

Relaxation in Noisy Quantum Maps versus Classical Iterative Algorithms: A Shared Decay Law with a Divergent Semiclassical Limit

Relaxation in Noisy Quantum Maps versus Classical Iterative Algorithms: A Shared Decay Law with a Divergent Semiclassical Limit

Abstract

Iterative numerical algorithms and noisy quantum dynamics both converge to a stationary state by exponential damping, and practitioners who port workloads between classical solvers and quantum co-processors need to know whether the two relaxation processes obey the same or different timing laws. The specific gap addressed here is that relaxation time has been formalized for noisy quantum maps on tori, where it was shown to be sensitive to classical chaos in a joint semiclassical and weak-noise limit, while the corresponding convergence-time laws of classical stationary iterations are usually analyzed separately, so no side-by-side scaling comparison with explicit arithmetic exists. We formulate both processes as linear contractions toward a fixed point, derive a common step-count law $t_\delta = \ln(\delta/D_0)/\ln(1 - \varepsilon)$ for depolarizing strength ε and tolerance δ , and evaluate it explicitly: at $\varepsilon = 0.01$ and $\delta = 10^{-6}$ both a qubit under depolarizing noise and a classical stationary iteration with convergence factor 0.99 need 1374.6 steps. We further derive that a logical fidelity threshold of 0.874, taken from a published quantum-error-correction and Quantum-Darwinism no-go result, is crossed after 28.9 steps at $\varepsilon = 0.01$. The principal result is that the finite- ε relaxation law does not distinguish quantum from classical relaxation; any operational difference must enter through the joint $\hbar \rightarrow 0, \varepsilon \rightarrow 0$ limit. This redirects engineering attention from per-step decay rates to limit-ordering questions and yields a concrete falsification protocol.

1. Introduction

Computational facilities that mix classical solvers with quantum processors face a scheduling problem that is currently solved by guesswork. A relaxation-type workload — driving a system from an arbitrary initial state to a target stationary or near-stationary state — can run as a classical stationary iteration on a CPU cluster or as a repeatedly applied noisy quantum map on a quantum device, and the operator must decide how many steps, and therefore how much wall-clock time and energy, each route needs. If the two routes obey different relaxation laws, one route can be ruled out by scaling alone before any hardware is bought; if they obey the same finite-step law, the decision hinges on quantities such as noise strength and limit ordering instead. Today no such side-by-side law is available in the literature in explicitly comparable form, so procurement and scheduling decisions rest on unquantified intuitions about "quantum speedup".

The cost of this gap is concrete. A team that assumes a quadratic-type quantum advantage for a relaxation task, by analogy with phase and amplitude estimation where a quadratic speedup is known, may commit to hardware whose actual relaxation time scales no better than a classical iteration at the relevant noise level. Conversely, a team that assumes equivalence may miss regimes, identified in the quantum-maps literature, where relaxation time depends on the chaotic character of an underlying classical system and therefore cannot be matched by any classical iteration with a fixed convergence factor.

This paper answers the question: does relaxation in noisy quantum maps differ from relaxation in classical algorithms? We answer it by constructing the simplest common mathematical frame — linear contraction toward a fixed point — deriving the step-count law both processes obey at fixed noise, evaluating that law with full

arithmetic, and then identifying the one place, the joint semiclassical and weak-noise limit, where the quantum case is known to behave differently. We also connect the relaxation clock to an independently published fidelity threshold, and we state a pre-registered falsification condition in the style of prior pre-registered falsification work on measurement-triggered relaxation.

The contributions are:

- A unified contraction-frame derivation of the relaxation step-count law $t_\delta = \ln(\delta/D_0)/\ln(1 - \varepsilon)$, valid identically for a depolarizing qubit channel and for a classical stationary iteration with convergence factor $1 - \varepsilon$, with explicit evaluation at $\varepsilon = 0.01$, $\delta = 10^{-6}$ giving 1374.6 steps (Section 4.1).
- A crossover analysis showing the law is insensitive to ε at the logarithmic level: reducing ε by a factor of ten multiplies the step count by only about 10.04, computed explicitly (Section 4.2).
- A coupling of the relaxation clock to the logical-fidelity threshold $F_L > 0.874$ from the quantum-error-correction and Quantum-Darwinism no-go theorem, yielding 28.9 steps at $\varepsilon = 0.01$ before the threshold is crossed (Section 4.3).
- A limit-ordering argument, grounded in the relaxation-time results for quantized toral maps, that locates the only known quantum-classical divergence in the joint $\hbar \rightarrow 0$, $\varepsilon \rightarrow 0$ limit, and a corresponding falsification criterion (Sections 4.4 and 6).

2. Background and Related Work

The closest formal antecedent to the present question is the relaxation time introduced for noisy quantum maps on the $2d$ -dimensional torus, which generalizes an earlier notion of dissipation time [6]. That work shows that relaxation time is sensitive to the chaotic behavior of the corresponding classical system when one simultaneously takes the semiclassical limit ($\hbar \rightarrow 0$) and the limit of small noise strength ($\varepsilon \rightarrow 0$). Two features of this result drive our analysis: first, relaxation time is a property of the map's approach to stationarity, exactly the quantity a classical numerical analyst calls convergence time; second, the sensitivity to chaos appears only in a joint limit, which means that at fixed $\hbar$ and fixed ε no statement of difference is available from that result. Our paper supplies the fixed- ε comparison that this result deliberately does not address, and then uses the joint-limit sensitivity as the candidate locus of genuine divergence.

On the classical side, the parallel Jacobi family of one-sided trigonometric and hyperbolic Jacobi-type algorithms for Hermitian eigenvalue problems provides the canonical example of an iterative classical relaxation process [4]. The supplied summary states that these algorithms compute eigenvalues and eigenvectors with high relative accuracy when the matrices permit, and that novel parallel variants are presented; the summary gives no convergence-rate numbers, so we use Jacobi iteration only as the named representative of the class of stationary iterations whose convergence factor ρ we treat as a free parameter. Our derivation in Section 4.1 is therefore a comparison of structural laws, not of specific implementations.

The quantum side of the speedup question is bounded by known results on phase and amplitude estimation, which admit a quadratic quantum speedup [5]. The same source records that a continuous trade-off between quantum speedup and circuit depth exists through a family of algorithms known as α -QPE, and that the scaling of α -QPE can be derived within a quantum singular value transformation framework. This matters here because it shows that even for the best-understood quantum speedups the advantage is not a step-count constant but a trade-off curve; a relaxation comparison that reports only asymptotic exponents would therefore overstate the practical difference. Our analysis keeps the comparison at the level of explicit step counts at fixed parameters.

A hard constraint on any quantum relaxation scheme that also aims to preserve quantum information comes from the QEC-Darwinism tradeoff: it has been proved that quantum error correction and Quantum Darwinism cannot coexist above a critical logical fidelity $F_L > 0.874$, described as a tight, model-independent no-go theorem establishing an exact tradeoff between protected quantum information and emergent classical objectivity [11]. We import the single number 0.874 from this result and ask how many relaxation steps at a given noise strength separate a pure initial state from this threshold; the arithmetic appears in Section 4.3. We do not extend or re-derive the no-go theorem itself.

The methodology of stating a falsification condition in advance is taken from pre-registered falsification work on deterministic measurement-triggered relaxation [12]. The supplied entry for this work gives no abstract text beyond the title, so we use it only in the limited role its title supports: as precedent for pre-registering a falsification criterion for a relaxation mechanism, which is the role our Section 4.4 and Section 6 adopt. Similarly, the QNFO corpus item "Programming the Vacuum" is supplied without any abstract text [9]; we note its presence in the corpus but draw no claim from it, because no content is available to support one.

Three further works shape the engineering framing. The distributed maximum-clique work defines the "costup" — a significant speedup achieved with minimal development effort — and obtains it by a simple modification of a state-of-the-art exact algorithm that allows distribution over an ad-hoc cluster of machines [1]. We adopt costup as the figure of merit for any proposed quantum-relaxation deployment: a route wins only if its step-count advantage survives the development effort of realizing it. The submodular maximization literature shows a complementary pattern: a problem with polynomial-time approximation algorithms still requires, in practice, frequently querying the submodular function, so query cost dominates runtime [2]; the analogue for relaxation is that per-step cost, not step count alone, determines wall-clock time. The TSP-with-neighborhoods work illustrates that generalizations of classical problems (there, Euclidean TSP to TSPN, which is NP-hard) admit constant-factor and other approximation results for structured cases [3]; by analogy, relaxation-time results for structured maps (tori) need not transfer to unstructured instances, a limitation we carry into Section 6. Finally, the consensus-tree algorithms show that optimal $O(kn)$ constructions are attainable and provably optimal because the input size is $\Omega(kn)$ [7], and the odd-gracefulness work shows that for graph labeling there are few algorithms despite many published papers on the labels themselves [8]; both are reminders that existence of a quantity (a consensus tree, a graceful labeling, a relaxation time) does not imply availability of an efficient procedure to compute or control it, which is precisely the gap this paper addresses for relaxation-time comparison.

3. Methods

3.1 Common contraction frame

Both processes under comparison are discrete-time linear contractions toward a fixed point. Let x_t be the state after t steps, x_* the fixed point, and suppose

$$\|x_{t+1} - x_*\| \leq \rho \|x_t - x_*\|,$$

with contraction factor $0 \leq \rho < 1$. Iterating gives

$$\|x_t - x_*\| \leq \rho^t \|x_0 - x_*\|.$$

Define the initial distance $D_0 = \|x_0 - x_*\|$ and a tolerance δ with $0 < \delta < D_0$. The step count t_δ at which the distance first falls to δ satisfies $\rho^{t_\delta} D_0 = \delta$, hence

$$t_\delta = \frac{\ln(\delta/D_0)}{\ln \rho}.$$

This frame is method-level and parameter-level only; it assumes linearity of the contraction, time-independence of ρ , and a norm in which the contraction inequality holds. No hardware, dataset, or simulation is used anywhere in this paper; every number in Section 5 is computed by hand arithmetic in Section 4 from the stated inputs.

3.2 Quantum instantiation: depolarizing noise on a qubit

For a single qubit undergoing a depolarizing channel of strength ε per step, the state updates as

$$\rho_{t+1} = (1 - \varepsilon) \rho_t + \varepsilon \frac{\mathbb{I}}{2},$$

so that after t steps

$$\rho_t = (1 - \varepsilon)^t \rho_0 + (1 - (1 - \varepsilon)^t) \frac{\mathbb{I}}{2}.$$

The maximally mixed state $\mathbb{I}/2$ is the fixed point, and the trace distance from the fixed point contracts as

$$D(\rho_t, \mathbb{I}/2) = (1 - \varepsilon)^t D(\rho_0, \mathbb{I}/2),$$

which is the contraction frame with $\rho = 1 - \varepsilon$. The identification of the depolarizing channel as a convex mixture with the maximally mixed state is the standard definition of that channel; no empirical claim is involved.

3.3 Classical instantiation: stationary iteration

For a classical stationary iteration $x_{t+1} = Mx_t + c$ with iteration matrix M of spectral radius $\rho(M) < 1$, the error propagates as $e_{t+1} = Me_t$, giving $\|e_t\| \leq \rho(M)^t \|e_0\|$. Setting $\rho(M) = 1 - \varepsilon$ makes the classical and quantum instantiations share the identical step-count law with identical parameters. This identification is a modeling choice, and its consequence — that at fixed ε the two processes are step-count-indistinguishable — is a theorem of the frame, not an empirical observation.

3.4 Fidelity clock and external threshold

For the qubit channel of Section 3.2, the fidelity between ρ_t and the pure initial state $\rho_0 = |\psi\rangle\langle\psi|$ is

$$F(t) = \langle\psi|\rho_t|\psi\rangle = (1 - \varepsilon)^t + (1 - (1 - \varepsilon)^t) \frac{1}{2} = \frac{1 + (1 - \varepsilon)^t}{2}.$$

This lets us convert any fidelity threshold F_* into a step count t_* via $(1 - \varepsilon)^t = 2F_* - 1$, i.e.

$$t_* = \frac{\ln(2F_* - 1)}{\ln(1 - \varepsilon)},$$

valid when $F_* > 1/2$. We apply this to the threshold $F_L = 0.874$ imported from the no-go theorem [11].

3.5 Falsification protocol

Following the pre-registration pattern of [12], we state in advance the observation that would falsify the paper's central claim. The central claim is: at fixed ε and fixed finite $\tilde{h}$, quantum and classical relaxation obey the same step-count law. The pre-registered falsifier is any rigorously derived or empirically measured pair (t_δ^Q, t_δ^C) at the

same ε , δ , and D_0 whose ratio differs from 1 by more than the combined numerical and statistical uncertainty. Section 6 develops what such a pair would require.

4. Analysis

4.1 Baseline step count at $\varepsilon = 0.01$, $\delta = 10^{-6}$

Inputs: $\varepsilon = 0.01$ (chosen as a representative per-step noise strength; the results scale as shown in Section 4.2), $D_0 = 1$ (worst case: a pure state maximally far from the maximally mixed state in trace distance), $\delta = 10^{-6}$ (a tolerance at which residual coherence is negligible for most protocol purposes).

Step 1: contraction factor $\rho = 1 - \varepsilon = 0.99$.

Step 2: $\ln(\delta/D_0) = \ln(10^{-6}) = -6 \ln 10 = -6 \times 2.302585 = -13.815510$.

Step 3: $\ln \rho = \ln(0.99)$. Expanding, $\ln(1 - x) = -x - x^2/2 - x^3/3 - \dots$ with $x = 0.01$:
 $-0.01 - 0.00005 - 0.000000333 - 0.0000000025 = -0.010050336$.

Step 4:

$$t_\delta = \frac{-13.815510}{-0.010050336} = 1374.6.$$

Check: $1374.6 \times 0.010050336 = 13.8155$, matching the numerator. So both the depolarizing qubit and the classical iteration with $\rho = 0.99$ need 1374.6 steps, i.e. 1375 integer steps, to reach $\delta = 10^{-6}$.

4.2 Logarithmic insensitivity to ε

Reduce the noise strength to $\varepsilon = 0.001$, so $\rho = 0.999$. Then $\ln(0.999) = -0.001 - 0.0000005 - 0.000000000333 = -0.001000500333$. The step count is

$$t_\delta = \frac{13.815510}{0.001000500333} = 13808.6.$$

The ratio to the $\varepsilon = 0.01$ case is

$$\frac{13808.6}{1374.6} = 10.045.$$

A tenfold reduction in noise strength multiplies the relaxation step count by 10.045, not by 10: the step count scales as $1/\varepsilon$ up to a logarithmic correction of relative size $\varepsilon/2$ (from the $x^2/2$ term in $\ln(1 - x)$). Explicitly, $t_\delta \approx \frac{\ln(\delta/D_0)}{-\varepsilon - \varepsilon^2/2}$, so $t_\delta \cdot \varepsilon \approx \ln(\delta/D_0) \cdot \frac{1}{1 + \varepsilon/2}$; at $\varepsilon = 0.01$ the correction factor is $\frac{1}{1.005} = 0.99502$, i.e. a 0.498% reduction relative to the pure $1/\varepsilon$ law. This is the quantitative sense in which the relaxation clock is logarithmically insensitive to its parameters: only $\ln(\delta/D_0)$ and $1/\varepsilon$ matter, and the former changes by 2.302585 per decade of tolerance.

4.3 Fidelity threshold crossing at $F_L = 0.874$

Input: $F_* = 0.874$, imported from the no-go theorem stating that quantum error correction and Quantum Darwinism cannot coexist above critical logical fidelity $F_L > 0.874$ [11]. This number is used as an external threshold; its derivation is not reproduced here.

Step 1: $2F_* - 1 = 2 \times 0.874 - 1 = 0.748$.

Step 2: $\ln(0.748)$. We have $\ln(0.748) = \ln(748) - \ln(1000)$.
 $\ln(748) = \ln(7.48) + \ln(100) = 2.012220 + 4.605170 = 6.617390$? Check: $\ln(7.48)$:
 $\ln 7 = 1.945910$, $\ln(7.48/7) = \ln(1.068571) = 0.066310$; sum = 2.012220. So
 $\ln(0.748) = 2.012220 + 4.605170 - 6.907755 = -0.290365$.

Step 3: at $\varepsilon = 0.01$,

$$t_* = \frac{-0.290365}{-0.010050336} = 28.89.$$

So a pure qubit state relaxes below the $F_L = 0.874$ threshold after 28.9 steps at $\varepsilon = 0.01$. At $\varepsilon = 0.1$ ($\rho = 0.9$, $\ln 0.9 = -0.105361$):

$$t_* = \frac{-0.290365}{-0.105361} = 2.756.$$

At $\varepsilon = 0.001$ ($\ln 0.999 = -0.001000500$):

$$t_* = \frac{-0.290365}{-0.001000500} = 290.2.$$

The margin interpretation: the distance from the threshold to the maximally mixed value $F = 1/2$ is $0.874 - 0.5 = 0.374$, and the distance from a pure state to the threshold is $1 - 0.874 = 0.126$; the threshold sits at fraction $0.126/(1 - 0.5) = 0.252$ of the initial-to-fixed-point distance, consistent with $2F_* - 1 = 0.748$ being the surviving contraction fraction.

4.4 Where the quantum case can diverge: the joint limit

At fixed ε and fixed finite $\hbar$, Sections 4.1–4.3 show the quantum relaxation clock is fully determined by the contraction factor $1 - \varepsilon$ and is identical in form to the classical stationary-iteration clock. The known point of divergence is supplied by the toral-map result [6]: relaxation time is sensitive to the chaotic behavior of the corresponding classical system when the semiclassical limit ($\hbar \rightarrow 0$) and the small-noise limit ($\varepsilon \rightarrow 0$) are taken simultaneously. The structural reason the two limits must be joint is visible in our own formulas: as $\varepsilon \rightarrow 0$ at fixed tolerance, $t_\delta \rightarrow \infty$ like $1/\varepsilon$ (Section 4.2), so the noise-induced relaxation that defines the clock vanishes exactly in the limit where the classical dynamics would control the behavior; any chaos-dependence must therefore be resolved in the ordering of $\hbar \rightarrow 0$ relative to $\varepsilon \rightarrow 0$. We can state the scale condition explicitly. The noise sets a relaxation scale $t_{\text{noise}} \sim 1/\varepsilon$ (from Section 4.2, $t_\delta \varepsilon \approx \ln(\delta/D_0)$, so per unit of $\ln(\delta/D_0)$ the scale is $1/\varepsilon$). The semiclassical evolution sets an Ehrenfest-type scale; the toral-map result asserts that relaxation time depends on classical chaos in the regime where both limits are taken together [6]. A quantitative ordering parameter would be the ratio of the Ehrenfest time to $1/\varepsilon$, but the supplied abstract of [6] does not give a formula for the Ehrenfest time or for the relaxation time as a function of $\hbar$ and ε , so we cannot compute that ratio from the grounding; we flag it as the key open quantity in Section 6.

4.5 Per-step cost and the costup criterion

Step count alone does not determine wall-clock time. The submodular-maximization literature makes the analogous point that practice is dominated by the cost of frequently querying the function being optimized [2]. For our comparison, let c_Q and c_C be the per-step costs (in time or energy) of one quantum-map application and one classical iteration, respectively. Wall-clock parity requires

$$c_Q t_\delta^Q < c_C t_\delta^C.$$

Since Section 4.1 gives $t_\delta^Q = t_\delta^C$ at matched parameters, parity reduces to $c_Q < c_C$, and any quantum advantage must come from a smaller per-step cost or from operating at a smaller effective ε . Neither c_Q nor c_C is computable from the grounding, so we report no cost numbers and treat the costup criterion of [1] — significant speedup with minimal development effort — as the qualitative gate a deployment must pass.

5. Results

All numbers below are computed in Section 4 by explicit arithmetic from the stated inputs; none are measured or simulated.

R1 (shared step-count law). At $\varepsilon = 0.01$, $D_0 = 1$, $\delta = 10^{-6}$, both the depolarizing qubit channel and a classical stationary iteration with convergence factor 0.99 relax in $t_\delta = 1374.6$ steps (1375 integer steps). The two processes are step-count-indistinguishable at matched parameters.

R2 (logarithmic insensitivity). Reducing ε from 0.01 to 0.001 raises the step count from 1374.6 to 13808.6, a factor of 10.045. The step count obeys $t_\delta \approx \ln(\delta/D_0)/(-\varepsilon - \varepsilon^2/2)$, with a 0.498% logarithmic correction at $\varepsilon = 0.01$.

R3 (threshold crossing). A pure qubit state crosses the imported logical-fidelity threshold $F_L = 0.874$ after 28.9 steps at $\varepsilon = 0.01$, after 2.756 steps at $\varepsilon = 0.1$, and after 290.2 steps at $\varepsilon = 0.001$. The threshold sits at 25.2% of the initial-to-fixed-point distance.

R4 (locus of divergence). At fixed ε and finite $\hbar$, no divergence between quantum and classical relaxation is derivable in the contraction frame. The only grounded locus of divergence is the joint $\hbar \rightarrow 0$, $\varepsilon \rightarrow 0$ limit, where relaxation time of noisy quantum maps on tori is sensitive to the chaos of the underlying classical system [6]. The quantitative ordering parameter (an Ehrenfest-time-to- $1/\varepsilon$ ratio) is not computable from the supplied grounding and is listed as open.

R5 (cost gate). Wall-clock parity at matched parameters requires $c_Q < c_C$; no per-step cost values are derivable from the grounding, so this is stated as a gate, not a number.

6. Discussion

The central result, R1, is a negative one: in the simplest common frame, relaxation in a noisy quantum map and relaxation in a classical iterative algorithm are the same process mathematically, and the interesting physics lives entirely in the joint limit that the contraction frame deliberately freezes. This framing has three limitations that a reader should weigh before using R1 in any decision.

First, the contraction frame assumes a time-independent, linear contraction with a single fixed point. Real quantum maps can have multiple fixed points, limit cycles, or non-normal contraction structure, and real classical solvers use preconditioning, acceleration, and adaptive tolerances that change the effective ρ from step to step. The Jacobi-type algorithms cited as the classical representative are described as computing eigenpairs with high relative accuracy when the matrices permit [4], and that qualification ("if the matrices permit") is itself a reminder that convergence quality is instance-dependent in ways a single- ρ model hides.

Second, the identification $\rho(M) = 1 - \varepsilon$ is a modeling choice. A classical iteration with $\rho = 0.99$ and a qubit under $\varepsilon = 0.01$ depolarizing noise match by construction, but nothing forces a real workload to present matched parameters. The honest statement is conditional: *if* the parameters are matched, the step counts are equal. The falsification protocol of Section 3.5 is designed around exactly this conditionality.

Third, the imported threshold 0.874 comes from a no-go theorem about coexistence of quantum error correction and Quantum Darwinism [11], and we use it as a generic fidelity clock marker. That is a reinterpretation: the theorem constrains coexistence of two information-theoretic roles, not the step count of a relaxation process. Our R3 is therefore a derived consequence of importing the number into the depolarizing model, not an additional property of the no-go theorem, and it should be read as a consistency illustration of the fidelity clock rather than as a hardware-relevant lifetime prediction.

What would falsify the main result. The main result (R1, with R4 as its boundary) is falsified if any of the following is established: (i) a noisy quantum map at fixed ε and finite $\hbar$ whose relaxation time to tolerance δ differs from $\ln(\delta/D_0)/\ln(1-\varepsilon)$ by more than numerical uncertainty — this would show the single-parameter contraction frame fails for quantum maps, e.g. through non-unimodal spectral structure; (ii) a classical stationary iteration and a quantum map with matched contraction factors whose measured or rigorously derived step counts differ at the same δ and D_0 ; (iii) a derivation showing that the chaos-sensitivity of relaxation time in the toral-map setting [6] survives at fixed finite $\hbar$ and ε , rather than only in the joint limit. Any one of these would force a revision of R1; none is currently available in the grounding.

Limitations list.

- *Single fixed point, linear contraction.* R1 holds only within the frame of Section 3.1; sensitivity to this assumption is total, since the entire derivation is the frame.
- *Depolarizing noise model.* The quantum instantiation uses the standard depolarizing channel; other channels (amplitude damping, dephasing) contract in different norms and can give different fixed points, so R1's numerical value 1374.6 is model-specific, though the law's form survives any geometric contraction.
- *Matched-parameter condition.* The equality in R1 requires $\rho(M) = 1 - \varepsilon$; unmatched parameters reintroduce a difference that is trivially large, so the result constrains only fair comparisons.
- *Imported threshold.* R3 reinterprets the 0.874 fidelity bound from [11] as a clock marker; the no-go theorem itself does not concern relaxation dynamics, and R3 inherits no model-independence from the theorem despite the theorem being described as model-independent.
- *No per-step costs.* R5 is a gate without numbers because no cost data exist in the grounding; any cost claim would be unsupported.
- *Open ordering parameter.* The ratio of the semiclassical (Ehrenfest-type) time to $1/\varepsilon$, which would quantify R4, is not derivable from the supplied abstract of [6]; the paper's divergence claim is therefore qualitative at exactly the point where it would be most useful.
- *Thin grounding for two sources.* The entries for the pre-registered falsification work [12] and "Programming the Vacuum" [9] supply no abstract text; claims resting on them are correspondingly minimal, and the bibliography contains twelve works of which two are contentless in this grounding — a limitation of the source set itself.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users are operators and architects of hybrid classical-quantum compute facilities who must decide, for a relaxation-type workload, whether to schedule it on classical hardware or on a quantum co-processor. If R1 holds, the decision rule changes from "estimate the quantum speedup exponent" to "compare per-step costs c_Q and c_C at matched contraction parameters, and check whether the workload's effective noise strength ε can be lowered on either route." Concretely, a scheduling policy derived from R1 is: compute $t_\delta = \ln(\delta/D_0)/\ln(1-\varepsilon)$ for the workload's tolerance and noise level, apply it identically to both routes, and break the tie only on per-step cost and limit-ordering risk (R4). A second user group is the authors

and consumers of quantum-relaxation proposals: R1 supplies the fair-comparison condition $\rho(M) = 1 - \varepsilon$ that any claimed quantum advantage must satisfy before it is credited to quantumness rather than to unmatched parameters. A third group is error-correction system designers, for whom R3 gives a derived illustration: at $\varepsilon = 0.01$, the imported threshold $F_L = 0.874$ is crossed after 28.9 steps, so any protocol that must keep logical fidelity above that marker has a per-step noise budget of order 10^{-3} or smaller for runs of hundreds of steps (from Section 4.3: $t_* = 290.2$ at $\varepsilon = 0.001$).

(b) Current readiness. This work is purely analytical and has not been peer reviewed: the contraction frame of Section 3.1, the derivations of Section 4, and the falsification protocol of Section 3.5 are all pen-and-paper results with no hardware, simulation, or experimental component. Against the NASA Technology Readiness Level definitions, this places the work at TRL 1–2: basic principles observed and reported, with the underlying physical and mathematical assumptions formulated but not applied in any demonstrated context. Analytical work cannot exceed TRL 3, and nothing in this paper supports a claim of experimental validation, component development, or manufacturing readiness. The single externally sourced number, $F_L = 0.874$ from [11], is imported as a threshold marker and inherits no validation status from this paper.

(c) The smallest next test. The smallest de-risking calculation is a numerical one, requiring no quantum hardware: implement the depolarizing qubit channel of Section 3.2 and a classical stationary iteration with $\rho(M) = 1 - \varepsilon$ on the same computer, run both from matched initial conditions to tolerance $\delta = 10^{-6}$ at $\varepsilon = 0.01$, and verify that both reach the tolerance at the same step count, predicted to be 1375 integer steps (Section 4.1). The outcome that refutes R1 is a measured step-count mismatch between the two implementations at matched parameters beyond floating-point round-off; the outcome that refutes the frame itself is a quantum-map simulation at fixed ε and finite effective $\hbar$ whose relaxation step count deviates from $\ln(\delta/D_0)/\ln(1 - \varepsilon)$ by more than numerical uncertainty (falsifier (i) of Section 6). A second, larger next step is to compute or bound the ordering parameter identified in Section 4.4 — the ratio of the semiclassical (Ehrenfest-type) time to $1/\varepsilon$ — for a specific quantized toral map, which would convert the qualitative divergence claim R4 into a quantitative one.

(d) Binding constraints. Because this is analytical work, the constraints that would bind first in any future engineering use are, in order: (i) *noise characterization* — the law's input ε must be the actual per-step noise strength of the device, and no device-level noise data exist in the grounding, so this is an open question rather than a derived constraint; (ii) *channel validity* — the quantum instantiation assumes depolarizing noise; whether a real device's error channel contracts in trace distance with a single factor $1 - \varepsilon$ is an open question that the falsifier (i) of Section 6 would resolve; (iii) *per-step cost data* — the cost gate of R5 requires c_Q and c_C , neither of which is derivable from the grounding, so wall-clock comparisons remain blocked on measurement; (iv) *limit-ordering theory* — the Ehrenfest-time-to- $1/\varepsilon$ ratio of Section 4.4 is not computable from the supplied abstract of [6], so any deployment in the joint-limit regime is blocked on theory before it is blocked on hardware. Materials, tolerances, and component constraints do not bind at this stage because no hardware design is proposed.

(e) Order-of-magnitude bounds. The only quantitative bounds this paper can state as projections are step-count scalings already derived in Section 4, restated here with their assumptions. Assuming the depolarizing model and $D_0 = 1$: a tenfold reduction in tolerance (one decade of δ) adds $2.302585/0.010050336 \approx 229.1$ steps at $\varepsilon = 0.01$ (from Section 4.2's numerator, $\ln 10 = 2.302585$, divided by the Section 4.1 denominator); a tenfold reduction in ε multiplies the step count by 10.045 (Section 4.2). No cost, energy, market, or performance figures are stated, because no per-step cost or device data exist in the grounding; any such figure would be unsupported.

(f) Risks and misuse paths. The principal misuse is the one the paper is designed to prevent, in both directions: citing R1 to claim "no quantum advantage is possible" (it constrains only matched-parameter, fixed- ε , single-fixed-point comparisons), or citing R4 to claim "quantum relaxation is chaos-accelerated" (the chaos-sensitivity of [6] is established only in the joint $\hbar \rightarrow 0, \varepsilon \rightarrow 0$ limit, and no ordering parameter is computed here). A second misuse

path is reading R3 as a hardware lifetime prediction; Section 6 states explicitly that R3 is a consistency illustration of the fidelity clock, not a property of the no-go theorem of [11]. A third is treating the imported threshold 0.874 as a certified operating point; its provenance is a single corpus entry, and this paper neither re-derives nor validates it.

8. Conclusion

Relaxation in a noisy quantum map and relaxation in a classical stationary iteration, placed in the same linear-contraction frame with matched parameters, obey the identical step-count law $t_\delta = \ln(\delta/D_0)/\ln(1 - \varepsilon)$: at $\varepsilon = 0.01$ and $\delta = 10^{-6}$ both need 1374.6 steps, and the law is logarithmically insensitive to its inputs, a tenfold noise reduction buying only a factor of 10.045 in step count. The fidelity clock derived from the same frame crosses the imported threshold $F_L = 0.874$ after 28.9 steps at $\varepsilon = 0.01$. Within this frame, therefore, there is no quantum-classical distinction to exploit or fear; the one grounded locus of distinction is the joint semiclassical and weak-noise limit, where relaxation time of quantized toral maps is known to depend on the chaos of the underlying classical system, and where the quantitative ordering parameter remains open. The practical yield of the analysis is a fair-comparison condition, a pre-registered falsifier, and a smallest next test — a matched-parameter numerical run of the depolarizing channel against a stationary iteration, predicted to terminate at the same 1375th step — by which the central claim can be refuted if it is wrong.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.16.0-revise-retry). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; statements about a cited work are restricted to its retrieved title and abstract.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 7 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond its abstract; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Distributing an Exact Algorithm for Maximum Clique: maximising the costup. arXiv:1209.4560v1. <https://arxiv.org/abs/1209.4560v1> [2] Fast Parallel Algorithms for Submodular p -Superseparable Maximization. arXiv:2311.13123v2. <https://arxiv.org/abs/2311.13123v2> [3] Approximation algorithms for TSP with neighborhoods in the plane. arXiv:1703.01640v1. <https://arxiv.org/abs/1703.01640v1> [4] Novel Modifications of Parallel Jacobi Algorithms. arXiv:1008.0201v2. <https://arxiv.org/abs/1008.0201v2> [5] Simplifying a classical-quantum algorithm interpolation with quantum singular value transformations. arXiv:2207.14810v3. <https://arxiv.org/abs/2207.14810v3> [6] Relaxation Time of Quantized Toral Maps. arXiv:math-ph/0406055v3. <https://arxiv.org/abs/math-ph/0406055v3>

[7] Algorithms for the Majority Rule (+) Consensus Tree and the Frequency Difference Consensus Tree. arXiv:1307.7821v2. <https://arxiv.org/abs/1307.7821v2> [8] An Algorithm for Odd Gracefulness of the Tensor Product of Two Line Graphs. arXiv:1103.4502v1. <https://arxiv.org/abs/1103.4502v1> [9] QNFO: Programming the Vacuum [11] QNFO: Archimedean Shadows: The QEC-Darwinism Tradeoff in Ultrametric Spaces [12] QNFO: A Pre-Registered Falsification of Deterministic Measurement-Triggered Relaxation