

---

# Countering Feedback Delays in Multi-Agent Learning

---

**Zhengyuan Zhou**  
Stanford University  
zyzhou@stanford.edu

**Panayotis Mertikopoulos**  
Univ. Grenoble Alpes, CNRS, Inria, LIG  
panayotis.mertikopoulos@imag.fr

**Nicholas Bambos**  
Stanford University  
bambos@stanford.edu

**Peter Glynn**  
Stanford University  
glynn@stanford.edu

**Claire Tomlin**  
UC Berkeley  
tomlin@eecs.berkeley.edu

## Abstract

In this paper, we consider a model of game-theoretic learning based on online mirror descent (OMD) with asynchronous and delayed information. Instead of focusing on a specific class of games (such as zero-sum or potential games), we introduce a general equilibrium stability notion for games with continuous action spaces, which we call *variational stability*. Our first contribution is that the “last iterate” (that is, the induced sequence of play) of OMD converges to variationally stable equilibria provided that the feedback delays faced by the players are synchronous and bounded. Subsequently, to tackle fully decentralized, asynchronous environments with unbounded feedback delays, we propose a variant of OMD which we call delayed mirror descent (DMD), and which relies on the repeated leveraging of past information. With this modification, the algorithm converges to variationally stable equilibria, with no feedback synchronicity assumptions, and even when the delays grow superlinearly relative to the game’s horizon.

## 1 Introduction

Online learning is a broad and powerful theoretical framework with widespread applications and great successes in machine learning, data science, networks, computer science, and many other fields [2, 5, 18]. The prototypical online learning problem may be described as follows: At each round  $t = 0, 1, \dots$ , a player selects an action  $x^t$  from some convex, compact set, and obtains a reward  $u^t(x^t)$  based on some *a priori* unknown payoff function  $u^t$ . Subsequently, the player receives some feedback (e.g. the past history of the reward functions or some restricted information thereof), and selects a new action  $x^{t+1}$  with the goal of maximizing the obtained reward. Aggregating over the rounds of the process, this is usually quantified by asking that the player’s (external) *regret*  $\text{Reg}(T) \equiv \max_{x \in \mathcal{X}} \sum_{t=1}^T [u^t(x) - u^t(x^t)]$  grow sublinearly with  $t$ , a property known as “no regret”.

One of the most widely used algorithmic schemes for learning in this context is the class of online mirror descent (OMD) algorithms [19]. Tracing its origins to [12] for offline optimization problems, OMD proceeds by taking a gradient step in the dual (gradient) space and projecting it back to the primal (decision) space via a *mirror map* generated by a strongly convex regularizer function (with different regularizers giving rise to different algorithms). In particular, OMD includes as special cases several seminal learning algorithms, such as Zinkevich’s online gradient descent (OGD) scheme [22], and the exponential weights (EW) algorithm [10]. Several variants of this class also exist and, perhaps unsurprisingly, they occur with a variety of different names – such as “Follow-the-Regularized-Leader” [7], dual averaging [13, 21], and so on.

When  $u^t$  is concave, OMD enjoys a sublinear  $\mathcal{O}(\sqrt{T})$  regret bound which is known to be universally tight.<sup>1</sup> A common instantiation of this is found in repeated multi-player games, where each player’s payoff function is determined by the actions of all the players via a fixed mechanism – the *stage game*. Even though this mechanism may be unknown to the players, this regret bound raises high expectations in terms of performance guarantees, so it is natural to assume that players adopt some variant of online mirror descent when faced with such online decision processes. This leads to the following central question: *if all players of a repeated game employ an OMD updating rule that leads to no regret, do their actions converge to a Nash equilibrium of the underlying one-shot game?*

Given the prominence of Nash equilibrium as a solution concept (compared to coarser notions such as correlated equilibria or the Hannan set), this problem lies at the heart of game-theoretic learning [3]. However, convergence to a Nash equilibrium is, in the words of [3], “considerably more difficult” than attaining a no-regret state for all players (which leads to other weaker equilibria mentioned above). A growing body of literature has focused on special classes of games (e.g. zero-sum games, routing games) and established the convergence of the “ergodic average”  $\frac{1}{T} \sum_{t=1}^T x^t$  of OMD [1, 8, 9]. In general however, the process may fail to converge altogether, even in simple, finite games [20].

A further major challenge arises in decentralized environments (such as transportation networks), where a considerable delay often occurs between a player’s action and the corresponding received feedback. To study learning in such settings, [14] recently introduced an elegant and general delay framework, where the gradient at round  $t$  is only available at round  $t + d^t - 1$ , with  $d^t$  being the delay associated with the player’s action at round  $t$ .<sup>2</sup> [14] then considered a very natural extension of OMD under delays: updating the set of gradients as they are received (see Algorithm 1 for details). If the total delay after time  $T$  is  $D(T) = \sum_{t=1}^T d^t$ , [14] showed that OMD enjoys an  $\mathcal{O}(D(T)^{1/2})$  regret bound. This natural extension has several strengths: no assumption is made on how the gradients are received (the delayed gradients can be received out-of-order); further, as pointed out in [4, 6], a gradient “does not need to be timestamped by the round  $s$  from which it originates, which is required by the pooling strategies of [6] and [4] in order to return the feedback to the appropriate learner”.

## 1.1 Our contributions

Our investigations here differ from existing work in the following aspects: First, we consider learning in games with asynchronous and delayed feedback following the general framework of [14]. Previous work on the topic has focused on the regret analysis of single-agent learning with delays, but the convergence properties of such processes in continuous games are completely unknown. Second, we focus throughout on the convergence of the sequence of play generated by OMD (its “last iterate”), as opposed to the algorithm’s ergodic average  $\frac{1}{T} \sum_{t=1}^T x^t$ . This last point is worth emphasizing for several reasons: *a)* the convergence of the last iterate is stronger and theoretically more appealing because it implies ergodic convergence; *b)* in a game-theoretic setting, payoffs are determined by the actual sequence of play, so ergodic convergence diminishes in value if it is not accompanied by similar conclusions for the players’ realized actions; and *c)* because there is no inherent averaging, the techniques used to prove convergence of  $x^t$  provide a much finer understanding of the evolution of OMD.

A further contribution of our paper is the introduction of an equilibrium stability notion which we call *variational stability*, and which is formally similar to the influential notion of *evolutionary stability* introduced in evolutionary game theory by Maynard Smith and Price [11]. By means of this notion, we are able to treat convergence questions in general games with continuous action spaces, without having to focus on a specific class of games, such as concave potential or strictly monotone games (though our analysis also covers such games). This stability notion is deeply related to monotone operators in variational analysis [15] and can be seen as a strict generalization of operator monotonicity in the current game-theoretic context.<sup>3</sup>

Our first result is that, assuming variational stability, the sequence of play induced by OMD converges to the game’s set of Nash equilibria, provided that the delays of all players are synchronous and bounded (see [Theorems 4.3](#) and [4.4](#)). As an inherited benefit, players adopting this learning algorithm

<sup>1</sup>In many formulations, a cost function (as opposed to a reward function) is used, in which case such cost functions need to be convex.

<sup>2</sup>Of course, taking  $d^t = 1$  yields the classical no-delay setting.

<sup>3</sup>In the appendix, we give two well-known classes of games that satisfy this equilibrium notion.

can receive gradients out-of-order and do not need to keep track of the timestamps from which the gradients originate. In fact, even in the special case of learning *without* delays, we are not aware of a similar convergence result for the actual sequence of play.

An important limitation of this result is that delays are assumed synchronous and bounded, an assumption which might not hold in large, decentralized environments. To lift this barrier, we introduce a modification of vanilla OMD which we call *delayed mirror descent* (DMD), and which leverages past information repeatedly, even in rounds where players receive no feedback. Thanks to this modification, play under DMD converges to variationally stable sets of Nash equilibria (Theorems 5.2 and 5.3), even if the players experience asynchronous and unbounded delays: in particular, delays could grow superlinearly in the game’s horizon, and DMD would still converge.

We mention that the convergence proofs for both OMD and DMD rely on designing a particular Lyapunov function, the so-called  $\lambda$ -Fenchel coupling which serves as a “primal-dual Bregman divergence” measure between actions and gradient variables. Thanks to its Lyapunov properties, the  $\lambda$ -Fenchel coupling provides a potent tool for proving convergence and we exploit it throughout. Further, we present a unified theoretical framework that puts the analysis of both algorithms under different delay assumptions on the same footing.

## 2 Problem setup

We introduce and formalize the framework of mirror descent learning on games with continuous actions under delays, where each player applies the low-regret learning algorithm online mirror descent (OMD) in a repeated game setting with feedback delays.

### 2.1 Games with Continuous Actions

We start with the definition of a game with continuous actions, which serves as a stage game and provides a reward function for each player in an online learning process.

**Definition 2.1.** A game with continuous actions  $\mathcal{G}$  is a tuple  $\mathcal{G} = (\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N)$ , where  $\mathcal{N}$  is the set of  $N$  players  $\{1, 2, \dots, N\}$ ,  $\mathcal{X}$  is the joint action space with  $\mathcal{X}_i$  a compact and convex subset of a finite-dimensional vector space  $\mathbb{R}^{d_i}$ , being the action space for player  $i$  and  $u_i : \mathcal{X} \rightarrow \mathbf{R}$  is the utility function for player  $i$ , such that the following hold:

1. For each  $i \in \mathcal{N}$ ,  $u_i(\mathbf{x})$  is continuous in  $\mathbf{x}$ .
2. For each  $i \in \mathcal{N}$ ,  $u_i$  is continuously differentiable in  $x_i$  and the partial gradient  $\nabla_{x_i} u_i(\mathbf{x})$  is Lipschitz continuous in  $\mathbf{x}$ .

Throughout the paper,  $\mathbf{x}_{-i}$  denotes the joint action of all players but player  $i$ . Consequently, the joint action<sup>4</sup>  $\mathbf{x}$  will frequently be written as  $(x_i, \mathbf{x}_{-i})$ . Two important quantities in the current context are:

**Definition 2.2.** We denote by  $\mathbf{v}(\mathbf{x})$  to be the column vector of partial gradients of the utility functions<sup>5</sup>:  $\mathbf{v}(\mathbf{x}) = (v_1(\mathbf{x}), \dots, v_N(\mathbf{x}))$ , where  $v_i(\mathbf{x}) \triangleq \nabla_{x_i} u_i(\mathbf{x})$ .

**Definition 2.3.** Given a game with continuous actions  $(\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N)$ ,  $\mathbf{x}^* \in \mathcal{X}$  is called a (pure-strategy) Nash equilibrium if for each  $i \in \mathcal{N}$ ,  $u_i(x_i^*, \mathbf{x}_{-i}^*) \geq u_i(x_i, \mathbf{x}_{-i}^*), \forall x_i \in \mathcal{X}_i$ .

### 2.2 Online Mirror Descent on Games under Delays

[14] provides a very general model of delays (also adopted here) and a natural extension of OMD under delays, which when applied to a repeated game, yields the following learning scheme:

Three comments are in order here. First, each  $h_i(\cdot)$  is a regularizer on  $\mathcal{X}_i$ , as defined next.

**Definition 2.4.** Let  $\mathcal{D}$  be a compact and convex subset of  $\mathbf{R}^m$ . We say that  $g : \mathcal{D} \rightarrow \mathbf{R}$  is a regularizer (with respect to some vector norm  $\|\cdot\|$ ) if  $g$  is continuous and  $g$  is  $K$ -strongly convex

<sup>4</sup>Note that boldfaced letters are only used to denote joint actions. In particular,  $x_i$  is a vector even though it is not boldfaced.

<sup>5</sup>Note that per the last assumption in the definition of a concave game (Definition 2.1), the gradient  $\mathbf{v}(\mathbf{x})$  always exists and is a continuous function on the joint action space  $\mathcal{X}$ .

---

**Algorithm 1** Online Mirror Descent on Games under Delays
 

---

```

1: Each player  $i$  chooses an initial  $y_i^0$ .
2: for  $t = 0, 1, 2, \dots$  do
3:   for  $i = 1, \dots, N$  do
4:      $x_i^t = \arg \max_{x_i \in \mathcal{X}_i} \{ \langle y_i^t, x_i \rangle - h_i(x_i) \}$ 
5:      $y_i^{t+1} = y_i^t + \alpha^t \sum_{s \in \mathcal{G}_i^t} v_i(\mathbf{x}^s)$ 
6:   end for
7: end for

```

---

with respect to  $\|\cdot\|$ : there exists some  $K > 0$  such that  $\forall t \in [0, 1], \forall \mathbf{d}, \mathbf{d}' \in \mathcal{D}: g(t\mathbf{d} + (1-t)\mathbf{d}') \leq tg(\mathbf{d}) + (1-t)g(\mathbf{d}') - \frac{1}{2}Kt(1-t)\|\mathbf{d}' - \mathbf{d}\|^2$ .

Second, the gradient step size  $\alpha^t$  in Algorithm 1 can be any positive and non-increasing sequence that satisfies the standard summability assumption:  $\sum_{t=0}^{\infty} \alpha^t = \infty, \sum_{t=0}^{\infty} (\alpha^t)^2 < \infty$ .

Thirdly, in regards to the delay model: in Algorithm 1,  $\mathcal{G}_i^t$  denotes the set of rounds whose gradients become available for player  $i$  at the current round  $t$ . Denote player  $i$ 's delay of the gradient at round  $s$  to be  $d_i^s$  (a positive integer), then this gradient  $v_i(\mathbf{x}^s)$  will be available at round  $s + d_i^s - 1$ :  $s \in \mathcal{G}_i^{s+d_i^s-1}$ . In particular, if  $d_i^s = 1, \forall s$ , then this corresponds to the case where player  $i$  doesn't have any delays. Note that each player can receive out-of-order gradients: this can happen if the gradient at an earlier round has a much larger delay than that of the gradient at a later round.

### 3 Variational Stability: A Key Criterion

In this section, we define a key equilibrium notion, called  $\lambda$ -variational stability and study the structural properties of a  $\lambda$ -variationally stable set. This equilibrium notion allows us to obtain general last-iterate convergence results, as opposed to results that only hold in specific classes of games. We also give in the appendix two classes of games (convex potential games and linear Cournot games) that satisfy the variationally stable criterion.

#### 3.1 Variational Stability

**Definition 3.1.** Given a game with continuous actions ( $\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N$ ), a set  $\mathcal{C} \subset \mathcal{X}$  is called:  $\lambda$ -variationally stable for some  $\lambda \in \mathbf{R}_{++}^N$  if it holds that (with equality if and only if  $\mathbf{x} \in \mathcal{C}$ ):

$$\sum_{i=1}^N \lambda_i \langle v_i(\mathbf{x}), x_i - x_i^* \rangle \leq 0, \forall \mathbf{x} \in \mathcal{X}, \forall \mathbf{x}^* \in \mathcal{C}.$$

We emphasize that in a game setting,  $\lambda$ -variational stability is more general than an important concept called operator monotonicity in variational analysis. Specifically,  $v(\cdot)$  is called a monotone operator [16] if the following holds (with equality if and only if  $\mathbf{x} = \tilde{\mathbf{x}}$ ):

$$\langle v(\mathbf{x}) - v(\tilde{\mathbf{x}}), \mathbf{x} - \tilde{\mathbf{x}} \rangle \triangleq \sum_{i=1}^N \langle v_i(\mathbf{x}) - v_i(\tilde{\mathbf{x}}), x_i - \tilde{x}_i \rangle \leq 0, \forall \mathbf{x}, \tilde{\mathbf{x}} \in \mathcal{X}. \quad (3.1)$$

It is known that ([17]) that if  $v(\cdot)$  is monotone, then a unique Nash equilibrium  $\mathbf{x}^*$  exists and (per the property of a Nash equilibrium) satisfies  $\langle v(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle \leq 0$ . Consequently, if  $v(\cdot)$  is a monotone operator, by expanding Equation 3.1, it then follows that  $\langle v(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle \leq \langle v(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle \leq 0$ , where equality is achieved if and only if  $\mathbf{x} = \mathbf{x}^*$ . This implies that when  $v(\mathbf{x})$  is a monotone operator, the singleton set of the unique Nash equilibrium is  $\mathbf{1}$ -variationally stable, where  $\mathbf{1}$  is the all-ones vector. The converse is not true: when  $v(\mathbf{x})$  is not a monotone operator, we can still have a unique Nash equilibrium that is  $\lambda$ -variationally stable, or more generally, have a  $\lambda$ -variationally stable set  $\mathcal{C}$ .

#### 3.2 Properties of Variational Stability

**Lemma 3.2.** Give a game with continuous actions ( $\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N$ ). If  $\mathcal{C}$  is a non-empty  $\lambda$ -variationally stable set, then  $\mathcal{C}$  is a closed and convex set of all Nash equilibria of the game.

The following lemma gives us a convenient sufficient condition ensuring that a singleton  $\lambda$ -variationally stable set  $\{\mathbf{x}^*\}$  exists; in this case, we simply say that  $\mathbf{x}^*$  is  $\lambda$ -variationally stable.

**Lemma 3.3.** *Given a game with continuous actions  $(\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N)$ , where each  $u_i$  is twice continuously differentiable. For each  $\mathbf{x} \in \mathcal{X}$ , define the  $\lambda$ -weighted Hessian matrix  $H^\lambda(\mathbf{x})$  as follows:*

$$H_{ij}^\lambda(\mathbf{x}) = \frac{1}{2}\lambda_i \nabla_{x_j} v_i(\mathbf{x}) + \frac{1}{2}\lambda_j (\nabla_{x_i} v_j(\mathbf{x}))^T. \quad (3.2)$$

If  $H^\lambda(\mathbf{x})$  is negative-definite for every  $\mathbf{x} \in \mathcal{X}$ , then the game admits a unique Nash equilibrium  $\mathbf{x}^*$  that is  $\lambda$ -globally variational stable.

*Remark 3.1.* It is important to note that the Hessian matrix so defined is a block matrix: each  $H_{ij}^\lambda(\mathbf{x})$  is a  $d_i \times d_j$  matrix. Writing it in terms of the utility function, we have  $H_{ij}^\lambda(\mathbf{x}) = \frac{1}{2}\lambda_i \nabla_{x_j} \nabla_{x_i} u_i(\mathbf{x}) + \frac{1}{2}\lambda_j (\nabla_{x_i} \nabla_{x_j} u_j(\mathbf{x}))^T$ .

## 4 Convergence to Nash equilibrium under synchronous and bounded delays

In this section, we tackle the convergence of the last iterate of OMD under delays. We start by defining an important divergence measure,  $\lambda$ -Fenchel coupling, that generalizes Bregman divergence. We then establish its useful properties that play an indispensable role in both this and next sections.

### 4.1 $\lambda$ -Fenchel Coupling

**Definition 4.1.** Fix a game with continuous action spaces  $(\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N)$  and for each player  $i$ , let  $h_i : \mathcal{X}_i \rightarrow \mathbf{R}$  be a regularizer with respect to the norm  $\|\cdot\|_i$  that is  $K_i$ -strongly convex.

1. The convex conjugate function  $h_i^* : \mathbb{R}^{d_i} \rightarrow \mathbf{R}$  of  $h_i$  is defined as:

$$h_i^*(y_i) = \max_{x_i \in \mathcal{X}_i} \{\langle x_i, y_i \rangle - h_i(x_i)\}.$$

2. The choice function  $C_i : \mathbb{R}^{d_i} \rightarrow \mathcal{X}_i$  associated with regularizer  $h_i$  for player  $i$  is defined as:

$$C_i(y_i) = \arg \max_{x_i \in \mathcal{X}_i} \{\langle x_i, y_i \rangle - h_i(x_i)\}.$$

3. For a  $\lambda \in \mathbb{R}_{++}^N$ , the  $\lambda$ -Fenchel coupling  $F^\lambda : \mathcal{X} \times \mathbb{R}^{\sum_{i=1}^N d_i} \rightarrow \mathbf{R}$  is defined as:

$$F^\lambda(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^N \lambda_i (h_i(x_i) - \langle x_i, y_i \rangle + h_i^*(y_i)).$$

Note that although the domain of  $h_i$  is  $\mathcal{X}_i \subset \mathbb{R}^{d_i}$ , the domain of its conjugate (gradient space)  $h_i^*$  is  $\mathbb{R}^{d_i}$ . The two key properties of  $\lambda$ -Fenchel coupling that will be important in establishing the convergence of OMD are given next.

**Lemma 4.2.** *For each  $i \in \{1, \dots, N\}$ , let  $h_i : \mathcal{X}_i \rightarrow \mathbf{R}$  be a regularizer with respect to the norm  $\|\cdot\|_i$  that is  $K_i$ -strongly convex and let  $\lambda \in \mathbb{R}_{++}^N$ . Then  $\forall \mathbf{x} \in \mathcal{X}, \forall \tilde{\mathbf{y}}, \mathbf{y} \in \mathbb{R}^{\sum_{i=1}^N d_i}$ :*

1.  $F^\lambda(\mathbf{x}, \mathbf{y}) \geq \frac{1}{2} \sum_{i=1}^N K_i \lambda_i \|C_i(y_i) - x_i\|_i^2 \geq \frac{1}{2} (\min_i K_i \lambda_i) \sum_{i=1}^N \|C_i(y_i) - x_i\|_i^2$ .
2.  $F^\lambda(\mathbf{x}, \tilde{\mathbf{y}}) \leq F^\lambda(\mathbf{x}, \mathbf{y}) + \sum_{i=1}^N \lambda_i \langle \tilde{y}_i - y_i, C_i(y_i) - x_i \rangle + \frac{1}{2} (\max_i \frac{\lambda_i}{K_i}) \sum_{i=1}^N (\|\tilde{y}_i - y_i\|_i^*)^2$ , where  $\|\cdot\|_i^*$  is the dual norm of  $\|\cdot\|_i$  (i.e.  $\|y_i\|_i^* = \max_{\|x_i\|_i \leq 1} \langle x_i, y_i \rangle$ ).

*Remark 4.1.* Collecting each individual choice map into a vector, we obtain the aggregate choice map  $C : \mathbb{R}^{\sum_{i=1}^N d_i} \rightarrow \mathcal{X}$ , with  $C(\mathbf{y}) = (C_1(y_1), \dots, C_N(y_N))$ . Since each space  $\mathcal{X}_i$  is endowed with norm  $\|\cdot\|_i$ , we can define the induced aggregate norm  $\|\cdot\|$  on the joint space  $\mathcal{X}$  as follows:  $\|\mathbf{x}\| = \sum_{i=1}^N \|x_i\|_i$ . We can also similarly define the aggregate dual norm:  $\|\mathbf{y}\|^* = \sum_{i=1}^N \|y_i\|_i^*$ . Henceforth, it shall be clear that the convergence in the joint space (e.g.  $C(\mathbf{y}^t) \rightarrow \mathbf{x}, \mathbf{y}^t \rightarrow \mathbf{y}$ ) will be defined under the respective aggregate norm. Finally, we assume throughout the paper that the

choice maps are regular in the following (very weak) sense: a choice map  $C(\cdot)$  is said to be  $\lambda$ -Fenchel coupling conforming if

$$C(\mathbf{y}^t) \rightarrow \mathbf{x} \text{ implies } F^\lambda(\mathbf{x}, \mathbf{y}^t) \rightarrow 0 \text{ as } t \rightarrow \infty. \quad (4.1)$$

Unless one aims for relatively pathological cases, choice maps induced by typical regularizers are always  $\lambda$ -Fenchel coupling conforming: examples include the Euclidean and entropic regularizers.

## 4.2 Convergence of OMD to Nash equilibrium

We start by characterizing the assumption on the delay model:

**Assumption 1.** The delays are assumed to be:

1. Synchronous:  $\mathcal{G}_i^t = \mathcal{G}_j^t, \forall i, j, \forall t$ .
2. Bounded:  $d_i^t \leq D, \forall i, \forall t$  (for some positive integer  $D$ ).

**Theorem 4.3.** Fix a game with continuous action spaces  $(\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N)$  that admits  $\mathbf{x}^*$  as the unique Nash equilibrium that is  $\lambda$ -variationally stable. Under Assumption 1, the OMD iterate  $\mathbf{x}^t$  given in Algorithm 1 converges to  $\mathbf{x}^*$ , irrespective of the initial point  $\mathbf{x}^0$ .

*Remark 4.2.* The proof is rather long and involved. To aid the understanding and enhance the intuition, we break it down into four main steps, each of which will be proved in the appendix in detail.

1. Since the delays are synchronous, we denote by  $\mathcal{G}^t$  the common set and  $d^t$  the common delay at round  $t$ . The gradient update in OMD under delays can then be written as:

$$y_i^{t+1} = y_i^t + \alpha^t \sum_{s \in \mathcal{G}^t} v_i(\mathbf{x}^s) = y_i^t + \alpha^t \left\{ |\mathcal{G}^t| v_i(\mathbf{x}^t) + \sum_{s \in \mathcal{G}^t} \{v_i(\mathbf{x}^s) - v_i(\mathbf{x}^t)\} \right\}. \quad (4.2)$$

Define  $b_i^t = \sum_{s \in \mathcal{G}^t} \{v_i(\mathbf{x}^s) - v_i(\mathbf{x}^t)\}$ . We show  $\lim_{t \rightarrow \infty} \|b_i^t\|_i^* = 0$  for each player  $i$ .

2. Define  $\mathbf{b}^t = (b_1^t, \dots, b_N^t)$  and we have  $\lim_{t \rightarrow \infty} \mathbf{b}^t = 0$  per Claim 1. Since each player's gradient update can be written as  $y_i^{t+1} = y_i^t + \alpha^t (|\mathcal{G}^t| v_i(\mathbf{x}^t) + b_i^t)$  per Claim 1, we can then write the joint OMD update (of all players) as:

$$\begin{aligned} \mathbf{x}^t &= C(\mathbf{y}^t), & (4.3) \\ \mathbf{y}^{t+1} &= \mathbf{y}^t + \alpha^t \{|\mathcal{G}^t| v(\mathbf{x}^t) + \mathbf{b}^t\}. & (4.4) \end{aligned}$$

Let  $B(\mathbf{x}^*, \epsilon) \triangleq \{\mathbf{x} \in \mathcal{X} \mid \|\mathbf{x} - \mathbf{x}^*\| < \epsilon\}$  be the open ball centered around  $\mathbf{x}^*$  with radius  $\epsilon$ . Then, using  $\lambda$ -Fenchel coupling as a "energy" function and leveraging the handle on  $\mathbf{b}^t$  given by Claim 1, we can establish that, for any  $\epsilon > 0$  the iterate  $\mathbf{x}^t$  will eventually enter  $B(\mathbf{x}^*, \epsilon)$  and visit  $B(\mathbf{x}^*, \epsilon)$  infinitely often, no matter what the initial point  $\mathbf{x}^0$  is. Mathematically, the claim is that  $\forall \epsilon > 0, \forall \mathbf{x}^0, |\{t \mid \mathbf{x}^t \in B(\mathbf{x}^*, \epsilon)\}| = \infty$ .

3. Fix any  $\delta > 0$  and consider the set  $\tilde{B}(\mathbf{x}^*, \delta) \triangleq \{C(\mathbf{y}) \mid F^\lambda(\mathbf{x}^*, \mathbf{y}) < \delta\}$ . In other words,  $\tilde{B}(\mathbf{x}^*, \delta)$  is some "neighborhood" of  $\mathbf{x}^*$ , which contains every  $\mathbf{x}$  that is an image of some  $\mathbf{y}$  (under the choice map  $C(\cdot)$ ) that is within  $\delta$  distance of  $\mathbf{x}^*$  under the  $\lambda$ -Fenchel coupling "metric". Although  $F^\lambda(\mathbf{x}^*, \mathbf{y})$  is not a metric,  $\tilde{B}(\mathbf{x}^*, \delta)$  contains an open ball within it. Mathematically, the claim is that for any  $\delta > 0, \exists \epsilon(\delta) > 0$  such that:  $B(\mathbf{x}^*, \epsilon) \subset \tilde{B}(\mathbf{x}^*, \delta)$ .
4. For any "neighborhood"  $\tilde{B}(\mathbf{x}^*, \delta)$ , after long enough rounds, if  $\mathbf{x}^t$  ever enters  $\tilde{B}(\mathbf{x}^*, \delta)$ , it will be trapped inside  $\tilde{B}(\mathbf{x}^*, \delta)$  thereafter. Mathematically, the claim is that for any  $\delta > 0, \exists T(\delta)$ , such that for any  $t \geq T(\delta)$ , if  $\mathbf{x}^t \in \tilde{B}(\mathbf{x}^*, \delta)$ , then  $\mathbf{x}^{\tilde{t}} \in \tilde{B}(\mathbf{x}^*, \delta), \forall \tilde{t} \geq t$ .

Putting all four elements above together, we note that the significance of Claim 3 is that, since the iterate  $\mathbf{x}^t$  will enter  $B(\mathbf{x}^*, \epsilon)$  infinitely often (per Claim 2),  $\mathbf{x}^t$  must enter  $\tilde{B}(\mathbf{x}^*, \delta)$  infinitely often. It therefore follows that, per Claim 4, starting from iteration  $t$ ,  $\mathbf{x}^t$  will remain in  $\tilde{B}(\mathbf{x}^*, \delta)$ . Since this is true for any  $\delta > 0$ , we have  $F^\lambda(\mathbf{x}^*, \mathbf{y}^t) \rightarrow 0$  as  $t \rightarrow \infty$ . Per Statement 1 in Lemma 4.2, this leads to that  $\|C(\mathbf{y}^t) - \mathbf{x}^*\| \rightarrow 0$  as  $t \rightarrow \infty$ , thereby establishing that  $\mathbf{x}^t = C(\mathbf{y}^t) \rightarrow \mathbf{x}^*$  as  $t \rightarrow \infty$ .

In fact, the result generalizes straightforwardly to multiple Nash equilibria. The proof of the convergence to the set case is line-by-line identical, provided we redefine, in a standard way, every quantity that measures the distance between two points to the corresponding quantity that measures the distance between a point and a set (by taking the infimum over the distances between the point and a point in that set). We directly state the result below.

**Theorem 4.4.** *Fix a game with continuous action spaces  $(\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N)$  that admits  $\mathcal{X}^*$  as a  $\lambda$ -variationally stable set (of necessarily all Nash equilibria), for some  $\lambda \in \mathbf{R}_{++}^N$ . Under Assumption 1, the OMD iterate  $\mathbf{x}^t$  given in Algorithm 1 satisfies  $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}^t, \mathcal{X}^*) = 0$ , irrespective of  $\mathbf{x}^0$ , where  $\text{dist}(\cdot, \cdot)$  is the standard point-to-set distance function induced by the norm  $\|\cdot\|$ .*

## 5 Delayed Mirror Descent: Asynchronous and unbounded delays

The synchronous and bounded delay assumption in Assumption 1 is fairly strong. In this section, by a simple modification of OMD, we propose a new learning algorithm called Delayed Mirror Descent (DMD), that allows the last-iterate convergence-to-Nash result to be generalized to cases with arbitrary asynchronous delays among players as well as unbounded delay growth.

### 5.1 Delayed Mirror Descent on Games

The main idea for the modification is that when player  $i$  doesn't receive any gradient on round  $t$ , instead of not doing any gradient updates as in OMD, he uses the most recent set of gradients to perform updates. More formally, define the most recent information set<sup>6</sup> as:

$$\tilde{\mathcal{G}}_i^t = \begin{cases} \mathcal{G}_i^t, & \text{if } \mathcal{G}_i^t \neq \emptyset \\ \tilde{\mathcal{G}}_i^{t-1}, & \text{if } \mathcal{G}_i^t = \emptyset. \end{cases}$$

Under this definition, Delayed Mirror Descent is (note that  $\tilde{\mathcal{G}}_i^t$  is always non-empty here):

---

#### Algorithm 2 Delayed Mirror Descent on Games

---

- 1: Each player  $i$  chooses an initial  $y_i^0$ .
  - 2: **for**  $t = 0, 1, 2, \dots$  **do**
  - 3:   **for**  $i = 1, \dots, N$  **do**
  - 4:      $x_i^t = \arg \max_{x_i \in \mathcal{X}_i} \{\langle y_i^t, x_i \rangle - h_i(x_i)\}$
  - 5:      $y_i^{t+1} = y_i^t + \frac{\alpha^t}{|\tilde{\mathcal{G}}_i^t|} \sum_{s \in \tilde{\mathcal{G}}_i^t} v_i(\mathbf{x}^s)$
  - 6:   **end for**
  - 7: **end for**
- 

### 5.2 Main Delay Assumption

We only make the following assumption on the delays:

**Assumption 2.** For each player  $i$ ,  $\lim_{t \rightarrow \infty} \sum_{s=\min \tilde{\mathcal{G}}_i^t}^t \alpha^s = 0$ .

This assumption essentially requires that no player's delays grow too fast. Note that in particular, players delays can be arbitrarily asynchronous. To make this assumption more concrete, we next give two more explicit delay conditions that satisfy the main delay assumption. As made formal by the following lemma, if the delays are bounded (but not necessarily synchronous), then Assumption 2 is satisfied. Furthermore, by appropriately choosing the sequence  $\alpha^t$ , Assumption 2 can accommodate delays that are unbounded and grow super-linearly.

**Lemma 5.1.** *Let  $\{d_i^s\}_{s=1}^\infty$  be the delay sequences for player  $i$ .*

---

<sup>6</sup>There may not be any gradient information in the first few rounds due to delays. Without loss of generality, we can always start at the first round when there is non-empty gradient information, or equivalently, assume that some gradient is available at  $t = 0$ .

1. If each player  $i$ 's delay is bounded (i.e.  $\exists d \in \mathbb{Z}, d_i^s \leq d, \forall s$ ), then Assumption 2 is satisfied for any positive, non-increasing, not-summable-but-square-summable sequence  $\{\alpha^t\}$ .
2. There exists a positive, non-increasing, not-summable-but-square-summable sequence (e.g.  $\alpha^t = \frac{1}{t \log t \log \log t}$ ) such that if  $d_i^s = O(s \log s), \forall i$ , then Assumption 2 is satisfied.

*Remark 5.1.* The proof to the second claim of Lemma 5.1 indicates that one can also easily obtain slightly larger delay growth rates:  $O(t \log t \log \log t), O(t \log t \log \log t \log \log \log t)$  and so on, by choosing the corresponding step size sequences. Further, it is conceivable that one can identify meaningfully larger delay growth rates that still satisfy Assumption 2, particularly under more restrictions on the degree of delay asynchrony among the players. We leave that for future work.

### 5.3 Convergence of DMD to Nash equilibrium

**Theorem 5.2.** Fix a game with continuous action spaces  $(\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N)$  that admits  $\mathbf{x}^*$  as the unique Nash equilibrium that is  $\lambda$ -variationally stable. Under Assumption 2, the DMD iterate  $\mathbf{x}^t$  given in Algorithm 2 converges to  $\mathbf{x}^*$ , irrespective of the initial point  $\mathbf{x}^0$ .

*Remark 5.2.* The proof here uses a similar framework as the one in Remark 4.2, although the details are somewhat different. Building on the notation and arguments given in Remark 4.2, we again outline three main ingredients that together establish the result. Each of the three pieces will again be proved in detail in the appendix.

1. The gradient update in DMD can be rewritten as:

$$y_i^{t+1} = y_i^t + \frac{\alpha^t}{|\tilde{\mathcal{G}}_i^t|} \sum_{s \in \tilde{\mathcal{G}}_i^t} v_i(\mathbf{x}^s) = y_i^t + \alpha^t v_i(\mathbf{x}^t) + \alpha^t \sum_{s \in \tilde{\mathcal{G}}_i^t} \frac{v_i(\mathbf{x}^s) - v_i(\mathbf{x}^t)}{|\tilde{\mathcal{G}}_i^t|}.$$

By defining:  $b_i^t = \sum_{s \in \tilde{\mathcal{G}}_i^t} \frac{v_i(\mathbf{x}^s) - v_i(\mathbf{x}^t)}{|\tilde{\mathcal{G}}_i^t|}$ , we can write player  $i$ 's gradient update as:

$$y_i^{t+1} = y_i^t + \alpha^t (v_i(\mathbf{x}^t) + b_i^t).$$

By bounding  $b_i^t$ 's magnitude using the delay sequence, Assumption 2 allows us to establish that  $b_i^t$  has negligible impact over time. Mathematically, the claim is that  $\lim_{t \rightarrow \infty} \|b_i^t\|_i^* = 0$ .

2. The joint DMD update can be written as:

$$\mathbf{x}^t = C(\mathbf{y}^t), \tag{5.1}$$

$$\mathbf{y}^{t+1} = \mathbf{y}^t + \alpha^t (v(\mathbf{x}^t) + \mathbf{b}^t). \tag{5.2}$$

Here again using  $\lambda$ -Fenchel coupling as a ‘‘energy’’ function and leveraging the handle on  $\mathbf{b}^t$  given by Claim 1, we show that for any  $\epsilon > 0$  the iterate  $\mathbf{x}^t$  will eventually enter  $B(\mathbf{x}^*, \epsilon)$  and visit  $B(\mathbf{x}^*, \epsilon)$  infinitely often, no matter what the initial point  $\mathbf{x}^0$  is. Furthermore, per Claim 3 in Remark 4.2,  $B(\mathbf{x}^*, \epsilon) \subset \tilde{B}(\mathbf{x}^*, \delta)$ . This implies that  $\mathbf{x}^t$  must enter  $\tilde{B}(\mathbf{x}^*, \delta)$  infinitely often.

3. Again using  $\lambda$ -Fenchel coupling, we show that under DMD, for any ‘‘neighborhood’’  $\tilde{B}(\mathbf{x}^*, \delta)$ , after long enough iterations, if  $\mathbf{x}^t$  ever enters  $\tilde{B}(\mathbf{x}^*, \delta)$ , it will be trapped inside  $\tilde{B}(\mathbf{x}^*, \delta)$  thereafter.

Combining the above three elements, it follows that under DMD, starting from iteration  $t$ ,  $\mathbf{x}^t$  will remain in  $\tilde{B}(\mathbf{x}^*, \delta)$ . Since this is true for any  $\delta > 0$ , we have  $F^\lambda(\mathbf{x}^*, \mathbf{y}^t) \rightarrow 0$  as  $t \rightarrow \infty$ , thereby establishing that  $\mathbf{x}^t = C(\mathbf{y}^t) \rightarrow \mathbf{x}^*$  as  $t \rightarrow \infty$ .

Here again, the result generalizes straightforwardly to the multiple Nash equilibria case (with identical proofs modulo using point-to-set distance metric):

**Theorem 5.3.** Consider a game with continuous action spaces  $(\mathcal{N}, \mathcal{X} = \prod_{i=1}^N \mathcal{X}_i, \{u_i\}_{i=1}^N)$  that admits  $\mathcal{X}^*$  as a  $\lambda$ -variationally stable set (of necessarily all Nash equilibria), for some  $\lambda \in \mathbf{R}_{++}^N$ . Under Assumption 2, the DMD iterate  $\mathbf{x}^t$  given in Algorithm 2 satisfies  $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}^t, \mathcal{X}^*) = 0$ , irrespective of  $\mathbf{x}^0$ , where  $\text{dist}(\cdot, \cdot)$  is the standard point-to-set distance function induced by the norm  $\|\cdot\|$ .



## 6 Conclusion

We examined a model of game-theoretic learning based on OMD with asynchronous and delayed information. By focusing on games with continuous action spaces whose equilibria are variationally stable, we showed that the sequence of play induced by the OMD algorithm converges whenever the feedback delays faced by the players are synchronous and bounded. Subsequently, to tackle fully decentralized, asynchronous environments with unbounded feedback delays (possibly growing sublinearly in the game’s horizon), we showed that our convergence result still holds under delayed mirror descent, a variant of vanilla OMD that leverages past information even in rounds where no feedback is received. To further enhance the distributed aspect of the algorithm, in future work we intend to focus on the case where the players’ gradient input is not only delayed, but also subject to stochastic imperfections – or, taking this to its logical extreme, when players only have observations of their in-game payoffs, and have no gradient information.

## References

- [1] M. BALANDAT, W. KRICHENE, C. TOMLIN, AND A. BAYEN, *Minimizing regret on reflexive banach spaces and learning nash equilibria in continuous zero-sum games*, arXiv preprint arXiv:1606.01261, (2016).
- [2] A. BLUM, *On-line algorithms in machine learning*, in Online algorithms, Springer, 1998, pp. 306–325.
- [3] N. CESA-BIANCHI AND G. LUGOSI, *Prediction, learning, and games*, Cambridge university press, 2006.
- [4] T. DESAUTELS, A. KRAUSE, AND J. W. BURDICK, *Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization.*, Journal of Machine Learning Research, 15 (2014), pp. 3873–3923.
- [5] E. HAZAN, *Introduction to Online Convex Optimization*, Foundations and Trends(r) in Optimization Series, Now Publishers, 2016.
- [6] P. JOULANI, A. GYORGY, AND C. SZEPESVÁRI, *Online learning under delayed feedback*, in Proceedings of the 30th International Conference on Machine Learning (ICML-13), 2013, pp. 1453–1461.
- [7] A. KALAI AND S. VEMPALA, *Efficient algorithms for online decision problems*, Journal of Computer and System Sciences, 71 (2005), pp. 291–307.
- [8] S. KRICHENE, W. KRICHENE, R. DONG, AND A. BAYEN, *Convergence of heterogeneous distributed learning in stochastic routing games*, in Communication, Control, and Computing (Allerton), 2015 53rd Annual Allerton Conference on, IEEE, 2015, pp. 480–487.
- [9] K. LAM, W. KRICHENE, AND A. BAYEN, *On learning how players learn: estimation of learning dynamics in the routing game*, in Cyber-Physical Systems (ICCPS), 2016 ACM/IEEE 7th International Conference on, IEEE, 2016, pp. 1–10.
- [10] N. LITTLESTONE AND M. K. WARMUTH, *The weighted majority algorithm*, INFORMATION AND COMPUTATION, 108 (1994), pp. 212–261.
- [11] J. MAYNARD SMITH AND G. R. PRICE, *The logic of animal conflict*, Nature, 246 (1973), pp. 15–18.
- [12] A. S. NEMIROVSKI AND D. B. YUDIN, *Problem Complexity and Method Efficiency in Optimization*, Wiley, New York, NY, 1983.
- [13] Y. NESTEROV, *Primal-dual subgradient methods for convex problems*, Mathematical Programming, 120 (2009), pp. 221–259.
- [14] K. QUANRUD AND D. KHASHABI, *Online learning with adversarial delays*, in Advances in Neural Information Processing Systems, 2015, pp. 1270–1278.
- [15] R. T. ROCKAFELLAR AND R. J. B. WETS, *Variational Analysis*, vol. 317 of A Series of Comprehensive Studies in Mathematics, Springer-Verlag, Berlin, 1998.
- [16] R. T. ROCKAFELLAR AND R. J.-B. WETS, *Variational analysis*, vol. 317, Springer Science & Business Media, 2009.
- [17] J. B. ROSEN, *Existence and uniqueness of equilibrium points for concave n-person games*, Econometrica: Journal of the Econometric Society, (1965), pp. 520–534.

- [18] S. SHALEV-SHWARTZ ET AL., *Online learning and online convex optimization*, Foundations and Trends® in Machine Learning, 4 (2012), pp. 107–194.
- [19] S. SHALEV-SHWARTZ AND Y. SINGER, *Convex repeated games and Fenchel duality*, in Advances in Neural Information Processing Systems 19, MIT Press, 2007, pp. 1265–1272.
- [20] Y. VIOSSAT AND A. ZAPECHELNYUK, *No-regret dynamics and fictitious play*, Journal of Economic Theory, 148 (2013), pp. 825–842.
- [21] L. XIAO, *Dual averaging methods for regularized stochastic learning and online optimization*, Journal of Machine Learning Research, 11 (2010), pp. 2543–2596.
- [22] M. ZINKEVICH, *Online convex programming and generalized infinitesimal gradient ascent*, in ICML '03: Proceedings of the 20th International Conference on Machine Learning, 2003, pp. 928–936.