Prospective Tests of Bandwidth-Limited Observer Dynamics

Experimental Protocol for the Ignorant Observer Framework

Aernoud Dekker

May 2026
Version 2.5

DOI: 10.17605/OSF.IO/2QJNE

Abstract

This document specifies a single prospective experiment: a mesoscopic visibility-loss test in the regime where Bandwidth-Limited Quantum Control (BLQC) and Penrose Objective Reduction (OR) can predict comparable timescales. Penrose OR assigns the loss timescale to gravitational self-energy, hence to mass geometry. BLQC assigns the loss timescale to finite-rate basis-reference tracking, hence to the deficit

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

where h_{KS} is the entropy-rate or instability proxy of the relevant basis-reference dynamics and C_{eff} is the useful information rate available to track or stabilize that reference.

The core question is therefore direct: at fixed thermal, readout, pulse, latency, and plant conditions, does the measured visibility-loss timescale follow mass/separation as Penrose OR predicts, or does it move with C_{eff} and h_{KS} as BLQC predicts? The protocol defines the required platform, control variables, confound controls, statistical tests, and decision rules for that discrimination.

Scope of this protocol. This protocol specifies a prospective Penrose-overlap test within the Ignorant Observer Framework (IOF). BLQC is the finite-rate basis-tracking mechanism evaluated here. The experiment is designed to distinguish a Penrose-style mass-geometry timescale from a BLQC-style capacity/instability timescale in the same mesoscopic apparatus.

A separate bridge-validation module tests whether BLQC effective capacity acts in Fisher-distinguishability units, which is required for the Paper A Born-rule bridge but not for the narrower BLQC/Penrose discrimination. A QGEM-style platform addendum maps the abstract protocol variables onto one candidate mesoscopic implementation.

Contents

1 The Experimental Question

The experiment targets the regime in which a mesoscopic quantum system exhibits a visibility-loss or collapse-like timescale of order milliseconds to hundreds of milliseconds. This is the regime where Penrose OR can become relevant for sufficiently massive or spatially separated superpositions, and where BLQC can also predict loss if the apparatus reference frame is maintained by a finite-rate tracking loop close to its instability boundary.

The decisive question is whether, in a mesoscopic experiment where Penrose OR is a serious candidate mechanism, the loss timescale follows:

- **mass geometry**, as Penrose OR predicts; or
- **basis-reference tracking capacity**, as BLQC predicts.

The same apparatus must expose both sets of variables: the mass-geometry variables entering the Penrose estimate and the basis-reference variables entering the BLQC estimate.

2 Competing Predictions

2.1 Penrose Objective Reduction

Penrose OR predicts a collapse-like timescale set by gravitational self-energy:

$$\tau_{\mathrm{OR}} \approx \frac{\hbar}{E_G}, \tag{1}$$

where E_G depends on the mass distribution mismatch between the branches of a spatial superposition. In simplified scaling form, for mass m and separation s,

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

up to geometry-dependent factors and saturation behavior at large separations.

Penrose OR therefore predicts that the loss timescale should change when mass, separation, or mass distribution changes, and should not have a leading dependence on the classical controller's useful basis-tracking capacity once ordinary experimental confounds are controlled.

2.2 BLQC

BLQC treats the measurement basis as a physical reference variable $\theta(t)$ implemented by apparatus: phase reference, local oscillator, pulse axis, field direction, path-separation reference, timing system, or equivalent control state. If this reference has instability or entropy-rate proxy h_{KS} and is tracked through useful capacity C_{eff}, BLQC defines

$$\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2. \tag{3}$$

In the chaos-wins regime ($\kappa > 0$), BLQC predicts growth of unresolved basis uncertainty:

$$\sigma_\theta^2(t) = \sigma_0^2 e^{2\kappa t}. \tag{4}$$

If visibility is reduced by Gaussian basis uncertainty, then

$$V(t) = \exp\left[-\frac{1}{2}\sigma_0^2 e^{2\kappa t}\right]. \tag{5}$$

For a chosen threshold V_*, the breakdown time is

$$t_{\text{break}} = \frac{1}{2\kappa} \ln\left(\frac{-2\ln V_*}{\sigma_0^2}\right). \tag{6}$$

BLQC therefore predicts that, at fixed mass geometry and fixed environmental controls, increasing useful C_{eff} delays loss, while increasing h_{KS} accelerates loss.

2.3 Discrimination Table

Controlled change	**Penrose OR**	**BLQC**
Increase mass at fixed C_{eff}, h_{KS}, and separation	Faster loss	No leading BLQC change unless reference dynamics also change
Increase separation at fixed C_{eff}, h_{KS}, and mass	Geometry-dependent OR shift	No leading BLQC change unless reference dynamics also change
Increase C_{eff} at fixed mass, separation, temperature, and readout	No leading OR change	Delayed loss
Increase h_{KS} at fixed mass, separation, temperature, and readout	No leading OR change	Faster loss

The primary discriminator is not the absolute timescale. It is the derivative of that timescale with respect to independently controlled variables.

3 Required Platform

A suitable platform must be mesoscopic enough for Penrose OR estimates to be meaningful and instrumented enough for BLQC variables to be exposed. Candidate classes include optomechanical, levitated-mass, cold-atom, interferometric, and QGEM-style pathfinder systems. The platform must preserve a usable visibility signal while also exposing the reference-tracking channel as an experimental variable.

Required capabilities:

- Create or probe a mesoscopic superposition or visibility signal with a measurable loss timescale.
- Vary mass, spatial separation, or mass distribution over a range large enough to change the Penrose OR estimate.
- Identify the physical basis-reference variable $\theta(t)$ used to define the measurement.

- Estimate or impose the reference instability/entropy-rate proxy h_{KS}.
- Vary useful basis-tracking capacity C_{eff} without changing mass geometry.
- Keep temperature, readout SNR, pulse/actuator behavior, sequence duration, and environmental coupling under monitored control.
- Preserve high-resolution reference logs whenever possible, so observer-relative reference loss can be separated from irreversible physical decoherence.

Exclusion criteria:

- The measurement basis is treated only as an external setting and cannot be monitored as a physical reference.
- Capacity variation changes temperature, readout backaction, or sequence duration in an uncontrolled way.
- Mass-geometry variation also changes the reference-tracking channel in a way that cannot be measured or modeled.
- The visibility-loss timescale is dominated by ordinary environmental decoherence with no usable dynamic range for either OR or BLQC variables.

4 Operational Variables

4.1 Mass-Geometry Variables

The Penrose side of the experiment requires independently specified mass-geometry variables:

- m: effective mass participating in the superposition or interferometric path distinction.
- s: spatial separation or equivalent branch-distinguishability scale.
- Geometry: shape and mass distribution relevant to E_G.

These variables must be changed while holding C_{eff}, h_{KS}, temperature, readout SNR, and sequence duration fixed or explicitly modeled.

4.2 Effective Basis-Tracking Capacity

The BLQC side requires a useful tracking-capacity variable:

$$C_{eff} = r\, b\, f, \tag{7}$$

where r is the useful update rate, b is the useful bits per update constraining the reference, and f is the fraction of updates that survive loss, rejection, latency, filtering, and estimator overhead.

C_{eff} is not raw controller power, raw ADC bandwidth, or a Landauer upper bound. It is the information rate by which the actual basis reference is constrained in the experiment.

Acceptable ways to vary C_{eff} include:

- changing accepted reference-update rate;
- changing useful bit depth;
- imposing calibrated packet/dropout schedules;
- changing estimator bandwidth or model order;
- changing filtering while preserving final readout SNR.

Artificial delay may be used only with latency-matched controls, because extra waiting time can produce ordinary coherence loss.

4.3 Reference Instability

The reference instability variable is h_{KS} or a calibrated proxy, such as a Lyapunov rate or prediction-error growth rate of the relevant basis-reference dynamics. It must be estimated from the same apparatus used in the Penrose-overlap test.

Acceptable estimation methods include:

- prediction-error growth from logged controller/plant states;
- Lyapunov exponent estimation for a fitted reference-dynamics model;
- calibrated injected reference instability on the target apparatus, with the injection path included in the platform model.

5 Experimental Design

The experiment should be run as a crossed or partially crossed design over mass-geometry and BLQC variables.

5.1 Baseline

For each platform, establish a baseline visibility curve $V(t)$ under stable reference tracking, stable temperature, stable readout SNR, and fixed mass geometry. The baseline must have enough visibility and dynamic range to identify a breakdown threshold V_*.

5.2 Capacity Sweep at Fixed Geometry

At fixed mass, separation, temperature, readout SNR, pulse/actuator behavior, sequence duration, and estimated h_{KS}, vary C_{eff} across a pre-registered range. Extract t_{break} for each condition.

BLQC predicts:

$$\frac{\partial t_{\mathrm{break}}}{\partial C_{\mathrm{eff}}} > 0 \quad (\kappa > 0). \tag{8}$$

Penrose OR predicts no leading dependence on C_{eff} when mass geometry and ordinary confounds are fixed.

5.3 Reference-Instability Sweep at Fixed Geometry

At fixed mass, separation, temperature, readout SNR, pulse/actuator behavior, sequence duration, and C_{eff}, vary h_{KS} or its calibrated proxy.

BLQC predicts:

$$\frac{\partial t_{\text{break}}}{\partial h_{\text{KS}}} < 0 \quad (\kappa > 0). \tag{9}$$

Penrose OR predicts no leading dependence on reference instability when mass geometry and ordinary confounds are fixed.

5.4 Mass-Geometry Sweep at Fixed Tracking Variables

At fixed C_{eff}, h_{KS}, temperature, readout SNR, pulse/actuator behavior, and sequence duration, vary mass, separation, or geometry.

Penrose OR predicts a geometry-dependent shift in τ_{OR}. BLQC predicts no leading shift unless the geometry change also alters C_{eff} or h_{KS}.

5.5 Randomization and Replication

Run order should be randomized across capacity, instability, and geometry settings. Each condition must include enough repetitions to estimate visibility with the pre-registered precision target. Replication should include at least one repeated condition after the full parameter sweep to detect slow drift.

6 Confound Controls

Thermal control. Temperature must be monitored at the mixing chamber and, where possible, through chip/platform proxies such as frequency drift, baseline coherence, or mechanical mode behavior. Capacity sweeps must not be accepted if they introduce uncontrolled heating.

Readout control. Readout SNR, detection efficiency, and readout-induced backaction must be monitored and included as nuisance variables. A capacity effect that tracks readout degradation is not evidence for BLQC.

Latency and wait-time control. Sequence duration and idle time must be matched across capacity conditions. If lower capacity merely adds waiting time, ordinary coherence loss explains the result.

Pulse and actuator control. Pulse amplitude, phase, actuator response, and plant diagnostics must remain within pre-registered tolerances. A capacity effect that tracks actuator distortion is confounded.

Offline reference logs (Shadow Channel). Where technically possible, the realized reference should be logged at higher resolution than the online tracker receives. The high-resolution reference log must be a passive, causally isolated shadow channel: it may be used for retrospective reconstruction but must not contribute to the online C_{eff} available during the run. Retrospective recovery of visibility via the R_{rec} statistic classifies the effect as observer-relative reference loss rather than irreversible physical decoherence.

7 Statistical Analysis

7.1 Primary Observable

The primary observable is the breakdown time t_{break}, defined by a pre-registered visibility threshold:

$$V(t_{\text{break}}) = V_*. \tag{10}$$

Secondary observables include fitted decay-rate parameters, recovered visibility after offline reference correction, and model residuals across the parameter grid.

7.2 Models

Fit at least the following model classes:

BLQC model:

$$t_{\text{break}} = \frac{A}{h_{\text{KS}} - C_{\text{eff}} \ln 2} + B, \tag{11}$$

or the threshold-derived logarithmic form when σ_0 is estimated.

Penrose OR model:

$$t_{\text{break}} = A\, \tau_{\text{OR}}(m, s, \text{geometry}) + B. \tag{12}$$

Combined-rate model:

$$\frac{1}{t_{\text{break}}} = \alpha\, \frac{1}{\tau_{\text{OR}}} + \beta\, \kappa + \gamma, \tag{13}$$

allowing both mechanisms or mechanism-like rates to contribute.

Nuisance models: standard exponential, Gaussian, stretched-exponential, thermal/readout, latency, and actuator-response models.

7.3 Geometry–Capacity Covariance and Mediation Check

In addition to fitting the additive combined-rate model,

$$\frac{1}{t_{\text{break}}} = \alpha\, \frac{1}{\tau_{\text{OR}}} + \beta\, \kappa + \gamma, \tag{14}$$

the analysis must test whether changes in mass geometry alter the empirical BLQC variables themselves.

For each mass-geometry condition, estimate:

$$C_{\text{eff}}^{\text{emp}}, \quad h_{\text{KS}}^{\text{emp}}, \quad I(\theta), \quad \kappa_{\text{eff}} = h_{\text{KS}}^{\text{emp}} - C_{\text{eff}}^{\text{emp}} \ln 2. \tag{15}$$

Then test whether the Penrose variable $1/\tau_{\text{OR}}$ predicts these quantities:

$$\frac{1}{\tau_{\text{OR}}} \to C_{\text{eff}}^{\text{emp}},\ h_{\text{KS}}^{\text{emp}},\ I(\theta),\ \kappa_{\text{eff}}. \tag{16}$$

If $1/\tau_{\text{OR}}$ remains an independent predictor of $1/t_{\text{break}}$ after controlling for κ_{eff}, the result supports an additive two-mechanism interpretation.

If the apparent Penrose dependence is mediated by changes in κ_{eff}, $I(\theta)$, or Fisher-unit tracking capacity, the result supports a unified geometry-capacity interpretation.

7.4 Decision Rules

Support for BLQC in the Penrose-overlap regime requires all of the following:

- At fixed mass geometry and fixed ordinary confounds, t_{break} increases with C_{eff}.
- At fixed mass geometry and fixed ordinary confounds, t_{break} decreases with h_{KS}.
- Breakdown times collapse better against κ or $\rho = C_{\text{eff}} \ln 2 / h_{\text{KS}}$ than against raw power, temperature, readout SNR, elapsed time, or actuator diagnostics.
- The capacity/instability effect is not fully explained by latency-matched ordinary coherence loss or offline-recoverable reference bookkeeping alone.

Support for Penrose OR or another mass-geometry mechanism requires:

- t_{break} follows mass/separation/geometry according to the OR estimate or another specified geometry-dependent model.
- t_{break} is independent of C_{eff} and h_{KS} within the experiment's sensitivity after ordinary confounds are controlled.

A combined result is possible. If both mass geometry and κ improve prediction, the result should be modeled as overlapping rates rather than assigned to either mechanism by inspection.

A dual-success result must also be classified as either additive or mediated:

- **Additive dual success:** Both $1/\tau_{\text{OR}}$ and κ independently improve prediction of $1/t_{\text{break}}$, with no significant covariance between mass geometry and empirical BLQC variables.
- **Mediated dual success:** Mass geometry significantly changes $C_{\text{eff}}^{\text{emp}}$, $h_{\text{KS}}^{\text{emp}}$, $I(\theta)$, or κ_{eff}, and the independent Penrose term weakens after those variables are included.

8 Bridge-Validation Module: Fisher Homogeneity

This module tests the specific theoretical bridge connecting BLQC to the Born probability law. The bridge requires that the useful component of C_{eff} operates as Fisher-distinguishability capacity, and that a globally scalar threshold κ corresponds to a constant Fisher information across the calibrated basis reference dial.

8.1 Operational Record Calibration

Before executing the capacity and instability sweeps, the experimenter must map the operational record family

$$p(o|\theta) \tag{17}$$

across a fine-grained, calibrated grid of the physical basis setting θ.

8.2 Empirical Fisher Information

From the empirical probability records, compute the observed Fisher information for the apparatus reference dial:

$$I(\theta) = \sum_{o} \frac{1}{p(o|\theta)} \left(\frac{\partial p(o|\theta)}{\partial \theta} \right)^2 . \tag{18}$$

8.3 Homogeneity Test

Test whether the Fisher information remains approximately constant over the relevant operating range:

$$I(\theta) \approx \alpha^2. \tag{19}$$

A strictly position-dependent $I(\theta)$ without a correspondingly position-dependent hardware capacity allocation violates the scalar-homogeneity premise.

8.4 Fisher-Unit Error Reporting

During the capacity sweeps in Section 5.2, unresolved reference error must be reported not just in physical coordinate units σ_θ, but in invariant Fisher distinguishability units:

$$\sigma_s^2 = I(\theta)\sigma_\theta^2. \tag{20}$$

For larger intervals, report

$$s(\theta) = \int^{\theta} \sqrt{I(u)}\, du. \tag{21}$$

The bridge predicts that visibility-loss breakdown t_{break} collapses best when the limiting tolerance is defined in these s-units.

9 Falsification Criteria

BLQC is weakened or falsified in the tested regime if:

1. C_{eff} variation at fixed geometry and fixed confounds produces no measurable change in t_{break} under adequate sensitivity.
2. h_{KS} variation at fixed geometry and fixed confounds produces no measurable change in t_{break} under adequate sensitivity.
3. Apparent capacity dependence is fully explained by temperature, readout SNR, latency, pulse distortion, actuator nonlinearity, or post-selection.
4. Mass-geometry variables explain the data and κ adds no predictive value.

Penrose OR is weakened in this setup if:

1. mass/separation variation fails to move the loss timescale under adequate sensitivity;
2. C_{eff} or h_{KS} moves the timescale at fixed mass geometry in a way not explained by standard experimental confounds.

The Fisher-capacity bridge is weakened or falsified if:

1. Apparent capacity dependence scales with $\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2$ at fixed mass geometry, but the operational records reveal that $I(\theta)$ is strongly position-dependent without a correspondingly local $\kappa(\theta)$, or the tracking capacity strictly reduces Euclidean coordinate error but not Fisher distinguishability error. In this event, BLQC is validated as a physical mechanism for observer-relative visibility loss, but the derivation of the binary Born rule via Cencov geometry is falsified.

10 Minimum Reporting Requirements

Any report of the experiment should include:

- platform description and mesoscopic mass/separation estimates;
- Penrose OR timescale estimate and uncertainty;
- operational definition of C_{eff};
- operational definition or estimate of h_{KS};
- full thermal/readout/latency/pulse/actuator diagnostics;
- visibility curves for all parameter conditions;
- model-comparison statistics for BLQC, Penrose OR, combined-rate, and nuisance models;
- empirical $I(\theta)$ and Fisher-homogeneity analysis;
- Fisher-unit reference-error reporting, using σ_s^2 or $s(\theta)$;
- estimates of $C_{\mathrm{eff}}^{\mathrm{emp}}$, $h_{\mathrm{KS}}^{\mathrm{emp}}$, and κ_{eff} for each mass-geometry condition;
- geometry-capacity covariance and mediation statistics;
- offline reference-log recovery analysis where available;
- recovery statistic R_{rec} when a shadow reference channel is implemented;
- pre-registered exclusion and confound criteria.

11 Conclusion

The Penrose-overlap experiment is the clean test of BLQC. It asks whether a mesoscopic visibility-loss timescale that could otherwise be attributed to gravitational objective reduction instead moves with finite-rate basis-reference tracking. If the timescale follows mass geometry and ignores C_{eff} and h_{KS}, Penrose OR or another geometry-dependent mechanism remains the better explanation. If the timescale follows $\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2$ at fixed mass geometry and after ordinary confounds are controlled, BLQC becomes a serious alternative or companion mechanism in the mesoscopic regime.

The Fisher-homogeneity module separately tests the Paper A bridge. A BLQC-positive but Fisher-negative result validates the finite-rate tracking mechanism while falsifying the proposed Born-rule bridge.

A QGEM Pathfinder Implementation Addendum

This addendum maps the protocol variables onto a QGEM-style single-mass mesoscopic spatial-superposition pathfinder. The purpose is not to privilege one platform, but to make the abstract variables experimentally concrete.

A.1 Physical Plant and Mass-Geometry Variables

In a QGEM-style pathfinder, an embedded spin, such as an NV center, is prepared in a superposition, and a magnetic field gradient entangles this spin with the center-of-mass spatial position of a levitated nanodiamond.

- **Mass m:** The mass of the levitated nanodiamond, with a representative target range of 10^{-15} to 10^{-14} kg.
- **Separation s:** The maximum spatial splitting between the center-of-mass branches, controlled by the amplitude and duration of the magnetic-gradient pulses $\partial_z B$.
- **Geometry sweep:** The Penrose OR timescale τ_{OR} is modified by varying the gradient-pulse current, changing s, or by using different nanodiamond masses, changing m.

A.2 Basis Reference and Tracking Variables

The measurement basis $\theta(t)$ is the physical microwave local-oscillator phase used to drive the final recombination and readout pulses of the spin-interferometry sequence.

- **Tracking channel:** The FPGA-based digital control loop that maintains the LO phase relative to the mechanical trap center and timing sequence.
- **Effective capacity C_{eff}:** Modulated directly within the FPGA by introducing an intentional, calibrated packet-drop or update-rejection rate on the digital feedback lines controlling the local-oscillator phase synthesizer, without altering raw clock speed or cryostat thermal environment.
- **Reference instability h_{KS}:** Controlled by injecting calibrated chaotic voltage noise into the local-oscillator phase modulation input. The entropy-rate proxy h_{KS} is extracted offline from the digital error-signal logs.

A.3 Order-of-Magnitude Capacity Estimate

The discrimination is only worth running if the throttled reference channel can be driven near the BLQC boundary $\kappa \gtrsim 0$. Counted as raw hardware it cannot: the full FPGA/DDS chain and master clock place C_{eff} many orders of magnitude above any plausible h_{KS}, giving $\kappa \ll 0$ (deep capacity-wins, $V \approx 1$, no overlap with τ_{OR}). This is the correct behaviour of an unthrottled chain and is not the regime under test.

The relevant C_{eff} is the useful rate that actually constrains θ online during the mesoscopic sequence, $C_{\mathrm{eff}} = r\,b\,f$, imposed here by the calibrated update-rejection schedule of the preceding subsection. Representative values:

Regime	r (Hz)	b	f	$C_{\text{eff}} \ln 2$	κ
Sparse online tracker	10	2	0.5	6.9	$+23$ to $+63$
Transitional	50	2	0.5	34.7	-4.7 to $+35$
Cleaner, throttled	100	2	0.5	69.3	-39 to $+0.7$
Unthrottled FPGA loop	1000	8	0.5	2773	$\ll 0$

$C_{\text{eff}} \ln 2$ in nats/s; $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ in s^{-1}, shown across a calibrated band $h_{\text{KS}} = 30$–70 nats/s.

A sparse-to-transitional tracker places $C_{\text{eff}} \ln 2$ in the 7–70 nats/s band. For a calibrated reference-instability rate h_{KS} of order tens of nats/s — imposed via the injected-noise path of the preceding subsection, not assumed as a native property of the controller — κ is transitional rather than deep-negative, and $t_{\text{break}} \sim 1/\kappa$ falls in the tens-of-ms window where τ_{OR} also lives for $m \sim 10^{-15}$ kg. Both κ knobs are set by design: C_{eff} by the imposed update-rejection schedule, h_{KS} by the injected chaotic phase modulation. The estimate is illustrative — b and f are order-of-magnitude choices — and the point is that the boundary is reachable by deliberate throttling across roughly a decade of r, not that any single row is the operating point. It is explicitly not a claim that unthrottled QGEM control electronics are intrinsically bandwidth-limited; it is the quantitative justification for sweeping C_{eff} (Section 5.2) rather than assuming the system sits at one end of the κ axis.

A.4 Fisher-Homogeneity Implementation

The NV-center spin readout provides the operational record family $p(o|\theta)$.

- **Calibration:** Sweep the intended phase θ_0 of the final microwave readout pulse from 0 to 2π. Map the resulting probability $p(|0\rangle \mid \theta)$ of detecting the NV center in the $|0\rangle$ state.
- **Fisher-homogeneity test:** Compute $I(\theta)$ from this curve and verify whether the empirical Fisher information is flat over the active tracking range.
- **Visibility loss:** Track the breakdown of the spin-readout contrast $V(t)$ as a function of free-evolution time.
- **Fisher-unit error reporting:** During the FPGA capacity sweeps, the unresolved phase error of the microwave local oscillator (σ_θ) must be converted into Fisher distinguishability units using the empirical Fisher information derived from the NV spin calibration:

  $$\sigma_s^2 = I(\theta)\, \sigma_\theta^2.$$

  The threshold for the breakdown time t_{break} must be evaluated against this s-unit tolerance.

A.5 Geometry–Capacity Covariance Challenge

The most significant experimental challenge in a QGEM-style pathfinder is the geometry-capacity covariance check. Increasing the magnetic gradient $\partial_z B$ to increase separation s typically requires higher coil current, which can induce thermal noise, mechanical vibration, or flux noise that couples back into the NV spin phase.

The protocol therefore requires verifying that increasing the spatial separation s does not implicitly increase the empirical reference instability $h_{\text{KS}}^{\text{emp}}$ of the microwave controller. If s strongly drives $h_{\text{KS}}^{\text{emp}}$, the result must be classified as mediated dual success rather than as a clean additive Penrose-plus-BLQC result.

A.6 Shadow Reference Channel

To distinguish observer-relative visibility loss from irreversible physical decoherence, the experiment requires a passive, high-bandwidth recording of the full reference dynamics.

The apparatus should log four synchronized streams:

- **Microwave LO phase:** $\phi_{\mathrm{LO}}(t)$ digitized through an IQ phase monitor referenced to the master clock.
- **FPGA control state:** DDS phase accumulator, accepted and rejected update packets, estimator state, and pulse trigger timestamps.
- **Plant state:** Gradient-drive waveform $\partial_z B(t)$, coil current $I_{\mathrm{coil}}(t)$, and trap-position proxy $x_{\mathrm{trap}}(t)$.
- **Shot-level outcomes:** NV readout outcomes timestamped against the master clock.

Crucial design rule. The online tracker receives only the throttled C_{eff} stream. The offline recorder receives the full-resolution streams but must be causally isolated from the feedback loop.

During offline analysis, reconstruct the realized basis phase for each shot k:

$$\theta_k^{\mathrm{real}} = \theta_k^{\mathrm{commanded}} + \delta\theta_k^{\mathrm{log}}. \tag{22}$$

Then compute the recovered visibility $V_{\mathrm{corr}}(t)$ using the reconstructed phases, and compare it to the raw online visibility $V_{\mathrm{raw}}(t)$ through the recovery statistic:

$$R_{\mathrm{rec}} = \frac{V_{\mathrm{corr}} - V_{\mathrm{raw}}}{V_{\mathrm{baseline}} - V_{\mathrm{raw}}}. \tag{23}$$

If $R_{\mathrm{rec}} \approx 1$, the loss is observer-relative reference ignorance. If $R_{\mathrm{rec}} \approx 0$, the loss is physical decoherence.

B The Combined-Rate Model in IOF Context

This appendix is strictly interpretive. It introduces no new measurements, alters no decision rules, and changes none of the falsification criteria specified in the body. Its purpose is to make the connection between the combined-rate regression of Section 7.2 and the broader Ignorant Observer Framework (IOF) corpus explicit, so that the experiment's outcomes can be read directly against the corpus's existing positions on the Heisenberg cut.

B.1 Non-Exclusivity and the Heisenberg Cut

The combined-rate model

$$\frac{1}{t_{\mathrm{break}}} = \alpha\,\frac{1}{\tau_{\mathrm{OR}}} + \beta\,\kappa + \gamma \tag{24}$$

is a joint fit, not a binary discriminator. The regression admits simultaneous nonzero α and β, and the protocol's decision rules (Section 7.4) explicitly classify such an outcome as "dual success," further partitioned into additive and mediated subcases by the geometry–capacity covariance analysis of Section 7.3.

The IOF reading of this structure is that the Heisenberg cut is not a single threshold owned by a single mechanism. IOF describes the *character* of the cut: an operational, observer-relative boundary set by what the apparatus can resolve, with $\kappa > 0$ the operational marker of crossing. Penrose Objective Reduction (OR), where active, is one of several physical decoherence channels that contribute to pushing a system across the cut. The two claims are not in direct metaphysical competition. They answer different questions: "what physical channels decohere heavy superpositions?" and "what makes a phenomenon count as having crossed the cut?".

B.2 Four Outcomes and Their IOF Readings

The regression admits four additive outcomes, plus a fifth degenerate case treated separately in Section B.4.

Outcome 1 — Pure BLQC ($\alpha \approx 0$, $\beta > 0$). The cut is operational and informational in this regime. Penrose OR, if real elsewhere, is empirically inactive at the masses tested. The corpus reading is straightforwardly IOF-positive.

Outcome 2 — Pure Penrose ($\alpha > 0$, $\beta \approx 0$). Mass-geometry sets the timescale and BLQC contributes nothing measurable in this regime. The corpus reading weakens BLQC in the Penrose-overlap band and supports a geometry-dependent mechanism. This is the principal BLQC falsification outcome listed in Section 9.

Outcome 3 — Additive Dual Success ($\alpha > 0$, $\beta > 0$, no mediation). Mass geometry and information capacity contribute as independent decoherence channels. Both are real; neither owns the cut singularly. The corpus reading absorbs OR as a parallel physical channel alongside the operational κ channel: IOF still describes the cut's character, OR contributes one of several channels that drive crossings.

Outcome 4 — Mediated Dual Success ($\alpha > 0$, $\beta > 0$, mass geometry mediates through empirical BLQC variables). The apparent gravitational signal is mediated by changes in $C_{\text{eff}}^{\text{emp}}$, $h_{\text{KS}}^{\text{emp}}$, $I(\theta)$, or κ_{eff}, and the independent Penrose term weakens once these are included. The corpus reading hands the cut cleanly to IOF: the apparent OR effect is an engineering correlation between mass and apparatus complexity, not a separate gravitational mechanism.

B.3 The Reversibility Axis (R_{rec})

The shadow-channel recovery statistic R_{rec} (Appendix A.5) runs orthogonal to the regression. It does not adjudicate which mechanism is bigger; it adjudicates the *character* of the loss attributed to either component.

- $R_{\text{rec}} \to 1$ on a component indicates observer-relative reference loss, consistent with BLQC/IOF and inconsistent with genuine OR collapse.
- $R_{\text{rec}} \to 0$ on a component indicates irreversible physical decoherence, consistent with OR or any other genuine collapse channel.

This second axis sharpens the IOF reading of Outcomes 3 and 4. An additive dual-success result with $R_{\text{rec}} \approx 1$ on the apparent Penrose component would be surprising, and would weaken the OR interpretation even when $\alpha > 0$: the mass-correlated signal would then be observer-relative in character, not genuinely irreversible.

B.4 The Collinearity Signature: Bridge Ansatz from *The Creation of Duality*

A fifth outcome is admitted by the regression but lies outside the additive partition above: a rank-deficient design matrix in which $1/\tau_{\mathrm{OR}}$ and κ are not independent regressors.

The Creation of Duality[1] proposes the **Bridge Ansatz**: a calibration that identifies the gravitational self-energy scale with the information-deficit rate via Margolus–Levitin saturation,

$$E_G = \frac{\pi}{2}\,\hbar\,\kappa_{\mathrm{info}}, \qquad \kappa_{\mathrm{info}} \equiv h_{\mathrm{KS}} - C_{\mathrm{eff}}\ln 2. \tag{25}$$

Under this Ansatz, the two timescales are proportional with the fixed ratio

$$\tau_{\mathrm{OR}} \;\approx\; \frac{2}{\pi}\,\tau_{\mathrm{loss}} \;\approx\; 0.637\,\tau_{\mathrm{loss}}, \tag{26}$$

and the combined-rate regressors collapse onto a single direction in design space: $1/\tau_{\mathrm{OR}} = (\pi/2)\,\kappa$.

The empirical signature is therefore not an additive fit at all, but a rank-1 design matrix. In regression diagnostics this appears as a large variance inflation factor and a near-singular condition number on $(1/\tau_{\mathrm{OR}}, \kappa)$. Only the combination $\alpha(\pi/2) + \beta$ is identifiable; α and β cannot be separately fit. The specific quantitative confirmation the Ansatz predicts is recovery of the ratio $\tau_{\mathrm{OR}}/\tau_{\mathrm{loss}} \approx 2/\pi$ within experimental tolerance.

The protocol does not assume the Ansatz; it provides the empirical conditions under which the Ansatz can be supported, weakened, or falsified.

B.5 Conditional Bridges to the Corpus

The protocol's outcomes map onto the corpus's positions on gravity–information interplay through three conditional statements.

Conditional 1 — Bridge Ansatz supported. If the design matrix on $(1/\tau_{\mathrm{OR}}, \kappa)$ is rank-deficient within experimental tolerance *and* the quantitative ratio $\tau_{\mathrm{OR}}/\tau_{\mathrm{loss}} \approx 2/\pi$ is recovered, the Bridge Ansatz of *The Creation of Duality* is supported. Gravity and information appear as two views of a single information-deficit budget, with the Margolus–Levitin coefficient as the calibration constant. This is the strongest unification reading available within the corpus.

Conditional 2 — Additive parallelism. If the regression yields cleanly identifiable $\alpha > 0$ and $\beta > 0$ with low collinearity and no significant mediation, the Bridge Ansatz is weakened. Gravity and information must then be absorbed as parallel decoherence channels operating in the same regime. IOF continues to describe the cut's character; OR is admitted as an independent physical contributor.

Conditional 3 — Mediated reading. If the apparent Penrose dependence is mediated by empirical BLQC variables (mass geometry pushes κ_{eff} rather than acting independently), the Bridge Ansatz is weakened from the opposite direction. The cut belongs operationally to IOF, the gravitational signal dissolves into apparatus engineering, and the corpus's deflationary reading of OR is supported.

[1] Dekker, A. (2026). *The Creation of Duality: A Speculative Extension on Appearance, Gravity, and Information.* Version 1.3. DOI: 10.17605/OSF.IO/C5EN8. The document self-identifies as the most speculative in the IOF corpus; its scientific weight is conditional on the experimental signatures discussed in this appendix.

Each of these conditionals is empirically reachable by the protocol as written. None requires modification of Sections 5–9 or of the falsification criteria. The Ansatz becomes a fifth named outcome alongside the four additive outcomes of Section B.2, with its own statistical signature (rank-deficiency, fixed ratio) and its own conditional consequence for the corpus.

B.6 Scope of this Appendix

This appendix does not change the protocol. Specifically:

- The primary discriminator (the derivative of t_{break} with respect to independently controlled variables, Section 2.3) is unchanged.
- The decision rules in Section 7.4 are unchanged. The Bridge Ansatz signature is supplementary: a rank-deficiency check is added to the design-matrix diagnostics, but the additive outcomes are still classified by the rules already specified.
- The falsification criteria in Section 9 are unchanged. The Ansatz is not a new falsifiable hypothesis the protocol is committed to testing; it is a corpus claim whose statistical signature happens to lie within the same experiment.
- No new measurements, platform requirements, or confound controls are introduced.

The appendix is a reading guide between the protocol's outputs and the IOF corpus's foundational claims. The protocol stands on its own for an experimentalist who does not engage with the corpus; for one who does, this appendix names the corpus's positions and shows which experimental outcomes support, weaken, or falsify each.