

A Capacity-Calibrated Protocol for Testing Penrose Objective Reduction

A mesoscopic visibility-loss design with a benchmarked finite-rate reference channel

Aernoud Dekker

June 2026

DOI: 10.17605/OSF.IO/2QJNE

Abstract

This document specifies a layered prospective test of Penrose Objective Reduction (OR) in the mesoscopic regime. The primary observable is a capacity-independent, *unrecoverable* visibility floor that follows mass geometry according to a Penrose-type estimate, measured with the engineered reference channel off and tracking capacity maximal.

The design's distinctive feature is its reference channel. The physical basis reference $\theta(t)$ is maintained by a finite-rate tracking loop whose control variables are set by design— C_{eff} by a calibrated update-rejection schedule, h_{KS} by injected, certified-expanding reference modulation—and whose behaviour is benchmarked in advance against the calibrated deficit

$$\kappa_{\text{cal}} = h_{\text{KS}} - \eta C_{\text{eff}} \ln 2,$$

with the coding efficiency η measured once in an electronics-only calibration arm and frozen, following the companion Bandwidth-Limited Quantum Control (BLQC) benchmark. Because reference physics is measured and imposed rather than estimated after the fact, it ceases to be a nuisance background and becomes the experiment's calibration arm. The same machinery supplies the classifier that separates outcomes: loss that is recoverable from a passive shadow reference $\log(R_{\text{rec}} \approx 1)$ is observer-relative reference bookkeeping within standard quantum mechanics; a floor that is unrecoverable ($R_{\text{rec}} \approx 0$), follows mass geometry, persists at maximal tracking capacity, and exceeds the logging-fidelity budget is the OR signature.

The protocol defines the required platform, experimental arms, confound controls, statistical models, tiered decision rules, and sensitivity requirements, and a QGEM-style addendum maps the abstract variables onto one candidate implementation. The operational benchmark of the reference channel itself—and the registered strong-tier (chaotic-corner) test that rides on it—are carried by the companion BLQC paper; this protocol's decision space lies along the objective-reduction axis.

Scope of this protocol. This protocol specifies a prospective OR-discrimination experiment. It presupposes the finite-rate reference-channel benchmark specified in the companion BLQC paper [1]—the electronics-only calibration arm (coding efficiency η), the expanding-dynamics gate, the dynamic-range budget, the capacity and instability sweeps, the shadow-channel recovery statistic R_{rec} , and the noisy-reference null model V_{null} —which must be passed on the target platform's reference channel before the arms below are run. The benchmark's outputs—the frozen coding efficiency η , the expanding-dynamics certificates, the dynamic-range budget, and the logging-fidelity budget—enter this protocol as fixed calibration inputs. The reference-channel hypothesis frame is restated in one paragraph in Section 1.2; its full statement, decision rules, sensitivity requirements, and the registered strong-tier test live in BLQC.

Contents

1	Scope and Logical Status of the Experiment	4
1.1	The Object of the Test	4
1.2	The Reference-Channel Hypotheses in Brief	4
1.3	Why a Calibrated Reference Channel	5
1.4	Relation to Existing Discrimination Methodologies	5
1.5	The Two Axes	5
1.6	Relation to the IOF Corpus	6
1.7	Beyond Penrose: The Architecture Is Rival-Agnostic	6
2	The Experimental Question	6
3	Competing Predictions	7
3.1	Penrose Objective Reduction	7
3.2	The Engineered Reference Channel	7
3.3	Discrimination Table and Its Limits	8
4	Required Platform	8
5	Operational Variables	9
5.1	Mass-Geometry Variables	9
5.2	The Calibrated Reference Channel (Imported)	10
6	Experimental Arms	10
6.1	Arm 0: Reference-Channel Benchmark (Prerequisite)	10
6.2	Arm 1: κ -Off Geometry Sweep (Primary OR Test)	11
6.3	Arm 2: Combined Crossed Sweep	11
6.4	Randomization and Replication	11
7	Confound Controls	11
7.1	Recoverability as Classifier	12
8	Statistical Analysis	12
8.1	Primary Observables	12
8.2	Models	13

8.3	Geometry–Capacity Covariance and Mediation Check	13
8.4	Decision Rules	13
9	Sensitivity and Power Requirements	14
9.1	Error Model	15
9.2	Pre-Registered Requirements	15
10	Falsification Criteria	15
11	Staged Feasibility Path	16
12	Minimum Reporting Requirements	16
13	Conclusion	17
A	QGEM Pathfinder Implementation Addendum (Stage 2)	17
A.1	Physical Plant and Mass-Geometry Variables	18
A.2	Basis Reference and Tracking Variables	18
A.3	Order-of-Magnitude Capacity Estimate	18
A.4	Geometry–Capacity Covariance Challenge	19
A.5	Shadow Reference Channel	19
B	The Combined-Rate Model in IOF Context	20
B.1	Non-Exclusivity and the Heisenberg Cut	20
B.2	Four Outcomes and Their IOF Readings	21
B.3	The Reversibility Axis (R_{rec})	21
B.4	The Collinearity Signature	21
B.5	Scope of this Appendix	22
	References	23

1 Scope and Logical Status of the Experiment

1.1 The Object of the Test

Penrose OR predicts a collapse-like timescale set by gravitational self-energy [2]:

$$\tau_{\text{OR}} \approx \frac{\hbar}{E_G}, \quad (1)$$

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 ,

$$\tau_{\text{OR}} \propto \frac{\hbar s}{Gm^2}, \quad (2)$$

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

The hypothesis under test is:

H_{OR} — **Objective reduction.** A κ -independent, unrecoverable ($R_{\text{rec}} \approx 0$, Section 7.1) visibility floor follows mass geometry according to a Penrose-type estimate.

E_G carries a well-known regularization ambiguity: the point-particle and nuclear-smearred mass-density prescriptions differ by orders of magnitude in τ_{OR} . The protocol therefore treats the two prescriptions as separate model variants in the statistical analysis (Section 8.2); the geometry *scaling* across the sweep, which is more robust than the absolute timescale, is the primary OR-model content.

Penrose OR predicts that the loss timescale should change when mass, separation, or mass distribution changes, and that the OR mechanism itself has no dependence on the classical controller’s basis-tracking capacity.

1.2 The Reference-Channel Hypotheses in Brief

The reference channel that calibrates this experiment carries its own hypothesis frame, stated and adjudicated in BLQC. In one paragraph: H_{null} (standard quantum mechanics with a noisy finite-bandwidth reference) holds that the full-resolution shadow log predicts the observed visibility loss by ordinary phase averaging; H_{op} (operational BLQC) holds that the same loss, measured online by the throttled observer, collapses against the calibrated deficit $\kappa_{\text{cal}} = h_{\text{KS}} - \eta C_{\text{eff}} \ln 2$ —a control-theoretic refinement of H_{null} , consistent with it by construction, since the data-rate theorem grounding the law is classical control theory; and H_{strong} (a residual, κ -scaled deficit below the shadow-log prediction) would be new physics—a reading raised and excluded as analysis in Part III of the foundational IOF paper [3] against existing recoverability experiments (logged-settings Bell tests [4], randomized-measurement and classical-shadow protocols [5, 6], and correlation spectroscopy [7, 8]), and surviving only as a registered test for certified deterministic-chaotic reference dynamics, with the prior strongly against. BLQC carries the full statements, the benchmark decision rules (κ_{cal} collapse, V_{null} agreement, $R_{\text{rec}} \approx 1$), and the registered chaotic-corner module. For the purposes of this protocol, the reference channel is settled physics: a calibrated, recoverable, observer-relative loss mechanism that standard quantum mechanics fully describes.

1.3 Why a Calibrated Reference Channel

A mesoscopic OR search must separate a candidate gravitational floor from the reference physics of its own apparatus: phase noise, finite tracking bandwidth, and partially corrected drift all reduce visibility, and at OR-relevant sensitivities they are not negligible. The usual treatment estimates these contributions as a nuisance budget. This protocol instead *engineers* them: the reference channel’s deficit is imposed and calibrated in advance, its loss is verified to collapse against κ_{cal} , and—decisively—its loss is verified to be *recoverable* from the passive shadow log. Reference physics thereby acquires a signature that distinguishes it from any collapse-type mechanism, on the same records: recoverability. The OR observable is defined against that calibrated background, and the technical-noise confound that would otherwise dominate the error budget is closed by measurement rather than by estimate.

1.4 Relation to Existing Discrimination Methodologies

Existing strategies for separating collapse-type loss from ordinary decoherence fall into two families. The first builds a calibrated environmental noise budget—gas collisions, blackbody radiation, photon recoil—and infers a collapse contribution as statistical excess over that baseline [9, 10]. The second looks for functional signatures distinctive of the collapse mechanism itself: mass scaling and size saturation of the excess decoherence rate, frequency-dependent collapse heating, or predicted spontaneous radiation [11, 12].

This protocol is designed to compose with those strategies, not replace them; what it adds is a third, complementary element: a *parameter-free conditioning classifier*. Instead of asking whether the observed loss exceeds a modeled budget, it asks whether the loss survives conditioning on a passive, full-resolution record of the realized reference—a question answered by the data themselves, with no noise model fit. None of the ingredients is new physics, and most are established practice. The recoverable/irreversible distinction is as old as the spin echo [13], and echo techniques have long been used to identify the source of a dephasing channel [14]; the recovery statistic R_{rec} is the conditioning-based analogue—where an echo physically refocuses reversible phase spread, conditioning analytically restores contrast lost to a logged reference, which applies equally when the loss cannot be re-evolved. Likewise, conditioning quantum records on passive witness logs is standard metrology: logged-settings Bell tests [4], seismometer-corrected atom interferometry [15], and correlation spectroscopy [7, 8]. The contribution of this protocol is composition and pre-registration: the reference channel is engineered and calibrated to a data-rate law *in advance* (Section 5.2), and recoverability against the passive log is promoted from noise hygiene to the registered classifier, bounded by a quantified logging-fidelity budget (Section 7.1). The claim is methodological discipline, not new physics.

1.5 The Two Axes

The central structural fact about this design must be stated openly: **the design discriminates along the OR axis but is blind along the BLQC-versus-null axis.**

Both κ knobs are set by design. C_{eff} is imposed by a calibrated update-rejection schedule on the tracking loop; h_{KS} is imposed by injected, certified-expanding reference modulation (certified under BLQC’s expanding-dynamics gate). Given that construction, H_{null} and H_{op} predict identical raw visibility, identical recoverability, and identical κ -scaling in every arm: throttling a loop that corrects an injected expanding disturbance increases residual reference error, and standard quantum mechanics converts residual reference error into fringe-visibility loss. No sweep of C_{eff}

or h_{KS} , however well confounds are controlled, can separate the two.

What the design *can* discriminate is the OR axis: whether a κ -independent, unrecoverable mass-geometry floor exists when the engineered channel is off (Arm 1, Section 6.2). Residuals against the noisy-reference null—the strong axis—are adjudicated under BLQC’s registered strong-tier rules, not under this protocol’s decision space.

Accordingly, the protocol does not claim that κ -dependence in the Penrose-overlap regime would make BLQC a serious physical alternative to standard quantum mechanics. κ -dependence validates the reference-channel calibration; the physically decisive observable here is the κ -off, unrecoverable floor.

1.6 Relation to the IOF Corpus

The IOF corpus’s keystone claim—that wavefunction collapse is finite-rate reference bookkeeping by an embedded observer—is interpretive: it is consistent with standard quantum mechanics by construction, and the corpus itself predicts that the engineered loss is recoverable ($R_{\text{rec}} \approx 1$). The corpus is therefore supported by *demonstration and consistency* (the operational tier, benchmarked in BLQC), not by discrimination against quantum mechanics. The live discriminating physics in this design belongs to Penrose OR: either outcome of the geometry sweep is informative about gravitational objective reduction, and neither outcome discriminates IOF from quantum mechanics. The corpus’s methodological contribution is the calibrated reference channel and the recoverability classifier, which together close the technical-noise confound on the OR axis.

1.7 Beyond Penrose: The Architecture Is Rival-Agnostic

This protocol is formulated against the Penrose schedule, but nothing in its architecture is specific to gravitational collapse. The conditioning classifier makes no reference to what causes a putative unrecoverable residual; it certifies only that the residual does not belong to the observer’s reference channel. Every model in the dynamical-collapse family predicts such a residual, and what distinguishes the rivals is the scaling variable of whatever floor survives the subtraction: mass and superposition geometry for gravitational reduction (the present target), amplification with nucleon number at a universal rate and length scale for spontaneous-localization (GRW/CSL) models, and noise–decoherence trade-offs for classical-channel gravity [16]. Retargeting the protocol therefore means replacing the predicted schedule in the geometry sweep (Arm 1, Section 6.2) and its decision rules; the platform, the calibrated reference channel, the recoverability classifier, and the audit structure carry over unchanged. Penrose OR is the flagship target because its mass–geometry schedule lands in an experimentally accessible window on this platform and carries a live public debate, not because the method is bespoke to it. The scope boundary of Section 1.4 still applies: the classifier governs interferometric-visibility loss, and collapse models are bounded in parallel through non-interferometric channels—spontaneous radiation, collapse heating—where recoverability is not the observable.

2 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 an engineered finite-rate tracking loop can be tuned to produce loss on comparable timescales.

The protocol splits the experimental question in two, matching the two axes of Section 1.5:

1. **OR question (primary physics question).** With the engineered channel off and tracking capacity maximal, does a residual, unrecoverable visibility floor follow mass geometry as Penrose OR predicts?
2. **Calibration question (control-physics question).** With mass geometry fixed, does the online observer-relative loss collapse against the calibrated deficit κ_{cal} , and is it fully recoverable from the shadow log? This is the BLQC benchmark, run on the target platform’s reference channel as a prerequisite (Arm 0, Section 6.1).

The same apparatus must expose both sets of variables: the mass-geometry variables entering the Penrose estimate and the basis-reference variables entering the calibration. But the two questions are answered by different arms, and conflating them would degrade both.

3 Competing Predictions

3.1 Penrose Objective Reduction

The OR prediction is stated in Section 1.1: a loss timescale $\tau_{\text{OR}} \approx \hbar/E_G$ that follows mass, separation, and mass distribution, evaluated under both E_G regularization prescriptions, with no dependence on the classical controller’s basis-tracking capacity. The empirically sharp form of the gravitational-collapse idea is the Diósi–Penrose model, whose parameter-free version is already excluded by underground searches for the predicted spontaneous radiation [12]; the surviving parameter space makes the mass-geometry *scaling* of any observed floor, rather than its absolute timescale, the primary model content (Section 8.2).

3.2 The Engineered Reference Channel

The reference channel 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} delivered with coding efficiency η , the calibrated deficit is

$$\kappa_{\text{cal}} = h_{\text{KS}} - \eta C_{\text{eff}} \ln 2. \quad (3)$$

In the chaos-wins regime ($\kappa_{\text{cal}} > 0$), unresolved basis uncertainty grows as $\sigma_{\theta}^2(t) = \sigma_0^2 e^{2\kappa_{\text{cal}} t}$ within the dynamic-range window, Gaussian basis uncertainty reduces visibility as $V(t) = \exp[-\frac{1}{2}\sigma_0^2 e^{2\kappa_{\text{cal}} t}]$, and for a chosen threshold V_* the breakdown time is

$$t_{\text{break}} = \frac{1}{2\kappa_{\text{cal}}} \ln \left(\frac{-2 \ln V_*}{\sigma_0^2} \right). \quad (4)$$

The derivation, validity conditions, operational estimators, and benchmark procedure for this law are given in BLQC; this protocol imports the law and its frozen calibration. At fixed mass geometry and fixed environmental controls, increasing useful C_{eff} delays online loss and increasing h_{KS} accelerates it—predictions shared with H_{null} (Section 1.5), and therefore calibration content, not discrimination content.

3.3 Discrimination Table and Its Limits

Controlled change	Penrose OR (mechanism)	Reference channel (mechanism)
Increase mass at fixed C_{eff} , h_{KS} , and separation	Faster loss	No leading change unless reference dynamics also change
Increase separation at fixed C_{eff} , h_{KS} , and mass	Geometry-dependent OR shift	No leading 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 table states *mechanism-level* dependencies. At the level of observed data the engineered κ channel and any OR channel add as rates,

$$\frac{1}{t_{\text{break}}^{\text{obs}}} \approx \frac{1}{\tau_{\text{OR}}} + \Gamma_{\kappa} + \Gamma_{\text{env}}, \quad (5)$$

so whenever the engineered channel is active, a capacity-dependent shift in $t_{\text{break}}^{\text{obs}}$ is predicted by *both* hypotheses. The derivative of t_{break} with respect to C_{eff} , taken alone, therefore discriminates nothing.

The discriminating observable of this protocol is the residual, κ -off, unrecoverable mass-geometry floor (Arm 1, Section 6.2). Recoverability against the passive shadow log is the classifier separating reference physics from collapse-type loss (Section 7.1).

4 Required Platform

A suitable platform must be mesoscopic enough for Penrose OR estimates to be meaningful and instrumented enough for the reference-channel 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. The staged feasibility path (Section 11) records which requirements are executable on near-term platforms and which await a future mesoscopic platform.

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} , with expanding-dynamics certification per BLQC.
- Vary useful basis-tracking capacity C_{eff} without changing mass geometry.
- Support an electronics-only calibration of the tracking loop (BLQC's Arm 0): the loop, throttle, injection path, and a reference-grade phase monitor must be operable with the quantum system removed or idle.
- Suppress the engineered channel on demand: in the κ -off configuration, the predicted residual engineered loss rate must fall below a pre-registered fraction of the target OR rate (Section 9).
- Keep temperature, readout SNR, pulse/actuator behavior, sequence duration, and environmental coupling under monitored control.
- Preserve high-resolution reference logs (shadow channel) with auditable fidelity, 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 engineered channel cannot be suppressed well enough to expose the κ -off floor (Arm 1 impossible).
- The visibility-loss timescale is dominated by ordinary environmental decoherence with no usable dynamic range for either OR or reference-channel variables.

5 Operational Variables

5.1 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 , reported under both regularization prescriptions.

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

5.2 The Calibrated Reference Channel (Imported)

The reference-channel variables enter this protocol as calibrated imports from the BLQC benchmark:

- **Effective capacity** $C_{\text{eff}} = r b f$, where r is the fraction of updates that survive loss, rejection, latency, filtering, and estimator overhead, b is the useful bits per update constraining the reference, and f is the useful update rate. 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, imposed here by a calibrated update-rejection schedule. Artificial delay may be used only with latency-matched controls, because extra waiting time produces ordinary coherence loss.
- **Coding efficiency** η , measured once in the electronics-only calibration arm and frozen before any quantum data are taken (BLQC’s “no free capacity scale” rule). It may not be refit per condition, per arm, or per platform run; inferring C_{eff} or η from the quantum-visibility data under test is prohibited.
- **Reference instability** h_{KS} , imposed by injected reference modulation and certified as intrinsically expanding—a positive largest Lyapunov exponent of the realized trajectory, with surrogate-data null tests rejecting stochastic diffusion—under BLQC’s expanding-dynamics gate. Gate-failure accounting (every attempted configuration reported) is part of the benchmark record.
- **Dynamic-range budget**: the exponential law holds only over a finite window bounded by Gaussian validity ($\sigma_\theta \lesssim 1$ rad), phase wrapping ($\sigma_\theta \approx \pi$), and attractor saturation; the pre-registered budget (minimum e-folds inside the validity window, visibility threshold V_* chosen accordingly) is fixed in the benchmark and inherited here.

The full definitions, estimators, calibration rules, and pre-registration conditions for these quantities are specified in BLQC and are not duplicated here.

6 Experimental Arms

The experiment is organized as three arms, run in order.

6.1 Arm 0: Reference-Channel Benchmark (Prerequisite)

Before any OR data are taken, the target platform’s reference channel must pass the BLQC benchmark: the electronics-only calibration (coding efficiency η measured and frozen; expanding-dynamics certificates issued; dynamic-range budget registered; logging-fidelity budget quantified by the injected-test-signal audit), followed by the capacity and instability sweeps on the platform’s reference channel, with breakdown times collapsing against κ_{cal} , raw visibility matching the parameter-free noisy-reference null V_{null} within the error budget, and full recoverability from the shadow log ($R_{\text{rec}} \approx 1$).

Classical control theory predicts that the reference channel passes; that is expected, and it is the point: **Arm 0 is calibration, not a test of new physics.** Its outputs are frozen before any OR data are taken, and they fix the prediction machinery—the same logs and the same variance

estimator—against which the engineered channel’s contribution is computed in the arms below. A platform whose reference channel fails the benchmark has a calibration or apparatus problem and is not ready for Arm 1; benchmark failure modes and their interpretation are specified in BLQC.

6.2 Arm 1: κ -Off Geometry Sweep (Primary OR Test)

With injection off, tracking capacity maximal (unthrottled), and the engineered channel’s predicted residual rate verified below the pre-registered suppression bound (Section 9), sweep mass, separation, or geometry at fixed C_{eff} , temperature, readout SNR, pulse/actuator behavior, and sequence duration.

The primary OR observable is a residual visibility floor in this arm that:

- follows mass/separation/geometry according to an OR-type estimate (either E_G regularization);
- is independent of C_{eff} and h_{KS} within sensitivity; and
- is unrecoverable from the shadow log ($R_{\text{rec}} \approx 0$).

This arm is the clean OR test. Embedding the OR comparison inside the engineered- κ regime would deliberately tune a comparable additive rate into the same time window and thereby degrade OR sensitivity; the suppressed-channel configuration is therefore the primary OR arm.

6.3 Arm 2: Combined Crossed Sweep

Run a crossed or partially crossed design over mass-geometry and engineered- κ variables, to test additive versus mediated structure (Section 8.3): whether mass-geometry changes also move the empirical tracking variables ($C_{\text{eff}}^{\text{emp}}$, $h_{\text{KS}}^{\text{emp}}$, κ_{eff}), and whether an OR-type term retains independent predictive value once they are included.

6.4 Randomization and Replication

Run order should be randomized across capacity, instability, and geometry settings within each arm. Each condition must include enough repetitions to meet the sensitivity requirements of Section 9. Replication should include at least one repeated condition after the full parameter sweep to detect slow drift; the drift bound is part of the sensitivity budget.

7 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 geometry or capacity effect that tracks readout degradation is not evidence for any hypothesis.

Latency and wait-time control. Sequence duration and idle time must be matched across conditions. If a configuration 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. An effect that tracks actuator distortion is confounded. In particular, the gradient drives that implement the geometry sweep must be audited for back-coupling into the reference channel (the geometry–capacity covariance check, Section 8.3 and Appendix A).

7.1 Recoverability as Classifier

The realized reference must be logged at higher resolution than the online tracker receives, by 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. The shadow-channel design rules, the per-shot reconstruction, and the recovery statistic

$$R_{\text{rec}} = \frac{V_{\text{corr}} - V_{\text{raw}}}{V_{\text{baseline}} - V_{\text{raw}}} \quad (6)$$

(recovered minus raw visibility, normalized to the baseline) are specified in BLQC; the QGEM addendum (Appendix A) maps the required log streams onto one platform.

R_{rec} is this protocol’s classifier. On the engineered channel, $R_{\text{rec}} \approx 1$ is the Arm-0 success condition: it certifies that the calibrated loss is online observer-relative reference ignorance, as standard quantum mechanics predicts. On an Arm-1 floor component, $R_{\text{rec}} \approx 0$ is the OR-positive condition: it certifies that the loss is not reference bookkeeping.

Logging-fidelity budget. The classification is interpretable only to the fidelity of the shadow log, in both directions: unlogged classical reference noise depresses R_{rec} and thereby *mimics an unrecoverable floor*. The injected-test-signal audit of the BLQC benchmark quantifies the log’s error budget; an Arm-1 floor whose unrecoverable component is smaller than the logging-fidelity budget is not a result (requirement R3, Section 9).

8 Statistical Analysis

8.1 Primary Observables

Three observables carry the protocol’s claims:

- the κ -off floor rate in Arm 1, as a function of mass geometry;
- the recovery statistic R_{rec} on the floor component;
- the breakdown time t_{break} per condition, defined by the pre-registered visibility threshold $V(t_{\text{break}}) = V_*$ (with V_* constrained by the imported dynamic-range budget, Section 5.2).

Secondary observables include fitted decay-rate parameters, recovered visibility after offline reference correction, and model residuals across the parameter grid. Residuals of raw visibility against the noisy-reference null V_{null} are computed throughout (the machinery is already in place from Arm 0) and analyzed under BLQC’s registered strong-tier rules; they do not enter this protocol’s decision space.

8.2 Models

Fit at least the following model classes:

Noisy-reference null: the parameter-free $V_{\text{null}}(t)$ computed per condition from the shadow log with the Arm-0-frozen estimator (specified in BLQC); the reference model for all reference-channel effects.

Calibrated reference-channel model:

$$t_{\text{break}} = \frac{A}{h_{\text{KS}} - \eta C_{\text{eff}} \ln 2} + B, \quad (7)$$

or the threshold-derived logarithmic form (Eq. 4) when σ_0 is estimated, with η frozen from Arm 0.

Penrose OR model:

$$t_{\text{break}} = A \tau_{\text{OR}}(m, s, \text{geometry}) + B, \quad (8)$$

fit separately under the point-particle and nuclear-smearred E_G prescriptions.

Combined-rate model:

$$\frac{1}{t_{\text{break}}} = \alpha \frac{1}{\tau_{\text{OR}}} + \beta \kappa_{\text{cal}} + \gamma, \quad (9)$$

allowing both mechanisms or mechanism-like rates to contribute.

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

8.3 Geometry–Capacity Covariance and Mediation Check

In addition to fitting the additive combined-rate model (Eq. 9), the analysis must test whether changes in mass geometry alter the empirical reference-channel variables themselves.

For each mass-geometry condition, estimate:

$$C_{\text{eff}}^{\text{emp}}, \quad h_{\text{KS}}^{\text{emp}}, \quad \kappa_{\text{eff}} = h_{\text{KS}}^{\text{emp}} - \eta C_{\text{eff}}^{\text{emp}} \ln 2. \quad (10)$$

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

$$\frac{1}{\tau_{\text{OR}}} \rightarrow C_{\text{eff}}^{\text{emp}}, \quad h_{\text{KS}}^{\text{emp}}, \quad \kappa_{\text{eff}}. \quad (11)$$

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} or tracking capacity, the result supports a unified geometry-capacity interpretation.

8.4 Decision Rules

The decision rules are tiered. Tier 1 (the reference-channel benchmark) and Tier 4 (the registered strong-tier residual) are adjudicated under BLQC’s rules; this protocol’s decision space is Tiers 2

and 3, plus the combined outcomes. Each tier has its own reading of the recovery statistic R_{rec} (Section 7.1).

Tier 1 — Reference-channel benchmark passed (prerequisite). Adjudicated in Arm 0 under BLQC’s decision rules: κ_{cal} collapse with frozen η , raw visibility matching the parameter-free V_{null} within the error budget, and $R_{\text{rec}} \approx 1$. Tier-1 success validates the calibration and is explicitly not evidence beyond H_{null} (Section 1.5).

Tier 2 — OR-positive. Requires all of:

- In the κ -off/high-capacity arm (Arm 1), a residual visibility floor follows mass/separation/geometry according to an OR-type estimate (either E_G regularization).
- The floor is independent of C_{eff} and h_{KS} within the experiment’s sensitivity after ordinary confounds are controlled.
- The floor is unrecoverable: $R_{\text{rec}} \approx 0$ on the floor component, with the unrecoverable component exceeding the logging-fidelity budget by the pre-registered margin (requirement R3). **Here unrecoverability is the success condition:** it certifies the loss is not online reference ignorance.

Tier 3 — No-OR / reference-only. The geometry sweep produces no residual floor at the pre-registered sensitivity, and all κ -on loss is fully recoverable. This weakens Penrose OR in the tested regime and leaves the calibrated reference channel as the complete account of the observed loss.

Tier 4 — Strong-tier residual (registered in BLQC, not asserted). A statistically significant deficit of V_{raw} below V_{null} , surviving the logging-fidelity audit and scaling with κ_{cal} , is analyzed under BLQC’s registered chaotic-corner rules. It is recorded here only because the shadow-channel design produces the relevant residuals at no extra cost; it does not alter the OR decision space.

Combined outcomes. If both mass geometry and κ_{cal} improve prediction in Arm 2, the result should be modeled as overlapping rates rather than assigned to either mechanism by inspection, and must be classified as either additive or mediated:

- **Additive dual success:** Both $1/\tau_{\text{OR}}$ and κ_{cal} independently improve prediction of $1/t_{\text{break}}$, with no significant covariance between mass geometry and empirical reference-channel variables. The κ -on contribution remains recoverable reference loss riding on top of any Tier-2 floor; it does not by itself compete with OR.
- **Mediated dual success:** Mass geometry significantly changes $C_{\text{eff}}^{\text{emp}}$, $h_{\text{KS}}^{\text{emp}}$, or κ_{eff} , and the independent Penrose term weakens after those variables are included.

9 Sensitivity and Power Requirements

“Adequate sensitivity” appears throughout the falsification criteria; this section defines it. All numerical targets below are defaults to be instantiated per platform at pre-registration; the structural requirement is that each target be fixed, with its power calculation, before data collection.

9.1 Error Model

Per visibility estimate from N shots at fringe visibility V , the shot-noise standard error is

$$\sigma_V \approx \sqrt{\frac{1 - V^2}{N}}. \quad (12)$$

The propagation of σ_V into breakdown-time uncertainty, and the worked precision examples for the reference-channel sweeps, are given with the benchmark in BLQC.

9.2 Pre-Registered Requirements

R1 — Benchmark sensitivity (Tier 1). The Arm-0 benchmark must meet BLQC’s registered sensitivity requirement for the capacity sweep (default: the smallest relative change in t_{break} detectable at 90% power across the swept range is 20%) on the target platform’s reference channel.

R2 — OR-floor sensitivity (Tier 2). In the κ -off arm, the standard error of the fitted floor rate per geometry condition must satisfy

$$\sigma(\Gamma_{\text{floor}}) < \frac{1}{3} \Delta \left(\frac{1}{\tau_{\text{OR}}} \right) \quad (13)$$

across the geometry sweep, evaluated under both E_G regularizations (3σ resolution of the predicted geometry-driven shift). In the same arm, the predicted residual engineered-channel rate (from the Arm-0 calibration) must satisfy

$$\Gamma_{\kappa}^{\text{resid}} < 0.1 \times \min_{\text{sweep}} \frac{1}{\tau_{\text{OR}}}. \quad (14)$$

R3 — Floor-classification budget (Tier 2). Quantify the total systematic error on the recoverability classification: the logging-fidelity budget (from the injected-test-signal audit of the Arm-0 benchmark), calibration drift, and estimator bias. The unrecoverable floor component $1 - R_{\text{rec}}$ claimed under Tier 2 must exceed this budget by a factor of at least 3. (The same budget bounds any Tier-4 claim; that requirement lives with the registered module in BLQC.)

R4 — Drift bound. The repeated-condition drift check (Section 6) must bound slow drift over the full campaign to less than one third of the smallest registered minimum detectable effect.

A verdict under any tier is valid only if the corresponding requirement was met; a null result obtained below these sensitivities is an unmet precondition, not a falsification.

10 Falsification Criteria

The criteria apply only when the corresponding sensitivity requirement of Section 9 was met.

Penrose OR is weakened in this setup if:

1. mass/separation variation in the κ -off arm fails to produce a residual floor (requirement R2 met), under both E_G regularizations;
2. any observed floor is recoverable from the shadow log ($R_{\text{rec}} \approx 1$), classifying it as reference bookkeeping rather than collapse.

The protocol’s preconditions fail (no OR verdict possible) if:

1. the reference channel fails the Arm-0 benchmark—a calibration or apparatus problem under BLQC’s failure analysis, to be investigated as such before any OR data are interpreted;
2. the engineered channel cannot be suppressed below the R2 bound in the κ -off configuration;
3. the logging-fidelity budget is too large to support the R3 classification margin at the target floor sensitivity.

Reference-channel criteria. The operational law’s own falsification criteria (capacity dependence, data collapse with frozen η , chaos dependence, functional form), and the registered strong-tier criteria, are specified in BLQC.

11 Staged Feasibility Path

The full Penrose-overlap experiment requires a platform that does not yet exist. The protocol is honest about this and registers a staged path, so that decision rules predate data at every stage.

Stages 0–1 — The BLQC benchmark (executable now and near-term). The electronics-only calibration and the standalone reference-channel demonstrator on an existing qubit or interferometric platform are specified and adjudicated in BLQC. No mesoscopic mass is required. Their deliverables—frozen η , expanding-dynamics certification practice, dynamic-range budgeting, the V_{null} machinery, demonstrated logging fidelity, and a first bound on the strong tier—are the calibration inputs this protocol presupposes.

Stage 2 — Penrose-overlap experiment (future platform; this protocol). Arms 0–2 on a QGEM-class mesoscopic platform (Appendix A). The required masses (10^{-15} – 10^{-14} kg) in spatial superposition at 100 nm–1 μ m splitting, held for tens to hundreds of milliseconds with environmental decoherence subdominant, exceed the demonstrated state of the art by many orders of magnitude (large-molecule interference has reached masses of order 10^{-22} kg; levitated-nanoparticle ground-state cooling is a stepping stone, not a spatial superposition). Stage 2 is contingent on that platform class maturing; the protocol’s role now is to fix the decision rules in advance.

A report from any stage must identify its stage and claim only that stage’s tiers.

12 Minimum Reporting Requirements

Any report of the experiment should include:

- the stage (Section 11) and the tiers claimed;
- platform description and mesoscopic mass/separation estimates;
- Penrose OR timescale estimates and uncertainties under both E_G regularizations;
- the Arm-0 benchmark report per BLQC’s minimum reporting requirements: operational definition of C_{eff} , the frozen η (value, method, uncertainty), expanding-dynamics certificates and the gate-failure rate, the dynamic-range budget and observed saturation onset, the κ_{cal} collapse, the V_{null} agreement, and the logging-fidelity audit;

- full thermal/readout/latency/pulse/actuator diagnostics;
- visibility curves for all parameter conditions, with the parameter-free V_{null} overlay and residuals;
- the fitted κ -off floor rate per geometry condition, under both E_G regularizations;
- model-comparison statistics for the noisy-reference null, calibrated reference-channel, Penrose OR, combined-rate, and nuisance models;
- estimates of $C_{\text{eff}}^{\text{emp}}$, $h_{\text{KS}}^{\text{emp}}$, and κ_{eff} for each mass-geometry condition;
- geometry-capacity covariance and mediation statistics;
- offline reference-log recovery analysis and the recovery statistic R_{rec} , reported per tier;
- the power analysis for requirements R1–R4 and whether each was met;
- pre-registered exclusion and confound criteria.

13 Conclusion

This protocol specifies a test of Penrose Objective Reduction whose error budget is closed by calibration rather than estimate.

The **calibration arm** (Arm 0; Stages 0–1, carried by BLQC) establishes the apparatus: a reference channel whose finite-rate loss collapses against the frozen deficit $\kappa_{\text{cal}} = h_{\text{KS}} - \eta C_{\text{eff}} \ln 2$, matches the parameter-free noisy-reference null, and is fully recoverable from the passive shadow log. This is classical control physics riding on standard quantum mechanics, and it is the point: reference physics, once calibrated and classified by recoverability, can no longer masquerade as collapse.

The **objective-reduction result** (Arm 1) is the search for an unrecoverable, κ -independent mass-geometry floor with the engineered channel off. That floor is the one observable in this design where a genuinely surprising number can appear: present, it supports a Penrose-type mechanism; absent at adequate sensitivity, it weakens gravitational objective reduction in the tested regime. The crossed sweep (Arm 2) then separates additive from mediated structure when both variables move.

The decisive classifier throughout is recoverability against the passive reference log: $R_{\text{rec}} \approx 1$ marks observer-relative reference bookkeeping; $R_{\text{rec}} \approx 0$, beyond the logging-fidelity budget, marks loss that no reference accounting can explain. Either outcome of the geometry sweep is informative about objective reduction; neither outcome discriminates the Ignorant Observer Framework from quantum mechanics, and the protocol is designed so that this distinction cannot be blurred.

A QGEM Pathfinder Implementation Addendum (Stage 2)

This addendum maps the protocol variables onto a QGEM-style single-mass mesoscopic spatial-superposition pathfinder [17]. The purpose is not to privilege one platform, but to make the abstract variables experimentally concrete. Everything in this appendix is a Stage-2 target (Section 11): the platform class it describes does not yet exist at the required mass, separation, and isolation, and the parameter values are design targets, not demonstrated capabilities.

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.
- **Coding efficiency η :** Calibrated per BLQC's Arm 0 by operating the identical FPGA loop, throttle schedules, and injection paths against the IQ phase monitor of the shadow channel (Section A.5) with the NV system idle; the achieved contraction performance fixes η , which is then frozen for all quantum runs.
- **Reference instability h_{KS} :** Controlled by driving the local-oscillator phase modulation input with a calibrated *deterministic chaotic* source (e.g. a Lyapunov-positive analog or digital chaotic oscillator), *not* additive stochastic noise—the realized reference trajectory must satisfy BLQC's expanding-dynamics gate, and the modulation depth must satisfy the imported dynamic-range budget (attractor scale σ_{sat} above the fit window; threshold crossing inside Gaussian validity). The entropy-rate proxy h_{KS} is extracted offline from the digital error-signal logs, and the positive largest Lyapunov exponent of the realized trajectory is certified from the same logs (with surrogate-data null tests) before the crossed sweep is interpreted.

A.3 Order-of-Magnitude Capacity Estimate

The engineered-channel configurations are only worth running if the throttled reference channel can be driven near the boundary $\kappa_{\text{cal}} \gtrsim 0$. 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—it is, however, exactly the configuration required for the κ -off OR arm (Arm 1).

The relevant C_{eff} for the engineered configurations is the useful rate that actually constrains θ online during the mesoscopic sequence, $C_{\text{eff}} = r b f$, imposed here by the calibrated update-rejection schedule of the preceding subsection. Representative values (pre- η ; the calibrated efficiency rescales the capacity column):

Regime	f (Hz)	b	r	$C_{\text{eff}} \ln 2$	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 $h_{\text{KS}} = 30\text{--}70$ 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 chaotic path of the preceding subsection, not assumed as a native property of the controller— κ is transitional rather than deep-negative, and $t_{\text{break}} \propto 1/\kappa$ falls in the tens-of-ms window where τ_{OR} also lives for $m \approx 10^{-15}$ kg. Both κ knobs are set by design: C_{eff} by the imposed update-rejection schedule, h_{KS} by the injected chaotic phase modulation. This is precisely why the engineered regime cannot discriminate the reference channel from the noisy-reference null (Section 1.5), and why Arm 1 runs with this channel off. The estimate is illustrative— b and r are order-of-magnitude choices—and the point is that the boundary is reachable by deliberate throttling across roughly a decade of f , 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 the engineered-channel configurations of the crossed sweep.

A.4 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-reference-channel result.

A.5 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, implementing the shadow-channel design rules of the BLQC benchmark.

The apparatus should log four synchronized streams:

- **Microwave LO phase:** $\phi_{\text{LO}}(t)$ 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 $I_{\text{coil}}(t)$, and trap-position proxy $x_{\text{trap}}(t)$.
- **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. The recorder’s fidelity must be demonstrated by the injected-test-signal audit of the Arm-0 benchmark; both the Tier-2 floor classification and any Tier-4 claim are bounded by that audit, since any unlogged reference noise depresses R_{rec} .

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

$$\theta_k^{\text{real}} = \theta_k^{\text{commanded}} + \delta\theta_k^{\text{log}}, \quad (15)$$

compute the recovered visibility $V_{\text{corr}}(t)$ using the reconstructed phases, and form the recovery statistic R_{rec} of Section 7.1. R_{rec} is read per tier (Section 8.4): on the engineered channel, $R_{\text{rec}} \approx 1$ is the Arm-0 success condition (observer-relative reference ignorance, as claimed); on an Arm-1 floor component, $R_{\text{rec}} \approx 0$ is the Tier-2 success condition (irreversible loss, not bookkeeping); and a robust R_{rec} deficit on the engineered channel beyond the logging-fidelity budget is the Tier-4 observable, referred to BLQC’s registered module.

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 8.2 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

$$\frac{1}{t_{\text{break}}} = \alpha \frac{1}{\tau_{\text{OR}}} + \beta \kappa_{\text{cal}} + \gamma \quad (16)$$

is a joint fit, not a binary discriminator. The regression admits simultaneous nonzero α and β , and the protocol’s decision rules (Section 8.4) explicitly classify such an outcome as “dual success,” further partitioned into additive and mediated subcases by the geometry–capacity covariance analysis of Section 8.3.

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 operational ($\alpha \approx 0, \beta > 0$). The cut is operational and informational in this regime. Penrose OR, if real elsewhere, is empirically inactive at the masses tested. The corpus’s operational-cut reading is the natural description of this regime—with the Section 1.5 caveat that the β term, on its own, is the engineered channel and carries no evidence beyond H_{null} .

Outcome 2 — Pure Penrose ($\alpha > 0, \beta \approx 0$). Mass-geometry sets the timescale and the engineered channel contributes nothing measurable in this regime. The corpus reading weakens the operational tier’s relevance in the Penrose-overlap band and supports a geometry-dependent mechanism. With $R_{\text{rec}} \approx 0$ on the floor, this is the Tier-2 outcome of Section 8.4.

Outcome 3 — Additive Dual Success ($\alpha > 0, \beta > 0$, no mediation). Mass geometry and information capacity contribute as independent loss 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. The recoverable κ -on contribution rides on top of the floor and does not by itself compete with OR.

Outcome 4 — Mediated Dual Success ($\alpha > 0, \beta > 0$, mass geometry mediates through empirical reference-channel variables). The apparent gravitational signal is mediated by changes in $C_{\text{eff}}^{\text{emp}}$, $h_{\text{KS}}^{\text{emp}}$, or κ_{eff} , and the independent Penrose term weakens once these are included. The corpus’s operational-cut reading is the natural description here: 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} (Section 7.1) runs orthogonal to the regression. It does not adjudicate which mechanism is bigger; it adjudicates the *character* of the loss attributed to either component, and it is read per tier (Section 8.4).

- $R_{\text{rec}} \rightarrow 1$ on a component indicates observer-relative reference loss, consistent with the operational tier and inconsistent with genuine OR collapse.
- $R_{\text{rec}} \rightarrow 0$ 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

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

inflation, a near-singular condition number, with only a combination of α and β identifiable and the two not separately fit.

In a controlled design both regressors are experimenter-set, so the design matrix’s rank is chosen, not discovered: collinearity can emerge only through a failure of experimental control, namely mass geometry driving the empirical tracking variables ($C_{\text{eff}}^{\text{emp}}$, $h_{\text{KS}}^{\text{emp}}$, κ_{eff}). The crossed design can expose whether mass geometry and the empirical tracking variables become statistically entangled, but rank deficiency alone cannot distinguish any deeper common-scale reading from geometry-induced apparatus mediation. Rank deficiency is therefore a diagnostic, not a discovery: it flags the entanglement and hands the interpretive question to the mediation analysis of Section 8.3.

The numerical coincidence that motivates the overlap design—an engineered deficit κ_{cal} of order tens of s^{-1} and the Penrose timescale τ_{OR} for $m \approx 10^{-15}$ kg both landing in the tens-of-milliseconds window (Appendix A)—is exactly that: a numerical coincidence that locates the experimentally interesting regime. The protocol attaches no mechanistic significance to it, and no outcome of the regression should be read as adjudicating one.

B.5 Scope of this Appendix

This appendix does not change the protocol. Specifically:

- The discriminating observable (the κ -off, unrecoverable floor, Section 3.3) is unchanged.
- The decision rules in Section 8.4 are unchanged. 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 10 are unchanged.
- 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.

References

- [1] Aernoud Dekker. Bandwidth-limited quantum control: A finite-rate phase-reference control law and its operational benchmark, 2026. OSF Preprints.
- [2] Roger Penrose. On gravity’s role in quantum state reduction. *Gen. Relativ. Gravit.*, 28(5): 581–600, 1996.
- [3] Aernoud Dekker. The ignorant observer: Formalizing ignorance as a physical constraint on observation, 2026. OSF Preprints.
- [4] Gregor Weihs, Thomas Jennewein, Christoph Simon, Harald Weinfurter, and Anton Zeilinger. Violation of bell’s inequality under strict einstein locality conditions. *Phys. Rev. Lett.*, 81:5039–5043, 1998. doi: 10.1103/PhysRevLett.81.5039.
- [5] Tiff Brydges, Andreas Elben, Petar Jurcevic, Benoît Vermersch, Christine Maier, Ben P. Lanyon, Peter Zoller, Rainer Blatt, and Christian F. Roos. Probing rényi entanglement entropy via randomized measurements. *Science*, 364:260–263, 2019. doi: 10.1126/science.aau4963.
- [6] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nat. Phys.*, 16:1050–1057, 2020. doi: 10.1038/s41567-020-0932-7.
- [7] C. W. Chou, D. B. Hume, M. J. Thorpe, D. J. Wineland, and T. Rosenband. Quantum coherence between two atoms beyond $Q = 10^{15}$. *Phys. Rev. Lett.*, 106:160801, 2011. doi: 10.1103/PhysRevLett.106.160801.
- [8] E. R. Clements, M. E. Kim, K. Cui, A. M. Hankin, S. M. Brewer, J. Valencia, J.-S. Chen, D. B. Hume, D. R. Leibbrandt, D. J. Wineland, et al. Lifetime-limited interrogation of two independent $^{27}\text{Al}^+$ clocks using correlation spectroscopy. *Phys. Rev. Lett.*, 125:243602, 2020. doi: 10.1103/PhysRevLett.125.243602.
- [9] Ridha Horchani. Experimental blueprint for distinguishing decoherence from objective collapse, 2025. arXiv:2512.02838.
- [10] Shaun Laing and James Bateman. Bayesian inference for near-field interferometric tests of collapse models. *Phys. Rev. A*, 110:012214, 2024. doi: 10.1103/PhysRevA.110.012214.
- [11] Germain Tobar, Stefan Forstner, Arkady Fedorov, and Warwick P. Bowen. Testing spontaneous wavefunction collapse with quantum electromechanics. *Quantum Sci. Technol.*, 8(4):045003, 2023. doi: 10.1088/2058-9565/ace2e5.
- [12] Sandro Donadi, Kristian Piscicchia, Catalina Curceanu, Lajos Diósi, Matthias Laubenstein, and Angelo Bassi. Underground test of gravity-related wave function collapse. *Nat. Phys.*, 17(1):74–78, 2021.
- [13] Erwin L. Hahn. Spin echoes. *Phys. Rev.*, 80:580–594, 1950. doi: 10.1103/PhysRev.80.580.
- [14] Yasunobu Nakamura, Yuri A. Pashkin, Tsuyoshi Yamamoto, and Jaw-Shen Tsai. Charge echo in a cooper-pair box. *Phys. Rev. Lett.*, 88:047901, 2002. doi: 10.1103/PhysRevLett.88.047901.
- [15] S. Merlet, J. Le Gouët, Q. Bodart, A. Clairon, A. Landragin, F. Pereira Dos Santos, and P. Rouchon. Operating an atom interferometer beyond its linear range. *Metrologia*, 46(1): 87–94, 2009. doi: 10.1088/0026-1394/46/1/011.

- [16] Angelo Bassi, Kinjalk Lochan, Seema Satin, Tejinder P. Singh, and Hendrik Ulbricht. Models of wave-function collapse, underlying theories, and experimental tests. *Reviews of Modern Physics*, 85:471–527, 2013. doi: 10.1103/RevModPhys.85.471.
- [17] Sougato Bose, Anupam Mazumdar, Marko Schut, and Marko Toroš. Mechanism for the quantum natured gravitons to entangle masses. *Phys. Rev. D*, 105(10):106028, 2022. See also QGEM Collaboration, “Pathfinder design for tabletop tests of quantum gravity,” arXiv:2509.01586 (2025).