Where Did the Measurement Basis Come From?

Finite Basis-Tracking and the Measurement Problem

`https://ignorantobserver.xyz`

Aernoud Dekker

May 2026

Companion to BLQC v1.7

DOI: 10.17605/OSF.IO/XAK6R

Abstract

Why is the measurement basis treated as if it came from outside physics? Quantum measurement is usually discussed only after the basis has already been chosen. This paper starts one step earlier: the basis is not an externally supplied classical parameter but a physical reference variable $\theta(t)$, generated and tracked by a finite observer-apparatus system inside the same causal history as the measured system.

This reframes apparent collapse as a limit of embedded self-knowledge. The observer can record which basis and outcome occurred, but cannot reconstruct the full causal ancestry that produced them together. The proposed operational handle is finite basis tracking. If the basis-producing dynamics generate information at rate h_{KS} while the useful tracking channel has capacity C_{eff}, the relevant deficit is

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

When $\kappa < 0$, ordinary quantum statistics are recovered. When $\kappa > 0$, IOF predicts an additional visibility factor,

$$V_{\mathrm{obs}} = V_{\mathrm{std}} \exp\left(-\tfrac{1}{2}\sigma_\theta^2\right),$$

in the Gaussian unresolved-basis limit, distinct from ordinary environmental decoherence. The claim is falsifiable: if visibility timing does not track κ under controlled variation of C_{eff} and h_{KS}, IOF fails in that regime.

Scope. This paper is a conceptual bridge between the full Ignorant Observer Framework (IOF) [1] and the technical BLQC protocol. It does not attempt to derive quantum mechanics, Bell correlations, or a new micro-dynamics. Its claim is narrower: the external-basis idealization hides a finite physical tracking problem, and BLQC tests whether that problem leaves a capacity-dependent visibility signature.

The Move IOF Makes

Quantum mechanics predicts probabilities before measurement and definite records after measurement. A particle is prepared in a state that allows several possible outcomes. A measurement is performed. One outcome is recorded.

The measurement problem asks what, if anything, physically happens in this passage from possibility to fact.

The usual answers take different routes. Collapse theories add a physical interruption of the wavefunction. Many-worlds denies collapse and lets every outcome occur in a different branch. Hidden-variable theories keep one outcome, but add an unseen state that determines it.

IOF starts one step earlier.

It asks where the measurement question itself came from.

A spin measurement is not the question, "what is the spin?" in the abstract. It is the more specific question, "relative to this axis, what result is registered?" The axis, phase, threshold, timing, or reference frame is the measurement basis. In standard use, that basis is treated as an externally supplied classical parameter. Once it is given, quantum mechanics supplies the probabilities.

IOF removes that idealization. The basis is not an outside input. It is a physical state of the observer-apparatus system.

The experimenter's selected angle, the controller's register state, the magnet orientation, the optical phase reference, the detector threshold, and the timing gate are all physical variables. They have histories. They are produced by prior conditions. They are maintained by finite hardware. They are never causally floating above the experiment.

This is the central move:

> The measurement setting and the measured system are not assumed to be ancestrally independent. They are descendants of one physical history.

That does not mean the apparatus secretly sends a signal to the particle, or that the particle secretly sends a signal to the apparatus. It means the pair consisting of "this setting" and "this outcome" belongs to one consistent causal history. The setting is part of the world being measured, not an exception to it.

In this reading, the apparent mystery of collapse is displaced. The problem is no longer, "how does an indeterminate world become definite?" The problem is, "why does a definite causal history appear to an embedded observer as a probability distribution?"

IOF's answer is finite self-knowledge.

An observer can record what happened. The basis was 37 degrees. The outcome was spin down. But the observer cannot reconstruct the full causal ancestry by which that basis and that outcome arose together. The causal chain is too deep, too entangled with the observer's own state, and too expensive to track from within.

This is why the resolution is not merely a control problem.

The control problem is real: a finite apparatus must track and stabilize its own basis. If the basis variable is θ, the observer's internal representation is $\hat{\theta}$, and the tracking error is $\delta\theta$, then finite tracking produces a variance in $\delta\theta$. That variance can reduce measured visibility.

But the deeper point is not just that controllers are imperfect. It is that the controller, the basis, the particle state, and the outcome all sit inside one causal history. The ordinary assumption that the measurement setting is statistically independent of the measured system is therefore not a metaphysical guarantee. It is an operational idealization.

Bell-style reasoning usually assumes measurement independence: the hidden state of the measured system is statistically independent of the later measurement setting [4]. IOF questions that assumption at its root. If both the system state and the basis-setting process descend from the same past, then correlation between them need not be conspiratorial. In a single-history reading it is not an added coordination, but a structural possibility.

The reason this correlation does not become ordinary hidden predictability is that the ancestral path cannot be traced in principle by the embedded observer.

This is also where IOF must be stated carefully. The framework does not claim that ordinary observers could uncover hidden variables if only they had better instruments. Nor does this short argument derive the Born rule from scratch. It gives a physical account of why a single-history, no-collapse embedding can appear probabilistic to an observer inside that history. The probability rule still has to be hosted by the chosen ontology, or taken as the operational quantum rule that bounded observers use. IOF's added claim is about the origin of the observer's unavoidable ignorance: the basis-setting process is part of the same inaccessible ancestry as the outcome.

Why the Observer Cannot Track Its Own Basis

So far the move is purely structural: the basis is a physical variable, and it shares ancestry with the system. The next step is to show why that shared ancestry is inaccessible from the inside—why a finite observer cannot, even in principle for a given configuration, keep a faithful internal account of the basis it is using. Two distinct facts about the embedded observer combine to force this, and they must be kept separate.

The first is structural. The measurement context—basis, phase, timing, orientation—is generated by the physical state of the observer-apparatus itself. To track that context is therefore to track a thing whose dynamics include the dynamics of the tracker. The target includes the tracker. The recursion is not an infinite regress; it is a feedback problem: the observer updates a model of a context that includes the observer's own updating state. A complete self-account would require the observer to include its own record-forming activity inside the record. This is what defines the demand on self-tracking; it is not, by itself, a thermodynamic claim.

The second is physical. The observer is a finite physical system. Tracking is not contemplation; it is the encoding, updating, comparison, and correction of information in physical registers, each step costing power, memory, bandwidth, and time. The Landauer bound gives one principled floor—the irreducible cost of erasing a bit—but it is only a ceiling on what is physically possible, typically far above the modest rate any one tracking loop actually devotes to the basis. The operative limit is the useful rate the loop actually achieves, set by all of its finite physical resources.

Two rates set the comparison.

First, the basis-tracking loop has a finite *useful* information rate. A system devotes only so much effective capacity to constraining its own basis. Call that effective basis-tracking capacity C_{eff}. It is the rate that genuinely constrains θ—accepted updates, useful bits per update, surviving fraction after latency and filtering—not the device's raw power or the Landauer ceiling.

Second, the observer-apparatus system has internal dynamics that generate fresh unpredictability. If those dynamics are chaotic, the relevant rate is the Kolmogorov-Sinai entropy rate, h_{KS}. If they are diffusive, the relevant quantity is a diffusion rate. Either way, the basis-producing state is not static. It keeps generating information that must be tracked.

Self-reference says what must be tracked. Finite physics says how fast. The deficit between the required rate, h_{KS}, at which the basis-producing dynamics generate fresh relevant information about the measurement context, and the available rate, $C_{\mathrm{eff}} \ln 2$, is the single physical number that governs the framework:

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

The sign of κ is the entire content of the two regimes. Where $\kappa < 0$, capacity wins: the represented basis stays locked to the implemented basis, IOF predicts no new probability law, and ordinary Born-rule statistics are recovered to high accuracy. Where $\kappa > 0$, chaos wins: the basis-producing dynamics outrun the observer's self-tracking, lossless self-tracking necessarily falls behind, the represented basis drifts from the implemented one, and the framework predicts a specific attenuation of visibility,

$$V_{\mathrm{obs}} = V_{\mathrm{std}} \, \exp\bigl(-\tfrac{1}{2}\sigma_\theta^2\bigr) \, ,$$

in which σ_θ^2 is the variance of the residual basis-tracking error. This factor is the single observable signature of finite self-tracking: unity when tracking keeps up, falling below unity exactly when the deficit goes positive. Everything the BLQC experiment tests lives in this one factor.

1 First Objection: Does This Just Move the Mystery?

A physicist will object at this point.

If the measurement setting and the system state share ancestry, have we really explained anything? Or have we simply hidden the Born rule inside an unknown global history?

This objection is fair. IOF should not answer it by pretending to derive more than it has derived.

The framework does not claim, in this document, to derive the Born rule from ancestral correlation alone. Within IOF, a separate route—treated in its own companion paper—attempts a conditional derivation of the binary Born form under bandwidth-limited quantum control. The present document takes the Born rule as an empirical constraint and asks a different question.

IOF adds a different claim.

It says that the measurement basis is not an external free parameter. It is a physical variable with causal ancestry. Therefore the usual assumption that the setting is statistically independent of the system is not a metaphysical necessity. It is an operational idealization.

The role of IOF is not to use hidden correlations as an adjustable explanation for every quantum statistic. That would be the empty, conspiratorial form of superdeterminism: correlations tuned after the fact to fit any result.

It is also not a completed deterministic theory. IOF does not rely on 't Hooft-style structural determinism [6] as a load-bearing premise. It does not claim that a deterministic cellular or global state-space dynamics has already been specified and shown to reproduce quantum theory.

The relevant claim sits between those two—narrower than a completed theory, and more specific than a bare appeal to hidden correlations. In the technical vocabulary of Bell's theorem it is a violation of *measurement independence*: the setting and the system share causal ancestry, so $P(\xi \mid \theta) \neq P(\xi)$. That places IOF on the measurement-dependence route, which is sometimes filed under “superdeterminism.” The label is accurate in the technical sense and misleading in the colloquial one, so it must be qualified precisely:

> **Epistemically bounded ancestral correlation**: non-conspiratorial measurement dependence in which the shared ancestry is real but cannot be reconstructed by the embedded observer as a predictive ledger.

Two qualifications carry the weight. *Non-conspiratorial*: following Palmer [5], a violation of measurement independence need not be fine-tuned or manipulative. In a single globally consistent history the correlation is a structural feature of that history, not a coordination arranged between settings and hidden states. *Epistemically bounded*: the embedded observer cannot reconstruct the shared ancestry, so the correlation is not a knob available for prediction or curve fitting. It is a limit on what an observer inside the history can know about the history that produced both the question and the answer.

The role of IOF is then to identify a specific physical limitation inside the observer-apparatus system: finite basis self-tracking. Its visibility consequences were established above—ordinary quantum statistics where $\kappa < 0$, and the tracking-loss factor $\exp(-\frac{1}{2}\sigma_\theta^2)$ where $\kappa > 0$.

For entangled measurements this means the ideal quantum correlation is not replaced by an arbitrary hidden-variable curve. It is multiplied by that same tracking-loss factor:

$$E_{\text{measured}}(a, b) = E_{\text{QM}}(a, b)\ \exp\bigl(-\tfrac{1}{2}(\sigma_A^2 + \sigma_B^2)\bigr)\ .$$

So if quantum mechanics predicts

$$E_{\text{QM}}(a, b) = -\cos(a - b),$$

IOF predicts, in the Gaussian basis-error regime,

$$E_{\text{measured}}(a, b) = -\cos(a - b)\ \exp\bigl(-\tfrac{1}{2}(\sigma_A^2 + \sigma_B^2)\bigr)\ .$$

The cosine is not fitted by ancestral correlation. It is the quantum correlation, inherited from the hosting embedding and recovered in the capacity-wins limit. IOF’s distinct claim is the extra, capacity-dependent visibility factor.

This is what prevents the framework from becoming unfalsifiable.

If changing effective basis-tracking capacity does nothing, while ordinary decoherence variables are controlled, IOF fails in that regime.

If changing C_{eff} shifts the visibility break according to $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$, then finite self-tracking is doing physical work.

The objection therefore sharpens the framework rather than defeating it.

IOF is not saying: “anything can happen because everything is ancestrally correlated.”

It is saying: “standard quantum statistics are recovered when basis tracking is stable; measurable departures should appear only when the observer-apparatus can no longer track the physical basis it is using.”

Two Threads

There are therefore two threads in the argument.

Thread A is the ontological or ancestral thread. It says that the basis, the system, and the outcome are not metaphysically separate ingredients inserted into the experiment from outside. They are descendants of one physical history. This is the thread that reframes the conceptual measurement problem by rejecting the external-basis idealization. It is interpretation, and its physical force is borrowed entirely from Thread B.

Thread B is the operational or experimental thread. It says that the observer-apparatus has a finite useful capacity for tracking the physical basis-producing state, and that stressed basis tracking should produce a visibility timing law governed by

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

The two threads are related, but they do not have the same evidential status. Thread A is the interpretation. Thread B is the falsifiable handle. A physicist may test the capacity-scaling claim without accepting the whole single-history ontology. Conversely, the ontology gains scientific pressure only if the operational κ-scaling survives ordinary control, noise, and recovery explanations.

2 Second Objection: Is This Just Correctable Reference Noise?

A physicist will now ask a sharper experimental question.

If the measurement basis drifts, jitters, or is tracked with finite bandwidth, why is this more than ordinary classical reference error? If the actual basis is logged with a better instrument, should the lost visibility not come back after re-binning the data?

The answer is: yes, for the IOF channel it should—and that is the point, not an embarrassment.

IOF does not predict irreversible physical collapse. The finite-tracking channel is observer-relative by construction: it describes what a given online controller, with a given capacity, can resolve in real time. A loss of that kind is *expected* to be recoverable from a sufficiently high-resolution offline reference log, because the realized basis existed all along; it simply was not delivered to the online loop in time. Recoverability therefore confirms the observer-relative character of the channel rather than refuting it. In the protocol's terms, a recovery statistic $R_{\mathrm{rec}} \to 1$ on the IOF component is the IOF-positive signature, and $R_{\mathrm{rec}} \to 0$ (irreversible loss) is the signature of genuine decoherence or objective collapse [3].

What would deflate the claim is not recoverability. It is the absence of the right *dependence*. IOF must distinguish three cases.

First, there is ordinary environmental decoherence. Here visibility is lost because the system becomes physically entangled with uncontrolled environmental degrees of freedom. Better bookkeeping cannot restore it ($R_{\mathrm{rec}} \to 0$), and the loss tracks thermal and coupling variables, not κ.

Second, there is generic reference noise. Here the apparatus used one basis while the analysis assumed another, but the residual is arbitrary—it does not scale with any independently imposed information deficit. Visibility may be partly recoverable, but its timing is governed by whatever phase noise happens to be present, not by κ.

Third, there is the falsifiable claim of IOF. For a given online observer-controller, unresolved basis uncertainty should not be arbitrary. It should scale with the deficit

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

where C_{eff} is the useful information rate actually constraining the basis reference, and h_{KS} is the instability or entropy-production rate of that reference dynamics.

The prediction is not merely "visibility goes down when the reference is noisy."

The prediction is that, in the chaos-wins regime,

$$\sigma_\theta^2(t) = \sigma_0^2 \, e^{2\kappa t},$$

and therefore, in the Gaussian basis-error regime,

$$V(t) = \exp\left(-\tfrac{1}{2}\sigma_0^2 \, e^{2\kappa t}\right).$$

It is often clearer, especially for a reader arriving from decoherence theory, to read the observed visibility as a product of two multiplicative channels:

$$V_{\mathrm{obs}} \approx V_{\mathrm{std}} \, V_{\mathrm{IOF}}, \qquad V_{\mathrm{IOF}} = \exp\left(-\tfrac{1}{2}\sigma_\theta^2\right),$$

where V_{std} is the ordinary environmental and decoherence channel—the visibility standard quantum mechanics already predicts—and V_{IOF} is the finite basis-tracking channel. IOF does not deny V_{std}. It claims only that, in stressed regimes, part of the observed visibility loss may belong to V_{IOF}, and may be misassigned to standard decoherence if the capacity-instability coordinate κ is never varied and tested.

Distinguishing From a Lindblad Description

This visibility curve must also be stated carefully. By itself, a curve of this form may be representable as a time-dependent dephasing or master-equation noise model. In that sense, the formula alone is not automatically distinguishable from a suitably chosen Lindblad-style effective description [8].

The proposed discriminator is not the functional form alone. It is the causal accounting behind the parameters: whether the fitted loss rate collapses against an independently calibrated κ rather than against ordinary decoherence variables. The recovery axis runs orthogonal to this: it classifies the *character* of the loss (observer-relative versus irreversible), not whether IOF wins. (The full operational protocol—moving C_{eff} and h_{KS} against fixed confounds, and computing R_{rec} from a passive high-resolution shadow log—is set out in the companion protocol [3].)

If the whole effect is captured by an ordinary Lindblad or phase-noise model whose parameters track temperature, idle time, actuator distortion, pulse noise, or environmental coupling better than κ, IOF has no independent win. IOF gains force only if the capacity-instability coordinate predicts visibility timing after the standard master-equation and reference-noise explanations have been given every chance to win.

So the primary observable is not an absolute visibility level. It is the movement of the breakdown time:

$$t_{\mathrm{break}} \propto \frac{1}{h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2},$$

under controlled variation of C_{eff} and h_{KS}.

This is where the experiment becomes nontrivial.

If changing C_{eff} only changes temperature, readout signal-to-noise, pulse shape, actuator behavior, latency, or idle time, then the result is confounded.

If changing C_{eff} produces a visibility shift whose timing does not collapse against κ, then the result is only a generic reference-noise benchmark.

But if, in a mesoscopic visibility experiment, after thermal, readout, latency, pulse, actuator, and offline-recovery controls are included, the loss timescale still collapses against

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

better than against ordinary decoherence variables, then IOF has found a real observer-side visibility channel.

The point is not that an external computer's ignorance causes physical decoherence.

The point is that the measurement basis is a physical reference variable, and a finite online observer-controller may lose operational access to that reference in a way that has a precise, falsifiable scaling law.

The decisive test is therefore:

- increase C_{eff} at fixed mass geometry, temperature, readout signal-to-noise, latency, pulse behavior, and plant dynamics;
- verify that t_{break} moves later;
- increase h_{KS} at fixed C_{eff} and fixed ordinary confounds;
- verify that t_{break} moves earlier;
- check whether breakdown times collapse against κ;
- reject the effect if it is fully explained by ordinary decoherence, latency, actuator distortion, or heating—or if its timing shows no κ-dependence once those are controlled.

This objection is not a threat to the framework. It is the experimental filter the framework must pass.

IOF survives only if the capacity-instability coordinate predicts visibility timing after the usual reference-error explanation has been given every chance to win.

3 Why the Mesoscopic Regime Matters

At this point, a physicist may concede the experiment but still resist the interpretation.

They may say: "This is meaningful, but it is metrology. You have shown that a finite-bandwidth reference loop can affect visibility. You have not shown that this belongs to quantum foundations."

That is the right place to introduce the mesoscopic regime.

The reason IOF points there is not that ordinary reference tracking becomes magically foundational at larger mass. The reason is that, in the mesoscopic regime, another serious proposal already predicts collapse-like visibility loss on comparable timescales: Penrose objective reduction [7].

Penrose's proposal assigns the loss timescale to gravitational self-energy. In simplified form,

$$\tau_{\mathrm{OR}} \sim \frac{\hbar\, s}{G\, m^2},$$

where m is the mass in superposition and s is the branch separation. For suitable mesoscopic masses and separations, this can fall in the millisecond to hundred-millisecond range.

IOF assigns a possible observer-side loss timescale to finite-rate basis tracking:

$$\tau_{\mathrm{loss}} \sim \frac{1}{h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2},$$

up to threshold factors set by the chosen visibility criterion and initial basis uncertainty.

The numerical example here is illustrative, not a universal IOF prediction. In the low-bandwidth regime used in the technical protocol [2], values such as

$$h_{\mathrm{KS}} \approx 50 \text{ nats/s}, \qquad C_{\mathrm{eff}} \approx 10 \text{ bits/s}$$

give

$$C_{\mathrm{eff}} \ln 2 \approx 6.93 \text{ nats/s}, \quad \kappa \approx 43.1 \text{ s}^{-1}, \quad 1/\kappa \approx 23 \text{ ms}.$$

The often-quoted 50–70 millisecond figure includes the threshold/log factor from the chosen visibility criterion and the assumed initial basis uncertainty. Different thresholds or initial uncertainties move that number. The point is not the exact illustrative value; it is that plausible low-useful-capacity, high-instability reference dynamics can put the IOF tracking-loss window in the same broad 10–100 millisecond mesoscopic range often associated with Penrose-style objective reduction.

This numerical overlap does not prove IOF. It is not evidence by itself.

Its value is experimental. It creates a regime where two different explanations can predict similar absolute loss times, but depend on different physical knobs.

Penrose-style objective reduction predicts that the loss time should move when mass, separation, or mass distribution changes. It should not have a leading dependence on the useful information rate of the basis-tracking loop once ordinary confounds are controlled.

IOF predicts the reverse for its observer-side channel. At fixed mass geometry, temperature, readout signal-to-noise, latency, pulse behavior, actuator response, and plant dynamics, the loss time should move with C_{eff} and h_{KS}.

The discriminator is therefore not the absolute timescale. It is the derivative:

- change mass/separation at fixed C_{eff} and h_{KS};
- change C_{eff} or h_{KS} at fixed mass/separation.

If the effect follows mass geometry and ignores basis-tracking capacity, IOF's falsifiable claim fails in that regime, and a geometry-dependent mechanism remains the better explanation.

If the effect follows C_{eff} and h_{KS} at fixed mass geometry, after ordinary reference-error explanations have been controlled, then the result is no longer just a generic statement that "bad references reduce contrast." It says that the finite self-tracking of the measurement basis predicts the timing of visibility loss in the same regime where objective-collapse physics was expected to matter.

The two mechanisms are not framed as mutually exclusive. The combined-rate treatment—additive, mediated, and collinear (the speculative Bridge-Ansatz) outcomes—is developed in the protocol [3]; here it suffices that the regime can discriminate the leading dependence.

The interpretation is still earned only after the controls. But the experiment is no longer merely a calibration exercise. It becomes a direct comparison between two proposed coordinates for the boundary: mass geometry, or finite basis access.

4 Control Is the Handle, Not the Whole Claim

This creates one possible confusion that should be removed explicitly.

Because the falsifiable claim is stated through tracking capacity, entropy rate, and visibility loss, IOF can sound like it has reduced the measurement problem to a control problem.

That is not the intended claim.

The control problem is the experimental handle. It is the part of the framework that can be isolated, varied, and falsified in a laboratory.

The deeper claim concerns epistemic randomness.

In IOF, the measurement basis and the measured outcome are not treated as causally independent facts dropped into the experiment from separate origins. They are read as descendants of one physical history. The observer can record the surface facts: this basis, this outcome. But the observer cannot reconstruct the full causal ancestry by which this basis and this outcome arose together.

That inaccessible ancestry is not the same thing as ordinary controller noise.

Controller noise is a local engineering imperfection. It can often be calibrated, corrected, or modeled away.

Causal self-opacity is structural. It concerns the observer's inability, as an embedded finite system, to stand outside the causal process that produces its own measurement question.

The experiment does not measure that full ancestry directly. It cannot. Instead, it measures an accessible proxy: how well the observer-apparatus system can track the physical basis-producing state through which that ancestry becomes an implemented measurement setting.

So the hierarchy is:

- the deep claim is causal opacity from within a single history;
- the epistemic consequence is apparent randomness;
- the operational proxy is finite basis tracking;
- the falsifiable prediction is capacity-dependent visibility timing.

If the experiment fails, the proxy fails and IOF loses its scientific force in that regime.

If the experiment succeeds, the result should not be described merely as "control bandwidth affects visibility." That would be true but incomplete. Nor should it be over-stated as a direct demonstration of structural self-opacity: what the laboratory delivers is an *operational proxy* for that opacity—a measurable signature that observer-side basis access constrains quantum visibility. The deliberately throttled loop is an engineered, recoverable analogue of the in-principle opacity Thread A describes, not the opacity itself. The mechanism is what becomes visible; the metaphysics remains interpretation.

The experiment uses control theory, but the target is not control theory. The target is the origin of epistemic randomness in an embedded observer.

5 Where the Heisenberg Cut Sits

There is a complaint that runs through every interpretation of quantum mechanics. The Heisenberg cut—the boundary between the quantum description used for the measured system and the classical description used for the apparatus and the record—is treated as a convention. Von Neumann showed that the cut can be moved without changing predictions. Decoherence sharpens the picture but locates the cut by an external property, the rate of environmental coupling. Objective-collapse proposals fix the cut universally, at a mass or geometry scale, without reference to who is observing.

IOF's reframing places the cut elsewhere: where the observer-apparatus system's *useful* basis-tracking rate runs out relative to its basis-producing dynamics. The cut sits where the deficit vanishes,

$$h_{\mathrm{KS}} = C_{\mathrm{eff}} \ln 2 \qquad (\kappa = 0).$$

It is essential to be precise about what C_{eff} is here, because the cut's location depends on it. C_{eff} is *not* the Landauer ceiling. The Landauer bound,

$$C \le \frac{P}{k_B T \ln 2},$$

is a thermodynamic upper limit on irreversible bookkeeping. In cryogenic electronics it is typically enormous compared with the few-bit-per-second rate actually assigned to a particular basis-tracking loop [2]. The operative quantity is the effective rate $C_{\mathrm{eff}} = r\, b\, f$ that genuinely constrains the reference: the accepted update rate r, the useful bits per update b, and the surviving fraction f after latency, rejection, and filtering. The Landauer bound enters only as a consistency ceiling, $C_{\mathrm{eff}} \le C$, never as the cut's location.

This has a consequence that a purely thermodynamic framing would obscure: the cut is set largely by *design*, not handed down by thermodynamics. The experimenter can move it deliberately—by throttling or widening the tracking loop, changing estimator bandwidth or model order, or imposing a calibrated packet-drop schedule—at fixed temperature and power. That is exactly what makes the cut an experimental variable rather than a philosophical posit: the BLQC protocol moves C_{eff} and watches t_{break} move with it.

The cut is therefore observer-relative without being subjective. Two apparatuses tracking the same basis, with different loop designs, power budgets, or operating temperatures, will place their cuts at different points. But for a *given* configuration the cut is fixed by that configuration, and any observer inspecting the same hardware agrees on where it sits. What the experimenter controls is the configuration, not the verdict the configuration then yields.

This places the cut on an operational diagram, not on an interpretive convention. And it predicts something conventional cut placement does not: the cut moves. Cooling the apparatus, increasing the power actually delivered to useful tracking, or improving the controller raises C_{eff}, and the cut shifts outward, toward more chaotic basis-producing dynamics. The BLQC test is, in this language, an experiment that measures the motion of the cut.

The measurement problem has historically taken its sharpest form because the Heisenberg cut was treated as floating—a matter of descriptive convenience. IOF's claim is narrower and testable: for a given finite apparatus the cut is not floating but located, by the basis-tracking budget that apparatus actually devotes to its reference. The standard interpretations were not reading that ledger.

6 Bell, Entanglement, and Ontology: What Is Being Claimed?

There is one more place where the argument can be misunderstood.

IOF speaks of a single history, shared ancestry, and the failure of measurement independence as a fundamental assumption. A physicist may hear this as a claim to have solved Bell's theorem, or as an unrestricted appeal to hidden correlations.

That is not the claim of this document.

The falsifiable claim of IOF does not require a complete hidden-variable theory. For the experiment, the required assumptions are narrower:

- the measurement basis is a physical reference variable;
- the observer-apparatus system has finite access to that basis-producing state;
- unresolved basis uncertainty can reduce visibility;
- the visibility timing should scale with κ when ordinary confounds are controlled.

The broader single-history reading adds an interpretation. In that reading, the system state, the measurement basis, and the outcome belong to one physical history. The setting is not an external free parameter added from outside the world. It is a physical variable with causal ancestry.

This means measurement independence is not treated as a metaphysical necessity. It becomes an operational idealization: often excellent, but not fundamental in the single-history interpretation.

This is the same position set out under the First Objection, and it should be named without euphemism. It is *superdeterminism in the technical sense*: the statistical-independence (measurement-independence) premise of Bell's theorem is not imposed, because the setting and the system share ancestry. Following Palmer [5], this is the *non-conspiratorial* form—in a single globally consistent history the correlation is structural, not a fine-tuned coordination—and it is logically distinct from the colloquial sense in which "superdeterminism" is taken to mean a cosmic conspiracy. What the qualifier *epistemically bounded ancestral correlation* adds is that the shared root is not available to the embedded observer as a predictive hidden-variable ledger. IOF adopts the technical label and rejects the conspiratorial one; it does not pretend to stand outside the measurement-dependence family altogether.

This does not refute Bell's theorem. It changes which premise the interpretation accepts.

Bell's theorem shows that no theory satisfying its full set of assumptions can reproduce the observed quantum correlations [4]. IOF's single-history reading does not keep all those assumptions. In particular, it does not assume that the hidden state of the measured system and the later measurement setting are ancestrally independent in the strongest sense.

Nor does IOF claim that arbitrary hidden correlations can explain anything. That would make the framework empty. The only reason the ontology discussion has scientific force here is that it is tied to a separate, risky prediction: capacity-dependent visibility timing.

Entanglement should be read in the same restrained way.

In the single-history reading, entangled correlations need not be pictured as a signal sent from one wing of the experiment to the other. They can be read as correlations inside one globally consistent history, where outcomes are contextual: they depend on the system state and the physical measurement basis together.

But this document does not derive Bell correlations from first principles. It assumes that ordinary quantum correlations are recovered in the capacity-wins limit—inherited from the hosting embedding—and asks whether finite basis access adds a measurable visibility factor when tracking is stressed.

So the boundary is clear:

- Bell is not "solved" here;
- entanglement is not explained by a new signal;
- the standard quantum correlations are hosted, and the Born rule is not derived in this document (its binary form is conditionally reconstructed in a companion paper);
- ontology is not used as a substitute for prediction;
- the experimental claim remains the capacity-instability scaling of visibility.

If that scaling fails, the single-history interpretation gains no independent support from this document.

If that scaling succeeds, the ontology becomes harder to dismiss, because observer-side causal opacity would have acquired a measurable proxy.

The IOF Account of Quantum Randomness

This gives the IOF account of quantum randomness.

In the single-history IOF embedding, the outcome may be definite in the underlying history. The basis may also be definite. But the embedded observer lacks access to the joint causal ancestry of the system state and the basis-setting state. Because that ancestry cannot be reconstructed, the observer must represent the situation probabilistically.

Probability is therefore not inserted because nature failed to decide. It appears because the observer cannot know how the decision was already embedded in the total causal history.

On this view, collapse is not a physical event added to unitary dynamics. Collapse is an update in the observer's representation.

Before measurement, the observer has limited access to the relevant causal variables. It does not know the system's full ontic state. It does not know the full basis-producing state. It does not know the ancestral correlation between them. The best available description is therefore a probability distribution.

After measurement, the observer has a record. The basis was this. The outcome was that. The observer's epistemic state has changed. But the framework does not require the physical world to jump from many realities into one. One realized history was always the case; the observer's access to it changed.

This reframes the measurement problem in a precise, framework-internal sense. The reframing is conceptual, and it earns physical—as opposed to merely economical—standing only if the capacity-dependent visibility channel is found. Note also what is *not* produced here: the probability law itself. The standard quantum correlations are supplied by the chosen no-collapse embedding—primarily the single deterministic history of this document, with pilot-wave and Everett as alternatives that recover the same correlations by other routes; the single-system Born weight is conditionally derived in a companion paper and taken here as a constraint. IOF inherits these in the capacity-wins limit and contributes only the sub-unity visibility factor when tracking is stressed. IOF is a modifier on an existing interpretation, not a replacement for one.

What it does, in that bounded sense, is dissolve the external-basis idealization:

It does not do so by adding a collapse trigger.

It does not do so by multiplying worlds.

It does not do so by pretending that ordinary hidden variables can be freely inspected.

It rejects the external-basis idealization: the measurement basis is a physical variable with causal ancestry, the outcome depends on the system in that context, the setting and the system need not be ancestrally independent, and the embedded observer cannot trace the ancestry that binds them. The physical content of that move is staked, in full, on the single falsifiable channel.

What appears as collapse is the moment a finite observer updates from a probability distribution to a record.

What appears as randomness is causal opacity from within a single history.

What appears as loss of quantum visibility, in the regimes where IOF departs from standard quantum mechanics, is finite-rate basis tracking becoming experimentally visible.

The clean experimental question is therefore not whether the experimenter had metaphysical freedom. It is whether visibility depends on the effective capacity with which the observer-apparatus system can track its own measurement basis.

If visibility loss follows ordinary environmental decoherence only, IOF's physical claim fails.

If, with ordinary confounds controlled, visibility also follows the deficit

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

then the measurement problem has not been solved by a new collapse mechanism. It has been reframed as a limit of embedded self-knowledge.

The observer asks a question whose causal origin it cannot fully know.

The world answers once.

The observer calls the answer random.

References

[1] A. Dekker, "The Ignorant Observer Framework," OSF Preprints (2026). DOI: 10.17605/OSF.IO/FCDSN

[2] A. Dekker, "Bandwidth-Limited Quantum Control: A finite-rate phase-reference test in the Penrose-overlap regime," OSF Preprints (2026). DOI: 10.17605/OSF.IO/G5WRH

[3] A. Dekker, "Prospective Tests of Bandwidth-Limited Observer Dynamics: Experimental Protocol for the Ignorant Observer Framework," OSF Preprints (2026). DOI: 10.17605/OSF.IO/2QJNE

[4] J.S. Bell, "On the Einstein-Podolsky-Rosen paradox," *Physics Physique Fizika* **1**, 195–200 (1964).

[5] T. Palmer, "Superdeterminism without conspiracy," *Universe* **10**(1), 47 (2024). DOI: 10.3390/universe10010047

[6] G. 't Hooft, *The Cellular Automaton Interpretation of Quantum Mechanics*, Fundamental Theories of Physics 185, Springer (2016).

[7] R. Penrose, "On Gravity's Role in Quantum State Reduction," *Gen. Relativ. Gravit.* **28**, 581–600 (1996).

[8] G. Lindblad, "On the generators of quantum dynamical semigroups," *Commun. Math. Phys.* **48**, 119–130 (1976).