
Convergence to Nash Equilibrium and No-regret Guarantee in (Markov) Potential Games

Jing Dong¹

Baoxiang Wang¹

Yaoliang Yu²

¹ Chinese University of Hong Kong, Shenzhen

² University of Waterloo

Abstract

In this work, we study potential games and Markov potential games under stochastic cost and bandit feedback. We propose a variant of the Frank-Wolfe algorithm with sufficient exploration and recursive gradient estimation, which provably converges to the Nash equilibrium while attaining sublinear regret for each individual player. Our algorithm simultaneously achieves a Nash regret and a regret bound of $O(T^{4/5})$ for potential games, which matches the best available result, without using additional projection steps. Through carefully balancing the reuse of past samples and exploration of new samples, we then extend the results to Markov potential games and improve the best available Nash regret from $O(T^{5/6})$ to $O(T^{4/5})$. Moreover, our algorithm requires no knowledge of the game, such as the distribution mismatch coefficient, which provides more flexibility in its practical implementation. Experimental results corroborate our theoretical findings and underscore the practical effectiveness of our method.

1 Introduction

One of the most popular solution concepts in multi-agent settings is Nash equilibrium, which originates from game theory to describe the behaviors of rational, selfish players (Roughgarden, 2010). It characterizes a stable state between the players, where no individual has any incentive to unilaterally deviate from their chosen strategy. It is worth noting that the computational complexity of finding a Nash equilibrium

in a general game is known to be PPAD-Hard (Chen et al., 2009; Daskalakis, 2013). Nonetheless, it has been established that the computation of a Nash equilibrium becomes more feasible in specific game contexts, such as potential games (Monderer and Shapley, 1996), where a potential function is available to quantify how individual strategy changes affect collective utility.

In this paper, we study the potential game and its stochastic multi-step extension, the Markov potential game (Shapley, 1953; Monderer and Shapley, 1996). We focus on the decentralized dynamic between the players under repeated games with bandit feedback. Specifically, over a time horizon, each player independently determines their strategy at each step and subsequently receives cost feedback based on their chosen strategies. Then they use this cost feedback information to refine their strategy. Under this decentralized dynamic, it is not clear what is the optimal way for a selfish player to choose their strategy. Thus, beyond the objective of arriving at a Nash equilibrium, a player should try to minimize the cost experienced while assuming the other players are malicious. One notion that evaluates the performance of an algorithm under possible malicious adversaries is the regret, which originates from the online learning community (Shalev-Shwartz et al., 2012).

We are interested in designing learning algorithms that simultaneously enable the players to converge to a Nash equilibrium fast and with small regret. To measure the speed of convergence to Nash equilibrium, we use the notion of Nash regret, which is the cumulative gap to the Nash equilibrium. Designing such an algorithm confronts two main challenges. One is the nonstationarity for individual players, which arises as the environment changes when other players update their strategy, and consequently, the costs change. The other challenge is due to the bandit cost feedback. Since each player can only receive the cost feedback related to their chosen strategy, they must explore sufficiently alternative strategies to learn the Nash equilibrium.

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

Table 1: Summary of Nash regret bounds and regret for Markov potential game. Here, n denotes the number of players, m is the maximum size of action space among all players, $\max_{i \in [n]} |\mathcal{A}_i| = m$, S is the size of the state space, κ is the minimum stopping probability of the game, D_∞ captures the distribution mismatch and T is the time horizon. Note that when $S = 1$, the results are applicable to the potential game.

	Nash regret	Regret	Require Projection
Giannou et al. (2022)	Asymptotic	N/A	Yes
Zhang et al. (2021)	$\tilde{O}(n^{5/6} \text{poly}(\kappa^{-1}, S, m) T^{5/6})$	N/A	Yes
Leonardos et al. (2022)	$O\left(\frac{m^{5/6} D_\infty^{10/3} S^{5/3} T^{5/6}}{\kappa^{5/2}}\right)$	N/A	Yes
Ours	$O\left(\frac{n^{3/2} m D_\infty T^{4/5}}{\kappa^3} + \frac{n^{5/8} m^{3/4} S D_\infty T^{4/5}}{\kappa^{8/3}}\right)$	$O\left(\frac{mn^{3/2} S T^{4/5}}{\kappa^3}\right)$	No

For potential games, previous works have shown that no-regret learning algorithms can converge to the equilibrium fast when additional information about the cost function is assumed (Panageas et al., 2023) or when additional restrictions on the initial strategies of the players are imposed (Cohen et al., 2016; Giannou et al., 2021; Dong et al., 2023). For the stochastic multi-step extension of the potential game, the Markov potential game, it remains uncertain whether such an algorithm can be designed. Prior attempts for the Markov potential game resort to the projected gradient descent algorithm, which either relies on a sophisticated analysis through proximal point or additional assumptions on the cost function to bound the gradient estimation error (Leonardos et al., 2022; Ding et al., 2022). In both cases, the gradient estimation error can be large and make it difficult to balance fast convergence and exploration.

In this work, we introduce a variant of the Frank-Wolfe algorithm with exploration and recursive gradient estimation. The use of a recursive gradient reduces the estimation error, enabling our algorithm to converge to the Nash equilibrium fast and ensuring each player has sublinear regret. Our algorithm can be extended to the Markov potential game and attains both Nash regret and regret of $O(T^{4/5})$, where T is the length of the time horizon. Compared to previous decentralized algorithms for the Markov potential game, our algorithm improves the previous Nash regret of $O(T^{5/6})$ to $O(T^{4/5})$. Moreover, different from the previous algorithms, which are often based on projected gradient methods, our approach requires no additional projection step and therefore enjoys better computational complexity. Table 1 summarizes the results for Nash regret and regret for (Markov) potential games. The algorithm is then evaluated empirically through a Markov congestion game task, which validates our theoretical findings and demonstrates the practical effectiveness of our method.

2 Related Works

Potential game and congestion game Potential game and congestion game were introduced by Monderer and Shapley (1996) and Rosenthal (1973), respectively. Notably, these two classes of games can be discussed interchangeably due to their equivalence. For potential (congestion) games, there is a long line of work that studies the dynamic of no-regret learning algorithms (Cominetti et al., 2010; Palaiopanos et al., 2017; Heliou et al., 2017; Daskalakis et al., 2011; Syrgkanis et al., 2015; Chen and Peng, 2020; Hsieh et al., 2021). In particular, under the bandit feedback model, Heliou et al. (2017) showed the asymptotic last iterate convergence for the potential game and Cominetti et al. (2010); Palaiopanos et al. (2017) showed similar result for the congestion game. It has been shown that algorithms that are based on entropic regularization (e.g., exponential weights) can converge fast to the equilibria (Cohen et al., 2016; Giannou et al., 2021; Dong et al., 2023), albeit with certain requirements on the initialization. Recent works (Cui et al., 2022; Panageas et al., 2023) then established sublinear Nash regret for congestion game, with the latter additionally guaranteeing sublinear regret for each player.

Markov potential game The study of Markov games, which is a stochastic multi-step extension of the one-step game, is initiated by Shapley (1953). In stochastic optimal control, the study of the Markov potential game, which is the multi-step extension of the potential game, can be dated back to Dechert and O'Donnell (2006); González-Sánchez and Hernández-Lerma (2013). Recently there have been several works that establish the sample complexity of Nash equilibrium or Nash regret for the Markov potential game. Song et al. (2021) builds on the single agent algorithm Nash-VI and provides a $\Omega(\epsilon^{-2})$ sample complexity lower bound for finding ϵ -approximate Nash equilibrium. They then introduced Nash coordinate

ascent and established a $O(\epsilon^{-3})$ upper bound. However, their algorithm requires some coordination between the players. When each player updates their strategies independently and the exact gradient information is available (which may require knowledge of the cost function), Chen et al. (2022); Zhang et al. (2022); Guo et al. (2023); Sun et al. (2023) showed a sample complexity bound of $O(\epsilon^{-2})$ based on variations of gradient descent algorithms, which can be translated to $O(\sqrt{T})$ Nash regret. When the gradient information can only be estimated through interacting with the other players and the environment, Leonardos et al. (2022); Ding et al. (2022) gave sample complexity bounds of $O(\epsilon^{-6})$ (or $O(T^{5/6})$ Nash regret) and $O(\epsilon^{-5})$ (or $O(T^{4/5})$ Nash regret), respectively. However, both of their methods require solving a projection problem at each step as they are both based on the projected gradient descent. Moreover, the algorithm presented by Ding et al. (2022) relies on a regression subroutine and assumptions on cost structures, which is not realizable in the tabular Markov congestion game.

We summarize the existing algorithms when no exact gradient information is available in Table 1.

3 No-regret Learning for Potential Games

In this section, we discuss the formulation of potential games and our algorithm for no-regret learning. Some of the techniques are used in the next section, where we generalize the results to Markov potential games.

3.1 Potential Games

Let Γ be a strategic game with a finite number of players, $\mathcal{N} = \{1, 2, \dots, n\}$. Each player has a finite set of strategies, and we denote the strategy set of player i as \mathcal{A}_i . The cost function associated with player i is denoted by $c_i : \mathcal{A} \rightarrow [0, 1]$, where $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_n$ denotes the set of joint strategy profiles. Denote $m = \max_{i \in [n]} |\mathcal{A}_i|$. Given a joint strategy $a \in \mathcal{A}$, we write $a = (a_i, a_{-i})$, where a_i is the strategy of the i -th player and a_{-i} is the strategy of all other players. We refer to the deterministic strategy profile introduced above as a pure strategy. Alternatively, each player i can also play a randomized (mixed) strategy π_i according to playing probability distributions over their strategy set $\Delta(\mathcal{A}_i)$. Upon playing a mixed strategy $\pi = \pi_1 \times \dots \times \pi_n$, the expected cost incurred for player i is $\mathbb{E}_{a_i \sim \pi_i, a_{-i} \sim \pi_{-i}} [c_i(a_i, a_{-i})] = \langle \pi_i, c_i(\cdot, \pi_{-i}) \rangle =: c_i(\pi_i, \pi_{-i})$, where $c_i(\cdot, \pi_{-i}) = \mathbb{E}_{a_{-i} \sim \pi_{-i}} [c_i(\cdot, a_{-i})]$. In the stochastic setting, player i receives a sample $C_i(a)$ from a distribution with mean $c_i(a_i, a_{-i})$.

In potential games, Γ admits a potential function $\Phi :$

$\mathcal{A} \rightarrow \mathbb{R}$ such that for all $a_i, a'_i \in \mathcal{A}_i$, $a_{-i} \in \mathcal{A}_{-i}$,

$$c_i(a_i, a_{-i}) - c_i(a'_i, a_{-i}) = \Phi(a_i, a_{-i}) - \Phi(a'_i, a_{-i}). \quad (1)$$

We slightly abuse the notations and denote $\Phi(\pi) = \mathbb{E}_{a \sim \pi} [\Phi(a)]$. By the definition of the potential function, we have $\nabla_{\pi_i} c_i(\pi) = c_i(\cdot, \pi_{-i}) = \nabla_{\pi_i} \Phi(\pi)$, and $c_i(\pi_i, \pi_{-i}) - c_i(\pi'_i, \pi_{-i}) = \Phi(\pi_i, \pi_{-i}) - \Phi(\pi'_i, \pi_{-i})$. Clearly, every global minimum of Φ is a Nash equilibrium. As such, there always exists a pure strategy Nash equilibrium in a potential game.

Noticing that $\nabla_{\pi}^2 \Phi(\pi)$ is continuous, and followed by the compactness of $\Delta(\mathcal{A})$, the potential function is smooth.

Lemma 3.1 (Smoothness). *For any $\pi, \pi' \in \Delta(\mathcal{A})$, there exists an L such that $\|\nabla \Phi(\pi) - \nabla \Phi(\pi')\|_2 \leq L \|\pi - \pi'\|_2$.*

Learning protocol We study the potential game in the repeated game setting with bandit feedback. Throughout a time horizon T , at each time step $t \in [T]$, all players take a strategy a^t based on their mixed strategy π^t . Subsequently, they receive some feedback on their chosen strategies. Then they use this cost information to update their strategy. With bandit feedback, when the players take action $a^t \sim \pi^t$, they observe $C_i(a_i^t, a_{-i}^t)$ only.

Solution concepts A popular solution concept of the potential game is Nash equilibrium, which is defined as the following.

Definition 3.1 (Nash equilibrium). *A strategy $\pi^* = (\pi_1^*, \dots, \pi_n^*)$ is called a **Nash equilibrium** if for all players $i \in [n]$, it holds $c_i(\pi_i^*, \pi_{-i}^*) \leq c_i(\pi_i, \pi_{-i}^*)$, $\forall \pi_i \in \Delta(\mathcal{A}_i)$.*

In learning algorithms, the agents progressively find the Nash equilibrium by improving the approximation to the Nash equilibrium. The approximate Nash equilibrium is defined as follows.

Definition 3.2 (ϵ -approximate Nash equilibrium). *A strategy $\pi^* = (\pi_1^*, \dots, \pi_n^*)$ is called an **ϵ -approximate Nash equilibrium** if for all players $i \in [n]$, it holds $c_i(\pi_i^*, \pi_{-i}^*) \leq c_i(\pi_i, \pi_{-i}^*) + \epsilon$, $\forall \pi_i \in \Delta(\mathcal{A}_i)$.*

Such approximate Nash equilibria can be obtained by achieving a sublinear Nash regret, which is defined below.

Definition 3.3 (Nash regret). *Let $\{\pi^t\}_{t=1}^T$ be a sequence of mixed strategies played across a time horizon of T , the Nash regret after T time step is defined as*

$$\text{Nash-Regret}(T)$$

$$= \sum_{t=1}^T \max_{i \in [n]} \left(c_i(\pi_i^t, \pi_{-i}^t) - \min_{\pi_i \in \Delta(\mathcal{A}_i)} c_i(\pi_i, \pi_{-i}^t) \right).$$

By finding a sublinear expected Nash regret, one can converge to the Nash equilibrium in the sense of *best-iterate convergence*. This means there exists some $t \in [T]$, which π^t forms an approximate equilibrium in expectation. If one further finds the high probability of Nash regret, then it also implies that almost all iterations are ϵ -Nash equilibrium. We note that this concept of convergence is commonly used in the literature (Anagnostides et al., 2022; Leonardos et al., 2022; Ding et al., 2022; Cui et al., 2022; Panageas et al., 2023).

From an individual player's perspective, it is unclear how a selfish player should update their strategy. Thus, a player should try to minimize the cost experienced while assuming the worst costs possible with malicious adversaries. The performance of an online algorithm under malicious adversaries is captured by the notion of regret. We refer to the algorithms that achieve sublinear regret as no-regret algorithms.

Definition 3.4 (Regret of the i -th player). *For an arbitrary sequence of $\{\pi_{-i}^t\}_{t=1}^{T-1}$, the regret of an online algorithm is defined as*

$$\text{Regret}(T) = \sum_{t=1}^T c_i(\pi_i^t, \pi_{-i}^t) - \min_{\pi_i^* \in \Delta(\mathcal{A}_i)} \sum_{t=1}^T c_i(\pi_i^*, \pi_{-i}^t).$$

3.2 Algorithm and analysis for Potential Games

We first introduce a general algorithm we proposed for potential games. Our algorithm is largely based on the stochastic Frank-Wolfe method with a recursive gradient estimator (Zhang et al., 2020), but with a carefully tuned exploration step to achieve sublinear regret. At each step $t \in [T]$, player i plays a mixed strategy $\tilde{\pi}^t$, which is formed by mixing π^t with a uniformly random distribution s_i at a probability of μ (Line 3). Subsequently, upon receiving the instantaneous cost C_i^t , player i uses the important sampling method to construct a one-sample gradient estimate of $\nabla \pi_i \Phi(\tilde{\pi}^t)$ (Line 5). This estimate is then used to create a recursive gradient estimator (Line 6), which is then used in the subsequent Frank-Wolfe update steps (Line 7 and 8).

Previous efforts have explored the application of the Frank-Wolfe method to games (Cui et al., 2022). An inherent advantage of the Frank-Wolfe method lies in its ability to address the Frank-Wolfe gap, defined as $G(\pi^t) = \max_{\pi} (\pi^t - \pi)^\top \nabla \Phi(\pi^t)$. It minimizes this gap for each step, while a small cumulative Frank-Wolfe

Algorithm 1: Stochastic Frank-Wolfe algorithm with exploration for potential game

Input: Parameters η_t, ρ_t, μ .

- 1 Initialize π_i^1, d_i^0 , for $i \in [n]$
 - 2 **for** $t = 1, \dots, T$ **do**
 - 3 $\tilde{\pi}_i^t = (1 - \mu)\pi_i^t + \mu s_i$, where s_i is a uniformly random distribution on $\Delta(\mathcal{A}_i)$
 - 4 Play $a_i^t \sim \tilde{\pi}_i^t$ and receive cost C_i^t
 - 5 Let $(\hat{C}_i^t)_{a_i} = \frac{C_i^t \mathbb{1}\{a_i = a_i^t\}}{(\tilde{\pi}_i^t)_{a_i}}$, $\forall a_i \in \mathcal{A}_i$
 - 6 Update $d_i^t = (1 - \rho_t)d_i^{t-1} + \rho_t \hat{C}_i^t$
 - 7 Update $\hat{\pi}_i^{t+1} = \text{argmin}_{\pi_i \in \Delta(\mathcal{A}_i)} \langle \pi_i, d_i^t \rangle$
 - 8 Update $\pi_i^{t+1} = (1 - \eta_t)\pi_i^t + \eta_t \hat{\pi}_i^{t+1}$
-

gap implies a small Nash regret, because

$$\begin{aligned} & \max_{\pi'_i \in \Delta(\mathcal{A}_i)} \Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi'_i, \pi_{-i}^t) \\ &= \max_{\pi'_i \in \Delta(\mathcal{A}_i)} (\pi_i^t - \pi'_i)^\top \nabla_{\pi_i} \Phi(\pi_i^t, \pi_{-i}^t) \\ &\leq \max_{\pi'} (\pi^t - \pi')^\top \nabla \Phi(\pi^t) = G(\pi^t). \end{aligned}$$

In the case of bandits feedback, analyzing the Frank-Wolfe method typically requires bounding the gradient estimation error, i.e., $\|\nabla \Phi(\pi^t) - \hat{\nabla} \Phi(\pi^t)\|_2$. One of the most straightforward methods for obtaining a gradient estimate is through the importance sampling method (Line 5). Yet this only provides an error bound of $\|\nabla \Phi(\pi^t) - \hat{\nabla} \Phi(\pi^t)\|_2 = O\left(\frac{1}{\sqrt{\mu}}\right)$, where μ is the exploration parameter. The error would have accumulated to $O\left(\frac{T}{\sqrt{\mu}}\right)$ and resulted in large Nash regret. To overcome this, Cui et al. (2022) uses the regression method step to estimate the gradients, which results in an estimation error of roughly $\|\nabla \Phi(\pi^t) - \hat{\nabla} \Phi(\pi^t)\|_\infty = O(t^{-1/6})$ and a Nash regret of $O(T^{5/6})$. Our method instead uses a recursive gradient estimator d^t (Lines 6 and 7) that reuses past gradient estimates. This enables us to enjoy a smaller estimation error of $\|\nabla \Phi(\pi^t) - d^t\|_2 = O(t^{-1/5})$ and a subsequently smaller Nash regret of $O(T^{4/5})$. We also show that a $O(T^{4/5})$ regret can be attained simultaneously for each individual player.

Another way to address the gradient estimation error is to use the projected gradient descent method, which only requires bounded gradient estimates $\hat{\nabla} \Phi(\pi^t)$ (Panageas et al., 2023). It has been shown that a $O(T^{4/5})$ Nash regret and $O(T^{4/5})$ regret for each individual player can be attained simultaneously. However, when this method is extended to the Markov potential game, the Nash regret is only proved for $O(T^{5/6})$ and it is unclear whether individual players can enjoy sublinear regret simultaneously (Leonardos

et al., 2022). It can also be shown that the utilization of a recursive gradient estimate, which may seem advantageous in its ability to reuse previous gradients for more informative updates, offers no improvement when used alongside the projected gradient descent method (Yuan et al., 2016). Moreover, the projected gradient descent method requires an additional projection step at each iteration, which can be computationally expensive.

The following theorem establishes the performance guarantee for Algorithm 1.

Theorem 3.1 (Nash regret). *If all players run Algorithm 1 with $\eta_t = \frac{1}{\sqrt{nt^{4/5}}}$, $\mu = \min\{\frac{1}{mn}, \frac{m}{n^{1/8}T^{1/5}}\}$, $\rho_t = \frac{4\mu^{1/3}n^{1/3}m^{1/3}}{t^\alpha}$, $T \geq mn^{7/8}$, then $\mathbb{E}[\text{Nash-regret}(T)] = O((mn^{15/8}L + n^{5/2})T^{4/5})$.*

Remark 3.1. *When $T \geq O(n^{75/8}m^5L^5\epsilon^{-5})$, there exists $t \in \{1, \dots, T\}$, such that π^t is an ϵ -approximate Nash equilibrium.*

Remark 3.2. *Using Markov inequality on Theorem 3.1, it is immediate that the sublinear regret also holds with high probability.*

Theorem 3.2 investigates the case where one player chooses to play according to Algorithm 1 while other players play with arbitrary actions. It establishes an upper bound of the player's regret.

Theorem 3.2 (Regret for i -th player). *Let $\{\tilde{\pi}_{-i}^t\}_{t=1}^T$ be an arbitrary sequence of strategies for all players except for player i . For player i , running Algorithm 1 with $\eta_t = \frac{1}{\sqrt{nt^{4/5}}}$, $\mu = \min\{\frac{1}{mn}, \frac{m}{n^{1/8}T^{1/5}}\}$, $\rho_t = \frac{4\mu^{1/3}n^{1/3}m^{1/3}}{t^{8/15}}$ and $T \geq mn^{7/8}$, we can obtain a regret at most*

$$\mathbb{E} \left[\sum_{t=1}^T c_i(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \min_{\pi_i^* \in \Delta(\mathcal{A}_i)} \sum_{t=1}^T c_i(\pi_i^*, \tilde{\pi}_{-i}^t) \right] \leq O(mnLT^{4/5}).$$

Extension to congestion game We remark that both Theorem 3.1 and Theorem 3.2 can be extended to congestion games, which is a potential game with exponentially many actions (each action is a combination of some resources). In such case, Panageas et al. (2023) introduced an algorithm that can simultaneously obtain a Nash regret and a regret of $O(\log(m)T^{4/5})$, under an additional condition of the feedback (the semi-bandit feedback model). To obtain an efficient algorithm with Nash regret and regret guarantees that are only logarithmically dependent on m (which means polynomial in the number of resources), we can use the Carathéodory decomposition method to perform our algorithm on a polytope that can be described with $O(\log(m))$ fractional numbers. Additionally, under the bandit feedback model, we can estimate the

gradients with existing estimators from the combinatorial bandit literature (Combes et al., 2015). With these adjustments, our algorithm can obtain a regret of $O(\log(m)T^{4/5})$ for both Nash regret and regret for individual players.

While we defer the proofs of both theorems to the appendix, we present an outline of the proof.

Proof outline. [Theorem 3.1] By the definition of potential functions, we can first decompose the Nash regret into two parts, where the first part concerns the performance of the optimization algorithm, and the second part concerns regret caused by the exploration. That is,

$$\begin{aligned} & \text{Nash-regret}(T) \\ & \leq \sum_{i=1}^n \sum_{t=1}^T \max_{\pi_i' \in \Delta(\mathcal{A}_i)} \Phi(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i', \tilde{\pi}_{-i}^t) \\ & \leq \max_{\pi_i'} \Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi_i', \pi_{-i}^t) + 2n\mu, \end{aligned}$$

where the inequality is by the Lipschitzness of the potential function and the update rules of Algorithm 1.

By the smoothness of the potential function, one can show the following descent inequality

$$\begin{aligned} & \Phi(\pi^{t+1}) \\ & \leq \Phi(\pi^t) - \eta_t G(\pi^t) + 2\eta_t \sqrt{n} \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{\eta_t^2 nL}{2}, \end{aligned}$$

where $G(\pi^t) = \max_{\pi} (\pi^t - \pi)^\top \nabla \Phi(\pi^t)$.

By rearranging the terms and taking summation over T , we can upper bound $\sum_{t=1}^T G(\pi^t)$ as,

$$\begin{aligned} & \sum_{t=1}^T G(\pi^t) \\ & \leq \frac{2n}{\eta_{T+1}} + 2\sqrt{n} \sum_{t=1}^T \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{nL}{2} \sum_{t=1}^T \eta_t. \end{aligned}$$

Notice that this is the upper bound of the Nash regret, as $\max_{\pi_i' \in \Delta(\mathcal{A}_i)} \Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi_i', \pi_{-i}^t) \leq G(\pi^t)$. It now only amounts to bounding the gradient estimation error $\|\nabla \Phi(\pi^t) - d^t\|_2$, which is derived in the following lemma through an induction argument.

Lemma 3.2. *Let $\eta_t = \frac{1}{\sqrt{nt^{3\alpha/2}}}$, $\rho_t = \frac{4\mu^{1/3}n^{1/3}m^{1/3}}{t^\alpha}$, $\mu \geq \frac{1}{mn}$, $\alpha, \beta \in (0, 1)$ in Algorithm 1. We have*

$$\mathbb{E} [\|\nabla \Phi(\pi^t) - d^t\|_2] \leq O\left(\mu\sqrt{n}L + \frac{m^{4/3}n^{4/3}}{\mu^{1/3}(t+1)^{\alpha/2}}\right).$$

The theorem follows by selecting the appropriate parameters α, β . \square

Proof outline. [Theorem 3.2] Using the definition of potential functions and the update rules of Algorithm 1, we can decompose the regret into two parts, where one is due to the exploration and one is due to suboptimal strategies, as

$$\begin{aligned} & \sum_{t=1}^T c_i(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - c_i(\pi_i^*, \tilde{\pi}_{-i}^t) \\ & \leq \mu T + (\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)). \end{aligned}$$

Then, by the smoothness and linearity of Φ , we obtain the following recurrence relationship

$$\begin{aligned} & \Phi(\pi_i^{t+1}, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t) \\ & \leq (1 - \eta_t) (\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)) \\ & \quad + 2\eta_t \|\nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - d_i^t\|_2 + \frac{\eta_t^2 L}{2}. \end{aligned}$$

Using proof by induction and by Lemma 3.2, we show that $\mathbb{E} [\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)] = O(1/t^{1/5})$. Lastly, putting everything together yields Theorem 3.2. \square

4 No-regret Learning for Markov Potential Games

In this section, we extend the algorithm and results to the Markov potential game, which is a natural stochastic extension of the potential game (Shapley, 1953; Leonardos et al., 2022).

In a Markov potential game, n agents interact in a finite horizon MDP $\langle \mathcal{S}, (\mathcal{A}_i, c_i)_{i=1}^n, P, \mu_0, \kappa \rangle$. Here, \mathcal{S} represents the state space, and for each player i , \mathcal{A}_i and c_i denote the strategy space and cost functions, respectively. We define the joint action space as $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$, and the transition function is given by $P : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$. This means that $P(\cdot | s, a)$ represents a probability measure over the state space \mathcal{S} when the agents at state $s \in \mathcal{S}$ jointly take action $a \in \mathcal{A}$. Additionally, μ_0 denotes the initial state distribution and $\kappa = \min_{s,a} \{\kappa_{s,a}\}$, where $\kappa_{s,a}$ is the probability that the game stops at state-action pair (s, a) . We assume $\kappa > 0$.

For each player $i \in [n]$, a stochastic policy is denoted as $\pi_i \in \Delta(\mathcal{A}_i)^{\mathcal{S}}$, which is a mapping that associates a state $s \in \mathcal{S}$ with a probability distribution over actions in $\Delta(\mathcal{A}_i)$. Lastly, we define m as the maximum cardinality among the strategy spaces of the individual players, i.e., $m = \max_{i \in [n]} |\mathcal{A}_i|$.

Learning protocol The players learn their policy through T episodes. Within each episode, at step h , all players observe the state s^h . Then, each player i plays $a_i^h \sim \pi_i^t$ and receives a cost C_i^h sampled from a distribution with mean $c_i(s^h, a_i^h, a_{-i}^h)$. Without loss

of generality, we assume the costs C_i^h are bounded between $[0, 1]$. Based on $a^h = (a_1^h, \dots, a_n^h)$, the players transit to the next state $s^{h+1} \sim P(\cdot | s^h, a^h)$. With probability κ_{s^h, a^h} , the game stops at step h . As $\kappa > 0$, the game will eventually stop with probability 1. We define the random variable H to be the final step before the game terminates. Additionally, we use H_t to indicate the number of steps taken in episode t .

Value function and potential function For player $i \in [n]$, the value function $V^i(s) : \Pi \rightarrow \mathbb{R}$ represents the expected cost when $s_0 = s$ and the players play strategy, $(a^t)_{t \geq 0}$, $a^t \sim \pi^t, \pi^t = (\pi_i^t, \pi_{-i}^t)$. Specifically, it is defined as $V_i^\pi(s) := \mathbb{E}_{\pi, P, H} \left[\sum_{h=0}^H C_i^h | s_0 = s \right]$. We slightly abuse the notation and denote $V_i^\pi(\mu) = \mathbb{E}_{s \sim \mu} [V_i^\pi(s)]$.

We define the state-action value function for player i as $Q_i^\pi(s, a)$, which denotes the expected cost when started from $s_0 = s, a_0 = a$,

$$Q_i^\pi(s, a) = \mathbb{E}_{\pi, P, H} \left[\sum_{h=0}^H C_i^h | s_0 = s, a_0 = a \right].$$

We also denote $Q_i^\pi(\mu, a) = \mathbb{E}_{s \sim \mu} [Q_i^\pi(s, a)]$. Notice that the value function can be equivalently expressed as

$$V_i^\pi(s) = \sum_{a \in \mathcal{A}} \pi(a | s) Q_i^\pi(s, a).$$

In the Markov potential game, there exists a global potential function Φ , that captures the incentive of all players to change their strategies at any state. Specifically, at any state $s \in \mathcal{S}$, $\pi_i, \pi'_i, i \in [n]$, the potential function satisfies, $V_i^{\pi_i, \pi_{-i}}(s) - V_i^{\pi'_i, \pi_{-i}}(s) = \Phi_s(\pi_i, \pi_{-i}) - \Phi_s(\pi'_i, \pi_{-i})$. By linearity of expectations, we also have $V_i^{\pi_i, \pi_{-i}}(\mu) - V_i^{\pi'_i, \pi_{-i}}(\mu) = \Phi_\mu(\pi_i, \pi_{-i}) - \Phi_\mu(\pi'_i, \pi_{-i})$, where $\Phi_\mu(\pi_i, \pi_{-i}) = \mathbb{E}_{s \sim \mu} [\Phi_s(\pi_i, \pi_{-i})]$.

Solution concepts Similar to the one-step potential game, the approximate Nash equilibrium is defined as follows

Definition 4.1 (ϵ -Nash equilibrium). *A strategy $\pi^* = (\pi_1^*, \dots, \pi_n^*)$ is called an ϵ -approximate Nash equilibrium if for all player $i \in [n]$, it holds $V_i(\pi_i^*, \pi_{-i}^*) \leq V_i(\pi_i, \pi_{-i}^*)$, $\forall \pi_i \in \Delta(\mathcal{A}_i)^{\mathcal{S}}$. When $\epsilon = 0$, we call it a **Nash equilibrium**.*

Such approximate Nash equilibrium can be obtained by achieving sublinear Nash regret, which is defined as follows.

Definition 4.2 (Nash regret in MPG). *Let $\{\pi^t\}$ be a sequence of stochastic policies played by the players across T episodes, the Nash regret after T episodes are defined as $\text{Nash-regret}(T) = \sum_{t=1}^T \max_{i \in [n]} \left(V_i^{\pi_i^t, \pi_{-i}^t}(\mu) - \min_{\pi_i} V_i^{\pi_i, \pi_{-i}^t}(\mu) \right)$.*

For a selfish individual player, another important metric for playing against potentially malicious adversaries is regret, which is defined as follows.

Definition 4.3 (Regret in MPG). *For any $i \in [n]$, and an arbitrary sequence of $\{\pi_{-i}^t\}_{t=1}^T$, the individual regret of player i is defined as $\sum_{t=1}^T \left(V_i^{\pi_i^t, \pi_{-i}^t}(\mu) - \min_{\pi_i} \sum_{t=1}^T V_i^{\pi_i, \pi_{-i}^t}(\mu) \right)$.*

4.1 Algorithm for Markov potential game

To apply Algorithm 1 to the Markov potential game, we need to build an unbiased one-sample gradient estimate. This may be constructed from a single or more steps of actions. We use the following estimator, which can be demonstrated to be unbiased and has bounded norms:

$$\widehat{\nabla}_{\pi_i^t} = \sum_{h=0}^{H_t} C_i^h \sum_{h=0}^{H_t} \nabla \log \pi_i^t(a_i^h | s^h). \quad (2)$$

With this, we can construct the recursive gradient estimator and perform the Frank-Wolfe update step. The overall algorithm is described in Algorithm 2.

Algorithm 2: Stochastic Frank-Wolfe with exploration for Markov potential game

Input: Parameters η_t, ρ_t, μ .

- 1 Initialize π_i^1, d_i^0 , for $i \in [n]$
 - 2 **for** $t = 1, \dots, T$ **do**
 - 3 $\tilde{\pi}_i^t = (1 - \mu)\pi_i^t + \mu s$, where s is a uniformly random distribution on $\Delta(\mathcal{A}_i)$
 - 4 **for** $h = 0, \dots, H$ **do**
 - 5 Play $a_i^h \sim \tilde{\pi}_i^t$ and receive rewards C_i^h
 - 6 $\widehat{\nabla}_{\pi_i^t} = \sum_{h=0}^{H_t} C_i^h \sum_{h=0}^{H_t} \nabla \log \pi_i^t(a_i^h | s^h)$
 - 7 Update $d_i^t = (1 - \rho_t)d_i^{t-1} + \rho_t \widehat{\nabla}_{\pi_i^t}$
 - 8 Update $\hat{\pi}_i^{t+1} = \operatorname{argmin}_{\pi_i \in \Delta(\mathcal{A}_i)} \langle \pi_i, d_i^t \rangle$
 - 9 Update $\pi_i^{t+1} = (1 - \eta_t)\pi_i^t + \eta_t \hat{\pi}_i^{t+1}$
-

Equipped with Algorithm 2, the following theorem establishes the Nash regret for Markov potential game.

Theorem 4.1 (Nash regret for Markov potential game). *If all players run Algorithm 2 with $\eta_t = \frac{\kappa}{n^{3/2}m^{4/5}}$, $\rho_t = \frac{4\mu^{1/3}}{n^{1/3}m^{2/3}t^{8/15}}$ and $\mu = \min\{\frac{1}{m^2n}, \frac{\kappa^{4/3}}{n^{7/8}m^{1/4}T^\beta}\}$, when $T \geq m^{7/4}n^{1/8}\kappa^{4/3}$, we have*

$$\begin{aligned} & \mathbb{E} [\text{Nash-regret}(T)] \\ & \leq O\left(\frac{n^{3/2}mD_\infty T^{4/5}}{\kappa^3} + \frac{n^{5/8}m^{3/4}SD_\infty T^{4/5}}{\kappa^{25/9}}\right), \end{aligned}$$

where $D_\infty = \max_i \left\| \frac{d^{\pi_i^*, \pi_{-i}}(\mu)}{\mu} \right\|_\infty$.

Remark 4.1. *When $T \geq O\left(\frac{n^{17/2}m^5D_\infty^5}{\kappa^{15}\epsilon^5} \frac{n^{15/8}m^{15/4}D_\infty^5S^5}{\kappa^{25/3}\epsilon^5}\right)$, there exists $t \in \{1, \dots, T\}$, such that π^t is an ϵ -equilibrium.*

Remark 4.2. *Using Markov inequality, it is immediate that Theorem 4.1 also holds with high probability.*

Our bound relies on a distribution mismatch coefficient $D_\infty = \left\| \frac{d^{\pi_i^*, \pi_{-i}}(\mu)}{\mu} \right\|_\infty$, which is an extension to the distribution mismatch coefficient ($D_\infty := \max_s (d_{s_0}^{\pi_s^*}(s)/\mu(s))$) (Agarwal et al., 2021). In the single-agent case, the coefficient is large whenever an MDP is difficult to explore, which may manifest in two cases: 1) a state s is important and the optimal policy needs to visit the state a lot, while the starting distribution rarely starts from this, or 2) a state s is unimportant and the optimal policy rarely spends time on this state, but it is likely to start from this state. These are extended to the multi-agent setting with $d_{s_0}^{\pi_s^*, \pi_{-i}}(s)$ defined instead, which means the time spent on state s by player i taking the optimal policy π_i^* while others play π_{-i} , to capture the exploration difficulty for individual player.

Compared to previous results of $O(\epsilon^{-6})$ on Markov potential game, our results improve them by $O(\epsilon^{-1})$ (Leonardos et al., 2022). Assuming the value function admits a linear instruction (i.e. $Q_i^\pi(s, a) = \langle \phi(s, a), \omega_i \rangle$, where ϕ is a known feature mapping and ω_i is a learnable parameter), Ding et al. (2022) established a $O(\epsilon^{-5})$ result under a regression oracle. However, the oracle can only be implemented for such a result when $\phi^\top \phi$ is invertible. With a tabular Markov potential game (where the states and actions are finite), ϕ may be a one-hot basis vector, and thus the assumption is not satisfied.

We also remark that our choice of parameters η_t, ρ_t, μ does not depend on D_∞ and hence requires no knowledge of the exploration difficulty. In contrast, the previous bounds crucially rely on the knowledge of D_∞ to choose the parameters (Zhang et al., 2021; Leonardos et al., 2022; Ding et al., 2022).

The following theorem establishes the regret guarantee for individual players, even when the other players are playing arbitrarily.

Theorem 4.2 (Regret for the i -th player in Markov potential game). *Let $\{\tilde{\pi}_{-i}^t\}_{t=1}^T$ be an arbitrary sequence of strategies for all players except for player i . For player i , run Algorithm 2 with $\eta_t = \frac{\kappa}{n^{3/2}m^{4/5}}$, $\rho_t = \frac{4\mu^{1/3}}{n^{1/3}m^{2/3}t^{8/15}}$ and $\mu = \min\{\frac{1}{m^2n}, \frac{\kappa^{4/3}}{n^{7/8}m^{1/4}T^\beta}\}$, when $T \geq m^{7/4}n^{1/8}\kappa^{4/3}$. Then the regret of player i is at most*

$$\mathbb{E} \left[\sum_{t=1}^T \left(V_i^{\tilde{\pi}_{-i}^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^*, \tilde{\pi}_{-i}^t}(\mu) \right) \right]$$

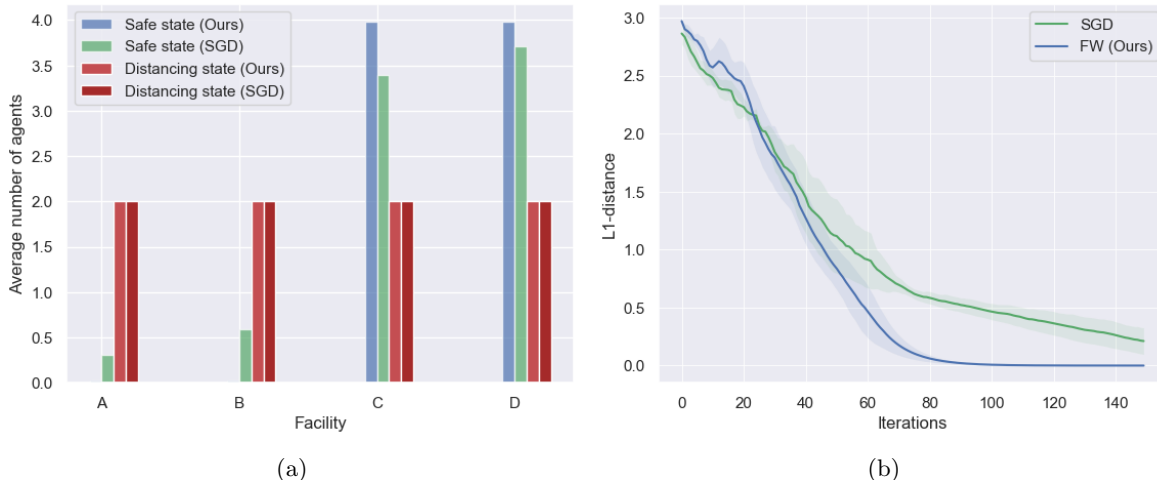


Figure 1: Figure 1a shows the final converged policy on each of the states. Figure 1b shows the convergence of the algorithms by L_1 distance to the final strategy.

$$\leq O\left(\frac{m^{4/3}n^{27/24}ST^{4/5}}{\kappa^{7/3}}\right).$$

5 Experiments

To validate our theoretical results, we implemented our algorithm on an example of the Markov congestion game, taken from Leonardos et al. (2022).

Experiments setup In the example game, we have $n = 8$ players, $\mathcal{A}_i = 4$ actions for player i to choose from, and $\mathcal{S} = \{\text{safe, distancing}\}$ as the set of states. In both states, each player’s best strategy is to be in the same facility with other players to minimize the cost. In particular, the cost of each facility is determined by the number of players choosing it, whereas the facility preference is $D > C > B > A$ as they have increasing costs in that order. However, when more than $n/2 = 4$ players choose the same facility, the state will be transitioned to the distancing state. The cost structure is the same in the distancing state, but each player will be penalized more by a constant amount (by 100 times more). To return to the safe state, the players will need to learn to distribute evenly across the facilities, with no more than $n/4 = 2$ players on the same facility. The game stops at each step with a probability 0.99. However, we limit the maximum length of the game to $H_t = 20$. When the transition is deterministic, it is clear that the optimal strategy for the players is to choose facility C and D with 4 players each in the safe state and to choose each of the facilities evenly in the distancing state.

We note that all the results shown are the averaged results obtained from 5 runs with different random seeds; the shades in Figure 1b denote 1 standard deviation

from the mean. All experiments are conducted with a 10 core CPU and 16 GB RAM.

Implementation and results We implemented the projected stochastic gradient descent method (SGD) proposed by Leonardos et al. (2022) and our stochastic Frank-Wolfe with exploration. The algorithms are run with $T = 150$ iterations, and each update is performed with 10 trajectories of fixed length $H_t = 20$. For the SGD method, we kept its learning rate as $\eta = 0.0001$ as specified in their paper. For stochastic FW with exploration, we choose $\eta = 0.1$, $\rho = \frac{0.9}{(t+1)^{3/5}}$ and $\mu = 0.001$.

Figure 1a shows the final learned strategy of each of the algorithms. As shown in the figure, our method learns the optimal strategy in both states after the algorithm iterations. In Figure 1b, we show the convergence of the algorithms by plotting the distance (in L_1 distance) of the strategy to the final strategy. In comparison, our algorithm converges much faster than the SGD method.

6 Conclusion

In this work, we studied potential games and Markov potential games with stochastic cost and bandit feedback and introduced an algorithm based on the celebrated Frank-Wolfe algorithm. Our algorithm converges to Nash equilibrium fast with a Nash regret of $O(T^{4/5})$, while achieving sublinear regret for each player. This improves over the previous Nash regret result for the Markov potential game and is the first algorithm to simultaneously achieve small Nash regret and regret. Experimental results validate the effectiveness of our algorithms.

Acknowledgement

We thank the reviewers and the area chair for constructive feedback. YY gratefully acknowledges funding support from NSERC and the Canada CIFAR AI chairs program. Baoxiang Wang is partially supported by the National Natural Science Foundation of China (62106213, 72150002, 72394361) and the Shenzhen Science and Technology Program (RCBS20210609104356063, JCYJ20210324120011032).

References

- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. (2021). On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *The Journal of Machine Learning Research*, 22(1):4431–4506.
- Anagnostides, I., Panageas, I., Farina, G., and Sandholm, T. (2022). On last-iterate convergence beyond zero-sum games. In *International Conference on Machine Learning*.
- Chen, D., Zhang, Q., and Doan, T. T. (2022). Convergence and price of anarchy guarantees of the softmax policy gradient in markov potential games. In *International Conference on Machine Learning: Workshop AI for Agent-Based Modelling*.
- Chen, X., Deng, X., and Teng, S.-H. (2009). Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57.
- Chen, X. and Peng, B. (2020). Hedging in games: Faster convergence of external and swap regrets. In *Advances in Neural Information Processing Systems*.
- Cohen, J., Héliou, A., and Mertikopoulos, P. (2016). Exponentially fast convergence to (strict) equilibrium via hedging. arXiv preprint arXiv:1607.08863.
- Combes, R., Talebi Mazraeh Shahi, M. S., Proutiere, A., et al. (2015). Combinatorial bandits revisited. In *Advances in neural information processing systems*.
- Cominetti, R., Melo, E., and Sorin, S. (2010). A payoff-based learning procedure and its application to traffic games. *Games and Economic Behavior*, 70(1):71–83.
- Cui, Q., Xiong, Z., Fazel, M., and Du, S. S. (2022). Learning in congestion games with bandit feedback. In *Advances in Neural Information Processing Systems*.
- Daskalakis, C. (2013). On the complexity of approximating a Nash equilibrium. *ACM Transactions on Algorithms (TALG)*, 9(3):1–35.
- Daskalakis, C., Deckelbaum, A., and Kim, A. (2011). Near-optimal no-regret algorithms for zero-sum games. In *ACM-SIAM Symposium on Discrete Algorithms*.
- Daskalakis, C., Foster, D. J., and Golowich, N. (2020). Independent policy gradient methods for competitive reinforcement learning. In *Advances in neural information processing systems*.
- Dechert, W. D. and O’Donnell, S. (2006). The stochastic lake game: A numerical solution. *Journal of Economic Dynamics and Control*, 30(9-10):1569–1587.
- 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*.
- Dong, J., Wu, J., Wang, S., Wang, B., and Chen, W. (2023). Taming the exponential action set: Sub-linear regret and fast convergence to nash equilibrium in online congestion games. arXiv preprint arXiv:2306.13673.
- Giannou, A., Lotidis, K., Mertikopoulos, P., and Vlatakis-Gkaragkounis, E.-V. (2022). On the convergence of policy gradient methods to Nash equilibria in general stochastic games. *Advances in Neural Information Processing Systems*.
- Giannou, A., Vlatakis-Gkaragkounis, E.-V., and Mertikopoulos, P. (2021). On the rate of convergence of regularized learning in games: From bandits and uncertainty to optimism and beyond. In *Advances in Neural Information Processing Systems*.
- González-Sánchez, D. and Hernández-Lerma, O. (2013). *Discrete-time Stochastic Control and Dynamic Potential Games: the Euler–Equation approach*. Springer Science & Business Media.
- Guo, X., Li, X., Maheshwari, C., Sastry, S., and Wu, M. (2023). Markov α -potential games: Equilibrium approximation and regret analysis. arXiv preprint arXiv:2305.12553.
- Héliou, A., Cohen, J., and Mertikopoulos, P. (2017). Learning with bandit feedback in potential games. In *Advances in Neural Information Processing Systems*.
- Hsieh, Y.-G., Antonakopoulos, K., and Mertikopoulos, P. (2021). Adaptive learning in continuous games: Optimal regret bounds and convergence to Nash equilibrium. In *Conference on Learning Theory*.
- Leonardos, S., Overman, W., Panageas, I., and Piliouras, G. (2022). Global convergence of multi-agent policy gradient in Markov potential games. In *International Conference on Learning Representations*.

- Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and economic behavior*, 14(1):124–143.
- Palaiopanos, G., Panageas, I., and Piliouras, G. (2017). Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. In *Advances in Neural Information Processing Systems*.
- Panageas, I., Skoulakis, S., Viano, L., Wang, X., and Cevher, V. (2023). Semi bandit dynamics in congestion games: Convergence to Nash equilibrium and no-regret guarantees. In *International Conference on Machine Learning*.
- Rosenthal, R. W. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67.
- Roughgarden, T. (2010). Algorithmic game theory. *Communications of the ACM*, 53(7):78–86.
- Shalev-Shwartz, S. et al. (2012). Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194.
- Shapley, L. S. (1953). Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100.
- Song, Z., Mei, S., and Bai, Y. (2021). When can we learn general-sum markov games with a large number of players sample-efficiently? In *International Conference on Learning Representations*.
- Sun, Y., Liu, T., Zhou, R., Kumar, P. R., and Shahrampour, S. (2023). Provably fast convergence of independent natural policy gradient for markov potential games. In *Conference on Neural Information Processing Systems*.
- Syrkkanis, V., Agarwal, A., Luo, H., and Schapire, R. E. (2015). Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems*.
- Yuan, K., Ying, B., and Sayed, A. H. (2016). On the influence of momentum acceleration on online learning. *The Journal of Machine Learning Research*, 17(1):6602–6667.
- Zhang, M., Shen, Z., Mokhtari, A., Hassani, H., and Karbasi, A. (2020). One sample stochastic Frank-Wolfe. In *International Conference on Artificial Intelligence and Statistics*.
- Zhang, R., Mei, J., Dai, B., Schuurmans, D., and Li, N. (2022). On the global convergence rates of decentralized softmax gradient play in Markov potential games. In *Advances in Neural Information Processing Systems*.
- Zhang, R., Ren, Z., and Li, N. (2021). Gradient play in stochastic games: Stationary points, convergence, and sample complexity. arXiv preprint arXiv:2106.00198.

Checklist

The checklist follows the references. For each question, choose your answer from the three possible options: Yes, No, Not Applicable. You are encouraged to include a justification to your answer, either by referencing the appropriate section of your paper or providing a brief inline description (1-2 sentences). Please do not modify the questions. Note that the Checklist section does not count towards the page limit. Not including the checklist in the first submission won't result in desk rejection, although in such case we will ask you to upload it during the author response period and include it in camera ready (if accepted).

In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes/No/Not Applicable]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Appendix for “Convergence to Nash Equilibrium and No-regret Guarantee in (Markov) Potential Games”

A Analysis for potential game

As each player shares the same configurations of parameters, we drop the index i and rewrite the update steps as

$$\begin{aligned}
\tilde{\pi}^t &= (1 - \mu)\pi^t + \mu s \\
(\hat{C}^t)_a &= \frac{C^t \cdot \mathbb{I}\{a = a^t\}}{(\tilde{\pi}^t)_a}, \quad \forall a \in \mathcal{A}, \quad d^t = (1 - \rho_t)d^{t-1} + \rho_t \hat{C}^t \\
\hat{\pi}^{t+1} &= \operatorname{argmin}_{\pi \in \Delta(\mathcal{A})} \langle \pi, d^t \rangle \\
\pi^{t+1} &= (1 - \eta_t)\pi^t + \eta_t \hat{\pi}^{t+1}.
\end{aligned}$$

A.1 Proof of Theorem 3.1

Theorem 3.1 (Nash regret). *If all players run Algorithm 1 with $\eta_t = \frac{1}{\sqrt{nt^{4/5}}}$, $\mu = \min\left\{\frac{1}{mn}, \frac{m}{n^{1/8}T^{1/5}}\right\}$, $\rho_t = \frac{4\mu^{1/3}n^{1/3}m^{1/3}}{t^\alpha}$, $T \geq mn^{7/8}$, then $\mathbb{E}[\text{Nash-regret}(T)] = O((mn^{15/8}L + n^{5/2})T^{4/5})$.*

Proof. We start by decomposing the regret as,

$$\begin{aligned}
\text{Nash-regret}(T) &= \sum_{t=1}^T \max_{i \in [n]} \left(c_i(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \min_{\pi'_i \in \Delta(\mathcal{A}_i)} c_i(\pi'_i, \tilde{\pi}_{-i}^t) \right) \\
&\leq \sum_{i=1}^n \sum_{t=1}^T \max_{\pi'_i \in \Delta(\mathcal{A}_i)} (c_i(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - c_i(\pi'_i, \tilde{\pi}_{-i}^t)) \\
&= \sum_{i=1}^n \sum_{t=1}^T \max_{\pi'_i \in \Delta(\mathcal{A}_i)} \Phi(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi'_i, \tilde{\pi}_{-i}^t),
\end{aligned}$$

where the last equality is by the definition of potential function.

Denote $\pi_i^{t'} = \operatorname{argmax}_{\pi'_i \in \Delta(\mathcal{A}_i)} \Phi(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi'_i, \tilde{\pi}_{-i}^t)$, we have

$$\begin{aligned}
&\max_{\pi'_i} \Phi(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi'_i, \tilde{\pi}_{-i}^t) \\
&= (\Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi_i^{t'}, \pi_{-i}^t)) + (\Phi(\tilde{\pi}^t) - \Phi(\pi^t)) + (\Phi(\pi_i^{t'}, \pi_{-i}^t) - \Phi(\pi_i^{t'}, \tilde{\pi}_{-i}^t)) \\
&\leq \max_{\pi'_i} \Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi'_i, \pi_{-i}^t) + (\Phi(\tilde{\pi}^t) - \Phi(\pi^t)) + (\Phi(\pi_i^{t'}, \pi_{-i}^t) - \Phi(\pi_i^{t'}, \tilde{\pi}_{-i}^t)) \\
&\leq \max_{\pi'_i} \Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi'_i, \pi_{-i}^t) + \sqrt{n}\|\tilde{\pi}^t - \pi^t\|_2 + \sqrt{n}\|\pi_{-i}^t - \tilde{\pi}_{-i}^t\|_2 \\
&\leq \max_{\pi'_i} \Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi'_i, \pi_{-i}^t) + \mu\sqrt{n}\|\pi^t - s\|_2 + \mu\sqrt{n}\|\pi_{-i}^t - s_{-i}\|_2 \\
&\leq \max_{\pi'_i} \Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi'_i, \pi_{-i}^t) + 2n\mu,
\end{aligned}$$

where the second inequality is by the Smoothness of Φ and the third inequality is by the update rule $\tilde{\pi}^t = (1 - \mu)\pi^t + \mu s$.

Define $G(\pi^t) = \max_{\pi} (\pi^t - \pi)^\top \nabla \Phi(\pi^t)$. By the smoothness of Φ , we have

$$\begin{aligned}
 \Phi(\pi^{t+1}) &\leq \Phi(\pi^t) + (\pi^{t+1} - \pi^t)^\top \nabla \Phi(\pi^t) + \frac{L}{2} \|\pi^{t+1} - \pi^t\|_2^2 \\
 &\leq \Phi(\pi^t) + \eta_t (\hat{\pi}^{t+1} - \pi^t)^\top \nabla \Phi(\pi^t) + \frac{\eta_t^2 L}{2} \|\hat{\pi}_i^{t+1} - \pi^t\|_2^2 \\
 &= \Phi(\pi^t) + \eta_t (\hat{\pi}^{t+1} - \pi^t)^\top d^t + \eta_t (\hat{\pi}^{t+1} - \pi^t)^\top (\nabla \Phi(\pi^t) - d^t) + \frac{\eta_t^2 L}{2} \|\hat{\pi}_i^{t+1} - \pi^t\|_2^2 \\
 &\leq \Phi(\pi^t) + \eta_t (\hat{\pi}^{t+1} - \pi^t)^\top d^t + \eta_t \|\hat{\pi}^{t+1} - \pi^t\|_2 \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{\eta_t^2 L}{2} \|\hat{\pi}_i^{t+1} - \pi^t\|_2^2 \\
 &\leq \Phi(\pi^t) + \eta_t (\hat{\pi}^{t+1} - \pi^t)^\top d^t + \eta_t \sqrt{n} \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{\eta_t^2 n L}{2}.
 \end{aligned}$$

where the second inequality is by update rule $\hat{\pi}^{t+1} = \operatorname{argmin}_{\pi \in \Delta(\mathcal{A})} \langle \pi, d^t \rangle$, the third inequality is by Cauchy-Schwarz inequality, and the last inequality is by Cauchy-Schwarz and that $\|\pi\|_2 \leq \sqrt{n}$ for any $\pi \in \Delta(\mathcal{A})$.

Let $\bar{\pi}^{t+1} = \operatorname{argmin}_{\pi} \langle \pi, \nabla \Phi(\pi^t) \rangle$. Then

$$\begin{aligned}
 &\Phi(\pi^{t+1}) \\
 &\leq \Phi(\pi^t) + \eta_t (\hat{\pi}^{t+1} - \pi^t)^\top d^t + \eta_t \sqrt{n} \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{\eta_t^2 n L}{2} \\
 &\leq \Phi(\pi^t) + \eta_t (\bar{\pi}^{t+1} - \pi^t)^\top d^t + \eta_t \sqrt{n} \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{\eta_t^2 n L}{2} \\
 &= \Phi(\pi^t) + \eta_t (\bar{\pi}^{t+1} - \pi^t)^\top \nabla \Phi(\pi^t) + \eta_t (\bar{x}_i^{t+1} - \pi^t)^\top (d^t - \nabla \Phi(\pi^t)) + \eta_t \sqrt{n} \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{\eta_t^2 n L}{2} \\
 &\leq \Phi(\pi^t) - \eta_t G(\pi^t) + 2\eta_t \sqrt{n} \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{\eta_t^2 n L}{2},
 \end{aligned}$$

where the second inequality is by the update rule, and the third inequality is by the definition of $G(\pi^t)$ and Cauchy-Schwarz.

Rearrange the terms gives

$$G(\pi^t) \leq \frac{\Phi(\pi^t) - \Phi(\pi^{t+1})}{\eta_t} + 2\sqrt{n} \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{\eta_t n L}{2}.$$

Summing over T gives

$$\begin{aligned}
 \sum_{t=1}^T G(\pi^t) &\leq \sum_{t=1}^T \frac{\Phi(\pi^t) - \Phi(\pi^{t+1})}{\eta_t} + 2\sqrt{n} \sum_{t=1}^T \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{nL}{2} \sum_{t=1}^T \eta_t \\
 &= \sum_{t=1}^T \left(\left(\frac{\Phi(\pi^t)}{\eta_t} - \frac{\Phi(\pi^{t+1})}{\eta_{t+1}} \right) + \left(\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right) \Phi(\pi^{t+1}) \right) + 2\sqrt{n} \sum_{t=1}^T \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{nL}{2} \sum_{t=1}^T \eta_t \\
 &\leq \left(\frac{\Phi(\pi^1)}{\eta_1} - \frac{\Phi(\pi^{T+1})}{\eta_{T+1}} \right) + \frac{n}{\eta_{T+1}} + 2\sqrt{n} \sum_{t=1}^T \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{nL}{2} \sum_{t=1}^T \eta_t \\
 &\leq \frac{2n}{\eta_{T+1}} + 2\sqrt{n} \sum_{t=1}^T \|\nabla \Phi(\pi^t) - d^t\|_2 + \frac{nL}{2} \sum_{t=1}^T \eta_t,
 \end{aligned}$$

where the second inequality is because $\sum_{t=1}^T \frac{1}{\eta_{t+1}} - \frac{1}{\eta} \leq \frac{1}{\eta_{T+1}}$, and the last inequality is by $\Phi(\pi) \in [0, n]$.

By Lemma A.2, we have

$$\begin{aligned}
 \max_{\pi'_i \in \Delta(\mathcal{A}_i)} \Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi'_i, \pi_{-i}^t) &= \max_{\pi'_i \in \Delta(\mathcal{A}_i)} (\pi_i^t - \pi'_i)^\top \nabla_{\pi_i} \Phi(\pi_i^t, \pi_{-i}^t) \\
 &\leq \max_{\pi'} (\pi^t - \pi')^\top \nabla \Phi(\pi^t)
 \end{aligned}$$

$$= G(\pi^t).$$

Thus, for a fixed $i \in [m]$, we have

$$\sum_{t=1}^T \max_{\pi'_i \in \Delta(\mathcal{A}_i)} \Phi(\pi_i^t, \pi_{-i}^t) - \Phi(\pi'_i, \pi_{-i}^t) \leq \frac{2n}{\eta_{T+1}} + 2\sqrt{n} \sum_{t=1}^T \|\nabla \Phi(\tilde{\pi}^t) - d^t\| + \frac{nL}{2} \sum_{t=1}^T \eta_t.$$

Take $\eta_t = \frac{1}{\sqrt{nt^{3\alpha/2}}}$ and $\rho_t = \frac{4\mu^{1/3}n^{1/3}m^{1/3}}{t^\alpha}$, by Lemma 3.2, we have

$$\sqrt{n} \sum_{t=1}^T \mathbb{E} [\|\nabla \Phi(\pi^t) - d^t\|_2] \leq O\left(\mu n L T + \frac{m^{4/3}n^{7/6}T^{1-\alpha/2}}{\mu^{1/3}}\right),$$

where C is an absolute constant.

Putting everything together, and take $\mu = \min\left\{\frac{1}{mn}, \frac{m}{n^{1/8}T^\beta}\right\}$ we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \max_{i \in [n]} \left(c_i(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \min_{\pi'_i \in \Delta(\mathcal{A}_i)} c_i(\pi'_i, \tilde{\pi}_{-i}^t) \right) \right] \\ &= O\left(\mu n^2 L T + \frac{m^{4/3}n^{11/6}T^{1-\alpha/2}}{\mu^{1/3}} + n^2 T^{3\alpha/2} + n^2 L T^{1-3\alpha/2}\right) \\ &= O\left(mn^{15/8} L T^{1-\beta} + mn^{15/8} T^{1-\alpha/2+\beta/3} + n^{2.5} T^{3\alpha/2} + n^{2.5} L T^{1-3\alpha/2}\right). \end{aligned}$$

Take $\beta = 1/5$, $\alpha = 8/15$, we have

$$\mathbb{E} \left[\sum_{t=1}^T \max_{i \in [n]} \left(c_i(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \min_{\pi'_i \in \Delta(\mathcal{A}_i)} c_i(\pi'_i, \tilde{\pi}_{-i}^t) \right) \right] = O\left(\left(mn^{15/8}L + n^{2.5}\right)T^{4/5}\right).$$

□

A.2 Proof of Theorem 3.2

Theorem 3.2 (Regret for i -th player). *Let $\{\tilde{\pi}_{-i}^t\}_{t=1}^T$ be an arbitrary sequence of strategies for all players except for player i . For player i , running Algorithm 1 with $\eta_t = \frac{1}{\sqrt{nt^{4/5}}}$, $\mu = \min\left\{\frac{1}{mn}, \frac{m}{n^{1/8}T^{1/5}}\right\}$, $\rho_t = \frac{4\mu^{1/3}n^{1/3}m^{1/3}}{t^{8/15}}$ and $T \geq mn^{7/8}$, we can obtain a regret at most*

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T c_i(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \min_{\pi_i^* \in \Delta(\mathcal{A}_i)} \sum_{t=1}^T c_i(\pi_i^*, \tilde{\pi}_{-i}^t) \right] \\ & \leq O\left(mnLT^{4/5}\right). \end{aligned}$$

Proof. Denote $\pi^* = \operatorname{argmin}_{\pi} \sum_{t=1}^T c_i(\pi, \tilde{\pi}_{-i}^t)$, by the definition of the potential function, we have

$$\begin{aligned} \sum_{t=1}^T c_i(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - c_i(\pi_i^*, \tilde{\pi}_{-i}^t) &= \sum_{t=1}^T \Phi(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t) \\ &= \sum_{t=1}^T (\Phi(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^t, \tilde{\pi}_{-i}^t)) + (\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)) \\ &\leq \sum_{t=1}^T \|\tilde{\pi}_i^t - \pi_i^t\|_2 + (\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)) \\ &\leq \mu T + (\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)), \end{aligned}$$

where the last inequality is by the update rule $\tilde{\pi}_i^t = (1 - \mu)\pi_i^t + \mu s_i$ and Lemma A.2.

By the smoothness of Φ , we have

$$\begin{aligned} & \Phi(\pi_i^{t+1}, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t) \\ &= \Phi(\pi_i^t + \eta_t(\hat{\pi}_i^{t+1} - \pi_i^t), \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t) \\ &\leq (\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)) + \eta_t \langle \hat{\pi}_i^{t+1} - \pi_i^t, \nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) \rangle + \frac{\eta_t^2 L}{2} \|\hat{\pi}_i^{t+1} - \pi_i^t\|^2 \\ &\leq (\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)) + \eta_t \langle \hat{\pi}_i^{t+1} - \pi_i^t, \nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) \rangle + \frac{\eta_t^2 L}{2}, \end{aligned}$$

where the last inequality is because π is the concatenation of n probability vectors.

For the second term, notice that

$$\begin{aligned} \langle \hat{\pi}_i^{t+1} - \pi_i^t, \nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) \rangle &= \langle \hat{\pi}_i^{t+1} - \pi_i^t, d_i^t \rangle + \langle \hat{\pi}_i^{t+1} - \pi_i^t, \nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - d_i^t \rangle \\ &\leq \langle \pi_i^* - \pi_i^t, d_i^t \rangle + \|\nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - d_i^t\|_2 \\ &\leq \langle \pi_i^* - \pi_i^t, \nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) \rangle + \langle \pi_i^* - \pi_i^t, d_i^t - \nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) \rangle + \|\nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - d_i^t\|_2 \\ &\leq (\Phi(\pi_i^*, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^t, \tilde{\pi}_{-i}^t)) + 2\|\nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - d_i^t\|_2, \end{aligned}$$

where the first inequality is by the update rule, $\hat{\pi}_i^{t+1} = \operatorname{argmin}_{\pi} \langle \pi, d_i^t \rangle$, and the last inequality is because of the definition of Φ

Therefore, we have

$$\Phi(\pi_i^{t+1}, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t) \leq (1 - \eta_t) (\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)) + 2\eta_t \|\nabla_{\pi_i} \Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - d_i^t\|_2 + \frac{\eta_t^2 L}{2}.$$

By Lemma 3.2, take $\eta_t = \frac{2}{\sqrt{nt^{3\alpha/2}}}$, $\rho_t = \frac{4\mu^{1/3}}{n^{2/3}m^{2/3}t^{8/15}}$, and $\mu = \min\left\{\frac{1}{mn}, \frac{m}{n^{1/8}T^{1/5}}\right\}$, we have

$$\begin{aligned} & \mathbb{E} [\Phi(\pi_i^{t+1}, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)] \\ & \leq \left(1 - \frac{2}{\sqrt{nt^{3\alpha/2}}}\right) \mathbb{E} [\Phi(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)] + O\left(\frac{\eta_t^2 L}{2} + \eta_t \mu \sqrt{n} L + \frac{\eta_t m^{4/3} n^{4/3}}{\mu^{1/3} (t+1)^{\alpha/2}}\right) \end{aligned}$$

Denote $a_{t+1} = \mathbb{E} [\Phi(\pi_i^{t+1}, \tilde{\pi}_{-i}^t) - \Phi(\pi_i^*, \tilde{\pi}_{-i}^t)]$, we have

$$a_{t+1} \leq \left(1 - \frac{2}{\sqrt{nt}^{3\alpha/2}}\right) a_t + O\left(\frac{L}{t^{3\alpha}} + \frac{mn^{3/8}L}{t^{3\alpha/2+1/5}} + \frac{mn}{t^{2\alpha-1/15}}\right)$$

Take $\alpha = 8/15$, $\beta = 1/5$, we have

$$a_{t+1} \leq \left(1 - \frac{2}{\sqrt{nt}^{4/5}}\right) a_t + O\left(\frac{mnL}{t}\right).$$

We now show that $a_t \leq Dt^{-1/5}$ by induction, where $D = \max\{a_1, mnL\}$.

When $t = 1$, we have $a_1 \leq D/1$, which is true by the definition of D . Assume $a_t \leq Dt^{-1/5}$, for $t \geq 2$. Then for $t + 1$, we have

$$\begin{aligned} a_{t+1} &\leq \left(1 - \frac{2}{\sqrt{nt}^{4/5}}\right) \frac{D}{t^{1/5}} + O\left(\frac{mnL}{t}\right) \\ &\leq \left(1 - \frac{2}{t^{4/5}}\right) \frac{mnL}{t^{1/5}} + O\left(\frac{mnL}{t}\right) \\ &= O\left(\frac{mnL}{t^{1/5}} - \frac{mnL}{t}\right) = O\left(\frac{mnL}{t^{1/5}} \left(\frac{t^{4/5} - 1}{t^{4/5}}\right)\right) = O\left(\frac{mnL}{t^{1/5}} \left(\frac{1}{t^{4/5} + 1}\right)\right) \\ &= O\left(\frac{mnL}{t^{1/5} + t}\right) \leq O\left(\frac{mnL}{t^{1/5} + 1}\right) \leq O\left(\frac{mnL}{(t+1)^{1/5}}\right) \leq O\left(\frac{D}{(t+1)^{1/5}}\right). \end{aligned}$$

Putting everything together, for $T \geq mn^{7/8}$, we have

$$\mathbb{E} \left[\sum_{t=1}^T c_i(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - c_i(\pi_i^*, \tilde{\pi}_{-i}^t) \right] \leq O\left(mT^{4/5} + mnL \sum_{t=1}^T t^{-1/5}\right) = O\left(mnLT^{4/5}\right).$$

□

A.3 Proof of Lemma 3.2

Lemma 3.2. Let $\eta_t = \frac{1}{\sqrt{nt^{3\alpha/2}}}$, $\rho_t = \frac{4\mu^{1/3}n^{1/3}m^{1/3}}{t^\alpha}$, $\mu \geq \frac{1}{mn}$, $\alpha, \beta \in (0, 1)$ in Algorithm 1. We have

$$\mathbb{E} \left[\|\nabla\Phi(\pi^t) - d^t\|_2 \right] \leq O \left(\mu\sqrt{n}L + \frac{m^{4/3}n^{4/3}}{\mu^{1/3}(t+1)^{\alpha/2}} \right).$$

Proof. We start by decomposing the term as

$$\begin{aligned} \|\nabla\Phi(\pi^t) - d^t\|_2 &\leq \|\nabla\Phi(\pi^t) - \nabla\Phi(\tilde{\pi}^t)\|_2 + \|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2 \\ &\leq L\|\pi^t - \tilde{\pi}^t\|_2 + \|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2 \\ &= \mu L\|\pi^t - s\|_2 + \|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2 \\ &\leq \mu L\sqrt{n} + \|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2, \end{aligned}$$

where the second inequality is by the smoothness of Φ , the equality is by the update rule $\tilde{\pi}^t = (1 - \mu)\pi^t + \mu s$ and the last inequality is due to π being a n -dimensional vector where each entry is bounded between 0 and 1.

For the second term, by the update rule of the algorithm ($d^t = (1 - \rho_t)d^{t-1} + \rho_t\hat{C}^t$), we have

$$\begin{aligned} &\|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2^2 \\ &= \left\| \nabla\Phi(\tilde{\pi}^t) - (1 - \rho_t)d^{t-1} - \rho_t\hat{C}^t \right\|_2^2 \\ &= \left\| \rho_t \left(\nabla\Phi(\tilde{\pi}^t) - \hat{C}^t \right) + (1 - \rho_t) \left(\nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1} \right) + (1 - \rho_t) \left(\nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right) \right\|_2^2 \\ &= \rho_t^2 \left\| \nabla\Phi(\tilde{\pi}^t) - \hat{C}^t \right\|_2^2 + (1 - \rho_t)^2 \left\| \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1} \right\|_2^2 + (1 - \rho_t)^2 \left\| \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\|_2^2 \\ &\quad + 2\rho_t(1 - \rho_t) \left\langle \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1}, \nabla\Phi(\tilde{\pi}^t) - \hat{C}^t \right\rangle + 2\rho_t(1 - \rho_t) \left\langle \nabla\Phi(\tilde{\pi}^t) - \hat{C}^t, \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\rangle \\ &\quad + 2(1 - \rho_t)^2 \left\langle \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1}, \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\rangle. \end{aligned}$$

Let $\mathbb{E}_t[\cdot]$ denote the conditional expectation given the history up to time t . Taking expectations on both sides and recalling \hat{C}^t is an unbiased estimate of the gradient, we have

$$\begin{aligned} &\mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2^2 \right] \\ &= \rho_t^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \hat{C}^t \right\|_2^2 \right] + (1 - \rho_t)^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1} \right\|_2^2 \right] + (1 - \rho_t)^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\|_2^2 \right] \\ &\quad + 2\rho_t(1 - \rho_t) \mathbb{E} \left[\mathbb{E}_t \left[\left\langle \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1}, \nabla\Phi(\tilde{\pi}^t) - \hat{C}^t \right\rangle \right] \right] \\ &\quad + 2\rho_t(1 - \rho_t) \mathbb{E} \left[\mathbb{E}_t \left[\left\langle \nabla\Phi(\tilde{\pi}^t) - \hat{C}^t, \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\rangle \right] \right] \\ &\quad + 2(1 - \rho_t)^2 \mathbb{E} \left[\left\langle \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1}, \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\rangle \right] \\ &= (1 - \rho_t)^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1} \right\|_2^2 \right] + \rho_t^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \hat{C}^t \right\|_2^2 \right] + (1 - \rho_t)^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\|_2^2 \right] \\ &\quad + 2(1 - \rho_t)^2 \mathbb{E} \left[\left\langle \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1}, \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\rangle \right]. \end{aligned}$$

Using $2\langle a, b \rangle \leq \beta_t \|a\|_2^2 + \|b\|_2^2/\beta_t$, for $\beta_t > 0$, we have

$$\begin{aligned} \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - d^t \right\|_2^2 \right] &\leq (1 - \rho_t)^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1} \right\|_2^2 \right] + \rho_t^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \hat{C}^t \right\|_2^2 \right] \\ &\quad + (1 - \rho_t)^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\|_2^2 \right] + \beta_t(1 - \rho_t)^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1} \right\|_2^2 \right] \\ &\quad + (1 - \rho_t)^2/\beta_t \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\|_2^2 \right] \\ &\leq (1 - \rho_t)(1 + \beta_t) \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1} \right\|_2^2 \right] + \rho_t^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \hat{C}^t \right\|_2^2 \right] \end{aligned}$$

$$+ (1 - \rho_t)(1 + 1/\beta_t)\mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1})\|_2^2 \right].$$

Set $\beta_t = \rho_t/2$, we have

$$\begin{aligned} \mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2^2 \right] &\leq (1 - \rho_t/2)\mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1}\|_2^2 \right] + \rho_t^2 \mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^t) - \hat{C}^t\|_2^2 \right] \\ &\quad + (1 + 2/\rho_t)\mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1})\|_2^2 \right] \\ &\leq (1 - \rho_t/2)\mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1}\|_2^2 \right] + \rho_t^2 \left(m^2 n^2 + \frac{n^2 m}{\mu} \right) + (1 + 2/\rho_t)L^2 \mathbb{E} \left[\|\tilde{\pi}^t - \tilde{\pi}^{t-1}\|_2^2 \right] \\ &\leq (1 - \rho_t/2)\mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1}\|_2^2 \right] + \frac{2\rho_t^2 m^2 n^2}{\mu} + (1 + 2/\rho_t)L^2 \mathbb{E} \left[\|\pi^t - \pi^{t-1}\|_2^2 \right], \end{aligned}$$

where the second inequality is by the smoothness of Φ , Lemma A.2 and Lemma A.1.

By the update rule $\pi^{t+1} = (1 - \eta_t)\pi^t + \eta_t \hat{\pi}^{t+1}$, we have $\|\pi_i^t - \pi_i^{t-1}\|_2^2 \leq \eta_t^2 \|\hat{\pi}_i^t - \pi_i^{t-1}\|_2^2 \leq \eta_t^2 n$.

Therefore, we have

$$\mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2^2 \right] \leq (1 - \rho_t/2)\mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^{t-1}) - d^{t-1}\|_2^2 \right] + \frac{2\rho_t^2 m^2 n^2}{\mu} + \frac{3\eta_t^2 n}{\rho_t}.$$

Define $a_t = \mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2^2 \right]$ and $\eta_t = \frac{1}{\sqrt{nt^{3\alpha/2}}}$, we have

$$\begin{aligned} a_t &\leq \left(1 - \frac{\rho_t}{2}\right) a_{t-1} + \frac{2\rho_t^2 m^2 n^2}{\mu} + \frac{3\eta_t^2 n}{\rho_t} \\ &\leq \left(1 - \frac{\rho_t}{2}\right) a_{t-1} + \frac{2\rho_t^2 m^2 n^2}{\mu} + \frac{3}{\rho_t t^{3\alpha}}. \end{aligned}$$

Take $\rho_t = \frac{4\mu^{1/3} n^{1/3} m^{1/3}}{t^\alpha}$. We want to show $a_t = \frac{C}{(t+1)^\alpha}$, $C = \max \left\{ a_1 2^\alpha, \frac{35m^{8/3} n^{8/3}}{\mu^{2/3}} \right\}$. We show this by induction.

The base case $t = 1$ clearly holds due to our choice of C_1 . When $\mu \geq 1/mn$, assume the $a_{\ell-1} \leq \frac{C}{\ell^\alpha}$, $2 \leq \ell \leq T$, we have

$$\begin{aligned} a_\ell &\leq \left(1 - \frac{2\mu^{1/3} n^{1/3} m^{1/3}}{\ell^\alpha}\right) \frac{C}{\ell^\alpha} + \frac{32m^{8/3} n^{8/3}}{\mu^{1/3} \ell^{2\alpha}} + \frac{3n^{8/3} m^{8/3}}{4\mu^{1/3} \ell^{2\alpha}} \\ &\leq \left(1 - \frac{2}{\ell^\alpha}\right) \frac{C}{\ell^\alpha} + \frac{C}{\ell^{2\alpha}} \leq C \left(\frac{\ell^\alpha - 1}{\ell^{2\alpha}} \right) \\ &= \left(\frac{C}{\ell^\alpha + 1} \right) \leq \left(\frac{C}{(\ell + 1)^\alpha} \right). \end{aligned}$$

Taking square root on both sides yield,

$$\mathbb{E} \left[\|\nabla\Phi(\tilde{\pi}^t) - d^t\|_2 \right] \leq \frac{\max\{C_1, \sqrt{35}m^{4/3}n^{4/3}\}}{\mu^{1/3}(t+1)^{\alpha/2}},$$

where C_1 is a constant. Combining everything yields the claimed result. \square

A.4 Auxiliary Lemmas

Lemma A.1. *Let $\mathbb{E}_t[\cdot]$ denote the conditional expectation given the history up to time t , we have*

- $\mathbb{E}_t[(\hat{C}_i^t)_a] = (\nabla_{\pi_i} \Phi(\tilde{\pi}))_a$.
- $\mathbb{E}_t[\|\hat{C}_i^t\|^2] = \frac{m}{\mu}$.

Proof. Let $\mathbb{E}_t[\cdot]$ denote the conditional expectation given the history up to time t . Then we have

$$\mathbb{E}_t[(\hat{C}_i^t)_a] = \mathbb{E}_t \left[\frac{C_i^t \cdot \mathbb{I}\{a = a_i^t\}}{(\tilde{\pi}_i^t)_a} \right] = \frac{c_i(a, \tilde{\pi}_{-i})}{(\tilde{\pi}_i^t)_a} \mathbb{E}_t[\mathbb{I}\{a = a_i^t\}] = \frac{c_i(a, \tilde{\pi}_{-i})}{(\tilde{\pi}_i^t)_a} (\pi_i^t)_a = c_i(a, \tilde{\pi}_{-i}) = (\nabla_{\pi_i} \Phi(\tilde{\pi}))_a.$$

We also have

$$\mathbb{E}_t[\|\hat{C}_i^t\|^2] \leq \sum_{a \in \mathcal{A}_i} \mathbb{E}_t \left[\left(\frac{C_i^t \cdot \mathbb{I}\{a = a_i^t\}}{(\tilde{\pi}_i^t)_a} \right)^2 \right] = \sum_{a \in \mathcal{A}_i} c_i(a, \tilde{\pi}_{-i})^2 \mathbb{E}_t \left[\frac{1}{(\tilde{\pi}_i^t)_a^2} \right] = \sum_{a \in \mathcal{A}_i} c_i(a, \tilde{\pi}_{-i})^2 \frac{(\tilde{\pi}_i^t)_a}{(\tilde{\pi}_i^t)_a^2} \leq \frac{m}{\mu},$$

where the last inequality is by our assumption $c_i(a, \tilde{\pi}_{-i}) \in [0, 1]$. □

Lemma A.2 (Properties of the potential function). *For any π, π' , we have*

1. $\|\Phi(\pi) - \Phi(\pi')\|_2 \leq \sqrt{n} \|\pi - \pi'\|_2$.
2. $\|\nabla_{\pi_i} \Phi(\pi)\|_2 \leq m$.
3. $\Phi(\pi_i, \pi_{-i}) - \Phi(\pi'_i, \pi_{-i}) = (\pi_i - \pi'_i) \nabla_{\pi_i} \Phi(\pi_i, \pi_{-i})$.

Proof. The first one is true as $\|\nabla_{\pi} \Phi(\pi)\|_2$ is bounded by \sqrt{n} , due to the boundedness of the cost function. The second one is due to $\nabla_{\pi_i} \Phi(\pi) = \nabla_{\pi_i} c_i(\pi)$ and $\|c_i(\cdot, \pi_{-i})\|_2 \leq m$. The third one is by noticing $\Phi(\pi)$ is linear in π_i for any $i \in [n]$. □

B Analysis for Markov potential game

B.1 Proof of Theorem 4.1

Theorem 4.1 (Nash regret for Markov potential game). *If all players run Algorithm 2 with $\eta_t = \frac{\kappa}{n^{3/2}m^{4/5}}$, $\rho_t = \frac{4\mu^{1/3}}{n^{1/3}m^{2/3}t^{8/15}}$ and $\mu = \min\{\frac{1}{m^2n}, \frac{\kappa^{4/3}}{n^{7/8}m^{1/4}T^\beta}\}$, when $T \geq m^{7/4}n^{1/8}\kappa^{4/3}$, we have*

$$\begin{aligned} & \mathbb{E}[\text{Nash-regret}(T)] \\ & \leq O\left(\frac{n^{3/2}mD_\infty T^{4/5}}{\kappa^3} + \frac{n^{5/8}m^{3/4}SD_\infty T^{4/5}}{\kappa^{25/9}}\right), \end{aligned}$$

where $D_\infty = \max_i \left\| \frac{d^{\pi_i^*, \pi_{-i}(\mu)}}{\mu} \right\|_\infty$.

Proof. Denote $\pi_i^{',t} = \min_{\pi_i} V_i^{\pi_i, \tilde{\pi}_{-i}^t}(\mu)$, we first decompose the regret as

$$\begin{aligned} \text{Nash-regret}(T) &= \sum_{t=1}^T \max_{i \in [n]} \left(V_i^{\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t}(\mu) - \min_{\pi_i} V_i^{\pi_i, \tilde{\pi}_{-i}^t}(\mu) \right) \\ &= \sum_{i=1}^n \sum_{t=1}^T \left(V_i^{\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^t, \tilde{\pi}_{-i}^t}(\mu) \right) + \sum_{i=1}^n \sum_{t=1}^T \left(V_i^{\pi_i^{',t}, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^t, \tilde{\pi}_{-i}^t}(\mu) \right) \\ &\quad + \sum_{i=1}^n \sum_{t=1}^T \left(V_i^{\pi_i^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^{',t}, \tilde{\pi}_{-i}^t}(\mu) \right) \\ &\leq \frac{n}{\kappa} \sum_{i=1}^n \sum_{t=1}^T \|\tilde{\pi}^t - \pi^t\|_2 + \sum_{i=1}^n \sum_{t=1}^T \left(V_i^{\pi_i^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^{',t}, \tilde{\pi}_{-i}^t}(\mu) \right) \\ &\leq \frac{n\mu}{\kappa} \sum_{i=1}^n \sum_{t=1}^T \|s - \pi^t\|_2 + \sum_{i=1}^n \sum_{t=1}^T \left(V_i^{\pi_i^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^{',t}, \tilde{\pi}_{-i}^t}(\mu) \right) \\ &\leq \frac{\mu n^{3/2} \sqrt{ST}}{\kappa} + \sum_{i=1}^n \sum_{t=1}^T \left(V_i^{\pi_i^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^{',t}, \tilde{\pi}_{-i}^t}(\mu) \right), \end{aligned}$$

where the first inequality is by Lemma 18 of Ding et al. (2022), and the last inequality is by $\|s - \pi^t\|_2 \leq \sqrt{nS}$.

Similar to the proof of Theorem 4.1, and by Lemma B.4, we have

$$\begin{aligned} \max_{\pi_i} \left(V_i^{\pi_i^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i, \tilde{\pi}_{-i}^t}(\mu) \right) &\leq \frac{1}{\kappa} \left\| \frac{d^{\pi_i^*, \pi_{-i}(\mu)}}{\mu} \right\|_\infty \max_{\pi_i'} (\pi_i' - \pi_i)^\top \nabla_{\pi_i} \Phi^\pi(\mu) \\ &\leq \frac{1}{\kappa} \left\| \frac{d^{\pi_i^*, \pi_{-i}(\mu)}}{\mu} \right\|_\infty G(\pi^t), \end{aligned}$$

where $G(\pi^t) = \max_{\pi} (\pi^t - \pi)^\top \nabla \Phi_\mu(\pi^t)$.

We bound $\sum_{t=1}^T \mathbb{E}[G(\pi^t)]$ in a similar fashion as Theorem 4.1.

Notice that for any π, π' , we have

$$\mathbb{E}[\Phi(\pi)] \leq \mathbb{E}[(H+1)] \leq \sum_{h=0}^{\infty} (1-\kappa)^h \kappa (h+1) \leq \frac{\kappa}{1-\kappa} \sum_{h=0}^{\infty} (1-\kappa)^h h = \frac{1}{\kappa},$$

Therefore, take $\eta_t = \frac{\kappa}{n^{3/2}m^{3\alpha/2}}$, $\rho_t = \frac{4\mu^{1/3}}{n^{1/3}m^{2/3}t^\alpha}$ and $\mu = \min\{\frac{1}{m^2n}, \frac{\kappa^{4/3}}{n^{7/8}m^{1/4}T^\beta}\}$, and by Lemma B.1 we have

$$\sum_{t=1}^T \mathbb{E}[G(\pi^t)] \leq \frac{2}{\kappa \eta_{T+1}} + \sqrt{nS} \sum_{t=1}^T \|\nabla \Phi_\mu(\pi^t) - d^t\|_2 + \frac{2nS(1-\kappa)}{\kappa^3} \sum_{t=1}^T \eta_t$$

$$\leq O\left(\frac{n^{3/2}mT^{3\alpha/2}}{\kappa^2} + \frac{n^{3/2}mST^{1-\beta}}{\kappa^{5/3}} + \frac{n^{5/8}m^{3/4}ST^{1-\alpha/2+\beta/3}}{\kappa^{25/9}} + \frac{nST^{1-3\alpha/2}}{\kappa^2}\right)$$

Take $\alpha = 8/15$, $\beta = 1/5$, we have

$$\sum_{t=1}^T \mathbb{E}[G(\pi^t)] \leq O\left(\frac{n^{3/2}mT^{4/5}}{\kappa^2} + \frac{n^{5/8}m^{3/4}ST^{4/5}}{\kappa^{25/9}}\right).$$

Combine everything, we have

$$\mathbb{E}[\text{Nash-regret}(T)] \leq O\left(\frac{n^{3/2}mD_\infty T^{4/5}}{\kappa^3} + \frac{n^{5/8}m^{3/4}SD_\infty T^{4/5}}{\kappa^{25/9}}\right),$$

where $D_\infty = \max_i \left\| \frac{d^{\pi_i^*, \pi_{-i}}(\mu)}{\mu} \right\|_\infty$.

□

B.2 Proof of Lemma B.1

Lemma B.1. *With $\eta_t = \frac{1}{n^{3/2}mt^{3\alpha/2}}$, $\rho_t = \frac{4\mu^{1/3}n^{1/3}m^{2/3}}{\kappa^{2/3}t^\alpha}$ and $\mu \geq \frac{1}{m^2n}$, we have*

$$\mathbb{E} [\|\nabla\Phi_\mu(\pi^t) - d^t\|_2] \leq O\left(\frac{n^{3/2}m\mu\sqrt{S}(1-\kappa)}{\kappa^3} + \frac{n^{1/3}m^{2/3}\sqrt{S}}{\mu^{1/3}\kappa^{7/3}(t+1)^{\alpha/2}}\right),$$

Proof. First notice that

$$\mathbb{E} [\nabla\Phi(\tilde{\pi}^t)] \leq \mathbb{E}[(H+1)] \leq \sum_{h=0}^{\infty} (1-\kappa)^h \kappa(h+1) \leq \frac{\kappa}{1-\kappa} \sum_{h=0}^{\infty} (1-\kappa)^h h = \frac{1}{\kappa}.$$

The proof is similar to Lemma 3.2. Similarly let $a_t = \mathbb{E} [\|\nabla\Phi_\mu(\tilde{\pi}^t) - d^t\|_2^2]$ and take $\eta_t = \frac{\kappa}{n^{3/2}mt^{3\alpha/2}}$, we have

$$\begin{aligned} a_t &\leq (1-\rho_t/2)a_{t-1} + \rho_t^2 \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \widehat{\nabla}_{\pi_i^t} \right\|_2^2 \right] + (1+2/\rho_t) \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\|_2^2 \right] \\ &\leq (1-\rho_t/2)a_{t-1} + \rho_t^2 \left(\mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) \right\|_2^2 \right] + \mathbb{E} \left[\left\| \widehat{\nabla}_{\pi_i^t} \right\|_2^2 \right] \right) + (1+2/\rho_t) \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\|_2^2 \right] \\ &\leq (1-\rho_t/2)a_{t-1} + \rho_t^2 \left(\frac{1}{\kappa^2} + \frac{24nm^2}{\mu\kappa^4} \right) + (1+2/\rho_t) \mathbb{E} \left[\left\| \nabla\Phi(\tilde{\pi}^t) - \nabla\Phi(\tilde{\pi}^{t-1}) \right\|_2^2 \right] \\ &\leq (1-\rho_t/2)a_{t-1} + \frac{25nm^2\rho_t^2}{\mu\kappa^4} + \frac{16n^2m^2(1-\kappa)^2(1+2/\rho_t)}{\kappa^6} \mathbb{E} \left[\left\| \tilde{\pi}^t - \tilde{\pi}^{t-1} \right\|_2^2 \right] \\ &\leq (1-\rho_t/2)a_{t-1} + \frac{25nm^2\rho_t^2}{\mu\kappa^4} + \frac{36n^3Sm^2\eta_t^2}{\kappa^6\rho_t} \\ &\leq (1-\rho_t/2)a_{t-1} + \frac{25nm^2\rho_t^2}{\mu\kappa^4} + \frac{36S}{\kappa^4\rho_t} \end{aligned}$$

where the second inequality is by Lemma B.2 and the definition of the potential function.

Take $\rho_t = \frac{4n^{1/3}m^{2/3}\mu^{1/3}}{t^\alpha}$. We want to show $a_t = \frac{C}{(t+1)^\alpha}$, $C = \max\left\{a_12^\alpha, \frac{61n^{5/3}m^{7/3}S}{\mu^{2/3}\kappa^4}\right\}$. We show this by induction.

The base case $t = 1$ clearly holds due to our choice of C_1 .

Assume the $a_{\ell-1} \leq \frac{C}{\ell^\alpha}$, $2 \leq \ell \leq T$, and notice that as $n, m \geq 1$, $\mu \leq 1$, we have

$$\begin{aligned} a_\ell &\leq \left(1 - \frac{2n^{1/3}m^{2/3}\mu^{1/3}}{t^\alpha}\right) \frac{C}{\ell^\alpha} + \frac{61n^{5/3}m^{7/3}S}{\mu^{1/3}\kappa^4\ell^{2\alpha}} \\ &\leq C \left(\frac{\ell^\alpha - 1}{\ell^{2\alpha}}\right) \\ &\leq \left(\frac{C}{\ell^\alpha + 1}\right) \leq \left(\frac{C}{(\ell+1)^\alpha}\right). \end{aligned}$$

Therefore

$$\begin{aligned} \mathbb{E} [\|\nabla\Phi_\mu(\pi^t) - d^t\|_2] &\leq \mathbb{E} [\|\nabla\Phi_\mu(\pi^t) - \nabla\Phi_\mu(\tilde{\pi}^t)\|_2] + \mathbb{E} [\|\nabla\Phi_\mu(\tilde{\pi}^t) - d^t\|_2] \\ &\leq \frac{4nm(1-\kappa)}{\kappa^3} \mathbb{E} [\|\pi^t - \tilde{\pi}^t\|_2] + \mathbb{E} [\|\nabla\Phi_\mu(\tilde{\pi}^t) - d^t\|_2] \\ &\leq \frac{4nm\mu(1-\kappa)}{\kappa^3} \mathbb{E} [\|\pi^t - s\|_2] + \mathbb{E} [\|\nabla\Phi_\mu(\tilde{\pi}^t) - d^t\|_2] \\ &\leq O\left(\frac{n^{3/2}m\mu\sqrt{S}(1-\kappa)}{\kappa^3} + \frac{n^{5/6}m^{7/6}\sqrt{S}}{\mu^{1/3}\kappa^2(t+1)^{\alpha/2}}\right). \end{aligned}$$

□

B.3 Proof of Theorem 4.2

Theorem 4.2 (Regret for the i -th player in Markov potential game). *Let $\{\tilde{\pi}_{-i}^t\}_{t=1}^T$ be an arbitrary sequence of strategies for all players except for player i . For player i , run Algorithm 2 with $\eta_t = \frac{\kappa}{n^{3/2}mt^{4/5}}$, $\rho_t = \frac{4\mu^{1/3}}{n^{1/3}m^{2/3}t^{8/15}}$ and $\mu = \min\{\frac{1}{m^2n}, \frac{\kappa^{4/3}}{n^{7/8}m^{1/4}T^\beta}\}$, when $T \geq m^{7/4}n^{1/8}\kappa^{4/3}$. Then the regret of player i is at most*

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \left(V_i^{\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^*, \tilde{\pi}_{-i}^t}(\mu) \right) \right] \\ & \leq O \left(\frac{m^{4/3}n^{27/24}ST^{4/5}}{\kappa^{7/3}} \right). \end{aligned}$$

Proof. Denote $\pi^* = \operatorname{argmin}_{\pi} \sum_{t=1}^T V_i^{\pi_i, \tilde{\pi}_{-i}^t}(\mu)$ by the definition of the potential function, we have

$$\begin{aligned} \sum_{t=1}^T \left(V_i^{\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^*, \tilde{\pi}_{-i}^t}(\mu) \right) &= \sum_{t=1}^T \Phi_{\mu}(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^*, \tilde{\pi}_{-i}^t) \\ &= \sum_{t=1}^T \left(\Phi_{\mu}(\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^t, \tilde{\pi}_{-i}^t) \right) + \left(\Phi_{\mu}(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^*, \tilde{\pi}_{-i}^t) \right) \\ &\leq \sum_{t=1}^T \|\tilde{\pi}_i^t - \pi_i^t\| + \left(\Phi_{\mu}(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^*, \tilde{\pi}_{-i}^t) \right) \\ &\leq \frac{4m\mu\sqrt{nST}}{\kappa^3} + \left(\Phi_{\mu}(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^*, \tilde{\pi}_{-i}^t) \right), \end{aligned}$$

where the last inequality is by the update rule and $\|\pi_i^t - s\|_2 \leq \sqrt{S}$.

Similar to the proof of Theorem 3.1, by the smoothness of Φ , $\eta_t = \frac{\kappa}{n^{3/2}mt^{3\alpha/2}}$, $\rho_t = \frac{4\mu^{1/3}n^{1/3}m^{2/3}}{\kappa^{2/3}t^\alpha}$ and $\mu = \min\{\frac{1}{m^2n}, \frac{\kappa^{4/3}}{n^{7/8}m^{1/4}T^\beta}\}$, and $\alpha = 8/15, \beta = 1/5$ we have

$$\begin{aligned} & \Phi_{\mu}(\pi_i^{t+1}, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^*, \tilde{\pi}_{-i}^t) \\ & \leq (1 - \eta_t) \left(\Phi_{\mu}(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^*, \tilde{\pi}_{-i}^t) \right) + 2\eta_t\sqrt{S}\|\nabla_{\pi_i}\Phi_{\mu}(\pi_i^t, \tilde{\pi}_{-i}^t) - d_i^t\| + \frac{8\eta_t^2n^{3/2}m\sqrt{S}(1-\kappa)}{2\kappa^3} \\ & \leq \left(1 - \frac{\kappa}{n^{3/2}mt^{3\alpha/2}} \right) \left(\Phi_{\mu}(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^*, \tilde{\pi}_{-i}^t) \right) + O \left(\frac{S(1-\kappa)}{t^{3\alpha/2+\beta}\kappa^3} + \frac{n^{27/24}m^{5/4}S}{\kappa^{7/3}t^{2\alpha-\beta/3}} + \frac{\sqrt{S}}{\kappa t^{3\alpha}} \right) \\ & \leq \left(1 - \frac{\kappa}{n^{3/2}mt^{3\alpha/2}} \right) \left(\Phi_{\mu}(\pi_i^t, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^*, \tilde{\pi}_{-i}^t) \right) + O \left(\frac{n^{27/24}m^{5/4}S}{\kappa^{7/3}t} \right). \end{aligned}$$

Using the same induction proof as the proof of Theorem 4.2, we have

$$\Phi_{\mu}(\pi_i^{t+1}, \tilde{\pi}_{-i}^t) - \Phi_{\mu}(\pi_i^*, \tilde{\pi}_{-i}^t) \leq \frac{D}{(t+1)^{1/5}},$$

where $D = \max\{D_1, \frac{m^{4/3}n^{27/24}S}{\kappa^{7/3}}\}$, and D_1 are absolute constants. Therefore, we have

$$\sum_{t=1}^T \left(V_i^{\tilde{\pi}_i^t, \tilde{\pi}_{-i}^t}(\mu) - V_i^{\pi_i^*, \tilde{\pi}_{-i}^t}(\mu) \right) \leq O \left(\frac{m^{4/3}n^{27/24}ST^{4/5}}{\kappa^{7/3}} \right).$$

□

B.4 Auxiliary Lemmas

Lemma B.2 ((Daskalakis et al., 2020, Lemma 2)). *Let \mathbb{E}_t denote the conditional expectation on all history up to time t , then*

1. $\mathbb{E}_t \left[\widehat{\nabla}_{\pi_i^t} \right] = \nabla_{\pi_i} \Phi^{\pi_i^t}(\mu)$.
2. $\mathbb{E}_t \left[\left\| \widehat{\nabla}_{\pi_i^t} \right\|_2^2 \right] \leq \frac{24m^2}{\mu\kappa^4}$.

Lemma B.3 ((Daskalakis et al., 2020, Lemma 4)). *For any π, π' , $\|\nabla\Phi_\mu(\pi) - \nabla\Phi_\mu(\pi')\|_2 \leq \frac{4nm(1-\kappa)}{\kappa^3} \|\pi - \pi'\|_2$.*

Lemma B.4 ((Leonardos et al., 2022, Lemma D.3)). *Fix any player i , let π_i^* be an optimal policy for agent i in the single agent MDP in which the rest of the agents are fixed to choose π_{-i} . Then,*

$$\Phi^{\pi_i^*, \pi_{-i}}(\mu) - \Phi^\pi(\mu) \leq \frac{1}{\kappa} \left\| \frac{d^{\pi_i^*, \pi_{-i}}(\mu)}{\mu} \right\|_\infty \max_{\pi_i'} (\pi_i' - \pi_i)^\top \nabla_{\pi_i} \Phi^\pi(\mu),$$

where $d^\pi(\mu) = \mathbb{E} \left[\sum_{h=0}^H P^\pi (s^h = s \mid s_0) \right]$.