

The Calibration Loop as Fixed Point: Ontological Underdetermination of the Qubit by Its Own Metrology

The Calibration Loop as Fixed Point: Ontological Underdetermination of the Qubit by Its Own Metrology

Abstract

Every engineered qubit is certified by a closed calibration chain: drive pulses are calibrated by readout, readout is calibrated by states prepared with those drives, and tomography validates the whole loop using measurements calibrated within it. The physical referent of the qubit label — a collective circuit mode, an electronic transition, a spin — is therefore never fixed by any single external measurement. We formalize this chain as a self-consistent fixed-point problem, derive its contraction condition and iteration count explicitly, and exhibit two distinct physical ontologies of the same device — a strict two-level abstraction and a three-level anharmonic carrier with leakage amplitude c — that are observationally identical (total-variation distance exactly zero) under any measurement chain confined to the computational subspace. Using the single-shot readout fidelity of 99.991(1)% reported for trapped-ion qubits over 10^6 trials, we compute the per-shot error scale 9×10^{-5} , the calibration iterations needed to reach it, and the shot budget ($N \geq 2500$ at one sigma for $|c|^2 = 0.01$) required to break the degeneracy once a leakage-resolving observable exists. We conclude that circularity is broken only by measurements reaching outside the calibrated subspace or by cross-platform consistency requirements, and we state the conditions under which these claims would fail.

1. Introduction

When an experimentalist "drives a qubit," something physical is excited; when the same experimentalist "reads out the qubit," something physical is measured. The engineering question is blunt: if the billiard-ball picture is rejected, what exactly do the pulses manipulate, and what does the readout measure — excitations of what medium? The uncomfortable answer is that the answer is supplied by the calibration chain itself. The drive frequency is calibrated by finding a resonance; the Rabi rate is calibrated by fitting oscillations in a signal; the readout discriminator is calibrated by assigning outcomes to states prepared by the already-calibrated drive; tomography validates the loop using measurements whose calibration rests on the loop. Each step certifies the others. The loop is closed, and the physical ontology of the qubit — collective mode, spin, circuit degree of freedom, or a two-level abstraction layered over a multi-level device — appears only as an internal bookkeeping choice within the loop.

This paper takes that closure seriously as a formal problem. We make four contributions. First, we model the calibration hierarchy as an iterated map and derive its contraction condition and convergence rate in closed form (Sections 3.1, 4.3). Second, we exhibit two physically distinct ontologies of the same device — a strict two-level abstraction and a three-level anharmonic ladder — and prove that, under any measurement chain that does not resolve the third level, their outcome distributions are identical with total-variation distance exactly zero (Sections 3.2, 4.4). Third, we anchor the analysis quantitatively in the readout fidelity reported for trapped-ion qubits, deriving the error scale, the calibration iteration count needed to reach that scale, and the shot budget that would resolve the two ontologies once a leakage-resolving measurement is available (Section 4). Fourth, we identify the structural features — leakage-resolving spectroscopy and cross-platform consistency — that break the circularity, and we state explicitly what would falsify the thesis (Sections 4.5, 6).

The epistemic status of our claims is deliberately modest: everything quantitative in this paper is derived here from stated inputs with shown arithmetic, or is a clearly labeled projection. We do not claim that any real experiment is ontologically confused; we claim that the standard certification pipeline, taken on its own terms, cannot distinguish certain ontologies, and we quantify exactly how indistinguishable they are.

2. Background and Related Work

The literature bearing on this question falls into three groups: works that treat self-referential reasoning as a formal method, works on quantum readout, and control-theoretic treatments of systems whose working models are truncations of richer dynamics.

Self-referential diagnosis. The closest structural analogue is adaptive fault diagnosis [1], which studies determining which processors in a remote location are reliable by asking them "Yes or No" questions, where the processors are of three types: those that always tell the truth, those that always lie, and those that sometimes do either. Its central move, as stated in its abstract, is that using self-referential reasoning, along with earlier techniques, one can regard both the truth-tellers and the liars as reliable. This is exactly the logical move our calibration loop performs: a readout channel whose "honesty" is in question is calibrated by sources whose calibration depends on the readout, and the fixed point legitimizes the whole chain at once. The supplied summary of [1] gives no further technical detail; we use only the structural parallel — reliability of witnesses established by a closed chain of testimony — and note that in [1]'s setting the identity of the reporters is given, whereas our open question is whether self-referential consistency fixes the *referent* of the reports at all.

Readout as the anchor of the loop. The trapped-ion readout demonstration of [2] is our principal quantitative anchor. It demonstrates single-shot qubit readout with fidelity sufficient for fault-tolerant quantum computation, for two types of qubit stored in single trapped calcium ions; for an optical qubit stored in the $(4S_{1/2}, 3D_{5/2})$ levels of $^{40}\text{Ca}^+$ it achieves 99.991(1)% average readout fidelity in one million trials using time-resolved photon counting, and an adaptive measurement technique allows 99.99% fidelity. Two features matter for us. First, the readout is photonic: the measured medium (scattered photons) is visibly distinct from the qubit carrier (electronic levels) — precisely the referent-resolving structure whose absence in other platforms motivates our underdetermination question. Second, the quoted fidelities supply the error scale we use in Section 4. The abstract of [2] reports the fidelity but not how the assignment map was calibrated, which is exactly the gap our formalization targets.

Control of distributed and approximate models. A second group treats control of systems whose working model is a truncation of a richer reality — the same epistemic situation as the two-level abstraction of a multi-level device. Reference [3] models a nonhomogeneous flexible wing under unsteady aerodynamic loads as a distributed parameter system of two coupled partial differential equations, actuated at the tip by flaps, and investigates well-posedness of the underlying Cauchy problem before constructing the boundary control law. The methodological pattern — establish that the truncated closed-loop model is well-posed before designing the controller — mirrors our discipline of proving that the calibration fixed point exists before asking what it refers to. Reference [6] treats exponential input-to-state stabilization of diagonal infinite-dimensional systems with uncertain, time-varying input delays, using a constant-delay predictor feedback controller designed on a truncated finite-dimensional model; the truncation there is explicitly a modeling choice whose stability consequences must be bounded, exactly as leakage out of the two-level abstraction must be bounded in our model. Reference [7] proposes a control-oriented notion of finite state approximation, approximating discrete-time plants with finite-valued sensors and actuators by deterministic finite memory systems for certified-by-design controller synthesis; its framing — the approximation is

built for the controller, not derived from the physics — is the control-theoretic statement of our claim that the qubit ontology is defined by the control loop rather than discovered by it. Certification for control, we emphasize, is weaker than certification of ontology.

Hierarchical and constrained control. Reference [4] describes the control of a rigid-wing airborne wind energy system across all operational phases — take-off, transition to power generation, pumping energy generation cycles, transition to hovering, and landing — via a hierarchical control structure with control design at all levels. This is the architectural template for a calibration hierarchy: distinct phases (drive calibration, readout calibration, tomography), each with its own control law, coupled sequentially. Reference [5] introduces Input Constrained Control Barrier Functions (ICCBFs) for synthesizing safety-critical controllers for nonlinear control-affine systems with input constraints, identifying a subset of the safe set and constructing a controller rendering it forward invariant, with the feedback controller given as the solution to a quadratic program. We borrow the forward-invariance idea as a question we pose but do not fully answer: is there a forward-invariant "ontological safe set" — a region of parameter space in which the two-level abstraction is certified valid — and what barrier function enforces it?

Quantum-compliant formulations. Reference [8] formulates network epidemic control — isolating locations by banning movement, modeled as link removal — in a quantum-compliant way, building on discrete-time susceptible-infected-susceptible and susceptible-infected-removed network models. Its relevance is methodological: it shows the quantum formalism applied as a modeling layer to a non-quantum control loop, supporting our view that the quantum-classical distinction is not what generates the self-referential structure — the closure of the calibration chain is.

QNFO corpus. The corpus document "No Thing There: Control, Readout, and Self-Referential Metrology in Engineered Quantum Systems" [9] is a pedagogical bridge between the qubit ontology critique and working physics, explaining what control pulses actually manipulate and what readout measures across superconducting transmon, trapped-ion, and spin qubit platforms, and including an electrical seesaw analogy and a self-referential metrology formalization. Our Sections 3–4 can be read as a formal companion to [9]: where [9] explains, we derive. Two further corpus entries, "Stabilization of Gottesman-Kitaev-Preskill States" [10] and "A Pre-Registered Falsification of Deterministic Measurement-Triggered Relaxation" [11], are supplied with summaries too thin to support substantive statements; we cite them only as adjacent corpus context and draw no claims from them, noting this as an explicit limitation in Section 6.

3. Methods

3.1 The calibration loop as an iterated map

Let $\mathcal{H}_{\text{sub}}$ denote the computational subspace (dimension $d = 2$) and $\mathcal{H}_{\text{full}}$ the full Hilbert space of the physical carrier. The calibration chain consists of stages indexed $j = 1, \dots, J$: drive calibration, Rabi calibration, readout discriminator calibration, and tomographic validation. Each stage j has a parameter vector $\theta_j \in \mathbb{R}^{n_j}$ and a validation statistic r_j computed from data produced using the current settings of all other stages. Collecting $\Theta = (\theta_1, \dots, \theta_J)$, the update is

$$\Theta_{k+1} = F(\Theta_k), \quad F : \mathbb{R}^n \rightarrow \mathbb{R}^n.$$

A calibrated apparatus is a fixed point Θ_* with $F(\Theta_*) = \Theta_*$. The self-referential character is that F 's definition at every stage involves measurements taken through the apparatus being calibrated — the situation of [1], where the reliability of the witnesses is established by the testimony chain itself.

Near a fixed point, write $\Theta_k = \Theta_\star + \delta_k$. For a scalar gain model of one stage (the case relevant to drive-amplitude calibration), the update is

$$x_{k+1} = x_k + g(r_k - r_\star),$$

where x_k is the drive parameter, g the loop gain, r_k the observed statistic, and $r_\star$ the target. Linearizing, $r_k - r_\star = s(x_k - x_\star)$ with sensitivity s , gives the error recursion

$$e_{k+1} = (1 - g s) e_k, \quad e_k \equiv x_k - x_\star.$$

The map is a contraction whenever $|1 - g s| < 1$, and then $e_k = e_0 |1 - g s|^k$. This is the formal content of "the loop converges": the fixed point is attractive and the chain self-consistent — but nothing in the contraction argument refers to the physical substrate. The fixed point is a property of the loop, not of $\mathcal{H}_{\text{full}}$. (A degenerate special case arises if the loop is instead written as a pure gain product $A \mapsto k_1 k_2 k_3 k_4 A$: then a nontrivial fixed point requires $k_1 k_2 k_3 k_4 = 1$, and if that holds, *every* amplitude is a fixed point — a continuum of self-consistent calibrations. We adopt the contraction model as primary and treat the gain-product condition as the marginal, non-contractive limit; see Appendix A.)

3.2 Two ontologies, one loop

We formalize the underdetermination conjecture with two models:

- **Model A (two-level abstraction):** $\mathcal{H}_A = \text{span}\{|0\rangle, |1\rangle\}$, with calibrated POVM $\{\Pi_0^{(A)}, \Pi_1^{(A)}\}$ acting on $\mathcal{H}_A$.
- **Model B (anharmonic three-level carrier):** $\mathcal{H}_B = \text{span}\{|0\rangle, |1\rangle, |2\rangle\}$, with the same calibrated POVM extended so that $\Pi_1^{(B)}$ projects onto the span of $\{|1\rangle, |2\rangle\}$ — i.e., the readout does not resolve the third level, the generic situation when the readout transition addresses only the $|0\rangle \leftrightarrow |1\rangle$ distinction.

A control pulse with residual leakage amplitude c produces, in Model B, the state

$$|\psi_B\rangle = \sqrt{1 - |c|^2} |1\rangle + c |2\rangle,$$

while Model A assigns the pure state $|1\rangle$. The outcome distributions are

$$p_A(1) = 1, \quad p_B(1) = \langle \psi_B | \Pi_1^{(B)} | \psi_B \rangle = (1 - |c|^2) + |c|^2 = 1.$$

The total-variation distance is

$$D_{\text{TV}}(p_A, p_B) = \frac{1}{2} \sum_m |p_A(m) - p_B(m)| = 0.$$

Both ontologies are observationally identical under the full calibrated measurement set whenever the chain does not resolve level $|2\rangle$. The tomographic reconstruction — a function only of outcome distributions — is the same object in both cases. This is the precise sense in which the qubit ontology is underdetermined by its metrology: the loop fixes the *operational* qubit, and the operational qubit is compatible with multiple substrates.

Connection to weak-drive perturbation theory. The leakage amplitude is not a free parameter: in a weakly anharmonic ladder driven near the $|0\rangle \leftrightarrow |1\rangle$ transition with Rabi amplitude Ω and anharmonicity-scale detuning Δ on the $|1\rangle \leftrightarrow |2\rangle$ transition, adiabatic elimination of $|2\rangle$ (setting $\dot{c}_2 \approx 0$ in $\dot{c}_2 = -i\Delta c_2 - i\frac{\Omega}{2}c_1$) gives $c_2 \approx \frac{\Omega}{2\Delta}c_1$, so $|c|^2 \lesssim \left(\frac{\Omega}{2\Delta}\right)^2$. We use this to relate the abstract leakage amplitude to an operating point in Section 4.6.

3.3 Breaking the degeneracy

If the measurement chain is extended so that it resolves $|2\rangle$ (a leakage-resolving measurement, with Π_1 projecting on $|1\rangle$ alone and a third outcome for $|2\rangle$), the distributions become

$$p_A(1) = 1, \quad p_B(1) = 1 - |c|^2, \quad p_B(2) = |c|^2,$$

and

$$D_{\text{TV}}(p_A, p_B) = \frac{1}{2} (|c|^2 + |c|^2) = |c|^2.$$

The degeneracy is broken exactly to the extent that the measurement reaches outside the computational subspace. With N independent shots, resolving a distribution difference of order D_{TV} requires roughly $N \sim D_{\text{TV}}^{-2}$ shots (standard binomial scaling, stated as a modeling assumption for the projections below).

3.4 Cross-platform constraint

A second route out of circularity is structural rather than statistical: two platforms whose calibration loops close differently (photonic readout as in [2], circuit readout, spin resonance) cannot both be saved by a single shared ontology unless that ontology is substrate-specific. Requiring one consistent assignment of "what is excited" across platforms with different loop topologies is a constraint no single loop can generate — the analogue, in our setting, of the multi-witness consistency that lets the self-referential diagnosis of [1] classify its processors.

4. Analysis

Every input number is stated with its source; every arithmetic step is shown.

Input 1 (from [2]): average readout fidelity $f_{\text{ro}} = 99.991\% = 0.99991$ for the optical qubit in $^{40}\text{Ca}^+$, over $N_{\text{trials}} = 10^6$ trials.

Input 2 (from [2]): the adaptive measurement technique allows fidelity $f_{\text{ad}} = 99.99\% = 0.99999$.

Input 3 (model assumption, Section 3.2): leakage amplitude $|c| = 0.1$, i.e., leakage probability $|c|^2 = 0.01$.

Input 4 (model assumption, Section 3.1): loop gain-sensitivity product $g s = 0.5$, giving contraction factor $|1 - g s| = 0.5$; initial calibration error $e_0 = 0.1$ in units of the drive parameter.

Input 5 (model assumption, Section 3.2): weak-drive operating point $\Omega/\Delta = 0.1$; tomographic shot budget $N_{\text{tom}} = 10^4$ shots per element.

4.1 Readout error scale

The per-shot readout error probability is

$$\varepsilon_{\text{ro}} = 1 - f_{\text{ro}} = 1 - 0.99991 = 9 \times 10^{-5}.$$

Over $N_{\text{trials}} = 10^6$ trials, the expected number of misidentifications is

$$\mu = N_{\text{trials}} \times \varepsilon_{\text{ro}} = 10^6 \times 9 \times 10^{-5} = 90.$$

The one-standard-error uncertainty on f_{ro} over 10^6 trials is

$$\sigma_f = \sqrt{\frac{f_{\text{ro}}(1 - f_{\text{ro}})}{N_{\text{trials}}}} = \sqrt{\frac{0.99991 \times 9 \times 10^{-5}}{10^6}} = \sqrt{8.99919 \times 10^{-11}} \approx 9.49 \times 10^{-6},$$

i.e., about 10^{-6} in fidelity units, consistent with the "(1)" in the last quoted digit of 99.991(1)%.

4.2 Adaptive-technique improvement

The adaptive-technique error is

$$\varepsilon_{\text{ad}} = 1 - f_{\text{ad}} = 1 - 0.99999 = 1 \times 10^{-5},$$

and the ratio of adaptive to time-resolved error is

$$\frac{\varepsilon_{\text{ad}}}{\varepsilon_{\text{ro}}} = \frac{1 \times 10^{-5}}{9 \times 10^{-5}} = \frac{1}{9} \approx 0.11111.$$

The adaptive technique reduces the error to roughly one-ninth of the time-resolved value, per the two fidelities quoted in [2].

4.3 Calibration-loop convergence

With contraction factor $|1 - g s| = 0.5$ and $e_0 = 0.1$, the error after k iterations is $e_k = 0.1 \times 0.5^k$. Taking as threshold the readout error scale $\varepsilon_{\text{ro}} = 9 \times 10^{-5}$ (calibration error should not exceed the noise floor it is calibrated against), we require

$$0.1 \times 0.5^k \leq 9 \times 10^{-5} \implies 0.5^k \leq 9 \times 10^{-4}.$$

Taking logarithms:

$$k \geq \frac{\ln(9 \times 10^{-4})}{\ln(0.5)} = \frac{2.19722 - 9.21034}{-0.69315} = \frac{-7.01312}{-0.69315} = 10.118.$$

Since k is an integer, we verify discretely: $0.5^{10} = 9.7656 \times 10^{-4} > 9 \times 10^{-4}$, while $0.5^{11} = 4.8828 \times 10^{-4} \leq 9 \times 10^{-4}$. Hence

$$k_{\text{min}} = 11 \text{ iterations,}$$

with $e_{11} = 0.1 \times 4.8828 \times 10^{-4} = 4.8828 \times 10^{-5} < 9 \times 10^{-5} < e_{10} = 9.7656 \times 10^{-5}$.

4.4 Ontological distinguishability

From Section 3.3, with $|c|^2 = 0.01$ (Input 3):

$$D_{\text{TV}}(p_A, p_B) = |c|^2 = 0.01 \quad (\text{leakage-resolving measurement}),$$

$$D_{\text{TV}}(p_A, p_B) = 0 \quad (\text{subspace-restricted measurement}).$$

4.5 Shot budget and noise-floor margin

To resolve a distribution difference $\Delta p = |c|^2 = 0.01$ at one-standard-error significance, the binomial standard error $\sqrt{p(1-p)/N}$ must be $\leq \Delta p$. Taking the conservative $p \approx 1/2$ bound, $\sqrt{1/(4N)} \leq 0.01$ gives

$$N \geq \frac{1}{4 \times (0.01)^2} = \frac{1}{4 \times 10^{-4}} = 2500 \text{ shots.}$$

For a three-sigma requirement, $3\sqrt{1/(4N)} \leq 0.01$ gives

$$N \geq \frac{9}{4 \times 10^{-4}} = 22500 \text{ shots.}$$

Both figures are far below the 10^6 trials already demonstrated in [2], so the shot budget is not the obstacle; the obstacle is whether the measurement chain *implements* a leakage-resolving observable at all. The margin of a leakage signal over the readout noise floor is

$$\frac{|c|^2}{\varepsilon_{\text{ro}}} = \frac{0.01}{9 \times 10^{-5}} = 111.1\bar{1} \approx 111,$$

so a leakage-resolving extension of a [2]-class readout would see the three-level ontology clearly, provided the extension exists.

4.6 Weak-drive consistency check (complementary model)

At the operating point $\Omega/\Delta = 0.1$ (Input 5), the adiabatic bound of Section 3.2 gives

$$\left(\frac{\Omega}{2\Delta}\right)^2 = \left(\frac{0.1}{2}\right)^2 = (0.05)^2 = 2.5 \times 10^{-3},$$

i.e., $|c|^2 \lesssim 2.5 \times 10^{-3}$, tenfold below the Input-3 value of 0.01 — so the shot budgets of Section 4.5 are conservative for weak drives. Conversely, with $N_{\text{tom}} = 10^4$ shots per tomographic element, the worst-case binomial resolution at $\hat{p} = 1/2$ is

$$\sigma = \sqrt{\frac{0.5 \times 0.5}{10^4}} = \sqrt{2.5 \times 10^{-5}} = 5.0 \times 10^{-3},$$

which exceeds 2.5×10^{-3} by a factor of 2: at this operating point and budget, population tomography cannot resolve the leakage. Resolving 2.5×10^{-3} at one sigma would require

$$N \geq \frac{0.25}{(2.5 \times 10^{-3})^2} = \frac{0.25}{6.25 \times 10^{-6}} = 4.0 \times 10^4 \text{ shots per element,}$$

a $4\times$ increase over the assumed budget. At a stronger drive, $\Omega/\Delta = 0.4$, the bound becomes $(0.2)^2 = 4.0 \times 10^{-2}$, an order of magnitude above the 5.0×10^{-3} floor — strong-drive tomography would distinguish the ontologies. These are projections under the stated assumptions (Ω/Δ values, worst-case $\hat{p} = 1/2$, no systematic errors).

5. Results

All numbers below are computed in Section 4 with shown arithmetic, or are labeled projections. No experimental or simulated data are reported in this paper.

- **R1 (readout error scale, computed from [2]).** Per-shot error $\varepsilon_{\text{ro}} = 9 \times 10^{-5}$; expected misidentifications in 10^6 trials: $\mu = 90$; one-standard-error on the fidelity $\approx 9.4864 \times 10^{-6}$, consistent with the reported 99.991(1)%.
- **R2 (adaptive improvement, computed from [2]).** $\varepsilon_{\text{ad}} = 1 \times 10^{-5}$, a factor ≈ 0.11111 (one-ninth) of the time-resolved error.
- **R3 (calibration convergence, computed under stated assumptions $gs = 0.5$, $e_0 = 0.1$).** The loop reaches calibration error below the readout noise floor after $k_{\text{min}} = 11$ iterations ($e_{11} = 4.8828 \times 10^{-5}$).

- **R4 (ontological degeneracy, computed under stated model).** Under subspace-restricted measurements, $D_{TV} = 0$ exactly: the two-level and three-level ontologies are observationally identical, and tomography cannot distinguish them. Under a leakage-resolving measurement with $|c| = 0.1$, $D_{TV} = 0.01$.
- **R5 (shot budget, computed under binomial scaling).** $N \geq 2500$ shots for one-sigma resolution of $\Delta p = 0.01$; $N \geq 22500$ for three-sigma; margin over the readout error floor $|c|^2/\varepsilon_{ro} \approx 111$.
- **R6 (weak-drive projection, labeled).** At $\Omega/\Delta = 0.1$ and $N_{tom} = 10^4$, the leakage bound 2.5×10^{-3} sits a factor of 2 below the statistical floor 5.0×10^{-3} ; one-sigma resolution requires $N \geq 4.0 \times 10^4$ shots per element. At $\Omega/\Delta = 0.4$ the bound 4.0×10^{-2} exceeds the floor and strong-drive tomography distinguishes the ontologies. Uncertainty is dominated by the assumed operating points and neglected systematic errors.
- **R7 (structural projection, labeled).** If a platform's full calibrated measurement set is confined to the computational subspace — the generic case for dispersive readout used only for qubit-state discrimination — then no finite shot budget distinguishes the two ontologies on that platform alone, because $D_{TV} = 0$ identically. The empirical risk of this projection lies in whether real readout chains are truly subspace-confined, a platform-specific question not settled here.

6. Discussion

What the result does and does not show. R4 shows that a specific pair of ontologies is indistinguishable by a specific measurement class (subspace-confined population tomography). It does not show that qubit ontology is unknowable, nor that all ontologies are equivalent. The equivalence is quantitative and degrades predictably: R6 shows that at larger Ω/Δ the leakage grows as $(\Omega/2\Delta)^2$ and becomes resolvable. The underdetermination is therefore a statement about the weak-drive, high-fidelity regime that fault-tolerant operation favors, not a universal claim.

Limitations of the model. The adiabatic-elimination step is a second-order approximation valid for $\Omega \ll \Delta$; it neglects counter-rotating terms, pulse-shape effects (a shaped pulse has a different leakage spectrum than a square pulse), and levels beyond $|2\rangle$. The calibration model of Section 3.1 is scalar and linear; a real calibration vector has cross-couplings between stages, and the contraction analysis does not establish uniqueness of the fixed point in the coupled case. The noise-floor comparisons treat assignment error and shot noise as independent; correlated errors (e.g., drift during a tomography run) would inflate the floor, strengthening R4, but we have not modeled them. The leakage amplitude $|c| = 0.1$ and the operating point $\Omega/\Delta = 0.1$ are assumptions of this analysis, not measurements; results R3–R6 scale with them.

Failure modes of the thesis. The claim "tomography cannot pin the referent" would be falsified by a demonstration that standard population tomography, at ordinary shot budgets, systematically distinguishes the two-level and three-level ontologies — i.e., if measured leakage residuals exceeded $(\Omega/2\Delta)^2$ by an unmodeled mechanism. It would also be weakened if the calibration loop were shown to have a unique fixed point that depends on substrate-specific parameters in a way that tomography could extract; in that case the loop itself would carry ontological information, contradicting our underdetermination claim. Finally, the cross-platform argument of Section 3.4 assumes that platforms with different loop topologies cannot share a substrate-neutral ontology; a demonstration of a single ontology consistently certified across photonic-readout and circuit-readout platforms would weaken it.

Open questions. Three questions remain open. (i) Does a forward-invariant "ontological safe set" exist in the sense of Section 2's borrowing from [5], i.e., a region of parameter space in which the two-level abstraction is certified valid, enforced by a barrier-type condition on the leakage amplitude? (ii) Does the coupled, multi-stage calibration

map admit multiple fixed points, and if so, do distinct fixed points correspond to distinct ontologies rather than distinct calibrations of one ontology? (iii) Are real readout chains on any platform strictly subspace-confined, or do weak out-of-subspace couplings always leak ontological information at some shot budget?

7. Conclusion

We modeled the closed calibration chain of an engineered qubit as an iterated map with a contraction condition, and showed that its fixed point is a property of the loop, not of the physical substrate. Two ontologies of the same device — a strict two-level abstraction and a three-level anharmonic carrier with leakage amplitude $|c| = 0.1$ — are observationally identical with total-variation distance exactly zero under any measurement chain confined to the computational subspace. Anchoring the analysis in the trapped-ion readout fidelity of $99.991(1)\%$ reported in [2], we computed the per-shot error scale $\varepsilon_{\text{ro}} = 9 \times 10^{-5}$, the calibration iteration count $k_{\text{min}} = 11$ needed to reach that scale under the assumed gain $gs = 0.5$, and the shot budgets $N \geq 2500$ (one sigma) and $N \geq 22500$ (three sigma) that would break the degeneracy once a leakage-resolving observable exists, with a signal-to-floor margin $|c|^2/\varepsilon_{\text{ro}} \approx 111$. The circularity of the calibration loop is broken only by measurements reaching outside the calibrated subspace or by cross-platform consistency requirements; within the loop, the qubit's ontology is a bookkeeping choice. We stated the conditions under which each claim would fail, and we regard the falsification paths of Section 6 as the proper test of the thesis.

References

- [1] Adaptive Fault Diagnosis using Self-Referential Reasoning. arXiv:1407.6255v1. <https://arxiv.org/abs/1407.6255v1>
- [2] High-fidelity readout of trapped-ion qubits. arXiv:0802.1684v2. <https://arxiv.org/abs/0802.1684v2>
- [3] Boundary Control of a Nonhomogeneous Flexible Wing with Bounded Input Disturbances. arXiv:1709.00759v2. <https://arxiv.org/abs/1709.00759v2>
- [4] Control of a Rigid Wing Pumping Airborne Wind Energy System in all Operational Phases. arXiv:2006.11141v1. <https://arxiv.org/abs/2006.11141v1>
- [5] Safe Control Synthesis via Input Constrained Control Barrier Functions. arXiv:2104.01704v1. <https://arxiv.org/abs/2104.01704v1>
- [6] Exponential input-to-state stabilization of a class of diagonal boundary control systems with delay boundary control. arXiv:2003.05711v1. <https://arxiv.org/abs/2003.05711v1>
- [7] A Control-Oriented Notion of Finite State Approximation. arXiv:1105.3788v3. <https://arxiv.org/abs/1105.3788v3>
- [8] A Quantum-Compliant Formulation for Network Epidemic Control. arXiv:2509.00337v1. <https://arxiv.org/abs/2509.00337v1>
- [9] QNFO: No Thing There: Control, Readout, and Self-Referential Metrology in Engineered Quantum Systems
- [10] QNFO: Stabilization of Gottesman-Kitaev-Preskill States
- [11] QNFO: A Pre-Registered Falsification of Deterministic Measurement-Triggered Relaxation