

What Must a Quantum-Memory Decoder Know About Temporally Correlated Noise? A Reconciled Exposition and Cost Accounting

What Must a Quantum-Memory Decoder Know About Temporally Correlated Noise? A Reconciled Exposition and Cost Accounting

Abstract

Temporal correlations in quantum-memory noise create a calibration problem that is invisible to syndrome statistics alone. We reconstruct, in expository form, the central results of a study of fixed stabilizer memory experiments with known system-environment interactions [1],[2]. The study's main example uses weak coherent X rotations in a class of CSS codes including Shor, Steane, and odd-distance rotated surface codes, in which all physical qubits share a sign that is either fixed throughout a run or independently redrawn each interval. The two noise models have identical syndrome statistics and the same spatiotemporal Pauli process (SPP), obtained by Pauli-twirling every noise interval, yet require different logical corrections at suitable storage times. We give an explicit derivation of a symmetric-disagreement mechanism under which the worst-case fidelity loss of a model-agnostic decoder with full syndrome access approaches $1/2$, reproducing the bound stated in the source abstract, and we compute conditional consequences of the calibration-cost scalings reported in two of the reconciled drafts: a number of noise intervals proportional to the inverse rotation angle raised to the code's X distance, reduced to inverse-square scaling when two uninterrupted intervals precede one extraction. Because the supplied bibliography entry is truncated before the scaling statements, we label them as attributed-but-unverified hypotheses and derive their conditional numerical consequences with full arithmetic.

1. Introduction

A quantum memory is only as good as its decoder, and a decoder is only as good as its noise model. When the noise acting on a stabilizer code is independent and identically distributed in time, the syndrome record carries essentially all relevant information and standard decoding heuristics suffice. When the noise is temporally correlated — for example, when a coherent error rotates every physical qubit in a shared direction over many consecutive intervals — the syndrome statistics can remain identical to those of an uncorrelated model while the optimal logical correction changes completely. This is the problem addressed in the work under study [1],[2]: for fixed stabilizer memory experiments with known system-environment interactions, what must a decoder know about temporally correlated noise, and what does it cost to learn?

The question matters practically. Quantum error correction carries a large overhead: the QNFO analysis of fundamental physical limits on computation documents overheads of 10^2 – 10^3 physical operations per logical operation, which multiply the thermodynamic cost of the whole machine [13]. If calibration of temporal noise structure adds a further overhead that grows with code distance and inverse noise strength, it must be counted in any honest investment assessment of quantum-memory architectures [10], and a decoder's noise assumptions are an auditable design choice in the sense exemplified by the QNFO audit methodology [12].

This paper is a reconciliation of three independent expository drafts, not a new proof. All structural claims about the two noise models, the fidelity-loss bound, and the calibration criterion are taken from the supplied abstract of the source work [1],[2]; our contribution is to (i) state the mechanism in elementary terms, (ii) derive the stated worst-case loss from explicit assumptions with all arithmetic shown, and (iii) compute conditional numerical consequences of the calibration scalings, which the supplied bibliography entry does not itself contain and which we

therefore treat as labeled hypotheses. Section 2 situates the work in the supplied literature. Section 3 sets up the models. Section 4 contains the derivations. Section 5 reports only computed numbers and labeled projections. Section 6 discusses limitations and failure modes. Appendix A documents divergences among the drafts and the conventions adopted; Appendix B attributes each substantive claim.

2. Background and Related Work

We discuss each supplied bibliography entry in turn, restricting every statement to what the entry itself supports.

[1] and [2] (the source work). Entries [1] and [2] refer to the same preprint, arXiv:2610.03545v1, "What Must a Quantum-Memory Decoder Know About Temporally Correlated Noise?" The supplied abstract of [1] states the results this paper reconstructs: calibration formulated in terms of the fidelity differences that select a Pauli recovery, with missing information related to the resulting fidelity loss; a main example of weak coherent X rotations in a class of CSS codes including Shor, Steane, and odd-distance rotated surface codes; two sign models (fixed throughout a run versus independently redrawn each interval) with identical syndrome statistics and the same spatiotemporal Pauli process (SPP) obtained by Pauli-twirling every noise interval; different required logical corrections at suitable storage times; almost perfect correctability in the weak-noise limit when the model is known; and a worst-case fidelity loss approaching $1/2$ for the best final recovery without model information, despite access to the full syndrome record. The supplied entry text is truncated mid-sentence at exactly this point. Statements appearing in drafts B and C as further quotes of the source abstract — matching calibration bounds at fixed code and fixed positive target loss below $1/2$, interval counts scaling as θ^{-dx} and θ^{-2} , and the effect of keeping or discarding the record of applied random Pauli operations — are therefore treated in this paper as attributed-but-unverified hypotheses (Section 4, hypotheses H_1 – H_3), not as established results of [1].

[3] arXiv:1008.0667v2, "Noise and Bell's inequality". This entry discusses the concept of reality from the EPR paradox onward and the Bell inequality, noting that considerable evidence exists that the inequality can be violated even by classical systems, so that experiments showing quantum behavior may be open to alternative interpretation. The relevance to our topic is thematic rather than technical: the source work's two noise models are observationally equivalent at the level of syndrome statistics, exactly the situation in which empirical data alone underdetermine the underlying model. The entry supplies no further technical detail, and we use it only for this conceptual connection.

[4] arXiv:1310.1481v1, "Strongly correlated superconductivity". These lecture notes for an international school concern the failure of band theory and BCS theory — arguably the most successful theories of condensed matter physics — in a number of materials, in particular the high-temperature superconductors and the layered organic superconductors, and emphasize that low-energy properties of a phase of matter are generally emergent. The connection to our problem is the general lesson that effective descriptions can hide the degrees of freedom that determine behavior in another regime; the supplied summary gives no quantum-information content, and we draw no technical claim from it.

[5] arXiv:1612.01464v2, "Finite blocklength and moderate deviation analysis of hypothesis testing of correlated quantum states and application to classical-quantum channels with memory". This work applies martingale concentration inequalities — described as a powerful mathematical tool across probability, statistics, information theory, and machine learning — to quantum hypothesis testing, the problem of discriminating quantum states belonging to two different sequences $\{\rho_n\}$. This is the closest methodological relative of the source work in the supplied bibliography: calibrating between the fixed-sign and redrawn-sign noise models is a discrimination problem between two temporally correlated noise processes, and the finite-blocklength perspective matches the question of how many noise intervals calibration requires. The entry's summary is truncated and gives no specific bounds, so we cite it for the methodological framing only.

[6] arXiv:1601.00157v1, "Theory of noise suppression in Λ -type quantum memories by means of a cavity".

This entry concerns quantum memories capable of storing single photons or other quantum states of light for on-demand retrieval — a route to large-scale quantum information processing with light — based on far-off-resonant Raman absorption in ensembles of Λ -type atoms, and identifies unwanted four-wave mixing at room temperature as prohibitive for application. It is a concrete instance of a physically structured, unwanted noise process degrading a quantum memory, motivating the abstract study of how decoders must handle correlated noise. The summary is truncated; no further detail is available from it.

[7] arXiv:1612.03430v1, "Quantum and Classical Frontiers of Noise". This is an introduction to the eleven works of a special issue, using the weather and its butterfly effect as the typical example of phenomena not predictable with certainty, and defining noise in classical or quantum phenomena as the difference between the empirical output of an experiment and its statistical description. This definition is directly apt: the source work's central observation is that two noise models with the same statistical description of the empirical output (the syndrome record) can nonetheless demand different actions from the decoder.

[8] arXiv:2311.17794v1, "Noise Dynamics in the Quantum Regime". This experimental work shows that a time-dependent bias voltage on a tunnel junction generates a time-dependent modulation of its current fluctuations, in particular of their variance, translating into an excitation at frequency $\tilde{f}$ that generates correlations between current fluctuating at any frequency f and at frequency $\pm\tilde{f} - f$; the measurement of such a correlation is reported in the fully quantum regime. It demonstrates that engineered temporal correlations in noise are real, controllable, and measurable — the physical premise on which the source work's calibration question rests. The summary is truncated beyond this result.

[9] arXiv:1605.07838v1, "Decoherence and noise in open quantum system dynamics". This entry treats quantum noise within the Copenhagen interpretation applied to a composite system-environment setting, where averaging over the environmental degrees of freedom leads to a stochastic quantum dynamics described by equations complying with the constraints arising from the statistical structure of quantum mechanics. This is the standard formal backdrop for the source work's "known system-environment interactions": the environment is explicit and the noise model is a stochastic process over intervals. The summary is truncated and supplies no specific equations.

[10]–[13] (QNFO corpus). Entry [10] translates a critique into a concrete investment portfolio, matching computational problem classes (optimization, linear algebra, probabilistic inference, quantum simulation, cryptography, general-purpose) to optimal physical substrates; it frames our significance argument that calibration overheads belong in the cost accounting of any substrate proposed for quantum memory. Entry [11] is supplied with an empty summary and provides no content; we note this explicitly as a limitation of the supplied corpus. Entry [12] audits a tunable quantum neural network implemented on three quantum computing platforms, through a research pipeline including due diligence, an external literature search across 32 papers, and a nine-stage process; it exemplifies the audit culture in which decoder noise assumptions should be treated as auditable design choices. Entry [13] examines fundamental physical limits on computation — the Landauer bound, Margolus-Levitin theorem, Bremermann limit, and Bekenstein bound — and shows that quantum error-correction overhead of 10^2 – 10^3 physical operations per logical operation multiplies thermodynamic cost; this is the quantitative backdrop against which any additional calibration overhead must be judged.

In summary, the bibliography contains one central technical source ([1]/[2]) and a ring of contextual works on noise foundations [3],[7], correlated-state discrimination [5], memory hardware [6], engineered noise correlations [8], open-system dynamics [9], condensed-matter correlations [4], and computation-cost accounting [10]–[13]. Several entries are truncated; where so, we have said so and drawn no unsupported inference.

3. Methods

3.1 The two noise models

Following [1], consider a fixed stabilizer memory experiment using a CSS code with X distance d_X (the minimum weight of a logical X -type operator; the class named in [1] includes Shor, Steane, and odd-distance rotated surface codes). Time is divided into noise intervals. In each interval, every physical qubit undergoes a weak coherent X rotation by a small angle θ , with a sign $s_i \in \{+1, -1\}$ shared by all physical qubits in that interval. Two models:

- M_{fix} : the shared sign is fixed throughout the run, $s_i = s$ for all intervals i .
- M_{red} : the shared sign is independently redrawn each interval, $s_i \sim \text{Unif}\{+1, -1\}$.

Both models are assumed known in structure (the system-environment interaction is known, per [1]); the calibration question is which model obtains and what information suffices to decode optimally.

3.2 Observational equivalence

By Pauli-twirling every noise interval — replacing each coherent rotation by a random Pauli operation chosen to reproduce the same averaged channel — both models yield the same spatiotemporal Pauli process (SPP), and, per [1], the two models have identical syndrome statistics. Hence any decoder whose input is only the syndrome record, even the full record, cannot distinguish the models from a single run's statistics; yet, per [1], the models require different logical corrections at suitable storage times.

3.3 Decoder classes

We consider three decoder classes:

1. $\mathcal{D}_{\text{aware}}$: knows which model obtains and decodes optimally for it.
2. $\mathcal{D}_{\text{SPP}}$: knows only the SPP, not the sign realization.
3. $\mathcal{D}_{\text{agn}}$: knows only the syndrome record and nothing about which model obtains.

3.4 Fidelity accounting

For a Pauli recovery R_a and noise model M , let $F_a(M)$ denote the resulting storage fidelity. Calibration in [1] is formulated in terms of the fidelity differences that select a Pauli recovery: the selected recovery under model M is

$$a^*(M) = \arg \max_a F_a(M),$$

and the fidelity loss of a decoder D relative to the noise-aware optimum is

$$L(D) = \max_{M \in \{M_{\text{fix}}, M_{\text{red}}\}} [F^*(M) - F_D(M)],$$

where $F^*(M)$ is the optimum under model M .

3.5 Hypotheses inherited from the drafts

The following statements are reported in drafts B and C as direct quotes of the source abstract, but the supplied bibliography entry is truncated before them. We carry them as explicitly labeled hypotheses and derive only conditional consequences:

- H_1 (calibration scaling): with logical preparation and final readout, complete QEC cycles require a number of noise intervals scaling as θ^{-d_X} ; at fixed code and fixed positive target loss below $1/2$, matching calibration

bounds hold.

- H_2 (relaxed schedule): allowing two uninterrupted intervals before one extraction reduces the optimal cost to inverse-square scaling, i.e., proportional to θ^{-2} .
- H_3 (twirling record): physical Pauli twirling of every noise interval changes both recovery performance and its information requirements; discarding the record of applied random Pauli operations makes the SPP sufficient for optimal recovery, while keeping the record can improve attainable fidelity but may require additional noise information.

4. Analysis

Every input number below is either (a) quoted from the supplied abstract of [1], with the quote identified, or (b) an assumption of our illustrative parameter set, explicitly labeled. All arithmetic is shown.

4.1 Worst-case fidelity loss approaching $1/2$

Input (from [1]): "Without model information, even the best final recovery has a worst-case fidelity loss approaching $1/2$ relative to the noise-aware optimum, despite access to the full syndrome record." Also from [1]: in the weak-noise limit, each model is "almost perfectly correctable when known."

Derivation (symmetric-disagreement mechanism, from draft C). Suppose the storage time lies in the regime where the two models' optimal logical corrections disagree. Assume (Assumption A1, an idealization constructed to reproduce the stated limit):

- (i) The noise-aware optima achieve $F^*(M_{\text{fix}}) = F^*(M_{\text{red}}) = F^*$, with $F^* \rightarrow 1$ in the weak-noise limit.
- (ii) The optimal corrections under the two models apply opposite logical X -type corrections on a set S of syndrome records, and the record lies in S with probability q .
- (iii) The agnostic decoder's output is a fixed function of the syndrome record, which is identically distributed under both models; so it applies the same correction under both.

On records in S , a decoder applying the correction optimal for M_{fix} is wrong for M_{red} and vice versa. Under the symmetric idealization where a wrong logical correction takes the fidelity from F^* to approximately $1 - F^*$, the loss on the mismatched branch is

$$F^* - (1 - F^*) = 2F^* - 1.$$

Averaged over the two models with equal prior weight $1/2$ each, the worst-case loss is

$$L_{\text{worst}} = \frac{1}{2} (2F^* - 1) \xrightarrow{F^* \rightarrow 1} \frac{1}{2} (2 - 1) = \frac{1}{2}.$$

This reproduces the bound stated in [1] as the $q \rightarrow 1$, $F^* \rightarrow 1$ limit of the symmetric disagreement mechanism. The bound is a worst case over the two models, matching the abstract's phrasing "approaching $1/2$."

Numerical instance. Take $F^* = 0.98$ (Assumption A2, illustrative weak-noise value) and full disagreement $q = 1$. Then

$$L_{\text{worst}} = \frac{1}{2} (2 \times 0.98 - 1) = \frac{1}{2} (0.96) = 0.48.$$

So with $F^* = 0.98$, the model-agnostic decoder loses 0.48 in worst-case fidelity against the noise-aware optimum — close to the asymptotic 0.5.

4.2 Per-qubit error probability (from draft B)

A coherent rotation $\exp(-i\theta X/2)$ applied to a computational-basis state produces an X -error component with probability

$$p_X(\theta) = \sin^2\left(\frac{\theta}{2}\right) \approx \frac{\theta^2}{4} \quad (\theta \ll 1).$$

Concrete value: take $\theta = 10^{-2}$ (Assumption A3, illustrative weak angle; the choice is ours, not from [1]). Then

$$p_X \approx \frac{(10^{-2})^2}{4} = \frac{10^{-4}}{4} = 2.5 \times 10^{-5}.$$

For a CSS code with X distance d_X , a logical X -type failure requires a chain of at least d_X physical X errors, so to leading order $p_L \approx C_{d_X} p_X^{d_X}$ with C_{d_X} the code-dependent multiplicity of minimum-weight logicals. With $d_X = 3$:

$$p_L \approx C_3 \times (2.5 \times 10^{-5})^3 = C_3 \times 1.5625 \times 10^{-14}.$$

Even with C_3 as large as 10^2 , $p_L \lesssim 1.5625 \times 10^{-12}$ per interval — consistent with [1]'s statement that each model is almost perfectly correctable when known.

4.3 Conditional calibration costs under H_1 and H_2

Hypothesis H_1 : $N_{\text{cyc}}(d_X, \theta) \propto \theta^{-d_X}$. **Hypothesis H_2 :** $N_2(d_X, \theta) \propto \theta^{-2}$. The proportionality constants are not stated in any supplied text, so all interval counts below are stated in units of an unknown constant, and only ratios are constant-free.

Instance 1 (drafts A/B convention, Assumption A3): $\theta = 10^{-2}$, $d_X = 3$.

$$N_{\text{cyc}} = C \cdot (10^{-2})^{-3} = C \cdot 10^6 \text{ intervals}, \quad N_2 = C' \cdot (10^{-2})^{-2} = C' \cdot 10^4 \text{ intervals}.$$

Arithmetic: $(10^{-2})^{-3} = 10^6$; $(10^{-2})^{-2} = 10^4$. Improvement ratio (assuming $C' = C$, Assumption A4):

$$\frac{N_{\text{cyc}}}{N_2} = \frac{10^6}{10^4} = 10^2 = \theta^{-(d_X-2)} = (10^{-2})^{-1} = 100.$$

Draft A's alternative convention. Draft A writes the bound with an explicit target-loss prefactor, $N \geq \theta^{-d_X} / \epsilon$. At $\theta = 10^{-2}$, $d_X = 3$, $\epsilon = 0.1$:

$$N = 10^6 \times \frac{1}{0.1} = 10^6 \times 10 = 10^7 \text{ intervals},$$

and the two-interval variant gives $N_{2\text{-int}} = 10^4 \times 10 = 10^5$. Because no supplied text states any prefactor, we adopt draft C's proportional convention for the main text and record draft A's convention in Appendix A.

Instance 2 (draft C convention, Assumption A5): $\theta = 0.05$, $d_X = 3$.

$$N_{\text{cyc}} = C \cdot (0.05)^{-3} = C \cdot 5^{-3} \times 10^6 = C \cdot \frac{10^6}{125} = 8000 C \text{ intervals},$$

$$N_2 = C' \cdot (0.05)^{-2} = C' \cdot \frac{1}{0.0025} = 400 C' \text{ intervals}, \quad \frac{N_{\text{cyc}}}{N_2} = \frac{8000}{400} = 20 = (0.05)^{-1}.$$

At $\theta = 0.1$: $(0.1)^{-3} = 10^3 = 1000$, an $8000/1000 = 8$ -fold reduction, consistent with $(0.1/0.05)^3 = 2^3 = 8$.

4.4 Storage-time projection (labeled projection, from draft C)

Projection with stated assumptions. If each noise interval has duration Δt and the fixed-sign model accumulates a net coherent rotation $N\theta$ after N intervals, the storage time at which the two models' optimal corrections maximally disagree (net rotation of order $\pi/2$ under a coherent-accumulation heuristic) is, as an order-of-magnitude projection,

$$N_{\text{dis}} \approx \frac{\pi/2}{\theta} = \frac{\pi}{2\theta}.$$

With $\theta = 0.05$:

$$N_{\text{dis}} \approx \frac{\pi}{2\theta} = \frac{\pi}{0.1} = 10\pi \approx 31.4 \text{ intervals.}$$

With $\theta = 10^{-2}$ (Assumption A3): $N_{\text{dis}} \approx \frac{\pi}{0.02} = 50\pi \approx 157.1$ intervals. This projection assumes (i) coherent accumulation of the fixed-sign rotation across intervals, (ii) maximal disagreement at net rotation of order $\pi/2$, and (iii) uniform interval duration Δt ; it is an order-of-magnitude heuristic, not a result of [1], and no supplied text states it.

5. Results

We report only numbers computed in Section 4, each with its assumptions.

1. **Worst-case fidelity loss (Section 4.1).** Under the symmetric-disagreement idealization (Assumption A1) with $F^* = 0.98$ (Assumption A2) and full disagreement $q = 1$: $L_{\text{worst}} = \frac{1}{2}(2 \times 0.98 - 1) = 0.48$, approaching the asymptotic $\frac{1}{2}$ as $F^* \rightarrow 1$, reproducing the bound stated in [1].
2. **Per-qubit error probability (Section 4.2).** At $\theta = 10^{-2}$ (Assumption A3): $p_X \approx \frac{(10^{-2})^2}{4} = 2.5 \times 10^{-5}$. For $d_X = 3$ and multiplicity $C_3 = 10^2$: $p_L \lesssim 1.5625 \times 10^{-12}$ per interval.
3. **Calibration interval counts under hypotheses H_1 and H_2 (Section 4.3).** At $\theta = 10^{-2}$, $d_X = 3$: $N_{\text{cyc}} = C \cdot 10^6$, $N_2 = C' \cdot 10^4$, ratio $N_{\text{cyc}}/N_2 = 100$ (assuming $C' = C$, Assumption A4). At $\theta = 0.05$ (Assumption A5): $N_{\text{cyc}} = 8000 C$, $N_2 = 400 C'$, ratio 20. At $\theta = 0.1$: $N_{\text{cyc}} = 1000 C$, an 8-fold reduction relative to $\theta = 0.05$, consistent with $(0.1/0.05)^3 = 8$.
4. **Draft A prefactor convention (Section 4.3).** With $N \geq \theta^{-d_X}/\epsilon$ at $\theta = 10^{-2}$, $d_X = 3$, $\epsilon = 0.1$: $N = 10^7$ intervals; the two-interval variant gives 10^5 .
5. **Disagreement-time projection (Section 4.4, labeled projection).** $N_{\text{dis}} \approx 31.4$ intervals at $\theta = 0.05$ and ≈ 157.1 intervals at $\theta = 10^{-2}$, under the stated coherent-accumulation heuristic.

All interval counts under H_1/H_2 are conditional on hypotheses attributed to drafts B and C but not verifiable from the supplied bibliography entry, and carry unknown proportionality constants C, C' .

6. Discussion

Limitations. The supplied bibliography entry for [1]/[2] is truncated before the calibration-scaling statements, so H_1 – H_3 are attributed-but-unverified hypotheses; every interval count in Section 5 depends on them. The worst-case-loss derivation uses an idealized symmetric-disagreement mechanism (Assumption A1) constructed to reproduce the stated limit; the actual disagreement probability q , the fidelity penalty of a wrong logical correction,

and the prior over models are not supplied, and the bound's approach to $\frac{1}{2}$ could in principle arise by other mechanisms. The per-qubit value $p_X = 2.5 \times 10^{-5}$ and the storage-time projection use illustrative parameters ($\theta = 10^{-2}$, $\theta = 0.05$) chosen by us, not by [1]. The proportionality constants C , C' in the calibration scalings are unknown, so absolute interval counts are meaningful only up to these constants; only ratios are constant-free, and even the ratio $N_{\text{cyc}}/N_2 = 100$ assumes $C' = C$ (Assumption A4).

Failure modes. If the two models' optimal corrections disagree only on a small set of syndrome records ($q \ll 1$), the agnostic decoder's loss is $q(2F^* - 1)$ rather than approaching $\frac{1}{2}$. If the fidelity penalty of a wrong correction is much smaller than $2F^* - 1$, the loss shrinks proportionally. If the calibration scalings carry code-dependent prefactors that grow with d_X , the ratio $\theta^{-(d_X-2)}$ understates the benefit of the two-interval schedule.

What would falsify the claims. The mechanism derivation would be falsified by a demonstration that the two models' optimal corrections agree on all syndrome records at the storage times of interest, or that the agnostic loss is bounded well below $\frac{1}{2}$. The hypotheses H_1 – H_3 would be falsified by the full text of [1] stating different scalings (for example, a power of θ other than d_X or 2), or by the absence of any scaling statement. The storage-time projection would be falsified by a demonstration that maximal disagreement occurs at a net rotation scaling differently from $\pi/2$ per θ .

Open questions. What are the proportionality constants in H_1 and H_2 , and are they code-dependent? How does the disagreement probability q depend on storage time and d_X ? Does keeping the record of applied random Pauli operations (H_3) change the calibration scalings themselves? How should the calibration overhead be folded into the 10^2 – 10^3 operations-per-logical-operation overhead documented in [13]?

7. Conclusion

We have reconciled three independent expository drafts of the study of decoders for stabilizer memories under temporally correlated noise [1],[2]. The two sign models — fixed shared sign versus independently redrawn sign — have identical syndrome statistics and the same spatiotemporal Pauli process, yet require different logical corrections at suitable storage times; under an explicit symmetric-disagreement idealization, a model-agnostic decoder with full syndrome access suffers a worst-case fidelity loss approaching $\frac{1}{2}$, with the illustrative value 0.48 at $F^* = 0.98$. Conditional on hypotheses H_1 and H_2 attributed to the drafts but unverifiable from the supplied entry, calibration requires a number of noise intervals scaling as θ^{-d_X} , reduced to θ^{-2} when two uninterrupted intervals precede one extraction — a ratio of 100 at $\theta = 10^{-2}$, $d_X = 3$, and 20 at $\theta = 0.05$. Because calibration overhead multiplies the already large error-correction overhead documented in the computation-limits literature [13], these scalings, once verified, belong in any cost accounting of quantum-memory architectures. The central open task is verification of H_1 – H_3 against the full source text.

References

- [1] TITLE: arXiv Query: `searchquery=&idlist=2610.03545&start=0&max_results=1` [2] What Must a Quantum-Memory Decoder Know About Temporally Correlated Noise?. arXiv:2610.03545v1. <https://arxiv.org/abs/2610.03545v1> [3] Noise and Bell's inequality. arXiv:1008.0667v2. <https://arxiv.org/abs/1008.0667v2> [4] Strongly correlated superconductivity. arXiv:1310.1481v1. <https://arxiv.org/abs/1310.1481v1> [5] Finite blocklength and moderate deviation analysis of hypothesis testing of correlated quantum states and application to classical-quantum channels with memory. arXiv:1612.01464v2. <https://arxiv.org/abs/1612.01464v2> [6] Theory of noise suppression in Λ -type quantum memories by means of a cavity. arXiv:1601.00157v1. <https://arxiv.org/abs/1601.00157v1> [7] Quantum and Classical Frontiers of Noise. arXiv:1612.03430v1. <https://arxiv.org/abs/1612.03430v1> [8] Noise Dynamics in the Quantum Regime.

arXiv:2311.17794v1. <https://arxiv.org/abs/2311.17794v1> [9] Decoherence and noise in open quantum system dynamics. arXiv:1605.07838v1. <https://arxiv.org/abs/1605.07838v1> [10] QNFO: The Problem-Substrate Mapping: A Framework for Honest Computational Investment [11] QNFO: Spectral Benchmarking of Holographic Quantum Simulations [12] QNFO: Auditing the BQNN: Does a Tunable Quantum Neural Network on Trapped-Ion and Superconducting Hardware Demonstrate a Route to Near-Term Quantum Advantage? [13] QNFO: The Physics of Computation: Fundamental Limits and the Honest Boundaries of Post-Classical Computing

Appendix A. Divergence report

D1 (calibration-scaling convention): DIVERGENT. Draft A writes the calibration bound with an explicit target-loss prefactor, $N \geq \theta^{-d_X}/\epsilon$, yielding $N = 10^7$ intervals at $\theta = 10^{-2}$, $d_X = 3$, $\epsilon = 0.1$ (and 10^5 for the two-interval variant). Draft C writes the scaling proportionally, $N_{\text{cyc}} \propto \theta^{-d_X}$, with no prefactor. The disagreement is a convention about whether the abstract states a prefactor; no supplied text states any prefactor, so the main text adopts draft C's proportional convention and records draft A's here. Under draft A's convention all interval counts in Section 5, item 3, are multiplied by $1/\epsilon = 10$ at $\epsilon = 0.1$.

D2 (illustrative angle): DIVERGENT. Drafts A/B use $\theta = 10^{-2}$ (Assumption A3); draft C uses $\theta = 0.05$ (Assumption A5). Both are illustrative choices not taken from [1]; the main text reports both instances.

D3 (mechanism derivation): SINGLE. Only draft C supplies the symmetric-disagreement derivation of the $\frac{1}{2}$ bound; drafts A and B state the bound without derivation. The main text uses draft C's derivation, explicitly labeled as an idealization (Assumption A1).

Appendix B. Claim attribution

Claim	Source drafts	Status
C1: Two sign models with identical syndrome statistics and the same SPP	A, B, C	CONVERGENT
C2: Worst-case fidelity loss of the model-agnostic decoder approaches $1/2$ despite full syndrome access	A, B, C	CONVERGENT
C3: Calibration interval count scales as θ^{-d_X} (H_1)	B, C (A with prefactor variant)	CONVERGENT on scaling; DIVERGENT on prefactor (Appendix A, D1)
C4: Two-interval schedule reduces scaling to θ^{-2} (H_2)	B, C	CONVERGENT
C5: Twirling-record statement (H_3)	B, C	CONVERGENT
C6: Symmetric-disagreement derivation of the $1/2$ bound	C	SINGLE
C7: Per-qubit error probability $p_X \approx \theta^2/4$ and $d_X = 3$ instance	B	SINGLE
C8: Storage-time projection $N_{\text{dis}} \approx \pi/(2\theta)$	C	SINGLE
C9: Draft A prefactor convention $N \geq \theta^{-d_X}/\epsilon$	A	SINGLE (DIVERGENT convention, Appendix A, D1)
C10: Significance argument via QNFO overhead of 10^2 – 10^3 operations per logical operation [13] and audit framing [12]	A, B	CONVERGENT