

Exploring the Gap between Tolerant and Non-tolerant Distribution Testing*

Sourav Chakraborty[†] Eldar Fischer[‡] Arijit Ghosh[§] Gopinath Mishra[¶]
Sayantan Sen^{||}

Abstract

The framework of distribution testing is currently ubiquitous in the field of property testing. In this model, the input is a probability distribution accessible via independently drawn samples from an oracle. The testing task is to distinguish a distribution that satisfies some property from a distribution that is far in some distance measure from satisfying it. The task of tolerant testing imposes a further restriction, that distributions close to satisfying the property are also accepted.

This work focuses on the connection between the sample complexities of non-tolerant testing of distributions and their tolerant testing counterparts. When limiting our scope to label-invariant (symmetric) properties of distributions, we prove that the gap is at most quadratic, ignoring poly-logarithmic factors. Conversely, the property of being the uniform distribution is indeed known to have an almost-quadratic gap.

When moving to general, not necessarily label-invariant properties, the situation is more complicated, and we show some partial results. We show that if a property requires the distributions to be non-concentrated, that is, the probability mass of the distribution is sufficiently spread out, then it cannot be non-tolerantly tested with $o(\sqrt{n})$ many samples, where n denotes the universe size. Clearly, this implies at most a quadratic gap, because a distribution can be learned (and hence tolerantly tested against any property) using $\mathcal{O}(n)$ many samples. Being non-concentrated is a strong requirement on properties, as we also prove a close to linear lower bound against their tolerant tests.

Apart from the case where the distribution is non-concentrated, we also show if an input distribution is very concentrated, in the sense that it is mostly supported on a subset of size s of the universe, then it can be learned using only $\mathcal{O}(s)$ many samples. The learning procedure adapts to the input, and works without knowing s in advance.

*A preliminary version of this paper has been accepted in RANDOM 2022.

[†]Indian Statistical Institute, Kolkata, India. Email: chakraborty.sourav@gmail.com.

[‡]Technion - Israel Institute of Technology, Israel. Email: eldar@cs.technion.ac.il.

[§]Indian Statistical Institute, Kolkata, India. Email: arijitiitkgpster@gmail.com.

[¶]University of Warwick, UK. Email: gopianjan117@gmail.com. Research supported in part by the Centre for Discrete Mathematics and its Applications (DIMAP) and by EPSRC award EP/V01305X/1.

^{||}Indian Statistical Institute, Kolkata, India. Email: sayantan789@gmail.com.

1 Introduction

Let D be a distribution over a finite set Ω , and \mathcal{P} be a property, that is, a set of distributions over Ω . Given access to independent random samples from Ω according to the distribution D , we are interested in the problem of distinguishing whether the distribution D is η -close to having the property \mathcal{P} , or is ε -far from having the property \mathcal{P} , where η and ε are two fixed proximity parameters such that $0 \leq \eta < \varepsilon \leq 2$. The distance of the distribution D from the property \mathcal{P} is defined as $\min_{D' \in \mathcal{P}} \|D - D'\|_1$, where $\|D - D'\|_1$ denotes the ℓ_1 -distance between the distributions D and D' ¹. A distribution D is said to be η -close to \mathcal{P} if the distance of D from \mathcal{P} is at most η . Similarly, D is said to be ε -far from \mathcal{P} , if the distance of D from \mathcal{P} is at least ε . The goal is to design a tester that uses as few samples as possible. For $\eta > 0$, the problem of distinguishing the two cases is referred to as the *tolerant distribution testing* problem of \mathcal{P} , and the particular case where $\eta = 0$ is referred to as the *non-tolerant distribution testing* problem of \mathcal{P} . The sample complexity (tolerant and non-tolerant testing) is the number of samples required by the best algorithm that can distinguish with high probability (usually with probability at least $\frac{2}{3}$) whether the distribution D is η -close to having the property \mathcal{P} , or is ε -far from having the property \mathcal{P} .

While results and techniques from distribution testing are already interesting in their own right, they have also found numerous applications in central problems in Theoretical Computer Science, and in particular in property testing, e.g. graph isomorphism testing [FM08, Gol19] and function isomorphism testing [ABC⁺13], learning theory [BBC⁺10, DKS17, DK16], and differential privacy [ADKR19, GKK⁺20, Zha21, ACF⁺21]. Thus, understanding the tolerant and non-tolerant sample complexity of distribution testing is a central problem in theoretical computer science.

There have been extensive studies of non-tolerant and tolerant testing of some specific distribution properties like uniformity, identity with a fixed distribution, equality of two distributions and independence of a joint distribution [BFR⁺00, BFF⁺01, Pan08, Val11, VV11, VV17a]. Various other specific distribution properties have also been studied [BC17, DKS18]. Then, some works investigated general tests for the large class of all shape-restricted properties of distributions, which contains properties like monotonicity, log-concavity, modality etc. [CDGR18, FLV17]. This paper proves general results about the gap between tolerant and non-tolerant distribution testing that hold for large classes of properties.

1.1 Our results

We now informally present our results. The formal definitions are presented in Section 2. We assume that the distributions are supported over a set $\Omega = [n] = \{1, 2, \dots, n\}$. We first prove a result about label-invariant distribution properties (properties that are invariant under all permutations of Ω). We show that, for any label-invariant distribution property, there is at most a quadratic blowup in its tolerant sample complexity as compared to its non-tolerant counterpart, ignoring poly-logarithmic factors.

Theorem 1.1 (Informal). *Any label-invariant distribution property that can be non-tolerantly tested using Λ samples, can also be tolerantly tested using $\tilde{\mathcal{O}}(\min\{\Lambda^2, n\})$ samples, where n is the size of the support of the distribution².*

¹Strictly speaking it is an infimum, but since all properties we consider are compact sets, it is equal to the minimum.

² $\tilde{\mathcal{O}}(\cdot)$ hides a poly-logarithmic factor.

This result gives a unified way for obtaining tolerant testers from their non-tolerant counterparts. The above result will be stated and proved formally in Section 4. Moreover, in Section 5, we give a constructive variant of the tolerant tester of Theorem 1.1, when the property can be expressed as the feasible solution to a set of linear inequalities.

Theorem 1.2 (Informal). *Any label-invariant distribution property that can be non-tolerantly tested using Λ samples and can be expressed as a feasible solution to m linear inequalities, can also be tolerantly tested using $\tilde{\mathcal{O}}(\min\{\Lambda^2, n\})$ samples and in time polynomial in m and n , where n is the size of the support of the distribution.*

Note that if $\Lambda = \Omega(\sqrt{n})$, Theorem 1.1 is obvious. It is only interesting if $\Lambda = o(\sqrt{n})$. Now we present a property for which this connection is useful. Consider a natural distribution property: given a distribution D and a parameter k , we want to decide whether the support size of D is at most k or ε -far from having support at most k . If $k = o(\sqrt{n})$, the query complexity for testing this problem is $\mathcal{O}(\frac{k}{\log k})$ [VV17b].

It is a natural question to investigate the extent to which the above theorem can be generalized. Though we are not resolving this question completely, as a first step in the direction of extending the above theorem for properties that are not necessarily label-invariant, we consider the notion of *non-concentrated* properties. By the notion of a non-concentrated distribution, intuitively, we mean that there is no significant portion of the base set of the distribution that carries only a negligible weight, making the probability mass of the distribution well distributed among its indices. Specifically, any subset $X \subseteq [n]$, for which $|X|$ is above some threshold (say βn with $\beta \in (0, \frac{1}{2})$), has probability mass of at least another threshold (say α with $\alpha \in (0, \frac{1}{2})$). A property is said to be non-concentrated if only non-concentrated distributions can satisfy the property. We prove a lower bound on the testing of any non-concentrated property (not necessarily label-invariant).

Theorem 1.3 (Informal). *In order to non-tolerantly test any non-concentrated distribution property, $\Omega(\sqrt{n})$ samples are required, where n is the size of the support of the distribution.*

The quadratic gap between tolerant testing and non-tolerant testing for any non-concentrated property follows from the above theorem, since by a folklore result, only $\mathcal{O}(n)$ many samples are required to learn any distribution approximately.

The proof of Theorem 1.3 for label-invariant non-concentrated properties is a generalization of the proof of the $\Omega(\sqrt{n})$ lower bound for classical uniformity testing, while for the whole theorem, that is, for the general (not label-invariant) non-concentrated properties, a more delicate argument is required. The formal proof is presented in Section 7.

The next natural question is about the sample complexity of any tolerant tester for non-concentrated properties. We address this question for label-invariant non-concentrated properties by proving the following theorem in Section 6.2. However, the question is left open for non-label-invariant properties.

Theorem 1.4 (Informal). *The sample complexity for tolerantly testing any non-concentrated label-invariant distribution property is $\Omega(n^{1-o(1)})$, where n is the size of the support of the distribution.*

A natural question related to tolerant testing is:

How many samples are required to learn a distribution?

As pointed out earlier, any distribution can be learnt using $\mathcal{O}(n)$ samples. But what if the distribution happens to be *very concentrated*? We present an upper bound result for learning a distribution, in which the sample complexity depends on the minimum cardinality of any set $S \subseteq [n]$ over which the unknown distribution is concentrated.

Theorem 1.5 (Informal). *To learn a distribution approximately, $\mathcal{O}(|S|)$ samples are enough, where $S \subseteq [n]$ is an unknown set of minimum cardinality whose mass is close to 1. Note that $|S|$ is also unknown, and the algorithm adapts to it.*

Observe that we cannot learn a distribution supported on the set S using $o(|S|)$ samples, so the above result is essentially tight.

1.2 Related works

Several forms of distribution testing have been investigated for over a century in statistical theory [Kin97, CF14], while combinatorial properties of distributions have been explored over the last two decades in Algorithm Theory, Machine Learning and Information Theory [Gol17, Mac03, CT01, BY22]. In Algorithm Theory, the investigation into testing properties of distributions started with the work of Goldreich and Ron [GR00], even though it was not directly stated there in these terms. Batu, Fortnow, Rubinfeld, Smith and White [BFR⁺00] launched the intensive study of property testing of distributions with the problem of equivalence testing³. Later, Batu, Fischer, Fortnow, Kumar, Rubinfeld and White [BFF⁺01] studied the problems of identity and independence testing of distributions⁴. Since then there has been a flurry of interesting works in this model. For example, Paninski [Pan08] proved tight bounds on uniformity testing, Valiant and Valiant [VV11] resolved the tolerant sample complexity for a large class of label-invariant properties that includes uniformity testing, Acharya, Daskalakis and Kamath [ADK15] proved various optimal testing results under several distance measures, and Valiant and Valiant [VV17a] studied the sample complexity of instance optimal identity testing. In [BC17], Batu and Cannone studied the problem of *generalized uniformity testing*, where the distribution is promised to be supported on an unknown set S , and proved a tight bound of $\tilde{\Theta}(|S|^{2/3})$ samples for non-tolerant uniformity testing. This is in contrast to the non-tolerant uniformity testing of a distribution supported over $[n]$, whose sample complexity is $\Theta(\sqrt{n})$, ignoring the dependence on the proximity parameter. Daskalakis, Kamath and Wright [DKW18] studied the problem of tolerant testing under various distance measures. Very recently, Canonne, Jain, Kamath and Li [CJKL22] revisited the problem of determining the sample complexity of tolerant identity testing, where they proved the optimal dependence on the proximity parameters. Going beyond studying specific properties, Canonne, Diakonikolas, Gouleakis and Rubinfeld [CDGR18] studied the class of *shape-restricted* properties of a distribution, a condition general enough to contain several interesting properties like monotonicity, log-concavity, t -modality etc. Their result was later improved by Fischer, Lachish and Vasudev [FLV17]. See the surveys of Cannone [Can20b, Can22] for a more exhaustive list.

While the most studied works concentrate on non-tolerant testing of distributions, a natural extension is to test such properties tolerantly. Since the introduction of tolerant testing in the

³Given two unknown probability distributions that can be accessed via samples from their respective oracles, equivalence testing refers to the problem of distinguishing whether they are same or far from each other.

⁴Given an unknown distribution accessible via samples, the problem of identity testing refers to the problem of distinguishing whether it is identical to a known distribution or far from it.

pioneering work of Parnas, Ron and Rubinfeld [PRR06], that defined this notion for classical (non-distribution) property testing, there have been several works in this framework. Note that it is nontrivial in many cases to construct tolerant testers from their non-tolerant counterparts, as in the case of tolerant junta testing [BCE⁺19] for example. In a series of works, it has been proved that tolerant testing of the most natural distribution properties, like uniformity, requires an almost linear number of samples [Val11, VV11]⁵. Now a natural question arises about how the sampling complexity of tolerant testing is related to non-tolerant testing of distributions in general. To the best of our knowledge, there is no known example with more than a quadratic gap.

It would also be interesting to bound the gap for sample-based testing as defined in the work of Goldreich and Ron [GR16]. This model was investigated further in the work of Fischer, Lachish and Vasudev [FLV15], where a general upper bound for non-tolerant sample-based testing of strongly testable properties was provided.

Organization of the paper Section 2 contains the definitions used throughout the paper. The high-level overview of the proofs of all our theorems are presented in Section 3. Section 4 contains the formal statement and proof of Theorem 1.1. Theorem 1.2 has been formally stated and proved in Section 5. Finally Theorem 1.3, Theorem 1.4 and Theorem 1.5 are formally stated and proved in Section 6, Section 7 and Section 8 respectively.

2 Notation and definitions

A probability distribution D over a universe $\Omega = [n]$ is a non-negative function $D : \Omega \rightarrow [0, 1]$ such that $\sum_{i \in \Omega} D(i) = 1$. For $S \subseteq \Omega$, the mass of S is defined as $D(S) = \sum_{i \in S} D(i)$, where $D(i)$ is the mass of i in D . The support of a probability distribution D on Ω is denoted by $\text{SUPP}(D)$. For any distribution D , by the top t elements of D , we refer to the first t elements in the support of D when the elements in the support are sorted according to a non-increasing order of their probability masses in D . When we write $\tilde{O}(\cdot)$, it sometimes suppresses a poly-logarithmic term in n and the inverse of the proximity parameter(s), as well as the inverse of the difference of two proximity parameters. Although there are several other distance measures, in this work, we mainly focus on the ℓ_1 distance.

Definition 2.1 (Distribution property). Let \mathcal{D} denote the set of all distributions over Ω . A *distribution property* \mathcal{P} is a topologically closed subset of \mathcal{D} ⁶. A distribution $D \in \mathcal{P}$ is said to be *in the property* or to *satisfy* the property. Otherwise, D is said to be *not in the property* or to *not satisfy* the property.

Definition 2.2 (Label-invariant property). Let us consider a property \mathcal{P} . For a distribution D and a permutation $\sigma : \Omega \rightarrow \Omega$, consider the distribution D_σ defined as $D_\sigma(\sigma(i)) = D(i)$ (equivalently, $D_\sigma(i) = D(\sigma^{-1}(i))$) for each $i \in \Omega$. If for every distribution D in \mathcal{P} , D_σ is also in \mathcal{P} for every permutation σ , then the property \mathcal{P} is said to be *label-invariant*.

Definition 2.3 (Distance between two distributions). The distance between two distributions D_1 and D_2 over Ω is the standard ℓ_1 distance between them, which is defined as $\|D_1 - D_2\|_1 :=$

⁵To be precise, the exact lower bounds for non-tolerant uniformity testing is $\Omega(\sqrt{n})$, and for tolerant uniformity testing it is $\Omega(\frac{n}{\log n})$, where n is the support size of the distribution and the proximity parameter ε is constant.

⁶We put this restriction to avoid formalism issues. In particular, the investigated distribution properties that we know of (such as monotonicity and being a k-histogram) are topologically closed.

$\sum_{i \in \Omega} |D_1(i) - D_2(i)|$. For $\eta \in [0, 2]$, D_1 and D_2 are said to be η -close to each other if $\|D_1 - D_2\|_1 \leq \eta$. Similarly, for $\varepsilon \in [0, 2]$, D_1 and D_2 are said to be ε -far from each other if $\|D_1 - D_2\|_1 \geq \varepsilon$.

Definition 2.4 (Distance of a distribution from a property). The distance of a distribution D from a property \mathcal{P} is the minimum ℓ_1 -distance between D and any distribution in \mathcal{P} . For $\eta \in [0, 2]$, a distribution D is said to be η -close to \mathcal{P} if the distance of D from \mathcal{P} is at most η . Analogously, for $\varepsilon \in [0, 2]$, a distribution D is said to be ε -far from \mathcal{P} if the distance of D from \mathcal{P} is at least ε .

Definition 2.5 ((η, ε) -tester). An (η, ε) -tester for a distribution property is a randomized algorithm that has sample access to the unknown distribution (upon query it can receive elements of Ω , each drawn according to the unknown distribution, independently of any previous query or the algorithm's private coins), and distinguishes whether the distribution is η -close to the property or ε -far from the property, with probability at least $\frac{2}{3}$, where η and ε are proximity parameters such that $0 \leq \eta < \varepsilon \leq 2$. The tester is said to be *tolerant* when $\eta > 0$, and *non-tolerant* when $\eta = 0$.

Now we define the notions of non-concentrated distributions and non-concentrated properties.

Definition 2.6 (Non-Concentrated distribution). A distribution D over the domain $\Omega = [n]$ is said to be (α, β) -non-concentrated if for any set $S \subseteq \Omega$ with size βn , the probability mass on S is at least α , where α and β are two parameters such that $0 < \alpha \leq \beta < \frac{1}{2}$.

Definition 2.7 (Non-Concentrated property). Let $0 < \alpha \leq \beta < \frac{1}{2}$. A distribution property \mathcal{P} is defined to be (α, β) -non-concentrated, if all distributions in \mathcal{P} are (α, β) -non-concentrated.

Note that the uniform distribution is (α, α) -non-concentrated for every α , and so is the property of being identical to the uniform distribution. Also, for any $0 < \alpha < \frac{1}{2}$ such that αn is an integer, the uniform distribution is the only (α, α) -non-concentrated one. Finally, observe that any arbitrary distribution is both $(0, \beta)$ -non-concentrated and $(\alpha, 1)$ -non-concentrated, for any $\alpha, \beta \in (0, 1)$.

2.1 Useful concentration bounds

Now we state some large deviation results that will be used throughout the paper.

Lemma 2.8 (Multiplicative Chernoff bound [DP09]). Let X_1, \dots, X_n be independent random variables such that $X_i \in [0, 1]$. For $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X]$, the following holds for any $0 \leq \delta \leq 1$.

$$\mathbb{P}(|X - \mu| \geq \delta\mu) \leq 2 \exp\left(-\frac{\mu\delta^2}{3}\right)$$

Lemma 2.9 (Additive Chernoff bound [DP09]). Let X_1, \dots, X_n be independent random variables such that $X_i \in [0, 1]$. For $X = \sum_{i=1}^n X_i$ and $\mu_l \leq \mathbb{E}[X] \leq \mu_h$, the following hold for any $\delta > 0$.

$$(i) \quad \mathbb{P}(X \geq \mu_h + \delta) \leq \exp\left(-\frac{2\delta^2}{n}\right).$$

$$(ii) \quad \mathbb{P}(X \leq \mu_l - \delta) \leq \exp\left(-\frac{2\delta^2}{n}\right).$$

3 Overview of our results

In this section, we give an overview of our results as follows:

3.1 Overview of proof of Theorem 1.1 and Theorem 1.2

We first show that for any label-invariant distribution property, the sample complexities of tolerant and non-tolerant testing are separated by at most a quadratic factor, ignoring poly-logarithmic terms. More specifically, in Theorem 1.1, we prove that for any label-invariant distribution property \mathcal{P} that has a non-tolerant tester with sample complexity Λ , there exists a tolerant tester for \mathcal{P} that uses $\tilde{\mathcal{O}}(\Lambda^2)$ samples, ignoring poly-logarithmic factors. Since we can learn a distribution using $\mathcal{O}(n)$ samples, our proof is particularly useful when $\Lambda = o(\sqrt{n})$, where n is the size of the support of the distribution that is being tested.

To prove Theorem 1.1 (restated as Theorem 4.1), we provide an algorithm for tolerant testing of \mathcal{P} with sample complexity $\tilde{\mathcal{O}}(\Lambda^2)$, based on the existence of a non-tolerant tester of \mathcal{P} with sample complexity $\mathcal{O}(\Lambda)$. Given the existence of such a non-tolerant tester with sample complexity $\mathcal{O}(\Lambda)$, one crucial observation that we use here is that there cannot be two distributions D_1 and D_2 that are identical on the elements with mass $\Omega(\frac{1}{\Lambda^2})$ (we call them *high* elements), where D_1 is in the property \mathcal{P} while D_2 is far from \mathcal{P} . This is formally stated as Lemma 4.3.

Given that the two distributions D_1 and D_2 are identical on all elements with mass $\Omega(\frac{1}{\Lambda^2})$, by the birthday paradox, we can say that $\mathcal{O}(\Lambda)$ samples are not enough to obtain any *low* elements, that is, elements with mass $o(\frac{1}{\Lambda^2})$, that appear more than once. Since the property \mathcal{P} is label-invariant, we can apply uniformly random permutations over the low elements of both D_1 and D_2 , making the samples obtained from both D_1 and D_2 appear as two uniformly random sequences. Thus, from the view of any tester that takes only $\mathcal{O}(\Lambda)$ samples, D_1 and D_2 will appear the same, which would contradict the existence of a non-tolerant tester that distinguishes D_1 from D_2 using $\mathcal{O}(\Lambda)$ samples. At this point, we would like to point out that the proof of Lemma 4.3 only assumes the existence of a non-tolerant tester, and is oblivious to its internal details. Later, in Lemma 4.4, we generalize this idea to show that when D_1 and D_2 are close with respect to the high elements, it cannot be the case that D_1 is in the property \mathcal{P} , while D_2 is far from \mathcal{P} . Although the proof follows a similar line to that of Lemma 4.3, more careful analysis is required to prove Lemma 4.4. Note that Lemma 4.4 is the main technical lemma required to prove Theorem 1.1.

Once we have Lemma 4.3 and Lemma 4.4, we can describe the algorithm of Theorem 1.1. Broadly speaking, we show that *partial learning* of the distribution is sufficient for constructing a tolerant tester for any label-invariant property, as opposed to the more familiar paradigm of *testing by learning* [DLM⁺07, Ser10]. Using Lemma 4.4, we show that estimating the masses of only the high elements is enough for us, along with the fact that the property \mathcal{P} that we are testing is label-invariant. Roughly, the algorithm has three steps. In the first step, we identify and measure the high elements of the unknown distribution D . In the second step, we construct a new distribution \tilde{D} that adheres to the high mass elements obtained from the first step. Finally, in the third step, we check whether there exists any distribution D_1 in \mathcal{P} that is close to \tilde{D} . If such a distribution exists, we accept, and otherwise we reject. In the first step, we need $\tilde{\mathcal{O}}(\Lambda^2)$ many samples to correctly estimate the masses of the high elements, which dominates the sample complexity of our tolerant tester.

It is important to note that the *computational efficiency* of the tolerant tester depends on how fast we can check whether the distribution \tilde{D} (constructed by the algorithm) is *close* to a known

property \mathcal{P} , where we have the complete description of \tilde{D} . Later, in Theorem 1.2 (restated as Theorem 5.5), we show that when the property \mathcal{P} can be expressed as a feasible solution to a set of linear inequalities, there exists an algorithm that tolerantly tests for \mathcal{P} in time polynomial in the support size of the distribution and the number of linear inequalities required to represent it. The algorithm is similar to that of Theorem 1.1, whereas its polynomial running time follows by using the Ellipsoid method.

3.2 Overview of Theorem 1.3 and Theorem 1.4

In Theorem 1.3, we show that in order to non-tolerantly test any non-concentrated property (defined in Definition 2.7), $\Omega(\sqrt{n})$ samples are required, where n denotes the support size of the distribution. Before directly proceeding to prove the result, as a warm-up, we first show an analogous result for label-invariant non-concentrated properties in Theorem 6.1. To prove the theorem, for any distribution D_{yes} in the label-invariant non-concentrated property \mathcal{P} that we are testing, we construct a new distribution D_{no} that is far from \mathcal{P} , whose support is a subset of the support of D_{yes} . The two distributions are identical over their high probability elements, and they only differ in their *low* probability elements, where a low probability element is an element with mass $\mathcal{O}(\frac{1}{n})$. Since D_{yes} and D_{no} differ only on the elements with mass $\mathcal{O}(\frac{1}{n})$, by the birthday paradox and the fact that the property is label-invariant, any tester that takes $o(\sqrt{n})$ samples cannot distinguish between D_{yes} and D_{no} , and the result follows. We note that the proof of Theorem 6.1 is a generalization of the lower bound proof for uniformity testing.

Though the proof of Theorem 1.3 (restated as Theorem 7.1) follows similarly to that of Theorem 6.1, delicate analysis is required to take care of the fact that the properties are no longer label-invariant. We also discuss briefly the reason why the technique used to prove Theorem 6.1 does not work to prove Theorem 1.3, in the beginning of Section 7.

As a step further, in Theorem 1.4 (restated as Theorem 6.3), we show that $\Omega(n^{1-o(1)})$ samples are necessary to tolerantly test any non-concentrated label-invariant property. This proof follows from an application of the *low frequency blindness theorem* of Valiant [Val11]. The question of tolerant testing of general non-concentrated properties remains open.

3.3 Overview of Theorem 1.5

Finally, we consider the problem of learning a distribution D , where D is *concentrated* over a unknown set $S \subseteq \Omega$. In Theorem 1.5 (restated as Theorem 8.1), we give an algorithm that achieves this with $\mathcal{O}(|S|)$ many samples, even when $|S|$ is also unknown. Note that this problem is reminiscent of the folklore result of learning a distribution over any set S that takes $\mathcal{O}(|S|)$ samples. However, the folklore result holds only for the case where the set S is known ⁷.

Broadly, the algorithm iterates over possible values of $|S|$. Starting from $s = 1$, we first take s many samples from the unknown distribution D , and construct a new empirical distribution D_s based upon the samples obtained. Once we have the distribution D_s , we apply the result of Valiant and Valiant [VV11] to test whether the unknown distribution D is close to the newly constructed distribution D_s , by using number of samples that is slightly smaller than s . If D_s is close to D , we report the distribution D_s as the output and terminate the algorithm. Otherwise, we double the value of s and perform another iteration of the two steps as mentioned above. Finally, we show that

⁷There are also prior results where only $|S|$ is known, such as in the work of Acharya, Diakonikolas, Li and Schmidt [ADLS17].

when $s \geq |S|$, where S is the unknown set on which D is concentrated, D_s will be close to D with high probability, and we will output a distribution satisfying the statement of Theorem 1.5. To the best of our knowledge, this is the first result of a tester of this kind that adapts to an unknown support size $|S|$.

4 Non-tolerant vs. tolerant sample complexities of label-invariant properties (Proof of Theorem 1.1)

We will prove that for any label-invariant property, the sample complexities of tolerant and non-tolerant testing are separated by at most a quadratic factor (ignoring some poly-logarithmic factors). Formally, the result is stated as follows:

Theorem 4.1 (Theorem 1.1 formalized). *Let \mathcal{P} be a label-invariant distribution property, for which there exists an $(0, \varepsilon)$ -tester (non-tolerant tester) with sample complexity $\Lambda(n, \varepsilon)$, where $\Lambda \in \mathbb{N}$ and $0 < \varepsilon \leq 2$. Then for any γ_1, γ_2 with $\gamma_1 < \gamma_2$ and $0 < \gamma_2 + \varepsilon < 2$, there exists a $(\gamma_1, \gamma_2 + \varepsilon)$ -tester (tolerant tester) that has sample complexity $\mathcal{O}\left(\frac{1}{(\gamma_2 - \gamma_1)^3} \cdot \min\{\Lambda^2 \log^2 \Lambda, n\}\right)$, where $\Lambda = \Lambda(n, \varepsilon)$, and n is the size of the support of the distribution.*

Let us assume that D is the unknown distribution and $\Lambda(n, \varepsilon) \geq \Omega\left(\frac{1}{\varepsilon}\right)$ ⁸. First note that if $\Lambda = \Omega(\sqrt{n})$, then we can construct a distribution \widehat{D} such that $\|D - \widehat{D}\|_1 < \frac{\gamma_2 - \gamma_1 + \varepsilon}{2}$, by using $\mathcal{O}\left(\frac{n}{(\gamma_2 - \gamma_1 + \varepsilon)^2}\right)$ samples from D . Thereafter we can report D to be γ_1 -close to the property if and only if \widehat{D} is $\frac{\gamma_2 + \gamma_1 + \varepsilon}{2}$ -close to the property. In what follows, we discuss an algorithm with sample complexity $\widetilde{\mathcal{O}}(\Lambda^2)$ when $\Lambda = o(\sqrt{n})$. Also, we assume that n and Λ are larger than some suitable constant. Otherwise, the theorem becomes trivial.

The idea behind the proof is to classify the elements of Ω with respect to their masses in D into *high* and *low*, as formally defined below in Definition 4.2. We argue that since \mathcal{P} is $(0, \varepsilon)$ -testable using $\Lambda(n, \varepsilon) = \mathcal{O}(q)$ samples, there cannot be two distributions D_1 and D_2 that are identical on all elements whose probability mass is at least $\frac{1}{q^2}$, for $q = \theta(\Lambda)$ (the set \mathbf{High}_{1/q^2} defined below), where $D_1 \in \mathcal{P}$ but D_2 is ε -far from \mathcal{P} . We will formally show this in Lemma 4.3, where we will use the fact that \mathcal{P} is label-invariant. Using Lemma 4.3, we prove Lemma 4.4, that (informally) says that if two distributions are close with respect to the high mass elements, then it is not possible that one distribution is close to \mathcal{P} while the other one is far from it. In our algorithm, we intend to approximate the masses of the set \mathbf{High}_{1/q^2} , and the term Λ^2 in the query complexity of our algorithm corresponds to that.

Definition 4.2. For a distribution D over Ω and $0 < \kappa < 1$, we define

$$\mathbf{High}_\kappa(D) = \{x \in \Omega \mid D(x) \geq \kappa\}$$

Now we define a quantity $q \in \mathbb{N}$ where $q = \Theta(\Lambda)$ ⁹. Assume that c^* is a suitable large constant (independent of Λ) such that, if we take Λ many samples from a distribution, then with probability at least $\frac{3}{4}$, we will not get any sample x whose mass is at most $(\frac{c^*}{\Lambda})^2$ more than once. We define

$$q := \frac{\Lambda}{c^*}. \tag{1}$$

⁸This is a reasonable assumption for any non-trivial property.

⁹Note that q and Λ are of the same order of magnitude. We have introduced q for writing proofs more rigorously.

We will complete the proof of Theorem 4.1 by using the following two lemmas which we will prove later.

Lemma 4.3. *Let \mathcal{P} be a label-invariant property that is $(0, \varepsilon)$ -testable using $\Lambda(n, \varepsilon)$ samples and consider q as defined in Equation 1. Let D_1 and D_2 be two distributions such that $\text{High}_{1/q^2}(D_1) = \text{High}_{1/q^2}(D_2)$, and for all $x \in \text{High}_{1/q^2}(D_1)$, the probability of x is the same for both distributions, that is, $D_1(x) = D_2(x)$. Then it is not possible that D_1 satisfies \mathcal{P} while D_2 is ε -far from satisfying \mathcal{P} .*

Lemma 4.4. *Let \mathcal{P} be a label-invariant property that is $(0, \varepsilon)$ -testable using $\Lambda(n, \varepsilon)$ samples, and consider q as defined in Equation (1). Let \bar{D} and \tilde{D} be two distributions over Ω , where $|\Omega| > 4q^2$, and let H contain the top q^2 elements of \bar{D} . Also, assume that $\left| \tilde{D}(\Omega \setminus H) - \bar{D}(\Omega \setminus H) \right| \leq \gamma$. If*

$$\sum_{x \in H} \left| \bar{D}(x) - \tilde{D}(x) \right| \leq \alpha, \quad (2)$$

then the following hold:

1. If \bar{D} is β -close to \mathcal{P} , then there exists a distribution D_1 in \mathcal{P} such that $\text{High}_{1/q^2}(D_1) \subseteq H$ and

$$\sum_{x \in H} \left| D_1(x) - \tilde{D}(x) \right| + \left| D_1(\Omega \setminus H) - \tilde{D}(\Omega \setminus H) \right| \leq (\alpha + \beta + \gamma). \quad (3)$$

2. If \bar{D} is $(\varepsilon + 3\alpha + \beta + 2\gamma)$ -far from \mathcal{P} and D_1 is a distribution such that $\text{High}_{1/q^2}(D_1) \subseteq H$ and

$$\sum_{x \in H} \left| D_1(x) - \tilde{D}(x) \right| + \left| D_1(\Omega \setminus H) - \tilde{D}(\Omega \setminus H) \right| \leq (\alpha + \beta + \gamma), \quad (4)$$

then the distribution D_1 does not satisfy the property \mathcal{P} .

Using the above two lemmas, we will prove Theorem 4.1 in Section 4.1. We present the proofs of Lemma 4.3 and Lemma 4.4 in Section 4.2.

4.1 Proof of Theorem 4.1

Let D be the unknown distribution that we need to test, and assume that $\zeta = \gamma_1$, $\eta = \gamma_2 - \gamma_1$, and $\eta' = \frac{\eta}{64}$. We now provide a tolerant $(\gamma_1, \gamma_2 + \varepsilon)$ -tester, that is, a $(\zeta, \zeta + \varepsilon + \eta)$ -tester for the property \mathcal{P} , as follows:

1. Draw $W = \mathcal{O}\left(\frac{q^2}{\eta'} \log q\right)$ many samples from the distribution D . Let $S \subseteq \Omega$ be the set of (distinct) samples obtained.
2. Draw additional $\mathcal{O}\left(\frac{W}{\eta'^2} \log W\right)$ many samples Z to estimate the value of $D(x)$ for all $x \in S$ ¹⁰.
3. Construct a set H as the union of S and arbitrary q^2 many elements from $\Omega \setminus (S \cup Z)$.

¹⁰Instead of two sets of random samples (where the first one is to generate the set S and the other one is the multi-set Z), one can work with only one set of random samples. But in that case, the sample complexity becomes $\mathcal{O}(q^2 \log n)$, as opposed to $\mathcal{O}(q^2 \log q)$ that we are going to prove.

4. Define a distribution \tilde{D} such that, for $x \in H$,

$$\tilde{D}(x) = \frac{\# x \text{ in the multi-set } Z}{|Z|}.$$

And for each $x \in \Omega \setminus H$,

$$\tilde{D}(x) = \frac{1 - \sum_{x \in H} \tilde{D}(x)}{|\Omega| - |H|}.$$

5. If there exists a distribution D_1 in \mathcal{P} that satisfies both the following conditions:

(A) $\sum_{x \in H} \left| D_1(x) - \tilde{D}(x) \right| + |D_1(\Omega \setminus H) - \tilde{D}(\Omega \setminus H)| \leq 26\eta' + \zeta.$

(B) $\text{High}_{1/q^2}(D_1) \subseteq H.$

then ACCEPT D .

6. If there does not exist any D_1 in \mathcal{P} that satisfies both Conditions (A) and (B) above, then REJECT D .

Note that Step 5 as mentioned above is not completely constructive in a computational sense. In Section 5, we give a constructive variant of the tester where the property \mathcal{P} can be expressed as a set of linear inequalities. We also give an example of a natural property that can be expressed as a set of linear inequalities.

Sample Complexity. The sample complexity of the tester is $\mathcal{O}(\frac{q^2}{\eta'^3} \log^2 q) = \mathcal{O}(\frac{\Lambda^2 \log^2 \Lambda}{(\gamma_2 - \gamma_1)^3})$, which follows from the above description.

Correctness of the algorithm. The proof of correctness of our algorithm is divided into a sequence of lemmas.

Lemma 4.5. *The set H and the distribution \tilde{D} satisfies the following three properties:*

(i) *With probability at least $1 - \frac{1}{q}$, $\text{High}_{\eta'/q^2}(D) \subseteq S \subseteq H$.*

(ii) *For any $x \in H$, if $D(x) \geq \frac{\eta'}{10W}$, $(1 - \eta')D(x) \leq \tilde{D}(x) \leq (1 + \eta')D(x)$ holds with probability at least $1 - \frac{1}{q^4}$.*

(iii) *For any $x \in \Omega$ with $D(x) \leq \frac{\eta'}{10W}$, either $x \notin H$, or $\tilde{D}(x) \leq (1 + \eta')\frac{\eta'}{10W}$ holds with probability at least $1 - \frac{1}{q^4}$.*

Proof. Let us prove the three parts one by one:

- (i) Consider any $x \in \text{High}_{\eta'/q^2}(D)$, that is, $D(x) \geq \frac{\eta'}{q^2}$. Then the probability that $x \notin H$ is at most $(1 - \frac{\eta'}{q^2})^{|H|} \leq \frac{1}{q^4}$. Applying the union bound over all the elements in $\text{High}_{\eta'/q^2}(D)$ (at most $\frac{q^2}{\eta'} = \mathcal{O}(q^3)$ ¹¹ many elements), the claim follows.

¹¹This follows from the assumption that $\Lambda(n, \epsilon)$ is at least $\Omega(1/\epsilon)$.

- **(ii)** Since $|Z| = \mathcal{O}(\frac{W}{\eta'^2} \log W)$, applying Chernoff bound, we see that $(1 - \eta')D(x) \leq \tilde{D}(x) \leq (1 + \eta')D(x)$ does not hold with probability at most $\frac{1}{q^4}$.
- **(iii)** Since $|Z| = \mathcal{O}(\frac{W}{\eta'^2} \log W)$, if x is in H (otherwise, we are already done), applying Chernoff bound (only on one side), the bound follows.

□

We now bound the ℓ_1 -distance between D and \tilde{D} with respect to H .

Lemma 4.6. $\sum_{x \in H} |D(x) - \tilde{D}(x)| \leq 5\eta'(1 + \eta') \leq 10\eta'$ holds with probability at least $1 - \frac{3}{q}$.

Proof. Recall the definition of $\text{High}_{\eta'/10W}(D)$. Note that

$$\sum_{x \in H} |D(x) - \tilde{D}(x)| = \sum_{x \in \text{High}_{\eta'/10W}(D)} |D(x) - \tilde{D}(x)| + \sum_{x \in H \setminus \text{High}_{\eta'/10W}(D)} |D(x) - \tilde{D}(x)|$$

Applying Lemma 4.5 (ii) for each $x \in \text{High}_{\eta'/10W}(D)$, and then using union bound over all such $x \in \text{High}_{\eta'/10W}(D)$, the first term is bounded by η' with probability at least $1 - \frac{1}{q}$.

Now the second term, notice that for each $x \in H \setminus \text{High}_{\eta'/10W}(D)$, $D(x) \leq \frac{\eta'}{10W}$. By Lemma 4.5 (iii), and using the union bound over all elements in $H \setminus \text{High}_{\eta'/10W}(D)$ (note that $|H| \leq 2W = \mathcal{O}(q^3)$), with probability at least $1 - \frac{2}{q}$, $\tilde{D}(x) \leq \eta'(1 + \eta')/10W$ for all $x \in H \setminus \text{High}_{\eta'/10W}(D)$. Since $|H| \leq 2W$, the second term is bounded by $4\eta'(1 + \eta')$ with probability at least $1 - \frac{2}{q}$. □

Now we prove a lemma that shows that for every distribution D , there is another distribution \bar{D} that is “similar” to D , and for which H contains the top q^2 elements of \bar{D} .

Lemma 4.7. *There exists a distribution \bar{D} such that H contains the top q^2 elements of \bar{D} . Moreover, the following hold:*

- (i) $\|D - \bar{D}\|_1 \leq 2\eta'$, with probability at least $1 - \frac{2}{q}$.
- (ii) $\sum_{x \in H} |\bar{D}(x) - \tilde{D}(x)| \leq 12\eta'$, with probability at least $1 - \frac{5}{q}$.
- (iii) $|\bar{D}(\Omega \setminus H) - \tilde{D}(\Omega \setminus H)| \leq 12\eta'$, with probability at least $1 - \frac{5}{q}$.

Proof. Let T be the set of q^2 largest elements of D . If $T \subseteq S$, H (as $S \subset H$) contains the largest q^2 elements of D . In that case, setting \bar{D} to be D gives us the above results.

Now, let us consider the case where $T \not\subseteq S$. By Lemma 4.5 (part (i)), with probability at least $1 - \frac{2}{q}$, $\text{High}_{\eta'/q^2}(D) \subseteq S$. Thus for any $x \in H \setminus S$, $D(x) < \frac{\eta'}{q^2}$. Consider the set $U = T \setminus H$. Notice that since $|H \setminus S| = q^2$ and $|T| = q^2$, $|U| \leq |H \setminus (T \cup S)|$. Let $U = \{y_1, \dots, y_{|U|}\} \subset \Omega \setminus H$, and let $z_1, \dots, z_{|U|}$ be some $|U|$ elements of $H \setminus (T \cup S)$. Note, by definition of T and U , the set $\{z_1, \dots, z_{|U|}\}$ and the set $\{y_1, \dots, y_{|U|}\}$ are disjoint.

Consider the distribution \bar{D} defined as follows:

- For elements in $\{z_1, \dots, z_{|U|}\}$, we define $\bar{D}(z_i) = D(y_i)$.

- For elements in $\{y_1, \dots, y_{|U|}\}$, we define $\overline{D}(y_i) = D(z_i)$.
- For all other x , we define $\overline{D}(x) = D(x)$.

Note that since all the elements in the sets $\{z_1, \dots, z_{|U|}\}$ and $\{y_1, \dots, y_{|U|}\}$ were from $\Omega \setminus S$, from Lemma 4.5 (part (i)), with probability at least $1 - \frac{2}{q}$, $D(y_i) \leq \frac{\eta'}{q^2}$ and $D(z_i) \leq \frac{\eta'}{q^2}$, for all $i \in \{1, \dots, |U|\}$. Moreover, as $|U| \leq q^2$, we have condition (i) as well. Furthermore, H contains the largest q^2 elements of \overline{D} due to its construction.

Using the triangle inequality (relative to H) along with Lemma 4.6 and the above expression, we can say that, with probability at least $1 - \frac{5}{q}$, (ii) follows.

Let us now prove (iii). Since \overline{D} and \tilde{D} are distributions, $\sum_{x \in H} \overline{D}(x) + \sum_{x \in \Omega \setminus H} \overline{D}(x) = \sum_{x \in H} \tilde{D}(x) + \sum_{x \in \Omega \setminus H} \tilde{D}(x)$. Thus,

$$\left| \overline{D}(\Omega \setminus H) - \tilde{D}(\Omega \setminus H) \right| = \left| \sum_{x \in H} \tilde{D}(x) - \sum_{x \in H} \overline{D}(x) \right| \leq \sum_{x \in H} \left| \tilde{D}(x) - \overline{D}(x) \right| \leq 12\eta'$$

The last inequality follows from (ii). □

Now we finally establish the correctness of the algorithm.

Proof of correctness of the algorithm. For completeness, consider the case where D is ζ -close to \mathcal{P} . By Lemma 4.7 (i) and the triangle inequality, we know that there exists a distribution \overline{D} that is $(\zeta + 2\eta')$ -close to \mathcal{P} and H contain the largest q^2 elements of \overline{D} . Since $\sum_{x \in H} \left| \overline{D}(x) - \tilde{D}(x) \right| \leq 12\eta'$ and $\left| \overline{D}(\Omega \setminus H) - \tilde{D}(\Omega \setminus H) \right| \leq 12\eta'$ hold from Lemma 4.7 (ii) and (iii), following Lemma 4.4 for $\alpha = 12\eta'$, $\beta = \zeta + 2\eta'$ and $\gamma = 12\eta'$, we can say that there exists a distribution D_1 in \mathcal{P} satisfying Equation (3) (which is same as satisfying **Condition (A)** and **Condition (B)** in Step 5 of the algorithm). Hence, our algorithm accepts D in Step 5.

For soundness, consider a distribution D that is $(\varepsilon + \zeta + \eta)$ -far from \mathcal{P} . Then following Lemma 4.7 (i), we know that there exists a distribution \overline{D} that is $(\varepsilon + \zeta + \eta - 2\eta')$ -far from \mathcal{P} , that is, $(\varepsilon + 3\alpha + \beta + 2\gamma)$ -far from \mathcal{P} , where $\alpha = 12\eta'$, $\beta = \zeta + 2\eta'$. Here, we are using that $\eta = 64\eta'$ and $\gamma = 12\eta'$. Also Lemma 4.7 guarantees that H contains the top q^2 elements of \overline{D} . Following Lemma 4.4, we know that there does not exist any such distribution D_1 in \mathcal{P} that satisfies both **Condition (A)** and **Condition (B)** of Step 5 of the algorithm. Thus the algorithm will REJECT the distribution D in Step 6.

Note that the total failure probability of the algorithm is bounded by the probability that Lemma 4.7 does not hold, which is at most $\frac{12}{q}$. □

4.2 Proof of Lemma 4.3 and Lemma 4.4

Proof of Lemma 4.3. We will prove this by contradiction. Let us assume that there are two distributions D_{yes} and D_{no} such that

- $D_{yes} \in \mathcal{P}$;
- D_{no} is ε -far from \mathcal{P} ;

- $\text{High}_{1/q^2}(D_{yes}) = \text{High}_{1/q^2}(D_{no}) = A$;
- For all $x \in A$, $D_{yes}(x) = D_{no}(x)$.

Now, we argue that any $(0, \varepsilon)$ -non-tolerant tester requires more than $\Lambda(n, \varepsilon)$ samples from the unknown distribution D to distinguish whether D is in the property or ε -far from it.

Let D_Y be a distribution obtained from D_{yes} by permuting the labels of $\Omega \setminus A$ using a uniformly random permutation. Specifically, consider a random permutation $\pi : \Omega \setminus A \rightarrow \Omega \setminus A$. The distribution D_Y is as follows:

- $D_Y(x) = D_{yes}(x)$ for each $x \in A$ and
- $D_Y(\pi(x)) = D_{yes}(x)$ for each $x \in \Omega \setminus A$.

Similarly, consider the distribution D_N obtained from D_{no} by permuting the labels of $\Omega \setminus A$ using a uniformly random permutation. Note that D_Y is in \mathcal{P} , whereas D_N is ε -far from \mathcal{P} , which follows from \mathcal{P} being label-invariant.

We will now prove that D_Y and D_N provide similar distributions over sample sequences. More formally, we will prove that any algorithm that takes at most $\Lambda(n, \varepsilon)$ many samples, cannot distinguish D_Y from D_N with probability at least $\frac{2}{3}$. We argue that this claim holds even if the algorithm is provided with additional information about the input: Namely, for all $x \in A$, the algorithm is told the value of $D_Y(x)$ (which is the same as $D_N(x)$). When the algorithm is provided with this information, it can ignore all samples obtained from A .

By the definition of A , for all $x \in \Omega \setminus A$, both $D_Y(x)$ and $D_N(x)$ are at most $\frac{1}{q^2}$. Let S_Y be a sequence of samples drawn according to D_Y . If $|S_Y| \leq \Lambda(n, \varepsilon)$, then with probability at least $\frac{3}{4}$, the sequence $(\Omega \setminus A) \cap S_Y$ has no element that appears twice. In other words, the set $(\Omega \setminus A) \cap S_Y$ is a set of at most $\Lambda(n, \varepsilon)$ distinct elements from $\Omega \setminus A$. Since the elements of $\Omega \setminus A$ were permuted using a uniformly random permutation, with probability at least $\frac{3}{4}$, the sequence $(\Omega \setminus A) \cap S_Y$ is a uniformly random sequence of distinct elements from $\Omega \setminus A$. Similarly, if S_N is a sequence of samples drawn according to D_N , then with probability at least $\frac{3}{4}$, the sequence $(\Omega \setminus A) \cap S_N$ is a uniformly random sequence of distinct elements from $\Omega \setminus A$. Thus, the distributions over the received sample sequence obtained from D_Y or D_N are of distance $\frac{1}{4}$ of each other, which is strictly less than $\frac{1}{3}$.

Hence, if the algorithm obtains at most $\Lambda(n, \varepsilon)$ many samples from the unknown distribution D , it cannot distinguish, with probability at least $\frac{2}{3}$, whether the samples are coming from D_Y or D_N . \square

For the proof of Lemma 4.4, we will need the following simple claim.

Claim 4.8. *Let $\sigma : [n] \rightarrow [n]$ be a permutation and let a_1, a_2, \dots, a_n and b_1, b_2, \dots, b_n be two sets of n positive real numbers. If $a_1 \geq a_2 \geq \dots \geq a_n$ and $b_1 \geq b_2 \geq \dots \geq b_n$, then the sum $\sum_{i \in [n]} |a_i - b_{\sigma(i)}|$ is minimized when σ is the identity permutation.*

Proof. First observe that if a, b, c, d are four real numbers with $a \geq b$ and $c \geq d$, then the following holds:

$$|a - c| + |b - d| \leq |a - d| + |b - c|. \quad (5)$$

The above can be proved by checking all possible orderings of the numbers a, b, c, d .

Once we have the above observation, we can now proceed to prove the claim. Let us consider the set of permutations that minimize $\sum_{i \in [n]} |a_i - b_{\sigma(i)}|$. Let σ be one such minimizing permutation that also minimizes the size for the following set S :

$$S = \{(i, j) : i < j \text{ and } \sigma(i) > \sigma(j)\}$$

Let i be an index such that $\sigma(i) < \sigma(i+1)$ (such an index i exists unless σ is the identity permutation). Let σ' be the permutation obtained from σ by swapping $\sigma(i)$ and $\sigma(i+1)$. Then the sum $\sum_{i \in [n]} |a_i - b_{\sigma'(i)}|$ does not increase from $\sum_{i \in [n]} |a_i - b_{\sigma(i)}|$, because of Equation 5. However, the size of the set S with respect to the permutation σ' strictly decreases, and we have a contradiction. \square

Now we present the proof of Lemma 4.4.

Proof of Lemma 4.4. We consider the two cases separately.

(1) If \bar{D} is β -close to \mathcal{P} , there exists a distribution D_1 in \mathcal{P} such that $\sum_x |\bar{D}(x) - D_1(x)| \leq \beta$. Since \mathcal{P} is label-invariant, any permutation of D_1 is also in \mathcal{P} . Without loss of generality, let us assume that the domain Ω is a subset of $\{1, \dots, n\}$.

By Claim 4.8, the permutation σ that minimizes $\sum_x |\bar{D}(x) - D_1(\sigma(x))| \leq \beta$ is the one that orders the i -th largest element of D_1 with the i -th largest element of \bar{D} , that is, if x is the element with the i -th largest probability mass in D_1 , then $\sigma(x)$ has the i -th largest probability mass in \bar{D} . Consider the distribution D_1^σ that is defined by $D_1^\sigma(x) = D_1(\sigma(x))$. Clearly, H contains the largest q^2 elements of D_1^σ , and hence also $\text{High}_{1/q^2}(D_1^\sigma) \subseteq H$.

As $\sum_{x \in \Omega} |D_1^\sigma(x) - \bar{D}(x)| \leq \beta$, $\sum_{x \in H} |\bar{D}(x) - \tilde{D}(x)| \leq \alpha$ and $|\bar{D}(\Omega \setminus H) - \tilde{D}(\Omega \setminus H)| \leq \gamma$, by the triangle inequality, we obtain

$$\begin{aligned} & \sum_{x \in H} |D_1^\sigma(x) - \tilde{D}(x)| + |D_1^\sigma(\Omega \setminus H) - \tilde{D}(\Omega \setminus H)| \\ \leq & \sum_{x \in H} |D_1^\sigma(x) - \bar{D}(x)| + \sum_{x \in H} |\bar{D}(x) - \tilde{D}(x)| \\ & + |D_1^\sigma(\Omega \setminus H) - \bar{D}(\Omega \setminus H)| + |\bar{D}(\Omega \setminus H) - \tilde{D}(\Omega \setminus H)| \\ \leq & \sum_{x \in H} |D_1^\sigma(x) - \bar{D}(x)| + \sum_{x \in H} |\bar{D}(x) - \tilde{D}(x)| \\ & + \sum_{x \in \Omega \setminus H} |D_1^\sigma(x) - \bar{D}(x)| + |\bar{D}(\Omega \setminus H) - \tilde{D}(\Omega \setminus H)| \\ = & \sum_{x \in \Omega} |D_1^\sigma(x) - \bar{D}(x)| + \sum_{x \in H} |\bar{D}(x) - \tilde{D}(x)| + |\bar{D}(\Omega \setminus H) - \tilde{D}(\Omega \setminus H)| \\ \leq & \alpha + \beta + \gamma \end{aligned}$$

(2) We will prove this case by contradiction. Let $D_1 \in \mathcal{P}$ be a distribution such that $\text{High}_{1/q^2}(D_1) \subseteq H$ and $\sum_{x \in H} |D_1(x) - \tilde{D}(x)| + |D_1(\Omega \setminus H) - \tilde{D}(\Omega \setminus H)| \leq \alpha + \beta + \gamma$. Then, as

$\sum_{x \in H} \left| \overline{D}(x) - \widetilde{D}(x) \right| \leq \alpha$, by the triangle inequality, we have

$$\sum_{x \in H} |D_1(x) - \overline{D}(x)| + \left| D_1(\Omega \setminus H) - \widetilde{D}(\Omega \setminus H) \right| \leq 2\alpha + \beta + \gamma. \quad (6)$$

Consider the distribution \widehat{D} defined as follows:

- For all $x \in H$, $\widehat{D}(x) = D_1(x)$.
- If $D_1(H) \geq \overline{D}(H)$, then for all $x \in \Omega \setminus H$,

$$\widehat{D}(x) = \overline{D}(x) \cdot \phi,$$

where $\phi = \frac{1 - D_1(H)}{1 - \overline{D}(H)}$. Notice that in this case $\phi \leq 1$.

- If $D_1(H) \leq \overline{D}(H)$, then pick the set $T \subset \Omega \setminus H$ with $|T| = 2q^2$ that minimizes $\overline{D}(T)$. Then for all $x \in T$,

$$\widehat{D}(x) = \overline{D}(x) + \frac{\overline{D}(H) - D_1(H)}{2q^2}$$

and for all $x \in \Omega \setminus (T \cup H)$, $\widehat{D}(x) = \overline{D}(x)$

Let us first prove that $\text{High}_{1/q^2}(\widehat{D}) \subseteq H$. In the case where $D_1(H) \geq \overline{D}(H)$, for all $x \in \Omega \setminus H$, $\widehat{D}(x) \leq \overline{D}(x)$. Since $\text{High}_{1/q^2}(\overline{D}) \subseteq H$, $\text{High}_{1/q^2}(\widehat{D}) \subseteq H$. Now, in the case where $D_1(H) \leq \overline{D}(H)$, the only $x \in \Omega \setminus H$ for which $\widehat{D}(x) > \overline{D}(x)$ are those in T . Since $|\Omega| > 4q^2$, the lowest $2q^2$ elements on \overline{D} must each have mass less than $\frac{1}{2q^2}$. So even if we add $\frac{1}{2q^2}$ for any element $x \in T$, $\widehat{D}(x) < 1/q^2$. Hence in this case also $\text{High}_{1/q^2}(\widehat{D}) \subseteq H$ since $\text{High}_{1/q^2}(\overline{D}) \subseteq H$ and $\text{High}_{1/q^2}(D_1) \subseteq H$.

Now let us bound the ℓ_1 distance between \widehat{D} and \overline{D} . Observe that

$$\sum_{x \in \Omega \setminus H} \left| \widehat{D}(x) - \overline{D}(x) \right| = \left| \widehat{D}(\Omega \setminus H) - \overline{D}(\Omega \setminus H) \right|.$$

This is because, in the case where $\widehat{D}(H) \geq \overline{D}(H)$, we have $\widehat{D}(x) = \phi \cdot \overline{D}(x) \leq \overline{D}(x)$ for all $x \in \Omega \setminus H$. On the other hand, in the case where $\widehat{D}(H) \leq \overline{D}(H)$ then for all $x \in \Omega \setminus H$, $\widehat{D}(x) \geq \overline{D}(x)$. Thus,

$$\begin{aligned} \sum_{x \in \Omega \setminus H} \left| \widehat{D}(x) - \overline{D}(x) \right| &= \left| \widehat{D}(\Omega \setminus H) - \overline{D}(\Omega \setminus H) \right| \\ &\leq \left| \widehat{D}(\Omega \setminus H) - \widetilde{D}(\Omega \setminus H) \right| + \left| \overline{D}(\Omega \setminus H) - \widetilde{D}(\Omega \setminus H) \right| \\ &\leq \left| \widehat{D}(\Omega \setminus H) - \widetilde{D}(\Omega \setminus H) \right| + \gamma \end{aligned}$$

Also note that, from the construction of \widehat{D} , we have for all $x \in H$, $\widehat{D}(x) = D_1(x)$ and thus

$\widehat{D}(\Omega \setminus H) = D_1(\Omega \setminus H)$. Thus,

$$\begin{aligned}
\|\widehat{D} - \overline{D}\|_1 &= \sum_{x \in H} |\widehat{D}(x) - \overline{D}(x)| + \sum_{x \in \Omega \setminus H} |\widehat{D}(x) - \overline{D}(x)| \\
&\leq \sum_{x \in H} |\widehat{D}(x) - \overline{D}(x)| + \left| \widehat{D}(\Omega \setminus H) - \overline{D}(\Omega \setminus H) \right| + \gamma \\
&= \left(\sum_{x \in H} |D_1(x) - \overline{D}(x)| + \left| D_1(\Omega \setminus H) - \overline{D}(\Omega \setminus H) \right| \right) + \gamma \\
&\quad \text{(From the construction of } \widehat{D} \text{)} \\
&\leq 2\alpha + \beta + 2\gamma \quad \text{(By Equation (6))}
\end{aligned}$$

Moreover, as $\text{High}_{1/q^2}(D_1) \subseteq H$ and by the construction of \widehat{D} , we have $\text{High}_{1/q^2}(D_1) = \text{High}_{1/q^2}(\widehat{D})$ and for all $x \in \text{High}_{1/q^2}(D_1)$, $D_1(x) = \widehat{D}(x)$. Since we assumed that D_1 is in \mathcal{P} , using Lemma 4.3, \widehat{D} is ε -close to \mathcal{P} . And since $\|\widehat{D} - \overline{D}\|_1 \leq 2\alpha + \beta + 2\gamma$, we conclude that \overline{D} is $(\varepsilon + 2\alpha + \beta + 2\gamma)$ -close to \mathcal{P} , which is a contradiction. \square

5 Computationally efficient tolerant testers

In this section we present a constructive variant of the tolerant tester studied in Section 4.1. Here, for any two vectors $a, b \in \mathbb{R}^N$, we say that $a \leq b$ if $a_i \leq b_i$ holds for every $i \in [N]$. Now let us recall the definitions of *polyhedron* and *projection map*.

Definition 5.1 (Polyhedron). Let A be a $M \times N$ real matrix, $b \in \mathbb{R}^M$ be a real vector and $Ax \leq b$ be a system of linear inequalities. The solution set $\{x \in \mathbb{R}^N \mid Ax \leq b\}$ of the system of inequalities is called a polyhedron. The *complexity* of a polyhedron is defined as MN .

Definition 5.2 (Projection map). Let n be an integer. For all integers $N \geq n$, a *projection map* is denoted as $\pi_n : \mathbb{R}^N \rightarrow \mathbb{R}^n$ and is defined as the projection of the points in \mathbb{R}^N on the first n coordinates.

Before directly proceeding to our results, we first define two variants of distribution properties.

Definition 5.3 (Linear property). Without loss of generality, let us assume that $\Omega = [n]$. A distribution property \mathcal{P} is said to be a *linear property* if there exists a polyhedron $\mathcal{LP} = \{x \in \mathbb{R}^N \mid Ax \leq b\}$, where A is a $M \times N$ real matrix and $b \in \mathbb{R}^M$ be a real vector, and $\pi_n(\mathcal{LP})$ ¹² is the set of distributions satisfying the property \mathcal{P} , that is, for every $z := (z_1, \dots, z_n, \dots, z_N) \in \mathcal{LP}$, the distribution D_z , defined as

$$D_z(i) = z_i, \quad \forall i \in [n]$$

satisfies the property \mathcal{P} , and conversely, for every distribution D that satisfies \mathcal{P} , there exists some $z \in \mathcal{LP}$ such that $D = D_z$ as defined above. The *complexity* of \mathcal{P} is defined as $M \times N$.

¹²Note that $\pi_n(\mathcal{LP})$ will also be a polyhedron in \mathbb{R}^n , see, e.g., Corollary 2.5 in Chapter 2 from the book by Bertsimas and Tsitsiklis [BT97]. However, the number of linear inequalities defining the property, which affects the running time of the tester, can sometimes be greatly reduced by using a projection.

Similar to linear properties, we can also define properties that are feasible solutions to a system of convex constraints.

Definition 5.4 (Convex property). A distribution property \mathcal{P} is said to be a *convex property* if \mathcal{P} is the set of all feasible solutions to a system of convex constraints over $D(i)$ for $i \in \Omega$, where Ω is the sample space of D . In other words, the set \mathcal{P} forms a convex set.

Now we show that some well studied label-invariant distribution properties can be represented as linear or convex properties.

Remark 1 (An example of a linear property: Approximate uniformity property). A distribution D over $[n]$ is said to be uniform if $D(i) = \frac{1}{n}$ for all $i \in [n]$. Let the property $\mathcal{P}_{u,\varepsilon}$ denote the set of all distributions that are ε -close to the uniform distribution, where $\varepsilon \in (0, 1)$ is a parameter. Consider the following polyhedron $\mathcal{LP}_{u,\varepsilon}$ in \mathbb{R}^{2n} :

$$\begin{aligned} \sum_{i \in [n]} z_{n+i} &\leq \varepsilon \\ z_i &\geq 0 && \forall i \in [2n] \\ -z_{n+i} &\leq z_i - \frac{1}{n} \leq z_{n+i} && \forall i \in [n] \end{aligned}$$

Now, observe that $\pi_n(\mathcal{LP}_{u,\varepsilon})$ will give us the set of distributions that are ε -close to uniform, i.e., the set $\mathcal{P}_{u,\varepsilon}$ (this would serve as the linear transformation mentioned in Definition 5.3). Also, note that approximate uniformity property has complexity $\mathcal{O}(n)$.

Now we present a property that can be expressed as a feasible solution to a system of convex inequalities, but that cannot be expressed as feasible solution to a system of linear inequalities.

Remark 2 (An example of a convex property: Entropy property). Let D be a distribution supported on $[n]$. Given a parameter $k \in \mathbb{R}$, let $\mathcal{P}_{E,k}$ denote the set of all distributions with entropy at least k . $\mathcal{P}_{E,k}$ can be expressed as a convex inequality as follows:

$$\sum_{i \in [n]} D(i) \log \frac{1}{D(i)} \geq k.$$

For a distribution property \mathcal{P} , let $\mathcal{CP} \subset \mathbb{R}^n$ denote the *geometric representation* of the set of probability distributions over the set $[n]$ that satisfy \mathcal{P} by considering each distribution over $[n]$ as a point in \mathbb{R}^n .

For all $\beta \in [0, 1]$, $k \leq n$ and $a \in \mathbb{R}^n$, we define the following convex set:

$$\Delta(k, q, a, \beta) := \left\{ x \in \mathbb{R}^d : \sum_{i=1}^k |x_i - a_i| + \left| \sum_{j>k} x_j - \sum_{j>k} a_j \right| \leq \beta, \text{ and } \forall i > k \text{ we have } x_i < \frac{1}{q^2} \right\}$$

Before proceeding to prove Theorem 1.2, we will first prove a more general result. We will show that if we have access to an emptiness oracle that takes a specific convex set (defined in the following statement) as input, and decides whether the convex set is empty or not, then we can design a tolerant tester for any label-invariant distribution property that takes $\tilde{O}(\Lambda^2)$ samples and performs a single emptiness query to the oracle. The result is formally stated and proved below.

Theorem 5.5. *Let \mathcal{P} be a label-invariant distribution property. If there is a $(0, \varepsilon)$ -tester (non-tolerant tester) with sample complexity $\Lambda(n, \varepsilon)$, then for any γ_1, γ_2 with $\gamma_1 < \gamma_2$ and $0 < \gamma_1 < \gamma_2 + \varepsilon < 2$, there exists a $(\gamma_1, \gamma_2 + \varepsilon)$ -tester (tolerant tester) that takes $s = \tilde{\mathcal{O}}(\Lambda^2)$ many samples and makes a single emptiness query to the set $\mathcal{CP} \cap \Delta(\tilde{\mathcal{O}}(s), \Lambda, \tilde{D}, \beta)$, where \tilde{D} is a known probability distribution and $\beta = \gamma_1 + \frac{\gamma_2 - \gamma_1}{3}$.*

Proof. Recall that in Step 5 of the tolerant tester presented in Section 4.1, the tester checks whether there is any distribution $D_1 \in \mathcal{P}$ that satisfies the following two conditions:

$$\sum_{x \in H} \left| D_1(x) - \tilde{D}(x) \right| + \left| D_1(\Omega \setminus H) - \tilde{D}(\Omega \setminus H) \right| \leq 26\eta' + \zeta$$

and

$$\text{High}_{1/q^2}(D_1) \subseteq H$$

where $\zeta = \gamma_1$, $\eta = \gamma_2 - \gamma_1$, $\eta' = \frac{\eta}{64}$ and $\eta' = \frac{\eta}{64}$. The set H and the distribution \tilde{D} are defined in the tolerant tester presented in Section 4.1.

Without loss of generality, we can assume that $H = \{1, \dots, |H|\}$. Therefore, in order to perform Step 5 of the tolerant tester, the following equations are needed to be satisfied:

$$D_1 \in \mathcal{CP} \tag{7}$$

$$D_1 \in \Delta\left(|H|, q, \tilde{D}, 26\eta' + \zeta\right) \tag{8}$$

We now present the tolerant $(\gamma_1, \gamma_2 + \varepsilon)$ -tester in its entirety, that is, a $(\zeta, \zeta + \varepsilon + \eta)$ -tester for the property \mathcal{P} , where $\zeta = \gamma_1$, $\eta = \gamma_2 - \gamma_1$, and $\eta' = \frac{\eta}{64}$.

1. Draw $W = \mathcal{O}\left(\frac{q^2}{\eta'} \log q\right)$ many samples from the distribution D . Let $S \subseteq \Omega$ be the set of (distinct) samples obtained.
2. Draw additional $\mathcal{O}\left(\frac{W}{\eta'^2} \log W\right)$ many samples Z to estimate the value of $D(x)$ for all $x \in S$.
3. Construct a set H as the union of S and q^2 many arbitrary elements from $\Omega \setminus (S \cup Z)$.
4. Define a distribution \tilde{D} such that, for $x \in H$,

$$\tilde{D}(x) = \frac{\# x \text{ in the multi-set } Z}{|Z|}.$$

And for each $x \in \Omega \setminus H$,

$$\tilde{D}(x) = \frac{1 - \sum_{x \in H} \tilde{D}(x)}{|\Omega| - |H|}.$$

5. If there exists a distribution $D_1 \in \mathcal{CP} \cap \Delta\left(|H|, q, \tilde{D}, 26\eta' + \zeta\right)$, then ACCEPT D .
6. If there does not exist any distribution D_1 that passes Step 5, then REJECT D .

Observe that the sample complexity of the tester is $\mathcal{O}\left(\frac{q^2}{\eta'} \log^2 q\right) = \tilde{\mathcal{O}}(\Lambda^2)$ in addition to a single emptiness query to the set $\mathcal{P} \in \mathcal{CP} \cap \Delta\left(|H|, q, \tilde{D}, 26\eta' + \zeta\right)$ in Step 5. The correctness proof of the above tester follows from the correctness argument presented in Section 4.1. \square

5.1 Emptiness checking when \mathcal{P} is a linear property: Proof of Theorem 1.2

Now we proceed to analyze the time complexity of the $(\gamma_1, \gamma_2 + \varepsilon)$ -tester described in Theorem 5.5 when \mathcal{P} is also a linear property. Recall that as \mathcal{P} is a linear property, there exists a polyhedron $\mathcal{LP} = \{x \in \mathbb{R}^N \mid Ax \leq b\}$, where A is a $M \times N$ real matrix and $b \in \mathbb{R}^M$ be a real vector, and $\pi_n(\mathcal{LP})$ is the set of distributions satisfying the property \mathcal{P} . (See Definition 5.3)

Now in Observation 5.6, we show that checking emptiness of $\pi_n(\mathcal{LP}) \cap \Delta(|H|, q, \tilde{D}, 26\eta' + \zeta)$ is equivalent to testing the feasibility of a family of inequalities.

Observation 5.6. Without loss of generality, assume that $H = \{1, \dots, |H|\}$ and $\Omega = \{1, \dots, n\}$. Checking emptiness of $\pi_n(\mathcal{LP}) \cap \Delta(|H|, q, \tilde{D}, 26\eta' + \zeta)$ is equivalent to testing the feasibility of the following set of inequalities:

$$Az \leq b \tag{9}$$

$$\sum_{i=1}^{|H|} \left| z_i - \tilde{D}(i) \right| + \left| \sum_{i=|H|+1}^n z_i - \sum_{i=|H|+1}^n \tilde{D}(i) \right| \leq 26\eta' + \zeta \tag{10}$$

$$z_i < \frac{1}{q^2} \quad \forall i \in [n] \setminus \{1, \dots, |H|\} \tag{11}$$

Note that the inequality in Equation (10) can be expressed as the following set of linear inequalities using slack variables z_{N+i} for all $i \in [|H| + 1]$:

$$\begin{aligned} \sum_{i=1}^{|H|} z_{N+i} + z_{N+|H|+1} &\leq 26\eta' + \zeta \\ z_{N+i} &\geq 0 && \forall i \in [|H| + 1] \\ -z_{N+i} &\leq z_i - \tilde{D}(i) \leq z_{N+i} && \forall i \in [|H|] \\ -z_{N+|H|+1} &\leq \sum_{i=|H|+1}^n z_i - \sum_{i=|H|+1}^n \tilde{D}(i) \leq z_{N+|H|+1} \end{aligned}$$

Therefore checking the emptiness of $\pi_n(\mathcal{LP}) \cap \Delta(|H|, q, \tilde{D}, 26\eta' + \zeta)$ is equivalent to checking the feasibility of the following set of linear inequalities:

$$\begin{aligned} Az &\leq b \\ \sum_{i=1}^{|H|} z_{N+i} + z_{N+|H|+1} &\leq 26\eta' + \zeta \\ z_{N+i} &\geq 0 && \forall i \in [|H| + 1] \\ -z_{N+i} &\leq z_i - \tilde{D}(i) \leq z_{N+i} && \forall i \in [|H|] \\ -z_{N+|H|+1} &\leq \sum_{i=|H|+1}^n z_i - \sum_{i=|H|+1}^n \tilde{D}(i) \leq z_{N+|H|+1} \\ z_i &< \frac{1}{q^2} && \forall i \in [n] \setminus \{1, \dots, |H|\} \end{aligned}$$

The feasibility of the above set of linear inequalities can be solved in a polynomial time in the complexity of the polyhedron, that is, in a polynomial time in N and M , using the Ellipsoid Method, where recall that A is a $M \times N$ real matrix (see, e.g., [BT97, GM07]). Thus, we have an efficient $(\gamma_1, \gamma_2 + \varepsilon)$ -tester for \mathcal{P} , that runs in time polynomial in the complexity of the label-invariant linear property \mathcal{P} . This concludes the proof of Theorem 1.2.

6 Sample complexity of testing non-concentrated label-invariant properties

In this section we first prove a lower bound of $\Omega(\sqrt{n})$ on the sample complexity of non-tolerant testing of any non-concentrated label-invariant property. Then we proceed to prove a tolerant lower bound of $\Omega(n^{1-o(1)})$ samples for such properties in Section 6.2.

6.1 Non-tolerant lower bound (Proof of Theorem 1.3 for label-invariant properties)

Here we first prove a lower bound result analogous to Theorem 1.3, where the properties are non-concentrated and label-invariant. In Section 7, we discuss why the proof of Theorem 6.1 does not directly work for Theorem 1.3, and then prove Theorem 1.3 using a different argument.

Theorem 6.1 (Analogous result of Theorem 1.3 for non-concentrated label-invariant properties). *Let \mathcal{P} be any (α, β) -non-concentrated label-invariant distribution property, where $0 < \alpha \leq \beta < \frac{1}{2}$. For ε with $0 < \varepsilon < \alpha$, any $(0, \varepsilon)$ -tester for property \mathcal{P} requires $\Omega(\sqrt{n})$ many samples, where n is the size of the support of the distribution.*

Proof. Let us first consider a distribution D_{yes} that satisfies the property. Since \mathcal{P} is an (α, β) -non-concentrated property, by Definition 2.7, D_{yes} is an (α, β) -non-concentrated distribution. From D_{yes} , we generate a distribution D_{no} such that the support of D_{no} is a subset of that of D_{yes} , and D_{no} is ε -far from \mathcal{P} . Hence, if we apply a random permutation over the elements of Ω , we show that D_{yes} and D_{no} are indistinguishable, unless we query for $\Omega(\sqrt{n})$ many samples. Below we formally prove this idea.

We will partition the domain Ω into two parts, depending on the probability mass of D_{yes} on the elements of Ω . Given the distribution D_{yes} , let us first order the elements of Ω according to their probability masses. In this ordering, let L be the smallest $2\beta n$ elements of Ω . We denote $\Omega \setminus L$ by H . Before proceeding further, note that the following observation gives an upper bound on the probabilities of the elements in L .

Observation 6.2. For all $x \in L$, $D_{yes}(x) \leq \frac{1-2\alpha}{1-2\beta} \frac{1}{n}$.

Proof of Observation 6.2. By contradiction, assume that there exists $x \in L$ such that $D_{yes}(x) > \frac{1-2\alpha}{1-2\beta} \frac{1}{n}$. This implies, for every $y \in H$, that $D_{yes}(y) > \frac{1-2\alpha}{1-2\beta} \frac{1}{n}$. So,

$$1 = \sum_{x \in \Omega} D_{yes}(x) = \sum_{x \in L} D_{yes}(x) + \sum_{y \in H} D_{yes}(y) > D_{yes}(L) + |H| \frac{1-2\alpha}{1-2\beta} \frac{1}{n}.$$

As $|L| = 2\beta n$ and D_{yes} is an (α, β) -non-concentrated distribution, $D_{yes}(L) \geq 2\alpha$. Also, $|H| = (1-2\beta)n$. Plugging these into the above inequality, we get a contradiction. \square

Note that Observation 6.2 implies that if S is a multi-set of $o\left(\sqrt{\frac{1-2\beta}{1-2\alpha}}n\right)$ samples from D_{yes} , then with probability $1 - o(1)$, no element from L appears in S more than once. Now using the distribution D_{yes} and the set L , let us define a distribution D_{no} such that D_{no} is ε -far from \mathcal{P} . Note that D_{no} is a distribution that comes from a distribution over a set of distributions, all of which are not (α, β) -non-concentrated. The distribution D_{no} is generated using the following random process:

- We partition L randomly into two equal sets of size βn . Let the sets be $\{x_1, \dots, x_{\beta n}\}$ and $\{y_1, \dots, y_{\beta n}\}$. We pair the elements of L randomly into βn pairs. Let $(x_1, y_1), \dots, (x_{\beta n}, y_{\beta n})$ be a random pairing of the elements in L , which is represented as P_L , that is, $P_L = \{(x_1, y_1), \dots, (x_{\beta n}, y_{\beta n})\}$.
- The probability mass of D_{no} at z is defined as follows:
 - If $z \notin L$, then $D_{no}(z) = D_{yes}(z)$.
 - For every pair $(x_i, y_i) \in P_L$, $D_{no}(x_i) = D_{yes}(x_i) + D_{yes}(y_i)$, and $D_{no}(y_i) = 0$.

We start by observing that the distribution D_{no} constructed above is supported on a set of at most $(1 - \beta)n$ elements. So, any distribution D_{no} constructed using the above procedure is ε -far from satisfying the property \mathcal{P} for any $\varepsilon < \alpha$.

We will now prove that D_{yes} and D_{no} both have similar distributions over the sequences of samples. More formally, we will prove that any algorithm that takes $o(\sqrt{n})$ many samples, cannot distinguish between D_{yes} from D_{no} with probability at least $\frac{2}{3}$.

Since any D_{no} produced using the above procedure has exactly the same probability mass on the elements in H as D_{yes} , any tester that distinguishes between D_{yes} and D_{no} must rely on samples obtained from L . Recall that the algorithm is given a uniformly random permutation of the distribution. Since $\text{SUPP}(D_{no}) \subset \text{SUPP}(D_{yes})$ (particularly, $\text{SUPP}(D_{no}) \cap L \subset \text{SUPP}(D_{yes}) \cap L$), it is not possible to distinguish between D_{yes} and D_{no} , unless an element of L appears at least twice. Otherwise, as in the proof of Lemma 4.3, the elements drawn from L are distributed identically to a uniformly random non-repeating sequence. But observe that $D_{yes}(i) = \mathcal{O}(\frac{1}{n})$ and $D_{no}(i) = \mathcal{O}(\frac{1}{n})$ when i is in L . Thus any sequence of $o(\sqrt{n})$ samples will provide only a distance of $o(1)$ between the two distributions, completing the proof. \square

6.2 Tolerant lower bound (Proof of Theorem 1.4)

Theorem 6.3 (Theorem 1.4 formalized). *Let \mathcal{P} be any (α, β) -non-concentrated label-invariant distribution property, where $0 < \alpha \leq \beta < \frac{1}{2}$. For any constant ε_1 and ε_2 with $0 < \varepsilon_1 < \varepsilon_2 < \alpha$, any $(\varepsilon_1, \varepsilon_2)$ -tester for \mathcal{P} requires $\Omega(n^{1-o(1)})$ samples, where n is the size of the support of the distribution.*

To prove the above theorem, we recall some notions and a theorem from Valiant's paper on a lower bound for the sample complexity of tolerant testing of symmetric properties [Val11]. These definitions refer to invariants of distributions, which are essentially a generalization of properties.

Definition 6.4. Let $\Pi : \mathcal{D}_n \rightarrow \mathbb{R}$ denote a real-valued function over the set \mathcal{D}_n of all distributions over $[n]$.

1. Π is said to be *label-invariant* if for any $D \in \mathcal{D}_n$ the following holds: $\Pi(D) = \Pi(D_\sigma)$ for any permutation $\sigma : [n] \rightarrow [n]$.

2. For any γ, δ with $\gamma \geq 0$ and $\delta \in [0, 2]$, Π is said to be (γ, δ) -weakly-continuous if for all distributions p^+, p^- satisfying $\|p^+ - p^-\|_1 \leq \delta$, we have $|\Pi(p^+) - \Pi(p^-)| \leq \gamma$.

For a property \mathcal{P} of distributions, we define $\Pi_{\mathcal{P}} : \mathcal{D}_n \rightarrow [0, 2]$ with respect to property \mathcal{P} as follows:

$$\text{For } D \in \mathcal{D}_n, \Pi_{\mathcal{P}}(D) := \text{the distance of } D \text{ from } \mathcal{P}.$$

From the triangle inequality property of ℓ_1 distances, $\Pi_{\mathcal{P}}$ (which refers to the distance function from the property \mathcal{P}) is (γ, γ) -weakly continuous, for any $\gamma \in [0, 2]$.

Theorem 6.5 (Low Frequency Blindness [Val11]). *Consider a function $\Pi : \mathcal{D}_n \rightarrow \mathbb{R}$ that is label-invariant and (γ, δ) -weakly-continuous, where $\gamma \geq 0$ and $\delta \in [0, 2]$. Let there exist two distributions p^+ and p^- in \mathcal{D}_n with n being the size of their supports, such that $\Pi(p^+) > b$, $\Pi(p^-) < a$, and they are identical for any index occurring with probability at least $\frac{1}{n}$ in either distribution, where $a, b \in \mathbb{R}$. Then any tester that has sample access to an unknown distribution D and distinguishes between $\Pi(D) > b - \gamma$ and $\Pi(D) < a + \gamma$, requires $\Omega(n^{1-o_\delta(1)})$ many samples from D ¹³.*

Note that in Theorem 6.5, we have assumed that p^+ and p^- are identical for any index that has probability mass at least $\frac{1}{n}$. We can actually replace this condition to $\mathcal{O}(\frac{1}{n})$ by adding $\mathcal{O}(n)$ many “dummy elements” to the support of p^+ and p^- . Now we are ready to prove Theorem 6.3.

Proof of Theorem 6.3. Consider $\Pi_{\mathcal{P}}$ as defined above. As \mathcal{P} is a label-invariant property, the function $\Pi_{\mathcal{P}}$ is also label-invariant. We have already noted that $\Pi_{\mathcal{P}}$ is (γ, γ) -weakly continuous as “distance from a property” satisfies the triangle inequality, for any $\gamma \in [0, 2]$. Now recall that the distributions D_{yes} and D_{no} considered in the proof of Theorem 6.1. The probability mass of each element in the support of D_{yes} and D_{no} is $\mathcal{O}(\frac{1}{n})$. Note that D_{yes} is in \mathcal{P} and D_{no} is ε -far from \mathcal{P} , for any $\varepsilon < \alpha$, and both of them have a support size of $\Theta(n)$. Here we take $\varepsilon > \varepsilon_2$. Now, we apply Theorem 6.5 with $a = 0$, some $b < \varepsilon$ and γ with $\gamma < \min\{\varepsilon_1, \varepsilon - \varepsilon_2\}$. Observe that this completes the proof of Theorem 6.3. \square

7 Sample complexity of non-concentrated properties (Proof of Theorem 1.3)

Theorem 7.1 (Theorem 1.3 formalized). *Let \mathcal{P} be any (α, β) -non-concentrated distribution property for $0 < \alpha < \beta < \frac{1}{2}$. For any ε with $0 < \varepsilon < \alpha$, any $(0, \varepsilon)$ -tester for \mathcal{P} requires $\Omega(\sqrt{n})$ many samples, where n is the size of the support of the distribution.*

Why does the proof of Theorem 6.1 work only for label-invariant properties? Note that the proof of Theorem 6.1 crucially uses the fact that the property \mathcal{P} is label-invariant. Recall that, while constructing D_{no} from D_{yes} , for each $i \in [\beta n]$, moving the masses of both x_i and y_i in D_{yes} to x_i to produce D_{no} is possible as the property \mathcal{P} is label-invariant. Due to this feature, we can apply a random permutation over Ω , and still the permuted distribution will behave identically with respect to \mathcal{P} . After applying the random permutation, the samples coming from D_{yes} and D_{no} are indistinguishable as long as there are no collisions among the elements in L , which is the case when we take $o(\sqrt{n})$ samples. However, this technique does not work when the property is not

¹³ $o_\delta(\cdot)$ suppresses a term in δ .

label-invariant, as the value of the distribution with respect to \mathcal{P} may not be invariant under the random permutation over Ω . This requires a new argument; although the proof is similar in spirit to the proof of Theorem 6.1, there are some crucial differences, and we present the proof next. In order to prove Theorem 7.1, instead of moving the masses of both x_i and y_i in D_{yes} to x_i to produce D_{no} , we randomly move the sum to either x_i or y_i , with probability proportional to the masses of x_i and y_i .

7.1 Proof of Theorem 7.1

The proof of Theorem 7.1 starts off identically to the proof of Theorem 6.1, but there is a departure in the construction of D_{yes} and D_{no} .

Let us first consider D_{yes} , L and P_L as discussed in the proof of Theorem 6.1, only here we cannot and will not pass D_{yes} through a random permutation. The difference starts from the description of the distribution D_{no} . In fact, D_{no} will be randomly chosen according to a distribution over a set of distributions, all of which are not (α, β) -non-concentrated. The distribution D_{no} is generated using the following random process:

- We partition L arbitrarily into two equal sets of size βn . Let the sets be $\{x_1, \dots, x_{\beta n}\}$ and $\{y_1, \dots, y_{\beta n}\}$. We pair the elements of L arbitrarily into βn pairs. Let $(x_1, y_1), \dots, (x_{\beta n}, y_{\beta n})$ be an arbitrary pairing of the elements in L . Let P_L be the set of pairs. So $P_L = \{(x_1, y_1), \dots, (x_{\beta n}, y_{\beta n})\}$. We refer to x_i and y_i as the elements corresponding to each other with respect to P_L , and denote $\pi(x_i) = y_i$ and $\pi(y_i) = x_i$.
- The probability mass of D_{no} at z is defined as follows:
 - If $z \notin L$, then $D_{no}(z) = D_{yes}(z)$.
 - For every pair $(x_i, y_i) \in P_L$, use independent random coins and
 - * With probability $\frac{D_{yes}(x_i)}{D_{yes}(x_i) + D_{yes}(y_i)}$, set $D_{no}(x_i) = D_{yes}(x_i) + D_{yes}(y_i)$ and $D_{no}(y_i) = 0$.
 - * With the remaining probability, that is, with probability $\frac{D_{yes}(y_i)}{D_{yes}(x_i) + D_{yes}(y_i)}$, set $D_{no}(x_i) = 0$ and $D_{no}(y_i) = D_{yes}(x_i) + D_{yes}(y_i)$.

Observe that any D_{no} constructed by the above procedure is supported over a set of at most $(1 - \beta)n$ elements. So any distribution D_{no} constructed using the above procedure is ε -far from satisfying the property \mathcal{P} , for any $\varepsilon < \alpha$. But since any D_{no} produced using the above procedure has exactly the same probability mass on elements in H as D_{yes} , any tester that distinguishes between D_{yes} and D_{no} must rely on samples obtained from L . However, we can prove that unless we receive two samples from the same pair in L (which occurs with low probability), the sample sequence cannot distinguish D_{yes} from D_{no} .

Note that there is an upper bound of $\mathcal{O}(\frac{1}{n})$ on the probability mass of any element in L . In fact, for any pair $(x_i, y_i) \in P_L$, the total probability mass of the pair is at most $\mathcal{O}(\frac{1}{n})$.

Observation 7.2 (Follows immediately from Observation 6.2). For all pairs $(x_i, y_i) \in P_L$, $D_{no}(x_i) + D_{no}(y_i) \leq 2\frac{1-2\alpha}{1-2\beta} \frac{1}{n}$. Also note that $D_{no}(x_i) + D_{no}(y_i) = D_{yes}(x_i) + D_{yes}(y_i)$ with probability 1 over the construction of D_{no} .

From Observation 7.2, observe that if S is a multi-set of $o\left(\sqrt{\frac{1-2\beta}{1-2\alpha}n}\right)$ samples from D_{yes} , then with probability $1 - o(1)$, no two elements in S (identical or not) are from the same pair in P_L . The same holds for D_{no} as well. Given that no two elements in S are from the same pair in P_L , we will now prove that D_{yes} and D_{no} have the same distributions over sample sequences. This implies that, for a sequence of $o\left(\sqrt{\frac{1-2\beta}{1-2\alpha}n}\right)$ many samples, D_{yes} and D_{no} induce distributions over samples sequences that have $o(1)$ variation distance from each other.

Note that under the condition that at most one element is drawn from any pair $(x_i, y_i) \in P_L$, the probability that the sample is x_i instead of y_i is equal to $\frac{D_{yes}(x_i)}{D_{yes}(x_i)+D_{yes}(y_i)}$, irrespective of whether the distribution is D_{yes} or D_{no} . So, we have the following lemma.

Lemma 7.3. *Let a_1, \dots, a_q be a sequence of q elements, where no element of L appears twice, additionally containing no two elements from the same pair in P_L (elements of H can appear freely). Then*

$$\Pr_{s_1, \dots, s_q \sim D_{yes}} [(s_1, \dots, s_q) = (a_1, \dots, a_q)] = \Pr_{s_1, \dots, s_q \sim D_{no}} [(s_1, \dots, s_q) = (a_1, \dots, a_q)].$$

Proof. Let us begin by defining an event \mathcal{E} as follows:

$$\mathcal{E} := \text{no element of } L \text{ appears twice, and no two elements from the same pair appear.}$$

Observe that we will be done by proving

$$\Pr_{s_1, \dots, s_q \sim D_{yes}} [s_i = a_i \text{ for each } i \in [q] \mid \mathcal{E}] = \Pr_{s_1, \dots, s_q \sim D_{no}} [s_i = a_i \text{ for each } i \in [q] \mid \mathcal{E}]. \quad (12)$$

We will prove this by using induction over q . Let us assume that we have generated samples $s_1 = a_1, \dots, s_k = a_k$ from the unknown distribution, where $1 \leq k < q$. Let $X_k = \{s_1, \dots, s_k\} \cap L$ be the samples we have seen until now from L , and $X'_k = \{\pi(x) : x \in X_k\}$. By the induction hypothesis, assume that Equation (12) holds for each q with $q \leq k$. We will show that Equation (12) holds for $q = k + 1$.

To do so, let us now define two distributions D_{yes}^{k+1} and D_{no}^{k+1} as follows, and consider a claim (Claim 7.4) about them.

$$D_{yes}^{k+1}(x) = \Pr_{s_1, \dots, s_q \sim D_{yes}} [s_{k+1} = x \mid \mathcal{E} \text{ and } s_i = a_i \text{ for } i \leq k].$$

Similarly,

$$D_{no}^{k+1}(x) = \Pr_{s_1, \dots, s_q \sim D_{no}} [s_{k+1} = x \mid \mathcal{E} \text{ and } s_i = a_i \text{ for } i \leq k].$$

Claim 7.4. $D_{yes}^{k+1}(x) = D_{no}^{k+1}(x)$ for every $x \in \Omega$.

Proof. We prove the claim separately when $x \in X_k \cup X'_k \subseteq L$, $x \in L \setminus (X_k \cup X'_k)$, and $x \notin L$.

- (i) $x \in X_k \cup X'_k$: $D_{yes}^{k+1}(x) = D_{no}^{k+1}(x) = 0$. This follows from the condition that no element of L appears twice, additionally containing no two elements of the same pair.
- (ii) $x \in L \setminus (X_k \cup X'_k)$: As $D_{yes}^{k+1}(x) = D_{no}^{k+1}(x) = 0$ for every $x \in X_k \cup X'_k$, we have the followings for each $x \in L \setminus (X_k \cup X'_k)$.

Assume that $x = x_i \in L \setminus (X_k \cup X'_k)$ for some $i \in [\beta n]$ (using the notation defined for the partition of L into pairs while we have described the random process for generating D_{no}). The argument for the case where $x = y_j$ for some $j \in [\beta n]$ is analogous to this.

Under D_{yes} , a direct calculation gives the probability for obtaining $x = x_i \in L \setminus (X_k \cup X'_k)$ as the $(k+1)$ -th sample s_{k+1} .

$$\begin{aligned} D_{yes}^{k+1}(x) &= D_{yes}(x \mid x \notin X_k \cup X'_k) \\ &= \frac{D_{yes}(x)}{1 - \sum_{y \in X_k \cup X'_k} D_{yes}(y)} \\ &= \frac{D_{yes}(x)}{1 - \sum_{y \in X_k} (D_{yes}(y) + D_{yes}(\pi(y)))}, \end{aligned}$$

Let us now consider D_{no} . Note that $x_i \in L \setminus (X_k \cup X'_k)$, and neither x_i nor y_i is present in the set of first k samples $\{s_1, \dots, s_k\}$. So, the probability of getting s_1, \dots, s_k as the sequence of first k samples is completely independent of how $D_{no}(x_i)$ and $D_{no}(y_i)$ are assigned while generating D_{no} , that is, whether we chose $D_{no}(x_i)$ to be $D_{yes}(x_i) + D_{yes}(y_i)$, or chose it to be zero (and made $D_{no}(y_i)$ equal to $D_{yes}(x_i) + D_{yes}(y_i)$ instead). That is, even when conditioned on the event that s_1, \dots, s_k is the sequence of first k samples, the probability that $D_{no}(x_i)$ is $D_{yes}(x_i) + D_{yes}(y_i)$ is $\frac{D_{yes}(x_i)}{D_{yes}(x_i) + D_{yes}(y_i)}$. Note that $D_{no}(x_i)$ is 0 with probability $\frac{D_{yes}(y_i)}{D_{yes}(x_i) + D_{yes}(y_i)}$.

Now we can calculate the probability of obtaining $x = x_i \in L$ as the $(k+1)$ -th sample s_k from the corresponding conditional probabilities.

$$\begin{aligned} D_{no}^{k+1}(x) &= D_{no}(x \mid x \notin X_k \cup X'_k) \\ &= \frac{D_{yes}(x_i) + D_{yes}(y_i)}{1 - \sum_{y \in X_k \cup X'_k} D_{no}(y)} \cdot \frac{D_{yes}(x_i)}{D_{yes}(x_i) + D_{yes}(y_i)} \\ &= \frac{D_{yes}(x)}{1 - \sum_{y \in X_k} (D_{no}(y) + D_{no}(\pi(y)))}. \end{aligned}$$

From the construction of D_{yes} and D_{no} , for each $y \in L$, $D_{yes}(y) + D_{yes}(\pi(y)) = D_{no}(y) + D_{no}(\pi(y))$. As $X_k \subseteq L$,

$$\sum_{y \in X_k} (D_{yes}(y) + D_{yes}(\pi(y))) = \sum_{y \in X_k} (D_{no}(y) + D_{no}(\pi(y))).$$

Hence, we have $D_{yes}^{k+1}(x) = D_{no}^{k+1}(x)$.

- (iii) $x \notin L$: Recall that for any $x \notin L$, $D_{yes}(x) = D_{no}(x)$. Proceeding in similar fashion to $D_{yes}^{k+1}(x)$ in Case (ii), we conclude that $D_{yes}^{k+1}(x) = D_{no}^{k+1}(x)$.

□

Now we are ready to prove Equation (12) for $q = k + 1$.

$$\begin{aligned}
& \Pr_{s_1, \dots, s_{k+1} \sim D_{yes}} [s_i = a_i \text{ for each } i \in [k+1] \mid \mathcal{E}] \\
= & \Pr_{s_1, \dots, s_{k+1} \sim D_{yes}} [s_i = a_i \text{ for each } i \in [k] \mid \mathcal{E}] \cdot \Pr[s_{k+1} = a_{k+1} \mid \mathcal{E} \text{ and } s_i = a_i \text{ for each } i \in [k]] \\
= & \Pr_{s_1, \dots, s_k \sim D_{yes}} [s_i = a_i \text{ for each } i \in [k] \mid \mathcal{E}] \cdot D_{yes}^{k+1}(a_{k+1}) \quad (\text{By the definition of } D_{yes}^{k+1}) \\
= & \Pr_{s_1, \dots, s_k \sim D_{no}} [s_i = a_i \text{ for each } i \in [k] \mid \mathcal{E}] \cdot D_{no}^{k+1}(a_{k+1}) \\
& \quad (\text{By the induction hypothesis and Claim 7.4, respectively}) \\
= & \Pr_{s_1, \dots, s_k \sim D_{no}} [s_i = a_i \text{ for each } i \in [k] \mid \mathcal{E}] \cdot \Pr[s_{k+1} = a_{k+1} \mid \mathcal{E} \text{ and } s_i = a_i \text{ for each } i \in [k]] \\
= & \Pr_{s_1, \dots, s_{k+1} \sim D_{no}} [s_i = a_i \text{ for each } i \in [k+1] \mid \mathcal{E}].
\end{aligned}$$

□

Following the construction of D_{yes} and D_{no} , we know that the two distributions differ only on the elements of L . Moreover, following Observation 7.2, we know that if we take $o\left(\sqrt{\frac{1-2\beta}{1-2\alpha}n}\right)$ many samples, then with probability $1 - o(1)$, neither any element of L will appear more than once nor two elements of same pair in P_L will appear. Under these two conditions, Lemma 7.3 states that D_{yes} and D_{no} will appear to be the same. Thus we can say that any $(0, \epsilon)$ -tester that receives $o\left(\sqrt{\frac{1-2\beta}{1-2\alpha}n}\right)$ samples cannot distinguish between D_{yes} and D_{no} , and obtain Theorem 7.1.

8 Learning a distribution (Proof of Theorem 1.5)

In this section, we prove an upper bound related to the tolerant testing of more general properties. Following a folklore result, when provided with oracle access to an unknown distribution D , we can always construct a distribution D' , such that the ℓ_1 distance between D' and D (the unknown distribution) is at most ϵ , by using $\mathcal{O}\left(\frac{n}{\epsilon^2}\right)$ samples from D ¹⁴. In this section, we provide a procedure that can be used for tolerant testing of properties, and in particular hints at how general tolerance gap bounds could be proved in the future. Our algorithm learns an unknown distribution approximately with high probability, adapting to the input, using as few samples as possible. Specifically, we prove that given a distribution D , if there exists a subset $S \subseteq [n]$ which holds most of the total probability mass of D , then the distribution D can be learnt using $\mathcal{O}(|S|)$ samples, even if the algorithm is unaware of $|S|$ in advance. Our result is formally stated as follows:

Theorem 8.1 (Theorem 1.5 formalized). *Let D denote the unknown distribution over $\Omega = [n]$, and assume that there exists a set $S \subseteq [n]$ with $D(S) \geq 1 - \frac{\eta}{2}$ ¹⁵, where $\eta \in [0, 2)$ is known but S and $|S|$ are unknown. Then there exists an algorithm that takes $\delta \in (0, 2]$ as input and constructs a distribution D' satisfying $\|D - D'\|_1 \leq \eta + \delta$ with probability at least $\frac{2}{3}$. Moreover, the algorithm uses, in expectation, $\mathcal{O}\left(\frac{|S|}{\delta^2}\right)$ many samples from D .*

¹⁴There is a writeup of this folklore result by Cannone [Can20a].

¹⁵Recall that the variation distance between two distribution is half than that of the ℓ_1 distance between them. So, we take $D(S) \geq 1 - \frac{\eta}{2}$ (with $\eta \in [0, 2)$) instead of $D(S) \geq 1 - \eta$ (with $\eta \in [0, 1)$).

Note that in the above theorem, the algorithm has no prior knowledge of $|S|$. Before directly proving the above, we first show that if $|S|$ is known, then $\mathcal{O}(|S|)$ many samples are enough to approximately learn the distribution D . We would like to point out that a similar question has been studied under the local differential privacy model with communication constraints, by Acharya, Kairouz, Liu and Sun [AKLS21] and by Chen, Kairouz and Özgür [CKÖ20].

Lemma 8.2 (Theorem 8.1 when $|S|$ is known). *Let D be the unknown distribution over $\Omega = [n]$ such that there exists a set $S \subseteq [n]$ with $|S| = s$, and $\eta \in [0, 2)$ such that $D(S) \geq 1 - \frac{\eta}{2}$, where $s \in [n]$ and $\eta \in (0, 1)$ are known. Then there exists an algorithm that takes $\delta \in (0, 2]$ as an input and constructs a distribution D' satisfying $\|D - D'\|_1 \leq \eta + \delta$ with probability at least $\frac{9}{10}$. Moreover, the algorithm uses $\mathcal{O}\left(\frac{s}{\delta^2}\right)$ many samples from D .*

We note that Lemma 8.2 can be obtained from the work of Acharya, Diakonikolas, Li and Schmidt [ADLS17, Theorem 2]. For completeness, we give a self-contained proof for this lemma below.

We later adapt the algorithm of Lemma 8.2 to give a proof to the scenario where $|S|$ is unknown, using a guessing technique. The idea is to guess $|S| = s$ starting from $s = 1$, and then to query for $\mathcal{O}(s)$ many samples from the unknown distribution D . From the samples obtained, we construct a distribution D_s , and use Lemma 8.3 presented below to distinguish whether D_s and D are close or far. We argue that, for $s \geq |S|$, D_s will be close to D with probability at least $\frac{9}{10}$. We bound the total probability for the algorithm reporting a D' that is too far from D (for example when terminating before $s \geq |S|$), and also bound the probability of the algorithm not terminating in time when s becomes at least as large as $|S|$.

Lemma 8.3 ([VV11]). *Let D_u and D_k denote unknown (input) and known (given in advance) distributions respectively over $\Omega = [n]$, such that the support of D_u is a set of s elements of $[n]$. Then there exists an algorithm $\text{TOL-ALG}(D_u, D_k, \varepsilon_1, \varepsilon_2, \kappa)$ that takes the full description of D_k , two proximity parameters $\varepsilon_1, \varepsilon_2$ with $0 \leq \varepsilon_1 < \varepsilon_2 \leq 2$ and $\kappa \in (0, 1)$ as inputs, queries $\mathcal{O}\left(\frac{1}{(\varepsilon_2 - \varepsilon_1)^2} \frac{s}{\log s} \log \frac{1}{\kappa}\right)$ many samples from D_u , and distinguishes whether $\|D_u - D_k\|_1 \leq \varepsilon_1$ or $\|D_u - D_k\|_1 \geq \varepsilon_2$ with probability at least $1 - \kappa$ ¹⁶.*

Note that Theorem 8.1 talks about learning a distribution with $\mathcal{O}(|S|)$ samples, where there exists an unknown set S with $D(S) \geq 1 - \eta/2$. To prove Theorem 8.1, we use Lemma 8.3, that crucially uses less than s queries for tolerant identity testing (as opposed to learning).

The original bound following the paper of Valiant and Valiant [VV11] is $\mathcal{O}\left(\frac{1}{(\varepsilon_2 - \varepsilon_1)^2} \frac{n}{\log n}\right)$, which holds for any general distributions D_u and D_k with constant success probability. When deploying Lemma 8.3, we “contract” the set $\Omega \setminus \text{SUPP}(D_k)$ to a single element, which allows us to substitute $s+1$ for n . Note that this does not change the distance between D_k and D_u . Hence, $\mathcal{O}\left(\frac{1}{(\varepsilon_2 - \varepsilon_1)^2} \frac{s}{\log s}\right)$ samples from D_u are enough for constant success probability. Following a recent work of Cannone, Jain, Kamath and Li [CJKL22], the dependence on the proximity parameters can be slightly improved. However we are not using that result since the focus of this work is different.

We first prove Lemma 8.2, and then proceed to prove Theorem 8.1.

Proof of Lemma 8.2. Let Z be a multi-set of $\mathcal{O}\left(\frac{s}{\delta^2}\right)$ samples taken from D . The algorithm constructs

¹⁶The multiplicative factor $\log \frac{1}{\kappa}$ is for amplifying the success probability from $\frac{2}{3}$ to $1 - \kappa$.

a distribution $D' : [n] \rightarrow [0, 1]$ such that

$$D'(x) = \frac{\# \text{ times } x \text{ appears in } Z}{|Z|}.$$

Observe that $\|D - D'\|_1 = 2 \max_{E \subseteq [n]} |D(E) - D'(E)|$. So, we will be done by showing the following:

$$\text{With probability at least } \frac{9}{10}, |D(E) - D'(E)| \leq \frac{\eta + \delta}{2} \text{ for all } E \subseteq [n] \quad (13)$$

Note that there are 2^n possibilities for E . So, a direct application of the union bound would require a failure probability of at most $\mathcal{O}(\frac{1}{2^n})$ for each E not satisfying $|D(E) - D'(E)| \leq \frac{\eta + \delta}{2}$, that is, $\mathcal{O}(\frac{n}{\delta^2})$ samples would be needed. Assuming that D is concentrated ($D(S) \geq 1 - \frac{\eta}{2}$), we argue below that it is enough to have a failure probability of $\mathcal{O}(\frac{1}{2^s})$ for each T not satisfying $|D(T) - D'(T)| \leq \frac{\delta}{4}$, but first we show that this is indeed the probability that we achieve.

Observation 8.4. Consider $T \subseteq [n]$. $|D(T) - D'(T)| \leq \frac{\delta}{4}$ holds with probability at least $1 - \frac{1}{100 \cdot 2^s}$.

Proof. Let X_i denote the binary random variable that takes value 1 if and only if the i -th sample in Z is an element of T , where $i \in [|Z|]$. So, $D'(T) = \frac{1}{|Z|} \sum_{i=1}^{|Z|} X_i$.

Observe that the expectation of $D'(T)$ is $\mathbb{E}[D'(T)] = D(T)$. Applying Chernoff bound (Lemma 2.9), we get the desired result. \square

By the above observation for every subset of S , applying the union bound over all possible subsets of S , we have $|D(T) - D'(T)| \leq \frac{\delta}{4}$ for every $T \subseteq S$ with probability at least $\frac{99}{100}$. Further applying the observation for $T = \Omega \setminus S$, we have $|D(\Omega \setminus S) - D'(\Omega \setminus S)| \leq \frac{\delta}{4}$ with probability at least $1 - \frac{1}{100 \cdot 2^s}$.

Let \mathcal{E} be the event that $|D(T) - D'(T)| \leq \frac{\delta}{4}$ for every $T \subseteq S$, and $|D(\Omega \setminus S) - D'(\Omega \setminus S)| \leq \frac{\delta}{4}$. Note that $\Pr(\mathcal{E}) \geq \frac{9}{10}$. So, to prove Equation (13) and conclude the proof of Lemma 8.2, we show that $|D(E) - D'(E)| \leq \frac{\eta + \delta}{2}$ holds, in the conditional probability space when \mathcal{E} occurs, for any $E \subseteq [n]$.

$$\begin{aligned} |D(E) - D'(E)| &\leq |D(E \cap S) - D'(E \cap S)| + |D(E \cap (\Omega \setminus S)) - D'(E \cap (\Omega \setminus S))| \\ &\leq \frac{\delta}{4} + \max \{D(\Omega \setminus S), D'(\Omega \setminus S)\} \\ &\leq \frac{\delta}{4} + D(\Omega \setminus S) + \frac{\delta}{4} \\ &\leq \frac{\eta + \delta}{2}. \end{aligned}$$

\square

Proof of Theorem 8.1. The algorithm is as follows:

1. Set $s = 1$.
2. Query for a multi-set Z_s of $\mathcal{O}(\frac{s}{\delta^2})$ many samples from D .

3. Construct a distribution $D_s : [n] \rightarrow [0, 1]$ such that

$$D_s(x) = \frac{\# \text{ times } x \text{ appears in } Z_s}{|Z_s|}$$

4. Call the algorithm TOL-ALG $\left(D_s, D, \eta + \frac{\delta}{2}, \eta + \delta, \frac{1}{100 \log^2 s}\right)$ (corresponding to Lemma 8.3) to distinguish whether $\|D - D_s\|_1 \leq \eta + \frac{\delta}{2}$ or $\|D - D_s\|_1 \geq \eta + \delta$. If we get $\|D - D_s\|_1 \leq \eta + \frac{\delta}{2}$ as the output of TOL-ALG, then we report D' as the output and QUIT. Otherwise, we double the value of s . If $s \leq 2n$, go back to Step 2. Otherwise, report FAILURE¹⁷.

Let \mathcal{S} denote the event that the algorithm quits with the desired output. We first show that $\Pr(\mathcal{S}) \geq \frac{2}{3}$. Then we analyze the expected sample complexity of the algorithm.

Observe that the algorithm quits after an iteration with guess s such that ALG-TOL reports $\|D - D_s\|_1 \leq \eta + \frac{\delta}{2}$. So, in that case, the probability that the algorithm exits with an output not satisfying $\|D - D_s\|_1 \leq \eta + \delta$ is at most $\frac{1}{100 \log^2 s}$. When summing this up over all possible s (all powers of k , even up to infinity), the probability that the algorithm does not produce the desired output, given that it quits, is at most $\sum_{k=1}^{\infty} \frac{1}{100k^2} \leq \frac{1}{10}$. So, denoting \mathcal{Q} as the event that the algorithm quits without reporting FAILURE, $\Pr(\mathcal{S} \mid \mathcal{Q}) \geq \frac{9}{10}$.

For the lower bound on $\Pr(\mathcal{Q})$, consider the case where $s \geq |S|$. In this case, $\|D_s - D\|_1 \leq \eta + \frac{\delta}{2}$ with probability at least $\frac{9}{10}$, and TOL-ALG quits by reporting D_s as the output with probability at least $1 - \frac{1}{100 \log^2 s}$. So, for any guess $s \geq |S|$, the algorithm quits and reports the desired output with probability at least $\frac{4}{5}$. So, the probability that the algorithm quits without reporting failure is at least the probability that the algorithm quits with a desired output at some iteration with a guess $s \geq |S|$, which is at least $1 - \left(\frac{1}{5}\right)^{(\log n - \log |S| + 1)}$. That is, $\Pr(\mathcal{Q}) \geq \frac{4}{5}$.

Hence, the success probability of the algorithm can be lower-bounded as

$$\Pr(\mathcal{S}) \geq \Pr(\mathcal{Q}) \cdot \Pr(\mathcal{S} \mid \mathcal{Q}) \geq \frac{9}{10} \cdot \frac{4}{5} > \frac{2}{3}.$$

Now, we analyze the sample complexity of the algorithm. The algorithm queries for $\mathcal{O}(s)$ samples when it runs the iteration whose guess is s . The algorithm goes to the iteration with guess $s > |S|$ if all prior iterations which guessed more than $|S|$ failed, which holds with probability at most $\mathcal{O}\left(\left(\frac{1}{5}\right)^{\lfloor \log s / |S| \rfloor}\right)$. Hence the expected sample complexity of the algorithm is at most

$$\sum_{k:s=2^k < |S|} \mathcal{O}\left(\frac{s}{\delta^2}\right) + \sum_{k:s=2^k \geq |S|} \mathcal{O}\left(\left(\frac{1}{5}\right)^{\lfloor \log(s/|S|) \rfloor} \cdot \frac{s}{\delta^2}\right) = \mathcal{O}\left(\frac{|S|}{\delta^2}\right).$$

To explain the above equality, note that in the LHS of the above equation, each term of the second sum is bounded by $\mathcal{O}\left(\left(\frac{1}{5}\right)^{(k - \log |S|)} \cdot 2^{(k - \log |S|)} \cdot \frac{|S|}{\delta^2}\right)$. Thus, substituting $k - \log(|S|)$ by r , we see that the second part of the LHS is upper bounded by $\sum_{r \geq 0} \mathcal{O}\left(\left(\frac{2}{5}\right)^r \cdot \frac{|S|}{\delta^2}\right)$ which is clearly $\mathcal{O}\left(\frac{|S|}{\delta^2}\right)$. Thus we have the above bound. \square

¹⁷By Lemma 8.3, this step uses $\mathcal{O}\left(\frac{s}{\delta^2}\right)$ samples.

References

- [ABC⁺13] Noga Alon, Eric Blais, Sourav Chakraborty, David García-Soriano, and Arie Matsliah. Nearly tight bounds for testing function isomorphism. *SIAM J. Comput.*, 42(2):459–493, 2013.
- [ACF⁺21] Jayadev Acharya, Clément L. Canonne, Cody Freitag, Ziteng Sun, and Himanshu Tyagi. Inference under information constraints III: local privacy constraints. *IEEE J. Sel. Areas Inf. Theory*, pages 253–267, 2021.
- [ADK15] Jayadev Acharya, Constantinos Daskalakis, and Gautam Kamath. Optimal testing for properties of distributions. In *NIPS*, pages 3591–3599, 2015.
- [ADKR19] Maryam Aliakbarpour, Ilias Diakonikolas, Daniel Kane, and Ronitt Rubinfeld. Private testing of distributions via sample permutations. In *NeurIPS*, pages 10877–10888, 2019.
- [ADLS17] Jayadev Acharya, Ilias Diakonikolas, Jerry Li, and Ludwig Schmidt. Sample-optimal density estimation in nearly-linear time. In *SODA*, pages 1278–1289, 2017.
- [AKLS21] Jayadev Acharya, Peter Kairouz, Yuhan Liu, and Ziteng Sun. Estimating sparse discrete distributions under privacy and communication constraints. In *ALT*, pages 79–98, 2021.
- [BBC⁺10] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Mach. Learn.*, pages 151–175, 2010.
- [BC17] Tugkan Batu and Clément L. Canonne. Generalized uniformity testing. In *FOCS*, pages 880–889, 2017.
- [BCE⁺19] Eric Blais, Clément L. Canonne, Talya Eden, Amit Levi, and Dana Ron. Tolerant junta testing and the connection to submodular optimization and function isomorphism. *ACM Trans. Comput. Theory*, pages 24:1–24:33, 2019.
- [BFF⁺01] Tugkan Batu, Lance Fortnow, Eldar Fischer, Ravi Kumar, Ronitt Rubinfeld, and Patrick White. Testing random variables for independence and identity. In *FOCS*, pages 442–451, 2001.
- [BFR⁺00] Tugkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D. Smith, and Patrick White. Testing that distributions are close. In *FOCS*, pages 259–269, 2000.
- [BT97] Dimitris Bertsimas and John N Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.
- [BY22] Arnab Bhattacharyya and Yuichi Yoshida. *Property Testing - Problems and Techniques*. Springer, 2022.
- [Can20a] Clément L Canonne. A short note on learning discrete distributions. <https://github.com/ccanonne/probabilitydistributiontoolbox/blob/master/learning.pdf>, 2020.

- [Can20b] Clément L. Canonne. *A Survey on Distribution Testing: Your Data is Big. But is it Blue?* Theory of Computing Library, 2020.
- [Can22] Clément L. Canonne. Topics and Techniques in Distribution Testing: A Biased but Representative Sample. March 2022.
- [CDGR18] Clément L. Canonne, Ilias Diakonikolas, Themis Gouleakis, and Ronitt Rubinfeld. Testing shape restrictions of discrete distributions. *Theory Comput. Syst.*, pages 4–62, 2018.
- [CF14] Gregory W Corder and Dale I Foreman. *Nonparametric statistics: A step-by-step approach*. John Wiley & Sons, 2014.
- [CJKL22] Clément L. Canonne, Ayush Jain, Gautam Kamath, and Jerry Li. The price of tolerance in distribution testing. In *COLT*, pages 573–624, 2022.
- [CKÖ20] Wei-Ning Chen, Peter Kairouz, and Ayfer Özgür. Breaking the communication-privacy-accuracy trilemma. In *NeurIPS*, 2020.
- [CT01] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2001.
- [DK16] Ilias Diakonikolas and Daniel M. Kane. A new approach for testing properties of discrete distributions. In *FOCS*, pages 685–694, 2016.
- [DKS17] Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *FOCS*, pages 73–84, 2017.
- [DKS18] Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. Sharp bounds for generalized uniformity testing. In *NeurIPS*, pages 6204–6213, 2018.
- [DKW18] Constantinos Daskalakis, Gautam Kamath, and John Wright. Which distribution distances are sublinearly testable? In *SODA*, pages 2747–2764, 2018.
- [DLM⁺07] Ilias Diakonikolas, Homin K. Lee, Kevin Matulef, Krzysztof Onak, Ronitt Rubinfeld, Rocco A. Servedio, and Andrew Wan. Testing for concise representations. In *FOCS*, pages 549–558, 2007.
- [DP09] Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- [FLV15] Eldar Fischer, Oded Lachish, and Yadu Vasudev. Trading query complexity for sample-based testing and multi-testing scalability. In *FOCS*, pages 1163–1182, 2015.
- [FLV17] Eldar Fischer, Oded Lachish, and Yadu Vasudev. Improving and extending the testing of distributions for shape-restricted properties. In *STACS*, pages 31:1–31:14, 2017.
- [FM08] Eldar Fischer and Arie Matsliah. Testing graph isomorphism. *SIAM J. Comput.*, pages 207–225, 2008.
- [GKK⁺20] Sivakanth Gopi, Gautam Kamath, Janardhan Kulkarni, Aleksandar Nikolov, Zhiwei Steven Wu, and Huanyu Zhang. Locally private hypothesis selection. In *COLT*, pages 1785–1816, 2020.

- [GM07] Bernd Gärtner and Jirí Matousek. *Understanding and using linear programming*. Springer, 2007.
- [Gol17] Oded Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017.
- [Gol19] Oded Goldreich. Testing isomorphism in the bounded-degree graph model. *Electron. Colloquium Comput. Complex.*, page 102, 2019.
- [GR00] Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. *Electron. Colloquium Comput. Complex.*, 7(20), 2000.
- [GR16] Oded Goldreich and Dana Ron. On sample-based testers. *ACM Trans. Comput. Theory*, pages 7:1–7:54, 2016.
- [Kin97] Terry King. *A guide to chi-squared testing*. Taylor & Francis, 1997.
- [Mac03] David J. C. MacKay. *Information theory, inference, and learning algorithms*. Cambridge University Press, 2003.
- [Pan08] Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Trans. Inf. Theory*, pages 4750–4755, 2008.
- [PRR06] Michal Parnas, Dana Ron, and Ronitt Rubinfeld. Tolerant property testing and distance approximation. *J. Comput. Syst. Sci.*, pages 1012–1042, 2006.
- [Ser10] Rocco A. Servedio. Testing by implicit learning: A brief survey. In *Property Testing - Current Research and Surveys*, pages 197–210. 2010.
- [Val11] Paul Valiant. Testing symmetric properties of distributions. *SIAM J. Comput.*, pages 1927–1968, 2011.
- [VV11] Gregory Valiant and Paul Valiant. The power of linear estimators. In *FOCS*, pages 403–412, 2011.
- [VV17a] Gregory Valiant and Paul Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM J. Comput.*, pages 429–455, 2017.
- [VV17b] Gregory Valiant and Paul Valiant. Estimating the unseen: Improved estimators for entropy and other properties. *J. ACM*, pages 37:1–37:41, 2017.
- [Zha21] Huanyu Zhang. Statistical inference in the differential privacy model. *CoRR*, abs/2108.05000, 2021.