-\gamma)^3}{72|\log(\underline{\pi}_{\text{ref}})|}, \frac{\epsilon B \lambda(1-\gamma)^5 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2 \exp\left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2\rho_{\min}}\right)}{15552|\log(\underline{\pi}_{\text{ref}})|^2}\right) \\ &= \frac{\epsilon B \lambda(1-\gamma)^5 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}{15552|\log(\underline{\pi}_{\text{ref}})|^2} \exp\left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2\rho_{\min}}\right), \end{aligned} \quad (88)$$

where in the last identity, we used the fact that $\epsilon < (1-\gamma)^{-1}$ and that

$$B \leq \frac{216|\log(\underline{\pi}_{\text{ref}})|}{\epsilon \lambda(1-\gamma)^2 \underline{\pi}_{\text{ref}}^2 \rho_{\min}^2} \exp\left(\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2\rho_{\min}}\right).$$

This shows that our condition on η guarantees that the one set in Corollary E.7 is satisfied. Finally using (88) and (87), we have

$$\begin{aligned} \frac{4}{\underline{\mu}_f} \max\left(2L_f, \frac{18\sigma_f^2}{\epsilon B \underline{\mu}_f}\right) &\leq \frac{4}{\underline{\mu}_f} \frac{15552|\log(\underline{\pi}_{\text{ref}})|^2}{\epsilon B \lambda(1-\gamma)^5 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2} \exp\left(\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2\rho_{\min}}\right) \\ &\leq \frac{559872|\log(\underline{\pi}_{\text{ref}})|^2}{\lambda^2 \epsilon B (1-\gamma)^6 \rho_{\min}^4 \underline{\pi}_{\text{ref}}^4} \exp\left(\frac{96|\log(\underline{\pi}_{\text{ref}})|}{\lambda(1-\gamma)^2\rho_{\min}}\right), \end{aligned}$$

which concludes the proof. \square

Corollary F.3. Assume that, for some $\underline{\pi}_{\text{ref}} \in (0, 1/4]$, $\underline{\pi}_{\text{ref}}$ satisfy $\mathbf{P}(\underline{\pi}_{\text{ref}})$. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_ρ and fix f to be the Kullback-Leibler divergence generator, i.e. $f(u) = u \log(u)$. Fix any $\epsilon \in (0, (1-\gamma)^{-1}]$, such that

$$\epsilon < \frac{16}{(1-\gamma)^3 \rho_{\min} (\log(2/\underline{\pi}_{\text{ref}}) + 1)}, \quad \text{and set } \lambda = \frac{(1-\gamma)\epsilon}{12 \log(|\underline{\pi}_{\text{ref}}|)}. \quad (89)$$

Additionally set any B such that

$$B \leq \frac{1}{\epsilon^2 (1-\gamma)^3 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2} \exp\left(\frac{576|\log(\underline{\pi}_{\text{ref}})|}{\epsilon(1-\gamma)^3 \rho_{\min}}\right). \quad (90)$$

Setting

$$H \geq \frac{1}{1-\gamma} \log\left(\frac{|\log(\underline{\pi}_{\text{ref}})|}{\epsilon(1-\gamma)^6 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}\right) + \frac{586|\log(\underline{\pi}_{\text{ref}})|}{\epsilon(1-\gamma)^4 \rho_{\min}},$$

and

$$\eta \leq \frac{\epsilon^2 (1-\gamma)^6 B \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}{93312|\log(\underline{\pi}_{\text{ref}})|} \exp\left(\frac{-576|\log(\underline{\pi}_{\text{ref}})|}{\epsilon(1-\gamma)^3 \rho_{\min}}\right), \quad (91)$$

and

$$T \geq \frac{644972544|\log(\underline{\pi}_{\text{ref}})|^2}{\epsilon^3 (1-\gamma)^8 \rho_{\min}^4 \underline{\pi}_{\text{ref}}^4 B} \exp\left(\frac{1152|\log(\underline{\pi}_{\text{ref}})|}{\epsilon(1-\gamma)^3 \rho_{\min}}\right) \log\left(\frac{6(v_\star^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon}\right), \quad (92)$$

guarantees that

$$v_*(\rho) - \mathbb{E}[v_{\theta_t}(\rho)] \leq \epsilon .$$

Thus, the sample complexity of f -PG, where f is the Kullback-Leibler divergence generator, to learn an ϵ -solution of the non-regularized problem is

$$TBH \approx \frac{|\log(\pi_{\text{ref}})|^3}{\epsilon^4(1-\gamma)^{12}\rho_{\min}^5\pi_{\text{ref}}^4} \exp\left(\frac{|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^3\rho_{\min}}\right)$$

Proof. This result follows from Corollary E.10, whose assumptions we check now. First, note that (89) implies that

$$\epsilon < \frac{16}{(1-\gamma)^3\rho_{\min}} \min\left(\frac{4}{|f'(\ell_f)|}, \frac{1}{|f'(\frac{1}{2})|}, \frac{4}{|f'(\frac{1}{2}\pi_{\text{ref}})|}\right) .$$

Additionally, we can rewrite the constants from Corollary E.10 using Lemma F.1, which gives

$$\begin{aligned} C_f^{(1)} &\leq 48|\log(\pi_{\text{ref}})|, & C_f^{(2)} &\leq 48, & C_f^{(3)} &= 10368|\log(\pi_{\text{ref}})|, \\ d(\epsilon) &\geq \exp\left(\frac{-576|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^3\rho_{\min}}\right)/9. \end{aligned} \quad (93)$$

where $C_f^{(1)}, C_f^{(2)}, C_f^{(3)}$, and $d(\epsilon)$ are defined in (78) and (77). Next, the condition on H in Corollary E.10 holds since

$$\begin{aligned} &\frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log\left(\frac{297\omega_f^2 C_1^f}{\epsilon d(\epsilon)(1-\gamma)^6 \rho_{\min}^2 \pi_{\text{ref}}^2}\right) \\ &\leq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log\left(\frac{128304|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^6 \rho_{\min}^2 \pi_{\text{ref}}^2}\right) + \frac{576|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^4 \rho_{\min}} \\ &\leq \frac{1}{1-\gamma} \log\left(\frac{|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^6 \rho_{\min}^2 \pi_{\text{ref}}^2}\right) + \frac{586|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^4 \rho_{\min}} \leq H , \end{aligned}$$

where in the second to last inequality, we used that $\pi_{\text{ref}} < 1/4$ and that $\log(128304) \leq 6$. Using (93), we have

$$\begin{aligned} &\min\left(\frac{(1-\gamma)^3}{C_f^{(2)}}, \frac{\epsilon^2 d(\epsilon)(1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}{C_f^{(3)}}\right) \\ &\geq \min\left(\frac{(1-\gamma)^3}{48}, \frac{\epsilon^2(1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}{93312|\log(\pi_{\text{ref}})|} \exp\left(\frac{-576|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^3\rho_{\min}}\right)\right) \\ &\geq \frac{\epsilon^2(1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}{93312|\log(\pi_{\text{ref}})|} \exp\left(\frac{-576|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^3\rho_{\min}}\right) , \end{aligned} \quad (94)$$

where in the last inequality, we used the condition on B introduced in (90). Hence our condition on the step-size ensures that the one assumed in Corollary E.10 is satisfied. Next, using (93) and (94) yields

$$\begin{aligned} &\frac{16C_f^{(1)}\ell(\epsilon)}{\epsilon d(\epsilon)(1-\gamma)^2 \rho_{\min}^2 \pi_{\text{ref}}^2} \max\left(\frac{C_f^{(2)}}{(1-\gamma)^3}, \frac{C_f^{(3)}}{\epsilon^2 d(\epsilon)(1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}\right) \\ &\leq \frac{16C_f^{(1)}\ell(\epsilon)}{\epsilon d(\epsilon)(1-\gamma)^2 \rho_{\min}^2 \pi_{\text{ref}}^2} \cdot \frac{93312|\log(\pi_{\text{ref}})|}{\epsilon^2(1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2} \exp\left(\frac{576|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^3\rho_{\min}}\right) \\ &\leq \frac{6912|\log(\pi_{\text{ref}})|\ell(\epsilon)}{\epsilon(1-\gamma)^2 \rho_{\min}^2 \pi_{\text{ref}}^2} \cdot \frac{93312|\log(\pi_{\text{ref}})|}{\epsilon^2(1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2} \exp\left(\frac{1152|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^3\rho_{\min}}\right) \\ &\leq \frac{644972544|\log(\pi_{\text{ref}})|^2}{\epsilon^3(1-\gamma)^8 \rho_{\min}^4 \pi_{\text{ref}}^4 B} \exp\left(\frac{1152|\log(\pi_{\text{ref}})|}{\epsilon(1-\gamma)^3\rho_{\min}}\right) \log\left(\frac{6(v_*^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon}\right) , \end{aligned}$$

which proves that under our condition on T the one assumed by Corollary E.10 is satisfied. \square

F.2 α -TSALLIS

Lemma F.4. Assume that, for some $\pi_{\text{ref}} \in (0, 1/4]$, π_{ref} satisfy $\mathbf{P}(\pi_{\text{ref}})$. For any $\alpha \in (0; 1)$, the function f_α defined by

$$f_\alpha(u) = \frac{u^\alpha - \alpha u + \alpha - 1}{\alpha(\alpha - 1)},$$

satisfies $\mathbf{A}_f(\pi_{\text{ref}})$, with

$$\omega_f = \pi_{\text{ref}}^{\alpha-1}, \quad \kappa_f = 2\pi_{\text{ref}}^{\alpha-1}, \quad \iota_f = 1.$$

Additionally, under the condition that

$$\lambda \leq \frac{4}{(1-\gamma)^2 \rho_{\min}} \cdot \frac{1-\alpha}{(\pi_{\text{ref}}/2)^{\alpha-1} - 1}, \quad (95)$$

we have

$$\begin{aligned} \zeta_f &= 1, \quad d_f \leq \frac{4|\log(\pi_{\text{ref}})|}{\alpha^2}, \quad y_f \leq 4|\log(\pi_{\text{ref}})|, \quad L_f = \frac{16\pi_{\text{ref}}^{\alpha-1}}{(1-\gamma)^3} + 180\lambda \frac{\pi_{\text{ref}}^{2\alpha-2} |\log(\pi_{\text{ref}})|}{\alpha^2(1-\gamma)^3}, \\ \beta_f &= \frac{4\gamma^H(H+1)}{(1-\gamma)^2} \left[1 + 6\lambda \frac{|\log(\pi_{\text{ref}})|}{\alpha^2} \right], \quad \sigma_f^2 \leq \frac{12\pi_{\text{ref}}^{3\alpha-3}}{(1-\gamma)^4} \left[1 + \frac{16\lambda^2}{\alpha^4} \log(\pi_{\text{ref}}|^2) \right], \\ \underline{\mu}_f &\geq \lambda(1-\gamma)\pi_{\text{ref}}^{2-2\alpha} \rho_{\min}^2 \pi_{\text{ref}}^2 \exp_\alpha \left(-\frac{16+32\gamma\lambda|\log(\pi_{\text{ref}})|/\alpha^2}{\lambda(1-\gamma)^2 \rho_{\min}} \right)^{4-2\alpha}. \end{aligned}$$

Proof. Fix any $\alpha \in (0, 1)$ and set $f = f_\alpha$. Firstly, note that we have

$$f(u) = \frac{u^\alpha - \alpha u + \alpha - 1}{\alpha(\alpha - 1)}, \quad f'(u) = \frac{u^{\alpha-1} - 1}{\alpha - 1}, \quad f''(u) = u^{\alpha-2}, \quad f'''(u) = (\alpha - 2)u^{\alpha-3}.$$

Satisfying $\mathbf{A}_f(\pi_{\text{ref}})$. Observe that (i) and (ii) are immediately valid from the expression above of the derivatives of f . Moreover, we have

$$1/(uf_\alpha''(u)) = u^{1-\alpha}, \quad \frac{|f_\alpha'''(u)|}{f_\alpha''(u)^2} = |\alpha - 2|u^{1-\alpha},$$

showing that (iii) of $\mathbf{A}_f(\pi_{\text{ref}})$, is satisfied with $\omega_f = \pi_{\text{ref}}^{\alpha-1}$, and $\kappa_f = 2\pi_{\text{ref}}^{\alpha-1}$. Finally, as f'' is a strictly decreasing function on \mathbb{R}_+ then (iv) is valid with $\iota_f = 1$. Next, we bound sequentially each of the constants that appear in the statement of the lemma.

Bounding ζ_f . For any $s \in \mathcal{S}$ and $\nu \in \mathcal{P}(\mathcal{A})$, using Jensen's inequality we have that

$$\sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right)} = \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s) \left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \right)^{2-\alpha} \geq \left(\sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s) \frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \right)^{2-\alpha} = 1,$$

Thus using (7), we have that $\zeta_f = 1$.

Bounding d_f . For any state $s \in \mathcal{S}$ and $\nu \in \mathcal{P}(\mathcal{A})$, it holds that

$$D^f(\nu || \pi_{\text{ref}}(\cdot|s)) = \frac{1}{\alpha(\alpha-1)} \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s) \left[\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \right)^\alpha - \alpha \frac{\nu(a)}{\pi_{\text{ref}}(a|s)} - (1-\alpha) \right].$$

Next, for $y \in]0; \frac{1}{\pi_{\text{ref}}}]$, define the function $p_y: [\alpha; 1] \rightarrow \mathbb{R}$ which satisfies

$$p_y(\beta) = y^\beta - \beta y - (1-\beta), \quad p_y'(\beta) = \log(y)y^\beta - y + 1.$$

It holds that

$$\frac{y^\alpha - \alpha y - (1-\alpha)}{\alpha-1} = \frac{p_y(\alpha) - p_y(1)}{\alpha-1} \leq \sup_{\beta \in [\alpha, 1]} |p_y'(\beta)| \leq y+1+y|\log(y)|\mathbf{1}_{y \geq 1} + |\log(y)|y^\alpha \mathbf{1}_{y \leq 1}.$$

Applying the previous inequality with $y = \nu(a)/\pi_{\text{ref}}(a|s)$ yields

$$\begin{aligned} D^f(\nu \parallel \pi_{\text{ref}}(\cdot|s)) &\leq \frac{1}{\alpha} \sum_{a \in \mathcal{A}} \pi_{\text{ref}}(a|s) \left[1 + \frac{\nu(a)}{\pi_{\text{ref}}(a|s)} + \frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \left| \log\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) \right| \mathbf{1}_{\nu(a)/\pi_{\text{ref}}(a|s) \geq 1} \right. \\ &\quad \left. + \left| \log\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) \right| \left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \right)^\alpha \mathbf{1}_{\nu(a)/\pi_{\text{ref}}(a|s) \leq 1} \right] \\ &= \frac{1}{\alpha} \sum_{a \in \mathcal{A}} \nu(a) \log(\nu(a)) \mathbf{1}_{\nu(a)/\pi_{\text{ref}}(a|s) \geq 1} - \frac{1}{\alpha} \sum_{a \in \mathcal{A}} \nu(a) \log(\pi_{\text{ref}}(a|s)) \mathbf{1}_{\nu(a)/\pi_{\text{ref}}(a|s) \geq 1} \\ &\quad + \frac{1}{\alpha} \max_{a \in \mathcal{A}} \left| \log\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) \right| \left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \right)^\alpha \mathbf{1}_{\nu(a)/\pi_{\text{ref}}(a|s) \leq 1} + \frac{2}{\alpha} , \end{aligned}$$

where in the last equality, we used that if $x \geq 1$, then $\log(x) \geq 0$. Next using that $\log(\nu(a)) \leq 0$ yields

$$\begin{aligned} D^f(\nu \parallel \pi_{\text{ref}}(\cdot|s)) &\leq -\frac{1}{\alpha} \sum_{a \in \mathcal{A}} \nu(a) \log(\pi_{\text{ref}}(a|s)) + \frac{1}{\alpha} \max_{x \in [0;1]} |\log(x)| x^\alpha + \frac{2}{\alpha} , \\ &\leq \frac{1}{\alpha} |\log(\pi_{\text{ref}})| + \frac{1}{\alpha} \max_{x \in [0;1]} |\log(x)| x^\alpha + \frac{2}{\alpha} , \end{aligned}$$

where in the last inequality, we used that the entropy of a probability measure on \mathcal{A} is bounded by $|\log(\pi_{\text{ref}})|$. Next using that $\max_{x \in [0;1]} |\log(x)| x^\alpha \leq e^{-1}/\alpha$ combined with $\max(1, \log(|\mathcal{A}|)) \leq |\log(\pi_{\text{ref}})|$ gives

$$D^f(\nu \parallel \pi_{\text{ref}}(\cdot|s)) \leq \frac{3|\log(\pi_{\text{ref}})|}{\alpha} + \frac{1}{\alpha^2} \leq \frac{4|\log(\pi_{\text{ref}})|}{\alpha^2} .$$

Thus, it holds that

$$d_f \leq \frac{4|\log(\pi_{\text{ref}})|}{\alpha^2} .$$

Bounding y_f . For any policy $\nu \in \mathcal{P}(\mathcal{A})$ and $s \in \mathcal{S}$, we have

$$\sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right)} \left| f' \left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \right) \right| = \left| \frac{1}{1-\alpha} \sum_{a \in \mathcal{A}} \frac{\nu(a)^{2-\alpha}}{\pi_{\text{ref}}(a|s)^{1-\alpha}} \left| \left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \right)^{\alpha-1} - 1 \right| \right| .$$

Next, define the function g_y for $y \in]0; \frac{1}{\pi_{\text{ref}}}]$ which satisfies

$$g_y(\beta) = y^{\beta-1} , \quad g'_y(\beta) = \log(y) y^{\beta-1} .$$

It holds that

$$\left| \frac{y^{\alpha-1} - 1}{1-\alpha} \right| = \left| \frac{g_y(\alpha) - g_y(1)}{\alpha-1} \right| \leq \sup_{\beta \in [\alpha, 1]} |g'_y(\beta)| \leq |\log(y)| \mathbf{1}_{y \leq 1} + |\log(y) y^{\alpha-1}| \mathbf{1}_{y \geq 1} ,$$

Hence, applying the previous inequality with $y = \nu(a)/\pi_{\text{ref}}(a|s)$ gives

$$\begin{aligned} &\sum_{a \in \mathcal{A}} \frac{\pi_{\text{ref}}(a|s)}{f''\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right)} \left| f' \left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \right) \right| \\ &\leq \sum_{a \in \mathcal{A}} \frac{\nu(a)^{2-\alpha}}{\pi_{\text{ref}}(a|s)^{1-\alpha}} \left[\left| \log\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) \right| \mathbf{1}_{\nu(a) \leq \pi_{\text{ref}}(a|s)} + \left| \log\left(\frac{\pi(a|s)}{\pi_{\text{ref}}(a|s)}\right) \right| \left(\frac{\pi_{\text{ref}}(a|s)}{\nu(a)} \right)^{1-\alpha} \right] \\ &\leq \sum_{a \in \mathcal{A}} \left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)} \right)^{1-\alpha} \nu(a) \left| \log\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) \right| \mathbf{1}_{\nu(a) \leq \pi_{\text{ref}}(a|s)} + \sum_{a \in \mathcal{A}} \nu(a) \left| \log\left(\frac{\pi(a|s)}{\pi_{\text{ref}}(a|s)}\right) \right| \\ &\leq \sum_{a \in \mathcal{A}} \nu(a) \left| \log\left(\frac{\nu(a)}{\pi_{\text{ref}}(a|s)}\right) \right| + 2|\log(\pi_{\text{ref}})| , \end{aligned}$$

where in the last inequality, we used for the first term that for any $u \in \mathbb{R}$, $u \mathbf{1}_{u \leq 1} \leq 1$ that the entropy of a probability distribution on \mathcal{A} is bounded by $\log(|\mathcal{A}|)$, the fact that $\pi_{\text{ref}} \leq 1/|\mathcal{A}|$. Using the same argument again to bound the first term gives

$$y_f \leq 4|\log(\pi_{\text{ref}})| .$$

Bounding L_f . Next, using Theorem B.12 and the bound on the constants previously computed, we have

$$\begin{aligned} L_f &= \frac{8\omega_f(\gamma\omega_f + (1-\gamma)\kappa_f)}{(1-\gamma)^3} + 4\lambda \frac{2\gamma^2\omega_f^2 d_f + 2\gamma(1-\gamma)\omega_f[\kappa_f d_f + y_f] + (1-\gamma)^2[\omega_f + 2\kappa_f y_f]}{(1-\gamma)^3} \\ &\leq \frac{16\pi_{\text{ref}}^{\alpha-1}}{(1-\gamma)^3} + 180\lambda \frac{\pi_{\text{ref}}^{2\alpha-2} |\log(\pi_{\text{ref}})|}{\alpha^2(1-\gamma)^3}, \end{aligned}$$

where in the last inequality, we used that $\pi_{\text{ref}} < 1/4$.

Bounding β_f . Using Lemma E.3 gives

$$\beta_f = \frac{2\gamma^H(H+1)}{(1-\gamma)^2} \omega_f [2 + 2\lambda d_f + \lambda(1-\gamma)y_f] \leq \frac{4\gamma^H(H+1)}{(1-\gamma)^2} \left[1 + 6\lambda \frac{|\log(\pi_{\text{ref}})|}{\alpha^2} \right].$$

Bounding σ_f . Using Lemma E.4 gives

$$\sigma_f^2 = \frac{12}{(1-\gamma)^4} [\omega_f^3 + \lambda^2 \gamma^2 \omega_f^3 d_f^2 + \lambda^2 (1-\gamma)^2 \omega_f^2 y_f^2] \leq \frac{12\pi_{\text{ref}}^{3\alpha-3}}{(1-\gamma)^4} \left[1 + \frac{16\lambda^2}{\alpha^4} \log(\pi_{\text{ref}}|^2) \right].$$

Bounding μ_f . Next, note that as f'_α is an increasing function then $f'_\alpha(\pi_{\text{ref}}/2) \leq f'_\alpha(1/2) \leq f'_\alpha(t_f) = f'_\alpha(1) = 0$. Thus, we have $|f'_\alpha(t_f)| \leq |f'_\alpha(1/2)| \leq |f'_\alpha(\pi_{\text{ref}}/2)|$. This proves that (95), guarantees that

$$\lambda \leq \frac{4}{(1-\gamma)^2 \rho_{\min}} \min \left(\frac{4}{|f'_\alpha(t_f)|}, \frac{1}{|f'_\alpha(1/2)|}, \frac{4}{|f'_\alpha(\pi_{\text{ref}}/2)|} \right)$$

Thus, using Lemma E.5, we have

$$\underline{\mu}_f = \frac{\lambda(1-\gamma)\rho_{\min}^2 \zeta_f^2}{\omega_f^2} \pi_{\text{ref}}^2 (f^*)'' \left(-\frac{16 + 8\gamma\lambda d_f}{\lambda(1-\gamma)^2 \rho_{\min}} \right)^2. \quad (96)$$

Next, recall from proposition 8 of Roulet et al. (2025) that

$$f_\alpha^*(x) = \frac{(1 + (\alpha-1)x)^{\frac{\alpha}{\alpha-1}} - 1}{\alpha}, \quad \text{for } x \leq \frac{1}{1-\alpha}.$$

Thus, we have

$$(f_\alpha^*)'(x) = (1 + (\alpha-1)x)^{\frac{1}{\alpha-1}} = \exp_\alpha(x),$$

where we have originally defined \exp_α in Section 2. Finally, it holds that

$$(f_\alpha^*)''(x) = \exp_\alpha(x)^{2-\alpha}.$$

Thus,

$$\underline{\mu}_f \geq \lambda(1-\gamma)\pi_{\text{ref}}^{2-2\alpha} \rho_{\min}^2 \pi_{\text{ref}}^2 \exp_\alpha \left(-\frac{16 + 32\gamma\lambda |\log(\pi_{\text{ref}})|/\alpha^2}{\lambda(1-\gamma)^2 \rho_{\min}} \right)^{4-2\alpha},$$

□

In the next two corollaries, we apply Corollary E.7 and Corollary E.10 to get more explicitly the sample complexity of f -PG with entropy regularization.

Corollary E.5. Assume that, for some $1/4 \geq \pi_{\text{ref}} > 0$, π_{ref} satisfy $\mathbf{P}(\pi_{\text{ref}})$. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_ρ . Fix any $(1-\gamma)^{-1} \geq \epsilon > 0$, $\alpha \in (0, 1)$, $\lambda > 0$, and B such that

$$\lambda \leq \min \left(\frac{4}{(1-\gamma)^2 \rho_{\min}} \cdot \frac{1-\alpha}{(\pi_{\text{ref}}/2)^{\alpha-1} - 1}, 1 \right), \quad B \leq \frac{\exp_\alpha \left(-\frac{48 |\log(\pi_{\text{ref}})| \max(\lambda, \alpha^2)}{\lambda \alpha^2 (1-\gamma)^2 \rho_{\min}} \right)^{2\alpha-4}}{\pi_{\text{ref}}^2 \rho_{\min}^2}.$$

Setting

$$H \geq \frac{1}{1-\gamma} \log \left(\frac{28152\pi_{\text{ref}}^{4\alpha-4} |\log(\pi_{\text{ref}})|^2}{\epsilon\lambda\alpha^4(1-\gamma)^5 \rho_{\min}^2 \pi_{\text{ref}}^2} \max(\alpha^2, \lambda)^2 \right) + \frac{196 |\log(\pi_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^3 \rho_{\min}} \max(\alpha^2, \lambda) \quad (97)$$

and

$$\eta \leq \frac{\epsilon B \lambda (1-\gamma)^5 \alpha^4 \pi_{\text{ref}}^{5-5\alpha} \rho_{\min}^2 \pi_{\text{ref}}^2}{3672 \log(\pi_{\text{ref}})^2 \max(\alpha^2, \lambda)^2} \exp_{\alpha} \left(-\frac{48 |\log(\pi_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2 \rho_{\min}} \max(\lambda, \alpha^2) \right)^{4-2\alpha},$$

and

$$T \geq \frac{14688 \log(\pi_{\text{ref}})^2 \max(\alpha^2, \lambda)^2}{\lambda^2 \epsilon B (1-\gamma)^6 \alpha^4 \pi_{\text{ref}}^{7-7\alpha} \rho_{\min}^4 \pi_{\text{ref}}^4} \exp_{\alpha} \left(-\frac{48 |\log(\pi_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2 \rho_{\min}} \max(\lambda, \alpha^2) \right)^{4\alpha-8},$$

guarantees that

$$v_{\star}^f(\rho) - \mathbb{E} \left[v_{\theta_t}^f(\rho) \right] \leq \epsilon,$$

where $f = f_{\alpha}$ is the α -Csizsár–Cressie–Read divergence generator. Thus, the sample complexity of f -PG to learn an ϵ -solution of the α -Tsallis regularized problem is

$$TBH \approx \frac{|\log(\pi_{\text{ref}})|^3 \max(\alpha^{-6}, \lambda^{-3})}{\epsilon B (1-\gamma)^9 \pi_{\text{ref}}^{7-7\alpha} \rho_{\min}^5 \pi_{\text{ref}}^4} \exp_{\alpha} \left(-\frac{|\log(\pi_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2 \rho_{\min}} \max(\lambda, \alpha^2) \right)^{4\alpha-8}.$$

Proof. To prove this corollary, we will show that under the conditions of this corollary, the assumptions of Theorem E.6 holds. Firstly, note that by using Lemma F.4, the assumption $\mathbf{A}_f(\pi_{\text{ref}})$ holds. Secondly, using Lemma F.4 note that

$$\begin{aligned} & \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log \left(\frac{216\omega_f^2}{\epsilon\mu_f(1-\gamma)^4} [4 + 4\lambda^2 d_f^2 + \lambda^2(1-\gamma)^2 y_f^2] \right) \\ & \leq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log \left(\frac{216\pi_{\text{ref}}^{2\alpha-2}}{\epsilon\mu_f(1-\gamma)^4} \left[4 + 32\lambda^2 \frac{|\log(\pi_{\text{ref}})|^2}{\alpha^4} \right] \right) \\ & \leq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log \left(\frac{864\pi_{\text{ref}}^{4\alpha-4} [1 + 32\lambda^2 |\log(\pi_{\text{ref}})|^2 / \alpha^4]}{\epsilon\lambda(1-\gamma)^5 \rho_{\min}^2 \pi_{\text{ref}}^2} \right) \\ & \quad - \frac{1}{1-\gamma} \log \left(\exp_{\alpha} \left(-\frac{16 + 32\gamma\lambda |\log(\pi_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2 \rho_{\min}} \right)^{4-2\alpha} \right), \end{aligned}$$

where in the last inequality, we used the lower bound on μ_f provided in Lemma F.4. Next using the definition of \exp_{α} (see Section 2) and the fact that $\lambda \leq 1$, we have

$$\begin{aligned} & \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log \left(\frac{216\omega_f^2}{\epsilon\mu_f(1-\gamma)^4} [4 + 4\lambda^2 d_f^2 + \lambda^2(1-\gamma)^2 y_f^2] \right) \\ & \leq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log \left(\frac{28152\pi_{\text{ref}}^{4\alpha-4} |\log(\pi_{\text{ref}})|^2}{\epsilon\lambda\alpha^4(1-\gamma)^5 \rho_{\min}^2 \pi_{\text{ref}}^2} \max(\alpha^2, \lambda)^2 \right) \\ & \quad + \frac{4-2\alpha}{1-\alpha} \frac{1}{1-\gamma} \log \left(1 + (1-\alpha) \left(\frac{48 |\log(\pi_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2 \rho_{\min}} \max(\alpha^2, \lambda) \right) \right) \\ & \leq \frac{1}{1-\gamma} \log \left(\frac{28152\pi_{\text{ref}}^{4\alpha-4} |\log(\pi_{\text{ref}})|^2}{\epsilon\lambda\alpha^4(1-\gamma)^5 \rho_{\min}^2 \pi_{\text{ref}}^2} \max(\alpha^2, \lambda)^2 \right) + \frac{196 |\log(\pi_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^3 \rho_{\min}} \max(\alpha^2, \lambda), \end{aligned}$$

where in the last inequality, we used that for $x \geq 0$, we have $\log(1+x) \leq x$. This shows that our condition on H guarantees that the one set in Theorem E.9 is satisfied. Next, using again Lemma F.4 observe that

$$L_f = 196 \frac{\pi_{\text{ref}}^{2\alpha-2} |\log(\pi_{\text{ref}})|}{\alpha^2(1-\gamma)^3} \max(\alpha^2, \lambda), \quad \sigma_f^2 \leq \frac{204\pi_{\text{ref}}^{3\alpha-3} \log(\pi_{\text{ref}})^2}{\alpha^4(1-\gamma)^4} \max(\alpha^2, \lambda)^2 \quad (98)$$

$$\underline{\mu}_f \geq \lambda(1-\gamma)\underline{\pi}_{\text{ref}}^{2-2\alpha} \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2 \exp_{\alpha} \left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2\rho_{\min}} \max(\lambda, \alpha^2) \right)^{4-2\alpha}. \quad (99)$$

Hence, we have that

$$\begin{aligned} & \min \left(\frac{1}{2L_f}, \frac{\epsilon B \underline{\mu}_f}{18\sigma_f^2} \right) \\ & \geq \min \left(\frac{\alpha^2(1-\gamma)^3 \min(\alpha^{-2}, \lambda^{-1})}{392\underline{\pi}_{\text{ref}}^{2\alpha-2} |\log(\underline{\pi}_{\text{ref}})|}, \right. \\ & \quad \left. \frac{\epsilon B \lambda(1-\gamma)^5 \alpha^4 \underline{\pi}_{\text{ref}}^{5-5\alpha} \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}{3672 \log(\underline{\pi}_{\text{ref}})^2 \max(\alpha^2, \lambda)^2} \exp_{\alpha} \left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2\rho_{\min}} \max(\lambda, \alpha^2) \right)^{4-2\alpha} \right) \end{aligned} \quad (100)$$

$$= \frac{\epsilon B \lambda(1-\gamma)^5 \alpha^4 \underline{\pi}_{\text{ref}}^{5-5\alpha} \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}{3672 \log(\underline{\pi}_{\text{ref}})^2 \max(\alpha^2, \lambda)^2} \exp_{\alpha} \left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2\rho_{\min}} \max(\lambda, \alpha^2) \right)^{4-2\alpha}, \quad (101)$$

where in the last identity, we used the fact that $\epsilon < (1-\gamma)^{-1}$ and that

$$B \leq \frac{1}{\underline{\pi}_{\text{ref}}^2 \rho_{\min}^2} \exp_{\alpha} \left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2\rho_{\min}} \max(\lambda, \alpha^2) \right)^{2\alpha-4}.$$

This shows that our condition on η guarantees that the one set in Theorem E.9 is satisfied. Finally using (101) and (99), we have

$$\begin{aligned} & \frac{\underline{\mu}_f}{\underline{\mu}_f} \max \left(2L_f, \frac{18\sigma_f^2}{\epsilon B \underline{\mu}_f} \right) \\ & \leq \frac{4}{\underline{\mu}_f} \frac{3672 \log(\underline{\pi}_{\text{ref}})^2 \max(\alpha^2, \lambda)^2}{\epsilon B \lambda(1-\gamma)^5 \alpha^4 \underline{\pi}_{\text{ref}}^{5-5\alpha} \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2} \exp_{\alpha} \left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2\rho_{\min}} \max(\lambda, \alpha^2) \right)^{2\alpha-4} \\ & \leq \frac{14688 \log(\underline{\pi}_{\text{ref}})^2 \max(\alpha^2, \lambda)^2}{\lambda^2 \epsilon B (1-\gamma)^6 \alpha^4 \underline{\pi}_{\text{ref}}^{7-7\alpha} \rho_{\min}^4 \underline{\pi}_{\text{ref}}^4} \exp_{\alpha} \left(-\frac{48|\log(\underline{\pi}_{\text{ref}})|}{\lambda\alpha^2(1-\gamma)^2\rho_{\min}} \max(\lambda, \alpha^2) \right)^{4\alpha-8}, \end{aligned}$$

which concludes the proof. \square

Corollary F.6. Assume that, for some $\underline{\pi}_{\text{ref}} \in (0, 1/4]$, $\underline{\pi}_{\text{ref}}$ satisfy $\mathbf{P}(\underline{\pi}_{\text{ref}})$. Assume in addition that the initial distribution ρ satisfies \mathbf{A}_{ρ} and fix f to be the α -Csiszár–Cressie–Read divergence generator, i.e.

$$f(u) = f_{\alpha}(u) = \frac{u^{\alpha} - \alpha u + \alpha - 1}{\alpha(\alpha - 1)}.$$

Fix any $\epsilon \in (0, (1-\gamma)^{-1}]$, such that

$$\epsilon < \frac{16}{(1-\gamma)^3 \rho_{\min}} \cdot \frac{1-\alpha}{(\underline{\pi}_{\text{ref}}/2)^{\alpha-1} - 1}, \quad \text{and set } \lambda = \frac{(1-\gamma)\alpha^2\epsilon}{16|\log(\underline{\pi}_{\text{ref}})|}. \quad (102)$$

Additionally set any B such that

$$B \leq \frac{1}{\epsilon^2 \alpha^2 (1-\gamma)^3 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2} \exp_{\alpha} \left(-\frac{384|\log(\underline{\pi}_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right)^{2\alpha-4} \quad (103)$$

Setting

$$H \geq \frac{1}{1-\gamma} \log \left(\frac{19008 \underline{\pi}_{\text{ref}}^{4\alpha-2} |\log(\underline{\pi}_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^6 \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2} \right) + \frac{1540 |\log(\underline{\pi}_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^4 \rho_{\min}},$$

and

$$\eta \leq \frac{\epsilon^2 \alpha^2 (1-\gamma)^6 B \rho_{\min}^2 \underline{\pi}_{\text{ref}}^2}{13824 \underline{\pi}_{\text{ref}}^{5\alpha-5} |\log(\underline{\pi}_{\text{ref}})|} \exp_{\alpha} \left(-\frac{384|\log(\underline{\pi}_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right)^{4-2\alpha}, \quad (104)$$

and

$$T \geq \frac{12^7 \pi_{\text{ref}}^{7\alpha-7} |\log(\pi_{\text{ref}})|^2}{\epsilon^3 \alpha^4 (1-\gamma)^8 \rho_{\min}^4 \pi_{\text{ref}}^4 B} \exp_{\alpha} \left(-\frac{384 |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right)^{4\alpha-8} \log \left(\frac{6(v_{\star}^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon} \right), \quad (105)$$

guarantees that

$$v_{\star}(\rho) - \mathbb{E}[v_{\theta_t}(\rho)] \leq \epsilon.$$

Thus, the sample complexity of f -PG, where f is the α -Csiszár–Cressie–Read divergence generator, to learn an ϵ -solution of the non-regularized problem is

$$TBH \approx \frac{\pi_{\text{ref}}^{7\alpha-7} |\log(\pi_{\text{ref}})|^3}{\epsilon^4 \alpha^6 (1-\gamma)^{12} \rho_{\min}^5 \pi_{\text{ref}}^4} \exp_{\alpha} \left(-\frac{384 |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right)^{4\alpha-8}.$$

Proof. To prove this corollary, we will show that under the conditions of this corollary, the assumptions of Theorem E.9 holds. Firstly, note that (102) implies that

$$\epsilon < \frac{16}{(1-\gamma)^3 \rho_{\min}} \min \left(\frac{4}{|f'(\iota_f)|}, \frac{1}{|f'(\frac{1}{2})|}, \frac{4}{|f'(\frac{1}{2} \pi_{\text{ref}})|} \right), \quad \lambda = \frac{(1-\gamma)\epsilon}{4} c_f,$$

with $c_f = \alpha^2/4 |\log(\pi_{\text{ref}})|$. Additionally, observe using Lemma F.4 that we have

$$C_f^{(1)} \leq \frac{64 \pi_{\text{ref}}^{2\alpha-2} |\log(\pi_{\text{ref}})|}{\alpha^2}, \quad C_f^{(2)} \leq 96 \pi_{\text{ref}}^{2\alpha-2}, \quad C_f^{(3)} \leq \frac{13824 \pi_{\text{ref}}^{5\alpha-5} |\log(\pi_{\text{ref}})|}{\alpha^2},$$

$$d(\epsilon) = \exp_{\alpha} \left(-\frac{384 |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right)^{4-2\alpha}, \quad (106)$$

where $C_f^{(1)}, C_f^{(2)}, C_f^{(3)}$, and $d(\epsilon)$ are defined in (78) and (77). Next, observe that

$$\begin{aligned} & \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log \left(\frac{297 \omega_f^2 C_1^f}{\epsilon d(\epsilon) (1-\gamma)^6 \rho_{\min}^2 \pi_{\text{ref}}^2} \right) \\ & \leq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log \left(\frac{19008 \pi_{\text{ref}}^{4\alpha-2} |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^6 \rho_{\min}^2 \pi_{\text{ref}}^2} \right) \\ & \quad + \frac{4-2\alpha}{1-\gamma} \log \left(\exp_{\alpha} \left(-\frac{384 |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right) \right) \\ & \leq \frac{4}{(1-\gamma)^2} + \frac{1}{1-\gamma} \log \left(\frac{19008 \pi_{\text{ref}}^{4\alpha-2} |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^6 \rho_{\min}^2 \pi_{\text{ref}}^2} \right) + \frac{1536 |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^4 \rho_{\min}}, \end{aligned}$$

where in the last inequality, we used that $\log(1+u) \leq x$ for $u > -1$. This shows that our condition on H implies the one assumed in Theorem E.9. Using (106), we have

$$\begin{aligned} & \min \left(\frac{(1-\gamma)^3}{C_f^{(2)}}, \frac{\epsilon^2 d(\epsilon) (1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}{C_f^{(3)}} \right) \\ & \geq \min \left(\frac{(1-\gamma)^3}{96 \pi_{\text{ref}}^{2\alpha-2}}, \frac{\epsilon^2 \alpha^2 (1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}{13824 \pi_{\text{ref}}^{5\alpha-5} |\log(\pi_{\text{ref}})|} \exp_{\alpha} \left(-\frac{384 |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right)^{4-2\alpha} \right) \\ & \geq \frac{\epsilon^2 \alpha^2 (1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2}{13824 \pi_{\text{ref}}^{5\alpha-5} |\log(\pi_{\text{ref}})|} \exp_{\alpha} \left(-\frac{384 |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right)^{4-2\alpha}, \quad (107) \end{aligned}$$

where in the last inequality, we used the condition on B introduced in (103). Hence our condition on the step-size ensures that the one assumed in Theorem E.9 is satisfied. Next, using (106) and (107) yields

$$\frac{16 C_f^{(1)} \ell(\epsilon)}{\epsilon d(\epsilon) (1-\gamma)^2 \rho_{\min}^2 \pi_{\text{ref}}^2} \max \left(\frac{C_f^{(2)}}{(1-\gamma)^3}, \frac{C_f^{(3)}}{\epsilon^2 d(\epsilon) (1-\gamma)^6 B \rho_{\min}^2 \pi_{\text{ref}}^2} \right)$$

$$\begin{aligned}
&\leq \frac{16C_f^{(1)}\ell(\epsilon)}{\epsilon d(\epsilon)(1-\gamma)^2\rho_{\min}^2\pi_{\text{ref}}^2} \cdot \frac{13824\pi_{\text{ref}}^{5\alpha-5}|\log(\pi_{\text{ref}})|}{\epsilon^2\alpha^2(1-\gamma)^6B\rho_{\min}^2\pi_{\text{ref}}^2} \exp_{\alpha}\left(-\frac{384|\log(\pi_{\text{ref}})|}{\epsilon\alpha^2(1-\gamma)^3\rho_{\min}}\right)^{2\alpha-4} \\
&\leq \frac{1024\pi_{\text{ref}}^{2\alpha-2}|\log(\pi_{\text{ref}})|\ell(\epsilon)}{\epsilon\alpha^2(1-\gamma)^2\rho_{\min}^2\pi_{\text{ref}}^2} \cdot \frac{13824\pi_{\text{ref}}^{5\alpha-5}|\log(\pi_{\text{ref}})|}{\epsilon^2\alpha^2(1-\gamma)^6B\rho_{\min}^2\pi_{\text{ref}}^2} \exp_{\alpha}\left(-\frac{384|\log(\pi_{\text{ref}})|}{\epsilon\alpha^2(1-\gamma)^3\rho_{\min}}\right)^{4\alpha-8} \\
&\leq \frac{12^7\pi_{\text{ref}}^{7\alpha-7}|\log(\pi_{\text{ref}})|^2}{\epsilon^3\alpha^4(1-\gamma)^8\rho_{\min}^4\pi_{\text{ref}}^4B} \exp_{\alpha}\left(-\frac{384|\log(\pi_{\text{ref}})|}{\epsilon\alpha^2(1-\gamma)^3\rho_{\min}}\right)^{4\alpha-8} \log\left(\frac{6(v_{\star}^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon}\right),
\end{aligned}$$

which proves that under our condition on T the one assumed by Theorem E.9 is satisfied. \square

Corollary F.7. *Assume the same condition of Corollary F.6. For any $(1-\gamma)^{-1} > \epsilon > 0$ and $\alpha \in (0, 1)$, denote respectively by $T(\epsilon, \alpha)$, $B(\epsilon, \alpha)$, and $H(\epsilon, \alpha)$, the thresholds set in Corollary F.6 on T , B , and H , to learn an ϵ -solution of the unregularized problem. Additionally, denote by $\alpha^*(\epsilon)$ the minimizer of $T(\epsilon, \alpha)B(\epsilon, \alpha)H(\epsilon, \alpha)$. It holds that*

$$\alpha^*(\epsilon) = \frac{11}{2} \cdot \frac{1}{\log(1/\epsilon)} + o\left(\frac{1}{\log(1/\epsilon)}\right).$$

Additionally, for $\epsilon < e^{-11}$, it holds that

$$T(\epsilon, \alpha^*(\epsilon))B(\epsilon, \alpha^*(\epsilon))H(\epsilon, \alpha^*(\epsilon)) = \tilde{O}\left(\frac{|\log(\pi_{\text{ref}})|^{11}}{\epsilon^{12}(1-\gamma)^{36}\rho_{\min}^{14}\pi_{\text{ref}}^{11}} \log\left(\frac{(v_{\star}^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon}\right)\right).$$

Proof. We first provide an equivalent of the α that optimises the sample complexity provided in Corollary F.6 and then bound the sample complexity obtained by using this α .

Finding the best α . Firstly, note that using Corollary F.6, we have

$$\begin{aligned}
S(\epsilon, \alpha) &:= \log(T(\epsilon, \alpha)B(\epsilon, \alpha)H(\epsilon, \alpha)) \\
&= \log\left(\frac{12^7\pi_{\text{ref}}^{7\alpha-7}|\log(\pi_{\text{ref}})|^2}{\epsilon^3\alpha^4(1-\gamma)^8\rho_{\min}^4\pi_{\text{ref}}^4}\right) \\
&\quad + \frac{4\alpha-8}{\alpha-1} \log\left(1 + (1-\alpha)\left(\frac{384|\log(\pi_{\text{ref}})|}{\epsilon\alpha^2(1-\gamma)^3\rho_{\min}}\right)\right) \\
&\quad + \log\left(\log\left(\frac{6(v_{\star}^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon}\right)\right) \\
&\quad + \log\left(\frac{1}{1-\gamma} \log\left(\frac{19008\pi_{\text{ref}}^{4\alpha-2}|\log(\pi_{\text{ref}})|}{\epsilon\alpha^2(1-\gamma)^6\rho_{\min}^2\pi_{\text{ref}}^2}\right) + \frac{1540|\log(\pi_{\text{ref}})|}{\epsilon\alpha^2(1-\gamma)^4\rho_{\min}}\right).
\end{aligned}$$

Firstly, observe that for any function $k(\epsilon)$ which does converge to a different value from 0, we have

$$\lim_{\epsilon \rightarrow 0} \frac{S(\epsilon, 1/\log(1/\epsilon))}{S(\epsilon, k(\epsilon))} < 1,$$

which establishes that $\alpha^*(\epsilon) \rightarrow 0$. This allows, to rewrite $S(\epsilon, \alpha)$ as

$$S(\epsilon, \alpha) = \log\left(\frac{1}{\epsilon^4\alpha^6}\right) + \frac{8-4\alpha}{1-\alpha} \log\left(\frac{1}{\epsilon\alpha^2}\right) + \psi(\alpha, \epsilon),$$

where $\psi(\alpha, \epsilon)$ is defined as

$$\begin{aligned}
\psi(\alpha, \epsilon) &= \log\left(\frac{12^7\pi_{\text{ref}}^{7\alpha-7}|\log(\pi_{\text{ref}})|^2}{(1-\gamma)^8\rho_{\min}^4\pi_{\text{ref}}^4}\right) \\
&\quad + \frac{4\alpha-8}{\alpha-1} \left[\log\left(1 + (1-\alpha)\left(\frac{384|\log(\pi_{\text{ref}})|}{\epsilon\alpha^2(1-\gamma)^3\rho_{\min}}\right)\right) - \log\left(\frac{1}{\epsilon\alpha^2}\right) \right] \\
&\quad + \log\left(\log\left(\frac{6(v_{\star}^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon}\right)\right)
\end{aligned}$$

$$+ \log \left(\frac{\epsilon \alpha^2}{1-\gamma} \log \left(\frac{19008 \pi_{\text{ref}}^{4\alpha-2} |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^6 \rho_{\min}^2 \pi_{\text{ref}}^2} \right) + \frac{1540 |\log(\pi_{\text{ref}})|}{(1-\gamma)^4 \rho_{\min}} \right) .$$

Importantly, observe that $\psi(\alpha, \epsilon)$ is dominated by $\log(1/\epsilon)$ when $(\alpha, \epsilon) \rightarrow 0$ and that

$$\frac{\partial \psi(\alpha, \epsilon)}{\partial \alpha} = o \left(\frac{1}{\alpha \log(\frac{1}{\epsilon \alpha})} \right) .$$

Computing the derivative of this function with respect to α yields

$$\frac{\partial S(\epsilon, \alpha)}{\partial \alpha} = \frac{-6}{\alpha} + 4 \frac{1}{(1-\alpha)^2} \log \left(\frac{1}{\epsilon} \right) + 4 \frac{1}{(1-\alpha)^2} \log \left(\frac{1}{\alpha^2} \right) - \frac{8}{\alpha} \left(1 + \frac{1}{1-\alpha} \right) + \frac{\partial \psi(\epsilon, \alpha)}{\partial \alpha} \quad (108)$$

$$= \frac{-22}{\alpha} + 4 \log \left(\frac{1}{\epsilon} \right) + o \left(\frac{1}{\alpha \log(\frac{1}{\epsilon \alpha})} \right) . \quad (109)$$

As

$$\left. \frac{\partial S(\epsilon, \alpha)}{\partial \alpha} \right|_{\alpha=\alpha^*(\epsilon)} = 0 ,$$

Then this implies that

$$\alpha^*(\epsilon) = \frac{11}{2} \cdot \frac{1}{\log(1/\epsilon)} + o \left(\frac{1}{\log(1/\epsilon)} \right) .$$

Next, we provide a bound on the sample complexity given by this $\alpha^*(\epsilon)$.

Computing the sample complexity. Firstly, note that for $\alpha \leq 1/2$ and $\epsilon < 1$, we have

$$\begin{aligned} & T(\epsilon, \alpha) B(\epsilon, \alpha) H(\epsilon, \alpha) \\ & \leq \frac{12^{11} \pi_{\text{ref}}^{7\alpha-7} |\log(\pi_{\text{ref}})|^3}{\epsilon^4 \alpha^6 (1-\gamma)^{12} \rho_{\min}^5 \pi_{\text{ref}}^4} \exp_{\alpha} \left(-\frac{384 |\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right)^{4\alpha-8} \log \left(\frac{6(v_{\star}^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon} \right) \\ & \leq \frac{\pi_{\text{ref}}^{7\alpha-7} |\log(\pi_{\text{ref}})|^3}{\epsilon^4 \alpha^6 (1-\gamma)^{12} \rho_{\min}^5 \pi_{\text{ref}}^4} \left(\frac{|\log(\pi_{\text{ref}})|}{\epsilon \alpha^2 (1-\gamma)^3 \rho_{\min}} \right)^{(8-4\alpha)/(1-\alpha)} \log \left(\frac{6(v_{\star}^f(\rho) - v_{\theta_0}^f(\rho))}{\epsilon} \right) . \quad (110) \end{aligned}$$

Next, note that for $\epsilon < e^{-11}$, we have $\alpha = \frac{11}{2 \log(1/\epsilon)} \leq 1/2$. Thus, using that $1/(1-\alpha) \leq 1+2\alpha$ yield

$$\left(\frac{1}{\epsilon \alpha^2} \right)^{(8-4\alpha)/(1-\alpha)} = \exp \left(\frac{8-4\alpha}{1-\alpha} \cdot \log \left(\frac{1}{\epsilon \alpha^2} \right) \right) \leq \exp \left((8+8\alpha) \cdot \log \left(\frac{1}{\epsilon \alpha^2} \right) \right) \lesssim \frac{1}{\epsilon^8 \alpha^{16}} .$$

Plugging in the previous bound in (110) concludes the proof. \square

G DISCUSSION ON UNREGULARIZED POLICY GRADIENT WITH f -SOFTARGMAX PARAMETERIZATION

In this section, we show that it is possible to derive Non-Uniform Łojasiewicz inequalities on the unregularized objective $v_{\theta}(\rho) := v_{\pi_{\theta}^f}(\rho)$ under f -SoftArgmax Parameterization in the bandit setting. We expect the analysis to extend to the RL setting using similar arguments to those of [Mei et al. \(2020b;a\)](#); [Liu et al. \(2025\)](#).

Theorem G.1. *Consider the bandit case, i.e. $|\mathcal{S}| = 1$. Assume that, for some $\pi_{\text{ref}} > 0$, f and π_{ref} satisfy $\mathbf{A}_f(\pi_{\text{ref}})$ and $\mathbf{P}(\pi_{\text{ref}})$ respectively. We also assume that the function $1/f''$ is convex and that the initial distribution ρ satisfies \mathbf{A}_{ρ} . Then, it holds that*

$$\|\nabla v_{\theta}(\rho)\|_2^2 \geq \frac{w_{\theta}^f(a^*)^2}{\zeta_f} \frac{1}{v_{\star}(\rho) - v_{\pi_{\text{ref}}}(\rho)} \cdot f'' \left(\frac{v_{\star}(\rho) - v_{\theta}(\rho)}{v_{\star}(\rho) - v_{\pi_{\text{ref}}}(\rho)} \right) ,$$

where $v_{\star}(\rho) = \max_{\pi \in \mathcal{P}(\mathcal{A})^{\mathcal{S}}} v_{\pi}(\rho)$.

Proof. Subsequently, we drop the dependency on the state for more clarity. Using Lemma C.2 (with $\gamma = 0$ and $|\mathcal{S}| = 1$), for any $a \in \mathcal{A}$, we have

$$\frac{\partial v_\theta(\rho)}{\partial \theta(a)} = w_\theta^f(a) \cdot (r(a) - \sum_b w_\theta^f(b)r(b)).$$

Denote by a^* any optimal action, and $r^* = r(a^*)$. We have

$$\|\nabla v_\theta(\rho)\|_2^2 = \sum_a w_\theta^f(a)^2 \cdot \left(r(a) - \sum_b w_\theta^f(b)r(b) \right)^2 \geq w_\theta^f(a^*)^2 \left(\sum_b w_\theta^f(b)(r^* - r(b)) \right)^2,$$

Using that $\Delta(a) = r^* - r(a) \geq 0$, and also

$$w_\theta^f(a) = \frac{1}{W_\theta^f} \cdot \frac{\pi_{\text{ref}}(a)}{f''(\pi_\theta^f(a)/\pi_{\text{ref}}(a))}, \quad W_\theta^f := \sum_{b \in \mathcal{A}} \frac{\pi_{\text{ref}}(b)}{f''(\pi_\theta^f(b)/\pi_{\text{ref}}(b))},$$

yields

$$\sum_b w_\theta^f(b)\Delta(b) = \frac{1}{W_\theta^f} \cdot \sum_{b \in \mathcal{A}} \Delta(b)\pi_{\text{ref}}(b) \cdot \frac{1}{f''(\pi_\theta^f(b)/\pi_{\text{ref}}(b))}.$$

Let's introduce a distribution $\gamma(a) = \Delta(a)\pi_{\text{ref}}(a)/(\sum_{b \in \mathcal{A}} \Delta(b)\pi_{\text{ref}}(b))$, then we have

$$\sum_{b \in \mathcal{A}} w_\theta^f(b)\Delta(b) = \frac{1}{W_\theta^f \cdot \mathbb{E}_{c \sim \pi_{\text{ref}}}[\Delta(c)]} \mathbb{E}_{c \sim \gamma} \left[1/(f''(\pi_\theta^f(c)/\pi_{\text{ref}}(c))) \right].$$

Next, we use that a map $x \mapsto 1/f''(x)$ is convex, and thus we have

$$\begin{aligned} \mathbb{E}_{c \sim \gamma} \left[1/(f''(\pi_\theta^f(c)/\pi_{\text{ref}}(c))) \right] &\geq 1/f'' \left(\mathbb{E}_{c \sim \gamma} \left[\pi_\theta^f(c)/\pi_{\text{ref}}(c) \right] \right) \\ &= 1/f'' \left(\frac{\sum_{c \in \mathcal{A}} \Delta(c)\pi_\theta^f(c)}{\sum_{c \in \mathcal{A}} \Delta(c)\pi_{\text{ref}}(c)} \right). \end{aligned}$$

Overall, we have

$$\sum_{b \in \mathcal{A}} \pi_\theta^f(b)\Delta(b) \geq \frac{1}{W_\theta^f} \frac{1}{\left(\sum_{c \in \mathcal{A}} \Delta(c)\pi_{\text{ref}}(c) \right) f'' \left(\frac{\sum_{c \in \mathcal{A}} \Delta(c)\pi_\theta^f(c)}{\sum_{c \in \mathcal{A}} \Delta(c)\pi_{\text{ref}}(c)} \right)}.$$

Finally, using that $W_\theta^f \leq \zeta_f$ concludes the proof. \square

Similarly, to Mei et al. (2020b); Liu et al. (2024; 2025), this Łojasiewicz inequality depends on the probability of the optimal action, which is very restrictive. Although extending the analysis of Mei et al. (2020b); Liu et al. (2024; 2025) in the deterministic setting is possible, addressing the stochastic setting for this type of Łojasiewicz inequality appears very challenging. This justifies adding a regularizer to the objective to ensure better PL inequalities, and on which the minimal coefficient can be lower bounded on the trajectory by leveraging a proper projection operator.

H LINKS WITH MIRROR DESCENT

We stress that our proposed method is fundamentally different from mirror descent. Let us define a mapping $\Phi(\pi) = \sum_{s \in \mathcal{S}} D^f(\pi(\cdot|s) \|\pi_{\text{ref}}(\cdot|s))$. For the functions f that we consider, Φ is Legendre on the positive orthant and separable across states (Bubeck et al., 2015). In this case, the f -regularized value function $v_\pi^f(\rho)$ can be optimized directly in the policy space via the Lazy Mirror Descent algorithm (or dual averaging; see Nesterov (2009); Xiao (2009); Juditsky et al. (2023)) with Φ as mirror map. Denoting by π_t the policy at step t and $\tilde{\pi}_t$ by the unnormalized policy at step t , the lazy MD updates reads:

$$\nabla \Phi(\tilde{\pi}_{t+1}) = \nabla \Phi(\tilde{\pi}_t) + \eta \nabla_\pi v_\pi^f(\rho)|_{\pi=\pi_t}, \quad \pi_{t+1} = \arg \min_{\pi \in \Pi} B_\Phi(\pi \|\tilde{\pi}_{t+1}). \quad (111)$$

where $\Pi = \mathcal{P}(\mathcal{A})^{|\mathcal{S}|}$ is a policy space and $B_{\Phi}(\pi||\pi') = \Phi(\pi) - \Phi(\pi') - \langle \nabla\Phi(\pi'), \pi - \pi' \rangle$ is the corresponding Bregman divergence. Since Φ is separable over states, the Bregman projection can be written state-wise as $\pi_{t+1}(\cdot|s) = \text{f-softargmax}(\nabla\Phi(\tilde{\pi}_{t+1})(s, \cdot), \pi_{\text{ref}}(\cdot|s))$.

By denoting $\theta_t = \nabla\Phi(\tilde{\pi}_t)$, one obtains updates that resemble those of (10) (after the removal of \mathcal{T}), with one important difference: the gradient in (111) is taken with respect to the policy π whereas in (10) it is computed w.r.t the "dual" parameter θ (in the MD terminology). Even more important, the update (111) can be expressed as, by the chain rule

$$\theta_{t+1} = \theta_t + \eta \left[\frac{\partial \pi^f}{\partial \theta} \Big|_{\theta=\theta_t} \right]^{-1} \nabla_{\theta} J^f(\theta_t),$$

which have an additional preconditioning term given by the inverse of the policy Jacobian.

A crucial feature of (10) is that it performs a gradient ascent in the "dual" space directly. This algorithm can be extended in the non-tabular setting directly, by parameterizing the function $\theta(s, a)$, allowing extensions to deep RL. This is in contrast with Lazy-MD methods (Nesterov, 2009; Xiao, 2009; Juditsky et al., 2023), due to preconditioning, which cannot be expressed as direct parameter-space gradient steps. This remark has several important implications, which we leave for future work.

I TECHNICAL LEMMAS

Lemma I.1 (Lemma 1.2.3 in Nesterov (2004)). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be twice continuously differentiable. Suppose there exists $L \geq 0$ such that for all $x \in \mathbb{R}^d$ and $v \in \mathbb{R}^d$,*

$$|v^{\top} \nabla^2 f(x) v| \leq L \|v\|^2.$$

Then f has an L -Lipschitz continuous gradient (i.e., f is L -smooth); in particular,

$$\|\nabla f(y) - \nabla f(x)\| \leq L \|y - x\|,$$

and

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle - \frac{L}{2} \|y - x\|^2$$

for all $x, y \in \mathbb{R}^d$.

Lemma I.2. *Consider any two policies π_i , $i = 1, 2$. It holds that*

$$\|d_{\rho}^{\pi_1} - d_{\rho}^{\pi_2}\|_1 \leq \frac{\gamma}{1 - \gamma} \sup_{s \in \mathcal{S}} \|\pi_1(\cdot|s) - \pi_2(\cdot|s)\|_1 .$$

Proof. Let us start from the definition of flow conservation constraints for the discounted state visitation (Puterman, 1994), for $i \in \{1, 2\}$, we have

$$d_{\rho}^{\pi_i}(s) = (1 - \gamma)\rho(s) + \gamma \sum_{s'} \mathbf{P}_{\pi_i}(s|s') d_{\rho}^{\pi_i}(s') .$$

Then, we have

$$\begin{aligned} \sum_{s \in \mathcal{S}} |d_{\rho}^{\pi_2}(s) - d_{\rho}^{\pi_1}(s)| &\leq \gamma \sum_{(s', a')} \sum_s |\mathbf{P}(s|s', a') \pi_2(a'|s') d_{\rho}^{\pi_2}(s') - \mathbf{P}(s|s', a') \pi_1(a'|s') d_{\rho}^{\pi_1}(s')| \\ &\leq \gamma \sum_{s', a'} \sum_s \mathbf{P}(s|s', a') |\pi_2(a'|s') - \pi_1(a'|s')| d_{\rho}^{\pi_2}(s') \\ &\quad + \gamma \sum_{s', a'} \sum_s \mathbf{P}(s|s', a') \pi_1(a'|s') |d_{\rho}^{\pi_1}(s') - d_{\rho}^{\pi_2}(s')| \\ &\leq \gamma \sup_{s \in \mathcal{S}} \|\pi_1(\cdot|s) - \pi_2(\cdot|s)\|_1 + \gamma \sum_{s'} |d_{\rho}^{\pi_1}(s') - d_{\rho}^{\pi_2}(s')| , \end{aligned}$$

which concludes the proof. \square

Lemma I.3 (Performance Difference Lemma). *It holds that*

$$v_{\star}^f(\rho) - v_{\theta}^f(\rho) = \frac{1}{1 - \gamma} \sum_{s \in \mathcal{S}} d_{\rho}^{\pi_{\star}^f}(s) \left[\sum_{a \in \mathcal{A}} \pi_{\star}^f(a|s) q_{\theta}^f(s, a) - \lambda \mathbf{D}^f(\pi_{\star}^f(\cdot|s) || \pi_{\text{ref}}(\cdot|s)) - v_{\pi_{\theta}}^f(s) \right].$$

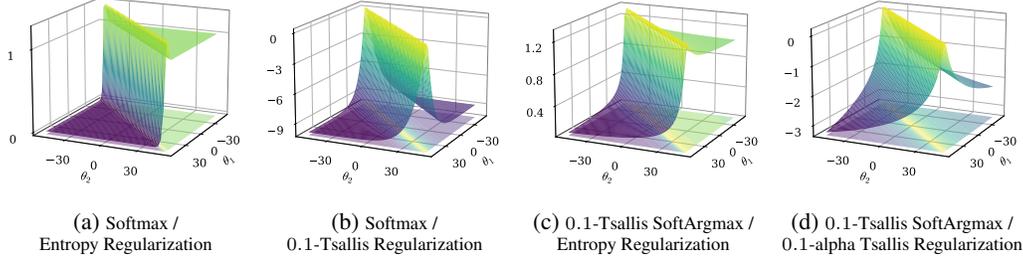


Figure 3: Regularized value landscapes for a one-state, two-action MDP (rewards 0, 1) for different coupling between parameterizations and regularizations.

Proof. Fix $\theta \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$ and any state $s \in \mathcal{S}$. It holds that

$$\begin{aligned}
v_\star^f(s) - v_{\pi_\theta}^f(s) &= \sum_{a \in \mathcal{A}} \pi_\star^f(a|s) q_\star^f(s, a) - \sum_{a \in \mathcal{A}} \pi_\theta^f(a|s) q_\theta^f(s, a) \\
&\quad - \lambda D^f(\pi_\star^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) + \lambda D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) \\
&= \sum_{a \in \mathcal{A}} \pi_\star^f(a|s) (q_\star^f(s, a) - q_\theta^f(s, a)) + \sum_{a \in \mathcal{A}} (\pi_\star^f(a|s) - \pi_\theta^f(a|s)) q_\theta^f(s, a) \\
&\quad - \lambda D^f(\pi_\star^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) + \lambda D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) \\
&= \gamma \sum_{a \in \mathcal{A}} \pi_\star^f(a|s) \sum_{s' \in \mathcal{S}} P(s'|s, a) (v_\star^f(s') - v_\theta^f(s')) + \sum_{a \in \mathcal{A}} (\pi_\star^f(a|s) - \pi_\theta^f(a|s)) q_\theta^f(s, a) \\
&\quad - \lambda D^f(\pi_\star^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) + \lambda D^f(\pi_\theta^f(\cdot|s) \| \pi_{\text{ref}}(\cdot|s)) ,
\end{aligned}$$

where in the last equality, we used the definition of the regularized Q-function (2). Expanding the recursion yields

$$\begin{aligned}
v_\star^f(s) - v_{\pi_\theta}^f(s) &= \frac{1}{1-\gamma} \sum_{s' \in \mathcal{S}} d_s^\star(s') \left[\sum_{a \in \mathcal{A}} (\pi_\star^f(a|s') - \pi_\theta^f(a|s')) q_\theta^f(s', a) \right] \\
&\quad + \frac{1}{1-\gamma} \sum_{s' \in \mathcal{S}} d_s^\star(s') \left[\lambda D^f(\pi_\theta^f(\cdot|s') \| \pi_{\text{ref}}(\cdot|s')) - \lambda D^f(\pi_\star^f(\cdot|s') \| \pi_{\text{ref}}(\cdot|s')) \right] \\
&= \frac{1}{1-\gamma} \sum_{s' \in \mathcal{S}} d_s^\star(s') \left[\sum_{a \in \mathcal{A}} \pi_\star^f(a|s') q_\theta^f(s', a) - \lambda D^f(\pi_\star^f(\cdot|s') \| \pi_{\text{ref}}(\cdot|s')) \right] \\
&\quad - \frac{1}{1-\gamma} \sum_{s' \in \mathcal{S}} d_s^\star(s') \left[\sum_{a \in \mathcal{A}} \pi_\theta^f(a|s') q_\theta^f(s', a) - \lambda D^f(\pi_\theta^f(\cdot|s') \| \pi_{\text{ref}}(\cdot|s')) \right] ,
\end{aligned}$$

which concludes the proof. \square

Lemma I.4 (Lemma 23 of Mei et al. (2020b)). Let $\pi \in \mathcal{P}(\mathcal{A})$. Denote $H(\pi) = \text{diag}(\pi) - \pi\pi^\top$. For any vector $x \in \mathbb{R}^{|\mathcal{A}|}$

$$\left\| H(\pi) \left(x - \frac{\langle x, \mathbf{1}_{|\mathcal{A}|} \rangle}{|\mathcal{A}|} \mathbf{1}_{|\mathcal{A}|} \right) \right\|_2 \geq \min_{a \in \mathcal{A}} \pi(a) \cdot \left\| x - \frac{\langle x, \mathbf{1}_{|\mathcal{A}|} \rangle}{|\mathcal{A}|} \mathbf{1}_{|\mathcal{A}|} \right\|_2 .$$

Lemma I.5 (Danskin 1966). Let $Z \subset \mathbb{R}^m$ be compact and let $\phi : \mathbb{R}^n \times Z \rightarrow \mathbb{R}$ be continuous. Define

$$f(x) = \max_{z \in Z} \phi(x, z), \quad Z_0(x) = \arg \max_{z \in Z} \phi(x, z).$$

Assume that for each fixed $z \in Z$, the map $x \mapsto \phi(x, z)$ is differentiable. If $Z_0(x) = \{\bar{z}\}$ and $x \mapsto \phi(x, \bar{z})$ is differentiable at x , then f is differentiable at x with

$$\frac{\partial f(x)}{\partial x} = \frac{\partial \phi(x, \bar{z})}{\partial x} .$$

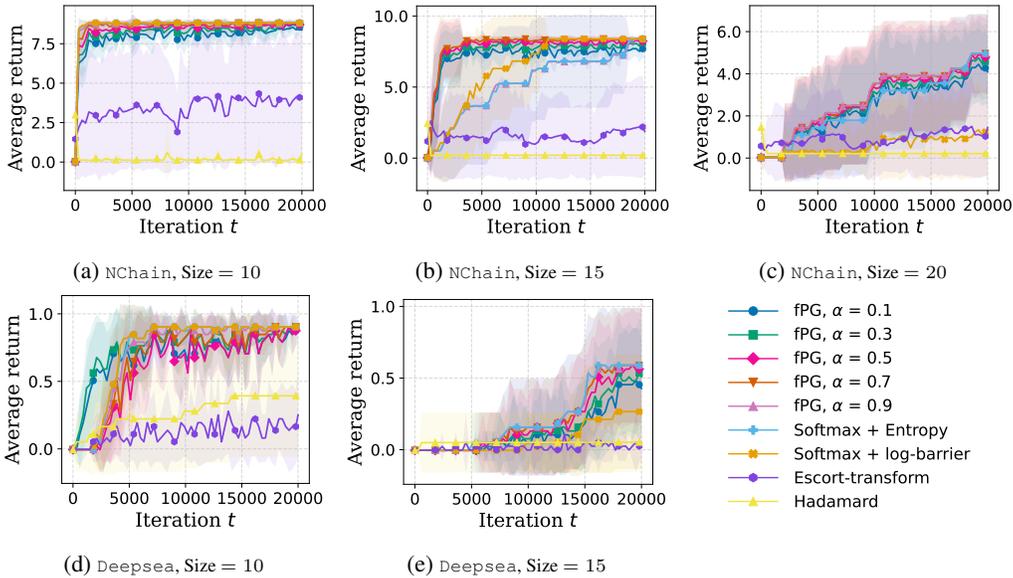


Figure 4: Average return as a function of training iterations on `NChain` (top row: sizes 10, 15, 20) and `DeepSea` (bottom row: sizes 10, 15). For f -PG, we report the best configuration for each divergence parameter α ; for all other methods, we report the best-performing configuration over their respective hyperparameters. Curves show the mean performance over 15 independent seeds, with shaded regions indicating one standard deviation.

J EXPERIMENTS

J.1 UNCOUPLING THE PARAMETERIZATION AND THE REGULARIZATION

Figure 3 compares the regularized value landscapes induced by different couplings between policy parameterizations and regularizers. Both `Softmax / Entropy` Regularization (Figure 3a) and `Softmax / 0.1-Tsallis` Regularization (Figure 3b) produce highly ill-conditioned objectives, characterized by wide flat plateaus separated by extremely sharp ridges. These geometries create large regions with vanishing gradients together with nearly singular directions, which are known to slow down and destabilize policy-gradient methods. Switching to the `Tsallis SoftArgmax` parameterization already improves the situation: under `Entropy` Regularization (Figure 3c), the flat directions are reduced and the basin around the optimum becomes more pronounced. However, the most favorable geometry is obtained when `Tsallis SoftArgmax` is coupled with `Tsallis` regularization (Figure 3d). In this matched `Tsallis-Tsallis` regime, the landscape becomes smooth, strongly curved, and well-conditioned, with a single broad basin leading to the optimum and no spurious flat regions or steep barriers. This alignment between the geometry induced by the parameterization and that of the regularizer yields an almost quadratic objective in logits, explaining why the `Tsallis-Tsallis` coupling provides the best convergence behavior.

J.2 TABULAR EXPERIMENTS

We evaluate the empirical performance of f -PG equipped with α -Tsallis regularization, with the goal of assessing how our coupled parameterization-regularization framework compares to the baselines summarized in Table 2. All methods are evaluated on the *unregularized* return objective, and we report learning curves as a function of training iterations. For each value of the Tsallis parameter α , we tune both the temperature parameter λ and the step-size η over the grid

$$\lambda \in \{10^{-3}, 10^{-2}, 10^{-1}, 1.0\}, \quad \eta \in \{10^{-4}, 3 \times 10^{-4}, 10^{-3}\}.$$

For the baseline methods, we analogously select the best-performing configuration over their respective hyperparameter grids. All curves are averaged over 15 independent random seeds, and shaded regions indicate one standard deviation.

NChain and DeepSea. We consider two canonical tabular exploration benchmarks. `NChain` is a long-horizon chain environment in which the agent must repeatedly move in one direction to reach a terminal state with a large +1 reward, while a small immediate reward equal to +0.01 is available for moving in the opposite direction. As the chain length increases, the probability of discovering the optimal policy decays exponentially unless sufficient structured exploration is induced. The `DeepSea` environment, described in Section 5, is a two-dimensional sparse-reward navigation task in which the agent must follow a precise sequence of actions to reach a distant rewarding state. Both environments therefore test the ability of a policy-gradient method to propagate credit over long horizons and through sparse feedback.

Results. Figure 4 reports learning curves on both environments. On `NChain` (top row, Figures 4a to 4c), the standard softmax-entropy policy gradient baseline performs competitively for the smallest instance (Size 10), but its performance degrades markedly as the chain length increases. In particular, for Size 15 a substantial performance gap opens up, and for Size 20 softmax converges slowly and remains far from optimal. In contrast, f -PG with $\alpha < 1$ achieves substantially higher returns and converges much faster for intermediate horizons, most notably for Size 1,5 where a clear performance gap with the softmax entropy-regularized policy gradient emerges. For the longest chain (Size 20), Tsallis regularization continues to outperform the other baselines but exhibits a similarly slow convergence trend to softmax-entropy, reflecting the difficulty of the problem at this scale. Moreover, alternative regularization schemes fail completely on the longest `NChain` of size 20. Additionally, Escort policy gradients and Hadamard parameterizations plateau at very low returns across all chain lengths, indicating the need for additional exploration.

A similar pattern is observed in `DeepSea` (bottom row, Figures 4d and 4e). For Size 10, Tsallis-regularized policies perform on par with the softmax-entropy baseline and in some cases converge faster. For the more challenging Size 15 instance, Tsallis policies with $\alpha < 1$ remain competitive. In this environment, f -PG dominates all other baselines, which fail to make meaningful progress toward high-return policies.

Although no single Tsallis parameter $\alpha < 1$ is uniformly optimal across all problem sizes, a clear and robust trend emerges across both `NChain` and `DeepSea`: *for every environment, there exists an $\alpha < 1$ that strictly outperforms all alternative schemes and matches or exceeds the performance of softmax-entropy.* These results corroborate the observations of Section 5 and confirm that jointly tailoring the policy parameterization and the regularization to the structure of the problem leads to improved empirical performance.

J.3 FULL DESCRIPTION OF THE DEEPRL EXPERIMENTS

We provide here a full description of the α -Tsallis PPO algorithm and the experimental setup of Section 5.

Algorithm description. α -Tsallis PPO follows the same overall algorithmic structure as standard PPO, with only minor modifications: the softmax policy parameterization and entropy regularization are replaced by their Tsallis α -softargmax and Tsallis α -regularization counterparts. Specifically, the policy network outputs unnormalised action logits, which are mapped to a probability distribution via the Tsallis α -softargmax instead of a usual softmax (see Section 3). At each interaction step, we compute the Tsallis divergence between the current policy and the uniform distribution over actions using the same coefficient α , and subtract $\lambda \cdot D^{f_\alpha}(\pi(\cdot|s)||\text{uniform})$ (see Table 2 for the expression of f_α) from the received reward.² All other components, including the clipped surrogate objective, value loss, and advantage normalization, remain unchanged.

Training pipeline. At each update, we collect trajectories from 16 parallel environments for 32 steps, followed by 16 epochs of PPO optimisation over 4 minibatches. We use a discounting factor

²This regularized reward is used in the computation of advantages and value targets for the `DeepSea` environment. The reason is that adding entropy to the rewards turns out to be critical for the method’s final performance, since the original softmax-PPO fails even in `DeepSea` of size 20. Notice that it contrasts with an original PPO that adds entropy regularization only to a loss function and not to a reward. For `Noisy Carpole`, we use a standard implementation of PPO as a baseline.

$\gamma = 0.99$ and GAE $\lambda = 0.95$, and apply gradient clipping at norm 0.5. Both actor and critic are two-layer multilayer perceptrons with 64 hidden units and \tanh activations. Optimization is performed using Adam (Kingma, 2014), and we perform a grid search over

$$\lambda \in \{10^{-3}, 10^{-2}, 10^{-1}, 1.0\}, \quad \eta \in \{10^{-4}, 3 \times 10^{-4}, 10^{-3}\}.$$

Each configuration is evaluated across 25 seeds. Episode returns are aggregated per configuration and reported as mean \pm standard error. For each α , we select the configuration that produces the highest last reward on average.