

---

# On the Global Convergence Rates of Decentralized Softmax Gradient Play in Markov Potential Games

---

**Runyu Zhang**  
Harvard University  
runyuzhang@fas.harvard.edu

**Jincheng Mei**  
Google Research, Brain Team  
jcmei@google.com

**Bo Dai**  
Google Research, Brain Team  
bodai@google.com

**Dale Schuurmans**  
University of Alberta  
Google Research, Brain Team  
schuurmans@google.com

**Na Li**  
Harvard University  
nali@seas.harvard.edu

## Abstract

Softmax policy gradient is a popular algorithm for policy optimization in single-agent reinforcement learning, particularly since projection is not needed for each gradient update. However, in multi-agent systems, the lack of central coordination introduces significant additional difficulties in the convergence analysis. Even for a stochastic game with identical interest, there can be multiple Nash Equilibria (NEs), which disables proof techniques that rely on the existence of a unique global optimum. Moreover, the softmax parameterization introduces non-NE policies with zero gradient, making it difficult for gradient-based algorithms in seeking NEs. In this paper, we study the finite time convergence of decentralized softmax gradient play in a special form of game, Markov Potential Games (MPGs), which includes the identical interest game as a special case. We investigate both gradient play and natural gradient play, with and without log-barrier regularization. The established convergence rates for the unregularized cases contain a trajectory dependent constant that can be *arbitrarily large*, whereas the log-barrier regularization overcomes this drawback, with the cost of slightly worse dependence on other factors such as the action set size. An empirical study on an identical interest matrix game confirms the theoretical findings.

## 1 Introduction

Multi-agent systems encounter vast application in real world scenarios, such as network routing [35, 8], social and economic decision making [36, 30], and robotic swarms [22, 14]. In these problems, a system consists of a group of agents interacting in a shared environment. Given the recent success of reinforcement learning (RL), increasing attention has been drawn to the possibility of applying RL algorithms, such as policy gradient, to multi-agent systems. However, the theoretical foundations for multi-agent reinforcement learning (MRL) remain limited. Unlike single-agent RL, the actions of other agents affect the dynamics and the decision making outcome for each individual in the system, raising additional theoretical challenges when analyzing joint performance.

The stochastic game (SG) is a classical multi-agent model that has received extensive attention in recent MRL studies. In a stochastic game, the environment is represented by a state space that evolves based on the joint actions of agents. Each agent in a stochastic game tries to maximize its own total reward by making decisions *independently*, based on state information shared between agents. The stochastic game model was first introduced in [31], with a series of followup works proposing NE-seeking algorithms, particularly in the RL setting (e.g. [21, 5, 32, 6, 17, 38] and citations therein). Given recent progress in the underlying theory of RL, many recent works have investigated finite

time iteration and sample complexity for learning NE or other general equilibria notions, such as correlated and coarse correlated equilibria (e.g. [33]).

There are different types of SGs, some with attributes that merit special attention; for example, two-player zero sum games [4, 9], which are widely used to model two player competitive games such as GO. In this paper, we will focus on another type of SG, the Markov potential game (MPG) [23, 27, 39, 18], which includes the identical interest game as a special case. The structure of a MPG enables efficient learning through the use of gradient-based algorithms such as gradient play. Recent work [39, 18] has focused on the iteration and sample complexity of finding a NE in an MPG under the *direct* policy parameterization, which is not practical in most real world scenarios, given the cost of projecting back to the probability simplex on every iteration. This drawback has motivated consideration of the *softmax* parameterization, which bypasses the projection step in the gradient update, and is perhaps the most popular approach to parameterizing policies in practice. [11] have studied natural gradient play for MPG under softmax parameterization, but only address asymptotic behavior and leave finite time complexity open.

From the perspective of analysis and practical performance, the extension from the direct to the softmax parameterization in policies is nontrivial. Even in the single agent case, as shown by [2, 24], there are policies in the softmax parameterization that have near-zero gradient and yet are far from being globally optimal, which creates difficulty for a gradient-based algorithm to escape suboptimal points. A similar issue exists for MPGs: due to the more complex interaction between agents, there is even a greater set of policies that obtain small gradient norm but are far from a NE. Based on our analysis and numerical results, even for natural gradient play—which is known to enjoy dimension free convergence in single agent learning [2]—we find in the multiagent setting that it can still become stuck in these undesirable regions. Such evidence suggests that preconditioning according to the Fisher information matrix [29, 3] is not sufficient to ensure fast convergence in multi-agent learning. A stronger form of regularization is required, which motivates the introduction of log-barrier regularization to avoid undesirable regions of policy space.

| Algorithm                                                             | Single-agent MDP                                                         | Multi-agent MPG                                                                                                    |
|-----------------------------------------------------------------------|--------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Gradient play,<br>direct parameterization                             | $O\left(\frac{ \mathcal{A} M^2}{(1-\gamma)^4\epsilon^2}\right)$<br>[2]   | $O\left(\frac{(\phi_{\max}-\phi_{\min})\sum_{i=1}^n \mathcal{A}_i M^2}{(1-\gamma)^4\epsilon^2}\right)$<br>[39, 18] |
| Gradient play,<br>softmax parameterization                            | $O\left(\frac{M^2}{(1-\gamma)^3c^2\epsilon}\right)$<br>[24]              | $O\left(\frac{n\max_i \mathcal{A}_i (\phi_{\max}-\phi_{\min})M^2}{(1-\gamma)^4c^2\epsilon^2}\right)^*$             |
| Natural gradient play,<br>softmax parameterization                    | $O\left(\frac{1}{(1-\gamma)^2\epsilon}\right)$<br>[2]                    | $O\left(\frac{n(\phi_{\max}-\phi_{\min})^2M}{(1-\gamma)^3c\epsilon^2}\right)^*$                                    |
| Gradient play + log-barrier reg.,<br>softmax parameterization         | $O\left(\frac{ \mathcal{A} ^2M^2}{(1-\gamma)^4\epsilon^2}\right)$<br>[2] | $O\left(\frac{n\max \mathcal{A}_i ^2(\phi_{\max}-\phi_{\min})M^2}{(1-\gamma)^4\epsilon^2}\right)$                  |
| Natural gradient play + log-barrier reg.,<br>softmax parameterization | Unknown                                                                  | $O\left(\frac{n\max_i \mathcal{A}_i (\phi_{\max}-\phi_{\min})M^2}{(1-\gamma)^4\epsilon^2}\right)$                  |

Table 1: Summary of known convergence rate results for gradient based methods in Markov decision processes (MDPs) and MPGs respectively. The new results proved in this paper for MPGs are displayed in bold font. Complexity bounds with ‘\*’ depend on an additional assumption on the MPG (See Theorem 4 and 5). The definitions of variables  $M$  and  $c$  appearing in some bounds can be found in (5) and (11). *Note that the definition of  $M$  is slightly different from the “distribution mismatch coefficient”  $D_\infty$  defined in [2] (see more details in descriptions that follows Assumption 1). To make the complexity results more comparable, we slightly modify and re-derive the results in [2, 24, 39, 18].*

**Our contribution:** In this paper, we provide *finite time* iteration complexity results for gradient play and natural gradient play under the softmax parameterization, considering both unregularized and log-barrier regularized dynamics. We summarize the convergence rates and compare them to existing results for the direct parameterization and to the corresponding single agent cases in Table 1. These findings suggest that regularization is crucial for obtaining fast convergence to a NE under the softmax parameterization in a MPG. In Table 1, the results for the two unregularized algorithms in the multi-agent case rely on the assumption that the set of stationary policies is isolated (which is also assumed in [11] when establishing the asymptotical convergence for natural policy gradient), and the corresponding complexity bounds contain an initialization dependent factor  $c$ . By contrast, the log-barrier regularized algorithms overcome both drawbacks, but as a tradeoff, their bounds

incur a slightly worse dependence on  $|\mathcal{A}_i|$  and  $M$ . We observe numerically that the log-barrier regularized algorithms are indeed more robust against becoming trapped near undesirable non-NE stationary points. To the best of our knowledge, the finite-time iteration complexity results are the first such results for MPGs under the softmax parameterization. Though the analysis for the gradient play follows their single-agent counterparts [2, 24], the results for natural gradient play are highly non-trivial, requiring very different analysis tools which have their own merits to the literature (see Remark 2 and Remark 3 for more details on the technical novelty in the analysis). Our results also convey the following two messages. First, finding the NE of a multi-agent MPG is harder than finding the global optimum for the single-agent case, because multi-agent learning suffers greater risk of becoming trapped near undesirable stationary points. This is reflected in the dependence of the complexity bounds on  $\epsilon$  in Table 1. Second, natural gradient play outperforms gradient play counterparts, suggesting that natural gradient play captures useful information about the geometry of the parameter space that accelerates the learning process.

## 2 Problem settings

We consider an infinite time horizon  $n$ -agent stochastic game (SG, [31])  $\mathcal{M} = (N, \mathcal{S}, \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n, P, r = (r_1, \dots, r_n), \gamma, \rho)$  which is specified by an agent set  $N = \{1, 2, \dots, n\}$ , a finite state space  $\mathcal{S}$ , a finite action space  $\mathcal{A}_i$  for each agent  $i \in N$ , a transition model  $P$  (such that  $P(s'|s, a) = P(s'|s, a_1, \dots, a_n)$  is the probability of transitioning into state  $s'$  upon taking action  $a := (a_1, \dots, a_n)$  in state  $s$  where  $a_i \in \mathcal{A}_i$  is action of agent  $i$ ), a reward function  $r_i : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  for each agent  $i$ , a discount factor  $\gamma \in [0, 1)$ , and an initial state distribution  $\rho$  over  $\mathcal{S}$ . We use  $s(t) \in \mathcal{S}$  to denote the state at time step  $t$ , and  $a(t) = (a_1(t), \dots, a_n(t)) \in \mathcal{A}$  to denote the total action.

A stochastic policy  $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$  (where  $\Delta(\mathcal{A})$  is the probability simplex over  $\mathcal{A}$ ) specifies a strategy, where agents choose their actions *jointly* based on the current state in a stochastic fashion; i.e.  $\Pr(a(t)|s(t)) = \pi(a(t)|s(t))$ . A *decentralized* stochastic policy is a special subclass of stochastic policies with  $\pi = \pi_1 \times \dots \times \pi_n$ , such that  $\pi_i : \mathcal{S} \rightarrow \Delta(\mathcal{A}_i)$ , where  $\pi_i$  is agent  $i$ 's own local policy. For decentralized stochastic policies, each agent takes its action based on the current state *s independently of other agents' action choices*; i.e.,

$$\Pr(a(t)|s(t)) = \pi(a(t)|s(t)) = \prod_{i=1}^n \pi_i(a_i(t)|s(t)).$$

For notation simplicity, we define  $\pi_I(a_I|s) := \prod_{i \in I} \pi_i(a_i|s)$ , where  $I \subseteq N$  is an index set. Further, we use the notation  $-i$  to denote the index set  $N \setminus \{i\}$ . In this paper we focus on tabular softmax parameterization for a policy, where policy  $\pi_\theta = (\pi_{\theta_1}, \dots, \pi_{\theta_n})$  is parameterized by a set of parameters  $\theta = (\theta_1, \dots, \theta_n)$ , with  $\theta_i = \{\theta_{s, a_i}\}_{s \in \mathcal{S}, a_i \in \mathcal{A}_i}$ , and where

$$\pi_{\theta_i}(a_i|s) = \frac{\exp(\theta_{s, a_i})}{\sum_{a'_i} \exp(\theta_{s, a'_i})}. \quad (1)$$

We denote agent  $i$ 's total reward starting from initial states  $s(0) \sim \rho$  as:  $J_i(\theta) := \mathbb{E}_{s(0) \sim \rho} [\sum_{t=0}^{\infty} \gamma^t r_i(s(t), a(t)) | \pi_\theta, s(0) = s]$ . Agent  $i$ 's objective is to maximize its own total reward  $J_i$ . A Nash equilibrium (NE) is often used to characterize the equilibrium (a joint policy) where no agent has a unilateral incentive to deviate from it.

**Definition 1.** (Nash equilibrium) A policy  $\theta^* = (\theta_1^*, \dots, \theta_n^*)$  is called a (Markov perfect) Nash equilibrium (NE) if

$$J_i(\theta_i^*, \theta_{-i}^*) \geq J_i(\theta'_i, \theta_{-i}^*), \quad \forall \theta'_i, \quad i \in N \quad (2)$$

Further, we define the 'NE-gap' of a policy  $\theta$  to be:

$$\text{NE-gap}_i(\theta) := \sup_{\theta'_i} J_i(\theta'_i, \theta_{-i}) - J_i(\theta_i, \theta_{-i}); \quad \text{NE-gap}(\theta) := \max_i \text{NE-gap}_i(\theta).$$

A policy  $\theta$  is an  $\epsilon$ -Nash equilibrium if:  $\text{NE-gap}(\theta) \leq \epsilon$ .

We define the value function with respect to stage cost  $r_i$  as:

$$V_i^\theta(s) := \mathbb{E} [\sum_{t=0}^{\infty} \gamma^t r_i(s(t), a(t)) | \pi_\theta, s(0) = s].$$

We define agent  $i$ 's  $Q$ -function and advantage function  $Q_i^\theta, A_i^\theta : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ ,

$$Q_i^\theta(s, a) := \mathbb{E} [\sum_{t=0}^{\infty} \gamma^t r_i(s(t), a(t)) | \pi_\theta, s(0) = s, a(0) = a], \quad A_i^\theta(s, a) := Q_i^\theta(s, a) - V_i^\theta(s).$$

We further define agent  $i$ 's ‘averaged’  $Q$ -function  $\overline{Q}_i^\theta : \mathcal{S} \times \mathcal{A}_i \rightarrow \mathbb{R}$  and ‘averaged’ advantage-function  $\overline{A}_i^\theta : \mathcal{S} \times \mathcal{A}_i \rightarrow \mathbb{R}$  as:

$$\overline{Q}_i^\theta(s, a_i) := \sum_{a_{-i}} \pi_{\theta_{-i}}(a_{-i}|s) Q_i^\theta(s, a_i, a_{-i}), \quad \overline{A}_i^\theta(s, a_i) := \sum_{a_{-i}} \pi_{\theta_{-i}}(a_{-i}|s) A_i^\theta(s, a_i, a_{-i}).$$

Finally, define the *discounted state visitation distribution*  $d_\theta$  of a policy  $\pi_\theta$  given an initial state distribution  $\rho$  as:

$$d_\theta(s) := \mathbb{E}_{s(0) \sim \rho} (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \Pr^\theta(s(t) = s | s(0)), \quad (3)$$

where  $\Pr^\theta(s(t) = s | s(0))$  is the state visitation probability that  $s(t) = s$  when executing  $\pi_\theta$  starting at state  $s(0)$ . From the policy gradient theorem [34], we have that (proof given in Appendix B):

$$\frac{\partial J_i(\theta)}{\partial \theta_{s, a_i}} = \frac{1}{1 - \gamma} d_\theta(s) \pi_{\theta_i}(a_i | s) \overline{A}_i^\theta(s, a_i). \quad (4)$$

For the remainder of the paper, we make the following assumptions on the stochastic games we study.

**Assumption 1.** *The stochastic game  $\mathcal{M}$  satisfies:  $\inf_\theta \min_{s \in \mathcal{S}} d_\theta(s) > 0$ .*

Assumption 1 requires that every state is visited with positive probability for any policy, which is a standard assumption for convergence proofs in the RL literature (e.g. [2, 24]). We will use  $M$  to denote the following quantity

$$M := \sup_\theta \max_s \frac{1}{d_\theta(s)}. \quad (5)$$

Note that  $M$  can be viewed as a measure of exploration sufficiency in the stochastic game, which is slightly different from the ‘distributional mismatch coefficient’ introduced in [2] defined by  $\sup_{\theta, \theta'} \max_s \frac{d_{\theta'}(s)}{d_\theta(s)}$ ; however, both can be upper bounded by  $\max_s \frac{1}{(1-\gamma)\rho(s)}$ .

We primarily focus on the following subclass of stochastic games in this paper:

**Definition 2.** *A stochastic game is called a Markov potential game (MPG, [37, 23, 39, 18, 26]) if there exists a potential function  $\phi : \mathcal{S} \times \mathcal{A}_1 \times \dots \times \mathcal{A}_n \rightarrow \mathbb{R}$  such that for any agent  $i$  and any pair of policy parameters  $(\theta'_i, \theta_{-i}), (\theta_i, \theta_{-i})$ :*

$$\begin{aligned} & \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t r_i(s(t), a(t)) \mid \pi = (\theta'_i, \theta_{-i}), s(0) = s \right] - \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t r_i(s(t), a(t)) \mid \pi = (\theta_i, \theta_{-i}), s(0) = s \right] \\ &= \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t \phi(s(t), a(t)) \mid \pi = (\theta'_i, \theta_{-i}), s(0) = s \right] - \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t \phi(s(t), a(t)) \mid \pi = (\theta_i, \theta_{-i}), s(0) = s \right], \quad \forall s. \quad (6) \end{aligned}$$

Without loss of generality, we assume that  $\phi_{\min} \leq \phi(s, a) \leq \phi_{\max}$  for all  $(s, a)$ . The definition of MPG is a generalization of the notion potential game in the one-shot setting [28]. Note that identical reward game where agents share a same reward function naturally satisfies the above condition and serves as one important special case of MPG. For non-identical reward settings, [23, 12] found that continuous MPGs can model applications such as the great fish war [19], the stochastic lake game [10], medium access control [23] etc. For tabular MPGs, [39, 18] also discuss necessary/sufficient conditions that implies a MPG, as well as its application and counterexamples.

Given a MPG, we define the *total potential function*  $\Phi$  as:

$$\Phi(\theta) := \mathbb{E}_{s(0) \sim \rho} \left[ \sum_{t=0}^{\infty} \gamma^t \phi(s(t), a(t)) \mid \pi_\theta, s(0) = s \right].$$

Given the property in (6), it is straightforward to verify that the NE condition (2) is equivalent to  $\Phi(\theta_i^*, \theta_{-i}^*) \geq \Phi(\theta'_i, \theta_{-i}^*), \forall \theta'_i, i \in N$  and that for the policy gradient,  $\frac{\partial J_i(\theta)}{\partial \theta_{s, a_i}} = \frac{\partial \Phi(\theta)}{\partial \theta_{s, a_i}}$  for all  $i, s, a_i$ .

**Remark 1 (Differences between MPG and single-agent/centralized MDP).** Because of the existence of the total potential function  $\Phi$ , it is natural to ask whether MPG renders the multi-agent policy gradient similar to single agent policy gradient and thus results and analysis tools developed for single agent policy gradient in e.g., [2, 24] would be easily extended to the multiagent case. Unfortunately, this is not the case. To illustrate how it differs from single-agent/centralized case, we can focus on the special type of MPGs where every agent has the same reward function, namely the identical interest case. In the single agent/centralized case, there is a unique global optimal solution which corresponds to the convergent stationary policy. However, in the multiagent case, even if the rewards are identical, because the policy is *decentralized*, i.e., agents taking *independent* policies  $\pi := \pi_1 \times \dots \times \pi_n$ , we loose the connection between stationary policies and optimal policies. As we shown later, the

convergent stationary policies are Nash equilibria, which are unfortunately non-unique even for the identical interest case. Moreover, a key condition that is used in establishing the convergence rate, Łojasiewicz condition (Lemma 1), is also much weaker for the multiagent case compared to single agent [24]: the left hand side is the *Nash gap*  $\max_{i, \theta_i^*} \Phi(\theta_i^*, \theta_{-i}) - \Phi(\theta)$  instead of the *optimality gap*  $\max_{\theta^*} \Phi(\theta^*) - \Phi(\theta)$ . Note that zero Nash gap does not imply zero optimality gap, as there exists *many* NEs of *different* values. These differences disable many proof technique used for single agent case and make the analysis harder and lead to different performance results, as demonstrated in the rest of the paper.

### 3 Relationship between first order stationary point and Nash equilibrium

Before studying convergence performance of gradient play algorithms, it is important to first understand the relationship between the stationary points and the NEs. Unfortunately, equivalence cannot be established in this setting. Standard optimization theory guarantees that all NEs are stationary points, but unfortunately not vice versa. Under softmax parameterization, there exist non-NE stationary points. For example, from the gradient formulation (4), it can be shown that any non-NE deterministic policies are also stationary points. However, the notion of NE and stationarity are indeed closely related. This section aims to characterize some differences between NE and non-NE stationary points. This differentiation of the NE and non-NE stationary points is established by the non-uniform Łojasiewicz condition (also known as gradient domination) for stochastic games.

**Lemma 1.** (*Non-uniform Łojasiewicz inequality; proof given in Appendix D*) *Define*

$$M(\theta) := \max_s \frac{1}{d_\theta(s)}, \quad c(\theta) := \min_s \sum_{a_i^* \in \arg\max_{a_i} \overline{Q}_i^\theta(s, a_i)} \pi_{\theta_i}(a_i^* | s). \quad (7)$$

Then we have that

$$\text{NE-gap}_i(\theta) \leq \frac{\sqrt{|\mathcal{A}_i|} M(\theta)}{c(\theta)} \|\nabla_{\theta_i} J_i(\theta)\|_2.$$

The Łojasiewicz condition (gradient domination) implies that the NE-gap of a policy can be bounded by the norm of its gradient, whereas the term ‘non-uniform’ refers to the factor  $\frac{\sqrt{|\mathcal{A}_i|} M(\theta)}{c(\theta)}$ , which cannot be bounded uniformly for all  $\theta$ . The counterpart of Lemma 1 for a single-agent MDP was first introduced in [24, Lemma 8]. One major difference between Lemma 1 and [24, Lemma 8] is how  $c(\theta)$  is defined. In [24],  $c(\theta) := \min_s \pi_\theta(a^*(s) | s)$ , where  $a^*(s)$  is the optimal action on state  $s$  (i.e.,  $a^* = \arg\max_a Q^*(s, a)$ ), whereas in MPG, because there’s no globally defined  $Q^*$ , the  $a_i^*$  in (7) is chosen as the *greedy* optimal action of the *current* averaged  $Q$ -function (i.e.,  $a_i^* \in \arg\max_{a_i} \overline{Q}_i^\theta(s, a_i)$ ).

Note that because  $c(\theta)$  on the denominator can be zero for certain policies (e.g. one can verify that any non-NE deterministic policy have  $c(\theta) = 0$ ), which implies that a  $\theta$  with gradient norm close to zero is not necessarily near a NE. Given this observation, we could differentiate the non-NE stationary points with NEs by whether  $c(\theta^*)$  equals to zero, which is formally stated in the following lemma:

**Lemma 2.** (*Proof given in Appendix D*) *Suppose  $\theta^*$  is a stationary point, i.e.  $\|\nabla \Phi(\theta^*)\| = 0$ , then  $\theta^*$  is a NE if and only if  $c(\theta^*) = 1$ ,  $\theta^*$  is not a NE if and only if  $c(\theta^*) = 0$ .*

### 4 Unregularized gradient play

We first investigate the convergence to NE for gradient and natural gradient play, respectively. Under the softmax parameterization, the two schemes are given by

$$\text{Gradient Play:} \quad \theta_i^{(t+1)} = \theta_i^{(t)} + \eta \nabla_{\theta_i} J_i(\theta_i^{(t)}) \quad (8)$$

$$\text{Natural Gradient Play:} \quad \theta_i^{(t+1)} = \theta_i^{(t)} + \eta F_i(\theta^{(t)})^\dagger \nabla_{\theta_i} J_i(\theta_i^{(t)}) \quad (9)$$

where  $\dagger$  denotes the Moore-Penrose inverse and  $F_i(\theta)$  is the Fisher information matrix for  $\pi_{\theta_i}$ :

$$F_i(\theta) := \mathbb{E}_{s \sim d_\theta(\cdot)} \mathbb{E}_{a_i \sim \pi_{\theta_i}(\cdot | s)} \left[ \nabla_{\theta_i} \log \pi_{\theta_i}(a_i | s) \nabla_{\theta_i} \log \pi_{\theta_i}(a_i | s)^\top \right].$$

For notational simplicity, we abbreviate the variables  $d_{\theta^{(t)}}$ ,  $A_i^{\theta^{(t)}}$  and  $\overline{A_i^{\theta^{(t)}}}$  as  $d^{(t)}$ ,  $A_i^{(t)}$  and  $\overline{A_i^{(t)}}$  respectively; and denote  $\pi_{\theta^{(t)}}(a|s)$  and  $\pi_{\theta^{(t)}}(a_i|s)$  as  $\pi^{(t)}(a|s)$  and  $\pi_i^{(t)}(a_i|s)$  respectively. For the softmax parameterization, we can establish the equivalence of natural gradient play and soft Q-learning [13], formally stated in the following lemma.

**Lemma 3.** (Proof given in Appendix C) *Natural gradient play is equivalent to*

$$\pi_i^{(t+1)}(a_i|s) \propto \pi_i^{(t)}(a_i|s) \exp\left(\frac{\overline{A_i^{(t)}}(s, a_i)}{(1-\gamma)}\right) \quad (10)$$

**Asymptotic convergence to Nash Equilibrium.** As stated in Section 3, there exist stationary points that are not NEs. It is not immediately obvious why running gradient methods can avoid converging to these points, thus before studying convergence rate to NE, it is necessary to first examine whether asymptotic convergence holds. Moreover, the asymptotic convergence result is used to establish the finite time convergence rate results later (see the subsection 4.1).

**Theorem 4.** (Proof given in Appendix E) *Suppose Assumption 1 holds and that the stationary policies are isolated, gradient play (8) with  $\eta \leq \frac{(1-\gamma)^3}{6n}$  guarantees that  $\lim_{t \rightarrow +\infty} \theta^{(t)} = \theta^{(\infty)}$ , where  $\theta^{(\infty)}$  is a NE. The same argument also holds for natural gradient play (10) with  $\eta \leq \frac{(1-\gamma)^2}{2n(\phi_{\max} - \phi_{\min})}$ .*

The proof of Theorem 4 resembles the technique used in [2] for the single agent case, where the additional assumption on the isolated stationary policies is introduced due to some specific technical difficulties encountered in multi-agent learning (see more discussion in Appendix E, which is also introduced in [11] for establishing the asymptotic convergence of NPG. We believe it is a conservative condition for ensuring the asymptotic convergence. It remains an interesting open question to establish convergence without this assumption.

#### 4.1 Finite time convergence rate

This section considers finite time convergence rate for gradient play and natural gradient play. Corresponding results for the single-agent setting can be found in [24] (for gradient play) and [2, 16, 25] (for natural gradient play). Some aspects of these analyses can be carried over to the multi-agent MPG setting; however, as will be discussed later, there are several fundamental differences that make the multi-agent case more challenging.

Our convergence results rely on the observation from Section 3 and the asymptotic convergence to NE. Combining Theorem 4 and Lemma 2, we know that  $c(\theta^{(t)})$  asymptotically converges to 1 for (natural) gradient play, and since  $c(\theta^{(t)}) > 0$  for any softmax policy (because  $\pi_{\theta^{(t)}}(a_i|s) > 0$ ),

$$c := \inf_t c(\theta^{(t)}) > 0. \quad (11)$$

We are now ready to give formal convergence rates for gradient and natural gradient play respectively.

**Theorem 5.** (Gradient play and natural gradient play; proof given in F) *Suppose Assumption 1 holds and that the stationary policies are isolated, gradient play (8) with  $\eta = \frac{(1-\gamma)^3}{6n}$  will guarantee that for all  $T$ ,*

$$\frac{\sum_{t=0}^{T-1} \text{NE-gap}(\theta^{(t)})^2}{T} \lesssim O\left(\frac{n \max_i |\mathcal{A}_i| (\phi_{\max} - \phi_{\min}) M^2}{(1-\gamma)^4 c^2 T}\right), \quad (12)$$

*Natural gradient play (10) with  $\eta = \frac{(1-\gamma)^2}{2n(\phi_{\max} - \phi_{\min})}$  will guarantee that for all  $T$ ,*

$$\frac{\sum_{t=0}^{T-1} \text{NE-gap}(\theta^{(t)})^2}{T} \lesssim O\left(\frac{n(\phi_{\max} - \phi_{\min})^2 M}{(1-\gamma)^3 c T}\right). \quad (13)$$

Here  $O(\cdot)$  hides constant factors,  $M$  and  $c$  are defined as in (5) and (11), respectively.

**Remark 2 (Proof sketch and novelty).** The proof for gradient play is relatively straightforward from the non-uniform Łojasiewicz inequality and standard non-convex optimization results, which we refer readers to the appendix for more details. However, the proof for natural gradient play is more involved and existing analysis on NPG cannot be generalized to this setting. For single-agent MDP, the analysis on NPG leverages the unique existence of optimal value function  $V^*$  so that similar analysis for mirror-descent can also carry over to NPG analysis, and thus obtain dimension

free convergence. However, in the multi-agent setting, there's no well-defined  $V^*$  as NEs can be non-unique with different potential values, thus, we need to further deploy additional structures of the total potential function  $\Phi$ . Our analysis rely on the sufficient ascent lemma (Lemma 20) that lower bounds the ascent amount  $\Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)})$  for each natural gradient step (we would like to further note that this sufficient ascent lemma cannot be trivially obtained by the smoothness of  $\Phi$ ). Then, we further lower bound the ascent amount in terms of NE-gap (Lemma 21). Lastly, the theorem follows by conducting standard telescoping techniques.

**Discussion on  $\frac{1}{c}$ :** The complexity results in Theorem 5 both depend on  $\frac{1}{c}$ . However, this term can become arbitrarily large. In fact, [20] show that  $c$  can be exponentially small in terms of the number of states  $|\mathcal{S}|$  for a general finite MDP, even under uniform initialization, hence convergence can be very slow. This conclusion is also confirmed by numerical evidence. As pointed out by [24], even for single agent settings, policy gradient can get stuck at regions with small gradient yet far from being global optimal. Similar or even worse phenomena can be observed for multi-agent MPG, as shown in Figure 1(a)-(c): even for a single state game ( $|\mathcal{S}| = 1$ ) with uniform initialization, unregularized gradient based algorithms can still enter regions with a relatively large NE-gap while the gradient norm and  $c(\theta)$  are close to zero.

**More comparison with learning for single-agent MDP:** For gradient play, we have established an iteration complexity of  $O\left(\frac{n \max_i |\mathcal{A}_i| (\phi_{\max} - \phi_{\min}) M^2}{(1-\gamma)^4 c^2 \epsilon^2}\right)$  to find an  $\epsilon$ -NE, whereas [24] show a complexity of  $O\left(\frac{(\phi_{\max} - \phi_{\min}) M^2}{(1-\gamma)^4 c^2 \epsilon}\right)$  to reach an  $\epsilon$ -global optimum for policy gradient in a single agent MDP. The dependence on  $\frac{1}{c}$  is better in the single agent case because of the existence of a global optimal policy  $\pi^*$  and optimal total reward  $V^*$ , which justify the definition of optimality gap  $\delta_t = V(\theta^{(t)}) - V^*$ . This, combined with the non-uniform Łojasiewicz condition which bounds  $\delta_t$  by the gradient norm, allows one to use techniques from convex smooth analysis to show that  $\delta_t$  is on the scale of  $\frac{1}{t}$ . By contrast, for multi-agent learning, there can be multiple NEs with different values, hence  $\delta_t$  is ill-defined. Further, note that the NE-gap is different from the optimality gap, hence gradient ascent no longer guarantees monotonic decreasing of NE-gap (Figure 1(a)), and we can only exploit non-convex optimization techniques that yield  $O(\frac{1}{c^2})$  complexities.

For the same reason, the rate of convergence we obtain for natural gradient play is  $O\left(\frac{n(\phi_{\max} - \phi_{\min})^2 M}{(1-\gamma)^3 c \epsilon^2}\right)$ , which is worse than the dimension free convergence rate of  $O\left(\frac{1}{(1-\gamma)^2 \epsilon}\right)$  given in [2] for single-agent MDPs. (A better exponential convergence rate for natural PG has also been proved in [16, 25] with the exponential factor being problem dependent.) Nevertheless, the dependence on  $\frac{1}{c}$ ,  $\frac{1}{1-\gamma}$  and  $M$  is better than gradient play, suggesting that the preconditioning of natural gradient play at least partially captures the geometry of the parameter space. We also note that the quadratic dependence on  $(\phi_{\max} - \phi_{\min})$  might be a proof artifact. It remains an open question whether this can be reduced to a linear dependence.

## 5 Gradient play with log-barrier regularization

The previous section has shown that, for unregularized objectives, the convergence rate for gradient based algorithms depends on a factor  $\frac{1}{c}$  that can be arbitrarily large for bad initializations. This motivates us to investigate regularization, in hopes of removing the dependence on  $\frac{1}{c}$ . For this purpose, we consider log-barrier regularization:

$$\tilde{J}_i(\theta) = J_i(\theta) + \lambda \sum_{s, a_i} \log \pi_{\theta_i}(a_i|s).$$

Define:

$$\tilde{\Phi}(\theta) = \Phi(\theta) + \lambda \sum_{i=1}^n \sum_{s, a_i} \log \pi_{\theta_i}(a_i|s). \quad (14)$$

It is not hard to verify that the gradient with respect to  $J_i$  is:

$$\frac{\partial \tilde{J}_i(\theta)}{\partial \theta_{s, a_i}} = \frac{\partial \tilde{\Phi}(\theta)}{\partial \theta_{s, a_i}} = \frac{1}{1-\gamma} d_\theta(s) \pi_{\theta_i}(a_i|s) \overline{A}_i^\theta(s, a_i) + \lambda - \lambda |\mathcal{A}_i| \pi_{\theta_i}(a_i|s).$$

**Discussion on the choice of the regularizer:** Before analyzing the resulting algorithm we first discuss the motivation for this regularizer. First, note that for each agent, the additional regularizer only depends on an agent’s own local policy, which is desirable for multiagent RL. As an alternative, one might impose regularization by choosing

$$\tilde{\Phi}(\theta) = \Phi(\theta) + \lambda \mathbb{E}_{s \sim d_\theta(\cdot)} \sum_{i=1}^n \sum_{a_i} \log \pi_{\theta_i}(a_i|s);$$

i.e., so that the regularization weight imposed on a state  $s$  depends on the state visitation probability  $d_\theta(s)$ . However, in this case the gradient of the  $i$ -th agent  $\nabla_{\theta_i} \tilde{\Phi}(\theta)$  will not only depend on its own policy parameter  $\theta_i$ , but also on other parameters of the other agents’ policies  $\theta_{-i}$ . Thus, running gradient based algorithms with such a regularization scheme can no longer be executed in a fully decentralized manner using local policy information. Therefore, we prefer regularization (14) which does not depend on  $d_\theta(s)$ . Secondly, we adopt the log-barrier instead of entropy regularization due to technical rather than practical considerations. Although entropy regularization achieves fast exponential convergence in single agent learning [7, 24], for multi-agent learning, we haven’t been able to obtain results as strong as the log-barrier regularization. Intuitively, the log-barrier regularized gradient field repels the trajectory from regions with small  $\pi_i(a_i|s)$  values (where the geometry becomes close to singular) more strongly, which enables us to obtain our current analysis. However, we emphasize that our result does not imply that log-barrier is better than entropy regularization in practice. It remains future work to determine whether entropy regularization, or other methods such as trust region based methods, can achieve the same, or even better convergence rates.

## 5.1 Gradient play

We first consider gradient play algorithm, i.e.,

$$\theta_i^{(t+1)} = \theta_i^{(t)} + \eta \nabla_{\theta_i} \tilde{J}_i(\theta^{(t)}). \quad (15)$$

Fortunately, similar analysis from [2] for single-agent MDP can be generalized to MPG with slight modifications. Here we only state the result and defer the proof to Appendix G.1.

**Theorem 6.** *Under Assumption 1, for  $\eta = \frac{(1-\gamma)^3}{6n+2\lambda \max_i |\mathcal{A}_i|(1-\gamma)^3}$ , and  $\lambda = \frac{\epsilon}{M \max_i |\mathcal{A}_i|}$ , let  $\theta^{(0)}$  be the uniform random policy, i.e.,  $\theta^{(0)} = \mathbf{0}$ , then running gradient play (15) for  $T$  steps, where  $T \gtrsim O\left(\frac{n \max_i |\mathcal{A}_i|^2 (\phi_{\max} - \phi_{\min}) M^2}{(1-\gamma)^4 \epsilon^2}\right)$  will guarantee that  $\min_{0 \leq t \leq T-1} \text{NE-gap}(\theta^{(t)}) \leq \epsilon$ .*

Note that compared to the unregularized case in Theorem 5, it only requires Assumption 1, while the convergence rate is accelerated by eliminating the dependence on  $\frac{1}{c}$ . However, as a (worthy) tradeoff, the dependence on the action space size  $\max_i |\mathcal{A}_i|$  now becomes quadratic. The key reason for these differences is that log-barrier regularization assures that any policy with sufficiently small gradient norm cannot be close to the boundary of the probability simplex where the non-uniform Łojasiewicz constant is large.

## 5.2 Natural gradient play

In the unregularized setting, we have seen that natural gradient play enjoys a better convergence rate than gradient play, which motivates us to consider whether a similar advantage still holds for the regularized case. In this section we consider natural gradient play

$$\theta_i^{(t+1)} = \theta_i^{(t)} + \eta F_i(\theta^{(t)})^\dagger \nabla_{\theta_i} \tilde{J}_i(\theta^{(t)}), \quad (16)$$

which is equivalent to (see the proof in Appendix C)

$$\pi_i^{(t+1)}(a_i|s) \propto \pi_i^{(t)}(a_i|s) \exp\left(\frac{\eta}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) + \frac{\eta \lambda}{d^{(t)}(s) \pi_i^{(t)}(a_i|s)} - \frac{\eta \lambda |\mathcal{A}_i|}{d^{(t)}(s)}\right). \quad (17)$$

**Theorem 7.** *(Proof given in Appendix G.2) Under Assumption 1, for*

$$\eta = \min\left\{\frac{1}{15\left(\frac{1}{(1-\gamma)^2} + \lambda |\mathcal{A}_i| M\right)}, \frac{1}{4\left(4\lambda \max_i |\mathcal{A}_i| M^2 + \frac{4M}{(1-\gamma)^2} + \frac{3nM}{(1-\gamma)^3}\right)}\right\}, \text{ the natural gradient play (17)}$$



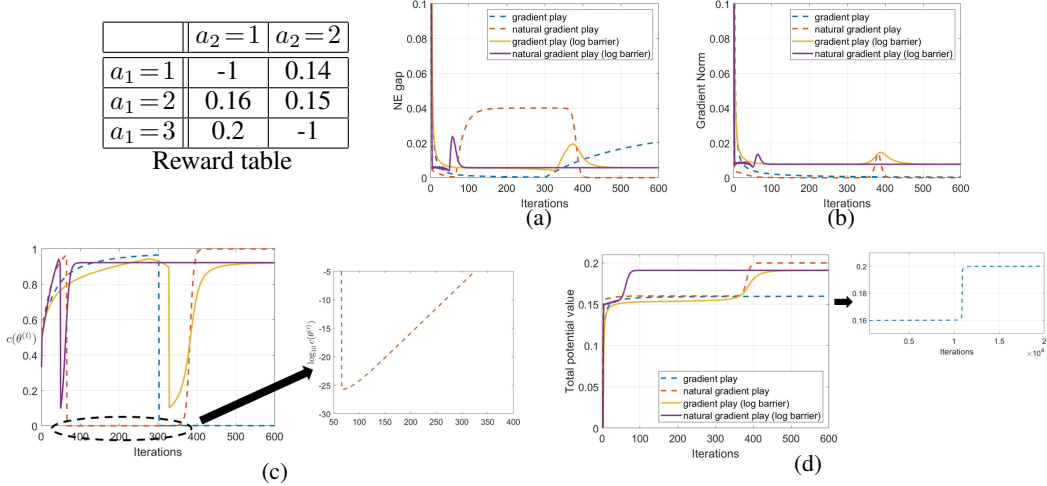


Figure 1: We consider a two-player identical reward matrix game as shown in the reward table. We run gradient play and natural gradient play (with and without log-barrier regularization) with initial policies being the uniform distribution (i.e.,  $\pi_1 = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$ ,  $\pi_2 = [\frac{1}{2}, \frac{1}{2}]$ ). The subfigures (a)-(d) show how the NE-gap( $\theta^t$ ),  $\|\nabla_{\theta}\Phi(\theta^t)\|_2$ ,  $c(\theta^t)$  (defined in (7)) and  $\Phi(\theta^t)$  change with each iteration respectively. In Figure (c), we zoom in on the  $\log_{10} c(\theta)$  factor for natural gradient play. In Figure (d), we also zoom out the trajectory for running gradient play to iteration  $2 \times 10^4$ . Here the step sizes were chosen to be  $\eta = 5$  while the regularization weight  $\lambda$  was chosen to be  $\lambda = 0.003$ . In consideration of numerical stability issues, we truncate the update step of natural gradient play with log-barrier regularization by a maximum absolute value of 1 for each entry. For more numerical results and corresponding analysis see Appendix A.

will guarantee that  $\frac{\sum_{t=0}^{T-1} \text{NE-gap}(\theta^t)}{T} \leq \frac{9(\tilde{\Phi}(\theta^T) - \tilde{\Phi}(\theta^0))}{\eta\lambda T} + \lambda \max_i |\mathcal{A}_i| M$ , Further, by setting  $\lambda = \frac{\epsilon}{2 \max_i |\mathcal{A}_i| M}$ ,  $\theta^0 = \mathbf{0}$ , for  $T \gtrsim O\left(\frac{n \max_i |\mathcal{A}_i| (\phi_{\max} - \phi_{\min}) M^2}{(1-\gamma)^4 \epsilon^2}\right)$ , we have  $\frac{\sum_{t=0}^{T-1} \text{NE-gap}(\theta^t)}{T} \leq \epsilon$ .

**Remark 3. (Proof sketch and novelty)** As also stated for unregularized natural gradient play, there's no direct analysis tools we could borrow from literature for the analysis of natural gradient play. Our analysis depends on two key lemmas. The first is a sufficient ascent lemma on  $\tilde{\Phi}(\theta^{t+1}) - \tilde{\Phi}(\theta^t)$  for each natural gradient step (Lemma 26). Another key lemma (Lemma 24) states that the algorithm implicitly ensures that the policies never go near the boundary of the probability simplex, i.e., it can be uniformly lower-bounded by  $\pi_i^{(t)}(a_i|s) \geq \frac{\lambda}{4(\lambda|\mathcal{A}_i|M + \frac{1}{(1-\gamma)^2})}$ ,  $\forall t$ . Combining the two lemmas, it

can be concluded that the ascent value  $\tilde{\Phi}(\theta^{t+1}) - \tilde{\Phi}(\theta^t)$  can be bounded by NE-gap( $\theta^t$ ) plus a  $\lambda \max_i |\mathcal{A}_i| M$  bias term (Lemma 27 and 28), thus the proof is finished by standard telescoping technique and choosing an appropriate  $\lambda$ .

Compared with gradient play, natural gradient play manages to reduce the time complexity by a  $\max_i |\mathcal{A}_i|$  factor. Further, gradient play only guarantees the minimal NE-gap smaller than  $\epsilon$ , while natural gradient play guarantees the average NE-gap along the trajectory smaller than  $\epsilon$ . To the best of our knowledge, this is the best time complexity bound for the softmax parameterization in a MPG.

## 6 An Illustrative example

This section aims to gain a better understanding of the four gradient play algorithms, (8), (10), (15), and (17). To better justify our theoretical results and provide additional insights, we choose a carefully designed simple two-player game so that our theoretical results can be easily revealed from the empirical observations. However the four algorithms also works for settings with more agents.<sup>1</sup> Due to space limits, we defer the simulation with more agents in Appendix A.

The reward table as well as the performance of the four algorithms are shown in Figure 1. Comparing the log-barrier regularized algorithms to the unregularized counterparts, one can see that the regularized dynamics converge faster but with a bias induced by the regularizer. This finding corroborates

<sup>1</sup>Code can be found in <https://github.com/DianYu420376/NeurIPS2022-softmax-MPG>

the analyses given in Theorem 6 and 7. By contrast, the unregularized dynamics are able to find a policy with zero NE-gap asymptotically, but tend to get stuck in regions where  $c(\theta^{(t)})$  is very close to zero, as illustrated in Fig 1(a)(b). Specifically unregularized natural gradient play gets stuck around iteration 100-400 in a region where the gradient norm and  $c(\theta^{(t)})$  are both close to zero while the NE-gap is not. This corroborates the finding in Lemma 1. Similar behavior can be observed for gradient play if we keep running the algorithm. In comparing the natural gradient play to gradient play algorithms, natural gradient play generally converges faster, which matches with our complexity analysis. However, natural gradient play with log-barrier regularization can suffer from numerical instability due to the  $1/\pi_i^{(t)}(a_i|s)$  term in the exponential factor. In this case, the stepsize needs to be chosen carefully. To bypass the numerical instability, we truncate the update step of natural gradient play with log-barrier regularization by a maximum absolute value of 1 for each entry.

## 7 Discussions and conclusions

We have established finite time iteration complexity bounds for gradient and natural gradient play under the softmax parameterization, considering both unregularized and log-barrier regularized dynamics, in the Markov potential game setting. To our best knowledge, these are the first finite time global convergence results for softmax gradient play for MPGs. However, our work suffers from the following limitations: firstly, the paper mainly focuses on MPG settings, which limits its application to general-sum Markov games; secondly, convergence results for the unregularized case relies on an extra assumption that the stationary points are isolated; thirdly, for the regularized case, we consider log-barrier regularization, which is admittedly a stronger regularization compared with entropy regularization which is more frequently used in practice. Some limitations are due to technical challenges, some might be caused by the fundamental difficulties of multi-agent learning. It remains interesting open questions to sharpen the analysis, derive similar or better bounds for other regularizations, and to develop more fundamental understandings of multi-agent learning.

## Acknowledgment

Runyu Zhang is supported by NSF AI institute: 2112085, ONR YIP: N00014-19-1-2217, NSF CNS: 2003111 and NSF CPS: 2038603.

## References

- [1] P.-A. Absil, R. Mahony, and B. Andrews. Convergence of the iterates of descent methods for analytic cost functions. *SIAM Journal on Optimization*, 16(2):531–547, 2005.
- [2] A. Agarwal, S. M. Kakade, J. D. Lee, and G. Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift, 2020.
- [3] S.-i. Amari. *Differential-geometrical methods in statistics*, volume 28. Springer Science & Business Media, 2012.
- [4] Y. Bai and C. Jin. Provable self-play algorithms for competitive reinforcement learning. In *International Conference on Machine Learning*, pages 551–560. PMLR, 2020.
- [5] M. Bowling and M. Veloso. An analysis of stochastic game theory for multiagent reinforcement learning. Technical report, Carnegie-Mellon Univ Pittsburgh Pa School of Computer Science, 2000.
- [6] L. Buşoniu, R. Babuška, and B. De Schutter. Multi-agent reinforcement learning: An overview. *Innovations in multi-agent systems and applications-1*, pages 183–221, 2010.
- [7] S. Cen, C. Cheng, Y. Chen, Y. Wei, and Y. Chi. Fast global convergence of natural policy gradient methods with entropy regularization. *Operations Research*, 2021.
- [8] R. Claes, T. Holvoet, and D. Weyns. A decentralized approach for anticipatory vehicle routing using delegate multiagent systems. *IEEE Transactions on Intelligent Transportation Systems*, 12(2):364–373, 2011.

- [9] C. Daskalakis, D. J. Foster, and N. Golowich. Independent policy gradient methods for competitive reinforcement learning. *arXiv preprint arXiv:2101.04233*, 2021.
- [10] W. D. Dechert and S. O’Donnell. The stochastic lake game: A numerical solution. *Journal of Economic Dynamics and Control*, 30(9-10):1569–1587, 2006.
- [11] R. Fox, S. McAleer, W. Overman, and I. Panageas. Independent natural policy gradient always converges in markov potential games. *CoRR*, abs/2110.10614, 2021.
- [12] D. González-Sánchez and O. Hernández-Lerma. *Discrete-time stochastic control and dynamic potential games: the Euler–Equation approach*. Springer Science & Business Media, 2013.
- [13] T. Haarnoja, H. Tang, P. Abbeel, and S. Levine. Reinforcement learning with deep energy-based policies. In *International Conference on Machine Learning*, pages 1352–1361. PMLR, 2017.
- [14] P. Iñigo-Blasco, F. Diaz-del Rio, M. C. Romero-Tertero, D. Cagigas-Muñiz, and S. Vicente-Diaz. Robotics software frameworks for multi-agent robotic systems development. *Robotics and Autonomous Systems*, 60(6):803–821, 2012.
- [15] S. M. Kakade and J. Langford. Approximately optimal approximate reinforcement learning. In C. Sammut and A. G. Hoffmann, editors, *Machine Learning, Proceedings of the Nineteenth International Conference (ICML 2002), University of New South Wales, Sydney, Australia, July 8-12, 2002*, pages 267–274. Morgan Kaufmann, 2002.
- [16] S. Khodadadian, P. R. Jhunjhunwala, S. M. Varma, and S. T. Maguluri. On the linear convergence of natural policy gradient algorithm. *arXiv preprint arXiv:2105.01424*, 2021.
- [17] M. Lanctot, V. Zambaldi, A. Gruslys, A. Lazaridou, K. Tuyls, J. Pérolat, D. Silver, and T. Graepel. A unified game-theoretic approach to multiagent reinforcement learning. *arXiv preprint arXiv:1711.00832*, 2017.
- [18] S. Leonardos, W. Overman, I. Panageas, and G. Piliouras. Global convergence of multi-agent policy gradient in markov potential games. *arXiv preprint arXiv:2106.01969*, 2021.
- [19] D. Levhari and L. Mirman. The great fish war: An example using a dynamic cournot-nash solution. *Bell Journal of Economics*, 11(1):322–334, 1980.
- [20] G. Li, Y. Wei, Y. Chi, Y. Gu, and Y. Chen. Softmax policy gradient methods can take exponential time to converge. *CoRR*, abs/2102.11270, 2021.
- [21] M. L. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994.
- [22] J. Liu and J. Wu. *Multiagent robotic systems*. CRC press, 2018.
- [23] S. V. Macua, J. Zazo, and S. Zazo. Learning parametric closed-loop policies for markov potential games. *CoRR*, abs/1802.00899, 2018.
- [24] J. Mei, C. Xiao, C. Szepesvari, and D. Schuurmans. On the global convergence rates of softmax policy gradient methods. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6820–6829. PMLR, 13–18 Jul 2020.
- [25] J. Mei, B. Dai, C. Xiao, C. Szepesvari, and D. Schuurmans. Understanding the effect of stochasticity in policy optimization. *Advances in Neural Information Processing Systems*, 34, 2021.
- [26] D. Mguni. Stochastic potential games. *arXiv preprint arXiv:2005.13527*, 2020.
- [27] D. Mguni, Y. Wu, Y. Du, Y. Yang, Z. Wang, M. Li, Y. Wen, J. Jennings, and J. Wang. Learning in nonzero-sum stochastic games with potentials. *arXiv preprint arXiv:2103.09284*, 2021.
- [28] D. Monderer and L. S. Shapley. Potential games. *Games and economic behavior*, 14(1): 124–143, 1996.
- [29] C. R. Rao. Information and the accuracy attainable in the estimation of statistical parameters. In *Breakthroughs in statistics*, pages 235–247. Springer, 1992.

- [30] M. Roscia, M. Longo, and G. C. Lazaroiu. Smart city by multi-agent systems. In *2013 International Conference on Renewable Energy Research and Applications (ICRERA)*, pages 371–376. IEEE, 2013.
- [31] L. S. Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10): 1095–1100, 1953.
- [32] Y. Shoham, R. Powers, and T. Grenager. Multi-agent reinforcement learning: a critical survey. Technical report, Technical report, Stanford University, 2003.
- [33] Z. Song, S. Mei, and Y. Bai. When can we learn general-sum markov games with a large number of players sample-efficiently?, 2021.
- [34] R. S. Sutton, D. A. McAllester, S. P. Singh, Y. Mansour, et al. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, volume 99, pages 1057–1063. Citeseer, 1999.
- [35] N. Tao, J. Baxter, and L. Weaver. A multi-agent, policy-gradient approach to network routing. In *In: Proc. of the 18th Int. Conf. on Machine Learning*. Citeseer, 2001.
- [36] A. G. Ventre, A. Maturo, Š. Hořková-Mayerová, and J. Kacprzyk. *Multicriteria and Multiagent Decision Making with Applications to Economics and Social Sciences*, volume 305. Springer, 2013.
- [37] S. Zazo, S. V. Macua, M. Sánchez-Fernández, and J. Zazo. Dynamic potential games with constraints: Fundamentals and applications in communications. *IEEE Transactions on Signal Processing*, 64(14):3806–3821, 2016.
- [38] K. Zhang, Z. Yang, and T. Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *arXiv preprint arXiv:1911.10635*, 2019.
- [39] R. Zhang, Z. Ren, and N. Li. Gradient play in multi-agent markov stochastic games: stationary points, convergence, and sample complexity. *CoRR*, abs/2106.00198, 2021.

## Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? **[Yes]** See Section ??.
- Did you include the license to the code and datasets? **[No]** The code and the data are proprietary.
- Did you include the license to the code and datasets? **[N/A]**

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]**
  - (b) Did you describe the limitations of your work? **[Yes]** See the contribution paragraph in Section 1, assumptions in theorems, and Section 7: Discussions and Conclusions.
  - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** As this work is primarily theoretical, we do not foresee any potentially negative societal impacts of this work.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**

2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
  - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The URL is provided in Section 6
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A] Our experiments are deterministic
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [N/A]
  - (b) Did you mention the license of the assets? [N/A]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

**Other notations:** We use the abbreviation  $\pi_{\theta_i, s}$  to denote the probability distribution  $\pi_{\theta_i}(\cdot|s)$  (as well as the corresponding  $|\mathcal{A}_i|$  dimensional vector). We use  $\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\infty$  to denote the  $\ell_1, \ell_2, \ell_\infty$  norm respectively.  $\text{KL}(\cdot|\cdot)$  is used to denote the KL divergence of two probability distributions. We also define the value function, Q-function, advantage function, averaged Q-function and averaged advantage function with respect to potential function  $\phi$  as

$$V_\phi^\theta(s) := \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t \phi(s(t), a(t)) \mid \pi_\theta, s(0) = s \right],$$

$$Q_\phi^\theta(s, a) := \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t \phi(s(t), a(t)) \mid \pi_\theta, s(0) = s, a(0) = a \right],$$

$$A_\phi^\theta(s, a) := Q_\phi^\theta(s, a) - V_\phi^\theta(s), \quad \overline{Q}_{i, \phi}^\theta(s, a_i) := \sum_{a_{-i}} \pi_{\theta_{-i}}(a_{-i} | s) Q_\phi^\theta(s, a_i, a_{-i}),$$

$$\overline{A}_{i, \phi}^\theta(s, a_i) := \sum_{a_{-i}} \pi_{\theta_{-i}}(a_{-i} | s) A_\phi^\theta(s, a_i, a_{-i}).$$

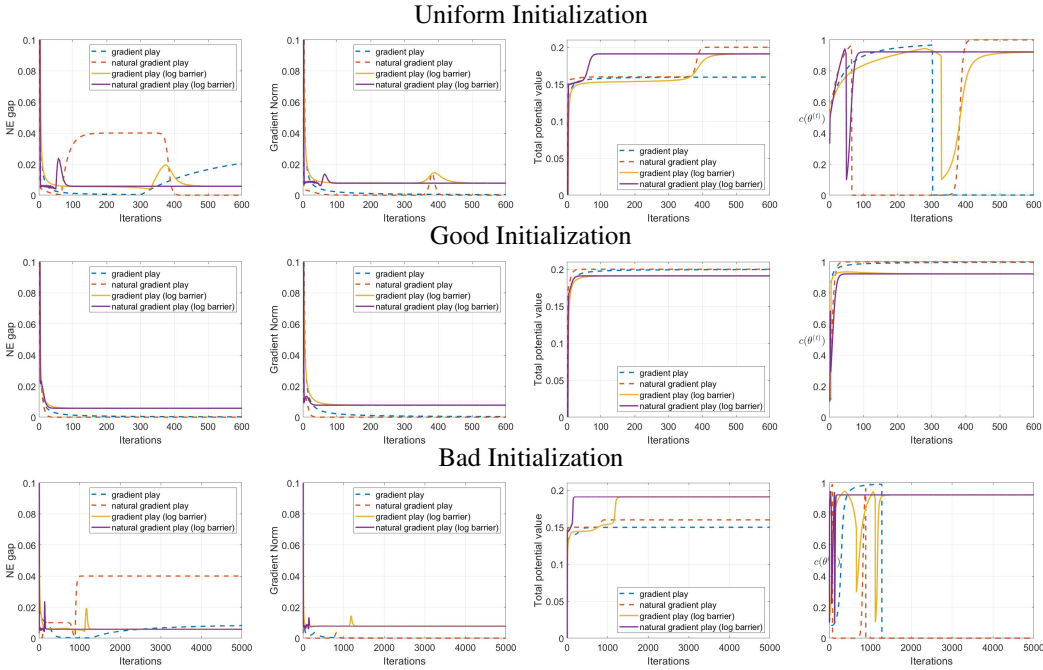


Figure 2: This set of figures shows how each algorithm performs for different initializations. Uniform initialization is the same as described in Figure 1; for good initialization, we choose the initial parameter as  $\pi_1 = [0.1, 0.1, 0.8], \pi_2 = [0.5, 0.5]$ ; for good initialization, we choose the initial parameter as  $\pi_1 = [0.8, 0.1, 0.1], \pi_2 = [0.5, 0.5]$ . Figures from top row to bottom plot out NE-gap( $\theta^t$ ),  $\|\nabla_\theta \Phi(\theta^t)\|_2$ ,  $\Phi(\theta^t)$  and  $c(\theta^t)$  respectively. Here we choose  $\eta = 5, \lambda = 0.003$ .




Figure 3: How  $\overline{Q}_i^{(t)}(a_i), \pi_i^{(t)}(a_i)$  change with time when running unregularized natural gradient play. The left two figures plots out  $\overline{Q}_1^{(t)}(a_1), \pi_1^{(t)}(a_1)$  and the right two figures plots out  $\overline{Q}_2^{(t)}(a_2), \pi_2^{(t)}(a_2)$ .

## A Numerical Simulations

This section provides more material for the numerical example shown in Section 6. Figure 2 displays numerical performance for different initialization policies. All four algorithms perform well given a good initialization, i.e., initial policy close to a stable NE. However for bad initialization that is close to a non-NE stationary point, log-barrier regularized algorithms can escape bad regions and converge to NE much faster than unregularized dynamics.

To examine why multi-agent learning suffers more from getting stuck at undesirable stationary points, we plot out the trajectory for  $\overline{Q}_i^{(t)}(a_i), \pi_i^{(t)}(a_i)$  for both agents in Figure 3. We will mainly focus our attention on the two plots on the left. Note that for the first few steps,  $\overline{Q}_1^{(t)}(a_1 = 2)$  is much larger than  $\overline{Q}_1^{(t)}(a_1 = 3)$ , thus the natural gradient play scheme (10) will drive  $\pi_1^{(t)}(a_1 = 2)$  close to 1 and  $\pi_1^{(t)}(a_1 = 3)$  close to 0 very quickly. However, at around iteration 70,  $\overline{Q}_1^{(t)}(a_1 = 3)$  becomes slightly larger than  $\overline{Q}_1^{(t)}(a_1 = 2)$ . Unfortunately, at this stage, most of the probability is assigned to the suboptimal action  $a_1 = 2$  and the optimal action receives  $\pi_1^{(t)}(a_1 = 3)$  close to zero. Thus it will take more steps to bring  $\pi_1^{(t)}(a_1 = 2)$  from 1 to 0 and  $\pi_1^{(t)}(a_1 = 3)$  from 0 to 1, which reflects as the trajectory being stuck at the non-NE stationary policy with  $\pi_1(a_1 = 3) = 1$  in numerical behavior. From this simulation, we may conclude that one important reason for natural gradient play to get stuck at undesirable stationary points is due to the fact that the value of averaged  $Q$ -functions  $\overline{Q}_i^{(t)}$ 's for different actions might switch order during the learning process. In contrast, for single agent bandit learning, the averaged  $Q$ -function as well as the  $Q$ -function itself is the same as the reward value of a certain action  $r(a)$ , and thus will not change order, which explains why it can achieve dimension free convergence in single agent learning.

Additionally, we would like to remark that the algorithms considered in this paper also generalizes to settings with more agents, and similar phenomenon will still be observed. See Figure 4 and 5 for numerical simulations on a 3-agent example and an 8-agent example.





Figure 4: A 3-agent Example Figure 5: An 8-agent Example

Running trajectories for all the four algorithms for one set of initializations takes approximately 2.04 seconds of CPU running time (Intel(R) Core(TM) i5-8250U CPU @ 1.60GHz 1.80 GHz).

## B Derivation of Gradient and Performance Difference Lemma

*Proof.* (of Equation 4) According to policy gradient theorem [34]:

$$\frac{\partial J_i(\theta)}{\partial \theta_{s,a_i}} = \frac{1}{1-\gamma} \sum_{s'} \sum_{a'} d_\theta(s') \pi_\theta(a'|s') \frac{\partial \log \pi_\theta(a'|s')}{\partial \theta_{s,a_i}} Q_i^\theta(s, a)$$

Since for softmax parameterization:

$$\frac{\partial \log \pi_\theta(a'|s')}{\partial \theta_{s,a_i}} = \frac{\partial \log \pi_{\theta_i}(a'_i|s')}{\partial \theta_{s,a_i}} = \mathbf{1}\{a'_i = a_i, s' = s\} - \mathbf{1}\{s' = s\} \pi_{\theta_i}(a_i|s)$$

Thus we have that:

$$\begin{aligned} \frac{\partial J_i(\theta)}{\partial \theta_{s,a_i}} &= \frac{1}{1-\gamma} \sum_{s'} \sum_{a'} d_\theta(s') \pi_\theta(a'|s') (\mathbf{1}\{a'_i = a_i, s' = s\} - \mathbf{1}\{s' = s\} \pi_{\theta_i}(a_i|s)) Q_i^\theta(s, a') \\ &= \frac{1}{1-\gamma} d_\theta(s) \pi_{\theta_i}(a_i|s) \sum_{a'_{-i}} \pi_{\theta_{-i}}(a'_{-i}|s) Q_i^\theta(s, a_i, a'_{-i}) - \frac{1}{1-\gamma} d_\theta(s) \pi_{\theta_i}(a_i|s) \sum_{a'} \pi_\theta(a'|s') Q_i^\theta(s, a') \\ &= \frac{1}{1-\gamma} d_\theta(s) \pi_{\theta_i}(a_i|s) \overline{Q}_i^\theta(s, a_i, a'_{-i}) - \frac{1}{1-\gamma} d_\theta(s) \pi_{\theta_i}(a_i|s) V_i^\theta(s) \end{aligned}$$

$$= \frac{1}{1-\gamma} d_\theta(s) \pi_{\theta_i}(a_i|s) \overline{A_i^\theta}(s, a_i)$$

□

We also introduce a useful lemma used throughout the proof which is derived from the performance difference lemma in MDP [15].

**Lemma 8.** Let  $\theta' = (\theta'_i, \theta_{-i})$ ,

$$J_i(\theta'_i, \theta_{-i}) - J_i(\theta_i, \theta_{-i}) = \frac{1}{1-\gamma} \sum_{s, a_i} d_{\theta'}(s) \pi_{\theta'_i}(a_i|s) \overline{A_i^\theta}(s, a_i)$$

*Proof.* From performance difference lemma [15]

$$\begin{aligned} J_i(\theta'_i, \theta_{-i}) - J_i(\theta_i, \theta_{-i}) &= \frac{1}{1-\gamma} \sum_{s, a} d_{\theta'}(s) \pi_{\theta'_i}(a|s) A_i^\theta(s, a) \\ &= \frac{1}{1-\gamma} \sum_{s, a_i} d_{\theta'}(s) \pi_{\theta'_i}(a_i|s) \sum_{a_{-i}} \pi_{\theta_{-i}}(a_{-i}|s) A_i^\theta(s, a_i, a_{-i}) \\ &= \frac{1}{1-\gamma} \sum_{s, a_i} d_{\theta'}(s) \pi_{\theta'_i}(a_i|s) \overline{A_i^\theta}(s, a_i). \end{aligned}$$

□

## C Derivation of Natural Gradient Play

**Lemma 9.**

$$\mathbb{E}_{a \sim \pi_\theta(\cdot|s)} [\nabla_{\theta_{i,s}} \log \pi_{\theta_i}(a_i|s) \nabla_{\theta_{i,s}} \log \pi_{\theta_i}(a_i|s)^\top] = \text{diag}\{\pi_{\theta_{i,s}}\} - \pi_{\theta_{i,s}} \pi_{\theta_{i,s}}^\top := F_{i,s}(\theta_{i,s}),$$

where  $\text{diag}\{\cdot\}$  denotes the diagonal matrix generated by the corresponding vector, and  $\pi_{\theta_{i,s}} \in \mathbb{R}^{|\mathcal{A}_i|}$  is the vector that denotes  $\pi_{\theta_i}(\cdot|s)$ . Further,  $F_{i,s}(\theta_{i,s})$  is a semi-positive definite matrix, where the eigenvalue 0 has the eigenspace of dimension 1 that is the span of the all one-vector  $\mathbf{1}$ .

*Proof.* Calculating the gradient using chain rule we have

$$\frac{\partial \log \pi_{\theta_i}(a_i|s)}{\partial \theta'_{a'_i, s}} = \mathbf{1}\{a'_i = a_i\} - \pi_{\theta_i}(a'_i|s).$$

Let  $\mathbf{1}_{a_i} \in \mathbb{R}^{|\mathcal{A}_i|}$  denote the vector where the entry corresponds to  $a_i$  is 1 and other entries are zero. Then

$$\begin{aligned} \nabla_{\theta_{i,s}} \log \pi_{\theta_i}(a|s) &= \mathbf{1}_{a_i} - \pi_{\theta_{i,s}} \\ \implies \nabla_{\theta_{i,s}} \log \pi_{\theta_i}(a_i|s) \nabla_{\theta_{i,s}} \log \pi_{\theta_i}(a_i|s)^\top &= \text{diag}\{\mathbf{1}_{a_i}\} - \pi_{\theta_{i,s}} \mathbf{1}_{a_i}^\top - \mathbf{1}_{a_i} \pi_{\theta_{i,s}}^\top + \pi_{\theta_{i,s}} \pi_{\theta_{i,s}}^\top \end{aligned}$$

Taking the expectation  $\mathbb{E}_{a \sim \pi_{\theta_i}(\cdot|s)}$  we have

$$\begin{aligned} \mathbb{E}_{a \sim \pi_{\theta_i}(\cdot|s)} [\nabla_{\theta_{i,s}} \log \pi_{\theta_i}(a|s) \nabla_{\theta_{i,s}} \log \pi_{\theta_i}(a|s)^\top] &= \text{diag}\{\pi_{\theta_{i,s}}\} - \pi_{\theta_{i,s}} \pi_{\theta_{i,s}}^\top - \pi_{\theta_{i,s}} \pi_{\theta_{i,s}}^\top + \pi_{\theta_{i,s}} \pi_{\theta_{i,s}}^\top \\ &= \text{diag}\{\pi_{\theta_{i,s}}\} - \pi_{\theta_{i,s}} \pi_{\theta_{i,s}}^\top \end{aligned}$$

Further, for softmax parameterization,  $\pi_{\theta_i}(a_i|s) > 0, \forall a_i$ . Thus  $F_{i,s}(\theta_{i,s})$  is a (non-strict) diagonally dominant matrix with diagonal entries all being positive and off-diagonal entries all being negative, in which case the all-one vector  $\mathbf{1}$  is the only eigenvector for eigenvalue 0. □

**Corollary 10.**

$$F_i(\theta) = \text{blkdiag}\{d_\theta(s) F_{i,s}(\theta_{i,s})\}_{s \in \mathcal{S}},$$

where  $\text{blkdiag}\{\cdot\}$  denotes the block-diagonal matrix generated by corresponding sub-matrices.



*Proof.* This is a direct corollary of Lemma 9, since

$$\frac{\partial \log \pi_{\theta_i}(a_i|s)}{\partial \theta_{a'_i, s'}} = 0, \text{ for } s' \neq s,$$

we have that

$$F_i(\theta) = \mathbb{E}_{s \sim d_\theta(\cdot)} \mathbb{E}_{a_i \sim \pi_{\theta_i}(\cdot|s)} [\nabla_{\theta_i} \log \pi_{\theta_i}(a_i|s) \nabla_{\theta_i} \log \pi_\theta(a_i|s)^\top] = \text{blkdiag}\{d_\theta(s) F_{i,s}(\theta_{i,s})\}$$

□

**Lemma 11.** For vector  $g : \mathcal{S} \times \mathcal{A}_i \rightarrow \mathbb{R}$ , with  $\sum_{a_i} g(s, a_i) = 0$ ,  $\forall s \in \mathcal{S}$ , we have that

$$[F_i(\theta)^\dagger g]_{(s, a_i)} = \frac{1}{d_\theta(s) \pi_{\theta_i}(a_i|s)} g(s, a_i) + c(s),$$

where  $c(s)$  is a function that depend on state  $s$  but not on  $a_i$ .

*Proof.* Since  $F_i(\theta)$  is a block diagonal matrix,

$$[F_i(\theta)^\dagger g]_{(s, \cdot)} = \frac{1}{d_\theta(s)} F_{i,s}(\theta_{i,s})^\dagger g(s, \cdot).$$

From Lemma 9, since  $F_{i,s}$  only has a one-dimensional eigenspace for eigenvalue 0, and the eigenspace is the span of the all-one vector  $\mathbf{1}$ , we have that

$$F_{i,s}(\theta_{i,s})^\dagger F_{i,s}(\theta_{i,s}) = I - \frac{1}{|\mathcal{A}_i|} \mathbf{1}\mathbf{1}^\top.$$

Let  $f(s, a_i) := \frac{1}{d_\theta(s) \pi_{\theta_i}(a_i|s)} g(s, a_i)$

$$\begin{aligned} d_\theta(s) [F_{i,s}(\theta_{i,s}) f(s, \cdot)]_{a_i} &= d_\theta(s) \left( \pi_{\theta_i}(a_i|s) f(s, a_i) - \pi_{\theta_i}(a_i|s) \sum_{a'_i} \pi_{\theta_i}(a'_i|s) f(s, a'_i) \right) \\ &= g(s, a_i) - \pi_{\theta_i}(a_i|s) \sum_{a'_i} g(s, a'_i) = g(s, a_i), \end{aligned}$$

i.e.,

$$\begin{aligned} d_\theta(s) F_{i,s}(\theta_{i,s}) f(s, \cdot) &= g(s, \cdot) \\ \implies \frac{1}{d_\theta(s)} F_{i,s}(\theta_{i,s})^\dagger g(s, \cdot) &= F_{i,s}(\theta_{i,s})^\dagger F_{i,s}(\theta_{i,s}) f(s, \cdot) \\ &= \left( I - \frac{1}{|\mathcal{A}_i|} \mathbf{1}\mathbf{1}^\top \right) f(s, \cdot) \\ &= f(s, \cdot) - c(s) \mathbf{1}, \end{aligned}$$

i.e.,

$$[F_i(\theta)^\dagger g]_{(s, a_i)} = f(s, a_i) - c(s),$$

which completes the proof. □

**Lemma 12.** Scheme (9) and (10) are equivalent. Similarly, (16) and (17) are equivalent.

*Proof.* It is not hard to check that  $\nabla_{\theta_i} J_i(\theta)$ ,  $\nabla_{\theta_i} \tilde{J}_i(\theta)$  satisfies

$$\sum_{a_i} [\nabla_{\theta_i} J_i(\theta)]_{(s, a_i)} = 0, \quad \sum_{a_i} [\nabla_{\theta_i} \tilde{J}_i(\theta)]_{(s, a_i)} = 0,$$

thus we can apply Lemma 11 and conclude

$$\begin{aligned} F_i(\theta^{(t)})^\dagger \nabla_{\theta_i} J_i(\theta_i^{(t)}) &= \frac{\overline{A_i^{(t)}}(s, a_i)}{1 - \gamma} + c(s) \\ F_i(\theta^{(t)})^\dagger \nabla_{\theta_i} \tilde{J}_i(\theta_i^{(t)}) &= \frac{\overline{A_i^{(t)}}(s, a_i)}{1 - \gamma} + \frac{\lambda}{d^{(t)}(s) \pi_i^{(t)}(a_i | s)} - \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} + c(s), \end{aligned}$$

which completes the proof.  $\square$

## D Proof of Lemma 1 and Lemma 2

**Lemma 13.**

$$\text{NE-gap}_i(\theta) \leq \frac{1}{1 - \gamma} \max_{s, a_i} \overline{A_i^\theta}(s, a_i), \quad \text{NE-gap}(\theta) \leq \frac{1}{1 - \gamma} \max_i \max_{s, a_i} \overline{A_i^\theta}(s, a_i).$$

*Proof.* From performance difference lemma

$$\begin{aligned} J_i(\theta'_i, \theta_{-i}) - J_i(\theta_i, \theta_{-i}) &= \frac{1}{1 - \gamma} \sum_{s, a_i} d_{\theta'}(s) \pi_{\theta'_i}(a_i | s) \overline{A_i^\theta}(s, a_i) \quad (\text{Lemma 8}) \\ &\leq \frac{1}{1 - \gamma} \sum_s d_{\theta'}(s) \max_{a_i} \overline{A_i^\theta}(s, a_i) \\ &\leq \frac{1}{1 - \gamma} \sum_s d_{\theta'}(s) \max_{a_i} \overline{A_i^\theta}(s, a_i) \\ &\leq \frac{1}{1 - \gamma} \max_{s, a_i} \overline{A_i^\theta}(s, a_i). \end{aligned}$$

Thus we have that

$$\text{NE-gap}_i(\theta) \leq \frac{1}{1 - \gamma} \max_{s, a_i} \overline{A_i^\theta}(s, a_i), \quad \text{NE-gap}(\theta) \leq \frac{1}{1 - \gamma} \max_i \max_{s, a_i} \overline{A_i^\theta}(s, a_i).$$

$\square$

*Proof of Lemma 1.* From Lemma 13 we have that

$$\text{NE-gap}_i(\theta) \leq \frac{1}{1 - \gamma} \max_{s, a_i} \overline{A_i^\theta}(s, a_i).$$

Since

$$\begin{aligned} \max_{a_i} \overline{A_i^\theta}(s, a_i) &\leq \frac{1}{\sum_{a_i^* \in \arg\max_{a_i} \overline{Q_i^\theta}(s, a_i)} \pi_{\theta_i}(a_i^* | s)} \sum_{a_i} |\pi_{\theta_i}(a_i | s) \overline{A_i^\theta}(s, a_i)| \\ &\leq \frac{\sqrt{|\mathcal{A}_i|}}{\sum_{a_i^* \in \arg\max_{a_i} \overline{Q_i^\theta}(s, a_i)} \pi_{\theta_i}(a_i^* | s)} \sqrt{\sum_{a_i} \left( \pi_{\theta_i}(a_i | s) \overline{A_i^\theta}(s, a_i) \right)^2} \\ &= \frac{\sqrt{|\mathcal{A}_i|}}{\sum_{a_i^* \in \arg\max_{a_i} \overline{Q_i^\theta}(s, a_i)} \pi_{\theta_i}(a_i^* | s)} \frac{1 - \gamma}{d_\theta(s)} \sqrt{\sum_{a_i} \left( \frac{1}{1 - \gamma} d_\theta(s) \pi_{\theta_i}(a_i | s) \overline{A_i^\theta}(s, a_i) \right)^2} \\ &\leq \frac{(1 - \gamma) M(\theta) \sqrt{|\mathcal{A}_i|}}{c(\theta)} \|\nabla_{\theta_i} J_i(\theta)\|_2. \end{aligned}$$

Thus

$$\text{NE-gap}_i(\theta) \leq \frac{1}{1 - \gamma} \max_{s, a_i} \overline{A_i^\theta}(s, a_i)$$

$$\leq \frac{\sqrt{|\mathcal{A}_i|}M(\theta)}{c(\theta)} \|\nabla_{\theta_i} J_i(\theta)\|_2$$

□

*Remark 4.* A similar bound to Lemma 1 can be obtained by leveraging equation (259) in [24]

$$\text{NE-gap}_i(\theta) \leq \frac{\sqrt{|\mathcal{S}||\mathcal{A}_i|}D_\infty}{c(\theta)} \|\nabla_{\theta_i} J_i(\theta)\|_2, \quad \text{where } D_\infty = \sup_{\theta, \theta'} \max_s \frac{d_{\theta'}(s)}{d_\theta(s)}. \quad (18)$$

Notice that there's an additional  $\sqrt{|\mathcal{S}|}$  dependency on the right hand side compared with Lemma 1, while replacing the term  $M(\theta)$  by  $D_\infty$ . We remark that there's no fundamental difference between these two bounds. There's no significant difference in the proof techniques and it is hard to tell which one is better. One can also easily re-derive the set of analysis in the paper using (18), with bounds that depends on  $D_\infty$  instead of  $M$  and slightly differs in the dependency on  $\mathcal{S}$  from our current result.

*Proof of Lemma 2.* Firstly, it is straightforward to see that if  $c(\theta^*) \neq 0$ ,  $\theta^*$  is a NE by applying Lemma 1. So we only need to focus on proving that if  $\theta^*$  is a NE, then  $c(\theta^*) = 1$ .

From performance difference lemma, let  $\theta' := (\theta'_i, \theta_{-i}^*)$

$$J_i(\theta'_i, \theta_{-i}^*) - J_i(\theta_i^*, \theta_{-i}^*) = \frac{1}{1-\gamma} \sum_{s, a_i} d_{\theta'}(s) \pi_{\theta'_i}(a_i|s) \overline{A_i^{\theta'}(s, a_i)}$$

Select  $a_i^*(s) \in \operatorname{argmax}_{a_i} \overline{A_i^{\theta^*}(s, a_i)}$  and set:

$$\pi_{\theta'_i}(a_i|s) = \mathbf{1}\{a_i = a_i^*(s)\},$$

then

$$\begin{aligned} J_i(\theta'_i, \theta_{-i}^*) - J_i(\theta_i^*, \theta_{-i}^*) &= \frac{1}{1-\gamma} \sum_{s, a_i} d_{\theta'}(s) \pi_{\theta'_i}(a_i|s) \overline{A_i^{\theta'}(s, a_i)} \\ &= \frac{1}{1-\gamma} \sum_s d_{\theta'}(s) \max_{a_i} \overline{A_i^{\theta'}(s, a_i)} \geq 0. \end{aligned}$$

Since  $\theta^*$  is a NE,

$$\implies \max_{a_i} \overline{A_i^{\theta^*}(s, a_i)} = 0, \quad \forall s, \forall i.$$

Let  $\Delta := \min_s \min_{a_i \notin \operatorname{argmax}_{a_i} \overline{A_i^{\theta^*}(s, a_i)}} |\overline{A_i^{\theta^*}(s, a_i)}|$ . Since  $\sum_{a_i} \pi_{\theta_i^*}(a_i|s) \overline{A_i^{\theta^*}(s, a_i)} = 0$

$$\begin{aligned} \implies 0 &= \sum_{a_i \in \operatorname{argmax}_{a_i} \overline{A_i^{\theta^*}(s, a_i)}} \pi_{\theta_i^*}(a_i|s) \max_{a_i} \overline{A_i^{\theta^*}(s, a_i)} + \sum_{a_i \notin \operatorname{argmax}_{a_i} \overline{A_i^{\theta^*}(s, a_i)}} \pi_{\theta_i^*}(a_i|s) \overline{A_i^{\theta^*}(s, a_i)} \\ &\leq -\Delta \sum_{a_i \notin \operatorname{argmax}_{a_i} \overline{A_i^{\theta^*}(s, a_i)}} \pi_{\theta_i^*}(a_i|s) \\ \implies \sum_{a_i \notin \operatorname{argmax}_{a_i} \overline{A_i^{\theta^*}(s, a_i)}} \pi_{\theta_i^*}(a_i|s) &= 0 \\ \implies \sum_{a_i \in \operatorname{argmax}_{a_i} \overline{A_i^{\theta^*}(s, a_i)}} \pi_{\theta_i^*}(a_i|s) &= 1 \\ \implies \sum_{a_i \in \operatorname{argmax}_{a_i} \overline{Q_i^{\theta^*}(s, a_i)}} \pi_{\theta_i^*}(a_i|s) &= 1 \\ \implies c(\theta^*) &= 1 \end{aligned}$$

□

## E Proof of Theorem 4

### E.0.1 Asymptotic convergence for gradient play

**Lemma 14.** For  $\eta \leq \frac{(1-\gamma)^3}{6n}$ , running scheme (8) will guarantee that  $\lim_{t \rightarrow +\infty} \nabla \Phi(\theta^{(t)}) = 0$ .

*Proof.* Since  $\Phi(\theta)$  is  $\beta$ -smooth w.r.t.  $\theta$ , where  $\beta = \frac{6n}{(1-\gamma)^3}$

$$\begin{aligned} \Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)}) &\geq \left\langle \nabla \Phi(\theta^{(t)}), \theta^{(t+1)} - \theta^{(t)} \right\rangle - \frac{\beta}{2} \|\theta^{(t+1)} - \theta^{(t)}\|_2^2 \\ &\geq \frac{\eta}{2} \|\nabla \Phi(\theta^{(t)})\|_2^2 \geq 0 \end{aligned}$$

which proves the monotonicity of  $\Phi(\theta^{(t)})$ . Since  $\phi$  is a bounded function, this gives:

$$\lim_{t \rightarrow +\infty} \|\nabla \Phi(\theta^{(t)})\|_2 = 0. \quad \square$$

From Lemma 14 and that the stationary points are isolated, we know that the limit for  $\theta^{(t)}$  exists, i.e., it is valid to define

$$\theta^{(\infty)} := \lim_{t \rightarrow +\infty} \theta^{(t)}.$$

We abbreviate the related functions with respect to  $\theta^{(\infty)}$  as follows:

$$\begin{aligned} Q_i^{(\infty)}(s, a) &:= Q_i^{\theta^{(\infty)}}(s, a), & V_i^{(\infty)}(s) &:= V_i^{\theta^{(\infty)}}(s), & A_i^{(\infty)}(s, a) &:= Q_i^{(\infty)}(s, a) - V_i^{(\infty)}(s) \\ \overline{Q_i^{(\infty)}}(s, a_i) &:= \sum_{a_{-i}} \pi_{-i}^{(\infty)}(a_{-i}|s) Q_i^{(\infty)}(s, a_i, a_{-i}), & \overline{A_i^{(\infty)}}(s, a_i) &:= \sum_{a_{-i}} \pi_{-i}^{(\infty)}(a_{-i}|s) A_i^{(\infty)}(s, a_i, a_{-i}) \end{aligned}$$

Since  $\theta^{(\infty)}$  is the limit of  $\theta^{(t)}$ , we have that:

$$\lim_{t \rightarrow +\infty} \overline{Q_i^{(t)}}(s, a_i) = \overline{Q_i^{(\infty)}}(s, a_i), \quad \lim_{t \rightarrow +\infty} \overline{A_i^{(t)}}(s, a_i) = \overline{A_i^{(\infty)}}(s, a_i) \quad (19)$$

Define:

$$\begin{aligned} I_0^{i,s} &:= \{a_i | \overline{Q_i^{(\infty)}}(s, a_i) = V_i^{(\infty)}(s)\} = \{a_i | \overline{A_i^{(\infty)}}(s, a_i) = 0\} \\ I_+^{i,s} &:= \{a_i | \overline{Q_i^{(\infty)}}(s, a_i) > V_i^{(\infty)}(s)\} = \{a_i | \overline{A_i^{(\infty)}}(s, a_i) > 0\} \\ I_-^{i,s} &:= \{a_i | \overline{Q_i^{(\infty)}}(s, a_i) < V_i^{(\infty)}(s)\} = \{a_i | \overline{A_i^{(\infty)}}(s, a_i) < 0\} \end{aligned}$$

Let

$$\Delta := \min_i \min_{\{s, a_i | \overline{A_i^{(\infty)}}(s, a_i) \neq 0\}} |\overline{A_i^{(\infty)}}(s, a_i)| \quad (20)$$

From Lemma 13, it is sufficient to show that  $I_+^{i,s} = \emptyset$ ,  $\forall i, s$ .

From the Lemma 14 and the above definitions we have the following corollaries:

**Corollary 15.** There exists  $T_1$ , such that  $\forall t > T_1$ ,  $\forall s \in \mathcal{S}$ ,  $\forall i \in \{1, 2, \dots, n\}$ ,

$$\begin{aligned} A_i^{(t)}(s, a_i) &< -\frac{\Delta}{4}, \quad \forall a_i \in I_-^{i,s} \\ A_i^{(t)}(s, a_i) &> \frac{\Delta}{4}, \quad \forall a_i \in I_+^{i,s} \\ |A_i^{(t)}(s, a_i)| &< \frac{\Delta}{4}, \quad \forall a_i \in I_0^{i,s} \end{aligned}$$

*Proof.* This is a direct corollary from (19) and (20). □

**Corollary 16.**

$$\begin{aligned}\lim_{t \rightarrow +\infty} \sum_{a_i \in I_0^{i,s}} \pi_i^{(t)}(a_i|s) &= 1 \\ \lim_{t \rightarrow +\infty} \sum_{a_i \in I_+^{i,s} \cup I_-^{i,s}} \pi_i^{(t)}(a_i|s) &= 0\end{aligned}$$

*Proof.* This is a direct corollary from Lemma 14,

$$\begin{aligned}\lim_{t \rightarrow +\infty} \nabla \Phi(\theta^{(t)}) &= 0 \\ \implies \lim_{t \rightarrow +\infty} \frac{\partial \Phi(\theta^{(t)})}{\partial \theta_{s,a_i}} &= \lim_{t \rightarrow +\infty} \frac{1}{1-\gamma} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) = 0 \\ \implies \lim_{t \rightarrow +\infty} \pi_i^{(t)}(a_i|s) \lim_{t \rightarrow +\infty} \overline{A_i^{(t)}}(s, a_i) &= 0 \\ \implies \lim_{t \rightarrow +\infty} \pi_i^{(t)}(a_i|s) &= 0, \quad \forall a_i \notin I_0^{i,s} \\ \implies \lim_{t \rightarrow +\infty} \sum_{a_i \in I_+^{i,s} \cup I_-^{i,s}} \pi_i^{(t)}(a_i|s) &= 0 \\ \implies \lim_{t \rightarrow +\infty} \sum_{a_i \in I_0^{i,s}} \pi_i^{(t)}(a_i|s) &= 1 - \lim_{t \rightarrow +\infty} \sum_{a_i \in I_+^{i,s} \cup I_-^{i,s}} \pi_i^{(t)}(a_i|s) = 1 \quad \square\end{aligned}$$

**Lemma 17.**  $\forall a_i \in I_+^{i,s}, \theta_{s,a_i}^{(t)}$  is bounded from below.  $\forall a_i \in I_-^{i,s}, \lim_{t \rightarrow +\infty} \theta_{s,a_i}^{(t)} = -\infty$ .

*Proof.* The first statement,  $\forall a_i \in I_+^{i,s}, \theta_{s,a_i}^{(t)}$  is bounded from below, is trivial from Corollary 15. We only need to prove the second statement. The key observation is that:

$$\sum_{a_i} \frac{\partial \Phi(\theta^{(t)})}{\partial \theta_{s,a_i}} = \frac{1}{1-\gamma} d^{(t)}(s) \sum_{a_i} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) = 0$$

Thus

$$\sum_{a_i} \theta_{s,a_i}^{(t)} = \sum_{a_i} \theta_{s,a_i}^{(0)}.$$

From Corollary 16, we have that

$$\begin{aligned}\lim_{t \rightarrow +\infty} \sum_{a_i \in I_+^{i,s} \cup I_-^{i,s}} \pi_i^{(t)}(a_i|s) &= 0 \\ \implies \exists a_i \in I_0^{i,s}, \text{ s.t. } \limsup_{t \rightarrow +\infty} \theta_{s,a_i}^{(t)} &= +\infty\end{aligned}$$

And since all  $\theta_{s,a_i}^{(t)}$  sum up to a constant and that  $\forall a_i \in I_+^{i,s}, \theta_{s,a_i}^{(t)}$  is bounded from below, we have that:

$$\exists \overline{a_i} \in I_0^{i,s} \cup I_-^{i,s}, \text{ s.t. } \liminf_{t \rightarrow +\infty} \theta_{s,\overline{a_i}}^{(t)} = -\infty. \quad (21)$$

From Corollary 15, for  $a_i \in I_-^{i,s}, \theta_{s,a_i}^{(t)}$  is monotonically decreasing for  $t > T_1$ , thus

$$\lim_{t \rightarrow +\infty} \theta_{s,a_i}^{(t)} := \theta_{s,a_i}^{(\infty)},$$

where  $\theta_{s,a_i}^{(\infty)}$  is either a constant or  $-\infty$ . We'll prove by contradiction. Suppose  $\theta_{s,a_i}^{(\infty)}$  is a constant, then for any  $\delta > 0$  there exists  $T_1' \geq T_1$  such that  $\forall t \geq T_1', |\theta_{s,a_i}^{(t)} - \theta_{s,a_i}^{(\infty)}| \leq \delta$ .

Let  $\bar{a}_i \in \mathcal{A}_i$  be defined as in (21), define:

$$\tau(t) := \begin{cases} t + 1, & \text{if } \theta_{s,\bar{a}_i}^{(t)} > \theta_{s,a_i}^{(\infty)} - \delta \\ \min_{t'} \{T_1 \leq t' \leq t \mid \theta_{s,\bar{a}_i}^{(\tau)} \leq \theta_{s,a_i}^{(\infty)} - \delta, \forall t' \leq \tau \leq t\}, & \text{otherwise} \end{cases}$$

We will focus on the set where  $\{t \mid \tau(t) \leq t\}$ . Since  $\liminf_{t \rightarrow +\infty} \theta_{s,\bar{a}_i}^{(t)} = -\infty$ , there are infinitely many elements in this set.

For all  $\tau(t) \leq \tau \leq t$ , we have that:

$$\begin{aligned} \left| \frac{\partial \Phi(\theta^{(\tau)})}{\partial \theta_{s,a_i}} \right| &= \left| \frac{\pi_i^{(\tau)}(a_i|s) \overline{A_i^{(\tau)}}(s, a_i)}{\pi_i^{(\tau)}(\bar{a}_i|s) \overline{A_i^{(\tau)}}(s, \bar{a}_i)} \right| = \exp(\theta_{s,a_i}^{(\tau)} - \theta_{s,\bar{a}_i}^{(\tau)}) \left| \frac{\overline{A_i^{(\tau)}}(s, a_i)}{\overline{A_i^{(\tau)}}(s, \bar{a}_i)} \right| \\ &\geq \frac{\overline{A_i^{(\tau)}}(s, a_i)}{\overline{A_i^{(\tau)}}(s, \bar{a}_i)} \geq \frac{\Delta(1-\gamma)}{4} \end{aligned}$$

Thus

$$\begin{aligned} \frac{\partial \Phi(\theta^{(\tau)})}{\partial \theta_{s,a_i}} &\leq \frac{\Delta(1-\gamma)}{4} \frac{\partial \Phi(\theta^{(\tau)})}{\partial \theta_{s,\bar{a}_i}}, \quad \tau(t) \leq \tau \leq t \\ \implies \frac{1}{\eta} (\theta_{s,a_i}^{(t+1)} - \theta_{s,a_i}^{(\tau(t))}) &= \sum_{\tau(t)}^t \frac{\partial \Phi(\theta^{(\tau)})}{\partial \theta_{s,a_i}} \leq \frac{\Delta(1-\gamma)}{4} \sum_{\tau(t)}^t \frac{\partial \Phi(\theta^{(\tau)})}{\partial \theta_{s,\bar{a}_i}} = \frac{\Delta(1-\gamma)}{4\eta} (\theta_{s,\bar{a}_i}^{(t+1)} - \theta_{s,\bar{a}_i}^{(\tau(t))}) \end{aligned} \quad (22)$$

Since:

$$\theta_{s,\bar{a}_i}^{(\tau(t))} \geq \theta_{s,\bar{a}_i}^{(\tau(t)-1)} - \eta \frac{1}{(1-\gamma)^2} \geq \theta_{s,a_i}^{(\infty)} - \delta - \eta \frac{1}{(1-\gamma)^2}$$

is bounded from below, and that  $\theta_{s,a_i}^{\tau(t)}$  is also bounded from above by  $\theta_{s,a_i}^{(T_1)}$ , thus taking  $\liminf_{t \rightarrow +\infty}$  on both sides of eq (22) will give

$$\begin{aligned} \liminf_{t \rightarrow +\infty} \theta_{s,a_i}^{(t+1)} - \theta_{s,a_i}^{(\tau(t))} &\leq \frac{\Delta(1-\gamma)}{4} \left( \liminf_{t \rightarrow +\infty} \theta_{s,\bar{a}_i}^{(t+1)} - \theta_{s,\bar{a}_i}^{(\infty)} + \delta + \eta \frac{1}{(1-\gamma)^2} \right) = -\infty \\ \implies \liminf_{t \rightarrow +\infty} \theta_{s,a_i}^{(t)} &= -\infty \end{aligned}$$

which contradicts the assumption that  $\theta_{s,a_i}^{(\infty)}$  is a constant, and thus we can conclude that

$$\lim_{t \rightarrow +\infty} \theta_{s,a_i}^{(t)} = -\infty, \quad \forall a_i \in I_-^{i,s}. \quad \square$$

**Lemma 18.**  $\forall a_i^+ \in I_+^{i,s}$ , for any  $a \in I_0^{i,s}$ , if there exists  $t \geq T_1$  such that  $\pi_i^{(t)}(a_i|s) \leq \pi_i^{(t)}(a_i^+|s)$ , then for all  $\tau \geq t$ ,  $\pi_i^{(\tau)}(a_i|s) \leq \pi_i^{(\tau)}(a_i^+|s)$

*Proof.* We will prove by induction. Suppose for a certain  $\tau \geq t$ , it holds that  $\pi_i^{(\tau)}(a_i|s) \leq \pi_i^{(\tau)}(a_i^+|s)$ , then:

$$\begin{aligned} \frac{\partial \Phi(\theta^{(\tau)})}{\partial \theta_{s,a_i^+}} &= \frac{1}{1-\gamma} d^{(\tau)}(s) \pi_i^{(\tau)}(a_i^+|s) \overline{A_i^{(\tau)}}(s, a_i^+) \\ &\geq \frac{1}{1-\gamma} d^{(\tau)}(s) \pi_i^{(\tau)}(a_i|s) \overline{A_i^{(\tau)}}(s, a_i^+) \\ &\geq \frac{1}{1-\gamma} d^{(\tau)}(s) \pi_i^{(\tau)}(a_i|s) \overline{A_i^{(\tau)}}(s, a_i) \\ &= \frac{\partial \Phi(\theta^{(\tau)})}{\partial \theta_{s,a_i}} \end{aligned}$$

Since  $\pi_i^{(\tau)}(a_i|s) \leq \pi_i^{(\tau)}(a_i^+|s) \implies \theta_{s,a_i}^{(\tau)} \leq \theta_{s,a_i^+}^{(\tau)}$ , we have:

$$\theta_{s,a_i^+}^{(\tau+1)} = \theta_{s,a_i^+}^{(\tau)} + \eta \frac{\partial \Phi(\theta^{(\tau)})}{\partial \theta_{s,a_i^+}} \geq \theta_{s,a_i}^{(\tau)} + \eta \frac{\partial \Phi(\theta^{(\tau)})}{\partial \theta_{s,a_i}} = \theta_{s,a_i}^{(\tau+1)}$$

Thus  $\pi_i^{(\tau+1)}(a_i|s) \leq \pi_i^{(\tau+1)}(a_i^+|s)$  also holds, which completes the proof.  $\square$

**Lemma 19.**  $I_+^{i,s} = \emptyset$ .

*Proof.* We will prove by contradiction. If  $I_+^{i,s} \neq \emptyset$ , select an arbitrary  $a_i^+ \in I_+^{i,s}$  and define

$$B_0^{i,s}(a_i^+) := \{a_i \in I_0^{i,s} \mid \pi_i^{(t)}(a_i|s) \leq \pi_i^{(t)}(a_i^+|s), \forall t \geq T_1\}.$$

From Lemma 17, we have that for any  $a_i \in I_-^{i,s}$   $\lim_{t \rightarrow +\infty} \frac{\pi_i^{(t)}(a_i|t)}{\pi_i^{(t)}(a_i^+|t)} = 0$ , thus there exists  $T_2 > T_1$  such that for any  $t \geq T_2$ ,

$$\frac{\pi_i^{(t)}(a_i|t)}{\pi_i^{(t)}(a_i^+|t)} \leq \frac{(1-\gamma)\Delta}{16|\mathcal{A}_i|}, \quad \forall a_i \in I_-^{i,s}.$$

Additionally, since for any  $a_i \in I_0^{i,s}$ ,  $\lim_{t \rightarrow +\infty} \overline{A_i^{(t)}}(s, a_i) = 0$ , there exists  $T_3 > T_1$  such that for any  $t \geq T_3$ ,

$$\overline{A_i^{(t)}}(s, a_i) \geq \frac{-\Delta}{16|\mathcal{A}_i|}, \quad \forall a_i \in I_0^{i,s}.$$

Thus, for  $t \geq \max\{T_2, T_3\}$ , from the fact that  $\sum_{a_i} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) = 0$ , we have:

$$\begin{aligned} 0 &= \sum_{a_i \in I_0^{i,s}} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) + \sum_{a_i \in I_+^{i,s}} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) + \sum_{a_i \in I_-^{i,s}} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) \\ &\geq \sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) + \sum_{a_i \in B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) \\ &\quad + \pi_i^{(t)}(a_i^+|s) \overline{A_i^{(t)}}(s, a_i^+) + \sum_{a_i \in I_-^{i,s}} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) \\ &\geq \sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) + \sum_{a_i \in B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s) \frac{-\Delta}{16|\mathcal{A}_i|} \\ &\quad + \pi_i^{(t)}(a_i^+|s) \overline{A_i^{(t)}}(s, a_i^+) + \sum_{a_i \in I_-^{i,s}} \frac{(1-\gamma)\Delta}{16|\mathcal{A}_i|} \pi_i^{(t)}(a_i^+|s) \overline{A_i^{(t)}}(s, a_i) \\ &\geq \sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) + |\mathcal{A}_i| \pi_i^{(t)}(a_i|s) \frac{-\Delta}{16|\mathcal{A}_i|} \\ &\quad + \pi_i^{(t)}(a_i^+|s) \frac{\Delta}{4} + |\mathcal{A}_i| \frac{(1-\gamma)\Delta}{16|\mathcal{A}_i|} \pi_i^{(t)}(a_i^+|s) \frac{-1}{1-\gamma} \\ &\geq \sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) + \pi_i^{(t)}(a_i^+|s) \frac{\Delta}{8} \\ \implies &\sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) < 0. \end{aligned}$$

Thus for  $t \geq \max\{T_2, T_3\}$ ,

$$\begin{aligned} \sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \theta_{s,a_i}^{(t+1)} &= \sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \theta_{s,a_i}^{(t)} + \eta \frac{1}{1-\gamma} d^{(t)}(s) \sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) \\ &< \sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \theta_{s,a_i}^{(t)}, \end{aligned}$$

which leads to the fact that  $\sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \theta_{s,a_i}^{(t)}$  is bounded from above. Further, from Lemma 17,  $\theta_{s,a_i^+}^{(t)}$  is bounded from below, thus the value

$$\frac{\sum_{a_i \in I_0^{i,s} \setminus B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s)}{\pi_i^{(t)}(a_i^+|s)}$$

is bounded from above. However from Corollary 16,

$$\lim_{t \rightarrow +\infty} \frac{\sum_{a_i \in I_0^{i,s}} \pi_i^{(t)}(a_i|s)}{\pi_i^{(t)}(a_i^+|s)} = +\infty.$$

Thus

$$\lim_{t \rightarrow +\infty} \frac{\sum_{a_i \in B_0^{i,s}(a_i^+)} \pi_i^{(t)}(a_i|s)}{\pi_i^{(t)}(a_i^+|s)} = +\infty,$$

which contradicts the fact that

$$\pi_i^{(t)}(a_i|s) \leq \pi_i^{(t)}(a_i^+|s), \quad \forall a_i \in B_0^{i,s}(a_i^+)$$

and finishes the proof by contradiction.  $\square$

Lemma 19 directly implies asymptotic convergence for gradient play as state in Theorem 4.

**Remark 5. (Discussion on the isolated stationary points assumption)** The proof of Theorem 4 resembles the technique used in [2] for the single agent case, which relies heavily on the fact that the sequence of  $Q$ -functions  $Q^{(t)}(s, a)$  obtains a limit  $Q^{(\infty)}(s, a)$ . The existence of such a limit in the single agent case follows from the monotonicity of the  $Q$ -functions. However, generalizing this proof to the multi-agent case requires the assumption that the sequence of averaged  $Q$ -functions  $\overline{Q_i^{(t)}}(s, a_i)$  (which can be non-monotonic, see, e.g., Figure 3 in Appendix) has a limit  $\overline{Q_i^{(\infty)}}(s, a_i)$ , which is not necessarily true in general. For instance, if the set of stationary policies  $\mathcal{SP} := \{\pi : \pi_i(a_i|s) \overline{A_i^\pi}(s, a_i) = 0, \forall s \in \mathcal{S}, a_i \in \mathcal{A}, i = 1, 2, \dots, n\}$  is not isolated, one cannot rule out the possibility that (natural) gradient play will not converge to a fixed point  $\pi^{(\infty)}$  (see e.g. [1] for counterexamples). Consequently,  $\overline{Q_i^{(t)}}(s, a_i)$  might not converge to a single value. For the above reasons, we assume the stationary policies are isolated to ensure that  $\pi^{(t)}$  converges to a fixed stationary policy  $\pi^{(\infty)}$  and thus  $\overline{Q_i^{(t)}}(s, a_i)$  obtains a limit. We believe that this assumption is a conservative condition that is sufficient to imply asymptotic convergence. It remains an interesting open question to establish convergence without this assumption.

## E.0.2 Asymptotic convergence for natural gradient play

The asymptotic convergence for natural gradient play is easier to establish compared with gradient play.

From Lemma 20 and the assumption that  $\phi(s, a)$  is upper-bounded, we know

$$\lim_{t \rightarrow +\infty} \sum_{a_i} \pi_i^{(t)}(a_i|s) \exp\left(\frac{\overline{\eta A_i^{(t)}}(s, a_i)}{1-\gamma}\right) = 1, \quad \forall s, i = 1, 2, \dots, n.$$



Since

$$\begin{aligned}
\sum_{a_i} \pi_i^{(t)}(a_i|s) \exp\left(\frac{\overline{\eta A_i^{(t)}}(s, a_i)}{1-\gamma}\right) &\geq \sum_{a_i} \pi_i^{(t)}(a_i|s) \left(1 + \left(\frac{\overline{\eta A_i^{(t)}}(s, a_i)}{1-\gamma}\right) + \frac{1}{4} \left(\frac{\overline{\eta A_i^{(t)}}(s, a_i)}{1-\gamma}\right)^2\right) \\
&\quad (\epsilon^x \geq 1 + x + \frac{x^2}{4} \text{ for } |x| \leq 1) \\
&= 1 + \frac{\eta^2}{4(1-\gamma)^2} \sum_{a_i} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i)^2 \\
&\implies \lim_{t \rightarrow +\infty} \sum_{a_i} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i)^2 = 0 \\
&\implies \lim_{t \rightarrow +\infty} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) = 0, \forall s, a_i, i = 1, 2, \dots, n \\
&\implies \lim_{t \rightarrow +\infty} \|\nabla_{\theta} \Phi(\theta^{(t)})\|_2 = 0
\end{aligned}$$

Similar to the proof for gradient play, from the assumption that stationary points are isolated, we can conclude that  $\pi^{(t)}$  converges to some stationary policy  $\pi^{(\infty)}$ , and we can define  $\overline{Q_i^{(\infty)}}(s, a_i), \overline{A_i^{(\infty)}}(s, a_i)$  accordingly. Asymptotic convergence is equivalent to

$$I_+^{i,s} := \left\{ a_i : \overline{A_i^{(\infty)}}(s, a_i) > 0 \right\} = \emptyset, \quad \forall s, i = 1, 2, \dots, n$$

We prove by contradiction. Suppose there exists  $a_i^+$  such that  $\overline{A_i^{(\infty)}}(s, a_i^+) > 0$ . From  $\lim_{t \rightarrow +\infty} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) = 0$ , we have that  $\lim_{t \rightarrow +\infty} \pi_i^{(t)}(a_i^+|s) = 0$ .

Select  $a_i^0$  such that  $\lim_{t \rightarrow +\infty} \pi_i^{(t)}(a_i^0|s) > 0$ . From  $\lim_{t \rightarrow +\infty} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) = 0$ , we have that  $\lim_{t \rightarrow +\infty} \overline{A_i^{(t)}}(s, a_i^0) = 0$ . Thus there exists  $\Delta > 0$  and  $T$  such that for  $t > T$ ,

$$\overline{A_i^{(t)}}(s, a_i^+) > \Delta, \overline{A_i^{(t)}}(s, a_i^0) < \frac{\Delta}{2}$$

Thus from natural gradient play scheme (10)

$$\frac{\pi_i^{(t)}(a_i^+|s)}{\pi_i^{(t)}(a_i^0|s)} = \frac{\pi_i^{(T)}(a_i^+|s)}{\pi_i^{(T)}(a_i^0|s)} \exp\left(\frac{\eta}{1-\gamma} \sum_{\tau=T}^{t-1} \overline{A_i^{(\tau)}}(s, a_i^+) - \overline{A_i^{(\tau)}}(s, a_i^0)\right) \geq \frac{\pi_i^{(T)}(a_i^+|s)}{\pi_i^{(T)}(a_i^0|s)},$$

which contradict the fact that  $\lim_{t \rightarrow +\infty} \frac{\pi_i^{(t)}(a_i^+|s)}{\pi_i^{(t)}(a_i^0|s)} = 0$ , and thus completes the proof.

## F Proof of Theorem 5

### F.1 Proof of Theorem 5 (Gradient play part)

*Proof of Theorem 5, gradient play.* From Lemma 29,  $\Phi$  is  $\beta$ -smooth with  $\beta = \frac{6n}{(1-\gamma)^3}$ , we have that:

$$\begin{aligned}
\Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)}) &\geq \left\langle \nabla \Phi(\theta^{(t)}), \theta^{(t+1)} - \theta^{(t)} \right\rangle - \frac{\beta}{2} \|\theta^{(t+1)} - \theta^{(t)}\|^2 \\
&= \left(\eta - \frac{\beta \eta^2}{2}\right) \|\nabla \Phi(\theta^{(t)})\|^2 \\
&\geq \frac{\eta}{2} \|\nabla \Phi(\theta^{(t)})\|^2
\end{aligned}$$

Summing over  $t$  we get:

$$\frac{\phi_{\max} - \phi_{\min}}{1 - \gamma} \geq \Phi(\theta^{(T)}) - \Phi(\theta^{(0)}) \geq \frac{\eta}{2} \sum_{t=0}^{T-1} \|\nabla \Phi(\theta^{(t)})\|^2$$

From Theorem 1 we have that

$$\|\nabla \Phi(\theta^{(t)})\| \geq \frac{c}{M \sqrt{\max_i |\mathcal{A}_i|}} \text{NE-gap}(\theta^{(t)})$$

Thus

$$\frac{1}{T} \sum_{t=0}^{T-1} \text{NE-gap}(\theta^{(t)})^2 \leq \frac{2 \max_i |\mathcal{A}_i| M^2 (\phi_{\max} - \phi_{\min})}{(1 - \gamma) c^2 \eta T}$$

which completes the proof.  $\square$

## F.2 Proof of Theorem 5 (Natural gradient play part)

**Lemma 20.** For  $\eta \leq \frac{(1-\gamma)^2}{2n(\phi_{\max} - \phi_{\min})}$ , running scheme (10) will guarantee that

$$\Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)}) \geq \frac{1}{\eta} \sum_{i=1}^n \sum_s d^{(t+1)}(s) \log Z_t^{i,s},$$

where  $Z_t^{i,s}$  is defined by

$$Z_t^{i,s} := \sum_{a_i} \pi_i^{(t)}(a_i|s) \exp\left(\frac{\eta \overline{A_{i,\phi}^{(t)}}(s, a_i)}{1 - \gamma}\right).$$

*Proof.* From performance difference lemma we have that

$$\Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)}) = \frac{1}{1 - \gamma} \sum_s d^{(t+1)}(s) \sum_a \left( \pi^{(t+1)}(a|s) - \pi^{(t)}(a|s) \right) A_\phi^{(t)}(s, a).$$

We define

$$\widetilde{A_{i,\phi}^{(t)}}(s, a_i) := \sum_{a_{-i}} \prod_{j=1}^{i-1} \pi_j^{(t+1)}(a_j|s) \prod_{j=i+1}^n \pi_j^{(t)}(a_j|s) A_\phi^{(t)}(s, a_i, a_{-i}). \quad (23)$$

Then

$$\begin{aligned} \Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)}) &= \frac{1}{1 - \gamma} \sum_s d^{(t+1)}(s) \sum_a \left( \pi^{(t+1)}(a|s) - \pi^{(t)}(a|s) \right) A_\phi^{(t)}(s, a) \\ &= \frac{1}{1 - \gamma} \sum_s d^{(t+1)}(s) \sum_a \sum_{i=1}^n \left( \prod_{j=1}^i \pi_j^{(t+1)}(a_j|s) \prod_{j=i+1}^n \pi_j^{(t)}(a_j|s) - \prod_{j=1}^{i-1} \pi_j^{(t+1)}(a_j|s) \prod_{j=i}^n \pi_j^{(t)}(a_j|s) \right) A_\phi^{(t)}(s, a) \\ &= \frac{1}{1 - \gamma} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \widetilde{A_{i,\phi}^{(t)}}(s, a_i) \\ &= \frac{1}{1 - \gamma} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \overline{A_{i,\phi}^{(t)}}(s, a_i) \\ &\quad + \frac{1}{1 - \gamma} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \left( \widetilde{A_{i,\phi}^{(t)}}(s, a_i) - \overline{A_{i,\phi}^{(t)}}(s, a_i) \right) \\ &= \underbrace{\frac{1}{1 - \gamma} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \sum_{a_i} \pi_i^{(t+1)}(a_i|s) \overline{A_{i,\phi}^{(t)}}(s, a_i)}_{\text{Part A}} \end{aligned}$$

$$+ \underbrace{\frac{1}{1-\gamma} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \left( \widetilde{A}_{i,\phi}^{(t)}(s, a_i) - \overline{A}_{i,\phi}^{(t)}(s, a_i) \right)}_{\text{Part B}}.$$

From scheme (10),

$$\overline{A}_{i,\phi}^{(t)}(s, a_i) = \frac{1-\gamma}{\eta} \left( \log \left( \frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} \right) + \log \left( Z_t^{i,s} \right) \right)$$

Substitute this into Part A, we have

$$\begin{aligned} \text{Part A} &= \frac{1}{\eta} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \sum_{a_i} \pi_i^{(t+1)}(a_i|s) \overline{A}_{i,\phi}^{(t)}(s, a_i) \\ &= \frac{1}{\eta} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \sum_{a_i} \pi_i^{(t+1)}(a_i|s) \left( \log \left( \frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} \right) + \log \left( Z_t^{i,s} \right) \right) \\ &= \frac{1}{\eta} \sum_s \sum_{i=1}^n d^{(t+1)}(s) \text{KL}(\pi_{i,s}^{(t+1)} || \pi_{i,s}^{(t)}) + \frac{1}{\eta} \sum_s \sum_{i=1}^n d^{(t+1)}(s) \log \left( Z_t^{i,s} \right). \end{aligned}$$

Further, we have that

$$\begin{aligned} & \left| \widetilde{A}_{i,\phi}^{(t)}(s, a_i) - \overline{A}_{i,\phi}^{(t)}(s, a_i) \right| \\ &= \left| \sum_{a_{-i}} \left( \prod_{j=1}^{i-1} \pi_j^{(t+1)}(a_j|s) - \prod_{j=1}^{i-1} \pi_j^{(t)}(a_j|s) \right) \prod_{j=i+1}^n \pi_j^{(t)}(a_j|s) A_\phi^{(t)}(s, a_i, a_{-i}) \right| \\ &\leq \frac{\phi_{\max} - \phi_{\min}}{1-\gamma} \sum_{j=1}^{i-1} \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \\ &\leq \frac{\phi_{\max} - \phi_{\min}}{1-\gamma} \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1. \end{aligned}$$

Thus

$$\begin{aligned} |\text{Part B}| &\leq \frac{1}{1-\gamma} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \sum_{a_i} \left| \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right| \left| \widetilde{A}_{i,\phi}^{(t)}(s, a_i) - \overline{A}_{i,\phi}^{(t)}(s, a_i) \right| \\ &\leq \frac{\phi_{\max} - \phi_{\min}}{(1-\gamma)^2} \sum_{i=1}^n \sum_s d^{(t+1)}(s) \sum_{a_i} \left| \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right| \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \\ &\leq \frac{\phi_{\max} - \phi_{\min}}{(1-\gamma)^2} \sum_s d^{(t+1)}(s) \left( \sum_{i=1}^n \|\pi_{i,s}^{(t+1)} - \pi_{i,s}^{(t)}\|_1 \right)^2 \\ &\leq \frac{n(\phi_{\max} - \phi_{\min})}{(1-\gamma)^2} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \|\pi_{i,s}^{(t+1)} - \pi_{i,s}^{(t)}\|_1^2 \\ &\leq \frac{2n(\phi_{\max} - \phi_{\min})}{(1-\gamma)^2} \sum_s d^{(t+1)}(s) \sum_{i=1}^n \text{KL}(\pi_{i,s}^{(t+1)} || \pi_{i,s}^{(t)}) \quad (\text{Pinsker's inequality}) \end{aligned}$$

Thus, when  $\eta \leq \frac{(1-\gamma)^2}{2n(\phi_{\max} - \phi_{\min})}$ , we have that

$$\begin{aligned} \Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)}) &= \text{Part A} + \text{Part B} \\ &\geq \left( \frac{1}{\eta} - \frac{2n(\phi_{\max} - \phi_{\min})}{(1-\gamma)^2} \right) \sum_s \sum_{i=1}^n d^{(t+1)}(s) \text{KL}(\pi_{i,s}^{(t+1)} || \pi_{i,s}^{(t)}) + \frac{1}{\eta} \sum_s \sum_{i=1}^n d^{(t+1)}(s) \log \left( Z_t^{i,s} \right) \end{aligned}$$

$$\geq \frac{1}{\eta} \sum_s \sum_{i=1}^n d^{(t+1)}(s) \log \left( Z_t^{i,s} \right),$$

which completes the proof.  $\square$

**Lemma 21.** For  $\eta \leq (1 - \gamma)^2$

$$\sum_{i=1}^n \sum_s d^{(t+1)}(s) \log Z_t^{i,s} \geq \frac{c\eta^2}{3M} \text{NE-gap}(\theta^{(t)})^2$$

*Proof.* From Lemma 13 we have that  $\text{NE-gap}(\theta) \leq \frac{1}{1-\gamma} \max_i \max_{s,a_i} \overline{A}_i^\theta(s, a_i)$ . On the other hand,

$$\begin{aligned} Z_t^{i,s} &= \sum_{a_i} \pi_i^{(t)}(a_i|s) \exp \left( \frac{\eta \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right) \\ &= \sum_{a_i \notin \arg\max_{a_i} \overline{Q}_i^{(t)}(s, a_i)} \pi_i^{(t)}(a_i|s) \exp \left( \frac{\eta \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right) + \sum_{a_i \in \arg\max_{a_i} \overline{Q}_i^{(t)}(s, a_i)} \pi_i^{(t)}(a_i|s) \exp \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right) \\ &\geq \sum_{a_i \notin \arg\max_{a_i} \overline{Q}_i^{(t)}(s, a_i)} \pi_i^{(t)}(a_i|s) \left( 1 + \frac{\eta \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right) \\ &\quad + \sum_{a_i \in \arg\max_{a_i} \overline{Q}_i^{(t)}(s, a_i)} \pi_i^{(t)}(a_i|s) \left( 1 + \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} + \frac{1}{2} \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right)^2 \right) \\ &= \sum_{a_i} \pi_i^{(t)}(a_i|s) + \frac{\eta}{1 - \gamma} \sum_{a_i} \pi_i^{(t)}(a_i|s) \overline{A}_i^{(t)}(s, a_i) + \frac{1}{2} \sum_{a_i \in \arg\max_{a_i} \overline{Q}_i^{(t)}(s, a_i)} \pi_i^{(t)}(a_i|s) \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right)^2 \\ &= 1 + \frac{1}{2} \sum_{a_i \in \arg\max_{a_i} \overline{Q}_i^{(t)}(s, a_i)} \pi_i^{(t)}(a_i|s) \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right)^2 \\ &\geq 1 + \frac{c}{2} \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right)^2. \end{aligned}$$

Thus

$$\log(Z_t^{i,s}) \geq \log \left( 1 + \frac{c}{2} \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right)^2 \right).$$

Because when  $\eta \leq (1 - \gamma)^2$ , we have  $\frac{c}{2} \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right)^2 \leq \frac{1}{2}$ , and that

$$\log(1 + x) \geq \frac{2}{3}x, \quad \text{for } 0 \leq x \leq \frac{1}{2},$$

thus

$$\log(Z_t^{i,s}) \geq \log \left( 1 + \frac{c}{2} \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right)^2 \right)$$

$$\geq \frac{c}{3} \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right)^2.$$

Thus

$$\begin{aligned} \sum_{i=1}^n \sum_s d^{(t+1)}(s) \log Z_t^{i,s} &\geq \frac{c}{3} \sum_{i=1}^n \sum_s d^{(t+1)}(s) \left( \frac{\eta \max_{a_i} \overline{A}_i^{(t)}(s, a_i)}{1 - \gamma} \right)^2 \\ &\geq \frac{c\eta^2}{3M(1-\gamma)^2} \max_i \max_s \max_{a_i} \overline{A}_i^{(t)}(s, a_i)^2 \\ &\geq \frac{c\eta^2}{3M} \text{NE-gap}(\theta^{(t)})^2. \end{aligned} \quad \square$$

We are now ready to prove the bound for natural gradient play in Theorem 5.

*Proof of Theorem 5, natural gradient play.* Combining Lemma 20 and 21 we have

$$\begin{aligned} \Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)}) &\geq \frac{1}{\eta} \sum_{i=1}^n \sum_s d^{(t+1)}(s) \log Z_t^{i,s} \\ &\geq \frac{c\eta}{3M} \text{NE-gap}(\theta^{(t)})^2 \end{aligned}$$

Summing over  $t$  we have

$$\frac{\phi_{\max} - \phi_{\min}}{1 - \gamma} \geq \Phi(\theta^{(T)}) - \Phi(\theta^{(0)}) \geq \frac{c\eta}{3M} \sum_{t=0}^{T-1} \text{NE-gap}(\theta^{(t)})^2,$$

thus

$$\frac{\sum_{t=0}^{T-1} \text{NE-gap}(\theta^{(t)})^2}{T} \leq \frac{3M(\phi_{\max} - \phi_{\min})}{(1 - \gamma)c\eta T},$$

which completes the proof.  $\square$

## G Proof for log-barrier regularization

### G.1 Proof of Theorem 6

We start with the following lemma:

**Lemma 22.** *Suppose  $\theta$  is such that  $\|\nabla_{\theta_i} \tilde{J}_i(\theta)\|_2 \leq \lambda$ , then  $\text{NE-gap}_i(\theta) \leq \lambda M |\mathcal{A}_i|$ , where  $M$  is defined as in Assumption 1.*

*Proof.* From  $\|\nabla_{\theta_i} \tilde{J}_i(\theta)\|_2 \leq \frac{\lambda}{2}$  we have that

$$\begin{aligned} \frac{\partial \tilde{J}_i(\theta)}{\partial \theta_{s, a_i}} &= \frac{1}{1 - \gamma} d_\theta(s) \pi_{\theta_i}(a_i | s) \overline{A}_i^\theta(s, a_i) + \lambda - \lambda |\mathcal{A}_i| \pi_{\theta_i}(a_i | s) \\ &= \pi_{\theta_i}(a_i | s) \left( \frac{1}{1 - \gamma} d_\theta(s) \overline{A}_i^\theta(s, a_i) - \lambda |\mathcal{A}_i| \right) + \lambda \leq \lambda \\ \implies \pi_{\theta_i}(a_i | s) \left( \frac{1}{1 - \gamma} d_\theta(s) \overline{A}_i^\theta(s, a_i) - \lambda |\mathcal{A}_i| \right) &\leq 0 \\ \implies \frac{1}{1 - \gamma} d_\theta(s) \overline{A}_i^\theta(s, a_i) - \lambda |\mathcal{A}_i| &\leq 0 \\ \implies \overline{A}_i^\theta(s, a_i) &\leq \frac{\lambda |\mathcal{A}_i| (1 - \gamma)}{d_\theta(s)} \leq \lambda |\mathcal{A}_i| (1 - \gamma) M \end{aligned}$$

Thus,

$$\begin{aligned}
\text{NE-gap}_i(\theta) &= \sup_{\theta_i^*} J_i(\theta_i^*, \theta_{-i}) - J_i(\theta_i, \theta_{-i}) = \frac{1}{1-\gamma} \sum_{s, a_i} d_{\theta^*}(s) \pi_{\theta_i^*}(a_i|s) \overline{A}_i^\theta(s, a_i) \\
&\leq \frac{1}{1-\gamma} \sum_{s, a_i} d_{\theta^*}(s) \max_{s, a_i} \overline{A}_i^\theta(s, a_i) \\
&\leq \frac{1}{1-\gamma} \sum_{s, a_i} d_{\theta^*}(s) \lambda |\mathcal{A}_i| (1-\gamma) M \\
&\leq \lambda |\mathcal{A}_i| M. \quad \square
\end{aligned}$$

Lemma 22 implies that any policy with gradient norm smaller than  $\lambda$  is also a  $\lambda M \max_i |\mathcal{A}_i|$ -NE. Thus by properly choosing  $\lambda$ , agents can find a  $\epsilon$ -NE by running gradient play.

We now prove Theorem 6.

*Proof of Theorem 6.* From Lemma 30,  $\tilde{\Phi}$  is  $\beta$ -smooth with  $\beta = \frac{6n}{(1-\gamma)^3} + 2\lambda \max_i |\mathcal{A}_i|$ , we have that:

$$\begin{aligned}
\tilde{\Phi}(\theta^{(t+1)}) - \tilde{\Phi}(\theta^{(t)}) &\geq \left\langle \nabla \tilde{\Phi}(\theta^{(t)}), \theta^{(t+1)} - \theta^{(t)} \right\rangle - \frac{\beta}{2} \|\theta^{(t+1)} - \theta^{(t)}\|^2 \\
&= \left( \eta - \frac{\beta \eta^2}{2} \right) \|\nabla \tilde{\Phi}(\theta^{(t)})\|^2 \\
&\geq \frac{\eta}{2} \|\nabla \tilde{\Phi}(\theta^{(t)})\|^2
\end{aligned}$$

For  $\theta^{(0)} = \mathbf{0}$ , summing over  $t$  we get:

$$\frac{\phi_{\max} - \phi_{\min}}{1-\gamma} \geq \tilde{\Phi}(\theta^{(T)}) - \tilde{\Phi}(\theta^{(0)}) \geq \frac{\eta}{2} \sum_{t=0}^{T-1} \|\nabla \tilde{\Phi}(\theta^{(t)})\|^2$$

Thus,

$$\min_{0 \leq t \leq T-1} \|\nabla \tilde{\Phi}(\theta^{(t)})\| \leq \frac{2(\phi_{\max} - \phi_{\min})}{(1-\gamma)\eta T}.$$

Thus for

$$\begin{aligned}
T &\geq \frac{2(\phi_{\max} - \phi_{\min})}{(1-\gamma)\eta \lambda^2} \\
&= \frac{2 \max_i |\mathcal{A}_i|^2 (\phi_{\max} - \phi_{\min}) M^2}{(1-\gamma)\eta \epsilon^2},
\end{aligned}$$

it can be guaranteed that

$$\min_{0 \leq t \leq T-1} \|\nabla \tilde{\Phi}(\theta^{(t)})\| \leq \lambda = \frac{\epsilon}{\max_i |\mathcal{A}_i| M}.$$

Then applying Lemma 22 completes the proof.  $\square$

## G.2 Proof of Theorem 7

For notational simplicity, we define the following variables:

$$\begin{aligned}
f_i^{(t)}(s, a_i) &:= \frac{1}{1-\gamma} \overline{A}_i^{(t)}(s, a_i) + \frac{\lambda}{d^{(t)}(s) \pi_i^{(t)}(a_i|s)} - \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} \\
Z_i^{i,s} &:= \sum_{a_i} \pi_i^{(t)}(a_i|s) \exp\left(\eta f_i^{(t)}(s, a_i)\right) \\
\Delta_i^{(t)}(s, a_i) &:= \frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} - 1
\end{aligned}$$

**Lemma 23.**

$$\begin{aligned}\sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i) &= 0 \\ \sum_{a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) &= 0 \\ Z_t^{i,s} &\geq 1\end{aligned}$$

*Proof.* From the definition of  $f_i^{(t)}(s, a_i)$ ,  $\Delta_i^{(t)}(s, a_i)$ ,

$$\begin{aligned}& \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i) \\ &= \frac{1}{1-\gamma} \sum_{a_i} \pi_i^{(t)}(a_i|s) \overline{A_i^{(t)}}(s, a_i) + \lambda \sum_{a_i} \pi_i^{(t)}(a_i|s) \frac{1}{d^{(t)}(s) \pi^{(t)}(a_i|s)} - \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} \sum_{a_i} \pi_i^{(t)}(a_i|s) \\ &= \frac{1}{d^{(t)}(s)} (\lambda |\mathcal{A}_i| - \lambda |\mathcal{A}_i|) = 0 \\ & \sum_{a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) = \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( \frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} - 1 \right) \\ &= \sum_{a_i} \pi_i^{(t+1)}(a_i|s) - \sum_{a_i} \pi_i^{(t)}(a_i|s) = 1 - 1 = 0\end{aligned}$$

Using the fact that  $e^x \geq 1 + x$ ,

$$\begin{aligned}Z_t^{i,s} &= \sum_{a_i} \pi_i^{(t)}(a_i|s) \exp\left(\eta f_i^{(t)}(s, a_i)\right) \\ &\geq \sum_{a_i} \pi_i^{(t)}(a_i|s) \left(1 + \eta f_i^{(t)}(s, a_i)\right) \\ &\geq \sum_{a_i} \pi_i^{(t)}(a_i|s) + \eta \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i) \geq 1.\end{aligned} \quad \square$$

**Lemma 24.** For  $\eta \leq \frac{1}{15 \left( \frac{1}{(1-\gamma)^2} + \lambda |\mathcal{A}_i| M \right)}$ ,  $\theta^{(0)} = \mathbf{0}$ , running scheme (17) will guarantee that

$$\pi_i^{(t)}(a_i|s) \geq \frac{\lambda}{4 \left( \lambda |\mathcal{A}_i| M + \frac{1}{(1-\gamma)^2} \right)}.$$

*Proof.* We will prove by induction. For  $\theta^0 = \mathbf{0}$  apparently  $\pi^{(0)}$  satisfies the lower bound. Suppose that

$$\pi_i^{(t)}(a_i|s) \geq \frac{\lambda}{4 \left( \lambda |\mathcal{A}_i| M + \frac{1}{(1-\gamma)^2} \right)},$$

then from the definition of  $f_i^{(t)}(s, a)$  we have that

$$-\frac{1}{(1-\gamma)^2} - \lambda |\mathcal{A}_i| M \leq f_i^{(t)}(s, a_i) \leq 5 \left( \frac{1}{(1-\gamma)^2} + \lambda |\mathcal{A}_i| M \right)$$

Thus

$$-\frac{1}{15} \leq \eta f_i^{(t)}(s, a_i) \leq \frac{1}{3},$$

which leads to the fact that

$$Z_t^{i,s} = \sum_{a_i} \pi_i^{(t)}(a_i|s) \exp\left(\eta f_i^{(t)}(s, a_i)\right)$$

$$\leq \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( 1 + \left( \eta f_i^{(t)}(s, a_i) \right) + \left( \eta f_i^{(t)}(s, a_i) \right)^2 \right) \quad (24)$$

$$(e^x \leq 1 + x + x^2, \text{ for } -\frac{1}{15} \leq x \leq \frac{1}{3})$$

$$= 1 + \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( \eta f_i^{(t)}(s, a_i) \right)^2 \quad (25)$$

$$\leq 1 + \frac{1}{3^2} = \frac{10}{9}.$$

Thus we have that

$$\frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} = \frac{\exp\left(\eta f_i^{(t)}(s, a_i)\right)}{Z_t^{i,s}} \geq \frac{1 + \eta f_i^{(t)}(s, a_i)}{Z_t^{i,s}} \geq \frac{1 - \frac{1}{15}}{\frac{10}{9}} = \frac{21}{25}.$$

Thus, for  $a_i$  such that  $\pi_i^{(t)}(a_i|s) \geq \frac{\lambda}{3\left(\lambda|\mathcal{A}_i|M + \frac{1}{(1-\gamma)^2}\right)}$ , we have

$$\pi_i^{(t+1)}(a_i|s) \geq \frac{21}{25} \frac{\lambda}{3\left(\lambda|\mathcal{A}_i|M + \frac{1}{(1-\gamma)^2}\right)} \geq \frac{\lambda}{4\left(\lambda|\mathcal{A}_i|M + \frac{1}{(1-\gamma)^2}\right)}.$$

On the other hand, for  $a_i$  such that  $\pi_i^{(t)}(a_i|s) < \frac{\lambda}{3\left(\lambda|\mathcal{A}_i|M + \frac{1}{(1-\gamma)^2}\right)}$ , we have

$$f_i^{(t)}(s, a_i) \geq -\frac{1}{(1-\gamma)^2} - \lambda|\mathcal{A}_i|M + 3\left(\frac{1}{(1-\gamma)^2} + \lambda|\mathcal{A}_i|M\right) = 2\left(\frac{1}{(1-\gamma)^2} + \lambda|\mathcal{A}_i|M\right),$$

From inequality (25) we have that

$$\begin{aligned} Z_{i,t} &\leq 1 + \eta^2 \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2 \\ &\leq 1 + 25\eta^2 \left( \frac{1}{(1-\gamma)^2} + \lambda|\mathcal{A}_i|M \right)^2 \\ &\leq 1 + \frac{5}{3}\eta \left( \frac{1}{(1-\gamma)^2} + \lambda|\mathcal{A}_i|M \right) \end{aligned}$$

Thus

$$\frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} = \frac{\exp\left(\eta f_i^{(t)}(s, a_i)\right)}{Z_t^{i,s}} \geq \frac{1 + \left(\eta f_i^{(t)}(s, a_i)\right)}{Z_t^{i,s}} \geq \frac{1 + 2\eta \left(\frac{1}{(1-\gamma)^2} + \lambda|\mathcal{A}_i|M\right)}{1 + \frac{5}{3}\eta \left(\frac{1}{(1-\gamma)^2} + \lambda|\mathcal{A}_i|M\right)} \geq 1,$$

then according to the induction assumption, we have

$$\pi_i^{(t+1)}(a_i|s) \geq \frac{\lambda}{4\left(\lambda|\mathcal{A}_i|M + \frac{1}{(1-\gamma)^2}\right)},$$

which completes the proof of the lemma.  $\square$

**Corollary 25.** Under the condition of Lemma 24, running (17) will guarantee that

$$-\frac{1}{15} \leq \eta f_i^{(t)}(s, a_i) \leq \frac{1}{3}, \quad Z_t^{i,s} \leq \frac{10}{9}, \quad -\frac{1}{5} \leq \Delta_i^{(t)}(s, a_i) \leq \frac{1}{2},$$

*Proof.* The first two inequalities are proved in the proof of Lemma 24, we only need to show  $-\frac{1}{5} \leq \Delta_i^{(t)}(s, a_i) \leq \frac{1}{2}$ . In the proof of Lemma 24, we have already shown that

$$\frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} = \frac{\exp\left(\eta f_i^{(t)}(s, a_i)\right)}{Z_t^{i,s}} \geq \frac{1 + \eta f_i^{(t)}(s, a_i)}{Z_t^{i,s}} \geq \frac{1 - \frac{1}{15}}{\frac{10}{9}} = \frac{21}{25} \geq \frac{4}{5},$$



thus

$$\Delta_i^{(t)}(s, a_i) \geq \frac{4}{5} - 1 = -\frac{1}{5}.$$

On the other hand,

$$\frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} = \frac{\exp\left(\eta f_i^{(t)}(s, a_i)\right)}{Z_t^{i,s}} \leq \exp\left(\eta f_i^{(t)}(s, a_i)\right) \leq \exp\left(\frac{1}{3}\right) \leq \frac{3}{2},$$

thus

$$\Delta_i^{(t)}(s, a_i) \leq \frac{3}{2} - 1 = \frac{1}{2},$$

which completes the proof of the corollary.  $\square$

**Lemma 26.**

$$\tilde{\Phi}(\theta^{(t+1)}) - \tilde{\Phi}(\theta^{(t)}) \geq \left( \frac{1}{2\eta} - 4\lambda \max_i |\mathcal{A}_i| M^2 - \frac{4M}{(1-\gamma)^2} - \frac{3nM}{(1-\gamma)^3} \right) \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2$$

*Proof.* Let  $\tilde{\theta}^{i,(t)}$  be defined as:

$$\tilde{\theta}^{i,(t)} := \left( \theta_1^{(t)}, \dots, \theta_{i-1}^{(t)}, \theta_i^{(t+1)}, \dots, \theta_n^{(t+1)} \right).$$

Then we have that

$$\begin{aligned} \Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)}) &= \sum_{i=1}^n \Phi(\tilde{\theta}^{i,(t)}) - \Phi(\tilde{\theta}^{i+1,(t)}) \\ &= \sum_{i=1}^n J_i(\tilde{\theta}^{i,(t)}) - J_i(\tilde{\theta}^{i+1,(t)}) \\ &= \frac{1}{1-\gamma} \sum_{i=1}^n \sum_s d_{\tilde{\theta}^{i,(t)}}(s) \sum_{a_i} \pi_i^{(t+1)}(a_i|s) \overline{A_i^{\tilde{\theta}^{i+1,(t)}}}(s, a_i) \quad (\text{Lemma 8}) \\ &= \frac{1}{1-\gamma} \sum_{i=1}^n \sum_s d_{\tilde{\theta}^{i,(t)}}(s) \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \overline{Q_i^{\tilde{\theta}^{i+1,(t)}}}(s, a_i). \end{aligned}$$

Thus

$$\begin{aligned} \tilde{\Phi}(\theta^{(t+1)}) - \tilde{\Phi}(\theta^{(t)}) &= \Phi(\theta^{(t+1)}) - \Phi(\theta^{(t)}) + \lambda \sum_i \sum_{s, a_i} \log \left( \frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} \right) \\ &= \frac{1}{1-\gamma} \sum_{i=1}^n \sum_s d_{\tilde{\theta}^{i,(t)}}(s) \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \overline{Q_i^{\tilde{\theta}^{i+1,(t)}}}(s, a_i) + \lambda \sum_i \sum_{s, a_i} \log \left( 1 + \Delta^{(t)}(s, a_i) \right) \\ &= \frac{1}{1-\gamma} \sum_{i=1}^n \sum_s d^{(t)}(s) \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \overline{Q_i^{(t)}}(s, a_i) \\ &\quad + \frac{1}{1-\gamma} \sum_{i=1}^n \sum_s d^{(t)}(s) \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \left( \overline{Q_i^{\tilde{\theta}^{i+1,(t)}}}(s, a_i) - \overline{Q_i^{(t)}}(s, a_i) \right) \\ &\quad + \frac{1}{1-\gamma} \sum_{i=1}^n \sum_s \left( d_{\tilde{\theta}^{i,(t)}}(s) - d^{(t)}(s) \right) \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \overline{Q_i^{\tilde{\theta}^{i+1,(t)}}}(s, a_i) \\ &\quad + \lambda \sum_i \sum_{s, a_i} \log \left( 1 + \Delta^{(t)}(s, a_i) \right) \end{aligned}$$

$$\begin{aligned}
&= \underbrace{\frac{1}{1-\gamma} \sum_{i=1}^n \sum_s d^{(t)}(s) \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \overline{Q_i^{(t)}}(s, a_i) + \lambda \Delta_i^{(t)}(s, a_i)}_{\text{Part A}} \\
&+ \underbrace{\lambda \sum_i \sum_{s, a_i} \log \left( 1 + \Delta_i^{(t)}(s, a_i) \right) - \Delta_i^{(t)}(s, a_i)}_{\text{Part B}} \\
&+ \underbrace{\frac{1}{1-\gamma} \sum_{i=1}^n \sum_s d^{(t)}(s) \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \left( \overline{Q_i^{\tilde{\theta}^{i+1, (t)}}}(s, a_i) - \overline{Q_i^{(t)}}(s, a_i) \right)}_{\text{Part C}} \\
&+ \underbrace{\frac{1}{1-\gamma} \sum_{i=1}^n \sum_s \left( d_{\tilde{\theta}^{i, (t)}}(s) - d^{(t)}(s) \right) \sum_{a_i} \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \overline{Q_i^{\tilde{\theta}^{i+1, (t)}}}(s, a_i)}_{\text{Part D}}.
\end{aligned}$$

We will now bound each part separately. We first get a lower bound for part A:

$$\begin{aligned}
\text{Part A} &= \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \left[ \left( \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \left( \frac{1}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) \right) + \frac{\lambda}{d^{(t)}(s)} \Delta_i^{(t)}(s, a_i) \right] \\
&= \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \left[ \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) \left( \frac{1}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) \right) + \frac{\lambda}{d^{(t)}(s)} \Delta_i^{(t)}(s, a_i) \right] \\
&= \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \left[ \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) \left( \frac{1}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) - \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} \right) + \frac{\lambda}{d^{(t)}(s)} \Delta_i^{(t)}(s, a_i) \right] \\
&= \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) \left( \frac{1}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) + \frac{\lambda}{d^{(t)}(s) \pi_i^{(t)}(a_i|s)} - \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} \right) \\
&= \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) f_i^{(t)}(s, a_i) \\
&= \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) \frac{1}{\eta} \left( \log \left( \frac{\pi_i^{(t+1)}(a_i|s)}{\pi_i^{(t)}(a_i|s)} \right) + \log(Z_t^{i, s}) \right) \\
&= \frac{1}{\eta} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) \log \left( 1 + \Delta_i^{(t)}(s, a_i) \right) \\
&\quad + \frac{1}{\eta} \sum_{i=1}^n \sum_s d^{(t)}(s) \log(Z_t^{i, s}) \sum_{a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) \\
&= \frac{1}{\eta} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i) \log \left( 1 + \Delta_i^{(t)}(s, a_i) \right) \\
&= \frac{1}{\eta} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \left| \Delta_i^{(t)}(s, a_i) \right| \left| \log \left( 1 + \Delta_i^{(t)}(s, a_i) \right) \right|
\end{aligned}$$

From the boundedness of  $\Delta_i^{(t)}(s, a_i)$  in Corollary 25, we have that

$$\left| \log \left( 1 + \Delta_i^{(t)}(s, a_i) \right) \right| \geq \frac{1}{2} \left| \Delta_i^{(t)}(s, a_i) \right|$$

Substitute this into the above inequalities, we get

$$\text{Part A} \geq \frac{1}{2\eta} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2$$

Now we will give a lower bound for part B. Similarly, from the boundedness of  $\Delta_i^{(t)}(s, a_i)$  in Corollary 25, we have that

$$\log \left( 1 + \Delta_i^{(t)}(s, a_i) \right) - \Delta_i^{(t)}(s, a_i) \geq -\Delta_i^{(t)}(s, a_i)^2.$$

Thus

$$\text{Part B} = \lambda \sum_i \sum_{s, a_i} \log \left( 1 + \Delta_i^{(t)}(s, a_i) \right) - \Delta_i^{(t)}(s, a_i) \geq -\lambda \sum_i \sum_{s, a_i} \Delta_i^{(t)}(s, a_i)^2.$$

Additionally, using Lemma 24,

$$\begin{aligned} \text{Part B} &\geq -\lambda \sum_i \sum_{s, a_i} \Delta_i^{(t)}(s, a_i)^2 \\ &\geq -4 \left( \lambda \max_i |\mathcal{A}_i| M + \frac{1}{(1-\gamma)^2} \right) \sum_i \sum_{s, a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2 \\ &\geq -4M \left( \lambda \max_i |\mathcal{A}_i| M + \frac{1}{(1-\gamma)^2} \right) \sum_i \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2. \end{aligned}$$

We will now move on to bound the absolute value of Part C.

$$|\text{Part C}| \leq \frac{1}{1-\gamma} \sum_{i=1}^n \sum_s d^{(t)}(s) \sum_{a_i} \left| \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right| \left| \overline{Q_i^{\tilde{\theta}^{i+1, (t)}}}(s, a_i) - \overline{Q_i^{(t)}}(s, a_i) \right|.$$

Since

$$\begin{aligned} \left| \overline{Q_i^{\tilde{\theta}^{i+1, (t)}}}(s, a_i) - \overline{Q_i^{(t)}}(s, a_i) \right| &\leq \sum_{a_{-i}} \pi_{\tilde{\theta}_{-i}^{i+1, (t)}}(a_{-i}|s) \left| Q^{\tilde{\theta}^{i+1, (t)}}(s, a_i, a_{-i}) - Q^{(t)}(s, a_i, a_{-i}) \right| \\ &\quad + \sum_{a_{-i}} \left| \pi_{\tilde{\theta}_{-i}^{i+1, (t)}}(a_{-i}|s) - \pi_{-i}^{(t)}(a_{-i}|s) \right| \left| Q^{(t)}(s, a_i) \right| \\ &\leq \max_a \left| Q^{\tilde{\theta}^{i+1, (t)}}(s, a) - Q^{(t)}(s, a) \right| + \frac{1}{1-\gamma} \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1. \end{aligned}$$

From Lemma 32,

$$\begin{aligned} \max_a \left| Q^{\tilde{\theta}^{i+1, (t)}}(s, a) - Q^{(t)}(s, a) \right| &\leq \frac{1}{(1-\gamma)^2} \max_s \|\pi_{\tilde{\theta}_s^{i+1, (t)}} - \pi_{\theta_s^{(t)}}\|_1 \\ &\leq \frac{1}{(1-\gamma)^2} \max_s \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \end{aligned}$$

Thus we have that

$$\begin{aligned} \left| \overline{Q_i^{\tilde{\theta}^{i+1, (t)}}}(s, a_i) - \overline{Q_i^{(t)}}(s, a_i) \right| &\leq \max_a \left| Q^{\tilde{\theta}^{i+1, (t)}}(s, a) - Q^{(t)}(s, a) \right| + \frac{1}{1-\gamma} \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \\ &\leq \frac{2}{(1-\gamma)^2} \max_s \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1. \end{aligned}$$

Thus

$$\begin{aligned}
|\text{Part C}| &\leq \frac{1}{1-\gamma} \sum_{i=1}^n \sum_s d^{(t)}(s) \sum_{a_i} \left| \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right| \frac{2}{(1-\gamma)^2} \max_s \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \\
&\leq \frac{2}{(1-\gamma)^3} \left( \max_s \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \right) \cdot \sum_{i=1}^n \sum_s d^{(t)}(s) \sum_{a_i} \left| \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right| \\
&\leq \frac{2}{(1-\gamma)^3} \left( \max_s \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \right) \cdot \sum_s d^{(t)}(s) \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \\
&\leq \frac{2}{(1-\gamma)^3} \left( \max_s \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \right)^2
\end{aligned}$$

From Cauchy-Schwarz inequality,

$$\begin{aligned}
\left( \sum_{j=1}^n \|\pi_{j,s}^{(t+1)} - \pi_{j,s}^{(t)}\|_1 \right)^2 &= \left( \sum_{j=1}^n \sum_{a_j} \pi_i^{(t)}(a_j|s) \left| \Delta_j^{(t)}(s, a_j) \right| \right)^2 \\
&\leq \left( \sum_{j=1}^n \sum_{a_j} \pi_i^{(t)}(a_j|s) \right) \left( \sum_{j=1}^n \sum_{a_j} \pi_j^{(t)}(a_j|s) \Delta_j^{(t)}(s, a_j)^2 \right) \\
&= n \sum_{j=1}^n \sum_{a_j} \pi_j^{(t)}(a_j|s) \Delta_j^{(t)}(s, a_j)^2
\end{aligned}$$

Thus

$$\begin{aligned}
|\text{Part C}| &\leq \frac{2n}{(1-\gamma)^3} \sum_{i=1}^n \sum_{a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2 \\
&\leq \frac{2nM}{(1-\gamma)^3} \sum_{i=1}^n \sum_{a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2
\end{aligned}$$

Lastly, we will bound the absolute value of Part D.

$$\begin{aligned}
|\text{Part D}| &= \left| \frac{1}{1-\gamma} \sum_{i=1}^n \sum_s \left( d_{\tilde{\theta}^{i,(t)}}(s) - d^{(t)}(s) \right) \sum_{a_i} \left( \pi^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right) \overline{Q_i^{\tilde{\theta}^{i+1,(t)}}(s, a_i)} \right| \\
&\leq \frac{1}{(1-\gamma)^2} \sum_{i=1}^n \sum_s \left| d_{\tilde{\theta}^{i,(t)}}(s) - d^{(t)}(s) \right| \sum_{a_i} \left| \pi^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right| \\
&\leq \frac{1}{(1-\gamma)^2} \sum_{i=1}^n \max_s \|\pi_{i,s}^{(t+1)} - \pi_{i,s}^{(t)}\|_1 \sum_s \left| d_{\tilde{\theta}^{i,(t)}}(s) - d^{(t)}(s) \right|.
\end{aligned}$$

From Corollary 34

$$\begin{aligned}
\sum_s \left| d_{\tilde{\theta}^{i,(t)}}(s) - d^{(t)}(s) \right| &\leq \frac{1}{1-\gamma} \max_s \left\| \pi_{\tilde{\theta}^{i,(t)}}(a|s) - \pi^{(t)}(a|s) \right\|_1 \\
&\leq \frac{1}{1-\gamma} \max_s \sum_{i=1}^n \|\pi_{i,s}^{(t+1)} - \pi_{i,s}^{(t)}\|_1.
\end{aligned}$$

Thus

$$\begin{aligned}
|\text{Part D}| &\leq \frac{1}{(1-\gamma)^2} \sum_{i=1}^n \max_s \|\pi_{i,s}^{(t+1)} - \pi_{i,s}^{(t)}\|_1 \sum_s \left| d_{\tilde{\theta}^{(t)}}(s) - d^{(t)}(s) \right| \\
&\leq \frac{1}{(1-\gamma)^3} \left( \sum_{i=1}^n \max_s \|\pi_{i,s}^{(t+1)} - \pi_{i,s}^{(t)}\|_1 \right) \left( \max_s \sum_{i=1}^n \|\pi_{i,s}^{(t+1)} - \pi_{i,s}^{(t)}\|_1 \right) \\
&\leq \frac{1}{(1-\gamma)^3} \left( \sum_{i=1}^n \max_s \|\pi_{i,s}^{(t+1)} - \pi_{i,s}^{(t)}\|_1 \right)^2
\end{aligned}$$

From Cauchy-Schwarz inequality

$$\begin{aligned}
\left( \sum_{i=1}^n \max_s \|\pi_{i,s}^{(t+1)} - \pi_{i,s}^{(t)}\|_1 \right)^2 &\leq n \sum_{i=1}^n \max_s \left( \sum_{a_i} \left| \pi_i^{(t+1)}(a_i|s) - \pi_i^{(t)}(a_i|s) \right| \right)^2 \\
&= n \sum_{i=1}^n \max_s \left( \sum_{a_i} \pi_i^{(t)}(a_i|s) \left| \Delta_i^{(t)}(a_i|s) \right| \right)^2 \\
&\leq n \sum_{i=1}^n \max_s \left( \sum_{a_i} \pi_i^{(t)}(a_i|s) \right) \left( \sum_{a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(a_i|s)^2 \right) \\
&\leq n \sum_{i=1}^n \max_s \sum_{a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(a_i|s)^2 \\
&\leq n \sum_{i=1}^n \sum_{s, a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(a_i|s)^2.
\end{aligned}$$

Thus

$$\begin{aligned}
|\text{Part D}| &\leq \frac{n}{(1-\gamma)^3} \sum_{i=1}^n \sum_{s, a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(a_i|s)^2 \\
&\leq \frac{nM}{(1-\gamma)^3} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(a_i|s)^2,
\end{aligned}$$

Combining the bounds on Part A,B,C,D we get

$$\begin{aligned}
\tilde{\Phi}(\theta^{(t+1)}) - \tilde{\Phi}(\theta^{(t)}) &= \text{Part A} + \text{Part B} + \text{Part C} + \text{Part D} \\
&\geq \left( \frac{1}{2\eta} - 4\lambda \max_i |\mathcal{A}_i| M^2 - \frac{4M}{(1-\gamma)^2} - \frac{3nM}{(1-\gamma)^3} \right) \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2,
\end{aligned}$$

which completes the proof.  $\square$

**Lemma 27.** *Under the condition as in Lemma 24,*

$$\sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2 \geq \frac{\eta^2}{9} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2$$

*Proof.* Recall from the definition of  $\Delta_i^{(t)}(s, a_i)$ :

$$\Delta_i^{(t)}(s, a_i) = \frac{\exp\left(\eta f_i^{(t)}(s, a_i)\right)}{Z_t^{i,s}} - 1.$$

Thus

$$\begin{aligned}
\sum_{a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2 &= \frac{1}{\left(Z_t^{i,s}\right)^2} \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( \exp\left(\eta f_i^{(t)}(s, a_i)\right) - Z_t^{i,s} \right)^2 \\
&= \frac{1}{\left(Z_t^{i,s}\right)^2} \left[ \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( \exp\left(\eta f_i^{(t)}(s, a_i)\right) - 1 \right)^2 \right. \\
&\quad \left. - 2 \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( \exp\left(\eta f_i^{(t)}(s, a_i)\right) - 1 \right) \left( Z_t^{i,s} - 1 \right) + \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( Z_t^{i,s} - 1 \right)^2 \right] \\
&= \frac{1}{\left(Z_t^{i,s}\right)^2} \left[ \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( \exp\left(\eta f_i^{(t)}(s, a_i)\right) - 1 \right)^2 - \left( Z_t^{i,s} - 1 \right)^2 \right]
\end{aligned}$$

Since  $|e^x - 1| \geq \frac{|x|}{2}$  for  $x \geq -1$ , we have that

$$\begin{aligned}
\sum_{a_i} \pi_i^{(t)}(a_i|s) \left( \exp\left(\eta f_i^{(t)}(s, a_i) - 1\right) \right)^2 &\geq \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( \frac{\eta}{2} f_i^{(t)}(s, a_i) \right)^2 \\
&\geq \frac{\eta^2}{4} \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2
\end{aligned}$$

Additionally, as is proved in Lemma 24,

$$\begin{aligned}
Z_t^{i,s} &= \sum_{a_i} \pi_i^{(t)}(a_i|s) \exp\left(\eta f_i^{(t)}(s, a_i)\right) \\
&\leq \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( 1 + \left(\eta f_i^{(t)}(s, a_i)\right) + \left(\eta f_i^{(t)}(s, a_i)\right)^2 \right) \\
&\leq 1 + \eta^2 \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2.
\end{aligned}$$

Thus

$$\begin{aligned}
\sum_{a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2 &= \frac{1}{\left(Z_t^{i,s}\right)^2} \left[ \sum_{a_i} \pi_i^{(t)}(a_i|s) \left( \exp\left(\eta f_i^{(t)}(s, a_i)\right) - 1 \right)^2 - \left( Z_t^{i,s} - 1 \right)^2 \right] \\
&\geq \frac{1}{\left(Z_t^{i,s}\right)^2} \left[ \frac{\eta^2}{4} \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2 - \left( \eta^2 \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2 \right)^2 \right] \\
&= \frac{1}{\left(Z_t^{i,s}\right)^2} \left( \eta^2 \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2 \right) \left( \frac{1}{4} - \eta^2 \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2 \right).
\end{aligned}$$

From Corollary 25

$$\begin{aligned}
-\frac{1}{15} &\leq \eta f_i^{(t)}(s, a_i) \leq \frac{1}{3} \\
\implies \eta^2 \sum_{a_i} \pi_i^{(t)}(s, a_i) f_i^{(t)}(s, a_i)^2 &\leq \frac{1}{9}.
\end{aligned}$$

Thus

$$\sum_{a_i} \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2 \frac{1}{\left(Z_t^{i,s}\right)^2} \left( \eta^2 \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2 \right) \left( \frac{1}{4} - \eta^2 \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2 \right)$$

$$\begin{aligned}
&\geq \frac{1}{\left(Z_t^{i,s}\right)^2} \left( \eta^2 \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2 \right) \left( \frac{1}{4} - \frac{1}{9} \right) \\
&\geq \left( \frac{9}{10} \right)^2 \left( \frac{1}{4} - \frac{1}{9} \right) \left( \eta^2 \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2 \right) \\
&\geq \frac{\eta^2}{9} \sum_{a_i} \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2,
\end{aligned}$$

Thus

$$\sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2 \geq \frac{\eta^2}{9} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2$$

which completes the proof.  $\square$

**Lemma 28.**

$$\text{NE-gap}(\theta^{(t)}) \leq \frac{\sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2}{4\lambda} + \lambda \max_i |\mathcal{A}_i| M,$$

where  $M = \sup_{\theta} \max_s \frac{1}{d_{\theta}(s)}$ .

*Proof.* We will now prove the lemma.

$$\begin{aligned}
d^{(t)}(s) \pi_i^{(t)} f_i^{(t)}(s, a_i)^2 &= d^{(t)}(s) \pi_i^{(t)}(a_i|s) \left( \frac{1}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) + \lambda \frac{1}{d^{(t)}(s) \pi_i^{(t)}(a_i|s)} - \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} \right)^2 \\
&= d^{(t)}(s) \pi_i^{(t)}(a_i|s) \left( \frac{1}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) - \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} \right)^2 + \frac{\lambda^2}{d^{(t)}(s) \pi_i^{(t)}(a_i|s)} + 2\lambda \left( \frac{1}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) - \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} \right) \\
&\geq 4\lambda \left( \frac{1}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) - \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} \right) \\
\implies \frac{1}{1-\gamma} \overline{A_i^{(t)}}(s, a_i) &\leq \frac{d^{(t)}(s) \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2}{4\lambda} + \frac{\lambda |\mathcal{A}_i|}{d^{(t)}(s)} \\
&\leq \frac{\sum_i \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2}{4\lambda} + \lambda \max_i |\mathcal{A}_i| M.
\end{aligned}$$

Thus from Lemma 13,

$$\text{NE-gap}_i(\theta^{(t)}) \leq \frac{1}{1-\gamma} \max_{s, a_i} \overline{A_i^{(t)}}(s, a_i) \leq \frac{\sum_i \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2}{4\lambda} + \lambda \max_i |\mathcal{A}_i| M,$$

which completes the proof.  $\square$

We are now ready to prove Theorem 7.

*Proof of Theorem 7.* From Lemma 26 we have that for

$$\eta \leq \min \left\{ \frac{1}{15 \left( \frac{1}{(1-\gamma)^2} + \lambda |\mathcal{A}_i| M \right)}, \frac{1}{4 \left( 4\lambda \max_i |\mathcal{A}_i| M^2 + \frac{4M}{(1-\gamma)^2} + \frac{3nM}{(1-\gamma)^3} \right)} \right\},$$

$$\tilde{\Phi}(\theta^{(t+1)}) - \tilde{\Phi}(\theta^{(t)}) \geq \frac{1}{4\eta} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2.$$

From Lemma 27,

$$\begin{aligned}\tilde{\Phi}(\theta^{(t+1)}) - \tilde{\Phi}(\theta^{(t)}) &\geq \frac{1}{4\eta} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) \Delta_i^{(t)}(s, a_i)^2 \\ &\geq \frac{\eta}{36} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2\end{aligned}$$

Thus by telescoping we have

$$\frac{\sum_{t=0}^{T-1} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2}{T} \leq \frac{36 \left( \tilde{\Phi}(\theta^{(T)}) - \tilde{\Phi}(\theta^{(0)}) \right)}{\eta T}.$$

From Lemma 28,

$$\begin{aligned}\frac{\sum_{t=0}^{T-1} \text{NE-gap}(\theta^{(t)})}{T} &\leq \frac{1}{4\lambda} \frac{\sum_{t=0}^{T-1} \sum_{i=1}^n \sum_{s, a_i} d^{(t)}(s) \pi_i^{(t)}(a_i|s) f_i^{(t)}(s, a_i)^2}{T} + \lambda \max_i |\mathcal{A}_i| M \\ &\leq \frac{9 \left( \tilde{\Phi}(\theta^{(T)}) - \tilde{\Phi}(\theta^{(0)}) \right)}{\eta \lambda T} + \lambda \max_i |\mathcal{A}_i| M.\end{aligned}$$

Specifically, set  $\lambda = \frac{\epsilon}{2 \max_i |\mathcal{A}_i| M}$  and  $\theta^{(0)} = \mathbf{0}$ , then for any

$$\begin{aligned}T &\geq \frac{18 \left( \tilde{\Phi}(\theta^{(T)}) - \tilde{\Phi}(\theta^{(0)}) \right)}{(1-\gamma)\eta\lambda\epsilon} = \frac{36 \max_i |\mathcal{A}_i| (\phi_{\max} - \phi_{\min}) M}{(1-\gamma)\eta\epsilon^2} \\ &\geq O\left(\frac{n \max_i |\mathcal{A}_i| (\phi_{\max} - \phi_{\min}) M^2}{(1-\gamma)^4 \epsilon^2}\right),\end{aligned}$$

we have

$$\frac{\sum_{t=0}^{T-1} \text{NE-gap}(\theta^{(t)})}{T} \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon,$$

which completes the proof.  $\square$

## H Smoothness Proofs

This section mainly focuses on the smoothness of  $\Phi$  and  $\tilde{\Phi}$ . We first state the smoothness results in Lemma 29 and Lemma 30. The auxiliary lemmas used during proof of the above two lemmas are stated in Lemma 31 and Appendix I.

**Lemma 29** (Smoothness of  $\Phi(\theta)$ ).

$$\|\nabla_{\theta} \Phi(\theta') - \nabla_{\theta} \Phi(\theta)\|_2 \leq \frac{6n}{(1-\gamma)^3} \|\theta' - \theta\|_2$$

*Proof.* From Lemma 31 we have that

$$\begin{aligned}\|\nabla_{\theta} \Phi(\theta') - \nabla_{\theta} \Phi(\theta)\|_2^2 &= \sum_{i=1}^n \|\nabla_{\theta_i} \Phi(\theta') - \nabla_{\theta_i} \Phi(\theta)\|_2^2 \\ &\leq \sum_{i=1}^n \|\nabla_{\theta_i} \Phi(\theta') - \nabla_{\theta_i} \Phi(\theta)\|_1^2 \\ &\leq \sum_{i=1}^n \left( \frac{6}{(1-\gamma)^3} \sum_{i=1}^n \|\theta'_i - \theta_i\|_2 \right)^2 \\ &= \frac{36n}{(1-\gamma)^6} \left( \sum_{i=1}^n \|\theta'_i - \theta_i\|_2 \right)^2\end{aligned}$$



$$\begin{aligned}
&\leq \frac{36n^2}{(1-\gamma)^6} \sum_{i=1}^n \|\theta'_i - \theta_i\|_2^2 \\
&= \frac{36n^2}{(1-\gamma)^6} \|\theta' - \theta\|_2^2,
\end{aligned}$$

thus

$$\|\nabla_{\theta} \Phi(\theta') - \nabla_{\theta} \Phi(\theta)\|_2 \leq \frac{6n}{(1-\gamma)^3} \|\theta' - \theta\|_2$$

□

**Lemma 30** (Smoothness of  $\tilde{\Phi}(\theta)$ ).

$$\left\| \nabla_{\theta} \tilde{\Phi}(\theta') - \nabla_{\theta} \tilde{\Phi}(\theta) \right\|_2 \leq \left( \frac{6n}{(1-\gamma)^3} + 2\lambda \max_i |\mathcal{A}_i| \right) \|\theta' - \theta\|_2$$

*Proof.* Since

$$\frac{\partial \sum_{i=1}^n \sum_{s, a_i} \log \pi_{\theta_i}(a_i|s)}{\partial \theta_{s, a_i}} = 1 - |\mathcal{A}_i| \pi_{\theta_i}(a_i|s)$$

we have that

$$\begin{aligned}
&\left\| \nabla_{\theta} \left( \sum_{i=1}^n \sum_{s, a_i} \log \pi_{\theta'_i}(a_i|s) - \sum_{i=1}^n \sum_{s, a_i} \log \pi_{\theta_i}(a_i|s) \right) \right\|_2^2 = \sum_{i=1}^n |\mathcal{A}_i|^2 \sum_s \sum_{a_i} (\pi_{\theta'_i}(a_i|s) - \pi_{\theta_i}(a_i|s))^2 \\
&\leq \sum_{i=1}^n |\mathcal{A}_i|^2 \sum_s \|\pi_{\theta'_i, s} - \pi_{\theta_i, s}\|_1^2 \stackrel{\text{(Corollary 37)}}{\leq} 4 \sum_{i=1}^n |\mathcal{A}_i|^2 \sum_s \|\theta'_{i, s} - \theta_{i, s}\|_2^2 \leq 4 \max_i |\mathcal{A}_i|^2 \|\theta' - \theta\|_2^2
\end{aligned}$$

Thus

$$\begin{aligned}
&\left\| \nabla_{\theta} \tilde{\Phi}(\theta') - \nabla_{\theta} \tilde{\Phi}(\theta) \right\|_2 \\
&\leq \|\nabla_{\theta} \Phi(\theta') - \nabla_{\theta} \Phi(\theta)\|_2 + \lambda \left\| \nabla_{\theta} \left( \sum_{i=1}^n \sum_{s, a_i} \log \pi_{\theta'_i}(a_i|s) - \sum_{i=1}^n \sum_{s, a_i} \log \pi_{\theta_i}(a_i|s) \right) \right\|_2 \\
&\leq \left( \frac{6n}{(1-\gamma)^3} + 2\lambda \max_i |\mathcal{A}_i| \right) \|\theta' - \theta\|_2.
\end{aligned}$$

□

**Lemma 31.**

$$\|\nabla_{\theta_i} \Phi(\theta') - \nabla_{\theta_i} \Phi(\theta)\|_1 \leq \frac{6}{(1-\gamma)^3} \sum_{i=1}^n \|\theta'_i - \theta_i\|_2$$

*Proof.*

$$\begin{aligned}
\|\nabla_{\theta_i} \Phi(\theta') - \nabla_{\theta_i} \Phi(\theta)\|_1 &= \frac{1}{1-\gamma} \sum_{s, a_i} \left| d_{\theta'}(s) \pi_{\theta'_i}(a_i|s) \overline{A_i^{\theta'}(s, a_i)} - d_{\theta}(s) \pi_{\theta_i}(a_i|s) \overline{A_i^{\theta}(s, a_i)} \right| \\
&= \frac{1}{1-\gamma} \sum_{s, a_i} \left| d_{\theta'}(s) \pi_{\theta'_i}(a_i|s) \sum_{a_{-i}} \pi_{\theta'_{-i}}(a_{-i}|s) A_i^{\theta'}(s, a_i, a_{-i}) - d_{\theta}(s) \pi_{\theta_i}(a_i|s) \sum_{a_{-i}} \pi_{\theta_{-i}}(a_{-i}|s) A_i^{\theta}(s, a_i) \right| \\
&\leq \frac{1}{1-\gamma} \sum_{s, a} \left| d_{\theta'}(s) \pi_{\theta'}(a|s) A_i^{\theta'}(s, a) - d_{\theta}(s) \pi_{\theta}(a|s) A_i^{\theta}(s, a) \right| \\
&\leq \frac{1}{1-\gamma} \left( \sum_{s, a} |d_{\theta'}(s) \pi_{\theta'}(a|s) - d_{\theta}(s) \pi_{\theta}(a|s)| |A_i^{\theta'}(s, a)| + \sum_{s, a} d_{\theta}(s) \pi_{\theta}(a|s) |A_i^{\theta'}(s, a) - A_i^{\theta}(s, a)| \right) \\
&\leq \frac{1}{1-\gamma} \left( \frac{1}{1-\gamma} \sum_{s, a} |d_{\theta'}(s) \pi_{\theta'}(a|s) - d_{\theta}(s) \pi_{\theta}(a|s)| + \max_{s, a_i} |A_i^{\theta'}(s, a) - A_i^{\theta}(s, a)| \right)
\end{aligned}$$

From Lemma 32 and Corollary 35, we have that

$$\begin{aligned}\|\nabla_{\theta_i} \Phi(\theta') - \nabla_{\theta_i} \Phi(\theta)\|_1 &\leq \frac{3}{(1-\gamma)^3} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1 \\ &\leq \frac{3}{(1-\gamma)^3} \max_s \sum \|\pi_{\theta'_{i,s}} - \pi_{\theta_{i,s}}\|_1\end{aligned}$$

From Corollary 37 we have that

$$\begin{aligned}\|\nabla_{\theta_i} \Phi(\theta') - \nabla_{\theta_i} \Phi(\theta)\|_1 &\leq \frac{3}{(1-\gamma)^3} \max_s \sum \|\pi_{\theta'_{i,s}} - \pi_{\theta_{i,s}}\|_1 \\ &\leq \frac{6}{(1-\gamma)^3} \sum_{i=1}^n \|\theta_i - \theta'_i\|_2.\end{aligned}\quad \square$$

## I Some Useful Lemmas

**Lemma 32.**

$$\begin{aligned}|Q^{\theta'}(s, a) - Q^\theta(s, a)| &\leq \frac{1}{(1-\gamma)^2} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1 \\ |V^{\theta'}(s) - V^\theta(s)| &\leq \frac{1}{(1-\gamma)^2} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1,\end{aligned}$$

and thus

$$|A^{\theta'}(s, a) - A^\theta(s, a)| \leq \frac{2}{(1-\gamma)^2} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1$$

*Proof.* From performance difference lemma we have that

$$\begin{aligned}&|Q^{\theta'}(s, a) - Q^\theta(s, a)| \\ &= \left| \sum_{t=1}^{+\infty} \gamma^t \sum_{s'} \Pr^{\theta'}(s(t) = s' | s(0) = s, a(0) = a) \sum_{a'} (\pi_{\theta'}(a' | s') - \pi_\theta(a' | s')) Q^\theta(s', a') \right| \\ &\leq \left| \sum_{t=1}^{+\infty} \gamma^t \sum_{s'} \Pr^{\theta'}(s(t) = s' | s(0) = s, a(0) = a) \sum_{a'} |\pi_{\theta'}(a' | s') - \pi_\theta(a' | s')| |Q^\theta(s', a')| \right| \\ &\leq \left| \sum_{t=1}^{+\infty} \gamma^t \max_{s'} \sum_{a'} |\pi_{\theta'}(a' | s') - \pi_\theta(a' | s')| \frac{1}{1-\gamma} \right| \\ &= \frac{1}{(1-\gamma)^2} \max_s \sum_a |\pi_{\theta'}(a | s) - \pi_\theta(a | s)| \\ &= \frac{1}{(1-\gamma)^2} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1\end{aligned}$$

Same argument also holds for  $|V^{\theta'}(s) - V^\theta(s)|$ , and thus

$$|A^{\theta'}(s, a) - A^\theta(s, a)| \leq |Q^{\theta'}(s, a) - Q^\theta(s, a)| + |V^{\theta'}(s) - V^\theta(s)| \leq \frac{2}{(1-\gamma)^2} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1.\quad \square$$

**Lemma 33.**

$$\frac{1}{1-\gamma} \sum_{s,a} (d_{\theta'}(s) \pi_{\theta'}(a | s) - d_\theta(s) \pi_\theta(a | s)) r(s, a) \leq \frac{1}{(1-\gamma)^2} \|r\|_\infty \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1,$$

where  $\|r\|_\infty = \max_{s,a} |r(s, a)|$ .

*Proof.* For any reward function  $r(s, a)$ , we can define its value function  $V^\theta(s)$  and  $Q^\theta(s, a)$  correspondingly. Using performance difference lemma we have that

$$\begin{aligned}
& \frac{1}{1-\gamma} \sum_{s,a} (d_{\theta'}(s)\pi_{\theta'}(a|s) - d_\theta(s)\pi_\theta(a|s)) r(s, a) \\
&= \sum_s \rho(s)(V^{\theta'}(s) - V^\theta(s)) \\
&= \frac{1}{1-\gamma} \sum_s d_{\theta'}(s) \sum_a (\pi_{\theta'}(a|s) - \pi_\theta(a|s)) Q^\theta(s, a) \\
&\leq \frac{1}{(1-\gamma)^2} \|r\|_\infty \sum_s d_{\theta'}(s) \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1 \\
&\leq \frac{1}{(1-\gamma)^2} \|r\|_\infty \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1. \quad \square
\end{aligned}$$

We have the following two corollaries for Lemma 33.

**Corollary 34.**

$$\frac{1}{1-\gamma} \sum_s |d_{\theta'}(s) - d_\theta(s)| \leq \frac{1}{(1-\gamma)^2} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1$$

*Proof.*

$$\begin{aligned}
\frac{1}{1-\gamma} \sum_s |d_{\theta'}(s) - d_\theta(s)| &= \max_{-1 \leq r(s) \leq 1} \frac{1}{1-\gamma} \sum_{s,a} (d_{\theta'}(s)\pi_{\theta'}(a|s) - d_\theta(s)\pi_\theta(a|s)) r(s) \\
&\leq \frac{1}{(1-\gamma)^2} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1. \quad \square
\end{aligned}$$

**Corollary 35.**

$$\frac{1}{1-\gamma} \sum_s |d_{\theta'}(s)\pi_{\theta'}(a|s) - d_\theta(s)\pi_\theta(a|s)| \leq \frac{1}{(1-\gamma)^2} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1.$$

*Proof.*

$$\begin{aligned}
& \frac{1}{1-\gamma} \sum_s |d_{\theta'}(s)\pi_{\theta'}(a|s) - d_\theta(s)\pi_\theta(a|s)| \\
&= \max_{-1 \leq r(s,a) \leq 1} \frac{1}{1-\gamma} \sum_{s,a} (d_{\theta'}(s)\pi_{\theta'}(a|s) - d_\theta(s)\pi_\theta(a|s)) r(s, a) \\
&\leq \frac{1}{(1-\gamma)^2} \max_s \|\pi_{\theta'_s} - \pi_{\theta_s}\|_1. \quad \square
\end{aligned}$$

**Lemma 36.**

$$\sum_{a_i} (\pi_{\theta'_i}(a_i|s) - \pi_{\theta_i}(a_i|s)) f(a_i) \leq 2\|f\|_\infty \|\theta'_{i,s} - \theta_{i,s}\|_2$$

*Proof.* It suffices to show that

$$\left\| \nabla_{\theta_{i,s}} \sum_{a_i} \pi_{\theta_i}(a_i|s) f(a_i) \right\|_2 \leq 2\|f\|_\infty, \quad \forall \theta,$$

then by Lagrange mean value theorem,

$$\sum_{a_i} (\pi_{\theta'_i}(a_i|s) - \pi_{\theta_i}(a_i|s)) f(a_i)$$

$$\leq \max_{t, \bar{\theta} = t\theta + (1-t)\theta'} \left\| \nabla_{\theta_{i,s}} \sum_{a_i} \pi_{\bar{\theta}_i}(a_i|s) f(a_i) \right\|_2 \|\theta'_{i,s} - \theta_{i,s}\|_2 \leq 2\|f\|_\infty \|\theta'_{i,s} - \theta_{i,s}\|_2.$$

Since

$$\frac{\partial \sum_{a_i} \pi_{\theta_i}(a_i|s) f(a_i)}{\partial \theta_{a_i,s}} = \pi_{\theta_i}(a_i|s) (f(a_i) - \bar{f}), \text{ where } \bar{f} = \sum_{a_i} \pi_{\theta_i}(a_i|s) f(a_i),$$

we have

$$\left\| \nabla_{\theta_{i,s}} \sum_{a_i} \pi_{\theta_i}(a_i|s) f(a_i) \right\|_2^2 = \sum_{a_i} \pi_{\theta_i}(a_i|s)^2 (f(a_i) - \bar{f})^2 \leq \sum_{a_i} \pi_{\theta_i}(a_i|s)^2 (2\|f\|_\infty)^2 \leq 4\|f\|_\infty^2,$$

which completes the proof.  $\square$

**Corollary 37.** (of Lemma 36)

$$\|\pi_{\theta'_{i,s}} - \pi_{\theta_{i,s}}\|_1 \leq 2\|\theta'_{i,s} - \theta_{i,s}\|_2 \leq 2\|\theta'_i - \theta_i\|_2$$

*Proof.*

$$\|\pi_{\theta'_{i,s}} - \pi_{\theta_{i,s}}\|_1 = \max_{f: \|f\|_\infty \leq 1} \sum_{a_i} (\pi_{\theta'_i}(a_i|s) - \pi_{\theta_i}(a_i|s)) f(a_i) \leq 2\|\theta'_{i,s} - \theta_{i,s}\|_2. \quad \square$$