

Random Restrictions of High-Dimensional Distributions and Uniformity Testing with Subcube Conditioning

Clément L. Canonne* Xi Chen† Gautam Kamath‡ Amit Levi§
Erik Waingarten¶

November 17, 2019

Abstract

We give a nearly-optimal algorithm for testing uniformity of distributions supported on $\{-1, 1\}^n$, which makes $\tilde{O}(\sqrt{n}/\varepsilon^2)$ queries to a subcube conditional sampling oracle (Bhattacharyya and Chakraborty (2018)). The key technical component is a natural notion of random restriction for distributions on $\{-1, 1\}^n$, and a quantitative analysis of how such a restriction affects the mean vector of the distribution. Along the way, we consider the problem of *mean testing* with independent samples and provide a nearly-optimal algorithm.

*IBM Research. ccanonne@cs.columbia.edu. Part of this work was conducted while a Motwani Postdoctoral Fellow at Stanford University.

†Columbia University. xichen@cs.columbia.edu. Supported by NSF IIS-1838154 and NSF CCF-1703925.

‡Cheriton School of Computer Science, University of Waterloo. g@csail.mit.edu. Supported by a University of Waterloo startup grant. Part of this work was conducted while the author was supported by a Microsoft Research Fellowship, as part of the Simons-Berkeley Research Fellowship program, and while visiting Microsoft Research, Redmond.

§Cheriton School of Computer Science, University of Waterloo. amit.levi@uwaterloo.ca. Research supported by the David R. Cheriton Graduate Scholarship.

¶Columbia University. eaw@cs.columbia.edu. Supported by the NSF Graduate Research Fellowship (Grant No. DGE-16-44869)

Contents

1	Introduction	1
1.1	Technical Overview	4
1.1.1	A Robust Pisier’s Inequality	4
1.1.2	Warmup: A Linear Query Algorithm	5
1.1.3	A Weaker Version of Theorem 1.5	6
1.2	Related Work	7
1.3	Notation and Preliminaries	8
2	The Algorithm	8
3	Proof of Theorem 1.5	12
3.1	Robust Pisier Inequality	13
3.2	Plan of the Proof of Lemma 3.1	16
3.3	From Total Variation to Directed Graphs	16
3.4	Bucketing	19
3.5	From Directed Graphs to Mean Vectors	20
3.6	Case 1	21
3.7	Case 2	24
4	Mean Testing	27
4.1	Application to Gaussian Mean Testing	31
A	Proof of Lemma 4.1	37

1 Introduction

The focus of this paper is high-dimensional distribution testing. The algorithmic problem is the following: we are granted oracle access (the type of which we will specify shortly) to a probability distribution p on $\Sigma = \{-1, 1\}^n$, and must distinguish with probability at least $2/3$ between the case where p is the uniform distribution, and that where p is ε -far from uniform in total variation distance. The classical works of distribution testing [GGR96, GR00, BFR⁺00] study the above question in the standard statistical setting, where the oracle provides independent samples from p . In this case, the hallmark results are an algorithm and a matching lower bound, showing that $\Theta(\sqrt{|\Sigma|}/\varepsilon^2)$ independent samples are necessary and sufficient for testing uniformity [Pan08, VV14]. When studying distributions supported on high-dimensional domains, unfortunately, this implies that the complexity of sample-optimal algorithms scales exponentially with the dimension, effectively making the problem intractable. To circumvent this issue, recent work has proceeded by either restricting the class of input distributions (e.g., restricting p to be a product distribution; see Section 1.2), or by allowing stronger oracle access. We take the latter approach, and consider an oracle access which is particularly well-suited to the high-dimensional structure: the *subcube conditional query model*.

Subcube conditional query access, first suggested in [CRS15] and studied in [BC18], allows algorithms to specify a subcube of the high-dimensional domain and request a sample from the distribution *conditioned* on the sample lying in the subcube specified — equivalently, to request samples after fixing some of their variables. The operation is akin to the notion of *restriction* in the analysis of Boolean functions. Specifically, we identify the distribution by its probability mass function $p: \{-1, 1\}^n \rightarrow \mathbb{R}_+$. An algorithm may then specify a subcube by a string $\rho \in \{-1, 1, *\}^n$, where $*$'s denote free variables and non- $*$'s denote the values of restricted variables. Calling the oracle on such a ρ results in a sample from the distribution $p|_\rho$ (now supported on $\{-1, 1\}^{\text{stars}(\rho)}$) given by restricting the function $p|_\rho: \{-1, 1\}^{\text{stars}(\rho)} \rightarrow \mathbb{R}_+$ and re-normalizing it, so that it represents a distribution $p|_\rho$.¹

Our main results are two-fold: (i) We define a natural notion of random restriction for high-dimensional distributions and analyze the behavior of the mean vector of a distribution under such random restrictions; (ii) Leveraging this analysis, we obtain a nearly-optimal algorithm for testing uniformity over $\{-1, 1\}^n$ with subcube conditioning. As we will show, subcube conditioning allows us to go from $\Theta(2^{n/2}/\varepsilon^2)$ sample complexity to $\tilde{\Theta}(\sqrt{n}/\varepsilon^2)$: that is, with subcube conditional queries, testing uniformity over arbitrary distributions is no harder than that over the much more restricted class of *product* ones.

Theorem 1.1 (Main Result: Uniformity Testing). *There exists an algorithm which, given subcube conditional query access to a distribution p supported on $\{-1, 1\}^n$ and a distance parameter $\varepsilon \in (0, 1)$, makes $\tilde{O}(\sqrt{n}/\varepsilon^2)$ queries and can distinguish with probability at least $2/3$ between the case when p is uniform, and when p is ε -far from uniform in total variation distance.*

Along the way, we study the *mean testing* problem, a natural variant of uniformity testing for high-dimensional distributions, and obtain a nearly-optimal bound for this question in the standard sampling model.

Product Distributions and Mean Distance. The simplest class of distributions on $\{-1, 1\}^n$ is arguably the class of *product distributions*, where all coordinates are independent. This setting

¹When conditioning on a subcube with zero support, one may consider models where the oracle returns a uniform sample [CFGM16] or outputs “error” [CRS15]. We note that our algorithm will never run into this scenario.

was studied in [CDKS17], and is particularly nice to analyze due to the relation between the total variation distance between distributions and the ℓ_2 distance between their mean vectors. Specifically, let p be a product distribution supported on $\{-1, 1\}^n$, and $\mu(p) \in [-1, 1]^n$ be its *mean vector*,

$$\mu(p) = \mathbf{E}_{\mathbf{x} \sim p}[\mathbf{x}] \in [-1, 1]^n.$$

It is not hard to show that if p is ε -far from uniform in total variation distance, then $\|\mu(p)\|_2 \gtrsim \varepsilon$. Hence, for product distributions, large total variation distance to uniformity implies large mean vector in ℓ_2 norm. Given this fact, Canonne, Diakonikolas, Kane, and Stewart [CDKS17] design an algorithm based on estimating the norm of the mean vector, and show that $O(\sqrt{n}/\varepsilon^2)$ many samples from product distributions suffice to test uniformity.²

However, the relationship observed between distance to uniformity in total variation and *mean distance* (i.e., the ℓ_2 norm of the mean vector) for product distributions is not true in general. A simple example is the uniform distribution supported on just two vectors $\{x, -x\}$, which is very far from uniform yet has mean vector 0. Towards relating these two notions for general distributions, we define our notion of random restriction.

Random Restrictions. For any $\sigma \in [0, 1]$, we write \mathcal{S}_σ for the distribution supported on subsets of $[n]$ given by letting $\mathbf{S} \sim \mathcal{S}_\sigma$ include each index $i \in \mathbf{S}$ independently with probability σ . Given any distribution p supported on $\{-1, 1\}^n$, let $\mathcal{D}_\sigma(p)$ be the distribution, supported on $\{-1, 1, *\}^n$, of *random restrictions* of p . In order to sample a random restriction $\rho \sim \mathcal{D}_\sigma(p)$, we sample a set $\mathbf{S} \sim \mathcal{S}_\sigma$ and a sample $\mathbf{x} \sim p$; then, we let ρ_i be set according to:

$$\rho_i = \begin{cases} * & \text{if } i \in \mathbf{S} \\ \mathbf{x}_i & \text{if } i \notin \mathbf{S} \end{cases}. \quad (1)$$

For any $\rho \in \{-1, 1, *\}^n$, we denote by $p|_\rho$ the distribution on $\{-1, 1\}^{\text{stars}(\rho)}$ given by $\mathbf{x}_{\text{stars}(\rho)}$ where \mathbf{x} is drawn from p *conditioned* on every $i \notin \text{stars}(\rho)$ being set to ρ_i . This defines the restriction of a distribution p . Another operation on distributions we will consider is that of *projection*; for any set $S \subset [n]$, we write $\bar{S} = [n] \setminus S$, and the distribution $p_{\bar{S}}$ supported on $\{-1, 1\}^{\bar{S}}$ is given by letting $\mathbf{y} \sim p_{\bar{S}}$ be $\mathbf{y} = \mathbf{x}_{\bar{S}}$ for $\mathbf{x} \sim p$.

Let $d_{\text{TV}}(p, \mathcal{U})$ denote the total variation distance between p and the uniform distribution of the same dimension. At a high level, our main technical result shows that *the mean distance of random restrictions is implied by total variation distance of random projections*.

Theorem 1.2 (Informal version; see Theorem 1.5). *Let p be any distribution over $\{-1, 1\}^n$. Then,*

$$\mathbf{E}_{\rho \sim \mathcal{D}_\sigma(p)} \left[\|\mu(p|_\rho)\|_2 \right] \geq \sigma \cdot \mathbf{E}_{\mathbf{S} \sim \mathcal{S}_\sigma} [d_{\text{TV}}(p_{\bar{\mathbf{S}}}, \mathcal{U})]. \quad (2)$$

Although the above differs from Theorem 1.5 in certain respects (the inequality in Theorem 1.5 incurs additional poly-logarithmic factors as well as a small additive error), (2) captures the key relationship between the total variation distance and the mean distance that we leverage, and will provide intuition for the introduction.

²In fact, they consider the more general problem of identity testing of product distributions, where they obtain analogous results.

Mean Testing. The above discussion naturally leads to the following problem, which we refer to as *mean testing*. Given sample access to a distribution p supported on $\{-1, 1\}^n$, we seek to distinguish with probability at least $2/3$ between the case where p is uniform and that where p has a large mean vector, i.e., $\|\mu(p)\|_2 \geq \varepsilon\sqrt{n}$. When p is assumed to be a product distribution, [CDKS17, DDK18] showed that for $\varepsilon \leq 1/\sqrt{n}$ the sample complexity of the problem is $\Theta(1/(\varepsilon^2\sqrt{n}))$.³ The idea is that the empirical mean of a product distribution will have norm concentrated around its true value, and thus it suffices to look at this empirical estimate. In our case, however, additional care is needed, as there can be arbitrary correlations between coordinates of a sample and this concentration does not hold in general. Nevertheless, we present an algorithm for mean testing which is optimal up to a triply-logarithmic small loss in the sample complexity.

Theorem 1.3 (Mean Testing). *There exists an algorithm which given*

$$O\left(\max\left\{\frac{1}{\varepsilon^2\sqrt{n}}, \frac{1}{\varepsilon}\right\} \cdot \log \log \log n\right)$$

i.i.d. samples from an arbitrary distribution p on $\{-1, 1\}^n$ and a parameter $\varepsilon \in (0, 1]$ can distinguish with probability at least $2/3$ between (i) p is the uniform distribution, and (ii) $\|\mu(p)\|_2 \geq \varepsilon\sqrt{n}$.

Moreover, as detailed in Section 4.1, the above immediately implies a similar sample complexity for *Gaussian* mean testing, where one is given i.i.d. samples from a multivariate normal distribution $p = \mathcal{G}(\mu, \Sigma)$ and must distinguish between $p = \mathcal{G}(0, I)$ and $\|\mu\|_2 \geq \varepsilon\sqrt{n}$.

Subcube Conditional Sampling. In view of the above, we aim to use random restrictions and Theorem 1.3 to test uniformity with subcube conditional queries. The final technical ingredient is the following inequality, very similar to the “chain rule” of Bhattacharyya and Chakraborty [BC18] suited for uniformity testing on $\{-1, 1\}^n$.

Lemma 1.4. *Let p be a distribution supported on $\{-1, 1\}^n$. Then, for any $\sigma \in [0, 1]$,*

$$d_{\text{TV}}(p, \mathcal{U}) \leq \mathbf{E}_{\mathbf{s} \sim \mathcal{S}_\sigma} [d_{\text{TV}}(p_{\mathbf{s}}, \mathcal{U})] + \mathbf{E}_{\rho \sim \mathcal{D}_\sigma(p)} [d_{\text{TV}}(p|_\rho, \mathcal{U})].$$

This lemma naturally leads to a recursive approach. Given a distribution p which is ε -far from uniform, either the total variation of random projections is large, or the total variation of random restrictions is large. In the former case, we apply (2), which allows us to reduce the problem to that of mean testing, and invoke Theorem 1.3. In the latter case, we take a random restriction and recurse (but on far fewer variables).

Quantitatively, Theorem 1.1 is a quartic improvement over the result of [BC18], who proved an upper bound of $\tilde{O}(n^2/\varepsilon^2)$ for uniformity testing with subcube conditional queries. We thus show that the complexity of the problem is not only polynomial in the dimension, but strongly sublinear; moreover, Theorem 1.1 is tight up to poly-logarithmic factors. Indeed, as observed in [BC18], the sample complexity lower bound of $\Omega(\sqrt{n}/\varepsilon^2)$ of [CDKS17, DDK18] for testing uniformity of product distributions carries over to subcube conditional sampling. The reason is that, for product distributions, the coordinates are already independent, and therefore conditioning on subcubes does not grant any additional power.

³This is implicit in [CDKS17, Theorem 4.1]. As stated, their result is suited for distinguishing between the mean vector having norm 0 or at least ε' , where $\varepsilon' \in (0, 1]$: we re-parameterize with $\varepsilon' = \varepsilon\sqrt{n}$.

1.1 Technical Overview

We now formally state and explain the intuition behind our main theorem, relating the distance to uniformity of random projections to the mean distance after random restrictions.

Theorem 1.5. *Let p be any distribution over $\{-1, 1\}^n$ and $\sigma \in [0, 1]$. Then,*

$$\mathbf{E}_{\rho \sim \mathcal{D}_\sigma(p)} \left[\|\mu(p|_\rho)\|_2 \right] \geq \frac{\sigma}{\text{poly}(\log n)} \cdot \tilde{\Omega} \left(\mathbf{E}_{\mathbf{s} \sim \mathcal{S}_\sigma} [d_{\text{TV}}(p_{\overline{\mathbf{s}}}, \mathcal{U})] - 2e^{-\min(\sigma, 1-\sigma)n/10} \right).$$

We encourage the reader to think of applying Theorem 1.5 to a distribution p which is ε -far from uniform, and to think of the case when the parameter σ is a small constant (or inverse of a poly-logarithmic factor), and $\mathbf{E}_{\mathbf{s} \sim \mathcal{S}_\sigma} [d_{\text{TV}}(p_{\overline{\mathbf{s}}}, \mathcal{U})] = \Theta(\varepsilon)$. Specifically, consider a parameter setting where $e^{-\min(\sigma, 1-\sigma)n/10} = o(\varepsilon)$ so that the right-hand side of the expression in Theorem 1.5 becomes $\varepsilon\sigma/\text{poly}(\log n, \log(1/\varepsilon))$.

The proof of Theorem 1.5 proceeds by proving a lemma (Lemma 3.1) which captures the behavior of random restrictions with t stars versus those with $t+1$ stars. We prove a robust version of Pisier's inequality [Pis86], an inequality which was first studied in the geometry of Banach spaces. When the projections $p_{\overline{\mathbf{s}}}$ are far from uniform in total variation, the robust version of Pisier's inequality will help us lower bound certain quantities of a collection of directed graphs defined using $p_{\overline{\mathbf{s}}}$ on the subcube $\{-1, 1\}^{\overline{\mathbf{s}}}$. The coordinate values of the mean vector of $p|_\rho$ will depend on whether random vertices and directions induced from ρ are directed edges of one of the graphs in the collection. Hence, the lower bound on expected norm of the mean vector will follow from studying structures of directed graphs in the collection.

We first introduce the robust version of Pisier's inequality. As a warm-up and in order to show the usefulness of this inequality, we give an algorithm making $\tilde{O}(n/\varepsilon^2)$ queries which we refer to as an *edge tester*. The reason for this name is that the algorithm samples a random restriction ρ with exactly one star, so that the distribution $p|_\rho$ is supported on an edge of the hypercube. Then, the algorithm tests whether $p|_\rho$ is uniform. Lastly, we provide a proof sketch for a weaker version of Theorem 1.5 which lower bounds the expectation of $\|\mu(p|_\rho)\|_2^2$, rather than $\|\mu(p|_\rho)\|_2$. The weaker inequality is insufficient for our purposes, but the proof is conceptually much simpler (since $\|\cdot\|_2^2$ is additive over coordinates and thus, we may use linearity of expectation). To prove Theorem 1.5, attempting to lower bound $\|\mu(p|_\rho)\|_2$ reveals further challenges which necessitate additional care.

1.1.1 A Robust Pisier's Inequality

We let \mathcal{H} denote the set of undirected edges of the hypercube $\{-1, 1\}^n$. For a function $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ and $i \in [n]$, we write $L_i f(x) \stackrel{\text{def}}{=} (f(x) - f(x^{(i)}))/2$. The following inequality is known as Pisier's inequality [Pis86]. We state it below in a way most closely to how it will be applied in this paper; the inequality holds in much larger generality for functions $f: \{-1, 1\}^n \rightarrow X$ over general Banach spaces X (see, in particular, [NS02]).

Theorem 1.6 (Pisier's inequality). *Let $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ be a function with $\mathbf{E}_x[f(x)] = 0$. Then,*

$$\mathbf{E}_{x \sim \{-1, 1\}^n} [|f(x)|] \lesssim \log n \cdot \mathbf{E}_{x, y \sim \{-1, 1\}^n} \left[\left| \sum_{i=1}^n y_i x_i L_i f(x) \right| \right].$$

In this paper, we will need a *robust* version of Pisier's inequality in order to derive Theorem 1.5. The notion of robustness is equivalent to the notion considered in [KMS18], who proved robust (and

directed) versions of Talagrand’s inequality for Boolean functions, and part of our proof utilizes the robustness of the inequality in a similar way. Specifically, we consider an arbitrary orientation of the edges of the hypercube and sum the values of $L_i f(\mathbf{x})$ only when the edge $\{\mathbf{x}, \mathbf{x}^{(i)}\}$ is oriented from \mathbf{x} to $\mathbf{x}^{(i)}$.

Theorem 1.7 (Robust version of Pisier’s inequality). *Let $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ be a function satisfying $\mathbf{E}_{\mathbf{x}}[|f(\mathbf{x})|] = 0$. Let $G = (\{-1, 1\}^n, E)$ be any orientation of the hypercube. Then,*

$$\mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} [|f(\mathbf{x})|] \lesssim \log n \cdot \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^n} \left[\left| \sum_{\substack{i \in [n] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in E}} \mathbf{y}_i \mathbf{x}_i L_i f(\mathbf{x}) \right| \right].$$

The proof follows the template of [NS02, Theorem 2], and checks that the necessary changes, even after considering directed edges, still give the desired inequality.

Remark 1.8. We note that Talagrand [Tal93] prove that the $\log n$ factor in Pisier’s inequality (Theorem 1.6) is unnecessary for real-valued functions, but that it is for general Banach spaces. While one may follow Talagrand’s proof to remove the $\log n$ factor in Theorem 1.7, it is not immediately clear whether this approach can handle arbitrary orientations.

1.1.2 Warmup: A Linear Query Algorithm

To see why Theorem 1.7 is helpful, we give a simple (albeit suboptimal) $\tilde{O}(n/\varepsilon^2)$ -query algorithm for testing uniformity. Suppose p is a distribution supported on $\{-1, 1\}^n$ which is ε -far from uniform in total variation distance. We apply Theorem 1.7 to the function $f(x) \stackrel{\text{def}}{=} 2^n \cdot p(x) - 1$, where the directed graph $G = (\{-1, 1\}^n, E)$ is given by letting

$$E = \left\{ (x, x^{(i)}) : \{x, x^{(i)}\} \text{ is an edge of } \{-1, 1\}^n, \text{ and } p(x) \geq p(x^{(i)}) \right\}.$$

Theorem 1.7 applied on f implies that⁴

$$\mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} [|f(\mathbf{x})|] \lesssim \log n \cdot \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^n} \left[\left| \sum_{i=1}^n \mathbf{y}_i \mathbf{x}_i (f(\mathbf{x}) - f(\mathbf{x}^{(i)}))^+ \right| \right]. \quad (3)$$

Given (3), we start by applying the fact that the left-hand side of the inequality is at least 2ε . From there, the following three (in)equalities are obtained via Khintchine’s inequality, importance sampling, and Jensen’s inequality, respectively. (We also use the convention that $0/0 = 0$.)

$$\begin{aligned} \frac{\varepsilon}{\log n} &\lesssim \mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} \left[\sqrt{\sum_{i=1}^n \left((f(\mathbf{x}) - f(\mathbf{x}^{(i)}))^+ \right)^2} \right] = \mathbf{E}_{\mathbf{x} \sim p} \left[\sqrt{\sum_{i=1}^n \left(\frac{(f(\mathbf{x}) - f(\mathbf{x}^{(i)}))^+}{1 + f(\mathbf{x})} \right)^2} \right] \\ &\leq \left(\mathbf{E}_{\mathbf{x} \sim p} \left[\sum_{i=1}^n \left(\frac{(p(\mathbf{x}) - p(\mathbf{x}^{(i)}))^+}{p(\mathbf{x})} \right)^2 \right] \right)^{1/2}. \quad (4) \end{aligned}$$

⁴Our proof of Theorem 1.5 will rely on this strengthening of Pisier’s inequality for multiple reasons, as our approach crucially uses the ability to pick any orientation of the edges. Even for this simple application one can observe that, without the strengthening, the right-hand side of (3) would be replaced by the same expression but without the $^+$. The discussion below (4) and the derivation of (5) explains why having the $^+$ in the expression is crucial.

Notice that for every x and i , either $(p(x) - p(x^{(i)}))^+ = 0$ or $p(x) > p(x^{(i)}) \geq 0$. First, this makes sure that $a/0$ with $a > 0$ never occurs in (4). Additionally,

$$0 \leq \frac{1}{2} \cdot \frac{(p(x) - p(x^{(i)}))^+}{p(x)} \leq \frac{|p(x) - p(x^{(i)})|}{p(x) + p(x^{(i)})}, \quad (5)$$

and the quantity on the right-hand side is exactly the bias on the distribution of the i th bit of a draw of p conditioning on all $i' \neq i$ bits set according to x . In particular, a standard averaging/bucketing argument shows that there exists $\beta \geq \varepsilon^2/(n \log^2 n)$ such that

$$\Pr_{\substack{\mathbf{x} \sim p \\ \mathbf{i} \sim [n]}} \left[\left(\frac{|p(\mathbf{x}) - p(\mathbf{x}^{(\mathbf{i})})|}{p(\mathbf{x}) + p(\mathbf{x}^{(\mathbf{i})})} \right)^2 \gtrsim \frac{\varepsilon^2}{\beta \cdot n \log^2 n \log(n/\varepsilon)} \right] \geq \beta. \quad (6)$$

The above lower bound suggests the following ‘‘edge tester:’’ for all $h \in \{0, \dots, O(\log(n/\varepsilon))\}$ where $2^{-h} \gtrsim \varepsilon^2/(n \log^2 n)$, independently sample $O(2^h \log(n/\varepsilon))$ pairs (\mathbf{x}, \mathbf{i}) (\mathbf{x} from p , and $\mathbf{i} \sim [n]$). For each pair (\mathbf{x}, \mathbf{i}) , we consider the distribution supported on $\{-1, 1\}$ given by sampling $\mathbf{y} \sim p$ conditioned on every $i' \neq i$ having $\mathbf{y}_{i'} = \mathbf{x}_{i'}$, and use $\tilde{O}(2^{-h} n/\varepsilon^2)$ queries to estimate the bias of the conditional distribution up to error $(\varepsilon \cdot (2^h/n)^{1/2})/\text{polylog}(n/\varepsilon)$ with high probability. If p was uniform, every conditional distribution considered will be uniform; however, if p is ε -far from uniform in total variation distance, (6) implies that some setting of h will reveal a large bias with high probability. The query complexity of this algorithm is

$$\sum_{h=0}^{O(\log(n/\varepsilon))} O\left(2^h \cdot \log\left(\frac{n}{\varepsilon}\right)\right) \cdot \tilde{O}\left(\frac{2^{-h} n}{\varepsilon^2}\right) = \tilde{O}\left(\frac{n}{\varepsilon^2}\right).$$

1.1.3 A Weaker Version of Theorem 1.5

In order to highlight some of the conceptual ideas involved in proving Theorem 1.5, we sketch how one may prove the following weaker inequality, which lower bounds the expected squared ℓ_2 norm instead of the expected ℓ_2 norm:

$$\mathbf{E}_{\rho \sim \mathcal{D}(t+1, p)} \left[\|\mu(p|_{\rho})\|_2^2 \right] \gtrsim \frac{1}{\log^2 n} \cdot \frac{t+1}{n-t} \cdot \mathbf{E}_{\substack{\mathbf{T} \subset [n] \\ |\mathbf{T}|=t}} \left[d_{\text{TV}}(p_{\overline{\mathbf{T}}}, \mathcal{U})^2 \right] \quad (7)$$

where $\mathcal{D}(t+1, p)$ denotes the distribution over random restrictions where we enforce $\text{stars}(\rho) = t+1$ (to compare this to Theorem 1.5, one should think of σ as t/n). Specifically, we sample a random set $\mathbf{S} \subset [n]$ of size $t+1$ and $\mathbf{x} \sim p$, then, we set ρ as in (1). Consider imposing a random order $\pi: [t+1] \rightarrow \mathbf{S}$ and apply linearity of expectation to re-write the left-hand side of (7) as

$$\sum_{i=1}^{t+1} \mathbf{E}_{\rho, \pi} \left[(\mu(p|_{\rho})_{\pi(i)})^2 \right] = (t+1) \mathbf{E}_{\rho, \pi} [\mu(p|_{\rho})_{\pi(t+1)}^2]. \quad (8)$$

Toward lower bounding (8), consider the following way of sampling ρ and π . We first sample t random indices $\pi(1), \dots, \pi(t)$ uniformly from $[n]$ without replacement and let $\mathbf{T} = \{\pi(1), \dots, \pi(t)\}$. Then, we sample $\mathbf{x} \sim p$. Finally, we sample $\pi(t+1)$ uniformly from the set $\overline{\mathbf{T}} = [n] \setminus \mathbf{T}$. We

may write $\mathbf{S} = \mathbf{T} \cup \{\pi(t+1)\}$ and similarly have ρ be set according to (1). Consider the function $\ell: \{-1, 1\}^{\overline{\mathbf{T}}} \rightarrow [-1, \infty)$ which is given by

$$\ell(y) = 2^{|\overline{\mathbf{T}}|} \Pr_{y \sim p_{\overline{\mathbf{T}}}}[\mathbf{y} = y] - 1. \quad (9)$$

We notice that ℓ has mean zero (since $p_{\overline{\mathbf{T}}}$ is a probability distribution), similarly to Section 1.1.2, we orient the edges of the hypercube $\{-1, 1\}^{\overline{\mathbf{T}}}$ from the endpoint with higher value of ℓ to lower value of ℓ . Applying Theorem 1.7 to $f = \ell$ and this orientation of the hypercube edges, the left-hand side is exactly $2d_{\text{TV}}(p_{\overline{\mathbf{T}}}, \mathcal{U})$. Further, we write the random variable $\mathbf{z} = \mathbf{x}_{\overline{\mathbf{T}}}$, which is distributed exactly as $\mathbf{z} \sim p_{\overline{\mathbf{T}}}$, and the random variable $\mathbf{j} = \pi(t+1)$, which is distributed uniformly from $\overline{\mathbf{T}}$. We have, similarly to (4) but with $f = \ell$ rather than p ,

$$\frac{d_{\text{TV}}(p_{\overline{\mathbf{T}}}, \mathcal{U})}{\log n} \lesssim \left(|\overline{\mathbf{T}}| \cdot \mathbf{E}_{\substack{\mathbf{z} \sim p_{\overline{\mathbf{T}}} \\ \mathbf{j} \sim \overline{\mathbf{T}}}} \left[\left(\frac{(\ell(\mathbf{z}) - \ell(\mathbf{z}^{\mathbf{j}}))^+}{1 + \ell(\mathbf{z})} \right)^2 \right] \right)^{1/2}. \quad (10)$$

One crucial observation is that the inner value of the expectation in the right-hand side of (10) is, up to a factor of at most 4, $\mu(p|\rho)_{\pi(t+1)}^2$, where ρ uses $\mathbf{T} \cup \{\mathbf{j}\}$ as stars and non-star values are set to \mathbf{z}_{-j} . Plugging this back into (8) gives (7).

1.2 Related Work

The seminal works of Goldreich, Goldwasser, and Ron [GGR96], Goldreich and Ron [GR00], and Batu, Fortnow, Rubinfeld, Smith, and White [BFR⁺00] initiated the study of distribution testing, viewing probability distributions as a natural application for property testing (see [Gol17] for coverage of this much broader field). Since these works, distribution testing has enjoyed a wealth of study, resulting in a thorough understanding of the complexity of testing many distribution properties (see, e.g., [BFF⁺01, BKR04, Pan08, ADJ⁺11, BFRV11, Val11, ILR12, DDS⁺13, CDVV14, VV17, Wag15, ADK15, BV15, DKN15, DK16, Can16, BCG17, BC17, DKW18, DGPP18], and [Rub12, Can15b, BW18, Kam18] for recent surveys). As a result, sample-optimal algorithms are known for a number of core problems.

However, the known sample complexity lower bounds, while sublinear, typically still involve a polynomial dependence on the domain size, suggesting that new models are needed for studying distribution testing on high-dimensional domains. One approach is to assume additional *structure* from the input distributions. Settings were studied where the distribution is known to be monotone [RS09], a low-degree Bayesian Network [CDKS17, DP17, ABDK18], a Markov Random Field [DDK18, GLP18, BBC⁺19], or having some “flat” histogram structure [DKP19].

The other approach is to allow stronger oracle access. The subcube conditional sampling model, which is the focus of this work, is a variant of the general conditional sampling model particularly apt for the study of high-dimensional distributions. Bhattacharyya and Chakraborty [BC18], who initiated the systematic study of this variant, showed that for many problems of interest such as uniformity, identity, and closeness testing, subcube conditional queries enabled one to avoid the curse of dimensionality, and established sample complexity upper bounds *polynomial* in the dimension (albeit superlinear). The conditional sampling model itself, which was introduced simultaneously by Chakraborty, Fischer, Goldhirsh, and Matsliah [CFGM13, CFGM16], and Canonne, Ron, and Servedio [CRS14, CRS15], allows more general queries: namely, the algorithm may specify an

arbitrary subset of the domain and request a sample conditioned on it lying in the subset. In many cases, the conditional sampling model circumvents sample-complexity lower bounds. Since its introduction, there has been significant study into the complexity of testing a number of properties of distributions under conditional samples, in both adaptive and nonadaptive settings [Can15a, FJO⁺15, ACK15b, FLV17, SSJ17, BCG17, BC18, KT19]. Beyond distribution testing, this model of conditional sampling has found applications in group testing [ACK15a], sublinear algorithms [GTZ17], and crowdsourcing [GTZ18]. Other ways to augment the power of distribution testing algorithms include letting the algorithm query the probability density function (PDF) or cumulative distribution function (CDF) of the distribution [BDKR05, GMV06, RS09, CR14], or giving it probability-revealing samples [OS18].

1.3 Notation and Preliminaries

We use boldface symbols to represent random variables, and non-boldface symbols for fixed values (potentially realizations of these random variables) — see, e.g., $\boldsymbol{\rho}$ versus ρ . Given a set $S \subseteq [n]$, we let \mathcal{U}_S denote the uniform distribution over $\{-1, 1\}^S$. Usually, as the support of \mathcal{U}_S will be clear from the context, we will drop the subscript and simply write \mathcal{U} . We write $f(n) \lesssim g(n)$ if, for some $c > 0$, $f(n) \leq c \cdot g(n)$ for all $n \geq 0$ (the \gtrsim symbol is defined similarly). $f(n) \asymp g(n)$ if $f(n) \lesssim g(n)$ and $f(n) \gtrsim g(n)$. We use the notation $\tilde{O}(f(n))$ as $O(f(n)\text{polylog}(f(n)))$, and $\tilde{\Omega}(f(n))$ to denote $\Omega(f(n)/(1 + |\text{polylog}(f(n))|))$. The notation $[k]$ denotes the set of integers $\{1, \dots, k\}$.

The Frobenius norm of a matrix $M \in \mathbb{R}^{d_1 \times d_2}$ is

$$\|M\|_F = \left(\sum_{i \in [d_1]} \sum_{j \in [d_2]} M_{ij}^2 \right)^{1/2}.$$

For a string $x \in \{-1, 1\}^n$, we use $x^{(i)}$ to denote the string that is identical to x but with coordinate i flipped, i.e., $x_j^{(i)} = x_j$ for all $j \neq i$, and $x_i^{(i)} = -x_i$.

We formally define the subcube conditional query access, which was suggested in Canonne, Ron, Servedio [CRS15] as an instance of the general conditional sampling oracle [CRS15, CFGM16], and first explicitly studied in Bhattacharyya and Chakraborty [BC18].

Definition 1.9. A *subcube conditional sampling* (SCOND) oracle for a distribution p supported on $\{-1, 1\}^n$ is an oracle which accepts a query subcube $\rho \in \{-1, 1, *\}^n$, and outputs a sample from the distribution $\mathbf{x} \sim p$ conditioned on every $i \notin \text{stars}(\rho)$ having $\mathbf{x}_i = \rho_i$. We use the convention that if the algorithm considers a restriction with zero support, the oracle outputs a uniform sample.

2 The Algorithm

We prove our theorem for uniformity testing with subcube conditioning, restated below:

Theorem 2.1. *There exists an algorithm SUBCONDUNI which, given $n \geq 1$, a subcube oracle to a distribution p over $\{-1, 1\}^n$ and a distance parameter ε with $\varepsilon \in (0, 1)$, has the following guarantees. The algorithm makes $\tilde{O}(\sqrt{n}/\varepsilon^2)$ many calls to the oracle and satisfies the following two conditions:*

- (i) *If p is uniform, then the algorithm returns **accept** with probability at least $2/3$.*
- (ii) *If $d_{\text{TV}}(p, \mathcal{U}) \geq \varepsilon$, then the algorithm returns **reject** with probability at least $2/3$.*

The rest of this section is devoted to the proof of Theorem 2.1. For our convenience of working with $\log(1/\varepsilon)$ in the proof, we assume below that $\varepsilon \leq 1/2$ in the input of SUBCONDUNI. This way $\log(c/\varepsilon)$ can be treated as $O(\log(1/\varepsilon))$ whenever $c \geq 1$ is a fixed constant. As discussed earlier in Section 1, SUBCONDUNI (see Algorithm 1) is based on Theorem 1.5 and Lemma 1.4, which we now prove.

Proof of Lemma 1.4: Fix any subset $S \subseteq [n]$ of size t . Given $u \in \{-1, 1\}^{\bar{S}}$, we let $p_{|\rho(S,u)}$ denote the distribution supported on $\{-1, 1\}^S$ given by drawing $\mathbf{x} \sim p$ conditioned on $\mathbf{x}_{\bar{S}} = u$. By unrolling the definition of total variation distance, we have

$$\begin{aligned}
2d_{\text{TV}}(p, \mathcal{U}) &= \sum_{x \in \{-1, 1\}^n} |p(x) - 1/2^n| \\
&= \sum_{u \in \{-1, 1\}^{\bar{S}}} \sum_{v \in \{-1, 1\}^S} \left| \Pr_{\mathbf{x} \sim p} [\mathbf{x}_{\bar{S}} = u \wedge \mathbf{x}_S = v] - 1/2^n \right| \\
&= \sum_{u \in \{-1, 1\}^{\bar{S}}} \sum_{v \in \{-1, 1\}^S} \left| \Pr_{\mathbf{u} \sim p_{\bar{S}}} [\mathbf{u} = u] \cdot \Pr_{\mathbf{x} \sim p} [\mathbf{x}_S = v \mid \mathbf{x}_{\bar{S}} = u] - 1/2^n \right| \\
&\leq \sum_{u \in \{-1, 1\}^{\bar{S}}} \sum_{v \in \{-1, 1\}^S} \left(\left| p_{\bar{S}}(u) \cdot \Pr_{\mathbf{x} \sim p} [\mathbf{x}_S = v \mid \mathbf{x}_{\bar{S}} = u] - \frac{p_{\bar{S}}(u)}{2^t} \right| + \left| \frac{p_{\bar{S}}(u)}{2^t} - \frac{1/2^{n-t}}{2^t} \right| \right) \\
&= \sum_{u \in \{-1, 1\}^{\bar{S}}} p_{\bar{S}}(u) \cdot 2d_{\text{TV}}(p_{|\rho(S,u)}, \mathcal{U}) + \sum_{v \in \{-1, 1\}^S} \frac{1}{2^t} \cdot 2d_{\text{TV}}(p_{\bar{S}}, \mathcal{U}) \\
&= \sum_{u \in \{-1, 1\}^{\bar{S}}} p_{\bar{S}}(u) \cdot 2d_{\text{TV}}(p_{|\rho(S,u)}, \mathcal{U}) + 2d_{\text{TV}}(p_{\bar{S}}, \mathcal{U}).
\end{aligned}$$

Taking the expectation of the inequality over the choice of $\mathbf{S} \sim \mathcal{S}_\sigma$ yields the lemma. ■

Let p be a distribution over $\{-1, 1\}^n$ with $d_{\text{TV}}(p, \mathcal{U}) \geq \varepsilon$. Let

$$\sigma \stackrel{\text{def}}{=} \sigma(\varepsilon) = \frac{1}{C_0 \cdot \log^4(16/\varepsilon)}$$

where $C_0 > 0$ is an absolute constant. (The value of C_0 is only used at the end of this section, where setting $C_0 = 10^{11}$ is good enough.) Let us further assume that n and ε together satisfy

$$e^{-\sigma n/10} \leq \varepsilon/8. \tag{11}$$

Violation of (11) implies that

$$n = O\left(\frac{1}{\sigma} \cdot \log\left(\frac{1}{\varepsilon}\right)\right) = O\left(\log^5\left(\frac{1}{\varepsilon}\right)\right) \tag{12}$$

and we handle this case by applying the linear-query tester described in Section 1.1.2 with query complexity $\tilde{O}(n/\varepsilon^2) = \tilde{O}(1/\varepsilon^2)$ using (12).

For the general case with (11) satisfied, consider a distribution p with $d_{\text{TV}}(p, \mathcal{U}) \geq \varepsilon$. Lemma 1.4 implies that either $\mathbf{E}_{\mathbf{S} \sim \mathcal{S}_\sigma} [d_{\text{TV}}(p_{\bar{S}}, \mathcal{U})] \geq \varepsilon/2$ or

$$\mathbf{E}_{\rho \sim \mathcal{D}_\sigma(p)} [d_{\text{TV}}(p_{|\rho}, \mathcal{U})] \geq \varepsilon/2. \tag{13}$$

Algorithm 1 SUBCONDUNI(n, p, ε)

Require: Dimension n , oracle access to distribution p over $\{-1, 1\}^n$, and parameter $\varepsilon \in (0, 1/2]$

- 1: **StartBaseCase** ▷ Base case: violation of (11)
- 2: **if** n and ε violate (11) **then**
- 3: Run the linear tester described in Section 1.1.2 and **return** the same answer
- 4: **EndBaseCase**
- 5: **StartMainCase** ▷ General case: (11) satisfied
- 6: Let $L = L(n, \varepsilon) = \tilde{O}(\sqrt{n}/\varepsilon)$ be as defined in (15) to simplify (14)
- 7: **for** $j = 1, 2, \dots, \lceil \log 2L \rceil$ **do** ▷ Test (14), via bucketing
- 8: Sample $s_j = 8L \log(2L) \cdot 2^{-j}$ restrictions from $\mathcal{D}_\sigma(p)$
- 9: **for** every restriction ρ sampled with $|\text{stars}(\rho)| > 0$ **do**
- 10: Run TESTMEAN($|\text{stars}(\rho)|, p_{|\rho}, 2^{-j}$) for $r = O(\log(n/\varepsilon))$ times
- 11: **return** reject if the majority of calls return reject
- 12: **for** $j = 1, 2, \dots, \lceil \log(4/\varepsilon) \rceil$ **do** ▷ Test the second part of (13), recursively
- 13: Sample $s'_j = (32/\varepsilon) \log(4/\varepsilon) \cdot 2^{-j}$ restrictions from $\mathcal{D}_\sigma(p)$
- 14: **for** each restriction ρ sampled satisfying $0 < |\text{stars}(\rho)| \leq 2\sigma n$ **do**
- 15: Run SUBCONDUNI($|\text{stars}(\rho)|, p_{|\rho}, 2^{-j}$) for $t = 100 \log(16/\varepsilon)$ times
- 16: **return** reject if the majority of calls return reject
- 17: **EndMainCase**
- 18: **return** accept

Assuming the former and using (11), we have from Theorem 1.5 that (using $\sigma = 1/\text{polylog}(1/\varepsilon)$)

$$\mathbf{E}_{\rho \sim \mathcal{D}_\sigma(p)} \left[\|\mu(p_{|\rho})\|_2 / \sqrt{n} \right] \geq \tilde{\Omega} \left(\frac{\varepsilon}{\sqrt{n}} \right). \quad (14)$$

This naturally lends itself to a recursive approach: for the general case when n and ε satisfy (11), we use TESTMEAN from Theorem 1.3 to test (14), and recursive calls to SUBCONDUNI to test (13).

The description of SUBCONDUNI is given in Algorithm 1. For convenience, we let

$$L \stackrel{\text{def}}{=} L(n, \varepsilon) = \tilde{O}(\sqrt{n}/\varepsilon) \quad (15)$$

such that the right hand side of (14) can be replaced by $1/L$.

We are now ready to prove Theorem 2.1.

Proof of Theorem 2.1: We start with (i) (the completeness) and prove by induction on n that, when p is uniform, SUBCONDUNI(n, p, ε) returns **accept** with probability at least $2/3$. The base case when $n = 1$ is trivial since (11) is violated and we just run the linear tester. Assuming that the statement is true for all dimensions smaller than n , we focus on the case when (11) is satisfied; the case when it is violated is again trivial. Since p is uniform, the random restriction $p_{|\rho}$ is uniform with probability one. On the other hand, note that the total number of restrictions ρ we draw in line 8 is $\sum_j s_j = O(L \log L) = \tilde{O}(\sqrt{n}/\varepsilon^2)$. As a result, one can set the constant hidden in the choice of r in line 10 to be sufficiently large so that SUBCONDUNI rejects in line 11 with probability no larger than $1/6$. Similarly, by invoking the inductive hypothesis, one can show that SUBCONDUNI rejects in line 16 with probability no larger than $1/6$. It follows from a union bound that SUBCONDUNI rejects with probability at most $1/3$ and this finishes the induction step.

We now turn towards establishing (ii) (the soundness), and hereafter assume that $d_{\text{TV}}(p, \mathcal{U}) \geq \varepsilon$. Again we prove by induction on n that $\text{SUBCONDUNI}(n, p, \varepsilon)$ rejects with probability at least $2/3$. For the general case in the induction step, it follows from our discussion earlier that either (14) or (13) holds (and in particular, the right hand side of (14) can be replaced by $1/L$).

For the first case, we note that the left hand side of (14) is the expectation of a random variable that takes values between 0 and 1 (and the right hand side can be replaced by $1/L$). Thus, with a simple bucketing argument there must exist a $j \in [\lceil \log 2L \rceil]$ such that

$$\Pr_{\rho \sim \mathcal{D}_{\sigma}(p)} \left[\frac{\|\mu(p|\rho)\|_2}{\sqrt{n}} \geq 2^{-j} \right] \geq \frac{2^{j-1}}{L \lceil \log 2L \rceil} \geq \frac{2^j}{4L \log 2L}.$$

It follows that one of the restrictions sampled satisfies the condition above with probability at least $1 - e^{-2} > 5/6$. When this happens, each call to TESTMEAN in line 15 rejects with probability at least $2/3$. As a result SUBCONDUNI rejects in line 11 with probability at least $2/3$.

For the second case, again by the bucketing argument, there exists a $j \in [\lceil \log(4/\varepsilon) \rceil]$ such that

$$\Pr_{\rho \sim \mathcal{D}_{\sigma}(p)} \left[d_{\text{TV}}(p|\rho, \mathcal{U}) \geq 2^{-j} \right] \geq \frac{\varepsilon 2^j}{4 \lceil \log(4/\varepsilon) \rceil} \geq \frac{\varepsilon 2^j}{8 \log(4/\varepsilon)}.$$

By (11), the probability of $|\text{stars}(\rho)| > 2\sigma n$ is by a Chernoff bound at most $e^{-\sigma n/3} < (\varepsilon/8)^3$. Thus,

$$\Pr_{\rho \sim \mathcal{D}_{\sigma}(p)} \left[d_{\text{TV}}(p|\rho, \mathcal{U}) \geq 2^{-j} \text{ and } 0 < |\text{stars}(\rho)| \leq 2\sigma n \right] \geq \frac{\varepsilon 2^j}{16 \log(4/\varepsilon)}.$$

Given our choice of s'_j , the probability that at least one of ρ satisfies $d_{\text{TV}}(p|\rho, \mathcal{U}) \geq 2^{-j}$ is at least $5/6$. The probability that the majority of calls reject is also at least $5/6$ (for this case we just need t to be a large enough constant). Thus, SUBCONDUNI rejects with probability at least $2/3$.

Finally, we bound the query complexity of SUBCONDUNI . Let $\Phi(n, \varepsilon)$ denote the complexity of $\text{SUBCONDUNI}(n, p, \varepsilon)$. We will show by induction on n that

$$\Phi(n, \varepsilon) \leq C \cdot \frac{\sqrt{n}}{\varepsilon^2} \cdot \log^c \left(\frac{n}{\varepsilon} \right) \tag{16}$$

for some absolute constants $C, c > 0$. To fix C and c , we let C_1 and c_1 be two constants such that the query complexity of the linear tester running on n and ε that violate (11) can be bounded by

$$C_1 \cdot \frac{1}{\varepsilon^2} \cdot \log^{c_1} \left(\frac{1}{\varepsilon} \right).$$

We also let C_2 and c_2 be two constants such that the query complexity of SUBCONDUNI between line 7 and 11 (the non-recursive part) can be bounded by

$$C_2 \cdot \frac{\sqrt{n}}{\varepsilon^2} \cdot \log^{c_2} \left(\frac{n}{\varepsilon} \right).$$

For this, note that the query complexity of this part is (recall that $L = \tilde{O}(\sqrt{n}/\varepsilon)$)

$$\sum_{j=1}^{\lceil \log 2L \rceil} \frac{L \log L}{2^j} \cdot \tilde{O} \left(\max \left\{ \frac{2^{2j}}{\sqrt{n}}, 2^j \right\} \right) \cdot \log \left(\frac{n}{\varepsilon} \right) = \tilde{O} \left(\frac{\sqrt{n}}{\varepsilon^2} \right).$$

We then set $C \stackrel{\text{def}}{=} 2 \max(C_1, C_2)$ and $c \stackrel{\text{def}}{=} \max(c_1, c_2)$.

To prove (16), the base case when $n = 1$ is trivial. For the induction step, the special case when (11) is violated is also trivial. For the general case, we have (by the choice of C and c)

$$\Phi(n, \varepsilon) \leq \frac{C}{2} \cdot \frac{\sqrt{n}}{\varepsilon^2} \cdot \log^c \left(\frac{n}{\varepsilon} \right) + \sum_{j=1}^{\lceil \log(4/\varepsilon) \rceil} s'_j \cdot 100 \log \left(\frac{16}{\varepsilon} \right) \cdot \Phi(2\sigma n, 2^{-j}).$$

Using the inductive hypothesis (and our choice of σ), each term in the second sum becomes

$$C \cdot \frac{32}{\varepsilon} \cdot \log \left(\frac{4}{\varepsilon} \right) \cdot 100 \cdot \log \left(\frac{16}{\varepsilon} \right) \cdot \sqrt{2\sigma n} \cdot 2^j \cdot \log^c (\sigma n 2^{j+1}) \leq \frac{C}{32} \cdot \frac{\sqrt{n}}{\varepsilon} \cdot 2^j \cdot \log^c \left(\frac{n}{\varepsilon} \right),$$

using $C_0 \geq (32^2 \cdot 100)^2 \cdot 2$. We can then finish the induction by using

$$\sum_{j=1}^{\lceil \log(4/\varepsilon) \rceil} 2^j < 2^{\lceil \log(4/\varepsilon) \rceil + 1} \leq \frac{16}{\varepsilon}.$$

This concludes the proof of the theorem. ■

3 Proof of Theorem 1.5

Let $\mathcal{S}(t)$ be the uniform distribution supported on all subsets of $[n]$ of size t . Given a distribution p supported on $\{-1, 1\}^n$, we let $\mathcal{D}(t, p)$ be the distribution supported on restrictions $\{-1, 1, *\}^n$ given in a similar fashion to that of $\mathcal{D}_\sigma(p)$, except we use sets of size t ; we sample $\mathbf{S} \sim \mathcal{S}(t)$ and $\mathbf{x} \sim p$, and we let $\rho_i = *$ if $i \in \mathbf{S}$ and \mathbf{x}_i otherwise. The bulk of the work goes into proving the following lemma. After the statement, we show the lemma implies Theorem 1.5.

Lemma 3.1. *Let p be a distribution supported on $\{-1, 1\}^n$, $t \in [n - 1]$, and denote*

$$\alpha \stackrel{\text{def}}{=} \mathbf{E}_{\mathbf{T} \sim \mathcal{S}(t)} \left[d_{\text{TV}}(p_{\overline{\mathbf{T}}}, \mathcal{U}) \right] \geq 0.$$

Then,

$$\mathbf{E}_{\rho \sim \mathcal{D}(t, p)} \left[\|\mu(p|_\rho)\|_2 \right] + \mathbf{E}_{\rho \sim \mathcal{D}(t+1, p)} \left[\|\mu(p|_\rho)\|_2 \right] \gtrsim \frac{t}{n} \cdot \frac{\alpha}{\log^2 n \cdot \log(n/\alpha) \cdot \log(1/\alpha)}. \quad (17)$$

Proof of Theorem 1.5 assuming Lemma 3.1: We assume that $5 \leq \sigma n \leq n - 5$ in the proof; otherwise the statement is trivially satisfied. Let $\delta = \min(\sigma, 1 - \sigma)/2 > 0$.

For each $t \in [n - 1]$, we write for the sake of this proof

$$\alpha_t \stackrel{\text{def}}{=} \mathbf{E}_{\mathbf{T} \sim \mathcal{S}(t)} \left[d_{\text{TV}}(p_{\overline{\mathbf{T}}}, \mathcal{U}) \right].$$

Notice that $\rho \sim \mathcal{D}_\sigma(p)$ may be sampled by drawing $\mathbf{k} \sim \text{Bin}(n, \sigma)$ and $\rho \sim \mathcal{D}(\mathbf{k}, p)$. We write β_t for the probability over $\mathbf{k} \sim \text{Bin}(n, \sigma)$ that $\mathbf{k} = t$. Let

$$B := \left\{ t \in [n - 1] : t/n \in [\sigma - \delta, \sigma + \delta] \right\}.$$

Then, by Chernoff's bound we have $\sum_{t \in B} \beta_t \geq 1 - 2e^{-\delta n/5}$ and thus,

$$\sum_{t \in B} \beta_t \cdot \alpha_t \geq \mathbf{E}_{\mathbf{T} \sim \mathcal{S}_\sigma} \left[d_{\text{TV}}(p_{\mathbf{T}}, \mathcal{U}) \right] - 2e^{-\delta n/5}. \quad (18)$$

Then, we have

$$\begin{aligned} 2 \cdot \mathbf{E}_{\rho \sim \mathcal{D}_\sigma(p)} \left[\|\mu(p|\rho)\|_2 \right] &\geq \sum_{t \in B} \left(\beta_t \cdot \mathbf{E}_{\rho \sim \mathcal{D}(t,p)} \left[\|\mu(p|\rho)\|_2 \right] + \beta_{t+1} \cdot \mathbf{E}_{\rho \sim \mathcal{D}(t+1,p)} \left[\|\mu(p|\rho)\|_2 \right] \right) \\ &\gtrsim \sum_{t \in B} \beta_t \left(\mathbf{E}_{\rho \sim \mathcal{D}(t,p)} \left[\|\mu(p|\rho)\|_2 \right] + \mathbf{E}_{\rho \sim \mathcal{D}(t+1,p)} \left[\|\mu(p|\rho)\|_2 \right] \right) \end{aligned} \quad (19)$$

$$\gtrsim \frac{\sigma}{\log^2 n} \cdot \sum_{t \in B} \beta_t \cdot \frac{\alpha_t}{\log(n/\alpha_t) \log(1/\alpha_t)} \quad (20)$$

$$\gtrsim \frac{\sigma}{\text{polylog}(n)} \cdot \tilde{\Omega} \left(\mathbf{E}_{\mathbf{T} \sim \mathcal{S}_\sigma} \left[d_{\text{TV}}(p_{\mathbf{T}}, \mathcal{U}) \right] - 2e^{-\delta n/5} \right). \quad (21)$$

In (19), we used $t/n \in [\sigma - \delta, \sigma + \delta]$, $\delta = \min(\sigma, 1 - \sigma)/2$ and $\sigma \geq 5/n$ to have

$$\frac{\beta_{t+1}}{\beta_t} = \frac{n-t}{t+1} \cdot \frac{\sigma}{1-\sigma} \geq \frac{(1-\sigma)/2}{(3\sigma/2) + (1/n)} \cdot \frac{\sigma}{1-\sigma} \gtrsim 1.$$

In (20) we applied Lemma 3.1 on each $t \in B$. In (21) we applied Jensen's inequality (as the function $f(a) = a/(\log(n/a) \log(1/a))$ when $a \neq 0$ and $f(0) = 0$ is convex in $[0, 1]$); and (18).

This finishes the proof of Theorem 1.5. ■

3.1 Robust Pisier Inequality

In this section, we prove the robust version of Pisier's inequality. Robustness here is equivalent to that of [KMS18] where we will consider a function $f: \{-1, 1\}^n \rightarrow \mathbb{R}$, and we will lower bound a functional on the values of the edges of the hypercube after assigning them directions.

Formally, fix $n \in \mathbb{N}$ and let \mathcal{H} be the undirected graph over the hypercube $\{-1, 1\}^n$ that consists of undirected edges $\{x, x^{(i)}\}$ with $x \in \{-1, 1\}^n$ and $i \in [n]$. For $i \in [n]$, recall that $L_i f: \{-1, 1\}^n \rightarrow \mathbb{R}$ is the linear operator given by

$$L_i f(x) = \frac{f(x) - f(x^{(i)})}{2}.$$

Notice that in Theorem 1.6, every edge $\{x, x^{(i)}\}$ in \mathcal{H} for a fixed value of \mathbf{y} is counted twice in the right-hand side; once for the endpoint x and once for the endpoint $x^{(i)}$. In the robust version, we may arbitrarily choose, for each edge $\{x, x^{(i)}\}$, whether to "charge" the edge to x or to $x^{(i)}$. For this purpose, we consider an orientation G of \mathcal{H} (so for each $\{x, x^{(i)}\}$ in \mathcal{H} , G contains either $(x, x^{(i)})$ or $(x^{(i)}, x)$), and charge an edge $\{x, x^{(i)}\}$ to x if $(x, x^{(i)})$ is in G and to $x^{(i)}$ otherwise. For convenience, we will abuse the notation in the rest of the section to use the name of a directed graph (such as G) to denote its edge set as well, since its vertex set is usually clear from the context. So we will write $(u, v) \in G$ if (u, v) is a directed edge in G .

We are now ready to state the robust version of Pisier's inequality, which generalizes Theorem 1.7.

Theorem 3.2 (Robust Pisier's inequality). *Let $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ be a function with*

$$\mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} [f(\mathbf{x})] = 0 \quad (22)$$

and let G be an orientation of \mathcal{H} . Then, for any $s \in [1, \infty)$ we have

$$\left(\mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} \left[|f(\mathbf{x})|^s \right] \right)^{1/s} \lesssim \log n \cdot \left(\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^n} \left[\left| \sum_{\substack{i \in [n] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G}} \mathbf{y}_i \mathbf{x}_i L_i f(\mathbf{x}) \right|^s \right] \right)^{1/s}.$$

The proof itself follows the template of [NS02, Theorem 2] and checks that the necessary changes still give the desired inequality. Before beginning with the proof, we recall some basic notions of Fourier analysis on the hypercube, and define some elements which appear in the argument. Recall that any function $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ has a unique Fourier expansion $f(x) = \sum_{S \subseteq [n]} \hat{f}(S) \chi_S(x)$, where $\chi_S(x) = \prod_{i \in S} x_i$ are the Fourier characters, and $\hat{f}(S) = \mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} [f(\mathbf{x}) \chi_S(x)]$. For $\rho > 0$ and $x \in \{-1, 1\}^n$, we let $N_\sigma(x)$ be the distribution supported on $\{-1, 1\}^n$ given sampling $\mathbf{y} \sim N_\sigma(x)$, where for each $i \in [n]$, we set $\mathbf{y}_i = \mathbf{x}_i$ with probability $1 - \sigma$, and a uniform random bit with probability σ . We denote $T_\rho f(x) = \mathbf{E}_{\mathbf{y} \sim N_\rho(x)} [f(\mathbf{y})]$, and we have for every $S \subseteq [n]$, $\widehat{T_\rho f}(S) = \rho^{|S|} \hat{f}(S)$. In a slight abuse of notation, we consider for any $x, y \in \{-1, 1\}^n$ and $t \in [0, 1]$, the distribution $N_{t, 1-t}(x, y)$, supported on $\{-1, 1\}^n$, to be the distribution given by letting $\mathbf{z} \sim N_{t, 1-t}(x, y)$ have each $i \in [n]$ set to $\mathbf{z}_i = \mathbf{x}_i$ with probability t and $\mathbf{z}_i = \mathbf{y}_i$ otherwise. For a function $g: \{-1, 1\}^n \rightarrow \mathbb{R}$ and $t \in [0, 1]$, the function $g: \{-1, 1\}^n \times \{-1, 1\}^n \rightarrow \mathbb{R}$ is given by letting $g_{t, 1-t}(x, y) = \mathbf{E}_{\mathbf{z} \sim N_{t, 1-t}(x, y)} [g(\mathbf{z})] = \sum_{S \subseteq [n]} \hat{g}(S) \prod_{i \in S} (tx_i + (1-t)y_i)$. Lastly, for any $\gamma > 0$, we let $\Delta^\gamma f$ be the linear operator given by $\Delta^\gamma f(x) = \sum_{S \subseteq [n]} \hat{f}(S) |S|^\gamma \chi_S(x)$.

Proof of Theorem 3.2: Let $\rho = 1 - 1/(n+1)$ and $q \in [1, \infty]$ such that $\frac{1}{s} + \frac{1}{q} = 1$. Fix the function $g: \{-1, 1\}^n \rightarrow \mathbb{R}$ with $\|g\|_q = 1$ satisfying $\langle T_\rho f, g \rangle = \|T_\rho f\|_s$. We have

$$\begin{aligned} \rho^n \|f\|_s &\leq \|T_\rho f\|_s = \langle T_\rho f, g \rangle = \sum_{\substack{S \subseteq [n] \\ S \neq \emptyset}} \rho^{|S|} \hat{f}(S) \hat{g}(S) \\ &= \frac{1}{\Gamma(1+\gamma)} \int_0^\rho \left(\sum_{\substack{S \subseteq [n] \\ S \neq \emptyset}} t^{|S|-1} |S|^{\gamma+1} \hat{f}(S) \hat{g}(S) \right) (\log(\rho/t))^\gamma dt \end{aligned} \quad (23)$$

$$= \frac{1}{\Gamma(1+\gamma)} \int_0^\rho \frac{1}{1-t} \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^n} \left[g_{t, 1-t}(\mathbf{x}, \mathbf{y}) \sum_{i=1}^n \mathbf{y}_i \mathbf{x}_i L_i \Delta^\gamma f(\mathbf{x}) \right] (\log(\rho/t))^\gamma dt, \quad (24)$$

where (23) follows from writing $\rho^{|S|} = \frac{1}{\Gamma(1+\gamma)} \int_0^\rho t^{|S|-1} |S|^{\gamma+1} \log(\rho/t)^\gamma dt$, and (24) follows from the

Fourier expansion of $g_{t,1-t}(x, y) \sum_{i=1}^n y_i x_i L_i \Delta^\gamma f(x)$. The main step is obtaining (25) below:

$$\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^n} \left[g_{t,1-t}(\mathbf{x}, \mathbf{y}) \sum_{i=1}^n \mathbf{y}_i \mathbf{x}_i L_i \Delta^\gamma f(\mathbf{x}) \right] = 2 \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^n} \left[g_{t,1-t}(\mathbf{x}, \mathbf{y}) \sum_{\substack{i \in [n] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G}} \mathbf{y}_i \mathbf{x}_i L_i \Delta^\gamma f(\mathbf{x}) \right] \quad (25)$$

$$\leq 2 \left(\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^n} \left[\left| \sum_{\substack{i \in [n] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G}} \mathbf{y}_i \mathbf{x}_i L_i \Delta^\gamma f(\mathbf{x}) \right|^s \right] \right)^{1/s}, \quad (26)$$

since (26) follows from (25) by noting that $\|g_{t,1-t}\|_q \leq 1$. We may now proceed substituting (26) into (23). For $\rho = 1 - 1/(n+1)$, we have that for γ approaching 0, we have

$$\frac{1}{\Gamma(1+\gamma)} \int_0^\rho \frac{\log(\rho/t)^\gamma}{1-t} dt \lesssim \log n,$$

so we obtain

$$\rho^n \|f\|_s \lesssim \log n \left(\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^n} \left[\left| \sum_{\substack{i \in [n] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G}} \mathbf{y}_i \mathbf{x}_i L_i \Delta^\gamma f(\mathbf{x}) \right|^s \right] \right)^{1/s},$$

and the right-hand side approaches the desired quantity, while the left-hand side is independent of γ . It thus remains to show (25). For simplicity, let $h_i = L_i \Delta^\gamma f$. We consider summing the edges of \mathcal{H} according to the orientation of the edge:

$$\begin{aligned} & \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^n} \left[g_{t,1-t}(\mathbf{x}, \mathbf{y}) \sum_{i=1}^n \mathbf{y}_i \mathbf{x}_i h_i(\mathbf{x}) \right] \quad (27) \\ &= \mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} \left[\sum_{\substack{i \in [n] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G}} \mathbf{E}_{\mathbf{y} \sim \{-1, 1\}^n} \left[\mathbf{y}_i \left(g_{t,1-t}(\mathbf{x}, \mathbf{y}) h_i(\mathbf{x}) - g_{t,1-t}(\mathbf{x}^{(i)}, \mathbf{y}) h_i(\mathbf{x}^{(i)}) \right) \right] \right] \end{aligned}$$

Now, for a fixed x and i , we have

$$\begin{aligned} & \mathbf{E}_{\mathbf{y} \sim \{-1, 1\}^n} \left[\mathbf{y}_i \left(g_{t,1-t}(x, \mathbf{y}) h_i(x) - g_{t,1-t}(x^{(i)}, \mathbf{y}) h_i(x^{(i)}) \right) \right] \\ &= \mathbf{E}_{\mathbf{y} \sim \{-1, 1\}^n} \left[\mathbf{y}_i g_{t,1-t}(x, \mathbf{y}) \left(h_i(x) - h_i(x^{(i)}) \right) \right]. \end{aligned}$$

This is because when expanding the terms in $g_{t,1-t}(x, \mathbf{y}) = \mathbf{E}_{\mathbf{z} \sim N_{t,1-t}(x, \mathbf{y})} [g(\mathbf{z})]$ two things may happen. 1) Either $\mathbf{z}_i = \mathbf{y}_i$ (for $(1-t)$ -fraction of terms), in which case the terms in $\mathbf{E}_{\mathbf{z} \sim N_{t,1-t}(x, \mathbf{y})} [g(\mathbf{z})]$ and $\mathbf{E}_{\mathbf{z} \sim N_{t,1-t}(x^{(i)}, \mathbf{y})} [g(\mathbf{z})]$ are the same, and these are scaled by $h_i(x)$ and $-h_i(x^{(i)})$, respectively.

Or, 2) $\mathbf{z}_i = x^{(i)}$, in which case terms in $\mathbf{E}_{\mathbf{z} \sim N_{t,1-t}(x,\mathbf{y})}[g(\mathbf{z})]$ and $\mathbf{E}_{\mathbf{z} \sim N_{t,1-t}(x^{(i)},\mathbf{y})}[g(\mathbf{z})]$ are scaled by $h_i(x)$ and $-h_i(x^{(i)})$, respectively. However, in this case, these terms are all *independent* of \mathbf{y}_i , and thus are added and subtracted for an overall contribution of zero. With these considerations, (27) becomes

$$\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1,1\}^n} \left[g_{t,1-t}(\mathbf{x}, \mathbf{y}) \sum_{i=1}^n \mathbf{y}_i \mathbf{x}_i h_i(\mathbf{x}) \right] = \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1,1\}^n} \left[\sum_{\substack{i \in [n] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G}} \mathbf{y}_i g_{t,1-t}(\mathbf{x}, \mathbf{y}) \left(h_i(\mathbf{x}) - h_i(\mathbf{x}^{(i)}) \right) \right]$$

and since $h_i(\mathbf{x}) - h_i(\mathbf{x}^{(i)}) = 2\Delta^\gamma L_i f(\mathbf{x})$, we obtain the desired bound. \blacksquare

3.2 Plan of the Proof of Lemma 3.1

Let $t \in [n-1]$ be the parameter in the statement of Lemma 3.1. For clarity, the rest of this section will always use T to denote a size- t subset of $[n]$ and S to denote a size- $(t+1)$ subset of $[n]$. For each size- t subset T of $[n]$, we write $\alpha(T) \stackrel{\text{def}}{=} d_{\text{TV}}(p_{\bar{T}}, \mathcal{U})$. Then $\alpha = \mathbf{E}_{\mathbf{T} \sim S(t)}[\alpha(\mathbf{T})]$. We also write $\mathcal{H}(T)$ to denote the undirected graph over the hypercube $\{-1,1\}^{\bar{T}}$ that consists of undirected edges $\{x, x^{(i)}\}$ with $x \in \{-1,1\}^{\bar{T}}$ and $i \in \bar{T}$. Again we will abuse the notation in the rest of the section to refer to $\mathcal{H}(T)$ as its edge set as well.

The proof of Lemma 3.1 consists of two steps. For each t -subset T , we first classify undirected edges $\{x, x^{(i)}\}$ in $\mathcal{H}(T)$ into different types according its *weight* defined as

$$w(\{x, x^{(i)}\}) \stackrel{\text{def}}{=} \frac{|p_{\bar{T}}(x) - p_{\bar{T}}(x^{(i)})|}{\max\{p_{\bar{T}}(x), p_{\bar{T}}(x^{(i)})\}}.$$

For each type of edges in $\mathcal{H}(T)$, we describe a method in Section 3.3 to assign each edge a direction. This then leads to a sequence of directed graphs over $\{-1,1\}^{\bar{T}}$, one for each type of edges in $\mathcal{H}(T)$, and their union is an orientation G of $\mathcal{H}(T)$ over $\{-1,1\}^{\bar{T}}$. At the end of Section 3.3 we apply the robust Pisier inequality (Theorem 3.2) on a shifted and scaled version of the probability mass function of $p_{\bar{T}}$ with the orientation G and $s = 1$. The result will be Lemma 3.6, which says there is a type of edges in $\mathcal{H}(T)$ such that its corresponding directed graph has

$$\mathbf{E}_{\mathbf{x} \sim p_{\bar{T}}} \left[\sqrt{\text{out-degree of } \mathbf{x}} \right] \tag{28}$$

bounded from below by a quantity that is linear in $\alpha(T)$; see Lemma 3.6 for details.

In the second step, we use this family of directed graphs promised by Lemma 3.6, one for each t -subset T , with the desired bound on (28) to finish the proof of Lemma 3.1. To this end we first apply standard bucketing arguments in Section 3.4 to simplify the situation, by focusing on one specific type of directed edges that makes the most significant contribution in the family. The final connection from these directed graphs to the mean vectors of randomly restricted distributions is made in Sections 3.5, 3.6, and 3.7. There will be two cases, depending on whether the type of edges we consider has large or small weights. They are handled in Sections 3.6 and 3.7, respectively.

3.3 From Total Variation to Directed Graphs

Let ℓ be a probability distribution over $\{-1,1\}^m$ with $m = n - t$. (Later we will identify ℓ as $p_{\bar{T}}$ for some t -subset T of $[n]$, and $\{-1,1\}^m$ as $\{-1,1\}^{\bar{T}}$.) Let \mathcal{H} denote the undirected graph over

$\{-1, 1\}^m$ that consists of undirected edges $\{x, x^{(i)}\}$ for all $x \in \{-1, 1\}^m$ and $i \in [m]$. Looking ahead, the purpose of this section is to construct an orientation G of \mathcal{H} for our application of Theorem 3.2 later with $s = 1$ and the function $f: \{-1, 1\}^m \rightarrow [-1, \infty)$ given by:

$$f(y) = 2^m \cdot \ell(y) - 1. \quad (29)$$

Note that $\mathbf{E}_y[f(\mathbf{y})] = 0$ and the left hand side of the robust Pisier inequality is $2d_{\text{TV}}(\ell, \mathcal{U})$.

For the construction of G , we start with a classification of undirected edges in \mathcal{H} .

Definition 3.3. An undirected edge $\{x, x^{(i)}\} \in \mathcal{H}$ is said to be a *zero edge* if $\ell(x) = \ell(x^{(i)})$ (they are called zero edges because the difference $\ell(x) - \ell(x^{(i)}) = 0$).

For each nonzero edge $\{x, x^{(i)}\} \in \mathcal{H}$, we define its *weight* as

$$w(\{x, x^{(i)}\}) \stackrel{\text{def}}{=} \frac{|\ell(x) - \ell(x^{(i)})|}{\max\{\ell(x), \ell(x^{(i)})\}}.$$

Note that the weight of an undirected edge is always in $(0, 1]$. We say a nonzero edge is *uneven* if its weight is at least $2/3$; otherwise, we call it an *even* edge (i.e., any nonzero edge with weight smaller than $2/3$). We say an even edge is at *scale* κ for some $\kappa \geq 1$ if

$$2^{-\kappa} < w(\{x, x^{(i)}\}) \leq 2^{-\kappa+1}.$$

We write $\mathcal{H}^{[z]}$ to denote the set of all zero edges, $\mathcal{H}^{[u]}$ to denote the set of all uneven edges, and $\mathcal{H}^{[\kappa]}$ for each $\kappa \geq 1$ to denote the set of even edges at scale κ . Hence $\mathcal{H}^{[z]}$, $\mathcal{H}^{[u]}$ and $\mathcal{H}^{[\kappa]}$ with $\kappa \geq 1$ together form a partition of \mathcal{H} ; we also view them as undirected graphs over $\{-1, 1\}^m$.

Next we construct a sequence of directed graphs $G^{[z]}$, $G^{[u]}$ and $G^{[\kappa]}$, $\kappa \geq 1$, as orientations of $\mathcal{H}^{[z]}$, $\mathcal{H}^{[u]}$ and $\mathcal{H}^{[\kappa]}$, respectively. We start with $G^{[z]}$ and $G^{[u]}$. For each zero edge $\{x, x^{(i)}\} \in \mathcal{H}^{[z]}$, we orient it arbitrarily in $G^{[z]}$. Next for each uneven edge $\{x, x^{(i)}\} \in \mathcal{H}^{[u]}$, we orient it from x to $x^{(i)}$ if $\ell(x) > \ell(x^{(i)})$ (note that if $\ell(x) = \ell(x^{(i)})$ then it is a zero edge).

Orientations of even edges at scale κ in $G^{[\kappa]}$ are more involved. For a fixed $\kappa \geq 1$, we consider $\mathcal{H}^{[\kappa]}$ as an undirected graph over $\{-1, 1\}^m$. We will consider a bijection $\varrho_\kappa: \{-1, 1\}^m \rightarrow [2^m]$ as an ordering of vertices in $\{-1, 1\}^m$ (so x is the $\varrho_\kappa(x)$ -th vertex in the ordering) such that the following property holds: For every $i \in [2^m - 1]$, the degree of $\varrho_\kappa^{-1}(i)$ has the largest degree among all vertices in the subgraph of $\mathcal{H}^{[\kappa]}$ induced by $\{\varrho_\kappa^{-1}(j) : j \geq i\}$. Such a bijection exists, e.g., by keeping deleting vertices one by one and each time deleting the one with the largest degree in the remaining graph (with tie breaking done arbitrarily). We fix such a bijection ϱ_κ and use it to orient edges in $G^{[\kappa]}$ as follows: For each undirected $\{x, x^{(i)}\} \in \mathcal{H}^{[\kappa]}$, we orient it from x to $x^{(i)}$ if $\varrho_\kappa(x) < \varrho_\kappa(x^{(i)})$ and orient it from $x^{(i)}$ to x otherwise. As a result, every $(x, x^{(i)}) \in G^{[\kappa]}$ has $\varrho_\kappa(x) < \varrho_\kappa(x^{(i)})$.

The above orientation will effectively streamline an argument from [KMS18, Section 6]. We record the property needed later for the orientation $G^{[\kappa]}$ of $\mathcal{H}^{[\kappa]}$:

Lemma 3.4. *Let U be a set of vertices in $\{-1, 1\}^m$ and let $v \in \{-1, 1\}^m \setminus U$. If the out-degree of every vertex $u \in U$ in $G^{[\kappa]}$ is bounded from above by a positive integer g , then the number of directed edges (u, v) from a vertex $u \in U$ to v in $G^{[\kappa]}$ is also at most g .*

Proof: Consider the vertex with the smallest $\varrho_\kappa(\cdot)$ value in $U \cup \{v\}$. If it is v , then every undirected $\{u, v\}$ with $u \in U$, if any, in $\mathcal{H}^{[\kappa]}$ is oriented as (v, u) in $G^{[\kappa]}$. So the number we care about is 0.

Otherwise, let $u \in U$ be the vertex with the smallest $\varrho_\kappa(\cdot)$ value among $U \cup \{v\}$. Then at the time when u is picked, all vertices $U \cup \{v\}$ remain in current undirected subgraph of $\mathcal{H}^{[\kappa]}$, denoted by H . At this moment, the degree of u in H is exactly its out-degree in $G^{[\kappa]}$, which by assumption is at most g . On the other hand, by the choice of u , v has degree at most g in H . Since the whole set U remains in H , the number of undirected edges $\{u, v\}$, $u \in U$, in $\mathcal{H}^{[\kappa]}$ is at most g . Even if all of them are oriented towards v in $G^{[\kappa]}$, the number we care about in the lemma is at most g . \blacksquare

With $G^{[z]}$, $G^{[u]}$ and $G^{[\kappa]}$ ready, we finally define G to be the union of these graphs, which is an orientation of \mathcal{H} over $\{0, 1\}^m$. Applying the robust Pisier inequality on f , G and $s = 1$, we have

$$\frac{d_{\text{TV}}(\ell, \mathcal{U})}{\log n} \lesssim \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \{-1, 1\}^m} \left[\left\| \sum_{\substack{i \in [m] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G}} \mathbf{y}_i \mathbf{x}_i L_i f(\mathbf{x}) \right\| \right] \leq \mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^m} \left[\sqrt{\sum_{\substack{i \in [m] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G}} (L_i f(\mathbf{x}))^2} \right]. \quad (30)$$

The second inequality is Khintchine's, which implies that for any vector $a \in \mathbb{R}^m$,

$$\mathbf{E}_{\mathbf{y} \sim \{-1, 1\}^m} \left[\left\| \sum_{i \in [m]} \mathbf{y}_i a_i \right\| \right] \leq \sqrt{\sum_{i \in [m]} a_i^2}.$$

Letting G' be the directed graph that contains the union of edges in $G^{[u]}$ and $G^{[\kappa]}$, $\kappa \in \mathbb{Z}_{\geq 0}$, but not those in $G^{[z]}$, we can continue the inequality above to have

$$\mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^m} \left[\sqrt{\sum_{\substack{i \in [m] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G}} (L_i f(\mathbf{x}))^2} \right] = \mathbf{E}_{\mathbf{x} \sim \ell} \left[\sqrt{\sum_{\substack{i \in [m] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G'}} \left(\frac{L_i f(\mathbf{x})}{1 + f(\mathbf{x})} \right)^2} \right] = \mathbf{E}_{\mathbf{x} \sim \ell} \left[\sqrt{\sum_{\substack{i \in [m] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G'}} \left(\frac{L_i \ell(\mathbf{x})}{\ell(\mathbf{x})} \right)^2} \right].$$

For the first equation we note that zero edges do not contribute anything and utilize importance sampling, by noting that $\ell(x) = (1 + f(x))/2^m$. Also note we never run into a situation of 0/0 in the second expectation because if $(x, x^{(i)}) \in G'$ has $\ell(x) = 0$, then either $\ell(x^{(i)}) = 0$ and it is a zero edge that should have been excluded from G' , or $\ell(x^{(i)}) > 0$ and $\{x, x^{(i)}\}$ is uneven. Then by the construction of $G^{[u]}$ we have $(x^{(i)}, x) \in G$ instead of $(x, x^{(i)})$.

The next lemma connects the sum for each $x \in \{-1, 1\}^m$ in the last expectation with its out-degrees in the directed graphs $G^{[u]}$ and $G^{[\kappa]}$ constructed.

Lemma 3.5. *For every $x \in \{-1, 1\}^m$,*

$$\sum_{\substack{i \in [m] \\ (\mathbf{x}, \mathbf{x}^{(i)}) \in G'}} \left(\frac{\ell(\mathbf{x}) - \ell(\mathbf{x}^{(i)})}{\ell(\mathbf{x})} \right)^2 \leq \text{outdeg}(x, G^{[u]}) + \sum_{\kappa \geq 1} 2^{-2\kappa+6} \cdot \text{outdeg}(x, G^{[\kappa]}).$$

Proof: First note that each edge $(x, x^{(i)}) \in G'$ is nonzero and either lies in $G^{[u]}$ or $G^{[\kappa]}$ for some $\kappa \geq 1$. If $(x, x^{(i)})$ is in $G^{[u]}$, then by the way we orient edges in $G^{[u]}$, we have $\ell(x) > \ell(x^{(i)})$ and this implies that the contribution of each such edge to the sum is at most 1.

Next assume that $(x, x^{(i)}) \in G^{[\kappa]}$ for some $\kappa \geq 1$. Since $\{x, x^{(i)}\}$ is even, we have $\ell(x), \ell(x^{(i)}) > 0$ (otherwise it is either zero or uneven). Using $w(\{x, x^{(i)}\}) < 2/3$ (otherwise it is uneven), we have

$$\frac{\max\{\ell(x), \ell(x^{(i)})\}}{\min\{\ell(x), \ell(x^{(i)})\}} \leq 3.$$

As a result, we have

$$\frac{|\ell(x) - \ell(x^{(i)})|}{\ell(x)} \leq w(\{x, x^{(i)}\}) \cdot \frac{\max\{\ell(x), \ell(x^{(i)})\}}{\min\{\ell(x), \ell(x^{(i)})\}} \leq 3w(\{x, x^{(i)}\}) \leq 2^{-\kappa+3}.$$

So the contribution of each such edge is at most $2^{-2\kappa+6}$ and we thus obtain the desired bound. \blacksquare

We are now ready to prove the main lemma of this subsection:

Lemma 3.6. *Letting $\beta = d_{\text{TV}}(\ell, \mathcal{U})$, one of the following two conditions must hold:*

- *Either the directed graph $G^{[u]}$ of uneven edges satisfies:*

$$\mathbf{E}_{x \sim \ell} \left[\sqrt{\text{outdeg}(\mathbf{x}, G^{[u]})} \right] \gtrsim \frac{\beta}{\log n}$$

- *Or, there exists a $\kappa \in [O(\log(n/\beta))]$ such that the directed graph $G^{(\kappa)}$ satisfies:*

$$\mathbf{E}_{x \sim \ell} \left[\sqrt{\text{outdeg}(\mathbf{x}, G^{(\kappa)})} \right] \gtrsim \frac{2^\kappa \cdot \beta}{\log n \cdot \log(n/\beta)}.$$

Proof: It follows from Lemma 3.5 that

$$\begin{aligned} \frac{\beta}{\log n} &\lesssim \mathbf{E}_{x \sim \ell} \left[\sqrt{\text{outdeg}(\mathbf{x}, G^{[u]}) + \sum_{\kappa \geq 1} 2^{-2\kappa+2} \cdot \text{outdeg}(\mathbf{x}, G^{[\kappa]})} \right] \\ &\leq \mathbf{E}_{x \sim \ell} \left[\sqrt{\text{outdeg}(\mathbf{x}, G^{[u]})} \right] + \sum_{\kappa=1}^{O(\log(n/\beta))} 2^{-\kappa+1} \cdot \mathbf{E}_{x \sim \ell} \left[\sqrt{\text{outdeg}(\mathbf{x}, G^{[\kappa]})} \right] + o\left(\frac{\beta}{\log n}\right). \end{aligned} \quad (31)$$

In (31) we used the concavity of $\sqrt{\cdot}$ as well as the fact that the degrees are always bounded by n . Lemma 3.6 then follows from (31). \blacksquare

3.4 Bucketing

We now start the proof of Lemma 3.1. For each t -subset T of $[n]$, let $\alpha(T) = d_{\text{TV}}(p_{\overline{T}}, \mathcal{U})$ and thus, $\alpha = \mathbf{E}_{\mathbf{T} \sim \mathcal{S}(t)}[\alpha(T)]$. For each T , we partition undirected edges in $\mathcal{H}(T)$ into $\mathcal{H}^{[z]}(T)$ (zero edges), $\mathcal{H}^{[u]}(T)$ (uneven edges), and $\mathcal{H}^{[\kappa]}(T)$ (even edges at scale $\kappa \geq 1$). We orient these edges to obtain directed graphs $G^{[u]}(T)$ and $G^{[\kappa]}(T)$. We apply Lemma 3.6 on $p_{\overline{T}}$ to conclude that one of the two conditions holds for either $G^{[u]}(T)$ or one of the graphs $G^{[\kappa]}(T)$, $\kappa \in [O(\log(n/\alpha(T)))]$.

Since $\alpha(T) \in [0, 1]$, there exists a $\zeta > 0$ such that with probability at least ζ over $\mathbf{T} \sim \mathcal{S}(t)$,

$$\alpha(\mathbf{T}) \gtrsim \frac{\alpha}{\zeta \log(1/\alpha)}.$$

Therefore, via another bucketing argument and Lemma 3.6, there exist two cases:

- **Case 1:** With probability at least $\zeta/2$ over the draw of $\mathbf{T} \sim \mathcal{S}(t)$, we have that the directed graph $G^{[u]}(\mathbf{T})$ of uneven edges of $p_{\overline{\mathbf{T}}}$ over $\{-1, 1\}^{\overline{\mathbf{T}}}$ satisfies

$$\mathbf{E}_{x \sim p_{\overline{\mathbf{T}}}} \left[\sqrt{\text{outdeg}(x, G^{[u]}(\mathbf{T}))} \right] \gtrsim \frac{\alpha}{\zeta \log n \log(1/\alpha)}.$$

Since the out-degree is always between 0 and n , there exist two parameters $d \in [n]$ and $\xi > 0$ such that with probability $\zeta/(2 \log n)$ over the draw of $\mathbf{T} \sim \mathcal{S}(t)$, we have

$$\mathbf{Pr}_{x \sim p_{\overline{\mathbf{T}}}} \left[d \leq \text{outdeg}(x, G^{[u]}(\mathbf{T})) \leq 2d \right] \geq \xi$$

and ξ satisfies

$$\sqrt{d} \cdot \xi \gtrsim \frac{\alpha}{\zeta \log^2 n \log(1/\alpha)}. \quad (32)$$

- **Case 2:** There exists a parameter $\kappa \in [O(\log(n/\alpha))]$ (using $\zeta \leq 1$) such that with probability at least $\zeta/(2 \cdot O(\log(n/\alpha)))$ over the draw of $\mathbf{T} \sim \mathcal{S}(t)$, the directed graph $G^{[\kappa]}(\mathbf{T})$ of even edges at scale κ of $p_{\overline{\mathbf{T}}}$ over $\{-1, 1\}^{\overline{\mathbf{T}}}$ satisfies

$$\mathbf{E}_{x \sim p_{\overline{\mathbf{T}}}} \left[\sqrt{\text{outdeg}(x, G^{[\kappa]}(\mathbf{T}))} \right] \gtrsim \frac{\alpha \cdot 2^\kappa}{\zeta \log n \log(n/\alpha) \log(1/\alpha)}.$$

By a bucketing argument again, there exist parameters $d \in [n]$ and $\xi > 0$ such that with probability at least $\zeta/(2 \cdot \log n \cdot O(\log(n/\alpha)))$ over the draw of $\mathbf{T} \sim \mathcal{S}(t)$, we have

$$\mathbf{Pr}_{x \sim p_{\overline{\mathbf{T}}}} \left[d \leq \text{outdeg}(x, G^{[u]}(\mathbf{T})) \leq 2d \right] \geq \xi$$

and ξ satisfies

$$\sqrt{d} \cdot \xi \gtrsim \frac{\alpha \cdot 2^\kappa}{\zeta \log^2 n \log(n/\alpha) \log(1/\alpha)}. \quad (33)$$

3.5 From Directed Graphs to Mean Vectors

In this section, we will show the crucial connection between analyzing the family of graphs defined in Section 3.3 and 3.4 and mean vectors of restrictions of the distribution. Consider a fixed distribution p supported on $\{-1, 1\}^n$ and let $t \in [n-1]$. We consider the family of directed graphs $G^{[u]}(T)$ and $G^{[\kappa]}(T)$, $\kappa \geq 1$, for each t -subset T of $[n]$.

It will be convenient to represent directed edges of these directed graphs as $(y, i) \in \{-1, 1\}^{\overline{\mathbf{T}}} \times \overline{\mathbf{T}}$: we say (y, i) is in a graph if it is the case for $(y, y^{(i)})$. Let $\pi = (\pi(1), \dots, \pi(t+1))$ be an (ordered) sequence of $t+1$ distinct indices from $[n]$. We use $S(\pi)$ to denote the corresponding $(t+1)$ -subset $\{\pi(1), \dots, \pi(t+1)\}$. Given π and $y \in \{-1, 1\}^n$, we define a restriction $\rho(\pi, y) \in \{-1, 1, *\}^n$ as

$$\rho(\pi, y)_i = \begin{cases} * & i = \pi(j) \text{ for some } j \in [t+1] \\ y_i & \text{otherwise} \end{cases}.$$

We will also consider sequences $\tau = (\tau(1), \dots, \tau(t))$ of t distinct indices from $[n]$; its corresponding set $S(\tau)$ and the restriction $\rho(\tau, y)$ given $y \in \{-1, 1\}^n$ are defined similarly.

We consider a slightly different but equivalent way of drawing ρ from $\mathcal{D}(t+1, p)$ and ρ' from $\mathcal{D}(t, p)$ which we will use to analyze Case 1 and Case 2 as specified in Section 3.4. We consider sampling $\rho \sim \mathcal{D}(t+1, p)$ according to the following procedure:

1. First, sample a sequence of $t + 1$ random indices $\boldsymbol{\pi} = (\boldsymbol{\pi}(1), \dots, \boldsymbol{\pi}(t + 1))$ uniformly from $[n]$ without replacements (so the set $S(\boldsymbol{\pi})$ can be viewed equivalently as drawn from $\mathcal{S}(t + 1)$).
2. Then, sample $\mathbf{y} \sim p$.
3. Finally, set $\boldsymbol{\rho} = \rho(\boldsymbol{\pi}, \mathbf{y})$.

Similarly we consider sampling $\boldsymbol{\rho}' \sim \mathcal{D}(t, p)$ according to the following procedure:

1. First, sample a sequence of t random indices $\boldsymbol{\tau} = (\boldsymbol{\tau}(1), \dots, \boldsymbol{\tau}(t))$ uniformly from $[n]$ without replacements (so the set $S(\boldsymbol{\tau})$ can be viewed equivalently as drawn from $\mathcal{S}(t)$).
2. Then, sample $\mathbf{y} \sim p$.
3. Finally, set $\boldsymbol{\rho}' = \rho(\boldsymbol{\tau}, \mathbf{y})$.

A useful observation is that for each $i \in [t + 1]$, the distribution of $(\boldsymbol{\pi}_{-i}, \mathbf{y})$ is the same as that of $(\boldsymbol{\tau}, \mathbf{y})$, where $\boldsymbol{\pi}_{-i}$ denotes the t -sequence obtained from $\boldsymbol{\pi}$ after removing its i th entry. As a result, $\rho(\boldsymbol{\pi}_{-i}, \mathbf{y})$ is distributed according to $\mathcal{D}(t, p)$.

Next we state the lemma that gives the connection between graphs and mean vectors.

Lemma 3.7. *Let $\boldsymbol{\pi}$ be a $(t + 1)$ -sequence of distinct elements in $[n]$ and $\mathbf{y} \in \{-1, 1\}^n$. Then for every $i \in [t + 1]$, we have*

$$\begin{aligned} \left| \mu(p_{|\rho(\boldsymbol{\pi}, \mathbf{y})})_{\boldsymbol{\pi}(i)} \right| &\geq \frac{1}{3} \cdot \mathbf{1} \left\{ (y_{S(\boldsymbol{\pi}_{-i})}, \boldsymbol{\pi}(i)) \in G^{[u]}(S(\boldsymbol{\pi}_{-i})) \right\} \\ &\quad + \sum_{k \geq 1} 2^{-k-1} \cdot \mathbf{1} \left\{ (y_{S(\boldsymbol{\pi}_{-i})}, \boldsymbol{\pi}(i)) \in G^{[k]}(S(\boldsymbol{\pi}_{-i})) \right\}. \end{aligned} \quad (34)$$

Proof: First recall that $G^{[u]}(S(\boldsymbol{\pi}_{-i}))$ and $G^{[k]}(S(\boldsymbol{\pi}_{-i}))$ are orientations of disjoint undirected edges. So the right-hand side of (34) is non-zero for at most one value.

Let $\ell = p_{|\rho(\boldsymbol{\pi}, \mathbf{y})}$, $T = S(\boldsymbol{\pi}_{-i})$, $z = y_T$ and $z' = z^{(\boldsymbol{\pi}(i))}$. Writing

$$a = \mathbf{Pr}_{\mathbf{x} \sim p} [\mathbf{x}_T = z] \quad \text{and} \quad a' = \mathbf{Pr}_{\mathbf{x} \sim p} [\mathbf{x}_T = z'],$$

we have $|\mu(\ell)_{\boldsymbol{\pi}(i)}| = |a - a'| / (a + a')$. The weight $w(\{z, z'\})$, defined as $|a - a'| / \max\{a, a'\}$, is at most $2 \cdot |\mu(\ell)_{\boldsymbol{\pi}(i)}|$. If $(z, \boldsymbol{\pi}(i)) \in G^{[u]}(T)$ is uneven, then the weight is at least $2/3$ and thus, $|\mu(\ell)_{\boldsymbol{\pi}(i)}| \geq 1/3$. If $(z, \boldsymbol{\pi}(i)) \in G^{[k]}(T)$ for some k , then the weight is at least 2^{-k} and $|\mu(\ell)_{\boldsymbol{\pi}(i)}| \geq 2^{-k-1}$. ■

3.6 Case 1

In this section, we prove Lemma 3.1 assuming that we are in the first case outlined in Section 3.4. As per Section 3.5 and the first case described in Section 3.4, we consider a distribution p supported on $\{-1, 1\}^n$, a parameter $t \in [n - 1]$, and we focus on the family of directed graphs $G^{[u]}(T)$, one for each size- t subset T of $[n]$. We assume that there are parameters ζ', ξ , and $d \geq 1$ such that with probability at least ζ' over the draw of $\mathbf{T} \sim \mathcal{S}(t)$,

$$\mathbf{Pr}_{\mathbf{x} \sim p_{\mathbf{T}}} [d \leq \text{outdeg}(\mathbf{x}, G^{(u)}(\mathbf{T})) \leq 2d] \geq \xi. \quad (35)$$

Notice that we will set $\zeta' = \zeta / (2 \log n)$, so that (32) implies

$$\sqrt{d} \cdot \xi \gtrsim \frac{\alpha}{\zeta' \log^3 n \log(1/\alpha)}. \quad (36)$$

In this case, we will show that (17) holds. The following definition and lemma will be useful:

Definition 3.8. Let τ be a t -sequence of distinct indices from $[n]$ and let $y \in \{-1, 1\}^n$. We say the pair (τ, y) is t -contributing if the restricted distribution $p_{|\rho(\tau, y)}$ satisfies

$$\left\| \mu(p_{|\rho(\tau, y)}) \right\|_2 \geq \frac{\sqrt{d}}{20}, \quad (37)$$

and we say (τ, y) is $(t+1)$ -contributing otherwise.

Note that the definition implies that a pair is either t -contributing or $(t+1)$ -contributing.

Lemma 3.9. Let π be a $(t+1)$ -sequence of distinct indices from $[n]$, and let $y \in \{-1, 1\}^n$. If there are $d+1$ distinct indices $i_1, \dots, i_{d+1} \in [t+1]$ such that

$$\left(\overline{y_{S(\pi_{-i_k})}}, \pi(i_k) \right) \in G^{[u]}(S(\pi_{-i_k})) \quad (38)$$

for all $k \in [d+1]$, then (π_{-i_k}, y) is a t -contributing pair for all $k \in [d+1]$.

Proof: We prove the lemma for $k=1$. For convenience we refer to i_1 as i and i_2, \dots, i_{d+1} as j_1, \dots, j_d . We show that for every $j = j_1, \dots, j_d$, $|\mu(p_{|\rho(\pi_{-i}, y)})_{\pi(j)}| \geq 1/20$. The lemma then follows.

To simplify notation, we $R = \overline{S(\pi)} \cup \{\pi(i), \pi(j)\}$. Let $z = y_R \in \{-1, 1\}^R$. Let

$$\begin{aligned} a_1 &= \Pr_{x \sim p} [\mathbf{x}_R = z] & a_2 &= \Pr_{x \sim p} [\mathbf{x}_R = z^{(\pi(i))}] \\ a_3 &= \Pr_{x \sim p} [\mathbf{x}_R = z^{(\pi(j))}] & a_4 &= \Pr_{x \sim p} [\mathbf{x}_R = z^{(\{\pi(i), \pi(j)\})}] \end{aligned}$$

Notice that since $(\overline{y_{S(\pi_{-i})}}, \pi(i))$ and $(\overline{y_{S(\pi_{-j})}}, \pi(j))$ are uneven edges, we have

$$a_1 + a_3 \geq 3(a_2 + a_4) \quad \text{and} \quad a_1 + a_2 \geq 3(a_3 + a_4). \quad (39)$$

We use these two inequalities to show that

$$\left| \mu(p_{|\rho(\pi_{-i}, y)})_{\pi(j)} \right| = \frac{|a_1 - a_3|}{a_1 + a_3} \geq \frac{1}{20}.$$

Suppose first that $a_3 \leq 9a_1/10$, then

$$\frac{|a_1 - a_3|}{a_1 + a_3} \geq \frac{a_1/10}{a_1 + a_3} \geq \frac{1}{10} \cdot \frac{1}{2} \geq \frac{1}{20}.$$

If $a_3 \geq 9a_1/10$, then by the second equality in (39), $a_2 \geq 17a_1/10 + 3a_4$. Furthermore, substituting into the first inequality of (39), $a_1 + a_3 \geq 3(17a_1/10 + 3a_4) + 3a_4$, which implies $a_3 \geq 4a_1$, so

$$\frac{|a_1 - a_3|}{a_1 + a_3} \geq \frac{3a_3/4}{a_1 + a_3} \geq \frac{3}{4} \cdot \frac{1}{2} \geq \frac{1}{20}.$$

This finishes the proof of the lemma. ■

Now we start to lower bound the expectation of $\|\mu(p_{|\rho})\|_2$ as $\rho \sim \mathcal{D}(t+1, p)$, or equivalently as ρ is drawn as $\rho(\boldsymbol{\pi}, \mathbf{y})$. For this purpose we introduce an indicator random variable \mathbf{X}_i for each $i \in [t+1]$: \mathbf{X}_i is 1 if the following event holds:

$$\left(\overline{\mathbf{y}_{S(\pi_{-i})}}, \boldsymbol{\pi}(i) \right) \in G^{[u]}(S(\pi_{-i})) \text{ and } (\pi_{-i}, \mathbf{y}) \text{ is } (t+1)\text{-contributing} \quad (40)$$

On the one hand, it follows from Lemma 3.7 (and the first part of the event above) that

$$\|\mu(p_{|\rho(\pi, \mathbf{y})})\|_2 \gtrsim \sqrt{\mathbf{X}_1 + \cdots + \mathbf{X}_{t+1}}.$$

On the other hand, it follows from Lemma 3.9 and the second part of (40) that $\mathbf{X}_1 + \cdots + \mathbf{X}_{t+1}$ is at most d with probability 1. As a result, we have

$$\mathbf{E}_{\pi, \mathbf{y}} \left[\|\mu(p_{|\rho(\pi, \mathbf{y})})\|_2 \right] \gtrsim \mathbf{E}_{\pi, \mathbf{y}} \left[\mathbf{X}_1 + \cdots + \mathbf{X}_{t+1} \right] / \sqrt{d}$$

This simplifies the task now to bound the probability of $\mathbf{X}_i = 1$. We claim that for each $i \in [t+1]$,

$$\Pr_{\pi, \mathbf{y}} [\mathbf{X}_i = 1] \geq \left(\zeta' \cdot \xi - \Pr_{\pi, \mathbf{y}} [(\pi_{-i}, \mathbf{y}) \text{ is } t\text{-contributing}] \right) \cdot \frac{d}{n}, \quad (41)$$

To see this is the case, we consider drawing π and \mathbf{y} by drawing \mathbf{y} and π_{-i} first and then $\pi(i)$. We consider the following event F over \mathbf{y} and π_{-i} :

Event \mathbf{F} : $S(\pi_{-i})$ as \mathbf{T} and $\mathbf{y}_{S(\pi_{-i})}$ as \mathbf{x} satisfy (35) and (π_{-i}, \mathbf{y}) is $(t+1)$ -contributing.

Note that it follows from our assumption at the beginning of this section that the first part of \mathbf{F} occurs with probability at least $\zeta' \cdot \xi$. It follows that the probability of \mathbf{F} is at least

$$\zeta' \cdot \xi - \Pr_{\pi, \mathbf{y}} [(\pi_{-i}, \mathbf{y}) \text{ is } t\text{-contributing}].$$

Finally, conditioning on π_{-i} and \mathbf{y} satisfying \mathbf{F} , $\pi(i)$ (together with π_{-i} and \mathbf{y}) leads to $\mathbf{X}_i = 1$ if it is one of the at least d edges of $\mathbf{y}_{S(\pi_{-i})}$ in $G^{[u]}(S(\pi_{-i}))$. This occurs with probability at least d/n .

Continuing from (41), we note that the probability of (π_{-i}, \mathbf{y}) being t -contributing is the same as (τ, \mathbf{y}) being t -contributing, where τ is drawn uniformly from all t -sequences of distinct indices. Combining everything,

$$\mathbf{E}_{\pi, \mathbf{y}} \left[\|\mu(p_{|\rho(\pi, \mathbf{y})})\|_2 \right] \gtrsim \frac{1}{\sqrt{d}} \cdot (t+1) \cdot \left(\zeta' \cdot \xi - \Pr_{\tau, \mathbf{y}} [(\tau, \mathbf{y}) \text{ is } t\text{-contributing}] \right) \cdot \frac{d}{n}. \quad (42)$$

Thus, either the probability of (τ, \mathbf{y}) being t -contributing is at least $\zeta' \cdot \xi/2$, in which case we have

$$\mathbf{E}_{\rho \sim \mathcal{D}(t, p)} \left[\|\mu(p_{|\rho})\|_2 \right] = \mathbf{E}_{\tau, \mathbf{y}} \left[\|\mu(p_{|\rho(\tau, \mathbf{y})})\|_2 \right] \gtrsim \zeta' \cdot \xi \cdot \sqrt{d} \gtrsim \frac{\alpha}{\log^3 n \log(1/\alpha)},$$

using (36) for the last inequality; or (42) is lower bounded by

$$\frac{1}{\sqrt{d}} \cdot (t+1) \cdot \frac{\zeta' \cdot \xi}{2} \cdot \frac{d}{n} \gtrsim \frac{t}{n} \cdot \frac{\alpha}{\log^3 n \log(1/\alpha)},$$

which shows (17) holds since $t \leq n-1$.

Remark 3.10. Note that if $dt/n \geq 1$, one may set the right hand side of (37) to be $\sqrt{dt/n}/20$ instead of $\sqrt{d}/20$, and this would result in a stronger lower bound in (17) where we can replace the t/n by $\sqrt{t/n}$. However, we do not have any control on the parameter d , which could be 1 in the worst case.

3.7 Case 2

In this section, we prove Lemma 3.1 assuming that we are in the second case outlined in Section 3.4. As per Section 3.5 and the second case of Section 3.4, we consider a distribution p over $\{-1, 1\}^n$, a parameter $t \in [n-1]$, and a parameter $\kappa \in [O(\log(n/\alpha))]$; we focus on the family of directed graphs $G^{[\kappa]}(T)$ for each t -subset T of $[n]$.

From Case 2 of Section 3.4, we assume there are parameters ζ' , ξ , and d such that with probability at least ζ' over the draw of $\mathbf{T} \sim \mathcal{S}(t)$, we have

$$\Pr_{\mathbf{x} \sim p_{\overline{\mathbf{T}}}} \left[d \leq \text{outdeg}(\mathbf{x}, G^{[\kappa]}(\mathbf{T})) \leq 2d \right] \geq \xi, \quad (43)$$

where we set $\zeta' = \zeta / (2 \log n \cdot O(\log(n/\alpha)))$. Thus, (33) implies

$$\sqrt{d} \cdot \xi \gtrsim \frac{\alpha \cdot 2^\kappa}{\zeta' \log^3 n \log^2(n/\alpha) \log(1/\alpha)}.$$

The goal is to lower bound the expectation of $\|\mu(p|_\rho)\|_2$ with $\rho \sim \mathcal{D}(t, p)$ or $\rho \sim \mathcal{D}(t+1, p)$. For this, we consider a similar notion of bad pairs to that of Section 3.6, and start with a lemma that will be useful to the analysis.

Definition 3.11. Let $\gamma \geq 1$.⁵ We say a restriction ρ with t stars is t -contributing if

$$\|\mu(p|_\rho)\|_2 \geq \frac{\sqrt{\gamma}}{2^{\kappa+1}},$$

and we say that ρ is $(t+1)$ -contributing otherwise.

Similar to before, a restriction is either t -contributing or $(t+1)$ -contributing.

Lemma 3.12. Let π be a $(t+1)$ -sequence of distinct indices from $[n]$ and let $y \in \{-1, 1\}^n$. Suppose that $i \in [t+1]$ is such that both restrictions $\rho(\pi_{-i}, y)$ and $\rho(\pi_{-i}, y^{(\pi(i))})$ are $(t+1)$ -contributing. Then we have

$$\sum_{j \in [t+1] \setminus \{i\}} \left(\mu(p|_{\rho(\pi, y)})_{\pi(j)} \right)^2 < \frac{\gamma}{2^{2\kappa+2}}.$$

Proof: Let a be the following conditional probability:

$$a = \Pr_{\mathbf{x} \sim p} \left[\mathbf{x}_{\pi(i)} = y_{\pi(i)} \mid \mathbf{x}_{S(\pi)} = y_{S(\pi)} \right],$$

and notice that for all $j \in [t+1] \setminus \{i\}$, we have

$$\mu(p_{\rho(\pi, y)})_{\pi(j)} = a \cdot \mu(p|_{\rho(\pi_{-i}, y)})_{\pi(j)} + (1-a) \cdot \mu(p|_{\rho(\pi_{-i}, y^{(\pi(i))})})_{\pi(j)}.$$

By Jensen's inequality,

$$\sum_{j \in [t+1] \setminus \{i\}} \left(\mu(p|_{\rho(\pi, y)})_{\pi(j)} \right)^2 \leq \sum_{j \in [t+1] \setminus \{i\}} \left(a \cdot \left(\mu(p|_{\rho(\pi_{-i}, y)})_{\pi(j)} \right)^2 + (1-a) \cdot \left(\mu(p|_{\rho(\pi_{-i}, y^{(\pi(i))})})_{\pi(j)} \right)^2 \right)$$

⁵We will set $\gamma = d$ at the end but would like to keep it as a parameter for the discussion in Remark 3.13.

which less than $\gamma/2^{2\kappa+2}$ since $\rho(\pi_{-i}, \mathbf{y})$ and $\rho(\pi_{-i}, \mathbf{y}^{(\pi(i))})$ are both $(t+1)$ -contributing. \blacksquare

We proceed in a similar fashion to Case 1. To bound the expectation of $\|\mu(p|\rho)\|_2$ as $\rho \sim \mathcal{D}(t+1, p)$, or equivalently as ρ is drawn as $\rho(\boldsymbol{\pi}, \mathbf{y})$, we introduce an indicator random variable \mathbf{X}_i for each $i \in [t+1]$: \mathbf{X}_i is 1 if the following event holds:

$$\left(\overline{\mathbf{y}_{S(\boldsymbol{\pi}_{-i})}}, \boldsymbol{\pi}(i)\right) \in G^{[\kappa]}(S(\boldsymbol{\pi}_{-i})) \text{ and both } \rho(\boldsymbol{\pi}_{-i}, \mathbf{y}) \text{ and } \rho(\boldsymbol{\pi}_{-i}, \mathbf{y}^{(\pi(i))}) \text{ are } (t+1)\text{-contributing} \quad (44)$$

First, it follows from Lemma 3.7 (and the first part of the event above) that

$$\left\|\mu(p|\rho(\boldsymbol{\pi}, \mathbf{y}))\right\|_2 \geq \frac{1}{2^{\kappa+1}} \cdot \sqrt{\mathbf{X}_1 + \cdots + \mathbf{X}_{t+1}}.$$

But this, along with Lemma 3.12 and the second part of (44), implies that $\mathbf{X}_1 + \cdots + \mathbf{X}_{t+1}$ is at most γ with probability 1. As a result, we have

$$\mathbf{E}_{\boldsymbol{\pi}, \mathbf{y}} \left[\left\|\mu(p|\rho(\boldsymbol{\pi}, \mathbf{y}))\right\|_2 \right] \gtrsim \frac{1}{2^{\kappa+1}} \cdot \mathbf{E}_{\boldsymbol{\pi}, \mathbf{y}} [\mathbf{X}_1 + \cdots + \mathbf{X}_{t+1}] \cdot \frac{1}{\sqrt{\gamma}}$$

This reduces the task now to bounding the probability of $\mathbf{X}_i = 1$.

We start with some notation. Let T be a t -subset of $[n]$ and let $z \in \{-1, 1\}^{\bar{T}}$. We use $\rho(z)$ to denote the restriction $\rho \in \{-1, 1, *\}^n$ with $\rho_i = z_i$ for all $i \in \bar{T}$ and $\rho_i = *$ for all $i \in T$. To lower bound the probability of $\mathbf{X}_i = 1$, we define the following two disjoint subsets of $\{-1, 1\}^{\bar{T}}$ for each t -subset T :

$$\begin{aligned} A_T &= \left\{z \in \{-1, 1\}^{\bar{T}} : d \leq \text{outdeg}(z, G^{[\kappa]}(T)) \leq 2d \text{ and } \rho(z) \text{ is } (t+1)\text{-contributing}\right\} \\ B_T &= \left\{w \in \{-1, 1\}^{\bar{T}} : \rho(w) \text{ is } t\text{-contributing}\right\}. \end{aligned}$$

Then the probability of $\mathbf{X}_i = 1$ (i.e., the event described in (44)) is at least the probability of the following event \mathbf{E} , where we first draw a t -subset \mathbf{T} , then \mathbf{z} from $p_{\bar{\mathbf{T}}}$, and finally draw \mathbf{i} from $\bar{\mathbf{T}}$:

$$\text{Event } \mathbf{E}: (\mathbf{z}, \mathbf{i}) \in G^{[\kappa]}(\mathbf{T}), \mathbf{z} \in A_{\mathbf{T}} \text{ and } \mathbf{z}^{(\mathbf{i})} \notin B_{\mathbf{T}}$$

We write the probability of \mathbf{E} as

$$\begin{aligned} \mathbf{Pr}_{\mathbf{T}, \mathbf{z}, \mathbf{i}} [\mathbf{E}] &= \mathbf{Pr}_{\mathbf{T}, \mathbf{z}, \mathbf{i}} \left[(\mathbf{z}, \mathbf{i}) \in G^{[\kappa]}(\mathbf{T}) \wedge \mathbf{z} \in A_{\mathbf{T}} \right] \\ &\quad - \mathbf{Pr}_{\mathbf{T}, \mathbf{z}, \mathbf{i}} \left[(\mathbf{z}, \mathbf{i}) \in G^{[\kappa]}(\mathbf{T}) \wedge \mathbf{z} \in A_{\mathbf{T}} \wedge \mathbf{z}^{(\mathbf{i})} \in B_{\mathbf{T}} \right]. \end{aligned} \quad (45)$$

The first probability on the right hand side of (45) is at least

$$\left(\zeta' \cdot \xi - \mathbf{Pr}_{\rho \sim \mathcal{D}(t, p)} [\rho \text{ is } t\text{-contributing}] \right) \cdot \frac{d}{n-t}.$$

To see this we first draw \mathbf{T} and \mathbf{z} and impose the condition that $\mathbf{z} \in A_{\mathbf{T}}$. Similar to the analysis in Case 1, this happens with probability at least

$$\zeta' \cdot \xi - \mathbf{Pr}_{\mathbf{T}, \mathbf{z}} [\mathbf{z} \in B_{\mathbf{T}}] = \zeta' \cdot \xi - \mathbf{Pr}_{\rho \sim \mathcal{D}(t, p)} [\rho \text{ is } t\text{-contributing}].$$

Then we draw \mathbf{i} and the probability we get an outgoing edge in $G^{[\kappa]}(\mathbf{T})$ is at least $d/(n-t)$.

The probability we subtract on the right hand side of (45) can be written as

$$\begin{aligned}
& \frac{1}{\binom{n}{t}} \sum_{T \in \mathcal{P}(t)} \sum_{z \in A_T} \Pr_{\mathbf{x} \sim p_{\bar{T}}} [\mathbf{x} = z] \cdot \sum_{i \in \bar{T}} \frac{1}{|\bar{T}|} \cdot \mathbf{1} \left\{ (z, i) \in G^{[\kappa]}(T) \wedge z^{(i)} \in B_T \right\} \\
& \leq \frac{3}{\binom{n}{t}} \cdot \frac{1}{n-t} \sum_{T \in \mathcal{P}(t)} \sum_{z \in A_T} \sum_{i \in \bar{T}} \Pr_{\mathbf{x} \sim p_{\bar{T}}} [\mathbf{x} = z^{(i)}] \cdot \mathbf{1} \left\{ (z, i) \in G^{[\kappa]}(T) \wedge z^{(i)} \in B_T \right\} \\
& = \frac{3}{\binom{n}{t}} \cdot \frac{1}{n-t} \sum_{T \in \mathcal{P}(t)} \sum_{w \in B_T} \sum_{i \in \bar{T}} \Pr_{\mathbf{x} \sim p_{\bar{T}}} [\mathbf{x} = w] \cdot \mathbf{1} \left\{ (w^{(i)}, i) \in G^{[\kappa]}(T) \wedge w^{(i)} \in A_T \right\} \\
& = \frac{3}{\binom{n}{t}} \cdot \frac{1}{n-t} \sum_{T \in \mathcal{P}(t)} \sum_{w \in B_T} \Pr_{\mathbf{x} \sim p_{\bar{T}}} [\mathbf{x} = w] \cdot \left[\text{number of edges from } A_T \text{ to } w \text{ in } G^{[\kappa]} \right].
\end{aligned}$$

where the first inequality used $\Pr_{\mathbf{x}}[\mathbf{x} = z] \leq 3 \cdot \Pr_{\mathbf{x}}[\mathbf{x} = z^{(i)}]$ because each $\{z, z^{(i)}\}$ is an even edge, and the second equation is just a change of variable. Given that every vertex in A_T has out-degree at most $2d$ in $G^{[\kappa]}$, we can now apply Lemma 3.4 to conclude that the number of edges from A_T to every $w \in B_T$ is at most $2d$. As a result, the probability we subtract can be bounded from above by

$$\frac{3}{\binom{n}{t}} \cdot \frac{1}{n-t} \sum_{T \in \mathcal{P}(t)} \sum_{w \in B_T} \Pr_{\mathbf{x} \sim p_{\bar{T}}} [\mathbf{x} = w] \cdot 2d \leq \frac{6d}{n-t} \cdot \Pr_{\mathbf{T}, \mathbf{z}} [z \in B_{\mathbf{T}}].$$

Therefore, we have (45) that

$$\Pr_{\pi, \mathbf{y}} [\mathbf{X}_i = 1] \geq \Pr_{\mathbf{T}, \mathbf{z}, \mathbf{k}} [\mathbf{E}] \geq \left(\zeta' \cdot \xi - 7 \cdot \Pr_{\rho \sim \mathcal{D}(t, p)} [\rho \text{ is } t\text{-contributing}] \right) \cdot \frac{d}{n-t}.$$

To conclude the proof of Lemma 3.1 for Case 2, we set $\gamma = d$. Then either we have

$$\Pr_{\rho \sim \mathcal{D}(t, p)} [\rho \text{ is } t\text{-contributing}] \geq \frac{\zeta' \cdot \xi}{14},$$

which implies

$$\mathbf{E}_{\rho \sim \mathcal{D}(t, p)} \left[\left\| \mu(p|\rho') \right\|_2 \right] \geq \frac{\zeta' \cdot \xi}{14} \cdot \frac{\sqrt{d}}{2^{\kappa+1}} \gtrsim \frac{\alpha}{\log^3 n \log^2(n/\alpha) \log(1/\alpha)},$$

or we have

$$\mathbf{E}_{\rho \sim \mathcal{D}(t+1, p)} \left[\left\| \mu(p|\rho) \right\|_2 \right] \gtrsim \frac{t+1}{2^\kappa \cdot \sqrt{d}} \cdot \frac{d}{n-t} \cdot \frac{\zeta' \cdot \xi}{14} \gtrsim \frac{t}{n} \cdot \frac{\alpha}{\log^3 n \log^2(n/\alpha) \log(1/\alpha)}.$$

Taking the sum always gives us (17).

Remark 3.13. Similar to Remark 3.10, note that if $dt/n \geq 1$, one may set $\gamma = dt/n$ instead of d , and this would result in a stronger lower bound in (17) where we can replace the t/n by $\sqrt{t/n}$. Similarly to Remark 3.10, we do not have control over the parameter d , which could be 1 in the worst case.

4 Mean Testing

For any $m \in \mathbb{N}$ and any distribution p supported on $\{-1, 1\}^m$, we consider $\mathbf{X} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(q)})$, $\mathbf{Y} = (\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(q)})$, a set of $2q$ i.i.d. samples from p , and let

$$\bar{\mathbf{X}} \stackrel{\text{def}}{=} \frac{1}{q} \sum_{i=1}^q \mathbf{x}^{(i)}, \quad \bar{\mathbf{Y}} \stackrel{\text{def}}{=} \frac{1}{q} \sum_{i=1}^q \mathbf{y}^{(i)}$$

be the empirical means in \mathbb{R}^m . Our core test statistic will take $2q$ i.i.d. samples from p , and compute the expression

$$\mathbf{Z} \stackrel{\text{def}}{=} \langle \bar{\mathbf{X}}, \bar{\mathbf{Y}} \rangle. \quad (46)$$

We will write $\mu(p) = \mathbf{E}_{\mathbf{x} \sim p}[\mathbf{x}] \in [-1, 1]^m$ as the mean vector and $\Sigma(p) \in \mathbb{R}^{m \times m}$ as the symmetric matrix with $\Sigma(p)_{ij} = \mathbf{E}_{\mathbf{x} \sim p}[\mathbf{x}_i \mathbf{x}_j]$ for all $i, j \in [m]$.⁶

Lemma 4.1. *The random variable \mathbf{Z} obtained from two tuples of q samples from p satisfies*

$$\mathbf{E}[\mathbf{Z}] = \langle \mathbf{E}[\bar{\mathbf{X}}], \mathbf{E}[\bar{\mathbf{Y}}] \rangle = \langle \mu(p), \mu(p) \rangle = \|\mu(p)\|_2^2, \quad (47)$$

$$\mathbf{Var}[\mathbf{Z}] \leq \frac{1}{q^2} \|\Sigma(p)\|_F^2 + \frac{4}{q} \|\mu(p)\|_2^2 \|\Sigma(p)\|_F. \quad (48)$$

The proof of (47) follows from the fact that $\bar{\mathbf{X}}$ and $\bar{\mathbf{Y}}$ are independent, and (48) is a computation which appears in Appendix A. We define the *blowup distribution* $\odot(p)$ as the distribution on $\{-1, 1\}^{m^2}$ such that, ordering $[m] \times [m]$ in the lexicographic way, $\odot(p)$ is the distribution of the vector

$$(\mathbf{x}_i \mathbf{x}_j)_{(i,j) \in [m] \times [m]} = (\mathbf{x}_1 \mathbf{x}_1, \mathbf{x}_1 \mathbf{x}_2, \dots, \mathbf{x}_1 \mathbf{x}_m, \mathbf{x}_2 \mathbf{x}_1, \mathbf{x}_2 \mathbf{x}_m, \dots, \mathbf{x}_m \mathbf{x}_m) \quad (49)$$

when $\mathbf{x} \sim p$. For $k \in \mathbb{N}$, we let $\odot^{(k)}(p)$ denote the distribution over $\{-1, 1\}^{m^{2^k}}$ obtained by iterating this process k times, so that in particular $\odot^{(0)}(p) = p$. Note that, for any $k \geq 0$, a sample from $\odot^{(k)}(p)$ can be obtained from a sample from p in time $O(m^{2^k})$.

Fact 4.2. *For any distribution p over $\{-1, 1\}^m$ and $k \in \mathbb{N}$, we have $\|\mu(\odot^{(k+1)}(p))\|_2^2 = \|\Sigma(\odot^{(k)}(p))\|_F^2$.*

Consider the following threshold test:

Algorithm 2 THRESHOLDZTEST(τ, S)

Require: A threshold $\tau > 0$ and a (multi)set $S = \{x_1, \dots, x_q, y_1, \dots, y_q\}$.

- 1: Compute Z from the samples in S according to (46).
 - 2: **if** $Z > \tau$ **then return reject**
 - 3: **return accept**
-

Hereafter, for $\varepsilon \in (0, 1]$, we consider the sequence $(\tau_k)_{k \in \mathbb{Z}_{\geq 0}}$ of real numbers, defined recursively as

$$\tau_k \stackrel{\text{def}}{=} \begin{cases} \frac{\varepsilon^2 n}{2} & \text{if } k = 0 \\ \frac{1}{5000} \cdot q^2 \tau_{k-1}^2 & \text{if } k \geq 1 \end{cases} \quad (50)$$

⁶In particular, note that due to the diagonal terms we have $\|\Sigma(p)\|_F \geq \sqrt{m}$ for all p .

In particular, writing $a \stackrel{\text{def}}{=} 1/5000$, we have

$$\tau_k = \frac{1}{aq^2} \cdot \left(\frac{aq^2 \varepsilon^2 n}{2} \right)^{2^k} \quad (51)$$

for $k \geq 0$. We are now ready to state the main testing algorithm. The algorithm takes as input a distribution p which is supported on $\{-1, 1\}^n$, and is written with two unspecified parameters, k_0 and q . The parameter k_0 denotes the number of rounds and q denotes the sample complexity.

Algorithm 3 MEANTESTER(p, ε)

Require: A distribution p supported on $\{-1, 1\}^n$.

- 1: For $T = O(\log k_0)$, draw T sets $\mathbf{S}_1, \dots, \mathbf{S}_T$ of $2q'$ i.i.d random samples from p , where $q' = q/(2T)$.
 - 2: **for all** $0 \leq k \leq k_0$ **do**
 - 3: Set τ_k as in (50), and $N_k \leftarrow 0$.
 - 4: **for all** $1 \leq t \leq T$ **do**
 - 5: Convert the $2q'$ samples from \mathbf{S}_t to a (multi)set $\mathbf{S}_t^{(k)}$ of samples from $\odot^{(k)}(p)$ as in (49).
 - 6: **if** THRESHOLDZTEST($\tau_k, \mathbf{S}_t^{(k)}$) returns reject **then** $N_k \leftarrow N_k + 1$.
 - 7: **if** $N_k > T/2$ **then return** reject ▷ Reject if the majority vote is reject
 - 8: **return** accept ▷ All $k_0 + 1$ tests were successful
-

Theorem 4.3. *Fix any $k_0 \in \mathbb{N}$. There exists an algorithm (Algorithm 3) which, given sample access to an arbitrary distribution p on $\{-1, 1\}^n$ and a parameter $\varepsilon \in (0, 1]$, has the following behavior:*

- *If p is the uniform distribution, the algorithm outputs accept with probability at least $2/3$;*
- *If p satisfies $\|\mu(p)\|_2 \geq \varepsilon\sqrt{n}$, the algorithm outputs reject with probability at least $2/3$.*

These guarantees hold as long as

$$q \gtrsim \max \left\{ \frac{1}{\varepsilon^2 \sqrt{n}}, \left(\frac{1}{\varepsilon^2} \right)^{\frac{2^{k_0+1}}{2^{k_0+2}-2}} \right\} \cdot \log \max(2, k_0).$$

The algorithm runs in time $O(q \cdot n^{2^{k_0}})$.

In particular, by setting $k_0 = \log \log n$, we obtain an algorithm for distinguishing the uniform distribution from a distribution p on $\{-1, 1\}^n$ with $\|\mu(p)\|_2 \geq \varepsilon\sqrt{n}$ which runs in time $n^{\Theta(\log n)}$ and has sample complexity

$$O \left(\max \left\{ \frac{1}{\varepsilon^2 \sqrt{n}}, \frac{1}{\varepsilon} \right\} \cdot \log \log \log n \right).$$

Proof of Theorem 4.3: The proof will proceed as follows: we first show that, when p is the uniform distribution \mathcal{U} (the completeness case), then all $k_0 + 1$ tests, when run on Lines 4–7 of Algorithm 3 return **accept** with high probability. To do so, notice that Line 6 of Algorithm 3 considers samples from $\odot^{(k)}(\mathcal{U})$. Hence, we analyze the mean and variance of each $\odot^{(k)}(\mathcal{U})$, and apply Chebyshev's inequality to show that each call to Line 6 then returns **accept** with probability at least $2/3$.

Since we are performing a total of k_0 outer loop iterations, the inner loop of $T = O(\log k_0)$ independent repetitions of the inner tester THRESHOLDZTEST along with the majority vote ensure

that each of the k_0 tests then returns `accept` with probability at least $1 - 1/(3k_0)$, and therefore by a union bound that all k_0 outcomes are simultaneously `accept` with probability at least $2/3$.

Lemma 4.4. *For the uniform distribution \mathcal{U} over $\{-1, 1\}^n$ and $k \in \mathbb{N}$, we have*

$$\|\Sigma(\odot^{(k)}(\mathcal{U}))\|_F^2 \leq 4 \cdot 2^k (n2^k)^{2^k}.$$

Proof: Let $K \stackrel{\text{def}}{=} 2^{k+1}$. Since whenever $\mathbf{x}_i^2 = 1$ for all $i \in [n]$ when $\mathbf{x} \sim \mathcal{U}$, and the \mathbf{x}_i 's are independent and have zero-mean, we have that, for any ordered K -tuple $(i_1, \dots, i_K) \in [n]^K$

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{U}} \left[\prod_{\ell=1}^K \mathbf{x}_{i_\ell} \right] = \begin{cases} 1 & \text{if each } i_\ell \text{ appears an even number of times} \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, $\|\Sigma(\odot^{(k)}(\mathcal{U}))\|_F^2$ is upper bounded by the number of ordered K -tuples of $[n]$ in which each index appears an even number of times. This in turn is at most

$$n(K-1) \cdot n(K-3) \cdot \dots \cdot n = \left(\frac{n}{2}\right)^{K/2} \cdot \frac{K!}{(K/2)!} \leq \left(\frac{n}{2}\right)^{K/2} \cdot \frac{eK}{2} \cdot K^{K/2},$$

which one can see as follows: from K variables, we may choose the value of the first (there are n choices), and then pair this variables with an other one to which we assign the same value (there are $K-1$ such choices). We then recurse on the remaining $K-2$ variables: this process, though it may lead to double-counting, ensures that each value appears an even number of times in the resulting K -tuple – as we always assign any chosen value to two variables. \blacksquare

Combining Fact 4.2 and Lemma 4.4, this implies

$$\|\mu(\odot^{(k)}(\mathcal{U}))\|_2^2 \leq 4 \cdot 2^{k-1} (n2^{k-1})^{2^{k-1}} = 2 \cdot 2^k \sqrt{(n2^k/2)^{2^k}}. \quad (52)$$

Lemma 4.5 (Completeness). *There exists a large enough universal constant $C > 0$, such that for any $k \in \mathbb{N}$ where $2^k \leq \log^2 n$. If $q \geq C/\varepsilon^2 \sqrt{n}$, then letting $\mathbf{S} = \{\mathbf{x}_1, \dots, \mathbf{x}_q, \mathbf{y}_1, \dots, \mathbf{y}_q\}$ be $2q$ i.i.d samples from $\odot^{(k)}(\mathcal{U})$. Then,*

$$\Pr_{\mathbf{S}} [\text{THRESHOLDZTEST}(\tau_k, \mathbf{S}) \text{ outputs } \text{accept}] \geq \frac{2}{3}. \quad (53)$$

Proof: Given \mathbf{S} , let \mathbf{Z} be the random variable given by (46), so $\text{THRESHOLDZTEST}(\tau_k, \mathbf{S})$ rejects whenever $\mathbf{Z} > \tau_k$. From Lemma 4.1, (51) and (52), for all $k \in \mathbb{N}$, $\mathbf{E}[\mathbf{Z}] = \|\mu(\odot^{(k)}(\mathcal{U}))\|_2^2$ satisfies

$$\|\mu(\odot^{(k)}(\mathcal{U}))\|_2^2 \leq \frac{\tau_k}{10}. \quad (54)$$

The case of $k = 0$ is trivial, and the case $k = 1$ follows from our choice of $q > C/\varepsilon^2 \sqrt{n}$. As a result, the left-hand side of (53) is at least

$$\Pr_{\mathbf{S}} [\mathbf{Z} > \tau_k] \leq \Pr_{\mathbf{S}} \left[|\mathbf{Z} - \mathbf{E}[\mathbf{Z}]| > \frac{\tau_k}{2} \right] \leq \frac{4 \text{Var}[\mathbf{Z}]}{\tau_k^2} \quad (55)$$

$$\leq \frac{4}{\tau_k^2} \left(\frac{1}{q^2} \|\Sigma(\odot^{(k)}(\mathcal{U}))\|_F^2 + \frac{4}{q} \|\mu(\odot^{(k)}(\mathcal{U}))\|_2^2 \|\Sigma(\odot^{(k)}(\mathcal{U}))\|_2 \right) \quad (56)$$

$$\leq \frac{4}{\tau_k^2 q^2} \cdot \frac{\tau_{k+1}}{10} + \frac{16}{\tau_k^2 q} \cdot \frac{\tau_k}{10} \cdot \sqrt{\frac{\tau_{k+1}}{10}} < \frac{1}{3}, \quad (57)$$

where we used Chebyshev's inequality in (55), Lemma 4.1 in (56), and Fact 4.2 and $\tau_{k+1} = a\tau_k^2 q^2$, as well as (54) in (57). \blacksquare

Finally, note that when executing $\text{MEANTESTER}(\mathcal{U}, \varepsilon)$, the algorithm performs k_0 checks in Line 6–7. By Lemma 4.5 and the fact that Line 7 outputs accept unless the majority of $T = O(\log k_0)$ iterations reject, $\text{MEANTESTER}(\mathcal{U}, \varepsilon)$ outputs accept with probability at least $2/3$. For the soundness case, the following lemma will be useful.

Lemma 4.6. *Let p be a distribution supported on $\{-1, 1\}^m$, satisfying*

1. $\|\mu(p)\|_2^2 > 2\tau$, and
2. When $\mathbf{S} = \{\mathbf{x}_1, \dots, \mathbf{x}_q, \mathbf{y}_1, \dots, \mathbf{y}_q\}$ is set to $2q$ i.i.d samples from p ,

$$\Pr_{\mathbf{S}} [\text{THRESHOLDZTEST}(\tau, \mathbf{S}) \text{ outputs accept}] \geq \frac{1}{3}.$$

Then $\|\mu(\odot(p))\|_2^2 \geq \frac{1}{48^2} \cdot \tau^2 q^2$.

Proof: Recall that $\text{THRESHOLDZTEST}(\tau, \mathbf{S})$ outputs accept if $\mathbf{Z} \leq \tau$, where \mathbf{Z} is set according to \mathbf{S} by (46). Lemma 4.1 implies $\mathbf{E}[\mathbf{Z}] = \|\mu(p)\|_2^2 > 2\tau$, so

$$\frac{1}{3} \leq \Pr_{\mathbf{Z}} [\mathbf{Z} \leq \tau] \leq \Pr_{\mathbf{Z}} \left[|\mathbf{Z} - \mathbf{E}[\mathbf{Z}]| \geq \frac{\mathbf{E}[\mathbf{Z}]}{2} \right] \leq \frac{4}{\|\mu(p)\|_2^4} \left(\frac{1}{q^2} \|\Sigma(p)\|_F^2 + \frac{4}{q} \|\mu(p)\|_2^2 \|\Sigma(p)\|_F \right), \quad (58)$$

where we used Chebyshev's inequality for the last inequality. Hence, writing $\|\mu(\odot(p))\|_2^2$ for $\|\Sigma(p)\|_F^2$ by Fact 4.2. In particular, at least one of the two terms in the right-most side of (58) is at least $1/6$, and thus either $\|\mu(\odot(p))\|_2^2 \geq \tau^2 q^2 / 6$ or $\|\mu(\odot(p))\|_2 \geq \tau q / 48$. \blacksquare

Lemma 4.7 (Soundness). *There exists a large enough $C > 0$ such that setting*

$$q \geq \left(\frac{C}{\varepsilon^2} \right)^{\frac{2^{k_0+1}}{2^{k_0+2}-2}},$$

the following holds. For any distribution p on $\{-1, 1\}^n$ with $\|\mu(p)\|_2 > \varepsilon\sqrt{n}$, there is some $k \in \{0, \dots, k_0\}$ such that letting $\mathbf{S} = \{\mathbf{x}_1, \dots, \mathbf{x}_q, \mathbf{y}_1, \dots, \mathbf{y}_q\}$ be $2q$ i.i.d samples from $\odot^{(k)}(p)$,

$$\Pr_{\mathbf{S}} [\text{THRESHOLDZTEST}(\tau_k, \mathbf{S}) \text{ outputs reject}] \geq \frac{2}{3}.$$

Proof: Assume for the sake of contradiction that for every $k \in \{0, \dots, k_0\}$, drawing \mathbf{S} from $\odot^{(k)}(p)$ satisfies $\text{THRESHOLDZTEST}(\tau_k, \mathbf{S})$ outputs accept with probability at least $1/3$. Then, by Lemma 4.6 and a simple induction using the definition of τ_k in (50) and the fact that $a < 1/(48^2 \cdot 2)$, every $k \in \{0, \dots, k_0 + 1\}$ satisfies $\|\mu(\odot^{(k)}(p))\|_2^2 \geq 2\tau_k$. Furthermore, we always have $\|\mu(\odot^{(k)}(p))\|_2^2 \leq n^{2^k}$, so by (51), we may lower bound $\|\mu(\odot^{(k_0+1)}(p))\|_2^2$ as

$$\frac{2}{aq^2} \left(\frac{aq^2 \varepsilon^2 n}{2} \right)^{2^{k_0+1}} \geq q^{2^{k_0+2}-2} (\varepsilon\sqrt{n})^{2^{k_0+2}} \left(\frac{a}{2} \right)^{2^{k_0+1}-1},$$

which contradicts the upper bound of $n^{2^{k_0+1}}$ for the setting of q when $C > 0$ is a large enough constant. \blacksquare

As per the foregoing discussion and Lemma 4.7, there exists a parameter $0 \leq k \leq k_0$ such that $\text{THRESHOLDZTEST}(\tau_k, \mathbf{S})$ returns `reject` with probability at least $2/3$ when \mathbf{S} is drawn from $\odot^{(k)}(p)$. For this setting of k , the algorithm will return `reject` in Line 7 with probability at least $1 - 1/(3T) > 2/3$. Finally, from the above analysis, the sample complexity is $2T$ times q' , which is set high enough to satisfy the constraints of Lemma 4.5 and Lemma 4.7. \blacksquare

4.1 Application to Gaussian Mean Testing

Theorem 4.8. *There exists an algorithm which, given q i.i.d. samples from an arbitrary Gaussian distribution p on \mathbb{R}^n and a distance parameter $\varepsilon \in (0, 1]$, has the following behavior:*

- If p is the standard Gaussian $\mathcal{G}(0_n, \mathbf{I}_n)$, then it outputs `accept` with probability at least $2/3$;
- If p is some $\mathcal{G}(\mu, \Sigma)$ with $\|\mu\|_2 > \varepsilon$ (and any Σ), then it outputs `reject` with probability at least $2/3$.

These guarantees hold as long as

$$q \geq C \cdot \frac{\sqrt{n}}{\varepsilon^2} \cdot \log \log \log n,$$

where $C > 0$ is an absolute constant, and the algorithm runs in time $\text{poly}(q, n^{\log n})$. Moreover, any algorithm for this task must have sample complexity $\Omega(n^{1/2}/\varepsilon^2)$.

Proof: The idea is to reduce the above question to ℓ_2 mean testing over $\{-1, 1\}^n$, and invoke Theorem 4.3. The natural approach is, given a sample $x \in \mathbb{R}^n$ from the unknown Gaussian p , to convert it to $y \in \{-1, 1\}^n$ by setting $y_i \stackrel{\text{def}}{=} \text{sign}(x_i)$ for all $i \in [n]$. This idea, used e.g., in [CKM⁺19] in the case of identity-covariance Gaussian distributions, clearly maps the standard Gaussian $\mathcal{G}(0_n, \mathbf{I}_n)$ to the uniform distribution \mathcal{U} on $\{-1, 1\}^n$; therefore, the crux is to argue about the soundness case, when $\|\mu(p)\|_2 > \varepsilon$: we will obtain a distribution q on $\{-1, 1\}^n$ with arbitrary covariance matrix, and need to show that $\|\mu(q)\|_2 \gtrsim \varepsilon$.

Let q denote the distribution on $\{-1, 1\}^n$ obtained in the above fashion from $p = \mathcal{G}(\mu, \Sigma)$, and note that the linear-time transformation maps a sample from p to a sample from q . Further, it is easy to check that

$$\|\mu(q)\|_2^2 = \sum_{i=1}^n \mathbf{E}_{y \sim q} [y_i]^2 = \sum_{i=1}^n (2 \Pr_{x \sim p}[x_i > 0] - 1)^2 = \sum_{i=1}^n \left(\text{Erf}\left(\frac{\mu_i}{\sqrt{2\Sigma_{ii}}}\right) \right)^2$$

and, from a relatively straightforward analysis of the error function Erf , we get that $\text{Erf}(x)^2 \geq \text{Erf}(1)^2 \cdot \min(x^2, 1) > \frac{2}{3} \min(x^2, 1)$ for all $x \in \mathbb{R}$. Therefore, we get, whenever $\|\mu(p)\|_2 > \varepsilon$, that

$$\|\mu(q)\|_2^2 > \frac{1}{\max_{1 \leq i \leq n} \Sigma_{ii}^2} \cdot \frac{\varepsilon^2}{3}.$$

Thus, it suffices to check (with high probability, say $11/12$) that (i) all $\Sigma_{ii} = \mathbf{Var}_p[x_i] \leq \mathbf{E}_{x \sim p}[x_i^2]$ are at most 2, and (ii) call the mean testing algorithm (Algorithm 3) on q with parameter $\varepsilon/(2\sqrt{3n})$. The first step can be done coordinate-wise (checking each $\mathbf{E}_{x \sim p}[x_i^2]$ by taking the empirical median, and concluding overall by a union bound over the n estimates)⁷ with $O(\log n)$ samples in total;

⁷In more detail, for an arbitrary univariate Gaussian $X \sim \mathcal{G}(\mu, \sigma)$, testing $\mathbf{E}[X^2] \leq 1$ vs. $\mathbf{E}[X^2] > 2$ only takes, by considering the median, a constant number of samples (independent of μ, σ).

the second will cost $O(\frac{\sqrt{n}}{\varepsilon^2} \cdot \log \log \log n)$ samples and be correct with probability at least $2/3$ by the foregoing discussion, so overall the test is correct with probability at least $7/12$. Repeating independently a constant number of times and taking a majority vote then bring the success probability to the desired $2/3$. ■

References

- [ABDK18] Jayadev Acharya, Arnab Bhattacharyya, Constantinos Daskalakis, and Saravanan Kandasamy. Learning and testing causal models with interventions. In *Advances in Neural Information Processing Systems 31*, NeurIPS '18. Curran Associates, Inc., 2018. [1.2](#)
- [ACK15a] Jayadev Acharya, Clément L. Canonne, and Gautam Kamath. Adaptive estimation in weighted group testing. In *Proceedings of the 2015 IEEE International Symposium on Information Theory*, ISIT '15, pages 2116–2120, Washington, DC, USA, 2015. IEEE Computer Society. [1.2](#)
- [ACK15b] Jayadev Acharya, Clément L. Canonne, and Gautam Kamath. A chasm between identity and equivalence testing with conditional queries. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques.*, RANDOM '15, pages 449–466, Dagstuhl, Germany, 2015. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. [1.2](#)
- [ADJ⁺11] Jayadev Acharya, Hirakendu Das, Ashkan Jafarpour, Alon Orlitsky, and Shengjun Pan. Competitive closeness testing. In *Proceedings of the 24th Annual Conference on Learning Theory*, COLT '11, pages 47–68, 2011. [1.2](#)
- [ADK15] Jayadev Acharya, Constantinos Daskalakis, and Gautam Kamath. Optimal testing for properties of distributions. In *Advances in Neural Information Processing Systems 28*, NIPS '15, pages 3577–3598. Curran Associates, Inc., 2015. [1.2](#)
- [BBC⁺19] Ivona Bezakova, Antonio Blanca, Zongchen Chen, Daniel Štefankovič, and Eric Vigoda. Lower bounds for testing graphical models: Colorings and antiferromagnetic Ising models. In *Proceedings of the 32nd Annual Conference on Learning Theory*, COLT '19, pages 283–298, 2019. [1.2](#)
- [BC17] Tuğkan Batu and Clément L. Canonne. Generalized uniformity testing. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '17, pages 880–889, Washington, DC, USA, 2017. IEEE Computer Society. [1.2](#)
- [BC18] Rishiraj Bhattacharyya and Sourav Chakraborty. Property testing of joint distributions using conditional samples. *Transactions on Computation Theory*, 10(4):16:1–16:20, 2018. [1](#), [1](#), [1.2](#), [1.3](#)
- [BCG17] Eric Blais, Clément L. Canonne, and Tom Gur. Distribution testing lower bounds via reductions from communication complexity. In *Proceedings of the 32nd Computational Complexity Conference*, CCC '17, pages 28:1–28:40, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. [1.2](#)

- [BDKR05] Tuğkan Batu, Sanjoy Dasgupta, Ravi Kumar, and Ronitt Rubinfeld. The complexity of approximating the entropy. *SIAM Journal on Computing*, 35(1):132–150, 2005. 1.2
- [BFF⁺01] Tuğkan Batu, Eldar Fischer, Lance Fortnow, Ravi Kumar, Ronitt Rubinfeld, and Patrick White. Testing random variables for independence and identity. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, FOCS '01, pages 442–451, Washington, DC, USA, 2001. IEEE Computer Society. 1.2
- [BFR⁺00] Tuğkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D. Smith, and Patrick White. Testing that distributions are close. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, FOCS '00, pages 259–269, Washington, DC, USA, 2000. IEEE Computer Society. 1, 1.2
- [BFRV11] Arnab Bhattacharyya, Eldar Fischer, Ronitt Rubinfeld, and Paul Valiant. Testing monotonicity of distributions over general partial orders. In *Proceedings of the 2nd Conference on Innovations in Computer Science*, ICS '11, pages 239–252, Beijing, China, 2011. Tsinghua University Press. 1.2
- [BKR04] Tuğkan Batu, Ravi Kumar, and Ronitt Rubinfeld. Sublinear algorithms for testing monotone and unimodal distributions. In *Proceedings of the 36th Annual ACM Symposium on the Theory of Computing*, STOC '04, New York, NY, USA, 2004. ACM. 1.2
- [BV15] Bhaswar Bhattacharya and Gregory Valiant. Testing closeness with unequal sized samples. In *Advances in Neural Information Processing Systems 28*, NIPS '15, pages 2611–2619. Curran Associates, Inc., 2015. 1.2
- [BW18] Sivaraman Balakrishnan and Larry Wasserman. Hypothesis testing for high-dimensional multinomials: A selective review. *The Annals of Applied Statistics*, 12(2):727–749, 2018. 1.2
- [Can15a] Clément L. Canonne. Big data on the rise? - testing monotonicity of distributions. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming*, ICALP '15, pages 294–305, 2015. 1.2
- [Can15b] Clément L. Canonne. A survey on distribution testing: Your data is big. but is it blue? *Electronic Colloquium on Computational Complexity (ECCC)*, 22(63), 2015. 1.2
- [Can16] Clément L. Canonne. Are few bins enough: Testing histogram distributions. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '16, pages 455–463, New York, NY, USA, 2016. ACM. 1.2
- [CDKS17] Clément L. Canonne, Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. Testing Bayesian networks. In *Proceedings of the 30th Annual Conference on Learning Theory*, COLT '17, pages 370–448, 2017. 1, 1, 1, 3, 1.2
- [CDVV14] Siu On Chan, Ilias Diakonikolas, Gregory Valiant, and Paul Valiant. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 1193–1203, Philadelphia, PA, USA, 2014. SIAM. 1.2

- [CFG13] Sourav Chakraborty, Eldar Fischer, Yonatan Goldhirsh, and Arie Matsliah. On the power of conditional samples in distribution testing. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, ITCS '13, pages 561–580, New York, NY, USA, 2013. ACM. 1.2
- [CFG16] Sourav Chakraborty, Eldar Fischer, Yonatan Goldhirsh, and Arie Matsliah. On the power of conditional samples in distribution testing. *SIAM Journal on Computing*, 45(4):1261–1296, 2016. 1, 1.2, 1.3
- [CKM⁺19] Clément L. Canonne, Gautam Kamath, Audra McMillan, Jonathan Ullman, and Lydia Zakyntinou. Private identity testing for high-dimensional distributions. *arXiv preprint arXiv:1905.11947*, 2019. 4.1
- [CR14] Clément L. Canonne and Ronitt Rubinfeld. Testing probability distributions underlying aggregated data. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming*, ICALP '14, pages 283–295, 2014. 1.2
- [CRS14] Clément L. Canonne, Dana Ron, and Rocco A. Servedio. Testing equivalence between distributions using conditional samples. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 1174–1192, Philadelphia, PA, USA, 2014. SIAM. 1.2
- [CRS15] Clément L. Canonne, Dana Ron, and Rocco A. Servedio. Testing probability distributions using conditional samples. *SIAM Journal on Computing*, 44(3):540–616, 2015. 1, 1, 1.2, 1.3
- [DDK18] Constantinos Daskalakis, Nishanth Dikkala, and Gautam Kamath. Testing Ising models. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, pages 1989–2007, Philadelphia, PA, USA, 2018. SIAM. 1, 1, 1.2
- [DDS⁺13] Constantinos Daskalakis, Ilias Diakonikolas, Rocco A. Servedio, Gregory Valiant, and Paul Valiant. Testing k-modal distributions: Optimal algorithms via reductions. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 1833–1852, Philadelphia, PA, USA, 2013. SIAM. 1.2
- [DGPP18] Ilias Diakonikolas, Themis Gouleakis, John Peebles, and Eric Price. Sample-optimal identity testing with high probability. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming*, ICALP '18, pages 41:1–41:14, 2018. 1.2
- [DK16] Ilias Diakonikolas and Daniel M. Kane. A new approach for testing properties of discrete distributions. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '16, pages 685–694, Washington, DC, USA, 2016. IEEE Computer Society. 1.2
- [DKN15] Ilias Diakonikolas, Daniel M. Kane, and Vladimir Nikishkin. Testing identity of structured distributions. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 1841–1854, Philadelphia, PA, USA, 2015. SIAM. 1.2

- [DKP19] Ilias Diakonikolas, Daniel M. Kane, and John Peebles. Testing identity of multidimensional histograms. In *Proceedings of the 32nd Annual Conference on Learning Theory*, COLT '19, pages 1107–1131, 2019. [1.2](#)
- [DKW18] Constantinos Daskalakis, Gautam Kamath, and John Wright. Which distribution distances are sublinearly testable? In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, pages 2747–2764, Philadelphia, PA, USA, 2018. SIAM. [1.2](#)
- [DP17] Constantinos Daskalakis and Qinxuan Pan. Square Hellinger subadditivity for Bayesian networks and its applications to identity testing. In *Proceedings of the 30th Annual Conference on Learning Theory*, COLT '17, pages 697–703, 2017. [1.2](#)
- [FJO⁺15] Moein Falahatgar, Ashkan Jafarpour, Alon Orlitsky, Venkatadheeraj Pichapati, and Ananda Theertha Suresh. Faster algorithms for testing under conditional sampling. In *Proceedings of the 28th Annual Conference on Learning Theory*, COLT '15, pages 607–636, 2015. [1.2](#)
- [FLV17] Eldar Fischer, Oded Lachish, and Yadu Vasudev. Improving and extending the testing of distributions for shape-restricted properties. In *Proceedings of the 34th Symposium on Theoretical Aspects of Computer Science*, STACS '17, pages 31:1–31:14, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. [1.2](#)
- [GGR96] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '96, pages 339–348, Washington, DC, USA, 1996. IEEE Computer Society. [1](#), [1.2](#)
- [GLP18] Reza Gheissari, Eyal Lubetzky, and Yuval Peres. Concentration inequalities for polynomials of contracting Ising models. *Electronic Communications in Probability*, 23(76):1–12, 2018. [1.2](#)
- [GMV06] Sudipto Guha, Andrew McGregor, and Suresh Venkatasubramanian. Streaming and sublinear approximation of entropy and information distances. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '06, pages 733–742, Philadelphia, PA, USA, 2006. SIAM. [1.2](#)
- [Gol17] Oded Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017. [1.2](#)
- [GR00] Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, 7(20), 2000. [1](#), [1.2](#)
- [GTZ17] Themistoklis Gouleakis, Christos Tzamos, and Manolis Zampetakis. Faster sublinear algorithms using conditional sampling. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages 1743–1757, Philadelphia, PA, USA, 2017. SIAM. [1.2](#)

- [GTZ18] Themis Gouleakis, Christos Tzamos, and Manolis Zampetakis. Certified computation from unreliable datasets. In *Proceedings of the 31st Annual Conference on Learning Theory*, COLT '18, pages 3271–3294, 2018. [1.2](#)
- [ILR12] Piotr Indyk, Reut Levi, and Ronitt Rubinfeld. Approximating and testing k-histogram distributions in sub-linear time. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '12, pages 15–22, New York, NY, USA, 2012. ACM. [1.2](#)
- [Kam18] Gautam Kamath. *Modern Challenges in Distribution Testing*. PhD thesis, Massachusetts Institute of Technology, September 2018. [1.2](#)
- [KMS18] Subhash Khot, Dor Minzer, and Muli Safra. On monotonicity testing and Boolean isoperimetric-type theorems. *SIAM Journal on Computing*, 47(6):2238–2276, 2018. [1.1.1](#), [3.1](#), [3.3](#)
- [KT19] Gautam Kamath and Christos Tzamos. Anaconda: A non-adaptive conditional sampling algorithm for distribution testing. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, Philadelphia, PA, USA, 2019. SIAM. [1.2](#)
- [NS02] Assaf Naor and Gideon Schechtman. Remarks on non-linear type and Pisier’s inequality. *Journal fur die Reine und Angewandte Mathematik*, 552:213–236, 2002. [1.1.1](#), [1.1.1](#), [3.1](#)
- [OS18] Krzysztof Onak and Xiaorui Sun. Probability-revealing samples. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, AISTATS '18, pages 84:2018–2026. JMLR, Inc., 2018. [1.2](#)
- [Pan08] Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008. [1](#), [1.2](#)
- [Pis86] Gilles Pisier. Probabilistic methods in the geometry of Banach spaces. In Giorgio Letta and Maurizio Pratelli, editors, *Probability and Analysis*, pages 167–241. Springer, 1986. [1.1](#), [1.1.1](#)
- [RS09] Ronitt Rubinfeld and Rocco A. Servedio. Testing monotone high-dimensional distributions. *Random Structures and Algorithms*, 34(1):24–44, 2009. [1.2](#)
- [Rub12] Ronitt Rubinfeld. Taming big probability distributions. *XRDS*, 19(1):24–28, 2012. [1.2](#)
- [SSJ17] Imdad S. B. Sardharwalla, Sergii Strelchuk, and Richard Jozsa. Quantum conditional query complexity. *Quantum Information & Computation*, 17(7& 8):541–566, 2017. [1.2](#)
- [Tal93] Michel Talagrand. Isoperimetry, logarithmic Sobolev inequalities on the discrete cube, and Margulis’ graph connectivity theorem. *Geometric & Functional Analysis*, 3(3):295–314, 1993. [1.8](#)
- [Val11] Paul Valiant. Testing symmetric properties of distributions. *SIAM Journal on Computing*, 40(6):1927–1968, 2011. [1.2](#)

- [VV14] Gregory Valiant and Paul Valiant. An automatic inequality prover and instance optimal identity testing. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '14, pages 51–60, Washington, DC, USA, 2014. IEEE Computer Society. [1](#)
- [VV17] Gregory Valiant and Paul Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017. [1.2](#)
- [Wag15] Bo Waggoner. l_p testing and learning of discrete distributions. In *Proceedings of the 6th Conference on Innovations in Theoretical Computer Science*, ITCS '15, pages 347–356, New York, NY, USA, 2015. ACM. [1.2](#)

A Proof of Lemma 4.1

We start by computing the second moment of the statistic:

$$\begin{aligned}
\mathbf{E}[(\mathbf{Z}(p))^2] &= \mathbf{E}[\langle \bar{\mathbf{X}}, \bar{\mathbf{Y}} \rangle \langle \bar{\mathbf{X}}, \bar{\mathbf{Y}} \rangle] = \sum_{1 \leq i, j \leq m} \mathbf{E}[\bar{\mathbf{X}}_i \bar{\mathbf{X}}_j] \mathbf{E}[\bar{\mathbf{Y}}_i \bar{\mathbf{Y}}_j] = \sum_{1 \leq i, j \leq m} \mathbf{E}[\bar{\mathbf{X}}_i \bar{\mathbf{X}}_j]^2 \\
&= \sum_{i=1}^m \mathbf{E}[\bar{\mathbf{X}}_i^2]^2 + 2 \sum_{i < j} \mathbf{E}[\bar{\mathbf{X}}_i \bar{\mathbf{X}}_j]^2 = \frac{m}{q^2} + 2 \frac{q-1}{q^2} \|\mu(p)\|_2^2 + \frac{(q-1)^2}{q^2} \|\mu(p)\|_4^4 + 2 \sum_{i < j} \mathbf{E}[\bar{\mathbf{X}}_i \bar{\mathbf{X}}_j]^2 \\
&\leq \frac{m}{q^2} + \frac{2}{q} \|\mu(p)\|_2^2 + \|\mu(p)\|_4^4 + 2 \sum_{i < j} \mathbf{E}[\bar{\mathbf{X}}_i \bar{\mathbf{X}}_j]^2. \tag{59}
\end{aligned}$$

The final equality follows since, for all $i \in [m]$,

$$q^2 \mathbf{E}[\bar{\mathbf{X}}_i^2] = \sum_{k=1}^q \mathbf{E}[\mathbf{X}_i^{(k)2}] + 2 \sum_{1 \leq k < \ell \leq q} \mathbf{E}[\mathbf{X}_i^{(k)} \mathbf{X}_i^{(\ell)}] = q + q(q-1)\mu(p)_i^2 \tag{60}$$

For any $i < j$, let $\sigma(p)_{ij} \stackrel{\text{def}}{=} \mathbf{E}_{x \sim p}[x_i x_j]$. We have

$$\begin{aligned}
\mathbf{E}[\bar{\mathbf{X}}_i \bar{\mathbf{X}}_j] &= \frac{1}{q^2} \sum_{1 \leq k, \ell \leq q} \mathbf{E}[\mathbf{X}_i^{(k)} \mathbf{X}_j^{(\ell)}] = \frac{1}{q^2} \left(\sum_{k=1}^q \mathbf{E}[\mathbf{X}_i^{(k)} \mathbf{X}_j^{(k)}] + 2 \sum_{1 \leq k < \ell \leq q} \mathbf{E}[\mathbf{X}_i^{(k)}] \mathbf{E}[\mathbf{X}_j^{(\ell)}] \right) \\
&= \frac{1}{q} \sigma(p)_{ij} + \frac{q-1}{q} \mu(p)_i \mu(p)_j,
\end{aligned}$$

so that

$$\begin{aligned}
2 \sum_{i < j} \mathbf{E}[\bar{\mathbf{X}}_i \bar{\mathbf{X}}_j]^2 &= \frac{1}{q^2} \sum_{i \neq j} \sigma(p)_{ij}^2 + \frac{(q-1)^2}{q^2} \sum_{i \neq j} \mu(p)_i^2 \mu(p)_j^2 + \frac{2(q-1)}{q^2} \sum_{i \neq j} \sigma(p)_{ij} \mu(p)_i \mu(p)_j \\
&\leq \frac{1}{q^2} \sum_{i \neq j} \sigma(p)_{ij}^2 + \|\mu(p)\|_2^4 - \|\mu(p)\|_4^4 + \frac{2}{q} \sum_{i \neq j} \sigma(p)_{ij} \mu(p)_i \mu(p)_j.
\end{aligned}$$

Combining this with (47) and (59), we have

$$\mathbf{Var}[\mathbf{Z}(p)] = \mathbf{E}[(\mathbf{Z}(p))^2] - \mathbf{E}[\mathbf{Z}(p)]^2 \leq \frac{m}{q^2} + \frac{2}{q} \|\mu(p)\|_2^2 + \frac{1}{q^2} \sum_{i \neq j} \sigma(p)_{ij}^2 + \frac{2}{q} \sum_{i \neq j} \sigma(p)_{ij} \mu(p)_i \mu(p)_j \tag{61}$$

We use the Cauchy-Schwarz inequality to simplify the last term:

$$\begin{aligned}
\mathbf{Var}[\mathbf{Z}(p)] &\leq \frac{m}{q^2} + \frac{2}{q} \|\mu(p)\|_2^2 + \frac{1}{q^2} \sum_{i \neq j} \sigma(p)_{ij}^2 + \frac{2}{q} \sqrt{\sum_{i \neq j} \sigma(p)_{ij}^2} \sqrt{\sum_{i \neq j} \mu(p)_i^2 \mu(p)_j^2} \\
&= \frac{m}{q^2} + \frac{2}{q} \|\mu(p)\|_2^2 + \frac{1}{q^2} \sum_{i \neq j} \sigma(p)_{ij}^2 + \frac{2}{q} \sqrt{\sum_{i \neq j} \sigma(p)_{ij}^2} \sqrt{\|\mu(p)\|_2^4 - \|\mu(p)\|_4^4} \\
&\leq \frac{m}{q^2} + \frac{2}{q} \|\mu(p)\|_2^2 + \frac{1}{q^2} \sum_{i \neq j} \sigma(p)_{ij}^2 + \frac{2}{q} \|\mu(p)\|_2^2 \sqrt{\sum_{i \neq j} \sigma(p)_{ij}^2}
\end{aligned}$$

Letting $\Sigma(p) \stackrel{\text{def}}{=} \mathbf{E}_{x \sim p} [xx^T]$, we have

$$\Sigma(p)_{ij} = \begin{cases} 1 & \text{if } i = j \\ \sigma(p)_{ij} & \text{otherwise} \end{cases}$$

and thus can rewrite the above as

$$\begin{aligned}
\mathbf{Var}[\mathbf{Z}(p)] &\leq \frac{1}{q^2} \|\Sigma(p)\|_F^2 + \frac{2}{q} \|\mu(p)\|_2^2 + \frac{2}{q} \|\mu(p)\|_2^2 \sqrt{\|\Sigma(p)\|_F^2 - \text{Tr}[\Sigma(p)]} \\
&\leq \frac{1}{q^2} \|\Sigma(p)\|_F^2 + \frac{2}{q} \|\mu(p)\|_2^2 + \frac{2}{q} \|\mu(p)\|_2^2 \|\Sigma(p)\|_F \\
&\leq \frac{1}{q^2} \|\Sigma(p)\|_F^2 + \frac{4}{q} \|\mu(p)\|_2^2 \|\Sigma(p)\|_F,
\end{aligned}$$

as desired.