

On the Complexity of Bandit Linear Optimization

Ohad Shamir

Weizmann Institute of Science, Israel

OHAD.SHAMIR@WEIZMANN.AC.IL

Abstract

We study the attainable regret for online linear optimization problems with bandit feedback, where unlike the full-information setting, the player can only observe its own loss rather than the full loss vector. We show that the price of bandit information in this setting can be as large as d , disproving the well-known conjecture (Dani et al. (2007)) that the regret for bandit linear optimization is at most \sqrt{d} times the full-information regret. Surprisingly, this is shown using “trivial” modifications of standard domains, which have no effect in the full-information setting. This and other results we present highlight some interesting differences between full-information and bandit learning, which were not considered in previous literature.

1. Introduction

We consider the problem of bandit linear optimization, which is a repeated game between a player and an adversary. At each round $t = 1, \dots, T$ the player chooses a point \mathbf{w}_t from a compact subset \mathcal{W} of \mathbb{R}^d , and simultaneously the adversary chooses a loss vector $\mathbf{x}_t \in \mathbb{R}^d$ (under some regularity assumptions described in Sec. 2). The player incurs a loss $\langle \mathbf{x}_t, \mathbf{w}_t \rangle$, and can observe the loss but not the loss vector \mathbf{x}_t . The player’s goal is to minimize its expected cumulative loss, $\mathbb{E} \left[\sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w}_t \rangle \right]$, where the expectation is with respect to the player’s and adversary’s possible randomization. The performance of the player is measured in terms of expected regret, defined as

$$\mathbb{E} \left[\sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w}_t \rangle \right] - \min_{\mathbf{w} \in \mathcal{W}} \mathbb{E} \left[\sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w} \rangle \right]. \quad (1)$$

Bandit optimization has proven to be a useful abstraction of sequential decision-making problems under uncertainty, such as multi-armed bandits and online routing (see Bubeck and Cesa-Bianchi (2012) for a survey). Moreover, using online-to-batch conversion techniques, algorithms for this setting can be readily applied to derivative-free stochastic optimization, where our goal is to stochastically optimize an unknown function given only noisy views of its values at various points.

The attainable regret in the bandit setting can be compared to the attainable regret in the full-information setting, where \mathbf{x}_t is revealed to the player after each round. Clearly, since the player receives less information in the bandit setting, the attainable regret will be larger. This degradation is known as the “price of bandit information” (Dani et al. (2007)), and characterizing it for general domains \mathcal{W} has remained an open problem. However, the standard conjecture and common wisdom (as articulated in Dani et al. (2007)) is that for linear optimization, this price is at most a multiplicative \sqrt{d} factor, where d is the dimension. Indeed, as far as we know, this holds for all domains that have been previously studied in the literature:

- When the domain \mathcal{W} is the corners of the d -dimensional simplex (a.k.a. multi-armed bandits setting), the minimax optimal regret is $\Theta(\sqrt{dT})$, vs. $\Theta(\sqrt{\log(d)T})$ in the full-information setting.
- When the domain \mathcal{W} is the boolean hypercube $\{-1, +1\}^d$, the minimax optimal regret is $\Theta(d\sqrt{T})$, vs. $\Theta(\sqrt{dT})$ in the full-information setting (Dani et al. (2007); Audibert et al. (2011)).
- When the domain \mathcal{W} is the unit Euclidean ball, there is an algorithm with regret $\mathcal{O}(\sqrt{dT})$, vs. $\mathcal{O}(\sqrt{T})$ in the full-information setting (Bubeck et al. (2012)).
- There is an $\Omega(d\sqrt{T})$ lower bound for a certain non-convex subset of the hypercube (a cartesian product of d 2-dimensional spheres) (Dani et al. (2008)). The corresponding full-information regret bound is $\mathcal{O}(\sqrt{dT})$.

We note in passing that there are other partial-information settings where the situation is different, but these are distinct from the bandit linear optimization setting we focus here (e.g. non-linear bandit optimization (Shamir (2013)) or using different information feedback, e.g. Cesa-Bianchi and Lugosi (2012)).

The main contribution of this paper is disproving this conjecture, and showing that the price of bandit information for online linear optimization can be as large as d rather than \sqrt{d} . We do this by proving that the $\mathcal{O}(\sqrt{dT})$ regret upper bound for the Euclidean ball is surprisingly brittle, and the attainable regret becomes $\Omega(d\sqrt{T})$ after “trivial” changes of the domain which do not matter at all in the full-information setting. These changes include (1) Shifting the ball away from the origin, and (2) Taking a simple convex subset of the Euclidean ball with a “flat” boundary, such as a cylinder or a capped Euclidean ball¹. We also explain how our techniques can be potentially applicable to other domains. Since the full-information regret in these cases is $\mathcal{O}(\sqrt{T})$, this establishes a price of bandit information on the order of d . This gap is tight as worst-case over all domains (or all convex domains), because for any domain, it is possible to get $\mathcal{O}(d\sqrt{T})$ regret (Bubeck et al. (2012); Hazan et al. (2014)), and the attainable regret in the full-information case is generally at least $\Omega(\sqrt{T})$.

We note that our lower bounds hold even against a stochastic adversary, which chooses loss vectors i.i.d. from a given distribution. In such a setting, any algorithm attaining $\mathcal{O}(\sqrt{d^\alpha T})$ regret can be used to find an ϵ -optimal point after $\mathcal{O}(d^\alpha/\epsilon^2)$ rounds (as described later). This allows us to re-phrase our result in the following manner, which might be conceptually interesting: If the price of bandit information was \sqrt{d} , then to find an ϵ -optimal point, we would need d times more rounds in the bandit setting, compared to the full-information setting (e.g. d/ϵ^2 vs. $1/\epsilon^2$). This is intuitively very appealing: Each round, we get to see only “ $1/d$ as much information” in the bandit setting (a single scalar compared to a d -dimensional vector), hence we need d times more rounds to get the same amount of information and return a point of similar accuracy. Unfortunately, our results show that there are cases where the price of viewing scalars vs. d -dimensional vectors is *quadratic* in d (e.g. d^2/ϵ^2 vs. $1/\epsilon^2$). So in some sense, the number of rounds required and the amount of information per round cannot be traded-off without significant loss.

1. We note that there are (inefficient) algorithmic techniques whose performance is not affected by such domain changes, e.g. discretizing the domain and running a standard multi-armed bandit algorithm on the discretized set, in which case the regret depends on the size of the set rather than its geometry. However, the dependence on the dimension d with such a technique is exponential, hence highly suboptimal to begin with.

A second contribution of our paper lies in disproving another common intuition: Namely, that any regret lower bounds with respect to a given domain \mathcal{W} automatically extend to its convex hull $\text{conv}(\mathcal{W})$. For example, this has been implicitly used to argue that the $\Omega(d\sqrt{T})$ regret lower bound for the boolean hypercube $\{-1, +1\}^d$ in Audibert et al. (2011) extends to the convex hypercube $[-1, +1]^d$ (Bubeck and Cesa-Bianchi (2012); Hazan et al. (2014)). Again, this is generally true for online linear optimization in the full-information setting, since the optimal points in $\text{conv}(\mathcal{W})$ lie in \mathcal{W} , and the player can always simulate playing over $\text{conv}(\mathcal{W})$ via randomization over \mathcal{W} . However, in the bandit setting the change in domain also changes the information feedback structure in non-intuitive ways. For example, if the loss vectors are binary and \mathcal{W} is the corners of the simplex (a.k.a. multi-armed bandits), there is a well-known $\Omega(\sqrt{dT})$ regret lower bound (Auer et al. (2002)). However, when we convexify the domain and take \mathcal{W} to be the simplex, then it is possible to get $\mathcal{O}(\sqrt{\log(d)T})$ regret - same as in the full-information setting! In fact, such a result can be shown to hold for most continuous domains. We emphasize that these regret upper bounds are achieved only for discrete losses, and in a manner which is of little practical value. However, they do demonstrate that one has to be careful when extending bandit optimization results from a finite domain to its convex hull, and that loss vectors supported on a finite set are not appropriate to prove bandit lower bounds for continuous domains. In any case, for completeness, we provide minimax optimal regret bounds for the simplex and the hypercube, using the techniques we develop here. For the hypercube, this bound is $\Theta(d\sqrt{T})$, and formally establishes that this is the minimax optimal regret for bandit linear optimization over convex domains.

1.1. Lower Bound Techniques and Main Ideas

In this subsection, we informally sketch the technical ideas used to get our lower bounds. The various domains we discuss below are sketched in Figure 1. The result on continuous domains and discrete losses is a bit orthogonal, and described separately in Sec. 5.

As mentioned earlier, our lower bounds are actually proven for the easier setting of *stochastic* bandit linear optimization, where the adversary is constrained to sample each loss vector \mathbf{x}_t i.i.d. from the same distribution. Moreover, some of the bounds also apply to the easier goal of minimizing error rather than regret, which may be of independent interest: In this case, the player may choose $\mathbf{w}_1, \dots, \mathbf{w}_T$ arbitrarily, observing $\langle \mathbf{x}_1, \mathbf{w}_1 \rangle, \dots, \langle \mathbf{x}_T, \mathbf{w}_T \rangle$, and then needs to return a vector $\hat{\mathbf{w}}$ which minimizes the expected error

$$\mathbb{E} [\langle \bar{\mathbf{x}}, \hat{\mathbf{w}} \rangle] - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle. \quad (2)$$

where $\bar{\mathbf{x}} = \mathbb{E}[\mathbf{x}]$. This goal is more relevant in a stochastic optimization setting, where we attempt to optimize a stochastic linear function based on querying its values $\langle \mathbf{x}_t, \mathbf{w}_t \rangle$ at various points. It is easy to show that any algorithm, attaining expected regret of R after T rounds, can attain an R/T error by returning the vector $\bar{\mathbf{w}} = \frac{1}{T} \sum_{t=1}^T \mathbf{w}_t$. Thus, any lower bound for minimizing error in a stochastic setting implies the same lower bound on minimizing regret and a possible non-stochastic setting, times a T factor. For example, a $\sqrt{d/T}$ error lower bound implies a \sqrt{dT} regret lower bound.

For the goal of minimizing error in a stochastic setting, it turns out that there is a simple strategy attaining $\sqrt{d/T}$ error under appropriate conditions: In each round, we randomly sample from some hypercube centered around the origin, and use it to create an unbiased estimate of the loss vector, with variance $\mathcal{O}(d)$. Repeating this for T rounds and averaging, we get an estimate $\tilde{\mathbf{x}}$ of




Figure 1: Various domains considered in the paper: Origin-centered Euclidean unit ball; shifted Euclidean unit ball; cylinder; and capped Euclidean ball. For a domain such as the unit ball, and for minimizing error, we can get a bounded-variance estimate of the expected loss vector by querying at random from an origin-centered hypercube. However, this is not possible in the other settings considered here, either because the domain is bounded away from the origin, or because the player needs to query far and to one side of the origin to attain small regret.

the expected loss vector $\bar{\mathbf{x}}$, which is unbiased and with variance $\mathcal{O}(d/T)$. We then return the point $\arg \min_{\mathbf{w} \in \mathcal{W}} \langle \tilde{\mathbf{x}}, \mathbf{w} \rangle$, leading to an $\mathcal{O}(\sqrt{d/T})$ error bound.

Unfortunately, this method breaks down if we cannot sample from such an origin-centered hypercube. In particular, if the variance of $\langle \mathbf{w}, \mathbf{x} \rangle$ for any \mathbf{w} in the domain is lower bounded by a constant (independent of the dimension), then the variance of the estimator becomes $\mathcal{O}(d^2)$ rather than $\mathcal{O}(d)$. Clearly, this is not true when \mathbf{w} can be arbitrarily close to the origin: If $\mathbf{w} = \mathbf{0}$, then the variance of $\langle \mathbf{w}, \mathbf{x} \rangle$ is always zero. However, it can happen when the domain is bounded away from the origin. Using this key observation together with canonical information-theoretic tools (i.e. reduction to a hypothesis testing problem), we can show a $\Omega(d/\sqrt{T})$ error lower bound for a simple domain bounded away from the origin, such as a shifted Euclidean ball. Due to the relationship between error and regret, this leads to an $\Omega(d\sqrt{T})$ regret lower bound for such domains.

Another way to force $\langle \mathbf{w}, \mathbf{x} \rangle$ to have constant variance, even if the domain contains the origin, is when the player must pick points far away from the origin most of the time. This cannot be enforced when the goal is minimizing error. However, when the goal is minimizing regret, then we can construct a situation where any point close to the origin leads to a large regret. In that case, either the player picks points close to the origin, and gets large regret, or picks points far from the origin, leading to the variance of $\langle \mathbf{w}_t, \mathbf{x}_t \rangle$ being large. This is a lose-lose situation, and carefully formalizing it leads to a $\Omega(d\sqrt{T})$ regret lower bound, similar to the shifted Euclidean ball setting. We formally show this for a cylindrical domain (see figure 1) as well as a hypercube, but the proof

technique appears applicable to other domains, such as a capped Euclidean ball, or more generally any domain with a flat surface orthogonal to the origin. In all these cases, the adversary strategy is to choose an expected loss vector at random from the flat surface, and introduce a constant amount of stochastic noise in the orthogonal dimension. Since the flat surface is far from the origin, the player must deal with constant variance if it wishes its regret to be small.

Although our lower bounds are shown for particular domains, we believe that the properties we identified – distance from origin, variance of the losses, and flatness of parts of the domain boundary – can play a key role in characterizing the attainable regret for any given domain.

2. Preliminaries

We use bold-faced letters to denote vectors (e.g. $\mathbf{x} = (x_1, \dots, x_d)$). We let $\|\cdot\|$ denote the Euclidean norm, and $\|\cdot\|_p$ the p -norm. We also define the function $\|\cdot\|_* : \mathbb{R}^d \mapsto [0, \infty)$ as

$$\|\mathbf{x}\|_* = \max_{\mathbf{w} \in \mathcal{W}} |\langle \mathbf{w}, \mathbf{x} \rangle|.$$

When \mathcal{W} is a symmetric convex set, then $\|\cdot\|_*$ is known as the support function of \mathcal{W} (Rockafellar (1970)), or equivalently, the dual norm to the norm whose unit ball is \mathcal{W} , hence the notation. However, we will use this notation even when \mathcal{W} is not necessarily convex or symmetric. We also define $\mathcal{N}(\mathbf{a}, \Sigma)$ to be the multivariate Gaussian distribution with mean \mathbf{a} and covariance matrix Σ , and let I denote the identity matrix.

As discussed in the introduction, our results focus on bandit linear optimization in a stochastic setting, where the loss vectors are assumed to be drawn i.i.d. from an unknown distribution \mathcal{D} with mean $\bar{\mathbf{x}}$. In this setting, the expected regret (as defined in Eq. (1)) can be equivalently written as

$$\mathbb{E} \left[\left\langle \bar{\mathbf{x}}, \sum_{t=1}^T \mathbf{w}_t \right\rangle \right] - T \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle. \quad (3)$$

Clearly, if any distribution is allowed, then the inner products above can be unboundedly large, and no interesting regret bound is possible. To prevent this, it is standard in the literature to assume that $\langle \mathbf{w}, \mathbf{x} \rangle$ is essentially bounded. Formally, we assume that \mathcal{D} satisfies the following two conditions:

- $\|\bar{\mathbf{x}}\|_* \leq 1$
- $\Pr_{\mathbf{x} \sim \mathcal{D}} (\|\mathbf{x}\|_* > z) \leq 2 \exp(-z^2/2)$ for all $z \geq 1$.

We denote any such distribution as a “valid” distribution. The latter condition requires $\|\mathbf{x}\|_*$ to have sub-Gaussian tails (or more precisely, tails dominated by a standard Gaussian random variable), and is a slight relaxation of the standard ‘dual’ setting (e.g. Bubeck et al. (2012); Audibert et al. (2011)), where it is assumed that $\|\mathbf{x}\|_* \leq 1$ with probability 1. We choose this purely for technical convenience, since it allows us to use Gaussian distributions for \mathcal{D} , and does not materially affect algorithmic approaches we’re aware of. Moreover, in terms of lower bounds, we lose almost nothing by this relaxation: Any lower bound for our setting can be transformed into an equivalent lower bound in the standard dual setting (where $\|\mathbf{x}\|_* \leq 1$ with probability 1), at the cost of a $\sqrt{\log(T)}$ factor. See Appendix B for details.

3. Expected Error

We begin by considering the attainable performance in terms of expected error (Eq. (2)), recalling that error lower bounds immediately transfer to regret lower bounds.

To motivate our results, let us show that it is easy to obtain $\mathcal{O}(\sqrt{d/T})$ error upper bounds under a mild condition. To do so, suppose that \mathcal{W} contains the set $\{-\mu, \mu\}^d$ for some $\mu > 0$, and consider the following simple player strategy: In each round t , the player draws a vector $\boldsymbol{\sigma} \in \{-1, +1\}^d$ uniformly at random, computes $v = \langle \mathbf{x}_t, \mu \boldsymbol{\sigma} \rangle$ (possible since $\mu \boldsymbol{\sigma} \in \mathcal{W}$), and computes the estimator $\tilde{\mathbf{x}}_t = \frac{1}{\mu} v \boldsymbol{\sigma} = \langle \mathbf{x}_t, \boldsymbol{\sigma} \rangle \boldsymbol{\sigma}$. It is easy to verify that

$$\mathbb{E}[\tilde{\mathbf{x}}_t | \mathbf{x}_t] = \mathbf{x}_t, \quad \mathbb{E}[\|\tilde{\mathbf{x}}_t\|_2^2 | \mathbf{x}_t] = d \|\mathbf{x}_t\|_2^2, \quad (4)$$

so $\tilde{\mathbf{x}}_t$ is an unbiased estimator of \mathbf{x}_t with variance bounded by $\mathcal{O}(d)$. After T rounds, the player computes $\tilde{\mathbf{x}} = \frac{1}{T} \sum_{t=1}^T \tilde{\mathbf{x}}_t$, and returns $\hat{\mathbf{w}} = \arg \min_{\mathbf{w} \in \mathcal{W}} \langle \tilde{\mathbf{x}}, \mathbf{w} \rangle$.

Theorem 1 *Suppose that \mathcal{W} is a subset of the unit Euclidean ball $\{\mathbf{w} : \|\mathbf{w}\|_2 \leq 1\}$, and contains $\{-\mu, \mu\}^d$ for some $\mu > 0$. Then for any distribution over loss vectors \mathbf{x} , such that $\mathbb{E}[\|\mathbf{x}\|_2^2] \leq c^2$ and with mean $\bar{\mathbf{x}}$, the player strategy described above satisfies*

$$\mathbb{E}[\langle \tilde{\mathbf{x}}, \hat{\mathbf{w}} \rangle] - \min_{\mathbf{w} \in \mathcal{W}} \langle \tilde{\mathbf{x}}, \mathbf{w} \rangle \leq 2c \sqrt{\frac{d}{T}}.$$

We note that [Bubeck et al. \(2012\)](#) provides a more sophisticated algorithm with $\mathcal{O}(\sqrt{dT})$ regret, even against a non-stochastic adversary.

Proof The idea of the proof is that since $\{\tilde{\mathbf{x}}_t\}_{t=1}^T$ are unbiased estimate of \mathbf{x}_t , which are themselves random variables with mean $\bar{\mathbf{x}}$, then their average $\tilde{\mathbf{x}}$ will be a good estimate of $\bar{\mathbf{x}}$, hence minimizing $\langle \tilde{\mathbf{x}}, \mathbf{w} \rangle$ will approximately minimize $\langle \bar{\mathbf{x}}, \mathbf{w} \rangle$.

More formally, let $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle$. Also, recall that any $\mathbf{w} \in \mathcal{W}$ satisfies $\|\mathbf{w}\|_2 \leq 1$, and therefore $\langle \mathbf{x}, \mathbf{w} \rangle - \langle \mathbf{x}', \mathbf{w} \rangle \leq \|\mathbf{x} - \mathbf{x}'\|_2$ for all $\mathbf{w} \in \mathcal{W}$ by the Cauchy-Schwartz inequality. Finally, recall that $\hat{\mathbf{w}} = \arg \min_{\mathbf{w} \in \mathcal{W}} \langle \tilde{\mathbf{x}}, \mathbf{w} \rangle$. Using these observations, we have

$$\langle \bar{\mathbf{x}}, \hat{\mathbf{w}} \rangle \leq \langle \tilde{\mathbf{x}}, \hat{\mathbf{w}} \rangle + \|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\| \leq \langle \tilde{\mathbf{x}}, \mathbf{w}^* \rangle + \|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\| \leq \langle \bar{\mathbf{x}}, \mathbf{w}^* \rangle + 2 \|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\|,$$

and therefore

$$\mathbb{E}[\langle \bar{\mathbf{x}}, \hat{\mathbf{w}} \rangle] - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle = \mathbb{E}[\langle \bar{\mathbf{x}}, \hat{\mathbf{w}} \rangle] - \langle \bar{\mathbf{x}}, \mathbf{w}^* \rangle \leq 2\mathbb{E}[\|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\|] \leq 2\sqrt{\mathbb{E}[\|\bar{\mathbf{x}} - \tilde{\mathbf{x}}\|^2]}, \quad (5)$$

where the last step is by Jensen's inequality. Recalling the definition of $\tilde{\mathbf{x}}$, this equals

$$2\sqrt{\mathbb{E}\left[\left(\frac{1}{T} \sum_{t=1}^T (\tilde{\mathbf{x}}_t - \bar{\mathbf{x}})\right)^2\right]}.$$

According to equation Eq. (4), $\mathbb{E}[\tilde{\mathbf{x}}_t | \mathbf{x}_t] = \mathbf{x}_t$, but we also have $\mathbb{E}[\mathbf{x}_t] = \bar{\mathbf{x}}$, and therefore $\mathbb{E}[\tilde{\mathbf{x}}_t - \bar{\mathbf{x}}] = \bar{\mathbf{x}} - \bar{\mathbf{x}} = 0$, so each summand in the equation above is zero-mean. Moreover, they are i.i.d. since each $\tilde{\mathbf{x}}_t$ is computed based on the independent realization \mathbf{x}_t . Therefore, the equation above equals

$$2\sqrt{\frac{1}{T} \mathbb{E}[(\tilde{\mathbf{x}}_1 - \bar{\mathbf{x}})^2]} \leq 2\sqrt{\frac{1}{T} \mathbb{E}[\|\tilde{\mathbf{x}}_1\|^2]},$$

which by Eq. (4) is at most $2\sqrt{\frac{1}{T}d \mathbb{E}[\|\mathbf{x}_1\|_2^2]} \leq 2c\sqrt{d/T}$ as required. \blacksquare

For example, a special case of the theorem above implies an $\mathcal{O}(\sqrt{d/T})$ error upper bound for the origin-centered unit Euclidean ball. The key property used was the ability to query at a random point in $\{-\mu, +\mu\}^d$ (a scaled origin-centered Boolean hypercube), which allowed computing an unbiased estimator of each \mathbf{x}_t with variance $\mathcal{O}(d)$.

This observation leads us to guess that when we cannot query from such a set, the learning problem may be harder. One case in which querying such a set is impossible is when \mathcal{W} is convex and bounded away from the origin. In a full-information learning setting, shifting \mathcal{W} away from the origin doesn't increase the learning complexity in general. Surprisingly, in the bandit setting this turns out to make a huge difference: Even for the unit Euclidean ball, shifting it so it doesn't include the origin is sufficient to make the attainable error jump from $\sqrt{d/T}$ to d/\sqrt{T} :

Theorem 2 *Suppose that $d > 1$, and let $\mathcal{W} = \{\mathbf{a} + \mathbf{w} : \|\mathbf{w}\|_2 \leq 1\}$, where $\mathbf{a} = (2, 0, \dots, 0) \in \mathbb{R}^d$. Then for any player strategy returning $\hat{\mathbf{w}} \in \mathcal{W}$, there exists a valid distribution over loss vectors with mean $\bar{\mathbf{x}}$ such that*

$$\mathbb{E}[\langle \bar{\mathbf{x}}, \mathbf{w} \rangle] - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle \geq c \min \left\{ 1, \frac{d-1}{\sqrt{T}} \right\},$$

where c is a positive universal constant.

The formal proof appears in Appendix A.1, and we note that due to the rotational symmetry of our linear optimization setting, the same lower bound would hold for a Euclidean ball shifted in any other direction. Although the result is specifically for a shifted Euclidean ball, we conjecture that the result can be generalized to other domains as well, as long as they are sufficiently large and at a constant distance from the origin.

Although the notion of a domain not containing the origin may appear unusual at first, it can actually be necessary to model some common situations. One example is when the domain has a linear constraint bounding it away from the origin, such as the probability simplex or a convex hull of points lying in the same quadrant. As another example, suppose our losses actually take the form $\langle \mathbf{w}_t, \mathbf{x}_t \rangle + \xi_t$ where ξ_t is some noise or other term not controlled by the player. This is possible to model in the bandit setting, by adding a dimension and playing over the domain $\mathcal{W} \times \{1\}$, using loss vectors of the form $(x_{t,1}, x_{t,2}, \dots, x_{t,d}, \xi_t)$. But this leads to a domain bounded away from the origin. In fact, if \mathcal{W} is the origin-centered Euclidean ball, then the techniques of Thm. 2 readily imply a $\Omega(d/\sqrt{T})$ error lower bound for the domain $\mathcal{W} \times \{1\}$.

As discussed earlier, the key idea in proving Thm. 2 is that we can construct a distribution where the variance of $\langle \mathbf{w}, \mathbf{x} \rangle$ for any $\mathbf{w} \in \mathcal{W}$ is lower bounded by a positive constant, independent of the norm of \mathbf{w} or the dimension. The effect of this is best seen in the estimation procedure described before Eq. (4): Recall that there we chose a query point $\mathbf{w}_t = \mu\boldsymbol{\sigma}$ where $\boldsymbol{\sigma} \in \{-1, +1\}^d$, in which case the conditional second moment of the estimator $\tilde{\mathbf{x}}_t = \frac{1}{\mu}v\boldsymbol{\sigma}$, where $v = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$, satisfies

$$\mathbb{E}[\|\tilde{\mathbf{x}}_t\|_2^2 | \mathbf{x}_t] = \frac{1}{\mu^2} \mathbb{E}[v^2 \|\boldsymbol{\sigma}\|_2^2] = \frac{d \mathbb{E}[v^2]}{\mu^2}. \quad (6)$$

Since $v = \langle \mathbf{x}_t, \mathbf{w}_t \rangle = \mu \langle \mathbf{x}_t, \boldsymbol{\sigma} \rangle$, we see that $\mathbb{E}[v^2]$ scales with μ^2 , and therefore Eq. (6) is $\mathcal{O}(d)$, independent of μ . In contrast, if $\mathbb{E}[v^2] = \mathbb{E}[\langle \mathbf{x}_t, \mathbf{w}_t \rangle^2]$ was forced to have constant positive variance,

then Eq. (6) is at least $\Omega(d/\mu^2)$. Moreover, if \mathcal{W} has bounded norm, then $\mu \leq \mathcal{O}(1/\sqrt{d})$, so Eq. (6) would scale as $\mathcal{O}(d^2)$ rather than $\mathcal{O}(d)$, eventually leading to an error bound scaling as d/\sqrt{T} rather than $\sqrt{d/T}$.

4. From Error to Regret

Having considered domains bounded away from the origin, it is natural to ask whether this is the only condition leading to $\Omega(d\sqrt{T})$ regret lower bounds. If the domain does contain the origin, then because of Thm. 1, it seems unlikely to show such lower bounds by proving error lower bounds. However, we can exploit the fact that attaining small regret is harder than attaining small error. In this section, we show how in fact it can be strictly harder: We identify a situation which leads to $\Omega(d\sqrt{T})$ regret lower bounds, even if the domain contains the origin, and even though better error upper bounds are possible.

Specifically, the following theorem demonstrates such a lower bound for a domain consisting of an origin-centered cylinder.

Theorem 3 *Suppose that $d > 1$, and let $\mathcal{W} = [-1, 1] \times \{\mathbf{w} \in \mathbb{R}^{d-1} : \|\mathbf{w}\|_2 \leq 1\}$. Then for any player strategy, there exists a valid distribution over loss vectors with mean $\bar{\mathbf{x}}$ such that*

$$\mathbb{E} \left[\sum_{t=1}^T \langle \bar{\mathbf{x}}, \mathbf{w}_t \rangle \right] - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \langle \bar{\mathbf{x}}, \mathbf{w} \rangle \geq c(d-1)\sqrt{T}$$

for any $T \geq d^4/16$, where c is a positive universal constant.

Note that the bound is a bit weaker than Thm. 2, in that it only holds for sufficiently large T . However, since this is a lower bound, it is still sufficient for proving that no algorithm for this setting will attain regret better than $\mathcal{O}(d\sqrt{T})$ in general.

The formal proof appears in Appendix A.2, and relies on the following construction: The adversary chooses a vector $\boldsymbol{\sigma} \in \{-1, +1\}^d$ uniformly at random, and constructs a loss vector distribution, whose expectation $\bar{\mathbf{x}}$ is $(-\frac{1}{4}, \mu\sigma_1, \dots, \mu\sigma_d)$ (with μ being a small scaling factor, on the order of $\sqrt{d/T}$ if T is sufficiently large), and where the first coordinate has constant Gaussian noise. It is possible to show that to get small regret, the player essentially needs to identify $\sigma_1, \dots, \sigma_d$. This would have been possible if the player queried at points \mathbf{w}_t whose first coordinate is 0. However, since $\bar{\mathbf{x}}$ has a large negative weight on the first coordinate, then $w_{t,1}$ must be large to get small regret. But because of the stochastic noise in the first coordinate, this means that $\langle \mathbf{w}_t, \bar{\mathbf{x}} \rangle$ will have large variance. A formal analysis of this leads to a regret lower bound of the form

$$\Omega \left(\sum_{t=1}^T \mathbb{E}[1 - w_{t,1}] + \mu\sqrt{dT} \left(1 - \mu \sqrt{\frac{1}{d} \sum_{t=1}^T \mathbb{E} \left[\frac{1}{w_{t,1}^2 + \frac{1}{d}} \right]} \right) \right).$$

The idea now is that no matter how the first coordinates $\{w_{t,1}\}_{t=1}^T$ are chosen by the player, the regret will be large: If they are significantly smaller than 1, then the first term above will be on the order of T . But if $w_{t,1}$ is almost 1, then the square root term will be small, again leading to a large regret. A careful analysis shows that the regret will always be $\Omega(d\sqrt{T})$.

Although a precise characterization is non-trivial, the proof technique appears to be potentially applicable to any domain with a flat surface, which is sufficiently large to contain the points

$(-a, \mu\sigma_1, \dots, \mu\sigma_d)$ for some constant $a \in (0, 1)$ and suitable scaling factor μ (possibly after rotating the domain around the origin). For example, the same proof technique would apply to a capped Euclidean ball, $\mathcal{W} = \{\mathbf{w} : \|\mathbf{w}\| \leq 1, w_1 \leq c\}$ for some constant $c \in (0, 1)$. The flatness of the surface seems to be important. To see this, let us understand why the proof breaks down, when instead of a cylinder or a capped ball, our domain is the origin-centered Euclidean unit ball (for which we know from [Bubeck et al. \(2012\)](#) that we can get $\mathcal{O}(\sqrt{dT})$ regret). First, when T is sufficiently large compared to d , we choose μ on the order of $\sqrt{d/T}$ (this is the regime which makes detecting $\sigma_1, \dots, \sigma_d$ information-theoretically hard). In that case, we have $\bar{\mathbf{x}} = \left(-a, \sqrt{\frac{d}{T}}\sigma_1, \dots, \sqrt{\frac{d}{T}}\sigma_d\right)$, and the optimal play is $\mathbf{w}^* = -\frac{1}{\|\bar{\mathbf{x}}\|}\bar{\mathbf{x}} = \frac{1}{\sqrt{a^2 + \frac{d^2}{T}}}\left(a, -\sqrt{\frac{d}{T}}\sigma_1, \dots, -\sqrt{\frac{d}{T}}\sigma_d\right)$. However, with an origin-centered ball, the player doesn't need to detect $\sigma_1, \dots, \sigma_d$ to get small regret: The player can go "further" along the first coordinate, and just play the fixed point $\mathbf{w} = (1, 0, \dots, 0)$ every round. The total expected regret is then

$$\begin{aligned}
 T(\langle \bar{\mathbf{x}}, \mathbf{w} \rangle - \langle \bar{\mathbf{x}}, \mathbf{w}^* \rangle) &= T\left(\bar{x}_1 + \frac{1}{\|\bar{\mathbf{x}}\|}\|\bar{\mathbf{x}}\|^2\right) = T(\bar{x}_1 + \|\bar{\mathbf{x}}\|) \\
 &= T\left(-a + \sqrt{a^2 + \frac{d^2}{T}}\right) \leq T\left(-a + a + \frac{1}{2a}\frac{d^2}{T}\right) = \frac{1}{2a}d^2 = \mathcal{O}(d^2),
 \end{aligned}$$

where we used the fact (immediate by Taylor expansion) that $\sqrt{a^2 + x} \leq a + \frac{1}{2a}x$ for all $x \geq 0$. So, we see that instead of $\Omega(d\sqrt{T})$ regret as in the case of a cylinder or capped ball, here the player can achieve a much smaller $\mathcal{O}(d^2)$ regret (assuming T is sufficiently large). In fact, this $\mathcal{O}(d^2)$ regret bound is tight for our construction – carrying through the lower bound derivation as in [Thm. 3](#), but for the origin-centered ball, yields an $\Omega(d^2)$ regret lower bound for sufficiently large T .

5. Loss Vectors from a Finite Set

Many of the regret lower bounds shown in the literature for finite domains (such as for the corners of the simplex and for the boolean hypercube ([Auer et al. \(2002\)](#); [Dani et al. \(2007\)](#); [Audibert et al. \(2011\)](#))) use loss vectors from a finite set (e.g. where each entry is binary). Based on analogues to the full information setting, it is often argued that these lower bound constructions also extend to the (continuous) convex hull of these domains. In this section, we point out that these analogues can be dangerous and generally do not hold.

In particular, the theorem below demonstrates a possibly surprising fact: If we perform bandit linear optimization over a continuous domain, and the loss vectors come from a finite set, then we can actually get the same regret as in the full information case – generally much smaller than the lower bound one would hope to achieve. This also holds for non-stochastic adversaries. We hasten to emphasize that the way we achieve this is "cheating" and not very useful in practice, since it heavily relies on the ability to perform arbitrary-precision computations. Nevertheless, it demonstrates that one has to be careful when extending bandit optimization results from a finite domain to its convex hull. Moreover, it shows that loss vectors supported on a finite set are not suitable to prove bandit lower bounds for continuous domains.

For simplicity, we will prove the result for the simplex and for binary-valued loss vectors, but from the proof it is easily seen to be extendable to other continuous domains in general (such as the hypercube $[-1, +1]^d$), as well as the loss vectors coming from any finite set.

Theorem 4 Suppose $\mathcal{W} = \{\mathbf{w} : \sum_i w_i = 1, \forall_i w_i \geq 0\}$ is the probability simplex in \mathbb{R}^d , and suppose the loss vectors are from the set $\{0, 1\}^d$. Then there exist a deterministic player strategy in the bandit setting, such that for any adversarial strategy for choosing loss vectors $\mathbf{x}_1, \dots, \mathbf{x}_T$,

$$\sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w}_t \rangle - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w} \rangle \leq \mathcal{O}\left(\sqrt{\log(d)T}\right)$$

Proof The proof uses the following observation: In the setting described above, it is possible for the player to determine \mathbf{x}_t precisely by perturbing its chosen point by an arbitrarily small amount, and feed it into a deterministic full-information algorithm for this domain, such as hedge (Freund and Schapire (1997); Cesa-Bianchi and Lugosi (2006)). Such algorithms attain $\mathcal{O}(\sqrt{\log(d)T})$ regret, from which the result follows.

More precisely, suppose that at the beginning of round t , the full-information algorithm determines that one should play point $\hat{\mathbf{w}}_t$. Let $p = \lceil \log_{10}(T) \rceil$, and define

$$\mathbf{w}'_t = \text{clip}_p(\hat{\mathbf{w}}_t) + (10^{-p-1}, 10^{-p-2}, \dots, 10^{-p-d}),$$

where $\text{clip}_p(\mathbf{w})$ clips every entry of \mathbf{w} (in its decimal representation) to p places after the decimal point. Note that for any $\mathbf{x} \in \{0, 1\}^d$, if we get $\langle \mathbf{x}, \mathbf{w}'_t \rangle = \sum_{i=1}^d x_i w'_{t,i}$, then we can determine \mathbf{x} precisely: We just need to look at the digits in locations $p+1, p+2, \dots, p+d$ after the decimal point. A small technical issue is that \mathbf{w}'_t does not lie in the simplex \mathcal{W} , so it cannot be chosen by the player. To fix this, the player picks $\mathbf{w}_t = \frac{1}{\|\mathbf{w}'_t\|_1} \mathbf{w}'_t$, which indeed lies in the simplex. Given the loss $\langle \mathbf{x}_t, \mathbf{w}_t \rangle$, we just multiply it by the known quantity $\|\mathbf{w}'_t\|_1$ to get $\langle \mathbf{x}_t, \mathbf{w}'_t \rangle$, from which we can read off \mathbf{x}_t and feed it back to the full-information algorithm.

By standard regret guarantees (e.g. Freund and Schapire (1997); Cesa-Bianchi and Lugosi (2006))², we therefore get that

$$\sum_{t=1}^T \langle \mathbf{x}_t, \hat{\mathbf{w}}_t \rangle - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w} \rangle \leq \mathcal{O}\left(\sqrt{\log(d)T}\right) \quad (7)$$

To get a regret bound using the actual plays \mathbf{w}_t , we note that we can lower bound the left hand side of Eq. (7) by

$$\sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w}_t \rangle - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w} \rangle - \sum_{t=1}^T |\langle \mathbf{x}_t, \hat{\mathbf{w}}_t - \mathbf{w}_t \rangle|, \quad (8)$$

and that

$$\begin{aligned} |\langle \mathbf{x}_t, \hat{\mathbf{w}}_t - \mathbf{w}_t \rangle| &\leq \|\hat{\mathbf{w}}_t - \mathbf{w}_t\|_1 = \left\| \frac{1}{\|\mathbf{w}'_t\|_1} \mathbf{w}'_t - \mathbf{w}_t \right\|_1 \leq \|\mathbf{w}'_t - \mathbf{w}_t\|_1 + \left\| \mathbf{w}'_t - \frac{1}{\|\mathbf{w}'_t\|_1} \mathbf{w}'_t \right\|_1 \\ &= \|\mathbf{w}'_t - \mathbf{w}_t\|_1 + \left| 1 - \frac{1}{\|\mathbf{w}'_t\|_1} \right| \|\mathbf{w}'_t\|_1 = \|\mathbf{w}'_t - \mathbf{w}_t\|_1 + \left| \|\mathbf{w}'_t\|_1 - 1 \right|. \end{aligned}$$

2. Strictly speaking, these algorithms are phrased so as to minimize $\max_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w} \rangle - \sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w}_t \rangle$, but they can be easily converted to our notion of regret by feeding them with $\mathbf{1} - \mathbf{x}_t$ where $\mathbf{1}$ is the all-ones vector.

By definition of \mathbf{w}'_t and the fact that \mathbf{w}_t is on the simplex, it is easy to verify that $\|\mathbf{w}'_t\|_1 \in [1 - \frac{1}{T}, 1 + \frac{1}{T}]$ and that $\|\mathbf{w}'_t - \mathbf{w}_t\|_1 \leq \frac{1}{T}$, so the above equals at most $2/T$, and we can lower bound Eq. (8) by

$$\sum_{t=1}^T \langle \mathbf{x}_t, \hat{\mathbf{w}}_t \rangle - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w} \rangle - 2$$

Combining this with Eq. (7), we get that

$$\sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w}_t \rangle - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \langle \mathbf{x}_t, \mathbf{w} \rangle \leq \mathcal{O}(\sqrt{\log(d)T}) + 2 = \mathcal{O}(\sqrt{\log(d)T}).$$

as required. ■

We note that this result depends on both the domain being continuous and the losses being discrete, and the regret can be $\Omega(\sqrt{dT})$ if we remove any of these requirements³

In light of this type of result, we are not aware of explicit minimax regret bounds for the simplex (or unit 1-norm ball) and the hypercube in the literature. For completeness, we provide such bounds below. Besides relying on continuous-valued rewards, the lower bounds require the techniques of Thm. 2 and Thm. 3: For the simplex, we provide an error lower bound relying on the fact that the simplex is bounded away from the origin; whereas for the hypercube, we rely on it having a flat surface orthogonal to the origin. The constructions we use are slightly different though, due to the different shapes of the domains. The formal proofs appear in Appendix A.3 and Appendix A.4.

Theorem 5 *Suppose that $d > 1$, and let \mathcal{W} be either the simplex $\{\mathbf{w} : \forall i w_i \geq 0, \|\mathbf{w}\|_1 = 1\}$ or the unit 1-norm ball $\{\mathbf{w} : \|\mathbf{w}\|_1 \leq 1\}$. Then for any player strategy returning some $\hat{\mathbf{w}} \in \mathcal{W}$, there exists a valid distribution over loss vectors with mean $\bar{\mathbf{x}}$ such that*

$$\mathbb{E}[\langle \bar{\mathbf{x}}, \mathbf{w} \rangle] - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle \geq c \min \left\{ 1, \sqrt{\frac{d}{T}} \right\},$$

where c is a positive universal constant.

This leads to a $\Omega(\sqrt{dT})$ regret lower bound, which matches the $\mathcal{O}(\sqrt{dT})$ upper bound for the simplex attained with multi-armed bandit algorithms such as EXP3 (Auer et al. (2002)).

Theorem 6 *Suppose that $d > 1$, and let $\mathcal{W} = [-1, 1]^d$. Then for any player strategy, there exists a valid distribution over loss vectors with mean $\bar{\mathbf{x}}$ such that*

$$\mathbb{E} \left[\sum_{t=1}^T \langle \bar{\mathbf{x}}, \mathbf{w}_t \rangle \right] - \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \langle \bar{\mathbf{x}}, \mathbf{w} \rangle \geq c(d-1)\sqrt{T}$$

for any $T \geq d^4/4$, where c is a positive universal constant.

This $\Omega(d\sqrt{T})$ regret lower bound matches the $\mathcal{O}(d\sqrt{T})$ upper bound attained in Bubeck et al. (2012) (using an algorithm which actually plays only on the corners of the hypercube).

3. Auer et al. (2002) show such a lower bound when the domain \mathcal{W} is the corners of the simplex, even though the losses are discrete and binary-valued. Thm. 5 shows such a lower bound using continuous-valued losses, even though the domain is the continuous probability simplex.

6. Conclusion

In this paper, we studied the attainable regret for bandit linear optimization problems, focusing on lower bounds, and made several observations:

- We formally proved that the price of bandit information for linear optimization can be as large as d (rather than \sqrt{d}), settling the conjecture made in Dani et al. (2007). In a stochastic optimization setting, this implies that even though we get “ $1/d$ ” as much information in the bandit setting, compared to the full-information setting, we might need d^2 (rather than d) more queries to achieve a similar error.
- In a bandit setting, the dependence of the regret on the domain geometry can be quite different than in the full-information setting, sometimes in counter-intuitive ways. Essentially, this is because changing the domain also changes the information feedback structure. In particular, supposedly “trivial” changes to the domain, such as shifting it away from the origin or intersecting it with a halfspace, can increase the regret from $\mathcal{O}(\sqrt{dT})$ to $\Omega(d\sqrt{T})$. Moreover, regret lower bounds for a given domain \mathcal{W} don’t necessarily extend to its convex hull.

Our results build on the following observation: At least for stochastic adversaries, and domains with an Euclidean-norm geometry (such as an Euclidean sphere), the question of whether the regret scales as \sqrt{d} or d crucially depends on whether the player can control the variance in the observed feedback. In particular, our results indicate that when the feedback is inherently noisy, e.g. of the form $\langle \mathbf{x}_t, \mathbf{w}_t \rangle + \xi_t$, where ξ_t is independent sub-Gaussian noise, then even for a benign domain such as an origin-centered Euclidean sphere, the regret is $\Theta(d\sqrt{T})$ (see Thm. 2 and the following discussion). We conjecture that with such noisy feedback, $\Theta(d\sqrt{T})$ is the attainable regret for any domain which is “non-degenerate” under the Euclidean norm. However, proving this conjecture, and understanding how the regret scales for arbitrary domains, is left to future work.

Acknowledgments

This research was supported by an Israel Science Foundation grant 425/13 and an FP7 Marie Curie CIG grant. We thank Sébastien Bubeck for several illuminating discussions which helped initiate the research presented here.

References

- J.-Y. Audibert, S. Bubeck, and G. Lugosi. Minimax policies for combinatorial prediction games. In *COLT*, 2011.
- P. Auer, N. Cesa-Bianchi, Y. Freund, and R. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- S. Bubeck, N. Cesa-Bianchi, and S. Kakade. Towards minimax policies for online linear optimization with bandit feedback. In *COLT*, 2012.

- N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- N. Cesa-Bianchi and G. Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5), 2012.
- V. Dani, S. Kakade, and T. Hayes. The price of bandit information for online optimization. In *NIPS*, 2007.
- V. Dani, T. Hayes, and S. Kakade. Stochastic linear optimization under bandit feedback. In *COLT*, 2008.
- O. Dekel, J. Ding, T. Koren, and Y. Peres. Bandits with switching costs: $t^{2/3}$ regret. 2014.
- Y. Freund and R. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.*, 55(1):119–139, 1997.
- E. Hazan, Z. Karnin, and R. Mehta. Volumetric spanners: an efficient exploration basis for learning. In *COLT*, 2014.
- R. T. Rockafellar. *Convex analysis*. Princeton University Press, 1970.
- O. Shamir. On the complexity of bandit and derivative-free stochastic convex optimization. In *COLT*, 2013.

Appendix A. Proofs

A.1. Proof of Thm. 2

To simplify notation, suppose that the game takes place in \mathbb{R}^{d+1} for some $d > 0$. We denote the first coordinate as coordinate 0, and the other coordinates as $1, 2, \dots, d$.

By Yao’s minimax principle, it is sufficient to provide a randomized strategy to choose a loss vector distribution \mathcal{D} , such that for any deterministic player, the expected error is as defined in the theorem. In particular, we use the following strategy:

- Choose $\sigma \in \{-1, +1\}^d$ uniformly at random.
- Use the distribution \mathcal{D}_σ over loss vectors \mathbf{x} , defined as follows: (x_1, \dots, x_d) is fixed to be $\mu\sigma$ (where $\mu \leq \frac{1}{2\sqrt{d}}$ is a parameter to be chosen later), and x_0 has a Gaussian distribution $\mathcal{N}(0, \frac{1}{36})$.

First, let us verify that any \mathcal{D}_σ is a valid distribution. We have

$$\mathbb{E}[\|\bar{\mathbf{x}}\|_*] = \sup_{\mathbf{w} \in \mathcal{W}} |\langle \mathbb{E}[\mathbf{x}], \mathbf{w} \rangle| = \sup_{\mathbf{w} \in \mathcal{W}} \left| 0 + \mu \sum_{i=1}^d \sigma_i w_i \right| = \mu\sqrt{d} \leq \frac{1}{2},$$

and moreover, for any \mathbf{x} in the support of \mathcal{D}_σ and for any $\mathbf{w} \in \mathcal{W}$,

$$|\langle \mathbf{x}, \mathbf{w} \rangle| \leq |3x_0| + \left| \sum_{i=1}^d x_i w_i \right| \leq |3x_0| + \mu \sum_{i=1}^d |w_i| \leq |3x_0| + \mu\sqrt{d} \leq |3x_0| + \frac{1}{2},$$

by the fact that $\|\cdot\|_1 \leq \sqrt{d}\|\cdot\|_2$ and our assumption on μ . $3x_0$ is normally distributed with mean zero and variance $9 * \frac{1}{36} = 1/4$, from which it is easily verified that $\Pr(\sup_{\mathbf{w}} |\langle \mathbf{x}, \mathbf{w} \rangle| > z) \leq 2 \exp(-z^2/2)$ for all $z \geq 1$ as required.

We now start the proof. For any fixed \mathcal{D}_σ , recall that $\bar{\mathbf{x}} = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_\sigma}[\mathbf{x}] = (0, \mu\sigma_1, \dots, \mu\sigma_d)$, and define

$$\mathbf{w}^* = \left(0, -\frac{\sigma_1}{\sqrt{d}}, \dots, -\frac{\sigma_d}{\sqrt{d}}\right) = -\frac{1}{\mu\sqrt{d}}\bar{\mathbf{x}}.$$

It is easily verified that \mathbf{w}^* is a minimizer of $\langle \bar{\mathbf{x}}, \mathbf{w} \rangle$ over \mathcal{W} . Therefore,

$$\begin{aligned} \mathbb{E} \left[\langle \bar{\mathbf{x}}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle \right] &= \mathbb{E}[\langle \bar{\mathbf{x}}, \hat{\mathbf{w}} - \mathbf{w}^* \rangle] = -\mu\sqrt{d} \mathbb{E}[\langle \mathbf{w}^*, \hat{\mathbf{w}} - \mathbf{w}^* \rangle] \\ &= \mu\sqrt{d} \mathbb{E}[\|\mathbf{w}^*\|_2^2 - \langle \mathbf{w}^*, \hat{\mathbf{w}} \rangle] = \frac{\mu\sqrt{d}}{2} \mathbb{E}[\|\mathbf{w}^*\|_2^2 + \|\mathbf{w}^*\|_2^2 - 2\langle \mathbf{w}^*, \hat{\mathbf{w}} \rangle]. \end{aligned}$$

Since $\sum_{i=1}^d \hat{w}_i^2 \leq 1 = \|\mathbf{w}^*\|_2^2$ and $w_0^* = 0$, we can lower bound the above by

$$\begin{aligned} &\frac{\mu\sqrt{d}}{2} \mathbb{E} \left[\sum_{i=1}^d (w_i^*)^2 + \sum_{i=1}^d \hat{w}_i^2 - 2 \sum_{i=1}^d w_i^* \hat{w}_i \right] = \frac{\mu\sqrt{d}}{2} \mathbb{E} \left[\sum_{i=1}^d (w_i^* - \hat{w}_i)^2 \right] \\ &= \frac{\mu\sqrt{d}}{2} \sum_{i=1}^d \mathbb{E} \left[\left(\frac{\sigma_i}{\sqrt{d}} + \hat{w}_i \right)^2 \right] \geq \frac{\mu\sqrt{d}}{2} \sum_{i=1}^d \left(\frac{1}{\sqrt{d}} \right)^2 \Pr(\sigma_i \hat{w}_i \geq 0) \\ &= \frac{\mu}{2\sqrt{d}} \sum_{i=1}^d \Pr(\sigma_i \hat{w}_i \geq 0), \end{aligned} \tag{9}$$

where $\Pr(\cdot)$ is the probability with respect to the joint randomness of $\boldsymbol{\sigma}$ and $\hat{\mathbf{w}}$. Since each σ_i is uniformly distributed on $\{-1, +1\}$, we can lower bound each probability term as follows:

$$\begin{aligned} \Pr(\sigma_i \hat{w}_i \geq 0) &= \frac{1}{2} (\Pr(\hat{w}_i \geq 0 | \sigma_i = 1) + \Pr(\hat{w}_i \leq 0 | \sigma_i = -1)) \\ &\geq \frac{1}{2} (\Pr(\hat{w}_i > 0 | \sigma_i = 1) + 1 - \Pr(\hat{w}_i > 0 | \sigma_i = -1)) \\ &= \frac{1}{2} (1 - (\Pr(\hat{w}_i > 0 | \sigma_i = -1) - \Pr(\hat{w}_i > 0 | \sigma_i = 1))) \\ &\geq \frac{1}{2} (1 - |\Pr(\hat{w}_i > 0 | \sigma_i = -1) - \Pr(\hat{w}_i > 0 | \sigma_i = 1)|). \end{aligned}$$

Plugging this back to Eq. (9), we get the lower bound

$$\begin{aligned} &\frac{\mu}{4\sqrt{d}} \sum_{i=1}^d (1 - |\Pr(\hat{w}_i > 0 | \sigma_i = -1) - \Pr(\hat{w}_i > 0 | \sigma_i = 1)|) \\ &= \frac{\mu\sqrt{d}}{4} \left(1 - \frac{1}{d} \sum_{i=1}^d |\Pr(\hat{w}_i > 0 | \sigma_i = -1) - \Pr(\hat{w}_i > 0 | \sigma_i = 1)| \right). \end{aligned}$$

To continue, recall that we assume the player is deterministic, in which case $\mathbf{w} = \mathbf{w}(\mathbf{v})$ is a function of the sequence of losses $\mathbf{v} = (v_1, \dots, v_T)$ observed by the player over T rounds. Letting p denote

the probability density function over \mathbf{v} , and $\mathbf{1}_A$ denote the indicator of the event A , we can rewrite the above as

$$\begin{aligned}
 & \frac{\mu\sqrt{d}}{4} \left(1 - \frac{1}{d} \sum_{i=1}^d \left| \int_{\mathbf{v}} \mathbf{1}_{w_i(\mathbf{v}) > 0} p(\mathbf{v} | \sigma_i = -1) d\mathbf{v} - \int_{\mathbf{v}} \mathbf{1}_{w_i > 0} p(\mathbf{v} | \sigma_i = 1) d\mathbf{v} \right| \right) \\
 &= \frac{\mu\sqrt{d}}{4} \left(1 - \frac{1}{d} \sum_{i=1}^d \left| \int_{\mathbf{v}} \mathbf{1}_{w_i(\mathbf{v}) > 0} (p(\mathbf{v} | \sigma_i = -1) - p(\mathbf{v} | \sigma_i = 1)) d\mathbf{v} \right| \right) \\
 &\geq \frac{\mu\sqrt{d}}{4} \left(1 - \frac{1}{d} \sum_{i=1}^d \int_{\mathbf{v}} |p(\mathbf{v} | \sigma_i = -1) - p(\mathbf{v} | \sigma_i = 1)| d\mathbf{v} \right). \tag{10}
 \end{aligned}$$

Each integral represents the total variation distance between the densities $p(\cdot | \sigma_i = 1)$ and $p(\cdot | \sigma_i = -1)$. By Pinsker's inequality, it can be upper bounded as follows:

$$\int_{\mathbf{v}} |p(\mathbf{v} | \sigma_i = -1) - p(\mathbf{v} | \sigma_i = 1)| d\mathbf{v} \leq \sqrt{2 D_{kl}(p(\mathbf{v} | \sigma_i = -1) || p(\mathbf{v} | \sigma_i = 1))}, \tag{11}$$

where $D_{kl}(p||q)$ is the Kullback-Leibler (KL) divergence between p and q . By the chain rule, this can be upper bounded in turn by

$$\sqrt{2 \sum_{t=1}^T D_{kl}(p(v_t | \sigma_i = -1, v_1, \dots, v_{t-1}) || p(v_t | \sigma_i = 1, v_1, \dots, v_{t-1}))}. \tag{12}$$

Since the player is deterministic, any values v_1, \dots, v_{t-1} determine the point \mathbf{w}_t that the player will choose on round t . Thus, the distribution of $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ depends on the distribution of \mathbf{x}_t , as determined by \mathcal{D}_σ . By definition of \mathcal{D}_σ :

- Under the condition $\sigma_i = 1$, $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ has a Gaussian distribution with mean $\mu w_{t,i} + \sum_{j \in \{1 \dots d\} \setminus \{i\}} (\mu \sigma_j w_{t,j})$, and variance $\frac{1}{36} w_{t,0}^2 \geq \frac{1}{36}$ (note that here we crucially use the domain assumptions on \mathcal{W} , which imply $w_{t,0} \geq 1$).
- Under the condition $\sigma_i = -1$, $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ has a Gaussian distribution with mean $-\mu w_{t,i} + \sum_{j \in \{1 \dots d\} \setminus \{i\}} (\mu \sigma_j w_{t,j})$, and variance $\frac{1}{36} w_{t,0}^2 \geq \frac{1}{36}$.

By a standard result on the KL divergence of two Gaussian distributions, we therefore have that Eq. (12) is at most

$$\sqrt{2 \sum_{t=1}^T \mathbb{E} \left[72 \mu^2 w_{t,i}^2 | \sigma_i = -1 \right]}.$$

Recalling that this constitutes an upper bound on the left hand side of Eq. (11), we get

$$\int_{\mathbf{v}} |p(\mathbf{v} | \sigma_i = -1) - p(\mathbf{v} | \sigma_i = 1)| d\mathbf{v} \leq \sqrt{2 \sum_{t=1}^T \mathbb{E} \left[72 \mu^2 w_{t,i}^2 | \sigma_i = -1 \right]}. \tag{13}$$

To get rid of the conditioning on $\sigma_i = -1$, note that the total variation is symmetric, so again using Pinsker's inequality, we can get a variant of Eq. (11) of the form

$$\int_{\mathbf{v}} |p(\mathbf{v} | \sigma_i = -1) - p(\mathbf{v} | \sigma_i = 1)| d\mathbf{v} \leq \sqrt{2 D_{kl}(p(\mathbf{v} | \sigma_i = 1) || p(\mathbf{v} | \sigma_i = -1))}.$$

Using the same derivation as above, this leads to the upper bound

$$\int_{\mathbf{v}} |p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1)| d\mathbf{v} \leq \sqrt{2 \sum_{t=1}^T \mathbb{E} [72\mu^2 w_{t,i}^2 | \sigma_i = 1]}. \quad (14)$$

Combining Eq. (13) and Eq. (14), and using the elementary inequality $\min\{\sqrt{a}, \sqrt{b}\} \leq \sqrt{\frac{a+b}{2}}$ and the fact that σ_i is uniformly distributed on $\{-1, +1\}$, we get

$$\begin{aligned} \int_{\mathbf{v}} |p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1)| d\mathbf{v} &\leq \min \left\{ \sqrt{2 \sum_{t=1}^T \mathbb{E} [72\mu^2 w_{t,i}^2 | \sigma_i = -1]}, \sqrt{2 \sum_{t=1}^T \mathbb{E} [72\mu^2 w_{t,i}^2 | \sigma_i = 1]} \right\} \\ &\leq \sqrt{144\mu^2 \sum_{t=1}^T \frac{1}{2} (\mathbb{E} [w_{t,i}^2 | \sigma_i = -1] + \mathbb{E} [w_{t,i}^2 | \sigma_i = 1])} \\ &= \mu \sqrt{144 \sum_{t=1}^T \mathbb{E} [w_{t,i}^2]}. \end{aligned} \quad (15)$$

Plugging this back into Eq. (10), we get an expected error lower bound of

$$\frac{\mu\sqrt{d}}{4} \left(1 - \frac{1}{d} \sum_{i=1}^d \mu \sqrt{144 \sum_{t=1}^T \mathbb{E} [w_{t,i}^2]} \right),$$

which by Jensen's inequality is at least

$$\frac{\mu\sqrt{d}}{4} \left(1 - \mu \sqrt{\frac{144}{d} \sum_{i=1}^d \sum_{t=1}^T \mathbb{E} [w_{t,i}^2]} \right) = \frac{\mu\sqrt{d}}{4} \left(1 - \mu \sqrt{\frac{144}{d} \sum_{t=1}^T \mathbb{E} \left[\sum_{i=1}^d w_{t,i}^2 \right]} \right).$$

By definition of \mathcal{W} , we have that $\sum_{i=1}^d w_{t,i}^2 \leq 1$ always, so we can lower bound this by

$$\frac{\mu\sqrt{d}}{4} \left(1 - \mu \sqrt{\frac{144T}{d}} \right). \quad (16)$$

Let us now consider two cases:

- If $\frac{1}{\sqrt{d}} > \sqrt{\frac{d}{144T}}$, then we choose $\mu = \frac{1}{2} \sqrt{\frac{d}{144T}}$ (which ensures the condition $\mu \leq \frac{1}{2\sqrt{d}}$), and Eq. (16) equals $\frac{\mu\sqrt{d}}{4} * \frac{1}{2} = \frac{d}{16\sqrt{144T}}$.
- If $\frac{1}{\sqrt{d}} \leq \sqrt{\frac{d}{144T}}$, then we choose $\mu = \frac{1}{2\sqrt{d}}$, and Eq. (16) is lower bounded by $\frac{\mu\sqrt{d}}{4} * \frac{1}{2} = \frac{1}{16}$.

Therefore, by choosing μ appropriately, we can lower bound Eq. (16) by $\min \left\{ \frac{1}{16}, \frac{d}{16\sqrt{144T}} \right\} \geq 0.005 \min \left\{ 1, \frac{d}{\sqrt{T}} \right\}$. Since this lower bound holds for a domain in \mathbb{R}^{d+1} for any $d > 0$, we get a $0.005 \min \left\{ 1, \frac{d-1}{\sqrt{T}} \right\}$ lower bound for a domain in \mathbb{R}^d for any $d > 1$, as required.

A.2. Proof of Thm. 3

As in the proof of Thm. 2, to simplify notation, we assume the game takes place in \mathbb{R}^{d+1} for some $d > 0$, so the domain is $\mathcal{W} = [-1, +1] \times \{\mathbf{w} \in \mathbb{R}^d : \|\mathbf{w}\|_2 \leq 1\}$. We denote the first coordinate as coordinate 0, and the other coordinates as 1, 2, \dots , d .

The proof uses a somewhat different loss vector distribution than that of Thm. 2, designed to force the algorithm to choose points far from and to one side of the origin to ensure small regret.

By Yao's minimax principle, it is sufficient to provide a randomized strategy to choose a loss vector distribution \mathcal{D} , such that for any deterministic player, the expected error is as defined in the theorem. In particular, we use the following strategy:

- Choose $\boldsymbol{\sigma} \in \{-1, +1\}^d$ uniformly at random.
- Use the distribution $\mathcal{D}_{\boldsymbol{\sigma}}$ over loss vectors \mathbf{x} , defined as follows: (x_1, \dots, x_d) has a Gaussian distribution $\mathcal{N}(\mu\boldsymbol{\sigma}, \frac{1}{16d}I)$ (where $\mu \leq \frac{1}{4\sqrt{d}}$ is a parameter to be chosen later), and x_0 is chosen independently according to a Gaussian distribution $\mathcal{N}(-\frac{1}{4}, \frac{1}{16})$. Note that this is different than the construction in the proof of Thm. 2, and the distribution of x_0 is not zero-mean. The idea is that $w_{t,0}$ will have to be close to 1 most of the time to get low regret, hence the points queried are far from and to one side of the origin.

First, let us verify that any $\mathcal{D}_{\boldsymbol{\sigma}}$ is a valid distribution. We have

$$\mathbb{E}[\|\bar{\mathbf{x}}\|_*] = \sup_{\mathbf{w} \in \mathcal{W}} |\langle \mathbb{E}[\mathbf{x}], \mathbf{w} \rangle| = \sup_{\mathbf{w} \in \mathcal{W}} \left| \frac{1}{4} + \mu \sum_{i=1}^d \sigma_i w_i \right| = \frac{1}{4} + \mu\sqrt{d} \leq \frac{1}{4} + \frac{1}{4} < 1.$$

Moreover, for any \mathbf{x} in the support of $\mathcal{D}_{\boldsymbol{\sigma}}$ and for any $\mathbf{w} \in \mathcal{W}$, by the Cauchy-Schwartz inequality,

$$|\langle \mathbf{x}, \mathbf{w} \rangle| \leq |x_0| + \left| \sum_{i=1}^d x_i w_i \right| \leq |x_0| + \sqrt{\sum_{i=1}^d x_i^2} \sqrt{\sum_{i=1}^d w_i^2} \leq |x_0| + \sqrt{\sum_{i=1}^d x_i^2}.$$

x_0 is normally distributed with mean $-1/4$ and variance $1/16$, and each x_i is independently normally distributed with mean at most $1/4\sqrt{d}$ and variance $1/16d$, from which it can be verified using Gaussian tail bounds that

$$\Pr\left(\sup_{\mathbf{w}} |\langle \mathbf{x}, \mathbf{w} \rangle| > z\right) \leq \Pr\left(|x_0| + \sqrt{\sum_{i=1}^d x_i^2} > z\right) \leq 2 \exp(-z^2/2)$$

for all $z \geq 1$ as required.

We now start the proof. Define the scalar and d -dimensional vectors

$$\hat{w}_0 = \frac{1}{T} \sum_{t=1}^T w_{t,0}, \quad \hat{\mathbf{w}} = \frac{1}{T} \sum_{t=1}^T (w_{t,1}, w_{t,2}, \dots, w_{t,d}), \quad \mathbf{w}^* = -\frac{1}{\sqrt{d}} \boldsymbol{\sigma}.$$

For any fixed $\mathcal{D}_{\boldsymbol{\sigma}}$, recall that $\bar{\mathbf{x}} = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\boldsymbol{\sigma}}}[\mathbf{x}] = (-\frac{1}{4}, \mu\sigma_1, \dots, \mu\sigma_d) = -\left(\frac{1}{4}, \mu\sqrt{d}w_1^*, \dots, \mu\sqrt{d}w_d^*\right)$. It is therefore easily verified that $(1, w_1^*, \dots, w_d^*)$ is a minimizer of $\langle \bar{\mathbf{x}}, \mathbf{w} \rangle$ over \mathcal{W} . Recalling the

formulation of regret in Eq. (3), we can write it as

$$\begin{aligned}
T \mathbb{E} \left[\left\langle \bar{\mathbf{x}}, \frac{1}{T} \sum_{t=1}^T \mathbf{w}_t \right\rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle \right] &= T \mathbb{E} \left[\frac{1}{4} (1 - \hat{w}_0) - \langle \mu\sqrt{d}\mathbf{w}^*, \hat{\mathbf{w}} \rangle + \langle \mu\sqrt{d}\mathbf{w}^*, \mathbf{w}^* \rangle \right] \\
&= T \mathbb{E} \left[\frac{1}{4} (1 - \hat{w}_0) + \mu\sqrt{d} \langle \mathbf{w}^*, \mathbf{w}^* - \hat{\mathbf{w}} \rangle \right] = T \mathbb{E} \left[\frac{1}{4} (1 - \hat{w}_0) + \mu\sqrt{d} \left(\|\mathbf{w}^*\|_2^2 - \langle \mathbf{w}^*, \hat{\mathbf{w}} \rangle \right) \right] \\
&= \frac{T}{4} \mathbb{E} [1 - \hat{w}_0] + \frac{\mu\sqrt{d}T}{2} \mathbb{E} [\|\mathbf{w}^*\|_2^2 + \|\mathbf{w}^*\|_2^2 - 2 \langle \mathbf{w}^*, \hat{\mathbf{w}} \rangle]. \tag{17}
\end{aligned}$$

By definition of the domain, each $(w_{t,1}, \dots, w_{t,d})$ has norm at most 1, and therefore $\hat{\mathbf{w}}$, which is their average, also has norm at most 1. Therefore, $\sum_{i=1}^d \hat{w}_i^2 \leq 1 = \|\mathbf{w}^*\|_2^2$, so we can lower bound the second expectation above by

$$\begin{aligned}
\mathbb{E} \left[\sum_{i=1}^d (w_i^*)^2 + \sum_{i=1}^d \hat{w}_i^2 - 2 \sum_{i=1}^d w_i^* \hat{w}_i \right] &= \mathbb{E} \left[\sum_{i=1}^d (w_i^* - \hat{w}_i)^2 \right] = \sum_{i=1}^d \mathbb{E} \left[\left(\frac{\sigma_i}{\sqrt{d}} + \hat{w}_i \right)^2 \right] \\
&\geq \sum_{i=1}^d \left(\frac{1}{\sqrt{d}} \right)^2 \Pr(\sigma_i \hat{w}_i \geq 0) = \frac{1}{d} \sum_{i=1}^d \Pr(\sigma_i \hat{w}_i \geq 0),
\end{aligned}$$

where $\Pr(\cdot)$ is the probability with respect to the joint randomness of σ and $\hat{\mathbf{w}}$. Plugging this back into Eq. (17), we get a regret lower bound of the form

$$\frac{T}{4} \mathbb{E} [1 - \hat{w}_0] + \frac{\mu T}{2\sqrt{d}} \sum_{i=1}^d \Pr(\sigma_i \hat{w}_i \geq 0).$$

These probabilities can now be lower bounded as in the proof of Thm. 2 (see the derivation following Eq. (9)), which imply that

$$\begin{aligned}
&\frac{T}{4} \mathbb{E} [1 - \hat{w}_0] + \frac{\mu T}{2\sqrt{d}} \sum_{i=1}^d \Pr(\sigma_i \hat{w}_i \geq 0) \\
&\geq \frac{T}{4} \mathbb{E} [1 - \hat{w}_0] + \frac{\mu\sqrt{d}T}{4} \left(1 - \frac{1}{d} \sum_{i=1}^d \int_{\mathbf{v}} |p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1)| d\mathbf{v} \right), \tag{18}
\end{aligned}$$

where $\mathbf{v} = (v_1, \dots, v_T)$ are the sequence of losses observed by the player over the T rounds, and p is the probability density function over \mathbf{v} . As in the proof of Thm. 2, we can now use Pinsker's inequality and the chain rule to upper the integral as follows:

$$\int_{\mathbf{v}} |p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1)| d\mathbf{v} \leq \sqrt{2 \sum_{t=1}^T D_{kl} (p(v_t|\sigma_i = 1, v_1 \dots v_{t-1}) || p(v_t|\sigma_i = -1, v_1 \dots v_{t-1}))}, \tag{19}$$

where D_{kl} is the KL divergence. Since the player is deterministic, any values v_1, \dots, v_{t-1} determine the point \mathbf{w}_t that the player will choose on round t . Thus, the distribution of $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ depends on the distribution of \mathbf{x}_t , as determined by \mathcal{D}_σ . By definition of \mathcal{D}_σ :

- Under the condition $\sigma_i = 1$, $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ has a Gaussian distribution with mean $-\frac{1}{4}w_{t,0} + \mu w_{t,i} + \sum_{j \in \{1 \dots d\} \setminus \{i\}} (\mu \sigma_j w_{t,j})$, and variance $\frac{1}{16}w_{t,0}^2 + \sum_{i=1}^d \frac{1}{16d}w_{t,i}^2$.
- Under the condition $\sigma_i = -1$, $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ has a Gaussian distribution with mean $-\frac{1}{4}w_{t,0} - \mu w_{t,i} + \sum_{j \in \{1 \dots d\} \setminus \{i\}} (\mu \sigma_j w_{t,j})$, and variance $\frac{1}{16}w_{t,0}^2 + \sum_{i=1}^d \frac{1}{16d}w_{t,i}^2$.

By a standard result on the KL divergence of two Gaussian distributions, we therefore have that Eq. (19) is at most

$$\sqrt{2 \sum_{t=1}^T \mathbb{E} \left[\frac{32\mu^2 w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \middle| \sigma_i = 1 \right]}. \quad (20)$$

As in the proof of Thm. 2, we can also upper bound Eq. (19) using the reverse order of probabilities:

$$\int_{\mathbf{v}} |p(\mathbf{v} | \sigma_i = -1) - p(\mathbf{v} | \sigma_i = 1)| d\mathbf{v} \leq \sqrt{2 \sum_{t=1}^T D_{kl}(p(v_t | \sigma_i = -1, v_1 \dots v_{t-1}) || p(v_t | \sigma_i = 1, v_1 \dots v_{t-1}))},$$

which leads to Eq. (19) being upper bounded by

$$\sqrt{2 \sum_{t=1}^T \mathbb{E} \left[\frac{32\mu^2 w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \middle| \sigma_i = -1 \right]}. \quad (21)$$

Combining Eq. (20) and Eq. (21) and using the elementary inequality $\min\{\sqrt{a}, \sqrt{b}\} \leq \sqrt{\frac{a+b}{2}}$ and the fact that σ_i is uniformly distributed on $\{-1, +1\}$, we get

$$\begin{aligned} & \int_{\mathbf{v}} |p(\mathbf{v} | \sigma_i = -1) - p(\mathbf{v} | \sigma_i = 1)| d\mathbf{v} \\ & \leq \min \left\{ \sqrt{2 \sum_{t=1}^T \mathbb{E} \left[\frac{32\mu^2 w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \middle| \sigma_i = -1 \right]}, \sqrt{2 \sum_{t=1}^T \mathbb{E} \left[\frac{32\mu^2 w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \middle| \sigma_i = 1 \right]} \right\} \\ & \leq \sqrt{64\mu^2 \sum_{t=1}^T \frac{1}{2} \left(\mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \middle| \sigma_i = -1 \right] + \mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \middle| \sigma_i = 1 \right] \right)} \\ & = 8\mu \sqrt{\sum_{t=1}^T \mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \right]}. \quad (22) \end{aligned}$$

Plugging this back into Eq. (18), we get an expected error lower bound of

$$\frac{T}{4} \mathbb{E} [1 - \hat{w}_0] + \frac{\mu \sqrt{dT}}{4} \left(1 - \frac{1}{d} \sum_{i=1}^d 8\mu \sqrt{\sum_{t=1}^T \mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \right]} \right),$$

which by Jensen's inequality is at least

$$\begin{aligned}
& \frac{T}{4} \mathbb{E}[1 - \hat{w}_0] + \frac{\mu\sqrt{dT}}{4} \left(1 - 8\mu \sqrt{\frac{1}{d} \sum_{i=1}^d \sum_{t=1}^T \mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \right]} \right) \\
&= \frac{T}{4} \mathbb{E}[1 - \hat{w}_0] + \frac{\mu\sqrt{dT}}{4} \left(1 - 8\mu \sqrt{\frac{1}{d} \sum_{t=1}^T \mathbb{E} \left[\frac{\sum_{i=1}^d w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d} w_{t,i}^2} \right]} \right) \\
&= \frac{T}{4} \mathbb{E}[1 - \hat{w}_0] + \frac{\mu\sqrt{dT}}{4} \left(1 - 8\mu \sqrt{\frac{1}{d} \sum_{t=1}^T \mathbb{E} \left[\frac{1}{\frac{w_{t,0}^2}{\sum_{i=1}^d w_{t,i}^2} + \frac{1}{d}} \right]} \right).
\end{aligned}$$

By definition of \mathcal{W} , we have that $\sum_{i=1}^d w_{t,i}^2 \leq 1$ always. Using this and recalling that $\hat{w}_0 = \frac{1}{T} \sum_{t=1}^T w_{t,0}$, we can lower bound the above by

$$\frac{1}{4} \sum_{t=1}^T \mathbb{E}[1 - w_{t,0}] + \frac{\mu\sqrt{dT}}{4} \left(1 - 8\mu \sqrt{\frac{1}{d} \sum_{t=1}^T \mathbb{E} \left[\frac{1}{w_{t,0}^2 + \frac{1}{d}} \right]} \right). \quad (23)$$

The trick now is to argue that not matter what are the values of $w_{t,0}$, this lower bound will be large: Either $w_{t,0}$ will tend to be bounded away from 1, and then Eq. (23) will be large due to the first term; or that $w_{t,0}$ will tend to be very close to 1, but then Eq. (23) will be large due to the second term. To make this precise, we use the following technical lemma:

Lemma 7 *For any integer $d \geq 1$, and for any $w \in [-1, 1]$, it holds that $\frac{1}{w^2 + \frac{1}{d}} \leq d(1 - |w|) + 1$.*

Proof Since both sides of the inequality are the same for w and $-w$, we can assume without loss of generality, that $w \in [0, 1]$, and prove that $\frac{1}{w^2 + \frac{1}{d}} \leq d(1 - w) + 1$ over this domain.

By algebraic manipulations, this inequality is equivalent to the assertion that $(d+1)w^2 - dw^3 - w \geq -\frac{1}{d}$, or equivalently,

$$w(1-w)(dw-1) \geq -\frac{1}{d}.$$

Considering the domain of w , the left hand side is non-positive only when $w \in [0, 1/d]$, because of the third term, but in that regime $(dw-1) \in [-1, 0]$ and $(1-w) \in [0, 1]$, and therefore

$$w(1-w)(dw-1) \geq w * 1 * (-1) = -w \geq -\frac{1}{d}$$

as required. ■

Applying this inequality, we can lower bound Eq. (23) by

$$\begin{aligned}
 & \frac{1}{4} \sum_{t=1}^T \mathbb{E} [1 - |w_{t,0}|] + \frac{\mu\sqrt{dT}}{4} \left(1 - 8\mu \sqrt{\frac{1}{d} \sum_{t=1}^T \mathbb{E} [d(1 - |w_{t,0}|) + 1]} \right) \\
 &= \frac{1}{4} \sum_{t=1}^T \mathbb{E} [1 - |w_{t,0}|] + \frac{\mu\sqrt{dT}}{4} \left(1 - 8\mu \sqrt{\frac{T}{d} + \sum_{t=1}^T \mathbb{E} [1 - |w_{t,0}|]} \right) \\
 &\geq \min_{z \in [0, T]} \left(\frac{z}{4} + \frac{\mu\sqrt{dT}}{4} \left(1 - 8\mu \sqrt{\frac{T}{d} + z} \right) \right),
 \end{aligned}$$

Recalling that μ is a free parameter, let us choose $\mu = \frac{1}{16} \sqrt{\frac{d}{T}}$ (which satisfies the assumption $\mu \leq \frac{1}{4\sqrt{d}}$, since we assume $T \geq d^4/16$). We therefore get

$$\frac{1}{4} \min_{z \in [0, T]} \left(z + \frac{1}{16} d\sqrt{T} \left(1 - \frac{1}{2} \sqrt{1 + \frac{d}{T}z} \right) \right). \quad (24)$$

Note that this expression is convex with respect to z , and by differentiating, has an extremal point at

$$z = 2^{-12} d^3 - \frac{T}{d}.$$

Since we assume $T \geq d^4/16$, this point is non-positive, and therefore the minimum in Eq. (24) is attained at $z = 0$, where it equals

$$\frac{1}{64} d\sqrt{T} \left(1 - \frac{1}{2} \right) = \frac{1}{128} d\sqrt{T}.$$

Since this lower bound holds for a domain in \mathbb{R}^{d+1} for any $d > 0$, we get a $\frac{1}{128}(d-1)\sqrt{T}$ lower bound for a domain in \mathbb{R}^d for any $d > 1$, as required.

A.3. Proof of Thm. 5

By Yao's minimax principle, it is sufficient to provide a randomized strategy to choose a loss vector distribution \mathcal{D} , such that for any deterministic player, the expected error is as defined in the theorem. In particular, we use the following strategy:

- Choose $J \in \{1, \dots, d\}$ uniformly at random.
- Use the distribution \mathcal{D}_J over loss vectors \mathbf{x} , defined as $\mathcal{N}(-\mu \mathbf{e}_J, \frac{1}{4}I)$ (where $\mu \leq \frac{1}{2}$ is a parameter to be chosen later).

First, let us verify that any \mathcal{D}_J is a valid distribution. We have

$$\mathbb{E}[\|\bar{\mathbf{x}}\|_*] = \sup_{\mathbf{w} \in \mathcal{W}} |\langle \mathbb{E}[\mathbf{x}], \mathbf{w} \rangle| = \sup_{\mathbf{w} \in \mathcal{W}} |\mu w_j| \leq \mu \leq \frac{1}{2},$$

and moreover, for any \mathbf{x} in the support of \mathcal{D}_σ and for any $\mathbf{w} \in \mathcal{W}$, $\langle \mathbf{x}, \mathbf{w} \rangle$ is Gaussian with mean $\mu w_j \in [-\frac{1}{2}, \frac{1}{2}]$ and variance $\frac{1}{4} \sum_{i=1}^d w_i^2 \leq \frac{1}{4} \sum_{i=1}^d |w_i| \leq \frac{1}{4}$, from which it is easily verified that $\Pr(\sup_{\mathbf{w}} |\langle \mathbf{x}, \mathbf{w} \rangle| > z) \leq 2 \exp(-z^2/2)$ for all $z \geq 1$ as required.

We now start the proof. For any fixed \mathcal{D}_σ , recall that $\bar{\mathbf{x}} = -\mu \mathbf{e}_j$, and define $\mathbf{w}^* = \mathbf{e}_j$. It is easily verified that \mathbf{w}^* is a minimizer of $\langle \bar{\mathbf{x}}, \mathbf{w} \rangle$ over \mathcal{W} . Also, let \mathbb{E}_j denote expectation assuming that $J = j$, and recall that J is chosen uniformly at random from $\{1, \dots, d\}$. Therefore,

$$\mathbb{E} \left[\langle \bar{\mathbf{x}}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle \right] = \frac{1}{d} \sum_{j=1}^d \mathbb{E}_j [\mu - \mu \mathbb{E}[\hat{w}_j]] = \mu - \frac{\mu}{d} \sum_{j=1}^d \mathbb{E}_j[\hat{w}_j].$$

Define the reference distribution $\mathcal{D}_0 = \mathcal{N}(\mathbf{0}, \frac{1}{2}I)$ over loss vectors, which has zero-mean, and let \mathbb{E}_0 denote expectation assuming the loss vectors are chosen from that distribution. Then we can lower bound the above by

$$\begin{aligned} \mu - \frac{\mu}{d} \sum_{j=1}^d (\mathbb{E}_0[\hat{w}_j] + |\mathbb{E}_0[\hat{w}_j] - \mathbb{E}_j[\hat{w}_j]|) &\geq \mu - \frac{\mu}{d} \left(\mathbb{E}_0 \left[\sum_{j=1}^d |\hat{w}_j| \right] + \sum_{j=1}^d |\mathbb{E}_0[\hat{w}_j] - \mathbb{E}_j[\hat{w}_j]| \right) \\ &\geq \left(1 - \frac{1}{d}\right) \mu - \frac{\mu}{d} \sum_{j=1}^d |\mathbb{E}_0[\hat{w}_j] - \mathbb{E}_j[\hat{w}_j]|, \end{aligned} \quad (25)$$

where in the last inequality we used the fact that \mathbf{w} has 1-norm at most 1. Now, recall that we assume that the player is deterministic, hence $\hat{\mathbf{w}} = \hat{\mathbf{w}}(\mathbf{v})$ is a deterministic function of the sequence of loss observations $\mathbf{v} = (v_1, \dots, v_T)$ made by the player over T rounds. Letting $p_0(\mathbf{v})$ and $p_j(\mathbf{v})$ denote the density functions with respect to $\mathcal{D}_0, \mathcal{D}_j$ respectively, we can upper bound the expectation differences in Eq. (25) as follows:

$$\begin{aligned} \sum_{j=1}^d |\mathbb{E}_0[\hat{w}_j] - \mathbb{E}_j[\hat{w}_j]| &= \sum_{j=1}^d \left| \int_{\mathbf{v}} \hat{w}_j(\mathbf{v}) (p_0(\mathbf{v}) - p_j(\mathbf{v})) d\mathbf{v} \right| \\ &\leq \sum_{j=1}^d \int_{\mathbf{v}} |\hat{w}_j(\mathbf{v})| |p_0(\mathbf{v}) - p_j(\mathbf{v})| d\mathbf{v} \\ &\leq \sum_{j=1}^d \int_{\mathbf{v}} |p_0(\mathbf{v}) - p_j(\mathbf{v})| d\mathbf{v}. \end{aligned}$$

By Pinsker's inequality, and the chain rule, this can be upper bounded by

$$\sum_{j=1}^d \sqrt{2 \sum_{t=1}^T D_{kl}(p_0(v_t|v_1 \dots v_{t-1}) || p_j(v_t|v_1 \dots v_{t-1}))}, \quad (26)$$

where $D_{kl}(p||q)$ is the Kullback-Leibler (KL) divergence between p and q . Since the player is deterministic, any values v_1, \dots, v_{t-1} determine the point \mathbf{w}_t that the player will choose on round t . Thus, the distribution of $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ depends on the distribution of \mathbf{x}_t :

- Under p_j (corresponding to \mathcal{D}_j), $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ has a Gaussian distribution with mean $-\mu w_{t,j}$ and variance $\frac{1}{4} \sum_{i=1}^d w_{t,i}^2 = \frac{1}{4} \|\mathbf{w}_t\|_2^2$.

- Under p_0 (corresponding to \mathcal{D}_0), $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ has a Gaussian distribution with mean 0 and variance $\frac{1}{4} \|\mathbf{w}_t\|_2$.

By a standard result on the KL divergence of two Gaussian distributions, we therefore have that Eq. (26) is at most

$$\sum_{j=1}^d \sqrt{2 \sum_{t=1}^T \mathbb{E}_0 \left[\frac{2\mu^2 w_{t,j}^2}{\|\mathbf{w}_t\|_2} \right]}.$$

By Jensen's inequality, this is at most

$$\begin{aligned} d \frac{1}{d} \sum_{j=1}^d \sqrt{2 \sum_{t=1}^T \mathbb{E}_0 \left[\frac{2\mu^2 w_{t,j}^2}{\|\mathbf{w}_t\|_2} \right]} &\leq d \sqrt{\frac{2}{d} \sum_{j=1}^d \sum_{t=1}^T \mathbb{E}_0 \left[\frac{2\mu^2 w_{t,j}^2}{\|\mathbf{w}_t\|_2} \right]} \\ &= d \sqrt{\frac{2}{d} \sum_{t=1}^T \mathbb{E}_0 \left[\frac{2 \sum_{j=1}^d \mu^2 w_{t,j}^2}{\|\mathbf{w}_t\|_2} \right]} = d \sqrt{\frac{2}{d} \sum_{t=1}^T (2\mu^2)} = 2\mu \sqrt{dT}. \end{aligned}$$

Plugging this back into Eq. (25), we get an error lower bound of

$$\left(1 - \frac{1}{d}\right) \mu - 2\mu^2 \sqrt{\frac{T}{d}} = \mu \left(1 - \frac{1}{d} - \mu \sqrt{\frac{T}{d}}\right) \geq \mu \left(\frac{1}{2} - \mu \sqrt{\frac{T}{d}}\right), \quad (27)$$

where we used the fact that $d > 1$. Let us now consider two cases:

- If $T \geq \frac{d}{4}$, then we choose $\mu = \frac{1}{4} \sqrt{\frac{d}{T}}$ (which ensures the condition $\mu \leq \frac{1}{2}$), and Eq. (27) equals $\frac{1}{4} \sqrt{\frac{d}{T}} \left(\frac{1}{2} - \frac{1}{4}\right) = \frac{1}{16} \sqrt{\frac{d}{T}}$.
- If $T < \frac{d}{4}$, then we choose $\mu = \frac{1}{2}$, and Eq. (27) is lower bounded by $\frac{1}{2} \left(\frac{1}{2} - \frac{1}{2} \sqrt{\frac{1}{4}}\right) = \frac{1}{8}$.

Therefore, by choosing μ appropriately, we can lower bound Eq. (16) by $\min \left\{ \frac{1}{8}, \frac{1}{16} \sqrt{\frac{d}{T}} \right\} \geq \frac{1}{16} \min \left\{ 1, \sqrt{\frac{d}{T}} \right\}$ as required.

A.4. Proof of Thm. 6

As in the proof of Thm. 3, to simplify notation, we assume the game takes place in \mathbb{R}^{d+1} for some $d > 0$, so the domain is $\mathcal{W} = [-1, +1] \times [-1, +1]^d$. We denote the first coordinate as coordinate 0, and the other coordinates as 1, 2, ..., d .

The proof uses a similar loss vector distribution as that of Thm. 3, but with a different scaling to ensure a valid distribution.

By Yao's minimax principle, it is sufficient to provide a randomized strategy to choose a loss vector distribution \mathcal{D} , such that for any deterministic player, the expected error is as defined in the theorem. In particular, we use the following strategy:

- Choose $\sigma \in \{-1, +1\}^d$ uniformly at random.

- Use the distribution \mathcal{D}_σ over loss vectors \mathbf{x} , defined as follows: (x_1, \dots, x_d) has a Gaussian distribution $\mathcal{N}(\mu\sigma, \frac{1}{16d^2}I)$ (where $\mu \leq \frac{1}{4d}$ is a parameter to be chosen later), and x_0 is chosen independently according to a Gaussian distribution $\mathcal{N}(-\frac{1}{4}, \frac{1}{16})$.

First, let us verify that any \mathcal{D}_σ is a valid distribution. We have

$$\mathbb{E}[\|\bar{\mathbf{x}}\|_*] = \sup_{\mathbf{w} \in \mathcal{W}} |\langle \mathbb{E}[\mathbf{x}], \mathbf{w} \rangle| = \sup_{\mathbf{w} \in \mathcal{W}} \left| \frac{1}{4} + \mu \sum_{i=1}^d \sigma_i w_i \right| = \frac{1}{4} + \mu d \leq \frac{1}{4} + \frac{1}{4} < 1.$$

Moreover, for any \mathbf{x} in the support of \mathcal{D}_σ and for any $\mathbf{w} \in \mathcal{W}$, by Hölder's inequality,

$$|\langle \mathbf{x}, \mathbf{w} \rangle| \leq |x_0| + \sum_{i=1}^d |x_i|.$$

x_0 is normally distributed with mean $-1/4$ and variance $1/16$, and each x_i is independently normally distributed with mean at most $1/4d$ and variance $1/16d^2$, from which it can be verified using Gaussian tail bounds that

$$\Pr \left(\sup_{\mathbf{w}} |\langle \mathbf{x}, \mathbf{w} \rangle| > z \right) \leq \Pr \left(|x_0| + \sum_{i=1}^d |x_i| > z \right) \leq 2 \exp(-z^2/2)$$

for all $z \geq 1$ as required.

We now start the proof. Define the scalar and d -dimensional vectors

$$\hat{w}_0 = \frac{1}{T} \sum_{t=1}^T w_{t,0}, \quad \hat{\mathbf{w}} = \frac{1}{T} \sum_{t=1}^T (w_{t,1}, w_{t,2}, \dots, w_{t,d}), \quad \mathbf{w}^* = -\boldsymbol{\sigma}.$$

For any fixed \mathcal{D}_σ , recall that $\bar{\mathbf{x}} = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_\sigma}[\mathbf{x}] = (-\frac{1}{4}, \mu\sigma_1, \dots, \mu\sigma_d)$. It is therefore easily verified that $(1, \sigma_1, \dots, \sigma_d)$ is a minimizer of $\langle \bar{\mathbf{x}}, \mathbf{w} \rangle$ over \mathcal{W} . Recalling the formulation of regret in Eq. (3), we can write it as

$$\begin{aligned} T \mathbb{E} \left[\left\langle \bar{\mathbf{x}}, \frac{1}{T} \sum_{t=1}^T \mathbf{w}_t \right\rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}, \mathbf{w} \rangle \right] &= T \mathbb{E} \left[\frac{1}{4} (1 - \hat{w}_0) + \mu d - \mu \sum_{i=1}^d (-\sigma_i) \hat{w}_i \right] \\ &= \frac{T}{4} \mathbb{E}[1 - \hat{w}_0] + \mu T \left(d - \sum_{i=1}^d \mathbb{E}[(-\sigma_i) \hat{w}_i] \right) \end{aligned} \quad (28)$$

Recalling that each σ_i is uniformly distributed on $\{-1, +1\}$, the sum of expectations can be upper bounded as follows:

$$\begin{aligned} \sum_{i=1}^d \mathbb{E}[(-\sigma_i) \hat{w}_i] &= \frac{1}{2} \sum_{i=1}^d (\mathbb{E}[\hat{w}_i | \sigma_i = -1] + \mathbb{E}[-\hat{w}_i | \sigma_i = 1]) \\ &= \frac{1}{2} \sum_{i=1}^d (\mathbb{E}[\hat{w}_i | \sigma_i = -1] - \mathbb{E}[\hat{w}_i | \sigma_i = 1]). \end{aligned}$$

Since the player's strategy is assumed to be deterministic, then $\hat{w}_i = \hat{w}_i(\mathbf{v})$ is a deterministic function of the sequence of losses $\mathbf{v} = (v_1, \dots, v_T)$ observed by the player over the T rounds. Letting $p(\cdot)$ denote the probability density function of \mathbf{v} , we can rewrite the above as

$$\begin{aligned} \frac{1}{2} \sum_{i=1}^d \int_{\mathbf{v}} \hat{w}_i(\mathbf{v})(p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1))d\mathbf{v} &\leq \frac{1}{2} \sum_{i=1}^d \int_{\mathbf{v}} |\hat{w}_i(\mathbf{v})|(p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1))|d\mathbf{v} \\ &\leq \frac{1}{2} \sum_{i=1}^d \int_{\mathbf{v}} |(p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1))|d\mathbf{v}, \end{aligned} \quad (29)$$

where we used the fact that $\hat{w}_i \in [-1, +1]$ by the domain assumptions. We can now use Pinsker's inequality and the chain rule to upper the integral as follows:

$$\int_{\mathbf{v}} |p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1)|d\mathbf{v} \leq \sqrt{2 \sum_{t=1}^T D_{kl}(p(v_t|\sigma_i = 1, v_1 \dots v_{t-1}) || p(v_t|\sigma_i = -1, v_1 \dots v_{t-1}))}, \quad (30)$$

where D_{kl} is the KL divergence. Since the player is deterministic, any values v_1, \dots, v_{t-1} determine the point \mathbf{w}_t that the player will choose on round t . Thus, the distribution of $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ depends on the distribution of \mathbf{x}_t , as determined by \mathcal{D}_σ . By definition of \mathcal{D}_σ :

- Under the condition $\sigma_i = 1$, $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ has a Gaussian distribution with mean $-\frac{1}{4}w_{t,0} + \mu w_{t,i} + \sum_{j \in \{1 \dots d\} \setminus \{i\}} (\mu \sigma_j w_{t,j})$, and variance $\frac{1}{16}w_{t,0}^2 + \sum_{i=1}^d \frac{1}{16d^2}w_{t,i}^2$.
- Under the condition $\sigma_i = -1$, $v_t = \langle \mathbf{x}_t, \mathbf{w}_t \rangle$ has a Gaussian distribution with mean $-\frac{1}{4}w_{t,0} - \mu w_{t,i} + \sum_{j \in \{1 \dots d\} \setminus \{i\}} (\mu \sigma_j w_{t,j})$, and variance $\frac{1}{16}w_{t,0}^2 + \sum_{i=1}^d \frac{1}{16d^2}w_{t,i}^2$.

By a standard result on the KL divergence of two Gaussian distributions, we therefore have that Eq. (30) is at most

$$\sqrt{2 \sum_{t=1}^T \mathbb{E} \left[\frac{32\mu^2 w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \middle| \sigma_i = 1 \right]}. \quad (31)$$

Applying Pinsker's inequality on the reverse order of probabilities in Eq. (30), we can also upper Eq. (30) by

$$\int_{\mathbf{v}} |p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1)|d\mathbf{v} \leq \sqrt{2 \sum_{t=1}^T D_{kl}(p(v_t|\sigma_i = -1, v_1 \dots v_{t-1}) || p(v_t|\sigma_i = 1, v_1 \dots v_{t-1}))},$$

which leads to Eq. (30) being upper bounded by

$$\sqrt{2 \sum_{t=1}^T \mathbb{E} \left[\frac{32\mu^2 w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \middle| \sigma_i = -1 \right]}. \quad (32)$$

Combining Eq. (31) and Eq. (32) and using the elementary inequality $\min\{\sqrt{a}, \sqrt{b}\} \leq \sqrt{\frac{a+b}{2}}$ and the fact that σ_i is uniformly distributed on $\{-1, +1\}$, we get

$$\begin{aligned}
& \int_{\mathbf{v}} |p(\mathbf{v}|\sigma_i = -1) - p(\mathbf{v}|\sigma_i = 1)| d\mathbf{v} \\
& \leq \min \left\{ \sqrt{2 \sum_{t=1}^T \mathbb{E} \left[\frac{32\mu^2 w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \middle| \sigma_i = -1 \right]}, \sqrt{2 \sum_{t=1}^T \mathbb{E} \left[\frac{32\mu^2 w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \middle| \sigma_i = 1 \right]} \right\} \\
& \leq \sqrt{64\mu^2 \sum_{t=1}^T \frac{1}{2} \left(\mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \middle| \sigma_i = -1 \right] + \mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \middle| \sigma_i = 1 \right] \right)} \\
& = 8\mu \sqrt{\sum_{t=1}^T \mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \right]}.
\end{aligned}$$

Plugging this back into Eq. (29), and that in turn into Eq. (28), we get an expected error lower bound of

$$\frac{T}{4} \mathbb{E}[1 - \hat{w}_0] + \mu T \left(d - 4\mu \sum_{i=1}^d \sqrt{\sum_{t=1}^T \mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \right]} \right),$$

which by Jensen's inequality is at least

$$\begin{aligned}
& \frac{T}{4} \mathbb{E}[1 - \hat{w}_0] + \mu T \left(d - 4\mu d \sqrt{\frac{1}{d} \sum_{i=1}^d \sum_{t=1}^T \mathbb{E} \left[\frac{w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \right]} \right) \\
& = \frac{T}{4} \mathbb{E}[1 - \hat{w}_0] + \mu T d \left(1 - 4\mu \sqrt{\frac{1}{d} \sum_{t=1}^T \mathbb{E} \left[\frac{\sum_{i=1}^d w_{t,i}^2}{w_{t,0}^2 + \sum_{i=1}^d \frac{1}{d^2} w_{t,i}^2} \right]} \right) \\
& = \frac{T}{4} \mathbb{E}[1 - \hat{w}_0] + \mu T d \left(1 - 4\mu \sqrt{\frac{1}{d} \sum_{t=1}^T \mathbb{E} \left[\frac{1}{\frac{w_{t,0}^2}{\sum_{i=1}^d w_{t,i}^2} + \frac{1}{d^2}} \right]} \right).
\end{aligned}$$

By definition of \mathcal{W} , we have that $\sum_{i=1}^d w_{t,i}^2 \leq \sum_{i=1}^d |w_{t,i}| \leq d$ always. Using this and recalling that $\hat{w}_0 = \frac{1}{T} \sum_{t=1}^T w_{t,0}$, we can lower bound the above by

$$\frac{1}{4} \sum_{t=1}^T \mathbb{E}[1 - w_{t,0}] + \mu T d \left(1 - 4\mu \sqrt{\sum_{t=1}^T \mathbb{E} \left[\frac{1}{w_{t,0}^2 + \frac{1}{d}} \right]} \right).$$

Using Lemma 7 from the proof of Thm. 3, this expression can be lower bounded by

$$\begin{aligned}
 & \frac{1}{4} \sum_{t=1}^T \mathbb{E} [1 - |w_{t,0}|] + \mu T d \left(1 - 4\mu \sqrt{\sum_{t=1}^T \mathbb{E} [d(1 - |w_{t,0}|) + 1]} \right) \\
 &= \frac{1}{4} \sum_{t=1}^T \mathbb{E} [1 - |w_{t,0}|] + \mu T d \left(1 - 4\mu \sqrt{T + d \sum_{t=1}^T \mathbb{E} [1 - |w_{t,0}|]} \right) \\
 &\geq \min_{z \in [0, T]} \left(\frac{z}{4} + \mu T d \left(1 - 4\mu \sqrt{T + dz} \right) \right),
 \end{aligned}$$

Recalling that μ is a free parameter, let us choose $\mu = \frac{1}{8} \sqrt{\frac{1}{T}}$ (which satisfies the assumption $\mu \leq \frac{1}{4d}$, since we assume $T \geq d^4/4$). We therefore get

$$\min_{z \in [0, T]} \left(\frac{z}{4} + \frac{1}{8} d \sqrt{T} \left(1 - \frac{1}{2} \sqrt{1 + \frac{d}{T} z} \right) \right). \quad (33)$$

Note that this expression is convex with respect to z , and by differentiating, has an extremal point at

$$z = \frac{1}{64} d^3 - \frac{T}{d}.$$

Since we assume $T \geq d^4/4$, this point is non-positive, and therefore the minimum in Eq. (33) is attained at $z = 0$, where it equals

$$\frac{1}{8} d \sqrt{T} \left(1 - \frac{1}{2} \right) = \frac{1}{16} d \sqrt{T}.$$

Since this lower bound holds for a domain in \mathbb{R}^{d+1} for any $d > 0$, we get a $\frac{1}{16}(d-1)\sqrt{T}$ lower bound for a domain in \mathbb{R}^d for any $d > 1$, as required.

Appendix B. From Sub-Gaussian to Bounded Distributions

As discussed in Sec. 2, our results use sub-Gaussian distributions rather than the more standard bounded distribution assumption. In this appendix, we explain why this is really without any loss of generality, and that lower bounds for the sub-Gaussian setting can be readily converted to the bounded setting, at the cost of a $\sqrt{\log(T)}$ factor.

In particular, suppose there is a sub-Gaussian distribution \mathcal{D} over the cost vectors, for which any player's strategy incurs expected error/regret at least R . The trick is to consider a ‘‘shrunk’’ distribution, which simply re-scales all cost vectors by $\Theta(1/\sqrt{\log(T)})$. Then the regret will be $\Omega(R/\sqrt{\log(T)})$, and because the distribution is sub-Gaussian, then with very high probability all the T cost vectors $\mathbf{x}_1, \dots, \mathbf{x}_T$ satisfy $\|\mathbf{x}_t\|_* \leq 1$. This means that even if we modify the distribution to force it to be bounded, the expected error regret will still be $\Omega(R/\sqrt{\log(T)})$. We note that a similar technique was used implicitly in Dekel et al. (2014).

A more formal result can be stated as follows:

Theorem 8 *Suppose that there exists a distribution \mathcal{D} satisfying $\|\mathbb{E}_{\mathbf{x} \sim \mathcal{D}}[\mathbf{x}]\|_* = 1$ and $\Pr_{\mathbf{x} \sim \mathcal{D}}(\|\mathbf{x}\|_* > z) \leq 2 \exp(-z^2/2)$ for all $z \geq 1$, such that for any player's strategy, the expected error/regret is at least R . Assuming $T > 1$, there exists a distribution \mathcal{D}' satisfying $\Pr_{\mathbf{x} \sim \mathcal{D}'}(\|\mathbf{x}\|_* \leq 1) = 1$, for which the expected error/regret of any player is at least $c \left(\frac{R}{\sqrt{\log(T)}} - \frac{1}{T^4} \right)$ for some universal constant c .*

Since the regret R is virtually always at least $\Omega(1/T)$, this means that we get the same error/regret up to constant and $\sqrt{\log(T)}$ factors. The $1/T^4$ term can be replaced by $1/T^s$ for arbitrarily large s , at the cost of affecting the c constant.

Proof [Proof Sketch] Let p be a parameter to be determined later. Given the distribution \mathcal{D} , we algorithmically define \mathcal{D}' and an auxiliary distribution $\hat{\mathcal{D}}$ as follows:

- $\hat{\mathcal{D}}$ samples \mathbf{x} according to \mathcal{D} and returns $\frac{1}{p\sqrt{\log(T)}}\mathbf{x}$. Let $\bar{\mathbf{x}}_{\hat{\mathcal{D}}} = \mathbb{E}_{\hat{\mathcal{D}}}[\mathbf{x}]$.
- \mathcal{D}' samples \mathbf{x} according to $\hat{\mathcal{D}}$, and returns \mathbf{x} if $\|\mathbf{x}\|_* \leq \frac{1}{2}$, and otherwise returns $\mathbb{E}_{\hat{\mathcal{D}}}[\mathbf{x} \mid \|\mathbf{x}\|_* \geq \frac{1}{2}]$. Let $\bar{\mathbf{x}}_{\mathcal{D}'} = \mathbb{E}_{\mathcal{D}'}[\mathbf{x}]$.

It is easily verified that $\bar{\mathbf{x}}_{\hat{\mathcal{D}}} = \bar{\mathbf{x}}_{\mathcal{D}'}$. Moreover, if p is large enough, then \mathcal{D}' always returns a vector such that $\|\mathbf{x}\|_* \leq 1$, and thus satisfies the lemma's requirement. In that case, we also have $\|\bar{\mathbf{x}}_{\hat{\mathcal{D}}}\|_* = \|\bar{\mathbf{x}}_{\mathcal{D}'}\|_* \leq 1$.

It remains to show the error/regret lower bound. We sketch the proof for the expected error - the proof for the expected regret is identical by replacing $\hat{\mathbf{w}}$ by $\frac{1}{T} \sum_{t=1}^T \mathbf{w}_t$. By the lemma's assumptions, we know that

$$\mathbb{E}_{\mathcal{D}} \left[\langle \bar{\mathbf{x}}_{\mathcal{D}}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\mathcal{D}}, \mathbf{w} \rangle \right] \geq R,$$

where $\mathbb{E}_{\mathcal{D}}$ signifies expectation with respect to drawing cost vectors from \mathcal{D} , and $\bar{\mathbf{x}}_{\mathcal{D}} = \mathbb{E}_{\mathcal{D}}[\mathbf{x}]$. Since $\hat{\mathcal{D}}$ simply scales the vectors drawn from \mathcal{D} by a fixed factor $\frac{1}{p\sqrt{\log(T)}}$, we have

$$\mathbb{E}_{\hat{\mathcal{D}}} \left[\langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \mathbf{w} \rangle \right] \geq \frac{R}{p\sqrt{\log(T)}}.$$

Now, let A be the event that for all cost vectors $\mathbf{x}_1, \dots, \mathbf{x}_T$, we have $\|\mathbf{x}_t\|_* \leq \frac{1}{2}$. By the assumptions on \mathcal{D} , $\hat{\mathcal{D}}$ and a union bound, $\Pr_{\hat{\mathcal{D}}}(\neg A) \leq 2T^{1-p^2/8}$ for all sufficiently large p . From the displayed equation above, it follows that

$$\begin{aligned} \frac{R}{p\sqrt{\log(T)}} &\leq \mathbb{E}_{\hat{\mathcal{D}}} \left[\langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \mathbf{w} \rangle \right] \\ &= \Pr_{\hat{\mathcal{D}}}(A) \mathbb{E}_{\hat{\mathcal{D}}} \left[\langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \mathbf{w} \rangle \middle| A \right] + \Pr_{\hat{\mathcal{D}}}(\neg A) \mathbb{E}_{\hat{\mathcal{D}}} \left[\langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \mathbf{w} \rangle \middle| \neg A \right] \\ &\leq \mathbb{E}_{\hat{\mathcal{D}}} \left[\langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \mathbf{w} \rangle \middle| A \right] + 2T^{1-p^2/8} * 2 \max_{\mathbf{w} \in \mathcal{W}} |\langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \mathbf{w} \rangle| \\ &\leq \mathbb{E}_{\hat{\mathcal{D}}} \left[\langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\hat{\mathcal{D}}}, \mathbf{w} \rangle \middle| A \right] + 4T^{1-p^2/8}, \end{aligned}$$

where we used the assumption that $\max_{\mathbf{w} \in \mathcal{W}} |\langle \mathbf{w}, \bar{\mathbf{x}}_{\hat{\mathcal{D}}} \rangle| = \|\bar{\mathbf{x}}_{\hat{\mathcal{D}}}\|_* \leq 1$. Switching sides and using the assumption that $\bar{\mathbf{x}}_{\hat{\mathcal{D}}} = \bar{\mathbf{x}}_{\mathcal{D}'}$, we get that

$$\mathbb{E}_{\hat{\mathcal{D}}} \left[\langle \bar{\mathbf{x}}_{\mathcal{D}'}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\mathcal{D}'}, \mathbf{w} \rangle \middle| A \right] \geq \frac{R}{p\sqrt{\log(T)}} - 4T^{1-p^2/8}.$$

But conditioned on A , the distribution of $\mathbf{x}_1, \dots, \mathbf{x}_T$ is the same under $\hat{\mathcal{D}}$ and \mathcal{D}' , and therefore

$$\mathbb{E}_{\mathcal{D}'} \left[\langle \bar{\mathbf{x}}_{\mathcal{D}'}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\mathcal{D}'}, \mathbf{w} \rangle \middle| A' \right] \geq \frac{R}{p\sqrt{\log(T)}} - 4T^{1-p^2/8},$$

where A' is the event that all cost vectors $\mathbf{x}_1, \dots, \mathbf{x}_T$ were drawn based on $\hat{\mathcal{D}}$ without modification. Since the error term is non-negative, this implies that

$$\begin{aligned} \mathbb{E}_{\mathcal{D}'} \left[\langle \bar{\mathbf{x}}_{\mathcal{D}'}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\mathcal{D}'}, \mathbf{w} \rangle \right] &\geq \Pr_{\mathcal{D}'}(A') \mathbb{E}_{\mathcal{D}'} \left[\langle \bar{\mathbf{x}}_{\mathcal{D}'}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\mathcal{D}'}, \mathbf{w} \rangle \middle| A' \right] \\ &\geq \Pr_{\mathcal{D}'}(A') \left(\frac{R}{p\sqrt{\log(T)}} - 4T^{1-p^2/8} \right). \end{aligned}$$

Finally, we have that $\Pr_{\mathcal{D}'}(A') = \Pr_{\hat{\mathcal{D}}}(A) \geq 1 - 2T^{1-p^2/8}$, so we get

$$\mathbb{E}_{\mathcal{D}'} \left[\langle \bar{\mathbf{x}}_{\mathcal{D}'}, \hat{\mathbf{w}} \rangle - \min_{\mathbf{w} \in \mathcal{W}} \langle \bar{\mathbf{x}}_{\mathcal{D}'}, \mathbf{w} \rangle \right] \geq \left(1 - 2T^{1-p^2/8}\right) \left(\frac{R}{p\sqrt{\log(T)}} - 4T^{1-p^2/8} \right).$$

Picking p sufficiently large, the result follows. ■