Where Did the Measurement Basis Come From?

Finite Basis-Tracking and the Measurement Problem

https://ignorantobserver.xyz

Aernoud Dekker

June 2026

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

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

When $\kappa < 0$, textbook fixed-basis quantum statistics are operationally available. When $\kappa > 0$, standard quantum mechanics still holds, but the observer's reference frame is unresolved and the raw record carries an additional visibility factor,

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

in the Gaussian unresolved-basis limit. This factor is observer-relative and recoverable—reference-frame physics within quantum mechanics, not a new decoherence channel; the stronger, unrecoverable reading is excluded by existing recoverability experiments, as the foundational paper shows [1]. What remains is a calibrated control law for when a finite observer can keep its own measurement question resolved—and an account of why, when it cannot, a definite history reads from the inside as probability.

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 [2] 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 [3]. Hidden-variable theories keep one outcome, but add an unseen state that determines it [4].

The Ignorant Observer Framework (IOF) [1]—which reads apparent collapse as the tracking limit of a finite observer—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 much is not new—and that is the point. That a reference frame is a physical system with finite resources, rather than an external abstraction, is already the premise of the quantum-reference-frame programme [5], which treats the frame as a bounded resource subject to superselection. That programme is IOF's honest lineage, precisely because it is standard quantum mechanics. What IOF adds is dynamical: the basis must be *tracked* by a finite-rate loop, and where its instability outruns that capacity, the observer's unconditioned records lose contrast on a calibrated schedule—capacity-dependent and, like all reference-frame dephasing, recoverable when the missing reference information is supplied. Quantum reference frames reorganize standard quantum mechanics and predict no departure from it, and neither does IOF: the contribution is the control-theoretic threshold κ and the tracking dynamics built on it, not a predicted deviation. The new ingredient is the finite-rate tracking, not the recognition that the frame is physical.

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 necessity. 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 [6, 7]. 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 conspiracy worry this raises is a fair one; it is taken up directly in the first objection below.)

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 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 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 structural quantum formalism still has to be hosted by the chosen ontology, or taken as the operational quantum rule that bounded observers use—though the single-system binary Born weight is itself conditionally derived from finite observer records in a companion paper [8], narrowing what the host must postulate.

IOF's own 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.

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 [9] 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, $C_{\text{eff}} \ln 2$, is the single physical number that governs the framework:

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

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, and textbook fixed-basis quantum mechanics is operationally available. Where $\kappa > 0$, chaos wins: standard quantum mechanics still holds, but the basis-producing dynamics outrun the observer's self-tracking; no tracking strategy, however efficient, can keep up; the represented basis drifts from the implemented one, and the observer's unconditioned record carries a specific attenuation of visibility,

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

in which V_{obs} is the observed fringe visibility and σ_θ^2 is the variance of the residual basis-tracking error. A compact derivation of this visibility law is given in Appendix A. The formula separates the loss of visibility into two multiplicative channels. The first, V_{std}, is the visibility ordinary quantum mechanics already predicts—the familiar attenuation of fringe contrast as the system couples to its environment (environmental decoherence)—and IOF leaves it untouched. The second, the exponential factor, is what IOF names: the observable mark of finite self-tracking in the unconditioned record, unity when tracking keeps up and falling below unity exactly when the deficit goes positive. Unlike the first channel, it is recoverable—conditioning the records on an independent log of the realized basis restores the contrast—which identifies it as reference-frame physics within standard quantum mechanics rather than a new decoherence mechanism; the stronger, unrecoverable reading of the same factor is stated and excluded by existing experiments in the foundational paper [1]. Because the two channels multiply, the self-tracking factor can in principle be isolated from ordinary decoherence, and everything exercised by the Bandwidth-Limited Quantum Control (BLQC) experiment [10]—a finite-rate phase-reference test developed in a companion paper—lives in that second factor.

Step back from the formula, though, and the whole account rests on a single conceptual move: the measurement basis and the system it acts on are not independent, because both inherit the same inaccessible ancestry. That move is what makes the explanation work—and it is also exactly the kind of move that can look like relocating a mystery rather than dissolving one. Before trusting it, that suspicion deserves a direct answer.

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?

IOF does not derive the Born rule from ancestral correlation; it takes the rule as an empirical constraint and asks a different question.

The role of IOF is not to use hidden correlations as an ad hoc 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 [11] 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, the framework declines to assume *measurement independence*: because the setting and the system share causal ancestry, $P(\xi \mid \theta) = P(\xi)$ is an additional assumption rather than a theorem, and the framework neither needs nor presumes it—the embedding permits the correlation without forcing it. No-signalling is preserved exactly, and the setting–system dependence is concealed—unreadable and unexploitable from within the history, its only operational fingerprint being the standard quantum statistics.

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 [12], 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:

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

So if quantum mechanics predicts

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

IOF predicts, in the Gaussian basis-error regime,

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

The factor touches only the joint correlation term: each wing's single-party marginals remain exactly 50/50, independent of the other wing's capacity, so nothing here signals—throttling one observer's tracking loop changes nothing the other can measure.

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. The extra factor is the capacity-dependent contrast of the *raw, unconditioned* record—and its status must be stated plainly: it is what standard quantum mechanics predicts for an online observer averaging over an unresolved reference. Logged-setting Bell tests have in effect already run the conditioned version of this experiment: when the realized settings are recorded and the data sorted offline, the full quantum correlation returns [1]. The factor is therefore a calibrated control law, not a deviation from quantum mechanics.

If changing effective basis-tracking capacity does nothing to the raw record, while ordinary decoherence variables are controlled, the apparatus model or the calibration of C_{eff} and h_{KS} is wrong—the law itself is classical control physics.

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 measurable work in the record—as expected, and now quantified.

The objection therefore sharpens the framework rather than defeating it.

IOF is not saying: "shared ancestry can be invoked to explain away any result whatsoever."

It is saying: "textbook statistics are operationally available when basis tracking is stable; where the observer-apparatus can no longer track the physical basis it is using, its unconditioned records lose contrast on a calibrated, recoverable schedule."

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 like all interpretations it is judged by coherence and economy, not by a discriminating experiment.

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

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

The two threads are related, but they do not have the same status, and the relation must be stated without overdraft. Thread B is true, measurable, and classical: it is the data-rate theorem applied to a reference loop, fully within standard quantum mechanics. It therefore cannot lend Thread A discriminating force—no experiment on Thread B chooses IOF over other interpretations. What Thread B does supply is a way to make Thread A's vocabulary *operational.* "Self-ignorance" is not a metaphor here; it is a calibrated rate with units, knobs, and a laboratory realization. A physicist may use Thread B without accepting the ontology; the ontology, in turn, is the reading under which Thread B's quantities are about something.

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 benchmark's terms, a recovery statistic $R_{\text{rec}} \to 1$ on this component is the signature that the loss belongs to the reference channel, exactly as the framework classifies it; $R_{\text{rec}} \to 0$ (irreversible loss) is the signature of genuine decoherence [10].

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

First, there is ordinary environmental decoherence [13, 14]. Here visibility is lost because the system becomes physically entangled with uncontrolled environmental degrees of freedom. Better bookkeeping cannot restore it ($R_{\text{rec}} \to 0$), 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 operational claim of IOF. For a given online observer-controller, unresolved basis uncertainty should not be arbitrary. It should scale with the deficit

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

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 same rate Zurek and Paz find saturating the entropy production of a chaotic system under decoherence, independent of the environmental coupling strength [15, 16]. Because that coupling never enters κ, the framework inherits the coupling-independence rather than reversing it: C_{eff} is a distinct tracking channel that subtracts from h_{KS}, recovering the passive rate as $C_{\text{eff}} \to 0$; the foundational paper [1] develops the reconciliation.

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,

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

For a reader arriving from decoherence theory, it is convenient to name the two channels introduced earlier:

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

IOF does not deny $V_{\rm std}$. It claims only that, in stressed regimes, part of the observed visibility loss may belong to $V_{\rm IOF}$, 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 [17].

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 benchmark—moving $C_{\rm eff}$ and $h_{\rm KS}$ against fixed confounds, and computing $R_{\rm rec}$ from a passive high-resolution shadow log—is set out in the companion BLQC paper [10].)

If the fitted loss tracks temperature, idle time, actuator distortion, pulse noise, or environmental coupling better than κ, the loss belongs to the first two cases and the tracking loop is not the bottleneck. If it collapses against an independently calibrated κ, the observer-side channel is calibrated—which is the expected outcome, since the law is classical control physics; what the collapse buys is not a win over quantum mechanics but a demonstrated, quantitative grip on the reference channel that the discrimination experiments of the protocol then rely on.

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

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

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

This is where the experiment becomes nontrivial.

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

If changing $C_{\rm eff}$ 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 visibility experiment, after thermal, readout, latency, pulse, actuator, and offline-recovery controls are included, the loss timescale still collapses against

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

better than against ordinary decoherence variables, then the observer-side visibility channel is real, calibrated, and correctly classified: reference physics under the observer's own control law.

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, calibrated 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 classification the framework itself insists on: recoverable, κ-scaled loss is the observer-relative channel working as described; anything unrecoverable belongs to decoherence or—if it tracked κ—to physics beyond quantum mechanics, a reading the foundational paper states and shows is excluded by existing data [1].

3 Control Is the Handle, Not the Whole Claim

There is one possible confusion here that should be removed explicitly.

Because the operational 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.

And the handle is not peripheral: it is the framework's center of gravity, while the hosting ontology is scaffolding—swap the host and both the machinery and its exposure to refutation are unchanged.

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 process that produces its own measurement question. Concretely, the experiment asks: if we deliberately degrade that access—the observer-apparatus's grip on the basis-producing process—does the fringe visibility change in the way $\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2$ predicts?

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 proxy's measurable form is capacity-dependent visibility timing in the unconditioned record—a classical control law, calibrated rather than discriminating.

If the calibration fails, the proxy is not under experimental command in that regime, and the framework's vocabulary loses its operational anchor there—though the law itself, being standard physics, is not in question.

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.

4 Where the Heisenberg Cut Sits

The measurement problem takes its sharpest form because the Heisenberg cut—the boundary between the quantum description of the measured system and the classical description of the apparatus and record—has been treated as floating. Von Neumann showed it can be moved without changing predictions [18]. Decoherence locates it by an external property, the rate of environmental coupling. Objective-collapse proposals fix it at a mass or geometry scale, without reference to who is observing.

IOF places the cut elsewhere: where the observer-apparatus system's *useful* basis-tracking rate runs out relative to its basis-producing dynamics—where the deficit vanishes, $\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2 = 0$. The quantity that matters is the effective tracking rate actually devoted to the reference, not the Landauer ceiling, which is typically enormous by comparison and enters only as a consistency bound.

This has a consequence a mass-based framing would miss: the cut is set largely by *design*. The experimenter can move it—by throttling or widening the tracking loop—at fixed temperature, power, and mass. It is observer-relative without being subjective: two apparatuses tracking the same basis with different loop designs place their cuts differently, but for a given configuration the cut is fixed, and any observer inspecting the same hardware agrees where it sits. In this language the BLQC test [10] measures the motion of the cut.

IOF's claim here is narrow and operational—observer-relative bookkeeping within standard quantum mechanics: 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.

5 Is the Boundary Set by Mass, or by the Observer?

That moving cut brings IOF into contact with an older proposal about where classical behaviour enters.

Penrose suggested that a superposition of two mass distributions is unstable on its own [19], collapsing after a time fixed by the gravitational self-energy E_G of the difference between them, $\tau_{\rm OR} = \hbar/E_G$. This is a gravity-side timescale: it depends on mass and geometry, and says nothing about who is watching.

IOF supplies a second timescale, from the other side. The observer loses track of the basis on a timescale fixed by its information deficit, $\tau_\kappa = 1/\kappa$. This one depends on the apparatus's tracking budget, not on the mass.

For mesoscopic masses the two can land in the same 10–100 ms window. The coincidence is numerical—the two estimates come from unrelated principles—but it locates the regime where the framework's control variables become genuinely useful experimental axes: the same mesoscopic apparatus that throttles a tracking loop can hold mass geometry fixed, and vice versa, which is exactly the crossed design a clean test of objective reduction requires.

The test is the one the cut already hands us, set out in full in the companion protocol [20]. On an apparatus that exposes both knobs—a QGEM-class experiment is the natural setting, with mesoscopic masses on one side and a real control loop on the other—the decisive classifier is *recoverability* against a passive log of the realized reference. The possible outcomes are clean:

- The raw-record loss moves with κ and is fully recoverable from the logged reference: the control law is calibrated—the expected outcome, and the demonstration that the apparatus resolves the relevant physics.
- A κ-independent, *unrecoverable* visibility floor follows mass geometry with tracking at maximal capacity: evidence for Penrose-type objective reduction. This is the live discrimination the design exists to make.
- No unrecoverable floor at sensitivities where OR predicts one: objective reduction is constrained at those parameters.
- A κ-scaled loss that survives conditioning on the logged reference: this would realize the stronger reading of the visibility law—excluded by existing data in the diffusive and externally randomized regimes, and carried only as a registered test for certified-chaotic basis dynamics, with the prior strongly against it [1, 10].

The crossed design turns one piece of the measurement problem into something an experiment can answer:

> Is there an unrecoverable classicalization floor set by mass geometry—or is every observed loss, once the reference is restored, the recoverable bookkeeping of a finite observer?

In one line: the definiteness an observer meets is objective *indexed to its configuration*—irreducible at that access, found the same way by every inspector there—and Penrose's proposal is that the index can be dropped. Recoverability is that question made operational: *does the index drop?*

That is a sharp fork—and the reason the two timescales meeting in the same window is worth an experiment. The fork is also wider than its flagship: any dynamical-collapse rival—spontaneous-localization (GRW/CSL-type) models as much as gravitational reduction—predicts some unrecoverable floor with its own scaling variable, and the same calibrated subtraction reads against whichever schedule is on trial; Penrose's is simply the case where the two timescales meet in an accessible window.

A third scale is worth noting, but only as context. Using representative biological-scale estimates [21, 22], $h_{\mathrm{KS}} \approx 50\ \mathrm{s}^{-1}$ and $C_{\mathrm{eff}} \approx 10$ bit/s give $\kappa \approx 43\ \mathrm{s}^{-1}$, hence an intrinsic tracking e-folding time $\tau_\kappa \approx 23$ ms. With the logarithmic threshold factor needed for an observable visibility or perceptual-integration break, the corresponding operational window can move into the 50–70 ms range. Thus a gravitational collapse scale, a laboratory tracking-loss scale, and the timescale on which a biological observer integrates experience can all fall within the same broad decade.

This convergence is suggestive, not evidential. It does not support the framework by itself, because the biological estimate is another use of the same κ-scaling, not an independent test of it. A finite observer built from biological-scale capacities should naturally classicalize on biological-scale times. The empirical weight remains where it belongs: on whether the measured visibility-loss timescale moves with κ under controlled variation of C_{eff} and h_{KS}. The fact that information, gravity, and biological observation may meet near one boundary is worth recording; what, if anything, it means is laid out in the interpretive layer of the corpus [1].

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 operational 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 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 [6]. 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.

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.

IOF 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;
- no explicit hidden-variable model $f(\xi, \theta)$ reproducing the singlet correlations is constructed here;
- the standard quantum correlations are hosted, and the Born rule is not derived in this document (its binary form is conditionally derived in a companion paper [8]);
- ontology is not used as a substitute for prediction;
- the operational content remains the capacity-instability scaling of unconditioned-record visibility—a calibrated control law within standard quantum mechanics.

The interpretation does not gain discriminating support from that law; nothing can choose between empirically equivalent interpretations. What the law gives the ontology is an operational anchor: observer-side causal opacity has a measurable proxy, with units, knobs, and a laboratory realization.

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—a coarse-graining over the inaccessible ancestral variables, forced by the observer's embedding rather than chosen for convenience.

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 its standing is interpretive: the capacity-dependent visibility channel is real and measurable, but it is recoverable reference physics within standard quantum mechanics, so the reframing is judged by its economy and coherence rather than by a pending discriminator. Note also what is *not* produced here: the full probability law itself. The division of labor is as follows. The hosting embedding—primarily the single deterministic history of this document, with pilot-wave and Everett as alternatives that recover the same correlations by other routes—supplies the structural formalism: the admissible-history measure, Hilbert space, composition, the multi-outcome rule, and the quantum correlations.

One element of that formalism is *derived* rather than postulated: the single-system binary Born weight $\cos^2(\theta/2)$ is conditionally derived from finite observer records in a companion paper [8], retiring it as a separate ontic postulate instead of inheriting it. This is not a separate concession but the same ingredient IOF rests on, seen from another side—the observer's finite information capacity. As finite *resolution* of records, that one capacity is what the companion paper turns into the binary weight itself; as finite *tracking* of the basis, it is what sits on either side of that law—upstream, the opacity account of *why* a definite history must read as probability at all, and downstream, the sub-unity visibility factor V_{IOF} when tracking is stressed. In the capacity-wins limit IOF reproduces the hosted statistics exactly: a modifier on an existing interpretation, not a replacement for one—but one whose host is asked to postulate one thing fewer.

The binary Born derivation is not a toy example. Many of the cleanest quantum measurements are physically binary: spin up/down along a chosen axis, qubit (0/1) readout, photon polarization (H/V), or the two output ports of an interferometer. In such cases the apparatus supplies a two-record structure, but the probability weights over those two records still require a law. The companion binary Born derivation addresses exactly that step: given a physical two-outcome basis, why should the weights take the $\cos^2(\theta/2)$ and $\sin^2(\theta/2)$ form rather than arbitrary values?

The distinction is important. IOF does not claim here to derive the existence of a two-channel apparatus from nothing; that belongs to the physical observable, basis, and record structure of the experiment. What the binary Born result contributes is the probability law over that already-implemented binary question. In that sense it gives quantitative form to the idea of epistemic randomness: the observer's ignorance is not merely a verbal appeal to hidden ancestry, but can, in the binary case, produce the standard quantum weighting under the stated finite-record assumptions.

The full multi-outcome Born rule still belongs to the host formalism, and a companion paper [23] shows this is a clean boundary rather than a gap left for future work: the finite-record route that yields the binary weight provably cannot be extended to the full rule, which must be inherited rather than derived.

Read qualitatively, that boundary is the framework's hinge. The binary weight is the last structure a finite observer's own records generate; the lawful architecture beyond it—relative phase, interference, the full Hilbert formalism—is not authored by the observer but received, the substrate it finds itself within. The binary case is therefore not a foothold from which the rest of quantum mechanics will eventually be derived. It is the point at which the framework stops trying to ground the substrate and turns to its actual question: not what the world is made of, but how a single definite history comes to appear, to a finite embedded observer, as a world of definite outcomes and classical records. The limit theorem marks where that turn is forced—the threshold past which the substrate is received, and the rising of appearance becomes the only thing left to explain.

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 operational face of that move is the tracking law; its interpretive face is the single-history reading; and the paper keeps separate books for the two.

The measurement basis is therefore not selected from outside the theory by fiat. It is the physically implemented context $\theta(t)$, accessed only through the finite observer-apparatus channel. When the channel can track this context, the basis remains operationally sharp; when it cannot, unresolved basis uncertainty is coarse-grained into the observer's record.

The particle is not what the quantum state literally becomes. The observed particle-like event is the finite record registered when a finite observer-apparatus channel samples the underlying quantum structure through a contextual basis. This leaves open whether the particle has any further ontic status; what IOF claims is only that, for any finite observer, the particle appears as a stable bounded record.

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 tracking fails, is finite-rate basis tracking becoming experimentally visible in the unconditioned record—recoverable, observer-relative, and fully inside standard quantum mechanics.

The clean experimental question is therefore not whether the experimenter had metaphysical freedom. It is whether visibility in the raw record depends on the effective capacity with which the observer-apparatus system can track its own measurement basis—and it does, as a matter of classical control physics, now stated as a calibrated law.

The measurement problem has not been solved here by a new collapse mechanism, and no new mechanism was needed. It has been reframed as a limit of embedded self-knowledge—a reading whose operational vocabulary can be built, throttled, and measured, and whose metaphysics remains, openly, an interpretation.

Where to go next. The corpus reads in this order: the foundational paper [1] for the full framework, its ontology, and the exclusion of the stronger reading; BLQC [10] for the operational benchmark—calibration arm, capacity and instability sweeps, and the recoverability classifier; the experimental protocol [20] for the mesoscopic objective-reduction discrimination built on that benchmark; and the conditional Born-rule derivation [8], with its limit theorem [23], for the foundations module. This document is the doorway; those four carry the weight.

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

The world answers once.

The observer calls the answer random.

Appendix A: Minimal Formal Model of Finite Basis Tracking

This appendix collects the minimal formal skeleton behind the conceptual argument, so that the central claim can be assessed without reference to the companion papers. It is deliberately compact: only the quantities and relations needed to show that “finite basis tracking” has a precise operational form are included. The full control-theoretic development is given in IOF [1] and BLQC [10], and the experimental design in the protocol [20].

A.1 Basis variable, estimate, and tracking error

Let $\theta(t)$ be the physical measurement-basis variable—the axis, phase, threshold, or reference frame that fixes the question the apparatus poses. In IOF it is a state of the observer–apparatus system, not an external parameter. Let $\hat{\theta}(t)$ be the observer's internal estimate of that variable, maintained by a finite-rate control loop, and let

$$\delta\theta(t) = \theta(t) - \hat{\theta}(t)$$

be the residual tracking error. Its variance $\sigma_\theta^2(t) = \langle \delta\theta(t)^2 \rangle$ is the single quantity that couples to visibility.

A.2 Effective tracking capacity

The loop maintaining $\hat{\theta}$ has a finite *useful* information rate,

$$C_{\text{eff}} = r\, b\, f \quad [\text{bits/s}],$$

where r is the fraction of updates surviving loss, rejection, latency, filtering, and estimator overhead, b the useful bits per update that genuinely constrain the reference, and f the update rate. C_{eff} is the rate that actually constrains θ, not the raw device power or the Landauer ceiling [9], which is typically far larger.

A.3 Self-ignorance rate

The basis-producing dynamics generate fresh relevant information at the Kolmogorov–Sinai rate h_{KS} (nats/s). The governing quantity is the deficit between generation and useful tracking,

$$\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2 \quad [\text{nats/s}],$$

where $\ln 2$ converts C_{eff} from bits to nats. Its sign separates the regimes: $\kappa < 0$ (capacity-wins), tracking is maintained and ordinary quantum statistics are recovered; $\kappa > 0$ (chaos-wins), the loop cannot keep up.

A.4 Error growth in the chaos-wins regime

For $\kappa > 0$, uncorrected basis error inherits the exponential separation of the chaotic reference dynamics, reduced by the rate of corrective updates:

$$\frac{d}{dt} \ln \sigma_\theta^2 \approx 2\big(h_{\text{KS}} - C_{\text{eff}} \ln 2\big) = 2\kappa,$$

so that

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

with σ_0^2 the initial basis uncertainty at $t = 0$ set by calibration. Equivalently, the error amplitude grows as $\sigma_\theta(t) = \sigma_0\, e^{\kappa t}$, so the intrinsic amplitude e-folding time is $\tau_\kappa \equiv 1/\kappa$.

A.5 Visibility suppression

The observer intends basis θ_0 but realizes $\theta = \theta_0 + \delta\theta$ with $\delta\theta \sim \mathcal{N}(0, \sigma_\theta^2)$. Averaging the interference term over this Gaussian basis uncertainty,

$$\langle \cos(\phi - \theta) \rangle = \cos(\phi - \theta_0)\, e^{-\sigma_\theta^2/2},$$

gives the IOF visibility factor

$$V_{\mathrm{IOF}}(t) = \exp\!\left(-\tfrac{1}{2}\sigma_\theta^2(t)\right) = \exp\!\left(-\tfrac{1}{2}\sigma_0^2\, e^{2\kappa t}\right),$$

valid in the small-noise regime $\sigma_\theta \lesssim 1$ rad. This double-exponential decay is qualitatively distinct from ordinary exponential ($e^{-\gamma t}$) or Gaussian decoherence. The measured visibility is the two-channel product

$$V_{\mathrm{obs}}(t) = V_{\mathrm{std}}(t)\, V_{\mathrm{IOF}}(t),$$

with V_{std} the ordinary environmental-decoherence factor the framework does not modify.

A.6 Breakdown time

The primary observable is not an absolute visibility but the time at which it crosses a chosen threshold $V_* \in (0, 1)$. Setting $V_{\mathrm{IOF}}(t_{\mathrm{break}}) = V_*$ (equivalently $\sigma_{\mathrm{tol}}^2 = -2\ln V_*$),

$$t_{\mathrm{break}} = \frac{1}{2\kappa} \ln\!\left(\frac{-2\ln V_*}{\sigma_0^2}\right) \qquad (\kappa > 0).$$

The breakdown time is set by $\tau_\kappa = 1/\kappa$ up to a slowly varying logarithmic threshold factor; hence $t_{\mathrm{break}} \propto 1/\kappa$.

A.7 Validity conditions

The model is validated through the dependence of t_{break} on the control variables. With thermal load, readout SNR, pulse and actuator conditions, latency, and plant dynamics held fixed, the law requires

$$\frac{\partial t_{\mathrm{break}}}{\partial C_{\mathrm{eff}}} > 0, \qquad \frac{\partial t_{\mathrm{break}}}{\partial h_{\mathrm{KS}}} < 0.$$

These are validity conditions for the operational control law, not tests of quantum mechanics: the law follows from standard quantum mechanics plus the data-rate theorem, so a failure of t_{break} to shift with $\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2$—or a decay fitting ordinary exponential or Gaussian form rather than the double-exponential closure above—indicts the apparatus model or the calibration of C_{eff} and h_{KS} in that regime, not the underlying physics. The full benchmark procedure, including the calibration arm and the recoverability classifier, is developed in the companion BLQC paper [10]; the platform requirements, confound controls, and decision rules for the mesoscopic discrimination experiment are developed in the companion protocol [20].

References

[1] Aernoud Dekker. The ignorant observer: Avidya as a physical limit. *OSF Preprints*, 2026. doi: 10.17605/OSF.IO/FCDSN.

[2] Tim Maudlin. Three measurement problems. *Topoi*, 14(1):7–15, 1995.

[3] David Wallace. *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford University Press, Oxford, 2012.

[4] Detlef Dürr, Sheldon Goldstein, and Nino Zanghì. *Quantum Physics Without Quantum Philosophy*. Springer, Berlin, 2013.

[5] Stephen D. Bartlett, Terry Rudolph, and Robert W. Spekkens. Reference frames, superselection rules, and quantum information. *Rev. Mod. Phys.*, 79(2):555–609, 2007.

[6] John S. Bell. On the einstein-podolsky-rosen paradox. *Physics Physique Fizika*, 1:195–200, 1964.

[7] John S. Bell. *Speakable and Unspeakable in Quantum Mechanics*. Cambridge University Press, Cambridge, 1987. 2nd ed. 2004.

[8] Aernoud Dekker. The born rule from finite observation: A conditional derivation of the binary born form. *OSF Preprints*, 2026. doi: 10.17605/OSF.IO/U5RDE.

[9] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.*, 5:183–191, 1961. doi: 10.1147/rd.53.0183.

[10] Aernoud Dekker. Bandwidth-limited quantum control: A finite-rate phase-reference control law and its operational benchmark. *OSF Preprints*, 2026. doi: 10.17605/OSF.IO/G5WRH.

[11] Gerard 't Hooft. *The Cellular Automaton Interpretation of Quantum Mechanics*. Fundamental Theories of Physics 185. Springer, 2016.

[12] Tim Palmer. Superdeterminism without conspiracy. *Universe*, 10(1):47, 2024. doi: 10.3390/universe10010047.

[13] Wojciech H. Zurek. Decoherence, einselection, and the quantum origins of the classical. *Rev. Mod. Phys.*, 75:715–775, 2003.

[14] Maximilian Schlosshauer. Decoherence, the measurement problem, and interpretations of quantum mechanics. *Rev. Mod. Phys.*, 76:1267–1305, 2005.

[15] Wojciech H. Zurek and Juan Pablo Paz. Decoherence, Chaos, and the Second Law. *Physical Review Letters*, 72(16):2508–2511, 1994. doi: 10.1103/PhysRevLett.72.2508.

[16] Wojciech H. Zurek. Decoherence, Chaos, Quantum-Classical Correspondence, and the Algorithmic Arrow of Time. *Physica Scripta*, T76:186–198, 1998. doi: 10.1238/Physica.Topical.076a00186. Also Acta Phys. Pol. B 29(12), 3689 (1998).

[17] Göran Lindblad. On the generators of quantum dynamical semigroups. *Commun. Math. Phys.*, 48:119–130, 1976.

[18] John von Neumann. *Mathematical Foundations of Quantum Mechanics*. Princeton University Press, Princeton, 1955. Translation of *Mathematische Grundlagen der Quantenmechanik*, Springer, Berlin (1932).

[19] Roger Penrose. On gravity's role in quantum state reduction. *Gen. Relativ. Gravit.*, 28: 581–600, 1996.

[20] Aernoud Dekker. A capacity-calibrated protocol for testing penrose objective reduction. *OSF Preprints*, 2026. doi: 10.17605/OSF.IO/2QJNE.

[21] Jieyu Zheng and Markus Meister. The unbearable slowness of being: Why do we live at 10 bits/s? *Neuron*, 113:192–204, 2024. doi: 10.1016/j.neuron.2024.11.008.

[22] Kannathal Natarajan, U. Rajendra Acharya, Fadhilah Alias, Thelma Tiamklang, and Sadasivan K. Puthusserypady. Nonlinear analysis of eeg signals at different mental states. *BioMedical Engineering OnLine*, 3:7, 2004. doi: 10.1186/1475-925X-3-7.

[23] Aernoud Dekker. Why fisher geometry gives binary born but not quantum phase: A limitative theorem on markov-invariant record geometry. *OSF Preprints*, 2026. doi: 10.17605/OSF.IO/N3SU9.