

Counting from Quantum Subset States: Sharp Copy Thresholds and Critical Profiles

Shiva Kintali*

September 5, 2026

Abstract

We study cardinality decisions when the only input is independent copies of the uniform quantum state of an unknown subset. The receiver is unrestricted, but has no membership oracle, input-dependent reflection or preparation unitary, retained classical input, or other input-dependent work. For a universe of size $M \geq 2k$, distinguishing cardinalities k and $k - 1$ with fixed worst-case error requires and suffices to use $k \log k + O(k)$ copies. More generally, the threshold is $k \log(k/g^2) + O(k)$ for gaps $1 \leq g \leq \sqrt{k}$, uniformly in the universe; in a universe of size $2k$, the all-gap order is $\Theta(k \log(1 + k/g^2))$. For every fixed positive integer gap and every fixed offset in the order- k critical window, we identify the limiting optimal average and worst-case errors through an explicit modified-Bessel integral. The resulting quantum error is strictly smaller than the optimal equiprior average error obtained by measuring each copy in the element basis, at every finite offset. The proofs use a uniform comparison with anchored phase states, an explicit spin decomposition, coherent Poissonization, controlled spectral limits, and a strict trace-norm comparison.

1 Introduction

1.1 Motivation and problem overview

A quantum sample of a finite set X is the unit vector

$$|X\rangle = \frac{1}{\sqrt{|X|}} \sum_{i \in X} |i\rangle.$$

Measuring this vector in its defining basis yields one uniform element of X and discards all coherence. A joint measurement on several copies can instead use the off-diagonal terms $|i\rangle\langle j|$, so the number of copies required to infer a global property of X need not equal the ordinary sampling complexity. Uniform subset states are therefore a particularly clean setting in which to separate three issues: the information in a coherent sample, the value of collective measurements, and the stronger power supplied by an oracle or a reusable state-preparation circuit.

We study the most local cardinality decision in this model. An unknown nonempty set $X \subseteq [M]$ has size either k or $k - g$, and the receiver receives only independent copies of $|X\rangle$. We ask:

1. What is the exact leading copy threshold for the adjacent promise $g = 1$?
2. How does that threshold change with a growing gap g ?

*Email: shiva.kintali@gmail.com Homepage: shivakintali.github.io

3. What is the limiting optimal error inside the order- k critical window?
4. Does retaining coherence strictly improve on measuring every copy in the element basis?

The answers expose two different scales. Very small gaps are governed by a coupon-collection threshold. In the dense regime $M = 2k$, sufficiently large gaps are governed instead by an inverse-square estimation scale. At half density and fixed gap, the transition between indistinguishability and reliable decision has a nonclassical limiting profile described by modified Bessel functions.

One motivation comes from recent work on shallow quantum computation with many copies of a classical input. Grier, Morris and Wu prove $\text{TC}^0 \subseteq \text{BQAC}^0 \circ \text{NC}^0$ when polynomially many input copies are supplied; in particular, their construction computes arbitrary symmetric Boolean functions using $n^2/\text{polylog } n$ input copies [1]. The present work isolates a related information question—how many copies carry enough information to decide a symmetric cardinality property—but deliberately removes circuit-depth restrictions and changes the physical encoding to $|X\rangle$. Our results are therefore neither shallow-circuit constructions nor lower bounds on the number of supplied classical strings. They provide an information-theoretic benchmark that any efficient or shallow implementation would still have to respect in this state-copy experiment.

Notation. All logarithms are natural unless a base is displayed. Constants implicit in O_ε and Ω_ε may depend on the fixed error $\varepsilon \in (0, 1/2)$, but not on the universe size or the stated varying cardinality parameters. For nonnegative integer a , set $\binom{a}{b} = 0$ when $b < 0$ or $b > a$. The falling factorial is $(z)_r = z(z-1)\cdots(z-r+1)$, with $(z)_0 = 1$. A sign string's weight is its number of negative signs, equivalently the Hamming weight of the binary string in the encoding $(-1)^{x_i}$. Pauli matrices use the convention

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Integers such as n, m, s are local dimension or weight indices where declared; they do not denote additional copies unless explicitly stated.

1.2 Access model and error criteria

Fix a known universe $[M] = \{1, \dots, M\}$ and integers $1 \leq g < k \leq M$. A single unknown set $X \subseteq [M]$ is promised to have size either k or $k - g$. With T copies the receiver receives exactly

$$(|X\rangle\langle X|)^{\otimes T}.$$

It may append fixed ancillas and perform any quantum channel and final measurement; equivalently, it may choose any binary effect $0 \leq E \leq I$ on the T -copy space. We interpret outcome E as the answer k . Its errors on the two classes are

$$e_k(E, X) = 1 - \text{Tr}\left(E(|X\rangle\langle X|)^{\otimes T}\right), \quad e_{k-g}(E, X) = \text{Tr}\left(E(|X\rangle\langle X|)^{\otimes T}\right).$$

The minimax error is the infimum, over E , of the maximum of these errors over both promised classes. The optimal copy complexity $T_{\text{opt}}(M, k, g, \varepsilon)$ is the smallest nonnegative integer T for which this maximum is at most ε .

There is no retained classical description of X , no query to its membership function, no input-dependent reflection, and no state-preparation circuit or inverse. The receiver also receives no auxiliary record correlated with X . In particular, an isometry used to compare statistical experiments below is not being claimed to have constant-depth or efficient circuit complexity.

For $1 \leq s \leq M$, write

$$\rho_{M,s}^{(T)} = \binom{M}{s}^{-1} \sum_{\substack{X \subseteq [M] \\ |X|=s}} (|X\rangle\langle X|)^{\otimes T}, \quad D(\rho, \sigma) = \frac{1}{2} \|\rho - \sigma\| * 1.$$

We omit M when it is fixed by the surrounding discussion. The equiprior average criterion first chooses one of the two sizes with probability one half and then chooses a uniformly random set of that size. Its optimum is

$$e * \text{av} = \frac{1 - D(\rho_{M,k}^{(T)}, \rho_{M,k-g}^{(T)})}{2}.$$

The corresponding two-mixture minimax value is

$$e_{\text{mm}} = \inf_{0 \leq E \leq I} \max \left\{ 1 - \text{Tr}(E \rho_{M,k}^{(T)}), \text{Tr}(E \rho_{M,k-g}^{(T)}) \right\}.$$

Permutation symmetrization shows that e_{mm} is exactly the optimal worst-case error over all promised subsets. It need not equal the equiprior value e_{av} at finite parameters.

Lemma 1.1 (Decision facts used below). *For trace-one density operators ρ, σ and an effect $0 \leq E \leq I$,*

$$|\text{Tr } E(\rho - \sigma)| \leq D(\rho, \sigma), \quad \max_{0 \leq E \leq I} \text{Tr } E(\rho - \sigma) = D(\rho, \sigma).$$

Trace distance cannot increase under a common quantum channel. For unit vectors u, v ,

$$D(uu^*, vv^*) = \sqrt{1 - |\langle u, v \rangle|^2}.$$

More generally, suppose every state ρ_z in a binary decision experiment is replaced, on the same Hilbert space, by a corresponding state $\tilde{\rho}_z$ with the same correct label and $D(\rho_z, \tilde{\rho}_z) \leq \delta$. For any fixed prior, the two optimal average errors differ by at most δ , and the two optimal worst-case errors differ by at most δ .

Proof. Put $\Delta = \rho - \sigma$ and write its Jordan decomposition as $\Delta = \Delta_+ - \Delta_-$. Since $\text{Tr } \Delta = 0$,

$$\text{Tr } \Delta_+ = \text{Tr } \Delta_- = \frac{1}{2} \|\Delta\|_1 = D(\rho, \sigma).$$

For every effect E , $-\text{Tr } \Delta_- \leq \text{Tr}(E\Delta) \leq \text{Tr } \Delta_+$, and the projection onto the support of Δ_+ attains the upper endpoint. This proves the two effect inequalities and, by applying the maximizing effect to the two equiprobable hypotheses, the formula for e_{av} .

If $\mathcal{N}$ is a quantum channel, then its adjoint is positive and unital, so it maps effects to effects. Therefore

$$D(\mathcal{N}(\rho), \mathcal{N}(\sigma)) = \max_{0 \leq E \leq I} \text{Tr}(\mathcal{N}^*(E)(\rho - \sigma)) \leq D(\rho, \sigma).$$

Restricting $uu^* - vv^*$ to $\text{span}\{u, v\}$ gives two opposite eigenvalues of magnitude $\sqrt{1 - |\langle u, v \rangle|^2}$, which proves the pure-state identity. In the final assertion, every fixed effect has outcome probabilities differing by at most δ on corresponding inputs. Thus its pointwise errors, their prior average, and their supremum each change by at most δ . Taking the infimum over effects gives one inequality between the two optima; interchanging the two experiments gives the reverse inequality.

For completeness, let U_π denote the single-copy unitary induced by a permutation $\pi \in S_M$, put $V_\pi = U_\pi^{\otimes T}$, and set

$$\bar{E} = \frac{1}{M!} \sum_{\pi \in S_M} V_\pi E V_\pi^\dagger.$$

On each cardinality orbit, the error of $\bar{E}$ is constant and equals the orbit-average error of E , namely $1 - \text{Tr}(E\rho_{M,k}^{(T)})$ or $\text{Tr}(E\rho_{M,k-g}^{(T)})$. Hence the worst-case error of $\bar{E}$ is exactly the maximum of these two mixture errors. Conversely, each mixture error of any E is an average of its pointwise errors and is therefore at most its worst-case error. Taking infima in the two directions proves that the original worst-case optimum equals e_{mm} . $\square$

The *element-basis restriction* requires every supplied copy first to be measured in the basis $\{|i\rangle : i \in [M]\}$, followed by arbitrary joint classical processing of the complete transcript. Our strict comparison is with this restriction, not with every adaptive or separable quantum measurement.

1.3 Past work and relation to the literature

The phrase *quantum approximate counting* covers several inequivalent input models. The distinctions below are essential: a consumable copy of $|X\rangle$ does not in general implement a membership query, a reflection about $|X\rangle$, or a state-preparation unitary and its inverse.

Oracle approximate counting and access models. Brassard, Høyer and Tapp introduced quantum counting with coherent queries to the membership predicate [2]; Nayak and Wu proved polynomial-method lower bounds for approximate counting and related symmetric statistics in the quantum query model [3]. Belovs later formalized several interfaces for quantum access to classical probability distributions—including ordinary sampling, coherent queries, and state-generation access—and characterized distribution distinguishability in oracle models through Hellinger distance [4]. Aaronson, Kothari, Kretschmer and Thaler initiated combined-resource lower bounds involving membership queries, QSamples and quantum witnesses. For their factor-two approximate-counting promise on a universe of size M with smaller weight k , their tradeoff implies that either the number of membership queries is $\Omega(\sqrt{M/k})$ or the number of QSamples/state-generation accesses is $\Omega(\min\{k^{1/3}, \sqrt{M/k}\})$ [5]. These are important predecessors, but their unknown-input-dependent oracles are absent here.

Copies-only approximate counting. The closest theorem is Belovs and Rosmanis’s resource-sensitive analysis of sizes a and $(1 + \epsilon)a$ [6]. In its copies-only row, for $M \geq 5a$ and $1/a \leq \epsilon \leq 1$, their accessible author version proves

$$\Omega\left(\min\left\{a, \frac{\sqrt{a}}{\epsilon}, \frac{M}{a\epsilon^2}\right\}\right).$$

Substituting $a = k - g$ and $\epsilon = g/(k - g)$ gives, up to constants when $g \leq k/2$,

$$\Omega\left(\min\left\{k, \frac{k^{3/2}}{g}, \frac{Mk}{g^2}\right\}\right).$$

For $g = 1$ this is only $\Omega(k)$, while their coupon argument gives $O(k \log k)$; they explicitly identify the resulting logarithmic discrepancy. Our adjacent result closes this same-task gap and extends uniformly to the denser range $M \geq 2k$. For $g = o(\sqrt{k})$, our lower bound gains the diverging factor $\log(k/g^2)$. For $g \gg \sqrt{k}$, however, the earlier bound can be stronger in common large-universe regimes, so our all-gap statement is not a uniform improvement over every parameter range.

Learning the whole set. Quantum coupon collection assumes that $|X| = k$ is known and asks for the complete set X , rather than one bit about its cardinality. Arunachalam, Belovs, Childs, Kothari, Rosmanis and de Wolf prove

$$\Theta(k \log \min\{k, M - k + 1\})$$

copies for exact reconstruction with high probability [7]. Bab Hadiashar, Nayak and Sinha sharpen the lower leading term to $(1 - o_k(1))k \log \min\{k, M - k + 1\}$ for sufficiently small fixed error [8], and Nayak and Sinha give a matching leading-term learning upper bound across the density range and analyze padded samples [9]. Thus a $k \log k$ quantum lower bound and leading coefficient one were already known for *reconstruction*. Those results do not imply the same lower bound for binary counting: identifying one of $\binom{M}{k}$ sets may be strictly harder than distinguishing two cardinality orbits.

Classical support size and occupancy. The matching upper bound for adjacent counting is classical in origin. Measuring in the element basis produces uniform draws from X , and the $k \log k + ck$ coupon-collection window and its Poisson/Gumbel missing-coupon limit go back to classical occupancy theory [10]. More general support-size estimation allows nonuniform atoms. Raskhodnikova, Ron, Shpilka and Smith prove strong moment-matching lower bounds under a minimum-mass assumption [11]; Wu and Yang obtain near-minimax distinct-elements bounds by polynomial approximation [12]; and Ferreira Pinto and Harms study testing support at most n versus total-variation fairness, nearly matching its general-distribution sample complexity [13]. These works provide statistical precedents for logarithmic support-size scales, but their models do not give a lower bound against a collective receiver for the exact flat promise considered here.

Subset-state testing, pseudorandomness and certificates. Jeronimo, Magrafta and Wu show that random subset states in an intermediate support range are information-theoretically indistinguishable from Haar-random states given polynomially many copies [14]. Jeronimo, Magrafta, Slote and Wu subsequently prove copies-only support-size indistinguishability for very sparse, exponentially separated random subsets and study how structured subset-state certificates change testability [15]. Their representative promise compares supports $2^{n/8}$ and $2^{n/4}$ in dimension 2^n , rather than dense adjacent sizes, and the certificate results supply extra input correlated with X . Their Johnson-scheme and moment methods are close technical relatives, but neither theorem subsumes our critical regime.

Earlier coherent-sample advantages and alternative encodings. This work is not the first evidence that a coherent subset sample can outperform an element sample. Quantum coupon collection improves sharply near full density [7]; complement sampling outputs an element outside a half-density set from one coherent copy, whereas bounded-error classical sampling needs linear sample complexity [16]. Zhang, He and Zhang obtain logarithmic copy complexity for one missing element using a deliberately different mixed-state encoding and Bell measurements [17]. The last result underscores that copy complexity depends on the physical encoding and is not an upper bound for copies of the pure state $|X\rangle$.

Asymptotic quantum statistics and similarly named problems. The operational use of trace distance is standard binary quantum hypothesis testing [18]. Quantum local asymptotic normality gives a broad precedent for Gaussian or oscillator limits of many-copy experiments in fixed dimension [19]. Here the dimension grows, the hypotheses are permutation mixtures on cardinality

orbits, and the limiting blocks arise from a spin/Johnson decomposition, so the finite-dimensional local-normality theorem does not supply our Bessel profile. Finally, Anshu, Haferkamp, Hwang and Nguyen's *quantum approximate counting* concerns the dimension of an accepting subspace accessed through reflection oracles and complexity classes such as QXC [20]; it is unrelated to copies of uniform subset states.

1.4 Main results and novel contributions

The following statements are proved in the indicated sections.

1. For all $M \geq 2k$,

$$T_{\text{opt}}(M, k, 1, \varepsilon) = k \log k + O_\varepsilon(k).$$

The adjacent lower bound already applies to the two uniform cardinality mixtures; the matching upper bound observes enough distinct elements (Theorem 3.1).

2. Uniformly for $M \geq 2k$ and $1 \leq g \leq \sqrt{k}$,

$$T_{\text{opt}}(M, k, g, \varepsilon) = k \log(k/g^2) + O_\varepsilon(k).$$

For all $1 \leq g < k$ there is a lower bound $\Omega_\varepsilon(k \log(1 + k/g^2))$, matched up to constants when $M = 2k$ (Theorem 3.3). The leading-term formula is informative when $g = o(\sqrt{k})$; its additive term is not negligible throughout the full displayed range.

3. For fixed $g \geq 1$, $c \in \mathbb{R}$, $M = 2k$, and $T_k = \lfloor k \log k + ck \rfloor$, let $\lambda = e^{-c}$. Then both optimal quantum error criteria converge to $(1 - \Phi_g(\lambda))/2$, where

$$\Phi_g(\lambda) = e^{-\lambda} \int_0^\infty 2r e^{-r^2} \sqrt{I_0(2\sqrt{\lambda}r)^2 - I_g(2\sqrt{\lambda}r)^2} dr.$$

The modified Bessel functions are defined in Theorem 4.1. The adjacent proof is expanded in full; Corollary 4.3 supplies every change for a general fixed gap.

4. The element-basis average-error limit is $(1 - C_g(\lambda))/2$, with

$$C_g(\lambda) = \max_{h \geq 0} \Pr\{h - g + 1 \leq Z \leq h\}, \quad Z \sim \text{Pois}(\lambda).$$

For every fixed $g \geq 1$ and finite $\lambda > 0$, $\Phi_g(\lambda) > C_g(\lambda)$ (Corollary 4.4). Neither profile convergence is asserted uniformly for growing g .

5. For the auxiliary phase-program experiment

$$|\psi_x\rangle = N^{-1/2} \sum_{i=1}^N (-1)^{x_i} |i\rangle,$$

the minimum number of consumed programs needed to distinguish every balanced string from every unbalanced string is

$$T_{\text{opt}}(N, \varepsilon) = \frac{N}{4} \log N + O_\varepsilon(N)$$

(Theorem 2.1). This is a program-state discrimination theorem, not the oracle version of a balancedness problem.

6. Under perfect completeness, every reject effect for balanced phase programs has rejection probability at most

$$1 - \frac{r(1 - 2/N)^{2T}}{1 + (r - 1)(1 - 4/N)^T}, \quad r = N/2 + 1,$$

on either adjacent slice. More precisely, rejection probability at least any fixed $\eta \in (0, 1)$ forces

$$T \geq \frac{\log((1 + \eta(r - 1))/(1 - \eta))}{-\log(1 - 4/N)} = \frac{N}{4} \log N + O_\eta(N).$$

Moreover, for fixed $c > 0$ and $\eta \in (0, 1)$, an OR of perfect-completeness tests on disjoint fresh groups, each containing at most cN programs, requires $\Omega_{c,\eta}(N^2)$ programs to reject an adjacent input with probability at least η (Theorem ??). This auxiliary result is separate from the bounded-error lower bound.

What is novel. Subject to the version-of-record and external-priority qualifications stated in Section ??, the principal new result is the adjacent copies-only lower bound $k \log k - O_\epsilon(k)$, which closes the explicit Belovs–Rosmanis logarithmic gap and yields the exact leading coefficient. The small-growing-gap interpolation $k \log(k/g^2) + O_\epsilon(k)$, the half-density all-gap characterization, the fixed-gap Bessel trace-distance profile, and the asymptotic coincidence of the optimal average and minimax profiles appear new for this experiment. The strict inequality $\Phi_g > C_g$ appears to give a new optimal-error separation for fixed-gap support-size decision. The phase-program threshold and its perfect-completeness span corollary also appear to be new auxiliary discrimination statements, although their proof language has close antecedents in symmetric-state and Johnson-scheme analyses [7, 15].

What is derived, standard, or previously known. The $O(k \log k)$ adjacent upper bound is the classical coupon-collector test already recorded in the closest predecessor [10, 6]; our novelty is its matching quantum lower bound. The element-basis Poisson law is derived from the classical missing-coupon limit, and the formula C_g is its shifted binary-testing specialization. The large-gap $O(k^2/g^2)$ algorithm at $M = 2k$ is an elementary Bernoulli-mean test and is not presented as a new quantum mechanism. Helstrom’s operational trace-distance formula, data processing, permutation symmetrization, harmonic/Schur decompositions, and the use of spin or Johnson blocks are established tools [18, 7, 19, 15]. What is new is the quantitative way these ingredients are combined to obtain the sharp lower threshold, the controlled fixed-window limit, the minimax transfer, and the strict coherence witness. Likewise, the paper does not claim the first general advantage of coherent subset samples, the first $k \log k$ subset-state lower bound, an efficient implementation of the optimal collective receiver, or a shallow-circuit result.

1.5 Overview of the proof strategy

Our argument has five layers.

1. Symmetry and the phase-program core. Permutation averaging reduces the average problem to the trace distance between the two uniform cardinality mixtures and allows a worst-case test to be symmetrized without increasing its error. We first study a related phase-state experiment. After coherent Poissonization, tuples are grouped by the parity of their coordinate occupation numbers. Poisson splitting makes these parity fibers independent, producing an exact isometry to a

tensor-product qubit encoding with a tunable Bloch bias. This is a mathematical comparison of experiments; it is not an asserted physical preprocessing circuit.

2. Spin decomposition and the leading lower bound. Uniform mixtures over adjacent Hamming slices are decomposed into harmonic chains, equivalently irreducible collective-spin blocks. Each block has rank one, and a ratio identity for the central and adjacent Hadamard/Krawtchouk columns reduces their fidelity to three collective-spin moments. The resulting trace-distance bound remains small until $T = (N/4) \log N - O(N)$. A parity-fiber Fourier measurement supplies the matching unrestricted phase-program receiver.

3. Transfer to subset counting and gap amplification. For the subset lower bound we restrict to a known $2k$ -element portion of the universe, embed it into a larger known space, and apply a reflection about a fixed uniform vector. On the size- k slice the balanced phase state is converted exactly to $-|X\rangle$, while on the size- $(k-1)$ slice it is uniformly $O(1/k)$ -close to $-|X\rangle$. Contractivity and a tensor-power perturbation bound transfer the phase lower bound to subset states. For a general gap, a common random universe-embedding channel places all intermediate cardinality mixtures in the same space, after which the triangle inequality telescopes the g -gap distance into g adjacent distances. Requiring constant total advantage in the resulting exponential trace-distance bound yields the logarithm $\log(k/g^2)$. The matching small-gap upper bound counts distinct observed elements; for $M = 2k$ and larger gaps, projection onto the known uniform vector completes the all-gap upper bound.

4. Critical-window spectral limit. To resolve $T = k \log k + ck$, the proof keeps the full spin-block trace norm rather than bounding it only through fidelity. A two-sided comparison between fixed copy number and coherent Poissonization is proved with uniform errors. On the central spin annulus, the rescaled tridiagonal spin generator converges to the bilateral nearest-neighbor operator; its exponential has matrix entries $I_j(2t)$. Local binomial limits turn the block weights into a Rayleigh density, producing the displayed modified-Bessel integral. Separate central and tail estimates justify passage to the integral. Finally, the two hypothesis block-mass distributions are asymptotically equalized, which turns the equiprior Helstrom limit into the same worst-case minimax limit.

5. Classical baseline and strict coherence. Under element-basis measurement, the only relevant statistic is the number of unseen coupons; its fixed-window limit is Poisson and yields $C_g(\lambda)$. In each limiting quantum block the hypothesis difference is rank two. Dephasing removes a nonzero off-diagonal coherence, and a strict trace-norm comparison shows that this loss is strict before integration. Hence $\Phi_g(\lambda) > C_g(\lambda)$ for every fixed $g \geq 1$ and finite positive λ .

1.6 Organization of the paper

Section 2 proves the sharp phase-program threshold: first the harmonic-chain spin-fidelity estimate, then coherent Poissonization, and finally a matching receiver. Section 3 transfers that bound to adjacent subset counting and proves the growing-gap and half-density all-gap results. Section 4 develops the fixed-window spin limit, the Bessel profile, the element-basis Poisson benchmark, and the strict collective advantage.

2 The sharp leading threshold for phase programs

This section changes the resource model. For $x \in \{0, 1\}^N$, a receiver is given independent copies of

$$|\psi_x\rangle = N^{-1/2} \sum_{i=1}^N (-1)^{x_i} |i\rangle,$$

but has no further access to x or any other input-dependent work. The receiver and final measurement are unrestricted. In particular, an abstract measurement here is not assumed to have a shallow physical implementation. Since $|\psi_{x \oplus 1^N}\rangle = -|\psi_x\rangle$, the physical input determines x only up to complementation. Balancedness is nevertheless well defined on the resulting ray because $|x| = N/2$ if and only if $|x \oplus 1^N| = N/2$. Write $T_{\text{opt}}(N, \varepsilon)$ for the least number of copies with which a single binary measurement has error at most ε on every balanced and every unbalanced input.

Theorem 2.1 (Program-only balancedness). *For fixed $0 < \varepsilon < 1/2$, as $N \rightarrow \infty$ through even integers, the minimum number of these phase programs needed to distinguish every balanced input from every unbalanced input, with worst-case error at most ε , is*

$$T_{\text{opt}}(N, \varepsilon) = \frac{N}{4} \log N + O_\varepsilon(N).$$

The lower bound allows nonzero error on balanced inputs and already applies to balanced versus one adjacent slice. This is not a lower bound on supplied classical input copies.

2.1 A direct spin-fidelity estimate

The following estimate is the representation-theoretic core of the lower bound.

Lemma 2.2 (Central-slice product estimate). *Let $n = 2m$, let $0 < a < 1$, and encode a bit b by*

$$|\phi_b\rangle = \sqrt{\frac{1+a}{2}} |0\rangle + (-1)^b \sqrt{\frac{1-a}{2}} |1\rangle.$$

With $|\Phi_x\rangle = \bigotimes_{i=1}^n |\phi_{x_i}\rangle$, let

$$\sigma_w = \binom{n}{w}^{-1} \sum_{|x|=w} |\Phi_x\rangle\langle\Phi_x|$$

be the uniform mixture over weight- w inputs. Then

$$\frac{1}{2} \|\sigma_m - \sigma_{m-1}\|_1 \leq \min \left\{ 1, \frac{\sqrt{24}}{a\sqrt{n}} \right\}. \quad (1)$$

Proof. Here is the required block decomposition, including its normalization. On the n -qubit product space, let Π_w project onto weight w , let $H_n = H^{\otimes n}$ for $H = 2^{-1/2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$, and let F_a be diagonal with entry $\sqrt{p(s)}$ on weight s , where

$$p(s) = \left(\frac{1+a}{2} \right)^{n-s} \left(\frac{1-a}{2} \right)^s.$$

The computational-basis amplitudes give the exact identity $|\Phi_x\rangle = 2^{n/2} F_a H_n |x\rangle$. Therefore

$$\sigma_w = \frac{2^n}{\binom{n}{w}} F_a H_n \Pi_w H_n F_a.$$

Put $L_+ = \sum_i |1\rangle\langle 0|_i$ and $L_- = L_+^\dagger$, and let $\hat{w}$ be the Hamming-weight operator. Their commutator is $[L_-, L_+] = nI - 2\hat{w}$. If a vector v has weight $j - 1$ with $j \leq n/2$ and $L_+v = 0$, then

$$0 = \|L_+v\|^2 = \|L_-v\|^2 + (n - 2j + 2)\|v\|^2,$$

so $v = 0$. Thus L_+ is injective from weight $j - 1$ to weight j , and its adjoint L_- is surjective in the reverse direction. Hence the weight- j kernel of L_- has dimension

$$d_j = \binom{n}{j} - \binom{n}{j-1}.$$

An orthonormal kernel vector generates the normalized ladder

$$|h, s\rangle = \frac{L_+^s h}{\sqrt{s!(n-2j)!/(n-2j-s)!}}, \quad 0 \leq s \leq n-2j.$$

Repeated use of the commutator gives

$$\|L_+^s h\|^2 = s! \frac{(n-2j)!}{(n-2j-s)!},$$

which verifies the normalization. The same commutations show that ladders from orthogonal kernel vectors of the same weight are orthogonal. If their starting weights differ, moving lowering operators to the right leaves a positive power of L_- acting on a kernel vector, so the ladders are again orthogonal. Finally,

$$\sum_{j=0}^m \left(\binom{n}{j} - \binom{n}{j-1} \right) (n-2j+1) = 2^n;$$

this follows by shifting the second sum and using $\sum_{j=0}^{m-1} \binom{n}{j} + \frac{1}{2} \binom{n}{m} = 2^{n-1}$. The ladders therefore exhaust the space. The collective raising, lowering, and z -spin operators preserve each ladder, with exactly the spin- $r/2$ matrix entries of the symmetric subspace on $r = n - 2j$ qubits. The collective Hadamard is the same fixed spin rotation on every multiplicity copy of this irreducible representation; any overall phase cancels in the density matrices.

In the spin- $r/2$ ladder, let the basis index s count raising steps and let b denote the input column. Counting the overlap signs by their intersection size gives the normalized Dicke matrix entry

$$h^{(r)} * s, b = 2^{-r/2} \sqrt{\frac{\binom{r}{b}}{\binom{r}{s}}} K_s^{(r)}(b), \quad K_s^{(r)}(b) = [z^{sb}(1+z)^{r-b}].$$

The restriction of Π_w to a ladder selects $b = w - j$, the collective Hadamard contributes $h_{s,b}^{(r)}$, and F_a contributes $\sqrt{p(j+s)}$ on the output row. Consequently each mixture block has rank one, with vector

$$u * w, j(s) = \sqrt{\frac{2^n}{\binom{n}{w}}} \sqrt{p(j+s)} h^{(r)} * s, w-j,$$

or zero if w is outside the ladder. Thus

$$\sigma_w = \bigoplus_j u * w, j u_{w,j}^\dagger \otimes I_{d_j}, \quad \sum_j d_j \|u_{w,j}\|^2 = 1.$$

This does not assert that mixtures of different weights commute.

Write $\ell = m - j$, so $r = 2\ell$. The central Hadamard column vanishes at odd s . For $\ell \geq 1$ and even s , the two generating polynomials give

$$K_s^{(2\ell)}(\ell) = (-1)^{s/2} \binom{\ell}{s/2}, \quad K_s^{(2\ell)}(\ell-1) = \frac{\ell-s}{\ell} K_s^{(2\ell)}(\ell).$$

Combining this coefficient ratio with $\binom{2\ell}{\ell-1}/\binom{2\ell}{\ell} = \ell/(\ell+1)$ yields, on the support of the central column,

$$\frac{h^{(2\ell)} * s, \ell-1}{h^{(2\ell)} * s, \ell} = \frac{\ell-s}{\sqrt{\ell(\ell+1)}}.$$

The neighboring slice has no vector when $\ell = 0$. Let $Z = \ell - s$ and $L = \sqrt{\ell(\ell+1)}$, with $Z/L = 0$ at that endpoint. On the central support, the definitions above imply

$$u_{m-1,j}(s) = \sqrt{\frac{m+1}{m}} \frac{Z}{L} u_{m,j}(s).$$

Measuring the commuting pair of total spin and J_z in σ_m gives the probability distribution $d_j |u_{m,j}(s)|^2$. Therefore

$$\sum_j d_j \langle u_{m,j}, u_{m-1,j} \rangle = \sqrt{\frac{m+1}{m}} \mathbb{E}[Z/L].$$

Purify the block decomposition using orthogonal registers for j and for its d_j multiplicity copies. Multiplying the neighboring vector in each j -block by a phase makes the corresponding inner product nonnegative, so the two purifications have overlap

$$F = \sum_j d_j |\langle u_{m,j}, u_{m-1,j} \rangle|.$$

Tracing out the purifying registers produces σ_m and σ_{m-1} . Contractivity and the pure-state trace-distance formula therefore show that the left side of (1) is at most $\sqrt{1-F^2}$. Individual block overlaps need not be nonnegative; the phase choices and absolute values are essential.

Three elementary spin moments control F . With $J_\alpha = \frac{1}{2} \sum_i \sigma_\alpha^{(i)}$ for $\alpha \in \{x, y, z\}$, the local Bloch vector is $((-1)^{x_i} \sqrt{1-a^2}, 0, a)$. For a uniformly random central-slice string, the product of two distinct signs has expectation $-1/(n-1)$. Since Z is the J_z eigenvalue and $L^2 = \ell(\ell+1)$ is the $J_x^2 + J_y^2 + J_z^2$ eigenvalue, expanding these operators gives

$$\mathbb{E}Z = ma, \quad \text{Var } Z = \frac{m(1-a^2)}{2}, \quad \mathbb{E}(L^2 - Z^2) = \frac{m(1+a^2)}{2} \leq m.$$

The first two identities follow by summing the one-qubit J_z moments. Moreover,

$$\mathbb{E}J_x^2 = \frac{ma^2}{2}, \quad \mathbb{E}J_y^2 = \frac{m}{2},$$

because $\sum_{i \neq q} (-1)^{x_i + x_q} = -n$ for every central-slice string. Adding these identities proves the formula for $\mathbb{E}(L^2 - Z^2)$. On $Z \geq ma/2$,

$$1 - Z/L = \frac{L^2 - Z^2}{L(L+Z)} \leq \frac{2(L^2 - Z^2)}{m^2 a^2}.$$

Let $G = \{Z \geq ma/2\}$. Chebyshev's inequality gives

$$\Pr(G^c) \leq \frac{4 \text{Var } Z}{m^2 a^2} = \frac{2(1-a^2)}{ma^2} \leq \frac{2}{ma^2}.$$

On G^c the deficit $1 - Z/L$ is at most two. Consequently

$$1 - \mathbb{E}(Z/L) \leq \frac{2\mathbb{E}(L^2 - Z^2)}{m^2 a^2} + 2\Pr(G^c) \leq \frac{6}{ma^2},$$

and therefore

$$F \geq \max\{0, 1 - 6/(ma^2)\}.$$

If $1 - 6/(ma^2) > 0$, then $\mathbb{E}(Z/L) > 0$ and $F \geq |\sum_j d_j \langle u_{m,j}, u_{m-1,j} \rangle| \geq 1 - 6/(ma^2)$; otherwise $F \geq 0$ suffices. Hence, when the nontrivial bound is below one,

$$\frac{1}{2} \|\sigma_m - \sigma_{m-1}\|_1 \leq \sqrt{1 - F^2} \leq \sqrt{2(1 - F)} \leq \sqrt{\frac{12}{ma^2}} = \frac{\sqrt{24}}{a\sqrt{n}}.$$

The trace distance is always at most one, which completes the proof. $\square$

2.2 Coherent Poissonization and a fixed copy count

Let $n = 2m$ and $h > 0$. Consider the anchored family

$$|\psi_x^{(h)}\rangle = \frac{\sqrt{h}|0\rangle + \sum_{i=1}^n (-1)^{x_i} |i\rangle}{\sqrt{n+h}}, \quad x \in \{0, 1\}^n.$$

For $0 \leq w \leq n$, define

$$\rho_{w,T}^{(h)} = \binom{n}{w}^{-1} \sum_{|x|=w} (|\psi_x^{(h)}\rangle \langle \psi_x^{(h)}|)^{\otimes T}, \quad D_T^{(h)} = D(\rho_{m,T}^{(h)}, \rho_{m-1,T}^{(h)}).$$

Lemma 2.3 (Anchored fixed-copy estimate). *For every even $n \geq 2$, every $h > 0$, and every integer $0 \leq T \leq n^2 - n$,*

$$D_T^{(h)} \leq \frac{2\sqrt{24}}{\sqrt{n}} \exp\left(\frac{2(T+n)}{n+h}\right). \quad (2)$$

If $h \in \{1, 2\}$, then the following holds for every integer $T \geq 0$:

$$D_T^{(h)} \leq \min\left\{1, \frac{C e^{2T/(n+h)}}{\sqrt{n}}\right\}, \quad C = 4\sqrt{6} e^2. \quad (3)$$

Proof. For a Poisson mean $\lambda > 0$, introduce, only as a mathematical comparison, the normalized pure state

$$|\Xi_x\rangle = \bigoplus_{s \geq 0} \sqrt{\Pr[\text{Pois}(\lambda) = s]} |\psi_x^{(h)}\rangle^{\otimes s}.$$

Let $D_{\text{coh}}^{(h)}(\lambda)$ be the trace distance between the central and adjacent mixtures of these coherent-count states. Dephasing the count register gives the block-diagonal distance $\sum_{s \geq 0} \Pr[\text{Pois}(\lambda) = s] D_s^{(h)}$. Moreover, discarding one program is a common channel, so $D_s^{(h)}$ is nondecreasing in s . Contractivity under dephasing therefore gives

$$D_T^{(h)} \Pr[\text{Pois}(\lambda) \geq T] \leq D_{\text{coh}}^{(h)}(\lambda).$$

To identify the coherent-count experiment exactly, expand $|\Xi_x\rangle$ in ordered tuples and group a tuple according to the set $S \subseteq [n]$ of nonanchor modes having odd occupation. The input-dependent

phase is constant on each group and equals $(-1)^{x \cdot S}$. Distinct groups are orthogonal, and normalizing their positive-amplitude sums gives input-independent unit vectors $|\eta_S\rangle$ such that

$$|\Xi_x\rangle = \sum_{S \subseteq [n]} \sqrt{p_\lambda(S)} (-1)^{x \cdot S} |\eta_S\rangle.$$

Poisson splitting says that the n nonanchor occupation numbers are independent $\text{Pois}(\lambda/(n+h))$ variables; the anchor has mean $\lambda h/(n+h)$ and carries no input phase. Hence the parity indicators are independent and

$$\Pr(i \in S) = \frac{1 - e^{-2\lambda/(n+h)}}{2} = \frac{1-a}{2}, \quad a = e^{-2\lambda/(n+h)}.$$

The isometry $|\eta_S\rangle \mapsto |S\rangle$ consequently sends every $|\Xi_x\rangle$ to

$$\bigotimes_{i=1}^n \left(\sqrt{\frac{1+a}{2}} |0\rangle + (-1)^{x_i} \sqrt{\frac{1-a}{2}} |1\rangle \right).$$

Thus Lemma 2.2 bounds $D_{\text{coh}}^{(h)}(\lambda)$. This is an exact equivalence of statistical experiments, not an asserted efficient physical transformation.

Now assume $T \leq n^2 - n$ and choose $\lambda = T + n$. If $S \sim \text{Pois}(\lambda)$, Cantelli's inequality gives

$$\Pr(S < T) \leq \Pr(S - \lambda \leq -n) \leq \frac{\lambda}{\lambda + n^2} \leq \frac{1}{2}.$$

It follows from the preceding comparison and Lemma 2.2 that

$$D_T^{(h)} \leq 2D_{\text{coh}}^{(h)}(T+n) \leq \frac{2\sqrt{24}}{\sqrt{n}} \exp\left(\frac{2(T+n)}{n+h}\right),$$

which proves (2). If $h \in \{1, 2\}$, then $e^{2n/(n+h)} \leq e^2$, so the same estimate and the trivial bound $D_T^{(h)} \leq 1$ give (3) throughout the finite range. When $T > n^2 - n$, the nontrivial expression in (3) exceeds one: for $n = 2$ this is immediate, and for even $n \geq 4$ one has $2(n^2 - n)/(n+2) \geq n/2 \geq \frac{1}{2} \log n$. Thus the minimum in (3) is one, and the inequality remains valid. $\square$

2.3 Matching unrestricted receiver

Lemma 2.4 (Uniform phase-ray reconstruction). *Let $N \geq 2$ be even, let $0 < \varepsilon < 1/2$, and let T be a nonnegative integer. If*

$$T \geq \frac{N}{4} \log \frac{2N}{\varepsilon^2},$$

then a single POVM, independent of x , identifies the representative of the ray $\{x, x \oplus 1^N\}$ whose last bit is zero with probability at least $1 - \varepsilon$, uniformly in x .

Proof. Choose the representative of each global sign pair whose last sign is positive, equivalently whose last bit is zero. Relabel that last coordinate as an anchor and put $n = N - 1$. On T programs, expand the tensor power in ordered mode tuples and group a tuple according to the set $S \subseteq [n]$ of nonanchor modes with odd occupation. If $p_T(S)$ is the probability of this parity set under T independent uniform draws from the N modes, and $|\eta_S\rangle$ is the normalized uniform vector over its tuple fiber, then

$$|\Omega_x\rangle = \sum_S \sqrt{p_T(S)} (-1)^{x \cdot S} |\eta_S\rangle.$$

The nonempty-fiber vectors are orthonormal and independent of x . For $y \in \{0, 1\}^n$, define

$$|f_y\rangle = 2^{-n/2} \sum_{S: p_T(S) > 0} (-1)^{y \cdot S} |\eta_S\rangle.$$

Character orthogonality gives

$$\sum_{y \in \{0, 1\}^n} |f_y\rangle\langle f_y| = \sum_{S: p_T(S) > 0} |\eta_S\rangle\langle \eta_S| =: P.$$

Thus these rank-one effects, together with the failure effect $I - P$, form a POVM. On input x , the probability of the correct representative outcome is

$$|\langle f_x | \Omega_x \rangle|^2 = 2^{-n} \left(\sum_S \sqrt{p_T(S)} \right)^2.$$

Let $u(S) = 2^{-n}$ be uniform on subsets of $[n]$, and define $\chi^2(p\|u) = \sum_S (p(S) - u(S))^2 / u(S)$. The parity sets follow the walk that flips each specified coordinate with probability $1/N$ and is idle, through the anchor, with probability $1/N$. For a character indexed by $A \subseteq [n]$,

$$\mathbb{E}[(-1)^{|A \cap S|}] = (1 - 2|A|/N)^T.$$

Parseval's identity for the Boolean cube therefore gives

$$\begin{aligned} \chi^2(p_T\|u) &= \sum_{j=1}^{N-1} \binom{N-1}{j} (1 - 2j/N)^{2T} \\ &\leq \exp(Ne^{-4T/N}) - 1. \end{aligned}$$

To prove the inequality, pair each $1 \leq j < N/2$ with $N - j$. Their binomial coefficients sum to $\binom{N-1}{j} + \binom{N-1}{j-1} = \binom{N}{j}$, and their eigenvalues have the same absolute value. The central eigenvalue is zero because the displayed lower bound on T is positive. Since $1 - z \leq e^{-z}$,

$$\chi^2(p_T\|u) \leq \sum_{j=1}^N \binom{N}{j} e^{-4Tj/N} = (1 + e^{-4T/N})^N - 1 \leq \exp(Ne^{-4T/N}) - 1.$$

At the stated T , put $q = Ne^{-4T/N} \leq \varepsilon^2/2 < 1$. The elementary bound $e^q - 1 \leq 2q$ for $0 \leq q \leq 1$ gives $\chi^2(p_T\|u) \leq \varepsilon^2$. Hence

$$\text{TV}(p_T, u) \leq \frac{1}{2} \sqrt{\chi^2(p_T\|u)} \leq \varepsilon/2.$$

The affinity satisfies $\sum_S \sqrt{p_T(S)u(S)} \geq \sum_S \min\{p_T(S), u(S)\} = 1 - \text{TV}(p_T, u)$. The correct-outcome probability is the square of this affinity, and is therefore at least $(1 - \varepsilon/2)^2 \geq 1 - \varepsilon$. $\square$

Proof of Theorem 2.1. First suppose $N \geq 4$ and put $n = N - 2 = 2m$. For the lower bound, restrict the promise by fixing one coordinate sign to $+1$ and another to -1 . A fixed unitary maps the normalized vector supported on those two coordinates with amplitudes $(+1, -1)$ to a single anchor basis vector; the orthogonal combination has zero amplitude. The remaining state is therefore the anchored family with $h = 2$. Balanced inputs have remaining weight m , while inputs on the lower adjacent slice have remaining weight $m - 1$.

Let $\gamma = 1 - 2\varepsilon$. A measurement with error at most ε on every input in these two slices also has both mixture errors at most ε , so Lemma 1.1 gives $D_T^{(2)} \geq \gamma$. Combining this with (3) yields

$$T \geq \frac{n+2}{4} \log n - \frac{n+2}{2} \log \frac{C}{\gamma} = \frac{N}{4} \log N - O_\varepsilon(N).$$

For the upper bound, apply Lemma 2.4 with $T = \lceil (N/4) \log(2N/\varepsilon^2) \rceil \leq (N/4) \log N + O_\varepsilon(N)$. Complementation preserves balancedness, so deciding whether the reconstructed representative has weight $N/2$ solves the promised problem with the same error.

When $N = 2$, the balanced ray is spanned by $(|1\rangle - |2\rangle)/\sqrt{2}$ and the unbalanced ray by $(|1\rangle + |2\rangle)/\sqrt{2}$. They are orthogonal, so one copy is sufficient; zero copies have minimax error $1/2$, so one copy is also necessary. This endpoint is absorbed by the asserted $O_\varepsilon(N)$ term. $\square$

The reconstruction POVM may be computationally expensive and deep. A circuit retaining classical inputs could repeatedly access them or invert an input-dependent preparation; those resources are specifically absent from this program-only model.

3 Adjacent counting from uniform-subset states

A second statistical model gives the receiver independent copies of

$$|X\rangle = |X|^{-1/2} \sum_{i \in X} |i\rangle$$

for an unknown nonempty subset X of a known universe. There is no membership oracle, no input-dependent reflection or state-preparation unitary, and no retained classical input. Distinguishing these resources is essential in quantum counting. See also the resource distinctions in [6].

We retain the notation $\rho_{M,s}^{(T)}$ and $T_{\text{opt}}(M, k, g, \varepsilon)$ from Section 1.2, and abbreviate

$$D_M(r, s; T) = D(\rho_{M,r}^{(T)}, \rho_{M,s}^{(T)}).$$

Theorem 3.1 (Sharp adjacent subset-count threshold). *Let $k \geq 2$, let the universe have size $M \geq 2k$, and fix $0 < \varepsilon < 1/2$. With an unrestricted receiver, the minimum number of independent states needed to distinguish every $|X| = k$ input from every $|X| = k - 1$ input with error at most ε is*

$$T_{\text{opt}}(M, k, 1, \varepsilon) = k \log k + O_\varepsilon(k),$$

uniformly over $M \geq 2k$. This is a state-copy theorem, not a supplied-classical-copy or shallow-decoder theorem.

Proof. For the lower bound, it suffices to restrict X to a fixed $2k$ -element subset of the original universe. Relabel those elements as the first $2k$ modes of a known $4k$ -dimensional space, and put

$$|u\rangle = (4k)^{-1/2} \sum_{i=1}^{4k} |i\rangle, \quad |\psi_X\rangle = |u\rangle - \sqrt{|X|/k} |X\rangle,$$

and let $R_u = I - 2|u\rangle\langle u|$. If $s = |X|$, then $\langle u|X\rangle = \sqrt{s/(4k)}$, and consequently $\|\psi_X\| = 1$. More explicitly, every mode in X has amplitude $-1/(2\sqrt{k})$ and every other mode has amplitude $1/(2\sqrt{k})$. Thus $|\psi_X\rangle$ is exactly a $4k$ -mode phase state whose last $2k$ signs are known positive. A fixed unitary combines those known modes into one anchor of amplitude $\sqrt{2k}$, so the comparison experiment of Section 2 has $n = h = 2k$.

Since $\langle u|\psi_X\rangle = 1 - s/(2k)$, direct calculation gives

$$R_u|\psi_X\rangle = (s/k - 1)|u\rangle - \sqrt{s/k}|X\rangle.$$

At $s = k$ this is exactly $-|X\rangle$. At $s = k - 1$ its squared overlap with $-|X\rangle$ is

$$(1 - 1/k)(1 + 1/(2k))^2 = 1 - \frac{3}{4k^2} - \frac{1}{4k^3}.$$

Put $\delta_k = 3/(4k^2) + 1/(4k^3) \leq 1/k^2$. The pure-state formula and Bernoulli's inequality give

$$D\left((R_u|\psi_X\rangle\langle\psi_X|R_u)^{\otimes T}, (|X\rangle\langle X|)^{\otimes T}\right) = \sqrt{1 - (1 - \delta_k)^T} \leq \sqrt{T\delta_k} \leq \frac{\sqrt{T}}{k}$$

when $s = k - 1$; the two states agree up to phase when $s = k$. Averaging over each slice, using convexity of the trace norm, and applying the triangle inequality therefore yields

$$D_{2k}(k, k - 1; T) \leq \frac{C_0 e^{T/(2k)}}{\sqrt{2k}} + \frac{\sqrt{T}}{k}, \quad C_0 = 2\sqrt{24}e, \quad T \leq 4k^2 - 2k, \quad (4)$$

where (2) with $n = h = 2k$ gives the first term. The direction of this inequality matters: it upper-bounds subset distinguishability. All embeddings, the anchor compression, and R_u are fixed and independent of X ; no input-dependent operation has been introduced. The auxiliary $4k$ -dimensional space is allowed even when the original universe has fewer than $4k$ elements.

Write $\gamma = 1 - 2\varepsilon > 0$. A worst-case test on the original universe also works on the restricted $2k$ -element subuniverse and induces errors at most ε on its two uniform mixtures. Lemma 1.1 therefore requires $D_{2k}(k, k - 1; T) \geq \gamma$. If $T > 2k \log k$, the desired lower bound is immediate. Otherwise $T \leq 4k^2 - 2k$ for $k \geq 2$, so (4) applies. For $k \geq K_\gamma := \lceil 64/\gamma^4 \rceil$, the inequalities $\log k \leq \sqrt{k}$ and $T \leq 2k \log k$ give $\sqrt{T}/k \leq \gamma/2$. Consequently

$$T \geq k \log k - 2k \log(\sqrt{2}C_0/\gamma).$$

For $2 \leq k < K_\gamma$, enlarge the coefficient of k to at least $\log K_\gamma$, making the same lower bound nonpositive. This proves it for all $k \geq 2$ with constants independent of M .

For the upper bound measure every program in the element basis. Output size k only after observing k distinct elements. On the smaller slice this never errs; on the larger slice its error is at most

$$k(1 - 1/k)^T \leq k e^{-T/k}.$$

Hence $T = \lceil k \log(k/\varepsilon) \rceil$ suffices. This receiver is allowed arbitrary classical storage and processing. $\square$

There is also a sharp threshold window in the following sense. If $\omega(k) \rightarrow \infty$ and $T \leq k \log k - k\omega(k)$, (4) makes the mixture trace distance tend to zero. No receiver's average success on these two uniform promises can then exceed one half by a fixed positive constant. If $T \geq k \log k + k\omega(k)$, the coupon receiver has worst-case error tending to zero. This conclusion concerns the leading threshold and its order- k window; it does not identify the limiting optimal error at a fixed offset.

3.1 A growing cardinality gap

The argument extends quantitatively when the two sizes are k and $k - g$. Write $T_{\text{opt}}(M, k, g, \varepsilon)$ for the required number of independent states in a universe of size M .

We first record an endpoint-sensitive refinement of Lemma 2.2. It will supply a polynomial rather than logarithmic lower bound once the gap reaches order $\sqrt{k}$.

Lemma 3.2 (Endpoint-sensitive product estimate). *In the notation of Lemma 2.2,*

$$D(\sigma_m, \sigma_{m-1}) \leq \min \left\{ 1, \sqrt{\frac{40}{n}} \sqrt{a^{-2} - 1} \right\}.$$

Proof. Use the joint total-spin and J_z distribution from Lemma 2.2, and retain the notation $\ell = m - j$, $Z = \ell - s$, and $L = \sqrt{\ell(\ell + 1)}$, with $Z/L = 0$ when $\ell = 0$. Put

$$c = \sqrt{\frac{m+1}{m}}, \quad H = L^2 - Z^2 - Z = (\ell - Z)(\ell + Z + 1) \geq 0.$$

The moments established in Lemma 2.2 give

$$\mathbb{E}H = \mathbb{E}(L^2 - Z^2) - \mathbb{E}Z = \frac{m(1-a)^2}{2}.$$

On $G = \{Z \geq ma/2\}$, in particular $\ell \geq Z > 0$, and

$$\begin{aligned} 1 - \frac{cZ}{L} &= 1 - \sqrt{\frac{(m+1)\ell}{m(\ell+1)}} + \frac{c(\ell-Z)}{L}, \\ 1 - \sqrt{\frac{(m+1)\ell}{m(\ell+1)}} &\leq \frac{m-\ell}{m(\ell+1)} \leq \frac{2(m-Z)}{m^2a}, \\ \frac{c(\ell-Z)}{L} &= \frac{cH}{L(\ell+Z+1)} \leq \frac{2\sqrt{2}H}{m^2a^2}. \end{aligned}$$

Here the second line uses $1 - \sqrt{q} \leq 1 - q$ and $\ell + 1 \geq ma/2$, while the third uses $c \leq \sqrt{2}$, $L \geq Z \geq ma/2$, and $\ell + Z + 1 \geq 2Z \geq ma$. On G^c the deficit is at most $1 + c < 3$. Chebyshev's inequality, as in Lemma 2.2, gives

$$\Pr(G^c) \leq \frac{2(1-a^2)}{ma^2}.$$

Since $\mathbb{E}(m - Z) = m(1 - a)$, averaging the preceding pointwise bounds yields

$$\begin{aligned} 1 - \mathbb{E}(cZ/L) &\leq \frac{2\mathbb{E}(m-Z)}{m^2a} + \frac{2\sqrt{2}\mathbb{E}H}{m^2a^2} + 3\Pr(G^c) \\ &\leq \frac{10(1-a^2)}{ma^2}. \end{aligned}$$

For the last line, relative to $(1 - a^2)/(ma^2)$ the three terms are at most 1, $\sqrt{2}$, and 6, respectively; their sum is less than 10. The phase-aligned block purifications in Lemma 2.2 have overlap $F \geq |\mathbb{E}(cZ/L)|$. Whenever the displayed nontrivial bound in the lemma is below one, one has $10(1 - a^2)/(ma^2) < 1/2$. Thus $\mathbb{E}(cZ/L) > 0$, and therefore

$$D(\sigma_m, \sigma_{m-1}) \leq \sqrt{1 - F^2} \leq \sqrt{2(1 - F)} \leq \sqrt{\frac{20(1 - a^2)}{ma^2}} = \sqrt{\frac{40}{n}} \sqrt{a^{-2} - 1}.$$

In the complementary parameter range, the trivial bound $D(\sigma_m, \sigma_{m-1}) \leq 1$ proves the claim. $\square$

Theorem 3.3 (Cardinality-gap dependence). *Fix $0 < \varepsilon < 1/2$. For integers $k \geq 2$, $M \geq 2k$, and $1 \leq g \leq \sqrt{k}$,*

$$T_{\text{opt}}(M, k, g, \varepsilon) = k \log(k/g^2) + O_\varepsilon(k),$$

uniformly in M and g . For every $M \geq 2k$ and $1 \leq g \leq k-1$,

$$T_{\text{opt}}(M, k, g, \varepsilon) = \Omega_\varepsilon(k \log(1 + k/g^2)).$$

When $M = 2k$, the latter bound is attained up to an ε -dependent constant for the full gap range. The leading-coefficient statement is informative when $g = o(\sqrt{k})$; at $g \asymp \sqrt{k}$ its additive $O(k)$ term must not be discarded.

Proof. First, mixture trace distance cannot increase when the universe is enlarged. Fix $r \leq q \leq M$. Choose a uniformly random q -subset $Q \subseteq [M]$ and a uniformly random bijection from $[q]$ to Q , and apply the induced isometry to all T registers. The same random bijection must be used on every register. After averaging over this randomness, the image of a uniformly random r -subset of $[q]$ is uniform among the r -subsets of $[M]$: for each fixed $R \subseteq [M]$ of size r , its probability is

$$\frac{\binom{M-r}{q-r}}{\binom{M}{q} \binom{q}{r}} = \binom{M}{r}^{-1},$$

where the equality is equivalent to $\binom{M}{q} \binom{q}{r} = \binom{M}{r} \binom{M-r}{q-r}$. This defines a common quantum channel taking $\rho_{q,r}^{(T)}$ to $\rho_{M,r}^{(T)}$, and contractivity gives

$$D_M(r, r-1; T) \leq D_q(r, r-1; T).$$

Apply this with $q = 2s$ for $s = k - g + 1, \dots, k$. The triangle inequality, (4) with k replaced by s , and the fact that both $s^{-1/2}e^{T/(2s)}$ and s^{-1} decrease with s give, for $s_0 = k - g + 1$,

$$D_M(k, k-g; T) \leq \frac{C_0 g}{\sqrt{2s_0}} e^{T/(2s_0)} + \frac{g\sqrt{T}}{s_0}, \quad T \leq 4s_0^2 - 2s_0. \quad (5)$$

Put $\gamma = 1 - 2\varepsilon$, $u = \log(k/g^2)$, and $C = 64/\gamma^2$. We claim $T \geq k(u - C)$ for $g \leq \sqrt{k}$. The claim is trivial when $u \leq C$. Otherwise suppose $T \leq k(u - C)$. Then $u > C \geq 64$, so k is certainly at least four, and $s_0 \geq k/2$. Moreover, $T \leq ku \leq k \log k \leq 4s_0^2 - 2s_0$, so (5) applies. For the last inequality, use $s_0 \geq k/2$, hence $4s_0^2 - 2s_0 \geq k^2 - k$, together with $\log k \leq k-1$. Its second term is at most

$$\frac{g\sqrt{T}}{s_0} \leq 2\sqrt{ug^2/k} = 2\sqrt{ue^{-u}} \leq 2\sqrt{Ce^{-C}},$$

because ue^{-u} decreases for $u \geq 1$. Its first term is at most

$$C_0 \exp\left(\frac{(k-s_0)u - Ck}{2s_0}\right) \leq C_0 e^{1-C/2},$$

using $k - s_0 = g - 1 \leq g$, $2s_0 \geq k$, $s_0 \leq k$, and $gu/k \leq \log k/\sqrt{k} \leq 2/e < 1$. Both error terms are smaller than $\gamma/3$. Indeed $C \geq 64$, $\gamma = 8/\sqrt{C}$, $C_0 = 2\sqrt{24e} < 30$, and

$$C_0 e^{1-C/2} \leq \frac{720}{C^2} < \frac{8}{3\sqrt{C}}, \quad 2\sqrt{Ce^{-C}} \leq \frac{2\sqrt{6}}{C} < \frac{8}{3\sqrt{C}},$$

where $e^{C/2} \geq C^2/8$ and $e^C \geq C^3/6$ were used. A worst-case test would require $D_M(k, k-g; T) \geq \gamma$ by Lemma 1.1, so (5) gives a contradiction. This proves the uniform lower estimate $T \geq k \log(k/g^2) - Ck$.

For the upper bound, measure in the element basis and let D_s count distinct observed elements on a size- s input. If I_i indicates that element i is missing, then for distinct i, j , $\text{Cov}(I_i, I_j) = (1 - 2/s)^T - (1 - 1/s)^{2T} \leq 0$, since $0 \leq 1 - 2/s \leq (1 - 1/s)^2$. Hence

$$f_T(s) := \mathbb{E}D_s = s(1 - (1 - 1/s)^T), \quad \text{Var } D_s \leq ke^{-T/k} \quad (s \leq k).$$

Here $\text{Var } D_s = \text{Var}(\sum_i I_i) \leq s(1 - 1/s)^T \leq se^{-T/s} \leq ke^{-T/k}$; the last function is increasing in s . For real $2 \leq s \leq k$ and $T \geq 4k$, differentiation gives

$$f'_T(s) = 1 - (1 - 1/s)^T - (T/s)(1 - 1/s)^{T-1} \geq 1 - (1 + 2T/s)e^{-T/s} \geq \frac{1}{2}.$$

The first inequality uses $(1 - 1/s)^T \leq e^{-T/s}$ and $(1 - 1/s)^{T-1} \leq 2e^{-T/s}$; the second follows because $(1 + 2x)e^{-x} < 1/2$ for $x \geq 4$. Integrating the derivative shows that the two promised means differ by at least $g/2$ when $k - g \geq 2$. If the smaller size is one, the same conclusion follows directly from $f_T(k) - f_T(1) \geq k(1 - e^{-4}) - 1 \geq (k - 1)/2 = g/2$. Use the threshold $\tau = (f_T(k) + f_T(k - g))/2$. Under either promise, an error requires a deviation of at least $g/4$, so Chebyshev's inequality bounds each error by $16ke^{-T/k}/g^2$. Consequently

$$T = \left\lceil k \log(k/g^2) + k \max\{4, \log(16/\varepsilon)\} \right\rceil$$

suffices for $g \leq \sqrt{k}$, in every universe: the displayed choice has $T \geq 4k$ and makes the last error bound at most ε .

For a second lower estimate, assume $T \geq 1$; zero copies have error $1/2$ and cannot meet the stated error criterion. In the size- s adjacent phase experiment, take $n = h = 2s$ and coherently Poissonize with mean $2T$. If $Y \sim \text{Pois}(2T)$, Cantelli's inequality gives

$$\Pr(Y \geq T) \geq 1 - \frac{2T}{2T + T^2} \geq \frac{1}{3}.$$

Let $D_{\text{ph},s}(T)$ denote the trace distance between the central and adjacent mixtures in this anchored phase experiment. The dephasing comparison in the proof of Lemma 2.3, now with Poisson mean $2T$, reads

$$D_{\text{ph},s}(T) \Pr(Y \geq T) \leq D(\sigma_s, \sigma_{s-1}),$$

where the parity-product parameter is $a = e^{-T/s}$. Lemma 3.2 therefore gives

$$D_{\text{ph},s}(T) \leq \frac{3\sqrt{20}}{\sqrt{s}} \sqrt{e^{2T/s} - 1}.$$

The phase-to-subset comparison used in Theorem 3.1 adds at most $\sqrt{T}/s$ on the smaller slice. Consequently

$$D_{2s}(s, s-1; T) \leq \frac{3\sqrt{20}}{\sqrt{s}} \sqrt{e^{2T/s} - 1} + \frac{\sqrt{T}}{s} \leq \frac{40\sqrt{T}}{s} \quad (T \leq s).$$

For the last inequality, apply convexity to obtain $e^{2x} - 1 \leq (e^2 - 1)x < 8x$ on $0 \leq x \leq 1$; then $3\sqrt{160} + 1 < 40$. Universe enlargement and telescoping therefore imply that any successful test must satisfy

$$T \geq \min \left\{ s_0, \frac{\gamma^2 s_0^2}{1600g^2} \right\}.$$

Indeed, if $T \leq s_0$, the preceding adjacent bound and the same universe enlargement and triangle-inequality telescope give $\gamma \leq 40g\sqrt{T}/s_0$; if $T > s_0$, the first alternative already holds.

We now verify that the two lower estimates give the claimed order throughout the gap range. If $g \leq \sqrt{k}$, then $s_0 \geq k/2$ and the last display gives $T \geq \gamma^2 k/6400$. When $u = \log(k/g^2) \geq 2C$, the earlier estimate $T \geq k(u-C)$ is $\Omega(ku)$; when $0 \leq u < 2C$, the new $\Omega_\varepsilon(k)$ estimate is $\Omega_\varepsilon(k \log(1+e^u))$. If $\sqrt{k} < g \leq k/2$, then the polynomial estimate gives

$$T \geq \frac{\gamma^2 k^2}{6400g^2} = \Omega_\varepsilon(k \log(1+k/g^2)),$$

using $\log(1+x) \leq x$. Finally, if $g > k/2$, then $k \log(1+k/g^2) \leq k^2/g^2 < 4$, while at least one copy is necessary. This proves the all-gap lower bound.

Finally let $M = 2k$ and write $|u_{2k}\rangle = (2k)^{-1/2} \sum_{i=1}^{2k} |i\rangle$. Measuring the effect $|u_{2k}\rangle\langle u_{2k}|$ on $|X\rangle$ produces a Bernoulli variable of mean $|X|/(2k)$. The two promised means differ by $g/(2k)$. Thresholding the empirical mean halfway between them and applying Hoeffding's inequality bounds either error by $\exp(-Tg^2/(8k^2))$. Thus, for $g > \sqrt{k}$,

$$T = \left\lceil \frac{8k^2}{g^2} \log \frac{1}{\varepsilon} \right\rceil$$

suffices. For $g \leq \sqrt{k}$, the distinct-count estimator above uses $O_\varepsilon(k \log(1+k/g^2))$ copies; for $g > \sqrt{k}$, the last display has this same order because $x/2 \leq \log(1+x) \leq x$ for $0 \leq x \leq 1$. This completes the matching half-density upper bound. No matching all-gap upper bound for arbitrarily larger M is asserted. $\square$

4 The adjacent critical window and collective advantage

Theorem 3.1 identifies the leading threshold. We now determine its fixed-offset error profile when the universe has size exactly $2k$. The resources remain independent subset states and an unrestricted receiver, with no classical input, membership oracle, or input-dependent preparation unitary. In particular, this section provides no shallow decoder or supplied-classical-copy bound.

We use modified Bessel functions with the normalization

$$I_j(z) = \frac{1}{2\pi} \int_0^{2\pi} e^{z \cos \theta} e^{ij\theta} d\theta.$$

For every integer $j \geq 0$ and $t \geq 0$,

$$I_j(2t) = \sum_{r=0}^{\infty} \frac{t^{2r+j}}{r!(r+j)!}. \quad (6)$$

To verify the series, expand $e^{2t \cos \theta} = e^{te^{i\theta}} e^{te^{-i\theta}}$. The two absolutely and uniformly convergent exponential series may be multiplied and integrated termwise. The surviving Fourier indices differ by j , giving (6). In particular all series used below have nonnegative terms for $t \geq 0$.

Theorem 4.1 (Critical adjacent-subset profile). *Let $\rho_s^{(T)}$ be the uniform mixture of $(|X\rangle\langle X|)^{\otimes T}$ over all s -subsets of $[2k]$. Fix $c \in \mathbb{R}$. For all sufficiently large k , put*

$$T_k = \lfloor k \log k + ck \rfloor, \quad \lambda = e^{-c}.$$

With $D(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$,

$$\lim_{k \rightarrow \infty} D(\rho_k^{(T_k)}, \rho_{k-1}^{(T_k)}) = \Phi(\lambda) := e^{-\lambda} \int_0^\infty 2re^{-r^2} \sqrt{I_0(2\sqrt{\lambda}r)^2 - I_1(2\sqrt{\lambda}r)^2} dr, \quad (7)$$

Both the optimal equiprior average error, with uniform priors within the two cardinalities, and the optimal worst-case error over all promised subsets converge to

$$\frac{1 - \Phi(\lambda)}{2}.$$

These two error criteria are not asserted to agree at finite k .

Proof. Subset comparison. Use the fixed $4k$ -mode embedding from Theorem 3.1: with $|u\rangle$ uniform on those modes, set

$$|\psi_X\rangle = |u\rangle - \sqrt{s/k}|X\rangle, \quad s = |X|, \quad R_u = I - 2|u\rangle\langle u|.$$

The last $2k$ modes are known positive anchors. At $s = k$, $R_u|\psi_X\rangle = -|X\rangle$ exactly; at $s = k - 1$ its squared overlap with $-|X\rangle$ is

$$(1 - 1/k)(1 + 1/(2k))^2 = 1 - \frac{3}{4k^2} - \frac{1}{4k^3}.$$

Let $D_{\text{set}}(T) = D(\rho_k^{(T)}, \rho_{k-1}^{(T)})$, and let $D_{\text{phase}}(T)$ be the trace distance between the corresponding uniform mixtures of the anchored phase states. The T -copy pure-state distance on the smaller slice is at most $\sqrt{T}/k$. Taking $T = T_k$, convexity and the reverse triangle inequality give the two-sided comparison

$$|D_{\text{set}}(T_k) - D_{\text{phase}}(T_k)| \leq \sqrt{T_k}/k = o(1). \quad (8)$$

Only the smaller class is perturbed. This estimate holds separately and uniformly for every input in that class, which will also allow transfer of worst-case measurements.

Canonical fixed-count and Poisson families. For the remainder of the proof set $T = T_k$ and write $n = 2k$. Group the anchored phase tensor power into its input-independent occupation-parity fibers, as in Section 2. A fixed isometry on their span gives the canonical vector

$$|\Omega_x^{\text{fix}}\rangle = \sum_{S \subseteq [n]} \sqrt{p_T(S)} (-1)^{x \cdot S} |S\rangle.$$

The parity walk has idle probability $1/2$ and flip probability $1/(2n)$ per coordinate. Its character eigenvalue on a j -set is $1 - j/n$. Coherently replacing the count by $\text{Pois}(T)$ gives an analogous canonical vector with probabilities p_P and eigenvalues $e^{-Tj/n}$. The count information remains inside the input-independent fibers; it is not discarded to assert this isometry. Poisson splitting makes p_P a product distribution with coordinate bias $a_k = e^{-T/n}$.

Parseval, with unnormalized probability Fourier coefficients, gives

$$2^n \sum_S (p_T(S) - p_P(S))^2 = \sum_{j=0}^n \binom{n}{j} [(1 - j/n)^T - e^{-Tj/n}]^2 = o(1). \quad (9)$$

Here is a summable justification of the limit, including the dependence of $T = T_k$ and $n = 2k$. For a fixed positive integer j and all sufficiently large k ,

$$T[\log(1 - j/n) + j/n] = -\frac{Tj^2}{2n^2} + O_j(T/n^3) = O_j(\log k/k) = o(1).$$

Factoring $e^{-Tj/n}$ from the bracket in (9), its remaining factor therefore tends to zero. The $j = 0$ term vanishes. For all indices $1 \leq j \leq n$, including $j = n$, positivity gives $0 \leq (1 - j/n)^T \leq e^{-Tj/n}$ once $T > 0$. Every summand is consequently at most

$$\binom{n}{j} e^{-2Tj/n} \leq \frac{(ne^{-2T/n})^j}{j!}, \quad ne^{-2T/n} \rightarrow 2\lambda.$$

Choose a finite B exceeding $ne^{-2T/n}$ for all sufficiently large k . For each fixed j , the preceding factorization and $\binom{n}{j} e^{-2Tj/n} \leq B^j/j!$ show that the summand tends to zero. After extending the summands by zero for $j > n$, dominated convergence applies to the summable majorant $B^j/j!$, proving (9). The nonnegative spectrum from the heavy anchors is used here.

Cauchy–Schwarz gives explicitly

$$\text{TV}(p_T, p_P) = \frac{1}{2} \sum_S |p_T(S) - p_P(S)| \leq \frac{1}{2} \sqrt{2^n \sum_S (p_T(S) - p_P(S))^2} = o(1).$$

Their affinity is at least one minus total variation, since $\sqrt{uv} \geq \min\{u, v\}$ for nonnegative u, v . The corresponding canonical pure vectors thus have distance at most the square root of twice that variation, uniformly in x : their inner product is exactly this affinity because the common input signs cancel. Both entire individual-state families are uniformly close. The Poisson canonical family is the product encoding

$$\bigotimes_{i=1}^{2k} \left(\sqrt{\frac{1+a_k}{2}} |0\rangle + (-1)^{x_i} \sqrt{\frac{1-a_k}{2}} |1\rangle \right), \quad ka_k^2 \rightarrow \lambda.$$

Let D_{prod} denote the trace distance between the uniform weight- k and weight- $(k-1)$ mixtures of this product family. The uniform individual-state comparison above gives $|D_{\text{prod}} - D_{\text{phase}}(T_k)| = o(1)$. Moreover $a_k \rightarrow 0$ and $\text{atanh}(a_k)/a_k \rightarrow 1$, so for $\beta_k = \text{atanh}(a_k)$ one has $\beta_k \sqrt{k} \rightarrow \sqrt{\lambda}$.

The exact spin formula. Use the harmonic-chain decomposition proved in Section 2, now indexed by spin $l = 0, \dots, k$. Its multiplicity and normalized weight are

$$d_l = \binom{2k}{k-l} - \binom{2k}{k-l-1}, \quad w_{k,l} = \frac{d_l}{\binom{2k}{k}}, \quad b_k = \frac{k+1}{k}.$$

For $a = a_k$ and $\beta = \text{atanh}(a)$, the local filter $F = \text{diag}(\sqrt{1+a}, \sqrt{1-a})$ satisfies

$$F^{\otimes 2k} = (1-a^2)^{k/2} e^{\beta J_z}.$$

After the collective Hadamard rotation, the two slice columns are standard J_z -basis vectors of magnetic numbers 0 and 1, while their squared-filter Gram operator is $(1-a^2)^k M_l$, where $M_l = e^{2\beta J_x^{(l)}}$. Thus the filter generator is tridiagonal in the chosen basis; we are not simultaneously diagonalizing J_x and its columns. For $l \geq 1$ put

$$A_l = \langle 0 | M_l | 0 \rangle, \quad B_l = \langle 1 | M_l | 1 \rangle, \quad C_l = \langle 0 | M_l | 1 \rangle.$$

Let v_l and z_l denote the central and neighboring rank-one block vectors in one multiplicity copy. The normalized slice factors in the block decomposition of Lemma 2.2 give

$$\|v_l\|^2 = \frac{(1-a^2)^k}{\binom{2k}{k}} A_l, \quad \|z_l\|^2 = \frac{(1-a^2)^k}{\binom{2k}{k}} b_k B_l, \quad \langle v_l, z_l \rangle = \frac{(1-a^2)^k}{\binom{2k}{k}} \sqrt{b_k} C_l.$$

There are d_l identical multiplicity copies. At $l = 0$ only the central column exists, and its squared norm is $(1 - a^2)^k / \binom{2k}{k}$. The rank-one difference identity $\|vv^\dagger - zz^\dagger\|_1 = \sqrt{(\|v\|^2 + \|z\|^2)^2 - 4|\langle v, z \rangle|^2}$ then gives exactly

$$D_{\text{prod}} = (1 - a^2)^k \left[\frac{w_{k,0}}{2} + \sum_{l=1}^k \frac{w_{k,l}}{2} \sqrt{(A_l + b_k B_l)^2 - 4b_k |C_l|^2} \right]. \quad (10)$$

The neighboring slice has no $l = 0$ column. The prefactor comes from all $2k$ filtered qubits and is independent of l . No commutativity of the two mixtures is assumed.

Uniform annuli and both tails. For $l \geq 1$, embed $K_l = J_x^{(l)}/l$ into $\ell^2(\mathbb{Z})$, zero off $[-l, l]$. Its norm is one and its adjacent entries are

$$\langle q+1 | K_l | q \rangle = \frac{\sqrt{l(l+1) - q(q+1)}}{2l} \longrightarrow \frac{1}{2}.$$

Here are quantitative details of the uniformity needed for the exponential. Let $K = (S + S^*)/2$, where $S|q\rangle = |q+1\rangle$ is the bilateral shift. Fix an integer $P \geq 0$. A path contributing to a matrix entry of K_l^P or K^P between coordinates $0, 1$, for $p \leq P$, stays within $[-P-1, P+1]$. On this finite set the displayed coefficients converge uniformly as $l \rightarrow \infty$, and the boundary of $[-l, l]$ is absent once $l > P+1$. If the largest coefficient discrepancy there is $\delta_{P,l}$, expanding the products along at most 2^P paths bounds the entry difference by $p2^P \delta_{P,l}$, since the individual coefficients have absolute value at most one. Thus every polynomial of degree at most P has uniformly convergent relevant entries when $l \geq \delta\sqrt{k}$.

On the annulus $\delta \leq l/\sqrt{k} \leq R$, set $z_{k,l} = 2\beta_k l$. The scalars are bounded by a fixed Z for all sufficiently large k , and

$$\sup_{l \leq R\sqrt{k}} |z_{k,l} - 2\sqrt{\lambda} l/\sqrt{k}| \leq 2R|\beta_k \sqrt{k} - \sqrt{\lambda}| \longrightarrow 0.$$

Because $\|K_l\|, \|K\| \leq 1$, truncating either exponential series above degree P has norm error at most $e^Z Z^{P+1}/(P+1)!$, independently of k, l . First choose P to make this error small; then use the finite-path convergence and scalar convergence just proved. This gives, uniformly on the compact annulus,

$$A_l, B_l \longrightarrow I_0(2\sqrt{\lambda}r), \quad C_l \longrightarrow I_1(2\sqrt{\lambda}r), \quad r = l/\sqrt{k}.$$

The Fourier transform of $\ell^2(\mathbb{Z})$ sends K to multiplication by $\cos\theta$, so its exponential entries are exactly the displayed Fourier integrals. This proves the asserted Bessel-entry limits without an unquantified interchange of an operator limit and an infinite series.

Set $b_{k,l} = \binom{2k}{k-l}/\binom{2k}{k}$ and $b_{k,k+1} = 0$. The weights satisfy $w_{k,l} = b_{k,l} - b_{k,l+1}$, so they form a probability law with survival $\Pr(L \geq l) = b_{k,l}$. Explicitly,

$$b_{k,l} = \prod_{q=0}^{l-1} \frac{k-q}{k+q+1} = \prod_{q=0}^{l-1} \left(1 - \frac{2q+1}{k+q+1} \right).$$

Taylor expansion of the logarithm, uniformly for $l = O(\sqrt{k})$, and the bound $1 - x \leq e^{-x}$ give, respectively,

$$\log b_{k,l} = -l^2/k + O(l^3/k^2) \quad (l = O(\sqrt{k})), \quad b_{k,l} \leq e^{-l^2/(k+l)} \leq e^{-l^2/(2k)}.$$

Consequently $R_k = L/\sqrt{k}$ converges in distribution to a nonnegative random variable with density $2re^{-r^2}$, and globally $\Pr(R_k \geq r) \leq e^{-r^2/2}$ for $r \geq 0$.

Define $f_k(0) = 1/2$ and, for $l \geq 1$, set

$$f_k(l) = \frac{1}{2} \sqrt{(A_l + b_k B_l)^2 - 4b_k |C_l|^2}.$$

Since $A_l, B_l, |C_l| \leq \|M_l\| = e^{2\beta_k l}$ and $1 < b_k \leq 3/2$ for $k \geq 2$, positivity gives $0 \leq f_k(l) \leq (5/4)e^{2\beta_k l} \leq C_1 e^{C_2 R_k}$ for constants depending only on c . For each fixed $C > 0$,

$$\mathbb{E}[e^{CR_k} \mathbf{1}_{R_k > R}] \leq e^{CR - R^2/2} + C \int_R^\infty e^{Cu - u^2/2} du \longrightarrow 0$$

uniformly in k . At zero, f_k is bounded on $R_k \leq \delta$, whose limiting mass is $1 - e^{-\delta^2}$; also $w_{k,0} = 1/(k+1)$. These estimates control both omitted ends of a compact annulus. On that annulus f_k converges uniformly to $\sqrt{I_0(2\sqrt{\lambda}r)^2 - I_1(2\sqrt{\lambda}r)^2}$. The radicand is nonnegative: it is the determinant of the positive-semidefinite 2×2 principal matrix obtained as the limit of M_l on coordinates 0, 1. Weak convergence, the two tail bounds, and $(1 - a_k^2)^k \rightarrow e^{-\lambda}$ prove (7). They also prove finiteness of the integral; its value lies in $[0, 1]$ as a trace-distance limit.

Minimax by midpoint block masses. The two product mixtures have spin-block masses

$$\mu_l = (1 - a^2)^k w_{k,l} A_l, \quad \nu_l = (1 - a^2)^k w_{k,l} b_k B_l \quad (l \geq 1),$$

with $\mu_0 = (1 - a^2)^k w_{k,0}$ and $\nu_0 = 0$. Each sequence sums to one. The same annular limits and exponential domination show $L_k := \sum_l |\mu_l - \nu_l| \rightarrow 0$. More explicitly,

$$L_k = (1 - a^2)^k \left[w_{k,0} + \sum_{l=1}^k w_{k,l} |A_l - b_k B_l| \right].$$

On every compact annulus the summand after $w_{k,l}$ tends uniformly to zero because $A_l, B_l \rightarrow I_0$ and $b_k \rightarrow 1$. Its small- l and large- l contributions tend uniformly to zero by the two tail estimates used for (7). This proves the asserted convergence rather than assuming it. Replace both block masses by $\omega_l = (\mu_l + \nu_l)/2$, retaining their normalized rank-one chain vectors and the normalized identity on multiplicity space. In the absent neighboring $l = 0$ block choose the central vector. Each modified state is at distance $L_k/4$ from its original.

In each block an involutive unitary exchanges the two normalized vectors, after phase adjustment. Their direct sum exchanges the modified pair. Symmetrizing an optimal equiprior POVM with its exchanged, outcome-reversed version equalizes the two errors without changing their mean. Thus the modified pair's minimax and average errors agree. If $\delta_k = L_k/4$, each original state is within trace distance δ_k of its modified counterpart. Lemma 1.1 changes either optimal error criterion by at most δ_k , and hence

$$e_{\text{av}} \leq e_{\text{mm}} \leq e_{\text{av}} + L_k/2.$$

The fixed isometries and uniform individual-state comparisons above transfer this limiting statement to the subset experiment. Finally permutation-averaging a POVM makes its error constant within each cardinality orbit without changing either class-average error. Hence two-mixture minimax is exactly worst-case-over-subsets minimax. Both limits claimed in the theorem follow. $\square$

Corollary 4.2 (Element-basis baseline and strict collective advantage). *If every supplied subset state is first measured in the element basis, followed by arbitrary classical processing, the optimal equiprior average error at T_k converges to*

$$\frac{1 - \max_{j \geq 0} e^{-\lambda} \lambda^j / j!}{2}.$$

For $0 < \lambda \leq 1$, the quantum worst-case limit is strictly smaller, since

$$\Phi(\lambda) \geq e^{-\lambda}(1 + \lambda/2) > e^{-\lambda}.$$

In particular at $c = 0$, its error is at most $(1 - 3/(2e))/2$ in the limit, whereas the classical element-measurement average optimum tends to $(1 - 1/e)/2$.

Proof. Under the uniform prior on the size- s slice, an ordered transcript y with D distinct labels has exact probability

$$P_s(y) = \frac{\binom{2k-D}{s-D}}{\binom{2k}{s} s^{T_k}}$$

when $D \leq s$, and zero otherwise. Conditional on D , the transcript is uniform on a set independent of s . Thus D is sufficient, including for total variation of the two experiments.

For the unseen count $U_s = s - D$, its falling-factorial moments satisfy

$$\mathbb{E}(U_s)_r = (s) * r(1 - r/s)^{T_k} \longrightarrow \lambda^r \quad (s = k, k-1), \quad 0 \leq \mathbb{E}(U_s)r \leq [se^{-T_k/s}]^r \leq B^r,$$

where $B < \infty$ is independent of all sufficiently large k . For fixed r , the moment limit follows from $\log s - T_k/s \rightarrow -c$ and an error $O_r(1/s + T_k/s^2) = o(1)$ in its logarithm. Moments with $r > s$ are zero. For fixed j , finite inclusion-exclusion gives

$$\Pr(U_s = j) = \frac{1}{j!} \sum r \geq 0 \frac{(-1)^r}{r!} \mathbb{E}(U_s) * j + r.$$

The summable majorant $B^{j+r}/(j!r!)$ proves convergence to the Poisson mass. Pointwise convergence to a probability mass function implies total variation convergence by a finite-cutoff argument. Hence the common statistic $k - D$ has limiting laws $\text{Pois}(\lambda)$ and $1 + \text{Pois}(\lambda)$. The total variation between a unimodal mass sequence $p_j = e^{-\lambda} \lambda^j / j!$ and its shift is $\frac{1}{2} \sum_{j \geq 0} |p_j - p_{j-1}| = \max_j p_j$, with $p_{-1} = 0$. This proves the classical formula. It is an average optimum, not an assertion of equal classical conditional errors.

Finally, Vandermonde's identity applied to the absolutely convergent Bessel series yields

$$I_0(2z)^2 - I_1(2z)^2 = \sum_{m \geq 0} \frac{1}{m+1} \binom{2m}{m} \frac{z^{2m}}{(m!)^2}.$$

Indeed, the coefficient of z^{2m} in $I_0(2z)^2$ is $\binom{2m}{m}/(m!)^2$, and the corresponding coefficient in $I_1(2z)^2$ is $m/(m+1)$ times this quantity. Its nonnegative terms through $m = 2$ are $1 + z^2 + z^4/2$, so its square root is at least $1 + z^2/2$ for $z \geq 0$. Substitute $z = \sqrt{\lambda}r$ into (7); the Rayleigh law has second moment one. For $0 < \lambda \leq 1$, the largest Poisson mass is $e^{-\lambda}$, proving strict advantage. This compares a collective state receiver with the specified element-basis measurement restriction, not with unrestricted classical input access or arbitrary classical encodings. $\square$

4.1 Every fixed integer gap

Corollary 4.3 (Fixed-gap critical profiles). *Fix a positive integer g and a real c . Let $k \rightarrow \infty$ with $k > g$, in the universe $[2k]$, and, for all sufficiently large k , set $T_k = \lfloor k \log k + ck \rfloor$. Put $\lambda = e^{-c}$. Then*

$$D(\rho_k^{(T_k)}, \rho_{k-g}^{(T_k)}) \longrightarrow \Phi_g(\lambda) = e^{-\lambda} \int_0^\infty 2re^{-r^2} \sqrt{I_0(2\sqrt{\lambda}r)^2 - I_g(2\sqrt{\lambda}r)^2} dr.$$

Optimal equiprior average error and optimal worst-case error both tend to $(1 - \Phi_g(\lambda))/2$. For the element-basis measurement restriction, the optimal average error tends to $(1 - C_g(\lambda))/2$, where

$$C_g(\lambda) = \max_{h \geq 0} \Pr\{h - g + 1 \leq Z \leq h\}, \quad Z \sim \text{Pois}(\lambda).$$

Neither limit is asserted uniformly for growing g .

Proof. Here are all changes needed in the preceding proof. In the subset comparison, write $t = g/k$. The smaller-slice squared overlap becomes

$$(1 - t)(1 + t/2)^2 = 1 - 3t^2/4 - t^3/4 \geq 1 - t^2.$$

The individual T_k -copy comparison error is at most $g\sqrt{T_k}/k = o(1)$. The canonical fixed-count/Poisson comparison is uniform in the input string and hence independent of its promised weight; it applies unchanged to both slices.

In the product experiment, use the fixed magnetic columns $0, g$ and put

$$\gamma_{k,g} = \frac{\binom{2k}{k}}{\binom{2k}{k-g}}, \quad A_l = \langle 0|M_l|0 \rangle, \quad B_l = \langle g|M_l|g \rangle, \quad C_l = \langle 0|M_l|g \rangle,$$

where the last two entries exist for $l \geq g$. The exact distance is

$$\frac{(1 - a^2)^k}{2} \left[\sum_{l=0}^{g-1} w_{k,l} A_l + \sum_{l=g}^k w_{k,l} \sqrt{(A_l + \gamma_{k,g} B_l)^2 - 4\gamma_{k,g} |C_l|^2} \right].$$

In particular the absent-slice terms retain A_l ; they are not all identically one. Since $A_l \leq e^{2\beta l}$ and $(1 - a^2)^k \leq 1$, their contribution to the distance is at most

$$\frac{1}{2} e^{2\beta g} \left(1 - \frac{\binom{2k}{k-g}}{\binom{2k}{k}} \right) = O_g(k^{-1}).$$

Here the last estimate follows directly from the fixed- g product formula for the binomial ratio, while $\beta = O(k^{-1/2})$. For fixed g , $\gamma_{k,g} \rightarrow 1$. The same finite-power argument for the tridiagonal matrices now uses fixed indices $0, g$, giving annular limits I_0, I_0, I_g . The Rayleigh multiplicities and Gaussian tails are unchanged; the bounded ratio $\gamma_{k,g}$ preserves the common exponential envelope. This proves the integral limit.

The total two-class spin-block mass discrepancy is explicitly

$$(1 - a^2)^k \left[\sum_{l=0}^{g-1} w_{k,l} A_l + \sum_{l=g}^k w_{k,l} |A_l - \gamma_{k,g} B_l| \right].$$

It tends to zero: on compact annuli the diagonal entries have the same I_0 limit and $\gamma_{k,g} \rightarrow 1$; the two tails remain controlled by the common exponential envelope; and the missing $l < g$ mass vanishes by the preceding estimate. Midpoint mass replacement, blockwise swaps, uniform individual-state comparison, and permutation symmetrization therefore give the same minimax conclusion, without assuming finite- k equality of errors.

Finally $\log s - T_k/s \rightarrow -c$ for both $s = k, k - g$. The bounded-factorial-moment and inclusion-exclusion proof of total-variation Poisson convergence thus applies to both missing counts. The sufficient statistic $k - D$ tends to Z versus $g + Z$. Writing $p_j = e^{-\lambda} \lambda^j / j!$ and $p_j = 0$ for $j < 0$, their total variation is $\frac{1}{2} \sum_j |p_j - p_{j-g}|$. For $j \geq g$, $p_j / p_{j-g} = \lambda^g / (j)_g$ strictly decreases, so the signed differences have only one crossing. The sums of the positive differences telescope to the largest mass of a window of g consecutive integers, which is $C_g(\lambda)$. $\square$

4.2 Strict collective advantage at every finite offset

Corollary 4.4. *For every fixed integer $g \geq 1$ and every finite $\lambda > 0$,*

$$\Phi_g(\lambda) > C_g(\lambda).$$

Consequently, in the fixed-gap critical window, the limiting optimal quantum worst-case error is strictly smaller than even the optimal element-basis receiver's equiprior average error. The receiver and input-access qualifications of Corollary 4.3 remain in force.

Proof. For $t > 0$ define two auxiliary vectors in $\ell^2(\mathbb{N}_0)$ by

$$u_j = \frac{t^j}{j!}, \quad v_j = \begin{cases} 0, & j < g, \\ t^{j-g}/(j-g)!, & j \geq g. \end{cases}$$

Their squared norms are both $I_0(2t)$, their inner product is $I_g(2t)$, and they are linearly independent since $u_0 = 1, v_0 = 0$. Thus $\Delta_t = uu^* - vv^*$ is a trace-zero rank-two trace-class operator with

$$\frac{1}{2}\|\Delta_t\|_1 = \sqrt{I_0(2t)^2 - I_g(2t)^2}.$$

Let $\mathcal{D}$ remove off-diagonal entries in this auxiliary basis. We prove strict trace-norm loss, not just contractivity. Let E be the diagonal projection onto coordinates with positive diagonal entry of Δ_t . It attains $\text{Tr}(E\Delta_t) = \|\mathcal{D}(\Delta_t)\|_1/2$. At coordinate zero the diagonal is positive. For some sufficiently large j , it is negative, since

$$v_j/u_j = (j)_g/t^g \longrightarrow \infty.$$

But $(\Delta_t)_{0j} = u_j > 0$, so $[E, \Delta_t]_{0j} = u_j \neq 0$.

An effect attaining $\|\Delta_t\|_1/2$ must commute with Δ_t . Indeed, write its two nonzero eigenvalues as $a, -a$ with unit eigenvectors w_+, w_- . For $0 \leq Q \leq I$, attaining the maximum requires $\langle w_+, Qw_+ \rangle = 1$ and $\langle w_-, Qw_- \rangle = 0$. Positivity of Q and $I - Q$ then implies $Qw_+ = w_+$ and $Qw_- = 0$. Self-adjointness prevents couplings from these vectors to the kernel, so $[Q, \Delta_t] = 0$. The displayed noncommutation of E proves

$$\frac{1}{2}\|\mathcal{D}(\Delta_t)\|_1 < \frac{1}{2}\|\Delta_t\|_1 \quad (t > 0).$$

Now set $t = \sqrt{\lambda}r$ and integrate with weight $e^{-\lambda}2re^{-r^2}dr$. Both half norms are at most $I_0(2t)$, and

$$e^{-\lambda} \int_0^\infty 2re^{-r^2} I_0(2\sqrt{\lambda}r) dr = 1$$

by termwise integration of its nonnegative series. The gap is measurable (in fact continuous) and positive for every $r > 0$, so its integral is strictly positive. Individual diagonal entries integrate to

$$e^{-\lambda} \int_0^\infty 2re^{-r^2} (u_j^2 - v_j^2) dr = p_j - p_{j-g}, \quad p_j = e^{-\lambda} \lambda^j / j!,$$

where negative indices mean zero. Tonelli and the triangle inequality therefore give

$$\Phi_g(\lambda) > e^{-\lambda} \int_0^\infty 2re^{-r^2} \frac{1}{2} \|\mathcal{D}(\Delta_{\sqrt{\lambda}r})\| * 1 dr \geq \frac{1}{2} \sum *j \geq 0 |p_j - p_{j-g}| = C_g(\lambda).$$

These auxiliary vectors establish an analytic comparison of the two limiting error formulas; they do not presume a physical constant-depth dephasing or collective decoder. $\square$

Use of AI: We prepared a list of problems and proof strategies and techniques related to Copy-complexity (how many copies of a classical input make shallow quantum computation powerful). We fed this to multiple AI tools (Claude, ChatGPT and two open-weight models with open-source harnesses) and prompted these AI tools to improve the existing bounds. When an AI tool was hallucinating, stuck, or going in the wrong direction, we stopped that AI tool, corrected the intermediate .md files, added more specific new ideas and switched to a different AI tool with more detailed fresh prompts. Equations written by one AI tool were critiqued, verified and expanded by a completely different AI tool in an iterative manner, till we were satisfied with their correctness and details. The author is solely responsible for the correctness of the proofs.

References

- [1] Daniel Grier, Jackson Morris and Kewen Wu. QAC^0 Contains TC^0 (with Many Copies of the Input). [arXiv:2601.03243](#) (2026).
- [2] Gilles Brassard, Peter Høyer and Alain Tapp. Quantum counting. In *Automata, Languages and Programming (ICALP 1998)*, Lecture Notes in Computer Science **1443**, 820–831 (1998). [doi:10.1007/BFb0055105](#).
- [3] Ashwin Nayak and Felix Wu. The quantum query complexity of approximating the median and related statistics. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, 384–393 (1999). [doi:10.1145/301250.301349](#).
- [4] Aleksandrs Belovs. Quantum algorithms for classical probability distributions. In *27th Annual European Symposium on Algorithms (ESA 2019)*, LIPIcs **144**, 16:1–16:11 (2019). [doi:10.4230/LIPIcs.ESA.2019.16](#).
- [5] Scott Aaronson, Robin Kothari, William Kretschmer and Justin Thaler. Quantum lower bounds for approximate counting via Laurent polynomials. In *35th Computational Complexity Conference (CCC 2020)*, LIPIcs **169**, 7:1–7:47 (2020). [doi:10.4230/LIPIcs.CCC.2020.7](#).
- [6] Aleksandrs Belovs and Ansis Rosmanis. Tight Quantum Lower Bound for Approximate Counting with Quantum States. *Computational Complexity* **35**, article 2 (2026). [doi:10.1007/s00037-025-00282-7](#). Accessible author version: [arXiv:2002.06879v2](#), 7 May 2024. Theorem 1.1, Table 1 and Appendix A.
- [7] Srinivasan Arunachalam, Aleksandrs Belovs, Andrew M. Childs, Robin Kothari, Ansis Rosmanis and Ronald de Wolf. Quantum Coupon Collector. In *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, LIPIcs **158**, 10:1–10:17 (2020). [doi:10.4230/LIPIcs.TQC.2020.10](#).
- [8] Shima Bab Hadiashar, Ashwin Nayak and Pulkit Sinha. Optimal lower bounds for Quantum Learning via Information Theory. *IEEE Transactions on Information Theory* **70**(3), 1876–1896 (2024). [doi:10.1109/TIT.2023.3324527](#). [arXiv:2301.02227v3](#).
- [9] Ashwin Nayak and Pulkit Sinha. Proper versus Improper Quantum PAC Learning. [arXiv:2403.03295](#) (2024). Theorems 1.1–1.4.
- [10] Paul Erdős and Alfréd Rényi. On a classical problem of probability theory. *Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei* **6**, 215–220 (1961). [repository record](#).
- [11] Sofya Raskhodnikova, Dana Ron, Amir Shpilka and Adam Smith. Strong lower bounds for approximating distribution support size and the distinct elements problem. *SIAM Journal on Computing* **39**(3), 813–842 (2009). [doi:10.1137/070701649](#).
- [12] Yihong Wu and Pengkun Yang. Sample complexity of the distinct elements problem. *Mathematical Statistics and Learning* **1**, 37–72 (2018). [doi:10.4171/MSL/1-1-2](#).
- [13] Renato Ferreira Pinto Jr. and Nathaniel Harms. Testing support size more efficiently than learning histograms. *TheoretCS* **5**, article 10 (2026); preliminary version in STOC 2025. [doi:10.46298/theoretics.26.10](#).

- [14] Fernando Granha Jeronimo, Nir Magrafta and Pei Wu. Pseudorandom and pseudoentangled states from subset states. [arXiv:2312.15285v2](#) (2024).
- [15] Fernando Granha Jeronimo, Nir Magrafta, Joseph Slote and Pei Wu. Coherence in Property Testing: Quantum-Classical Collapses and Separations. [arXiv:2411.15148](#) (2024).
- [16] Marcello Benedetti, Harry Buhrman and Jordi Weggemans. Provable and Verifiable Quantum Advantage in Sample Complexity. *Physical Review Letters* **136**, 040601 (2026). [doi:10.1103/q55v-wm7y](#). Author version titled *Complement Sampling: Provable, Verifiable and NISQable Quantum Advantage in Sample Complexity*, [arXiv:2502.08721v2](#).
- [17] Jing-Peng Zhang, Min-Quan He and Dan-Bo Zhang. Quantum coupon collector with mixed-state encoding. [arXiv:2406.10988](#) (2024).
- [18] Carl W. Helstrom. *Quantum Detection and Estimation Theory*. Academic Press, New York (1976).
- [19] Jonas Kahn and Mădălin Guță. Local asymptotic normality for finite dimensional quantum systems. *Communications in Mathematical Physics* **289**, 597–652 (2009). [doi:10.1007/s00220-009-0787-3](#). [arXiv:0804.3876](#).
- [20] Anurag Anshu, Jonas Haferkamp, Yeongwoo Hwang and Quynh T. Nguyen. On the complexity of unique quantum witnesses and quantum approximate counting. In *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, LIPIcs **362**, 10:1–10:20 (2026). [doi:10.4230/LIPIcs.ITCS.2026.10](#).