

# Operational Circularity of Qubit Ontology: Excitations Defined by Their Own Readout

## Operational Circularity of Qubit Ontology: Excitations Defined by Their Own Readout

### Abstract

Is the qubit a pre-existing object that control pulses and readout merely probe, or a two-level subsystem constituted by the very operations that define it? We formalize this question as a self-referential calibration problem: an ontology map  $\Phi$  assigns Hilbert-space identifications to operational data, and calibration seeks a fixed point of the composite loop  $T = \Gamma \circ \Phi$ , in which states are defined by measurements that are themselves defined by states. We distinguish two regimes. When  $T$  is a contraction with Lipschitz constant  $L < 1$ , Banach's fixed-point theorem guarantees a unique self-consistent identification and the circularity is benign, in the same sense that ordinary instrument calibration is benign. When the loop admits multiple fixed points or a continuous equivariant symmetry, physically inequivalent ontologies—an energy eigenmode, a field excitation, a collective variable—remain compatible with identical operational data, and the circularity masks genuine ontological underdetermination. We exhibit an explicit two-branch toy model in which the final ontology is fixed by the initial guess rather than by the data. Using the single-shot readout fidelity of 99.991% reported for trapped-ion optical qubits, we compute the shot budget required to resolve ontology-discriminating discrepancies:  $N \approx 2.65 \times 10^6$  shots for a discrepancy  $\delta = 10^{-3}$  at 1% failure probability. We conclude that uniqueness is a property of the calibration dynamics, not of the qubit, and state falsifiable conditions for both outcomes.

### 1. Introduction

Quantum information processing speaks the language of qubits: two-level systems that are initialized, manipulated by control pulses, and measured. The working assumption of much of the field is that the qubit is a pre-existing object—a particular two-dimensional subspace of some Hilbert space—that control and readout merely access. This paper questions that assumption in a specific, formalizable way.

The operational loop is unavoidable. To implement a control pulse, one needs a model of the qubit: its transition frequency, its coupling to the drive, its relaxation channels. To validate that model, one runs tomography; but tomography is defined by a measurement apparatus whose calibration itself presupposes a model of the states being measured. States are defined by measurements, and measurements are defined by states. We call this **operational circularity**. The question is not whether the circle exists—it does—but when it is benign and when it is not.

We distinguish two regimes:

1. **Benign circularity (calibration self-consistency).** The loop converges to a unique operational identification: a single assignment of Hamiltonian parameters, readout observables, and error channels that is stable under further calibration. This is the regime of ordinary instrument calibration, and it is formally a contraction fixed-point iteration.
1. **Ontological underdetermination.** Multiple, physically inequivalent identifications of the controlled entity—e.g., "the qubit is an energy eigenmode," "the qubit is a localized field excitation," "the qubit is a collective variable of many degrees of freedom"—are each exactly self-consistent with the same operational data. In this regime, no amount of calibration *inside the loop* distinguishes them.

Our contribution is threefold. First, we give a formal definition of the calibration fixed-point problem (Section 3). Second, we derive the contraction condition under which the fixed point is unique, and exhibit an explicit toy model in which two inequivalent ontologies are both exact fixed points (Section 4), with all arithmetic shown. Third, we compute, from a published readout fidelity figure [4], the experimental shot budget needed to test whether ontology-discriminating observables can be resolved at all. Section 6 discusses limitations and what would falsify each claim.

The framing draws on the QNFO corpus, particularly the treatment of self-referential metrology across transmon, trapped-ion, and spin platforms [9], the critique of particle ontology projected onto field-theoretic reality [12], and the constructive question of what minimal ontological commitments survive that critique [10]. The methodological stance—falsifiability as a pledge rather than a slogan—is adopted from [11]: every quantitative claim below is either derived with shown arithmetic or explicitly labeled a projection with stated assumptions.

## 2. Background and Related Work

---

The relevant literature falls into two groups: formal ontology engineering, which supplies the vocabulary of "ontology" we adapt, and the QNFO corpus, which supplies the physical critique; one experimental result supplies the empirical anchor.

**Ontology as a formal artifact.** The Gene Ontology (GO) project is described in [1] as the largest resource for cataloguing gene function, combining solid conceptual underpinnings with practical features that have made it widely adopted and essential for data analysis; the entry supplies a primer introducing the structure of the ontology. We borrow from this tradition the idea that an ontology is an explicit, structured assignment of entities and relations—not a metaphysical thesis but a representational commitment. Our "ontology map"  $\Phi$  (Section 3) is exactly such a commitment, applied to Hilbert-space identifications rather than gene functions.

Crucially, ontologies can be used without understanding their construction, and this is a double-edged property: [2] states that the GO is sufficiently simple that it can be used without deep understanding of its structure or development, which is both a strength and a weakness, and discusses common misinterpretations, biases, and remedies. This is precisely the failure mode we attribute to qubit ontology: the two-level abstraction is usable without knowing what it commits us to, and the misuse is invisible until the abstraction is pushed beyond its calibration domain.

Ontologies are not static. [5] analyzes how mappings between popular life-science ontologies change, noting that ontologies change continuously and that the evolution of ontology mappings had received little attention. This supports our dynamic treatment: the ontology identification of a qubit is not fixed once but is the output of an evolving calibration process, and the stability of that process is the object of study.

Ontology matching—the problem of finding correspondences between entities in different ontologies so that they can be aligned or merged—is addressed by [3], which presents GraphMatcher, an ontology matching system using a graph attention approach. Our multi-platform question (does calibration on a transmon, an ion, and a spin qubit converge to the *same* identification?) is structurally an ontology-matching problem: we ask whether the operational ontologies of different platforms have a common alignment or only a family of partial correspondences.

That ontologies can be built for domains far from biology is shown by [7], which presents a community-developed domain ontology for magnetic materials, motivated by complex multiscale behaviour and the coexistence of multiple unit systems that pose persistent challenges for data exchange and interpretation. The unit-system problem is a direct analogue of our problem: the same physical magnet can be described in multiple unit conventions, and the ontology must manage the equivalence. Similarly, [6] proposes an ontology-based approach to formally define qualitative conversational features by deriving quantitative definitions from linguistic descriptors—an example of pinning down an entity (a conversational feature) that has no pre-existing sharp definition, exactly the situation of the qubit. Finally, [8] develops an ontological model for business intelligence in which relations are mined from web semantics using decision trees; the summary supplied gives no further technical detail, but it illustrates the general pattern that an ontology can be *derived from data* rather than stipulated—the direction of constitution our fixed-point formalism makes precise.

**The physical critique.** [9] is described as 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 a self-referential metrology formalization. Our Section 3 is a sharpening of that program; the supplied summary does not state its formal results in detail, so we treat it as motivation rather than as a source of theorems. [12] supplies the polemical thesis that the qubit-gate-circuit model is an epistemic failure—a projection of particle ontology onto relational, field-theoretic reality; its supplied entry states only a revision note, so we cite it as the hypothesis our formalism is meant to discipline rather than as a source of technical claims. [10], the companion paper, asks what comes next if that diagnosis holds, stripping quantum mechanics to its minimal ontological commitments and surveying constructive paradigms for post-particle computation; our uniqueness/underdetermination dichotomy is a tool for deciding whether the diagnosis holds. [11] states five principles, a concrete portfolio, and a Falsification Pledge; we adopt falsifiability as a design constraint on our claims (Section 6).

**Experimental anchor.** [4] 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  $^{40}\text{Ca}^+$ , an average readout fidelity of 99.991(1)% over one million trials is achieved using time-resolved photon counting, with an adaptive technique allowing 99.99% fidelity. We use this figure as the empirical input for our shot-budget computation (Section 4.3): it fixes the per-shot noise floor against which ontology-discriminating signals must compete.

### 3. Methods

---

#### 3.1 Operational ontology maps

Let  $\mathcal{D}$  be the space of operational data: sequences of control waveforms, their applied settings, and readout records. Let  $\mathcal{O}$  be the space of Hilbert-space identifications: assignments of the form "the manipulated degree of freedom is the two-level system spanned by  $\{|0\rangle, |1\rangle\}$  with Hamiltonian  $H = \frac{\hbar}{2}(\omega_0\sigma_z + \Omega_x\sigma_x + \Omega_y\sigma_y)$ , readout observable  $M$ , error channels  $\{L_k\}$ ," together with the physical interpretation of each symbol (eigenmode, field excitation, collective variable).

An **ontology map** is a function

$$\Phi : \mathcal{D} \rightarrow \mathcal{O},$$

which takes a body of operational data and returns an identification. Calibration is the search for a **self-consistent datum**  $x^* \in \mathcal{D}$  satisfying

$$x^* = \Gamma(\Phi(x^*)),$$

where  $\Gamma : \mathcal{O} \rightarrow \mathcal{D}$  is the **realization map**: given an identification, it produces the operational data that the identified system would generate under the calibrated protocol. A pair  $(x^*, \Phi(x^*))$  satisfying this equation is a **fixed point of the operational loop**.

#### 3.2 Benign circularity: the contraction condition

Let  $d_{\mathcal{D}}$  be a metric on  $\mathcal{D}$  (e.g., total-variation distance between predicted readout distributions plus weighted parameter distance). Define the composite calibration operator

$$T = \Gamma \circ \Phi : \mathcal{D} \rightarrow \mathcal{D}.$$

**Proposition 1 (Benign circularity).** If  $T$  is a contraction with Lipschitz constant  $L < 1$ ,

$$d_{\mathcal{D}}(T(x), T(x')) \leq L d_{\mathcal{D}}(x, x') \quad \text{for all } x, x' \in \mathcal{D},$$

then by the Banach fixed-point theorem  $T$  has a unique fixed point  $x^*$ , and iteration from any initial guess  $x_0$  converges to it, with

$$d_{\mathcal{D}}(x_n, x^*) \leq L^n d_{\mathcal{D}}(x_0, x^*).$$

Convergence to tolerance  $\varepsilon$  requires

$$n \geq \frac{\ln(\varepsilon/d_{\mathcal{D}}(x_0, x^*))}{\ln L}.$$

Since  $0 < L < 1$ ,  $\ln L < 0$ ; with  $\varepsilon < d_{\mathcal{D}}(x_0, x^*)$  the numerator is negative and  $n$  is positive, as required. In this regime the circularity is benign: the loop defines a unique operational identification regardless of the starting ontology.

#### 3.3 Ontological underdetermination

If  $T$  is not globally contractive, uniqueness can fail. The physically interesting failure is not chaos but **multiplicity**: several fixed points, each corresponding to a different physical interpretation of the same symbols. We say the qubit ontology is **underdetermined** on a data set  $x^*$  if there exist two ontology maps  $\Phi_A, \Phi_B$  and realization maps  $\Gamma_A, \Gamma_B$  such that

$$x^* = \Gamma_A(\Phi_A(x^*)) = \Gamma_B(\Phi_B(x^*)),$$

with  $\Phi_A(x^*)$  and  $\Phi_B(x^*)$  physically inequivalent (differing in the interpretation of the controlled entity, not merely in parameter values). A second, continuous route to underdetermination is **equivariance**: if a group  $G$  acts on the identification space and  $T$  commutes with the action, then  $g \cdot x^*$  is a fixed point whenever  $x^*$  is, and no iteration inside the loop selects a member of the orbit. For a qubit, the natural candidate is  $SO(3)$  acting on the Bloch-sphere axis labeling;  $\dim SO(3) = 3$ , so three continuous descriptive parameters (two axis angles plus one basis phase) are invisible to any equivariant operational statistic.

### 3.4 Discriminating observable and shot budget

To escape the loop, one needs an observable whose expectation differs between candidate ontologies by an amount  $\delta$ , resolved against readout noise. Given single-shot readout fidelity  $F$ , the per-shot error probability is  $p_e = 1 - F$ . By Hoeffding's inequality, estimating a Bernoulli mean to additive accuracy  $\delta$  with failure probability  $\alpha$  requires

$$N \geq \frac{\ln(2/\alpha)}{2\delta^2}$$

shots. This gives the minimal experimental budget for any ontology-discriminating test, instantiated in Section 4.3 with the fidelity from [4].

## 4. Analysis

All numbers in this section are computed here from stated inputs. The only empirical input is the readout fidelity from [4]; everything else is a derivation or an explicitly labeled projection.

### 4.1 Readout error probability and error counts

**Input 1 (from [4]).** Average single-shot readout fidelity for the optical qubit in  $^{40}\text{Ca}^+$ :  $F = 99.991\%$ , stated uncertainty (1) in the last digit, over  $N_{\text{trials}} = 10^6$  trials.

**Derivation 4.1.** Per-shot error probability:

$$p_e = 1 - F = 1 - 0.99991 = 9 \times 10^{-5}.$$

One readout shot in  $1/p_e = 1/(9 \times 10^{-5}) \approx 1.1111 \times 10^4$  is wrong on average. Over the reported  $10^6$  trials, the expected number of wrong single-shot outcomes is

$$N_{\text{err}} = p_e N_{\text{trials}} = 9 \times 10^{-5} \times 10^6 = 90.$$

The quoted uncertainty (1) corresponds to  $\pm 1 \times 10^{-5}$  in  $F$ , i.e.  $\pm 10$  counts in  $N_{\text{err}}$ ; a binomial fluctuation of order  $\sqrt{N_{\text{err}}} = \sqrt{90} \approx 9.49$  is of the same order, so the quoted fidelity uncertainty is consistent with pure counting statistics.

### 4.2 Statistical floor and the contraction regime

**Derivation 4.2 (tomographic statistical floor; projection).** Assume  $N = 10^4$  shots per tomographic setting (a stated assumption, not a measurement) and worst-case Bernoulli parameter  $p = 1/2$ :

$$\sigma_p = \sqrt{\frac{p(1-p)}{N}} = \sqrt{\frac{0.25}{10^4}} = \sqrt{2.5 \times 10^{-5}} \approx 5.00 \times 10^{-3}.$$

The readout systematic floor is  $p_e = 9 \times 10^{-5}$ , smaller by the factor

$$\frac{\sigma_p}{p_e} = \frac{5.00 \times 10^{-3}}{9 \times 10^{-5}} \approx 55.6.$$

So at  $10^4$  shots, statistics dominate readout error by a factor of about 56. The crossover shot count at which the readout floor takes over is

$$N \gtrsim \frac{p(1-p)}{p_e^2} = \frac{0.25}{(9 \times 10^{-5})^2} = \frac{0.25}{8.1 \times 10^{-9}} \approx 3.09 \times 10^7$$

shots per setting.

**Derivation 4.3 (contraction convergence; projection).** Model a single-parameter calibration loop as the linear contraction  $T(x) = x^* + L(x - x^*)$  with labeled illustrative inputs  $L = 0.6$ , initial discrepancy  $d_0 = d_{\mathcal{D}}(x_0, x^*) = 1$  in normalized units, and target tolerance  $\varepsilon = 10^{-3}$ . Then

$$n \geq \frac{\ln(\varepsilon/d_0)}{\ln L} = \frac{\ln(10^{-3})}{\ln(0.6)}.$$

Evaluate the logarithms:

$$\begin{aligned}\ln(10^{-3}) &= -3 \ln 10 \approx -3 \times 2.302585 = -6.907755, \\ \ln(0.6) &= \ln 6 - \ln 10 \approx 1.791759 - 2.302585 = -0.510826.\end{aligned}$$

Divide:

$$n \geq \frac{-6.907755}{-0.510826} \approx 13.524,$$

so  $n = 14$  rounds suffice. Check by direct evaluation:  $0.6^2 = 0.36$ ,  $0.6^4 = 0.1296$ ,  $0.6^8 = 0.01679616$ , and

$$0.6^{14} = 0.6^8 \times 0.6^4 \times 0.6^2 = 0.01679616 \times 0.1296 \times 0.36 \approx 7.836 \times 10^{-4} < 10^{-3},$$

while  $0.6^{13} = 0.6^{14}/0.6 \approx 1.306 \times 10^{-3} > 10^{-3}$ , so  $n = 14$  is exactly the minimum. A variant projection with a smaller initial discrepancy  $d_0 = 5 \times 10^{-2}$  and the readout floor  $\varepsilon = p_e = 9 \times 10^{-5}$  as target gives  $0.6^n \leq 1.8 \times 10^{-3}$ , i.e.

$$n \geq \frac{\ln(1.8 \times 10^{-3})}{\ln 0.6} = \frac{-6.3208}{-0.5108} \approx 12.37,$$

so  $n = 13$  cycles reach the readout-limited floor under those assumptions.

#### 4.3 Shot budget for ontology discrimination

**Input:**  $F = 0.99991$  from [4], so  $p_e = 9 \times 10^{-5}$  (Derivation 4.1). Suppose two candidate ontologies differ in the expectation value of some readout-accessible observable by  $\delta$ . At confidence level  $\alpha = 0.01$ , Hoeffding's inequality requires

$$N \geq \frac{\ln(2/\alpha)}{2\delta^2} = \frac{\ln 200}{2\delta^2}.$$

Evaluate  $\ln 200 = \ln 2 + \ln 100 \approx 0.693147 + 4.605170 = 5.298317$ . For a discriminability target  $\delta = 10^{-3}$ :

$$N \geq \frac{5.298317}{2 \times 10^{-6}} = 2.649159 \times 10^6.$$

Compare with the empirical trial count in [4],  $N_{\text{rep}} = 10^6$ :

$$\frac{N}{N_{\text{rep}}} = \frac{2.649159 \times 10^6}{10^6} \approx 2.65.$$

For a harder target  $\delta = 10^{-4}$ :

$$N \geq \frac{5.298317}{2 \times 10^{-8}} = 2.649159 \times 10^8,$$

about 265 times the demonstrated trial count. The per-shot error floor also bounds resolvability: discrepancies  $\delta \lesssim p_e = 9 \times 10^{-5}$  are comparable to the readout noise per shot.

#### 4.4 A two-ontology toy model with two fixed points

We now exhibit the underdetermined regime explicitly. Let the operational datum be a single scalar  $x \in \mathbb{R}$  (a normalized discrepancy between predicted and observed readout statistics), with two candidate ontologies  $A$  and  $B$ :

$$T_A(x) = 0.5x + 0.5a, \quad T_B(x) = 0.5x + 0.5b,$$

with ontology-dependent offsets  $a \neq b$ , and a selection rule that keeps the branch whose prediction is closer. Solving  $x = T_A(x)$ :  $x = 0.5x + 0.5a \Rightarrow 0.5x = 0.5a \Rightarrow x_A^* = a$ ; similarly  $x_B^* = b$ . Each branch is contractive with  $L = 0.5 < 1$ , so both fixed points are locally attracting; the failure of uniqueness comes from the multiplicity of branch fixed points, not from divergence.

**Concrete numbers.** Take  $a = 0$ ,  $b = 1$ , and initial guess  $x_0 = 0.6$ . The basin boundary is the midpoint  $m = (a + b)/2 = 0.5$ ; since  $x_0 > m$ , the iteration follows branch  $B$ :

$$x_1 = 0.5 \times 0.6 + 0.5 \times 1 = 0.8,$$

$$x_2 = 0.5 \times 0.8 + 0.5 = 0.9,$$

$$x_3 = 0.5 \times 0.9 + 0.5 = 0.95,$$

with closed form  $x_n = 1 - 0.4 \times 0.5^n$  (check:  $n = 1$ :  $1 - 0.2 = 0.8$ ;  $n = 2$ :  $1 - 0.1 = 0.9$ ;  $n = 3$ :  $1 - 0.05 = 0.95$ ). The iteration converges to  $x_B^* = 1$  and never visits branch  $A$ . The two ontologies are physically inequivalent—say, an energy eigenmode with offset  $a = 0$  versus a collective variable with offset  $b = 1$ —yet each is exactly self-consistent on its own basin, and the final ontology is determined by the *initial guess*, not by the data.

#### 4.5 Symmetry-breaking sensitivity

**Derivation 4.5 (projection).** If the readout channel had a small non-equivariant component—an asymmetry  $\epsilon$  between the two readout axes—the smallest detectable asymmetry is bounded by the total estimation uncertainty:

$$\epsilon_{\min} \approx \sqrt{\sigma_p^2 + p_e^2}.$$

With  $N = 10^4$  (Derivation 4.2):

$$\epsilon_{\min} \approx \sqrt{(5.00 \times 10^{-3})^2 + (9 \times 10^{-5})^2} = \sqrt{2.5 \times 10^{-5} + 8.1 \times 10^{-9}} = \sqrt{2.5081 \times 10^{-5}} \approx 5.008 \times 10^{-3}.$$

At the crossover count  $N = 3.09 \times 10^7$ :  $\sigma_p = \sqrt{0.25/(3.09 \times 10^7)} \approx 9.0 \times 10^{-5}$ , so

$$\epsilon_{\min} \approx \sqrt{(9.0 \times 10^{-5})^2 + (9 \times 10^{-5})^2} = \sqrt{1.62 \times 10^{-8}} \approx 1.27 \times 10^{-4}.$$

Any genuine ontological signature smaller than  $\epsilon_{\min}$  is operationally invisible.

## 5. Results

All numbers below are computed in Section 4; assumptions are restated where the input was assumed rather than sourced.

**R1 (Convergence in the contraction regime; projection).** With Lipschitz constant  $L = 0.6$  (illustrative input) and initial discrepancy  $d_0 = 1$ , calibration reaches tolerance  $\epsilon = 10^{-3}$  in exactly  $n = 14$  rounds (Derivation 4.3; verified by  $0.6^{14} \approx 7.836 \times 10^{-4} < 10^{-3} < 0.6^{13} \approx 1.306 \times 10^{-3}$ ). A variant with  $d_0 = 5 \times 10^{-2}$  and target  $p_e = 9 \times 10^{-5}$  needs  $n = 13$  cycles. In this regime the operational circularity is benign: unique fixed point, initial-ontology independence.

**R2 (Multiplicity in the two-branch model).** In the toy model of Section 4.4 with offsets  $a = 0$ ,  $b = 1$ , there are two exact fixed points  $x_A^* = 0$  and  $x_B^* = 1$ , separated by the basin boundary  $m = 0.5$ ; iteration from  $x_0 = 0.6$  converges to  $x_B^* = 1$  via  $x_1 = 0.8$ ,  $x_2 = 0.9$ ,  $x_3 = 0.95$ , with closed form  $x_n = 1 - 0.4 \times 0.5^n$ . The final ontology is determined by the initial guess, not the data: ontological underdetermination is realizable in a locally contractive loop.

**R3 (Readout error floor).** From the fidelity  $F = 0.99991$  in [4]:  $p_e = 9 \times 10^{-5}$ , i.e. about one error per  $1.11 \times 10^4$  shots; over the reported  $10^6$  trials,  $N_{\text{err}} = 90 \pm 10$  wrong outcomes expected.

**R4 (Shot budget; projection).** Resolving an ontology-discriminating signal of size  $\delta = 10^{-3}$  at confidence  $\alpha = 0.01$  requires  $N = 2.649159 \times 10^6$  shots, approximately 2.65 times the  $10^6$ -trial count demonstrated in [4]. For  $\delta = 10^{-4}$ , the requirement rises to  $N = 2.649159 \times 10^8$  shots. These assume the fidelity of [4], readout noise only, and Hoeffding's bound.

**R5 (Statistical floor and symmetry sensitivity; projections).** At  $N = 10^4$  shots per setting,  $\sigma_p \approx 5.00 \times 10^{-3}$ , exceeding the readout floor by a factor  $\approx 55.6$ ; the crossover is  $N \approx 3.09 \times 10^7$  shots. The smallest detectable readout asymmetry is  $\epsilon_{\min} \approx 5.008 \times 10^{-3}$  at  $10^4$  shots and  $\epsilon_{\min} \approx 1.27 \times 10^{-4}$  at the crossover count.

**R6 (Structural dichotomy).** The equivariant symmetry of Section 3.3 leaves  $d_{\text{orbit}} = 3$  continuous descriptive parameters (two axis angles, one basis phase) unconstrained by any equivariant operational statistic; the family of data-compatible ontologies is uncountable. The cross-platform question reduces to ontology alignment in the sense of [3]: either a correspondence preserving physical interpretation exists across platform fixed points, or it does not. The formalism does not decide which obtains—that is an empirical question with the budget of R4.

**Headline claim.** The circularity of qubit ontology is benign exactly when the calibration operator  $T$  is a contraction (unique fixed point, convergence in a small number of rounds in the illustrative model), and problematic exactly to the extent that multiplicity or an equivariant symmetry survives.

## 6. Discussion

**Limitations.** The contraction analysis (R1) uses an illustrative Lipschitz constant; real calibration loops may have state-dependent or iteration-dependent  $L$ , and the Banach uniqueness guarantee applies only to the idealized operator  $T = \Gamma \circ \Phi$  with exact realization maps. Real  $\Gamma$  includes unmodeled drift, and drift can violate the contraction condition between rounds. The two-branch model (R2) is deliberately minimal: one scalar datum, two branches. Real ontology spaces are high-dimensional, and multiplicity there may take forms—continuous families of fixed points, limit cycles—not captured here. The shot-budget result (R4) assumes the discriminating observable is readout-accessible with the fidelity of [4] and that the only noise is readout noise; additional decoherence during the discriminating sequence would inflate  $N$  beyond the computed values. The dimension count behind R6 assumes perfect equivariance; any real platform has stray asymmetries (bias fields, crosstalk) that partially break  $SO(3)$ , and Derivation 4.5 quantifies the sensitivity only under the assumed shot counts and an idealized noise model—independent, stationary shots. Correlated readout errors would inflate  $\sigma_p$  and worsen  $\epsilon_{\min}$ .

**Failure modes.** The framework fails if calibration maps are chaotic rather than contractive—then no fixed point exists and "the qubit" has no self-consistent definition at all, a stronger conclusion than we claim. It also fails if the symmetry orbit is not a gauge: if two members of the orbit correspond to genuinely different physical couplings (e.g., single-ion transition versus collective motional mode) that happen to be related by the equivariance, then the "benign gauge" reading is wrong and the underdetermination is physical, not conventional. We cannot decide this from operational data by construction—that is the point—but an independent probe outside the loop (e.g., a measurement coupling to a third level, or to the environment) could.

**What would falsify the claims.** (i) A demonstration that a calibration loop on a real platform converges to a unique, platform-independent Hilbert-space identification including the axis-labeling degrees of freedom—i.e., that the orbit is in fact broken by operational data alone—would falsify the 3-dimensional-orbit claim (R6). (ii) A demonstration that two members of the orbit give different predictions for an observable accessible only outside the calibration loop (for example, a probe coupling to a third level or to the environment) would show that the equivariant symmetry is physical rather than conventional, falsifying the benign-gauge reading of R6. (iii) A demonstration that the shot budget of R4 is unattainable in practice—for instance, that correlated readout errors inflate the required  $N$  by orders of magnitude beyond the Hoeffding value—would undermine the empirical testability of the underdetermination claim rather than the claim itself, and would shift the burden to a different discriminating observable.

## 7. Conclusion

We asked whether the qubit is a pre-existing object that calibration merely probes, or an entity constituted by the operational loop that defines it. The answer is conditional. Formally, the operational circularity—the fact that states are defined by measurements whose calibration presupposes states—is captured by a fixed-point problem for the composite operator  $T = \Gamma \circ \Phi$ . When  $T$  is a contraction with Lipschitz constant  $L < 1$ , Banach's theorem guarantees a unique self-consistent identification, and the circularity is benign: in the illustrative model of Derivation 4.3,  $n = 14$  rounds suffice to reach tolerance  $\varepsilon = 10^{-3}$  from  $L = 0.6$  and  $d_0 = 1$ . When the loop admits multiple fixed points or a continuous equivariant symmetry, physically inequivalent ontologies remain compatible with identical operational data; the two-branch model of Section 4.4 realizes this with two exact fixed points  $x_A^* = 0$  and  $x_B^* = 1$ , and the final ontology is fixed by the initial guess rather than by the data.

The empirical anchor is the single-shot readout fidelity  $F = 99.991\%$  for the optical qubit in  $^{40}\text{Ca}^+$  reported in [4], giving a per-shot error floor  $p_e = 9 \times 10^{-5}$ . From it we computed the budget for any ontology-discriminating test:  $N = 2.649159 \times 10^6$  shots to resolve  $\delta = 10^{-3}$  at  $\alpha = 0.01$ , about 2.65 times the demonstrated  $10^6$ -trial count, rising to  $N = 2.649159 \times 10^8$  shots for  $\delta = 10^{-4}$ . The headline conclusion is that uniqueness is a property of the calibration dynamics, not of the qubit: whether operational circularity is benign or masks ontological underdetermination is decidable only by the contraction structure of  $T$  and by experiments whose budget we have quantified.

## References

[1] Primer on the Gene Ontology. arXiv:1602.01876v1. <https://arxiv.org/abs/1602.01876v1> [2] Gene Ontology: Pitfalls, Biases, Remedies. arXiv:1602.01875v1. <https://arxiv.org/abs/1602.01875v1> [3] GraphMatcher: A Graph Representation Learning Approach for Ontology Matching. arXiv:2404.14450v1. <https://arxiv.org/abs/2404.14450v1> [4] High-fidelity readout of trapped-ion qubits. arXiv:0802.1684v2. <https://arxiv.org/abs/0802.1684v2> [5] How do Ontology Mappings Change in the Life Sciences?. arXiv:1204.2731v1. <https://arxiv.org/abs/1204.2731v1> [6] Towards Ontology-Based Descriptions of Conversations with Qualitatively-Defined Concepts. arXiv:2509.04926v1. <https://arxiv.org/abs/2509.04926v1> [7] A Community-Developed Domain Ontology for Magnetic Materials. arXiv:2609.11464v1. <https://arxiv.org/abs/2609.11464v1> [8] A Framework for Business Intelligence Application using Ontological Classification. arXiv:1109.1088v1. <https://arxiv.org/abs/1109.1088v1> [9] DOI 10.5281/zenodo.21451776. QNFO: No Thing There: Control, Readout, and Self-Referential Metrology in Engineered Quantum Systems. [10] DOI 10.5281/zenodo.22753022. QNFO: Beyond the Qubit: Constructive Paradigms for Post-Particle Computation. [11] DOI 10.5281/zenodo.21299278. QNFO: Manifesto for Honest Computation. [12] DOI 10.5281/zenodo.21254143. QNFO: The Qubit Delusion: How Particle Ontology Sabotaged Quantum Computing.

## Appendix A. Divergence report

No unresolved divergences survived reconciliation of the independent drafts. Two convention choices are documented for transparency. First, the shot-budget bound is stated throughout using Hoeffding's inequality,  $N \geq \ln(2/\alpha)/(2\delta^2)$ ; a draft alternative using the binomial standard error would give a smaller constant but the same scaling in  $\delta^{-2}$ , and the Hoeffding convention was retained as the conservative choice. Second, the Lipschitz constant  $L = 0.6$  and initial discrepancy  $d_0 = 1$  in Derivation 4.3 are labeled illustrative inputs (a projection), not measurements; all drafts agreed on this labeling, and the main text marks every result depending on them as a projection.

## Appendix B. Claim attribution

| Claim | Substance                                                                                                                               | Source drafts | Status     |
|-------|-----------------------------------------------------------------------------------------------------------------------------------------|---------------|------------|
| C1    | Ontology map $\Phi : \mathcal{D} \rightarrow \mathcal{O}$ and fixed-point loop $T = \Gamma \circ \Phi$                                  | A, B, C       | CONVERGENT |
| C2    | Contraction condition ( $L < 1$ ) implies unique fixed point and benign circularity (Banach)                                            | A, B, C       | CONVERGENT |
| C3    | $p_e = 9 \times 10^{-5}$ from $F = 99.991\%$ in [4]; $N_{\text{err}} = 90 \pm 10$ over $10^6$ trials                                    | A, B, C       | CONVERGENT |
| C4    | Shot budget $N = 2.649159 \times 10^6$ for $\delta = 10^{-3}$ , $\alpha = 0.01$ (Hoeffding)                                             | A, B, C       | CONVERGENT |
| C5    | $N = 2.649159 \times 10^8$ for $\delta = 10^{-4}$ ; ratio $\approx 265$ to demonstrated trials                                          | A, B          | SINGLE     |
| C6    | Two-branch toy model with fixed points $x_A^* = 0$ , $x_B^* = 1$ ; outcome set by initial guess                                         | A, B, C       | CONVERGENT |
| C7    | Convergence count $n = 14$ for $L = 0.6$ , $d_0 = 1$ , $\varepsilon = 10^{-3}$ ; $n = 13$ for the readout-floor variant                 | A, B          | CONVERGENT |
| C8    | Statistical floor $\sigma_p \approx 5.00 \times 10^{-3}$ at $N = 10^4$ ; crossover $N \approx 3.09 \times 10^7$                         | A, B, C       | CONVERGENT |
| C9    | Symmetry sensitivity $\epsilon_{\text{min}} \approx 5.008 \times 10^{-3}$ ( $10^4$ shots) and $\approx 1.27 \times 10^{-4}$ (crossover) | B, C          | CONVERGENT |
| C10   | $SO(3)$ equivariance leaves $d_{\text{orbit}} = 3$ continuous parameters unconstrained                                                  | A, C          | CONVERGENT |
| C11   | Cross-platform question framed as ontology alignment in the sense of [3]                                                                | A, B          | CONVERGENT |
| C12   | Falsifiability conditions (i)–(iii) in Section 6                                                                                        | A, B, C       | CONVERGENT |