

A Mathematical Theory of Correlated Error Accumulation in Quantum Circuits

RamaKrishna Pasupuleti

workisfun415@gmail.com

Kakatiya University <https://orcid.org/0009-0008-8418-1430>

Research Article

Keywords: correlated error accumulation, regular variation, Euler–Maclaurin, NLS estimation, Bayes risk, stability, universality, quantum circuit diagnostics

Posted Date: August 11th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10654302/v1>

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

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

A Mathematical Theory of Correlated Error Accumulation in Quantum Circuits

Ramakrishna Pasupuleti

Independent Researcher, Suryapet, Telangana 508213, India

ORCID: 0009-0008-8418-1430 | workisfun415@gmail.com

Abstract

We develop a rigorous multi-level mathematical framework for correlated error accumulation in quantum circuits, inspired by the layered theoretical development of randomised benchmarking. Three explicit axioms (A1–A3) with measurable violation conditions replace ad hoc assumptions. The Correlated Error Accumulation Theorem (CEAT) proves that for $C(k) \sim k^{-\beta}$, the accumulation operator $T^n(C) = \sum C(|i-j|)$ satisfies $T^n(C) \sim n^\alpha$ with $\alpha = \max(1, 2-\beta)$, proved via Euler–Maclaurin expansion. A Stability Theorem with three cases (additive, multiplicative, and truncated perturbations) establishes robustness; a Universality Theorem defines the result precisely through the theory of regular variation. The perturbation bound $|C^{\text{GHZ}} - \phi^2/4| \leq \phi^4/48$ (proved via alternating series) and five explicit validity conditions B1–B5 connect the theorem to circuit infidelity. All five NLS regularity conditions are proved directly for the K–R model; consistency (Jennrich), asymptotic normality (White), and finite-n upward bias $O(n^{\beta-1})$ are established and Monte Carlo–verified. A Sample Complexity Theorem gives the minimum qubit count $N(\beta, \epsilon)$. Bayes-optimal regime boundaries are derived from simulation-calibrated distributions, yielding $\alpha^* = 1.57$ at $n = 2-10$. Appendix D characterises four failure modes for Axiom A2. Four figures verify all major results.

Keywords: correlated error accumulation, regular variation, Euler–Maclaurin, NLS estimation, Bayes risk, stability, universality, quantum circuit diagnostics

1. Introduction

1.1 *Why existing benchmarking does not expose correlation structure*

Randomised benchmarking (RB) [1,2] extracts a single exponential decay rate from Clifford circuit survival probability. The decay rate estimates the average depolarising error per gate under an approximately Markovian noise model. By construction, RB averages over all circuit instances and all error sources, collapsing the result to a scalar. This averaging is the source of its power — and its limitation: it cannot detect whether gate errors are correlated across a circuit, nor does it expose how infidelity grows as a function of circuit complexity. A device with strongly correlated (long-range) errors may produce the same RB rate as a device with independent errors but will behave very differently when circuits are scaled to larger qubit counts.

Quantum volume [10] and CLOPS [21] address whole-system performance but aggregate across all error sources into a single metric that grows slowly with improvements. Direct RB [4] and mirror circuits [8] extend benchmarking to multi-qubit settings but retain the averaging structure that obscures correlation. None of these methods answer the question: *how rapidly does infidelity grow as a function of qubit count, and what does the growth rate reveal about the noise mechanism?*

1.2 *Why a scaling exponent is useful*

If infidelity scales as n^α , the exponent α carries mechanistic information unavailable from a single-point measurement. A device with $\alpha=1$ has errors that accumulate linearly — consistent with independent gate errors. A device with $\alpha=1.5$ has errors that accumulate faster than linear, indicating correlations that span multiple gates (long-range). A device with $\alpha \approx 1.0$ after a recalibration event differs qualitatively from a device with $\alpha \approx 3.8$: the same T^2 coherence time can accompany either, making T^2 uninformative about the correlation structure. The present work shows, under stated assumptions, that α estimates the correlation decay exponent $\beta=2-\alpha$, connecting an easily measured observable to a physically meaningful noise property.

Importantly, the statistical properties of the α estimator (consistency, asymptotic normality, finite- n bias) are derived here from first principles, not just asserted. This distinguishes the present treatment from prior phenomenological uses of power-law fitting. Whether Axioms A1–A3 hold for a given device is an empirical question; the statistical theory holds conditional on those axioms, and the conditions can be tested.

1.3 *Why GHZ is the appropriate probe*

An n -qubit GHZ state $(|0\dots 0\rangle + |1\dots 1\rangle)/\sqrt{2}$ is prepared by a Hadamard on qubit 0 followed by $n-1$ CNOT gates in a cascade. Each CNOT couples the current qubit to the growing entangled state, so phase errors on successive gates contribute additively to a single global phase $\varphi = \sum \delta^k$. The infidelity of the output GHZ state equals $\sin^2(\varphi/2)$, which depends only on this total accumulated phase. This makes GHZ circuits a direct probe of cumulative correlated error: the scaling of infidelity with n reflects the structure of the sum $\sum \delta^k$, which is precisely $T^n(C)/4$ under Axiom A3.

Alternative circuit families (Clifford depth scaling, Hahn echo sequences) probe different aspects of the noise spectrum. The CEAT applies to any circuit family where infidelity is an accumulated covariance, but GHZ circuits have the cleanest analytical connection because the cascade structure makes the additive phase model exact.

1.4 Summary of Main Results and Assumptions

The paper proves four main theorems. Theorem 4.1 (CEAT): for any correlation sequence $C(k) \sim k^{-\beta}$, the accumulated double sum $T_n(C)$ scales as n^α with $\alpha = \max(1, 2-\beta)$, proved via Euler–Maclaurin expansion of generalised harmonic numbers (no quantum assumptions required). Theorem 5.2 (Stability): additive perturbations with $\gamma > \beta$, multiplicative slowly-varying factors, and finite truncations all leave α unchanged. Theorem 5.3 (Universality): the exponent depends only on the regular variation class of $C(k)$; quantum circuits are one application of a general analytical result. Theorem 6.1 (Perturbation bound): $|C_{\text{GHZ}}(\varphi) - \varphi^2/4| \leq \varphi^4/48$, proved via alternating series. All five NLS regularity conditions are proved directly for the K–R model $f(n; \theta) = C n^\alpha + R$; consistency and asymptotic normality follow from Jennrich (1969) and White (1980); the finite- n upward bias is $O(n^{\beta-1})$. The Bayes-optimal regime boundary at $n=2-10$ is $\alpha^*=1.57$, not the asymptotic value 1.0. Main assumptions: Axioms A1 (stationarity of gate errors), A2 (power-law correlation tail), and A3 (infidelity below 0.5). Whether A1–A3 hold for a specific device is an empirical question addressed in the companion experimental paper [6]. The K–R parameterisation $C n^\alpha + R$ is an application of the CEAT, not a consequence of it; this paper derives what the exponent α means, not the model itself.

2. Related Work

2.1 Benchmarking methods

RB [1,2], interleaved RB [3], direct RB [4], mirror circuits [8], cross-entropy benchmarking [9], and quantum volume [10] each measure different aggregate properties of circuit performance. Harper et al. [11] showed that noise can exhibit power-law structure detectable from circuit data. Fogarty et al. [5] demonstrated non-exponential RB decay under low-frequency noise. None provide the scaling-exponent analysis developed here.

2.2 Noise spectroscopy and 1/f noise [13,14,15]

Bylander et al. [14] characterised 1/f noise [13,14,15] in superconducting qubits by noise spectroscopy. Paladino et al. [15] reviewed the theoretical framework. The 1/f model motivates Axiom A2 but is not required; the CEAT holds for any correlation with the prescribed tail.

2.3 Power-law scaling in statistical physics

Long-range dependence in time series is a classical subject [19,20,22,23,28,29,30]. A stationary process whose covariance satisfies $C(k) \sim k^{-\beta}$ with $0 < \beta < 1$ is called *long-range dependent* [22,23]; its partial sums exhibit non-standard scaling. Fractional Brownian motion [24] provides a Gaussian model with this property. The Hurst exponent $H = 1 - \beta/2$ characterises the memory of such processes [20]. Regular variation theory [20] formalises the class of functions with power-law tails; the Karamata Tauberian theorem [20,25] relates spectral singularities to tail behaviour. Our CEAT is an application of these classical results to accumulated circuit errors, situating quantum hardware diagnostics within a well-established mathematical tradition.

2.4 The K–R diagnostic framework

Pasupuleti [6] introduced the K–R fitting model $C(d) = C^0 \cdot n^\alpha + R$ empirically. The present paper provides the theoretical foundation: α estimates $2 - \beta$ under stated conditions, with quantified finite- n bias.

3. Mathematical Background and Axioms

$\mathbb{N} = \{1, 2, \dots\}$. For sequences f, g : $f(n) \sim g(n)$ means $f/g \rightarrow 1$ as $n \rightarrow \infty$. $H(n, s) = \sum k^{-s}$ is the generalised harmonic number; $\zeta(s)$ is the Riemann zeta function; $\gamma^M \approx 0.5772$ is the Euler–Mascheroni constant. A function $L: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is *slowly varying* if $L(\lambda x)/L(x) \rightarrow 1$ for all $\lambda > 0$. A function $f(x) \sim x^\alpha \cdot L(x)$ for a slowly varying L is *regularly varying* with index α [20].

A1 (Stationarity). $\{\delta_i\}$ is zero-mean and weakly stationary: $E[\delta_i \delta_{i+k}] = C(k)$ depends only on k . Testable: $CV_{\alpha} \geq 8\%$ signals violation.

A2 (Regular variation). $C(k) \sim L \cdot k^{-\beta} \cdot \Lambda(k)$ as $k \rightarrow \infty$, where Λ is slowly varying and $L > 0$, $\beta > 0$.
Special cases: A2a: $\Lambda = 1$, $0 < \beta < 1$ (long-range dependent); A2b: $\sum |C(k)| < \infty$ (summable, $\beta > 1$); A2c: $C(k) \sim L/k$ (marginal). This is a gate-level hypothesis; deriving it from bath dynamics is the Level 2 Programme (Section 10).

A3 (Perturbative regime). $E[\varphi^2] = T_n(C) < \pi^2/4$. All measured infidelities below 0.5. Ensures $|C_{\text{GHZ}} - \varphi^2/4| \leq \varphi^4/48$ (Theorem 4.1).

4. The Correlated Error Accumulation Theorem

4.1 Definition and Lemmas

Definition 4.1. The accumulation operator $T_n(C) = \sum_i \sum_j C(|i-j|) = n \cdot C(0) + 2 \sum_{k=1}^n (n-k) C(k)$.

Lemma 4.1 (Euler–Maclaurin for $H(n,s)$). $H(n,s) = n^s \{1-s\} / (1-s) + \zeta(s) + O(n^{-s})$ for $s \neq 1$, $s > 0$;
 $H(n,1) = \ln n + \gamma^M + O(n^{-1})$. [Abramowitz–Stegun §23.2, ref. [19]]

Lemma 4.2 (Double sum decomposition). $S(n,\beta) = \sum_{k=1}^n (n-k) C(k) = n \cdot H(n-1,\beta) - H(n-1,\beta-1)$.

Lemma 4.3 (Regular variation extension). If $C(k) \sim L \cdot k^{-\beta} \cdot \Lambda(k)$ with Λ slowly varying, then $T_n(C) \sim T_n(L \cdot k^{-\beta})$ up to multiplicative slowly varying corrections. This follows from Karamata's theorem [20, §1.5] applied to the partial sums.

4.2 Main Theorem

Theorem 4.1 (CEAT). Under Assumption A ($C(k) \sim L k^{-\beta}$), as $n \rightarrow \infty$:

$$T_n(C) \sim 2L \cdot n^{\{2-\beta\}} / [(1-\beta)(2-\beta)] \quad (0 < \beta < 1) \quad (1a)$$

$$T_n(C) \sim 2L \cdot n \cdot \ln n \quad (\beta = 1) \quad (1b)$$

$$T_n(C) \sim n \cdot [C(0) + 2 \sum C(k)] \quad (\beta > 1) \quad (1c)$$

Equivalently: $T_n(C) \sim n^\alpha$ where $\alpha = \max(1, 2-\beta)$. (Proof: Appendix A.1.)

4.3 Corollary: Regime Classification

Corollary 4.1. (i) $\alpha > 1 \Leftrightarrow$ long-range dependent errors ($\beta < 1$, non-summable). (ii) $\alpha = 1 \Leftrightarrow$ summable errors ($\beta > 1$ or exponential tail). (iii) The boundary $\alpha = 1$ separates summable from non-summable correlations.

Note: $\beta > 1$ does NOT imply Markovian dynamics; it implies summable correlations. Many non-Markovian processes (e.g., Ornstein–Uhlenbeck below their correlation timescale) have summable correlations.

4.4 Interpretation

The exponent α measures how rapidly gate errors fail to decorrelate across the circuit. Under A2a ($\beta < 1$), the error on gate i remains correlated with gates arbitrarily far ahead; errors across the circuit do not average out but instead reinforce each other, producing super-linear infidelity growth. The quantity $2 - \beta = \alpha$ is therefore a direct measure of the strength of this long-range memory: the smaller β , the longer the memory, the larger α . One empirical model consistent with $T_n(C) \sim n^\alpha$ is the K–R parameterisation $C n^\alpha + R$; it is an application of the CEAT, not a derivation of it.

5. Two-Term Expansion, Stability, and Universality

5.1 Two-Term Expansion

Theorem 5.1 (Remainder). $T_n(C) = A \cdot n^{2-\beta} + B \cdot n + O(n^{1-\beta})$, where $A = 2L/[(1-\beta)(2-\beta)]$ and $B = L \cdot [1 + 2\zeta(\beta)]$. Accuracy: $< 0.01\%$ at $n \geq 1000$ for all $\beta \in (0, 1)$. Near $\beta = 1$, $\zeta(\beta) \rightarrow \infty$ causing slow convergence.

5.2 Stability Theorem

Theorem 5.2 (Stability — three cases).

(i) Additive perturbation: $C(k) = Lk^{-\beta} + \varepsilon(k)$, $|\varepsilon(k)| \leq ck^{-\gamma}$. If $\gamma > \beta$: $\alpha_{\text{eff}} = 2 - \beta + O(n^{\beta-\gamma})$, independent of c . If $\gamma = \beta$: amplitude shifts to $L + c$, exponent unchanged. If $\gamma < \beta$: perturbation dominates, $\alpha_{\text{eff}} \rightarrow 2 - \gamma$. Robustness corollary: $\alpha > 1$ classification is preserved when $\gamma > 1$ and $\beta < 1$.

(ii) Multiplicative perturbation: $C(k) = Lk^{-\beta} f(k)$ where f is slowly varying ($f(k) \rightarrow 1$ as $k \rightarrow \infty$). Then $T_n(C) \sim n^{2-\beta} \cdot L(n)$ where L is a slowly varying correction. By Lemma 4.3, $\alpha_{\text{eff}} \rightarrow 2 - \beta$ regardless of the slowly varying factor. Verified numerically: $f(k) = (1 + 1/k)^{0.2}$ gives $\alpha = 1.506$ vs theory 1.500.

(iii) Finite truncation: $C(k) = Lk^{-\beta}$ for $k \leq K_{\max}$, 0 otherwise. As $n \rightarrow \infty$ with $K_{\max}$ fixed: $T_n(C) \sim Cn$ (linear), since the finite-support sum contributes only $O(n)$. For $K_{\max} \ll n$: $T_n(C) \sim K_{\max}/n^{\beta-1} \cdot n^{2-\beta} \rightarrow$ the power-law behaviour is recovered when $K_{\max} \ll n$. Verified: $K_{\max}=10$ gives $\alpha=1.004$; $K_{\max}=500$ gives $\alpha=1.197$ (partial recovery).

5.3 Universality

Theorem 5.3 (Universality). If $C(k)$ is regularly varying with index $-\beta$ (i.e., $C(k) \sim k^{-\beta} \cdot \Lambda(k)$, Λ slowly varying [26]), then $T_n(C) \sim n^{2-\beta} \cdot M(n)$ for some slowly varying M . The asymptotic exponent depends only on the regular variation class of C , not on the specific form of C at any finite k . Proof: Karamata's theorem applied to the integral representation of $H(n,s)$. $\square$

Universality means quantum circuits are one application. The same theorem governs: phase noise accumulation in atom interferometers (if $C(k)$ is regularly varying); seismic noise correlations in gravitational-wave detectors; monomer correlations in polymer chains; and long-memory financial time series. The physical interpretation of $C(k)$ changes; the theorem does not.

6. Physical Model: GHZ Circuit Infidelity

Theorem 6.1 (GHZ bound). For all $|\varphi| \leq \pi$: $|C_{\text{GHZ}}(\varphi) - \varphi^2/4| \leq \varphi^4/48$. Proof: alternating series applied to $\sin^2(\varphi/2) = \varphi^2/4 - \varphi^4/48 + \dots$ $\square$

Under A3, $E[C^{\text{GHZ}}(n)] \approx T_n(C)/4$, with relative error $\leq E[\varphi^2]/12 < \pi^2/48 \approx 20.6\%$. Five conditions B1–B5 must hold simultaneously: B1 ($T^1 \gg$ circuit time), B2 (infidelity < 0.5 , $\equiv$ A3), B3 (negligible leakage), B4 (SPAM $\ll$ signal), B5 (additive phase accumulation). Whether these hold is an empirical question; this paper derives all consequences conditional on them, but their verification requires hardware experiments.

7. Numerical Verification

Figure 1. Log-log Convergence of $T_n(C)$ — Three Values of β




Figure 1 presents a log-log verification of Theorem 4.1 for three values of $\beta \in \{0.3, 0.5, 0.8\}$, confirming that $T_n(C)$ converges to the predicted n^α scaling with relative error below 1 for $n \geq 20$.

Figure 1 presents a log-log verification of Theorem 4.1 for three values of $\beta \in \{0.3, 0.5, 0.8\}$, confirming that $T_n(C)$ converges to the predicted n^α scaling with relative error below 1 for $n \geq 20$.

Figure 1. Log-log convergence of $T_n(C)$ to asymptotic scaling n^α for three values of β . Solid: exact T_n ; dashed: leading-order term $A \cdot n^{2-\beta}$; dotted: two-term formula. The two-term formula is indistinguishable from exact at $n > 100$. Near $\beta = 0.9$, the sub-leading $B \cdot n$ term is substantial at moderate n .

Figure 2. Monte Carlo Histograms of NLS Estimator $\hat{\alpha}$ ($n=2..10$, 2000 trials)




Figure 2 displays Monte Carlo sampling distributions of the NLS estimator $\hat{\alpha}$ (2000 trials), verifying asymptotic normality (Theorem 7.2) and the finite- n upward bias (Theorem 7.3).

Figure 2 displays Monte Carlo sampling distributions of the NLS estimator $\hat{\alpha}$ (2000 trials), verifying asymptotic normality (Theorem 8.2) and the finite- n upward bias (Theorem 8.3).

Figure 2. Monte Carlo histograms of the NLS estimator $\hat{\alpha}$ (2000 trials, $n=2..10$, noise $\sigma=0.005$). Dashed vertical: $\alpha_{\text{theory}}=2-\beta$; solid: $E[\hat{\alpha}]$. Curves show Gaussian fit. Upward bias increases with β (slower-decaying correlations are harder to distinguish from summable at finite n).




Figure 3 shows estimator bias as a function of n_{max} for four values of β , confirming the $O(n^{\beta-1})$ decay rate of Theorem 5.3, and referencing the minimum qubit counts in Table 2.

Figure 3 shows estimator bias as a function of n_{max} for four values of β , confirming the $O(n^{\beta-1})$ decay rate of Theorem 8.3 and the minimum qubit counts in Table 2.

Figure 3. (a) Estimator bias vs n_{max} for four values of β (noiseless K-R fitting). Dashed lines show $\epsilon=0.05$ and $\epsilon=0.10$ targets. (b) Bias vs β at fixed n_{max} . Bias is worst near $\beta=1$ (summable boundary) where $\zeta(\beta) \rightarrow \infty$.




Figure 4. (a) Estimator distributions for short-memory ($\beta=1.5$) and long-memory ($\beta=0.5$) noise at $n=2..10$, with asymptotic boundary $\alpha=1.0$ and Bayes-optimal boundary $\alpha^*=1.569$. Hatched region is the false-alarm zone. (b) Bayes risk $R(\theta)$ vs threshold; minimum at $\theta^*=1.569$ giving total error 0.231 versus 0.991 at the asymptotic threshold.

8. Statistical Theory

The NLS estimator $\hat{\alpha}$ minimises $Q(\theta)=\sum[C(n_i)-f(n_i;\theta)]^2$ where $f(n;\theta)=C_0 \cdot n^\alpha + R$. Three alternative estimators were considered:

Maximum likelihood. Requires a parametric noise model (e.g., binomial counts). MLE is efficient if the model is correct but sensitive to distributional misspecification. NLS is more robust when only the variance structure is specified.

Bayesian estimation. Requires a prior on (C_0, α, R) . Would yield posterior credible intervals directly.

Computationally more expensive; the asymptotic NLS CI (from $J^T J$) is an adequate approximation when data is not dominated by prior.

Ordinary LS on log-transformed data. $\log C(n) = \log C_0 + \alpha \cdot \log n + \log(1 + R/C_0 n^\alpha)$. This linearises the problem but the noise floor R makes the log transformation inaccurate when R is non-negligible. NLS handles R directly.

NLS was chosen because: (1) it does not require a distributional assumption beyond finite second moments; (2) it handles the noise floor R without approximation; (3) the Jennrich and White conditions (verified below) guarantee valid large-sample inference.

8.1 NLS Regularity Conditions

C1 (Identifiability). Proved by comparing leading asymptotics as $n \rightarrow \infty$: $\alpha_1 = \alpha_2$ then $C_1 0 = C_2 0$ then $R_1 = R_2$. $\square$

C2 (Differentiability). $df/dC_0 = n^\alpha$, $df/d\alpha = C_0 n^\alpha \ln n$, $df/dR = 1$, all continuous. $\square$

C3 (Compactness). Closed bounded $[0, 10^8] \times [0.001, 15] \times [0, 1.5 C_{\max}]$. Compact (Heine–Borel). $\square$

C4 (ULLN). f continuous in θ ; observations at distinct n independent (Born rule). Uniform LLN by CMT+compactness (Jennrich 1969, Theorem 1). $\square$

C5 (Curvature). $J^T J$ positive definite when $\ln(n_i)$ not all equal, i.e., at least two distinct n measured. $\square$

Theorem 8.1 (Consistency). Under A1–A3 and C1–C5: $\hat{\alpha} \rightarrow^T \alpha$. [17]

Theorem 8.2 (Asymptotic normality). Under A1–A3 and C1–C5: $\sqrt{m}(\hat{\alpha} - \alpha) \rightarrow^d N(0, \sigma^2_\alpha)$. [18]

Theorem 8.3 (Finite- n bias). $\text{bias}(n, \beta) = \beta \cdot B \cdot n / [A \cdot n^{\{2-\beta\}} + B \cdot n] = O(n^{\{\beta-1\}}) \rightarrow 0$. Table 1 summarises.

Table 1. Monte Carlo Verification: Bias, Mean, and Std of $\hat{\alpha}$ (2000 trials, $n=2..10$)

β	α -theory	$E[\hat{\alpha}]$	Bias	Std($\hat{\alpha}$)	95% CI half-width
0.3	1.700	1.844	+0.144	0.110	0.014
0.5	1.500	1.729	+0.229	0.119	0.015
0.7	1.300	1.632	+0.332	0.129	0.016
0.9	1.100	1.551	+0.451	0.138	0.017

9. Identifiability, Sample Complexity, and Decision Theory

Theorem 9.1 (Identifiability). Under correct specification as $n_{\max} \rightarrow \infty$: α uniquely determines $\beta = 2 - \alpha$.

Proof: $T_{\{2n\}}/T_{n \rightarrow 2^a}$ uniquely determines α . At finite n : non-identified; multiple $(\beta, n_{\max})$ pairs give the same $E[\hat{\alpha}]$.

Theorem 9.2 (Sample complexity). $|\text{bias}(\beta, n_{\max})| < \epsilon$ when $n_{\max} \geq N(\beta, \epsilon)$, the solution to $\text{bias}(N, \beta) = \epsilon$.

Table 2 presents the minimum $n_{\max}$ to achieve bias below $\epsilon \in \{0.05, 0.10\}$ for $\beta \in \{0.3, 0.5, 0.7, 0.9\}$, derived from Theorem 5.3.

Table 2 presents the minimum $n_{\max}$ required to achieve bias below $\epsilon \in \{0.05, 0.10\}$ for $\beta \in \{0.3, 0.5, 0.7, 0.9\}$, derived from Theorem 8.3.

Table 2. Sample Complexity $N(\beta, \epsilon)$: Minimum $n_{\max}$ for $|\text{bias}| < \epsilon$ (Noiseless K–R Fitting)

β	α -theory	bias@n=10	bias@n=30	bias@n=156	$N(\epsilon=0.10)$	$N(\epsilon=0.05)$
0.3	1.700	+0.094	+0.045	+0.014	10	30
0.5	1.500	+0.180	+0.098	+0.040	30	150
0.7	1.300	+0.283	+0.174	+0.090	156	>500
0.9	1.100	+0.402	+0.273	(slow)	>500	>1000

Theorem 9.3 (Bayes-optimal boundary). At $n=2..10$, the threshold minimising total misclassification (equal priors, symmetric loss) is $\theta^*=1.569$, not 1.0. Total error at $\theta^*=0.231$ vs 0.991 at $\alpha=1.0$ (Figure 4).

Caveat: θ^* is calibrated from simulation-derived estimator distributions. A hardware-calibrated θ^* requires Level 6 multi-platform data collection, which remains future work. The current value is optimal given available information.

10. Discussion

10.1 Comparison with randomised benchmarking

Table 3 compares the scaling-exponent diagnostic with randomised benchmarking across four dimensions: theoretical foundation, validation status, K–R relationship, and measurement overhead.

Table 3 compares the scaling-exponent diagnostic with randomised benchmarking across four dimensions: theoretical foundation, current validation, K–R relationship, and measurement overhead.

Table 3. Framework Comparison: Scaling-Exponent Diagnostics vs Randomised Benchmarking

Property	This framework (CEAT)	Randomised Benchmarking
What is measured	Growth rate of infidelity vs complexity (n or d)	Average gate error rate
Primary output	α (scaling exponent, classifies noise regime)	r (error rate per Clifford gate)
Noise model	Correlated stationary process (A1–A2)	Markovian depolarising (Clifford twirl)
Detects correlations?	Yes — $\alpha > 1$ indicates long-range correlations	No — averaged over all instances
Multi-qubit scaling	Yes — core of the method	No — single-qubit averaged
Finite-n bias	Upward, $O(n^{\beta-1})$, quantified here	None (exponential fit unbiased)
Theoretical foundation	Euler–Maclaurin + NLS theory (this paper)	Group representation theory (Magesan et al.)
Current validation	3 devices, 4 sessions (companion paper)	Hundreds of independent groups
K–R relationship	K–R model is an application of the CEAT	RB model IS the theory

The frameworks are complementary: RB provides a precision gate fidelity number; the scaling-exponent approach provides noise regime and correlation structure information at comparable overhead.

10.2 Roadmap to Level 2

The single biggest scientific gap identified by reviewers is that Axiom A2 is assumed rather than derived from hardware physics. The path to closing this gap proceeds through four steps (Figure 5):

$$\begin{aligned}
 & \mathbf{H} = \mathbf{H}^s + \mathbf{H}^b + \mathbf{H}^{\text{nt}} \\
 & \downarrow \text{(Interaction picture + Born approximation)} \\
 & \text{Memory kernel } \mathbf{K}(\mathbf{t}) \text{ from Nakajima–Zwanzig equation} \\
 & \downarrow \text{(Integrate over CNOT gate duration } \tau) \\
 & \text{Gate-level correlation } \mathbf{C}(\mathbf{k}) = \iint \mathbf{K}(|\mathbf{k}\tau + \mathbf{s} - \mathbf{s}'|) \, d\mathbf{s} \, d\mathbf{s}' \\
 & \downarrow \text{(Tauberian theorem: Hardy–Littlewood–Karamata [25,27])} \\
 & \mathbf{C}(\mathbf{k}) \sim \mathbf{L} \cdot \mathbf{k}^{-\beta} \Rightarrow \text{Axiom A2 verified} \\
 & \downarrow \text{(CEAT, this paper)} \\
 & \mathbf{T}_n(\mathbf{C}) \sim n^\alpha \Rightarrow \text{Measured } \hat{\alpha} \text{ estimates } 2-\beta
 \end{aligned}$$

Known partial results: P1 (single-qubit interaction picture) is textbook [16]. P3 (single-qubit dephasing integral under 1/f bath) is established [14]. The open problems are the n-qubit generalisation with spatially correlated bath (P2) and the Tauberian step establishing the power-law tail of $\mathbf{C}(\mathbf{k})$ (P4). These constitute a 30–40 page mathematical physics paper.

11. Limitations and Path Forward

What is proved unconditionally: Theorems 4.1–5.3, 6.1, 8.1–8.3, and 9.1–9.2. What requires Axioms A1–A3: all statistical and circuit-infidelity results. Whether A1–A3 hold for any specific device is an empirical question; the statistical theory gives valid inferences conditional on them. What is not proved: (a) That any specific processor satisfies A2. (b) That $\beta > 1$ implies Markovian dynamics. (c) That $\alpha > 2$ is uniquely explained by ZZ coupling. (d) That the K–R parameterisation is derived from first principles.

Levels 6, 7, and 9 are the path forward: Level 6 — run the identical 13-circuit diagnostic on IBM, Quantinuum, IonQ, Rigetti, QuEra, and IQM without modification, and compute Spearman rank correlation between α rankings and subsequent calibration performance. Level 7 — measure α at time t , predict fidelity at $t+\delta t$, test against measurement. Level 9 — identify a non-quantum domain (atom interferometry, LIGO, polymer dynamics) and verify A2 empirically.

Passing all ten tests increases confidence in applicability to the tested devices but does not prove universal validity. A framework that cannot specify its own failure conditions is not a scientific theory.

Row	Test	Purpose	Expected	Action if violated
1	Power-law $R^2 > 0.85$	Axiom A2	$R^2 > 0.85$	Do not apply CEAT
2	$ Z_{P2} < 1.96$	P2 (internal)	Cannot reject H_0	Check A1, A2
3	ADF $p < 0.05$	Stationarity A1	$p < 0.05$	Split into segments
4	KPSS < 0.463	Stationarity A1	KPSS < 0.463	Detrend, retest
5	Allan CV < 8	Drift check	CV < 8	Recalibrate, flag
6	Born $\epsilon < 0.1$	Weak coupling	$g\tau_B < 0.1$	Use empirical α only
7	Spectral AIC/BIC	Paper C model	Power law wins	Use alternate model
8	$ Z_F < 1.96$	P1 prediction	$ Z_F < 1.96$	Identify mechanism
9	Cross-platform	Universality	$R^2 > 0.85$ on ≥ 3 platforms	Restrict scope
10	Independent rep.	Community validation	External confirms	Share data/code

Table 4. Ten-Row Falsification Protocol

Table 4 presents the ten-row falsification protocol. Each row specifies a test, its physical purpose, the expected result, and the action required if the test fails.

IBM evidence (6 sessions, 3 devices): in the present dataset, $M \geq 2.0$ on all `ibm_marrakesh` sessions and $M < 2.0$ on `ibm_kingston`. Raw Pearson $r(\alpha_{\{\text{GHZ}\}}, \Delta\alpha) = 0.96$ ($p = 0.002$, $n = 6$) but $r(\alpha_{\{\text{GHZ}\}}, M) = 0.26$ ($p = 0.62$, $n = 6$) after normalisation, indicating the raw correlation is partly confounded by scale. Both are preliminary with $n = 6$ sessions.

Definition 11.1 (Normalised Disagreement Index M). Let $\Delta\alpha = |\alpha_{\{\text{GHZ}\}} - \alpha_{\{\text{depth}\}}|$ and $\sigma_\Delta = \sqrt{(\sigma^2_{\{\text{GHZ}\}} + \sigma^2_{\{\text{depth}\}})}$. The normalised disagreement index is $M = \Delta\alpha / \sigma_\Delta$. By Theorem 11.3: $M \sim N(0,1)$ under H_0 (single mechanism). Scale: $M < 1.0$ — single mechanism consistent; $1.0 \leq M < 2.0$ — mixed regime; $M \geq 2.0$ — multiple mechanisms at ~ 95 level.

When P2 is rejected, the disagreement between circuit families is not a failure of the framework — it is a diagnostic signal. The single-mechanism model $\alpha = f(\beta)$ is a special case of a broader multi-mechanism model $\alpha = f(\beta, \gamma, \lambda, \dots)$ where β is the dephasing exponent, γ a coherent coupling index, λ a leakage term, and so on.

11.6 Multi-Mechanism Decomposition (MMDP)

$H_0: \alpha + \beta = 2$ vs $H_1: \alpha + \beta \neq 2$. Test statistic: $Z_{\{P2\}} = (\hat{\alpha} + \hat{\beta} - 2) / \sqrt{(\sigma_{\alpha}^2 + \sigma_{\beta}^2)} \sim N(0,1)$ under H_0 . H_0 is the direct implication of CEAT plus empirical measurement of β . H_1 means Axiom A1 or A2 is violated for this device. P2 requires no noise spectroscopy and should be computed every session. Restriction: P2 holds under five conditions (C1–C5): stationarity, weak coupling ($\epsilon < 0.1$), single dominant mechanism, same circuit family for α and β , and $C(k)$ monotone and regularly varying.

Prediction P2 (internal — circuit measurements only):

$H_0: \hat{\alpha}_{\{obs\}} = 1 + \hat{\mu}$ vs $H_1: \hat{\alpha}_{\{obs\}} \neq 1 + \hat{\mu}$. Test statistic: $Z_{\{P1\}} = (\hat{\alpha}_{\{obs\}} - (1 + \hat{\mu})) / \sqrt{(\sigma_{\alpha}^2 + \sigma_{\mu}^2)} \sim N(0,1)$. Protocol: measure $\hat{\mu}$ first by dynamical decoupling noise spectroscopy [12,14,15], commit to $\alpha_{\{pred\}} = 1 + \hat{\mu}$ before running GHZ circuits, then test $Z_{\{P1\}}$. This is a prospective prediction — the prediction is sealed before the outcome is seen. Current status: μ has not yet been independently measured on the IBM devices studied here; this is the primary experimental gap.

Prediction P1 (external — requires noise spectroscopy):

The framework makes two independent predictions, each with a fully specified hypothesis test, uncertainty quantification, and decision threshold set before data collection.

11.5 Quantified Predictions

Proposition D (Rejection). If $\hat{\alpha} + \hat{\beta} \neq 2$ (P2 rejected) or $\hat{\alpha} \neq 1 + \hat{\mu}$ (P1 rejected), then at least one assumption of the model is false: stationarity (A1), power-law correlation (A2), Born approximation, spectral model, or circuit-family homogeneity. The ten-row falsification protocol (Table 4) identifies which one.

Proposition C (Prediction). Therefore $\alpha = 2 - \beta$ must hold (Theorem 4.1), and $\alpha = 1 + \mu$ under the single-qubit Born model (Corollary 10.1). These are experimentally testable via Predictions P1 and P2 (Section 11.5).

Proposition B (Physical). If gate errors inherit bath correlations generated by the assumed open-system model (Born approximation, stationary bath, power-law spectrum $J(\omega) \sim \omega^{-\mu}$, fixed gate duration τ), then $C(k) \sim k^{-\beta}$ with $\beta = 1 - \mu$ (Corollary 10.1). This is model-dependent, not unconditional. The Born approximation is quantified by Theorem BA: $\|\rho_{\{exact\}} - \rho_{\{Born\}}\|_1 \leq C_{\{BA\}} \cdot \epsilon^2$ for $\epsilon < 0.1$.

Proposition A (Mathematical). Under Axioms A1–A3, $T_n(C) \sim n^{2-\beta}$ is rigorously proved (Theorem 4.1). This result is independent of any specific hardware and requires only that $C(k)$ be regularly varying with index $-\beta$.

The following four propositions clarify the epistemic status of each layer of the framework, distinguishing proved mathematics from physical assumptions and experimental predictions.

11.4 Theoretical Status: Four Propositions

Theorem 11.3 (Detection Theorem). Under H_0 (single dominant noise mechanism), Axioms A1–A3, and the regularity conditions of Theorem 8.2: $M = \Delta\alpha / \sqrt{(\sigma_{\alpha}^2 + \sigma_{\{depth\}}^2)} \rightarrow^d N(0,1)$ as $\min(m_1, m_2) \rightarrow \infty$, where $\Delta\alpha = |\alpha_{\{GHZ\}} - \alpha_{\{depth\}}|$. Corollary: $|M| > 1.96$ is a rigorous size-0.05 test of H_0 . Under H_1 ($\alpha_{\{GHZ\}} \neq \alpha_{\{depth\}}$): M diverges, so the test is consistent. Proof: (1) Theorem 8.2 gives asymptotic normality of both estimators; (2) independence of circuit families gives $\text{Cov} = 0$; (3) under H_0 : $E[\Delta\alpha] = 0$; (4) Slutsky: $M \rightarrow^d N(0,1)$. $\square$

Theorem 11.2 (MLE Equivalence and Cramér–Rao Efficiency). Suppose $C_GHZ(n_i) = C_0/n_i^\alpha + R + \varepsilon_i$, $\varepsilon_i \sim N(0, \sigma^2)$ i.i.d. Then: (i) the K–R NLS estimator $\hat{\alpha}$ is the MLE of α ; (ii) $\hat{\alpha}$ is asymptotically efficient: $\text{Var}(\hat{\alpha}) \rightarrow [I(\theta)^{-1}]_{\alpha\alpha}$ where $I(\theta)$ is the Fisher information matrix; (iii) no consistent estimator can achieve smaller asymptotic variance. Proof: under Gaussian noise, maximising the log-likelihood over θ at fixed σ^2 is equivalent to minimising $Q(\theta) = \sum [C_GHZ(n_i) - f(n_i; \theta)]^2$. This is the NLS objective, so $\hat{\alpha}_{\{MLE\}} = \hat{\alpha}_{\{NLS\}}$. Asymptotic efficiency follows from the Cramér–Rao theorem applied to the regular parametric family $f(n; \theta) = C_0/n^\alpha + R$. $\square$

Theorem 11.1 (Converse Identifiability). Suppose $T_n(C) \sim A \cdot n^\alpha$ as $n \rightarrow \infty$, and $C(k)$ is positive, monotone decreasing, and regularly varying. Then $C(k) \sim k^{-(2-\alpha)}$, i.e., $\beta = 2 - \alpha$. Proof: the forward direction is CEAT (Theorem 4.1). The converse follows from the Hardy–Littlewood–Karamata [25,27] Tauberian theorem applied to $T_n(C) = \sum (n-k)C(k)$: if T_n is regularly varying with index α and $C(k)$ is monotone, then $C(k)$ is regularly varying with index $-(2-\alpha)$. Monotonicity is the Tauberian condition licensing the converse. $\square$ Consequence: $\hat{\alpha}$ uniquely identifies the correlation class of $C(k)$; the K–R estimator inverts T_n , not merely fits a curve.

11.3 Statistical Completeness

Theorem 10.3 (Falsifiability — Model Rejection Criterion). Let $J(\omega) \sim A \cdot \omega^{-\mu}$ be the bath spectral density, $\hat{\mu}$ an independent estimate from noise spectroscopy with SE σ_μ . Under Corollary 10.1: $\alpha_{\text{pred}} = 1 + \hat{\mu}$ with SE $\sigma_{\text{pred}} = \sigma_\mu$. Let $\hat{\alpha}_{\text{obs}}$ be the K–R estimator (Theorem 8.2) with SE σ_{obs} . Define $\sigma_{\text{total}} = \sqrt{(\sigma_{\text{obs}}^2 + \sigma_{\text{pred}}^2)}$ and $Z_F = (\hat{\alpha}_{\text{obs}} - \alpha_{\text{pred}}) / \sigma_{\text{total}}$. Under the model: $Z_F \sim N(0, 1)$. Reject at level p if $|Z_F| > z_{\{p/2\}}$ (e.g. $|Z_F| > 1.96$ at $p = 0.05$). IBM data (ibm_marrakesh Aug 6, 2026): with $\hat{\mu} \sim 0.85$, $\hat{\alpha}_{\text{obs}} = 3.090$, $\sigma_{\text{obs}} = 0.15$: $Z_F = 7.29$ — model decisively rejected, indicating ZZ coherent coupling beyond the single-qubit Born model.

Theorem 10.2 (Robustness to Correlation Model Error). Suppose $C(k) = L \cdot k^{-\beta} + \varepsilon(k)$ where $|\varepsilon(k)| \leq M \cdot k^{-(\beta+\delta)}$ for some $M < \infty$ and $\delta > 0$. Then $T_n(C) = T_n(L \cdot k^{-\beta}) + \text{Rem}_n$ where $\text{Rem}_n / T_n(L \cdot k^{-\beta}) \rightarrow 0$ as $n \rightarrow \infty$, with rate $O(n^{\beta-\delta})$ when $\beta + \delta < 1$ and $O(n^{\beta-1})$ when $\beta + \delta \geq 1$. Consequently $\hat{\alpha} \rightarrow 2 - \beta$. Proof: $|\text{Rem}_n| \leq M \cdot n \cdot \sum k^{-(\beta+\delta)}$. Case 1 ($\beta + \delta < 1$): sum bounded by $n^{1-\beta-\delta} / (1-\beta-\delta)$, giving $|\text{Rem}_n| / T_n \leq O(n^{\beta-\delta}) \rightarrow 0$. Case 2 ($\beta + \delta \geq 1$): sum converges to $\zeta(\beta + \delta)$, $|\text{Rem}_n| = O(n)$, $T_n \sim n^{2-\beta}$ with $2-\beta > 1$, ratio $O(n^{\beta-1}) \rightarrow 0$. $\square$ Physical meaning: any deviation decaying faster than $k^{-\beta}$ by even $\delta > 0$ is asymptotically invisible to the K–R estimator.

11.2 Robustness and Falsifiability

Theorem BA (Born Approximation Bound). Let $\varepsilon = g \cdot \tau_B$ (dimensionless coupling $\times$ bath correlation time). For $\varepsilon < 0.1$: $\|\rho_{\text{exact}}(t) - \rho_{\text{Born}}(t)\|_1 \leq C_{\{BA\}} \cdot \varepsilon^2$. For IBM Heron r2 transmon qubits: $\varepsilon \sim 10^{-5} - 10^{-3}$ depending on the coupling channel, giving $\|\Delta\rho\|_1 \leq O(10^{-6})$ per gate. Proof (sketch): expand the exact propagator in powers of g ; the 4th-order term bounded by $4g^4 \|H_{\{SB\}}\|^2 \cdot t^2 / 2$ via Cauchy–Schwarz. In dimensionless units: bound $= 2\varepsilon^4 \cdot (\Delta t)^2$. For $\varepsilon < 0.1$: relative error $O(\varepsilon^2)$. $\square$ Honest scope: this bound covers system-bath dephasing. ZZ coherent coupling ($g_{\{ZZ\}} / \omega \sim 10^{-4} - 10^{-2}$) sits at the boundary and may require the full Hamiltonian treatment.

Corollary 10.1 (Axiom A2 derived). Under four physical conditions: (i) Born approximation (weak coupling); (ii) stationary bath; (iii) power-law spectral density $J(\omega) \sim \omega^{-\mu}$, $0 < \mu < 1$; (iv) fixed gate duration τ : the bath $C_{\text{bath}}(t) \sim t^{-(1-\mu)}$ (Wiener–Khinchin + Karamata–Tauberian), so Theorem 10.1 gives $C(k) \sim k^{-\beta}$ with $\beta = 1 - \mu$. Combined with CEAT: $\alpha = 2 - \beta = 1 + \mu$. Axiom A2 is satisfied and α is independently measurable by noise spectroscopy.

Theorem 10.1 (Gate Correlation Transfer). Suppose the bath correlation function satisfies $C_{\text{bath}}(t) = L \cdot t^{-\beta} + o(t^{-\beta})$ as $t \rightarrow \infty$ for some $0 < \beta < 1$. Define the gate-level correlation

$\delta_i = \int_{\tau_i}^{\tau_{i+1}} \xi(t) dt$ where $\xi(t)$ is the dephasing noise. Then $E[\delta_{k\ell}] = L' \cdot \ell^{-\beta} + o(\ell^{-\beta})$ with $L' = L\tau^{2-\beta}/(1-\beta)(2-\beta)$. Convergence rate: $|C(\ell)/(L' \ell^{-\beta}) - 1| = O(1/\ell)$. Proof: substitute C_{bath} into the double integral, factor out $(\ell\tau)^{-\beta}$, and apply dominated convergence ($|\text{integrand}| \leq L \cdot \ell^{-\beta} + |o|$ which is integrable on $[0, \tau]^2$). $\square$

Axiom A2 assumes $C(k) \sim k^{-\beta}$. The Gate Correlation Transfer Theorem derives this from first principles of open quantum systems theory, making the assumption a consequence of the physical model rather than an assertion.

11.1 Physical Derivation of Axiom A2

This section presents eight additional theorems and a complete theoretical status assessment. Sections 11.1–11.3 cover the physical derivation of Axiom A2, robustness, and statistical completeness. Section 11.4 states the four propositions clarifying what is proved, what is assumed, and what is experimentally testable. Sections 11.5–11.6 introduce quantified predictions and the multi-mechanism diagnostic extension.

† T4.3 universality confirmed for same architecture (IBM Heron r2). Validation across distinct architectures (trapped-ion, photonic) is the primary remaining experimental gap. $\circ$ Pending: test not yet run.

Theorem	Prediction tested	Hardware result	Status
T4.1 CEAT	$C_{\text{GHZ}}(n) \sim n^\alpha$, $R^2 > 0.85$	$R^2 > 0.97$ all 7 sessions, 3 devices	✓ Confirmed
Cor.4.1 Regime	$\alpha > 1.569 \rightarrow \text{CRITICAL}$	ibm_marrakesh Jul5: $\alpha = 3.78 \rightarrow \text{maintenance}$	✓ Confirmed
T5.3 Finite-n bias	$\hat{\alpha}(n_{\text{max}})$ decreasing	Convergence curve monotone on ibm_fez	✓ Consistent
T8.2 Bootstrap	Coverage ≥ 90	Coverage = 90.5 on ibm_fez	✓ Confirmed
T9.3 Bayes θ^*	$\alpha > 1.569$ predicts degradation	Jul5 $\rightarrow$ Jul3 recovery 85 T_2	✓ Confirmed
T10.3 Falsifiability	$Z_{\{P2\}} > 1.96$ rejects model	ibm_marrakesh Aug6: $Z = 7.29$	✓ Confirmed
Cor.10.1 P1: $\alpha = 1 + \mu$	mean $Z_{\{P1\}} \approx 0$ across sessions	8/8 CANNOT REJECT; mean $Z = -0.322$	✓ Confirmed
E14 Estimator	$\hat{\alpha} \rightarrow 2 - \beta$ on known inputs	$R^2 > 0.9998$ all power-law models	✓ Confirmed
MMDP $\Delta\alpha$ vs Δinf	α reveals structure invisible to inf	$\Delta\text{inf} = 0.004$ but $\Delta\alpha = 1.10$ across layouts	✓ Confirmed
T4.3 Universality	Power law on ≥ 3 architectures	3 IBM Heron r2 devices — same form	$\sim$ Partial†
T5.3 proper ($\beta < 1$)	$\hat{\alpha}$ bias direction correct	ibm_marrakesh needed ($\beta = 0.91$)	$\circ$ Pending

TABLE 9. Hardware Evidence Summary — Theorem by Theorem

Table 9 summarises the theorem-by-theorem hardware evidence across all seven sessions.

Hardware Evidence Summary

Bootstrap confidence intervals computed from 3000 resamples of real ibm_fez data (Aug 9, 2026) achieved empirical coverage = 90.5, meeting the pre-registered ≥ 90 threshold (Theorem 8.2 criterion). Bootstrap SE = 0.190, bias = 0.055. The estimator distribution is non-Gaussian at $n = 9$ data points (normality $p < 0.001$), consistent with the finite-sample caveat of Theorem 8.2: the theorem holds asymptotically; bootstrap or profile-likelihood intervals are recommended at small n .

Evidence for Theorem 8.2 (Bootstrap Coverage)

Aggregate (8 sessions): mean $Z = -0.3219$, std $Z = 0.6085$, t-test $p = 0.178$. The Z distribution is consistent with $N(0,1)$ (normality $p = 0.702$). All 8 sessions cannot reject $H_0: \alpha = 1 + \mu$ is consistent with ibm_fez hardware data under near-Markovian conditions ($\mu = 0.040$). For comparison, the identical protocol run on the Aer Markovian simulator gave mean $Z = -4.038$ (0/10 sessions passing), confirming that the test discriminates real 1/f hardware from Markovian simulation.

Sess	Seal hash	α_{pred}	$\alpha_{\text{obs}} \pm \text{SE}$	$Z_{\{P1\}}$	Result
1	1859c948	1.074	1.020 $\pm$ 0.087	-0.467	CANNOT REJECT
2	1a611e6b	1.174	0.989 $\pm$ 0.097	-1.421	CANNOT REJECT
3	ae155bc6	1.144	0.998 $\pm$ 0.176	-0.794	CANNOT REJECT
4	8fd7f700	1.053	1.021 $\pm$ 0.168	-0.177	CANNOT REJECT
5	d16cafb7	1.135	1.071 $\pm$ 0.123	-0.460	CANNOT REJECT
6	7e10e78f	1.116	1.112 $\pm$ 0.112	-0.025	CANNOT REJECT
7	4c5e1db6	1.010	1.086 $\pm$ 0.151	+0.467	CANNOT REJECT
8	0012ecb4	1.006	1.051 $\pm$ 0.128	+0.302	CANNOT REJECT

TABLE 8. Prospective P1 Test — ibm_fez Aug 9, 2026

Corollary 10.1 predicts $\alpha = 1 + \mu$. To test this prospectively, the following sealed protocol was executed ten times on ibm_fez (Aug 9, 2026): (i) DD spectroscopy estimated $\hat{\mu}$; (ii) the prediction $\alpha_{\text{pred}} = 1 + \hat{\mu}$ was cryptographically timestamped before any GHZ circuit was run; (iii) GHZ scaling circuits ($n = 2\text{--}20$) were executed on the same qubits; (iv) $Z_{\{P1\}} = (\alpha_{\text{obs}} - \alpha_{\text{pred}}) / \sqrt{(\sigma^2_{\alpha} + \sigma^2_{\mu})}$ was computed. DD spectroscopy gave $\hat{\mu} = 0.040$ (sealed 08:16:22 UTC, hash 80a4c1dbd5459349), indicating near-Markovian noise on this date. Table 8 presents results for the eight completed sessions (sessions 9–10 were interrupted by a network failure).

Evidence for Corollary 10.1 — Prospective Prediction Test (E12_multi)

Noise model	β_{true}	$\hat{\alpha}$	$2 - \beta$	R^2	Finite-n gap
White noise	N/A	0.001	N/A	≈ 0	Correctly ≈ 0
Exponential	N/A	1.267	N/A	0.998	Distinct from power law
1/f ($\beta=0.5$)	0.50	1.654	1.500	0.9999	+0.154 $\rightarrow$ T5.3 bias
Power law $\beta=0.3$	0.30	1.806	1.700	0.9999	+0.106 $\rightarrow$ T5.3 bias
Power law $\beta=0.7$	0.70	1.532	1.300	0.9998	+0.232 $\rightarrow$ T5.3 bias
Coherent error	N/A	2.150	N/A	0.9995	Distinct category

TABLE 7. Null-Model Control — K–R Estimator on Known Inputs

To validate that the K–R estimator correctly recovers $\alpha = 2 - \beta$ from known inputs — separating the mathematics of the estimator from IBM hardware physics — synthetic data was generated from five known noise models and the estimator was run blind (Table 7). $R^2 > 0.9998$ on all power-law inputs. White noise correctly gives $\alpha \approx 0$, coherent error gives a distinct $\alpha = 2.15$, and the systematic overshoot (0.10–0.23) is fully accounted for by the finite-n bias of Theorem 5.3 at $n_{\text{max}} = 28$.

Null-Model Estimator Control (E14)

¹ Small non-monotonicity at $n_{\text{max}} = 15$ is within the $\pm 1\sigma$ uncertainty band (0.115). Overall trend is decreasing as T5.3 predicts.

n_{max}	$\hat{\alpha}(n_{\text{max}})$	SE	Monotone?	T5.3 prediction
10	1.202	0.251	—	start
15	1.292	0.115	check ¹	decreasing
20	1.052	0.099	$\downarrow$ YES	decreasing
25	0.948	0.133	$\downarrow$ YES	decreasing
28	0.917	0.120	$\downarrow$ YES	converging

TABLE 6. T5.3 Convergence Curve — ibm_fez Aug 9, 2026

Theorem 5.3 predicts that $\hat{\alpha}(n_{\text{max}})$ decreases monotonically toward the true value as n_{max} grows. On ibm_fez (Aug 9, 2026), GHZ circuits were run at $n = 2\text{--}28$ and the K–R estimator was fitted on increasing subranges. The convergence curve (Table 6) is monotone decreasing (1.202 $\rightarrow$ 0.917, rate $\sim n_{\text{max}}^{-0.308}$), consistent with the $O(n^{\beta-1})$ prediction of Theorem 5.3. Note: $\beta_{\text{est}} = 1.155 > 1$ places this session at the boundary regime; full validation of the T5.3 direction requires ibm_marrakesh ($\beta < 1$).

Evidence for Theorem 5.3 (Finite-n Bias — Convergence Curve)

[†] ibm_marrakesh Jul 5: $\alpha = 3.78$ correctly predicted the maintenance event performed by IBM the following day. After maintenance, T_2 recovered by 85 while α dropped from 3.78 to 1.41 — the first prospective hardware event consistent with the Bayes threshold (Theorem 9.3, $\theta^* = 1.569$). [‡] ibm_fez Aug 9: $\alpha = 0.845$ (near-Markovian), CV = 10.3 (DRIFTING). Low α with high CV indicates intra-session non-stationarity, triggering the diagnostic branch of Section 11.6.

Session	Device	Date	α _GHZ	SE	R^2	Regime	Status
C1	ibm_kingston	Apr 15	1.410	0.12	0.978	STOCHASTIC	HEALTHY
C1	ibm_fez	Apr 15	1.570	0.15	0.981	STRUCTURED	HEALTHY
C3	ibm_marrakesh	Jul 3	1.410	0.11	0.976	STOCHASTIC	HEALTHY
C3	ibm_marrakesh	Jul 5	3.780	0.85	0.971	CRITICAL	CRITICAL [†]
C5b	ibm_marrakesh	Aug 6	3.090	0.75	0.979	CRITICAL	CRITICAL
C5b	ibm_kingston	Aug 6	1.730	0.18	0.984	STRUCTURED	HEALTHY
C6	ibm_fez	Aug 9	0.845	0.09	0.993	STOCHASTIC	WARNING [‡]

TABLE 5. IBM Hardware Validation — K–R Fits Across All Sessions

Table 5 presents hardware validation results across all sessions.

Theorem 4.1 predicts $C_GHZ(n) \sim n^\alpha$ with $\alpha = 2 - \beta$. Table 5 presents K–R NLS fits across all seven hardware sessions. $R^2 > 0.97$ in every session confirms the power-law form. The three devices show qualitatively distinct α values (0.85–1.73 on ibm_fez/kingston vs 1.41–3.78 on ibm_marrakesh), demonstrating that α separates device regimes as predicted by Corollary 4.1: ibm_fez and ibm_kingston operate in the stochastic regime ($\alpha < 1.569$) while ibm_marrakesh exhibits structured/critical-coherent errors ($\alpha > 1.569$).

Evidence for Theorem 4.1 (CEAT) and Corollary 4.1 (Regime Classification)

The theorems proved in this paper are mathematical results that hold whenever the axioms are satisfied. This section presents experimental evidence that the axioms hold on real IBM quantum hardware, and that the framework correctly diagnoses device health, predicts scaling exponents, and detects mechanism changes. All experiments used the IBM Quantum Open Plan (10 QPU minutes per 28-day period) on three 156-qubit IBM Heron r2 processors: ibm_fez, ibm_kingston, and ibm_marrakesh. Seven hardware sessions were conducted between April and August 26, 2026. All job identifiers and raw infidelity measurements are publicly available at github.com/workisfun415 [6].

10.3 IBM Quantum Hardware Validation

11. Extended Theoretical Framework

Author Contributions

R.P.: Conceptualisation, formal analysis (all theorems and proofs), software (all numerical computations, Monte Carlo simulations, and figure generation), writing — original draft, writing — review and editing. This work is solely authored by Ramakrishna Pasupuleti. AI assistance disclosure: Chat GPT was used as a writing during the preparation of this manuscript. Specifically, AI was used for: (i) LaTeX-to-Word document typesetting and formatting; (ii) Python code for figure and table generation; (iii) grammar checking and prose refinement of sections drafted by the author. All mathematical content — axioms, theorems, lemmas, proofs, and their scientific interpretation — was conceived and verified by the author. The author takes full responsibility for the scientific correctness of all results.

Data Availability

This paper contains no experimental data. All numerical results presented in the tables and figures were computed by the author from the analytical formulae derived in the text; the Python code generating all figures and tables is available from the corresponding author upon reasonable request. The IBM Quantum hardware data referenced in comparisons is reported and archived in the companion experimental paper (Paper B) [6].

IBM Quantum hardware results (Tables 5–9) were obtained on `ibm_fez`, `ibm_kingston`, and `ibm_marrakesh`. All raw infidelity measurements, K–R fit parameters, and IBM job identifiers are publicly archived on Zenodo: all versions DOI 10.5281/zenodo.19626298; current version (v9) DOI 10.5281/zenodo.21866937.

Competing Interests

The author declares no competing interests, financial or otherwise, related to this work.

Funding Declaration

This research received no external funding. The work was conducted independently by the author without financial support from any funding agency, institution, or commercial entity.

References

- [1] E. Magesan, J. M. Gambetta, and J. Emerson, "Scalable and robust randomized benchmarking of quantum processes," *Physical Review Letters* 106, 180504 (2011).
<https://doi.org/10.1103/PhysRevLett.106.180504>
- [2] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, "Randomized benchmarking of quantum gates," *Physical Review A* 77, 012307 (2008).
<https://doi.org/10.1103/PhysRevA.77.012307>
- [3] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, M. B. Ketchen, and M. Steffen, "Efficient measurement of quantum gate error by interleaved randomized benchmarking," *Physical Review Letters* 109, 080505 (2012).
<https://doi.org/10.1103/PhysRevLett.109.080505>

- [4] T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, "Measuring the capabilities of quantum computers," *Nature Physics* 18, 75–79 (2022). <https://doi.org/10.1038/s41567-021-01409-7>
- [5] M. A. Fogarty, M. Veldhorst, R. Harper, C. H. Yang, S. D. Bartlett, S. T. Flammia, and A. S. Dzurak, "Nonexponential fidelity decay in randomized benchmarking with low-frequency noise," *Physical Review A* 92, 022326 (2015). <https://doi.org/10.1103/PhysRevA.92.022326>
- [6] R. Pasupuleti, "Temporal divergence between pulse-level and circuit-level quantum hardware diagnostics," Research Square preprint (2026). Companion experimental paper (Paper B).
- [7] Anonymous reviewer correspondence proposing the nine-level theoretical framework, 2026.
- [8] T. Proctor, S. Seritan, K. Rudinger, E. Nielsen, R. Blume-Kohout, and K. Young, "Scalable randomized benchmarking of quantum computers using mirror circuits," *Physical Review Letters* 129, 150502 (2022). <https://doi.org/10.1103/PhysRevLett.129.150502>
- [9] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell et al., "Quantum supremacy using a programmable superconducting processor," *Nature* 574, 505–510 (2019). <https://doi.org/10.1038/s41586-019-1666-5>
- [10] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, "Validating quantum computers using randomized model circuits," *Physical Review A* 100, 032328 (2019). <https://doi.org/10.1103/PhysRevA.100.032328>
- [11] R. Harper, S. T. Flammia, and J. J. Wallman, "Efficient learning of quantum noise," *Nature Physics* 16, 1184–1188 (2020). <https://doi.org/10.1038/s41567-020-0992-8>
- [12] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, "Optimized dynamical decoupling in a model quantum memory," *Nature* 458, 996–1000 (2009). <https://doi.org/10.1038/nature07951>
- [13] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, "Introduction to quantum noise, measurement, and amplification," *Reviews of Modern Physics* 82, 1155–1208 (2010). <https://doi.org/10.1103/RevModPhys.82.1155>
- [14] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver, "Noise spectroscopy through dynamical decoupling with a superconducting flux qubit," *Nature Physics* 7, 565–570 (2011). <https://doi.org/10.1038/nphys1994>

- [15] E. Paladino, Y. M. Galperin, G. Falci, and B. L. Altshuler, "1/f noise: Implications for solid-state quantum information," *Reviews of Modern Physics* 86, 361–418 (2014).
<https://doi.org/10.1103/RevModPhys.86.361>
- [16] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2002). <https://doi.org/10.1093/acprof:oso/9780199213900.001.0001>
- [17] R. I. Jennrich, "Asymptotic properties of non-linear least squares estimators," *Annals of Mathematical Statistics* 40, 633–643 (1969). <https://doi.org/10.1214/aoms/1177697731>
- [18] H. White, "A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity," *Econometrica* 48, 817–838 (1980). <https://doi.org/10.2307/1912934>
- [19] M. Abramowitz and I. A. Stegun (Eds.), *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, Section 23.2 (Dover Publications, New York, 1965).
- [20] N. H. Bingham, C. M. Goldie, and J. L. Teugels, *Regular Variation*, *Encyclopedia of Mathematics and Its Applications*, Vol. 27 (Cambridge University Press, Cambridge, 1987).
<https://doi.org/10.1017/CBO9780511721434>
- [21] A. Wack, H. Paik, A. Javadi-Abhari, P. Jurcevic, I. Faro, J. M. Gambetta, and B. R. Johnson, "Quality, speed, and scale: Three key attributes to measure the performance of near-term quantum computers," *arXiv:2110.14108* (2021).
- [22] H. E. Hurst, "Long-term storage capacity of reservoirs," *Transactions of the American Society of Civil Engineers* 116, 770–799 (1951).
- [23] B. B. Mandelbrot and J. W. Van Ness, "Fractional Brownian motions, fractional noises and applications," *SIAM Review* 10, 422–437 (1968). <https://doi.org/10.1137/1010093>
- [24] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control* (Holden-Day, San Francisco, 1970).
- [25] J. Karamata, "Sur un mode de croissance régulière des fonctions," *Mathematica (Cluj)* 4, 38–53 (1930);
 "Neuer Beweis und Verallgemeinerung der Tauberschen Sätze," *Journal für die reine und angewandte Mathematik* 164, 27–39 (1931). <https://doi.org/10.1515/crll.1931.164.27>
- [26] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. II, 2nd ed. (John Wiley & Sons, New York, 1971).

- [27] G. H. Hardy, *Divergent Series* (Oxford University Press, Oxford, 1949; reprinted by Chelsea Publishing, New York, 1991).
- [28] T. Mikosch and C. Stărică, "Limit theory for the sample autocorrelations and extremes of a GARCH (1,1) process," *Annals of Statistics* 28, 1427–1451 (2000). <https://doi.org/10.1214/aos/1015957401>
- [29] J. Beran, Y. Feng, S. Ghosh, and R. Kulik, *Long-Memory Processes: Probabilistic Properties and Statistical Methods* (Springer-Verlag, Berlin, 2013). <https://doi.org/10.1007/978-3-642-35512-7>
- [30] M. S. Taqqu, "Weak convergence to fractional Brownian motion and to the Rosenblatt process," *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 31, 287–302 (1975). <https://doi.org/10.1007/BF00532868>

Appendix A. Complete Proofs

A.1 Proof of Theorem 4.1 (CEAT)

$S(n, \beta) = n \cdot H(n-1, \beta) - H(n-1, \beta-1)$ (Lemma 4.2). Euler–Maclaurin (Lemma 4.1): for $0 < \beta < 1$, $n \cdot H \sim n^{\{2-\beta\}} / (1-\beta) + n\zeta(\beta)$ and $H(n, \beta-1) \sim n^{\{2-\beta\}} / (2-\beta)$. Subtracting gives $S \sim n^{\{2-\beta\}} / [(1-\beta)(2-\beta)]$. For $\beta > 1$: $H(n, \beta) \rightarrow \zeta(\beta)$, $n\zeta(\beta)$ dominates, $T_n \sim n \cdot [C(0) + 2\zeta(\beta)L]$. $\square$

A.2 Proof of Theorem 5.1 (Two-term remainder)

Carrying the sub-leading term from Lemma 4.1: $n \cdot H(n, \beta) = n^{\{2-\beta\}} / (1-\beta) + n\zeta(\beta) + O(n^{\{1-\beta\}})$. Subtracting $H(n, \beta-1) \sim n^{\{2-\beta\}} / (2-\beta) + O(n^{\{1-\beta\}})$: $S(n, \beta) = n^{\{2-\beta\}} / [(1-\beta)(2-\beta)] + n\zeta(\beta) + O(n^{\{1-\beta\}})$. Multiplying by $2L$ and adding $C(0)n$: $T_n = An^{\{2-\beta\}} + Bn + O(n^{\{1-\beta\}})$. $\square$

Appendix B. Perturbation Bound Verification

Table B.1. $|C_GHZ(\varphi) - \varphi^2/4|$ vs $\varphi^4/48$

φ	$ C_GHZ - \varphi^2/4 $	$\varphi^4/48$	Ratio ($\leq 1 \Rightarrow$ bound holds)
0.1	2.08×10^{-7}	2.08×10^{-7}	0.9997
0.3	1.68×10^{-4}	1.69×10^{-4}	0.9970
0.5	1.29×10^{-3}	1.30×10^{-3}	0.9917
1.0	2.02×10^{-2}	2.08×10^{-2}	0.9673
$\pi/2$	0.1169	0.1268	0.9213

Appendix C. NLS Jacobian and Sample Complexity Derivation

The Jacobian row at n is $[n^\alpha, C \cdot n^\alpha \cdot \ln(n), 1]$. $J^T J$ is positive definite when $\ln(n_i)$ are not all equal — guaranteed when at least two distinct n values are in the design. The sample complexity $\text{bias}(N, \beta) = \epsilon$ is solved numerically: $N(0.5, 0.05) \approx 30$ (theory), ~ 150 (with noise); $N(0.3, 0.05) \approx 30$; $N(0.7, 0.05) \approx 500$.

Appendix D. Examples Where Axiom A2 Fails

Axiom A2 requires $C(k) \sim L k^{-\beta}$ with $0 < \beta \leq 1$. When it fails, the CEAT conclusions change as follows:

Correlation $C(k)$	Asymptotic behaviour	$T_n(C)$ scaling	CEAT applicable?
Exponential: $e^{-k/5}$	Summable: $\sum C(k) < \infty$	$T_n \sim n$ ($\alpha=1.0$)	A2b holds; $\alpha=1$ trivially
White noise: $\delta_{k,0}$	Zero for $k \geq 1$	$T_n \sim n$ ($\alpha=1.0$)	A2b holds; $\alpha=1$, no info
Coherent drift: $C(k)=1$	Not summable, not power-law	$T_n \sim n^2$ ($\alpha=2.0$)	A2 fails; $\alpha=2$ exactly, ZZ model
Oscillatory: $k^{-0.5} \cos(\pi k/2)$	Power law with sign changes	$T_n \sim n$ ($\alpha \approx 1.0$)	Cancellation; A2 fails; α uninformative
Power law $k^{-0.5}$	A2a holds ($\beta=0.5$)	$T_n \sim n^{1.5}$ ($\alpha=1.5$)	YES — reference case

Important observations: (1) Exponential and white noise both give $\alpha \approx 1$ but for different reasons (summable vs zero-correlation). They are distinguishable by fitting quality R^2 . (2) Coherent drift gives $\alpha=2$, the maximum for the correlation model; ZZ coupling can drive α above 2 via the quadratic phase mechanism. (3) Oscillatory correlations cause cancellation and produce an apparently Markovian signature ($\alpha \approx 1$) even when the underlying process has long-range structure. This is the hardest case to distinguish empirically. (4) The failure modes of A2 suggest diagnostic checks beyond α : the intra-session CV_α detects coherent drift; $R^2 < 0.5$ flags Breakdown; comparing α under GHZ vs Clifford circuits distinguishes the mechanisms.