

Testing Support Size More Efficiently Than Learning Histograms

Received Oct 16, 2025
Revised Mar 1, 2026
Accepted Mar 21, 2026
Published May 21, 2026

Key words and phrases
Distribution testing, property testing, support size, Chebyshev polynomials

Renato Ferreira Pinto Jr.^a ✉ 

Nathaniel Harms^b ✉ 

^a Columbia University, USA

^b University of British Columbia, Canada

ABSTRACT. Consider two problems about an unknown probability distribution p :

1. How many samples from p are required to test if p is supported on n elements or not? Specifically, given samples from p , determine whether it is supported on at most n elements, or it is “ ε -far” (in total variation distance) from being supported on n elements.
2. Given m samples from p , what is the largest lower bound on its support size that we can produce?

The best known upper bound for problem (1) uses a general algorithm for learning the histogram of the distribution p , which requires $\Theta(\frac{n}{\varepsilon^2 \log n})$ samples. We show that testing can be done more efficiently than learning the histogram, using only $O(\frac{n}{\varepsilon \log n} \log(1/\varepsilon))$ samples, nearly matching the best known lower bound of $\Omega(\frac{n}{\varepsilon \log n})$. This algorithm also provides a better solution to problem (2), producing larger lower bounds on support size than what follows from previous work. The proof relies on an analysis of Chebyshev polynomial approximations *outside* the range where they are designed to be good approximations, and the paper is intended as an accessible self-contained exposition of the Chebyshev polynomial method.

This work was done while Renato Ferreira Pinto Jr. was a student at the University of Waterloo, and Nathaniel Harms was a postdoc at EPFL. Renato Ferreira Pinto Jr. was supported by an NSERC Canada Graduate Scholarship. Nathaniel Harms was supported by the Swiss State Secretariat for Education, Research, and Innovation (SERI) under contract number MB22.00026, an NSERC Postdoctoral Fellowship, and a Simons Fellowship. Thanks to Amit Levi for a discussion on the bounded support size assumption, and to the anonymous reviewers for their comments and suggestions. Some of this work was done during the Sublinear Algorithms program at the Simons Institute. We also thank Michael Kapralov for funding a research visit for Renato to EPFL. A preliminary version of this work appeared at STOC 2025.

1. Introduction

Sadly, it is often necessary to make decisions based on probability distributions. For example, you may need to make a decision based on the support size of an unknown distribution:

EXAMPLE 1.1. There is a population of fish in the lake, which may be described as a probability distribution p over the set of possible fish species. You want to preserve a specimen of each species. You have no idea how many fish or fish species there are, and you only have 10000 sample jars. Is that enough, or should you buy more? You would like to decide between:

1. There are at most 10000 species of fish in the lake. (You have enough jars.)
2. Any collection of 10000 species will miss at least 0.1% of the population, i.e. 0.1% of the fish will belong to species *not* in the collection. (You need more jars.)

The naïve strategy is to sample random fish until you have filled all the jars, and if you continue to see new species, buy more jars. Of course, you prefer to predict in advance if you need more jars. How many fish do you need to sample before you can make this decision (and be correct with probability at least 99%)? ◆

The best known algorithms for making this type of decision work by learning the histogram of the distribution (the unordered multiset of nonzero probabilities densities $\{p_i | p_i \neq 0\}$). In this paper we show there is a more sample-efficient way to do it. Formally, the problem is to distinguish between distributions with support size at most n , and those which are ε -far from having support size at most n . We say p is ε -far from having support size at most n when, for every probability distribution q with support size $|\text{supp}(q)| \leq n$, the total variation (TV) distance between p and q is $\text{dist}_{\text{TV}}(p, q) > \varepsilon$. Then:

DEFINITION 1.2 (Testing Support Size). A *support-size tester* with sample complexity $m(n, \varepsilon, \sigma)$ is an algorithm A which takes as input the parameters $n \in \mathbb{N}$, $\varepsilon \in (0, 1)$, and $\sigma \in (0, 1)$. It draws a multiset \mathbf{S} of $m = m(n, \varepsilon, \sigma)$ independent samples from an unknown probability distribution p over \mathbb{N} , and its output must satisfy:

1. If $|\text{supp}(p)| \leq n$ then $\mathbb{P}_{\mathbf{S}} [A(\mathbf{S}) \text{ outputs Accept}] \geq \sigma$; and
2. If p is ε -far from having support size at most n , then $\mathbb{P}_{\mathbf{S}} [A(\mathbf{S}) \text{ outputs Reject}] \geq \sigma$.

Unless otherwise noted, we set $\sigma = 3/4$.

Testing support size is a basic statistical decision problem (see e.g. the textbook [12, §11.4] and recent work [18, 2, 3, 26]) that underlies many other commonly studied tasks. For example, it is a decision version of the problem of *estimating* support size – a problem whose history dates back to Fisher [11], Goodman [20], and Good & Turing [19]; see [7] for a survey, and more recent work including [30, 33, 34, 37, 25]. Testing support size is also important for understanding the *testing vs. learning* question in property testing.

1.1 Testing vs. Learning Distributions

Testing support size is a *distribution testing* problem, which is a type of *property testing*. A basic technique in property testing is the testing-by-learning approach of [15], where the tester learns an approximation of the input and makes its decision based on this approximation. One of the guiding questions in property testing is: When can this technique be beaten?

It requires $\Theta(n/\varepsilon^2)$ samples to *learn* a distribution p over domain $[n]$, up to TV distance ε (see e.g. the survey [8]). Remarkable and surprising recent work [33, 34, 21, 24] proves that the *histogram* of distributions p over $[n]$ can be learned with only $\Theta(\frac{n}{\varepsilon^2 \log n})$ samples — a vanishing fraction of the domain. This means one can test the support size – or perform several other tasks – using $O(\frac{n}{\varepsilon^2 \log n})$ samples to learn the histogram. As noted in the textbook [12] and recent work [18, 3], this testing-by-learning algorithm gives the best known upper bound for testing support size when the true support size is promised to satisfy $|\text{supp}(p)| = O(n)$. Without this promise there does not appear to be a simple way to obtain a similar bound from known results (see [18, pp.21] and the discussion in Appendix C, where we sketch arguments to obtain bounds of $O(\frac{n}{\varepsilon^3 \log n})$ from [35]). We show that the testing-by-learning algorithm can be beaten, while also removing any restriction on the true support size:

THEOREM 1.3. *For all $n \in \mathbb{N}$ and $\varepsilon \in (0, 1)$, the sample complexity of testing support size of an unknown distribution p (over any countable domain) is at most*

$$m(n, \varepsilon) = O\left(\frac{n}{\varepsilon \log n} \cdot \min\{\log(1/\varepsilon), \log n\}\right).$$

In terms of Example 1.1, if you have n jars, you can decide whether to buy more jars after taking $O(\frac{n}{\varepsilon \log n} \log(1/\varepsilon))$ samples, filling less than a sublinear $O(\frac{\log(1/\varepsilon)}{\varepsilon \log n})$ -fraction of the jars. Theorem 1.3 nearly matches the best known lower bound of $\Omega(\frac{n}{\varepsilon \log n})$ (which can be deduced from lower bounds of [33, 37]). An important note about this problem is that, in the setting of Theorem 1.3, the related problem of *estimating* the support size is impossible: the unknown support size of p is unbounded and the probability densities p_i can be arbitrarily small. If one makes assumptions to avoid these issues, then there are tight bounds of $\Theta(\frac{k}{\log k} \log^2(1/\varepsilon))$ for estimating the support size up to $\pm \varepsilon k$ when each nonzero density p_i is promised to satisfy $p_i > 1/k$ [37, 25] (see also [30, 33, 1, 34]). Recalling Example 1.1, k would be the total number of *fish* in the lake, whereas in our results n is the number of *species*.

1.2 Good Lower Bounds on Support Size

Quoting I. J. Good in [7], “I don’t believe it is usually possible to estimate the number of unseen species [i.e. support size]...but only an approximate lower bound to that number.” So, in lieu of an estimate, what is the best lower bound we can get?

QUESTION 1.4. Given m samples from a distribution p and parameter $\varepsilon \in (0, 1)$, what is the biggest number \widehat{S} that we can output, while still satisfying $\widehat{S} \leq (1 + \varepsilon)|\text{supp}(p)|$?

Our algorithm can produce lower bounds as large as $\Omega(\frac{\varepsilon}{\log(1/\varepsilon)} m \log m)$ out of only m samples. Let us formalize the quality of these lower bounds. A reasonable target for such a “Good” lower bound \widehat{S} is that it should exceed the ε -effective support size, the smallest number of elements covering $1 - \varepsilon$ probability mass:

DEFINITION 1.5 (Effective support size). For any probability distribution p over \mathbb{N} and any $\varepsilon \in (0, 1)$, we define the ε -effective support size $\text{eff}_\varepsilon(p)$ as the smallest number $k \in \mathbb{N}$ such that there exists distribution q with $|\text{supp}(q)| = k$ and $\text{dist}_{\text{TV}}(p, q) \leq \varepsilon$.

With no assumptions on the distribution, the effective support size is a more natural target for estimation, because it accounts for the fact that an arbitrarily large number of elements in the support may comprise an arbitrarily small probability mass. Effective support size was used in [9, 4], and algorithms for estimating it were analyzed in [14, 27, 13], but these algorithms assume the algorithm can learn the probability densities p_i by quering element i . Our algorithm provides a Good lower bound using only samples:

COROLLARY 1.6. For all $n \in \mathbb{N}$ and $\varepsilon \in (0, 1)$, there is an algorithm which draws at most

$$O\left(\frac{n}{\varepsilon \log n} \cdot \min\{\log(1/\varepsilon), \log n\}\right)$$

samples from an arbitrary distribution p , and outputs a number \widehat{S} which satisfies (with probability at least $3/4$)

$$\min\{\text{eff}_\varepsilon(p), n\} \leq \widehat{S} \leq (1 + \varepsilon)|\text{supp}(p)|.$$

We are not aware of any prior work focusing on this type of guarantee, although it seems quite natural: the value \widehat{S} output by the algorithm can be used for problems like Example 1.1 to predict the resources required for a task depending on support size. The algorithm will tell us either:

1. If $\widehat{S} < n$ (which would be the case if we were promised, say, $|\text{supp}(p)| < n/2$), then \widehat{S} units (jars) suffice to cover the distribution up to the removal of ε mass, and fewer than $\approx \widehat{S}$ units will fail to cover the whole distribution. Furthermore, we obtain this estimate using fewer samples than required to learn the histogram.
2. If $\widehat{S} \geq n$, then we have learned that we need more than $\approx \widehat{S}$ units to cover the distribution.

Let us translate this result into an answer for Question 1.4. For a fixed sample size m , the naïve approach, taking \widehat{S} to be the number of unique elements appearing in the sample, produces a lower bound $\widehat{S} \leq |\text{supp}(p)|$ that exceeds $\min\{\text{eff}_\varepsilon(p), \Theta(\varepsilon m)\}$ (see Proposition 2.2). The histogram learners (if they were proved to succeed on this task without an assumption on true support

size) produce a lower bound $\widehat{S} \leq (1 + \varepsilon)|\text{supp}(p)|$ that exceeds $\min\{\text{eff}_\varepsilon(p), \Theta(\varepsilon^2 m \log m)\}$. Our algorithm improves this lower bound up to $\min\{\text{eff}_\varepsilon(p), \Theta(\frac{\varepsilon}{\log(1/\varepsilon)} m \log m)\}$.

1.3 Testing vs. Learning Boolean Functions

Testing support size of distributions is also important for understanding testing vs. learning of Boolean functions. Arguably the most well understood model of learning is the distribution-free sample-based PAC model: The algorithm receives samples of the form $(x, f(x))$ where $x \sim p$ is drawn from an unknown distribution, and labeled by an unknown function $f : \mathcal{X} \rightarrow \{0, 1\}$, which is promised to belong to a hypothesis class \mathcal{H} . The algorithm should output a function $g \in \mathcal{H}$ such that $\mathbb{P}_{x \sim p}[f(x) \neq g(x)] \leq \varepsilon$. (Requiring $g \in \mathcal{H}$ is called *proper* PAC learning.)

Whereas the PAC learner assumes the input f belongs to the hypothesis class \mathcal{H} , a *tester* for \mathcal{H} is an algorithm which tests this assumption:

DEFINITION 1.7 (Distribution-free sample-based testing; see formal Definition 5.1). Given labeled samples $(x, f(x))$ with $x \sim p$ drawn from unknown distribution p and labeled by unknown function f , decide whether (1) $f \in \mathcal{H}$ or (2) f is ε -far from all functions $g \in \mathcal{H}$, meaning $\mathbb{P}_{x \sim p}[f(x) \neq g(x)] \geq \varepsilon$.

Testing \mathcal{H} can be done by running a proper PAC learner and checking if it worked [15]. One motivation for testing algorithms is to aid in model selection, i.e. choosing an appropriate hypothesis class \mathcal{H} for learning. To be useful for this, the tester should be more efficient than the learner. The sample complexity of proper PAC learning is between $\Omega(\text{VC}/\varepsilon)$ and $O(\frac{\text{VC}}{\varepsilon} \log(1/\varepsilon))$, where VC is the VC dimension of \mathcal{H} (and there are examples where either bound is tight) [6, 23, 22]. But there is no similar characterization of the sample complexity for *testing*, and it is not known when it is possible to do better than the testing-by-PAC-learning method.

In fact, very little is known about property testing in the distribution-free sample-based setting, and there are not many positive results (see e.g. [16, 12, 31, 5, 32]). For example, the basic class \mathcal{H}_n of functions $f : \mathbb{N} \rightarrow \{0, 1\}$ with $|f^{-1}(1)| \leq n$ has a tight bound of $\Theta(\text{VC}/\varepsilon)$ for proper learning with samples, but for testing with samples, tight bounds in terms of both VC and ε are not known! Using the lower bounds of [34, 37] for support size estimation, [5] show lower bounds of the form $\Omega(\frac{\text{VC}}{\varepsilon \log \text{VC}})$ for this class and several others (e.g. halfspaces, k -alternating functions &c.). This still leaves room for interesting upper bounds: any bound of $o(\text{VC}/\varepsilon)$ means that the tester is doing something significantly different from testing-by-PAC-learning.

Indeed, [16] show an upper bound of $O(\frac{\text{VC}}{\varepsilon^2 \log \text{VC}})$ samples for testing \mathcal{H}_n (with the extra promise that $f^{-1}(1) \subset [2n]$), by learning the *distribution* instead of the *function*: their algorithm learns the histogram of the underlying distribution p on the subdomain $f^{-1}(1)$, with some adjustments to handle the fact that the probability mass of p inside $f^{-1}(1)$ may be small. This

1 If the domain \mathbb{N} is replaced with $[Cn]$ for constant C , and queries are allowed, then $O(1/\varepsilon^2)$ queries suffice.

shows that the lower bound of [5] from support-size estimation is sometimes tight in terms of the VC dimension, and hints at a closer connection to distribution testing. We give a simple proof of a tighter and cleaner relationship – that testing support size of distributions and functions are essentially the same problem.

THEOREM 1.8 (Equivalence of testing support size for distributions and functions). *Let $m^{\text{DIST}}(n, \varepsilon, \sigma)$ be the sample complexity for testing support size with success probability σ . Let $m^{\text{FUN}}(n, \varepsilon, \sigma)$ be the sample complexity for distribution-free sample-based testing \mathcal{H}_n with success probability σ . Then $\forall n \in \mathbb{N}, \varepsilon, \sigma \in (0, 1)$, and $\xi \in (0, 1 - \sigma)$,*

$$m^{\text{DIST}}(n, \varepsilon, \sigma) \leq m^{\text{FUN}}(n, \varepsilon, \sigma) \leq m^{\text{DIST}}(n, \varepsilon, \sigma + \xi) + O\left(\frac{\log(1/\xi)}{\varepsilon}\right).$$

With Theorem 1.3, this improves the best upper bound for testing \mathcal{H}_n from $O(\frac{\text{VC}}{\varepsilon^2 \log \text{VC}})$ [16] to $O(\frac{\text{VC}}{\varepsilon \log \text{VC}} \log(1/\varepsilon))$. So, not only can \mathcal{H}_n be tested more efficiently than learned, but even more efficiently than the histogram of the underlying distribution can be learned.

We note also that there are reductions from testing the class \mathcal{H}_n to testing several other classes like k -alternating functions, halfspaces, and decision trees [5, 10], which emphasizes the role of testing support size of distributions to the testing vs. learning question for Boolean functions.

1.4 Techniques and Open Problems

The first part of our proof closely follows the optimal upper bound of [37] for the support size estimation problem. In short, if domain element i has probability density p_i and appears N_i times in the sample, the estimator of [37] outputs $\sum_i f(N_i)$, where $f(N_i)$ is constructed such that $\mathbb{E}[f(N_i)] = Q(p_i)$ where $Q(p_i)$ is a function satisfying $Q(p_i) \approx \mathbb{1}[p_i > 0]$. If this can be done, then

$$\mathbb{E}\left[\sum_i f(N_i)\right] = \sum_i Q(p_i) \approx |\text{supp}(p)|.$$

They make the necessary assumption that every nonzero density satisfies $p_i > 1/n$, allowing them to use Chebyshev polynomials to define a suitable Q where $Q(p_i) \approx \mathbb{1}[p_i > 1/n]$. Section 2 explains why Chebyshev polynomials arise naturally.

We follow the same strategy, although we no longer have any assumption on the nonzero densities. We construct a Chebyshev polynomial approximation $Q(p_i) \approx \mathbb{1}[p_i > \ell]$ for carefully chosen ℓ (depending on ε and n). The main challenge is to account for the densities $0 < p_i < \ell$, where the Chebyshev polynomial is not “designed” to be a good approximation, i.e. $Q(p_i) \neq \mathbb{1}[p_i > 0]$.

Main idea: We establish a tradeoff between the number of “heavy elements” with $p_i > \ell$ (on which the polynomial approximation satisfies $Q(p_i) \approx 1$), and the value of the estimator on the

“light elements” with $p_i \leq \ell$. This is proved in Section 3.6, with an easier version giving a worse bound in Section 3.5. The easier version uses concavity of Q on $(0, \ell)$ to show $Q(p_i) \geq p_i/\ell$ (Proposition 3.14). In this case, we use $\ell \approx \varepsilon/n$. This gives us a tradeoff of the form

$$\sum_i Q(p_i) \approx \#\{\text{heavy elements}\} + \frac{n}{\varepsilon} \sum_{\text{light } i} p_i. \quad (1)$$

Then, what remains to show is that, when p is ε -far from $|\text{supp}(p)| \leq n$, this quantity is large (and then bound the variance of the estimator, which depends on the degree of the polynomial approximation). To get a better bound on the sample size and prove our main theorem, we use a different setting of parameters in our polynomial approximation, and we establish a more complicated tradeoff than Equation (1) (see Proposition 3.25), essentially

$$\sum_i Q(p_i) \geq \left(n + \frac{\varepsilon}{p^*}\right) Q(p^*), \quad (2)$$

where p^* is the n^{th} largest density and ε is the distance to $|\text{supp}(p)| \leq n$, and we characterize the worst-case behaviour of this tradeoff.

Open questions. The sample complexity upper bound from Theorem 1.3 leaves a small gap compared to the best known lower bound of $\Omega(\frac{n}{\varepsilon \log n})$ from [33, 37]. Can we close this gap? We remark that our proof strategy does not allow us to improve any of our parameters (see Remark 3.31), and therefore we suspect that our upper bound is tight, unless there is a better algorithmic strategy (say, a non-linear estimator).

We may also ask for a *tolerant* tester [29], i. e. an algorithm that accepts not only when p has support size at most n , but also when p is ε' -close to such a distribution for some $\varepsilon' < \varepsilon$. Equivalently, we wish to distinguish between $\text{eff}_{\varepsilon'}(p) \leq n$ and $\text{eff}_{\varepsilon}(p) > n$. Our tester from Theorem 1.3, like any sample-based tester, enjoys the very small amount of tolerance $\varepsilon' = \Theta(1/m(n, \varepsilon))$. Under the promise that the true support size satisfies $|\text{supp}(p)| = O(n)$, algorithms for learning the histogram [33, 34, 21, 24] imply an upper bound of $O\left(\frac{n}{(\varepsilon - \varepsilon')^2 \log n}\right)$ for tolerant testing of support size, and in Appendix C we discuss how the techniques of [35] may plausibly allow one to drop the support size assumption at the cost of higher dependence on $\varepsilon - \varepsilon'$. Can we achieve $\frac{n}{\log n} \cdot \tilde{O}\left(\frac{1}{(\varepsilon - \varepsilon')^2}\right)$?

1.5 Organization

One may treat the first part of our paper (up to Remark 3.8) as an alternative self-contained exposition of [37], and we intend the paper to be as clear and accessible an explanation of the Chebyshev polynomial method as possible (we also refer the reader to the textbook [38] on polynomial methods in statistics).

Section 2 sets up the testing algorithm using generic parameters.

Section 3 places Constraints I to IV on the parameters in order to guarantee the success of the tester, and then in Section 3.5 we optimize the parameters to complete a weaker (but simpler) version of Theorem 1.3. The stronger version follows by replacing Constraint IV with the looser Constraint IVb in Section 3.6.

Section 4 proves that the algorithm outputs a lower bound on support size, Corollary 1.6.

Section 5 proves equivalence of testing support size of distributions and functions, Theorem 1.8.

2. Defining a Test Statistic

We begin by defining a “test statistic” \widehat{S} which the tester will use to make its decision – the tester will output Accept if \widehat{S} is small and output Reject if it is large. But our test statistic will require $\log(1/\varepsilon) < \log(n)$, so we first handle the case where the parameter ε is very small.

2.1 Small ε

If ε is small enough that $\log(1/\varepsilon) = \Omega(\log n)$, we will use the following simple tester. This result is folklore (appearing e.g. in [18, 2, 3]) but we include a proof for the sake of completeness. We start with the following simple observation:

OBSERVATION 2.1. *For any $n \in \mathbb{N}$, $\varepsilon \in (0, 1)$, and probability distribution p , p is ε -far from having support size at most n if and only if $\text{eff}_\varepsilon(p) > n$.*

Using this observation, we give

PROPOSITION 2.2. *There is an algorithm which, given inputs $n \in \mathbb{N}$, $\varepsilon \in (0, 1)$, and sample access to an arbitrary distribution p , draws at most $O(n/\varepsilon)$ samples from p and outputs a number \widehat{S} which (with probability at least $3/4$) satisfies*

$$\min\{\text{eff}_\varepsilon(p), n\} \leq \widehat{S} \leq |\text{supp}(p)|.$$

In particular, there is a support size tester using $O(n/\varepsilon)$ samples, obtained by running this algorithm with parameters $n + 1, \varepsilon$ and outputting Accept if and only if $\widehat{S} \leq n$.

PROOF. Write $k := \min\{n, \text{eff}_\varepsilon(p)\}$. Let $m := 10n/\varepsilon$. The algorithm draws m independent random samples from p and outputs \widehat{S} , defined as the number of distinct domain elements in the sample.

It is clear that $\widehat{S} \leq |\text{supp}(p)|$ with probability 1. To show that $\widehat{S} \geq \min\{\text{eff}_\varepsilon(p), n\}$, suppose we draw an infinite sequence of independent random samples $\mathbf{S} = \{\mathbf{s}_1, \mathbf{s}_2, \dots\}$ from p . For each $i \in \mathbb{N}$, define the random variable \mathbf{t}_i as the smallest index $\mathbf{t}_i \in [m]$ such that the multiset $\{\mathbf{s}_1, \dots, \mathbf{s}_{\mathbf{t}_i}\}$ is supported on i unique elements. Note that $\mathbf{t}_1 = 1$.

CLAIM 2.3. For all $i \in \{2, \dots, k\}$, $\mathbb{E}[\mathbf{t}_i - \mathbf{t}_{i-1}] \leq 1/\varepsilon$.

PROOF OF CLAIM. Write $\mathbf{S}_t := \{\mathbf{s}_1, \dots, \mathbf{s}_t\}$ as the prefix of \mathbf{S} up to element t . Fix any t_{i-1} . Conditional on the event $\mathbf{t}_{i-1} = t_{i-1}$, the multiset $\mathbf{S}_{t_{i-1}} = \{\mathbf{s}_1, \dots, \mathbf{s}_{t_{i-1}}\}$ is supported on $i - 1 < k$ elements. Then $p(\text{supp}(\mathbf{S}_{t_{i-1}})) < 1 - \varepsilon$ since p is ε -far from having support size at most $k - 1$ by Observation 2.1. Therefore for every $t > t_{i-1}$ we have $\mathbb{P}[\mathbf{s}_t \notin \mathbf{S}_{t_{i-1}} \mid \mathbf{t}_{i-1} = t_{i-1}] > \varepsilon$, so $\mathbb{E}[\mathbf{t}_i - \mathbf{t}_{i-1}] \leq 1/\varepsilon$. ■

From this claim we may deduce $\mathbb{E}[\mathbf{t}_k] \leq \mathbf{t}_1 + (k - 1)/\varepsilon \leq n/\varepsilon$, so by Markov's inequality

$$\mathbb{P}[\mathbf{t}_k > 10n/\varepsilon] \leq 1/10.$$

Therefore if $m \geq 10n/\varepsilon$ the tester will see at least k unique elements in a sample of size m , with probability at least $9/10$. ■

2.2 The test statistic

We now turn to the case where ε is large. Specifically, we fix a small universal constant $a \in (0, 1)$ (e.g. $a = 1/128$ suffices; we prioritize clarity of exposition over optimizing constants) and assume:

ASSUMPTION 2.4. For every $n \in \mathbb{N}$, we assume $n^{-a} < \varepsilon < 1/3$, so that $\log(1/\varepsilon) < a \log n$.

We will write $\mathbf{S} \subset \mathbb{N}$ for the (random) multiset of samples received by the algorithm.

DEFINITION 2.5 (Sample Histogram). For a fixed multiset $S \subset \mathbb{N}$, the *sample histogram* is the sequence N_i where N_i is the number of times $i \in \mathbb{N}$ appears in S . If \mathbf{S} is a random multiset then we write N_i .

The goal will be to find the best function f to plug in to the following definition of our test statistic; we will show how to choose f in Section 2.3 and the remainder of the paper, but for now we leave it undetermined.

DEFINITION 2.6 (Test statistic). For a given function $f : \{0\} \cup \mathbb{N} \rightarrow \mathbb{R}$ (which is required to satisfy $f(0) = -1$), we define a test statistic as follows. On random sample \mathbf{S} ,

$$\widehat{\mathbf{S}} := \sum_{i \in \mathbb{N}} (1 + f(N_i)).$$

Given an appropriate function f and resulting test statistic, we define our support-size tester:

DEFINITION 2.7 (Support-Size Tester). Given parameters $n \in \mathbb{N}$ and $\varepsilon \in (0, 1)$ our tester chooses m independent samples \mathbf{S} from the distribution p , and outputs *Accept* if and only if

$$\widehat{\mathbf{S}} < (1 + \varepsilon/2)n.$$

REMARK 2.8. Our test statistic is a *linear estimator* (see e.g. [33, 36, 34, 37]) because it can be written as

$$\widehat{\mathbf{S}} = \sum_{j=1}^m F_j \cdot (1 + f(j)),$$

where F_j is the number of elements $i \in \mathbb{N}$ which occur j times in the sample \mathbf{S} , i.e. $N_i = j$. The sequence F_1, F_2, \dots is known as the *fingerprint*.

A standard trick to ease the analysis is to consider instead the *Poissonized* version of the algorithm.

DEFINITION 2.9 (Poissonized Support-Size Tester). Given parameters $n \in \mathbb{N}$ and $\varepsilon \in (0, 1)$ our tester chooses $\mathbf{m} \sim \text{Poi}(m)$ and then takes \mathbf{m} independent samples \mathbf{S} from the distribution p , and outputs Accept if and only if

$$\widehat{\mathbf{S}} < (1 + \varepsilon/2)n.$$

The advantages of Poissonization are these (see e.g. the survey [8]):

FACT 2.10. Let \mathbf{S} be the sample of size $\mathbf{m} \sim \text{Poi}(m)$ drawn by the Poissonized tester. Then the random variables N_i are mutually independent and are distributed as

$$N_i \sim \text{Poi}(mp_i).$$

FACT 2.11. If there is a Poissonized support-size tester with sample complexity $m(n, \varepsilon)$ and success probability σ , then there is a standard support-size tester with sample complexity at most $O(m(n, \varepsilon))$ and success probability 0.99σ .

2.3 Choosing a function f

Now we see how to choose the function f to complete Definition 2.6. To motivate the choice, compute the expected value of the statistic:

PROPOSITION 2.12. For any given finitely supported $f: \{0\} \cup \mathbb{N} \rightarrow \mathbb{R}$ with $f(0) = -1$ and parameter m , the (Poissonized) test statistic $\widehat{\mathbf{S}}$ satisfies

$$\mathbb{E}[\widehat{\mathbf{S}}] = \sum_{i \in \mathbb{N}} (1 + e^{-mp_i} P(p_i)) \quad (3)$$

where $P(x)$ is the polynomial

$$P(x) = \sum_{j=0}^{\infty} f(j) \frac{m^j}{j!} x^j.$$

PROOF. Using Fact 2.10,

$$\begin{aligned}\mathbb{E}[\widehat{\mathbf{S}}] &= \sum_{i \in \mathbb{N}} \mathbb{E}[1 + f(\mathbf{N}_i)] = \sum_{i \in \mathbb{N}} \left(1 + \sum_{j \geq 0} \mathbb{P}[\mathbf{N}_i = j] f(j)\right) = \sum_i \left(1 + \sum_{j \geq 0} \frac{e^{-mp_i} (mp_i)^j}{j!} f(j)\right) \\ &= \sum_i \left(1 + e^{-mp_i} \sum_{j \geq 0} \frac{m^j f(j)}{j!} p_i^j\right) = \sum_i (1 + e^{-mp_i} P(p_i)).\end{aligned}$$

For convenience, we define

DEFINITION 2.13. For given polynomial $P(x)$, define

$$Q(x) := 1 + e^{-mx} P(x),$$

so that $Q(p_i)$ quantifies the contribution of element $i \in \mathbb{N}$ to the expected value of $\widehat{\mathbf{S}}$:

$$\mathbb{E}[\widehat{\mathbf{S}}] = \sum_{i \in \mathbb{N}} Q(p_i). \quad (4)$$

If we could set $Q(x) = \mathbb{1}[x > 0]$ then $\mathbb{E}[\widehat{\mathbf{S}}] = |\text{supp}(p)|$. But this means $P(0) = -1$ while $P(x) = 0$ for all $x > 0$, which is not a polynomial. Our goal is therefore to choose a polynomial P which approximates this function as closely as possible.

Specifically, if we choose any coefficients a_1, \dots, a_d , we may then define

$$f(j) := \begin{cases} a_j \frac{j!}{m^j} & \text{if } j \in [d] \\ -1 & \text{if } j = 0 \\ 0 & \text{if } j > d \end{cases}$$

to obtain the degree- d polynomial

$$P_d(x) := \sum_{j=0}^{\infty} f(j) \frac{m^j}{j!} x^j = \left(\sum_{j=1}^d a_j x^j \right) - 1,$$

which satisfies the necessary condition $P_d(0) = -1$. The other condition that we want $P_d(x)$ to satisfy is that $P_d(x) \approx 0$ for $x > 0$. This cannot be achieved by a low-degree polynomial, but we can instead ask for $P_d(x) \approx 0$ inside a chosen “safe interval” $[\ell, r]$. In other words, we want a polynomial which satisfies $P_d(0) = -1$ and

$$\max_{x \in [\ell, r]} |P_d(x)| \leq \delta,$$

for some small δ . The lowest-degree polynomials which satisfy these conditions are known as the (shifted and scaled) *Chebyshev polynomials*, which we define in Section 2.4.

REMARK 2.14. In the support-size *estimation* problem, there is a lower bound $p_i \geq 1/n$ on the nonzero densities, so [37] choose a “safe interval” $[\ell, r]$ with $\ell = 1/n$, so that there is no

density p_i to the left of the interval, $p_i \in (0, \ell)$. In our problem, we do not have this guarantee, and therefore we must handle values p_i that do not fall in the “safe” interval $[\ell, r]$.

2.4 Chebyshev polynomials and the definition of P_d

For a given degree d , the Chebyshev polynomial $T_d(x)$ is the polynomial which grows fastest on $x > 1$ while satisfying $|T_d(x)| \leq 1$ in the interval $x \in [-1, 1]$.

DEFINITION 2.15 (Chebyshev polynomials). We write $T_d(x)$ for the degree- d Chebyshev polynomial. These can be defined recursively as

$$\begin{aligned} T_0(x) &:= 1 \\ T_1(x) &:= x \\ T_d(x) &:= 2x \cdot T_{d-1}(x) - T_{d-2}(x). \end{aligned}$$

$T_d(x)$ may also be computed via the following closed-form formula: For all $d \in \mathbb{N}$ and $x \in \mathbb{R}$,

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

We want the polynomials $P_d(x)$ to be bounded in $[\ell, r]$ and to grow fast on $x \in [0, \ell)$ so that $P_d(0) = -1$. So we define a map ψ which translates $[\ell, r]$ to $[-1, 1]$,

$$\psi(x) := -\frac{2x - r - \ell}{r - \ell}$$

so $\psi(0) = \frac{r+\ell}{r-\ell}$. For convenience we write

$$\alpha := \frac{\ell}{r}$$

so that $\psi(0) = 1 + \frac{2\alpha}{1-\alpha}$. We now define

DEFINITION 2.16 (Shifted & Scaled Chebyshev Polynomials). For any d , define

$$P_d(x) := -\delta T_d(\psi(x)),$$

where $\delta := \frac{1}{T_d(1 + \frac{2\alpha}{1-\alpha})}$.

This way, as required,

$$P_d(0) = -\delta T_d(\psi(0)) = -1, \quad \text{and} \quad |P_d(x)| \leq \delta \text{ for } x \in [\ell, r].$$

We illustrate the Chebyshev polynomial $T_d(x)$ and the shifted and scaled result $Q(p_i)$ in Figure 1. The resulting function values $f(N_i)$ which define our test statistic are given in Proposition A.4.




Figure 1. The polynomial $T_{11}(x)$ and the resulting $Q(p_i) = 1 + e^{-mp_i} P_d(p_i)$ for $d = 11$ and certain choices of ℓ, r, n, m . The ‘safe’ interval $[\ell, r]$ is between the two vertical lines in Figure 1b. See that $Q(p_i)$ is an approximation of the ‘idealized’ function $Q(p_i) = \mathbb{1}[p_i > 0]$.

3. Choosing Parameters

Our goal is now:

GOAL 3.1. Choose degree d and “safe interval” $[\ell, r]$ such that the (Poissonized) support-size tester in Definition 2.9, instantiated with the (shifted and scaled) Chebyshev polynomial P_d , works with the smallest sample size parameter m , even though densities p_i may not belong to $[\ell, r]$.

To do this, we will set up a series of constraints that these parameters must satisfy in order for the tester to succeed, and then minimize m with respect to these parameters.

3.1 Constraint I: Choose d such that $\delta \leq \epsilon$

We will want that $\delta \leq \epsilon$ so that for every density p_i in the “safe interval” $[\ell, r]$ we have $Q(p_i) \approx 1$, specifically $|Q(p_i) - 1| \leq \epsilon$. We will write this constraint in terms of ℓ and r :

CONSTRAINT I. For constant $C_d := 4 \ln(2)$, we will require for any choice of ℓ, r, ϵ that $r \geq 3\ell$ and that the degree d satisfies

$$d \geq C_d \sqrt{\frac{r - \ell}{2\ell}} \log\left(\frac{20}{\epsilon}\right).$$

PROPOSITION 3.2. Assume Constraint I. Then $\delta \leq \epsilon/20$ and for all $x \in [\ell, r]$,

$$|Q(x) - 1| \leq \delta \leq \frac{\epsilon}{20}.$$

To prove this, we require a lower bound on the growth rate of the Chebyshev polynomials. This is well known but we provide a proof for the sake of completeness:

PROPOSITION 3.3. *There is a universal constant $c := \frac{1}{2\ln(2)}$ such that, for all $d \in \mathbb{N}$ and $\gamma \in [0, 1]$,*

$$T_d(1 + \gamma) \geq 2^{c \cdot d \sqrt{\gamma} - 1}.$$

PROOF. Using Definition 2.15, we have

$$\begin{aligned} T_d(1 + \gamma) &= \frac{1}{2} \left(\left((1 + \gamma) + \sqrt{(1 + \gamma)^2 - 1} \right)^d + \left((1 + \gamma) - \sqrt{(1 + \gamma)^2 - 1} \right)^d \right) \\ &\geq \frac{1}{2} \left(1 + \sqrt{2\gamma} \right)^d = \frac{1}{2} e^{d \ln(1 + \sqrt{2\gamma})}. \end{aligned}$$

Using the inequality $\ln(1 + x) \geq \frac{x}{1+x}$, which is valid for $x > -1$, together with the assumption that $\gamma \leq 1$, we obtain

$$T_d(1 + \gamma) \geq \frac{1}{2} e^{d \cdot \frac{\sqrt{2\gamma}}{1 + \sqrt{2\gamma}}} \geq \frac{1}{2} e^{d \cdot \frac{\sqrt{2\gamma}}{2\sqrt{2}}} = \frac{1}{2} e^{\frac{1}{2} d \sqrt{\gamma}}. \quad \blacksquare$$

PROOF OF PROPOSITION 3.2. Recall $\alpha := \ell/r$ and $\psi(0) = 1 + \frac{2\alpha}{1-\alpha}$. It suffices to show

$$T_d \left(1 + \frac{2\alpha}{1-\alpha} \right) \geq \frac{20}{\varepsilon}$$

so that, for $x \in [\ell, r]$ (which satisfies $\psi(x) \in [-1, 1]$ and therefore $T_d(\psi(x)) \in [-1, 1]$),

$$P_d(x) = -\delta T_d(\psi(x)) = -\frac{T_d(\psi(x))}{T_d(\psi(0))}$$

is within the interval $[-\varepsilon/20, \varepsilon/20]$. By Constraint I, $\frac{2\alpha}{1-\alpha} \leq 1$ since $\alpha \leq 1/3$, so we may apply Proposition 3.3 to get

$$T_d \left(1 + \frac{2\alpha}{1-\alpha} \right) \geq 2^{c \cdot d \sqrt{\frac{2\alpha}{1-\alpha}} - 1},$$

where $c = \frac{1}{2\ln(2)}$ is the constant from Proposition 3.3. Therefore when $d \geq \frac{2}{c} \sqrt{\frac{r-\ell}{2\ell}} \log(20/\varepsilon) = \frac{2}{c} \sqrt{\frac{1-\alpha}{2\alpha}} \log(20/\varepsilon) \geq \frac{1}{c} \sqrt{\frac{1-\alpha}{2\alpha}} (1 + \log(20/\varepsilon))$ we have the result with $C_d = \frac{2}{c} = 4 \ln(2)$. ■

3.2 Constraint II: Choose m such that $Q(p_i) \approx 1$ when $p_i > r$

Next we need to choose large enough m such that the term e^{-mp_i} cancels out the growth of $P_d(p_i)$ to the right of the “safe interval”, leading to the desired $Q(p_i) = 1 + e^{-mp_i} P_d(p_i) \approx 1$ in this case. (Figure 2 shows $1 + P_d(p_i)$ without the e^{-mp_i} term.)

We place the following constraint:

CONSTRAINT II. We require $m \geq 5.5 \cdot d/(r - \ell)$.

PROPOSITION 3.4. *Assume Constraint II. Then*

$$\forall x \in (r, 1] : |1 - Q(x)| \leq \delta.$$




Figure 2. Untamed right tail.

PROOF. Write $\gamma = 2\frac{x-r}{r-\ell}$ so that

$$-\psi(x) = \frac{2x - r - \ell}{r - \ell} = \frac{r - \ell - 2r + 2x}{r - \ell} = 1 + \frac{2x - 2r}{r - \ell} = 1 + \gamma.$$

Note that $|T_d(z)| = |T_d(-z)|$ for all z . We need to bound $|1 - Q(x)| = e^{-mx} |P_d(x)| = \delta e^{-mx} |T_d(\psi(x))|$ for $x > r$, and $|T_d(\psi(x))| = T_d(-\psi(x)) = T_d(1 + \gamma)$. So we require $e^{mx} \geq T_d(1 + \gamma)$. Bound $T_d(1 + \gamma)$ by

$$\begin{aligned} T_d(1 + \gamma) &= \frac{1}{2} \left((1 + \gamma + \sqrt{2\gamma + \gamma^2})^d + (1 + \gamma - \sqrt{2\gamma + \gamma^2})^d \right) \\ &\leq (1 + \gamma + \sqrt{2\gamma + \gamma^2})^d \leq e^{d(\gamma + \sqrt{2\gamma + \gamma^2})}. \end{aligned}$$

If $\gamma < 1$ then $\gamma + \sqrt{2\gamma + \gamma^2} < 1 + \sqrt{3}$, and therefore since $\frac{x}{r-\ell} > \frac{x}{r} > 1$ we have

$$d(\gamma + \sqrt{2\gamma + \gamma^2}) \leq d(1 + \sqrt{3}) \leq d(1 + \sqrt{3}) \frac{x}{r - \ell} \leq mx.$$

If $\gamma \geq 1$ then $\gamma + \sqrt{2\gamma + \gamma^2} \leq (1 + \sqrt{3})\gamma$, so

$$d(\gamma + \sqrt{2\gamma + \gamma^2}) \leq d(1 + \sqrt{3}) \cdot 2\frac{x-r}{r-\ell} \leq d(1 + \sqrt{3}) \cdot 2\frac{x}{r-\ell} \leq mx.$$

We then conclude

$$T_d(1 + \gamma) \leq e^{d(\gamma + \sqrt{2\gamma + \gamma^2})} \leq e^{mx}. \quad \blacksquare$$

3.3 Constraint III: Variance

We will need the test statistic $\widehat{\mathbf{S}}$ to satisfy two conditions. First, when $|\text{supp}(p)| \leq n$ we need

$$\mathbb{P} \left[\widehat{\mathbf{S}} > (1 + \varepsilon/2)n \right] \leq 1/4.$$

Second, when p is ε -far from having support size n , we need

$$\mathbb{P} \left[\widehat{\mathbf{S}} < (1 + \varepsilon/2)n \right] \leq 1/4.$$

We will bound $\mathbb{E}[\widehat{\mathbf{S}}]$ in each of these cases in Section 3.4. But first we will bound the variance. We impose the following constraint:

CONSTRAINT III. We require $m \leq \frac{1}{4^4} \varepsilon^2 n^2$, and $d^6 9^d \left(\frac{r+\ell}{r-\ell}\right)^{2d-2} \leq \frac{1}{4} m(r-\ell)^2 n^2$.

PROPOSITION 3.5. Assume Constraints I to III. Then $\text{Var}[\widehat{\mathbf{S}}] \leq \varepsilon^2 \frac{n^2}{4^3}$, and therefore (by Chebyshev's inequality),

$$\mathbb{P}\left[|\widehat{\mathbf{S}} - \mathbb{E}[\widehat{\mathbf{S}}]| > \frac{\varepsilon}{4} n\right] \leq \frac{1}{4}.$$

To prove Proposition 3.5, we will require bounds on the values of $f(N_i)$ in the definition of

$$\widehat{\mathbf{S}} = \sum_{i \in \mathbb{N}} (1 + f(N_i)).$$

For this we do some tedious calculations using the coefficients of the polynomial $P_d(x)$. We put these in Appendix A. The result is:

PROPOSITION 3.6. For any d and any $k \in [d]$,

$$|f(k)| \leq \delta \cdot d^2 \cdot 3^d \cdot \left(\frac{2d}{m(r-\ell)}\right)^k \left(\frac{r+\ell}{r-\ell}\right)^{d-k}.$$

We may now estimate the variance.

PROOF OF PROPOSITION 3.5. Recall that each N_i is an independent Poisson random variable $N_i \sim \text{Poi}(mp_i)$. Then

$$\begin{aligned} \text{Var}[\widehat{\mathbf{S}}] &= \sum_{i \in \mathbb{N}} \text{Var}[1 + f(N_i)] = \sum_{i \in \mathbb{N}} \text{Var}[(1 + f(N_i)) \mathbb{1}[N_i > 0]] \\ &\leq \sum_{i \in \mathbb{N}} \mathbb{E}[(1 + f(N_i))^2 \mathbb{1}[N_i > 0]] \leq \max_{j \in [d]} (1 + f(j))^2 \cdot \sum_{i \in \mathbb{N}} \mathbb{E}[\mathbb{1}[N_i > 0]] \\ &= \max_{j \in [d]} (1 + f(j))^2 \cdot \sum_{i \in \mathbb{N}} (1 - e^{-mp_i}) \leq \max_{j \in [d]} (1 + f(j))^2 \cdot m \sum_{i \in \mathbb{N}} p_i = m \cdot \max_{j \in [d]} (1 + f(j))^2. \end{aligned}$$

From Constraint II, which implies $m(r-\ell) > 2d$, the upper bound in Proposition 3.6 is maximized at $k = 1$, so we have

$$\forall j \in [d] : |f(j)| \leq \delta \cdot d^2 \cdot 3^d \cdot \frac{2d}{m(r-\ell)} \cdot \left(\frac{r+\ell}{r-\ell}\right)^{d-1}.$$

If this is at most 1, then the first part of Constraint III gives the desired bound.

$$\text{Var}[\widehat{\mathbf{S}}] \leq 4m \leq \frac{\varepsilon^2 n^2}{4^3}$$

Otherwise, $(1 + |f(j)|)^2 \leq 4f(j)^2$, so (using Constraint I and Proposition 3.2 to ensure $\delta \leq \varepsilon/20 \leq \varepsilon/4^2$), we get from Constraint III that

$$\begin{aligned} \text{Var} \left[\widehat{\mathcal{S}} \right] &\leq 4^2 m \delta^2 \cdot d^6 \cdot 9^d \cdot \frac{1}{m^2 (r - \ell)^2} \left(\frac{r + \ell}{r - \ell} \right)^{2d-2} \\ &\leq \frac{1}{4^2} \varepsilon^2 \cdot d^6 \cdot 9^d \cdot \frac{1}{m (r - \ell)^2} \left(\frac{r + \ell}{r - \ell} \right)^{2d-2} \leq \varepsilon^2 \frac{n^2}{4^3}. \quad \blacksquare \end{aligned}$$

3.4 Constraint IV: $\mathbb{E}[\widehat{\mathcal{S}}] \geq (1 + \varepsilon)n$ when p is ε -far

The above Constraints I and II ensure that $Q(x) \in 1 \pm \delta$ when $x \in [\ell, 1]$. The properties of the Chebyshev polynomial also ensure that $Q(x) \leq 1 + \delta$ for $x < \ell$ (see the behaviour of $Q(x)$ for $x \leq \ell$ in Figure 3). This means $\mathbb{E}[\widehat{\mathcal{S}}] = \sum_{i:p_i>0} Q(p_i)$ is essentially an underestimate of $|\text{supp}(p)|$, which is enough to guarantee that the tester will accept p when $|\text{supp}(p)| \leq n$ with high probability:

LEMMA 3.7. *Assume Constraints I and II. Then*

$$\mathbb{E}[\widehat{\mathcal{S}}] \leq (1 + \delta) |\text{supp}(p)| < (1 + \varepsilon/4) |\text{supp}(p)|.$$

PROOF. First note that $\psi(x) > 1$ for all $x \in (0, \ell)$, so that $P_d(x) = -\delta T_d(\psi(x)) < 0$. Hence $Q(x) = 1 + e^{-mx} P_d(x) < 1$. Combining with Propositions 3.2 and 3.4 we conclude that, for all $x \in (0, 1]$, $Q(x) \leq 1 + \delta$. We also have $Q(0) = 0$. Recalling (4), we obtain

$$\mathbb{E}[\widehat{\mathcal{S}}] = \sum_{i \in \mathbb{N}} Q(p_i) \leq \sum_{i \in \mathbb{N}: p_i > 0} (1 + \delta) = (1 + \delta) |\text{supp}(p)| < (1 + \varepsilon/4) |\text{supp}(p)|,$$

the last inequality since $\delta < \varepsilon/4$ by Proposition 3.2. ■

REMARK 3.8. If we were guaranteed that all nonzero densities satisfied $p_i \geq \ell$, then Lemma 3.7 would also show $\mathbb{E}[\widehat{\mathcal{S}}] \geq (1 - \delta) |\text{supp}(p)|$. In particular, for the support-size *estimation* task, where the nonzero densities satisfy $p_i > 1/n$, we could simply set $\ell = 1/n$, skip the remainder of this section, and recover the optimal $O\left(\frac{n}{\log n} \log^2(1/\varepsilon)\right)$ bound of [37].

Now we need to ensure that $\mathbb{E}[\widehat{\mathcal{S}}]$ will be large when p is ε -far from having support size n . The difficulty is that the densities p_i may be arbitrarily small and lie outside the “safe interval” $[\ell, r]$ where the polynomial approximation guarantees $Q(p_i) \approx 1$. So we will require properties of the Chebyshev polynomial approximation outside the interval $[\ell, r]$. We impose the following constraint:

CONSTRAINT IV. We require $\ell \leq C_\ell \frac{\varepsilon}{n}$, where C_ℓ is any constant $C_\ell \leq 1/20$.

The constraint will give us the desired guarantee:




Figure 3. $Q(p_i)$ for $p_i < \ell$, and the linear lower bound $Q(p_i) \geq (1 - \delta)\frac{p_i}{\ell}$ in Proposition 3.14.

LEMMA 3.9. *Assume Constraints I, II and IV. Then for every $k \leq n$, if p is ε -far from being supported on k elements,*

$$\mathbb{E} \left[\widehat{\mathbf{S}} \right] > (1 + 3\varepsilon/4)k.$$

REMARK 3.10. Constraint IV and Lemma 3.9 lead to an upper bound of $O\left(\frac{n}{\varepsilon \log n} \log^2(1/\varepsilon)\right)$ instead of $O\left(\frac{n}{\varepsilon \log n} \log(1/\varepsilon)\right)$. We put them here to ease the exposition, since the argument is simpler. The argument for the better bound (using a weaker constraint $\ell \leq C_\ell \frac{\varepsilon}{n} \log(1/\varepsilon)$) is in Section 3.6.

To prove Lemma 3.9, we will use the concavity of $P_d(x)$ on $[0, \ell]$ to lower bound $Q(p_i)$ by the line between the points $(0, 0)$ and $(\ell, Q(\ell))$. This is illustrated in Figure 3, which shows Q compared to the linear lower bound. First, we need a standard fact about the roots of Chebyshev polynomials.

FACT 3.11. *For each $d \in \mathbb{N}$, $T_d(x)$ has d distinct roots, all within $[-1, 1]$.*

COROLLARY 3.12. *For each $d \in \mathbb{N}$, $T_d(x)$ is convex increasing in $[1, +\infty)$.*

PROOF. Since $T_d(x)$ has a positive leading coefficient, all of its first d derivatives are eventually positive nondecreasing as $x \rightarrow +\infty$. Moreover, all of its derivatives have their roots within $[-1, 1]$ by Fact 3.11 and the Gauss-Lucas theorem. Hence all of its first d derivatives are positive on $(1, +\infty)$. ■

COROLLARY 3.13. *P_d is concave increasing on $[0, \ell]$.*

PROOF. This follows from the Corollary 3.12 and the fact that $P_d = -\delta T_d(\psi(x))$, with ψ an affine function mapping $[0, \ell]$ onto $\left[1, 1 + \frac{2\alpha}{1-\alpha}\right]$. ■

PROPOSITION 3.14. *For each $x \in [0, \ell]$, we have $(1 - \delta)\frac{x}{\ell} \leq Q(x) \leq 1$.*

PROOF. Recall that $P_d(0) = -1$ by construction. It is also immediate to check that $T_d(1) = 1$ and hence $P_d(\ell) = -\delta T_d(\psi(\ell)) = -\delta T_d(1) = -\delta$. By Corollary 3.13, $P_d(x)$ is concave increasing (and hence negative) on $[0, \ell]$. Thus for any $x \in [0, \ell]$ we have

$$Q(x) = 1 + e^{-mx} P_d(x) \leq 1 + e^{-mx} \cdot 0 = 1$$

and, on the other hand,

$$Q(x) = 1 + e^{-mx} P_d(x) \geq 1 + P_d\left(\frac{x}{\ell} \cdot \ell + \left(1 - \frac{x}{\ell}\right) \cdot 0\right) \geq 1 + \left(\frac{x}{\ell} \cdot (-\delta) + \left(1 - \frac{x}{\ell}\right) \cdot (-1)\right) = (1 - \delta) \frac{x}{\ell}.$$

■

We now calculate the lower bound on the test statistic by partitioning the densities p_i into weight classes.

DEFINITION 3.15 (Heavy and light elements). An element $i \in \text{supp}(p)$ is called *heavy* if $p_i \geq \ell$; otherwise, if $p_i \in (0, \ell)$, it is called *light*.

LEMMA 3.16 (Refinement of Lemma 3.9). Assume Constraints I, II and IV. For any distribution p , let n_H be the number of heavy elements in $\text{supp}(p)$, and let μ_L be the total probability mass of the light elements in $\text{supp}(p)$. Then

$$\mathbb{E}[\widehat{\mathbf{S}}] \geq (1 - \delta) \left(n_H + \frac{\mu_L}{\ell} \right). \quad (5)$$

In particular, for any $k \leq n$, if p is ε -far from being supported on k elements, then

$$\mathbb{E}[\widehat{\mathbf{S}}] > (1 + 3\varepsilon/4)k.$$

PROOF. Let $H \subset \text{supp}(p)$ denote the set of heavy elements, so that $n_H = |H|$. Recall that, by (4),

$$\mathbb{E}[\widehat{\mathbf{S}}] = \sum_{i \in \mathbb{N}} Q(p_i) = \sum_{i \in H} Q(p_i) + \sum_{i \in \text{supp}(p) \setminus H} Q(p_i).$$

By Propositions 3.2 and 3.4, the first summation is

$$\sum_{i \in H} Q(p_i) \geq \sum_{i \in H} (1 - \delta) = (1 - \delta)n_H.$$

By Proposition 3.14, the second summation is

$$\sum_{i \in \text{supp}(p) \setminus H} Q(p_i) \geq \sum_{i \in \text{supp}(p) \setminus H} (1 - \delta) \frac{p_i}{\ell} = (1 - \delta) \frac{\mu_L}{\ell}.$$

This establishes (5). Now, suppose p is ε -far from being supported on $k \leq n$ elements. We consider two cases.

Case 1. Suppose $\mu_L > \varepsilon/10$. Then, using $\delta \leq \varepsilon/20$ (from Proposition 3.2) and $\ell \leq \frac{\varepsilon}{20n} \leq \frac{\varepsilon}{20k}$ (from Constraint IV),

$$\mathbb{E}[\widehat{\mathbf{S}}] \geq (1 - \delta) \frac{\mu_L}{\ell} > (1 - \varepsilon/20) \cdot \frac{\varepsilon/10 \cdot 20k}{\varepsilon} = (1 - \varepsilon/20) \cdot 2k > (1 + 3\varepsilon/4)k.$$

Case 2. Suppose $\mu_L \leq \varepsilon/10$. Then there must exist a partition $H = R \cup T$ of the heavy elements such that $|R| = k$ and $p(T) > 9\varepsilon/10$, since otherwise p would be ε -close to some distribution supported on at most k elements. In fact, we can choose this partition such that $p(i) > p(j)$ for every $i \in R$ and $j \in T$. It follows that $p_j \leq 1/k$ for every $j \in T$, since otherwise we would have $p(R) > \frac{1}{k} \cdot |R| = 1$. Hence

$$|T| \geq \frac{p(T)}{1/k} > \frac{9\varepsilon/10}{1/k} = \frac{9k\varepsilon}{10}.$$

Thus we have

$$\begin{aligned} \mathbb{E} \left[\widehat{\mathbf{S}} \right] &\geq (1 - \delta)n_H \geq (1 - \varepsilon/20) (|R| + |T|) > (1 - \varepsilon/20) \left(k + \frac{9k\varepsilon}{10} \right) \\ &= (1 - \varepsilon/20) \cdot (1 + 9\varepsilon/10)k > (1 + 9\varepsilon/10 - \varepsilon/20 \cdot 2)k > (1 + 3\varepsilon/4)k. \quad \blacksquare \end{aligned}$$

3.5 Correctness of the tester and optimizing the parameters

Let us repeat our list of constraints:

- Constraint I: $r \geq 3\ell$ and $d \geq C_d \sqrt{\frac{r-\ell}{2\ell}} \log\left(\frac{20}{\varepsilon}\right)$, where $C_d = 4 \ln(2)$.
- Constraint II: $m \geq 5.5d/(r - \ell)$.
- Constraint III: $m \leq \frac{1}{4^4} \varepsilon^2 n^2$ and $d^6 9^d \left(\frac{r+\ell}{r-\ell}\right)^{2d-2} \leq \frac{1}{4} m(r - \ell)^2 n^2$.
- Constraint IV: $\ell \leq C_\ell \frac{\varepsilon}{n}$ for any $C_\ell \leq 1/20$.

We show in Section 3.5.1 that if these constraints are satisfied, then the tester defined in Definition 2.7 will be correct. Then we optimize m according to these constraints in Section 3.5.2.

3.5.1 Correctness

LEMMA 3.17 (Correctness (Poissonized)). *Suppose n, ε satisfy Assumption 2.4, and that Constraints I to IV are satisfied with some sample size $m = m(n, \varepsilon)$. Then the Poissonized test statistic $\widehat{\mathbf{S}}$ in Definition 2.6 satisfies*

$$(1 + 3\varepsilon/4) \min\{\text{eff}_\varepsilon(p) - 1, n\} < \mathbb{E} \left[\widehat{\mathbf{S}} \right] < (1 + \varepsilon/4) |\text{supp}(p)|$$

and

$$\mathbb{P} \left[\left| \widehat{\mathbf{S}} - \mathbb{E} \left[\widehat{\mathbf{S}} \right] \right| > \varepsilon n/4 \right] \leq 1/4.$$

In particular, if p is ε -far from having support size at most n , then $\text{eff}_\varepsilon(p) > n$ (Observation 2.1), so $\widehat{\mathbf{S}} > (1 + \varepsilon/2)n$ with probability at least $3/4$; and if $|\text{supp}(p)| \leq n$ then $\widehat{\mathbf{S}} < (1 + \varepsilon/2)n$ with probability at least $3/4$. So the algorithm in Definition 2.9 is a correct (Poissonized) support-size tester with sample complexity m .

PROOF. Lemma 3.7 gives that $\mathbb{E} \left[\widehat{\mathbf{S}} \right] < (1 + \varepsilon/4) |\text{supp}(p)|$, while Lemma 3.9 gives that $\mathbb{E} \left[\widehat{\mathbf{S}} \right] > (1 + 3\varepsilon/4) \min\{\text{eff}_\varepsilon(p) - 1, n\}$. Combining these with Chebyshev's inequality via Proposition 3.5 gives the result. ■

Translating to the standard (non-Poissonized) testing model using Fact 2.11:

COROLLARY 3.18 (Correctness (Standard)). *Suppose that, for n, ε satisfying Assumption 2.4, Constraints I to IV are satisfied with some sample size $m = m(n, \varepsilon)$. Then (also for n, ε satisfying Assumption 2.4), there is a support-size tester with sample complexity $O(m(n, \varepsilon))$.*

3.5.2 Optimizing the parameters

Now we complete Goal 3.1 and find the setting of parameters which minimizes m while satisfying all of our constraints. The (weaker version of the) main Theorem 1.3 follows from Proposition 3.19, using Corollary 3.18, and Proposition 2.2 to handle small ε that do not satisfy Assumption 2.4.

PROPOSITION 3.19. *Under Assumption 2.4, we may satisfy all of the constraints in such a way that*

$$m(n, \varepsilon) = O\left(\frac{n}{\varepsilon \log n} \log^2\left(\frac{1}{\varepsilon}\right)\right).$$

We first sketch some rough calculations to motivate the asymptotic dependence that each of ℓ, r, d and m should have on n and ε . Combining Constraints I and II, we obtain

$$m \gtrsim \frac{d}{r} \gtrsim \sqrt{\frac{1}{\ell r}} \log\left(\frac{1}{\varepsilon}\right). \quad (6)$$

Thus, to allow m to be as small as possible, we should make ℓ and r as large as possible. We already have the bound $\ell \lesssim \frac{\varepsilon}{n}$ from Constraint IV. As we make r larger, d needs to grow due to Constraint I, which can only happen as long as the bound $d^6 9^d \left(\frac{r+\ell}{r-\ell}\right)^{2d-2} \leq \frac{1}{4} m (r-\ell)^2 n^2$ is preserved. The left-hand side of this inequality is exponential in d while the right-hand side is polynomial in n , suggesting that we need $d = O(\log n)$. Since $d \gtrsim \sqrt{\frac{r}{\ell}} \log\left(\frac{1}{\varepsilon}\right)$, this suggests setting

$$r \approx \ell \cdot \frac{\log^2 n}{\log^2(1/\varepsilon)}.$$

Plugging back into (6), we obtain

$$m \gtrsim \frac{1}{\ell} \cdot \frac{\log^2(1/\varepsilon)}{\log n} \approx \frac{n}{\varepsilon \log n} \log^2\left(\frac{1}{\varepsilon}\right). \quad (7)$$

REMARK 3.20. In Section 3.6 we replace Constraint IV with the looser Constraint IVb that $\ell \leq C_\ell \frac{\varepsilon}{n} \log(1/\varepsilon)$, which we can see from Equation (7) should improve m by a $\log(1/\varepsilon)$ factor.

Let us now make this plan rigorous. (We have not tried too hard to optimize the constants.)

PROOF. We require Assumption 2.4 with $a := 1/128$. We set

— $\ell := C_\ell \frac{\varepsilon}{n}$ for $C_\ell := 1/20$.

- $r := C_r \frac{\varepsilon \log^2 n}{n \log^2(1/\varepsilon)}$, for constant $C_r := 4a^2 C_\ell$.
- $d := \lceil C_d \sqrt{\frac{r-\ell}{2\ell}} \log\left(\frac{20}{\varepsilon}\right) \rceil$, recalling that $C_d = 4 \ln(2)$.
- $m := 5.5 \frac{d}{r-\ell}$.

We claim that Constraints I to IV are satisfied. Constraints II and IV are immediately satisfied by the choices of m and ℓ , while Constraint I is satisfied by the choice of d , and because $r \geq 4\ell$ since

$$\frac{r}{\ell} = \frac{C_r \log^2 n}{C_\ell \log^2(1/\varepsilon)} > 4a^2 \cdot \frac{\log^2 n}{(a \log n)^2} = 4.$$

Let us now verify Constraint III. Start with a calculation of $\frac{d}{r-\ell}$, which also gives the claimed bound of $m = O\left(\frac{n}{\varepsilon \log n} \log^2(1/\varepsilon)\right)$:

$$\begin{aligned} \frac{d}{r-\ell} &\leq \frac{2C_d \sqrt{\frac{r-\ell}{2\ell}} \log\left(\frac{20}{\varepsilon}\right)}{r-\ell} = \frac{2C_d \log\left(\frac{20}{\varepsilon}\right)}{\sqrt{2\ell}(r-\ell)} \leq \frac{2C_d \log\left(\frac{20}{\varepsilon}\right)}{\sqrt{2\ell} \cdot r/2} = \frac{2C_d \log\left(\frac{20}{\varepsilon}\right)}{\sqrt{\frac{\varepsilon}{n}} \cdot C_\ell \cdot C_r \frac{\varepsilon \log^2 n}{n \log^2(1/\varepsilon)}} \\ &= \frac{2C_d}{\sqrt{C_\ell C_r}} \cdot \frac{n}{\varepsilon} \cdot \frac{\log(1/\varepsilon) \log(20/\varepsilon)}{\log n} \leq \frac{12C_d}{\sqrt{C_\ell C_r}} \cdot \frac{n}{\varepsilon} \cdot \frac{\log^2(1/\varepsilon)}{\log n} = \frac{6C_d}{C_\ell a} \cdot \frac{n}{\varepsilon} \cdot \frac{\log^2(1/\varepsilon)}{\log n}, \end{aligned}$$

the last inequality because $\varepsilon < 1/2$ by Assumption 2.4, so $\log(1/\varepsilon) > 1$ and hence $\log(20/\varepsilon) = \log(20) + \log(1/\varepsilon) < 5 + \log(1/\varepsilon) < 6 \log(1/\varepsilon)$.

To verify Constraint III, we start with checking the inequality $d^6 9^d \left(\frac{r+\ell}{r-\ell}\right)^{2d-2} \leq \frac{1}{4} m(r-\ell)^2 n^2$. Using $m = 5.5 \frac{d}{r-\ell}$, the right-hand side is at least

$$\frac{1}{4} m(r-\ell)^2 n^2 \geq \frac{1}{4} \cdot 5.5 \cdot d(r-\ell) n^2,$$

so it suffices to prove

$$\frac{d}{r-\ell} d^4 9^d \left(\frac{r+\ell}{r-\ell}\right)^{2d-2} \leq \frac{5.5}{4} n^2.$$

Using the upper bound on $\frac{d}{r-\ell}$ and the upper bound $\frac{r+\ell}{r-\ell} \leq \frac{(5/4)r}{(3/4)r} = 5/3$, it suffices to prove

$$d^4 5^{2(d-1)} \frac{6C_d}{C_\ell} \frac{1}{a} \cdot \frac{n \log^2(1/\varepsilon)}{\varepsilon \log n} \leq \frac{5.5}{4 \cdot 9} n^2 \iff 5^{2(d-1)} \leq \frac{5.5 C_\ell}{4 \cdot 9 \cdot 6C_d} \cdot a \cdot \frac{\varepsilon n}{d^4} \cdot \frac{\log n}{\log^2(1/\varepsilon)}. \quad (8)$$

Upper bound d by

$$\begin{aligned} d-1 &\leq C_d \sqrt{\frac{r-\ell}{2\ell}} \log\left(\frac{20}{\varepsilon}\right) \leq \frac{C_d}{\sqrt{2}} \sqrt{\frac{r}{\ell}} \cdot \log\left(\frac{20}{\varepsilon}\right) \leq \frac{C_d}{\sqrt{2}} \sqrt{\frac{C_r \log(n)}{C_\ell \log(1/\varepsilon)}} \cdot \log\left(\frac{20}{\varepsilon}\right) \\ &= \frac{2C_d}{\sqrt{2}} \cdot a \cdot \frac{\log(n)}{\log(1/\varepsilon)} \cdot \log\left(\frac{20}{\varepsilon}\right) \leq 12 \frac{C_d}{\sqrt{2}} \cdot a \cdot \log n. \end{aligned}$$

Lower bound the right-hand side of (8) using the inequality above via $d \leq 2(d-1)$, and $\varepsilon \geq n^{-a}$, which implies that $\log(1/\varepsilon) \leq a \log n$, to obtain

$$\begin{aligned} \frac{5.5C_\ell}{4 \cdot 9 \cdot 6C_d} \cdot a \cdot \frac{\varepsilon n}{d^4} \cdot \frac{\log n}{\log^2(1/\varepsilon)} &\geq \frac{5.5C_\ell}{4 \cdot 9 \cdot 6C_d} \cdot a \cdot \frac{n^{1-a}}{\left(2 \cdot 12 \frac{C_d}{\sqrt{2}} \cdot a \cdot \log n\right)^4} \cdot \frac{\log n}{a^2 \log^2 n} \\ &= K \cdot \frac{1}{a^5} \cdot n^{1-a} \cdot \frac{1}{\log^5 n} \end{aligned}$$

for constant $K = \frac{5.5 \cdot 4C_\ell}{4 \cdot 9 \cdot 6 \cdot 2^4 \cdot 12^4 C_d^5}$. Hence, returning to (8), it suffices to establish

$$n^{\frac{24 \log(5)}{\sqrt{2}} C_d a} \leq K \cdot \frac{1}{a^5} \cdot n^{1-a} \cdot \frac{1}{\log^5 n}.$$

We have $\frac{24 \log(5)}{\sqrt{2}} C_d a \leq 0.8536$ while $\frac{K}{a^5} \geq 3.2$, and indeed the inequality

$$n^{0.8536} \leq 3.2 \cdot \frac{n^{127/128}}{\log^5 n}$$

holds for sufficiently large n (say, $n \geq 5.4 \cdot 10^{84}$).

It remains to verify the inequality $m \leq \frac{1}{4^4} \varepsilon^2 n^2$. We have

$$m = 5.5 \frac{d}{r-\ell} \leq 5.5 \cdot \frac{6C_d}{C_\ell} \cdot \frac{1}{a} \cdot \frac{n}{\varepsilon} \cdot \frac{\log^2(1/\varepsilon)}{\log n},$$

so (since $\varepsilon \geq n^{-a}$ and hence $\log(1/\varepsilon) \leq a \log n$) it suffices to have

$$\begin{aligned} 5.5 \cdot \frac{6C_d}{C_\ell} \cdot \frac{1}{a} \cdot \frac{n}{\varepsilon} \cdot \frac{\log^2(1/\varepsilon)}{\log n} \leq \frac{1}{4^4} \varepsilon^2 n^2 &\iff \varepsilon^3 \geq 4^4 \cdot 5.5 \cdot 6 \cdot \frac{C_d}{C_\ell} \cdot \frac{1}{a} \cdot \frac{1}{n} \cdot \frac{a^2 \log^2 n}{\log n} \\ &\iff \frac{n^{1-3a}}{\log n} \geq a \cdot 4^4 \cdot 5.5 \cdot 6 \cdot \frac{C_d}{C_\ell} \iff \frac{n^{125/128}}{\log n} \geq 3660, \end{aligned}$$

which holds for sufficiently large n (say, $n \geq 8 \cdot 10^4$). ■

3.6 Constraint IVb: Improvements from loosening constraint IV

In this section we loosen Constraint IV and obtain the improved bound of $O\left(\frac{n}{\varepsilon \log n} \log(1/\varepsilon)\right)$. We replace the earlier constraint of $\ell \leq C_\ell \frac{\varepsilon}{n}$ with:

CONSTRAINT IVb. We require $\ell \leq C_\ell \frac{\varepsilon}{n} \log(1/\varepsilon)$ for any constant C_ℓ satisfying $C_\ell \leq \min\left\{\frac{C_d}{4\sqrt{3}}, \frac{1}{3}\right\}$.

To get the improved upper bound from this constraint, it is sufficient to replace the earlier Lemma 3.9 with:

LEMMA 3.21. *Assume Constraints I, II and IVb. Then for all $\varepsilon \leq 1/3$, and all $k \leq n$, if p is ε -far from being supported on k elements,*

$$\mathbb{E} \left[\widehat{\mathbf{S}} \right] > (1 + 3\varepsilon/4)k.$$

Following the same arguments as in Section 3.5 then leads to our main result:

THEOREM 3.22. *For all $n \in \mathbb{N}$ and $\varepsilon \in (0, 1)$, the sample complexity of testing support size is at most*

$$m(n, \varepsilon) = O\left(\frac{n}{\varepsilon \log n} \cdot \min\left\{\log \frac{1}{\varepsilon}, \log n\right\}\right).$$

PROOF. The proof of correctness under Constraints I to III and IVb is the same as Lemma 3.17, replacing Lemma 3.9 with Lemma 3.21.

The proof of the improved sample complexity follows from the same calculations as in Proposition 3.19: starting with the choice of parameters from Proposition 3.19, we

1. Scale ℓ and r up by a factor of $\log(1/\varepsilon)$.
2. Scale m down by a factor of $\log(1/\varepsilon)$.
3. Leave d unchanged.

It is then straightforward to see that each of Constraints I to III is still satisfied under the new choices of parameters, and that Constraint IVb is indeed also satisfied. ■

To show that Constraint IVb is sufficient, we first define a function $\Phi(\lambda)$ to quantify the worst-case behaviour of the sum $\sum_i Q(p_i)$ (Section 3.6.1), and then prove a lower bound on $\Phi(\lambda)$ (Section 3.6.2).

Some notation. For $x \in (0, \ell)$, we will frequently write $x = \lambda\ell$ for $\lambda \in (0, 1)$ and

$$\gamma = \gamma(\lambda) = \frac{2\alpha}{1-\alpha}(1-\lambda)$$

which satisfies

$$\psi(x) = \frac{r + \ell - 2x}{r - \ell} = 1 + \frac{2\ell - 2\lambda\ell}{r - \ell} = 1 + \frac{2\ell(1-\lambda)}{r - \ell} = 1 + \frac{2\alpha}{1-\alpha}(1-\lambda) = 1 + \gamma. \quad (9)$$

To simplify the analysis, we will define and analyze an auxiliary function Q^* that lower bounds Q . Recall $Q(x) = 1 + e^{-mx}P_d(x)$. We define the function

$$Q^*(x) := \begin{cases} 1 + P_d(x) & \text{if } x < \ell \\ 1 - \delta & \text{if } x \geq \ell, \end{cases}$$

which satisfies:

PROPOSITION 3.23 (Properties of Q^*). *Q^* is continuous, concave and non-decreasing on $[0, 1]$, and for all $x \in (0, \ell]$, $Q(x) \geq Q^*(x)$. Furthermore, if Constraints I and II hold, then $Q(x) \geq Q^*(x)$ for all $x \in (0, 1]$.*

PROOF. For $x \leq \ell$, $P_d(x)$ is concave and increasing (Corollary 3.13), so $Q^*(x) = 1 + P_d(x)$ is also concave increasing. To ensure continuity, it suffices to check the value of Q^* at $x = \ell$. Observe that

$$P_d(\ell) = -\delta T_d(\psi(\ell)) = -\delta T_d(1) = -\delta,$$

so $Q^*(\ell) = 1 - \delta = 1 + P_d(\ell)$, which makes Q^* continuous. Then since $Q^*(x)$ is concave increasing on $x \leq \ell$, and constant $Q^*(x) = Q^*(\ell)$ on $x \geq \ell$, it is concave and non-decreasing on $x \in (0, 1)$.

Now we show $Q(x) \geq Q^*(x)$. If $x \in (0, \ell)$, write $x = \lambda\ell$ so that $\psi(x) = 1 + \gamma$. Then

$$P_d(x) = -\delta T_d(\psi(x)) = -\delta T_d(1 + \gamma),$$

so $P_d(x) \leq 0$ since $T_d(1 + \gamma) \geq 0$ for $\gamma \geq 0$. Since $mx \geq 0$,

$$Q(x) = 1 + e^{-mx} P_d(x) \geq 1 + P_d(x) = Q^*(x).$$

If Constraints I and II hold, then for $x \geq \ell$, by Propositions 3.2 and 3.4,

$$Q(x) \geq 1 - \delta = Q^*(x). \quad \blacksquare$$

Since the quantity $\sum_i Q^*(p_i)$ is invariant under permutations of p , we may assume:

ASSUMPTION 3.24. Without loss of generality, we assume p is sorted, so that $p_1 \geq p_2 \geq \dots$.

3.6.1 Worst-case behaviour and the function $\Phi(\lambda)$

We first show that the worst-case behaviour of Q^* is captured by two variables: the n^{th} -largest density p_n , and the total mass of elements lighter than p_n . Informally, the worst-case distributions for our test statistic are those concentrating essentially all of their mass at p_n .

PROPOSITION 3.25. Let $n \in \mathbb{N}$ and let p be a (sorted) distribution with $p_n > 0$. Let $\mu := \sum_{i>n} p_i$. Then

$$\sum_i Q^*(p_i) \geq \left(n + \frac{\mu}{p_n}\right) Q^*(p_n).$$

PROOF. Recall that Q^* is concave and non-decreasing by Proposition 3.23. Since Q^* is non-decreasing and $p_i \geq p_n$ for any $i \leq n$, we have $Q^*(p_i) \geq Q^*(p_n)$ for any $i \leq n$. On the other hand, for any $i > n$, we have $p_i \leq p_n$ and then, using the concavity of Q^* ,

$$Q^*(p_i) = Q^*\left(\frac{p_i}{p_n} \cdot p_n + \left(1 - \frac{p_i}{p_n}\right) \cdot 0\right) \geq \frac{p_i}{p_n} \cdot Q^*(p_n) + \left(1 - \frac{p_i}{p_n}\right) \cdot Q^*(0) = \frac{p_i}{p_n} \cdot Q^*(p_n),$$

the last equality since $Q^*(0) = 1 + P_d(0) = 0$. We conclude that

$$\sum_i Q^*(p_i) \geq \sum_{i \leq n} Q^*(p_n) + \sum_{i > n} \frac{p_i}{p_n} \cdot Q^*(p_n) = \left(n + \frac{\mu}{p_n}\right) Q^*(p_n). \quad \blacksquare$$

We can now transform the problem of lower bounding $\sum_i Q^*(p_i)$ into the problem of lower bounding another function $\Phi(\lambda)$ defined as:

DEFINITION 3.26. For $\lambda \in (0, 1]$, define

$$\Phi(\lambda) := \left(1 + \frac{\varepsilon}{\lambda \ell n}\right) Q^*(\lambda \ell).$$

We will show a lower bound of $\Phi(\lambda) \geq 1 + 3\varepsilon/4$ below, in Section 3.6.2. Here we show why it works:

LEMMA 3.27. *Assume Constraint I. Suppose $\Phi(\lambda) \geq 1 + 3\varepsilon/4$ for all $\lambda \in (0, 1)$. Then for any $k \leq n$ and any p that is ε -far from having $|\text{supp}(p)| \leq k$,*

$$\sum_i Q^*(p_i) > (1 + 3\varepsilon/4)k.$$

PROOF. Since p is sorted and ε -far from having $|\text{supp}(p)| \leq k$, we have $p_k > 0$ and $\sum_{i>k} p_i > \varepsilon$. We also have $p_k \leq 1/k$ since p is a probability distribution. We consider two cases: $p_k \geq \ell$ and $p_k < \ell$.

First, suppose $p_k \geq \ell$. Then $Q^*(p_k) = 1 - \delta$ by definition, so by Proposition 3.25,


$$\sum_i Q^*(p_i) > \left(k + \frac{\varepsilon}{1/k}\right) (1 - \delta) = (1 + \varepsilon)(1 - \delta)k \geq (1 + 3\varepsilon/4)k,$$

the last inequality since $\delta \leq \varepsilon/20 < \varepsilon/8$ by Proposition 3.2. Otherwise, if $p_k = \lambda\ell$ for some $\lambda \in (0, 1)$, then Proposition 3.25 gives


$$\sum_i Q^*(p_i) > \left(k + \frac{\varepsilon}{\lambda\ell}\right) Q^*(\lambda\ell) = k \left(1 + \frac{\varepsilon}{\lambda k \ell}\right) Q^*(\lambda\ell) \geq \Phi(\lambda) \cdot k \geq (1 + 3\varepsilon/4)k. \quad \blacksquare$$

3.6.2 Lower bound on $\Phi(\lambda)$

Our goal is to prove $\Phi(\lambda) \geq 1 + 3\varepsilon/4$. Figure 4 shows a plot of $\Phi(\lambda)$ compared to $1 + \varepsilon$ for certain example settings of parameters.



(a) $\Phi(\lambda)$ when d is too small.



(b) $\Phi(\lambda)$ under Constraints I and IVb.

Figure 4. The function $\Phi(\lambda)$ with bad and good parameters.

Note that $\Phi(\lambda) = \left(1 + \frac{\varepsilon}{\lambda \ell n}\right) Q^*(\lambda \ell)$ is undefined on $\lambda = 0$. We will proceed in two steps:

1. Show that $\Phi(\lambda) \geq 1 + 3\varepsilon/4$ at $\lambda = 1$ and in the limit $\lambda \rightarrow 0^+$;
2. Show that $\Phi(\lambda) \geq 1 + 3\varepsilon/4$ within $\lambda \in (0, 1)$ by studying the derivative $\Phi'(\lambda)$.

The case $\lambda = 1$ is the easiest to handle:

PROPOSITION 3.28. *Assume Constraints I and IVb. Then for all $n \in \mathbb{N}$ and $\varepsilon \in (0, 1)$, $\Phi(1) \geq 1 + 3\varepsilon/4$.*

PROOF. Note that, by Proposition 3.2, we have $\delta \leq \varepsilon/20 < \varepsilon/8$. When $\lambda = 1$, we have $Q^*(\lambda\ell) = 1 - \delta$ by definition and hence, using the inequality $\ell \leq \frac{\varepsilon}{n} \log(1/\varepsilon)$ obtained from Constraint IVb,

$$\Phi(1) = \left(1 + \frac{\varepsilon}{\ell n}\right) Q^*(\ell) \geq \left(1 + \frac{1}{\log(1/\varepsilon)}\right) (1 - \delta) > (1 + \varepsilon)(1 - \delta) \geq 1 + 3\varepsilon/4. \quad \blacksquare$$

To reason about $\lim_{\lambda \rightarrow 0^+} \Phi(\lambda)$, we will use L'Hôpital's rule, for which we will need bounds on the derivative $T'_d(x)$ of the Chebyshev polynomials. We state the following bound, whose proof proceeds by a direct calculation and is deferred to Appendix B.

PROPOSITION 3.29. *For any $d \in \mathbb{N}$ and $\gamma \in (0, 1)$, we have*

$$T'_d(1 + \gamma) \geq \frac{d}{\sqrt{3\gamma}} (T_d(1 + \gamma) - 1).$$

We may now complete Step 1.

PROPOSITION 3.30. *Assume Constraints I and IVb. Then the limit $\lim_{\lambda \rightarrow 0^+} \Phi(\lambda)$ exists, is finite, and satisfies*

$$\lim_{\lambda \rightarrow 0^+} \Phi(\lambda) \geq 2. \quad (10)$$

PROOF. Recall from Equation (9) that $Q^*(\lambda\ell) = 1 + P_d(\lambda\ell) = 1 - \delta T_d(1 + \gamma)$. Then

$$\begin{aligned} \lim_{\lambda \rightarrow 0^+} \Phi(\lambda) &= \lim_{\lambda \rightarrow 0^+} \left(1 + \frac{\varepsilon}{\lambda \ell n}\right) (1 - \delta T_d(1 + \gamma)) \\ &= \lim_{\lambda \rightarrow 0^+} [1 - \delta T_d(1 + \gamma)] + \frac{\varepsilon}{\ell n} \cdot \lim_{\lambda \rightarrow 0^+} \left[\frac{1 - \delta T_d(1 + \gamma)}{\lambda} \right]. \end{aligned}$$

Since $\gamma = \frac{2\alpha}{1-\alpha}(1 - \lambda)$, and by the definition of δ , the first limit evaluates to

$$1 - \delta T_d\left(1 + \frac{2\alpha}{1-\alpha}\right) = 1 - \delta \cdot \frac{1}{\delta} = 0.$$

As for the second limit, note that both the numerator (for the same reason) and denominator go to 0 as $\lambda \rightarrow 0^+$. Applying L'Hôpital's rule,

$$\begin{aligned} \lim_{\lambda \rightarrow 0^+} \Phi(\lambda) &= \frac{\varepsilon}{\ell n} \cdot \lim_{\lambda \rightarrow 0^+} \frac{\frac{d}{d\lambda}(1 - \delta T_d(1 + \gamma))}{\frac{d}{d\lambda}\lambda} = -\frac{\delta\varepsilon}{\ell n} \cdot \lim_{\lambda \rightarrow 0^+} \frac{d}{d\lambda} T_d\left(1 + \frac{2\alpha}{1-\alpha}(1 - \lambda)\right) \\ &= \frac{\delta\varepsilon}{\ell n} \cdot \frac{2\alpha}{1-\alpha} \cdot \lim_{\lambda \rightarrow 0^+} T'_d\left(1 + \frac{2\alpha}{1-\alpha}(1 - \lambda)\right) = \frac{\delta\varepsilon}{\ell n} \cdot \frac{2\alpha}{1-\alpha} \cdot T'_d\left(1 + \frac{2\alpha}{1-\alpha}\right), \end{aligned}$$

where the last equality holds since $T'_d(x)$ is a polynomial and hence continuous, and this also establishes the existence and finiteness of the limit. Using Proposition 3.29 and the definition

of δ ,

$$\lim_{\lambda \rightarrow 0^+} \Phi(\lambda) \geq \frac{\delta \varepsilon}{\ell n} \cdot \frac{2\alpha}{1-\alpha} \cdot \frac{d}{\sqrt{3 \cdot \frac{2\alpha}{1-\alpha}}} \cdot \left(\frac{1}{\delta} - 1\right) = \sqrt{\frac{2\alpha}{1-\alpha}} \cdot \frac{\delta \varepsilon d}{\ell n \sqrt{3}} \cdot \left(\frac{1}{\delta} - 1\right).$$

Since $\delta \leq 1/2$ by Proposition 3.2, we have $\frac{1}{\delta} - 1 \geq \frac{1}{2\delta}$. Using Constraints I and IVb and the definition $\alpha = \ell/r$, we obtain

$$\lim_{\lambda \rightarrow 0^+} \Phi(\lambda) \geq \sqrt{\frac{2\alpha}{1-\alpha}} \cdot \frac{\delta \varepsilon}{n \sqrt{3}} \cdot \frac{n}{C_{\ell \varepsilon} \log(1/\varepsilon)} \cdot C_d \sqrt{\frac{1-\alpha}{2\alpha}} \log(1/\varepsilon) \cdot \frac{1}{2\delta} \geq 2. \quad \blacksquare$$

REMARK 3.31. The last step in the above proof shows that Constraint IVb is the best possible relaxation of our constraint on ℓ for the present proof strategy: if we had $\ell \gg \frac{\varepsilon}{n} \log(1/\varepsilon)$, then we would only have obtained a $o(1)$ lower bound on $\lim_{\lambda \rightarrow 0^+} \Phi(\lambda)$, but we require a $1 + 3\varepsilon/4$ lower bound.

Thanks to Proposition 3.30, we hereafter consider the continuous extension $\Phi : [0, 1] \rightarrow \mathbb{R}$, i.e. we define $\Phi(0) := \lim_{\lambda \rightarrow 0^+} \Phi(\lambda)$.

The next step is to show that Φ satisfies a certain differential inequality, which will help us conclude that $\Phi(\lambda) \geq 1 + 3\varepsilon/4$ at any critical points where $\Phi'(\lambda) = 0$. It is convenient to define $L := \frac{\ell n}{\varepsilon}$, so that

$$\Phi(\lambda) = \left(1 + \frac{1}{L\lambda}\right) (1 - \delta T_d(1 + \gamma)).$$

LEMMA 3.32. Assume Constraint I. Define $A := \sqrt{\frac{1}{3}} \cdot C_d \log\left(\frac{1}{\varepsilon}\right)$. Then for all $\lambda \in (0, 1)$,

$$\Phi'(\lambda) \geq -\Phi(\lambda) \left(A + \frac{1}{\lambda(L\lambda + 1)}\right) + (1 - \delta)A \left(1 + \frac{1}{L\lambda}\right).$$

PROOF. Recall that $\gamma = \frac{2\alpha}{1-\alpha}(1 - \lambda)$ so $\frac{d}{d\lambda}\gamma = -\frac{2\alpha}{1-\alpha}$. Then

$$\begin{aligned} \Phi'(\lambda) &= \frac{d}{d\lambda} \left[\left(1 + \frac{1}{L\lambda}\right) (1 - \delta T_d(1 + \gamma)) \right] \\ &= -\frac{1}{L\lambda^2} (1 - \delta T_d(1 + \gamma)) + \left(1 + \frac{1}{L\lambda}\right) \cdot (-\delta) \cdot T'_d(1 + \gamma) \cdot \left(-\frac{2\alpha}{1-\alpha}\right) \\ &= \frac{-\frac{1}{L\lambda^2}}{1 + \frac{1}{L\lambda}} \Phi(\lambda) + \delta \cdot \frac{2\alpha}{1-\alpha} \cdot \left(1 + \frac{1}{L\lambda}\right) \cdot T'_d(1 + \gamma) \\ &= -\frac{1}{\lambda(L\lambda + 1)} \Phi(\lambda) + \delta \cdot \frac{2\alpha}{1-\alpha} \cdot \left(1 + \frac{1}{L\lambda}\right) \cdot T'_d(1 + \gamma). \end{aligned}$$

Using Proposition 3.29 and Constraint I and recalling that $T_d(1 + \gamma) \geq 1$,

$$\begin{aligned} \Phi'(\lambda) &\geq -\frac{1}{\lambda(L\lambda + 1)}\Phi(\lambda) + \delta \cdot \frac{2\alpha}{1 - \alpha} \cdot \left(1 + \frac{1}{L\lambda}\right) \cdot \frac{d}{\sqrt{3 \cdot \frac{2\alpha}{1 - \alpha}(1 - \lambda)}} (T_d(1 + \gamma) - 1) \\ &\geq -\frac{1}{\lambda(L\lambda + 1)}\Phi(\lambda) + \delta \cdot \sqrt{\frac{2\alpha}{1 - \alpha}} \cdot \left(1 + \frac{1}{L\lambda}\right) \cdot \frac{C_d \sqrt{\frac{1 - \alpha}{2\alpha}} \log\left(\frac{1}{\varepsilon}\right)}{\sqrt{3}} (T_d(1 + \gamma) - 1) \\ &= -\frac{1}{\lambda(L\lambda + 1)}\Phi(\lambda) + \delta \cdot \left(1 + \frac{1}{L\lambda}\right) \cdot \frac{C_d \log\left(\frac{1}{\varepsilon}\right)}{\sqrt{3}} (T_d(1 + \gamma) - 1) \\ &= -\frac{1}{\lambda(L\lambda + 1)}\Phi(\lambda) + A \left(1 + \frac{1}{L\lambda}\right) \cdot \delta T_d(1 + \gamma) - \delta A \left(1 + \frac{1}{L\lambda}\right). \end{aligned}$$

Now, we may use the definition of Φ to rewrite the middle term:

$$\Phi(\lambda) = \left(1 + \frac{1}{L\lambda}\right) (1 - \delta T_d(1 + \gamma)) \implies \delta T_d(1 + \gamma) = 1 - \frac{\Phi(\lambda)}{1 + \frac{1}{L\lambda}}.$$

Therefore

$$\begin{aligned} \Phi'(\lambda) &\geq -\frac{1}{\lambda(L\lambda + 1)}\Phi(\lambda) + A \left(1 + \frac{1}{L\lambda}\right) \left(1 - \frac{\Phi(\lambda)}{1 + \frac{1}{L\lambda}}\right) - \delta A \left(1 + \frac{1}{L\lambda}\right) \\ &= -\frac{1}{\lambda(L\lambda + 1)}\Phi(\lambda) + A \left(1 + \frac{1}{L\lambda} - \Phi(\lambda)\right) - \delta A \left(1 + \frac{1}{L\lambda}\right) \\ &= -\Phi(\lambda) \left(A + \frac{1}{\lambda(L\lambda + 1)}\right) + (1 - \delta)A \left(1 + \frac{1}{L\lambda}\right). \quad \blacksquare \end{aligned}$$

We may now conclude the argument.

LEMMA 3.33. *Assume Constraints I and IVb. Then for all $n \in \mathbb{N}$ and $\varepsilon \leq 1/3$,*

$$\forall \lambda \in [0, 1] : \Phi(\lambda) \geq 1 + 3\varepsilon/4.$$

PROOF. Note that Propositions 3.28 and 3.30 already give the result for $\lambda = 0$ and $\lambda = 1$, and moreover Φ is continuous on $[0, 1]$ (by Proposition 3.23) and differentiable in $(0, 1)$ (by the differentiability of Q^* in $(0, \ell)$). Thus it remains to check points $\lambda \in (0, 1)$ for which $\Phi'(\lambda) = 0$, if any exist. Suppose $\lambda \in (0, 1)$ satisfies $\Phi'(\lambda) = 0$. Then Lemma 3.32 implies that

$$0 = \Phi'(\lambda) \geq -\Phi(\lambda) \left(A + \frac{1}{\lambda(L\lambda + 1)}\right) + (1 - \delta)A \left(1 + \frac{1}{L\lambda}\right),$$

where again $A := \sqrt{\frac{1}{3}} \cdot C_d \log\left(\frac{1}{\varepsilon}\right)$. Rearranging, we obtain

$$\Phi(\lambda) \geq \frac{(1 - \delta)A \left(1 + \frac{1}{L\lambda}\right)}{\left(A + \frac{1}{\lambda(L\lambda + 1)}\right)} = \frac{(1 - \delta)A \left(\frac{L\lambda + 1}{L\lambda}\right)}{\left(\frac{A\lambda(L\lambda + 1) + 1}{\lambda(L\lambda + 1)}\right)} = \frac{(1 - \delta)(AL^2\lambda^2 + 2AL\lambda + A)}{AL^2\lambda^2 + AL\lambda + L}.$$

Take $K := A/L \geq \frac{C_d}{\sqrt{3}C_\ell}$ which satisfies $K \geq 4$ under Constraint IVb. Then the right-hand side is

$$(1 - \delta) \left(1 + \frac{KL^2\lambda + (K-1)L}{KL^3\lambda^2 + KL^2\lambda + L} \right) = (1 - \delta) \left(1 + \frac{KL\lambda + (K-1)}{KL^2\lambda^2 + KL\lambda + 1} \right).$$

If $L\lambda < 1/K$ then, by preserving only the term $(K-1)$ in the numerator and using the bound $KL^2\lambda^2 \leq KL\lambda \leq 1$ in the denominator, this is at least

$$(1 - \delta) \left(1 + \frac{K-1}{3} \right) \geq (1 - \delta) \cdot 2 \geq 1 + 3\varepsilon/4,$$

the last inequality since $\delta \leq \varepsilon/20$ by Proposition 3.2. Otherwise, by preserving only the term $KL\lambda$ in the numerator and upper bounding each term in the denominator by the maximum of the three, it is at least

$$(1 - \delta) \left(1 + KL\lambda \cdot \min \left\{ \frac{1}{3KL^2\lambda^2}, \frac{1}{3KL\lambda}, \frac{1}{3} \right\} \right) \geq (1 - \delta) \left(1 + \min \left\{ \frac{1}{3L}, \frac{1}{3} \right\} \right) \\ \geq (1 - \delta)(1 + \varepsilon) \geq 1 + 3\varepsilon/4,$$

where the first inequality holds since $3KL\lambda \geq 3$ and $\lambda \leq 1$, and the second inequality holds since $L \leq C_\ell \log(1/\varepsilon) \leq \frac{1}{3\varepsilon}$ by Constraint IVb. ■

We now have all the ingredients to conclude the soundness of the improved tester, and hence finish the proof of the main theorem.

PROOF OF LEMMA 3.21. Combine Equation (4), Proposition 3.23, and Lemmas 3.27 and 3.33 to obtain

$$\mathbb{E} [\widehat{\mathcal{S}}] = \sum_{i \in \mathbb{N}} Q(p_i) \geq \sum_{i \in \mathbb{N}} Q^*(p_i) > (1 + 3\varepsilon/4)k. \quad \blacksquare$$

4. An Effective Lower Bound on Support Size

To prove Corollary 1.6 from the introduction, which we restate here for convenience, we need to essentially do a binary search for the correct parameters, in order to control the variance (since the variance depends on the parameter n).

COROLLARY 4.1. *There exists an algorithm A which, given inputs $n \in \mathbb{N}$, $\varepsilon \in (0, 1/3)$ and sample access to unknown distribution p over \mathbb{N} , draws $O\left(\frac{n}{\varepsilon \log n} \cdot \min\{\log(1/\varepsilon), \log n\}\right)$ independent samples from p and outputs a number $\widehat{\mathcal{S}}$ which satisfies (with probability at least $3/4$ over the samples)*

$$\min\{\text{eff}_\varepsilon(p), n\} \leq \widehat{\mathcal{S}} \leq (1 + \varepsilon)|\text{supp}(p)|.$$

PROOF. The algorithm proceeds as follows:

1. For $i = 0$ to $\log n$, perform the following.
 - a. Set $n_i := n/2^i$ and $\delta_i = \frac{1}{4 \cdot 2^{i+1}}$.

- b. If n_i, ε do not satisfy Assumption 2.4, use $O(m(n_i, \varepsilon) \cdot \log(1/\delta_i))$ samples to obtain an estimate $\widehat{\mathbf{S}}_i$ satisfying the conditions of Proposition 2.2 with probability at least $1 - \delta_i$ (using the median trick to boost the error). Output $\widehat{\mathbf{S}}_i$ and terminate.
 - c. Otherwise, if n_i, ε satisfy Assumption 2.4, use $O(m(n_i, \varepsilon) \cdot \log(1/\delta_i))$ samples obtain an estimate of the test statistic $\widehat{\mathbf{S}}_i$ from Definition 2.6, which (again using the median trick) satisfies the conditions of Lemma 3.17 with probability at least $1 - \delta_i$. If $\widehat{\mathbf{S}}_i \geq n_{i+1}$ output $\widehat{\mathbf{S}}_i$ and terminate. Otherwise continue.
2. If the algorithm did not terminate in $\log n$ steps, then output 1.

The number of samples used by the algorithm is at most (for some constant $C > 0$):

$$\begin{aligned}
& C \cdot \sum_{i=0}^{\log n} \frac{n_i}{\log(n_i)} \log\left(\frac{1}{\delta_i}\right) \frac{1}{\varepsilon} \cdot \min\{\log(1/\varepsilon), \log(n_i)\} \\
&= C \cdot \frac{1}{\varepsilon} \sum_{i=0}^{\frac{1}{2}\log n} \frac{n/2^i}{\log(n) - i} (i+3) \cdot \min\{\log(1/\varepsilon), \log(n) - i\} \\
&\quad + C \cdot \frac{1}{\varepsilon} \sum_{i=1+\frac{1}{2}\log n}^{\log n} \frac{n/2^i}{\log(n) - i} (i+3) \cdot \min\{\log(1/\varepsilon), \log(n) - i\} \\
&\leq C \cdot \frac{2n}{\varepsilon \log n} \min\{\log(1/\varepsilon), \log(n)\} \sum_{i=0}^{\frac{1}{2}\log n} \frac{1}{2^i} (i+3) + 2C \cdot \frac{1}{\varepsilon} \sqrt{n} \log n \\
&= O\left(\frac{n}{\varepsilon \log n} \min\{\log(1/\varepsilon), \log n\}\right).
\end{aligned}$$

The probability that any of the estimates $\widehat{\mathbf{S}}_i$ fails to satisfy the conditions in Lemma 3.17 or Proposition 2.2 (whichever corresponds to the estimator used in the i^{th} step) is at most $\sum_{i=0}^{\log n} \delta_i \leq \frac{1}{4}$. Assuming that every $\widehat{\mathbf{S}}_i$ satisfies its corresponding condition, we separately prove the upper and lower bounds on the output of the algorithm.

CLAIM 4.2. *The output satisfies $\widehat{\mathbf{S}} \leq (1 + \varepsilon) |\text{supp}(p)|$.*

Proof. If the algorithm reaches its last step and outputs 1, there is nothing to show. Similarly, if it outputs an estimate $\widehat{\mathbf{S}}_i$ coming from Proposition 2.2, then $\widehat{\mathbf{S}}_i \leq |\text{supp}(p)|$ and we are done.

The remaining case is that the algorithm terminates at some step i^* by outputting an estimate $\widehat{\mathbf{S}}_{i^*}$ satisfying the conditions of Lemma 3.17. We claim that $n_{i^*} \leq 3|\text{supp}(p)|$. Indeed, assuming that $n_{i^*} > 3|\text{supp}(p)|$ for a contradiction, Lemma 3.17 implies that

$$\begin{aligned}
\widehat{\mathbf{S}}_{i^*} &< (1 + \varepsilon/4) |\text{supp}(p)| + \frac{\varepsilon n_{i^*}}{4} < (1 + \varepsilon/4) \cdot \frac{n_{i^*}}{3} + \frac{\varepsilon n_{i^*}}{4} \\
&= n_{i^*} \left(\frac{1}{3} + \frac{\varepsilon}{12} + \frac{\varepsilon}{4} \right) = n_{i^*} \left(\frac{1}{3} + \frac{\varepsilon}{3} \right) < \frac{n_{i^*}}{2} = n_{i^*+1},
\end{aligned}$$

contradicting the assumption that the algorithm terminates at step i^* . Hence $n_{i^*} \leq 3|\text{supp}(p)|$. Thus Lemma 3.17 gives

$$\widehat{\mathcal{S}}_i < (1 + \varepsilon/4)|\text{supp}(p)| + \frac{\varepsilon n_{i^*}}{4} \leq (1 + \varepsilon/4)|\text{supp}(p)| + \frac{3\varepsilon}{4}|\text{supp}(p)| = (1 + \varepsilon)|\text{supp}(p)|. \quad \blacklozenge$$

CLAIM 4.3. *The output satisfies $\widehat{\mathcal{S}} \geq \min\{\text{eff}_\varepsilon(p), n\}$.*

Proof. We first consider the edge case where $\text{eff}_\varepsilon(p) = 1$. In this case, the output satisfies $\widehat{\mathcal{S}} \geq \text{eff}_\varepsilon(p)$ since the algorithm always outputs at least 1. Going forward, suppose $\text{eff}_\varepsilon(p) \geq 2$.

Now let i^* be the smallest non-negative integer such that $\text{eff}_\varepsilon(p) - 1 \geq n_{i^*+1}$, which exists because $n_{\log(n)+1} = \frac{1}{2} \leq \text{eff}_\varepsilon(p) - 1$, so $i^* = \log n$ satisfies the condition. We consider two cases.

First, suppose that $\text{eff}_\varepsilon(p) - 1 \geq n_{i^*}$, which implies that $i^* = 0$ by the minimality of i^* . Then $n = n_0 \leq \text{eff}_\varepsilon(p) - 1$, so p is ε -far from having support size n by Observation 2.1, and we conclude that the algorithm terminates with output satisfying $\widehat{\mathcal{S}}_0 \geq (1 + \varepsilon/2)n$ (by Lemma 3.17) or $\widehat{\mathcal{S}}_0 \geq n$ (by Proposition 2.2).

Otherwise, we have that $n_{i^*} > \text{eff}_\varepsilon(p) - 1 \geq n_{i^*+1}$. For all $j < i^*$, if the algorithm terminates on loop j then it outputs $\widehat{\mathcal{S}}_j \geq n_{j+1} \geq n_{i^*} > \text{eff}_\varepsilon(p) - 1$, so that $\widehat{\mathcal{S}}_j \geq \text{eff}_\varepsilon(p)$.

Finally, suppose the algorithm reaches loop i^* . If n_{i^*}, ε do not satisfy Assumption 2.4, then the output $\widehat{\mathcal{S}}_{i^*}$ satisfies $\widehat{\mathcal{S}}_{i^*} \geq \min\{\text{eff}_\varepsilon(p), n_{i^*}\} = \text{eff}_\varepsilon(p)$ by Proposition 2.2.

Otherwise, n_{i^*}, ε satisfy Assumption 2.4, which implies that $\text{eff}_\varepsilon(p) - 1 \geq n_{i^*+1} > 4/\varepsilon$; then, since $\min\{\text{eff}_\varepsilon(p) - 1, n_{i^*}\} = \text{eff}_\varepsilon(p) - 1$, we conclude from Lemma 3.17 that

$$\begin{aligned} \widehat{\mathcal{S}}_{i^*} &> (1 + 3\varepsilon/4)(\text{eff}_\varepsilon(p) - 1) - \varepsilon n_{i^*}/4 = (1 + 3\varepsilon/4)(\text{eff}_\varepsilon(p) - 1) - \varepsilon n_{i^*+1}/2 \\ &\geq (1 + \varepsilon/4)(\text{eff}_\varepsilon(p) - 1) > \text{eff}_\varepsilon(p) - 1 + \frac{\varepsilon}{4} \cdot \frac{4}{\varepsilon} = \text{eff}_\varepsilon(p), \end{aligned}$$

which is larger than n_{i^*+1} , so the algorithm terminates with output satisfying $\widehat{\mathcal{S}}_{i^*} > \text{eff}_\varepsilon(p)$. \blacklozenge

This completes the proof of the corollary. \blacksquare

5. Testing Support Size of Functions

For any $n \in \mathbb{N}$, we write \mathcal{H}_n for the set of functions $f: \mathbb{N} \rightarrow \{0, 1\}$ which satisfy $|f^{-1}(1)| \leq n$. We will refer to pairs (f, p) , consisting of a function $f: \mathbb{N} \rightarrow \{0, 1\}$ and a probability distribution p over \mathbb{N} , as *function-distribution* pairs. A function-distribution pair (f, p) is ε -far from \mathcal{H}_n if

$$\forall h \in \mathcal{H}_n : \mathbb{P}_{x \sim p} [f(x) \neq h(x)] \geq \varepsilon.$$

For any multiset $S \subset \mathbb{N}$ and function $f: \mathbb{N} \rightarrow \{0, 1\}$, we will write S_f for the labeled multiset

$$S_f := \{(x, f(x)) : x \in S\}.$$

DEFINITION 5.1 (Testing Support Size of Functions). A support-size tester for functions, with sample complexity $m(n, \varepsilon, \sigma)$, is an algorithm B which receives as input the parameters $n \in \mathbb{N}$, $\varepsilon \in (0, 1)$, and $\sigma \in (0, 1)$, and is required to satisfy the following. For any function-distribution pair (f, p) consisting of a function $f: \mathbb{N} \rightarrow \{0, 1\}$ and a probability distribution p over \mathbb{N} , if \mathbf{S}_f is a labeled sample of size $m = m(n, \varepsilon, \sigma)$ drawn from p , then

1. If $f \in \mathcal{H}_n$ then $\mathbb{P} [B(\mathbf{S}_f) \text{ outputs Accept}] \geq \sigma$; and
2. If (f, p) is ε -far from \mathcal{H}_n then $\mathbb{P} [B(\mathbf{S}_f) \text{ outputs Reject}] \geq \sigma$.

We write $m^{\text{FUN}}(n, \varepsilon, \sigma)$ for the optimal sample complexity of a support-size tester for functions.

We will also write $m^{\text{DIST}}(n, \varepsilon, \sigma)$ for the optimal sample complexity of support-size testing of distributions, with parameters n and ε , and success probability σ (replacing $3/4$ in Definition 1.2). The next two propositions establish Theorem 1.8 from the introduction.

PROPOSITION 5.2. For any $n \in \mathbb{N}$ and $\varepsilon \in (0, 1)$, and any $\sigma \in (0, 1)$,

$$m^{\text{DIST}}(n, \varepsilon, \sigma) \leq m^{\text{FUN}}(n, \varepsilon, \sigma).$$

PROOF. We reduce testing support size of distributions, to testing support size of functions. Let B be a distribution-free sample-based support-size tester for functions on domain \mathbb{N} , with sample complexity $m = m^{\text{FUN}}(n, \varepsilon)$. Then, given n and ε , the support-size tester for input distribution p is as follows.

1. Let \mathbf{S} be m independent random samples from p .
2. Output $B(\mathbf{S}_f)$, where the sample \mathbf{S} is labeled by the constant function $f(x) = 1$.

Suppose that p is ε -far from having $|\text{supp}(p)| \leq n$. Then for any function $g: \mathbb{N} \rightarrow \{0, 1\}$ with $|g^{-1}(1)| \leq n$, it must be the case that $\sum_{i \notin g^{-1}(1)} p_i \geq \varepsilon$. So

$$\mathbb{P}_{\mathbf{S}} [B(\mathbf{S}_f) \text{ outputs Reject}] \geq 3/4.$$

Now suppose p has $|\text{supp}(p)| \leq n$. Let $g: \mathbb{N} \rightarrow \{0, 1\}$ be the function $g(i) = \mathbb{1} [i \in \text{supp}(p)]$, so $|g^{-1}(1)| \leq n$. Then

$$\mathbb{P}_{\mathbf{S}} [B(\mathbf{S}_g) \text{ outputs Accept}] \geq 3/4.$$

Since $\sum_{i \notin g^{-1}(1)} p_i = 0$, the labeled sample \mathbf{S}_f has the same distribution as the labeled sample \mathbf{S}_g . So

$$\mathbb{P}_{\mathbf{S}} [B(\mathbf{S}_f) \text{ outputs Accept}] \geq 3/4. \quad \blacksquare$$

PROPOSITION 5.3. For any $n \in \mathbb{N}$ and $\varepsilon \in (0, 1)$, and any $\sigma, \xi \in (0, 1)$ which satisfy $\sigma + \xi < 1$,

$$m^{\text{FUN}}(n, \varepsilon, \sigma) \leq m^{\text{DIST}}(n, \varepsilon, \sigma + \xi) + \frac{\ln(2/\xi)}{\varepsilon}.$$

PROOF. Let A be the distribution support-size tester with sample complexity m^{DIST} . On input parameters n and ε , given sample access to the function-distribution pair (f, p) , we define the tester B which performs the following.

1. Take $m_1 = \frac{\ln(2/\xi)}{\varepsilon}$ random labeled samples $\mathbf{S}_f^{(1)}$. If $\mathbf{S}_f^{(1)} \cap f^{-1}(1) = \emptyset$, output Accept.
2. Let $\mathbf{z} \sim \mathbf{S}^{(1)} \cap f^{-1}(1)$ be chosen uniformly at random from the elements of the multiset $\mathbf{S}^{(1)} \cap f^{-1}(1)$.
3. Take $m_2 = m^{\text{DIST}}(n, \varepsilon, \sigma + \xi)$ random labeled samples $\mathbf{S}_f^{(2)}$ and let \mathbf{T} be the multiset obtained by replacing each element of $\mathbf{S}^{(2)} \cap f^{-1}(0)$ with \mathbf{z} ; i.e.

$$\mathbf{T} := (\mathbf{S}^{(2)} \cap f^{-1}(1)) \cup \mathbf{Z}$$

where \mathbf{Z} is the multiset containing \mathbf{z} with multiplicity $|\mathbf{S}^{(2)} \cap f^{-1}(0)|$.

4. Output $A(\mathbf{T})$.

To establish correctness of this tester, note that, conditional on the event $E := (\mathbf{S}^{(1)} \cap f^{-1}(1) \neq \emptyset)$, \mathbf{T} is distributed as m_2 independent samples from the distribution $p^{(\mathbf{z})}$ defined as

$$p_i^{(\mathbf{z})} := \begin{cases} p_i & \text{if } i \in f^{-1}(1) \setminus \{\mathbf{z}\} \\ p_i + \sum_{j \in f^{-1}(0)} p_j & \text{if } i = \mathbf{z} \\ 0 & \text{if } i \in f^{-1}(0). \end{cases}$$

If $|f^{-1}(1)| \leq n$ then it is clear that $|\text{supp}(p^{(\mathbf{z})})| \leq n$. If (f, p) is ε -far from satisfying $|f^{-1}(1)| \leq n$, then the sum of the largest n values p_i on domain $i \in f^{-1}(1)$ is at most $\sum_{j \in f^{-1}(1)} p_j - \varepsilon$. Therefore the sum of the largest n values of $p_i^{(\mathbf{z})}$ is at most

$$\sum_{j \in f^{-1}(0)} p_j + \sum_{j \in f^{-1}(1)} p_j - \varepsilon \leq 1 - \varepsilon,$$

so $p^{(\mathbf{z})}$ is ε -far from having $|\text{supp}(p^{(\mathbf{z})})| \leq n$.

Suppose that $|f^{-1}(1)| \leq n$. Write $\rho_E := \mathbb{P}_{\mathbf{S}^{(1)}}[E]$. Then

$$\mathbb{P}_{\mathbf{S}^{(1)}, \mathbf{S}^{(2)}}[B \text{ outputs Accept}] = (1 - \rho_E) + \rho_E \cdot \mathbb{P}_{\mathbf{S}^{(2)}}[A(\mathbf{S}^{(2)}) \text{ outputs Accept}] \geq \sigma + \xi.$$

Now suppose (f, p) is ε -far from satisfying $|f^{-1}(1)| \leq n$. Then

$$\mathbb{P}_{\mathbf{S}^{(1)}, \mathbf{S}^{(2)}}[B \text{ outputs Reject}] = \rho_E \cdot \mathbb{P}_{\mathbf{S}^{(2)}}[A(\mathbf{S}^{(2)}) \text{ outputs Reject}] \geq \rho_E \cdot (\sigma + \xi).$$

Since (f, p) is ε -far from satisfying $|f^{-1}(1)| \leq n$, it must be the case that $\sum_{j \in f^{-1}(1)} p_j \geq \varepsilon$, so

$$\rho_E \geq 1 - (1 - \varepsilon)^{m_1} \geq 1 - e^{-\varepsilon m_1} \geq 1 - \xi/2.$$

Then

$$\mathbb{P}_{\mathbf{S}^{(1)}, \mathbf{S}^{(2)}}[B \text{ outputs Reject}] \geq (1 - \xi/2)(\sigma + \xi) \geq \sigma. \quad \blacksquare$$

References

- [1] Jayadev Acharya, Hirakendu Das, Alon Orlitsky, and Ananda Theertha Suresh. A unified maximum likelihood approach for estimating symmetric properties of discrete distributions. *Proceedings of the International Conference on Machine Learning (ICML)*, volume 70, pages 11–21, 2017. [URL](#) (3)
- [2] Tomer Adar and Eldar Fischer. Refining the adaptivity notion in the huge object model. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, 2024. [DOI](#) (2, 8)
- [3] Tomer Adar, Eldar Fischer, and Amit Levi. Support testing in the huge object model. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, 2024. [DOI](#) (2, 3, 8)
- [4] Eric Blais, Clément L Canonne, and Tom Gur. Distribution testing lower bounds via reductions from communication complexity. *ACM Transactions on Computation Theory (TOCT)*, 11(2):1–37, 2019. [DOI](#) (4)
- [5] Eric Blais, Renato Ferreira Pinto Jr, and Nathaniel Harms. VC dimension and distribution-free sample-based testing. *Proceedings of the ACM SIGACT Symposium on Theory of Computing (STOC)*, 2021. [DOI](#) (5, 6)
- [6] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM (JACM)*, 36(4):929–965, 1989. [DOI](#) (5)
- [7] John Bunge and Michael Fitzpatrick. Estimating the number of species: a review. *Journal of the American statistical Association*, 88(421):364–373, 1993. [DOI](#) (2, 3)
- [8] Clément L Canonne. A Survey on Distribution Testing: Your Data is Big. But is it Blue?, number 9 in Graduate Surveys. *Theory of Computing Library*, 2020. [DOI](#) (3, 10)
- [9] Clément L Canonne, Ilias Diakonikolas, and Alistair Stewart. Testing for families of distributions via the fourier transform. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. [URL](#) (4)
- [10] Renato Ferreira Pinto Jr. and Nathaniel Harms. Distribution testing under the parity trace. *arXiv:2304.01374*, 2023. [URL](#) (6)
- [11] Ronald A Fisher, A Steven Corbet, and Carrington B Williams. The relation between the number of species and the number of individuals in a random sample of an animal population. *The Journal of Animal Ecology*:42–58, 1943. [DOI](#) (2)
- [12] Oded Goldreich. *Introduction to property testing*. Cambridge University Press, 2017. [DOI](#) (2, 3, 5, 41)
- [13] Oded Goldreich. On the complexity of estimating the effective support size. Oded Goldreich, editor, *Computational Complexity and Local Algorithms : On the Interplay Between Randomness and Computation*, pages 204–236. Springer Nature Switzerland, Cham, 2025. [DOI](#) (4)
- [14] Oded Goldreich. Testing bipartiteness in an augmented VDF bounded-degree graph model. *arXiv:1905.03070*, 2019. [URL](#) (4)
- [15] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *Journal of the ACM (JACM)*, 45(4):653–750, 1998. [DOI](#) (3, 5)
- [16] Oded Goldreich and Dana Ron. On sample-based testers. *ACM Transactions on Computation Theory (TOCT)*, 8(2):1–54, 2016. [DOI](#) (5, 6)
- [17] Oded Goldreich and Dana Ron. On the relation between the relative earth mover distance and the variation distance (an exposition). Oded Goldreich, editor, *Computational Complexity and Property Testing: On the Interplay Between Randomness and Computation*, pages 141–151. Springer Cham, 2020. [DOI](#) (41)
- [18] Oded Goldreich and Dana Ron. Testing distributions of huge objects. *TheoretCS*, 2, 2023. [DOI](#) (2, 3, 8, 40)
- [19] I.J. Good. The population frequencies of species and the estimation of population parameters. *Biometrika*:237–264, 1953. [DOI](#) (2)
- [20] Leo A Goodman. On the estimation of the number of classes in a population. *The Annals of Mathematical Statistics*, 20(4):572–579, 1949. [DOI](#) (2)
- [21] Yanjun Han, Jiantao Jiao, and Tsachy Weissman. Local moment matching: a unified methodology for symmetric functional estimation and distribution estimation under Wasserstein distance. *Proceedings of the Conference on Learning Theory (COLT)*, pages 3189–3221. PMLR, 2018. [URL](#) (3, 7, 40)
- [22] Steve Hanneke. CSTheory stackexchange answer, 2019. [URL](#) (5)
- [23] Steve Hanneke. The optimal sample complexity of PAC learning. *Journal of Machine Learning Research*, 17(38):1–15, 2016. [URL](#) (5)
- [24] Yi Hao and Alon Orlitsky. The broad optimality of profile maximum likelihood. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. [URL](#) (3, 7)
- [25] Yi Hao and Alon Orlitsky. Unified sample-optimal property estimation in near-linear time. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. [URL](#) (2, 3)
- [26] Esty Kelman, Ephraim Linder, and Sofya Raskhodnikova. Online versus offline adversaries in property testing. *Proceedings of the Innovations in Theoretical Computer Science Conference (ITCS)*, 65:1–65:18, 2025. [DOI](#) (2)

- [27] Shyam Narayanan and Jakub Tětek. Estimating the effective support size in constant query complexity. *Symposium on Simplicity in Algorithms (SOSA)*, pages 242–252, 2023. DOI (4)
- [28] OEIS. OEIS sequence A008310. URL (37)
- [29] Michal Parnas, Dana Ron, and Ronitt Rubinfeld. Tolerant property testing and distance approximation. *Journal of Computer and System Sciences*, 72(6):1012–1042, 2006. DOI (7)
- [30] Sofya Raskhodnikova, Dana Ron, Amir Shpilka, and Adam Smith. Strong lower bounds for approximating distribution support size and the distinct elements problem. *SIAM Journal on Computing*, 39(3):813–842, 2009. DOI (2, 3)
- [31] Dana Ron and Asaf Rosin. Almost optimal distribution-free sample-based testing of k -modality. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, 2020. DOI (5)
- [32] Dana Ron and Asaf Rosin. Optimal distribution-free sample-based testing of subsequence-freeness with one-sided error. *ACM Transactions on Computation Theory (TOCT)*, 14(1):1–31, 2022. DOI (5)
- [33] Gregory Valiant and Paul Valiant. Estimating the unseen: an $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new CLTs. *Proceedings of the ACM SIGACT Symposium on Theory of Computing (STOC)*, 2011. DOI (2, 3, 7, 10, 40)
- [34] Gregory Valiant and Paul Valiant. Estimating the unseen: improved estimators for entropy and other properties. *Journal of the ACM (JACM)*, 2017. DOI (2, 3, 5, 7, 10, 40)
- [35] Gregory Valiant and Paul Valiant. Instance optimal learning of discrete distributions. *Proceedings of the ACM SIGACT Symposium on Theory of Computing (STOC)*, 2016. DOI (3, 7, 41)
- [36] Gregory Valiant and Paul Valiant. The power of linear estimators. *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*, 2011. DOI (10)
- [37] Yihong Wu and Pengkun Yang. Chebyshev polynomials, moment matching, and optimal estimation of the unseen. *The Annals of Statistics*, 47(2):857–883, 2019. DOI (2, 3, 5–7, 10, 11, 17, 38)
- [38] Yihong Wu and Pengkun Yang. Polynomial methods in statistical inference: theory and practice. *Foundations and Trends in Communications and Information Theory*. now Publishers, 2020. DOI (7)

A. Coefficients of P_d and Values of f

PROPOSITION A.1 (Chebyshev Polynomial Coefficients [28]). *The j^{th} coefficient b_j of $T_d(x)$ is 0 when j has opposite parity as d . Otherwise*

$$b_j = 2^{j-1} \cdot d \cdot (-1)^{\frac{d-j}{2}} \cdot \frac{\left(\frac{d+j}{2} - 1\right)!}{\left(\frac{d-j}{2}\right)! \cdot j!}.$$

We put a bound on the largest coefficient of $T_d(x)$. The leading coefficient is 2^{d-1} , so the bound will necessarily be exponential in d .

PROPOSITION A.2. *For every $d \geq 1$ and every $j \leq d$, the j^{th} coefficient of $T_d(x)$ satisfies $|b_j| \leq d \cdot 3^d \leq 9^d$.*

PROOF. Let $b_j^{(d)}$ be the coefficient of x^j in $T_d(x)$, with $b_{-1}^{(d)} := 0$. We will show by induction that for every $d \geq 0$, $|b_j^{(d)}| \leq \max\{1, d\} \cdot 3^d$. In the base case we have $b_j^{(0)}, b_j^{(1)} \in \{0, 1\}$. By the recursive definition of $T_d(x)$,

$$\begin{aligned} b_j^{(d)} &= 2 \cdot b_{j-1}^{(d-1)} - b_j^{(d-2)} \\ &\leq 2 \cdot |b_{j-1}^{(d-1)}| + |b_j^{(d-2)}| \\ &\leq 2 \cdot \max\{1, d-1\} \cdot 3^{d-1} + \max\{1, d-2\} \cdot 3^{d-2} \\ &\leq 2 \cdot d3^{d-1} + d3^{d-2} = \left(2 + \frac{1}{3}\right) d3^{d-1} \leq d3^d. \end{aligned}$$

The same argument shows $b_j^{(d)} \geq -d3^d$. Now $d \leq 3^d$, so $|b_j^{(d)}| \leq 3^{2d} = 9^d$. ■

We can now calculate the coefficients of P_d in terms of δ and the coefficients b_j of T_d .

PROPOSITION A.3 (Coefficients of P_d). *For given ℓ, r , the polynomial $P_d(x) = \left(\sum_{k=1}^d a_k x^k\right) - 1$ defined by $P_d(x) := -\delta \cdot T_d(\psi(x))$ has coefficients*

$$a_k = (-1)^{k+1} \cdot \delta \cdot 2^k \sum_{j=k}^d b_j \cdot \frac{1}{(r-\ell)^j} \cdot \binom{j}{k} \cdot (r+\ell)^{j-k},$$

where b_j is the j^{th} coefficient of $T_d(x)$.

PROOF. For convenience, write $\alpha := r + \ell$ and $\beta := r - \ell$. Observe that

$$\begin{aligned} \frac{1}{\delta} \cdot P_d(x) &= -T_d(\psi(x)) = -\sum_{j=0}^d b_j \cdot (\psi(x))^j \\ &= -\sum_{j=0}^d b_j \left(-\frac{2x - r - \ell}{r - \ell} \right)^j = -\sum_{j=0}^d b_j \left(\frac{-2x + \alpha}{\beta} \right)^j \\ &= -\sum_{j=0}^d b_j \sum_{k=0}^j \binom{j}{k} \left(\frac{-2x}{\beta} \right)^k \left(\frac{\alpha}{\beta} \right)^{j-k}. \end{aligned}$$

For each fixed k , the coefficient of x^k in this polynomial is

$$-(-2)^k \sum_{j=k}^d b_j \cdot \frac{1}{\beta^j} \cdot \binom{j}{k} \alpha^{j-k} = (-1)^{k+1} 2^k \sum_{j=k}^d b_j \cdot \frac{1}{(r - \ell)^j} \cdot \binom{j}{k} (r + \ell)^{j-k}. \quad \blacksquare$$

We combine these calculations into a formula for the values $f(k)$ in the test statistic (Definition 2.6), for the sake of completeness and to generate Figure 5. A similar figure appears in [37] for their support size estimator.




Figure 5. Example values of $1 + f(N_i)$ in the estimator.

PROPOSITION A.4. For given ℓ, r, m and d , the function f in Definition 2.6 obtained from the polynomial P_d satisfies $f(0) = -1$, $f(k) = 0$ for $k > d$, and for all $k \in [d]$,

$$f(k) = (-1)^{k+1} \cdot \delta \cdot d \cdot \frac{1}{m^k} \sum_{\substack{j=k \\ j \equiv d \pmod{2}}}^d (-1)^{\frac{d-j}{2}} \left(2^{k+j-1} \cdot \frac{\left(\frac{d+j}{2} - 1 \right)!}{\left(\frac{d-j}{2} \right)! (j-k)!} \right) \cdot \frac{(r + \ell)^{j-k}}{(r - \ell)^j}.$$

PROOF.

$$\begin{aligned}
 f(k) &= a_k \frac{k!}{m^k} = (-1)^{k+1} \cdot \delta \cdot 2^k \frac{k!}{m^k} \sum_{j=k}^d b_j \cdot \frac{1}{(r-\ell)^j} \cdot \binom{j}{k} \cdot (r+\ell)^{j-k} \\
 &= (-1)^{k+1} \cdot \delta \cdot 2^k \frac{k!}{m^k} \sum_{\substack{j=k \\ j \equiv d \pmod{2}}}^d (-1)^{\frac{d-j}{2}} \left(2^{j-1} \cdot d \cdot \frac{\left(\frac{d+j}{2} - 1\right)!}{\left(\frac{d-j}{2}\right)! \cdot j!} \right) \cdot \frac{(r+\ell)^{j-k}}{(r-\ell)^j} \cdot \binom{j}{k} \\
 &= (-1)^{k+1} \cdot \delta \cdot d \cdot \frac{1}{m^k} \sum_{\substack{j=k \\ j \equiv d \pmod{2}}}^d (-1)^{\frac{d-j}{2}} \left(2^{k+j-1} \cdot \frac{\left(\frac{d+j}{2} - 1\right)!}{\left(\frac{d-j}{2}\right)! (j-k)!} \right) \cdot \frac{(r+\ell)^{j-k}}{(r-\ell)^j}. \quad \blacksquare
 \end{aligned}$$

PROPOSITION 3.6. (Restated) For any d and any $k \in [d]$,

$$|f(k)| \leq \delta \cdot d^2 \cdot 3^d \cdot \left(\frac{2d}{m(r-\ell)} \right)^k \left(\frac{r+\ell}{r-\ell} \right)^{d-k}.$$

PROOF. Using Proposition A.3 with b_j being the j^{th} coefficient of $T_d(x)$, and using $|b_j| \leq d3^d$ from Proposition A.2, we have

$$\begin{aligned}
 |f(k)| &:= |a_k| \frac{k!}{m^k} = \delta \cdot 2^k \frac{k!}{m^k} \cdot \left| \sum_{j=k}^d b_j \cdot \binom{j}{k} \cdot \frac{(r+\ell)^{j-k}}{(r-\ell)^j} \right| \leq \delta \cdot 2^k \frac{k!}{m^k} \cdot \sum_{j=k}^d |b_j| \cdot \binom{j}{k} \cdot \frac{(r+\ell)^{j-k}}{(r-\ell)^j} \\
 &\leq \delta \cdot d \cdot 3^d \cdot 2^k \frac{k!}{m^k} \cdot \sum_{j=k}^d \binom{j}{k} \cdot \frac{(r+\ell)^{j-k}}{(r-\ell)^j} \\
 &= \delta \cdot d \cdot 3^d \cdot 2^k \frac{k!}{m^k (r-\ell)^k} \sum_{j=k}^d \binom{j}{k} \left(\frac{r+\ell}{r-\ell} \right)^{j-k} \\
 &= \delta \cdot d \cdot 3^d \cdot 2^k \frac{1}{m^k (r-\ell)^k} \sum_{j=k}^d \frac{j!}{(j-k)!} \left(\frac{r+\ell}{r-\ell} \right)^{j-k} \\
 &\leq \delta \cdot d \cdot 3^d \cdot 2^k \frac{1}{m^k (r-\ell)^k} (d-k+1) \frac{d!}{(d-k)!} \left(\frac{r+\ell}{r-\ell} \right)^{d-k} \\
 &\leq \delta \cdot d^2 \cdot 3^d \cdot \left(\frac{2d}{m(r-\ell)} \right)^k \left(\frac{r+\ell}{r-\ell} \right)^{d-k}. \quad \blacksquare
 \end{aligned}$$

B. Derivative of T_d

The derivative of T_d is often expressed in terms of the so-called *Chebyshev polynomials of the second kind*. However, since we only require a simple lower bound, we opt for a direct calculation instead.

PROPOSITION 3.29. (Restated) For any $d \in \mathbb{N}$ and $\gamma \in (0, 1)$, we have

$$T'_d(1+\gamma) \geq \frac{d}{\sqrt{3\gamma}} (T_d(1+\gamma) - 1).$$

PROOF. Note that $T'_d(1 + \gamma) = \frac{d}{d\gamma} T_d(1 + \gamma)$. Using the closed-form formula for T_d , we have

$$\begin{aligned} \frac{d}{d\gamma} T_d(1 + \gamma) &= \frac{1}{2} \left[\frac{d}{d\gamma} \left(1 + \gamma + \sqrt{2\gamma + \gamma^2}\right)^d + \frac{d}{d\gamma} \left(1 + \gamma - \sqrt{2\gamma + \gamma^2}\right)^d \right] \\ &= \frac{1}{2} \left[d \left(1 + \gamma + \sqrt{2\gamma + \gamma^2}\right)^{d-1} \left(1 + \frac{2 + 2\gamma}{2\sqrt{2\gamma + \gamma^2}}\right) \right. \\ &\quad \left. + d \left(1 + \gamma - \sqrt{2\gamma + \gamma^2}\right)^{d-1} \left(1 - \frac{2 + 2\gamma}{2\sqrt{2\gamma + \gamma^2}}\right) \right] \\ &= \frac{d}{2\sqrt{2\gamma + \gamma^2}} \left[\left(1 + \gamma + \sqrt{2\gamma + \gamma^2}\right)^d - \left(1 + \gamma - \sqrt{2\gamma + \gamma^2}\right)^d \right] \\ &\geq \frac{d}{\sqrt{3\gamma}} \left[\frac{\left(1 + \gamma + \sqrt{2\gamma + \gamma^2}\right)^d + \left(1 + \gamma - \sqrt{2\gamma + \gamma^2}\right)^d}{2} - \left(1 + \gamma - \sqrt{2\gamma + \gamma^2}\right)^d \right] \\ &\geq \frac{d}{\sqrt{3\gamma}} (T_d(1 + \gamma) - 1), \end{aligned}$$

where the inequalities used the fact that $\gamma \in (0, 1)$ and hence $1 + \gamma - \sqrt{2\gamma + \gamma^2} \in (0, 1)$. ■

C. Remark on the Bounded Support Size Assumption

As noted in [18, pp. 21], there seems to be no simple way to obtain a support-size tester with sample complexity $O\left(\frac{n}{\varepsilon^2 \log n}\right)$ from the prior literature, without a promise that the true support size is at most $O(n)$. By “simple”, we mean using the known results as a black box without carefully analyzing how the algorithms work. Let us elaborate on this.

C.1 Using histogram learner with TV distance guarantee.

Prior works [33, 34, 21] prove the following (informal) statement: There is an algorithm A which draws m samples from an arbitrary probability distribution p and outputs a sorted distribution q with $q_1 \geq q_2 \geq \dots$, with the following guarantee. If the support size of p satisfies

$$|\text{supp}(p)| \leq O\left(\varepsilon^2 m \log m\right),$$

then the output q satisfies $\|q - p^*\|_1 < \varepsilon$ with probability at least $3/4$, where p^* is the sorted copy of p , i.e. it is a permutation of p such that $p_1^* \geq p_2^* \geq \dots$. In particular, this will hold if $|\text{supp}(p)| \leq n$ and we draw $m = \Theta\left(\frac{n}{\varepsilon^2 \log n}\right)$ samples. Using the typical testing-by-learning technique, this suggests the following tester:

1. Draw $m = O\left(\frac{n}{\varepsilon^2 \log n}\right)$ samples and produce a sorted distribution q . If input p satisfies $|\text{supp}(p)| \leq n$ then (with high probability) we have $\|q - p^*\|_1 < \varepsilon/2$. Output Reject if q is not $\varepsilon/2$ -close to having support size at most n .

2. Verify that q is indeed close to p^* and output Accept if this is the case. In particular, using a *tolerant*² tester to test whether p is ε -close to the set of permutations of q . Unfortunately, the only tolerant tester for this property that we are aware of is itself obtained by the same testing-by-learning approach of approximating the histogram (see e.g. [12]), which is only guaranteed to work when $|\text{supp}(p)| \leq n$.

Algorithms for learning the histogram *should* be able to test support size with no promise on the true support size. One could hope for the following property of the algorithms: Given $m = O\left(\frac{n}{\varepsilon^2 \log n}\right)$ samples, if p is ε -far from having support size at most n , then with high probability the output q is $\Omega(\varepsilon)$ -far from having support size at most n . This does not appear to follow from the proven properties of these algorithms without further analysis.

C.2 Using the histogram learner with relative Earth-mover distance guarantee.

In [35], the assumption of bounded support size is removed from the histogram learner. However, by necessity, this new algorithm's guarantee is not in terms of the TV distance but rather the *truncated relative Earth-mover distance* $R_\tau(\cdot, \cdot)$. We will not formally define this metric here, but we note two of its properties:

1. R_τ ignores densities less than τ , i.e. if p, p' and q, q' are pairs of distributions such that $p_i = p'_i$ whenever $p_i, p'_i > \tau$, and $q_i = q'_i$ whenever $q_i, q'_i > \tau$, then $R_\tau(p, q) = R_\tau(p', q')$.
2. $\text{dist}_{\text{TV}}(p^*, q^*) \leq R_0(p, q)$ (see also the exposition [17]).

The main result in [35] is

THEOREM C.1 (Theorem 2 of [35]). *There exists an algorithm satisfying the following for some absolute constant c , sufficiently large m , and any $w \in [1, \log m]$. Given m independent draws from a distribution p with histogram p^* , with high probability the algorithm outputs a generalized³ histogram q^* satisfying*

$$R_{\frac{w}{m \log m}}(p^*, q^*) \leq \frac{c}{\sqrt{w}}.$$

As an application, [35] give a procedure to estimate the expected number of unique elements that would be seen in a sample of size $m \log m$, given a sample of size m :

THEOREM C.2 (Proposition 1 of [35]). *Given m samples from an arbitrary distribution p , with high probability over the randomness of the samples, one can estimate the expected number of unique elements that would be seen in a set of k samples drawn from p , to within error $k \cdot c \sqrt{\frac{k}{m \log m}}$ for some universal constant c .*

2 A tolerant tester should Accept if the distribution p is ε_1 -close to the class, and Reject if p is ε_2 -far from the set, for $\varepsilon_1 < \varepsilon_2$.

3 A standard histogram maps each density value α to the number of elements i satisfying $p_i = \alpha$. A generalized histogram may map α to non-integral values as well.

We sketch two ways to obtain a support-size tester with sample complexity $\text{poly}(1/\varepsilon) \cdot O(\frac{n}{\log n})$ using these theorems.

1. Estimating Unique Elements. Combining Theorem C.2 with the naïve tester from Section 2.1, which makes its decision based on the number of unique elements observed, gives a support size tester with sample complexity $O(\frac{n}{\varepsilon^5 \log n})$. The idea is to set $m := \Theta(\frac{n}{\varepsilon^5 \log n})$, use Theorem C.2 to obtain an estimate of the expected number of unique elements that would be seen in a sample of size $k := \Theta(n/\varepsilon)$ to within error $\varepsilon n/8$, and accept if and only if this estimate is at most $(1 + \varepsilon/8)n$.

When $|\text{supp}(p)| \leq n$, the expected number of unique elements is at most n , so the tester accepts. Now, suppose p is ε -far from being supported on n elements. We claim that any set of size at most $(1 + \varepsilon/2)n$ misses at least $\varepsilon/2$ mass from p . Note that the n^{th} largest probability mass in p is at most $1/n$, and the remaining elements of smaller probability mass (the *light elements*) make up at least ε mass; thus there are at least εn light elements. Hence the most probability mass that a set of size $(1 + \varepsilon/2)n$ can cover comes from picking the n elements of largest mass plus $\varepsilon n/2$ light elements, for a total of at most $1 - \varepsilon + \frac{\varepsilon n}{2} \cdot \frac{1}{n} = 1 - \varepsilon/2$ mass, i. e. $\varepsilon/2$ mass is missed.

Therefore, an argument similar to the proof of Proposition 2.2 shows that the expected number of unique elements observed in a sample of size $\Theta(n/\varepsilon)$ is at least $(1 + \varepsilon/4)n$, so the tester rejects.

2. Learning the Histogram. Instead of applying the naïve tester on top of the unique elements estimator, one may hope for a more efficient tester by directly analyzing the learned histogram q^* from Theorem C.1. We expect that one could obtain a support size tester with sample complexity $O(\frac{n}{\varepsilon^3 \log n})$ as follows. By setting $w := \Theta(1/\varepsilon^2)$ and $m := \Theta(\frac{n}{\varepsilon^3 \log n})$ (at least when this satisfies the constraint $w \leq \log m$ from Theorem C.1), the algorithm from Theorem C.1 yields q^* satisfying

$$R_{\Theta(\varepsilon/n)}(p^*, q^*) \leq \Theta(\varepsilon). \quad (11)$$

It appears that this suffices for testing support size. However, we will not attempt to prove this, since the bound is worse than our Theorem 1.3, and it is not possible to improve this bound while treating the algorithm as a black box. This is because, if $\tau \gg \varepsilon/n$, we can construct distributions p, q whose histograms satisfy $R_\tau(p^*, q^*) < \varepsilon$ and yet $|\text{supp}(p)| \leq n$ while q is ε -far from having $|\text{supp}(p)| \leq n$.