

MONOTONE NEAR-ZERO-SUM GAMES: A GENERALIZATION OF CONVEX-CONCAVE MINIMAX

Ruichen Luo*

IST Austria

3400 Klosterneuburg, Austria

Sebastian U. Stich

CISPA Helmholtz Center

66386 St. Ingbert, Germany

Krishnendu Chatterjee

IST Austria

3400 Klosterneuburg, Austria

ABSTRACT

Zero-sum and non-zero-sum (aka general-sum) games are relevant in a wide range of applications. While general non-zero-sum games are computationally hard, researchers focus on the special class of monotone games for gradient-based algorithms. However, there is a substantial gap between the gradient complexity of monotone zero-sum and monotone general-sum games. Moreover, in many practical scenarios of games the zero-sum assumption needs to be relaxed. To address these issues, we define a new intermediate class of *monotone near-zero-sum games* that contains monotone zero-sum games as a special case. Then, we present a novel algorithm that transforms the near-zero-sum games into a sequence of zero-sum subproblems, improving the gradient-based complexity for the class. Finally, we demonstrate the applicability of this new class to model practical scenarios of games motivated from the literature.

1 INTRODUCTION

Two-player zero-sum games (also known as strictly competitive games) and their generalization to non-zero-sum games (Nash, 1951; Rosen, 1965) are crucial in domains like economics (von Neumann & Morgenstern, 1947), artificial intelligence (Yannakakis & Togelius, 2018), and biology in the form of evolutionary game theory (Weibull, 1997; Smith, 1982). As initiated in Rosen (1965); Tseng (1995); Nemirovski (2004), this paper focuses on the computationally tractable class of *monotone games* with compact convex strategy spaces.

Early work studies the restrictive classes of games where the two players have *equal conditioning*: the seminal paper of Nesterov (2005) established the complexity for bilinearly-coupled zero-sum games, which was later extended by Nemirovski (2004) to general-sum classes. Later on, the field explores the broader concept of *general conditioning*, which allows for more detailed portrayal of the two players (Chambolle & Pock, 2011; Lin et al., 2020). Early explorations of Chambolle & Pock (2011); Chen et al. (2014) focus on the bilinearly-coupled cases; starting from Lin et al. (2020), there has been extensive research into the generally-coupled, generally-conditioned, and zero-sum games (Yang et al., 2020; Wang & Li, 2020; Kovalev & Gasnikov, 2022; Lan & Li, 2023; Bo̧ et al., 2023). However, it has been an interesting open question whether these recent developments can be further generalized to the broader classes of generally-coupled, generally-conditioned, and *non-zero-sum* games.

On the theory side, there is a substantial gap between the gradient complexity of monotone zero-sum games and that of monotone general-sum games. The recent developments in minimax optimization (Lin et al., 2020; Kovalev & Gasnikov, 2022; Lan & Li, 2023) establish much better complexity results for the zero-sum cases compared to the long-standing bounds of general-sum games (Rockafellar, 1976; Tseng, 1995).

On the application side, strictly competitive scenarios, modeled by monotone zero-sum games, are often insufficient. Real-world game settings frequently involve factors such as transaction fees or semi-cooperation (Kalai & Kalai, 2013; Halpern & Rong, 2013), necessitating a relaxation of the zero-sum assumption in the modeling.

*This work was partially done during RL's stay at CISPA.

Our contributions Our work makes the first step towards bridging the gap between the monotone zero-sum and general-sum classes. For this purpose, we introduce a new intermediate class of monotone games, present a novel algorithm for this class, and show the applicability of this new class. In detail:

- *Theoretical motivation.* Our work extends the recent studies of monotone zero-sum games to the more general non-zero-sum settings. Specifically, we define a new intermediate class of games called *monotone near-zero-sum games*, characterized by a smoothness parameter δ describing the game’s proximity to a zero-sum game. This new class of games presents a natural interpolation between monotone zero-sum games and a class of monotone general-sum games based on the near-zero-sum parameter δ , and thus, it partially bridges the gap of the monotone zero-sum and general-sum classes.
- *Main theoretical result.* We propose a novel algorithm, Iterative Coupling Linearization (ICL), that provides a black-box reduction from monotone near-zero-sum games to zero-sum games. It converges to an ε -Nash equilibrium within $\tilde{\mathcal{O}}\left(\left(\frac{L}{\sqrt{\mu\nu}} + \frac{L}{\min\{\mu,\nu\}} \cdot \min\left\{1, \sqrt{\frac{\delta}{\mu+\nu}}\right\}\right) \cdot \log^2\left(\frac{D^2}{\varepsilon}\right)\right)$ gradient queries,¹ where L is the smoothness parameter, μ and ν are the strong concavity parameters of the two players, δ is the near-zero-sum parameter, and D is the diameter. When δ is small, our results improve the long-standing complexity results of Tseng (1995); Nemirovski (2004) for the first time in non-zero-sum classes.
- *Practical applications.* Besides the theoretical motivation, we demonstrate the practical relevance of this new class of games. We consider *regularized matrix games* and *competitive games with small additional incentives*. These games are not zero-sum but naturally have a near-zero-sum structure, where our methods are applied to achieve provably faster rates.

2 DEFINITIONS, PREVIOUS RESULTS, AND THE NEW PROBLEM CLASS

2.1 BASIC DEFINITIONS

This paper studies the *Nash Equilibrium Problem* (NEP) for two-person general-sum games, in which Player 1 wants to maximize its utility function $u_1(\mathbf{x}, \mathbf{y})$ over $\mathbf{x} \in X$ and Player 2 wants to maximize its utility function $u_2(\mathbf{x}, \mathbf{y})$ over $\mathbf{y} \in Y$. Here, X and Y are compact and convex sets, and $u_1(\cdot, \cdot): X \times Y \rightarrow \mathbb{R}$ and $u_2(\cdot, \cdot): X \times Y \rightarrow \mathbb{R}$ are smooth functions. A pair of decisions $(\mathbf{x}^*, \mathbf{y}^*) \in X \times Y$ is a *Nash equilibrium* if

$$u_1(\mathbf{x}^*, \mathbf{y}^*) \geq u_1(\mathbf{x}, \mathbf{y}^*), \text{ for all } \mathbf{x} \in X, \text{ and } u_2(\mathbf{x}^*, \mathbf{y}^*) \geq u_2(\mathbf{x}^*, \mathbf{y}), \text{ for all } \mathbf{y} \in Y.$$

A pair of decisions $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in X \times Y$ is an ε -*accurate Nash equilibrium* if there exists a Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ such that $\|\bar{\mathbf{x}} - \mathbf{x}^*\|^2 + \|\bar{\mathbf{y}} - \mathbf{y}^*\|^2 \leq \varepsilon$. A pair of decisions $(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \in X \times Y$ is an ε -*approximate Nash equilibrium*, if $u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \geq u_1(\mathbf{x}, \hat{\mathbf{y}}) - \varepsilon$ for all $\mathbf{x} \in X$ and $u_2(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \geq u_2(\hat{\mathbf{x}}, \mathbf{y}) - \varepsilon$ for all $\mathbf{y} \in Y$. The relation between accurate and approximate Nash equilibria is discussed in Appendix B. The goal of this paper is to find an ε -accurate (or an ε -approximate) Nash equilibrium by iterative algorithms which subsequently query the gradients of the utility functions.

Notations Let \mathcal{X} and \mathcal{Y} be Euclidean spaces. In the space $\mathcal{X} \times \mathcal{Y}$, for all $\mathbf{z} = (\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$ and $\mathbf{z}' = (\mathbf{x}', \mathbf{y}') \in \mathcal{X} \times \mathcal{Y}$, define $\langle \mathbf{z}', \mathbf{z} \rangle \stackrel{\text{def}}{=} \langle \mathbf{x}', \mathbf{x} \rangle + \langle \mathbf{y}', \mathbf{y} \rangle$. For all these spaces, the norms are those induced by inner products. Assume that the diameter of $X \subseteq \mathcal{X}$ is bounded by D_X and the diameter of $Y \subseteq \mathcal{Y}$ is bounded by D_Y . Let $D = \sqrt{D_X^2 + D_Y^2}$. Assume that $u_1(\cdot, \cdot)$ and $u_2(\cdot, \cdot)$ are L -smooth, that is,

$$\|\nabla u_1(\mathbf{z}') - \nabla u_1(\mathbf{z})\| \leq L \|\mathbf{z}' - \mathbf{z}\|, \quad \|\nabla u_2(\mathbf{z}') - \nabla u_2(\mathbf{z})\| \leq L \|\mathbf{z}' - \mathbf{z}\|, \text{ for all } \mathbf{z}, \mathbf{z}' \in X \times Y.$$

To facilitate our analysis, we adopt the following formulation that decomposes the game into a coupling part and a zero-sum part. Denote

$$g = -\frac{1}{2}(u_1 + u_2), \quad h = \frac{1}{2}(-u_1 + u_2), \quad \mathcal{H} = (\nabla_{\mathbf{x}} h, -\nabla_{\mathbf{y}} h), \quad \mathcal{F} = -(\nabla_{\mathbf{x}} u_1, \nabla_{\mathbf{y}} u_2),$$

¹In the $\tilde{\mathcal{O}}(\cdot)$ notations, the poly-logarithm terms are omitted.

where g is the coupling part, h is the zero-sum part, \mathcal{H} is the operator corresponding to the zero-sum part h , and \mathcal{F} is the operator corresponding to the game. Then, we have

$$u_1 = -g - h, \quad u_2 = -g + h, \quad \mathcal{F} = \nabla g + \mathcal{H}.$$

Since the utilities u_1 and u_2 are both L -smooth, we have the functions g and h are both L -smooth, and the operators \mathcal{H} and \mathcal{F} are both L -Lipschitz continuous. While similar decompositions can be found in the literature of variational inequalities and game theory (Nemirovski, 1995; Halpern & Rong, 2013; Chen et al., 2017; Hwang & Rey-Bellet, 2020), we emphasize that this notation is particularly suited to characterize our near-zero-sum games (to be defined later) for explicitly separating the non-zero-sum coupling part.

2.2 PROBLEM CLASSES

Monotone general-sum games The seminal work of Nash (1951); Rosen (1965) establishes the existence of Nash equilibrium for concave games. But to obtain a tractable class of NEPs, further restrictions need to be considered (Rosen, 1965). Specifically, we make the following assumptions:

Assumption 1 (Convex-concave zero-sum part). *There exists $(\mu, \nu) \in [0, L] \times [0, L]$ such that the function $h(\mathbf{x}, \mathbf{y}) - \frac{\mu}{2} \|\mathbf{x}\|^2$ is convex in \mathbf{x} for any fixed $\mathbf{y} \in Y$, and the function $h(\mathbf{x}, \mathbf{y}) + \frac{\nu}{2} \|\mathbf{y}\|^2$ is concave in \mathbf{y} for any fixed $\mathbf{x} \in X$.*

Assumption 2 (jointly convex coupling part). *The function $g(\cdot, \cdot)$ is jointly convex.*

The operator $\mathcal{H} = (\nabla_{\mathbf{x}} h, -\nabla_{\mathbf{y}} h)$ is monotone with modulus $\min\{\mu, \nu\}$ under Assumption 1, and the operator ∇g is monotone under Assumption 2. Hence, under Assumptions 1 and 2, the game (or the operator $\mathcal{F} = \nabla g + \mathcal{H}$) is monotone with modulus $\min\{\mu, \nu\}$ (Rosen, 1965; Nemirovski, 1995), that is, $\langle \mathcal{F}(\mathbf{z}') - \mathcal{F}(\mathbf{z}), \mathbf{z}' - \mathbf{z} \rangle \geq \min\{\mu, \nu\} \cdot \|\mathbf{z}' - \mathbf{z}\|^2$, for all $\mathbf{z}, \mathbf{z}' \in X \times Y$.

In this paper, we refer to a game as a *monotone (general-sum) game* if it satisfies Assumptions 1 and 2, and refer to a game as a *strongly monotone* game if it is a monotone game with modulus $\mu, \nu > 0$. It is known that there exists a unique Nash equilibrium for strongly monotone games (Rosen, 1965).

Monotone zero-sum games (convex-concave minimax optimization) We now consider a subclass: monotone zero-sum games. A two-person game is *zero-sum* if $g = 0$. A game is said to be a *monotone zero-sum game* if it is zero-sum and satisfies Assumption 1. Note that monotone zero-sum games trivially satisfy Assumption 2 (since $g = 0$ is convex), and therefore form a subclass of monotone general-sum games. By Sion’s minimax theorem (Sion, 1958), the NEP for monotone zero-sum games is equivalent to *convex-concave minimax optimization*, that is, finding or approaching a saddle point of the function $h(\cdot, \cdot)$.

2.3 PREVIOUS RESULTS

We begin with a historical overview of the related study of NEPs for monotone games. Many prior studies focus on restrictive cases. Early seminal work by Nesterov and Nemirovski in the early 2000s primarily addressed the restrictive classes of games with equal conditioning ($\mu = \nu$). Nesterov (2005) initially studied the bilinearly-coupled, equally-conditioned, and zero-sum cases, which Nemirovski (2004) later generalized to cover *generally-coupled* and *general-sum* settings. Subsequently, research has shifted towards the broader concept of *general conditioning* ($\mu \neq \nu$). This shift is motivated by the need for more flexible modeling of player behaviors in realistic problems and can lead to substantially faster convergence rates, particularly when one of the player has a better conditioning (Chambolle & Pock, 2011; Lin et al., 2020). Early explorations focus on the bilinearly-coupled and zero-sum cases, developing various primal-dual algorithms in different oracle settings (Chambolle & Pock, 2011; Chen et al., 2014; Kolmogorov & Pock, 2021; Thekumparampil et al., 2022). More recently, the seminal work of Lin et al. (2020) spurred extensive research into *generally-coupled*, *generally-conditioned*, and zero-sum games (see Yang et al. (2020); Wang & Li (2020); Zhang et al. (2022); Kovalev & Gasnikov (2022); Bōt et al. (2023); Lan & Li (2023); Lin et al. (2025), among others).

Despite these advancements, we are not aware of any study addressing the more general settings of *generally-coupled*, *generally-conditioned*, and *non-zero-sum* games. For this particularly challenging class, the only established results are the long-standing bounds from Tseng (1995); Nemirovski

(2004). These bounds, however, remain a huge gap from the optimal rates achieved for the zero-sum cases (Lin et al., 2020; Kovalev & Gasnikov, 2022; Lan & Li, 2023).

Now, we formally outline the state-of-the-art gradient complexity results of monotone general-sum and zero-sum games within the general settings of *general couplings* and *general conditioning*. To simplify the presentation, we assume strong monotonicity for now in the Section 2.3, while the results for non-strongly monotone games are indeed similar (as we will discuss later). For general-sum games, the NEPs can be solved using variational inequality methods for the operator \mathcal{F} , leading to the following long-standing gradient complexity:

Proposition 1 (Tseng (1995)). *For strongly monotone general-sum games, an ε -accurate Nash equilibrium can be found with the number of gradient queries bounded by $\mathcal{O}\left(\frac{L}{\min\{\mu, \nu\}} \cdot \log\left(\frac{D^2}{\varepsilon}\right)\right)$.*

For the zero-sum cases, the gradient complexity can be significantly improved due to recent advances in minimax optimization.

Proposition 2 (Lin et al. (2020); Kovalev & Gasnikov (2022); Zhang et al. (2022); Lan & Li (2023)). *For strongly monotone zero-sum games, an ε -accurate Nash equilibrium can be found with the number of gradient queries bounded by $\mathcal{O}\left(\frac{L}{\sqrt{\mu\nu}} \cdot \log\left(\frac{D^2}{\varepsilon}\right)\right)$. This rate is minimax optimal, as $\Omega\left(\frac{L}{\sqrt{\mu\nu}} \cdot \log\left(\frac{D^2}{\varepsilon}\right)\right)$ gradient queries are required in general.*

2.4 THE NEW PROBLEM CLASS

Theoretical motivation As shown in Propositions 1 and 2, a huge gap exists in the gradient complexities for solving NEPs for monotone general-sum games versus zero-sum games. This motivates the exploration of an intermediate problem class that partially bridges this gap.

Monotone near-zero-sum games We introduce the class of *monotone δ -near-zero-sum games*, which naturally interpolates between monotone zero-sum ($\delta = 0$)² and general-sum ($\delta = L$) games.

Assumption 3 (Near-zero-sum). *Let $\delta \in [0, L]$ such that the function $g(\cdot, \cdot)$ is δ -smooth.*

Definition (MONOTONE NEAR-ZERO-SUM GAMES). *If a two-person general-sum game satisfies Assumptions 1 to 3, we call it a monotone δ -near-zero-sum game.*

3 ALGORITHM AND CONVERGENCE ANALYSIS

We first focus on the algorithm for strongly monotone near-zero-sum games in Sections 3.1 and 3.2. Then, in Section 3.3, we present the results for (non-strongly) monotone near-zero-sum games.

3.1 ALGORITHM

For the zero-sum classes, Lin et al. (2020); Kovalev & Gasnikov (2022); Carmon et al. (2022); Lan & Li (2023) obtained the optimal convergence rate by the Catalyst method (Lin et al., 2018). However, for the more general classes of non-zero-sum games, the application of Catalyst is complicated by the fact that the regularized minimization transforms the problem into a Stackelberg game, whose solution deviates significantly from a Nash equilibrium (see Appendix C for more details). Thus, we are not aware of how the similar smoothing techniques can be applied directly to non-zero-sum games.

This raises the challenge: *can we leverage the off-the-shelf algorithms designed for zero-sum games to solve the non-zero-sum problems of interest?* Now, we introduce our novel algorithm, Iterative Coupling Linearization (ICL), which overcomes the aforementioned challenge and presents a clean black-box framework to solve near-zero-sum games by using zero-sum algorithms as an oracle.

²In 0-near-zero-sum game, let Player 1 maximize $\mathbf{a}_1^\top \mathbf{x} + \mathbf{b}_1^\top \mathbf{y} - h(\mathbf{x}, \mathbf{y})$ and Player 2 maximize $\mathbf{a}_2^\top \mathbf{x} + \mathbf{b}_2^\top \mathbf{y} + h(\mathbf{x}, \mathbf{y})$, respectively. The Nash equilibrium in the above game is the same as that in the following zero-sum game: Player 1 maximizes $\mathbf{a}_1^\top \mathbf{x} - \mathbf{b}_2^\top \mathbf{y} - h(\mathbf{x}, \mathbf{y})$ and Player 2 maximizes $-\mathbf{a}_1^\top \mathbf{x} + \mathbf{b}_2^\top \mathbf{y} + h(\mathbf{x}, \mathbf{y})$.

Algorithm 1 Iterative Coupling Linearization (ICL)**Require:** $\mathbf{z}_0 = (\mathbf{x}_0, \mathbf{y}_0) \in X \times Y$.1: **for** $t = 0, 1, \dots, T-1$ **do**2: Let $\varphi_t(\mathbf{x}, \mathbf{y}) \stackrel{\text{def}}{=} \langle \nabla_{\mathbf{x}} g(\mathbf{x}_t, \mathbf{y}_t), \mathbf{x} \rangle + \frac{1}{2\eta_t} \|\mathbf{x} - \mathbf{x}_t\|^2 + h(\mathbf{x}, \mathbf{y}) - \langle \nabla_{\mathbf{y}} g(\mathbf{x}_t, \mathbf{y}_t), \mathbf{y} \rangle - \frac{1}{2\eta_t} \|\mathbf{y} - \mathbf{y}_t\|^2$. (1)3: Find an inexact solution $\mathbf{z}_{t+1} = (\mathbf{x}_{t+1}, \mathbf{y}_{t+1}) \in X \times Y$ to $\min_{\mathbf{x} \in X} \max_{\mathbf{y} \in Y} \varphi_t(\mathbf{x}, \mathbf{y})$ such that

$$\langle \nabla_{\mathbf{x}} \varphi_t(\mathbf{z}_{t+1}), \mathbf{x}_{t+1} - \mathbf{x} \rangle - \langle \nabla_{\mathbf{y}} \varphi_t(\mathbf{z}_{t+1}), \mathbf{y}_{t+1} - \mathbf{y} \rangle \leq \varepsilon_t, \text{ for all } \mathbf{x} \in X, \mathbf{y} \in Y. \quad (2)$$
4: **end for****Potential function** Our algorithm leverages a natural potential function $\Delta: X \times Y \rightarrow \mathbb{R}$ defined as:

$$\Delta(\mathbf{z}) = \max_{\tilde{\mathbf{z}} = (\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in X \times Y} \underbrace{g(\mathbf{z}) - g(\tilde{\mathbf{z}})}_{\text{jointly convex coupling}} + \underbrace{h(\mathbf{x}, \tilde{\mathbf{y}}) - h(\tilde{\mathbf{x}}, \mathbf{y})}_{\text{convex-concave zero-sum}}, \text{ for all } \mathbf{z} = (\mathbf{x}, \mathbf{y}) \in X \times Y.$$

This potential function decomposes into a jointly convex coupling part and a convex-concave zero-sum part. We show below in Propositions 3 and 4 that minimizing this potential function $\Delta(\cdot)$ is sufficient for finding a Nash equilibrium (with detailed proofs in Appendix D.1):**Proposition 3.** *For any $\mathbf{z} = (\mathbf{x}, \mathbf{y}) \in X \times Y$, we have $\Delta(\mathbf{z}) \geq 0$ and*

$$2\Delta(\mathbf{z}) \geq \max_{\tilde{\mathbf{z}} = (\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in X \times Y} u_1(\tilde{\mathbf{x}}, \mathbf{y}) - u_1(\mathbf{x}, \mathbf{y}) + u_2(\mathbf{x}, \tilde{\mathbf{y}}) - u_2(\mathbf{x}, \mathbf{y}).$$

Proposition 4. *Let $\mathbf{z}^* \in X \times Y$. In monotone games, \mathbf{z}^* is the Nash equilibrium iff $\Delta(\mathbf{z}^*) = 0$.***Algorithm description** Our ICL algorithm solves the strongly monotone near-zero-sum game by iteratively linearizing the coupling part, thereby transforming the non-zero-sum game into a sequence of strongly monotone zero-sum subproblems. The pseudocode is presented in Algorithm 1.Specifically, at every iteration t , we linearize the coupling part in the potential function at \mathbf{z}_t :

$$\min_{\mathbf{z} \in X \times Y} \Delta(\mathbf{z}) \rightsquigarrow \min_{\mathbf{z} \in X \times Y} \max_{\tilde{\mathbf{z}} \in X \times Y} \langle \nabla g(\mathbf{z}_t), \mathbf{z} - \tilde{\mathbf{z}} \rangle + \frac{1}{2\eta_t} \left(\|\mathbf{z} - \mathbf{z}_t\|^2 - \|\tilde{\mathbf{z}} - \mathbf{z}_t\|^2 \right) + h(\mathbf{x}, \tilde{\mathbf{y}}) - h(\tilde{\mathbf{x}}, \mathbf{y}),$$

and note that this minimax optimization can be fully decomposed into two separate problems:

$$\begin{aligned} \min_{\mathbf{x} \in X} \max_{\tilde{\mathbf{y}} \in Y} \langle \nabla g(\mathbf{z}_t), (\mathbf{x}, -\tilde{\mathbf{y}}) \rangle + h(\mathbf{x}, \tilde{\mathbf{y}}) + \frac{1}{2\eta_t} \left(\|\mathbf{x} - \mathbf{x}_t\|^2 - \|\tilde{\mathbf{y}} - \mathbf{y}_t\|^2 \right) &= \varphi_t(\mathbf{x}, \tilde{\mathbf{y}}), \text{ and} \\ \min_{\mathbf{y} \in Y} \max_{\tilde{\mathbf{x}} \in X} \langle \nabla g(\mathbf{z}_t), (-\tilde{\mathbf{x}}, \mathbf{y}) \rangle - h(\tilde{\mathbf{x}}, \mathbf{y}) + \frac{1}{2\eta_t} \left(-\|\tilde{\mathbf{x}} - \mathbf{x}_t\|^2 + \|\mathbf{y} - \mathbf{y}_t\|^2 \right) &= -\varphi_t(\tilde{\mathbf{x}}, \mathbf{y}). \end{aligned}$$

Moreover, by Sion's minimax theorem (Sion, 1958), (after simple substitutions) these two separate problems unify into a *single* saddle point problem of $\min_{\mathbf{x} \in X} \max_{\mathbf{y} \in Y} \varphi_t(\mathbf{x}, \mathbf{y})$, where φ_t is defined in Equation (1). The update $\mathbf{z}_{t+1} = (\mathbf{x}_{t+1}, \mathbf{y}_{t+1})$ is then computed by inexactly solving this *unified* saddle point problem, where the inexactness condition is specified in Equation (2). Our algorithm, thus, provides a clean and black-box reduction from near-zero-sum games to zero-sum games.

3.2 CONVERGENCE ANALYSIS

Throughout Section 3.2, let $\mathbf{z}^* = (\mathbf{x}^*, \mathbf{y}^*)$ be the (unique) Nash equilibrium for the game. We only present the main proof ideas in this section, and the detailed proofs can be found in Appendix D.2. The core of our convergence analysis is to use the properties of the potential function and derive the following descent lemma, based on which we have the outer loop convergence.**Lemma 5** (Descent lemma). *In monotone δ -near-zero-sum games, for $\eta_t \leq \frac{1}{\delta}$, Algorithm 1 ensures*

$$\left(\frac{1}{2\eta_t} + \frac{\min\{\mu, \nu\}}{2} \right) \|\mathbf{z}_{t+1} - \mathbf{z}^*\|^2 \leq \frac{1}{2\eta_t} \|\mathbf{z}_t - \mathbf{z}^*\|^2 + \varepsilon_t.$$

Lemma 6 (Outer loop convergence). *Let $\eta_t = \eta \in (0, \frac{1}{\delta}]$, for all $t \in [0, T-1] \cap \mathbb{Z}$. Denote $\theta = \frac{\min\{\mu, \nu\}}{\eta^{-1} + \min\{\mu, \nu\}}$. Let $\varepsilon_t \leq \frac{\theta\varepsilon}{4\eta}$, for all $t \in [0, T-1] \cap \mathbb{Z}$. For strongly monotone δ -near-zero-sum games, if the outer loop iterate $t \geq \frac{1}{\theta} \log \frac{2D^2}{\varepsilon}$, then Algorithm 1 converges to an ε -accurate Nash equilibrium, that is, $\|\mathbf{z}_t - \mathbf{z}^*\|^2 \leq \varepsilon$.*

For the inner loop, any optimal gradient method for zero-sum games can be used:

Lemma 7 (Inner loop complexity (Kovalev & Gasnikov, 2022; Carmon et al., 2022; Thekumpampil et al., 2022; Lan & Li, 2023)). *Under Assumption 1 with modulus $\mu, \nu > 0$, at each iteration $t \in [0, T-1] \cap \mathbb{Z}$, for $\eta_t \geq \frac{1}{L}$, the inexact solution $(\mathbf{x}_{t+1}, \mathbf{y}_{t+1})$ in Equation (2) of Algorithm 1 can be found with the number of gradient queries bounded by $\mathcal{O}\left(\frac{L}{\sqrt{(\eta_t^{-1} + \mu)(\eta_t^{-1} + \nu)}} \cdot \log\left(\frac{LD^2}{\varepsilon_t}\right)\right)$.*

Combining the outer loop convergence result (Lemma 6) with the inner loop complexity (Lemma 7), we obtain the main theoretical result of this paper, the overall gradient complexity of Algorithm 1.

Theorem 1 (Main theoretical result). *Denote $\eta = \min\left\{\frac{1}{\delta}, \frac{1}{\min\{\mu, \nu\}}\right\}$ and $\theta = \frac{\min\{\mu, \nu\}}{\eta^{-1} + \min\{\mu, \nu\}}$. Let $\eta_t = \eta$ and $\varepsilon_t = \frac{\theta\varepsilon}{4\eta}$, for all $t \in [0, T-1] \cap \mathbb{Z}$. For strongly monotone δ -near-zero-sum games, for $T \geq \frac{1}{\theta} \log \frac{2D^2}{\varepsilon}$, the outer loop iterates of Algorithm 1 converge to an ε -accurate Nash equilibrium with the number of gradient queries bounded by*

$$\mathcal{O}\left(\left(\frac{L}{\sqrt{\mu\nu}} + \frac{L}{\min\{\mu, \nu\}} \cdot \min\left\{1, \sqrt{\frac{\delta}{\mu + \nu}}\right\}\right) \cdot \log\left(\frac{LD^2}{\min\{\mu, \nu\} \cdot \varepsilon}\right) \log\left(\frac{D^2}{\varepsilon}\right)\right).$$

Finally, we highlight the conditions under which Algorithm 1 achieves a faster convergence rate compared to variational inequality methods (Tseng, 1995).

REMARK 1 (Acceleration in strongly monotone near-zero-sum games). *Consider strongly monotone near-zero-sum games where $\min\{\mu, \nu\} + \delta = o(\max\{\mu, \nu\})$. The gradient complexity of Algorithm 1 is $\tilde{\mathcal{O}}\left(\left(\frac{L}{\sqrt{\mu\nu}} + \frac{L}{\min\{\mu, \nu\}} \cdot \sqrt{\frac{\delta}{\mu + \nu}}\right) \cdot \log^2\left(\frac{D^2}{\varepsilon}\right)\right)$, which (ignoring logarithms) improves upon the $\mathcal{O}\left(\frac{L}{\min\{\mu, \nu\}} \cdot \log\left(\frac{D^2}{\varepsilon}\right)\right)$ rate of variational inequality methods as stated in Proposition 1. We also remark that for the special case of zero-sum games ($\delta = 0$), our rate recovers the optimal $\mathcal{O}\left(\frac{L}{\sqrt{\mu\nu}} \cdot \log\left(\frac{D^2}{\varepsilon}\right)\right)$ rate as stated in Proposition 2 up to a logarithm term.*

3.3 ACCELERATION IN NON-STRONGLY MONOTONE NEAR-ZERO-SUM GAMES

We first state the known results for non-strongly monotone games in literature.

Proposition 8 (Nemirovski (2004)). *In monotone general-sum games where $\mu = 0$ or $\nu = 0$, an ε -approximate Nash equilibrium can be found within $\mathcal{O}\left(\frac{LD^2}{\varepsilon}\right)$ gradient queries.*

Then, we provide our result, which is obtained by a similar reduction as in Lin et al. (2020); Wang & Li (2020); Thekumpampil et al. (2022). The proof can be found in Appendix D.3.

Corollary 9. *In monotone δ -near-zero-sum games with modulus μ and ν :*

- (a) *For $\mu > 0$ and $\nu = 0$, an ε -approximate Nash equilibrium can be obtained within $\tilde{\mathcal{O}}\left(\left(\frac{LD_Y}{\sqrt{\mu\varepsilon}} + \frac{LD_Y^2}{\varepsilon}\right) \cdot \min\left\{1, \sqrt{\frac{\delta}{\mu}}\right\}\right) \cdot \log^2\left(\frac{LD^2}{\varepsilon}\right)$ gradient queries;*
- (b) *For $\mu = 0$ and $\nu > 0$, an ε -approximate Nash equilibrium can be obtained within $\tilde{\mathcal{O}}\left(\left(\frac{LD_X}{\sqrt{\nu\varepsilon}} + \frac{LD_X^2}{\varepsilon}\right) \cdot \min\left\{1, \sqrt{\frac{\delta}{\nu}}\right\}\right) \cdot \log^2\left(\frac{LD^2}{\varepsilon}\right)$ gradient queries; and*
- (c) *For $\mu = 0$ and $\nu = 0$, an ε -approximate Nash equilibrium can be obtained within $\tilde{\mathcal{O}}\left(\left(\frac{LD_X D_Y}{\varepsilon} + \frac{LD^2}{\varepsilon}\right) \cdot \min\left\{1, \sqrt{\frac{\delta}{\varepsilon/D_X^2 + \varepsilon/D_Y^2}}\right\}\right) \cdot \log^2\left(\frac{LD^2}{\varepsilon}\right)$ gradient queries.*

REMARK 2 (Acceleration in non-strongly monotone near-zero-sum games). *For non-strongly monotone general-sum games, our rate of finding an ε -approximate Nash equilibrium (ignoring logarithms) is faster than the $\mathcal{O}\left(\frac{LD^2}{\varepsilon}\right)$ rate in literature: (a) when $\delta = o(\mu)$ and $\nu = 0$; (b) when $\mu = 0$ and $\delta = o(\nu)$; or (c) when $\mu = \nu = 0$ and $\varepsilon/D^2 + \delta = o(\varepsilon/D_X^2 + \varepsilon/D_Y^2)$. We also remark that for the special case of zero-sum games ($\delta = 0$), our rates recover the optimal rates of Lin et al. (2020); Wang & Li (2020) up to logarithm terms.*

Technical novelty Our technical contributions include a novel coupling linearization technique and the derivation of a descent lemma (Lemma 5). These elements combine to form a clean, general, and powerful black-box reduction. This reduction’s key advantage is its ability to solve monotone non-zero-sum problems by treating any off-the-shelf zero-sum algorithm as an oracle. Moreover, the inherent black-box nature of our method also allows for deriving complexity results for other oracle settings. For instance, we present some additional results with proximal oracles in Appendix E.

4 APPLICATION EXAMPLES

In this section, we present practical examples of monotone near-zero-sum games. We focus on the application of our approach, while the proof details are presented in Appendix D.4.

4.1 OUR APPROACH FOR REGULARIZED MATRIX GAMES

Regularized matrix games We demonstrate the applicability of Iterative Coupling Linearization to regularized matrix games. Let $X \subseteq \mathbb{R}^n$ and $Y \subseteq \mathbb{R}^m$ be compact convex sets, and let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times n}$ with $\|\mathbf{A}\| \leq L$, $\|\mathbf{B}\| \leq L$, and $\left\|\frac{\mathbf{A}+\mathbf{B}}{2}\right\| \leq \beta$. Let $\mathcal{R}: X \times Y \rightarrow \mathbb{R}$ be an L -smooth regularizer that is μ -strongly concave- ν -strongly convex. Let Player 1 maximize

$$u_1(\mathbf{x}, \mathbf{y}) = \langle \mathbf{A}\mathbf{x}, \mathbf{y} \rangle + \mathcal{R}(\mathbf{x}, \mathbf{y})$$

over $\mathbf{x} \in X$ and Player 2 maximize

$$u_2(\mathbf{x}, \mathbf{y}) = \langle \mathbf{B}\mathbf{x}, \mathbf{y} \rangle - \mathcal{R}(\mathbf{x}, \mathbf{y})$$

over $\mathbf{y} \in Y$. Assume that $\beta \leq \frac{1}{2}\sqrt{\mu\nu}$, then the game is $\min\{\frac{\mu}{2}, \frac{\nu}{2}\}$ -strongly monotone, so classic variational inequality methods yield an ε -accurate Nash equilibrium within $\mathcal{O}\left(\frac{L}{\min\{\mu, \nu\}} \cdot \log\left(\frac{D^2}{\varepsilon}\right)\right)$ gradient queries (Tseng, 1995).

Now, we show that our ICL algorithm can be applied to get a faster rate by leveraging the near-zero-sum structure. Since $-\frac{1}{2}(u_1 + u_2)(\cdot, \cdot)$ is not jointly convex, violating Assumption 2, we first apply a “convex reformulation” technique.

Convex reformulation technique Specifically, we choose the parameters β_1 and β_2 based on the relationship between 2β , μ , and ν : (i) If $2\beta \leq \mu$ and $2\beta \leq \nu$, let $\beta_1 = \beta_2 = \beta$; (ii) if $\mu \leq 2\beta \leq \nu$, let $\beta_1 = \frac{\mu}{2}$ and $\beta_2 = \frac{2\beta^2}{\mu}$; and (iii) if $\nu \leq 2\beta \leq \mu$, let $\beta_1 = \frac{2\beta^2}{\nu}$ and $\beta_2 = \frac{\nu}{2}$. With these choices, we always have $\beta_1 \leq \frac{\mu}{2}$, $\beta_2 \leq \frac{\nu}{2}$, and $\sqrt{\beta_1\beta_2} = \beta$.

We then reformulate the problem as follows: Player 1 maximizes $\tilde{u}_1(\mathbf{x}, \mathbf{y}) = u_1(\mathbf{x}, \mathbf{y}) - \beta_2 \|\mathbf{y}\|^2$ over $\mathbf{x} \in X$, and Player 2 maximizes $\tilde{u}_2(\mathbf{x}, \mathbf{y}) = u_2(\mathbf{x}, \mathbf{y}) - \beta_1 \|\mathbf{x}\|^2$ over $\mathbf{y} \in Y$. This reformulated NEP has the same Nash equilibrium as the original. Let

$$\tilde{g}(\mathbf{x}, \mathbf{y}) = -\left\langle \left(\frac{\mathbf{A} + \mathbf{B}}{2}\right) \mathbf{x}, \mathbf{y} \right\rangle + \left(\frac{\beta_1}{2} \|\mathbf{x}\|^2 + \frac{\beta_2}{2} \|\mathbf{y}\|^2\right)$$

and

$$\tilde{h}(\mathbf{x}, \mathbf{y}) = -\left\langle \left(\frac{\mathbf{A} - \mathbf{B}}{2}\right) \mathbf{x}, \mathbf{y} \right\rangle - \left(\mathcal{R}(\mathbf{x}, \mathbf{y}) + \frac{\beta_1}{2} \|\mathbf{x}\|^2 - \frac{\beta_2}{2} \|\mathbf{y}\|^2\right).$$

Then, $\tilde{u}_1 = -\tilde{g} - \tilde{h}$ and $\tilde{u}_2 = -\tilde{g} + \tilde{h}$. Since $\beta \leq \frac{1}{2}\sqrt{\mu\nu}$, by Cauchy-Schwartz inequality, $\tilde{g}(\cdot, \cdot)$ is jointly convex. Further, by the choices of β_1 and β_2 , we have $\tilde{g}(\cdot, \cdot)$ is $(\beta + \max\{\beta_1, \beta_2\})$ -smooth, and $\tilde{h}(\cdot, \cdot)$ is $\frac{\mu}{2}$ -strongly convex- $\frac{\nu}{2}$ -strongly concave.

Our approach applied to reformulated games Now applying Algorithm 1 to the reformulated NEP, by Theorem 1, we obtain an ε -accurate Nash equilibrium with the number of gradient queries bounded by

$$\tilde{\mathcal{O}}\left(\left(\frac{L}{\sqrt{\mu\nu}} + \frac{L}{\min\{\mu, \nu\}} \cdot \frac{\beta}{\sqrt{\mu\nu}}\right) \cdot \log^2\left(\frac{D^2}{\varepsilon}\right)\right).$$

When $\min\{\mu, \nu\} + \beta = o\left(\frac{1}{2}\sqrt{\mu\nu}\right)$, this rate surpasses the best-known $\mathcal{O}\left(\frac{L}{\min\{\mu, \nu\}} \cdot \log\left(\frac{D^2}{\varepsilon}\right)\right)$ gradient complexity of variational inequality methods (Tseng, 1995). This acceleration leveraging the near-zero-sum structure seems to be a new result even in the well-studied context of matrix games.

EXAMPLE 1 (Matrix games with transaction fees). *Consider regularized matrix games with transaction fees. Let $X = \mathcal{P}_n \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathbb{R}_{\geq 0}^n \mid x_1 + \dots + x_n = 1\}$ and $Y = \mathcal{P}_m \stackrel{\text{def}}{=} \{\mathbf{y} \in \mathbb{R}_{\geq 0}^m \mid y_1 + \dots + y_m = 1\}$. Let $\mathbf{M} \in \mathbb{R}^{m \times n}$ be the payoff matrix of Player 1 without transaction fee, with $-\mathbf{M}$ as the payoff matrix of Player 2 without transaction fee. Assume $\|\mathbf{M}\| \leq L$. Denote $\mathbf{M}_+ \stackrel{\text{def}}{=} \frac{1}{2}(\mathbf{M} + \text{abs}(\mathbf{M}))$ and $\mathbf{M}_- \stackrel{\text{def}}{=} \frac{1}{2}(-\mathbf{M} + \text{abs}(\mathbf{M}))$.*³

Now, suppose there is a transaction fee of $\rho \in [0, 1]$ charged by some third party on every payment. Then, the payoff matrices of Player 1 and Player 2 with transaction fees are

$$\mathbf{A} = (1 - \rho)\mathbf{M}_+ - \mathbf{M}_- \quad \text{and} \quad \mathbf{B} = -\mathbf{M}_+ + (1 - \rho)\mathbf{M}_-.$$

Let $\mathcal{R}: X \times Y \rightarrow \mathbb{R}$ be an L -smooth regularizer that is μ -strongly concave- ν -strongly convex. Assume that $\rho \|\text{abs}(\mathbf{M})\| = o\left(\sqrt{\mu\nu}\right)$.

Let Player 1 maximize $u_1(\mathbf{x}, \mathbf{y}) = \langle \mathbf{A}\mathbf{x}, \mathbf{y} \rangle + \mathcal{R}(\mathbf{x}, \mathbf{y})$ over $\mathbf{x} \in X$, and Player 2 maximize $u_2(\mathbf{x}, \mathbf{y}) = \langle \mathbf{B}\mathbf{x}, \mathbf{y} \rangle - \mathcal{R}(\mathbf{x}, \mathbf{y})$ over $\mathbf{y} \in Y$. Applying the convex reformulation and then Algorithm 1, we obtain an ε -accurate Nash equilibrium with the number of gradient queries bounded by

$$\tilde{\mathcal{O}}\left(\left(\frac{L}{\sqrt{\mu\nu}} + \frac{L}{\min\{\mu, \nu\}} \cdot \frac{\rho \|\text{abs}(\mathbf{M})\|}{\sqrt{\mu\nu}}\right) \cdot \log^2\left(\frac{D^2}{\varepsilon}\right)\right).$$

Matrix games with transaction fees are analogous to many practical scenarios, for instance, (a) betting sites often set odds such that there is overall no loss-gain (zero-sum game) but charge a small transaction fee; or (b) negotiation between two clients via a trusted-third party that charges a brokerage fee is another example of a case with small transaction fee.

4.2 OUR APPROACH FOR COMPETITIVE GAMES WITH SMALL ADDITIONAL INCENTIVES

We show the applicability of Iterative Coupling Linearization to competitive games with small additional incentives. Let X and Y be compact convex sets in Euclidean spaces. Let $h: X \times Y \rightarrow \mathbb{R}$ be the competition payoff function, which is L -smooth and μ -strongly convex- ν -strongly concave. Let $g: X \times Y \rightarrow \mathbb{R}$ be the additional incentive function, which is β -smooth with $\beta \leq L$. Let Player 1 maximize $u_1 = -g - h$ over $\mathbf{x} \in X$, and Player 2 maximize $u_2 = -g + h$ over $\mathbf{y} \in Y$.

We explore two scenarios where the games are $\min\{\frac{\mu}{2}, \frac{\nu}{2}\}$ -strongly monotone, to which our ICL algorithm as well as the classic variational inequalities (Tseng, 1995) can be applied:

1. If $g(\cdot, \cdot)$ is jointly convex and $\beta = o(\max\{\mu, \nu\})$, applying Algorithm 1 directly yields an ε -accurate Nash equilibrium with the number of gradient queries bounded by

$$\tilde{\mathcal{O}}\left(\left(\frac{L}{\sqrt{\mu\nu}} + \frac{L}{\min\{\mu, \nu\}} \cdot \sqrt{\frac{\beta}{\mu + \nu}}\right) \cdot \log^2\left(\frac{D^2}{\varepsilon}\right)\right).$$

2. If $\beta = o\left(\frac{1}{2}\min\{\mu, \nu\}\right)$, we first apply the “convex reformulation” technique. We reformulate the problem as follows: Player 1 maximizes $\tilde{u}_1(\mathbf{x}, \mathbf{y}) = u_1(\mathbf{x}, \mathbf{y}) - \beta \|\mathbf{y}\|^2$ over $\mathbf{x} \in X$, and Player 2 maximizes $\tilde{u}_2(\mathbf{x}, \mathbf{y}) = u_2(\mathbf{x}, \mathbf{y}) - \beta \|\mathbf{x}\|^2$ over $\mathbf{y} \in Y$. This reformulated NEP has the same Nash equilibrium as the original. Let $\tilde{g} = -\frac{1}{2}(\tilde{u}_1 + \tilde{u}_2)$ and $\tilde{h} = \frac{1}{2}(-\tilde{u}_1 + \tilde{u}_2)$. Then, $\tilde{h}(\cdot, \cdot)$ is $\frac{\mu}{2}$ -strongly convex- $\frac{\nu}{2}$ -strongly concave, and $\tilde{g}(\cdot, \cdot)$ is jointly convex and 2β -smooth.

³ $\text{abs}(\mathbf{M})$ represents a matrix of the same dimensions as \mathbf{M} where each element is the absolute value of the corresponding element in \mathbf{M} . A more concrete illustration is given in Appendix F.

Applying Algorithm 1 to the reformulated NEP, we obtain an ε -accurate Nash equilibrium with the number of gradient queries bounded by

$$\tilde{\mathcal{O}}\left(\frac{L}{\sqrt{\mu\nu}} \cdot \log^2\left(\frac{D^2}{\varepsilon}\right)\right).$$

In both scenarios 1 and 2, our gradient queries are fewer than the $\mathcal{O}\left(\frac{L}{\min\{\mu, \nu\}} \cdot \log\left(\frac{D^2}{\varepsilon}\right)\right)$ gradient queries of the classic variational inequality methods (Tseng, 1995).

EXAMPLE 2 (Competitive games with small cooperation incentives). *Consider the games where cooperation coexists with competition. Let $X \subseteq X_a \times X_b$ and $Y \subseteq Y_a \times Y_b$ be compact convex sets in Euclidean spaces. For $\mathbf{x} = (x_a, x_b) \in X$, $x_a \in X_a$ represents Player 1's effort in cooperation, and $x_b \in X_b$ represents Player 1's effort in competition (and similarly for $\mathbf{y} = (y_a, y_b) \in Y$). Let $f_a: X_a \times Y_a \rightarrow \mathbb{R}$ be the cooperation incentive function given by*

$$f_a(\mathbf{x}_a, \mathbf{y}_a) = \mathcal{R}_1(\mathbf{x}_a) + \tilde{g}(\mathbf{x}_a, \mathbf{y}_a) + \mathcal{R}_2(\mathbf{y}_a),$$

where the regularizer $\mathcal{R}_1: X_a \rightarrow \mathbb{R}$ is μ -strongly convex and L -smooth, the function $\tilde{g}: X_a \times Y_a \rightarrow \mathbb{R}$ is jointly convex and β -smooth, and the regularizer $\mathcal{R}_2: Y_a \rightarrow \mathbb{R}$ is ν -strongly convex and L -smooth. Let $f_b: X_b \times Y_b \rightarrow \mathbb{R}$ be the competition payoff function, which is L -smooth and μ -strongly convex- ν -strongly concave. Assume that $\beta = o(\max\{\mu, \nu\})$.

Let Player 1 maximize $u_1(\mathbf{x}, \mathbf{y}) = -f_a(\mathbf{x}_a, \mathbf{y}_a) - f_b(\mathbf{x}_b, \mathbf{y}_b)$ over $\mathbf{x} \in X$, and Player 2 maximize $u_2(\mathbf{x}, \mathbf{y}) = -f_a(\mathbf{x}_a, \mathbf{y}_a) + f_b(\mathbf{x}_b, \mathbf{y}_b)$ over $\mathbf{y} \in Y$. Denoting $\tilde{h}(\mathbf{x}, \mathbf{y}) = \mathcal{R}_1(\mathbf{x}_a) + f_b(\mathbf{x}_b, \mathbf{y}_b) - \mathcal{R}_2(\mathbf{y}_a)$, the NEP can be reformulated as Player 1 maximizing $\tilde{u}_1(\mathbf{x}, \mathbf{y}) = -\tilde{g}(\mathbf{x}_a, \mathbf{y}_a) - \tilde{h}(\mathbf{x}, \mathbf{y})$ over $\mathbf{x} \in X$, and Player 2 maximizing $\tilde{u}_2(\mathbf{x}, \mathbf{y}) = -\tilde{g}(\mathbf{x}_a, \mathbf{y}_a) + \tilde{h}(\mathbf{x}, \mathbf{y})$ over $\mathbf{y} \in Y$. Applying Algorithm 1 as detailed above, we obtain an ε -accurate Nash equilibrium with the number of gradient queries bounded by

$$\tilde{\mathcal{O}}\left(\left(\frac{L}{\sqrt{\mu\nu}} + \frac{L}{\min\{\mu, \nu\}} \cdot \sqrt{\frac{\beta}{\mu + \nu}}\right) \cdot \log^2\left(\frac{D^2}{\varepsilon}\right)\right).$$

As a final remark, the modeling of the coexistence of competition and cooperation has been well-researched (for instance, see the studies of Nash (1950; 1953); Selten (1960); Raiffa (1952); Kalai & Rosenthal (1978); Kalai & Kalai (2013); Halpern & Rong (2013) on semi-cooperative games). Indeed, these theories are often applied to the scenarios where cooperation dominates, and optimization techniques have been used to accelerate the dominant cooperation part (Chen et al., 2017). Our work contributes to this line of research on semi-cooperative games where competition dominates, yet there is a small cooperation incentive. For example, repeated Prisoners' dilemma with a stochastic number of rounds with small benefit-to-cost ratio is an example of competition with small cooperation incentive (as benefit-to-cost ratio b/c is small) (Nowak, 2006a; b; Sigmund, 2010). We give a more concrete illustration of donation games in Appendix F.

5 BASIC NUMERICAL EXPERIMENTS

We conducted basic numerical experiments to validate our theoretical results, focusing on matrix games with transaction fees as in Example 1. We set $n = m = 10000$, $\mu = 10^{-4}$, $\nu = 1$, and $\varepsilon = 10^{-7}$. A sparse, random matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$ such that $\|\mathbf{M}\| = 1$ was generated. The regularizer was defined as $\mathcal{R}(\mathbf{x}, \mathbf{y}) = -\frac{\mu}{2} \|\mathbf{x}\|^2 + \frac{\nu}{2} \|\mathbf{y}\|^2$. We varied the transaction fee ρ from $\{0.00\%, 0.03\%, \dots, 0.18\%\}$. Our implementation of ICL (Algorithm 1), detailed in Section 4.1, used the Lifted Primal-Dual method (Thekumparampil et al., 2022) for the inner loop. We compared ICL against the Optimistic Gradient Descent Ascent (OGDA) (Popov, 1980) and Extra-Gradient (EG) (Korpelevich, 1976) methods for variational inequalities. More details and additional experiments are provided in Appendix G, and our code can be found here: https://github.com/riekenluo/Monotone_Near_Zero_Sum_Games.

The numerical results, summarized in Table 1, demonstrate that: (i) our ICL algorithm converges faster when the game is closer to a zero-sum one; (ii) the classic variational inequality methods (EG or OGDA) do not benefit from the near-zero-sum structure; and (iii) our ICL algorithm is faster than the classic variational inequality methods when the game is sufficiently near-zero-sum. In particular,

Table 1: Gradient query counts (in thousands) to converge to an ε -accurate Nash equilibrium under various transaction fees. Error bars indicate 2-sigma variations across 10 independent runs.

Methods \ Transaction Fee ρ	0.00%	0.03%	0.06%	0.09%	0.12%	0.15%	0.18%
ICL (Algorithm 1)	9.1 ± 0.0	22.6 ± 0.4	42.2 ± 0.3	65.0 ± 0.3	75.7 ± 0.3	113.7 ± 0.7	123.8 ± 0.6
OGDA (Popov, 1980)	93.9 ± 0.5	93.9 ± 0.5	93.9 ± 0.5	93.9 ± 0.5	93.9 ± 0.5	94.0 ± 0.6	94.0 ± 0.6
EG (Korpelevich, 1976)	132.9 ± 0.8	132.9 ± 0.8	132.9 ± 0.8	132.9 ± 0.8	132.9 ± 0.8	132.9 ± 0.8	132.9 ± 0.8

in the experiments, ICL requires fewer gradient queries to converge to an ε -accurate Nash equilibrium when the transaction fee $\rho \leq 0.12\%$. This empirical observation aligns with our theoretical prediction in Example 1, which suggests that ICL converges faster when $\rho \|\text{abs}(\mathbf{M})\| \ll \sqrt{\mu\nu} = 1\%$.

6 CONCLUSIONS, LIMITATIONS, AND FUTURE WORK

In this work we consider the class of monotone games and present a condition that naturally interpolates between the zero-sum and a non-zero-sum class. We develop an efficient gradient-based approach and show its applicability with several examples motivated from the literature.

There are some limitations of our work: (a) in our complexity there is a $\log^2(\frac{D^2}{\varepsilon})$ dependency rather than a single logarithm dependency, and whether this double logarithm dependency can be removed is an interesting question; and (b) whether lower-bound results can be obtained for the new class also remains a challenging question and may involve difficult construction of non-quadratic functions, especially given that the lower bound for the general-sum classes remains widely open relative to the long-standing upper bound established in Tseng (1995); Nemirovski (2004).

The current work focuses on two-player monotone non-zero-sum games, which constitutes a crucial first step in generalizing two-player monotone zero-sum games with general conditioning. A natural, though highly non-trivial, extension involves generalizing our fast convergence rates to the multiplayer near-zero-sum setting. We notice that, as a preliminary step, the fast rate in zero-sum games (Proposition 2) have not yet been extended to multiplayer settings. Indeed, this preliminary step is conceptually challenging, as the zero-sum condition in multiplayer games does not imply the strict competition found in two-player games. Furthermore, classic reductions in von Neumann & Morgenstern (1947) show that an n -player general-sum game can be reduced to an $(n+1)$ -player zero-sum game by letting the $(n+1)$ th player take the negative of the summation of the first n players. This implies that three-player zero-sum games are inherently no easier than two-player general-sum games. These results suggest that achieving a fast convergence rate in the multiplayer setting will likely require significantly stronger structural assumptions than the simple summation of utilities to zero. Hence, we leave the full generalization to the multiplayer setting for future research.

In addition to the above theoretical aspects, there are several other interesting directions as well: for example, (a) exploring other applications of regularized matrix games with near-zero-sum payoff matrices is an interesting direction; and (b) in the research of semi-cooperative games where competition dominates, applying our methods in more practical examples is another fruitful direction for future research.

ACKNOWLEDGMENTS

The authors thank for the helpful discussion with Anton Rodomanov during the initial preparation of this project.

RL and KC acknowledge the support of ERC CoG 863818 (ForM-SMArt) and Austrian Science Fund (FWF) 10.55776/COE12.

Ethics statement

This is a theoretical research without any ethics concern. All the experiments are based on randomly generated data. The only LLM usage is to aid or polish writing.

Reproducibility statement

We include the detailed proofs of all the results in the Appendix. We include the source codes of all the experiments, including the scripts for the generation of random data. The codes are based on standard Python libraries (like `numpy` or `scipy`) and are easy to run.

REFERENCES

Bharat Adsul, Jugal Garg, Ruta Mehta, Milind Sohoni, and Bernhard von Stengel. Fast algorithms for rank-1 bimatrix games. *Operations Research*, 69(2):613–631, 2021.

Radu Ioan Boț, Ernö Robert Csetnek, and Michael Sedlmayer. An accelerated minimax algorithm for convex-concave saddle point problems with nonsmooth coupling function. *Computational Optimization and Applications*, 86(3):925–966, 2023.

Ozan Candogan, Asuman Ozdaglar, and Pablo A Parrilo. Dynamics in near-potential games. *Games and Economic Behavior*, 82:66–90, 2013.

Yair Carmon, Arun Jambulapati, Yujia Jin, and Aaron Sidford. Recapp: Crafting a more efficient catalyst for convex optimization. In *International Conference on Machine Learning*, pp. 2658–2685. PMLR, 2022.

Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.

Yunmei Chen, Guanghui Lan, and Yuyuan Ouyang. Optimal primal-dual methods for a class of saddle point problems. *SIAM Journal on Optimization*, 24(4):1779–1814, 2014.

Yunmei Chen, Guanghui Lan, and Yuyuan Ouyang. Accelerated schemes for a class of variational inequalities. *Mathematical Programming*, 165:113–149, 2017.

Joseph Halpern and Nan Rong. Cooperative equilibrium: A solution predicting cooperative play. In *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, pp. 319–326, 2013.

Aamal Hussain, Francesco Belardinelli, and Georgios Piliouras. Beyond strict competition: approximate convergence of multi-agent q-learning dynamics. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI '23*, 2023. ISBN 978-1-956792-03-4. doi: 10.24963/ijcai.2023/16. URL <https://doi.org/10.24963/ijcai.2023/16>.

Sung-Ha Hwang and Luc Rey-Bellet. Strategic decompositions of normal form games: Zero-sum games and potential games. *Games and Economic Behavior*, 122:370–390, 2020.

Adam Kalai and Ehud Kalai. Cooperation in strategic games revisited. *The Quarterly Journal of Economics*, 128(2):917–966, 2013.

Ehud Kalai and Robert W Rosenthal. Arbitration of two-party disputes under ignorance. *International Journal of Game Theory*, 7:65–72, 1978.

Ravi Kannan and Thorsten Theobald. Games of fixed rank: A hierarchy of bimatrix games. *Economic Theory*, 42:157–173, 2010.

Vladimir Kolmogorov and Thomas Pock. One-sided frank-wolfe algorithms for saddle problems. In *International Conference on Machine Learning*, pp. 5665–5675. PMLR, 2021.

Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.

Dmitry Kovalev and Alexander Gasnikov. The first optimal algorithm for smooth and strongly-convex-strongly-concave minimax optimization. *Advances in Neural Information Processing Systems*, 35:14691–14703, 2022.

Guanghui Lan and Yan Li. A novel catalyst scheme for stochastic minimax optimization. *arXiv preprint arXiv:2311.02814*, 2023.

Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. Catalyst acceleration for first-order convex optimization: from theory to practice. *Journal of Machine Learning Research*, 18(212):1–54, 2018.

Tianyi Lin, Chi Jin, and Michael I Jordan. Near-optimal algorithms for minimax optimization. In *Conference on Learning Theory*, pp. 2738–2779. PMLR, 2020.

Tianyi Lin, Chi Jin, and Michael I Jordan. Two-timescale gradient descent ascent algorithms for nonconvex minimax optimization. *Journal of Machine Learning Research*, 26(11):1–45, 2025.

Ruta Mehta. Constant rank two-player games are ppad-hard. *SIAM Journal on Computing*, 47(5):1858–1887, 2018.

Aryan Mokhtari, Asuman E Ozdaglar, and Sarath Pattathil. Convergence rate of $o(1/k)$ for optimistic gradient and extragradient methods in smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 30(4):3230–3251, 2020.

John Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2), 1951.

John Nash. Two-person cooperative games. *Econometrica: Journal of the Econometric Society*, pp. 128–140, 1953.

Arkadi Nemirovski. Information-based complexity of convex programming. *Lecture notes*, 834, 1995.

Arkadi Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.

A. S. Nemirovsky and D. B. Yudin. *Problem complexity and method efficiency in optimization*. Wiley-Interscience, 1983.

Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103:127–152, 2005.

Yurii E. Nesterov. *Introductory Lectures on Convex Optimization - A Basic Course*, volume 87 of *Applied Optimization*. Springer, 2004. ISBN 978-1-4613-4691-3. doi: 10.1007/978-1-4419-8853-9. URL <https://doi.org/10.1007/978-1-4419-8853-9>.

Martin A Nowak. *Evolutionary dynamics: exploring the equations of life*. Harvard university press, 2006a.

Martin A Nowak. Five rules for the evolution of cooperation. *science*, 314(5805):1560–1563, 2006b.

Leonid Denisovich Popov. A modification of the Arrow-Hurwitz method of search for saddle points. *Mat. Zametki*, 28(5):777–784, 1980.

Howard Raiffa. *Arbitration schemes for generalized two-person games*. University of Michigan, 1952.

R Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.

J Ben Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, pp. 520–534, 1965.

Reinhard Selten. Bewertung strategischer spiele. *Zeitschrift für die gesamte Staatswissenschaft/Journal of Institutional and Theoretical Economics*, H(2):221–282, 1960.

Karl Sigmund. *The Calculus of Selfishness*. Princeton University Press, Princeton, NJ, 2010. ISBN 978-0-691-14275-5.

Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 03 1958. ISSN 0030-8730. doi: 10.2140/pjm.1958.8.171. URL <https://cir.nii.ac.jp/crid/1362825894260931328>.

John Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.

Kiran K. Thekumparampil, Niao He, and Sewoong Oh. Lifted primal-dual method for bilinearly coupled smooth minimax optimization. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera (eds.), *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 4281–4308. PMLR, 28–30 Mar 2022.

Paul Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60(1-2):237–252, 1995.

John von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton university press, 1947.

Yuanhao Wang and Jian Li. Improved algorithms for convex-concave minimax optimization. *Advances in Neural Information Processing Systems*, 33:4800–4810, 2020.

Jörgen W Weibull. *Evolutionary game theory*. MIT press, 1997.

Junchi Yang, Siqi Zhang, Negar Kiyavash, and Niao He. A catalyst framework for minimax optimization. *Advances in Neural Information Processing Systems*, 33:5667–5678, 2020.

Georgios N Yannakakis and Julian Togelius. *Artificial intelligence and games*, volume 2. Springer, 2018.

Junyu Zhang, Mingyi Hong, and Shuzhong Zhang. On lower iteration complexity bounds for the convex concave saddle point problems. *Mathematical Programming*, 194(1):901–935, 2022.

A RELATED WORK

We discuss other classes of games in the literature that bridge the gap between zero-sum and general-sum games.

Near network zero-sum games Near network zero-sum games (Hussain et al., 2023) define a class of games that is close to network zero-sum games in terms of maximum pairwise difference (Candogan et al., 2013; Hussain et al., 2023). Limited to the setting of two-person games, monotone near-zero-sum games considered in this paper differ in three aspects: (i) The utility functions in this paper can be general functions, rather than bilinear functions; (ii) the difference between near-zero-sum games and zero-sum games in this paper is characterized by (higher-order) smoothness parameter, rather than by function values; and (iii) the solution of near network zero-sum games is taken directly from the zero-sum case, which only guarantees convergence to a neighborhood of the Nash equilibrium.

Rank- k games In the setting of matrix games, one of the most significant attempts on bridging the gap between zero-sum and non-zero-sum games is the class of Rank- k games introduced in Kannan & Theobald (2010). As a generalization of zero-sum matrix games, Kannan & Theobald (2010) study matrix games where $\text{rank}(\mathbf{A} + \mathbf{B}) = k$, where \mathbf{A} and \mathbf{B} are the payoff matrices of the two players. To find an approximate Nash equilibrium, an FPTAS exists when k is small (Kannan & Theobald, 2010); to find an exact Nash equilibrium, Rank-1 games can be solved in polynomial time (Adsul et al., 2021), while Rank-3 games are already PPAD-hard (Mehta, 2018). It is crucial to emphasize that monotone near-zero-sum games, as considered in this paper, are fundamentally distinct from Rank- k games. Specifically: (i) The utility functions in this paper can be general functions, rather than bilinear functions; (ii) matrix games can be sufficiently near-zero-sum but still have full rank; and (iii) the focus of this paper is on gradient-based algorithms and complexity within the Nemirovsky-Yudin optimization model (Nemirovsky & Yudin, 1983), while the study of Rank- k games focuses on algorithms and complexity on Turing machines.

B RELATIONS TO APPROXIMATE NASH EQUILIBRIUM

Indeed, an approximate Nash equilibrium can be obtained from an accurate Nash equilibrium (Nemirovski, 2004). Below, we include this result for self-consistency.

Proposition 10 (Nemirovski (2004)). *In a monotone general-sum game, let $\mathbf{z}^* = (\mathbf{x}^*, \mathbf{y}^*)$ be the Nash equilibrium. Let $\bar{\mathbf{z}} = (\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in X \times Y$ and $\gamma \in (0, \frac{1}{\sqrt{2}L}]$. We have*

$$\begin{aligned} & \max_{\mathbf{x} \in X, \mathbf{y} \in Y} u_1(\mathbf{x}, \hat{\mathbf{y}}) - u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}) + u_2(\hat{\mathbf{x}}, \mathbf{y}) - u_2(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \\ & \leq \max_{\mathbf{x} \in X, \mathbf{y} \in Y} \langle \mathcal{F}(\hat{\mathbf{x}}, \hat{\mathbf{y}}), (\hat{\mathbf{x}}, \hat{\mathbf{y}}) - (\mathbf{x}, \mathbf{y}) \rangle \\ & \leq \frac{2}{\gamma} \sqrt{D_X^2 + D_Y^2} \|\bar{\mathbf{z}} - \mathbf{z}^*\|, \end{aligned}$$

where $\hat{\mathbf{x}} \stackrel{\text{def}}{=} \Pi_X(\bar{\mathbf{x}} + \gamma \nabla_{\mathbf{x}} u_1(\bar{\mathbf{x}}, \bar{\mathbf{y}}))$ and $\hat{\mathbf{y}} \stackrel{\text{def}}{=} \Pi_Y(\bar{\mathbf{y}} + \gamma \nabla_{\mathbf{y}} u_2(\bar{\mathbf{x}}, \bar{\mathbf{y}}))$.⁴

Proof. Denote $\mathbf{x}_+ \stackrel{\text{def}}{=} \Pi_X(\bar{\mathbf{x}} + \gamma \nabla_{\mathbf{x}} u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}))$, $\mathbf{y}_+ \stackrel{\text{def}}{=} \Pi_Y(\bar{\mathbf{y}} + \gamma \nabla_{\mathbf{y}} u_2(\hat{\mathbf{x}}, \hat{\mathbf{y}}))$, and $\mathbf{z}_+ \stackrel{\text{def}}{=} (\mathbf{x}_+, \mathbf{y}_+)$. Consider any $\tilde{\mathbf{x}} \in X$ and $\tilde{\mathbf{y}} \in Y$. By the assignment of $\hat{\mathbf{x}}$, we have

$$\langle \nabla_{\mathbf{x}} u_1(\bar{\mathbf{x}}, \bar{\mathbf{y}}), \hat{\mathbf{x}} \rangle - \frac{1}{2\gamma} \|\hat{\mathbf{x}} - \bar{\mathbf{x}}\|^2 \geq \langle \nabla_{\mathbf{x}} u_1(\bar{\mathbf{x}}, \bar{\mathbf{y}}), \mathbf{x}_+ \rangle - \frac{1}{2\gamma} \|\mathbf{x}_+ - \bar{\mathbf{x}}\|^2 + \frac{1}{2\gamma} \|\mathbf{x}_+ - \hat{\mathbf{x}}\|^2. \quad (3)$$

By the assignment of \mathbf{x}_+ , we have

$$\langle \nabla_{\mathbf{x}} u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}), \mathbf{x}_+ \rangle - \frac{1}{2\gamma} \|\mathbf{x}_+ - \bar{\mathbf{x}}\|^2 \geq \langle \nabla_{\mathbf{x}} u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}), \tilde{\mathbf{x}} \rangle - \frac{1}{2\gamma} \|\tilde{\mathbf{x}} - \bar{\mathbf{x}}\|^2 + \frac{1}{2\gamma} \|\mathbf{x}_+ - \tilde{\mathbf{x}}\|^2. \quad (4)$$

⁴In a Euclidean space \mathcal{Q} , for a non-empty, closed, and convex set $Q \subseteq \mathcal{Q}$ and a vertex $\mathbf{u} \in \mathcal{Q}$, let $\Pi_Q(\mathbf{u})$ denote the projection of \mathbf{u} onto Q , that is, $\Pi_Q(\mathbf{u}) \stackrel{\text{def}}{=} \arg \min_{\mathbf{v} \in Q} \|\mathbf{u} - \mathbf{v}\|$.

In view of

$$\begin{aligned}
& \langle \nabla_{\mathbf{x}} u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}), \tilde{\mathbf{x}} - \hat{\mathbf{x}} \rangle \\
&= \langle \nabla_{\mathbf{x}} u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}), \mathbf{x}_+ - \hat{\mathbf{x}} \rangle + \langle \nabla_{\mathbf{x}} u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}), \tilde{\mathbf{x}} - \mathbf{x}_+ \rangle \\
&\stackrel{(4)}{\leq} \langle \nabla_{\mathbf{x}} u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}), \mathbf{x}_+ - \hat{\mathbf{x}} \rangle - \frac{1}{2\gamma} \|\mathbf{x}_+ - \bar{\mathbf{x}}\|^2 + \frac{1}{2\gamma} \|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\|^2 - \frac{1}{2\gamma} \|\mathbf{x}_+ - \tilde{\mathbf{x}}\|^2 \\
&= \langle \nabla_{\mathbf{x}} u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}) - \nabla_{\mathbf{x}} u_1(\bar{\mathbf{x}}, \bar{\mathbf{y}}), \mathbf{x}_+ - \hat{\mathbf{x}} \rangle + \langle \nabla_{\mathbf{x}} u_1(\bar{\mathbf{x}}, \bar{\mathbf{y}}), \mathbf{x}_+ - \hat{\mathbf{x}} \rangle - \frac{1}{2\gamma} \|\mathbf{x}_+ - \bar{\mathbf{x}}\|^2 \\
&\quad + \frac{1}{2\gamma} \|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\|^2 - \frac{1}{2\gamma} \|\mathbf{x}_+ - \tilde{\mathbf{x}}\|^2 \\
&\stackrel{(3)}{\leq} \langle \nabla_{\mathbf{x}} u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}) - \nabla_{\mathbf{x}} u_1(\bar{\mathbf{x}}, \bar{\mathbf{y}}), \mathbf{x}_+ - \hat{\mathbf{x}} \rangle - \frac{1}{2\gamma} \|\hat{\mathbf{x}} - \bar{\mathbf{x}}\|^2 - \frac{1}{2\gamma} \|\mathbf{x}_+ - \hat{\mathbf{x}}\|^2 \\
&\quad + \frac{1}{2\gamma} \|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\|^2 - \frac{1}{2\gamma} \|\mathbf{x}_+ - \tilde{\mathbf{x}}\|^2 \\
&\leq L \|(\hat{\mathbf{x}}, \hat{\mathbf{y}}) - \bar{\mathbf{z}}\| \cdot \|\mathbf{x}_+ - \hat{\mathbf{x}}\| - \frac{1}{2\gamma} \|\hat{\mathbf{x}} - \bar{\mathbf{x}}\|^2 - \frac{1}{2\gamma} \|\mathbf{x}_+ - \hat{\mathbf{x}}\|^2 \\
&\quad + \frac{1}{2\gamma} \|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\|^2 - \frac{1}{2\gamma} \|\mathbf{x}_+ - \tilde{\mathbf{x}}\|^2 \\
&\leq \frac{L}{2\sqrt{2}} \|(\hat{\mathbf{x}}, \hat{\mathbf{y}}) - \bar{\mathbf{z}}\|^2 - \frac{1}{2\gamma} \|\hat{\mathbf{x}} - \bar{\mathbf{x}}\|^2 + \frac{1}{2\gamma} \|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\|^2 - \frac{1}{2\gamma} \|\mathbf{x}_+ - \tilde{\mathbf{x}}\|^2
\end{aligned} \tag{5}$$

(where we have used $\gamma \leq \frac{1}{\sqrt{2}L}$ in the last inequality), and similarly,

$$\langle \nabla_{\mathbf{y}} u_2(\hat{\mathbf{x}}, \hat{\mathbf{y}}), \tilde{\mathbf{y}} - \hat{\mathbf{y}} \rangle \leq \frac{L}{2\sqrt{2}} \|(\hat{\mathbf{x}}, \hat{\mathbf{y}}) - \bar{\mathbf{z}}\|^2 - \frac{1}{2\gamma} \|\hat{\mathbf{y}} - \bar{\mathbf{y}}\|^2 + \frac{1}{2\gamma} \|\bar{\mathbf{y}} - \tilde{\mathbf{y}}\|^2 - \frac{1}{2\gamma} \|\mathbf{y}_+ - \tilde{\mathbf{y}}\|^2, \tag{6}$$

we have

$$\begin{aligned}
& \langle \mathcal{F}(\hat{\mathbf{x}}, \hat{\mathbf{y}}), (\hat{\mathbf{x}}, \hat{\mathbf{y}}) - (\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \rangle \\
&\stackrel{(5)(6)}{=} \left(\frac{L}{\sqrt{2}} - \frac{1}{2\gamma} \right) \|(\hat{\mathbf{x}}, \hat{\mathbf{y}}) - \bar{\mathbf{z}}\|^2 + \frac{1}{2\gamma} \|\bar{\mathbf{z}} - (\tilde{\mathbf{x}}, \tilde{\mathbf{y}})\|^2 - \frac{1}{2\gamma} \|\mathbf{z}_+ - (\tilde{\mathbf{x}}, \tilde{\mathbf{y}})\|^2 \\
&\leq \frac{1}{2\gamma} \|\bar{\mathbf{z}} - (\tilde{\mathbf{x}}, \tilde{\mathbf{y}})\|^2 - \frac{1}{2\gamma} \|\mathbf{z}_+ - (\tilde{\mathbf{x}}, \tilde{\mathbf{y}})\|^2,
\end{aligned} \tag{7}$$

where we have used $\gamma \leq \frac{1}{\sqrt{2}L}$ in the last inequality.

Taking $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) := (\mathbf{x}^*, \mathbf{y}^*)$ in Equation (7) for the moment, we get

$$\|\mathbf{z}_+ - \mathbf{z}^*\|^2 \stackrel{(7)}{\leq} \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - 2\gamma \langle \mathcal{F}(\hat{\mathbf{x}}, \hat{\mathbf{y}}), (\hat{\mathbf{x}}, \hat{\mathbf{y}}) - \mathbf{z}^* \rangle \leq \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2. \tag{8}$$

Finally, in view of

$$\begin{aligned}
& \langle \mathcal{F}(\hat{\mathbf{x}}, \hat{\mathbf{y}}), (\hat{\mathbf{x}}, \hat{\mathbf{y}}) - (\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \rangle \\
&\stackrel{(7)}{\leq} \frac{1}{2\gamma} \|\bar{\mathbf{z}} - (\tilde{\mathbf{x}}, \tilde{\mathbf{y}})\|^2 - \frac{1}{2\gamma} \|\mathbf{z}_+ - (\tilde{\mathbf{x}}, \tilde{\mathbf{y}})\|^2 \\
&= \frac{1}{2\gamma} \left(\|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\|^2 - \|\mathbf{x}_+ - \tilde{\mathbf{x}}\|^2 + \|\bar{\mathbf{y}} - \tilde{\mathbf{y}}\|^2 - \|\mathbf{y}_+ - \tilde{\mathbf{y}}\|^2 \right) \\
&\leq \frac{1}{2\gamma} (\|\bar{\mathbf{x}} - \tilde{\mathbf{x}} + \mathbf{x}_+ - \tilde{\mathbf{x}}\| \cdot \|\bar{\mathbf{x}} - \mathbf{x}_+\| + \|\bar{\mathbf{y}} - \tilde{\mathbf{y}} + \mathbf{y}_+ - \tilde{\mathbf{y}}\| \cdot \|\bar{\mathbf{y}} - \mathbf{y}_+\|) \\
&\leq \frac{1}{\gamma} (D_X \|\bar{\mathbf{x}} - \mathbf{x}_+\| + D_Y \|\bar{\mathbf{y}} - \mathbf{y}_+\|) \\
&\leq \frac{1}{\gamma} \sqrt{D_X^2 + D_Y^2} \cdot \sqrt{\|\bar{\mathbf{x}} - \mathbf{x}_+\|^2 + \|\bar{\mathbf{y}} - \mathbf{y}_+\|^2} \\
&\leq \frac{1}{\gamma} \sqrt{D_X^2 + D_Y^2} \cdot \sqrt{2 \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 + 2 \|\mathbf{z}_+ - \mathbf{z}^*\|^2} \\
&\stackrel{(8)}{\leq} \frac{2}{\gamma} \sqrt{D_X^2 + D_Y^2} \cdot \|\bar{\mathbf{z}} - \mathbf{z}^*\|,
\end{aligned} \tag{9}$$

we have

$$\begin{aligned} u_1(\tilde{\mathbf{x}}, \hat{\mathbf{x}}) - u_1(\hat{\mathbf{x}}, \hat{\mathbf{y}}) + u_2(\hat{\mathbf{x}}, \tilde{\mathbf{y}}) - u_2(\hat{\mathbf{x}}, \hat{\mathbf{y}}) &\leq \langle \mathcal{F}(\hat{\mathbf{x}}, \hat{\mathbf{y}}), (\hat{\mathbf{x}}, \hat{\mathbf{y}}) - (\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \rangle \\ &\stackrel{(9)}{\leq} \frac{2}{\gamma} \sqrt{D_X^2 + D_Y^2} \cdot \|\bar{\mathbf{z}} - \mathbf{z}^*\|, \end{aligned}$$

and the desired bound follows because $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$ can take arbitrary points in X and Y , respectively. \square

We also state the following sufficient condition for the accurate Nash equilibrium, which can be used as stopping criterion for the optimization algorithms. Similar results can be found, for instance, in Nemirovski (2004); Yang et al. (2020).

Proposition 11 (Stopping criterion Nemirovski (2004); Yang et al. (2020)). *In a monotone general-sum game, let $\mathbf{z}^* = (\mathbf{x}^*, \mathbf{y}^*)$ be the Nash equilibrium. Let $\bar{\mathbf{z}} = (\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in X \times Y$, $\gamma \in (0, \frac{1}{2L}]$, and $\mu = \min\{\mu, \nu\}$. We have*

$$\|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 \leq \left(\frac{4}{\mu^2 \gamma^2} - \frac{2}{\mu \gamma} + 16 \right) \|\mathbf{z}_+ - \bar{\mathbf{z}}\|^2,$$

where $\mathbf{z}_+ = \Pi_Z(\bar{\mathbf{z}} - \gamma \mathcal{F}(\bar{\mathbf{z}}))$, in which $\hat{\mathbf{z}} = \Pi_Z(\bar{\mathbf{z}} - \gamma \mathcal{F}(\bar{\mathbf{z}}))$.

Proof. We have

$$\begin{aligned} &\left(1 - \frac{\mu \gamma}{2}\right) \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - \left(\frac{2}{\mu \gamma} - 1\right) \|\mathbf{z}_+ - \bar{\mathbf{z}}\|^2 \\ &\leq \|\mathbf{z}_+ - \mathbf{z}^*\|^2 \\ &\stackrel{(7)}{\leq} \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - 2\gamma \langle \mathcal{F}(\hat{\mathbf{z}}), \hat{\mathbf{z}} - \mathbf{z}^* \rangle \\ &\leq \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - 2\gamma \langle \mathcal{F}(\hat{\mathbf{z}}) - \mathcal{F}(\mathbf{z}^*), \hat{\mathbf{z}} - \mathbf{z}^* \rangle \\ &\leq \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - 2\mu \gamma \|\hat{\mathbf{z}} - \mathbf{z}^*\|^2 \\ &\leq \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - \mu \gamma \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 + 2\mu \gamma \|\hat{\mathbf{z}} - \bar{\mathbf{z}}\|^2 \\ &= (1 - \mu \gamma) \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - 2\mu \gamma \|\hat{\mathbf{z}} - \bar{\mathbf{z}}\|^2 + 4\mu \gamma \|\hat{\mathbf{z}} - \bar{\mathbf{z}}\|^2 \\ &\leq (1 - \mu \gamma) \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - 2\mu \gamma \|\hat{\mathbf{z}} - \bar{\mathbf{z}}\|^2 + 8\mu \gamma \|\mathbf{z}_+ - \bar{\mathbf{z}}\|^2 + 8\mu \gamma \|\mathbf{z}_+ - \hat{\mathbf{z}}\|^2 \\ &\leq (1 - \mu \gamma) \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - 2\mu \gamma \|\hat{\mathbf{z}} - \bar{\mathbf{z}}\|^2 + 8\mu \gamma \|\mathbf{z}_+ - \bar{\mathbf{z}}\|^2 + 8\mu \gamma \|\bar{\mathbf{z}} - \gamma \mathcal{F}(\hat{\mathbf{z}}) - \bar{\mathbf{z}} - \gamma \mathcal{F}(\bar{\mathbf{z}})\|^2 \\ &\leq (1 - \mu \gamma) \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 - 2\mu \gamma \|\hat{\mathbf{z}} - \bar{\mathbf{z}}\|^2 + 8\mu \gamma \|\mathbf{z}_+ - \bar{\mathbf{z}}\|^2 + 8\mu L^2 \gamma^3 \|\hat{\mathbf{z}} - \bar{\mathbf{z}}\|^2 \\ &\leq (1 - \mu \gamma) \|\bar{\mathbf{z}} - \mathbf{z}^*\|^2 + 8\mu \gamma \|\mathbf{z}_+ - \bar{\mathbf{z}}\|^2, \end{aligned}$$

where in the second to last inequality we use $\gamma \leq \frac{1}{2L}$. Finally, the desired bound follows from rearrangement. \square

C DISCUSSIONS ON THE CATALYST METHODS

In this section, we present the intuition of most existing algorithms for convex-concave minimax optimization considering general conditioning, and explain why similar idea may not work directly when generalized to monotone near-zero-sum games.

Most of existing algorithms for minimax optimization with general conditioning are based on Catalyst (Lin et al., 2018). In minimax optimization, we have $u_1 + u_2 = 0$. Assume without loss of generality that $\mu \leq \nu$. The function $f(\mathbf{x}) \stackrel{\text{def}}{=} -u_1(\mathbf{x}, \mathbf{y}(\mathbf{x}))$ is μ -strongly convex over $\mathbf{x} \in X$, in which $\mathbf{y}(\mathbf{x}) \stackrel{\text{def}}{=} \arg \max_{\mathbf{y} \in Y} u_2(\mathbf{x}, \mathbf{y})$. At the core of these algorithms, they build a function \hat{f}_t and get an inexact solution $\hat{\mathbf{x}}_{t+1}$ at each iteration t :

$$\hat{\mathbf{x}}_{t+1} \approx \arg \min_{\mathbf{x} \in X} \left[\hat{f}_t(\mathbf{x}) \stackrel{\text{def}}{=} f(\mathbf{x}) + \frac{\nu}{2} \|\mathbf{x} - \hat{\mathbf{x}}_t\|^2 \right]. \quad (10)$$

The outer loop is an inexact accelerated proximal point algorithm with $\tilde{\mathcal{O}}\left(\sqrt{\frac{\nu}{\mu}} \cdot \log\left(\frac{1}{\varepsilon}\right)\right)$ iterations (Nesterov, 2005; Lin et al., 2018; Carmon et al., 2022), and the inner loop of solving the smoothed Equation (10) can be any method with the number of gradient queries $\tilde{\mathcal{O}}\left(\frac{L}{\nu} \cdot \log\left(\frac{1}{\varepsilon}\right)\right)$ (Tseng, 1995). So, the total gradient complexity is⁵

$$\tilde{\mathcal{O}}\left(\sqrt{\frac{\nu}{\mu}} \cdot \log\left(\frac{1}{\varepsilon}\right)\right) \cdot \tilde{\mathcal{O}}\left(\frac{L}{\nu} \cdot \log\left(\frac{1}{\varepsilon}\right)\right) = \tilde{\mathcal{O}}\left(\frac{L}{\sqrt{\mu\nu}} \cdot \log^2\left(\frac{1}{\varepsilon}\right)\right).$$

However, if we try to apply the above Catalyst method to monotone non-zero-sum games, the algorithm may only converge to a Stackelberg solution, which can be very different from the Nash equilibrium in non-zero-sum games.

EXAMPLE 3 (Stackelberg solution). *Consider the case where $X = [0, 1] \times [1, 2] \subseteq \mathbb{R}^2$ and $Y = [-1, 0] \subseteq \mathbb{R}$. Let Player 1 maximize*

$$u_1(\mathbf{x}, y) = -\frac{1}{2}(x_1 - 1)^2 - \frac{1}{2}(x_2 - 1)^2 + \frac{1}{2}x_1y$$

over $\mathbf{x} \in X$, and Player 2 maximize

$$u_2(\mathbf{x}, y) = \frac{1}{2}x_2y - (y + 1)^2$$

over $y \in Y$. Then, the Catalyst minimization of $f(x) = -u_1(\mathbf{x}, y(\mathbf{x}))$ will lead to the Stackelberg solution $(\mathbf{x} = (\frac{40}{63}, \frac{68}{63}), y = -\frac{46}{63})$, which is different from the Nash equilibrium $(\mathbf{x} = (\frac{5}{8}, 1), y = -\frac{3}{4})$.

Therefore, we are not aware of how the Catalyst methods for convex-concave minimax optimization can be applied in NEPs for non-zero-sum games.

D PROOF DETAILS

D.1 PROOFS FOR THE RESULTS IN SECTION 3.1

Proof of Proposition 3. For any $\mathbf{z} = (\mathbf{x}, \mathbf{y}) \in X \times Y$,

$$\Delta(\mathbf{z}) \geq g(\mathbf{z}) - g(\mathbf{z}) + h(\mathbf{x}, \mathbf{y}) - h(\mathbf{x}, \mathbf{y}) = 0,$$

and for all $\tilde{\mathbf{z}} = (\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in X \times Y$, we have

$$\begin{aligned} \Delta(\mathbf{z}) &\geq \frac{1}{2}[g(\mathbf{z}) - g(\mathbf{x}, \tilde{\mathbf{y}}) + h(\mathbf{x}, \tilde{\mathbf{y}}) - h(\mathbf{x}, \mathbf{y})] + \frac{1}{2}[g(\mathbf{z}) - g(\tilde{\mathbf{x}}, \mathbf{y}) + h(\mathbf{x}, \mathbf{y}) - h(\tilde{\mathbf{x}}, \mathbf{y})] \\ &= \frac{1}{2}[2g(\mathbf{z}) + u_2(\mathbf{x}, \tilde{\mathbf{y}}) + u_1(\tilde{\mathbf{x}}, \mathbf{y})] \\ &= \frac{1}{2}[u_1(\tilde{\mathbf{x}}, \mathbf{y}) - u_1(\mathbf{x}, \mathbf{y}) + u_2(\mathbf{x}, \tilde{\mathbf{y}}) - u_2(\mathbf{x}, \mathbf{y})]. \end{aligned}$$

□

Proof of Proposition 4. The (if) part follows directly from Proposition 3. Now we prove the (only if) part. Suppose $\mathbf{z}^* = (\mathbf{x}^*, \mathbf{y}^*)$ is the Nash equilibrium. For all $\tilde{\mathbf{z}} = (\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in X \times Y$,

$$g(\mathbf{z}^*) - g(\tilde{\mathbf{z}}) + h(\mathbf{x}^*, \tilde{\mathbf{y}}) - h(\tilde{\mathbf{x}}, \mathbf{y}^*) \leq \langle \nabla g(\mathbf{z}^*), \mathbf{z}^* - \tilde{\mathbf{z}} \rangle + \langle \mathcal{H}(\mathbf{z}^*), \mathbf{z}^* - \tilde{\mathbf{z}} \rangle \leq 0,$$

where in the first inequality we use Assumptions 1 and 2. Then, we have $\Delta(\mathbf{z}^*) = 0$. □

⁵The double logarithm term may be avoided by combining this algorithmic idea with some complicated techniques (Kovalev & Gasnikov, 2022; Carmon et al., 2022), which we omit here for the simplicity of presentation.

D.2 PROOFS FOR THE RESULTS IN SECTION 3.2

The main technical work in the convergence analysis is to use the properties of our potential function and prove the descent lemma (Lemma 5).

Proof of Lemma 5. By Assumption 1, we can upper bound the convex-concave zero-sum part

$$\begin{aligned}
h(\mathbf{x}_{t+1}, \mathbf{y}^*) - h(\mathbf{x}^*, \mathbf{y}_{t+1}) &= h(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}) - h(\mathbf{x}^*, \mathbf{y}_{t+1}) + h(\mathbf{x}_{t+1}, \mathbf{y}^*) - h(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}) \\
&\leq \langle \nabla_{\mathbf{x}} h(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}), \mathbf{x}_{t+1} - \mathbf{x}^* \rangle - \frac{\mu}{2} \|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 \\
&\quad - \langle \nabla_{\mathbf{y}} h(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}), \mathbf{y}_{t+1} - \mathbf{y}^* \rangle - \frac{\nu}{2} \|\mathbf{y}_{t+1} - \mathbf{y}^*\|^2 \\
&= \langle \mathcal{H}(\mathbf{z}_{t+1}), \mathbf{z}_{t+1} - \mathbf{z}^* \rangle - \frac{\mu}{2} \|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 - \frac{\nu}{2} \|\mathbf{y}_{t+1} - \mathbf{y}^*\|^2.
\end{aligned} \tag{11}$$

By Assumptions 2 and 3, we can upper bound the jointly convex coupling part

$$\begin{aligned}
g(\mathbf{z}_{t+1}) - g(\mathbf{z}^*) &= g(\mathbf{z}_{t+1}) - g(\mathbf{z}_t) + g(\mathbf{z}_t) - g(\mathbf{z}^*) \\
&\leq \langle \nabla g(\mathbf{z}_t), \mathbf{z}_{t+1} - \mathbf{z}_t \rangle + \frac{\delta}{2} \|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2 + \langle \nabla g(\mathbf{z}_t), \mathbf{z}_t - \mathbf{z}^* \rangle \\
&= \langle \nabla g(\mathbf{z}_t), \mathbf{z}_{t+1} - \mathbf{z}^* \rangle + \frac{\delta}{2} \|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2.
\end{aligned} \tag{12}$$

In view of

$$\langle \mathbf{z}_{t+1} - \mathbf{z}_t, \mathbf{z}_{t+1} - \mathbf{z}^* \rangle = \frac{1}{2} \|\mathbf{z}_{t+1} - \mathbf{z}^*\|^2 - \frac{1}{2} \|\mathbf{z}_t - \mathbf{z}^*\|^2 + \frac{1}{2} \|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2, \tag{13}$$

and

$$\begin{aligned}
&\left\langle \nabla g(\mathbf{z}_t) + \mathcal{H}(\mathbf{z}_{t+1}) + \frac{1}{\eta_t} (\mathbf{z}_{t+1} - \mathbf{z}_t), \mathbf{z}_{t+1} - \mathbf{z}^* \right\rangle \\
&\leq \langle \nabla_{\mathbf{x}} \varphi_t(\mathbf{z}_{t+1}), \mathbf{x}_{t+1} - \mathbf{x}^* \rangle - \langle \nabla_{\mathbf{y}} \varphi_t(\mathbf{z}_{t+1}), \mathbf{y}_{t+1} - \mathbf{y}^* \rangle \\
&\stackrel{(2)}{\leq} \varepsilon_t,
\end{aligned} \tag{14}$$

we have

$$\begin{aligned}
0 &= -\Delta(\mathbf{z}^*) \leq g(\mathbf{z}_{t+1}) - g(\mathbf{z}^*) + h(\mathbf{x}_{t+1}, \mathbf{y}^*) - h(\mathbf{x}^*, \mathbf{y}_{t+1}) \\
&\stackrel{(11)(12)}{\leq} \langle \nabla g(\mathbf{z}_t) + \mathcal{H}(\mathbf{z}_{t+1}), \mathbf{z}_{t+1} - \mathbf{z}^* \rangle - \frac{\mu}{2} \|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 - \frac{\nu}{2} \|\mathbf{y}_{t+1} - \mathbf{y}^*\|^2 + \frac{\delta}{2} \|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2 \\
&= \left\langle \nabla g(\mathbf{z}_t) + \mathcal{H}(\mathbf{z}_{t+1}) + \frac{1}{\eta_t} (\mathbf{z}_{t+1} - \mathbf{z}_t), \mathbf{z}_{t+1} - \mathbf{z}^* \right\rangle - \frac{1}{\eta_t} \langle \mathbf{z}_{t+1} - \mathbf{z}_t, \mathbf{z}_{t+1} - \mathbf{z}^* \rangle \\
&\quad - \frac{\mu}{2} \|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 - \frac{\nu}{2} \|\mathbf{y}_{t+1} - \mathbf{y}^*\|^2 + \frac{\delta}{2} \|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2 \\
&\stackrel{(13)(14)}{\leq} \varepsilon_t + \frac{1}{2\eta_t} \|\mathbf{x}_t - \mathbf{x}^*\|^2 - \left(\frac{1}{2\eta_t} + \frac{\mu}{2} \right) \|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 \\
&\quad + \frac{1}{2\eta_t} \|\mathbf{y}_t - \mathbf{y}^*\|^2 - \left(\frac{1}{2\eta_t} + \frac{\nu}{2} \right) \|\mathbf{y}_{t+1} - \mathbf{y}^*\|^2 - \left(\frac{1}{2\eta_t} - \frac{\delta}{2} \right) \|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2,
\end{aligned}$$

where the first equality follows from Proposition 4 and the first inequality follows from the definition of $\Delta(\cdot)$. Finally, the desired bound follows from $\eta_t \leq \frac{1}{\delta}$. \square

With Lemma 5, we are ready to prove the complexity of the outer loop (Lemma 6).

Proof of Lemma 6. For monotone δ -nearly-zero-sum games and $\eta \leq \frac{1}{\delta}$, by Lemma 5, for any $k \in [0, t-1] \cap \mathbb{Z}$, we have

$$\|\mathbf{z}_{k+1} - \mathbf{z}^*\|^2 \leq (1 - \theta) \|\mathbf{z}_k - \mathbf{z}^*\|^2 + 2\eta\varepsilon_k.$$

Then, unrolling this recursion (from $k = t - 1, t - 2, \dots$, to 0) yields

$$\begin{aligned}
\|\mathbf{z}_t - \mathbf{z}^*\|^2 &\leq (1 - \theta)^t \|\mathbf{z}_0 - \mathbf{z}^*\|^2 + 2\eta \sum_{k=0}^{t-1} (1 - \theta)^{t-k-1} \varepsilon_k \\
&\leq (1 - \theta)^t (D_X^2 + D_Y^2) + \frac{2\eta}{\theta} \cdot \max_{k \in [0, t-1] \cap \mathbb{Z}} \varepsilon_k \\
&\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} \\
&= \varepsilon,
\end{aligned}$$

where the last inequality follows from $t \geq \frac{1}{\theta} \log \frac{2(D_X^2 + D_Y^2)}{\varepsilon}$ and $\varepsilon_t \leq \frac{\theta \varepsilon}{4\eta}$. \square

Below, we also include the proof of the gradient complexity of the inner loops for completeness. This result of the inner loops is heavily based on the previous results of optimal gradient methods in minimax optimization (see, for instance, Kovalev & Gasnikov (2022); Carmon et al. (2022); Thekumpampil et al. (2022); Lan & Li (2023)).

Proof of Lemma 7. Let $\mathbf{z}_{t+1}^* = (\mathbf{x}_{t+1}^*, \mathbf{y}_{t+1}^*) \in X \times Y$ denote the saddle point of $\varphi_t(\cdot, \cdot)$. Denote

$$\bar{\varepsilon}_t = \frac{\varepsilon_t^2}{8L^2(D_X^2 + D_Y^2)}.$$

By Proposition 10, an inexact solution in Equation (2) of Algorithm 1 can be obtained from a pair of decisions $\bar{\mathbf{z}}_{t+1} = (\bar{\mathbf{x}}_{t+1}, \bar{\mathbf{y}}_{t+1}) \in X \times Y$ that satisfies $\|\bar{\mathbf{z}}_{t+1} - \mathbf{z}_{t+1}^*\|^2 \leq \bar{\varepsilon}_t$.

The function $\varphi_t(\cdot, \cdot)$ is $(\eta_t^{-1} + \mu)$ -strongly convex- $(\eta_t^{-1} + \nu)$ -strongly concave and $2L$ -smooth, where the $2L$ -smoothness follows from $\eta_t \geq \frac{1}{L}$. Hence, by Kovalev & Gasnikov (2022); Carmon et al. (2022); Thekumpampil et al. (2022); Lan & Li (2023), the aforementioned pair of decisions $\bar{\mathbf{z}}_{t+1}$ can be found within

$$\mathcal{O} \left(\frac{L}{\sqrt{(\eta_t^{-1} + \mu)(\eta_t^{-1} + \nu)}} \cdot \log \left(\frac{D_X^2 + D_Y^2}{\bar{\varepsilon}_t} \right) \right)$$

gradient queries. Finally, after substituting the $\bar{\varepsilon}_t$, the desired bound follows. \square

Finally, we prove Theorem 1, our main theoretical result.

Proof of Theorem 1. The overall gradient complexity is given by the multiplication of outer loop iterations (Lemma 6) and inner loop gradient complexity (Lemma 7):

$$\begin{aligned}
&\mathcal{O} \left(\frac{\eta^{-1} + \min\{\mu, \nu\}}{\min\{\mu, \nu\}} \cdot \log \frac{2(D_X^2 + D_Y^2)}{\varepsilon} \right) \cdot \mathcal{O} \left(\frac{L}{\sqrt{(\eta^{-1} + \mu)(\eta^{-1} + \nu)}} \cdot \log \frac{L(D_X^2 + D_Y^2)}{\varepsilon_t} \right) \\
&= \mathcal{O} \left(\frac{\delta + \min\{\mu, \nu\}}{\min\{\mu, \nu\}} \cdot \log \frac{D_X^2 + D_Y^2}{\varepsilon} \right) \cdot \mathcal{O} \left(\frac{L}{\sqrt{(\delta + \mu)(\delta + \nu)}} \cdot \log \left(\frac{L(D_X^2 + D_Y^2)}{\min\{\mu, \nu\} \cdot \varepsilon} \right) \right) \\
&= \mathcal{O} \left(\frac{L}{\min\{\mu, \nu\}} \cdot \sqrt{\frac{\delta + \min\{\mu, \nu\}}{\delta + \max\{\mu, \nu\}}} \cdot \log \left(\frac{L(D_X^2 + D_Y^2)}{\min\{\mu, \nu\} \cdot \varepsilon} \right) \log \left(\frac{D_X^2 + D_Y^2}{\varepsilon} \right) \right) \\
&= \mathcal{O} \left(\left(\frac{L}{\sqrt{\mu\nu}} + \frac{L}{\min\{\mu, \nu\}} \cdot \min \left\{ 1, \sqrt{\frac{\delta}{\mu + \nu}} \right\} \right) \cdot \log \left(\frac{L(D_X^2 + D_Y^2)}{\min\{\mu, \nu\} \cdot \varepsilon} \right) \log \left(\frac{D_X^2 + D_Y^2}{\varepsilon} \right) \right),
\end{aligned}$$

where the first relation follows from $\eta = \min \left\{ \frac{1}{\delta}, \frac{1}{\min\{\mu, \nu\}} \right\}$. \square

D.3 PROOFS FOR THE RESULT IN SECTION 3.3

We state a more general formulation of Corollary 9.

Corollary 12. *For monotone δ -near-zero-sum games where $\mu = 0$ or $\nu = 0$, an ε -approximate Nash equilibrium can be found within*

$$\mathcal{O} \left(\left(\frac{L}{\sqrt{\bar{\mu}\bar{\nu}}} + \frac{L}{\min\{\bar{\mu}, \bar{\nu}\}} \cdot \min \left\{ 1, \sqrt{\frac{\delta}{\bar{\mu} + \bar{\nu}}} \right\} \right) \cdot \log \left(\frac{L^2 D^2}{\min\{\bar{\mu}, \bar{\nu}\} \cdot \varepsilon} \right) \log \left(\frac{LD^2}{\varepsilon} \right) \right)$$

gradient queries, where $\bar{\mu} = \mu + \min \left\{ \frac{\varepsilon}{2D_X^2}, L \right\}$ and $\bar{\nu} = \nu + \min \left\{ \frac{\varepsilon}{2D_Y^2}, L \right\}$.

Proof of Corollary 12. We consider the reduced game where Player 1 maximizes

$$\hat{u}_1 = u_1 - \min \left\{ \frac{\varepsilon}{4D_X^2}, \frac{L}{2} \right\} \|\mathbf{x}\|^2 + \min \left\{ \frac{\varepsilon}{4D_Y^2}, \frac{L}{2} \right\} \|\mathbf{y}\|^2$$

over $\mathbf{x} \in X$ and Player 2 maximizes

$$\hat{u}_2 = u_2 + \min \left\{ \frac{\varepsilon}{4D_X^2}, \frac{L}{2} \right\} \|\mathbf{x}\|^2 - \min \left\{ \frac{\varepsilon}{4D_Y^2}, \frac{L}{2} \right\} \|\mathbf{y}\|^2$$

over $\mathbf{y} \in Y$. Any $\frac{\varepsilon}{2}$ -approximate Nash equilibrium of the reduced game is an ε -approximate Nash equilibrium in the original game. This reduction is similar to the ones used in Lin et al. (2020); Wang & Li (2020).

Denote $\hat{g} \triangleq -\frac{1}{2}(\hat{u}_1 + \hat{u}_2)$ and $\hat{h} \triangleq \frac{1}{2}(-\hat{u}_1 + u_2)$. Then, we have $\hat{h} = h + \left\{ \frac{\varepsilon}{4D_X^2}, \frac{L}{2} \right\} \|\mathbf{x}\|^2 - \left\{ \frac{\varepsilon}{4D_Y^2}, \frac{L}{2} \right\} \|\mathbf{y}\|^2$, which is $2L$ -smooth and $\bar{\mu}$ -strongly convex- $\bar{\nu}$ -strongly concave. We also have $\hat{g} = -\frac{1}{2}(\hat{u}_1 + \hat{u}_2) = -\frac{1}{2}(u_1 + u_2) = g$, which is jointly convex and δ -smooth. By Theorem 1, we obtain the number of gradient queries for an $\frac{\varepsilon^2}{32L^2D^2}$ -accurate Nash equilibrium in the reduced game:

$$\mathcal{O} \left(\left(\frac{L}{\sqrt{\bar{\mu}\bar{\nu}}} + \frac{L}{\min\{\bar{\mu}, \bar{\nu}\}} \cdot \min \left\{ 1, \sqrt{\frac{\delta}{\bar{\mu} + \bar{\nu}}} \right\} \right) \cdot \log \left(\frac{L^2 D^2}{\min\{\bar{\mu}, \bar{\nu}\} \cdot \varepsilon} \right) \log \left(\frac{LD^2}{\varepsilon} \right) \right).$$

Finally, following from Proposition 10, we obtain the desired $\frac{\varepsilon}{2}$ -approximate Nash equilibrium of the reduced game by taking an extragradient step from the $\frac{\varepsilon^2}{32L^2D^2}$ -accurate Nash equilibrium. \square

D.4 PROOFS FOR THE RESULTS IN SECTION 4

Proposition 13 (Convex reformulation in bilinear coupling). *For $\beta_1, \beta_2 \geq 0$ and $\mathbf{M} \in \mathbb{R}^{m \times n}$ such that $\sqrt{\beta_1\beta_2} \geq \|\mathbf{M}\|$, the function $\tilde{g}(\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ defined as*

$$\tilde{g}(\mathbf{x}, \mathbf{y}) = \frac{\beta_1}{2} \|\mathbf{x}\|^2 + \langle \mathbf{M}\mathbf{x}, \mathbf{y} \rangle + \frac{\beta_2}{2} \|\mathbf{y}\|^2$$

is jointly convex.

Proof. The quadratic function $\tilde{g}(\cdot, \cdot)$ is bounded below: for all $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{y} \in \mathbb{R}^m$,

$$\begin{aligned} \tilde{g}(\mathbf{x}, \mathbf{y}) &\geq \frac{\beta_1}{2} \|\mathbf{x}\|^2 - \|\mathbf{M}\mathbf{x}\| \|\mathbf{y}\| + \frac{\beta_2}{2} \|\mathbf{y}\|^2 \\ &\geq \frac{\beta_1}{2} \|\mathbf{x}\|^2 - \sqrt{\beta_1\beta_2} \|\mathbf{x}\| \|\mathbf{y}\| + \frac{\beta_2}{2} \|\mathbf{y}\|^2 \\ &\geq 0, \end{aligned}$$

where in the first inequality we used the Cauchy-Schwarz inequality. Therefore, $\tilde{g}(\cdot, \cdot)$ is jointly convex. \square

Proposition 14 (Convex reformulation in general coupling). *For $\beta \geq 0$ and $g : X \times Y \rightarrow \mathbb{R}$ such that $g(\cdot, \cdot)$ is β -smooth, the function $\tilde{g}(\cdot, \cdot) : X \times Y \rightarrow \mathbb{R}$ defined as*

$$\tilde{g}(\mathbf{x}, \mathbf{y}) = \frac{\beta}{2} \|\mathbf{x}\|^2 + g(\mathbf{x}, \mathbf{y}) + \frac{\beta}{2} \|\mathbf{y}\|^2$$

is jointly convex.

Proof. For all $\mathbf{z} = (\mathbf{x}, \mathbf{y}) \in X \times Y$ and $\mathbf{z}' = (\mathbf{x}', \mathbf{y}') \in X \times Y$, we have

$$\begin{aligned} \langle \nabla \tilde{g}(\mathbf{z}') - \nabla \tilde{g}(\mathbf{z}), \mathbf{z}' - \mathbf{z} \rangle &= \beta \|\mathbf{z}' - \mathbf{z}\|^2 + \langle \nabla g(\mathbf{z}') - \nabla g(\mathbf{z}), \mathbf{z}' - \mathbf{z} \rangle \\ &\geq \beta \|\mathbf{z}' - \mathbf{z}\|^2 - \beta \|\mathbf{z}' - \mathbf{z}\|^2 \\ &= 0, \end{aligned}$$

where the first inequality follows from the β -smoothness of $g(\cdot, \cdot)$. Therefore, the function $\tilde{g}(\cdot, \cdot)$ is jointly convex (Nesterov, 2004, Theorem 2.1.3). \square

E ADDITIONAL RESULTS FOR OTHER ORACLE AND FUNCTION CLASSES

In this section, we consider a different class of Nash equilibrium problem and demonstrate the applicability of our ICL framework. In Bōt et al. (2023), they considered a zero-sum (or strictly competitive) game where the two players have proximal oracle and gradient oracle, respectively. We now consider the generalization where an additional incentive is added. To prevent ambiguity, we will define the problem class in a self-contained way.

In this section, we are interested in the Nash equilibrium problem, or equivalently, the variational inequality problem given by operator \mathcal{F} defined on $X \times Y$:

$$\mathcal{F}(\mathbf{x}, \mathbf{y}) = \nabla g(\mathbf{x}, \mathbf{y}) + (\nabla_{\mathbf{x}} h(\mathbf{x}, \mathbf{y}), -\nabla_{\mathbf{y}} h(\mathbf{x}, \mathbf{y}) + \nabla \psi(\mathbf{y})), \quad (\mathbf{x}, \mathbf{y}) \in X \times Y,$$

where

1. The sets X and Y are compact convex sets in Euclidean spaces. The diameter of X is bounded by D_X , and the diameter of Y is bounded by D_Y .
2. The function $g : X \times Y \rightarrow \mathbb{R}$ is δ -smooth and convex.
3. The function $\psi : Y \rightarrow \mathbb{R} \cup \{+\infty\}$ is proper, lower semicontinuous, ν -strongly convex, and with domain $\text{dom } \psi = \{\mathbf{y} \in Y \mid \psi(\mathbf{y}) < +\infty\}$.
4. For all $\mathbf{y} \in \text{dom } \psi$, the function $h(\cdot, \mathbf{y}) : X \rightarrow \mathbb{R} \cup \{+\infty\}$ is proper, lower semi-continuous, and μ -strongly convex.
5. For all $\mathbf{x} \in \Pi_X(\text{dom } h) \triangleq \{\mathbf{u} \in X \mid \exists \mathbf{v} \in Y \text{ such that } (\mathbf{u}, \mathbf{v}) \in \text{dom } h\}$, we have that $\text{dom } h(\mathbf{x}, \cdot) = Y$ and the function $h(\mathbf{x}, \cdot) : Y \rightarrow \mathbb{R}$ is concave and continuously differentiable. Moreover, $\Pi_X(\text{dom } h)$ is closed.
6. There exists $L_{yx}, L_{yy} \geq 0$ such that for all $(x, y), (x', y') \in \Pi_X(\text{dom } h) \times \text{dom } \psi$,

$$\|\nabla_{\mathbf{y}} h(\mathbf{x}, \mathbf{y}) - \nabla_{\mathbf{y}} h(\mathbf{x}', \mathbf{y}')\| \leq L_{yx} \|\mathbf{x} - \mathbf{x}'\| + L_{yy} \|\mathbf{y} - \mathbf{y}'\|.$$

We assume the players can query the gradient ∇g , the proximal operator of $h(\cdot, \mathbf{y})$ for any fixed $\mathbf{y} \in Y$, the partial gradient $\nabla_{\mathbf{y}} h(\cdot, \cdot)$, and the proximal oracle of $\psi(\cdot)$. The problem studied in Bōt et al. (2023) corresponds to a special case of the additional incentive $\delta = 0$, while we generalize their results to $\delta \neq 0$.

We will use the complexity results in Bōt et al. (2023) as a black box. We cite their results below.

Lemma 15 (Bōt et al. (2023), Theorem 14). *For $\delta = 0$ and $\mu > 0$, there exists an algorithm which returns an ε -accurate Nash equilibrium with the number of partial gradient queries to $\nabla_{\mathbf{y}} h(\cdot)$ and the number of proximal oracle queries to $\text{prox}_{h(\cdot, \mathbf{y})}(\cdot)$ and $\text{prox}_{\psi}(\cdot)$ bounded by*

$$\mathcal{O} \left(\left(1 + \frac{L_{yx}}{\sqrt{\mu\nu}} + \frac{L_{yy}}{\nu} \right) \cdot \log \left(\frac{D^2}{\varepsilon} \right) \right).$$

By applying our ICL algorithm (Algorithm 1), we obtain the following complexity result:

Theorem 2. Assume h is L -smooth over $\Pi_X(\text{dom } h) \times \text{dom } \psi$, and ψ is L -smooth over $\text{dom } \psi$. For $\mu > 0$, there exists an algorithm which returns an ε -accurate Nash equilibrium with the number of partial gradient queries to $\nabla_y h(\cdot)$ and the number of proximal oracle queries to $\text{prox}_{h(\cdot, y)}(\cdot)$ and $\text{prox}_\psi(\cdot)$ bounded by

$$\mathcal{O} \left(\left(1 + \frac{\delta}{\min\{\mu, \nu\}} \right) \left(1 + \frac{L_{yx}}{\sqrt{(\delta + \mu)(\delta + \nu)}} + \frac{L_{yy}}{\delta + \nu} \right) \cdot \log \left(\frac{LD^2}{\min\{\mu, \nu\} \cdot \varepsilon} \right) \log \left(\frac{D^2}{\varepsilon} \right) \right),$$

and with the number of gradient queries to $\nabla g(\cdot)$ bounded by

$$\mathcal{O} \left(\left(1 + \frac{\delta}{\min\{\mu, \nu\}} \right) \cdot \log \left(\frac{D^2}{\varepsilon} \right) \right).$$

Proof. The result follows from multiplying the outer loop iterations in Lemma 6 and the inner loop complexity in Lemma 15. \square

F ILLUSTRATION OF THE APPLICATION EXAMPLES

Matrix games with transaction fees We give a concrete illustration for matrix games with transaction fees. Let the payoff matrices of Player 1 and Player 2 without transaction fees be

$$\mathbf{M} = \begin{bmatrix} 300 & -200 \\ -100 & 400 \end{bmatrix} \quad \text{and} \quad -\mathbf{M} = \begin{bmatrix} -300 & 200 \\ 100 & -400 \end{bmatrix},$$

respectively. Then,

$$\text{abs}(\mathbf{M}) = \begin{bmatrix} 300 & 200 \\ 100 & 400 \end{bmatrix}, \quad \mathbf{M}_+ = \begin{bmatrix} 300 & 0 \\ 0 & 400 \end{bmatrix}, \quad \text{and} \quad \mathbf{M}_- = \begin{bmatrix} 0 & 200 \\ 100 & 0 \end{bmatrix}.$$

Let 1% of transaction fees be imposed on every payment. Then, the payoff matrices of Player 1 and Player 2 with transaction fees are

$$\mathbf{A} = \begin{bmatrix} 297 & -200 \\ -100 & 396 \end{bmatrix} \quad \text{and} \quad \mathbf{B} = \begin{bmatrix} -300 & 198 \\ 99 & -400 \end{bmatrix},$$

respectively. We also draw the Table 2 for easier comparisons.

Table 2: An illustration of matrix games with transaction fee $\rho = 0.01$.

300/-300	-200/200
-100/100	400/-400

297/-300	-200/198
-100/99	396/-400

Competitive games with small cooperation incentives We give a concrete illustration for competitive games with small cooperation incentives. The donation games are a canonical model in evolutionary game theory for studying the altruistic collaboration (Nowak, 2006a; Sigmund, 2010).

Let us consider a simplified version of donation games played for one round, as shown in Table 3. In this game, a player can choose to incur a personal cost, c , to provide a larger benefit, b , to another player. Consider a concrete example where the cost to donate is $c = 50$ and the benefit conferred is only slightly higher at $b = 51$.

In this scenario, if both players cooperate, they each pay a cost of 50 and receive a benefit of 51, resulting in a modest net payoff of 1. However, the temptation to defect is substantial: by withholding their own donation while still receiving the benefit from the other player, a defector can achieve a payoff of 51. Conversely, the cooperating player who is defected upon is left with a “sucker’s payoff” of -50 . This low benefit-to-cost ratio ($b/c \approx 1.02$) creates a very small incentive for cooperation.

Table 3: An illustration of donation games.

		Player 2: Cooperate	Player 2: Defect
		(1, 1)	(-50, 51)
		(51, -50)	(0, 0)
Player 1: Cooperate	Player 1: Defect		

G MORE EXPERIMENTAL DETAILS

G.1 IMPLEMENTATION DETAILS

We generate the sparse matrix \mathbf{M} following the procedures outlined in Nemirovski (2004); Nesterov (2005): (i) The random seeds are set from 0, 111, 222, ..., and 999; (ii) 100000 coordinates of \mathbf{M} are chosen uniformly at random; (iii) each chosen coordinate is assigned a random value independently drawn from a uniform distribution between $[-1, 1]$; (iv) all remaining coordinates are set to 0.

We implement our ICL method as described in Algorithm 1. The classic OGDA and classic EG methods are implemented as outlined in Popov (1980) and Korpelevich (1976), respectively. All solvers are initialized at $(\mathbf{x}_0, \mathbf{y}_0) = (\mathbf{1}_n/n, \mathbf{1}_m/m)$, where $\mathbf{1}_k \in \mathbb{R}^k$ denotes the vector of size k where every element in the vector is equal to 1. The setup for ICL is detailed in Theorem 1. The stepsize for OGDA is set to $\frac{1}{2L}$ following Popov (1980); Mokhtari et al. (2020), and for EG is set to $\frac{1}{\sqrt{2L}}$ following Korpelevich (1976); Nemirovski (2004). For the inner loop, the Lifted Primal Dual method (Thekumparampil et al., 2022) is used, with the theoretical setup maintained as specified in (Thekumparampil et al., 2022, Theorem 2).

G.2 MORE DETAILS OF THE EXPERIMENT RUNS

We conducted our experiments on e2-highcpu vCPUs within the Google Cloud environment. The memory requirement of our experiments is quite modest, requiring only sufficient RAM for a few 10000×10000 sparse matrices (that is, about 60 MB). Each independent run completes within about 3 minutes.

We plot the convergence behaviors in Figure 1. Note that for ICL, only iterates within the outer loop are plotted. Figure 1 shows results for a single seed (seed 0), as plotting all seeds in a single figure would introduce excessive visual complexity due to the unaligned x -axis representing the counts of gradient queries in the outer loop. Nonetheless, we observed consistent convergence patterns across different seeds: (i) Transaction fee changes have little impact on the convergence of OGDA and EG, but significantly accelerate the convergence of ICL as ρ decreases; (ii) ICL converges fastest when $\rho \leq 0.12\%$; and (iii) OGDA converges fastest when $\rho \geq 0.15\%$.

We should clarify that all 21 entries (7 entries for each of the three algorithms) are indeed plotted in Figure 1. The perceived ‘lack of distinct lines’ is actually a crucial empirical result supporting our theoretical claims:

- For OGDA (Red) and EG (Green): The seven lines corresponding to different parameter settings for OGDA and EG almost completely merge and overlap. This empirical observation is fully consistent with our theory, which states that the convergence rate of classical variational inequality methods does not benefit from the near-zero-sum structure. Their rates are determined solely by $\frac{L}{\min\{\mu, \nu\}}$, regardless of the δ parameter.
- For the proposed ICL algorithm (Blue): The seven lines are clearly separated, showing a distinct dependency on the δ parameter. This separation empirically validates our key claim that our algorithm successfully harnesses the near-zero-sum structure to achieve a faster convergence rate.

Finally, we report CPU times of experiment runs to converge to an ε -accurate Nash equilibrium in Table 4, with error bars indicating 2-sigma variations across 10 independent runs using randomly generated matrices. Table 4 shows that ICL achieves the shortest CPU time when $\rho \leq 0.12\%$, while OGDA achieves the shortest CPU time when $\rho \geq 0.15\%$.

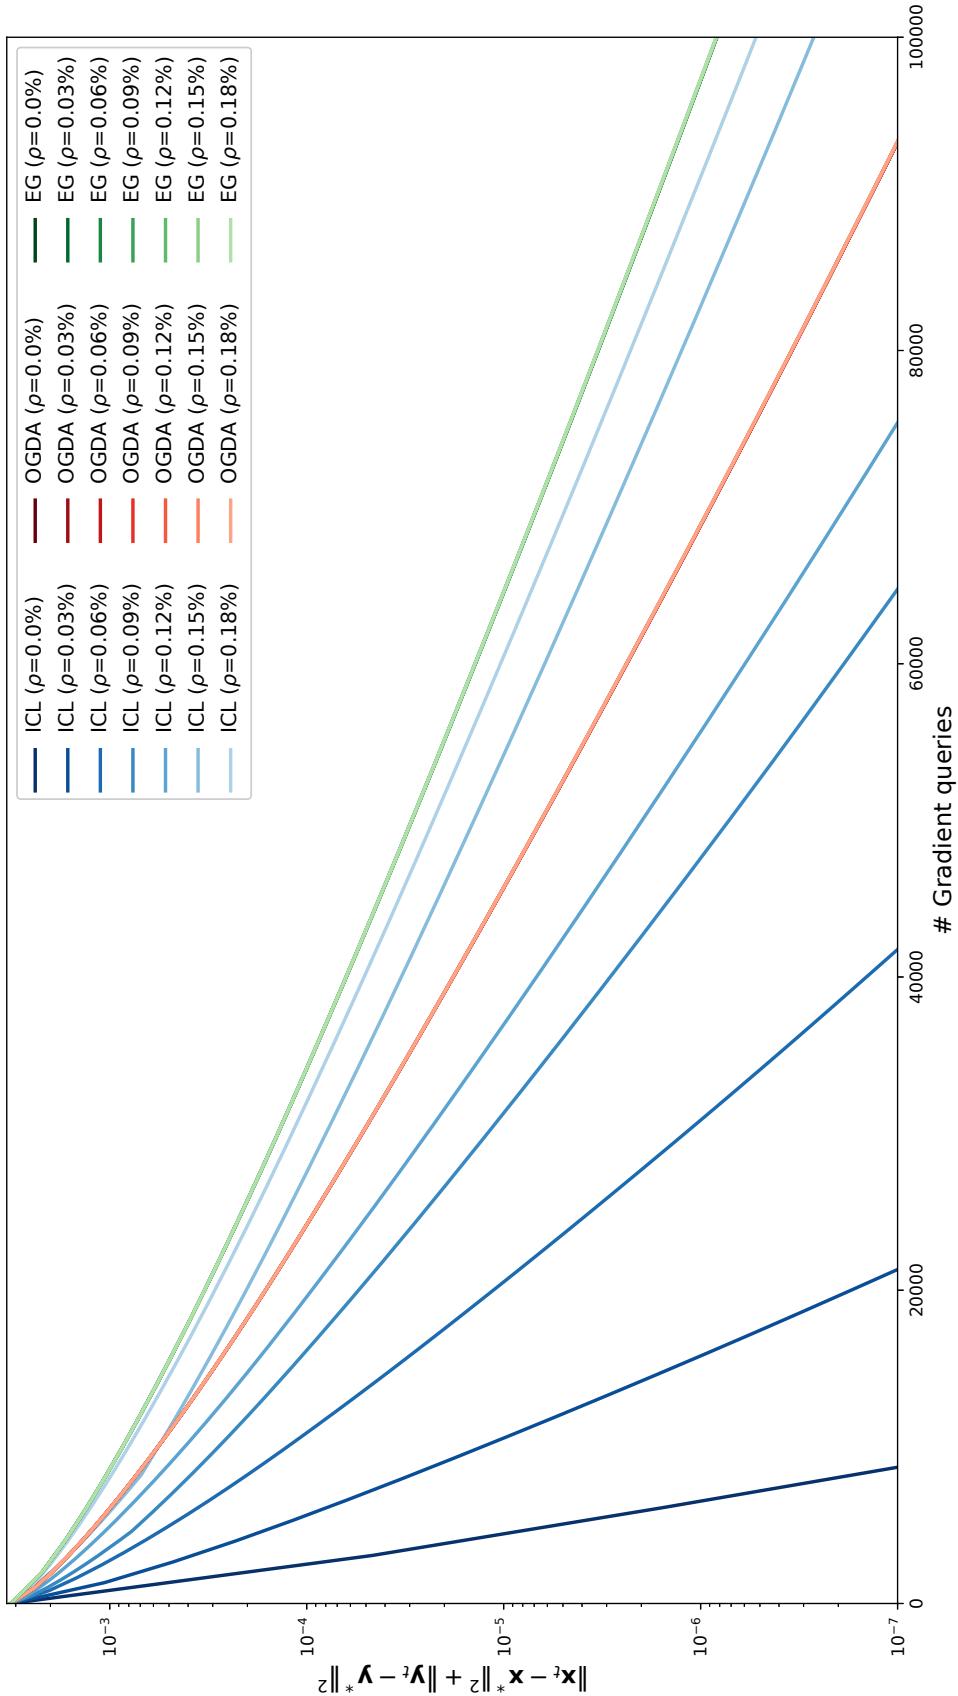


Figure 1: Comparisons of the convergence of the ICL, OGDA, and EG methods with respect to the gradient query counts.

Table 4: The CPU times (in seconds) of the algorithms to converge to an ε -accurate Nash equilibrium under various transaction fees. The error bars indicate 2-sigma variations across the independent runs with 10 randomly generated matrices.

Methods \ Transaction fee ρ	0.00%	0.03%	0.06%	0.09%	0.12%	0.15%	0.18%
ICL (Algorithm 1)	19 ± 0	49 ± 0	93 ± 0	142 ± 1	167 ± 0	247 ± 2	264 ± 3
OGDA (Popov, 1980)	186 ± 1	185 ± 1	185 ± 1	185 ± 0	185 ± 1	185 ± 1	186 ± 1
EG (Korpelevich, 1976)	258 ± 1	256 ± 2	258 ± 2	257 ± 3	257 ± 2	257 ± 2	257 ± 2

G.3 ADDITIONAL RUNS UNDER DIFFERENT PARAMETER SETTING

In this section, we run additional numerical experiments under different parameter setting. We change the parameter $\nu = 0.01$, and we vary the transaction fee δ from $\{0.0\%, 0.3\%, \dots, 1.8\%\}$. We keep the other parameter settings unchanged.

The results, summarized in Table 5, demonstrate that ICL requires fewer gradient queries to converge to an ε -accurate Nash equilibrium when the transaction fee ρ is below 1.2%. This empirical observation aligns with our theoretical prediction in Example 1, which suggests that ICL converges faster when $\rho \|\text{abs}(\mathbf{M})\| \ll \sqrt{\mu\nu} = 10\%$.

Table 5: Gradient query counts to converge to an ε -accurate Nash equilibrium under various transaction fees. Error bars indicate 2-sigma variations across the independent runs with 10 randomly generated matrices.

Methods \ Transaction fee ρ	0.0%	0.3%	0.6%	0.9%	1.2%	1.5%	1.8%
ICL (Algorithm 1)	924 ± 0	924 ± 0	824 ± 0	1030 ± 0	1236 ± 0	1648 ± 0	2060 ± 0
OGDA (Popov, 1980)	1364 ± 0	1364 ± 0	1364 ± 0	1361 ± 4	1359 ± 5	1353 ± 6	1350 ± 6
EG (Korpelevich, 1976)	1848 ± 0	1848 ± 0	1848 ± 0	1848 ± 0	1848 ± 0	1848 ± 0	1848 ± 0

We also observe that the CPU times in this setting are within 5 seconds for all independent runs.