
Forward Looking Best-Response Multiplicative Weights Update Methods for Bilinear Zero-sum Games

Michail Fasoulakis^{1,2} Evangelos Markakis² Yannis Pantazis³ Constantinos Varsos^{1,4}

¹Institute of Computer Science, Foundation for Research and Technology - Hellas

²Department of Informatics, Athens University of Economics and Business

³Institute of Applied and Computational Mathematics, Foundation for Research and Technology - Hellas

⁴Computer Science Department, University of Crete

Abstract

Our work focuses on extra gradient learning algorithms for finding Nash equilibria in bilinear zero-sum games. The proposed method, which can be formally considered as a variant of Optimistic Mirror Descent (Mertikopoulos et al., 2019), uses a large learning rate for the intermediate gradient step which essentially leads to computing (approximate) best response strategies against the profile of the previous iteration. Although counter-intuitive at first sight due to the irrationally large, for an iterative algorithm, intermediate learning step, we prove that the method guarantees last-iterate convergence to an equilibrium. Particularly, we show that the algorithm reaches first an $\eta^{1/\rho}$ -approximate Nash equilibrium, with $\rho > 1$, by decreasing the Kullback-Leibler divergence of each iterate by at least $\Omega(\eta^{1+\frac{1}{\rho}})$, for sufficiently small learning rate η , until the method becomes a contracting map, and converges to the exact equilibrium. Furthermore, we perform experimental comparisons with the optimistic variant of the multiplicative weights update method, by (Daskalakis and Panageas, 2019) and show that our algorithm has significant practical potential since it offers substantial gains in terms of accelerated convergence.

1 INTRODUCTION

Our work focuses on the problem of designing learning algorithms for finding Nash equilibria in zero-sum

games. Zero-sum games form a fundamental class of bimatrix games, where the two players need to solve a max-min and a min-max optimization problem respectively, with a bilinear objective function. It is well known by von Neumann’s *minmax theorem*, that these two problems have the same optimum. Apart from their role in the development of game theory, zero-sum games also have a prominent role in optimization, as any linear program can be recast to solving such a game (Adler, 2013). More recently, there has also been a renewed interest in the learning theory community for zero-sum games, given their applications on boosting and reinforcement learning (see (Dai et al., 2018)), and their relevance in formulating GANs in deep learning (as they capture the interaction between the Generator and the Discriminator, see (Goodfellow et al., 2014)).

Although one can solve a zero-sum game by centralized linear programming algorithms, the application areas above highlight the importance of developing fast, iterative learning algorithms. Several approaches have been proposed throughout the past decades starting with *fictitious play* (Robinson, 1951). Recently, some of the more standard methodologies include the family of no-regret algorithms as well as several classes of first-order methods. To mention a few examples, the important class of *Multiplicative Weights Update* (MWU) algorithms (Littlestone and Warmuth, 1994; Freund and Schapire, 1999), together with Gradient Descent, Mirror Descent, and Extra Gradient methods (for a survey see (Bubeck, 2015)), all fall within the above approaches.

In this work, we are interested in methods that exhibit *last-iterate* convergence, a property most desirable from an application point of view, meaning that the strategy profile (x^t, y^t) , reached at iteration t of an iterative algorithm, converges to the actual equilibrium as $t \rightarrow \infty$. Unfortunately, many of the methods mentioned above do not satisfy this. No-regret algorithms, like the MWU method, are known to converge only in

an average sense, resulting in an ε -Nash equilibrium in expectation (see (Arora et al., 2012)) for $\varepsilon > 0$. In fact, it was shown in (Bailey and Piliouras, 2018) that several MWU variants do not satisfy last-iterate convergence. Similarly, the same can be shown for many descent-based methods (see e.g., (Mertikopoulos et al., 2019)).

Driven by these negative results, recent works have focused on certain *optimistic* variations of well known optimization methods. In particular, (Daskalakis and Panageas, 2019) studied a variant of MWU, referred to as the *Optimistic Multiplicative Weights Update* method (OMWU), where an extra negative momentum term is added to correct the dynamics behavior. Their main result is that for zero-sum games with a unique Nash equilibrium, OMWU exhibits last-iterate convergence. Even further, (Mertikopoulos et al., 2019) considered an extra gradient method, named *Optimistic Mirror Descent* (OMD), where last-iterate convergence for a more general class of min-max optimization problems is established. These positive results have generated more interest on the behavior and limitations of such approaches, which is not yet fully understood. Namely, they give rise to further questions, such as: (i) can we prove last-iterate convergence for other related dynamics, and (ii) can we establish faster convergence rates? These questions are the main focus of our work.

1.1 Our Contribution

We introduce a simple yet substantially different variant of Optimistic Mirror Descent method with entropy regularization (Mertikopoulos et al., 2019), for the case of zero-sum games. OMD is an extra gradient method, i.e., it contains an intermediate gradient step before the final update step, and each iteration is characterized by its learning rate parameter, which is the same for both steps (and often the same across all iterations). Our tweak is that the intermediate step uses a different learning rate parameter from the update step in each iteration. In fact, we set this to be sufficiently large, which yields a *game-theoretic interpretation*, namely that we compute (approximate) best response strategies against the profile of the previous iteration, as a look ahead move. Then, during the final update step, we apply multiplicative weights updates by rewarding more the pure strategies that perform better against the best responses that we found in the intermediate step. Consequently, we refer to this OMD variant as Forward-Looking Best-Response - Multiplicative Weights Update (FLBR-MWU) method.

At first sight, this may look counter-intuitive, since learning rates are usually kept small in classic MWU algorithms and, more generally, in any kind of iterative gradient-type optimization algorithms (apart from

the notable exception of (Bailey and Piliouras, 2019)). However, our theoretical and experimental study reveal the following promising findings:

- In Section 3, we investigate theoretically the convergence properties of FLBR-MWU. If η is the standard learning rate parameter used in the update step, and ξ is the corresponding parameter in the intermediate step, then FLBR-MWU exhibits last-iterate convergence for games with a unique equilibrium, when ξ is sufficiently large and $\eta\xi < 1$. Our proof employs a similar methodology to (Daskalakis and Panageas, 2019), adapting convergence tools from the field of dynamical systems. Our method also appears to attain faster convergence, quantified in terms of η , compared to OMD and OMWU. In particular, we prove that the decrease in the divergence from the equilibrium is at least $\Omega(\eta^{1+1/\rho})$ per iteration, for any $\rho > 1$, until we reach an approximate $O(\eta^{1/\rho})$ -equilibrium, by which time, our rule becomes a contraction map (see also Figure 1). This improves on the $\Omega(\eta^3)$ bound established for OMWU in (Daskalakis and Panageas, 2019). Although our bounds do not translate into bounds with respect to time, we suspect a linear convergence rate is highly likely (supported also by our experiments). This has been recently established for OMWU in (Wei et al., 2021), and is left as an open problem for FLBR-MWU.
- In Section 4, we perform numerical experiments, using randomly generated data, comparing FLBR-MWU with OMWU¹. Our experiments reveal that in practice our method achieves indeed a much faster convergence rate, showing an average speedup by a factor of 10 for small size games and up to hundreds, or even higher, for larger games compared to OMWU.

1.2 A Revealing Example

The upper plot in Figure 1 attempts to demonstrate in a qualitative manner the differences we observed in convergence between the proposed FLBR-MWU and the OMWU dynamics. The two phases of the learning dynamics (decrease of divergence, followed by contraction), are highlighted along with the regions of convergence. A quantitative presentation is shown in the lower plots of Figure 1 which depict the convergence behavior of MWU (blue lines), OMWU (red lines) and FLBR-MWU (black lines) for a random realization of a 10×10 payoff matrix with learning rate $\eta = 0.1$. We

¹We note that for the case of zero-sum games, it has been shown in (Wei et al., 2021) that OMWU can be seen as a variant of OMD with entropy regularization.

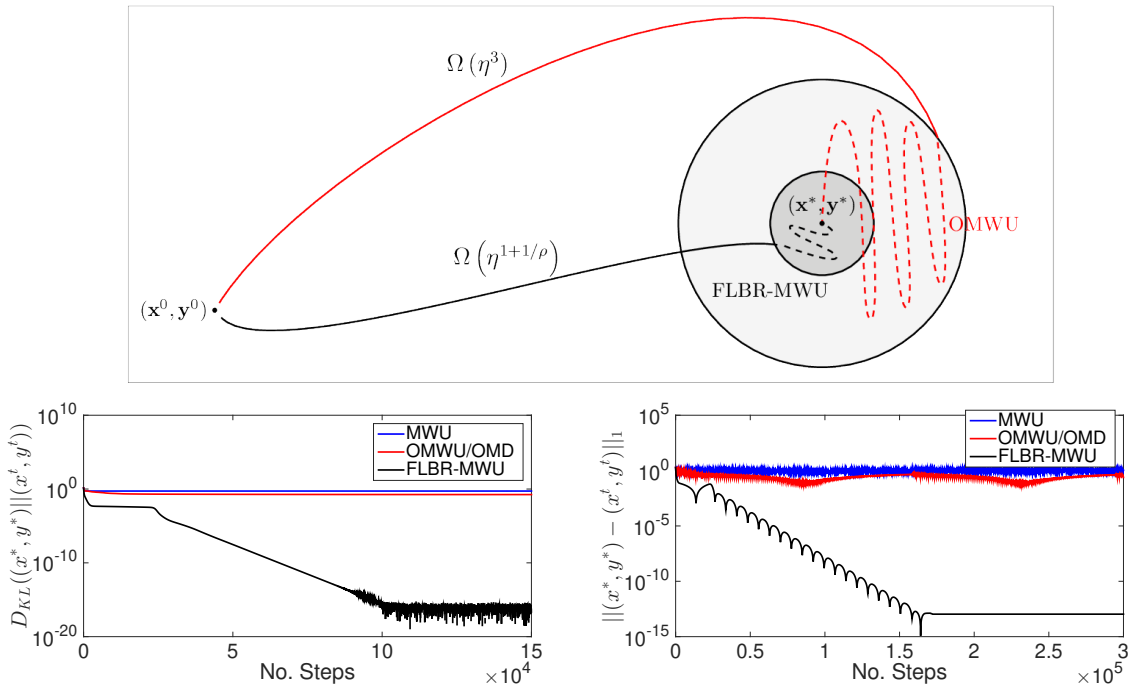


Figure 1: Upper plot: Schematic representation of the convergence path of OMWU (red) and FLBR-MWU (black). Lower plots: A random realization of the learning dynamics for three variants of MWU. The convergence rate for the proposed algorithm (FLBR-MWU) is significantly faster than the existing state-of-the-art algorithms.

provide two measures of convergence, the Kullback-Leibler Divergence (D_{KL}) from the Nash equilibrium (lower left panel), and the respective l_1 norm difference (lower right panel), which reveal different aspects of the dynamics.

As anticipated, MWU fails to converge and a smaller learning rate η would not fix this issue. OMWU does converge but in a very slow pace requiring an enormous number of steps, whereas OMD behaves almost in the same manner as OMWU as expected by (Wei et al., 2021) (OMD is not explicitly shown here nor in Section 4; see Section C in the supplementary material for longer, and more detailed simulations). On the other hand, FLBR-MWU converges up to machine precision, as revealed by both D_{KL} and l_1 metrics. Indeed, FLBR-MWU is able to escape from the D_{KL} plateau (seen in the lower left panel), where the dynamics are moving towards a direction with slow D_{KL} decline, and ultimately converges to the equilibrium in an oscillatory manner with decreasing amplitude (damped oscillations), as is evident from the l_1 -norm difference (lower right panel). Overall, FLBR-MWU has more than one order of magnitude faster convergence rate relative to OMWU and furthermore tolerates larger values for the learning rate, thus the speed of equilibrium computation is significantly accelerated.

1.3 Related Work

To position our paper within the existing literature, the works most related to ours are (Daskalakis and Panageas, 2019) and (Mertikopoulos et al., 2019), regarding the OMWU and the OMD methods, respectively. Conceptually, the definition of our method is closer to (Mertikopoulos et al., 2019) since it uses an extra gradient step. Technically, however, our analysis is based on the mathematical arsenal used in (Daskalakis and Panageas, 2019).

We also overview other relevant works on optimization methods for learning problems. It is well known that most of the standard no-regret algorithms exhibit convergence only in an average sense. Hence, the solution at iteration t , as $t \rightarrow \infty$, may diverge or enter a limit cycle. Such behaviors can be observed, among others, for Gradient Descent/Ascent (GDA) in unconstrained optimization, as well as for MWU methods for constrained problems, see e.g., (Bailey and Piliouras, 2018). Given the importance of achieving last-iterate convergence for applications on learning, such as training GANs, (Daskalakis et al., 2018) and (Liang and Stokes, 2019) studied an optimistic variant of GDA, referred to as OGDA, which has also been considered in previous works, e.g., (Rakhlin and Sridharan, 2013). Their main result states that OGDA exhibits last-iterate convergence for the unconstrained

minmax problem with bilinear functions.

Driven by this, (Daskalakis and Panageas, 2019) proposed to study the constrained version of minmax problems, that corresponds to finding equilibria in zero-sum games. They proposed an optimistic variant of MWU, termed OMWU, and proved that for games with a unique Nash equilibrium, it converges in the last-iterate sense. The sequence of approximations in OMWU uses two previous steps in order to compute the next update, where the extra term corrects the behaviour of the MWU dynamics. Moreover, the performance of OMWU provides a strengthening to the supporting experimental evidence in (Syrkkanis et al., 2015). Very recently, an analysis on the number of required steps, both for OMWU and OGDA, was provided in (Wei et al., 2021), establishing a linear rate of convergence. Finally, an alternative view on the behavior of OMWU by studying volume contraction is given in (Cheung and Piliouras, 2020), and further generalizations have been obtained in (Lei et al., 2021) for convex-concave landscapes.

In parallel to the study of OMWU, the work of (Mertikopoulos et al., 2019) considered a method, where they agglomerate an intermediate approximation with the former state in order to compute the next state. This method is known as Optimistic Mirror Descent (OMD) (Chiang et al., 2012; Rakhlin and Sridharan, 2013), or Mirror-Prox (Nemirovski, 2004). The main result of (Mertikopoulos et al., 2019) is that OMD attains last-iterate convergence for a quite general class of problems that encompasses zero-sum games. A similar approach was introduced by (Gidel et al., 2019) to cope with computational issues. All these techniques fall under the umbrella of *extra-gradient* methods, whose origins date back to (Korpelevich, 1976) (for more details see (Facchinei and Pang, 2003; Bubeck, 2015)).

Several other streams of works have also focused on convergence properties of extra-gradient methods. In (Liang and Stokes, 2019), a linear convergence rate was proved for OMD under the assumption that the game matrix is square and full rank. Under the same assumption, the work of (Mokhtari et al., 2020) considered more general forms of saddle-point problems. Recently, in a different direction, (Cen et al., 2021) developed extra-gradient methods for Quantal Response equilibria, which can be used also for finding an approximate Nash equilibrium. Furthermore, (Hsieh et al., 2019) investigated asymptotic last-iterate convergence for variants of extra-gradient algorithms.

To our knowledge, the idea of using different rates in the intermediate and the update steps of extra gradient methods, has also been used in (Azizian et al., 2020). There are however substantial differences with our work.

Most importantly, (Azizian et al., 2020) involves the unconstrained bilinear case. Even further, their result holds under certain spectral assumptions, and shows only local convergence (starting from a point near the fixed point), whereas we do not need such a condition.

2 BASIC DEFINITIONS

2.1 Zero-sum Games and Approximate Equilibria

We consider finite 2-player zero-sum games, defined by a matrix² $R \in (0, 1]^{n \times n}$, where without loss of generality, we assume both players have n pure strategies. We refer to the two players as the *row player* and the *column player* respectively. If the row player plays the i -th row and the column player plays the j -th column, then the payoff of the row player is R_{ij} , and the payoff of the column player is $-R_{ij}$. We also allow *mixed strategies* as probability distributions (column vectors) on the pure strategies. E.g., a mixed strategy for the row player will be denoted as $x = (x_1, \dots, x_n)$, where x_i is the probability of playing the i -th row. For convenience, we will denote the i -th pure strategy of a player by the unit vector e_i , which has probability one in its i -th coordinate and 0 elsewhere.

A pair (x, y) , where x, y are mixed strategies for the row and the column player respectively, is called a strategy profile. Given such a profile, the expected payoff of the row player is $x^T R y$, whereas for the column player, it is $-x^T R y$. This is obviously a bilinear function, since it is equivalent to $\sum_{i,j} R_{ij} x_i y_j$.

The fundamental solution concept in game theory is that of Nash equilibrium, stating that no player has an incentive to deviate to another strategy.

Definition 1. A strategy profile (x^*, y^*) is a Nash equilibrium in the zero-sum game defined by matrix R , if and only if, for any $i, j \in [n]$,

$$(x^*)^T R y^* \geq e_i^T R y^* \text{ and } (x^*)^T R e_j \geq (x^*)^T R y^*,$$

The payoff of the row player at an equilibrium, $v = (x^*)^T R y^*$, is referred to as the *value* of the game. It is well known that the value of a game and its equilibrium strategies are given by the solution of the following max-min (saddle-point) problem, over the n -dimensional simplex Δ_n : $v = \max_{x \in \Delta_n} \min_{y \in \Delta_n} x^T R y = \min_{y \in \Delta_n} \max_{x \in \Delta_n} x^T R y$.

A useful concept in the analysis of games is the *support* of a mixed strategy x , which is the set of pure strategies that have a positive probability under x , i.e., $\text{supp}(x) =$

²We can always scale appropriately so that the entries are in $(0, 1]$, without affecting the equilibrium strategies.

$\{i : x_i > 0\}$. It is easy to see that at an equilibrium (x^*, y^*) , any pure strategy e_i , with $i \in \text{supp}(x^*)$, is a best response against y^* (resp. for any $j \in \text{supp}(y^*)$, e_j is a best response to x^*).

In our work, we will also need to argue about approximate equilibria to establish convergence. We start with defining approximate best responses. Given a profile (x, y) , we say that a strategy x' is an ε -best-response strategy to y with $\varepsilon \in [0, 1]$, if it yields a payoff that is at most ε less than the best-response payoff. We can define now an approximate equilibrium, as a profile (x, y) where x and y are both approximate best responses to each other. This is precisely the standard notion of additive, approximate equilibria (Nisan et al., 2007).

Definition 2. *A strategy profile (x^*, y^*) is an ε -Nash equilibrium in the zero-sum game defined by matrix R , if and only if, for any i, j*

$$(x^*)^T R y^* \geq e_i^T R y^* - \varepsilon \text{ and } (x^*)^T R y^* \leq (x^*)^T R e_j + \varepsilon$$

2.2 Relevant MWU Variants

One of the standard versions of multiplicative weights update methods results from the FTRL dynamics (Follow-The-Regularized-Leader), when the regularizer is the negative entropy function, (see e.g., (Hoi et al., 2018)). MWU rewards better the pure strategies that perform well against the previous iteration. In particular, if (x^{t-1}, y^{t-1}) is the profile at the end of iteration $t-1$, and η is the learning rate parameter, then for $i \in [n]$, x_i^t is set to be analogous to $x_i^{t-1} \cdot e^{\eta e_i^T R y^{t-1}}$ (with appropriate normalization).

In the remaining paper, we often make comparisons or references to the optimistic variant proposed by (Daskalakis and Panageas, 2019), referred to as OMWU. The idea of "optimism" here is to take into account two previous iterations in order to compute the next update, where the extra term can be seen as a negative momentum, correcting the behaviour of MWU dynamics. The dynamics of OMWU are described below for all $i, j \in [n]$.

$$\begin{aligned} x_i^t &= x_i^{t-1} \cdot \frac{e^{2\eta e_i^T R y^{t-1} - \eta e_i^T R y^{t-2}}}{\sum_{j=1}^n x_j^{t-1} e^{2\eta e_j^T R y^{t-1} - \eta e_j^T R y^{t-2}}}, \\ y_j^t &= y_j^{t-1} \cdot \frac{e^{-2\eta e_j^T R^T x^{t-1} + \eta e_j^T R^T x^{t-2}}}{\sum_{i=1}^n y_i^{t-1} e^{-2\eta e_i^T R^T x^{t-1} + \eta e_i^T R^T x^{t-2}}}. \end{aligned} \quad (1)$$

3 FORWARD LOOKING BEST RESPONSE MULTIPLICATIVE WEIGHTS UPDATE METHOD (FLBR-MWU)

3.1 Definition of the Dynamics

We now present the method studied in this work, which we refer to as Forward Looking Best-Response Multiplicative Weights Update method (FLBR-MWU). We provide first a short description of the main idea behind the dynamics. This is an extra gradient method and each iteration has an intermediate and a final step. Suppose that starting from some initial profile, we reach the profile (x^{t-1}, y^{t-1}) by the end of iteration $t-1$. In the intermediate step of iteration t , we compute a strategy \hat{x}^t for the row player (resp. \hat{y}^t for the column player), which is an approximate best-response strategy to y^{t-1} (resp. to x^{t-1}). This serves as a look ahead step of what would be the currently optimal choices. In the final step of iteration t , we compute the new mixed strategy x^t for the row player, by performing multiplicative weights updates, but after assuming that the opponent was playing \hat{y}^t .

Formally, the first step of the dynamics, denoted as the intermediate best response (IBR) step, is defined below, at iteration t , and for all $i, j \in [n]$, given a non-negative parameter $\xi \in \mathbb{R}^+$ (ξ will be chosen sufficiently large, as will become clear from Lemma 1).

$$\begin{aligned} \hat{x}_i^t &= x_i^{t-1} \cdot \frac{e^{\xi e_i^T R y^{t-1}}}{\sum_{j=1}^n x_j^{t-1} e^{\xi e_j^T R y^{t-1}}}, \\ \hat{y}_j^t &= y_j^{t-1} \cdot \frac{e^{-\xi e_j^T R^T x^{t-1}}}{\sum_{i=1}^n y_i^{t-1} e^{-\xi e_i^T R^T x^{t-1}}}. \end{aligned} \quad (2)$$

The second step, which updates the profile (x^{t-1}, y^{t-1}) to (x^t, y^t) is below, given the learning rate parameter $\eta \in (0, 1)$. We assume that we use the same fixed constants η and ξ in all iterations³.

$$\begin{aligned} x_i^t &= x_i^{t-1} \cdot \frac{e^{\eta e_i^T R \hat{y}^t}}{\sum_{j=1}^n x_j^{t-1} e^{\eta e_j^T R \hat{y}^t}}, \\ y_j^t &= y_j^{t-1} \cdot \frac{e^{-\eta e_j^T R^T \hat{x}^t}}{\sum_{i=1}^n y_i^{t-1} e^{-\eta e_i^T R^T \hat{x}^t}}. \end{aligned} \quad (3)$$

Remark 1. *By setting $\xi = \eta$ in Equation (2) above, the proposed method becomes the same as OMD with entropic regularization (Mertikopoulos et al., 2019), which can also be viewed as OMWU (Wei et al., 2021). In our method however, η and ξ differ substantially across both our theoretical and experimental results.*

³It is an interesting topic for future work, to examine adaptive schemes for ξ and η throughout the iterations.

3.2 Main Results

We consider games with a unique Nash equilibrium, as in (Daskalakis and Panageas, 2019), since it has been argued that the set of zero-sum games with non-unique equilibrium has *Lebesgue measure* equal to zero (van Damme, 1991). For convenience, we also assume that the initial strategy profile consists of the uniform distribution for each player. However, our results hold for any fully-mixed initial profile, with a non-zero probability to all pure strategies.

The main result of our work is the following theorem.

Theorem 1. *Consider a zero-sum game with a unique Nash equilibrium (x^*, y^*) . Starting with the uniform distribution for each player, the FLBR-MWU dynamics attain last-iterate convergence to the Nash equilibrium, i.e., $\lim_{t \rightarrow \infty} (x^t, y^t) = (x^*, y^*)$, when η is chosen sufficiently small, and for big enough ξ , so that $\eta\xi < 1$.*

The goal of the remaining section is to establish the proof of Theorem 1. Towards this, we start with the choice of ξ . The next lemma provides the important observation, that as $\xi \rightarrow \infty$, the strategy \hat{x}^t , computed in the first step of iteration t , becomes a best response against y^{t-1} (analogously for \hat{y}^t).

Lemma 1. *Given any $t > 0$, let \hat{x}^t, \hat{y}^t , be the strategies produced by the first step of iteration t . As $\xi \rightarrow +\infty$, then \hat{x}^t converges to a best-response strategy against y^{t-1} (similarly for \hat{y}^t against x^{t-1}).*

The proof of the lemma can be found in the supplementary material. In the sequel, we assume that ξ has been chosen sufficiently large, so that \hat{x}^t is an ϵ -best response with $\epsilon \rightarrow 0$. For appropriate choices of ξ in practice, we refer to the discussion in Section 4.

The proof of Theorem 1 is split into 3 parts. The first part establishes that after a certain number of iterations, the dynamics reach a profile (x^t, y^t) , that is an $O(\eta^{1/\rho})$ -Nash equilibrium with $\rho > 1$. The second part shows that the profile (x^t, y^t) lies within a neighborhood of the actual equilibrium (x^*, y^*) . Finally, the last part shows that the update rule of FLBR-MWU is a contracting map, i.e., once we are within a neighborhood of (x^*, y^*) , the dynamics converge to their fixed point, which directly implies last-iterate convergence. These three parts are established in Theorems 2, 3 and 5 respectively. The structure of the proof is similar to the convergence proof of OMWU in (Daskalakis and Panageas, 2019). There are however differences in various parts of the analysis. Most importantly, in the first part, we are able to establish a better convergence rate to an approximate equilibrium, whereas OMWU achieves an $\Omega(\eta^3)$ decrease rate. Furthermore, in the third part, the analysis of our Jacobian matrix (proof of Theorem 5) is also different since we are analyzing

sufficiently different dynamics.

To proceed with the first part of the proof, we will use the *Kullback-Leibler (KL) divergence* as a measure of progress. The KL divergence quantifies the similarity between two distributions, and here we will consider the divergence between a profile (x^t, y^t) and the equilibrium (x^*, y^*) , which equals:

$$D_{KL}((x^*, y^*) || (x^t, y^t)) = \sum_{i=1}^n x_i^* \ln(x_i^*/x_i^t) + \sum_{j=1}^n y_j^* \ln(y_j^*/y_j^t). \quad (4)$$

Note that by the initialization and the definition of the dynamics, $x_i^t > 0, y_j^t > 0$ for any given t , and any i, j , so that the logarithmic terms above are well-defined.

Theorem 2. *Consider a zero-sum game with a unique Nash equilibrium (x^*, y^*) . Assume that we run the FLBR-MWU dynamics with the uniform distribution as the initial strategy for both players, and using a sufficiently small η and a big enough ξ . Then, for any $\rho > 1$, the KL divergence $D_{KL}((x^*, y^*) || (x^t, y^t))$ decreases at every iteration with a rate of at least $\Omega(\eta^{1+1/\rho})$, until we reach an $O(\eta^{1/\rho})$ -Nash equilibrium of the game.*

Proof. Let (x^*, y^*) be the Nash equilibrium of the game, and let v be the value of the game, $v = (x^*)^T R y^*$. We take the difference of the KL divergence between two consecutive iterations:

$$D_{KL}((x^*, y^*) || (x^t, y^t)) - D_{KL}((x^*, y^*) || (x^{t-1}, y^{t-1})) = - \left(\sum_{i=1}^n x_i^* \ln(x_i^t/x_i^{t-1}) + \sum_{j=1}^n y_j^* \ln(y_j^t/y_j^{t-1}) \right).$$

We show that this difference is negative and we quantify the decrease in the KL divergence, until we reach an $O(\eta^{1/\rho})$ -Nash equilibrium. Analytically, we have that

$$\begin{aligned} & D_{KL}((x^*, y^*) || (x^t, y^t)) - D_{KL}((x^*, y^*) || (x^{t-1}, y^{t-1})) \\ &= - \sum_{i=1}^n x_i^* \ln e^{\eta e_i^T R \hat{y}^t} + \ln \left(\sum_{i=1}^n x_i^{t-1} e^{\eta e_i^T R \hat{y}^t} \right) \\ &\quad - \sum_{j=1}^n y_j^* \ln e^{-\eta e_j^T R^T \hat{x}^t} + \ln \left(\sum_{j=1}^n y_j^{t-1} e^{-\eta e_j^T R^T \hat{x}^t} \right) \\ &= -\eta (x^*)^T R \hat{y}^t + \eta (y^*)^T R^T \hat{x}^t \\ &\quad + \eta (x^{t-1})^T R y^{t-1} - \eta (y^{t-1})^T R^T x^{t-1} \\ &\quad + \ln \left(\sum_{i=1}^n x_i^{t-1} e^{\eta e_i^T R \hat{y}^t - \eta (x^{t-1})^T R y^{t-1}} \right) \\ &\quad + \ln \left(\sum_{j=1}^n y_j^{t-1} e^{-\eta e_j^T R^T \hat{x}^t + \eta (y^{t-1})^T R^T x^{t-1}} \right). \end{aligned}$$

Notice that in the last expression above, the third term, $\eta (x^{t-1})^T R y^{t-1}$, cancels out with the fourth term. Also, since (x^*, y^*) is an equilibrium, it holds that $x^{*T} R \hat{y}^t \geq v$ and $(y^*)^T R^T \hat{x}^t \leq v$. Therefore, the first

and second terms also cancel out and yield an upper bound with the two logarithmic terms.

We now apply the Taylor expansion of e^x . For convenience, let also $p_i(\eta) = \eta(e_i^T R \hat{y}^t - (x^{t-1})^T R y^{t-1})$, and let $q_j(\eta) = \eta(-e_j^T R^T \hat{x}^t + (x^{t-1})^T R^T y^{t-1})$. Using these abbreviations, the difference of the KL divergences is upper bounded by

$$\begin{aligned} & \ln \left(1 + \eta((x^{t-1})^T R \hat{y}^t - (x^{t-1})^T R y^{t-1}) \right) \\ & + \sum_{i=1}^n x_i^{t-1} \sum_{k=2}^{\infty} \frac{(p_i(\eta))^k}{k!} + \ln \left(1 + \eta(-(y^{t-1})^T R^T \hat{x}^t \right. \\ & \left. + (y^{t-1})^T R^T x^{t-1}) + \sum_{j=1}^n y_j^{t-1} \sum_{k=2}^{\infty} \frac{(q_j(\eta))^k}{k!} \right). \end{aligned}$$

It is easy to see that $|p_i(\eta)| \leq \eta$ and $|q_j(\eta)| \leq \eta$. This means that for any $k \geq 2$ (i.e., for both odd and even values of k), $(p_i(\eta))^k \leq \eta^k$ and $(q_j(\eta))^k \leq \eta^k$. By using the geometric series, we have that $\sum_{k=2}^{\infty} \frac{(p_i(\eta))^k}{k!} \leq \eta^2/(1-\eta)$, and similarly for the series concerning $q_j(\eta)$. If we also use the inequality $\ln(x) \leq x - 1$, we obtain the following sequence of steps.

$$\begin{aligned} & D_{KL}((x^*, y^*) || (x^t, y^t)) - D_{KL}((x^*, y^*) || (x^{t-1}, y^{t-1})) \\ & \leq \ln \left(1 + \eta((x^{t-1})^T R \hat{y}^t - (x^{t-1})^T R y^{t-1}) + \frac{\eta^2}{(1-\eta)} \right) \\ & + \ln \left(1 + \eta(-(y^{t-1})^T R^T \hat{x}^t + (y^{t-1})^T R^T x^{t-1}) \right. \\ & \left. + \frac{\eta^2}{(1-\eta)} \right) \leq \eta((x^{t-1})^T R \hat{y}^t - (x^{t-1})^T R y^{t-1} \\ & + (x^{t-1})^T R y^{t-1} - (\hat{x}^t)^T R y^{t-1}) + 2 \frac{\eta^2}{(1-\eta)} \\ & = -\eta(\varepsilon_1 + \varepsilon_2) + 4\eta^2 \leq -\eta(\max\{\varepsilon_1, \varepsilon_2\}) + 4\eta^2, \end{aligned} \quad (5)$$

where $\varepsilon_1 = (\hat{x}^t)^T R y^{t-1} - (x^{t-1})^T R y^{t-1}$, $\varepsilon_2 = (x^{t-1})^T R \hat{y}^t - (x^{t-1})^T R y^{t-1}$, and the last inequality holds because $\eta \leq 1/2$. Let us look now more carefully at ε_1 (an analogous argument holds for ε_2). The term ε_1 expresses the additional benefit for the row player, if at the profile (x^{t-1}, y^{t-1}) , she deviates to \hat{x}^t . By Lemma 1, we know that as $\xi \rightarrow \infty$, then \hat{x}^t tends to her best response against y^{t-1} . Hence when we select ξ sufficiently large, ε_1 tends to the best possible deviation gain of the row player at the profile (x^{t-1}, y^{t-1}) (resp. for ε_2 and the column player).

To finish the proof, suppose that the profile (x^{t-1}, y^{t-1}) is not an $O(\eta^{1/\rho})$ -Nash equilibrium. Then there exists a deviation that provides additional gain of $\Omega(\eta^{1/\rho})$ to one of the players. This implies that $\max\{\varepsilon_1, \varepsilon_2\} = \Omega(\eta^{1/\rho})$. Hence, by (5), and since $\eta < 1$, we can see that as long as we have not reached an $O(\eta^{1/\rho})$ -Nash

equilibrium, the KL divergence will keep decreasing by at least $\eta \Omega(\eta^{1/\rho}) - 4\eta^2 = \Omega(\eta^{1+1/\rho})$. As the KL divergence cannot decrease forever, eventually, our dynamics will reach an $O(\eta^{1/\rho})$ -Nash equilibrium. \square

Consider now the first iteration t of the dynamics, where (x^t, y^t) forms an $O(\eta^{1/\rho})$ -Nash equilibrium for some fixed $\rho > 1$. The next step is to show that if we make η small enough, this profile falls within a neighborhood of the equilibrium (x^*, y^*) .

Theorem 3. *Let (x^*, y^*) be the unique Nash equilibrium of the zero-sum game, and let (x^t, y^t) be the first profile reached by the dynamics that forms an $O(\eta^{1/\rho})$ -Nash equilibrium for some $\rho > 1$. Then*

$$\lim_{\eta \rightarrow 0} \|(x^*, y^*) - (x^t, y^t)\|_1 = 0,$$

The proof of Theorem 3 can be found at the supplementary material, Section A.

The next and final step of our proof is to show that our dynamics induce a contracting map. An update rule with a fixed point x is called a contraction, if there exists a region U around x , such that for any starting point in U , the rule converges to its fixed point as $t \rightarrow \infty$. In our case, the Nash equilibrium (x^*, y^*) of the game is a fixed point of the FLBR-MWU dynamics and Theorem 3 guarantees that we can reach a neighborhood around (x^*, y^*) . To proceed, we state a sufficient condition for a dynamical system to converge to its fixed point.

Theorem 4 (see (Galor, 2007)). *Let x^* be a fixed point for the dynamical system $x^{(t+1)} = g(x^{(t)})$. If all eigenvalues of the Jacobian matrix of g at x^* have absolute value less than one, then there exists a neighborhood U of x^* such that for all $x \in U$, g converges to x^* , starting from x .*

Using Theorem 4, we show the following theorem, whose proof can be found at the supplementary material, Section A.

Theorem 5. *The update rule of FLBR-MWU is a contraction, as long as $\eta\xi < 1$, i.e. $\lim_{t \rightarrow \infty} (x^t, y^t) = (x^*, y^*)$.*

4 NUMERICAL EXPERIMENTS

We note that additional supporting figures and elaboration on our experiments can be found in the supplementary material (Section C).

Nash equilibrium estimation. In order to make comparisons, we need first to compute the equilibria of the generated instances. Instead of using a linear programming solver, the equilibrium computation is performed using the proposed FLBR-MWU

algorithm with $\eta = 0.05$. FLBR-MWU is an iterative approach thus a convergence criterion to ensure that the Nash equilibrium has been reached is required. We propose as a convergence criterion the D_{KL} between the update step and the IBR step of our dynamics: $D_{KL}((x^t, y^t) || (\hat{x}^t, \hat{y}^t))$. This metric is sufficient because the best response strategy at Nash equilibrium is exactly the equilibrium strategy, thus $\lim_{t \rightarrow \infty} D_{KL}((x^t, y^t) || (\hat{x}^t, \hat{y}^t)) = 0$ (for small enough η). We return the solution when the convergence criterion becomes 10^{-15} , which is approximately the machine’s arithmetic precision, or when the maximum number of steps, denoted by t_{\max} , –typically millions of steps– has been reached. In the infrequent latter case (it happened in less than 0.1% of the instances), we discard the returned solution.

Effect of the intermediate rate (ξ). In our learning dynamics, the best response strategy is approximated by the softmax function (a.k.a. the normalized exponential function or the Gibbs measure in statistical physics). Sending ξ to infinity, one out of the potentially many best response strategies is obtained as intermediate dynamics by Equation (2). However, ξ should be finite from a practical point of view. Since it appears at the exponentials’ argument, very high values of ξ may result in arithmetic imprecision. Therefore, we conducted a numerical study to assess the effect of ξ on the convergence of the algorithm. Table 1 presents various statistics about the number of steps required for several values of ξ and for two values of the size of the payoff matrix R , with $\eta = 0.1$. We average over 10^3 repetitions using random payoff matrices, whose elements are iid sampled from $\mathcal{U}([0, 1])$. Evidently, as ξ increases, the FLBR-MWU dynamics require fewer steps in order to reach a specific threshold of accuracy (set to 10^{-10} for the D_{KL} between the Nash equilibrium and the FLBR-MWU dynamics). However, the solution occasionally produces ‘NaN’ for values of ξ above 200, due to overflow in the exponentials⁴. Overall, values between 50 and 100 are a sufficient compromise between the best response approximation and machine precision trade-off. We also note that even when we select values that violate the condition $\eta\xi < 1$ (which we needed for our theoretical results), we still attain convergence in most cases. In the remaining experiments of this section, we set $\xi = 100$, even though larger values can be tolerated especially when both $n \gg 1$ and $x_i^*, y_j^* \ll 1$ hold.

Effect of the learning rate (η). The first row of panels in Figure 2 shows the D_{KL} between the Nash equilibrium and the FLBR-MWU dynamics for the same payoff matrix instance as in Figure 1 and for

Table 1: Statistics on the number of steps till convergence for various values of ξ and n . The maximum number of steps was set to $t_{\max} = 10^6$.

Matrix size	Statistic	$\xi = 20$	$\xi = 50$	$\xi = 100$	$\xi = 200$
$n = 10$	Mean	85.9K	52.7K	41.6K	57.6K
	Median	32.0K	22.3K	19.8K	18.0K
	t_{\max} was hit	0.9%	0.0%	0.0%	2.2%
$n = 20$	Mean	352.1K	233.3K	173.6K	141.4K
	Median	225.0K	123.4K	82.2K	65.6K
	t_{\max} was hit	13.5%	5.9%	3.5%	2.5%

various values of the learning rate, η . The difference between the left and right panels is that for the right column of panels, the x-axis has been rescaled by multiplying each run with the respective learning rate. A linear scaling is numerically observed showing that the number of steps is effectively of order $O(\eta^{-1})$ for a fixed accuracy level. This inversely-proportional behavior is observed not only during the convergence to the approximate Nash equilibrium, but also during the contraction period. As a rule of thumb, we propose to increase the rate η , because it accelerates the convergence, but with caution since a very large η might result in an oscillatory solution, thus failing to converge (blue line in second row of panels).




Figure 2: The D_{KL} between the Nash equilibrium and the FLBR-MWU dynamics for two instances, with no rescaling of the x-axis (left panels) and with rescaling (right panels). The relationship between the number of steps and learning rate is inversely proportional.

Effect of the payoff matrix size (n). The rate of convergence is sensitive to the size of the payoff matrix and the number of steps is expected to substantially increase on average as the size of the game increases. We performed a numerical comparison between FLBR-MWU and OMWU to evaluate the number of steps required to achieve a predefined level of accuracy. Table 2 presents statistics on the number of steps for each learning algorithm computed on 100 repetitions

⁴Overflow can be easily fixed by subtracting the maximum value but with an increased underflow risk.

using element-wise uniformly-sampled and iid random payoff matrices. The learning rate was set to $\eta = 0.1$. Given that FLBR-MWU requires almost twice as many calculations per iteration, relative to OMWU, it is fair to multiply the number of steps of FLBR-MWU with two and then compare it with the number of steps of OMWU. We observe that FLBR-MWU is approximately 15 times faster on average when $n = 5$. As the size of the payoff matrix increases, the performance gap in convergence rate as measured by the number of steps also increases. Indeed, even for $n = 10$, OMWU requires more than 4.2M steps in half of the runs, while the respective number for FLBR-MWU is 16.3K, implying that FLBR-MWU is 100 times faster than OMWU in the median sense. Larger game sizes make OMWU essentially impractical while FLBR-MWU is still able to converge in less than 5M steps.

Table 2: Statistics on the number of steps till convergence for various sizes of the game. The maximum number of steps was set to $t_{\max} = 5 \times 10^6$.

Learning alg.	Statistic	$n = 5$	$n = 10$	$n = 50$
FLBR-MWU	Mean	33.7K	103.3K	984.9K
	Median	9.8K	16.3K	409.3K
	t_{\max} was hit	0.0%	0.0%	1.0%
OMWU	Mean	1088.8K	3323.2K	5000.0K
	Median	353.8K	4208.1K	5000.0K
	t_{\max} was hit	9.0%	46.0%	100.0%

Acknowledgements

This work has been supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the "First Call for H.F.R.I. Research Projects to support faculty members and researchers and the procurement of high-cost research equipment" grant (Project Number: HFRI-FM17-3512). We are also grateful to the anonymous reviewers of AISTATS 2022 for their important feedback on improving the presentation of our work.

References

- Ilan Adler. The equivalence of linear programs and zero-sum games. *Int. J. Game Theory*, 42(1):165–177, 2013.
- Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory Comput.*, 8(1):121–164, 2012.
- Waïss Azizian, Damien Scieur, Ioannis Mitliagkas, Simon Lacoste-Julien, and Gauthier Gidel. Accelerating smooth games by manipulating spectral shapes. In *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020*, volume 108, pages 1705–1715. PMLR, 2020.
- James P. Bailey and Georgios Piliouras. Multiplicative weights update in zero-sum games. In *Proceedings of the Conference on Economics and Computation (EC’18)*, pages 321–338, 2018.
- 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: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019*, pages 12977–12987, 2019.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.*, 8(3-4): 231–357, 2015.
- Shicong Cen, Yuting Wei, and Yuejie Chi. Fast policy extragradient methods for competitive games with entropy regularization. *ArXiv*, abs/2105.15186, 2021.
- Yun Kuen Cheung and Georgios Piliouras. Chaos, extremism and optimism: Volume analysis of learning in games. In *Proceedings of Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, (NeurIPS’20)*, 2020.
- Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Proceedings of COLT’12 - The 25th Annual Conference on Learning Theory*, pages 6.1–6.20, 2012.
- Bo Dai, Albert Shaw, Lihong Li, Lin Xiao, Niao He, Zhen Liu, Jianshu Chen, and Le Song. SBEED: convergent reinforcement learning with nonlinear function approximation. In *Proceedings of the 35th International Conference on Machine Learning, ICML’18*, pages 1133–1142, 2018.
- Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *Proceedings of the ITCS’19*, 2019.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training GANs with optimism. In *Proceedings of the International Conference on Learning Representations (ICLR’18)*, 2018.
- Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM J. Comput.*, 39(6):2531–2597, 2010.
- Francisco Facchinei and J. S. Pang. Finite-dimensional variational inequalities and complementarity problems. 2003.
- Yoan Freund and Robert E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.

- Oded Galor. *Discrete Dynamical Systems*. Springer-Verlag, 2007.
- Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. In *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Proceedings of Annual Conference on Neural Information Processing Systems (NIPS '14)*, pages 2672–2680, 2014.
- Steven C. H. Hoi, Doyen Sahoo, Jing Lu, and Peilin Zhao. Online learning: A comprehensive survey. *CoRR*, abs/1802.02871, 2018.
- Yu-Guan Hsieh, F. Iutzeler, J. Malick, and P. Mertikopoulos. On the convergence of single-call stochastic extra-gradient methods. In *NeurIPS*, 2019.
- G. M. Korpelevich. The extragradient method for finding saddle points and other problems. 1976.
- Qi Lei, Sai Ganesh Nagarajan, Ioannis Panageas, and Xiao Wang. Last iterate convergence in no-regret learning: constrained min-max optimization for convex-concave landscapes. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics, AISTATS'21*, pages 1441–1449, 2021.
- Tengyuan Liang and James Stokes. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *Proceedings of The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS'19*, pages 907–915, 2019.
- Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Inf. Comput.*, 108(2): 212–261, 1994.
- 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 *Proceedings of the 7th International Conference on Learning Representations, ICLR'19*, 2019.
- 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 *Proceedings of The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS'20*, pages 1497–1507, 2020.
- Arkadi Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM J. Optim.*, 15: 229–251, 2004.
- Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge Univ. Press, 2007.
- Henk Norde. Bimatrix games have quasi-strict equilibria. *Math. Program.*, 85(1):35–49, 1999.
- Alfio Quarteroni, Riccardo Sacco, and Fausto Saleri. *Numerical Mathematics (Texts in Applied Mathematics)*. Springer-Verlag, 2006.
- Sasha Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 3066–3074, 2013.
- Julia Robinson. An iterative method of solving a game. *Annals of Mathematics*, 54(2):296–301, 1951.
- Gilbert Strang. *Introduction to Linear Algebra*. Wellesley-Cambridge Press, 2009.
- Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E. Schapire. Fast convergence of regularized learning in games. In *Proceedings of Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015*, pages 2989–2997, 2015.
- Eric van Damme. *Stability and Perfection of Nash Equilibria*. Springer-Verlag, 1991.
- Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. In *Proceedings of the 9th International Conference on Learning Representations ICLR '21*, 2021.

Supplementary Material: Forward Looking Best-Response Multiplicative Weights Update Methods for Bilinear Zero-sum Games

A MISSING PROOFS

A.1 Proof of Lemma 1

Fix t and let us consider the formula that produces the coordinates of \hat{x}^t , given x^{t-1}, y^{t-1} . For simplicity in writing, we drop the superscript $t-1$ and refer to x, y as the strategies of the two players computed at the end of iteration $t-1$. Focusing on the row player (the same argument follows for the column player too), we know that for every $i \in [n]$,

$$\hat{x}_i^t = x_i \cdot \frac{e^{\xi e_i^T R y}}{\sum_{j=1}^n x_j e^{\xi e_j^T R y}}.$$

We want to compute for every $i \in [n]$, the limit $\lim_{\xi \rightarrow \infty} \hat{x}_i^t$. We distinguish two cases, depending on whether e_i is a best-response strategy against y or not. Denote by $B(y)$ the set of the pure best-response strategies of the row player against y , i.e., $B(y) = \{i : e_i \text{ is a best response against } y\}$.

We start with the first case, where $i \in B(y)$. This means that $e_i^T R y$ is a best-response payoff against y and therefore $e^{\xi e_i^T R y} = e^{\xi e_j^T R y}$ for any $j \in B(y)$. Based on this, we can now write the limit as

$$\lim_{\xi \rightarrow \infty} \hat{x}_i^t = x_i \cdot \frac{e^{\xi e_i^T R y}}{\sum_{j \in B(y)} x_j e^{\xi e_j^T R y} + \sum_{j \notin B(y)} x_j e^{\xi e_j^T R y}} = x_i \cdot \frac{1}{\sum_{j \in B(y)} x_j + \sum_{j \notin B(y)} x_j e^{\xi(e_j^T R y - e_i^T R y)}},$$

Obviously, we can see that $(e_j^T R y - e_i^T R y) < 0$ for any $j \notin B(y)$. Hence, as ξ goes to infinity, it holds that $\lim_{\xi \rightarrow \infty} = \sum_{j \notin B(y)} x_j e^{\xi(e_j^T R y - e_i^T R y)} = 0$. Thus, we have that

$$\lim_{\xi \rightarrow \infty} \hat{x}_i^t = \frac{x_i}{\sum_{j \in B(y)} x_j}.$$

Consider now the second case where $i \notin B(y)$. Let $A(y)$ be the set of indices j , such that $e^{\xi e_j^T R y} > e^{\xi e_i^T R y}$. Since $i \notin B(y)$, we know that $A(y) \neq \emptyset$. Then, we have

$$\lim_{\xi \rightarrow \infty} \hat{x}_i^t = x_i \cdot \frac{e^{\xi e_i^T R y}}{\sum_{j \in A(y)} x_j e^{\xi e_j^T R y} + \sum_{j \notin A(y)} x_j e^{\xi e_j^T R y}} = x_i \cdot \frac{1}{\sum_{j \in A(y)} x_j e^{\xi(e_j^T R y - e_i^T R y)} + \sum_{j \notin A(y)} x_j e^{\xi(e_j^T R y - e_i^T R y)}}.$$

As ξ goes to infinity, it holds that $\lim_{\xi \rightarrow \infty} = \sum_{j \in A(y)} x_j e^{\xi(e_j^T R y - e_i^T R y)} = +\infty$. For the second term of the denominator, we can see that for the indices $j \notin A(y)$, for which $e^{\xi e_j^T R y} < e^{\xi e_i^T R y}$, it holds that $\lim_{\xi \rightarrow \infty} = \sum_j x_j e^{\xi(e_j^T R y - e_i^T R y)} = 0$. For the remaining indices j , for which $e^{\xi e_j^T R y} = e^{\xi e_i^T R y}$, the limit is $\sum_j x_j \leq 1$. Thus, in total, we have that

$$\lim_{\xi \rightarrow \infty} \hat{x}_i^t = 0.$$

Hence, when $\xi \rightarrow \infty$, the strategy \hat{x} will eventually contain in its support only best responses to y , and therefore \hat{x} will form a best response to y as well.

A.2 Proof of Theorem 3

The proof is based on the following lemma, shown in (Etessami and Yannakakis, 2010), which we state here for the case of zero-sum games:

Lemma 3. *Consider a zero-sum game given by matrix R with a unique Nash equilibrium (x^*, y^*) , and let $|R|$ be the number of bits needed for the representation of R . There exists a polynomial p such that for every $\delta > 0$, every ε -Nash equilibrium (x, y) satisfies that $|x_i^* - x_i| < \delta$, as long as $\varepsilon \leq 1/2^{p(|R|+size(\delta))}$, where $size(\delta) = O(\log(1/\delta))$ is the number of bits needed for representing δ .*

By the assumptions in the statement of Theorem 3, we fix $\varepsilon = c \cdot \eta^{1/\rho}$, for some constant c , so that (x^t, y^t) is an ε -Nash equilibrium. We claim that there exists $\delta(\eta)$ such that ε and $\delta(\eta)$ satisfy the inequality stated in Lemma 3. In particular, by looking more carefully at the desired inequality and solving with respect to δ , one can construct a function $\delta(\eta)$, such that for the given ε we have selected, it holds that

$$\varepsilon \leq 1/2^{p(|R|+size(\delta(\eta)))} \text{ and } \lim_{\eta \rightarrow 0} \delta(\eta) = 0.$$

Hence, we can now apply Lemma 3 and obtain that for any ε -Nash equilibrium (x, y) we have that $|x_i^* - x_i| \leq \delta(\eta)$ and $|y_i^* - y_i| \leq \delta(\eta)$. The proof now of Theorem 3 is immediate, since $\|(x^*, y^*) - (x^t, y^t)\|_1 = \sum_{i=1}^n |x_i^* - x_i| + \sum_{i=1}^n |y_i^* - y_i| \leq 2n \cdot \delta(\eta)$, which goes to 0 as $\eta \rightarrow 0$.

A.3 Proof of Theorem 5

To prove the theorem, we describe first a discrete dynamical system that captures the FLBR-MWU dynamics, and we will prove that for an appropriate norm of the Jacobian matrix of the system, its value is less than one⁵. The update rule φ of FLBR-MWU is

$$\begin{aligned} \varphi(x, y) &= (\varphi_1(x, y), \varphi_2(x, y)), \text{ where} \\ \varphi_{1,i}(x, y) &= (\varphi_1(x, y))_i = x_i \frac{e^{\eta e_i^T R f(x, y)}}{\sum_{\ell} x_{\ell} e^{\eta e_{\ell}^T R f(x, y)}}, \\ \varphi_{2,i}(x, y) &= (\varphi_2(x, y))_i = y_i \frac{e^{-\eta e_i^T R^T h(x, y)}}{\sum_{\ell} y_{\ell} e^{-\eta e_{\ell}^T R^T h(x, y)}}, \end{aligned} \quad (6)$$

where $f(x, y)$ and $h(x, y)$ are column vectors with $(f(x, y))_i = y_i \frac{e^{-\xi e_i^T R^T x}}{\sum_{\ell} y_{\ell} e^{-\xi e_{\ell}^T R^T x}}$, and $(h(x, y))_i = x_i \frac{e^{\xi e_i^T R y}}{\sum_{\ell} x_{\ell} e^{\xi e_{\ell}^T R y}}$, for all $i \in \{1, \dots, n\}$.

Clearly, the dynamics of FLBR-MWU are captured by $(x^{t+1}, y^{t+1}) = \varphi(x^t, y^t)$. The Jacobian of φ is a $2n \times 2n$ matrix, which can be written in the form of a 2×2 block matrix, as follows:

$$J = \begin{pmatrix} \frac{\partial \varphi_1}{\partial \mathbf{x}} & \frac{\partial \varphi_1}{\partial \mathbf{y}} \\ \frac{\partial \varphi_2}{\partial \mathbf{x}} & \frac{\partial \varphi_2}{\partial \mathbf{y}} \end{pmatrix}. \quad (7)$$

In order to use Theorem 4 and prove that ϕ is a contraction, we need to argue about the eigenvalues of J at the equilibrium (x^*, y^*) . Towards this, in Section B, we provide the exact form of each entry of J at (x^*, y^*) (after some simplification steps by exploiting the fact that (x^*, y^*) is an equilibrium).

We analyze first the eigenvalues that are derived by the rows of J that correspond to $\varphi_{1,i}$ for some $i \notin \text{supp}(x^*)$ and to $\varphi_{2,i}$ for some $i \notin \text{supp}(y^*)$. Let $x^{*T} R y^* = v$ be the value of the game. By referring to Subsection B, we have that for any $i \notin \text{supp}(x^*)$:

$$\frac{\partial \varphi_{1,i}}{\partial x_i}(x^*, y^*) = \frac{e^{\eta e_i^T R y^*}}{e^{\eta v}}, \quad \frac{\partial \varphi_{1,i}}{\partial x_j}(x^*, y^*) = 0 \text{ for any } i \neq j, \text{ and } \frac{\partial \varphi_{1,i}}{\partial y_j}(x^*, y^*) = 0, \text{ for any } j.$$

Hence, the i -th row of the upper block of J has only one non-zero entry, namely, the diagonal element, provided that $i \notin \text{supp}(x^*)$. Thus, $\frac{e^{\eta e_i^T R y^*}}{e^{\eta v}}$ is an eigenvalue of J at (x^*, y^*) . We note also that⁶ $e_i^T R y^* < v$ for $i \notin \text{supp}(x^*)$,

⁵Besides (Galor, 2007), readers could also refer to Chapter 7 (Quarteroni et al., 2006).

⁶A unique Nash equilibrium of a zero-sum game is also a quasi-strict equilibrium (Theorem 1 in (Norde, 1999)), meaning that strategies that are not in the support of the equilibrium have strictly less payoff than the best-response payoff.

hence $|\frac{\partial \varphi_{1,i}}{\partial x_i}(x^*, y^*)| < 1$. Analogously, for $i \notin \text{supp}(y^*)$ we have that $\frac{\partial \varphi_{2,i}}{\partial y_i}(x^*, y^*) = \frac{e^{-\eta e_i^T R^T x^*}}{e^{-\eta v}}$, whereas all other partial derivatives of $\varphi_{2,i}$ are zero. Thus, $\frac{e^{-\eta e_i^T R^T x^*}}{e^{-\eta v}}$ is also an eigenvalue of J , with $|\frac{e^{-\eta e_i^T R^T x^*}}{e^{-\eta v}}| < 1$, since $e_i^T R^T x^* > v$ for $i \notin \text{supp}(y^*)$, by Footnote 2.

We now focus on the rows and columns that correspond to the support of x^* and y^* . We denote this submatrix as \tilde{J} , with $k_1 = |\text{supp}(x^*)|$, $k_2 = |\text{supp}(y^*)|$ and $k = k_1 + k_2$. Thus, $\tilde{J} \in \mathbb{R}^{k \times k}$. It can be seen that J has eigenvalues with absolute value less than one iff the same holds for \tilde{J} as well.

Using equations (11) and computing $((\mathbf{1}_{k_1}, \mathbf{0}_{k_2})^T \cdot \tilde{J})_j$ for an arbitrary coordinate j , we end up with the quantity $\sum_i x_i^* \sum_k R_{ik} y_k^* R_{kj}^T - \sum_i x_i^* \sum_k x_k^* \sum_l R_{kl} y_l^* R_{lj}^T$, that equals zero. Thus, $(\mathbf{1}_{k_1}, \mathbf{0}_{k_2})$ is a left eigenvector of \tilde{J} corresponding to the zero eigenvalue. Using the same argumentation we have that $(\mathbf{0}_{k_1}, \mathbf{1}_{k_2})$ is also a left eigenvector of \tilde{J} with eigenvalue zero.

We will make use of the following claim, regarding orthogonal pairs of eigenvectors.

Claim 1. *Consider a matrix $A \in \mathbb{R}^{n \times n}$, an eigenvalue λ and a left eigenvector u^T , corresponding to λ . Then for every right eigenvector v that does not correspond to λ , it holds that $u^T v = 0$.*

The proof of the claim, which is a simple linear algebra exercise, is at the end of this section. From Claim 1, it follows that for any right eigenvector (\tilde{x}, \tilde{y}) corresponding to a nonzero eigenvalue, we have

$$\tilde{x}^T \mathbf{1}_{k_1} = 0 \text{ and } \tilde{y}^T \mathbf{1}_{k_2} = 0. \quad (8)$$

With that in hand, let us now rewrite \tilde{J} , as $\tilde{J} = J' + A$, where J' is produced by deleting the term $-x_i^*$ (resp. $-y_i^*$) from every element of the upper left (resp. lower right) block of \tilde{J} . I.e., A contains $-x_i^*$ in all entries of the i -th row in the upper left block, and $-y_i^*$ in all entries of the $(k_1 + i)$ -th row in the bottom right block. The other two blocks of A contain only zeros. Using (8), we can see that for every non-zero eigenvalue λ of \tilde{J} , that corresponds to a right eigenvector (\tilde{x}, \tilde{y}) , it holds that $A \cdot (\tilde{x}, \tilde{y}) = 0$, thus λ is also an eigenvalue of the matrix J' . By the equations in Subsection B, we can write J' as a 2×2 block matrix, as follows.

$$J' = \begin{pmatrix} I_{k_1 \times k_1} + \eta \xi D^{xx} & \eta D^{xy} \\ \eta D^{yx} & I_{k_2 \times k_2} + \eta \xi D^{yy} \end{pmatrix},$$

with

$$D_{ij}^{xx} = -x_i^* \left(\sum_k R_{ik} y_k^* R_{kj}^T - \sum_k x_k^* \sum_l R_{kl} y_l^* R_{lj}^T \right), \text{ with } i, j \in [k_1],$$

$$D_{ij}^{yy} = -y_i^* \left(\sum_k R_{kj} R_{ik}^T x_k^* - \sum_k y_k^* \sum_l R_{kl}^T x_l^* R_{lj} \right), \text{ with } i, j \in [k_2],$$

$$D_{ij}^{yx} = -y_i^* \left(R_{ij}^T - e_j^T R y^* \right) \frac{e^{\xi e_j^T R y^*}}{e^{\xi v}}, \text{ with } i \in [k_2], j \in [k_1],$$

$$D_{ij}^{xy} = x_i^* \left(R_{ij} - e_j^T R^T x^* \right) \frac{e^{-\xi e_j^T R^T x^*}}{e^{-\xi v}}, \text{ with } i \in [k_1], j \in [k_2].$$

We observe that all the entries of the matrices $D^{xx}, D^{yy}, D^{yx}, D^{xy}$ are within the interval $[-1, 1]$. Furthermore, for the remainder of the proof, and without loss of generality, we assume that (x^*, y^*) is a mixed strategy profile, i.e., both x^* and y^* are mixed.⁷

We now consider the diagonal elements of D^{xx} and D^{yy} . For D^{xx} , and for any i , its i -th element along the diagonal is

$$-x_i^* \left(\sum_l R_{il}^2 y_l^* - \sum_k x_k^* \sum_l R_{kl} y_l^* R_{li}^T \right)$$

We establish the following useful property.

⁷If exactly one of x^*, y^* were mixed, this would also imply the existence of a pure equilibrium, contradicting our uniqueness assumption. If (x^*, y^*) is a pure strategy profile, then \tilde{J} is a 2×2 block matrix, where each block is a single element. Using the equations of (11) in Section B, then the matrix \tilde{J} only has the zero eigenvalue and the proof of Theorem 5 follows directly.

Lemma 4. For any $i \in [k_1]$, $D_{ii}^{xx} < 0$, and for any $j \in [k_2]$, $D_{jj}^{yy} < 0$.

Proof of Lemma 4. We first prove that for any i , $D_{ii}^{xx} \leq 0$. For the sake of contradiction, assume that there exists an index i , such that $D_{ii}^{xx} > 0$. This means that

$$\sum_l R_{il}^2 y_l^* < \sum_k x_k^* \sum_l R_{kl} y_l^* R_{il}.$$

To proceed, we claim that

$$v \leq \sum_l R_{il} z_l, \tag{9}$$

where $z_l = \frac{R_{il} y_l^*}{v}$, and $z = (z_l)_{l \in [n]}$. To see this, it is crucial to notice first that both y^* and z are probability vectors and also that $v = \sum_l R_{il} y_l^*$. Hence, the LHS and the RHS of Equation (9) are two different convex combinations of the R_{il} values. To go from the LHS to the RHS, we simply replace y_l^* by z_l . For each R_{il} that is itself less than v , the coefficient y_l^* is replaced by a smaller coefficient, since $z_l < y_l^*$ in this case (by the definition of z_l). On the contrary, for each R_{il} with $R_{il} > v$, it holds that $z_l > y_l^*$ (and we also have $z_l = y_l^*$ when $R_{il} = v$). Hence, we can think of the move from the LHS to the RHS of (9), as transferring probability mass from the lowest valued R_{il} 's to the highest ones. Let Δ be the total amount of probability mass that was transferred. Then $\Delta = \sum_{l: R_{il} < v} (y_l^* - z_l) \geq 0$. Note that it also holds that $\Delta = \sum_{l: R_{il} > v} (z_l - y_l^*)$. If we compare now the LHS with the RHS, the RHS has a deficit of a total value of at most $\Delta \cdot v$ from the terms with $R_{il} < v$, compared to the corresponding terms of the LHS. At the same time, it has a surplus of at least $\Delta \cdot v$ from the terms with $R_{il} > v$. Combining the deficit and the surplus, this proves Equation (9).

Using (9), we can now obtain the following contradiction:

$$v \leq \sum_l R_{il} z_l < x^{*T} R z = v,$$

where the strict inequality above follows by the condition stated just before Equation (9), and the final equality holds since $x^{*T} R e_j = v$ for any $j \in \text{supp}(y^*)$ (and so for any $j \in \text{supp}(z)$).

Thus, we have reached a contradiction, which means that $D_{ii}^{xx} \leq 0$ for every $i \in [k_1]$. In addition, it is not difficult to see that in case $D_{ii}^{xx} = 0$ for some i , the strategy profile (i, y^*) is also a Nash equilibrium. But this would imply that there also exists a pure equilibrium formed by i and a pure best response (from the support of y^*), contradicting the fact that we have a unique equilibrium. Hence, D_{ii}^{xx} is strictly negative for every $i \in [k_1]$.

Similarly, the same analysis holds for the matrix D^{yy} , completing the proof of the lemma. \square

To finish the proof, we estimate an upper bound on the p -norm of J' for $p \in \mathbb{N}$. We have that

$$\begin{aligned} \|J'\|_p^p &= \sum_j \left(\sum_i |J'_{ij}|^p \right) \leq k \max_j \left(\sum_i |J'_{ij}|^p \right) \\ &\leq k \left(|1 + \eta \xi D_{\ell\ell}^{xx}|^p + \eta^p \xi^p \sum_{i=1}^{k_1} |D_{i\ell}^{xx}|^p + \eta^p \sum_{i=1}^{k_2} |D_{i\ell}^{yx}|^p \right) \\ &\leq k \left(|1 + \eta \xi D_{\ell\ell}^{xx}|^p + \eta^p \xi^p k_1 + \eta^p k_2 \right), \end{aligned}$$

where ℓ is the column of J' that achieves the maximum sum, i.e., $\ell = \arg \max_{j \in [k_1+k_2]} \sum_i |J'_{ij}|^p$, and we assumed without loss of generality that ℓ belongs to $\{1, \dots, k_1\}$. We can now see that since $D_{\ell\ell}^{xx}$ is negative, and both $|D_{i\ell}^{xx}|$, $|D_{i\ell}^{yx}|$ are at most equal to one, then if $\eta \xi < 1$, and η is sufficiently small, there exists an appropriate p so that $\|J'\|_p^p < 1$. However, it is well known that the maximum absolute value of an eigenvalue of a matrix is bounded by the induced matrix norms, therefore it suffices to check that $\|J'\| < 1$ for some matrix norm, see (Quarteroni et al., 2006). Thus, the absolute value of the maximum eigenvalue of J' is less than one, and this concludes our proof.

Proof of Claim 1. Consider two distinct eigenvalues of A λ_1 and λ_2 , such that v is the corresponding to λ_1 left eigenvector, while u is the corresponding to λ_2 right eigenvector ((Strang, 2009)). In other words, v is the corresponding to λ_1 right eigenvector for A^T . We observe that, $v^T(A^T u) = (v^T A^T)u = (Av)^T u$. So, $\lambda_1 v^T u = (A^T v)^T u = v^T(Au) = v^T \lambda_2 u = \lambda_2 v^T u$. Thus, $v^T u = 0$. \square

B EQUATIONS OF THE JACOBIAN ENTRIES

Recall the form of the Jacobian of our dynamical system in Equation (7).

We compute the form of each entry of J at the point (x, y) . Let $Q_x = \sum_{\ell} x_{\ell} e^{\eta e_{\ell}^T R f(x, y)}$, $Q_y = \sum_{\ell} y_{\ell} e^{-\eta e_{\ell}^T R^T h(x, y)}$, $S_x = \sum_{\ell} x_{\ell} e^{\xi e_{\ell}^T R y}$, and $S_y = \sum_{\ell} y_{\ell} e^{-\xi e_{\ell}^T R^T x}$.

$$\begin{aligned}
\frac{\partial \varphi_{1,i}}{\partial x_i} &= e^{\eta e_i^T R f(x, y)} \frac{Q_x \left(1 + \eta x_i \frac{\partial}{\partial x_i} (e_i^T R f(x, y)) \right) - x_i \frac{\partial}{\partial x_i} Q_x}{Q_x^2}, \quad i \in [n], \\
\frac{\partial \varphi_{1,i}}{\partial x_j} &= x_i e^{\eta e_i^T R f(x, y)} \frac{\eta Q_x \frac{\partial}{\partial x_j} (e_i^T R f(x, y)) - \frac{\partial}{\partial x_j} Q_x}{Q_x^2}, \quad i, j \in [n] \text{ and } i \neq j, \\
\frac{\partial \varphi_{1,i}}{\partial y_j} &= x_i e^{\eta e_i^T R f(x, y)} \frac{\eta Q_x \frac{\partial}{\partial y_j} (e_i^T R f(x, y)) - \frac{\partial}{\partial y_j} Q_x}{Q_x^2}, \quad i, j \in [n], \\
\frac{\partial \varphi_{2,i}}{\partial x_j} &= y_i e^{-\eta e_i^T R^T h(x, y)} \frac{-\eta Q_y \frac{\partial}{\partial x_j} (e_i^T R^T h(x, y)) - \frac{\partial}{\partial x_j} Q_y}{Q_y^2}, \quad i, j \in [n], \\
\frac{\partial \varphi_{2,i}}{\partial y_i} &= e^{-\eta e_i^T R^T h(x, y)} \frac{Q_y \left(1 - \eta y_i \frac{\partial}{\partial y_i} (e_i^T R^T h(x, y)) \right) - y_i \frac{\partial}{\partial y_i} Q_y}{Q_y^2}, \quad i \in [n], \\
\frac{\partial \varphi_{2,i}}{\partial y_j} &= y_i e^{-\eta e_i^T R^T h(x, y)} \frac{-\eta Q_y \frac{\partial}{\partial y_j} (e_i^T R^T h(x, y)) - \frac{\partial}{\partial y_j} Q_y}{Q_y^2}, \quad i, j \in [n] \text{ and } i \neq j.
\end{aligned} \tag{10}$$

At the point (x^*, y^*) , after exploiting the fact that this is an equilibrium profile, and simplifying some of the calculations, we obtain the following forms.

$$\begin{aligned}
\frac{\partial \varphi_{1,i}}{\partial x_i} &= 1 - x_i^* \left(\eta \xi \left(\sum_k R_{ik}^2 y_k^* - \sum_k x_k^* \sum_l R_{kl} y_l^* R_{li}^T \right) + 1 \right), \quad i \in \text{supp}(x^*), \\
\frac{\partial \varphi_{1,i}}{\partial x_i} &= \frac{e^{\eta e_i^T R y^*}}{e^{\eta v}}, \quad i \notin \text{supp}(x^*), \\
\frac{\partial \varphi_{1,i}}{\partial x_j} &= -x_i^* \left(\eta \xi \left(\sum_k R_{ik} y_k^* R_{kj}^T - \sum_k x_k^* \sum_l R_{kl} y_l^* R_{lj}^T \right) + 1 \right), \quad i \in \text{supp}(x^*), i \neq j, \\
\frac{\partial \varphi_{1,i}}{\partial x_j} &= 0, \quad i \notin \text{supp}(x^*) \text{ and } i \neq j, \\
\frac{\partial \varphi_{1,i}}{\partial y_j} &= x_i^* \eta (R_{ij} - e_j^T R^T x^*) \frac{e^{-\xi e_j^T R^T x^*}}{e^{-\xi v}}, \quad \text{for all } i \in \text{supp}(x^*), \\
\frac{\partial \varphi_{1,i}}{\partial y_j} &= 0, \quad i \notin \text{supp}(x^*), \\
\frac{\partial \varphi_{2,i}}{\partial x_j} &= -y_i^* \eta (R_{ij}^T - e_j^T R y^*) \frac{e^{\xi e_j^T R y^*}}{e^{\xi v}}, \quad \text{for all } i \in \text{supp}(y^*), \\
\frac{\partial \varphi_{2,i}}{\partial x_j} &= 0, \quad i \notin \text{supp}(y^*), \\
\frac{\partial \varphi_{2,i}}{\partial y_i} &= 1 - y_i^* \left(\eta \xi \left(\sum_k (R_{ik}^T)^2 x_k^* - \sum_k y_k^* \sum_l R_{ki}^T x_l^* R_{li} \right) + 1 \right), \quad i \in \text{supp}(y^*), \\
\frac{\partial \varphi_{2,i}}{\partial y_i} &= \frac{e^{-\eta e_i^T R^T x^*}}{e^{-\eta v}}, \quad i \notin \text{supp}(y^*), \\
\frac{\partial \varphi_{2,i}}{\partial y_j} &= -y_i^* \left(\eta \xi \left(\sum_k R_{kj} R_{ik}^T x_k^* - \sum_k y_k^* \sum_l R_{kl}^T x_l^* R_{lj} \right) + 1 \right), \quad i \in \text{supp}(y^*), i \neq j, \\
\frac{\partial \varphi_{2,i}}{\partial y_j} &= 0, \quad i \notin \text{supp}(y^*) \text{ and } i \neq j,
\end{aligned} \tag{11}$$

C ADDITIONAL EXPERIMENTS

In this section, we demonstrate the properties of the FLBR-MWU algorithm using additional metrics and perform further comparisons.

Convergence to the value of the game. Figure 3 shows the evolution of the current value of the game at each iteration, with the same payoff matrix as that used in the example of Figure 1, in the main body of the

paper. The current value of the game at iteration t is defined as $v^t = x^{tT} R y^t$, and it serves as another convergence measure to Nash equilibrium. MWU (blue) oscillates around the true value of the game ($v = 0.529677$) without converging, while OMWU (red) oscillates with decreasing amplitude and eventually it converges to the true value. The current game value for the FLBR-MWU dynamics (black) converges much faster requiring only a few thousand steps.




Figure 3: The value of the game as a function of the number of steps for the three MWU variants.

Dynamics trajectories. Figures 4 and 5 show the trajectories of the row player (i.e., x_i^t for $i = 1, \dots, 10$) in linear and log scale, respectively. Similarly, Figures 6 and 7 show the trajectories of the column player (i.e., y_i^t). Again, the payoff matrix is the same as in Figure 1 from the main paper, and the Nash equilibrium is estimated as:

$$\begin{pmatrix} x^{*T} \\ y^{*T} \end{pmatrix} = \begin{pmatrix} 0.126766, 0.276988, 0, 0.22506, 0.081435, 0, 0.191705, 0, 0.098045, 0 \\ 0, 0.058227, 0, 0.298188, 0.213176, 0, 0, 0.283403, 0.000376, 0.146628 \end{pmatrix}.$$

First, we note that for all pure strategies that do not belong to the support of x^* or y^* , the corresponding probabilities in x^t and y^t converge to 0 under FLBR-MWU, after a few thousand steps. Additionally, we observe interesting patterns during the evolution of the learning dynamics in both scales which are intimately connected with the KL divergence trajectory shown in Figure 1 of the main paper. Indeed, it is worth looking at y_9^t (log scale; Figure 7), which shows the most interesting pattern. Initially it seems that this is not a surviving strategy of the dynamics and its probability decreases for the first 10K steps. However, and, despite its very low value, it recovers to the actual Nash equilibrium value. Similarly, we observe that the non-zero elements of x^t (linear scale; Figure 4) are linearly evolving for several thousands of steps. Those changes in the dynamics correspond to the plateau of the KL divergence observed in Figure 1 of the main text. Our explanation of the dynamics trajectories is as follows: starting from the uniform state, the FLBR-MWU algorithm first finds an approximate Nash equilibrium with a value close to the true value of the game but then escapes from it until it eventually converges to the actual Nash equilibrium.

Another interesting observation is that the dynamics of the IBR step (recall Equation (3) in the main paper) drive the FLBR-MWU dynamics in the sense that when the IBR dynamics are above the FLBR-MWU dynamics, then the corresponding probabilities in the update step of FLBR-MWU increase, while the opposite is true when the IBR dynamics are below the FLBR-MWU dynamics.

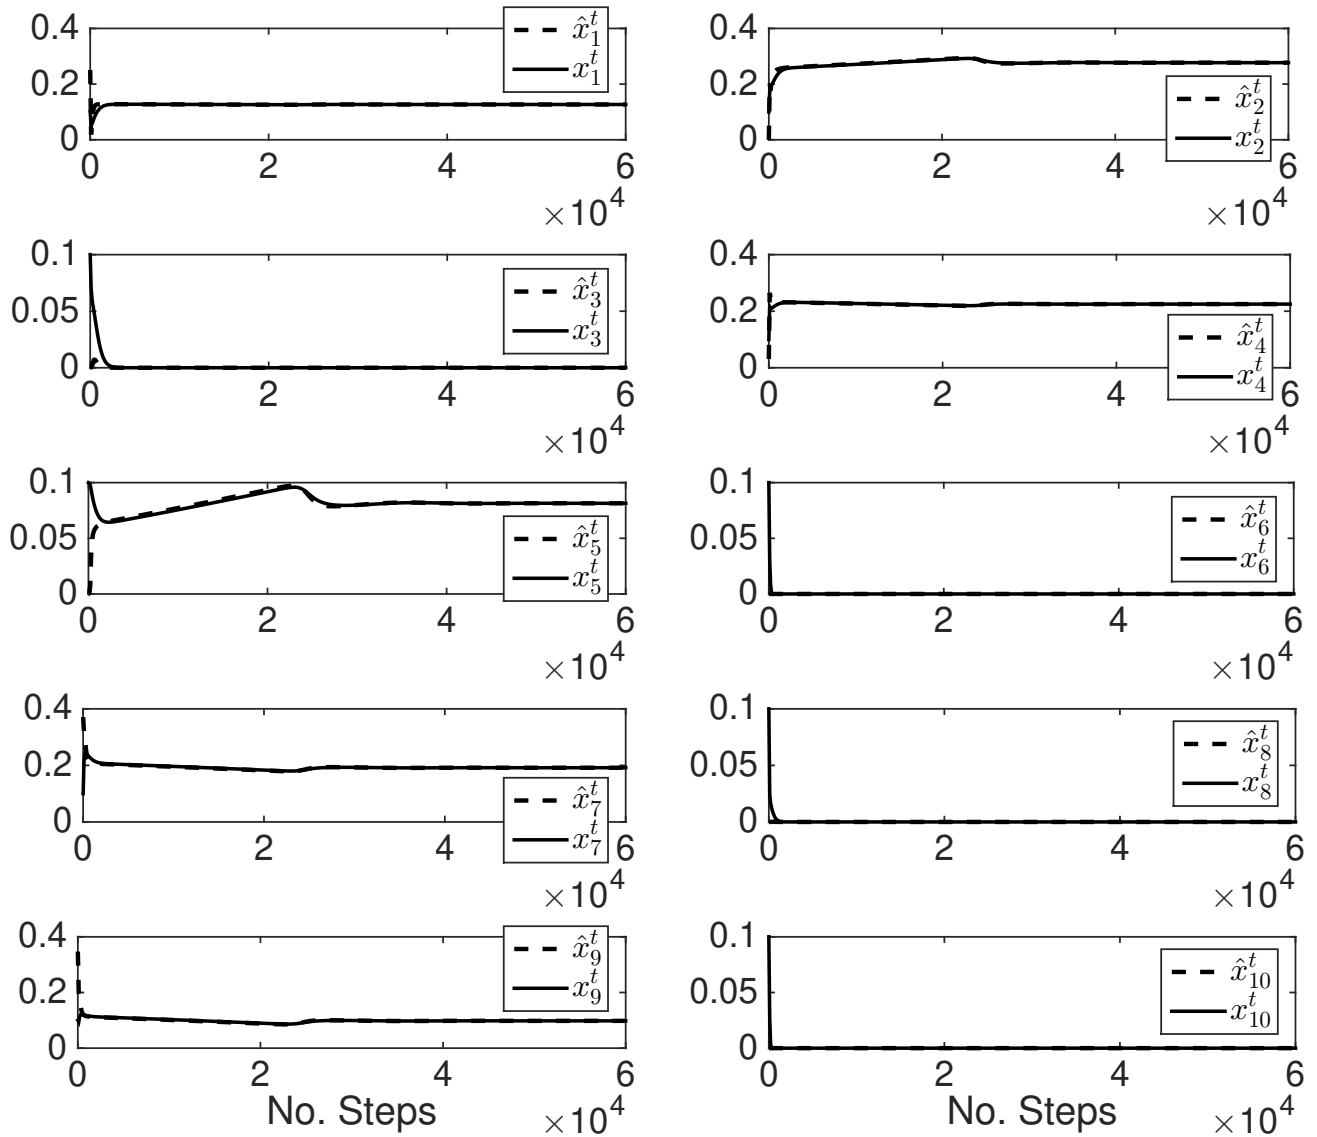


Figure 4: The dynamics of the update step per coordinate (solid), as well as the IBR step (dashed) for the row player. For the equilibrium strategy x^* , it holds that $\text{supp}(x^*) = \{1, 2, 4, 5, 7, 9\}$. Note that x^t converges to the same support.

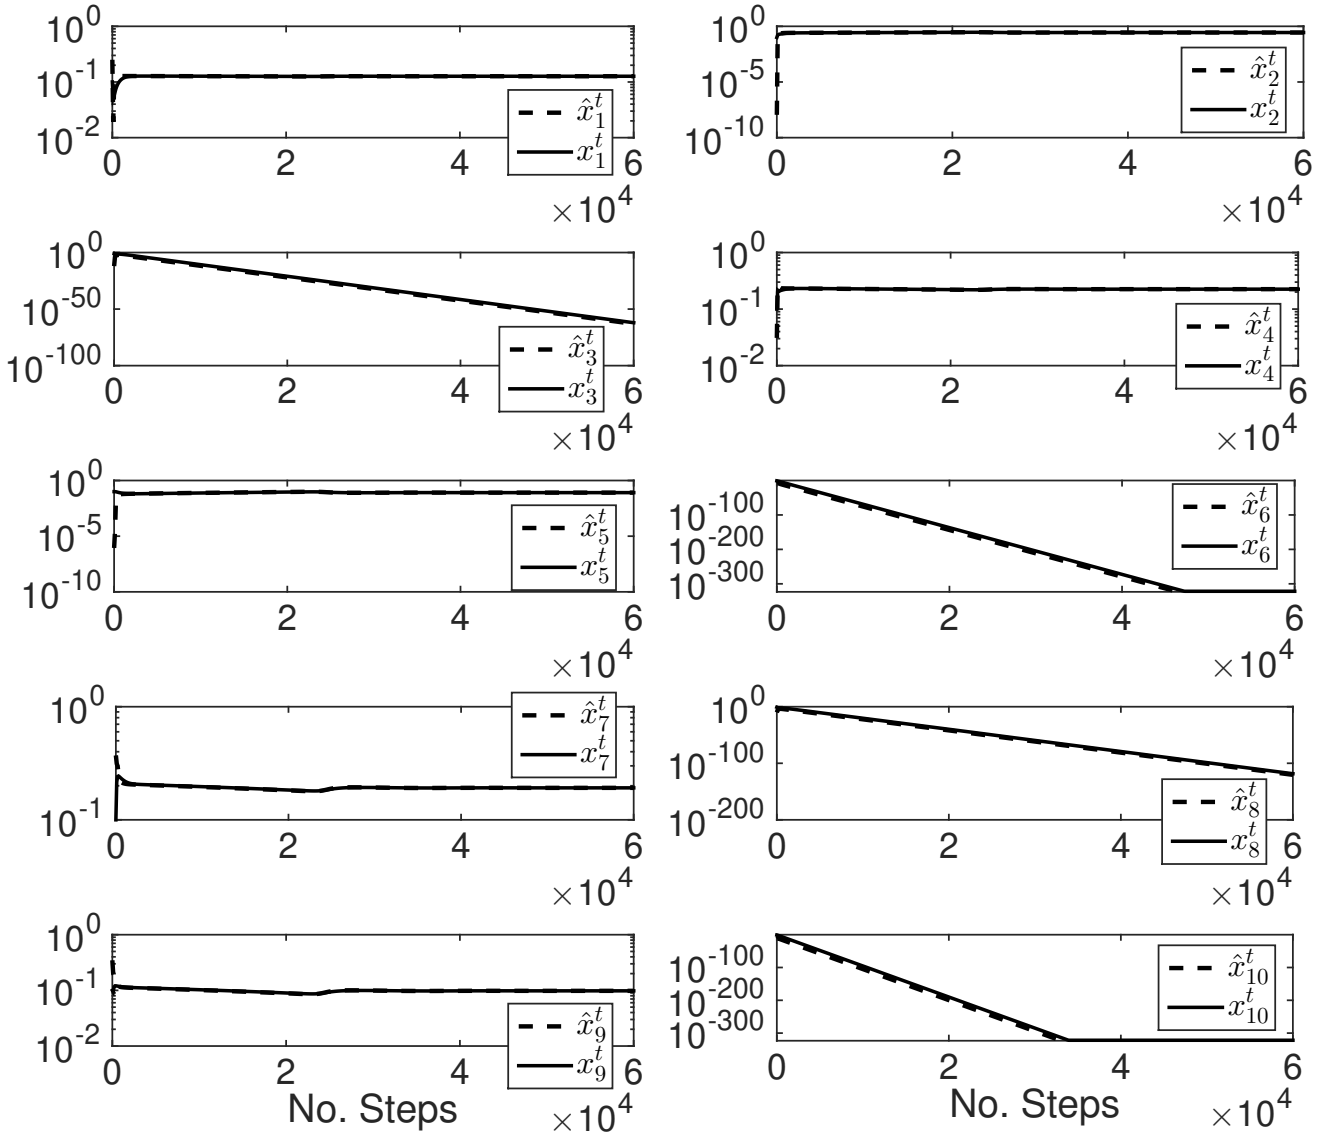


Figure 5: Same as Figure 4, but in logarithmic scale.

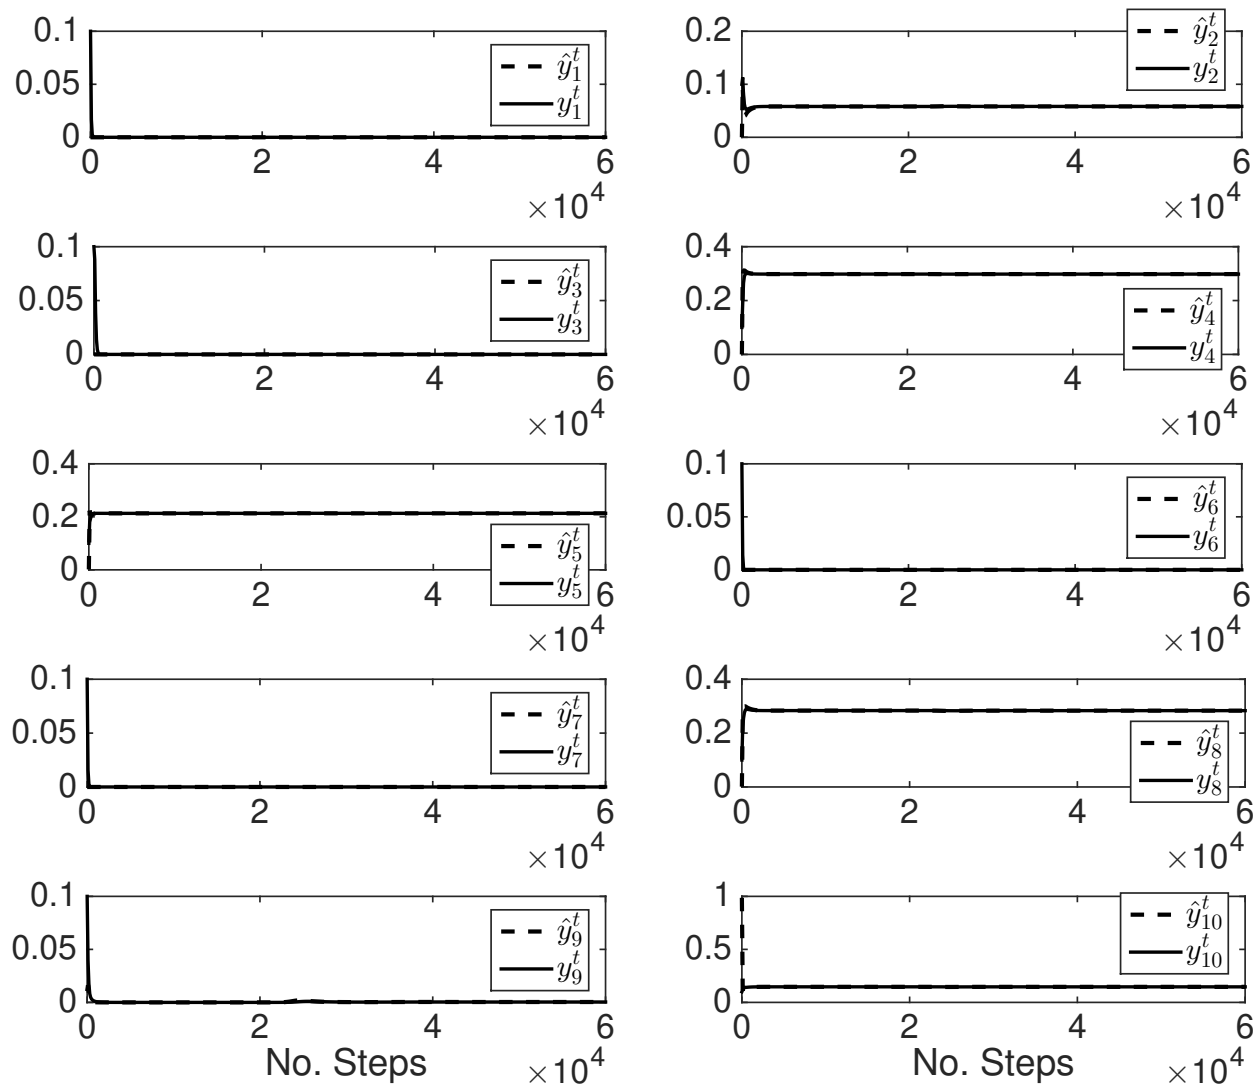


Figure 6: The dynamics of the update step per coordinate (solid), as well as the IBR step (dashed) for the column player. For the equilibrium strategy y^* , it holds that $\text{supp}(y^*) = \{2, 4, 5, 8, 9, 10\}$.

Effect of the intermediate rate (ξ). We present further statistical information on the effect of ξ . Figure 8 shows the distribution of the number of steps as a boxplot for $n = 10$ (left) and $n = 20$ (right). The red line in the boxplot corresponds to the median value while the blue box corresponds to the area covered by the 2nd and 3rd quantiles. The distribution of the number of steps till convergence is positively (or right) skewed. Therefore we also report the statistics of the right tail in Table 3. The presented results further validate the suggested value for ξ in Section 4 of the main paper. We also remark that the product $\eta\xi$ is not always less than 1 in our experiments. Hence, although we needed the condition $\eta\xi < 1$ to prove our theoretical result in Section 3, the numerical evidence shows that the product can take values greater than 1 and still attain convergence (however $\eta\xi$ should not become arbitrarily large).

Number of steps. Moving on, we present additional comparisons between FLBR-MWU and OMWU. Figure 9 demonstrates the distribution of the number of steps till convergence for FLBR-MWU (left) and OMWU (right). Interestingly, the distribution for payoff matrix size $n = 50$ with the FLBR-MWU algorithm is similar to the

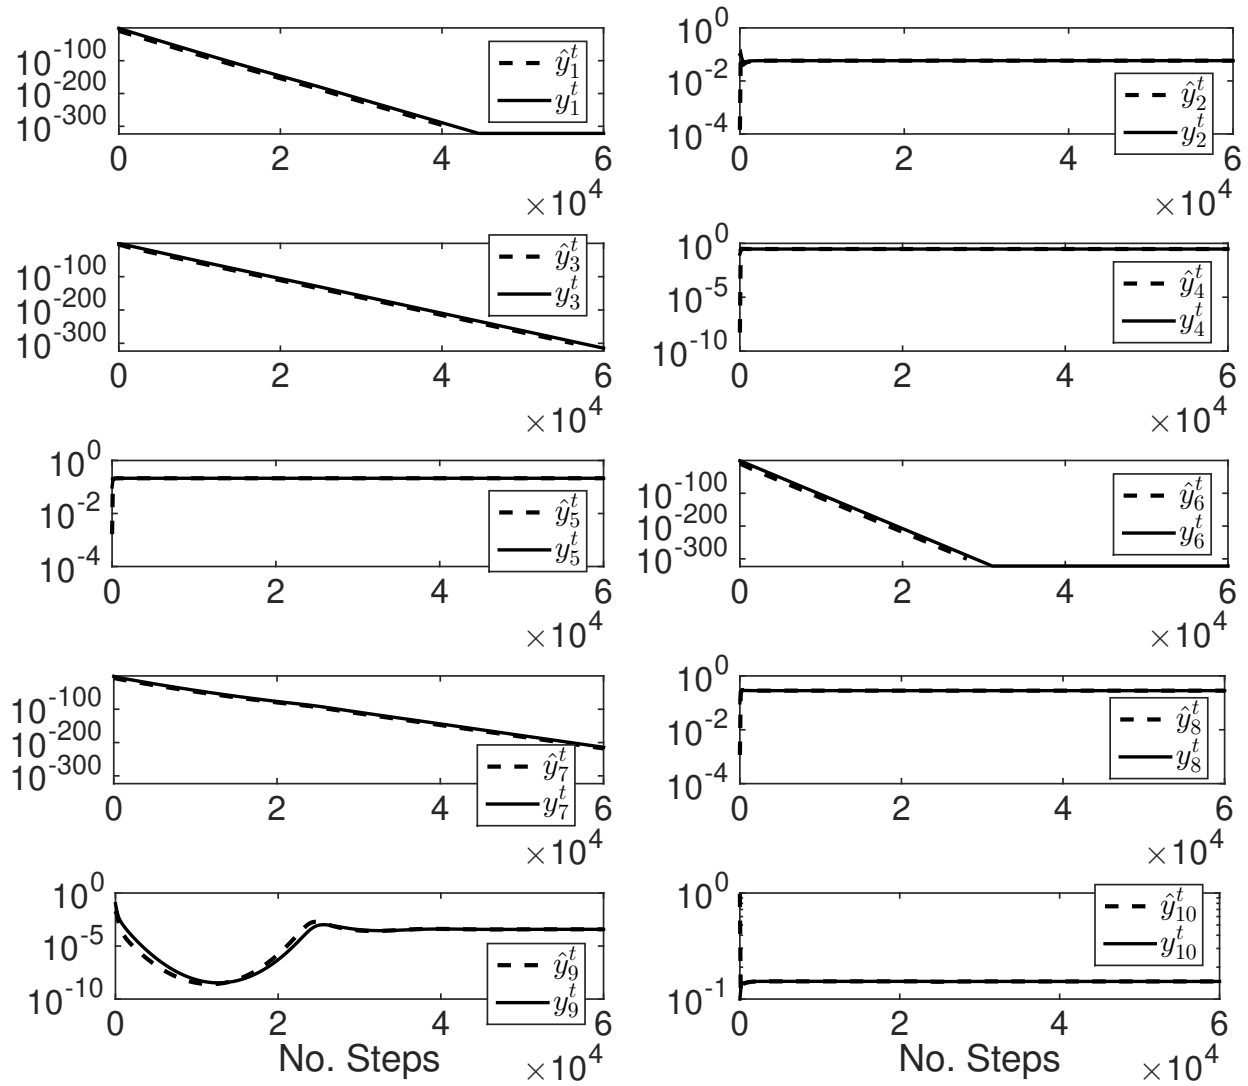



Figure 7: Same as Figure 6, but in logarithmic scale.

Figure 8: Boxplots for the number of steps until convergence for various values of ξ and two payoff matrix sizes.Table 3: Quantile statistics on the number of steps till convergence for various values of ξ and n . The maximum number of steps was set to $t_{\max} = 2 \times 10^6$.

Matrix size	Quantile	$\xi = 10$	$\xi = 20$	$\xi = 50$	$\xi = 100$	$\xi = 200$
$n = 10$	75%	127.1K	83.3K	54.5K	44.1K	43.6K
	90%	346.7K	209.1K	137.9K	111.3K	110.3K
	97.5%	1035.2K	640.3K	322.1K	228.7K	372.8K
$n = 20$	75%	1957.2K	1127.2K	576.9K	441.3K	342.7K
	90%	2000.0K	2000.0K	1644.8K	1076.0K	830.8K
	97.5%	2000.0K	2000.0K	2000.0K	2000.0K	2000.0K

distribution for $n = 5$ with the OMWU algorithm. The computational gains are expected to be even more dramatic for larger games.




Figure 9: Boxplots for the number of steps until convergence for various payoff matrix sizes under FLBR-MWU (left) and OMWU/OMD (right). The computational gains when FLBR-MWU is used are striking.

Comparisons between MWU, OMWU, and OMD. We also present a comparison among the MWU, OMWU and OMD dynamics (where for OMD we implemented the version of (Mertikopoulos et al., 2019) with entropy regularization). Figure 10 shows the evolution of a long run of 5 million steps and two values for the learning rate, η . We use the same payoff matrix as in Figure 1 of the main text and recall that the proposed FLBR-MWU method converged after only 100K steps (see Figure 1 in the main text). It is evident from the KL divergence in Figure 10 (leftmost panels) that the OMWU and OMD algorithms have almost the same behavior, as expected by (Wei et al., 2021), and they both converge, but in a very slow pace. The oscillatory behavior is prominent even after a large number of steps, as quantified by the l_1 norm difference (rightmost panels of Figure 10).




Figure 10: KL divergence and l_1 norm difference for $t_{\max} = 5 \times 10^6$ and two values for the learning rate: $\eta = 0.1$ (upper row of panels) and $\eta = 0.02$ (lower row of panels).

Finally, we report in Table 4 several convergence statistics between OWMU, OMD and FLBR-MWU algorithms with $\eta = 0.1$. This table is an extension of Table 2 from the main text (in Section 4). Once again, the proposed FLBR-MWU algorithm is orders of magnitude faster while the closeness of the statistics between OWMU and OMD reveals the (almost) equivalence between the two algorithms.

Table 4: Statistics on the number of steps till convergence for OWMU, OMD and FLBR-MWU and various payoff matrix sizes. The maximum number of steps was set to $t_{\max} = 5 \times 10^6$.

Matrix size	Statistic	$n = 5$	$n = 10$	$n = 20$	$n = 50$
OWMU	Mean	1287.3K	3280.9K	4997.8K	5000.0K
	Median	631.9K	3697.8K	5000.0K	5000.0K
	t_{\max}	12.0	44.0	98.0	100.0
OMD	Mean	1287.6K	3292.9K	4997.8K	5000.0K
	Median	631.9K	3629.1K	5000.0K	5000.0K
	t_{\max}	12.0	44.0	98.0	100.0
FLBR-MWU	Mean	18.8K	45.9K	267.1K	1130.8K
	Median	8.0K	21.4K	64.0K	701.3K
	t_{\max}	0.0	0.0	0.0	2.0