
Symmetric (Optimistic) Natural Policy Gradient for Multi-agent Learning with Parameter Convergence

Sarath Pattathil
MIT

Kaiqing Zhang
University of Maryland, College Park

Asuman Ozdaglar
MIT

Abstract

Multi-agent interactions are increasingly important in the context of reinforcement learning, and the theoretical foundations of policy gradient methods have attracted surging research interest. We investigate the global convergence of natural policy gradient (NPG) algorithms in multi-agent learning. We first show that vanilla NPG may not have *parameter convergence*, i.e., the convergence of the vector that parameterizes the policy, even when the payoffs are regularized (which enabled strong convergence guarantees in the *policy space* in the literature). This non-convergence of parameters leads to stability issues in learning, which becomes especially relevant in the function approximation setting, where we can only operate on low-dimensional parameters, instead of the high-dimensional policy. We then propose variants of the NPG algorithm, for several standard multi-agent learning scenarios: two-player zero-sum matrix and Markov games, and multi-player monotone games, with global last-iterate parameter convergence guarantees. We also generalize the results to certain function approximation settings. Note that in our algorithms, the agents take *symmetric* roles. Our results might also be of independent interest for solving nonconvex-nonconcave minimax optimization problems with certain structures. Simulations are also provided to corroborate our theoretical findings.

1 Introduction

Policy gradient (PG) methods have served as the workhorse of modern reinforcement learning (RL) (Schulman et al.,

2015, 2017; Haarnoja et al., 2018), and enjoy the desired properties of being scalable to large state-action spaces, stability with function approximation, as well as sample efficiency. In fact, policy gradient methods have achieved impressive empirical performance in multi-agent RL (Lowe et al., 2017; Yu et al., 2021), the regime where many RL’s recent successes are pertinent to (Silver et al., 2017; OpenAI, 2018; Shalev-Shwartz et al., 2016).

Despite the tremendous empirical successes, theoretical foundations of PG methods, even for the single-agent setting, have not been uncovered until recently (Fazel et al., 2018; Agarwal et al., 2019; Zhang et al., 2019a; Wang et al., 2019; Mei et al., 2020; Cen et al., 2021a). The theoretical understanding of PG methods for multi-agent RL remains largely elusive, except for several recent attempts (Daskalakis et al., 2020; Zhao et al., 2021; Wei et al., 2021; Cen et al., 2021b). The key challenge is that in the policy parameter space, even for the basic two-player zero-sum matrix game, the problem becomes *nonconvex-nonconcave* and is computationally intractable in general (Daskalakis et al., 2021).

In this paper, we aim to fill in the gap by studying the global convergence of natural PG (NPG) (Kakade, 2002), which forms the basis for many popular PG algorithms (e.g., Proximal Policy Optimization (PPO)/Trust Region Policy Optimization (TRPO)), in the parameter space and for multi-agent learning. We are interested in the setting where the agents take *symmetric* roles and operate *independently*, as it does not require a central coordinator and it scales favorably with the number of agents. Analysis of this setting is challenging precisely because the concurrent updates of the agents makes the learning environment *non-stationary* from one agent’s perspective. With *asymmetric* update rules among agents, the non-stationarity issue can be mitigated, and the global convergence of PG methods has been established lately in (Zhang et al., 2019b; Daskalakis et al., 2020; Zhao et al., 2021). However, though being valid as an optimization scheme, asymmetric update-rules might be hard to justify in game-theoretic multi-agent learning with symmetric players. It is thus desirable to develop provably convergent PG methods with symmetric update rules.

1.1 Related work

Policy gradient RL methods for games. Gradient-descent-ascent (GDA) with projection on simplexes can be viewed as symmetric policy gradient methods for solving matrix games with *direct policy parameterization* (Agarwal et al., 2019), which enjoys an average-iterate convergence (Cesa-Bianchi and Lugosi, 2006). Such a guarantee is shared with multiplicative weight update (MWU), which is especially suitable for repeated matrix games, and equivalent to natural PG method with tabular softmax policy parameterization (Agarwal et al., 2019). To the best of our knowledge, however, neither parameter convergence nor function approximation has been studied in this context. For Markov games, Zhang et al. (2019b); Bu et al. (2019); Hambly et al. (2021) have studied global convergence of PG methods for those with a linear quadratic structure; for zero-sum Markov games, Daskalakis et al. (2020) established global convergence of independent PG with two-timescale stepsizes for the tabular setting; Zhao et al. (2021) studied a double-loop natural PG algorithm with function approximation; more recently, Alacaoglu et al. (2022) proposed a framework of natural actor-critic algorithms. No last-iterate convergence to the Nash equilibrium was established in these works, and these update-rules were all *asymmetric*. Qiu et al. (2021) developed symmetric policy optimization methods for certain zero-sum Markov games with structured transitions. Concurrently, Zhang et al. (2022) proposed a policy optimization framework with fast average-iterate convergence guarantees for finite-horizon Markov games. Finally, Leonardos et al. (2021); Zhang et al. (2021); Ding et al. (2022) have studied global convergence of symmetric PG methods in Markov potential games recently, not focused on last-iterate or parameter convergence.

Last-iterate convergence in constrained multi-agent learning. Several papers including Tseng (1995); Azizian et al. (2020); Ibrahim et al. (2020); Fallah et al. (2020) and references therein studied the *last-iterate* behavior of strongly monotone games. Furthermore, Golowich et al. (2020a,b) extended this analysis to the monotone game setting. However, these papers did not consider *parametrized policies*. More specifically, in the matrix game setting with a simplex constraint, papers including Daskalakis and Panageas (2019); Wei et al. (2020); Cen et al. (2021b) showed the last-iterate policy convergence of optimistic methods. However, these papers did not consider policy parameterization or the function approximation settings, and some papers required the assumption that the NE is unique (Daskalakis and Panageas, 2019; Wei et al., 2020). For Markov games, Zhao et al. (2021) established last-iterate convergence, but not to NE due to asymmetric update; Wei et al. (2021); Cen et al. (2021b) were, to the best of our knowledge, the only last-iterate policy convergence results in Markov games with symmetric updates. However, these

works did not study the function approximation setting, or monotone games beyond the two-player zero-sum case. Also, though having greatly inspired our work, the regularization idea in Cen et al. (2021b) alone cannot prevent the non-convergence issue of the policy parameters from happening (see §2). Our goal, in contrast, is to study the (last-iterate) convergence of the actual policy parameters, and for more general multi-agent learning settings beyond the tabular zero-sum one.

Nonconvex-nonconcave minimax optimization. It is shown that for general nonconvex-nonconcave minimax optimization, even *local* solution concepts (Jin et al., 2020) may not exist, and finding them can be intractable (Daskalakis et al., 2021). Thus, specific structural properties have to be exploited to design efficient algorithms with *global* convergence. Lin et al. (2019); Thekumparampil et al. (2019); Nouiehed et al. (2019); Yang et al. (2020) have studied the nonconvex-(strongly)-concave or the nonconvex-Polyak-Lojasiewicz (PL) or PL-PL settings, with global convergence rate guarantees. The algorithms in these papers are all *asymmetric* in that they run the inner loop (which solves the maximization problem) multiple times (or on a faster timescale with larger stepsizes) to reach an approximate solution of the inner optimization problem, and then run one step of descent on the outer problem (or on a slower timescale with smaller stepsizes). Closely related to one motivation of our work, Vlatakis-Gkaragkounis et al. (2019); Flokas et al. (2021); Mladenovic et al. (2021) studied nonconvex-nonconcave minimax problem with *hidden convexity* structures, and show that GDA can fail to converge globally even so. Interestingly, the benefit of regularization (more generally, strict convexity), and natural gradient flow under Fisher information geometry, were also examined in Flokas et al. (2021); Mladenovic et al. (2021) to establish some positive convergence results. Different from our work, the dynamics there are in *continuous-time*, and the parameterization in Flokas et al. (2021) is decoupled by dimension, and the convergence rate in the perturbed game is not global. These conditions prevent the application of their results and proof techniques to our setting directly. Also, last-iterate finite-time rates of the iterates were not established in Mladenovic et al. (2021).

We focus on the *last-iterate* convergence of the *policy parameters*, which is critical to establish in order to avoid stability issues during learning. For example, if the norm of the parameters blow up, we might end up with precision issues in computing the gradient and updating the parameters. This is particularly relevant to the setting with function approximation, where we can only operate on *low-dimensional* parameters of the policy, instead of the high-dimensional policy per se. Indeed, we aim to explore the convergence property in the function approximation setting to handle large state-action spaces. Finally, our results are

also motivated from the study of *nonconvex-nonconcave minimax* optimization problems, especially those with certain structures that yield global convergence of gradient-based methods. We aim to explore such structures in multi-agent learning with *parameterized* policies. We summarize our contributions as follows.

Contributions. Our contributions are three-fold. First, we identify the non-convergence issue in the policy parameter space of natural PG methods for RL. We show that this issue persists even with entropy regularized rewards. Second, we develop symmetric variants of the natural PG method, i.e., both without and with the optimistic updates (Rakhlin and Sridharan (2013)) and establish the last-iterate global convergence to the Nash equilibrium in the policy parameter space. Third, we generalize the scope of symmetric PG methods in game-theoretic multi-agent learning, including two-player zero-sum matrix and Markov games (MGs), multi-player monotone games, and the corresponding linear function approximation settings under certain assumptions, in order to handle enormously large state-action spaces, all with last-iterate parameter convergence rate guarantees. We have also provided numerical experiments to validate the effectiveness of our algorithms.

Notation. For vector $v \in \mathbb{R}^d$, we use $[v]_a$ with $a \in \{1, 2, \dots, d\}$ to denote the a -th element of v . We use $\|v\|$ to denote the ℓ_2 -Euclidean norm of a vector v and $\|Q\|$ to denote the ℓ_2 -induced norm of matrix Q . We also use $\|Q\|_\infty$ to denote the infinity norm and $\|Q\|_F$ to denote the Frobenius norm of matrix Q . For a finite-set \mathcal{S} , we use $\Delta(\mathcal{S})$ to denote the simplex over \mathcal{S} . We use $\mathbb{1}$ to denote the matrix of all ones of appropriate dimension. For any positive integer n , we use $[n]$ to denote the set $\{1, \dots, n\}$. We use the subscript $-i$ to denote the quantities of all players other than player i . $\text{KL}(p\|q)$ denotes the KL divergence between two probability distributions p and q . For a matrix C , we use $C = [A \mid B]$ to denote the concatenation of the component matrices A and B . For two vectors $x, y \in \mathbb{R}^d$, $x \cdot y$ denotes their inner-product, i.e., $x^\top y$. We use I_d to denote an identity matrix of dimension d .

2 Motivation & Background

In this section, we introduce the background of the natural PG methods we study, with two-player¹ zero-sum matrix games being a motivating example.

Zero-sum matrix games. Two-player zero-sum matrix games are characterized by a tuple $(\mathcal{A}, \mathcal{B}, Q)$, where $Q \in \mathbb{R}^{n \times n}$ denotes the cost² matrix, \mathcal{A} and \mathcal{B} denote the action

spaces of players 1 and 2, respectively. For notational simplicity, we assume both action spaces have cardinality n , i.e., $|\mathcal{A}| = |\mathcal{B}| = n$. Note that our results can be readily generalized to the setting with different action-space cardinalities. For convenience, we use *indices* of the actions to denote the actions, i.e., $\mathcal{A} = \mathcal{B} = \{1, 2, \dots, n\}$, without loss of generality. Note that the *actual actions* of both players for the same index need not to be the same, and the cost matrix Q needs not to be symmetric. The problem can thus be formulated as a minimax (i.e., saddle-point optimization) problem

$$\min_{g \in \Delta(\mathcal{A})} \max_{h \in \Delta(\mathcal{B})} f(g, h) := g^\top Q h, \quad (2.1)$$

where g and h are referred to as the policies/strategies of the players. By Minimax Theorem (Neumann, 1928), the min and max operators in (2.1) can be interchanged, and the solution concept of *Nash equilibrium* (NE), which is defined as a pair of policies (g^*, h^*) such that

$$f(g, h^*) \geq f(g^*, h^*) \geq f(g^*, h), \quad \text{for any } (g, h) \in \Delta(\mathcal{A}) \times \Delta(\mathcal{B})$$

always holds. In particular, at the Nash equilibrium, the players execute the best-response policies of each other, and have no incentive to deviate from it.

Policy parameterization. To develop policy gradient methods for multi-player learning, the policies $(g, h) \in \Delta(\mathcal{A}) \times \Delta(\mathcal{B})$ are parameterized by some parameters θ and ν . Specifically, consider the following softmax parameterization that is common in practice: for any $a \in \mathcal{A}$ and $b \in \mathcal{B}$

$$g_\theta(a) = \frac{e^{p_\theta(a)}}{\sum_{a' \in \mathcal{A}} e^{p_\theta(a')}}, \quad h_\nu(b) = \frac{e^{q_\nu(b)}}{\sum_{b' \in \mathcal{B}} e^{q_\nu(b')}}, \quad (2.2)$$

where $\theta, \nu \in \mathbb{R}^d$ for some integer $d > 0$, $p_\theta, q_\nu : \mathbb{R}^d \rightarrow \mathbb{R}^n$ are two differentiable functions. Note that $g_\theta(a), h_\nu(b) > 0$ for any bounded p_θ, q_ν , and $\sum_{a \in \mathcal{A}} g_\theta(a) = \sum_{b \in \mathcal{B}} h_\nu(b) = 1$. This parameterization gives the following minimax problem for the zero-sum matrix game:

$$\min_{\theta \in \mathbb{R}^n} \max_{\nu \in \mathbb{R}^n} f(\theta, \nu) := g_\theta^\top Q h_\nu, \quad (2.3)$$

where by a slight abuse of notation, we use $f(\theta, \nu)$ to denote $f(g_\theta, h_\nu)$. In this section and §3, we consider the tabular softmax parameterization where $p_\theta = \theta \in \mathbb{R}^n$ and $q_\nu = \nu \in \mathbb{R}^n$. In §C, we consider the setting with function approximation where $d < n$.

The benefits of softmax parameterization are that: 1) it transforms a *constrained* problem over simplexes to an *unconstrained* one, making it easier to implement; 2) it readily incorporates function approximation to deal with large spaces (see §C). On the other hand, this policy parameterization makes the optimization problem (2.3) more challenging to solve. Indeed, the minimax problem (2.3) becomes a nonconvex-nonconcave problem in θ and ν , even with the tabular parameterization as we will show later in Lemma 2.2.

¹Hereafter, we use *player* and *agent* interchangeably.

²Note that we can also model it as a payoff, with a negative sign.

Remark 2.1 (Hidden bilinear problem). Note that Problem (2.3), which resembles a bilinear zero-sum game, in fact falls into the class of *hidden bilinear* minimax problems discussed in Vlatakis-Gkaragkounis et al. (2019) (or more generally the *hidden convex-concave* games studied in Flokas et al. (2021); Mladenovic et al. (2021)). It was shown in Flokas et al. (2021) that for general smooth functions of g_θ and h_ν , vanilla gradient descent-ascent exhibits a variety of behaviors antithetical to convergence to the solutions. We here instead, show that for the specific *softmax parameterization*, and for certain variants of the vanilla gradient-descent-ascent method, the last-iterate convergence rate of the parameters θ and ν can be established.

Natural PG & Non-convergence pitfall. Before proceeding further, we first introduce the *regularized* game:

$$\min_{\theta \in \mathbb{R}^n} \max_{\nu \in \mathbb{R}^n} f_\tau(\theta, \nu) := g_\theta^\top Q h_\nu - \tau \mathcal{H}(g_\theta) + \tau \mathcal{H}(h_\nu), \quad (2.4)$$

where the cost of both players is regularized by the Shannon entropy of the policies, with $\tau > 0$ being the regularization parameter, and $\mathcal{H}(\pi) = -\sum_{a \in \mathcal{A}} \pi(a) \log(\pi(a))$ for π on a simplex. The entropy regularization, which is commonly used in single-player RL, enjoys the benefits of both encouraging exploration and accelerating convergence (Neu et al., 2017; Mei et al., 2020). Our hope is also to exploit the benefits of entropy regularization in the multi-player setting. Indeed, the regularized cost traces its source in the game theory literature (McKelvey and Palfrey, 1995), to model the imperfect knowledge of the cost matrix Q . In the next lemma, we show that the problem in Equation 2.4 can be of the nonconvex-nonconcave type:

Lemma 2.2. The minimax problem (2.4) is nonconvex in θ and nonconcave in ν , even if $p_\theta = \theta$ and $q_\nu = \nu$.

Note the nonconvexity in the parameters remains even when we regularize with the entropy of the policy, i.e., when $\tau > 0$.

Motivated by the successes of *natural policy gradient* (Kakade, 2002) and its variants, as PPO/TRPO (Schulman et al., 2015, 2017), in RL practice, we consider the natural PG descent-ascent update for (2.4), which is given by

$$\begin{aligned} \theta_{t+1} &= \theta_t - \eta \cdot F_\theta^\dagger(\theta_t) \cdot \frac{\partial f_\tau(\theta_t, \nu_t)}{\partial \theta} \\ &= (1 - \eta\tau)\theta_t - \eta Q h_{\nu_t} + \eta\tau \left(\log \sum_{a' \in \mathcal{A}} e^{\theta_t(a')} - 1 \right), \end{aligned} \quad (2.5)$$

$$\begin{aligned} \nu_{t+1} &= \nu_t + \eta \cdot F_\nu^\dagger(\nu_t) \cdot \frac{\partial f_\tau(\theta_t, \nu_t)}{\partial \nu} \\ &= (1 - \eta\tau)\nu_t + \eta Q^\top g_{\theta_t} - \eta\tau \left(\log \sum_{b' \in \mathcal{B}} e^{\nu_t(b')} - 1 \right), \end{aligned} \quad (2.6)$$

where $F_\theta(\theta) = \mathbb{E}_{a \sim g_\theta} [(\nabla_\theta \log g_\theta(a))(\nabla_\theta \log g_\theta(a))^\top]$ and $F_\nu(\nu) = \mathbb{E}_{b \sim h_\nu} [(\nabla_\nu \log h_\nu(b))(\nabla_\nu \log h_\nu(b))^\top]$ are the Fisher information matrices, M^\dagger denotes the pseudo-inverse of the matrix M , and $\eta > 0$ is the stepsize. The derivations for natural policy gradient can be found in §A.2 for completeness.

Unfortunately, the vanilla NPG update (2.5)-(2.6) may fail to converge in the parameter space for *any* stepsize $\eta > 0$. The key reason for the failure is that the mappings represented in (2.5)-(2.6) may not have a fixed point for a general Q and τ (which could be the only limit point for this dynamics). In fact, this issue persists even when the regularization parameter $\tau = 0$. We formalize this pitfall in the following lemma, with its proof deferred to appendix.

Lemma 2.3 (Pitfall of vanilla NPG). There exists a game (2.4) with $\tau \geq 0$ and $|\mathcal{B}| = 1$, such that the updates (2.5)-(2.6) do not converge for any $\eta > 0$.

Remarkably, we emphasize that our Lemma 2.3, by construction, also even applies to the single-agent setting, with a regularized cost and the NPG update, as studied recently in Cen et al. (2021a); Lan (2021); Zhan et al. (2021). These works only focus on the convergence in the policy space, which does not imply the desired convergence in the policy parameter space. The later becomes especially relevant in the function approximation setting, as we will study later. Finally, we remark that, the non-convergence here also should not be confused with the *last-iterate non-convergence* of no-regret learning algorithms for solving bilinear zero-sum games (Daskalakis and Panageas, 2018; Bailey and Piliouras, 2018), as our example is essentially a single-agent case. We summarize the importance and motivation of establishing parameter convergence as follows.

Importance of Parameter Convergence:

Numerical instability: Prior works were only able to show the convergence of *values* and/or *policies*, and the convergence behavior of policy parameters was *unclear* (or overlooked). Arguably, having (last-iterate) parameter convergence is the *strongest* type of convergence among the three. In practice, having parameters blow-up to infinity can cause numerical issues. For example, once the size of the parameter crosses a threshold (say 2^{64} for an integer in a 64-bit operating system), there would be overflow issues, and the stored parameter would be void, and NaN (not a number) would be returned by the program. This blow-up would then cause trouble in recovering the policy, or approximating the policy with arbitrary accuracy. In order to circumvent this issue, a common practice in Neural Network training is to do *Clipping/Projection*. In fact, ensuring the *stability* of the model is very important in deep learning. Specifically, consider $\max_{\theta \in \mathbb{R}^n} q^\top g_\theta + \tau \mathcal{H}(g_\theta)$, where we know that under the NPG updates, $g_\theta \rightarrow g^*$ while θ could blow up to infinity. One could clip the

parameter θ to some large constant θ_{\max} , i.e., solving $\max_{\|\theta\|_{\infty} \leq \theta_{\max}} q^{\top} g_{\theta} + \tau \mathcal{H}(g_{\theta})$ instead. For concreteness, let $n = 2$, $\tau = 1$, $\theta_{\max} = 80$ and $q = [-2, -3]$. The optimal solution is then given by $g_i^* \propto \exp(q_i)$ for $i = 1, 2$. On running the vanilla NPG algorithm, since we do weight clipping, the algorithm converges to $\theta = [\theta_{\max}, \theta_{\max}]$ corresponding to the distribution $[1/2, 1/2] \neq g^*$. Meanwhile, the modified NPG we propose converges to $\theta = [-2, -3]$ (see Theorem 3.4) which exactly corresponds to the optimal solution g^* . Hence, in practice where the norm of the (neural network) parameters is bounded, one might not obtain policy convergence using vanilla NPG as desired, while our proposed algorithm works.

Nonconvex-nonconcave minimax optimization: The second reason comes from a minimax optimization perspective of solving (2.3). We view optimization over the parameter space as an interesting *nonconvex-nonconcave* minimax optimization problem with a *hidden structure* (See Lemma 2.2). To the best of our knowledge, this is the first paper to provide a *symmetric* discrete-time algorithm to solve certain nonconvex-nonconcave problem (and more generally non-monotone variational inequalities) with last-iterate convergence rates, even including the specialized settings (like the ones with Polyak-Łojasiewicz condition (Nouiehed et al., 2019; Yang et al., 2020)).

Function approximation (FA): Parameter convergence becomes crucial in FA settings (used in practice). Here, the policy lies in a high-dimensional space (or even an infinite-dimensional space if the actions are *continuous*), which we simply do not have access to and/or cannot operate on. The way practitioners run PG methods is to just operate on the low-dimensional policy parameter space. Thus, parameter convergence is necessary to design meaningful stopping criteria for optimization algorithms. If parameters explode to infinity, we cannot decide on how close we are to convergence, and the numerical issue mentioned before would cause trouble in recovering the policy.

3 Warm-up: (Optimistic) NPG for Matrix Games

To address the pitfall above about parameter convergence, we introduce two variants of the vanilla NPG (2.5)-(2.6), and show their convergence for solving matrix games.

3.1 NPG for Matrix Games

We first introduce the following variant of the vanilla NPG update:

$$\theta_{t+1} = (1 - \eta\tau)\theta_t - \eta Q h_{\nu_t}, \quad (3.1)$$

$$\nu_{t+1} = (1 - \eta\tau)\nu_t + \eta Q^{\top} g_{\theta_t}, \quad (3.2)$$

where we removed the last term in (2.5)-(2.6), respectively. Note that these updates correspond to the popular Multiplicative Weights Update (MWU) for the regularized game in policy space (we succinctly represent g_{θ_t} and h_{ν_t} as g_t and h_t , respectively), i.e.,

$$\begin{aligned} g_{t+1}(a) &\propto g_t(a)^{(1-\eta\tau)} e^{-\eta[Qh_t]_a}, \\ h_{t+1}(b) &\propto h_t(b)^{(1-\eta\tau)} e^{\eta[Q^{\top}g_t]_b}. \end{aligned} \quad (3.3)$$

First, we provide a convergence result for the updates in Equations (3.1)-(3.2), the non-optimistic version, both in terms of policy as well as parameters. In order to do so, we need to first show that the iterates of the regularized MWU in Equations (3.1)-(3.2) ensure that the policies stay bounded away from the boundary of the simplex. We show this in the following lemma:

Lemma 3.1. The policies corresponding to the iterates of regularized MWU in Equations (3.1)-(3.2) with stepsize $\eta < 1/\tau$ stay within a set $\Delta' \subset \Delta$ which is bounded away from the boundary of the simplex, i.e., $\forall x \in \Delta'$ with $x = (x_1, \dots, x_n)^{\top}$, and for all $i \in [n]$, $x_i \geq \delta > 0$ for some δ .

Since the iterates of the policies lie within Δ' , a closed and bounded set, and the regularized cost is continuously differentiable with respect to the policies, we let L denote the smoothness constant of the regularized cost in the policy space, i.e. $\forall z_1, z_2 \in \mathcal{Z}'$,

$$\|(Mz_1 + \tau \nabla \mathcal{H}(z_1)) - (Mz_2 + \tau \nabla \mathcal{H}(z_2))\| \leq L \|z_1 - z_2\|, \quad (3.4)$$

where $z = [g; h]$, $\mathcal{Z}' \in \Delta' \times \Delta'$ and with a slight abuse of notation, we define $\nabla \mathcal{H}(z) = [\nabla_g \mathcal{H}(g); \nabla_h \mathcal{H}(h)]$, and also, we define the matrix $M = \begin{pmatrix} 0 & Q \\ -Q^{\top} & 0 \end{pmatrix}$.

Finally, the entropy regularized optimization problem in the policy space can be formulated as:

$$\min_{g \in \Delta(\mathcal{A})} \max_{h \in \Delta(\mathcal{B})} = g^{\top} Q h - \tau \mathcal{H}(g) + \tau \mathcal{H}(h). \quad (3.5)$$

Now, we use this to derive the policy and parameter convergence of the MWU updates in Equations (3.1)-(3.2) in the following theorem.

Theorem 3.2. The solution (g^*, h^*) to the problem (3.5) is unique. Furthermore, let $\theta^* = \frac{-Qh^*}{\tau}$ and $\nu^* = \frac{Q^{\top}g^*}{\tau}$. Then on running Equations (3.1)-(3.2) with the stepsize satisfying $0 < \eta \leq \tau/L^2$, we have:

$$\text{KL}(z^* \| z_{t+1}) \leq \left(1 - \frac{\eta\tau}{2}\right) \text{KL}(z^* \| z_t), \quad (3.6)$$

and

$$\begin{aligned} &\|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 \\ &\leq (1 - \eta\tau/4)^t \left(\|\theta_0 - \theta^*\|^2 + \|\nu_0 - \nu^*\|^2 + \frac{4C}{\eta\tau} \right), \end{aligned} \quad (3.7)$$

where $C = \left(1 + \frac{1}{\eta\tau}(1 - \eta\tau)^2\right) 4\eta^2 \|Q\|_\infty^2 \text{KL}(z^* \| z_0)$. Here $z_0 = (g_{\theta_0}, h_{\nu_0})$ and $z^* = (g^*, h^*)$.

To the best of our knowledge, this is the first policy convergence guarantee for NPG (without optimism) for regularized matrix games. As the theorem above shows, parameter convergence can be established as well. In other words, though the vanilla NPG descent-ascent diverges for Problem (2.4), the variant we propose (Equations (3.1)-(3.2)) *implicitly regularizes* the parameter iterates to converge to a particular solution, in last-iterate. Recall from Lemma 2.2 that Problem (2.4) is a nonconvex-nonconcave minimax problem and there have been no convergence guarantees, to the best of our knowledge, of any symmetric and simultaneous-update algorithms in general (see Yang et al. (2020); Rafique et al. (2018); Lin et al. (2019) and references therein as examples for some structured nonconvex-nonconcave problems). From Theorem 3.2, the rate we can hope to achieve with NPG is $\mathcal{O}(\kappa^2 \log(1/\epsilon))$ (number of steps to reach a point ϵ close to the solution), where κ is the condition number L/τ of the problem. In the next subsection, we see how adding optimism (Rakhlin and Sridharan (2013)) will improve this rate of convergence.

3.2 Optimistic NPG (ONPG) for Matrix Games

In this subsection, we study the variant of the NPG updates in Equations (3.1)-(3.2) along with optimism (Popov, 1980; Rakhlin and Sridharan, 2013; Daskalakis and Panageas, 2018, 2019; Mokhtari et al., 2020b). In particular, we introduce the intermediate iterates $(\bar{\theta}_t, \bar{\nu}_t)$, and the following optimistic variant of the NPG (here $(\bar{\theta}_0, \bar{\nu}_0)$ is initialized as (θ_0, ν_0)):

$$\begin{aligned} \bar{\theta}_{t+1} &= (1 - \eta\tau)\theta_t - \eta Q h_{\bar{\nu}_t}, & \theta_{t+1} &= (1 - \eta\tau)\theta_t - \eta Q h_{\bar{\nu}_{t+1}} \\ \bar{\nu}_{t+1} &= (1 - \eta\tau)\nu_t + \eta Q^\top g_{\bar{\theta}_t}, & \nu_{t+1} &= (1 - \eta\tau)\nu_t + \eta Q^\top g_{\bar{\theta}_{t+1}}. \end{aligned}$$

This optimistic update is motivated by the success of optimistic gradient methods in saddle-point problems recently analyzed in several papers including Hsieh et al. (2019); Mokhtari et al. (2020a,b). Note that our algorithm is symmetric in θ and ν updates and both players update simultaneously. The update rules are also tabulated in Algorithm 1.

Remark 3.3 (Connections to the literature). Note that the natural PG update rule in Equations (3.1)-(3.2) has a close relationship to the multiplicative weight update rule (Freund and Schapire, 1997; Arora et al., 2012) in the policy space (g_θ, h_ν) , see Section C.3 in Agarwal et al. (2019) for a detailed discussion. Similarly, the optimistic NPG update in Algorithm 1 also relates to the optimistic MWU update (Daskalakis and Panageas, 2019). In fact, recent works Wei et al. (2020); Cen et al. (2021b) have shown the last-iterate *policy convergence* of OMWU for zero-sum matrix games (with Wei et al. (2020) relying on the uniqueness assumption of the NE). Our goal, in contrast, is to study the (last-

iterate) convergence behavior of the actual *policy parameters* (θ, ν) , and go beyond the tabular zero-sum setting.

Inspired by the results which show that adding optimistic updates improves convergence rates (Rakhlin and Sridharan, 2013), we next explore our modified NPG updates with optimism, and show that the convergence rate does in fact improve (in line with the comparison of the performance of GDA and optimistic GDA, see e.g., Fallah et al. (2020)).

In the next theorem, we show that our Optimistic NPG Algorithm (Algorithm 1) in fact converges linearly in last-iterate to a unique point in the set of NE in the parameter space, at a faster rate³ than the non-optimistic counterpart in Equations (3.1)-(3.2). The results are formally stated below.

Theorem 3.4. Let $\theta^* = \frac{-Qh^*}{\tau}$ and $\nu^* = \frac{Q^\top g^*}{\tau}$, where g^* and h^* are the solutions to the regularized game (2.4). Then on running Algorithm 1 with the stepsize satisfying $0 < \eta \leq \min\left\{\frac{1}{2\tau+2\|Q\|_\infty}, \frac{1}{4\|Q\|_\infty}\right\}$, we have:

$$\|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 \leq (1 - \eta\tau/2)^t V_0, \quad (3.8)$$

where $V_0 = \|\theta_0 - \theta^*\|^2 + \|\nu_0 - \nu^*\|^2 + \frac{2C}{\eta\tau}$ and $C = \left(1 + \frac{1}{\eta\tau}(1 - \eta\tau)^2\right) 4\eta^2 \|Q\|_\infty^2 \text{KL}(z^* \| z_0)$. Here $z_0 = (g_{\theta_0}, h_{\nu_0})$ and $z^* = (g^*, h^*)$.

This result shows that the specific hidden bilinear minimax problem we are dealing with does not fall into the spurious categories discussed in Vlatakis-Gkaragkounis et al. (2019), if we resort to the (optimistic) natural PG update. Note that achieving parameter convergence is a non-trivial task since we are dealing with a nonconvex-nonconcave minimax problem (see Lemma 2.2) The proof relies on the specific structure of the softmax policy parametrization and the construction of a novel Lyapunov function (see §B for more details). Next, we show how the optimistic NPG algorithm solves the original matrix game without regularization.

Corollary 3.5. If we run the non-optimistic variant of NPG in Equations (3.1)-(3.2) or the optimistic version in Algorithm 1 for time $T = \mathcal{O}\left(\frac{\log n}{\eta\epsilon} \log\left(\frac{1}{\epsilon}\right)\right)$ and set $\tau = \epsilon/(8 \log n)$, we have that the output (θ_T, ν_T) is an ϵ -NE of the original unregularized Problem (2.3).

We extend these results to simple function approximation settings in the Section C.

4 Multi-player Monotone Games

Monotone games. Consider a multi-player continuous game over simplexes, which strictly generalizes the zero-

³Note that we say that the optimistic version achieves a faster rate, since the range of stepsizes which permits convergence is much larger for the optimistic variant. This can be noted from the fact that L in equation (3.4) will be larger than $\|Q\|_\infty + \tau$.

sum matrix game in §2. The game is characterized by $(\mathcal{N}, \{\mathcal{A}_i\}_{i \in [N]}, \{f_i\}_{i \in [N]})$, where $\mathcal{N} = [N]$ is the set of players. Without loss of generality, we assume $|\mathcal{A}_i| = n$ for all $i \in [N]$. For notational convenience, let Δ denote the simplex over \mathcal{A}_i , and $z := (g_1, g_2, \dots, g_N) \in \Delta^N$ denote the strategy profile of all N players, with each $g_i \in \Delta$. We define the *pseudo-gradient* operator $F : \Delta^N \rightarrow \mathbb{R}^{nN}$ as $F(z) := [\nabla_{g_i} f_i(g_i, g_{-i})]_{i=1}^N$. To make the N -player game tractable, we make the following standard assumptions on F (Rosen, 1965; Nemirovski, 2004; Facchinei and Pang, 2007).

Assumption 4.1 (Monotonicity & Smoothness). The operator F is monotone and smooth, i.e., $\forall z, z' \in \Delta^N$

$$\langle F(z) - F(z'), z - z' \rangle \geq 0, \quad \|F(z) - F(z')\| \leq L \cdot \|z - z'\|,$$

where $L > 0$ is the Lipschitz constant of the operator F .

The goal is to find the NE, given by strategy z^* such that $f_i(z_i^*, z_{-i}^*) \leq f_i(z_i, z_{-i}^*)$, $\forall z_i \in \Delta$, $i \in [N]$. Under Assumption 4.1, it is known that the NE exists (Rosen, 1965).

Policy parameterization & regularized game. To develop policy gradient methods, we parameterize each policy $g_i \in \Delta$ by g_{θ_i} in the softmax form as before, i.e., for any $a_i \in \mathcal{A}_i$, $g_{\theta_i}(a_i) = e^{p_{\theta_i}(a_i)} \cdot (\sum_{a'_i \in \mathcal{A}_i} e^{p_{\theta_i}(a'_i)})^{-1}$, where $\theta_i \in \mathbb{R}^d$, and we consider both the tabular case with $p_{\theta_i} = \theta_i$ and the linear function approximation case with $p_{\theta_i}(a_i) = \phi_{a_i}^\top \theta_i$. This parameterization leads to the following set of optimization problems:

$$\min_{\theta_i \in \mathbb{R}^n} f_i(g_{\theta_i}, g_{\theta_{-i}}), \quad \forall i \in [N], \quad (4.1)$$

whose solution $(\theta_1^*, \theta_2^*, \dots, \theta_N^*)$, if exists, corresponds to the Nash equilibrium under this parameterization. Note that (4.1) can also be viewed as a nonconvex game (Lemma 2.2 is a special case) with a “hidden” *monotone variational inequality* structure, which generalizes the class of hidden convex-concave problems discussed in Flokas et al. (2021); Mladenovic et al. (2021).

Motivated by §2, we also consider the regularized game in hope of stronger convergence guarantees for solving (4.1). Specifically, the players solve

$$\min_{\theta_i \in \mathbb{R}^n} f_i(g_{\theta_i}, g_{\theta_{-i}}) - \tau \mathcal{H}(g_{\theta_i}), \quad \forall i \in [N], \quad (4.2)$$

where $\tau > 0$ and \mathcal{H} is the entropy function. With a small enough τ , the solution to (4.2) approximates that to (4.1).

4.1 Softmax parameterization

We first consider the tabular softmax parameterization with $p_{\theta_i} = \theta_i \in \mathbb{R}^n$ for all $i \in [N]$. In this case, the Nash equilibrium $\theta^* = (\theta_1^*, \theta_2^*, \dots, \theta_N^*)$ of the regularized monotone game (4.2) satisfies the following property.

Lemma 4.2. The NE of the game (4.2) exists. A vector $\theta^* = (\theta_1^*, \theta_2^*, \dots, \theta_N^*)$ is a NE of (4.2) if and only if: $g_{\theta_i^*}(a) \propto \exp\left(\frac{-[\nabla_{g_{\theta_i}} f_i(g_{\theta_i^*}, g_{\theta_{-i}^*})]_a}{\tau}\right)$ and the vector $(g_{\theta_1^*}, g_{\theta_2^*}, \dots, g_{\theta_N^*})$ is unique. We denote $g_i^* := g_{\theta_i^*}$.

Note that although the NE policy $(g_{\theta_1^*}, g_{\theta_2^*}, \dots, g_{\theta_N^*})$ is unique, the NE parameter $(\theta_1^*, \theta_2^*, \dots, \theta_N^*)$ is not necessarily the case. Motivated by §2 and §3, we propose the following update-rule for solving (4.2): \forall players $i \in [N]$,

$$\begin{aligned} \bar{\theta}_i^{t+1} &= (1 - \eta\tau)\theta_i^t - \eta \nabla_{g_{\theta_i}} f_i(g_{\bar{\theta}_i^t}, g_{\bar{\theta}_{-i}^t}), \\ \theta_i^{t+1} &= (1 - \eta\tau)\theta_i^t - \eta \nabla_{g_{\theta_i}} f_i(g_{\bar{\theta}_i^{t+1}}, g_{\bar{\theta}_{-i}^{t+1}}). \end{aligned}$$

We refer to the update-rule as *optimistic NPG* (as summarized in Algorithm 3), as it corresponds to the optimistic version of the (specific instance of) natural PG direction for the regularized objective (4.2). We choose this specific instance of NPG due to the pitfall discussed in §2; and the optimistic update is meant to obtain fast last-iterate convergence. See §E.2 for a detailed derivation of the update rule.

As shown in §2, the problem (4.1) is nonconvex in the policy parameter space, and can be challenging in general. Our strategy is to show that our algorithm solves the regularized problem (4.2) fast, with last-iterate parameter convergence (see Theorem 4.3), which, with small enough τ , also solves the nonconvex game (4.1) (see Corollary 4.5).

Theorem 4.3. Let $z^* = (g_i^*)_{i=1}^N$ be the unique Nash equilibrium given in Lemma 4.2. Also, we denote $z_t = (g_{\theta_i^t})_{i=1}^N$. Then for Algorithm 3 with stepsize $0 < \eta < \frac{1}{2(N+4)L+2\tau}$, we have:

$$\begin{aligned} \max \{ \text{KL}(z^* \| z_t), \text{KL}(z^* \| \bar{z}_{t+1}) \} &\leq (1 - \eta\tau)^t 2\text{KL}(z^* \| z_0), \\ \|\theta_{t+1} - \theta^*\|^2 &\leq (1 - \eta\tau/2)^t V_0, \end{aligned} \quad (4.3)$$

where $\theta_i^* = \frac{-\nabla_{g_{\theta_i}} f_i(g_i^*, g_{-i}^*)}{\tau}$, $V_0 = \|\theta^t - \theta^*\|^2 + \frac{2NC}{\eta\tau}$, and $C = 4\eta^2 L^2 \left(1 + \frac{1}{\eta\tau}(1 - \eta\tau)^2\right) \text{KL}(z^* \| z_0)$.

The proof follows by first showing convergence in the policy space, in which we are dealing with a strongly convex problem under convex constraints. We then use this result, along with a novel Lyapunov function to demonstrate the convergence in the parameter space, in which it is a nonconvex problem. The proof technique might of independent interest, and might be generalized to showing convergence in other nonconvex games with a hidden monotonicity structure.

Remark 4.4. The proof for Theorem 4.3 follows by first showing the convergence rate of the Proximal Point method, and then observing that Optimistic methods approximate this method and could potentially achieve the same convergence rates (see Mokhtari et al. (2020b) for a unified analysis). We provide a convergence analysis for the Proximal Point and Extragradient methods in §E.5.

Now, we present the convergence of Algorithm 3 to an ϵ -NE of the un-regularized problem (4.1).

Corollary 4.5. If we run Algorithm 3 for time $T = \mathcal{O}\left(\frac{N \log n}{\eta \epsilon} \log\left(\frac{1}{\epsilon}\right)\right)$ and set the regularization parameter $\tau = \epsilon/(4N \log n)$, we have that $\theta^T = [\theta_1^T, \theta_2^T, \dots, \theta_N^T]$, the iterate at time T , is an ϵ -NE of Problem (4.1).

We extend the results to certain function approximation settings in §E.6 in the Appendix.

5 Optimistic NPG for Markov Games

We now generalize our results in previous sections to the sequential decision-making case of Markov games. We provide a brief description of Markov games in §F.1.

Policy parameterization. Following the matrix game setting, we also use the softmax parameterization of the policies. Specifically, for any $\theta, \nu \in \mathbb{R}^d$, $(s, a, b) \in \mathcal{S} \times \mathcal{A} \times \mathcal{B}$,

$$g_\theta(a|s) = \frac{e^{p_\theta(s,a)}}{\sum_{a' \in \mathcal{A}} e^{p_\theta(s,a')}} \quad h_\nu(b|s) = \frac{e^{q_\nu(s,b)}}{\sum_{b' \in \mathcal{B}} e^{q_\nu(s,b')}} \quad (5.1)$$

Note that for any $s \in \mathcal{S}$, $\sum_a g_\theta(a|s) = \sum_b h_\nu(b|s) = 1$. We will consider both 1) the tabular case with $p_\theta = \theta$ and $q_\nu = \nu$, where $\theta \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{S}|}$ and $\nu \in \mathbb{R}^{|\mathcal{B}| \times |\mathcal{S}|}$; and 2) the linear function approximation case with $p_\theta(s, a) = \theta \cdot \phi(s, a)$ and $q_\nu(s, a) = \nu \cdot \phi(s, a)$ (See §F.3 for more details)⁴.

The parameterization thus leads to the following definition of the solution concept.

Definition 5.1 (ϵ -in-class-NE for Markov games). The policy parameter pair $(\tilde{\theta}, \tilde{\nu})$ is an ϵ -in-class Nash equilibrium if it satisfies that for all $s \in \mathcal{S}$,

$$V^{\tilde{\theta}, \nu}(s) - \epsilon \leq V^{\tilde{\theta}, \tilde{\nu}}(s) \leq V^{\theta, \tilde{\nu}}(s) + \epsilon, \forall \theta, \nu \in \mathbb{R}^d, \quad (5.2)$$

where $V^{\theta, \nu}(s) = V^{g_\theta, h_\nu}(s)$ denotes the value of the parameterized policy pair (g_θ, h_ν) . Note that if we are in the tabular setting, we will have $d = n$, and the definition covers that of the standard ϵ -NE for Markov games. We also define the ϵ -in-class NE Q -value accordingly.

Given that matrix games considered in §2 is a special case of the Markov games with $|\mathcal{S}| = 1$ and $\gamma = 0$, Lemma 2.2 implies that finding the NE in Definition 5.1 is nonconvex-nonconcave in general, and can be challenging to solve.

⁴Once again, note that we use the same features for both players for notational convenience. Our results continue to hold when the two players have different features ϕ .

We show in the following lemma that, for the tabular and linear function approximation settings we consider, such a parameterized NE exists.

Lemma 5.2 (Existence of parameterized/in-class NE). Under policy parameterization (5.1) with tabular parameterization, the in-class NE defined in Definition 5.1 exists.

Motivated by §3, we consider the modified version of NPG with optimism to solve this problem.

Optimistic NPG. Following §2, we also consider the *regularized* Markov games (Geist et al., 2019; Zhang et al., 2020; Cen et al., 2021b), in hope of favorable convergence guarantees. Define the regularized value functions as

$$V_\tau^{\theta, \nu}(s) := \mathbb{E}_{a_t \sim g_\theta(\cdot | s_t), b_t \sim h_\nu(\cdot | s_t)} \left[\sum_{t=0}^{\infty} \gamma^t (r_t - \tau \log g_\theta(a_t | s_t) + \tau \log h_\nu(b_t | s_t)) \mid s_0 = s \right], \quad (5.3)$$

where $r_t = r(s_t, a_t, b_t)$, $\tau < 1$, and

$$Q_\tau^{\theta, \nu}(s, a, b) := r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [V_\tau^{\theta, \nu}(s')]. \quad (5.4)$$

We denote by V_τ^* and Q_τ^* , the NE value and Q -functions respectively, for the regularized Markov game (“regularized NE”), i.e., $V_\tau^* = \min_\theta \max_\nu V_\tau^{\theta, \nu}$ and Q_τ^* is the corresponding Q -function. Note that their existence follows along similar lines as Lemma 5.2. As a generalization of regularized matrix games in §2, the non-convergence pitfall of vanilla NPG also occurs. We also define⁵ $f_\tau(Q(s); g_\theta(\cdot | s), h_\nu(\cdot | s)) := -g_\theta(\cdot | s)^\top Q(s) h_\nu(\cdot | s) - \tau \mathcal{H}(g_\theta(\cdot | s)) + \tau \mathcal{H}(h_\nu(\cdot | s))$

5.1 Convergence guarantee

To stabilize the algorithm, we propose the update rule where the parameters (θ, ν) for all states are updated at a faster time scale, and the Q matrix is updated at a slower time scale. To be more precise, at every time t of the outer loop, we solve the matrix game

$$\min_{\theta \in \mathbb{R}^n} \max_{\nu \in \mathbb{R}^n} f_\tau(Q(s); g_\theta(\cdot | s), h_\nu(\cdot | s)), \quad (5.5)$$

for each state $s \in \mathcal{S}$ by running T_{inner} iterations of the Optimistic NPG algorithm (Algorithm 1). At the end of each inner loop, the outer loop updates the Q matrix for each state $s \in \mathcal{S}$ as $Q_{t+1}(s, a, b) = r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [f_\tau(Q_t(s'); g_{\theta_{T_{inner}}}(\cdot | s'), h_{\nu_{T_{inner}}}(\cdot | s'))]$.

The complete algorithm is presented in Algorithm 6. Note that we use the name ONPG for Markov games because the

⁵We add a minus sign in front of Q in order to continue working with costs, instead of rewards, as done in the previous sections.

inner matrix game is solved using the ONPG updates. The two-timescale-type update rule (between the policy and value updates) for solving infinite-horizon Markov games has also been used before in Sayin et al. (2021); Cen et al. (2021b); Wei et al. (2021).

Next, we provide a convergence result for the performance of Algorithm 6 for the regularized Markov game.

Theorem 5.3. Let Q_τ^* be the NE Q -value of the regularized Markov Game under the tabular parametrization. Choose the stepsize $\eta = \frac{1-\gamma}{2(1+\tau(\log n+1-\gamma))}$ for the inner loop in Algorithm 6. Let T denote the total number of iterations ($T_{outer} \cdot T_{inner}$). Then, after

$$\begin{aligned} T_{inner} &= \mathcal{O}\left(\frac{1}{\eta\tau} \left(\log \frac{1}{\epsilon} + \log \frac{1}{1-\gamma} + \log \log n + \log \frac{1}{\eta}\right)\right), \\ T_{outer} &= \mathcal{O}\left(\frac{1}{1-\gamma} \left(\log \frac{1}{\epsilon} + \log \left(\frac{8}{\tau} (1 + C^2 \|Q^*\|_\infty^2)\right) + \log \frac{1 + \tau \log n}{1-\gamma}\right)\right), \end{aligned} \quad (5.6)$$

iterations, we have $\|Q_T - Q_\tau^*\|_F \leq \epsilon$ and $\max\{\|\theta_T - \theta^*\|, \|\nu_T - \nu^*\|\} \leq \epsilon$ where (Q_T, θ_T, ν_T) is the output of Algorithm 6 after T iterations, and (θ^*, ν^*) are defined in Equation (F.8).

Finally, we show how the optimistic NPG algorithm solves the original Markov game without regularization.

Corollary 5.4. If we run Algorithm 6 for time $T_{inner} = \mathcal{O}\left(\frac{\log n}{(1-\gamma)^2 \epsilon} \log\left(\frac{1}{\epsilon}\right)\right)$, $T_{outer} = \mathcal{O}\left(\frac{1}{(1-\gamma)} \log\left(\frac{1}{\epsilon}\right)\right)$, and setting $\tau = \mathcal{O}((1-\gamma)\epsilon/\log n)$, the output (θ_T, ν_T) will be an ϵ -in-class NE (Definition 5.1) of the original unregularized Markov game.

Remark 5.5. We remark that Theorem 5.3, to the best of our knowledge, is the first to show parameter convergence in Markov games with policy parametrization, and is different from the policy convergence results of several recent works Zhao et al. (2021); Wei et al. (2021); Cen et al. (2021b) (see §1.1 for a detailed comparison).

We extend these results to simple function approximation settings in §F.3 and have simulations in §G in the Appendix.

Acknowledgements

S.P. acknowledges support from MathWorks Engineering Fellowship. A.O and K.Z. were supported by MIT-DSTA grant 031017-00016. K.Z. also acknowledges support from Simons-Berkeley Research Fellowship.

References

Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. (2019). Optimality and approximation with policy gradient methods in Markov decision processes. *arXiv preprint arXiv:1908.00261*.

Alacaoglu, A., Viano, L., He, N., and Cevher, V. (2022). A natural actor-critic framework for zero-sum Markov games. In *International Conference on Machine Learning*, pages 307–366. PMLR.

Arora, S., Hazan, E., and Kale, S. (2012). The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing*, 8(1):121–164.

Azizian, W., Mitliagkas, I., Lacoste-Julien, S., and Gidel, G. (2020). A tight and unified analysis of gradient-based methods for a whole spectrum of differentiable games. In *International Conference on Artificial Intelligence and Statistics*, pages 2863–2873. PMLR.

Bailey, J. P. and Piliouras, G. (2018). Multiplicative weights update in zero-sum games. In *ACM Conference on Economics and Computation*, pages 321–338.

Bauschke, H. H., Borwein, J. M., and Combettes, P. L. (2003). Bregman monotone optimization algorithms. *SIAM Journal on control and optimization*, 42(2):596–636.

Beck, A. and Teboulle, M. (2003). Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175.

Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.

Branavan, S. R., Chen, H., Zettlemoyer, L. S., and Barzilay, R. (2009). Reinforcement learning for mapping instructions to actions. Association for Computational Linguistics.

Bu, J., Ratliff, L. J., and Mesbahi, M. (2019). Global convergence of policy gradient for sequential zero-sum linear quadratic dynamic games. *arXiv preprint arXiv:1911.04672*.

Cen, S., Cheng, C., Chen, Y., Wei, Y., and Chi, Y. (2021a). Fast global convergence of natural policy gradient methods with entropy regularization. *Operations Research*.

Cen, S., Wei, Y., and Chi, Y. (2021b). Fast policy extragradient methods for competitive games with entropy regularization. *arXiv preprint arXiv:2105.15186*.

Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press.

Daskalakis, C., Foster, D. J., and Golowich, N. (2020). Independent policy gradient methods for competitive reinforcement learning. In *Advances in Neural Information Processing Systems*.

Daskalakis, C. and Panageas, I. (2018). The limit points of (optimistic) gradient descent in min-max optimization. In *Neural Information Processing Systems*.

Daskalakis, C. and Panageas, I. (2019). Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *Innovations in Theoretical Computer Science (ITCS)*.

- Daskalakis, C., Skoulakis, S., and Zampetakis, M. (2021). The complexity of constrained min-max optimization. In *Proceedings of Symposium on Theory of Computing*, pages 1466–1478.
- Ding, D., Wei, C.-Y., Zhang, K., and Jovanovic, M. (2022). Independent policy gradient for large-scale Markov potential games: Sharper rates, function approximation, and game-agnostic convergence. In *International Conference on Machine Learning*, pages 5166–5220. PMLR.
- Facchinei, F. and Pang, J.-S. (2007). *Finite-Dimensional Variational Inequalities and Complementarity Problems*. Springer Science & Business Media.
- Fallah, A., Ozdaglar, A., and Pattathil, S. (2020). An optimal multistage stochastic gradient method for minimax problems. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 3573–3579. IEEE.
- Fazel, M., Ge, R., Kakade, S. M., and Mesbahi, M. (2018). Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*.
- Filar, J. and Vrieze, K. (2012). *Competitive Markov Decision Processes*. Springer Science & Business Media.
- Flokas, L., Vlatakis-Gkaragkounis, E.-V., and Piliouras, G. (2021). Solving min-max optimization with hidden structure via gradient descent ascent. *arXiv preprint arXiv:2101.05248*.
- Freund, Y. and Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139.
- Geist, M., Scherrer, B., and Pietquin, O. (2019). A theory of regularized Markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169.
- Golowich, N., Pattathil, S., and Daskalakis, C. (2020a). Tight last-iterate convergence rates for no-regret learning in multi-player games. *arXiv preprint arXiv:2010.13724*.
- Golowich, N., Pattathil, S., Daskalakis, C., and Ozdaglar, A. (2020b). Last iterate is slower than averaged iterate in smooth convex-concave saddle point problems. In *Conference on Learning Theory*, pages 1758–1784. PMLR.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*.
- Hambly, B. M., Xu, R., and Yang, H. (2021). Policy gradient methods find the Nash equilibrium in n-player general-sum linear-quadratic games. *Available at SSRN 3894471*.
- Hsieh, Y.-G., Iutzeler, F., Malick, J., and Mertikopoulos, P. (2019). On the convergence of single-call stochastic extra-gradient methods. *arXiv preprint arXiv:1908.08465*.
- Ibrahim, A., Azizian, W., Gidel, G., and Mitliagkas, I. (2020). Linear lower bounds and conditioning of differentiable games. In *International Conference on Machine Learning*, pages 4583–4593. PMLR.
- Jin, C., Netrapalli, P., and Jordan, M. (2020). What is local optimality in nonconvex-nonconcave minimax optimization? In *International Conference on Machine Learning*, pages 4880–4889. PMLR.
- Kakade, S. M. (2002). A natural policy gradient. In *Advances in Neural Information Processing Systems*, pages 1531–1538.
- Krantz, S. G. and Parks, H. R. (2012). *The Implicit Function Theorem: History, Theory, and Applications*. Springer Science & Business Media.
- Lan, G. (2021). Policy mirror descent for reinforcement learning: Linear convergence, new sampling complexity, and generalized problem classes. *arXiv preprint arXiv:2102.00135*.
- Leonardos, S., Overman, W., Panageas, I., and Piliouras, G. (2021). Global convergence of multi-agent policy gradient in Markov potential games. *arXiv preprint arXiv:2106.01969*.
- Lin, T., Jin, C., and Jordan, M. I. (2019). On gradient descent ascent for nonconvex-concave minimax problems. *arXiv preprint arXiv:1906.00331*.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, O. P., and Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, pages 6379–6390.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38.
- Mei, J., Xiao, C., Szepesvari, C., and Schuurmans, D. (2020). On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning*, pages 6820–6829. PMLR.
- Mertikopoulos, P. and Sandholm, W. H. (2016). Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324.
- Mladenovic, A., Sakos, I., Gidel, G., and Piliouras, G. (2021). Generalized natural gradient flows in hidden convex-concave games and gans. In *International Conference on Learning Representations*.
- Mokhtari, A., Ozdaglar, A., and Pattathil, S. (2020a). A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *International Conference on Artificial Intelligence and Statistics*, pages 1497–1507. PMLR.

- Mokhtari, A., Ozdaglar, A. E., and Pattathil, S. (2020b). Convergence rate of $\mathcal{O}(1/k)$ for optimistic gradient and extragradient methods in smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 30(4):3230–3251.
- Nemirovski, A. (2004). Prox-method with rate of convergence $o(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251.
- Neu, G., Jonsson, A., and Gómez, V. (2017). A unified view of entropy-regularized Markov decision processes. *arXiv preprint arXiv:1705.07798*.
- Neumann, J. v. (1928). Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320.
- Nouiehed, M., Sanjabi, M., Huang, T., Lee, J. D., and Razaviyayn, M. (2019). Solving a class of non-convex min-max games using iterative first order methods. In *Advances in Neural Information Processing Systems*, pages 14905–14916.
- OpenAI (2018). Openai five. <https://blog.openai.com/openai-five/>.
- Popov, L. D. (1980). A modification of the arrow-hurwicz method for search of saddle points. *Mathematical notes of the Academy of Sciences of the USSR*, 28(5):845–848.
- Qiu, S., Wei, X., Ye, J., Wang, Z., and Yang, Z. (2021). Provably efficient fictitious play policy optimization for zero-sum Markov games with structured transitions. In *International Conference on Machine Learning*, pages 8715–8725. PMLR.
- Rafique, H., Liu, M., Lin, Q., and Yang, T. (2018). Non-convex min-max optimization: Provable algorithms and applications in machine learning. *arXiv preprint arXiv:1810.02060*.
- Rakhlin, A. and Sridharan, K. (2013). Optimization, learning, and games with predictable sequences. In *Neural Information Processing Systems*, pages 3066–3074.
- Rosen, J. B. (1965). Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, pages 520–534.
- Sayin, M., Zhang, K., Leslie, D., Basar, T., and Ozdaglar, A. (2021). Decentralized q-learning in zero-sum markov games. *Advances in Neural Information Processing Systems*, 34:18320–18334.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015). Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shalev-Shwartz, S., Shammah, S., and Shashua, A. (2016). Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*.
- Shapley, L. S. (1953). Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359.
- Sokota, S., D’Orazio, R., Kolter, J. Z., Loizou, N., Lantot, M., Mitliagkas, I., Brown, N., and Kroer, C. (2022). A unified approach to reinforcement learning, quantal response equilibria, and two-player zero-sum games. *arXiv preprint arXiv:2206.05825*.
- Thekumparampil, K. K., Jain, P., Netrapalli, P., and Oh, S. (2019). Efficient algorithms for smooth minimax optimization. *arXiv preprint arXiv:1907.01543*.
- Tseng, P. (1995). On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60(1-2):237–252.
- Vlatakis-Gkaragkounis, E.-V., Flokas, L., and Piliouras, G. (2019). Poincaré recurrence, cycles and spurious equilibria in gradient-descent-ascent for non-convex non-concave zero-sum games. In *Advances in Neural Information Processing Systems*, pages 10450–10461.
- Wang, L., Cai, Q., Yang, Z., and Wang, Z. (2019). Neural policy gradient methods: Global optimality and rates of convergence. *arXiv preprint arXiv:1909.01150*.
- Wei, C.-Y., Lee, C.-W., Zhang, M., and Luo, H. (2020). Linear last-iterate convergence in constrained saddle-point optimization.
- Wei, C.-Y., Lee, C.-W., Zhang, M., and Luo, H. (2021). Last-iterate convergence of decentralized optimistic gradient descent/ascent in infinite-horizon competitive Markov games. *arXiv preprint arXiv:2102.04540*.
- Yang, J., Kiyavash, N., and He, N. (2020). Global convergence and variance-reduced optimization for a class of nonconvex-nonconcave minimax problems. *arXiv preprint arXiv:2002.09621*.
- Yu, C., Velu, A., Vinitsky, E., Wang, Y., Bayen, A., and Wu, Y. (2021). The surprising effectiveness of MAPPO in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*.
- Zhan, W., Cen, S., Huang, B., Chen, Y., Lee, J. D., and Chi, Y. (2021). Policy mirror descent for regularized reinforcement learning: A generalized framework with linear convergence. *arXiv preprint arXiv:2105.11066*.
- Zhang, K., Kakade, S. M., Başar, T., and Yang, L. F. (2020). Model-based multi-agent RL in zero-sum Markov games with near-optimal sample complexity. *arXiv preprint arXiv:2007.07461*.

- Zhang, K., Koppel, A., Zhu, H., and Başar, T. (2019a). Global convergence of policy gradient methods to (almost) locally optimal policies. *arXiv preprint arXiv:1906.08383*.
- Zhang, K., Yang, Z., and Başar, T. (2019b). Policy optimization provably converges to Nash equilibria in zero-sum linear quadratic games. In *Advances in Neural Information Processing Systems*.
- Zhang, R., Liu, Q., Wang, H., Xiong, C., Li, N., and Bai, Y. (2022). Policy optimization for markov games: Unified framework and faster convergence. *arXiv preprint arXiv:2206.02640*.
- Zhang, R., Ren, Z., and Li, N. (2021). Gradient play in multi-agent Markov stochastic games: Stationary points and convergence. *arXiv preprint arXiv:2106.00198*.
- Zhao, Y., Tian, Y., Lee, J. D., and Du, S. S. (2021). Provably efficient policy gradient methods for two-player zero-sum markov games. *arXiv preprint arXiv:2102.08903*.

Supplementary Materials for “Symmetric (Optimistic) Natural Policy Gradient for Multi-agent Learning with Parameter Convergence”

A Missing Definitions and Proofs in §2

A.1 Proof of Lemma 2.2

We show that the problem

$$\min_{\theta \in \mathbb{R}^n} \max_{\nu \in \mathbb{R}^n} g_{\theta}^{\top} Q h_{\nu}, \quad (\text{A.1})$$

is nonconvex-nonconcave.

Let

$$Q = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \quad (\text{A.2})$$

Consider $\theta_1 = (0, 0)$ and $\theta_2 = (\log 4, \log 9)$. This implies $g_{\theta_1} = (1/2, 1/2)^{\top}$ and $g_{\theta_2} = (4/13, 9/13)^{\top}$. Also, from the form of Q , we have $Q h_{\nu} = [h_{\nu}(1), h_{\nu}(2)]^{\top}$, $\forall \nu$. Now, for $[h_{\nu}(1), h_{\nu}(2)] = [1/3, 2/3]$, we have

$$\frac{1}{2}(g_{\theta_1}^{\top} Q h_{\nu} + g_{\theta_2}^{\top} Q h_{\nu}) < g_{(\theta_1 + \theta_2)/2} Q h_{\nu}, \quad (\text{A.3})$$

which implies nonconvexity in θ .

Similarly, taking $\nu_1 = (0, 0)$ and $\nu_2 = (\log 4, \log 9)$ (which implies $h_{\nu_1} = (1/2, 1/2)^{\top}$ and $h_{\nu_2} = (4/13, 9/13)^{\top}$, and taking $g_{\theta} = (2/3, 1/3)^{\top}$, we have

$$\frac{1}{2}(g_{\theta}^{\top} Q h_{\nu_1} + g_{\theta}^{\top} Q h_{\nu_2}) > g_{\theta} Q h_{(\nu_1 + \nu_2)/2}, \quad (\text{A.4})$$

which implies nonconcavity in ν .

Note that adding regularization does not get rid of this convexity. For example consider the specific case when h_{ν} is a constant policy and the matrix Q is 0. We show the nonconvexity of the function $-\tau \mathcal{H}(g_{\theta})$ in θ next. Consider $\theta_1 = (0, 0)$ and $\theta_2 = (\log 4, \log 9)$. We have $g_{\theta_1} = (1/2, 1/2)$ and $g_{\theta_2} = (4/13, 9/13)$. Furthermore, we have $g_{(\theta_1 + \theta_2)/2} = (1/3, 2/3)$, and we can see that:

$$\frac{-\tau}{2}(\mathcal{H}(g_{\theta_1}) + \mathcal{H}(g_{\theta_2})) < -\tau \mathcal{H}(g_{(\theta_1 + \theta_2)/2}) \quad (\text{A.5})$$

which shows nonconvexity of $-\tau \mathcal{H}(g_{\theta})$. This completes the proof of the lemma. \square

A.2 Vanilla NPG for matrix games

Next, we compute the Fisher Information Matrix $F_{\theta}(\theta) = \mathbb{E}_{a \sim g_{\theta}} [(\nabla_{\theta} \log g_{\theta}(a))(\nabla_{\theta} \log g_{\theta}(a))^{\top}]$. For the softmax parametrization, we have:

$$\nabla_{\theta} \log g_{\theta}(a) = \nabla_{\theta} \left(\theta(a) - \log \left(\sum_{a' \in \mathcal{A}} e^{\theta(a')} \right) \right) = [-g_{\theta}(1), -g_{\theta}(2), \dots, 1 - g_{\theta}(a), \dots, -g_{\theta}(n)]^{\top}. \quad (\text{A.6})$$

Now, consider the $(i, i)^{th}$ element of the Fisher information matrix. We have:

$$[F_{\theta}(\theta)]_{ii} = g_{\theta}(i)(1 - g_{\theta}(i))(1 - g_{\theta}(i)) + \sum_{j \neq i} g_{\theta}(j)g_{\theta}(i)^2 = g_{\theta}(i)(1 - g_{\theta}(i)). \quad (\text{A.7})$$

Similarly, we have the $(i, j)^{th}$ element, where $i \neq j$ is given by:

$$[F_\theta(\theta)]_{ij} = (1 - g_\theta(i) - g_\theta(j))g_\theta(i)g_\theta(j) - g_\theta(i)(1 - g_\theta(i))g_\theta(j) - g_\theta(j)(1 - g_\theta(j))g_\theta(i) = -g_\theta(i)g_\theta(j). \quad (\text{A.8})$$

Therefore, the matrix $F_\theta(\theta)$ can be succinctly written as:

$$F_\theta(\theta) = \text{diag}(g_\theta) - g_\theta g_\theta^\top, \quad (\text{A.9})$$

where $\text{diag}(g_\theta)$ is a diagonal matrix with entries g_θ . Note that this is in fact $\nabla_\theta g_\theta$ (see [Mei et al. \(2020\)](#)), i.e.,

$$\nabla_\theta g_\theta = \text{diag}(g_\theta) - g_\theta g_\theta^\top. \quad (\text{A.10})$$

Therefore, we have:

$$F_\theta^\dagger(\theta) \nabla_\theta g_\theta = I. \quad (\text{A.11})$$

The update of the vanilla NPG thus simplifies to the following:

$$\begin{aligned} \theta_{t+1} &= \theta_t - \eta \cdot F_\theta^\dagger(\theta_t) \cdot \frac{\partial f_\tau(\theta_t, \nu_t)}{\partial \theta} = \theta_t - \eta \frac{\partial f_\tau(\theta_t, \nu_t)}{\partial g_\theta} \\ &= \theta_t - \eta (Qh_{\nu_t} + \tau(\mathbf{1} + \log g_{\theta_t})). \end{aligned} \quad (\text{A.12})$$

However, since $g_\theta(a) = \frac{e^{\theta(a)}}{\sum_{a' \in \mathcal{A}} e^{\theta(a')}}$, we have

$$\theta_{t+1}(a) = (1 - \eta\tau)\theta_t(a) - \eta \left([Qh_{\nu_t}]_a + \tau - \tau \log \sum_{a' \in \mathcal{A}} e^{\theta_t(a')} \right). \quad (\text{A.13})$$

A similar update for ν leads to the updates in Equations (2.5)-(2.6). Note that when we write a constant in the update, we mean a constant vector with all elements being the same.

A.3 Proof of Lemma 2.3

We restate the lemma here first for convenience:

Lemma A.1 (Pitfall of vanilla NPG). There exists a game (2.4) with $\tau \geq 0$ (we allow for unregularized games as well) and a dummy player 2, i.e., $|\mathcal{B}| = 1$, for which the updates (2.5)-(2.6) do not converge for any $\eta > 0$.

Proof. Consider the θ update under NPG:

$$\theta_{t+1}(a) = (1 - \eta\tau)\theta_t(a) - \eta \left([Qh_{\nu_t}]_a + \tau - \tau \log \sum_{a' \in \mathcal{A}} e^{\theta_t(a')} \right). \quad (\text{A.14})$$

From here, it is easy to see that it need not converge for the case where $\tau = 0$, since this would require $Qh_{\nu_t} = 0$ which need not be the case (For example consider $Q = [1 \mid 1]^\top$, and $h_\nu = [1]$. In this case, $h_{\nu_t} = 1$ for any parameter ν_t).

Next, we consider the case where $\tau > 0$. Suppose θ converges to some point θ^* . Since $|\mathcal{B}| = 1$, we have $h_{\nu_t} = [1]$. Substituting the point θ^* into the update we have:

$$\theta^*(a) = (1 - \eta\tau)\theta^*(a) - \eta \left([Q]_a + \tau - \tau \log \sum_{a' \in \mathcal{A}} e^{\theta^*(a')} \right). \quad (\text{A.15})$$

This implies:

$$\eta\tau\theta^*(a) = -\eta\tau \left(\frac{[Q]_a}{\tau} + 1 \right) + \eta\tau \log \sum_{a' \in \mathcal{A}} e^{\theta^*(a')}. \quad (\text{A.16})$$

This leads to:

$$\log e^{\theta^*(a)} - \log \sum_{a' \in \mathcal{A}} e^{\theta^*(a')} = -\frac{[Q]_a}{\tau} - 1. \quad (\text{A.17})$$

Algorithm 1 Optimistic NPG

Initialize: $\theta_0 = 0$ and $\nu_0 = 0$.
for $t = 1, 2, \dots$ **do**
 $\bar{\theta}_{t+1} = (1 - \eta\tau)\theta_t - \eta Q h_{\bar{\nu}_t}$
 $\bar{\nu}_{t+1} = (1 - \eta\tau)\nu_t + \eta Q^\top g_{\bar{\theta}_t}$

 $\theta_{t+1} = (1 - \eta\tau)\theta_t - \eta Q h_{\bar{\nu}_{t+1}}$
 $\nu_{t+1} = (1 - \eta\tau)\nu_t + \eta Q^\top g_{\bar{\theta}_{t+1}}$
end for

However,

$$\log e^{\theta^*(a)} - \log \sum_{a' \in \mathcal{A}} e^{\theta^*(a')} = \log \frac{e^{\theta^*(a)}}{\sum_{a' \in \mathcal{A}} e^{\theta^*(a')}} = \log g_{\theta^*}(a). \quad (\text{A.18})$$

Substituting this in Equation (A.17), we have:

$$g_{\theta^*}(a) = \exp\left(-\frac{[Q]_a}{\tau} - 1\right). \quad (\text{A.19})$$

This need not be a valid probability measure. For example, consider $Q = [-2 \mid 2]^\top$ and $\tau = 1$, we have:

$$g_{\theta^*}(1) = e > 1, \quad (\text{A.20})$$

which contradicts the fact that g_θ is a probability measure. This implies that the original NPG updates cannot have a fixed point, and therefore does not converge for any stepsize $\eta > 0$. \square

B Missing Details and Proofs in §3

Remark B.1. We note that all results presented in this section also follow for the case where the action spaces for both players are asymmetric. However, we stick to the case where the number of actions is the same for both players, for ease of exposition.

B.1 Proof of Theorem 3.2

B.1.1 Policy convergence

Consider the following modified NPG updates for the regularized game:

$$\theta_{t+1} = (1 - \eta\tau)\theta_t - \eta Q h_{\nu_t}, \quad (\text{B.1})$$

$$\nu_{t+1} = (1 - \eta\tau)\nu_t + \eta Q^\top g_{\theta_t}. \quad (\text{B.2})$$

Note that these updates correspond to the popular Multiplicative Weights Update (Freund and Schapire, 1997; Arora et al., 2012) for the regularized game in policy space (we succinctly represent g_{θ_t} and h_{ν_t} as g_t and h_t , respectively), i.e.,

$$\begin{aligned} g_{t+1}(a) &\propto g_t(a)^{(1-\eta\tau)} e^{-\eta[Qh_t]_a}, \\ h_{t+1}(b) &\propto h_t(b)^{(1-\eta\tau)} e^{\eta[Q^\top g_t]_b}. \end{aligned} \quad (\text{B.3})$$

We can write these updates as a mirror descent update with Bregman function given by the negative entropy (i.e., the corresponding Bregman distance is the KL divergence) as follows:

$$\begin{aligned} g_{t+1} &= \operatorname{argmin}_{g \in \Delta} \{ \langle Qh_t + \tau \nabla_g \mathcal{H}(g_t), g \rangle + \text{KL}(g \| g_t) \}, \\ h_{t+1} &= \operatorname{argmax}_{h \in \Delta} \{ \langle Q^\top g_t - \tau \nabla_h \mathcal{H}(h_t), h \rangle - \text{KL}(h \| h_t) \}. \end{aligned} \quad (\text{B.4})$$

Note that we can write these updates succinctly as one Mirror Descent update in the following form:

$$z_{t+1} = \operatorname{argmin}_{z \in \mathcal{Z}} \{ \langle Mz_t + \tau \nabla \mathcal{H}(z_t), z \rangle + \text{KL}(z \| z_t) \}, \quad (\text{B.5})$$

where $z = [g; h]$, $\mathcal{Z} \in \Delta \times \Delta$ and with slight abuse of notation, we define $\nabla \mathcal{H}(z) = [\nabla_g \mathcal{H}(g); \nabla_h \mathcal{H}(h)]$. Also, we define the matrix

$$M = \begin{pmatrix} 0 & Q \\ -Q^\top & 0 \end{pmatrix}. \quad (\text{B.6})$$

We can now use properties of mirror decent to analyze the iterates of MWU.

First, we have the following two lemmas which follow from [Bauschke et al. \(2003\)](#), Proposition 2.3, and Lemma D.4 in [Sokota et al. \(2022\)](#) which will be used to derive the final convergence rate:

Lemma B.2. For all $z \in \mathcal{Z}$, we have

$$\eta \langle Mz_t + \tau \nabla \mathcal{H}(z_t), z_{t+1} - z \rangle \leq \text{KL}(z \| z_t) - \text{KL}(z \| z_{t+1}) - \text{KL}(z_{t+1} \| z_t). \quad (\text{B.7})$$

Lemma B.3. For all $z \in \mathcal{Z}$, we have

$$\eta \langle Mz + \tau \nabla \mathcal{H}(z), z^* - z \rangle \leq -\eta \tau (\text{KL}(z \| z^*) + \text{KL}(z^* \| z)). \quad (\text{B.8})$$

In the next lemma, we show that the iterates of MWU on the regularized problem will be bounded away from the boundary of the simplex.

Lemma B.4. For $\eta < 1/\tau$, the iterates of regularized MWU stay within a set $\Delta' \subset \Delta$ which is bounded away from the boundary of the simplex, i.e., $x_i \geq \delta > 0$ for some $\delta > 0$, $\forall x \in \Delta'$.

Proof. Consider the update of g . We have the following property from Mirror Descent (see [Beck and Teboulle \(2003\)](#))

$$\begin{aligned} \text{KL}(g^* \| g_{t+1}) &\leq (1 - \eta\tau) \text{KL}(g^* \| g_t) - (1 - \eta\tau) \text{KL}(g_{t+1} \| g_t) - \eta\tau \mathcal{H}(g_{t+1}) + \eta\tau \mathcal{H}(g^*) \\ &\quad - \eta \langle g_{t+1} - g^*, Qh_t \rangle \\ &\leq (1 - \eta\tau) \text{KL}(g^* \| g_t) + 2\eta\tau \log n + 2\eta \|Q\|_\infty. \end{aligned} \quad (\text{B.9})$$

This implies that

$$\text{KL}(g^* \| g_t) \leq 2 \log n + \frac{2 \|Q\|_\infty}{\tau} + \text{KL}(g^* \| g_0). \quad (\text{B.10})$$

From the definition of the KL divergence and Equation (B.10), we have:

$$g_{t,i} \geq \exp \left(\frac{-1}{g_{\min}^*} \left(2 \log n + \frac{2 \|Q\|_\infty}{\tau} + \text{KL}(g^* \| g_0) + \mathcal{H}(g^*) \right) \right) > 0, \quad \forall i, \text{ and } \forall t. \quad (\text{B.11})$$

Here g_{\min}^* is the smallest value of the Nash Equilibrium policy (which is greater than 0 since the NE policy of the regularized game is in the interior of the simplex). This completes the proof of the Lemma. \square

Since the iterates lie within Δ' , we let L denote the Lipschitz constant of $Mz + \tau \nabla \mathcal{H}(z)$ (a continuous function over $\Delta \times \Delta$ whose norm approaches infinity as z approaches the boundary) in the set \mathcal{Z}' (where $\mathcal{Z}' = \Delta' \times \Delta'$), i.e.,

$$\| (Mz_1 + \tau \nabla \mathcal{H}(z_1)) - (Mz_2 + \tau \nabla \mathcal{H}(z_2)) \| \leq L \|z_1 - z_2\|, \quad \forall z_1, z_2 \in \mathcal{Z}'. \quad (\text{B.12})$$

Now, we use these lemmas to derive the convergence rate of MWU for the regularized problem.

Theorem B.5. Consider the modified NPG updates in Equation (B.1)-(B.2) with stepsize satisfying $0 \leq \eta \leq \tau/L^2$. We then have:

$$\text{KL}(z^* \| z_{t+1}) \leq \left(1 - \frac{\eta\tau}{2} \right) \text{KL}(z^* \| z_t). \quad (\text{B.13})$$

Proof. Note that the constraint $\eta \leq \tau/L^2$ will automatically satisfy $\eta < 1/\tau$ (as needed by Lemma B.4) since $\tau \leq L$. We have the following string of inequalities:

$$\begin{aligned}
 \text{KL}(z^* \| z_{t+1}) &\stackrel{*1}{\leq} \text{KL}(z^* \| z_t) - \text{KL}(z_{t+1} \| z_t) + \eta \langle F(z_t) + \tau \nabla \mathcal{H}(z_t), z^* - z_{t+1} \rangle \\
 &= \text{KL}(z^* \| z_t) - \text{KL}(z_{t+1} \| z_t) + \eta \langle F(z_{t+1}) + \tau \nabla g(z_{t+1}), z^* - z_{t+1} \rangle \\
 &\quad + \eta \langle (F(z_t) + \tau \nabla \mathcal{H}(z_t)) - (F(z_{t+1}) + \tau \nabla \mathcal{H}(z_{t+1})), z^* - z_{t+1} \rangle \\
 &\stackrel{*2}{\leq} \text{KL}(z^* \| z_t) - \text{KL}(z_{t+1} \| z_t) - \eta \tau (\text{KL}(z_{t+1} \| z^*) + \text{KL}(z^* \| z_{t+1})) \\
 &\quad + \eta \langle (F(z_t) + \tau \nabla \mathcal{H}(z_t)) - (F(z_{t+1}) + \tau \nabla \mathcal{H}(z_{t+1})), z^* - z_{t+1} \rangle \\
 &\stackrel{*3}{\leq} \text{KL}(z^* \| z_t) - \text{KL}(z_{t+1} \| z_t) - \eta \tau (\text{KL}(z_{t+1} \| z^*) + \text{KL}(z^* \| z_{t+1})) \\
 &\quad + \eta L \|z_{t+1} - z_t\| \|z^* - z_{t+1}\| \\
 &\stackrel{*4}{\leq} \text{KL}(z^* \| z_t) - \text{KL}(z_{t+1} \| z_t) - \eta \tau (\text{KL}(z_{t+1} \| z^*) + \text{KL}(z^* \| z_{t+1})) \\
 &\quad + \frac{1}{2} \|z_{t+1} - z_t\|^2 + \frac{\eta^2 L^2}{2} \|z^* - z_{t+1}\|^2 \\
 &\stackrel{*5}{\leq} \text{KL}(z^* \| z_t) - \text{KL}(z_{t+1} \| z_t) - \text{KL}(z_{t+1} \| z^*) + \eta^2 L^2 \text{KL}(z_{t+1} \| z^*) \\
 &\quad - \eta \tau \text{KL}(z_{t+1} \| z^*) - \eta \tau \text{KL}(z^* \| z_{t+1}) \\
 &\stackrel{*6}{\leq} \text{KL}(z^* \| z_t) - \eta \tau \text{KL}(z^* \| z_{t+1}). \tag{B.14}
 \end{aligned}$$

Here (*1) follows from Lemma B.2, (*2) follows from Lemma B.3, (*3) follows from Equation (B.12), (*4) follows from Young's inequality, (*5) follows from Pinskers inequality and (*6) follows from $\eta \leq \tau/L^2$. Therefore, we have

$$\text{KL}(z^* \| z_{t+1}) \leq \frac{1}{1 + \eta \tau} \text{KL}(z^* \| z_t) \leq \left(1 - \frac{\eta \tau}{2}\right) \text{KL}(z^* \| z_t), \tag{B.15}$$

which completes the proof. \square

B.1.2 Parameter convergence

In this section, we show convergence of the policy parameters. We have the following theorem.

Theorem B.6. Consider the modified NPG updates in Equation (B.1)-(B.2) with stepsize satisfying $0 \leq \eta \leq \tau/L^2$. We then have:

$$\|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 \leq (1 - \eta \tau / 4)^t \left(\|\theta_0 - \theta^*\|^2 + \|\nu_0 - \nu^*\|^2 + \frac{4C}{\eta \tau} \right), \tag{B.16}$$

where $\theta^* = \frac{-Qh^*}{\tau}$, $\nu^* = \frac{Q^\top g^*}{\tau}$ and $C = \left(1 + \frac{1}{\eta \tau} (1 - \eta \tau)^2\right) 4\eta^2 \|Q\|_\infty^2 \text{KL}(z^* \| z_0)$.

Proof. We begin by first providing the intuition of the proof, when the opponent is playing the NE strategy. We denote the NE strategies of the players as g^* and h^* . Then, the NPG update has the following form:

$$\theta_{t+1} = (1 - \eta \tau) \theta_t - \eta Q h^*. \tag{B.17}$$

We know that the NE satisfy:

$$g^*(a) = \frac{e^{-[Qh^*]_a/\tau}}{\sum_{a' \in \mathcal{A}} e^{-[Qh^*]_{a'}/\tau}} = \frac{e^{-[Qh^*]_a/\tau}}{K}, \tag{B.18}$$

where we define $K := \sum_{a' \in \mathcal{A}} e^{-[Qh^*]_{a'}/\tau}$. Taking log on both sides, we have:

$$-\eta Q h^* = \eta \tau \log g^* + \eta \tau \log K. \tag{B.19}$$

Substituting this back into the θ update in (B.17), we have:

$$\begin{aligned}
 \theta_{t+1} &= (1 - \eta\tau)\theta_t + \eta\tau \log g^* + \eta\tau \log K \\
 &= \theta_t - \eta\tau\theta_t + \eta\tau \log g^* + \eta\tau \log K - \eta\tau \log Z_{\theta_t} + \eta\tau \log Z_{\theta_t} \\
 &= \theta_t + \eta\tau \log g^* - \eta\tau \log \left(\frac{e^{\theta_t}}{Z_{\theta_t}} \right) + \eta\tau \log \left(\frac{K}{Z_{\theta_t}} \right) \\
 &= \theta_t + \eta\tau \log g^* - \eta\tau \log g_{\theta_t} + \eta\tau \log \left(\frac{K}{Z_{\theta_t}} \right) = \theta_t + \eta\tau \log \left(\frac{g^*}{g_{\theta_t}} \right) + \eta\tau \log \left(\frac{K}{Z_{\theta_t}} \right) \\
 &= \theta_t + \eta\tau \log \left(\frac{g^* K}{g_{\theta_t} Z_{\theta_t}} \right). \tag{B.20}
 \end{aligned}$$

We can further simplify this as:

$$\theta_{t+1} = \theta_t + \eta\tau \log \left(\frac{e^{-[Qh^*]/\tau}}{e^{\theta_t}} \right), \tag{B.21}$$

which is nothing but:

$$\theta_{t+1} = \theta_t + \eta\tau \left(\frac{-Qh^*}{\tau} - \theta_t \right). \tag{B.22}$$

Note that this is the Gradient Descent update on the strongly convex function $\frac{1}{2} \left\| \frac{-Qh^*}{\tau} - \theta \right\|^2$ with stepsize $\eta\tau$. This update leads to the following convergence guarantees:

$$\|\theta_{t+1} - \theta^*\|^2 \leq (1 - \eta\tau)^2 \|\theta_t - \theta^*\|^2, \tag{B.23}$$

where $\theta^* = \frac{-Qh^*}{\tau}$.

The analysis above shows that if one of the players is already at the NE strategy, the parameters of the second player converges to the NE at a linear rate. However, the original NPG update for θ is given by

$$\theta_{t+1} = (1 - \eta\tau)\theta_t - \eta Q h^* + \eta Q (h^* - h_{\nu_t}). \tag{B.24}$$

Since $h_{\nu_{t+1}}$ converges to h^* at a linear rate (from Theorem B.5), we expect the term

$$\varepsilon_t = \eta Q (h^* - h_{\nu_t}), \tag{B.25}$$

to be small, and goes to 0. This is formalized in what follows.

The NPG update for θ can be re-written using ε_t as:

$$\theta_{t+1} = \theta_t + \eta\tau \left(\frac{-Qh^*}{\tau} - \theta_t \right) + \varepsilon_t. \tag{B.26}$$

Once again, defining $\theta^* = \frac{-Qh^*}{\tau}$, we have:

$$\begin{aligned}
 \|\theta_{t+1} - \theta^*\|^2 &= \|\theta_t + \eta\tau(\theta^* - \theta_t) + \varepsilon_t - \theta^*\|^2 = \|(1 - \eta\tau)(\theta_t - \theta^*) + \varepsilon_t\|^2 \\
 &= (1 - \eta\tau)^2 \|\theta_t - \theta^*\|^2 + 2(1 - \eta\tau)(\theta_t - \theta^*)^\top \varepsilon_t + \|\varepsilon_t\|^2 \\
 &\leq^{*1} (1 - \eta\tau)^2 \|\theta_t - \theta^*\|^2 + \eta\tau \|\theta_t - \theta^*\|^2 + \frac{1}{\eta\tau} (1 - \eta\tau)^2 \|\varepsilon_t\|^2 + \|\varepsilon_t\|^2 \\
 &= (1 - \eta\tau + \eta^2\tau^2) \|\theta_t - \theta^*\|^2 + \left(1 + \frac{1}{\eta\tau} (1 - \eta\tau)^2 \right) \|\varepsilon_t\|^2, \tag{B.27}
 \end{aligned}$$

where *1 follows from Young's inequality.

Next, we analyze the error term $\|\varepsilon_t\|$. We have:

$$\|\varepsilon_t\|^2 = \|\eta Q (h^* - h_{\nu_t})\|^2 \leq \eta^2 \|Q\|_\infty^2 \|h^* - h_{\nu_t}\|_1^2 \leq^{*1} 2\eta^2 \|Q\|_\infty^2 (\text{KL}(h^* \| h_{\nu_t})). \tag{B.28}$$

Here *1 follows from Pinsker's Inequality. Now, writing the same inequality for ν , we have:

$$\begin{aligned} \|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 &\leq (1 - \eta\tau + \eta^2\tau^2)(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) \\ &\quad + \left(1 + \frac{1}{\eta\tau}(1 - \eta\tau)^2\right) 2\eta^2\|Q\|_\infty^2 \text{KL}(z^*\|z_t) \\ &\leq (1 - \eta\tau + \eta^2\tau^2)(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + \left(1 + \frac{1}{\eta\tau}(1 - \eta\tau)^2\right) 4\eta^2C\|Q\|_\infty^2 \left(1 - \frac{\eta\tau}{2}\right)^t \text{KL}(z^*\|z_0). \end{aligned} \quad (\text{B.29})$$

Define:

$$C = \left(1 + \frac{1}{\eta\tau}(1 - \eta\tau)^2\right) 4\eta^2\|Q\|_\infty^2 \text{KL}(z^*\|z_0). \quad (\text{B.30})$$

Substituting back in Equation (B.51), along with the corresponding expression for ν , we have:

$$\|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 \leq (1 - \eta\tau + \eta^2\tau^2)(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + C \left(1 - \frac{\eta\tau}{2}\right)^t. \quad (\text{B.31})$$

For $\eta\tau < 1/2$ we have:

$$\|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 \leq (1 - \eta\tau/4)(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + C(1 - \eta\tau/2)^t. \quad (\text{B.32})$$

Consider the Lyapunov function:

$$V_{t+1} = \|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 + \frac{4C}{\eta\tau}(1 - \eta\tau/2)^{t+1}. \quad (\text{B.33})$$

We have:

$$\begin{aligned} V_{t+1} &\leq (1 - \eta\tau/4)(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + C(1 - \eta\tau/2)^t + \frac{4C}{\eta\tau}(1 - \eta\tau/2)(1 - \eta\tau/2)^t \\ &= (1 - \eta\tau/4)(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + \frac{4C}{\eta\tau}(1 - \eta\tau)^t(1 - \eta\tau/4) \\ &= (1 - \eta\tau/4) \left(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2 + \frac{4C}{\eta\tau}(1 - \eta\tau)^t \right) = (1 - \eta\tau/4)V_t. \end{aligned} \quad (\text{B.34})$$

This shows linear convergence of the parameter θ to θ^* since:

$$\|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 \leq V_{t+1} \leq (1 - \eta\tau/4)^t V_0. \quad (\text{B.35})$$

This completes the proof. \square

B.2 Proof of Theorem 3.4

We first prove the following result which follows from Theorem 1 in Cen et al. (2021b).

Lemma B.7. [Theorem 1, Cen et al. (2021b)] Consider Algorithm 1. Suppose that the learning rates satisfy:

$$0 < \eta \leq \min \left\{ \frac{1}{2\tau + 2\|Q\|_\infty}, \frac{1}{4\|Q\|_\infty} \right\}. \quad (\text{B.36})$$

Let $z_t = (g_{\theta_t}, h_{\nu_t})$ and $\bar{z}_t = (g_{\bar{\theta}_t}, h_{\bar{\nu}_t})$. Then:

$$\max \left\{ \text{KL}(z^*\|z_t), \frac{1}{2}\text{KL}(z^*\|\bar{z}_{t+1}) \right\} \leq (1 - \eta\tau)^t \text{KL}(z^*\|z_0). \quad (\text{B.37})$$

Proof. Let g_t and \bar{g}_t denote g_{θ_t} and $g_{\bar{\theta}_t}$ respectively. Also, for any parameter θ , we denote Z_θ to be the normalizing

constant $\sum_{a' \in \mathcal{A}} e^{\theta(a')}$ (Define Z_ν similarly). We have:

$$\begin{aligned}
 \bar{g}_{t+1}(a) &= \frac{e^{\bar{\theta}_{t+1}(a)}}{\sum_{a' \in \mathcal{A}} e^{\bar{\theta}_{t+1}(a')}} \\
 &\propto e^{\bar{\theta}_{t+1}(a)} = e^{(1-\eta\tau)\theta_t(a) - \eta[Qh_{\bar{v}_t}]_a} = e^{(1-\eta\tau)\theta_t(a) + \log Z_{\theta_t} - \log Z_{\theta_t} - \eta[Qh_{\bar{v}_t}]_a} \\
 &= e^{(1-\eta\tau) \log e^{\theta_t(a)} + \log Z_{\theta_t} - \log Z_{\theta_t} - \eta[Qh_{\bar{v}_t}]_a} = e^{(1-\eta\tau) \log \left(\frac{e^{\theta_t(a)}}{Z_{\theta_t}} \right) + \log Z_{\theta_t} - \eta[Qh_{\bar{v}_t}]_a} \\
 &\propto e^{(1-\eta\tau) \log \left(\frac{e^{\theta_t(a)}}{Z_{\theta_t}} \right) - \eta[Qh_{\bar{v}_t}]_a} = e^{(1-\eta\tau) \log g_t(a) - \eta[Qh_{\bar{v}_t}]_a} = e^{\log g_t(a)^{(1-\eta\tau)} - \eta[Qh_{\bar{v}_t}]_a} \\
 &= g_t(a)^{(1-\eta\tau)} e^{-\eta[Qh_{\bar{v}_t}]_a}.
 \end{aligned} \tag{B.38}$$

Therefore:

$$\bar{g}_{t+1}(a) \propto g_t(a)^{(1-\eta\tau)} e^{-\eta[Qh_{\bar{v}_t}]_a}. \tag{B.39}$$

Similarly, we have:

$$\begin{aligned}
 g_{t+1}(a) &\propto g_t(a)^{(1-\eta\tau)} e^{-\eta[Qh_{\bar{v}_{t+1}}]_a}, \\
 \bar{h}_{t+1}(a) &\propto h_t(a)^{(1-\eta\tau)} e^{\eta[Q^\top g_{\bar{h}_t}]_a}, \\
 h_{t+1}(a) &\propto h_t(a)^{(1-\eta\tau)} e^{\eta[Q^\top g_{\bar{h}_{t+1}}]_a},
 \end{aligned} \tag{B.40}$$

which is the same as the OMW updates for the regularized problem in [Cen et al. \(2021b\)](#). Therefore, by Theorem 1 in [Cen et al. \(2021b\)](#), we have convergence of g_θ and h_ν to the solution of the regularized min-max problem. \square

We begin by first providing the intuition of the proof, when the opponent is playing the NE strategy. We denote the NE strategies of the players as g^* and h^* . Then, the optimistic NPG update has the following form:

$$\theta_{t+1} = (1 - \eta\tau)\theta_t - \eta Qh^*. \tag{B.41}$$

From [Mertikopoulos and Sandholm \(2016\)](#), we know that the NE satisfy:

$$g^*(a) = \frac{e^{-[Qh^*]_a/\tau}}{\sum_{a' \in \mathcal{A}} e^{-[Qh^*]_{a'}/\tau}} = \frac{e^{-[Qh^*]_a/\tau}}{K}, \tag{B.42}$$

where we define $K := \sum_{a' \in \mathcal{A}} e^{-[Qh^*]_{a'}/\tau}$. Taking log on both sides, we have:

$$-\eta Qh^* = \eta\tau \log g^* + \eta\tau \log K. \tag{B.43}$$

Substituting this back into the θ update in (B.41), we have:

$$\begin{aligned}
 \theta_{t+1} &= (1 - \eta\tau)\theta_t + \eta\tau \log g^* + \eta\tau \log K \\
 &= \theta_t - \eta\tau\theta_t + \eta\tau \log g^* + \eta\tau \log K - \eta\tau \log Z_{\theta_t} + \eta\tau \log Z_{\theta_t} \\
 &= \theta_t + \eta\tau \log g^* - \eta\tau \log \left(\frac{e^{\theta_t}}{Z_{\theta_t}} \right) + \eta\tau \log \left(\frac{K}{Z_{\theta_t}} \right) \\
 &= \theta_t + \eta\tau \log g^* - \eta\tau \log g_{\theta_t} + \eta\tau \log \left(\frac{K}{Z_{\theta_t}} \right) = \theta_t + \eta\tau \log \left(\frac{g^*}{g_{\theta_t}} \right) + \eta\tau \log \left(\frac{K}{Z_{\theta_t}} \right) \\
 &= \theta_t + \eta\tau \log \left(\frac{g^* K}{g_{\theta_t} Z_{\theta_t}} \right).
 \end{aligned} \tag{B.44}$$

We can further simplify this as:

$$\theta_{t+1} = \theta_t + \eta\tau \log \left(\frac{e^{-[Qh^*]_a/\tau}}{e^{\theta_t}} \right), \tag{B.45}$$

which is nothing but:

$$\theta_{t+1} = \theta_t + \eta\tau \left(\frac{-Qh^*}{\tau} - \theta_t \right). \quad (\text{B.46})$$

Note that this is the Gradient Descent update on the strongly convex function $\frac{1}{2} \left\| \frac{-Qh^*}{\tau} - \theta \right\|^2$ with stepsize $\eta\tau$. Note that this update leads to the following convergence guarantees:

$$\|\theta_{t+1} - \theta^*\|^2 \leq (1 - \eta\tau)^2 \|\theta_t - \theta^*\|^2, \quad (\text{B.47})$$

where $\theta^* = \frac{-Qh^*}{\tau}$.

The analysis above shows that if one of the players is already at the NE strategy, the parameters of the second player converges to the NE at a linear rate. However, the original OGDA update for θ is given by

$$\theta_{t+1} = (1 - \eta\tau)\theta_t - \eta Qh^* + \eta Q(h^* - h_{\bar{\nu}_{t+1}}). \quad (\text{B.48})$$

Since $h_{\bar{\nu}_{t+1}}$ converges to h^* at a linear rate (from Lemma B.7), we expect the term

$$\varepsilon_t = \eta Q(h^* - h_{\bar{\nu}_{t+1}}), \quad (\text{B.49})$$

to be small, and goes to 0. This is formalized in what follows.

The OGDA update for θ can be re-written using ε_t as:

$$\theta_{t+1} = \theta_t + \eta\tau \left(\frac{-Qh^*}{\tau} - \theta_t \right) + \varepsilon_t. \quad (\text{B.50})$$

Once again, defining $\theta^* = \frac{-Qh^*}{\tau}$, we have:

$$\begin{aligned} \|\theta_{t+1} - \theta^*\|^2 &= \|\theta_t + \eta\tau(\theta^* - \theta_t) + \varepsilon_t - \theta^*\|^2 = \|(1 - \eta\tau)(\theta_t - \theta^*) + \varepsilon_t\|^2 \\ &= (1 - \eta\tau)^2 \|\theta_t - \theta^*\|^2 + 2(1 - \eta\tau)(\theta_t - \theta^*)^\top \varepsilon_t + \|\varepsilon_t\|^2 \\ &\leq^* (1 - \eta\tau)^2 \|\theta_t - \theta^*\|^2 + \eta\tau \|\theta_t - \theta^*\|^2 + \frac{1}{\eta\tau} (1 - \eta\tau)^2 \|\varepsilon_t\|^2 + \|\varepsilon_t\|^2 \\ &= (1 - \eta\tau + \eta^2\tau^2) \|\theta_t - \theta^*\|^2 + \left(1 + \frac{1}{\eta\tau} (1 - \eta\tau)^2\right) \|\varepsilon_t\|^2, \end{aligned} \quad (\text{B.51})$$

where *1 follows from Young's inequality.

Next, we analyze the error term $\|\varepsilon_t\|$. We have:

$$\|\varepsilon_t\|^2 = \|\eta Q(h^* - h_{\bar{\nu}_{t+1}})\|^2 \leq \eta^2 \|Q\|_\infty^2 \|(h^* - h_{\bar{\nu}_{t+1}})\|_1^2 \leq^* 2\eta^2 \|Q\|_\infty^2 (\text{KL}(h^* \| h_{\bar{\nu}_{t+1}})). \quad (\text{B.52})$$

Here *1 follows from Pinsker's Inequality. Now, writing the same inequality for ν , we have:

$$\begin{aligned} \|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 &\leq (1 - \eta\tau + \eta^2\tau^2) (\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) \\ &\quad + \left(1 + \frac{1}{\eta\tau} (1 - \eta\tau)^2\right) 2\eta^2 \|Q\|_\infty^2 \text{KL}(z^* \| \bar{z}_{t+1}) \\ &\leq (1 - \eta\tau + \eta^2\tau^2) (\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + \left(1 + \frac{1}{\eta\tau} (1 - \eta\tau)^2\right) 4\eta^2 C \|Q\|_\infty^2 (1 - \eta\tau)^t \text{KL}(z^* \| z_0). \end{aligned} \quad (\text{B.53})$$

Define:

$$C = \left(1 + \frac{1}{\eta\tau} (1 - \eta\tau)^2\right) 4\eta^2 \|Q\|_\infty^2 \text{KL}(z^* \| z_0). \quad (\text{B.54})$$

This gives us (Using Lemma B.7):

$$\|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 \leq (1 - \eta\tau + \eta^2\tau^2) (\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + C(1 - \eta\tau)^t. \quad (\text{B.55})$$

For $\eta\tau < 1/2$ we have:

$$\|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 \leq (1 - \eta\tau/2)(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + C(1 - \eta\tau)^t. \quad (\text{B.56})$$

Consider the Lyapunov function:

$$V_{t+1} = \|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 + \frac{2C}{\eta\tau}(1 - \eta\tau)^{t+1}. \quad (\text{B.57})$$

We have:

$$\begin{aligned} V_{t+1} &\leq (1 - \eta\tau/2)(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + C(1 - \eta\tau)^t + \frac{2C}{\eta\tau}(1 - \eta\tau)(1 - \eta\tau)^t \\ &= (1 - \eta\tau/2)(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2) + \frac{2C}{\eta\tau}(1 - \eta\tau)^t(1 - \eta\tau/2) \\ &= (1 - \eta\tau/2) \left(\|\theta_t - \theta^*\|^2 + \|\nu_t - \nu^*\|^2 + \frac{2C}{\eta\tau}(1 - \eta\tau)^t \right) \\ &= (1 - \eta\tau/2)V_t. \end{aligned} \quad (\text{B.58})$$

This shows linear convergence of the parameter θ to θ^* since:

$$\|\theta_{t+1} - \theta^*\|^2 + \|\nu_{t+1} - \nu^*\|^2 \leq V_{t+1} \leq (1 - \eta\tau/2)^t V_0. \quad (\text{B.59})$$

This completes the proof. \square

Note that proof of Corollary 3.5 follows from Remark 4 and the preceding discussion in [Cen et al. \(2021b\)](#).

C Matrix Games with Function Approximation

To handle games with excessively large action spaces, we resort to policy parameterization with function approximation. In particular, consider the following problem:

$$\min_{\theta \in \mathbb{R}^d} \max_{\nu \in \mathbb{R}^d} g_\theta^\top Q h_\nu, \quad (\text{C.1})$$

where g_θ and h_ν are both parameterized in a softmax way as in (2.2), with the linear function class $p_\theta(a) = \phi_a^\top \theta$ (also called *log-linear* policy in [Agarwal et al. \(2019\)](#)), where $\phi_a \in \mathbb{R}^d$ is a low-dimensional feature representation of the action (see [Branavan et al. \(2009\)](#); [Agarwal et al. \(2019\)](#)) (Note that usually $d < n$). We define:

$$\Phi = [\phi_1, \phi_2, \dots, \phi_n] \in \mathbb{R}^{d \times n}. \quad (\text{C.2})$$

Assumption C.1. Φ is a full rank matrix. In particular, assume that $\Phi = [M \mid 0]$, where $M \in \mathbb{R}^{d \times d}$ is an invertible $d \times d$ square matrix.

Note that the full-rankness of Φ is a standard assumption (see Assumption 6.2 in [Agarwal et al. \(2019\)](#)). It essentially requires the features to be the bases of some low-dimensional space. Furthermore, the results also extend to the case where the matrix Φ is of the form $[M \mid c\mathbb{1}]$ where $\mathbb{1}$ is the matrix of all 1s of appropriate dimension, and c is any constant. This particular structure of the feature matrix, though being restrictive, ensures that the constraint set of policies is convex, as shown next, otherwise the minimax theorem of $\min \max = \max \min$ might not hold, i.e., the Nash equilibrium for the parameterized game does not exist. Moreover, the assumption is also not as restrictive as it seems. For example, in applications of self-driving car and robotics, only a subset of actions (steering angles) is essential in controlling the agent, with other actions being insignificant/redundant. Our feature matrix encodes patterns like these. Moreover, as a first step studying policy optimization in multi-agent learning with function approximation, we start with this simpler setting. Extending the ideas to more general FA settings is an interesting direction worth exploring.

Remark C.2. Note that the results in this section are presented for the case where the feature matrix is identical for both players purely for simplification of notation. The results continue to hold for the case with asymmetric features as well (as long as the feature matrix also satisfies Assumption C.1).

As motivated in the previous section, we study the following regularized problem in order to solve (C.1) efficiently:

$$\min_{\theta \in \mathbb{R}^d} \max_{\nu \in \mathbb{R}^d} g_\theta^\top Q h_\nu - \tau \mathcal{H}(g_\theta) + \tau \mathcal{H}(h_\nu), \quad (\text{C.3})$$

where \mathcal{H} denotes the entropy function and $\tau > 0$ is the regularization parameter. Note that this problem can still be nonconvex-nonconcave in general, given the example in §3 as a special case.

We define the solution to this problem next, the Nash equilibrium in the parameterized policy classes, i.e., in-class NE.

Definition C.3 (ϵ -in-class Nash equilibrium). The policy parameter $(\tilde{\theta}, \tilde{\nu})$ is an ϵ -Nash equilibrium of the matrix game with function approximation (or ϵ -in-class NE), if it satisfies that for all $i \in [N]$,

$$g_{\tilde{\theta}}^\top Q h_\nu - \epsilon \leq g_{\tilde{\theta}}^\top Q h_{\tilde{\nu}} \leq g_{\tilde{\theta}}^\top Q h_{\tilde{\nu}} + \epsilon, \quad \forall \theta, \nu \in \mathbb{R}^d. \quad (\text{C.4})$$

Furthermore, when $\epsilon = 0$, we refer to it as the *in-class Nash Equilibrium*.

C.1 Equivalent problem characterization

In this subsection, we study the regularized problem (C.3) under a log-linear parametrization and find an equivalent problem in the tabular case.

First, in the following lemma, we characterize the set of distributions covered by this parametrization, and study the equivalent problem in the space of probability vectors.

Lemma C.4. Under Assumption C.1, the log-linear parametrization in Equation (2.2) covers all distributions in the following *convex* set:

$$\tilde{\Delta} = \{\mu : \mu \in \Delta, \mu_{d+1} = \mu_{d+2} = \dots = \mu_n\}, \quad (\text{C.5})$$

and Problem (C.3) is equivalent to

$$\min_{g_\theta \in \tilde{\Delta}} \max_{h_\nu \in \tilde{\Delta}} g_\theta^\top Q h_\nu - \tau \mathcal{H}(g_\theta) + \tau \mathcal{H}(h_\nu). \quad (\text{C.6})$$

Lemma C.4 characterizes the set of distributions which can be represented by log-linear parametrization. Therefore, when we try to solve the matrix game with such function approximation, the best we can hope for is to find an equilibrium within the set $\tilde{\Delta}$.

Next, we characterize the Nash equilibrium of the regularized Problem (C.3) in the function approximation setting, and show its equivalence to another problem in the tabular softmax setting.

Theorem C.5. An in-class Nash equilibrium (θ^*, ν^*) of Problem (C.3) under the function approximation setting exists, and any such in-class NE satisfies:

$$g_{\theta^*}(a) = \frac{e^{-\frac{[\Psi^\top Q \Psi h_{\nu^*}]_a}{\tau}}}{\sum_{a'} e^{-\frac{[\Psi^\top Q \Psi h_{\nu^*}]_{a'}}{\tau}}}, \quad h_{\nu^*}(a) = \frac{e^{\frac{[\Psi^\top Q^\top \Psi g_{\theta^*}]_a}{\tau}}}{\sum_{a'} e^{\frac{[\Psi^\top Q^\top \Psi g_{\theta^*}]_{a'}}{\tau}}},$$

where $\Psi \in \mathbb{R}^{n \times n}$ is defined as:

$$\Psi = \begin{pmatrix} I_d & \mathbf{0} & & \\ & \frac{1}{n-d} & \dots & \frac{1}{n-d} \\ \mathbf{0} & \dots & \dots & \dots \\ & \frac{1}{n-d} & \dots & \frac{1}{n-d} \end{pmatrix}. \quad (\text{C.7})$$

Furthermore, Problem (C.3) is equivalent to⁶:

$$\min_{g_\theta \in \tilde{\Delta}} \max_{h_\nu \in \tilde{\Delta}} g_\theta^\top \Psi^\top Q \Psi h_\nu - \tau \mathcal{H}(g_\theta) + \tau \mathcal{H}(h_\nu). \quad (\text{C.8})$$

Note that the matrix Ψ defined in Theorem C.5 is the invariant matrix for the set $\tilde{\Delta}$, i.e., $\Psi \mu = \mu, \forall \mu \in \tilde{\Delta}$.

⁶By equivalent, we mean that the two problems have the same value at the NE.

Algorithm 2 Optimistic NPG (Function Approximation)

Initialize: $\theta_0 = 0$ and $\nu_0 = 0$.

for $t = 1, 2, \dots$ **do**

$$\bar{\theta}_{t+1} = (1 - \eta\tau)\theta_t - \eta[(M^\top)^{-1} | 0] \tilde{P}Qh_{\bar{\nu}_t}$$

$$\bar{\nu}_{t+1} = (1 - \eta\tau)\nu_t + \eta[(M^\top)^{-1} | 0] \tilde{P}Q^\top g_{\bar{\theta}_t}$$

$$\theta_{t+1} = (1 - \eta\tau)\theta_t - \eta[(M^\top)^{-1} | 0] \tilde{P}Qh_{\bar{\nu}_{t+1}}$$

$$\nu_{t+1} = (1 - \eta\tau)\nu_t + \eta[(M^\top)^{-1} | 0] \tilde{P}Q^\top g_{\bar{\theta}_{t+1}}$$

end for

C.2 Optimistic NPG algorithm

From Section 3, we see that the optimistic version of NPG leads to faster convergence (of both policy and parameters). This motivates us to focus on the optimistic version of the methods. In this subsection, (and the ones that follow), we focus on optimistic methods instead of their non-optimistic counterparts. Note that similar to Section 3, we can derive convergence rates for the non-optimistic versions as well, which would be slower than the corresponding optimistic versions.

As Theorem C.5 characterizes the solution to the function approximation setting to that of the tabular softmax setting, we modify the algorithm for function approximation setting as follows:

$$\bar{\theta}_{t+1} = (1 - \eta\tau)\theta_t - \eta[(M^\top)^{-1} | 0] \tilde{P}Qh_{\bar{\nu}_t}, \quad \theta_{t+1} = (1 - \eta\tau)\theta_t - \eta[(M^\top)^{-1} | 0] \tilde{P}Qh_{\bar{\nu}_{t+1}}, \quad (\text{C.9})$$

and a similar update for ν to reach the solution of the regularized problem under a log-linear parametrization. Here, the matrix \tilde{P} is defined as

$$\tilde{P} = \begin{pmatrix} I_d & \frac{-1}{n-d} & \cdots & \frac{-1}{n-d} \\ & \cdots & \cdots & \cdots \\ & \frac{-1}{n-d} & \cdots & \frac{-1}{n-d} \\ \mathbf{0} & & \mathbf{0} & \end{pmatrix}.$$

The additional term involving the inverse of the feature matrix arises due to the nature of the log-linear function approximation. We make this formal in the following proposition.

Proposition C.6. Consider Algorithm 2 used to solve Problem (C.3) under the log-linear parametrization in Equation (2.2) under Assumption C.1. Then the iterates of Algorithm 2 have the same guarantees provided in Theorem 3.4. Here, the NE parameter value to which the algorithm converges to is given by:

$$\theta^* = \frac{-[(M^\top)^{-1} | 0] \tilde{P}Qh_{\nu^*}}{\tau}, \quad \nu^* = \frac{[(M^\top)^{-1} | 0] \tilde{P}Q^\top g_{\theta^*}}{\tau}.$$

Next, we show how the optimistic NPG algorithm solves the original matrix game without regularization.

Corollary C.7. If we run Algorithm 2 for time $T = \mathcal{O}\left(\frac{\log n}{\eta\epsilon} \log\left(\frac{1}{\epsilon}\right)\right)$ and set the regularization parameter $\tau = \epsilon/(8 \log n)$, we have that the output (θ_T, ν_T) is an ϵ -in-class NE (Definition C.3) of the unregularized Problem (C.1).

D Missing Details and Proofs in §C

Remark D.1. We note that all results present in this section also follow for the case where the cardinality of the action spaces for both players are unequal. However, we stick to the case where the number of actions is the same for both players for ease of exposition.

D.1 Proof of Lemma C.4

The first part of the lemma describes the set of distributions in the m -dimensional simplex covered by this parametrization. Since the set of distributions covered by the log-linear parametrization would be the same for all invertible M^7 , for

⁷To see this, consider $\tilde{\theta} = M\theta$. Since M is invertible, there is a 1-1 correspondence between $\tilde{\theta}$ and θ . The distribution parametrized by θ under the function approximation matrix M , is the same as the distribution parametrized by $\tilde{\theta}$ under the function approximation matrix I . Therefore it is enough to consider the special case of $M = I$.

simplicity, we study the case where $M = I$. Note that this would imply:

$$\begin{aligned} g_\theta(a) &\propto e^{\theta(a)} & \forall a \in \{1, 2, \dots, d\}, \\ g_\theta(a) &\propto 1 & \forall a \in \{d+1, d+2, \dots, n\}. \end{aligned} \quad (\text{D.1})$$

Similarly for h_ν . Therefore, according to this parametrization the first d elements can be chosen freely, and the rest $n - d$ parameters have to be equal. In other words, this parametrization covers the following set of distributions:

$$\tilde{\Delta} = \{\mu : \mu \in \Delta, \mu_{d+1} = \mu_{d+2} = \dots = \mu_n\}, \quad (\text{D.2})$$

which is a closed convex subset of the n -dimensional simplex. To see that any element of $\tilde{\Delta}$ can be represented by the log-linear parametrization, we can take the parameters $\theta(a) = \log \mu(a)$ for $a = 1, 2, \dots, d$ ⁸. This would be a valid parametrization under the log-linear function approximation setting, and therefore all elements of $\tilde{\Delta}$ can be represented in this manner. Therefore, since the two sets are equivalent, we can rewrite the problem in terms of the policy vectors lying in the constraint set $\tilde{\Delta}$ which completes the proof of the lemma.

D.2 Proof of Theorem C.5

For the regularized game (C.6), let player 2 play the NE strategy h_ν^* . Then player 1's optimization problem is given by:

$$\min_{g_\theta \in \tilde{\Delta}} g_\theta^\top Q h_\nu^* - \tau \mathcal{H}(g_\theta). \quad (\text{D.3})$$

Define the following Lagrange multipliers (and associated constraints):

$$\begin{aligned} \lambda &: \sum_{a=1}^n g_\theta(a) = 1, \\ \beta_1 &: g_\theta(d+1) = g_\theta(d+2), \\ \beta_2 &: g_\theta(d+2) = g_\theta(d+3), \\ &\dots \\ \beta_{n-d-1} &: g_\theta(n-1) = g_\theta(n). \end{aligned} \quad (\text{D.4})$$

Therefore, for the optimal Lagrange multipliers, taking the first-order optimality conditions (with respect to the variable g_θ), we have:

$$\begin{aligned} [Qh_\nu^*]_a + \tau(\log g_\theta(a) + 1) + \lambda &= 0 & \forall a \in \{1, 2, \dots, d\}, \\ [Qh_\nu^*]_a + \tau(\log g_\theta(a) + 1) + \lambda + \beta_{a-d} - \beta_{a-d-1} &= 0 & \forall a \in \{d+1, d+2, \dots, n\}, \end{aligned} \quad (\text{D.5})$$

where β_0 and β_{n-d} are defined to be equal to 0. This gives us:

$$g_\theta^*(a) \propto e^{-[Qh_\nu^*]_a/\tau} \quad \forall a \in \{1, 2, \dots, d\}. \quad (\text{D.6})$$

For actions with indices $a > d$, the equality constraints give us:

$$[Qh_\nu^*]_a + \beta_{a-d} - \beta_{a-d-1} = C \quad \forall a \in \{d+1, d+2, \dots, n\}, \quad (\text{D.7})$$

for some constant C . On solving these equations for C , we have:

$$C = \frac{1}{n-d} \sum_{a=d+1}^n [Qh_\nu^*]_a, \quad (\text{D.8})$$

which gives us

$$g_\theta^*(a) \propto e^{-C/\tau} \quad \forall a \in \{d+1, \dots, n\}. \quad (\text{D.9})$$

⁸Note that if $\mu(a) = 0$, the corresponding parameter would be $-\infty$.

Now, consider the symmetric matrix $\Psi \in \mathbb{R}^{n \times n}$ defined as:

$$\Psi = \begin{pmatrix} I_d & \mathbf{0} & & \\ & \frac{1}{n-d} & \cdots & \frac{1}{n-d} \\ & \mathbf{0} & \cdots & \\ & & \frac{1}{n-d} & \cdots & \frac{1}{n-d} \end{pmatrix}. \quad (\text{D.10})$$

Note that

$$\begin{aligned} [\Psi^\top Q h_\nu^*]_a &= [Q h_\nu^*]_a \quad \forall a \in \{1, 2, \dots, d\}, \\ [\Psi^\top Q h_\nu^*]_a &= C \quad \forall a \in \{d+1, d+2, \dots, n\}. \end{aligned} \quad (\text{D.11})$$

Therefore, we can succinctly write the optimal distribution as:

$$g_\theta^*(a) \propto e^{-[\Psi^\top Q h_\nu^*]_a / \tau} \quad \forall a \in \{1, 2, \dots, n\}. \quad (\text{D.12})$$

Also, we have $\Psi \mu = \mu$, $\forall \mu \in \tilde{\Delta}$.

Now, since h_ν^* is the optimal distribution for Player 2, we must have $h_\nu^* \in \tilde{\Delta}$ which implies:

$$\Psi h_\nu^* = h_\nu^*. \quad (\text{D.13})$$

Doing a similar calculation for g_ν^* , we have that the NE satisfy:

$$g_\theta^*(a) = \frac{e^{-[\Psi^\top Q \Psi h_\nu^*]_a / \tau}}{\sum_{a'} e^{-[\Psi^\top Q \Psi h_\nu^*]_{a'} / \tau}}, \quad h_\nu^*(a) = \frac{e^{[\Psi^\top Q^\top \Psi g_\theta^*]_a / \tau}}{\sum_{a'} e^{[\Psi^\top Q^\top \Psi g_\theta^*]_{a'} / \tau}}. \quad (\text{D.14})$$

Now, from this characterization of the NE, we see that these solutions are also the same as those of the following problem:

$$\min_{g \in \Delta} \max_{h \in \Delta} g^\top \Psi^\top Q \Psi h - \tau \mathcal{H}(g) + \tau \mathcal{H}(h). \quad (\text{D.15})$$

Note that in the analysis above we have found a solution in the (relative) interior of the constraint set. For example, if for some action a , we have $f_\theta(a) = 0$, then the term $\log f_\theta(a)$ is not defined and the Lagrangian would be different. However, since this is a strongly convex strongly concave minimax problem over a convex compact set, there is a unique solution (see [Facchinei and Pang \(2007\)](#)). As we have already found a solution in the interior of the constraint set, by the argument above we have that this is the unique solution. This allows us to solve the first-order KKT optimality conditions (see [Boyd and Vandenberghe \(2004\)](#)) to find the solution. Also, since all terms satisfy $f_\theta(a) > 0$ (similarly for $g_\nu(a)$) we do not need to explicitly write down the Lagrange multipliers for the non-negativity constraint. This completes the proof. \square

D.3 Proof of Proposition C.6

The algorithm is similar to the one proposed for the tabular case. However, since the exponents are elements of $\Phi^\top \theta$ instead of just θ as was in the case of tabular softmax, we have the updates modified as well (we write the updates here for the combined update instead of the two step update for ease of presentation):

$$\begin{aligned} \Phi^\top \theta_{t+1} &= (1 - \eta\tau) \Phi^\top \theta_t - (2 - \eta\tau) \eta \Psi^\top \tilde{P} Q \Psi h_{\nu_t} + (1 - \eta\tau) \eta \Psi^\top \tilde{P} Q \Psi h_{\nu_{t-1}}, \\ \Phi^\top \nu_{t+1} &= (1 - \eta\tau) \Phi^\top \nu_t + (2 - \eta\tau) \eta \Psi^\top \tilde{P} Q^\top \Psi g_{\theta_t} - (1 - \eta\tau) \eta \Psi^\top \tilde{P} Q^\top \Psi g_{\theta_{t-1}}, \end{aligned} \quad (\text{D.16})$$

where Φ is a full-rank feature matrix. Note that the additional matrix \tilde{P} is to ensure that the probability vector at each step of the algorithm satisfies the function approximation constraint⁹. Since Φ is full-rank, we can explicitly write this update for θ and ν as follows:

$$\begin{aligned} \theta_{t+1} &= [\Phi \Phi^\top]^{-1} \Phi \left((1 - \eta\tau) \Phi^\top \theta_t - (2 - \eta\tau) \eta \Psi^\top \tilde{P} Q \Psi h_{\nu_t} + (1 - \eta\tau) \eta \Psi^\top \tilde{P} Q \Psi h_{\nu_{t-1}} \right), \\ \nu_{t+1} &= [\Phi \Phi^\top]^{-1} \Phi \left((1 - \eta\tau) \Phi^\top \nu_t + (2 - \eta\tau) \eta \Psi^\top \tilde{P} Q^\top \Psi g_{\theta_t} - (1 - \eta\tau) \eta \Psi^\top \tilde{P} Q^\top \Psi g_{\theta_{t-1}} \right), \end{aligned}$$

⁹This is based on the fact that for a probability vector $x_i \propto e^{y_i}$, we will also have $x_i \propto e^{y_i - k}$ for any constant k independent of the index i .

which can be simplified to:

$$\begin{aligned}\theta_{t+1} &= (1 - \eta\tau)\theta_t - (2 - \eta\tau)\eta[\Phi\Phi^\top]^{-1}\Phi\Psi^\top\tilde{P}Q\Psi h_{\nu_t} + (1 - \eta\tau)\eta[\Phi\Phi^\top]^{-1}\Phi\Psi^\top\tilde{P}Q\Psi h_{\nu_{t-1}}, \\ \nu_{t+1} &= (1 - \eta\tau)\nu_t + (2 - \eta\tau)\eta[\Phi\Phi^\top]^{-1}\Phi\Psi^\top\tilde{P}Q^\top\Psi g_{\theta_t} - (1 - \eta\tau)\eta[\Phi\Phi^\top]^{-1}\Phi\Psi^\top\tilde{P}Q^\top\Psi g_{\theta_{t-1}}.\end{aligned}\quad (\text{D.17})$$

Note that for $\Phi = [M \mid 0]$, we have

$$[\Phi\Phi^\top]^{-1}\Phi = ([M|0][M|0]^\top)^{-1}[M|0] = (MM^\top)^{-1}[M|0] = [(M^\top)^{-1}|0]. \quad (\text{D.18})$$

From the previous discussion, since all the terms g_{θ_t} and h_{ν_t} lie in the set $\tilde{\Delta}$, we have $\Psi g_{\theta_t} = g_{\theta_t}$ and $\Psi h_{\nu_t} = h_{\nu_t}$. Also, we have:

$$[\Phi\Phi^\top]^{-1}\Phi\Psi^\top = [(M^\top)^{-1}|0]\Psi^\top = [(M^\top)^{-1}|0], \quad (\text{D.19})$$

from the structure of Ψ . Therefore, the update can be written as:

$$\begin{aligned}\theta_{t+1} &= (1 - \eta\tau)\theta_t - (2 - \eta\tau)\eta[(M^\top)^{-1}|0]\tilde{P}Qh_{\nu_t} + (1 - \eta\tau)\eta[(M^\top)^{-1}|0]\tilde{P}Qh_{\nu_{t-1}}, \\ \nu_{t+1} &= (1 - \eta\tau)\nu_t + (2 - \eta\tau)\eta[(M^\top)^{-1}|0]\tilde{P}Q^\top g_{\theta_t} - (1 - \eta\tau)\eta[(M^\top)^{-1}|0]\tilde{P}Q^\top g_{\theta_{t-1}},\end{aligned}\quad (\text{D.20})$$

which completes the proof. \square

D.4 Simulation to show convergence of ONPG in the function approximation setting

In this section, we show the performance of the ONPG algorithm under Function Approximation (Algorithm 2). The policies have a log linear parametrization, with the feature matrix $\Phi = [I \mid 0] \in \mathbb{R}^{10 \times 100}$, and the cost matrix Q is chosen to be a random matrix of dimension 100×100 .




Figure 1: Behavior of Algorithm 2 in the matrix game under the log-linear function approximation setting.

E Missing Definitions and Proofs in §4

Remark E.1. We note that all results presented in this section also follow for the case where the number of possible actions for each player can be different. However, we stick to the case where the number of action is the same for both players for ease of exposition. Note that the actual action spaces need not be identical, but only their cardinalities.

Algorithm 3 Optimistic NPG for monotone games

Initialize: $\theta_i^0 = 0$ for all players i .
for $t = 1, 2, \dots$ **do**
 $\tilde{\theta}_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_{\theta_i}} f_i(g_{\tilde{\theta}_i^t}, g_{\tilde{\theta}_{-i}^t}) \quad \forall i \in [N]$.
 $\theta_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_{\theta_i}} f_i(g_{\tilde{\theta}_i^{t+1}}, g_{\tilde{\theta}_{-i}^{t+1}}) \quad \forall i \in [N]$.
end for

Definition E.2 (In-class Nash equilibrium for a monotone game). The policy parameter $\theta^* = [\theta_1^*, \theta_2^*, \dots, \theta_N^*]$ is an NE under function approximation, i.e., (in-class NE) of the monotone game, if it satisfies that for all $i \in [N]$,

$$f_i(g_{\theta_i^*}, g_{\theta_{-i}^*}) \leq f_i(g_{\theta_i}, g_{\theta_{-i}^*}), \quad \forall \theta_i \in \mathbb{R}^d. \quad (\text{E.1})$$

Note that if we are in the tabular setting, we will have $d = n$.

Definition E.3 (ϵ -in-class Nash equilibrium for a monotone game). The policy parameter $(\tilde{\theta}_1, \dots, \tilde{\theta}_N)$ is an ϵ -Nash equilibrium under function approximation (or *in-class* ϵ -NE) of the monotone game if it satisfies that for all $i \in [N]$,

$$f_i(g_{\tilde{\theta}_i}, g_{\tilde{\theta}_{-i}}) - \epsilon \leq f_i(g_{\theta_i}, g_{\tilde{\theta}_{-i}}), \quad \forall \theta_i \in \mathbb{R}^d. \quad (\text{E.2})$$

Note that if we are in the tabular setting, we will have $d = n$.

We will also use the notation $f_i^r(g_i, g_{-i}) = f_i(g_i, g_{-i}) - \tau\mathcal{H}(g_i)$.

E.1 Proof of Lemma 4.2

We can follow the analysis for a two player game from [Mertikopoulos and Sandholm \(2016\)](#) to write down the solution form for the N-player monotone setting. We let g_i and g_{-i} denote g_{θ_i} and $g_{\theta_{-i}}$ respectively. First, note that the solutions in the policy space exists for the unregularized game, since we are solving a monotone VI over a convex compact set, and this solution is unique if we regularize the game, since in this case we are solving a stringly monotone VI over a convex compact set. See [Facchinei and Pang \(2007\)](#).

Consider player i 's optimization problem when other players play the equilibrium strategies:

$$\min_{g_i \in \Delta} f_i(g_i, g_{-i}^*) - \tau\mathcal{H}(g_i). \quad (\text{E.3})$$

Since we are in the monotone setting with a strongly convex regularizer, the first order Karush–Kuhn–Tucker (KKT) conditions are both necessary and sufficient. The first order KKT conditions are:

$$[\nabla_{g_i} f_i(g_i^*, g_{-i}^*)]_a + \tau(\log g_i^* + 1) + \lambda = 0,$$

where λ is the Lagrange multiplier corresponding to the simplex constraint. This implies:

$$g_i^*(a) \propto e^{\frac{-[\nabla_{g_i} f_i(g_i^*, g_{-i}^*)]_a}{\tau}}, \quad (\text{E.4})$$

which shows the NE in the policy space. Now, to complete the proof of the lemma, we need to find parameters θ^* which leads to this distribution. This can be easily seen by setting $\theta_i^* = \frac{-[\nabla_{g_i} f_i(g_i^*, g_{-i}^*)]}{\tau}$, thereby completing the proof of the lemma. \square

E.2 (Optimistic) NPG for monotone games

As noted in [§A.2](#), we have that the Fisher Information matrix $F_\theta(\theta) = \nabla_\theta g_\theta$. Therefore, the NPG update for player i can be simplified as:

$$\begin{aligned}
 \theta_i^{t+1} &= \theta_i^t - \eta \cdot F_\theta^\dagger(\theta_i^t) \cdot \frac{\partial f_i^r(g_{\theta_i^t}, g_{\theta_{-i}^t})}{\partial \theta} = \theta_t - \eta \frac{\partial f_i^r(g_{\theta_i^t}, g_{\theta_{-i}^t})}{\partial g_{\theta_i}} \\
 &= \theta_t - \eta \left(\nabla_{g_{\theta_i}} f_i(g_{\theta_i^t}, g_{\theta_{-i}^t}) + \tau(\mathbf{1} + \log g_{\theta_i^t}) \right). \quad (\text{E.5})
 \end{aligned}$$

This can be simplified as:

$$\theta_i^{t+1}(a) = (1 - \eta\tau)\theta_i^t(a) - \eta[\nabla_{g_{\theta_i}} f_i(g_{\theta_i^t}, g_{\theta_{-i}^t})]_a + \eta\tau(\log Z_{\theta_i^t} - 1), \quad (\text{E.6})$$

where $Z_{\theta_i^t} = \sum_{a' \in \mathcal{A}} e^{\theta_i^t(a')}$. Note that this update will have the same pitfall of parameter divergence as the NPG update for the matrix game, since a matrix game is a special case of the monotone game. Therefore, we propose the following modified version of NPG, as done for the matrix game:

$$\theta_i^{t+1}(a) = (1 - \eta\tau)\theta_i^t(a) - \eta[\nabla_{g_{\theta_i}} f_i(g_{\theta_i^t}, g_{\theta_{-i}^t})]_a. \quad (\text{E.7})$$

This leads to the modified NPG dynamics for the monotone game. Now, similar to the matrix game, we analyze the optimistic version of this algorithm with updates:

$$\bar{\theta}_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_{\theta_i}} f_i(g_{\bar{\theta}_i^t}, g_{\bar{\theta}_{-i}^t}), \quad \theta_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_{\theta_i}} f_i(g_{\bar{\theta}_i^{t+1}}, g_{\bar{\theta}_{-i}^{t+1}}),$$

in §4.

E.3 Proof of Theorem 4.3

We prove the following Lemma first:

Lemma E.4. For any $z = (g_{\theta_i}, g_{\theta_{-i}}) \in \Delta^N$, consider an update of the form:

$$\theta_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_{\theta_i}} f_i(g_{\theta_i}, g_{\theta_{-i}}) \quad \forall i \in [N]. \quad (\text{E.8})$$

We have:

$$\langle \log z_{t+1} - (1 - \eta\tau) \log z_t - \eta\tau \log z^*, z - z^* \rangle \leq 0, \quad (\text{E.9})$$

where $z^* = (g_i^*)_{i=1}^N$.

Proof. In the proof below, we define $g_i := g_{\theta_i}$ and $g_i^t := g_{\theta_i^t}$. From the update sequence in Equation (E.8), we have

$$\log g_i^{t+1} = (1 - \eta\tau) \log g_i^t - \eta\nabla_{g_i} f_i(g_i, g_{-i}) + c \cdot \mathbf{1}, \quad (\text{E.10})$$

where c is the normalization constant. This implies:

$$\begin{aligned} \langle \log g_i^{t+1} - (1 - \eta\tau) \log g_i^t, g_i - g_i^* \rangle &= \langle -\eta\nabla_{g_i} f_i(g_i, g_{-i}) + c \cdot \mathbf{1}, g_i - g_i^* \rangle \\ &= \langle -\eta\nabla_{g_i} f_i(g_i, g_{-i}), g_i - g_i^* \rangle. \end{aligned} \quad (\text{E.11})$$

Note that $\langle c \cdot \mathbf{1}, g_i - g_i^* \rangle = 0$, since $g_i, g_i^* \in \Delta$. Since this is true for all players i , we have:

$$\langle \log z_{t+1} - (1 - \eta\tau) \log z_t, z - z^* \rangle = -\eta \sum_i \langle \nabla_{g_i} f_i(g_i, g_{-i}), g_i - g_i^* \rangle = -\eta \langle F(z), z - z^* \rangle. \quad (\text{E.12})$$

Now, from the properties of the NE, we have:

$$\eta\tau \log g_i^* = -\eta\nabla_{g_i} f_i(g_i^*, g_{-i}^*) + c \cdot \mathbf{1}, \quad (\text{E.13})$$

which gives us:

$$\langle \eta\tau \log g_i^*, g_i - g_i^* \rangle = \langle -\eta\nabla_{g_i} f_i(g_i^*, g_{-i}^*), g_i - g_i^* \rangle. \quad (\text{E.14})$$

Since this is true for all players i , we have

$$\langle \eta\tau \log z^*, z - z^* \rangle = -\eta \langle F(z^*), z - z^* \rangle. \quad (\text{E.15})$$

From Equations (E.12) and (E.15), we have:

$$\langle \log z_{t+1} - (1 - \eta\tau) \log z_t - \eta\tau \log z^*, z - z^* \rangle = -\eta \langle F(z) - F(z^*), z - z^* \rangle \leq 0, \quad (\text{E.16})$$

where the last step uses the monotonicity assumption of F . This completes the proof of the lemma. \square

Lemma E.5. For updates of the form:

$$\theta_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_{\theta_i}} f_i(g_{\bar{\theta}_i^{t+1}}, g_{\bar{\theta}_{-i}^{t+1}}) \quad \forall i \in [N], \quad (\text{E.17})$$

we have

$$(1 - \eta\tau)\text{KL}(z^* \| z_t) \geq (1 - \eta\tau)\text{KL}(\bar{z}_{t+1} \| z_t) + \eta\tau\text{KL}(\bar{z}_{t+1} \| z^*) + \text{KL}(z_{t+1} \| \bar{z}_{t+1}) \\ - \langle \log \bar{z}_{t+1} - \log z_{t+1}, \bar{z}_{t+1} - z_{t+1} \rangle + \text{KL}(z^* \| z_{t+1}), \quad (\text{E.18})$$

and

$$\text{KL}(z^* \| \bar{z}_{t+1}) = \text{KL}(z^* \| z_{t+1}) - \text{KL}(\bar{z}_{t+1} \| z_{t+1}) - \langle z^* - \bar{z}_{t+1}, \log \bar{z}_{t+1} - \log z_{t+1} \rangle. \quad (\text{E.19})$$

Proof. From the definition of KL divergence we have:

$$-\langle \log z_{t+1} - (1 - \eta\tau)\log z_t - \eta\tau \log z^*, z^* \rangle = -(1 - \eta\tau)\text{KL}(z^* \| z_t) + \text{KL}(z^* \| z_{t+1}). \quad (\text{E.20})$$

Next, note that:

$$\langle \log z_{t+1} - (1 - \eta\tau)\log z_t - \eta\tau \log z^*, \bar{z}_{t+1} \rangle \\ = \langle \log \bar{z}_{t+1} - (1 - \eta\tau)\log z_t - \eta\tau \log z^*, \bar{z}_{t+1} \rangle + \langle \log \bar{z}_{t+1} - \log z_{t+1}, z_{t+1} \rangle \\ - \langle \log \bar{z}_{t+1} - \log z_{t+1}, \bar{z}_{t+1} - z_{t+1} \rangle \\ = (1 - \eta\tau)\text{KL}(\bar{z}_{t+1} \| z_t) + \eta\tau\text{KL}(\bar{z}_{t+1} \| z^*) + \text{KL}(z_{t+1} \| \bar{z}_{t+1}) \\ - \langle \log \bar{z}_{t+1} - \log z_{t+1}, \bar{z}_{t+1} - z_{t+1} \rangle. \quad (\text{E.21})$$

Now, substituting $z = \bar{z}_{t+1}$ in Lemma E.4 we have:

$$\langle \log z_{t+1} - (1 - \eta\tau)\log z_t - \eta\tau \log z^*, \bar{z}_{t+1} - z^* \rangle \leq 0. \quad (\text{E.22})$$

Substituting Equations (E.20) and (E.21) in (E.22), we get the Inequality (E.18).

Inequality (E.19) follows from the properties of KL divergence. \square

From the ONPG updates in Algorithm (3), we have:

$$\log \bar{g}_i^{t+1} - \log g_i^{t+1} = -\eta (\nabla_{g_i} f_i(\bar{g}_i^t, \bar{g}_{-i}^t) - \nabla_{g_i} f_i(\bar{g}_i^{t+1}, \bar{g}_{-i}^{t+1})) + c \cdot \mathbf{1}. \quad (\text{E.23})$$

This implies:

$$\langle \log \bar{g}_i^{t+1} - \log g_i^{t+1}, \bar{g}_i^{t+1} - g_i^{t+1} \rangle \\ = -\eta \langle \nabla_{g_i} f_i(\bar{g}_i^t, \bar{g}_{-i}^t) - \nabla_{g_i} f_i(g_i^t, g_{-i}^t) + \nabla_{g_i} f_i(g_i^t, g_{-i}^t) - \nabla_{g_i} f_i(\bar{g}_i^{t+1}, \bar{g}_{-i}^{t+1}), \bar{g}_i^{t+1} - g_i^{t+1} \rangle \\ \leq^{*1} \eta (\|\nabla_{g_i} f_i(\bar{g}_i^t, \bar{g}_{-i}^t) - \nabla_{g_i} f_i(g_i^t, g_{-i}^t)\| + \|\nabla_{g_i} f_i(g_i^t, g_{-i}^t) - \nabla_{g_i} f_i(\bar{g}_i^{t+1}, \bar{g}_{-i}^{t+1})\|) \|\bar{g}_i^{t+1} - g_i^{t+1}\| \\ \leq^{*2} \eta L (\|\bar{g}_i^t - g_i^t\| + \|\bar{g}_{-i}^t - g_{-i}^t\| + \|\bar{g}_i^{t+1} - g_i^t\| + \|\bar{g}_{-i}^{t+1} - g_{-i}^t\|) \|\bar{g}_i^{t+1} - g_i^{t+1}\| \\ \leq^{*3} \frac{1}{2} \eta L (\|\bar{g}_i^t - g_i^t\|^2 + \|\bar{g}_{-i}^t - g_{-i}^t\|^2 + \|\bar{g}_i^{t+1} - g_i^t\|^2 + \|\bar{g}_{-i}^{t+1} - g_{-i}^t\|^2 + 4\|\bar{g}_i^{t+1} - g_i^{t+1}\|^2), \quad (\text{E.24})$$

where (*1) follows from the fact that $a^\top b \leq \|a\| \|b\|$, (*2) follows from Assumption 4.1, and (*3) follows from $x \cdot y \leq \frac{1}{2}(x^2 + y^2)$. Since this is true for all players i , we have:

$$\langle \log \bar{z}_{t+1} - \log z_{t+1}, \bar{z}_{t+1} - z_{t+1} \rangle \\ \leq \frac{1}{2} \eta L \sum_{i=1}^N (\|\bar{g}_i^t - g_i^t\|^2 + \|\bar{g}_{-i}^t - g_{-i}^t\|^2 + \|\bar{g}_i^{t+1} - g_i^t\|^2 + \|\bar{g}_{-i}^{t+1} - g_{-i}^t\|^2 + 4\|\bar{g}_i^{t+1} - g_i^{t+1}\|^2) \\ \leq \frac{1}{2} \eta L (N\|\bar{z}_t - z_t\|^2 + N\|\bar{z}_{t+1} - z_t\|^2 + 4\|\bar{z}_{t+1} - z_{t+1}\|^2) \\ \leq^{*1} \eta L (N\text{KL}(z_t \| \bar{z}_t) + N\text{KL}(\bar{z}_{t+1} \| z_t) + 4\text{KL}(z_{t+1} \| \bar{z}_{t+1})), \quad (\text{E.25})$$

where (*1) follows from Pinsker's Inequality and the fact that the l_1 norm is an upper bound for the l_2 norm. Substituting this in Equation (E.18), we have

$$\begin{aligned} \text{KL}(z^* \| z_{t+1}) &\leq (1 - \eta\tau)\text{KL}(z^* \| z_t) - (1 - \eta\tau - N\eta L)\text{KL}(\bar{z}_{t+1} \| z_t) - \eta\tau\text{KL}(\bar{z}_{t+1} \| z^*) \\ &\quad - (1 - 4\eta L)\text{KL}(z_{t+1} \| \bar{z}_{t+1}) + N\eta L\text{KL}(z_t \| \bar{z}_t). \end{aligned} \quad (\text{E.26})$$

For $\eta < \frac{1}{2(N+4)L+2\tau}$, we have:

$$N\eta L \leq (1 - \eta\tau)(1 - 4\eta L). \quad (\text{E.27})$$

This gives us:

$$\begin{aligned} \text{KL}(z^* \| z_{t+1}) + (1 - 4\eta L)\text{KL}(z_{t+1} \| \bar{z}_{t+1}) &\leq (1 - \eta\tau)\text{KL}(z^* \| z_t) + N\eta L\text{KL}(z_t \| \bar{z}_t) \\ &\leq (1 - \eta\tau) (\text{KL}(z^* \| z_t) + (1 - 4\eta L)\text{KL}(z_t \| \bar{z}_t)). \end{aligned} \quad (\text{E.28})$$

Define:

$$V_t = \text{KL}(z^* \| z_t) + (1 - 4\eta L)\text{KL}(z_t \| \bar{z}_t). \quad (\text{E.29})$$

Then we have:

$$V_{t+1} \leq (1 - \eta\tau)V_t, \quad (\text{E.30})$$

and therefore:

$$\text{KL}(z^* \| z_{t+1}) \leq V_{t+1} \leq (1 - \eta\tau)^t V_0 = (1 - \eta\tau)^t \text{KL}(z^* \| z_0), \quad (\text{E.31})$$

which shows convergence of $\text{KL}(z^* \| z_{t+1})$. Next we show convergence of $\text{KL}(z^* \| \bar{z}_{t+1})$.

Similar to derivation of Equation (E.25), we have

$$-\langle z^* - \bar{z}_{t+1}, \log \bar{z}_{t+1} - \log z_{t+1} \rangle \leq \eta L (N\text{KL}(z_t \| \bar{z}_t) + N\text{KL}(\bar{z}_{t+1} \| z_t) + 4\text{KL}(z^* \| \bar{z}_{t+1})). \quad (\text{E.32})$$

Substituting this in Equation (E.19), we have:

$$(1 - 4\eta L)\text{KL}(z^* \| \bar{z}_{t+1}) \leq \text{KL}(z^* \| z_{t+1}) + \eta L (N\text{KL}(z_t \| \bar{z}_t) + N\text{KL}(\bar{z}_{t+1} \| z_t)). \quad (\text{E.33})$$

Now, using Equation (E.26), we have:

$$\begin{aligned} (1 - 4\eta L)\text{KL}(z^* \| \bar{z}_{t+1}) &\leq (1 - \eta\tau)\text{KL}(z^* \| z_t) - (1 - \eta\tau - 2N\eta L)\text{KL}(\bar{z}_{t+1} \| z_t) - \eta\tau\text{KL}(\bar{z}_{t+1} \| z^*) \\ &\quad - (1 - 4\eta L)\text{KL}(z_{t+1} \| \bar{z}_{t+1}) + 2N\eta L\text{KL}(z_t \| \bar{z}_t) \\ &\leq (1 - \eta\tau)\text{KL}(z^* \| z_t) + 2N\eta L\text{KL}(z_t \| \bar{z}_t) \\ &\leq \text{KL}(z^* \| z_t) + (1 - 4\eta L)\text{KL}(z_t \| \bar{z}_t) := V_t, \end{aligned} \quad (\text{E.34})$$

which gives us:

$$\text{KL}(z^* \| \bar{z}_{t+1}) \leq 2V_t \leq 2(1 - \eta\tau)^t V_0 = 2(1 - \eta\tau)^t \text{KL}(z^* \| z_0). \quad (\text{E.35})$$

This completes the proof of the first part of the theorem.

Next, we prove parameter convergence. We have from Equation (B.51):

$$\|\theta_i^{t+1} - \theta_i^*\|^2 = (1 - \eta\tau + \eta^2\tau^2)\|\theta_i^t - \theta_i^*\|^2 + \left(1 + \frac{1}{\eta\tau}(1 - \eta\tau)^2\right)\|\varepsilon_t\|^2, \quad (\text{E.36})$$

where $\theta_i^* = \frac{-\nabla_{g_i} f(g_i^*, g_{-i}^*)}{\tau}$ (Note that the term ε_t however is different here from definition of ε_t in Equation (B.51)). Now, for ONPG, we have:

$$\|\varepsilon_t\|^2 = \eta^2 \|\nabla_{g_i} f_i(g_i^*, g_{-i}^*) - \nabla_{g_i} f_i(\bar{g}_i^{t+1}, \bar{g}_{-i}^{t+1})\|^2 \leq \eta^2 L^2 \|\bar{z}_{t+1} - z^*\|^2 \leq 2\eta^2 L^2 \text{KL}(z^* \| z_{t+1}). \quad (\text{E.37})$$

This gives us:

$$\|\theta_i^{t+1} - \theta_i^*\|^2 \leq (1 - \eta\tau + \eta^2\tau^2)\|\theta_i^t - \theta_i^*\|^2 + C(1 - \eta\tau)^t, \quad (\text{E.38})$$

where

$$C = 4\eta^2 L^2 \left(1 + \frac{1}{\eta\tau} (1 - \eta\tau)^2 \right) \text{KL}(z^* \| z_0), \quad (\text{E.39})$$

from the first part of the Theorem proved above. For $\eta\tau < 1/2$, this reduces to:

$$\|\theta_i^{t+1} - \theta_i^*\|^2 \leq (1 - \eta\tau/2)\|\theta_i^t - \theta_i^*\|^2 + C(1 - \eta\tau)^t. \quad (\text{E.40})$$

Now, consider the Lyapunov function:

$$V_{t+1} = \|\theta_i^{t+1} - \theta_i^*\|^2 + \frac{2C}{\eta\tau} (1 - \eta\tau)^{t+1}. \quad (\text{E.41})$$

We have:

$$\begin{aligned} V_{t+1} &\leq (1 - \eta\tau/2)\|\theta_i^t - \theta_i^*\|^2 + C(1 - \eta\tau)^t + \frac{2C}{\eta\tau} (1 - \eta\tau)(1 - \eta\tau)^t \\ &= (1 - \eta\tau/2)\|\theta_i^t - \theta_i^*\|^2 + \frac{2C}{\eta\tau} (1 - \eta\tau)^t (1 - \eta\tau + \eta\tau/2) \\ &= (1 - \eta\tau/2)\|\theta_i^t - \theta_i^*\|^2 + \frac{2C}{\eta\tau} (1 - \eta\tau)^t (1 - \eta\tau/2) \\ &= (1 - \eta\tau/2) \left(\|\theta_i^t - \theta_i^*\|^2 + \frac{2C}{\eta\tau} (1 - \eta\tau)^t \right) \\ &= (1 - \eta\tau/2)V_t. \end{aligned} \quad (\text{E.42})$$

This shows linear convergence of the parameter θ_i to θ_i^* since:

$$\|\theta_i^{t+1} - \theta_i^*\|^2 \leq V_{t+1} \leq (1 - \eta\tau/2)^t V_0. \quad (\text{E.43})$$

Merging these inequalities for all players i completes the proof of the Theorem. \square

E.4 Proof of Corollary 4.5

We define $g_i := g_{\theta_i}$. The Duality gap for the regularized game is given by:

$$\begin{aligned} \text{DG}_\tau(g_1, g_2, \dots, g_N) &= \sum_{i=1}^N \left[f_i^T(g_i, g_{-i}) - \min_{\tilde{g}_i} f_i^T(\tilde{g}_i, g_{-i}) \right] = \max_{\tilde{g}_1, \tilde{g}_2, \dots, \tilde{g}_N} \sum_{i=1}^N [f_i^T(g_i, g_{-i}) - f_i^T(\tilde{g}_i, g_{-i})] \\ &= \max_{\tilde{g}_1, \tilde{g}_2, \dots, \tilde{g}_N} \sum_{i=1}^N [f_i^T(g_i, g_{-i}) - f_i^T(g_i, g_{-i}^*) + f_i^T(g_i, g_{-i}^*) - f_i^T(\tilde{g}_i, g_{-i}^*) \\ &\quad + f_i^T(\tilde{g}_i, g_{-i}^*) - f_i^T(\tilde{g}_i, g_{-i})] \\ &\leq \max_{\tilde{g}_1, \tilde{g}_2, \dots, \tilde{g}_N} \sum_{i=1}^N [f_i^T(g_i, g_{-i}) - f_i^T(g_i, g_{-i}^*) + f_i^T(g_i, g_{-i}^*) - f_i^T(g_i^*, g_{-i}^*) \\ &\quad + f_i^T(\tilde{g}_i, g_{-i}^*) - f_i^T(\tilde{g}_i, g_{-i})]. \end{aligned} \quad (\text{E.44})$$

Algorithm 4 Proximal Point Method

Initialize: $\theta_i^0 = 0$ for all players i .
for $t = 1, 2, \dots$ **do**
 $\theta_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_{\theta_i}} f_i(g_{\theta_i^{t+1}}, g_{\theta_i^{t+1}}) \quad \forall i \in [N]$.
end for

Next, we note that:

$$\begin{aligned}
 & \sum_{i=1}^N f_i^T(g_i, g_{-i}) - f_i^T(g_i, g_{-i}^*) + f_i^T(g_i, g_{-i}^*) - f_i^T(g_i^*, g_{-i}^*) + f_i^T(\tilde{g}_i, g_{-i}^*) - f_i^T(\tilde{g}_i, g_{-i}) \\
 & \leq \sum_{i=1}^N [\|f_i^T(g_i, g_{-i}) - f_i^T(g_i, g_{-i}^*)\| + \|f_i^T(g_i, g_{-i}^*) - f_i^T(g_i^*, g_{-i}^*)\| + \|f_i^T(\tilde{g}_i, g_{-i}^*) - f_i^T(\tilde{g}_i, g_{-i})\|] \\
 & \leq \sum_{i=1}^N [\|f_i^T(g_i, g_{-i}) - f_i^T(g_i, g_{-i}^*)\| + \|f_i^T(g_i, g_{-i}^*) - f_i^T(g_i^*, g_{-i}^*)\| + \|f_i^T(\tilde{g}_i, g_{-i}^*) - f_i^T(\tilde{g}_i, g_{-i})\|] \\
 & \leq C_1 \sum_{i=1}^N [\|g_i - g_i^*\| + \|g_{-i} - g_{-i}^*\|] \leq C_2 N \sqrt{\text{KL}(z^* \| z_t)}. \tag{E.45}
 \end{aligned}$$

This follows by noting that the functions f_i^T are Lipschitz since they are smooth functions defined on a compact domain. The last step follows from Pinsker's inequality and the fact that the l_1 norm is an upper bound for the l_2 norm.

Combining the two inequalities (E.44) and (E.45) we have:

$$\text{DG}_\tau(g_1, g_2, \dots, g_N) \leq C_2 N \sqrt{\text{KL}(z^* \| z_t)}. \tag{E.46}$$

Let DG denote the Duality gap of the unregularized problem. Then we have:

$$\text{DG}(g_1, g_2, \dots, g_N) \leq \text{DG}_\tau(g_1, g_2, \dots, g_N) + 2N\tau \log n. \tag{E.47}$$

Therefore, setting $\tau = \frac{\epsilon}{4N \log n}$ and solving the regularized problem to an accuracy of $\frac{\epsilon^2}{4C_2^2 N^2}$ in terms of KL divergence, we have that:

$$\text{DG}(g_1, g_2, \dots, g_N) \leq \epsilon, \tag{E.48}$$

completing the proof. \square

E.5 Proximal point and Extragradient methods to solve multi-player monotone games

We define $g_i := g_{\theta_i}$ and $g_i^t := g_{\theta_i^t}$ for simplicity.

E.5.1 Proximal-point updates

In this subsection, we show the convergence of the Proximal Point (PP) updates to the NE of the regularized N-player monotone game. The PP algorithm is presented in Algorithm 4

Theorem E.6. Let $z^* = (g_i^*)_{i=1}^N$ be the Nash equilibrium of Problem (4.2). Also, we denote $z_t = (g_i^t)_{i=1}^N$. Define $\theta_i^* := \frac{-\nabla_{g_i} f_i(g_i^*, g_{-i}^*)}{\tau}$. Then for updates in Algorithm 4, we have for $\eta\tau < 1/2$:

- $\text{KL}(z^* \| z_{t+1}) \leq (1 - \eta\tau)^t \text{KL}(z^* \| z_0)$.
- $\|\theta_{t+1} - \theta^*\|^2 \leq (1 - \eta\tau/2)^t V_0$, where,

$$V_0 = \|\theta^t - \theta^*\|^2 + \frac{2NC}{\eta\tau}, \quad C = 2\eta^2 L^2 \left(1 + \frac{4}{\eta\tau} (1 - \eta\tau)^2\right) \text{KL}(z^* \| z_0).$$

Proof. From the definition of the KL divergence we have:

$$-(\log z_{t+1} - (1 - \eta\tau) \log z_t - \eta\tau \log z^*, z^*) = -(1 - \eta\tau)\text{KL}(z^* \| z_t) + \text{KL}(z^* \| z_{t+1}), \quad (\text{E.49})$$

and

$$\langle \log z_{t+1} - (1 - \eta\tau) \log z_t - \eta\tau \log z^*, z_{t+1} \rangle = (1 - \eta\tau)\text{KL}(z_{t+1} \| z_t) + \text{KL}(z_{t+1} \| z^*). \quad (\text{E.50})$$

Substituting in Lemma E.4 with $z = z_{t+1}$, we have:

$$\langle \log z_{t+1} - (1 - \eta\tau) \log z_t - \eta\tau \log z^*, z_{t+1} - z^* \rangle \leq 0, \quad (\text{E.51})$$

and using the two equalities, we get:

$$-(1 - \eta\tau)\text{KL}(z^* \| z_t) + \text{KL}(z^* \| z_{t+1}) + (1 - \eta\tau)\text{KL}(z_{t+1} \| z_t) + \text{KL}(z_{t+1} \| z^*) \leq 0. \quad (\text{E.52})$$

This implies:

$$\text{KL}(z^* \| z_{t+1}) \leq (1 - \eta\tau)\text{KL}(z^* \| z_t). \quad (\text{E.53})$$

This shows linear convergence of the KL divergence to the Nash equilibrium for the proximal point method which completes the proof of the first part of the Theorem.

When the agents are playing the NE strategy, the updates reduce to

$$\theta_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_i} f_i(g_i^*, g_{-i}^*). \quad (\text{E.54})$$

From Lemma E.4, we know that the NE satisfy:

$$g_i^*(a) = \frac{e^{-[\nabla_{g_i} f(g_i^*, g_{-i}^*)]_a}}{K}, \quad (\text{E.55})$$

where $K = \sum_{a' \in \mathcal{A}_i} e^{-[\nabla_{g_i} f(g_i^*, g_{-i}^*)]_{a'}}$. Taking log on both sides, we have:

$$-\eta\nabla_{g_i} f(g_i^*, g_{-i}^*) = \eta\tau \log g_i^* + \eta\tau \log K. \quad (\text{E.56})$$

Substituting this back into the θ update, we have:

$$\begin{aligned} \theta_i^{t+1} &= (1 - \eta\tau)\theta_i^t + \eta\tau \log g_i^* + \eta\tau \log K = \theta_i^t - \eta\tau\theta_i^t + \eta\tau \log g_i^* + \eta\tau \log K - \eta\tau \log Z_{\theta_i^t} + \eta\tau \log Z_{\theta_i^t} \\ &= \theta_i^t + \eta\tau \log g_i^* - \eta\tau \log \left(\frac{e^{\theta_i^t}}{Z_{\theta_i^t}} \right) + \eta\tau \log \left(\frac{K}{Z_{\theta_i^t}} \right) = \theta_i^t + \eta\tau \log g_i^* - \eta\tau \log g_{\theta_i^t} + \eta\tau \log \left(\frac{K}{Z_{\theta_i^t}} \right) \\ &= \theta_i^t + \eta\tau \log \left(\frac{g_i^*}{g_{\theta_i^t}} \right) + \eta\tau \log \left(\frac{K}{Z_{\theta_i^t}} \right) = \theta_i^t + \eta\tau \log \left(\frac{g_i^* K}{g_{\theta_i^t} Z_{\theta_i^t}} \right). \end{aligned} \quad (\text{E.57})$$

We can further simplify this as:

$$\theta_i^{t+1} = \theta_i^t + \eta\tau \log \left(\frac{e^{-[\nabla_{g_i} f(g_i^*, g_{-i}^*)] / \tau}}{e^{\theta_i^t}} \right). \quad (\text{E.58})$$

This is nothing but:

$$\theta_i^{t+1} = \theta_i^t + \eta\tau \left(\frac{-[\nabla_{g_i} f(g_i^*, g_{-i}^*)]}{\tau} - \theta_i^t \right). \quad (\text{E.59})$$

Now the Proximal Point updates can be written as:

$$\theta_i^{t+1} = \theta_i^t + \eta\tau \left(\frac{-\nabla_{g_i} f(g_i^*, g_{-i}^*)}{\tau} - \theta_i^t \right) + \varepsilon_t, \quad (\text{E.60})$$

where $\varepsilon_t = \eta\nabla_{g_i} f_i(g_i^*, g_{-i}^*) - \eta\nabla_{g_i} f_i(g_i^{t+1}, g_{-i}^{t+1})$.

Defining $\theta_i^* = \frac{-\nabla_{g_i} f(g_i^*, g_{-i}^*)}{\tau}$, we have:

$$\begin{aligned} \|\theta_i^{t+1} - \theta_i^*\|^2 &= \|\theta_i^t + \eta\tau (\theta_i^* - \theta_i^t) + \varepsilon_t - \theta_i^*\|^2 = \|(1 - \eta\tau)(\theta_i^t - \theta_i^*) + \varepsilon_t\|^2 \\ &= (1 - \eta\tau)^2 \|\theta_i^t - \theta_i^*\|^2 + 2(1 - \eta\tau)(\theta_i^t - \theta_i^*)^\top \varepsilon_t + \|\varepsilon_t\|^2 \\ &\leq^{*1} (1 - \eta\tau)^2 \|\theta_i^t - \theta_i^*\|^2 + \eta\tau \|\theta_i^t - \theta_i^*\|^2 + \frac{4}{\eta\tau} (1 - \eta\tau)^2 \|\varepsilon_t\|^2 + \|\varepsilon_t\|^2 \\ &= (1 - \eta\tau + \eta^2\tau^2) \|\theta_i^t - \theta_i^*\|^2 + \left(1 + \frac{4}{\eta\tau} (1 - \eta\tau)^2\right) \|\varepsilon_t\|^2, \end{aligned} \quad (\text{E.61})$$

where *1 follows from Young's inequality.

Now, for the proximal point methods, we have:

$$\|\varepsilon_t\|^2 = \eta^2 \|\nabla_{g_i} f_i(g_i^*, g_{-i}^*) - \nabla_{g_i} f_i(g_i^{t+1}, g_{-i}^{t+1})\|^2 \leq \eta^2 L^2 \|z_{t+1} - z^*\|^2 \leq 2\eta^2 L^2 \text{KL}(z^* \| z_{t+1}), \quad (\text{E.62})$$

where it follows from Pinsker's inequality and the fact that the l_1 norm is an upper bound for the l_2 norm.

This gives us:

$$\|\theta_i^{t+1} - \theta_i^*\|^2 \leq (1 - \eta\tau + \eta^2\tau^2) \|\theta_i^t - \theta_i^*\|^2 + C(1 - \eta\tau)^t, \quad (\text{E.63})$$

where

$$C = 2\eta^2 L^2 \left(1 + \frac{4}{\eta\tau} (1 - \eta\tau)^2\right) \text{KL}(z^* \| z_0). \quad (\text{E.64})$$

For $\eta\tau < 1/2$, this reduces to:

$$\|\theta_i^{t+1} - \theta_i^*\|^2 \leq (1 - \eta\tau/2) \|\theta_i^t - \theta_i^*\|^2 + C(1 - \eta\tau)^t. \quad (\text{E.65})$$

Now, consider the Lyapunov function:

$$V_{t+1} = \|\theta_i^{t+1} - \theta_i^*\|^2 + \frac{2C}{\eta\tau} (1 - \eta\tau)^{t+1}. \quad (\text{E.66})$$

We have:

$$\begin{aligned} V_{t+1} &\leq (1 - \eta\tau/2) \|\theta_i^t - \theta_i^*\|^2 + C(1 - \eta\tau)^t + \frac{2C}{\eta\tau} (1 - \eta\tau)(1 - \eta\tau)^t \\ &= (1 - \eta\tau/2) \|\theta_i^t - \theta_i^*\|^2 + \frac{2C}{\eta\tau} (1 - \eta\tau)^t (1 - \eta\tau + \eta\tau/2) \\ &= (1 - \eta\tau/2) \|\theta_i^t - \theta_i^*\|^2 + \frac{2C}{\eta\tau} (1 - \eta\tau)^t (1 - \eta\tau/2) = (1 - \eta\tau/2) \left(\|\theta_i^t - \theta_i^*\|^2 + \frac{2C}{\eta\tau} (1 - \eta\tau)^t \right) \\ &= (1 - \eta\tau/2) V_t. \end{aligned} \quad (\text{E.67})$$

This shows linear convergence of the parameter θ_i to θ_i^* since:

$$\|\theta_i^{t+1} - \theta_i^*\|^2 \leq V_{t+1} \leq (1 - \eta\tau/2)^t V_0. \quad (\text{E.68})$$

Merging these inequalities for all players i completes the proof of the theorem. \square

Algorithm 5 Extragradient Method

Initialize: g_0 and h_0 .

for $t = 1, 2, \dots$ **do**

$$\bar{\theta}_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_{\theta_i}} f_i(g_{\theta_i}^t, g_{\theta_{-i}}^t) \quad \forall i \in [N].$$

$$\theta_i^{t+1} = (1 - \eta\tau)\theta_i^t - \eta\nabla_{g_{\theta_i}} f_i(g_{\bar{\theta}_i^{t+1}}, g_{\bar{\theta}_{-i}^{t+1}}) \quad \forall i \in [N].$$

end for

E.5.2 Extragradient updates

In this subsection, we show the convergence of the Extragradient method to the NE of the regularized N-player monotone game. The EG algorithm is presented in Algorithm 5.

Theorem E.7. Let $z^* = (g_i^*)_{i=1}^N$ be the unique Nash equilibrium of Problem (4.2). Also, we denote $z_t = (g_{\theta_i}^t)_{i=1}^N$. Then for updates in Algorithm 5, we have for stepsize satisfying $0 < \eta < \frac{1}{2NL + \tau}$:

- Convergence of distributions:

$$\max \{ \text{KL}(z^* \| z_t), \text{KL}(z^* \| \bar{z}_{t+1}) \} \leq (1 - \eta\tau)^t 2\text{KL}(z^* \| z_0). \quad (\text{E.69})$$

- Convergence of parameters:

$$\|\theta_{t+1} - \theta^*\|^2 \leq (1 - \eta\tau/2)^t V_0, \quad (\text{E.70})$$

where,

$$V_0 = \|\theta^t - \theta^*\|^2 + \frac{2NC}{\eta\tau}, \quad C = 4\eta^2 L^2 \left(1 + \frac{4}{\eta\tau} (1 - \eta\tau)^2 \right) \text{KL}(z^* \| z_0).$$

Proof. From the EG updates in Algorithm 5, we have:

$$\log \bar{g}_i^{t+1} - \log g_i^{t+1} = -\eta (\nabla_{g_i} f_i(g_i^t, g_{-i}^t) - \nabla_{g_i} f_i(\bar{g}_i^{t+1}, \bar{g}_{-i}^{t+1})) + c \cdot \mathbf{1}. \quad (\text{E.71})$$

This implies:

$$\begin{aligned} \langle \log \bar{g}_i^{t+1} - \log g_i^{t+1}, \bar{g}_i^{t+1} - g_i^{t+1} \rangle &= -\eta \langle \nabla_{g_i} f_i(g_i^t, g_{-i}^t) - \nabla_{g_i} f_i(\bar{g}_i^{t+1}, \bar{g}_{-i}^{t+1}), \bar{g}_i^{t+1} - g_i^{t+1} \rangle \\ &\leq^{*1} \eta \|\nabla_{g_i} f_i(g_i^t, g_{-i}^t) - \nabla_{g_i} f_i(\bar{g}_i^{t+1}, \bar{g}_{-i}^{t+1})\| \|\bar{g}_i^{t+1} - g_i^{t+1}\| \\ &\leq^{*2} \eta L (\|\bar{g}_i^{t+1} - g_i^t\| + \|\bar{g}_{-i}^{t+1} - g_{-i}^t\|) \|\bar{g}_i^{t+1} - g_i^{t+1}\| \\ &\leq^{*3} \frac{1}{2} \eta L (\|\bar{g}_i^{t+1} - g_i^t\|^2 + 2\|\bar{g}_i^{t+1} - g_i^{t+1}\|^2 + \|\bar{g}_{-i}^{t+1} - g_{-i}^t\|^2), \end{aligned} \quad (\text{E.72})$$

where (*1) follows from the fact that $a^\top b \leq \|a\| \|b\|$, (*2) follows from Assumption 4.1, and (*3) follows from $x \cdot y \leq \frac{1}{2}(x^2 + y^2)$. Since this is true for all players, we have:

$$\begin{aligned} \langle \log \bar{z}_{t+1} - \log z_{t+1}, \bar{z}_{t+1} - z_{t+1} \rangle &\leq \frac{1}{2} \eta L \sum_{i=1}^N (\|\bar{g}_i^{t+1} - g_i^t\|^2 + 2\|\bar{g}_i^{t+1} - g_i^{t+1}\|^2 + \|\bar{g}_{-i}^{t+1} - g_{-i}^t\|^2) \\ &\leq \frac{1}{2} \eta L (N \|\bar{z}_{t+1} - z_t\|^2 + 2\|\bar{z}_{t+1} - z_{t+1}\|^2) \\ &\leq \frac{N}{2} \eta L (\|\bar{z}_{t+1} - z_t\|^2 + \|\bar{z}_{t+1} - z_{t+1}\|^2) \\ &\leq^{*1} N \eta L (\text{KL}(\bar{z}_{t+1} \| z_t) + \text{KL}(z_{t+1} \| \bar{z}_{t+1})), \end{aligned} \quad (\text{E.73})$$

where (*1) follows from Pinsker's Inequality and the fact that the l_1 norm is an upper bound for the l_2 norm. Substituting this in Equation (E.18), we have

$$\begin{aligned} \text{KL}(z^* \| z_{t+1}) &\leq (1 - \eta\tau) \text{KL}(z^* \| z_t) - (1 - \eta\tau - N\eta L) \text{KL}(\bar{z}_{t+1} \| z_t) \\ &\quad - \eta\tau \text{KL}(\bar{z}_{t+1} \| z^*) - (1 - N\eta L) \text{KL}(z_{t+1} \| \bar{z}_{t+1}). \end{aligned} \quad (\text{E.74})$$

If $\eta \leq \frac{1}{\tau + NL}$, we have:

$$\text{KL}(z^* \| z_{t+1}) \leq (1 - \eta\tau) \text{KL}(z^* \| z_t). \quad (\text{E.75})$$

Similar to the derivation of Equation (E.73), we have:

$$-\langle \log \bar{z}_{t+1} - \log z_{t+1}, z^* - \bar{z}_{t+1} \rangle \leq N\eta L (\text{KL}(\bar{z}_{t+1} \| z_t) + \text{KL}(z^* \| \bar{z}_{t+1})). \quad (\text{E.76})$$

Substituting this in Equation (E.19) we have:

$$(1 - N\eta L) \text{KL}(z^* \| \bar{z}_{t+1}) \leq \text{KL}(z^* \| z_{t+1}) + N\eta L \text{KL}(\bar{z}_{t+1} \| z_t). \quad (\text{E.77})$$

Now, plugging Inequality (E.74) in (E.77) we have:

$$\begin{aligned} (1 - N\eta L) \text{KL}(z^* \| \bar{z}_{t+1}) &\leq (1 - \eta\tau) \text{KL}(z^* \| z_t) \\ &\quad - (1 - \eta\tau - 2N\eta L) \text{KL}(\bar{z}_{t+1} \| z_t) - \eta\tau \text{KL}(\bar{z}_{t+1} \| z^*) - (1 - N\eta L) \text{KL}(z_{t+1} \| \bar{z}_{t+1}). \end{aligned} \quad (\text{E.78})$$

With stepsize $\eta < \frac{1}{\tau + 2NL}$ we have:

$$\text{KL}(z^* \| \bar{z}_{t+1}) \leq 2 \text{KL}(z^* \| z_t) \leq 2(1 - \eta\tau)^t \text{KL}(z^* \| z_0), \quad (\text{E.79})$$

which completes the first part of the proof.

The proof of parameter convergence follows exactly from the proof of Theorem 4.3 and we avoid rewriting it here. \square

E.6 Parameterization with function approximation

In this section, we discuss the monotone game setting with function approximation for policy parameterization, as discussed for matrix games in §C. In this setting, the regularized problem each player faces is:

$$\min_{\theta_i \in \mathbb{R}^d} f_i(g_{\theta_i}, g_{\theta_{-i}}) - \tau \mathcal{H}(g_{\theta_i}), \quad (\text{E.80})$$

where g_{θ_i} is a log-linear policy parametrization. In the next lemma, we show the existence of a NE in this setting as well as the equivalence of this problem to one on the entire simplex, in the same spirit as §C for matrix games:

Lemma E.8. The in-class NE (Definition E.2) for the unregularized (and regularized) monotone game under the log-linear policy parametrization exists. Also, under Assumption C.1, Problem (E.80) is equivalent to:

$$\min_{g_{\theta_i} \in \Delta} f_i(\Psi g_{\theta_i}, \Psi g_{\theta_{-i}}) - \tau \mathcal{H}(g_{\theta_i}), \quad (\text{E.81})$$

where Ψ is defined in Equation (C.7).

Now, using Lemma E.8, and Proposition C.6, we describe an algorithm to solve this problem in the following corollary.

Corollary E.9. The update rule:

$$\begin{aligned} \bar{\theta}_i^{t+1} &= (1 - \eta\tau) \theta_i^t - [(M^\top)^{-1} | 0] \eta \tilde{P} \nabla_{g_{\theta_i}} f_i(g_{\bar{\theta}_i^t}, g_{\bar{\theta}_{-i}^t}) \\ \theta_i^{t+1} &= (1 - \eta\tau) \theta_i^t - [(M^\top)^{-1} | 0] \eta \tilde{P} \nabla_{g_{\theta_i}} f_i(g_{\bar{\theta}_i^{t+1}}, g_{\bar{\theta}_{-i}^{t+1}}), \end{aligned}$$

solves Problem (E.80) with similar guarantees given by Theorem 4.3. Here, the NE parameter value to which the algorithm converges to is given by $\theta_i^* = \frac{-[(M^\top)^{-1} | 0] \tilde{P} \nabla_{g_{\theta_i}} f_i(g_i^*, g_{-i}^*)}{\tau}$. Furthermore, by choosing the regularization parameter τ small enough, like in Corollary 4.5, we reach an ϵ -in-class NE (Definition E.3) of the unregularized monotone game under the function approximation setting.

E.7 Proof of Lemma E.8

The proof of this Lemma follows along the lines of Lemma C.4 and Theorem C.5. The key here is to notice that:

$$\nabla_{g_{\theta_i}} f_i(\Psi g_{\theta_i}, \Psi g_{\theta_{-i}}) = \Psi^\top \nabla_{\Psi g_{\theta_i}} f_i(\Psi g_{\theta_i}, \Psi g_{\theta_{-i}}), \quad (\text{E.82})$$

and that:

$$\Psi \mu = \mu \quad \forall \mu \in \tilde{\Delta}. \quad (\text{E.83})$$

The rest of the proof is identical to the proof of Lemma C.4. \square

Algorithm 6 Optimistic NPG for Markov Games

Initialize: $Q_0 = 0$
for $t = 1, 2, \dots, T_{outer}$ **do**
 for $s = 1, 2, \dots, |\mathcal{S}|$ **do**
 Let $Q_t(s, a, b) := r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [V_t(s')]$
 Solve $\min_{\theta \in \mathbb{R}^n} \max_{\nu \in \mathbb{R}^n} f_\tau(Q(s); g_\theta, h_\nu)$ by running the Optimistic NPG algorithm (Algorithm 1) for T_{inner} iterations and return the last iterates $(\theta_{T_{inner}}, \nu_{T_{inner}})$.
 Set $V_{t+1}(s) = f_\tau(Q_t(s); g_{\theta_{T_{inner}}}(\cdot | s), h_{\nu_{T_{inner}}}(\cdot | s))$
 end for
end for

F Missing Definitions and Proofs in §5

F.1 Zero-Sum Markov Games

A two-player zero-sum Markov game is characterized by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{B}, P, r, \gamma)$, where \mathcal{S} is the state space; \mathcal{A}, \mathcal{B} are the action spaces of players 1 and 2, respectively; $P : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow \Delta(\mathcal{S})$ denotes the transition probability of states; $r : \mathcal{S} \times \mathcal{A} \times \mathcal{B} \rightarrow [0, 1]$ denotes the bounded reward function of player 1 (thus $-r$ is the reward function of player 2); and $\gamma \in [0, 1)$ is the discount factor. The goal of player 1 (player 2) is to minimize (maximize) the long-term accumulative discounted reward.

Specifically, at each time t , player 1 (player 2) has a Markov stationary policy $g : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ ($h : \mathcal{S} \rightarrow \Delta(\mathcal{B})$). The state makes a transition from s_t to s_{t+1} following the probability distribution $P(\cdot | s_t, a_t, b_t)$, given (a_t, b_t) . As in the Markov decision process model, one can define the *state-value function* under a pair of joint policies (g, h) as

$$V^{g,h}(s) := \mathbb{E}_{a_t \sim g(\cdot | s_t), b_t \sim h(\cdot | s_t)} \left[\sum_{t \geq 0} \gamma^t r(s_t, a_t, b_t) \mid s_0 = s \right].$$

Also, the *state-action/Q-value function* under (g, h) is defined as

$$Q^{g,h}(s, a, b) := \mathbb{E}_{a_t \sim g(\cdot | s_t), b_t \sim h(\cdot | s_t)} \left[\sum_{t \geq 0} \gamma^t r(s_t, a_t, b_t) \mid s_0 = s, a_0 = a, b_0 = b \right].$$

Similar as the matrix game setting, a common solution concept in Markov game is also the (Markov perfect) *Nash equilibrium* policy pair (g^*, h^*) , which satisfies the following saddle-point inequality:

$$V^{g,h^*}(s) \geq V^{g^*,h^*}(s) \geq V^{g^*,h}(s), \quad \forall s \in \mathcal{S}.$$

It follows from [Shapley \(1953\)](#); [Filar and Vrieze \(2012\)](#) that there exists a Nash equilibrium $(g^*, h^*) \in \Delta(\mathcal{A})^{|\mathcal{S}|} \times \Delta(\mathcal{B})^{|\mathcal{S}|}$ for two-player discounted zero-sum MGs. The state-value $V^* := V^{g^*,h^*}$ is referred to as the *value of the game*. The corresponding Q-value function is denoted by Q^* .

We focus on the softmax parameterization g_θ and h_ν of the policies g and h , respectively.

Remark F.1. We note that all results presented in this section also follow for the case where the cardinality of the action spaces for both players are asymmetric. However, we stick to the case where the number of action is the same for both players in all states for ease of exposition.

F.2 Proof of Theorem 5.3

We first have the following lemma which shows the smoothness property of the NE policy with respect to the game matrix Q .

Lemma F.2. Consider the following entropy regularized game:

$$\min_{x \in \tilde{\Delta}} \max_{y \in \tilde{\Delta}} x^\top Q y - \tau \mathcal{H}(x) + \tau \mathcal{H}(y), \tag{F.1}$$

where $\tilde{\Delta} \subseteq \Delta$ is a convex compact subset of the simplex given by Equation (C.5). Let (x_Q^*, y_Q^*) denote the unique solution to this problem (note that this is unique since we have a strongly convex-strongly concave objective over a compact convex set). Then, we have:

$$\max \{ \|x_{Q_1}^* - x_{Q_2}^*\|, \|y_{Q_1}^* - y_{Q_2}^*\| \} \leq C \cdot \|Q_1 - Q_2\|_F, \quad (\text{F.2})$$

for some constant $C > 0$ and for any $Q_1, Q_2 \in \mathbb{R}^{n \times n}$.

Proof. First notice that by the proof of Theorem C.5, solving (F.1) is equivalent to solving

$$\min_{x \in \Delta} \max_{y \in \Delta} x^\top \Psi^\top Q \Psi y - \tau \mathcal{H}(x) + \tau \mathcal{H}(y), \quad (\text{F.3})$$

with Ψ being defined in Equation (C.7), i.e., (x_Q^*, y_Q^*) also solves (F.3). Also, notice that the solution to (F.3) always lies in the relative interior of Δ , given by (D.14). In other words, $[x_Q^*]_a > 0$ and $[y_Q^*]_a > 0$ for all a . Due to the simplex constraint, the free variable is of dimension $n - 1$, and the last dimension of x can be represented as $1 - \sum_{a=1}^{n-1} x_a > 0$. Let $\tilde{x} = (x_1, x_2, \dots, x_{n-1})^\top$, $\tilde{y} = (y_1, y_2, \dots, y_{n-1})^\top$, and $f(x, y) := x^\top \Psi^\top Q \Psi y - \tau \mathcal{H}(x) + \tau \mathcal{H}(y)$. Recall that $f(x, y)$ is strongly convex in x and strongly concave in y . Note that $f(x, y) = f(\Lambda(\tilde{x}), \Lambda(\tilde{y})) = \tilde{f}(\tilde{x}, \tilde{y}) :=$

$$\Lambda(\tilde{x})^\top \Psi^\top Q \Psi \Lambda(\tilde{y}) - \tau \mathcal{H}(\Lambda(\tilde{x})) + \tau \mathcal{H}(\Lambda(\tilde{y})), \text{ where } x = \Lambda(\tilde{x}) = \begin{bmatrix} I \\ -\mathbf{1}^\top \end{bmatrix} \tilde{x} + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}, \text{ and } \mathbf{1} \text{ denotes an all-one vector}$$

of proper dimension. Note that $\tilde{f}(\tilde{x}, \tilde{y})$ is also strongly convex in \tilde{x} and strongly concave in \tilde{y} , as for any \tilde{y} , the Hessian $\nabla_{\tilde{x}}^2 \tilde{f}(\tilde{x}, \tilde{y}) = [I \mid -\mathbf{1}] \nabla_x^2 f(\Lambda(\tilde{x}), \Lambda(\tilde{y})) \begin{bmatrix} I \\ -\mathbf{1}^\top \end{bmatrix} \succ 0$, so is the Hessian with respect to \tilde{y} for any \tilde{x} . Hence, the solution to the minimax problem

$$\min_{\{\tilde{x} \mid \tilde{x}_a \geq 0, 1 - \sum_{a=1}^{n-1} \tilde{x}_a \geq 0\}} \max_{\{\tilde{y} \mid \tilde{y}_a \geq 0, 1 - \sum_{a=1}^{n-1} \tilde{y}_a \geq 0\}} \tilde{f}(\tilde{x}, \tilde{y}) \quad (\text{F.4})$$

is given by $(\tilde{x}_Q^*, \tilde{y}_Q^*)$, where \tilde{x}_Q^* and \tilde{y}_Q^* are just the first $n - 1$ dimensions of x_Q^* and y_Q^* , satisfying $[\tilde{x}_Q^*]_a > 0$, $[\tilde{y}_Q^*]_a > 0$, and $1 - \sum_{a=1}^{n-1} [\tilde{x}_Q^*]_a > 0$, $1 - \sum_{a=1}^{n-1} [\tilde{y}_Q^*]_a > 0$, i.e., the constraints in (F.4) are not violated at $(\tilde{x}_Q^*, \tilde{y}_Q^*)$. By KKT conditions, it holds that at $(\tilde{x}_Q^*, \tilde{y}_Q^*)$

$$\tau \nabla_{\tilde{x}} \mathcal{H}(\Lambda(\tilde{x}_Q^*)) - [I \mid -\mathbf{1}] \Psi^\top Q \Psi \Lambda(\tilde{y}_Q^*) = 0 \quad (\text{F.5})$$

$$\tau \nabla_{\tilde{y}} \mathcal{H}(\Lambda(\tilde{y}_Q^*)) + [I \mid -\mathbf{1}] \Psi^\top Q^\top \Psi \Lambda(\tilde{x}_Q^*) = 0. \quad (\text{F.6})$$

Define operator $\mathcal{G}(\tilde{x}, \tilde{y}, \text{vec}(Q))$ as

$$\mathcal{G}(\tilde{x}, \tilde{y}, \text{vec}(Q)) := \begin{bmatrix} \tau \nabla_{\tilde{x}} \mathcal{H}(\Lambda(\tilde{x})) - [I \mid -\mathbf{1}] \Psi^\top Q \Psi \Lambda(\tilde{y}) \\ \tau \nabla_{\tilde{y}} \mathcal{H}(\Lambda(\tilde{y})) + [I \mid -\mathbf{1}] \Psi^\top Q^\top \Psi \Lambda(\tilde{x}) \end{bmatrix},$$

where $\text{vec}(Q)$ is the vectorization of Q . Then, $(\tilde{x}_Q^*, \tilde{y}_Q^*)$ is given by the solution to $\mathcal{G}(\tilde{x}, \tilde{y}, \text{vec}(Q)) = 0$. Note that the Jacobian of \mathcal{G} with respect to $[\tilde{x}^\top, \tilde{y}^\top]^\top$ is

$$\mathcal{M}(\tilde{x}, \tilde{y}, \text{vec}(Q)) := \begin{bmatrix} \frac{\partial \mathcal{G}}{\partial \tilde{x}} & \frac{\partial \mathcal{G}}{\partial \tilde{y}} \end{bmatrix} = \begin{bmatrix} \tau \nabla_{\tilde{x}}^2 \mathcal{H}(\Lambda(\tilde{x})) & -[I \mid -\mathbf{1}] \Psi^\top Q \Psi \begin{bmatrix} I \\ -\mathbf{1}^\top \end{bmatrix} \\ [I \mid -\mathbf{1}] \Psi^\top Q^\top \Psi \begin{bmatrix} I \\ -\mathbf{1}^\top \end{bmatrix} & \tau \nabla_{\tilde{y}}^2 \mathcal{H}(\Lambda(\tilde{y})) \end{bmatrix},$$

which is always invertible for any \tilde{x} and \tilde{y} belonging to the constraints in (F.4), due to the fact that $\tau \nabla_{\tilde{x}}^2 \mathcal{H}(\Lambda(\tilde{x}))$, $\tau \nabla_{\tilde{y}}^2 \mathcal{H}(\Lambda(\tilde{y})) \succ 0$, and $\mathcal{M}(\tilde{x}, \tilde{y}, \text{vec}(Q))$ is skew-symmetric, yielding the fact that the real parts of the eigenvalues of $\mathcal{M}(\tilde{x}, \tilde{y}, \text{vec}(Q))$, which are the eigenvalues of $(\mathcal{M}^\top + \mathcal{M})/2$, are always positive, and in fact uniformly lower bounded by some constant $\eta > 0$ for any \tilde{x} and \tilde{y} belong to the constraints in (F.4), due to the strong convexity of $\mathcal{H}(\Lambda(\tilde{x}))$ and $\mathcal{H}(\Lambda(\tilde{y}))$. Hence, we have

$$\|\mathcal{M}(\tilde{x}, \tilde{y}, \text{vec}(Q))^{-1}\|_2 \leq \frac{2}{\lambda_{\min}(\mathcal{M}(\tilde{x}, \tilde{y}, \text{vec}(Q)) + \mathcal{M}(\tilde{x}, \tilde{y}, \text{vec}(Q))^\top)} \leq \frac{1}{\eta}.$$

Due to the invertibility of $\mathcal{M}(\tilde{x}, \tilde{y}, \text{vec}(Q))$, we can apply implicit function theorem (Krantz and Parks, 2012) for any solution to $\mathcal{G}(\tilde{x}, \tilde{y}, \text{vec}(Q)) = 0$, and obtain that for any such a solution $(\tilde{x}_Q^*, \tilde{y}_Q^*, \text{vec}(Q))$, there exists a neighborhood of it such that for any $(\tilde{x}, \tilde{y}, \text{vec}(\tilde{Q}))$ in the neighborhood

$$\frac{\partial[\tilde{x}^\top, \tilde{y}^\top]^\top}{\partial \text{vec}(\tilde{Q})} = -\mathcal{M}(\tilde{x}, \tilde{y}, \text{vec}(\tilde{Q}))^{-1} \cdot \frac{\partial \mathcal{G}(\tilde{x}, \tilde{y}, \text{vec}(\tilde{Q}))}{\partial \text{vec}(\tilde{Q})}.$$

Notice that $\frac{\partial \mathcal{G}(\tilde{x}, \tilde{y}, \text{vec}(\tilde{Q}))}{\partial \text{vec}(\tilde{Q})}$ is uniformly bounded in norm on the constrained sets in (F.4), due to the boundedness of the sets. Hence, there exists a uniform constant $C' > 0$ such that

$$\left\| \frac{\partial[(\tilde{x}_Q^*)^\top, (\tilde{y}_Q^*)^\top]^\top}{\partial \text{vec}(Q)} \right\|_2 \leq \left\| \mathcal{M}(\tilde{x}_Q^*, \tilde{y}_Q^*, \text{vec}(Q))^{-1} \right\|_2 \cdot \left\| \frac{\partial \mathcal{G}(\tilde{x}_Q^*, \tilde{y}_Q^*, \text{vec}(Q))}{\partial \text{vec}(\tilde{Q})} \right\|_2 \leq C'$$

for any $(\tilde{x}_Q^*, \tilde{y}_Q^*, \text{vec}(Q))$. By the mean-value theorem, we know that

$$\left\| [(\tilde{x}_{Q_1}^*)^\top, (\tilde{y}_{Q_1}^*)^\top]^\top - [(\tilde{x}_{Q_2}^*)^\top, (\tilde{y}_{Q_2}^*)^\top]^\top \right\|_2 \leq C' \cdot \left\| \text{vec}(Q_1) - \text{vec}(Q_2) \right\|_2.$$

Finally, notice that

$$\left\| [(x_{Q_1}^*)^\top, (y_{Q_1}^*)^\top]^\top - [(x_{Q_2}^*)^\top, (y_{Q_2}^*)^\top]^\top \right\|_2 \leq \left\| \begin{bmatrix} I \\ -\mathbf{1}^\top \end{bmatrix} \right\|_2 \cdot \left\| [(\tilde{x}_{Q_1}^*)^\top, (\tilde{y}_{Q_1}^*)^\top]^\top - [(\tilde{x}_{Q_2}^*)^\top, (\tilde{y}_{Q_2}^*)^\top]^\top \right\|_2,$$

which completes the proof by the equivalence of norms. \square

F.2.1 Tabular case and proof of Theorem 5.3

We set the stepsize to be:

$$\eta = \frac{1 - \gamma}{2(1 + \tau(\log n + 1 - \gamma))}. \quad (\text{F.7})$$

The convergence of the Q -function follows from Lemma B.7, along with Theorem 2 in Cen et al. (2021b).

We show the convergence of the parameter next. We define:

$$\theta^*(s) = \frac{-Q_\tau^* h_{\nu^*}(\cdot | s)}{\tau}, \quad \nu^*(s) = \frac{Q_\tau^{*\top} g_{\theta^*}(\cdot | s)}{\tau}. \quad (\text{F.8})$$

We have:

$$\|\theta_t - \theta_{Q_t}^*\|^2 = \|\theta_t - \theta_{Q_t}^* + \theta_{Q_t}^* - \theta^*\|^2 \leq 2\|\theta_t - \theta_{Q_t}^*\|^2 + 2\|\theta_{Q_t}^* - \theta^*\|^2. \quad (\text{F.9})$$

Now, the first term converges approximately after the inner loop terminates. We can analyze the second term as follows:

$$\begin{aligned} \|\theta_{Q_t}^* - \theta^*\|^2 &=^{*1} \frac{1}{\tau} \|Q_t h_{\nu_{Q_t}^*} - Q^* h_{\nu_{Q_t}^*}\|^2 = \frac{1}{\tau} \|Q_t h_{\nu_{Q_t}^*} - Q^* h_{\nu_{Q_t}^*} + Q^* h_{\nu_{Q_t}^*} - Q^* h_{\nu_{Q_t}^*}\|^2 \\ &\leq \frac{2}{\tau} \left(\|Q_t h_{\nu_{Q_t}^*} - Q^* h_{\nu_{Q_t}^*}\|^2 + \|Q^* h_{\nu_{Q_t}^*} - Q^* h_{\nu_{Q_t}^*}\|^2 \right) \\ &\leq \frac{2}{\tau} \left(\|Q_t - Q^*\|_F^2 + \|Q^*\|_\infty^2 \|h_{\nu_{Q_t}^*} - h_{\nu_{Q_t}^*}\|_1^2 \right) \\ &\leq^{*2} \frac{2}{\tau} \left(\|Q_t - Q^*\|_F^2 + C \|Q^*\|_\infty^2 \|Q_t - Q^*\|_F^2 \right) = \frac{2}{\tau} (1 + C^2 \|Q^*\|_\infty^2) \|Q_t - Q^*\|_F^2, \end{aligned} \quad (\text{F.10})$$

where (*1) follows from the definition of θ^* and (*2) follows from Lemma F.2. Substituting this in Equation (F.9) we have:

$$\|\theta_t - \theta_{Q_t}^*\|^2 \leq 2\|\theta_t - \theta_{Q_t}^*\|^2 + \frac{4}{\tau} (1 + C^2 \|Q^*\|_\infty^2) \|Q_t - Q^*\|_F^2. \quad (\text{F.11})$$

Therefore, we have $\|\theta_t - \theta_{Q^*}^*\|^2 \leq \epsilon$ if

$$\|\theta_t - \theta_{Q_t}^*\|^2 \leq \frac{\epsilon}{4}, \quad \|Q_t - Q^*\|_F^2 \leq \frac{\epsilon}{\frac{8}{\tau}(1 + C^2\|Q^*\|_\infty^2)}. \quad (\text{F.12})$$

Note that the first term can be achieved by setting the inner-loop iterations T_{inner} to be the following (from Theorem 3.4):

$$T_{inner} = \mathcal{O}\left(\frac{1}{\eta\tau} \left(\log \frac{1}{\epsilon} + \log \frac{1}{1-\gamma} + \log \log n + \log \frac{1}{\eta}\right)\right), \quad (\text{F.13})$$

and the second term can be achieved by setting the outer-loop iterations T_{outer} to be:

$$T_{outer} = \mathcal{O}\left(\frac{1}{1-\gamma} \left(\log \frac{1}{\epsilon} + \log \left(\frac{8}{\tau}(1 + C^2\|Q^*\|_\infty^2)\right) + \log \frac{1 + \tau \log n}{1-\gamma}\right)\right). \quad (\text{F.14})$$

This completes the proof. \square

F.3 Function approximation setting

In this subsection, we discuss Markov games where the policies have a log-linear parametrization. The basic idea is to follow the tabular setting, but only for those states for which there is an action for which the feature vector corresponding to the state action pair is non-zero. The detailed description of the algorithm can be found in §F.3.1. We first make the following assumption on the feature matrix ϕ .

Assumption F.3. The feature matrix Φ is of the form $\Phi = [\phi_1, \phi_2, \dots, \phi_{|S| \times |A|}] = [I \mid 0]$.

This assumption plays a similar role as Assumption C.1 for the matrix game case, whose discussion has been provided after Assumption C.1. Essentially, it ensures that the feature matrix is not only full-rank (a standard assumption in single-agent RL), but also lead to a convex subset in the policy space, so that the in-class equilibrium exists (and thus the optimization problem is well-defined).

Theorem F.4. Let Q_τ^* be the in-class NE (Definition 5.1) Q -value defined in Equation (5.4) of the regularized problem, under the log-linear parametrization satisfying Assumption F.3. Choose the stepsize $\eta = \frac{1-\gamma}{2(1+\tau(\log n+1-\gamma))}$ for the inner loop in Algorithm 7. Then, after running Algorithm 7 for

$$\begin{aligned} T_{inner} &= \mathcal{O}\left(\frac{1}{\eta\tau} \left(\log \frac{1}{\epsilon} + \log \frac{1}{1-\gamma} + \log \log n + \log \frac{1}{\eta}\right)\right), \\ T_{outer} &= \mathcal{O}\left(\frac{1}{1-\gamma} \left(\log \frac{1}{\epsilon} + \log \left(\frac{8}{\tau}(1 + C^2\|Q^*\|_\infty^2)\right) + \log \frac{1 + \tau \log n}{1-\gamma}\right)\right), \end{aligned} \quad (\text{F.15})$$

iterations, we have $\|Q_T - Q_\tau^*\|_F \leq \epsilon$ and $\max\{\|\theta_T - \theta^*\|, \|\nu_T - \nu^*\|\} \leq \epsilon$ where (Q_T, θ_T, ν_T) is the output of Algorithm 6 after T iterations and (θ^*, ν^*) are defined in Equation (F.16).

Remark F.5. Note that Theorem F.4 provides the convergence rate for a two player Markov game under the function approximation setting. This covers the tabular case by setting the feature matrix, Φ (See Equation (C.2) in §C), to be equal to the identity matrix. In particular, making this substitution recovers the results of Theorem 5.3.

Remark F.6. We remark that Theorem F.4, to the best of our knowledge, provides the first symmetric algorithm with convergence guarantees for the Markov game under the function approximation setting. The only other existing result in this setting is Zhao et al. (2021), where the update is asymmetric, and one of the players performs multiple updates while the other player updates once. An asymmetric update-rule also appears in Daskalakis et al. (2020), without function approximation. Our results also improve over Wei et al. (2021); Cen et al. (2021b) by generalizing the results to the case of function approximation, as well as showing parameter convergence.

F.3.1 Proof of Theorem F.4

In this section, we consider the Markov game with log-linear parametrized policies. For sake of simplicity, we assume that the feature matrix $\Phi \in \mathbb{R}^{d \times |S|n}$ is of the form $\Phi = [I_d \mid 0]$ (Assumption F.3). Each column of Φ is a feature vector corresponding to some state action pair (s, a) . Note that for each state, there could be 0 to $\min\{n, d\}$ actions for which the feature vector is non-zero.

Algorithm 7 Optimistic NPG for Markov Games with Function Approximation

Initialize: $Q_0 = 0$

for $t = 1, 2, \dots, T_{outer}$ **do**

Let $Q_t(s, a, b) := r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [f_\tau(Q_t(s'); g_{\theta_{T_{inner}}}, h_{\nu_{T_{inner}}})]$

for $i = 1, 2, \dots, D$ **do**

Solve $\min_{\theta \in \mathbb{R}^{|A_{s_{d_i}}|}} \max_{\nu \in \mathbb{R}^{|B_{s_{d_i}}|}} f_\tau(Q(s_{d_i}); g_\theta, h_\nu)$ by running the Optimistic NPG algorithm (Algorithm 2) with feature matrix $\Phi_{s_{d_i}}$ for T_{inner} iterations and return the last iterates $(\theta_{T_{inner}}^{d_i}, \nu_{T_{inner}}^{d_i})$.

end for

Set $(\theta_{T_{inner}}, \nu_{T_{inner}}) = \left([\theta_{T_{inner}}^{d_1}, \theta_{T_{inner}}^{d_2}, \dots, \theta_{T_{inner}}^{d_D}], [\nu_{T_{inner}}^{d_1}, \nu_{T_{inner}}^{d_2}, \dots, \nu_{T_{inner}}^{d_D}] \right)$.

end for

Now, consider a state $s \in \mathcal{S}$. Define $A_s = \{a \in \mathcal{A} : \Phi_{s,a} \neq 0\}$ where $\Phi_{s,a}$ corresponds to the column in the feature matrix for state s , and Action a , and here 0 denotes the zero vector. Therefore A_s is the set of actions in state s for which the feature vector is non-zero. For sake of notational simplicity, let these be the actions $1, 2, \dots, |A_s|$. Note that A_s can be an empty set. We further assume that the first $|A_1|$ columns of Φ are corresponding to State 1, the next $|A_2|$ columns correspond to state 2 and so on. Note that we have $\sum_{i \in \mathcal{S}} |A_i| = d$.

For state s , if A_s is nonempty, define the feature matrix $\Phi_s = [I_{|A_s|} | 0] \in \mathbb{R}^{|A_s| \times n}$. Note that this would be the feature matrix corresponding to each state for the original feature matrix Φ .

Now, define $\tilde{\Delta}_s$ corresponding to each state s , and feature matrix Φ_s as in Equation (C.5). This corresponds to the set of permissible distributions under the function approximation setting for state s . If the set A_s is empty, we take $\tilde{\Delta}_s$ to be the singleton set with the uniform distribution.

Next, let the first d_1 columns of the Matrix Φ correspond to state s_{d_1} , i.e., $d_1 = |A_1|$, the next d_2 columns correspond to actions in state s_{d_2} and so on till finally the columns from $d - d_D + 1$ to column d correspond to state s_{d_D} , i.e., we partition the columns for which the feature vector is non-zero into the different states they correspond to. Therefore, D corresponds to the number of states for which there is at least one action for which the feature vector corresponding to the state action pair is non-zero. This means that the states $s_{d_1}, s_{d_2}, \dots, s_{d_D}$ are the only states for which there is at least one action with a nonzero feature vector, and therefore $\tilde{\Delta}_s$ is not a singleton set for these states. For all other states $s \in \mathcal{S} \setminus \{s_{d_1}, s_{d_2}, \dots, s_{d_D}\}$, we have that $\tilde{\Delta}_s$ is a singleton set containing the uniform distribution.

Now, the algorithm used to solve the Markov game in this function approximation setting, is similar to the tabular setting, except that we only have to run the inner iteration on the states s_{d_i} , $i = \{1, 2, \dots, D\}$. We describe the algorithm in detail in Algorithm 7 in §F. In order to characterize the point where the parameter converges to, we define:

$$\theta^* = [\theta_{s_{d_1}}^*, \theta_{s_{d_2}}^*, \dots, \theta_{s_{d_D}}^*], \quad \nu^* = [\nu_{s_{d_1}}^*, \nu_{s_{d_2}}^*, \dots, \nu_{s_{d_D}}^*], \quad (\text{F.16})$$

where

$$\theta_{s_{d_i}}^* = -[I_{|A_{s_{d_i}}|} | 0] \frac{Q_\tau^*(s_{d_i}) h_{\nu^*}(\cdot | s_{d_i})}{\tau} \in \mathbb{R}^{A_{s_{d_i}}}, \quad \nu_{s_{d_i}}^* = [I_{|A_{s_{d_i}}|} | 0] \frac{Q_\tau^*(s_{d_i})^\top g_{\theta^*}(\cdot | s_{d_i})}{\tau} \in \mathbb{R}^{A_{s_{d_i}}},$$

and Q_τ^* is the in-class NE Q-value (Definition 5.1).

The Nash equilibrium (θ^*, ν^*) under the function approximation setting satisfies:

$$V^{\theta^*, \nu}(s) \leq V^{\theta^*, \nu^*}(s) \leq V^{\theta, \nu^*}(s) \quad \forall \theta, \nu \in \mathbb{R}^d, \forall s \in \mathcal{S}. \quad (\text{F.17})$$

Note that this is equivalent to (from Lemma C.4):

$$V^{g_{\theta^*}, h_{\nu^*}}(s) \leq V^{g_{\theta^*}, h_{\nu^*}}(s) \leq V^{g_{\theta}, h_{\nu^*}}(s) \quad \forall g_{\theta}, h_{\nu} \in \tilde{\Delta}_s, \forall s \in \mathcal{S}. \quad (\text{F.18})$$

We denote the NE V (and Q) as V^* (and Q^*) (We use this notation for the general regularized version with $\tau \geq 0$, i.e., we allow for the unregularized game with $\tau = 0$), i.e.,

$$Q^*(s, a, b) = r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [V^*(s')]. \quad (\text{F.19})$$

Next, we define the soft-Bellman operator as:

$$\mathcal{T}_\tau(Q)(s, a, b) := r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} \left[\min_{\theta \in \mathbb{R}^{|A_s|}} \max_{\nu \in \mathbb{R}^{|A_s|}} f_\tau(Q(s'); g_\theta(\cdot | s'), h_\nu(\cdot | s')) \right], \quad (\text{F.20})$$

where

$$f_\tau(Q(s'); g_\theta(\cdot | s'), h_\nu(\cdot | s')) := g_\theta(\cdot | s)^\top Q(s) h_\nu(\cdot | s) - \tau \mathcal{H}(g_\theta(\cdot | s)) + \tau \mathcal{H}(h_\nu(\cdot | s)), \quad (\text{F.21})$$

and $Q(s)$ is the Q-value matrix at state s . Let θ_s, ν_s be the NE parameters at state s . Note that by *NE* we mean the solution to the min-max problem (which is the in class NE in the function approximation setting). Then the concatenation $\theta = [\theta_1^\top, \theta_2^\top, \dots, \theta_{|S|}^\top]^\top$ (and similarly for ν) denotes the parameters.

Note that the inner min max problem is equivalent to (from §C):

$$\min_{g_\theta \in \tilde{\Delta}_{s'}} \max_{h_\nu \in \tilde{\Delta}_{s'}} f_\tau(Q(s'); g_\theta(\cdot | s'), h_\nu(\cdot | s')). \quad (\text{F.22})$$

Consider the value iteration:

$$Q_{t+1} = \mathcal{T}_\tau(Q_t). \quad (\text{F.23})$$

We have $\|Q_t - Q^*\|_\infty \leq \gamma^t \|Q_0 - Q^*\|_\infty$, due to the non-expansiveness property of the min max operator and the contraction factor $\gamma < 1$. Note from Lemma 5.2, this fixed point in fact corresponds to the in-class NE Q -value matrix of the regularized Markov game. Thus, we have that the fixed point of the soft-Bellman operator defined in (F.20) is the NE Q -value matrix Q^* .

The inner problem, which solves the saddle point problem, is solved for each state with the input feature matrix Φ_s using Algorithm 2. The iteration complexity follows from a similar analysis to the tabular case. This completes the proof. \square

F.4 Proof of Lemma 5.2

Consider the operator

$$\mathcal{T}(V)(s) := \min_{\theta_s \in \mathbb{R}^{|A_s|}} \max_{\nu_s \in \mathbb{R}^{|A_s|}} \mathbb{E}_{a \sim g_{\theta_s}, b \sim h_{\nu_s}} \left[r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [V(s')] \right]. \quad (\text{F.24})$$

We put the subscript on the variables to show that it is for state s . Note that we state proof here for the unregularized game, i.e. with $\tau = 0$. The same proof also follows if we used the soft bellman operator defined in Equation (F.20) for regularized games.

Note that this is the operator for the value function V corresponding to the Q value operator in Equation (F.20) with $\tau = 0$. We have, from the nonexpansive property of the min-max operator (see for example Filar and Vrieze (2012)):

$$\|\mathcal{T}(V_1) - \mathcal{T}(V_2)\|_\infty \leq \gamma \|V_1 - V_2\|_\infty, \quad (\text{F.25})$$

which shows that this is a contracting operator and therefore has a unique fixed point by the Banach Fixed Point theorem. We show that this fixed point will lead to the in-class Nash equilibrium policy defined in Definition 5.1. Let the fixed point be denoted by V^* , and let (θ^*, ν^*) be the minimax policy parameters in Equation (F.24) when plugging in V^* . Note that $\theta^* = (\theta_s^*)_{s=1}^{|S|}$ and similarly ν^* . We will show that (θ^*, ν^*) is in fact the NE policy parameters. We have:

$$\begin{aligned} V^*(s) &= \mathbb{E}_{a \sim g_{\theta_s^*}, b \sim h_{\nu_s^*}} \left[r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [V^*(s')] \right] \\ &= \max_{\nu \in \mathbb{R}^d} \mathbb{E}_{a \sim g_{\theta_s^*}, b \sim h_{\nu_s}} \left[r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [V^*(s')] \right] \\ &\geq \mathbb{E}_{a \sim g_{\theta_s^*}, b \sim h_{\nu_s}} \left[r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [V^*(s')] \right] \quad \forall \nu \in \mathbb{R}^d. \end{aligned} \quad (\text{F.26})$$

Now applying Lemma 4.3.3 in Filar and Vrieze (2012) to Inequality (F.26), we have:

$$V^* \geq V^{\theta^*, \nu^*} \quad \forall \nu \in \mathbb{R}^d. \quad (\text{F.27})$$

Applying the same inequality for ν^* , we have:

$$V^{\theta^*, \nu} \leq V^* \leq V^{\theta, \nu^*} \quad \theta, \nu \in \mathbb{R}^d. \quad (\text{F.28})$$

Furthermore, since we have:

$$V^* = \mathbb{E}_{a \sim g_{\theta^*}, b \sim h_{\nu^*}} [r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [V^*(s')]]. \quad (\text{F.29})$$

Lemma 4.3.3 in [Filar and Vrieze \(2012\)](#) gives us that $V^* = V^{\theta^*, \nu^*}$. Combining this with Inequality (F.28), we have:

$$V^{\theta^*, \nu} \leq V^{\theta^*, \nu^*} \leq V^{\theta, \nu^*} \quad \forall \theta, \nu \in \mathbb{R}^d, \quad (\text{F.30})$$

which shows that (θ^*, ν^*) is the required NE. Finally, using Theorem 4.3.2 (iii) in [Filar and Vrieze \(2012\)](#), we can find the NE (θ^*, ν^*) for each state $s \in \mathcal{S}$ by solving the following matrix game:

$$\min_{\theta \in \mathbb{R}^{|A_s|}} \max_{\nu \in \mathbb{R}^{|A_s|}} \mathbb{E}_{a \sim g_{\theta}, b \sim h_{\nu}} [r(s, a, b) + \gamma \mathbb{E}_{s' \sim \mathbb{P}(\cdot | s, a, b)} [V^*(s')]].$$

This completes the proof. \square

G Simulations

We now provide simulation results to corroborate our theoretical results. First, we study matrix games under the tabular setting in Figure 2a. Here, we show the behavior of vanilla NPG and our proposed variant (Equations (3.1)-(3.2)). We plot the first element of the iterate, i.e., $\theta(1)$ on the y-axis. It is shown that even for vanishingly small stepsizes, vanilla NPG diverges, whereas the proposed variant converges even with reasonable step-size choices. The cost matrix Q is taken to be an identity matrix of dimension 5.

Next, we confirm the convergence of our variant of NPG in Figure 2b. In this figure, we compare the behavior of ONPG and NPG, and show that ONPG admits convergence for larger stepsizes. Smaller stepsizes that enables NPG convergence would lead to a slower convergence rate than ONPG. This is in line with our results in Theorems 3.2 and 3.4.

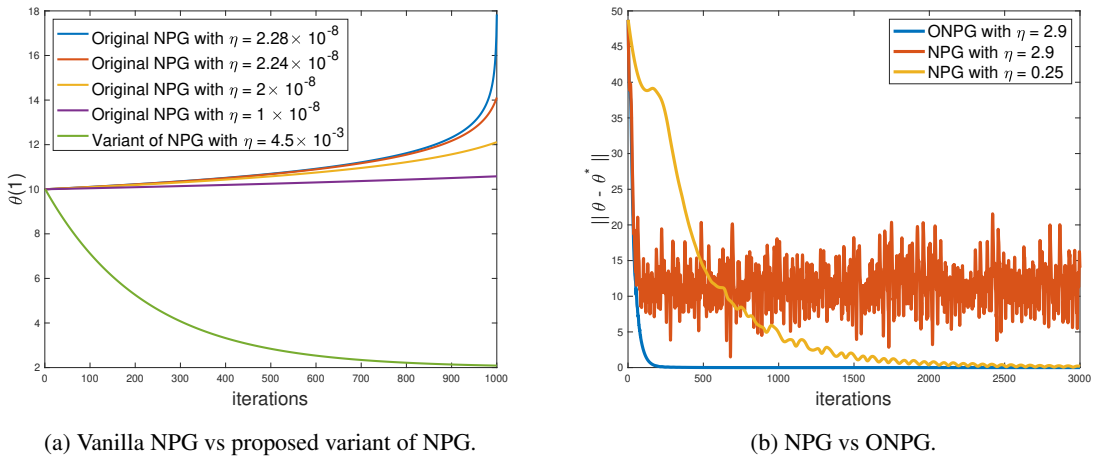
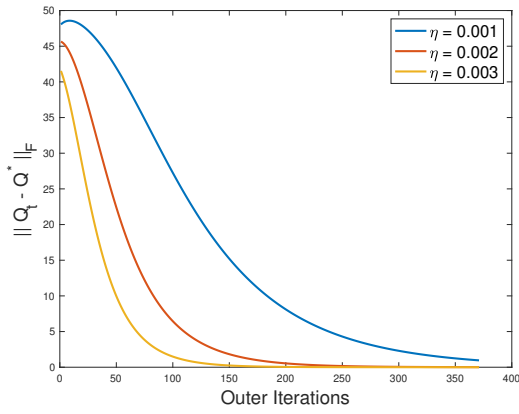


Figure 2: Comparison of vanilla NPG, proposed variant of NPG and ONPG in matrix game under the tabular setting, in terms of parameter convergence.


G.1 ONPG in Markov games with function approximation

Figures 3a and 3b study the behavior of Algorithm 7 in Markov games with log-linear function approximation, and corroborate the results of Theorem F.4. Here we take the feature matrix $\Phi \in \mathbb{R}^{10 \times 100}$, and $|\mathcal{S}| = 10$, i.e., there are 10 states. The first 10 columns correspond to the first action of each of the 10 states. This means that $\Phi_{(s,1)} = e_s$ for $s = \{1, 2, \dots, 10\}$ where e_s is a standard basis vector with element 1 at position s . We take the discount factor $\gamma = 0.8$. We take the transition

probability to be uniform for each state action pair, i.e., $\mathbb{P}(\cdot|s, a, b) = 1/10$ for all (s, a, b) , i.e., $\mathbb{P}(s'|s, a, b) = 1/10$ for all state $s' \in \mathcal{S}$. Finally, we take the regularization parameter $\tau = 0.1$.



(a) Q -matrix to in-class NE Q -matrix.



(b) Parameters to the in-class NE parameters.

Figure 3: Convergence in Markov Games with linear function approximation.