Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

How noisy quantum time loops select forecasts that match the future

Peter M. Catt

peter.catt@virtualblue.co.nz

Independent Researcher https://orcid.org/0009-0000-1324-0861

Research Article

Keywords: Closed timelike curves, Postselected teleportation, Amplitude damping, Quantum Doeblin coefficient, Semidefinite programming, Forecasting, Superconducting qubits, Qiskit

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11777101/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

How noisy quantum time loops select forecasts that match the future

Peter Maurice Catt

Auckland, New Zealand · www.theknowablefuture.com · ORCID 0009-0000-1324-0861

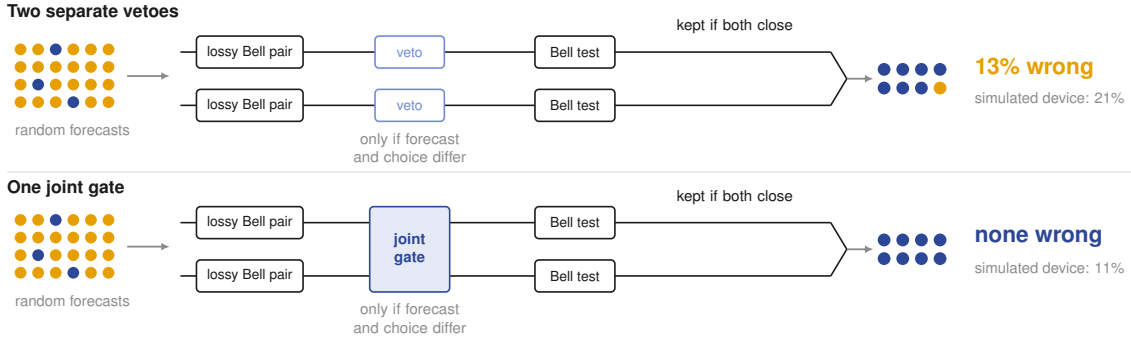
ABSTRACT

Used as a forecast test, a quantum time loop simulated by postselected teleportation selects the runs in which an early random forecast matches a later choice, and it selects well only if a later operation can stop the loop from closing when the two differ. We prove an exact formula for the least closing probability of n loops under energy loss. At any non-zero loss, one loop can never be stopped completely, so some wrong forecasts always survive. Two loops operated jointly can be stopped completely by a single entangling gate for loss up to 82.8%, and n loops for loss up to $1-(2^{1/n}-1)^2$. No wrong forecast then survives, for any number M of possible choices, but only a fraction $p_{\max}/M$ of runs survives. For two loops, complete stopping needs a nonseparable operation once loss exceeds 50%. Under a simulated noise model for a six-qubit superconducting circuit, the joint gate reaches 2.9 bits of contrast per loop against 2.4 for the best separate operations at 83% loss, and wrong three-bit forecasts fall from 21% to 11% of surviving runs. In a pilot on an IBM Heron processor at a programmed loss of 86%, the joint gate left both loops closed in 0.90% of runs against 1.19% for a separate veto on each loop, a little under half the simulated gain.

Keywords: Closed timelike curves, Postselected teleportation, Amplitude damping, Quantum Doeblin coefficient, Semidefinite programming, Forecasting, Superconducting qubits, Qiskit

Graphical Abstract: How noisy quantum time loops select forecasts that match the future

A veto acts on runs whose forecast and choice differ, and Bell tests keep a run only if both loops close



Loss: probability that a stored qubit gives up its excitation. Forecast: one of 8 values drawn at random. Dots show the mix of right and wrong forecasts at 80% loss (amplitude damping only), not run counts; only 3 to 4 runs in 100 are kept. Simulated device: Qiskit, typical superconducting noise, 83% loss. Nothing is sent back in time.

Peter Maurice Catt · theknowablefuture.com · 18 November 2026

Corresponding author: peter.catt@theknowablefuture.com

Introduction

Postselected teleportation simulates a closed timelike curve (Lloyd et al., 2011a,b; Svetlichny, 2009). A qubit is teleported without the usual correction, and in the runs where the Bell measurement returns the reference outcome, so that the loop closes, the qubit behaves as if it had travelled back to the moment the entangled pair was made. Such loops have been run on photonic, trapped-ion and superconducting processors (Huang et al., 2026; Lloyd et al., 2011b) and underlie a proposed advantage in metrology (Arvidsson-Shukur et al., 2023). Ji et al. (2026) recently defined the retrocausal capacity of a noisy loop. A later party who wants an earlier record to read m leaves the loop alone when the record already agrees and otherwise applies a veto, an operation chosen to stop the loop from closing. The record on the surviving runs is reliable to the extent that the largest closing probability, $p_{\max}$, exceeds the least closing probability, $p_{\min}$, that any veto can enforce. With $R = p_{\max}/p_{\min}$, the largest number of values of the later choice that the record can hold with error ε is the integer part of $\varepsilon R/(1-\varepsilon)+1$ (Ji et al., 2026). We call $\log_2 R$ the contrast of the loop. In the notation of Ji et al. (2026) it is the sum of the channel's max-information and Doeblin information. It is a logarithmic ratio in bits, distinct from the mutual information between record and later choice, which we report separately.

Read as a forecast test, the closing test favours the runs in which an early random record agrees with a later choice, and a wrong record survives only as often as the vetoed loops still close. The least closing probability therefore sets the error rate of the test.

For n loops built on a qubit channel, $p_{\min}$ is the least singlet fraction (Horodecki et al., 1999) that an operation on the stored qubits can leave, and $4^n p_{\min}$ is the quantum Doeblin coefficient of n copies of the channel (George et al., 2026). That coefficient is known not to be multiplicative, and amplitude damping, the standard model of energy loss (Nielsen & Chuang, 2010), has been reported as a numerical example (George et al., 2026). For amplitude damping, Ji et al. (2026) give the one-copy values and from them a lower bound on the asymptotic capacity. We are not aware of a closed form for more than one copy. Here we derive it for energy loss. From it we obtain the number of values that one loop, and several loops together, can hold. We simulate a six-qubit superconducting experiment that tests the result, and we report a pilot run of that experiment on an IBM processor.

Methods

Loop model and semidefinite programme

One loop is a Bell pair of a stored qubit and a reference qubit whose stored half has suffered amplitude damping of strength γ, the probability that an excitation is lost. A later operation of any kind acts on the stored qubits of n such loops, and each loop is then tested by a Bell measurement (Figure 1a). Amplitude damping has Kraus operators $K_0 = |0\rangle\langle 0| + s|1\rangle\langle 1|$ and $K_1 = \sqrt{\gamma}|0\rangle\langle 1|$, with $s = \sqrt{1-\gamma}$ (Nielsen & Chuang, 2010). The pair starts in $|\Phi\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$, and its stored qubit then passes through this channel. With the stored qubit written first, the state of the loop is $J = \frac{1}{2}\big(|v\rangle\langle v| + \gamma|01\rangle\langle 01|\big)$, where $|v\rangle = |00\rangle + s|11\rangle$. Let J_n be the state of n loops and Φ_n the projector onto n copies of $|\Phi\rangle$. A later channel $\mathcal{E}$ on the stored qubits gives closing probability $p(\mathcal{E}) = \mathrm{Tr}[\Phi_n(\mathcal{E}\otimes\mathrm{id})(J_n)] = \mathrm{Tr}[J_n\Omega]$, with $\Omega = (\mathcal{E}^\dagger\otimes\mathrm{id})(\Phi_n)$. Because $\mathcal{E}^\dagger$ is completely positive and unital, $\Omega \geq 0$ and $\mathrm{Tr}_R\Omega = I/2^n$, where R denotes the reference qubits and I the identity, and every such Ω arises from a channel. Hence

$$p_{\min}(n) = \min\big\{\mathrm{Tr}[J_n\Omega] :\ \Omega \geq 0,\ \mathrm{Tr}_R\Omega = I/2^n\big\}. \tag{1}$$

For Hermitian Y on the stored qubits with $Y\otimes I \leq J_n$, every admissible Ω gives $\mathrm{Tr}[J_n\Omega] \geq \mathrm{Tr}[(Y\otimes I)\Omega] = \mathrm{Tr}[Y]/2^n$. In the notation of George et al. (2026), $4^n p_{\min}(n)$ is the Doeblin coefficient of n copies of the channel, whose Choi operator is $2^n J_n$.

Circuit

The simulated device is a line of six qubits in the order environment, reference, stored, stored, reference, environment, so that the two stored qubits are neighbours. Each loop is prepared directly in the state $(|S_0R_0A_0\rangle + s|S_1R_1A_0\rangle + \sqrt{\gamma}|S_0R_1A_1\rangle)/\sqrt{2}$, where S, R and A are its stored, reference and environment qubits. This is a Bell pair whose stored half has undergone amplitude damping, with A holding the lost excitation. The preparation uses one two-qubit state preparation and two controlled-NOT gates. The later operation is nothing, a Pauli Z gate on each stored qubit (the best separate veto), or the joint veto, which consists of phase rotations $-t$ and $+t$ and one controlled-Z gate. The closing test is a controlled-NOT gate from stored to reference and a Hadamard gate on the stored qubit, followed by readout of both, and a run counts as closed when all four bits are 0. After transpilation to controlled-Z, $\sqrt{X}$ and virtual Z rotations the circuit has nine controlled-Z gates: six in the preparation, one in the joint veto and two in the closing test. Without device noise it reproduces the three closed forms to 10^{-15}.

Noise model

Simulations used Qiskit (Javadi-Abhari et al., 2024) (version 2.6.2) and the density-matrix method of Qiskit Aer (version 0.18.2), which returns exact outcome probabilities. Every X and $\sqrt{X}$ gate carries a depolarising error with average infidelity r_1, and every controlled-Z gate one with average infidelity r_2. Energy relaxation and dephasing with times T_1 and T_2 act during gates, with assumed durations of 32 ns for single-qubit gates and 68 ns for controlled-Z gates, and during the idle intervals found by scheduling each stage of the circuit. Readout has a symmetric assignment error e. We call these errors device noise, as distinct from the amplitude damping γ that defines each loop. The typical set is $T_1 = 260$ µs, $T_2 = 140$ µs, $r_1 = 2.3 \times 10^{-4}$, $r_2 = 2.0 \times 10^{-3}$ and $e = 0.9\%$, close to medians reported for ibm_kingston (Cobos et al., 2025). A favourable set (300 µs, 200 µs, 1.5×10^{-4}, 1.2×10^{-3}, 0.5%) and an unfavourable set (200 µs, 100 µs, 3.0×10^{-4}, 3.3×10^{-3}, 1.5%) bracket it. The model omits crosstalk, leakage, coherent gate errors, readout asymmetry and drift, so a device may perform worse. Reported probabilities include the readout error and use no error mitigation. Shot noise was estimated by binomial sampling of the exact probabilities.

Floor for separate operations

For each loop we computed the noisy pair state ρ at the end of the preparation and the noisy closing effect M, the operator for which $\mathrm{Tr}[\rho M]$ is the probability that the closing test reads 00. M was reconstructed from simulations of the test on 16 product input states. The least closing probability of one noisy loop over all channels $\mathcal{C}$ on its stored qubit, the minimum of $\mathrm{Tr}[(\mathcal{C} \otimes \mathrm{id})(\rho)M]$, is a semidefinite programme in the Choi operator of $\mathcal{C}$. Preparation and test act separately on the two loops in this model, so the product of the two single-loop minima bounds every pair of separate operations from below, with or without shared randomness. The Pauli Z veto attains it to three decimal places. The same programme with a joint channel on both stored qubits gives the best joint operation that the noisy states allow. Supplementary Note S2 lists the simulated values.

Forecast test protocol

An earlier record is drawn uniformly from K values, and a later message is drawn in the same way. The later party leaves the loops alone when the two agree and applies the veto otherwise. Among the runs in which both loops close, the record is wrong with probability $(K-1)p_{\mathrm{veto}}/[p_{\mathrm{close}} + (K-1)p_{\mathrm{veto}}]$, and the fraction of runs that survive is $[p_{\mathrm{close}} + (K-1)p_{\mathrm{veto}}]/K$. Supplementary Tables S6 to S8 give the wrong shares for one to three bits, the bits carried by a three-bit record with the surviving fraction, and the cost of waiting for a message that is generated during the run. The bits carried are the mutual information between record and later choice on surviving runs, $\log_2 K - h(w) - w\log_2(K-1)$ for a wrong share w, where h is the binary entropy. When the message is generated during the run, the loops wait through a readout and the feed-forward delay, and the whole test runs as one dynamic circuit on eight qubits. This circuit implements the one-bit task: a record qubit and a message qubit are each put in an equal superposition and measured, the record before the loops are prepared and the message after, and the veto is applied by feed-forward when the two recorded bits differ. The two- and three-bit values are not simulated as dynamic circuits; they follow from the formula applied to simulated closing probabilities. The record and message measurements carry the same assignment error as the loop readout, and the record is scored against the recorded message bit, which is also the bit that conditions the veto.

Numerical checks

Programme (1) was solved with CVXPY (Diamond & Boyd, 2016) (version 1.9.3) and the Clarabel solver for $n = 1$, 2 and 3, and the results agree with equation (2) to solver precision. The operator $J_n - Y \otimes I$, with Y the certificate of equation (4), has no eigenvalue below -10^{-16} for $n \leq 4$ on a grid of 50 values of γ. For depolarising noise, in which the stored qubit is replaced by the maximally mixed state with probability q, the two-loop minimum equalled the square of the one-loop minimum, $q/4$, at $q = 0.05$, 0.2 and 0.5, as the multiplicativity of the coefficient for such channels requires (George et al., 2026).

Pilot hardware run

The pilot was run on ibm_kingston, a 156-qubit IBM Heron processor, at $\gamma = 0.86$. Its design, stopping rule, analysis and predictions were fixed in a plan written on 14 November 2026, before any job for this paper had run, and the plan is deposited with the code. A line of six qubits was chosen from the calibration data of the day and kept for all jobs. Three circuits were run: no operation, the Pauli Z veto on each stored qubit, and the joint gate at $t = \pi$, the optimal phase for $\gamma > 0.828$ without device noise. The phase was not tuned on the device. The loss is programmed: each loop is prepared in the damped state with its environment qubit, as in the simulated circuit, and is distinct from the decoherence of the processor, which acts in addition. One block is a single job of five circuits of 10,000 shots, with the two veto circuits placed symmetrically about the no-operation circuit and their places exchanged in the second block. No error mitigation, readout correction, twirling or dynamical decoupling was used. The number of blocks depended on the processor time used and not on the results. The primary outcome is the difference between

the fractions of shots in which both loops close under the separate veto and under the joint gate, pooled over blocks, with a one-sided z test at the level 0.01. The pilot measures the closing fractions of these three circuits and does not run the forecast protocol; forecast shares quoted for it follow from the formula applied to the measured fractions. Supplementary Note S3 gives the jobs, the calibration values and the counts of every circuit.

Planned hardware test

The pilot compares the joint gate with one separate veto. The full test is planned as follows, with the protocol and shot numbers fixed before any of its data are taken. A line of six qubits is chosen on an IBM Heron processor from the calibration data of the day, for low controlled-Z error on its five couplings and low readout error on the four measured qubits. The two damped pairs are characterised by tomography with nine measurement settings, and the closing test on 16 product input states, so that the floor for separate operations comes from measured states and a measured test. The three arms (no operation, separate veto, joint gate) are run at $\gamma = 0.83$, 0.86 and 0.90 in interleaved blocks, and the gate phase is scanned at $\gamma = 0.80$. Every shot records both loops, so the same runs also give the closing frequency of each loop on its own under no operation and under the Pauli Z veto. These single-loop values are the baseline against which the joint gate is read. The primary outcome is the difference between the measured floor and the joint-gate closing probability at each γ. State tomography takes the measurement operators as known, and effect tomography takes the input states as known. Single-qubit gates and the readout assignment matrix, calibrated in the same session, are therefore treated as trusted, and states and effects are reconstructed by maximum likelihood under positivity constraints. The floor carries the sampling uncertainty of all of these, so the tomography, closing-test and calibration counts are resampled, the floor is recomputed for each resample, and the joint gate is compared with that distribution in a one-sided test at a level fixed in advance. The figure of 16 standard errors quoted for the simulation is shot noise alone. The achieved contrast per loop follows from the measured closing probabilities.

Statistics and reproducibility

Simulated probabilities are exact outputs of the density-matrix method and carry no sampling error, and agreement with the closed forms is reported as the largest absolute deviation. The shot-noise figures are binomial standard errors of the difference between two independent proportions at the stated number of shots. The simulation of the dynamic circuit used 150,000 sampled runs, and its uncertainties are binomial standard errors. In the pilot and in the planned test each closing probability is the fraction of shots in which all four measured bits read 0, with a binomial standard error. A large language model (Claude, Anthropic) assisted with the proof, code and text; the author verified all content and takes responsibility for it.

Results

Least closing probability of n lossy loops

With $s = \sqrt{1-\gamma}$, the least probability that all n loops close is

$$p_{\min}(n) = \frac{\max\{0,\ 2-(1+s)^n\}^2}{4^n}. \tag{2}$$

A diagonal phase gate attains this value. A loss event leaves its loop in a state that cannot close, so only the branch without loss matters. Its amplitude is a sum of 2^n terms of sizes $s^{|x|}$, one for each pattern x of excited stored qubits, with phases that the gate sets freely. The terms cancel exactly when the largest, 1, does not exceed the sum of the others, $(1+s)^n - 1$. No operation does better, as an explicit solution of the dual semidefinite programme shows. The largest closing probability is multiplicative, $p_{\max}(n) = (1+s)^{2n}/4^n$, as the additivity of the max-information implies (Ji et al., 2026), and is reached by leaving the loops alone. The contrast per loop is therefore $C_n = (2/n)\log_2\left[(1+s)^n/(2-(1+s)^n)\right]$ while $(1+s)^n < 2$, and it is infinite otherwise.

Proof of the upper bound

Let $\mathcal{E}$ be the diagonal unitary $D = \sum_x e^{i\theta(x)}|x\rangle\langle x|$, where x runs over the 2^n bit strings of the stored qubits and $|x|$ is the number of ones in x. Expanding J_n gives one term for each set of loops that carry $\gamma|01\rangle\langle 01|$. Such a loop has stored bit 0 and reference bit 1, D changes no bit value, and $|\Phi\rangle$ has no component on $|01\rangle$, so every term with a non-empty set contributes nothing. The remaining term is 2^{-n} times the projector onto the unnormalised vector $|v\rangle^{\otimes n}$, and $\langle\Phi|^{\otimes n}(D\otimes I)|v\rangle^{\otimes n} = 2^{-n/2}\sum_x s^{|x|}e^{i\theta(x)}$. Hence

$$p(D) = \frac{\left|\sum_x s^{|x|}e^{i\theta(x)}\right|^2}{4^n}. \tag{3}$$

The moduli $s^{|x|}$ have largest member 1, at $x = 0\cdots 0$, and the others sum to $(1+s)^n - 1$. Complex numbers with given moduli can be given phases that make them sum to zero exactly when no modulus exceeds the sum of the others. If

$(1+s)^n \geq 2$ the phases can therefore be chosen so that $p(D) = 0$. Otherwise $\theta(0\cdots0) = 0$ and $\theta(x) = \pi$ for every other x give the least modulus of the sum, $2-(1+s)^n$.

Proof of the lower bound

Nothing is to be shown when $(1+s)^n \geq 2$. Otherwise put $a = 1+s$ and $\tau = (2-a^n)/2^n > 0$, and take

$$Y = \tau\Big(|0\cdots0\rangle\langle 0\cdots0| - \sum_{x\neq 0\cdots0} s^{|x|}|x\rangle\langle x|\Big). \tag{4}$$

Then $\mathrm{Tr}[Y] = \tau\big(1-(a^n-1)\big) = 2^n\tau^2$, so $p_{\min}(n) \geq \tau^2$, the right side of equation (2), once $Y\otimes I \leq J_n$ is shown. Split the space of one loop into the span of $|00\rangle$ and $|11\rangle$, the line of $|01\rangle$ and the line of $|10\rangle$. J respects this splitting and vanishes on $|10\rangle$. Sorting the loops by these three types splits the space of n loops into sectors, labelled by the set T of loops in $|01\rangle$ and the set U of loops in $|10\rangle$. J_n is block diagonal over the sectors, and so is $Y\otimes I$, which is diagonal in the bit-string basis.

In a sector with U not empty, $J_n = 0$ and every basis state has a stored bit equal to 1, so its entry in $Y\otimes I$ is $-\tau s^{|x|} \leq 0$.

In a sector with U empty, let T have k elements and put $m = n-k$. The basis states are labelled by the strings z of m bits, each bit being the common value of the two qubits of a remaining loop, and the stored string has $|x| = |z|$. On this sector $J_n = 2^{-n}\gamma^k|w\rangle\langle w|$ with $w(z) = s^{|z|}$, and $Y\otimes I$ has entry τ at $z = 0\cdots0$ and $-\tau s^{|z|}$ elsewhere. For a vector c write c_0 for its entry at $z = 0\cdots0$ and $u = \sum_{z\neq 0\cdots0} s^{|z|}c_z$. The inequality to be shown is

$$2^{-n}\gamma^k|c_0+u|^2 - \tau|c_0|^2 + \tau\sum_{z\neq 0\cdots0} s^{|z|}|c_z|^2 \;\geq\; 0. \tag{5}$$

For $s > 0$ and $m \geq 1$ the Cauchy–Schwarz inequality gives $|u|^2 \leq (a^m-1)\sum_{z\neq 0\cdots0} s^{|z|}|c_z|^2$. So inequality (5) holds if the form $A|c_0+u|^2 - \tau|c_0|^2 + B|u|^2$ is non-negative for all c_0 and u, with $A = 2^{-n}\gamma^k$ and $B = \tau/(a^m-1)$. That is the case when $A \geq \tau$ and $(A-\tau)(A+B) \geq A^2$. The second condition is $A(2-a^m) \geq \tau$, and it implies the first because $0 < 2-a^m \leq 1$. For $m = 0$ or $s = 0$ the sum in (5) is empty or $u = 0$, and $A \geq \tau$ is enough. In every case it suffices that

$$\gamma^k\big(2-a^{n-k}\big) \;\geq\; 2-a^n, \qquad k = 0,1,\ldots,n. \tag{6}$$

Since $\gamma = a(2-a)$, direct expansion gives $\gamma(2-a^{j-1}) - (2-a^j) = (a-1)(a^j-2a+2)$ for $j \geq 1$. The right side is non-negative for $1 \leq a \leq 2$: it equals $(a-1)(2-a)$ at $j = 1$, and $a^j - 2a + 2 \geq (a-1)^2 + 1$ for $j \geq 2$. Using this for $j = n, n-1, \ldots, n-k+1$, each time multiplied by a power of γ, gives the chain $2-a^n \leq \gamma(2-a^{n-1}) \leq \cdots \leq \gamma^k(2-a^{n-k})$, which is (6).

Largest closing probability and contrast

The operator $Z = \mathrm{diag}\big((1+s)/2,\; s(1+s)/2\big)$ on the stored qubit satisfies $Z\otimes I \geq J$. Hence $Z^{\otimes n}\otimes I \geq J_n$, and every admissible Ω gives $\mathrm{Tr}[J_n\Omega] \leq (\mathrm{Tr}Z)^n/2^n = (1+s)^{2n}/4^n$. Doing nothing attains this value, since $\langle\Phi|J|\Phi\rangle = (1+s)^2/4$. The contrast per loop is $C_n = \log_2[p_{\max}(n)/p_{\min}(n)]/n$. For a noisy device we report the achieved contrast, computed in the same way from the closing probabilities with the loops left alone and with a given veto applied. It bounds the contrast of the device from below, because neither choice need be optimal.

One loop and several loops

One loop has finite contrast for every $\gamma > 0$, so it holds a finite alphabet at any fixed error. Its least closing probability is $(1-s)^2/4$, the known single-copy value (George et al., 2026), which is positive for every $\gamma > 0$ and is attained by the Pauli Z gate. Its contrast $C_1 = 2\log_2[(1+s)/(1-s)]$ is 4.3 bits at $\gamma = 0.6$, 2.8 bits at $\gamma = 0.8$ and 1.9 bits at $\gamma = 0.9$. Two loops operated jointly have more contrast per loop than one for every $0 < \gamma < 1$, and infinite contrast once $\gamma \leq 2\sqrt{2}-2 \approx 0.828$ (Figure 1b,c). Below that value they can be stopped completely, so a wrong record never survives, however many values M the later choice can take. The gate that does so is a controlled-Z gate with phase rotations $-t$ and $+t$ on the two stored qubits. In matrix form it is $D = \mathrm{diag}(1, e^{it}, e^{-it}, -1)$, and equation (3) gives $p(D) = (1-s^2+2s\cos t)^2/16$, the fringe traced by scanning t. It vanishes at $\cos t = -\gamma/(2s)$ whenever $\gamma \leq 2s$, that is $\gamma \leq 2\sqrt{2}-2$. Two separate operations, even with shared randomness, leave at least $(1-s)^4/16$ and so reach at most C_1 per loop.

In general n loops can be stopped completely for $\gamma \leq 1-(2^{1/n}-1)^2$, which is 0.828, 0.932, 0.964 and 0.978 for two to five loops. For every $\gamma < 1$ some finite number of loops, $n \geq \ln 2/\ln(1+s)$, therefore holds an alphabet of any size without error, and in the framework of Ji et al. (2026) the regularised Doeblin information and the asymptotic capacity of an amplitude-damping loop are infinite. The one-copy lower bound of Ji et al. (2026) is therefore not tight for any $0 < \gamma < 1$. The Doeblin information of n loops, $-\log_2[4^n p_{\min}(n)]$, equals $-2\log_2\big(2-(1+s)^n\big)$ while $(1+s)^n < 2$ and is infinite otherwise, so its regularisation is infinite for every $\gamma < 1$. The coefficient is also strictly

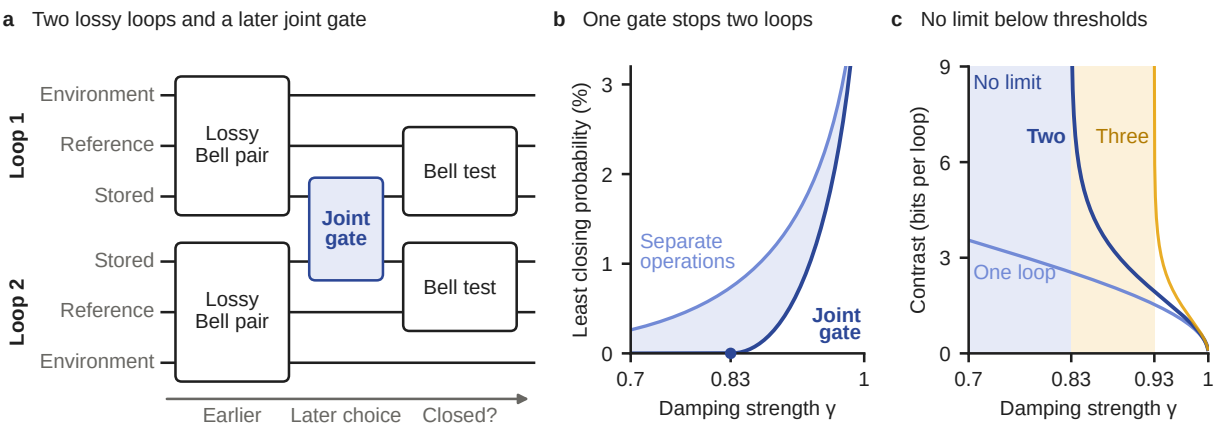


Figure 1. Lossy quantum time loops and the joint veto. a, Each loop is a Bell pair of a stored and a reference qubit whose stored half loses an excitation to an environment qubit with probability γ. A later operation acts on the stored qubits, and a Bell test on each loop decides whether it closes. **b**, Least probability that two loops both close, for separate operations, $(1-s)^4/16$, and for the joint gate, equation (2) with $n=2$. The joint value is zero up to $\gamma = 2\sqrt{2}-2 \approx 0.83$ (dot). **c**, Contrast per loop, C_n, for one, two and three loops. Shading marks where the contrast is infinite and no curve is drawn: up to $\gamma \approx 0.83$ for two loops (blue shading, where three loops also have none) and up to $\gamma \approx 0.93$ for three loops (orange shading).

submultiplicative: for $0 < \gamma < 1$, $p_{\min}(n+m) < p_{\min}(n)\, p_{\min}(m)$ whenever the right side is positive, because $(2-a^n)(2-a^m) - (2-a^{n+m}) = 2(a^n-1)(a^m-1) > 0$ with $a = 1+s$. Only a fraction $p_{\max}(n)/M$ of runs survives for an alphabet of M values.

Separable operations

A two-loop operation is separable when Ω is a sum of products of positive operators on the two loops. Local operations with classical communication are of this kind. A nonseparable operation is needed over part of the range found above. For $\gamma \le 1/2$ a separable operation already stops both loops. For $\gamma > 1/2$ no separable operation does, and so no protocol of local operations and classical communication (Supplementary Note S1). Operations with positive partial transpose (Rains, 2001) are those whose Ω stays positive under transposition of one loop. Adding this constraint to programme (1) gives a further result (Supplementary Table S1): from about $\gamma = 0.75$ upward such operations do no better than two separate vetoes, while the entangling gate reaches zero up to 0.828. The effect is analogous to entanglement distillation (Bennett et al., 1996), where for loss-type noise two imperfect pairs can yield a perfect one (Campbell & Benjamin, 2008).

Predicted behaviour on a six-qubit line

A hardware test needs a line of six superconducting qubits (Methods). Each loop takes a third qubit as the environment that receives the lost excitation, and the circuit needs nine controlled-Z gates between neighbours. A simulation in Qiskit predicts what the device should show. Its noise model has relaxation and dephasing times of 260 µs and 140 µs, a controlled-Z error of 0.2% and a readout error of 0.9%, close to calibration medians reported for ibm_kingston (Cobos et al., 2025). The joint gate is compared with a floor that no pair of separate operations can beat, computed by semidefinite programming from the noisy pair states and the noisy closing test. At $\gamma = 0.83$ the joint gate leaves both loops closed in 0.42% of runs against a floor of 0.90%, which lifts the achieved contrast from 2.4 to 2.9 bits per loop (Figure 2a,b). At $\gamma = 0.86$ the gain is 0.6 bits per loop, and the gap of 0.63 percentage points is about 16 standard errors of shot noise at 10^5 shots per setting. A less favourable parameter set still leaves a gain of 0.4 bits per loop there (Supplementary Tables S2 to S5). Under this noise model the implemented gate reaches at most about 2.9 bits of contrast per loop, and below $\gamma \approx 0.75$ the errors it adds outweigh the gain from interference. A scan of the phase t at $\gamma = 0.80$ dips below the floor near $t = 0.9\pi$ (Figure 2c).

Forecast test

With K equally likely values, the share of wrong records among the surviving runs is $(K-1)p_{\text{veto}}/[p_{\text{close}} + (K-1)p_{\text{veto}}]$, where p_{close} and p_{veto} are the closing probabilities with the loops left alone and with the veto applied. Applied to the simulated closing probabilities at $\gamma = 0.83$, the formula gives a three-bit record that is wrong in 11% of surviving runs with the joint gate against 21% with the best separate vetoes. The record then shares 2.2 bits of mutual information with the later choice against 1.7, and about 3% of runs survive (Supplementary Tables S6 and S7).

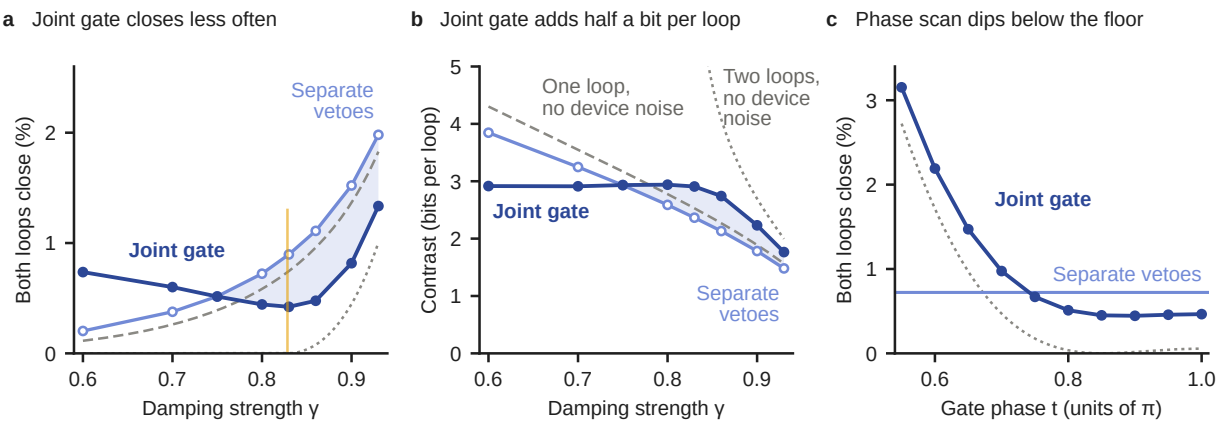


Figure 2. Predicted behaviour on a six-qubit line (simulation). **a**, Probability that both loops close under the typical noise set, for the joint gate at its best phase (filled circles) and for the best separate vetoes (open circles), which set the floor for all pairs of separate operations. Grey curves show the values without device noise (dashed, separate; dotted, joint), and the vertical line marks the threshold $\gamma = 0.828$. Values include readout error and use no error mitigation. **b**, Achieved contrast per loop, from the same closing probabilities. Grey curves show one loop (dashed) and two loops operated jointly (dotted) without device noise. **c**, Phase scan of the joint gate at $\gamma = 0.80$ against the separate-veto floor (horizontal line). The dotted curve is the fringe without device noise.

Pilot on ibm_kingston

A pilot at $\gamma = 0.86$ compared the joint gate, at the phase $t = \pi$ and without tuning, with the Pauli Z veto on each loop (Methods). Over two blocks, with 40,000 shots for each veto, both loops closed in $1.19 \pm 0.05\%$ of runs under the separate vetoes and in $0.90 \pm 0.05\%$ under the joint gate (Table 1). The difference of 0.30 ± 0.07 percentage points passes the test fixed in advance ($z = 4.1$, one-sided $p = 2 \times 10^{-5}$), and the two blocks, sent two days apart, gave 0.27 and 0.32 points. The simulation predicted 0.63 points, so the measured gain is a little under half of the predicted one: the separate veto closed about as often as predicted, and the joint gate about twice as often. The achieved contrast is 2.25 ± 0.04 bits per loop with the joint gate against 2.05 ± 0.03 with the separate vetoes, and by the formula of the forecast test a three-bit record would be wrong in 23.5% of surviving runs against 29.0%.

Table 1. Pilot on ibm_kingston at $\gamma = 0.86$: fraction of shots in which both loops close (percent, with binomial standard errors), achieved contrast per loop (bits), and the values predicted before the run

	No operation	Separate veto	Joint gate	Difference
Shots	20,000	40,000	40,000	
Measured	20.43 ± 0.29	1.193 ± 0.054	0.898 ± 0.047	$+0.295 \pm 0.072$
Block 1	20.29	1.275	1.005	+0.270
Block 2	20.56	1.110	0.790	+0.320
Predicted, typical noise set	21.33	1.110	0.477	+0.633
Predicted, calibration of the day	21.70	1.046	0.346	+0.700
Achieved contrast, measured		2.05 ± 0.03	2.25 ± 0.04	$+0.21 \pm 0.05$
Achieved contrast, predicted		2.13	2.74	+0.61

Two limits apply. First, the closing fraction varied between the four circuits of the separate veto by more than binomial sampling predicts (Supplementary Note S3). A test on the eight circuit fractions, which allows for that variation, gives $t = 2.5$ with five degrees of freedom and one-sided $p = 0.03$, which does not meet the level of 0.01 fixed in advance. The pilot is therefore preliminary evidence that the joint gate closes less often than the separate veto as implemented, and no more than that. Second, each loop on its own reached 2.00 ± 0.02 and 2.15 ± 0.02 bits under the Pauli Z veto, so the joint gate exceeds the better single loop by about two standard errors only. The pilot does not test the floor for all pairs of separate operations, and it covers one loss value and one phase.

Discussion

A forecast test of this kind sends nothing to the past. The complete early record is independent of the later choice, and the agreement appears only in the runs that the closing test selects. In forecasting terms (Catt, 2009), where

forecastability is the mutual information between a future outcome and the information available now at a given horizon (Catt, 2026), the forecastability of the later choice from the early record is therefore zero, and their mutual information is positive only among the runs that are kept. Nothing about the later choice is knowable at the earlier time: the record becomes informative only once the outcomes of the closing tests are known. The loop is a physical model of selection. A set of forecasts kept because they came true shows an agreement with the future that the forecasts never had. The theorem shows that, under energy loss, this selection can be made exact only by acting on several loops together.

The gain depends on the kind of noise. For depolarising noise the coefficient is multiplicative (George et al., 2026; Ji et al., 2026), so acting on the loops together brings no improvement. Loss with thermal excitation, for which the coefficient is known to be multiplicative over part of the parameter range (George et al., 2026), has not yet been solved. The six-qubit predictions rest on a noise model that omits crosstalk, leakage, coherent gate errors, readout asymmetry and drift, and the pilot shows the consequence: the joint gate closed about twice as often on the device as in the simulation. Its phase was not tuned, which is the first thing to examine. The planned test (Methods) scans the phase, fixes the protocol and shot numbers before any data are taken and compares the joint gate with a floor built from measured states.

Conclusion

Equation (2) gives the least closing probability of any number of loops under energy loss. One lossy loop holds only a finite alphabet at any fixed error. Loops operated together have more contrast per loop, and below a threshold that rises with their number they hold an alphabet of any size without error, so that a forecast test built on them passes no wrong forecast. Under the simulated noise model the implemented circuit reaches about 2.9 bits of contrast per loop. In a pilot on an IBM Heron processor the joint gate closed less often than a separate veto on each loop, with a gain a little under half of that simulated. The planned hardware test will measure the floor for all separate operations.

Acknowledgements

This study received no funding.

Data and code availability

The simulation outputs and pilot counts behind both figures and all tables, and the code that generates them (Python 3, Qiskit 2.6.2, Qiskit Aer 0.18.2, CVXPY 1.9.3), are deposited at Zenodo, https://doi.org/10.5281/zenodo.24326287.

Competing interests

The author declares no financial or non-financial competing interests.

References

Arvidsson-Shukur, D. R. M., McConnell, A. G., & Yunger Halpern, N. (2023). Nonclassical advantage in metrology established via quantum simulations of hypothetical closed timelike curves. *Physical Review Letters, 131*(15), Article 150202. https://doi.org/10.1103/PhysRevLett.131.150202

Bennett, C. H., Brassard, G., Popescu, S., Schumacher, B., Smolin, J. A., & Wootters, W. K. (1996). Purification of noisy entanglement and faithful teleportation via noisy channels. *Physical Review Letters, 76*(5), 722–725. https://doi.org/10.1103/PhysRevLett.76.722

Campbell, E. T., & Benjamin, S. C. (2008). Measurement-based entanglement under conditions of extreme photon loss. *Physical Review Letters, 101*(13), Article 130502. https://doi.org/10.1103/PhysRevLett.101.130502

Catt, P. M. (2009). Forecastability: Insights from physics, graphical decomposition, and information theory. *Foresight: The International Journal of Applied Forecasting, 13*, 24–33.

Catt, P. M. (2026). *Forecastability as an information-theoretic limit on prediction* (arXiv:2603.27074). arXiv. https://doi.org/10.48550/arXiv.2603.27074

Cobos, J., Fraxanet, J., Benito, C., di Marcantonio, F., Rivero, P., Kapás, K., Werner, M. A., Legeza, Ö., Bermudez, A., & Rico, E. (2025). *Real-time dynamics in a (2+1)-D gauge theory: The stringy nature on a superconducting quantum simulator* (arXiv:2507.08088). arXiv. https://doi.org/10.48550/arXiv.2507.08088

Diamond, S., & Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research, 17*(83), 1–5. https://jmlr.org/papers/v17/15-408.html

George, I., Hirche, C., Nuradha, T., & Wilde, M. M. (2026). Quantum Doeblin coefficients: Interpretations and applications. *Quantum, 10*, Article 2115. https://doi.org/10.22331/q-2026-05-22-2115

Horodecki, M., Horodecki, P., & Horodecki, R. (1999). General teleportation channel, singlet fraction, and quasidistillation. *Physical Review A, 60*(3), 1888–1898. https://doi.org/10.1103/PhysRevA.60.1888

Huang, Y.-T., Huang, H.-W., Lin, J.-D., Miranowicz, A., Lambert, N., Chen, G.-Y., Nori, F., & Chen, Y.-N. (2026). Experimental simulation of postselected closed timelike curves for decoding scrambled quantum information. *Physical Review Research, 8*(2), Article 023084. https://doi.org/10.1103/tm83-sxpm

Javadi-Abhari, A., Treinish, M., Krsulich, K., Wood, C. J., Lishman, J., Gacon, J., Martiel, S., Nation, P. D., Bishop, L. S., Cross, A. W., Johnson, B. R., & Gambetta, J. M. (2024). *Quantum computing with Qiskit* (arXiv:2405.08810). arXiv. https://doi.org/10.48550/arXiv.2405.08810

Ji, K., Lloyd, S., & Wilde, M. M. (2026). Retrocausal capacity of a quantum channel: Communicating through noisy closed timelike curves. *Physical Review Letters, 136*(23), Article 230801. https://doi.org/10.1103/znyd-npk5

Lloyd, S., Maccone, L., Garcia-Patron, R., Giovannetti, V., & Shikano, Y. (2011a). Quantum mechanics of time travel through post-selected teleportation. *Physical Review D, 84*(2), Article 025007. https://doi.org/10.1103/PhysRevD.84.025007

Lloyd, S., Maccone, L., Garcia-Patron, R., Giovannetti, V., Shikano, Y., Pirandola, S., Rozema, L. A., Darabi, A., Soudagar, Y., Shalm, L. K., & Steinberg, A. M. (2011b). Closed timelike curves via postselection: Theory and experimental test of consistency. *Physical Review Letters, 106*(4), Article 040403. https://doi.org/10.1103/PhysRevLett.106.040403

Nielsen, M. A., & Chuang, I. L. (2010). *Quantum computation and quantum information* (10th anniversary ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511976667

Rains, E. M. (2001). A semidefinite program for distillable entanglement. *IEEE Transactions on Information Theory, 47*(7), 2921–2933. https://doi.org/10.1109/18.959270

Svetlichny, G. (2009). *Effective quantum time travel* (arXiv:0902.4898). arXiv. https://doi.org/10.48550/arXiv.0902.4898

Supplementary Information

Supplementary Note S1. Separable operations and positive partial transpose

Proposition. Call a two-loop operation separable when Ω is a sum of products $P_i \otimes Q_i$ of positive operators on loop 1 and loop 2. A separable operation with closing probability zero exists if and only if $\gamma \leq 1/2$.

Proof. Zero closing probability forces every product to have P_i or Q_i supported on the kernel of J, which is spanned by $|10\rangle$ and $s|00\rangle - |11\rangle$. A positive operator supported there has a reduced operator ρ on the stored qubit with $\langle 0|\rho|0\rangle \leq s^2\langle 1|\rho|1\rangle$. Split the sum by which factor lies in the kernel. The products of the first kind contribute to the entry of $\mathrm{Tr}_R\Omega$ at $|00\rangle$ at most s^2 times what they contribute at $|10\rangle$, and those of the second kind at most s^2 times what they contribute at $|01\rangle$. Every diagonal entry of $\mathrm{Tr}_R\Omega$ equals $1/4$, so $1/4 \leq s^2/4 + s^2/4$, that is $\gamma \leq 1/2$.

Conversely, for $\gamma \leq 1/2$ let $P_a = (s|00\rangle - |11\rangle)(s\langle 00| - \langle 11|)$, $P_b = |10\rangle\langle 10|$, $Q_a = |00\rangle\langle 00|/(8s^2)$ and $Q_b = \big(1/4 - 1/(8s^2)\big)|00\rangle\langle 00| + |10\rangle\langle 10|/8$. Then $\Omega = P_a \otimes Q_a + P_b \otimes Q_b + Q_a \otimes P_a + Q_b \otimes P_b$ is positive, has $\mathrm{Tr}_R\Omega = I/4$ and gives closing probability zero. The construction was also checked numerically at $\gamma = 0.1$, 0.3 and 0.5.

Supplementary Table S1 gives the least probability that two loops both close when the operation is restricted to positive partial transpose, from programme (1) with that added constraint. Up to $\gamma = 1/2$ the value is zero, in line with the proposition. From about $\gamma = 0.75$ it equals the value for two separate operations, while the joint gate gives zero over the whole range of the table.

Supplementary Table S1. Least probability that two loops both close (percent, without device noise)

Damping γ	Positive partial transpose	Separate operations	Joint gate
0.30	0.0000	0.0044	0
0.50	0.0000	0.0460	0
0.55	0.0091	0.0734	0
0.60	0.0418	0.1141	0
0.65	0.1084	0.1739	0
0.70	0.2223	0.2615	0
0.75	0.3906	0.3906	0
0.80	0.5836	0.5836	0
0.828	0.7333	0.7333	0

Supplementary Note S2. Simulation tables

All values are percentages and include readout error unless stated. The floor is the least value that any pair of separate operations can reach on the simulated noisy device (Methods).

Supplementary Table S2. Probability that both loops close, typical noise set

Damping γ	No operation	Separate veto (floor)	Joint gate	Gap	Gap with ideal readout
0.70	34.03	0.377	0.601	−0.22	0.00
0.75	30.05	0.516	0.515	0.00	+0.17
0.80	26.11	0.722	0.443	+0.28	+0.40
0.83	23.78	0.896	0.422	+0.47	+0.57
0.86	21.33	1.110	0.477	+0.63	+0.72
0.90	18.03	1.522	0.817	+0.70	+0.77
0.93	15.43	1.981	1.335	+0.65	+0.70

From $\gamma = 0.83$ upward the joint gate is within 0.05 points of the best joint operation that the noisy states allow. At $\gamma = 0.80$ the best joint operation would reach 0.31% against 0.44% for the gate. At $\gamma = 0.86$ the gap of 0.63 points has a binomial standard error of 0.04 points at 10^5 shots per setting and 0.09 points at 2×10^4.

Supplementary Table S3. Gap between the floor and the joint gate (percentage points) under the three noise sets

Noise set	$\gamma = 0.80$	$\gamma = 0.83$	$\gamma = 0.86$	$\gamma = 0.90$
Favourable	+0.39	+0.58	+0.73	+0.79
Typical	+0.28	+0.47	+0.63	+0.70
Unfavourable	+0.13	+0.34	+0.51	+0.59

Supplementary Table S4. Contrast per loop without device noise, achieved contrast per loop with typical noise (bits), and the number of values of a later choice held at 10% error

Damping γ	One loop, no device noise	Two loops jointly, no device noise	Separate vetoes, typical noise	Joint gate, typical noise	Values at 10% error, separate	Values at 10% error, joint
0.70	3.55	no limit	3.25	2.91	11	7
0.75	3.17	no limit	2.93	2.93	7	7
0.80	2.78	no limit	2.59	2.94	5	7
0.83	2.53	8.53	2.36	2.91	3	7
0.86	2.27	4.08	2.13	2.74	3	5
0.90	1.89	2.69	1.78	2.23	2	3
0.93	1.56	2.00	1.48	1.77	1	2

Supplementary Table S5. Gain in achieved contrast per loop from the joint gate (bits) under the three noise sets

Noise set	$\gamma = 0.80$	$\gamma = 0.83$	$\gamma = 0.86$	$\gamma = 0.90$
Favourable	+0.65	+0.84	+0.85	+0.56
Typical	+0.35	+0.54	+0.61	+0.45
Unfavourable	+0.12	+0.30	+0.39	+0.33

Supplementary Table S6. Forecast test: share of wrong records among surviving runs (percent), typical noise set

Damping γ	1 bit, separate	1 bit, joint	2 bits, separate	2 bits, joint	3 bits, separate	3 bits, joint
0.80	2.7	1.7	7.7	4.8	16.2	10.6
0.83	3.6	1.7	10.2	5.1	20.9	11.0
0.86	4.9	2.2	13.5	6.3	26.7	13.5
0.90	7.8	4.3	20.2	12.0	37.1	24.1

With the favourable set the three-bit shares at $\gamma = 0.83$ are 19.5% and 7.0%. About 12% of runs survive in the one-bit test at $\gamma = 0.83$ and about 3% in the three-bit test.

Supplementary Table S7. Three-bit records under typical noise: wrong share (percent), bits carried about the later choice on surviving runs, and runs surviving (percent)

Damping γ	Wrong share, separate	Wrong share, joint	Bits carried, separate	Bits carried, joint	Runs surviving, separate	Runs surviving, joint
0.80	16.2	10.6	1.91	2.21	3.90	3.65
0.83	20.9	11.0	1.67	2.19	3.76	3.34
0.86	26.7	13.5	1.41	2.05	3.64	3.08
0.90	37.1	24.1	1.01	1.53	3.58	2.97

Supplementary Table S8. Waiting before the later operation in the one-bit test, $\gamma = 0.86$, typical noise set

Wait	Separate veto	Joint gate	1 bit wrong, separate	1 bit wrong, joint
none	1.110	0.477	4.9	2.2
1 µs	1.176	0.577	5.3	2.7
2 µs	1.243	0.677	5.6	3.1
4 µs	1.379	0.876	6.3	4.1

The dynamic circuit of the Methods implements the one-bit task only; the two- and three-bit values in Supplementary Tables S6 and S7 come from the formula applied to simulated closing probabilities. A sampled simulation of the dynamic circuit, with 150,000 runs at $\gamma = 0.86$ and a wait of 2 µs, gave 5.6 ± 0.2% wrong records with separate vetoes and 3.0 ± 0.1% with the joint gate.

Supplementary Note S3. Pilot run on ibm_kingston

The pilot used qubits 148 to 153 of ibm_kingston in the order environment, reference, stored, stored, reference, environment, chosen from 422 candidate lines by the script of the plan on 15 November 2026, from the calibration of that day and immediately before the first job. The plan itself was written on 14 November 2026 and was not changed. On that calibration the controlled-Z errors of the five couplings were 0.12% to 0.34%, the readout errors of the four measured qubits 0.24% to 0.73%, the relaxation times 210 to 388 µs and the dephasing times 96 to 573 µs. The transpiled circuits have 8 controlled-Z gates for no operation and for the separate veto and 9 for the joint gate. A simulation from that calibration, made before the first job, predicted closing fractions of 21.70%, 1.046% and 0.346% for no operation, the separate veto and the joint gate.

Block 1 is job dbsiu37v5hds73dj0vpg, sent on 15 November 2026, and block 2 is job dbu9vp2e8crs738b9n9g, sent on 17 November 2026. Each used 15 seconds of processor time. A second job was cancelled by the author while queued, returned no data, and was replaced by an identical job in the same circuit order. Supplementary Table S9 gives the count of every circuit in run order.

The four circuits of the separate veto closed in 0.92% to 1.38% of shots, a spread larger than binomial sampling predicts ($\chi^2 = 10.3$ with three degrees of freedom, $p = 0.02$); the four circuits of the joint gate are consistent with it ($\chi^2 = 5.4$, $p = 0.15$). The binomial standard error of the pooled difference therefore understates its uncertainty. Taking the eight circuit fractions as the observations gives a difference of 0.295 ± 0.119 points ($t = 2.5$, five degrees of freedom by Welch's formula, one-sided $p = 0.03$), and an exact permutation test over the circuits gives $p = 0.04$. Every shot records both loops. Loop 1 alone closed in 44.7% of shots under no operation and 11.1% under the Pauli Z veto, and loop 2 in 45.8% and 10.3%, against 47.2% and 9.8% without device noise. The single-loop contrasts are 2.00 ± 0.02 and 2.15 ± 0.02 bits. The product of the two single-loop fractions is 20.47% under no operation against 20.43% for both loops, and 1.153% against 1.193% under the separate veto, so the loops behave nearly independently. By the formula of the forecast test, the share of wrong records among surviving runs would be 5.5%, 14.9% and 29.0% for one, two and three bits with the separate vetoes, and 4.2%, 11.6% and 23.5% with the joint gate.

Supplementary Table S9. Pilot on ibm_kingston: shots in which both loops closed, of 10,000 per circuit, in run order

	Circuit 1	Circuit 2	Circuit 3	Circuit 4	Circuit 5
Block 1, arm	separate	joint	none	joint	separate
Block 1, closed	138	98	2,029	103	117
Block 2, arm	joint	separate	none	separate	joint
Block 2, closed	80	130	2,056	92	78