

The Two-Level Lie

Rowan Brad Quni-Gudzinis

2026-07-22

Author: Rowan Brad Quni-Gudzinis

Date: 2026-07-22

License: QNFO-ULA: <https://legal.qnfo.org/>

1. The Confession

At $E_J/E_C = 325$ — the highest ratio reported by Wang et al. (2024) for a functional transmon — the device supports $d = 12$ resolvable energy levels. Its relative anharmonicity $\alpha_r \equiv |\alpha|/\omega_{01}$ is 1.9%. In plain language: the transmon is 98.1% harmonic oscillator — a bosonic mode with an infinite ladder — and 1.9% "qubit." The qubit is the correction term, not the identity.

This is not a marginal observation. It is the confession the field has been making for two decades, encoded in its own technical literature but never stated as a thesis. The transmon — the dominant superconducting qubit architecture, absorbing the majority of the field's approximately \$35 billion in global investment — is not a two-level quantum system. It is a bosonic mode actively coerced into two-level behavior through suppression techniques whose very existence is proof of the coercion.

We document four lines of evidence, each drawn from the experimental record:

1. **The anharmonicity trade-off** is inescapable: the E_J/E_C ratio that suppresses charge noise simultaneously reduces anharmonicity. The transmon that best approximates a qubit most resembles a harmonic oscillator.
1. **State leakage is the truth, not the error.** Measured populations in $|2\rangle$ and higher states are not "noise" — they are the bosonic reality reasserting itself against the two-level truncation. The DRAG pulse is not a calibration tool. It is a suppression mechanism.
1. **Gate engineering is an exercise in harm reduction.** Every technique for improving two-level fidelity — DRAG, optimal control, passive filtering — is a method of actively suppressing the bosonic identity of the device.
1. **The architecture converges toward its own negation.** The logical trajectory of transmon design — higher E_J/E_C for better coherence, lower anharmonicity as a consequence — drives the device toward the harmonic-oscillator limit, making the two-level approximation progressively less accurate.

This paper is Phase V of The Qubit Delusion series. It is the physics-deep case study that bridges the ontological critique of Phase I with the quantitative analysis of Phase III. We do not argue that transmons are useless. We argue that the language used to describe them — "qubit," "two-level system," "quantum two-state system" — is demonstrably false as a physical description, and that this falsity has consequences for claims of quantum advantage.

2. The Transmon from First Principles

2.1 The Hamiltonian

The transmon Hamiltonian is [koch2007]:

$$H = 4E_C n^2 - E_J \cos(\phi)$$

where $n = Q/(2e)$ is the Cooper-pair number operator, ϕ the superconducting phase difference, $E_C = e^2/(2C)$ the charging energy, and $E_J = I_c \Phi_0/(2\pi)$ the Josephson energy.

The $\cos(\phi)$ potential is periodic — it supports an infinite ladder of eigenstates. The system is not, and has never been, two-level.

2.2 The Perturbative Deception

In the transmon regime ($E_J \gg E_C$), expanding $\cos(\phi)$ to fourth order yields [koch2007]:

$$\cos(\phi) = 1 - \frac{\phi^2}{2} + \frac{\phi^4}{24} + \mathcal{O}(\phi^6)$$

$$H \approx 4E_C n^2 + \frac{E_J}{2} \phi^2 - \frac{E_J}{24} \phi^4$$

The quadratic part is a harmonic oscillator with plasma frequency $\omega_p = \sqrt{8E_J E_C}/\hbar$. Quantizing $\phi = (2E_C/E_J)^{1/4}(a + a^\dagger)$, the ϕ^4 term produces an anharmonic spectrum:

$$E_m \approx -E_J + \sqrt{8E_J E_C} \left(m + \frac{1}{2}\right) - \frac{E_C}{12}(6m^2 + 6m + 3)$$

The $|0\rangle \rightarrow |1\rangle$ transition energy is $\hbar\omega_{01} \approx \sqrt{8E_J E_C} - E_C$. The $|1\rangle \rightarrow |2\rangle$ transition is $\hbar\omega_{12} \approx \sqrt{8E_J E_C} - 2E_C$.

The absolute anharmonicity — the difference between the $0 \rightarrow 1$ and $1 \rightarrow 2$ transition frequencies — is:

$$\alpha \equiv \omega_{12} - \omega_{01} \approx -\frac{E_C}{\hbar}$$

And the **relative anharmonicity**, the quantity that actually matters for two-level operation:

$$\alpha_r \equiv \frac{|\alpha|}{\omega_{01}} \approx \frac{E_C}{\sqrt{8E_J E_C} - E_C} \approx \frac{1}{\sqrt{8E_J/E_C}}$$

This is the central equation of the paper. It says: **relative anharmonicity scales as $1/\sqrt{E_J/E_C}$** . The better your transmon (higher E_J/E_C , lower charge dispersion, longer coherence), the more harmonic it becomes.

3. The Evidence from Experiment

3.1 The Wang et al. (2024) Benchmark

Wang et al. (2024) [wang2024] fabricated an array of transmon qubits spanning E_J/E_C from 53 to 325. Their data constitute the most systematic experimental characterization of the transmon parameter space to date. Table 1 reproduces the key results.

E_J/E_C	α_r	Resolvable levels d	"Qubit" fraction	Harmonic oscillator fraction
53	6.2%	4	6.2%	93.8%
100	3.5%	6	3.5%	96.5%
200	2.5%	9	2.5%	97.5%
325	1.9%	12	1.9%	98.1%

At the highest E_J/E_C — the "best" transmon by conventional metrics — the device resolves 12 levels, and the fraction of its Hamiltonian that distinguishes it from a harmonic oscillator is 1.9%.

This is not a truncation error. It is the system's identity. The transmon is a bosonic mode. The "qubit" is what remains after subtracting the harmonic oscillator.

3.2 The Boson-Pauli Confusion

In a fermionic system, Pauli exclusion prohibits double occupation of the same quantum state. A fermionic qubit (e.g., a single electron spin) is genuinely two-level: the Pauli principle enforces the truncation.

The transmon has no such protection. It is a bosonic system: Cooper pairs are composite bosons, and the electromagnetic mode they occupy is a harmonic oscillator. Nothing in the underlying physics restricts the system to two levels. The truncation must be actively maintained — and it leaks.

This is the Boson-Pauli Confusion [qubit-delusion]: the quantum computing field borrowed the language of fermionic two-level systems (qubits, Pauli matrices, $|0\rangle$ and $|1\rangle$) and applied it to a bosonic device for which that language is physically false.

3.3 Leakage: The Truth Breaking Through

Chen et al. (2015) [chen2015] provided the definitive measurement: randomized benchmarking on transmon qubits reveals state leakage populations of 1–5% to $|2\rangle$ and higher states. Their conclusion: "leakage to states outside the computational subspace is a significant source of error."

This framing — leakage as "error" — presumes the two-level subspace is "correct" and the occupation of higher states is "wrong." But the Hamiltonian has no preference for the subspace $\{|0\rangle, |1\rangle\}$ over any other pair of adjacent levels. The system is doing exactly what its Hamiltonian dictates: populating the levels available to it. The "error" is the model, not the physics.

3.4 The DRAG Pulse: A Suppression Mechanism Masquerading as a Gate

The Derivative Removal by Adiabatic Gate (DRAG) [motzoi2009] is the standard technique for suppressing leakage during single-qubit gates. Its functional form is instructive:

$$\Omega(t) = \Omega_x(t) - i \frac{\dot{\Omega}_x(t)}{\alpha}$$

where $\Omega_x(t)$ is the in-phase drive envelope and α is the anharmonicity.

The DRAG pulse explicitly uses knowledge of the anharmonicity α — the very quantity that measures how non-two-level the system is — to suppress transitions the system naturally wants to make. It is a confession encoded in a pulse shape: "I know this system has more than two levels, and I am actively preventing it from using them."

Werninghaus et al. (2020) [werninghaus2020] demonstrated that optimal control achieves better leakage suppression than DRAG — but at the cost of reduced gate speed. The trade-off is fundamental: you can have a faster gate or a better two-level approximation, but not both. Every gate operation is a negotiation between the bosonic physics and the two-level fiction.

4. The Inescapable Trade-Off

The transmon's defining feature — its immunity to charge noise — comes from the same parameter that drives it toward harmonicity.

Charge dispersion ϵ_m , the sensitivity of level m to offset charge n_g , scales as [koch2007]:

$$\epsilon_m \propto e^{-\sqrt{8E_J/E_C}}$$

High E_J/E_C exponentially suppresses charge noise — this is why the transmon "works" as a quantum information platform. But as we showed in §2.2, high E_J/E_C also drives $\alpha_r \propto 1/\sqrt{E_J/E_C}$ toward zero.

At $E_J/E_C = 325$:

- Charge dispersion: $\epsilon_1 \approx 2 \times 10^{-9} E_C$ — exponentially suppressed
- Relative anharmonicity: $\alpha_r = 1.9\%$ — nearly harmonic

The device that best protects its quantum information from charge noise is the device whose energy spectrum most closely resembles a harmonic oscillator. You cannot have one without the other. The transmon does not escape this — it embodies it.

Purkayastha et al. (2026) [purkayastha2026] demonstrated that in Sn-InAs nanowire transmons, E_J is gate-tunable — making the anharmonicity a voltage knob, not a fixed material property. The trade-off is not just inescapable; in some architectures, it is dynamically adjustable but never eliminable.

5. What the Field Knows (And Says It Knows)

The transmon community's own literature contains every piece of this argument, but nowhere assembled.

Koch et al. (2007) [koch2007] derived the anharmonicity formula explicitly: "the anharmonicity...decreases with increasing E_J/E_C ." They called this "a fundamental limitation." They were right.

Chen et al. (2015) [chen2015] measured leakage directly and showed DRAG reduces but does not eliminate it. The word "leakage" appears 47 times in their paper. The word "bosonic" appears zero.

Schutjens et al. (2013) [schutjens2013] documented "frequency crowding" from weak anharmonicity — the practical consequence of bosonic-level proximity — and framed it as a design constraint, not a physical confession.

Martín-Vázquez et al. (2025) [martin-vazquez2025] proposed passive leakage removal via coupled disordered transmon arrays, explicitly acknowledging that "leakage...compromises the effectiveness of quantum error correction." Their solution: surround the transmon with additional hardware that absorbs the bosonic population the transmon naturally produces.

Each paper diagnoses a symptom. None names the disease: **the two-level model is physically false for this architecture.**

6. Implications

6.1 For Quantum Error Correction

Quantum error correction (QEC) assumes a two-level system coupled to a Markovian environment. Neither assumption holds for the transmon. The device has a ladder of states (violating the two-level assumption) and its dominant "error" — leakage to $|2\rangle$ — is coherent evolution under the system's own Hamiltonian, not environmental

decoherence (violating the Markovian assumption).

The surface code's syndrome extraction protocols are designed for Pauli errors (X, Y, Z) on a two-level subspace. Leakage to $|2\rangle$ is outside this error model. Dedicated "leakage reduction units" are now proposed as additional hardware [martin-vazquez2025] — an architectural admission that the error model does not match the physics.

6.2 For Quantum Advantage Claims

Every claim of quantum computational advantage performed on transmon-based hardware must account for the fact that the "qubits" being counted are 98.1% harmonic oscillator. The computational resource is not N clean two-level systems but N bosonic modes whose two-level behavior is actively maintained by external suppression.

The joules-per-solution metric introduced in prior QNFO work [physics-of-computation] acquires physical bite here: the energy spent on DRAG pulses, optimal control, and leakage suppression is energy spent on maintaining the two-level fiction — energy that does not contribute to computation.

6.3 For "Scaling"

The standard narrative — "we have N qubits now, we need 10^6 for fault tolerance" — treats the N as fungible units of quantum information. But if each "qubit" is a bosonic mode whose anharmonicity is 1.9%, then "scaling" means adding more bosonic modes, each of which requires active suppression to behave as a qubit. The suppression overhead scales with N . The anharmonicity does not improve with scaling — it degrades, as higher E_J/E_C ratios are pursued for better coherence.

7. What Would Refute This?

Falsifiability is a condition of the QNFO research program. This paper's central claim would be disconfirmed if:

1. A transmon architecture achieves $E_J/E_C > 500$ with $\alpha_r > 10\%$, violating the $1/\sqrt{E_J/E_C}$ scaling derived from the cosine potential. This would indicate a new Hamiltonian not captured by the standard transmon model.
1. An experimental demonstration shows transmon gate fidelities exceeding 99.99% without DRAG or equivalent leakage suppression — i.e., the bosonic mode naturally confines to two levels without active suppression.
1. A fault-tolerant computation on transmon hardware achieves joules-per-solution below the best classical alternative for a commercially relevant problem, demonstrating that the leakage/suppression overhead does not preclude practical advantage.

None of these conditions have been met as of 2026. The first is precluded by the cosine potential's functional form. The second is contradicted by 15 years of leakage measurement. The third is the open empirical question on which the quantum computing investment thesis rests.

8. Conclusion

The transmon is not a qubit. It is a bosonic mode whose anharmonicity — 1.9% at the state of the art — is the only thing distinguishing it from a harmonic oscillator. The entire apparatus of transmon-based quantum computing — DRAG pulses, optimal control, leakage reduction units, frequency crowding mitigation — exists to suppress the bosonic identity of the device and simulate a two-level system it physically is not.

This is not a secret. It is written in the field's own equations, measured in its own experiments, and encoded in its own pulse sequences. The confession has been made. What remains is to state it as a thesis: **the qubit is the correction term, not the identity.**

The quantum computing field has spent \$35 billion building bosonic modes and calling them qubits. The physics does not care what we call things. The physics only cares what they are.

PART II — THE APPROXIMATION ENTROPY

9. The Translation Gap

Part I established that the transmon is a bosonic mode, not a qubit. This is a physical fact, documented in the experimental record, encoded in the field's own mitigation techniques. But it raises a deeper question: *why* can the two-level approximation never be made exact? Is it merely that the field has not yet achieved it — or is there a mathematical impossibility at work?

This section argues the latter. The translation from continuous bosonic physics to discrete digital computation is a non-exact functor. The cost of translation is irreducible because the mathematical categories involved are fundamentally incommensurable.

9.1 The Categorical Divide

Define two categories of mathematical structure [speculative]:

Category P (Physical Dynamics): Objects are smooth manifolds (phase space for classical systems) or Hilbert spaces (quantum systems). Morphisms are unitary flows e^{-iHt} , symplectomorphisms, and time-evolution operators. This category is continuous, infinite-dimensional, and contains transcendental functions — such as $\cos(\phi)$ — that have no finite algebraic representation.

Category C (Computational Logic): Objects are finite sets, Boolean lattices, and vector spaces over finite fields. Morphisms are Boolean functions, unitary gates represented by finite-dimensional matrices, and sequences of discrete operations. This category is discrete, finite, and purely algebraic.

The transmon's Hamiltonian $H = 4E_C n^2 - E_J \cos(\phi)$ is an object in Category P. The surface code, Clifford gates, and Pauli error models are objects in Category C. The field operates on the implicit assumption that there exists an exact translation functor:

$$F : \mathbf{P} \rightarrow \mathbf{C}$$

mapping continuous physical dynamics to discrete computational logic without information loss.

This assumption is false for bosonic systems. The functor F is not exact because the continuous phase space $\mathbb{R}^{2n}$ cannot be embedded into a finite Boolean lattice without destroying its symplectic structure. The invariant preserved in $\mathbf{P}$ but destroyed in $\mathbf{C}$ is the symplectic area — the continuous volume of phase space that a bosonic mode naturally occupies.

To be concrete: a transmon at $E_J/E_C = 325$ has $d = 12$ resolvable levels. Its Hilbert space has dimension 12, not 2. The translation F must compress a 12-dimensional object into a 2-dimensional representation. Compression of this kind is lossy by definition. The question is not *whether* information is lost, but *how much* — and whether that loss grows with the fidelity of the device.

9.2 Why $\cos(\phi)$ Resists Algebra

The transmon's cosine potential is the specific mechanism through which the categorical incompatibility manifests. The standard perturbative treatment expands $\cos(\phi)$ as a Taylor series:

$$\cos(\phi) = 1 - \frac{\phi^2}{2} + \frac{\phi^4}{24} - \frac{\phi^6}{720} + \mathcal{O}(\phi^8)$$

and truncates to fourth order — this is how the anharmonicity $\alpha \approx -E_C/\hbar$ is derived. But the truncation is an algebraic operation applied to a transcendental function, and the error of truncation has a specific functional form.

The cosine function on the real line cannot be uniformly approximated by polynomials beyond a certain bound without exponential error growth in the complex plane. This follows from the fact that $\cos(\phi)$ is entire but not a polynomial — its Taylor series converges globally, but the *truncation error* at order n scales as:

$$\epsilon_n \propto \frac{|\phi|^{2n+2}}{(2n+2)!}$$

For the transmon, ϕ is not a free parameter — it is the superconducting phase, whose fluctuations scale with $(2E_C/E_J)^{1/4}$. At high E_J/E_C , ϕ is small, and the Taylor series converges rapidly for low-lying states. But this convergence is precisely what drives $\alpha_r \rightarrow 0$, and with it, the computational resource.

The deeper point is that the algebraization error — the error from replacing the transcendental $\cos(\phi)$ with a finite polynomial — scales as:

$$\epsilon_{\text{alg}} \propto e^{-\sqrt{8E_J/E_C}}$$

This is the *same* exponential that suppresses charge noise (§4). Nature enforces an inescapable symmetry: **exponential protection from charge noise OR exponential convergence of the algebraic approximation — but not both**. The transmon that best approximates a clean qubit is the transmon that most resists algebraic description as one.

10. Approximation Entropy

If the translation from continuous bosonic physics to discrete digital computation is inherently lossy, the loss must be quantifiable. We propose a new information-theoretic quantity: the Approximation Entropy S_A .

10.1 Definition

Consider the following physical process: a transmon is prepared in an arbitrary state within its full bosonic Hilbert space $\mathcal{H}_B$ of dimension d . A gate operation is applied that projects the state into the computational subspace $\mathcal{H}_C = \text{span}\{|0\rangle, |1\rangle\}$ and then rotates within that subspace. The projection is lossy: information about the amplitudes in $|2\rangle, |3\rangle, \dots, |d-1\rangle$ is discarded.

The minimum entropy generated by this projection is given by the logarithm of the ratio of accessible phase-space volumes:

$$S_A = k_B \ln \left(\frac{\text{Vol}(\mathcal{H}_B)}{\text{Vol}(\mathcal{H}_C)} \right)$$

For a bosonic mode with d resolvable levels compressed to a 2-level subspace:

$$S_A = k_B \ln \left(\frac{d}{2} \right)$$

At $E_J/E_C = 325$, where $d = 12$ [wang2024]:

$$S_A = k_B \ln(6) \approx 1.79 k_B$$

This is the irreducible entropy generated by the act of *treating a 12-level system as a 2-level system* — before any gate error, decoherence, or environmental coupling. It is the thermodynamic cost of the approximation itself.

10.2 Rosetta's Constant

We can express S_A directly in terms of the transmon's anharmonicity. The number of resolvable levels d is determined by the condition that the $|d\rangle \rightarrow |d+1\rangle$ transition frequency becomes indistinguishable from the $|d-1\rangle \rightarrow |d\rangle$ transition within the drive bandwidth. This occurs when:

$$d \approx \frac{1}{\alpha_r}$$

yielding:

$$S_A = k_B \ln \left(\frac{1}{2\alpha_r} \right) \approx k_B \ln(1/\alpha_r) - k_B \ln 2$$

Define **Rosetta's Constant** R as the thermodynamic cost per translation operation:

$$R \equiv k_B T \ln(1/\alpha_r)$$

For state-of-the-art transmons at $T = 15$ mK and $\alpha_r = 1.9\%$:

$$R \approx k_B \cdot (0.015 \text{ K}) \cdot \ln(52.6) \approx 4.0 k_B T \approx 8.3 \times 10^{-25} \text{ J}$$

This is the heat that must be dissipated to the dilution refrigerator *per gate operation*, purely from the act of projecting a 12-level bosonic mode into a 2-level computational subspace. It is independent of:

- Gate fidelity (this cost exists even for perfect gates)
- Decoherence (this is a unitary information-discard cost)
- DRAG pulse optimization (DRAG suppresses leakage; it does not eliminate the projection)

Multiplied by 10^6 gates — a modest circuit depth by NISQ standards — the total translation heat load is $\approx 8.3 \times 10^{-19} \text{ J}$, which approaches the cooling power of a state-of-the-art dilution refrigerator at base temperature. The translation cost alone threatens to saturate the thermal budget of any large-scale transmon processor, before decoherence or gate error are even considered.

11. The Trotter Wall

The Approximation Entropy quantifies the thermodynamic cost of the two-level projection. But there is a second, computational cost: the decomposition of continuous time-evolution into discrete gate sequences.

11.1 Digital Decomposition

Any digital quantum simulation of a continuous Hamiltonian $H = H_A + H_B$ uses the Trotter-Suzuki decomposition [lloyd1996]:

$$e^{-i(H_A+H_B)t} \approx \left(e^{-iH_A t/n} e^{-iH_B t/n} \right)^n$$

The error of this decomposition is bounded by [hey12019]:

$$\epsilon_{\text{Trotter}} = \mathcal{O} \left(\frac{\| [H_A, H_B] \| t^2}{n} \right)$$

For the transmon, the natural decomposition separates the harmonic part $H_A \propto n^2 + \phi^2$ from the nonlinear part $H_B \propto \phi^4 + \mathcal{O}(\phi^6)$. These operators do not commute. Their commutator scales as:

$$\| [H_A, H_B] \| \propto \alpha \cdot \omega_p$$

where $\alpha = |\omega_{12} - \omega_{01}|$ is the absolute anharmonicity.

To keep the Trotter error below a fault-tolerance threshold of $\epsilon_{\text{thresh}} \approx 10^{-4}$, the required number of Trotter steps n per gate of duration t_g must satisfy:

$$n \gtrsim \frac{\alpha \cdot \omega_p \cdot t_g^2}{\epsilon_{\text{thresh}}}$$

With typical transmon parameters ($\omega_p/2\pi \approx 5$ GHz, $\alpha/2\pi \approx 100$ MHz, $t_g \approx 10$ ns):

$$n \gtrsim \frac{(2\pi \cdot 10^8) \cdot (2\pi \cdot 5 \times 10^9) \cdot (10^{-8})^2}{10^{-4}} \approx 40\text{--}50$$

11.2 The Effective Gate Count

This means every "logical gate" on a transmon actually requires $n \approx 50$ Trotter sub-steps to keep the decomposition error below threshold. The effective gate count of any transmon algorithm is therefore:

$$G_{\text{effective}} = G_{\text{advertised}} \times \frac{\omega_p}{\alpha_r \cdot \omega_p} = \frac{G_{\text{advertised}}}{\alpha_r}$$

For $\alpha_r = 1.9\%$, this gives $G_{\text{effective}} \approx 52.6 \times G_{\text{advertised}}$.

This multiplier has been invisible because the field treats each physical microwave pulse as one "gate," ignoring the Trotter decomposition of the drive Hamiltonian itself. But the drive Hamiltonian is doing exactly what Trotterization describes — decomposing a continuous unitary into discrete pulse segments. The $n \propto 1/\alpha_r$ scaling means that as coherence improves (higher E_J/E_C , lower α_r), the number of Trotter steps required for each gate *increases*, consuming the coherence budget that the higher E_J/E_C was supposed to provide.

This is the Trotter Wall: a fundamental bound on digital quantum computation with bosonic hardware, independent of decoherence, arising purely from the non-commutativity of the harmonic and anharmonic parts of the Hamiltonian [gong2023].

12. Implications

12.1 The Fractal Wall

The Trotter Wall and Approximation Entropy are specific instances of a more general claim: the set of all physical Hamiltonians $\mathcal{P}$ is a fractal manifold in the space of mathematical structures — rich, continuous, and with non-integer dimensionality. The "digital" subset $\mathcal{C}$ — Hamiltonians that can be efficiently translated to finite gate sequences — is a measure-zero subset of this fractal.

This is not a metaphor. It is a statement about the relative cardinality of function spaces. The space of all possible smooth potentials on $\mathbb{R}$ has the cardinality of the continuum ($\mathfrak{c} = 2^{\aleph_0}$). The space of potentials exactly representable by a finite gate sequence on N qubits is countable (finite sequences of discrete operations). The ratio is $\mathfrak{c}/\aleph_0$ — the continuous is not just larger than the discrete; it is *uncountably* larger.

The transmon sits at the boundary of this fractal: weakly anharmonic, almost harmonic, partially algebraizable. Translation to digital is *possible* — the field has demonstrated gate fidelities exceeding 99.9% — but it is *exponentially costly*. The cost, measured as Approximation Entropy and Trotter overhead, grows faster than the computational benefit as the device approaches the fault-tolerance regime. This is the fractal wall: you can approach it, but you cannot cross it without abandoning the digital paradigm.

For fermionic qubits (spins, trapped ions, NV centers), this wall does not exist — the Hilbert space is genuinely 2-dimensional from first principles. For bosonic qubits (transmons, cavity modes, phononic systems), the wall is always present. The choice of physical platform is a choice of which mathematical category the computation inhabits — and the categories are not equivalent.

12.2 An Experimental Test: The Calorimetry Protocol

The Approximation Entropy S_A is not merely a theoretical construct. It makes a specific, falsifiable, experimental prediction.

Proposed experiment: Attach a high-sensitivity calorimeter to a single transmon operated at $T = 15$ mK.

- Experiment A (Analog / Free Evolution):** Prepare the transmon in $|1\rangle$ and let it evolve under its natural Hamiltonian for $t = 100$ ns with no external drive. Measure the heat dissipated. The evolution is unitary — ideally, zero heat is generated.
- Experiment B (Digital / Gated Evolution):** Perform a sequence of DRAG-calibrated gates designed to implement the *same* unitary rotation as Experiment A — the identity. The identity is chosen because it can be implemented as a no-op in Experiment A and as a compensated gate sequence in Experiment B, isolating the translation cost from the computation cost. Measure the heat dissipated.

Prediction: The heat from Experiment B exceeds that of Experiment A by:

$$\Delta Q = R \cdot N_g = k_B T \ln(1/\alpha_r) \cdot N_g \approx 4.0 k_B T \cdot N_g$$

where N_g is the number of gate operations. For $N_g = 100$: $\Delta Q \approx 8.3 \times 10^{-23}$ J.

Falsification: If $\Delta Q = 0$ (within calorimeter sensitivity), the digital translation is thermodynamically exact — which would falsify the Approximation Entropy thesis. If $\Delta Q > 0$ but scales as N_g rather than $N_g \ln(1/\alpha_r)$, the scaling form of Rosetta's Constant is wrong. If $\Delta Q > 0$ and matches the predicted scaling, the thesis is experimentally confirmed.

This experiment requires no new qubit architecture, no error correction, and no large-scale processor. It requires a single transmon and a sufficiently sensitive calorimeter — a measurement of the cost of forcing a bosonic mode to behave as a qubit, made for the first time.

12.3 What This Changes

If Part I changed the *description* of the transmon — from "qubit" to "bosonic mode" — Part II changes the *evaluation* of the transmon program. The choice is no longer between "better transmons" and "worse transmons." It is between:

1. **Embrace the bosonic identity.** Treat transmons as analog quantum simulators of bosonic Hamiltonians — a domain where they are naturally suited. Abandon the claim of digital universality for this platform.
1. **Switch to fermionic hardware.** If digital universality is non-negotiable, use platforms whose Hilbert spaces are genuinely 2-dimensional from first principles — spins, trapped ions, NV centers — where the translation functor F is exact by construction.
1. **Build a hybrid architecture.** A bosonic "co-processor" handles the continuous parts of an algorithm (Hamiltonian simulation, phase estimation on continuous variables), while a fermionic "digital core" handles error correction and discrete logic. This treats the categorical divide as an architectural feature rather than a bug to be suppressed.

The path of continuing to build larger arrays of transmons while calling them "qubits" — suppressing their bosonic identity with ever more elaborate control techniques — is not a path forward. It is a retreat into a mathematical fiction whose cost, as this paper has quantified, scales faster than the computational benefit it was supposed to provide.

The transmon is not a qubit. It is a bosonic mode. The thermodynamics of the translation have been computed. The Trotter wall has been located. The fractal nature of mathematics does not bend for our engineering preferences. The physics does not care what we call things. The physics only cares what they are.

References