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.

Quantum amplitude amplification of improbable futures

Peter M. Catt

peter.catt@virtualblue.co.nz

Research Article

Keywords: Quantum computing, Amplitude amplification, Rare-event sampling, Markov processes, Entropy rate, Superconducting qubits, Qiskit

Posted Date: November 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11588189/v1

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

Additional Declarations: No competing interests reported.

Quantum amplitude amplification of improbable futures

Peter M. Catt*

Virtual Blue Limited, Auckland, New Zealand · ORCID 0009-0000-1324-0861

*Corresponding author: peter.catt@virtualblue.co.nz

Abstract

Some consequential futures consist of plausible steps arranged into an unlikely sequence, which simulation seldom produces. Here ideal amplitude amplification, which raises the chance of drawing such futures while preserving their relative likelihoods, becomes a gate-level circuit that marks a future by its coherently accumulated surprisal, with a finite-precision analysis of the marking and the circuit's logical cost. For a five-state Markov process six steps ahead, a 30-qubit circuit verified against exact enumeration raised the share of rare futures drawn from 2.5% to 98.0% in four iterations, leaving their relative likelihoods unchanged. The known process is an exactly verifiable testbed; the claim is a construction and a hardware demonstration, not a computational advantage. On IBM Kingston, compact reference circuits with matched controls showed the structure of amplification at 228 routed two-qubit gates ($z = 50$, per-future $\rho = 0.98$) but not at about 1,170, including after a pre-registered test of noise suppression.

Keywords: Quantum computing; Amplitude amplification; Rare-event sampling; Markov processes; Entropy rate; Superconducting qubits; Qiskit

Introduction

Many consequential outcomes are path dependent. Droughts and floods can arise from combinations of drivers that are individually unremarkable [1], financial distress

spreads through chains of exposure between institutions [2], and power-grid blackouts propagate through sequences of individual line failures [3]. Classical rare-event simulation, through importance sampling, multilevel splitting and genealogical particle methods, is highly effective once the event of interest has been specified [4–8]. Guo et al. [9] consider a different target: every outcome whose own probability is at most a threshold Δ, without naming those outcomes in advance. For stationary processes they set $\Delta = 2^{-\alpha L}$ for trajectories of length L, with α at least the entropy rate, so that a trajectory is rare when its surprisal per step is at least α. They show that a classical algorithm with sample access needs $\Omega(1/\Delta)$ samples to sample from the rare-conditioned distribution, whereas a quantum algorithm with coherent access to a state-preparation unitary and its inverse needs $\tilde{\mathcal{O}}\left(1/\sqrt{p_{\mathrm{rare}}\Delta}\right)$ applications of them, where p_{rare} is the total probability of the rare set. This work takes the constructive step: it realises that rarity test as an explicit circuit, verifiable against exact enumeration, and measures what the construction costs and how its amplification step behaves on superconducting hardware.

Guo et al. [9] build the marking of rare outcomes from the state-preparation unitary itself, by applying a polynomial threshold to outcome probabilities, so that rare outcomes never have to be identified, and establish the asymptotic sampling complexity of the approach. Gupta et al. [10] combine quantum stochastic models, which prepare sample states with less memory than any classical model of the process, with amplitude amplification, and establish advantages in both memory and time. Both results concern asymptotic complexity. When the process is known, the marking can be computed from its rules: for a Markov process with a known transition matrix, the surprisal of a trajectory is a sum of per-step terms, so it can be accumulated coherently and compared with the threshold, and amplitude

amplification [11, 12] then raises the probability of the marked futures. Amplitude amplification and estimation already underpin quantum approaches to Monte Carlo estimation and financial risk [13–16], and Grover search has been run on hardware at a few qubits [17, 18].

This work makes four contributions. First, an explicit rarity oracle for stochastic trajectories, which accumulates fixed-point transition surprisals in a Fourier-basis register and reduces the threshold test to the sign of one bit. Second, a finite-precision analysis that separates circuit error, which must be zero, from classification error caused by rounding, with a sufficient precision for exact classification (Proposition 1). Third, the logical cost of the complete circuit as the horizon and the precision grow, including the preparation and reflection steps. Fourth, controlled hardware experiments with matched oracle and control circuits, which detect the predicted amplification structure at 228 routed two-qubit gates but not at about 1,170, including after a pre-registered test of standard noise suppression.

Because the circuit is built from the transition matrix, a classical computer holding the same matrix can compute the rare-set probability, and sample exactly from the fixed-point rare-conditioned distribution, by dynamic programming over states and accumulated fixed-point surprisal; where the fixed-point classification is exact, as at $H = 6$ with $b = 7$, that is also the real-valued target. The known process is the testbed that makes every result checkable against exact values; no quantum advantage is claimed for it.

Methods

Stochastic process and future trajectories

The process has five states and transition matrix

$$T = \begin{pmatrix} 0.88 & 0.05 & 0.03 & 0.025 & 0.015 \\ 0.20 & 0.65 & 0.07 & 0.05 & 0.03 \\ 0.15 & 0.08 & 0.62 & 0.10 & 0.05 \\ 0.10 & 0.08 & 0.12 & 0.62 & 0.08 \\ 0.08 & 0.06 & 0.10 & 0.12 & 0.64 \end{pmatrix}. \quad (1)$$

Its stationary entropy rate is $h_\infty = 1.1535596$ bits per step: the uncertainty in each step that remains once the full past is known [19]. A future $x = (x_1, \ldots, x_H)$ is the sequence of the next H states given $X_0 = 1$, with probability $P_H(x) = \prod_{t=1}^{H} T_{x_{t-1}x_t}$. Following Guo et al. [9] with $\alpha = h_\infty + \delta$, the rare set, its probability and the threshold are

$$\mathcal{R}_\Delta = \{x : 0 < P_H(x) \le \Delta_H\},$$
$$p_{\mathrm{rare}} = \sum_{x \in \mathcal{R}_\Delta} P_H(x), \qquad \Delta_H = 2^{-H(h_\infty + \delta)}, \quad (2)$$

so a future is rare when its surprisal $-\log_2 P_H(x)$ is at least $S_\Delta = H(h_\infty + \delta)$; no future lies exactly on the threshold, so the choice between a strict and a non-strict inequality does not matter here. We set $\delta = 1.5$ bits per step to obtain a low-probability target while keeping the number of ideal amplification iterations modest: at $H = 6$, $p_{\mathrm{rare}} = 0.025$ and the first amplification maximum occurs at $r = 4$. Table S5 shows that smaller δ values produce substantially less rare targets, whereas larger values require materially more amplification iterations. Guo et al. treat stationary processes; here every future instead starts in state 1, the most persistent state, so the threshold is a fixed reference level set by the stationary entropy rate, not an estimate of the expected surprisal from state 1.

Rare futures as atypical sequences

The entropy rate of the chain is

$$h_\infty = -\sum_i \pi_i \sum_j T_{ij} \log_2 T_{ij}, \quad (3)$$

where π is its stationary distribution. By the Shannon–McMillan–Breiman theorem [19, 20], the per-step surprisal $-H^{-1}\log_2 P_H(X)$ of a future drawn from the process converges to h_∞ with probability one, from any starting state. A rare future has per-step surprisal of at least $h_\infty + \delta$, so the rare set lies outside the typical set and p_{rare} vanishes as the horizon grows. The rate is set by the scaled cumulant generating function of the surprisal $S_H = -\log_2 P_H(X)$,

$$\Lambda(\lambda) = \lim_{H\to\infty} \frac{1}{H} \log_2 \mathbb{E}\left[2^{\lambda S_H}\right] = \log_2 \rho\left(T^{(1-\lambda)}\right), \qquad \left(T^{(1-\lambda)}\right)_{ij} = T_{ij}^{1-\lambda}, \tag{4}$$

where ρ is the spectral radius and $\Lambda'(0) = h_\infty$. By the Gärtner–Ellis theorem [21],

$$p_{\mathrm{rare}} = 2^{-H\, I(h_\infty+\delta)+o(H)}, \qquad I(s) = \sup_{\lambda}[\lambda s - \Lambda(\lambda)]. \tag{5}$$

For the chain of equation (1), $I(h_\infty + 1.5) = 0.439$ bits per step, attained at $\lambda^* = 0.648$. Exact enumeration up to $H = 10$, and a transfer-matrix calculation on a grid of 2×10^{-4} bits beyond it, show the approach to this rate: $-H^{-1}\log_2 p_{\mathrm{rare}}$ falls from 0.887 at $H = 6$ to 0.611 at $H = 20$ and 0.478 at $H = 120$, approaching I from above as the sub-exponential factor in equation (5) loses weight. The rare futures are nevertheless the majority by count, 88.9% of the 15,625 futures at $H = 6$ (Fig. 1b): most sequences contain some unlikely transitions, while most of the probability lies on futures that remain in persistent states. The probability mass the algorithm targets therefore falls exponentially with the horizon, while ideal amplitude amplification changes the cost of reaching it from about $1/p_{\mathrm{rare}}$ draws to about $\pi/\left(4\sqrt{p_{\mathrm{rare}}}\right)$ coherent iterations (equation (13)).

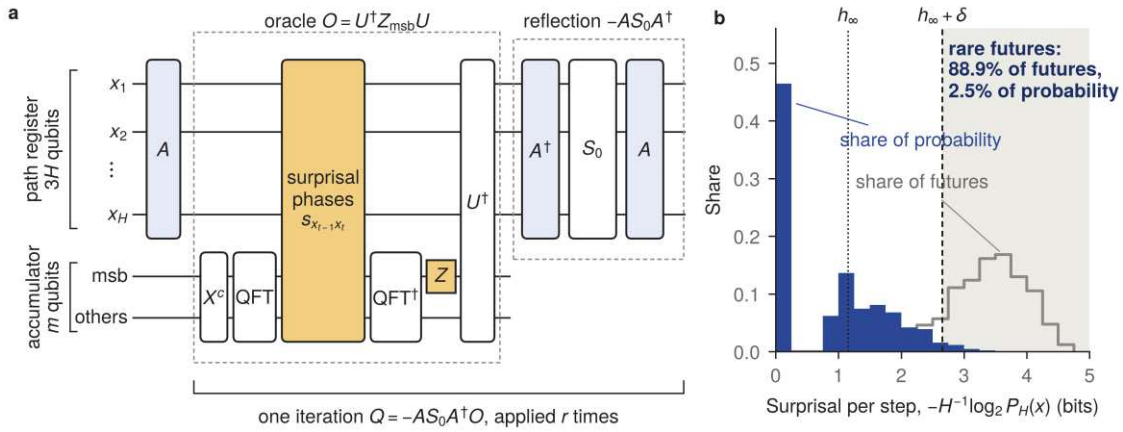


Figure 1

Figure 1. The oracle marks rare futures, which are most sequences but little probability. a, One amplification iteration $Q = -AS_0A^\dagger O$ acting on the path register (three qubits per step) and the accumulator (m qubits). A prepares $|P_H\rangle$. The oracle $O = U^\dagger Z_{\mathrm{msb}} U$ loads the offset c, applies the quantum Fourier transform, adds the fixed-point surprisal $s_{x_{t-1}x_t}$ of each transition as phase rotations conditioned on the step registers, returns to the computational basis, changes the sign when the most significant accumulator bit is 1, and uncomputes. The reflection about the prepared state applies $A^\dagger$, the reflection S_0 about the all-zero state, and A. Pale blue blocks prepare or unprepare the state; amber blocks mark rare futures. **b**, Per-step surprisal of the 15,625 futures at $H = 6$, as a share of futures (grey line) and as a share of probability (blue bars), in bins of 0.25 bits. The dotted line marks the entropy rate h_∞ and the dashed line the rare threshold $h_\infty + \delta$; the shaded region holds the rare futures.

Encoding and state preparation

Each step is held in three qubits, of which five basis states encode the process states and three are unused, so the path register has $3H$ qubits. The trajectory distribution is prepared as

$$|P_H\rangle = \sum_x \sqrt{P_H(x)}\ |x\rangle \tag{6}$$

by a cascade of uniformly controlled R_y rotations [22–24]. Within each step the three bits are prepared from most to least significant, each conditioned on the previous step's register and on the bits already prepared. Rotations conditioned on unused encodings are set to zero, so the prepared state has no amplitude on invalid encodings. At $H = 6$ the prepared probabilities matched the enumerated $P_6(x)$ to within 2.4×10^{-15} in every basis state.

Fixed-point surprisal

Transition surprisals and the threshold are held to b fractional bits,

$$s_{ij} = \mathrm{round}\left(2^b\left(-\log_2 T_{ij}\right)\right), \qquad T_b = \lceil 2^b S_\Delta \rceil. \tag{7}$$

With $S_b(x) = \sum_t s_{x_{t-1}x_t}$ the fixed-point surprisal of a future, the fixed-point test $S_b(x) \geq T_b$ can disagree with the real-valued definition for futures close to the threshold. Each s_{ij} lies within $1/2$ of $2^b\left(-\log_2 T_{ij}\right)$ and T_b lies within 1 of $2^b S_\Delta$, so

$$|2^{-b} S_b(x) + \log_2 P_H(x)| \leq H\, 2^{-(b+1)}, \qquad 0 \leq 2^{-b} T_b - S_\Delta < 2^{-b}. \tag{8}$$

Proposition 1 (precision for exact classification). *Let g be the smallest distance, in bits, between the threshold S_Δ and the surprisal of any valid future. If $b > \log_2[(H/2+1)/g]$, the fixed-point test $S_b(x) \geq T_b$ agrees with the real-valued definition for every future.*

Proof. By equation (8), the rescaled fixed-point surprisal $2^{-b} S_b(x)$ lies within $H\, 2^{-(b+1)}$ of the true surprisal, and the rescaled threshold $2^{-b} T_b$ lies within 2^{-b} above S_Δ. The order of the two can change only if the shifts together reach g, which needs $(H/2 + 1)\, 2^{-b} \geq g$. ▫

The condition is sufficient, not necessary. Below it, agreement depends on how the rounding errors fall and has to be checked for each b; at $H = 6$ the bound gives $b = 13$, whereas $b = 7$ already classifies every future correctly (Results). Two kinds of error are therefore recorded separately: a circuit error, where the circuit's result differs from the fixed-point test, which must be zero; and a precision error, where the fixed-point test differs from the real-valued rare set, reported as a count and as misclassified probability mass.

Rarity oracle

An accumulator register of m qubits is initialised to an offset c, transformed to the Fourier basis, and receives each transition surprisal as phase rotations [25, 26]: for accumulator qubit j, a diagonal gate [27] on that qubit and the step registers the transition depends on applies the phase $2\pi s/2^{j+1}$. The inverse transform then leaves $c + \sum s \bmod 2^m$ in the register. With

$$\begin{gathered} 2^{m-1} \geq \max(T_b,\ S_{\max} - T_b + 1), \\ c = 2^{m-1} - T_b, \end{gathered} \tag{9}$$

where $S_{\max}$ is the largest fixed-point surprisal over valid futures. In operator form, with U the accumulation step (offset, Fourier transform, phase rotations and inverse transform),

$$\begin{gathered} U|x\rangle|0\rangle_m = |x\rangle|c + S_b(x)\rangle_m, \qquad O = U^\dagger (I \otimes Z_{\text{msb}}) U, \\ O|x\rangle|0\rangle_m = (-1)^{[S_b(x) \geq T_b]}\,|x\rangle|0\rangle_m, \end{gathered} \tag{10}$$

where $[\cdot]$ is 1 when its argument holds and 0 otherwise. Equation (9) gives $0 \leq c + S_b(x) < 2^m$ for every valid future, so the sum never wraps and its most significant bit is $[S_b(x) \geq T_b]$: no comparator and no validity flag are needed. Unused encodings are given zero surprisal; they never carry amplitude, because the preparation assigns

them none and both the oracle and the reflection about $|P_H\rangle$ preserve the span of valid basis states.

Amplitude amplification

Write the prepared state of equation (6) as

$$|P_H\rangle = \sin\theta\,|\psi_{\mathcal{R}}\rangle + \cos\theta\,|\psi_{\bar{\mathcal{R}}}\rangle, \qquad |\psi_{\mathcal{R}}\rangle = \frac{1}{\sqrt{p}} \sum_{x\in\mathcal{R}_\Delta} \sqrt{P_H(x)}\;|x\rangle, \tag{11}$$

where $\sin^2\theta = p$ is the probability of the marked set and $|\psi_{\bar{\mathcal{R}}}\rangle$ is the normalised component on unmarked futures. With A the preparation, $S_0 = I - 2|0\rangle\langle 0|$ the reflection about the all-zero state of the path register and O the oracle, one iteration is $Q = -AS_0A^\dagger O$ (Fig. 1a). Q rotates the plane spanned by $|\psi_{\mathcal{R}}\rangle$ and $|\psi_{\bar{\mathcal{R}}}\rangle$ by 2θ [11, 12], so

$$Q^r|P_H\rangle = \sin[(2r+1)\theta]\,|\psi_{\mathcal{R}}\rangle + \cos[(2r+1)\theta]\,|\psi_{\bar{\mathcal{R}}}\rangle, \tag{12}$$

and after r iterations each future is drawn with probability

$$P_r(x) = \begin{cases} P_H(x)\,\dfrac{\sin^2[(2r+1)\theta]}{\sin^2\theta}, & x \in \mathcal{R}_\Delta, \\ P_H(x)\,\dfrac{\cos^2[(2r+1)\theta]}{\cos^2\theta}, & x \notin \mathcal{R}_\Delta. \end{cases} \tag{13}$$

The probability of a rare future is therefore $P_r = \sin^2[(2r+1)\theta]$, first maximal near $r^* = \lfloor \pi/(4\theta) \rfloor$ [12], and the distribution within the rare set, $P_r(x \mid \mathcal{R}_\Delta) = P_H(x)/p$, is the conditional distribution of the process for every r: amplification changes how often rare futures are drawn, not which of them are drawn. At every horizon studied here r^* is also the first integer at which P_r reaches a local maximum. At $r = 1$, equation (13) gives the ideal change in each future's probability used to test the hardware output, and the ideal hardware signal $\sin^2 3\theta - p$. Where the process can only be sampled, as in the setting of Guo et al. [9], a rare future takes $1/p$ draws on average, against

about $\pi/\left(4\sqrt{p}\right)$ iterations of Q; by equation (5), the first grows as 2^{HI} and the second as $2^{HI/2}$.

Validation

Three checks were applied at $H = 6$. First, for every one of the 15,625 valid futures, the path register was set to that basis state, the accumulator circuit was run in Qiskit Aer's matrix-product-state simulator [28], and the full accumulator register was read and compared with $c + \sum s \bmod 2^m$. Second, the complete amplification circuit, from preparation through four iterations, was simulated as a statevector at $b = 7$ and at $b = 4$, and after each iteration the probability of every path-register basis state and of the accumulator's zero state was recorded and compared with exact amplitude amplification of the enumerated distribution under the same fixed-point marking. Third, the probability of a rare future under the real-valued definition was computed directly from the recorded probabilities, never inferred from the analytical formula.

Resource accounting

Logical resources were counted after transpiling the same circuits to the basis $\{R_z, \sqrt{X}, X, \mathrm{CNOT}\}$ with all-to-all connectivity at optimisation level 3. Counts are reported at a fixed precision, $b = 4$, and at the lowest b with zero mismatches.

For a K-state process held in $q = \lceil \log_2 K \rceil$ qubits per step ($q = 3$ here), the circuit uses qH path qubits and m accumulator qubits, $qH + m$ in all, with m set by the threshold and the largest accumulated surprisal (equation (9)). A future of length H has H transitions. The first leaves the fixed start state, so its phase rotations are conditioned on one step register and cost m diagonal gates on $q + 1$ qubits; each of the remaining $H - 1$ transitions couples two consecutive step registers to an accumulator qubit, giving $(H - 1)m$ diagonal gates on $2q + 1$ qubits in the

computation and as many again in the uncomputation, so the oracle's cost grows as Hm before synthesis. For r iterations the two-qubit cost of the whole circuit is

$$C = C_A + r\,(C_O + C_R), \tag{14}$$

where C_A is the preparation, C_O the oracle and C_R the reflection about the prepared state, which holds two preparations and the reflection about zero. In the counts of Table S3, C_A grows linearly from 62 two-qubit gates at $H = 2$ to 286 at $H = 6$ and C_R from 152 to 696, so the oracle is more than 90% of every iteration and the preparation is under 1% of the total at $H = 6$.

Hardware reference and matched control

The full circuit is too deep for present hardware, so the amplification step was tested with compact references at $H = 2$ and $H = 3$. Each indexes complete futures in base 5, using 5 and 7 qubits respectively, prepares the trajectory distribution with a general state-preparation circuit, and marks rare futures with a diagonal gate built from the enumerated rare set. These references serve only to test amplification on hardware and do not scale. Because rare futures are the majority by count (21 of 25 at $H = 2$, 109 of 125 at $H = 3$), noise that flattens the output distribution raises the raw rare-hit rate on its own. The primary statistic is therefore the difference between each circuit and a matched control: the same circuit with a diagonal gate that is the identity on all valid futures and has random phases on unused indices, compiled to the same physical qubits. An identity would be removed by the compiler; the random phases, on indices that carry no amplitude in the ideal circuit, make it synthesise a diagonal gate on the same qubits, of comparable routed two-qubit cost. Noisy simulations used noise models built from IBM calibration snapshots of Fez, Kingston and Marrakesh, retrieved on 29 October 2026, with 3 repeats of 20,000 shots per arm.

The hardware test was fixed before any run: a one-sided, pooled two-proportion z test of oracle against control, passed when $z \geq 5$, using four alternating oracle and control pairs of 5,000 shots each, giving 20,000 shots per arm, submitted as eight primitive unified blocs (PUBs) of a single Sampler job for each horizon, on the backend chosen from the noisy-simulation results before submission. Control and oracle circuits must agree in routed two-qubit count to within 2%. Matching the routed two-qubit count controls a major contribution to each circuit's exposure to noise but does not make the physical noise identical, because gate placement, timing, coherent errors and drift can differ between the two circuits. The control is the identity on the valid subspace only in the noiseless circuit, where the unused indices carry no amplitude. Errors that populate those indices, 7 of 32 at $H = 2$ and 3 of 128 at $H = 3$, make the control's random phases dynamically relevant, so the measured difference is not a function of the rare-future marking alone. The size of that effect was tested by repeating the noisy simulation with 24 further independent phase seeds at each horizon, under criteria fixed and hashed before any result was seen and deposited with the data. Each control was compiled to the oracle's layout and kept only if it met the same 2% matching rule; all were kept at both horizons (Results, Table S6).

Software and hardware

Circuits were built, simulated and compiled for IBM hardware with Qiskit [29] and Qiskit Aer, and submitted with Qiskit IBM Runtime version 0.51.0. The 30-qubit statevector simulation at $b = 7$ used Qiskit version 2.5.2 and Qiskit Aer version 0.17.2 on an Apple-silicon Mac with 32 GB of memory and took 39 minutes; the calibration-snapshot noisy simulations, and the compilation and submission of the hardware jobs, used the same software on the same machine. All other computations used Qiskit version 2.6.1 and Qiskit Aer version 0.18.2 on two CPU cores of a Linux

machine, where the 27-qubit simulation at $b = 4$ took 17 minutes. NumPy computed the classical reference values: trajectory probabilities, the rare set and its mass. Hardware runs used IBM Kingston (ibm_kingston), a quantum processing unit (QPU) on the IBM Quantum Platform built on a 156-qubit superconducting IBM Heron r2 processor and hosted in the platform's Washington DC (us-east) region. IBM Fez and Marrakesh, QPUs with the same processor type, were used only in the noisy simulations that selected it. The main hardware run used the Sampler primitive's default options, under which no dynamical decoupling, twirling or readout mitigation is applied, so the measured signal reflects the processor's native noise. In the calibration in force when the jobs ran, Kingston's median CZ error was 1.8×10^{-3}, its median readout error 9.8×10^{-3}, and its median T_1 and T_2 were 235 and 150 μs; values for the qubits used are given with Table 2.

Noise-suppression run

A follow-up run on 30 October 2026 tested whether standard error suppression recovers the three-step signal. Its design and analysis were fixed in a pre-registration record written at 00:17:01 UTC on 30 October 2026, 17 minutes before submission, and deposited with the data. The record embeds the Kingston calibration that IBM served at 00:08:59 UTC, and IBM's job records show the first job submitted at 00:34:16 UTC, so the record is bracketed by two IBM-side timestamps and was prospectively fixed before the first hardware submission; it was not lodged with an external preregistration registry. The compiled circuits of the 29 October jobs were resubmitted unchanged, on the same physical qubits, in one IBM Batch under three conditions: native; dynamical decoupling [30, 31] with XpXm sequences scheduled as late as possible; and the same decoupling with Pauli twirling of the two-qubit gates [32, 33] (Qiskit's active-accum strategy, 20 randomisations of 250 shots).

Measurement twirling was off throughout. Twelve jobs ran in four rounds with the order of conditions rotated; each job held one oracle and control pair for each horizon at 5,000 shots per circuit, giving 20,000 shots per arm per horizon in each condition. IBM's records of all twelve jobs confirm the options of each condition, that the submitted circuits are identical to those of 29 October, and that the counts analysed are identical to the recorded shots.

Statistical analyses

All statistics were computed from the raw counts. At each horizon the amplification signal is the oracle-minus-control difference $d = P_{\mathrm{oracle}}(\mathrm{rare}) - P_{\mathrm{control}}(\mathrm{rare})$, where each probability is the fraction of shots that returned a rare future; raw rare-hit rates are reported but are not used as evidence of amplification (Results). The retention of the ideal signal, $R = d/d_{\mathrm{ideal}}$, is reported as a descriptive quantity with no test attached. Detection used the pre-specified one-sided shot-level pooled two-proportion z test fixed before the 29 October run, passed at $z \geq 5$; 95% intervals for the signal use the unpooled standard error. The statistic treats the shots within an arm as independent Bernoulli draws; the experimental units above the shot are the four oracle and control pairs within a run, and the runs themselves, so the pair-level agreement and the between-day comparison reported below are complementary evidence rather than substitutes for the pre-specified test. With k_O and k_C rare outcomes in n_O and n_C shots, $\hat{p}_O = k_O/n_O$, $\hat{p}_C = k_C/n_C$ and $\bar{p} = (k_O + k_C)/(n_O + n_C)$, the statistic and the interval are

$$z = \frac{\hat{p}_O - \hat{p}_C}{\sqrt{\bar{p}(1-\bar{p})(n_O^{-1} + n_C^{-1})}}, \qquad \hat{p}_O - \hat{p}_C \pm 1.96\sqrt{\frac{\hat{p}_O(1-\hat{p}_O)}{n_O} + \frac{\hat{p}_C(1-\hat{p}_C)}{n_C}}, \tag{15}$$

The consistency of the four pair signals within a run was assessed with

$$Q_\chi = \sum_{k=1}^{4} w_k \left(\hat{d}_k - \bar{d}\right)^2, \qquad \bar{d} = \frac{\sum_k w_k \hat{d}_k}{\sum_k w_k}, \tag{16}$$

where $\hat{d}_k$ is the signal of pair k and w_k the inverse of its unpooled variance, compared with a χ^2 distribution on three degrees of freedom. The per-future pattern was measured by the Pearson correlation ρ between each future's change in frequency from control to oracle and its change under ideal amplification (equation (13) at $r = 1$), over all futures and over the rare futures alone. The shape of the rare-conditioned output was measured by its total-variation distance from the target,

$$D_{\mathrm{TV}} = \frac{1}{2} \sum_{x \in \mathcal{R}_\Delta} \left| \frac{f(x)}{\sum_{y \in \mathcal{R}_\Delta} f\,(y)} - P(x \mid \mathcal{R}_\Delta) \right|, \tag{17}$$

where $f(x)$ is the number of shots that returned x; the reference value for sampling alone is the mean distance over 200 samples, each of the observed number of rare hits drawn from the target. In the noise-suppression run, the primary test was one-sided, for an increase in the three-step signal under each form of suppression relative to native, with an unpooled standard error and a threshold of $z \geq 3$ for each of the two comparisons. Three secondary checks were set: that the four round-matched changes share a sign; that the per-future correlation at three steps reaches $\rho \geq 0.30$ and exceeds native; and that the two-step signal under suppression passes $z \geq 5$ and does not fall below native by $z \geq 3$.

Results

Rare sets and precision

At $H = 6$, 13,888 of the 15,625 futures are rare, yet together they carry probability 0.0250 (Table S1 (Additional file 1), Fig. 1b). As the horizon grows their total probability falls towards the large-deviation rate of 0.439 bits per step (Methods). The nearest future lies 5.7×10^{-4} bits from the threshold, so the bound above guarantees

exact classification only from $b = 13$. In practice $b = 7$ already gives zero mismatches, while $b = 4$ misclassifies 9 futures carrying 1.4×10^{-4} of probability, 0.57% of the rare mass (Table S2). Below the guaranteed bound agreement is not monotone in b: $b = 9$ misclassifies two futures even though $b = 7$ and $b = 8$ misclassify none, so exactness at $b = 7$ is a verified property of that choice, not a guaranteed one.

Circuit validation

The accumulator returned the exact register value on all 15,625 valid futures at $b = 7$ and on all 15,625 at $b = 4$, with no errors. With Table S2 and the statevector check below, this means the $b = 7$ oracle marks exactly the real-valued rare set and the $b = 4$ oracle marks exactly its fixed-point set.

Exact agreement of the complete circuit at six steps

The complete 30-qubit circuit, simulated as a statevector at $b = 7$, reproduced exact amplitude amplification of the real-valued rare set future by future: every trajectory probability agreed to within 5.9×10^{-15} at every iteration, the accumulator returned to zero to within 4.2×10^{-14}, and the probability on unused encodings stayed below 2.5×10^{-31} (Fig. 2a, Table 1). After four iterations a rare future was drawn with probability 0.9798, against 0.0250 before amplification, and the amplified rare-conditioned distribution lay within 5.0×10^{-15} of its target in total variation, so the relative likelihoods of the rare futures were unchanged.

The 27-qubit simulation at $b = 4$ followed exact amplification of its fixed-point marked set equally closely, to within 4.4×10^{-15}. Measured against the real-valued target, its final probability was 0.9787 and its rare-conditioned distribution lay 5.7×10^{-3} from the target, the effect of the 9 futures misclassified at $b = 4$, not a circuit error.

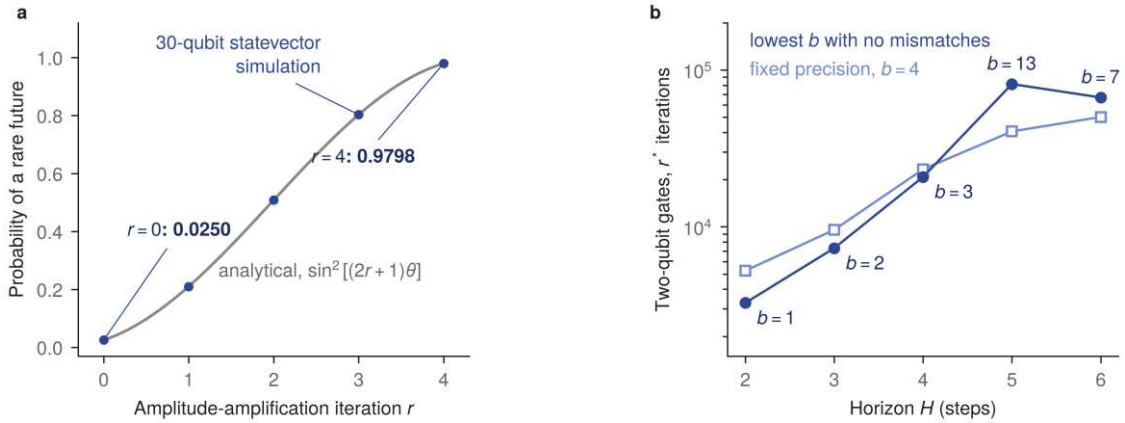


Figure 2

Figure 2. Four iterations raise the rare-future probability to 0.98, and cost grows with horizon. a, Probability of drawing a rare future against amplification iteration r at $H = 6$ ($\delta = 1.5$, $b = 7$). The grey line is $\sin^2[(2r+1)\theta]$ with $\sin^2\theta = p_{\mathrm{rare}}$. Blue points are the complete gate-level Qiskit circuit, simulated as a 30-qubit statevector and computed directly from the simulated probabilities of all 2^{18} path-register states; they lie on the curve to within 4.1×10^{-14}. **b**, Logical two-qubit gates for r^* iterations of the general circuit against horizon, at fixed precision $b = 4$ (grey squares) and at the lowest b with no mismatches (blue circles, labelled with b), after transpiling to $\{R_z, \sqrt{X}, X, \mathrm{CNOT}\}$ with all-to-all connectivity.

Resource scaling with horizon and precision

At fixed $b = 4$ the oracle grows from 2,441 to 11,777 two-qubit gates between $H = 2$ and $H = 6$ (Fig. 2b, Table S3). The cost is dominated by the diagonal gates that couple consecutive step registers to an accumulator qubit, $(H-1)m$ in the computation and as many again in the uncomputation, each on seven qubits and compiled to $2^7 - 2 = 126$ CNOTs [27]. With the Fourier transforms and the first step, the oracle costs 262 to 305 two-qubit gates per transition and accumulator qubit, so it grows in proportion to $(H-1)m$; the step between $H = 3$ and $H = 4$ reflects the accumulator growing from 8 to 9 qubits. The counts are specific to this compilation:

each diagonal gate is synthesised generically at $2^7 - 2$ CNOTs, whereas only $5^2 = 25$ of its 64 control patterns carry a non-zero phase, so sparser synthesis would lower the constant without changing the growth with $(H-1)m$. Exact marking changes the picture through the accumulator width: at $H = 5$ the lowest b with zero mismatches is 13, giving an oracle of 19,725 two-qubit gates, more than at $H = 6$ with $b = 7$ (15,933).

Predicted hardware regime

Noise erodes much of the ideal signal at two steps and most of it at three (Fig. 3, Table S4). At $H = 2$ with one iteration the predicted signal is +0.35 to +0.41 against an ideal +0.65, with z from 70 to 82. At $H = 3$ it falls to +0.039 to +0.060 against an ideal +0.52, with z from 8.6 to 12.5. On Fez and Marrakesh the $H = 3$ control missed the 2% matching tolerance, so only Kingston gives a valid three-step comparison, with a predicted z of 12.5. A second iteration raises the ideal success but lowers the predicted signal at $H = 2$ on all three backends, and at $H = 3$ on all three, although every three-step control at $r = 2$ misses the matching rule. One iteration is therefore the preferred design at both horizons, and Kingston, which gives the largest predicted signal at both, is the chosen backend.

The predicted signal is not materially sensitive to the particular random phases given to the control. Over 24 further phase seeds at each horizon, all of which met the matching rule, the predicted signal varied by a standard deviation of 0.0041 at $H = 2$ and 0.0029 at $H = 3$, 1.3 and 0.9 times the shot-noise standard error of a single arm, and no seed changed its sign (Table S6). Every seed gave a two-step signal above $z = 80$ and a three-step signal of at most +0.0553, under a fifth of the ideal +0.521; within this noise model, the qualitative result at both horizons is therefore robust to the phase choice.

Raw rare-hit rates cannot identify amplification

At $H = 2$ and $H = 3$ rare futures are most of the valid basis states even though they carry little probability under the process, so noise that flattens the output raises the raw rare-hit rate on its own: in the noisy simulations the control alone reaches well above p_{rare}, towards the 0.656 of a uniform output over all indices at $H = 2$ and 0.852 at $H = 3$. A raw rare-event frequency therefore cannot identify amplitude amplification, and every hardware conclusion below rests on the oracle-minus-control contrast and on the per-future pattern of the change.

Table 1. Amplification at $H = 6$ (Qiskit Aer statevector)

b	r	$P(\mathrm{rare})$, real definition	Exact amplification, fixed-point marking	Max error per trajectory	Accumulator leakage	D_{TV} to real-valued target
7	0	0.024977	0.024977	2.3×10^{-15}	n/a	4.6×10^{-15}
7	1	0.210066	0.210066	5.9×10^{-15}	1.3×10^{-14}	4.6×10^{-15}
7	2	0.508127	0.508127	3.8×10^{-15}	1.9×10^{-14}	4.7×10^{-15}
7	3	0.803021	0.803021	1.4×10^{-15}	3.2×10^{-14}	4.7×10^{-15}
7	4	0.979845	0.979845	5.3×10^{-16}	4.1×10^{-14}	4.9×10^{-15}
4	0	0.024977	0.024977	1.1×10^{-15}	n/a	4.6×10^{-15}
4	1	0.209060	0.209060	4.3×10^{-15}	6.0×10^{-15}	5.2×10^{-3}
4	2	0.505902	0.505902	7.5×10^{-16}	1.1×10^{-15}	5.6×10^{-3}
4	3	0.800485	0.800485	1.9×10^{-17}	1.5×10^{-14}	5.7×10^{-3}
4	4	0.978669	0.978669	7.2×10^{-16}	2.7×10^{-14}	5.7×10^{-3}

$b = 7$ uses 30 qubits and marks the real-valued rare set exactly; $b = 4$ uses 27 qubits and misclassifies 9 futures. $P(\mathrm{rare})$ is computed directly from the simulated probabilities of all 2^{18} path-register states. The exact-amplification column gives $P(\mathrm{rare})$, real definition, under exact amplification of the fixed-point marked set. Accumulator leakage is the magnitude of $1 - P(\mathrm{accumulator} = 0)$. D_{TV} is the total-variation distance of the amplified rare-conditioned distribution from its real-valued target.

Amplification signal on superconducting hardware

Both jobs ran on IBM Kingston on 29 October 2026 (UTC), with four alternating oracle and control pairs of 5,000 shots each (20,000 shots per arm), on the same

physical qubits as the noisy simulation (Table 2). Both comparisons met the matching rule, with 228 routed two-qubit gates in each arm at $H = 2$ and 1,168 and 1,166 at $H = 3$. IBM's job records retain the submitted circuits and the raw shots of every PUB; the counts analysed here are identical to those records, and in each job the four oracle circuits are identical, as are the four controls. At $H = 2$ the oracle circuit drew a rare future with probability 0.6836 and the control with 0.4351, a signal of +0.2485 (95% interval +0.2391 to +0.2579; $z = 50.1$), against +0.408 predicted and +0.65 ideal. At $H = 3$ the probabilities were 0.7353 and 0.7583, a signal of −0.0230 (95% interval −0.0315 to −0.0144; $z = -5.3$), against +0.060 predicted and +0.52 ideal. By the stated test (confirmatory), the two-step difference reaches the threshold of $z \geq 5$ set before the runs and the three-step difference does not reach it.

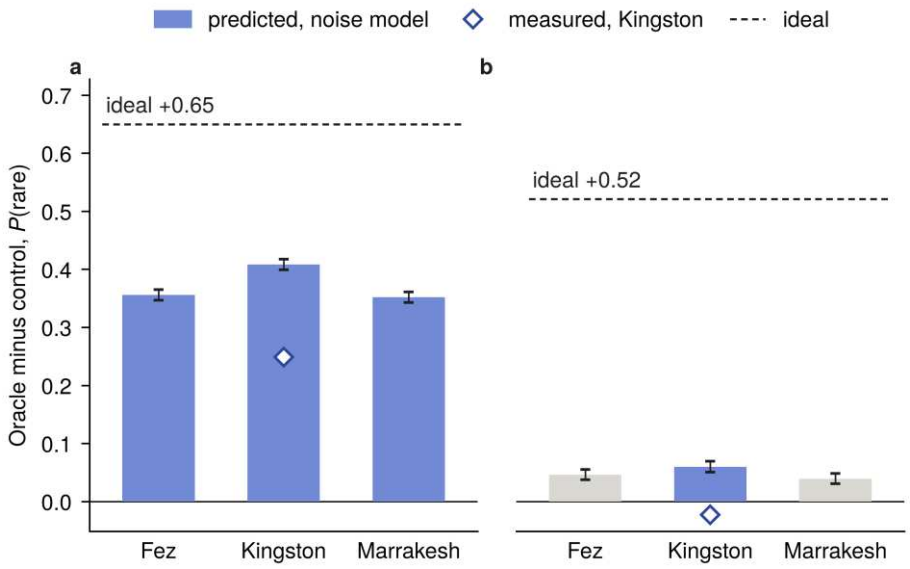


Figure 3

Figure 3. The two-step signal is detected on hardware, the three-step signal is not. Oracle-minus-control signal for one amplification iteration, predicted under noise models built from IBM calibration snapshots retrieved on 29 October 2026, and measured on Kingston. **a**, $H = 2$ (5 qubits). **b**, $H = 3$ (7 qubits). Bars show means over 3 simulations of 20,000 shots per arm; error bars show 95% binomial intervals

for a 20,000-shot hardware run. Grey bars fall below the $z \geq 5$ threshold or miss the 2% matching rule. The dashed line is the ideal signal, $\sin^2 3\theta - p_{\mathrm{rare}}$. Black diamonds are the signals measured on Kingston on 29 October, with 95% intervals (hidden by the symbol in **a**).

The following analyses are exploratory; they were not specified before the run. The four pairs gave signals from +0.2418 to +0.2564 at $H = 2$ and from −0.0268 to −0.0186 at $H = 3$, consistent within sampling error ($\chi^2_3 = 1.8$, $p = 0.61$, and $\chi^2_3 = 0.5$, $p = 0.92$). Across the 25 two-step futures, the change in each future's frequency from control to oracle tracked the change predicted by ideal amplification (Pearson $\rho = 0.98$; 0.81 over the 21 rare futures), and did so within each pair separately (ρ from 0.98 to 0.99; 0.74 to 0.84 over rare futures), so the two-step difference has the pattern of amplification, not only a higher rare-hit rate; across the 125 three-step futures the same correlation was $\rho = 0.04$.

At two steps the measured signal was 61% of its prediction, almost entirely because the control drew rare futures more often than predicted (0.435 against 0.300; oracle 0.684 against 0.708), consistent with stronger mixing than the noise model predicted, since the control's output, concentrated on the most likely future, is the more sensitive to mixing. Among the shots that returned a rare future, the spread over the 21 rare futures lay 0.141 from the target $P(x \mid \mathcal{R}_\Delta)$ in total variation for the oracle arm and 0.175 for the control, against 0.013 and 0.017 expected from sampling alone and 0.439 for a uniform spread. Noise therefore distorted the relative likelihoods in both arms by far more than sampling error, but the amplified arm was no further from the target than its control. In aggregate the oracle arm's rare-hit rate, 0.684, lies only 0.028 above the 0.656 of a uniform output, so the aggregate alone cannot separate

amplification from mixing; the oracle arm's spread over the rare futures, 0.141 from the target against 0.439 for a uniform spread, and its per-future correlation of 0.81 over the rare futures alone show that it kept the shape of the amplified distribution, which mixing would have erased.

At three steps the signal had the opposite sign to its prediction in all four pairs. Both arms drew rare futures far more often than $p_{\mathrm{rare}} = 0.086$, towards the 0.852 of a uniform output over all 128 indices (relaxation towards the all-zero index, a non-rare future, pulls the other way), and their spreads over rare futures (both 0.368 from the target, against 0.029 from sampling alone and 0.457 for a uniform spread) had lost most of the target's shape. The three-step outputs were therefore largely noise-dominated. Were both arms depolarised equally, the oracle arm would still lie above the control, so the consistent deficit of 0.0230 shows that the two compiled circuits responded differently to noise: the control departed further from its prediction (0.758 against 0.607) than the oracle (0.735 against 0.667). The submitted circuits, retrieved from IBM's job records and simulated without noise (400,000 shots each), reproduce the ideal values within sampling error (0.607 and 0.086 at $H = 3$, against 0.607 and 0.086; 0.772 and 0.122 at $H = 2$, against 0.773 and 0.123), confirming that the deficit does not arise from the submitted circuit logic; its physical mechanism was not identified. The noise models were optimistic at both horizons, more so for the deeper circuit. They used calibration data retrieved at least 6.9 hours before the jobs ran, Kingston was recalibrated 19 minutes before they ran, and they omit crosstalk, coherent errors, leakage and drift.

Table 2. Hardware run on IBM Kingston

H	r	Shots per arm	Routed 2q, oracle / control	$P(\text{rare})$, oracle	$P(\text{rare})$, control	Signal	95% interval	z	Predicted signal	Ideal signal	Retention
2	1	20000	228 / 228	0.6836	0.4351	+0.2485	[+0.2391, +0.2579]	50.1	+0.408	+0.650	0.38
3	1	20000	1168 / 1166	0.7353	0.7583	−0.0230	[−0.0315, −0.0144]	−5.3	+0.060	+0.521	−0.04

Four oracle and control pairs of 5,000 shots each per job, submitted alternately. z is the pooled two-proportion statistic of the stated test; the interval uses the unpooled standard error. Predicted signal is the mean of the Kingston noisy simulation (Table S4). Retention is the signal divided by the ideal signal. Jobs dbhgjlht9uhs73ca7e9g and dbhgk2hq260s73bv9100. Calibration in force at run time on the qubits used: $H = 2$: qubits 3, 24, 16, 22, 23, median CZ error 1.2×10^{-3}, readout error 1.2×10^{-2}, T_1 227 μs, T_2 118 μs. $H = 3$: qubits 75, 71, 93, 79, 74, 72, 73, median CZ error 1.2×10^{-3}, readout error 1.3×10^{-2}, T_1 248 μs, T_2 154 μs.

Standard suppression does not recover the three-step signal

To test whether standard error suppression could remove the three-step deficit, the same compiled circuits were rerun on the same physical qubits of Kingston on 30 October 2026 (Methods). The run used three conditions: native, dynamical decoupling, and dynamical decoupling with Pauli gate twirling. It had four rotated rounds in one batch, and the analysis was pre-registered before submission. Neither form of suppression raised the three-step signal (confirmatory, pre-registered: +0.013 native, −0.004 and +0.006 with suppression; changes $z = -2.8$ and -1.1, against a pre-registered threshold of $z \geq 3$ for an increase; Fig. 4, Table 3). Decoupling lowered it in all four rounds. Twirling moved both arms towards random output (rare-future frequencies 0.83 and 0.82, against 0.85 for a uniform output). The per-future correlation at three steps rose from 0.36 to 0.54 under decoupling, meeting the pre-registered criterion, but over the rare futures alone, a check that was not pre-registered, it was 0.10, against 0.17 native, so no amplification pattern among the rare futures was recovered. The two-step signal was detected in every condition ($z = 25$ to 40), and decoupling raised it from +0.135 to +0.195. The native runs did not repeat the 29 October values: the three-step signal was +0.013 against −0.023, and the two-step signal was about half as large (+0.135 against +0.249), although the reported two-qubit gate and readout errors of the qubits used were unchanged.

Taken together, the two campaigns bracket the practical regime for these reference circuits: the amplification signature was present at 228 routed two-qubit gates and absent at about 1,170. The two horizons differ in more than gate count, in qubit number, physical qubits, layout, depth and oracle, so these results bracket the regime for these implementations rather than establishing a universal gate-count threshold. The reproducible feature is the distinction between the two regimes, not

the magnitude of the two-step signal, which was +0.249 on 29 October and +0.135 on 30 October: at two steps the signal passed the detection threshold on both days under every condition, with each future's change tracking ideal amplification, while at three steps the signal was near zero or reversed on both days and under both forms of suppression.

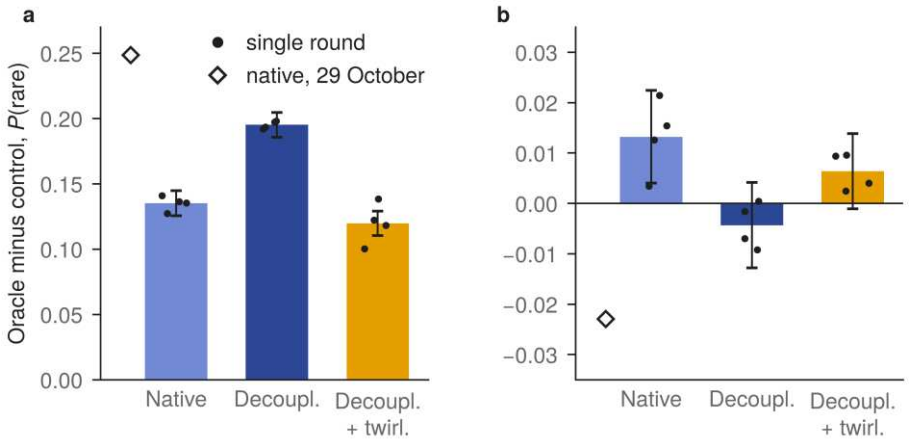


Figure 4

Figure 4. Noise suppression does not recover the three-step signal. Oracle-minus-control signal for one amplification iteration on IBM Kingston, 30 October 2026, under three conditions: native, dynamical decoupling (XpXm), and dynamical decoupling with Pauli twirling of the two-qubit gates. **a**, $H = 2$. **b**, $H = 3$. Bars are the signals pooled over four rounds of 5,000 shots per arm, with 95% intervals from the unpooled standard error; dots are the four single-round signals; the open diamond is the native signal measured on 29 October. The ideal signals are +0.650 at $H = 2$ and +0.521 at $H = 3$.

Table 3. Noise-suppression run on IBM Kingston, 30 October 2026

Condition	H	$P(\text{rare})$, oracle	$P(\text{rare})$, control	Signal	95% interval	z	Per-future ρ	Change against native (z)	Retention
Native	2	0.6343	0.4993	+0.1350	[+0.1253, +0.1446]	27.2	0.97	n/a	0.21
Decoupling	2	0.7004	0.5054	+0.1951	[+0.1857, +0.2044]	39.9	0.99	+0.0601 (+8.8)	0.30
Decoupling + twirling	2	0.6888	0.5691	+0.1197	[+0.1103, +0.1291]	24.8	0.99	−0.0153 (−2.2)	0.18
Native	3	0.6807	0.6675	+0.0132	[+0.0040, +0.0224]	2.8	0.36	n/a	0.03
Decoupling	3	0.7500	0.7543	−0.0044	[−0.0128, +0.0041]	−1.0	0.54	−0.0176 (−2.8)	−0.01
Decoupling + twirling	3	0.8271	0.8208	+0.0064	[−0.0011, +0.0138]	1.7	0.24	−0.0069 (−1.1)	0.01

The circuits and physical qubits of Table 2. Each condition comprises four oracle and control pairs of 5,000 shots, one in each of four rounds with the conditions rotated, giving 20,000 shots per arm. Decoupling: XpXm sequences scheduled as late as possible. Twirling: Pauli twirling of the two-qubit gates, 20 randomisations of 250 shots. z is the pooled two-proportion statistic; the interval and the change against native use unpooled standard errors. Retention is the signal divided by the ideal signal. Per-future ρ is the Pearson correlation between each future's measured and ideal change from control to oracle; over the rare futures alone it was 0.81, 0.86, 0.88 at $H = 2$ and 0.17, 0.10, 0.05 at $H = 3$ (native, decoupling, decoupling + twirling). Ideal signals are +0.650 at $H = 2$ and +0.521 at $H = 3$. Batch 7855afd6-8494-4677-b80a-698b9c7fc19c.

Discussion

What has been constructed. When the process is a known finite-state Markov chain, the rare set defined as in Guo et al. [9] can be marked by a circuit that is exactly verifiable at this scale. The circuit combines standard techniques: surprisal accumulated in the Fourier basis [25, 26], so each transition costs one diagonal gate per accumulator qubit; an offset preload, so the comparison reduces to the sign of one bit; and reliance on the prepared state's support, so unused encodings need no flag. What is new is the complete trajectory-level construction and its characterisation: the rarity of a future, defined by its surprisal against an entropy-rate threshold, becomes an executable oracle whose every component was checked against exact enumeration.

Precision and resources. Exact marking has a cost that does not follow the horizon smoothly. Proposition 1 guarantees classification from b of about $\log_2[(H/2+1)/g]$, and the gap g varies irregularly with the horizon (Table S1): the bound is 15 at $H = 5$ and 13 at $H = 6$. Below it, zero mismatches depend on how the rounding errors fall and must be checked for each b; at $H = 6$, $b = 7$ and $b = 8$ classify every future correctly while $b = 9$ misclassifies two (Table S2), which is a property of the rounding pattern, not of the circuit. The oracle dominates the cost at every horizon, with the preparation under 1% of the total at $H = 6$, and the general construction grows as Hm in conditioned phase rotations; at about 66,802 logical two-qubit gates for four iterations at $b = 7$ with this synthesis, the six-step circuit is far beyond present hardware.

What the hardware shows. Raw rare-hit rates cannot identify amplification on a noisy processor, because noise alone pushes the output towards the many rare

555 futures; the matched control is therefore essential, and its limits are stated in
556 Methods. Against that control, the two-step circuit showed the signature of amplitude
557 amplification on both days and under every condition: a signal that passed the
558 detection threshold set before the run, and a change in each future's frequency that
559 tracked ideal amplification with $\rho = 0.98$ across all 25 futures. At three steps, with
560 about 1,170 routed two-qubit gates, the signal was near zero or reversed on both
561 days, the per-future correlation collapsed, and the tested suppression techniques,
562 dynamical decoupling and Pauli gate twirling, did not recover it; twirling moved both
563 arms towards random output. The two circuits can also differ under noise for reasons
564 unrelated to amplification, as the reversed sign of 29 October shows, and a noiseless
565 replay of the submitted three-step circuits confirms that their intended ideal behaviour
566 is correct. Error mitigation by extrapolation has carried circuits of thousands of two-
567 qubit gates on IBM processors [34], and its measurement cost can be reduced with
568 learned surrogates [35], but it corrects estimated expectation values, such as the
569 probability of a rare future, not the individual samples that a sampler of rare futures
570 returns.

571 **Relation to the recent rare-event literature.** Guo et al. [9] and Gupta et al. [10]
572 establish asymptotic sampling advantages, the first from a coherent probability
573 threshold and the second from quantum stochastic models combined with
574 amplification. This paper's contribution lies elsewhere: an executable construction, its
575 finite-precision and resource characterisation, and a hardware test of the
576 amplification step with matched controls and pre-registered analysis.

577 **Limits and the path ahead.** The results have four limits. First, the transition matrix is
578 known, so the same information gives exact classical sampling of the rare-
579 conditioned distribution, and no quantum advantage is claimed. Second, the six-step

results are ideal simulations, beyond present hardware depth. Third, the hardware references take the rare set as a list, so they test amplification and the noise model, not the general oracle. Fourth, the calibration-derived noise models omit crosstalk, coherent errors, leakage and drift, and they predicted a detectable three-step signal that did not appear (the two-step signal was 61% of its prediction on 29 October and 33% on 30 October). Together these restrictions keep every result checkable while separating the correctness of the algorithm from what present hardware can run. The mechanism is aimed at models that can be run but whose probabilities are hard to calculate, such as large climate, epidemic and financial simulators whose observed paths aggregate many hidden ones. There the rule-based oracle demonstrated here cannot be built; it would be replaced by a rarity test built from the generator alone, as in Guo et al. [9]. That test needs the simulator itself, and its inverse, to run coherently as a quantum circuit, which is the dominant cost and remains open for simulators of this size.

Conclusions

A gate-level Qiskit circuit marks and amplifies the rare futures of a five-state Markov process over six steps, with every component verified against exact enumeration. At 7 fractional bits it marks the real-valued rare set exactly, and its 30-qubit simulation raises the probability of a rare future from 0.025 to 0.980 in four iterations. On IBM Kingston, a two-step reference circuit of 228 two-qubit gates, which marks rare futures from a precomputed list, drew rare futures more often than its matched control by +0.2485 ($z = 50.1$), passing the test set before the run; the difference was detected again on 30 October with and without noise suppression, while at three steps, with 1,168 two-qubit gates, the predicted signal was not seen on either day. A hardware signature consistent with amplitude amplification was therefore detected on

Kingston at 228 two-qubit gates but not at 1,168, and the complete six-step circuit, at tens of thousands of logical two-qubit gates, remains beyond current superconducting hardware.

List of abbreviations

CNOT: controlled-NOT gate; CZ: controlled-Z gate; PUB: primitive unified bloc; QPU: quantum processing unit.

Declarations

Availability of data and materials

The datasets generated and analysed during the current study, and the code, are available in the Zenodo repository, https://doi.org/10.5281/zenodo.23851083 (version 1.1.0), as two archives. The data archive holds the raw measurement counts for every circuit, IBM's records of the submitted circuits and results (jobs dbhgjlht9uhs73ca7e9g and dbhgk2hq260s73bv9100, and the twelve jobs of batch 7855afd6-8494-4677-b80a-698b9c7fc19c), the pre-registration record of the noise-suppression run, the calibration data in force when each job ran, the enumerated rare sets and all numerical results reported here. The code archive holds the Qiskit pipeline that builds, verifies and costs the circuits, the hardware submission and analysis scripts, and the scripts that generate every number, table and figure in this article (project name: qred; operating systems: macOS and Linux; programming language: Python 3.12 or later; requirements: Qiskit, Qiskit Aer, Qiskit IBM Runtime, NumPy, SciPy, Matplotlib; licence as stated on the Zenodo record; no restrictions on use by non-academics).

Competing interests

The author declares that he has no competing interests.

Funding

This study received no funding.

Authors' contributions

PMC conceived and designed the study, built and verified the circuits, ran the simulations and the hardware experiments, analysed the results and wrote the manuscript. The author read and approved the final manuscript.

Acknowledgements

We acknowledge the use of IBM Quantum services for this work. The hardware experiments ran on ibm_kingston, a QPU built on an IBM Heron r2 processor. The views expressed are those of the author, and do not reflect the official policy or position of IBM or the IBM Quantum team.

References

1. Zscheischler J, Westra S, van den Hurk BJJM, Seneviratne SI, Ward PJ, Pitman A, et al. Future climate risk from compound events. Nat Clim Change. 2018;8:469–477.

2. Haldane AG, May RM. Systemic risk in banking ecosystems. Nature. 2011;469:351–355.

3. Dobson I, Carreras BA, Lynch VE, Newman DE. Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization. Chaos. 2007;17:026103.

4. Asmussen S, Glynn PW. Stochastic simulation: algorithms and analysis. New York: Springer; 2007.

5. Bucklew JA. Introduction to rare event simulation. New York: Springer; 2004.

6. Cérou F, Guyader A. Adaptive multilevel splitting for rare event analysis. Stoch Anal Appl. 2007;25:417–443.

7. Del Moral P, Garnier J. Genealogical particle analysis of rare events. Ann Appl Probab. 2005;15:2496–2534.

8. Dupuis P, Wang H. Dynamic importance sampling for uniformly recurrent Markov chains. Ann Appl Probab. 2005;15:1–38.

9. Guo N, et al. Quantum enhanced rare event discovery and sampling. arXiv:2606.06316 [Preprint]. 2026. Available from: https://arxiv.org/abs/2606.06316.

10. Gupta A, et al. Spatiotemporal quantum advantages for rare-event sampling. Phys Rev Lett. 2026. https://doi.org/10.1103/7zh3-zvnd.

11. Grover LK. A fast quantum mechanical algorithm for database search. In: Proceedings of the 28th Annual ACM Symposium on Theory of Computing. New York: ACM; 1996. p. 212–219.

12. Brassard G, Høyer P, Mosca M, Tapp A. Quantum amplitude amplification and estimation. Contemp Math. 2002;305:53–74.

13. Montanaro A. Quantum speedup of Monte Carlo methods. Proc R Soc A. 2015;471:20150301.

14. Woerner S, Egger DJ. Quantum risk analysis. npj Quantum Inf. 2019;5:15.

15. Stamatopoulos N, Egger DJ, Sun Y, Zoufal C, Iten R, Shen N, et al. Option pricing using quantum computers. Quantum. 2020;4:291.

16. Egger DJ, García Gutiérrez R, Cahué Mestre J, Woerner S. Credit risk analysis using quantum computers. IEEE Trans Comput. 2021;70:2136–2145.

17. Figgatt C, Maslov D, Landsman KA, Linke NM, Debnath S, Monroe C. Complete 3-qubit Grover search on a programmable quantum computer. Nat Commun. 2017;8:1918.

18. Zhang K, Rao P, Yu K, Lim H, Korepin V. Implementation of efficient quantum search algorithms on NISQ computers. Quantum Inf Process. 2021;20:233.

19. Cover TM, Thomas JA. Elements of information theory. 2nd ed. Hoboken: Wiley; 2006.

20. Breiman L. The individual ergodic theorem of information theory. Ann Math Stat. 1957;28:809–811.

21. Dembo A, Zeitouni O. Large deviations techniques and applications. 2nd ed. Berlin: Springer; 2010.

22. Grover L, Rudolph T. Creating superpositions that correspond to efficiently integrable probability distributions. arXiv:quant-ph/0208112 [Preprint]. 2002. Available from: https://arxiv.org/abs/quant-ph/0208112.

23. Möttönen M, Vartiainen JJ, Bergholm V, Salomaa MM. Transformation of quantum states using uniformly controlled rotations. Quantum Inf Comput. 2005;5:467–473.

24. Möttönen M, Vartiainen JJ, Bergholm V, Salomaa MM. Quantum circuits for general multiqubit gates. Phys Rev Lett. 2004;93:130502.

25. Draper TG. Addition on a quantum computer. arXiv:quant-ph/0008033 [Preprint]. 2000. Available from: https://arxiv.org/abs/quant-ph/0008033.

26. Ruiz-Perez L, Garcia-Escartin JC. Quantum arithmetic with the quantum Fourier transform. Quantum Inf Process. 2017;16:152.

27. Shende VV, Bullock SS, Markov IL. Synthesis of quantum-logic circuits. IEEE Trans Comput Aided Des Integr Circuits Syst. 2006;25:1000–1010.

28. Vidal G. Efficient classical simulation of slightly entangled quantum computations. Phys Rev Lett. 2003;91:147902.

29. Javadi-Abhari A, Treinish M, Krsulich K, Wood CJ, Lishman J, Gacon J, et al. Quantum computing with Qiskit. arXiv:2405.08810 [Preprint]. 2024. Available from: https://arxiv.org/abs/2405.08810.

30. Viola L, Knill E, Lloyd S. Dynamical decoupling of open quantum systems. Phys Rev Lett. 1999;82:2417–2421.

31. Ezzell N, Pokharel B, Tewala L, Quiroz G, Lidar DA. Dynamical decoupling for superconducting qubits: a performance survey. Phys Rev Appl. 2023;20:064027.

32. Knill E. Quantum computing with realistically noisy devices. Nature. 2005;434:39–44.

33. Wallman JJ, Emerson J. Noise tailoring for scalable quantum computation via randomized compiling. Phys Rev A. 2016;94:052325.

34. Kim Y, Eddins A, Anand S, Wei KX, van den Berg E, Rosenblatt S, et al. Evidence for the utility of quantum computing before fault tolerance. Nature. 2023;618:500–505.

35. Liao WY, et al. Sample-efficient quantum error mitigation via classical learning surrogates. Commun Phys. 2026. https://doi.org/10.1038/s42005-026-02827-w.

Additional files

Additional file 1. Supplementary tables (PDF): Table S1, rare sets for horizons 2 to 6; Table S2, fixed-point precision at $H = 6$; Table S3, logical cost of the general circuit; Table S4, predicted oracle-minus-control signals under the calibration-derived noise models; Table S5, sensitivity of the six-step rare set to the rarity offset δ; Table S6, sensitivity of the predicted signal to the matched control's random phases.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- CattQuantumAmplificationEPJQuantumTechnologyadditionalfile1.pdf