
Optimistic Mirror Descent Either Converges to Nash or to Strong Coarse Correlated Equilibria in Bimatrix Games

Ioannis Anagnostides
Carnegie Mellon University
Pittsburgh, PA 15213
ianagnos@cs.cmu.edu

Gabriele Farina
Carnegie Mellon University
Pittsburgh, PA 15213
gfarina@cs.cmu.edu

Ioannis Panageas
University of California Irvine
Irvine, CA 92697
ipanagea@ics.uci.edu

Tuomas Sandholm
Carnegie Mellon University
Strategy Robot, Inc.
Optimized Markets, Inc.
Strategic Machine, Inc.
Pittsburgh, PA 15213
sandholm@cs.cmu.edu

Abstract

We show that, for any sufficiently small fixed $\epsilon > 0$, when both players in a *general-sum two-player (bimatrix) game* employ *optimistic mirror descent (OMD)* with smooth regularization, learning rate $\eta = O(\epsilon^2)$ and $T = \Omega(\text{poly}(1/\epsilon))$ repetitions, either the dynamics reach an ϵ -approximate Nash equilibrium (NE), or the average correlated distribution of play is an $\Omega(\text{poly}(\epsilon))$ -strong coarse correlated equilibrium (CCE): any possible unilateral deviation does not only leave the player worse, but will decrease its utility by $\Omega(\text{poly}(\epsilon))$. As an immediate consequence, when the iterates of OMD are bounded away from being Nash equilibria in a bimatrix game, we guarantee convergence to an *exact* CCE after only $O(1)$ iterations. Our results reveal that uncoupled no-regret learning algorithms can converge to CCE in general-sum games remarkably faster than to NE in, for example, zero-sum games. To establish this, we show that when OMD does not reach arbitrarily close to a NE, the (cumulative) regret of both players is not only negative, but decays linearly with time. Given that regret is the canonical measure of performance in online learning, our results suggest that cycling behavior of no-regret learning algorithms in games can be justified in terms of efficiency.

1 Introduction

The last few years have seen a tremendous amount of progress in computational game solving, witnessed over a series of breakthrough results on benchmark applications in AI [Bowling et al., 2015, Brown and Sandholm, 2017, Moravčík et al., 2017, Silver et al., 2016, Vinyals et al., 2019]. Most of these advances rely on algorithms for approximating a *Nash equilibrium (NE)* [Nash, 1950] in *two-player zero-sum games*. Indeed, in that regime it is by now well-understood how to compute a NE at scale. However, many real-world interactions are not zero-sum, but instead have *general-sum utilities*, and often more than two players. In such settings, Nash equilibria suffer from several drawbacks. First, finding even an approximate NE is *computationally intractable* [Daskalakis et al., 2008, Etessami and Yannakakis, 2007, Chen et al., 2009, Rubinstein, 2016]—subject to well-believed complexity-theoretic assumptions. Furthermore, even if we were to reach one, NE outcomes can be dramatically more inefficient in terms of the *social welfare* compared to other more permissive equilibrium concepts [Moulin and Vial, 1978]. Finally, NE suffer from *equilibrium selection* issues:

there can be a multitude of equilibria, and an equilibrium strategy may perform poorly against the “wrong” equilibrium strategies [Harsanyi et al., 1988, Harsanyi, 1995, Matsui and Matsuyama, 1995], thereby necessitating some form of communication between the players.

A competing notion of rationality is Aumann’s concept of *correlated equilibrium (CE)* [Aumann, 1974], generalizing Nash’s original concept. Unlike NE, a correlated equilibrium can be computed *exactly* in polynomial time [Papadimitriou and Roughgarden, 2008, Jiang and Leyton-Brown, 2015]. Further, CE can arise from simple *uncoupled learning dynamics*, overcoming the often unreasonable assumption that players have perfect knowledge over the game’s utilities. As such, correlated equilibria constitute a much more plausible outcome under independent agents with *bounded rationality*. Indeed, it is folklore that when all players in a general game employ a *no-regret learning algorithm* [Hart and Mas-Colell, 2000], the average correlated distribution of play converges to a *coarse correlated equilibrium (CCE)*—a further relaxation of CE [Moulin and Vial, 1978].

Now as it happens, there are specific classes of games, such as *strictly competitive games* [Adler et al., 2009] or *constant-sum polymatrix games* [Daskalakis and Papadimitriou, 2009a, Cai and Daskalakis, 2011, Cai et al., 2016], for which CCE “collapse” to NE. In fact, although computing Nash equilibria in such games is amenable to linear programming [Adler, 2013], the state of the art algorithms are based on uncoupled learning procedures, for reasons mostly relating to scalability. Nevertheless, such algorithms require $\Omega(\text{poly}(1/\epsilon))$ iterations to reach an ϵ -approximate Nash equilibrium [Daskalakis et al., 2011], meaning that the convergence is slow, particularly in the high precision regime—both in theory and in practice. Furthermore, prior literature treats convergence to NE in, for example, zero-sum games analogously to convergence to CCE in general-sum games. While this unifying treatment—which is an artifact of the no-regret framework—may seem compelling at first glance, it is unclear whether CCE share similar convergence properties to NE. Indeed, as we have alluded to, those equilibrium concepts are fundamentally different (in general games).

Our primary contribution is to challenge this traditional framework, establishing that *convergence to CCE can be remarkably faster than the convergence to NE*; this fundamental difference is showcased in Figure 1. In fact, we show that the only obstacle for converging *exactly* after only $O(1)$ iterations to a CCE in general-sum two-player games is reaching arbitrarily close to a NE. Our results also reveal an intriguing complementarity: the farther the dynamics are from NE, the faster the guarantee of convergence to CCE.

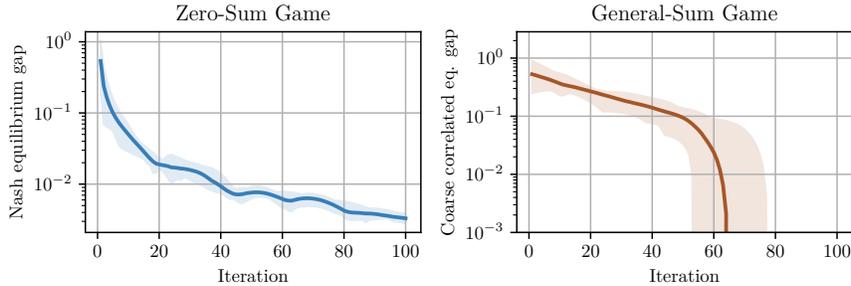


Figure 1: Convergence of uncoupled no-regret dynamics to NE in a zero-sum game (left image) versus convergence to CCE in a general-sum game (right image); further details are provided in Section 4.1.

1.1 Our Contributions

We study uncoupled no-regret learning dynamics in the fundamental class of two-player general-sum games (a.k.a. *bimatrix games*). Specifically, we focus on *optimistic mirror descent (OMD)* [Chiang et al., 2012, Rakhlin and Sridharan, 2013a], a variant of the standard mirror descent method which incorporates a *prediction* into the optimization step. Before we state our main result, let us introduce *strong CCE*, a *refinement* of CCE for which any unilateral deviation from a player is not only worse, but *decreases* its utility by an additive $\epsilon \geq 0$ (see Definition 2.4 for a formal description). In particular, we clarify that any strong CCE is (trivially) an *exact CCE*. We are now ready to state our main theorem.¹

¹For simplicity in the exposition of our results, we use the $O(\cdot)$ and $\Omega(\cdot)$ notation in the main body to suppress parameters that depend on the underlying game; precise statements are given in Appendix A.

Theorem 1.1 (Abridged; Full Version in Corollary A.2). *Fix any sufficiently small $\epsilon > 0$, and suppose that both players in a bimatrix game employ OMD with learning rate $\eta = O(\epsilon^2)$ and smooth regularizer for $T = \Omega(\text{poly}(1/\epsilon))$ repetitions. Then,*

- *Either the dynamics reach an ϵ -approximate Nash equilibrium;*
- *Or, otherwise, the average correlated distribution of play is an $\Omega(\text{poly}(\epsilon))$ -strong CCE.*

Here, the convergence to Nash equilibrium is implied in a *last-iterate* sense (and not a typical time-average): *i.e.*, there exists a time $t \in \llbracket T \rrbracket$ for which the dynamics reach an ϵ -approximate Nash equilibrium. Of course, in light of well-established impossibility results, OMD is certainly not going to reach an ϵ -approximate NE—for a sufficiently small *constant* $\epsilon > 0$ [Rubinstein, 2016]—in *every* bimatrix game. As a result, an immediate interesting implication of Theorem 1.1 is that, by setting ϵ to be sufficiently small, OMD either yields the best known polynomial-time approximation for Nash equilibria in bimatrix games, or, otherwise, $O(1)$ iterations suffice to obtain a strong CCE. It is intriguing that the farther the dynamics are from yielding a Nash equilibrium, the *stronger* the CCE predicted by Theorem 1.1. Indeed, the only impediment for reaching a strong CCE after only $O(1)$ iterations lies in reaching very close to a Nash equilibrium.²

To establish Theorem 1.1, we prove that when the iterates of OMD are bounded away from being Nash equilibria, the regret of *both* players is not only negative, but *decreases linearly* over time (Theorem 3.4); see Figure 3 for an illustration. In this way, our approach represents a substantial departure from prior work which endeavored to characterize no-regret learning dynamics from a unifying standpoint for all possible games. Furthermore, besides the folklore connection of regret with CCE (which allows us to establish Theorem 1.1), regret is the canonical measure of performance in online learning. Hence, our results can be construed as follows: when the dynamics do not reach a Nash equilibrium, we obtain remarkably stronger performance guarantees. This seems to suggest that *cycling behavior in games may not be a bug, but a feature.*³

From a technical standpoint, a central ingredient in our proof is a remarkable property—discovered in prior works—coined as the RVU bound [Syrgkanis et al., 2015, Rakhlin and Sridharan, 2013a]. Specifically, our argument crucially leverages the rather enigmatic negative term in the RVU bound (Proposition 2.1), which was treated as a “cancellation factor” in prior works. This allows us to establish (conditionally) linear decay for the regrets in Theorem 3.4. To this end, an important component of our proof is to connect the path lengths of the two players in Lemma 3.3 through a potential-type argument; notably, such properties break in arbitrary games with more than two players (see Remark 3.7). (We give a detailed sketch of our proof of Theorem 1.1 in Section 3.)

Importantly, our techniques apply under arbitrary convex and compact strategy spaces, thereby allowing for a direct extension to, for example, *normal-form coarse correlated equilibria* in *extensive-form games (EFGs)* [Moulin and Vial, 1978]. Another compelling aspect of our result is that both players employ a *constant* learning rate, a feature which has been extensively motivated in prior works (*e.g.*, see [Bailey and Piliouras, 2019, Piliouras et al., 2021, Golowich et al., 2020a]). Besides the improved performance guarantees observed in practice under a time-invariant learning rate, it arguably induces a more natural behavior from an economic standpoint. Finally, we corroborate our theoretical findings through experiments on standard benchmark extensive-form games in Section 4.2.

1.2 Related Work

Nash Equilibria in Bimatrix Games Bimatrix games is one of the most fundamental and well-studied classes of games. In a celebrated series of works, it was shown that computing a NE in bimatrix games does not admit a fully polynomial-time approximation scheme (FPTAS), unless every problem in PPAD [Papadimitriou, 1994] can be solved in polynomial time [Chen et al., 2009, Daskalakis et al.,

²Closeness to NE is always implied in terms of the best response gap, not the distance to the set of NE.

³To explain this point further, it is important to connect *last-iterate convergence* with NE. Namely, if all players employ no-regret learning algorithms and the dynamics converge pointwise, then it is easy to see that the limit point has to be a NE. On the other hand, correlated equilibrium concepts are, by definition, incompatible with last-iterate convergence—at least under uncoupled dynamics. Indeed, correlation inherently *requires* cycling behavior. In light of this, there are games for which *stability can be at odds with efficiency* since NE can be dramatically more inefficient than correlation schemes [Moulin and Vial, 1978]. We also refer to our example in Section 4 where this phenomenon becomes clear; *c.f.*, see [Kleinberg et al., 2011].

2009]. In light of this intractability result, there has been a substantial amount of attention in deriving polynomial-time approximation algorithms for NE in bimatrix games [Daskalakis et al., 2006a, 2007, Bosse et al., 2010, Tsaknakis and Spirakis, 2008, Kontogiannis and Spirakis, 2010, Daskalakis and Papadimitriou, 2009b, Deligkas et al., 2022a,b]. In fact, even computing a sufficiently small *constant* approximation essentially requires quasi-polynomial time [Rubinstein, 2016, Kothari and Mehta, 2018], subject to the exponential-time hypothesis for PPAD [Babichenko et al., 2016], matching the seminal QPTAS of Lipton et al. [2003]. It is also worth stressing that, in terms of computing Nash equilibria, there is a polynomial-time reduction from any game with a *succinct representation* to a bimatrix game [Daskalakis et al., 2006b]; this illustrates the generality of bimatrix games.

Near-Optimal Regret in Games In a pioneering work by Daskalakis et al. [2011], it was shown that there exist uncoupled no-regret learning dynamics such that when employed by *both* players in a *zero-sum* game, the cumulative regret incurred by each player is bounded by $O(\log T)$, substantially improving over the $\Omega(\sqrt{T})$ lower bound under *adversarial utilities*. Thereafter, there has been a considerable amount of effort in simplifying and extending the previous results along several lines [Rakhlin and Sridharan, 2013a, Syrgkanis et al., 2015, Chen and Peng, 2020, Foster et al., 2016, Wei and Luo, 2018, Farina et al., 2019a, Daskalakis and Golowich, 2022, Daskalakis et al., 2021, Piliouras et al., 2021]. This line of work was culminated in a recent result due to Daskalakis et al. [2021], establishing $O(\text{polylog}(T))$ individual regret when all players employ an *optimistic* variant of *multiplicative weights update* in general games, nearly-matching the lower bound in [Daskalakis et al., 2011]. Nevertheless, the optimality here is only *existential*: there *exist* games—in fact, zero-sum—for which the guarantee in [Daskalakis et al., 2021] is essentially unimprovable. But we argue that this is not a good enough justification for considering the problem of no-regret learning in games settled. As we show in this paper, substantially more refined guarantees are possible beyond zero-sum games.

Last-Iterate Convergence A folklore phenomenon in the theory of learning in games is that broad families of no-regret algorithms exhibit *cyclic* or even *chaotic behavior* [Sato et al., 2002, Sandholm, 2010, Mertikopoulos et al., 2018, Andrade et al., 2021, Cheung and Piliouras, 2020, Cheung and Tao, 2021], even in low-dimensional zero-sum games. A compelling approach to ameliorate this problem was proposed by Daskalakis et al. [2018], showing that an *optimistic* variant of gradient descent guarantees last-iterate convergence in unconstrained bilinear (zero-sum) games. Thereafter, there has been a tremendous amount of interest in strengthening and extending their result in different regimes; for a highly incomplete list, we refer to [Daskalakis and Panageas, 2019, 2018, Wei et al., 2021a, Golowich et al., 2020a,b, Hsieh et al., 2021, Zhou et al., 2017, Lin et al., 2020, Mertikopoulos et al., 2019, Liang and Stokes, 2019, Zhang and Yu, 2020, Mokhtari et al., 2020, Daskalakis et al., 2020, Wei et al., 2021b, Anagnostides and Panageas, 2021], and references therein. The precursor of our main result is the recent paper of Anagnostides et al. [2022a] that obtained similar results but for the *sum* of the players’ regrets, which does not have implications in terms of convergence to CCE.

2 Preliminaries

Notation and Conventions We let $\mathbb{N} := \{1, 2, \dots\}$ be the set of natural numbers. The (discrete) time index is represented exclusively via the variable t . We also use the shorthand notation $\llbracket p \rrbracket := \{1, 2, \dots, p\}$. We let $\|\cdot\|$ be a norm on \mathbb{R}^d , for some $d \in \mathbb{N}$, and $\|\cdot\|_*$ be the *dual norm* of $\|\cdot\|$; namely, $\|\mathbf{v}\|_* := \sup_{\|\mathbf{u}\| \leq 1} \langle \mathbf{u}, \mathbf{v} \rangle$. For a (nonempty) convex and compact set $\mathcal{X} \subseteq \mathbb{R}^d$, we denote by $\Omega_{\mathcal{X}} := \max_{\mathbf{x}, \hat{\mathbf{x}} \in \mathcal{X}} \|\mathbf{x} - \hat{\mathbf{x}}\|$ the *diameter* of \mathcal{X} with respect to $\|\cdot\|$; to lighten the notation, the underlying norm will be implicit in our notation. Moreover, we let $\|\mathcal{X}\| := \max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|$. For a matrix \mathbf{A} we let $\|\mathbf{A}\|_{\text{op}}$ be its *operator norm*: $\|\mathbf{A}\|_{\text{op}} := \sup_{\|\mathbf{u}\| \leq 1} \|\mathbf{A}\mathbf{u}\|_*$. We point out that $\|\mathbf{A}\|_{\text{op}} \neq 0$ when $\mathbf{A} \neq \mathbf{0}$. Finally, for a finite nonempty set \mathcal{A} , we let $\Delta(\mathcal{A}) := \left\{ \mathbf{x} \in \mathbb{R}_{\geq 0}^{|\mathcal{A}|} : \sum_{a \in \mathcal{A}} \mathbf{x}(a) = 1 \right\}$ be the *probability simplex* over \mathcal{A} , where $\mathbf{x}(a)$ is the coordinate of \mathbf{x} corresponding to $a \in \mathcal{A}$.

2.1 Online Learning and Regret

Let \mathcal{X} be a nonempty, convex and compact subset of \mathbb{R}^d , for some $d \in \mathbb{N}$. In the *online learning framework* the *learner* commits to a *strategy* $\mathbf{x}^{(t)} \in \mathcal{X}$ at every iteration $t \in \mathbb{N}$. Then, the learner receives as feedback from the environment a (linear) *utility function* $u^{(t)} : \mathcal{X} \ni \mathbf{x} \mapsto \langle \mathbf{x}, \mathbf{u}^{(t)} \rangle$, with $\mathbf{u}^{(t)} \in \mathbb{R}^d$, so that the utility received at time t is given by $\langle \mathbf{x}^{(t)}, \mathbf{u}^{(t)} \rangle$. Based on that feedback, the

learner may adapt its next strategy. The framework is *online* in the sense that no information about future utilities is available. The canonical objective in this framework is to minimize the *cumulative external regret* (or simply *regret*), defined for a *time horizon* $T \in \mathbb{N}$ as follows.

$$\text{Reg}^T := \max_{\mathbf{x}^* \in \mathcal{X}} \left\{ \sum_{t=1}^T \langle \mathbf{x}^*, \mathbf{u}^{(t)} \rangle \right\} - \sum_{t=1}^T \langle \mathbf{x}^{(t)}, \mathbf{u}^{(t)} \rangle, \quad (1)$$

That is, the performance is measured in terms of the optimal *fixed strategy in hindsight*.

2.2 Optimistic Regret Minimization

To leverage the additional structure in more benign environments, several *predictive* algorithms have been recently introduced [Hazan and Kale, 2011, Chiang et al., 2012, 2013, Rakhlin and Sridharan, 2013b,a, Syrgkanis et al., 2015, Farina et al., 2021]. In this paper we employ a predictive variant of MD, known as *optimistic mirror descent* [Chiang et al., 2012, Rakhlin and Sridharan, 2013a].

Optimistic Mirror Descent Let \mathcal{R} be a 1-strongly convex continuously differentiable regularizer (or *distance generating function (DGF)*) with respect to a norm $\|\cdot\|$ on \mathbb{R}^d . We say that \mathcal{R} is *G-smooth*, with $G > 0$, if for any $\mathbf{x}, \hat{\mathbf{x}} \in \mathcal{X}$,

$$\|\nabla \mathcal{R}(\mathbf{x}) - \nabla \mathcal{R}(\hat{\mathbf{x}})\|_* \leq G \|\mathbf{x} - \hat{\mathbf{x}}\|. \quad (2)$$

For example, the *Euclidean regularizer* $\mathcal{R}(\mathbf{x}) := \frac{1}{2} \|\mathbf{x}\|_2^2$, which is 1-strongly convex with respect to the Euclidean norm $\|\cdot\|_2$, trivially satisfies the smoothness condition of (2) with $G = 1$; we recall that $\|\cdot\|_2$ is self-dual. Moreover, we let $D_{\mathcal{R}}(\mathbf{x} \parallel \hat{\mathbf{x}}) := \mathcal{R}(\mathbf{x}) - \mathcal{R}(\hat{\mathbf{x}}) - \langle \nabla \mathcal{R}(\hat{\mathbf{x}}), \mathbf{x} - \hat{\mathbf{x}} \rangle$ be the *Bregman divergence* induced by \mathcal{R} [Rockafellar, 1970]. *Optimistic mirror descent (OMD)*⁴ is parameterized by a (dynamic) *prediction vector* $\mathbf{m}^{(t)} \in \mathbb{R}^d$, for every time $t \in \mathbb{N}$, and a *learning rate* $\eta > 0$, so that its update rule takes the following form for $t \in \mathbb{N}$:

$$\begin{aligned} \mathbf{x}^{(t)} &:= \arg \max_{\mathbf{x} \in \mathcal{X}} \left\{ \langle \mathbf{x}, \mathbf{m}^{(t)} \rangle - \frac{1}{\eta} D_{\mathcal{R}}(\mathbf{x} \parallel \hat{\mathbf{x}}^{(t-1)}) \right\}; \\ \hat{\mathbf{x}}^{(t)} &:= \arg \max_{\hat{\mathbf{x}} \in \mathcal{X}} \left\{ \langle \hat{\mathbf{x}}, \mathbf{u}^{(t)} \rangle - \frac{1}{\eta} D_{\mathcal{R}}(\hat{\mathbf{x}} \parallel \hat{\mathbf{x}}^{(t-1)}) \right\}. \end{aligned} \quad (\text{OMD})$$

Further, $\hat{\mathbf{x}}^{(0)} := \arg \min_{\hat{\mathbf{x}} \in \mathcal{X}} \mathcal{R}(\hat{\mathbf{x}})$; for convenience, we also let $\mathbf{x}^{(0)} := \hat{\mathbf{x}}^{(0)}$. We will refer to $\hat{\mathbf{x}}^{(0)}, \hat{\mathbf{x}}^{(1)}, \dots$ as the *secondary* sequence of OMD, while $\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \dots$ is the *primary* sequence. We also let $\Omega_{\mathcal{R}} := \max_{\mathbf{x} \in \mathcal{X}} D_{\mathcal{R}}(\mathbf{x} \parallel \hat{\mathbf{x}}^{(0)})$ denote the \mathcal{R} -diameter of \mathcal{X} . As usual, we consider the “one-recency bias” prediction mechanism [Syrgkanis et al., 2015], wherein $\mathbf{m}^{(t)} := \mathbf{u}^{(t-1)}$ for $t \in \mathbb{N}$. (To make $\mathbf{u}^{(0)}$ well-defined in games, each player initially receives the utility corresponding to the players’ strategies at $t = 0$; this is only made for convenience in the analysis.)

An important special case of (OMD) arises under the Euclidean regularizer $\mathcal{R}(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|_2^2$, in which case we refer to the dynamics as *optimistic gradient descent (OGD)*:

$$\begin{aligned} \mathbf{x}^{(t)} &:= \Pi_{\mathcal{X}} \left(\hat{\mathbf{x}}^{(t-1)} + \eta \mathbf{m}^{(t)} \right); \\ \hat{\mathbf{x}}^{(t)} &:= \Pi_{\mathcal{X}} \left(\hat{\mathbf{x}}^{(t-1)} + \eta \mathbf{u}^{(t)} \right). \end{aligned} \quad (\text{OGD})$$

Here, $\Pi_{\mathcal{X}}(\cdot)$ stands for the Euclidean projection to the set \mathcal{X} .

A fundamental property crystallized in [Syrgkanis et al., 2015], building on [Rakhlin and Sridharan, 2013a], is the *regret bounded by variation in utilities (RVU)*. For our purposes, we will employ a refinement of the RVU bound for OMD, which follows from [Rakhlin and Sridharan, 2013a, Lemma 1 in the Full Version]—this can be extracted en route to the proof of [Syrgkanis et al., 2015, Theorem 18].

Proposition 2.1 ([Rakhlin and Sridharan, 2013a, Syrgkanis et al., 2015]). *The regret of (OMD) with learning rate $\eta > 0$ can be bounded as*

$$\text{Reg}^T \leq \frac{\Omega_{\mathcal{R}}}{\eta} + \eta \sum_{t=1}^T \|\mathbf{u}^{(t)} - \mathbf{u}^{(t-1)}\|_*^2 - \frac{1}{4\eta} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t-1)}\|^2 \right).$$

⁴To avoid any confusion, we point out that OMD sometimes stands for *online* mirror descent in the literature.

For convenience, we will use a shorthand notation for the *second-order path length*:

$$\Sigma_{\mathcal{X}}^T := \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right). \quad (3)$$

2.3 No-Regret Learning and Coarse Correlated Equilibria

A folklore connection ensures that when all players in a general game employ a no-regret learning algorithm, the average correlated distribution of play converges to a *coarse correlated equilibrium* (CCE) [Aumann, 1974, Moulin and Vial, 1978]. Formally, let us first introduce the notion of a CCE in general *normal-form games* (NFGs). To this end, we consider a set of p players $\llbracket p \rrbracket$, with each player i having a set of available actions \mathcal{A}_i . The *utility* of player $i \in \llbracket p \rrbracket$ is a function $u_i : \times_{i=1}^p \mathcal{A}_i \ni \mathbf{a} \mapsto \mathbb{R}$ indicating the utility received by player i under the action profile $\mathbf{a} = (a_1, \dots, a_p)$.

Definition 2.2 (Approximate Coarse Correlated Equilibrium). A distribution $\boldsymbol{\mu}$ over the set $\times_{i=1}^p \mathcal{A}_i$ is an ϵ -*approximate coarse correlated equilibrium*, with $\epsilon \geq 0$, if for any player $i \in \llbracket p \rrbracket$ and any unilateral deviation $a'_i \in \mathcal{A}_i$,

$$\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\mu}} [u_i(a'_i, \mathbf{a}_{-i})] \leq \mathbb{E}_{\mathbf{a} \sim \boldsymbol{\mu}} [u_i(\mathbf{a})] + \epsilon.$$

We are now ready to state the fundamental theorem connecting no-regret learning with CCE; for completeness, we include the short proof in Appendix A.

Theorem 2.3 (Folklore). *Suppose that every player $i \in \llbracket p \rrbracket$ employs a no-regret learning algorithm with regret Reg_i^T up to time $T \in \mathbb{N}$. Moreover, let $\boldsymbol{\mu}^{(t)} := \mathbf{x}_1^{(t)} \otimes \dots \otimes \mathbf{x}_p^{(t)}$ be the correlated distribution of play at time $t \in \llbracket T \rrbracket$, and $\bar{\boldsymbol{\mu}} := \frac{1}{T} \sum_{t=1}^T \boldsymbol{\mu}^{(t)}$ be the average correlated distribution of play up to time T . Then,*

$$\mathbb{E}_{\mathbf{a} \sim \bar{\boldsymbol{\mu}}} [u_i(a'_i, \mathbf{a}_{-i})] \leq \mathbb{E}_{\mathbf{a} \sim \bar{\boldsymbol{\mu}}} [u_i(\mathbf{a})] + \frac{1}{T} \max_{i \in \llbracket p \rrbracket} \text{Reg}_i^T.$$

We also introduce a refinement of CCE which we refer to as *strong CCE*:

Definition 2.4 (Strong Coarse Correlated Equilibrium). A probability distribution $\boldsymbol{\mu}$ over the set $\times_{i=1}^p \mathcal{A}_i$ is an ϵ -*strong coarse correlated equilibrium*, with $\epsilon \geq 0$, if for any player $i \in \llbracket p \rrbracket$ and any unilateral deviation $a'_i \in \mathcal{A}_i$,

$$\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\mu}} [u_i(a'_i, \mathbf{a}_{-i})] \leq \mathbb{E}_{\mathbf{a} \sim \boldsymbol{\mu}} [u_i(\mathbf{a})] - \epsilon.$$

In a strong CCE any deviation from the mediator's recommendation is not only worse, but can decrease the player's utility by a significant amount. In that sense, a strong CCE can be a much more self-enforcing equilibrium outcome, and, as such, arguably more likely to occur. A strong CCE (Definition 2.4) can be thought of as a standard CCE (Definition 2.2) but with a "negative approximation". We include a discussion and an illustration of strong CCE in Section 4.1.

2.4 Bimatrix Games

A *bimatrix game* involves two players. Each player has a set of strategies $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$, respectively. The (expected) *payoffs* of each player under strategies $(\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$ is given by the bilinear forms $\mathbf{x}^\top \mathbf{A} \mathbf{y}$ and $\mathbf{x}^\top \mathbf{B} \mathbf{y}$, respectively. Here, $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times m}$ are the payoff matrices of the game.⁵ As an example, the special case where $\mathcal{X} = \Delta(\mathcal{A}_\mathcal{X})$ and $\mathcal{Y} = \Delta(\mathcal{A}_\mathcal{Y})$ corresponds to normal-form games, but our current formulation captures *extensive-form games* as well. By convention, we will refer to the two players as player \mathcal{X} and player \mathcal{Y} respectively. Furthermore, the underlying bimatrix game will be referred to as (\mathbf{A}, \mathbf{B}) , without specifying the strategy sets.

Definition 2.5 (Approximate Nash Equilibrium). A pair of strategies $(\mathbf{x}^*, \mathbf{y}^*) \in \mathcal{X} \times \mathcal{Y}$ is an ϵ -*approximate Nash equilibrium* of (\mathbf{A}, \mathbf{B}) , for $\epsilon \geq 0$, if for any $(\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$,

$$\begin{aligned} \mathbf{x}^\top \mathbf{A} \mathbf{y}^* &\leq (\mathbf{x}^*)^\top \mathbf{A} \mathbf{y}^* + \epsilon; \\ (\mathbf{x}^*)^\top \mathbf{B} \mathbf{y} &\leq (\mathbf{x}^*)^\top \mathbf{B} \mathbf{y}^* + \epsilon. \end{aligned} \quad (4)$$

To make Definition 2.5 meaningful, some normalization has to be imposed on the utilities. Here, we will assume that $\max_{\mathbf{y} \in \mathcal{Y}} \|\mathbf{A} \mathbf{y}\|_* \leq 1$ and $\max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{B}^\top \mathbf{x}\|_* \leq 1$.

⁵We assume that $\mathbf{A} \neq \mathbf{0}$ and $\mathbf{B} \neq \mathbf{0}$. In the contrary case Theorem 1.1 follows trivially.

3 Main Result

In this section we sketch the main ingredients required for the proof of Theorem 1.1; all the proofs are deferred to Appendix A. Moreover, a detailed version of Theorem 1.1 for normal-form games under Euclidean regularization is given in Corollary A.2. In the sequel, we assume that players \mathcal{X} and \mathcal{Y} employ regularizers $\mathcal{R}_{\mathcal{X}}$ and $\mathcal{R}_{\mathcal{Y}}$, respectively, so that the regret of each player enjoys an RVU bound with respect to the same pair of dual norms $(\|\cdot\|, \|\cdot\|_*)$. The first step is to cast the refined RVU bound of Proposition 2.1 for bimatrix games.

Corollary 3.1. *Suppose that both players employ (OMD) with learning rate $\eta > 0$. Then,*

$$\begin{aligned} \text{Reg}_{\mathcal{X}}^T &\leq \frac{\Omega_{\mathcal{R}_{\mathcal{X}}}}{\eta} + \eta \|\mathbf{A}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 - \frac{1}{4\eta} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right); \\ \text{Reg}_{\mathcal{Y}}^T &\leq \frac{\Omega_{\mathcal{R}_{\mathcal{Y}}}}{\eta} + \eta \|\mathbf{B}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\|^2 - \frac{1}{4\eta} \sum_{t=1}^T \left(\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t)}\|^2 + \|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 \right). \end{aligned}$$

The next critical step consists of showing that approximate fixed points of (OMD) under smooth regularization—in the sense of (2)—correspond to approximate Nash equilibria (recall Definition 2.5) of the underlying bimatrix game, as we formalize below.

Proposition 3.2 (Approximate Fixed Points of OMD). *Consider a bimatrix game (\mathbf{A}, \mathbf{B}) , and suppose that both players employ (OMD) with learning rate $\eta > 0$ and a G -smooth regularizer. Then, if $\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|, \|\widehat{\mathbf{x}}^{(t)} - \mathbf{x}^{(t)}\| \leq \epsilon\eta$ and $\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|, \|\widehat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}\| \leq \epsilon\eta$, the pair $(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})$ is a $(2\epsilon G \max\{\Omega_{\mathcal{X}}, \Omega_{\mathcal{Y}}\} + \epsilon\eta)$ -approximate Nash equilibrium of (\mathbf{A}, \mathbf{B}) .*

Indeed, we show that when the iterates of (OMD) do not change by much (relatively to the learning rate), each player is approximately best responding to the observed utility. The following ingredient is where we rely on the two-player aspect of the game. On a high level, we show that when the strategies of one of the players change fast over time, the other player ought to be “moving” rapidly as well.

Lemma 3.3. *Suppose that both players in a bimatrix game (\mathbf{A}, \mathbf{B}) employ (OMD) with learning rate $\eta > 0$. Then, for any $T \in \mathbb{N}$,*

$$\begin{aligned} \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\| &\geq \frac{1}{2\eta \|\mathcal{X}\| \|\mathbf{A}\|_{\text{op}}} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 + \|\widehat{\mathbf{x}}^{(t)} - \mathbf{x}^{(t)}\|^2 \right) - \frac{2}{\|\mathbf{A}\|_{\text{op}}}; \\ \sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\| &\geq \frac{1}{2\eta \|\mathcal{Y}\| \|\mathbf{B}\|_{\text{op}}} \sum_{t=1}^T \left(\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 + \|\widehat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}\|^2 \right) - \frac{2}{\|\mathbf{B}\|_{\text{op}}}. \end{aligned}$$

The intuition is that as long as one player is “moving” substantially faster than the other player, its utility will be monotonically increasing as the (repeated) game progresses. But this cannot occur for too long—by a potential argument—since the utility of each player is bounded. We are now ready to state the main technical theorem.

Theorem 3.4 (Linear Decay of Regret; Full Version in Theorem A.1). *Suppose that both players in a bimatrix game (\mathbf{A}, \mathbf{B}) employ (OMD) with smooth regularizer, learning rate $\eta = O(\epsilon^2)$ and $T = \Omega\left(\frac{1}{\epsilon^4 \eta^2}\right)$, for a sufficiently small fixed $\epsilon > 0$. Then, if the dynamics do not reach an $O(\epsilon)$ -approximate NE, then*

$$\max\{\text{Reg}_{\mathcal{X}}^T, \text{Reg}_{\mathcal{Y}}^T\} \leq -\Omega(\epsilon^4 \eta T).$$

By virtue of Theorem 2.3 and Proposition 3.2, Theorem 3.4 immediately implies Theorem 1.1. An illustration of the linear decay of regret in a bimatrix game can be seen in Figure 3. Before we sketch the proof of Theorem 3.4, let us point out the following useful lemma.

Lemma 3.5 (Stability of OMD). *Suppose that both players employ (OMD) with learning rate $\eta > 0$. Then, for any $t \in \mathbb{N}$,*

$$\begin{aligned} \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\| &\leq 3\eta; \\ \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\| &\leq 3\eta. \end{aligned}$$

Sketch Proof of Theorem 3.4. When the iterates of (OMD) are $\Omega(\epsilon)$ from being a Nash equilibrium, Proposition 3.2 implies that $\Sigma_{\mathcal{X}}^T + \Sigma_{\mathcal{Y}}^T = \Omega(\epsilon^2 \eta^2 T)$, where we used the notation of (3). Now suppose that $\Sigma_{\mathcal{X}}^T \geq \Sigma_{\mathcal{Y}}^T$. Then, using Corollary 3.1 we get that $\text{Reg}_{\mathcal{X}}^T = -\Omega(\epsilon^2 \eta T)$. For the regret of player \mathcal{Y} , we first use Lemma 3.3 to obtain that $\Sigma_{\mathcal{Y}}^T = \Omega(\epsilon^4 \eta^2 T)$. Finally, using Corollary 3.1 and Lemma 3.5 we can conclude that $\text{Reg}_{\mathcal{Y}}^T \leq \frac{\Omega_{\mathcal{R}_{\mathcal{Y}}}}{\eta} - \Omega(\epsilon^4 \eta T) = -\Omega(\epsilon^4 \eta T)$. \square

While the dichotomy of Theorem 1.1 is based on whether only a *single* iterate is an ϵ -approximate Nash equilibrium, our techniques also directly give analogous guarantees depending on whether *most*—say 99%—of the iterates are ϵ -approximate Nash equilibria, as we formalize in Corollary A.3.

Remark 3.6 (Extensive-Form Games). Theorem 1.1 has direct implications for *normal-form coarse correlated equilibria (NFCCE)* [Moulin and Vial, 1978] in extensive-form games using the *sequence-form strategy representation* [Romanovskii, 1962, von Stengel, 1996, Koller et al., 1996].

Remark 3.7 (Multiplayer Games). Theorem 1.1 does not extend to arbitrary games with $p \geq 3$. To see this, consider a 3-player game for which the utility of player 3 does not depend on the strategies of the other players. Then, the regret of player 3 will be strictly positive—as long as the initialization differs from the optimal strategy. So, even if the dynamics are far from Nash equilibria, the CCE gap will always be strictly positive. At a superficial level, this issue occurs in games where the strategic interactions form, in some sense, multiple “connected components”. Nevertheless, characterizing the multiplayer games for which Theorem 1.1 holds is an interesting question for the future.

4 Experiments

In this section we provide experiments supporting our theoretical findings. We start by analyzing the CCE and the behavior of (OGD) in a simple bimatrix NFG in Section 4.1, while in Section 4.2 we experiment with several benchmark games used in the EFG-solving literature.

4.1 An Illustrative Example

First, we study the 3×3 bimatrix normal-form game (\mathbf{A}, \mathbf{B}) , where

$$\mathbf{A} := \begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}; \quad \mathbf{B} := \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}. \quad (5)$$

It can be shown that this bimatrix game has a unique Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ such that $\mathbf{x}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and $\mathbf{y}^* = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ [Avis et al., 2010]. Moreover, $(\mathbf{x}^*, \mathbf{y}^*)$ secures a social welfare $\text{SW}(\mathbf{x}^*, \mathbf{y}^*) := (\mathbf{x}^*)^\top \mathbf{A} \mathbf{y}^* + (\mathbf{x}^*)^\top \mathbf{B} \mathbf{y}^* = \frac{1}{4} + \frac{1}{3} \approx 0.5833$. On the other hand, it is easy to see that there exists an exact CCE μ^* such that $\text{SW}(\mu^*) = 1$. In fact, this social welfare is optimal even without any incentive-compatibility constraints. Furthermore, using a linear programming solver, we find that the *strongest* CCE (in the sense of Definition 2.4) has a parameter of roughly 0.2083. The entire landscape of CCE associated with the bimatrix game (5) is illustrated in Figure 2. The blue region corresponds to strong CCE, under which *both* players obtain a high utility. On the other hand, configurations for which one of the players receives low utility are not incentive compatible.

Next, we focus on the behavior of the (OMD) dynamics. We let both players employ Euclidean regularization and learning rate $\eta := 0.1$. The convergence to CCE of the induced (OGD) dynamics is illustrated in Figure 3. In particular, after $T = 1000$ iterations the average

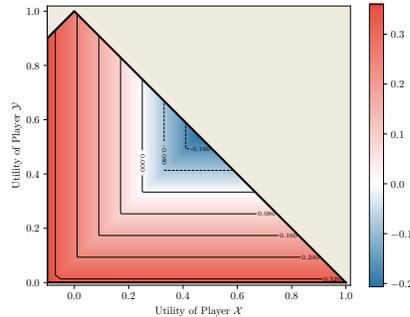


Figure 2: The maximum incentive-compatibility parameter of a CCE which guarantees a given pair of utilities.

correlated distribution of play $\bar{\mu}$ reads (with precision up to 4 decimal places)

$$\bar{\mu} \approx \begin{bmatrix} 0.1594 & 0.1778 & 0.0048 \\ 0.0029 & 0.1614 & 0.1607 \\ 0.1642 & 0.0075 & 0.1613 \end{bmatrix}.$$

This correlated distribution secures a social welfare of $\text{SW}(\bar{\mu}) \approx 0.4793 + 0.5027 = 0.9819$. Thus, $\bar{\mu}$ is near-optimal in terms of the obtained social welfare. As such, it substantially outperforms the efficiency of the Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$. Moreover, we see that $\bar{\mu}$ is (approximately) a 0.1525-strong CCE. Indeed, for player \mathcal{X} the maximum possible utility attainable from a unilateral deviation is roughly 0.3268, compared to 0.4793 obtained under $\bar{\mu}$; for player \mathcal{Y} the maximum utility from a unilateral deviation is roughly 0.3420, compared to 0.5027. It is worth noting that (OGD) does *not* converge to the strongest possible CCE of the game, even under different random initializations. The results described here are robust to different initializations—although under the definition of (OGD) each player should start from the uniform distribution.

Finally, let us elaborate on Figure 1. The left image illustrates the Nash gap of the average strategies of (OGD) (with $\eta := 0.1$) in the zero-sum game $(\mathbf{A}, -\mathbf{A})$, while the right image shows the CCE gap of the average correlated distribution of play in the bimatrix game (\mathbf{A}, \mathbf{B}) , as given in (5).

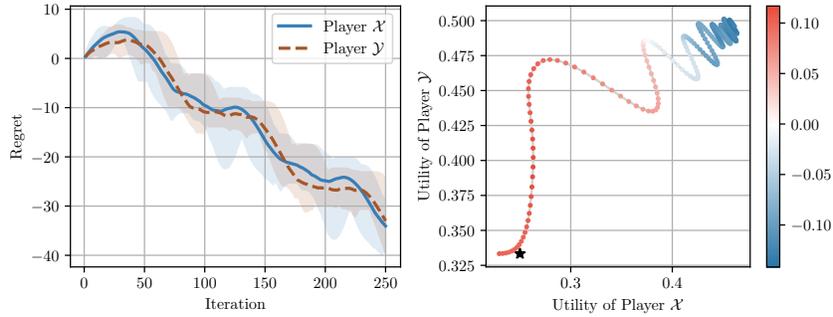


Figure 3: The convergence of (OGD) dynamics to CCE in the bimatrix game (\mathbf{A}, \mathbf{B}) introduced in (5). On the left image we plot the regret of each player under different random initializations; we recall that the maximum average regret is tantamount to the CCE gap. We see that after about 60 iterations both players experience *negative regret*—implying *exact* convergence to CCE. Furthermore, their regret decays linearly over time; this translates to convergence to strong CCE, as illustrated in the right image. More precisely, the color of each point (in the right image) corresponds to the CCE gap at the given iterate. Even when the initialization is “close” to the Nash equilibrium (represented by “*”), the dynamics manage to avoid it in search of more efficient outcomes. In fact, the limit CCE of the dynamics yields near-optimal social welfare.

4.2 Benchmark Games

Next, we illustrate the convergence of (OMD) on several benchmark bimatrix EFGs; namely: (i) *Sheriff* [Farina et al., 2019b]; (ii) *Liar’s Dice* [Lisý et al., 2015]; (iii) *Battleship* [Farina et al., 2019b]; and (iv) *Goofspiel* [Ross, 1971]. A detailed description of the game instances we used for our experiments is included in Appendix B.

We instantiated (OMD) with Euclidean regularization. After a very mild tuning process, we chose for all games a (time-invariant) learning rate of $\eta = (2 \max\{\|\mathbf{A}\|_2, \|\mathbf{B}\|_2\})^{-1}$; here, $\|\cdot\|_2$ stands for the spectral norm of the corresponding matrix. For all games the initialization is chosen so that $\hat{\mathbf{x}}^{(0)} := \arg \min_{\hat{\mathbf{x}} \in \mathcal{X}} \mathcal{R}_{\mathcal{X}}(\hat{\mathbf{x}})$ and $\hat{\mathbf{y}}^{(0)} := \arg \min_{\hat{\mathbf{y}} \in \mathcal{Y}} \mathcal{R}_{\mathcal{Y}}(\hat{\mathbf{y}})$, with the sole exception of Battleship for which that initialization is (virtually) already a Nash equilibrium. In light of this, for Battleship we initialized the dynamics in some arbitrary deterministic strategies; we stress that the conclusions derived here are robust to different initializations. Our results are summarized in Figure 4.

From these benchmark games, only Liar’s Dice is constant-sum—in fact, zero-sum. Hence, as expected, the NE gap of the (OGD) dynamics essentially converges to 0. Perhaps surprisingly, the same appears to hold for both Sheriff and Battleship. On the other hand, the dynamics exhibit a remarkably different behavior in Goofspiel. Indeed, although initially the dynamics appear to gradually converge to a Nash equilibrium, after about 600 iterations the NE gap of the last iterate

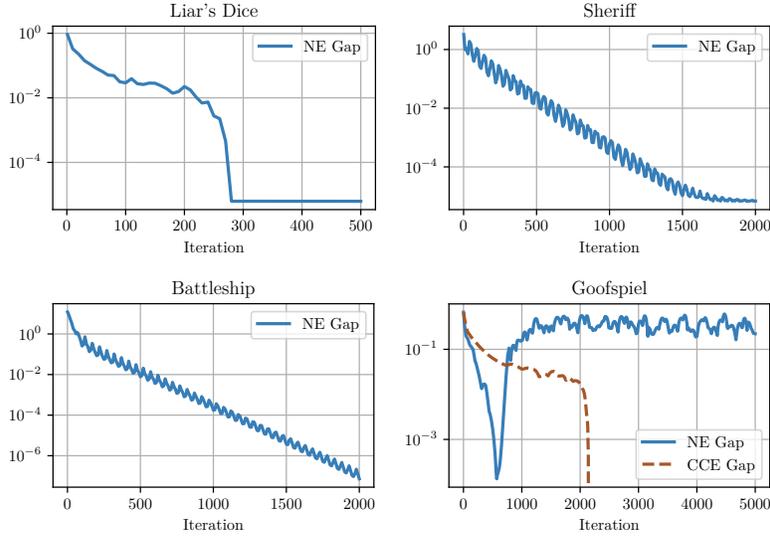


Figure 4: The NE gap of the last iterate and the CCE gap of the average correlated distribution of play under (OGD) in different benchmark games.

rapidly increases. Afterwards, the CCE gap starts to decay remarkably fast, eventually leading to a strong CCE. Specifically, after $T = 5000$ iterations the average correlated distribution of play is (roughly) a 0.0637-strong CCE. These results are consistent with the predictions of Theorem 1.1.

5 Discussion

Our primary contribution was to establish a new characterization for the convergence properties of OMD—an uncoupled no-regret learning algorithm—when employed by both players in a general-sum game: OMD either reaches arbitrarily close to a Nash equilibrium, or, otherwise, *both* players experience $-\Omega(T)$ regret. Our results open several interesting avenues for future research:

- Can we extend Theorem 1.1 from coarse correlated equilibria to *correlated equilibria*? At the very least, such an extension would be particularly challenging since all the known *no-internal-regret* algorithms used for converging to correlated equilibria involve the stationary distribution of a Markov chain. While the recent reduction in [Anagnostides et al., 2022b] may seem useful, their technique only applies for (OMD) under entropic regularization. In contrast, our current argument crucially relies on the smoothness of the regularizer.
- For which classes of multiplayer games would Theorem 1.1 hold? As we pointed out in Remark 3.7, while Theorem 1.1 does not extend in arbitrary multiplayer games, it is still interesting to give sufficient conditions under which our results would carry over.
- Can we characterize the bimatrix games for which (OMD) exhibits last-iterate convergence? It should be noted that recent results seem to suggest that such a characterization may be too hard to obtain in general [Andrade et al., 2021].
- Finally, can we improve Theorem 1.1 in terms of the dependence on T and η ? For example, it would be interesting to extend Theorem 1.1 under a learning rate that does not depend on ϵ . Some form of coupling in the spirit of *alternation* [Tammelin et al., 2015] could be useful in that direction.

Acknowledgements

We are grateful to anonymous NeurIPS reviewers for many helpful comments. Ioannis Panageas is supported by a start-up grant. Part of this work was done while Ioannis Panageas was visiting the Simons Institute for the Theory of Computing. Tuomas Sandholm is supported by the National Science Foundation under grants IIS-1901403 and CCF-1733556.

References

- Michael Bowling, Neil Burch, Michael Johanson, and Oskari Tammelin. Heads-up limit hold'em poker is solved. *Science*, 347(6218), January 2015.
- Noam Brown and Tuomas Sandholm. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, Dec. 2017.
- Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, May 2017.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- John Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36:48–49, 1950.
- Constantinos Daskalakis, Paul Goldberg, and Christos Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 2008.
- Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points (extended abstract). In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, pages 113–123, 2007.
- Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3):14:1–14:57, 2009.
- Aviad Rubinfeld. Settling the complexity of computing approximate two-player nash equilibria. In Irit Dinur, editor, *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS*, pages 258–265. IEEE Computer Society, 2016.
- Hervé Moulin and J-P Vial. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3):201–221, 1978.
- John C Harsanyi, Reinhard Selten, et al. A general theory of equilibrium selection in games. *MIT Press Books*, 1, 1988.
- John C. Harsanyi. A new theory of equilibrium selection for games with complete information. *Games and Economic Behavior*, 8(1):91–122, 1995.
- Akihiko Matsui and Kiminori Matsuyama. An approach to equilibrium selection. *Journal of Economic Theory*, 65(2):415–434, 1995.
- Robert Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
- Christos H. Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *J. ACM*, 55(3):14:1–14:29, 2008.
- Albert Xin Jiang and Kevin Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. *Games Econ. Behav.*, 91:347–359, 2015.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68:1127–1150, 2000.
- Ilan Adler, Constantinos Daskalakis, and Christos H. Papadimitriou. A note on strictly competitive games. In *Internet and Network Economics, 5th International Workshop, WINE 2009*, volume 5929 of *Lecture Notes in Computer Science*, pages 471–474. Springer, 2009.

- Constantinos Daskalakis and Christos H. Papadimitriou. On a network generalization of the minmax theorem. In *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009*, volume 5556 of *Lecture Notes in Computer Science*, pages 423–434. Springer, 2009a.
- Yang Cai and Constantinos Daskalakis. On minmax theorems for multiplayer games. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, pages 217–234. SIAM, 2011.
- Yang Cai, Ozan Candogan, Constantinos Daskalakis, and Christos Papadimitriou. Zero-sum polymatrix games: A generalization of minmax. *Mathematics of Operations Research*, 41(2):648–655, 2016.
- Ilan Adler. The equivalence of linear programs and zero-sum games. *Int. J. Game Theory*, 42(1): 165–177, 2013.
- Constantinos Daskalakis, Alan Deckelbaum, and Anthony Kim. Near-optimal no-regret algorithms for zero-sum games. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011.
- Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Conference on Learning Theory*, pages 6–1, 2012.
- Alexander Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In *Advances in Neural Information Processing Systems*, pages 3066–3074, 2013a.
- Robert D. Kleinberg, Katrina Ligett, Georgios Piliouras, and Éva Tardos. Beyond the nash equilibrium barrier. In Bernard Chazelle, editor, *Innovations in Computer Science - ICS 2011*, pages 125–140. Tsinghua University Press, 2011.
- Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems*, pages 2989–2997, 2015.
- James P. Bailey and Georgios Piliouras. Fast and furious learning in zero-sum games: Vanishing regret with non-vanishing step sizes. In *Advances in Neural Information Processing Systems 32*, pages 12977–12987, 2019.
- Georgios Piliouras, Ryann Sim, and Stratis Skoulakis. Optimal no-regret learning in general games: Bounded regret with unbounded step-sizes via clairvoyant mwu. *arXiv preprint arXiv:2111.14737*, 2021.
- Noah Golowich, Sarath Pattathil, and Constantinos Daskalakis. Tight last-iterate convergence rates for no-regret learning in multi-player games. In *Advances in Neural Information Processing Systems 2020*, 2020a.
- Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences*, 48(3):498–532, 1994.
- Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1), 2009.
- Constantinos Daskalakis, Aranyak Mehta, and Christos H. Papadimitriou. A note on approximate nash equilibria. In *Internet and Network Economics, Second International Workshop, WINE 2006*, volume 4286 of *Lecture Notes in Computer Science*, pages 297–306. Springer, 2006a.
- Constantinos Daskalakis, Aranyak Mehta, and Christos H. Papadimitriou. Progress in approximate nash equilibria. In *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007)*, pages 355–358. ACM, 2007.
- Hartwig Bosse, Jaroslaw Byrka, and Evangelos Markakis. New algorithms for approximate nash equilibria in bimatrix games. *Theoretical Computer Science*, 411(1):164–173, 2010.
- Haralampos Tsaknakis and Paul G. Spirakis. An Optimization Approach for Approximate Nash Equilibria. *Internet Mathematics*, 5(4):365 – 382, 2008.

- Spyros C Kontogiannis and Paul G Spirakis. Well supported approximate equilibria in bimatrix games. *Algorithmica*, 57(4):653–667, 2010.
- Constantinos Daskalakis and Christos H. Papadimitriou. On oblivious ptas’s for nash equilibrium. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, pages 75–84. ACM, 2009b.
- Argyrios Deligkas, Michail Fasoulakis, and Evangelos Markakis. A polynomial-time algorithm for 1/3-approximate nash equilibria in bimatrix games. In *30th Annual European Symposium on Algorithms, ESA 2022*, volume 244 of *LIPICs*, pages 41:1–41:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022a.
- Argyrios Deligkas, Michail Fasoulakis, and Evangelos Markakis. A polynomial-time algorithm for 1/2-well-supported nash equilibria in bimatrix games. *CoRR*, abs/2207.07007, 2022b.
- Pravesh Kothari and Ruta Mehta. Sum-of-squares meets nash: Optimal lower bounds for finding any equilibrium. *Electron. Colloquium Comput. Complex.*, page 126, 2018.
- Yakov Babichenko, Christos H. Papadimitriou, and Aviad Rubinfeld. Can almost everybody be almost happy? In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 1–9. ACM, 2016.
- Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, page 36–41. Association for Computing Machinery, 2003.
- Constantinos Daskalakis, Alex Fabrikant, and Christos H. Papadimitriou. The game world is flat: The complexity of nash equilibria in succinct games. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006*, volume 4051 of *Lecture Notes in Computer Science*, pages 513–524. Springer, 2006b.
- Xi Chen and Binghui Peng. Hedging in games: Faster convergence of external and swap regrets. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Dylan J Foster, Zhiyuan Li, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos. Learning in games: Robustness of fast convergence. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29, 2016.
- Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. In *Conference On Learning Theory, COLT 2018*, volume 75 of *Proceedings of Machine Learning Research*, pages 1263–1291. PMLR, 2018.
- Gabriele Farina, Christian Kroer, Noam Brown, and Tuomas Sandholm. Stable-predictive optimistic counterfactual regret minimization. In *International Conference on Machine Learning (ICML)*, 2019a.
- Constantinos Daskalakis and Noah Golowich. Fast rates for nonparametric online learning: from realizability to learning in games. In *STOC ’22: 54th Annual ACM SIGACT Symposium on Theory of Computing, 2022*, pages 846–859. ACM, 2022.
- Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. Near-optimal no-regret learning in general games. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021*, pages 27604–27616, 2021.
- Yuzuru Sato, Eizo Akiyama, and J. Dooyne Farmer. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99(7):4748–4751, 2002.
- William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010.
- Panayotis Mertikopoulos, Christos H. Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 2703–2717. SIAM, 2018.

- Gabriel P. Andrade, Rafael M. Frongillo, and Georgios Piliouras. Learning in matrix games can be arbitrarily complex. In Mikhail Belkin and Samory Kpotufe, editors, *Conference on Learning Theory, COLT 2021*, volume 134 of *Proceedings of Machine Learning Research*, pages 159–185. PMLR, 2021.
- Yun Kuen Cheung and Georgios Piliouras. Chaos, extremism and optimism: Volume analysis of learning in games. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020*, 2020.
- Yun Kuen Cheung and Yixin Tao. Chaos of learning beyond zero-sum and coordination via game decompositions. In *9th International Conference on Learning Representations, ICLR 2021*. OpenReview.net, 2021.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. In *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net, 2018.
- Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019*, volume 124 of *LIPICs*, pages 27:1–27:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In *NeurIPS 2018*, pages 9256–9266, 2018.
- Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. 2021a.
- Noah Golowich, Sarath Pattathil, Constantinos Daskalakis, and Asuman E. Ozdaglar. Last iterate is slower than averaged iterate in smooth convex-concave saddle point problems. In *Conference on Learning Theory, COLT 2020*, volume 125 of *Proceedings of Machine Learning Research*, pages 1758–1784. PMLR, 2020b.
- Yu-Guan Hsieh, Kimon Antonakopoulos, and Panayotis Mertikopoulos. Adaptive learning in continuous games: Optimal regret bounds and convergence to nash equilibrium. In *Conference on Learning Theory, COLT 2021, 15-19 August 2021, Boulder, Colorado, USA*, volume 134 of *Proceedings of Machine Learning Research*, pages 2388–2422. PMLR, 2021.
- Zhengyuan Zhou, Panayotis Mertikopoulos, Aris L. Moustakas, Nicholas Bambos, and Peter W. Glynn. Mirror descent learning in continuous games. In *56th IEEE Annual Conference on Decision and Control, CDC 2017, Melbourne, Australia, December 12-15, 2017*, pages 5776–5783. IEEE, 2017.
- Tianyi Lin, Zhengyuan Zhou, Panayotis Mertikopoulos, and Michael I. Jordan. Finite-time last-iterate convergence for multi-agent learning in games. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 6161–6171. PMLR, 2020.
- Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. In *7th International Conference on Learning Representations, ICLR 2019*. OpenReview.net, 2019.
- Tengyuan Liang and James Stokes. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019*, volume 89 of *Proceedings of Machine Learning Research*, pages 907–915. PMLR, 2019.
- Guojun Zhang and Yaoliang Yu. Convergence of gradient methods on bilinear zero-sum games. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.

- Aryan Mokhtari, Asuman E. Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020*, volume 108 of *Proceedings of Machine Learning Research*, pages 1497–1507. PMLR, 2020.
- Constantinos Daskalakis, Dylan J. Foster, and Noah Golowich. Independent policy gradient methods for competitive reinforcement learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020*, 2020.
- Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Last-iterate convergence of decentralized optimistic gradient descent/ascent in infinite-horizon competitive markov games. In *Conference on Learning Theory, COLT 2021*, volume 134 of *Proceedings of Machine Learning Research*, pages 4259–4299. PMLR, 2021b.
- Ioannis Anagnostides and Ioannis Panageas. *Frequency-Domain Representation of First-Order Methods: A Simple and Robust Framework of Analysis*, pages 131–160. 2021.
- Ioannis Anagnostides, Ioannis Panageas, Gabriele Farina, and Tuomas Sandholm. On last-iterate convergence beyond zero-sum games. In *International Conference on Machine Learning, ICML 2022*, volume 162 of *Proceedings of Machine Learning Research*, pages 536–581. PMLR, 2022a.
- Elad Hazan and Satyen Kale. Better algorithms for benign bandits. *J. Mach. Learn. Res.*, 12: 1287–1311, 2011.
- Chao-Kai Chiang, Chia-Jung Lee, and Chi-Jen Lu. Beating bandits in gradually evolving worlds. In *COLT 2013 - The 26th Annual Conference on Learning Theory*, volume 30 of *JMLR Workshop and Conference Proceedings*, pages 210–227. JMLR.org, 2013.
- Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Conference on Learning Theory*, pages 993–1019, 2013b.
- Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Faster game solving via predictive blackwell approachability: Connecting regret matching and mirror descent. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- I. Romanovskii. Reduction of a game with complete memory to a matrix game. *Soviet Mathematics*, 3, 1962.
- Bernhard von Stengel. Efficient computation of behavior strategies. *Games and Economic Behavior*, 14(2):220–246, 1996.
- Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior*, 14(2), 1996.
- David Avis, Gabriel D. Rosenberg, Rahul Savani, and Bernhard von Stengel. Enumeration of nash equilibria for two-player games. *Economic Theory*, 42(1):9–37, 2010.
- Gabriele Farina, Chun Kai Ling, Fei Fang, and Tuomas Sandholm. Correlation in extensive-form games: Saddle-point formulation and benchmarks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019b.
- Viliam Lisý, Marc Lanctot, and Michael Bowling. Online Monte Carlo counterfactual regret minimization for search in imperfect information games. In *Autonomous Agents and Multi-Agent Systems*, pages 27–36, 2015.
- Sheldon M Ross. Goofspiel—the game of pure strategy. *Journal of Applied Probability*, 8(3): 621–625, 1971.
- Ioannis Anagnostides, Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Noah Golowich, and Tuomas Sandholm. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, 2022*, pages 736–749. ACM, 2022b.

Oskari Tammelin, Neil Burch, Michael Johanson, and Michael Bowling. Solving heads-up limit Texas hold'em. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.

Francisco Facchinei and Jong-Shi Pang. *Finite-dimensional variational inequalities and complementarity problems*. Springer, 2003.

Checklist

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

A Omitted Proofs

In this section we include all the proofs deferred from the main body. For the convenience of the reader, all claims will be restated. Before we proceed with the proof of Theorem 2.3, let us point out some additional notational conventions. First, with a standard abuse of notation, we overload

$$u_i : \times_{i=1}^p \Delta(\mathcal{A}_i) \ni (\mathbf{x}_1, \dots, \mathbf{x}_p) \mapsto \mathbb{E}_{\mathbf{a} \sim \mathbf{x}} [u_i(\mathbf{a})] = \sum_{\mathbf{a} \in \mathcal{A}} u_i(a_1, \dots, a_p) \prod_{j \in [p]} \mathbf{x}_j(a_j)$$

to denote the mixed extension of player's $i \in [p]$ utility function, where $\mathcal{A} := \times_{j=1}^p \mathcal{A}_j$. Furthermore, for $a_i \in \mathcal{A}_i$ we let

$$u_i(a_i, \mathbf{x}_{-i}) := \sum_{\mathbf{a}_{-i} \in \mathcal{A}_{-i}} u_i(a_1, \dots, a_i, \dots, a_p) \prod_{j \neq i} \mathbf{x}_j(a_j),$$

where $\mathcal{A}_{-i} := \times_{j \neq i} \mathcal{A}_j$.

Theorem 2.3 (Folklore). *Suppose that every player $i \in [p]$ employs a no-regret learning algorithm with regret Reg_i^T up to time $T \in \mathbb{N}$. Moreover, let $\boldsymbol{\mu}^{(t)} := \mathbf{x}_1^{(t)} \otimes \dots \otimes \mathbf{x}_p^{(t)}$ be the correlated distribution of play at time $t \in [T]$, and $\bar{\boldsymbol{\mu}} := \frac{1}{T} \sum_{t=1}^T \boldsymbol{\mu}^{(t)}$ be the average correlated distribution of play up to time T . Then,*

$$\mathbb{E}_{\mathbf{a} \sim \bar{\boldsymbol{\mu}}} [u_i(a'_i, \mathbf{a}_{-i})] \leq \mathbb{E}_{\mathbf{a} \sim \bar{\boldsymbol{\mu}}} [u_i(\mathbf{a})] + \frac{1}{T} \max_{i \in [p]} \text{Reg}_i^T.$$

Proof. By definition of regret (1), it follows that for any player $i \in [p]$ and any possible deviation $a'_i \in \mathcal{A}_i$,

$$\sum_{t=1}^T u_i(a'_i, \mathbf{x}_{-i}^{(t)}) - \sum_{t=1}^T u_i(\mathbf{x}^{(t)}) \leq \text{Reg}_i^T. \quad (6)$$

Moreover,

$$\mathbb{E}_{\mathbf{a} \sim \bar{\boldsymbol{\mu}}} [u_i(\mathbf{a})] = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{a} \sim \boldsymbol{\mu}^{(t)}} [u_i(\mathbf{a})] = \frac{1}{T} \sum_{t=1}^T u_i(\mathbf{x}^{(t)});$$

and

$$\mathbb{E}_{\mathbf{a} \sim \bar{\boldsymbol{\mu}}} [u_i(a'_i, \mathbf{a}_{-i})] = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{a} \sim \boldsymbol{\mu}^{(t)}} [u_i(a'_i, \mathbf{a}_{-i})] = \frac{1}{T} \sum_{t=1}^T u_i(a'_i, \mathbf{x}_{-i}^{(t)}).$$

Thus, the theorem follows directly from (6). \square

Corollary 3.1. *Suppose that both players employ (OMD) with learning rate $\eta > 0$. Then,*

$$\begin{aligned} \text{Reg}_{\mathcal{X}}^T &\leq \frac{\Omega_{\mathcal{R}\mathcal{X}}}{\eta} + \eta \|\mathbf{A}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 - \frac{1}{4\eta} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t-1)}\|^2 \right); \\ \text{Reg}_{\mathcal{Y}}^T &\leq \frac{\Omega_{\mathcal{R}\mathcal{Y}}}{\eta} + \eta \|\mathbf{B}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\|^2 - \frac{1}{4\eta} \sum_{t=1}^T \left(\|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t)}\|^2 + \|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t-1)}\|^2 \right). \end{aligned}$$

Proof. First, the utility $\mathbf{u}_{\mathcal{X}}^{(t)}$ observed by player \mathcal{X} at time $t \geq 0$ is equal to $\mathbf{A}\mathbf{y}^{(t)}$. Thus, the claimed bound for $\text{Reg}_{\mathcal{X}}^T$ follows directly from Proposition 2.1 using the fact that $\|\mathbf{A}\mathbf{y}^{(t)} - \mathbf{A}\mathbf{y}^{(t-1)}\|_* \leq \|\mathbf{A}\|_{\text{op}} \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|$. Similarly, the utility $\mathbf{u}_{\mathcal{Y}}^{(t)}$ observed by player \mathcal{Y} at time $t \geq 0$ is tantamount to $\mathbf{B}^\top \mathbf{x}^{(t)}$. As a result, the bound for $\text{Reg}_{\mathcal{Y}}^T$ follows from Proposition 2.1 and the fact that $\|\mathbf{B}^\top \mathbf{x}^{(t)} - \mathbf{B}^\top \mathbf{x}^{(t-1)}\|_* \leq \|\mathbf{B}\|_{\text{op}} \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\|$, since $\|\mathbf{B}^\top\|_{\text{op}} = \|\mathbf{B}\|_{\text{op}}$. \square

Proposition 3.2 (Approximate Fixed Points of OMD). *Consider a bimatrix game (\mathbf{A}, \mathbf{B}) , and suppose that both players employ (OMD) with learning rate $\eta > 0$ and a G -smooth regularizer. Then, if $\|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t-1)}\|, \|\hat{\mathbf{x}}^{(t)} - \mathbf{x}^{(t)}\| \leq \epsilon\eta$ and $\|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t-1)}\|, \|\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}\| \leq \epsilon\eta$, the pair $(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})$ is a $(2\epsilon G \max\{\Omega_{\mathcal{X}}, \Omega_{\mathcal{Y}}\} + \epsilon\eta)$ -approximate Nash equilibrium of (\mathbf{A}, \mathbf{B}) .*

Proof. First, by definition of the Bregman divergence, the update rule of (OMD) can be equivalently expressed as

$$\begin{aligned}\mathbf{x}^{(t)} &:= \arg \max_{\mathbf{x} \in \mathcal{X}} \left\{ \langle \mathbf{x}, \mathbf{A}\mathbf{y}^{(t-1)} \rangle - \frac{1}{\eta} \mathcal{R}(\mathbf{x}) + \frac{1}{\eta} \langle \mathbf{x}, \nabla \mathcal{R}(\widehat{\mathbf{x}}^{(t-1)}) \rangle \right\}; \\ \widehat{\mathbf{x}}^{(t)} &:= \arg \max_{\widehat{\mathbf{x}} \in \mathcal{X}} \left\{ \langle \widehat{\mathbf{x}}, \mathbf{A}\mathbf{y}^{(t)} \rangle - \frac{1}{\eta} \mathcal{R}(\widehat{\mathbf{x}}) + \frac{1}{\eta} \langle \widehat{\mathbf{x}}, \nabla \mathcal{R}(\widehat{\mathbf{x}}^{(t-1)}) \rangle \right\}.\end{aligned}$$

Now the maximization problem associated with the update rule of the secondary sequence can be equivalently cast in a variational inequality form (e.g., see [Facchinei and Pang, 2003]):

$$\left\langle \widehat{\mathbf{x}} - \widehat{\mathbf{x}}^{(t)}, \mathbf{A}\mathbf{y}^{(t)} - \frac{1}{\eta} \left(\nabla \mathcal{R}(\widehat{\mathbf{x}}^{(t)}) - \nabla \mathcal{R}(\widehat{\mathbf{x}}^{(t-1)}) \right) \right\rangle \leq 0, \quad \forall \widehat{\mathbf{x}} \in \mathcal{X}.$$

Thus,

$$\begin{aligned}\langle \widehat{\mathbf{x}} - \widehat{\mathbf{x}}^{(t)}, \mathbf{A}\mathbf{y}^{(t)} \rangle &\leq \frac{1}{\eta} \langle \widehat{\mathbf{x}} - \widehat{\mathbf{x}}^{(t)}, \nabla \mathcal{R}(\widehat{\mathbf{x}}^{(t)}) - \nabla \mathcal{R}(\widehat{\mathbf{x}}^{(t-1)}) \rangle \\ &\leq \frac{1}{\eta} \|\widehat{\mathbf{x}} - \widehat{\mathbf{x}}^{(t)}\| \|\nabla \mathcal{R}(\widehat{\mathbf{x}}^{(t)}) - \nabla \mathcal{R}(\widehat{\mathbf{x}}^{(t-1)})\|_* \quad (7)\end{aligned}$$

$$\leq \frac{G}{\eta} \|\widehat{\mathbf{x}} - \widehat{\mathbf{x}}^{(t)}\| \|\widehat{\mathbf{x}}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| \quad (8)$$

$$\leq 2\epsilon G \Omega_{\mathcal{X}}, \quad (9)$$

for any $\widehat{\mathbf{x}} \in \mathcal{X}$, where (7) derives from the Cauchy-Schwarz inequality; (8) follows since $\mathcal{R}_{\mathcal{X}}$ is assumed to be G -smooth; and (9) uses that $\|\widehat{\mathbf{x}}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| \leq 2\epsilon\eta$, which in turn follows since $\|\widehat{\mathbf{x}}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| \leq \|\widehat{\mathbf{x}}^{(t)} - \mathbf{x}^{(t)}\| + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| \leq 2\epsilon\eta$ (triangle inequality), as well as the fact that, by definition, $\|\widehat{\mathbf{x}} - \widehat{\mathbf{x}}^{(t)}\| \leq \Omega_{\mathcal{X}}$ for any $\widehat{\mathbf{x}} \in \mathcal{X}$. As a result, we have shown that for any $\widehat{\mathbf{x}} \in \mathcal{X}$,

$$\langle \widehat{\mathbf{x}}^{(t)}, \mathbf{A}\mathbf{y}^{(t)} \rangle \geq \langle \widehat{\mathbf{x}}, \mathbf{A}\mathbf{y}^{(t)} \rangle - 2\epsilon G \Omega_{\mathcal{X}}. \quad (10)$$

Furthermore, by Cauchy-Schwarz inequality we have that

$$\langle \mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}, \mathbf{A}\mathbf{y}^{(t)} \rangle \geq -\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\| \|\mathbf{A}\mathbf{y}^{(t)}\|_* \geq -\epsilon\eta, \quad (11)$$

where we used the normalization assumption $\|\mathbf{A}\mathbf{y}^{(t)}\|_* \leq 1$. Thus, combing (11) with (10) yields that

$$\langle \mathbf{x}^{(t)}, \mathbf{A}\mathbf{y}^{(t)} \rangle \geq \langle \widehat{\mathbf{x}}^{(t)}, \mathbf{A}\mathbf{y}^{(t)} \rangle - \epsilon\eta \geq \langle \widehat{\mathbf{x}}, \mathbf{A}\mathbf{y}^{(t)} \rangle - 2\epsilon G \Omega_{\mathcal{X}} - \epsilon\eta, \quad (12)$$

for any $\widehat{\mathbf{x}} \in \mathcal{X}$. By symmetry, we analogously get that

$$\langle \mathbf{y}^{(t)}, \mathbf{B}^{\top} \mathbf{x}^{(t)} \rangle \geq \langle \widehat{\mathbf{y}}^{(t)}, \mathbf{B}^{\top} \mathbf{x}^{(t)} \rangle - \epsilon\eta \geq \langle \widehat{\mathbf{y}}, \mathbf{B}^{\top} \mathbf{x}^{(t)} \rangle - 2\epsilon G \Omega_{\mathcal{Y}} - \epsilon\eta, \quad (13)$$

for any $\widehat{\mathbf{y}} \in \mathcal{Y}$. Thus, recalling Definition 2.5, the claim follows from (12) and (13). \square

Lemma 3.3. *Suppose that both players in a bimatrix game (\mathbf{A}, \mathbf{B}) employ (OMD) with learning rate $\eta > 0$. Then, for any $T \in \mathbb{N}$,*

$$\begin{aligned}\sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\| &\geq \frac{1}{2\eta\|\mathcal{X}\|\|\mathbf{A}\|_{\text{op}}} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 + \|\widehat{\mathbf{x}}^{(t)} - \mathbf{x}^{(t)}\|^2 \right) - \frac{2}{\|\mathbf{A}\|_{\text{op}}}; \\ \sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\| &\geq \frac{1}{2\eta\|\mathcal{Y}\|\|\mathbf{B}\|_{\text{op}}} \sum_{t=1}^T \left(\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 + \|\widehat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}\|^2 \right) - \frac{2}{\|\mathbf{B}\|_{\text{op}}}.\end{aligned}$$

Proof. By 1-strong convexity of $\mathcal{R}_{\mathcal{X}}$ with respect to $\|\cdot\|$,

$$\langle \mathbf{x}^{(t)}, \mathbf{A}\mathbf{y}^{(t-1)} \rangle - \frac{1}{\eta} D_{\mathcal{R}_{\mathcal{X}}}(\mathbf{x}^{(t)} \parallel \widehat{\mathbf{x}}^{(t-1)}) - \langle \widehat{\mathbf{x}}^{(t-1)}, \mathbf{A}\mathbf{y}^{(t-1)} \rangle \geq \frac{1}{2\eta} \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2, \quad (14)$$

where we used the definition of the update rule of the primary sequence of (OMD). Similarly,

$$\langle \widehat{\mathbf{x}}^{(t)}, \mathbf{A}\mathbf{y}^{(t)} \rangle - \frac{1}{\eta} D_{\mathcal{R}_{\mathcal{X}}}(\widehat{\mathbf{x}}^{(t)} \parallel \widehat{\mathbf{x}}^{(t-1)}) - \langle \mathbf{x}^{(t)}, \mathbf{A}\mathbf{y}^{(t)} \rangle + \frac{1}{\eta} D_{\mathcal{R}_{\mathcal{X}}}(\mathbf{x}^{(t)} \parallel \widehat{\mathbf{x}}^{(t-1)}) \geq \frac{1}{2\eta} \|\widehat{\mathbf{x}}^{(t)} - \mathbf{x}^{(t)}\|^2. \quad (15)$$

Hence, summing (14) and (15) yields that

$$\langle \mathbf{x}^{(t)}, \mathbf{A}(\mathbf{y}^{(t-1)} - \mathbf{y}^{(t)}) \rangle \geq \frac{1}{2\eta} \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 + \|\widehat{\mathbf{x}}^{(t)} - \mathbf{x}^{(t)}\|^2 \right) - \langle \widehat{\mathbf{x}}^{(t)}, \mathbf{A}\mathbf{y}^{(t)} \rangle + \langle \widehat{\mathbf{x}}^{(t-1)}, \mathbf{A}\mathbf{y}^{(t-1)} \rangle,$$

where we used that $D_{\mathcal{R}_X}(\widehat{\mathbf{x}}^{(t)} \parallel \widehat{\mathbf{x}}^{(t-1)}) \geq 0$. Thus, a telescopic summation over all $t \in [T]$ implies that

$$\sum_{t=1}^T \langle \mathbf{x}^{(t)}, \mathbf{A}(\mathbf{y}^{(t-1)} - \mathbf{y}^{(t)}) \rangle \geq \frac{1}{2\eta} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 + \|\widehat{\mathbf{x}}^{(t)} - \mathbf{x}^{(t)}\|^2 \right) - 2\|\mathcal{X}\|, \quad (16)$$

since $-\langle \widehat{\mathbf{x}}^{(T)}, \mathbf{A}\mathbf{y}^{(T)} \rangle \geq -\|\widehat{\mathbf{x}}^{(T)}\| \|\mathbf{A}\mathbf{y}^{(T)}\|_* \geq -\|\mathcal{X}\|$ and $\langle \widehat{\mathbf{x}}^{(0)}, \mathbf{A}\mathbf{y}^{(0)} \rangle \geq -\|\widehat{\mathbf{x}}^{(0)}\| \|\mathbf{A}\mathbf{y}^{(0)}\|_* \geq -\|\mathcal{X}\|$, where we used the normalization assumption. Furthermore,

$$\langle \mathbf{x}^{(t)}, \mathbf{A}(\mathbf{y}^{(t-1)} - \mathbf{y}^{(t)}) \rangle \leq \|\mathbf{x}^{(t)}\| \|\mathbf{A}(\mathbf{y}^{(t-1)} - \mathbf{y}^{(t)})\|_* \leq \|\mathcal{X}\| \|\mathbf{A}\|_{\text{op}} \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|.$$

Thus, combining this inequality with (16) implies that

$$\sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\| \geq \frac{1}{2\eta \|\mathcal{X}\| \|\mathbf{A}\|_{\text{op}}} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 + \|\widehat{\mathbf{x}}^{(t)} - \mathbf{x}^{(t)}\|^2 \right) - \frac{2}{\|\mathbf{A}\|_{\text{op}}}.$$

This completes the first part of the claim. The second part follows analogously by symmetry. \square

Lemma 3.5 (Stability of OMD). *Suppose that both players employ (OMD) with learning rate $\eta > 0$. Then, for any $t \in \mathbb{N}$,*

$$\begin{aligned} \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\| &\leq 3\eta; \\ \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\| &\leq 3\eta. \end{aligned}$$

Proof. Fix any $t \in \mathbb{N}$. By definition of the primary sequence of (OMD),

$$\langle \mathbf{x}^{(t)}, \mathbf{A}\mathbf{y}^{(t-1)} \rangle - \frac{1}{\eta} D_{\mathcal{R}_X}(\mathbf{x}^{(t)} \parallel \widehat{\mathbf{x}}^{(t-1)}) - \langle \widehat{\mathbf{x}}^{(t-1)}, \mathbf{A}\mathbf{y}^{(t-1)} \rangle \geq \frac{1}{2\eta} \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2,$$

where we used the 1-strong convexity of the regularizer \mathcal{R}_X with respect to $\|\cdot\|$. In turn, this implies that

$$\langle \mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}, \mathbf{A}\mathbf{y}^{(t-1)} \rangle \geq \frac{1}{\eta} \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2,$$

since $D_{\mathcal{R}_X}(\mathbf{x}^{(t)} \parallel \widehat{\mathbf{x}}^{(t-1)}) \geq \frac{1}{2} \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2$ (by 1-strong convexity of \mathcal{R}_X). Thus, an application of Cauchy-Schwarz inequality yields that

$$\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \leq \eta \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| \|\mathbf{A}\mathbf{y}^{(t-1)}\|_* \implies \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| \leq \eta, \quad (17)$$

since $\|\mathbf{A}\mathbf{y}^{(t-1)}\|_* \leq 1$ by the normalization assumption. Similar reasoning applied for the secondary sequence of (OMD) implies that for any $t \in \mathbb{N}$,

$$\|\widehat{\mathbf{x}}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| \leq \eta. \quad (18)$$

Now if $t = 1$, it follows from (17) that $\|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\| = \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| \leq \eta$ since $\mathbf{x}^{(0)} = \widehat{\mathbf{x}}^{(0)}$. Otherwise, for $t \geq 2$, applying the triangle inequality yields that $\|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\| \leq \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| + \|\mathbf{x}^{(t-1)} - \widehat{\mathbf{x}}^{(t-2)}\| + \|\widehat{\mathbf{x}}^{(t-1)} - \widehat{\mathbf{x}}^{(t-2)}\| \leq 3\eta$ by (17) and (18). This completes the first part of the proof. Analogously, we conclude that $\|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\| \leq 3\eta$ for any $t \in \mathbb{N}$. \square

Theorem A.1 (Linear Decay of Regret; Full Version of Theorem 3.4). *Suppose that both players in a bimatrix game (\mathbf{A}, \mathbf{B}) employ (OMD) with G -smooth regularizer, learning rate $\eta > 0$ such that*

$$\eta \leq \min \left\{ \frac{1}{4 \max\{\|\mathbf{A}\|_{\text{op}}, \|\mathbf{B}\|_{\text{op}}\}}, \frac{\epsilon^2}{96 \|\mathbf{A}\|_{\text{op}} \|\mathbf{B}\|_{\text{op}} \max\{\|\mathcal{X}\|, \|\mathcal{Y}\|\}} \right\}$$

and

$$T \geq \max \left\{ \frac{16 \max\{\|\mathcal{X}\|, \|\mathcal{Y}\|\}}{\epsilon^2 \eta}, \frac{32 \max\{\Omega_{\mathcal{R}_X}, \Omega_{\mathcal{R}_Y}\}}{\epsilon^2 \eta^2}, \frac{2048 \max\{\Omega_{\mathcal{R}_Y} \|\mathcal{X}\|^2 \|\mathbf{A}\|_{\text{op}}^2, \Omega_{\mathcal{R}_X} \|\mathcal{Y}\|^2 \|\mathbf{B}\|_{\text{op}}^2\}}{\epsilon^4 \eta^2} \right\},$$

for some fixed $\epsilon > 0$. Then, if the dynamics do not reach a $(2\epsilon G \max\{\Omega_X, \Omega_Y\} + \epsilon\eta)$ -approximate NE, then

$$\max\{\text{Reg}_{\mathcal{X}}^T, \text{Reg}_{\mathcal{Y}}^T\} \leq -\min \left\{ \frac{\epsilon^2 \eta}{32}, \frac{\epsilon^4 \eta}{2048 \max\{\|\mathcal{X}\|^2 \|\mathbf{A}\|_{\text{op}}^2, \|\mathcal{Y}\|^2 \|\mathbf{B}\|_{\text{op}}^2\}} \right\} T.$$

Proof. Suppose that there exists $t \in \llbracket T \rrbracket$ such that

$$\left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right) + \left(\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t)}\|^2 + \|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 \right) \leq \epsilon^2 \eta^2.$$

This would imply that $\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|, \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\| \leq \epsilon \eta$ and $\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t)}\|, \|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\| \leq \epsilon \eta$. In turn, by Proposition 3.2 it follows that the pair of strategies $(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})$ is a $(2\epsilon G \max\{\Omega_{\mathcal{X}}, \Omega_{\mathcal{Y}}\} + \epsilon \eta)$ -approximate Nash equilibrium, contradicting our assumption. As a result, we conclude that for all $t \in \llbracket T \rrbracket$,

$$\left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right) + \left(\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t)}\|^2 + \|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 \right) \geq \epsilon^2 \eta^2.$$

Summing over all $t \in \llbracket T \rrbracket$ yields that

$$\sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right) + \sum_{t=1}^T \left(\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t)}\|^2 + \|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 \right) \geq \epsilon^2 \eta^2 T. \quad (19)$$

We distinguish between two cases. First, we treat the case where

$$\sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right) \geq \sum_{t=1}^T \left(\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t)}\|^2 + \|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 \right). \quad (20)$$

Then, by virtue of (19),

$$\sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right) \geq \frac{\epsilon^2 \eta^2}{2} T. \quad (21)$$

Further, by the triangle inequality and Young's inequality,

$$\|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 \leq 2\|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 + 2\|\widehat{\mathbf{y}}^{(t-1)} - \mathbf{y}^{(t-1)}\|^2,$$

and summing over all $t \in \llbracket T \rrbracket$ yields that

$$\begin{aligned} \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 &\leq 2 \sum_{t=1}^T \|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 + 2 \sum_{t=1}^T \|\widehat{\mathbf{y}}^{(t-1)} - \mathbf{y}^{(t-1)}\|^2 \\ &\leq 2 \sum_{t=1}^T \|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2 + 2 \sum_{t=1}^T \|\widehat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}\|^2, \end{aligned}$$

where the last inequality follows since $\widehat{\mathbf{y}}^{(0)} = \mathbf{y}^{(0)}$. Hence, combining the latter bound with (20) implies that

$$\sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right) \geq \frac{1}{2} \sum_{t=1}^T \|\mathbf{y}^{(t)} - \widehat{\mathbf{y}}^{(t-1)}\|^2. \quad (22)$$

Now we are ready to bound the regret of player \mathcal{X} . By Corollary 3.1,

$$\text{Reg}_{\mathcal{X}}^T \leq \frac{\Omega_{\mathcal{R}_X}}{\eta} + \eta \|\mathbf{A}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 - \frac{1}{4\eta} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right). \quad (23)$$

But (22) implies that

$$\begin{aligned} \eta \|\mathbf{A}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 - \frac{1}{8\eta} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \widehat{\mathbf{x}}^{(t-1)}\|^2 \right) \\ \leq \left(\eta \|\mathbf{A}\|_{\text{op}}^2 - \frac{1}{16\eta} \right) \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 \leq 0, \end{aligned}$$

since $\eta \leq \frac{1}{4\|\mathbf{A}\|_{\text{op}}}$. From this we conclude that

$$\text{Reg}_{\mathcal{X}}^T \leq \frac{\Omega_{\mathcal{R}_X}}{\eta} - \frac{1}{8\eta} \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t-1)}\|^2 \right) \leq \frac{\Omega_{\mathcal{R}_X}}{\eta} - \frac{\epsilon^2 \eta}{16} T \leq -\frac{\epsilon^2 \eta}{32} T, \quad (24)$$

for $T \geq \frac{32\Omega_{\mathcal{R}_X}}{\epsilon^2 \eta^2}$, where we used (21). Next, we focus on the regret of player \mathcal{Y} . By (21) and Lemma 3.3,

$$\sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\| \geq \frac{\epsilon^2 \eta}{4\|\mathcal{X}\|\|\mathbf{A}\|_{\text{op}}} T - \frac{2}{\|\mathbf{A}\|_{\text{op}}} \geq \frac{\epsilon^2 \eta}{8\|\mathcal{X}\|\|\mathbf{A}\|_{\text{op}}} T, \quad (25)$$

since $T \geq \frac{16\|\mathcal{X}\|}{\epsilon^2 \eta}$. Further, by Cauchy-Schwarz inequality,

$$\sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 \geq \frac{1}{T} \left(\sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\| \right)^2 \geq \frac{\epsilon^4 \eta^2}{64\|\mathcal{X}\|^2 \|\mathbf{A}\|_{\text{op}}^2} T. \quad (26)$$

Now from Corollary 3.1, the regret of player \mathcal{Y} can be bounded as

$$\begin{aligned} \text{Reg}_{\mathcal{Y}}^T &\leq \frac{\Omega_{\mathcal{R}_Y}}{\eta} + \eta \|\mathbf{B}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\|^2 - \frac{1}{4\eta} \sum_{t=1}^T \left(\|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t)}\|^2 + \|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t-1)}\|^2 \right) \\ &\leq \frac{\Omega_{\mathcal{R}_Y}}{\eta} + \eta \|\mathbf{B}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\|^2 - \frac{1}{8\eta} \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2, \end{aligned}$$

where we used that

$$\sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 \leq 2 \sum_{t=1}^T \|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t-1)}\|^2 + 2 \sum_{t=1}^T \|\hat{\mathbf{y}}^{(t)} - \mathbf{y}^{(t)}\|^2.$$

Furthermore, by Lemma 3.5 and (26),

$$\eta \|\mathbf{B}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\|^2 \leq 9\eta^3 \|\mathbf{B}\|_{\text{op}}^2 T \leq \frac{\epsilon^4 \eta}{1024\|\mathcal{X}\|^2 \|\mathbf{A}\|_{\text{op}}^2} T \leq \frac{1}{16\eta} \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2,$$

for $\eta \leq \epsilon^2 (96\|\mathcal{X}\|\|\mathbf{A}\|_{\text{op}}\|\mathbf{B}\|_{\text{op}})^{-1}$. As a result, (26) implies that

$$\text{Reg}_{\mathcal{Y}}^T \leq \frac{\Omega_{\mathcal{R}_Y}}{\eta} - \frac{1}{16\eta} \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 \leq \frac{\Omega_{\mathcal{R}_Y}}{\eta} - \frac{\epsilon^4 \eta}{1024\|\mathcal{X}\|^2 \|\mathbf{A}\|_{\text{op}}^2} T \leq -\frac{\epsilon^4 \eta}{2048\|\mathcal{X}\|^2 \|\mathbf{A}\|_{\text{op}}^2} T,$$

for $T \geq \frac{2048\Omega_{\mathcal{R}_Y} \|\mathcal{X}\|^2 \|\mathbf{A}\|_{\text{op}}^2}{\epsilon^4 \eta^2}$. Similarly, let us treat the case where

$$\sum_{t=1}^T \left(\|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t)}\|^2 + \|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t-1)}\|^2 \right) \geq \sum_{t=1}^T \left(\|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t)}\|^2 + \|\mathbf{x}^{(t)} - \hat{\mathbf{x}}^{(t-1)}\|^2 \right).$$

Then, for $\eta \leq \frac{1}{4\|\mathbf{B}\|_{\text{op}}}$ and $T \geq \frac{32\Omega_{\mathcal{R}_Y}}{\epsilon^2 \eta^2}$,

$$\text{Reg}_{\mathcal{Y}}^T \leq \frac{\Omega_{\mathcal{R}_Y}}{\eta} - \frac{1}{8\eta} \sum_{t=1}^T \left(\|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t)}\|^2 + \|\mathbf{y}^{(t)} - \hat{\mathbf{y}}^{(t-1)}\|^2 \right) \leq \frac{\Omega_{\mathcal{R}_Y}}{\eta} - \frac{\epsilon^2 \eta}{16} T \leq -\frac{\epsilon^2 \eta}{32} T.$$

Moreover, for $T \geq \frac{16\|\mathcal{Y}\|}{\epsilon^2 \eta}$,

$$\sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\|^2 \geq \frac{\epsilon^4 \eta^2}{64\|\mathcal{Y}\|^2 \|\mathbf{B}\|_{\text{op}}^2} T.$$

Thus, for $\eta \leq \epsilon^2 (96\|\mathcal{Y}\|\|\mathbf{A}\|_{\text{op}}\|\mathbf{B}\|_{\text{op}})^{-1}$,

$$\eta \|\mathbf{A}\|_{\text{op}}^2 \sum_{t=1}^T \|\mathbf{y}^{(t)} - \mathbf{y}^{(t-1)}\|^2 \leq 9\eta^3 \|\mathbf{A}\|_{\text{op}}^2 T \leq \frac{\epsilon^4 \eta}{1024\|\mathcal{Y}\|^2 \|\mathbf{B}\|_{\text{op}}^2} T \leq \frac{1}{16\eta} \sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\|^2.$$

Finally, for $T \geq \frac{2048\Omega_{\mathcal{R}_X}\|\mathcal{Y}\|^2\|\mathbf{B}\|_{\text{op}}^2}{\epsilon^4\eta^2}$,

$$\text{Reg}_{\mathcal{X}}^T \leq \frac{\Omega_{\mathcal{R}_X}}{\eta} - \frac{1}{16\eta} \sum_{t=1}^T \|\mathbf{x}^{(t)} - \mathbf{x}^{(t-1)}\|^2 \leq \frac{\Omega_{\mathcal{R}_X}}{\eta} - \frac{\epsilon^4\eta}{1024\|\mathcal{Y}\|^2\|\mathbf{B}\|_{\text{op}}^2} T \leq -\frac{\epsilon^4\eta}{2048\|\mathcal{Y}\|^2\|\mathbf{B}\|_{\text{op}}^2} T.$$

□

Next, we state the implication of Theorem A.1 in normal-form games under (OGD). In that setting, it holds that $\|\mathcal{X}\|, \|\mathcal{Y}\| = 1$; $\Omega_{\mathcal{X}}, \Omega_{\mathcal{Y}} \leq \sqrt{2}$; $\Omega_{\mathcal{R}_X}, \Omega_{\mathcal{R}_Y} \leq 1$; and $G = 1$. Thus, we obtain the following simplified statement.

Corollary A.2 (OGD in Normal-Form Games). *Suppose that both players in a bimatrix game (\mathbf{A}, \mathbf{B}) employ (OGD) with learning rate $\eta > 0$ such that*

$$\eta \leq \min \left\{ \frac{1}{4 \max\{\|\mathbf{A}\|_{\text{op}}, \|\mathbf{B}\|_{\text{op}}\}}, \frac{\epsilon^2}{96\|\mathbf{A}\|_{\text{op}}\|\mathbf{B}\|_{\text{op}}} \right\}$$

and

$$T \geq \max \left\{ \frac{16}{\epsilon^2\eta}, \frac{32}{\epsilon^2\eta^2}, \frac{2048 \max\{\|\mathbf{A}\|_{\text{op}}^2, \|\mathbf{B}\|_{\text{op}}^2\}}{\epsilon^4\eta^2} \right\},$$

for some fixed $\epsilon > 0$. Then,

- Either there exists $t \in \llbracket T \rrbracket$ such that the pair of strategies $(\mathbf{x}^{(t)}, \mathbf{y}^{(t)}) \in \mathcal{X} \times \mathcal{Y}$ constitutes an $\epsilon(3 + \eta)$ -approximate Nash equilibrium;
- Or, otherwise, the average correlated distribution of play after T repetitions of the game is a

$$\min \left\{ \frac{\epsilon^2\eta}{32}, \frac{\epsilon^4\eta}{2048 \max\{\|\mathbf{A}\|_{\text{op}}^2, \|\mathbf{B}\|_{\text{op}}^2\}} \right\} - \text{strong coarse correlated equilibrium.}$$

Finally, we state an extension of Theorem 1.1 that establishes a dichotomy based on whether *most* of the iterates are approximate Nash equilibria—not just a *single* iterate. The proof is almost identical to the argument of Theorem A.1, and is therefore omitted.

Corollary A.3. *Suppose that both players in a bimatrix game employ (OGD) with learning rate $\eta = O(\epsilon^2\delta)$ and $T = \Omega\left(\frac{1}{\eta^2\epsilon^4\delta^2}\right)$ repetitions, for a sufficiently small $\epsilon > 0$ and $\delta \in (0, 1)$. Then,*

- Either a $1 - \delta$ fraction of the iterates is an ϵ -approximate Nash equilibrium;
- Or, otherwise, the average correlated distribution of play is an $\Omega(\epsilon^4\eta\delta^2)$ -strong CCE.

B Description of the Game Instances

In this section we provide a detailed description of the game instances we used in our experiments in Section 4.2.

Liar’s Dice The first game we experimented on is *Liar’s dice*, a popular benchmark introduced by Lisý et al. [2015]. In our instantiation, each of the two players initially privately roles a *single* unbiased 4-face die. Then, the first player announces any face value up to 4, as well as the minimum number of dice the player believes have that value (among the dice of both players). Subsequently, each player in its own turn can either make a higher bid, or challenge the claim made by the previous player by declaring that player a “liar”. In particular, a bid is higher than the previous one if either the face value is higher, or if the claimed number of dices is greater. In case the current player challenges the previous bid, all dice have to be revealed. If the claim was valid, the last bidder wins and receives a reward of +1, while the challenger incurs a negative payoff of −1. Otherwise, the utilities obtained are reversed.

Sheriff Our second benchmark is a bargaining game inspired by the board game *Sheriff of Nottingham*, introduced by Farina et al. [2019b]. This game consists of two players: the *smuggler* and the *sheriff*. In our instantiation, the smuggler initially selects a number $n \in \{0, 1, 2, 3, 4, 5\}$ which corresponds to the number of *illegal items* to be loaded in the cargo. Each illegal item has a fixed value of 1. Next, 2 rounds of bargaining between the two players follow. At each round, the smuggler decides on a *bribe* ranging from 0 to $b := 3$ (inclusive), and the sheriff must decide whether or not the cargo will be inspected given the bribe amount. The sheriff’s decision is binding only in the *last round* of bargaining: if the sheriff accepts the bribe, the game stops with the smuggler obtaining a utility of n minus the bribe amount b proposed in the last bargaining round, while the sheriff receives a utility equal to b . In contrast, if the sheriff does not accept the bribe in last bargaining round and decides to inspect the cargo, there are two possible alternatives:

- If the cargo has no illegal items (*i.e.* $n = 0$), the smuggler receives the fixed amount of 3, while sheriff incurs a negative payoff of -3 ;
- Otherwise, the utility of the smuggler is set to $-2n$, while the utility of the Sheriff is $2n$.

Battleship Our next benchmark is *Battleship*, a parametric version of the popular board game introduced in [Farina et al., 2019b]. At the beginning, each player secretly places its ships on separate locations on a grid of size 2×2 . Every ship has size 1 and a value of 4, and the placement is such that there is no overlap with any other ship. After the placement, players take turns at “firing” at their opponent’s ships. The game proceeds until either one player has sunk all of the opponent’s ships, or each player has completed $r = 2$ rounds of firing. At the end of the game, each player’s payoff is the sum of the values of the opponent’s ships that were sunk, minus the sum of the values of the ships that the player has lost *multiplied by two*. The latter modification makes the game general-sum, and incentivizes players to be more risk-averse.

Goofspiel Our final benchmark is *Goofspiel*, introduced by Ross [1971]. In this game every player has a hand of cards numbered from 1 to h , where in our instantiation $h := 3$. An additional stack of h cards is shuffled and singled out as winning the current prize. In every turn a prize card is revealed, and players privately choose one of their cards to bid. The player with the highest card wins the current prize, while in case of a tie the prize card is discarded. Due to this tie-breaking mechanism, even two-player instances are general-sum. After the completion of h turns, players obtain the sum of the values of the prize cards they have won. Further, the instances we consider are of *limited information*—the actions of the other player are observed only at the end of the game. This makes the game strategically more involved as each player has less information about the opponent’s actions.