DIFFERENTIALLY PRIVATE EQUILIBRIUM FINDING IN POLYMATRIX GAMES

Anonymous authorsPaper under double-blind review

ABSTRACT

We study equilibrium finding in polymatrix games under differential privacy constraints. Prior work in this area fails to achieve both high-accuracy equilibria and a low privacy budget. To better understand the fundamental limitations of differential privacy in games, we show hardness results establishing that no algorithm can simultaneously obtain high accuracy and a vanishing privacy budget as the number of players tends to infinity. This impossibility holds in two regimes: (i) We seek to establish equilibrium approximation guarantees in terms of Euclidean *distance* to the equilibrium set, and (ii) The adversary has access to all communication channels. We then consider the more realistic setting in which the adversary can access only a bounded number of channels and propose a new distributed algorithm that: recovers strategies with simultaneously vanishing *Nash gap* (in expected utility, also referred to as *exploitability*) and *privacy budget* as the number of players increases. Our approach leverages structural properties of polymatrix games. To our knowledge, this is the first paper that can achieve this in equilibrium computation. Finally, we also provide numerical results to justify our algorithm.

1 Introduction

Many multi-agent settings are hard because everyone affects everyone else. However, in many settings of interest, the interaction is more localized and tractable. Polymatrix games provide a tractable and rich model that can capture such settings (Janovskaja, 1968; Cai et al., 2016; Deligkas et al., 2017). In polymatrix games, players engage in pairwise interactions defined by an underlying graph in which nodes represent players and edges model the interaction between the two players connected by an edge. Each player chooses the same strategy in each pairwise game she is involved with, and her utility is the sum of her utilities for each of these games.

In several settings, such as security games (De Nittis et al., 2018) and financial markets (Evangelista et al., 2022; Donmez et al., 2024), the players' utility functions may be sensitive and need to be kept private by the players as they update their behavior toward equilibrium. Typically, there are two approaches for finding equilibrium while keeping the utility functions private: centralized and distributed methods. In centralized methods, the players send their utility functions to a trusted central server (Kearns et al., 2014; Rogers & Roth, 2014; Cummings et al., 2015). When a trusted central server is infeasible, the only viable option for computing equilibrium is for the players to use a distributed algorithm where they make local computations and exchange information with their neighbors. Throughout the process, a malicious adversary may inspect the communication channel and infer the utility function, which might reveal sensitive information about individual preferences. To keep the computation secure and prevent the leakage of sensitive information, tools from differential privacy (DP) (Dwork et al., 2006) can be employed.

Several authors have considered the question of differentially private equilibrium computation in games. However, previous work was either unable to achieve high accuracy and low differential privacy budget simultaneously (Ye et al., 2021; Wang et al., 2022), or only achieved a weaker form of differential privacy where the adversary is still able to infer part of the private information (Wang & Başar, 2024; Wang & Nedić, 2024). Details can be found in Appendix A.

CONTRIBUTIONS AND TECHNIQUES

In this paper, we study the problem of distributed, differentially private equilibrium computation in polymatrix games. We build on the notion of adjacent distributed optimization problem defined in Huang et al. (2015). In particular, given a polymatrix game \mathcal{G} we define *adjacent games* as those that differ from \mathcal{G} on the utility function of a single player (c.f. Figure 1 (a) and Definition 2.1). Then, once the algorithm for equilibrium finding is differentially private, an adversary cannot distinguish, with high probability from his observations, the original game from any other games that are adjacent to it. Hence, when differential privacy holds, the adversary is unable to determine the utility function of any player with high confidence, ensuring the equilibrium computation process is privacy-preserving.

Impossibility. Using this concept of adjacency, we begin by showing that finding an approximate equilibrium with arbitrarily small differential privacy guarantees can be impossible, depending on the desiderata imposed on the process. Specifically, we show that if the adversary can access an arbitrary number of communication channels, or the target accuracy metric of interest is the Euclidean distance from the equilibrium set, then distributed computation of high-quality equilibria while providing vanishing privacy guarantees is impossible. In other words, one will inevitably suffer either a low approximation of the equilibrium, or a low level of privacy (high privacy budget). While these negative results serve as guardrails to guide algorithm design, it is important to realize that they do not preclude all paths to meaningful results. Indeed, it is important to realize that Euclidean distance to the equilibrium set is only one of the set of possible quality metrics for measuring equilibrium approximation. Another metric is *Nash gap* or *exploitability*, meaning how close a strategy is to equilibrium in terms of *expected utility* rather than metric distance to the equilibrium set in strategy space. This is particularly relevant in practice, where potential gains from deviation are more important than proximity in strategy space.

Positive Results. Finally, we complement our negative result with a positive result regarding exploitability guarantees, corroborated by experiments in Appendix E. In particular, we propose a new distributed algorithm for computing a coarse correlated equilibrium (c.f. Figure 1 (b)). As is typical in the differential privacy literature, our algorithm involves each player communicating a noisy version of their strategy to their neighbors and updating their strategy using a regularized proximal gradient step. A key novelty of the algorithm is to scale the regularizer proportional to the harmonic mean of the degrees divided by the degree of the player. The inverse proportionality with the degree ensures that more regularization is introduced for players with a lower degree (since a low-degree player's gradient is more sensitive to variations in the utility matrices of its neighbors).

High-Level Intuition. We propose Algorithm 1, which simultaneously achieves low exploitability and low DP budget, with guarantees that improve as the number of players N increases. Moreover, the algorithm achieves that whenever the associated graph with the polymatrix game is sparse or dense, by imposing an adaptive regularizer on players' utility functions. When the graph associated with the polymatrix game is sparse, we leverage the fact that for most nodes that are at least of distance t to the edge of the changed utility matrices, changes in the utility matrices between adjacent polymatrix games have minimal impact on them during the first t updates. In contrast, when the graph is dense, additional regularization, which is inversely proportional to the players' degrees, stabilizes the updates of players with low degrees so that the adversary's observation in two adjacent games will be similar. For players with high degrees, the aggregative nature of the utility function in polymatrix games mitigates the effect of changes in the utility matrices.

2 Preliminaries

We now review several key concepts related to differential privacy and polymatrix games.

Basics. For any integer N>0, let $[N]:=\{1,2,3,\cdots,N-1,N\}$. For any vector \boldsymbol{x} , we use $\|\boldsymbol{x}\|_p$ to denote its p-norm. By default, $\|\boldsymbol{x}\|$ represents the 2-norm. Let $\Delta^n:=\{\boldsymbol{x}\in[0,1]^n\colon\sum_{i=1}^nx_i=1\}$ denotes the n-1 dimensional probability simplex. Moreover, for any discrete set \mathcal{S} , we can define $\Delta^{\mathcal{S}}$ as the probability simplex over \mathcal{S} , with each index as the element in \mathcal{S} . For instance, $\Delta^{[N]}=\Delta^N$.

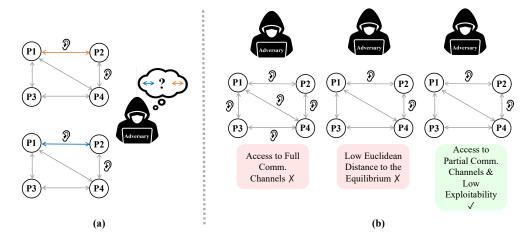


Figure 1 (a): An illustration of adjacent polymatrix games defined in Definition 2.1. The nodes represent the players, and the lines represent the edges of the polymatrix game. The two games differ at the blue/orange edge, and the adversary cannot differentiate these two games with certainty from his observations. If on an edge means that the communication channels between those two players can be accessed by the adversary.

Figure 1 (b): The illustration of results: the impossibility when the adversary can access all communication channels showed in Lemma 3.1 (left), the impossibility to achieve low Euclidean distance showed in Lemma 3.2 (middle), and the positive results showed in Theorem 7.1 (right).

For any convex set \mathcal{C} , we use $\operatorname{Proj}_{\mathcal{C}}(x') \coloneqq \operatorname{argmin}_{x \in \mathcal{C}} \|x - x'\|^2$ to denote the projection to \mathcal{C} from a vector x' with respect to Euclidean distance.

Polymatrix Games. Polymatrix games can be written as a tuple

$$\mathcal{G} := \left([N], E, \left\{ \mathcal{A}_i \right\}_{i \in [N]}, \left\{ U_{i,j} \right\}_{(i,j) \in E} \right),$$

where

- [N] is the set of players so there are N players.
- E is the set of edges, where each $(i, j) \in E$ indicates that players i, j interact with each other.
- $\{A_i\}_{i\in[N]}$ is the set of players' action set, which means player $i\in[N]$ chooses actions in A_i .
- Let A := max_{i∈[N]} |A_i| be the size of the largest action set.
 U_{i,j} ∈ [-1,1]^{A_i×A_j} is the utility matrix between player (i, j) ∈ E.

For each Player $i \in [N]$, we denote with the symbol $\mathcal{N}(i) \coloneqq \{j \in [N] \colon (i,j) \in E\}$ the set of her neighbors. Then, when each player $k \in [N]$ chooses her strategy as $\pi_k \in \Delta^{\mathcal{A}_k}$, the utility of player i is $\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \pi_i^\top U_{i,j} \pi_j^{-1}$. Moreover, when the strategy profile of all players is $\pi \coloneqq (\pi_1, \pi_2, \cdots, \pi_N)$, we can define the gradient of player i's strategy with respect to her loss function (the negative of utility function) as $g_i^\pi \coloneqq -\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} U_{i,j} \pi_j$.

Differential Privacy. We adopt the notion of adjacency in Huang et al. (2015). For any two polymatrix games, they are adjacent to each other if and only if they only differ from the utility matrices on a single edge. An intuitive illustration can be found in Figure 1 (a). Formally,

Definition 2.1 (Game adjacency). Given two polymatrix games $\mathcal{G} = ([N], E, \{A_i\}_{i \in [N]}, \{A_i\}_{i \in [N]})$ $\{U_{i,j}\}_{(i,j)\in E}$ and $\mathcal{G}'=\left([N'],E',\{\mathcal{A}_i'\}_{i\in[N']},\left\{U_{i,j}'\right\}_{(i,j)\in E'}\right)$, they are said to be *adjacent*, indicated as $\mathcal{G} \sim \mathcal{G}'$, if

¹Typically the definition of player *i*'s utility is $\sum_{j \in \mathcal{N}(i)} \pi_i^{\top} U_{i,j} \pi_j$. However, to ensure that the function and its gradient are bounded by constants regardless of the game size, we divide it by $|\mathcal{N}(i)|$, which aligns with the DP literature (Huang et al., 2015; Wei et al., 2020; Ye et al., 2021). Note that such modification will not change the equilibrium.

1. N = N', E = E' and $A_i = A'_i$ for any $i \in [N]$; and 2. except for an edge $(i,j) \in E$, $U_{i',j'} = U'_{i',j'}$ and $U_{j',i'} = U'_{j',i'}$ for any $(i',j') \in E \setminus \{(i,j),(i,j)\}$.

Consider now a generic iterative algorithm for finding equilibrium in polymatrix games, and denote with Θ the space of all possible internal states of the algorithm at any time. Then, for a given polymatrix game \mathcal{G} , when the algorithm runs t>0 timesteps, it will generate intermediate variables $\boldsymbol{\theta} \coloneqq \left(\theta^{(1)}, \theta^{(2)}, \cdots, \theta^{(t)}\right) \in \Theta^t$, which are called *executions* of the algorithm. Θ^t is the t-fold Cartesian product of Θ with itself. Due to the randomness in the algorithm, this will result in a distribution $\mathcal{P}_{\mathcal{G}}$ over Θ^t . Moreover, there is an adversary with access to the communication channels between the players. Therefore, for any execution $\boldsymbol{\theta} \in \Theta^t$, the adversary can observe $\mathcal{R}_{\mathcal{G}}(\boldsymbol{\theta}) = \left(o^{(1)}, o^{(2)}, \cdots, o^{(t)}\right) =: \boldsymbol{o} \in \mathcal{O}^t$, where \mathcal{O} is the space of all *observations* at a single timestep. Hence, we can define the (ϵ, δ) -differential privacy $((\epsilon, \delta)$ -DP) as follows.

Definition 2.2 $((\epsilon, \delta)$ -Differential Privacy). For an $\epsilon \geq 0$ and $\delta \geq 0$, an iterative distributed algorithm for finding equilibria is (ϵ, δ) -differentially private, if and only if for any two adjacent polymatrix game $\mathcal{G}, \mathcal{G}'$, any timestep t > 0 and any set of observations $\mathcal{S} \subseteq \mathcal{O}^t$,

$$\mathcal{P}_{\mathcal{G}}\left(\left\{\boldsymbol{\theta}\colon\mathcal{R}_{\mathcal{G}}(\boldsymbol{\theta})\in\mathcal{S}\right\}\right)\leq e^{\epsilon}\mathcal{P}_{\mathcal{G}'}\left(\left\{\boldsymbol{\theta}\colon\mathcal{R}_{\mathcal{G}'}(\boldsymbol{\theta})\in\mathcal{S}\right\}\right)+\delta. \tag{1}$$

Intuitively, Definition 2.2 guarantees that the algorithm will generate similar observations while deployed on adjacent games. Therefore, the adversary cannot differentiate $\mathcal{G}, \mathcal{G}'$ from the observation so he cannot infer the utility matrices with certainty from observations.

In this paper, we consider a generalization of Definition 2.2, the Rényi differential privacy (Rényi DP) (Mironov, 2017). The distribution of observations up to timestep t can be written as $\mu_{\mathcal{G}}^{(t)}(\boldsymbol{o}) \coloneqq \mathcal{P}_{\mathcal{G}}\left(\mathcal{R}_{\mathcal{G}}^{-1}\left(\boldsymbol{o}\right)\right)$ for $\boldsymbol{o} \in \mathcal{O}^t$ and thus we can define the (α, ϵ) -Rényi DP as follows.

Definition 2.3 ((α, ϵ) -Rényi Differential Privacy). For $\alpha > 1$ and $\epsilon \ge 0$, an iterative distributed algorithm for finding equilibria is (α, ϵ) -Rényi differentially private, if and only if for any two adjacent polymatrix game $\mathcal{G}, \mathcal{G}'$ and timestep t > 0,

$$D_{\alpha}\left(\mu_{\mathcal{G}}^{(t)}, \mu_{\mathcal{G}'}^{(t)}\right) := \frac{1}{\alpha - 1} \log \mathbb{E}_{\boldsymbol{o} \sim \mu_{\mathcal{G}'}^{(t)}} \left[\left(\frac{\mu_{\mathcal{G}}^{(t)}(\boldsymbol{o})}{\mu_{\mathcal{G}'}^{(t)}(\boldsymbol{o})}\right)^{\alpha} \right] \le \epsilon, \tag{2}$$

where ϵ is also called the *privacy budget*.

Definition 2.3 can be extended to $\alpha=1$ and $\alpha=+\infty$ by defining $D_1\left(\cdot,\cdot\right)=\lim_{\alpha\to 1^+}D_\alpha\left(\cdot,\cdot\right)$, which is KL-divergence (Mironov, 2017), and $D_\infty\left(\cdot,\cdot\right)\coloneqq\lim_{\alpha\to+\infty}D_\alpha\left(\cdot,\cdot\right)$. Moreover, Rényi DP can be converted to (ϵ,δ) -DP according to the following lemma from Mironov (2017).

Lemma 2.4. If an algorithm satisfies (α, ϵ) -Rényi DP for $\alpha > 1$, then for any $\delta \in (0, 1)$, the algorithm also satisfies $\left(\epsilon + \frac{\log(1/\delta)}{\alpha - 1}, \delta\right)$ -DP. Moreover, (∞, ϵ) -Rényi DP is equivalent to $(\epsilon, 0)$ -DP.

3 NECESSARY CONDITIONS FOR ACHIEVING HIGH ACCURACY AND DIFFERENTIAL PRIVACY SIMULTANEOUSLY

In this section, we will show the necessary conditions to achieve differential privacy and high accuracy in approximating equilibria in polymatrix games. Specifically, we will show that: (i) the adversary should not have access to all communication channels; and (ii) the accuracy metric cannot be the Euclidean distance to the equilibrium. If even one of the conditions is violated, one can find a constant $c_0 > 0$ so that when N is large enough, there exists a polymatrix game with N players so that the approximation error ζ and differential privacy budget ϵ cannot be guaranteed to be smaller than c_0 simultaneously in that game. The results are summarized in Figure 1 (b).

All hardness results in this section are constructed with zero-sum polymatrix games (Cai et al., 2016), a special class of polymatrix games, that is, polymatrix games $\mathcal{G} = ([N], E, \{\mathcal{A}_i\}_{i \in [N]}, \{U_{i,j}\}_{(i,j) \in E})$ such that, for any strategy profile π , $\sum_{i \in [N]} \sum_{j \in \mathcal{N}(i)} \pi_i^{\mathsf{T}} U_{i,j} \pi_j \equiv 0$.

One can see that when $U_{i,j} = -U_{j,i}$ for any $(i,j) \in E$, the polymatrix game is zero-sum. Since the coarse correlated equilibria (CCE) in zero-sum polymatrix games collapse to NE (Cai et al., 2016),

an algorithm that approximates CCE will approximate NE. Therefore, without loss of generality, in this section, we assume the algorithm will approximate the NE instead of CCE.

Assume that there is a distributed algorithm for finding equilibrium with the following guarantees of accuracy and privacy. In any polymatrix game \mathcal{G} with N players, the algorithm outputs a strategy profile π with distribution $\mu_{\mathcal{G}}$ on the observations. Here we omit the timestep superscript on $\mu_{\mathcal{G}}$ to emphasize that it is the distribution over all observations till termination of the algorithm. Then, there exists a parameter $\alpha \geq 1$ such that the algorithm satisfies the following accuracy and privacy guarantees in any N-player game \mathcal{G} :

Accuracy:
$$\frac{1}{N} \sum_{i=1}^{N} \mathbb{E} \left[\max_{\widehat{\pi}_i \in \Delta^{\mathcal{A}_i}} \left\langle \pi_i - \widehat{\pi}_i, \boldsymbol{g}_i^{\boldsymbol{\pi}} \right\rangle \right] \leq \zeta$$
 (3)

Privacy:
$$D_{\alpha}(\mu_{\mathcal{G}}, \mu_{\mathcal{G}'}) \leq \epsilon \quad \forall \mathcal{G} \sim \mathcal{G}'.$$
 (4)

When $\max_{\widehat{\pi}_i \in \Delta^{\mathcal{A}_i}} \langle \pi_i - \widehat{\pi}_i, \boldsymbol{g}_i^{\boldsymbol{\pi}} \rangle \leq \zeta$ for every player $i \in [N]$, the strategy profile $\boldsymbol{\pi}$ is called an ζ -approximate NE. The accuracy metric $\max_{\widehat{\pi}_i \in \Delta^{\mathcal{A}_i}} \langle \pi_i - \widehat{\pi}_i, \boldsymbol{g}_i^{\boldsymbol{\pi}} \rangle$ is also known as *exploitability*. It measures how much player i can benefit herself by unilaterally deviating from the current strategy profile $\boldsymbol{\pi}$, which should be 0 when $\boldsymbol{\pi}$ is the equilibrium. It is a weaker accuracy guarantee than the Euclidean distance to the equilibrium. Because when $\boldsymbol{\pi}$ is close to the equilibrium in terms of Euclidean distance, then the exploitability must be small. But the reverse does not always hold.

In the following, we will show that if the adversary has access to all communication channels, it is impossible to achieve low exploitability ζ and low privacy budget ϵ simultaneously.

Lemma 3.1. For any $N \ge 12$, there exists two zero-sum adjacent polymatrix games with N players so that for any algorithm guaranteeing (3) and (4), we have

$$\zeta \ge \min\left\{\frac{3\exp\left(-2\epsilon\right)}{112}, \frac{1}{112}\right\}. \tag{5}$$

Lemma 3.1 implies that even when N goes to infinity, once $\zeta \leq \frac{1}{112}$, we have $\epsilon \geq \frac{\log 3}{2}$. To illustrate the proof of Lemma 3.1, we can first consider a relaxation of (3). If $\mathbb{E}\left[\max_{\widehat{\pi}_i \in \Delta^{A_i}} \langle \pi_i - \widehat{\pi}_i, g_i^{\pi} \rangle\right] \leq \zeta$ holds for any player $i \in [N]$, then we can construct two adjacent polymatrix games, such that the approximate equilibrium any player $i \in [N]$ converges to differs much in two games, once the accuracy is lower than ζ in both games. Then, the adversary can distinguish those two games by computing the approximate equilibrium from his observations. While back to the original condition (3), we can use the pigeon-hole principle to construct a set of zero-sum polymatrix games, such that there exists two adjacent polymatrix games satisfying $\mathbb{E}\left[\max_{\widehat{\pi}_i \in \Delta^{A_i}} \langle \pi_i - \widehat{\pi}_i, g_i^{\pi} \rangle\right] \leq \mathcal{O}(\zeta)$ simultaneously for some player $i \in [N]$. The full proof is postponed to Appendix B.1, where we also provide a lower bound for (ϵ, δ) -DP.

Moreover, even when the adversary only has access to the observation of a single player, it is still impossible to find an accurate approximation of the equilibrium, in terms of Euclidean distance to the equilibrium, while guaranteeing privacy.

Let $\Delta^{\mathcal{A}_i,*}$ be the set of Nash equilibrium (NE) of player i in a zero-sum polymatrix game. Moreover, $\mu_{\mathcal{G},i}$ denotes marginal distribution on the observation of player i given the joint distribution $\mu_{\mathcal{G}}$ on all observations. Then, consider an algorithm with the following guarantee in any N-player game \mathcal{G} . There exists a parameter $\alpha \geq 1$ such that for any player $i \in [N]$, the algorithm satisfies that

Accuracy:
$$\mathbb{E}\left[\left\|\pi_{i} - \operatorname{Proj}_{\Delta^{\mathcal{A}_{i},*}}(\pi_{i})\right\|^{2}\right] \leq \zeta_{i}$$
 (6)

Privacy:
$$D_{\alpha}(\mu_{\mathcal{G}_i}, \mu_{\mathcal{G}_i}) < \epsilon_i \quad \forall \mathcal{G} \sim \mathcal{G}'.$$
 (7)

(6) guarantees that the output strategy of the algorithm is close to the set of NE, in terms of Euclidean distance. (7) states that for any two adjacent games $\mathcal{G} \sim \mathcal{G}'$, the Rényi divergence between the distribution of player i's observations is small. This implies that by accessing the communication channel of player i, the adversary cannot distinguish \mathcal{G} and \mathcal{G}' , so he cannot speculate the utility matrices.

Lemma 3.2. For any $N \ge 8$, there exists two zero-sum adjacent polymatrix games with N players so that for any algorithm guaranteeing (6) and (7), then

$$\frac{1}{N} \sum_{i=1}^{N} \zeta_i \ge \min \left\{ \frac{3}{8} \exp \left(-\frac{4}{N} \sum_{i=1}^{N} \epsilon_i \right), \frac{1}{16} \right\}. \tag{8}$$

Lemma 3.2 implies that once $\frac{1}{N}\sum_{i=1}^{N}\zeta_i\leq \frac{1}{16}$, the average differential privacy budget $\frac{1}{N}\sum_{i=1}^{N}\epsilon_i\geq \frac{\log 6}{4}$, even when N goes to infinity. The proof of Lemma 3.2 is based on the following observation:

for a two-player zero-sum game with row player's utility matrix as $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, when row player's

strategy changes from (0.51, 0.49) to (0.49, 0.51), the column player's best-response changes from (0,1) to (1,0). In this way, for two adjacent polymatrix games with the utility matrices on the edge (i,j) changes, player i,j's equilibria will also change. Then, by a careful construction of the utility matrices in the game, such changes in equilibria will gradually propagate to the rest of the players. Lastly, we can prove that in two adjacent games with the equilibrium differs much, it is impossible to compute the equilibria accurately while the communications between players (the adversary's observations) are similar. The full proof can be found in Appendix B.2.

4 METHODS

In this section, we will propose an iterative algorithm for finding coarse correlated equilibrium (CCE) in general-sum polymatrix games (Moulin & Vial, 1978). As we will show in Section 6, this algorithm guarantees differential privacy. We start by recalling a well-known connection between regret and convergence to CCE.

Proposition 4.1. For an iterative algorithm that produces strategy profile $\pi^{(t)}$ at timestep $t \in [T]$, its average strategy is the ζ -approximate CCE of a polymatrix game \mathcal{G} with N players if and only if for any player $i \in [N]$, $\max_{\widehat{\pi}_i \in \Delta^{\mathcal{A}_i}} \frac{1}{T} \sum_{t=1}^T \left\langle \boldsymbol{g}_i^{\boldsymbol{\pi}^{(t)}}, \pi_i^{(t)} - \widehat{\pi}_i \right\rangle \leq \zeta$.

By picking a timestep $t \in [T]$ uniformly at random and all players will play according to $\pi^{(t)}$, the expected increase of utility when player i deviates to $\widehat{\pi}_i$ is $\frac{1}{T} \sum_{t=1}^T \left\langle \boldsymbol{g}_i^{\boldsymbol{\pi}^{(t)}}, \pi_i^{(t)} - \widehat{\pi}_i \right\rangle$.

Let $\bar{g}_i^{(t)} := g_i^{\bar{\pi}^{(t)}}$, I^{A_i} as the identity matrix with index $A_i \times A_i$, and $\mathcal{N}\left(\mathbf{0}, \sigma^2 I^{A_i}\right)$ as the i.i.d. normal distribution over A_i with zero mean and σ^2 variance. Then, the algorithm is as follows: for any player $i \in [N]$, at each iteration $t \in \{0, 1, 2, \cdots, T-1, T\}$ for some T > 0 ($\pi_i^{(0)}$ is initialized as the uniform distribution over A_i), player i will sample a random noise $n_i^{(t)} \sim \mathcal{N}\left(\mathbf{0}, \sigma^2 I^{A_i}\right)$. Then, she will broadcast her strategy $\pi_i^{(t)}$ plus the noise to her neighbors. After receiving the noisy strategies from her neighbors, she will do one step of the projected gradient descent with the gradient $\bar{g}_i^{(t)}$, learning rate η , and weight-decay parameter τ_i . The complete algorithm is shown in Algorithm 1.

The intuition of $\tau_i \geq 0$ being inversely proportional to $|\mathcal{N}(i)|$ is that for players with fewer neighbors, her gradient is more sensitive to the variation of a single utility matrix, so more additional regularization (larger decay rate) is required to hide the difference of $\pi_i^{(t)}$ between adjacent games.

Moreover, to avoid observations from two adjacent games from increasingly diverging as updating more timesteps, the gradient descent of $\pi_i^{(t+1)}$ starts from $\bar{\pi}_i^{(t)}$ instead of $\pi_i^{(t)}$. However, we need to pay the price of suffering additional approximation error caused by the noise added to $\bar{\pi}_i^{(t)}$.

Remark 4.2. Let $\boldsymbol{n}^{(t)} \coloneqq \left(\boldsymbol{n}_1^{(t)}, \boldsymbol{n}_2^{(t)}, \cdots, \boldsymbol{n}_N^{(t)}\right)$. In Algorithm 1, the execution of the algorithm at timestep t is $\left(\boldsymbol{\pi}^{(t)}, \boldsymbol{n}^{(t)}\right)$. Then, $\mathcal{R}_{\mathcal{G}}\left(\left(\boldsymbol{\pi}^{(s)}, \boldsymbol{n}^{(s)}\right)_{s=1}^t\right) = \left(\bigcup_{i \in [N]} \left\{\pi_i^{(s)} + \boldsymbol{n}_i^{(s)}\right\}\right)_{s=1}^t$.

5 Convergence Analysis

This section will show the convergence of Algorithm 1 in general-sum polymatrix games.

Algorithm 1 Find CCE in polymatrix games with DP guarantee

```
Input: Player index i

Initialize \pi_i^{(0)} as uniform distribution over \mathcal{A}_i

Let \overline{\mathcal{N}} \coloneqq N \cdot \left(\sum_{i=1}^N \frac{1}{|\mathcal{N}(i)|}\right)^{-1} be the harmonic mean of players' degrees

Initialize \tau_i \leftarrow \frac{(\overline{\mathcal{N}})^{5/9}}{|\mathcal{N}(i)| \log N}

for t = 0, 1, ..., T do

Sample n_i^{(t)} \sim \mathcal{N}\left(\mathbf{0}, \sigma^2 \mathbf{I}^{A_i}\right)

Broadcast \pi_i^{(t)} + \mathbf{n}_i^{(t)} to neighbors j \in \mathcal{N}(i)

for j \in \mathcal{N}(i) do

Receive \pi_j^{(t)} + \mathbf{n}_j^{(t)} and compute \overline{\pi}_j^{(t)} \leftarrow \operatorname{Proj}_{\Delta^{A_j}}\left(\pi_j^{(t)} + \mathbf{n}_j^{(t)}\right)

end for

Gradient \overline{g}_i^{(t)} \leftarrow -\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} U_{i,j} \overline{\pi}_j^{(t)}

\pi_i^{(t+1)} \leftarrow \underset{\pi_i \in \Delta^{A_i}}{\operatorname{argmin}} \left\langle \pi_i, \overline{g}_i^{(t)} \right\rangle + \tau_i \|\pi_i\|^2 + \frac{1}{\eta} \|\pi_i - \overline{\pi}_i^{(t)}\|^2. (9)

{Note that (9) is equivalent to \pi_i^{(t+1)} \leftarrow \operatorname{Proj}_{\Delta^{A_i}}\left(\frac{\overline{\pi}_i^{(t)} - \eta \overline{g}_i^{(t)}}{1 + \eta \tau_i}\right).}
```

end for

Theorem 5.1. Consider Algorithm 1. Let $A = \max_{i \in [N]} |\mathcal{A}_i|$. The update-rule can achieve the following guarantee in any polymatrix game. For any T > 0, player $i \in [N]$ and strategy $\pi_i \in \Delta^{\mathcal{A}_i}$,

$$\frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbb{E}\left[\left\langle g_{i}^{\pi^{(t+1)}}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle\right] \\
\leq \frac{1}{\eta T} + A \frac{\sigma^{2}}{2\eta} + \left(\frac{2\eta^{2}}{\sigma} + \frac{7\sigma}{2}\right) A^{\frac{3}{2}} + \frac{1}{2\left(\overline{\mathcal{N}}\right)^{4/9} \log N} + \frac{2\eta\sqrt{A}}{\sigma\left(\overline{\mathcal{N}}\right)^{4/9} \log N}.$$
(10)

The proof sketch and the proof are postponed to Appendix C. Theorem 5.1 implies that the average strategy converges to coarse correlated equilibrium (CCE). Besides, it is noteworthy that when applying more noise for better differential privacy guarantees, *i.e.* larger σ , the convergence rate will worsen since the gradients provide less information.

With a fixed learning rate η and noise σ , the accuracy increases when the number of players N grows, while previous work (Gade et al., 2020; Ye et al., 2021; Wang et al., 2022; Shakarami et al., 2022) cannot, enabling achieving high accuracy and low differential budget simultaneously (see Section 7).

6 Analysis of Differential Privacy

In this section, we show that Algorithm 1 can simultaneously guarantee differential privacy. Since Lemma 3.1 states that the adversary should not have access to all communication channels, we will bound the Rényi DP when the adversary only has access to the communication channel of a single player, *i.e.* Rényi divergence of $\mu_{\mathcal{G},i}^{(t)}$, the marginal distribution of $\mu_{\mathcal{G}}^{(t)}$ over $\left(\pi_i^{(s)} + n_i^{(s)}\right)_{s=1}^t$. According to the composition rule of Rényi differential privacy (Mironov, 2017), when the adversary has access to multiple communication channels, the privacy budget is the product of the number of communication channels he has access to, and the privacy budget when he only has access to a single communication channel.

We break the analysis into two cases: dense graphs and sparse graphs. In the following, we will show that the privacy budget of Algorithm 1 is proportional to η^2, T but inversely proportional to σ^2 . Intuitively, when the learning rate η is larger or the algorithm runs more timesteps, the observation will diverge more due to the different utility matrices. But when we add more noise, *i.e.* larger σ , it will be harder to determine which distribution a sample $\pi_i^{(t)} + \boldsymbol{n}_i^{(t)}$ is from.

Theorem 6.1. Consider Algorithm 1 and any two adjacent polymatrix games $\mathcal{G}, \mathcal{G}'$ that differs at (v_1, v_2) . Let dist(i, j) be the length of the shortest path between players i, j, which is ∞ when i, j are not connected. The update-rule guarantees the following for any T > 0 and player $i \in [N]$,

$$\frac{1}{N} \sum_{i=1}^{N} D_{\alpha} \left(\mu_{\mathcal{G},i}^{(T)}, \mu_{\mathcal{G}',i}^{(T)} \right) \le \frac{\alpha \eta^{2}}{\sigma^{2}} \min \left\{ \clubsuit, \spadesuit \right\} T, \tag{11}$$

where

The proof sketch and a complete proof can be found in Appendix D. Theorem 6.1 implies that when the graph is dense, i.e. $\overline{\mathcal{N}} \geq N^p$ for some $p \in (0,1)$, will be sublinear in N. When the graph is sparse, i.e. $\mathcal{N}^{\max} \coloneqq \max_{i \in [N]} |\mathcal{N}(i)|$ is a constant, will be small since the most players will be distant from v_1, v_2 (approximately $\Omega(\log N)$). Therefore, even with a fixed learning rate η and noise σ , the privacy budget will decrease as the number of players grows, which is in stark contrast to previous work (Gade et al., 2020; Ye et al., 2021; Wang et al., 2022; Shakarami et al., 2022). The details are deferred to Section 7.

7 TRADE-OFF BETWEEN ACCURACY AND DIFFERENTIAL PRIVACY

According to Theorem 5.1 and Theorem 6.1, by picking η , σ and T carefully, when the number of players grows to infinity, we may obtain a zero privacy budget while converging to CCE.

Theorem 7.1. Consider Algorithm 1 with $\sigma = \frac{1}{\sqrt{T}}$ and any two adjacent polymatrix games $\mathcal{G}, \mathcal{G}'$ that differs at (v_1, v_2) . The update-rule guarantees the following,

$$\frac{1}{N} \sum_{i=1}^{N} \max \left\{ \frac{1}{T} \sum_{t=1}^{T} \mathbb{E} \left[\left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \right], 0 \right\} \\
\leq \frac{A+1}{\eta T} + 2A^{3/2} \eta^{2} \sqrt{T} + \frac{7A^{3/2}}{2\sqrt{T}} + \frac{1}{2(\overline{N})^{4/9} \log N} + \frac{2\eta \sqrt{AT}}{(\overline{N})^{4/9} \log N} \tag{14}$$

$$\frac{1}{N} \sum_{i=1}^{N} D_{\alpha} \left(\mu_{\mathcal{G},i}^{(T)}, \mu_{\mathcal{G}',i}^{(T)} \right) \le \alpha \eta^{2} \min \left\{ \clubsuit, \clubsuit \right\}. \tag{15}$$

The proof is simply substituting the values of σ into Theorem 5.1 and Theorem 6.1. When the graph is dense, i.e. $\overline{\mathcal{N}} \geq N^p$ for some p > 0, take $\eta = \frac{1}{T^{\frac{3}{4}1/3}}, T = \frac{N^{8p/9}}{(\log N)^4}$, then

$$\frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\mathbb{E}\left[\left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}},\boldsymbol{\pi}_{i}^{(t+1)}-\boldsymbol{\pi}_{i}\right\rangle \right]\leq\mathcal{O}\left(\frac{\left(\log N\right)^{2/3}}{N^{4p/27}}\right),$$

where the \mathcal{O} notation takes A as constant. Simultaneously, the differential budget guarantees that

$$\frac{1}{N} \sum_{i=1}^{N} D_{\alpha} \left(\mu_{\mathcal{G},i}^{(T)}, \mu_{\mathcal{G}',i}^{(T)} \right) \leq \mathcal{O} \left(\frac{\alpha \left(\log N \right)^{2/3}}{N^{4p/27}} \right).$$

Therefore, when the graph is dense, the accuracy and differential privacy budget of Algorithm 1 would both be better as the number of players increases.

Consider when the graph is sparse, *i.e.* \mathcal{N}^{\max} is a constant. If $\mathcal{N}^{\max} \geq 2$, let $\eta = \frac{1}{T \spadesuit^{1/3}}, T = (1 - \log_N \log N) \log_{\mathcal{N}^{\max}} N$, then $\spadesuit \leq \frac{4A}{\log N}$, since at most $\frac{2(\mathcal{N}^{\max})^T}{\mathcal{N}^{\max} - 1} \leq \frac{2N}{\log N}$ players satisfy

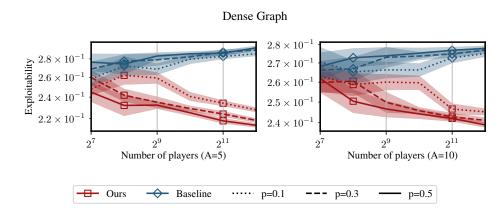


Figure 3: We evaluate the exploitability of our algorithm and a baseline on dense graphs. Each node (player) connects to another node independently with probability p, and duplicate edges are then removed. All players have action sets of size A. Both the baseline and our method are run under the same differential privacy budget.

$$\begin{split} \min \left\{ \mathrm{dist}(i,v_1), \mathrm{dist}(i,v_2) \right\} &< T. \text{ If } \mathcal{N}^{\max} = 1, \spadesuit \leq \frac{2}{N} \text{ for any } T > 0. \text{ Therefore,} \\ & \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbb{E} \left[\left\langle g_i^{\pi^{(t+1)}}, \pi_i^{(t+1)} - \pi_i \right\rangle \right] \leq \mathcal{O} \left(\frac{1}{(\log N)^{1/3}} \right) \\ & \frac{1}{N} \sum_{i=1}^{N} D_{\alpha} \left(\mu_{\mathcal{G},i}^{(T)}, \mu_{\mathcal{G}',i}^{(T)} \right) \leq \mathcal{O} \left(\frac{\alpha}{(\log N)^{1/3}} \right). \end{split}$$

Compared to previous work (Gade et al., 2020; Ye et al., 2021; Wang et al., 2022; Shakarami et al., 2022), the main contribution of Theorem 7.1 is that for any desired accuracy $\zeta_0 > 0$ and differential privacy budget $\epsilon_0 > 0$, we can ensure our accuracy and differential privacy budget to be lower than ζ_0 , ϵ_0 simultaneously when N is large enough. In contrast, to keep accuracy smaller than ζ_0 , previous work will suffer a differential privacy budget larger than ϵ_0 , no matter how large N is, and vice versa.

We emphasize that Theorem 7.1 bounds the differential privacy budget when the adversary can access a single communication channel. By the composition rule for Rényi differential privacy (Mironov, 2017) (see the beginning of Section 6), the overall privacy budget scales linearly with the number of channels the adversary accesses, specifically, it equals the number of accessed channels times the bound in Theorem 7.1.

8 EXPERIMENTS

In this section, we evaluate the exploitability $\left(\max_{\widehat{\pi}_i \in \Delta^{\mathcal{A}_i}} \langle \pi_i - \widehat{\pi}_i, \boldsymbol{g}_i^{\boldsymbol{\pi}} \rangle\right)$ of our method and a baseline algorithm adapted from Huang et al. (2015). Under the same differential-privacy budget, our method consistently attains lower exploitability. Moreover, as the number of players increases, the exploitability of our method decreases, whereas the baseline's does not. Additional experimental results are provided in Appendix E.

9 CONCLUSION AND LIMITATIONS

This paper shows the necessary conditions to find the equilibria of polymatrix games with DP guarantees. Moreover, to justify the conditions, we propose an algorithm to find the equilibria with DP guarantees under the conditions. Moreover, as the number of players in the game increases, the algorithm will gradually achieve perfect accuracy and DP guarantees. Lastly, we provide empirical evidence to justify the algorithm in Appendix E. The limitation of this work is that we only discuss polymatrix games. Therefore, an interesting future direction is to extend it to other games, such as normal-form games and extensive-form games.

10 ETHICS STATEMENT

This paper presents work that aims to advance the field of Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

11 REPRODUCIBILITY STATEMENT

We include code in the supplementary materials. Moreover, the proofs are presented in Appendices B to D.

12 USE OF LARGE LANGUAGE MODELS

In this paper, we use large language models (LLMs) to improve writing, e.g., by correcting grammatical errors.

REFERENCES

- Naman Agarwal and Karan Singh. The price of differential privacy for online learning. In *International Conference on Machine Learning (ICML)*, 2017.
- Raef Bassily, Cristóbal Guzmán, and Michael Menart. Differentially private algorithms for the stochastic saddle point problem with optimal rates for the strong gap. In *Conference on Learning Theory (COLT)*, 2023.
- Digvijay Boob and Cristóbal Guzmán. Optimal algorithms for differentially private stochastic monotone variational inequalities and saddle-point problems. *Mathematical Programming*, 204(1): 255–297, 2024.
- Yang Cai, Ozan Candogan, Constantinos Daskalakis, and Christos Papadimitriou. Zero-sum polymatrix games: A generalization of minmax. *Mathematics of Operations Research*, 41(2):648–655, 2016.
- Rachel Cummings, Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. Privacy and truthful equilibrium selection for aggregative games. In *Web and Internet Economics: 11th International Conference, WINE 2015, Amsterdam, The Netherlands, December 9-12, 2015, Proceedings 11*, pp. 286–299. Springer, 2015.
- Giuseppe De Nittis, Alberto Marchesi, and Nicola Gatti. Computing the strategy to commit to in polymatrix games. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- Argyrios Deligkas, John Fearnley, Rahul Savani, and Paul Spirakis. Computing approximate nash equilibria in polymatrix games. *Algorithmica*, 77(2):487–514, 2017.
- Ahmed Said Donmez, Yuksel Arslantas, and Muhammed O Sayin. Team collaboration vs competition: New fictitious play dynamics for multi-team zero-sum games. *arXiv preprint arXiv:2402.02147*, 2024.
 - Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pp. 265–284. Springer, 2006.
- David Evangelista, Yuri Saporito, and Yuri Thamsten. Price formation in financial markets: a game-theoretic perspective. *arXiv preprint arXiv:2202.11416*, 2022.
- Shripad Gade, Anna Winnicki, and Subhonmesh Bose. On privatizing equilibrium computation in aggregate games over networks. *IFAC-PapersOnLine*, 53(2):3272–3277, 2020.
- Manuel Gil, Fady Alajaji, and Tamas Linder. Rényi divergence measures for commonly used univariate continuous distributions. *Information Sciences*, 249:124–131, 2013.

- Abhradeep Guha Thakurta and Adam Smith. (nearly) optimal algorithms for private online learning in full-information and bandit settings. *Neural Information Processing Systems (NeurIPS)*, 2013.
- Zhenqi Huang, Sayan Mitra, and Nitin Vaidya. Differentially private distributed optimization. In
 International Conference on Distributed Computing and Networking (ICDCN), 2015.
 - Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially private online learning. In *Conference on Learning Theory (COLT)*, 2012.
 - Elena Janovskaja. Equilibrium points in polymatrix games. *Lithuanian Mathematical Journal*, 8(2): 381–384, 1968.
 - Michael Kearns, Mallesh Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: Incentives and privacy. In *Innovations in Theoretical Computer Science (ITCS)*, 2014.
 - Mingyang Liu, Gabriele Farina, and Asuman Ozdaglar. A policy-gradient approach to solving imperfect-information games with iterate convergence. *arXiv preprint arXiv:2408.00751*, 2024.
 - Ilya Mironov. Rényi differential privacy. In 2017 IEEE 30th computer security foundations symposium (CSF), pp. 263–275. IEEE, 2017.
 - H. Moulin and J.-P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *Int. J. Game Theory*, 7(3):201–221, September 1978. ISSN 1432-1270. doi: 10.1007/BF01769190.
 - Erfan Nozari, Pavankumar Tallapragada, and Jorge Cortés. Differentially private distributed convex optimization via functional perturbation. *IEEE Transactions on Control of Network Systems*, 5(1): 395–408, 2016.
 - Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Neural Information Processing Systems (NeurIPS)*, 2019.
 - Ryan M Rogers and Aaron Roth. Asymptotically truthful equilibrium selection in large congestion games. In *ACM Conference on Economics and Computation (EC)*, 2014.
 - Mehran Shakarami, Claudio De Persis, and Nima Monshizadeh. Distributed dynamics for aggregative games: Robustness and privacy guarantees. *International Journal of Robust and Nonlinear Control*, 32(9):5048–5069, 2022.
 - Jimin Wang, Ji-Feng Zhang, and Xingkang He. Differentially private distributed algorithms for stochastic aggregative games. *Automatica*, 142:110440, 2022.
 - Yongqiang Wang and Tamer Başar. Ensuring both almost sure convergence and differential privacy in nash equilibrium seeking on directed graphs. *IEEE Transactions on Automatic Control*, 2024.
 - Yongqiang Wang and Angelia Nedić. Differentially-private distributed algorithms for aggregative games with guaranteed convergence. *IEEE Transactions on Automatic Control*, 2024.
 - Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek, and H Vincent Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE transactions on information forensics and security*, 15:3454–3469, 2020.
 - Zhenhuan Yang, Shu Hu, Yunwen Lei, Kush R Vashney, Siwei Lyu, and Yiming Ying. Differentially private sgda for minimax problems. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2022.
 - Maojiao Ye, Guoqiang Hu, Lihua Xie, and Shengyuan Xu. Differentially private distributed nash equilibrium seeking for aggregative games. *IEEE Transactions on Automatic Control*, 67(5): 2451–2458, 2021.
 - Liang Zhang, Kiran K Thekumparampil, Sewoong Oh, and Niao He. Bring your own algorithm for optimal differentially private stochastic minimax optimization. *Neural Information Processing Systems (NeurIPS)*, 2022.

A RELATED WORK

In this section, we discuss prior work specifically targeted at distributed *equilibrium computation*. We omit a discussion of prior work on distributed (single-agent) optimization as it is hardly applicable to our purposes, for several important reasons. For one, distributed optimization with differential privacy guarantee (Huang et al., 2015; Nozari et al., 2016; Wei et al., 2020) aims to find the minimum of the weighted sum of players' utility functions. However, distributed equilibrium finding is more complex since it aims to find the equilibrium point of those utility functions. Furthermore, the target output of the distributed optimization is of constant dimensionality with respect to the number of agents involved in the optimization process, such as the best parameters for a neural network (Wei et al., 2020). However, the output's dimension of equilibrium finding scales linearly with respect to the number of agents, since we need to specify the strategy of each individual player.

Online Learning with Differential Privacy. Prior work, including Jain et al. (2012); Guha Thakurta & Smith (2013); Agarwal & Singh (2017), has studied no-regret learning with DP guarantee, focusing on ensuring that an adversary cannot determine the utility (loss) vector at any single timestep. However, that condition does not imply differential privacy with respect to the payoff function itself. For instance, in a two-player normal-form game, the utility vector at timestep t for player 1 is $U\pi_2^{(t)}$, where U is the utility matrix and $\pi_2^{(t)}$ is the strategy of player 2 at timestep t. In this case, an algorithm that reveals $U\pi_2^{(t)} + n^{(t)}$, where $n^{(t)}$ is sampled from a multivariate Gaussian, ensures differential privacy with respect to each individual utility vector. However, consider the case when $\pi_2^{(t)}$ is sampled uniformly on each axis for any timestep t. By taking the average over all utility vectors, the Gaussian noise cancels out, and we will get the average of U's columns. Therefore, U is not differentially private, while each utility vector is still private.

Differential Privacy in Minimax Optimization. Yang et al. (2022); Zhang et al. (2022); Bassily et al. (2023); Boob & Guzmán (2024) analyze differential privacy in minimax optimization, but their setting differs from ours in several aspects. First, they study two-player zero-sum games, whereas we consider multi-player general-sum polymatrix games (a class that subsumes two-player zero-sum). Second, they assume each player's utility decomposes as an average of multiple component utilities and aim to prevent an adversary from inferring any single component. This, however, does not protect the average itself, the true game utility, which is precisely what we seek to keep private.

Differential Privacy in Games. Other prior work, including Gade et al. (2020); Ye et al. (2021); Wang et al. (2022); Shakarami et al. (2022), has focused on finding equilibrium in aggregative games with differential privacy guarantees. However, they cannot achieve accuracy and differential privacy simultaneously (e.g., privacy budget $\mathcal{O}(\epsilon)$ and accuracy $\mathcal{O}\left(\frac{1}{\epsilon}\right)$). In other words, to achieve a fixed accuracy, the privacy budget will not continue to decrease while the number of players increases. The reason is that they did not exploit the game's structure to cancel out the noise imposed during the update. Therefore, accuracy will be harmed when achieving a low differential privacy budget by adding a large amount of noise. Furthermore, the privacy guarantee of Wang & Nedić (2024); Wang & Başar (2024) is weak because the adversary may determine the value of the utility function in a small region around the equilibrium point. Therefore, when the utility function is multilinear, such as polymatrix games, normal-form games, and extensive-form games, the adversary can fully determine the utility functions with certainty. Instead, in this paper, we exploit the structure of polymatrix games, whether the associated graph is dense or sparse, to mitigate the impact of the imposed noise, thereby achieving accuracy and differential privacy simultaneously. Specifically, we can achieve privacy budget $\mathcal{O}\left(\frac{\epsilon}{f(N)}\right)$ and accuracy $\mathcal{O}\left(\frac{1}{\epsilon g(N)}\right)$, where f(N), g(N) are monotone functions with respect to the number of players N and goes to infinity as $N \to \infty$. Details can be found in Theorem 7.1.

B PROOF OF SECTION 3

The section is arranged as follows. For both Lemma 3.1 and Lemma 3.2, we will show that there exists a polymatrix game, so that achieving accuracy and (ϵ, δ) -DP simultaneously is impossible. Then, we further show that achieving accuracy and Rényi DP is impossible by showing that for

 $(1,\epsilon)$ -Rényi DP. Because the α -Rényi divergence grows monotonically with respect to α (Mironov, 2017).

B.1 Proof of Lemma 3.1

Lemma 3.1. For any $N \ge 12$, there exists two zero-sum adjacent polymatrix games with N players so that for any algorithm guaranteeing (3) and (4), we have

$$\zeta \ge \min\left\{\frac{3\exp\left(-2\epsilon\right)}{112}, \frac{1}{112}\right\}. \tag{5}$$

Table 1: i is the row player and j is the column player, the utility in the table is $U_{i,j}(a_i, a_j) = -U_{i,i}(a_i, a_i)$.

	a_1	
a_1	1	
a_2	0	

	a_1
a_1	0
a_2	1

	a_1	a_2
a_1	0.5	0
a_2	0	0.5

Proof. Consider a zero-sum polymatrix game with 3N players, where its corresponding graph is a chain, i.e. the edge set $E = \{(i,i+1)\}_{i \in [3N-1]}$. For $i \in [N]$, $U_{3i-1,3i-2}$ is Table 1 (left) and $U_{3i,3i-1}$ is Table 1 (right). All other utility matrices are zero matrices. Therefore, the only equilibrium is all players playing a_1 . Let denote the game as \mathcal{G}^{\emptyset} .

In this case, for any $i \in [N]$, when changing $U_{3i-1,3i-2}$ to Table 1 (middle), the equilibrium of both 3i-1 and 3i will be choosing a_2 deterministically, and the equilibrium of other players keep unchanged. Let denote this new game as $\mathcal{G}^{\{i\}}$.

B.1.1 LOWERBOUNDS ON (ϵ, δ) -DP

This section will show that if the algorithm satisfies (ϵ, δ) -DP, then

$$\zeta \ge \frac{1 - \delta}{28\left(1 + e^{\epsilon}\right)}.\tag{16}$$

Suppose the distributed algorithm satisfies the following guarantee of accuracy and privacy. Let $g_i^{\pi}(\mathcal{G})$ be the gradient of player i in \mathcal{G} and $\pi^{\{i\}}$ be the output of the algorithm in $\mathcal{G}^{\{i\}}$. Then,

$$\forall i \in [N], \sum_{j=1}^{3N} \mathbb{E} \left[\max_{\widehat{\pi}_j \in \Delta^{\mathcal{A}_j}} \left\langle \pi_j^{\{i\}} - \widehat{\pi}_j, \boldsymbol{g}_j^{\boldsymbol{\pi}^{\{i\}}}(\mathcal{G}^{\{i\}}) \right\rangle \right] \leq N\zeta$$
 (Accuracy)

$$\forall i \in [N], \text{ set } \mathcal{C}, \quad \Pr(\boldsymbol{\pi}^{\emptyset} \in \mathcal{C}) \le e^{\epsilon} \Pr(\boldsymbol{\pi}^{\{i\}} \in \mathcal{C}) + \delta,$$
 (Privacy)

Moreover, it is easy to verify that for any strategy profile $\pi \in \times_{i=1}^{3N} \Delta^{\mathcal{A}_i}$ and $i \in [N], k \in \{0\} \cup [N]$ (let $\mathcal{G}^{\{0\}} = \mathcal{G}^{\emptyset}$ for simplicity), we have

$$\max_{\widehat{\pi}_{3i-1} \in \Delta^{\mathcal{A}_{3i-1}}} \left\langle \pi_{3i-1} - \widehat{\pi}_{3i-1}, \boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}}(\mathcal{G}^{\{k\}}) \right\rangle + \max_{\widehat{\pi}_{3i} \in \Delta^{\mathcal{A}_{3i}}} \left\langle \pi_{3i} - \widehat{\pi}_{3i}, \boldsymbol{g}_{3i}^{\boldsymbol{\pi}}(\mathcal{G}^{\{k\}}) \right\rangle \ge \begin{cases} 0.5\pi_{3i}(a_2) & i \neq k \\ 0.5\pi_{3i}(a_1) & i = k. \end{cases}$$

Since the algorithm is distributed, the findr for each player can be written as $f_i : \mathcal{O}^T \times \mathbb{R}^{\mathcal{A}_i \times \mathcal{A}_{i-1}} \times \mathbb{R}^{\mathcal{A}_i \times \mathcal{A}_{i+1}} \times \Delta^{\mathcal{A}_i} \to \mathbb{R}$ (let $\mathcal{A}_0 = \mathcal{A}_{N+1} = \emptyset$ for notational simplicity), which is a conditional distribution over all possible strategies in $\Delta^{\mathcal{A}_i}$, conditioned on all past observations and player's utility matrices.

Let $adv(\boldsymbol{o}, \pi_{3i}) := f_{3i}(\boldsymbol{o}_{3i}, \boldsymbol{U}_{3i,3i-1}, \pi_{3i})$ for $\boldsymbol{o} \in \mathcal{O}^T$ and $\pi_{3i} \in \Delta^{\mathcal{A}_{3i}}$, where $\boldsymbol{U}_{3i,3i-1}$ is the utility matrix between player 3i, 3i-1 in \mathcal{G}^{\emptyset} . Note that since the adversary cannot access utility matrices, $adv(\boldsymbol{o}, \pi_{3i})$ is fixed even though the observation now comes from games $\mathcal{G}^{\{k\}}$ for some $k \in [N]$.

Then, the adversary may sample $\pi_{3i}^{\text{adv},k}$ according to $\text{adv}(\boldsymbol{o}^k,\pi_{3i})$, where $\boldsymbol{o}^k \in \mathcal{O}^T$ is the observation in game $\mathcal{G}^{\{k\}}$. Because the utility matrices of player 3i in \mathcal{G}^{\emptyset} and $\mathcal{G}^{\{k\}}$ are identical, we have

$$\begin{aligned} 0.5 \mathbb{E} \left[\pi_{3i}^{\mathrm{adv},k}(a_2) \right] = & 0.5 \mathbb{E} \left[\pi_{3i}^{\{k\}}(a_2) \right] \\ \leq & \mathbb{E} \left[\max_{\widehat{\pi}_{3i-1} \in \Delta^{\mathcal{A}_{3i-1}}} \left\langle \pi_{3i-1}^{\{k\}} - \widehat{\pi}_{3i-1}, \boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\{k\}}}(\mathcal{G}^{\{k\}}) \right\rangle + \max_{\widehat{\pi}_{3i} \in \Delta^{\mathcal{A}_{3i}}} \left\langle \pi_{3i}^{\{k\}} - \widehat{\pi}_{3i}, \boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\{k\}}}(\mathcal{G}^{\{k\}}) \right\rangle \right]. \end{aligned}$$

Similar bounds also holds for $\mathbb{E}\left[\pi_{3i}^{\mathrm{adv}}(a_1)\right]$ when i=k.

Then, by Markov inequality and the definition of (ϵ, δ) -DP, we have

$$4\mathbb{E}\left[\max_{\widehat{\pi}_{3i-1}\in\Delta^{A_{3i-1}}}\left\langle\pi_{3i-1}^{\emptyset}-\widehat{\pi}_{3i-1},\boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\emptyset}}(\mathcal{G}^{\emptyset})\right\rangle + \max_{\widehat{\pi}_{3i}\in\Delta^{A_{3i}}}\left\langle\pi_{3i}^{\emptyset}-\widehat{\pi}_{3i},\boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\emptyset}}(\mathcal{G}^{\emptyset})\right\rangle\right]$$

$$\geq 2\mathbb{E}\left[\pi_{3i}^{\text{adv},\emptyset}(a_{2})\right]$$

$$\geq \Pr\left(\pi_{3i}^{\text{adv},\emptyset}(a_{2})\geq 0.5\right)$$

$$= \Pr\left(\pi_{3i}^{\text{adv},\emptyset}(a_{1})\leq 0.5\right)$$

$$\geq e^{-\epsilon}\Pr\left(\pi_{3i}^{\text{adv},\{i\}}(a_{1})\leq 0.5\right) - \delta e^{-\epsilon}$$

$$\geq e^{-\epsilon}\left(1-2\mathbb{E}\left[\pi_{3i}^{\text{adv},\{i\}}(a_{1})\right]\right) - \delta e^{-\epsilon}$$

$$\geq e^{-\epsilon}\left(1-4\mathbb{E}\left[\max_{\widehat{\pi}_{3i-1}\in\Delta^{A_{3i-1}}}\left\langle\pi_{3i-1}^{\{i\}}-\widehat{\pi}_{3i-1},\boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\{i\}}}(\mathcal{G}^{\{i\}})\right\rangle + \max_{\widehat{\pi}_{3i}\in\Delta^{A_{3i}}}\left\langle\pi_{3i}^{\{i\}}-\widehat{\pi}_{3i},\boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\{i\}}}(\mathcal{G}^{\{i\}})\right\rangle\right]\right) - \delta e^{-\epsilon}.$$

Moreover, for each $i \in [N]$ and $\mathcal{G}^{\{i\}}$, we can further define $\mathcal{G}^{\{i,j\}}$ for $i \neq j \in [N]$ as the adjacent game to $\mathcal{G}^{\{i\}}$ that differs from $\mathcal{G}^{\{i\}}$ only at $U_{3j-1,3j-2}$, by changing it from Table 1 (left) to Table 1 (middle). Also note that $\mathcal{G}^{\{i,j\}}$ is also adjacent to $\mathcal{G}^{\{j\}}$. Such a process can be recursively applied to $\mathcal{G}^{\{i,j\}}$ and finally we will get 2^N games $\{\mathcal{G}^{\mathcal{S}}\}_{\mathcal{S}\subset [N]}$.

Lemma B.1. Consider the 2^N games $\{\mathcal{G}^S\}_{S\subseteq [N]}$ constructed above. There must exist two adjacent games \mathcal{G}^S and $\mathcal{G}^{S\cup \{i\}}$ with $i\in [N]\setminus S$, so that

$$\mathbb{E}\left[\max_{\widehat{\pi}_{3i-1}\in\Delta^{\mathcal{A}_{3i-1}}}\left\langle \pi_{3i-1}^{\mathcal{S}} - \widehat{\pi}_{3i-1}, \boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\mathcal{S}}}(\mathcal{G}^{\mathcal{S}})\right\rangle + \max_{\widehat{\pi}_{3i}\in\Delta^{\mathcal{A}_{3i}}}\left\langle \pi_{3i}^{\mathcal{S}} - \widehat{\pi}_{3i}, \boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\mathcal{S}}}(\mathcal{G}^{\mathcal{S}})\right\rangle\right] \leq 7\zeta \qquad (18)$$

$$\mathbb{E}\left[\max_{\widehat{\pi}_{3i-1}\in\Delta^{\mathcal{A}_{3i-1}}}\left\langle \pi_{3i-1}^{\mathcal{S}\cup\{i\}} - \widehat{\pi}_{3i-1}, \boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\mathcal{S}\cup\{i\}}}(\mathcal{G}^{\mathcal{S}\cup\{i\}})\right\rangle + \max_{\widehat{\pi}_{3i}\in\Delta^{\mathcal{A}_{3i}}}\left\langle \pi_{3i}^{\mathcal{S}\cup\{i\}} - \widehat{\pi}_{3i}, \boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\mathcal{S}\cup\{i\}}}(\mathcal{G}^{\mathcal{S}\cup\{i\}})\right\rangle\right] \leq 7\zeta.$$

$$(19)$$

The proof is postponed to Appendix B.3. Since (17) also holds for the $\mathcal{G}^{\mathcal{S}}$ and $\mathcal{G}^{\mathcal{S} \cup \{i\}}$ in Lemma B.1, we have

$$28\left(e^{-\epsilon}+1\right)\zeta \ge e^{-\epsilon}-\delta e^{-\epsilon}.$$

Therefore, $\zeta \geq \frac{e^{-\epsilon} - \delta e^{-\epsilon}}{28(e^{-\epsilon} + 1)} = \frac{1 - \delta}{28(1 + e^{\epsilon})}$.

B.1.2 LOWERBOUNDS ON (α, ϵ) -RÉNYI DP

This section will show that if the algorithm satisfies (ϵ, δ) -DP, then

$$\frac{1}{N} \sum_{i=1}^{N} \zeta_i \ge \frac{1 - \max_{i \in [N]} \delta_i}{1 + \exp\left(\frac{1}{N} \sum_{i=1}^{N} \epsilon_i\right)} - \frac{2}{N}.$$
 (20)

By monotonicity of Rényi Divergence (Mironov, 2017), we only need to show the lowerbound of KL-divergence. By further post-processing on $\pi_{3i}^{\mathrm{adv},\mathcal{S}}$ for $i\in[N]$ and $\mathcal{S}\subseteq[N]$, the adversary can output whether $\pi_{3i}^{\mathrm{adv},\mathcal{S}}(a_2)\geq 0.5$. Let the distribution of the output be $\mu_{3i}^{\mathcal{S}}$. Then, for the adjacent games \mathcal{G}^{\emptyset} , $\mathcal{G}^{\{i\}}$, where $i\in[N]$, we have

$$\begin{split} \epsilon \geq \mathrm{KL}\left(\mu_{3i}^{\emptyset}, \mu_{3i}^{\{i\}}\right) &= \Pr\left(\pi_{3i}^{\mathrm{adv},\emptyset}(a_2) \geq 0.5\right) \log \frac{\Pr\left(\pi_{3i}^{\mathrm{adv},\emptyset}(a_2) \geq 0.5\right)}{\Pr\left(\pi_{3i}^{\mathrm{adv},\{i\}}(a_2) \geq 0.5\right)} \\ &+ \left(1 - \Pr\left(\pi_{3i}^{\mathrm{adv},\emptyset}(a_2) \geq 0.5\right)\right) \log \frac{1 - \Pr\left(\pi_{3i}^{\mathrm{adv},\emptyset}(a_2) \geq 0.5\right)}{1 - \Pr\left(\pi_{3i}^{\mathrm{adv},\{i\}}(a_2) \geq 0.5\right)} \\ &\stackrel{(i)}{\geq} (1 - \exp l_i^{\emptyset}) \log \frac{1 - \exp l_i^{\emptyset}}{\exp l_i^{\{i\}}} + \exp l_i^{\emptyset} \log \frac{\exp l_i^{\emptyset}}{1 - \exp l_i^{\emptyset}}, \end{split}$$

where

$$\operatorname{expl}_{i}^{\mathcal{S}} \coloneqq 4\mathbb{E}\left[\max_{\widehat{\pi}_{3i-1} \in \Delta^{\mathcal{A}_{3i-1}}} \left\langle \pi_{3i-1}^{\mathcal{S}} - \widehat{\pi}_{3i-1}, \boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\mathcal{S}}}(\mathcal{G}^{\mathcal{S}}) \right\rangle + \max_{\widehat{\pi}_{3i} \in \Delta^{\mathcal{A}_{3i}}} \left\langle \pi_{3i}^{\mathcal{S}} - \widehat{\pi}_{3i}, \boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\mathcal{S}}}(\mathcal{G}^{\mathcal{S}}) \right\rangle \right].$$

$$(i) \text{ is because } \Pr\left(\pi_{3i}^{\mathrm{adv},\emptyset}(a_2) \geq 0.5\right) \leq \exp \mathrm{l}_i^\emptyset \text{ and } \Pr\left(\pi_{3i}^{\mathrm{adv},\{i\}}(a_2) \geq 0.5\right) \geq 1 - \exp \mathrm{l}_i^{\{i\}}.$$

By Lemma B.1, there exists $S \subseteq [N]$ and $i \in [N] \setminus s$ so that $\exp l_i^S, \exp l_i^{S \cup \{i\}} \le 28\zeta$. Therefore,

$$\epsilon \ge (1 - 28\zeta) \log \frac{1 - 28\zeta}{28\zeta} + 28\zeta \log \frac{28\zeta}{1 - 28\zeta} = (1 - 56\zeta) \log \frac{1 - 28\zeta}{28\zeta}.$$

It implies that when $\zeta \leq \frac{1}{112}$, $\epsilon \geq \frac{1}{2} \log \frac{3}{112\zeta}$ so that $\zeta \geq \frac{3}{112 \exp(2\epsilon)}$.

B.2 Proof of Lemma 3.2

Lemma 3.2. For any $N \ge 8$, there exists two zero-sum adjacent polymatrix games with N players so that for any algorithm guaranteeing (6) and (7), then

$$\frac{1}{N} \sum_{i=1}^{N} \zeta_i \ge \min \left\{ \frac{3}{8} \exp \left(-\frac{4}{N} \sum_{i=1}^{N} \epsilon_i \right), \frac{1}{16} \right\}. \tag{8}$$

Proof. Consider a zero-sum polymatrix game, where its corresponding graph is a chain. Specifically, there is an edge $(i,j) \in E$, if and only if i = j - 1. Moreover, each player has two actions and $U_{i,i+1}$ is shown in Table 2 (left).

Table 2: i is the row player and j is the column player, the utility in the table is $U_i(a_i, a_j)$ and $U_i(a_i, a_i) = -U_i(a_i, a_i)$.

	a_1	a_2		
a_1	0.5	$0.5 - \epsilon_i$	a_1	0.
a_2	$0.5 - 3\epsilon_i$	$0.5-2\epsilon_i$	a_2	0

 $\begin{array}{c|cccc} & a_1 & a_2 \\ \hline a_1 & 0.5 - 2\epsilon_i & 0.5 - 3\epsilon_i \\ \hline a_2 & 0.5 - \epsilon_i & 0.5 \\ \end{array}$

Moreover, $0.1 = \epsilon_1 = 10\epsilon_2 \ge 10^2\epsilon_3 = \dots = 10^{N-1}\epsilon_N > 0$. Then, the Nash equilibrium (NE) is $\left(a_1, a_2, a_1, a_2, \cdots, a_{\frac{(-1)^N+3}{2}}\right)$. However, when the utility matrix between player (1, 2) changes to

Table 2 (right), the NE becomes $\left(a_2, a_1, a_2, a_1, \cdots, a_{\frac{(-1)^{N+1}+3}{2}}\right)$. Let the original game be denoted as \mathcal{G}^0 and the new game as \mathcal{G}^1 .

With the observation $o_i \in \mathcal{O}^T$ of any player i > 2, the adversary may sample strategy π_i^{adv} from the distribution $\text{adv}(o_i, \pi_i) \coloneqq f_i(o_i, U_{i,i-1}, U_{i,i+1}, \pi_i)$ to obtain the equilibrium of player i, where $U_{i,i-1}, U_{i,i+1}$ are both utility matrices in \mathcal{G}^0 .

B.2.1 LOWERBOUNDS ON (ϵ, δ) -DP

For each player $i \in [N] \setminus \{1, 2\}$, assume the distributed algorithm may achieve the following guarantee.

$$\forall k \in \{0,1\}, \quad \mathbb{E}\left[\left\|\pi_i^k - \pi_i^{*,k}\right\|^2\right] \le \zeta_i$$
 (Accuracy)

$$\forall \, \mathcal{C}_i \subseteq \Delta^{\mathcal{A}_i}, \quad \Pr(\pi_i^0 \in \mathcal{C}_i) \le e^{\epsilon_i} \Pr(\pi_i^1 \in \mathcal{C}_i) + \delta_i, \tag{Privacy}$$

where π_i^0, π_i^1 is the output of the algorithm when deploying on player i for $\mathcal{G}^0, \mathcal{G}^1$.

Let $\pi_i^{\mathrm{adv},k}$ be the random variable sampled from $\mathrm{adv}(o_i^k,\cdot)$, where o_i^k is the observation in \mathcal{G}^k with $k \in \{0,1\}$. Since post-processing will not decrease the strength of DP, observing $\pi_i^{\mathrm{adv},k}$ instead of o_i^k will leak less information. Moreover, since the utility matrices of players i>2 do not change, for any player i>2 and $k\in\{0,1\}$, we have

$$\mathbb{E}\left[\left\|\pi_i^k - \pi_i^{*,k}\right\|^2\right] = \mathbb{E}\left[\left\|\pi_i^{\mathrm{adv},k} - \pi_i^{*,k}\right\|^2\right].$$

Let the equilibrium of \mathcal{G}^0 be $\pi^{*,0} \coloneqq ((1,0),(0,1),\cdots)$ and the equilibrium of \mathcal{G}^1 be $\pi^{*,1} \coloneqq ((0,1),(1,0),\cdots)$. Then, by Markov inequality, for any $k \in \{0,1\}$ and i>2, we have

$$\Pr\left(\left\|\pi_i^{\text{adv},k} - \pi_i^{*,k}\right\|^2 \ge 1\right) \le \frac{\mathbb{E}\left[\left\|\pi_i^{\text{adv},k} - \pi_i^{*,k}\right\|^2\right]}{1} = \frac{\mathbb{E}\left[\left\|\pi_i^k - \pi_i^{*,k}\right\|^2\right]}{1} \le \zeta_i.$$

Therefore, according to the privacy guarantee, for any i > 2, we have

$$1 - \zeta_{i} \leq \Pr\left(\left\|\pi_{i}^{\text{adv},0} - \pi_{i}^{*,0}\right\|^{2} \leq 1\right) \stackrel{(i)}{\leq} \Pr\left(\left\|\pi_{i}^{\text{adv},0} - \pi_{i}^{*,1}\right\|^{2} \geq 1\right)$$

$$\stackrel{(ii)}{\leq} e^{\epsilon_{i}} \Pr\left(\left\|\pi_{i}^{\text{adv},1} - \pi_{i}^{*,1}\right\|^{2} \geq 1\right) + \delta_{i} \leq e^{\epsilon_{i}} \zeta_{i} + \delta_{i}.$$

(i) is because $\left\|\pi_i^{*,0} - \pi_i^{*,1}\right\|^2 = 2$. (ii) is because post-processing will not decrease the strength of DP.

Therefore, we have $\zeta_i \geq \frac{1-\delta_i}{1+e^{\epsilon_i}}$ for i>2. For $i\in\{1,2\}$, we have $1\geq \frac{1-\delta_i}{1+e^{\epsilon_i}}$ since $1-\delta_i\leq 1$.

Moreover, we have

$$\frac{1}{N}\left(2+\sum_{i=3}^{N}\zeta_i\right)\geq \frac{1}{N}\sum_{i=1}^{N}\frac{1-\delta_i}{1+e^{\epsilon_i}}\geq \frac{1}{N}\sum_{i=1}^{N}\frac{1-\max_{i\in[N]}\delta_i}{1+e^{\epsilon_i}}\stackrel{(i)}{\geq}\frac{1-\max_{i\in[N]}\delta_i}{1+\exp\left(\frac{1}{N}\sum_{i=1}^{N}\epsilon_i\right)}.$$

(i) uses Jensen's inequality.

B.2.2 LOWERBOUNDS ON (α, ϵ) -RÉNYI DP

By further post-processing on $\pi_i^{\mathrm{adv},k}$, we can get a random variable sampled from μ_i^k , which is the distribution over $\{0,1\}$ indicating whether $\left\|\pi_i^{\mathrm{adv},k}-\pi_i^{*,0}\right\|\leq 1$. Then, for any i>2, we have

$$\epsilon_i \overset{(i)}{\geq} \mathrm{KL}(\mu_i^0, \mu_i^1) \overset{(ii)}{\geq} (1 - \zeta_i) \log \frac{1 - \zeta_i}{\zeta_i} + \zeta_i \log \frac{\zeta_i}{1 - \zeta_i} = (1 - 2\zeta_i) \log \frac{1 - \zeta_i}{\zeta_i}.$$

(i) is because post-processing will not weaken Rényi DP (Mironov, 2017). (ii) is because

$$\begin{split} & \mu_i^0(0) = \Pr\left(\left\| \pi_i^{\text{adv},0} - \pi_i^{*,0} \right\| > 1 \right) \leq \zeta_i \\ & \mu_i^1(1) = \Pr\left(\left\| \pi_i^{\text{adv},1} - \pi_i^{*,0} \right\| \leq 1 \right) \leq \Pr\left(\left\| \pi_i^{\text{adv},1} - \pi_i^{*,1} \right\| \geq 1 \right) \leq \zeta_i. \end{split}$$

Therefore, the KL-divergence is lowerbounded by $(1 - \zeta_i) \log \frac{1 - \zeta_i}{\zeta_i} + \zeta_i \log \frac{\zeta_i}{1 - \zeta_i}$.

When $\zeta_i \leq \frac{1}{4}$, we have $\epsilon_i \geq \frac{1}{2}\log\frac{3}{4\zeta_i}$. Therefore, $\zeta_i \geq \frac{3}{4e^{2\epsilon_i}}$. Moreover, if $\frac{1}{N}\sum_{i=1}^N \zeta_i \leq \frac{1}{16}$, by pigeon-hole principle, there exists a set of players $\mathcal{I} \subseteq [N] \setminus \{1,2\}$ with $|\mathcal{I}| \geq \frac{N}{2}$, so that for any $i \in \mathcal{I}$, $\zeta_i \leq \frac{1}{4}$. This further implies that

$$\frac{1}{N} \sum_{i=1}^{N} \zeta_i \ge \frac{1}{2|\mathcal{I}|} \sum_{i \in \mathcal{I}} \zeta_i \qquad \qquad \frac{1}{N} \sum_{i=1}^{N} \epsilon_i \ge \frac{1}{2|\mathcal{I}|} \sum_{i \in \mathcal{I}} \epsilon_i,$$

since ζ_i , $\epsilon_i \geq 0$ for any $i \in [N]$.

Therefore, finally,

$$\frac{1}{N} \sum_{i=1}^{N} \zeta_i \ge \frac{1}{2|\mathcal{I}|} \sum_{i \in \mathcal{I}} \zeta_i \ge \frac{3}{8|\mathcal{I}|} \sum_{i \in \mathcal{I}} e^{-2\epsilon_i} \stackrel{(i)}{\ge} \frac{3}{8} \exp\left(-\frac{2}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \epsilon_i\right) \ge \frac{3}{8} \exp\left(-\frac{4}{N} \sum_{i=1}^{N} \epsilon_i\right).$$

(i) uses Jensen's inequality.

B.3 OMITTED PROOF OF LEMMAS

Lemma B.1. Consider the 2^N games $\{\mathcal{G}^S\}_{S\subseteq[N]}$ constructed above. There must exist two adjacent games \mathcal{G}^S and $\mathcal{G}^{S\cup\{i\}}$ with $i\in[N]\setminus S$, so that

$$\mathbb{E}\left[\max_{\widehat{\pi}_{3i-1}\in\Delta^{\mathcal{A}_{3i-1}}}\left\langle \pi_{3i-1}^{\mathcal{S}} - \widehat{\pi}_{3i-1}, \boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\mathcal{S}}}(\mathcal{G}^{\mathcal{S}})\right\rangle + \max_{\widehat{\pi}_{3i}\in\Delta^{\mathcal{A}_{3i}}}\left\langle \pi_{3i}^{\mathcal{S}} - \widehat{\pi}_{3i}, \boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\mathcal{S}}}(\mathcal{G}^{\mathcal{S}})\right\rangle\right] \leq 7\zeta \qquad (18)$$

$$\mathbb{E}\left[\max_{\widehat{\pi}_{3i-1}\in\Delta^{\mathcal{A}_{3i-1}}}\left\langle \pi_{3i-1}^{\mathcal{S}\cup\{i\}} - \widehat{\pi}_{3i-1}, \boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\mathcal{S}\cup\{i\}}}(\mathcal{G}^{\mathcal{S}\cup\{i\}})\right\rangle + \max_{\widehat{\pi}_{3i}\in\Delta^{\mathcal{A}_{3i}}}\left\langle \pi_{3i}^{\mathcal{S}\cup\{i\}} - \widehat{\pi}_{3i}, \boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\mathcal{S}\cup\{i\}}}(\mathcal{G}^{\mathcal{S}\cup\{i\}})\right\rangle\right] \leq 7\zeta.$$

$$(19)$$

Proof. By the accuracy condition and non-negativity of $\max_{\widehat{\pi}_j \in \Delta^{A_j}} \left\langle \pi_j^{\emptyset} - \widehat{\pi}_j, \boldsymbol{g}_j^{\boldsymbol{\pi}^{\emptyset}}(\mathcal{G}^{\emptyset}) \right\rangle$, we have

$$\sum_{i \in [N]} \mathbb{E} \left[\max_{\widehat{\pi}_{3i-1} \in \Delta^{\mathcal{A}_{3i-1}}} \left\langle \pi_{3i-1}^{\emptyset} - \widehat{\pi}_{3i-1}, \boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\emptyset}}(\mathcal{G}^{\emptyset}) \right\rangle + \max_{\widehat{\pi}_{3i} \in \Delta^{\mathcal{A}_{3i}}} \left\langle \pi_{3i}^{\emptyset} - \widehat{\pi}_{3i}, \boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\emptyset}}(\mathcal{G}^{\emptyset}) \right\rangle \right] \leq 3N\zeta.$$

Therefore, by the pigeon-hole principle, there must exist a subset $\mathcal{I}\subseteq [N]$ with $|\mathcal{I}|\geq \lceil\frac{N}{2}\rceil+1$, so that for any $i\in\mathcal{I}$, $\mathbb{E}\left[\max_{\widehat{\pi}_{3i-1}\in\Delta^{\mathcal{A}_{3i-1}}}\left\langle\pi^\emptyset_{3i-1}-\widehat{\pi}_{3i-1},\boldsymbol{g}^{\pi^\emptyset}_{3i-1}(\mathcal{G}^\emptyset)\right\rangle+\max_{\widehat{\pi}_{3i}\in\Delta^{\mathcal{A}_{3i}}}\left\langle\pi^\emptyset_{3i}-\widehat{\pi}_{3i},\boldsymbol{g}^{\pi^\emptyset}_{3i}(\mathcal{G}^\emptyset)\right\rangle\right]\leq 7\zeta.$ The same guarantee holds for any $\mathcal{G}^{\{i\}}$ with $i\in[N]$.

Then, we will form a meta graph, with each node as a subset of [N] so that there are 2^N nodes in total. For each two nodes $\mathcal{S}^1, \mathcal{S}^2$ with $|\mathcal{S}^1| \leq |\mathcal{S}^2|$, they are connected in the meta graph if and only if $\mathcal{G}^{\mathcal{S}^1}$ and $\mathcal{G}^{\mathcal{S}^2}$ are adjacent. In other words, $\mathcal{S}^2 = \mathcal{S}^1 \cup \{i\}$ and $i \in [N] \setminus \mathcal{S}^1$. The edge between $(\mathcal{S}^1, \mathcal{S}^2)$ is labeled as $i \in [N]$ if $\mathcal{S}^2 \setminus \mathcal{S}^1 = \{i\}$. Therefore, each node has N edges and their labels are different from each other. Then, there are 2^N nodes in the meta graph and $N2^{N-1}$ edges.

For each node $S \subseteq [N]$, the set \mathcal{I} constructed above can be viewed as selecting edges with labels in \mathcal{I} . Then, for each node S and its selected edge with label i, it is guaranteed that

$$\mathbb{E}\left[\max_{\widehat{\pi}_{3i-1} \in \Delta^{\mathcal{A}_{3i-1}}} \left\langle \pi_{3i-1}^{\mathcal{S}} - \widehat{\pi}_{3i-1}, \boldsymbol{g}_{3i-1}^{\boldsymbol{\pi}^{\mathcal{S}}}(\mathcal{G}^{\mathcal{S}}) \right\rangle + \max_{\widehat{\pi}_{3i} \in \Delta^{\mathcal{A}_{3i}}} \left\langle \pi_{3i}^{\mathcal{S}} - \widehat{\pi}_{3i}, \boldsymbol{g}_{3i}^{\boldsymbol{\pi}^{\mathcal{S}}}(\mathcal{G}^{\mathcal{S}}) \right\rangle \right] \leq 7\zeta. \quad (21)$$

Because $|\mathcal{I}| \geq \left\lceil \frac{N}{2} \right\rceil + 1$ and there are $N2^{N-1}$ edges in total, by pigeon-hole principle, there must exist an edge being selected twice. In other words, there exist two adjacent nodes $\mathcal{S}^1, \mathcal{S}^2$ that both select edge with label i. By definition, assume $|\mathcal{S}^1| \leq |\mathcal{S}^2|$ without loss of generality, then $\mathcal{S}^2 = \mathcal{S}^1 \cup \{i\}$. Therefore, (21) is satisfied for both $\mathcal{S}^1, \mathcal{S}^2$ on i.

C PROOF OF THEOREM 5.1

Theorem 5.1. Consider Algorithm 1. Let $A = \max_{i \in [N]} |\mathcal{A}_i|$. The update-rule can achieve the following guarantee in any polymatrix game. For any T > 0, player $i \in [N]$ and strategy $\pi_i \in \Delta^{\mathcal{A}_i}$,

$$\frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbb{E}\left[\left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle\right] \\
\leq \frac{1}{\eta T} + A \frac{\sigma^{2}}{2\eta} + \left(\frac{2\eta^{2}}{\sigma} + \frac{7\sigma}{2}\right) A^{\frac{3}{2}} + \frac{1}{2\left(\overline{\mathcal{N}}\right)^{4/9} \log N} + \frac{2\eta\sqrt{A}}{\sigma\left(\overline{\mathcal{N}}\right)^{4/9} \log N}.$$
(10)

We will give a proof sketch of Theorem 5.1 first.

C.1 PROOF SKETCH

To show the convergence of Algorithm 1, by Lemma D.4. in Liu et al. (2024), we have

$$\frac{\tau_{i}}{2} \left\| \pi_{i}^{(t+1)} \right\|^{2} - \frac{\tau_{i}}{2} \left\| \pi_{i} \right\|^{2} + \left\langle \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \\
\leq \frac{1}{2\eta} \left\| \pi_{i} - \bar{\pi}_{i}^{(t)} \right\|^{2} - \frac{1 + \eta \tau_{i}}{2\eta} \left\| \pi_{i} - \pi_{i}^{(t+1)} \right\|^{2} - \frac{1}{2\eta} \left\| \pi_{i}^{(t+1)} - \bar{\pi}_{i}^{(t)} \right\|^{2}.$$

Moreover, $\left|\frac{\tau_i}{2} \left\| \pi_i^{(t+1)} \right\|^2 - \frac{\tau_i}{2} \left\| \pi_i \right\|^2 \right| \leq \frac{\tau_i}{2}$. Then, by adding $\left\langle \boldsymbol{g}_i^{\boldsymbol{\pi}^{(t+1)}} - \bar{\boldsymbol{g}}_i^{(t)}, \pi_i^{(t+1)} - \pi_i \right\rangle$ on both sides, we have

$$\begin{split} & \left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}}, \boldsymbol{\pi}_{i}^{(t+1)} - \boldsymbol{\pi}_{i} \right\rangle - \frac{\tau_{i}}{2} \\ \leq & \frac{1}{2\eta} \left\| \boldsymbol{\pi}_{i} - \bar{\boldsymbol{\pi}}_{i}^{(t)} \right\|^{2} - \frac{1 + \eta \tau_{i}}{2\eta} \left\| \boldsymbol{\pi}_{i} - \boldsymbol{\pi}_{i}^{(t+1)} \right\|^{2} \\ & - \frac{1}{2\eta} \left\| \boldsymbol{\pi}_{i}^{(t+1)} - \bar{\boldsymbol{\pi}}_{i}^{(t)} \right\|^{2} + \left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}} - \bar{\boldsymbol{g}}_{i}^{(t)}, \boldsymbol{\pi}_{i}^{(t+1)} - \boldsymbol{\pi}_{i} \right\rangle. \end{split}$$

Moreover, $\mathbb{E}\left[\left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}} - \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i}\right\rangle\right] \leq \left(\frac{\eta^{2}}{\sigma} + \frac{7\sigma}{2}\right)A^{\frac{3}{2}} + \frac{2\eta\sqrt{A}}{|\mathcal{N}(i)|\sigma}\sum_{j\in\mathcal{N}(i)}\tau_{j}$, where the complete proof is postponed to Lemma C.1. Therefore, by taking expectation on both sides and telescoping, we have

$$\begin{split} &\sum_{t=1}^{T} \mathbb{E}\left[\left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}}, \boldsymbol{\pi}_{i}^{(t+1)} - \boldsymbol{\pi}_{i} \right\rangle\right] \\ \leq &\frac{1}{2\eta} \mathbb{E}\left[\left\|\boldsymbol{\pi}_{i} - \overline{\boldsymbol{\pi}}_{i}^{(1)}\right\|^{2}\right] + \left(\frac{\eta^{2}}{\sigma} + \frac{7\sigma}{2}\right) A^{\frac{3}{2}} T + \sum_{t=2}^{T} \mathbb{E}\left[\left\|\boldsymbol{\pi}_{i} - \overline{\boldsymbol{\pi}}_{i}^{(t)}\right\|^{2} - \left\|\boldsymbol{\pi}_{i} - \boldsymbol{\pi}_{i}^{(t)}\right\|^{2}\right] \\ &+ \frac{2\eta\sqrt{A}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_{j}. \end{split}$$

Further,
$$\mathbb{E}\left[\left\|\boldsymbol{\pi}_{i}-\boldsymbol{\bar{\pi}}_{i}^{(t)}\right\|^{2}-\left\|\boldsymbol{\pi}_{i}-\boldsymbol{\pi}_{i}^{(t)}\right\|^{2}\right] \text{ is bounded by,}$$

$$\mathbb{E}\left[\left\|\boldsymbol{\pi}_{i}-\operatorname{Proj}_{\Delta^{A_{i}}}\left(\boldsymbol{\pi}_{i}^{(t)}+\boldsymbol{n}_{i}^{(t)}\right)\right\|^{2}-\left\|\boldsymbol{\pi}_{i}-\boldsymbol{\pi}_{i}^{(t)}\right\|^{2}\right]$$

$$\leq \mathbb{E}\left[\left\|\boldsymbol{\pi}_{i}-\boldsymbol{\pi}_{i}^{(t)}-\boldsymbol{n}_{i}^{(t)}\right\|^{2}-\left\|\boldsymbol{\pi}_{i}-\boldsymbol{\pi}_{i}^{(t)}\right\|^{2}\right]$$

$$=\mathbb{E}\left[2\left\langle\boldsymbol{n}_{i}^{(t)},\boldsymbol{\pi}_{i}-\boldsymbol{\pi}_{i}^{(t)}\right\rangle+\left\|\boldsymbol{n}_{i}^{(t)}\right\|^{2}\right]$$

$$\stackrel{(i)}{=}\mathbb{E}\left[\left\|\boldsymbol{n}_{i}^{(t)}\right\|^{2}\right]=A\sigma^{2}.$$

(i) is because $\mathbb{E}\left[\boldsymbol{n}_{i}^{(t)}\right] = \mathbf{0}$ and $\boldsymbol{n}_{i}^{(t)}$ is independent of $\pi_{i} - \pi_{i}^{(t)}$. By aggregating the results above and substitute τ_{i} by its value in Algorithm 1, the proof is concluded.

C.2 FORMAL PROOF OF THEOREM 5.1

By Lemma D.4. in Liu et al. (2024), since Algorithm 1 is equivalent to $\pi_i^{(t+1)} = \operatorname{argmin}_{\pi_i \in \Delta^{\mathcal{A}_i}} \left\langle \pi_i, \bar{\boldsymbol{g}}_i^{(t)} \right\rangle + \frac{\tau_i}{2} \left\| \pi_i \right\|^2 + \frac{1}{2\eta} \left\| \pi_i - \bar{\pi}_i^{(t)} \right\|^2$, for any $\pi_i \in \Delta^{\mathcal{A}_i}$, we have

$$\frac{\eta \tau_{i}}{2} \left\| \pi_{i}^{(t+1)} \right\|^{2} - \frac{\eta \tau_{i}}{2} \left\| \pi_{i} \right\|^{2} + \eta \left\langle \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \\
\leq \frac{1}{2} \left\| \pi_{i} - \bar{\pi}_{i}^{(t)} \right\|^{2} - \frac{1 + \eta \tau_{i}}{2} \left\| \pi_{i} - \pi_{i}^{(t+1)} \right\|^{2} - \frac{1}{2} \left\| \pi_{i}^{(t+1)} - \bar{\pi}_{i}^{(t)} \right\|^{2}.$$

Then, by adding $\eta\left\langle ar{m{g}}_i^{(t+1)} - ar{m{g}}_i^{(t)}, \pi_i^{(t+1)} - \pi_i \right\rangle$ on both sides, we have

$$\frac{\eta \tau_{i}}{2} \left\| \pi_{i}^{(t+1)} \right\|^{2} - \frac{\eta \tau_{i}}{2} \left\| \pi_{i} \right\|^{2} + \eta \left\langle \bar{\boldsymbol{g}}_{i}^{(t+1)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \\
\leq \frac{1}{2} \left\| \pi_{i} - \bar{\pi}_{i}^{(t)} \right\|^{2} - \frac{1 + \eta \tau_{i}}{2} \left\| \pi_{i} - \pi_{i}^{(t+1)} \right\|^{2} - \frac{1}{2} \left\| \pi_{i}^{(t+1)} - \bar{\pi}_{i}^{(t)} \right\|^{2} + \eta \left\langle \bar{\boldsymbol{g}}_{i}^{(t+1)} - \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle.$$
(22)

Moreover,

$$\begin{split} \left\langle \bar{\boldsymbol{g}}_{i}^{(t+1)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle &= \left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle + \left\langle \bar{\boldsymbol{g}}_{i}^{(t+1)} - \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \\ &\geq \left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle - \left\| \bar{\boldsymbol{g}}_{i}^{(t+1)} - \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}} \right\|_{\infty} \cdot \left\| \pi_{i}^{(t+1)} - \pi_{i} \right\|_{1} \\ &\geq \left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle - 2 \left\| \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \boldsymbol{U}_{i,j} \left(\pi_{j}^{(t+1)} - \bar{\pi}_{j}^{(t+1)} \right) \right\|_{\infty} \\ &\stackrel{(i)}{\geq} \left\langle \boldsymbol{g}_{i}^{\boldsymbol{\pi}^{(t+1)}}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle - \frac{2}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \left\| \pi_{j}^{(t+1)} - \bar{\pi}_{j}^{(t+1)} \right\|_{1}. \end{split}$$

(i) is because $U_{i,j} \in [-1,1]^{A_i \times A_j}$. Recall that $A \coloneqq \max_{i \in [N]} |\mathcal{A}_i|$, we have

$$\left\| \boldsymbol{\pi}_{j}^{(t+1)} - \overline{\boldsymbol{\pi}}_{j}^{(t+1)} \right\|_{1} \leq \sqrt{A} \left\| \boldsymbol{\pi}_{j}^{(t+1)} - \overline{\boldsymbol{\pi}}_{j}^{(t+1)} \right\| = \sqrt{A} \left\| \boldsymbol{\pi}_{j}^{(t+1)} - \operatorname{Proj}_{\Delta^{\mathcal{A}_{i}}} \left(\boldsymbol{\pi}_{j}^{(t+1)} + \boldsymbol{n}_{j}^{(t+1)} \right) \right\| \leq \sqrt{A} \left\| \boldsymbol{n}_{j}^{(t+1)} \right\|.$$

By taking the expectation on both sides of (22), we have,

$$\mathbb{E}\left[\left(\frac{\eta\tau_{i}}{2} \left\|\pi_{i}^{(t+1)}\right\|^{2} - \frac{\eta\tau_{i}}{2} \left\|\pi_{i}\right\|^{2} + \eta \left\langle-\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \boldsymbol{U}_{i,j} \pi_{j}^{(t+1)}, \pi_{i}^{(t+1)} - \pi_{i}\right\rangle\right)\right]$$

$$\leq \mathbb{E}\left[\frac{1}{2} \left\|\pi_{i} - \bar{\pi}_{i}^{(t)}\right\|^{2} - \frac{1 + \eta\tau_{i}}{2} \left\|\pi_{i} - \pi_{i}^{(t+1)}\right\|^{2} - \frac{1}{2} \left\|\pi_{i}^{(t+1)} - \bar{\pi}_{i}^{(t)}\right\|^{2}\right]$$

$$+ \eta \mathbb{E}\left[\left\langle\bar{\boldsymbol{g}}_{i}^{(t+1)} - \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i}\right\rangle + \frac{2\sqrt{A}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \left\|\boldsymbol{n}_{j}^{(t+1)}\right\|^{2}\right].$$
(23)

By Jensen's Inequality,
$$\mathbb{E}\left[\left\|\boldsymbol{n}_{j}^{(t+1)}\right\|\right] \leq \sqrt{\mathbb{E}\left[\left\|\boldsymbol{n}_{j}^{(t+1)}\right\|^{2}\right]} = \sqrt{|\mathcal{A}_{j}|}\sigma$$
. Therefore,
$$\frac{2\sqrt{A}}{|\mathcal{N}(i)|}\sum_{j\in\mathcal{N}(i)}\mathbb{E}\left[\left\|\boldsymbol{n}_{j}^{(t+1)}\right\|\right] \leq 2A\sigma.$$

Moreover, since $\|\pi_i\|^2 \in \left[\frac{1}{|\mathcal{A}_i|}, 1\right]$, we have

1028
1029
$$\mathbb{E}\left[\eta\left\langle -\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} U_{i,j} \pi_j^{(t+1)}, \pi_i^{(t+1)} - \pi_i \right\rangle\right]$$
1031
1032
$$\leq \mathbb{E}\left[\frac{1}{2} \left\|\pi_i - \overline{\pi}_i^{(t)}\right\|^2 - \frac{1 + \eta \tau_i}{2} \left\|\pi_i - \pi_i^{(t+1)}\right\|^2 - \frac{1}{2} \left\|\pi_i^{(t+1)} - \overline{\pi}_i^{(t)}\right\|^2 + \eta \left\langle \overline{\boldsymbol{g}}_i^{(t+1)} - \overline{\boldsymbol{g}}_i^{(t)}, \pi_i^{(t+1)} - \pi_i \right\rangle\right]$$
1033
1034
$$+ 2A\eta\sigma + \frac{\eta}{2}\tau_i.$$
1035

To bound $\mathbb{E}\left[\left\langle ar{m{g}}_i^{(t+1)} - ar{m{g}}_i^{(t)}, \pi_i^{(t+1)} - \pi_i
ight
angle \right]$, we have the following lemma.

Lemma C.1. Consider Algorithm 1. For any player $i \in [N]$, we have

$$\mathbb{E}\left[\left\langle \bar{\boldsymbol{g}}_{i}^{(t+1)} - \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \right] \leq \left(\frac{2\eta^{2}}{\sigma} + \frac{3\sigma}{2}\right) A^{\frac{3}{2}} + \frac{2\eta\sqrt{A}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_{j}. \tag{24}$$

The proof is postponed to the end of this section

We can further bound
$$\mathbb{E}\left[\left\|\pi_i-\overline{\pi}_i^{(t)}\right\|^2-\left\|\pi_i-\pi_i^{(t)}\right\|^2
ight]$$
 by

$$\mathbb{E}\left[\left\|\boldsymbol{\pi}_{i}-\operatorname{Proj}_{\Delta^{\mathcal{A}_{i}}}\left(\boldsymbol{\pi}_{i}^{(t)}+\boldsymbol{n}_{i}^{(t)}\right)\right\|^{2}-\left\|\boldsymbol{\pi}_{i}-\boldsymbol{\pi}_{i}^{(t)}\right\|^{2}\right] \leq \mathbb{E}\left[\left\|\boldsymbol{\pi}_{i}-\boldsymbol{\pi}_{i}^{(t)}-\boldsymbol{n}_{i}^{(t)}\right\|^{2}-\left\|\boldsymbol{\pi}_{i}-\boldsymbol{\pi}_{i}^{(t)}\right\|^{2}\right]$$

$$=\mathbb{E}\left[2\left\langle\boldsymbol{n}_{i}^{(t)},\boldsymbol{\pi}_{i}^{(t)}-\boldsymbol{\pi}_{i}\right\rangle+\left\|\boldsymbol{n}_{i}^{(t)}\right\|^{2}\right]$$

$$\stackrel{(i)}{=}\mathbb{E}\left[\left\|\boldsymbol{n}_{i}^{(t)}\right\|^{2}\right]=A\sigma^{2}.$$

(i) is because $\mathbb{E}\left[m{n}_i^{(t)}
ight]=m{0}$ and $m{n}_i^{(t)}$ is independent of $\pi_i-\pi_i^{(t)}$.

Finally, by telescoping,

$$\sum_{t=1}^{T} \mathbb{E} \left[\left\langle -\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} U_{i,j} \pi_{j}^{(t+1)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \right] \\
\leq \frac{1}{2\eta} \left\| \pi_{1} - \overline{\pi}_{1}^{(1)} \right\|^{2} + A \frac{\sigma^{2}}{2\eta} T + \left(\frac{2\eta^{2}}{\sigma} + \frac{3\sigma}{2} \right) A^{\frac{3}{2}} T + 2A\sigma T + \frac{1}{2} \tau_{i} T + \frac{2\eta\sqrt{A}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_{j} T \\
\leq \frac{1}{\eta} + A \frac{\sigma^{2}}{2\eta} T + \left(\frac{2\eta^{2}}{\sigma} + \frac{7\sigma}{2} \right) A^{\frac{3}{2}} T + \frac{1}{2} \tau_{i} T + \frac{2\eta\sqrt{A}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_{j} T. \tag{25}$$

Further, by taking summation over all player $i \in [N]$ and averaging, we have

$$\frac{1}{N} \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbb{E} \left[\left\langle -\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} U_{i,j} \pi_{j}^{(t+1)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \right] \\
\leq \frac{1}{\eta} + A \frac{\sigma^{2}}{2\eta} T + \left(\frac{2\eta^{2}}{\sigma} + \frac{7\sigma}{2} \right) A^{\frac{3}{2}} T + \frac{T}{2N} \sum_{i=1}^{N} \tau_{i} + \frac{1}{N} \sum_{i=1}^{N} \frac{2\eta \sqrt{A}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_{j} T.$$

Then, we can further bound the summation over τ_i by Lemma C.2.

Lemma C.2. When $\tau_i = \frac{c}{|\mathcal{N}(i)|}$ for any $i \in [N]$, where c > 0 is a game-dependent constant, we have

$$\frac{1}{N} \sum_{i=1}^{N} \tau_i = \frac{c}{\overline{\mathcal{N}}} \qquad \frac{1}{N} \sum_{i=1}^{N} \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_j \le \frac{c}{\overline{\mathcal{N}}}.$$
 (26)

The proof is postponed to the end of this section. Therefore, since $c = \frac{(\overline{\mathcal{N}})^{5/9}}{\log N}$ in Algorithm 1, we have

$$\frac{1}{N} \sum_{i=1}^{N} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E} \left[\left\langle -\frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \boldsymbol{U}_{i,j} \pi_{j}^{(t+1)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \right] \\
\leq \frac{1}{\eta T} + A \frac{\sigma^{2}}{2\eta} + \left(\frac{2\eta^{2}}{\sigma} + \frac{7\sigma}{2} \right) A^{\frac{3}{2}} + \frac{1}{2 \left(\overline{\mathcal{N}} \right)^{4/9} \log N} + \frac{2\eta \sqrt{A}}{\sigma \left(\overline{\mathcal{N}} \right)^{4/9} \log N}. \qquad \square$$

C.3 OMITTED PROOF OF APPENDIX C

Lemma C.1. Consider Algorithm 1. For any player $i \in [N]$, we have

$$\mathbb{E}\left[\left\langle \bar{\boldsymbol{g}}_{i}^{(t+1)} - \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \right] \leq \left(\frac{2\eta^{2}}{\sigma} + \frac{3\sigma}{2}\right) A^{\frac{3}{2}} + \frac{2\eta\sqrt{A}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_{j}. \tag{24}$$

Proof. By Hölder's Inequality, we have

$$\mathbb{E}\left[\left\langle \bar{\boldsymbol{g}}_{i}^{(t+1)} - \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle\right] \leq \mathbb{E}\left[\left\|\bar{\boldsymbol{g}}_{i}^{(t+1)} - \bar{\boldsymbol{g}}_{i}^{(t)}\right\| \cdot \left\|\pi_{i}^{(t+1)} - \pi_{i}\right\|\right].$$

Moreover,

$$\begin{aligned} \left\| \bar{\boldsymbol{g}}_{i}^{(t+1)} - \bar{\boldsymbol{g}}_{i}^{(t)} \right\| &= \frac{1}{|\mathcal{N}(i)|} \left\| \sum_{j \in \mathcal{N}(i)} \boldsymbol{U}_{i,j} \left(\operatorname{Proj}_{\Delta^{\mathcal{A}_{j}}} \left(\pi_{j}^{(t+1)} + \boldsymbol{n}_{j}^{(t+1)} \right) - \bar{\pi}_{j}^{(t)} \right) \right\| \\ &\leq \frac{\sqrt{|\mathcal{A}_{i}|}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \left\| \operatorname{Proj}_{\Delta^{\mathcal{A}_{j}}} \left(\pi_{j}^{(t+1)} + \boldsymbol{n}_{j}^{(t+1)} \right) - \bar{\pi}_{j}^{(t)} \right\|_{1} \\ &\leq \frac{\sqrt{|\mathcal{A}_{i}|}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \sqrt{|\mathcal{A}_{j}|} \left\| \operatorname{Proj}_{\Delta^{\mathcal{A}_{j}}} \left(\pi_{j}^{(t+1)} + \boldsymbol{n}_{j}^{(t+1)} \right) - \bar{\pi}_{j}^{(t)} \right\| \\ &\leq \frac{\sqrt{|\mathcal{A}_{i}|}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \sqrt{|\mathcal{A}_{j}|} \left\| \pi_{j}^{(t+1)} + \boldsymbol{n}_{j}^{(t+1)} - \bar{\pi}_{j}^{(t)} \right\|. \end{aligned}$$

Therefore,

$$\begin{split} & \mathbb{E}\left[\left\langle \bar{\boldsymbol{g}}_{i}^{(t+1)} - \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle\right] \\ \leq & \frac{A}{\left|\mathcal{N}(i)\right|} \sum_{j \in \mathcal{N}(i)} \mathbb{E}\left[\left\|\boldsymbol{\pi}_{j}^{(t+1)} + \boldsymbol{n}_{j}^{(t+1)} - \bar{\boldsymbol{\pi}}_{j}^{(t)}\right\| \cdot \left\|\boldsymbol{\pi}_{i}^{(t+1)} - \pi_{i}\right\|\right] \\ \leq & \frac{A}{\left|\mathcal{N}(i)\right|} \sum_{j \in \mathcal{N}(i)} \left(\frac{1}{\sigma\sqrt{A}} \mathbb{E}\left[\left\|\boldsymbol{\pi}_{j}^{(t+1)} + \boldsymbol{n}_{j}^{(t+1)} - \bar{\boldsymbol{\pi}}_{j}^{(t)}\right\|^{2}\right] + \frac{\sigma\sqrt{A}}{4} \mathbb{E}\left[\left\|\boldsymbol{\pi}_{i}^{(t+1)} - \boldsymbol{\pi}_{i}\right\|^{2}\right]\right). \end{split}$$

Furthermore, since $\mathbb{E}\left[\boldsymbol{n}_{j}^{(t+1)}\right]=0$ and $\boldsymbol{n}_{j}^{(t+1)}$ is independent of $\pi_{j}^{(t+1)}, \overline{\pi}_{j}^{(t)}$, we have

$$\mathbb{E}\left[\left\|\boldsymbol{\pi}_{j}^{(t+1)} + \boldsymbol{n}_{j}^{(t+1)} - \overline{\boldsymbol{\pi}}_{j}^{(t)}\right\|^{2}\right] = \mathbb{E}\left[\left\|\boldsymbol{\pi}_{j}^{(t+1)} - \overline{\boldsymbol{\pi}}_{j}^{(t)}\right\|^{2}\right] + \mathbb{E}\left[\left\|\boldsymbol{n}_{j}^{(t+1)}\right\|^{2}\right]$$

$$\leq \mathbb{E}\left[\left\|\operatorname{Proj}_{\Delta^{A_{j}}}\left(\frac{\overline{\boldsymbol{\pi}}_{j}^{(t)} - \eta \boldsymbol{g}_{j}^{(t)}}{1 + \eta \tau_{j}}\right) - \overline{\boldsymbol{\pi}}_{j}^{(t)}\right\|^{2}\right] + A\sigma^{2}$$

$$\leq \mathbb{E}\left[\left\|\frac{\overline{\boldsymbol{\pi}}_{j}^{(t)} - \eta \boldsymbol{g}_{j}^{(t)}}{1 + \eta \tau_{j}} - \overline{\boldsymbol{\pi}}_{j}^{(t)}\right\|^{2}\right] + A\sigma^{2}$$

$$= \mathbb{E}\left[\left\|\frac{\eta \tau_{j}}{1 + \eta \tau_{j}} \overline{\boldsymbol{\pi}}_{j}^{(t)} - \frac{\eta}{1 + \eta \tau_{j}} \boldsymbol{g}_{j}^{(t)}\right\|^{2}\right] + A\sigma^{2}$$

$$\stackrel{(i)}{\leq} 2\eta \tau_{j} \mathbb{E}\left[\left\|\overline{\boldsymbol{\pi}}_{j}^{(t)}\right\|^{2}\right] + 2\eta^{2} \mathbb{E}\left[\left\|\boldsymbol{g}_{j}^{(t)}\right\|^{2}\right] + A\sigma^{2}$$

$$\stackrel{(ii)}{\leq} 2\eta \tau_{j} + 2\eta^{2} A + A\sigma^{2}.$$

 $\begin{array}{l} (i) \text{ is because } (a+b)^2 \leq 2a^2 + 2b^2 \text{ for any } a,b \in \mathbb{R} \text{ and } 1 + \eta \tau_j \geq \max{\{1,\eta \tau_j\}}. \ (ii) \text{ is because } \\ U_{i,j} \in [-1,1]^{\mathcal{A}_i \times \mathcal{A}_j}. \text{ Also, } \frac{\sigma \sqrt{A}}{4} \mathbb{E} \left[\left\| \pi_i^{(t+1)} - \pi_i \right\|^2 \right] \leq \frac{\sigma \sqrt{A}}{2}. \text{ Therefore,} \end{array}$

$$\mathbb{E}\left[\left\langle \bar{\boldsymbol{g}}_{i}^{(t+1)} - \bar{\boldsymbol{g}}_{i}^{(t)}, \pi_{i}^{(t+1)} - \pi_{i} \right\rangle \right] \leq \left(\frac{2\eta^{2}}{\sigma} + \frac{3\sigma}{2}\right) A^{\frac{3}{2}} + \frac{2\eta\sqrt{A}}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_{j}. \qquad \Box$$

Lemma C.2. When $\tau_i = \frac{c}{|\mathcal{N}(i)|}$ for any $i \in [N]$, where c > 0 is a game-dependent constant, we have

$$\frac{1}{N} \sum_{i=1}^{N} \tau_i = \frac{c}{\overline{\mathcal{N}}} \qquad \frac{1}{N} \sum_{i=1}^{N} \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_j \le \frac{c}{\overline{\mathcal{N}}}.$$
 (26)

Proof. By definition,

$$\frac{1}{N}\sum_{i=1}^{N}\tau_{i} = \frac{c}{N}\sum_{i=1}^{N}\frac{1}{|\mathcal{N}(i)|} = \frac{c}{\overline{\mathcal{N}}}.$$

Moreover.

$$\frac{1}{N} \sum_{i=1}^{N} \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \tau_j = \frac{c}{N} \sum_{i=1}^{N} \sum_{j \in \mathcal{N}(i)} \frac{1}{|\mathcal{N}(i)| \cdot |\mathcal{N}(j)|}$$

$$\leq \frac{c}{N} \sum_{i=1}^{N} \sum_{j \in \mathcal{N}(i)} \left(\frac{1}{2|\mathcal{N}(i)|^2} + \frac{1}{2|\mathcal{N}(j)|^2} \right)$$

$$\stackrel{(i)}{=} \frac{c}{N} \sum_{i=1}^{N} \frac{2|\mathcal{N}(i)|}{2|\mathcal{N}(i)|^2}$$

$$= \frac{c}{N}.$$

(i) is because each player contributes $\frac{1}{2|\mathcal{N}(i)|^2}$ to the summation 2 times on each edge, i.e. $2|\mathcal{N}(i)|$ times in total.

D Proof of Section 6

Theorem 6.1. Consider Algorithm 1 and any two adjacent polymatrix games $\mathcal{G}, \mathcal{G}'$ that differs at (v_1, v_2) . Let dist(i, j) be the length of the shortest path between players i, j, which is ∞ when i, j

are not connected. The update-rule guarantees the following for any T > 0 and player $i \in [N]$,

$$\frac{1}{N} \sum_{i=1}^{N} D_{\alpha} \left(\mu_{\mathcal{G},i}^{(T)}, \mu_{\mathcal{G}',i}^{(T)} \right) \le \frac{\alpha \eta^{2}}{\sigma^{2}} \min \left\{ \clubsuit, \spadesuit \right\} T, \tag{11}$$

where

We will give a proof sketch of Theorem 6.1 first.

D.1 PROOF SKETCH: THE CASE OF DENSE GRAPHS (\$\.*)

Firstly, we will introduce the chain rule of Rényi divergence.

Lemma D.1 (Chain Rule of Rényi divergence). For any distribution p, q over random variables X^1, X^2 , for any $\alpha \ge 1$, we have

$$D_{\alpha}\left(p(X^{1}, X^{2}), q(X^{1}, X^{2})\right) \leq D_{\alpha}\left(p(X^{1}), q(X^{1})\right) + \sup_{\widehat{x}^{1}} D_{\alpha}\left(p(X^{2} \mid X^{1} = \widehat{x}^{1}), q(X^{2} \mid X^{1} = \widehat{x}^{1})\right).$$
(27)

A proof of the above lemma is provided in Appendix D.5 for completeness. By Lemma D.1, the divergence between the distribution over all observations can be decomposed into the divergence between the distribution over the observation at a single timestep. Formally, for any $\alpha \geq 1$, adjacent games $\mathcal{G} \sim \mathcal{G}'$, and player $i \in [N]$,

$$D_{\alpha}\left(\mu_{\mathcal{G},i}^{(T)}, \mu_{\mathcal{G}',i}^{(T)}\right) \\ \leq \sum_{t=1}^{T} \sup_{\left(\widehat{o}^{(1)}, \dots, \widehat{o}^{(t-1)}\right)} D_{\alpha}\left(\mu_{\mathcal{G},i}^{(t)}(\cdot | \widehat{o}^{(1)}, \dots, \widehat{o}^{(t-1)}), \mu_{\mathcal{G}',i}^{(t)}(\cdot | \widehat{o}^{(1)}, \dots, \widehat{o}^{(t-1)})\right).$$

Note that the observation here $\widehat{o}^{(s)} = \pi_i^{(s)} + \boldsymbol{n}_i^{(s)}$. When the graph is dense, *i.e.* $\overline{\mathcal{N}} = N^p$ for some p > 0, the degree $|\mathcal{N}(i)|$ of most players will close to $\overline{\mathcal{N}}$.

Then, since post-processing will not increase the privacy budget, we can augment the set of observations to bound $D_{\alpha}\left(\mu_{\mathcal{G},i}^{(T)},\mu_{\mathcal{G}',i}^{(T)}\right)$. Specifically, for two adjacent games $\mathcal{G}\sim\mathcal{G}'$ that differs at the utility matrices on edge (v_1,v_2) , the augmented observation will include $\left\{\pi_j^{(t)}+\boldsymbol{n}_j^{(t)}\right\}_{j\in\mathcal{S}}$, where $\mathcal{S}\supseteq[N]\setminus\{v_1,v_2\}$ and v_1,v_2 will be included in \mathcal{S} if their degrees are no less than $\overline{\mathcal{N}}^{2/9}$. Let $\mu_{\mathcal{G},\mathcal{S}}^{(t)}$ denotes the marginal distribution of $\mu_{\mathcal{G}}^{(t)}$ on $\left\{\pi_j^{(t)}+\boldsymbol{n}_j^{(t)}\right\}_{j\in\mathcal{S}}$. We can see that for any $i\in\mathcal{S}$,

$$D_{\alpha}\left(\mu_{\mathcal{G},i}^{(T)},\mu_{\mathcal{G}',i}^{(T)}\right) \leq D_{\alpha}\left(\mu_{\mathcal{G},\mathcal{S}}^{(T)},\mu_{\mathcal{G}',\mathcal{S}}^{(T)}\right).$$

For $i \notin \mathcal{S}$, we can show that $D_{\alpha}\left(\mu_{\mathcal{G},i}^{(T)},\mu_{\mathcal{G}',i}^{(T)}\right)$ is bounded by some constant. Therefore, since $|\mathcal{S}| \geq N-2$, $\frac{1}{N}\sum_{i=1}^{N}D_{\alpha}\left(\mu_{\mathcal{G},i}^{(T)},\mu_{\mathcal{G}',i}^{(T)}\right)$ is bounded when $D_{\alpha}\left(\mu_{\mathcal{G},\mathcal{S}}^{(T)},\mu_{\mathcal{G}',\mathcal{S}}^{(T)}\right)$ is sublinear in N.

Note that $\mu_{\mathcal{G},\mathcal{S}}^{(t)}(\cdot | \widehat{o}^{(1)}, \cdots, \widehat{o}^{(t-1)})$ is a multivariate Gaussian distribution with mean $\left(\pi_i^{(t)}\right)_{i \in \mathcal{S}}$ and variance $\sigma^2 \boldsymbol{I}^{\times_{i \in \mathcal{S}} \mathcal{A}_i}$. Therefore, by the Rényi divergence of multivariate Gaussian (Gil et al., 2013), the divergence is bounded by $\frac{\alpha}{2\sigma^2} \sum_{i \in \mathcal{S}} \left\| \pi_i^{(t)} - \pi_i'^{(t)} \right\|^2$.

If S = [N], then for any player $i \in [N] \setminus \{v_1, v_2\}$, given all past observations $\left(\left\{\pi_j^{(s)} + \boldsymbol{n}_j^{(s)}\right\}_{i \in \mathcal{S}}\right)^{t-1} \text{ are identical in } \mathcal{G}, \mathcal{G}', \ \pi_i^{(t)} = {\pi'}_i^{(t)}. \ \text{For } i \in \{v_1, v_2\}, \ \text{given } \overline{\pi}_i^{(t-1)} = {\pi'}_i^{(t)}.$ $\overline{\pi}'_i^{(t-1)} \text{ and only } \boldsymbol{U}_{v_1,v_2} \neq \boldsymbol{U}'_{v_1,v_2}, \left\| \boldsymbol{\pi}_i^{(t)} - \boldsymbol{\pi}'_i^{(t)} \right\| \leq \mathcal{O}\left(\frac{1}{|\mathcal{N}(i)|}\right) \leq \mathcal{O}\left(\frac{1}{|\mathcal{N}|^{2/9}}\right) \text{ by definition of } \mathcal{S}.$ \mathcal{O} hides constants invariant to the number of players N, e.g. A, the size of the largest action set.

If $v_1, v_2 \notin \mathcal{S}$, then for any $i \in \mathcal{S} \setminus (\mathcal{N}(v_1) \cup \mathcal{N}(v_2)), \pi_i^{(t)} = \pi_i^{(t)}$. For $i \in \mathcal{N}(v_1) \cup \mathcal{N}(v_2)$, we will further augment the space of observations so that the adversary can observe $\left(\left\{n_{v_1}^{(s)},n_{v_2}^{(s)}\right\}\right)_{s=1}^{t-1}$. Assume $i \in \mathcal{N}(v_1) \setminus \mathcal{N}(v_2)$ for ease of exposition. For any $j \in \mathcal{N}(i) \setminus \{v_1, v_2\}$, we have $\overline{\pi}_j^{(t-1)} = 0$ $\overline{\pi}'_{j}^{(t-1)}$. Therefore, $\left\|\pi_{i}^{(t)} - {\pi'}_{i}^{(t)}\right\| \leq \mathcal{O}\left(\left\|\overline{\pi}_{v_{1}}^{(t-1)} - \overline{\pi}'_{v_{1}}^{(t-1)}\right\|\right) \leq \mathcal{O}\left(\left\|\pi_{v_{1}}^{(t-1)} - {\pi'}_{v_{1}}^{(t-1)}\right\|\right)$. Moreover, due to the additional regularization imposed on each player, we have the following lemma.

Lemma D.2. Consider the Algorithm 1. For any player $i \in [N]$ and timestep t > 0, by updat- $\textit{ing the strategy} \ \left\{\pi_i^{(s)}\right\}_{s=0}^t, \left\{\pi_i'^{(s)}\right\}_{s=0}^t \ \textit{with the same noise} \ \left\{n_i^{(s)}\right\}_{s=0}^t \ \textit{but different gradients}$ $\left\{g_{i}^{(s)}\right\}_{s=0}^{t}, \left\{g_{i}^{\prime(s)}\right\}_{s=0}^{t}$ individually, we have

$$\left\| \pi_i^{(t)} - {\pi'}_i^{(t)} \right\| \le \frac{2\sqrt{\mathcal{A}_i}}{\tau_i}. \tag{28}$$

The proof is postponed to Appendix D.5. Therefore, finally, $\left\|\pi_i^{(t)} - \pi_i^{\prime(t)}\right\| \leq \mathcal{O}\left(\frac{1}{\tau_{n_i}}\right) \leq$ $\mathcal{O}\left(\frac{\log N}{(\overline{\mathcal{N}})^{1/3}}\right)$ because $|\mathcal{N}(v_1)| \leq \left(\overline{\mathcal{N}}\right)^{2/9}$. Moreover, since $|\mathcal{N}(v_1)| \leq \left(\overline{\mathcal{N}}\right)^{2/9}$,

$$\sum_{i \in \mathcal{S}} \left\| \pi_i^{(t)} - {\pi'}_i^{(t)} \right\|^2 = \sum_{i \in \mathcal{S} \cap (\mathcal{N}(v_1) \cup \mathcal{N}(v_2))} \left\| \pi_i^{(t)} - {\pi'}_i^{(t)} \right\|^2 \le \mathcal{O}\left(\frac{(\log N)^2}{(\overline{\mathcal{N}})^{4/9}}\right).$$

For cases of $v_1 \in \mathcal{S}, v_2 \notin \mathcal{S}$ and $v_1 \notin \mathcal{S}, v_2 \in \mathcal{S}$, the proof is similar.

PROOF SKETCH: THE CASE OF SPARSE GRAPHS (A) D.2

For the sparse graph, the degree of all players is bounded by a constant \mathcal{N}^{\max} . For simplicity, let's consider two adjacent games $\mathcal{G} \sim \mathcal{G}'$ which differs at the utility matrices on edge $(v_1, v_2) \in E$, and $\sigma = 0$ in Algorithm 1, *i.e.* the noise-free case.

At timestep $t, \pi_i^{(t)}$ in \mathcal{G} and its counterpart $\pi_i^{\prime(t)}$ in \mathcal{G}' are identical if $t \leq \min \{ \operatorname{dist}(i, v_1), \operatorname{dist}(i, v_2) \}$. This can be proved by mathematical induction.

Therefore, Rényi divergence between $\mu^{(t)}_{\mathcal{G},i}(\cdot\,|\,\widehat{o}^{(1)},\cdots,\widehat{o}^{(t-1)})$ and $\mu^{(t)}_{\mathcal{G}',i}(\cdot\,|\,\widehat{o}^{(1)},\cdots,\widehat{o}^{(t-1)})$ is 0, since they are both normal distribution with mean $\pi_i^{(t)}$ and variance 0.

Recall that the graph is sparse and the degree of each player is bounded by a constant \mathcal{N}^{\max} . In such cases, most players are at least of distance $\mathcal{O}(\log_{\mathcal{N}^{\max}} N)$ to the edge (v_1, v_2) . Therefore, if T in Algorithm 1 is no larger than $\mathcal{O}(\log_{\mathcal{N}^{\max}} N)$, the Rényi divergence of most players at timestep t is 0. Therefore, by averaging across all players, the privacy budget will be 0 as the number of players goes to infinity.

The general proof for the cases when $\sigma > 0$ is shown in Appendix D.4.

In the following, we will introduce the proof for dense graphs (\$\\display\$) first and that for the sparse graphs (all later.

D.3 PROOF FOR DENSE GRAPHS ()

 Let $\overline{\mathcal{N}} \coloneqq \left(\frac{1}{N} \sum_{i \in [N]} \frac{1}{|\mathcal{N}(i)|}\right)^{-1}$ be the harmonic mean of players' degrees. When the graph is dense, $\overline{\mathcal{N}}$ will be relatively large. Now, we will bound the Rényi DP. Let (v_1, v_2) be the edge on which the utility matrix differs in \mathcal{G} and \mathcal{G}' . Then, we define the set \mathcal{S} as

$$\mathcal{S} := \begin{cases} [N] & |\mathcal{N}(v_1)|, |\mathcal{N}(v_2)| \ge \left(\overline{\mathcal{N}}\right)^{2/9} \\ [N] \setminus \{v_1\} & |\mathcal{N}(v_1)| < \left(\overline{\mathcal{N}}\right)^{2/9}, |\mathcal{N}(v_2)| \ge \left(\overline{\mathcal{N}}\right)^{2/9} \\ [N] \setminus \{v_2\} & |\mathcal{N}(v_2)| < \left(\overline{\mathcal{N}}\right)^{2/9}, |\mathcal{N}(v_1)| \ge \left(\overline{\mathcal{N}}\right)^{2/9} \\ [N] \setminus \{v_1, v_2\} & |\mathcal{N}(v_1)|, |\mathcal{N}(v_2)| < \left(\overline{\mathcal{N}}\right)^{2/9}. \end{cases}$$

Consider the marginal distribution $\mu_{\mathcal{G},\mathcal{S}}^{(t)}$ of $\mu_{\mathcal{G}}^{(t)}$ over $\left(\left\{\pi_i^{(s)}+n_i^{(s)}\right\}_{i\in\mathcal{S}}\right)_{s=1}^t$. Then, $D_{\alpha}\left(\mu_{\mathcal{G},\mathcal{S}}^{(t)},\mu_{\mathcal{G}',\mathcal{S}}^{(t)}\right)\geq D_{\alpha}\left(\mu_{\mathcal{G},i}^{(t)},\mu_{\mathcal{G}',i}^{(t)}\right)$ for any $i\in[N]\setminus\{v_1,v_2\}$ and t>0 since $i\in\mathcal{S}$ and enlarging the set of observations will not decrease the Rényi divergence. This is because post-processing (delete observations) will not increase Rényi divergence (Mironov, 2017).

D.3.1 $v_1, v_2 \in S$

By Lemma D.1, let $\pi_i^{(t)}$ be the counterpart of $\pi_i^{(t)}$ for any $i \in [N]$, we have

$$D_{\alpha}\left(\mu_{\mathcal{G},\mathcal{S}}^{(T)},\mu_{\mathcal{G}',\mathcal{S}}^{(T)}\right) \leq \sum_{t=1}^{T} \sup_{\left(\widehat{o}^{(1)},\widehat{o}^{(2)},\dots,\widehat{o}^{(t-1)}\right)} D_{\alpha}\left(\mu_{\mathcal{G},\mathcal{S}}(\cdot \mid \widehat{o}^{(1)},\widehat{o}^{(2)},\dots,\widehat{o}^{(t-1)}),\mu_{\mathcal{G}',\mathcal{S}}(\cdot \mid \widehat{o}^{(1)},\widehat{o}^{(2)},\dots,\widehat{o}^{(t-1)})\right)$$

$$\stackrel{(i)}{=} \frac{\alpha}{2\sigma^{2}} \sum_{t=1}^{T} \sum_{i \in [N]} \sup_{\left(\widehat{o}^{(1)},\widehat{o}^{(2)},\dots,\widehat{o}^{(t-1)}\right)} \left\|\pi_{i}^{(t)} - \pi_{i}^{\prime(t)}\right\|^{2}.$$

(i) is by the Rényi divergence of multi-variate Gaussian distribution (Gil et al., 2013), since the distribution of $\pi_i^{(t)} + n_i^{(t)}$ is a multivariate gaussian with mean $\pi_i^{(t)}$ and variance $\sigma^2 I_{\mathcal{A}_i}$ ($I_{\mathcal{A}_i}$ is the identity matrix indexed by $\mathcal{A}_i \times \mathcal{A}_i$). The inequality above implies that the Rényi divergence can be bounded by the squared 2-norm of $\pi_i^{(t)} - {\pi'}_i^{(t)}$, given all past observations $(\widehat{o}^{(1)}, \widehat{o}^{(2)}, \cdots, \widehat{o}^{(t-1)})$ are identical in $\mathcal{G}, \mathcal{G}'$.

Given the observations are identical in \mathcal{G},\mathcal{G}' , the gradient $\overline{\boldsymbol{g}}_i^{(t-1)}$ of all players except v_1,v_2 are identical in \mathcal{G},\mathcal{G}' since the gradient only depends on $\pi_i^{(t-1)}+\boldsymbol{n}_i^{(t-1)}$ and the utility matrices, which are identical for $i\in[N]\setminus\{v_1,v_2\}$. Therefore, $\pi_i^{(t)}=\pi_i'^{(t)}$ for any $i\in[N]\setminus\{v_1,v_2\}$. For $i\in\{v_1,v_2\}$, we have

$$\left\| \pi_{i}^{(t)} - \pi_{i}^{\prime(t)} \right\| = \left\| \operatorname{Proj}_{\Delta^{\mathcal{A}_{i}}} \left(\frac{\overline{\pi}_{i}^{(t-1)} - \eta \overline{g}_{i}^{(t-1)}}{1 + \eta \tau_{i}} \right) - \operatorname{Proj}_{\Delta^{\mathcal{A}_{i}}} \left(\frac{\overline{\pi}_{i}^{(t-1)} - \eta g_{i}^{\prime(t-1)}}{1 + \eta \tau_{i}} \right) \right\|$$

$$\leq \left\| \frac{\overline{\pi}_{i}^{(t-1)} - \eta \overline{g}_{i}^{(t-1)}}{1 + \eta \tau_{i}} - \frac{\overline{\pi}_{i}^{(t-1)} - \eta g_{i}^{\prime(t-1)}}{1 + \eta \tau_{i}} \right\|$$

$$\leq \eta \left\| \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \left(U_{i,j} - U_{i,j}^{\prime} \right) \overline{\pi}_{j}^{(t-1)} \right\|$$

$$\stackrel{(i)}{\leq} \frac{2\eta}{|\mathcal{N}(i)|} \sqrt{|\mathcal{A}_{i}|}.$$

(i) is because $U_{i,j} \in [-1,1]^{\mathcal{A}_i \times \mathcal{A}_j}$ and $U_{i,j} \neq U'_{i,j}$ only if $(i,j) = (v_1,v_2)$ or $(i,j) = (v_2,v_1)$. Therefore, since the distribution of $\pi_i^{(t)} + n_i^{(t)}$ is actually a multivariate gaussian with mean $\pi_i^{(t)}$ and variance $\sigma^2 I_{\mathcal{A}_i}$, so that the Rényi divergence is bounded by $\left(\frac{2\alpha\eta^2}{\sigma^2|\mathcal{N}(v_1)|^2}|\mathcal{A}_{v_1}| + \frac{2\alpha\eta^2}{\sigma^2|\mathcal{N}(v_2)|^2}|\mathcal{A}_{v_2}|\right)T$. By the definition of \mathcal{S} , it is further bounded by $\frac{4\alpha\eta^2}{\sigma^2(\overline{\mathcal{N}})^{4/9}}AT$.

D.3.2 $v_1, v_2 \notin \mathcal{S}$

When the adversary may only observe $\left(\pi_k^{(s)} + \boldsymbol{n}_k^{(s)}\right)_{s=1}^t$ for some $k \in \{v_1, v_2\}$, by Lemma D.1, we only need to bound $\left\|\pi_k^{(t)} - {\pi'}_k^{(t)}\right\|^2$ when all past observations $\left(\pi_k^{(s)} + \boldsymbol{n}_k^{(s)}\right)_{s=1}^{t-1}$ are the same in $\mathcal{G}, \mathcal{G}'$. Then,

$$\begin{aligned} \left\| \boldsymbol{\pi}_{k}^{(t)} - \boldsymbol{\pi'}_{k}^{(t)} \right\| &= \left\| \operatorname{Proj}_{\Delta^{\mathcal{A}_{k}}} \left(\frac{\overline{\boldsymbol{\pi}}_{k}^{(t-1)} - \eta \boldsymbol{g}_{k}^{(t-1)}}{1 + \eta \tau_{k}} \right) - \operatorname{Proj}_{\Delta^{\mathcal{A}_{k}}} \left(\frac{\overline{\boldsymbol{\pi}}_{k}^{(t-1)} - \eta \boldsymbol{g}_{k}^{\prime(t-1)}}{1 + \eta \tau_{k}} \right) \right\| \\ &\leq \left\| \frac{\overline{\boldsymbol{\pi}}_{k}^{(t-1)} - \eta \boldsymbol{g}_{k}^{(t-1)}}{1 + \eta \tau_{k}} - \frac{\overline{\boldsymbol{\pi}}_{k}^{(t-1)} - \eta \boldsymbol{g}_{k}^{\prime(t-1)}}{1 + \eta \tau_{k}} \right\| \\ &\leq \eta \left\| \boldsymbol{g}_{k}^{(t-1)} - \boldsymbol{g}_{k}^{\prime(t-1)} \right\| \\ &\leq 2\eta \sqrt{|\mathcal{A}_{k}|}. \end{aligned}$$

For $k \in [N] \setminus \{v_1, v_2\}$, similar to the proof in the previous section, we will augment the observations of the adversary first. The adversary may observe $\left(\left\{\boldsymbol{\pi}_i^{(s)} + \boldsymbol{n}_i^{(s)}\right\}_{i \in [N] \setminus \{v_1, v_2\}}\right)_{s=1}^t$ and $\left(\left\{\boldsymbol{n}_{v_1}^{(s)}, \boldsymbol{n}_{v_2}^{(s)}\right\}\right)_{s=1}^t$ at timestep t. Note that the actual observation of the adversary is still $\left(\boldsymbol{\pi}_k^{(s)} + \boldsymbol{n}_k^{(s)}\right)_{s=1}^t$. We augment the observations to simplify the proof, since the Rényi divergence of the the augmented observation's distribution upperbounds that of the actual observation.

Still, by the chain rule of Rényi divergence, we need to bound $\sum_{i \in [N] \setminus \{v_1, v_2\}} \left\| \pi_i^{(t)} - {\pi'}_i^{(t)} \right\|$ for any $t \in [T]$. For any $i \notin \mathcal{N}(v_1) \cup \mathcal{N}(v_2)$, similar to the discussion in previous section, the gradients $\bar{\boldsymbol{g}}_i^{(t)}$ and $\bar{\pi}_i^{(t-1)}$ are identical in $\mathcal{G}, \mathcal{G}'$, so that the $\pi_i^{(t)} = {\pi'}_i^{(t)}$.

For $i \in \mathcal{N}(v_1) \cup \mathcal{N}(v_2) \setminus \{v_1, v_2\}$, we have

$$\begin{aligned} \left\| \pi_{i}^{(t)} - \pi_{i}^{\prime(t)} \right\| &\leq \eta \left\| \bar{\mathbf{g}}_{i}^{(t-1)} - \mathbf{g}_{i}^{\prime(t-1)} \right\| = \frac{\eta}{|\mathcal{N}(i)|} \left\| \sum_{j \in \mathcal{N}(i)} \mathbf{U}_{i,j} \left(\bar{\pi}_{j}^{(t-1)} - \bar{\pi}_{j}^{\prime(t-1)} \right) \right\| \\ &\stackrel{(i)}{=} \frac{\eta}{|\mathcal{N}(i)|} \left\| \sum_{j \in \{v_{1}, v_{2}\} \cap \mathcal{N}(i)} \mathbf{U}_{i,j} \left(\bar{\pi}_{j}^{(t-1)} - \bar{\pi}_{j}^{\prime(t-1)} \right) \right\| \\ &\leq \frac{\eta \sqrt{|\mathcal{A}_{i}|}}{|\mathcal{N}(i)|} \sum_{j \in \{v_{1}, v_{2}\} \cap \mathcal{N}(i)} \left\| \bar{\pi}_{j}^{(t-1)} - \bar{\pi}_{j}^{\prime(t-1)} \right\|_{1} \\ &\leq \frac{\eta A}{|\mathcal{N}(i)|} \sum_{j \in \{v_{1}, v_{2}\} \cap \mathcal{N}(i)} \left\| \bar{\pi}_{j}^{(t-1)} - \bar{\pi}_{j}^{\prime(t-1)} \right\|. \end{aligned}$$

 $(i) \text{ uses the fact that } \overline{\pi}_i^{(t-1)} = \overline{\pi}_i'^{(t-1)} \text{ for } i \not\in \{v_1, v_2\}. \text{ Since, } \boldsymbol{n}_j^{(t-1)} = \boldsymbol{n}_j'^{(t-1)}, \text{ for any } j \in \{v_1, v_2\}, \text{ we have } \boldsymbol{n}_j^{(t-1)} = \boldsymbol{n}_j'^{(t-1)}, \text{ for any } \boldsymbol{n}_j^{(t-1)} = \boldsymbol{n}_j'^{(t-1)}, \text{$

$$\begin{split} \left\| \overline{\pi}_{j}^{(t-1)} - \overline{\pi}_{j}^{\prime(t-1)} \right\| &= \left\| \operatorname{Proj}_{\Delta^{A_{j}}} \left(\pi_{j}^{(t-1)} + \boldsymbol{n}_{j}^{(t-1)} \right) - \operatorname{Proj}_{\Delta^{A_{j}}} \left(\pi_{j}^{\prime(t-1)} + \boldsymbol{n}_{j}^{\prime(t-1)} \right) \right\| \\ &\leq \left\| \pi_{j}^{(t-1)} + \boldsymbol{n}_{j}^{(t-1)} - \pi_{j}^{\prime(t-1)} - \boldsymbol{n}_{j}^{\prime(t-1)} - \boldsymbol{n}_{j}^{\prime(t-1)} \right\| \\ &= \left\| \pi_{j}^{(t-1)} - \pi_{j}^{\prime(t-1)} \right\|. \end{split}$$

Moreover, due to the additional regularization, for $j \in \{v_1, v_2\}$, $\left\|\pi_j^{(t-1)} - \pi_j^{(t-1)}\right\|$ will be bounded as follows according to Lemma D.2.

$$\left\| \pi_j^{(t-1)} - {\pi'}_j^{(t-1)} \right\| \le \frac{2\sqrt{\mathcal{A}_j}}{\tau_j} \stackrel{(i)}{\le} 2\sqrt{\mathcal{A}_j} \frac{(\overline{\mathcal{N}})^{2/9} \log N}{(\overline{\mathcal{N}})^{5/9}} = \frac{2\sqrt{\mathcal{A}_j} \log N}{(\overline{\mathcal{N}})^{1/3}}.$$

(i) is because when $v_1 \notin \mathcal{S}$, $|\mathcal{N}(v_1)| < \left(\overline{\mathcal{N}}\right)^{2/9}$ so that $\tau_{v_1} \geq \frac{\left(\overline{\mathcal{N}}\right)^{5/9}}{\left(\overline{\mathcal{N}}\right)^{2/9}\log N}$ by definition. The same argument also holds for v_2 . Therefore, $\left\|\pi_i^{(t)} - {\pi'}_i^{(t)}\right\| \leq \frac{4\eta A^{3/2}\log N}{\left(\overline{\mathcal{N}}\right)^{1/3}|\mathcal{N}(i)|}$, where $A = \max_{i \in [N]} |\mathcal{A}_i|$. Furthermore, for $i \in [N] \setminus \{v_1, v_2\}$

$$D_{\alpha}\left(\mu_{\mathcal{G},i}^{(T)},\mu_{\mathcal{G}',i}^{(T)}\right) \leq D_{\alpha}\left(\mu_{\mathcal{G},\mathcal{S}}^{(T)},\mu_{\mathcal{G}',\mathcal{S}}^{(T)}\right) \leq \frac{\alpha}{2\sigma^{2}}\left|\mathcal{N}(v_{1})\cup\mathcal{N}(v_{2})\right| \cdot \left(\frac{16\eta^{2}A^{3}\left(\log N\right)^{2}}{\left(\overline{\mathcal{N}}\right)^{2/3}\min_{i\in\mathcal{N}(v_{1})\cup\mathcal{N}(v_{2})}\left|\mathcal{N}(i)\right|^{2}}\right)T$$

$$\leq \frac{\alpha}{2\sigma^{2}}\left(2\left(\overline{\mathcal{N}}\right)^{2/9}\right) \cdot \left(\frac{16\eta^{2}A^{3}\left(\log N\right)^{2}}{\left(\overline{\mathcal{N}}\right)^{2/3}\min_{i\in\mathcal{N}(v_{1})\cup\mathcal{N}(v_{2})}\left|\mathcal{N}(i)\right|^{2}}\right)T$$

$$= \frac{16\alpha\eta^{2}A^{3}\left(\log N\right)^{2}}{\sigma^{2}\left(\overline{\mathcal{N}}\right)^{4/9}\min_{i\in\mathcal{N}(v_{1})\cup\mathcal{N}(v_{2})}\left|\mathcal{N}(i)\right|^{2}}T.$$

Finally, since $\min_{i \in \mathcal{N}(v_1) \cup \mathcal{N}(v_2)} |\mathcal{N}(i)|^2 \geq 1$, we can conclude the proof. Moreover, $\min_{i \in \mathcal{N}(v_1) \cup \mathcal{N}(v_2)} |\mathcal{N}(i)|^2$ can be much larger than a constant in practice, *e.g.*, Erdős–Rényi graphs.

The proofs for the rest possibilities $(v_1 \in \mathcal{S}, v_2 \notin \mathcal{S} \text{ and } v_1 \notin \mathcal{S}, v_2 \in \mathcal{S})$ are similar.

D.4 PROOF FOR SPARSE GRAPHS ()

In this section, we show the DP guarantee when the associated graph of the polymatrix game is sparse. Formally, there exists a constant $\mathcal{N}^{\max} > 0$ so that $|\mathcal{N}(i)| \leq \mathcal{N}^{\max}$ for any $i \in [N]$.

We will augment the observations from $\left(\pi_i^{(s)} + \boldsymbol{n}_i^{(s)}\right)_{s=1}^{t-1}$ to $\left(\left\{\pi_j^{(s)} + \boldsymbol{n}_j^{(s)}\right\}_{j \in \mathcal{S}_i^{(s)}}\right)_{s=1}^t$, where

$$\mathcal{S}_i^{(s)} \coloneqq \begin{cases} \{j \colon \min \left\{ \mathsf{dist}(j, v_1), \mathsf{dist}(j, v_2) \right\} \geq s \} & \min \left\{ \mathsf{dist}(i, v_1), \mathsf{dist}(i, v_2) \right\} \geq s \\ \{i\} & \mathsf{Otherwise}. \end{cases}$$

When $s \leq \min \{ \operatorname{dist}(i, v_1), \operatorname{dist}(i, v_2) \}$, we will show in the following that with all past observations identical in $\mathcal{G}, \mathcal{G}'$, then $\pi_j^{(s)}$ and its counterpart $\pi'_j^{(s)}$ in \mathcal{G}' are identical for any $j \in \mathcal{S}_i^{(s)}$. When s = 0, $\mathcal{S}_i^{(0)} = [N]$ and $\pi_j^{(0)} = \pi'_j^{(0)}$ for any $j \in [N]$ since they are both initialized to uniform distribution over \mathcal{A}_j . When s > 0, for any $j \in \mathcal{S}_i^{(s)}$, we must have $\mathcal{N}(j) \subseteq \mathcal{S}_i^{(s-1)}$, otherwise $\min \{ \operatorname{dist}(j, v_1), \operatorname{dist}(j, v_2) \} < s$. Therefore, since all past observations are identical, we have $\overline{\pi}_k^{(s-1)}$ are identical in $\mathcal{G}, \mathcal{G}'$ for any $k \in \mathcal{N}(j)$ so that the gradient $\mathbf{g}_j^{(s-1)} = \mathbf{g}_j'^{(s-1)}$. Moreover, given $\overline{\pi}_j^{(s-1)}$ are identical, $\pi_j^{(s)}$ are identical in $\mathcal{G}, \mathcal{G}'$.

When $s > \min \{ \operatorname{dist}(i, v_1), \operatorname{dist}(i, v_2) \}$, $\left\| \pi_i^{(s)} - {\pi'}_i^{(s)} \right\| \leq 2\eta \sqrt{|\mathcal{A}_i|}$, given the past observation $\pi_i^{(s-1)} + \boldsymbol{n}_i^{(s-1)}$ is identical.

By Lemma D.1, for any T > 0,

$$\begin{split} D_{\alpha}\left(\mu_{\mathcal{G},i}^{(T)},\mu_{\mathcal{G}',i}^{(T)}\right) &\leq \max\left\{0,T-\min\left\{\operatorname{dist}(i,v_{1}),\operatorname{dist}(i,v_{2})\right\}\right\}\frac{2\alpha\eta^{2}}{\sigma^{2}}|\mathcal{A}_{i}|\\ &\leq \mathbb{1}\left(T>\min\left\{\operatorname{dist}(i,v_{1}),\operatorname{dist}(i,v_{2})\right\}\right)\frac{2\alpha\eta^{2}}{\sigma^{2}}|\mathcal{A}_{i}|T. \end{split}$$

D.5 OMITTED PROOF OF LEMMAS

Lemma D.2. Consider the Algorithm 1. For any player $i \in [N]$ and timestep t > 0, by updating the strategy $\left\{\pi_i^{(s)}\right\}_{s=0}^t$, $\left\{\pi_i^{(s)}\right\}_{s=0}^t$ with the same noise $\left\{n_i^{(s)}\right\}_{s=0}^t$ but different gradients $\left\{g_i^{(s)}\right\}_{s=0}^t, \left\{g_i'^{(s)}\right\}_{s=0}^t$ individually, we have

$$\left\| \pi_i^{(t)} - {\pi'}_i^{(t)} \right\| \le \frac{2\sqrt{\mathcal{A}_i}}{\tau_i}. \tag{28}$$

Proof.

$$\left\| \boldsymbol{\pi}_{i}^{(t+1)} - \boldsymbol{\pi}_{i}^{\prime(t+1)} \right\| = \left\| \operatorname{Proj}_{\Delta^{A_{i}}} \left(\frac{\operatorname{Proj}_{\Delta^{A_{i}}} \left(\boldsymbol{\pi}_{i}^{(t)} + \boldsymbol{n}_{i}^{(t)} \right) - \eta \bar{\boldsymbol{g}}_{i}^{(t)}}{1 + \eta \tau_{i}} \right) - \operatorname{Proj}_{\Delta^{A_{i}}} \left(\frac{\operatorname{Proj}_{\Delta^{A_{i}}} \left(\boldsymbol{\pi}_{i}^{\prime(t)} + \boldsymbol{n}_{i}^{(t)} \right) - \eta \boldsymbol{g}_{i}^{\prime(t)}}{1 + \eta \tau_{i}} \right) \right\|$$

$$\leq \left\| \frac{\operatorname{Proj}_{\Delta^{A_{i}}} \left(\boldsymbol{\pi}_{i}^{(t)} + \boldsymbol{n}_{i}^{(t)} \right) - \eta \bar{\boldsymbol{g}}_{i}^{(t)}}{1 + \eta \tau_{i}} - \frac{\operatorname{Proj}_{\Delta^{A_{i}}} \left(\boldsymbol{\pi}_{i}^{\prime(t)} + \boldsymbol{n}_{i}^{(t)} \right) - \eta \boldsymbol{g}_{i}^{\prime(t)}}{1 + \eta \tau_{i}} \right\|$$

$$\leq \frac{1}{1 + \eta \tau_{i}} \left\| \operatorname{Proj}_{\Delta^{A_{i}}} \left(\boldsymbol{\pi}_{i}^{(t)} + \boldsymbol{n}_{i}^{(t)} \right) - \operatorname{Proj}_{\Delta^{A_{i}}} \left(\boldsymbol{\pi}_{i}^{\prime(t)} + \boldsymbol{n}_{i}^{(t)} \right) \right\| + \frac{\eta}{1 + \eta \tau_{i}} \left\| \bar{\boldsymbol{g}}_{i}^{(t)} - \boldsymbol{g}_{i}^{\prime(t)} \right\|$$

$$\leq \frac{1}{1 + \eta \tau_{i}} \left\| \boldsymbol{\pi}_{i}^{(t)} - \boldsymbol{\pi}_{i}^{\prime(t)} \right\| + \frac{\eta}{1 + \eta \tau_{i}} \left\| \bar{\boldsymbol{g}}_{i}^{(t)} - \boldsymbol{g}_{i}^{\prime(t)} \right\|.$$

By recursively applying the inequality above, we ha

$$\left\| \pi_{i}^{(t+1)} - {\pi'}_{i}^{(t+1)} \right\| \leq \left(\frac{1}{1 + \eta \tau_{i}} \right)^{t+1} \left\| \pi_{i}^{(0)} - {\pi'}_{i}^{(0)} \right\| + \eta \sum_{s=0}^{t} \left(\frac{1}{1 + \eta \tau_{i}} \right)^{t-s+1} \left\| \boldsymbol{g}_{i}^{(s)} - {\boldsymbol{g'}}_{i}^{(s)} \right\|$$

$$\stackrel{(i)}{=} \eta \sum_{s=0}^{t} \left(\frac{1}{1 + \eta \tau_{i}} \right)^{t-s+1} \left\| \boldsymbol{g}_{i}^{(s)} - {\boldsymbol{g'}}_{i}^{(s)} \right\| .$$

(i) is because $\pi_i^{(0)}, {\pi'}_i^{(0)}$ are both initialized as uniform distribution over \mathcal{A}_i . Therefore, since each element of the gradient is bounded by [-1,1] by definition, we have

$$\left\| \pi_i^{(t+1)} - {\pi'}_i^{(t+1)} \right\| \le 2\eta \sum_{s=0}^t \left(\frac{1}{1 + \eta \tau_i} \right)^{t-s+1} \sqrt{\mathcal{A}_i} \le 2\eta \sqrt{\mathcal{A}_i} \frac{1}{(1 + \eta \tau_i) \left(1 - \frac{1}{1 + \eta \tau_i} \right)} = \frac{2\sqrt{\mathcal{A}_i}}{\tau_i}.$$

Lemma D.1 (Chain Rule of Rényi divergence). For any distribution p, q over random variables X^1, X^2 , for any $\alpha \geq 1$, we have

$$D_{\alpha}\left(p(X^{1}, X^{2}), q(X^{1}, X^{2})\right) \le D_{\alpha}\left(p(X^{1}), q(X^{1})\right) + \sup_{\widehat{x}^{1}} D_{\alpha}\left(p(X^{2} \mid X^{1} = \widehat{x}^{1}), q(X^{2} \mid X^{1} = \widehat{x}^{1})\right).$$
(27)

Proof. By definition, when $\alpha > 1$, we have

$$\begin{array}{ll} & D_{\alpha}\left(p(X^{1},X^{2}),q(X^{1},X^{2})\right) \\ & = \frac{1}{\alpha-1}\log\int_{x^{1},x^{2}}\left(p(x^{1},x^{2})\right)^{\alpha}\left(q(x^{1},x^{2})\right)^{1-\alpha}dx^{1}dx^{2} \\ & = \frac{1}{\alpha-1}\log\int_{x^{1}}\left(\int_{x^{2}}\left(p(x^{2}\,|\,x^{1})\right)^{\alpha}\left(q(x^{2}\,|\,x^{1})\right)^{1-\alpha}dx^{2}\right)\left(p(x^{1})\right)^{\alpha}\left(q(x^{1})\right)^{1-\alpha}dx^{1} \\ & = \frac{1}{\alpha-1}\log\sup_{\widehat{x}^{1}}\left(\int_{x^{2}}\left(p(x^{2}\,|\,x^{1})\right)^{\alpha}\left(q(x^{2}\,|\,x^{1})\right)^{1-\alpha}dx^{2}\right)\left(\int_{x^{1}}\left(p(x^{1})\right)^{\alpha}\left(q(x^{1})\right)^{1-\alpha}dx^{1}\right) \\ & = \frac{1}{\alpha-1}\sup\sup_{\widehat{x}^{1}}\log\int_{x^{2}}\left(p(x^{2}\,|\,\widehat{x}^{1})\right)^{\alpha}\left(q(x^{2}\,|\,\widehat{x}^{1})\right)^{1-\alpha}dx^{2} + \frac{1}{\alpha-1}\log\int_{x^{1}}\left(p(x^{1})\right)^{\alpha}\left(q(x^{1})\right)^{1-\alpha}dx^{1} \\ & = D_{\alpha}\left(p(X^{1}),q(X^{1})\right) + \sup_{x \in X}D_{\alpha}\left(p(X^{2}\,|\,X^{1})=\widehat{x}^{1}),q(X^{2}\,|\,X^{1}=\widehat{x}^{1})\right). \end{array}$$

(i) is by Tonelli's theorem. (ii) is by Hölder's Inequality.

When $\alpha = 1$, the corresponding Rényi divergence is KL-divergence. Therefore,

$$\begin{split} &D_{1}\left(p(X^{1},X^{2}),q(X^{1},X^{2})\right) \\ &= \int_{x^{1},x^{2}} p(x^{1},x^{2}) \log \frac{p(x^{1},x^{2})}{q(x^{1},x^{2})} dx^{1} dx^{2} \\ &= \int_{x^{1},x^{2}} p(x^{2} \mid x^{1}) p(x^{1}) \log \frac{p(x^{2} \mid x^{1}) p(x^{1})}{q(x^{2} \mid x^{1}) q(x^{1})} dx^{1} dx^{2} \\ &= \int_{x^{1},x^{2}} p(x^{2} \mid x^{1}) p(x^{1}) \log \frac{p(x^{1})}{q(x^{1})} dx^{1} dx^{2} + \int_{x^{1},x^{2}} p(x^{2} \mid x^{1}) p(x^{1}) \log \frac{p(x^{2} \mid x^{1})}{q(x^{2} \mid x^{1})} dx^{1} dx^{2} \\ &= D_{1}\left(p(X^{1}), q(X^{1})\right) + \mathbb{E}_{\widehat{x}^{1} \sim p(X^{1})}\left[D_{1}\left(p(X^{2} \mid X^{1} = \widehat{x}^{1}), q(X^{2} \mid X^{1} = \widehat{x}^{1})\right)\right] \\ &\leq D_{1}\left(p(X^{1}), q(X^{1})\right) + \sup_{\widehat{x}^{1}} D_{1}\left(p(X^{2} \mid X^{1} = \widehat{x}^{1}), q(X^{2} \mid X^{1} = \widehat{x}^{1})\right). \end{split}$$

E EXPERIMENTS

 The experimental results are shown in Figures 4 and 5. The baseline algorithm is adapted from Huang et al. (2015). Our implementation uses PyTorch (Paszke et al., 2019) to enable efficient, fully parallel updates of all players' strategies, and all runs were executed on $8 \times$ NVIDIA A100 GPUs. The error bars denote 1σ .

On dense graphs, as the number of players increases, our algorithm's exploitability and the privacy budget both decrease, whereas the baseline's exploitability increases. On sparse graphs, by contrast, only the privacy budget decreases. A plausible explanation is that the convergence rate on sparse graphs is $\mathcal{O}\left(\frac{1}{(\log N)^{1/3}}\right)$, which is too mild to overcome constant factors and stochastic noise for $N \leq 2^{16}$. Hyperparameters were chosen according to Theorem 7.1. The code is provided in the supplementary materials.

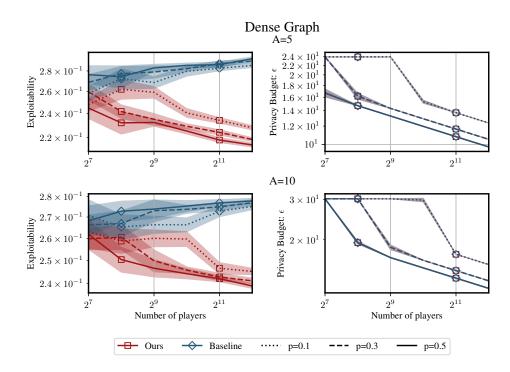


Figure 4: Experiment results of dense graphs. Each node (player) in the graph will be connected to another node with probability p, the connection is sampled i.i.d. for each node. Then, the duplicate edges will be removed. The action set sizes of all players are set to A. We can see from the result that the exploitability and the privacy budget both decrease as the number of players increases.

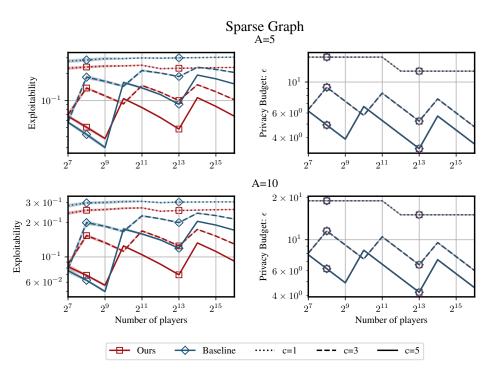


Figure 5: Experiment results of sparse graphs. We will randomly generate cN edges in total, with each node appearing c times. Then, duplicate edges and self-loops will be removed. The action set sizes of all players are set to A. The result shows that while the exploitability remains unchanged, the exploitability decreases as the number of players increases.