An Evaluation of the Ignorant Observer Framework

Control Theory, Quantum Measurement, Capacity–Backaction, and Contemplative Philosophy

Aernoud Dekker's multi-part theoretical and experimental architecture
Scientific editorial assessment for public communication
May 2026

Executive Summary

When a body of work attempts to connect quantum measurement, limits of classical control, the scaling of quantum computers, the structure of spacetime, and Advaita Vedanta, the first reaction of a rigorous scientist should be caution. The history of "quantum consciousness" speculation is crowded with claims that dissolve into metaphor, mysticism, or unfalsifiable assertion. This stack earns attention because it does something rarer: it places its most radical implications behind a concrete laboratory lever.

The *Ignorant Observer Framework* (IOF) is unusual because it does not ask the reader to accept metaphysics first. Its technical core, the *BLQC* paper, begins with a concrete physical question:

> *What happens to a quantum experiment when the physical system maintaining the measurement basis has finite information-processing capacity?*

That question is not mystical. It is a control-theoretic question. The framework treats the measurement basis as a real, maintained reference structure rather than an ideal mathematical input. The word "observer" does not mean a human mind imposing reality. It means the physical system that keeps the experimental coordinate frame, phase reference, pulse axis, or readout basis usable. If that reference structure drifts faster than the controller can track it, the observer does not merely lack information verbally or psychologically. The apparatus has a finite-rate physical blind spot.

The strongest feature of the project is its architecture. The speculative material is not used as evidence for the physics. Instead, the stack is organized around a falsifiable mechanism:

- a finite-rate basis-tracking claim, expressed through C_{eff}, h_{KS}, and the deficit rate $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$;
- an experimental protocol designed to discriminate BLQC capacity/instability scaling from Penrose-style mass-geometry scaling under shared mesoscopic conditions;
- a standalone quantum-engineering paper, *The Capacity–Backaction Frontier*, that formalizes useful syndrome capacity versus induced physical instability as a standard-QEC rate coordinate;
- a foundational extension, *The Creation of Duality*, that asks how subject, object, space, time, and gravity-like structure could arise for a finite observer;
- a forensic data-mining program that openly reports negative and non-causal findings;
- a philosophical and cosmological layer that is explicitly framed as interpretation, structural analogy, and speculation.

This architecture is the project's main strength. BLQC is the lab manual: narrow, operational, and falsifiable. *The Capacity–Backaction Frontier* is a standalone quantum-engineering paper: it translates finite useful capacity into a standard QEC rate-accounting coordinate without requiring acceptance of the broader IOF interpretation. *The Creation of Duality* is the bridge from finite tracking to subject/object structure. *The Ignorant Observer* supplies the philosophical vocabulary. The cosmological notes are the sandbox. Because the lab lever comes first, the speculative material is not asked to carry the physics. It becomes a conditional blast radius.

This makes the work scientifically interesting in a very specific sense: it can lose. If the Penrose-overlap experiment finds no capacity/instability dependence under adequate sensitivity and controlled confounds, the technical bridge fails in its most important regime. If it does, the consequences are large.

Part I: The Core Physics – Measurement as a Finite-Rate Control Problem

In textbook quantum mechanics, the measurement basis is normally treated as an ideal parameter: the phase of a laser, the angle of a polarizer, or the reference frame of an interferometer is assumed to be available without cost. BLQC asks what changes if that basis is maintained by real hardware with finite update rate, latency, drift, noise, and feedback bandwidth.

This is the clarifying move. The technical paper strips away the metaphysics and asks an almost mundane question: what happens to a quantum experiment if the computer maintaining the lasers, pulses, local oscillator, or reference phase cannot process the relevant information fast enough? The hidden assumption being challenged is not quantum theory itself, but the idealization that the basis parameter θ is simply available. In the proposed mechanism, the experimenter may analyze the data as if the apparatus measured at θ, while the physical basis has wandered to $\theta + \delta\theta$.

The BLQC mechanism is compact:

1. The physical reference structure that defines the measurement basis has internal instability or drift, summarized by an effective entropy or chaos rate h_{KS}.
2. The controller that tracks and corrects that drift has a finite useful capacity C_{eff}.
3. If $C_{\mathrm{eff}} \ln 2 < h_{\mathrm{KS}}$, the controller cannot keep up. The deficit rate
$$\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2$$
is positive.
4. In that regime, unresolved basis uncertainty grows exponentially, and the predicted visibility follows the distinctive double-exponential form
$$V(t) = \exp\left[-\frac{1}{2}\sigma_0^2 e^{2\kappa t}\right].$$

This is not presented as a replacement for quantum mechanics. It is a proposed observer-side visibility-loss channel: a way in which finite reference tracking can mimic collapse-like or decoherence-like behavior without changing Born-rule marginals or invoking a conscious mind as a physical force. If confirmed, the result would not be that "measurement is magic." It would be the opposite: the measurement basis is a physical, thermodynamic object, and its maintenance has to be counted as part of the experiment.

The Penrose connection is important but must be read precisely. Penrose Objective Reduction (OR) predicts a collapse timescale from gravitational self-energy. BLQC predicts a visibility-loss timescale from information deficit in basis tracking. These are different mechanisms with different knobs. The striking claim is that, for illustrative mesoscopic regimes, the two timescale windows can overlap. That overlap does not prove BLQC, and it does not disprove Penrose. It creates a dangerous experimental ambiguity: a collapse-like loss of visibility in the Penrose window might be gravitational, controller-limited, or a mixture of both.

This is the paper's most powerful experimental insight. The absolute time is not the discriminator; the derivative is. Penrose OR should track mass geometry. BLQC should track C_{eff} and h_{KS}. The practical challenge to the lab is to use a mesoscopic apparatus in which both sets of variables are exposed: mass, separation, and geometry on one side; useful reference-tracking capacity and reference instability on the other. If the loss timescale moves with the latter while the former and ordinary confounds are controlled, the apparatus was not merely a passive witness. Its finite information capacity was part of the measurement physics.

Part II: Quantum Engineering – The Capacity–Backaction Frontier

The Capacity–Backaction Frontier should be read as a standalone quantum-engineering paper. Its subject is standard fault-tolerant QEC: syndrome information must be extracted fast enough to protect a logical state, while the measurements, readout photons, amplification, decoding, and feedback activity that provide that information can themselves add heat, leakage, quasiparticles, photon shot-noise dephasing, crosstalk, latency, and correlated errors.

The paper does not propose a new backaction mechanism. Its contribution is an aggregate rate-accounting coordinate for a tradeoff that QEC teams already manage component by component. It defines the useful online correction capacity C_{eff} as the mutual-information rate by which syndrome data are measured, decoded, and incorporated into a timely Pauli-frame or feedback update. It defines the effective instability rate as

$$h_{\text{eff}}(N, C_{\text{eff}}) = h_0(N) + h_{\text{corr}}(N) + g_N(C_{\text{eff}}),$$

where h_0 is the baseline instability, h_{corr} captures correlated error generation, and g_N is the backaction induced by measurement and control activity. The diagnostic ratio is

$$\rho_{\text{CB}} = \frac{\epsilon_{\text{QEC}} C_{\text{eff}} \ln 2}{h_{\text{eff}}(N, C_{\text{eff}})}.$$

The ratio is not meant to replace component-level error budgets. It sits above them. Component budgets diagnose photon shot noise, resonator ringdown, leakage, quasiparticles, amplifier saturation, crosstalk, TLS drift, and decoder backlog. The aggregate coordinate asks whether operating points across those mechanisms can be compared in one rate language: useful syndrome information gained per unit physical instability induced.

That is the paper's scientific value. It gives QEC engineering a falsifiable reporting and prediction proposal: logical performance should be analyzed not only against raw readout power, raw bandwidth, assignment fidelity, or cycle time, but also against ρ_{CB} or the protection margin

$$M = \epsilon_{\text{QEC}} C_{\text{eff}} \ln 2 - h_{\text{eff}}$$

under matched physical-load controls. If those aggregate variables do not improve prediction beyond conventional error-budget variables, the coordinate is not useful in that regime. If they

do, the paper has earned its place as a compact engineering language for capacity, latency, backaction, and correlated instability.

The relationship to BLQC is therefore indirect and limited. Peer review of this coordinate would not support BLQC's visibility-loss prediction. It would support a more modest premise that BLQC also uses: finite useful information capacity can be treated as a measurable physical control resource in quantum systems. That would be valuable on its own merit, while leaving the BLQC experiment to stand or fall independently.

Part III: The Experimental Rigor – How the Framework Can Lose

The strongest scientific document in the stack is the experimental protocol. It is written in a preregistration spirit: define the knobs, define the confounds, define the statistical model, and define what counts against the theory. It reads less like manifesto and more like an instruction sheet for making the framework vulnerable.

The protocol focuses on the Penrose-overlap regime: a mesoscopic visibility-loss experiment in which Penrose OR and BLQC can predict comparable timescales while assigning those timescales to different variables. This is the right place to test the framework because the competing explanations touch the same observable but disagree about which physical knob controls it.

In that regime, Penrose OR predicts that the loss timescale should follow gravitational self-energy, hence mass, separation, and mass distribution. BLQC predicts that the same kind of timescale should move with the finite-rate basis-reference deficit

$$\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2.$$

The decisive experiment therefore asks for a derivative map on the same apparatus: vary mass geometry while holding reference tracking fixed, and vary C_{eff} or h_{KS} while holding mass geometry, temperature, readout, pulse behavior, latency, and plant dynamics fixed.

That is a high-value feature. A vague theory says "things get worse." A testable theory says *which* knob moves the timescale, which knobs do not, and which functional form should organize the data. BLQC expects unresolved basis variance to grow into a double-exponential visibility-loss geometry; Penrose OR expects a geometry-governed timescale. A mixed outcome is also allowed, but then it must be modeled as overlapping rates rather than claimed as a simple victory.

A successful Penrose-overlap test would prove something narrower and stronger than the philosophical language sometimes suggests. It would not prove Vedanta, cosmic non-duality, or that Penrose OR is globally wrong. It would show that finite-rate basis-reference tracking is an experimentally active visibility channel in the mesoscopic regime where gravitational objective reduction is also a serious candidate. A clean null result in that regime would directly weaken the technical bridge that gives the larger stack its scientific force.

The *Forensic Signatures* paper is also scientifically healthy because it does not overclaim. The public data search found:

- little or no BLQC-positive signal in the available Google and Chinese superconducting-qubit recovery datasets;
- Gompertz-compatible recovery signatures in LIGO glitch data;

- insufficient auxiliary controller data to attribute the LIGO signatures causally to finite-rate basis tracking.

This is exactly the right conclusion. Retrospective curve fitting cannot establish the mechanism. Gompertz-like curves occur widely in nature. A visible shape is not enough; the decisive question is whether changing effective controller capacity moves the system into or out of that regime under controlled conditions.

This is also an important credibility marker. A weaker project would hide the null Google and Chinese results, or convert the LIGO curve match into a premature victory. This project does the opposite: it publishes the negative and non-causal findings, then uses them to close the lazy route of retrospective proof. The existing data are suggestive enough to motivate the controlled experiment, but not strong enough to replace it.

Part IV: Philosophy and Cosmology – A Quarantined Sandbox

The philosophical layer is where the IOF could have failed by overreach, but the stack mostly avoids that failure by separating inspiration from proof. This is the epistemological quarantine. Advaita Vedanta supplies language and structure: ignorance as veiling, projection, finite perspective, and misidentification. The technical framework does not claim to prove Vedanta. It asks whether the Advaitic analysis of ignorance has a formal analogue in finite information processing.

This distinction matters. The framework does not say that human thoughts collapse wavefunctions. It says that any finite observer, including an apparatus, must maintain a basis through limited physical channels. If those channels cannot track the relevant degrees of freedom, the observer sees an effectively classical, definite, or degraded reality because the missing information is physically unavailable to that observer.

That is a structural bridge, not a mystical shortcut. Avidya becomes a name for a formal limit: the inability of a finite observer to track the full causal state from which appearances arise. The important move is structural isomorphism, not mystical causation. Sanskrit terms are not being smuggled in as physical forces; they function as a disciplined vocabulary for patterns already expressed in the mathematics of finite tracking.

The Creation of Duality is the document where this bridge is developed most boldly. It is not a proof that reality is "made of mind," and it is not a replacement for physics. Its scientific value is different: it gives a disciplined mathematical account of why a finite observer may be forced to render reality in the familiar categories of subject, object, space, time, and gravity-like constraint.

The strongest part of that paper is its use of established machinery. The Data-Rate Theorem supplies the threshold: a finite tracker loses a chaotic source when $C_{\text{eff}} \ln 2 < h$. The Information Bottleneck supplies the subject/object split: under finite capacity, a stable controller must compress the world into action-relevant structure, which naturally separates an internal control channel from an external modeled environment. Information geometry then gives a serious interpretation of space as distinguishability rather than primitive extension, while time can be read operationally as successive updating.

That is why the paper is not disposable metaphysics. It belongs near QBism, Relational Quantum Mechanics, information geometry, Wheeler's "It from Bit" program, and other epistemic approaches that ask how much of "objective reality" is the invariant structure of finite obser-

vation. The manuscript stretches further than those programs, but its engine is recognizable mathematics rather than decorative mysticism.

The risky move is the Bridge Ansatz. *The Creation of Duality* maps the information-deficit rate κ to an energy scale using Margolus–Levitin saturation, yielding the conditional correspondence

$$E_G = \frac{\pi}{2}\hbar\kappa_{\text{info}},$$

and therefore a proportionality between Penrose-style collapse time and IOF tracking-loss time. This is elegant and provocative, but it is not derived from first principles. It is a phenomenological bridge from finite-rate observation to gravitational self-energy. The paper is honest about that. The ansatz is where the work becomes high-risk physics.

Its saving grace is the same as BLQC's: it is experimentally vulnerable. If coherence time changes with effective tracking capacity under fixed mass, geometry, temperature, readout, and plant dynamics, the bridge becomes interesting. If collapse-like behavior remains strictly mass-geometric and independent of C_{eff}, the bridge fails. That is a clean scientific posture.

The cosmological notes are even more speculative, and they are correctly labeled as such. Their value is not that they establish new cosmology. Their value is that they explore the possible blast radius if the finite-rate observer principle is experimentally confirmed. The document is explicit that it is not a paper, not a derivation, and not a replacement for standard cosmology. Its "Speculation Ladder" is a useful safeguard: metaphor should become dictionary, dictionary should become toy model, toy model should become functional form, and functional form should eventually risk data.

Within that sandbox, the notes are genuinely fertile. Black holes become candidates for tracking horizons; the Big Bang can be re-read as an observational resolution boundary; dark-energy-like language becomes a metaphor for information pressure; MOND-like acceleration scaling becomes a tempting but explicitly unproven dimensional echo. The document even flags its own numerical coincidences as cheap unless they produce curves. That self-skepticism is exactly what keeps the playground productive rather than grandiose.

The important point is discipline. The cosmological document does not replace general relativity, ΛCDM, particle physics, or observational cosmology. It asks whether the IOF dictionary might generate toy models or functional forms that someday risk disagreement with data. That is the right status for the material: provocative, useful, and not yet established.

Conclusion

The *Ignorant Observer Framework* is ambitious in a way that modern science often finds uncomfortable. It connects laboratory quantum measurement, control theory, quantum-computing architecture, Penrose-style collapse physics, non-dual philosophy, and speculative cosmology. That breadth creates risk. It also creates the project's unusual power: a single operational claim supports, limits, or collapses the larger structure.

The reason the work deserves attention is that the technical core is not a metaphor. It is a proposed finite-rate law:

$$C_{\text{eff}} \ln 2 \quad \text{versus} \quad h_{\text{KS}}.$$

If the observer can track the basis, coherence remains visible. If the observer cannot, unresolved basis variance grows and visibility can disappear in a sharply structured way.

Everything else in the stack depends on that claim. If the BLQC experiment fails, the philosophical and cosmological extensions remain interesting essays, but their scientific force collapses. If the BLQC experiment succeeds, the consequences are large: quantum measurement, Penrose-overlap tests, QEC scaling, and the physical meaning of observer ignorance all need to be re-examined.

That is the mark of a serious research program. It does not ask for belief. It proposes a lever, a curve, and a way to be wrong. If the lever does nothing, the blast radius disappears. If the lever moves the curve, the physics community has a new variable to account for: the finite information capacity of the physical observer that maintains the measurement basis.

This assessment is intentionally promotional in tone. Its claims are conditional where the underlying document stack is conditional: BLQC is treated as a falsifiable technical proposal, the capacity–backaction frontier as a standalone standard-QEC rate-accounting paper, and the philosophical/cosmological material as interpretation and speculation rather than established physics.