



Sample-Optimal Identity Testing with High Probability

Ilias Diakonikolas*
 CS, USC
 diakonik@usc.edu

Themis Gouleakis†
 CSAIL, MIT
 tgoule@mit.edu

John Peebles‡
 CSAIL, MIT
 jpeebles@mit.edu

Eric Price
 CS, UT Austin
 ecprice@cs.utexas.edu

August 30, 2017

Abstract

We study the problem of testing identity against a given distribution (a.k.a. goodness-of-fit) with a focus on the high confidence regime. More precisely, given samples from an unknown distribution p over n elements, an explicitly given distribution q , and parameters $0 < \varepsilon, \delta < 1$, we wish to distinguish, with probability at least $1 - \delta$, whether the distributions are identical versus ε -far in total variation (or statistical) distance. Existing work has focused on the constant confidence regime, i.e., the case that $\delta = \Omega(1)$, for which the sample complexity of identity testing is known to be $\Theta(\sqrt{n}/\varepsilon^2)$.

Typical applications of distribution property testing require small values of the confidence parameter δ (which correspond to small “ p -values” in the statistical hypothesis testing terminology). Prior work achieved arbitrarily small values of δ via black-box amplification, which multiplies the required number of samples by $\Theta(\log(1/\delta))$. We show that this upper bound is suboptimal for any $\delta = o(1)$, and give a new identity tester that achieves the optimal sample complexity. Our new upper and lower bounds show that the optimal sample complexity of identity testing is

$$\Theta\left(\frac{1}{\varepsilon^2} \left(\sqrt{n \log(1/\delta)} + \log(1/\delta)\right)\right)$$

for any n, ε , and δ . For the special case of uniformity testing, where the given distribution is the uniform distribution U_n over the domain, our new tester is surprisingly simple: to test whether $p = U_n$ versus $d_{TV}(p, U_n) \geq \varepsilon$, we simply threshold $d_{TV}(\hat{p}, U_n)$, where \hat{p} is the empirical probability distribution. We believe that our novel analysis techniques may be useful for other distribution testing problems as well.

*Supported by NSF Award CCF-1652862 (CAREER) and a Sloan Research Fellowship.

†Supported by the NSF under Grant No. 1420692.

‡Supported by the NSF Graduate Research Fellowship under Grant No. 1122374, and by the NSF under Grant No. 1065125.

1 Introduction

1.1 Background

Distribution property testing [GR00, BFR⁺00, BFR⁺13] studies the following family of problems: given sample access to one or more unknown distributions, determine whether they satisfy some global property or are “far” from satisfying the property. This broad inference task originates from the field of statistics and has been extensively studied in hypothesis testing [NP33, LR05] with somewhat different formalism. In this work, we study the following standard formalization of this statistical task:

$\mathcal{T}(\mathcal{P}, n, \varepsilon, \delta)$: Given a family of distributions \mathcal{P} with domain D of size n , parameters $0 < \varepsilon, \delta < 1$, and sample access to an unknown distribution p over the same domain, we want to distinguish *with probability at least* $1 - \delta$ between the following cases:

- Completeness: $p \in \mathcal{P}$.
- Soundness: $d_{\text{TV}}(p, \mathcal{P}) \geq \varepsilon$.

We call this the problem of (ε, δ) -testing property \mathcal{P} .

We denote by $d_{\text{TV}}(p, q)$ the total variation distance or statistical distance between distributions p and q , i.e., $d_{\text{TV}}(p, q) \stackrel{\text{def}}{=} (1/2) \cdot \|p - q\|_1$. Similarly, we denote by $d_{\text{TV}}(p, \mathcal{P})$ the minimum total variation distance between p and any $q \in \mathcal{P}$. The goal is to characterize the sample complexity of the problem $\mathcal{T}(\mathcal{P}, n, \varepsilon, \delta)$, i.e., the number of samples that are necessary and sufficient to correctly distinguish between the completeness and soundness cases with high probability.

During the past two decades the above general problem has received significant attention within the computer science community, where the emphasis has been on pinning down the sample complexity of testing *in the constant probability of success regime*—i.e., the setting where $\delta = \Omega(1)$. See [GR00, BFR⁺00, BFF⁺01, Bat01, BDKR02, BKR04, Pan08, Val11, DDS⁺13, ADJ⁺11, LRR11, ILR12, CDVV14, VV14, DKN15b, DKN15a, CDGR16, DKN17, CDKS17, CDS17] for a sample of works, and [Rub12, Can15] for two surveys. The constant confidence regime is fairly well understood: for a range of fundamental properties we now have sample-optimal testers (and matching information-theoretic lower bounds) [Pan08, CDVV14, VV14, DKN15b, DKN15a, ADK15, DK16, DGPP16]. In sharp contrast, as we explain below, the high confidence regime—i.e., the case that $\delta = o(1)$ —is poorly understood even for the most basic properties.

Since testing is a decision problem, any algorithm that succeeds with probability $2/3$ can be used along with standard “BPP amplification” to boost the success probability to any desired accuracy. Specifically, if $S(n, \varepsilon)$ is a sample complexity upper bound for $\mathcal{T}(\mathcal{P}, n, \varepsilon, 1/3)$, this black-box method implies that $S(n, \varepsilon) \cdot \Theta(\log(1/\delta))$ is a sample complexity upper bound for $\mathcal{T}(\mathcal{P}, n, \varepsilon, \delta)$. For essentially all distribution properties studied in the literature, the standard amplification method is the only known way to achieve high confidence probability. This discussion naturally leads us to the following questions: *For a given distribution property \mathcal{P} , does black-box amplification give sample-optimal testers for $\mathcal{T}(\mathcal{P}, n, \varepsilon, \delta)$? Specifically, is the $\Theta(\log(1/\delta))$ multiplicative increase in the sample size the best possible? If not, can we design testers whose sample complexity is information-theoretically optimal, as a function of all parameters, i.e., n, ε , and δ ?*

We believe that this is a fundamental question that merits theoretical investigation in its own right. In most settings, the sample size m is known a priori, i.e., we are presented with a sample of size m and we are not able to draw more samples from the distribution. In such cases, we are interested in characterizing the *optimal tradeoff curve* between the proximity parameter ε and the error probability δ . Previous work in distribution testing characterized the very special case of this tradeoff corresponding to $\delta = \Omega(1)$. We note here that the analogous question in the context of *distribution learning* has been intensely studied in statistics and probability theory (see, e.g., [vdVW96, DL01]) and tight bounds are known in a range of settings.

Moreover, understanding the regime of “small δ ” is of critical practical importance in applications of hypothesis testing (e.g., in biology), as it corresponds to bounding the probabilities of Type I and Type II errors. Perhaps surprisingly, with one exception [HM13], this basic problem has not been previously investigated in the finite sample regime. A conceptual contribution of this work is to raise this problem as a fundamental goal in distribution property testing.

1.2 Our Results

In this paper, we characterize the aforementioned tradeoff for the problem of *identity testing* (a.k.a. goodness of fit). Identity testing corresponds to the property $\mathcal{P} = \{q\}$, where q is an explicitly given distribution on a domain of size n . The important special case that q is the uniform distribution over the domain is referred to as *uniformity testing*. Identity testing is arguably *the* most fundamental distribution testing problem. For this problem, we show that the black-box amplification is suboptimal for *any* $\delta = o(1)$, and give a new identity tester that achieves the optimal sample complexity. Specifically, we prove the following theorem:

Theorem 1 (Main Result). *There exists a computationally efficient (ε, δ) -identity tester for distributions of support size n with sample complexity*

$$\Theta\left(\frac{1}{\varepsilon^2}\left(\sqrt{n \log(1/\delta)} + \log(1/\delta)\right)\right). \quad (1)$$

Moreover, this sample size is information-theoretically optimal, up to a constant factor, for all n, ε, δ .

Since the sample complexity of identity testing is $\Theta(\sqrt{n}/\varepsilon^2)$ for $\delta = 1/3$ [VV14, DKN15b], standard amplification gives a sample upper bound of $\Theta(\sqrt{n} \log(1/\delta)/\varepsilon^2)$ for this problem. It is not hard to observe that this bound cannot be optimal for all values of δ . For example, in the extreme case that $\delta = 2^{-\Theta(n)}$, this gives a sample complexity of $\Theta(n^{3/2}/\varepsilon^2)$. On the other hand, for such values of δ , we can *learn* the underlying distribution (and therefore test for identity) with $O(n/\varepsilon^2)$ samples¹.

The case where $1 \gg \delta \gg 2^{-\Theta(n)}$ is more subtle, and it is not a priori clear how to improve upon naive amplification. Theorem 1 provides a smooth transition between the extremes of $\Theta(\sqrt{n}/\varepsilon^2)$ for constant δ and $\Theta(n/\varepsilon^2)$ for $\delta = 2^{-\Theta(n)}$. It thus provides a quadratic improvement in the dependence on δ over the naive bound for all $\delta \geq 2^{-\Theta(n)}$, and shows that this is the best possible. For $\delta < 2^{-\Theta(n)}$, it turns out that the additive $\Theta(\log(1/\delta)/\varepsilon^2)$ term is necessary, as outlined in Section 1.4, so learning the distribution is optimal.

Our main technical contribution here is to obtain the first sample-optimal *uniformity tester* for the high confidence regime. Our sample-optimal identity tester follows from our uniformity tester by applying the recent result of Goldreich [Gol16], which provides a black-box reduction of identity to uniformity. We also show a matching information-theoretic lower bound on the sample complexity.

The sample-optimal uniformity tester we introduce is remarkably simple: to distinguish between the cases that p is the uniform distribution U_n over n elements versus $d_{TV}(p, U_n) \geq \varepsilon$, we simply compute $d_{TV}(\hat{p}, U_n)$ for the empirical distribution \hat{p} . The tester accepts that $p = U_n$ if the value of this statistic is below some well-chosen threshold, and rejects otherwise. It should be noted that such a tester was not previously known to work with sub-learning sample complexity, even in the constant confidence regime.

Surprisingly, in a literature with several different uniformity testers [GR00, Pan08, VV14, DKN15b], no one previously proposed using the empirical total variation distance. In fact, it would be natural to assume—as was suggested in [BFR⁺00, BFR⁺13]—that such a tester cannot possibly work. A likely reason for this is the following observation: When the sample size m is smaller than the domain size n , the empirical total variation distance is very far from the true distance to uniformity. This suggests that the empirical distance statistic gives little, if any, information in this setting.

¹This follows from the fact that, for any distribution p over n elements, the empirical probability distribution \hat{p}_m obtained after $m = \Omega((n + \log(1/\delta))/\varepsilon^2)$ samples drawn from p is ε -close to p in total variation distance with probability at least $1 - \delta$.

We show in this paper that, in contrast to the above intuition, the natural “plug-in” estimator relying on the empirical distance from uniformity actually works for the following reason: the empirical distance from uniformity is *noticeably smaller* for the uniform distribution than for “far from uniform” distributions, *even with a sub-linear sample size*. Moreover, we obtain the stronger statement that the “plug-in” estimator is a sample-optimal uniformity tester for all parameters n, ε and δ .

1.3 Discussion and Prior Work

Uniformity testing is the first and one of the most well-studied problems in distribution testing [GR00, Pan08, VV14, DKN15b, DGPP16]. As already mentioned, the literature has almost exclusively focused on the case of constant error probability δ . The first uniformity tester, introduced by Goldreich and Ron [GR00], counts the number of collisions among the samples and was shown to work with $O(\sqrt{n}/\varepsilon^4)$ samples [GR00]. A related tester proposed by Paninski [Pan08], which relies on the number of distinct elements in the set of samples, was shown to have the optimal $m = \Theta(\sqrt{n}/\varepsilon^2)$ sample complexity, as long as $m = o(n)$. Recently, a χ^2 -based tester was shown in [VV14, DKN15b] to achieve the optimal $\Theta(\sqrt{n}/\varepsilon^2)$ sample complexity without any restrictions. Finally, the original collision-based tester of [GR00] was very recently shown to also achieve the optimal $\Theta(\sqrt{n}/\varepsilon^2)$ sample complexity [DGPP16]. Thus, the situation for constant values of δ is well understood.

Uniformity is, to our knowledge, the only property for which the high-probability regime has previously been considered. In [HM13], it was shown that the distinct-elements tester of [Pan08] achieves the optimal sample complexity of $m = \Theta(\sqrt{n \log(1/\delta)}/\varepsilon^2)$, as long as $m = o(n)$. We emphasize that for $m = \Omega(n)$, as is the case in many practically relevant settings (see, e.g., the Polish lottery example in [Rub14]), the tester of [Pan08] is known to fail completely even in the constant confidence regime (also see next paragraph).

It is important to note that *all* previously considered uniformity testers [GR00, Pan08, VV14, DKN15b] do not achieve the optimal sample complexity (as a function of all parameters, including δ), and this is *inherent*, i.e., not just a failure of previous analyses. Roughly speaking, since the collision statistic [GR00] and the χ^2 -based statistic [VV14, DKN15b] are not Lipschitz, it can be shown that their high-probability performance is poor. Specifically, in the completeness case ($p = U_n$), if many samples happen to land in the same bucket (domain element), these test statistics become quite large, leading to their suboptimal behavior *for all* $\delta = o(1)$. (For a formal justification, the reader is referred to Section V of [HM13]). On the other hand, the distinct-elements tester [Pan08] does not work *at all* for $m = \omega(n)$. For example, if ε or δ are sufficiently small to necessitate $m \gg n \log n$, then typically all n domain elements will appear in both the completeness and soundness cases, hence the test statistic provides no information.

Interestingly, our empirical total variation distance statistic can be seen to be closely related to the distinct-elements tester of [Pan08] for $m \ll n$. In this regime, we show that the empirical total variation distance can be written as a linear function of the number of elements that are never sampled. In contrast, the distinct-elements statistic of [Pan08] is the number of elements seen exactly once. However, as we prove, our tester is in fact optimal in all parameter settings, whereas the distinct elements tester is not.

The problem of identity testing against an arbitrary (explicitly given) distribution was studied in [BFF⁺01], who gave an $(\varepsilon, 1/3)$ -tester with sample complexity $\tilde{O}(n^{1/2})/\text{poly}(\varepsilon)$. The tight bound of $\Theta(n^{1/2}/\varepsilon^2)$ was first given in [VV14] using a chi-squared type tester (inspired by [CDVV14]). In subsequent work, a similar chi-squared tester that also achieves the same sample complexity bound was given in [ADK15]. (We note that the [VV14, ADK15] testers have sub-optimal sample complexity in the high confidence regime, even for the case of uniformity.) In a related work, [DKN15b] obtained a reduction of identity to uniformity that preserves the sample complexity, up to a constant factor, in the constant error probability regime. More recently, Goldreich [Gol16], building on [DK16], gave a different reduction of identity to uniformity that preserves the error probability. We use the latter reduction in this paper to obtain an optimal identity tester starting from our new optimal uniformity tester.

1.4 Our Techniques

Upper Bound for Uniformity Testing We would like to show that the test statistic $d_{TV}(\hat{p}, U_n)$ is with high probability larger when $d_{TV}(\hat{p}, U_n) \geq \varepsilon$ than when $p = U_n$. We start by showing that among all possible alternative distributions p with $d_{TV}(p, U_n) \geq \varepsilon$, it suffices to consider those in a very simple family. We then show that the test statistic is highly concentrated around its expectation, and that the expectations are significantly different in the two cases.

To simplify the structure of p , we show (Section 4) that if p “majorizes” another distribution q , then the test statistic $d_{TV}(\hat{p}, U_n)$ stochastically dominates $d_{TV}(\hat{q}, U_n)$. (In fact, this statement holds for *any* test statistic that is a convex symmetric function of the empirical histogram.) Therefore, for any p , if we average out the large and small entries of p , the test statistic becomes harder to distinguish from uniform. This lets us reduce to only considering p either U_n or of the form $\frac{1 \pm \varepsilon}{n}$ in each coordinate.

We remark that the aforementioned stochastic domination lemma may also be very useful in rigorous empirical comparisons of test statistics. A major difficulty in empirical studies of distribution testing is that the space of alternative hypotheses is large, and it is therefore impossible to experiment on all such distributions. Our structural lemma reduces the space dramatically for uniformity testing: for any convex symmetric test statistic (which includes all existing ones), the worst case distribution will have αn coordinates of value $(1 + \varepsilon/\alpha)/n$, and the rest of value $(1 - \varepsilon/(1 - \alpha))/n$, for some α . Hence, there are only n possible worst-case distributions for any ε . Notably, this reduction does not lose any absolute constants, so it could be used to identify the non-asymptotic optimal constants for a given set of parameters.

Returning to our uniformity tester, we now need to separate the expectation of the test statistic in our two situations. We achieve this by providing an explicit expression for the Hessian of this expectation, as a function of p . We note that the Hessian is diagonal, and for our two situations of $p_i \approx 1/n$ each entry is within constant factors of the same value, giving a lower bound on its eigenvalues. Since the expectation is minimized at $p = U_n$, we use strong convexity to show the desired expectation gap. Specifically, we prove that this gap is $\varepsilon^2 \cdot \min(m^2/n^2, \sqrt{m/n}, 1/\varepsilon)$.

Finally, we need to show that the test statistic concentrates about its expectation. For $m \geq n$, this follows from McDiarmid’s inequality: since the test statistic is $1/m$ -Lipschitz in the m samples, with probability $1 - \delta$ it lies within $\sqrt{\log(1/\delta)/m}$ of its expectation. When m is larger than the desired sample complexity (1), this is less than the expectation gap above. The concentration is trickier when $m < n$, since the expectation gap is smaller, so we need to establish tighter concentration. We get this by using a Bernstein variant of McDiarmid’s inequality, which is stronger than the standard version of McDiarmid in this context.

Upper Bound for Identity Testing In [Gol16], it was shown how to reduce ε -testing of an arbitrary distribution q over $[n]$ to $\varepsilon/3$ -testing of U_{6n} . This reduction preserves the error probability δ , so applying it gives an identity tester with the same sample complexity as our uniformity tester, up to constant factors.

Sample Complexity Lower Bound To match our upper bound (1), we need two lower bounds. The lower bound of $\Omega(\frac{1}{\varepsilon^2} \log(1/\delta))$ is straightforward from the same lower bound as for distinguishing a fair coin from an ε -biased coin, while the $\sqrt{n \log(1/\delta)}/\varepsilon^2$ bound is more challenging.

For intuition, we start with a $\sqrt{n \log(1/\delta)}$ lower bound for constant ε . When $p = U_n$, the chance that all m samples are distinct is at least $(1 - m/n)^m \approx e^{-m^2/n}$. Hence, if $m \ll \sqrt{n \log(1/\delta)}$, this would happen with probability significantly larger than 2δ . On the other hand, if p is uniform over a random subset of $n/2$ coordinates, the m samples will also all be distinct with probability $(1 - 2m/n)^m > 2\delta$. The two situations thus look the same with 2δ probability, so no tester could have accuracy $1 - \delta$.

This intuition can easily be extended to include a $1/\varepsilon$ dependence, but getting the desired $1/\varepsilon^2$ dependence requires more work. First, we Poissonize the number of samples, so we independently see $\text{Poi}(mp_i)$

samples of each coordinate i ; with exponentially high probability, this Poissonization only affects the sample complexity by constant factors. Then, in the alternative hypothesis, we set each p_i independently at random to be $\frac{1 \pm \varepsilon}{n}$. This has the unfortunate property that p no longer sums to 1, so it is a “pseudo-distribution” rather than an actual distribution. Still, it is exponentially likely to sum to $\Theta(1)$, and using techniques from [WY16, DK16] this is sufficient for our purposes.

At this point, we are considering a situation where the number of times we see each coordinate is either $\text{Poi}(m/n)$ or $\frac{1}{2}(\text{Poi}((1 - \varepsilon)\frac{m}{n}) + \text{Poi}((1 + \varepsilon)\frac{m}{n}))$, and every coordinate is independent of the others. These two distributions have Hellinger distance at least $\varepsilon^2 m/n$ in each coordinate. Then the composition property for Hellinger distance over n independent coordinates implies $m \geq \sqrt{n \log(1/\delta)}/\varepsilon^2$ is necessary for success probability $1 - \delta$.

1.5 Notation and Basic Definitions

We write $[n]$ to denote the set $\{1, \dots, n\}$. We consider discrete distributions over $[n]$, which are functions $p : [n] \rightarrow [0, 1]$ such that $\sum_{i=1}^n p_i = 1$. We use the notation p_i to denote the probability of element i in distribution p . For $S \subseteq [n]$, we will denote $p(S) = \sum_{i \in S} p_i$. We will also sometimes think of p as an n -dimensional vector. We will denote by U_n the uniform distribution over $[n]$.

For $r \geq 1$, the ℓ_r -norm of a distribution is identified with the ℓ_r -norm of the corresponding vector, i.e., $\|p\|_r = (\sum_{i=1}^n |p_i|^r)^{1/r}$. The ℓ_r -distance between distributions p and q is defined as the ℓ_r -norm of the vector of their difference. The total variation distance between distributions p and q is defined as $d_{\text{TV}}(p, q) \stackrel{\text{def}}{=} \max_{S \subseteq [n]} |p(S) - q(S)| = (1/2) \cdot \|p - q\|_1$. The Hellinger distance between p and q is $H(p, q) \stackrel{\text{def}}{=} (1/\sqrt{2}) \cdot \|\sqrt{p} - \sqrt{q}\|_2 = (1/\sqrt{2}) \cdot \sqrt{\sum_{i=1}^n (\sqrt{p_i} - \sqrt{q_i})^2}$. We denote by $\text{Poi}(\lambda)$ the Poisson distribution with parameter λ .

1.6 Structure of this Paper

In Section 2, we formally describe and analyze our sample-optimal uniformity tester. In Section 3, we give our matching sample complexity lower bound. Finally, Section 4 establishes our stochastic domination result that is crucial for the analysis of the soundness in Section 2, and may be useful in the rigorous empirical evaluation of test statistics.

2 Sample-Optimal Uniformity Testing

In this section, we describe and analyze our optimal uniformity tester. Given samples from an unknown distribution p over $[n]$, our tester returns “YES” with probability $1 - \delta$ if $p = U_n$, and “NO” with probability $1 - \delta$ if $d_{\text{TV}}(p, U_n) \geq \varepsilon$.

2.1 Our Test Statistic

We define a very natural statistic that yields a uniformity tester with optimal dependence on the domain size n , the proximity parameter ε , and the error probability δ . Our statistic is a thresholded version of the empirical total variation distance between the unknown distribution p and the uniform distribution. Our tester TEST-UNIFORMITY is described in the following pseudocode:

Algorithm TEST-UNIFORMITY($p, n, \varepsilon, \delta$)

Input: sample access to a distribution p over $[n]$, $\varepsilon > 0$, and $\delta > 0$.

Output: “YES” if $p = U_n$; “NO” if $d_{TV}(p, U_n) \geq \varepsilon$.

1. Draw $m = \Theta\left(\frac{1}{\varepsilon^2} \cdot \left(\sqrt{n \log(1/\delta)} + \log(1/\delta)\right)\right)$ iid samples from p .
2. Let $X = (X_1, X_2, \dots, X_n) \in \mathbb{Z}_{>0}^n$ be the histogram of the samples. That is, X_i is the number of times domain element i appears in the (multi-)set of samples.
3. Define the random variable $S = \frac{1}{2} \sum_{i=1}^n \left| \frac{X_i}{m} - \frac{1}{n} \right|$ and set the threshold

$$t = \mu(U_n) + C \cdot \begin{cases} \varepsilon^2 \cdot \frac{m^2}{n^2} & \text{for } m \leq n \\ \varepsilon^2 \cdot \sqrt{\frac{m}{n}} & \text{for } n < m \leq \frac{n}{\varepsilon^2} \\ \varepsilon & \text{for } \frac{n}{\varepsilon^2} \leq m \end{cases},$$

where C is a universal constant (derived from the analysis of the algorithm), and $\mu(U_n)$ is the expected value of the statistic in the completeness case. (We can compute $\mu(U_n)$ in $O(m)$ time using the procedure in Appendix A.)

4. If $S \geq t$ return “NO”; otherwise, return “YES”.

The main part of this section is devoted to the analysis of TEST-UNIFORMITY, establishing the following theorem:

Theorem 2. *There exists a universal constant $C > 0$ such that the following holds: Given*

$$m \geq C \cdot \frac{1}{\varepsilon^2} \left(\sqrt{n \log(1/\delta)} + \log(1/\delta) \right)$$

samples from an unknown distribution p , Algorithm TEST-UNIFORMITY is an (ε, δ) uniformity tester for p .

As we point out in Appendix A, the value $\mu(U_n)$ can be computed efficiently, hence our overall tester is computationally efficient. To prove correctness of the above tester, we need to show that the expected value of the statistic in the completeness case is sufficiently separated from the expected value in the soundness case, and also that the value of the statistic is highly concentrated around its expectation in both cases. In Section 2.2, we bound from below the difference in the expectation of our statistic in the completeness and soundness cases. In Section 2.3, we prove the desired concentration which completes the proof of Theorem 2.

2.2 Bounding the Expectation Gap

The expectation of the statistic in algorithm TEST-UNIFORMITY can be viewed as a function of the n variables p_1, \dots, p_n . Let $\mu(p) \stackrel{\text{def}}{=} \mathbb{E}[S(X_1, \dots, X_n)]$ be the aforementioned function when the samples are drawn from distribution p .

Our analysis has a number of complications for the following reason: the function $\mu(p) - \mu(U_n)$ is a linear combination of sums that have no indefinite closed form, even if the distribution p assigns only two possible probabilities to the elements of the domain. This statement is made precise in Appendix B. As such, we should only hope to obtain an approximation of this quantity.

A natural approach to try and obtain such an approximation would be to produce separate closed form approximations for $\mu(p)$ and $\mu(U_n)$, and combine these quantities to obtain an approximation for their

difference. However, one should not expect such an approach to work in our context. The reason is that the difference $\mu(p) - \mu(U_n)$ can be much smaller than $\mu(p)$ and $\mu(U_n)$; it can even be arbitrarily small. As such, obtaining separate approximations of $\mu(p)$ and $\mu(U_n)$ to any fixed accuracy would contribute to much error to their difference.

To overcome these difficulties, we introduce the following technique, which is novel in this context. We directly bound from below the difference $\mu(p) - \mu(U_n)$ using *strong convexity*. Specifically, we show that the function μ is strongly convex with appropriate parameters and use this fact to bound the desired expectation gap. The main result of this section is the following lemma:

Lemma 3. *Let p be a distribution over $[n]$ and $\varepsilon = d_{\text{TV}}(p, U_n)$. For all $m \geq 6$ and $n \geq 2$, we have that:*

$$\mu(p) - \mu(U_n) \geq \Theta(1) \cdot \begin{cases} \varepsilon^2 \cdot \frac{m^2}{n^2} & \text{for } m \leq n \\ \varepsilon^2 \cdot \sqrt{\frac{m}{n}} & \text{for } n < m \leq \frac{n}{\varepsilon^2} \\ \varepsilon & \text{for } \frac{n}{\varepsilon^2} \leq m \end{cases} .$$

We note that the bounds in the right hand side above are tight, up to constant factors. Any asymptotic improvement would yield a uniformity tester with sample complexity that violates our tight information-theoretic lower bounds.

The proof of Lemma 3 requires a number of intermediate lemmas. Our starting point is as follows: By the intermediate value theorem, we have the quadratic expansion

$$\mu(p) = \mu(U_n) + \nabla\mu(U_n)^\top(p - U_n) + \frac{1}{2}(p - U_n)^\top H_{p'}(p - U_n) ,$$

where $H_{p'}$ is the Hessian matrix of the function μ at some point p' which lies on the line segment between U_n and p . This expression can be simplified as follows: First, we show (Fact 18) that our μ is minimized over all probability distributions on input U_n . Thus, the gradient $\nabla\mu(U_n)$ must be orthogonal to being a direction in the space of probability distributions. In other words, $\nabla\mu(U_n)$ must be proportional to the all-ones vector. More formally, since μ is symmetric its gradient is a symmetric function, which implies it will be symmetric when given symmetric input. Moreover, $(p - U_n)$ is a direction within the space of probability distributions, and therefore sums to 0, making it orthogonal to the all-ones vector. Thus, we have that $\nabla\mu(U_n)^\top(p - U_n) = 0$, and we obtain

$$\mu(p) - \mu(U_n) = \frac{1}{2}(p - U_n)^\top H_{p'}(p - U_n) \geq \frac{1}{2}\|p - U_n\|_2^2 \cdot \sigma \geq \frac{1}{2}\|p - U_n\|_1^2/n \cdot \sigma , \quad (2)$$

where σ denotes the minimum eigenvalue of the Hessian of μ on the line segment between U_n and p .

The majority of this section is devoted to proving a lower bound for σ . Before doing so, however, we must first address a technical consideration. Because we are considering a function over the space of probability distributions—which is not full-dimensional—the Hessian and gradient of μ with respect to \mathbb{R}^n depend not only on the definition of our statistic S , but also its parameterization. For the purposes of this subsection, we parameterize S as $S(x) = \sum_{i=1}^n \max\{\frac{x_i}{m} - \frac{1}{n}, 0\}$.

In the analysis we are about to perform, it will be helpful to sometimes think of $\frac{1}{n}$ as a free parameter. Thus, we define

$$S_t(x) \triangleq \frac{1}{m} \sum_{i=1}^n \max\{x_i - t, 0\}$$

and

$$\mu_t(p) \triangleq \mathbb{E}_{x \sim \text{Multinomial}(m, p)}[S_t(x)] = \frac{1}{m} \sum_{i=1}^n \sum_{k=t}^m \binom{m}{k} p_i^k (1 - p_i)^{m-k} (k - t) . \quad (3)$$

Note that when $t = m/n$ we have $S_t = S$ and $\mu_t = \mu$. Also note that when we compute the Hessian of $\mu_t(p)$, we are treating $\mu_t(p)$ as a function of p and not of t . In the following lemma, we derive an exact expression for the entries of the Hessian. This result is perhaps surprising in light of the likely nonexistence of a closed form expression for $\mu(p)$. That is, while the expectation $\mu(p)$ may have no closed form, we prove that the Hessian of $\mu(p)$ does in fact have a closed form.

Lemma 4. *The Hessian of $\mu_t(p)$ viewed as a function of p is a diagonal matrix whose i th diagonal entry is given by*

$$h_{ii} = s_{t,i} ,$$

where we define $s_{t,i}$ as follows: Let Δt be the distance of t from the next largest integer, i.e., $\Delta t \triangleq [t] - t$. Then, we have that

$$s_{t,i} = \begin{cases} 0 & \text{for } t = 0 \\ (m-1) \binom{m-2}{t-1} p_i^{t-1} (1-p_i)^{m-t-1} & \text{for } t \in \mathbb{Z}_{>0} \\ \Delta t \cdot s_{[t],i} + (1-\Delta t) \cdot s_{\lceil t \rceil,i} & \text{for } t \geq 0 \text{ and } t \notin \mathbb{Z} \end{cases} .$$

In other words, we will derive the formula for integral $t \geq 1$ and then prove that the value for nonintegral $t \geq 0$ can be found by linearly interpolating between the closest integral values of t .

Proof. Note that because $S_t(x)$ is a separable function of x , $\mu_t(p)$ is a separable function of p , and hence the Hessian of $\mu_t(p)$ is a diagonal matrix. By Equation 3, the i -th diagonal entry of this Hessian can be written explicitly as the following expression:

$$s_{t,i} = \frac{\partial^2}{\partial p_i^2} \mu_t(p) = \frac{d^2}{d p_i^2} \frac{1}{m} \sum_{k=t}^m \binom{m}{k} p_i^k (1-p_i)^{m-k} (k-t) .$$

Notice that if we sum starting from $k = 0$ instead of $k = t$, then the sum equals the expectation of $\text{Bin}(m, p_i)$ minus t . That is, notice that:

$$\frac{d^2}{d p_i^2} \frac{1}{m} \sum_{k=0}^m \binom{m}{k} p_i^k (1-p_i)^{m-k} (k-t) = \frac{d^2}{d p_i^2} \frac{1}{m} (p_i m - t) = 0 .$$

By this observation and the fact that the summand is 0 when $k = t$, we can switch which values of k we are summing over to k from 0 through t if we negate the expression:

$$s_{t,i} = \frac{\partial^2}{\partial p_i^2} \mu_t(p) = \frac{1}{m} \frac{d^2}{d p_i^2} \sum_{k=0}^t \binom{m}{k} p_i^k (1-p_i)^{m-k} (t-k) .$$

We first prove the case when $t \in \mathbb{Z}_+$. In this case, we view $s_{t,i}$ as a sequence with respect to t (where i is fixed), which we denote s_t . We now derive a generating function for this sequence.² Observe that derivatives that are not with respect to the formal variable commute with taking generating functions. Then, the generating function for the sequence $\{s_t\}$ is

$$\frac{d^2}{d p_i^2} \frac{1}{m} \left(x \frac{d}{d x} \left(\frac{(p_i x + 1 - p_i)^m}{1-x} \right) - \frac{x \frac{d}{d x} (p_i x + 1 - p_i)^m}{1-x} \right) = (m-1)(p_i x + 1 - p_i)^{m-2} x .$$

²To avoid potential convergence issues, we view generating functions as formal polynomials from the ring of infinite formal polynomials. Under this formalism, there is no need to deal with convergence at all.

Note that the coefficient on x^0 is 0, so $s_{0,i} = 0$ as claimed. For $t \in \mathbb{Z}_{>0}$, the right hand side is the generating function of

$$(m-1) \binom{m-2}{t-1} p^{t-1} (1-p)^{m-t-1}.$$

Thus, this expression gives the i -th entry Hessian in the $t \in \mathbb{Z}_{\geq 0}$, as claimed.

Now consider the case when t is not an integer. In this case, we have:

$$\begin{aligned} s_{t,i} &\triangleq \frac{d^2}{dp_i^2} \frac{1}{m} \sum_{k=\lceil t \rceil}^m \binom{m}{k} p_i^k (1-p_i)^{m-k} (k-t) \\ &= \frac{d^2}{dp_i^2} \frac{1}{m} \sum_{k=\lceil t \rceil}^m \binom{m}{k} p_i^k (1-p_i)^{m-k} (k - \lceil t \rceil + \Delta t) \\ &= s_{\lceil t \rceil, i} + \Delta t \frac{d^2}{dp_i^2} \frac{1}{m} \sum_{k=\lceil t \rceil}^m \binom{m}{k} p_i^k (1-p_i)^{m-k} \\ &= s_{\lceil t \rceil, i} - \Delta t \frac{d^2}{dp_i^2} \frac{1}{m} \sum_{k=0}^{\lceil t \rceil - 1} \binom{m}{k} p_i^k (1-p_i)^{m-k}. \end{aligned}$$

The last equality is because if we change bounds on the sum so they are from 0 through m , we get 1 which has partial derivative 0. Thus, we can flip which terms we are summing over if we negate the expression.

Note that this expression we are subtracting above can be alternatively written as:

$$\Delta t \frac{d^2}{dp_i^2} \frac{1}{m} \sum_{k=0}^{\lceil t \rceil - 1} \binom{m}{k} p_i^k (1-p_i)^{m-k} = \Delta t \cdot (s_{\lceil t \rceil} - s_{\lfloor t \rfloor}).$$

Thus, we have

$$s_{t,i} = s_{\lceil t \rceil, i} - \Delta t \cdot (s_{\lceil t \rceil, i} - s_{\lfloor t \rfloor, i}) = \Delta t \cdot s_{\lfloor t \rfloor, i} + (1 - \Delta t) \cdot s_{\lceil t \rceil, i},$$

as desired. This completes the proof of Lemma 4. \square

It will be convenient to simplify the exact expressions of Lemma 4 into something more manageable. This is done in the following lemma:

Lemma 5. Fix any constant $c > 0$. The Hessian of $\mu(p)$, viewed as a function of p , is a diagonal matrix whose i -th diagonal entry is given by

$$h_{ii} = s_{t:=m/n, i} \geq \Theta(1) \cdot \begin{cases} \frac{m^2}{n} & \text{for } m \leq n \\ \sqrt{mn} & \text{for } n < m \leq c \cdot \frac{n}{\varepsilon^2} \end{cases},$$

assuming $p_i = \frac{1 \pm \varepsilon}{n}$, $m \geq 6$, $n \geq 2$, and $\varepsilon \leq 1/2$.

Similarly, these bounds are tight up to constant factors, as further improvements would violate our sample complexity lower bounds.

Proof. By Lemma 4, we have an exact expression $s_{t,i}$ for the i th entry of the Hessian of $\mu_t(p)$.

First, consider the case where $m \leq n$. Then we have

$$s_{t,i} = (1 - \Delta t) \cdot s_{\lceil t \rceil, i}.$$

Substituting $t = m/n$, $\lceil t \rceil = 1$, and $\Delta t = \lceil t \rceil - t = 1 - m/n$ gives

$$s_{t,i} = \frac{m}{n} \cdot (m-1)(1-p_i)^{m-2} = \Theta(1) \cdot \frac{m^2}{n}.$$

Now consider the case where $n < m \leq \Theta(1) \cdot \frac{n}{\varepsilon^2}$. Note that the case where $n < m < 2n$ follows from (1) the fact that $s_{t,i}$ for fractional t linearly interpolates between the value of $s_{t',i}$ the nearest two integral values of t' and (2) the analyses of the cases where $m \leq n$ and $2n \leq m \leq \Theta(1) \frac{n}{\varepsilon^2}$. Thus, all we have left to do is prove the case where $2n \leq m \leq \Theta(1) \cdot \frac{n}{\varepsilon^2}$.

Since $s_{t,i}$ is a convex combination of $s_{\lceil t \rceil,i}$ and $s_{\lfloor t \rfloor,i}$, it suffices to bound from below these quantities for $t = m/n$. Both of these tasks can be accomplished simultaneously by bounding from below the quantity $s_{t=m/n+\gamma,i}$ for arbitrary $\gamma \in [-1, 1]$.

We do this as follows: Let $t = m/n + \gamma$. Note that Stirling's approximation is tight up to constant factors as long as the number we are taking the factorial of is not zero. Note that $m-2 \geq 1$, $t-1 \geq 1$, and $m-t-1 \geq m/2-2 \geq 1$. Thus, if we apply Stirling's approximation to the factorials in the definition of the binomial coefficient and substitute $t = m/n + \gamma$, we obtain the following approximation, which is tight up to constant factors:

$$\begin{aligned} & (m-1) \binom{m-2}{t-1} p_i^{t-1} (1-p_i)^{m-t-1} \\ &= \Theta(1) \cdot (m-1) \cdot \sqrt{\frac{(m-m/n-1-\gamma)(m/n-1+\gamma)}{(m-2)^3}} \cdot \frac{(m-2)^m p_i^{m/n-1+\gamma} (1-p_i)^{m-m/n-1-\gamma}}{(m-m/n-1-\gamma)^{m-m/n-\gamma} (m/n+\gamma)^{m/n-1+\gamma}} \\ &= \Theta(1) \cdot \sqrt{\frac{m}{n}} \cdot \frac{m^m p_i^{m/n-1+\gamma} (1-p_i)^{m-m/n-1-\gamma}}{(m-m/n-1-\gamma)^{m-m/n-\gamma} (m/n-1+\gamma)^{m/n+\gamma}} \\ &= \Theta(1) \cdot \sqrt{mn} \cdot \frac{m^{m/n+\gamma} p_i^{m/n+\gamma} (1-p_i)^{m-m/n-1-\gamma}}{(1-\frac{1}{n}-\frac{1}{m})^{m-m/n-\gamma} (m/n)^{m/n+\gamma}} \\ &= \Theta(1) \cdot \sqrt{mn} \cdot \frac{(np_i)^{m/n+\gamma} (1-p_i)^{m-m/n-1-\gamma}}{(1-\frac{1}{n})^{m-m/n-\gamma}} \\ &= \Theta(1) \cdot \sqrt{mn} \cdot \frac{(1 \pm \varepsilon)^{m/n+\gamma} (1-\frac{1 \pm \varepsilon}{n})^{m-m/n-\gamma}}{(1-\frac{1}{n})^{m-m/n-\gamma}} \\ &= \Theta(1) \cdot \sqrt{mn} \cdot \frac{(1 \pm \varepsilon)^{m/n} (1-\frac{1 \pm \varepsilon}{n})^{m-m/n}}{(1-\frac{1}{n})^{m-m/n}} \\ &= \Theta(1) \cdot \sqrt{mn} \cdot (1 \pm \varepsilon)^{m/n} \left(1 \mp \frac{\varepsilon}{n-1}\right)^{m-m/n} \\ &\geq \Theta(1) \cdot \sqrt{mn} \cdot (1 \pm \varepsilon)^{m/n} (1 \mp \varepsilon)^{\frac{m}{n-1} - \frac{m}{n(n-1)}} \\ &= \Theta(1) \cdot \sqrt{mn} \cdot (1 \pm \varepsilon)^{m/n} (1 \mp \varepsilon)^{m/n} \\ &= \Theta(1) \cdot \sqrt{mn} \cdot (1 - \varepsilon^2)^{m/n} \\ &= \Theta(1) \cdot \sqrt{mn} \cdot e^{-\Theta(1) \cdot \varepsilon^2 m/n} \\ &\geq \Theta(1) \cdot \sqrt{mn}. \end{aligned}$$

This completes the proof of Lemma 5. □

We are now ready to prove the desired expectation gap.

Proof of Lemma 3: We start by reducing the soundness case to a much simpler setting. To do this, we use the following fact, established in Section 4:

Fact 6. *For any distribution p on $[n]$, there exists a distribution p' supported on $[n]$ whose probability mass values are in the set $\{\frac{1+\varepsilon'}{n}, \frac{1}{n}, \frac{1-\varepsilon'}{n}\}$ for some $\varepsilon' \geq d_{\text{TV}}(p, U_n)/2$, with at most one element having mass $\frac{1}{n}$, and such that $\mu(p') \leq \mu(p)$.*

Fact 6 is proven in Section 4. By Fact 6, there is a distribution p' that satisfies the conditions of Lemma 5, has total variation distance $\Theta(\varepsilon)$ to the uniform distribution, and $\mu(p') \leq \mu(p)$. Therefore, it suffices to prove a lower bound on the expectation gap between the completeness and soundness cases for distributions p of this form.

Note that all probability distributions on the line from p to U_n are also of this form for different (no larger) values of ε . Thus, Lemma 5 gives a lower bound on the diagonal entries of the Hessian at all points on this line. Since the Hessian is diagonal, this also bounds from below the minimum eigenvalue of the Hessian on this line. Therefore, by this and Equation (2), we obtain the first two cases of this lemma, as well as the third case for $\frac{n}{\varepsilon^2} \leq m \leq 4 \cdot \frac{n}{\varepsilon^2}$.

The final case of this lemma for $4 \cdot \frac{n}{\varepsilon^2} \leq m$ follows immediately from the folklore fact that if one takes at least this many samples, the empirical distribution approximates the true distribution with expected ℓ_1 error at most $\varepsilon/2$. For completeness, we give a proof. We have

$$\begin{aligned} \mathbb{E}[||X/m - p||_1] &= \sum_i \mathbb{E}[|X_i/m - p_i|] \leq \sum_i \sqrt{\text{Var}[X_i/m - p_i]} \\ &\leq \sum_i \sqrt{mp_i/m^2} \leq \sum_i \sqrt{m(1/n)/m^2} \\ &= \sqrt{n/m} \leq \varepsilon/2, \end{aligned} \tag{4}$$

where Equation (4) follows from the fact that the sum is a symmetric concave function of p , so it is maximized by setting all the p_i 's to be equal. \square

2.3 Concentration of Test Statistic: Proof of Theorem 2

Let the m samples be $Y_1, \dots, Y_m \in [n]$, and let $X_i, i \in [n]$, be the number of $j \in [m]$ for which $Y_j = i$. Let S be our empirical total variation test statistic, $S = \frac{1}{2} \sum_{i=1}^n |\frac{X_i}{m} - \frac{1}{n}|$. We prove the theorem in two parts, one when $m \geq n$, and one when $m \leq n$.

We will require a ‘‘Bernstein’’ form of the standard bounded differences (McDiarmid) inequality:

Lemma 7 (Bernstein version of McDiarmid’s inequality [Yin04]). *Let Y_1, \dots, Y_m be independent random variables taking values in the set \mathcal{Y} . Let $f : \mathcal{Y}^m \rightarrow \mathbb{R}$ be a function of y_1, \dots, y_m so that for every $j \in [m]$ and $y_1, \dots, y_m, y'_j \in \mathcal{Y}$, we have that:*

$$|f(y_1, \dots, y_j, \dots, y_m) - f(y_1, \dots, y'_j, \dots, y_m)| \leq B.$$

Then, we have:

$$\Pr[f(Y_1, \dots, Y_m) - \mathbb{E}[f] \geq z] \leq \exp\left(\frac{-2z^2}{mB^2}\right). \tag{5}$$

If in addition, for each $j \in [m]$ and $y_1, \dots, y_{j-1}, y_{j+1}, \dots, y_m$ we have that

$$\text{Var}_{Y_j}[f(y_1, \dots, Y_j, \dots, y_m)] \leq \sigma_j^2,$$

then we have

$$\Pr[f(Y_1, \dots, Y_m) - \mathbb{E}[f] \geq z] \leq \exp\left(\frac{-z^2}{2 \sum_{j=1}^m \sigma_j^2 + 2Bt/3}\right). \tag{6}$$

2.3.1 Case I: $m \geq n$

Since the Y_j 's are independent and S is $\frac{1}{m}$ -Lipschitz in them, the first form of McDiarmid's inequality implies that

$$\Pr[S - \mathbb{E}[S] \geq z] \leq \exp(-2mz^2) ,$$

and similarly, by applying it to $-S$, we have $\Pr[S - \mathbb{E}[S] \leq -z] \leq \exp(-2mz^2)$.

Let R be the right-hand side of the Equation in Lemma 3, so $\mu(p) - \mu(U_n) \geq R$ in the soundness case. Since we threshold the tester at $t = \mu(U_n) + R/2$, we find in both the completeness and soundness cases that the success probability will be at least

$$1 - \exp(-mR^2/2) ,$$

and hence we just need to show

$$mR^2/2 \geq \log(1/\delta) . \tag{7}$$

Since we are in the regime that $m \geq n$, there are two possible cases in Lemma 3.

For $n \leq m \leq n/\varepsilon^2$, we need that

$$\frac{m}{2} \cdot \Theta(1) \cdot \varepsilon^4 m/n \geq \log(1/\delta)$$

or

$$m \geq \Theta(1) \cdot \frac{\sqrt{n \log(1/\delta)}}{\varepsilon^2} .$$

For $m \geq n/\varepsilon^2$, we need that

$$\frac{m}{2} \cdot \Theta(1) \cdot \varepsilon^2 \geq \log(1/\delta)$$

or

$$m \geq \Theta(1) \cdot \frac{\log(1/\delta)}{\varepsilon^2} .$$

The theorem's assumption on m implies that both conditions hold, which completes the proof of Theorem 2 in this case.

2.3.2 Case II: $m \leq n$

To establish Theorem 2 for $m \leq n$, we will require the Bernstein form of McDiarmid's inequality (Equation (6) in Lemma 7).

To apply this form of Lemma 7, it suffices to compute B and σ_j for our test statistic as a function of the Y_j 's. Note that for $m \leq n$, $|\frac{X_i}{m} - \frac{1}{n}|$ is equal to $\frac{X_i}{m} - \frac{1}{n}$ whenever $X_i \neq 0$. In particular, this implies

$$\begin{aligned} S &= \frac{1}{2} \sum_{i=1}^n \left| \frac{X_i}{m} - \frac{1}{n} \right| \\ &= \frac{1}{2} \sum_{i=1}^n \left\{ \left(\frac{X_i}{m} - \frac{1}{n} \right) + \frac{2}{n} \cdot \mathbf{1}_{X_i=0} \right\} \\ &= \frac{1}{n} \cdot |\{i : X_i = 0\}| . \end{aligned}$$

Hence, the value of the parameter B for our test statistic is $1/n$, since each Y_j will affect the number of nonzero X_i 's by at most 1. In particular, the function value as Y_j varies and the other $Y_{j'}$'s are kept fixed can

be written as the sum of a deterministic quantity plus $(1/n) \cdot b$, where b is a Bernoulli random variable that is 1 if sample Y_j collides with another sample $Y_{j'}$ and 0 otherwise. Thus, the variance of S as Y_j varies and the other $Y_{j'}$'s are kept fixed is given by $\text{Var}[(1/n) \cdot b]$. This variance is $(1/n^2) \cdot r(1-r)$, where $r \leq m/n$ is the probability that Y_j collides with another $Y_{j'}$. Thus, the variance of S as Y_j varies and the other $Y_{j'}$'s are kept fixed is at most

$$\frac{1}{n^2} \cdot r(1-r) \leq r/n^2 \leq m/n^3 =: \sigma_j^2.$$

Applying Equation (6) of Lemma 7 we find

$$\Pr[|S - \mathbb{E}[S]| \geq z] \leq 2 \exp\left(\frac{-z^2}{2 \cdot m^2/n^3 + (2/3) \cdot z/n}\right).$$

By Lemma 3 in the soundness case we have expectation gap $\mu(p) - \mu(U_n) \geq R := C\varepsilon^2 m^2/n^2$ for some constant $C < 1$. Substituting $z = R/2$ in the above concentration inequality yields that our tester will be correct with probability $1 - \delta$ as long as

$$m \geq \Theta(1) \cdot \frac{1}{\varepsilon^2} \sqrt{n \log(2/\delta)},$$

for an appropriately chosen constant, which is true by assumption. This completes the proof of Theorem 2.

3 Matching Information-Theoretic Lower Bound

In this section, we prove our matching sample complexity lower bound. Namely, we prove:

Theorem 8. *Any algorithm that distinguishes with probability at least $1 - \delta$ the uniform distribution on $[n]$ from any distribution that is ε -far from uniform, in total variation distance, requires at least*

$$\Omega\left(\left(\sqrt{n \log(1/\delta)} + \log(1/\delta)\right) / \varepsilon^2\right)$$

samples.

Theorem 8 will immediately follow from separate sample complexity lower bounds of $\Omega(\log(1/\delta)/\varepsilon^2)$ and $\Omega(\sqrt{n \log(1/\delta)}/\varepsilon^2)$ that we will prove. We start with a simple sample complexity lower bound of $\Omega(\log(1/\delta)/\varepsilon^2)$:

Lemma 9. *For all n, ε , and δ , any (ε, δ) uniformity tester requires $\Omega(\log(1/\delta)/\varepsilon^2)$ samples.*

Proof. If n is odd, set the last probability to $1/n$, subtract 1 from n , and invoke the following lower bound instance on the remaining elements. If n is even, do the following. Consider the distribution p which has probability $p_i = \frac{1+\varepsilon}{n}$ for each element $1 \leq i \leq \frac{n}{2}$ and $p_i = \frac{1-\varepsilon}{n}$ for each element $\frac{n}{2} \leq i \leq n$. Clearly, $d_{\text{TV}}(p, U_n) = \varepsilon$. Note that the probability that a sample comes from the first half of the domain is $\frac{1+\varepsilon}{2}$ and the probability that it comes from the second half of the domain is $\frac{1-\varepsilon}{2}$. Therefore, distinguishing p from U_n is equivalent to distinguishing between a fair coin and an ε -biased coin. It is well-known (see, e.g., Chapter 2 of [BY02]) that this task requires $m = \Omega(\log(1/\delta)/\varepsilon^2)$ samples. \square

The rest of this section is devoted to the proof of the following lemma, which gives our desired lower bound:

Lemma 10. *For all n, ε , and δ , any (ε, δ) uniformity tester requires at least $\Omega(\sqrt{n \log(1/\delta)}/\varepsilon^2)$ samples.*

To prove Lemma 10, we will construct two indistinguishable families of *pseudo-distributions*. A pseudo-distribution w is a non-negative measure, i.e., it is similar to a probability distribution except that the “probabilities” may sum to something other than 1. We will require that our pseudo-distributions always sum to a quantity within a constant factor of 1. A pair of pseudo-distribution families is said to be δ -*indistinguishable using m samples* if no tester exists that can, for every pair of pseudo-distributions w, w' —one from each of the two families—distinguish the product distributions $\otimes \text{Poi}(mw_i)$ versus $\otimes \text{Poi}(mw'_i)$ with failure probability at most δ .

This technique is fairly standard and has been used in [WY16, VV14, DK16] to establish lower bounds for distribution testing problems. The benefit of the method is that it is much easier to show that pseudo-distributions are indistinguishable, as opposed to working with ordinary distributions. Moreover, lower bounds proven using pseudo-distributions imply lower bounds on the original distribution testing problem.

We will require the following lemma, whose proof is implicit in the analyses of [WY16, VV14, DK16]:

Lemma 11. *Let \mathcal{P}' be a property of distributions. We extend \mathcal{P}' to the unique property \mathcal{P} of pseudo-distributions which agrees with \mathcal{P}' on true distributions and is preserved under rescaling. Suppose we have two families $\mathcal{F}_1, \mathcal{F}_2$ of pseudo-distributions with the following properties:*

1. *All pseudo-distributions in \mathcal{F}_1 have property \mathcal{P} and all those in \mathcal{F}_2 are ε -far in total variation distance from any pseudo-distribution that has the property.*
2. *\mathcal{F}_1 and \mathcal{F}_2 are δ -indistinguishable using m samples.*
3. *Every pseudo-distribution in each family has ℓ_1 -norm within the interval $[\frac{1}{c_1}, c_2]$, for some constants $c_1, c_2 > 1$.*

Then there exist two families $\tilde{\mathcal{F}}_1, \tilde{\mathcal{F}}_2$ of probability distributions with the following properties:

1. *All distributions in $\tilde{\mathcal{F}}_1$ have property \mathcal{P} and all those in $\tilde{\mathcal{F}}_2$ are $\frac{\varepsilon}{c_2}$ -far in total variation distance from any distribution that has the property.*
2. *Any tester that can distinguish $\tilde{\mathcal{F}}_1$ and $\tilde{\mathcal{F}}_2$ has worst-case error probability $\geq \delta - 2^{-cm}$, for some constant $c > 0$ or requires $\Theta(1) \cdot m$ samples.*

In our case, the property of distributions \mathcal{P}' is simply being the uniform distribution. The families of pseudo-distributions we will use for our lower bound are the family \mathcal{F}_1 that only contains the uniform distribution and the family \mathcal{F}_2 of all pseudo-distributions of the form $w_i = \frac{1 \pm \varepsilon}{n}$ such that $|1 - \sum_i w_i| \leq \varepsilon/2$. Note that this constraint on the sum of the w_i 's ensures the first and last conditions needed to invoke Lemma 11³.

Furthermore, by Lemma 9, the required number of samples m satisfies $m \geq \Omega(\log(1/\delta))$. Ignoring constant factors, we may assume that $m \geq c' \log(1/\delta)$ for any constant $c' > 0$. In particular, by selecting c' appropriately, we can guarantee that $2^{-cm} \leq \delta/3$, where c is the constant in the last statement of Lemma 11. Thus, the error probability guaranteed by Lemma 11 for distinguishing the true distribution families is at least $(2/3)\delta$.

Thus, all that remains is to show that \mathcal{F}_1 and \mathcal{F}_2 are δ -indistinguishable using m samples. In order to show these families are indistinguishable, we show that it is impossible to distinguish whether the product distribution $\otimes \text{Poi}(mw_i)$ has w uniform or w generated according to the following random process: we pick each w_i independently by setting $w_i = \frac{1+\varepsilon}{n}$ or $w_i = \frac{1-\varepsilon}{n}$ each with probability $1/2$.

A distribution generated by this process has a small probability of not being in \mathcal{F}_2 . Specifically, this happens iff it fails to satisfy the constraint on having a sum within $1 \pm \varepsilon/2$. However, by an application of

³If we did not have this constraint, the first condition would not be satisfied (because e.g., a w such that $w_i = (1 + \varepsilon)/n$ for all i is not ε -far from being proportional to the all-ones vector).

the Chernoff bound, it follows that this happens with probability at most $\leq 2^{-\Theta(1) \cdot n}$. Since the bound in Lemma 9 is larger (up to constant factors) than the lower bound we presently wish to prove in the case that $\delta/3 > 2^{-\Theta(1) \cdot n}$, we will assume $\delta \leq 2^{-\Theta(1) \cdot n}$; in which case, the following lemma implies that we can still invoke Lemma 11, where the probability of not being in \mathcal{F}_2 is absorbed into our overall indistinguishability probability, and we get a final indistinguishability probability of at least $(2/3)\delta - \delta/3 = \delta/3$.

The following lemma is implicit in [WY16, VV14, DK16]:

Lemma 12. *Let \mathcal{P} be a property of pseudo-distributions. Suppose we have two families $\mathcal{F}_1, \mathcal{F}_2$ of pseudo-distributions and two distributions $\mathcal{D}_1, \mathcal{D}_2$ on \mathcal{F}_1 and \mathcal{F}_2 respectively with the following properties:*

1. *With probability each at least $1 - \delta_1$, a distribution output by \mathcal{D}_1 is in \mathcal{F}_1 and a distribution output by \mathcal{D}_2 is in \mathcal{F}_2 .*
2. *If we generate w according to \mathcal{D}_1 or \mathcal{D}_2 , then any algorithm for determining which family w came from given access to $\otimes \text{Poi}(mw_i)$ has worst case error probability at least δ_2 .*

Then \mathcal{F}_1 and \mathcal{F}_2 are $(\delta_2 - \delta_1)$ -indistinguishable using m samples.

Thus, we now simply need to show that \mathcal{D}_1 and \mathcal{D}_2 are hard to distinguish. Let X_i, X'_i be the random variables equal to the number of times the element i is sampled in the completeness and soundness cases respectively. We will require a technical lemma that will be used to bound the Hellinger distance between any pair of corresponding coordinates in the completeness and soundness cases. By $\frac{1}{2} \text{Poi}(\lambda_1) + \frac{1}{2} \text{Poi}(\lambda_2)$, we denote a uniform mixture of the corresponding distributions. We have

Fact 13 (Lemma 7 of [VV14]). *For any $\lambda > 0, \varepsilon < 1$ we have*

$$H^2 \left(\text{Poi}(\lambda), \frac{1}{2} \text{Poi}((1 + \varepsilon)\lambda) + \frac{1}{2} \text{Poi}((1 - \varepsilon)\lambda) \right) \leq C\lambda^2\varepsilon^4.$$

for some constant C .

We are now ready to prove Lemma 10.

Proof of Lemma 10: As follows from the preceding discussion, it suffices to show that \mathcal{D}_1 and \mathcal{D}_2 are hard to distinguish. We will use Fact 13 to show that the Hellinger distance between the overall distributions is small, which implies their total variation distance is small, and hence that they cannot be distinguished with probability better than δ .

Recall that each of the n coordinates of the vectors output by these distributions is distributed according to $\text{Poi}(m/n)$ for \mathcal{D}_1 vs. a uniform mixture of $\text{Poi}((1 \pm \varepsilon)m/n)$ for \mathcal{D}_2 . A single coordinate then has

$$H^2(X_1, X'_1) \leq C(m^2/n^2)\varepsilon^4,$$

so the collection of all coordinates has

$$H^2(X, X') \leq 1 - (1 - C(m^2/n^2)\varepsilon^4)^n \leq 1 - e^{-C(m^2/n^2)\varepsilon^4}.$$

For this latter quantity to be at least $1 - \delta$, we need

$$m = \Omega \left((1/\varepsilon^2) \cdot \sqrt{n \log(1/\delta)} \right)$$

samples, as desired. The result then follows from Lemma 12 and Lemma 11. This completes the proof. \square

4 Stochastic Domination for Statistics of the Histogram

In this section, we consider the set of statistics which are symmetric convex functions of the histogram (i.e., the number of times each domain element is sampled) of an arbitrary random variable Y . We start with the following definition:

Definition 14. Let $p = (p_1, \dots, p_n), q = (q_1, \dots, q_n)$ be probability distributions and $p^\downarrow, q^\downarrow$ denote the vectors with the same values as p and q respectively, but sorted in non-increasing order. We say that p majorizes q (denoted by $p \succ q$) if

$$\forall k : \sum_{i=1}^k p_i^\downarrow \geq \sum_{i=1}^k q_i^\downarrow. \quad (8)$$

The following theorem from [Arn12] gives an equivalent definition:

Theorem 15. [Arn12] Let $p = (p_1, \dots, p_n), q = (q_1, \dots, q_n)$ be any pair of probability distributions. Then, $p \succ q$ if and only if there exists a doubly stochastic matrix A such that $q = Ap$.

Remark: It is shown in [Arn12] that multiplying the distribution p by a doubly stochastic matrix is equivalent to performing a series of so called ‘‘Robin hood operations’’ and permutations of elements. Robin hood operations are operations in which probability mass is transferred from heavier to lighter elements. For more details, the reader is referred to [Arn12, MOA79].

Note that Definition 14 defines a partial order over the set of probability distributions. We will see that the uniform distribution is a minimal element for this partial order, which directly follows as a special case of the following lemma.

Lemma 16. Let p be a probability distribution over $[n]$ and $S \subseteq [n]$. Let q be the distribution which is identical to p on $[n] \setminus S$, and for every $i \in S$ we have $q_i = \frac{p(S)}{|S|}$, where $|S|$ denotes the cardinality of S . Then, we have that $p \succ q$.

Proof. Let $A = (a_{ij})$ be the doubly stochastic matrix $A = (a_{ij})$ with entries:

$$a_{ij} = \begin{cases} 1 & \text{if } i = j \notin S \\ \frac{1}{|S|} & \text{if } i \in S \wedge j \in S \\ 0 & \text{otherwise} \end{cases}.$$

Observe that $q = Ap$. Therefore, Theorem 15 implies that $p \succ q$. \square

In the rest of this section, we use the following standard terminology: We say that a real random variable A stochastically dominates a real random variable B if for all $x \in \mathbb{R}$ it holds $\Pr[A > x] \geq \Pr[B > x]$. We now state the main result of this section (see Section 4.1 for the proof):

Lemma 17. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a symmetric convex function and p be a distribution over $[n]$. Suppose that we draw m samples from p , and let X_i denote the number of times we sample element i . Let $g(p)$ be the random variable $f(X_1, X_2, \dots, X_n)$. Then, for any distribution q over $[n]$ such that $p \succ q$, we have that $g(p)$ stochastically dominates $g(q)$.

As a simple consequence of the above, we obtain the following:

Fact 18. Let p be a distribution on $[n]$ and $S \subseteq [n]$. Let p' be the distribution which is identical to p in $[n] \setminus S$ and the probabilities in S are averaged (i.e., $p'_i = \frac{p(S)}{|S|}$). Then, we have that $\mu(p') \leq \mu(p)$, where $\mu(\cdot)$ denotes the expectation of our statistic as defined in Section 2. In particular, $\mu(U_n) \leq \mu(p)$ for all p .

Proof. Recall that our statistic applies a symmetric convex function f to the histogram of the sampled distribution. Since p' is averaging the probability masses on a subset $S \subseteq [n]$, Lemma 16 gives us that $p \succ p'$. Therefore, by Lemma 17 we conclude that $g(p)$ stochastically dominates $g(p')$, which implies that: $\mu(p) = \mathbb{E}[g(p)] \geq \mathbb{E}[g(p')] = \mu(p')$, as was to be shown. \square

The following lemma shows that given an arbitrary distribution p over $[n]$ that is ε -far from the uniform distribution U_n , if we average the heaviest $\lfloor n/2 \rfloor$ elements and then the lightest $\lfloor n/2 \rfloor$ elements, we will get a distribution that is $\varepsilon' > \varepsilon/2$ -far from uniform.

Lemma 19. *Let p be a probability distribution and p' be the distribution obtained from p after averaging the $\lfloor \frac{n}{2} \rfloor$ heaviest and the $\lfloor \frac{n}{2} \rfloor$ lightest elements separately. Then, the following holds:*

$$\frac{\|p - U_n\|_1}{2} \leq \|p' - U_n\|_1 \leq \|p - U_n\|_1 .$$

We note that by doing the averaging as suggested by the above lemma, we obtain a distribution p' that is supported on the following set of three values: $\{\frac{1+\varepsilon'}{n}, \frac{1}{n}, \frac{1-\varepsilon'}{n}\}$, for some $\frac{\varepsilon}{2} \leq \varepsilon' \leq \varepsilon$. Hence, we can reduce the computation of the expectation gap for an arbitrary distribution p , to computing the gap for a distribution of this form. Fact 6 is an immediate corollary of Lemmas 16, 17, and 19.

4.1 Proof of Lemma 17

To establish Lemma 17, we are going to use the following intermediate lemmas:

Lemma 20. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a symmetric convex function, and $a, b, c \in \mathbb{R}$ such that $0 < a < b$ and $c > 0$. Then,*

$$f(a, b + c, x_3, \dots, x_n) \geq f(a + c, b, x_3, \dots, x_n) .$$

Proof. Consider the set of convex functions $f'_{x_3, \dots, x_n} : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined as:

$$f'_{x_3, \dots, x_n}(x_1, x_2) = f(x_1, x_2, x_3, \dots, x_n) .$$

We will show that for every possible choice of x_3, \dots, x_n it holds that:

$$f'_{x_3, \dots, x_n}(a, b + c) \geq f'_{x_3, \dots, x_n}(a + c, b) .$$

Since f is symmetric, so is f' . Therefore, we have that $f'_{x_3, \dots, x_n}(a, b + c) = f'_{x_3, \dots, x_n}(b + c, a)$. The 3 points: $P_1 = (a, b + c)$, $P_2 = (a + c, b)$, $P_3 = (b + c, a)$ are collinear since their coordinates satisfy the equation $x_1 + x_2 = a + b + c$.

We have that P_2 is between P_1 and P_3 since:

$$\langle P_1 \vec{P}_2, P_2 \vec{P}_3 \rangle = \langle (c, -c), (b - a, a - b) \rangle > 0 .$$

By applying Jensen's inequality, we get that

$$f'_{x_3, \dots, x_n}(a + c, b) \leq \frac{f'_{x_3, \dots, x_n}(a, b + c) + f'_{x_3, \dots, x_n}(b + c, a)}{2} = f'_{x_3, \dots, x_n}(a, b + c) .$$

as desired. \square

The stochastic domination between the two statistics is established in the following lemma:

Lemma 21. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a symmetric convex function, p be a distribution over $[n]$, and $a, b \in [n]$ be such that $p_a < p_b$. Also, let q be the distribution which is identical to p on $[n] \setminus \{a, b\}$, and for which:

$$\begin{pmatrix} q_a \\ q_b \end{pmatrix} = \begin{pmatrix} w & 1-w \\ 1-w & w \end{pmatrix} \begin{pmatrix} p_a \\ p_b \end{pmatrix}, \quad (9)$$

where $w \in [\frac{1}{2}, 1]$. Suppose we take m samples from p and let X_i denote the number of times we sample element i . Let $g(p)$ be the random variable $f(X_1, X_2, \dots, X_n)$. Then, $g(p)$ stochastically dominates $g(q)$.

Proof. To prove stochastic domination between $g(p)$ and $g(q)$, we are going to define a coupling under which it is always true that $g(p)$ takes a larger value than $g(q)$.

Initially, we define an auxiliary coupling between p and q as follows: To get a sample from q , we first sample from p and use it as our sample, unless the output is element “ b ”, in which case we output “ a ” with probability $\frac{(1-w)(p_b-p_a)}{p_b}$ and “ b ” otherwise⁴.

Suppose now that we draw m samples from p , which we also convert to samples from q using the above rule.

In relation to this coupling we define the following random variables:

- X_{low} : The number of times element “ a ” is sampled.
- X_{high} : The number of times element “ b ” is sampled and *is not* swapped for element “ a ” in q .
- X_{mid} : The number of times element “ b ” is sampled and *is* swapped for element “ a ” in q .

From the above, we have that:

$$X_a = X_{low}, \quad X_b = X_{high} + X_{mid}, \quad X'_a = X_{low} + X_{mid}, \quad X'_b = X_{high},$$

where X'_i is the number of times element i is sampled in q .

We want to show that $g(p)$ stochastically dominates $g(q)$. That is, we want to show that⁵:

$$\forall t : \Pr[f(X_{low}, X_{high} + X_{mid}, X_3, \dots, X_n) \geq t] \geq \Pr[f(X_{low} + X_{mid}, X_{high}, X_3, \dots, X_n) \geq t].$$

We now condition on the events $E_{\{y,z\}} : \{X_{low}, X_{high}\} = \{y, z\}$ and $E_{c,x_3,\dots,x_n} : X_{mid} = c \wedge X_3 = x_3 \wedge \dots \wedge X_n = x_n$, where $y \leq z$ without loss of generality. Let $B = E_{\{y,z\}} \wedge E_{c,x_3,\dots,x_n}$.

We have that:

$$\begin{aligned} & \Pr[f(X_{low}, X_{high} + X_{mid}, X_3, \dots, X_n) \geq t] \\ &= \sum_{y \leq z} \Pr[f(X_{low}, X_{high} + X_{mid}, X_3, \dots, X_n) \geq t \mid B] \Pr[B]. \end{aligned}$$

So, it suffices to show that for every $y, z, c, x_3, \dots, x_n, t$:

$$\begin{aligned} & \Pr[f(X_{low}, X_{high} + X_{mid}, X_3, \dots, X_n) \geq t \mid B] \\ & \geq \Pr[f(X_{low} + X_{mid}, X_{high}, X_3, \dots, X_n) \geq t \mid B]. \end{aligned} \quad (10)$$

At this point, we have conditioned on everything except which of X_{low} and X_{high} is y and which is z . That is, after conditioning on the event $B = E_{\{y,z\}} \wedge E_{c,x_3,\dots,x_n}$, we have that:

$$\{f(X_{low} + X_{mid}, X_{high}, X_3, \dots, X_n), f(X_{low}, X_{high} + X_{mid}, X_3, \dots, X_n)\} = \{u, w\},$$

⁴Note that this coupling does not fix the value of $g(q)$ given a fixed value for $g(p)$, and is defined for convenience. We still have to show stochastic domination for the coupled random variables using a second coupling.

⁵To simplify notation, we pick $a = 1$ and $b = 2$ without loss of generality.

where $u = f(y + c, z, x_3, \dots, x_n)$, $w = f(y, z + c, x_3, \dots, x_n)$. Since by assumption $y \leq z$, we have by Lemma 20 that $u \leq w$. Then (10) holds trivially as an equality for $t \leq u$ and for $t > w$. For the remaining values of t , it is equivalent to:

$$\Pr[f(X_{low}, X_{high} + X_{mid}, X_3, \dots, X_n) = w \mid B] \geq \Pr[f(X_{low} + X_{mid}, X_{high}, X_3, \dots, X_n) = w \mid B],$$

and hence to

$$\Pr[X_{low} = y, X_{high} = z \mid B] \geq \Pr[X_{low} = z, X_{high} = y \mid B].$$

Now, this is also equivalent to a version with less restricted conditioning,

$$\Pr[X_{low} = y, X_{high} = z \mid E_{c, x_3, \dots, x_n}] \geq \Pr[X_{low} = z, X_{high} = y \mid E_{c, x_3, \dots, x_n}],$$

because neither event occurs in the added regime where $E_{\{y, z\}}$ is false. But if we rethink how our samples were drawn, we find that this is equivalent to showing that

$$p_a^y q_b^z \geq p_a^z q_b^y.$$

This holds since $q_b^{z-y} > p_a^{z-y}$, concluding the proof. \square

Proof of Lemma 17: Since $p \succ q$, we have by Theorem 15 (and the remark that follows it) that q can be constructed from p by repeated applications of (9). Therefore, Lemma 21 and the fact that stochastic domination is transitive imply that $g(p)$ stochastically dominates $g(q)$. \square

4.2 Proof of Lemma 19

Recall that p_i^\downarrow denotes the vector p with entries rearranged in non-increasing order. Suppose that at least $n/2$ elements have at least $1/n$ probability mass⁶. Therefore, if p is not the uniform distribution, we have

$$\sum_{k=1}^{\lfloor n/2 \rfloor} p_k^\downarrow = \frac{1 + \varepsilon'}{2} > \frac{1}{2},$$

for some $\varepsilon' > 0$.

Thus, we have that

$$p_k^{\prime\downarrow} = \begin{cases} \frac{1+\varepsilon'}{n} & \text{for } k \leq \frac{n}{2} \\ \frac{1-\varepsilon'}{n} & \text{for } k > \frac{n}{2} \end{cases}$$

when n is even, and

$$p_k^{\prime\downarrow} = \begin{cases} \frac{1+\varepsilon'}{n} & \text{for } k \leq \frac{n-1}{2} \\ \frac{1}{n} & \text{for } k = \frac{n+1}{2} \\ \frac{1-\varepsilon'}{n} & \text{for } k > \frac{n+1}{2} \end{cases}$$

when n is odd.

Moreover, since we are just averaging, we have that $\sum_{k=1}^{\lfloor n/2 \rfloor} p_k^\downarrow = \sum_{k=1}^{\lfloor n/2 \rfloor} p_k^{\prime\downarrow} = \frac{1+\varepsilon'}{2}$. Since we have assumed that the majority of elements has mass at least $\frac{1}{n}$, we know that the total variation distance is given by:

$$d_{\text{TV}}(p, U_n) = \sum_{i: p_i > 1/n} (p_i - 1/n) \leq 2 \sum_{k=1}^{\lfloor n/2 \rfloor} (p_k^\downarrow - 1/n) = 2 \sum_{k=1}^{\lfloor n/2 \rfloor} (p_k^{\prime\downarrow} - 1/n) \leq 2 \sum_{i: p_i^\downarrow > 1/n} (p_i^\downarrow - 1/n) = 2d_{\text{TV}}(p', U_n).$$

Thus, $d_{\text{TV}}(p', U_n) \geq (1/2)d_{\text{TV}}(p, U_n)$ or $\|p - U_n\|_1/2 \leq \|p' - U_n\|_1$, as desired.

⁶This is without loss of generality, since we can use essentially the same argument in the other case.

5 Conclusions and Future Work

In this paper, we gave the first uniformity tester that is sample-optimal, up to constant factors, as a function of the confidence parameter. Our tester is remarkably simple and our novel analysis may be useful in other related settings. By using a known reduction of identity to uniformity, we also obtain the first sample-optimal identity tester in the same setting.

Our result is a step towards understanding the behavior of distribution testing problems in the high-confidence setting. We view this direction as one of fundamental theoretical and important practical interest. A number of interesting open problems remain. Perhaps the most appealing one is to design a *general technique* (see, e.g., [DK16]) that yields sample-optimal testers in the high confidence regime for a wide range of properties. From the practical standpoint, it would be interesting to perform a detailed experimental evaluation of the various algorithms (see, e.g., [HM13, BW17]).

References

- [ADJ⁺11] J. Acharya, H. Das, A. Jafarpour, A. Orlitsky, and S. Pan. Competitive closeness testing. *Journal of Machine Learning Research - Proceedings Track*, 19:47–68, 2011.
- [ADK15] J. Acharya, C. Daskalakis, and G. Kamath. Optimal testing for properties of distributions. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3591–3599, 2015.
- [Arn12] Barry Arnold. *Majorization and the Lorenz order: A brief introduction*, volume 43. Springer Science & Business Media, 2012.
- [Bat01] T. Batu. *Testing Properties of Distributions*. PhD thesis, Cornell University, 2001.
- [BDKR02] T. Batu, S. Dasgupta, R. Kumar, and R. Rubinfeld. The complexity of approximating entropy. In *ACM Symposium on Theory of Computing*, pages 678–687, 2002.
- [BFF⁺01] T. Batu, E. Fischer, L. Fortnow, R. Kumar, R. Rubinfeld, and P. White. Testing random variables for independence and identity. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, pages 442–451, 2001.
- [BFR⁺00] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing that distributions are close. In *IEEE Symposium on Foundations of Computer Science*, pages 259–269, 2000.
- [BFR⁺13] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing closeness of discrete distributions. *J. ACM*, 60(1):4, 2013.
- [BKR04] T. Batu, R. Kumar, and R. Rubinfeld. Sublinear algorithms for testing monotone and unimodal distributions. In *ACM Symposium on Theory of Computing*, pages 381–390, 2004.
- [BW17] S. Balakrishnan and L. A. Wasserman. Hypothesis testing for densities and high-dimensional multinomials: Sharp local minimax rates. *CoRR*, abs/1706.10003, 2017.
- [BY02] Z. Bar-Yossef. *The Complexity of Massive Data Set Computations*. PhD thesis, Berkeley, CA, USA, 2002.
- [Can15] C. L. Canonne. A survey on distribution testing: Your data is big, but is it blue? *Electronic Colloquium on Computational Complexity (ECCC)*, 22:63, 2015.

- [CDGR16] C. L. Canonne, I. Diakonikolas, T. Gouleakis, and R. Rubinfeld. Testing shape restrictions of discrete distributions. In *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016*, pages 25:1–25:14, 2016.
- [CDKS17] C. L. Canonne, I. Diakonikolas, D. M. Kane, and A. Stewart. Testing bayesian networks. In *Proceedings of the 30th Conference on Learning Theory, COLT 2017*, pages 370–448, 2017.
- [CDS17] C. L. Canonne, I. Diakonikolas, and A. Stewart. Fourier-based testing for families of distributions. *CoRR*, abs/1706.05738, 2017.
- [CDVV14] S. Chan, I. Diakonikolas, P. Valiant, and G. Valiant. Optimal algorithms for testing closeness of discrete distributions. In *SODA*, pages 1193–1203, 2014.
- [DDS⁺13] C. Daskalakis, I. Diakonikolas, R. Servedio, G. Valiant, and P. Valiant. Testing k -modal distributions: Optimal algorithms via reductions. In *SODA*, pages 1833–1852, 2013.
- [DGPP16] I. Diakonikolas, T. Gouleakis, J. Peebles, and E. Price. Collision-based testers are optimal for uniformity and closeness. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:178, 2016.
- [DK16] I. Diakonikolas and D. M. Kane. A new approach for testing properties of discrete distributions. In *FOCS*, pages 685–694, 2016. Full version available at abs/1601.05557.
- [DKN15a] I. Diakonikolas, D. M. Kane, and V. Nikishkin. Optimal algorithms and lower bounds for testing closeness of structured distributions. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015*, pages 1183–1202, 2015.
- [DKN15b] I. Diakonikolas, D. M. Kane, and V. Nikishkin. Testing identity of structured distributions. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 1841–1854, 2015.
- [DKN17] I. Diakonikolas, D. M. Kane, and V. Nikishkin. Near-optimal closeness testing of discrete histogram distributions. *CoRR*, abs/1703.01913, 2017. To appear in ICALP 2017.
- [DL01] L. Devroye and G. Lugosi. *Combinatorial methods in density estimation*. Springer Series in Statistics, Springer, 2001.
- [Gol16] O. Goldreich. The uniform distribution is complete with respect to testing identity to a fixed distribution. *ECCC*, 23, 2016.
- [Gos78] R. William Gosper. Decision procedure for indefinite hypergeometric summation. *Proceedings of the National Academy of Sciences*, 75(1):40–42, 1978.
- [GR00] O. Goldreich and D. Ron. On testing expansion in bounded-degree graphs. Technical Report TR00-020, Electronic Colloquium on Computational Complexity, 2000.
- [HM13] D. Huang and S. Meyn. Generalized error exponents for small sample universal hypothesis testing. *IEEE Trans. Inf. Theor.*, 59(12):8157–8181, December 2013.
- [ILR12] P. Indyk, R. Levi, and R. Rubinfeld. Approximating and Testing k -Histogram Distributions in Sub-linear Time. In *PODS*, pages 15–22, 2012.
- [LR05] E. L. Lehmann and J. P. Romano. *Testing statistical hypotheses*. Springer Texts in Statistics. Springer, 2005.

- [LRR11] R. Levi, D. Ron, and R. Rubinfeld. Testing properties of collections of distributions. In *ICS*, pages 179–194, 2011.
- [MOA79] A. W. Marshall, I. Olkin, and B. C. Arnold. *Inequalities: theory of majorization and its applications*, volume 143. Springer, 1979.
- [NP33] J. Neyman and E. S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694-706):289–337, 1933.
- [Pan08] L. Paninski. A coincidence-based test for uniformity given very sparsely-sampled discrete data. *IEEE Transactions on Information Theory*, 54:4750–4755, 2008.
- [PS95] P. Paule and M. Schorn. A mathematica version of zeilberger’s algorithm for proving binomial coefficient identities. *Journal of Symbolic Computation*, 20(5):673 – 698, 1995.
- [PWZ97] M. Petkovsek, H.S. Wilf, and D. Zeilberger. *A = B (Online Edition)*. Ak Peters Series. Taylor & Francis, 1997.
- [Rub12] R. Rubinfeld. Taming big probability distributions. *XRDS*, 19(1):24–28, 2012.
- [Rub14] R. Rubinfeld. Taming probability distributions over big domains. *Talk given at STOC’14 Workshop on Efficient Distribution Estimation*, 2014. Available at <http://www.iliasdiakonikolas.org/stoc14-workshop/rubinfeld.pdf>.
- [Val11] P. Valiant. Testing symmetric properties of distributions. *SIAM J. Comput.*, 40(6):1927–1968, 2011.
- [vdVW96] A. W. van der Vaart and J. A. Wellner. *Weak convergence and empirical processes*. Springer Series in Statistics. Springer-Verlag, New York, 1996. With applications to statistics.
- [VV14] G. Valiant and P. Valiant. An automatic inequality prover and instance optimal identity testing. In *FOCS*, 2014.
- [WY16] Y. Wu and P. Yang. Minimax rates of entropy estimation on large alphabets via best polynomial approximation. *IEEE Transactions on Information Theory*, 62(6):3702–3720, June 2016.
- [Yin04] Y. Ying. Mcdiarmid’s inequalities of bernstein and bennett forms. *City University of Hong Kong*, 2004.

Appendix

A Computation of the Expectation in Completeness Case

Our statistic can be written as: $S = \sum_{i=1}^n \max\{X_i - \frac{m}{n}, 0\}$. Therefore, by linearity of expectation, we get:

$$\mathbb{E}[S] = \sum_{i=1}^n \mathbb{E} \left[\max\{X_i - \frac{m}{n}, 0\} \right] = n \cdot \mathbb{E} \left[\max\{X_i - \frac{m}{n}, 0\} \right].$$

So, all we need to do is to compute: $\mathbb{E} \left[\max\{X_i - \frac{m}{n}, 0\} \right]$ for a single value of i in the completeness case.

Note that $X_i \sim \text{Bin}(m, \frac{1}{n})$ and that the above expectation can be written as:

$$\mathbb{E} \left[\max\{X_i - \frac{m}{n}, 0\} \right] = \sum_{k=\lceil \frac{m}{n} \rceil}^m \Pr[X_i = k] \left(k - \frac{m}{n}\right), \quad (11)$$

where $\Pr[X_i = k] = \frac{(1-\frac{1}{n})^{m-k}}{n^k}$. This is a sum of $O(m)$ terms each of which can be computed in constant time, giving an $O(m)$ runtime overall.

B Non-Existence of Indefinite Closed-Form for Components of Expectation

In this appendix, we formalize and prove our assertion from Section 2.2 that the function $\mu(p) - \mu(U_n)$ is a linear combination of sums each of which has no indefinite closed form.

Recall Equation (3) which says that the expectation is a linear combination of sums with summands of the form $\binom{m}{k} q^k (1-q)^{m-k} (k-t)$ for various values of q —where the values of q are themselves different variables that any closed form would need to depend on (in addition to the other variables). A sum is said to have an indefinite closed form if, when the upper and lower limits of the sum are replaced with new variables, the resulting sum has a closed form valid for all values of all variables.

By closed form, we mean a closed form as defined in [PWZ97, Definition 8.1.1] which, as far as we are aware, is the main formal sense in which the phrase is used in combinatorics. This definition of closed form says that a function can be written as a sum of a constant number of rational functions, where the numerator and denominator in each is a linear combination of a constant number of products of exponentials, factorials, and constant degree polynomials. An example of such a function is $\frac{1}{\binom{m}{k}} + 7 \cdot \frac{2^k k! + 2k^5}{3^k} + k$.

To prove that a sum with summands $\binom{m}{k} q^k (1-q)^{m-k} (k-t)$ has no indefinite closed form—where m, k, q, t , and the limits of the sum are the variables that the closed form would need to be a function of—one can run Gosper's algorithm on this summand—with k as the index of summation—and observe that it returns that there is no indefinite closed form solution in the sense we have described [PWZ97, Theorem 5.6.3], [Gos78, PS95].