

The query complexity of uniformity testing with purified unitary access

Abstract

We determine the bounded-error query complexity of uniformity testing when access to a distribution is provided by a preparation unitary with arbitrary garbage. For every integer $d \geq 2$ and $0 < \varepsilon < 1/2$, distinguishing the uniform distribution on d labels from distributions at total-variation distance greater than ε requires $\Omega(\min\{d^{1/3}\varepsilon^{-4/3}, \sqrt{d}/\varepsilon\})$ queries. This matches the upper bound of Canonne, Kothari, and O'Donnell (TQC, 2025) and remains valid when inverse and controlled queries are allowed. In particular, it determines the joint dependence on dimension and accuracy, including the transition at $\varepsilon = d^{-1/2}$. The main obstacle is to control information obtained from the action of the entire unitary, not merely from the state it prepares. We construct reflection oracles with finite-dimensional, phase-randomized garbage and compare the uniform distribution with a symmetric Dirichlet prior. Our adversary operator combines a quadratic cutoff on label degrees with normalized Dirichlet moments. An exact posterior covariance identity controls the radial directions orthogonal to these moment vectors. The resulting full-operator norm estimate yields both branches of the lower bound.

Note. This paper was generated entirely by AI, including MiMo, using an automated research pipeline developed by Chenchua Liu and Hanyu Li.

1 Introduction

A program that prepares a sample can be used coherently: a quantum algorithm may apply the preparation unitary, reverse it, and query it on states other than its designated input. Uniformity testing asks how many such queries are needed to distinguish the uniform distribution on d labels from a distribution at total-variation distance greater than ε . The problem measures the benefit of coherent access without assuming that the sampler discards its garbage or has a particular reversible implementation.

In the classical sampling model, Paninski's coincidence test [9] identifies the $\sqrt{d}/\varepsilon^2$ sample scale in the sparse regime: repeated observations reveal departures from uniformity before the full distribution can be learned. Bravyi, Harrow, and Hassidim [2] gave an $O(d^{1/3})$ quantum uniformity tester at constant accuracy for distributions supplied by reversible function oracles. Gilyén and Li [5] developed a framework for purified quantum access that treats classical distributions and quantum states together, with algorithms for closeness, independence, and entropy estimation. These access models differ in the side information exposed by the oracle, so lower bounds must keep the oracle specification explicit.

For purified unitary access, the closeness tester of Luo, Wang, and Li [8] gives an $O(\sqrt{d}/\varepsilon)$ uniformity tester. Canonne, Kothari, and O'Donnell [3] improved the upper bound to $O(\min\{d^{1/3}\varepsilon^{-4/3}, \sqrt{d}/\varepsilon\})$ and conjectured that it is optimal. More recently, Chen, Wang, and Zhang [4, Theorem 2.3] obtained the lower bound $\Omega(d^{1/3} + d^{1/4}/\varepsilon)$ by combining a collision lower bound with quantum sample-to-query lifting. That result captures the optimal dependence on either parameter when the other is fixed, but leaves a gap when both vary. For example, at $\varepsilon = d^{-1/2}$ it gives $\Omega(d^{3/4})$, whereas the upper bound is $O(d)$.

We prove the matching joint lower bound in the purified unitary model of [3, Section 3.2]. Thus we establish the purified-access assertion of their Conjecture 2. The same conjecture also asks for a lower bound in the stronger quantum string-oracle model, where the oracle evaluates a function reversibly on arbitrary inputs. Our result does not establish that stronger assertion.

The central difficulty is to limit the information available through the entire preparation unitary, rather than just through its prepared state. We choose reflection extensions of phase-randomized purifications and compare the uniform distribution with a symmetric Dirichlet prior. The independent phases provide Fourier indices, and Dirichlet moments specify one normalized radial vector in each Fourier sector. A quadratic cutoff on the label degrees penalizes repeated use of a single label more heavily than use of distinct labels. Spreading each label over many independently phased garbage coordinates makes degree-decreasing transitions sparse.

The proof follows the progress-measure principle of the adversary method [6]; adversary bounds for state-generating oracles and distributional input problems were developed by Belovs [1]. We give the progress argument directly on latent L^2 spaces. The problem-specific step is a full-operator norm estimate for the Fourier–Dirichlet construction. In particular, the radial complement of the chosen moment vectors cannot be discarded: an exact Dirichlet covariance identity bounds its contribution. Balancing the cutoff with posterior displacement and posterior fluctuations gives the two parameter regimes. Kutin’s collision lower bound [7] handles the remaining constant-accuracy range.

1.1 The model and the main result

Write $[d] = \{1, \dots, d\}$, let Δ_d be the probability simplex on $[d]$, and let $u_d = (1/d, \dots, 1/d)$. We use $d_{\text{TV}}(p, q) = \frac{1}{2} \sum_{i=1}^d |p_i - q_i|$. A valid preparation oracle is a unitary U with a designated zero state such that

$$U|0\rangle = \sum_{i=1}^d \sqrt{p_i} |i\rangle |\psi_i\rangle, \quad \|\psi_i\| = 1. \quad (1)$$

Measuring the first register samples p . The garbage states need not be orthogonal, and the action of U away from $|0\rangle$ is unrestricted. The algorithm may apply U and $U^\dagger$; all its other operations must be independent of the unknown oracle. We also consider algorithms allowed controlled versions of these queries. Each oracle call has unit cost, and oracle-independent computation is free.

Theorem 1. *There is an absolute constant $c > 0$ such that, for every integer $d \geq 2$ and $0 < \varepsilon < 1/2$, any quantum algorithm that distinguishes*

$$p = u_d \quad \text{from} \quad d_{\text{TV}}(p, u_d) > \varepsilon$$

with success probability at least 0.99 on every valid promised preparation oracle has worst-case query complexity at least

$$c \min \left\{ \frac{d^{1/3}}{\varepsilon^{4/3}}, \frac{\sqrt{d}}{\varepsilon} \right\}.$$

The conclusion remains valid when controlled U and controlled $U^\dagger$ are also available.

The success probability 0.99 can be replaced by any fixed constant strictly between $1/2$ and 1 , at the expense of changing c . Together with [3], the theorem determines the bounded-error query complexity up to absolute constant factors. The cube-root expression applies when $\varepsilon \geq d^{-1/2}$, and the square-root expression applies when $\varepsilon \leq d^{-1/2}$, in each case within the stated range $0 < \varepsilon < 1/2$. Both expressions equal d at the transition. The case $d = 1$ has no nonuniform alternative and is excluded.

The lower bound is a worst-case statement over preparation unitaries, including their garbage spaces. Our hard instances have finite garbage dimension and use only $O(\log d + \log(1/\varepsilon))$ oracle qubits. Every oracle is specified on its entire space, so the proof accounts for queries on arbitrary inputs.

2 A Fourier–Dirichlet adversary

Let M be a positive integer, to be chosen later, and set

$$\mathcal{H} = \mathbb{C}^d \otimes \mathbb{C}^M, \quad \mathcal{H}' = \text{span}\{|\perp\rangle\} \oplus \mathcal{H}.$$

For $p \in \Delta_d$ and $\theta = (\theta_{ij})_{i \in [d], j \in [M]} \in [0, 2\pi)^{dM}$, define

$$|v(p, \theta)\rangle = \sum_{i=1}^d \sum_{j=1}^M \sqrt{p_i/M} e^{i\theta_{ij}} |i, j\rangle, \quad (2)$$

$$|s_v\rangle = (|\perp\rangle + |v\rangle)/\sqrt{2}, \quad U_v = 2|s_v\rangle\langle s_v| - I_{\mathcal{H}'}. \quad (3)$$

The operator U_v is a Hermitian unitary and satisfies $U_v|\perp\rangle = |v\rangle$. Its conditional garbage state for label i is $M^{-1/2} \sum_j e^{i\theta_{ij}} |j\rangle$, which has norm 1. The state $|\perp\rangle$ can be encoded as the designated all-zero input, with an extra flag separating it from the output subspace. After embedding $\mathcal{H}'$ in a qubit space, assign the action $-I$ to unused dimensions.

The construction also works with pairwise orthogonal garbage states. Indeed, apply the known isometry $|i, j\rangle \mapsto |i\rangle|i\rangle|j\rangle$ to the output subspace and extend it to an isometric embedding of $\mathcal{H}'$. Conjugating U_v by this embedding and using $-I$ on the orthogonal complement gives such an oracle. This extension has zero query difference on the complementary subspace, so it does not affect the norm estimates below.

Let ν be normalized Haar measure on the phase torus, and let μ_a be the symmetric Dirichlet measure $\text{Dir}(a, \dots, a)$ on Δ_d . Under the null prior, $p = u_d$ and $\theta \sim \nu$. Under the alternative prior, $p \sim \mu_a$ and $\theta \sim \nu$ independently. We initially allow any $a > 0$; the application uses $a \geq 1$. The alternative prior need not be supported entirely on the promise set. In Section 4 we show that it assigns an absolute positive probability to distributions satisfying the strict separation condition.

To analyze an arbitrary tester, introduce inaccessible latent spaces

$$\mathcal{L}_0 = L^2(\nu), \quad \mathcal{L}_1 = L^2(\mu_a \otimes \nu).$$

The initial latent state on each side is the constant function $\mathbf{1}$. The oracle acts pointwise on these spaces. Thus the joint state is a square-integrable, vector-valued function of the oracle parameters, and tracing out the latent space gives the algorithm’s state averaged over the corresponding prior. This representation does not give the algorithm access to the parameters.

We prove the lower bound even with controlled queries available. Truncating a tester at its worst-case query bound on promised inputs and padding halted branches with inactive controlled queries gives a fixed number of query steps on all inputs, without changing its behavior on the promise. Intermediate measurements and randomness can be purified and retained in the workspace. We may therefore work with joint pure states throughout, including on inputs outside the promise.

Suppose $\Gamma : \mathcal{L}_1 \rightarrow \mathcal{L}_0$ is a bounded operator satisfying $\|\Gamma\| = A$ and $\Gamma\mathbf{1} = A\mathbf{1}$. For the joint state $|\Psi_b^t\rangle$ after t queries under prior b , set

$$W_t = \langle \Psi_0^t | (\Gamma \otimes I) | \Psi_1^t \rangle.$$

Initially $W_0 = A$. Oracle-independent operations commute with $\Gamma \otimes I$ and leave W_t unchanged. Write U_b for the pointwise oracle on side b . If

$$\|(\Gamma \otimes I)U_1 - U_0(\Gamma \otimes I)\| \leq B, \quad (4)$$

then a query changes W_t by at most B : the difference of the two progress operators is

$$U_0^\dagger(\Gamma \otimes I)U_1 - (\Gamma \otimes I) = U_0^\dagger((\Gamma \otimes I)U_1 - U_0(\Gamma \otimes I)).$$

Tensoring with an arbitrary workspace preserves the norm bound. For a controlled query, the inactive block has zero difference and the active block is exactly (4). Inverse queries coincide with forward queries for our Hermitian oracles.

Let r_b be the final probability of rejecting uniformity under prior b . Decomposing according to the final acceptance and rejection projections, which commute with $\Gamma \otimes I$, and applying Cauchy–Schwarz yields

$$|W_T| \leq A \left(\sqrt{(1-r_0)(1-r_1)} + \sqrt{r_0 r_1} \right). \quad (5)$$

A constant separation between these output distributions therefore forces a constant fractional decrease in the progress measure, and hence $T = \Omega(A/B)$. We will verify the required separation after choosing a .

We now construct Γ . For $\mathbf{n} \in \mathbb{Z}^{dM}$, define

$$e_{\mathbf{n}}(\theta) = e^{i \sum_{i,j} n_{ij} \theta_{ij}}, \quad k_i(\mathbf{n}) = \sum_{j=1}^M |n_{ij}|.$$

For a degree vector $\mathbf{k} \in \mathbb{Z}_{\geq 0}^d$, define $K(\mathbf{k}) = \sum_i k_i$ and $R(\mathbf{k}) = \sum_i k_i^2$. We abbreviate these by K and R when $\mathbf{k}$ is fixed; in particular, $K \leq R$. We use $\mathbf{e}_{ij}$ for a coordinate vector in $\mathbb{Z}^{dM}$ and $\mathbf{e}_i$ for one in $\mathbb{Z}^d$. The rising factorial is $(x)_r = x(x+1) \cdots (x+r-1)$ for $r \geq 1$, with $(x)_0 = 1$. The Dirichlet density is proportional to $\prod_i p_i^{a-1}$, so its normalizing integral gives the moments

$$m_{\mathbf{k}} = \mathbb{E}_{\mu_a} \prod_i p_i^{k_i} = \frac{\prod_i (a)_{k_i}}{(ad)_K}. \quad (6)$$

Set

$$w_{\mathbf{k}}(p) = \frac{\prod_i p_i^{k_i/2}}{\sqrt{m_{\mathbf{k}}}}, \quad f_{\mathbf{n}}(p, \theta) = e_{\mathbf{n}}(\theta) w_{\mathbf{k}(\mathbf{n})}(p).$$

The functions $e_{\mathbf{n}}$ form an orthonormal basis of $\mathcal{L}_0$, and the functions $f_{\mathbf{n}}$ form an orthonormal family in $\mathcal{L}_1$. Indeed, phase integration separates distinct indices, while $m_{\mathbf{k}}$ normalizes the radial vector in each sector. They are not a basis of $\mathcal{L}_1$: each Fourier sector also contains the radial directions orthogonal to $w_{\mathbf{k}}$. All inner products below are conjugate-linear in the first argument.

For $A \geq 1$, define

$$\alpha(\mathbf{k}) = \left(A - \sum_i k_i^2 \right)_+, \quad \Gamma = \sum_{\mathbf{n} \in \mathbb{Z}^{dM}} \alpha(\mathbf{k}(\mathbf{n})) |e_{\mathbf{n}}\rangle \langle f_{\mathbf{n}}|, \quad (7)$$

where $x_+ = \max\{x, 0\}$. Only finitely many indices satisfy $R < A$, so Γ is finite-rank. Orthogonality gives $\|\Gamma\| = A$ and $\Gamma \mathbf{1} = A \mathbf{1}$, as required. The sum of squared degrees controls both the total degree and the concentration of degree on any one label.

3 The full query norm bound

Lemma 1. *Let $d \geq 2$, $a > 0$, and $A \geq 1$. Set $L_A = (\sqrt{A} + 1)^2$ and choose an integer $M \geq 16A^2L_A$. The adversary operator (7) and the reflection oracles (3) satisfy*

$$B := \|(\Gamma \otimes I_{\mathcal{H}'})U_1 - U_0(\Gamma \otimes I_{\mathcal{H}'})\| \leq 20 \left(\sqrt{1 + A/d} + \frac{A^{3/2}}{a\sqrt{d}} + \frac{A}{\sqrt{ad}} \right). \quad (8)$$

Proof. We first reduce the full reflection to its preparation and erasure blocks. Let $V_b : \mathcal{L}_b \rightarrow \mathcal{L}_b \otimes \mathcal{H}$ be the isometry $h \mapsto h|v_b\rangle$, and write $\Gamma_H = \Gamma \otimes I_{\mathcal{H}}$. Relative to the decomposition into $|\perp\rangle$ and $\mathcal{H}$, the oracle is

$$U_b = \begin{pmatrix} 0 & V_b^\dagger \\ V_b & V_b V_b^\dagger - I \end{pmatrix}.$$

Define

$$D_+ = \Gamma_H V_1 - V_0 \Gamma, \quad D_- = \Gamma V_1^\dagger - V_0^\dagger \Gamma_H.$$

The lower-right block of the query difference is $D_+ V_1^\dagger + V_0 D_-$. Since V_0 and V_1 are isometries,

$$B \leq 2(\|D_+\| + \|D_-\|). \quad (9)$$

In particular, bounding both of these operators controls queries on arbitrary inputs, including inputs orthogonal to the prepared state.

Principal Fourier directions. Let $\mathcal{P} = \overline{\text{span}}\{f_{\mathbf{n}}\}$. For a fixed Fourier index $\mathbf{n}$ with degree vector $\mathbf{k}$, set

$$S = ad + K, \quad q_i = \frac{a + k_i}{S}, \quad b_i = \sqrt{q_i}.$$

The measure $w_{\mathbf{k}}^2 d\mu_a$ is $\text{Dir}(a + k_1, \dots, a + k_d)$, so q is its mean. Multiplication by $v_{ij} = M^{-1/2} \sqrt{p_i} e^{i\theta_{ij}}$ shifts the Fourier index from $\mathbf{n}$ to $\mathbf{n} + \mathbf{e}_{ij}$. If $n_{ij} \geq 0$, this increases k_i by 1; we call it a creation direction. The moment ratio in (6) gives the exact identity

$$\sqrt{p_i} w_{\mathbf{k}} = b_i w_{\mathbf{k} + \mathbf{e}_i}. \quad (10)$$

If $n_{ij} < 0$, the shift decreases k_i by 1; we call it an erasure direction. In this case define

$$b_i^- = \sqrt{\frac{a + k_i - 1}{ad + K - 1}}. \quad (11)$$

Here $k_i \geq 1$ and $K \geq 1$, so the ratio is well-defined and lies in $(0, 1]$. The identity $\sqrt{p_i} w_{\mathbf{k} - \mathbf{e}_i} = b_i^- w_{\mathbf{k}}$ shows that the projection of $\sqrt{p_i} w_{\mathbf{k}}$ onto $w_{\mathbf{k} - \mathbf{e}_i}$ has coefficient b_i^- .

For $D_+|_{\mathcal{P}}$, distinct input vectors $f_{\mathbf{n}}$ have orthogonal output columns: an output component has coordinates $(i, j, \mathbf{n} + \mathbf{e}_{ij})$, which determine $\mathbf{n}$ uniquely. For $D_-|_{\mathcal{P} \otimes \mathcal{H}}$, the row with output index $\mathbf{n}$ has input coordinates $(i, j, \mathbf{n} + \mathbf{e}_{ij})$, and distinct rows have disjoint input blocks. Consequently, the two norms are bounded by the supremum of the Euclidean norms of their column and row coefficients, respectively.

In a creation direction, these coefficients are, respectively,

$$\frac{1}{\sqrt{M}} \left(\alpha(\mathbf{k} + \mathbf{e}_i) b_i - \frac{\alpha(\mathbf{k})}{\sqrt{d}} \right), \quad \frac{1}{\sqrt{M}} \left(\alpha(\mathbf{k}) b_i - \frac{\alpha(\mathbf{k} + \mathbf{e}_i)}{\sqrt{d}} \right). \quad (12)$$

There are at most M such directions for each label i . A nonzero coefficient implies $R < A$. Splitting either coefficient into a change of weight times $1/\sqrt{d}$ and a weight times $b_i - 1/\sqrt{d}$ separates the two effects we need to bound. Since $|\alpha(\mathbf{k} + \mathbf{e}_i) - \alpha(\mathbf{k})| \leq 2k_i + 1$, the first contribution is at most

$$g := \left(\frac{1}{d} \sum_i (2k_i + 1)^2 \right)^{1/2} = \sqrt{1 + 4(K + R)/d} \leq \sqrt{1 + 8A/d}. \quad (13)$$

For the second, with square roots taken componentwise,

$$\begin{aligned} \|\sqrt{q} - \sqrt{u_d}\|_2^2 &\leq d \sum_i (q_i - 1/d)^2 \\ &= \frac{d}{(ad + K)^2} \sum_i (k_i - K/d)^2 \leq \frac{R}{a^2 d}. \end{aligned} \quad (14)$$

The first inequality follows by dividing $q_i - 1/d$ by $\sqrt{q_i} + 1/\sqrt{d} \geq 1/\sqrt{d}$. Since every weight is at most A , the second contribution is bounded by

$$h := \frac{A^{3/2}}{a\sqrt{d}}.$$

Thus the creation contribution to either principal norm is at most $\sqrt{1 + 8A/d} + h$.

For erasure directions, the coefficients are obtained from (12) by replacing $\mathbf{k} + \mathbf{e}_i$ with $\mathbf{k} - \mathbf{e}_i$ and b_i with b_i^- . There are at most K coordinates with $n_{ij} < 0$. If at least one adjacent weight is nonzero, then

$$R < L_A, \quad K \leq R < L_A.$$

To check the boundary case, if $R(\mathbf{k} - \mathbf{e}_i) < A$, then $k_i - 1 < \sqrt{A}$ and $R(\mathbf{k}) = R(\mathbf{k} - \mathbf{e}_i) + 2(k_i - 1) + 1 < A + 2\sqrt{A} + 1$. Each erasure coefficient has absolute value at most $2A/\sqrt{M}$. Its total contribution to either principal norm is therefore at most

$$\eta_M := 2A\sqrt{L_A/M}. \quad (15)$$

The radial complement. In the Fourier sector $\mathbf{n}$, the radial space of $\mathcal{L}_1$ is $L^2(\mu_a)$, whereas $\mathcal{P}$ retains only $w_{\mathbf{k}}$. Let $h_{\mathbf{n}} \perp w_{\mathbf{k}}$, suppressing its Fourier factor. For D_+ , a creation direction contributes the functional

$$\langle w_{\mathbf{k}+\mathbf{e}_i}, \sqrt{p_i} h_{\mathbf{n}} \rangle = \langle r_i, h_{\mathbf{n}} \rangle, \quad r_i = \frac{p_i - q_i}{\sqrt{q_i}} w_{\mathbf{k}}.$$

Indeed, (10) gives $\sqrt{p_i} w_{\mathbf{k}+\mathbf{e}_i} = p_i w_{\mathbf{k}} / \sqrt{q_i}$, and the constant part vanishes against $h_{\mathbf{n}}$. The posterior covariance, obtained from the first two Dirichlet moments, is

$$\text{Cov}(p_i, p_\ell) = \frac{q_i \delta_{i\ell} - q_i q_\ell}{S + 1}.$$

It follows that the residual vectors have the exact Gram matrix

$$(\langle r_i, r_\ell \rangle)_{i,\ell} = \frac{I - bb^\dagger}{S + 1}, \quad b = (\sqrt{q_i})_{i=1}^d, \quad \|b\| = 1. \quad (16)$$

Thus the map $h_{\mathbf{n}} \mapsto (\langle r_i, h_{\mathbf{n}} \rangle)_i$ has norm at most $(S + 1)^{-1/2}$. More explicitly, the squared norm of its weighted creation output is at most

$$\sum_{i,j: n_{ij} \geq 0} \frac{\alpha(\mathbf{k} + \mathbf{e}_i)^2}{M} |\langle r_i, h_{\mathbf{n}} \rangle|^2 \leq A^2 \sum_i |\langle r_i, h_{\mathbf{n}} \rangle|^2 \leq \frac{A^2}{S + 1} \|h_{\mathbf{n}}\|^2.$$

Erasure directions give zero on this complement, because $\sqrt{p_i} w_{\mathbf{k}-\mathbf{e}_i} = b_i^- w_{\mathbf{k}}$. The output blocks for distinct $\mathbf{n}$ are orthogonal, so

$$\|D_+|_{\mathcal{P}^\perp}\| \leq \frac{A}{\sqrt{ad}}. \quad (17)$$

For D_- , fix an output Fourier index $\mathbf{n}$ with degree vector $\mathbf{k}$. Its input in query coordinate (i, j) has Fourier index $\mathbf{n} + \mathbf{e}_{ij}$. If $n_{ij} \geq 0$, the input radial complement is orthogonal to $w_{\mathbf{k}+\mathbf{e}_i}$, so (10) makes its contribution zero. If $n_{ij} < 0$, the input degree vector is $\mathbf{k} - \mathbf{e}_i$. Put

$$q_i^- = (b_i^-)^2 = \frac{a + k_i - 1}{ad + K - 1}.$$

The residual of $\sqrt{p_i} w_{\mathbf{k}}$ orthogonal to $w_{\mathbf{k}-\mathbf{e}_i}$ is

$$s_i = \frac{p_i - q_i^-}{\sqrt{q_i^-}} w_{\mathbf{k}-\mathbf{e}_i}.$$

Under the Dirichlet posterior with degree vector $\mathbf{k} - \mathbf{e}_i$, the mean of p_i is q_i^- and the total parameter is $S - 1$. Hence

$$\|s_i\|^2 = \frac{\text{Var}(p_i)}{q_i^-} = \frac{1 - q_i^-}{S} \leq \frac{1}{ad + K}.$$

Different query coordinates are orthogonal input blocks, and at most K of them have $n_{ij} < 0$. The norm of this output row is at most

$$\alpha(\mathbf{k}) \left(\frac{1}{M} \sum_{i,j: n_{ij} < 0} \|s_i\|^2 \right)^{1/2} \leq A \sqrt{\frac{K}{M(ad + K)}} \leq \frac{A}{\sqrt{M}}.$$

Distinct output rows have disjoint input blocks, including their full radial spaces. Therefore

$$\|D_-|_{\mathcal{P}^\perp \otimes \mathcal{H}}\| \leq \frac{A}{\sqrt{M}}. \quad (18)$$

Combining the principal and radial estimates by the triangle inequality now gives

$$\|D_+\| \leq \sqrt{1 + 8A/d} + h + \eta_M + \frac{A}{\sqrt{ad}}, \quad (19)$$

$$\|D_-\| \leq \sqrt{1 + 8A/d} + h + \eta_M + \frac{A}{\sqrt{M}}. \quad (20)$$

Our choice $M \geq 16A^2 L_A$ ensures $\eta_M \leq 1/2$ and $A/\sqrt{M} \leq 1/(4\sqrt{L_A}) \leq 1/8$, since $A \geq 1$. Substituting into (9) yields

$$B \leq 4\sqrt{1 + 8A/d} + 4h + \frac{2A}{\sqrt{ad}} + \frac{9}{4}.$$

Finally, $\sqrt{1 + 8x} \leq 3\sqrt{1 + x}$ for $x \geq 0$ and $\sqrt{1 + A/d} \geq 1$ imply (8). $\square$

4 Completing the lower bound

To turn the operator estimate into a testing lower bound, we need enough mass on the strict alternative. The following dimension-uniform estimate allows us to take a proportional to ε^{-2} , even for small d .

Lemma 2. *If $d \geq 2$, $a \geq 1$, and $P \sim \text{Dir}(a, \dots, a)$, then*

$$\mathbb{P}\left[d_{\text{TV}}(P, u_d) \geq \frac{1}{24\sqrt{a}}\right] \geq \frac{1}{144}.$$

Proof. Fix i and let $X = P_i - 1/d$. Its variance is

$$\sigma^2 = \frac{d-1}{d^2(ad+1)} \geq \frac{1}{4ad^2},$$

where we used $d-1 \geq d/2$ and $ad+1 \leq 2ad$. The marginal P_i has distribution $\text{Beta}(a, b)$, with $b = (d-1)a \geq a \geq 1$. Expanding the centered fourth moment using $\mathbb{E}P_i^r = (a)_r/(a+b)_r$ for $1 \leq r \leq 4$ gives

$$\frac{\mathbb{E}X^4}{\sigma^4} = 3 + 6 \frac{(a-b)^2(a+b+1) - ab(a+b+2)}{ab(a+b+2)(a+b+3)} \leq 3 + \frac{6}{a} \leq 9.$$

For the first inequality, discard the negative term in the numerator, use $(a-b)^2 \leq b^2$, and note that $b(a+b+1) \leq (a+b+2)(a+b+3)$. Hölder's inequality gives $(\mathbb{E}X^2)^{3/2} \leq \mathbb{E}|X|(\mathbb{E}X^4)^{1/2}$, and hence

$$\mathbb{E}|X| \geq \frac{\sigma}{3} \geq \frac{1}{6d\sqrt{a}}.$$

For $Z = d_{\text{TV}}(P, u_d) = \frac{1}{2} \sum_i |P_i - 1/d|$, it follows that

$$\mathbb{E}Z \geq \frac{1}{12\sqrt{a}}, \quad \mathbb{E}Z^2 \leq \frac{d}{4} \sum_i \mathbb{E}(P_i - 1/d)^2 = \frac{d-1}{4(ad+1)} \leq \frac{1}{4a}.$$

The Paley–Zygmund inequality now yields

$$\mathbb{P}[Z \geq \tfrac{1}{2}\mathbb{E}Z] \geq \frac{(\mathbb{E}Z)^2}{4\mathbb{E}Z^2} \geq \frac{1}{144}.$$

Since $\frac{1}{2}\mathbb{E}Z \geq 1/(24\sqrt{a})$, this proves the claim. $\square$

Proof of Theorem 1. First suppose $0 < \varepsilon \leq 1/48$, and choose

$$a = (48\varepsilon)^{-2} \geq 1, \quad q_0 = 1/144.$$

Lemma 2 gives $\mathbb{P}[d_{\text{TV}}(P, u_d) \geq 2\varepsilon] \geq q_0$. In particular, the alternative prior assigns probability at least q_0 to the strict promise $d_{\text{TV}}(P, u_d) > \varepsilon$.

Let Q be the original algorithm's worst-case query count, and amplify its success probability by nine independent repetitions and majority vote. The error on any promised input is at most $2^9(10^{-2})^5 < (q_0/8)^2$. After truncation and padding as in Section 2, the amplified algorithm uses $T = 9Q$ query steps on all inputs. Put $\delta = (q_0/8)^2$. Its average rejection probabilities satisfy

$$r_0 \leq \delta, \quad r_1 \geq q_0(1 - \delta) \geq q_0/2.$$

We make no assumption about its answers outside the promise. By (5) and $\sqrt{1-x} \leq 1-x/2$,

$$|W_T| \leq A \left(\sqrt{1 - q_0/2} + \sqrt{\delta} \right) \leq A(1 - q_0/8).$$

Since $W_0 = A$ and each query changes the progress by at most B , this implies $T \geq q_0 A / (8B)$, and therefore $Q = \Omega(A/B)$.

It remains to choose the cutoff. If $1 \leq a \leq d$, take $A = (da^2)^{1/3}$. Then

$$\frac{A}{d} \leq 1, \quad \frac{A^{3/2}}{a\sqrt{d}} = 1, \quad \frac{A}{\sqrt{ad}} = (a/d)^{1/6} \leq 1.$$

Lemma 1 gives $B = O(1)$ and $Q = \Omega(d^{1/3}a^{2/3})$. If $a \geq d$, take $A = a$. Each term in (8) is then $O(\sqrt{a/d})$, giving $Q = \Omega(\sqrt{ad})$. Together these estimates show

$$Q = \Omega \left(\min \{ d^{1/3}a^{2/3}, \sqrt{ad} \} \right). \quad (21)$$

Substituting $a = (48\varepsilon)^{-2}$ yields

$$\begin{aligned} Q &= \Omega \left(\min \left\{ 48^{-4/3} d^{1/3} \varepsilon^{-4/3}, 48^{-1} \sqrt{d} \varepsilon^{-1} \right\} \right) \\ &= \Omega \left(\min \left\{ d^{1/3} \varepsilon^{-4/3}, \sqrt{d}/\varepsilon \right\} \right). \end{aligned}$$

The constants are absolute, including near $\varepsilon = d^{-1/2}$; absorbing the fixed powers of 48 does not require the two cutoffs to meet at exactly that value.

Now suppose $1/48 < \varepsilon < 1/2$. By Kutin's collision lower bound [7, Theorem 1.1], for $4 \mid D$, distinguishing a one-to-one function $f : [D] \rightarrow [D]$ from a four-to-one function requires $\Omega(D^{1/3})$ standard quantum function queries. For $d \geq 8$, set $D = 4\lfloor d/4 \rfloor$ and consider the preparation

$$U_f |0\rangle = \frac{1}{\sqrt{d}} \left(\sum_{x=1}^D |f(x)\rangle |x\rangle + \sum_{x=D+1}^d |x\rangle |x\rangle \right). \quad (22)$$

It is a valid oracle of the form (1). If f is one-to-one, the output distribution is exactly u_d . If f is four-to-one, $D/4$ labels have probability $4/d$, another $3D/4$ labels have probability zero, and the remaining $d - D$ labels have probability $1/d$. Its distance from uniform is therefore

$$d_{\text{TV}}(p, u_d) = \frac{3D}{4d} \geq \frac{6}{11} > \frac{1}{2} > \varepsilon.$$

The inequality $D/d \geq 8/11$ follows by writing $d = 4m + r$ with $m \geq 2$ and $0 \leq r \leq 3$.

For completeness, the full preparation unitary and its inverse and controlled versions all require only $O(1)$ function queries. Prepare a uniform superposition of $x \in [d]$ by a known unitary. Reversibly evaluate $f(x)$ for $x \leq D$ into a scratch register, using a fixed legal dummy input to f when $x > D$. Copy the appropriate value, either $f(x)$ or x , into the output register and uncompute the scratch register. This specifies a unitary circuit for (22), with clean scratch registers on its designated input. Reversing the circuit implements its inverse. To control a function evaluation, compute its value unconditionally into scratch space, conditionally add that value to the target register, and uncompute. Thus controlled evaluations also cost $O(1)$ standard queries; for an additive function oracle, its inverse is obtained by conjugating with negation of the target register. Consequently, a uniformity tester with the permitted query types would give a collision tester at constant-factor query overhead. Kutin's bound therefore implies $Q = \Omega(d^{1/3})$.

For $2 \leq d < 8$, uniformity and a point mass are distinct promised inputs, since the latter has distance $1 - 1/d \geq 1/2 > \varepsilon$ from uniform. No zero-query algorithm distinguishes them with success probability 0.99 on both, giving $Q \geq 1 = \Omega(d^{1/3})$ for these finitely many values of d . Finally, throughout the constant-accuracy range,

$$\min\{d^{1/3}\varepsilon^{-4/3}, \sqrt{d}/\varepsilon\} \leq 48^{4/3}d^{1/3},$$

so this completes the lower bound for all parameters.

In the Dirichlet construction we may choose $M = \lceil 16A^2(\sqrt{A} + 1)^2 \rceil = O(A^3)$. Our choices satisfy $A \leq \max\{d, a\}$, and $a = (48\varepsilon)^{-2}$, so encoding the label and garbage registers, including the optional label copy, takes $O(\log d + \log(1/\varepsilon))$ qubits. The collision construction uses $O(\log d)$ oracle qubits. This also verifies the stated finite-register realization of the hard instances. $\square$

References

- [1] Aleksandrs Belovs. Quantum algorithms for classical probability distributions. In *27th Annual European Symposium on Algorithms (ESA 2019)*, LIPIcs, volume 144, pages 16:1–16:11, 2019. doi:10.4230/LIPIcs.ESA.2019.16.
- [2] Sergey Bravyi, Aram W. Harrow, and Avinatan Hassidim. Quantum algorithms for testing properties of distributions. *IEEE Transactions on Information Theory*, 57(6):3971–3981, 2011. doi:10.1109/TIT.2011.2134250.
- [3] Clément L. Canonne, Robin Kothari, and Ryan O’Donnell. Uniformity testing when you have the source code. In *20th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2025)*, LIPIcs, volume 350, pages 7:1–7:20, 2025. doi:10.4230/LIPIcs.TQC.2025.7.
- [4] Kean Chen, Qisheng Wang, and Zhicheng Zhang. A list of complexity bounds for property testing by quantum sample-to-query lifting. Preprint, arXiv:2512.01971, 2025. arXiv:2512.01971.
- [5] András Gilyén and Tongyang Li. Distributional property testing in a quantum world. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, LIPIcs, volume 151, pages 25:1–25:19, 2020. doi:10.4230/LIPIcs.ITCS.2020.25.
- [6] Peter Høyer, Troy Lee, and Robert Špalek. Negative weights make adversaries stronger. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC 2007)*, pages 526–535, 2007. doi:10.1145/1250790.1250867.
- [7] Samuel Kutin. Quantum lower bound for the collision problem with small range. *Theory of Computing*, 1(2):29–36, 2005. doi:10.4086/toc.2005.v001a002.
- [8] Jingquan Luo, Qisheng Wang, and Lvzhou Li. Succinct quantum testers for closeness and k -wise uniformity of probability distributions. *IEEE Transactions on Information Theory*, 70(7):5092–5103, 2024. doi:10.1109/TIT.2024.3393756.
- [9] Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008. doi:10.1109/TIT.2008.928987.