

Improving Pearson’s chi-squared test: hypothesis testing of distributions - optimally*

Trung Dang^{1,2}, Walter McKelvie^{1,3}, Paul Valiant¹, and Hongao Wang¹

¹Purdue University

²UT Austin

³Columbia University

October 13, 2023

Abstract

Pearson’s chi-squared test, from 1900, is the standard statistical tool for “hypothesis testing on distributions”: namely, given samples from an unknown distribution Q that may or may not equal a hypothesis distribution P , we want to return “yes” if $P = Q$ and “no” if P is far from Q . While the chi-squared test is easy to use, it has been known for a while that it is not “data efficient”, it does not make the best use of its data. Precisely, for accuracy ε and confidence δ , and given n samples from the unknown distribution Q , a tester should return “yes” with probability $> 1 - \delta$ when $P = Q$, and “no” with probability $> 1 - \delta$ when $|P - Q| > \varepsilon$. The challenge is to find a tester with the *best* tradeoff between ε , δ , and n .

We introduce a new tester, efficiently computable and easy to use, which we hope will replace the chi-squared tester in practical use. Our tester is found via a new non-convex optimization framework that essentially seeks to “find the tester whose Chernoff bounds on its performance are as good as possible”. This tester is $1 + o(1)$ optimal, in that the number of samples n needed by the tester is within $1 + o(1)$ factor of the samples needed by *any* tester, even non-linear testers (for the setting: accuracy ε , confidence δ , and hypothesis P). We complement this algorithmic framework with matching lower bounds saying, essentially, that “our tester is instance-optimal, even to $1 + o(1)$ factors, to the degree that Chernoff bounds are tight”. Our overall non-convex optimization framework extends well beyond the current problem and is of independent interest.

1 Introduction

In this paper we consider the statistical problem of “hypothesis testing of distributions”, and provide a principled yet practical new approach to optimal performance. The standard approach in practice is Pearson’s chi-squared test (described below), invented in 1900, and remaining one of the most widely used tests in statistics, a bedrock of the field. Nonetheless, the chi-squared test was invented before the advent of principled methodical approaches to data-efficient statistics, and from this perspective looks a bit ad-hoc. Since 1900, many post-hoc justifications of the chi-squared test have been developed, as well as “patches” to tailor its performance to certain situations.

*Work partially funded by NSF award CCF-2127806




Figure 1: *Performance of Various Uniformity Testers.* Plots show the result of numerical experiments comparing the performance of our new tester to 3 standard testers, that respectively threshold the chi-squared metric, the total variation distance, and the number of singleton elements in the data. The left plot shows $\epsilon = 0.9$, with error δ (maximum of Type I and Type II errors) plotted as a function of $n = k$. The yellow dotted line, expressing a brute force computation of the ground truth optimal tester, is visually synonymous with our tester for the range of parameters for which we could compute it, $k = n \leq 50$. The right plot shows similar plots for $\epsilon = 0.7$, and $\epsilon = 0.45$, for larger ranges of $n = k$. In all cases in this paper, our tester was optimized and computed in $< 1s$ in Matlab. The choice of threshold for each tester significantly affects its performance, and in all cases (except [VV17] which specifies its thresholds) we were as generous as possible, choosing the threshold of lowest error rate.

The most standard TCS model that captures many of the use cases of the chi-squared test is as follows: we start with a known *hypothesis* distribution P of support size n ; we receive k samples from an unknown distribution Q , and wish to return a “yes” or “no” answer distinguishing, as well as possible, whether our hypothesis P is correct—namely $P = Q$ —versus far from correct. Explicitly, given an ℓ_1 distance bound ϵ , and a failure probability δ , we want to say “yes” with probability $\geq 1 - \delta$ when the samples are from P , and we want to say “no” with probability $\geq 1 - \delta$ when the samples are from *any* distribution Q such that $\|P - Q\|_1 \geq \epsilon$. The goal is to find the algorithm with the best tradeoff between the number of samples k , the accuracy threshold ϵ , and the confidence $1 - \delta$, for any given hypothesis P .

There are thus 3 equivalent ways of viewing our problem: given k data points and an allowed failure rate δ , find the tester that optimizes the accuracy ϵ ; alternatively, given k data points and a required accuracy ϵ , find the tester with the smallest failure rate δ ; or, finally, given an allowed failure probability δ and a required accuracy ϵ , find the minimum amount of data k needed. Algorithms for any of these formulations can be transparently reparameterized/reformulated to solve the problem from the other perspectives (via a black-box binary search reduction).

We set out to solve this bedrock problem of statistics, in a way that is practical, robust, and extensible.

Our algorithm: This paper presents a new algorithm for this problem, arising from a novel




Figure 2: Performance of Testers on a Nonuniform Distribution with a Single Heavy Element. *Plots show performance of testers on distributions composed of one element with weight $\frac{1}{2}$ and n elements each with weight $\frac{1}{2n}$, with $\epsilon = 0.9$ and k ranging over 20, 30, or 40.*

optimization framework. Our algorithm makes optimal use of its data, even up to $1 + o(1)$ factors: phrased in terms of optimizing the failure probability in terms of a given number of samples k and accuracy ϵ , we show that the log failure rate $\log \frac{1}{\delta}$ of the tester returned by our algorithm converges to within $1 + o(1)$ factor of optimal (see Theorem 4). This is not just an instance-optimal result, but a sub-constant-factor tight result. Complementing the theoretical analysis, simulations show that our tester outperforms all other testers in a variety of settings. Of particular interest is the simplest (and thus most widely studied) setting of the *uniform* distribution hypothesis. Despite our tester being designed for far more general settings, we significantly outperform all other testers proposed for this iconic setting—most notably the total variation (TV) tester, the collisions tester (which is synonymous with the chi-squared tester in this special case), and the singletons tester. In particular, for small uniform distributions, we can derive, via brute force numerics, the ground-truth optimal tester; as seen in Figure 1, the performance of our algorithm is visually *indistinguishable* from the optimal. The fact that our tester is visually indistinguishable from optimal on small inputs is particularly striking given that we analytically expect its performance to get ever more optimal for larger inputs.

While the chi-squared estimator—and many variants—perform fairly well on uniform distributions, the nonuniform case is far more challenging and exposes giant performance gulfs. See Figure 2, showing a massive practical advantage of our tester over chi-squared. When $n = 80$ and $k = 40$, effectively no meaningful information can be learned using a chi-squared test, yet our test’s error rate remains less than 10%.

Our new algorithmic approach: Our algorithmic approach is surprisingly direct: for the problem of hypothesis testing on distributions, we consider *all* testers from a flexible and general class—called “semilinear testers”, see Definition 1—and return the *best* one, via an optimization setup.

The direct approach is daunting, because describing the performance of an arbitrary algorithm—let alone optimizing it—is already challenging. For an algorithm to be effective, its error rate must be as small as possible, where by statistical convention, we say that a “type I error” is when we accidentally reject the hypothesis $P = Q$ when it is true, and a “type II error” is when we

accidentally accept the hypothesis when $|P - Q| \geq \varepsilon$. The failure rate of our tester is the maximum failure over these two cases. Within the case of type II error, there is a third nested optimization problem, we must necessarily consider the worst-case error over *all* distributions Q that are ε -far from our hypothesis. Finally, even fixing a certain tester, a certain error-type, and a certain alternative distribution Q , we bound the failure rate via a Chernoff bound, which is a 4th level of optimization. In short, we cleanly capture the challenge of designing an optimal tester via a 4-level non-convex optimization setup, which on its surface is not encouraging.

The main point of the paper is that, despite appearances, our optimization problem is 1) algorithmically tractable, admitting effective code; 2) partially analytically tractable, showing that optimal testers have an intriguing “log cosh” shape; 3) naturally provides matching lower bounds, showing that even though we only optimize over semilinear algorithms, we asymptotically achieve the best performance possible from *any* algorithm. The generality of this new framework suggests that such an optimization approach might provide canonical optimal solutions for a much wider array of related problems that we have not begun to explore yet.

Semilinear testers: We now introduce the class of “semilinear” estimators and testers. Semilinear estimators are simple yet flexible, expressing for each domain element, an arbitrary function of the number of times this element is sampled, while summing linearly across different domain elements. The goal of this paper is to find the “right” tester from this class, instead of an ad-hoc tester from a larger class.

Definition 1. Given a sample space (domain) indexed by j , a semilinear estimator is represented by a table with coefficients $c_{i,j}$. Given several samples from a distribution on this domain, we represent the samples by their histogram, with s_j counting the number of times element j was sampled. The semilinear estimator \mathbf{c} , on sample \mathbf{s} , will return $\sum_j c_{s_j,j}$. Namely, for each domain element j , the number s_j records the number of times this element occurred in the sample, and we look up, in the j^{th} column of our table \mathbf{c} , its s_j^{th} entry, and add up all these $c_{s_j,j}$. A semilinear *threshold* tester consists of a semilinear estimator and a threshold γ : it returns “yes” if the semilinear estimate is below the threshold γ , else “no”. Without loss of generality, we may set the threshold γ to be 0.

Example 1. The Pearson chi-squared test, for a given hypothesis P , and given k samples from some distribution Q , represented as a histogram \mathbf{s} , computes the estimate

$$\sum_j \frac{(s_j - kP_j)^2}{P_j}$$

and returns “yes” if it is below some threshold. Namely, it computes the squared difference between the number of times each domain element is seen, versus its expectation (if the hypothesis is true), normalized by its expectation. The Pearson chi-squared test is thus a semilinear tester.

As a related example (discussed more below), the uniformity tester of Gupta and Price [GP22] computes the “Huber statistic” of $s_j - kP_j$ instead of simply squaring it, but since this function can be represented as a lookup table for each potential value of s_j , this statistic is also semilinear. By contrast, the instance-optimal tester of Valiant and Valiant [VV17] is not quite semilinear, because it essentially computes *two* semilinear threshold tests and returns the OR of them.

1.1 Main Results

The main results are an algorithmic upper bound for hypothesis testing on distributions, and a lower bound that matches it, provided that “Chernoff bounds are tight”. While there is a long and celebrated history in statistics of analyzing “reverse” Chernoff bounds, under the name “the large deviations principle”, these tools typically give insight only in the limit as the number of samples goes to infinity, without commenting on the rate of convergence; and without further foundational statistical work in this direction, our lower bounds are necessarily also of this flavor. See Section 5 for more discussion.

Towards this end, we introduce the limit that we use for our lower bounds: we start with a fixed hypothesis distribution P of support size n , and consider taking $\text{Poi}(k)$ samples from it; in the limit, we essentially repeat this process s times, for some positive integer $s \rightarrow \infty$. Technically, this corresponds to taking $\text{Poi}(ks)$ samples from a version of P that has each domain element subdivided into s identical copies of itself:

Definition 2. Given a distribution P supported on n elements, define $P^{\text{sub}(s)}$, read as “ P subdivided s times” to be the distribution supported on ns elements, where each domain element of P is subdivided into s new domain elements of $P^{\text{sub}(s)}$, each of $\frac{1}{s}$ times the original probability mass.

Our main results are about the optimum of the overall optimization problem, which we denote $\Delta(P, k, \varepsilon)$:

Definition 3. Given a distribution P , a bound ε on the ℓ_1 distance that we wish to test for, and a number of samples k , let $\Delta(P, k, \varepsilon)$ be the optimal objective value of the nonconvex optimization problem of Equation 1.

Our main results—including both upper and lower bounds—are summarized in the following theorem.

Theorem 4. *Given a distribution P of support size n , a bound ε on the ℓ_1 distance that we wish to test for, and a number of samples k , then:*

- *In polynomial time, we can compute $\Delta(P, k, \varepsilon)$, along with (a representation of) coefficients $c_{i,j}$ to a semilinear threshold tester such that, given $\text{Poi}(k)$ samples from either P or an arbitrary distribution Q with ℓ_1 distance $\geq \varepsilon$ from P , the tester distinguishes these cases except with error probability $\leq e^{\Delta(P, k, \varepsilon)}$.*
- *For positive integer s , on the problem of testing the hypothesis $P^{\text{sub}(s)}$ with $\text{Poi}(k \cdot s)$ samples, the same tester as above, with coefficients $c_{i,j}$, will have error probability $\leq e^{s \cdot \Delta(P, k, \varepsilon)}$.*
- *By contrast, for **any** sequence of testers T_s indexed by s —including arbitrary nonlinear testers—letting δ_s denote the failure probability of tester s for testing the hypothesis $P^{\text{sub}(s)}$ with $\text{Poi}(k \cdot s)$ samples to accuracy ε , we have that*

$$\liminf_{s \rightarrow \infty} \frac{1}{s} \log \delta_s \geq \Delta(P, k, \varepsilon)$$

1.2 Perspective

1.2.1 The choice of ℓ_1 metric

We briefly comment on the choice of the ℓ_1 metric in defining the hypothesis testing problem: our tester is asked to reject the hypothesis P on any distribution Q such that $|P - Q|_1 \geq \varepsilon$. However many other distance metrics are possible, including the KL-divergence, Hellinger distance, or chi-squared divergence—all of which show up in related statistical contexts. We choose ℓ_1 for this paper both because it is standard for work in the statistical property testing field for related estimation tasks. Further, and importantly, since ℓ_1 distance is equivalent to total variation distance, it exactly captures how distinguishable a single sample from P versus Q is, thus making it a natural benchmark.

However, the optimization framework introduced in this paper is far more general than just the ℓ_1 distance distribution hypothesis testing framework we work with here. Many aspects of our analysis extend unchanged to other metrics, and understanding the implications of this would presumably yield interesting follow-up papers. In short, we choose the ℓ_1 metric but the approach we introduce here is not in any way wed to it.

1.2.2 Comparison with [VV17]

Our main result is a bit unusual in comparison with standard results about testing properties of distributions. Compare, for example, with the main theorem of [VV17] which describes a precise formula for the number of samples needed to run accurate hypothesis testing as a function of the distribution P , and the accuracy ε (that paper assumes $\delta = \Theta(1)$): the main result there characterizes the number of samples in terms of the $\ell_{\frac{2}{3}}$ norm of the hypothesis P . By contrast, our optimization framework is a rather opaque expression compared to the $\ell_{\frac{2}{3}}$ norm result of [VV17].

We view our current approach as furthering and complementing the goals of [VV17], for the following reasons. First, our results here are tighter, both in theory, and in practice. While [VV17] obtained performance that was tight to constant factors, the results here are $1 + o(1)$ tight, and in simulations (see figures in Section 1) are essentially identical to performance of the ground truth optimal algorithm in situations where the ground truth can be numerically ascertained. Constant factors are often critical in data-efficiency contexts, and a focus of increasing recent interest. Second, while the current paper does not provide a clean formula for the performance of our new algorithm—see Equation 1 and below, which presumably have no closed-form solution—the formula from [VV17] is already shown to be constant-factor tight by that paper, so may be used here as a closed-form estimate of the performance of our algorithm as well. The fact that our optimization framework does not appear to have a closed-form solution is not a fault of our paper but a fact of nature: this expression is optimal yet this expression is not simple. This paper provides access to optimality via an algorithm that efficiently finds the testing coefficients and the objective value; and this seems like the best that can be hoped for.

1.2.3 Comparison to [GP22]

This recent paper focuses on the distribution hypothesis testing problem, restricted to the case of *uniform* distributions. The paper shares our focus on $1 + o(1)$ -tight analysis. The main results are identifying an asymptotic regime where Pearson’s chi-squared test performs optimally, and then modifying the chi-squared test—in a semilinear way, to instead compute the Huber function on each

domain element—leading to $1 + o(1)$ -optimal performance across a fairly wide range of parameter space.

The main distinction between our paper and [GP22] is that we consider generic distributions instead of just uniform distributions. It is not at all clear how one could optimally combine uniformity testers to test more general distributions, although our optimization framework in some sense can be viewed as the prescription for this. We briefly point out that non-uniform distributions are (of course) more common than uniform distributions, so uniformity testing can be considered a special case, but whose nice mathematical structure does not always reflect real life. We briefly mention 3 examples where non-uniform hypotheses might arise. Pearson’s chi-squared test is often used to compare two empirically accessed distributions, and in many cases, one of these distributions may be much more expensive to sample than the other; and in this case, one can call the inexpensive distribution the “known” hypothesis P as it is so cheap to obtain farther samples. Beyond empirically obtained distributions, there are also many cases in science where prior knowledge leads one to hypothesize an explicit non-uniform distribution. For example, back in 1928, R.A. Fisher investigated Mendelian genetics and its prediction of non-uniform allele frequencies [FB28]. A much more modern example is the question of how to test whether a quantum computer is accurately sampling from the distribution of outcomes expected by the laws of quantum mechanics, and distinguishing this from the case that the quantum computer is sampling from an erroneous distribution [WZXK21].

As further comparison between our paper and results of [GP22], see the figures in Section 1 which compare the performance of our tester with several analyzed in [GP22]. Because of the asymptotic setup of [GP22], the parameters of the “Huber tester” they recommend are never specified—since the specifics do not matter in the limit considered in that paper. But we are thus unable to compare our algorithm to any specific proposal from [GP22].

At a higher level, the Huber function arises in [GP22] out of the dual requirements that the function is quadratic near the origin and linear away from it. But the Huber function is just one choice of many possible functions that fit these criteria equally well. Our paper shows that in fact the log cosh function is the correct (i.e., optimal) form in general.

Finally, we provide perspective on two of the quirks of our analysis.

1.2.4 Depoissonization

While our main theorem is phrased in terms of a Poisson-distributed number of samples $\text{Poi}(k)$ from a distribution, the more standard setting is to take some fixed number k of samples. As discussed above, we move to the “Poissonized” setting because here, the number of times that each domain element j is sampled becomes completely independent of the other domain elements, which is crucial for both the upper and lower bound analysis. However, the natural conclusion of a “Poissonized” analysis is “depoissonization”—showing how to move these results back to the original framework, with exactly k samples. Handling this in full would substantially extend the length of this paper, so we instead just outline the method here.

We could simply take the tester \mathbf{c} produced by our optimization and use it with exactly k samples, and its error probability would increase by at most a factor inversely proportional to the probability that $\text{Poi}(k)$ equals k . However, this would not be expected to match any corresponding lower bounds of the flavor of Theorem 4.

Instead, we adapt Equation 1 while preserving as much of its meaning as possible. Each of the two components of the outer max of Equation 1—representing type 1 and type 2 error—represents

a Chernoff bound on the failure probability, and thus the exponential of either of these can be considered as an expectation over the process of taking $\text{Poi}(k)$ samples from a distribution that is either P or Q . We consider “tilting” this expectation by multiplying the terms of the objective function by some exponential α^i so that the maximum contribution to either of the two terms comes from cases where exactly k samples are taken, instead of some other value in the support of $\text{Poi}(k)$. In this sense, we essentially add another constraint to the optimization, that we are looking for a solution “where $\text{Poi}(k)$ mimics the deterministic value k ”. After some simplification, it turns out that this new constraint can be equivalently reexpressed as a new constraint on the inner max, stipulating that the “worst case Q ” that we construct (described via histogram variable $h_x^{Q|y_j}$) must have the same total probability mass (namely, 1) as P . For technical reasons, such a constraint must be left *out* of the original Equation 1 in order for Theorem 4 to hold, but to produce the best tester for the deterministic k -sample regime, we add this extremely natural constraint to Equation 1. We point out that adding a constraint to this maximization can only *decrease* the overall objective value; thus we expect the $c_{i,j}$ customized to the deterministic case to outperform the Poissonized $c_{i,j}$, even when the Poissonized tester is evaluated on $\text{Poi}(k)$ samples. While we do not include the details, we expect *all* parts of Theorem 4 to extend to this “depoissonized” tester, including that, for fixed j , the coefficients $c_{i,j}$ will have the same log cosh functional form as before. The main difference will be that the additional constraint will induce a new dual variable—which we may call β . Thus when we express the overall optimization problem as a 2-level optimization, as in Equation 7, there will be a third variable β in the outermost optimization, joining α and u .

In short, the technical machinery of this paper extends naturally to “depoissonizing” the estimator of Theorem 4, with very little change or computational overhead.

1.2.5 “Chernoff bounds are tight”

The approach of this paper is predicated on the intuition that finding the algorithm with optimal Chernoff bound for the failure probability in our setting should be viewed as morally the same task as finding the optimal algorithm; this, combined with the versatility of manipulating Chernoff bounds enables the approach of our paper.

In statistics, the intuition that “Chernoff bounds are tight” is called the “large deviations principle”. We briefly comment on 3 old and extremely respected results establishing this intuition.

The first is Cramér’s theorem [Cra38], which considers the common case for Chernoff bounds where we wish to derive a tail bound on the mean of a large number of i.i.d. copies of a real-valued distribution X . Namely, for some threshold γ , we wish to derive a bound on $\Pr_{x_1, \dots, x_n \leftarrow X} [\frac{1}{n} \sum_i x_i \geq \gamma]$. The standard Chernoff bound on this (provided $\gamma \geq \mu(X)$) is $\min_{t \geq 0} \mathbb{E}_{x \leftarrow X} [e^{t(x-\gamma)}]^n$. Taking logarithms and normalizing by $\frac{1}{n}$, we have that the best Chernoff bound does not change with n . In the limit as $n \rightarrow \infty$, how good is this Chernoff bound? Cramér’s theorem actually says that the limit of $\frac{1}{n}$ times the log of the tail probability *exactly equals* the corresponding expression from the Chernoff (upper) bound. This theorem is extremely general, requiring almost no assumptions on the distribution X except that the relevant quantities in the equality claim are well defined. The cost of this generality is that the convergence rate is hard to control.

While Cramér’s theorem shows convergence in the *exponent* of the tail probability, a much more precise result is the Bahadur-Rao theorem [BR60], which actually shows multiplicative convergence in the probability itself, and not just in the exponent. The full formulation is complicated, but it essentially says that the true tail bound is not only better than the optimal Chernoff bound,

but in fact *quasilinear* in the optimal Chernoff bound. Intuitively, suppose we are trying to decide between two probabilistic processes A, B , and we have Chernoff bounds on their failure probabilities (in some sense) denoted respectively a, b ; then even in regimes where the Chernoff bounds a, b are *not* accurate enough, since both a, b are quasilinear in their true failure probabilities, we can correctly choose between A, B by choosing the better of a, b , trusting on quasilinearity to preserve the *ordering* of the failure probabilities, even when “distorted” by Chernoff bounds. In short, “because Chernoff bounds are quasilinear in the true tail bounds, *optimizing* over Chernoff bounds is as good as optimizing over the true tail probabilities”.

Finally, there is a long history of generalizing Cramér’s theorem to settings well beyond means of i.i.d. variables. One of the most natural and general variants is the Gärtner-Ellis theorem—see Section 5. While Cramér’s theorem talks about the convergence of tail bounds on a random variable Z_n equal to the mean of n i.i.d. copies of a distribution, Gärtner-Ellis by contrast says that, for almost any sequence of random variables Z_n , we can “pretend, in the style of Cramér’s theorem, that Z_n is the mean of n i.i.d. copies of a distribution”, and so long as the moment generating function of Z_n converges to something differentiable as $n \rightarrow \infty$, we can conclude convergence analogously to Cramér’s theorem.

In short, we leverage a long history of tools and intuition that “Chernoff bounds are tight”, and we hope this paper exposes the benefits of these intuitions to a wider audience.

1.3 Related work

We direct the reader to several notable surveys about distribution testing [Can20] and property testing [Ron08, Gol17], along with the paper [BFF⁺01] that initialized much of the work in this area.

A historically and practically important special case of our problem of hypothesis testing on distributions is the problem of *uniformity testing*, where the hypothesis is a uniform distribution. There is a long line of work on this topic, including relatively recent papers [Pan08, DGPP18, GP22]. We highlight this last paper which gets $1 + o(1)$ -tight bounds on the sample complexity of uniformity testing in various natural asymptotic regimes. This paper suggests modifying the chi-squared statistic into the Huber statistic, which is a function combining quadratic and linear regions. This paper demonstrates the value of looking at a richer class of functional forms for testers, and emphasizes the importance of $1 + o(1)$ -tight analysis when proposing new statistical testers.

The paper [VV17] focuses on a very similar setting to the one of the current paper, of hypothesis testing on distributions.. Its main feature is “instance-optimal” results, that, up to some constant factor, perform as well as is possible for the particular hypothesis distribution P . This work revealed that the difficulty of testing P depends on $\|P\|_{\frac{2}{3}}$, the $\frac{2}{3}$ norm of the distribution, an interesting structural insight. The instance-optimal analysis, however, is at the expense of some loss of constant factors, and it is unclear how well these algorithms would perform in practice.

Some classic results in the area, including parts of [VV17], were extended to the “tolerant testing” regime in [ADK15]. A tolerant tester is not only required to accept distributions that satisfy the given property, but also tolerate some corruption/error without rejecting. The paper [ADK15] showed several tolerant testing constructions based around a tester that rejects distributions with ℓ_1 distance $\geq \varepsilon$, while accepting distributions with chi-squared distance $\leq \varepsilon'$. This result shows how, in some cases, tolerant testing can be achieved without paying a high cost in the other parameters of the problem.

[Gol20] shows a general relation between hypothesis testing on distribution and the simpler uniformity testing problem, with a general constant-factor tight reduction.

The paper [DGK⁺21] encompasses several previously studied regimes of these problems, examining how the sample complexity of testing the independence and closeness of distributions depends on the desired failure probability δ . Many previous works took δ to be a constant—implicitly dealing with other δ with the black-box amplification process of taking a majority vote over independent repetitions of the tester on fresh data; by contrast, this paper examines when it is possible to combine the data in more subtle ways to get better performance as a function of δ .

Many variants of these problems exist. In the “closeness testing” problem we are given sample access to *two* distributions, P, Q (instead of being given P explicitly) and the aim is to distinguish when $P = Q$ from when $\|P - Q\|_1 \geq \varepsilon$ [BFR⁺00, ADJ⁺11, BFR⁺13, CDVV14].

Generalizing in a different direction, the Generalized Uniformity Testing model seeks to test uniformity without knowing the support of the distribution in advance [BC17, DKS18].

There are also several works [CFGM16, CRS15, ACK18] on distribution testing on conditional samples. In this model, the testing algorithm has a kind of “query access” to the distribution, in that the tester can name a subset as an input, and then receives samples *conditioned* on lying in the named subset. By choosing the subset properly, they can obtain much stronger results, achieving constant sample complexity rather than $O(\sqrt{n})$.

There are also many works [DKN15, DDS⁺, CDSS14] focusing on the identity testing problem for *structured distributions*, where the distribution is guaranteed to have certain global structural properties, like being monotonic, or k -modal. In many cases, this structure can be leveraged to yield dramatically (and sometimes exponentially) better performance.

Recent attention on quantum computing has led to several papers posing variants of the property testing problem under the quantum setting [OW15, BOW19, BCL20]. The goal is to seek algorithms that can distinguish whether the mixed state has some given property or is ϵ -far away from any mixed state possessing the property, with some constant confidence. Analogously to the classical case, the optimality of the algorithm is measured by “copy complexity”, asking how many copies of the same mixed state are required to conduct the test. This problem is closely related to identity testing or property testing problem in the classical setting, as every mixed state can be seen as a probability distribution on the support of pure states, and by taking a measurement on one mixed state, we can get a sample from that distribution; this builds a direct relationship between the copy complexity and the classical sample complexity.

2 Preliminaries

2.1 Notation

Globally, we use P to denote our hypothesis distribution, and we use j to index the *unique* probabilities of P —so that we can collectively deal with all identical domain elements at once; let y_j denote the probability of each of the domain elements indexed by j , and let $h_{y_j}^P$ denote the number of such domain elements. While each element of P that has probability y_j is symmetric, and should be treated symmetrically by our tester, different j should be treated differently. Thus when we consider an alternative probability distribution Q , we separately represent a histogram for each equivalence class j of our hypothesis distribution P . Namely, let $h_x^{Q|y_j}$ be the histogram of Q on the domain elements that have probability exactly y_j on P . When describing a Chernoff bound,

we often use the notation “ $\stackrel{\text{Ch}}{\leq}$ ” to emphasize that the inequality comes from a Chernoff bound.

2.2 The optimization problem

In this section we describe and derive the overall optimization problem of finding the semilinear testing coefficients \mathbf{c} that have the best possible Chernoff bounds on their performance.

Consider taking $\text{Poi}(k)$ samples from some distribution Q , and then, finding the best Chernoff bound on the probability that the test will mistakenly say “yes”. We let $F_{i,j}$ denote the *number* of different domain elements within the equivalence class j that have been seen exactly i times in the sample (F for “fingerprint”, as in [BFR⁺00]). Thus the semilinear tester will compute the quantity $\sum_{i,j} c_{i,j} F_{i,j}$ and return “yes” if this is < 0 . Thus by standard Chernoff bounds we have

$$\begin{aligned} \Pr_{F \leftarrow Q^{\text{Poi}(k)}} \left[\sum_{i,j} c_{i,j} F_{i,j} < 0 \right] &\stackrel{\text{Ch}}{\leq} \inf_{t \leq 0} \prod_{j,x} \left(\sum_i e^{tc_{i,j}} \text{poi}(kx, i) \right)^{h_x^{Q|y_j}} \\ &= \inf_{t \leq 0} \exp \left(\sum_{j,x} h_x^{Q|y_j} \log \sum_i e^{tc_{i,j}} \text{poi}(kx, i) \right), \end{aligned}$$

and conversely, when we draw samples from the hypothesis distribution P , we compute the best Chernoff bound on the probability that the tester mistakenly computes a statistic ≥ 0 and thus outputs “no”:

$$\begin{aligned} \Pr_{F \leftarrow P^{\text{Poi}(k)}} \left[\sum_{i,j} c_{i,j} F_{i,j} \geq 0 \right] &\stackrel{\text{Ch}}{\leq} \inf_{t' \geq 0} \prod_j \left(\sum_i e^{t'c_{i,j}} \text{poi}(ky_j, i) \right)^{h_{y_j}^P} \\ &= \inf_{t' \geq 0} \exp \left(\sum_j h_{y_j}^P \log \sum_i e^{t'c_{i,j}} \text{poi}(ky_j, i) \right). \end{aligned}$$

Therefore, the following is an upper bound of the log of the error probability δ obtained by semilinear testers:

$$\min_{\mathbf{c}} \max_{Q: |Q-P|_1 \geq \varepsilon} \log \text{err}(Q, \mathbf{c}) \stackrel{\text{Ch}}{\leq} \min_{\mathbf{c}} \max_{Q: |Q-P|_1 \geq \varepsilon} \max \left\{ \inf_{t \leq 0} \sum_{j,x} h_x^{Q|y_j} \log \sum_i e^{tc_{i,j}} \text{poi}(kx, i), \inf_{t' \geq 0} \sum_j h_{y_j}^P \log \sum_i e^{t'c_{i,j}} \text{poi}(ky_j, i) \right\}$$

The constraint that Q is a distribution such that $|Q - P|_1 \geq \varepsilon$ is informally stated, and must be rephrased in terms of the explicit optimization variables $h_x^{Q|y_j}$. We have 3 constraints on the histogram $h_x^{Q|y_j}$: the $\geq \varepsilon$ distance constraint, the fact that for each j , the total number of domain elements described by $h_x^{Q|y_j}$ for that j equals the corresponding quantity $h_{y_j}^P$, and the constraint that $h_x^{Q|y_j} \geq 0$; for the sake of efficient optimization, we must *relax* the problem by omitting any *integrality* constraint. Namely, even though $h_x^{Q|y_j}$ represents the *number* of domain

elements satisfying certain conditions, we do not enforce the restriction that it must have an integer value. (The question of whether this relaxation is “safe” is resolved, eventually, by the matching upper bound of Theorem 4.) We thus have our main expression upper-bounding the overall failure probability of our testing problem:

$$\min_{\mathbf{c}} \max \left\{ \begin{array}{l} \max_{\substack{h_x^{Q|y_j}: \sum_j \sum_x h_x^{Q|y_j} |x-y_j| \geq \varepsilon \\ \forall j, \sum_x h_x^{Q|y_j} = h_{y_j}^P \\ \forall j, \forall x, h_x^{Q|y_j} \geq 0}} \min_{t \leq 0} \sum_{x,j} h_x^{Q|y_j} \log \sum_i e^{tc_{i,j}} \text{poi}(kx, i), \min_{t' \geq 0} \sum_j h_{y_j}^P \log \sum_i e^{t'c_{i,j}} \text{poi}(ky_j, i) \end{array} \right\} \quad (1)$$

3 Upper bound

As a reminder, Equation 1 describes the best error bound of any semilinear tester, and we have denoted this quantity as $\Delta(P, k, \varepsilon)$. In this section, we analyze the structure of the coefficients $c_{i,j}$ corresponding to this best error bound, as well as showing other desirable properties of the optimal variable values of Equation 1. Later, in Section 4, we show how to use the optimum value of $h_x^{Q|y_j}$ to construct a distribution-over-distributions that will form the basis of the lower bound analysis, thus yielding the main theorem.

For the sake of applying Sion’s minimax theorem soon, we first *lower bound* Equation 1, by restricting the domain of the inner max to “ $= \varepsilon$ ” instead of “ $\geq \varepsilon$ ”. Later on, with Lemma 9, we actually show that this lower bound is tight, so we incur no loss here.

$$\min_{\mathbf{c}} \max \left\{ \begin{array}{l} \max_{\substack{h_x^{Q|y_j}: \sum_j \sum_x h_x^{Q|y_j} |x-y_j| = \varepsilon \\ \forall j, \sum_x h_x^{Q|y_j} = h_{y_j}^P \\ \forall j, \forall x, h_x^{Q|y_j} \geq 0}} \min_{t \leq 0} \sum_{x,j} h_x^{Q|y_j} \log \sum_i e^{tc_{i,j}} \text{poi}(kx, i), \min_{t' \geq 0} \sum_j h_{y_j}^P \log \sum_i e^{t'c_{i,j}} \text{poi}(ky_j, i) \end{array} \right\} \quad (2)$$

We aim to simplify this optimization problem by reducing the number of nested layers of min and max; and towards this end, we aim to swap the order of the max and the min of the first term. We will utilize Sion’s minimax theorem to achieve this. We first state the theorem for clarity.

Lemma 5 (Sion’s minimax theorem; [Sio58]). *Let X and Y be convex spaces, one of which is compact. Let $f : X \times Y \rightarrow \mathbb{R}$ be a function such that for all $x \in X, y \in Y$,*

- $f(x, \cdot)$ is upper semi-continuous and quasi-concave on Y ,
- $f(\cdot, y)$ is lower semi-continuous and quasi-convex on X ,

then

$$\sup_y \inf_x f(x, y) = \inf_x \sup_y f(x, y).$$

The function being optimized is a convex function of t because $e^{tc_{i,j}} \text{poi}(kx, i)$ is log-convex, and sums of log-convex functions are log-convex, and thus its logarithm is convex. The function being optimized is linear in the maximization variable $h_x^{Q|y_j}$. The domains of optimization are all convex as they are defined by linear constraints. We use limits to argue that the continuous domain of x can be discretized (in the limit) without changing the optimum. In the following lemma, we will show the compactness of the space of $h_x^{Q|y_j}$.

Lemma 6. *Focusing on some particular j , and considering histogram locations x that are multiples of some fixed spacing α , we claim that the set of histograms satisfying the relaxed constraints $\sum_x h_x |x - y_j| \leq \varepsilon$, and $\sum_x h_x = h_{y_j}^P$ and $\forall x, h_x \geq 0$ is compact when considered as a subset of the set of sequences, in the ℓ_1 norm.*

Proof. We reparameterize h_x to index by nonnegative integers i , so that h_i denotes what we previously referred to as $h_{\alpha i}$.

The constraints on h , reexpressed, now read: $\sum_{i=0}^{\infty} h_i |\alpha i - y_j| \leq \varepsilon$, and $\sum_{i=0}^{\infty} h_i = h_{y_j}^P$ and $\forall i \geq 0, h_i \geq 0$. We show that this is a compact set by equivalently showing that this set is closed, bounded, and equismall at infinity [Tre16]. The set is clearly closed; it is bounded since we explicitly have that the sum of the entries of h equals the fixed parameter $h_{y_j}^P$. Equismall at infinity means that for every $\beta > 0$, there exists an integer i_β such that $\sum_{i=i_\beta}^{\infty} |h_i| \leq \beta$ for all h . Explicitly, we take i_β such that $\alpha i_\beta - y_j \geq \frac{\varepsilon}{\beta}$. Thus for all $i \geq i_\beta$, we have that $|\alpha i - y_j| \geq \frac{\varepsilon}{\beta}$. Thus suppose, for the sake of contradiction, that $\sum_{i=i_\beta}^{\infty} |h_i| > \beta$; multiplying these last two inequalities (and using $h_i \geq 0$) yields that $\sum_{i=i_\beta}^{\infty} h_i |\alpha i - y_j| > \beta \frac{\varepsilon}{\beta} = \varepsilon$, contradicting the first constraint. Thus these constraints describe a compact set in the ℓ_1 topology. \square

Thus, we can apply Lemma 6 to each j separately, and conclude that the constraints on the histogram in Equation 2 describe a set that lies in the direct product (across the finite set of j) of compact sets, and thus is itself compact. Therefore, we can apply Sion's minimax theorem (Lemma 5).

Thus Equation 2 equals

$$\min_{\mathbf{c}} \max \left\{ \min_{t \leq 0} \max_{\substack{h^{Q|y_j}: \sum_j \sum_x h_x^{Q|y_j} |x - y_j| = \varepsilon \\ \forall j, \sum_x h_x^{Q|y_j} = h_{y_j}^P \\ \forall j, \forall x, h_x^{Q|y_j} \geq 0}} \sum_{x,j} h_x^{Q|y_j} \log \sum_i e^{tc_{i,j}} \text{poi}(kx, i), \min_{t' \geq 0} \sum_j h_{y_j}^P \log \sum_i e^{t'c_{i,j}} \text{poi}(ky_j, i) \right\} \quad (3)$$

We can pull both min's outside of the brackets, and also outside of the max, and hence merge them with the outer min to yield:

$$\min_{\mathbf{c}, t \leq 0, t' \geq 0} \max \left\{ \begin{array}{l} \max_{\substack{h^{q|y_j}: \sum_j \sum_x h_x^{Q|y_j} |x-y_j| = \varepsilon \\ \forall j, \sum_x h_x^{Q|y_j} = h_{y_j}^P \\ \forall j, \forall x, h_x^{Q|y_j} \geq 0}} \sum_{x,j} h_x^{Q|y_j} \log \sum_i e^{tc_{i,j}} \text{poi}(kx, i), \sum_j h_{y_j}^P \log \sum_i e^{t'c_{i,j}} \text{poi}(ky_j, i) \end{array} \right\} \quad (4)$$

The inner max is now a linear program over $h^{q|y_j}$, so we take its dual. Let α be the (scalar) dual variable for the first constraint; let γ be the dual variable for the second constraint, a vector with an entry for each j .

$$\min_{\mathbf{c}, t \leq 0, t' \geq 0} \max \left\{ \begin{array}{l} \min_{\alpha, \gamma:} \varepsilon \alpha + \sum_j \gamma_j h_{y_j}^P, \sum_j h_{y_j}^P \log \sum_i e^{t'c_{i,j}} \text{poi}(ky_j, i) \\ \forall j, \forall x \geq 0, \alpha |x-y_j| + \gamma_j \geq \log \sum_i e^{tc_{i,j}} \text{poi}(kx, i) \end{array} \right\} \quad (5)$$

We pull the min outside of the brackets, outside of the max, and merge it with the outer min:

$$\min_{\substack{\mathbf{c}, t \leq 0, t' \geq 0, \alpha, \gamma: \\ \forall j, \forall x \geq 0, \alpha |x-y_j| + \gamma_j \geq \log \sum_i e^{tc_{i,j}} \text{poi}(kx, i)}} \max \left\{ \varepsilon \alpha + \sum_j \gamma_j h_{y_j}^P, \sum_j h_{y_j}^P \log \sum_i e^{t'c_{i,j}} \text{poi}(ky_j, i) \right\} \quad (6)$$

Summarizing the above, we have:

Lemma 7. *Equation 2 equals Equation 6.*

Next, to understand how Equation 2 relates to Equation 1, we need to understand the role of the dual and the dual variable α .

Lemma 8. *The optimum of Equation 6 is ≤ 0 . If the optimum is 0, then there is an optimal solution with $\alpha = 0$; and if the optimum of Equation 6 is < 0 then $\alpha < 0$.*

Proof. If the value of Equation 6 equals 0, then, consider setting $\alpha = \gamma_j = t = t' = 0$, in which case all the exponentials will equal 1, and, since $\sum_i \text{poi}(kx, i) = 1$ for any x , both logarithm expression equal 0, and the constraints are satisfied and the overall expression has value 0. Thus we have constructed an optimal solution satisfying the first claims of the lemma. Further, the optimal value can never be > 0 because the solution above gives an objective value of 0, no matter what the the inputs $\varepsilon, h_{y_j}^P, y$ are.

For the last part of the lemma, suppose for the sake of contradiction $\alpha \geq 0$. Consider setting x from the constraints below the min to equal y_j .

Then, the first term in the max is at least

$$\sum_j \gamma_j h_{y_j}^P \geq \sum_j h_{y_j}^P \log \sum_i e^{tc_{i,j}} \text{poi}(ky_j, i)$$

The right hand side is convex as a function of t (since $e^{tc_{i,j}}$ is log-linear and thus log-convex, and the sum of log-convex functions is log-convex; hence its logarithm is convex, and the entire expression is a sum of convex functions and hence convex). Further, its value at $t = 0$ is 0. Thus, by convexity, its values at $t \leq 0$ and $t' \geq 0$ cannot *both* be negative. Therefore the maximum of these is ≥ 0 , contradicting our assumption that the objective is negative. \square

Lemma 9. *Equation 1 equals Equation 2.*

Proof. Recall that α in Equation 6 is the dual variable corresponding to the constraint of Equation 3 that the ℓ_1 distance from the hypothesis equals ε . Thus the fact that Equation 6 has an optimal solution with $\alpha \leq 0$, from Lemma 8, means that the optimum of Equation 3 would be unchanged if we changed this $= \varepsilon$ constraint to $\geq \varepsilon$. Equation 2 with a $\geq \varepsilon$ constraint is at most Equation 3 with a $\geq \varepsilon$ constraint (trivially, from the “weak” direction of minimax), which from the above equals Equations 2 and 3 as written. Conversely, Equation 2 with a $\geq \varepsilon$ bound is at least Equation 2 as written, since enlarging the domain of a maximum can only increase its value. Thus Equation 2 must equal Equation 2 with a $\geq \varepsilon$ bound, namely Equation 1. \square

Combining Lemmas 7 and 9 yields that Equation 1 equals Equation 6.

Looking closely at the inner max in Equation 6, we note that the first term is the dual of the type II error probability i.e, the probability of the tester accepting a sample from Q , while the second term is the type I error probability, i.e. the probability of the tester rejecting a sample from P . Intuitively, we expect both error probabilities to be equal in the optimum. We formalize this intuition with the following lemma.

Lemma 10. *There exists an optimum for Equation 6 such that the two terms in the max are equal.*

Proof. Letting $\alpha = \gamma_j = t = t' = 0$ yields a feasible point for Equation 6 with both terms in the max equal to 0.

The only remaining case is when the value of the optimum is < 0 . Observe that in this scenario, $t' > 0$ (otherwise, $t = 0$ and thus the second term of the max is equal to 0, and hence the overall value of the objective function is at least 0). Consider an optimal solution \mathbf{c}, γ . We separately consider the cases where the first term of the max is bigger, or smaller.

Suppose the first term is bigger. Then since $t' > 0$, find a nonzero hypothesis value $h_{y_j}^P$, and increase the corresponding $c_{i,j}$ until the second term of the max equals the first. The variable $c_{i,j}$ only occurs in one other place, namely the constraint of the min. Because $t \leq 0$, increasing $c_{i,j}$ will decrease the right-hand side of the inequality constraint, and thus will not violate it. Hence we have constructed a feasible solution with the same objective value but which now satisfies the claim of the lemma.

For the other case, suppose the first term is smaller. We increase γ until the first term equals the second term. The only other places γ appears are on the left-hand side of the inequality constraints below the min; increasing α can only increase this left-hand side. Hence we have constructed a feasible solution with the same objective value but which now satisfies the claim of the lemma. \square

We next split Equation 6 into a component for each j , expressing the overall optimization as a “nested” optimization, choosing α, t, t' on the outside, and coefficients $\kappa_{i,j}$ inside. Here we use $\kappa_{i,j}$ in place of $c_{i,j}$ and u to replace t and t' . The following lemma will show that this change does not have any effect on the results of our optimization problem. And then the further lemmas will rely on this nested optimization to derive clean structural properties of the optimum.

Lemma 11. *Equation 6 equals*

$$\min_{\substack{\alpha \leq 0 \\ u \in [0,1]}} \left(\varepsilon \alpha (1-u) + \sum_j \min_{\kappa_{i,j}} \left(u h_{y_j}^P \left(\log \sum_i e^{(1-u)\kappa_{i,j}} \text{poi}(ky_j, i) \right) + (1-u) h_{y_j}^P \max_{x \geq 0} \left(-\alpha |x - y_j| + \log \sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i) \right) \right) \right) \quad (7)$$

Proof. Given any optimal point of Equation 6, specified by $c_{i,j}, t \leq 0, t' \geq 0, \alpha, \gamma_j$, we consider it in the context of Equation 7 (noting, by Lemma 8 that $\alpha \leq 0$ at optimum), and setting $\kappa_{i,j} = (t' - t)c_{i,j}$ and $u = -\frac{t}{t' - t}$ (with u set arbitrarily if $t = t' = 0$). Substituting in the relation $t c_{i,j} = -u \kappa_{i,j}$ into the constraint in Equation 6, we have

$$\gamma_j \geq \max_{x \geq 0} \left(-\alpha |x - y_j| + \log \sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i) \right) \quad (8)$$

And thus, since $(1-u)\kappa_{i,j} = t'c_{i,j}$, the convex combination of the two terms in the max of the objective of Equation 6, with weights $1-u, u$, after substituting Equation 8 for each γ_j in the first term, is greater than or equal to the objective function of Equation 7. Thus by the triangle inequality, this max in Equation 6—which is its objective function value at optimum—is at least the feasible value we have found for Equation 7. Thus Equation 7 is less than or equal to Equation 6.

Conversely, given any feasible point of Equation 7, we construct the corresponding feasible point of Equation 6, by setting $t = -u, t' = 1 - u$ (so that $t' - t = 1$), $c_{i,j} = \kappa_{i,j} + s$ for a shift s to be determined later, and setting

$$\gamma_j = \max_{x \geq 0} \left(-\alpha |x - y_j| + \log \sum_i e^{t c_{i,j}} \text{poi}(kx, i) \right)$$

We note that γ_j depends linearly on s with slope t , and thus the first term of the max in the objective of Equation 6 depends linearly on s with slope t (since $\sum_j h_{y_j}^P = 1$). Similarly, the second term in the max in the objective of Equation 6 depends linearly on s with slope t' . Thus we pick s which makes these two terms equal.

Thus the objective value of Equation 6 here equals both of the two terms of its max (because we just chose s to make these terms equal); and in particular equals the convex combination of them with weights $1-u = t', u = -t$, which is thus exactly the objective value of Equation 7, since the contributions of s (whatever they are) exactly cancel with the weights. Thus Equation 6 is less than or equal to Equation 7.

Combining both parts yields the desired equality. \square

We now show the uniqueness of the optimal solutions for $\kappa_{i,j}$ in Equation 7, up to an irrelevant j -dependent additive shift.

Lemma 12. *The inner minimization in Equation 7, if $u \in (0, 1)$, has a unique solution for $\kappa_{i,j}$, for each j , up to an additive shift.*

Proof. Consider two different optima, $\kappa_{i,j}$ and $\kappa'_{i,j}$. Since Equation 7 is convex in κ , and in particular, is the sum of convex terms, we have that each term must in particular be linear at any convex combination of $\kappa_{i,j}$ and $\kappa'_{i,j}$. In particular, $\log \sum_i e^{(1-u)\kappa_{i,j}} \text{poi}(ky_j, i)$ must be linear, meaning that

$\sum_i e^{(1-u)\kappa_{i,j}} \text{poi}(ky_j, i)$ must be exponential as we interpolate between $\kappa_{i,j}$ and $\kappa'_{i,j}$. But the only way for the sum of several exponentials to equal an exponential is if all the exponentials have the same base; thus $\kappa_{i,j} = \kappa'_{i,j} + s$ for some additive shift, as claimed. \square

This shift-invariant structure of $\kappa_{i,j}$'s allows us an extra degree of freedom, as we can choose the shifting factor to further simplify the equation. In the following lemma, we reexpress the inner minimization in Equation 7 by taking its dual and then simplifying while utilizing the shift-invariant property; this also allows us to find the crucial and insightful closed form for $\kappa_{i,j}$.

Lemma 13. *The inner minimization in Equation 7, if $u \in (0, 1)$, can be expressed for any j as*

$$\max_{a_x \geq 0} h_{y_j}^P \log \frac{\sum_i (\sum_x a_x \text{poi}(kx, i))^{1-u} \text{poi}(ky_j, i)^u}{(\sum_x a_x e^{\alpha|x-y_j|})^{1-u}} \quad (9)$$

where the optimal variable values of Equation 7 can be expressed in terms of the optimal variables of Equation 9 as

$$\kappa_{i,j} = \log \sum_x a_x \frac{\text{poi}(kx, i)}{\text{poi}(ky_j, i)} \quad (10)$$

Proof. We first point out that because Equation 7 is invariant to shifts in $\kappa_{i,j}$ for each j , there is some shift of the optimal $\kappa_{i,j}$ that makes the inner max equal to 0. Thus the optimum is unchanged if we restrict the inner maximum to have value ≤ 0 and remove this term from the objective. Exponentiating both the objective and the new constraints—and for now omitting the scaling factor $uh_{y_j}^P$ from the objective function, since it merely scales the optimum—yields

$$\begin{aligned} \min_{\kappa_{i,j}} & \left(\sum_i e^{(1-u)\kappa_{i,j}} \text{poi}(ky_j, i) \right) \\ \text{such that } \forall x, & 1 \geq e^{-\alpha|x-y_j|} \left(\sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i) \right) \end{aligned} \quad (11)$$

The dual of this convex optimization is then:

$$\max_{\lambda_x \geq 0} \min_{\kappa_{i,j}} \left(\sum_i e^{(1-u)\kappa_{i,j}} \text{poi}(ky_j, i) \right) + \sum_x \lambda_x \left(-1 + e^{-\alpha|x-y_j|} \left(\sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i) \right) \right) \quad (12)$$

To solve the inner minimization, we compute the derivative with respect to a single $\kappa_{i,j}$:

$$\begin{aligned} & (1-u)e^{(1-u)\kappa_{i,j}} \text{poi}(ky_j, i) + \sum_x \lambda_x e^{-\alpha|x-y_j|} (-u)e^{-u\kappa_{i,j}} \text{poi}(kx, i) \\ & = e^{(1-u)\kappa_{i,j}} (1-u) \text{poi}(ky_j, i) + e^{-u\kappa_{i,j}} \sum_x \lambda_x e^{-\alpha|x-y_j|} (-u) \text{poi}(kx, i) \end{aligned}$$

Setting this to 0 yields $\kappa_{i,j} = \log \frac{u}{1-u} \sum_x \lambda_x e^{-\alpha|x-y_j|} \frac{\text{poi}(kx, i)}{\text{poi}(ky_j, i)}$

Plugging this into Equation 12 (to evaluate the dual problem), the objective function is:

$$\left(\sum_i \left(\frac{u}{1-u} \sum_x \lambda_x e^{-\alpha|x-y_j|} \frac{\text{poi}(kx, i)}{\text{poi}(ky_j, i)} \right)^{1-u} \text{poi}(ky_j, i) \right) + \sum_x \lambda_x \left(-1 + e^{-\alpha|x-y_j|} \left(\sum_i \left(\frac{u}{1-u} \sum_x \lambda_x e^{-\alpha|x-y_j|} \frac{\text{poi}(kx, i)}{\text{poi}(ky_j, i)} \right)^{-u} \text{poi}(kx, i) \right) \right)$$

which simplifies to

$$\sum_i \frac{1}{u} \left(\frac{u}{1-u} \sum_x \lambda_x e^{-\alpha|x-y_j|} \text{poi}(kx, i) \right)^{1-u} \text{poi}(ky_j, i)^u - \sum_x \lambda_x$$

Since this expression must be maximized in terms of λ_x , it must in particular be maximized with respect to scalings of λ_x . Replacing λ_x with $e^z \lambda_x$ for some parameter z , we calculate that the maximum over z of this expression equals

$$\left(\frac{\sum_i (\sum_x \lambda_x e^{-\alpha|x-y_j|} \text{poi}(kx, i))^{1-u} \text{poi}(ky_j, i)^u}{(\sum_x \lambda_x)^{1-u}} \right)^{\frac{1}{u}}$$

For convenience, we reexpress λ_x as $a_x e^{\alpha|x-y_j|}$, where for each x , we have a_x is nonzero if and only if λ_x is nonzero. Taking the logarithm of this, and multiplying back by the factor $u h_{y_j}^P$ that we dropped earlier yields that the inner minimization of the j^{th} term in Equation 7 equals the maximization

$$\max_{a_x \geq 0} h_{y_j}^P \log \frac{\sum_i (\sum_x a_x \text{poi}(kx, i))^{1-u} \text{poi}(ky_j, i)^u}{(\sum_x a_x e^{\alpha|x-y_j|})^{1-u}}$$

as claimed. Finally, we point out that, defining $\kappa_{i,j}$ in terms of a_x yields $\kappa_{i,j} = \log \frac{u}{1-u} \sum_x a_x \frac{\text{poi}(kx, i)}{\text{poi}(ky_j, i)}$; since $\kappa_{i,j}$ in Equation 7 is shift-invariant, and a_x is scale-invariant (since we have rescaled by the optimal scaling factor e^z), we can drop the $\frac{u}{1-u}$ factor without affecting anything, yielding Equation 10. \square

The above expression for the optimal values $\kappa_{i,j}$ can be viewed as representing $\kappa_{i,j}$, for any fixed j , as the sum of exponentials in i :

$$\kappa_{i,j} = \log \sum_x a_x \frac{\text{poi}(kx, i)}{\text{poi}(ky_j, i)} = \log \sum_x a_x e^{k(y_j-x)} (x/y_j)^i \quad (13)$$

We will now show a crucial property of the hardest distribution Q with respect to $\kappa_{i,j}$: each distinct probability mass y_j in P is replaced by *exactly* two probability masses, which we denote as $x_{1,j}$ and $x_{2,j}$. This property is known to be the worst case for ε -far uniformity testing for both the TV and collision tester [DGPP18]. In the context of Equation 10, for any fixed j , we would expect that exactly two a_x 's are non-zero, which is formally shown below in Lemma 16. En route to this lemma, we will show an additional property of κ as well as a technical lemma to aid with the final analysis of the two-point structure.

Lemma 14. *The optimal value of $\kappa_{i,j}$ given by Equation 10 is convex as a function of i (where Equation 10 applies, by Lemma 13, if $u \in (0, 1)$).*

Proof. We reexpress Equation 10 as Equation 13, and point out that exponential functions are log-convex, and sums of log-convex functions are log-convex, yielding that $\kappa_{i,j}$ is a convex function of i . \square

Lemma 15. *Suppose $g : \mathbb{Z}^+ \rightarrow \mathbb{R}$ is log-concave. Then $\sum_{k=0}^{\infty} g(k) \text{poi}(\lambda, k)$ is a log-concave function of λ for $\lambda \geq 0$; and if g is strictly log-concave somewhere, then this function is strictly log-concave for all $\lambda > 0$.*

Proof. We instead work with the function $f(\lambda) = e^\lambda \sum_{k=0}^{\infty} g(k) \text{poi}(\lambda, k) = \sum_{k=0}^{\infty} g(k) \lambda^k \frac{1}{k!}$, which differs from the desired function only in the e^λ factor, and thus is log-concave exactly when the original function is.

We have $f'(\lambda) = \sum_{k=0}^{\infty} g(k+1) \lambda^k \frac{1}{k!}$ and $f''(\lambda) = \sum_{k=0}^{\infty} g(k+2) \lambda^k \frac{1}{k!}$. We note that we can extend the sum down to, say, $k = -1$ under the convention that $(-1)! = \infty$.

The condition for log-concavity is: $f'(\lambda)^2 - f''(\lambda)f(\lambda) \geq 0$. Writing this out as a double sum:

$$f'(\lambda)^2 - f''(\lambda)f(\lambda) = \sum_{j,k=0}^{\infty} (g(j+1)g(k+1) - g(j+2)g(k)) \lambda^{j+k} \frac{1}{j!k!}$$

Replacing j by $j-1$ yields the identical sum

$$\sum_{j=1}^{\infty} \sum_{k=0}^{\infty} (g(j)g(k+1) - g(j+1)g(k)) \lambda^{j+k-1} \frac{1}{(j-1)!k!}$$

Adding this sum to the corresponding sum with j, k swapped (where we symmetrize the domain of the sum to the set of integers $j, k \geq 0$ where $j+k \geq 1$, which does not change the summation):

$$\sum_{\substack{j,k \geq 0 \\ j+k \geq 1}} \lambda^{j+k-1} (g(j)g(k+1) - g(j+1)g(k)) \left(\frac{1}{(j-1)!k!} - \frac{1}{(k-1)!j!} \right)$$

We now point out that this sum is nonnegative term-by-term: the expression $g(j)g(k+1) - g(j+1)g(k)$ is nonnegative for log-concave g when $j \geq k$, and symmetrically, nonpositive when $j \leq k$; the term $\left(\frac{1}{(j-1)!k!} - \frac{1}{(k-1)!j!} \right)$ is also nonnegative when $j \geq k$ and nonpositive when $j \leq k$; hence their product is always nonnegative, meaning the overall log-concavity condition is the sum of nonnegative terms, and is thus nonnegative, as desired.

To show strict log-concavity, we point out that the term $\left(\frac{1}{(j-1)!k!} - \frac{1}{(k-1)!j!} \right)$ is strictly positive whenever $j > k$; and thus if there is a location $\ell \geq 1$ (namely in the interior of \mathbb{Z}^+) such that $g(\ell+1)^2 > g(\ell)g(\ell+2)$ then, letting $j = \ell$ and $k = \ell - 1$ yields $g(j)g(k+1) - g(j+1)g(k) > 0$ and hence the overall log-concavity expression is strictly positive, multiplied by some power of $\lambda > 0$. \square

We are now ready to prove the two-point structure of q .

Lemma 16. *For any fixed $\alpha < 0, u \in (0, 1)$, and fixed j , the inner minimization of Equation 7 will have a solution $\kappa_{i,j}$ that, when expressed via Equation 10, has exactly two x such that $a_x \neq 0$, which we denote as $x_{1,j} < x_{2,j}$, and which satisfy $x_{1,j} < y_j < x_{2,j}$. We normalize so that we express $a_{x_{1,j}} = q_j$ and $a_{x_{2,j}} = 1 - q_j$, for some $q_j \in (0, 1)$ —taking advantage of the fact that Equation 7 is invariant to shifting $\kappa_{i,j}$, and thus we can scale a_x arbitrarily.*

Proof. If $\kappa_{i,j}$ is exactly linear as a function of i (for our fixed j), namely $\kappa_{i,j} = a + bi$, then, expressing $\kappa_{i,j}$ in the form of Equation 13, there must be a single nonzero a_x .

Otherwise, since by Lemma 14, $\kappa_{i,j}$ is a convex function of i , then $\kappa_{i,j}$ must be *strictly* convex (since we already covered the case where it is linear). Thus by Lemma 15, we have that the expression $\log \sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i)$ in the inner max of Equation 7 is a strictly concave function of x ; and thus must intersect with each branch of the absolute value function at most once, namely at most twice in total.

Thus, considering the dual form of Equation 7 analyzed in the proof of Lemma 13, by complementary slackness the dual variable λ_x must be nonzero in at most 2 locations. And by definition of a_x , it is nonzero if and only if the corresponding λ_x is.

Thus in all cases, the maximization $\max_x -\alpha|x - y_j| + \log \sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i)$ is tight for at most 2 points, with at most one intersection point per branch of $\alpha|x - y_j|$. Thus we call $x_{1,j}$ the (possible) intersection point below y_j and call $x_{2,j}$ the (possible) intersection point above y_j .

If there are no intersections, then $a_x = 0$ everywhere, and $\kappa_{i,j} = -\infty$ and the expression $\log \sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i)$ is infinite (since $t < 0$) and thus the optimization of Equation 11 is infeasible, so this cannot happen. If there is one intersection, then, there is at most one nonzero a_x (by complementary slackness), and thus $\kappa_{i,j}$ is linear in i by Equation 13. If $\kappa_{i,j}$ is linear in i then $\sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i)$ is linear in x . Now, if a linear function is bounded by the constraint $\alpha|x - y_j| + \beta$ and intersects it once, the intersection point must be at $x = 0$. Thus the only nonzero a_x can be a_0 . This leads to $\kappa_{i,j} = \log a_0 e^{k(y_j - 0)} (0/y_j)^i$, which when $i \geq 1$ is $\log 0 = -\infty$. As above, the expression $\log \sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i)$ is thus infinite (since $u > 0$) and thus the constraint of Equation 11 is infeasible, so this cannot happen.

Thus it must be that we have two intersection points, $x_{1,j}$ and $x_{2,j}$. Namely, $\alpha|x - y_j| + \gamma_j = \sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i)$ at both $x = x_{1,j}$ and $x = x_{2,j}$. Neither $x_{1,j}$ nor $x_{2,j}$ can equal y_j since $\sum_i e^{-u\kappa_{i,j}} \text{poi}(kx, i)$ is smooth, and thus the maximum of this smooth function plus the “v-shaped” function $-\alpha|x - y_j|$ cannot occur at the corner of the “v” (since $\alpha < 0$). \square

We next show, via a relatively straightforward calculation, that we can explicitly find a negative objective function value for our optimization problem, hence implying that the global optimum value is < 0 . We will then use this to rule out certain pathological behavior that could occur when certain variables equal 0.

Lemma 17. *The optimum value of Equation 7 is < 0 , attained in the limit as $\alpha \rightarrow 0$ from the negative side, and setting $u = \frac{1}{2}$, $\tau = \left(-\alpha \frac{y_j^2}{uk^2}\right)^{\frac{1}{3}}$, $x_{1,j} = y_j - \tau$, $x = y_j + \tau$, and $\kappa_{i,j} = \log \frac{\text{poi}(kx_{1,j}, i) + \text{poi}(kx_{2,j}, i)}{2\text{poi}(ky_j, i)}$.*

The final result of this section, and of the upper-bound portion of our analysis, puts together the pieces we have shown so far to summarize the properties of the optimum of Equation 7—and hence also Equation 1. We use the structural properties of the upper bound to reexpress our optimal

objective function value, via Equation 14, in a form that will show up crucially in our lower bound, in Lemma 30.

Lemma 18. *When Equation 7 is optimized for $\alpha, u, \kappa_{i,j}$, then $\alpha < 0$, $u \in (0, 1)$, and (the optimal) $\kappa_{i,j}$ can be expressed via Lemma 16 as $\kappa_{i,j} = \log \frac{q_j \text{poi}(kx_{1,j}, i) + (1 - q_j)x_{2,j}}{\text{poi}(ky_j, i)}$, and its value equals*

$$\varepsilon\alpha(1 - u) + \sum_j h_{y_j}^P \log \frac{\sum_i (q_j \text{poi}(kx_{1,j}, i) + (1 - q_j)\text{poi}(kx_{2,j}, i))^{1-u} \text{poi}(ky_j, i)^u}{(q_j e^{\alpha|x_{1,j} - y_j|} + (1 - q_j)e^{\alpha|x_{2,j} - y_j|})^{1-u}} \quad (14)$$

which by Lemmas 7, 9, and 11 equals Equation 1. Further, for these $\alpha, u, q_j, x_{1,j}, x_{2,j}$, the derivative of Equation 14 with respect to α or u or any q_j is 0.

Proof. We first show that we can apply Lemma 16. From Lemma 17, the objective function at optimum must be negative. We now point out that, since Equation 7 equals Equation 6 (by Lemma 11), then Lemma 8 yields that $\alpha < 0$. Next, we point out that $u \neq 1$ because if $u = 1$ then all the terms in Equation 7 vanish, leaving an objective of 0; but the objective is actually negative. Finally, we point out that $u \neq 0$, since otherwise, if $u = 0$, the inner minimization of Equation 7 contains the term $\max_{x \geq 0} -\alpha|x - y_j|$ which is infinite since $\alpha < 0$, and hence contradicts the fact that the optimum is negative. Thus the conditions of Lemma 16 apply, namely $\alpha < 0$ and $u \in (0, 1)$.

We can thus apply Lemma 13 for each j , where $\kappa_{i,j}$ is reexpressed using the 2-point form of Lemma 16, and we use the other part of Lemma 13 to say that the inner minimization of Equation 7 equals the inner maximization below, to conclude that Equation 7 equals

$$\min_{\alpha < 0, u \in (0, 1)} \left(\varepsilon\alpha(1 - u) + \sum_j \max_{q_j, x_{1,j}, x_{2,j}} h_{y_j}^P \log \frac{\sum_i (q_j \text{poi}(kx_{1,j}, i) + (1 - q_j)\text{poi}(kx_{2,j}, i))^{1-u} \text{poi}(ky_j, i)^u}{(q_j e^{\alpha|x_{1,j} - y_j|} + (1 - q_j)e^{\alpha|x_{2,j} - y_j|})^{1-u}} \right) \quad (15)$$

By Lemmas 7, 9, and 11 we have that Equation 7 equals Equation 1, so thus Equation 14 equals Equation 1, as desired.

Since the function being optimized in Equation 15 is smooth, and the inner maximization has a unique solution (by Lemma 12), and the optimum is attained for α, u, q_j in the interior of their domains (for q_j this is from Lemma 16), thus the derivatives of Equation 14 with respect to any of α, u, q_j are all 0, as desired. \square

4 Lower bound

Given a hypothesis distribution specified by P , and a number of samples k , we derive a particular distribution “ Q ” that is *far* from the hypothesis P , yet *hard to distinguish* from it, and we derive our lower bounds from this. If Q is ε -far from P and no k -sample tester can distinguish P from Q with success probability $\geq 1 - \delta$, then, a fortiori, no k -sample tester can distinguish P from the entire set of distributions $\geq \varepsilon$ distance from P , with success probability $\geq 1 - \delta$.

Instead of providing a single distribution Q , we instead provide a *distribution over distributions* $\mathcal{Q}_{[\varepsilon, \varepsilon']}$, each of whose members is itself $\geq \varepsilon$ -far from P , and such that no $\text{Poi}(k)$ -sample tester can distinguish P from a random distribution from $\mathcal{Q}_{[\varepsilon, \varepsilon']}$ with expected success probability $\geq 1 - \delta$.

Explicitly, for each domain element, we will flip a (weighted) coin between 2 probabilities; we call this a “coin flip distribution”; we then condition on the ℓ_1 distance of this distribution from our hypothesis P being in the interval $[\varepsilon, \varepsilon']$, and call $\mathcal{Q}_{[\varepsilon, \varepsilon']}$ a “conditional coin flip distribution”. We choose ε' later so that it converges to ε .

In this section we will find it convenient to regard j as representing a single domain element, instead of as an equivalence class of identical domain elements; thus y_j refers to the probability of domain element j in the hypothesis; and we avoid referring to $h_{y_j}^P$.

Definition 19. Let \mathcal{Q} be the distribution-over-distributions defined according to the following process. (We note that, because we are in the Poissonized setting, the total probability mass of a distribution output by \mathcal{Q} need not be exactly 1.) From Lemma 18, we take optimal coefficients $q_j, x_{1,j}, x_{2,j}$ for each j . For each domain element, with associated j , we flip a weighted coin, and with probability q_j set this domain element to have probability $x_{1,j}$ and otherwise, have probability $x_{2,j}$.

We further define $\mathcal{Q}_{[\varepsilon, \varepsilon']}$ to be a *conditional* distribution, where we sample a distribution q from \mathcal{Q} but proceed to return q *conditioned* on whether its ℓ_1 distance from P is $\in [\varepsilon, \varepsilon']$.

For a histogram h , recording for each domain element, the number of times it was sampled, we let $\mathcal{Q}(h), \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)$ denote the probability of h appearing as this histogram from $\text{Poi}(k)$ samples from q drawn respectively from \mathcal{Q} or $\mathcal{Q}_{[\varepsilon, \varepsilon']}$. We will also sometimes consider a realization r of the coin flip process in \mathcal{Q} , where r_j denotes whether the j^{th} coin is heads or tails. Let $\mathcal{Q}(h, r)$ denote the joint probability that the coin flips realized outcome r , and then that the histogram sampled from this was h . For the sake of convenience, we define the joint probability of observing a histogram h and where the coin flip realization has ℓ_1 distance ε from the hypothesis: $\mathcal{Q}(h, [\varepsilon, \varepsilon']) = \sum_{r: \sum_j |x_{r_j, j} - y_j| \in [\varepsilon, \varepsilon']} \mathcal{Q}(h, r)$; and we define the probability that the coin flip realization has ℓ_1 distance $\in [\varepsilon, \varepsilon']$ from the hypothesis, regardless of the histogram sampling process: $\mathcal{Q}([\varepsilon, \varepsilon']) = \sum_h \mathcal{Q}(h, [\varepsilon, \varepsilon'])$. Using this notation, we can express the law of conditional probability: $\mathcal{Q}_{[\varepsilon, \varepsilon']}(h) = \frac{\mathcal{Q}(h, [\varepsilon, \varepsilon'])}{\mathcal{Q}([\varepsilon, \varepsilon'])}$. □

For the sake of completeness, we state and prove the standard Neyman-Pearson Lemma, showing that, to distinguish 2 hypotheses, the optimal tester (over *all* possible testers!) is a log-likelihood threshold tester.

Lemma 20 (Neyman-Pearson). *The tester for distinguishing P from $\mathcal{Q}_{[\varepsilon, \varepsilon']}$ such that the max of type-1 and type-2 error is minimized, is defined by a log-likelihood threshold γ , and a tie-breaking probability λ , where the tester says “no” if $\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} > \gamma$, “yes” if the log-likelihood is $< \gamma$, and flips a λ -biased coin if the log-likelihood equals γ . Equivalently, the tester says “no” according to a λ -weighted coin flip between the output of whether $\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} \geq \gamma$ and the strict version $\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} > \gamma$.*

Proof. First, we will define the log-likelihood threshold tester as follows:

$$\Gamma(h) = \begin{cases} 1 & \text{if } \log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} > \gamma \\ 1 & \text{with probability } p_\gamma \text{ if } \log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} = \gamma \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

We will then consider the best log-likelihood threshold tester

$$\arg \min_{\Gamma} \max \left\{ \Pr_{h \leftarrow P} [\Gamma(h) = 1], \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 0] \right\}$$

We claim that this minimum is achieved when $\Pr_{h \leftarrow P} [\Gamma(h) = 1] = \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 0]$. Suppose for the sake of contradiction that when it takes its minimum, we have $\Pr_{h \leftarrow P} [\Gamma(h) = 1] < \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 0]$. As we can choose p_γ and γ as we wish, so we can increase $\Pr_{h \leftarrow P} [\Gamma(h) = 1]$ by increasing p_γ and γ , and by changing this, $\Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 0]$ will not increase. Thus, we can get a better tester unless $\Pr_{h \leftarrow P} [\Gamma(h) = 1] = \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 0]$, which contradicts the hypothesis that this is the best tester. A corresponding proof applies for the case $\Pr_{h \leftarrow P} [\Gamma(h) = 1] > \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 0]$. Thus, the optimal log-likelihood threshold tester $\Gamma(\cdot)$, where the max of type-1 error and type-2 error is minimized, has these two terms are equal to each other.

We set the tester $\Gamma(h)$ as the tester above and γ as the threshold in the tester above, we will show that this is the optimal tester that will minimize the max of type-1 error and type-2 error among all testers $\Gamma^*(h)$. In order to do that, notice the following:

$$\begin{aligned} & \mathbb{E}_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h)] - \mathbb{E}_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma^*(h)] - e^\gamma \left(\mathbb{E}_{h \leftarrow P} [\Gamma(h)] - \mathbb{E}_{h \leftarrow P} [\Gamma^*(h)] \right) \\ &= \sum_h (\Gamma(h) - \Gamma^*(h)) (Q_{[\varepsilon, \varepsilon']}(h) - e^\gamma P(h)) \\ &= \sum_{h \in L} (\Gamma(h) - \Gamma^*(h)) (Q_{[\varepsilon, \varepsilon']}(h) - e^\gamma P(h)) + \sum_{h \in M} (\Gamma(h) - \Gamma^*(h)) (Q_{[\varepsilon, \varepsilon']}(h) - e^\gamma P(h)) \\ &= \sum_{h \in L} (1 - \Gamma^*(h)) (Q_{[\varepsilon, \varepsilon']}(h) - e^\gamma P(h)) + \sum_{h \in M} (-\Gamma^*(h)) (Q_{[\varepsilon, \varepsilon']}(h) - e^\gamma P(h)) \end{aligned}$$

Here $L = \{h : Q_{[\varepsilon, \varepsilon']}(h) > e^\gamma P(h)\}$, $M = \{h : Q_{[\varepsilon, \varepsilon']}(h) < e^\gamma P(h)\}$. So we have $\Gamma(h) = 1$, when $h \in L$ and $\Gamma(h) = 0$, when $h \in M$. And notice that $\Gamma^*(h) \in [0, 1]$, so both components of the equation above are non-negative, which means the equation above is non-negative. And we notice that:

$$\begin{aligned} & \mathbb{E}_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h)] - \mathbb{E}_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma^*(h)] - e^\gamma \left(\mathbb{E}_{h \leftarrow P} [\Gamma(h)] - \mathbb{E}_{h \leftarrow P} [\Gamma^*(h)] \right) \\ &= \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 1] - \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma^*(h) = 1] - e^\gamma \left(\Pr_{h \leftarrow P} [\Gamma(h) = 1] - \Pr_{h \leftarrow P} [\Gamma^*(h) = 1] \right) \\ &= \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma^*(h) = 0] - \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 0] - e^\gamma \left(\Pr_{h \leftarrow P} [\Gamma(h) = 1] - \Pr_{h \leftarrow P} [\Gamma^*(h) = 1] \right) \end{aligned}$$

And by the definition of $\Gamma(h)$, we know that $\Pr_{h \leftarrow P} [\Gamma(h) = 1] = \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 0]$. Then suppose Γ^* is a better tester, which means $\Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma^*(h) = 0] < \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} [\Gamma(h) = 0]$ and $\Pr_{h \leftarrow P} [\Gamma^*(h) = 1] < \Pr_{h \leftarrow P} [\Gamma(h) = 1]$. However, this will make the equation above negative, which contradicts the fact that the equation is non-negative for any possible tester Γ^* . So there is no such a tester $\Gamma^*(h)$ with a better max of type-1 and type-2 error than $\Gamma(h)$. That finishes the proof. \square

Given that the optimal tester for distinguishing 2 hypotheses is a log-likelihood threshold tester, we now look to understand Chernoff bounds on the performance of such testers.

Lemma 21. *For any two distributions P, Q , we have that the log-likelihood tester with threshold γ has failure probability bounded by the Chernoff bound*

$$\Pr_{x \leftarrow P} \left[\log \frac{Q(x)}{P(x)} \geq \gamma \right] \stackrel{Ch}{\leq} \min_{u \leq 1} \sum_x P(x)^u Q(x)^{1-u} e^{-\gamma(1-u)}$$

Proof. We simply take the Chernoff bound of the distribution that takes value $\log \frac{Q(x)}{P(x)}$ with probability $P(x)$; for Chernoff parameter t , quick simplification gives $u = 1 - t$; hence $t \geq 0$ translates to the bound $u \leq 1$. \square

Plugging in $P, \mathcal{Q}_{[\varepsilon, \varepsilon']}$ to Lemma 21 with some threshold γ , and also, symmetrically, plugging in $\mathcal{Q}_{[\varepsilon, \varepsilon]}, P$ with the corresponding threshold $-\gamma$ yields

Corollary 22.

$$\min_{\gamma} \max \left\{ \Pr_{h \leftarrow P} \left[\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} \geq \gamma \right], \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon]}} \left[\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} \leq \gamma \right] \right\} \stackrel{Ch}{\leq} \min_{u \in [0, 1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$$

Proof. Explicitly plugging into Lemma 21 (where for the second invocation, the expression inside the probability is easily seen to be the negative of what is in Lemma 21, and hence equivalent) the left hand side is bounded by

$$\min_{\gamma} \max \left\{ \min_{u \leq 1} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u} e^{-\gamma(1-u)}, \min_{u \geq 0} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u} e^{\gamma u} \right\} \quad (17)$$

If $\gamma = 0$, then both components of the max are equal; and the min is attained at $u \in [0, 1]$ since $\sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$ equals 1 at $u = 0$ and $u = 1$ and is convex; and thus the overall expression equals $\min_{u \in [0, 1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$ as claimed.

Otherwise, without loss of generality, we consider the case $\gamma > 0$.

If 0 is not an optimal u for $\min_{u \geq 0} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u} e^{\gamma u}$ then the $e^{\gamma u}$ term is strictly greater than 1, and thus the min is greater than $\min_{u \in [0, 1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$, implying that γ is not actually the optimal choice of γ , a contradiction.

Otherwise, 0 is an optimal u , meaning that our expression is $\sum_h \mathcal{Q}_{[\varepsilon, \varepsilon']}(h) = 1$. And if 1 is the optimal value, then since this value is attained at both $u = 0$ and $u = 1$ and the expression $a^u b^{1-u}$ is convex in u for any $a, b \geq 0$, we have that the Equation 17 attains its minimum of 1, for $\gamma = 0$, equaling $\min_{u \in [0, 1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$ as desired. \square

We point out that the Chernoff bounds for a \leq tail are the same as those for a $<$ tail.

In light of Lemma 20, Corollary 22 is bounding the performance of the best distinguisher for P versus $\mathcal{Q}_{[\varepsilon, \varepsilon']}$.

And the best distinguisher of P from all distributions $\geq \varepsilon$ far from P must in particular distinguish P from $\mathcal{Q}_{[\varepsilon, \varepsilon]}$.

From Lemma 20, the error of the optimal tester is at least

$$\min_{\gamma} \max \left\{ \Pr_{h \leftarrow P} \left[\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon]}(h)}{P(h)} > \gamma \right], \Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon]}} \left[\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon]}(h)}{P(h)} < \gamma \right] \right\}$$

If the optimal γ is ≥ 0 then the max of these two expressions is at least the first one, which is at least $\Pr_{h \leftarrow P} \left[\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} > 0 \right]$; and the best Chernoff bound on this is $\min_{u \in [0, 1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$.

On the other hand, if the optimal γ is ≤ 0 , then the error of the optimal tester is at least the second term $\Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} \left[\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} < \gamma \right]$, which is at least $\Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} \left[\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} < 0 \right]$; and the best Chernoff bound on this is the expression from above $\min_{u \in [0, 1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$.

Definition 23. We thus define ch_1 to be the smaller of the ratio between $\Pr_{h \leftarrow P} \left[\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} > 0 \right]$ or $\Pr_{h \leftarrow \mathcal{Q}_{[\varepsilon, \varepsilon']}} \left[\log \frac{\mathcal{Q}_{[\varepsilon, \varepsilon']}(h)}{P(h)} < 0 \right]$ and the common Chernoff bound $\min_{u \in [0, 1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$. \square

Namely, the failure probability of the best tester is at least

$$ch_1 \cdot \min_{u \in [0, 1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$$

Towards understanding this bound and reexpressing it via a *second* layer of Chernoff bounds, we define the following “coin flip” random variables.

Definition 24. Let $q_{2,j} = 1 - q_{1,j} = q_j$ for convenience. Let h_j be the number of times domain element j is sampled. Define, for each $j \in \{1, \dots, n\}$ the independent random coin-flip variable X_j so that for each $r_j \in \{1, 2\}$ —denoting the outcome of the j^{th} coin flip—the variable X_j takes value $|x_{r_j, j} - y_j|$ with probability $\pi_{r_j, j}$ defined to be

$$\pi_{r_j, j} = \sum_{h_j} \frac{\text{poi}(k y_j, h_j)^u}{(q_1 \text{poi}(k x_{1,j}, h_j) + q_2 \text{poi}(k x_{2,j}, h_j))^u} q_{r_j, j} \text{poi}(k x_{r_j, j}, h_j)$$

where we emphasize that X_j is an unnormalized distribution, in that its probabilities may not sum to 1.

Lemma 25. For $u \in [0, 1]$, and using the probabilities π from Definition 24, we have

$$\sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u} \geq \frac{1}{\mathcal{Q}([\varepsilon, \varepsilon'])} \sum_{r: \sum_j |x_{r_j, j} - y_j| \in [\varepsilon, \varepsilon']} \prod_j \pi_{r_j, j} = \frac{1}{\mathcal{Q}([\varepsilon, \varepsilon'])} \Pr_{x_1 \leftarrow X_1, \dots, x_n \leftarrow X_n} \left[\sum_j x_j \in [\varepsilon, \varepsilon'] \right] \quad (18)$$

The best Chernoff bound for the right hand side is

$$\frac{1}{\mathcal{Q}([\varepsilon, \varepsilon'])} \min_{s \geq 0} e^{-s\varepsilon} \prod_j \left(e^{s|x_{1,j} - y_j|} \pi_{1,j} + e^{s|x_{2,j} - y_j|} \pi_{2,j} \right) \quad (19)$$

Proof. We have, since the joint probability $\mathcal{Q}(h, [\varepsilon, \varepsilon'])$ is at most the probability of the marginal

$\mathcal{Q}(h)$, that

$$\begin{aligned}
\sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u} &= \frac{1}{\mathcal{Q}([\varepsilon, \varepsilon'])^{1-u}} \sum_h P(h)^u \mathcal{Q}(h, [\varepsilon, \varepsilon'])^{1-u} \\
&\geq \frac{1}{\mathcal{Q}([\varepsilon, \varepsilon'])^{1-u}} \sum_h \frac{P(h)^u \mathcal{Q}(h, [\varepsilon, \varepsilon'])}{\mathcal{Q}(h)^u} \\
&= \frac{1}{\mathcal{Q}([\varepsilon, \varepsilon'])^{1-u}} \sum_h \left(\frac{P(h)^u}{\mathcal{Q}(h)^u} \sum_{r: \sum_j |x_{r,j} - y_j| \in [\varepsilon, \varepsilon']} \mathcal{Q}(h, r) \right) \quad (20)
\end{aligned}$$

Then we can easily compute $\mathcal{Q}(r, h) = \prod_j q_{r,j} \text{poi}(kx_{r,j}, h_j)$; $P(h) = \prod_j \text{poi}(ky_j, h_j)$; $\mathcal{Q}(h) = \prod_j (q_{1,j} \text{poi}(kx_{1,j}, h_j) + q_{2,j} \text{poi}(kx_{2,j}, h_j))$. Thus the right hand side of Equation 20 becomes, after pulling the r sum to the outside, and pulling the product over j before the sum over h :

$$\frac{1}{\mathcal{Q}([\varepsilon, \varepsilon'])^{1-u}} \sum_{r: \sum_j |x_{r,j} - y_j| \in [\varepsilon, \varepsilon']} \prod_j \sum_{h_j} \frac{\text{poi}(ky_j, h_j)^u}{(q_{1,j} \text{poi}(kx_{1,j}, h_j) + q_{2,j} \text{poi}(kx_{2,j}, h_j))^u} q_{r,j} \text{poi}(kx_{r,j}, h_j)$$

where the sum expression is seen to be exactly $\pi_{r,j}$ from Definition 24 as claimed.

This expression (ignoring the common $\frac{1}{\mathcal{Q}([\varepsilon, \varepsilon'])^{1-u}}$ factor out front) is seen to exactly compute the probability that the sum of the independent random coin-flips X_1, \dots, X_n is in the interval $[\varepsilon, \varepsilon']$. We can then apply standard Chernoff bounds, as desired. \square

Definition 26. Let ch_2 be the ratio between the right hand side of Equation 18 and its Chernoff bound, Equation 19, for the u defined as $\arg \min_{u \in [0,1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$

Combining the above results:

Lemma 27. *The failure probability of any algorithm for distinguishing $\text{Poi}(k)$ samples from P versus $\mathcal{Q}_{[\varepsilon, \varepsilon']}$ is at least $ch_1 \cdot ch_2 \cdot \frac{1}{\mathcal{Q}([\varepsilon, \varepsilon'])^{1-u}} \min_{s \geq 0} e^{-s\varepsilon} \prod_j (e^{s|x_{1,j} - y_j|} \pi_{1,j} + e^{s|x_{2,j} - y_j|} \pi_{2,j})$ for $u = \arg \min_{u \in [0,1]} \sum_h P(h)^u \mathcal{Q}_{[\varepsilon, \varepsilon']}(h)^{1-u}$.*

We bound $\mathcal{Q}([\varepsilon, \varepsilon']) \leq \mathcal{Q}(\leq \varepsilon')$ and bound this with a standard Chernoff bound, as

$$\mathcal{Q}(\leq \varepsilon') \leq e^{-\varepsilon' \alpha} \prod_j \left(q_j e^{\alpha|x_{1,j} - y_j|} + (1 - q_j) e^{\alpha|x_{2,j} - y_j|} \right)$$

where we will use Chernoff parameter α , with α from the optimization in Lemma 18. Since Chernoff bounds are upper bounds, and $\mathcal{Q}([\varepsilon, \varepsilon'])$ appears in the denominator of the expression of Lemma 27, we thus have a lower bound for the failure probability.

The expression from Lemma 27 is lower-bounded by its minimum over all u . Thus, from Lemma 27, expanding out the expressions for π from Definition 24 and substituting in the above Chernoff bound for $\mathcal{Q}([\varepsilon, \varepsilon'])$ we have:

Corollary 28. *The failure probability of any algorithm for distinguishing $\text{Poi}(k)$ samples from P versus $\mathcal{Q}_{[\varepsilon, \varepsilon']}$ is at least*

$$ch_1 \cdot ch_2 \cdot e^{\alpha(\varepsilon' - \varepsilon)} \cdot \min_{u \in [0,1], s \geq 0} e^{\varepsilon(-s+(1-u)\alpha)} \prod_j \sum_i \frac{poi(k y_j, i)^u (q_j poi(k x_{1,j}, i) e^{s|x_{1,j}-y_j|} + (1-q_j) poi(k x_{2,j}, i) e^{s|x_{2,j}-y_j|})}{(q_1 poi(k x_{1,j}, i) + q_2 poi(k x_{2,j}, i))^u (q_j e^{\alpha|x_{1,j}-y_j|} + (1-q_j) e^{\alpha|x_{2,j}-y_j|})^{1-u}} \quad (21)$$

Finally, we claim that, except for the Chernoff factors ch_1, ch_2 and the factor $e^{\alpha(\varepsilon' - \varepsilon)}$, Equation 21 is exactly equal to our upper bound, as expressed in either Lemma 18 or Equation 1. We will use the fact that Equation 29 below was optimized in Lemma 18 to show how to optimize this related equation with respect to s (deferring for the moment the optimization over u).

Lemma 29. *We claim that if $u \in (0, 1)$ and the derivative of this expression*

$$\log e^{\varepsilon(1-u)\alpha} \prod_j \left(\frac{\sum_i poi(y_j, i)^u (q_j poi(x_{1,j}, i) + (1-q_j) poi(x_{2,j}, i))^{1-u}}{(q_j e^{|x_{1,j}-y_j|^\alpha} + (1-q_j) e^{|x_{2,j}-y_j|^\alpha})^{1-u}} \right) \quad (22)$$

with respect to any q_j is 0, and the derivative with respect to α is also 0, then the below expression when minimized over s achieves its global minimum at $s = 0$:

$$\log e^{-s\varepsilon} \prod_j \left(\sum_i \left(\frac{poi(y_j, i)}{(q_j poi(x_{1,j}, i) + (1-q_j) poi(x_{2,j}, i))} \right)^u (q_j e^{s|x_{1,j}-y_j|} poi(x_{1,j}, i) + (1-q_j) e^{s|x_{2,j}-y_j|} poi(x_{2,j}, i)) \right) \quad (23)$$

Proof. Equation 23 is clearly convex in s , since log-convexity is preserved under sums and products; we thus show its s derivative is 0 at $s = 0$ to finish the proof of global optimality. The proof is ultimately straightforward, where we find expressions for the derivatives of Equation 22, and find the right linear combination of them to imply that the s derivative of Equation 23 is 0 at $s = 0$.

The q_j derivative of Equation 22, which equals 0, is

$$(1-u) \left(\frac{\sum_i \left(\frac{poi(y_j, i)}{q_j (poi(x_{1,j}, i) + (1-q_j) poi(x_{2,j}, i))} \right)^u (poi(x_{1,j}, i) - poi(x_{2,j}, i))}{\sum_i (q_j poi(x_{1,j}, i) + (1-q_j) poi(x_{2,j}, i))^{1-u} (poi(y_j, i))^u} - \frac{(e^{|x_{1,j}-y_j|^\alpha} - e^{|x_{2,j}-y_j|^\alpha})}{(q_j e^{|x_{1,j}-y_j|^\alpha} + (1-q_j) e^{|x_{2,j}-y_j|^\alpha})} \right)$$

which implies, since $u \in (0, 1)$, that

$$\frac{\sum_i \left(\frac{poi(y_j, i)}{q_j (poi(x_{1,j}, i) + (1-q_j) poi(x_{2,j}, i))} \right)^u (poi(x_{1,j}, i) - poi(x_{2,j}, i))}{\sum_i (q_j poi(x_{1,j}, i) + (1-q_j) poi(x_{2,j}, i))^{1-u} (poi(y_j, i))^u} = \frac{(e^{|x_{1,j}-y_j|^\alpha} - e^{|x_{2,j}-y_j|^\alpha})}{(q_j e^{|x_{1,j}-y_j|^\alpha} + (1-q_j) e^{|x_{2,j}-y_j|^\alpha})} \quad (24)$$

On the other hand, the α derivative of Equation 23, which is equals to 0, is

$$\varepsilon(1-u) - (1-u) \left(\sum_j \frac{(q_j |x_{1,j} - y_j| e^{|x_{1,j}-y_j|^\alpha} + (1-q_j) |x_{2,j} - y_j| e^{|x_{2,j}-y_j|^\alpha})}{(q_j e^{|x_{1,j}-y_j|^\alpha} + (1-q_j) e^{|x_{2,j}-y_j|^\alpha})} \right)$$

Therefore, since $u \in (0, 1)$, we have

$$\sum_j \frac{(q_j |x_{1,j} - y_j| e^{|x_{1,j}-y_j|^\alpha} + (1-q_j) |x_{2,j} - y_j| e^{|x_{2,j}-y_j|^\alpha})}{(q_j e^{|x_{1,j}-y_j|^\alpha} + (1-q_j) e^{|x_{2,j}-y_j|^\alpha})} = \varepsilon \quad (25)$$

Going back to Equation 23, the s derivative of the log of the j^{th} term of Equation 23, evaluated at $s = 0$ is

$$\frac{\sum_i \left(\frac{\text{poi}(y_j, i)}{q_j(\text{poi}(x_{1,j}, i) + (1-q_j)\text{poi}(x_{2,j}, i))} \right)^u (q_j |x_{1,j} - y_j| \text{poi}(x_{1,j}, i) + (1-q_j) |x_{2,j} - y_j| \text{poi}(x_{2,j}, i))}{\sum_i \text{poi}(y_j, i)^u (q_j \text{poi}(x_{1,j}, i) + (1-q_j) \text{poi}(x_{2,j}, i))^{1-u}}$$

Notice that the last term in the numerator can be decomposed as

$$(q_j |x_{1,j} - y_j| \text{poi}(x_{1,j}, i) + (1-q_j) |x_{2,j} - y_j| \text{poi}(x_{2,j}, i)) = A(\text{poi}(x_{1,j}, i) - \text{poi}(x_{2,j}, i)) + B(q_j \text{poi}(x_{1,j}, i) + (1-q_j) \text{poi}(x_{2,j}, i))$$

for $A = q_j(1-q_j)(|x_{1,j} - y_j| - |x_{2,j} - y_j|)$ and $B = q_j |x_{1,j} - y_j| + (1-q_j) |x_{2,j} - y_j|$.

Substituting back to the s derivative of the log of the j^{th} term of Equation 23, evaluated at $s = 0$, we have

$$\begin{aligned} & \frac{\sum_i \left(\frac{\text{poi}(y_j, i)}{q_j(\text{poi}(x_{1,j}, i) + (1-q_j)\text{poi}(x_{2,j}, i))} \right)^u A(\text{poi}(x_{1,j}, i) - \text{poi}(x_{2,j}, i))}{\sum_i \text{poi}(y_j, i)^u (q_j \text{poi}(x_{1,j}, i) + (1-q_j) \text{poi}(x_{2,j}, i))^{1-u}} \\ & + \frac{\sum_i \left(\frac{\text{poi}(y_j, i)}{q_j(\text{poi}(x_{1,j}, i) + (1-q_j)\text{poi}(x_{2,j}, i))} \right)^u B(q_j \text{poi}(x_{1,j}, i) + (1-q_j) \text{poi}(x_{2,j}, i))}{\sum_i \text{poi}(y_j, i)^u (q_j \text{poi}(x_{1,j}, i) + (1-q_j) \text{poi}(x_{2,j}, i))^{1-u}} \\ & = A \frac{(e^{|x_{1,j} - y_j| \alpha} - e^{|x_{2,j} - y_j| \alpha})}{(q_j e^{|x_{1,j} - y_j| \alpha} + (1-q_j) e^{|x_{2,j} - y_j| \alpha})} + B \quad (24) \\ & = \frac{((A + Bq_j) e^{|x_{1,j} - y_j| \alpha} + ((1-q_j)B - A) e^{|x_{2,j} - y_j| \alpha})}{(q_j e^{|x_{1,j} - y_j| \alpha} + (1-q_j) e^{|x_{2,j} - y_j| \alpha})} \\ & = \frac{(q_j |x_{1,j} - y_j| e^{|x_{1,j} - y_j| \alpha} + (1-q_j) |x_{2,j} - y_j| e^{|x_{2,j} - y_j| \alpha})}{(q_j e^{|x_{1,j} - y_j| \alpha} + (1-q_j) e^{|x_{2,j} - y_j| \alpha})} \end{aligned}$$

which is exactly the j -th term of the left hand side of Equation 25. Hence, summing up the previous expression for all j gives us ϵ . This shows that the derivative of s at $s = 0$ of Equation 23 is $-\epsilon + \epsilon = 0$, concluding the proof. \square

The next lemma concludes this section, summarizing that, except for the slack in the Chernoff bounds, and our choice of $\epsilon' > \epsilon$, our upper and lower bounds match.

Lemma 30. *Under the conditions of Lemma 18—namely, when the testing upper bound of Equation 7 is optimized—then we have a lower bound for any testing algorithm that equals our upper bound times $ch_1 \cdot ch_2 \cdot e^{\alpha(\epsilon' - \epsilon)}$.*

Proof. Corollary 28 provides our testing lower bound, since Lemma 18 guarantees its conditions are satisfied. We now show that the minimization expression in Equation 21 (namely, without the initial $ch_1 \cdot ch_2 \cdot e^{\alpha(\epsilon' - \epsilon)}$ terms) exactly equals our testing upper bound—as analyzed in Lemma 18.

From Corollary 28, the minimization expression in Equation 21 without the initial $e^{(1-u)\alpha}$ term and without the second term in the denominator, is minimized over s when $s = 0$; thus the min over u and s is just the min over u of the corresponding expression when $s = 0$ is substituted, namely

$$\min_{u \in [0,1]} e^{\epsilon(1-u)\alpha} \prod_j \sum_i \frac{\text{poi}(k y_j, i)^u (q_j \text{poi}(k x_{1,j}, i) + (1-q_j) \text{poi}(k x_{2,j}, i))^{1-u}}{(q_j e^{\alpha |x_{1,j} - y_j|} + (1-q_j) e^{\alpha |x_{2,j} - y_j|})^{1-u}}$$

The logarithm of the expression being minimized is exactly Equation 14, which Lemma 18 guarantees has u derivative equal to 0; further, the expression being minimized is log-convex in u since log-convexity is preserved under both sums and products. Thus its global minimum is attained for the u from the upper bound optimization, and its global minimum equals our upper bound, as expressed in Equation 14 or Equation 1. \square

5 Gärtner-Ellis Theorem

We introduce a special version of Gärtner-Ellis Theorem, to analyze the convergence of Chernoff bounds when the number of domain elements n goes to infinity.

The Gärtner-Ellis theorem is a generalization of Cramér’s theorem, where Cramér’s theorem essentially says that “in the limit as $n \rightarrow \infty$, Chernoff bounds on the mean of i.i.d. variables are tight, to $1 + o(1)$ factors in the exponent. If Z_n denotes the mean of n independent copies of some real-valued random variable, and letting

$$\Lambda_n(\lambda) = \log \mathbb{E}[e^{\lambda Z_n}]$$

denote its corresponding moment generating function, then the log tail probabilities $\Pr[Z_n \geq x]$ are bounded by Chernoff bounds as $\min_{\lambda \geq 0} \Lambda_n(\lambda) - \lambda x = \min_{\lambda \geq 0} \Lambda_n(n\lambda) - n\lambda x$. Since we assume that Z_n is the mean of i.i.d. random variables, then $\frac{1}{n}\Lambda_n(n\lambda)$ is identical for all n , and we may instead represent this as $\Lambda(\lambda)$, leading to a log Chernoff bound of $n \min_{\lambda \geq 0} \Lambda(\lambda) - \lambda x = -n\Lambda^*(x)$, where we define

$$\Lambda^*(x) = \sup_{\lambda \in \mathbb{R}} \lambda x - \Lambda(\lambda)$$

to be the Legendre-Fenchel transform of $\Lambda(\lambda)$. Cramér’s theorem says that this bound is tight in the limit as $n \rightarrow \infty$, where we scale the log probability by $\frac{1}{n}$ so that the right hand side does not go to infinity with n . Namely, while $\log \Pr[Z_n \geq x] \leq -n\Lambda^*(x)$ for any n , by Chernoff bounds, Cramér’s theorem says that the limit is tight, in the sense that $\lim_{n \rightarrow \infty} \frac{1}{n} \log \Pr[Z_n \geq x] = -\Lambda^*(x)$.

The Gärtner-Ellis theorem generalizes this to non-independent random variables. Namely, Z_n is no longer restricted to be the mean of n i.i.d. random variables but can now be any distribution, depending on n in an almost arbitrary way—provided that the log moment generating function has a limit as $n \rightarrow \infty$, and avoids a few other pathologies as specified in Theorem 32.

Assumption 31. *For each $\lambda \in \mathbb{R}$, the logarithmic moment generating function has a limit*

$$\Lambda(\lambda) = \lim_{n \rightarrow \infty} \frac{1}{n} \Lambda_n(n\lambda) \tag{26}$$

which is finite and differentiable.

Then we can introduce the Gärtner-Ellis Theorem as follows:

Theorem 32 (special case of Gärtner-Ellis Theorem). *If Assumption 31 holds for Z_n , and if x is not a “discrete point”, in the sense that Λ^* is continuous at x then we have:*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \Pr[Z_n > x] = \lim_{n \rightarrow \infty} \frac{1}{n} \log \Pr[Z_n \geq x] = -\Lambda^*(x) \tag{27}$$

if Λ^ is non-increasing at x , and otherwise*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \Pr[Z_n < x] = \lim_{n \rightarrow \infty} \frac{1}{n} \log \Pr[Z_n \leq x] = -\Lambda^*(x) \tag{28}$$

5.1 Identity testing in the limit $n \rightarrow \infty$

While we show upper and lower bounds for identity testing, we ultimately would like to compare these bounds to say that our tester is near-optimal. The Gärtner Ellis theorem provides a tool for doing this, in the limit as $n \rightarrow \infty$. Unfortunately, the Gärtner Ellis theorem does not give much insight into the rate of convergence, and a more targeted investigation of when “Chernoff bounds are tight” seems warranted. Nonetheless, we view the Gärtner Ellis theorem as saying, essentially, that we should expect our Chernoff bounds to be tight for finite sample cases, since asymptotically they are tight.

We specifically consider a limit where we send $n \rightarrow \infty$ and $k \rightarrow \infty$ proportionally to n , and where the upper bounds of Section 3 on the log failure probability are exactly proportional to n . Explicitly, for any hypothesis distribution P , we consider “subdividing it by factor s ” to mean the process where each domain element of P is split into s new domain elements each of $\frac{1}{s}$ the probability; this new distribution, which we denote $P^{sub(s)}$ will have $n \cdot s$ domain elements; to make up for the reduced mass of each domain elements, we will sample $Poi(k \cdot s)$ times instead of $Poi(k)$ times.

The basic claim is that the optimization of Section 3 does not depend on s , except for trivial scaling. Explicitly, consider Equation 1 as we modify s . The hypothesis histogram entries $h_{y_j}^P$ will scale proportionally to s , and we will thus choose to scale the histogram entries of the optimization variables $h_x^{Q|y_j}$ to scale with s ; the probability masses, as measured by y_j, x , scale inversely with s , and thus the Poisson parameters ky_j or kx are invariant to s ; the optimization variables \mathbf{c}, t, t' will be invariant to s ; the constraints will thus be invariant to s ; thus the overall objective function, since it has a factor of h in both the first and second term, will scale proportionally to s . Since this optimization computes an upper bound on the log failure probability of our tester, this means that the tester coefficients ($c_{i,j}$) are independent of s , and the failure probability gets raised to the s power (thus decaying exponentially with s).

Thus $\frac{1}{s}$ times the log failure probability is independent of s ; and our overall claim is that this exactly matches our lower bounds, in the limit $s \rightarrow \infty$ (meaning that $n, k \rightarrow \infty$).

Recall that we constructed a lower bound parameterized by ε' . We first show, by the Gärtner-Ellis theorem that, for each fixed $\varepsilon' > \varepsilon$, the two Chernoff factors ch_1, ch_2 vanish in the limit; then we take the limit $\varepsilon' \rightarrow \varepsilon$ so that the final remaining factor $e^{\alpha(\varepsilon' - \varepsilon)}$ vanishes, leading, by Lemma 30, to the desired conclusion.

Lemma 33. *Fixing a hypothesis distribution P , and a number of samples k , and considering the limit as $s \rightarrow \infty$ of $P^{sub(s)}$, while we take $Poi(ks)$ samples, then, for any fixed $\varepsilon' > \varepsilon$, we have $\lim_{s \rightarrow \infty} \frac{1}{s} \log ch_1(P^{sub(s)}, ks, \varepsilon, \varepsilon') = 0$ and $\lim_{s \rightarrow \infty} \frac{1}{s} \log ch_2(P^{sub(s)}, ks, \varepsilon, \varepsilon') = 0$*

Thus by Lemma 30 we have that for any sequence of testers, on input $P^{sub(s)}$ as $s \rightarrow \infty$, their log failure probability, times $\frac{1}{s}$, must have \liminf greater than or equal to our upper bound plus $\alpha(\varepsilon' - \varepsilon)$, where α is negative, and is evaluated from our upper bound at $s = 1$.

Thus sending $\varepsilon' \rightarrow \varepsilon$ yields that any sequence of testers must have log failure probability times $\frac{1}{s}$ that has \liminf (as $s \rightarrow \infty$) greater than or equal to our upper bound. Since our upper bound is explicitly an upper bound on testing performance, then there is a sequence of testers—namely, with coefficients $c_{i,j}$, independent of s , such that $\frac{1}{s}$ times the log of its failure probability has the limit exactly as specified by our upper bound, and no testers can do any better in this limit. This yields the final part of our main result, Theorem 4.

References

- [ACK18] Jayadev Acharya, Clément Canonne, and Gautam Kamath. A chasm between identity and equivalence testing with conditional queries. *Theory of Computing*, 14:1–46, 2018.
- [ADJ⁺11] Jayadev Acharya, Hirakendu Das, Ashkan Jafarpour, Alon Orlitsky, and Shengjun Pan. Competitive closeness testing. In Sham M. Kakade and Ulrike von Luxburg, editors, *Proceedings of the 24th Annual Conference on Learning Theory*, volume 19 of *Proceedings of Machine Learning Research*, pages 47–68, Budapest, Hungary, 09–11 Jun 2011. PMLR.
- [ADK15] Jayadev Acharya, Constantinos Daskalakis, and Gautam Kamath. Optimal testing for properties of distributions. *Advances in Neural Information Processing Systems*, 28, 2015.
- [BC17] T. Batu and C. L. Canonne. Generalized uniformity testing. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 880–889, Los Alamitos, CA, USA, oct 2017. IEEE Computer Society.
- [BCL20] Sebastien Bubeck, Sitan Chen, and Jerry Li. Entanglement is necessary for optimal quantum property testing. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 692–703. IEEE, 2020.
- [BFF⁺01] T. Batu, E. Fischer, L. Fortnow, R. Kumar, R. Rubinfeld, and P. White. Testing random variables for independence and identity. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 442–451, 2001.
- [BFR⁺00] Tugkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D Smith, and Patrick White. Testing that distributions are close. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 259–269. IEEE, 2000.
- [BFR⁺13] Tuğkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D. Smith, and Patrick White. Testing closeness of discrete distributions. *J. ACM*, 60(1), feb 2013.
- [BOW19] Costin Bădescu, Ryan O’Donnell, and John Wright. Quantum state certification. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 503–514, New York, NY, USA, 2019. Association for Computing Machinery.
- [BR60] Raghu Raj Bahadur and R Ranga Rao. On deviations of the sample mean. *The Annals of Mathematical Statistics*, 31(4):1015–1027, 1960.
- [Can20] 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.
- [CDSS14] Siu-On Chan, Ilias Diakonikolas, Rocco A. Servedio, and Xiaorui Sun. Efficient density estimation via piecewise polynomial approximation. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC ’14, page 604–613, New York, NY, USA, 2014. Association for Computing Machinery.

- [CDVV14] Siu-On Chan, Ilias Diakonikolas, Gregory Valiant, and Paul Valiant. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, page 1193–1203, USA, 2014. Society for Industrial and Applied Mathematics.
- [CFGM16] Sourav Chakraborty, Eldar Fischer, Yonatan Goldhirsh, and Arie Matsliah. On the power of conditional samples in distribution testing. *SIAM Journal on Computing*, 45(4):1261–1296, 2016.
- [Cra38] Harald Cramér. Sur un nouveau théoreme-limite de la théorie des probabilités. *Actual. Sci. Ind.*, 736:5–23, 1938.
- [CRS15] Clément L. Canonne, Dana Ron, and Rocco A. Servedio. Testing probability distributions using conditional samples. *SIAM Journal on Computing*, 44(3):540–616, 2015.
- [DDS⁺] Constantinos Daskalakis, Ilias Diakonikolas, Rocco A. Servedio, Gregory Valiant, and Paul Valiant. *Testing k -Modal Distributions: Optimal Algorithms via Reductions*, pages 1833–1852.
- [DGK⁺21] Ilias Diakonikolas, Themis Gouleakis, Daniel M. Kane, John Peebles, and Eric Price. Optimal testing of discrete distributions with high probability. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, page 542–555, New York, NY, USA, 2021. Association for Computing Machinery.
- [DGPP18] Ilias Diakonikolas, Themis Gouleakis, John Peebles, and Eric Price. Sample-optimal identity testing with high probability. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [DKN15] Ilias Diakonikolas, Daniel M. Kane, and Vladimir Nikishkin. Testing identity of structured distributions. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, page 1841–1854, USA, 2015. Society for Industrial and Applied Mathematics.
- [DKS18] Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. Sharp bounds for generalized uniformity testing. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [FB28] R. A. Fisher and Bhai Balmukand. The estimation of linkage from the offspring of selfed heterozygotes. *Journal of Genetics*, 20(1):79–92, Jul 1928.
- [Gol17] Oded Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017.
- [Gol20] Oded Goldreich. The uniform distribution is complete with respect to testing identity to a fixed distribution. In *Computational Complexity and Property Testing: On the Interplay Between Randomness and Computation*, pages 152–172. Springer, 2020.

- [GP22] Shivam Gupta and Eric Price. Sharp constants in uniformity testing via the Huber statistic. In Po-Ling Loh and Maxim Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 3113–3192. PMLR, 02–05 Jul 2022.
- [OW15] Ryan O’Donnell and John Wright. Quantum spectrum testing. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC ’15, page 529–538, New York, NY, USA, 2015. Association for Computing Machinery.
- [Pan08] Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.
- [Ron08] Dana Ron. Property testing: A learning theory perspective. *Foundations and Trends® in Machine Learning*, 1(3):307–402, 2008.
- [Sio58] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.
- [Tre16] François Trèves. *Topological Vector Spaces, Distributions and Kernels: Pure and Applied Mathematics, Vol. 25*, volume 25. Elsevier, 2016.
- [VV17] Gregory Valiant and Paul Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017.
- [WZXX21] Jiyuan Wang, Qian Zhang, Guoqing Harry Xu, and Miryung Kim. Qdiff: Differential testing of quantum software stacks. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 692–704. IEEE, 2021.