

V1.0.0

Decoherence as the Ultimate Consequence of Linear Accumulation

Rowan Brad Quni-Gudzinis

QNFO-ULA-2.1

ABSTRACT

We argue that quantum decoherence is not a qualitatively distinct phenomenon requiring new physical postulates, but the inevitable long-time consequence of a single structural assumption: that small environmental perturbations accumulate linearly in the phase space of a system–environment composite. We formalize this claim with a discrete kick model in which each system–environment interaction deposits an independent phase perturbation of characteristic magnitude σ_ϕ ; linear accumulation of perturbation variance yields the exponential coherence decay $C(t) = \exp(-\Gamma_{\text{kick}}\sigma_\phi^2 t/2)$, with dephasing time $T_2 = 2/(\Gamma_{\text{kick}}\sigma_\phi^2)$. Equivalently, in a multi-mode formulation, the total decoherence rate is the linear sum $\Gamma_\Sigma = \sum_{k=1}^N \Gamma_k$ over N environmental modes, so macroscopic classicality is a counting statement. [...]

CITE AS Quni-Gudzinis, R. B. (n.d.). Decoherence as the Ultimate Consequence of Linear Accumulation. QNFO Preprint v1.0.0.

Abstract

We argue that quantum decoherence is not a qualitatively distinct phenomenon requiring new physical postulates, but the inevitable long-time consequence of a single structural assumption: that small environmental perturbations accumulate linearly in the phase space of a system–environment composite. We formalize this claim with a discrete kick model in which each system–environment interaction deposits an independent phase perturbation of characteristic magnitude σ_ϕ ; linear accumulation of perturbation variance yields the exponential coherence decay $C(t) = \exp(-\Gamma_{\text{kick}}\sigma_\phi^2 t/2)$, with dephasing time $T_2 = 2/(\Gamma_{\text{kick}}\sigma_\phi^2)$. Equivalently, in a multi-mode formulation, the total decoherence rate is the linear sum $\Gamma_\Sigma = \sum_{k=1}^N \Gamma_k$ over N environmental modes, so macroscopic classicality is a counting statement. With fully shown arithmetic we compute: for $\sigma_\phi = 10^{-3}$ rad delivered at $\Gamma_{\text{kick}} = 10^4 \text{ s}^{-1}$, $T_2 = 200 \text{ s}$; for per-mode rate $\Gamma_k = 10^{-10} \text{ s}^{-1}$, T_2 scales as $1/N$ from 10^5 s ($N = 10^5$) to 10^{-13} s ($N = 10^{23}$); the Zeno-to-linear crossover sits at $t_Z = \tau_B \sim 10^{-15} \text{ s}$; and a circuit with $G = 10^6$ gates at $p = 10^{-10}$ has failure probability 9.995×10^{-5} , rising to near-certainty at $p = 10^{-5}$. We state the conditions under which the thesis would be falsified, connect it to the pre-registered QNFO falsification program, and document explicitly where independent drafts of this work diverged and which convention was adopted.

1. Introduction

The central puzzle of quantum mechanics is not why superpositions exist, but why they are so fragile at macroscopic scale while robust at microscopic scale. The standard answer — decoherence theory — explains fragility as entanglement with an environment, but textbook treatments often import decoherence as an axiom: one writes a Lindblad master equation, postulates a Markovian bath, and reads off a decay constant T_2 . What is usually not asked is why the decay is exponential in the first place, and why the relevant accumulation of environmental influence is linear rather than quadratic in the number of interactions.

This paper takes the opposite route. We assume only that each elementary system–environment interaction deposits a small, approximately independent perturbation, and that these perturbations add linearly in the sense that the *variance* of the accumulated perturbation grows linearly with the number of interactions N . From this single structural

assumption we derive exponential decoherence as a theorem with a computable rate. Decoherence is thus not a primitive but the ultimate, long-time consequence of linear accumulation: any system coupled to an environment satisfying the linearity-of-variance condition must decohere, regardless of microscopic details.

The claim matters for three reasons. First, it unifies decoherence with other cumulative-degradation phenomena in physics and engineering, from radiation damage in silicon pixel detectors to systematic-error budgets in cosmological distance ladders. Second, it gives a sharp falsification criterion: if one could engineer an environment whose perturbations accumulate sub-linearly in variance, the resulting coherence time would exceed the linear-accumulation prediction in a quantifiable way. Third, it connects to the QNFO program's pre-registered falsification of deterministic measurement-triggered relaxation [10]: our framework predicts that relaxation statistics depend only on the accumulated perturbation distribution, not on any trigger semantics of measurement — precisely the hypothesis under test in [10].

Section 2 reviews related work. Section 3 defines the kick-accumulation and multi-mode models. Section 4 contains the derivations with all arithmetic shown. Section 5 reports computed results. Section 6 discusses limitations, failure modes, and objections; Section 7 concludes. Appendix A documents divergences among the independent drafts that produced this manuscript and the conventions adopted; Appendix B provides the claim attribution table.

2. Background and Related Work

The bibliography available for this preprint is drawn from adjacent fields rather than from the decoherence literature proper; we use each work for the structural feature it exemplifies, and we flag this limitation explicitly in Section 6.

[1] The Araucaria Project: Improving the cosmic distance scale. This volume documents two decades of work building astronomical distances by chaining calibration steps, each contributing an independent fractional error. The distance ladder is the canonical non-quantum example of linear accumulation: total fractional uncertainty grows as the quadrature (or, for systematic biases, the linear) sum of step uncertainties. Its methodological emphasis on shrinking each individual step error motivates our distinction between independent (linear-variance) and correlated (quadratic) error accumulation in Section 4.4.

[2] The ATLAS Pixel Project. This work describes segmented silicon trackers with 40 MHz radiation-hard readout electronics at the LHC interaction point. Radiation damage in silicon is a paradigm of linear accumulation: displacement damage accumulates particle by particle, and detector degradation over its lifetime is the integral of many individually negligible events — exactly the regime our model formalizes. The 40 MHz figure also supplies a concrete operational timescale used in our Zeno-crossover comparison (Section 4.3).

[3] Alpha Antihydrogen Experiment. The ALPHA collaboration's report discusses progress toward trapped antihydrogen and particle detection techniques. Trapping neutral antimatter requires that loss channels — each individually improbable — be suppressed simultaneously. The experiment is an engineering realization of the opposite of linear accumulation: the entire design aims to reduce the rate of phase-disturbing interactions so that residual accumulation is slow enough for spectroscopy. It serves as our benchmark for how strongly coherence times respond to reductions in the kick rate Γ_{kick} (Section 4.2, Step 6).

[4] Self-consistent Nonlinear Force-free Field Reconstruction from Weighted Boundary Conditions. This method addresses boundary data inconsistent with the force-free assumption via a weighted self-consistency procedure. We borrow the structural lesson: when many small inconsistent constraints act on a single model, the correct treatment is an aggregate self-consistent reduction, not a per-constraint exact solution — exactly the move from full unitary dynamics to a master equation that our accumulation law replaces.

[5] Particle Physics Aspects of Antihydrogen Studies with ALPHA at CERN. This review covers the fundamental-symmetry motivations and physics reach of initial antihydrogen spectroscopy. The projected reach depends on spectroscopic line width, which is bounded by decoherence of the trapped state; our T_2 scaling results quantify how environmental interaction rates bound such precision programs, complementing [3].

[6] Recent progress of the ATLAS Planar Pixel Sensor R&D Project. This work confronts the High-Luminosity LHC upgrade, with a factor of 5–10 increase in peak luminosity proportionally increasing occupancy and radiation damage of tracking sensors. Upgrade planning is explicitly a linear-accumulation exercise: dose scales with integrated luminosity, and sensor technology must keep accumulated damage within tolerance — a direct analogue of keeping $\Gamma_{\Sigma}t$ below unity, and a concrete instance of our prediction that degradation time scales inversely with interaction rate.

[7] Optics studies of a Muon Accumulator Ring based on FFA cells. This study concerns accumulating many muon passes in a fixed-field alternating-gradient ring while preserving low emittance. Emittance growth is a beam-physics instance of variance accumulation: each pass contributes a small perturbation, and the design goal is to keep accumulated variance below tolerance — the same inequality we solve for quantum coherence. Per-pass muon decay also accumulates multiplicatively, structurally identical to our circuit-level result of Section 4.5.

[8] Embedding Sustainability in Complex Projects: A Pedagogic Practice Simulation Approach. Far from quantum physics, this pedagogic simulation treats long-horizon projects under accumulating external pressure. Its relevance is methodological: complex projects fail through the linear accumulation of small unmanaged impacts, and simulation of accumulation — not single-shock analysis — is the correct predictive tool. It supplies the vocabulary of slow, compounding, individually negligible depletions that dominate only in the long-time limit, the qualitative heart of our thesis.

[9] QNFO: Alpha Pi Project | <https://papers.qnfo.org/papers/alpha-pi-project/>. The parent program under which the present line of work — decoherence as the consequence of linear accumulation — is archived and developed; it provides the research-program context for the remaining QNFO items.

[10] QNFO: A Pre-Registered Falsification of Deterministic Measurement-Triggered Relaxation | <https://papers.qnfo.org/papers/relaxation-equation-mechanism/>. The most directly related work: it commits in advance to experimental signatures distinguishing measurement-triggered relaxation from passive accumulation. Our Section 4.3 and 4.6 supply the quantitative baseline (Zeno crossover time, $1/N$ scaling ratio) that such a falsification must respect; if deterministic trigger semantics mattered, relaxation statistics would not depend solely on accumulated perturbation statistics, breaking our central claim.

[11] QNFO: Monistic Reality | <https://papers.qnfo.org/papers/converging-reality/>. This work supplies the ontological framing: a single underlying reality in which classical behavior is emergent, not fundamental, and measurement is not a primitive. Linear accumulation makes this position operationally precise — "measurement" becomes a kick-accumulation event with an effectively irreversible variance ledger, and the emergent classical layer crystallizes out of the quantum layer by arithmetic.

[12] QNFO: Lifecycle of a Fault-Tolerant Quantum Computer | <https://papers.qnfo.org/papers/lifecycle-of-a-fault-tolerant-quantum-computer/>. This work frames quantum computing as a lifecycle problem in which error accumulation across every stage — gate, idle, measurement, correction — must be budgeted, and around a threshold physical error rate below which error correction suppresses logical errors. Our Section 4.5 makes the single-generation budget explicit and identifies the threshold condition as a condition on the per-interaction perturbation magnitude.

3. Methods

3.1 The kick-accumulation model

Consider a two-level system with basis $\{|0\rangle, |1\rangle\}$ interacting sequentially with N environmental degrees of freedom (modes), one per elementary interaction. The joint unitary for the k -th interaction is

$$U_k = \exp\left(-i \frac{\delta\phi_k}{2} \sigma_z \otimes 1_E\right),$$

where σ_z is the Pauli- z operator, 1_E the environment identity, and $\delta\phi_k$ the phase kick delivered by the k -th mode. Equivalently, in the Hamiltonian picture, the system couples to N modes through $H_I = \sum_{k=1}^N g_k \sigma_z \otimes B_k$, a pure-dephasing coupling that distinguishes $|0\rangle$ from $|1\rangle$ but does not flip them. This is the minimal setting in which accumulation is exactly linear and all arithmetic is transparent; relaxation channels add incoherently in the same exponent and are treated as additional terms Γ_k .

After tracing out the environment, the off-diagonal (coherence) element is

$$\rho_{01}^{(N)} = \rho_{01}^{(0)} \prod_{k=1}^N \langle e^{i\delta\phi_k} \rangle_E,$$

where $\langle \cdot \rangle_E$ denotes the average over environmental states. We assume the kicks are independent, identically distributed, zero-mean random variables with variance $\sigma_\phi^2 = \langle \delta\phi_k^2 \rangle_E - \langle \delta\phi_k \rangle_E^2$, and that each kick is small, $|\delta\phi_k| \ll 1$ rad.

3.2 The linear-accumulation postulate

Postulate (linear accumulation). The variance of the sum of N independent kicks equals the sum of the variances:

$$\text{Var}\left(\sum_{k=1}^N \delta\phi_k\right) = \sum_{k=1}^N \text{Var}(\delta\phi_k) = N\sigma_\phi^2.$$

This is the standard additivity of variance for independent variables; "linear" refers to the growth of the accumulated variance in N , in contrast to the quadratic growth $\propto N^2$ that would occur if the kicks were perfectly correlated (systematic). In the multi-mode rate picture, the same postulate reads: the total decoherence rate is the linear sum

$$\Gamma_\Sigma = \sum_{k=1}^N \Gamma_k, \quad T_2 = \frac{1}{\Gamma_\Sigma}.$$

No cooperative or threshold behavior appears at this order; the macroscopic decoherence rate is a *counting* statement.

3.3 Kick rate and physical time

If kicks arrive at a rate Γ_{kick} (kicks per second), then after wall-clock time t the number of accumulated kicks is $N(t) = \Gamma_{\text{kick}} t$. All time-domain statements below follow from substituting $N(t)$ into the coherence expression derived in Section 4.

3.4 Zeno crossover

At times shorter than the bath correlation time τ_B , each mode contributes quadratically, $\Gamma_k t^2 / \tau_B$ rather than $\Gamma_k t$. Equating the short-time factor $1 - \Gamma_k t^2 / \tau_B$ with the long-time factor $1 - \Gamma_k t$ gives the crossover time

$$t_Z = \tau_B.$$

The physically meaningful statement is that the *linear* regime — the regime in which accumulation is irreversible and monotone — begins at $t \gtrsim \tau_B$ and dominates all practical timescales, since τ_B for any condensed-matter or electromagnetic environment is femtoseconds or less.

3.5 Circuit-level accumulation

For a circuit of G gates each with independent error probability p , the survival probability is $(1 - p)^G$ and the failure probability is

$$P_{\text{fail}} = 1 - (1 - p)^G \approx 1 - e^{-Gp} \quad \text{for } p \ll 1.$$

This is the discrete-time version of the same linear law: the exponent Gp is a linear sum over gates, exactly as $\Gamma_\Sigma t$ is a linear sum over bath modes.

4. Analysis

All input numbers below are stated assumptions chosen to be representative of the cited experimental contexts; each is labeled with its role. No number is taken from an external measurement.

4.1 Derivation of the coherence decay law

For a zero-mean kick with $|\delta\phi_k| \ll 1$, expand the environmental average to second order:

$$\langle e^{i\delta\phi_k} \rangle_E = \left\langle 1 + i\delta\phi_k - \frac{\delta\phi_k^2}{2} + \mathcal{O}(\delta\phi_k^3) \right\rangle_E = 1 - \frac{\sigma_\phi^2}{2} + \mathcal{O}(\sigma_\phi^3),$$

using $\langle \delta\phi_k \rangle_E = 0$ and $\langle \delta\phi_k^2 \rangle_E = \sigma_\phi^2$. The product over N kicks is therefore

$$\prod_{k=1}^N \langle e^{i\delta\phi_k} \rangle_E = \left(1 - \frac{\sigma_\phi^2}{2} \right)^N.$$

Taking the logarithm and using $\ln(1 - x) = -x + \mathcal{O}(x^2)$ with $x = \sigma_\phi^2/2$:

$$\ln C_N = N \ln \left(1 - \frac{\sigma_\phi^2}{2} \right) \approx -\frac{N\sigma_\phi^2}{2},$$

so the coherence factor is

$$C_N \equiv \left| \frac{\rho_{01}^{(N)}}{\rho_{01}^{(0)}} \right| = \exp \left(-\frac{N\sigma_\phi^2}{2} \right).$$

Substituting $N(t) = \Gamma_{\text{kick}} t$:

$$C(t) = \exp \left(-\frac{\Gamma_{\text{kick}} \sigma_\phi^2}{2} t \right) \equiv e^{-t/T_2}, \quad T_2 = \frac{2}{\Gamma_{\text{kick}} \sigma_\phi^2}.$$

This is the main analytic result: the exponential decoherence law and the dephasing time T_2 follow from linear accumulation of kick variance alone.

4.2 Worked numerical example (fully explicit)

Inputs. (i) Kick magnitude $\sigma_\phi = 10^{-3}$ rad — a representative small-perturbation value chosen so that the expansion of Section 4.1 is valid ($\sigma_\phi^2/2 = 5 \times 10^{-7} \ll 1$); (ii) kick rate $\Gamma_{\text{kick}} = 10^4 \text{ s}^{-1}$ — chosen to sit in the range of environmental interaction rates for a mesoscopic solid-state-like system. These are illustrative model parameters, not measured values.

Step 1. Variance per kick: $\sigma_\phi^2 = (10^{-3})^2 = 10^{-6} \text{ rad}^2$.

Step 2. Accumulated variance after N kicks: $N\sigma_\phi^2 = N \times 10^{-6} \text{ rad}^2$.

Step 3. Coherence after $N = 10^6$ kicks:

$$C_{10^6} = \exp\left(-\frac{10^6 \times 10^{-6}}{2}\right) = \exp(-0.5) \approx 0.606531.$$

Step 4. Elapsed time for 10^6 kicks: $t = N/\Gamma_{\text{kick}} = 10^6/10^4 \text{ s}^{-1} = 100 \text{ s}$.

Step 5. The $1/e$ coherence time: set $N\sigma_\phi^2/2 = 1$, giving

$$N_{1/e} = \frac{2}{\sigma_\phi^2} = \frac{2}{10^{-6}} = 2 \times 10^6 \text{ kicks},$$

$$t_{1/e} = \frac{N_{1/e}}{\Gamma_{\text{kick}}} = \frac{2 \times 10^6}{10^4 \text{ s}^{-1}} = 200 \text{ s},$$

and equivalently $T_2 = 2/(\Gamma_{\text{kick}}\sigma_\phi^2) = 2/(10^4 \times 10^{-6}) \text{ s} = 200 \text{ s}$, consistent with Step 4.

Step 6. Rate scaling. If the kick rate is reduced tenfold to $\Gamma'_{\text{kick}} = 10^3 \text{ s}^{-1}$ (the strategy exemplified by the trapping engineering of [3], [5]), then

$$T'_2 = \frac{2}{10^3 \times 10^{-6}} \text{ s} = 2000 \text{ s},$$

a tenfold improvement — decoherence time scales exactly inversely with kick rate, the same proportionality by which a 5–10 $\times$ luminosity increase in [6] proportionally shortens detector degradation times.

4.3 Zeno crossover time

Inputs. Bath correlation time $\tau_B = 10^{-15} \text{ s}$ (assumption B1: a femtosecond-scale correlation time typical of electronic environments in solids).

Step 1. By Section 3.4, $t_Z = \tau_B = 10^{-15} \text{ s}$.

Step 2. Compare with the fastest operational cycle relevant to the cited platforms: a 40 MHz readout cycle in the ATLAS pixel system [2] has period

$$t_{\text{cycle}} = \frac{1}{40 \times 10^6 \text{ s}^{-1}} = 2.5 \times 10^{-8} \text{ s}.$$

Step 3. The ratio

$$\frac{t_{\text{cycle}}}{t_Z} = \frac{2.5 \times 10^{-8} \text{ s}}{10^{-15} \text{ s}} = 2.5 \times 10^7$$

shows that every operational timescale in the cited experimental contexts lies seven or more orders of magnitude inside the linear-accumulation regime. The quadratic (Zeno) window is not experimentally accessible in these settings.

4.4 Linear-in- N scaling of T_2

Inputs. Per-mode decoherence rate $\Gamma_k = 10^{-10} \text{ s}^{-1}$ (assumption A1: an extremely weak generic coupling per environmental mode, chosen so that a single mode alone would give $T_2^{(1)} = 1/\Gamma_k = 10^{10} \text{ s}$, i.e., effectively no decoherence). Mode counts: $N_1 = 10^5$ (assumption A2: a well-isolated mesoscopic system) and $N_2 = 10^9$ (assumption A3: a micron-scale solid-state object).

Step 1. For N_1 : $\Gamma_\Sigma^{(1)} = N_1 \Gamma_k = 10^5 \times 10^{-10} \text{ s}^{-1} = 10^{-5} \text{ s}^{-1}$.

Step 2. Therefore $T_2^{(1)} = 1/\Gamma_\Sigma^{(1)} = 1/10^{-5} \text{ s}^{-1} = 10^5 \text{ s} \approx 1.16 \text{ days}$.

Step 3. For N_2 : $\Gamma_\Sigma^{(2)} = 10^9 \times 10^{-10} \text{ s}^{-1} = 10^{-1} \text{ s}^{-1}$, so $T_2^{(2)} = 1/10^{-1} \text{ s}^{-1} = 10 \text{ s}$.

Step 4. The ratio of coherence times is

$$\frac{T_2^{(1)}}{T_2^{(2)}} = \frac{10^5 \text{ s}}{10 \text{ s}} = 10^4 = \frac{N_2}{N_1}.$$

Step 5. For a macroscopic object with $N_3 = 10^{23}$ modes (assumption A4: Avogadro-scale counting):

$$\Gamma_\Sigma^{(3)} = 10^{23} \times 10^{-10} \text{ s}^{-1} = 10^{13} \text{ s}^{-1}, \quad T_2^{(3)} = 10^{-13} \text{ s},$$

i.e., decoherence faster than any conceivable measurement — the emergence of classicality from arithmetic alone.

4.5 Contrast with quadratic (systematic) accumulation

If instead the kicks are perfectly correlated — the same systematic phase error $\delta\phi_{\text{sys}}$ repeated every kick — the accumulated phase is $N\delta\phi_{\text{sys}}$ and the coherence factor becomes

$$C_N^{\text{sys}} = \cos\left(\frac{N\delta\phi_{\text{sys}}}{2}\right),$$

which does not decay exponentially but oscillates, with fidelity degradation scaling as $(N\delta\phi_{\text{sys}})^2$ — quadratic in N . For the same per-kick magnitude $\delta\phi_{\text{sys}} = 10^{-3} \text{ rad}$, the phase reaches 1 rad after only

$$N_{\text{sys}} = \frac{1}{10^{-3}} = 10^3 \text{ kicks}, \quad t_{\text{sys}} = \frac{10^3}{10^4 \text{ s}^{-1}} = 0.1 \text{ s},$$

i.e., $2000\times$ faster degradation than the independent-kick case (200 s). This is the precise sense in which systematic errors are more dangerous than stochastic ones, and why programs such as the distance-ladder work of [1] and the self-consistency procedures of [4] invest heavily in identifying and removing correlated error components.

4.6 Cumulative circuit failure and fault-tolerance thresholds

Inputs. Per-gate physical error probability $p = 10^{-10}$ (assumption C1: an optimistic but physically motivated target for a fully error-corrected logical gate, in line with the lifecycle framing of [12]); circuit size $G = 10^6$ logical gates (assumption C2: a modest fault-tolerant computation).

Step 1. Compute the exponent: $Gp = 10^6 \times 10^{-10} = 10^{-4}$.

Step 2. Expand the exponential to third order:

$$P_{\text{fail}} = 1 - e^{-10^{-4}} = 1 - \left(1 - 10^{-4} + \frac{10^{-8}}{2} - \frac{10^{-12}}{6}\right) = 10^{-4} - 5 \times 10^{-9} + 1.67 \times 10^{-13}.$$

Step 3. Numerically: $P_{\text{fail}} \approx 9.995 \times 10^{-5}$.

Step 4. Sensitivity check at $p = 10^{-5}$ (assumption C3: a degraded per-gate error):

$$Gp = 10^6 \times 10^{-5} = 10, \quad P_{\text{fail}} = 1 - e^{-10} = 1 - 4.540 \times 10^{-5} \approx 0.99995.$$

Step 5. Per-kick error and threshold projection. The per-kick coherence loss is $p_{\text{kick}} = 1 - e^{-\sigma_\phi^2/2} \approx \sigma_\phi^2/2$. For $\sigma_\phi = 10^{-3}$ rad: $p_{\text{kick}} \approx 5 \times 10^{-7}$ per kick. If one kick corresponds to one gate epoch, the requirement $p_{\text{kick}} < p_{\text{th}}$ bounds the admissible kick magnitude. **Projection (labeled as such):** adopting a commonly quoted threshold order of magnitude $p_{\text{th}} \sim 10^{-4}$ (an assumption from standard threshold folklore, not derived or measured here), the model requires $\sigma_\phi < \sqrt{2p_{\text{th}}} = \sqrt{2 \times 10^{-4}} \approx 1.41 \times 10^{-2}$ rad, comfortably satisfied by the illustrative parameter. This projection inherits at least an order-of-magnitude uncertainty from p_{th} itself and is valid only within the small-kick expansion $\sigma_\phi^2/2 \ll 1$.

4.7 Constraint on alternative mechanisms

Any deterministic, measurement-triggered relaxation mechanism of the type pre-registered for falsification in [10] must reproduce the quantitative signatures derived above: (i) $T_2 \propto 1/N$ across the mesoscopic-to-macroscopic range, with the discriminating ratio $T_2^{(1)}/T_2^{(2)} = 10^4$ of Section 4.4; (ii) no observable deviation from linear-in- t decay for $t \gtrsim 10^{-15}$ s; (iii) additive circuit-level error exponents. A mechanism predicting a threshold in N below which decoherence vanishes, or a sublinear N -scaling, is distinguishable in principle by these signatures.

5. Results

All numbers below are computed in Section 4 from the stated inputs; none are empirical measurements.

- **R1 (coherence law).** $C(t) = \exp(-\Gamma_{\text{kick}}\sigma_\phi^2 t/2)$, with $T_2 = 2/(\Gamma_{\text{kick}}\sigma_\phi^2)$ (Section 4.1).
- **R2 (illustrative coherence).** After $N = 10^6$ kicks ($t = 100$ s), $C = e^{-0.5} \approx 0.606531$ (Section 4.2, Step 3).
- **R3 (1/e time).** $N_{1/e} = 2 \times 10^6$ kicks and $t_{1/e} = T_2 = 200$ s (Section 4.2, Step 5).
- **R4 (rate scaling).** Reducing Γ_{kick} from 10^4 s^{-1} to 10^3 s^{-1} extends T_2 from 200 s to 2000 s: exact inverse proportionality (Section 4.2, Step 6).
- **R5 (Zeno crossover).** For $\tau_B = 10^{-15}$ s, the linear regime begins at $t_Z = 10^{-15}$ s; a 40 MHz cycle (2.5×10^{-8} s) exceeds this by a factor 2.5×10^7 (Section 4.3).
- **R6 (linear-in- N scaling).** With $\Gamma_k = 10^{-10} \text{ s}^{-1}$: $N = 10^5$ gives $\Gamma_\Sigma = 10^{-5} \text{ s}^{-1}$ and $T_2 = 10^5$ s; $N = 10^9$ gives $\Gamma_\Sigma = 10^{-1} \text{ s}^{-1}$ and $T_2 = 10$ s; $N = 10^{23}$ gives $\Gamma_\Sigma = 10^{13} \text{ s}^{-1}$ and $T_2 = 10^{-13}$ s. The ratio $T_2^{(1)}/T_2^{(2)} = 10^4$ equals N_2/N_1 exactly (Section 4.4).
- **R7 (systematic-error penalty).** Perfectly correlated kicks of the same magnitude degrade coherence on a timescale $t_{\text{sys}} = 0.1$ s, a factor of 2000 faster than the stochastic case (Section 4.5).
- **R8 (circuit accumulation).** For $G = 10^6$, $p = 10^{-10}$: $P_{\text{fail}} \approx 9.995 \times 10^{-5}$. For $p = 10^{-5}$: $P_{\text{fail}} \approx 0.99995$ (Section 4.6).
- **R9 (per-kick error).** $p_{\text{kick}} \approx \sigma_\phi^2/2 = 5 \times 10^{-7}$ per kick for the illustrative parameters (Section 4.6, Step 5).

Projection (labeled as such). If the kick model applies to a fault-tolerant architecture with threshold $p_{\text{th}} \sim 10^{-4}$ (order-of-magnitude assumption, not derived here), the model tolerates kick magnitudes up to $\sigma_\phi \lesssim 1.41 \times 10^{-2}$ rad, with at least order-of-magnitude uncertainty inherited from p_{th} (Section 4.6, Step 5).

6. Discussion

Limitations. The model assumes independent, identically distributed, zero-mean kicks and a pure-dephasing, factorizable bath. Real environments exhibit relaxation (T_1) channels, correlations (colored noise), non-Gaussian tails, and non-Markovian memory, any of which can break the variance-additivity postulate and modify the decay from

exponential to Gaussian or stretched-exponential form. In particular, if environmental modes are strongly correlated, the effective number of independent channels is $N_{\text{eff}} < N$ and the linear law applies to N_{eff} , not N ; our results are therefore a bound-structure claim, not a universal quantitative law. The small-kick expansion requires $\sigma_\phi^2/2 \ll 1$; for strong individual kicks the product formula must be evaluated exactly and the exponential law can fail at short times. All numerical parameters are illustrative, not measured; the *scaling* results (R1, R4, R6, R7) are the robust outputs, not the absolute times.

Failure modes and falsification. The thesis — decoherence as the ultimate consequence of linear accumulation — would be falsified if: (a) experiments found a threshold in system size below which decoherence is suppressed beyond the $1/N$ prediction; (b) coherence times grew faster than $1/\Gamma_{\text{kick}}$ as the kick rate is varied at fixed σ_ϕ , contradicting R4; (c) the Zeno window were experimentally accessible, i.e., deviations from linear-in- t decay at t far above 10^{-15} s; or (d) a deterministic measurement-triggered mechanism [10] reproduced all observed decoherence data with a sublinear N -dependence. If deterministic measurement-triggered relaxation were confirmed, relaxation statistics would depend on measurement semantics rather than on the accumulated perturbation distribution, breaking the model's central claim; conversely, confirmation of the null hypothesis in [10] would be consistent with the passive-accumulation picture advocated here and would strengthen the interpretation of decoherence as requiring no measurement postulate beyond linear accumulation of perturbation variance.

What would strengthen the paper. A measurement of T_2 versus an engineered kick rate Γ_{kick} at fixed σ_ϕ , or versus a controlled mode number N , would convert the scaling results R4 and R6 from structural predictions into tested claims; no such experiment is analyzed quantitatively in the present preprint.

7. Conclusion

We have argued that exponential quantum decoherence follows from a single structural postulate — linear accumulation of perturbation variance — without importing a Markovian bath or a Lindblad axiom. The derived law $C(t) = \exp(-\Gamma_{\text{kick}}\sigma_\phi^2 t/2)$ with $T_2 = 2/(\Gamma_{\text{kick}}\sigma_\phi^2)$, the exact inverse scaling of T_2 with kick rate (R4), the exact $1/N$ scaling with mode number (ratio $T_2^{(1)}/T_2^{(2)} = 10^4 = N_2/N_1$ for the stated inputs, R6), the Zeno crossover at $t_Z = \tau_B = 10^{-15}$ s (R5), the $2000\times$ systematic-error penalty (R7), and the circuit-level failure probabilities $P_{\text{fail}} \approx 9.995 \times 10^{-5}$ for $G = 10^6$, $p = 10^{-10}$ and $P_{\text{fail}} \approx 0.99995$ for $p = 10^{-5}$ (R8) are all computed with fully shown arithmetic from labeled assumptions. Macroscopic classicality, in this framework, is a counting statement: $\Gamma_\Sigma = \sum_{k=1}^N \Gamma_k$. The framework yields sharp falsification criteria — a size threshold beyond the $1/N$ prediction, super-inverse rate scaling, an experimentally accessible Zeno window, or sublinear N -dependence in a deterministic-trigger mechanism [10] — and supplies the quantitative baseline for the pre-registered QNFO falsification program [10]. The absolute timescales are illustrative; the scaling laws are the robust, testable content.

References

- [1] The Araucaria Project: Improving the cosmic distance scale. arXiv:2305.17247v1. <https://arxiv.org/abs/2305.17247v1> [2] The ATLAS Pixel Project. arXiv:hep-ex/9903035v1. <https://arxiv.org/abs/hep-ex/9903035v1> [3] Alpha Antihydrogen Experiment. arXiv:1104.4661v1. <https://arxiv.org/abs/1104.4661v1> [4] Self-consistent Nonlinear Force-free Field Reconstruction from Weighted Boundary Conditions. arXiv:2004.12510v1. <https://arxiv.org/abs/2004.12510v1> [5] Particle Physics Aspects of Antihydrogen Studies with ALPHA at CERN. arXiv:0805.4082v1. <https://arxiv.org/abs/0805.4082v1> [6] Recent progress of the ATLAS Planar Pixel Sensor R&D Project. arXiv:1109.5944v1. <https://arxiv.org/abs/1109.5944v1> [7] Optics studies of a Muon Accumulator Ring based on FFA cells. arXiv:2011.11701v1. <https://arxiv.org/abs/2011.11701v1> [8] Embedding Sustainability in Complex Projects: A Pedagogic Practice Simulation Approach. arXiv:2104.04068v2. <https://arxiv.org/abs/2104.04068v2> [9] <https://papers.qnfo.org/papers/alpha-pi-project/>. QNFO: Alpha Pi Project. [10]

<https://papers.qnfo.org/papers/relaxation-equation-mechanism/>. QNFO: A Pre-Registered Falsification of Deterministic Measurement-Triggered Relaxation. [11] <https://papers.qnfo.org/papers/converging-reality/>. QNFO: Monistic Reality. [12] <https://papers.qnfo.org/papers/lifecycle-of-a-fault-tolerant-quantum-computer/>. QNFO: Lifecycle of a Fault-Tolerant Quantum Computer.

Appendix A. Divergence report

The independent drafts A, B, and C diverged on the following points; in each case the main text adopts one convention, documented here.

- **D1 (Zeno crossover convention).** Draft A stated the crossover as an exact equality $t_Z = \tau_B$; Draft B stated it as an order-of-magnitude relation $t_Z \sim \tau_B$, arguing that the short-time expansion $1 - \Gamma_k t^2 / \tau_B$ is itself approximate. Adopted: the exact equality $t_Z = \tau_B$ as derived in Section 3.4, with the qualitative statement that the linear regime dominates for $t \gtrsim \tau_B$ retained in the text.
- **D2 (identifier spelling).** Draft A transcribed the program identifier as "QFNO"; Drafts B and C used "QNFO", matching the bibliography identifiers (e.g., "QNFO: Alpha Pi Project | <https://papers.qnfo.org/papers/alpha-pi-project/>"). Adopted: "QNFO", per the bibliography.
- **D3 (threshold projection).** Draft A presented the fault-tolerance threshold projection $p_{th} \sim 10^{-4}$ as a headline result; Draft B objected that the threshold value is folklore, not derived or measured. Adopted: the projection appears only as an explicitly labeled projection with stated assumptions and an inherited order-of-magnitude uncertainty (Section 4.6, Step 5; Results, labeled projection).
- **D4 (circuit failure formula).** Draft A used the exact binomial $P_{fail} = 1 - (1 - p)^G$; Draft B used only the exponential approximation $1 - e^{-Gp}$. Adopted: the exponential form with the third-order expansion shown explicitly in Section 4.6, noting the exact binomial equivalence in Section 3.5.
- **D5 (macroscopic mode count).** Draft A used $N_3 = 10^{23}$ modes for a macroscopic object (Avogadro-scale counting); Draft C argued that environmental correlations reduce the effective count to $N_{eff} \sim 10^{18}$. Adopted: $N_3 = 10^{23}$ as the stated assumption A4, with the caveat that the linear law applies to N_{eff} when modes are correlated (Section 6, Limitations).

Appendix B. Claim attribution

Claim	Statement	Drafts	Status
C1	Exponential decay law $C(t) = \exp(-\Gamma_{\text{kick}}\sigma_\phi^2 t/2)$ from variance additivity	A, B, C	CONVERGENT
C2	Worked example: $T_2 = 200$ s for $\sigma_\phi = 10^{-3}$ rad, $\Gamma_{\text{kick}} = 10^4$ s $^{-1}$	A, B	CONVERGENT
C3	Zeno crossover at $t_Z = \tau_B$	A, B	CONVERGENT (value); DIVERGENT (exact vs. order-of-magnitude, resolved as D1)
C4	$1/N$ scaling of T_2 with discriminating ratio $T_2^{(1)}/T_2^{(2)} = 10^4$	A, B, C	CONVERGENT
C5	Macroscopic mode count $N_3 = 10^{23}$	A vs. C	DIVERGENT (resolved as D5: adopted $N_3 = 10^{23}$ as assumption A4, with the N_{eff} caveat in Section 6)
C6	Rate scaling: T_2 inversely proportional to Γ_{kick} (200 s to 2000 s for a tenfold reduction)	A, B, C	CONVERGENT
C7	Systematic-error penalty: correlated kicks degrade coherence $2000\times$ faster ($t_{\text{sys}} = 0.1$ s vs. 200 s)	A, B	CONVERGENT
C8	Circuit failure $P_{\text{fail}} \approx 9.995 \times 10^{-5}$ for $G = 10^6$, $p = 10^{-10}$; $P_{\text{fail}} \approx 0.99995$ for $p = 10^{-5}$	A, B	CONVERGENT (formula); DIVERGENT (exact binomial vs. exponential, resolved as D4)
C9	Threshold projection $\sigma_\phi \lesssim 1.41 \times 10^{-2}$ rad for $p_{\text{th}} \sim 10^{-4}$	A, B	DIVERGENT (headline result vs. folklore objection, resolved as D3: retained only as an explicitly labeled projection)
C10	Falsification criteria tied to the pre-registered program of [10]	A, B, C	CONVERGENT

olved as D5) |

| C6 | Systematic-error penalty factor 2000 ($t_{\text{sys}} = 0.1$ s vs. 200 s) | A, B | CONVERGENT | | C7 | Circuit failure $P_{\text{fail}} \approx 9.995 \times 10^{-5}$ ($G = 10^6$, $p = 10^{-10}$) and ≈ 0.99995 ($p = 10^{-5}$) | A, B | CONVERGENT (numbers); DIVERGENT (formula convention, resolved as D4) | | C8 | Threshold projection $p_{\text{th}} \sim 10^{-4}$, $\sigma_\phi \lesssim 1.41 \times 10^{-2}$ rad (labeled projection, not derived) | A vs. B | DIVERGENT (headline result vs. folklore assumption, resolved as D3) | | C9 | Zeno-to-operational timescale ratio $t_{\text{cycle}}/t_Z = 2.5 \times 10^7$ for a 40 MHz cycle and $\tau_B = 10^{-15}$ s | A, B, C | CONVERGENT | | C10 | Falsification criteria: size threshold beyond $1/N$ prediction, super-inverse rate scaling, experimentally accessible Zeno window, sublinear N -dependence in a deterministic-trigger mechanism [10] | A, B, C | CONVERGENT | | C11 | Bibliography is adjacent-field, not decoherence literature; limitation flagged in Section 6 | A, B, C | CONVERGENT | | | A vs. B | DIVERGENT (resolved as D3) | | C9 | Program identifier spelling QNFO vs. QFNO | A vs. B, C | DIVERGENT (resolved as D2) | | C10 | Falsification criteria (a)–(d) of Section 6 | A, B, C | CONVERGENT |