

One calibration query activates a finite adaptive advantage in frozen-versus-renewed dephasing tests

Lluis Eriksson

17 October 2026

Working-paper manuscript; not independently reviewed

Abstract

We quantify a finite calibration resource in one-sided discrimination between a frozen unknown qubit phase and independently renewed dephasing. With five signal queries, dephasing parameter $\eta = 7/10$, and one separate noiseless calibration query, an explicit causal strategy has failure probability 0.4657926731886833125. Every parallel five-signal strategy, even with perfect knowledge of the carrier, has failure probability at least 0.46603. The gap is certified by rational matrices and a solver-free verifier. Without calibration, no causal strategy improves the parallel minimax value, so one calibration query is sufficient and minimal at this signal budget. The resource distinction is essential: six parallel signal queries already achieve 0.38442745, and the displayed strategy therefore gives no advantage when calibration and signal queries can be freely exchanged at equal cost. A trajectory bound also rules out an improved asymptotic error exponent in this model. The contribution is a small explicit finite-resource construction, not a claim of a new phase-estimation framework or large scientific impact.

1 Experiment and resource accounting

Let $U_\theta = \text{diag}(1, e^{i\theta})$ and let

$$\mathcal{E}_\eta(\rho) = p\rho + qZ\rho Z, \quad p = (1 + \eta)/2, \quad q = (1 - \eta)/2, \quad 0 \leq \eta \leq 1.$$

The signal black box has two hypotheses. Under H_0 , every call is the same U_θ , with arbitrary unknown θ . Under H_1 , each call is independently $U_\phi \circ \mathcal{E}_\eta$, with arbitrary unknown common carrier ϕ . A separate calibration oracle implements U_ϕ under *both* hypotheses. Under H_0 , θ and ϕ range independently. In particular, the calibration oracle contains no direct label of the hypothesis.

A test must have rejection probability exactly zero under H_0 for every pair (θ, ϕ) . Its failure probability β is the probability of accepting H_0 under H_1 , maximized over ϕ . The resources are an ordered pair (n, L) : n signal queries and L calibration queries. Signal queries are ordinary forward calls, with arbitrary finite quantum memory and ideal controls. There is no environment access, coherent bypass of the signal, or indefinite causal order. Parallel strategies apply their signal calls in one layer; causal strategies may interleave controls and signal calls.

The calibration resource counts information about the *unknown* carrier. It does not charge for the laboratory phase convention, preparation and readout apparatus, classical precision, energy, or control synthesis. A complete physical implementation would need additional resource accounting. No statement about consciousness, an ether, or the ontology of time is involved.

One calibration call applied to $|+\rangle$ supplies

$$|r_\phi\rangle = 2^{-1/2}(|0\rangle + e^{i\phi}|1\rangle).$$

Measure it with the canonical covariant phase POVM $E(d\hat{\phi}) = |e_{\hat{\phi}}\rangle\langle e_{\hat{\phi}}| d\hat{\phi}/(2\pi)$, where $|e_t\rangle = |0\rangle + e^{it}|1\rangle$. This is a valid POVM because its integral is the identity. The error $\delta = \phi - \hat{\phi}$ has density and moments

$$f(\delta) = \frac{1 + \cos \delta}{2\pi}, \quad \mu_k = \int_0^{2\pi} e^{ik\delta} f(\delta) d\delta = \begin{cases} 1 & k = 0, \\ 1/2 & |k| = 1, \\ 0 & |k| \geq 2. \end{cases} \quad (1)$$

Before each signal call apply the known compensating gate $U_{-\hat{\phi}}$. Under H_0 this merely changes the frozen phase to $\theta - \hat{\phi}$, so a test that rejects no frozen phase remains one-sided for every calibration outcome. Under H_1 the residual carrier is δ and its law is independent of ϕ .

The construction does not require a continuous measurement outcome. Use the seven-outcome POVM $E_j = |e_{t_j}\rangle\langle e_{t_j}|/7$, $t_j = 2\pi j/7$, $0 \leq j < 7$, and compensate by U_{-t_j} . Its probabilities are $(1 + \cos(\phi - t_j))/7$. A five-query rejection function $s(\delta)$ has Fourier degree at most five. Therefore $(1 + \cos \delta)s(\delta)$ has degree at most six, and averaging it on this seven-point grid equals its circle average for every shift ϕ . The resulting success is exactly the value computed from (1). The conditional errors themselves need not have a carrier-independent distribution; the final success does.

2 Exact finite-resource separation

Theorem 1. *For $\eta = 7/10$ and resource budget $(n, L) = (5, 1)$ there is a causal test with exactly zero false positives and carrier-independent failure*

$$\beta_{\text{cert}} = \frac{7452682771018933}{16000000000000000}.$$

The optimal parallel failure with five signal queries equals $\beta_{\text{par}} = 46603/100000$, even if the carrier is supplied perfectly. Consequently

$$\beta_{\text{par}} - \beta_{\text{cert}} = \frac{3797228981067}{16000000000000000} = 0.0002373268113166875 > 0.$$

With no calibration calls and no other carrier information, the five-signal causal minimax optimum equals $46603/100000$. Thus one calibration query is the least integer calibration budget that permits a strict causal advantage for these fixed parameters. This statement does not assert optimality of the displayed causal test or minimality of five signal queries.

The gap is about 0.02373 percentage points of failure probability. The parallel comparison is at the same signal budget; giving the parallel side a perfect carrier is an upper bound on what any finite calibration assistance can accomplish.

2.1 Parallel benchmark and the zero-calibration bound

Write $x \in \{0, 1\}^n$, $w(x) = \sum_j x_j$, and $d(x, y)$ for Hamming distance. For words of fixed weight w , the matrix

$$K_w = (\eta^{d(x, y)})_{w(x)=w(y)=w}$$

is positive semidefinite, as a principal submatrix of $\begin{pmatrix} 1 & \eta \\ \eta & 1 \end{pmatrix}^{\otimes n}$. Its row sum is $r_w = \sum_j \binom{w}{j} \binom{n-w}{j} \eta^{2j}$. Let $N_w = \binom{n}{w}$ and $F_w = r_w/N_w$. Since the all-ones vector is an eigenvector, positivity on its orthogonal complement gives

$$K_w \succeq F_w \mathbf{1}\mathbf{1}^T. \quad (2)$$

For a parallel pure input $\sum_x |x\rangle|a_x\rangle$, the frozen output span is generated by the mutually orthogonal vectors $|\psi_w\rangle = \sum_{w(x)=w} |x\rangle|a_x\rangle$. One-sidedness forces the acceptance effect to be the

identity on this span. Put $t_x = \| |a_x\rangle \|^2$ and $t_w = \sum_{w(x)=w} t_x$. Projection onto the frozen span lower-bounds failure by

$$\sum_{w:t_w>0} \frac{1}{t_w} \sum_{x,y:w(x)=w(y)=w} t_x t_y \eta^{d(x,y)} \geq \sum_w t_w F_w \geq \min_w F_w.$$

The carrier cancels within every weight sector. Uniform input on a minimizing weight sector and projection onto its uniform vector attain equality. Purifying a mixed input covers arbitrary ancillas. This proves the benchmark even when ϕ is known.

For an arbitrary causal circuit without carrier information, purify all operations and expand its frozen final state as $\sum_x e^{iw(x)\theta} |a_x\rangle$. Set $|\psi_w\rangle = \sum_{w(x)=w} |a_x\rangle$. These vectors need not be orthogonal, but integration of the unit-norm frozen state gives $\sum_w \|\psi_w\|^2 = 1$. Averaging the renewed output over ϕ removes terms of unequal weight. Equation (2) implies

$$\bar{\rho}_1 \succeq \sum_w F_w |\psi_w\rangle \langle \psi_w|.$$

One-sidedness makes the acceptance effect equal to the identity on all ψ_w , so average failure is at least $\min_w F_w$. Worst-case failure is at least average failure, and the parallel construction attains this value for every ϕ .

For $n = 5$, the minimum is at $w = 2, 3$ for $0 \leq \eta \leq 1$, and

$$\beta_{\text{par}}(5, \eta) = \frac{1 + 6\eta^2 + 3\eta^4}{10}. \quad (3)$$

The zero-calibration and parallel arguments reproduce the structure of the earlier calibrated-noise work [1]; they are included here to state the comparison without relying on a numerical optimum.

2.2 Causal certificate and its realization

The attached files `Q.json` and `R.json` contain exact rational 32 by 32 real symmetric matrices, indexed by five-bit words in lexicographic order. Define $Q_5 = Q$ and obtain lower levels recursively by

$$(Q_{k-1})_{x,y} = (Q_k)_{x0,y0} + (Q_k)_{x1,y1}. \quad (4)$$

The verifier checks

$$Q \succeq 0, \quad R \succeq 0, \quad Q - R \succeq 0, \quad Q_0 = [1], \quad (5)$$

$$(Q_k)_{xa,yb} = 0 \quad (a \neq b), \quad k = 1, \dots, 5, \quad (6)$$

$$\sum_{y:w(y)=w} R_{x,y} = 0 \quad (x \in \{0,1\}^5, \quad w = 0, \dots, 5). \quad (7)$$

Lower Q_k are positive semidefinite because they are sums of principal submatrices. The exact ranks of $(Q, R, Q - R)$ are $(32, 26, 32)$.

For completeness these constraints have an ordinary-circuit interpretation. Choose trajectory vectors with Gram matrix Q_k . By (6), the two last-symbol sectors are orthogonal, and can be represented as query states $|0\rangle$ and $|1\rangle$ tensor memory. By (4), the sums of the two child vectors have the same Gram matrix as the parent vectors. An isometry therefore maps parents to these sums; extending the isometry to a unitary implements the next control. The diagonal query then multiplies each child by its corresponding phase. Iteration constructs a five-query causal circuit. A query qubit and memory of dimension at most 16 suffice. Finally $0 \preceq R \preceq Q$ realizes R as the Gram matrix of a rejection effect on the terminal vectors. Equation (7) annihilates every frozen weight sum, proving exact one-sidedness.

Under renewed dephasing and residual carrier δ , the rejection probability is

$$s(\delta) = \sum_{x,y} R_{x,y} \eta^{d(x,y)} e^{i(w(x)-w(y))\delta}.$$

Integrating with (1) gives the rational functional

$$s = \sum_{x,y} R_{x,y} \eta^{d(x,y)} \mu_{w(x)-w(y)} = \frac{8547317228981067}{16000000000000000}. \quad (8)$$

Equations (3) and (8) prove the strict gap.

The verifier uses only Python's standard-library integer fractions. It performs symmetric Schur elimination, rejecting a negative pivot or a nonzero remaining row at a zero pivot. It also replays the probability through all 32 independent I/Z noise histories, checks the six parallel weight values and the seven-point Fourier quadrature, and rejects four deliberately broken constraints. Numerical optimization was used to find a candidate, followed by rational projection and an interior perturbation; the theorem uses only the final exact certificate. No solver status is used as a proof of feasibility or optimality.

3 What the resource comparison rules out

Proposition 1. *The displayed construction does not outperform parallel testing under a freely reallocatable budget of six equally costly queries.*

Proof. Using all six calls on the signal, the balanced parallel test has

$$\beta_{\text{par}}(6, 7/10) = \frac{1 + 9(7/10)^2 + 9(7/10)^4 + (7/10)^6}{20} = \frac{7688549}{20000000} = 0.38442745 < \beta_{\text{cert}}.$$

This is a countercomparison to the displayed construction, not a theorem that all causal schemes lose at total budget six. $\square$

Proposition 2. *For fixed $n \geq 1$ signal calls, any one-sided causal protocol in this model, even with arbitrary noiseless carrier-calibration assistance, has $\beta \geq p^n + q^n$. Calibration cannot improve the optimal asymptotic error exponent per signal query beyond $-\log p$. The optimal exponent under a freely reallocatable total-query budget is also $-\log p$.*

Proof. In the renewed mixture, the all- I history has probability p^n . Its joint signal and calibration experiment is exactly the allowed null experiment with $\theta = \phi$. The all- Z history, of probability q^n , is the null experiment with $\theta = \phi + \pi$ and the same calibration oracle. Exact one-sidedness forces acceptance on both histories, including when calibration and signal operations are coherent and interleaved. Their probabilities give the lower bound.

Conversely, for any fixed-weight parallel input the sum over that weight is bounded by the sum over all words:

$$F_w \leq \frac{(1 + \eta)^n}{\binom{n}{w}}.$$

A largest binomial coefficient is at least $2^n/(n+1)$, so $\beta_{\text{par}} \leq (n+1)p^n$. The two bounds have the same exponent. For total budget $N = n + L$, the lower bound $\beta \geq p^n \geq p^N$ and the all-signal parallel construction with $n = N$ give the same conclusion. The case $p = 1$ has zero exponent. $\square$

The path argument extends an obstruction already present in Ref. [1]; it is not presented as a new general no-go principle. Finite-sample gains may still matter when signal access is scarce, calibration is reusable, or the two kinds of calls have different costs. Those possibilities require explicit new models and proofs.

4 Prior art, self-review, and open work

Finite quantum phase references and covariant phase estimation are established subjects [2, 3]. For a longer token $\sum_{j=0}^L c_j e^{ij\phi} |j\rangle$, its Fourier error moments are autocorrelations of the coefficients c_j . Optimizing an average phase-dependent payoff is a Toeplitz eigenvalue problem, a structure explicitly treated in recent phase-estimation work [4, 5]. We therefore make no novelty claim for that reduction. The operational value of asymmetry for discrimination is also known in broader settings [6]. Our narrow candidate contribution is the exact one-query activation at the stated one-sided signal-testing task and its explicit resource comparison.

In the first exploration, the earlier fixed five-query tester needed a longer reference to show an advantage. Reoptimizing the signal tester produced the present one-query certificate. This refutes any inference that the reference threshold of one fixed tester is a universal threshold. The current certificate was recalculated by a separate exact verification script, but all reasoning and code were produced in the same assistant session: this is *self-review*, not review by an independent researcher or model.

The strongest outstanding questions are quantitative and finite: the exact causal optimum with one calibration query, the minimum signal budget allowing activation, explicit low-memory circuits, comparisons with coherent rather than measured reference access, and robustness when both test classes are jointly optimized at a positive false-positive allowance. No significant impact rating or priority guarantee follows from the current evidence. Before admission as an accepted article, the scope and novelty need external scrutiny, and the physical cost of controls must be separated from the finite carrier-information budget. A failure of a more ambitious conjecture should be retained as a constraint on subsequent searches rather than hidden by changing the benchmark.

References

- [1] L. Eriksson, *Phase calibration unlocks adaptive tests for renewed quantum noise*, working paper, AIRR ARR-2026-33GYTDAT5C8THB3M, v1 (2026). [Permanent version](#). Deposited; not peer reviewed.
- [2] V. Buzek, R. Derka, and S. Massar, *Optimal Quantum Clocks* (1998), [arXiv:quant-ph/9808042](#).
- [3] S. D. Bartlett, T. Rudolph, and R. W. Spekkens, *Reference frames, superselection rules, and quantum information* (2006), [arXiv:quant-ph/0610030](#).
- [4] J. G. Smith, C. H. W. Barnes, and D. R. M. Arvidsson-Shukur, *Risk-minimizing states for the quantum-phase-estimation algorithm* (2025), [arXiv:2505.00764](#).
- [5] Q. Qian and C. N. Gagatsos, *Optimal Probe State for Phase Estimation Under Covariant Measurement* (2026), [arXiv:2606.18169](#).
- [6] M. Piani et al., *Robustness of asymmetry and coherence of quantum states* (2016), [arXiv:1601.03782](#).