The Ignorant Qubit

Information-Theoretic Bounds on Scalable Quantum Computing

Aernoud Dekker

December 2025
Version 1.0

DOI: 10.17605/OSF.IO/G293X

Abstract

Current approaches to Fault-Tolerant Quantum Computing (FTQC) assume that scaling is limited primarily by environmental decoherence and gate fidelity. Building on the Data-Rate Theorem from control theory and recent forensic analysis of quantum processor recovery dynamics, we **propose a hypothesis**: that a more fundamental constraint may exist—the finite information processing capacity of the classical control system. By modeling the Quantum-Classical Interface (QCI) as a rate-limited information channel, we derive a condition where the entropy generation rate of the N-qubit system (λ_N) exceeds the error-correction bandwidth ($C_N \ln 2$). Under this model, hardware platforms exhibit a **structural bifurcation**: the model forecasts a qubit-count wall for superconducting systems ($N_{\max} \sim 10^3$–10^6 depending on the scaling exponent p), while photonic systems would face a threshold at $N = 1$ but scale linearly thereafter. Trapped ions avoid both walls but face bandwidth latency limiting algorithm depth. The only theoretical escape—native topological qubits with sublinear scaling ($p < 1$)—remains experimentally unrealized; synthetic topological codes on standard hardware inherit the limits of the substrate. We present correlational evidence from Google Sycamore (6.5% stable delayed-geometry signatures under rigorous model-based classification, with 64% boundary events) and Chinese 63-qubit processors (uniformly fast recovery, serving as a qualitative capacity-wins baseline), consistent with the predicted two-regime structure. If the model is correct, the thermodynamic coupling between controller bandwidth and system entropy (the "Catch-22") would imply that no current hardware platform possesses the simultaneous C, λ, and p parameters required for cryptographic-scale fault tolerance. **We do not claim that fault-tolerant quantum computing is impossible in principle.** Rather, we test a specific bottleneck hypothesis: that for cryogenic platforms where controller bandwidth, heat load, and multi-qubit error correlations co-scale, the classical control loop may saturate before cryptographic-scale fault tolerance is reached. We propose a decisive experimental test—the Power-Scaling Test—that can confirm or falsify the hypothesis.

"Each individual Anu is called Avidya, Ignorance."
— Sri Yukteswar, *The Holy Science* (Sutra 4)

Contents

1 Introduction

- **The Problem:** The "Quantum Scaling Stagnation." Despite improvements in T_1/T_2, scaling logical qubits has proven non-linear and difficult. Increasing Code Distance (d) often yields diminishing returns.
- **The Standard View:** Scaling is limited by correlated noise, cosmic rays, and crosstalk. Solution: Better materials, shielding, and larger codes.
- **The Proposed View:** Scaling is limited by the **Thermodynamics of Control**. The classical controller must track the quantum state's phase evolution. This tracking requires a bandwidth that scales with system entropy.
- **Hypothesis:** There exists an **Ignorance Wall** where the complexity of the entangled state grows faster than the controller's ability to ingest syndrome data, leading to a "Blind Spot" where error correction fails.

Paper Structure. Section 2 reviews the Data-Rate Theorem and existing empirical evidence for capacity-limited behavior in quantum hardware. Section 3 derives the tracking equation and defines the stability condition. Section 4 applies this framework to predict scaling laws for different qubit technologies, identifying the "Ignorance Wall." Section 6 proposes experimental tests that could falsify the framework. Section 7 argues that the capacity constraint is fundamental, not merely technological. Section 8 summarizes implications and limitations.

Claims vs. Non-Claims

- **Claim (conditional):** Under the tracking ansatz (Eq. 5), platforms with super-linear λ_N and $\leq$linear C_N exhibit an $N_{\max}$.
- **Claim (testable):** At fixed chip temperature, increasing effective syndrome throughput should increase coherence/logical performance if the system is capacity-limited.
- **Non-claim:** We do not prove that QEC must fail for all architectures, nor that quantum computation is impossible in principle.

2 Prior Work: The Information-Theoretic Control Limit

The Ignorance Wall hypothesis builds upon the *Ignorant Observer Framework* (IOF), which applies the Data-Rate Theorem from control theory to quantum measurement [1]. This section summarizes the theoretical foundation and empirical validation that motivates the present analysis.

2.1 The Data-Rate Theorem

The Data-Rate Theorem [4,5] establishes a fundamental limit on stabilizing unstable dynamical systems over finite-capacity channels. For a system with Lyapunov exponent λ (measuring the rate of trajectory divergence), the minimum channel capacity required for stabilization is:

$$C_{\min} = \frac{\lambda}{\ln 2} \quad [\text{bits/s}] \tag{1}$$

Within its modeling assumptions, this is a mathematical theorem: if $C < \lambda/\ln 2$, tracking fails. We use this as a guiding bound for the QEC control loop, while acknowledging that the mapping to real QEC dynamics involves the ansatz developed in Section 3.

2.2 Two Regimes

The IOF identifies two distinct operational regimes separated by the critical inequality $\lambda \lessgtr C \ln 2$:

- **Capacity-Wins** ($C \ln 2 > \lambda$): The observer/controller has sufficient bandwidth to track the system. Standard quantum mechanics is recovered; coherence is maintained.
- **Chaos-Wins** ($\lambda > C \ln 2$): The system's complexity exceeds the controller's tracking capacity. The controller becomes "informationally blind," leading to loss of phase coherence.

The information deficit rate $\kappa = \lambda - C \ln 2$ (in nats/s) governs the timescale of tracking failure. Since amplitude σ grows as $e^{\kappa t}$ (variance σ^2 as $e^{2\kappa t}$), the amplitude e-folding time is:

$$\tau_{\text{loss}} = \frac{1}{\kappa} = \frac{1}{\lambda - C \ln 2} \quad (\kappa > 0) \tag{2}$$

Throughout this work, τ_{loss} refers to amplitude e-folding; the variance e-folding time is $\tau_{\text{var}} = 1/(2\kappa)$.

2.3 Empirical Validation: Forensic Signatures

The two-regime framework has been investigated through forensic analysis of existing experimental datasets [2]. The primary methodology uses **model-based classification**: fitting multiple functional forms (exponential, sigmoid, delayed-exponential) to recovery curves and selecting the best model via AICc. Events are classified across multiple analysis windows to identify **stable** populations (consistent classification) versus **boundary/flip** events (classification-sensitive).

- **Stable Fast** (capacity-wins geometry): Immediate exponential recovery; best-fit model is exponential or rational across all windows.
- **Stable Delayed** (chaos-wins geometry): Delayed recovery onset; best-fit model is sigmoid or delayed-exponential across all windows.
- **Flip/Boundary**: Classification varies with analysis window; events near the decision threshold with weak model discrimination.

Analysis of two quantum processor architectures revealed:

1. **Chinese 63-qubit processor** (Li et al., 2025): Uniformly fast exponential recovery (sub-millisecond), serving as a qualitative capacity-wins baseline. Due to differences in observables (charge-parity jumps vs. aggregate error counts) and sampling cadence, this is treated as a qualitative comparison rather than quantitative validation.

2. **Google Sycamore** (McEwen et al., 2022): Of 230 cosmic ray events, 29.6% showed stable fast geometry, 6.5% showed stable delayed geometry (consistent hesitation signature), and 63.9% were boundary/flip events with classification-dependent outcomes. The stable delayed population provides weak evidence for hesitation-like dynamics—sufficient to warrant prospective tests, but not strong enough to claim a robust signature.

Similar patterns have been observed in LIGO gravitational wave interferometers, where 33.6% of analyzable glitches show stable delayed geometry with strong curvature discrimination (AUC = 0.950) [2]. The full forensic analysis across multiple platforms is presented in [2].

Key Finding: Both regimes appear to exist in current hardware. The evidence is correlational and consistent with a physical phase boundary, but controlled experiments (varying bandwidth at constant temperature) are needed to establish causation. This paper extends this framework to analyze how the Ignorance Wall would scale with system size N if the hypothesis is correct.

3 Theoretical Framework

Notation: In this paper we write λ for the effective entropy-rate (corresponding to h_{KS} in the experimental protocol) and C for effective syndrome throughput in bits/s (corresponding to C_{eff} in the protocol). The deficit rate $\kappa = \lambda - C \ln 2$ plays the same role as κ in the main IOF framework.

3.1 The Controller as a Finite Observer

We model the quantum computer as a closed-loop control system:

1. **Plant:** The qubit array, whose quantum state evolves under both coherent dynamics and environmental perturbations.
2. **Observer:** The classical controller (FPGA/Cryo-CMOS), which estimates the system state from syndrome measurements.
3. **Channel:** The readout/control lines, characterized by finite Shannon capacity C (bits/s).

3.2 Derivation of the Tracking Equation

The tracking equation follows from standard results in stochastic control theory [4,5]. We present a **modeling ansatz**—not a theorem of nature—that connects these results to quantum error correction. Whether this model class accurately describes QEC dynamics is an empirical question (see Section 6).

What σ^2 Represents in QEC:

In quantum error correction, the decoder maintains an estimate of the *Pauli frame*—the cumulative record of which errors have occurred and been corrected. Let $\mathbf{e}(t)$ denote the true error configuration (a vector over the Pauli group) and $\hat{\mathbf{e}}(t)$ the decoder's estimate. We define $\sigma^2 = \mathbb{E}[\|\mathbf{e} - \hat{\mathbf{e}}\|^2]$ as the expected squared error in the Pauli frame estimate, normalized to be dimensionless.

This is *not* the same as the logical error rate—it is the decoder's *uncertainty* about the current error state, which determines whether future corrections will be applied correctly.

Step 1: Error Accumulation (Entropy Production)

In the absence of syndrome measurements, errors accumulate stochastically. For an N-qubit system with per-qubit error rate p and measurement cycle time τ_c, approximately Np/τ_c new error bits are generated per second. Each error bit adds ~ 1 nat of entropy to the Pauli frame. We model this as:

$$\left.\frac{d\sigma^2}{dt}\right|_{\text{errors}} \approx \lambda_{sys} \tag{3}$$

where λ_{sys} is the effective entropy generation rate. The identification $\lambda_0 \sim 1/T_2$ (base dephasing rate) is an **ansatz**: we posit that the intrinsic decoherence rate sets the baseline entropy generation. This is physically motivated—dephasing represents loss of phase information—but the precise coefficient is system-dependent and should be calibrated empirically.

Step 2: Syndrome Acquisition (Entropy Extraction)

Each syndrome measurement provides classical information about the error state. By Shannon's source coding theorem, a channel of capacity C bits/s can reduce entropy at most at rate C bits/s. For a Kalman-like estimator, variance reduction is multiplicative:

$$\left.\frac{d\sigma^2}{dt}\right|_{\text{syndromes}} = -C \ln 2 \cdot \sigma^2 \tag{4}$$

The factor $\ln 2$ converts bits to nats. The multiplicative dependence on σ^2 reflects that measurements are more informative when uncertainty is larger.

Step 3: The Combined Model

Combining error accumulation with syndrome-based correction yields the competition equation:

$$\frac{d\sigma^2}{dt} = \lambda_{sys} - C_{ctrl} \ln 2 \cdot \sigma^2 \tag{5}$$

This is a **linear variance-balance ansatz** inspired by Kalman filtering [11] and the Data-Rate Theorem [4,5]. We do not claim it is the unique or exact dynamics of QEC decoders—real decoders (e.g., MWPM, Union-Find) have discrete, nonlinear dynamics. Rather, we propose that Eq. (5) captures the *essential competition* between entropy production (λ) and entropy extraction (C), and that this competition determines whether tracking succeeds or fails.

Relation to Open-Loop Tracking (Section 2):

The QEC tracking equation differs structurally from the pure exponential growth $\sigma^2(t) = \sigma_0^2 e^{2\kappa t}$ described in Section 2. In that open-loop setting, errors *compound* multiplicatively (chaotic trajectory divergence), and capacity only reduces the growth *rate*. Here, errors *inject* at constant rate λ (additive), while syndrome measurements *reduce* variance proportionally (multiplicative feedback). The result is a steady-state rather than unbounded growth. Failure in QEC occurs when this steady-state exceeds the code threshold ($\sigma_{ss}^2 > \epsilon$), not from divergence. Crucially, the regime boundary $\lambda > C \ln 2$ is identical in both models—only the within-regime dynamics differ.

Units and Dimensional Consistency:

We work in a normalized coordinate system where σ^2 is **dimensionless**. In this convention:

- λ has units of [1/s] = [nats/s], interpretable as entropy generation rate

- C has units of [bits/s]; the factor $\ln 2$ converts to [nats/s]
- The ratio $\lambda/(C \ln 2)$ is dimensionless, as required for σ_{ss}^2

Connection to Classical Information Throughput:

For quantum systems, C_{ctrl} represents the *classical* information rate at which syndrome data can be extracted and processed. For syndrome extraction in QEC, a natural **entropy-rate upper bound heuristic** is:

$$C_{ctrl} \lesssim \sum_{i=1}^{N} r_i \cdot H(p_i) \tag{6}$$

where r_i is the measurement rate for syndrome bit i, and $H(p_i) = -p_i \log_2 p_i - (1-p_i) \log_2 (1-p_i)$ is the binary entropy of that syndrome outcome. This is not a rigorous theorem but a modeling approximation: we assume the classical information extracted per measurement is bounded by the entropy of the syndrome outcomes. The Holevo capacity (which bounds quantum-to-classical information transfer for general quantum ensembles) addresses a different question; here we are concerned with the rate of *classical* syndrome bits reaching the decoder.

This bound establishes that C_{ctrl} is not arbitrary but determined by measurable experimental parameters: the syndrome extraction rate and the entropy of syndrome outcomes.

Where is the Bottleneck? The capacity C in this framework refers to the rate at which *classical* information about the quantum state can be extracted and processed—not the quantum measurement bandwidth per se. In a QEC system, the bottleneck is typically the classical processing chain: syndrome decoding, error identification, and feedback computation. While quantum measurements may occur rapidly, the classical controller must interpret and act on this data. The limit is thus on *classical syndrome throughput*—how many bits of error information per second can be extracted, decoded, and acted upon.

3.3 The Stability Condition

The steady-state solution of Eq. (5) is:

$$\sigma_{ss}^2 = \frac{\lambda_{sys}}{C_{ctrl} \ln 2} \tag{7}$$

For QEC to succeed, this uncertainty must remain below the code's error threshold ϵ:

$$\sigma_{ss}^2 < \epsilon \quad \Longrightarrow \quad C_{ctrl} > \frac{\lambda_{sys}}{\epsilon \ln 2} \tag{8}$$

When $C_{ctrl} < \lambda_{sys}/(\epsilon \ln 2)$, the steady-state uncertainty exceeds the code's error threshold: tracking fails not because the variance diverges, but because it stabilizes at an unacceptably high value. This is the **capacity-limited** regime: the controller cannot reduce uncertainty below the level required for reliable error correction.

3.4 The Ignorant Observer Principle

The tracking equation (Eq. 5) is not an ad-hoc model for quantum control—it is a universal constraint on *any* physical observer. The underlying principle is simple [1]:

1. **The controller is a physical system** with internal dynamics characterized by a divergence rate λ (how fast internal states drift apart).

2. **The controller has finite capacity** C (bits/s)—no physical substrate can process unlimited information.

Any system tracking an internal variable (whether a quantum error syndrome, a measurement basis angle, or a classical reference signal) obeys the fundamental competition:

$$\frac{d\sigma^2}{dt} = \underbrace{\lambda}_{\text{chaos adds variance}} - \underbrace{C \ln 2 \cdot \sigma^2}_{\text{capacity reduces variance}}$$

This yields an irreducible steady-state uncertainty:

$$\sigma_{ss}^2 = \frac{\lambda}{C \ln 2} \tag{9}$$

Structural Correspondence 3.1 (Epistemic vs. Ontological Interpretation). The IOF interprets this limit as epistemic, not ontological: σ^2 reflects the controller's inability to trace the causal history of the system's state. Under this interpretation, the quantum system has a definite configuration at every moment; the controller simply cannot track it perfectly. This "informational blindness" is what the Ignorance Wall describes. **Note:** This interpretive stance is not required for the operational predictions of the model—readers may treat the framework instrumentally without adopting its ontological commitments.

Application to QEC: In quantum error correction, the "observer" is the classical control system (FPGAs, decoders, feedback loops). Its λ is set by the rate at which the quantum system generates entropy (errors); its C is set by the syndrome extraction and processing bandwidth. The Ignorant Observer Principle guarantees that *any* such controller faces the same fundamental trade-off, regardless of implementation details.

3.5 Scope and Assumptions

Operational Definitions for This Framework

- λ_N **(entropy generation rate):** The rate at which the Pauli frame uncertainty grows in the absence of syndrome measurements. Operationally: could be estimated from the rate of syndrome bit flips in a free-running (no correction) experiment, or from $\lambda_0 \approx 1/T_2$ scaled by the crosstalk exponent p.
- C_N **(syndrome extraction capacity):** The mutual information rate between the true error state and the classical syndrome stream. Operationally: $C_N \approx N \cdot r_{syn} \cdot H(p_{syn})$, where r_{syn} is the syndrome measurement rate and $H(p_{syn})$ is the entropy per syndrome bit.
- **Failure criterion:** The model predicts failure when $\sigma_{ss}^2 = \lambda_N/(C_N \ln 2) > \epsilon$, i.e., when steady-state uncertainty exceeds the code's error threshold. In practice, this manifests as logical error rate exceeding a threshold, or coherence time falling below the algorithm runtime.
- **Scaling exponent** p**:** Defined by $\lambda_N = \lambda_0 \cdot N^p$. Currently estimated from device architecture (e.g., $p \approx 1.5$ for surface codes with ZZ crosstalk), but should be measured directly from multi-qubit coherence scaling.

Key assumption: The linear variance-balance model (Eq. 5) captures the essential dynamics. This is a modeling ansatz, not a proven theorem.

4 Scaling Laws: The "Ignorance Wall"

4.1 Scaling of Chaos (λ)

The system Lyapunov exponent λ_N characterizes how fast unknown perturbations (dephasing, crosstalk, cosmic rays) corrupt the global phase reference. We consider three physical mechanisms:

1. Coupling-Induced Chaos: In a system of N qubits with pairwise interactions, the number of coupling terms scales as $\binom{N}{2} \propto N^2$ for fully connected graphs, or $\propto N$ for nearest-neighbor topologies. Each coupling channel can propagate errors, contributing to λ.

2. Entanglement-Enhanced Sensitivity: Entangled states are exponentially more sensitive to local perturbations than product states. A single-qubit error in a GHZ state corrupts the entire N-qubit superposition. This suggests λ grows with entanglement depth, not just qubit count.

3. Correlated Noise: Cosmic rays, TLS defects, and thermal fluctuations create spatially correlated errors that cannot be modeled as independent single-qubit noise. These correlations increase effective λ beyond naive single-qubit estimates.

The Role of QEC in Error Localization:

A key objection is that QEC codes are specifically designed to *localize* errors, potentially making λ scale sublinearly. We address this directly:

- **Surface codes** [13] localize errors to topological defects (anyons), but syndrome extraction

still requires tracking $O(N)$ measurement outcomes per cycle. The *information rate* to process these syndromes scales with N.

- **Error localization reduces the logical error rate**, but does not reduce the *entropy generation rate* that the controller must process. The controller still receives $\sim N$ bits of syndrome data per cycle.
- **When errors are correlated** (as in cosmic ray events), error localization partially fails. The effective λ increases because multiple stabilizers are triggered simultaneously.

We therefore parameterize λ_N as:

$$\lambda_N = \lambda_0 + \alpha N^p \tag{10}$$

where $p > 1$ reflects superlinear scaling due to correlations. The exponent p is architecture-dependent:

- $p \approx 1.2$–1.3: Well-isolated systems with minimal crosstalk
- $p \approx 1.5$: Typical surface code with residual ZZ coupling
- $p \approx 1.8$–2.0: Dense coupling or high crosstalk environments

When might λ scale sublinearly?

If future architectures achieve:

- Perfect error localization with no correlated noise
- Topological protection that intrinsically suppresses error propagation
- Quantum LDPC codes with constant-rate syndrome extraction

then p could approach or fall below 1, eliminating the ignorance wall. Recent advances (e.g., Google's Willow processor, December 2024 [15]) have demonstrated below-threshold operation where *logical* error rates decrease with code distance. However, this does not eliminate the capacity constraint: the controller still processes $O(N)$ syndrome bits per cycle, and rare correlated error bursts (e.g., cosmic ray events) may create an ultimate floor on achievable error rates. The entropy generation rate λ that the controller must track remains distinct from the logical error rate after decoding.

4.2 Scaling of Capacity (C)

The effective controller capacity C_N depends on the syndrome extraction architecture:

Baseline (Independent Readout):

$$C_N = N \cdot C_{qb} \tag{11}$$

where $C_{qb} \approx 10^6$ bits/s per qubit is an order-of-magnitude estimate for superconducting systems (actual values depend on readout architecture and are not precisely characterized in the literature).

Sublinear Scaling Factors:

- **Wiring constraints:** Physical routing limits the number of independent readout channels.
- **Multiplexing overhead:** Frequency-multiplexed readout reduces effective bandwidth per qubit.
- **Classical processing latency:** Syndrome decoding time grows with code size.
- **Heat load:** Controller power dissipation limits cryogenic bandwidth.

These factors suggest C_N scales at best linearly, and often sublinearly, with N.

4.3 The Inequality

The **Ignorance Wall** occurs when chaos generation exceeds capacity:

$$\lambda_N > C_N \ln 2 \quad \Longrightarrow \quad \lambda_0 + \alpha N^p > C_0 N \ln 2 \tag{12}$$

For $p > 1$, there exists a critical $N_{\max}$ beyond which the inequality is always satisfied. Beyond this point, adding qubits adds more entropy than the controller can drain. The logical qubit collapses not because of noise, but because the controller is saturated.

5 Empirical Evidence: The Signature of Saturation

5.1 Retrospective Analysis: Model-Based Classification

We present forensic analysis of high-energy events (cosmic rays) in superconducting processors using **model-based classification** [2]. The methodology fits multiple functional forms (exponential, sigmoid, delayed-exponential, rational) to each recovery curve and selects the best model via AICc (corrected Akaike Information Criterion). Events are classified across multiple analysis windows (60, 100, 150 ms) to identify stable populations.

Classification Categories:

- **Stable Fast:** Best-fit model is fast-geometry (exponential or rational) across all windows. Consistent with capacity-wins dynamics.
- **Stable Delayed:** Best-fit model is delayed-geometry (sigmoid or delayed-exponential) across all windows. Consistent with chaos-wins dynamics where $\lambda > C \ln 2$.
- **Flip/Boundary:** Classification varies across windows; events near the decision threshold with weak model discrimination ($|\Delta\text{AICc}| < 4$).

This approach avoids the "Average Fallacy" by revealing population structure rather than treating all events as a single distribution.

Methodological Note: The model-based approach supersedes earlier derivative-test heuristics (t_{peak} classification), which produced inflated estimates of the delayed-geometry fraction. The derivative test remains a useful intuitive entry point but is not the primary classifier. Formal mixture modeling (Hartigan's dip test) has been applied to LIGO data, confirming departure from unimodality in the pooled distribution [2].

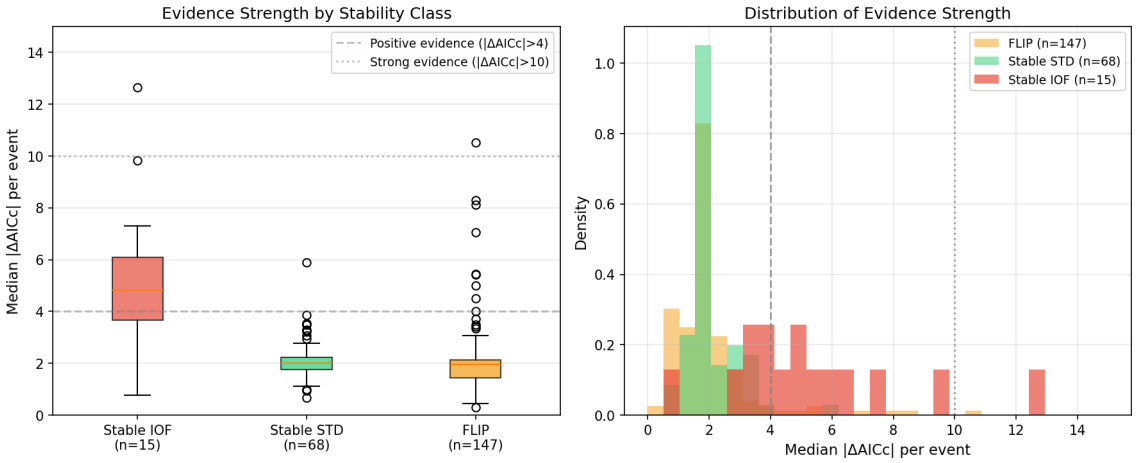

Figure 1: Evidence strength by stability class in Google Sycamore cosmic ray data (n=230 events). Left: Median $|\Delta\text{AICc}|$ per event by stability class. Right: Distribution of evidence strength. Stable IOF (delayed-geometry, n=15, 6.5%) shows stronger model discrimination than Stable STD (fast, n=68, 29.6%) or Flip events (n=147, 63.9%). Dashed lines indicate positive ($|\Delta\text{AICc}| > 4$) and strong ($|\Delta\text{AICc}| > 10$) evidence thresholds. Data from [6], analysis from [2].

5.2 Empirical Validation: Two Regimes Observed

Comparison of cosmic ray recovery data from two independent superconducting qubit experiments reveals patterns consistent with the Ignorance Wall framework:

System	Qubits	Recovery Time (τ)	Recovery Shape	IOF Interpretation
Google (McEwen 2022)	26	∼26 ms median (IQR: 24–28 ms)	**Mixed:** 29.6% Stable Fast 6.5% Stable Delayed 63.9% Boundary	Mixed regime (marginal C/λ)
Chinese (Li et al. 2025)	63	22–213 μs	**Uniformly Fast** (qualitative baseline)	Capacity-Wins ($C \gg \lambda/\ln 2$)

Table 1: Comparison of recovery regimes. Google results from model-based classification with 3-window stability analysis [2]. Chinese results treated as qualitative baseline due to different observables and pipeline.

Mixed Regime (Google): In the Google Sycamore processor, model-based classification identifies a small but stable delayed-geometry population (6.5% of 230 events). These events show nonzero onset delay ($D > 0$) before recovery begins. The majority (63.9%) are boundary/flip events where classification depends on analysis window choice, indicating they lie near the decision threshold. Notably, early-time curvature does *not* discriminate between populations in this platform (Mann-Whitney $p = 0.44$); the evidence rests on model geometry alone.

The stable delayed population provides weak evidence for hesitation-like dynamics—sufficient to warrant prospective tests that directly manipulate controller bandwidth, but not strong enough (absent classifier-external discrimination) to claim a robust IOF signature.

Capacity-Wins Baseline (Chinese): The 63-qubit processor showed uniformly fast exponential recovery (sub-millisecond). This is consistent with capacity-wins operation where the

effective monitoring bandwidth exceeds the perturbation rate. However, direct comparison is limited by differences in observables (charge-parity jumps vs. aggregate error counts) and sampling cadence (5.6 μs vs. 100 μs).

Interpretation: The contrast is *suggestive* of two-regime structure but not *confirmatory*. The Chinese system appears to operate in the capacity-wins regime, while a subset of Google events show dynamics consistent with the chaos-wins regime. Definitive validation requires controlled experiments where bandwidth is varied on a single system while holding other parameters constant (Section 6).

Caveats on Cross-System Comparison:

We acknowledge significant limitations in comparing these two systems [2]:

- **Different observables:** Google data tracks aggregate error counts; Chinese data tracks charge-parity jumps and bit-flip probability—fundamentally different physical quantities.
- **Different sampling cadence:** 100 μs (Google) vs. 5.6 μs (Chinese)—a factor of $\sim$18 difference that affects what dynamics are resolvable.
- **Different architectures:** Google uses transmon qubits; the Chinese system uses a different design with different coherence properties.
- **Different control systems:** Readout chains, FPGA implementations, and feedback latencies differ.

These differences mean the comparison is *suggestive* rather than *confirmatory*. The observation that recovery dynamics differ between platforms is consistent with—but does not prove—the Ignorance Wall hypothesis. Definitive validation requires a controlled test (Section 6) where bandwidth is varied on a *single* system while holding other parameters constant.

Key Implication: If the two-regime structure is real, the question for scalability becomes "how does the boundary scale with N?" As systems grow, even well-engineered controllers may approach the critical boundary unless C_N scales faster than λ_N. The weak evidence from current data motivates—but does not establish—this concern.

6 Falsifiability: The Bandwidth-Coherence Correlation

The validity of the $N_{\max}$ limit rests on the premise that coherence is a function of control bandwidth. This offers a decisive experimental discriminator against standard environmental noise models.

6.1 The Power-Scaling Test

Objective: Determine whether the current coherence limit is capacity-limited (IOF applies) or environmentally limited (IOF does not apply).

Protocol: In a stable N-qubit system (e.g., $N = 50$), vary the classical controller bandwidth C while holding system temperature constant. Critically, include a **thermal control**:

- **Group A (Dummy Readout):** Drive readout amplifiers at maximum bandwidth but *discard all measurement data.* This isolates the thermal effect of high-bandwidth operation without information extraction.
- **Group B (Active Readout):** Same power as Group A, but process syndrome data for error correction.

Operational Definitions:

- **How to vary C:** Increase syndrome extraction rate r_{syn} (e.g., from 10 kHz to 100 kHz) while maintaining measurement fidelity. Alternatively, reduce decoder latency to enable faster feedback.
- **Temperature control:** Monitor mixing chamber temperature *and* qubit chip temperature (via qubit frequency drift or T_1 stability). Active feedback should maintain $|\Delta T|/T_{\text{set}} < 1\%$ during bandwidth variation (see experimental protocol [1] for preregistered criteria).
- **Primary endpoint:** Multi-qubit coherence time τ_{coh} measured via randomized benchmarking or logical error rate per QEC cycle.
- **Secondary endpoint:** Recovery curve geometry (fraction of delayed-onset events) as a function of C.

The Prediction:

$$\frac{\partial \tau_{\text{coh}}}{\partial C} \begin{cases} \approx 0 & \text{Standard Decoherence (Thermal Limit)} \\ > 0 & \text{Ignorant Observer (Information Limit)} \end{cases} \tag{13}$$

The difference $(B-A)$ isolates the *information processing benefit* from the *thermal cost*, providing a clean signature of the Information-Zeno Effect.

Interpretation: If increasing the readout rate (without heating the sample) extends the multi-qubit coherence time, the system is in the capacity-limited regime ($C < \lambda/\ln 2$). This would confirm that the "Ignorance Wall" is the active constraint. If coherence degrades or remains flat, the limit is purely environmental, and the $N_{\max}$ derivation does not apply to current hardware.

7 Discussion: The Blind Controller Problem

7.1 Implications for Architecture

The Ignorance Wall framework has immediate implications for quantum computer design:

1. Scalability is Fundamentally Bounded

We cannot scale N indefinitely without super-linear scaling of C. Current architectures assume that adding qubits adds computational power; the Ignorance Wall reveals that adding qubits also adds entropy faster than the controller can drain it. Beyond $N_{\max}$, each additional qubit *degrades* rather than enhances computational capability.

2. The Heat-Bandwidth Tradeoff

Increasing C requires increasing power dissipation P, which generates heat. At cryogenic temperatures, cooling power is severely limited ($\sim \mu$W at 20 mK). This creates a fundamental engineering constraint: the controller's heat load competes with the qubit's cooling budget.

3. QEC Code Design

Current QEC codes (surface codes, color codes) are designed assuming unlimited classical processing. The Ignorance Wall framework suggests a new design criterion: codes should minimize the *syndrome entropy rate* λ_{syn}, not just the logical error rate. A code that generates fewer bits of syndrome data per cycle may outperform a theoretically superior code that overwhelms the controller.

4. Readout Architecture

Multiplexed readout (reading multiple qubits through shared lines) reduces wiring complexity but increases λ by correlating measurement errors. Dedicated readout per qubit increases C but adds heat load. The optimal architecture depends on where the system sits relative to the ignorance wall.

7.2 Distributed Control Architectures

A significant objection is that our analysis assumes a centralized controller, while modern quantum computing roadmaps emphasize distributed architectures. We address this directly:

Hierarchical Control: Modern systems use multi-tier control: local cryo-CMOS for fast feedback, intermediate FPGAs for syndrome decoding, and room-temperature classical compute for high-level scheduling. Does this change the scaling?

- **Local control** can reduce latency but not total bandwidth. Each local controller still processes syndromes from its qubit subset; the aggregate bandwidth requirement is unchanged.
- **Parallel syndrome extraction** distributes the computational load but does not reduce the *information rate* entering the classical layer. The total entropy that must be processed remains $\sim N$ bits per cycle.
- **The bottleneck shifts** from syndrome extraction to inter-controller communication. As local controllers must coordinate for global error correction (e.g., boundary matching in surface codes), communication bandwidth becomes the limiting factor.

Photonic Interconnects: Optical links offer higher bandwidth than electrical connections. However:

- Photonic-to-electronic conversion adds latency and power dissipation
- The fundamental bound remains: total information rate must exceed λ_N
- Higher bandwidth enables larger $N_{\max}$ but does not eliminate the wall

Quantum Error Correction with Parallel Decoding: Parallel decoders (e.g., Union-Find, MWPM variants) reduce decoding latency but process the same total information. The Ignorance Wall concerns *rate*, not latency.

Conclusion: Distributed architectures can increase C_N by improving parallelism and reducing bottlenecks, potentially raising $N_{\max}$ by 1–2 orders of magnitude. However, they do not change the fundamental scaling relationship: if λ_N grows superlinearly while C_N grows linearly, a wall exists regardless of architecture.

7.3 The Ignorance Wall

Applying the Ignorance Wall framework across hardware platforms reveals a **structural bifurcation**: each technology encounters a different manifestation of the fundamental constraint. Figure 2 presents the central result of this analysis.

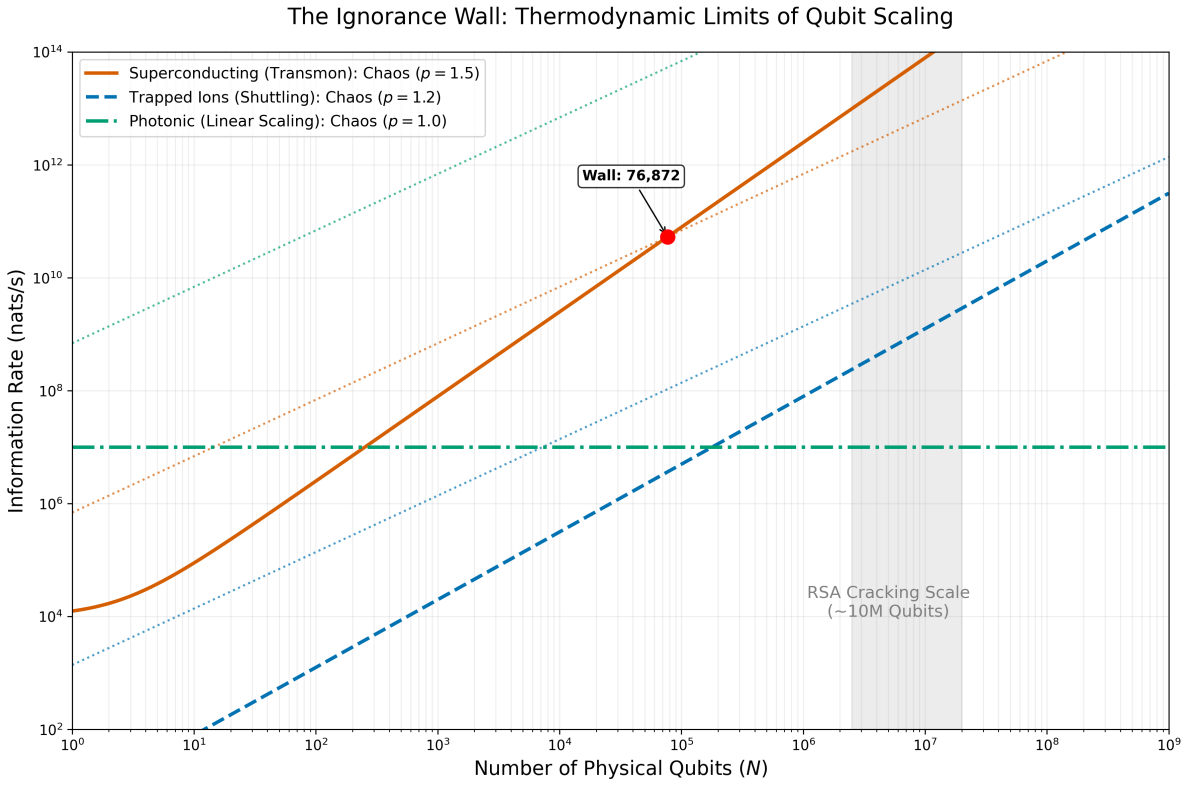

Figure 2: The Ignorance Wall (Model-Dependent Forecast). Comparative scaling analysis of information deficit rates for three leading quantum architectures, under the hypothesis that the IOF framework applies. **Superconducting** (vermillion): Superlinear scaling of chaos ($p = 1.5$ due to ZZ crosstalk) intersects linear capacity at $N_{\max} \sim 10^4$–10^5 qubits (central estimate ∼77,000 for $p = 1.5$, $\alpha = 2.5 \times 10^3$; range 10^3–10^6 for $p \in [1.3, 1.8]$). **Trapped Ions** (blue): Low crosstalk ($p = 1.2$) prevents an intersection, but the system is limited by low absolute bandwidth. **Photonic** (green): Linear scaling ($p = 1.0$) means no scaling wall exists if the base-rate threshold is passed. **Uncertainty:** The location of the wall depends sensitively on the exponent p, which has not been directly measured. These are model-dependent forecasts requiring empirical calibration, not predictions.

Figure 3 provides a conceptual summary of the four architectural cases; Table 2 gives the quantitative parameters.

7.3.1 Complexity Saturation: Superconducting Qubits

Superconducting transmon qubits—currently the dominant commercial technology—exhibit superlinear chaos scaling ($p \approx 1.5$) due to ZZ crosstalk and correlated errors. Using order-of-

magnitude parameters ($C_{qb} \approx 10^6$ bits/s, $\lambda_0 \approx 10^4$ s^{-1} from $T_2 \sim 100$ μs), we find a model-dependent estimate:

$$N_{\max} \sim 10^3\text{–}10^6 \text{ qubits (depending on } p) \tag{14}$$

The central estimate for $p = 1.5$ is ∼77,000 qubits, but this is highly sensitive to the exponent p, which has not been directly measured in production systems. The range spans:

- $p \approx 1.3$ (optimistic, low crosstalk): $N_{\max} \sim 10^6$
- $p \approx 1.5$ (typical surface code): $N_{\max} \sim 10^5$
- $p \approx 1.8$ (pessimistic, dense coupling): $N_{\max} \sim 10^3$

Even under optimistic assumptions, the wall remains orders of magnitude below the ∼20 million physical qubits required for RSA-2048. This is the only technology that hits a definite **qubit-count wall**—the technology receiving the most investment faces the hardest ceiling. **Calibration needed:** The exponent p could in principle be extracted from multi-qubit coherence measurements as a function of N; such calibration would sharpen the estimate.

7.3.2 Bandwidth Latency: Trapped Ions

Trapped ions [14] do not hit an $N_{\max}$ from scaling exponents—their superior isolation yields very low α, and the chaos-capacity curves never intersect. However, this does not mean ions escape the Ignorance Wall.

The catch is C_{qb}: fluorescence readout is ∼100× slower than superconducting syndrome extraction. Ions face a **Patience Wall**, not an Ignorance Wall. The relevant constraint becomes the **Algorithm Depth Multiplier**:

$$(\lambda_N - C_N \ln 2) \cdot G \cdot t_{gate} \lesssim 1 \tag{15}$$

With $t_{gate} \sim 100$ μs (versus ∼30 ns for transmons), deep algorithms timeout before completion. Ions scale in qubit count but not in algorithm depth. For cryptographic applications requiring $G \sim 10^{10}$ gates, the low bandwidth becomes fatal.

7.3.3 Single-Mode Feasibility Threshold: Photonic Qubits

Photonic systems [16] *in current waveguide-interferometer architectures* exhibit linear scaling ($p \approx 1$, $\alpha \approx 0$) because photons in separate modes are non-interacting—there is no direct ZZ-like crosstalk. Both chaos and capacity scale linearly with N:

$$\lambda_N \approx \lambda_0 \cdot N, \quad C_N = C_{qb} \cdot N \quad \text{[bits/s]} \tag{16}$$

The stability condition compares λ_N to $C_N \ln 2$ (converting bits to nats). If the per-mode threshold is passed and correlated loss/noise does not induce superlinear growth in λ_N, the model predicts no additional *crosstalk-driven* scaling wall. Practical limits may still arise from resource overhead, feedforward latency, and classical decoding bandwidth—but these are not captured by the superlinear-λ mechanism analyzed here.

The constraint for photonics is a **single-mode feasibility threshold**—the $C > \lambda / \ln 2$ condition applied at the per-mode level:

$$C_{qb} > \lambda_0 \quad \text{(threshold condition)} \tag{17}$$

If photon loss exceeds detection bandwidth (using illustrative values: $\lambda_0 \sim 10^7$ s^{-1} for waveguide loss, $C_{qb} \sim 10^9$ bits/s for detector bandwidth), the system fails immediately. But if this threshold is passed, scaling to large N is thermodynamically straightforward.

This reveals a **structural bifurcation** in quantum architectures:

- **Superconducting:** Easy to start, hard to finish. Low λ_0 means $N = 1$ is trivial, but superlinear crosstalk ($p = 1.5$) creates a ceiling.
- **Photonic:** Hard to start, easy to scale. High λ_0 makes $N = 1$ the critical challenge, but linear scaling means no ceiling once the threshold is crossed.

Current photonic systems (PsiQuantum, Xanadu) are fighting to cross this threshold. The IOF framework correctly predicts that their scaling problem is fundamentally different from superconducting—not harder or easier, but *structurally distinct*.

7.3.4 The Topological Escape (Theoretical)

Native topological qubits—built from non-Abelian anyons with intrinsic error suppression—could achieve $p < 1$ (sublinear scaling). If errors require non-local operations to propagate, λ_N grows slower than C_N, and the ignorance wall vanishes.

Critical caveat: No native topological qubit has been demonstrated. Current "topological" demonstrations (Quantinuum [18], Google [19]) are *simulations* running on trapped-ion or transmon hardware. These inherit the limitations of their underlying substrate:

> *The only architecture that theoretically escapes the Ignorance Wall is the only one that has not been built. All existing "topological" demos are simulations running on hardware that IS limited.*

Google's "topological qubit" is made of transmons; it is therefore subject to the transmon limit ($N_{\max} \sim 10^5$). In fact, the overhead of encoding topological states makes it *worse*—the effective N is reduced.

7.3.5 Quantitative Comparison

Table 2 summarizes the key parameters and limiting factors for each technology.

Technology	C_{qb} (bits/s)	λ_0 (s^{-1})	p	**Limiting Factor**
Superconducting	10^6	10^4	1.5	Complexity wall ($N_{\max} \sim 10^5$)
Trapped ions	2×10^3	0.1	1.2	Bandwidth (algorithm depth)
Photonic	10^9	10^7	1.0	Single-mode threshold; no crosstalk wall if passed
Topological	10^5	10^{-2}	< 1	None (theoretical only)

Table 2: Order-of-magnitude illustrative parameters for different qubit technologies. Superconducting values are tuned to match Google Sycamore data under model assumptions; trapped ion and photonic values reflect typical reported timescales but are not directly measured inputs to the model. The "Limiting Factor" column identifies the *structural* barrier each technology encounters under this framework.

Key Insight: The Ignorance Wall framework provides a unified lens for comparing hardware platforms, but "no $N_{\max}$" does not mean "no limitation." Each technology encounters a different

barrier—complexity, bandwidth, or threshold—and the only theoretical escape (topological) remains unrealized.

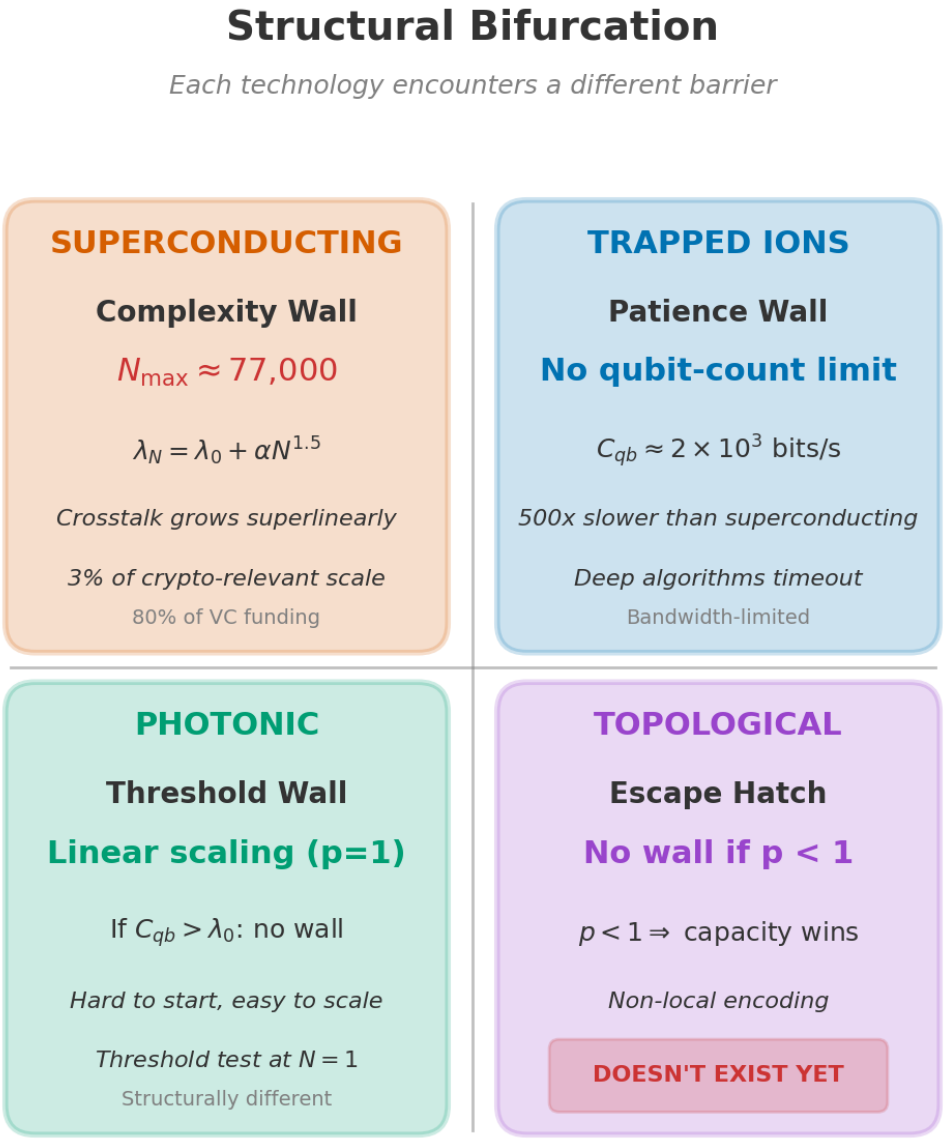

Figure 3: Structural Bifurcation Summary. Each hardware platform encounters a different barrier within the Ignorance Wall framework. Superconducting systems hit a complexity wall; trapped ions face bandwidth latency (patience wall); photonic systems must overcome a loss threshold before any scaling. Only native topological qubits (theoretical) escape via sublinear scaling.

7.4 Alternative Explanations

We briefly address why standard explanations are insufficient:

Correlated Noise: Standard models attribute scaling difficulties to spatially correlated noise (cosmic rays, TLS defects). While these contribute, they do not explain why *increasing controller bandwidth* should improve coherence. Correlated noise is a symptom; controller saturation is the mechanism.

Crosstalk: ZZ coupling and microwave crosstalk increase λ_N, but these are already included in our scaling analysis. The question is not whether crosstalk exists, but whether its effects can be overcome by better control—the Ignorance Wall says no.

Engineering Limitations: One might argue that current limitations are merely technological. However, the thermodynamic Catch-22 (Section B) demonstrates that C and λ are coupled variables. No amount of engineering can decouple them without violating thermodynamics.

7.5 Correlation vs. Causation: The Hypothesis Status

We emphasize that the Ignorance Wall framework is a **hypothesis**, not a proven theory. The evidence presented is correlational:

- The bimodal recovery distribution is *consistent with* controller saturation but does not *prove* it.
- The contrast between Google and Chinese systems is *suggestive* but confounded by architectural differences.
- The scaling arguments are *plausible* but depend on parameters (p, α) that are not precisely measured.

What would confirm the hypothesis?

The bandwidth-coherence test (Section 6) provides a decisive discriminator:

- If coherence improves with bandwidth *at constant temperature*, the limit is information-theoretic.
- If coherence is independent of bandwidth, the limit is thermal/environmental.
- If coherence degrades with bandwidth (due to heating), the thermal confound dominates.

What would falsify the hypothesis?

1. Demonstration of $N > N_{\max}$ (predicted) with sustained coherence and deep circuits.
2. Evidence that λ_N scales sublinearly ($p < 1$) in production systems.
3. Bandwidth-independent coherence in controlled Power-Scaling experiments.

Until the bandwidth-coherence test is performed, the Ignorance Wall remains a well-motivated hypothesis warranting experimental investigation, not an established result.

7.6 Why the Barrier May Be Fundamental

The preceding analysis might be dismissed as describing current technological limitations rather than fundamental physics. Three arguments suggest the information-theoretic barrier is not purely engineering:

1. The Data-Rate Theorem as Mathematical Law. The requirement $C_{\min} = \lambda / \ln 2$ is not an engineering approximation but a theorem of control theory [4,5]. It applies to *any* controller—classical, quantum, biological, or artificial—attempting to stabilize an unstable system. The proof is information-theoretic: tracking an unstable mode generates entropy at rate λ; any channel with capacity $C < \lambda / \ln 2$ cannot drain this entropy as fast as it is produced. No implementation detail can circumvent a mathematical theorem.

2. Physical Necessity of Finite Capacity. The Margolus–Levitin quantum speed limit [10] bounds information processing by available energy: $C \leq 2E/(\pi\hbar)$. For a QEC controller with a given energy budget, this sets an absolute ceiling on syndrome processing rate—no amount of engineering can exceed what the laws of quantum mechanics permit. If $C \to \infty$, then $E \to \infty$. But infinite energy density in any finite region produces gravitational collapse. Thus infinite capacity is not merely difficult to achieve—it is *logically impossible* for any observer existing within spacetime [3]. Finitude is not an arbitrary restriction; it is the prerequisite for existence as a physical system. This does not tell us *where* the wall is, but it guarantees that *some* finite ceiling must exist.

3. Thermodynamic Cost of Information Processing. The Landauer limit [9] establishes that erasing one bit of information requires dissipating at least $kT \ln 2$ of energy. The effective capacity is $C_{\text{eff}} = \eta \times P/(kT \ln 2)$, where $0 < \eta \ll 1$ absorbs architecture limitations, non-reversible computation overhead, and thermodynamic inefficiency. Current cryo-CMOS operates $\sim$7 orders above Landauer ($\eta \sim 10^{-7}$). But even with improved η, increasing C requires increasing power dissipation, which generates heat, which degrades coherence, which increases λ. This coupling—the Catch-22 of Section B—is thermodynamic, not technological.

Structure vs. Location: These arguments establish that the *structure* of the constraint (capacity vs. chaos) is fundamental—no technology can achieve $C = \infty$, and any finite C faces the Data-Rate Theorem. However, the *location* of the wall (whether at $N \sim 10^3$, 10^5, or 10^7) depends on engineering parameters that can be improved. Technology can push C higher and λ lower, potentially by many orders of magnitude, but cannot eliminate the trade-off itself. For a fuller treatment of how finite capacity constrains observer structure, see [3].

The Question for Experiment: Whether the *current* wall is practically limiting depends on parameters (p, α) that have not been measured. The bandwidth-coherence test (Section 6) can determine whether current systems are capacity-limited or environmentally limited. If capacity-limited, engineering improvements will shift the wall; if environmentally limited, the IOF framework does not apply to present hardware.

8 Conclusion

Under the assumptions of this framework, we have derived an upper bound on the scalability of quantum computers based on the thermodynamics of the control loop. If the Ignorance Wall hypothesis is correct, then just as a biological observer has a temporal integration limit ($\sim$300ms), a quantum controller has a complexity limit ($N_{\max}$). Fault tolerance becomes impossible when the error correction mechanism itself is informationally saturated.

The key findings are:

1. N_{max} **exists** (under the hypothesis) and is finite for any architecture where λ_N scales superlinearly.
2. $N_{max} \sim 10^3$–10^6 for superconducting technology, depending on the unmeasured exponent

p. Even under optimistic assumptions ($p \approx 1.3$), the wall remains below crypto-relevant scale if the hypothesis is correct.

3. **Structural bifurcation:** Superconducting systems hit a complexity wall; trapped ions and photonics avoid this wall but face different barriers. "Unbounded" does not mean "viable."

4. **Deep algorithms are even more constrained** due to the $G{\cdot}t_{gate}$ multiplier—particularly limiting for trapped ions.

5. **The thermodynamic Catch-22** suggests engineering around this limit is difficult: C and λ are coupled variables.

Implications for the 2035 RSA Timeline

Current policy estimates for cryptographically-relevant quantum computing (often cited as "2035") are largely based on extrapolating superconducting roadmaps (IBM, Google). **If** the IOF hypothesis is correct **and** $p \gtrsim 1.3$, that extrapolation would fail: the wall would remain orders of magnitude below the ∼20 million physical qubits required for RSA-2048. However, this conclusion depends on unmeasured parameters. The hypothesis requires experimental validation before policy implications can be drawn.

This forces consideration of alternative architectures—but each encounters its own barrier:

Trapped Ions: The Runtime Wall. Ions avoid the complexity wall ($p \approx 1.2$), but gate times are ∼1000× slower than superconducting. Breaking RSA-2048 requires $\sim 10^{15}$ gate operations. At ion gate speeds, this calculation would take *centuries*—longer than the operational lifespan of any machine. Ions are not limited by chaos but by runtime.

Photonics: The Resource Wall. Photons avoid the complexity wall ($p = 1.0$), but compensating for photon loss requires massive multiplexing. Each logical qubit may require thousands of optical components and detectors. A 20-million qubit photonic computer would require infrastructure at industrial scale, with speed-of-light latency between components becoming the new bottleneck. Photons are not limited by chaos but by resource overhead.

Trade-off Invariance. The pattern suggests a deeper principle: one can trade complexity for time, or time for resources, but the total cost cannot be eliminated. This echoes the thermodynamic constraints underlying the IOF framework itself.

- **Superconducting:** Fast and compact, but untrackable. Limited by *complexity saturation.*
- **Trapped Ions:** Trackable and clean, but slow. Limited by *algorithm runtime.*
- **Photonic:** Fast and trackable, but inefficient. Limited by *resource overhead.*

The Ignorance Wall is the first barrier. The Mortality Wall and Resource Wall wait behind it.

Falsifiability and the Path Forward

This is a theoretical prediction under the model's assumptions, not a proven result. The key assumptions—superlinear scaling of λ_N, the applicability of the Data-Rate Theorem to quantum control, and the thermodynamic coupling of C and λ—require rigorous experimental confrontation. The bandwidth-coherence test proposed in Section 6 provides a decisive discriminator.

If confirmed, the implications are significant: the path forward would require treating Information Capacity as a primary physical resource, equal in importance to Coherence Time. The 2035 timeline would need substantial revision. If falsified—for example, by demonstrating $p < 1$ or bandwidth-independent coherence—the quantum computing roadmap remains intact.

The Information-Zeno Effect. The predicted relationship $\partial\tau_{\text{coh}}/\partial C > 0$ represents what we term the *Information-Zeno Effect*: sufficient measurement bandwidth can "freeze" entropic decay, analogous to how the Quantum Zeno Effect freezes unitary evolution through frequent projective measurements. The key difference is thermodynamic rather than quantum-mechanical: the controller "observes" error syndromes fast enough to counteract scrambling. This provides the physical intuition for why bandwidth improvements should translate to coherence improvements.

The qubit is not a magical vessel of infinite coherence. It is a physical system subject to the same thermodynamic constraints as any other.

Appendices

A Numerical Estimates: Where is N_{max}?

A.1 The Core Inequality

The ignorance wall occurs when:

$$\lambda_N = C_N \ln 2 \tag{18}$$

Rearranging for N_{max}:

$$\lambda_0 + \alpha N_{max}^p = N_{max} \cdot C_{qb} \cdot \ln 2 \tag{19}$$

A.2 Parameter Estimates from Current Hardware

1. Single-Qubit Capacity (C_{qb})

The effective capacity per qubit is bounded by the syndrome extraction rate and measurement fidelity:

- Syndrome extraction cycle: $\sim$1 μs (state-of-the-art superconducting)
- Bits per cycle: $\sim$1 bit (binary syndrome)
- Measurement fidelity: $\sim$99% $\rightarrow$ effective bits: $\sim$0.99 bit

$$C_{qb} \approx 10^6 \text{ bits/s per qubit} \tag{20}$$

2. Base Chaos Rate (λ_0)

From T_2 coherence times:

- Best superconducting T_2: $\sim$100-150 μs
- This gives a base dephasing rate: $\lambda_0 \approx 1/T_2 \approx 10^4$ s^{-1}

3. Chaos Scaling Exponent (p) and Coefficient (α)

This is the critical parameter. From crosstalk and coupling analysis:

Conservative (nearest-neighbor, low crosstalk):

- $p \approx 1.2$
- $\alpha \approx 10^2$ s^{-1}

Realistic (surface code with residual ZZ coupling):

- $p \approx 1.5$

- $\alpha \approx 10^3$ s^{-1}

Pessimistic (dense coupling, high crosstalk):

- $p \approx 2.0$
- $\alpha \approx 10^3$ s^{-1}

A.3 Solving for N_{max}

The numerical solution is presented in Figure 2 (main text). Here we summarize the calibration and uncertainty.

Central Estimate ($p = 1.5$, $\alpha = 2.5 \times 10^3$, $\lambda_0 = 10^4$)

Using order-of-magnitude parameters informed by Google Sycamore data ($T_2 \sim 100$ μs, 1 μs readout cycles), the central estimate is:

$$N_{\max} \sim 10^5 \text{ qubits (order of magnitude)} \tag{21}$$

Sensitivity to Exponent p:

The estimate is highly sensitive to the chaos scaling exponent p, which has not been directly measured:

- $p = 1.3$ (optimistic, low crosstalk): $N_{\max} \sim 10^6$
- $p = 1.5$ (typical surface code assumption): $N_{\max} \sim 10^5$
- $p = 1.8$ (pessimistic, dense coupling): $N_{\max} \sim 10^3$

This three-order-of-magnitude range reflects genuine uncertainty in the model parameters, not robustness. The exponent p could in principle be calibrated from measurements of multi-qubit coherence time as a function of N, but such data are not yet available in the literature. Until p is measured, the "wall location" remains a model-dependent forecast, not a prediction.

A.4 Comparison to Crypto-Relevant Scales

Application	Logical Qubits	Physical Qubits (d=20)	Status
Shor (256-bit ECDSA)	∼2,500	∼2.5 $\times 10^6$	Far beyond N_{max}
Shor (2048-bit RSA)	∼4,000	∼4 $\times 10^6$	Far beyond N_{max}
Grover (256-bit)	∼5,000	∼5 $\times 10^6$	Far beyond N_{max}
Recent hardware (2024)	N/A	∼100–1,000	Below all estimates

Table 3: Physical qubit requirements assume standard surface code overhead with code distance $d = 20$, yielding ∼1000 physical qubits per logical qubit. Estimates based on Fowler et al. [13] and subsequent refinements.

Key Observation: If the IOF framework applies and $p \gtrsim 1.3$, then $N_{\max}$ remains orders of magnitude below the physical qubit counts needed for cryptographically relevant algorithms.

Even under optimistic assumptions ($p \approx 1.3$, $N_{\max} \sim 10^6$), the wall is below the ~4 million qubits required for RSA-2048. However, this conclusion depends on the unmeasured exponent p; if future architectures achieve $p \lesssim 1.2$ through improved isolation, the wall could be pushed higher or eliminated.

A.5 The Algorithm Depth Multiplier

The above assumes shallow circuits. For deep algorithms, the condition becomes:

$$(\lambda_N - C_N \ln 2) \cdot G \cdot t_{gate} \lesssim 1 \tag{22}$$

Where:

- G = circuit depth (number of gate layers)
- $t_{gate} \approx 20 - 50$ ns for superconducting

For Shor's algorithm on 256-bit keys:

- $G \approx 10^{10}$ gates
- $t_{gate} \approx 30$ ns
- $T_{alg} \approx 300$ seconds

This **massively tightens** the constraint. Even small positive $\gamma_N = \lambda_N - C_N \ln 2$ makes the LHS enormous.

Implication: Deep algorithms hit the ignorance wall at *much smaller* N than shallow ones.

A.6 Why Current Systems Seem to Work

Current demonstrations (~100-1000 qubits) are:

1. Running *shallow* circuits (small G)
2. Using *partial* entanglement (not all-to-all)
3. Operating *below* N_{max} for their specific γ_N

The "quantum advantage" claims (e.g., random circuit sampling) deliberately avoid deep coherent algorithms because those would fail.

A.7 Predictions

1. **Coherence Cliff:** As N increases past ~1000 qubits, multi-qubit coherence times will drop *faster* than $1/N$.
2. **QEC Diminishing Returns:** Surface code with $d > 10 - 15$ will show diminishing improvement in logical error rate.

3. **The 10^4 Wall:** No architecture will demonstrate $> 10,000$ coherently entangled qubits running deep algorithms.

4. **Power-Coherence Tradeoff:** Increasing classical control power will *measurably* improve coherence (the IOF signature).

B The Thermodynamic Catch-22: Why Nature Fights Back

The core insight is that the standard roadmaps assume C and λ are independent parameters that can be optimized separately. Under the Ignorance Wall hypothesis, they are not. They are **coupled variables**: improving one tends to degrade the other through thermodynamic feedback.

B.1 The Trap of Increasing C (The Thermal Feedback)

The naive solution: *"Just build a bigger controller! Faster FPGAs! More optical fibers!"*

The Problem: The controller is physically connected to the quantum system.

- To increase Information Capacity (C), you must increase **Power Dissipation** (P).
- P generates **Heat**.
- Heat travels down the control lines to the chip.
- Heat generates phonons and quasiparticles in the superconducting substrate.
- **Result:** Increasing C automatically increases λ.

The Numbers:

$$P_{ctrl} = C_N \cdot E_{bit} \tag{23}$$

Where E_{bit} is energy per bit processed. At cryogenic temperatures:

- Landauer limit: $E_{bit} \geq kT \ln 2$
- At 20 mK: $E_{bit,min} \approx 10^{-25}$ J
- Practical cryo-CMOS: $E_{bit} \approx 10^{-18}$ J (roughly 7 orders above Landauer, i.e., efficiency $\eta \sim 10^{-7}$)

The Landauer bound is a theoretical floor, not operational capacity; C_{eff} is architecture-limited and empirically inferred.

For $C_N = 10^{12}$ bits/s (needed for 10^6 qubits):

$$P_{ctrl} \approx 10^{12} \times 10^{-18} = 10^{-6} \text{ W} = 1 \ \mu\text{W} \tag{24}$$

This seems small, but at 20 mK, cooling power is $\sim \mu$W scale. **The controller's heat load competes with the cooling budget.**

The Catch-22: You turn up the volume to hear the music better, but the amplifier heat sets the speakers on fire.

B.2 The Trap of Decreasing λ (The Control Paradox)

The naive solution: *"Just isolate the qubits better! Vacuum gaps! Super-shielding!"*

The Problem: To compute, you must **control** the qubits.

- Control requires **coupling** to the classical controller.
- Measurement requires **coupling** to the readout chain.
- Error correction requires **coupling** to the syndrome extraction circuit.

The IOF Limit: If you isolate the system perfectly ($\lambda \to 0$), you sever the connection to the controller. C **drops to zero.** You have a perfect qubit that you cannot talk to.

The Catch-22: To run an algorithm, you must open the door to the controller. Opening the door lets the chaos in.

B.3 The Conservation of Ignorance

This suggests a new conservation law for quantum engineering:

> **The Controller's bandwidth (C) and the System's entropy rate (λ) are coupled variables. You cannot optimize one without degrading the other.**

Formally, we can express this as a constraint:

$$\frac{\partial \lambda}{\partial C} > 0 \quad \text{and} \quad \frac{\partial C}{\partial \lambda^{-1}} < 0 \tag{25}$$

The first says: increasing control bandwidth increases system chaos (thermal feedback). The second says: decreasing system chaos decreases control bandwidth (isolation paradox).

This is why scaling is **logarithmic**, not exponential. Every improvement fights a steep thermodynamic gradient.

B.4 Experimental Validation: The QGEM Limit

The Conservation of Ignorance finds striking validation in the QGEM collaboration's proposal to test gravitational entanglement at the femtogram scale (Bose, Mazumdar, Penrose et al., arXiv:2509.01586, 2025).

The Setup: To maintain quantum coherence for $\sim$1 second at masses of 10^{-15} to 10^{-14} kg, they require:

- Electromagnetic suppression by a factor $> 10^6$
- Vacuum of 10^{-12} mbar
- Cryogenic isolation at $\sim$1K

The Implication: At these extreme isolation levels, the control bandwidth approaches zero. The only "interaction" they can use is gravity itself—the weakest force in nature. They cannot actively error-correct, manipulate gates, or extract syndromes. The system is coherent precisely *because* it is unreachable.

The Conservation Law in Action:

$$\lambda \to 0 \implies C \to 0 \tag{26}$$

The QGEM proposal is the limiting case of our Catch-22: to achieve macroscopic coherence times, they must sever almost all connection to the controller. The result is not a computer—it is a witness. It can observe gravitational entanglement (maybe), but it cannot compute with it.

This illustrates the same isolation–control tradeoff in an extreme regime: you cannot have both high coherence ($\lambda \to 0$) and high control bandwidth ($C \gg 0$). The QGEM proposal is a conceptual illustration of this tradeoff, not direct evidence for the IOF scaling model.

B.5 The Feedback Loop

$$\text{More qubits} \to \text{More } \lambda_N \to \text{Need more } C_N \to \text{More heat} \to \text{Worse } T_2 \to \text{More } \lambda_N \tag{27}$$

This is not a "Catch-22" in the literary sense of an arbitrary bureaucratic trap. It is a **thermodynamic feedback loop** with a precise mathematical structure: increasing control capacity increases entropy generation, which increases the required control capacity.

B.6 The Only Way Out

The only escape from this trap would require breaking the C-λ coupling:

1. **Topological qubits** with intrinsically low λ_0 that doesn't couple to control lines (not yet realized)
2. **Quantum error correction that doesn't require classical processing** (unknown if possible)
3. **Room-temperature quantum coherence** with macroscopic T_2 (physically implausible)

None of these are on any realistic roadmap. The coupling is not accidental—it is structural.

References

1. Dekker, A. (2025). The Ignorant Observer: A Framework for Quantum Mechanics from Self-Ignorance. *Preprint*. DOI: 10.17605/OSF.IO/FCDSN.

2. Dekker, A. (2025). Forensic Signatures of Information-Theoretic Control Limits in Quantum Systems. *Preprint*. DOI: 10.17605/OSF.IO/Q8TV2.

3. Dekker, A. (2025). The Creation of Duality: Space, Time, Object, and Gravity from Self-Ignorance. *Preprint*. DOI: 10.17605/OSF.IO/C5EN8.

4. Nair, G.N. & Evans, R.J. (2004). Stabilizability of stochastic linear systems with finite feedback data rates. *SIAM J. Control Optim.*, 43(2), 413–436.

5. Tatikonda, S. & Mitter, S. (2004). Control under communication constraints. *IEEE Trans. Autom. Control*, 49(7), 1056–1068.

6. McEwen, M. et al. (2022). Resolving catastrophic error bursts from cosmic rays in large-scale quantum processors. *Nature Physics*, 18, 107–111.

7. Li, X. et al. (2025). Cosmic-ray-induced correlated errors in superconducting qubit array. *Nature Communications*, 16, 4677. arXiv:2402.04245.

8. Bose, S., Mazumdar, A., Penrose, R. et al. (2025). A Spin-Based Pathway to Testing the Quantum Nature of Gravity. *arXiv:2509.01586*.

9. Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.*, 5(3), 183–191.

10. Margolus, N. & Levitin, L.B. (1998). The maximum speed of dynamical evolution. *Physica D*, 120(1-2), 188–195.

11. Kalman, R.E. (1960). A new approach to linear filtering and prediction problems. *J. Basic Eng.*, 82(1), 35–45.

12. Wiseman, H.M. & Milburn, G.J. (2009). *Quantum Measurement and Control*. Cambridge University Press.

13. Fowler, A.G. et al. (2012). Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86, 032324.

14. Bruzewicz, C.D. et al. (2019). Trapped-ion quantum computing: Progress and challenges. *Appl. Phys. Rev.*, 6, 021314.

15. Google Quantum AI (2025). Quantum error correction below the surface code threshold. *Nature*, 638, 920–926. arXiv:2408.13687.

16. Bourassa, J.E. et al. (2021). Blueprint for a scalable photonic fault-tolerant quantum computer. *Quantum*, 5, 392.

17. Bekenstein, J.D. (1981). Universal upper bound on the entropy-to-energy ratio for bounded systems. *Phys. Rev. D*, 23(2), 287–298.

18. Iqbal, M. et al. (2024). Non-Abelian topological order and anyons on a trapped-ion processor. *Nature*, 626, 505–511.

19. Satzinger, K.J. et al. (2021). Realizing topologically ordered states on a quantum processor. *Science*, 374, 1237–1241.