Optimal Unateness Testers for Real-Valued Functions: Adaptivity Helps*

Roksana Baleshzar [†]   Deeparnab Chakrabarty [‡]   Ramesh Krishnan S. Pallavoor [†]

Sofya Raskhodnikova [†]   C. Seshadhri [§]

Abstract

We study the problem of testing unateness of functions $f : \{0,1\}^d \to \mathbb{R}$. We give a $O(\frac{d}{\varepsilon} \cdot \log \frac{d}{\varepsilon})$-query nonadaptive tester and a $O(\frac{d}{\varepsilon})$-query adaptive tester and show that both testers are optimal for a fixed distance parameter ε. Previously known unateness testers worked only for Boolean functions, and their query complexity had worse dependence on the dimension both for the adaptive and the nonadaptive case. Moreover, no lower bounds for testing unateness were known[1]. We also generalize our results to obtain optimal unateness testers for functions $f : [n]^d \to \mathbb{R}$.

Our results establish that adaptivity helps with testing unateness of real-valued functions on domains of the form $\{0,1\}^d$ and, more generally, $[n]^d$. This stands in contrast to the situation for monotonicity testing where there is no adaptivity gap for functions $f : [n]^d \to \mathbb{R}$.

1 Introduction

We study the problem of testing whether a given real-valued function f on domain $[n]^d$, where $n, d \in \mathbb{N}$, is unate. A function $f : [n]^d \to \mathbb{R}$ is *unate* if for every coordinate $i \in [d]$, the function is either nonincreasing in the i^{th} coordinate or nondecreasing in the i^{th} coordinate. Unate functions naturally generalize monotone functions, which are nondecreasing in all coordinates, and **b**-monotone functions, which have a particular direction in each coordinate (either nondecreasing or nondecreasing), specified by a bit-vector $\mathbf{b} \in \{0,1\}^d$. More precisely, a function is **b**-monotone if it is nondecreasing in coordinates i with $\mathbf{b}_i = 0$ and nonincreasing in the other coordinates. Observe that a function f is unate iff there exists some $\mathbf{b} \in \{0,1\}^d$ for which f is **b**-monotone.

A *tester* [36, 27] for a property $\mathcal{P}$ of a function f is an algorithm that gets a distance parameter $\varepsilon \in (0,1)$ and query access to f. It has to accept with probability at least $2/3$ if f has property $\mathcal{P}$ and reject with probability at least $2/3$ if f is ε-far (in Hamming distance) from $\mathcal{P}$. We say that f is ε-far from $\mathcal{P}$ if at least an ε fraction of values of f must be modified to make f satisfy $\mathcal{P}$. A tester has *one-sided error* if it always accepts a function satisfying $\mathcal{P}$, and has *two-sided error* otherwise.

---

[1]Concurrent work by Chen et al. [19] proves an $\Omega(d/\log^2 d)$ lower bound on the nonadaptive query complexity of testing unateness of Boolean functions. Our stronger lower bounds are for real valued functions.

A *nonadaptive* tester makes all its queries at once, while an *adaptive* tester can make queries after seeing answers to the previous ones.

Testing of various properties of functions, including monotonicity (see, e.g., [26, 21, 22, 33, 24, 23, 28, 1, 29, 3, 8, 7, 10, 13, 9, 6, 14, 15, 12, 18, 17, 31, 4, 5, 20, 34] and recent surveys [35, 11]), the Lipschitz property [30, 13, 9], bounded-derivative properties [12], and unateness [26, 32], has been studied extensively over the past two decades. Even though unateness testing was initially discussed in the seminal paper by Goldreich et al. [26] that gave first testers for properties of functions, relatively little is known about testing this property. All previous work on unateness testing focused on the special case of Boolean functions on domain $\{0,1\}^d$. The domain $\{0,1\}^d$ is called the *hypercube* and the more general domain $[n]^d$ is called the *hypergrid*. Goldreich et al. [26] provided a $O(\frac{d^{3/2}}{\varepsilon})$-query nonadaptive tester for unateness of Boolean functions on the hypercube. Recently, Khot and Shinkar [32] improved the query complexity to $O(\frac{d \log d}{\varepsilon})$, albeit with an adaptive tester.

In this paper, we improve upon both these works, and our results hold for a more general class of functions. Specifically, we show that unateness of real-valued functions on hypercubes can be tested nonadaptively with $O(\frac{d}{\varepsilon} \log \frac{d}{\varepsilon})$ queries and adaptively with $O(\frac{d}{\varepsilon})$ queries. More generally, we describe a $O(\frac{d}{\varepsilon} \cdot (\log \frac{d}{\varepsilon} + \log n))$-query nonadaptive tester and a $O(\frac{d \log n}{\varepsilon})$-query adaptive tester of unateness of real-valued functions over hypergrids.

In contrast to the state of knowledge for unateness testing, the complexity of testing monotonicity of real-valued functions over the hypercube and the hypergrid has been resolved. For constant distance parameter ε, it is known to be $\Theta(d \log n)$. Moreover, this bound holds for all *bounded-derivative* properties [12], a large class that includes **b**-monotonicity and some properties quite different from monotonicity, such as the Lipschitz property. Amazingly, the upper bound for all these properties is achieved by the same simple and, in particular, nonadaptive, tester. Even though proving lower bounds for adaptive testers has been challenging in general, a line of work, starting from Fischer [23] and including [8, 14, 12], has established that adaptivity does not help for this large class of properties. Since unateness is so closely related, it is natural to ask whether the same is true for testing unateness.

We answer this in the negative: we prove that any nonadaptive tester of real valued functions over the hypercube (for some constant distance parameter) must make $\Omega(d \log d)$ queries. More generally, it needs $\Omega(d(\log d + \log n))$ queries for the hypergrid domain. These lower bounds complement our algorithms, completing the picture for unateness testing of real-valued functions. From a property testing standpoint, our results establish that unateness is different from monotonicity and, more generally, any derivative-bounded property.

1.1 Formal Statements and Technical Overview

Our testers are summarized in the following theorem, stated for functions over the hypergrid domains. (Recall that the hypercube is a special case of the hypergrid with $n = 2$.)

Theorem 1.1. *Consider functions $f : [n]^d \to \mathbb{R}$ and a distance parameter $\varepsilon \in (0, 1/2)$.*
1. *There is a nonadaptive unateness tester that makes $O(\frac{d}{\varepsilon}(\log \frac{d}{\varepsilon} + \log n))$ queries*[2].
2. *There is an adaptive unateness tester that makes $O(\frac{d \log n}{\varepsilon})$ queries.*

---

[2]For many properties, when the domain is extended from the hypercube to the hypergrid, testers incur an extra multiplicative factor of $\log n$ in the query complexity. This is the case for our adaptive tester. However, note that the complexity of nonadaptive unateness testing (for constant ε) is $\Theta(d(\log d + \log n))$ rather than $\Theta(d \log d \log n)$.

Both testers have one-sided error.

Our main technical contribution is the proof that the extra $\Omega(\log d)$ is needed for nonadaptive testers. This result demonstrates a gap between adaptive and nonadaptive unateness testing.

Theorem 1.2. *Any nonadaptive unateness tester (even with two-sided error) for real-valued functions $f : \{0,1\}^d \to \mathbb{R}$ with distance parameter $\varepsilon = 1/8$ must make $\Omega(d \log d)$ queries.*

The lower bound for adaptive testers is an easy adaptation of the monotonicity lower bound in [14]. We state this theorem for completeness and prove it in Appendix A.1.

Theorem 1.3. *Any unateness tester for functions $f : [n]^d \to \mathbb{R}$ with distance parameter $\varepsilon \in (0, 1/4)$ must make $\Omega\left(\frac{d \log n}{\varepsilon} - \frac{\log 1/\varepsilon}{\varepsilon}\right)$ queries.*

Theorems 1.2 and 1.3 directly imply that our nonadaptive tester is optimal for constant ε, even for the hypergrid domain. The details appear in Appendix A.2.

1.1.1 Overview of Techniques

We first consider the hypercube domain. For each $i \in [d]$, an *i-edge* of the hypercube is a pair (x, y) of points in $\{0,1\}^d$, where $x_i = 0, y_i = 1$, and $x_j = y_j$ for all $j \in ([d] \setminus \{i\})$. Given an input function $f : \{0,1\}^d \to \mathbb{R}$, we say an *i*-edge (x, y) is *increasing* if $f(x) < f(y)$, *decreasing* if $f(x) > f(y)$, and *constant* if $f(x) = f(y)$.

Our nonadaptive unateness tester on the hypercube uses the work investment strategy from [6] (also refer to Section 8.2.4 of Goldreich's book [25]) to "guess" a good dimension where to look for violations of unateness (specifically, both increasing and decreasing edges). For all $i \in [d]$, let α_i be the fraction of the *i*-edges that are decreasing, β_i be the fraction of the *i*-edges that are increasing, and $\mu_i = \min(\alpha_i, \beta_i)$. The dimension reduction theorem from [12] implies that if the input function is ε-far from unate, then the average of μ_i over all dimensions is at least $\frac{\varepsilon}{4d}$. If the tester knew which dimension had $\mu_i = \Omega(\varepsilon/d)$, it could detect a violation with high probability by querying the endpoints of $O(1/\mu_i) = O(d/\varepsilon)$ uniformly random edges. However, the tester does not know which μ_i is large and, intuitively, nonadaptively checks the following $\log d$ different scenarios, one for each $k \in [\log d]$: exactly 2^k different μ_i's are $\varepsilon/2^k$, and all others are 0. This leads to the query complexity of $O(\frac{d \log d}{\varepsilon})$.

With adaptivity, this search through $\log d$ different scenarios is not required. A pair of queries in each dimensions detects influential coordinates (i.e., dimensions with many non-constant edges), and the algorithm focuses on finding violations among those coordinates. This leads to the query complexity of $O(d/\varepsilon)$, removing the $\log d$ factor.

It is relatively easy to extend (both adaptive and nonadaptive) testers from hypercubes to hypergrids by incurring an extra factor of $\log n$ in the query complexity. The role of *i*-edges is now played by *i-lines*. An *i*-line is a set of n domain points that differ only on coordinate i. The domain $[n]$ is called a line. Monotonicity on the line (a.k.a. sortedness) can be tested with $O(\frac{\log n}{\varepsilon})$ queries, using, for example, the classical *tree tester* from [22]. Instead of sampling a random *i*-edge, we sample a random *i*-line ℓ and run the tree tester on the restriction $f_{|\ell}$ of function f to the line ℓ. This is optimal for adaptive testers, but, interestingly, not for nonadaptive testers. We show that for each function f on the line that is ε-far from unateness, one of the two scenarios happen: (1) the tree tester is likely to find a violation of unateness; (2) function f is increasing (and also

3

decreasing) on a constant fraction of pairs in $[n]$. This new angle on the classical tester allows us to replace the factor $(\log d)(\log n)$ with $\log d + \log n$ in the query complexity. Thus, the nonadaptive complexity becomes $O(d(\log d + \log n))$, which we show is optimal.

The nonadaptive lower bound. Our most significant finding is the $\log d$ gap in the query complexity between adaptive and nonadaptive testing of unateness. By previous work [23, 14], it suffices to prove lower bounds for *comparison-based* testers, i.e., testers that can only perform comparisons of the function values at queried points, but cannot use the values themselves. Our main technical contribution is the $\Omega(d \log d)$ lower bound for nonadaptive comparison-based testers of unateness on hypercube domains.

Intuitively, we wish to construct $K = \Theta(\log d)$ families of functions where, for each $k \in [K]$, functions in the k^{th} family have 2^k dimensions i with $\mu_i = \Theta(1/2^k)$, while $\mu_i = 0$ for all other dimensions. What makes the construction challenging is the existence of a *single, universal* nonadaptive $O(d)$-tester for all **b**-monotonicity properties, proven in [12]. In other words, there is a single distribution on $O(d)$ queries that defines a nonadaptive property tester for **b**-monotonicity, regardless of **b**. Since unateness is the union of all **b**-monotonicity properties, our construction must be able to fool such algorithms. Furthermore, nonadaptivity must be critical, since we obtained a $O(d)$-query adaptive tester for unateness.

Another obstacle is that once a tester finds a non-constant edge in each dimension, the problem reduces to testing **b**-monotonicity for a vector **b** determined by the directions (increasing or decreasing) of the non-constant edges. That is, intuitively, most edges in our construction must be constant. This is one of the main technical challenges. The previous lower bound constructions for monotonicity testing [8, 14] crucially used the fact that all edges in the hard functions were non-constant.

We briefly describe how we overcome the problems mentioned above. By Yao's minimax principle, it suffices to construct **Yes** and **No** distributions that a deterministic nonadaptive tester cannot distinguish. First, for some parameter m, we partition the hypercube into m subcubes based of the first $\log_2 m$ most significant coordinates. Both distributions, **Yes** and **No**, sample a uniform k from $[K]$, where $K = \Theta(\log d)$, and a set $R \subseteq [d]$ of cardinality 2^k. Furthermore, each subcube $j \in [m]$ selects an "action dimension" $r_j \in R$ uniformly at random. For both distributions, in any particular subcube j, the function value is completely determined by the coordinates *not in R*, and the random coordinate $r_j \in R$. Note that all the i-edges for $i \in (R \setminus \{r_j\})$ are constant. Within the subcube, the function is a linear function with exponentially increasing coefficients. In the **Yes** distribution, any two cubes j, j' with the same action dimension orient the edges in that dimension the same way (both increasing or both decreasing), while in the **No** distribution each cube decides on the orientation independently. The former correlation maintains unateness while the latter independence creates distance to unateness. We prove that to distinguish the distributions, any comparison-based nonadaptive tester must find two distinct subcubes with the same action dimension r_j and, furthermore, make a specific query (in both) that reveals the coefficient of r_j. We show that, with $o(d \log d)$ queries, the probability of this event is negligible.

2 Upper Bounds

In this section, we prove parts 1-2 of Theorem 1.1, starting from the hypercube domain.

Recall the definition of i-edges and i-lines from Section 1.1.1 and what it means for an edge to be increasing, decreasing, and constant.

The starting point for our algorithms is the dimension reduction theorem from [12]. It bounds the distance of $f : [n]^d \to \mathbb{R}$ to monotonicity in terms of average distances of restrictions of f to one-dimensional functions.

Theorem 2.1 (Dimension Reduction, Theorem 1.8 in [12])**.** *Fix a bit vector $\mathbf{b} \in \{0,1\}^d$ and a function $f : [n]^d \to \mathbb{R}$ which is ε-far from $\mathbf{b}$-monotonicity. For all $i \in [d]$, let μ_i be the average distance of $f_{|\ell}$ to $\mathbf{b}_i$-monotonicity over all i-lines ℓ. Then,*

$$\sum_{i=1}^{d} \mu_i \geq \frac{\varepsilon}{4}.$$

For the special case of the hypercube domains, i-lines become i-edges, and the average distance μ_i to $\mathbf{b}_i$-monotonicity is the fraction of i-edges on which the function is not $\mathbf{b}_i$-monotone.

2.1 The Nonadaptive Tester over the Hypercube

We now describe Algorithm 1, the nonadaptive tester for unateness over the hypercubes.

Algorithm 1: The Nonadaptive Unateness Tester over Hypercubes

    input : distance parameter $\varepsilon \in (0, 1/2)$; query access to a function $f : \{0,1\}^d \to \mathbb{R}$.

1 for $r = 1$ *to* $\lceil 3 \log(4d/\varepsilon) \rceil$ **do**

2   **repeat** $s_r = \lceil \frac{16d \ln 4}{\varepsilon \cdot 2^r} \rceil$ **times**

3   Sample a dimension $i \in [d]$ uniformly at random.

4   Sample $3 \cdot 2^r$ i-edges uniformly and independently at random and **reject** if there exists an increasing edge and a decreasing edge among the sampled edges.

5 accept

It is evident that Algorithm 1 is a nonadaptive, one-sided error tester. Furthermore, its query complexity is $O\left(\frac{d}{\varepsilon} \log \frac{d}{\varepsilon}\right)$. It suffices to prove the following.

Lemma 2.2. *If f is ε-far from unate, Algorithm 1 rejects with probability at least $2/3$.*

Proof. Recall that α_i is the fraction of i-edges that are decreasing, β_i is the fraction of i-edges that are increasing and $\mu_i = \min(\alpha_i, \beta_i)$.

Define the d-dimensional bit vector $\mathbf{b}$ as follows: for each $i \in [d]$, let $\mathbf{b}_i = 0$ if $\alpha_i < \beta_i$ and $\mathbf{b}_i = 1$ otherwise. Observe that the average distance of f to $\mathbf{b}_i$-monotonicity over a random i-edge is precisely μ_i. Since f is ε-far from being unate, f is also ε-far from being $\mathbf{b}$-monotone. By Theorem 2.1, $\sum_{i \in [d]} \mu_i \geq \frac{\varepsilon}{4}$. Hence, $\mathbf{E}_{i \in [d]}[\mu_i] \geq \frac{\varepsilon}{4d}$. We now apply the work investment strategy due to Berman et al. [6] to get an upper bound on the probability that Algorithm 1 fails to reject.

Theorem 2.3 ([6])**.** *For a random variable $X \in [0, 1]$ with $\mathbf{E}[X] \geq \mu$ for $\mu < \frac{1}{2}$, let $p_r = \Pr[X \geq 2^{-r}]$ and $\delta \in (0, 1)$ be the desired error probability. Let $s_r = \frac{4 \ln 1/\delta}{\mu \cdot 2^r}$. Then,*

$$\prod_{i=1}^{\lceil 3 \log(1/\mu) \rceil} (1 - p_r)^{s_r} \leq \delta.$$

Consider running Algorithm 1 on a function f that is ε-far from unate. Let $X = \mu_i$ where i is sampled uniformly at random from $[d]$. Then $\mathbf{E}[X] \geq \frac{\varepsilon}{4d}$. Applying the work investment strategy (Theorem 2.3) on X with $\mu = \frac{\varepsilon}{4d}$, we get that the probability that, in some iteration, Step 3 samples a dimension i such that $\mu_i \geq 2^{-r}$ is at least $1 - \delta$. We set $\delta = 1/4$. Conditioned on sampling such a dimension, the probability that Step 4 fails to obtain an increasing edge and a decreasing edge among its $3 \cdot 2^r$ samples is at most $2 \left(1 - 2^{-r}\right)^{3 \cdot 2^r} \leq 2e^{-3} < 1/9$, as the fraction of both increasing and decreasing edges in the dimension is at least 2^{-r}. Hence, the probability that Algorithm 1 rejects f is at least $\frac{3}{4} \cdot \frac{8}{9} = \frac{2}{3}$, which completes the proof of Lemma 2.2. $\qquad\square$

2.2 The Adaptive Tester over the Hypercube

We now describe Algorithm 2, an adaptive tester for unateness over the hypercube domain with good expected query complexity. The final tester is obtained by repeating this tester and accepting if the number of queries exceeds a specified bound.

Algorithm 2: The Adaptive Unateness Tester over Hypercubes

    input : distance parameter $\varepsilon \in (0, 1/2)$; query access to a function $f : \{0,1\}^d \to \mathbb{R}$.

1 repeat $10/\varepsilon$ **times**
2   **for** $i = 1$ *to* d **do**
3   Sample an i-edge e_i uniformly at random.
4   **if** *e_i is non-constant (i.e., increasing or decreasing)* **then**
5   Sample i-edges uniformly at random till we obtain a non-constant edge e_i'.
6   **reject** if one of the edges e_i, e_i' is increasing and the other is decreasing.
7 accept

Claim 2.4. *The expected number of queries made by Algorithm 2 is $40d/\varepsilon$.*

Proof. Consider one iteration of the **repeat**-loop in Step 1. We prove that the expected number of queries in this iteration is $4d$. The total number of queries in Step 3 is $2d$, as 2 points per dimension are queried. Let E_i be the event that edge e_i is non-constant and T_i be the random variable for the number of i-edges sampled in Step 5. Then $\mathbf{E}[T_i] = \frac{1}{\alpha_i + \beta_i} = \frac{1}{\Pr[E_i]}$. Therefore, the expected number of all edges sampled in Step 5 is $\sum_{i=1}^{d} \Pr[E_i] \cdot \mathbf{E}[T_i] = \sum_{i=1}^{d} \Pr[E_i] \cdot \frac{1}{\Pr[E_i]} = d$. Hence, the expected number of queries in Step 5 is $2d$. Since there are $10/\varepsilon$ iterations in Step 1, the expected number of queries in Algorithm 2 is $40d/\varepsilon$. $\qquad\square$

Claim 2.5. *If f is ε-far from unate, Algorithm 2 accepts with probability at most $1/6$.*

Proof. First, we bound the probability that a violation of unateness is detected in some dimension $i \in [d]$ in one iteration of the **repeat**-loop. Consider the probability of finding a decreasing i-edge in Step 3, and an increasing i-edge in Step 5. The former is exactly α_i, and the latter is $\frac{\beta_i}{\alpha_i + \beta_i}$. Therefore, the probability we detect a violation from dimension i is $\frac{2\alpha_i\beta_i}{\alpha_i + \beta_i} \geq \min(\alpha_i, \beta_i) = \mu_i$. The probability that we fail to detect a violation in any of the d dimensions is at most $\prod_{i=1}^{d}(1 - \mu_i) \leq \exp\left(-\sum_{i=1}^{d} \mu_i\right)$, which is at most $e^{-\varepsilon/4}$ by Theorem 2.1 (Dimension Reduction). By Taylor expansion of $e^{-\varepsilon/4}$, the probability of finding a violation in one iteration is at least $1 - e^{-\varepsilon/4} \geq \frac{\varepsilon}{4} - \frac{\varepsilon^2}{32} > \frac{\varepsilon}{5}$. The probability that Algorithm 2 does not reject in any iteration is at most $(1 - \varepsilon/5)^{10/\varepsilon} < 1/6$. $\qquad\square$

6

Proof of Theorem 1.1, Part 2 (for the special case of the hypercube domain). We run Algorithm 2, aborting and accepting if we ever make more than $240d/\varepsilon$ queries. By Markov's inequality, the probability of aborting is at most $1/6$. By Claim 2.5, if f is ε-far from unate, Algorithm 2 accepts with probability at most $1/6$. The theorem follows by a union bound. □

2.3 Extension to Hypergrids

We start by establishing terminology for lines and pairs. Consider a function $f : [n]^d \to \mathbb{R}$. Recall the definition of i-lines from Section 1.1.1. A pair of points that differ only in coordinate i is called an i-pair. An i-pair (x, y) with $x_i < y_i$ is called *increasing* if $f(x) < f(y)$, *decreasing* if $f(x) > f(y)$, and *constant* if $f(x) = f(y)$.

Algorithm 3: Tree Tester

    input : Query access to a function $h : [n] \mapsto \mathbb{R}$.

1 Pick $x \in [n]$ uniformly at random.
2 Let $Q_x \subseteq [n]$ be the set of points visited in a binary search for x. Query h on all points in Q_x.
3 If there is an increasing pair in Q_x, set $\mathsf{dir} \leftarrow \{\uparrow\}$; otherwise, $\mathsf{dir} \leftarrow \emptyset$.
4 If there is a decreasing pair in Q_x, update $\mathsf{dir} \leftarrow \mathsf{dir} \cup \{\downarrow\}$.
5 **Return** dir.

The main tool for extending Algorithms 1 and 2 to work on hypergrids is the *tree tester*, designed by Ergun et al. [22] to test monotonicity of functions $h : [n] \to \mathbb{R}$. We modify the tree tester to return information about directions it observed instead of just accepting or rejecting. See Algorithm 3. The following lemma summarizes the guarantee of the tree tester.

Lemma 2.6 ([22, 12]). *If $h : [n] \mapsto \mathbb{R}$ is ε-far from monotone (respectively, antimonotone), then the output of Algorithm 3 on h contains $\downarrow$ (respectively, $\uparrow$) with probability at least ε.*

Algorithm 4: The Adaptive Unateness Tester over Hypergrids

    input : distance parameter $\varepsilon \in (0, 1/2)$; query access to a function $f : [n]^d \to \mathbb{R}$.

1 **repeat** $10/\varepsilon$ **times**
2   **for** $i = 1$ *to* d **do**
3   Sample an i-line ℓ_i uniformly at random.
4   Let dir_i be the output of Algorithm 3 on $f_{|\ell_i}$.
5   **if** $\mathsf{dir}_i \neq \emptyset$ **then**
6   Sample i-lines uniformly at random and run Algorithm 3 on f restricted to each line until it returns a non-empty set. Call it dir'_i.
7   If $\mathsf{dir}_i \cup \mathsf{dir}'_i = \{\uparrow, \downarrow\}$, **reject**.
8 **accept**

Our hypergrid testers are stated in Algorithms 4 and 5. Next, we explain how Lemma 2.6 and Theorem 2.1 are used in the analysis of the adaptive tester. For a dimension $i \in [d]$, let α_i and β_i denote the average distance of $f_{|\ell}$ to monotonicity and antimonotonicity, respectively, over all

---

Algorithm 5: The Nonadaptive Unateness Tester over Hypergrids

> **input** : distance parameter $\varepsilon \in (0, 1/2)$; query access to a function $f : [n]^d \to \mathbb{R}$.

1 repeat $220/\varepsilon$ **times**

2   **for** $i = 1$ *to* d **do**

3   Sample an i-line ℓ uniformly at random.

4   **Reject** if Algorithm 3, on input $f_{|\ell}$, returns $\{\uparrow, \downarrow\}$.

5 for $r = 1$ *to* $\lceil 3 \log(200d/\varepsilon) \rceil$ **do**

6   **repeat** $s_r = \lceil \frac{800d \ln 4}{\varepsilon \cdot 2^r} \rceil$ **times**

7   Sample a dimension $i \in [d]$ uniformly at random.

8   Sample $3 \cdot 2^r$ i-pairs uniformly and independently at random.

9   If we find an increasing and a decreasing pair among the sampled pairs, **reject**.

10 accept

---

i-lines ℓ. Then $\mu_i := \min(\alpha_i, \beta_i)$ is the average fraction of points per i-line that needs to change to make f unate. Define the **b**-vector with $\mathbf{b}_i = 0$ if $\alpha_i < \beta_i$, and $\mathbf{b}_i = 1$ otherwise. By Theorem 2.1, if f is ε-far from unate, and thus ε-far from **b**-monotone, then $\sum_{i=1}^{d} \mu_i \geq \varepsilon/4$. By Lemma 2.6, the probability that the output of Algorithm 3 on $f_{|\ell}$ contains $\downarrow$ (respectively, $\uparrow$), where ℓ is a uniformly random i-line, is at least α_i (respectively, β_i). The rest of the analysis of Algorithm 4 is similar to that in the hypercube case.

Proof of Theorem 1.1, Part 2. The tester is Algorithm 4. As in the proof of Claim 2.4, the expected running time of Algorithm 4 is at most $(40d \log n)/\varepsilon$. The proof of Claim 2.5 carries over almost word-to-word. Fix dimension i. The probability that $\downarrow \in \mathsf{dir}_i$ in Step 4 is at least α_i. The probability that $\uparrow \in \mathsf{dir}'_i$ in Step 6 is at least $\frac{\beta_i}{\alpha_i + \beta_i}$. The rest of the calculation is identical to that of the proof of Claim 2.5. $\qquad\square$

To analyze the nonadaptive tester, we prove Lemma 2.7, which demonstrates the power of the tree tester and may be of independent interest.

Lemma 2.7. *Consider a function $h : [n] \to \mathbb{R}$ which is ε-far from monotone (respectively, anti-monotone). At least one of the following holds:*

    1. $\Pr[\text{Algorithm 3, on input } h, \text{ returns } \{\uparrow, \downarrow\}] \geq \varepsilon/25$.

    2. $\Pr_{u,v \in [n]}[(u, v) \text{ is a decreasing (respectively, increasing) pair}] \geq \varepsilon/25$.

Proof. Let T be a balanced binary search tree consisting of elements in $[n]$, such that the set of points visited in a binary search for some $x \in [n]$ corresponds to a path from the root to the node containing x in T. Let Q_x denote the set of points visited in a binary search for $x \in [n]$. For $x, y \in [n]$, denote the least common ancestor of x and y by $\mathsf{lca}(x, y)$.

Let $W_{\uparrow\downarrow}$ be a set of points x such that Q_x contains both an increasing and a decreasing pair (with respect to h). If $|W_{\uparrow\downarrow}| \geq \frac{\varepsilon n}{10}$, then Case 1 of Lemma 2.7 holds. We may therefore assume that $|W_{\uparrow\downarrow}| < \frac{\varepsilon n}{10}$. Let $\mathcal{E}$ be the event that for any $u, v \in [n]$ such that $u < v$, the pair (u, v) is decreasing. We will prove that $\Pr[\mathcal{E}] \geq \varepsilon/25$.

Let $W_{\downarrow}$ be that set of points $x \in [n]$ such that Q_x contains a decreasing pair. Similarly, define the set $W_{\uparrow}$. Let W_c denote the set of points x such that $h_{|Q_x}$ is constant.

Claim 2.8 ([22]). *The function h restricted to the set $W_{\uparrow} \cup W_c$ is monotone.*

8

Proof. The proof is by contradiction. Suppose $x, y \in (W_\uparrow \cup W_c)$ such that $x < y$, but $h(x) > h(y)$. Consider $z = \mathsf{lca}(x, y)$. Either $h(x) > h(z)$ or $h(z) > h(y)$, contradicting the fact that $x, y \in W_\uparrow \cup W_c$. $\qquad\square$

By symmetry, the function h restricted to the set $W_\downarrow \cup W_c$ is antimonotone.

A priori, points in $W_\uparrow$ and $W_\downarrow$ could be interspersed. The next claim shows that they are in different halves of the tree T.

Claim 2.9. *If $x \in W_\downarrow$ and $y \in W_\uparrow$, then $\mathsf{lca}(x, y)$ is the root of T (which is equal to $\lceil n/2 \rceil$).*

Proof. Suppose not. Let $z := \mathsf{lca}(x, y)$ and w be the parent of z. Consider the case where z is the left child of w, x lies in the left subtree of z and y lies in the right subtree of z. (All the other cases have analogous proofs.) Observe that all points in Q_y lie in the interval $[z, w]$. Both w and z are in Q_x as well as in Q_y. As $x \in W_\uparrow$ and $y \in W_\downarrow$, it must be the case that $h(w) = h(z)$. Since $y \notin W_{\uparrow\downarrow}$, for all $p \in Q_y$, we have $h(p) = h(w)$. This contradicts the fact that $y \in W_\uparrow$.

In all cases, we conclude that either $x \notin W_\downarrow$ or $y \notin W_\uparrow$. Thus, z cannot have a parent, and $z = \lceil n/2 \rceil$. $\qquad\square$

Claim 2.10. *Let $g : [n] \mapsto \mathbb{R}$ be an antimonotone function and $\mathsf{dist}(g, \text{constant})$ denote the fraction of points that need to be changed so that g is a constant function. If g is antimonotone, and $\mathsf{dist}(g, \text{constant}) \geq \rho$, where $\rho \leq \frac{1}{2}$, then*

$$\Pr_{u,v \in [n]:u<v}[(u, v) \text{ is decreasing}] \geq \frac{\rho}{2}.$$

Proof. The probability that $g(u) \neq g(v)$ is at least $\rho(1 - \rho)$ which is at least $\frac{\rho}{2}$ when $\rho \leq \frac{1}{2}$. Since g is antimonotone, (u, v) is a decreasing pair. $\qquad\square$

Let L (respectively, R) be the set of points in $[n] \setminus W_{\uparrow\downarrow}$ in the left (respectively, right) subtree of the root. Define $\mu_L := |L|/n$; similarly, define μ_R. Observe that both μ_L and μ_R are at least $\frac{1}{2} - \frac{\varepsilon}{10}$. By Claims 2.8 and 2.9, $h_{|L}$ (and $h_{|R}$) is either monotone or antimonotone. Now, if any of these two functions were antimonotone and $\frac{\varepsilon}{2}$-far from being constant (w.l.o.g., assume $h_{|L}$ satisfies the condition), then by Claim 2.10, we would have

$$\Pr[\mathcal{E}] \geq \Pr_{u<v}[(u, v) \text{ is decreasing and } u, v \in L] \geq \frac{\varepsilon}{4} \cdot \left(\frac{1}{2} - \frac{\varepsilon}{10}\right)^2 \geq \frac{\varepsilon}{25}.$$

Assume that this doesn't occur. We have two cases.

Case 1. Both $h_{|L}$ and $h_{|R}$ are $\frac{\varepsilon}{2}$-close[3] to being constant. In this case, at least $(1 - \frac{\varepsilon}{2})|L|$ points of L evaluate to a constant C_1, and at least $(1 - \frac{\varepsilon}{2})|R|$ points of R evaluate to constant C_2. We must have $C_1 > C_2$, for otherwise, we can make h monotone by changing only $\frac{\varepsilon}{2} \cdot (|R| + |L|) + \frac{\varepsilon n}{10} < \varepsilon n$ points, which is a contradiction. Hence,

$$\Pr[\mathcal{E}] \geq \Pr_{u<v}[h(u) = C_1 \text{ and } h(v) = C_2] \geq \left(1 - \frac{\varepsilon}{2}\right)^2 \mu_L \mu_R > \frac{1}{4} \cdot \left(\frac{1}{2} - \frac{\varepsilon}{10}\right)^2 \geq \frac{\varepsilon}{25}.$$

---

[3]A function h is ε-close to a property $\mathcal{P}$ if it is sufficient to change at most ε-fraction of values in h to make it satisfy $\mathcal{P}$.

Case 2. At least one of the functions is $\frac{\varepsilon}{2}$-far from being constant and is monotone. W.l.o.g., assume $h_{|L}$ satisfies this condition. Note that all points in L are only in $W_\uparrow \cup W_c$, and so, all points in R must be in $W_\downarrow \cup W_c$. This implies that $h_{|R}$ is antimonotone. (Note that a constant function is also antimonotone.) But then, $h_{|R}$ must be $\frac{\varepsilon}{2}$-close to being constant. Then at least $(1 - \frac{\varepsilon}{2})|R|$ points in R evaluate to a constant, say C. Let U denote the set of points in L whose values are strictly greater than C. Since $h_{|L}$ is monotone, we can make h monotone by deleting all points in $U, W_{\uparrow\downarrow}$, and the points in R that do not evaluate to C. The total number of points to be deleted is at most $|U| + \frac{\varepsilon n}{10} + \frac{\varepsilon n}{2}$, which must be at least εn, as h is ε-far from monotone. Hence, $|U| > \varepsilon n/3$. Therefore,

$$\Pr[\mathcal{E}] \geq \Pr_{u < v}\left[u \in U \text{ and } h(v) = C\right] \geq \frac{\varepsilon}{3} \cdot \left(1 - \frac{\varepsilon}{2}\right) \mu_R > \frac{\varepsilon}{25}.$$

This completes the proof of Lemma 2.7. $\qquad\square$

We now analyze Algorithm 5. It is evident that it has one-sided error and makes $O(\frac{d}{\varepsilon}(\log n + \log \frac{d}{\varepsilon}))$ queries. It suffices to prove the following.

Theorem 2.11. *If $f : [n]^d \mapsto \mathbb{R}$ is ε-far from unate, then Algorithm 5 rejects with probability at least $2/3$.*

Proof. For any line ℓ, we define the following quantities.
- α_ℓ: the distance of $f_{|\ell}$ to monotonicity.
- β_ℓ: the distance of $f_{|\ell}$ to antimonotonicity.
- σ_ℓ: the probability that Algorithm 3, on input $f_{|\ell}$, returns $\{\uparrow, \downarrow\}$.
- δ_ℓ: the probability that a uniformly random pair in ℓ is decreasing.
- λ_ℓ: the probability that a uniformly random pair in ℓ is increasing.

Let L_i be the set of i-lines. By Theorem 2.1,

$$\frac{1}{n^{d-1}} \sum_{i=1}^{d} \min\left(\sum_{\ell \in L_i} \alpha_\ell, \sum_{\ell \in L_i} \beta_\ell\right) \geq \frac{\varepsilon}{4}.$$

By Lemma 2.7, for every line ℓ, we have $\sigma_\ell + \delta_\ell \geq \alpha_\ell/25$ and $\sigma_\ell + \lambda_\ell \geq \beta_\ell/25$. Also note,

$$\frac{1}{n^{d-1}} \sum_{i=1}^{d} \left[\sum_{\ell \in L_i} \sigma_\ell + \min\left(\sum_{\ell \in L_i} \delta_\ell, \sum_{\ell \in L_i} \lambda_\ell\right)\right] \geq \frac{1}{n^{d-1}} \sum_{i=1}^{d} \min\left(\sum_{\ell \in L_i}(\sigma_\ell + \delta_\ell), \sum_{\ell \in L_i}(\sigma_\ell + \lambda_\ell)\right)$$

Combining these bounds, we obtain that the LHS is at least $\varepsilon/100$. Note that the first term, which is equal to $\sum_{i=1}^{d} \mathbf{E}_{\ell \in L_i}[\sigma_\ell]$, is the expected number of times a single iteration of Steps 2-4 rejects. If this quantity is at least $\varepsilon/200$, then the tester rejects with probability at least $2/3$. If not, then we have $n^{-(d-1)} \sum_{i=1}^{d} \min(\sum_{\ell \in L_i} \delta_\ell, \sum_{\ell \in L_i} \lambda_\ell) \geq \varepsilon/200$. Using a calculation identical to that of the proof of Lemma 2.2, the probability that Step 9 rejects in some iteration is at least $2/3$. $\qquad\square$

3 The Lower Bound for Nonadaptive Testers over Hypercubes

In this section, we prove Theorem 1.2, which gives a lower bound for nonadaptive unateness testers for functions over the hypercube.

Previous work of [14] on lower bounds for monotonicity testing shows that, for a special class of properties, which includes unateness, it is sufficient to prove lower bounds for *comparison-based testers*. Comparison-based testers base their decisions only on the *order* of the function values at queried points, and not on the values themselves.

We first state the reduction to comparison-based testers from [14]. Let a (t, ε, δ)-tester for a property $\mathcal{P}$ be a t-query tester, with distance parameter ε, that errs with (two-sided) probability at most δ. Consider functions of the form $f : D \to \mathbb{R}$, where D is an arbitrary partial order (in particular the hypergrid/cube). A property $\mathcal{P}$ is invariant under monotone transformations if, for all strictly increasing maps $\phi : \mathbb{R} \to \mathbb{R}$ and all functions f, $\mathsf{dist}(f, \mathcal{P}) = \mathsf{dist}(\phi \circ f, \mathcal{P})$. In particular, unateness is invariant under monotone transformations.

Theorem 3.1 (implicit in Theorem 2.1 of [14]). *Let $\mathcal{P}$ be a property invariant under monotone transformations. Suppose there exists a nonadaptive (resp., adaptive) (t, ε, δ)-tester for $\mathcal{P}$. Then there exists a nonadaptive (resp., adaptive) comparison-based $(t, \varepsilon, 2\delta)$-tester for $\mathcal{P}$.*

Our main lower bound theorem is stated next. In the light of the previous discussion, it implies Theorem 1.2.

Theorem 3.2. *Any nonadaptive comparison-based tester for unateness of functions $f : \{0, 1\}^d \to \mathbb{R}$ must make $\Omega(d \log d)$ queries.*

By Theorem 3.1 and Yao's minimax principle [37], it suffices to prove the lower bound for deterministic, nonadaptive, comparison-based testers over a known distribution of functions. It may be useful for the reader to recall the sketch of the main ideas given in Section 1.1.1. For convenience, assume d is a power of 2 and let $d' := d + \log_2 d$. We will focus on functions $h : \{0, 1\}^{d'} \to \mathbb{R}$, and prove the lower bound of $\Omega(d \log d)$ for this class of functions, as $\Omega(d \log d) = \Omega(d' \log d')$.

3.1   The Hard Distributions

We first partition $\{0, 1\}^{d'}$ into d subcubes based on the most significant $\log_2 d$ bits. More precisely, for $i \in [d]$, the i^{th} subcube is defined as

$$C_i := \{x \in \{0, 1\}^{d'} : \mathsf{val}(x_{d'} x_{d'-1} \cdots x_{d+1}) = i - 1\},$$

where $\mathsf{val}(z)$ denotes the integer equivalent to the binary string z. Specifically, $\mathsf{val}(z_p z_{p-1} \ldots z_1) = \sum_{i=1}^{p} z_i 2^{i-1}$.

Let $m = d$. We denote the set of indices of the subcube by $[m]$ and the set of dimensions by $[d]$. We use $i, j \in [m]$ to index subcubes, and $a, b \in [d]$ to index dimensions. We now define a series of random variables, where each subsequent variable may depend on the previous ones.

- k: a number picked uniformly at random from $\left[\frac{1}{2} \log_2 d\right]$.
- R: a uniformly random subset of $[d]$ of size 2^k.
- r_i: for each $i \in [m]$, r_i is picked from R uniformly and independently at random.
- α_b: for each $b \in [d]$, α_b is picked from $\{-1, +1\}$ uniformly and independently at random. (Note: α_b only needs to be defined for each $b \in R$. We define it over $[d]$ just so that it is independent of R.)
- β_i: for each $i \in [m]$, β_i is picked from $\{-1, +1\}$ uniformly and independently at random.

11

We denote by $\boldsymbol{S}$ the tuple $(k, R, \{r_i\})$, also referred to as the *shared randomness*. We use $\boldsymbol{T}$ to refer to the entire set of random variables. Given $\boldsymbol{T}$, define the following functions:

$$f_{\boldsymbol{T}}(x) := \sum_{b \in [d'] \setminus R} x_b 3^b + \alpha_{r_i} \cdot x_{r_i} 3^{r_i}, \text{ where } i \text{ is the subcube with } x \in C_i.$$

$$g_{\boldsymbol{T}}(x) := \sum_{b \in [d'] \setminus R} x_b 3^b + \beta_i \cdot x_{r_i} 3^{r_i}, \text{ where } i \text{ is the subcube with } x \in C_i.$$

The distribution **Yes** generates $f_{\boldsymbol{T}}$ and the distribution **No** generates $g_{\boldsymbol{T}}$.

In all cases, the function restricted to any subcube C_i is linear. Consider some dimension $b \in R$. There can be numerous r_i's that are equal to b. For $f_{\boldsymbol{T}}$, in all of these subcubes, the coefficient of x_{r_i} has the same sign, namely α_{r_i}. For $g_{\boldsymbol{T}}$, the coefficient β_i is potentially different, as it depends on the actual subcube.

Claim 3.3. *Every $f \in \mathsf{supp}(\mathbf{Yes})$ is unate.*

Proof. Fix some $f \in \mathsf{supp}(\mathbf{Yes})$. Since f restricted to any C_i is linear, it suffices to argue that the coefficient of any x_b (when it is non-zero) has the same sign, in all C_i's. For any $b \in [d'] \setminus R$, the coefficient of x_b is always 3^b. If $b \in R$, then the coefficient is either 0 or $3^b \alpha_b$. $\square$

Claim 3.4. *A function $g \sim \mathbf{No}$ is $\frac{1}{8}$-far from unate with probability at least $9/10$.*

Proof. Fix $\boldsymbol{T} = \mathcal{T}$. Condition on any choice of k and R. Note that $|R| \leq \sqrt{d}$. For any $r \in R$, let $A_r = \{i : r_i = r\}$ denote the set of subcube indices with $r_i = r$. Observe that $\mathbf{E}[|A_r|] \geq m/\sqrt{d} = \sqrt{d}$. By a Chernoff bound and union bound, for all $r \in R$, we have $|A_r| \geq \sqrt{d}/2$ with probability at least $1 - d \exp(-\sqrt{d}/8)$.

Condition on the event that $|A_r| \geq \sqrt{d}/2$ for all $r \in R$. For each $i \in A_r$, there is a random choice of β_i. Partition A_r into A_r^+ and A_r^-, depending on whether β_i is $+1$ or -1. Again, by a Chernoff bound and union bound, for all $r \in R$, we have $\min(|A_r^+|, |A_r^-|) \geq |A_r|/4$ with probability at least $1 - d \exp(-\sqrt{d}/32)$. Thus, we can assume the above event holds with probability at least $1 - d(\exp(-\sqrt{d}/8) + \exp(-\sqrt{d}/32))$, which is at least $9/10$, for large enough d and for any choice of k and R.

Denote the size of any subcube C_i by s. In $g_{\mathcal{T}}$, for all $i \in A_r^+$, all r-edges in C_i are increasing, whereas, for all $j \in A_r^-$, all r-edges in C_j are decreasing. To make $g_{\mathcal{T}}$ unate, we must make all these edges have the same direction (i.e., increasing or decreasing). This requires modifying at least $\frac{s}{2} \cdot \min(|A_r^+|, |A_r^-|) \geq \frac{s|A_r|}{8}$ values in $g_{\mathcal{T}}$. Summing over all r, we need to change at least $\frac{s}{8} \sum_r |A_r|$ values. Since the A_r's partition the set of subcubes, this corresponds to at least a $\frac{1}{8}$-fraction of the domain. $\square$

3.2 From Functions to Signed Graphs that are Hard to Distinguish

For convenience, denote $x \prec y$ if $\mathsf{val}(x) < \mathsf{val}(y)$. Note that $\prec$ forms a total ordering on $\{0,1\}^{d'}$. Given $x \prec y \in \{0,1\}^{d'}$ and a function $h : \{0,1\}^{d'} \to \mathbb{R}$, define $\mathsf{sgn}_h(x, y)$ to be 1 if $h(x) < h(y)$, 0 if $h(x) = h(y)$, and -1 if $h(x) > h(y)$.

Any deterministic, nonadaptive, comparison-based tester is defined as follows: It makes a set of queries Q and decides whether or not the input function h is unate depending on the $\binom{|Q|}{2}$-comparisons in Q. More precisely, for every pair $(x, y) \in Q \times Q$, $x \prec y$, we insert an edge labeled

with $\mathsf{sgn}_h(x,y)$. Let this signed graph be called G_h^Q. Any nonadaptive, comparison-based algorithm can be described as a partition of the universe of all signed graphs over Q into $\mathcal{G}_Y$ and $\mathcal{G}_N$. The algorithm accepts the function h iff $G_h^Q \in \mathcal{G}_Y$.

Let $\boldsymbol{G}_Y^Q$ be the distribution of the signed graphs G_h^Q when $h \sim \textbf{Yes}$. Similarly, define $\boldsymbol{G}_N^Q$ when $h \sim \textbf{No}$. Our main technical theorem is Theorem 3.5, which is proved in Section 3.3.

Theorem 3.5. *For small enough $\delta > 0$ and large enough d, if $|Q| \leq \delta d \log d$, then $\|\boldsymbol{G}_Y^Q - \boldsymbol{G}_N^Q\|_{\mathrm{TV}} = O(\delta)$.*

We now prove that Theorem 3.5 implies Theorem 3.2, the main lower bound.

Proof of Theorem 3.2. Consider the distribution over functions where with probability $1/2$, we sample from **Yes** and with the remaining probability we sample from **No**. By Theorem 3.1 and Yao's minimax principle, it suffices to prove that any deterministic, nonadaptive, comparison-based tester making at most $\delta d \log d$ queries (for small enough $\delta > 0$) errs with probability at least $1/3$. Now, note that

$$\Pr[\text{error}] = \frac{1}{2} \cdot \Pr_{h \sim \textbf{Yes}}[G_h^Q \in \mathcal{G}_N] + \frac{1}{2} \cdot \Pr_{h \sim \textbf{No}}[G_h^Q \in \mathcal{G}_Y \text{ and } h \text{ is } \frac{1}{8}\text{-far from unate}].$$

By Theorem 3.5, the first term is at least $\frac{1}{2} \cdot \left(\Pr_{h \sim \textbf{No}}[G_h^Q \in \mathcal{G}_N] - O(\delta) \right)$, and by Claim 3.4, the second term is at least $\frac{1}{2} \cdot \left(\Pr_{h \sim \textbf{Yes}}[G_h^Q \in \mathcal{G}_Y] - O(\delta) - \frac{1}{10} \right)$. Summing them up, we get $\Pr[\text{error}] \geq \frac{1}{2} - O(\delta) - \frac{1}{20}$ which is at least $\frac{1}{3}$ for small enough δ. $\square$

The proof of Theorem 3.5 is naturally tied to the behavior of sgn_h. Ideally, we would like to say that $\mathsf{sgn}_h(x,y)$ is almost identical regardless of whether $h \sim \textbf{Yes}$ or $h \sim \textbf{No}$. Towards this, we determine exactly the set of pairs (x,y) that potentially differentiate **Yes** and **No**.

Claim 3.6. *For all $h \in \mathsf{supp}(\textbf{Yes}) \cup \mathsf{supp}(\textbf{No})$, for all $x \in C_i$ and $y \in C_j$ such that $i < j$, we have $\mathsf{sgn}_h(x,y) = 1$.*

Proof. For any h, we can write $h(x)$ as $\sum_{b>d} 3^b \cdot x_b + \sum_{b \leq d} c_b(x) \cdot 3^b \cdot x_b$, where $c_b : \{0,1\}^{d'} \to \{-1, 0, +1\}$. Thus, $h(y) - h(x) = \sum_{b>d} 3^b(y_b - x_b) + \sum_{b \leq d} 3^b(c_b(y) \cdot y_b - c_b(x) \cdot x_b)$. Recall that $x \in C_i, y \in C_j$, and $j > i$. Let q denote the most significant bit of difference between x and y. We have $q > d$, and $y_q = 1$ and $x_q = 0$. Note that for $b \leq d$, $|c_b(y) \cdot y_b - c_b(x) \cdot x_b)| \leq 2$. Thus, $h(y) - h(x) \geq 3^q - 2\sum_{b<q} 3^b > 0$. $\square$

Thus, comparisons between points in different subcubes reveal no information about which distribution h was generated from. Therefore, the "interesting" pairs that can distinguish whether $h \sim \textbf{Yes}$ or $h \sim \textbf{No}$ must lie in the same subcube. The next claim shows a further criterion that is needed for a pair to be interesting. We first define another notation.

Definition 3.7. *For any setting of the shared randomness $\boldsymbol{S}$, subcube C_i, and points $x, y \in C_i$, we define $t_{\boldsymbol{S}}^i(x,y)$ to be the most significant coordinate of difference (between x, y) in $([d] \setminus R) \cup \{r_i\}$.*

Note that $\boldsymbol{S}$ determines R and $\{r_i\}$. For any $\boldsymbol{T}$ that extends $\boldsymbol{S}$ and any function, the restriction to C_i is unaffected by the coordinates in $R \setminus r_i$. Thus, $t_{\boldsymbol{S}}^i(x,y)$ is the first coordinate of difference that is influential in C_i.

Claim 3.8. *Fix some $\boldsymbol{S}$, subcube C_i, and points $x, y \in C_i$. Let $c = t^i_{\boldsymbol{S}}(x, y)$, and assume $x \prec y$. For any $\boldsymbol{T}$ that extends $\boldsymbol{S}$:*

- *If $c \neq r_i$, then $\mathsf{sgn}_{f_{\boldsymbol{T}}}(x, y) = \mathsf{sgn}_{g_{\boldsymbol{T}}}(x, y) = 1$.*
- *If $c = r_i$, $\mathsf{sgn}_{f_{\boldsymbol{T}}}(x, y) = \alpha_c$ and $\mathsf{sgn}_{g_{\boldsymbol{T}}}(x, y) = \beta_i$.*

Proof. Assume $x \in C_i$. Recall that $f_{\boldsymbol{T}}(x) = \sum_{b \in [d'] \setminus R} x_b 3^b + \alpha_{r_i} \cdot x_{r_i} 3^{r_i}$ and $g_{\boldsymbol{T}}(x) = \sum_{b \in [d'] \setminus R} x_b 3^b + \beta_i \cdot x_{r_i} 3^{r_i}$.

First, consider the case $c \neq r_i$. Thus, $c \notin R$. Observe that $x_b = y_b$, for all $b > c$ such that $b \notin R$. Furthermore, $x_c = 0$ and $y_c = 1$. Thus, $f_{\boldsymbol{T}}(y) - f_{\boldsymbol{T}}(x) > 3^c - \sum_{b < c} 3^b > 0$. An identical argument holds for $g_{\boldsymbol{T}}$.

Now, consider the case $c = r_i$. Thus, $f_{\boldsymbol{T}}(y) - f_{\boldsymbol{T}}(x) = \alpha_c 3^c + \sum_{b < c, b \notin R} (y_b - x_b) 3^b$. Using the same geometric series arguments as above, $\mathsf{sgn}_{f_{\boldsymbol{T}}}(x, y) = \alpha_c$. An analogous argument shows that $\mathsf{sgn}_{g_{\boldsymbol{T}}}(x, y) = \beta_i$ $\qquad \square$

3.3 Proving Theorem 3.5: Good and Bad Events

For a given Q, we first identify certain "bad" values for $\boldsymbol{S}$, on which Q could potentially distinguish between $f_{\boldsymbol{S}}$ and $g_{\boldsymbol{S}}$. We will prove that the probability of a bad $\boldsymbol{S}$ is small for a given Q. Furthermore, we show that Q cannot distinguish between $f_{\boldsymbol{S}}$ and $g_{\boldsymbol{S}}$ for any good $\boldsymbol{S}$. We set up some definitions.

Definition 3.9. *Given a pair (x, y), define $\mathsf{cap}(x, y)$ to be the 5 most significant coordinates[4] in which they differ. We say (x, y) captures these coordinates. For any set $S \subseteq \{0, 1\}^{d'}$, define $\mathsf{cap}(S) := \bigcup_{x, y \in S} \mathsf{cap}(x, y)$ to be the coordinates captured by the set S.*

Fix any Q. We set $Q_i := Q \cap C_i$. We define two bad events for $\boldsymbol{S}$.

- **Abort Event $\mathcal{A}$:** There exists $x, y \in Q$ with $\mathsf{cap}(x, y) \subseteq R$.

- **Collision Event $\mathcal{C}$:** There exists $i, j \in [d]$ with $r_i = r_j$, $r_i \in \mathsf{cap}(Q_i)$ and $r_j \in \mathsf{cap}(Q_j)$.

If the abort event doesn't occur, then for any pair (x, y), the sign $\mathsf{sgn}_h(x, y)$ is determined by $\mathsf{cap}(x, y)$ for any $h \in \mathsf{supp}(\mathbf{Yes}) \cup \mathsf{supp}(\mathbf{No})$. The heart of the analysis lies in Theorem 3.10, which states that the bad events happen rarely. Theorem 3.10 is proved in Section 3.4.

Theorem 3.10. *If $|Q| \leq \delta d \log d$, then $\Pr[\mathcal{A} \cup \mathcal{C}] = O(\delta)$.*

When neither the abort nor the collision events happen, we say $\boldsymbol{S}$ is good for Q. Next, we show that conditioned on a good $\boldsymbol{S}$, the set Q cannot distinguish $f \sim \mathbf{Yes}$ from $g \sim \mathbf{No}$.

Lemma 3.11. *For any signed graph G over Q,*

$$\Pr_{f \sim \mathbf{Yes}}[G^Q_f = G | \boldsymbol{S} \text{ is good}] = \Pr_{g \sim \mathbf{No}}[G^Q_g = G | \boldsymbol{S} \text{ is good}].$$

Proof. We first describe the high level ideas in the proof. As stated above, when the abort event doesn't happen, the sign $\mathsf{sgn}_h(x, y)$ is determined by $\mathsf{cap}(x, y)$ for any $h \in \mathsf{supp}(\mathbf{Yes}) \cup \mathsf{supp}(\mathbf{No})$. Furthermore, a pair (x, y) has a possibility of distinguishing (that is, the pair is interesting) only if $x, y \in C_i$ and $r_i \in \mathsf{cap}(x, y)$. Focus on such interesting pairs. For such a pair, both $\mathsf{sgn}_{f_{\boldsymbol{T}}}(x, y)$

---

[4]There is nothing special about the constant 5. It just needs to be sufficiently large.

14

and $\mathsf{sgn}_{g_T}(x, y)$ are equally likely to be $+1$ or -1. Therefore, to distinguish, we would need two interesting pairs, $(x, y) \in C_i$ and $(x', y') \in C_j$ with $i \neq j$. Note that, when $g \sim \mathbf{No}$, the signs $\mathsf{sgn}_{g_T}(x, y)$ and $\mathsf{sgn}_{g_T}(x', y')$ are independently set, whereas when $f \sim \mathbf{Yes}$, the signs are either the same when $r_i = r_j$, or independently set. But if the collision event doesn't occur, we have $r_i \neq r_j$ for interesting pairs in different subcubes. Therefore, the probabilities are the same.

We now prove the lemma formally. Condition on a good $\boldsymbol{S}$. Note that the probability of the $\mathbf{Yes}$ distribution depends solely on $\{\alpha_b\}$ and that of the $\mathbf{No}$ distribution depends solely on $\{\beta_i\}$.

Consider any pair $(x, y) \in Q \times Q$ with $x \prec y$. We can classify it into three types: (i) x and y are in different subcubes, (ii) x and y are both in C_i, and $t_{\boldsymbol{S}}^i(x, y) \neq r_i$, (iii) x and y are both in C_i, and $t_{\boldsymbol{S}}^i(x, y) = r_i$. For convenience, we refer to the third type as *interesting pairs*. Let $h \in \mathsf{supp}(\mathbf{Yes}|\boldsymbol{S}) \cup \mathsf{supp}(\mathbf{No}|\boldsymbol{S})$. For the first and second types of pairs, by Claim 3.6 and Claim 3.8, we have $\mathsf{sgn}_h(x, y) = 1$. For interesting pairs, by Claim 3.8, $\mathsf{sgn}_h(x, y)$ must have the same label for all pairs in $Q_i \times Q_i$. Thus, any G whose labels disagree with the above can never be G_f^Q or G_g^Q.

Fix a signed graph G. For any pair $(x, y) \in Q \times Q$, where $x \prec y$, let $w(x, y)$ be the label in G. Furthermore, for all interesting pairs in the same Q_i, $w(x, y)$ has the same label, denoted w_i. Let I denote the set of subcubes with interesting pairs. At this point, all of our discussion depends purely on $\boldsymbol{S}$ and involves no randomness.

Now we focus on $g \sim (\mathbf{No}|\boldsymbol{S})$.

$$
\begin{aligned}
\Pr_{g \sim (\mathbf{No}|\boldsymbol{S})}[G_g^Q = G] &= \Pr\Big[\bigwedge_{i \in I} \bigwedge_{\substack{x, y \in Q_i \\ t_{\boldsymbol{S}}^i(x,y) = r_i}} (w(x, y) = \mathsf{sgn}_{g_T}(x, y))\Big] \\
&= \Pr\Big[\bigwedge_{i \in I} \bigwedge_{\substack{x, y \in Q_i \\ t_{\boldsymbol{S}}^i(x,y) = r_i}} (w(x, y) = \beta_i)\Big] \quad \text{(by Claim 3.8)} \\
&= \Pr\Big[\bigwedge_{i \in I} (w_i = \beta_i)\Big]
\end{aligned}
$$

Observe that each β_i is chosen uniformly and independently at random from $\{-1, +1\}$, and so this probability is exactly $2^{-|I|}$.

The analogous expressions for $f \sim (\mathbf{Yes}|\boldsymbol{S})$ yield:

$$
\Pr_{f \sim (\mathbf{Yes}|\boldsymbol{S})}[G_f^Q = G] = \Pr\Big[\bigwedge_{i \in I} (w_i = \alpha_{r_i})\Big]
$$

Note the difference here: if multiple r_i's are the same, the individual events are not independent over different subcubes. This is precisely what the abort and collision events capture. We formally argue below.

Consider an interesting pair $(x, y) \in Q_i \times Q_i$. Since the abort event $\mathcal{A}$ does not happen, $\mathsf{cap}(x, y) \not\subseteq R$. If $t_{\boldsymbol{S}}^i(x, y) = r_i \notin \mathsf{cap}(x, y)$, then there is a coordinate of $\overline{R}$ that is more significant that $t_{\boldsymbol{S}}^i(x, y)$. This contradicts the definition of the latter; so $r_i \in \mathsf{cap}(x, y) \subseteq \mathsf{cap}(Q_i)$. Equivalently, a subcube index $i \in I$ iff $r_i \in \mathsf{cap}(Q_i)$.

Since the collision event $\mathcal{C}$ does not happen, for any $j \in [m]$ where $r_j = r_i, r_j \notin \mathsf{cap}(Q_j)$. Alternately, for $i, i' \in I$, $r_i \neq r_{i'}$. Thus, $\Pr[\bigwedge_{i \in I}(w_i = \alpha_{r_i})] = \prod_{i \in I} \Pr[w_i = \alpha_{r_i}] = 2^{-|I|}$. $\quad\square$

Now, we are armed to prove Theorem 3.5.

Proof of Theorem 3.5. Given any subset of signed graphs, $\mathcal{G}$, it suffices to upper bound

$$\left| \Pr_{f \sim \mathbf{Yes}}[G_f^Q \in \mathcal{G}] - \Pr_{f \sim \mathbf{No}}[G_f^Q \in \mathcal{G}] \right| \leq \sum_{\text{good } \mathbf{S}} \left| \Pr[\mathbf{S}] \cdot \left(\Pr_{f \sim \mathbf{Yes}}[G_f^Q \in \mathcal{G} | \mathbf{S}] - \Pr_{f \sim \mathbf{No}}[G_f^Q \in \mathcal{G} | \mathbf{S}] \right) \right|$$

$$+ \sum_{\text{bad } \mathbf{S}} \left| \Pr[\mathbf{S}] \cdot \left(\Pr_{f \sim \mathbf{Yes}}[G_f^Q \in \mathcal{G} | \mathbf{S}] - \Pr_{f \sim \mathbf{No}}[G_f^Q \in \mathcal{G} | \mathbf{S}] \right) \right|.$$

The first term of the RHS is 0 by Lemma 3.11. The second term is at most the probability of bad events, which is $O(\delta)$ by Theorem 3.10. $\qquad\square$

3.4 Bounding the Probability of Bad Events: Proof of Theorem 3.10

We prove Theorem 3.10 by individually bounding $\Pr[\mathcal{A}]$ and $\Pr[\mathcal{C}]$.

Lemma 3.12. *If $|Q| \leq \delta d \log d$, then $\Pr[\mathcal{A}] \leq d^{-1/4}$.*

Proof. Fix any choice of k (in $\mathbf{S}$). For any pair of points $x, y \in Q$, we have $\Pr[\mathtt{cap}(x,y) \subseteq R] \leq (\frac{2^k}{d-5})^5$. Since $d - 5 \geq d/2$ for all $d \geq 10$ and $k \leq (\log_2 d)/2$, the probability is at most $32 d^{-5/2}$. For a large enough d, a union bound over all pairs in $Q \times Q$, which are at most $d^2 \log^2 d$ in number, completes the proof. $\qquad\square$

The collision event is more challenging to bound, and is actually the heart of the lower bound. We start by showing that, if each Q_i captures few coordinates, then the collision event has low probability. A critical point is the appearance of $d \log d$ in this bound.

Lemma 3.13. *If $\sum_i |\mathtt{cap}(Q_i)| \leq M$, then $\Pr[\mathcal{C}] = O\left(\frac{M}{d \log d} \right)$.*

Proof. For any $r \in [d]$, define $A_r := \{j : r \in \mathtt{cap}(Q_j)\}$ to be the set of indices of Q_j's that capture coordinate r. Let $a_r := |A_r|$. Define $n_\ell := |\{r : a_r \in (2^{\ell-1}, 2^\ell]\}|$. Observe that $\sum_{\ell \leq \log_2 d} n_\ell 2^\ell \leq 2 \sum_{r \in [d]} a_r \leq 2M$.

Fix k. For $r \in [d]$, we say the event $\mathcal{C}_r$ occurs if (a) $r \in R$, and (b) there exists $i, j \in [d]$ such that $r_i = r_j = r$, and $r_i \in \mathtt{cap}(Q_i)$ and $r_j \in \mathtt{cap}(Q_j)$. By the union bound, $\Pr[\mathcal{C}|k] \leq \sum_{r=1}^d \Pr[\mathcal{C}_r|k]$.

Let us now compute $\Pr[\mathcal{C}_r|k]$. Only sets Q_j's with $j \in A_r$ are of interest, since the others do not capture r. Event $\mathcal{C}_r$ occurs if at least two of these sets have $r_i = r_j = r$. Hence,

$$\Pr[\mathcal{C}_r|k] = \Pr[r \in R] \cdot \Pr[\exists i, j \in A_r : r_i = r_j = r \mid r \in R]$$

$$= \frac{2^k}{d} \cdot \sum_{c \geq 2} \binom{a_r}{c} \left(\frac{1}{2^k} \right)^c \left(1 - \frac{1}{2^k} \right)^{a_r - c}. \tag{1}$$

A fixed r is in R with probability $\binom{d-1}{2^k-1}/\binom{d}{2^k} = \frac{2^k}{d}$. Given that $|R| = 2^k$, the probability that $r_i = r$ is precisely 2^{-k}.

If $a_r \geq \frac{2^k}{4}$, then we simply upper bound (1) by $\frac{2^k}{d}$. For $a_r < \frac{2^k}{4}$, we upper bound (1) by

$$\frac{2^k}{d} \left(1 - \frac{1}{2^k} \right)^{a_r} \sum_{c \geq 2} \left(a_r \cdot \frac{1}{2^k} \cdot \left(1 - \frac{1}{2^k} \right)^{-1} \right)^c \leq \frac{2^k}{d} \sum_{c \geq 2} \left(\frac{a_r}{2^{k-1}} \right)^c \leq \frac{8 a_r^2}{2^k d}.$$

16

Summing over all r and grouping according to n_ℓ, we get

$$\Pr[\mathcal{C}|k] \leq \sum_{r=1}^{d} \Pr[\mathcal{C}_r|k] \leq \sum_{r:a_r \geq 2^{k-2}} \frac{2^k}{d} + \frac{8}{d} \sum_{r:a_r < 2^{k-2}} \frac{a_r^2}{2^k} \leq \frac{2^k}{d} \sum_{\ell > k-2} n_\ell + \frac{8}{d} \sum_{\ell=1}^{k-2} n_\ell 2^{2\ell-k}.$$

Averaging over all k, we get

$$\Pr[\mathcal{C}] = \frac{2}{\log_2 d} \sum_{k=1}^{(\log_2 d)/2} \Pr[\mathcal{C}|k] \quad \leq \quad \frac{16}{d\log_2 d} \sum_{k=1}^{(\log_2 d)/2} \left(\sum_{\ell=1}^{k-2} n_\ell 2^{2\ell-k} + \sum_{\ell > k-2} n_\ell 2^k \right)$$

$$= \frac{16}{d\log_2 d} \left(\sum_{\ell=1}^{(\log_2 d)/2} n_\ell \sum_{k \geq \ell+2} 2^{2\ell-k} + \sum_{\ell=1}^{\log_2 d} n_\ell \sum_{k < \ell+2} 2^k \right). \tag{2}$$

Now, $\sum_{k \geq \ell+2} 2^{2\ell-k} \leq 2^\ell$ and $\sum_{k < \ell+2} 2^k \leq 4 \cdot 2^\ell$. Substituting, $\Pr[\mathcal{C}] \leq \frac{80}{d\log_2 d} \sum_{\ell=1}^{\log_2 d} n_\ell 2^\ell \leq \frac{160M}{d\log_2 d}$, proving the lemma. $\qquad\square$

We are now left to bound $\sum_i |\mathtt{cap}(Q_i)|$. This is done by the following combinatorial lemma.

Lemma 3.14. *Let V be a set of vectors over an arbitrary alphabet and any number of dimensions. For any natural number c and $x, y \in V$, let $\mathtt{cap}_c(x,y)$ denote the (set of) first c coordinates at which x and y differ. Then $|\mathtt{cap}_c(V)| \leq c(|V| - 1)$.*

Proof. We construct c different edge-colored graphs $G_1, \ldots, G_c$ over the vertex set V. For every coordinate $i \in \mathtt{cap}_c(V)$, there must exist at least one pair of vectors x, y such that $i \in \mathtt{cap}_c(x,y)$. Thinking of each $\mathtt{cap}_c(x,y)$ as an ordered set, find a pair (x,y) where i appears "earliest" in $\mathtt{cap}_c(x,y)$. Let the position of i in this $\mathtt{cap}_c(x,y)$ be denoted t. We add edge (x,y) to G_t, and color it i. Note that the same edge (x,y) cannot be added to G_t with multiple colors, and hence all G_t's are simple graphs. Furthermore, observe that each color is present only once over all G_t's.

We claim that each G_t is acyclic. Suppose not. Let there be a cycle C and let (x,y) be the edge in C with the smallest color i. Clearly, $x_i \neq y_i$ since $i \in \mathtt{cap}_c(x,y)$. There must exist another edge (u,v) in C such that $u_i \neq v_i$. Furthermore, the color of (u,v) is $j > i$. Thus, j is the t^{th} entry in $\mathtt{cap}_c(u,v)$. Note that $i \in \mathtt{cap}_c(u,v)$ and must be the s^{th} entry for some $s < t$. But this means that the edge (u,v) colored i should be in G_s, contradicting the presence of $(x,y) \in G_t$. $\qquad\square$

We wrap up the bound now.

Lemma 3.15. *If $|Q| \leq \delta d \log d$, then $\Pr[\mathcal{C}] = O(\delta)$.*

Proof. Lemma 3.14 applied to each Q_i, yields $\sum_i |\mathtt{cap}(Q_i)| \leq 5|Q_i| = 5|Q|$. An application of Lemma 3.13 completes the proof. $\qquad\square$

4   Acknowledgments

References

[1] Nir Ailon and Bernard Chazelle. Information theory in property testing and monotonicity testing in higher dimension. *Inf. Comput.*, 204(11):1704–1717, 2006.

[2] Roksana Baleshzar, Meiram Murzabulatov, Ramesh Krishnan S. Pallavoor, and Sofya Raskhodnikova. Testing unateness of real-valued functions. *CoRR*, abs/1608.07652, 2016.

[3] Tugkan Batu, Ronitt Rubinfeld, and Patrick White. Fast approximate PCPs for multidimensional bin-packing problems. *Inf. Comput.*, 196(1):42–56, 2005.

[4] Aleksandrs Belovs and Eric Blais. Quantum algorithm for monotonicity testing on the hypercube. *Theory of Computing*, 11:403–412, 2015.

[5] Aleksandrs Belovs and Eric Blais. A polynomial lower bound for testing monotonicity. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 1021–1032, 2016.

[6] Piotr Berman, Sofya Raskhodnikova, and Grigory Yaroslavtsev. L_p-testing. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 164–173, 2014.

[7] Arnab Bhattacharyya, Elena Grigorescu, Kyomin Jung, Sofya Raskhodnikova, and David P. Woodruff. Transitive-closure spanners. *SIAM J. Comput.*, 41(6):1380–1425, 2012.

[8] Eric Blais, Joshua Brody, and Kevin Matulef. Property testing lower bounds via communication complexity. *Computational Complexity*, 21(2):311–358, 2012.

[9] Eric Blais, Sofya Raskhodnikova, and Grigory Yaroslavtsev. Lower bounds for testing properties of functions over hypergrid domains. In *Proceedings, IEEE Conference on Computational Complexity (CCC)*, pages 309–320, 2014.

[10] Jop Briët, Sourav Chakraborty, David García-Soriano, and Arie Matsliah. Monotonicity testing and shortest-path routing on the cube. *Combinatorica*, 32(1):35–53, 2012.

[11] Deeparnab Chakrabarty. Monotonicity testing. In *Encyclopedia of Algorithms*, pages 1352–1356. Springer, 2016.

[12] Deeparnab Chakrabarty, Kashyap Dixit, Madhav Jha, and C. Seshadhri. Property testing on product distributions: Optimal testers for bounded derivative properties. In *Proceedings, ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1809–1828, 2015.

[13] Deeparnab Chakrabarty and C. Seshadhri. Optimal bounds for monotonicity and Lipschitz testing over hypercubes and hypergrids. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 419–428, 2013.

[14] Deeparnab Chakrabarty and C. Seshadhri. An optimal lower bound for monotonicity testing over hypergrids. *Theory of Computing*, 10:453–464, 2014.

[15] Deeparnab Chakrabarty and C. Seshadhri. An $o(n)$ monotonicity tester for boolean functions over the hypercube. *SIAM J. Comput.*, 45(2):461–472, 2016.

[16] Deeparnab Chakrabarty and C. Seshadhri. A $\widetilde{O}(n)$ non-adaptive tester for unateness. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:133, 2016. Also appeared as arXiv report 1608.06980.

[17] Xi Chen, Anindya De, Rocco A. Servedio, and Li-Yang Tan. Boolean function monotonicity testing requires (almost) $O(n^{1/2})$ non-adaptive queries. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 519–528, 2015.

[18] Xi Chen, Rocco A. Servedio, and Li-Yang Tan. New algorithms and lower bounds for monotonicity testing. In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 286–295, 2014.

[19] Xi Chen, Erik Waingarten, and Jinyu Xie. Beyond talagrand functions: New lower bounds for testing monotonicity and unateness. *CoRR*, abs/1702.06997, 2017. To appear in STOC 2017.

[20] Kashyap Dixit, Sofya Raskhodnikova, Abhradeep Thakurta, and Nithin M. Varma. Erasure-resilient property testing. In *Proceedings, International Colloquium on Automata, Languages and Processing (ICALP)*, pages 91:1–91:15, 2016.

[21] Yevgeny Dodis, Oded Goldreich, Eric Lehman, Sofya Raskhodnikova, Dana Ron, and Alex Samorodnitsky. Improved testing algorithms for monotonicity. *Proceedings, International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*, pages 97–108, 1999.

[22] Funda Ergün, Sampath Kannan, Ravi Kumar, Ronitt Rubinfeld, and Mahesh Viswanathan. Spot-checkers. *J. Comput. System Sci.*, 60(3):717–751, 2000.

[23] Eldar Fischer. On the strength of comparisons in property testing. *Inf. Comput.*, 189(1):107–116, 2004.

[24] Eldar Fischer, Eric Lehman, Ilan Newman, Sofya Raskhodnikova, Ronitt Rubinfeld, and Alex Samorodnitsky. Monotonicity testing over general poset domains. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, pages 474–483, 2002.

[25] Oded Goldreich. *Introduction to Property Testing (working draft)*. 2015. URL: `www.wisdom.weizmann.ac.il/~oded/PDF/pt-v1.pdf`.

[26] Oded Goldreich, Shafi Goldwasser, Eric Lehman, Dana Ron, and Alex Samorodnitsky. Testing monotonicity. *Combinatorica*, 20:301–337, 2000.

[27] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *J. ACM*, 45(4):653–750, 1998.

[28] Shirley Halevy and Eyal Kushilevitz. Distribution-free property-testing. *SIAM J. Comput.*, 37(4):1107–1138, 2007.

[29] Shirley Halevy and Eyal Kushilevitz. Testing monotonicity over graph products. *Random Struct. Algorithms*, 33(1):44–67, 2008.

[30] Madhav Jha and Sofya Raskhodnikova. Testing and reconstruction of Lipschitz functions with applications to data privacy. *SIAM J. Comput.*, 42(2):700–731, 2013.

[31] Subhash Khot, Dor Minzer, and Muli Safra. On monotonicity testing and boolean isoperimetric type theorems. In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 52–58, 2015.

[32] Subhash Khot and Igor Shinkar. An $\widetilde{O}(n)$ queries adaptive tester for unateness. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2016*, pages 37:1–37:7, 2016.

[33] Eric Lehman and Dana Ron. On disjoint chains of subsets. *J. Combin. Theory Ser. A*, 94(2):399–404, 2001.

[34] Ramesh Krishnan S. Pallavoor, Sofya Raskhodnikova, and Nithin Varma. Parameterized property testing of functions. In *Proceedings, Innovations in Theoretical Computer Science (ITCS)*, 2017.

[35] Sofya Raskhodnikova. Testing if an array is sorted. In *Encyclopedia of Algorithms*, pages 2219–2222. Springer, 2016.

[36] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM J. Comput.*, 25(2):252–271, 1996.

[37] Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity (extended abstract). In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 222–227, 1977.

A Missing Details from the Main Body

A.1 The Lower Bound for Adaptive Testers over Hypergrids

We show that every unateness tester for functions $f : [n]^d \mapsto \mathbb{R}$ requires $\Omega\left(\frac{d \log n}{\varepsilon} - \frac{\log 1/\varepsilon}{\varepsilon}\right)$ queries for $\varepsilon \in (0, 1/4)$ and prove Theorem 1.3.

Proof of Theorem 1.3. By Yao's minimax principle and the reduction to testing with comparison-based testers from [14] (stated for completeness in Theorem 3.1), it is sufficient to give a hard input distribution on which every deterministic comparison-based tester fails with probability more than 2/3. We use the hard distribution constructed by Chakrabarty and Seshadhri [14] to prove the same lower bound for testing monotonicity. Their distribution is a mixture of two distributions, **Yes** and **No**, on positive and negative instances, respectively. Positive instances for their problem are functions that are monotone and, therefore, unate; negative instances are functions that are ε-far from monotone. We show that their **No** distribution is supported on functions that are ε-far from unate, i.e., negative instances for our problem. Then the required lower bound for unateness follows from the fact that every deterministic comparison-based tester needs the stated number of queries to distinguish **Yes** and **No** distributions with high enough probability.

We start by describing the **Yes** and **No** distribution used in [14]. We will define them as distributions on functions over the hypercube domain. Next, we explain how to convert functions over hypercubes to functions over hypergrids.

Without loss of generality, assume n is a power of 2 and let $\ell := \log_2 n$. For any $z \in [n]$, let $bin(z)$ denote the binary representation of $z - 1$ as an ℓ-bit vector $(z_1, \ldots, z_\ell)$, where z_1 is the least significant bit.

We now describe the mapping used to convert functions on hypergrids to functions on hypercubes. Let $\phi : [n]^d \to \{0,1\}^{d\ell}$ be the mapping that takes $y \in [n]^d$ to the concatenation of $bin(y_1), \ldots, bin(y_d)$. Any function $f : \{0,1\}^{d\ell} \mapsto \mathbb{R}$ can be easily converted into a function $\widetilde{f} : [n]^d \mapsto \mathbb{R}$, where $\widetilde{f}(y) := f(\phi(y))$.

Let $m := d\ell$. For $x \in \{0,1\}^m$, let $\mathtt{val}(x) = \sum_{i=1}^m x_i 2^{i-1}$ denote the value of the binary number represented by vector x. For simplicity, assume $1/\varepsilon$ is a power of 2. Partition the set of points $x \in \{0,1\}^m$ according to the most significant $\log(1/2\varepsilon)$ dimensions. That is, for $k \in \{1, 2, \ldots, 1/2\varepsilon\}$, let

$$S_k := \{x : \mathtt{val}(x) \in [(k-1) \cdot \varepsilon 2^{m+1}, k \cdot \varepsilon 2^{m+1} - 1]\}.$$

The hypercube is partitioned into $1/2\varepsilon$ sets S_k of equal size, and each S_k forms a subcube of dimension $m' = m - \log(1/\varepsilon) + 1$.

We now describe the **Yes** and **No** distributions for functions on hypercubes. The **Yes** distribution consists of a single function $f(x) = 2\mathtt{val}(x)$. The **No** distribution is uniform over $m'/2\varepsilon$ functions $g_{j,k}$, where $j \in [m']$ and $k \in [1/2\varepsilon]$, defined as follows:

$$g_{j,k}(x) = \begin{cases} 2\mathtt{val}(x) - 2^j - 1 & \text{if } x_j = 1 \text{ and } x \in S_k; \\ 2\mathtt{val}(x), & \text{otherwise.} \end{cases}$$

To get the **Yes** and **No** distributions for the hypergrid, we convert f to $\widetilde{f}$ and each function $g_{j,k}$ to $\widetilde{g_{j,k}}$, using the transformation defined before.

Chakrabarty and Seshadhri [14] proved that f is monotone and each function $\widetilde{g_{j,k}}$ is ε-far from monotone. It remains to show that functions $\widetilde{g_{j,k}}$ are also ε-far from unate.

Claim A.1. *Each function $\widetilde{g_{j,k}}$ is ε-far from unate.*

Proof. To prove that $\widetilde{g_{j,k}}$ is ε-far from unate, it suffices to show that there exists a dimension i, such that there are at least $\varepsilon 2^{d\ell}$ increasing i-pairs and at least $\varepsilon 2^{d\ell}$ decreasing i-pairs w.r.t. $\widetilde{g_{j,k}}$ and that all of these i-pairs are disjoint. Let $u, v \in [n]^d$ be two points such that $\phi(u)$ and $\phi(v)$ differ only in the j^{th} bit. Clearly, u and v form an i-pair, where $i = \lceil j/\ell \rceil$. Now, if $\phi(u), \phi(v) \in S_k$ and $u \prec v$, then $\widetilde{g_{j,k}}(v) = \widetilde{g_{j,k}}(u) - 1$. So, the i-pair (u, v) is decreasing. The total number of such i-pairs is $2^{d\ell - \log(1/2\varepsilon) - 1} = \varepsilon 2^{d\ell}$. If $\phi(u), \phi(v) \in S_{k'}$ where $k' \neq k$, then the i-pair (u, v) is increasing. Clearly, there are at least $\varepsilon 2^{d\ell}$ such i-pairs. All the i-pairs we mentioned are disjoint. Hence, $\widetilde{g_{j,k}}$ is ε-far from unate. $\square$

This completes the proof of Theorem 1.3. $\square$

A.2 The Lower Bound for Nonadaptive Testers over Hypergrids

The lower bound for nonadaptive testers over hypergrids follows from a combination of the lower bound for nonadaptive testers over hypercube and the lower bound for adaptive testers over hypergrids.

Theorem A.2. *Any nonadaptive unateness tester (even with two-sided error) for real-values functions $f : [n]^d \mapsto \mathbb{R}$ must make $\Omega(d(\log n + \log d))$ queries.*

Proof. Fix $\varepsilon = 1/8$. The proof consists of two parts. The lower bound for adaptive testers is also a lower bound for nonadaptive tester, and so, the bound of $\Omega(d \log n)$ holds. Next, we extend the $\Omega(d \log d)$ lower bound for hypercubes. Assume n to be a power of 2. Define function $\psi : [n] \mapsto \{0, 1\}$ as $\psi(a) := \mathbb{1}[a > n/2]$ for $a \in [n]$. For $x = (x_1, x_2, \ldots, x_d) \in [n]^d$, define the mapping $\Psi : [n]^d \mapsto \{0, 1\}^d$ as $\Psi(x) := (\psi(x_1), \psi(x_2), \ldots, \psi(x_d))$. Any function $f : \{0, 1\}^d \mapsto \mathbb{R}$ can be extended to $\tilde{f} : [n]^d \mapsto \mathbb{R}$ using the mapping $\tilde{f}(x) = f(\Psi(x))$ for all $x \in [n]^d$. The proof of Theorem 3.2 goes through for hypergrids as well, and so we have an $\Omega(d \log d)$ lower bound. Combining the two lower bounds, we get a bound of $\Omega(d \cdot \max\{\log n, \log d\})$, which is asymptotically equal to $\Omega(d(\log n + \log d))$. $\qquad \square$