

Retrospective Forensic Signatures in Transient Recovery Dynamics: A Screening Protocol for Bandwidth-Limited Control Hypotheses

Aernoud Dekker

May 2026
Version 1.4

DOI: 10.17605/OSF.IO/M5TB9

Abstract

We present a model-agnostic screening protocol for detecting recovery transients compatible with Bandwidth-Limited Quantum Control (BLQC)-style dynamics. BLQC, the finite-rate basis-tracking module of IOF, motivates double-exponential visibility decay $V(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$, which linearizes as $\ln(-\ln V) \propto t$. The protocol tests this against the power-law alternative $\ln(-\ln V) \propto \ln t$ using a dual R^2 statistic.

We apply this protocol to three precision instrument platforms: Chinese 63-qubit processor cosmic ray events, Google Sycamore qubit cosmic ray recovery, and LIGO gravitational wave detector glitches. Results show: (1) Chinese data shows power-law dominance with mean $\Delta R^2 = -0.21$ and 0.0–0.1% joint Gompertz-consistent events; (2) Google shows power-law dominance with mean $\Delta R^2 = -0.12$ and 0.0% joint Gompertz-consistent events; (3) LIGO shows 43% Gompertz-consistent events, and among well-fit events ($R_t^2 \geq 0.90$), 88% prefer t -linear over $\ln t$ -linear scaling. A public H1 AUX_CLN cross-check finds no corrected-significant BLQC-positive versus BLQC-negative label separation in available auxiliary channels, but a controller-regime reanalysis finds that LIGO's fitted κ has a strong raw association with ASC recovery-stress ranks (Spearman $\rho = 0.123$, BH $q \approx 5.2 \times 10^{-6}$), which is absorbed by morphology controls. Capacity-sweep simulations show the exact expected BLQC signal shape: decreasing effective controller capacity increases κ , advances the visibility break, and moves traces into the t -linear double-log regime. These results do not establish causal attribution to BLQC, but they support treating LIGO as a serious candidate regime and motivate the proposed controlled-capacity experiment.

Status of this document. This paper is a retrospective forensic analysis, not a validation of the Ignorant Observer Framework (IOF) or BLQC as a physical mechanism. IOF is the full framework; BLQC is the technical finite-rate basis-tracking module used here as the target model. The results should be read as a screening exercise: they ask whether existing recovery traces contain the double-log geometry predicted by a bandwidth-limited tracking model. The public H1 auxiliary-channel cross-check reported below narrows the present claim: it does not identify a corrected-significant BLQC-specific label separator in the available public auxiliary data. The newer controller-regime analysis is nevertheless supportive rather than contradictory: it shows that fitted BLQC rate κ is linked to public controller-family recovery proxies, especially ASC, before ordinary glitch morphology is controlled. Attribution to BLQC would require prospective intervention, especially controlled variation of effective controller capacity C_{eff} while holding thermal and plant-level confounds fixed.

Dependency on BLQC. This screening protocol assumes the BLQC finite-rate basis-tracking hypothesis as its target model. Compatible retrospective signatures are not standalone proof of the framework; their scientific weight depends on whether the bandwidth-dependent visibility law is confirmed prospectively.

1 Introduction

Bandwidth-Limited Quantum Control (BLQC), the technical module of IOF, motivates a specific visibility decay law for recovery transients in precision instruments:

$$V(t) = \exp\left(-\frac{\sigma_0^2}{2}e^{2\kappa t}\right), \quad \kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2 \quad (1)$$

where κ is the proposed “ignorance rate”—the net rate at which controller state uncertainty grows when system entropy production (h_{KS}) exceeds effective channel capacity (C_{eff}).

This double-exponential form has a testable signature: taking double logarithms yields

$$\ln(-\ln V) = \ln\left(\frac{\sigma_0^2}{2}\right) + 2\kappa t \quad (2)$$

If the BLQC-style visibility law applies, $\ln(-\ln V)$ versus t is linear. The alternative—power-law variance growth $\sigma^2(t) \propto t^\alpha$ —yields $\ln(-\ln V) \propto \ln t$ instead.

This paper presents a model-agnostic protocol for discriminating between these two fitted growth laws and applies it to recovery transients from three precision instrument platforms: superconducting qubits (Chinese 63-qubit and Google Sycamore processors) and LIGO gravitational wave detector glitches. Because all datasets are retrospective and were collected for other purposes, the protocol can identify compatibility or incompatibility with the proposed signature, but it cannot by itself establish physical cause.

2 Rationale

Why should recovery transients follow the BLQC visibility law? The argument has three steps.

2.1 A Perturbation Creates a State-Estimation Problem

After a disturbance, the controller does not simply “return to baseline.” It must reconstruct the relevant system state (or error state) from measurements and then apply corrective actions. Recovery is a closed-loop inference-plus-actuation process. Even when the underlying plant is deterministic, the controller operates with a belief distribution over possible states because it has finite sensing resolution, finite bandwidth, noise, latency, and model mismatch.

2.2 Uncorrected State Uncertainty Expands Exponentially

For many high-dimensional, nonlinear systems, small state errors amplify roughly as $\delta x(t) \sim \delta x(0)e^{h_{\text{KS}}t}$, where h_{KS} is an effective entropy rate (a Kolmogorov–Sinai-like rate). This is the control-theory version of positive Lyapunov exponents: if you do not track the state finely enough, prediction error grows exponentially. In terms of error variance,

$$\sigma^2(t) \propto \sigma_0^2 e^{2h_{\text{KS}}t} \quad (3)$$

in the absence of sufficient corrective information.

2.3 The Controller Can Only Remove Uncertainty at a Finite Rate

A real controller has limited effective information throughput C_{eff} (bits/s): finite sensor bandwidth, finite compute, finite telemetry rate, finite actuator update rate. In the BLQC framing, this means the best you can do is reduce uncertainty growth by a rate proportional to $C_{\text{eff}} \ln 2$ (nats/s). When the plant’s internal entropy production exceeds this, the net ignorance rate

$$\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2 \quad (4)$$

is positive and the controller’s state uncertainty still grows exponentially:

$$\sigma^2(t) = \sigma_0^2 e^{2\kappa t} \quad (5)$$

2.4 Why Visibility Becomes Double-Exponential

In many measurement settings, the observable “visibility” is an overlap or coherence-like quantity that decays with the variance of an underlying phase or error variable. A generic and widely used form is

$$V(t) = \exp\left(-\frac{1}{2}\sigma^2(t)\right) \quad (6)$$

which is the standard Gaussian-overlap relationship: as uncertainty over a phase/state spreads, the overlap with the initial coherent state decreases.

Substituting Eq. (5) gives

$$V(t) = \exp\left(-\frac{\sigma_0^2}{2}e^{2\kappa t}\right) \quad (7)$$

—double-exponential decay.

This form has a distinctive “flat-then-steep” geometry: early on, $e^{2\kappa t}$ is modest, so $V(t)$ stays near 1 (apparent plateau). Once uncertainty has compounded sufficiently, the exponential in the exponent dominates and $V(t)$ drops rapidly. This is the pattern expected in a model where a controller initially maintains an adequate state estimate, but its internal estimate eventually loses the race against instability and finite information throughput.

2.5 Domain of Validity

This prediction applies when the recovery observable is dominated by controller state-tracking failure rather than by passive multi-channel relaxation in the plant. If the transient is governed by distributed relaxation times or mixture-of-modes physics, one should expect stretched exponentials or sums of exponentials instead. This is not a contradiction; it is a domain-of-validity statement about when recovery is *observer-limited* versus *plant-limited*.

The dual R^2 test presented in this paper discriminates precisely this: $\ln(-\ln V) \propto t$ (BLQC, observer-limited) versus $\ln(-\ln V) \propto \ln t$ (power-law, typically plant-limited relaxation).

3 Protocol

The BLQC detection protocol consists of six sequential stages. Signal preprocessing and eligibility gates are platform-adapted because the raw observables have different polarity, sampling, and noise structure. The invariant part of the protocol is the post-eligibility construction of visibility and the dual- R^2 test.

3.1 Step 0: Signal Preprocessing

Raw instrument data is converted to a normalized decay signal $E(t)$:

1. **Bandpass filtering** (if applicable): Remove out-of-band noise while preserving transient dynamics.
2. **Envelope extraction**: For oscillatory signals (e.g., LIGO strain), compute the Hilbert envelope $E(t) = |H[s(t)]|$.
3. **Peak identification**: Locate the primary peak within ± 500 ms of the reported event time, constrained to avoid secondary features.
4. **Time centering**: Define $t = 0$ at peak location; all subsequent analysis uses peak-relative time.

3.2 Step 1: Eligibility Filtering

Events must pass phenomenological criteria with *no BLQC assumptions*. For the LIGO envelope pipeline, the eligibility gate is:

1. **Single dominant pulse**: Secondary peaks $< 80\%$ of primary amplitude.
2. **Monotonic approach to baseline**: Theil-Sen slope checks for a consistent trend (increasing or decreasing) toward baseline in the analysis window.
3. **Sufficient dynamic range**: $|\text{peak} - \text{baseline}|/\text{baseline} > 0.5$ (excursion amplitude).
4. **Signal-to-noise**: Peak amplitude exceeds the platform threshold above baseline noise (LIGO: 5σ ; qubit recovery data: 3σ).
5. **Baseline stability**: Coefficient of variation in tail region < 0.5 .

Events failing any criterion are excluded from BLQC testing but may be analyzed separately. For the Google and Chinese frozen recovery datasets, the runners use the corresponding polarity-adapted amplitude, SNR, recovery-trend, and finite-data checks before applying the same onset, visibility, and dual- R^2 stages.

3.3 Step 2: Onset Time Estimation

The BLQC signature depends critically on correct time anchoring. We estimate the onset time t_{inf} using a 3-window model fitting approach:

1. For each window duration $w \in \{w_1, w_2, w_3\}$:
 - (a) Extract post-peak data: $t \in [0, w]$
 - (b) Compute baseline from window tail: $b = \text{median}(E[-n/5 :])$
 - (c) Transform to recovery coordinate: $z = 1 - (E - b)/(E_{\text{peak}} - b)$
 - (d) Fit sigmoid model: $z(t) = 1/(1 + e^{-k(t-t_0)})$
 - (e) Fit delayed-exponential model: $z(t) = \begin{cases} 0 & t < d \\ 1 - e^{-(t-d)/\tau} & t \geq d \end{cases}$
 - (f) Select winner by AICc; extract $t_{\text{inf}} = t_0$ or $t_{\text{inf}} = d$
2. Average t_{inf} across windows with valid fits; if exactly one valid onset estimate remains, use that estimate, and if no valid onset estimate is available, fall back to peak anchoring.

The onset time shifts the analysis origin: $t_{\text{onset}} = t_{\text{peak}} + t_{\text{inf}}$.

Note on circularity. The use of sigmoid and delayed-exponential models for onset estimation does *not* bias the subsequent BLQC test. These models operate in the z -coordinate to identify a single time shift; they do not impose Gompertz or double-exponential structure on the visibility. The onset estimate merely aligns the time axis so that the dual R^2 test is not penalized by arbitrary peak-centering. The BLQC hypothesis is tested independently in Step 4 via t -linearity of $\ln(-\ln V)$.

3.4 Step 3: Visibility Construction

Construct the normalized visibility signal for dual R^2 testing:

1. Anchor time to onset: $t' = t - t_{\text{onset}}$
2. Extract analysis window: $t' \in [0, T_{\text{fit}}]$
3. Compute baseline from window tail: $b = \text{median}(E[-n/5 :])$
4. Construct visibility:

$$V(t') = \frac{E(t') - b}{E_{\text{peak}} - b} \quad (8)$$

where E_{peak} is the envelope value at the original peak time (before onset shifting). This ensures consistent normalization independent of the onset estimate. As a robustness check, we also tested $V_{\text{onset}} = (E(t') - b)/(E(t' = 0) - b)$ using the first post-onset sample; results were qualitatively unchanged.

5. Apply validity bounds: retain only $t' \geq t_{\text{min}}$ and $V \in (0.01, 0.99)$

3.5 Step 4: Dual R^2 Test

The core BLQC signature test compares two linearizations:

1. Compute double-log transform: $y = \ln(-\ln V)$
2. **Test 1 (BLQC hypothesis):** Linear regression of y versus t'
 - Obtain R_t^2 (coefficient of determination)
 - Extract $\kappa = \text{slope}/2$
3. **Test 2 (Power-law hypothesis):** Linear regression of y versus $\ln t'$
 - Obtain $R_{\ln t}^2$
4. Compute discriminant: $\Delta R^2 = R_t^2 - R_{\ln t}^2$
 - $\Delta R^2 > 0$: Exponential variance growth (Gompertz-consistent)
 - $\Delta R^2 < 0$: Power-law variance growth
 - $\Delta R^2 \approx 0$: Indeterminate

Interpretation of κ . The slope-derived quantity $\kappa = \text{slope}/2$ is interpretable as an effective ignorance rate only for Gompertz-consistent events ($R_t^2 \geq 0.90$ and $\Delta R^2 > 0$). For events with $\Delta R^2 < 0$, the t -linear fit is misspecified and κ is merely a regression coefficient without physical meaning.

Algorithm: Dual R^2 Test

Input: Time array t , visibility array V , minimum time t_{\min}

1. Filter: keep indices where $t \geq t_{\min}$ and $0.01 < V < 0.99$
2. Compute $y \leftarrow \ln(-\ln V)$
3. Fit: $y = a_1 + b_1 \cdot t$ (BLQC model)
4. Compute R_t^2 from residuals
5. Fit: $y = a_2 + b_2 \cdot \ln t$ (Power-law model)
6. Compute $R_{\ln t}^2$ from residuals
7. *Return:* $R_t^2, R_{\ln t}^2, \Delta R^2 = R_t^2 - R_{\ln t}^2, \kappa = b_1/2$

3.6 Step 5: Classification

Events are classified as **Gompertz-consistent** (exhibiting the t -linear double-log signature) if they satisfy *both*:

1. $R_t^2 \geq 0.90$ (the double-log transform is well-linearized by *some* monotone function)
2. $\Delta R^2 > 0$ (exponential variance growth favored over power-law)

The R_t^2 criterion is essential: events where neither model fits well (low R_t^2) have ΔR^2 dominated by noise. Restricting to $R_t^2 \geq 0.90$ ensures we only interpret ΔR^2 for events where the linearization is meaningful.

As a robustness check, we also report results at a strict threshold ($\Delta R^2 \geq 0.05$).

Null calibration (pipeline-matched). We validated classification thresholds via 4,000 synthetic null transients processed through the *exact same pipeline* as real data, including three-window onset estimation. This addresses the “look-elsewhere effect” critique that onset selection could artificially inflate linearity.

Null models included two-exponential mixtures, stretched exponentials ($\beta < 1$), and adversarial nulls: compressed stretched exponentials ($\beta > 1$) and saturated/clipped transients. Under the null:

- Only 3.3% achieved $R_t^2 \geq 0.90$
- Of those, only 6.1% had $\Delta R^2 > 0$
- Overall false-positive rate: 0.2%

This suggests that the joint criterion ($R_t^2 \geq 0.90$ AND $\Delta R^2 > 0$) has strong specificity against the synthetic non-Gompertz transients considered here. It does not exhaust all possible instrumental or physical confounds.

4 Capacity-Sweep Simulation

The retrospective datasets can show whether BLQC-compatible recovery geometry is present, but they cannot by themselves show what the operational control variable should be. To clarify the expected signal, we ran a capacity-sweep simulation that asks a sharper question: does changing effective controller capacity move a recovery trace into or out of the t -linear double-log regime?

Two simulations were used.

1. **Ideal finite-rate bound.** The visibility law is generated directly from the finite-rate uncertainty balance

$$\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2. \quad (9)$$

When $C_{\text{eff}} \ln 2 < h_{\text{KS}}$, variance grows exponentially and $V(t) = \exp(-\frac{1}{2}\sigma^2(t))$ produces a straight line in $\ln(-\ln V)$ versus t . When $C_{\text{eff}} \ln 2 \geq h_{\text{KS}}$, the controller-capacity condition is no longer violated and the BLQC-positive regime should disappear.

2. **Packetized servo analogue.** An ensemble of residual phase errors expands under an unstable plant and is corrected only when enough channel bits accumulate to send a fixed-size quantized correction packet. Visibility is computed from the simulated ensemble variance; the Gompertz curve is not imposed. This deliberately crude model tests whether finite-bit feedback can generate BLQC-positive traces without hard-coding the target curve.

The simulation shows that the exact target shape is not merely a delayed or Gompertz-looking curve. The target is ordinary-time linearity after the double-log transform:

$$\ln(-\ln V(t)) = \text{const} + 2\kappa t. \quad (10)$$

In the ideal finite-rate model, the dual- R^2 test recovers the imposed κ exactly within numerical precision. As capacity approaches the threshold $\rho = C_{\text{eff}} \ln 2 / h_{\text{KS}} = 1$, the fitted κ decreases and the visibility break time moves later. Above threshold, no BLQC-positive trace is expected. The packetized-servo analogue reproduces the qualitative capacity response but with biased κ estimates and distorted threshold behavior, illustrating what real controllers may do: they can generate the same diagnostic regime without producing a perfect textbook curve.

BLQC operational signature: capacity moves traces into or out of the double-log regime

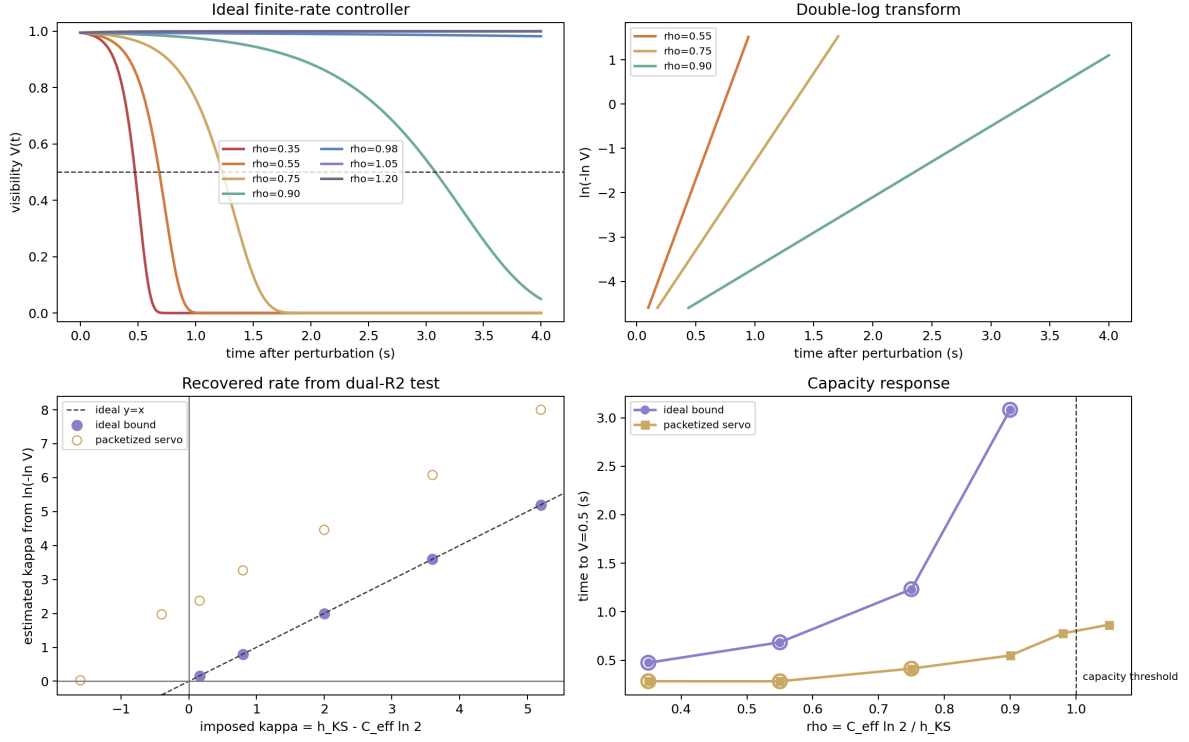


Figure 1: Capacity-sweep simulation. Top-left: ideal finite-rate visibility traces for several capacity ratios $\rho = C_{\text{eff}} \ln 2 / h_{\text{KS}}$. Top-right: the BLQC-positive traces are linear in ordinary time after the double-log transform. Bottom-left: the dual- R^2 test recovers the imposed κ for the ideal bound; the packetized-servo analogue produces the same qualitative regime with biased rates. Bottom-right: increasing effective capacity delays the $V = 0.5$ break and eventually moves the system out of the BLQC-positive regime.

This simulation does not prove that LIGO glitches are BLQC-generated. It specifies what the proposed experiment must test. If BLQC is the right operational account, controlled reduction of C_{eff} should increase fitted κ , shift visibility breaks earlier, and increase the fraction of traces satisfying the dual- R^2 criterion. Controlled increase of C_{eff} should do the opposite. This capacity dependence is the operational bridge that retrospective datasets alone cannot supply.

5 Platform-Specific Adaptations

The protocol parameters are adapted for each platform's characteristic time scales and signal format:

Parameter	Chinese	Google	LIGO
Sampling rate	179 kHz	10 MHz	4096 Hz
Signal type	Probability	Error counts	Hilbert envelope
Signal direction	Decays to baseline	Rises to baseline	Decays to baseline
Fit window T_{fit}	0.5 ms	150 μs	60 ms
Onset windows (post-peak)	[0.2, 0.3, 0.5] ms	[60, 100, 150] μs	[60, 100, 150] ms
Minimum time t_{min}	0.01 ms	0.05 μs	0.5 ms

Table 1: Protocol parameters by platform.

Chinese. Probability data decays from peak to baseline. No bandpass filtering required.

Google. Error count data shows a *dip* during cosmic ray impact, then *rises* back to baseline. Visibility is inverted: $V = (b - y)/(b - y_{\text{min}})$ where y is error count and b is baseline.

LIGO. Raw strain data undergoes 10–500 Hz bandpass filtering and Hilbert transform to extract the amplitude envelope. The envelope decays from peak toward baseline.

6 Results

6.1 Chinese 63-Qubit Processor

We analyzed cosmic ray events from the 63-qubit superconducting processor.

SI_Fig8a (Charge-Parity). 78 events.

- Raw t -linear preference ($\Delta R^2 > 0$): 3.8%
- **Gompertz-consistent** ($R_t^2 \geq 0.90$ AND $\Delta R^2 > 0$): 0.0%
- Mean $\Delta R^2 = -0.220 \pm 0.102$

SI_Fig12a (Bitflip). 1437 events (of 2137 total).

- Well-fit ($R_t^2 \geq 0.90$): 111 events (7.7%)
- Raw t -linear preference ($\Delta R^2 > 0$): 0.1%
- **Gompertz-consistent** ($R_t^2 \geq 0.90$ AND $\Delta R^2 > 0$): 0.1% (1 event)
- Mean $\Delta R^2 = -0.205 \pm 0.093$

Interpretation. Chinese data shows strong power-law dominance ($\Delta R^2 \ll 0$). The $\ln t$ -linear model consistently outperforms t -linear, indicating variance growth inconsistent with BLQC exponential dynamics.




Figure 2: Chinese 63-qubit SI_Fig8a (charge-parity) analysis. Left: ΔR^2 distribution showing 3.8% raw $\Delta R^2 > 0$ but 0.0% joint BLQC-positive events. Center: Joint R_t^2 vs ΔR^2 scatter (red = joint Gompertz-consistent) with no events satisfying the joint criterion. Right: Model shape distribution (secondary).




Figure 3: Chinese 63-qubit SI_Fig12a (bitflip) analysis. Left: ΔR^2 distribution showing only 0.1% raw $\Delta R^2 > 0$ and 0.1% joint BLQC-positive events. Center: Joint R_t^2 vs ΔR^2 scatter (red = joint Gompertz-consistent). Right: Model shape distribution (secondary).

6.2 Google Sycamore

We analyzed cosmic ray recovery events from the 26-qubit Sycamore processor.

Dataset. 277 events; 276 passed eligibility filtering with valid dual R^2 .

BLQC Signature Test.

- Raw t -linear preference ($\Delta R^2 > 0$): 4.0%
- **Gompertz-consistent** ($R_t^2 \geq 0.90$ AND $\Delta R^2 > 0$): 0.0%
- Strict joint threshold ($R_t^2 \geq 0.90$ AND $\Delta R^2 \geq 0.05$): 0.0%
- Mean $\Delta R^2 = -0.116 \pm 0.062$

Stability Class Comparison.

- `stable_iof` (delayed geometry): 0% raw $\Delta R^2 > 0$, 0% joint BLQC-positive, mean $\Delta R^2 = -0.142$
- `stable_std` (fast geometry): 4.4% raw $\Delta R^2 > 0$, 0% joint BLQC-positive, mean $\Delta R^2 = -0.101$

Interpretation. Google shows power-law dominance similar to Chinese data, with no BLQC signatures in either geometry class.

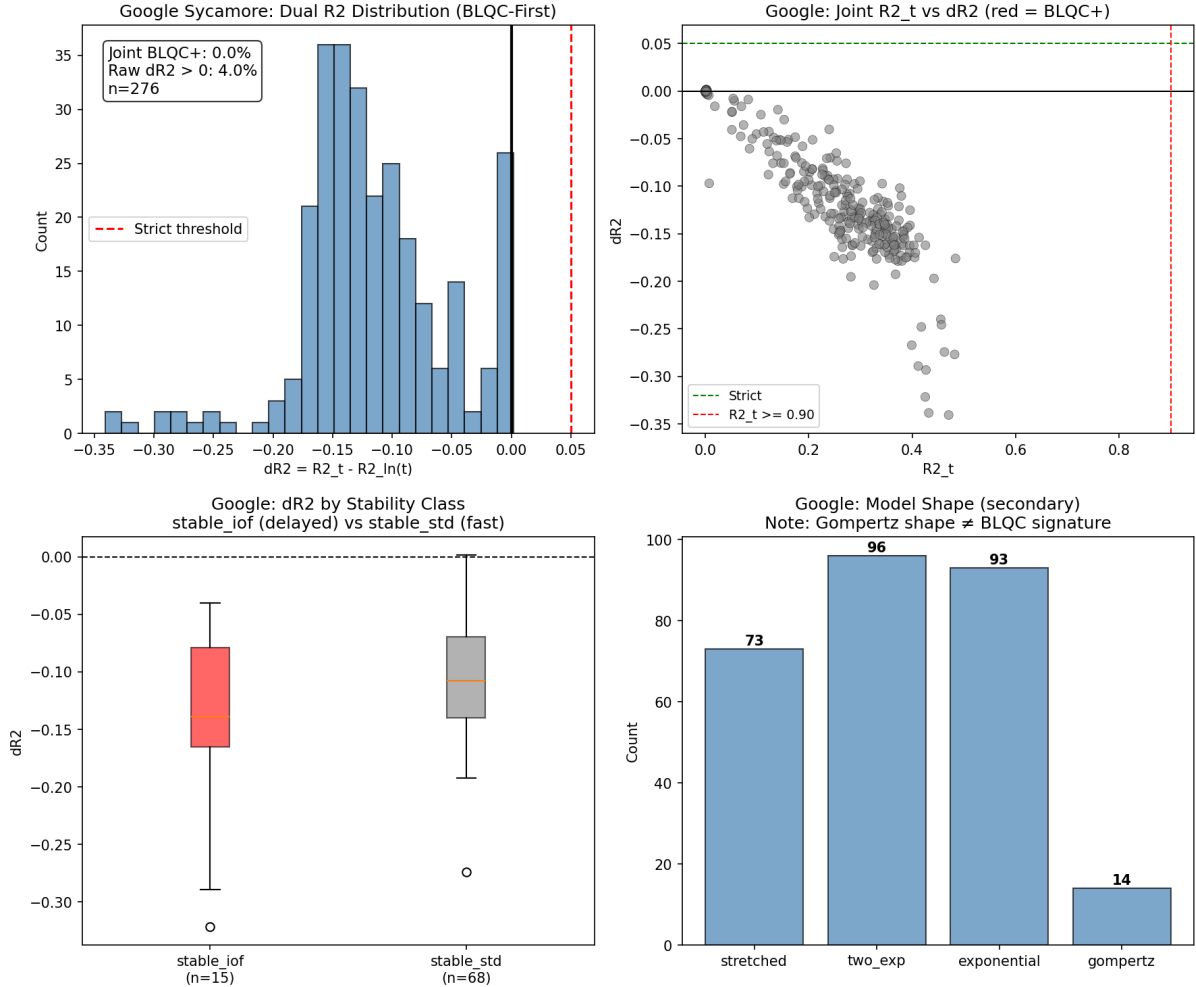


Figure 4: Google Sycamore analysis. Top-left: ΔR^2 distribution showing 4.0% raw $\Delta R^2 > 0$ but 0.0% joint BLQC-positive events. Top-right: Joint R_t^2 vs ΔR^2 scatter (red = joint Gompertz-consistent). Bottom-left: ΔR^2 by stability class. Bottom-right: Model shape distribution (secondary).

6.3 LIGO H1 Glitches

We analyzed “Extremely Loud” glitches from LIGO H1 during the O3a observing run.

Dataset. 6914 events with $\text{SNR} \geq 50$; 5268 passed eligibility filtering (76.2%).

BLQC Signature Test.

- Valid dual R^2 : 5263 events
- Well-fit ($R_t^2 \geq 0.90$): 2566 events (48.8%)
- Of well-fit events, $\Delta R^2 > 0$: 2256 (87.9%)
- **Gompertz-consistent** ($R_t^2 \geq 0.90$ AND $\Delta R^2 > 0$): 42.9% of all valid events
- Strict joint threshold ($R_t^2 \geq 0.90$ AND $\Delta R^2 \geq 0.05$): 38.3% of all valid events

Note on the raw ΔR^2 distribution. The overall ΔR^2 distribution has mean -0.011 and is centered near zero with a 50/50 split above/below. However, this includes events where neither model fits well ($R_t^2 < 0.90$), for which ΔR^2 is dominated by noise due to collinearity between t and $\ln t$ on the 60 ms analysis window (correlation ≈ 0.90). The meaningful comparison is among well-fit events, where the signal strongly favors t -linearity.

Interpretation. Among LIGO glitches where the double-log transform linearizes well ($R_t^2 \geq 0.90$), 88% prefer t -linear over $\ln t$ -linear. This is qualitatively different from qubit platforms, where even well-fit events prefer $\ln t$ -linear. The asymmetry is compatible with a difference in underlying variance-growth dynamics, but retrospective data alone cannot determine whether the cause is BLQC, another exponential-error mechanism, or an instrumental artifact.

Public auxiliary-channel cross-check. To test whether the LIGO strain-envelope signature can be tied directly to controller-channel activity, we joined the BLQC classifications to public H1 O3a AUX_CLN frame data where available. The public AUX_CLN coverage intersects 2048 of the 5263 valid LIGO events: 863 Gompertz-consistent and 1185 non-Gompertz-consistent. Of these, 762 are strict Gompertz-consistent events. For each event we extracted ± 10 s windows from 16 public auxiliary channels spanning alignment (ASC), length-control (LSC), seismic-isolation (HPI), and suspension-control (SUS) readouts, and computed robust peak- z , post-peak RMS, post/pre RMS ratio, and post-peak impulse features.

The first AUX analysis asked a simple label-separation question: do BLQC-positive events have larger public auxiliary excursions? The answer is no at corrected significance. The auxiliary-channel join produced 32,672 usable channel windows out of 32,768 requested. Across event-level AUX summaries, BLQC-positive events show only a very small upward shift: rank-AUC values range from 0.508 to 0.513, with no event-level metric significant after Benjamini-Hochberg correction ($q = 0.514$ across the tested metrics). Channel-level tests likewise show no corrected significant separation; the smallest channel-level corrected value is $q = 0.779$. The strongest raw AUX correlations are instead with catalog-level glitch duration and SNR.

We then reframed the question in the form suggested by the BLQC mechanism: not “do BLQC-positive events have bigger AUX spikes?”, but “are fitted BLQC rates related to controller-regime proxies?” Each AUX metric was converted to a within-channel percentile rank, then aggregated by event across all channels and by controller family (ASC, LSC, SUS, HPI). These features estimate pre-event control burden, post-event recovery stress, and combined controller stress without comparing raw channel amplitudes across channels. This produced 360 controller-regime features over 2042 events with usable AUX data.

This controller-regime analysis finds a supportive but non-decisive signal. The strongest raw association is between fitted κ and ASC recovery-stress rank:

$$\rho_{\text{Spearman}} = 0.123, \quad p = 2.47 \times 10^{-8}, \quad q \approx 5.2 \times 10^{-6}. \quad (11)$$

The association is unchanged by time control and remains present after controlling for SNR and duration ($\rho \approx 0.091$, corrected $q \approx 8.9 \times 10^{-4}$). However, it is absorbed by morphology

controls (duration, peak frequency, and bandwidth; $\rho \approx 0.003$) and by the full control set. Label separation remains weak: the largest BLQC-positive versus negative separation is an HPI pre-burden feature with AUC = 0.462 and corrected $q = 0.49$.

The evidential status is therefore more precise than a simple negative or positive result. The public AUX data does not supply a corrected-significant BLQC-specific label separator, but it does show that the fitted BLQC rate is linked to controller-family recovery stress before ordinary glitch morphology is regressed out. This supports further investigation and does not contradict BLQC; it also shows why the prospective experiment must vary effective controller capacity directly instead of relying on retrospective public auxiliary channels.

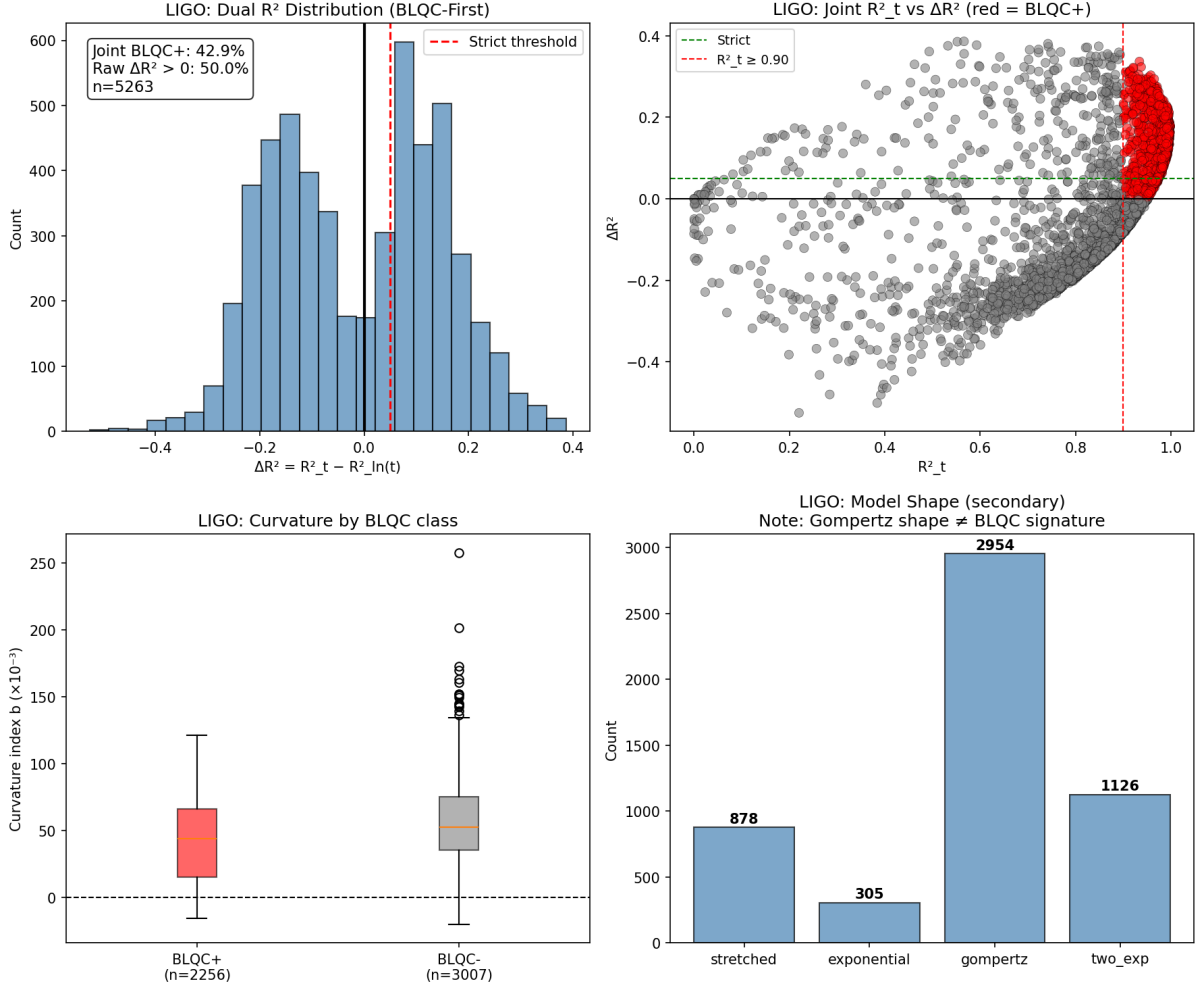


Figure 5: LIGO H1 analysis. Top-left: ΔR^2 distribution centered near zero. Top-right: Joint R_t^2 vs ΔR^2 scatter showing 88% of well-fit events ($R_t^2 \geq 0.90$, right of vertical line) have $\Delta R^2 > 0$ (red points in upper-right quadrant). Bottom-left: Curvature by BLQC class. Bottom-right: Model shape distribution (secondary).

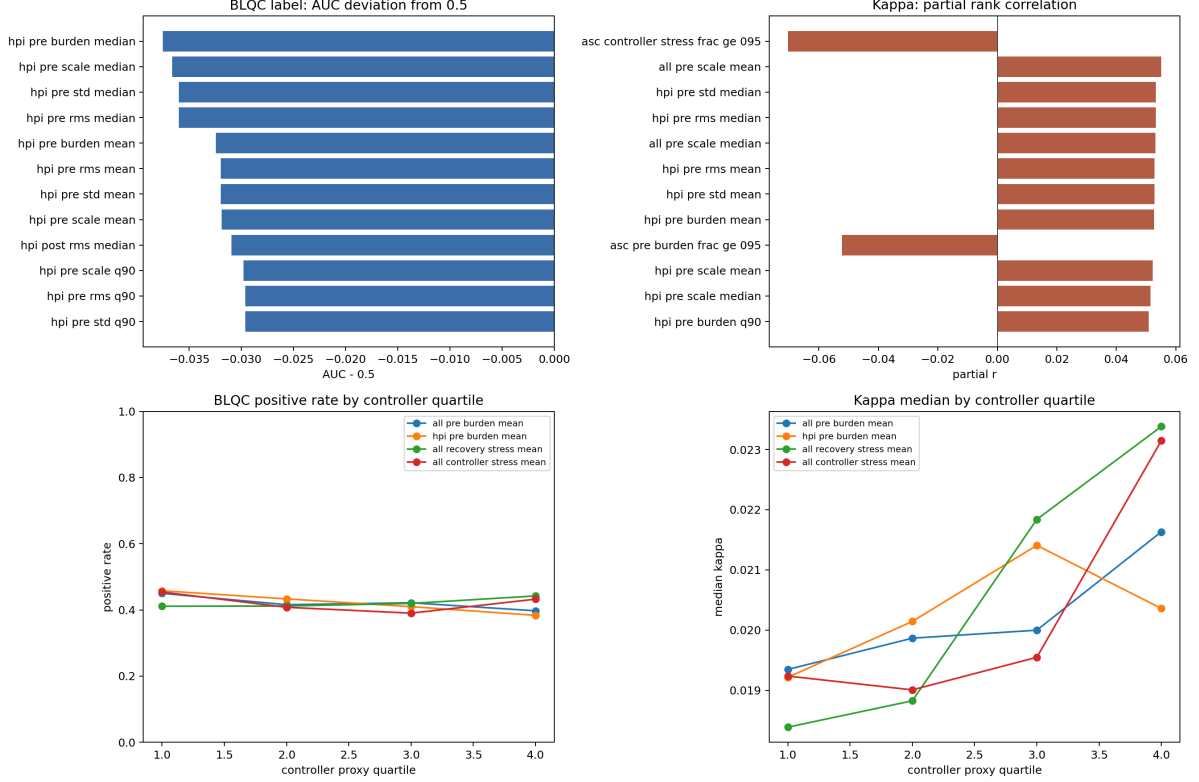


Figure 6: Public H1 AUX_CLN controller-regime diagnostics. Top-left: BLQC-positive versus BLQC-negative label separation is weak after multiple-testing correction. Top-right: the largest partial κ correlations after full controls are small. Bottom-left: BLQC-positive rate changes little across broad controller-proxy quartiles. Bottom-right: median fitted κ increases across several recovery-stress quartiles before morphology controls are applied, motivating a controlled-capacity experiment rather than retrospective attribution.

7 Cross-Platform Comparison

Platform	N	$R_t^2 \geq 0.90$ (%)	Of those, $\Delta R^2 > 0$ (%)	Gompertz-consistent (%)
Chinese (Fig8a)	78	0.0	–	0.0
Chinese (Fig12a)	1437	7.7	0.9	0.1
Google Sycamore	276	0.0	–	0.0
LIGO H1	5263	48.8	87.9	42.9

Table 2: Gompertz-consistent scaling test across platforms. Events are counted as Gompertz-consistent only when both $R_t^2 \geq 0.90$ and $\Delta R^2 > 0$ hold. For LIGO, among events where the double-log transform linearizes well, 88% prefer t -linear. Qubit platforms show uniformly negative mean ΔR^2 and essentially no joint-positive population.

The cross-platform pattern reveals an asymmetric screening result:

- **Chinese:** Strong power-law preference ($\Delta R^2 \approx -0.21$); the $\ln t$ -linear model dominates.
- **Google:** Power-law preference ($\Delta R^2 \approx -0.12$); variance growth inconsistent with Gompertz scaling.

- **LIGO**: Among events where the test has discriminatory power ($R_t^2 \geq 0.90$), 88% prefer t -linear—a candidate signal for Gompertz-consistent scaling. Public auxiliary-channel analysis does not yield a corrected-significant BLQC label separator, but fitted κ is raw-correlated with ASC recovery-stress ranks, suggesting that the strain-channel signature is at least coupled to controller-regime observables.

8 Discussion

8.1 Delayed Geometry versus BLQC Signature

An important distinction must be made between *delayed geometry* and *BLQC signature*. Delayed geometry—a flat-then-accelerating decay profile—can arise from multiple physical mechanisms: thermal relaxation with multiple time constants, saturation effects, nonlinear dynamics, or BLQC chaos-wins dynamics. The Gompertz function, which captures this shape, appears throughout nature in contexts unrelated to bandwidth-limited control.

The dual R^2 test specifically discriminates exponential variance growth ($\sigma^2 \propto e^{2\kappa t}$) from a power-law alternative ($\sigma^2 \propto t^\alpha$) within the chosen visibility construction. A delayed recovery shape with $\Delta R^2 < 0$ indicates that the shape is not captured by the BLQC-style growth law. A positive ΔR^2 is a necessary compatibility check for the proposed signature, not confirmation of the BLQC mechanism.

The Google data illustrates this: events classified as “delayed geometry” (`stable_iof`) show $\Delta R^2 = -0.14$, firmly negative. The delayed shape is present, but the BLQC signature is absent. This demonstrates that phenomenological geometry classification cannot substitute for the dual R^2 test.

8.2 Interpretation of Cross-Platform Results

Both quantum computing platforms (Chinese and Google) show negative mean ΔR^2 (-0.21 and -0.12 respectively), indicating power-law variance growth dominates. The $\ln t$ -linear model consistently outperforms t -linear.

LIGO shows qualitatively different behavior. While the raw ΔR^2 distribution is centered near zero (due to collinearity between t and $\ln t$ making ΔR^2 noisy for poorly-fit events), the meaningful comparison is among well-fit events. For events with $R_t^2 \geq 0.90$, 88% prefer t -linear—an asymmetry absent in qubit data. This suggests that the LIGO strain-channel subset deserves targeted follow-up: it is Gompertz-consistent in this protocol, while the qubit datasets are power-law dominated.

The public AUX_CLN cross-check adds a useful qualification. It does not demonstrate that BLQC-positive labels are uniquely identifiable from public auxiliary channels. However, the controller-regime reanalysis shows that fitted κ is raw-correlated with ASC recovery stress. That is the direction expected if recovery geometry is connected to finite-bandwidth control, but the dependence is not yet separable from ordinary glitch morphology. The result therefore supports further investigation without overstating attribution.

8.3 Specificity of the BLQC Signature

The BLQC prediction— $\ln(-\ln V) \propto t$ —is specific at the level of the fitted growth law. It requires visibility to follow $V(t) = \exp(-\frac{1}{2}\sigma^2(t))$ with variance growing *exponentially*: $\sigma^2(t) = \sigma_0^2 e^{2\kappa t}$. Power-law variance growth ($\sigma^2 \propto t^\alpha$) yields $\ln(-\ln V) \propto \ln t$ instead. Logarithmic or linear variance growth produce different signatures entirely.

The presence of Gompertz-consistent events (88% of well-fit LIGO glitches) warrants consideration of alternative explanations. Several non-BLQC mechanisms can produce Gompertz-like decay shapes; however, the dual R^2 test specifically targets the variance *growth law* (t -linear

vs $\ln t$ -linear), not merely the decay profile. Candidate confounds to investigate in future work include:

- **Multiplicative noise:** Geometric Brownian motion produces exponential variance growth, though its physical basis in interferometer glitches requires investigation.
- **Lyapunov instability:** Lyapunov instability can generate exponential growth of state-estimation error. If the measured recovery observable is an overlap-like function of that error variance, then a BLQC-like double-log linearity can appear even without invoking an explicit bandwidth-limited controller model. BLQC’s added content is the prediction that the effective growth rate depends systematically on the controller’s information capacity (C_{eff}).
- **Measurement artifacts:** Saturation or nonlinear detector response could potentially mimic the signature.

Distinguishing these alternatives from BLQC would require controlled experiments varying system parameters. The dual R^2 test provides a necessary but not sufficient condition for BLQC; positive events should be treated as candidates for follow-up. The capacity-sweep simulation clarifies what follow-up must look for: κ and the visibility-break time should move systematically with effective controller capacity, not merely with uncontrolled glitch morphology.

8.4 Limitations

1. **Onset time sensitivity:** The BLQC signature depends on correct time anchoring. Systematic errors in t_{inf} estimation could bias ΔR^2 .
2. **Visibility bounds:** The restriction $V \in (0.01, 0.99)$ excludes early and late dynamics where the double-log transform is numerically unstable.
3. **Retrospective analysis:** All datasets analyzed here were collected for other purposes. Dedicated experiments with controlled C_{eff} variation would provide stronger tests.
4. **Auxiliary-channel specificity:** The public H1 AUX_CLN cross-check does not show a corrected significant BLQC-positive versus BLQC-negative label separation. The controller-regime reanalysis does show a raw κ -ASC recovery-stress association, but that association is absorbed by morphology controls. This limits present attribution while still motivating targeted capacity-controlled tests.

8.5 Prospective Test

The decisive operational test of BLQC would be intervention: systematically varying C_{eff} while holding h_{KS} and thermal conditions fixed should shift ΔR^2 predictably. Decreasing controller bandwidth should increase fitted κ , move the visibility break earlier, and increase the fraction of Gompertz-consistent traces. Increasing controller bandwidth should reduce fitted κ and, in the model, suppress the Gompertz-consistent regime. Candidate experiments include:

- Varying LIGO servo loop gain during controlled glitch injection
- Adjusting qubit feedback controller bandwidth
- Comparing identical qubits with different readout architectures

9 Conclusion

We present a model-agnostic protocol for testing Gompertz-consistent scaling in transient recovery dynamics. The dual R^2 test discriminates between double-exponential visibility decay ($\ln(-\ln V) \propto t$) and power-law decay ($\ln(-\ln V) \propto \ln t$). Events are classified as Gompertz-consistent only if $R_t^2 \geq 0.90$ (the linearization is well-fit) *and* $\Delta R^2 > 0$ (exponential variance growth preferred).

Applied to three precision instrument platforms:

1. Chinese 63-qubit data: 0.0–0.1% joint Gompertz-consistent, mean $\Delta R^2 = -0.21$
2. Google qubit recovery: 0.0% joint Gompertz-consistent, mean $\Delta R^2 = -0.12$
3. LIGO glitches: 43% Gompertz-consistent; among well-fit events ($R_t^2 \geq 0.90$), 88% prefer t -linear

Both qubit platforms show power-law variance growth ($\Delta R^2 \ll 0$). LIGO shows a strong preference for t -linear scaling among events where the test has discriminatory power. This asymmetry—88% t -linear in LIGO vs. uniformly $\ln t$ -linear in qubits—is compatible with, but not uniquely explained by, the BLQC finite-rate basis-tracking model.

A public H1 AUX_CLN cross-check does not yet supply the missing causal bridge. In 2048 covered LIGO events, simple auxiliary-channel amplitude features do not significantly distinguish BLQC-positive from BLQC-negative events after correction for multiple comparisons, and the strongest raw AUX structure tracks ordinary glitch duration and SNR. A controller-regime reanalysis sharpens the picture: fitted κ is strongly raw-correlated with ASC recovery-stress ranks, but the association is absorbed by morphology controls. The current evidential status is therefore precise: the LIGO strain-envelope result is a strong phenomenological anomaly relative to the qubit platforms; public auxiliary data provides a suggestive controller-regime link but not unique BLQC attribution.

Importantly, even events with “delayed geometry” decay profiles can show negative ΔR^2 , demonstrating that phenomenological decay shape does not imply Gompertz-consistent dynamics. The dual R^2 test is essential for distinguishing variance growth laws. Capacity-sweep simulations show what an operational BLQC signal should do: lower effective controller capacity increases κ and advances the visibility break, while higher capacity suppresses the regime. Definitive attribution to BLQC would require controlled experiments varying C_{eff} .

Data Availability

LIGO data from the Gravitational Wave Open Science Center (<https://gwosc.org>). Google data from Figshare DOI <https://doi.org/10.6084/m9.figshare.16673851>. Chinese data from Figshare DOI <https://doi.org/10.6084/m9.figshare.28815434>. Analysis code available at <https://osf.io/m5tb9>.

Acknowledgments

This work uses data from the LIGO Scientific Collaboration, Google Quantum AI, and the Chinese Academy of Sciences quantum computing group.

References

Dekker, A. (2025). Bandwidth-Limited Quantum Control: A Control-Theoretic Framework for Measurement and Decoherence. *OSF Preprints*. <https://doi.org/10.17605/OSF.IO/G5WRH>

- Li, Y., et al. (2025). Cosmic-ray-induced correlated errors in superconducting qubit array. *Nature Communications*, 16, 4677.
- McEwen, M., et al. (2022). Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits. *Nature Physics*, 18, 107–111.