

Improved Bounds for High-Dimensional Equivalence and Product Testing using Subcube Queries

Tomer Adar*

Eldar Fischer[†]

Amit Levi[‡]

July 9, 2024

Abstract

We study property testing in the subcube conditional model introduced by Bhattacharyya and Chakraborty (2017). We obtain the first equivalence test for n -dimensional distributions that is quasi-linear in n , improving the previously known $\tilde{O}(n^2/\varepsilon^2)$ query complexity bound to $\tilde{O}(n/\varepsilon^2)$. We extend this result to general finite alphabets with logarithmic cost in the alphabet size.

By exploiting the specific structure of the queries that we use (which are more restrictive than general subcube queries), we obtain a cubic improvement over the best known test for distributions over $\{1, \dots, N\}$ under the interval querying model of Canonne, Ron and Servedio (2015), attaining a query complexity of $\tilde{O}((\log N)/\varepsilon^2)$, which for fixed ε almost matches the known lower bound of $\Omega((\log N)/\log \log N)$. We also derive a product test for n -dimensional distributions with $\tilde{O}(n/\varepsilon^2)$ queries, and provide an $\Omega(\sqrt{n}/\varepsilon^2)$ lower bound for this property.

*Technion - Israel Institute of Technology, Israel. Email: tomer-adar@campus.technion.ac.il.

[†]Technion - Israel Institute of Technology, Israel. Email: eldar@cs.technion.ac.il. Research supported by an Israel Science Foundation grant number 879/22.

[‡]University of Haifa and Technion - Israel Institute of Technology, Israel. Email: alevi@ds.haifa.ac.il.

1 Introduction

Property testing seeks to efficiently distinguish between objects that have some property and objects that are ε -far from any object that has it, with respect to a predefined metric and a proximity parameter ε . Property testing of functions, and in particular string testing, was initiated in [RS96, BLR93]. The term “efficiently” usually refers to a sublinear amount of resources for moderately-sized tasks, but for high-dimensional inputs practicality mandates to restrict this amount further to the poly-logarithmic scale.

The study of distribution testing was implicitly initiated in [GR11] (motivated by a problem in graph testing), and formally defined as the *sampling model* in [BFR⁺00, BFF⁺01]. In this model, the algorithm gets access to a sequence of independent unconditional samples from the input distribution, and then decides whether to accept or reject based on them. A major topic of investigation concerns testing that a distribution over $\{1, \dots, n\}$ is the uniform one. A long string of results, starting with the original [GR11] and culminating in [Pan08], has reached the tight bound of $\Theta(\sqrt{n}/\varepsilon^2)$.

A square-root lower bound on one of the most basic of properties is impractical in many real-world settings. For example, the square-root sample complexity of uniformity testing is impractical when the input ranges over 100-bit binary strings, an example which is still smaller than real-world inputs. There are three approaches to overcome this problem: the first, on which we expand below, is moving to stronger query models. The second is narrowing the scope of the allowable inputs, such as *product distributions* and *Bayesian networks* [CDKS17]. The third is moving to a model where the metric is more lax (usually coupled with an even weaker query model), such as the Huge Object Model defined in [GR23] that is concerned with the earth-mover distance metric.

We focus here on the first (and most common) approach to the scaling problem, that of considering a model with stronger queries. In the distribution testing setting, this usually focuses on *conditional sampling*. Instead of drawing a sequence of independent unconditional samples, we allow the algorithm to choose a subset of the possible outcomes and draw a sample conditioned on belonging to it. The models within this paradigm differ in the subset conditions they allow.

The first investigation of such a model was the fully conditional model [CFG13, CRS15]. In this model, the algorithm can choose *any* subset to condition on. This model is very powerful but not realistic. If we are able to restrict the input distribution to any subset, we probably already have access to an explicit description of the distribution and thus have no need to sample it.

Further studies consider restricted forms of the conditional model. For example, [CRS15] considers two additional models: the pair model and the interval model. In the pair model, the algorithm can still draw unconditional samples, and additionally it can draw conditional samples from any subset of two elements. Uniformity testing in the pair model requires $\tilde{\Theta}(1/\varepsilon^2)$ samples. In the interval model, we test distributions over $\Omega = \{1, \dots, N\}$ (for some N), and the algorithm can condition on interval sets, which are sets of the form $\{a \leq x \leq b\}$. The lower bound for uniformity testing in the interval model is $\tilde{\Omega}(\log N)$. The best known upper bound is $\tilde{O}((\log N)^3/\varepsilon^2)$.

The subcube conditional model [BC17] is motivated by database analysis. In this model, we test distributions over the set $\prod_{i=1}^n \Omega_i$, and the algorithm can query *subcube subsets*, which are sets of the form $\prod_{i=1}^n A_i$ where $A_i \subseteq \Omega_i$ for every $1 \leq i \leq n$. While not being extremely restrictive, its queries correspond to selection by attribute values, which is common in practice. Some of the prior work refers to a weaker variant of this model, where each A_i is either trivial ($A_i = \Omega_i$) or a singleton

($A_i = \{a_i\}$ for some $a_i \in \Omega_i$). In this paper we mainly deal with the strong model, where general $A_i \subseteq \Omega_i$ are allowed. In the most investigated setting, the binary setting where $\Omega_i = \{0, 1\}$ for all i , the two models are the same.

Uniformity testing [CCK⁺21, CM24] can be done in the weak model using $\tilde{O}(m^{21}\sqrt{n}/\varepsilon^2)$ queries where $m = \max_i |\Omega_i|$, and requires at least $\Omega(\sqrt{mn}/\varepsilon^2)$ queries [BGKV21]. Other properties studied under the subcube model include identity to some fixed distribution [BC17, BCŠV23] and having a probability density function supported on a low-dimensional subspace [CJLW21].

Two related properties are of particular interest in distribution testing, the identity property and the equivalence property. In the identity property, a *reference distribution* is given to the algorithm in advance, and the task is to check whether the input distribution is identical to it. The equivalence property (henceforth: EQUIVALENCE) has an input consisting of *two* distributions, both of which accessible through the testing model, and the task is to check whether they are equal to each other.

In the subcube conditional model, since the input distribution is defined over a set of tuples, another natural property is that of being a product distribution (henceforth: PRODUCT). A distribution over tuples is called a product if its entries are independently distributed.

Our main upper bound result is a test for EQUIVALENCE that for the binary setting ($\Omega_i = \{0, 1\}$) uses only $\tilde{O}(n/\varepsilon^2)$ subcube queries, which improves on the previously known result of $\tilde{O}(n^2/\varepsilon^2)$ from [BC17]. Additionally our result uses only *prefix queries* from one distribution and *marginal prefix queries* from the other, which we define below and are rather restricted forms of subcube queries. One can think of prefix queries as the queries that can be made fast when the database is sorted according to a concatenation of its attributes in a pre-defined order (functioning as its primary key). Importantly, the use of restricted queries allows us to derive an improved test also for PRODUCT, and to generalize the test to general $\Omega_1, \dots, \Omega_n$ with a logarithmic cost in the alphabet size.

The restricted form of our queries also allows us to tighten the previously known upper bound on equivalence testing (and through it the special case of uniformity testing) in the interval conditional model, obtaining an upper bound of $\tilde{O}(\log N/\varepsilon^2)$ interval queries, which matches the lower bound for every fixed ε up to poly-double-logarithmic factors in N .

We complement our upper bound for PRODUCT with an $\Omega(\sqrt{n}/\varepsilon^2)$ lower bound. The question of whether we can go below $O(n)$ for testing our properties (even for the binary setting and a fixed ε) remains open.

2 Organization of the paper

In Section 3 we summarize our contributions. After some preliminaries in Section 4, we provide the core of our main proofs, followed by the more technical details. In Section 5 we provide the EQUIVALENCE testing upper bound in the binary setting, and a short proof of the corollary about interval queries. In Section 6 we provide the PRODUCT testing lower bound.

The technical part of the paper follows. In Section 7 we prove the technical lemmas whose proofs were deferred from Section 5. Then we prove the theorems derived from this upper bound: in Section 8 we extend the EQUIVALENCE test to the non-binary setting, and in Section 9 we derive a test for PRODUCT. In Section 10 we prove the technical lemmas from Section 6.

All upper bound proofs implicitly construct their algorithms. For reference, explicit representations of the binary setting algorithms for EQUIVALENCE and PRODUCT are given in the appendix.

3 Our results

We improve on the previously known result of $\tilde{O}(n^2/\varepsilon^2)$ queries for equivalence testing of two distributions over $\{0, 1\}^n$ [BC17]. Our methods can also supersede the $\tilde{O}(n/\varepsilon)$ algorithm of [BCŠV23] which tests identity with a distribution given in advance that belongs to a very restricted class of inputs (their parameter refers to KL-divergence, which indeed incurs a quadratic gap when converted to total-variation distance). We provide more details on the latter below (before Lemma 5.13).

Theorem 5.8. *Let τ, μ be two distributions over $\{0, 1\}^n$, where τ is accessible through the prefix oracle access and μ is accessible through the marginal prefix oracle access. For every $\varepsilon > 0$ we can distinguish between $\tau = \mu$ and $d_{\text{TV}}(\tau, \mu) > \varepsilon$ using $\tilde{O}(n/\varepsilon^2)$ queries.*

The prefix queries in the statement above are a special case of subcube queries, but they can also be seen as interval queries, allowing us to prove the following corollary:

Corollary 5.9. *Let τ, μ be two distributions over $\{1, \dots, N\}$, both accessible through the interval oracle. Then we can distinguish between $\tau = \mu$ and $d_{\text{TV}}(\tau, \mu) > \varepsilon$ using $\tilde{O}((\log N)/\varepsilon^2)$ queries.*

We show a lower bound for testing a distribution μ over $\{0, 1\}^n$ for being a product. To the best of our knowledge, it is the first lower bound for product testing in the binary setting. Our construction is similar to the lower bound for uniformity testing of [CDKS17].

Theorem 6.1. *Every ε -test for PRODUCT must make at least $\Omega(\sqrt{n}/\varepsilon^2)$ subcube queries.*

We generalize our upper bound for equivalence testing to strings of size n over larger alphabets. Our result is incomparable with the previously known result of $\tilde{O}(\frac{n^5}{\varepsilon^5} \log \log |\Omega|)$ [BC17]. Note that the cited result refers to strings over a single alphabet (that is, Ω^n), whereas our result refers to strings over mixed alphabets (that is, $\prod_{i=1}^n \Omega_i$). Additionally, our result is more efficient when $|\Omega|$ is not very large with respect to n .

Theorem 8.5. *Let μ and τ be two distributions over $\prod_{i=1}^n \Omega_i$, where $\Omega_1, \dots, \Omega_n$ are all finite. If μ is accessible through the marginal prefix oracle and τ is accessible through the prefix oracle, then we can distinguish between $\tau = \mu$ and $d_{\text{TV}}(\tau, \mu) > \varepsilon$ using $\tilde{O}(\sum_{i=1}^n \log_2 |\Omega_i|/\varepsilon^2)$ queries.*

We apply the same generalization for product testing as well. In the binary setting, it also improves on the previously known result of $\tilde{O}(n^2/\varepsilon^2)$ [BC17].

Theorem 9.4. *Let μ be a distribution over $\prod_{i=1}^n \Omega_i$. For every $0 < \varepsilon < 1$, we can distinguish between the case where μ is a product distribution and the case where it is ε -far from every product distribution at the cost of $\tilde{O}(\sum_{i=1}^n \log |\Omega_i|/\varepsilon^2)$ subcube queries. Moreover, if $|\Omega_i| = 2$ for every $1 \leq i \leq n$, then all these queries are prefix queries.*

4 Preliminaries

For brevity, for $m \in \mathbb{N}$ we let $[m]$ denote the set $\{1, \dots, m\}$.

Definition 4.1 (Bernoulli distribution). Let $0 \leq p \leq 1$. The *Bernoulli distribution with parameter p* , denoted by $\text{Ber}(p)$, is the distribution over $\{0, 1\}$ whose probability to draw 1 is p .

Definition 4.2 (Conditional distribution). Let μ be a distribution over Ω , and let $B \subseteq \Omega$ be an event. The corresponding *conditional distribution* is denoted by $\mu|_B$ and is defined by $\mu|_B(x) = 0$ for $x \notin B$ and $\mu|_B(x) = \mu(x)/\Pr_\mu[B]$ for $x \in B$.

Definition 4.3 (Index-restricted distribution). Let μ be a distribution over $\prod_{i=1}^n \Omega_i$, and let $I \subseteq [n]$ be a set of indices. We use $\mu|_I$ to denote the distribution that draws $x \sim \mu$, and returns the restricted string $x|_I \in \prod_{i \in I} \Omega_i$.

Note that the correct way to parse the overloaded notation $\mu|_I^B$ is “ $(\mu|_B)|_I$ ”, that is, we first apply the condition and then restrict the indices.

Definition 4.4 (Common statistical divergence measures). Let μ and τ be two distributions over a finite set Ω . We use two well known divergence measures, namely the total-variation distance $d_{\text{TV}}(\mu, \tau) = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \tau(x)| = \max_{E \subseteq \Omega} |\mu(E) - \tau(E)|$ and the KL-divergence $D_{\text{KL}}(\mu, \tau) = \mathbb{E}_{x \sim \mu} \left[\log_2 \frac{\mu(x)}{\tau(x)} \right]$.

Here we define the query model for distributions over binary strings. The definitions for general alphabets appear in Section 8.

Definition 4.5 (Subcube oracle access). Let μ be an unknown distribution over $\{0, 1\}^n$. The *subcube oracle* has a set $I \subseteq [n]$ and a string $w \in \{0, 1\}^I$ as input, and its output distributes as $\mu|^{x_I=w}$. For technical reasons, if $\Pr_\mu[x_I = w] = 0$, then the oracle indicates an error. Note that the answers of the oracle are fully independent of the answers that were given to previous queries.

We note that the “error behavior” in the above definition does not really affect our results, since all our upper bounds use only restrictions to guaranteed positive probability outcomes, while our lower bounds use distributions that have no zero-probability elements.

Definition 4.6 (Prefix oracle access). Let μ be an unknown distribution over $\{0, 1\}^n$. The *prefix oracle* is a restricted case of the subcube oracle, where $I = [k]$ for some $0 \leq k \leq n - 1$.

Definition 4.7 (Marginal subcube oracle access). Let μ be an unknown distribution over $\{0, 1\}^n$. The *marginal subcube oracle* has a set $I \subseteq [n]$, an index $i \in [n] \setminus I$ and a string $w \in \{0, 1\}^I$ as input, and its output is a single bit that distributes as $\mu|_i^{x_I=w}$. For technical reasons, if $\Pr_\mu[x_I = w] = 0$, then the oracle indicates an error.

Definition 4.8 (Marginal prefix oracle access). Let μ be an unknown distribution over $\{0, 1\}^n$. The *marginal prefix oracle* is a restricted case of the marginal subcube oracle, where $I = [i - 1]$.

We now define the interval querying model for which we derive a new bound.

Definition 4.9 (Interval oracle access). Let μ be an unknown distribution over $[N]$. The *interval oracle* has two elements $1 \leq a \leq b \leq N$ as input, and its output distributes as $\mu|_{\{a, \dots, b\}}$.

We next define our properties for distributions over general alphabet strings, and what it means to test for a property.

Definition 4.10 (EQUIVALENCE). A pair of distributions μ, τ over $\prod_{i=1}^n \Omega_i$ belongs to the equivalence property if $\mu = \tau$. The distance of a given pair (μ, τ) from EQUIVALENCE is explicitly defined

as $d_{\text{TV}}(\mu, \tau)$ instead of the natural $\inf_{\nu}(d_{\text{TV}}(\mu, \nu) + d_{\text{TV}}(\tau, \nu))$. The two quantities are easily seen to be identical using the triangle inequality.

Definition 4.11 (PRODUCT). A distribution μ over $\prod_{i=1}^n \Omega_i$ is called a *product distribution* if there exist distributions μ_1, \dots, μ_n over $\Omega_1, \dots, \Omega_n$ respectively, for which $\mu \sim \mu_1 \times \dots \times \mu_n$. We denote by PRODUCT the set of all product distributions over $\prod_{i=1}^n \Omega_i$.

Definition 4.12 (ε -test). For some property \mathcal{P} of distributions (which is a set of distributions) and $\varepsilon > 0$, we say that an algorithm \mathcal{A} is an ε -test for \mathcal{P} if:

- For every input distribution $\mu \in \mathcal{P}$, \mathcal{A} accepts with probability at least $\frac{2}{3}$.
- For every input distribution μ for which $d_{\text{TV}}(\mu, \nu) > \varepsilon$ for all $\nu \in \mathcal{P}$, \mathcal{A} rejects with probability at least $\frac{2}{3}$.

5 Linear algorithm for EQUIVALENCE

To construct our improvement on the algorithm of [BC17], we first define some templates for analyzing and comparing differences between distributions.

Definition 5.1 (Single-bit divergence). We call $d : [0, 1] \times [0, 1] \rightarrow [0, \infty)$ a *single-bit divergence* if:

- For every $p, q \in [0, 1]$, $d(p, q) = 0$ if and only if $p = q$ (*Positivity*).
- For every $p' \leq p \leq q \leq q'$, $d(p, q) \leq d(p', q')$ and $d(q, p) \leq d(q', p')$ (*Monotonicity*).

Three useful single-bit divergences are $\text{TV}(p, q) = |p - q|$, $\text{KL}(p, q) = p \log_2 \frac{p}{q} + (1 - p) \log_2 \frac{1-p}{1-q}$ and $\chi^2(p, q) = \frac{(p-q)^2}{(p+q)(2-(p+q))}$. Note that all three are derived from their probability-theoretic counterparts for $\text{Ber}(p)$ and $\text{Ber}(q)$. We have two motivations to prefer this non-standard form of χ^2 : first, the symmetry matches the idea of two unknown distributions, which is not the case in standard χ^2 -tests, and second, it is bounded by 1, which makes the analysis more similar to total-variation distance than to KL-divergence.

Definition 5.2 (Slice-wise divergence). Let d be a single-bit divergence. For every n and two distributions τ, μ over $\{0, 1\}^n$, the *slice-wise divergence* of τ and μ with respect to d is:

$$\Delta_d(\tau, \mu) = \sum_{i=1}^n \mathbb{E}_{w \sim \tau} \left[d(\tau|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i^{x_{[i-1]}=w_{[i-1]}}(1)) \right]$$

(note that the d -divergence is fed the probabilities of the two single-bit distributions to draw 1).

Recall the $\tilde{O}(n^5/\text{poly}(\varepsilon))$ algorithm of [BC17], which is actually $\tilde{O}(n^2/\varepsilon^2)$ for the binary setting. As a motivation, the paper uses (and has a self-contained proof of) the following bound:

$$d_{\text{TV}}(\tau, \mu) \leq \sum_{i=1}^n \mathbb{E}_{w \sim \tau} \left[d_{\text{TV}}(\tau|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i^{x_{[i-1]}=w_{[i-1]}}(1)) \right] = \Delta_{\text{TV}}(\tau, \mu)$$

The algorithm then distinguishes between $\Delta_{\text{TV}}(\tau, \mu) = 0$ (which always holds if $\tau = \mu$) and $\Delta_{\text{TV}}(\tau, \mu) > \varepsilon$ (which always holds if $d_{\text{TV}}(\tau, \mu) > \varepsilon$).

To improve their test, inside the slice-wise divergence expression we substitute the single-bit total-variation distance with the single-bit χ^2 -distance, and analyze the more convenient $\Delta_{\chi^2}(\tau, \mu)$.

The following is immediate, and in fact holds for every slice-wise divergence:

Observation 5.3. *Let τ, μ be two distributions over $\{0, 1\}^n$. If $\tau = \mu$ then $\Delta_{\chi^2}(\tau, \mu) = 0$.*

The following lemma, which we prove in Subsection 5.1, provides the conversion from the total variation distance to the slice-wise chi-square divergence.

Lemma 5.4. *Let τ, μ be two distributions over $\{0, 1\}^n$. If $\tau \neq \mu$, then $\Delta_{\chi^2}(\tau, \mu) \geq \frac{(d_{\text{TV}}(\tau, \mu))^2}{24 \log(2n/d_{\text{TV}}(\tau, \mu))}$.*

The following lemma, which we prove in Section 7, states that we can distinguish between single bit distributions using linearly many samples with respect to the inverse of their χ^2 -divergence.

Lemma 5.5. *Let $p, q \in [0, 1]$ be two probabilities. Given unconditional sampling access to $\text{Ber}(p)$ and $\text{Ber}(q)$, we can distinguish, with probability $\frac{2}{3}$, between the case where $p = q$ and the case where $\chi^2(p, q) > \varepsilon$, at the cost of $O(1/\varepsilon)$ samples from each of them.*

We use a common variant of the Levin's work balance method:

Lemma 5.6 ([Lev85], optimization exercise in [Gol17]). *Let X be a non-negative random variable that is bounded by 1. Assume that there exists some random variable Y such that for every $y \in \text{supp}(Y)$ and every $\varepsilon' > 0$, we can distinguish between $\mathbb{E}[X|Y = y] = 0$ and $\mathbb{E}[X|Y = y] > \varepsilon'$ using some black-box algorithm whose resource cost is $O(1/\varepsilon')$. Also, assume that we can draw independent unconditional samples from Y at resource cost $O(1)$ per sample. Then we can distinguish between $\mathbb{E}[X] = 0$ and $\mathbb{E}[X] > \varepsilon$ at a total resource cost of $O(\varepsilon^{-1} \log^2(1/\varepsilon))$.*

Using the above we can efficiently detect a large slicewise χ^2 divergence between two distributions.

Lemma 5.7. *Let τ, μ be two distributions over $\{0, 1\}^n$. Then for every $\rho > 0$, we can distinguish between $\Delta_{\chi^2}(\tau, \mu) = 0$ and $\Delta_{\chi^2}(\tau, \mu) > \rho$ using $O(\frac{n}{\rho} \log^2 \frac{n}{\rho})$ prefix queries.*

Proof. We normalize the divergence:

$$\begin{aligned} \frac{1}{n} \Delta_{\chi^2}(\tau, \mu) &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{w \sim \tau} \left[\chi^2(\tau|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i^{x_{[i-1]}=w_{[i-1]}}(1)) \right] \\ &= \mathbb{E}_{\substack{w \sim \tau \\ i \sim [n]}} \left[\chi^2(\tau|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i^{x_{[i-1]}=w_{[i-1]}}(1)) \right] \end{aligned}$$

Then we apply Levin's method, Lemma 5.6, with the following input. Y is a random variable that receives a value (i, w) where i is uniformly drawn from $[n]$ and w is drawn from τ (independently of i). The parameter ε is set to ρ/n , and the random variable X is defined as a function of $Y = (i, w)$ (meaning that it is constant when conditioned on a specific value of Y) by

$$X(i, w) = \chi^2 \left(\mu|_i^{x_{[i-1]}=w_{[i-1]}}(1), \tau|_i^{x_{[i-1]}=w_{[i-1]}}(1) \right)$$

The distinction between the cases $X(Y) = 0$ and $X(Y) > \varepsilon'$ is performed using $O(1/\varepsilon')$ many queries (our resource cost) through the single-bit χ^2 -test of Lemma 5.5. \square

Our main theorem follows immediately from the above statements:

Theorem 5.8. *Let τ, μ be two distributions over $\{0, 1\}^n$, where τ is accessible through the prefix oracle access and μ is accessible through the marginal prefix oracle access. For every $\varepsilon > 0$ we can distinguish between $\tau = \mu$ and $d_{\text{TV}}(\tau, \mu) > \varepsilon$ using $\tilde{O}(n/\varepsilon^2)$ queries.*

Proof. Apply Lemma 5.7 using $\rho = \frac{\varepsilon^2}{24 \log(2n/\varepsilon)}$. Completeness follows from Observation 5.3 and soundness follows from Lemma 5.4. \square

Corollary 5.9. *Let τ, μ be two distributions over $\{1, \dots, N\}$, both accessible through the interval oracle. Then we can distinguish between $\tau = \mu$ and $d_{\text{TV}}(\tau, \mu) > \varepsilon$ using $\tilde{O}((\log N)/\varepsilon^2)$ queries.*

Proof. Without loss of generality, assume that $N = 2^\ell$ for some integer ℓ (otherwise we just pad the two distributions with zero-probability elements). For every $i \geq 0$, let $\text{bin}_i : \{0, \dots, 2^i - 1\} \rightarrow \{0, 1\}^i$ and $\text{unbin}_i : \{0, 1\}^i \rightarrow \{0, \dots, 2^i - 1\}$ be the mappings between small integers and their representation as i -bit strings. That is, $\text{unbin}_i(x_1, \dots, x_i) = \sum_{j=1}^i 2^{i-j} x_j$ and $\text{bin}_i = (\text{unbin}_i)^{-1}$.

To apply our equivalence algorithm on distributions over $[N]$, every $t \in [N]$ is interpreted as an ℓ -bit string (using the map $t \rightarrow \text{bin}_\ell(t-1)$). Then every prefix query $\tau|^{x_{[i-1]}=w}$ (respectively $\mu|^{x_{[i-1]}=w}$) is simulated using the interval query $\tau|^{a, \dots, b}$ (respectively $\mu|^{a, \dots, b}$), where $a = 2^{N-i} \text{unbin}_{i-1}(w) + 1$ and $b = 2^{N-i} (\text{unbin}_{i-1}(w) + 1)$. Since the simulated algorithm uses $\tilde{O}(\ell/\varepsilon^2)$ prefix queries, the simulation uses $\tilde{O}((\log N)/\varepsilon^2)$ interval queries. \square

5.1 Proof of Lemma 5.4

The proof works by comparing the TV-distance and the slice-wise chi-square divergence with the KL-divergence, which is equal to its slice-wise version as per the following lemmas.

Lemma 5.10 (Basic chain rule for D_{KL} , folklore). *Let μ and τ be distributions over Ω . Then for every random variable $X : \Omega \rightarrow R$, we have the equality $D_{\text{KL}}(\mu, \tau) = D_{\text{KL}}(X(\mu), X(\tau)) + \mathbb{E}_{x \sim X(\tau)} [D_{\text{KL}}(\mu|^{X=x}, \tau|^{X=x})]$.*

Lemma 5.11 (Chain rule for D_{KL} , repeated form). *Let μ and τ be distributions over $\{0, 1\}^n$. Then $D_{\text{KL}}(\mu, \tau) = \sum_{i=1}^n \mathbb{E}_{w \sim \mu} [D_{\text{KL}}(\mu|_i^{x_{[i-1]}=w_{[i-1]}}, \tau|_i^{x_{[i-1]}=w_{[i-1]}})] = \Delta_{\text{KL}}(\mu, \tau)$.*

Proof. We obtain this result by repeatedly applying the basic chain rule using induction. \square

The following technical lemma, proved in Section 7, provides some connection between the chi-square distance and the KL-divergence of Bernoulli distributions.

Lemma 5.12. *Let $p, q \in [0, 1]$. Then $\chi^2(p, q) \geq \frac{1}{12} \text{KL}(p, q) / \log \max\{\frac{1}{q}, \frac{1}{1-q}\}$.*

At this point we prove a version of Lemma 5.4 for a parameterized restricted case. We then reduce the general case to this lemma using an appropriate parameter. We note here that if we were to use the following parameterized lemma directly instead of Lemma 5.4, we would have obtained a direct generalization of the binary setting part of [BCŠV23, Theorem 4.1].

Lemma 5.13. *Let $0 < q < \frac{1}{2}$. Assume that for every $1 \leq i \leq n$ and for every condition $w \in \{0, 1\}^{i-1}$, the prefix marginal $\mu|_i^{x_{[i-1]}=w}$ is equivalent to $\tau|_i^{x_{[i-1]}=w}$ or equivalent to $\text{Ber}(p_{i,w})$ for some $q \leq p_{i,w} \leq 1 - q$ (or both). Then $\Delta_{\chi^2}(\tau, \mu) \geq \frac{1}{6}(d_{\text{TV}}(\tau, \mu))^2 / \log q^{-1}$.*

Proof. By Pinsker's inequality and Lemma 5.11 we obtain:

$$2(d_{\text{TV}}(\tau, \mu))^2 \leq D_{\text{KL}}(\tau, \mu) = \sum_{i=1}^n \mathbb{E}_{w \sim \tau} \left[\text{KL} \left(\tau|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i^{x_{[i-1]}=w_{[i-1]}}(1) \right) \right]$$

By our assumption, for every $1 \leq i \leq n$, if $D_{\text{KL}}(\tau|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i^{x_{[i-1]}=w_{[i-1]}}(1)) \neq 0$ then the probability of $\mu|_i^{x_{[i-1]}=w_{[i-1]}}$ to draw 1 is between q and $1 - q$. By Lemma 5.12 we obtain:

$$2(d_{\text{TV}}(\tau, \mu))^2 \leq \sum_{i=1}^n \mathbb{E}_{w \sim \tau} \left[\chi^2 \left(\tau|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i^{x_{[i-1]}=w_{[i-1]}}(1) \right) \cdot 12 \log q^{-1} \right]$$

That is,

$$\frac{(d_{\text{TV}}(\tau, \mu))^2}{6 \log q^{-1}} \leq \sum_{i=1}^n \mathbb{E}_{w \sim \tau} \left[\chi^2 \left(\tau|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i^{x_{[i-1]}=w_{[i-1]}}(1) \right) \right] = \Delta_{\chi^2}(\tau, \mu) \quad \square$$

We could bound the KL-divergence of single bits using the χ^2 distance only because we assumed that the marginal probabilities of μ are not too close to 0 or 1 (unless they are equal to their counterparts in τ). In the general case we cannot assume it, hence we need to instead consider a distribution μ' which is close to μ while satisfying this assumption.

We consider μ as a ‘‘probability tree’’, where the root represents the empty string, every edge represents an additional bit, and every leaf represents a complete sample. This tree (and hence the distribution μ) is fully determined using probabilities of the form $\Pr_{x \sim \mu}[x_i = 1 | x_{[i-1]} = w]$, where $1 \leq i \leq n$ and $w \in \{0, 1\}^{i-1}$.

We construct another distribution μ' based on such a tree pattern. For every $1 \leq i \leq n$ and $w \in \{0, 1\}^{i-1}$, we set $\Pr_{x \sim \mu'}[x_i = 1 | x_{[i-1]} = w]$ as follows:

$$\begin{aligned} & \min \left\{ \frac{d_{\text{TV}}(\tau, \mu)}{2n}, \Pr_{x \sim \tau}[x_i = 1 | x_{[i-1]} = w] \right\} & \text{if } \Pr_{x \sim \mu}[x_i = 1 | x_{[i-1]} = w] < \frac{d_{\text{TV}}(\tau, \mu)}{2n} \\ 1 - \min \left\{ \frac{d_{\text{TV}}(\tau, \mu)}{2n}, \Pr_{x \sim \tau}[x_i = 0 | x_{[i-1]} = w] \right\} & \text{if } \Pr_{x \sim \mu}[x_i = 1 | x_{[i-1]} = w] > 1 - \frac{d_{\text{TV}}(\tau, \mu)}{2n} \\ & \Pr_{x \sim \mu}[x_i = 1 | x_{[i-1]} = w] & \text{otherwise} \end{aligned}$$

Observe that for every $1 \leq i \leq n$ and $w \in \{0, 1\}^{i-1}$, $d_{\text{TV}} \left(\mu|_i^{x_{[i-1]}=w}, \mu'|_i^{x_{[i-1]}=w} \right) \leq \frac{d_{\text{TV}}(\tau, \mu)}{2n}$. Hence, $d_{\text{TV}}(\mu, \mu') \leq \Delta_{\text{TV}}(\mu, \mu') \leq \frac{1}{2}d_{\text{TV}}(\tau, \mu)$. By the triangle inequality, $d_{\text{TV}}(\tau, \mu') \geq \frac{1}{2}d_{\text{TV}}(\tau, \mu)$.

Since the assumptions of Lemma 5.13 hold for μ' (with $q = \frac{1}{2n}d_{\text{TV}}(\tau, \mu)$), we can now conclude the proof of Lemma 5.4:

$$\Delta_{\chi^2}(\tau, \mu) = \sum_{i=1}^n \mathbb{E}_{w \sim \tau} \left[\chi^2 \left(\tau|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i^{x_{[i-1]}=w_{[i-1]}}(1) \right) \right]$$

$$\begin{aligned}
[\text{Monotonicity of } \chi^2] &\geq \sum_{i=1}^n \mathbb{E}_{w \sim \tau} \left[\chi^2 \left(\tau|_i^{x^{[i-1]}=w^{[i-1]}}(1), \mu'|_i^{x^{[i-1]}=w^{[i-1]}}(1) \right) \right] \\
[\text{Lemma 5.13 with } q = \frac{1}{2n} d_{\text{TV}}(\tau, \mu)] &\geq (d_{\text{TV}}(\tau, \mu'))^2 / 6 \log \frac{2n}{d_{\text{TV}}(\tau, \mu)} \\
[d_{\text{TV}}(\tau, \mu') \geq \frac{1}{2} d_{\text{TV}}(\tau, \mu)] &\geq (d_{\text{TV}}(\tau, \mu))^2 / 24 \log \frac{2n}{d_{\text{TV}}(\tau, \mu)} \quad \square
\end{aligned}$$

6 Lower bound for PRODUCT

This section is devoted to the following lower bound:

Theorem 6.1. *Every ε -test for PRODUCT must make at least $\Omega(\sqrt{n}/\varepsilon^2)$ subcube queries.*

Let π_n denote the uniform distribution over $\{0, 1\}^n$. We show that distinguishing between π_n (which is in particular a product distribution) and a distribution that is ε -far from every product distribution requires $\tilde{\Omega}(\sqrt{n}/\varepsilon^2)$ many queries.

Before we present our construction, we cite the corresponding lower bound for uniformity of a product distribution [CDKS17]:

Lemma 6.2 ([CDKS17]). *Let \mathcal{N} be the following distribution over inputs: draw $b_1, \dots, b_n \sim \{+1, -1\}$ uniformly and independently, and return the distribution $\prod_{i=1}^n \text{Ber}\left(\frac{1}{2} + b_i \frac{\varepsilon}{\sqrt{n}}\right)$. Then the drawn input is always $\Omega(\varepsilon)$ -far from π_n , and any unconditional sampling algorithm that distinguishes between inputs drawn from \mathcal{N} and π_n must take at least $\Omega(\sqrt{n}/\varepsilon^2)$ many samples.*

In our construction, instead of adding a random bias for each coordinate, we partition the coordinates into pairs, and in each pair introduce a random “anti-product bias” as follows.

For $b \in \{0, +1, -1\}$, let ν_b be the following distribution over $\{0, 1\}^2$:

$$\begin{aligned}
\nu_b(00) &= \frac{1}{4} + b_i \frac{\varepsilon}{\sqrt{n}} & \nu_b(01) &= \frac{1}{4} - b_i \frac{\varepsilon}{\sqrt{n}} \\
\nu_b(10) &= \frac{1}{4} - b_i \frac{\varepsilon}{\sqrt{n}} & \nu_b(11) &= \frac{1}{4} + b_i \frac{\varepsilon}{\sqrt{n}}
\end{aligned}$$

That is, ν_0 is the uniform distribution over two bits, and ν_{+1}, ν_{-1} are non-product distributions over two bits.

Let \mathcal{Y} be the distribution over inputs that always returns the uniform distribution over $\{0, 1\}^n$, and let \mathcal{N} be the following distribution over “bad” inputs over $\{0, 1\}^n$: we partition $[n]$ into fixed pairs (for example, $(2i-1, 2i)$ for every $1 \leq i \leq \lfloor n/2 \rfloor$). For every pair we draw $b_i \in \{+1, -1\}$ uniformly and independently. The distribution of the i th pair is ν_{b_i} . All pairs are independent of each other. We assume that n is even, since for an odd n we can just assume that the n th bit is distributed uniformly and independently of the rest.

The following lemma, whose proof appears in Section 10, states that \mathcal{N} always draws a distribution far from PRODUCT.

Lemma 6.3. *Every input distribution drawn from \mathcal{N} is $\Omega(\varepsilon)$ -far from any product distribution.*

We now prove the indistinguishability of \mathcal{N} from the uniform distribution, starting with the sampling model, to which we later show a reduction from the subcube conditional model.

Lemma 6.4. *An unconditional sampling algorithm that distinguishes between the uniform distribution and inputs that are drawn from \mathcal{N} must make $\Omega(\sqrt{n}/\varepsilon^2)$ many queries.*

Proof. Without loss of generality, assume that n is even. Consider the following bijection over $\{0, 1\}^2$: $(x, y) \rightarrow (x \oplus y, y)$. For $b \in \{+1, 0, -1\}$, Let ν'_b be the distribution that draws $(x, y) \sim \nu_b$ and returns $(x \oplus y, y)$. Observe that ν'_b is identical to the product distribution $\text{Ber}(\frac{1}{2} + \frac{2b\varepsilon}{\sqrt{n}}) \times \text{Ber}(\frac{1}{2})$. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a bijection that applies the above mapping for every pair individually.

Let \mathcal{N}' be the distribution over inputs that draws $\mu \sim \mathcal{N}$, and then returns the distribution that draws $x \sim \mu$ and returns $f(x)$. Note also that the distribution that drawn $x \sim \pi_n$ and returns $f(x)$ is identical to π_n . Since f is a bijection, this means that the number of samples needed to distinguish between the uniform distribution and \mathcal{N} is exactly the same as the number of samples needed to distinguish between the uniform distribution and \mathcal{N}' .

Observe that the form of a distribution resulting from \mathcal{N}' is $\prod_{i=1}^{n/2} \left(\text{Ber} \left(\frac{1}{2} + \frac{2b_i\varepsilon}{\sqrt{n}} \right) \times \text{Ber} \left(\frac{1}{2} \right) \right)$. The restriction of \mathcal{N}' to odd indexes is identical to the construction of Lemma 6.2 (with parameters $\sqrt{2}\varepsilon$ and $n/2$), which requires at least $\Omega(\sqrt{n}/\varepsilon^2)$ many unconditional samples to distinguish between it and the uniform distribution. Hence, this is also a lower bound for the number of samples needed to distinguish between the uniform distribution and inputs drawn from \mathcal{N} . \square

Lemma 6.5. *Let $b \in \{0, +1, -1\}$. Then for every subcube restriction $q \in \{0, 1, *\}^2$, we can simulate a q -conditioned sample from ν_b by drawing a single unconditional sample from ν_b , without having any knowledge about b .*

Proof. Let $q \in \{0, +1, *\}^2$ and let $b \in \{0, +1, -1\}$. We first draw an unconditional sample $x \sim \nu_b$. There are three kinds of subcube restrictions:

- Trivial: $q \in \{00, 01, 10, 11\}$. To simulate such a query, we simply return q , ignoring the unconditional sample we have.
- Unconditional: $q = **$. To simulate such a query, we simply return x .
- Single restriction: $q \in \{0*, 1*, *0, *1\}$. We denote the parity of x by $p = x|_1 \oplus x|_2$, and construct the output as the concatenation of the two bits out_1 and out_2 , based on the following table:

q	0*	1*	*0	*1
out_1	0	1	p	$1 \oplus p$
out_2	p	$1 \oplus p$	0	1

The correctness of the output in the trivial case and the unconditional case is trivial. We prove the correctness of the single-restriction case only for $q = 0*$, since the other cases are analogous.

$$\Pr_{\nu_b}[00|0*] = \frac{\Pr_{\nu_b}[00]}{\Pr_{\nu_b}[0*]} = \frac{\frac{1}{4} + \frac{b_i\varepsilon}{\sqrt{n}}}{\frac{1}{4} + \frac{b_i\varepsilon}{\sqrt{n}} + \frac{1}{4} - \frac{b_i\varepsilon}{\sqrt{n}}} = \frac{1}{2} + 2\frac{b_i\varepsilon}{\sqrt{n}} = \Pr_{\nu_b}[00 \vee 11] = \Pr_{x \sim \nu_b}[x|_1 \oplus x|_2 = 0] \quad \square$$

At this point we can prove Theorem 6.1.

Proof of Theorem 6.1. By Lemma 6.3, \mathcal{N} draws an input distribution that is $\Omega(\varepsilon)$ -far from any product distribution. Observe that since the structure of the pairs is known in advance and since they are independent, we can simulate every q -subcube-query algorithm using a q -unconditional-sample algorithm: for each query we draw a single sample and then use the simulation procedure of Lemma 6.5 for every pair in itself.

Since an unconditional test requires $\Omega(\sqrt{n}/\varepsilon^2)$ queries to distinguish inputs drawn from \mathcal{N} from the uniform distribution by Lemma 6.4, and every subcube query to the uniform distribution or an input drawn from \mathcal{N} can be simulated using a single unconditional query, the lower bound holds for subcube algorithms as well. \square

7 Technical proofs for EQUIVALENCE testing

Here we prove the lemmas deferred from Section 5. We start with some helper lemmas.

Lemma 7.1. *Let $0 \leq p \leq q \leq 1$ for which $p + q \leq 1$ and let X be the sum of N independent bits drawn from $\text{Ber}(p)$. Then,*

$$\Pr[X \geq \frac{1}{2}(p + q)N] \leq e^{-\frac{1}{12}\chi^2(p,q)}$$

Proof. Let $\delta = q - p$. If $\delta \leq 2p$ then by Chernoff's bound:

$$\begin{aligned} \Pr[X \geq \frac{1}{2}(p + q)N] &= \Pr[X \geq (1 + \frac{\delta}{2p})\mathbb{E}[X]] \\ &\leq e^{-\frac{\delta^2}{12p^2}pN} = e^{-\frac{\delta^2}{12p}N} \leq e^{-\frac{(p-q)^2}{12(p+q)(2-(p+q))}N} = e^{-\frac{1}{12}\chi^2(p,q)N} \end{aligned}$$

If $\delta > 2p$ then:

$$\begin{aligned} \Pr[X \geq \frac{1}{2}(p + q)N] &= \Pr[X \geq (1 + \frac{\delta}{2p})\mathbb{E}[X]] \\ &\leq e^{-\frac{\delta}{6p}pN} = e^{-\frac{\delta}{6}N} \leq e^{-\frac{(q-p)^2}{6(p+q)}} \leq e^{-\frac{1}{12}\chi^2(p,q)N} \end{aligned}$$

\square

Lemma 7.2. *Let $0 \leq p \leq q \leq 1$ for which $p + q \leq 1$ and let X be the sum of N independent bits drawn from $\text{Ber}(q)$. Then,*

$$\Pr[X \leq \frac{1}{2}(p + q)N] \leq e^{-\frac{1}{8}\chi^2(p,q)} \leq e^{-\frac{1}{12}\chi^2(p,q)}$$

Proof. Let $\delta = q - p \leq q$. By Chernoff's bound:

$$\begin{aligned} \Pr[X \leq \frac{1}{2}(p + q)N] &= \Pr[X \leq (1 - \frac{\delta}{2q})\mathbb{E}[X]] \\ &\leq e^{-\frac{\delta^2}{8q^2}qN} = e^{-\frac{\delta^2}{8q}N} \leq e^{-\frac{(p-q)^2}{8(p+q)(2-(p+q))}N} = e^{-\frac{1}{8}\chi^2(p,q)N} \end{aligned}$$

\square

Lemma 7.3. *Let $p, q \in [0, 1]$. Let X be the sum of N independent bits drawn from $\text{Ber}(p)$ and Y be the sum of N independent bits drawn from $\text{Ber}(q)$. Then, with probability at least $\left(1 - e^{-\frac{1}{12}\chi^2(p,q)N}\right)^2$, the sign of $X - Y$ matches the sign of $p - q$.*

Proof. Without loss of generality, $p + q \leq 1$ (otherwise use $1 - q$ and $1 - p$ instead of p and q , noting that $d(1 - q, 1 - p) = d(p, q)$ and that the sign of $(1 - q) - (1 - p)$ matches the sign of $p - q$).

If $p = q$ then the lemma is vacuously correct. If $p < q$, then by Lemma 7.1,

$$\begin{aligned} \Pr[X - Y < 0] = \Pr[X < Y] &\geq \Pr[X < \frac{1}{2}(p + q)N] \Pr[Y > \frac{1}{2}(p + q)N] \\ &\geq \left(1 - e^{-\frac{1}{12}\chi^2(p,q)N}\right)^2 \end{aligned}$$

The case where $p > q$ is analogous, using Lemma 7.2. □

We recall and prove Lemma 5.5:

Lemma 5.5. *Let $p, q \in [0, 1]$ be two probabilities. Given unconditional sampling access to $\text{Ber}(p)$ and $\text{Ber}(q)$, we can distinguish, with probability $\frac{2}{3}$, between the case where $p = q$ and the case where $\chi^2(p, q) > \varepsilon$, at the cost of $O(1/\varepsilon)$ samples from each of them.*

Proof. We repeat the following procedure 64 times: let $N = \lceil 24/\varepsilon \rceil$. Also, let X be the sum of N independent samples drawn from $\text{Ber}(p)$ and Y be the sum of N independent samples drawn from $\text{Ber}(q)$. To conclude a single trial we check whether $A > B$ or $A < B$ (or neither) holds.

Let A be the number of trials with $X > Y$ and B be the number of trials with $X < Y$. If $|A|, |B| \leq 40$ we accept ($p = q$), and otherwise we reject ($\chi^2(p, q) > \varepsilon$).

If $p = q$, then by symmetry, $\mathbb{E}[X < Y] = \mathbb{E}[X > Y] \leq \frac{1}{2}$. That is, $\mathbb{E}[A] = \mathbb{E}[B] \leq 32$. By Chernoff's bound, $\Pr[A \geq 41] < e^{-2 \cdot 9^2/64} < \frac{1}{6}$ and $\Pr[B \geq 41] < \frac{1}{6}$. Hence, the probability to accept is at least $1 - \frac{2}{6} = \frac{2}{3}$.

If $\chi^2(p, q) > \varepsilon$, then by Lemma 7.3, one of $\Pr[X < Y]$ and $\Pr[X > Y]$ is at least $(1 - e^{-\frac{1}{12}\chi^2(p,q)N})^2 \geq (1 - e^{-2})^2 > 0.74$. Without loss of generality, we assume that $p < q$. In this case, $\mathbb{E}[A] > 47.36$. By the Chernoff bound, the probability to reject is at least $1 - \Pr[A \leq 40] \geq 1 - e^{-2 \cdot 7.36^2/64} > \frac{2}{3}$. □

Lemma 7.4 (Well known). *For $p, q \in [0, 1]$, $\text{KL}(p, q) \leq \frac{(p-q)^2}{q(1-q)}$. More formally, in $(0, 1) \times (0, 1)$, the ratio between these expressions is a non-negative continuous function that is bounded by 1.*

Lemma 7.5 (Direct corollary). *Let $p, q \in [0, 1]$. If $p = aq$ for some real a , then*

$$\text{KL}(aq, q) \leq \frac{(a-1)^2}{1-q} q$$

Finally, we recall and prove Lemma 5.12:

Lemma 5.12. *Let $p, q \in [0, 1]$. Then $\chi^2(p, q) \geq \frac{1}{12} \text{KL}(p, q) / \log \max\{\frac{1}{q}, \frac{1}{1-q}\}$.*

Proof. We actually prove that $\chi^2(p, q)/\text{KL}(p, q) \geq \frac{1}{12 \log \max\{\frac{1}{q}, \frac{1}{1-q}\}}$ for every $p \in [0, 1]$ and $q \in (0, 1)$. For the edge cases of q (which we do not use in our proofs anyway except when $p = q$), the bound remains correct by considering the limit of $\text{KL}(p, q)/\log \max\{\frac{1}{q}, \frac{1}{1-q}\}$.

Without loss of generality, $q \leq \frac{1}{2}$. We can assume so since $\chi^2(p, q) = \chi^2(1-p, 1-q)$, $\text{KL}(p, q) = \text{KL}(1-p, 1-q)$ and $\max\{\log 1/q, \log 1/(1-q)\} = \max\{\log 1/(1-q), \log 1/(1-(1-q))\}$. Based on this assumption it is sufficient to show that $\frac{\chi^2(p, q)}{\text{KL}(p, q)} \geq \frac{1}{12 \log q^{-1}}$.

Let $a = p/q$ ($0 < a \leq 1/q$). If $a \geq 2$:

$$\begin{aligned} \chi^2(aq, q) &\geq \frac{(a-1)^2}{2(a+1)}q \\ \text{KL}(aq, q) &\leq p \log \frac{p}{q} \leq a \log q^{-1} \cdot q \\ \frac{\chi^2(aq, q)}{\text{KL}(aq, q)} &\geq \frac{(a-1)^2}{2a(a+1) \log q^{-1}} \geq \frac{1}{12 \log q^{-1}} \end{aligned}$$

If $a \leq 2$:

$$\begin{aligned} \chi^2(aq, q) &\geq \frac{(a-1)^2}{2(a+1)}q \\ \text{KL}(aq, q) &\leq \frac{(a-1)^2}{1-q}q \\ \frac{\chi^2(aq, q)}{\text{KL}(aq, q)} &\geq \frac{1-q}{2(a+1)} \geq \frac{1}{12 \log q^{-1}} \end{aligned}$$

where the very last inequality uses the assumption that $q \leq \frac{1}{2}$. □

8 Extending the EQUIVALENCE test to general alphabets

We extend the definitions of the prefix oracle to non-binary settings.

Definition 8.1 (Subcube oracle access in non-binary strings). Let μ be an unknown distribution over $\prod_{i=1}^n \Omega_n$, where $\Omega_1, \dots, \Omega_n$ are all finite. The *subcube oracle* has as input a tuple (A_1, \dots, A_n) where $A_i \subseteq \Omega_i$ for every $1 \leq i \leq n$. The output distributes as $\mu|_{\prod_{i=1}^n A_i}$. For technical reasons, if $\Pr_\mu[\prod_{i=1}^n A_i] = 0$, then the oracle indicates an error. Note that the answers of the oracle are fully independent of the answers that were given to previous queries.

In the corresponding definition for a prefix oracle, we still demand that until the “break-off index” i all conditions force single outcomes from the sets Ω_j , while after the break-off index there are no restrictions at all. However, at index i we allow conditions to any subset of Ω_i to take place. There is no such distinction in the binary case, where $|\Omega_i| = 2$ and hence all non-trivial conditions are to a single outcome.

Definition 8.2 (Prefix oracle access in non-binary settings). Let μ be an unknown distribution over $\prod_{i=1}^n \Omega_n$, where $\Omega_1, \dots, \Omega_n$ are all finite. The input of the *prefix oracle* consists of an index $1 \leq i \leq n$, which we refer to as the *index of the prefix*, elements $a_j \in \Omega_j$ for every $1 \leq j \leq i-1$, and a condition $A \subseteq \Omega_i$. The output of the oracle distributes like $\mu|_{\{x: x_i \in A \wedge x_1 = a_1, \dots, x_{i-1} = a_{i-1}\}}$.

Definition 8.3 (Marginal subcube oracle access in non-binary settings). Let μ be an unknown distribution over $\prod_{i=1}^n \Omega_n$, where $\Omega_1, \dots, \Omega_n$ are all finite. The *marginal subcube oracle* has as input an index i and a tuple (A_1, \dots, A_n) where $A_j \subseteq \Omega_j$ for every $1 \leq j \leq n$. The output distributes as $\mu|_i^{\prod_{j=1}^n A_j}$. For technical reasons, if $\Pr_\mu[\prod_{i=1}^n A_i] = 0$, then the oracle indicates an error.

Definition 8.4 (Marginal prefix oracle access in non-binary settings). Let μ be an unknown distribution over $\prod_{i=1}^n \Omega_n$, where $\Omega_1, \dots, \Omega_n$ are all finite. The input of the *marginal prefix oracle* consists of an index $1 \leq i \leq n$, which we refer to as the *index of the prefix*, elements $a_j \in \Omega_j$ for every $1 \leq j \leq i-1$, and a condition $A \subseteq \Omega_i$. The output of the oracle distributes like $\mu|_i^{\{x: x_i \in A \wedge x_1 = a_1, \dots, x_{i-1} = a_{i-1}\}}$.

Based on these definitions, we state Theorem 8.5:

Theorem 8.5. *Let μ and τ be two distributions over $\prod_{i=1}^n \Omega_i$, where $\Omega_1, \dots, \Omega_n$ are all finite. If μ is accessible through the marginal prefix oracle and τ is accessible through the prefix oracle, then we can distinguish between $\tau = \mu$ and $d_{TV}(\tau, \mu) > \varepsilon$ using $\tilde{O}(\sum_{i=1}^n \log_2 |\Omega_i| / \varepsilon^2)$ queries.*

Before we prove Theorem 8.5, we need the following.

Lemma 8.6 (Binary form of a composite distribution). *Let μ be a distribution over $\prod_{i=1}^n \Omega_i$, where $\Omega_1, \dots, \Omega_n$ are non-empty finite sets. There exists a distribution μ^* over $\{0, 1\}^{\sum_{i=1}^n \lceil \log_2 |\Omega_i| \rceil}$ that is equivalent to μ up to relabeling, for which every subcube (respectively prefix) query to μ^* can be simulated using a single subcube (respectively prefix) query to μ , and every marginal subcube (respectively prefix) query to μ^* can be simulated using a single marginal subcube (respectively prefix) query to μ .*

Proof. For every $1 \leq i \leq n$, let $f_i : \Omega_i \rightarrow \{0, 1\}^{\lceil \log_2 |\Omega_i| \rceil}$ be an arbitrary injective function from Ω_i to $\{0, 1\}^{\lceil \log_2 |\Omega_i| \rceil}$. Let $f : \prod_{i=1}^n \Omega_i \rightarrow \{0, 1\}^{\sum_{i=1}^n \lceil \log_2 |\Omega_i| \rceil}$ be the concatenation of these mappings. More precisely, $f((x_1, \dots, x_n)) = \text{concatenate}(f_1(x_1), \dots, f_n(x_n))$. Let $N = \sum_{i=1}^n \lceil \log_2 |\Omega_i| \rceil$.

For every $1 \leq i \leq n$, we define a projection function $g_i : \{0, 1\}^{\sum_{j=1}^n \lceil \log_2 |\Omega_j| \rceil} \rightarrow \Omega_i$ by

$$g_i(x) = (f_i)^{-1} \left(x|_{\{\sum_{j=1}^{i-1} \lceil \log_2 |\Omega_j| \rceil + 1, \dots, \sum_{j=1}^i \lceil \log_2 |\Omega_j| \rceil\}} \right)$$

where $g_i(x)$ is defined arbitrarily if the corresponding binary string is not in f_i 's image, as this will be a zero-probability event. Observe that for every $(x_1, \dots, x_n) \in \prod_{i=1}^n \Omega_i$ and for every $1 \leq i \leq n$, $x_i = g_i(f(x_1, \dots, x_n))$.

Let μ^* be the distribution over $\{0, 1\}^{\sum_{i=1}^n \lceil \log_2 |\Omega_i| \rceil}$ that draws $x \sim \mu$ and returns $f(x)$. Since f is an injective function, μ^* is equivalent to μ up to relabeling.

For simulating subcube queries consider some $I^* \subseteq [N]$ and $w^* \in \{0, 1\}^{I^*}$. For every $1 \leq i \leq n$ and $x^* \in \{0, 1\}^N$, let $x_i = g_i(x^*)$ (that is, we decompose x^* using $x^* = f(x_1, \dots, x_n)$). Also, for every $1 \leq i \leq n$, let $A_i(I^*, w^*) = \{g_i(s^*) : s^* \in \{0, 1\}^N \wedge s^*|_{I^*} = w^*\}$. Based on this composition, we obtain:

$$\{x^* \in \{0, 1\}^N : x^*|_{I^*} = w^*\} = \left\{ f(x_1, \dots, x_n) : \bigwedge_{i=1}^n (x_i \in A_i(I^*, w^*)) \right\}$$

The last expression is the f -image of all elements in the μ -subcube condition $\prod_{i=1}^n A_i(I^*, w^*)$. Hence, every subcube query of μ^* can be simulated using a single subcube query to μ .

Observe that this construction preserves prefix queries. That is, if a subcube query to μ^* is a prefix query, then the simulated subcube query to μ is a prefix query as well. Note that this argument holds for marginal queries as well, since we can extract the relevant bit of the sampled coordinate. \square

Theorem 8.5 now follows.

Proof of Theorem 8.5. Let μ and τ be two distributions over $\prod_{i=1}^n \Omega_i$, where $\Omega_1, \dots, \Omega_n$ are all finite. We use Lemma 8.6 to define two distributions μ^*, τ^* that are identical to μ, τ respectively up to relabeling (which is the same in both constructions). Based on this lemma:

- $d_{\text{TV}}(\tau^*, \mu^*) = d_{\text{TV}}(\tau, \mu)$, since μ^*, τ^* are the same as μ, τ up to relabeling (which is the same for both constructions $\mu \rightarrow \mu^*$ and $\tau \rightarrow \tau^*$).
- Every prefix query to τ^* can be simulated using a single prefix query to τ .
- Every marginal prefix query to μ^* can be simulated using a single marginal prefix query to μ .

Hence, we can distinguish between $\tau = \mu$ and $d_{\text{TV}}(\tau, \mu) > \varepsilon$ using Theorem 5.8 with the input $(\mu^*, \tau^*, \varepsilon)$ by simulating every query to τ^* or μ^* through a single query to the corresponding input distribution τ or μ . \square

9 Upper bound for PRODUCT

We reduce a test for PRODUCT to a test for EQUIVALENCE, based on the following key observation:

Observation 9.1. *Let μ be a distribution over $\prod_{i=1}^n \Omega_i$. If μ is ε -far from any product distribution, then in particular it is ε -far from the product of its marginals, $\prod_{i=1}^n \mu|_i$.*

In the binary setting, simulating a marginal prefix query to the product of marginals is pretty straightforward and can be done using one unconditional query to μ , which is in particular a prefix query:

Observation 9.2. *Let μ be a distribution over $\{0, 1\}^n$, and let $\mu' = \prod_{i=1}^n \mu|_i$ be the product of μ 's marginals. Then we can simulate every marginal prefix query to μ' using one unconditional sample from μ (and keeping its i th entry).*

In the general setting, we need the stronger subcube access. The reason is that when taking a prefix marginal query at index i from the binary representation of $\mu' = \prod_{j=1}^n \mu|_j$, it may be that that i is “in the middle” of the $\lceil \log |\Omega_j| \rceil$ -bit representation of j th coordinate of μ' , and then this query must translate to a non-trivial set $A_j \subseteq \Omega_j$ when simulating it using query access to μ .

Observation 9.3. *Let μ be a distribution over $\prod_{i=1}^n \Omega_i$, and let $\mu' = \prod_{i=1}^n \mu|_i$ be the product of μ 's marginals. Then we can simulate every marginal prefix query to μ' using one subcube query to μ .*

At this point, we can prove Theorem 9.4.

Theorem 9.4. *Let μ be a distribution over $\prod_{i=1}^n \Omega_i$. For every $0 < \varepsilon < 1$, we can distinguish between the case where μ is a product distribution and the case where it is ε -far from every product*

distribution at the cost of $\tilde{O}(\sum_{i=1}^n \log |\Omega_i|/\varepsilon^2)$ subcube queries. Moreover, if $|\Omega_i| = 2$ for every $1 \leq i \leq n$, then all these queries are prefix queries.

Proof. Let $\mu' = \prod_{i=1}^n \mu|_i$ be the product of μ 's marginals. If μ is a product distribution then $\mu' = \mu$, and if μ is ε -far from every product distribution, then in particular $d_{\text{TV}}(\mu, \mu') > \varepsilon$.

Since μ is accessible through the subcube oracle, we can use Observation 9.3 to simulate every marginal prefix query to μ' at the cost of one subcube query to μ . If $|\Omega_i| = 2$ for every $1 \leq i \leq n$, then we can use an unconditional sample instead of a subcube query by Observation 9.2.

Hence, we can reduce the ε -test of μ for PRODUCT to an ε -test of the equivalence of μ' and μ , which we perform using Theorem 8.5. As noted above, this produces subcube queries for general alphabets, and only prefix queries for the binary setting. \square

10 Technical proofs for the PRODUCT lower bound

Recall that we denote the uniform distribution over $\{0, 1\}^n$ by π_n .

Before we prove Lemma 6.3, we need the following technical lemmas.

Lemma 10.1. *For a string $y \in \{0, 1\}^n$, let $C_1(y)$ be the number of 1s in y . If $n \geq 70$, then $\Pr_{y \sim \pi_n}[C_1(y) > \frac{1}{2}n + \frac{1}{4}\sqrt{n}] \geq \frac{1}{4}$.*

Proof. We use the well-known bound $\binom{n}{\lfloor n/2 \rfloor} \leq 2^n \cdot \sqrt{\frac{2}{\pi n}}$ (for all $n \geq 1$) to obtain for $n \geq 70$:

$$\begin{aligned} \Pr_{y \sim \pi_n} \left[\left| C_1(y) - \frac{1}{2}n \right| \leq \frac{1}{4}\sqrt{n} \right] &= \sum_{k=\lceil \frac{1}{2}n - \frac{1}{4}\sqrt{n} \rceil}^{\lfloor \frac{1}{2}n + \frac{1}{4}\sqrt{n} \rfloor} \Pr_{y \sim \pi_n} [C_1(y) = k] \\ [\text{Since } \binom{n}{k} &\leq \binom{n}{\lfloor n/2 \rfloor}] &\leq \left(2 \cdot \frac{1}{4}\sqrt{n} + 1 \right) 2^{-n} \binom{n}{\lfloor n/2 \rfloor} \\ &\leq \left(\frac{1}{2}\sqrt{n} + 1 \right) \frac{\sqrt{2}}{\sqrt{\pi n}} \leq \frac{1}{2} \end{aligned}$$

By symmetry reasons, if $n \geq 70$ then:

$$\begin{aligned} \Pr_{y \sim \pi_n} \left[C_1(y) > \frac{1}{2}n + \frac{1}{4}\sqrt{n} \right] &= \frac{1}{2} \Pr_{y \sim \pi_n} \left[\left| C_1(y) - \frac{1}{2}n \right| > \frac{1}{4}\sqrt{n} \right] \\ &= \frac{1}{2} \left(1 - \Pr_{y \sim \pi_n} \left[\left| C_1(y) - \frac{1}{2}n \right| \leq \frac{1}{4}\sqrt{n} \right] \right) \geq \frac{1}{4} \end{aligned}$$

\square

Lemma 10.2. *Let $0 < \delta < \frac{1}{4\sqrt{n}}$ and τ be a product distribution over $\{0, 1\}^n$. For sufficiently large n , if $|\Pr_{\tau}[x_i = 1] - \frac{1}{2}| > \delta$ for every $1 \leq i \leq n$, then the distance of τ from the uniform distribution over $\{0, 1\}^n$ is at least $\frac{1}{16}\delta\sqrt{n}$.*

Proof. Without loss of generality, we can assume that $\Pr_\tau[x_i = 1] \geq \frac{1}{2}$ for every $1 \leq i \leq n$. Otherwise, we can negate the “wrong” bits while preserving the distance from the uniform distribution. Based on this assumption, we have a product distribution whose probability to draw 1 at any individual index is at least $\frac{1+\delta}{2}$.

Let τ' be the product distribution whose probability to draw 1 at any index is exactly $\frac{1+\delta}{2}$. Observe that the distance of τ' from the uniform distribution is a lower bound of the distance of τ from the uniform distribution.

Let $y \in \{0, 1\}^n$ be a string, and let $C_1(y)$ (respectively $C_0(y)$) be the number of 1s (respectively 0s) in y . If $C_1(y) \geq \frac{1}{2}n + \frac{1}{4}\sqrt{n}$, then:

$$\begin{aligned}
\frac{\Pr_{\tau'}[y]}{\Pr_{\pi_n}[y]} &= (1 + \delta)^{C_1(y)} \cdot (1 - \delta)^{n - C_1(y)} \\
&= (1 + \delta)^{C_1(y) - C_0(y)} \cdot ((1 + \delta)(1 - \delta))^{C_0(y)} \\
&\geq (1 + \delta)^{\sqrt{n}/2} \cdot ((1 + \delta)(1 - \delta))^{\frac{1}{2}n - \frac{1}{4}\sqrt{n}} \\
&= (1 + \delta)^{\sqrt{n}/2} \cdot (1 - \delta^2)^{\frac{1}{2}n - \frac{1}{4}\sqrt{n}} \\
&\geq (1 + \delta)^{\sqrt{n}/2} \cdot (1 - \delta^2)^{\frac{1}{2}n} \\
[(1 + a)^b \geq 1 + ab \text{ for } b \geq 1, |a| < 1] &\geq \left(1 + \frac{1}{2}\delta\sqrt{n}\right) \cdot \left(1 - \frac{1}{2}\delta^2 n\right) \\
[\delta < \frac{1}{4\sqrt{n}}] &\geq 1 + \frac{1}{4}\delta\sqrt{n}
\end{aligned}$$

In particular,

$$\frac{\Pr_{y \sim \tau'} [C_1(y) > \frac{1}{2}n + \frac{1}{4}\sqrt{n}]}{\Pr_{y \sim \pi_n} [C_1(y) > \frac{1}{2}n + \frac{1}{4}\sqrt{n}]} \geq 1 + \frac{1}{4}\delta\sqrt{n}$$

Since $\Pr_{y \sim \pi_n} [C_1(y) > \frac{1}{2}n + \frac{1}{4}\sqrt{n}] \geq \frac{1}{4}$ by Lemma 10.1, we obtain that $d_{\text{TV}}(\tau', \pi_n) \geq \frac{1}{4}\delta\sqrt{n} \cdot \frac{1}{4} = \frac{1}{16}\delta\sqrt{n}$. \square

We recall and prove Lemma 6.3:

Lemma 6.3. *Every input distribution drawn from \mathcal{N} is $\Omega(\varepsilon)$ -far from any product distribution.*

Proof. Without loss of generality, we assume that n is even. Let $b_1, \dots, b_{n/2} \in \{+1, -1\}$ be the parameters of the construction, that is, the drawn input is $\mu = \prod_{i=1}^{n/2} \nu_{b_i}$. Recall that for $b \in \{+1, -1\}$, $\Pr_{\nu_b}[x_1 = 1] = \Pr_{\nu_b}[x_2 = 1] = \frac{1}{2}$ and $\Pr_{\nu_b}[x_1 \oplus x_2 = 1] = \frac{1}{2} - \frac{2b\varepsilon}{\sqrt{n}}$. We assume that n is sufficiently large so that $\lfloor n/6 \rfloor$ satisfies the constraints of Lemma 10.2.

For a given product distribution τ , for every $1 \leq i \leq n/2$, let:

$$\begin{aligned}
\delta_{i,0} &= \Pr_\tau[x_{2i} = 1] - \frac{1}{2} \\
\delta_{i,1} &= \Pr_\tau[x_{2i-1} = 1] - \frac{1}{2} \\
\delta_{i,2} &= \Pr_\tau[x_{2i} \oplus x_{2i-1} = 1] - \frac{1}{2}
\end{aligned}$$

Let I_0 be the set of indexes for which $|\delta_{i,0}| > \frac{\varepsilon}{10\sqrt{n}}$. Let I_1 be defined analogously for $\delta_{i,1}$, and let $I_2 = [n/2] \setminus (I_0 \cup I_1)$ be the set of all other indexes. Since $|I_0| + |I_1| + |I_2| = n/2$, at least one of them contains at least $\frac{1}{6}n$ elements.

Case I. $|I_0| \geq \frac{1}{6}n$ Let $I = \{2i : i \in I_0\}$. Observe that $\mu|_I$ is uniform over $\{0, 1\}^{n/2}$, since μ is the product of $n/2$ independent distributions over pairs whose marginals are $\frac{1}{2}$, and every pair contributes exactly one index. According to Lemma 10.2,

$$d_{\text{TV}}(\tau, \mu) \geq d_{\text{TV}}(\tau|_I, \mu|_I) \geq \frac{1}{16} \cdot \frac{\varepsilon}{10\sqrt{n}} \cdot \sqrt{\frac{1}{6}n} = \frac{1}{160\sqrt{6}}\varepsilon > \frac{1}{400}\varepsilon$$

Case II. $|I_1| \geq \frac{1}{6}n$ Completely analogous to Case I. Again $d_{\text{TV}}(\tau, \mu) > \frac{1}{400}\varepsilon$.

Case III. $|I_2| \geq \frac{1}{6}n$ For every $i \in I_2$, $|\delta_{i,0}|, |\delta_{i,1}| \leq \frac{\varepsilon}{10\sqrt{n}}$. Hence,

$$\begin{aligned} |\delta_{i,2}| &= \left| \Pr_{\tau}[x_{2i-1} \oplus x_{2i} = 1] - \frac{1}{2} \right| = \left| \left(\frac{1}{2} - \delta_{i,0} \right) \left(\frac{1}{2} + \delta_{i,1} \right) + \left(\frac{1}{2} + \delta_{i,0} \right) \left(\frac{1}{2} - \delta_{i,1} \right) - \frac{1}{2} \right| \\ &= 2|\delta_{i,0}||\delta_{i,1}| \leq \frac{\varepsilon^2}{50n} \end{aligned}$$

Let $I = \{2i : i \in I_2\} \cup \{2i - 1 : i \in I_2\}$ Let $f : \{0, 1\}^I \rightarrow \{0, 1\}^{|I_2|}$ be the function that maps every pair in I_2 to its parity bit. In other words, for every $i \in I_2$, the bits x_{2i-1}, x_{2i} are mapped to a single bit $x_{2i-1} \oplus x_{2i}$. Let μ' (respectively τ') be the distribution over $\{0, 1\}^{|I_2|}$ that draws a sample $x \sim \mu$ (respectively $x \sim \tau$) and returns $f(x)$. Observe that both μ' and τ' are product distributions over $\{0, 1\}^{|I_2|}$, since the pairs are independent and every pair is mapped into a single bit.

Note that $d_{\text{TV}}(\tau', \pi_{|I_2|}) \leq \Delta_{\text{TV}}(\tau', \pi_{|I_2|}) = \sum_{i=1}^{|I_2|} d_{\text{TV}}(\tau'|_i, \text{Ber}(1/2)) \leq |I_2| \cdot \frac{\varepsilon^2}{50n} \leq \frac{1}{50}\varepsilon^2$.

In μ' , by definition of the ν_b s, all marginals have the form $\frac{1}{2} \pm \frac{\varepsilon}{\sqrt{n}}$, hence by Lemma 10.2,

$$d_{\text{TV}}(\mu', \pi_{|I_2|}) \geq \frac{1}{16} \cdot \frac{\varepsilon}{\sqrt{n}} \cdot \sqrt{\frac{1}{6}n} = \frac{\varepsilon}{16\sqrt{6}} \geq \frac{1}{40}\varepsilon$$

By a data processing inequality and the triangle inequality,

$$d_{\text{TV}}(\mu, \tau) \geq d_{\text{TV}}(\mu', \tau') \geq d_{\text{TV}}(\mu', \pi_{|I_2|}) - d_{\text{TV}}(\tau', \pi_{|I_2|}) \geq \frac{1}{40}\varepsilon - \frac{1}{50}\varepsilon^2 > \frac{1}{400}\varepsilon \quad \square$$

References

- [BC17] Rishiraj Bhattacharyya and Sourav Chakraborty. Property testing of joint distributions using conditional samples. *CoRR*, abs/1702.01454, 2017.
- [BCŠV23] Antonio Blanca, Zongchen Chen, Daniel Štefankovič, and Eric Vigoda. Complexity of high-dimensional identity testing with coordinate conditional sampling. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 1774–1790. PMLR, 2023.

- [BFF⁺01] Tugkan Batu, Eldar Fischer, Lance Fortnow, Ravi Kumar, Ronitt Rubinfeld, and Patrick White. Testing random variables for independence and identity. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 442–451. IEEE, 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.
- [BGKV21] Arnab Bhattacharyya, Sutanu Gayen, Saravanan Kandasamy, and NV Vinodchandran. Testing product distributions: A closer look. In *Algorithmic Learning Theory*, pages 367–396. PMLR, 2021.
- [BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comput. Syst. Sci.*, 47(3):549–595, 1993.
- [CCK⁺21] Clément L Canonne, Xi Chen, Gautam Kamath, Amit Levi, and Erik Waingarten. Random restrictions of high dimensional distributions and uniformity testing with subcube conditioning. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 321–336. SIAM, 2021.
- [CDKS17] Clément L Canonne, Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Testing bayesian networks. In *Conference on Learning Theory*, pages 370–448. PMLR, 2017.
- [CFGM13] Sourav Chakraborty, Eldar Fischer, Yonatan Goldhirsh, and Arie Matsliah. On the power of conditional samples in distribution testing. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 561–580, 2013.
- [CJLW21] Xi Chen, Rajesh Jayaram, Amit Levi, and Erik Waingarten. Learning and testing junta distributions with sub cube conditioning. In *Conference on Learning Theory*, pages 1060–1113. PMLR, 2021.
- [CM24] Xi Chen and Cassandra Marcussen. Uniformity testing over hypergrids with subcube conditioning. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4338–4370. SIAM, 2024.
- [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.
- [Gol17] Oded Goldreich. *Introduction to property testing*. Cambridge University Press, 2017.
- [GR11] Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 68–75, 2011.
- [GR23] Oded Goldreich and Dana Ron. Testing distributions of huge objects. *TheoretCS*, 2, 2023.
- [Lev85] Leonid A Levin. One-way functions and pseudorandom generators. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 363–365, 1985.
- [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.

[RS96] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.

A Explicit algorithms

We provide here explicit representations for the binary setting algorithms constructed in this paper.

Levin’s work-balance method For a random variable X that is bounded between 0 and 1, it distinguishes between $E[X] = 0$ and $E[X] > \varepsilon$. To do that, we use another random variable Y , for which we can distinguish between $E[X|Y = y] = 0$ and $E[X|Y = y] > \rho$ for every $0 < \rho < 1$, at a cost of $O(1/\rho)$. Overall, the cost of the algorithm is $O(\log^2 \varepsilon^{-1}/\varepsilon)$. The original construction is found in [Lev85] and the following optimized version appears as an exercise in [Gol17].

The black box is run a logarithmic number of times for every y that we draw (as opposed to once in, for example, [Gol17, Exercise 8.4]) since it refers here to a procedure with two-sided error.

Algorithm 1 Levin’s work-balance procedure

input Y – a random variable, accessible through unconditional sampling.

input ε – a threshold parameter.

input A random black box that, for every $y \in \text{supp}(Y)$ and $0 < \varepsilon' < 1$:

completeness if $E[X|Y = y] = 0$, it accepts with at least probability $2/3$.

soundness if $E[X|Y = y] > \varepsilon'$, it rejects with probability at least $2/3$.

resource cost complexity $O(1/\varepsilon')$.

completeness If $E[X] = 0$, then the output is ACCEPT with probability at least $2/3$.

soundness If $E[X] > \varepsilon$, then the output is REJECT with probability at least $2/3$.

resource cost complexity $O(\log^2 \varepsilon^{-1}/\varepsilon)$.

for t from 1 to $\lceil \log_2(2/\varepsilon) \rceil$ **do**

let $\varepsilon' \leftarrow 2^{-t}$.

for $\lceil 2^{3-t}\varepsilon^{-1} \rceil$ **times do**

draw $y \sim Y$.

set $r \leftarrow 0$.

for $\lceil 64(\log_2 \varepsilon^{-1} + 2) \rceil$ **times do**

run the black box with (y, ε') .

if the black box accepts **then**

set $r \leftarrow r + 1$.

else

set $r \leftarrow r - 1$.

if $r < 0$ **then**

return REJECT.

return ACCEPT.

χ^2 -test of single-bit distributions The following is the algorithm that detects the difference between two Bernoulli distributions as it was described in Lemma 5.5.

Algorithm 2 χ^2 -test of single-distributions

input Two Bernoulli distributions $\text{Ber}(p)$ and $\text{Ber}(q)$, accessible through samples.
completeness If $p = q$, then the algorithm accepts with probability at least $2/3$.
soundness If $\chi^2(p, q) > \varepsilon$, then the algorithm rejects with probability at least $2/3$.
complexity $O(1/\varepsilon)$ samples.
let $N \leftarrow \lceil 16/\varepsilon \rceil$.
set $A, B \leftarrow 0$.
for 64 **times do**
 draw N independent samples from $\text{Ber}(p)$, let X be their sum.
 draw N independent samples from $\text{Ber}(q)$, let Y be their sum.
 if $X > Y$ **then**
 set $A \leftarrow A + 1$.
 if $X < Y$ **then**
 set $B \leftarrow B + 1$.
if $A \leq 40$ and $B \leq 40$ **then**
 return ACCEPT.
else
 return REJECT.

Testing EQUIVALENCE The proof of Theorem 5.8 translates to the following explicit algorithm.

Algorithm 3 ε -test for binary-alphabet EQUIVALENCE

input $n, \varepsilon > 0$, two distributions μ, τ over $\{0, 1\}^n$.
 μ is accessible through the marginal prefix oracle.
 τ is accessible through the prefix oracle.
completeness If $\tau = \mu$, then the output is ACCEPT with probability at least $\frac{2}{3}$.
soundness If $d_{\text{TV}}(\tau, \mu) > \varepsilon$, then the output is REJECT with probability at least $\frac{2}{3}$.
let π be the uniform distribution over $[n]$.
let Y be a random variable that distributes as $\pi \times \tau$.
let X be a random variable defined as a function of $Y = (i, w)$:

$$X(i, w) = \chi^2 \left(\mu|_i^{x_{[i-1]}=w_{[i-1]}}(1), \tau|_i^{x_{[i-1]}=w_{[i-1]}}(1) \right)$$

let $\rho \leftarrow \frac{\varepsilon^2}{24 \log(\varepsilon/2n)}$.

run Levin's procedure (Algorithm 1), where its input consists of $Y, \rho/n$, and the black box $((i, w), \varepsilon') \rightarrow (\text{Algorithm 2 with input } \mu|_i^{x_{[i-1]}=w_{[i-1]}}, \tau|_i^{x_{[i-1]}=w_{[i-1]}} \text{ and } \varepsilon')$.
if Levin's procedure accepts **then**
 return ACCEPT.
else
 return REJECT.

Testing PRODUCT The proof of Theorem 9.4 translates in the binary setting to the following explicit algorithm. Note that it is almost identical to Algorithm 3, since we only substitute the marginal prefix oracle $\tau|_i^{x_{[i-1]}=w_{[i-1]}}$ with the marginal oracle of μ .

Algorithm 4 ε -test for binary-alphabet PRODUCT

input $n, \varepsilon > 0$, a distribution μ over $\{0, 1\}^n$.

μ is accessible through the prefix oracle.

completeness If $\mu \in \text{PRODUCT}$, then the output is ACCEPT with probability at least $\frac{2}{3}$.

soundness If $\min_{\tau \in \text{PRODUCT}} d_{\text{TV}}(\mu, \tau) > \varepsilon$, then the output is REJECT with probability at least $\frac{2}{3}$.

let π be the uniform distribution over $[n]$.

let Y be a random variable that distributes as $\pi \times \mu$.

let X be a random variable defined as a function of $Y = (i, w)$:

$$X(i, w) = \chi^2 \left(\mu|_i^{x_{[i-1]}=w_{[i-1]}}(1), \mu|_i(1) \right)$$

let $\rho \leftarrow \frac{\varepsilon^2}{24 \log(\varepsilon/2n)}$.

run Levin's procedure (Algorithm 1), where its input consists of $Y, \rho/n$, and the black box $((i, w), \varepsilon') \rightarrow (\text{Algorithm 2 with input } \mu|_i^{x_{[i-1]}=w_{[i-1]}}, \mu|_i \text{ and } \varepsilon')$.

if Levin's procedure accepts **then**

return ACCEPT.

else

return REJECT.
