
Convex Games in Banach Spaces

Karthik Sridharan
TTI-Chicago
karthik@ttic.edu

Ambuj Tewari
TTI-Chicago
tewari@ttic.edu

Abstract

We study the regret of an online learner playing a multi-round game in a Banach space \mathfrak{B} against an adversary that plays a convex function at each round. We characterize the minimax regret when the adversary plays linear functions in terms of the Martingale type of the dual of \mathfrak{B} . The cases when the adversary plays bounded and uniformly convex functions respectively are also considered. Our results connect online convex learning to the study of the geometry of Banach spaces. We also show that appropriate modifications of the Mirror Descent algorithm from convex optimization can be used to achieve our regret upper bounds. Finally, we provide a version of Mirror Descent that adapts to the changing exponent of uniform convexity of the adversary's functions. This adaptive mirror descent strategy provides new algorithms even for the more familiar Hilbert space case where the loss functions on each round have varying exponents of uniform convexity (curvature).

1 Introduction

Online convex optimization [1, 2, 3] has emerged as an abstraction that allows a unified treatment of a variety of online learning problems where the underlying loss function is convex. In this abstraction, a T -round game is played between the learner (or the player) and an adversary. At each round $t \in \{1, \dots, T\}$, the player makes a move \mathbf{w}_t in some set \mathcal{W} . In the learning context, the set \mathcal{W} will represent some hypothesis space. Once the player has made his choice, the adversary then picks a *convex* function ℓ_t from some set \mathcal{F} and the player suffers "loss" $\ell_t(\mathbf{w}_t)$. In the learning context, the adversary's move ℓ_t encodes the data seen at time t and the loss function used to measure the performance of \mathbf{w}_t on that data. As with any abstraction, on one hand, we lose contact with the concrete details of the problem at hand, but on the other hand, we gain the ability to study related problems from a unified point of view. An added benefit of this abstraction is that it connects online learning with geometry of convex sets, theory of optimization and game theory.

An important notion in the online setting is that of the cumulative *regret* incurred by the player which is the difference between the cumulative loss of the player and the cumulative loss for the best fixed move in hindsight. The goal of regret minimizing algorithms is to control the growth rate of the regret as a function of T . There has been a huge amount of work characterizing the best regret rates possible under a variety of assumptions on the player's and adversary's sets \mathcal{W} and \mathcal{F} . With a few exceptions that we mention later, most of the work has been in the setting where these sets live in some Euclidean space \mathbb{R}^d . Whenever the results do not explicitly involve the dimensionality d , they are also usually applicable in any Hilbert space \mathfrak{H} . There also has been a lot of work dealing with ℓ_p spaces for $p \in [1, 2]$. But here, the fact exploited is that a strongly convex function, (with dimension-independent constant of strong convexity) is available in these Banach spaces. There has been less work dealing with arbitrary Banach spaces (where strongly convex functions might not exist). Our focus in this paper is to extend the study of optimal regret rates to the case when the set \mathcal{W} lives in a general Banach space \mathfrak{B} .

Before we explain what our specific contributions are, let us briefly mention two examples to show why one might want to move beyond Hilbert spaces and consider general Banach spaces. A first family of examples is ℓ_p spaces with $p \geq 2$. As p grows, the ℓ_p balls become larger and thus have better approximation properties. On the other hand, as we show below, the cumulative regret rate when competing against a fixed size ball in ℓ_p is $O(T^{1/p})$. So, there is a trade-off here between

approximation properties and the regret rate (or the estimation error in a stochastic setting). In high-dimensions, one can easily construct examples where the approximation property dominates the trade-off and it is advantageous to use $p > 2$ even though the estimation error suffers (see appendix for a worked out example). Another example is prediction with squared loss where the learner is trying to predict a signal y_t given input x_t . At each step, the learner chooses a *function* f_t and suffers the loss $(y_t - f_t(x_t))^2$. Here, the viewpoint of considering f_t as a “point” in a function space is very fruitful and it very natural to assume that the space of functions that the learner can use is a Banach space of functions. For more details, see [4].

In the Hilbert space setting, it is known that the “degree of convexity” or “curvature” of the functions ℓ_t played by the adversary has a significant impact on the achievable regret rates. For example, if the adversary can play arbitrary convex and Lipschitz functions, the best regret possible is $O(\sqrt{T})$. However, if the adversary is constrained to play *strongly convex* and Lipschitz functions, the regret can be brought down to $O(\log T)$. Further, it is also known, via minimax lower bounds [5], that these are the best possible rates in these situations. In a general Banach space, strongly convex functions might not even exist. We will, therefore, need a generalization of strong convexity called *q-uniform convexity* (strong convexity is 2-uniform convexity). There will, in general, be a number $q^* \in [2, \infty)$ such that q^* -uniformly convex functions are the “most curved” functions available on \mathfrak{B} . There are, again, two extremes: the adversary can play either arbitrary convex-Lipschitz functions or q^* -uniformly convex functions. We show that the minimax optimal rates in these two situations are of the order $\Theta^*(T^{1/p^*})$ and $\Theta^*(T^{2-q^*})$ respectively¹ where p^* is the Martingale type of the dual \mathfrak{B}^* of \mathfrak{B} . A Hilbert space has $p^* = q^* = 2$. We also give upper and lower bounds for the intermediate case when the adversary plays q -uniformly convex functions for $q > q^*$. This case, as far as we know, has not been analyzed *even in the Hilbert space setting*.

Another natural game that we have not seen analyzed before is the convex-bounded game: here the adversary plays convex and bounded functions. Of course, being Lipschitz on a bounded domain implies boundedness but the reverse implication is false: a bounded function can have arbitrarily bad Lipschitz constant. For the convex-bounded game, we do not have a tight characterization but we can give non-trivial upper bounds. However, these upper bounds suffice to prove, for example, that the following three properties of \mathfrak{B} are equivalent: (1) the convex-bounded game when the player plays in the unit ball of \mathfrak{B} has non-trivial (i.e. $o(T)$) minimax regret; (2) the corresponding convex-Lipschitz game has non-trivial minimax regret; and (3) the Banach space \mathfrak{B} is super-reflexive.

We further describe player strategies that achieve the optimal rates for these convex games. These strategies are all based on the *Mirror Descent* algorithm that originated in the convex optimization literature [6]. Usually Mirror Descent is run with a strongly convex function but it turns out that it can also be analyzed in our Banach space setting if it is run with a q -uniformly convex function Ψ . Moreover, with the correct choice of Ψ , it achieves all the upper bounds presented in this paper. Thus, part of our contribution is also to show the remarkable properties of the Mirror Descent algorithm.

Our final contribution is an adaptive algorithm, building on previous work, that *adapts* to the exponent of uniform convexity in the adversary’s functions. Our results have novel implications even in a Hilbert space. For example, [7] showed how to adapt to an adversary that mixes linear and strong convex functions in its moves. We can now allow this mix to also consist of functions with intermediate degrees of uniform convexity.

Related work The idea of exploiting minimax-maximin duality to analyze optimal regret rates also appears in the recent work of Abernethy et al. [8]. The earliest papers we know of that explore the connection of the type of a Banach space to learning theory are those of Donahue et al. [9] and Gurvits [10]. Mendelson and Schechtman [11] gave estimates of the fat-shattering dimension of linear functionals on a Banach space in terms of its type. In the context of online regression with squared loss, Vovk [4] also gives rates worse than $O(\sqrt{T})$ when the class of functions one is competing against is not in a Hilbert space, but in some Banach space. He also mentions “Banach Learning” as an open problem in his *online prediction wiki*². For recent work exploring Banach spaces for learning applications, see [12, 13, 14]. These papers also give more reasons for considering general Banach spaces in Learning Theory.

Outline The rest of the paper is organized as follows. In Section 2, we formally define the minimax and maximin values of the game between a player and an adversary. We also introduce the notions of Martingale type and uniform convexity from functional analysis. Section 3 considers convex-Lipschitz and linear games. A key result in that section is Theorem 4 which gives a characterization

¹Our informal $\Theta^*(\cdot)$ notation hides factors that are $o(T^\epsilon)$ for every $\epsilon > 0$.

²<http://onlineprediction.net/?n=Open.BanachLearning>

of the minimax regret of these games in terms of the type. Convex-bounded games are treated in Section 4 and the equivalence of super-reflexivity of the space to the existence of player strategies achieving non-trivial regret guarantees is established (Corollary 6). We next consider the case when the adversary plays “curved” functions in Section 5. Here the regret depends on the exponent of uniform convexity of the functions played by the adversary (Theorem 8). In Section 6, we describe player strategies based on the Mirror Descent algorithm that achieve the upper bounds presented in Section 3 and Section 4. In Section 7, using the techniques of Bartlett et al. [7], we give a player strategy that adapts to the exponent of uniform convexity of the functions being played by the adversary. We conclude in Section 8 by exploring directions for future work including a discussion on how the ideas in this paper might lead to practical algorithms. All proofs omitted from the main body of the paper can be found in the appendix.

2 Preliminaries

2.1 Regret and Minimax Value

Our primary objects of study are certain T -round games where the player \mathcal{P} makes moves in convex set \mathcal{W} contained in some (real) separable Banach space \mathfrak{B} . The adversary \mathcal{A} plays bounded continuous convex functions on \mathcal{W} chosen from a fixed function class \mathcal{F} . topology induced by the norm $\|\ell\| := \sup_{\mathbf{w} \in \mathcal{W}} |\ell(\mathbf{w})|$. The game proceeds as follows.

For $t = 1$ to T

- \mathcal{P} plays $\mathbf{w}_t \in \mathcal{W}$,
- \mathcal{A} plays $\ell_t \in \mathcal{F}$,
- \mathcal{P} suffers $\ell_t(\mathbf{w}_t)$.

For given sequences $\mathbf{w}_{1:T}, \ell_{1:T}$, we define the *regret* of \mathcal{P} as,

$$\text{Reg}(\mathbf{w}_{1:T}, \ell_{1:T}) := \sum_{t=1}^T \ell_t(\mathbf{w}_t) - \inf_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \ell_t(\mathbf{w}) .$$

Given the tuple $(T, \mathfrak{B}, \mathcal{W}, \mathcal{F})$, we can define the minimax value of the above game as follows.

Definition 1. Given $T \geq 1$ and $\mathfrak{B}, \mathcal{W}, \mathcal{F}$ satisfying the conditions above, define the minimax value,

$$V_{T, \mathfrak{B}}(\mathcal{W}, \mathcal{F}) := \inf_{\mathbf{w}_1 \in \mathcal{W}} \sup_{\ell_1 \in \mathcal{F}} \cdots \inf_{\mathbf{w}_T \in \mathcal{W}} \sup_{\ell_T \in \mathcal{F}} \text{Reg}(\mathbf{w}_{1:T}, \ell_{1:T}) .$$

When T and \mathfrak{B} are clear from context, we will simply denote the minimax value by $V(\mathcal{W}, \mathcal{F})$. A *player strategy* (or \mathcal{P} -strategy) W is a sequence (W_1, \dots, W_T) of functions such that $W_t : \mathcal{F}^{t-1} \rightarrow \mathcal{W}$. For a strategy W , we define the regret as,

$$\text{Reg}(W, \ell_{1:T}) := \sum_{t=1}^T \ell_t(W_t(\ell_{1:t-1})) - \inf_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \ell_t(\mathbf{w}) .$$

In terms of player strategies, the minimax value takes a simpler form,

$$V_{T, \mathfrak{B}}(\mathcal{W}, \mathcal{F}) = \inf_W \sup_{\ell_{1:T}} \text{Reg}(W, \ell_{1:T}) ,$$

where the supremum is over all sequences $\ell_{1:T} \in \mathcal{F}^T$. Let Q denote distributions over \mathcal{F}^T . We can define the maximin value,

$$U_{T, \mathfrak{B}}(\mathcal{W}, \mathcal{F}) := \sup_Q \inf_W \mathbb{E}_{\ell_{1:T} \sim Q} [\text{Reg}(W, \ell_{1:T})] .$$

One has that

$$V_{T, \mathfrak{B}}(\mathcal{W}, \mathcal{F}) \geq U_{T, \mathfrak{B}}(\mathcal{W}, \mathcal{F}) . \tag{1}$$

This inequality will be the starting point for our subsequent analysis. For more about the minimax and maximin values refer, [15, p. 31].

2.2 Martingale type and Uniform Convexity

One of the goals of this paper is to characterize the minimax value in terms of the geometric properties of the space \mathfrak{B} and “degree of convexity” inherent in the functions in \mathcal{F} . Among the geometric characteristics of a Banach space \mathfrak{B} , the most useful for us is the notion of *Martingale*

type (or *M-type*) of \mathfrak{B} . A Banach space \mathfrak{B} has *M-type* p if there is some constant C such that for any $T \geq 1$ and martingale difference sequence $\mathbf{d}_1, \dots, \mathbf{d}_T$ with values in \mathfrak{B} ,

$$\mathbb{E} \left[\left\| \sum_{t=1}^T \mathbf{d}_t \right\|^2 \right] \leq C \left(\sum_{t=1}^T \mathbb{E} [\|\mathbf{d}_t\|^p] \right)^{2/p}. \quad (2)$$

We also define the best M-type possible for a Banach space,

$$p^*(\mathfrak{B}) := \sup \{p : \mathfrak{B} \text{ has M-type } p\}. \quad (3)$$

A Banach space \mathfrak{B} has *M-cotype* q if there is some constant C such that for any $T \geq 1$ and martingale difference sequence $\mathbf{d}_1, \dots, \mathbf{d}_T$ with values in \mathfrak{B} ,

$$\left(\sum_{t=1}^T \mathbb{E} [\|\mathbf{d}_t\|^q] \right)^{1/q} \leq C \mathbb{E} \left[\left\| \sum_{t=1}^T \mathbf{d}_t \right\|^2 \right]^{1/2}. \quad (4)$$

A closely related notion in Banach space theory is that of super-reflexivity. Refer [16] for more details.

Definition 2. A Banach space \mathfrak{B} is *super-reflexive* if no non-reflexive space is finitely representable in \mathfrak{B} .

A result of Pisier [16] shows that a Banach space \mathfrak{B} has non-trivial M-type ($p^* > 1$) (or equivalently non-trivial Martingale co-type) if and only if it is super-reflexive.

To measure the “degree of convexity” of the functions played by the adversary, we need the notion of uniform convexity. Let $\|\cdot\|$ be the norm associated with a Banach space \mathfrak{B} . A function $\ell : \mathfrak{B} \rightarrow \mathbb{R}$ is said to be (C, q) -uniformly convex on \mathfrak{B} if there is some constant $C > 0$ such that, for any $\mathbf{v}_1, \mathbf{v}_2 \in \mathfrak{B}$ and any $\theta \in [0, 1]$,

$$\ell(\theta \mathbf{v}_1 + (1 - \theta) \mathbf{v}_2) \leq \theta \ell(\mathbf{v}_1) + (1 - \theta) \ell(\mathbf{v}_2) - \frac{C\theta(1 - \theta)}{q} \|\mathbf{v}_1 - \mathbf{v}_2\|^q.$$

If $C \geq 1$ we simply say that the function ℓ is q -uniformly convex.

The following remarkable theorem of Pisier [17] shows that the concept of M-types and existence of uniformly convex functions in the Banach space are intimately connected.

Theorem (Pisier). A Banach space \mathfrak{B} has M-cotype q iff there exists a q -uniformly convex function on \mathfrak{B} .

Now consider some $p \in [1, p_M^*(\mathfrak{B}^*)]$. Then, by definition, \mathfrak{B}^* has M-type p . It is a fact that \mathfrak{B}^* has M-type p iff \mathfrak{B} has M-cotype $\frac{p}{p-1}$ [16, Chapter 6]. Thus, \mathfrak{B} has M-cotype $\frac{p}{p-1}$. Now, Pisier’s theorem guarantees the existence of a $\frac{p}{p-1}$ -uniformly convex function on \mathfrak{B} .

For a convex function $\ell : \mathfrak{B} \rightarrow \mathbb{R}$, its *subdifferential* at a point \mathbf{v} is defined as,

$$\partial \ell(\mathbf{v}) = \{\boldsymbol{\lambda} \in \mathfrak{B}^* : \forall \mathbf{v}', \ell(\mathbf{v}') \geq \ell(\mathbf{v}) + \boldsymbol{\lambda}(\mathbf{v}' - \mathbf{v})\},$$

where \mathfrak{B}^* denotes the *dual space* of \mathfrak{B} . This consists of all continuous linear functions on \mathfrak{B} with norm defined as $\|\ell\|_* := \sup_{\mathbf{w}: \|\mathbf{w}\| \leq 1} \ell(\mathbf{w})$. If $\partial \ell(\mathbf{v})$ is a singleton then we say ℓ is differentiable at \mathbf{v} and denote the unique member of $\partial \ell(\mathbf{v})$ by $\nabla \ell(\mathbf{v})$. If ℓ is differentiable at \mathbf{v}_1 , define the *Bregman divergence* associated with ℓ as,

$$\Delta_\ell(\mathbf{v}_1, \mathbf{v}_2) = \ell(\mathbf{v}_1) - \ell(\mathbf{v}_2) - \nabla \ell(\mathbf{v}_2)(\mathbf{v}_1 - \mathbf{v}_2).$$

Recall that a function $\ell : \mathfrak{B} \rightarrow \mathbb{R}$ is L -Lipschitz on some set $\mathcal{W} \in \mathfrak{B}$ if for any $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{W}$, we have $\ell(\mathbf{v}_1) - \ell(\mathbf{v}_2) \leq L \|\mathbf{v}_1 - \mathbf{v}_2\|$. Given a set \mathcal{W} in a Banach space \mathfrak{B} , we define the following natural sets of convex functions on \mathcal{W} ,

$$\begin{aligned} \text{lin}(\mathcal{W}) &:= \{\ell : \ell \text{ is linear and 1-Lipschitz on } \mathcal{W}\}, \\ \text{cvx}(\mathcal{W}) &:= \{\ell : \ell \text{ is convex and 1-Lipschitz on } \mathcal{W}\}, \\ \text{bdd}(\mathcal{W}) &:= \{\ell : \ell \text{ is convex and bounded by 1 on } \mathcal{W}\}, \\ \text{cvx}_{q,L}(\mathcal{W}) &:= \{\ell : \ell \text{ is } q\text{-uniformly convex and } L\text{-Lipschitz on } \mathcal{W}\}. \end{aligned}$$

In the following sections, we will analyze the minimax value $V(\mathcal{W}, \mathcal{F})$ when the adversary’s set of moves is one these 4 sets defined above. For readability, we will drop the dependence of these sets on \mathcal{W} when it is clear from context. For example, we will refer to $V(\mathcal{W}, \text{cvx}(\mathcal{W}))$ simply as $V(\mathcal{W}, \text{cvx})$.

3 Convex-Lipschitz and Linear Games

Given a Banach space \mathfrak{B} with a norm $\|\cdot\|$, denote its unit ball by $U(\mathfrak{B}) := \{\mathbf{v} \in \mathfrak{B} : \|\mathbf{v}\| \leq 1\}$. Consider the case when the \mathcal{P} 's set \mathcal{W} is the unit ball $U(\mathfrak{B})$ for some \mathfrak{B} . This setting is not as restrictive as it sounds since any bounded *symmetric* convex set K in a vector space \mathbf{V} gives a Banach space $\mathfrak{B} = (\mathbf{V}, \|\cdot\|_K)$, where we equip \mathbf{V} with the norm,

$$\|\mathbf{v}\|_K := \inf \{\alpha > 0 : \mathbf{v} \in \alpha K\} . \quad (5)$$

Moreover, (the closure of) K is the unit ball of this Banach space.

So, fix \mathfrak{B} and consider the case $\mathcal{W} = U(\mathfrak{B})$, $\mathcal{F} = \text{cvx}(\mathcal{W})$. Theorem 14 given in [5] gives us $V(\mathcal{W}, \text{cvx}) = V(\mathcal{W}, \text{lin})$. We are therefore led to consider the case $\mathcal{W} = U(\mathfrak{B})$, $\mathcal{F} = \text{lin}(\mathcal{W})$. Note that $\text{lin}(\mathcal{W})$ is simply the unit ball $U(\mathfrak{B}^*)$. The theorem below relates the minimax value $V(U(\mathfrak{B}), \text{lin})$ to the behaviour of martingale difference sequences in \mathfrak{B}^* .

Theorem 3. *The minimax value $V(U(\mathfrak{B}), \text{lin})$ of the linear game is bounded as,*

$$V(U(\mathfrak{B}), \text{lin}) \geq \sup_{\mathbf{M}} \mathbb{E} \left[\left\| \sum_{t=1}^T \ell_t \right\|_{\star} \right] .$$

where the supremum is over distributions \mathbf{M} of martingale difference $(\ell_t)_{t=1}^T$ such that each $\ell_t \in U(\mathfrak{B}^*)$.

Proof. Recall that we denote a general distribution over \mathcal{A} 's sequences by Q and \mathcal{P} -strategies by W . Equation (1) gives us,

$$V(\mathcal{W}, \mathcal{F}) \geq \sup_Q \inf_W \mathbb{E}_{\ell_{1:T} \sim Q} [\text{Reg}(W, \ell_{1:T})] .$$

If we define $V_Q := \inf_W \mathbb{E}_{\ell_{1:T} \sim Q} [\text{Reg}(W, \ell_{1:T})]$ we can succinctly write, $V(\mathcal{W}, \mathcal{F}) = \sup_Q V_Q$.

Now, let us fix a distribution Q and denote the conditional expectation w.r.t. $\ell_{1:t}$ by $\mathbb{E}_t[\cdot]$ and the full expectation w.r.t. $\ell_{1:T}$ by $\mathbb{E}[\cdot]$. Substituting the definition of regret and noting that the infimum in its definition does not depend on the strategy W , we get

$$V_Q \geq \inf_W \left(\mathbb{E} \left[\sum_{t=1}^T \ell_t(W_t(\ell_{1:t-1})) \right] \right) - \mathbb{E} \left[\inf_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \ell_t(\mathbf{w}) \right] . \quad (6)$$

Let us simplify the infimum over \mathcal{P} -strategies as follows,

$$\begin{aligned} \inf_W \mathbb{E} \left[\sum_{t=1}^T \ell_t(W_t(\ell_{1:t-1})) \right] &= \inf_W \sum_{t=1}^T \mathbb{E} [\mathbb{E}_{t-1} [\ell_t(W_t(\ell_{1:t-1}))]] = \sum_{t=1}^T \inf_{W_t} \mathbb{E} [\mathbb{E}_{t-1} [\ell_t(W_t(\ell_{1:t-1}))]] \\ &= \sum_{t=1}^T \mathbb{E} \left[\inf_{\mathbf{w}_t \in \mathcal{W}} \mathbb{E}_{t-1} [\ell_t(\mathbf{w}_t)] \right] . \end{aligned}$$

Substituting this into (6), we get,

$$\begin{aligned} V_Q &\geq \sum_{t=1}^T \mathbb{E} \left[\inf_{\mathbf{w}_t \in \mathcal{W}} \mathbb{E}_{t-1} [\ell_t(\mathbf{w}_t)] \right] - \mathbb{E} \left[\inf_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \ell_t(\mathbf{w}) \right] \\ &= \mathbb{E} \left[\sup_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T -\ell_t(\mathbf{w}) \right] - \sum_{t=1}^T \mathbb{E} \left[\sup_{\mathbf{w}_t \in \mathcal{W}} \mathbb{E}_{t-1} [-\ell_t(\mathbf{w}_t)] \right] . \end{aligned} \quad (7)$$

Since the losses ℓ_t are linear and \mathcal{W} is the unit ball, we can re-write the above as

$$V_Q \geq \mathbb{E} \left[\left\| \sum_{t=1}^T -\ell_t(\mathbf{w}) \right\| \right] - \sum_{t=1}^T \mathbb{E} [\|\mathbb{E}_{t-1} [-\ell_t(\mathbf{w}_t)]\|]$$

If we restrict ourselves to only distribution Q such that (ℓ_1, \dots, ℓ_T) are martingale difference sequences then clearly $\mathbb{E}_{t-1} [-\ell_t(\mathbf{w}_t)] = 0$ and so

$$\sup_Q V_Q \geq \sup_{\mathbf{M}} \mathbb{E} \left[\left\| \sum_{t=1}^T -\ell_t(\mathbf{w}) \right\| \right]$$

and so we get the lower bound. \square

Given the above result, we can now characterize the minimax value $V(U(\mathfrak{B}), \text{lin})$ in terms of the $p^*(\mathfrak{B}^*)$ where \mathfrak{B}^* is the dual space of \mathfrak{B} .

Theorem 4. *For all p, p' such that $p < p^*(\mathfrak{B}^*) < p'$, there exists a constant C such that,*

$$\Omega\left(T^{1/p'}\right) = V(U(\mathfrak{B}), \text{lin}) = V(U(\mathfrak{B}), \text{cvx}) \leq CT^{1/p}. \quad (8)$$

If the supremum in (3) is achieved, the upper bound also holds for $p = p^(\mathfrak{B}^*)$.*

Proof. To prove the lower bound, note that for any finite dimensional Banach space has $p^*(\mathfrak{B}^*) = 2$ with a possibly dimension dependent constant. In this case, the lower bound of \sqrt{T} for the linear game is easy: pick some non-zero vector, say $\ell \in U(\mathfrak{B}^*)$, and use ℓ or $-\ell$ at random in Theorem 3 and the lower bound follows. Therefore, assume \mathfrak{B} is infinite dimensional. The lower bound in this case is proved using Lemma 12 which in turn is proved using ideas from [16]. Lemma 12 shows that for any $p' > p^*(\mathfrak{B}^*)$,

$$\frac{\sup_{\mathbf{M}} \mathbb{E} \left[\left\| \sum_{t=1}^T \mathbf{d}_t \right\| \right]}{T^{1/p'}} \rightarrow \infty$$

However we have from Theorem 3 that

$$V(U(\mathfrak{B}), \text{lin}) \geq \sup_{\mathbf{M}} \mathbb{E} \left[\left\| \sum_{t=1}^T \mathbf{d}_t \right\| \right]$$

Hence we can conclude that for any $p' > p^*$ asymptotically $T^{1/p'}$ is dominated by $V(U(\mathfrak{B}), \text{lin})$ and hence the lower bound.

As for the upper bounds, if $p^*(\mathfrak{B}^*) = 1$ then the upper bound is trivial. On the other hand, when M-type is non-trivial, then M-type p implies M-cotype $q = \frac{p}{p-1}$. Therefore, for each $p \in (1, p^*(\mathfrak{B}^*))$, \mathfrak{B} is of M-cotype q . By Theorem (Pisier), there exists a q -uniformly convex function on \mathfrak{B} . Using this function in the Mirror Descent algorithm, Proposition 9 yields the required upper bound. \square

Although we have stated the above theorem for the special case when $\mathcal{W} = U(\mathfrak{B})$ and $\mathcal{F} = \text{lin}(U(\mathfrak{B}))$, it actually gives us regret rates when \mathcal{P} plays in $rU(\mathfrak{B})$ and \mathcal{A} plays L -Lipschitz linear functions via the following equality which is easy to prove from first principles,

$$V(rU(\mathfrak{B}), L \text{lin}(U(\mathfrak{B}))) = r \cdot L \cdot V(U(\mathfrak{B}), \text{lin}(U(\mathfrak{B}))) . \quad (9)$$

For the special case when $(\mathfrak{B}, \mathfrak{B}^*) = (\ell_q, \ell_p)$ for $1/p + 1/q = 1$, the rates given by the theorem above become $\Theta(T^{1/p})$ and $\Theta(T^{1/2})$ when $p \in (1, 2]$ and $p \in [2, \infty)$ respectively.

4 Convex-Bounded Games

Another natural game we consider is one in which \mathcal{P} plays from some convex set \mathcal{W} and \mathcal{A} plays some convex function *bounded* by 1. In the following theorem, we bound the value of such a game.

Theorem 5. *For all p, p' such that $p < p^*(\mathfrak{B}^*) < p'$, there is a constant C such that,*

$$\Omega\left(T^{1/p'}\right) \leq V(U(\mathfrak{B}), \text{bdd}) \leq CT^{1/p+1/2q} . \quad (10)$$

where $q = \frac{p}{p-1}$. If the supremum in (3) is achieved, the upper bound also holds for $p = p^(\mathfrak{B}^*)$.*

Proof of Theorem 5. Let us actually consider the case $\mathcal{W} = rU(\mathfrak{B})$ and $\mathcal{F} = \text{bdd}(rU(\mathfrak{B}))$. The bounds will turn out to be independent of r . Note that we have the inclusion, $\text{bdd}(rU(\mathfrak{B})) \supseteq \frac{1}{r} \text{lin}(U(\mathfrak{B}))$ which implies

$$V(rU(\mathfrak{B}), \text{bdd}(rU(\mathfrak{B}))) \geq V\left(rU(\mathfrak{B}), \frac{1}{r} \text{lin}(U(\mathfrak{B}))\right) .$$

The lower bound is now immediate due to lower bound on linear game on unit ball (Theorem 4) and property (9).

For the upper bound, note that any convex function bounded by 1 on the scaled ball $rU(\mathfrak{B})$ is $\frac{2}{\epsilon r}$ Lipschitz on the ball of radius $(1 - \epsilon)r$ [18]. Hence, by upper bound in Theorem 4 and property

(9), we see that there exists a strategy say W whose regret on the ball of radius $r(1 - \epsilon)$ is bounded by $\frac{C}{\epsilon} T^{\frac{1}{p}}$ for any $p \in [1, p^*(\mathfrak{B}^*)]$. That is

$$\sum_{t=1}^T \ell_t(W_t) - \operatorname{argmin}_{\mathbf{w} \in r(1-\epsilon)U(\mathfrak{B})} \sum_{t=1}^T \ell_t(\mathbf{w}) \leq \frac{C}{\epsilon} T^{1/p}, \quad \forall p \in [1, p^*(\mathfrak{B}^*)] \quad (11)$$

Let $\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w} \in rU(\mathfrak{B})} \sum_{t=1}^T \ell_t(\mathbf{w})$. Now we consider two cases, first when $\mathbf{w}^* \in (1 - \epsilon)rU(\mathfrak{B})$. In this case the regret of the strategy on the unit ball is bounded by $\frac{C}{\epsilon} T^{1/p}$ for all $p \in [1, p^*(\mathfrak{B}^*)]$. On the other hand if $\mathbf{w}^* \notin (1 - \epsilon)rU(\mathfrak{B})$, then define $\mathbf{w}_\epsilon^* = \frac{r(1-\epsilon)\mathbf{w}^*}{\|\mathbf{w}^*\|}$. Note that $\mathbf{w}_\epsilon^* \in r(1 - \epsilon)U(\mathfrak{B})$. In this case by convexity of $\sum_{t=1}^T \ell_t(\mathbf{w})$, we have that

$$\sum_{t=1}^T \ell_t(\mathbf{w}_\epsilon^*) \leq \frac{r(1-\epsilon)}{\|\mathbf{w}^*\|} \sum_{t=1}^T \ell_t(\mathbf{w}^*) + \left(1 - \frac{r(1-\epsilon)}{\|\mathbf{w}^*\|}\right) \sum_{t=1}^T \ell_t(0)$$

Hence, we have that

$$\sum_{t=1}^T \ell_t(\mathbf{w}_\epsilon^*) - \sum_{t=1}^T \ell_t(\mathbf{w}^*) \leq \left(\frac{r(1-\epsilon)}{\|\mathbf{w}^*\|} - 1\right) \sum_{t=1}^T \ell_t(\mathbf{w}^*) + \left(1 - \frac{r(1-\epsilon)}{\|\mathbf{w}^*\|}\right) T \leq 2 \left(1 - \frac{r(1-\epsilon)}{\|\mathbf{w}^*\|}\right) T$$

However, since $\|\mathbf{w}^*\| \leq r$ we see that $\sum_{t=1}^T \ell_t(\mathbf{w}_\epsilon^*) - \sum_{t=1}^T \ell_t(\mathbf{w}^*) \leq 2\epsilon T$. Combining with (11) we see that for any $p \in [1, p^*(\mathfrak{B}^*)]$, $\sum_{t=1}^T \ell_t(\mathbf{w}_t) - \sum_{t=1}^T \ell_t(\mathbf{w}^*) \leq \frac{C}{\epsilon} T^{1/p} + 2\epsilon T$. Choosing $\epsilon = \sqrt{\frac{C}{2T^{1/q}}}$ we get the required upper bound. \square

Although we have stated the above result for the unit ball, the proof given above shows that the bounds are independent of the radius of the ball in which the player is playing.

Theorems 4 and 5 imply the following interesting corollary.

Corollary 6. *The following statements are equivalent :*

1. $V(U(\mathfrak{B}), \text{bdd}) = o(T)$.
2. $V(U(\mathfrak{B}), \text{cvx}) = o(T)$.
3. \mathfrak{B}^* has non-trivial M -type
4. Both \mathfrak{B} and \mathfrak{B}^* are super-reflexive.

Proof of Corollary 6. The implications $\mathbf{3} \Rightarrow \mathbf{1}$ and $\mathbf{3} \Rightarrow \mathbf{2}$ follow from the upper bounds in Theorems 4 and 5. The reverse implications $\mathbf{1} \Rightarrow \mathbf{3}$ and $\mathbf{2} \Rightarrow \mathbf{3}$, in turn, follow from the lower bounds in those theorems. The equivalence of $\mathbf{3}$ and $\mathbf{4}$ is due to deep results of Pisier [19]. \square

The convex-Lipschitz games (and q -uniformly convex-Lipschitz games considered below) depend, by definition, not only on the player's set \mathcal{W} but also on the norm $\|\cdot\|$ of the underlying Banach space \mathfrak{B} . This is because \mathcal{A} 's functions are required to be Lipschitz w.r.t. $\|\cdot\|$. However, note that the convex-bounded game can be defined only in terms of the player set \mathcal{W} . Hence, one would expect the value of the game to be characterized solely by properties of set \mathcal{W} . This is what the following corollary confirms.

Corollary 7. *Let \mathcal{W} be any symmetric bounded convex subset of a vector space \mathbf{V} . The value of the bounded convex game on \mathcal{W} is non-trivial (i.e. $o(T)$) iff the Banach space $(\mathbf{V}, \|\cdot\|_{\mathcal{W}})$ (where $\|\cdot\|_{\mathcal{W}}$ is defined as in (5)) is super-reflexive.*

5 Uniformly Convex-Lipschitz Games

For any Hilbert space \mathfrak{H} , it is known that $V(U(\mathfrak{H}), \text{cvx}_{2,L})$ is much smaller than $V(U(\mathfrak{H}), \text{lin})$, i.e. the game is much easier for \mathcal{P} if \mathcal{A} plays 2-uniformly convex (also called *strongly convex*) functions. In fact, it is known that $V(U(\mathfrak{H}), \text{cvx}_{2,L}) = \Theta(L^2 \log T)$ while $V(U(\mathfrak{H}), \text{lin}) = \Theta(\sqrt{T})$. This suggests that we should get a rate between $\log(T)$ and \sqrt{T} if \mathcal{A} plays q -uniformly convex functions in a Hilbert space \mathfrak{H} for some $q > 2$. As far as we know, there is no characterization of the achievable rates for these intermediate situations even for Hilbert spaces. Our next result provides upper and lower bounds for $V(U(\mathfrak{B}), \text{cvx}_{q,L})$ in a Banach space, when the exponent of \mathcal{A} 's uniform convexity lies in an intermediate range between its minimum possible value q^* and its maximum value ∞ . It is easy to see that the minimum possible value q^* is $p^*(\mathfrak{B}^*)/(p^*(\mathfrak{B}^*) - 1)$.

Theorem 8. Let $q^* = \frac{p^*(\mathfrak{B}^*)}{p^*(\mathfrak{B}^*)-1}$ and $q \in (q^*, \infty)$. Let $p = q/(q-1)$ be the dual exponent of q . Then, as long as $\text{cvx}_{q,L}$ is non-empty, there exists K that depends on L such that for all $p' > p^*(\mathfrak{B}^*)$,

$$\Omega \left(\left(1 - \frac{1}{L}\right) \frac{1}{p} T^{1-p+\frac{p}{p'}} \right) \leq V(U(\mathfrak{B}), \text{cvx}_{q,L}) \leq KT^{\min\{2-p, 1/p^*(\mathfrak{B}^*)\}}. \quad (12)$$

Proof. We start by proving the lower bound. To this end note that if $\text{cvx}_{q,L}$ is non-empty, then the adversary plays L -Lipschitz, q -uniformly convex loss functions. Note that given such a function, there exists a norm $|\cdot|$ such that $|\cdot| \leq \|\cdot\| \leq L|\cdot|$ (ie. an equivalent norm) and $\frac{1}{q}|\cdot|^q$ is a q -uniformly convex function [20]. Given this we consider a game where adversary plays only functions from

$$\text{lincvx}_{q,L}(\mathcal{W}) := \left\{ \ell(\mathbf{w}) = \langle \mathbf{w}, \mathbf{x} \rangle + \frac{1}{q} |\mathbf{w}|^q : |x|_* \leq L - 1 \right\}$$

Note that since the above is L -Lipschitz w.r.t. $|\cdot|$, it is automatically L -Lipchitz w.r.t. $\|\cdot\|$. Hence $\text{lincvx}_{q,L} \subseteq \text{cvx}_{q,L}$, and so we have that $V(U(\mathfrak{B}), \text{lincvx}_{q,L}) \leq V(U(\mathfrak{B}), \text{cvx}_{q,L})$ However note that

$$V(U(\mathfrak{B}), \text{lincvx}_{q,L}) \geq \inf_W \sup_P \mathbb{E}_{\ell_{1:T} \sim P} \text{Reg}(W, \ell_{1:T}) \quad (13)$$

Also note that

$$\begin{aligned} \text{Reg}(W, \ell_{1:T}) &= \sum_{t=1}^T \left(\langle \mathbf{x}_t, \mathbf{w}_t \rangle + \frac{|\mathbf{w}_t|^q}{q} \right) - \inf_{\mathbf{w} \in U(\mathfrak{B})} \sum_{t=1}^T \left(\langle \mathbf{x}_t, \mathbf{w} \rangle + \frac{|\mathbf{w}|^q}{q} \right) \\ &= \sum_{t=1}^T \left(\langle \mathbf{x}_t, \mathbf{w}_t \rangle + \frac{|\mathbf{w}_t|^q}{q} \right) + T \sup_{\mathbf{w} \in U(\mathfrak{B})} \left(\left\langle -\frac{1}{T} \sum_{t=1}^T \mathbf{x}_t, \mathbf{w} \right\rangle - \frac{|\mathbf{w}|^q}{q} \right) \\ &\geq \sum_{t=1}^T \left(\langle \mathbf{x}_t, \mathbf{w}_t \rangle + \frac{\|\mathbf{w}_t\|^q}{Lq} \right) + T \sup_{\mathbf{w} \in U(\mathfrak{B})} \left(\left\langle -\frac{1}{T} \sum_{t=1}^T \mathbf{x}_t, \mathbf{w} \right\rangle - \frac{\|\mathbf{w}\|^q}{q} \right) \\ &= \sum_{t=1}^T \left(\langle \mathbf{x}_t, \mathbf{w}_t \rangle + \frac{\|\mathbf{w}_t\|^q}{Lq} \right) + \frac{T \left\| -\frac{1}{T} \sum_{t=1}^T \mathbf{x}_t \right\|_*^p}{p} \end{aligned}$$

Where the last step is by definition of convex dual of $\frac{\|\cdot\|^q}{q}$. Now note that since we have a supremum over distribution in (13), and so we can lower bound the value by picking distribution such that $\mathbf{d}_1, \dots, \mathbf{d}_T$ are martingale difference sequences and $\mathbf{d}_t \in \frac{L-1}{L} U(\mathfrak{B}^*)$. Thus we see that

$$\begin{aligned} V(U(\mathfrak{B}), \text{lincvx}_{q,L}) &\geq \sup_{\mathbf{M}} \left\{ \inf_W \mathbb{E} \left[\sum_{t=1}^T \left(\langle \mathbf{d}_t, \mathbf{w}_t \rangle + \frac{\|\mathbf{w}_t\|^q}{Lq} \right) + \frac{T^{1-p} \left\| -\sum_{t=1}^T \mathbf{d}_t \right\|_*^p}{p} \right] \right\} \\ &= \sup_{\mathbf{M}} \left\{ \inf_W \mathbb{E} \left[\sum_{t=1}^T \frac{\|\mathbf{w}_t\|^q}{Lq} \right] + \frac{T^{1-p} \mathbb{E} \left[\left\| -\sum_{t=1}^T \mathbf{d}_t \right\|_*^p \right]}{p} \right\} \\ &= \sup_{\mathbf{M}} \left\{ \frac{T^{1-p} \mathbb{E} \left[\left\| -\sum_{t=1}^T \mathbf{d}_t \right\|_*^p \right]}{p} \right\} \geq \frac{T^{1-p} \left(\sup_{\mathbf{M}} \mathbb{E} \left[\left\| -\sum_{t=1}^T \mathbf{d}_t \right\|_*^p \right] \right)^p}{p} \quad (14) \end{aligned}$$

where the first equality is because \mathbf{w}_t is only dependent on the history and so the conditional expectations over \mathbf{d}_t are 0 and the last step is due to Jensen's inequality. Now note that using Lemma 12 we see that for any $p' > p^*$ we have that

$$V(U(\mathfrak{B}), \text{lincvx}_{q,L}) = \Omega \left(\left(1 - \frac{1}{L}\right) \frac{T^{1-p+\frac{p}{p'}}}{p} \right) \quad (15)$$

(the $1 - 1/L$ term above comes from the fact that the martingale differences come from ball of radius $1 - 1/L$ while Lemma 12 is over unit ball). Note that the above lower bound becomes 0 when $L = 1$ but however in that case it means that the adversary is forced to play 1-Lipschitz, q -uniformly convex function. However since from each q -uniformly convex L -Lipschitz convex function we can

build an equivalent norm with distortion $1/L$, this means that the function the adversary plays can be used to construct the original norm itself. From the construction in [20] it becomes clear that the functions the adversary can play can be merely the norm plus some constant and so the lower bound of 0 is real.

Now we turn to proving the upper bound. Consider the regret of the mirror descent algorithm when we run it using a q^* -uniformly convex function Ψ that is C -Lipschitz on the unit ball. Here, for simplicity, we assume that the supremum is achieved in (3) (otherwise we can pick a Ψ that is q' -uniformly convex for $q' = p'/(p' - 1)$ where $p' = p^*(\mathfrak{B}^*) - 1/\log T$ and pay a constant factor). Note that in the case when $q > q^* + 1$, we have that each $\sigma_t^* = 0$ and so by Theorem 10,

$$\text{Reg}(\mathbf{w}_{1:T}, \ell_{1:T}) \leq 2 \min_{\lambda_1, \dots, \lambda_T} \sum_{t=1}^T \frac{(L+C)^{p^*}}{\left(\sum_{j \leq t} \lambda_j\right)^{p^*-1}} + 2 \sum_{t=1}^T \lambda_t C \leq 2 \min_{\lambda} \frac{(L+C)^{p^*} T^{2-p^*}}{\lambda^{p^*-1}} + 2T\lambda C$$

Using $\lambda = \frac{L+C}{T^{1/q^*}(2C)^{1/p^*}}$ we see that $\text{Reg}(\mathbf{w}_{1:T}, \ell_{1:T}) \leq 8(2C)^{1/q^*} (L+C) T^{1/p^*}$. On the other hand when $q \leq q^* + 1$, using the upper bound in the theorem with $\lambda_t = 0$ for all t we see that since all $q_t = q$ and all $\sigma_t^* = 1$ and $L_t = L$ we find that the regret of the adaptive algorithm is bounded as

$$\text{Reg}(\mathbf{w}_{1:T}, \ell_{1:T}) \leq \sum_{t=1}^T \left(\frac{2(L+C)^p}{t^{p-1}} + \frac{2(L+C)^{p^*}}{t^{p^*-1}} \right) \leq \sum_{t=1}^T \frac{4(L+C)^p}{t^{p-1}} \leq 4(L+C)^p \int_1^T \frac{1}{t^{p-1}} dt$$

Hence we see that for $p < p^*(\mathfrak{B}^*) \leq 2$, $\text{Reg}(\mathbf{w}_{1:T}, \ell_{1:T}) \leq \frac{4(L+C)^p}{2-p} T^{2-p}$. Since the regret of the adaptive algorithm bounds the value of the game, we see that by picking constant $K = \max\{\frac{4(L+C)^p}{2-p}, 8(2C)^{1/q^*} (L+C)\}$ we get the required upper bound. \square

The upper and lower bounds do not match in general and it is an interesting open problem to remove this gap. Note that the upper and lower bounds do match for the two extreme cases $q \rightarrow q^*$ and $q \rightarrow \infty$. When $q \rightarrow q^*$, then both lower and upper bound exponents tend to $2 - p^*(\mathfrak{B}^*)$. On the other hand, when $q \rightarrow \infty$, both exponents tend to $1/p^*(\mathfrak{B}^*)$.

6 Strategy for the Player

In this section, we provide a strategy known as Mirror Descent and is given as Algorithm 1 below which uses uniformly convex function Ψ as internal regularizer and is guaranteed to achieve low regret.

Algorithm 1 Mirror Descent (Parameters : $\eta > 0$, $\Psi : \mathfrak{B} \rightarrow \mathbb{R}$ which is uniformly convex)

for $t = 1$ to T **do**
 Play \mathbf{w}_t and receive ℓ_t
 $\mathbf{w}'_{t+1} \leftarrow \nabla \Psi^*(\nabla \Psi(\mathbf{w}_t) - \eta \boldsymbol{\lambda}_t)$ where $\boldsymbol{\lambda}_t \in \partial \ell_t(\mathbf{w}_t)$
 Update $\mathbf{w}_{t+1} \leftarrow \underset{\mathbf{w} \in \mathcal{W}}{\text{argmin}} \Delta_{\Psi}(\mathbf{w}, \mathbf{w}'_{t+1})$
end for

Example : As a more concrete example for the above strategy, if we consider the case where the player plays from the unit ball of a d dimensional ℓ_q space. In this case $\Psi(\mathbf{w}) = \frac{1}{q'} \|\mathbf{w}\|_q^{q'}$ where $q' = q$ whenever $q > 2$ and $q' = 2$ otherwise. The corresponding dual then is $\Psi^*(\mathbf{x}) = \frac{1}{p'} \|\mathbf{x}\|_p^{p'}$ where $p' = p$ if $q > 2$ and $p' = 2$ otherwise. Correspondingly we get

$$\begin{aligned} \nabla \Psi(\mathbf{w}) &= \|\mathbf{w}\|^{q'-q} (\text{sign}(\mathbf{w}_1)|\mathbf{w}_1|^{q-1}, \dots, \text{sign}(\mathbf{w}_d)|\mathbf{w}_d|^{q-1}) \quad \text{and} \\ \nabla \Psi^*(\mathbf{x}) &= \|\mathbf{x}\|^{q'-q} (\text{sign}(\mathbf{x}_1)|\mathbf{x}_1|^{q-1}, \dots, \text{sign}(\mathbf{x}_d)|\mathbf{x}_d|^{q-1}) \end{aligned}$$

We can use this to update \mathbf{w}'_{t+1} and get a concrete algorithm from the above strategy.

The following proposition gives a regret bound for Mirror Descent.

Proposition 9. *Suppose $\mathcal{W} \subseteq \mathfrak{B}$ is such that $\|\mathbf{w}\| \leq B$. Let MD denote the \mathcal{P} -strategy obtained by running Mirror Descent with a function Ψ that is q -uniformly convex on \mathfrak{B} and C -Lipschitz on \mathcal{W} , and the learning rate $\eta = (BC/T)^{1/p} \cdot (1/L)$. Here, $p = q/(q-1)$ is the dual exponent of q . Then, for all sequences $\ell_{1:T}$ such that ℓ_t is L -Lipschitz on \mathcal{W} , we have,*

$$\text{Reg}(\text{MD}, \ell_{1:T}) = O\left((BC)^{1/q} \cdot L \cdot T^{1/p}\right)$$

Proof. For $\lambda \in \mathfrak{B}^*$, $\mathbf{w} \in \mathfrak{B}$ we denote the pairing $\lambda(\mathbf{w})$ by $\langle \lambda, \mathbf{w} \rangle$ where $\langle \cdot, \cdot \rangle : \mathfrak{B}^* \times \mathfrak{B} \rightarrow \mathbb{R}$. This pairing is bilinear but not symmetric. We will first show that, for any $\mathbf{w} \in \mathcal{W}$,

$$\eta \langle \lambda_t, \mathbf{w}_t - \mathbf{w} \rangle \leq \Delta_\Psi(\mathbf{w}, \mathbf{w}_t) - \Delta_\Psi(\mathbf{w}, \mathbf{w}_{t+1}) + \frac{\eta^p}{p} \|\lambda_t\|_*^p, \quad (16)$$

where $p = q/(q-1)$. We have,

$$\begin{aligned} \eta \langle \lambda_t, \mathbf{w}_t - \mathbf{w} \rangle &= \langle \eta \lambda_t, \mathbf{w}_t - \mathbf{w}_{t+1} + \mathbf{w}_{t+1} - \mathbf{w} \rangle \\ &= \underbrace{\langle \eta \lambda_t, \mathbf{w}_t - \mathbf{w}_{t+1} \rangle}_{s_1} + \underbrace{\langle \eta \lambda_t + \nabla \Psi(\mathbf{w}_{t+1}) - \nabla \Psi(\mathbf{w}_t), \mathbf{w}_{t+1} - \mathbf{w} \rangle}_{s_2} \\ &\quad + \underbrace{\langle \nabla \Psi(\mathbf{w}_t) - \nabla \Psi(\mathbf{w}_{t+1}), \mathbf{w}_{t+1} - \mathbf{w} \rangle}_{s_3} \end{aligned} \quad (17)$$

Now, by definition of the dual norm and the fact that $ab \leq \frac{a^p}{p} + \frac{b^q}{q}$ for any $a, b \geq 0$, we get

$$s_1 \leq \|(\mathbf{w}_t - \mathbf{w}_{t+1})\| \cdot \|\eta \lambda\|_* \leq \frac{1}{q} \|\mathbf{w}_t - \mathbf{w}_{t+1}\|^q + \frac{1}{p} \|\eta \lambda_t\|_*^p.$$

By the definition of the update, \mathbf{w}_{t+1} minimizes $\langle \eta \lambda_t - \nabla \Psi(\mathbf{w}_t), \mathbf{w} \rangle + \Psi(\mathbf{w})$ over $\mathbf{w} \in \mathcal{W}$. Therefore, $s_2 \leq 0$. Using simple algebraic manipulations, we get

$$s_3 = \Delta_\Psi(\mathbf{w}, \mathbf{w}_t) - \Delta_\Psi(\mathbf{w}, \mathbf{w}_{t+1}) - \Delta_\Psi(\mathbf{w}_{t+1}, \mathbf{w}_t).$$

Plugging this into (17), we get

$$\eta \langle \lambda_t, \mathbf{w}_t - \mathbf{w} \rangle \leq \Delta_\Psi(\mathbf{w}, \mathbf{w}_t) - \Delta_\Psi(\mathbf{w}, \mathbf{w}_{t+1}) + \frac{\eta^p}{p} \|\lambda_t\|_*^2 + \underbrace{\frac{1}{q} \|\mathbf{w}_t - \mathbf{w}_{t+1}\|^q - \Delta_\Psi(\mathbf{w}_{t+1}, \mathbf{w}_t)}_{s_4}$$

Using that Ψ is q -uniformly convex on \mathfrak{B} implies that $s_4 \leq 0$. So, we get (16).

We can now bound the regret as follows. For any $\mathbf{w} \in \mathcal{W}$, since $\lambda_t \in \partial \ell_t(\mathbf{w}_t)$, we have, $\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{w}) \leq \langle \lambda_t, \mathbf{w}_t - \mathbf{w} \rangle$. Combining this with (16) and summing over t gives,

$$\sum_{t=1}^T (\ell_t(\mathbf{w}_t) - \ell_t(\mathbf{w})) \leq \frac{\Delta_\Psi(\mathbf{w}, \mathbf{w}_1) - \Delta_\Psi(\mathbf{w}, \mathbf{w}_{T+1})}{\eta} + \frac{\eta^{p-1}}{p} \sum_{t=1}^T \|\lambda_t\|_*^p.$$

Now, $\Delta_\Psi(\mathbf{w}, \mathbf{w}_{T+1}) \geq 0$ and $\Delta_\Psi(\mathbf{w}, \mathbf{w}_1) \leq 2BC$. Further $\|\lambda_t\|_* \leq L$ since ℓ_t is L -Lipschitz. Plugging these above and optimizing over η gives the required upper bound. \square

Note that all the above algorithm needed to achieve low regret was a uniformly convex function Ψ . Theorem (Pisier) [16] gives us exactly this, it guarantees existence of a $\frac{p}{p-1}$ -uniformly convex function on a given Banach space for any $p \in [1, p^*(\mathfrak{B}^*))$ thus making sure that the above mirror descent algorithm with this choice of Ψ gives us the optimal rate for convex lipschitz games.

7 Adaptive Player Strategy

A natural extension of the q -uniformly convex lipschitz game is a game where at round t , \mathcal{A} plays q_t uniformly convex functions. In this section, we give an adaptive player strategy for such games that achieves the upper bound in Theorem 8 whenever the adversary plays only q -uniformly convex functions on all rounds and in general gets intermediate rates when the modulus of convexity on each round is different.

Now for the sake of readability, we assume that the supremum in (3) is achieved. The following theorem states that the same adaptive algorithm achieves the upper bound suggested in Theorem 8 for various q -uniformly convex games. Further the algorithm adjusts itself to the scenario when \mathcal{A} plays a different (σ_t, q_t) -uniformly convex function at each round t . To see this let, $\sigma_j^* = \sigma_j \mathbb{1}_{\{q_j < q^* + 1\}}$. In the above algorithm we set at each round λ_t that satisfies,

$$2C\lambda_t = \sum_{i \leq t} \frac{\frac{\sigma_i^*}{M_i^{q_i}}}{\left(\sum_{j \leq t} \left[\frac{\sigma_j^*}{M_j^{q_j}} + \frac{\lambda_j}{M_j^{q^*}} \right] \right)^{p_i}} + \frac{1}{\left(\sum_{j \leq t} \left[\frac{\sigma_j^*}{M_j^{q_j}} + \frac{\lambda_j}{M_j^{q^*}} \right] \right)^{p^* - 1}}, \quad (18)$$

where $M_t = L_t + C$. The following theorem which upper bounds the regret of the adaptive algorithm.

Algorithm 2 Adaptive Mirror Descent (Parameters : $\Psi : \mathfrak{B} \rightarrow \mathbb{R}$ which is q^* -uniformly convex)

$C \leftarrow$ Lipschitz constant of Ψ on $U(\mathfrak{B})$, $\mathbf{w}_1 \leftarrow \mathbf{0}$, $\Phi_1 \leftarrow \mathbf{0}$

for $t = 1$ to T **do**

Play \mathbf{w}_t and receive ℓ_t which is L_t -Lipschitz and (σ_t, q_t) -uniformly convex

Pick λ_t that satisfies (18)

$\Phi_{t+1} \leftarrow \Phi_t + \ell_t + \lambda_t \Psi$

$\mathbf{w}'_{t+1} \leftarrow \nabla \Phi_{t+1}^* (\nabla \Phi_t(\mathbf{w}_t))$

Update $\mathbf{w}_{t+1} \leftarrow \operatorname{argmin}_{\mathbf{w} \in \mathcal{W}} \Delta_{\Phi_{t+1}}(\mathbf{w}, \mathbf{w}'_{t+1})$

end for

Theorem 10. Let $\mathcal{W} = U(\mathfrak{B})$. Let AMD denote the \mathcal{P} -strategy obtained by running Adaptive Mirror Descent with a Ψ which is $q^* = \frac{p^*(\mathfrak{B}^*)}{p^*(\mathfrak{B}^*)-1}$ uniformly convex. Then, for all sequences $\ell_{1:T}$ such that ℓ_t is L_t -Lipschitz and (σ_t, q_t) -uniformly convex, we have,

$$\operatorname{Reg}(\text{AMD}, \ell_{1:T}) \leq \min_{\lambda_{1:T}} \sum_{t=1}^T \left\{ \left(\sum_{i \leq t} \frac{\frac{2\sigma_i^*}{M_t^{q_i}}}{\left(\sum_{j \leq t} \left[\frac{\sigma_j^*}{M_t^{q_j}} + \frac{\lambda_j}{M_t^{q^*}} \right] \right)^{p_i}} \right) + \frac{2}{\left(\sum_{j \leq t} \left[\frac{\sigma_j^*}{M_t^{q_j}} + \frac{\lambda_j}{M_t^{q^*}} \right] \right)^{p^*-1}} + 2\lambda_t C \right\}$$

Proof. Note that $f_t = \ell_t + \lambda_t \Psi$, Ψ is q^* -uniformly convex and ℓ_t is (σ_t, q_t) -uniformly convex. Hence,

$$\Delta_{f_t}(\mathbf{w}'_{t+1}, \mathbf{w}_t) \geq \frac{\sigma_t^*}{q_t} \|\mathbf{w}'_{t+1} - \mathbf{w}_t\|^{q_t} + \frac{\lambda_t}{q^*} \|\mathbf{w}'_{t+1} - \mathbf{w}_t\|^{q^*}$$

Where $\sigma_t^* = \sigma_t \mathbb{I}_{\{q_t < q^*+1\}}$. Now since $\Phi_{t+1} = \sum_{i \leq t} f_t$, we see that

$$\begin{aligned} \Delta_{\Phi_{t+1}}(\mathbf{w}_t, \mathbf{w}'_{t+1}) &= \langle \nabla \Phi_{t+1}(\mathbf{w}_t) - \nabla \Phi_{t+1}(\mathbf{w}'_{t+1}), \mathbf{w}_t - \mathbf{w}'_{t+1} \rangle - \Delta_{\Phi_{t+1}}(\mathbf{w}'_{t+1}, \mathbf{w}_t) \\ &= \langle \nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}'_{t+1} \rangle - \Delta_{\Phi_{t+1}}(\mathbf{w}'_{t+1}, \mathbf{w}_t) \\ &\leq \langle \nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}'_{t+1} \rangle - \sum_{i=1}^t \frac{\sigma_i^*}{q_i} \|\mathbf{w}'_{i+1} - \mathbf{w}_i\|^{q_i} - \sum_{i=1}^t \frac{\lambda_i}{q^*} \|\mathbf{w}'_{i+1} - \mathbf{w}_i\|^{q^*} \end{aligned}$$

Now consider any arbitrary sequence $\beta_1, \dots, \beta_{2t}$ of non-negative numbers such that $\sum_{i=1}^{2t} \beta_i = 1$. In this case note that by Fenchel-Young inequality,

$$\begin{aligned} \Delta_{\Phi_{t+1}}(\mathbf{w}_t, \mathbf{w}'_{t+1}) &\leq \sum_{i=1}^{2t} \langle \beta_i \nabla f_t(\mathbf{w}_t), \mathbf{w}_t - \mathbf{w}'_{t+1} \rangle - \sum_{i=1}^t \frac{\sigma_i^*}{q_i} \|\mathbf{w}'_{i+1} - \mathbf{w}_i\|^{q_i} - \sum_{i=1}^t \frac{\lambda_i}{q^*} \|\mathbf{w}'_{i+1} - \mathbf{w}_i\|^{q^*} \\ &\leq \sum_{i=1}^t \left(\frac{\beta_i^{p_i} \|\nabla f_t(\mathbf{w}_t)\|_*^{p_i}}{p_i (\sigma_i^*)^{p_i/q_i}} + \frac{\beta_i^{p^*} \|\nabla f_t(\mathbf{w}_t)\|_*^{p^*}}{p^* \lambda_i^{p^*/q^*}} \right) \leq \sum_{i=1}^t \left(\frac{\beta_i^{p_i} (L_t + C)^{p_i}}{p_i (\sigma_i^*)^{p_i/q_i}} + \frac{\beta_i^{p^*} (L_t + C)^{p^*}}{p^* \lambda_i^{p^*/q^*}} \right) \end{aligned}$$

In the above we used the fact that since ℓ_t is L_t -Lipschitz and Ψ is C -Lipschitz so that, $\|\nabla f_t(\mathbf{w}_t)\|_* \leq (L_t + C)$ (we were able to get rid of the λ_t because we use $\lambda_t \leq 1$ and so $L_t + \lambda_t C \leq L_t + C$). Now choosing $\forall i \leq t$, $\beta_i \propto \frac{\sigma_i^*}{(L_t + C)^{q_i}}$ and $\forall t < i \leq 2t$, $\beta_i \propto \frac{\lambda_i}{(L_t + C)^{q^*}}$ we see that

$$\begin{aligned} \Delta_{\Phi_{t+1}}(\mathbf{w}_t, \mathbf{w}'_{t+1}) &\leq \sum_{i=1}^t \left(\frac{\frac{\sigma_i^*}{(L_t + C)^{q_i}}}{p_i \left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t + C)^{q_j}} + \frac{\lambda_j}{(L_t + C)^{q^*}} \right) \right)^{p_i}} + \frac{\frac{\lambda_i}{(L_t + C)^{q^*}}}{p^* \left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t + C)^{q_j}} + \frac{\lambda_j}{(L_t + C)^{q^*}} \right) \right)^{p^*}} \right) \\ &\leq \sum_{i=1}^t \left(\frac{\frac{\sigma_i^*}{(L_t + C)^{q_i}}}{\left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t + C)^{q_j}} + \frac{\lambda_j}{(L_t + C)^{q^*}} \right) \right)^{p_i}} + \frac{\frac{\lambda_i}{(L_t + C)^{q^*}} + \frac{\sigma_i^*}{(L_t + C)^{q_i}}}{\left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t + C)^{q_j}} + \frac{\lambda_j}{(L_t + C)^{q^*}} \right) \right)^{p^*}} \right) \\ &= \left(\sum_{i=1}^t \frac{\frac{\sigma_i^*}{(L_t + C)^{q_i}}}{\left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t + C)^{q_j}} + \frac{\lambda_j}{(L_t + C)^{q^*}} \right) \right)^{p_i}} \right) + \frac{1}{\left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t + C)^{q_j}} + \frac{\lambda_j}{(L_t + C)^{q^*}} \right) \right)^{p^*-1}} \end{aligned}$$

where in the first step we used the fact that $p^*, p_i \geq 1$ to remove them from the denominator. Thus using Lemma 13 we conclude that

$$\operatorname{Reg}(\mathbf{w}_{1:T}, \ell_{1:T}) \leq \sum_{t=1}^T \left(\sum_{i=1}^t \frac{\frac{\sigma_i^*}{(L_t + C)^{q_i}}}{\left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t + C)^{q_j}} + \frac{\lambda_j}{(L_t + C)^{q^*}} \right) \right)^{p_i}} + \frac{1}{\left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t + C)^{q_j}} + \frac{\lambda_j}{(L_t + C)^{q^*}} \right) \right)^{p^*-1}} + 2C\lambda_t \right)$$

Since we choose λ_t 's that satisfy Equation 18, using Lemma 14 we get the required statement. \square

Using the above regret bound we get the following corollary showing that the Adaptive Mirror Descent algorithm can be used to achieve all upper bounds on the regret presented in the paper.

Corollary 11. *There exists a Ψ which is q^* -uniformly convex function, and using this function with the Adaptive Mirror Descent (AMD) algorithm, we have the following.*

1. *Regret of AMD for convex-Lipschitz game matches upper bound in (8).*
2. *Regret of AMD for q -uniformly convex game matches upper bound in (12).*
3. *For the bounded convex game, there exists a $C > 0$ such that using AMD on $1 - CT^{-\frac{1}{2q^*}}$ ball achieves the upper bound in (10) for the game played on the unit ball.*

Proof. Claim 2 is shown in the constructive proof of the upper bound of Theorem 8. As for claim 1, note that this is the case of linear functions and so it is the same as adversary picking each $\sigma_t = 0$. Regret in this case again can be found in the proof of the upper bound of Theorem 8 and so claim 1 also holds. As for the last claim, given claim 1, it is evident from proof of Theorem 5. \square

When $q_1, \dots, q_T = 2$, AMD enjoys the same guarantee as Algorithm 4 in [7] (see Theorem 4.2).

8 Discussion

In future work, we also plan to convert the player strategies given here into implementable algorithms. Online learning algorithms can be implemented in infinite dimensional reproducing kernel Hilbert spaces [21] by exploiting the representer theorem and duality. We can, therefore, hope to implement online learning algorithms in infinite dimensional Banach spaces where some analogue of the representer theorem is available. Der and Lee [12] have made progress in this direction using the notion of *semi-inner products*. For $L_q(\Omega, \mu)$ spaces with q even, they showed how the problem of finding a maximum margin linear classifier can be reduced to a finite dimensional convex program using “moment functions”. The types (and their associated constants) of L_q spaces are well known from classical Banach space theory. So, we can use their ideas to get online algorithms in these spaces with provable regret guarantees. Vovk [4] also defines “Banach kernels” for certain Banach spaces of real valued functions and gives an implementable algorithm assuming the Banach kernel is efficiently computable. His interest is in prediction with the squared loss. It will be interesting to explore the connection of his ideas with the setting of this paper.

Using online-to-batch conversions, our results also imply error bounds for the estimation error in the batch setting. If $p^*(\mathfrak{B}^*) < 2$ then we get a rate worse than $O(T^{-1/2})$. However, we get the ability to work with richer function classes. This can decrease the approximation error. The study of this trade-off can be helpful.

We would also like to improve our lower and/or upper bounds where they do not match. In this regard, we should mention that the upper bound for the convex-bounded game given in Theorem 5 is not tight for a Hilbert space. Our upper bound is $O(T^{3/4})$ but it can be shown that using the self-concordant barrier $\log(1 - \|\mathbf{w}\|^2)$ for the unit ball, we get an upper bound of $O(T^{2/3})$.

Acknowledgments We thank the Colt 2010 reviewers for their helpful comments. We would also like to thank Maxim Raginsky for pointing out certain subtle issues in our earlier version.

References

- [1] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on Machine Learning*, 2003.
- [2] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2–3):169–192, 2007.
- [3] S. Shalev-Shwartz. *Online Learning: Theory, Algorithms, and Applications*. PhD thesis, Hebrew University of Jerusalem, 2007.
- [4] V. Vovk. Competing with wild prediction rules. *Machine Learning*, 69(2–3):193–212, 2007.
- [5] J. Abernethy, P. L. Bartlett, A. Rakhlin, and A. Tewari. Optimal strategies and minimax lower bounds for online convex games. In *Proceedings of the Nineteenth Annual Conference on Computational Learning Theory*, 2008.
- [6] A. Nemirovski and D. Yudin. *Problem complexity and method efficiency in optimization*. Nauka Publishers, Moscow, 1978.
- [7] P. L. Bartlett, E. Hazan, and A. Rakhlin. Adaptive online gradient descent. In *Advances in Neural Information Processing Systems 21*, 2007.

- [8] J. Abernethy, A. Agarwal, P. L. Bartlett, and A. Rakhlin. A stochastic view of optimal regret through minimax duality. In *Proceedings of the Twentieth Annual Conference on Computational Learning Theory*, 2009.
- [9] M. J. Donahue, L. Gurvits, C. Darken, and E. Sontag. Rates of convex approximation in non-Hilbert spaces. *Constructive Approximation*, 13:187–220, 1997.
- [10] L. Gurvits. A note on a scale-sensitive dimension of linear bounded functionals in Banach spaces. *Theoretical Computer Science*, 261(1):81–90, 2001.
- [11] S. Mendelson and G. Schechtman. The shattering dimension of sets of linear functionals. *The Annals of Probability*, 32(3):1746–1770, 2004.
- [12] R. Der and D. Lee. Large-margin classification in Banach spaces. In *Processing of Eleventh International Conference on Artificial Intelligence and Statistics*, 2007.
- [13] Haizhang Zhang, Yuesheng Xu, and Jun Zhang. Reproducing kernel Banach spaces for machine learning. *Journal of Machine Learning Research*, 10:2741–2775, 2009.
- [14] Fedor Zhdanov, Alexey Chernov, and Yuri Kalnishkan. Aggregating algorithm competing with Banach lattices, 2010. arXiv preprint [arXiv:1002.0709](https://arxiv.org/abs/1002.0709) available at <http://arxiv.org/abs/1002.0709>.
- [15] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [16] G. Pisier. Probabilistic methods in the geometry of Banach spaces. In G. Letta and M. Pratelli, editors, *Probability and Analysis*, volume 1206 of *Lecture Notes in Mathematics*, pages 167–241. Springer, 1986.
- [17] G. Pisier. Martingales with values in uniformly convex spaces. *Israel Journal of Mathematics*, 20(3–4):326–350, 1975.
- [18] C. Zălinescu. *Convex analysis in general vector spaces*. World Scientific Publishing Co. Inc., River Edge, NJ, 2002.
- [19] G. Pisier. Sur les espaces de Banach qui ne contiennent pas uniformément de ℓ_n^1 . *C. R. Acad. Sci. Paris Sér. A-B*, 277:A991–A994, 1973.
- [20] J. Borwein, A. J. Guirao, P. Hájek, and J. Vanderwerff. Uniformly convex functions on Banach spaces. *Proceedings of the American Mathematical Society*, 137(3):1081–1091, 2009.
- [21] J. Kivinen, A. J. Smola, and R. C. Williamson. Online learning with kernels. *IEEE Transactions on Signal Processing*, 52(8):2165–2176, 2002.
- [22] K. Sridharan and A. Tewari. Convex games in banach spaces. *Technical Report no. TTIC-TR-2009-6*, Toyota Technological Institute at Chicago, 2009.

Appendix

Lemma 12. For any Banach space \mathfrak{B}^* and any $p' > p^*(\mathfrak{B}^*)$ we have that

$$\frac{\sup_{\mathbf{M}} \mathbb{E} \left[\left\| \sum_{t=1}^T \mathbf{d}_t \right\| \right]}{T^{1/p'}} \rightarrow \infty$$

as $T \rightarrow \infty$, where \mathbf{M} refers to distributions over martingale difference sequences $(\mathbf{d}_t)_{t=1}^T$ such that each $\mathbf{d}_t \in U(\mathfrak{B}^*)$

Lemma 13. For the Adaptive Mirror Descent Algorithm we have that

$$\text{Reg}(\mathbf{w}_{1:T}, \ell_{1:T}) \leq \sum_{t=1}^T (\Delta_{\Phi_{t+1}}(\mathbf{w}_t, \mathbf{w}'_{t+1}) + 2\lambda_t C)$$

Lemma 14. Define for any sequence $\lambda_1, \dots, \lambda_S$ of any size S ,

$$O_S(\lambda_1, \dots, \lambda_S) = \sum_{t=1}^S \left(\sum_{i=1}^t \left(\frac{\frac{\sigma_i^*}{(L_t+C)^{q_i}}}{\left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t+C)^{q_j}} + \frac{\lambda_j}{(L_t+C)^{q^*}} \right) \right)^{p_i}} \right) + \frac{1}{\left(\sum_{j=1}^t \left(\frac{\sigma_j^*}{(L_t+C)^{q_j}} + \frac{\lambda_j}{(L_t+C)^{q^*}} \right) \right)^{p^*-1}} + 2C\lambda_t \right)$$

Then as long as we pick λ_t that satisfies Equation 18, we have that for any T

$$O_T(\lambda_1, \dots, \lambda_T) \leq 2 \min_{\lambda_1, \dots, \lambda_T} O\{\lambda_1, \dots, \lambda_T\}$$

For proofs of the above two lemma's refer to a longer version of this paper at [22].