

Beyond the Qubit, Part II: A Constructive Research Programme for Post-Particle Computation, with Quantified Instrument Budgets, a Fano-Scale Admissibility Gate, and Embodied-Verification Pathways

Beyond the Qubit, Part II: A Constructive Research Programme for Post-Particle Computation, with Quantified Instrument Budgets, a Fano-Scale Admissibility Gate, and Embodied-Verification Pathways

Abstract

Research programmes that frame computation as a circuit of qubits face a growing epistemic critique: the companion analysis argues that the qubit-gate-circuit model projects a particle ontology onto systems better described relationally and field-theoretically [9]. The specific gap is that critiques of the qubit model stop at diagnosis: they name no measurement standard to inherit, no error budget a successor must satisfy, and no embodied test protocol by which a successor formalism could be falsified. This paper supplies the constructive half of the answer. We strip the computational description to its minimal relational content and show that for $n = 50$ constituent units a pairwise relational parametrization carries 1225 parameters against $2^{50} \approx 1.1259 \times 10^{15}$ Hilbert-space amplitudes, a compression factor of 9.19×10^{11} . We convert the NEXT-100 design target of 0.5–1% FWHM energy resolution at the double-beta-decay energy of ^{136}Xe [5] into a one-sigma admissibility band $\sigma_E/Q_{\beta\beta} \in [2.1233 \times 10^{-3}, 4.2466 \times 10^{-3}]$, derive a root-sum-square error-composition rule (five independent stages at 10^{-2} each compose to 2.2361×10^{-2} statistically versus 5×10^{-2} worst-case), and derive a tracking-drift bound of 6.0×10^{-2} m per minute under stated illustrative assumptions. We then map immersive-interface capabilities [1], [2], [3], [4], [6], [7], [8] onto a verification protocol in which a human operator is the readout instrument, and size a falsification pilot at 63 participants per group for 80% power against a medium effect. The work is analytical and unreviewed (TRL 1–2); its practical effect is to convert a philosophical dispute into a numerical gate that any laboratory can apply.

1. Introduction

The people who face the problem this paper bears on are working physicists, instrument builders, immersive-interface laboratory operators, and research programme funders who must decide where to spend the next decade of scarce experimental effort. Quantum-computing roadmaps assume the qubit-gate-circuit model is the correct target; the companion analysis argues instead that this model is an epistemic failure — a projection of particle ontology onto relational, field-theoretic reality — and asks what comes next [9]. If that diagnosis is right, the cost of continuing on the qubit path is not merely engineering effort on a contested abstraction but the absence of any prepared alternative: no accepted measurement standard to inherit, no error budget a successor must meet, and no test protocol that could falsify a candidate formalism. Today a researcher who accepts the diagnosis has nowhere to go.

The so-what is concrete. A funder holding two competing post-qubit proposals has no shared yardstick and must judge them by narrative. A laboratory that wants to test a relational or field-theoretic formalism has no specification of how precise its apparatus must be. A sceptic who wants to kill the whole post-particle programme has no single number to attack. If the results of this paper hold, the go/no-go decision on candidate formalisms becomes a

numerical gate; an engineer can build the measurement instrument on hardware that already ships; and a theorist can state, in advance and in public, what observation would falsify the programme — the discipline that the Manifesto for Honest Computation demands of its own portfolio through its Falsification Pledge [12].

The strategy is deliberately conservative. We do not propose a new physics theory. We take the minimal constructive stance announced in the companion paper — strip quantum mechanics to its minimal core and rebuild [9] — and we ask what engineering scaffolding any such rebuild must satisfy to count as science rather than as manifesto. The scaffolding has three parts: (i) a precision standard inherited from the NEXT programme, whose Phase I design anticipates an energy resolution nearing the Fano factor in xenon, 0.5–1% FWHM at $Q_{\beta\beta,^{136}\text{Xe}}$, with a time-projection-chamber design that tracks events and identifies their double-end-blob topology [5]; (ii) an error-composition rule, derived in Section 4, that converts per-stage errors into an end-to-end pass/fail figure; and (iii) a human-in-the-loop verification layer grounded in the immersive-interface literature, in which skeleton-fitted self-avatars [1], instant motion tracking [6], calibrated depth perception [8], real-time physiological monitoring [7], projection-mapped room re-specification [3], shared multi-user analytic spaces [2], and spatially identical remote views [4] together make a human operator into a calibrated, auditable measurement channel.

Two objections must be met at the outset. First, that immersive interfaces are an application domain, not a computational substrate; our reply is that the programme's claim is epistemic — that relational observables are the correct primitive level of description — and an epistemic claim requires an instrument that measures relational observables, which is what these interfaces are. Second, that the programme is unfalsifiable philosophy; our reply is the falsification conditions stated explicitly in Sections 5 and 6 and the pilot design sized by the power calculation of Section 4.

The contributions of this paper are:

- A minimal relational parametrization of an n -unit computational description, with an explicit count showing $n(n-1)/2$ pairwise parameters against 2^n Hilbert-space amplitudes, yielding a compression factor of 9.19×10^{11} at $n = 50$, with a memory projection of 18.014 PB versus 1.96×10^4 bytes under a stated 16-byte-per-amplitude assumption (derived in Section 4.1, Result R1).
- A derived one-sigma equivalent of the NEXT Fano-scale resolution standard, $\sigma_E/Q_{\beta\beta} \in [2.1233 \times 10^{-3}, 4.2466 \times 10^{-3}]$, computed from the stated 0.5–1% FWHM figure by the Gaussian FWHM-to-sigma conversion (Section 4.2, Result R2).
- A derived error-composition rule with a worked example: for n independent stages each at relative error ε , the worst-case stack is $n\varepsilon$ and the statistical stack is $\sqrt{n}\varepsilon$; for $n = 5$, $\varepsilon = 10^{-2}$, these are 5×10^{-2} and 2.2361×10^{-2} respectively (Section 4.3, Result R3).
- A derived drift-growth bound for anchored tracking in an embodied verification rig: random-walk accumulation of per-frame anchor error gives $\delta_{\text{drift}} = \varepsilon_f \sqrt{N} = 6.0 \times 10^{-2}$ m over a one-minute session under stated illustrative assumptions, against 3.6 m under a constant per-frame bias (Section 4.4, Result R4).
- Derived latency, persistence, and tracking budgets for the embodied instrument: a single 90 Hz frame consumes 55.6% of an assumed 20 ms motion-to-photon budget; one-frame persistence blurs a 360 deg/s head motion by 4.0° ; inter-sample angular travel is 6.0° at 60 Hz versus 0.36° at 1 kHz, a factor of 16.7 (Section 4.5, Results R5–R7).
- A sized falsification experiment: 63 participants per group for 80% power against a medium effect at $\alpha = 0.05$ two-sided (Section 4.6, Result R8).
- A mapped verification protocol and engineering pathway, in which each embodied capability is assigned a role in the rig, with a TRL-honest readiness statement and the smallest next test (Sections 3 and 7).

2. Background and Related Work

This section reviews the supplied bibliography in its exact numbering. Eight entries are research papers with substantive summaries; four are programme documents, two of which carry only titles. We discuss each and state plainly where the supplied summary gives no further detail.

Precision measurement as the inherited standard. The NEXT experiment entry describes Phase I of the NEXT-100 neutrinoless double-beta-decay experiment, anticipating an energy resolution nearing the Fano factor in xenon of 0.5–1% FWHM at $Q_{\beta\beta,^{136}\text{Xe}}$, with a time-projection-chamber design allowing tracking and identification of the double-end-blob feature of the decay signature [5]. Two features of this entry matter for our argument. First, the resolution figure is a design *anticipation*, not a measured result; we treat it throughout as a target standard, not as achieved performance. Second, the entry shows what a mature experimental paradigm looks like from the outside: a quoted resolution and a topology identifier that turns an energy measurement into a topology-verified one. Section 4 converts the quoted FWHM figure into the one-sigma scale we adopt as our admissibility gate.

Embodiment and avatar calibration. The fitted-avatars paper addresses automatic skeleton adjustment for self-avatars in virtual reality, motivated by the role self-avatars play in enhancing presence and providing embodiment, and in collaborative virtual reality in improving communication through gestures; it notes that whether one uses complex motion capture or a few trackers with inverse kinematics, a good skeleton fit is essential [1]. The supplied summary is truncated mid-sentence and gives no quantitative results; we therefore use this work only for its architectural claim — that embodiment quality depends on a calibration step that can be automated — which is exactly the role we assign to avatar fitting in the verification rig of Section 3.

Multi-user analytics as a model of shared instrument state. The interactive multi-user 3D visual analytics project developed a prototype AR data-analytics application with an interactive 3D interface, hand-gesture-based controls, and multi-user support for a shared experience, built to explore the advantages and disadvantages of AR for visual data analytics [2]. The relevance to our programme is structural: a shared, gesture-controlled volumetric interface is a working precedent for several operators inspecting the same computational state simultaneously, which is the collaboration mode our verification protocol requires when a formalism's predictions must be checked by more than one observer.

Redirected locomotion as relational perturbation. The Reality Distortion Room is a proof-of-concept system using projection mapping and unencumbered interaction with the Microsoft RoomAlive system to study users' locomotive responses to visual effects that seemingly transform the physical room; it presents five effects that augment the appearance of a physical room to subtly encourage user motion [3]. The supplied summary is truncated before its results, so we cite no outcome numbers from it. What we take from it is the existence proof that a physical room can be perceptually re-specified by controlled visual offsets while the physical room is held fixed — the programme's clearest precedent for a relational perturbation experiment — and the concrete count of five effect families, which we use as the stage count $n = 5$ in the worked error-budget example of Section 4.3, an explicitly labelled modelling assumption rather than a finding of [3].

Remote guidance as the expert-operator channel. The mixed-reality remote-assistance system is a volumetric communication system for procedural tasks in which a remote expert visually guides a local operator; the two parties share a view that is spatially identical — the real object for the local operator and a mixed-reality "hologram" for the remote expert [4]. This is the communication topology our rig inherits: a formalism's predicted state can be rendered as the hologram layer while the apparatus's actual state is the local layer, and disagreement between expert view and operator view is the observable that constitutes a test.

Tracking as the anchoring primitive. The instant motion tracking paper identifies proper anchoring of virtual content to the real world — tracking — as a key challenge in enabling AR features, and presents a motion-tracking system that, per the summary, leverages advances in computer vision and mobile computing to support AR across platforms and adoption in major products [6]. The summary is truncated before method details, so we take from it only the problem statement: anchoring error is the fundamental per-frame quantity in any AR system. Section 4.4 builds a drift model on exactly this quantity, with the per-frame error value stated as an assumption of this paper, not of [6].

Real-time physiological monitoring as the operator readout. The simulator-sickness paper states that simulator sickness remains a significant problem for wider VR adoption, that the standard Simulator Sickness Questionnaire is a subjective measurement inadequate for real-time applications such as VR games, and that the work therefore targets real-time detection of simulator sickness from players' psychophysiological data during gameplay [7]. The summary is truncated before its detection results. Its role in our programme is methodological: it establishes that operator state can be read continuously from physiological signals rather than by post-hoc questionnaires, which is what makes a human operator usable as an in-loop instrument rather than as an end-of-session report.

Depth perception as the calibration target. The depth-perception study investigates how object shape, fidelity, color, and luminance affect depth perception in handheld mobile AR, noting that many AR applications require depth matching of real and virtual objects and that color and luminance are depth cues traditionally studied in two-dimensional objects, with little prior research on these factors in the mobile AR setting described [8]. The summary is truncated before its findings. We use it as the citation for the claim that depth matching between real and virtual objects is an active, factor-decomposable calibration problem — the same problem our rig faces when a rendered hologram layer must be co-registered with the physical apparatus.

The programme documents. The companion paper "Beyond the Qubit: Constructive Paradigms for Post-Particle Computation" is described as the companion to "The Qubit Delusion" (v1.1), posing the question of what comes next if the qubit-gate-circuit model is an epistemic failure — a projection of particle ontology onto relational, field-theoretic reality — and announcing the stripping of quantum mechanics to its minimal core [9]. This paper is the constructive continuation of that question; the supplied summary states the framing without technical detail, so we use it as the premise this paper constructs from, not as a source of derived results. The "Ultrametric Physics Research Plan" [10] and "Simulated Stability and the Superdeterministic Quantum Machine" [11] entries carry titles only; the supplied summaries give no further detail, so we relate them to our argument only as named strands of the same programme and draw no findings from them. Finally, the "Manifesto for Honest Computation" entry states that it contains a Phase VI (Final Synthesis) with 5 principles, a concrete portfolio, and the Falsification Pledge [12]. We adopt the Falsification Pledge as a normative commitment: every acceptance criterion in this paper is stated so that a named observation would refute it. We do not reproduce the five principles or the portfolio, since the supplied summary does not state their content.

What the supplied literature does not contain is any prior work connecting relational computation formalisms to embodied-interface measurement, or any arithmetic connecting per-stage errors to a single admissibility number. That connection is the gap this paper fills, and because no supplied entry states it, every claim of the connection made here is derived within this paper rather than attributed.

3. Methods

The method is constructive derivation plus protocol mapping, with no new data collection. It has four components.

Component 1: minimal relational parametrization. Following the stripping move of [9], we posit that the minimal descriptive content of an n -unit computational system is the set of pairwise relations among units, not the exponential amplitude vector of the Hilbert-space description. We define

$$R(n) = \binom{n}{2} = \frac{n(n-1)}{2}$$

as the number of independent pairwise parameters, against

$$H(n) = 2^n$$

for the Hilbert-space description of n qubits. The comparison ratio is $\rho(n) = H(n)/R(n)$. This is a parametrization-count comparison, not a claim that the two descriptions are empirically equivalent; the equivalence question is exactly what the programme's experiments are designed to test.

Component 2: standard extraction and unit conversion. We take the NEXT design target — energy resolution nearing the Fano factor in xenon, 0.5–1% FWHM at $Q_{\beta\beta, 136\text{Xe}}$ [5] — and convert it from FWHM to one-sigma form under the assumption, stated here and carried through the paper, that the resolution figure behaves as a Gaussian width for the purpose of budget composition. The conversion constant is $2\sqrt{2\ln 2}$. We emphasize that [5] states this figure as an anticipation for Phase I, not a measurement; our gate inherits it as a target, and Section 6 treats this inheritance as a load-bearing assumption.

Component 3: error-composition and instrument-budget arithmetic. We model any candidate computational pipeline — from formalism to rendered prediction to physical readout — as n independent stages, each contributing a relative error ε_i . Two composition rules bound the end-to-end error: the worst-case linear stack $\sum_i \varepsilon_i$, appropriate when stage errors are correlated or adversarially aligned, and the root-sum-square (RSS) stack $\sqrt{\sum_i \varepsilon_i^2}$, appropriate when stage errors are independent and zero-mean. For the embodied instrument we model a display-plus-tracking loop with three budgeted quantities: frame period $T_f = 1/f$ at frame rate f ; motion-to-photon latency budget B ; and display persistence p expressed as a fraction of T_f . For a head or hand rotating at angular velocity ω , persistence blur is

$$\theta_{\text{blur}} = \omega p T_f,$$

and inter-sample angular travel at tracker update rate f_t is

$$\theta_{\text{sample}} = \frac{\omega}{f_t}.$$

All numeric inputs to the instrument model are stated as explicit assumptions in Section 4, because the supplied literature establishes that tracking is the key anchoring challenge [6] and that sickness is a real-time detection problem [7], but does not supply numeric latency figures.

Component 4: experiment sizing and protocol mapping. For a two-group behavioral comparison (relational-perturbation interface versus matched baseline) we use the standard group-size formula for a two-sided test of a difference in means with equal groups:

$$n_{\text{group}} = \frac{2(z_{1-\alpha/2} + z_{1-\beta})^2}{d^2},$$

where d is the standardized effect size (Cohen's d), α the two-sided significance level, and $1 - \beta$ the target power. The normal quantiles $z_{0.975} = 1.95996$ and $z_{0.80} = 0.84161$ are standard statistical constants used as stated inputs. The verification rig itself is defined as follows: a candidate formalism's predictions are rendered as a mixed-

reality hologram layer over the physical apparatus, following the shared-view topology of the remote-assistance system [4]. The rig's calibration channels are: automatic skeleton fitting of the operator's self-avatar [1]; instant motion tracking for anchoring the hologram to the apparatus [6]; depth matching between hologram and apparatus, decomposed by the shape/fidelity/color/luminance factors of the mobile-AR depth study [8]; continuous physiological monitoring of the operator in place of retrospective questionnaires [7]; and, where the apparatus occupies a room-scale volume, projection-mapped room re-specification of the kind demonstrated by the Reality Distortion Room [3]. Multi-operator review uses the shared gesture-controlled 3D interface pattern of the multi-user analytics prototype [2]. The portfolio is organized under the five principles, concrete portfolio, and Falsification Pledge of [12]: each line item must name, in advance, the observation that would falsify it.

No empirical measurement is reported in this paper. Every number in Sections 4 and 5 is either (a) quoted from a bibliography entry and cited, (b) derived from such a quote by shown arithmetic, or (c) explicitly labelled a projection under stated assumptions.

4. Analysis

Every input number below is either a standard mathematical or statistical constant, an explicit assumption of the model stated in place, or a figure quoted from a supplied bibliography entry, and each is labelled as such. All arithmetic is shown.

4.1 Relational compression (Result R1)

Take $n = 50$ constituent units, chosen as a round representative scale for a mid-sized register; the choice of n is an assumption and the scaling is given for general n by the formulas of Section 3.

$$H(50) = 2^{50} = 1,125,899,906,842,624 \approx 1.1259 \times 10^{15}.$$

$$R(50) = \frac{50 \times 49}{2} = \frac{2450}{2} = 1225.$$

$$\rho(50) = \frac{2^{50}}{1225} = \frac{1.1258999068 \times 10^{15}}{1225} = 9.19102 \times 10^{11}.$$

Check: $1225 \times 9.19102 \times 10^{11} = 1.125900 \times 10^{15}$, consistent to the displayed digits. We adopt $\rho(50) = 9.19 \times 10^{11}$.

Memory consequence, as a projection with stated assumptions: if one complex amplitude is stored in 16 bytes (two 8-byte double-precision floats — an assumption about representation), the Hilbert description of $n = 50$ requires

$$16 \times 2^{50} = 1.80144 \times 10^{16} \text{ bytes} = 18.014 \text{ PB},$$

using $1 \text{ PB} = 10^{15}$ bytes, while the pairwise description requires

$$16 \times 1225 = 19,600 \text{ bytes} = 1.96 \times 10^4 \text{ bytes}.$$

The ratio is the same $\rho(50) = 9.19 \times 10^{11}$.

4.2 The Fano-gate band from the NEXT resolution target (Result R2)

Input (cited). The NEXT Phase I design anticipates an energy resolution nearing the Fano factor in xenon of 0.5–1% FWHM at $Q_{\beta\beta,^{136}\text{Xe}}$ [5].

Step 1: express the FWHM band in relative form. The stated band is

$$\frac{\text{FWHM}_E}{Q_{\beta\beta}} \in [0.005, 0.01].$$

Step 2: convert FWHM to one sigma. For a Gaussian, $\text{FWHM} = 2\sqrt{2 \ln 2} \sigma$. The constant is

$$2\sqrt{2 \ln 2} = 2\sqrt{2 \times 0.693147} = 2\sqrt{1.386294} = 2 \times 1.177410 = 2.354820.$$

Therefore $\sigma_E = \text{FWHM}_E / 2.354820$.

Step 3: compute the relative one-sigma band. At the 1% FWHM end:

$$\frac{\sigma_E}{Q_{\beta\beta}} = \frac{0.01}{2.354820} = 4.2466 \times 10^{-3}.$$

Arithmetic: $0.01 / 2.354820 = 4.24661 \times 10^{-3}$, rounded to 4.2466×10^{-3} . At the 0.5% FWHM end:

$$\frac{\sigma_E}{Q_{\beta\beta}} = \frac{0.005}{2.354820} = 2.1233 \times 10^{-3}.$$

Arithmetic: $0.005 / 2.354820 = 2.12323 \times 10^{-3}$, rounded to 2.1233×10^{-3} .

Result 4.2. The Fano-gate band is

$$\frac{\sigma_E}{Q_{\beta\beta}} \in [2.1233 \times 10^{-3}, 4.2466 \times 10^{-3}]$$

under the stated Gaussian-width assumption. Expressed symbolically, the resolved energy window is $\Delta E = (0.005 \text{ to } 0.01) \times Q_{\beta\beta}$; the entry does not state a numerical value of $Q_{\beta\beta}$, so no absolute energy is computed here — the fractional statement is the full content imported from [5].

4.3 Error-composition rule and the worked five-stage example (Result R3)

Inputs (assumptions of this paper, stated). $n = 5$ independent stages; each stage contributes relative error $\varepsilon = 10^{-2}$. The stage count is motivated by the five effect families demonstrated in the Reality Distortion Room study [3]; the per-stage value 10^{-2} is chosen to sit at the coarse end of the Fano-gate band derived in Section 4.2, so that the example is diagnostic rather than flattering.

Step 1: worst-case linear stack. If stage errors align in sign,

$$\varepsilon_{\text{WC}} = \sum_{i=1}^n \varepsilon_i = n\varepsilon = 5 \times 10^{-2}.$$

Step 2: statistical RSS stack. If stage errors are independent and zero-mean,

$$\varepsilon_{\text{RSS}} = \sqrt{\sum_{i=1}^n \varepsilon_i^2} = \sqrt{n \varepsilon^2} = \varepsilon \sqrt{n} = 10^{-2} \sqrt{5}.$$

Computing: $\sqrt{5} = 2.236068$, so $\varepsilon_{\text{RSS}} = 2.2361 \times 10^{-2}$.

Step 3: the ratio and its meaning.

$$\frac{\varepsilon_{\text{WC}}}{\varepsilon_{\text{RSS}}} = \frac{5 \times 10^{-2}}{2.2361 \times 10^{-2}} = \sqrt{5} = 2.2361.$$

Result 4.3. For the stated assumptions, the same five one-percent stages compose to 5×10^{-2} worst-case or 2.2361×10^{-2} statistically — a factor of $\sqrt{5} \approx 2.2361$ apart. Whether a candidate paradigm is buildable can hinge entirely on whether its stage errors are independent: a pipeline admissible under the RSS rule fails the identical per-stage numbers under the linear rule. Any admissibility claim must therefore state, and ideally test, the independence assumption.

4.4 Tracking drift-growth bound (Result R4)

Inputs (assumptions of this paper, stated). The instant-tracking literature identifies anchoring of virtual content to the real world as the key enabling challenge for AR features [6]; it supplies no per-frame error figure in the summary provided. We therefore assume, for illustration only: per-frame anchor error $\varepsilon_f = 10^{-3}$ m (one millimetre, a round figure at the scale of handheld-AR depth-matching studies [8]); session length $T = 60$ s; frame rate $f = 60$ Hz, giving $N = fT = 3600$ frames. We assume frame-to-frame anchor errors are independent and zero-mean, so drift grows as a random walk.

Step 1: frame count.

$$N = fT = 60 \times 60 = 3600.$$

Step 2: random-walk drift. For N independent zero-mean steps each of magnitude ε_f , the root-mean-square accumulated displacement is

$$\delta_{\text{drift}} = \varepsilon_f \sqrt{N} = 10^{-3} \times \sqrt{3600} = 10^{-3} \times 60 = 6.0 \times 10^{-2} \text{ m}.$$

Step 3: worst-case contrast. If per-frame errors align (e.g., a constant bias in the anchor estimate), the linear stack is

$$\delta_{\text{lin}} = \varepsilon_f N = 10^{-3} \times 3600 = 3.6 \text{ m},$$

which is sixty times the random-walk figure ($3.6/0.06 = 60$).

Result 4.4. Under the stated assumptions, one minute of anchored tracking accumulates 6.0×10^{-2} m of random-walk drift but 3.6 m of drift if a constant bias is present. Bias elimination, not variance reduction, is the binding requirement on the tracking channel of the verification rig: a bias of even 10^{-3} m per frame destroys room-scale co-registration within a minute, while the same variance is tolerable.

4.5 Instrument latency, persistence, and tracking budgets (Results R5–R7)

D2 (Frame period and latency budget, Result R5). Assumption: frame rate $f = 90$ Hz, a common head-mounted display rate (stated as a model assumption; the supplied entries do not fix a rate). Then

$$T_f = \frac{1}{90} \text{ s} = 0.011111 \text{ s} = 11.111 \text{ ms}.$$

Assumption: motion-to-photon budget $B = 20$ ms, stated as a model assumption for interactive comfort. The fraction of the budget consumed by one frame period alone is

$$\frac{T_f}{B} = \frac{11.111}{20} = 0.55556 = 55.6\%.$$

Result 4.5 (R5). A single 90 Hz frame period alone consumes 55.6% of a 20 ms motion-to-photon budget, leaving $20 - 11.111 = 8.889$ ms for sensing, tracking, and rendering. Any candidate formalism whose verification protocol requires more than one frame of pipeline depth at 90 Hz must therefore either relax the frame

rate or accept a latency budget that is already more than half spent before computation begins.

D3 (Persistence blur, Result R6). Assumptions: display persistence of one frame, $p = 1$; head rotation at $\omega = 360 \text{ deg/s}$, a round figure for a brisk head turn (stated as a model assumption; the supplied entries do not fix a rotation rate). Then

$$\theta_{\text{blur}} = \omega p T_f = 360 \text{ deg/s} \times 1 \times 0.011111 \text{ s} = 4.0^\circ.$$

Result 4.6 (R6). One-frame persistence at 90 Hz blurs a 360 deg/s head motion by 4.0° . At the Fano-gate scale derived in Section 4.2, a 4.0° angular smear is the entire visual budget of a co-registration task: if the hologram layer must be read against apparatus features to a precision comparable to the gate, persistence must be reduced below one frame or the rotation rate of the reading motion must be restricted.

D4 (Inter-sample angular travel, Result R7). Assumptions: the same $\omega = 360 \text{ deg/s}$; tracker update rates $f_t = 60 \text{ Hz}$ (a baseline consumer rate) and $f_t = 1000 \text{ Hz}$ (a high-rate target). Then

$$\theta_{\text{sample}}(60) = \frac{360}{60} = 6.0^\circ, \quad \theta_{\text{sample}}(1000) = \frac{360}{1000} = 0.36^\circ.$$

The ratio is

$$\frac{\theta_{\text{sample}}(60)}{\theta_{\text{sample}}(1000)} = \frac{6.0}{0.36} = 16.7.$$

Result 4.7 (R7). Between consecutive tracker samples, a 360 deg/s head motion sweeps 6.0° at a 60 Hz update rate but only 0.36° at 1 kHz — a factor of 16.7. The tracking channel of the verification rig must therefore operate at or near the kilohertz class if the operator's reading of the hologram layer is to be attributed to a known head pose rather than to inter-sample pose uncertainty.

D5 (Experiment sizing, Result R8). Inputs (standard statistical constants and stated assumptions): two-sided $\alpha = 0.05$, so $z_{1-\alpha/2} = z_{0.975} = 1.95996$; target power $1 - \beta = 0.80$, so $z_{1-\beta} = z_{0.80} = 0.84161$; standardized effect size $d = 0.5$ (a medium effect, stated as a design assumption). Then

$$z_{0.975} + z_{0.80} = 1.95996 + 0.84161 = 2.80157,$$

$$(z_{0.975} + z_{0.80})^2 = 2.80157^2 = 7.8488,$$

$$n_{\text{group}} = \frac{2 \times 7.8488}{0.5^2} = \frac{15.6976}{0.25} = 62.79.$$

Rounding up to the next integer participant, $n_{\text{group}} = 63$.

Result 4.8 (R8). A two-group falsification pilot comparing a relational-perturbation interface against a matched baseline requires 63 participants per group (126 total) for 80% power against a medium effect ($d = 0.5$) at $\alpha = 0.05$ two-sided. If the true effect is small ($d = 0.2$), the same formula gives $n_{\text{group}} = 2 \times 7.8488/0.04 = 392.44$, i.e. 393 per group — the pilot's feasibility is therefore highly sensitive to the effect size, and Section 6 treats $d = 0.5$ as a load-bearing assumption.

5. Results

Every number in this section is either carried from a derivation in Section 4 with its inputs and arithmetic shown there, or labelled a projection with stated assumptions. No empirical measurement is reported.

R1 (Relational compression, Section 4.1). For $n = 50$ constituent units, the pairwise relational parametrization carries $R(50) = 1225$ parameters against $H(50) = 2^{50} = 1.1259 \times 10^{15}$ amplitudes, a compression factor $\rho(50) = 9.19 \times 10^{11}$. Under the stated 16-byte-per-amplitude assumption, the Hilbert description requires 18.014 PB and the pairwise description 1.96×10^4 bytes. This is a parametrization-count result, not a claim of empirical equivalence between the two descriptions.

R2 (Fano-gate band, Section 4.2). The NEXT Phase I design target of 0.5–1% FWHM at $Q_{\beta\beta, 136}\text{Xe}$ [5], converted under the stated Gaussian-width assumption, yields the one-sigma admissibility band

$$\frac{\sigma_E}{Q_{\beta\beta}} \in [2.1233 \times 10^{-3}, 4.2466 \times 10^{-3}] .$$

No absolute energy is computed, because the supplied entry does not state a numerical value of $Q_{\beta\beta}$.

R3 (Error-composition rule, Section 4.3). For $n = 5$ independent stages each at relative error $\varepsilon = 10^{-2}$, the worst-case stack is 5×10^{-2} and the statistical RSS stack is 2.2361×10^{-2} , a factor of $\sqrt{5} \approx 2.2361$ apart. Admissibility of a candidate pipeline can flip depending solely on whether stage errors are independent.

R4 (Tracking drift bound, Section 4.4). Under the stated illustrative assumptions ($\varepsilon_f = 10^{-3}$ m per frame, $f = 60$ Hz, $T = 60$ s, independent zero-mean frame errors), random-walk drift is $\delta_{\text{drift}} = 6.0 \times 10^{-2}$ m per minute, against 3.6 m under a constant per-frame bias — a factor of 60. Bias elimination, not variance reduction, is the binding requirement on the rig's tracking channel.

R5 (Latency budget, Section 4.5). At $f = 90$ Hz, one frame period $T_f = 11.111$ ms consumes 55.6% of an assumed 20 ms motion-to-photon budget, leaving 8.889 ms for all other pipeline stages.

R6 (Persistence blur, Section 4.5). One-frame persistence at 90 Hz blurs a 360 deg/s head motion by 4.0° .

R7 (Inter-sample travel, Section 4.5). At $\omega = 360$ deg/s, inter-sample angular travel is 6.0° at a 60 Hz tracker and 0.36° at 1 kHz, a factor of 16.7.

R8 (Pilot sizing, Section 4.5). The falsification pilot requires 63 participants per group (126 total) for 80% power against $d = 0.5$ at $\alpha = 0.05$ two-sided; against $d = 0.2$ the requirement rises to 393 per group.

6. Discussion

The results convert a philosophical dispute into a set of numbers, but each number rests on assumptions that a critic can attack directly. This section states the limitations, the failure modes, and the falsification conditions, and argues against the programme's own claims.

The compression result is the weakest link. R1 compares parameter counts of two descriptions of the same system, but a smaller parametrization is not automatically a correct or complete one. The $n(n-1)/2$ pairwise count assumes that pairwise relations are the minimal descriptive content — the stripping move of [9] — and the supplied summary of [9] states this framing without technical detail, so the assumption cannot be checked against the companion paper's text as supplied here. A Hilbert-space description with 2^n amplitudes is over-complete for many purposes (most amplitudes are never individually resolved), so the honest statement is that the pairwise parametrization is *sufficient* if the relational thesis holds, and the compression factor measures the stakes of that thesis, not its truth. What would falsify the relational thesis at this level: an experimental demonstration that a computational-relevant observable of an n -unit system depends on structure that no pairwise parametrization can express, with the dependence surviving all reductions of the description.

The gate inherits an anticipation, not a measurement. R2 is derived from a design *anticipation* in the NEXT Phase I entry [5], not from achieved performance; the entry describes a system scheduled for data taking, and the supplied summary states no measured resolution. If the anticipated 0.5–1% FWHM is not reached, the gate band shifts upward and every downstream budget loosens proportionally. The Gaussian-width assumption is also a modelling choice: the entry does not state that the resolution figure is Gaussian, and non-Gaussian resolution functions would change the FWHM-to-sigma conversion.

The error-composition example is diagnostic, not predictive. R3's inputs ($n = 5$, $\varepsilon = 10^{-2}$) are assumptions of this paper; the stage count is motivated by the five effect families of [3] as an explicitly labelled modelling analogy, and [3]'s supplied summary is truncated before its results, so no outcome of that study supports the number. The rule itself (linear versus RSS stacking) is standard and exact under its stated independence condition; what is missing is any measurement of stage errors in a real candidate pipeline, which no supplied entry provides.

The instrument budgets are illustrative. R4–R7 rest on round-number assumptions ($\varepsilon_f = 10^{-3}$ m, $f = 60$ Hz, $f = 90$ Hz, $B = 20$ ms, $\omega = 360$ deg/s) chosen to be representative, not measured; the supplied entries establish that tracking is the key anchoring challenge [6] and that sickness detection is a real-time problem [7], but supply no numeric latency or error figures in the summaries provided. The structural conclusions (bias dominates variance for drift; frame period dominates the latency budget; tracker rate sets inter-sample uncertainty) follow from the algebra and are robust to the specific round numbers, but the absolute magnitudes scale linearly with the assumed inputs.

The pilot sizing is sensitive to the effect size. R8 assumes $d = 0.5$. If the true effect of a relational perturbation on operator performance is small ($d = 0.2$), the required sample rises from 63 to 393 per group, a factor of 6.2 — derived as $0.5^2/0.2^2 = 6.25$ in the group-size formula. A pilot designed at $d = 0.5$ that observes a smaller effect would be underpowered, and a null result would then be ambiguous between "no effect" and "smaller effect than assumed."

Falsification of the main result. The main constructive result of this paper is the claim that a numerical admissibility gate can be stated at all — specifically, that the Fano-gate band of R2 plus the composition rule of R3 together form a pass/fail criterion a laboratory can apply. This claim is falsified if either (i) the NEXT design target is withdrawn or revised such that no fractional resolution figure exists to convert, in which case the gate has no anchor; or (ii) a candidate formalism is shown to make predictions whose precision cannot be expressed as a per-stage relative-error budget, in which case the composition rule does not apply and the gate is inapplicable rather than merely stringent. Observation (ii) is the more likely attack surface, and the paper's response is that the burden then shifts to the formalism's proponents to state what error model does apply.

Limitations list.

- The relational parametrization's sufficiency is assumed from [9]'s framing, which the supplied summary states without technical detail; the compression factor is sensitive to this assumption in the sense that the entire R1 result is conditional on it.
- The Fano-gate band inherits an unmeasured design anticipation from [5]; the band endpoints scale linearly with the assumed FWHM figure.
- The FWHM-to-sigma conversion assumes Gaussian resolution behaviour, which [5] does not state.
- The error-composition example's $n = 5$ and $\varepsilon = 10^{-2}$ are modelling assumptions; the RSS rule's validity depends on stage-error independence, which is untested here.
- The drift bound's $\varepsilon_f = 10^{-3}$ m, frame rate, and session length are illustrative; the drift magnitude scales as $\varepsilon_f \sqrt{fT}$.

- The latency, persistence, and inter-sample budgets use assumed rates (90 Hz, 60 Hz, 1 kHz, 20 ms, 360 deg/s) not fixed by any supplied entry.
- The pilot size assumes $d = 0.5$; the required n scales as d^{-2} .
- The bibliography contains only twelve entries, four of which are programme documents, two carrying titles only [10], [11]; the literature base for the embodied-interface side is eight research papers with truncated summaries, and no supplied entry connects relational formalisms to embodied measurement — that connection is derived here, not cited.

Open questions. Whether pairwise relations suffice for computational observables; whether stage errors in a candidate pipeline are in fact independent; what error model applies if the RSS rule does not; whether operator physiological state [7] can be calibrated tightly enough to serve as an in-loop instrument rather than a coarse indicator; and whether the Fano-scale target is achieved by NEXT Phase I as anticipated.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The users are, first, programme funders choosing between post-qubit research proposals: for them, the Fano-gate band of R2 and the composition rule of R3 convert a narrative judgement into a numerical one — a proposal is admissible only if it states its per-stage error model and shows the composed error falls within the gate. Second, laboratory instrument builders: the budgets of R4–R7 specify the tracking channel's requirements (bias below 10^{-3} m per frame, tracker rate near 1 kHz, persistence below one frame) on hardware whose component capabilities the immersive-interface literature already demonstrates in kind — anchoring systems [6], skeleton-fitted self-avatars [1], shared multi-user analytic spaces [2], spatially identical remote views [4], depth-matching factor studies [8], and room-scale projection re-specification [3]. Third, experimental groups in the precision-measurement community, who gain a reusable statement of what their resolution standard means in one-sigma budget terms. If no user can yet be named for the relational formalism itself — because no candidate formalism has been built — the work is foundational in the direct sense: it builds the acceptance infrastructure that any future candidate must pass through, which is a prerequisite the field currently lacks.

(b) Current readiness. This paper is analytical work that has not been peer reviewed: against the NASA Technology Readiness Level definitions it sits at TRL 1–2 (basic principles observed and formulated; concept and application formulated). The component technologies cited — tracking [6], avatars [1], shared analytics [2], remote assistance [4] — may individually sit higher in their own fields, but the verification rig that composes them for formalism testing is a paper concept only. No claim of manufacturing readiness is made or supportable; analytical work of this kind does not exceed TRL 3 even after review.

(c) The smallest next test. The smallest de-risking experiment is a single-operator bench pilot: render a known geometric prediction as a hologram layer over a physical fixture using an existing anchoring system [6], have one operator judge co-registration under controlled head motions, and measure (i) accumulated anchor drift over a one-minute session against the R4 prediction of 6.0×10^{-2} m under zero-bias conditions, and (ii) the presence of any constant per-frame bias, which R4 predicts would produce metre-scale drift. The outcome that refutes the drift model is a measured drift substantially exceeding 6.0×10^{-2} m in a session verified to be bias-free, or a measured bias floor above 10^{-3} m per frame that no calibration removes — either would show the rig's tracking channel cannot hold room-scale co-registration and the embodied-verification pathway fails at its first gate.

(d) Binding constraints. In the order they bind: (i) tracking bias, per R4 — a constant per-frame bias of 10^{-3} m destroys co-registration in one minute, so bias calibration, not variance, is the first engineering problem; (ii) tracker update rate, per R7 — inter-sample uncertainty of 6.0° at 60 Hz is incompatible with precision reading, so the

kilohertz tracking class is required; (iii) latency budget, per R5 — with 55.6% of a 20 ms budget consumed by the frame period alone, the rendering pipeline has under 9 ms, which constrains the complexity of the rendered formalism prediction; (iv) persistence, per R6 — one-frame persistence blurs fast reading motions by 4.0°; (v) participant throughput, per R8 — a 126-participant pilot is a recruitment and ethics-approval constraint before it is a scientific one. Whether operator physiological monitoring [7] can be calibrated to in-loop standard, and whether depth-matching factors [8] transfer from handheld AR to the rig's display class, are open questions the supplied summaries do not resolve.

(e) Order-of-magnitude bounds (projections). The only cost-relevant quantity derivable from this paper's own arithmetic is the pilot's participant count: 126 participants at R8's sizing, or 786 at $d = 0.2$ — labelled projections under the stated effect-size assumptions, with no per-participant cost stated here because no supplied entry provides one. The memory projection of R1 (18.014 PB versus 1.96×10^4 bytes at $n = 50$, under the 16-byte assumption) bounds the storage difference between descriptions; no compute, energy, or hardware cost is derivable from the supplied material and none is asserted.

(f) Risks and misuse paths. The gate as constructed can be misused rhetorically: a proposal could quote the RSS composition rule while its stage errors are in fact correlated, passing a gate it would fail under the linear rule — which is why R3 concludes that any admissibility claim must state and test the independence assumption. Conversely, a sceptic could apply the worst-case rule to an adversary's proposal and the RSS rule to a favoured one; the mitigation is that the choice of rule must be declared in advance under the Falsification Pledge discipline of [12]. The compression figure of R1 could be quoted as evidence that the relational programme is already validated; it is a parametrization-count comparison under an assumed thesis, and quoting it otherwise misrepresents the result.

8. Conclusion

This paper supplies the constructive half of the post-qubit question posed by the companion analysis [9]: not whether the qubit-gate-circuit model fails, but what any successor must satisfy to count as testable science. The answer is a three-part scaffolding, each part stated as a number. A precision standard inherited from the NEXT programme's anticipated Fano-scale resolution [5], converted to a one-sigma admissibility band $\sigma_E/Q_{\beta\beta} \in [2.1233 \times 10^{-3}, 4.2466 \times 10^{-3}]$. An error-composition rule showing that five one-percent stages compose to 5×10^{-2} worst-case or 2.2361×10^{-2} statistically, so that admissibility can hinge on the independence of stage errors alone. And an embodied-verification layer, grounded in the immersive-interface literature [1]–[8], with derived budgets — 6.0×10^{-2} m per minute of bias-free tracking drift against 3.6 m under bias, 55.6% of a 20 ms latency budget consumed by one 90 Hz frame, 4.0° of persistence blur, a factor-16.7 inter-sample advantage at kilohertz tracking — and a falsification pilot sized at 63 participants per group. The parametrization-count comparison at $n = 50$ (1225 pairwise parameters against 2^{50} amplitudes, a factor of 9.19×10^{11}) states the stakes of the relational thesis without claiming its truth. Every acceptance criterion is stated so that a named observation would refute it, in keeping with the Falsification Pledge of [12]. The work is analytical and unreviewed, TRL 1–2; its immediate product is not a device but a gate — a numerical standard that turns the question "what comes after the qubit?" into a question a laboratory can answer with an instrument.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.16.0-revise-retry). Language models were tools in that process; they are not authors. The named author is accountable for the content.

- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; statements about a cited work are restricted to its retrieved title and abstract.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 16 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond its abstract; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] Fitted avatars: automatic skeleton adjustment for self-avatars in virtual reality. arXiv:2307.09558v1. <https://arxiv.org/abs/2307.09558v1> [2] Interactive Multi-User 3D Visual Analytics in Augmented Reality. arXiv:2002.05305v1. <https://arxiv.org/abs/2002.05305v1> [3] Reality Distortion Room: A Study of User Locomotion Responses to Spatial Augmented Reality Effects. arXiv:2510.23840v1. <https://arxiv.org/abs/2510.23840v1> [4] Remote Assistance with Mixed Reality for Procedural Tasks. arXiv:2208.03261v1. <https://arxiv.org/abs/2208.03261v1> [5] The NEXT experiment. Towards phase I. arXiv:1310.7054v2. <https://arxiv.org/abs/1310.7054v2> [6] Instant Motion Tracking and Its Applications to Augmented Reality. arXiv:1907.06796v1. <https://arxiv.org/abs/1907.06796v1> [7] Real-Time Detection of Simulator Sickness in Virtual Reality Games Based on Players' Psychophysiological Data during Gameplay. arXiv:2010.06152v1. <https://arxiv.org/abs/2010.06152v1> [8] The Effects of Object Shape, Fidelity, Color, and Luminance on Depth Perception in Handheld Mobile Augmented Reality. arXiv:2008.05505v1. <https://arxiv.org/abs/2008.05505v1> [9] QNFO: Beyond the Qubit: Constructive Paradigms for Post-Particle Computation [10] QNFO: Ultrametric Physics Research Plan [11] QNFO: Simulated Stability and the Superdeterministic Quantum Machine [12] QNFO: Manifesto for Honest Computation