

Online Vector Balancing and Geometric Discrepancy

Nikhil Bansal

CWI Amsterdam & TU Eindhoven, Netherlands
N.Bansal@cwi.nl

Sahil Singla

Institute for Advanced Study & Princeton University, USA
singla@cs.princeton.edu

Haotian Jiang

University of Washington, USA
jhtdavid@cs.washington.edu

Makrand Sinha

CWI Amsterdam, Netherlands
makrand@cwi.nl

ABSTRACT

We consider an online vector balancing question where T vectors, chosen from an arbitrary distribution over $[-1, 1]^n$, arrive one-by-one and must be immediately given a \pm sign. The goal is to keep the discrepancy—the ℓ_∞ -norm of any signed prefix-sum—as small as possible. A concrete example of this question is the online interval discrepancy problem where T points are sampled one-by-one uniformly in the unit interval $[0, 1]$, and the goal is to immediately color them \pm such that every sub-interval remains always nearly balanced. As random coloring incurs $\Omega(T^{1/2})$ discrepancy, while the worst-case offline bounds are $\Theta(\sqrt{n \log(T/n)})$ for vector balancing and 1 for interval balancing, a natural question is whether one can (nearly) match the offline bounds in the online setting for these problems. One must utilize the stochasticity as in the worst-case scenario it is known that discrepancy is $\Omega(T^{1/2})$ for any online algorithm.

In a special case of online vector balancing, Bansal and Spencer [BS19] recently show an $O(\sqrt{n \log T})$ bound when each coordinate is *independently* chosen. When there are *dependencies* among the coordinates, as in the interval discrepancy problem, the problem becomes much more challenging, as evidenced by a recent work of Jiang, Kulkarni, and Singla [JKS19] that gives a non-trivial $O(T^{1/\log \log T})$ bound for online interval discrepancy. Although this beats random coloring, it is still far from the offline bound.

In this work, we introduce a new framework that allows us to handle online vector balancing even when the input distribution has dependencies across coordinates. In particular, this lets us obtain a $\text{poly}(n, \log T)$ bound for online vector balancing under arbitrary input distributions, and a $\text{polylog}(T)$ bound for online interval discrepancy. Our framework is powerful enough to capture other well-studied geometric discrepancy problems; e.g., we obtain a $\text{poly}(\log^d(T))$ bound for the online d -dimensional Tusnády’s problem. All our bounds are tight up to polynomial factors.

A key new technical ingredient in our work is an *anti-concentration* inequality for sums of pairwise uncorrelated random variables, which might also be of independent interest.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '20, June 22–26, 2020, Chicago, IL, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6979-4/20/06...\$15.00

<https://doi.org/10.1145/3357713.3384280>

CCS CONCEPTS

• **Theory of computation** → **Randomness, geometry and discrete structures; Online algorithms; Discrete optimization; Algorithmic game theory.**

KEYWORDS

online vector balancing; geometric discrepancy; anti-concentration; envy minimization

ACM Reference Format:

Nikhil Bansal, Haotian Jiang, Sahil Singla, and Makrand Sinha. 2020. Online Vector Balancing and Geometric Discrepancy. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC '20)*, June 22–26, 2020, Chicago, IL, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3357713.3384280>

1 INTRODUCTION

Consider the following online vector balancing question, originally proposed by Spencer [Spe77]: vectors $v_1, v_2, \dots, v_T \in [-1, 1]^n$ arrive online, and upon the arrival of v_t , a sign $\varepsilon_t \in \{\pm 1\}$ must be chosen irrevocably, so that the ℓ_∞ -norm of the signed sum $d_t = \varepsilon_1 v_1 + \dots + \varepsilon_t v_t$ remains as small as possible. That is, find the smallest B such that $\max_{t \in [T]} \|d_t\|_\infty \leq B$. As we shall see later, the problem arises naturally in various contexts where one wants to divide an incoming stream of objects, so that the split is as even as possible along each of the various dimensions that one might care about.

A naïve algorithm is to pick each sign ε_t randomly and independently, which by standard tail bounds gives $B = \Theta((T \log n)^{1/2})$ with high probability. In most of the interesting settings, $T \gg n$, and a natural question is whether the dependence on T can be improved from $T^{1/2}$ to say, $\log T$, or removed altogether (possibly with a worse dependence on n).

Offline setting. The offline version of the problem, where the vectors v_1, \dots, v_T are given in advance and the goal is to minimize $\max_{t \in [T]} \|d_t\|_\infty$, is known as the signed-series problem. It was first studied by Spencer [Spe77], who obtained a bound independent of T , but exponential in n . This was later improved by Bárány and Grinberg [BG81] to $B \leq 2n$. Chobanyan [Cho94] showed a beautiful connection between the signed-series problem and the classic Steinitz problem on the rearrangement of vector sequences—any upper bound on B also holds for the latter problem. Steinitz problem has a much longer history, originating from a question of Riemann and Lévy in the 19th century (c.f. the survey [Bár08] for some fascinating history). A long-standing conjecture for both the problems, still open, is that $B = O(n^{1/2})$. Another notable bound is

due to Banaszczyk [Ban12], who showed that $B = O((n \log T)^{1/2})$. While the original argument in [Ban12] was non-constructive, a polynomial time algorithm to find such a signing was recently given in [BG17].

In general, there has been extensive work on various offline discrepancy problems over last several decades, and several powerful techniques such as the partial coloring method [Spe85] and convex geometric methods [Gia97, Ban98, Ban12, MNT14] have been developed, which significantly improve upon the bounds given by random coloring. While these initial methods were mostly non-algorithmic, several new algorithmic techniques and insights have been developed in recent years [Ban10, LM15, Rot14, ES18, BDG16, LRR17, BDGL18, DNTT18].

Online setting. The online setting was first studied in the 70's and 80's, but it did not receive much interest later as it was realized that the best guarantees are already achieved by trivial algorithms. In particular, the $T^{1/2}$ dependence on T achieved by random coloring cannot be improved [Spe77]. See [Spe87, Bár79] for even more specific lower bounds. The difficulty is that the all-powerful adversary, upon seeing the signs chosen by the algorithm until time $t - 1$, can choose the next input vector v_t to be *orthogonal* to d_{t-1} . Now, irrespective of the choice of the sign ε_t , the resulting signed sum d_t satisfies

$$\begin{aligned} \|d_t\|_2^2 &= \|d_{t-1} + \varepsilon_t v_t\|_2^2 = \|d_{t-1}\|_2^2 + 2\varepsilon_t \langle d_{t-1}, v_t \rangle + \|v_t\|_2^2 \\ &= \|d_{t-1}\|_2^2 + \|v_t\|_2^2. \end{aligned} \quad (1)$$

For any d_{t-1} , one can always pick v_t with¹ $\|v_t\|_\infty \leq 1$ and $\|v_t\|_2^2 \geq n - 1$, resulting in $\|d_t\|_2^2 \geq (n - 1)t$, and hence $\|d_t\|_\infty = \Omega(t^{1/2})$ for all $t \in [T]$ (as long as $n > 1$).

It is therefore natural to ask if relaxing the power of the adversary, or making additional assumptions on the input sequence, can lead to interesting new ideas and to algorithms that perform much better, and in particular, give bounds that only mildly depend on T .

A natural assumption is that of *stochasticity*: if the arriving vectors are chosen in an i.i.d. manner from some distribution p , can we maintain that the ℓ_∞ norm of the current signed-sum d_t —henceforth, referred to as *discrepancy*—is poly(n) or poly($n, \log T$)?

Previous work and challenges. Recently, this stochastic setting was studied by Bansal and Spencer [BS19], where they considered the case where p is the uniform distribution on all $\{-1, 1\}^n$ vectors. They give an online algorithm achieving a bound of $O(\sqrt{n})$ on the *expected* discrepancy, matching the best possible offline bound, and an $O(\sqrt{n} \log T)$ discrepancy bound at all times $t \in [T]$, with high probability.

In general, the algorithmic discrepancy approaches developed in the last decade do not seem to help in the online setting. This is because in the offline setting, the algorithms can ensure that the discrepancy stays low by *simultaneously* updating the colors of various elements in a correlated way. In the online setting, however, the discrepancy must necessarily rise (in the ℓ_2 sense) whenever the incoming vector v_t is almost orthogonal to d_{t-1} , which can happen quite often. The only thing that the online algorithm can

do is to *actively* try to *cancel* this increase, whenever possible, by choosing the sign ε_t cleverly.

The algorithm of [BS19] crucially uses that if the coordinates of v_t are independently distributed and mean-zero², then for any d_{t-1} the incoming vector v_t will typically be far from being orthogonal to d_{t-1} . More quantitatively, the *anti-concentration* property for independent random variables gives that for any $d_{t-1} = (d_1, \dots, d_n)$, the random vector $v_t = (X_1, \dots, X_n)$ with X_1, \dots, X_n being independent and mean-zero satisfies

$$\mathbb{E}_v \left[|\langle d_{t-1}, v_t \rangle| \right] = \Omega \left(\left(\sum_{i=1}^n d_i^2 \cdot \mathbb{E}[X_i]^2 \right)^{1/2} \right).$$

Whenever $|\langle d_{t-1}, v_t \rangle|$ is large, the algorithm can choose ε_t appropriately to create a *negative drift* in (1), to offset the increase due to the $\|v_t\|_2^2$ term. We give a more detailed description below in §2.1.

In many interesting settings, however, the X_i 's can be *dependent*. For example, motivated by an envy minimization problem, Jiang, Kulkarni, and Singla [JKS19] considered the following natural online interval discrepancy problem: points x_1, \dots, x_T arrive uniformly in the interval $[0, 1]$, and the goal is to assign them signs online to minimize the discrepancy of every sub-interval of $[0, 1]$. (For adversarial arrivals, [JKS19] show poly(T) lower bounds.) Viewing the sub-intervals (after proper discretization) as coordinates, this becomes a stochastic online vector balancing problem, but where the random variables X_i corresponding to the various sub-intervals are dependent (details in §2.2). They give a non-trivial algorithm that achieves $T^{1/\log \log T}$ discrepancy, which is much better than the $T^{1/2}$ bound obtained by random coloring, but still substantially worse than polylog(T).

In general, the difficulty with dependent coordinates X_i is that even a small correlation can destroy anti-concentration, which makes it difficult to create a negative drift. For example, suppose the distribution p is mostly supported on vectors with an equal number of $+1$ and -1 coordinates. Now if d has the form $d = c(1, \dots, 1)$, then the incoming vector v_t is almost always orthogonal to it, and $\|d_T\|_2$ can potentially increase as fast as $\Omega(T^{1/2})$.

In this paper, we focus on the stochastic setting where the coordinates have dependencies, and give several results both for specific geometric problems and for general vector balancing under arbitrary distributions. In general, there are various other ways in which one can relax the power of the adversary, and in §8 we describe several interesting open questions and directions in this area.

1.1 Our Discrepancy Bounds

We first consider the following interval discrepancy problem. Let $x = x_1, \dots, x_T$ be a sequence of points drawn uniformly in $[0, 1]$ and let $\varepsilon_1, \dots, \varepsilon_T \in \{\pm 1\}$ be a signing. For an interval $I \subseteq [0, 1]$, let $\mathbf{1}_I$ denote the indicator function of the interval I . For any time $t \in [T]$, we define the discrepancy of interval I to be

$$\text{disc}_t(I) := \left| \varepsilon_1 \mathbf{1}_I(x_1) + \dots + \varepsilon_t \mathbf{1}_I(x_t) \right|.$$

We show the following bounds on discrepancy.

¹For any $d \in \mathbb{R}^n$, any basic feasible solution to $\langle d, x \rangle = 0$ with $x \in [-1, 1]^n$ has at least $n - 1$ coordinates ± 1 .

²Note that this holds in the case of uniform distribution over $\{-1, 1\}^n$.

THEOREM 1.1 (INTERVAL DISCREPANCY). *There is an online algorithm which selects signs $\varepsilon_t \in \{\pm 1\}$ such that, with high probability³, for every interval $I \subseteq [0, 1]$ we have $\max_{t \in [T]} \text{disc}_t(I) = O(\log^3 T)$. Moreover, with constant probability, for any online algorithm, $\max_{I \subseteq [0, 1]} \max_{t \in [T]} \text{disc}_t(I) = \Omega(\sqrt{\log T})$.*

This gives an exponential improvement over the $T^{1/\log \log T}$ bound of [JKS19], and is tight up to polynomial factors. The lower bound also improves a previous bound of $\Omega(\log^{1/4} T)$ of [JKS19].

There are two natural d -dimensional generalizations of the interval discrepancy problem, and our framework, which we will describe in §1.2, can handle both of them.

d -dimensional Online Interval Discrepancy: Consider a sequence of points x_1, \dots, x_T drawn uniformly from the unit cube $[0, 1]^d$. The goal is to simultaneously minimize the discrepancy of every interval for all the d -coordinates. In other words, to minimize the following for every interval I and every coordinate $i \in [d]$:

$$\text{disc}_t^i(I) := \left| \varepsilon_1 \mathbf{1}_I(x_1(i)) + \dots + \varepsilon_t \mathbf{1}_I(x_t(i)) \right|.$$

The offline version of this problem for $d \geq 2$ is equivalent to the classic d -permutations problem, where an upper bound of $O(\sqrt{d} \log T)$ [SST97] and a breakthrough lower bound of $\Omega(\log T)$ [NNN12, Fra18] for $d \geq 3$, and $\Omega(\sqrt{d})$ in general is known for the worst-case placement of points.

We show the following generalization of Theorem 1.1 that matches the best offline bounds, up to polynomial factors.

THEOREM 1.2 (d -DIMENSIONAL INTERVAL DISCREPANCY). *There is an online algorithm which selects signs $\varepsilon_t \in \{\pm 1\}$ such that, with high probability, for each $i \in [d]$ and $I \subseteq [0, 1]$, we have $\max_{t \in [T]} \text{disc}_t^i(I) = O(d \log^3 T)$. Moreover, with constant probability, for any online algorithm there exists an interval I and a coordinate $i \in [d]$, such that $\max_{t \in [T]} \text{disc}_t^i(I) = \Omega(\sqrt{d} \log(T/d))$.*

Previously, Jiang et al. [JKS19] could extend their analysis for online interval discrepancy to the $d = 2$ case and prove the same $T^{1/\log \log T}$ bound. However, their proof is rather ad-hoc and does not seem to generalize to higher d . In contrast, our bound holds for any d , and is tight up to polynomial factors.

The second natural generalization of interval discrepancy is to d -dimensional axis-parallel boxes, which gives the following online version of the extensively studied Tuszady’s Problem.

d -dimensional Online Tuszady’s Problem: Consider a sequence of points x_1, \dots, x_T drawn uniformly from the unit cube $[0, 1]^d$. The goal is to simultaneously minimize the discrepancy of all axis-parallel boxes. In other words, to minimize the following for every box B :

$$\text{disc}_t(B) := \left| \varepsilon_1 \mathbf{1}_B(x_1) + \dots + \varepsilon_t \mathbf{1}_B(x_t) \right|.$$

The (offline) Tuszady’s problem has a fascinating history (see [Mat09] and references there in), and after a long line of work, it is known that for the worst-case placement of points, the offline discrepancy is at most $O_d(\log^{d-\frac{1}{2}} T)$ [Nik17] and at least

$\Omega_d(\log^{d-1} T)$ [MN15]. We show the following result in the online setting, which is tight to within polynomial factors.

THEOREM 1.3 (TUSNADY’S PROBLEM). *There is an online algorithm which selects signs $\varepsilon_t \in \{\pm 1\}$ such that, with high probability, for every axis-parallel box B , we have $\max_{t \in [T]} \text{disc}_t(B) = O_d(\log^{2d+1} T)$. Moreover, for any online algorithm, with constant probability, there exists a box B such that $\max_{t \in [T]} \text{disc}_t(B) = \Omega_d(\log^{d/2} T)$.*

In contrast, the proof approach of [JKS19] completely breaks down for the Tuszady’s problem even in two dimensions and does not give any better lower bounds in terms of d . We recently learned that results similar to Theorems 1.1 and 1.3 were also obtained by Dwivedi et al. [DFGGR19], in the context of understanding the power of online thinning in reducing discrepancy.

Remark: Although all the problems above are stated for uniform distributions, one can use the probability integral transformation to reduce any product distribution to the uniform distribution without increasing the discrepancy, so our results in Theorems 1.2 and 1.3 also apply to any product distribution over $[0, 1]^d$.

Finally, note that Theorem 1.1 follows as a direct corollary of either of the above theorems.

General distributions. We now consider the setting of arbitrary distributions for the online vector balancing problem. Here we need to tackle the orthogonality issue which gave $\Omega(T^{1/2})$ lower bounds discussed in (1). As discussed earlier, for the uniform distribution over $\{-1, +1\}^n$, Bansal and Spencer [BS19] get around this issue since this does not happen for the uniform distribution reasonably often, and hence, $\mathbb{E}[\langle d_{t-1}, v_t \rangle]$ is large for any vector d_{t-1} . Using this, they obtain the bound $O(n^{1/2} \log T)$. Our next result shows that such a $\text{poly}(n, \log T)$ upper bound is possible even for arbitrary distributions.

THEOREM 1.4. (Vector balancing under dependencies) *For any sequence of vectors $v_1, \dots, v_T \in [-1, 1]^n$ sampled i.i.d. from some arbitrary distribution p , there is an online algorithm which selects signs $\varepsilon_t \in \{\pm 1\}$ such that, with high probability, we have*

$$\max_{t \in [T]} \left\| \varepsilon_1 v_1 + \dots + \varepsilon_t v_t \right\|_\infty = O(n^2(\log T + \log n)).$$

In §4.2 we show that the dependencies on n and $\log T$ in this theorem are tight up to polynomial factors as there is an $\Omega(n^{1/2} + (\log T / \log \log T)^{1/2})$ lower bound.

All of the above results follow from a general framework that we discuss next. In addition to the framework below, the key new technical ingredient is an anti-concentration inequality for dependent random variables, which we describe below in Theorem 1.5. This may be of independent interest.

1.2 Our Framework

To tackle the orthogonality issue, one of our key idea is to work with a *different basis* for the discrepancy vectors. More specifically, instead of maintaining bounds on the individual coordinate discrepancies $d_t(i)$, we maintain bounds on suitable linear combinations of them. This basis ensures that the (new) coordinates of the incoming vector are *uncorrelated*, i.e., $\mathbb{E}[X(i) \cdot X(j)] = \mathbb{E}[X(i)] \cdot \mathbb{E}[X(j)]$ for distinct coordinates i, j . Note that this condition is only on the

³Throughout the paper, “with high probability” means with $1 - 1/\text{poly}(n, T)$ probability where the exponent of the polynomial can be made as large as desired, depending on the constant in the discrepancy upper bound.

expected values, and is much weaker, e.g., even pairwise independence. Once one finds a suitable new basis, which turns out to be an eigenbasis of the covariance matrix, the anti-concentration bound for such random variables (proved below in Theorem 1.5), together with the standard exponential penalty based framework used in previous works [BS19, JKS19], gives Theorem 1.4.

For our results on geometric discrepancy problems, there is an additional challenge, we cannot afford to lose a $\text{poly}(n)$ factor, as in Theorem 1.4 above, since the dimension $n = \Theta(T)$. In this case, however, the update vectors are $(\log T)$ -sparse in the original basis (see §2) and one could hope to utilize this sparsity. Yet another challenge in this case is that bounding the discrepancy in a new basis preserves ℓ_2 -discrepancy in the original basis, but could lead to a \sqrt{n} loss in ℓ_∞ -discrepancy. To get $\text{polylog}(T)$ bounds, we use a natural basis from wavelet theory, called the *Haar system*, which simultaneously has sparsity, uncorrelation, and avoids the ℓ_2 to ℓ_∞ loss. This also easily extends to higher dimensions as these wavelets can be tensorized in a natural way to get a suitable basis for higher dimensional versions of the problems. A more detailed description of our framework is given in §2. Next we discuss our anti-concentration results.

1.3 Our Anti-Concentration Results for Non-Independent Random Variables

Suppose X_1, \dots, X_n are independent $\{-1, +1\}$ random variables with mean zero. Then, it is well-known that $|\sum_i X_i|$ has mean $\Theta(n^{1/2})$, and moreover, this value is at least $\Omega(n^{1/2})$ with constant probability.

Now, on the other hand, consider the following distribution. Let H_n be $n \times n$ Hadamard matrix and let $H_n(i)$ denote its i -th row for $i \in [n]$. Consider the random vector $X = (X_1, \dots, X_n)$, where $X = \xi \cdot H_n(i)$ for a Rademacher random variable $\xi \in \{-1, +1\}$ and a uniformly chosen $i \in [n]$. Then the X_i 's are still mean-zero and $\{-1, +1\}$, and in fact, pairwise independent. However, the magnitude of the sum $|\sum_i X_i|$ behaves very differently from the i.i.d. setting above. It takes value n with probability only $1/n$ (if $X = \xi \cdot H_n(1)$, the row of all 1's) and is 0 otherwise. In particular the mean is $\mathbb{E}[|\sum_i X_i|] = 1$ (instead of $n^{1/2}$ above), and moreover the entire contribution to the mean comes from an event with probability only $1/n$.

Nevertheless, we can say interesting things about the anti-concentration of sums of such random variables. In particular, we show the following results for uncorrelated or pairwise independent random variables.

THEOREM 1.5. (Uncorrelated anti-concentration) *For any vector $(a_1, \dots, a_n) \in \mathbb{R}^n$, let X_1, \dots, X_n be uncorrelated random variables that are bounded $|X_i| \leq c$, satisfy $\mathbb{E}[X_i X_j] = 0$ for all $i \neq j$, and have sparsity s (the number of non-zero X_i 's in any outcome). Then*

$$\mathbb{E}\left[\left|\sum_i a_i X_i\right|\right] \geq \mathbb{E}\left[\sum_i |a_i| X_i^2\right] \cdot \frac{1}{cs}. \quad (2)$$

Moreover, this bound is tight, even for pairwise independent random variables.

The tightness holds for the Hadamard example above, where $\mathbb{E}[\sum_i X_i] = 1$, $s = n$, $c = 1$, and $\mathbb{E}[\sum_i X_i^2] = n$.

THEOREM 1.6. (Pairwise independent anti-concentration) *For any $(a_1, \dots, a_n) \in \mathbb{R}^n$, let X_1, \dots, X_n be mean-zero pairwise independent random variables with sparsity $s \leq n$. Then*

$$\mathbb{E}\left[\left|\sum_i a_i X_i\right|\right] \geq \mathbb{E}\left[\sum_i |a_i X_i|\right] \cdot \frac{1}{s}. \quad (3)$$

Note that this bound is also tight for the Hadamard example. In general, the bound (3) is stronger than in (2); and a simple example in §3.2 shows that (3) cannot hold for uncorrelated random variables.

Although the anti-concentration properties and the small-ball probabilities for independent variables have been extensively studied (c.f. [NV13]), the uncorrelated and pairwise independent setting does not seem to have been studied before, and Theorems 1.5 and 1.6 do not seem to be known, to the best of our knowledge.

1.4 Applications to Envy Minimization

A classic measure of fairness in the field of fair division is envy [Fol67, TV85, LMMS04, Bud11]. A recent work of Benade et al. [BKPP18] introduced the *online envy minimization* problem where T items arrive one-by-one. In the two player setting, on arrival of item $t \in \{1, \dots, T\}$ we get to see the valuations $v_{it} \in [0, 1]$ for both the players $i \in \{1, 2\}$. The goal is to immediately and irrevocably allocate the item to one of the players while minimizing the maximum *envy*. There are two natural notions of envy: cardinal and ordinal (see §7 for definitions). Benade et al. [BKPP18] show an $\Omega(T^{1/2})$ lower bound for online envy minimization in the *adversarial* model—the reason is similar to Bárány's [Bár79] lower bound for online discrepancy. Can we obtain better bounds when the player valuations are drawn from a distribution?⁴

In the special case of product distributions (each player independently draws their value), Jiang et al. [JKS19] observed that the 2-dimensional interval discrepancy bounds also hold for online envy minimization. In particular, they obtained a $T^{1/\log \log T}$ bound on the ordinal envy. Our new interval discrepancy bound from Theorem 1.2 immediately improves this to an $O(\log^3 T)$ bound on ordinal envy. Moreover, we use our vector balancing result to obtain an $O(\log T)$ bound on the cardinal envy even for general distributions.

Corollary 1.7. *Suppose valuations of two players are drawn i.i.d. from some distribution p over $[0, 1] \times [0, 1]$. Then, for an arbitrary distribution p (i.e., player valuations for the same item could be correlated), the online cardinal envy is $O(\log T)$. Moreover, if p is a product distribution (i.e., player valuations for the same item are independent) then the online ordinal envy is also $O(\log^3 T)$.*

Paper Organization. The rest of the paper is organized as follows: in §2, we give an overview of previous challenges and our main ideas. In §3, we prove our key anti-concentration theorems that are necessary for our upper bounds on discrepancy. In §4, we give upper and lower bounds for online discrepancy under certain “uncorrelation” assumptions on the distribution. Then, we apply these bounds in §5 to obtain our vector balancing result (Theorem 1.4). In §6, we again apply these bounds to obtain our geometric discrepancy results

⁴If we make a simplifying assumption that the distribution does not depend on the time horizon T , better bounds are known [ZP19, DGK⁺14].

(Theorems 1.2 and 1.3). In §7, we show why our results immediately apply to online envy minimization. Finally, in §8 we end with some discussion of open problems and directions.

2 PROOF OVERVIEW

Let us start by reviewing the approach considered by Bansal and Spencer [BS19] in the case of independent coordinates. We also discuss the challenges involved in extending it to the setting of dependent coordinates.

2.1 Independent Coordinates: Bansal and Spencer

Consider the online vector balancing problem, when each arriving vector is uniformly chosen from $\{\pm 1\}^n$, so that all the coordinates are independent. To design an online algorithm, it is natural to keep a potential function that keeps track of the discrepancy and chooses a sign ε_t for the current vector v_t that minimizes the increase in the potential. Formally, let $d_t = \varepsilon_1 v_1 + \dots + \varepsilon_t v_t$ denote the discrepancy vector at time t . For a parameter $0 < \lambda < 1$, define the potential function

$$\Phi_t = \sum_{i \in [n]} \cosh(\lambda d_t(i)),$$

where $d_t(i)$ denotes the i th coordinate of d_t and $\cosh(x) = \frac{1}{2} \cdot (e^x + e^{-x})$ for all $x \in \mathbb{R}$. One should think of the above potential function as a proxy for the maximum discrepancy as Φ_t is dominated by the maximum discrepancy: $\Phi_t \approx e^{\lambda \|d_t\|_\infty}$.

On the arrival of vector v_t , the algorithm chooses a sign $\varepsilon_t \in \{\pm 1\}$, which updates the discrepancy vector to $d_t = d_{t-1} + \varepsilon_t v_t$ and changes the potential from Φ_{t-1} to Φ_t . If we can show that whenever $\Phi_t > 2n$, the drift $\Delta\Phi_t := \Phi_t - \Phi_{t-1}$ is negative in expectation for the sign ε_t chosen by the algorithm, then we can say that the potential after T arrivals, Φ_T , is bounded by $\text{poly}(nT)$ with high probability. This implies $\cosh(\lambda \|d_T\|_\infty)$ is bounded by $\text{poly}(nT)$, which means a bound of $O(\lambda^{-1} \log T)$ on the maximum discrepancy.

Let us try to compute the expected drift. Define $d = d_{t-1}$. By considering the Taylor expansion, we get $\cosh(x + \delta) \leq \cosh(x) + \sinh(x)\delta + \cosh(x)\delta^2$ where $\sinh(x) = \frac{1}{2} \cdot (e^x - e^{-x})$ for all $x \in \mathbb{R}$. So,

$$\begin{aligned} \Delta\Phi_t &\approx \sum_{i \in [n]} \left(\lambda \sinh(\lambda d(i)) \cdot (\varepsilon_t v_t(i)) + \lambda^2 \cosh(\lambda d(i)) \cdot (\varepsilon_t v_t(i))^2 \right) \\ &= \varepsilon_t \lambda L + \lambda^2 Q, \end{aligned}$$

where $L = \sum_{i \in [n]} \sinh(\lambda d(i)) \cdot v_t(i)$ is the *linear* term and $Q = \sum_{i \in [n]} \cosh(\lambda d(i))$ is the *quadratic* term from the Taylor expansion (note that $(\varepsilon_t v_t(i))^2 = 1$). Since the algorithm is free to choose the sign ε_t to minimize the drift, $\Delta\Phi_t \approx -\lambda |L| + \lambda^2 Q$. Now if one can show that $\mathbb{E}_{v_t} [|L|] \geq \frac{\mathbb{E}[Q]}{2\lambda}$, we would get that the expected drift $\mathbb{E}[\Delta\Phi_t] < 0$, and this would translate to a good discrepancy bound of $O(\lambda^{-1} \log T)$ if λ is large as described above.

Since $\cosh(x)$ and $|\sinh(x)|$ only differ by at most 1, we can make the approximation $Q \approx \sum_{i \in [n]} |\sinh(\lambda d(i))|$ up to some small error. So, denoting $\beta = 1/\lambda$ and $a_i = \sinh(\lambda d(i))$, our task reduces to proving the following anti-concentration statement:

Question. Let X_1, \dots, X_n be independent random variables with $|X_i| \leq 1$. What is the smallest β such that the following holds:

$$\mathbb{E} \left[\left| \sum_{i \in [n]} a_i X_i \right| \right] \geq \frac{1}{\beta} \cdot \mathbb{E} \left[\sum_{i \in [n]} |a_i| X_i^2 \right]. \quad (4)$$

In the case where the X_i 's are independent Rademacher (± 1) random variables, classical Khintchine's inequality and Cauchy-Schwarz tell us that

$$\begin{aligned} \mathbb{E} \left[\left| \sum_{i \in [n]} a_i X_i \right| \right] &\geq \frac{1}{\sqrt{2}} \cdot \left(\sum_{i \in [n]} a_i^2 \right)^{1/2} \\ &\geq \frac{1}{\sqrt{2n}} \left(\sum_{i \in [n]} |a_i| \right) = \frac{1}{\sqrt{2n}} \cdot \mathbb{E} \left[\sum_{i \in [n]} |a_i| X_i^2 \right], \end{aligned}$$

so $\beta = O(\sqrt{n})$, which suffices for the discrepancy application. In general, when X_i 's are not Rademacher but are still bounded ($|X_i| \leq 1$), mean-zero, and *independent*, then following [BS19] one can still show that $\beta = O(\sqrt{n})$.

The above gives a bound of $O(\sqrt{n} \log T)$ on the maximum discrepancy at every time $t \in [T]$. However, when the input distribution has dependencies across coordinates, i.e. the X_i 's are dependent, one can not take β to be small in general. For example, $\beta \rightarrow \infty$ when all a_i 's are one and a random set of coordinates $S \subset [n]$ of size $n/2$ (say n is even) take value $+1$ and the remaining coordinates in $[n] \setminus S$ take value -1 .

Next we discuss the simplest geometric discrepancy problem—the interval discrepancy problem in one dimension—where such a situation already arises if we use the same approach as above.

2.2 Interval Discrepancy: Previous Barriers

Recall, we have T points x_1, \dots, x_T chosen uniformly from $[0, 1]$ which need to be given ± 1 signs online. Consider the *dyadic* intervals $I_{j,k} := [k2^{-j}, (k+1)2^{-j}]$ where $0 \leq k < 2^j$ and $0 \leq j \leq \log T$. For intuition, imagine embedding the unit interval on a complete binary tree of height $\log T$; now sub-intervals corresponding to every node of the binary tree are dyadic intervals. Note that the smallest dyadic interval has size $2^{-\log T} = 1/T$. By a standard reduction, every sub-interval of $[0, 1]$ is contained in a union of some $O(\log T)$ dyadic intervals, so it suffices to track the discrepancy of these dyadic intervals.

Denoting by $\mathbf{1}_I$ the indicator function for an interval I , define

$$d_t(I) := \varepsilon_1 \mathbf{1}_I(x_1) + \dots + \varepsilon_t \mathbf{1}_I(x_t).$$

Note that $|d_t(I_{j,k})|$ is the discrepancy of the interval $I_{j,k}$ at time t . A natural choice of algorithm is to use the potential function

$$\Phi_t = \sum_{j,k} \cosh(\lambda d_t(I_{j,k})),$$

which is a proxy for the maximum discrepancy of any dyadic interval. Ideally, we want to set $0 < \lambda < 1$ as large as possible. Defining $d_{j,k} = d_{t-1}(I_{j,k})$, and doing a similar analysis as before, we derive

$$\Delta\Phi_t \approx \varepsilon_t \lambda L + \lambda^2 Q,$$

where $L = \sum_{j,k} \sinh(\lambda d_{j,k}) \cdot \mathbf{1}_{I_{j,k}}(x_t)$ and $Q = \sum_{j,k} \cosh(\lambda d_{j,k}) \cdot \mathbf{1}_{I_{j,k}}(x_t)^2$. The problem again reduces to showing an anti-concentration statement as in Eq. (4) with X_i 's being the indicators $\mathbf{1}_{I_{j,k}}$ for all j, k . It turns out that the smallest β one can hope for

this setting is exponential in the height of the tree (see the full version [BJS19] for an example), which for binary trees of height $\log T$ only yields a $\text{poly}(T)$ bound on the discrepancy.

One can still leverage something out of this approach—letting $B = T^{1/\log \log T}$, it was shown by Jiang, Kulkarni, and Singla [JKS19] that by embedding B -adic intervals on a B -ary tree of height $\log \log T$, the above approach gives a sub-polynomial $T^{1/\log \log T}$ bound for the interval discrepancy problem. However, this cannot be pushed to give a $\text{polylog}(T)$ bound because the above obstruction does not allow us to handle trees of height $\log T$.

2.3 Interval Discrepancy: A New Potential and the BDG Inequality

To get around the previous problem, we take a different approach and instead of directly using the discrepancies in the potential Φ_t , we work with linear combinations of discrepancies with the following desirable properties. First, if there is a bound on these linear combinations then it should imply a bound on the original discrepancies. Second, and more importantly, the term L in $\Delta\Phi_t$ can be viewed as a martingale, which leads to much better anti-concentration properties, i.e., smaller β in (4).

More specifically, consider the previous embedding of the dyadic intervals of length at least $1/T$ on the complete binary tree of depth $\log T$. For any interval $I_{j,k}$, let the left half interval be $I_{j,k}^l$ and the right half interval be $I_{j,k}^r$, and consider the difference of their discrepancies

$$d_t^-(I_{j,k}) := d_t(I_{j,k}^l) - d_t(I_{j,k}^r).$$

Note that if $|d_t(I_{j,k})| \leq \alpha$ and also $|d_t^-(I_{j,k})| \leq \alpha$, then both $|d_t(I_{j,k}^l)| \leq \alpha$ and $|d_t(I_{j,k}^r)| \leq \alpha$. A simple inductive argument now shows that if $|d_t([0, 1])| \leq \alpha$ and the differences of discrepancy for every dyadic interval $I_{j,k}$ satisfies $|d_t^-(I_{j,k})| \leq \alpha$, then every dyadic interval also has discrepancy at most α , thus satisfying the first property above. So let us consider a different potential function:

$$\Xi_t := \cosh(\lambda d_t(I_{0,0})) + \sum_{j,k} \cosh(\lambda d_t^-(I_{j,k}))$$

with j, k ranging over all the dyadic intervals (corresponding to internal nodes of the tree) and $0 < \lambda < 1$ is a parameter that we want to set as large as possible. Denoting $d_{j,k}^- = d_{t-1}^-(I_{j,k})$, as before, we can write $\Delta\Xi_t \approx \varepsilon_t \lambda L + \lambda^2 Q$, with

$$L = \sinh(\lambda d_t(I_{0,0})) + \sum_{j,k} \sinh(\lambda d_{j,k}^-) \cdot X_{j,k}(x_t) \quad \text{and}$$

$$Q = \cosh(\lambda d_t(I_{0,0})) + \sum_{j,k} \cosh(\lambda d_{j,k}^-) \cdot X_{j,k}(x_t)^2,$$

where $X_{j,k} = \mathbf{1}_{I_{j,k}^l} - \mathbf{1}_{I_{j,k}^r}$ for any interval $I_{j,k}$. Note that $X_{j,k}$ takes value 1 on the left half of $I_{j,k}$, and -1 on the right half of $I_{j,k}$, and is zero otherwise.

Anti-concentration via Martingale analysis. Now we show how the random variable L can be viewed as a $(\log T)$ -step martingale. Let us view a uniform point $x \in [0, 1]$ as being sampled one bit at a time, starting with the most significant bit. At any point where j bits of x have been revealed, the interval $I_{j,k}$ on the j^{th} level of the

dyadic tree is determined. Now, consider the process that starts with the value $Y_0 = \sinh(\lambda d_{0,0})$ at the root and at any time $0 \leq j \leq \log T$, the process is on some node of the j^{th} level. Conditioned on this node being $I_{j,k}$, the payoff $Y_j := a_j X_j$ where $a_j = \sinh(d_{j,k}^-)$ and X_j equals 1 if the process moves to the left child and equals -1 otherwise. Defining $L_j = Y_0 + Y_1 + \dots + Y_j$, it follows that the sequence $L_0, \dots, L_{\log T}$ is a martingale and $L = L_{\log T}$.

Moreover, by the approximation $\cosh(x) \approx |\sinh(x)|$, we get that $Q = |Y_0| + |Y_1| + \dots + |Y_{\log T}|$. Letting $a_0 = Y_0$ and $X_0 = 1$, the question then becomes—what is the smallest β such that the following holds:

$$\mathbb{E} \left[\sum_{i=0}^{\log T} a_i X_i \right] \geq \frac{1}{\beta} \cdot \mathbb{E} \left[\sum_{i=0}^{\log T} |a_i| X_i^2 \right] = \frac{1}{\beta} \cdot \mathbb{E} \left[\sum_{i=0}^{\log T} |a_i| \right].$$

For martingales, a statement similar to Khintchine's inequality is implied by the well-known Burkholder-Davis-Gundy (BDG) inequality:

$$\mathbb{E} \left[\max_{t \leq \log T} \left| \sum_{i=0}^t a_i X_i \right| \right] \geq c \cdot \mathbb{E} \left[\left(\sum_{i=0}^t a_i^2 \right)^{1/2} \right]$$

for a positive constant c . One can also prove (see the appendix in the full version) that

$$(1 + \log T) \cdot \mathbb{E} \left[\sum_{i=0}^{\log T} a_i X_i \right] \geq \mathbb{E} \left[\max_{t \leq \log T} \left| \sum_{i=0}^t a_i X_i \right| \right].$$

Then, similar to the analysis for independent Rademacher random variables, using Cauchy-Schwarz,

$$(1 + \log T) \cdot \mathbb{E} \left[\sum_{i=0}^{\log T} a_i X_i \right] \geq c \cdot \mathbb{E} \left[\left(\sum_{i=0}^{\log T} a_i^2 \right)^{1/2} \right] \\ \geq \frac{c}{\sqrt{\log T}} \cdot \mathbb{E} \left[\sum_{i=0}^{\log T} |a_i| \right].$$

So we can conclude that $\beta = \text{polylog}(T)$, which gives a $\text{polylog}(T)$ bound on interval discrepancy.

How to extend this analysis to d -dimensional Tusnády's problem? The martingale analysis above strongly relied on the interval structure of the problem, which is not clear even for the two-dimensional Tusnády's problem. To answer this question, we take a much more general view of our online discrepancy problem.⁵

2.4 A More General View of Changing Basis

One can also view the above analysis of the interval discrepancy problem as a more general underlying principle—that of working with a different *basis*. For example, let us take a linear algebraic approach to interval discrepancy and consider it as a vector balancing problem in $\mathbb{R}^{\mathcal{D}}$, where $\mathcal{D} = \{I_{j,k} \mid 0 \leq j \leq \log T, 0 \leq k < 2^j\}$ is the set of all dyadic intervals. When a new point $x \in [0, 1]$ arrives, the coordinate $I \in \mathcal{D}$ of the update vector v_t is given by

$$v_t(I) = \mathbf{1}_I(x).$$

⁵ The more general view in fact gives a (slightly) better bound for interval discrepancy than the martingale based argument above. However, we include this martingale argument here, as it is insightful and could be useful for other problems.

Note that the update v_t lives in a T -dimensional subspace \mathcal{V} of the $(2T - 1)$ -dimensional space \mathbb{R}^D since the T -intervals, $I_{\log T, k}$, at the bottom layer determine the rest of the coordinates.

The original potential function Φ from §2.1 corresponded to working with the original basis, but with the potential function Ξ from §2.3, our approach consisted of bounding the ℓ_∞ -discrepancy in a different basis of the subspace \mathcal{V} . In general, we may choose any basis and then define a potential function as the sum of hyperbolic-cosines of the coordinates. To choose the right basis, we need several properties from it, but most importantly we need uncorrelation.

Uncorrelation and anti-concentration via the Eigenbasis. Recall that we say random variables X, Y are *uncorrelated* if $\mathbb{E}[XY] = \mathbb{E}[X] \cdot \mathbb{E}[Y]$, which is a condition only on the expected values of the random variables. Using Theorem 1.5, to show anti-concentration it suffices that the coordinates in the *new basis* are mean-zero and uncorrelated, i.e., $\mathbb{E}_v[v(i)v(j)] = 0$ for distinct coordinates i, j .

For our vector balancing results under arbitrary distributions in Theorem 1.4, we work in an *eigenbasis* of the covariance matrix. As will be shown in the proof later, standard results from linear algebra imply that the coordinates are uncorrelated in any eigenbasis. Our next lemma uses this anti-concentration (along with the hyperbolic cosine potential) to bound discrepancy in the *new basis* in terms of *sparsity*—number of non-zero coordinates—of the incoming vectors.

Lemma 2.1. (Bounded discrepancy) *Let p be a distribution supported over s -sparse vectors in $[-1, 1]^n$ satisfying $\mathbb{E}_{v \sim p}[v(i)v(j)] = 0$ for all $i \neq j \in [n]$. Then for vectors v_1, \dots, v_T sampled i.i.d. from p , there is an online algorithm that maintains $O(s(\log n + \log T))$ discrepancy with high probability.*

Even though this lemma implies low discrepancy in the new basis, we need to be careful in bounding discrepancy in the original basis.

Sparsity and going back to the original basis. As discussed briefly in §1.2, although working in an eigenbasis allows us to obtain polynomial bounds for vector balancing, this is apriori not sufficient for our polylogarithmic geometric discrepancy bounds. There are two main challenges—firstly, working in a new basis might lose any sparsity that we might have in the original basis; e.g., in the one-dimensional interval discrepancy problem the arriving vectors are $(\log T)$ -sparse (dyadic intervals) in the original basis, but could be $\Omega(T)$ -sparse in the new basis; and secondly, even if one can find a new basis where the coordinates are uncorrelated and have low sparsity, Lemma 2.1 only implies low ℓ_∞ -discrepancy in the new basis. So going back to the original basis might lose us a factor \sqrt{n} more (we can only claim ℓ_2 -discrepancy is the same). Recall, when we view interval discrepancy as vector balancing, $n = \Theta(T)$, so we cannot afford losing \sqrt{n} . Fortunately, there is a special basis consisting of *Haar wavelets* that allows us to prove $\text{polylog}(T)$ geometric discrepancy bounds.

2.5 Haar Wavelets: Polylogarithmic Geometric Discrepancy

There is a natural orthogonal basis associated with the unit interval—the basis of Haar wavelet functions. These consist of functions $\Psi_{j,k}$ (formally defined in §6.1). Together these functions are known to

form an orthogonal basis for functions on the unit interval with bounded L_2 -norm.

Associated with the one-dimensional Haar wavelets is a natural martingale, which is the same martingale that our previous analysis in §2.3 relied on (e.g., $X_{j,k} = \Psi_{j+1,k}$ in the notation of §2.3.). It turns out that the Haar wavelets have nice orthogonality and sparsity properties that allow us to use Lemma 2.1—in particular, $\mathbb{E}_x[h(x)h'(x)] = 0$ for distinct Haar wavelet functions $h \neq h'$ and x sampled uniformly from $[0, 1]$. Moreover, moving from the basis of Haar wavelets to the original basis does not incur any additional loss in the discrepancy bound, since for any dyadic interval I , one can show that its discrepancy satisfies $|d_t(I)| \leq \alpha \|\widehat{1}_I\|_1$, where α is a bound on the discrepancy in the Haar basis and $\|\widehat{1}_I\|_1$ is the ℓ_1 -norm of the function 1_I in the Haar basis. We prove that this ℓ_1 -norm is one, so $|d_t(I)| \leq \alpha$. This gives a more direct proof of the $\text{polylog}(T)$ interval discrepancy bound and also extends easily to the d -dimensional interval discrepancy problem.

Tusánady’s problem. Given the above framework of working in the Haar basis, our extension to the d -dimensional Tusánady’s problem now naturally follows. For example, in two dimensions, we work with the basis of Haar wavelet functions which is formed by a taking tensor product $\Psi_{j,k} \times \Psi_{j',k'}$ of the one dimensional wavelets. These functions form an orthogonal basis for all bounded product functions over $[0, 1]^2$ and have nice sparsity properties. Moreover, we prove that for any axis-parallel box, the ℓ_1 -norm of the Haar basis coefficients is one, so we do not lose any additional factor in the discrepancy bound while moving from the Haar basis to the original basis. This gives a polylogarithmic bound for two-dimensional Tusánady’s problem, and also extends easily to higher dimensions.

Notation. All logarithms in this paper will be base two. For any integer k , throughout the paper $[k]$ will denote the set $\{1, \dots, k\}$. For a vector $u \in \mathbb{R}^d$, we use $u(i)$ to denote the i^{th} coordinate of u for $i \in [d]$. Given another vector $v \in \mathbb{R}^d$, the notation $u \leq v$ denotes that $u(i) \leq v(i)$ for each $i \in [d]$. The all ones vector is denoted by $\mathbf{1}$. Given a distribution p , we use the notation $x \sim p$ to denote an element x sampled from the distribution p . For a real function f , we will write $\mathbb{E}_{x \sim p}[f(x)]$ to denote the expected value of $f(x)$ under x sampled from p . If the distribution is clear from the context, then we will abbreviate the above as $\mathbb{E}_x[f(x)]$.

3 ANTI-CONCENTRATION ESTIMATES

In this section we prove the anti-concentration results: we first prove it for uncorrelated random variables, and then give an improved bound for pairwise independent random variables. Although in the rest of this paper we only use the weaker bound for uncorrelated random variables, we think the improved anti-concentration for pairwise independent random variables is of independent interest and will find applications in the future.

3.1 Pairwise Uncorrelated Random Variables

The following anti-concentration bound will be used in our discrepancy applications.

THEOREM 1.5. (Uncorrelated anti-concentration) *For any vector $(a_1, \dots, a_n) \in \mathbb{R}^n$, let X_1, \dots, X_n be uncorrelated random variables*

that are bounded $|X_i| \leq c$, satisfy $\mathbb{E}[X_i X_j] = 0$ for all $i \neq j$, and have sparsity s (the number of non-zero X_i 's in any outcome). Then

$$\mathbb{E} \left[\left| \sum_i a_i X_i \right| \right] \geq \mathbb{E} \left[\sum_i |a_i| X_i^2 \right] \cdot \frac{1}{cs}. \quad (2)$$

Moreover, this bound is tight, even for pairwise independent random variables.

Note that if we have pairwise uncorrelated mean-zero random variables X_1, \dots, X_n , then we get $\mathbb{E}[X_i X_j] = \mathbb{E}[X_i] \cdot \mathbb{E}[X_j] = 0$, so the above lemma implies anti-concentration in this case. The bound in the above lemma is tight because of the Hadamard example described previously in §1.3.

The following is the main claim in the proof of Theorem 1.5. Roughly it says that $\mathbb{E} \left[\left| \sum_i a_i X_i \right| \right] \geq \frac{1}{c} \cdot \max_{k \in [n]} \mathbb{E}[|a_k| X_k^2]$. Combined with the observation that

$$\max_{k \in [n]} \mathbb{E}[|a_k| X_k^2] \geq \frac{1}{n} \cdot \mathbb{E} \left[\sum_k |a_k| X_k^2 \right],$$

this implies Theorem 1.5 when sparsity $s = n$. However, to get inequality (2) in terms of sparsity s , the statement of the claim has to be more refined.

Claim 3.1. For any $(a_1, \dots, a_n) \in \mathbb{R}^n$ and random variables X_1, \dots, X_n satisfying $|X_i| \leq c$ and $\mathbb{E}[X_i X_j] = 0$ for distinct i, j , the following holds for any $k \in [n]$,

$$\mathbb{E} \left[\left| \sum_i a_i X_i \right| \cdot \mathbf{1}_{X_k \neq 0} \right] \geq \frac{1}{c} \cdot \mathbb{E}[|a_k| X_k^2].$$

PROOF. Using that $|X_k| \leq c$, we have

$$\begin{aligned} c \cdot \mathbb{E} \left[\left| \sum_i a_i X_i \right| \cdot \mathbf{1}_{X_k \neq 0} \right] &\geq \mathbb{E} \left[\left| \sum_i a_i X_i \right| \cdot |X_k| \right] \\ &\geq \mathbb{E} \left[\text{sign}(a_k) \left(a_k X_k^2 + \sum_{i \neq k} a_i X_i X_k \right) \right]. \end{aligned}$$

Since $\mathbb{E}[X_i X_k] = 0$ for $i \neq k$, it follows that

$$\begin{aligned} c \cdot \mathbb{E} \left[\left| \sum_i a_i X_i \right| \cdot \mathbf{1}_{X_k \neq 0} \right] &\geq \sum_i a_i \cdot \text{sign}(a_k) \cdot \mathbb{E}[X_i X_k] \\ &= \mathbb{E}[|a_k| X_k^2]. \quad \square \end{aligned}$$

When combined with the following easy claim, this will prove Theorem 1.5.

Claim 3.2. Let Y_1, \dots, Y_n be correlated random variables such that for any outcome at most s of them are non-zero. Moreover, suppose there is a random variable L which satisfies

$$\mathbb{E}[|L| \cdot \mathbf{1}_{Y_k \neq 0}] \geq \mathbb{E}[|Y_k|] \quad \text{for all } k \in [n].$$

Then, $\mathbb{E}[|L|] \geq \frac{1}{s} \sum_k \mathbb{E}[|Y_k|]$.

PROOF. Sum the given inequality for all $k \in [n]$ to get

$$\sum_k \mathbb{E}[|Y_k|] \leq \sum_k \mathbb{E}[|L| \cdot \mathbf{1}_{Y_k \neq 0}] = \mathbb{E}[|L| \cdot \sum_k \mathbf{1}_{Y_k \neq 0}] \leq \mathbb{E}[|L| \cdot s]. \quad \square$$

PROOF OF THEOREM 1.5. Applying Claim 3.1 and Claim 3.2 (with $L = \sum_i a_i X_i$ and $Y_i = \frac{1}{c} \cdot |a_i| X_i^2$), we get that

$$\mathbb{E} \left[\left| \sum_i a_i X_i \right| \right] \geq \mathbb{E} \left[\sum_k |a_k| X_k^2 \right] \cdot \frac{1}{cs}. \quad \square$$

3.2 Pairwise Independent Random Variables

In the special case of pairwise independent random variables, it is possible to obtain an improved inequality over Theorem 1.5.

THEOREM 1.6. (Pairwise independent anti-concentration) For any $(a_1, \dots, a_n) \in \mathbb{R}^n$, let X_1, \dots, X_n be mean-zero pairwise independent random variables with sparsity $s \leq n$. Then

$$\mathbb{E} \left[\left| \sum_i a_i X_i \right| \right] \geq \mathbb{E} \left[\sum_i |a_i| X_i^2 \right] \cdot \frac{1}{s}. \quad (3)$$

Notice, (3) immediately implies (2) for mean-zero pairwise independent random variables with $|X_i| \leq c$. One cannot hope to prove the stronger statement (3) for uncorrelated random variables. See the full version [BJS19] for an example and the proof of Theorem 1.5.

4 ONLINE DISCREPANCY UNDER UNCORRELATED ARRIVALS

In this section we consider the vector balancing problem in the special case when the input distribution has uncorrelated coordinates. All our upper and lower bounds will then follow from choosing a suitable basis to reduce the original problem to a basis with uncorrelated coordinates.

4.1 Upper Bounds

We say a vector in \mathbb{R}^d is s -sparse if it has at most s non-zero coordinates. The following lemma bounds the discrepancy for uncorrelated sparse distributions.

Lemma 2.1. (Bounded discrepancy) Let p be a distribution supported over s -sparse vectors in $[-1, 1]^n$ satisfying $\mathbb{E}_{v \sim p}[v(i)v(j)] = 0$ for all $i \neq j \in [n]$. Then for vectors v_1, \dots, v_T sampled i.i.d. from p , there is an online algorithm that maintains $O(s(\log n + \log T))$ discrepancy with high probability.

PROOF OF LEMMA 2.1. Our algorithm will use the same potential function approach described in §2, and uses our anti-concentration lemma from §3 to argue that the potential always remains polynomially bounded.

Algorithm. At any time step t , let $d_t = \varepsilon_1 v_1 + \dots + \varepsilon_t v_t$ denote the current discrepancy vector after the signs $\varepsilon_1, \dots, \varepsilon_t \in \{\pm 1\}$ have been chosen. Set $\lambda = \frac{1}{2s}$ and define the potential function

$$\Phi_t := \sum_{i \in [n]} \cosh(\lambda d_t(i)).$$

When the vector v_t arrives, the algorithm chooses the sign ε_t that minimizes the increase $\Phi_t - \Phi_{t-1}$.

Bounded Positive Drift. Let us fix a time t . To simplify the notation, let $\Delta\Phi = \Phi_t - \Phi_{t-1}$, let $d = d_{t-1}$, and let $v = v_t$.

After choosing the sign ε_t , the discrepancy vector $d_t = d + \varepsilon_t v$. To bound the change $\Delta\Phi$, since $\cosh'(x) = \sinh(x)$ and $\sinh'(x) = \cosh(x)$, using Taylor expansion

$$\Delta\Phi \leq \sum_i \left(\lambda \sinh(\lambda d(i)) \cdot (\varepsilon_t v(i)) + \lambda^2 \cosh(\lambda d(i)) \cdot (\varepsilon_t v(i))^2 \right),$$

where the last inequality follows since $|\sinh(x)| \leq \cosh(x)$ for all $x \in \mathbb{R}$, and since $|\varepsilon_t v(i)| \leq 1$ and $\lambda < 1$, the higher order terms in the Taylor expansion are dominated by the first and second order terms.

Set $L = \sum_i \sinh(\lambda d(i))v(i)$, and $Q^* = \sum_i \cosh(\lambda d(i))v(i)^2$, and $Q = \sum_i |\sinh(\lambda d(i))|v(i)^2$. Since $\cosh(x) \leq |\sinh(x)| + 1$ for $x \in \mathbb{R}$ and $|v(i)| \leq 1$, we have $Q^* \leq Q + n$. Therefore,

$$\Delta\Phi \leq \varepsilon_t \cdot \lambda \cdot L + \lambda^2 \cdot Q + \lambda^2 n.$$

Since, the algorithm chooses ε_t to minimize the increase in the potential:

$$\Delta\Phi \leq -\lambda \cdot |L| + \lambda^2 \cdot Q + \lambda^2 n.$$

Now, since $\mathbb{E}_v[v(i)v(j)] = 0$ for all $i, j \in [n]$, we can apply Theorem 1.5 with $X_i = v(i)$ and $a_i = \sinh(\lambda d(i))$ to get that $\mathbb{E}_v[|L|] \geq \frac{1}{s} \cdot \mathbb{E}[Q] = 2\lambda \cdot \mathbb{E}[Q]$, which yields that

$$\mathbb{E}_v[\Delta\Phi] \leq -\lambda \cdot \mathbb{E}_v[|L|] + \lambda^2 \cdot \mathbb{E}_v[Q] + \lambda^2 n \leq -\lambda^2 \cdot \mathbb{E}_v[Q] + \lambda^2 n \leq n.$$

Discrepancy Bound. The above implies that for any time $t \in [T]$, the expectation $\mathbb{E}[\Phi_t] \leq nT$. By Markov's inequality and a union bound over the T time steps, with probability at least $1 - T^{-2}$, the potential $\Phi_t \leq nT^4$ for every time $t \in [T]$. Since at any time t , we have $\cosh(\lambda \|d_t\|_\infty) \leq \Phi_t$, this implies that with probability at least $1 - T^{-2}$, the discrepancy at every time is

$$O\left(\frac{\log(nT^4)}{\lambda}\right) = O(s(\log n + \log T)),$$

which finishes the proof of Lemma 2.1. \square

4.2 Lower Bounds

We now show that the dependence on s and $\log T$ in Lemma 2.1, cannot be improved up to polynomial factors. In particular, a lower bound of $\Omega(s^{1/2})$, even when the time horizon is $T = n$, follows directly from the following more general statement for the vector balancing problem under distributions with uncorrelated coordinates. This general version will later also imply our lower bounds for geometric discrepancy.

Lemma 4.1. *Let p be a distribution supported over vectors in $[-1, 1]^n$ with ℓ_2 -norm k , such that for every $i \neq j \in [n]$ we have $\mathbb{E}_{v \sim p}[v(i)v(j)] = 0$. Then, for any online algorithm that receives as input vectors v_1, \dots, v_n sampled i.i.d. from p , with probability at least $3/4$, the discrepancy is $\Omega(k)$ at some time $t \in [n]$.*

We remark that the above lower bound may not hold if the algorithms are offline.

PROOF OF LEMMA 4.1. Since the distribution p over inputs is fixed, we may assume that the algorithm is deterministic. Let

$d_t = \varepsilon_1 v_1 + \dots + \varepsilon_t v_t$ denote the discrepancy vector at any time $t \in [n]$. Consider the quadratic potential function:

$$\Phi_t := \|d_t\|_2^2 = \sum_{i \in [n]} |d_t(i)|^2.$$

We will need the following claim that shows Φ_t increases in expectation for any online algorithm. Let us define $\Delta\Phi_t = \Phi_t - \Phi_{t-1}$.

Claim 4.2. *Conditioned on any v_1, \dots, v_{t-1} and signs $\varepsilon_1, \dots, \varepsilon_{t-1}$ such that $\|d_{t-1}\|_\infty \leq \frac{k}{4}$, we have*

$$\mathbb{E}_{v_t}[\Delta\Phi_t] \geq k^2/2 \quad (5)$$

where the expectation is taken only over the update $v_t \sim p$.

PROOF. Set $\Delta\Phi = \Delta\Phi_t$, vector $v = v_t$, and $d = d_{t-1}$. When the update v arrives, note that $d_t = d + \varepsilon_t v$. Therefore, the increase in the potential is given by

$$\Delta\Phi = \sum_{i=1}^n \left(2d(i) \cdot \varepsilon_t v(i) + (\varepsilon_t v(i))^2 \right) \quad (6)$$

$$= 2\varepsilon_t \left(\sum_{i=1}^n d(i)v(i) \right) + \|v\|_2^2 = 2L + k^2, \quad (7)$$

where $L = \varepsilon_t \left(\sum_{i=1}^n d(i)v(i) \right)$.

To bound the expected value of L , we use Jensen's inequality and $\mathbb{E}_v[v(i)v(j)] = 0$ for $i \neq j$ to get:

$$\begin{aligned} (\mathbb{E}_v[L])^2 &\leq \mathbb{E}_v[L^2] \\ &= \sum_{i=1}^n |d(i)|^2 \cdot \mathbb{E}_v[v(i)^2] + \sum_{i \neq j} d(i)d(j) \cdot \mathbb{E}_v[v(i)v(j)] \\ &= \sum_{i=1}^n |d(i)|^2 \cdot \mathbb{E}_v[v(i)^2] \leq \|d\|_\infty^2 \cdot \sum_{i=1}^n \mathbb{E}_v[v(i)^2] \\ &= \|d\|_\infty^2 k^2 \leq \frac{k^4}{16}. \end{aligned}$$

Therefore, plugging the above in (6), we get

$$\mathbb{E}_v[\Delta\Phi] \geq -2 \cdot |\mathbb{E}_v[L]| + k^2 \geq -2 \cdot \left(\frac{k^4}{16}\right)^{1/2} + k^2 \geq \frac{k^2}{2}. \quad \square$$

To prove Lemma 4.1 using the last claim, we define τ to be the first time that $\|d_\tau\|_\infty > k/4$ if such a τ exists, or $\tau = n$ otherwise. Let us define a new potential Φ_t^* which remains the same as Φ_t for $t \leq \tau$ and increases by $k^2/2$ deterministically for every $t > \tau$.

Note that for all possible random choices,

$$\Phi_n^* \leq \Phi_{\tau-1} + \frac{nk^2}{2} \leq \frac{nk^2}{16} + \frac{nk^2}{2},$$

where the second inequality holds since $\|d_{\tau-1}\|_\infty \leq k/4$ and therefore, $\Phi_{\tau-1} \leq \frac{1}{16} \cdot nk^2$.

Moreover, let \mathcal{E} be the event that $\|d_t\|_\infty \leq k/4$ for every $t \leq n$. Note that when \mathcal{E} occurs then the final potential $\Phi_n^* \leq \frac{1}{16} \cdot nk^2$. Defining $p = \mathbb{P}[\mathcal{E}]$, we have

$$\mathbb{E}[\Phi_n^*] \leq p \cdot \frac{nk^2}{16} + (1-p) \left(\frac{nk^2}{16} + \frac{nk^2}{2} \right) = \frac{nk^2}{16} + (1-p) \frac{nk^2}{2}. \quad (8)$$

Moreover, from Claim 4.2 and the definition of Φ_t^* , it follows that $\mathbb{E}[\Phi_n^*] \geq \frac{1}{2} \cdot nk^2$. Comparing this with (8) yields that $p \leq 1/8$. Hence, with probability at least $7/8$, the discrepancy must be $k/4$ at some point. \square

Dependence on T . We next show that the discrepancy must be $\Omega((\log T / \log \log T)^{1/2})$ with high probability even when $n = O(1)$ (we assume $n \geq 2$ throughout this discussion). We only sketch the proof here as the arguments are standard. The idea is that for large T , there is a high probability of getting a long enough run of consecutive vectors with each v_t almost orthogonal to d_{t-1} .

Let \mathfrak{p} be the uniform distribution⁶ over vectors on the unit sphere S^{n-1} . For any vector $u \in \mathbb{R}^n$, and v sampled from \mathfrak{p} , there is a universal constant c so that for all $\delta \leq 1$, we have $\mathbb{P}[|\langle u, v \rangle| \leq \delta \|u\|_2 / n^{1/2}] \geq c\delta$.

Let $\beta \geq 1$ be some parameter that we optimize later. Setting $\delta = 1/(4\beta)$ gives that whenever $\|d_{t-1}\|_2 \leq \beta n^{1/2}$, there is at least $c/(4\beta)$ probability that $|\langle d_{t-1}, v_t \rangle| \leq 1/4$, and hence irrespective of the sign ε_t ,

$$\|d_t\|_2^2 \geq \|d_{t-1}\|_2^2 - 2|\langle d_{t-1}, v_t \rangle| + \|v_t\|_2^2 \geq \|d_{t-1}\|_2^2 + 1/2.$$

So for any τ consecutive steps, with at least $(c/4\beta)^\tau$ probability, this happens at every step (or the ℓ_2 -discrepancy already exceeds $\beta n^{1/2}$ at some step), and hence the discrepancy has ℓ_2 -norm at least $\Omega(\tau^{1/2})$.

Partitioning the time horizon T into T/τ disjoint blocks, and setting $\beta = \log(T)$, and $\tau = \Omega(\log T / \log \log T)$, the probability such a run does not occur in any block is at most $(1 - (c/4\beta)^\tau)^{(T/\tau)} = T^{-\Omega(1)}$ by our choice of the parameters. This gives the claimed lower bound.

5 ONLINE VECTOR BALANCING: POLYNOMIAL BOUNDS

In this section, we prove our vector balancing result for arbitrary distributions.

THEOREM 1.4. (Vector balancing under dependencies) *For any sequence of vectors $v_1, \dots, v_T \in [-1, 1]^n$ sampled i.i.d. from some arbitrary distribution \mathfrak{p} , there is an online algorithm which selects signs $\varepsilon_t \in \{\pm 1\}$ such that, with high probability, we have*

$$\max_{t \in [T]} \left\| \varepsilon_1 v_1 + \dots + \varepsilon_t v_t \right\|_\infty = O(n^2(\log T + \log n)).$$

PROOF OF THEOREM 1.4. Without loss of generality, we may assume that the distribution \mathfrak{p} is *symmetric*, i.e. both v and $-v$ have the same probability density, since we can always multiply the incoming vector v with a Rademacher ± 1 random variable without changing the problem. Let $P \in \mathbb{R}^{d \times d}$ denote the covariance matrix of our input distribution, and since \mathfrak{p} is symmetric, we get $P = \mathbb{E}_{v \sim \mathfrak{p}}[v v^T]$. Let U denote the orthogonal matrix whose columns u_1, \dots, u_n form an eigenbasis for P . Note that in terms of its spectral decomposition, $P = \sum_{k=1}^n \lambda_k u_k u_k^T$ for $\lambda_k \in \mathbb{R}$.

To prove our discrepancy bound, instead of working in the original basis, we will view our problem as a vector balancing problem in the basis given by the columns of U . Now the update sequence is given by w_1, \dots, w_T where $w_t = \frac{1}{\sqrt{n}} \cdot U^T v$ is the normalized update vector in the basis U .

⁶Our argument works for a wide class of distributions \mathfrak{p} , as long as for any $d_{t-1} \in \mathbb{R}^n$, the random incoming vector v_t sampled from \mathfrak{p} has a non-trivial probability of having a small inner product with d_{t-1} . We only give the argument for the uniform distribution on the unit sphere for simplicity.

Since $\|v\|_2 \leq \sqrt{n}$ and orthogonal matrices preserve ℓ_2 -norm, we have $\|U^T v\|_2 = \|v\|_2 \leq \sqrt{n}$. It follows that for any t , we have $\|w_t\|_\infty \leq \|w_t\|_2 = \frac{1}{\sqrt{n}} \cdot \|U^T v\|_2 \leq 1$. Furthermore, any two coordinates of the update vectors w_t 's are uncorrelated, i.e., for any $i \neq j \in [n]$ we have

$$\mathbb{E}[w_t(i) \cdot w_t(j)] = \frac{1}{n} \mathbb{E}[\langle u_i, v \rangle \langle u_j, v \rangle] = \frac{1}{n} u_i^T P u_j = 0,$$

where the last equality holds since $P = \sum_{k=1}^n \lambda_k u_k u_k^T$.

Thus, we can use the online algorithm from Lemma 2.1 to select signs $\varepsilon_1, \dots, \varepsilon_T \in \{\pm 1\}$. Let $d_t = \varepsilon_1 v_1 + \dots + \varepsilon_t v_t$ denote the discrepancy in the original basis. Now using the trivial bound of $s \leq n$ on sparsity in Lemma 2.1, we get that with high probability,

$$\frac{1}{\sqrt{n}} \left\| U^T d_t \right\|_\infty = O(n(\log n + \log T)).$$

Again, using that orthogonal matrices preserve ℓ_2 -norm,

$$\begin{aligned} \|d_t\|_\infty &\leq \|d_t\|_2 = \|U^T d_t\|_2 \\ &\leq \sqrt{n} \cdot \left\| U^T d_t \right\|_\infty = O(n^2(\log n + \log T)). \quad \square \end{aligned}$$

6 ONLINE GEOMETRIC DISCREPANCY: POLYLOGARITHMIC BOUNDS

In this section, we will prove our results on geometric discrepancy problems. For this, we will need a special basis of orthogonal functions on the unit interval called the Haar system. We briefly review its properties.

6.1 Preliminaries: Haar System

Let $\Psi : \mathbb{R} \rightarrow \mathbb{R}$ denote the *mother wavelet* function

$$\Psi(x) = \begin{cases} 1 & \text{if } 0 \leq x < \frac{1}{2} \\ -1 & \text{if } \frac{1}{2} \leq x < 1 \\ 0 & \text{otherwise.} \end{cases}$$

The *unnormalized* Haar wavelet functions are defined as follows: let $\Psi_{0,0}(x) = 1$ for all $x \in \mathbb{R}$, and for any $j \in \mathbb{N}^*$ and $0 \leq k < 2^{j-1}$ define

$$\Psi_{j,k}(x) := \Psi(2^{j-1}x - k).$$

We call j as the *scale* and k as the *shift* of the wavelet.

The Haar wavelet functions have nice orthogonality properties. In particular, let x be drawn uniformly from the unit interval $[0, 1]$. Then, one can easily check that

$$\begin{aligned} \mathbb{E}_x[\Psi_{j,k}(x)^2] &= 2^{-(j-1)} && \text{for } j > 0, \\ \mathbb{E}_x[\Psi_{j,k}(x)] &= 0 && \text{for } j > 0, \\ \mathbb{E}_x[\Psi_{j,k}(x)\Psi_{j',k'}(x)] &= 0 && \text{unless } j = j' \text{ and } k = k'. \end{aligned} \quad (9)$$

The Haar wavelet functions are not just orthogonal, but they form an orthogonal basis (not orthonormal), called the *Haar system*, for the class of functions on the unit interval with bounded L_2 -norm. In particular, we have the following proposition where for $j \in \mathbb{Z}_{\geq 0}$ we denote $\mathcal{H}_j = \bigcup_{0 \leq k < 2^{j-1}} \{\Psi_{j,k}\}$ and let $\mathcal{H} = \bigcup_{j \geq 0} \mathcal{H}_j$.

Proposition 6.1 ([Wal04], Chapter 5). *For any $f : [0, 1] \rightarrow \mathbb{R}$ such that $\mathbb{E}_x[f(x)^2] < \infty$, we have*

$$f = \sum_{h \in \mathcal{H}} \widehat{f}(h) \cdot h(x)$$

where $\widehat{f}(h) = \frac{\mathbb{E}_x[f(x)h(x)]}{\mathbb{E}_x[h(x)^2]}$ is the corresponding coefficient in the Haar system basis for $h \in \mathcal{H}$.

Indeed, since the Haar system forms an orthogonal basis, we also have that

$$\mathbb{E}_x[f(x)^2] = \sum_{h \in \mathcal{H}} \widehat{f}(h)^2 \cdot \mathbb{E}_x[h(x)^2].$$

A simple corollary of Proposition 6.1 is that $\mathcal{H}^{\otimes d}$ is an orthogonal basis for the linear space spanned by all functions over the unit cube $[0, 1]^d$ that have a product structure and bounded L_2 -norm. In particular, let $\mathbf{h} = (h_1, \dots, h_d)$ be an element of $\mathcal{H}^{\otimes d}$ which we will view as a function from $[0, 1]^d \rightarrow \mathbb{R}$ by defining $\mathbf{h}(x) = \prod_{i=1}^d h_i(x(i))$ for $x \in [0, 1]^d$. Note that distinct \mathbf{h} and \mathbf{h}' are orthogonal since for x drawn uniformly from $[0, 1]^d$,

$$\mathbb{E}_x[\mathbf{h}(x)\mathbf{h}'(x)] = \prod_{i=1}^d \mathbb{E}_{x(i)}[h_i(x(i))h'_i(x(i))] = 0. \quad (10)$$

Moreover, any product function can be expressed by functions in $\mathcal{H}^{\otimes d}$ as given in the following proposition⁷.

Proposition 6.2. For any $f : [0, 1]^d \rightarrow \mathbb{R}$ such that $f(x) = \prod_{i=1}^d f_i(x(i))$ for some $f_i : [0, 1] \rightarrow \mathbb{R}$ satisfying $\mathbb{E}_{x(i)}[f_i(x(i))^2] < \infty$, we have that

$$f = \sum_{\mathbf{h} \in \mathcal{H}^{\otimes d}} \widehat{f}(\mathbf{h})\mathbf{h},$$

where $\widehat{f}(\mathbf{h}) = \frac{\mathbb{E}_x[f(x)\mathbf{h}(x)]}{\mathbb{E}_x[\mathbf{h}(x)^2]}$.

PROOF. Expressing each f_i in the Haar system basis using Proposition 6.1, we get the statement of the proposition by tensoring. \square

Let $\mathcal{H}_{\leq j} = \bigcup_{j' \leq j} \mathcal{H}_{j'}$, and define $\mathcal{H}_{< j}, \mathcal{H}_{> j}, \mathcal{H}_{\geq j}$ analogously. Then, we have the following lemma about the Haar system decomposition of indicator functions of dyadic intervals.

Proposition 6.3. Let $\mathbf{1}_{I_{\ell, m}}$ denote the indicator function for the interval $I_{\ell, m} = [m2^{-\ell}, (m+1)2^{-\ell})$. Then,

$$\begin{aligned} \sum_{\mathbf{h} \in \mathcal{H}_0} |\widehat{\mathbf{1}}_{I_{\ell, m}}(\mathbf{h})| &= 2^{-\ell}, \\ \sum_{\mathbf{h} \in \mathcal{H}_j} |\widehat{\mathbf{1}}_{I_{\ell, m}}(\mathbf{h})| &= 2^{-(\ell+1-j)} \text{ for any } 1 \leq j \leq \ell \text{ and} \\ \widehat{\mathbf{1}}_{I_{\ell, m}}(\mathbf{h}) &= 0 \text{ for any } \mathbf{h} \in \mathcal{H}_{> \ell}. \end{aligned}$$

In particular, we have $\sum_{\mathbf{h} \in \mathcal{H}} |\widehat{\mathbf{1}}_{I_{\ell, m}}(\mathbf{h})| = \sum_{\mathbf{h} \in \mathcal{H}_{\leq \ell}} |\widehat{\mathbf{1}}_{I_{\ell, m}}(\mathbf{h})| = 1$.

See the full version [BJSS19] for a proof of Proposition 6.3.

We also get a similar proposition about dyadic boxes. In particular, let $\boldsymbol{\ell} = (\ell_1, \dots, \ell_d)$ for non-negative integers ℓ_i 's and let $\mathbf{m} = (m_1, \dots, m_d)$ for integers $0 \leq m_i < 2^{\ell_i}$. Let $\mathcal{H}_{\leq \boldsymbol{\ell}}^{\otimes d} = \mathcal{H}_{\leq \ell_1} \times \dots \times \mathcal{H}_{\leq \ell_d}$. Then, for the dyadic box

$$I_{\boldsymbol{\ell}, \mathbf{m}} = I_{\ell_1, m_1} \times \dots \times I_{\ell_d, m_d},$$

we have the following proposition. Below we write $\min\{\mathbf{e}, \mathbf{f}\}$ to denote the vector whose i^{th} coordinate is $\min\{e(i), f(i)\}$ for $\mathbf{e}, \mathbf{f} \in \mathbb{R}^d$.

⁷More generally, Proposition 6.2 holds for any L_2 -integrable function $f \in L_2([0, 1]^d)$, as the linear span of product functions with domain $[0, 1]^d$ is dense in $L_2([0, 1]^d)$.

Proposition 6.4. Let $\mathbf{1}_{I_{\ell, m}}$ denote the indicator function for the dyadic box $I_{\ell, m}$. Then,

$$\begin{aligned} \sum_{\mathbf{h} \in \mathcal{H}_j^{\otimes d}} |\widehat{\mathbf{1}}_{I_{\ell, m}}(\mathbf{h})| &= 2^{-\|\min\{\boldsymbol{\ell}, \boldsymbol{\ell}+1-j\}\|_1} \text{ for any } j \leq \ell \text{ and} \\ \widehat{\mathbf{1}}_{I_{\ell, m}}(\mathbf{h}) &= 0 \text{ for any } \mathbf{h} \notin \mathcal{H}_{\leq \ell}. \end{aligned}$$

In particular, we have $\sum_{\mathbf{h} \in \mathcal{H}^{\otimes d}} |\widehat{\mathbf{1}}_{I_{\ell, m}}(\mathbf{h})| = \sum_{\mathbf{h} \in \mathcal{H}_{\leq \ell}^{\otimes d}} |\widehat{\mathbf{1}}_{I_{\ell, m}}(\mathbf{h})| = 1$.

The proof of the above proposition follows from Proposition 6.3 by tensoring.

6.2 Online Interval Discrepancy Problem

Now we prove Theorem 1.2 for the d -dimensional interval discrepancy problem. Let $\mathbf{x} = (x_1, \dots, x_T)$ be a sequence of points in $[0, 1]^d$ and let $\varepsilon \in \{\pm 1\}^T$ be a signing. For any interval $I \subseteq [0, 1]$ and time $t \in [T]$, recall that the discrepancy of interval I along coordinate direction i at time t is denoted

$$\text{disc}_i^t(I, \mathbf{x}, \varepsilon) := \left| \varepsilon_1 \mathbf{1}_I(x_1(i)) + \dots + \varepsilon_t \mathbf{1}_I(x_t(i)) \right|.$$

We will just write $\text{disc}_i^t(I)$ when the input sequence and signing is clear from the context.

6.2.1 Upper Bounds. To maintain the discrepancy of all intervals, it will suffice to bound the discrepancy of every dyadic interval $I_{j, k} = [k2^{-j}, (k+1)2^{-j})$ of length at least $1/T$ along every coordinate direction i . Let $\mathcal{D} = \{I_{j, k} \mid 0 \leq j \leq \log T, 0 \leq k < 2^j\}$. Then, we prove the following.

Lemma 6.5. Given any sequence x_1, \dots, x_T sampled independently and uniformly from $[0, 1]^d$, there is an online algorithm that chooses a signing such that w.h.p. for every time $t \in [T]$, we have

$$\max_{i \in [d]} \text{disc}_i^t(I) = O(d \log^2 T) \text{ for all } I \in \mathcal{D}.$$

Before proving Lemma 6.5, we first show why it implies the upper bound in Theorem 1.2.

PROOF OF THE UPPER BOUND IN THEOREM 1.2. Without loss of generality, it suffices to consider half-open intervals. Every half-open interval $I \subseteq [0, 1]$ can be decomposed as a union of at most $2 \log T$ disjoint dyadic intervals in \mathcal{D} and two intervals $I_1 \subseteq I_{\log T, k}$ and $I_2 \subseteq I_{\log T, k'}$ for some $0 \leq k, k' < T$. Note that the length of I_1 and I_2 is at most $2^{-\log T} = 1/T$. We can then write,

$$\text{disc}_i^t(I) \leq (2 \log T) \cdot \max_{I \in \mathcal{D}} \text{disc}_i^t(I) + \text{disc}_i^t(I_1) + \text{disc}_i^t(I_2).$$

Applying the algorithm from Lemma 6.5, the discrepancy of every dyadic interval can be bounded w.h.p. by $O(d \log^2 T)$. The last two terms can be bounded by N_1 and N_2 respectively where N_1 (resp. N_2) is the number of points whose projections on any of the i coordinates is in I_1 (resp. I_2).

The probability that a random point z drawn uniformly from $[0, 1]^d$ has some coordinate $z(i)$ for $i \in [d]$ in I_1 or I_2 is at most $2d/T$. It follows that $\mathbb{E}[N_1 + N_2] \leq 2d$, so by Chernoff bounds, with probability at least $1 - T^{-4}$, the number $N_1 + N_2 \leq 4d \log T$.

Overall, w.h.p. for any interval I , we have

$$\max_{i \in [d]} \text{disc}_i^t(I) \leq 2 \log T \cdot (d \log^2 T) + 4d \log T = O(d \log^3 T). \quad \square$$

Next, we prove the missing Lemma 6.5.

PROOF OF LEMMA 6.5. We will consider the d -dimensional interval discrepancy problem as a vector balancing problem in $[d] \times \mathcal{H}_{\leq \log T}$ dimensions, where $\mathcal{H}_{\leq \log T}$ are the Haar wavelet functions with scale parameter at most $\log T$. Note that $|\mathcal{H}_{\leq \log T}| = T$, so the update vector in the vector balancing version will be Td -dimensional. Let us abbreviate $\mathcal{H}' = \mathcal{H}_{\leq \log T}$.

At any time when the point $x_t \in [0, 1]$ arrives, then the (i, h) coordinate of the update vector $v_t \in [-1, 1]^{d \times \mathcal{H}'}$ is given by

$$v_t(i, h) = h(x_t(i)).$$

Note that all the coordinates $(i, \Psi_{0,0})$ for $i \in [d]$ will always have the same value where $\Psi_{0,0}$ is constant Haar wavelet. So, to apply the online algorithm given by Lemma 2.1 we will only consider the subspace spanned by the coordinates (i, h) where $i \in [d]$ and $h \neq \Psi_{0,0}$ and the extra coordinate $(1, \Psi_{0,0})$.

Let us check first that we satisfy the conditions Lemma 2.1. First, note that the $\|v_t\|_\infty \leq 1$ and the vector v_t has at most $d \log T + 1$ non-zero coordinates, since for any fixed scale $0 \leq j \leq \log T$ and any point $z \in [0, 1]$, all but one of the values $\{h(z)\}_{h \in \mathcal{H}_j}$ are zero. The last condition to check is that the coordinates of the vector v_t are uncorrelated. This is a consequence of (9), since whenever coordinates (i, h) and (i', h') satisfy $i \neq i'$ or $h \neq h'$, we have

$$\mathbb{E}_{v_t}[v_t(i, h) \cdot v_t(i', h')] = \mathbb{E}_{x_t}[h(x_t(i)) \cdot h'(x_t(i'))] = 0.$$

To elaborate more, first note that we cannot have $h = h' = \Psi_{0,0}$ since we are working in the aforementioned subspace. Now, if $i \neq i'$ then the coordinates $x_t(i)$ and $x_t(i')$ are sampled independently from $[0, 1]$, and $\mathbb{E}_z[h(z)] = 0$ for $h \neq \Psi_{0,0}$ when z is drawn uniformly from $[0, 1]$. Otherwise, for $i = i'$ but $h \neq h'$, it follows from the orthogonality of the Haar system that $\mathbb{E}_z[h(z)h'(z)] = 0$.

Next, applying the online algorithm from Lemma 2.1, we select signs $\varepsilon_1, \dots, \varepsilon_T$ such that we get an ℓ_∞ bound on the vector $d_t = \sum_{l \leq t} \varepsilon_l v_l$. In particular, with high probability we have

$$|d_t(i, h)| = \left| \sum_{l \leq t} \varepsilon_l h(x_l(i)) \right| = O(d \log^2 T) \text{ for any } i \in [d], h \in \mathcal{H}'.$$

Note that the bound on $|d_t(i, \Psi_{0,0})|$ for $i \neq 1$ then follows because $|d_t(i, \Psi_{0,0})| = |d_t(1, \Psi_{0,0})|$.

To finish the proof, we need to bound the discrepancy of every dyadic interval in terms of $\|d_t\|_\infty$. Note that for any dyadic interval $I \in \mathcal{D}$, its coefficients in the Haar system basis $\widehat{1}_I(h) = 0$ for $h \in \mathcal{H}_{> \log T}$ using Proposition 6.3. Now, for any $i \in [d]$ and dyadic interval $I \in \mathcal{D}$, we can write

$$\begin{aligned} \text{disc}_t^i(I) &= \left| \sum_{l \leq t} \varepsilon_l \widehat{1}_I(x_l(i)) \right| = \left| \sum_{l \leq t} \varepsilon_l \sum_{h \in \mathcal{H}'} \widehat{1}_I(h) h(x_l(i)) \right| \\ &= \left| \sum_{h \in \mathcal{H}'} \widehat{1}_I(h) \left(\sum_{l \leq t} \varepsilon_l h(x_l(i)) \right) \right| = \left| \sum_{h \in \mathcal{H}'} \widehat{1}_I(h) d_t(i, h) \right| \\ &\leq \|d_t\|_\infty \cdot \left(\sum_{h \in \mathcal{H}'} |\widehat{1}_I(h)| \right) \leq \|d_t\|_\infty = O(d \log^2 T), \end{aligned}$$

where the second last inequality follows again from Proposition 6.3. \square

6.2.2 Lower Bounds.

PROOF OF THE LOWER BOUND IN THEOREM 1.2. Set $A = T/d$. We will again consider the d -dimensional interval discrepancy problem as a vector balancing problem in $[d] \times \mathcal{H}_{\leq \log A}$ dimensions where $\mathcal{H}_{\leq \log A}$ are the Haar wavelet functions with scale parameter at most $\log A$. Note that $|\mathcal{H}_{\leq \log T}| = A$, so the update vector in the vector balancing version will be T -dimensional. Let us abbreviate $\mathcal{H}' = \mathcal{H}_{\leq \log A}$.

At any time when the point $x_t \in [0, 1]^d$ arrives, then the (i, h) coordinate of the update vector v_t is given by

$$v_t(i, h) = \begin{cases} 0 & \text{if } h = \Psi_{0,0} \\ h(x_t(i)) & \text{otherwise.} \end{cases}$$

Here we are essentially ignoring the coordinates (i, h) with $h = \Psi_{0,0}$. Since for any fixed scale $0 < j \leq \log A$ and any point $z \in [0, 1]$, all but one of the values $\{h(z)\}_{h \in \mathcal{H}_j}$ are zero, the vector v_t has $d \log A$ non-zero coordinates all of which take value ± 1 . It follows that the Euclidean norm of any update vector v_t is $\sqrt{d \log A}$.

Furthermore, from the orthogonality of the Haar system, it follows that the coordinates of the vector v_t are uncorrelated:

$$\mathbb{E}_{v_t}[v_t(i, h)v_t(i', h')] = \mathbb{E}_{x_t}[h(x_t(i))h'(x_t(i'))] = 0.$$

Then, applying Lemma 4.1, we get that with probability at least $3/4$, there is a $t \in [T]$ and a coordinate (i, h) with $h \neq \Psi_{0,0}$ such that $|d_t(i, h)| = \Omega(\sqrt{d \log A})$.

Let $h = \Psi_{j,k}$ for some j, k where $j > 0$ (recall that coordinates (i, h) where $h = \Psi_{0,0}$ are always 0). Then, by definition $h = \mathbf{1}_{I_1} - \mathbf{1}_{I_2}$ where I_1 and I_2 are the first and second halves of the interval $I_{j-1,k}$. In this case,

$$\begin{aligned} |d_t(i, h)| &= \left| \left(\sum_{s \leq t} \varepsilon_s \mathbf{1}_{I_1}(x_s) \right) - \left(\sum_{s \leq t} \varepsilon_s \mathbf{1}_{I_2}(x_s) \right) \right| \\ &\leq 2 \max \left\{ |\text{disc}_t(I_1)|, |\text{disc}_t(I_2)| \right\}. \end{aligned}$$

Therefore, substituting $A = T/d$, there exists an interval I such that $\text{disc}_t^i(I) = \Omega \left(\sqrt{d \log \left(\frac{T}{d} \right)} \right)$. \square

6.3 Online Tusnády's Problem

Given the above framework of working in the Haar basis, our extension to the d -dimensional Tusnády's problem naturally follows by taking tensor products of the one dimensional wavelets. See the full version [BJS19] for full proofs.

7 APPLICATIONS TO ONLINE ENVY MINIMIZATION

In this section we use our vector balancing and two-dimensional interval discrepancy results to bound online envy. Let us first give the formal definition of envy.

Recall that there are two players and T items where for item $t \in \{1, \dots, T\}$, the valuation of the player $i \in \{1, 2\}$ is $v_{it} \in [0, 1]$. The *cardinal envy* is the standard notion of envy studied in fair division, which is the max over every player the difference between the player's valuation for the other player's allocation and the player's valuation for their own allocation [LMMS04, Bud11]. Formally, if

Player i is allocated set S_i by an algorithm, the cardinal envy is defined as

$$\text{envy}_C(\mathbf{v}_1, \mathbf{v}_2, S_1, S_2) := \max \left\{ \sum_{t \in S_2} v_{1t} - \sum_{t \in S_1} v_{1t}, \sum_{t \in S_1} v_{2t} - \sum_{t \in S_2} v_{2t} \right\}.$$

The notion of ordinal envy is defined ignoring the precise item valuations, but only with respect to the relative ordering of the items. Roughly, it is the worst possible cardinal envy for $[0, 1]$ valuations consistent with any given relative ordering. Thus for valuations in $[0, 1]$ the ordinal envy is always at least the cardinal envy [JKS19]. For $i \in \{1, 2\}$, let π_i denote the decreasing order with respect to the valuations v_{it} . Denote π_i^t the first t items in the order π_i . If Player i is allocated set S_i , the ordinal envy is defined as

$$\text{envy}_O(\pi_1, \pi_2, S_1, S_2) := \max_{t \geq 0} \left\{ |S_2 \cap \pi_1^t| - |S_1 \cap \pi_1^t|, |S_1 \cap \pi_2^t| - |S_2 \cap \pi_2^t| \right\}.$$

Jiang et al. [JKS19] discuss three equivalent definitions of ordinal envy.

Next, we prove Corollary 1.7, which is restated below.

Corollary 1.7. *Suppose valuations of two players are drawn i.i.d. from some distribution p over $[0, 1] \times [0, 1]$. Then, for an arbitrary distribution p (i.e., player valuations for the same item could be correlated), the online cardinal envy is $O(\log T)$. Moreover, if p is a product distribution (i.e., player valuations for the same item are independent) then the online ordinal envy is also $O(\log^3 T)$.*

PROOF. When the player valuations are drawn independently in $[0, 1]$, the “moreover” part is immediate from the following lemma of [JKS19] along with our Theorem 1.2 for 2-dimensional interval discrepancy.

Lemma 7.1 (Lemma 26 in [JKS19]). *For two players with independent valuations, any upper bound for 2-dimensional interval discrepancy problem also holds for 2-player online ordinal envy minimization.*

Next, we bound online cardinal envy under arbitrary distributions. In the following lemma we reduce this problem to 2-dimensional vector balancing.

Lemma 7.2. *For two players taking values from an arbitrary distribution p over $[0, 1] \times [0, 1]$, any upper bound for 2-dimensional vector balancing problem also holds for 2-player online cardinal envy minimization.*

PROOF. For $i \in \{1, 2\}$, let u_{it} denote the valuation of Player i for t^{th} item. We define the corresponding vector $v_t = (u_{1t}, -u_{2t})$. If our online vector balancing algorithm assigns the next vector v_t a + sign, we give the item to Player 2, and otherwise we give it to Player 1. The crucial observation is that $d_t(1)$ and $d_t(2)$ capture precisely the cardinal envy of Players 1 and 2, respectively. Thus, any bound $\|d_t\|_\infty$ implies a bound on the maximum cardinal envy. \square

The last lemma when combined with Theorem 1.4 finishes the proof of Corollary 1.7. \square

8 OPEN PROBLEMS AND DIRECTIONS

We close this paper by mentioning some interesting open problems that seem to require fundamental new techniques, and new directions in online discrepancy that remain unexplored.

Improving the dependence on n for general distributions. Theorem 1.4 gives a bound of $O(n^2 \log T)$ for online vector balancing problem under inputs sampled from an arbitrary distribution. However, an optimal dependence of $O(n^{1/2})$ on n is achievable in the special case where the distribution has independent coordinates [BS19], and also in the offline setting with worst-case inputs [Ban12]. This motivates the following question.

Question 1. *Given an arbitrary distribution p supported over $[-1, 1]^n$, is there an online algorithm that maintains discrepancy $\sqrt{n} \cdot \text{polylog}(T)$ on a sequence of T inputs sampled i.i.d. from p ?*

As the anti-concentration bound in Theorem 1.5 for uncorrelated variables is a $n^{1/2}$ factor worse than that for independent random variables, even getting a dependence of $n \cdot \text{polylog}(T)$ is an interesting first step.

Bounds in terms of sparsity. Several natural problems such as the d -dimensional interval discrepancy and d -dimensional Tusnády’s problem are best viewed as vector balancing problems where the input vectors are sparse. This motivates the following online version of the Beck-Fiala problem, where the online sequence x_1, \dots, x_T is chosen independently from some distribution p supported over s -sparse n -dimensional vectors over $[-1, 1]^n$. In the offline setting with worst-case inputs (and where we care about the discrepancy of every prefix), the methods of Banaszczyk [Ban12] give a bound of $(s \log T)^{1/2}$.

Question 2. *Given an arbitrary distribution p supported over s -sparse vectors in $[-1, 1]^n$, is there an online algorithm that maintains discrepancy $\text{poly}(s, \log T, \log n)$ on a sequence of T inputs sampled i.i.d. from p ?*

Resolving the above question would imply polylogarithmic bounds for Tusnády’s problem in d -dimensions (similar to that in Theorem 1.3) in the much more general setting where the points x_T are sampled from an arbitrary distribution over points in $[0, 1]^d$. Currently, Theorems 1.2 and 1.3 only hold when the points x_t are sampled from a product distribution on $[0, 1]^d$.

Prophet model. The last decade has seen several online problems being studied in the prophet model where the online inputs are sampled independently from known *non-identical* distributions (see, e.g., [Luc17]). The model clearly generalizes the i.i.d. model and for point mass distributions it captures the offline problem. This model becomes useful for online problems where the adversarial arrival guarantees are weak, which raises the following question.

Question 3. *Given arbitrary distributions p_1, \dots, p_T supported over vectors in $[-1, 1]^n$, is there an online algorithm that maintains discrepancy $\text{poly}(n, \log T, \log n)$ on a sequence of T inputs where vector v_t is sampled independently from p_t ?*

The techniques in Theorem 1.4 do not work since the eigenbasis may change with each arrival. It will be also interesting to study this prophet model for distributions over s -sparse vectors.

Oblivious adversary model. A very interesting direction that is strictly harder than the above stochastic settings is to understand

online vector balancing when the adversary is oblivious or non-adaptive, i.e., the adversary chooses the entire input sequence (without any stochastic assumptions) beforehand and is *not allowed* to change the inputs later based on the execution of the algorithm.

Recall that if the adversary is fully adaptive, then one cannot hope to prove a bound better than $\Theta(T^{1/2})$, but this might be possible for oblivious adversaries.

Question 4. *Is there an online algorithm that maintains discrepancy $\text{poly}(n, \log T)$ on any sequence of T vectors in $[-1, 1]^n$ chosen by an oblivious adversary?*

One could also consider the same question in the Beck-Fiala setting, and ask if better bounds are possible when there is sparsity.

Question 5. *Is there an online algorithm that maintains discrepancy $\text{poly}(s, \log T, \log n)$ on any sequence of T vectors in $[-1, 1]^n$ that are s -sparse and chosen by an oblivious adversary?*

Resolving Questions 4 and 5 would also have implications for both online geometric discrepancy and online envy minimization problems in the oblivious adversary setting.

ACKNOWLEDGMENTS

The full version of this paper is accessible at [BJS19]. The authors are thankful to Janardhan Kulkarni for several discussions on this project. The authors would also like to thank the anonymous referees of STOC 2020 for their helpful comments and suggestions. N.B. was supported by the ERC Consolidator Grant 617951 and the NWO VICI grant 639.023.812. H.J. was supported in part by the National Science Foundation, Grant Number CCF-1749609, CCF-1740551, DMS-1839116. S.S. was supported in part by the Schmidt Foundation and part of the work done while visiting CWI Amsterdam. M.S. was supported by the Netherlands Organization for Scientific Research, Grant Number 617.001.351 and the NWO VICI grant 639.023.812.

REFERENCES

- [Ban98] Wojciech Banaszczyk. Balancing vectors and Gaussian measures of n -dimensional convex bodies. *Random Struct. Algorithms*, 12(4):351–360, 1998.
- [Ban10] Nikhil Bansal. Constructive Algorithms for Discrepancy Minimization. In *Proceedings of FOCS 2010*, pages 3–10, 2010.
- [Ban12] Wojciech Banaszczyk. On series of signed vectors and their rearrangements. *Random Struct. Algorithms*, 40(3):301–316, 2012.
- [Bár79] Imre Bárány. On a Class of Balancing Games. *J. Comb. Theory, Ser. A*, 26(2):115–126, 1979.
- [Bár08] Imre Bárány. On the Power of Linear Dependencies. In *Building Bridges*, pages 31–45. Springer, 2008.
- [BDG16] Nikhil Bansal, Daniel Dadush, and Shashwat Garg. An Algorithm for Komlós Conjecture Matching Banaszczyk’s Bound. In *Proceedings of FOCS 2016*, pages 788–799, 2016.
- [BDGL18] Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. The Gram-Schmidt Walk: A Cure for the Banaszczyk Blues. In *Proceedings of STOC 2018*, pages 587–597, 2018.
- [BG81] Imre Bárány and Victor S Grinberg. On some combinatorial questions in finite-dimensional spaces. *Linear Algebra and its Applications*, 41:1–9, 1981.
- [BG17] Nikhil Bansal and Shashwat Garg. Algorithmic discrepancy beyond partial coloring. In *Proceedings of STOC 2017*, pages 914–926, 2017.
- [BJS19] Nikhil Bansal, Haotian Jiang, Sahil Singla, and Makrand Sinha. Online vector balancing and geometric discrepancy. *CoRR*, abs/1912.03350, <http://arxiv.org/abs/1912.03350>, 2019.
- [BKPP18] Gerdus Benade, Aleksandr M. Kazachkov, Ariel D. Procaccia, and Christos-Alexandros Psomas. How to Make Envy Vanish Over Time. In *Proceedings of EC 2018*, pages 593–610, 2018.
- [BS19] Nikhil Bansal and Joel H. Spencer. On-Line Balancing of Random Inputs. *CoRR*, abs/1903.06898, 2019.
- [Bud11] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Polit. Econ.*, 119(6):1061–1103, 2011.
- [Cho94] Sergej Chobanyan. Convergence as of rearranged random series in Banach space and associated inequalities. In *Probability in Banach Spaces*, 9, pages 3–29. Springer, 1994.
- [DFGGR19] Raaz Divedi, Ohad N. Feldheim, Ori Gurel-Gurevich, and Aaditya Ramdas. The power of online thinning in reducing discrepancy. *Probability Theory and Related Fields*, 174:103–131, 2019.
- [DGK⁺14] John P. Dickerson, Jonathan R. Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. The computational rise and fall of fairness. In *Proceedings of AAAI*, pages 1405–1411, 2014.
- [DNTT18] Daniel Dadush, Aleksandar Nikolov, Kunal Talwar, and Nicole Tomczak-Jaegermann. Balancing Vectors in Any Norm. In *Proceedings of FOCS 2018*, pages 1–10, 2018.
- [ES18] Ronen Eldan and Mohit Singh. Efficient algorithms for discrepancy minimization in convex sets. *Random Struct. Algorithms*, 53(2):289–307, 2018.
- [Fol67] Duncan K Foley. Resource allocation and the public sector. *Yale Econ Essays*, 7:45–98, 1967.
- [Fra18] Cole Franks. A simplified disproof of Beck’s three permutations conjecture and an application to root-mean-squared discrepancy. *arXiv:1811.01102*, 2018.
- [Gia97] Apostolos A Giannopoulos. On some vector balancing problems. *Studia Mathematica*, 122(3):225–234, 1997.
- [JKS19] Haotian Jiang, Janardhan Kulkarni, and Sahil Singla. Online Geometric Discrepancy for Stochastic Arrivals with Applications to Envy Minimization. *arXiv:1910.01073*, 2019.
- [LM15] Shachar Lovett and Raghu Meka. Constructive Discrepancy Minimization by Walking on the Edges. *SIAM J. Comput.*, 44(5):1573–1582, 2015.
- [LMMS04] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of EC 2004*, pages 125–131, 2004.
- [LRR17] Avi Levy, Harishchandra Ramadas, and Thomas Rothvoss. Deterministic Discrepancy Minimization via the Multiplicative Weight Update Method. In *Proceedings of IPCO 2017*, pages 380–391, 2017.
- [Luc17] Brendan Lucier. An economic view of prophet inequalities. *SIGecom Exchanges*, 16(1):24–47, 2017.
- [Mat09] Jiri Matousek. *Geometric discrepancy: An illustrated guide*, volume 18. Springer Science & Business Media, 2009.
- [MN15] Jiří Matoušek and Aleksandar Nikolov. Combinatorial discrepancy for boxes via the γ_2 norm. In *Proceedings of SoCG 2015*, pages 1–15, 2015.
- [MNT14] Jiří Matoušek, Aleksandar Nikolov, and Kunal Talwar. Factorization norms and hereditary discrepancy. *CoRR*, abs/1408.1376, 2014.
- [Nik17] Aleksandar Nikolov. Tighter bounds for the discrepancy of boxes and polytopes. *CoRR*, abs/1701.05532, 2017.
- [NNN12] Alantha Newman, Ofer Neiman, and Aleksandar Nikolov. Beck’s three permutations conjecture: A counterexample and some consequences. In *Proceedings of FOCS 2012*, pages 253–262, 2012.
- [NV13] Hoi H. Nguyen and Van H. Vu. *Small Ball Probability, Inverse Theorems, and Applications*, pages 409–463. Springer Berlin Heidelberg, 2013.
- [Rot14] Thomas Rothvoß. Constructive Discrepancy Minimization for Convex Sets. In *Proceedings of FOCS 2014*, pages 140–145, 2014.
- [Spe77] Joel Spencer. Balancing games. *J. Comb. Theory, Ser. B*, 23(1):68–74, 1977.
- [Spe85] Joel Spencer. Six standard deviations suffice. *Trans. Am. Math. Soc.*, 289(2):679–706, 1985.
- [Spe87] Joel H. Spencer. *Ten lectures on the probabilistic method*, volume 52. Society for Industrial and Applied Mathematics Philadelphia, 1987.
- [SST97] Joel H. Spencer, Aravind Srinivasan, and Prasad Tetali. The discrepancy of permutation families. In *Proceedings of SODA*, 1997.
- [TV85] William Thomson and Hal Varian. Theories of justice based on symmetry. *Social goals and social organizations: essays in memory of Elisha Pazner*, 126, 1985.
- [Wal04] David Walnut. *An Introduction to Wavelet Analysis*. Applied and Numerical Harmonic Analysis. Birkhäuser Basel, 1 edition, 1 2004.
- [ZP19] David Zeng and Alexandros Psomas. Fairness-efficiency tradeoffs in dynamic fair division. *CoRR*, abs/1907.11672, 2019.