Purifying an unknown qubit with stabilizer operations

Yuxuan Zhang

24 September 2026

Institute of Physics, École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
Department of Physics, Princeton University, Princeton, New Jersey 08544, USA

Working manuscript. The proof was generated with AI assistance and passed internal mathematical review and exact arithmetic checks. External expert confirmation and literature priority remain unverified.

Abstract

We give a five-copy stabilizer protocol that increases the fidelity of every depolarized pure qubit at noise strength $\delta = 2/3$. It uses Pauli measurements, stabilizer preparation, classical randomization, and an untouched input as a fallback. Its fidelity is at least $11693/17496 > 2/3$ for every pure target, and its Haar average is $325/486$. Success is deterministic. A separate four-copy postselection protocol has Haar-conditional fidelity $547/726 > 3/4$ at $\delta = 1/2$. Either construction refutes the universally quantified equality in QIQCOP problem `op_a64dc63d6ae49127`. The constructions leave the two-copy no-go theorem intact; no claim of globally optimal fidelity is needed or made.

The question and the fidelity convention

For a unit Bloch vector $r \in \mathbb{R}^3$, write

$$\psi(r) = \frac{I + r \cdot \sigma}{2}, \qquad \rho_t(r) = \frac{I + tr \cdot \sigma}{2}, \qquad t = 1 - \delta.$$

Retaining one input has fidelity $f_t = (1+t)/2$ with the pure target. The QIQCOP question [1] asks whether classically simulable operations can ever exceed this value, with arbitrarily many copies. In the qubit case, the admissible maps are completely stabilizer preserving, completely positive, and trace nonincreasing. For a probabilistic map $\mathcal{E}$, the relevant fidelity is

$$F(\mathcal{E}) = \frac{\int \mathrm{Tr}[\psi(r)\mathcal{E}(\rho_t(r)^{\otimes n})]\, dr}{\int \mathrm{Tr}[\mathcal{E}(\rho_t(r)^{\otimes n})]\, dr},$$

where dr is normalized Haar measure. This is the ratio of the two integrals, as in He et al. [2], not an unweighted average of conditional fidelities.

The distinction matters for the four-copy example. The five-copy construction below has success probability one and improves every target separately, so it does not rely on conditioning the target distribution.

A deterministic protocol with five copies

Choose uniformly a signed Pauli axis

$$v \in \{e_x, -e_x, e_y, -e_y, e_z, -e_z\}.$$

Measure the first four inputs in that Pauli basis. If all four outcomes equal v, prepare the corresponding pure stabilizer state. Otherwise output the untouched fifth input. Discard the records.

Every branch is an ordinary stabilizer operation: Pauli measurement, conditional preparation or retention, and discarding. These operations preserve the stabilizer cone even in the presence of an arbitrary reference system. Classical mixing preserves that property. The branches exhaust all outcomes, making the protocol trace preserving. No step depends on the unknown target.

For a fixed r and a fixed chosen v, the replacement event has probability $((1+tv\cdot r)/2)^4$. The replacement output has target fidelity $(1+v\cdot r)/2$. On the other outcomes the fifth input still has fidelity f_t, since the input is a product state conditional on r. Thus the gain, averaged over the protocol's own randomness, is

$$G_t(r) = \frac{1}{6}\sum_v \left(\frac{1+tv\cdot r}{2}\right)^4 \frac{v\cdot r - t}{2}.$$

Pair the two signs of each axis. For any real x,

$$\begin{aligned}&(1+tx)^4(x-t) + (1-tx)^4(-x-t)\\ &\quad = 2\left[-t + (4t-6t^3)x^2 + (4t^3-t^5)x^4\right].\end{aligned}$$

Using $r_x^2 + r_y^2 + r_z^2 = 1$ gives the exact expression

$$G_t(r) = \frac{t - 6t^3 + (4t^3 - t^5)S_4}{96}, \qquad S_4 = r_x^4 + r_y^4 + r_z^4.$$

Cauchy–Schwarz gives $S_4 \geq 1/3$. At $t = 1/3$, the coefficient of S_4 is $35/243 > 0$, hence

$$G_{1/3}(r) \geq \frac{1/9 + 35/729}{96} = \frac{29}{17496} > 0.$$

We have proved the following statement for every pure qubit target:

$$F_{\text{det}}(r) \geq \frac{2}{3} + \frac{29}{17496} = \frac{11693}{17496} > \frac{2}{3}.$$

For the Haar average, $\mathbb{E}[S_4] = 3/5$, giving

$$F_{\text{Haar}} = \frac{2}{3} + \frac{1}{486} = \frac{325}{486}.$$

The same protocol has a strictly positive pointwise gain whenever $0 < t < \sqrt{2\sqrt{13} - 7}$, because its lower bound has numerator proportional to $3 - 14t^2 - t^4$. The rational point $t = 1/3$ alone suffices to refute the universal equality.

Discarding additional inputs extends the construction to every $n \geq 5$. Independently accepting the output with any fixed probability $0 < s \leq 1$ gives that success probability for every target and leaves its conditional fidelity unchanged.

Four-copy postselection

A shorter counterexample for the Haar-conditional criterion uses the single Kraus operator

$$K = |0\rangle\langle 0000|, \qquad \mathcal{E}(X) = KXK^\dagger.$$

It is implemented by measuring all four inputs in the computational basis, accepting only outcome 0000, and preparing $|0\rangle$. Since $K^\dagger K \leq I$, it is trace nonincreasing. Its Choi projector is a five-qubit all-zero stabilizer projector, and the measurement implementation proves complete stabilizer preservation.

For Haar-random r, its z coordinate is uniform on $[-1, 1]$. The success probability for a fixed target is $((1 + tz)/2)^4$, and the output fidelity is $(1 + z)/2$. Define $D = 1 + 2t^2 + t^4/5$. Direct polynomial integration gives

$$s = \frac{D}{16}, \qquad N = \frac{D + 4t/3 + 4t^3/5}{32}, \qquad F = \frac{N}{s}.$$

Subtracting the one-copy fidelity yields

$$F - \frac{1+t}{2} = \frac{t(5 - 18t^2 - 3t^4)}{30D}.$$

At $t = 1/2$, these quantities are exactly

$$s = \frac{121}{1280}, \quad N = \frac{547}{7680}, \quad F = \frac{547}{726} = \frac{3}{4} + \frac{5}{1452}.$$

Equivalently, $N - (3/4)s = 1/3072 > 0$. This example does not improve every target; it exploits the conditional distribution of accepted inputs. That is allowed by the displayed Haar-conditional definition, but should not be confused with the pointwise theorem above.

There is also a fixed-axis deterministic variant with five inputs: use this replacement event and retain the fifth input on failure. At $t = 1/2$, its Haar fidelity is $3/4 + 1/3072 = 2305/3072$. Unlike the signed-axis protocol at $t = 1/3$, this fixed-axis variant is not claimed to improve every target.

Scope, verification, and prior work

The universal equality in [1] is false if either admissible example is correct. Determining optimal fidelity at each copy number is a different problem. We make no claim about the best four-copy deterministic protocol, the optimal five-copy fidelity, or arbitrary odd dimensions.

He et al. [2] prove an analytical two-copy no-go theorem and report numerical no-go results for three and four copies. The constructions here do not contradict their two-copy theorem. The four-copy example conflicts with the numerical four-copy assertion under the published arXiv definition. We have not established the cause of that discrepancy or checked the subscription-only final article and supplement in full; we therefore do not attribute an error to a particular implementation. The later resource-law paper [3] treats two-copy tradeoffs and does not establish an arbitrary-copy obstruction.

The archived checker evaluates the four-copy integrals, the five-copy polynomial and decisive rational gaps exactly, and checks 1024 Pauli Choi coefficients for the four-copy map. Internal review also examined complete stabilizer preservation and the pointwise quantifiers; arithmetic reproduction alone would not establish either. The frozen proof and its review are included separately so that this exposition can be checked against the actual research record. An early arithmetic error in the five-copy Haar gain was caught before the frozen review; the accepted value is $1/486$.

The constructions and proof were prepared in round 35 of an AI-assisted research campaign. This publication is a presentation of that result, not a new discovery or evidence of independent human refereeing. K-Dense's scientific-writing guidance [4] was used for evidence tracking and exposition. Author affiliations follow the author's instructions; no unrecorded funding or conflict declaration is inferred.

References

1. QIQCOP Zoo, *Universal purification with classically simulable operations*, problem `op_a64dc63d6ae49127`. https://qiqc-op.com/problem/op_a64dc63d6ae49127/. Statement checked 24 September 2026.
2. K. He, C. Zhu, H. Yao, J. Liu, Y. Li, and X. Wang, *No-Go Theorems for Universal Quantum State Purification via Classically Simulable Operations*, arXiv:2504.10516v2. https://arxiv.org/abs/2504.10516v2. Equations (2)–(3), Appendix B, and the stated numerical extensions.
3. *A Nonstabilizerness Resource Law for Universal Quantum State Purification*, arXiv:2607.08626v1. https://arxiv.org/abs/2607.08626v1. Two-copy resource tradeoffs.
4. T. Kassis, V. Agarwal, Y. He, D. Patel, and A. M. Brueckner, *Scientific Agent Skills: A Library of Procedural Knowledge for Research Agents* (2026). https://doi.org/10.48550/arXiv.2609.00065. Local scientific-writing skill at commit `330c8e764435`.