

Article

The Energy-Advantage Frontier of Random Circuit Sampling Across the Simulability Transition

Hikaru Wakaura *  and Taiki Tanimae

QIRI (Quantum Integrated Research Institute Inc.), 1–16–3 Akasaka, Minato-ku, Tokyo 107-0061, Japan;
t.tanimae@deeptell.jp or t.tanimae@qiri.co.jp

* Correspondence: h.wakaura@deeptell.jp or h.wakaura@qiri.co.jp

Abstract

We introduce the *energy-advantage frontier* of random circuit sampling (RCS)—the locus in the noise-rate/size plane where the classical and quantum energy-to-solution are equal—and locate it relative to the noise-induced classical-simulability transition. Our firm, model-free result is that the classical sampling-energy floor collapses sharply across the transition, exhibiting a finite-size scaling collapse over exact systems $n \leq 12$ (with a collapse exponent of order unity). Built on it, the energy-advantage frontier is *distinct* from the computational (wall-clock) frontier: in the exascale-classical, few-femtojoule regime, the energy frontier is the more noise-robust, so an energy advantage can persist where a time advantage does not—a projection, not a deployment claim, whose direction reverses outside that regime. We give a conditional Landauer lower bound for sampling, $R_E \equiv E_{\text{cl}}/E_q \gtrsim 2^{(1.0 \pm 0.2)n}/n$ under standard RCS hardness, package R_E as a single figure of merit computed by a released open tool (exs), and map an energy-normalized cross-entropy witness within a cost model. A run on a 156-qubit IBM Heron processor tests the fidelity model (no energy is measured on hardware) and shows real noise exceeds the depolarizing idealization, so quoted frontier locations are upper bounds.

Keywords: random circuit sampling; quantum advantage; tensor networks; Landauer principle; energy complexity; cross-entropy benchmarking

MSC: 81P68; 68Q12; 81P45

1. Introduction

Quantum computational advantage has been demonstrated through random circuit sampling (RCS) on superconducting [1] and other platforms, with the linear cross-entropy benchmark (XEB) as the witness. The story is, however, incomplete in two respects. First, recent experiments and theory have shown that the noisy RCS device sits near a *phase transition* in classical simulability: cross-entropy benchmarking resolves a noise-induced transition between a globally correlated, hard-to-simulate phase and a factorized, efficiently simulable one [2,3], and a polynomial-time classical algorithm exists above a critical noise rate [4,5]. Second, “advantage” has been quantified almost exclusively in runtime, even as energy consumption becomes the dominant practical constraint on both quantum and classical computing [6,7].

These two threads have not been connected. On the energy side, the sharpest result—an exponential quantum energy-consumption advantage proven via Landauer’s principle and energy conservation—applies only to the query/decision setting (Simon’s



Academic Editor: Younghun Kwon

Received: 4 June 2026

Revised: 24 June 2026

Accepted: 26 June 2026

Published: 30 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

problem) [8]. The single study of *sampling* energy treats boson sampling with engineering metric–noise–resource accounting, and addresses neither a phase transition nor spoofing [9]. The intersection—energy complexity of qubit-based RCS, tied to the noise-induced simulability transition, and its consequences for benchmarking—is untouched.

Here we map that intersection. We treat the noise rate p , depth d , and size n as the axes of a phase diagram and ask three questions: (a) Where does the *energy*-advantage frontier ($E_{cl} = E_q$) lie, and does it coincide with the *computational* frontier ($T_{cl} = T_q$)? (b) Is an energy-normalized XEB harder to spoof than XEB itself? (c) Can the classical energy be lower-bounded?

Our *central, defensible claim* is conceptual and structural. That two different cost measures cross a fixed target at different noise rates is, by itself, a near-tautology; the non-trivial content is threefold. (i) *Both* frontiers are pinned to the *same* physical object—the simulability transition—because both the classical energy and the classical runtime are reductions of the same contraction cost, which collapses at p_c ; the energy question therefore inherits a sharp phase structure rather than a smooth trade-off. (ii) The *gap* between the frontiers is not a free parameter but is fixed by one dimensionless ratio of classical-to-quantum energy-per-operation to throughput-per-gate, so it can be estimated from hardware numbers rather than assumed. (iii) In the physically relevant regime (exascale-class classical hardware, few-femtojoule gates) that ratio places the energy frontier at *higher* noise than the computational one, i.e., the energy advantage is the more noise-robust—a statement with operational content for benchmarking, not a dimensional artifact, though its direction reverses outside that regime (Appendix B). This linkage of an energy-complexity statement to the simulability transition is, to our knowledge, new for RCS. We are explicit about what we do *not* originate: that noise destroys the correlations’ underlying hardness due to Refs. [4,5]; the Landauer/energy-conservation lower-bound technique is from Ref. [8]; XEB spoofing is from Refs. [10,11]. Our contribution is to (1) recast that mechanism as an *energy* order parameter and map its collapse exactly, (2) expose the energy/computational frontier decoupling and package it as a single reportable figure of merit $R_E \equiv E_{cl}/E_q$ with an open tool (*exs*) for any platform, and (3) propose the energy-charged witness—an assembly with a new quantitative axis, not a new mechanism. Formally, the contributions are Definition 1 (the energy-advantage ratio R_E and its frontier), Proposition 1 (a conditional exponential lower bound on R_E), and a finite-size scaling analysis of the underlying operator-entanglement transition; the remainder is a numerical mapping of the frontier and an open benchmark tool. The firm, model-free part of the story is the collapse of the exact classical energy floor across the transition (Section 3); the frontier separation built on it is a projection (Section 4). Everything else is either a corollary or a model-dependent quantification of it: the frontier *locations* are projections (Section 4), the spoofing threshold is a consequence within a cost model (Section 5), and the lower bound is a reduction to existing hardness (Section 6). Relative to prior work, we go beyond the decision-problem setting of Ref. [8] and the boson-sampling engineering accounting of Ref. [9] by treating qubit RCS and, crucially, by locating the energy statement *relative* to the noise-induced phase transition of Refs. [2,3]. We are careful to separate unconditional, conditional, and model-dependent claims, and not to present a floor inequality as a deployment energy saving.

2. Model and Energy Accounting

We use Google Sycamore-style circuits on an $L \times L$ lattice ($n = L^2$): each cycle applies random single-qubit gates from $\{\sqrt{X}, \sqrt{Y}, \sqrt{W}\}$ and two-qubit $\text{fSim}(\pi/2, \pi/6)$ gates on one of four bond patterns (ABCD) that together tile every nearest-neighbor bond. Noise is depolarizing with the two-qubit rate $p \equiv p_2$ and $p_1 = p_2/10$.

Classical energy. The best classical sampler is a tensor-network contraction. We charge it to the Landauer floor: each irreversible elementary operation dissipates $\geq k_B T \ln 2$ at the classical machine's temperature T [12,13]. The operation count is set by the contraction cost. For exact small systems we read the cost off the operator-space entanglement entropy (OSEE) S_{op} of the noisy density matrix across a balanced cut—the bond dimension of the minimal matrix-product operator is $\chi_{\text{eff}} = 2^{S_{\text{op}}}$ —and for a large n , from the single-amplitude contraction treewidth $\text{tw}(n, d)$ measured with COTENGR [14]. These are two distinct proxies for the *logarithm of the dominant classical cost*: OSEE is the bond dimension of the noisy-state matrix-product operator, and treewidth is the largest intermediate of a single-amplitude contraction. They are not identical, and we do not claim either is the *optimal* classical cost; what we use is that both *scale* as the cost and set the relevant bond/intermediate dimension. They agree on *scaling*—both $O(n)$ for deep two-dimensional circuits—but differ by an $O(1)$ factor at the overlap ($S_{\text{op}} \approx 7.4$ vs. $\text{tw} \approx 10$ – 15 at $n = 9$), and the treewidth itself varies by $\sim 30\%$ with the contraction optimizer (greedy vs. hyper-optimized); we use the hyper-optimized value and propagate the optimizer spread into the frontier uncertainty. Each proxy is the relevant exponent for its regime: the exact energy collapse (Section 3) uses OSEE, and the large- n frontier (Section 4) is the treewidth. This join is calibrated, not proven (Section 8); crucially, the *firm* result—the energy collapse—is read from OSEE directly on exact states and does not depend on it. For the frontier we use a per-sample cost $N n 2^{\text{tw}_{\text{eff}}}$, an upper estimate for *faithful* sampling; a frugal adversary targeting only a fixed XEB score is cheaper (Section 5). Both S_{op} and tw collapse above the noise threshold, locking E_{cl} to the simulability transition.

Remark 1 (Why the two proxies scale together). *The agreement is not coincidental: both quantities are controlled by the same minimal cut C of the circuit across a balanced bipartition. The contraction treewidth is the log-size of the largest intermediate when sweeping the amplitude network across C , which, for a geometrically-ordered contraction, equals the number of bonds crossing C up to a constant, $\text{tw} = \Theta(|\partial C|)$ (the contraction-width = min-cut relation [15]). The operator-space entanglement is the operator Schmidt entropy of ρ across the same cut, bounded by $S_{\text{op}} \leq 2 S_{\text{vN}}(|\psi\rangle) \leq 2|\partial C|$ and generically $\Theta(|\partial C|)$ before noise. Hence both are $\Theta(|\partial C|) = \Theta(\min(d, L) L)$ for an $L \times L$ depth- d circuit, differing only by the $O(1)$ factor separating an operator (ρ) from a state ($|\psi\rangle$) and a gate bond dimension from a qubit. Depolarizing noise shrinks the effective cut (correlations decay past a finite length), so both collapse together; this is the structural reason that the OSEE transition and the contraction-cost transition coincide, beyond the numerical calibration above.*

Quantum energy. We distinguish two figures. The *logical* floor charges only irreversible operations; since unitary gates are reversible, this is the readout, $E_q^{\text{logical}} = k_B T \ln 2 (n + n_{\text{reset}}) = O(n)$. The *practical* energy $E_q^{\text{prac}} = N_{\text{gates}} e_{\text{gate}}$ uses a transparent per-gate energy $e_{\text{gate}} \in [1 \text{ aJ}, 1 \text{ pJ}]$.

What “*advantage*” means here. Our primary comparison is deliberately a *two-sided Landauer floor*: the classical lower bound (Section 6) against the quantum logical floor. This isolates the computational thermodynamics from engineering overhead and is the only level at which a rigorous one-sided statement exists. It is *not* a claim about deployed machines: real CMOS dissipates $\sim 10^{-13}$ – 10^{-11} J per FLOP—some 8–11 orders above $k_B T \ln 2$ —while a superconducting quantum processor carries a cryogenic overhead of order 10^3 – 10^4 W of wall power per watt of useful cooling. Both inflations are large and act in opposite directions, so the *practical* energy verdict is set by the dimensionless ratio $\rho = (\text{classical J/op})/(\text{quantum J/op including overhead})$. We do not build ρ into the floor-level frontiers; instead, our sensitivity analysis (Appendix B) sweeps the per-operation energies and bit-per-op factors over their full plausible ranges, and Section 8 gives a con-

create ρ estimate anchored to a measured classical run. All “frontier” statements below are floor-level unless explicitly tagged as “practical”.

We can now state the central objects precisely.

Definition 1 (Energy- and time-advantage ratios and frontiers). *For a circuit family on n qubits at depth d and a two-qubit noise rate p , let $E_{\text{cl}}(n, d, p)$ and $T_{\text{cl}}(n, d, p)$ be the classical energy and wall-clock cost of the optimal sampler, and E_q, T_q are the corresponding quantum costs. Define the energy-advantage ratio and time-advantage ratio:*

$$R_E(n, d, p) \equiv \frac{E_{\text{cl}}}{E_q}, \quad R_T(n, d, p) \equiv \frac{T_{\text{cl}}}{T_q}. \quad (1)$$

The energy-advantage frontier is the locus $\{(n, d, p) : R_E = 1\}$, and the computational frontier is $\{(n, d, p) : R_T = 1\}$; a quantum energy (resp. time) advantage holds where $R_E > 1$ (resp. $R_T > 1$). Both ratios are reductions in the same contraction cost, through the Landauer floor and a finite throughput respectively.

3. Classical Energy Collapses Across the Transition

This section is the firm, model-free anchor of this paper: it uses exact density matrices only ($n \leq 12$), no extrapolation, and no cost model beyond the identification $\chi_{\text{eff}} = 2^{S_{\text{op}}}$. The energy-collapse data (Figure 1) are averaged over three random circuit instances, and the scaling analysis (Figure 2) are averaged over five; quoted uncertainties are standard deviations over instances. Figure 1 shows that the operator-space entanglement saturates with depth at $p = 0$ and then collapses sharply as p crosses a depth-dependent threshold. Fitting $S_{\text{op}}(p)$ to a logistic form gives a collapse midpoint $p_c \approx c/d$ and a width that narrows monotonically with depth— $w = 0.054, 0.021, 0.012$ at $d = 4, 8, 12$ for $n = 9$, i.e., the transition sharpens as $\sim 1/d$. Fitting the midpoint across a grid of twelve exact cells ($n = 8, 9, 10, 12$ at $d = 4, 8, 12$) pins the per-cycle error budget at the threshold to a near-constant $c \equiv p_c d = 1.072 \pm 0.069$, so the $p_c \approx c/d$ law is calibrated over four sizes and three depths rather than extrapolated from one. We stress that this operator-entanglement transition is, *a priori*, a distinct object from the XEB transition of Ref. [2] and the efficient-sampling threshold of Ref. [4]: the three need not share a critical rate. That all three sit at a comparable per-cycle error budget $O(1)$ is consistent with a common origin but is not a proof of equality, which we do not claim; what we establish directly is the OSEE transition, and it is the OSEE that sets the classical (MPO) energy.

Finite-size scaling collapse. The sharpening also holds in *system size*: measuring the order parameter $S_{\text{op}}/S_{\text{sat}}$ on four exact sizes $n = 8, 9, 10, 12$ (lattices $2 \times 4, 3 \times 3, 2 \times 5, 3 \times 4$; five circuit instances each; $d = 8$), the width narrows from $w = 0.027$ at $n = 8$ to 0.017 at $n = 12$, and the curves collapse onto a single scaling function $S_{\text{op}}/S_{\text{sat}} = f[(p - p_c(n))n^{1/\nu}]$ with an order-unity exponent and small collapse residual (Figure 2). The data collapse—over admittedly a narrow size range ($n = 8$ – 12 , all that exact density matrices allow)—is, together with the depth sharpening, our cleanest *model-free* evidence for a genuine transition rather than a crossover. We are explicit that ν is *not* a claimed universal critical exponent, and we now quantify how loosely it is constrained. The collapse optimum lies near $\nu \approx 0.85$ – 0.95 , but the residual landscape is shallow: the collapse residual stays within twice its minimum over the entire range $\nu \in [0.4, 2.0]$, and a leave-one-size-out jackknife gives $\nu = 1.3 \pm 0.4$. The exponent is therefore pinned only to order unity over this size range; we report it as a goodness-of-collapse summary, not a measured universality class. Equally, the collapse demonstrates self-similarity over this window but does not establish a thermodynamic-limit transition; the midpoint is nearly size-independent here ($p_c = 0.146, 0.135, 0.139, 0.124$), so we attempt no $n \rightarrow \infty$ extrapolation. The mechanism—noise erasing the correlations that make

RCS hard—is itself known [4,5]; what is new here is its sharp, model-free rendering as an *energy* collapse with quantified finite-size sharpening. The implied classical sampling-energy floor falls by $\sim$ six orders of magnitude across the transition (Figure 1b), e.g., from 5×10^{-6} J to 2×10^{-12} J for $n = 12$, $d = 12$. The low-noise plateau grows steeply with size ($\chi_{\text{eff}} \approx 150$ at $n = 9$ vs. 1400 at $n = 12$); the χ^3 growth that extrapolated (Section 4) opens the energy advantage at scale.




Figure 1. Noise collapses the classical sampling energy. (a) Operator-space entanglement entropy of the exact noisy density matrix saturates with depth and collapses across a depth-dependent noise threshold. (b) Landauer-floor classical energy E_{cl} vs. noise, with quantum-energy bands; E_{cl} falls by orders of magnitude through the transition. The absolute scale is the thermodynamic floor ($k_B T \ln 2$ per operation) and lies $\sim 10^{10}$ below the energy a real machine of this size would spend (Section 8); only the *collapse*, not the absolute joules, is physical here.




Figure 2. Finite-size scaling collapse of the exact OSEE transition. (a) Order parameter $S_{\text{op}}/S_{\text{sat}}$ vs. noise for four exact sizes ($n = 8, 9, 10, 12$, $d = 8$, five instances each); the transition sharpens with size. (b) The curves collapse onto $f[(p - p_c(n)) n^{1/\nu}]$ with an order-unity exponent ($\nu \approx 0.85\text{--}0.95$; jackknife 1.3 ± 0.4 ; see text— ν is a collapse-quality summary, not a measured critical exponent).

4. Energy Frontier vs. Computational Frontier

The same contraction cost converts to energy via the Landauer floor and to wall-clock via a classical throughput R_{cl} ; the quantum side converts via e_{gate} and gate time t_{gate} . We summarize the comparison by two dimensionless figures of merit: the *energy-to-solution ratio* $R_E \equiv E_{\text{cl}}/E_q$ and the *time-to-solution ratio* $R_T \equiv T_{\text{cl}}/T_q$, with a quantum advantage when the corresponding ratio exceeds unity. R_E is the reportable quantity we advocate alongside XEB; in conditional-floor accounting (Section 6), it takes the platform-agnostic per-sample form $R_E \sim 2^{\text{tw}_{\text{eff}}}$, and it is computed for an arbitrary device specification by the released exs tool (see Code and tool availability, Section 8). The two frontiers ($R_E = 1$ and $R_T = 1$) therefore need not coincide. We stress that what follows is a *projection*, not a prediction: the noiseless cost is anchored in the measured treewidth (which for

deep circuits saturates at $tw \rightarrow n$, Section 6), but the noise-dependence enters through a calibrated collapse model, $p_c = c/d$ with $c = 1.0$ fit to exact systems of only $n \in \{9, 12\}$ and $d \in \{4, 8, 12\}$. Absolute frontier locations are therefore uncertain at the level of this calibration; from the sensitivity analysis (Appendix B), a $\pm 50\%$ uncertainty in c moves all thresholds proportionally (e.g., p_E^* over 0.033–0.13 at $n = 49, d = 16$), while their *ordering* and the existence of an energy-only window are preserved. The structural result, robust to this uncertainty, is that the two frontiers do not coincide.

Within the projection (Figure 3, $d = 16$): the energy frontier appears already at $n = 25$ ($p_E^* \approx 0.060$, rising to ≈ 0.072 at $n = 49$), whereas the computational frontier appears only at $n = 49$ ($p_C^* \approx 0.041$). Propagating the *measured* calibration uncertainty $c = 1.072 \pm 0.069$ (Section 3) through the frontier solver gives, at $n = 49$, $p_E^* = 0.072 \pm 0.005$, $p_C^* = 0.041 \pm 0.003$, and a gap $\Delta p = 0.030 \pm 0.002$ —*resolved*, and clearly distinguishable from zero, under the data uncertainty. The gap only becomes ambiguous ($\Delta p \sim 0.015$ – 0.061) under the deliberate factor-of-two stress of p_c in our sensitivity sweep (Appendix B), which is a worst-case knob range, not the calibration error. Robust at both levels are the *sign* ($p_E^* > p_C^*$) and the *trend* (the energy frontier appears two sizes earlier, at $n = 25$ vs. $n = 49$). Moreover $n = 49$ —the only size carrying both frontiers—is the *least* validated point: it is beyond exact density matrices and beyond converged TN sampling (Appendix C), and its absolute p_c is further an upper bound under realistic noise (Section 7). The quantitative ordering therefore leans on the size *trend* ($n = 25 \rightarrow 36 \rightarrow 49$), not on the $n = 49$ gap; establishing it firmly needs exact systems beyond $n = 12$ or converged large- n sampling, left to future work. The qualitative statement—an energy-only regime exists where the device wins on (floor) energy but not on time—is the RCS analogue of the energy-before-computation ordering reported for boson sampling [9], here tied to the noise transition, and already holds unambiguously at the exactly-simulable sizes $n \leq 12$, where the computational frontier is absent entirely while the energy frontier is finite.




Figure 3. Projected energy vs. computational advantage frontier (floor accounting, $d = 16$). Red (●): the energy frontier p_E^* , the locus $R_E = 1$ above which classical sampling is cheaper in energy; it emerges at $n = 25$. Blue (■): the computational frontier p_C^* , with the locus $R_T = 1$ (wall-clock), which emerges only at $n = 49$. Shaded red/blue bands: the $\pm 1\sigma$ spread from the *measured* calibration $c = 1.072 \pm 0.069$ (Section 3), propagated through the frontier solver. Green region between the curves: the energy-only window, where the device wins on (floor) energy but not on time. The $n = 49$ point is the only size carrying both frontiers and is the least validated (beyond exact and converged-TN reach); the quantitative claim rests on the size trend, not on this single point.

5. Energy-Charged XEB and Spoofing

Linear XEB is classically spoofable [10,11,16]. We ask whether charging a spoofer its own classical energy restores the witness. Define the energy-normalized witness $W = \text{XEB}/\text{joule}$ and compare the honest device against the cheaper of the two spoofers: a faithful noisy-distribution sampler, and a fidelity- F sparse spoofer that hits $\text{XEB} = F$ at cost $\sim F \times (\text{full simulation})$ [16]. The benchmark is spoof-resistant where the spoofer's energy exceeds E_q . Figure 4 shows that for $n \leq 16$, the full classical simulation is cheaper than the device even at unit fidelity, so energy charging cannot help; but for $n \geq 25$, a finite spoof-resistant window $p_2 < p_{\text{spoofer}} \approx 0.03\text{--}0.04$ opens, with $p_{\text{spoofer}} \approx p_C^*$. We are explicit that this coincidence is largely a *consequence of the cost model*: both the spoofer and the honest classical reference are charged the same contraction cost, so the spoofer's energy crosses E_q near where classical sampling becomes cheap by construction. The result should therefore be read as a proposal—an energy-normalized witness, and the regime in which it can help *within a given cost model*—rather than as a proof of spoof-resistance. In particular, the strongest known XEB spoofers need not faithfully sample and may achieve a target score below the contraction-cost we assume [10,11]. To probe this we charge an adversary the lowest-energy classical method achieving the same XEB, modelled as reaching the target at a *reduced* effective treewidth $\beta \text{tw}_{\text{eff}}$ ($\beta < 1$ captures exploiting the XEB structure rather than reconstructing the distribution). The resistant threshold is only logarithmically sensitive to β : at $n = 49$, $d = 16$, it moves from $p_{\text{spoofer}} = 0.072$ ($\beta = 1$) to 0.064 ($\beta = 0.7$) and 0.053 ($\beta = 0.5$)—a $\sim 25\%$ shrink even when the adversary needs only half the treewidth, because $E_{\text{spoofer}} \propto 2^{\beta \text{tw}_{\text{eff}}}$ and tw_{eff} collapses steeply across the transition. The window therefore persists against sub-full-simulation adversaries, but *closes* for a spoofer whose cost does not track the treewidth at all ($\beta \rightarrow 0$, e.g., a trivial frugal spoof), ruling out that class is beyond our cost model and is the central open question for the witness.




Figure 4. Spoofing resistance of an energy-charged XEB. (a) XEB-per-joule of the honest device vs. the best spoofer ($n = 49, d = 16$); the device wins only in the shaded low-noise window. (b) The spoof-resistance threshold p_{spoofer} opens at $n \geq 25$ and saturates near the computational frontier.

6. A Conditional Landauer Lower Bound for Sampling

The energy floor here is, under a standard hardness assumption, more than a model; we state it as a conditional proposition with the assumption made explicit. It is a reduction inheriting the hardness conjecture and its hidden constants, not a self-contained theorem. Any classical algorithm sampling $p_U^{(p)}$ to total-variation distance δ dissipates to $E_{\text{cl}} \geq k_B T \ln 2 \cdot R_{\text{min}}$, with R_{min} as the minimal irreversible-operation count.

Assumption 1 (RCS contraction hardness). *No classical algorithm δ samples the output distribution p_U of the Sycamore-type ensemble in the time sub-exponential in the contraction treewidth*

$\text{tw}(n, d)$, for δ below a constant $\delta_0 > 0$. This strengthens the average-case hardness underlying RCS supremacy claims [17,18].

Under Assumption 1, $R_{\min} \geq 2^{\Omega(\text{tw})}$ (the $\Omega(\cdot)$ constant and δ -dependence are not controlled here). We make the exponent explicit: fitting the measured treewidth to $\text{tw} \approx c \min(d, L) L$ gives, for deep circuits $d \gtrsim L$, saturation at $\text{tw} \rightarrow cn$ with $c = 1.07 \pm 0.16$ (Appendix A), the maximum for a single amplitude—consistent with the known L^2 scaling of 2D supremacy-circuit contraction [15].

Proposition 1 (Conditional two-sided energy separation). *Under Assumption 1, for deep circuits $d \gtrsim \sqrt{n}$ and noise $p < p_c$, the classical sampling-energy floor satisfies $E_{\text{cl}} \geq k_B T \ln 2 \cdot 2^{(1.0 \pm 0.2)n}$. Paired with the quantum logical floor $E_q^{\text{logical}} = O(n)$ (unitary gates are reversible and cost zero; only readout and resets are irreversible), the energy-advantage ratio of Definition 1 obeys*

$$R_E = \frac{E_{\text{cl}}}{E_q^{\text{logical}}} \geq \frac{2^{(1.0 \pm 0.2)n}}{n + n_{\text{reset}}}. \quad (2)$$

This is the sampling analogue of the query-setting result of Ref. [8].

Remark 2. *Three caveats are essential. (i) Both sides are Landauer floors of idealized, error-free machines, and the asymmetry is delicate: $E_q^{\text{logical}} = O(n)$ holds only for a hypothetical noiseless physical device, because fault tolerance carries its own Landauer cost—syndrome extraction and the reset of measured ancillas are irreversible—so a fault-tolerant realization at the sizes where the bound bites would dissipate far more than $O(n) k_B T \ln 2$ is plausibly enough to erode the separation even at the floor level. (ii) The bound is stated at the δ implicit in the treewidth (high-fidelity reproduction); a frontier at a fixed, larger δ would be cheaper classically and is not mapped here. (iii) The exponent c is fitted to single-amplitude treewidths on $L \leq 7$ lattices; for $n \leq 9$, we checked it against the independently-measured OSEE ($S_{\text{op}} \approx 7$ vs. $\text{tw} = 8$ at $n = 9$), but OSEE (an MPO cost) and treewidth (an amplitude-contraction cost) are distinct proxies whose quantitative join is not validated across the full size range.*

Crucially, above p_c , the noisy distribution is efficiently samplable [4], so $R_{\min} = \text{poly}(n)$ unconditionally, and the bound collapses: the energy lower bound itself is an order parameter for the transition—exponential (conditional) below p_c , polynomial (unconditional) above it. The exponential part is necessarily conditional, as an unconditional version would separate complexity classes; the unconditional content is the $O(n)$ readout floor and the above-threshold collapse.

7. Hardware Run on a Superconducting Processor

We executed our circuits on IBM's 156-qubit Heron processor `ibm_marrakesh`, sampling 3×3 RCS at depths $d = 2$ –10 (4096 shots each); we are explicit that this validates the *fidelity model* and the noise-axis placement, not the energy accounting (no energy is measured on hardware). Two points follow (Figure 5). First, in the anti-concentrated regime ($d \geq 6$), the measured XEB sits ~ 20 –40% below the digital-depolarizing model evaluated at the device-median rates ($F_{\text{XEB}} = 0.40$ vs. model 0.49 at $d = 8$): real hardware carries coherent, leakage, and correlated errors beyond pure depolarizing, so the simple digital model *over-estimates* fidelity. This has a direct consequence for our phase diagram: since real noise degrades coherence faster than depolarizing at an equal p_2 , the true transition occurs at a *lower* effective noise rate, so the depolarizing-computed p_c and all frontier locations (Sections 3–5) should be read as *upper bounds* on the real-hardware values—the qualitative ordering of frontiers is unaffected, but their absolute positions shift down under

realistic noise by an amount our depolarizing model cannot capture. (At $d \leq 4$ the circuits are not yet anti-concentrated and $F_{\text{XEB}} > 1$, as expected.) Second, the device's two-qubit error ($p_2 = 0.0027$, a Heron-generation figure well below the 0.78% of earlier Eagle parts) lies far below the projected frontier thresholds, so *on the noise axis*, present hardware is deep inside the advantage regime; the binding constraint is connectivity and decomposition—a fSim layer has no native heavy-hex [19], two-qubit gate, so a 9-qubit depth-10 circuit inflates to 162 cz + 970 single-qubit gates after routing (Figure 5b). The near-term obstacle to an energy advantage is thus architectural overhead rather than raw gate fidelity.




Figure 5. Live run on *ibm_marrakesh* (156-qubit Heron, $p_2 = 0.0027$). (a) Measured XEB (4096 shots, one circuit instance per depth) vs. the digital-depolarizing model; in the anti-concentrated regime ($d \geq 6$), the device falls below the model (real noise exceeds depolarizing). (b) Heavy-hex transpilation overhead.

8. Robustness and Discussion

A sensitivity sweep over eight modelling and hardware assumptions (Appendix B) confirms the qualitative ordering $p_C^* \lesssim p_{\text{spoofer}} < p_E^*$ and the energy-only window survive throughout. The thresholds respond to physically distinct knobs— e_{gate} moves only the energy and spoofing thresholds, R_{cl} moves only the computational one, and the shot budget moves none—while the dominant uncertainty in the absolute position is the calibration of p_C , which scales all thresholds and is anchored here to exact small-system data.

Realistic-energy anchor. To connect the Landauer floor to deployed machines we anchor to a concrete, measured classical run: sampling 10^6 bitstrings from the 53-qubit, 20-cycle Sycamore circuit took ≈ 15 h on 512 GPUs [7,16]—of order 10^{10} J of wall energy ($\sim 10^3$ kWh). The underlying contraction is $\sim 10^{19}$ FLOPs (an exaFLOPS machine would need tens of seconds), whose $k_B T \ln 2$ floor is only ~ 1 J; the $\sim 10^{10}$ gap is the $\sim 10^4 k_B T$ real cost of a 64-bit FLOP. A superconducting processor of this class draws ~ 26 kW (dilution refrigerator plus control and chilled-water cooling; this is the value reported for the same Sycamore apparatus [1] and for IBM Quantum System Two), so, over a minute-long job, it dissipates $\sim 10^6$ J, so the realistic ratio favors the quantum device by $\rho \sim 10^3$ – 10^4 on energy, even where the time advantage is contested. These remain order-of-magnitude figures but are now anchored to a measured classical run, rather than to the abstract floor. We state plainly the limit of this argument: the *sign* of the practical energy advantage is not something we *prove*—it rests on the measured classical wall energy ($\sim 10^{10}$ J) on one side and an *estimated* quantum draw (~ 25 kW) on the other, with the latter unmeasured here. What we prove is the floor-level structure (Sections 3 and 6); what we argue, from the measured classical run, is that realistic ρ points the same way. A definitive sign would require an end-to-end energy audit of a quantum sampler at this scale, which does not yet exist.

Practical energy frontier. To make this quantitative we recompute the frontier with *both* sides charged at engineering rates: classical energy $E_{\text{cl}}^{\text{prac}} = (\text{contraction FLOPs}) \times 10^{-11}$ J (a GPU figure), and quantum energy $E_q^{\text{prac}} = N_{\text{gates}} e_{\text{gate}} N + P_{\text{cryo}} t_{\text{wall}}$ with $P_{\text{cryo}} = 26$ kW

(cryogenics, control, and chilled-water cooling) over the wall time t_{wall} of the shot budget. The cryogenic term dominates ($E_q^{\text{prac}} \sim 10^5\text{--}10^6$ J at $N = 10^6$), yet the energy-advantage frontier *survives*: at $d = 16$, it sits at $p_E^{\text{prac}} \approx 0.050$ ($n = 36$) and 0.058 ($n = 49$)—close to the floor-level projection ($p_E^* \approx 0.069\text{--}0.072$) and is absent for $n \leq 25$, where the device is simply too small. Thus the cryogenic/control overhead shifts the frontier modestly and pushes its onset to slightly larger sizes but does not erase the energy-only structure; error-correction overhead, omitted here, would push it further and is the dominant open caveat at fault-tolerant scale.

This realistic anchor and the floor-level “energy-only window” of Section 4 are *not* in tension once read correctly, and the distinction is important. The window is a statement about the *relative location of two frontiers* at a fixed scale: as noise increases at a fixed (n, d) , the time advantage is lost first (at p_C^*) and the energy advantage is lost last (at $p_E^* > p_C^*$), so a band of high noise exists where only the energy advantage survives. It is *not* the claim that quantum saves no time at low noise—at low noise and large scale, both advantages hold, exactly as the realistic anchor indicates. In short, the realistic numbers say the energy advantage is real at scale; the floor analysis adds that, *in this parameter regime*, it is more robust to noise than the time advantage—an ordering that, as noted above, is set by the energy/throughput ratio and is not universal. The deployment claim and the structural claim are different statements, and we make only the latter.

Limitations. Our conclusions are bounded by five assumptions, stated here so they are not mistaken for results. (1) *Noise model*: purely depolarizing ($p_1 = p_2/10$); coherent errors, leakage, crosstalk, and non-Markovian noise can shift both the transition location and the OSEE collapse, and the claim that the energy frontier tracks the transition is established only for depolarizing noise. (2) *Accuracy*: the cost proxies target high-fidelity reproduction; the δ -dependence of all frontiers is unmapped. (3) *Cost proxies*: OSEE (small n) and treewidth (large n) are distinct and only loosely cross-checked. A validated Liouville-MPS sampler (Appendix C) extends the OSEE transition to $n = 16, 20$ and confirms it persists, but is not χ -converged there, so it does not yet pin the large- n critical rate or the OSEE-treewidth join. (4) *Idealization*: the two-sided floor compares error-free reversible machines and omits quantum error-correction overhead, which would dominate any fault-tolerant realization. (5) *Extrapolation*: the quantitative frontier ordering rests on a single size ($n = 49$) within a calibrated model.

In sum, the firm result is model-free: the exact classical energy floor of RCS collapses across the noise-induced simulability transition (Section 3). Built on it, within clearly-bounded floor accounting, the energy and computational advantage frontiers are distinct, and in the exascale-class/few-fJ regime, the energy frontier is the more noise-robust; an energy-normalized XEB witness is spoof-resistant only in a model-dependent low-noise window; a conditional Landauer lower bound for sampling inherits the transition. These would be *confirmed* by an end-to-end approximate-TN sampler reproducing the energy collapse at fixed δ for $n \gg 12$ and a fault-tolerant quantum energy budget and *falsified* if, under such accounting, the two frontiers coincided or the energy floor failed to collapse at p_C . Pursuing these—together with non-depolarizing noise and an explicit charging of the strongest known XEB spoofers—is the central open task.

Bottom line. One result is firm and new: the classical *energy* of sampling RCS, computed exactly for $n \leq 12$, collapses sharply across the noise-induced simulability transition, with finite-size sharpening in both depth and size. Everything built on it—the frontier decoupling, the energy-charged witness, and the conditional lower bound—is a framework and a set of projections whose reliability we have bounded explicitly. We have deliberately applied the tightest scrutiny to the new energy claims (whose sign under realistic accounting we argue but cannot prove) and acknowledged that the underlying *mechanism* is not ours.

We present the work in that spirit: a quantitative energy axis on a known transition, honestly delimited, rather than a closed claim of quantum energy advantage.

Author Contributions: Conceptualization, H.W. and T.T.; methodology, software, and formal analysis, H.W.; validation and investigation, H.W. and T.T.; writing—original draft, H.W.; writing—review and editing, H.W. and T.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original data presented in this study are openly available at [GitHub]: [https://github.com/deeptell-inc/noisy_rcs_phase.git] (accessed on 25 June 2026).

Acknowledgments: We thank the developers of QISKIT, QISKIT-AER, QUIMB, and COTENGRA.

Conflicts of Interest: The authors Hikaru Wakaura and Taiki Tanimae are chief officers of Quantum Integrated Research Institute Inc. The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as potential conflicts of interest.

Appendix A. Treewidth Scaling and the Explicit Constant

Figure A1 shows the single-amplitude contraction treewidth $\text{tw}(n, d)$ measured with COTENGRA for $L \times L$ lattices, $L = 3-7$, and depths $d = 8-20$. For deep circuits ($d \gtrsim L$) the treewidth saturates at the spatial cross-section, $\text{tw} \rightarrow cn$ with $c = 1.07 \pm 0.16$ (measured $\text{tw}/n \in [0.86, 1.02]$ for $d \geq 16$); for shallow circuits, $\text{tw} \approx d\sqrt{n}$. The saturated value $\sim n$ is the maximum possible for a single amplitude, fixing the exponent of the conditional lower bound of Section 6.




Figure A1. (a) Treewidth vs. n saturates at the line $\text{tw} = n$ for deep circuits. (b) Lattice scaling collapse against $\min(d, L) L$.

Appendix B. Sensitivity Analysis

Figure A2 sweeps eight modelling and hardware assumptions about the baseline $(n, d) = (49, 16)$. The ordering $p_C^* \lesssim p_{\text{spoof}} < p_E^*$ is robust throughout. Each threshold responds to physically distinct knobs: the per-gate energy e_{gate} moves p_E^* and p_{spoof} but not p_C^* ; the classical throughput R_{cl} moves only p_C^* ; the shot budget moves none. The dominant driver of the absolute position is the calibration of p_c , anchored here to exact small-system data, and p_{spoof} is most sensitive to the fidelity–cost exponent of the sparse spoofer.




Figure A2. Range of each threshold (p_E^* , p_C^* , p_{spoof}) induced by sweeping each assumption; the dashed line is the baseline.

Appendix C. Tensor-Network Sampler Beyond Exact Systems

To probe the transition past the $n \leq 12$ reach of dense density matrices, we evolve ρ as a vectorized state $|\rho\rangle\rangle$ in a matrix-product form with physical dimension four (per-site index 2 ket + bra); unitaries and depolarizing channels act as Liouville superoperators, with the bond dimension truncated to χ . The operator-space entanglement is read from the canonicalized Schmidt spectrum, normalized by the purity. The sampler is validated against exact dense evolution to $<10^{-2}$ in OSEE at $n = 9$ for all noise rates (large χ , no truncation).

Applied to $n = 16$ (4×4) and $n = 20$ (4×5), the noise-induced transition *persists* with a sharp, narrow collapse (fitted widths $w \approx 0.023$ – 0.026 , comparable to the exact $n \leq 12$ values); see Figure A3. We deliberately do *not* quote refined large- n critical rates: the operator entanglement in the transition region does not converge in the bond dimension. At $n = 16$, $p_2 = 0.08$ the OSEE rises monotonically, $3.96 \rightarrow 4.36 \rightarrow 5.04$, as $\chi = 256 \rightarrow 512 \rightarrow 1024$ —the increments *grow* rather than shrink, and the bond saturates at every χ —so, even $\chi = 1024$ (already ~ 13 min per point) is far from the saturated plateau ($S_{op} \sim 8$). Because under-resolved OSEE biases the apparent p_c low, the $\chi = 256$ midpoints are lower bounds on the truncation error, not physical critical rates. A converged large- n scaling analysis would require χ of order thousands—exponentially costly and beyond our present budget. What the sampler does establish, robustly, is that the transition and its sharpening survive to $n = 20$, and that the machinery to reach advantage-regime sizes is in hand; pinning the large- n exponent is left to GPU-scale contraction.




Figure A3. OSEE transition at $n = 16, 20$ via the validated TN sampler ($\chi = 256$). The transition persists with a narrow width. Triangles: $n = 16$ at $\chi = 512$, lying above the $\chi = 256$ curve—OSEE is not yet converged, so the absolute p_c is biased low and not quoted.

References

1. Arute, F.; Arya, K.; Babbush, R.; Bacon, D.; Bardin, J.C.; Barends, R.; Biswas, R.; Boixo, S.; Brandao, F.G.; Buell, D.A.; et al. Quantum supremacy using a programmable superconducting processor. *Nature* **2019**, *574*, 505–510. [[CrossRef](#)] [[PubMed](#)]
2. Morvan, A.; Villalonga, B.; Mi, X.; Mandrà, S.; Bengtsson, A.; Klimov, P.V.; Chen, Z.; Hong, S.; Erickson, C.; Drozdov, I.K.; et al. Phase transitions in random circuit sampling. *Nature* **2024**, *634*, 328–333. [[CrossRef](#)] [[PubMed](#)]
3. Ware, B.; Deshpande, A.; Hangleiter, D.; Niroula, P.; Fefferman, B.; Gorshkov, A.V.; Gullans, M.J. A sharp phase transition in linear cross-entropy benchmarking. *arXiv* **2023**, arXiv:2305.04954. [[CrossRef](#)]
4. Aharonov, D.; Gao, X.; Landau, Z.; Liu, Y.; Vazirani, U. A polynomial-time classical algorithm for noisy random circuit sampling. *arXiv* **2023**, arXiv:2211.03999. [[CrossRef](#)]
5. Dalzell, A.M.; Hunter-Jones, N.; Brandão, F.G.S.L. Random quantum circuits transform local noise into global white noise. *Commun. Math. Phys.* **2021**, *405*, 78. [[CrossRef](#)]
6. Auffèves, A. Quantum Technologies Need a Quantum Energy Initiative. *PRX Quantum* **2022**, *3*, 020101. [[CrossRef](#)]
7. Zhao, X.H.; Zhong, H.S.; Pan, F.; Chen, Z.H.; Fu, R.; Su, Z.; Xie, X.; Zhao, C.; Zhang, P.; Ouyang, W.; et al. Leapfrogging Sycamore: Harnessing 1432 GPUs for 7× Faster Quantum Random Circuit Sampling. *arXiv* **2024**, arXiv:2406.18889. [[CrossRef](#)]
8. Meier, F.; Yamasaki, H. Energy-Consumption Advantage of Quantum Computation. *PRX Energy* **2025**, *4*, 023008. [[CrossRef](#)]
9. Soret, A.; Dridi, N.; Wein, S.C.; Giesz, V.; Mansfield, S.; Emeriau, P.E. Quantum Energetic Advantage before Computational Advantage in Boson Sampling. *arXiv* **2026**, arXiv:2601.08068. [[CrossRef](#)]
10. Aaronson, S.; Gunn, S. On the Classical Hardness of Spoofing Linear Cross-Entropy Benchmarking. *arXiv* **2020**, arXiv:1910.12085. [[CrossRef](#)]
11. Gao, X.; Kalinowski, M.; Chou, C.N.; Lukin, M.D.; Barak, B.; Choi, S. Limitations of Linear Cross-Entropy as a Measure for Quantum Advantage. *PRX Quantum* **2024**, *5*, 010334. [[CrossRef](#)]
12. Landauer, R. Irreversibility and Heat Generation in the Computing Process. *IBM J. Res. Dev.* **1961**, *5*, 183–191. [[CrossRef](#)]
13. Bennett, C.H. Logical Reversibility of Computation. *IBM J. Res. Dev.* **1973**, *17*, 525–532. [[CrossRef](#)]
14. Gray, J.; Kourtis, S. Hyper-optimized tensor network contraction. *Quantum* **2021**, *5*, 410. [[CrossRef](#)]
15. Boixo, S.; Isakov, S.V.; Smelyanskiy, V.N.; Babbush, R.; Ding, N.; Jiang, Z.; Bremner, M.J.; Martinis, J.M.; Neven, H. Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **2018**, *14*, 595–600. [[CrossRef](#)]
16. Pan, F.; Chen, K.; Zhang, P. Solving the sampling problem of the Sycamore quantum circuits. *Phys. Rev. Lett.* **2022**, *129*, 090502. [[CrossRef](#)] [[PubMed](#)]
17. Bouland, A.; Fefferman, B.; Nirkhe, C.; Vazirani, U. On the complexity and verification of quantum random circuit sampling. *Nat. Phys.* **2019**, *15*, 159–163. [[CrossRef](#)]
18. Dalzell, A.M.; Hunter-Jones, N.; Brandão, F.G.S.L. Random quantum circuits anticconcentrate in log depth. *PRX Quantum* **2022**, *3*, 010333. [[CrossRef](#)]
19. Chamberland, C.; Zhu, G.; Yoder, T.J.; Hertzberg, J.B.; Cross, A.W. Topological and Subsystem Codes on Low-Degree Graphs with Flag Qubits. *Phys. Rev. X* **2020**, *10*, 011022. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.