Bandwidth-Limited Quantum Control

A finite-rate phase-reference control law and its operational benchmark

https://ignorantobserver.xyz

Aernoud Dekker

June 2026

DOI: 10.17605/OSF.IO/G5WRH

Abstract

We propose Bandwidth-Limited Quantum Control (BLQC), an operational benchmark for testing whether a physical measurement reference can become the limiting variable in an interferometric or qubit experiment. The target is not a modification of the Schrödinger equation, a derivation of the Born rule, or a solution to Bell's theorem, but a narrower control question: whether observer-relative online visibility loss, when the phase or basis reference $\theta(t)$ is tracked through a finite-rate loop, scales with the calibrated deficit

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

where h_{KS} is the certified reference-instability rate, C_{eff} is the useful tracking capacity, and η is measured once in an electronics-only calibration arm and frozen.

In the chaos-wins regime, the benchmark predicts a double-exponential visibility signature

$$V_{\mathrm{IOF}}(t) = \exp[-\tfrac{1}{2}\sigma_0^2 e^{2\kappa_{\mathrm{cal}} t}],$$

with breakdown time controlled by κ_{cal} inside a pre-registered dynamic-range window. A passive high-resolution shadow log then classifies the loss: recovery after conditioning identifies reference-frame bookkeeping within standard quantum mechanics, while any residual below the shadow-log null belongs to a separate strong-tier test. The main use of BLQC is therefore calibration: it provides a recoverable reference-channel subtrahend for mesoscopic objective-reduction experiments, including Penrose-style mass-geometry tests.

Terminology convention. In this project, *The Ignorant Observer Framework* (IOF) refers to the full framework: finite self-knowledge, structural ignorance, basis tracking, philosophical inspiration, and the broader interpretation stack. *Bandwidth-Limited Quantum Control* (BLQC) is the technical module that realizes IOF's κ control law as a finite-rate basis-tracking experiment. The control law originates in IOF [1]; this paper states its operational benchmark and carries the registered strong-tier (chaotic-corner) module, and the accompanying experimental protocol uses the benchmarked channel to test Penrose objective reduction in the mesoscopic regime [2].

Keystone claim. The central claim of this paper is not merely that finite controllers introduce noise. It is that, when the physical variable defining the measurement basis is itself tracked through a finite-capacity channel, the unresolved online basis uncertainty should scale with the self-ignorance rate

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

with the ideal data-rate expression recovered when $\eta = 1$. If this scaling is confirmed in a standalone online-control benchmark, observer-side basis tracking becomes a calibrated variable rather than a purely external idealization. The scaling itself is classical control physics—data-rate logic combined with phase averaging—so confirming it calibrates the apparatus and lends no evidence against standard quantum mechanics. The paper's live experimental value therefore lies in two places: the calibrated control law itself, and the Penrose-overlap regime, where the benchmarked channel calibrates a discrimination of mass-geometry (objective-reduction) scaling from tracking-capacity scaling [2]. The stronger hypothesis—an unrecoverable, capacity-dependent visibility channel beyond standard quantum mechanics—is analyzed and excluded for diffusive and externally randomized reference dynamics in the foundational IOF paper [1]; it survives only as a registered test for certified-chaotic reference dynamics (Section 10.3), with the prior strongly against it.

Scope of the proposal. The proposed test can be stated without adopting a hidden-variable model, a metaphysical interpretation of quantum mechanics, or any claim about consciousness. A reference-channel demonstrator reproduces and calibrates the expected finite-rate online tracking boundary; that recoverable, observer-relative control law is the operational claim of this paper. A stronger reading—that a controlled mesoscopic experiment contains an additional, unrecoverable visibility channel not exhausted by ordinary latency, pulse-error, and offline-recoverable tracking artifacts—is analyzed in the foundational IOF paper [1] and excluded by existing data for diffusive and externally randomized reference dynamics; it remains open only for certified-chaotic reference dynamics, as the registered module of Section 10.3, with the prior strongly against it. Broader interpretive consequences should be weighed against this operational status.

1 Introduction

This paper proposes a narrow experimental question. In an interferometer, spin experiment, or qubit platform, the measurement basis is not only a symbol in a Hilbert-space calculation. It is also a physical reference frame: a phase reference, field direction, pulse phase, local oscillator, or calibrated control axis. The experimenter normally treats this reference as externally known. The present proposal asks what happens when the physical reference itself is treated as an object that must be tracked by a finite-rate controller.

The motivation for treating the reference this way is structural. Any laboratory implementation includes an embedded controller that cannot stand outside the dynamics producing its own measurement question. The basis used to define the measurement is normally absorbed into a textbook idealization—"we know the setting"—but that idealization is itself an engineering claim about a finite tracking loop. The present paper asks what happens to quantum visibility when that idealization is deliberately stressed.

The central hypothesis is simple. If the reference dynamics generate unresolved information at rate h_{KS} while the tracking loop can constrain the reference only at effective rate C_{eff}, then the unresolved basis uncertainty should scale with the deficit

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

where η is the calibrated coding efficiency of the actual estimator/actuator chain. In the chaos-wins regime ($\kappa_{\mathrm{cal}} > 0$), the predicted visibility loss has a characteristic double-exponential form and a breakdown time proportional to $1/\kappa_{\mathrm{cal}}$.

This is not presented as a completed theory of quantum measurement. It is a proposed operational bridge between finite-rate control theory and quantum visibility. A low-risk reference-

channel demonstrator may deliberately throttle the basis-tracking channel and check whether online breakdown times collapse as a function of the calibrated deficit $h_{\mathrm{KS}} - \eta C_{\mathrm{eff}} \ln 2$. That exercise is best understood as a calibration of the observer-relative control boundary. The further physics test is the mesoscopic Penrose-overlap experiment, where the same reference-channel variable can be compared against Penrose-style mass-geometry dependence on the same target platform.

The mesoscopic regime is strategically important because Penrose objective reduction (OR) predicts collapse-like timescales of the same order for certain masses and separations. The two loss classes are not the same: Penrose OR depends on gravitational self-energy, whereas the reference-channel contribution depends on tracking bandwidth and is recoverable when the realized reference is logged. Their overlap means a single setup may discriminate between mass-geometry scaling and tracking-capacity scaling.

Sections 2–3 define the reference-tracking variables and the finite-rate lemma; Sections 4–6 derive the visibility prediction and its no-signaling form; Section 7 states the testable predictions. Sections 8–10 then give the benchmark itself: the hypothesis frame and its logical status, the benchmark protocol (calibration arm, sweeps, shadow channel, noisy-reference null), and the decision rules, sensitivity requirements, and registered strong-tier module. Section 11 states the Born-bridge validation module, Section 12 the staged feasibility path, Section 13 the Penrose-overlap motivation, and Section 14 the falsification criteria.

2 The Measurement Basis as Physical Variable

In standard treatments, the measurement basis θ is treated as a known external parameter. But in the laboratory θ is implemented by apparatus: a polarizer orientation, a magnetic field direction, a microwave pulse phase, a timing reference, or a local oscillator. The "choice" of θ is therefore also a physical process involving control electronics, calibration records, feedback loops, and finite precision.

This is a small reframing with a large consequence. The basis θ is not an external free parameter applied to the experiment from outside the world. It is a physical variable with its own causal history: produced by the controller, maintained by finite hardware, perturbed by chaotic or diffusive internal dynamics, and continually corrected by feedback. The setting and the system therefore share an ancestry in the same physical history. This does not require any new signal between them; it only acknowledges that the measurement question is itself an object of physics, not a free variable inserted from outside.

What follows from this is operational, not metaphysical. If the measurement basis is an internal physical variable, then maintaining it costs information processing rate, and losing track of it has experimental consequences. The standard treatment absorbs both costs into the idealization $\theta = \theta_0$. The present proposal asks under what conditions that idealization holds, and what visibility law replaces it when it does not. A philosophical and interpretive treatment of this move—and its connection to single-history ontologies, Bell-type constraints, and the broader IOF reading—is given in the accompanying document *Where Did the Measurement Basis Come From?*; the present paper restricts itself to the operational control law and its benchmark.

The proposal below does not require a new Bell model. It asks a more basic control question: under what conditions can the apparatus maintain and know its own reference frame well enough that the ideal basis parameter used in the quantum calculation remains experimentally valid? If finite-rate reference tracking contributes to visibility loss, the effect should be visible even

in ordinary interferometric or qubit coherence experiments, without first resolving foundational debates about nonlocality.

Relation to quantum reference frames. Treating the basis as a physical system rather than an external dial is also the starting point of the quantum-reference-frame programme [3], which studies a reference frame as a bounded physical resource—depleted by use, subject to superselection when two parties lack a shared frame, and quantified by a resource theory of asymmetry. The present proposal shares that starting point but differs in object and consequence. Quantum reference frames reorganize the frame-relative content of standard quantum mechanics and predict no departure from it; neither does BLQC. BLQC adds the *finite-rate tracking dynamics* of the frame—the competition between an instability rate h_{KS} and a useful tracking capacity C_{eff}—and predicts observer-relative, recoverable visibility behavior, governed by κ, that a static resource accounting does not describe. The new ingredient is the data-rate dynamics on the frame, not the recognition that the frame is physical.

Relation to adaptive phase estimation. A finite-rate loop tracking a phase reference is, in form, the problem of adaptive quantum phase estimation and quantum feedback control [4, 5]. That literature asks how *well* a phase can be estimated given quantum resources—the standard-quantum and Heisenberg limits mark best-case performance. The present proposal is complementary rather than a special case: it does not pursue optimal estimation but characterizes the *breakdown* regime, in which a deliberately throttled tracking channel can no longer hold the reference and visibility decays with the data-rate deficit κ. Metrology describes the best a controller can do; BLQC describes what happens, and with what signature, when it is forced to do far less.

Relation to the quantum Zeno effect. Repeated observation already has a familiar grip on coherence: frequent projective measurement can freeze unitary evolution—the quantum Zeno effect [6]—or, in a different spectral regime, accelerate decay—the anti-Zeno effect [7]. BLQC uses "observation rate" in a different sense. The Zeno and anti-Zeno effects act through measurement back-action on the system, reshaping its own evolution, and the rate that governs them is the rate of projective measurement. The rate here, C_{eff}, is that of a classical loop tracking the apparatus reference $\theta(t)$: it imprints no projective back-action on the system, and raising it sharpens the basis estimate rather than freezing or hastening the system. The contrast is sharpest against the anti-Zeno effect, where faster observation *shortens* coherence; BLQC predicts the opposite monotonicity, since more tracking capacity lengthens t_{break}. Back-action-rate physics proper—where faster measurement does perturb the system—is treated in the companion capacity–backaction analysis [8].

The following sections develop a quantitative framework: we first define the reference-tracking capacity and the apparatus information-generation rate, then show how their interplay determines when tracking fails and what observable consequences follow.

3 The Tracking Problem

3.1 Definitions

We define an **observer/controller** as any physical system that:

1. Maintains or estimates a reference variable $\theta(t)$ used to define a measurement basis
2. Has finite effective channel capacity C_{eff} devoted to that reference-tracking task

The Landauer limit bounds the maximum physically possible information processing rate given available power P and temperature T:

$$C = \frac{P}{kT \ln 2} \quad \text{(bits/s)} \tag{1}$$

This is an upper bound, not an estimator of the tested quantity. In cryogenic electronics the Landauer bound can be enormous compared with the few-bit or few-hertz effective rate intentionally assigned to a particular tracking loop. The experimentally relevant C_{eff} is the information rate that actually constrains θ, preferably imposed directly by a digital rate limit or update-rate throttle. Power and temperature enter as thermodynamic consistency checks and as practical implementation knobs, not as substitutes for measuring or imposing C_{eff}.

3.2 Apparatus Dynamics: h_{KS}

We characterize the observer's apparatus by its Kolmogorov-Sinai entropy rate h_{KS}—the information-production rate of the classical degrees of freedom (voltage references, timing circuits, feedback loops) that define and maintain the measurement basis.

Even deterministic feedback loops can exhibit chaotic sensitivity: tiny untracked differences grow exponentially, characterized by positive Lyapunov exponents and quantified by h_{KS}. Measurement apparatus naturally tends toward high h_{KS} because amplification from microscopic quantum events to macroscopic records requires nonlinear or threshold dynamics, which generically produce chaos. In contrast, purely diffusive dynamics ($h_{\text{KS}} \to 0$) produce only slow, linear growth of uncertainty—a limiting case in which the unconditioned-record contrast settles to the steady-state diffusive value $V_{\text{IOF}} = \exp[-D_\theta/(2C_{\text{eff}} \ln 2)]$ rather than exhibiting a tracking-failure threshold.

Quantum error correction can be understood as engineering that suppresses effective instability through finite-rate syndrome extraction, decoding, and feedback. The companion capacity–backaction analysis treats this as a standard quantum-engineering tradeoff: useful control capacity and measurement-induced physical instability must be accounted for together, without treating the quantum-computing case as an independent test of BLQC [8].

3.3 The Data-Rate Theorem: Scope and Application

The Data-Rate Theorem (DRT) establishes a necessary condition for mean-square stabilization: for a scalar unstable linear plant with Lyapunov exponent λ, stabilization over a finite-capacity channel requires $C > \lambda/\ln 2$ [9, 10]. For nonlinear systems, related results connect stabilizing data rates to topological feedback entropy and invariance entropy [11, 12], and the closely related task of observing or synchronizing a chaotic system over a rate-limited channel has been treated directly, with the minimal rate set by the sum of positive Lyapunov exponents [13, 14]—the same h_{KS} that enters below.

The present application is narrower than a full nonlinear-control theorem. We use DRT logic to motivate a finite-rate phase-reference lemma that can be tested directly.

Toy lemma (finite-rate phase-reference tracking). Consider a one-dimensional reference error interval Δ_n for θ sampled every Δt. During one step, unresolved reference dynamics expand the interval by approximately $e^{\lambda \Delta t}$. A channel delivering $R\Delta t$ bits about the reference can reduce the interval by at most a factor $2^{-R\Delta t}$ under ideal coding. Thus, up to constants set

by the coding scheme and noise model,

$$\Delta_{n+1} \propto e^{\lambda \Delta t} 2^{-R\Delta t} \Delta_n. \tag{2}$$

Taking the continuous-time limit gives

$$\frac{d}{dt} \ln \Delta \approx \lambda - R \ln 2. \tag{3}$$

Since variance scales as squared error amplitude,

$$\frac{d}{dt} \ln \sigma_\theta^2 \approx 2(\lambda - R \ln 2). \tag{4}$$

For a multi-mode chaotic reference, we replace λ by the relevant entropy-rate proxy h_{KS} and R by the effective tracking-channel rate C_{eff}. This replacement is the empirical bridge tested below. The proposal is not that Eq. (4) is a new theorem of quantum mechanics; it is that phase-reference visibility should reveal whether this finite-rate tracking law is operationally active for the online observer/controller, and whether that same coordinate remains predictive in the Penrose-overlap regime.

3.4 The Self-Ignorance Rate

We first define the ideal **self-ignorance rate** κ as the gap between information generation and ideal processing capacity:

$$\kappa = h_{\mathrm{KS}} - C_{\mathrm{eff}} \ln 2 \tag{5}$$

This is the $\eta = 1$ limit of the benchmark quantity. For a real estimator/actuator chain, the experimentally tested rate is

$$\kappa_{\mathrm{cal}} = h_{\mathrm{KS}} - \eta\, C_{\mathrm{eff}} \ln 2, \tag{6}$$

with η measured and frozen in Arm 0 (Section 9.1). Unless a derivation is explicitly using the ideal toy law, empirical predictions and benchmark decisions below use κ_{cal}.

Following the DRT condition, the sign of the relevant rate defines two regimes:

- **Capacity-wins** ($\kappa_{\mathrm{cal}} < 0$ in the benchmark): Observer can in principle keep up with basis tracking; textbook fixed-basis quantum mechanics is operationally available
- **Chaos-wins** ($\kappa_{\mathrm{cal}} > 0$ in the benchmark): Tracking error is predicted to grow until the basis is no longer resolved to the required tolerance; standard quantum mechanics still holds, with an unstable, unresolved reference frame for the online observer

3.5 Operational Estimators

For the framework to be testable, C_{eff} and h_{KS} must be operationally defined.

Effective capacity C_{eff}: We define C_{eff} as the effective information rate in the basis-tracking loop (not the whole device):

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

where $r \in (0, 1]$ is the fraction of updates that actually constrain θ after overhead, latency, queueing, rejected samples, algorithmic processing limits, and filtering, b is the effective bits per update constraining θ, and f is the update rate (Hz). In the cleanest experiment, C_{eff} is not inferred from total device power or raw Shannon bandwidth; it is imposed by a known rate limiter

on the reference-tracking data stream or inferred from useful online estimator performance. In QEC-like settings the analogous capacity includes decoder throughput, clock-cycle budget, queue depth, and timely Pauli-frame update rate. Varying f or another useful-throughput parameter while holding temperature, readout power, plant dynamics, latency controls, and signal-to-noise ratio fixed is the preferred knob.

Metric entropy rate h_{KS}: For chaotic systems, h_{KS} equals the sum of positive Lyapunov exponents (Pesin identity). We estimate h_{KS} from the rate at which prediction error grows with prediction horizon: given controller state x_t, track how $|x_{t+n\Delta} - g^n(x_t)|$ increases with n, where g^n denotes n iterations of a fitted model. The exponential growth rate estimates $\lambda_{\max}$, which provides a practical proxy (and lower bound) for h_{KS}. Standard algorithms apply to logged digital controller states (e.g., FPGA-based readout). The experiment should verify $h_{\mathrm{KS}} > 0$ before testing capacity dependence. Prediction-error growth is necessary but not sufficient for this: a purely stochastic reference—a random walk, or noise driving a stable loop—can also show error growing with prediction horizon while producing only bounded or diffusive uncertainty, not the multiplicative $e^{2\kappa t}$ separation the law requires. The realized basis-reference trajectory must therefore be certified as genuinely *expanding*—a positive largest Lyapunov exponent of the realized closed-loop trajectory, discriminated from stochastic diffusion by surrogate-data (phase-randomized) null tests—before any capacity or instability sweep is read as a test of the chaos-wins regime. This certification requirement is the expanding-dynamics gate of the benchmark below (Section 9.2).

4 From Phase-Reference Tracking to Visibility

The experimental claim does not require a hidden-variable model. It only requires a signal whose ideal quantum prediction depends on a physical reference phase or basis variable. Write the ideal interference contribution as

$$I_{\mathrm{int}} \propto \cos(\phi - \theta), \tag{8}$$

where ϕ is the system phase and θ is the apparatus reference used to define the measurement basis.

The controller maintains an estimate $\hat{\theta}$ of the realized reference. The physical reference may differ from the intended value by

$$\theta = \theta_0 + \delta\theta, \tag{9}$$

where $\delta\theta$ is not assumed to be fundamental quantum noise. It is the unresolved tracking error of the reference loop. If the experimenter analyzes data as if $\theta = \theta_0$, this unresolved reference error appears as visibility loss.

The bridge hypothesis is therefore:

> If the reference-tracking loop is in the chaos-wins regime, the variance of $\delta\theta$ follows the finite-rate law derived in Eq. (4), with λ replaced by the measured entropy-rate proxy h_{KS} and R by the imposed or estimated C_{eff}.

This framing makes the experiment conservative. A positive result would show that reference tracking contributes to visibility. A null result would rule out the proposed mechanism in the tested regime without requiring a verdict on broader interpretations of quantum mechanics.

5 Basis Uncertainty

5.1 Basis Tracking Error

Due to finite capacity, the controller's estimate $\hat{\theta}$ of the reference basis has uncertainty σ_θ. The realized $\theta(t_m)$ may differ from the intended value θ_0. The key distinction is between the reference frame used in data analysis and the physical frame implemented by the apparatus. The experimenter records outcomes as if $\theta = \theta_0$, but the actual measurement may have been made in basis $\theta = \theta_0 + \delta\theta$.

Ordinary phase noise can also reduce visibility, so the mere presence of visibility loss would not support the model. The distinctive claim is the *scaling*: double-exponential decay with breakdown time $t_{\text{break}} \propto 1/\kappa_{\text{cal}}$, where $\kappa_{\text{cal}} = h_{\text{KS}} - \eta C_{\text{eff}} \ln 2$ is built from independently measured or imposed control variables and a frozen coding efficiency. If present, this structure could be missed when data are analyzed only with ordinary exponential or Gaussian dephasing models; this is a methodological concern rather than evidence that the effect has already been observed.

The framework does not require that capacity limits be uncircumventable in principle. It requires that *for a given online observer/controller with given capacity*, the limit exists and has predictable consequences. If a higher-capacity retrospective observer can recover the visibility from a high-resolution reference log, that recovery should be interpreted as observer-relative reference-frame loss, not irreversible physical decoherence. If C_{eff} increases, t_{break} should increase accordingly (Eq. 15). The claim is that *wherever* the capacity limit lies, it determines online visibility loss through the calibrated tracking variable. Whether one labels this "fundamental" or "technical" is mostly a terminological choice for a standalone demonstrator; in either case the loss is recoverable, observer-relative reference physics. What the Penrose-overlap regime decides is a different axis: whether coherence loss tracks mass geometry (Penrose OR) or tracking variables (κ_{cal}).

5.2 Dynamics

In the chaos-wins regime, we model the basis estimation error as inheriting the exponential separation rate of the uncontrolled reference dynamics, reduced by the effective information rate of corrective updates. For a chaotic mode with Lyapunov exponent λ, uncorrected estimation error scales as $e^{\lambda t}$; therefore its *variance* scales as $e^{2\lambda t}$. In the ideal toy law, finite-rate correction reduces the effective exponent by $C_{\text{eff}} \ln 2$, yielding:

$$\frac{d}{dt} \ln \sigma_\theta^2 \approx 2(h_{\text{KS}} - C_{\text{eff}} \ln 2) = 2\kappa \tag{10}$$

In the benchmark fits, this ideal κ is replaced by the calibrated κ_{cal} of Section 9.1.

This is the finite-rate tracking hypothesis on which the paper rests. It is the minimal application of data-rate logic to observer-side basis tracking: instability expands unresolved basis uncertainty, finite capacity contracts it, and the net observable rate is their difference. The law is classical control physics—data-rate logic combined with phase averaging—not a modification of quantum mechanics. The proposed experiment tests exactly this scaling.

Integrating Eq. (10):

$$\sigma_\theta^2(t) = \sigma_0^2 \, e^{2\kappa t} \tag{11}$$

The intrinsic timescale of this growth is the amplitude e-folding time $\tau_\kappa \equiv 1/\kappa$. The **threshold-crossing time** for uncertainty to grow from σ_0^2 to tolerance σ_{tol}^2 adds a threshold-dependent log

factor:

$$t_{\text{break}} = \frac{1}{2\kappa} \ln\left(\frac{\sigma_{\text{tol}}^2}{\sigma_0^2}\right) = \tau_\kappa \cdot \tfrac{1}{2} \ln\left(\frac{\sigma_{\text{tol}}^2}{\sigma_0^2}\right) \quad (\kappa > 0) \tag{12}$$

5.3 Visibility Decay

When the observer has Gaussian uncertainty in its measurement basis, the observed visibility (fringe contrast) decays as:

$$V_{\text{IOF}}(t) = \exp\left(-\frac{\sigma_0^2}{2} e^{2\kappa t}\right) \tag{13}$$

where σ_0^2 is the initial basis uncertainty at $t = 0$, set by apparatus calibration and initialization quality.

Validity: This assumes Gaussian angle uncertainty and is accurate for $\sigma_\theta \lesssim 1$ rad. For larger uncertainties, higher cumulants contribute and suppression saturates. A diffusive analogue of this phase-averaging law already describes the raw, unconditioned signal in optical-clock correlation spectroscopy, where cross-record analysis recovers the lost contrast; that recovery adjudicates the suppression as observer-relative standard physics, not a new channel (see the foundational IOF paper [1]).

Derivation: The observer intends basis θ_0 but realizes $\theta = \theta_0 + \delta\theta$ where $\delta\theta \sim \mathcal{N}(0, \sigma_\theta^2)$. Averaging the interference term $\cos(\phi - \theta)$ over this distribution:

$$\langle \cos(\phi - \theta) \rangle = \cos(\phi - \theta_0) \cdot e^{-\sigma_\theta^2/2} \tag{14}$$

5.4 Observable Breakdown Time

Equivalently, for a visibility threshold $V_* \in (0, 1)$ (so $\sigma_{\text{tol}}^2 = -2 \ln V_*$), the **breakdown time** t_{break} satisfies $V(t_{\text{break}}) = V_*$:

$$t_{\text{break}} = \frac{1}{2\kappa} \ln\left(\frac{-2 \ln V_*}{\sigma_0^2}\right) \quad (\kappa > 0) \tag{15}$$

This is the primary experimental observable.

6 Correlation Attenuation and No-Signaling Check

The core experiment is an interferometric visibility test, not a Bell-theorem replacement. If the same reference-tracking mechanism is tested in a Bell-type platform, its immediate prediction is conservative: basis uncertainty attenuates correlations below the ideal quantum value while preserving single-party marginals.

With Gaussian basis uncertainty σ^2 that is unbiased and outcome-independent, marginals remain exactly 50/50 (no signaling), but joint correlations are attenuated:

$$\langle E \rangle = -\cos(\theta_A - \theta_B) \cdot e^{-(\sigma_A^2 + \sigma_B^2)/2} \tag{16}$$

This is the independent-error form (independent Gaussian basis errors on the two wings); with correlated reference drift the exponent becomes $-\frac{1}{2} \operatorname{Var}(\delta\theta_A - \delta\theta_B)$, recovering the above when

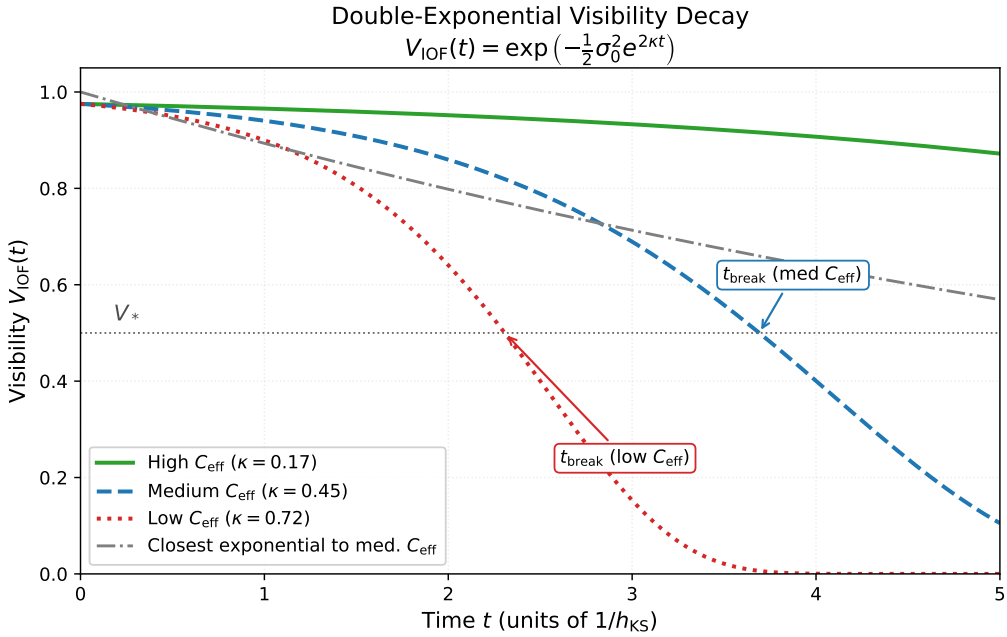

Figure 1: Double-exponential visibility decay for three capacity levels. Visibility $V_{\mathrm{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$ is plotted for high capacity (green, $\kappa = 0.17$), medium capacity (blue, $\kappa = 0.45$), and low capacity (red, $\kappa = 0.72$). Here κ denotes the plotted net rate; in the benchmark it is replaced by the calibrated κ_{cal}. Lower observer capacity yields a higher calibrated net rate, causing earlier breakdown. The arrows mark t_{break}, the time when visibility crosses the threshold $V_* = 0.5$. For comparison, the gray dash-dot line is the closest ordinary exponential ($e^{-\gamma t}$) to the medium-κ curve over its pre-break window: the two closely track each other on the plateau and become statistically separable only at the knee, so the double-exponential shape is distinguishable only once the observation window reaches t_{break}. The prediction: reducing C_{eff} shifts t_{break} leftward.

the cross-covariance vanishes. As uncertainty grows, the CHSH parameter $|S|$ decreases from $2\sqrt{2}$ toward zero—correlations fall *below* ideal QM predictions. Any systematic change in one-party marginals should be treated as an instrumental bias or failed control condition, not as support for the present model.

The distinctive prediction is the κ-dependent scaling: $\sigma^2(t) = \sigma_0^2 e^{2\kappa t}$, yielding double-exponential visibility decay $V_{\mathrm{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$ with breakdown time $t_{\mathrm{break}} \propto 1/\kappa$. Here V_{IOF} is the observer-relative, capacity-dependent factor: phase averaging of the observer's unconditioned records over unresolved reference error, within standard quantum mechanics and recoverable when the missing reference information is supplied—the hallmark of reference-frame physics [3]. The observed raw-record visibility is the product $V_{\mathrm{obs}} = V_{\mathrm{std}}\, V_{\mathrm{IOF}}$, with V_{std} the ordinary physical visibility, including environmental decoherence, which the framework does not modify. Figure 1 illustrates this.

7 Testable Predictions

The framework makes specific predictions that can be compared with standard dephasing and standard online-control failure. The cleanest reference-channel demonstrator is deliberate throttling of the reference-tracking channel. At fixed temperature, fixed readout signal-to-noise ratio, fixed plant dynamics, fixed sequence duration, fixed environmental coupling, and frozen coding

efficiency, increasing the imposed tracking rate C_{eff} should delay online breakdown:

$$\frac{\partial t_{\text{break}}}{\partial C_{\text{eff}}} > 0 \qquad (\kappa_{\text{cal}} > 0). \tag{17}$$

More specifically, the measured breakdown times should collapse against

$$\kappa_{\text{cal}} = h_{\text{KS}} - \eta\, C_{\text{eff}} \ln 2, \tag{18}$$

with

$$t_{\text{break}} = \frac{1}{2\kappa_{\text{cal}}} \ln\left(\frac{-2\ln V_*}{\sigma_0^2}\right). \tag{19}$$

Power is a secondary implementation knob. If increasing controller power increases the effective tracking rate while temperature and readout SNR are clamped, extra useful online capacity should extend observer-relative visibility time rather than shorten it—an operational statement about reference quality in a closed tracking loop, expected from classical control, not a signature distinguishing the framework from standard quantum mechanics or from decoherence. If power is increased without increasing the imposed or measured C_{eff}, no improvement is predicted. This distinction is important because the Landauer expression is only an upper bound on possible processing, not the actual rate assigned to the basis-reference loop.

Observer-relative visibility should follow the double-exponential decay $V_{\text{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa_{\text{cal}}t})$, qualitatively distinct from standard exponential ($e^{-\gamma t}$) or Gaussian (e^{-t^2}) decoherence. Higher apparatus entropy rate h_{KS} should produce faster online visibility loss (smaller t_{break}), independent of environmental decoherence channels, when C_{eff} and η are held fixed. A fixed-duration loss that is fully recoverable from a withheld high-resolution reference log is a successful observer-relative tracking benchmark, but not evidence of irreversible physical collapse.

Qualitative distinctness is not the same as measurability. With finite shot noise and a finite observation window, the double-exponential bend must be separated not from a fixed exponential but from the *closest* ordinary-decay curve. Figure 2 quantifies this for the strongest such single adversary, the best-fitting stretched exponential $\exp[-(\Gamma t)^\beta]$, which contains pure-exponential ($\beta = 1$) and Gaussian ($\beta = 2$) decay as special cases. Modelling each visibility estimate with fringe shot noise (variance $\approx (1-V^2)/N$ per point) and refitting both forms to every dataset, the receiver-operating-characteristic area (AUC) for distinguishing them rises from chance (0.5) to near-certain (1) only once the observation window passes t_{break}: below roughly half of t_{break} the bend has not yet developed and the two forms are statistically indistinguishable, while past the knee a few hundred shots per point at percent-level fringe precision suffice for clean separation. The operational consequence is that the wrong-functional-form falsification criterion stated below has discriminating power only when the window reaches the breakdown knee; an experiment that stops short of t_{break} cannot separate the IOF signature from ordinary dephasing at any shot count. This is a statement about measurement budget, not evidence for the framework; the full method and code accompany the online verification materials.

The benchmark protocol that operationalizes these predictions—calibration arm, certified sweeps, shadow channel, noisy-reference null, decision rules, and sensitivity requirements—is specified in Sections 8–10. The mesoscopic objective-reduction experiment that uses the benchmarked channel is specified separately [2]. QGEM-style pathfinder apparatuses [15], designed for extreme isolation and cryogenic operation with mesoscopic masses, may provide useful future platforms if they expose the required reference-tracking and mass-geometry knobs.

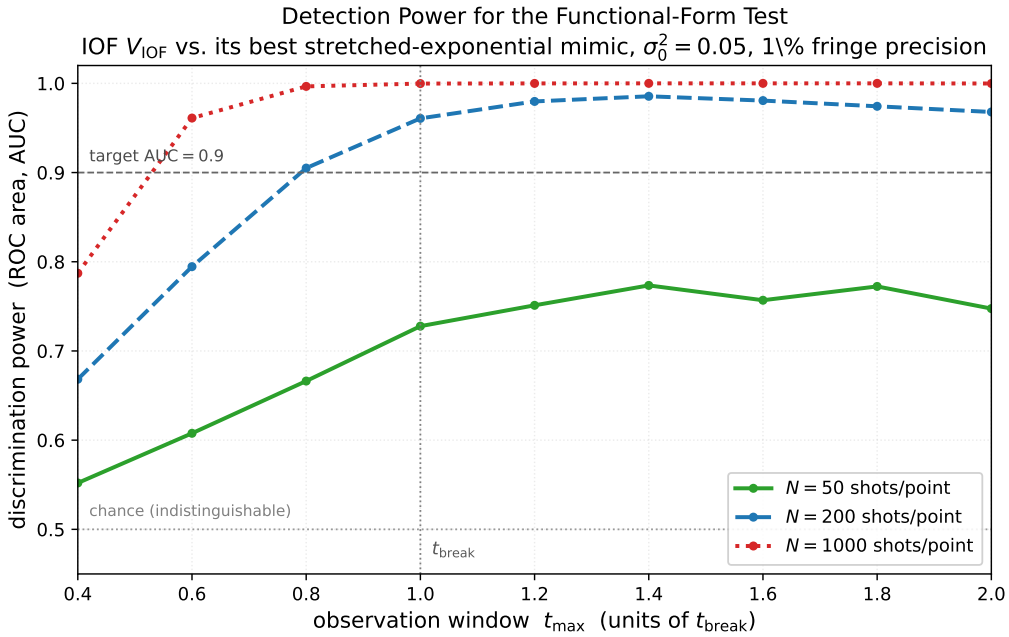

Figure 2: Detection power for the functional-form test. ROC area (AUC) for distinguishing the IOF double-exponential $V_{\mathrm{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$ from its best stretched-exponential mimic $\exp[-(\Gamma t)^\beta]$, versus the observation window t_{max} (in units of t_{break}), for three shot budgets per time point (green $N = 50$, blue $N = 200$, red $N = 1000$; $\sigma_0^2 = 0.05$, 1% fringe precision). AUC = 0.5 is chance (statistically indistinguishable); AUC $\to 1$ is clean separation. The signature becomes resolvable only as the window passes t_{break}, and beyond a few hundred shots measurement precision—not shot count—sets the limit. The axis is normalised to t_{break}, so the curve is independent of κ and of the absolute timescale. A statement about experimental budget, not evidence for the framework.

8 Hypotheses and Logical Status

The benchmark distinguishes four hypotheses about the origin of a measured visibility-loss timescale.

H_{null} — **Standard QM with a noisy finite-bandwidth reference.** The measured trajectory of the basis reference, recorded at full resolution by a passive shadow channel (Section 9.7), predicts the observed visibility loss by ordinary phase averaging:

$$V_{\mathrm{null}}(t) = V_{\mathrm{baseline}}(t)\, \exp\!\left[-\tfrac{1}{2}\,\mathrm{Var}\big(\delta\theta_{\log}(t)\big)\right]. \tag{20}$$

Nothing beyond textbook quantum mechanics and classical control theory is involved.

H_{op} — **Operational (weak) BLQC.** The same reference-channel loss, measured online by the throttled observer, collapses against the calibrated deficit $\kappa_{\mathrm{cal}} = h_{\mathrm{KS}} - \eta\, C_{\mathrm{eff}} \ln 2$ (with the coding efficiency η of Section 9.1) built from independently measured or imposed quantities. This establishes that online observer-relative visibility is governed by a finite-rate tracking boundary—a calibrated, quantitative control law. H_{op} is consistent with H_{null}; in the engineered regime it is the control-theoretic refinement of H_{null}, not an alternative to it.

H_{strong} — **Strong BLQC.** Visibility loss remains beyond the H_{null} prediction computed from the full-resolution reference log, and the residual still scales with κ_{cal}. This would be new physics: a departure from standard quantum mechanics conditional on the measured reference trajectory. This reading is raised and excluded as analysis in the foundational IOF paper [1] (Part III): existing experiments that sort outcomes against passively logged references—logged-settings Bell tests [16], randomized-measurement and classical-shadow protocols [17, 18], and

correlation spectroscopy [19, 20]—recover the coherent structure that the diffusive/stochastic form of H_{strong} would forbid. What remains untested is the deterministic-chaotic-basis corner (certified positive largest Lyapunov exponent, Section 9.2); this paper registers that corner (Section 10.3) with the prior strongly against.

H_{OR} — **Objective reduction.** A κ-independent, unrecoverable ($R_{\text{rec}} \approx 0$, Section 9.7) visibility floor follows mass geometry according to a Penrose-type estimate. H_{OR} is not adjudicated by this benchmark: it requires a mesoscopic platform with the engineered channel off and tracking capacity maximal, and its arms, decision rules, and sensitivity requirements are specified in the companion protocol [2].

One structural relation among the non-null hypotheses is worth recording. H_{strong} and H_{OR} agree on the signature class—an unrecoverable loss, $R_{\text{rec}} \approx 0$ beyond the logging-fidelity budget—and are rivals only on the controlling variable: observer-side κ against gravity-side mass geometry. A confirmed strong-tier residual would therefore weigh against the gravitational reading as much as against standard quantum mechanics, since it would locate the unrecoverable floor in the observer's tracking ledger rather than in spacetime; conversely, a mass-geometry floor at maximal tracking capacity owes nothing to κ. The crossed design of the companion protocol [2] discriminates the two rather than conflating them.

The two axes. The central structural fact about the engineered regime must be stated openly: **it is blind along the BLQC-versus-null axis.** Both κ knobs are set by design—C_{eff} by a calibrated update-rejection schedule, h_{KS} by injected, certified-expanding reference modulation. Given that construction, H_{null} and H_{op} predict identical raw visibility, identical recoverability, and identical κ-scaling in every configuration: 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. The data-rate theorem that grounds the BLQC law is itself a result of classical control theory. No sweep of C_{eff} or h_{KS}, however well confounds are controlled, can separate the two. What the benchmark *can* establish is the operational tier (κ_{cal} collapse, V_{null} agreement, recoverability), and what it can *test* is the strong axis: whether any loss survives the parameter-free H_{null} correction computed from the shadow log (Section 9.8). The OR axis belongs to the companion protocol.

9 The Benchmark Protocol

9.1 Calibrated Coding Efficiency η

The idealized data-rate contraction $C_{\text{eff}} \ln 2$ assumes that every delivered bit halves the reference-error interval. No real estimator/actuator chain attains that bound: the achieved contraction depends on quantizer design, latency, model mismatch, and noise. The empirically tested law is therefore

$$\kappa_{\text{cal}} = h_{\text{KS}} - \eta\, C_{\text{eff}} \ln 2, \qquad 0 < \eta \leq 1, \tag{21}$$

where η is the loop's coding efficiency.

Calibration rule (pre-registration condition). η must be measured *once*, in the electronics-only calibration arm (Arm 0, Section 9.4), from the observed contraction performance of the actual estimator/actuator chain across the full throttle and injection schedules—and then frozen. 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: that would make the κ-collapse circular. The factors b and r are measured with uncertainties (known bit depth; logged accepted/rejected-update fraction) rather than fit, and the coefficient $\ln 2$ is held fixed at its information-theoretic

value. This is the "no free capacity scale" rule: one independently calibrated efficiency, fixed before the quantum test, and no further freedom.

9.2 The Expanding-Dynamics Gate

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, estimated preferably from the full-resolution shadow log rather than the throttled online stream. Acceptable estimation methods include prediction-error growth from logged controller/plant states, Lyapunov exponent estimation for a fitted reference-dynamics model, and calibrated injected reference instability with the injection path included in the apparatus model.

Gate (pre-registration condition). The chaos-wins law $\sigma_\theta^2(t) = \sigma_0^2 e^{2\kappa t}$ requires the realized basis-reference trajectory to be *intrinsically expanding*—a positive largest Lyapunov exponent of the realized reference trajectory, after the actual feedback and injection paths are included. Additive stochastic noise driving an otherwise stable tracking loop is *not* an acceptable realization of h_{KS} for this law: it produces a bounded, stationary error distribution (ordinary dephasing), not the multiplicative $e^{2\kappa t}$ separation. Because prediction-error growth alone can also be produced by stochastic diffusion (a random walk, or transient growth in a stable loop), the realized trajectory must be *certified* as expanding—a positive largest-Lyapunov-exponent estimate from the error-signal logs, together with surrogate-data (phase-randomized) null tests that reject the stochastic-diffusion hypothesis—before the capacity and instability sweeps of Section 9.5 are interpreted. A configuration that fails this gate has not realized the chaos-wins regime and cannot be used as a falsification test of the capacity- or instability-dependence criteria: a bounded-error run is an unmet precondition, not a null. The gate also delimits the one strong-tier corner left open by existing data: diffusive/stochastic reference dynamics are already adjudicated against the strong reading (Sections 8 and 9.8), so only certified-expanding configurations bear on the registered module of Section 10.3.

Gate-failure accounting. To prevent the gate from functioning as an unfalsifiability loophole, every attempted configuration must be reported, gate-passing or not. If more than a pre-registered fraction (default: one third) of attempted instability configurations fail certification, the report must state that the chaos-wins regime was not adequately realized, and may not present the surviving subset as an unbiased test of the κ law.

9.3 Dynamic-Range Budget

The exponential law holds only over a finite window, and the experimenter must budget that window explicitly before data collection. Three scales bound it from above:

- **Gaussian validity:** the visibility law $V = \exp(-\sigma_\theta^2/2)$ assumes Gaussian angle uncertainty and is accurate for $\sigma_\theta \lesssim 1$ rad; beyond this, higher cumulants contribute and suppression saturates.
- **Phase wrapping:** for a phase variable, error growth is capped near $\sigma_\theta \approx \pi$ regardless of the injected attractor size; the distribution approaches uniform-on-circle and visibility approaches a hard floor.
- **Attractor saturation:** a bounded chaotic source produces exponential separation only up to its attractor diameter σ_{sat}; growth saturates there.

Because Gaussian validity (≈ 1 rad) and wrapping ($\approx \pi$ rad) are separated by less than a decade, the entire fit window—including the region around t_{break} where the double-exponential bend is resolved—must lie below the Gaussian-validity scale.

Pre-registered budget. Before the sweeps, the experimenter must register:

$$\sigma_0 \ll \sigma_{\text{tol}} \lesssim 1 \text{ rad} < \min(\pi, \sigma_{\text{sat}}), \qquad N_{\text{efold}} \equiv \ln\frac{\sigma_{\text{tol}}}{\sigma_0} \geq N_{\min}, \tag{22}$$

with a default minimum of $N_{\min} = 4$ amplitude e-folds inside the validity window. Since $\sigma_{\text{tol}}^2 = -2\ln V_*$, this constrains the visibility threshold from below: $V_* \geq 0.8$ gives $\sigma_{\text{tol}} \approx 0.67$ rad, comfortably inside Gaussian validity, while $V_* = 0.05$ would place the threshold crossing outside it. Example: $\sigma_0 = 1$ mrad and $V_* = 0.8$ give $N_{\text{efold}} = \ln(668) \approx 6.5$. Fits of the double-exponential form must be restricted to the pre-saturation window, and the report must state where saturation onset was observed.

9.4 Arm 0: Electronics-Only Calibration

With the quantum system removed or idle, operate the full tracking loop—throttle schedules, injection path, estimator, actuator—against a reference-grade phase monitor (IQ phase meter referenced to the master clock; in practice, the shadow channel of Section 9.7). Measure:

- the residual-error statistics $\text{Var}(\delta\theta(t))$ under every throttle/injection schedule to be used in the quantum runs;
- the coding efficiency η (Section 9.1), from the observed contraction performance;
- h_{KS} and the largest-Lyapunov-exponent certification for every injection configuration (expanding-dynamics gate);
- the realized dynamic-range budget (Section 9.3).

Classical control theory predicts that $\text{Var}(\delta\theta(t))$ in this arm already follows the κ_{cal} law; that is expected, and it is the point: **Arm 0 is calibration, not a test of new physics.** Its outputs (η, h_{KS}, gate certificates, dynamic-range budget) are frozen before any quantum data are taken. Arm 0 also fixes the H_{null} prediction machinery: the same logs and the same variance estimator later used to compute V_{null} in the quantum runs.

9.5 Arm 1: Capacity and Instability Sweeps

On the quantum platform, at fixed plant dynamics and fixed environmental confounds, run the engineered-channel sweeps against the frozen Arm-0 calibration:

Capacity sweep. At fixed h_{KS} (certified), fixed temperature, readout SNR, pulse/actuator behavior, and sequence duration, vary C_{eff} across a pre-registered range. Extract t_{break} per condition. The operational prediction is

$$\frac{\partial t_{\text{break}}}{\partial C_{\text{eff}}} > 0 \quad (\kappa_{\text{cal}} > 0), \tag{23}$$

with collapse against κ_{cal} under the frozen η.

Instability sweep. At fixed C_{eff} and fixed ordinary confounds, vary h_{KS} across certified-expanding configurations. The operational prediction is

$$\frac{\partial t_{\text{break}}}{\partial h_{\text{KS}}} < 0 \quad (\kappa_{\text{cal}} > 0). \tag{24}$$

Both predictions are shared with H_{null} (Section 8); this arm benchmarks the operational tier and feeds the strong-axis residual analysis (Section 9.8). It requires no mesoscopic mass and can be run on a standalone qubit or interferometric platform (Stage 1, Section 12).

Run order should be randomized across capacity and instability settings, each condition must include enough repetitions to meet the sensitivity requirements of Section 10.2, and replication should include at least one repeated condition after the full parameter sweep to detect slow drift.

9.6 Confound Controls

Thermal control. Temperature must be monitored at the mixing chamber and, where possible, through 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 any tier.

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.

The full platform-requirements and confound-controls core for the mesoscopic objective-reduction experiment—including geometry–capacity covariance auditing—is specified once, in the companion protocol [2]; the controls above are the subset a standalone reference-channel demonstrator requires.

9.7 The Shadow Reference Channel and the Recovery Statistic

The realized reference must be logged at higher resolution than the online tracker receives, by a passive shadow channel.

Crucial design rule. The online tracker receives only the throttled C_{eff} stream. The offline recorder receives the full-resolution streams—the reference phase through an independent monitor, the controller state (accepted and rejected updates, estimator state, trigger timestamps), plant diagnostics, and shot-level outcomes, all synchronized to the master clock—but must be causally isolated from the feedback loop: it may be used for retrospective reconstruction but must not contribute to the online C_{eff} available during the run.

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

$$\theta_k^{\text{real}} = \theta_k^{\text{commanded}} + \delta\theta_k^{\text{log}}. \tag{25}$$

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

$$R_{\text{rec}} = \frac{V_{\text{corr}} - V_{\text{raw}}}{V_{\text{baseline}} - V_{\text{raw}}}, \tag{26}$$

where $V_{\text{baseline}}(t)$ is the measured visibility with the engineered channel off. $R_{\text{rec}} \approx 1$ classifies the loss as online observer-relative reference ignorance—recoverable reference-frame physics; $R_{\text{rec}} \approx 0$ classifies it as irreversible physical loss. On the engineered channel, $R_{\text{rec}} \approx 1$ is the benchmark success condition (Section 10.1); a robust deficit beyond the logging-fidelity budget is the registered strong-tier observable (Section 10.3); and on a mass-geometry floor in the companion protocol, $R_{\text{rec}} \approx 0$ is the objective-reduction signature [2].

Logging-fidelity audit. The strong tier and the R_{rec} classification are interpretable only if the shadow log's fidelity is independently demonstrated: any unlogged classical reference noise mimics a strong-tier residual. Known test signals (amplitude, spectrum, and timing preregistered) must be injected into the reference path and recovered from the shadow log within a quantified error budget. That budget propagates into the V_{null} systematic error (Section 10.2); a claimed residual smaller than the logging-fidelity budget is not a result.

9.8 The Noisy-Reference Null Model

The reference model against which all reference-channel effects are judged is the parameter-free H_{null} prediction computed per condition from the shadow log:

$$V_{\text{null}}(t) = V_{\text{baseline}}(t)\, \exp\!\left[-\tfrac{1}{2}\,\mathrm{Var}\big(\delta\theta_{\log}(t)\big)\right], \tag{27}$$

where $\mathrm{Var}(\delta\theta_{\log}(t))$ is estimated from the full-resolution reference log using the estimator frozen in Arm 0, and $V_{\text{baseline}}(t)$ is the measured engineered-channel-off baseline. No parameters are fit.

Two distinct statements can then be registered:

- **Weak (operational) BLQC:** $\mathrm{Var}(\delta\theta_{\log}(t))$ itself follows the calibrated law $\sigma_0^2 e^{2\kappa_{\text{cal}} t}$ within the dynamic-range window, with η frozen—and the raw quantum visibility matches V_{null} within the error budget.
- **Strong BLQC:** a statistically significant deficit $V_{\text{null}}(t) - V_{\text{raw}}(t) > 0$ remains after the logging-fidelity budget and all confound audits, and that residual scales with κ_{cal}.

The second is a far harder claim. It is the strong reading of the corpus's visibility law, raised and excluded in the foundational IOF paper [1] (Part III) against existing recoverability experiments; in particular, correlation-spectroscopy data [19, 20] already show that diffusive reference suppression of the raw unconditioned signal is fully recovered by cross-record sorting, adjudicating the diffusive form of the law (operational layer real, strong layer false). The strong tier is therefore confined to the certified deterministic-chaotic corner (Section 9.2) and is registered without being asserted (Section 10.3), with the prior strongly against.

10 Decision Rules, Sensitivity, and the Registered Strong-Tier Module

10.1 Benchmark Success (Operational Tier)

The benchmark succeeds when all of the following hold:

- At fixed plant and fixed ordinary confounds, t_{break} increases with C_{eff} and decreases with h_{KS} (Section 9.5).
- Breakdown times collapse against κ_{cal} or $\rho = \eta\, C_{\text{eff}} \ln 2/h_{\text{KS}}$, with η frozen from Arm 0, better than against raw power, temperature, readout SNR, elapsed time, or actuator diagnostics.
- Raw visibility matches the parameter-free V_{null} within the error budget, and $R_{\text{rec}} \approx 1$. **Here recoverability is a success condition**: it certifies that the loss was online observer-relative reference ignorance, as the operational tier claims.

Success validates the finite-rate reference-channel model as an operational mechanism—a calibrated control law. **It is explicitly not evidence beyond** H_{null} (Section 8): the κ-scaling is expected classical control physics, and passing the benchmark calibrates the apparatus for downstream use—most importantly as the calibration arm of the objective-reduction protocol [2]—without discriminating anything from standard quantum mechanics.

10.2 Sensitivity and Power Requirements

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.

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

$$\sigma_V \approx \sqrt{\frac{1 - V^2}{N}}. \tag{28}$$

Near threshold, the local slope of the double-exponential law gives the propagated breakdown-time uncertainty

$$\sigma(t_{\text{break}}) \approx \frac{\sigma_V}{\kappa_{\text{cal}}\, \sigma_{\text{tol}}^2\, V_*}, \qquad \sigma_{\text{tol}}^2 = -2 \ln V_*. \tag{29}$$

Worked example: $V_* = 0.8$, $N = 200$ shots per time point gives $\sigma_V \approx 0.042$ and $\sigma(t_{\text{break}}) \approx 0.12/\kappa_{\text{cal}}$, against $t_{\text{break}} \approx 6.5/\kappa_{\text{cal}}$ for $\sigma_0 = 1$ mrad—about 2% relative precision per condition.

R1 — Sweep sensitivity. Register the smallest relative change in t_{break} detectable at 90% power across the swept C_{eff} range. Default target: 20%.

R2 — Strong-tier budget. Quantify the total systematic error on V_{null}: logging-fidelity budget (from the injected-test-signal audit), calibration drift, and estimator bias. The minimum claimed deficit $1 - R_{\text{rec}}$ must exceed this budget by a factor of at least 3.

R3 — Drift bound. The repeated-condition drift check (Section 9.5) must bound slow drift over the full campaign to less than one third of the smallest registered minimum detectable effect.

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

10.3 The Registered Strong-Tier Module (Chaotic Corner)

The strong reading of the visibility law—that contrast lost to an unresolved basis is unrecoverable, a departure from standard quantum mechanics conditional on the measured reference trajectory—is excluded for diffusive and externally randomized reference dynamics by existing experiments (Section 8). One corner has never been directly probed: recoverability under *certified deterministic-chaotic* basis dynamics, with a positive largest Lyapunov exponent certified by the expanding-dynamics gate. Nothing in the formalism motivates a difference between chaotic and diffusive unresolved variation at the level of record statistics, so the prior is strongly against; the shadow-channel design tests the corner almost for free, and this module registers that test without asserting its outcome. Observing a strong-tier residual would overturn the exclusion rather than confirm the corpus's surviving claims.

The registered positive criterion requires all of:

- a statistically significant deficit of V_{raw} below V_{null}, in gate-certified configurations, that survives the logging-fidelity audit (Section 9.7) and all confound controls;
- the deficit scales with κ_{cal};
- the deficit exceeds the total systematic budget on V_{null} by the pre-registered margin (requirement R2).

This is an extraordinary-claim module: any unlogged classical noise mimics it, so the evidentiary bar is set by the demonstrated logging fidelity, not by statistical significance alone. A null—raw visibility matching V_{null} within the systematic budget everywhere—is not merely the default expectation: it is the outcome already indicated by the existing logged-settings, randomized-measurement, and correlation-spectroscopy data, and it is fully consistent with the IOF corpus, whose operational and interpretive claims do not depend on this module.

11 Born-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; it is required for the corpus's conditional Born-rule derivation [21], not for the benchmark above or for the OR discrimination [2]. One asymmetry is registered up front: a flat $I(\theta)$ is also what standard quantum mechanics predicts for an ideal binary readout, so a homogeneous result confirms a prediction the bridge shares with quantum mechanics and lends no positive evidential weight to the bridge. The module is a falsification-only check on the bridge premise.

11.1 Operational Record Calibration

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

$$p(o|\theta) \tag{30}$$

across a fine-grained, calibrated grid of the physical basis setting θ—for example, by sweeping the intended phase θ_0 of the final readout pulse across the full dial and mapping the outcome probabilities. The full-dial map serves calibration; the homogeneity verdict uses only the active range below.

11.2 Empirical Fisher Information

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

$$I(\theta) = \sum_o \frac{1}{p(o|\theta)} \left(\frac{\partial p(o|\theta)}{\partial \theta} \right)^2 . \tag{31}$$

11.3 Homogeneity Against the Finite-Visibility Null

The homogeneity test must be specified against the correct null, because finite readout visibility produces inhomogeneity by itself. For the ideal binary Born fringe, $p(0|\theta) = \cos^2(\theta/2)$, the binary Fisher information is exactly flat: $I(\theta) = 1$ (the apparent $0/0$ at the fringe extrema resolves to 1 in the limit). For finite visibility $V < 1$, $p(0|\theta) = \frac{1}{2}(1 + V\cos\theta)$ gives

$$I(\theta) = \frac{V^2 \sin^2\theta}{1 - V^2\cos^2\theta}, \tag{32}$$

which is *not* flat: it vanishes at the fringe extrema, with dips of angular width of order $\sqrt{1-V^2}$, purely as a consequence of readout imperfection.

The module therefore pre-registers:

- a **finite-V null model**: the $I(\theta)$ profile predicted from the independently measured readout visibility and noise;
- an **active operating range**: homogeneity is evaluated over the range of θ actually used by the tracking loop (in practice, a neighborhood of mid-fringe), not blindly over 0 to 2π;
- an **acceptance criterion**: the bridge premise is challenged only by inhomogeneity *in excess of* the finite-V null over the active range. Readout-induced non-flatness is apparatus imperfection, not bridge failure.

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

11.4 Fisher-Unit Error Reporting

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

$$\sigma_s^2 = I(\theta)\sigma_\theta^2. \tag{33}$$

For larger intervals, report

$$s(\theta) = \int^{\theta} \sqrt{I(u)}\, du. \tag{34}$$

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

12 Staged Feasibility Path

Stage 0 — Electronics-only calibration (executable now). Arm 0 in full: the tracking loop, throttle schedules, injection paths, expanding-dynamics certification, η calibration, and dynamic-range budget, against a reference-grade phase monitor. No quantum system required. Classical control theory predicts the κ_{cal} law here; Stage 0 verifies and calibrates it.

Stage 1 — Standalone reference-channel demonstrator (near-term). The sweeps of Section 9.5 plus the Fisher module on an existing qubit or interferometric platform (NV center, trapped ion, atom interferometer). No mesoscopic mass is required, because the operational tier never depended on one. Deliverables: the benchmark verdict (κ_{cal} collapse, V_{null} agreement, $R_{\text{rec}} \approx 1$), the Fisher-homogeneity analysis, and a first bound on the strong tier from the achieved logging fidelity.

Stage 2 — Penrose-overlap experiment (future platform). The mesoscopic objective-reduction discrimination, with this benchmark as its calibration arm, is specified in the companion protocol [2], including the required platform, the κ-off geometry sweep, the crossed design, and a QGEM-style implementation addendum.

A report from any stage must identify its stage and claim only that stage's results.

13 Penrose-Overlap Regime

The framework's predictions can be compared with other approaches to quantum state reduction. Penrose's gravitational Objective Reduction (OR) provides a natural comparison because both predict characteristic waveform-collapse or visibility-loss timescales in a mesoscopic regime, while depending on different physical variables. This overlap is not used as evidence for the present mechanism. It is used to identify an experimental regime in which two candidate mechanisms can be discriminated.

Penrose proposes that quantum superpositions become unstable due to gravitational self-energy, giving a characteristic timescale [22]:

$$\tau_{\text{OR}} = \frac{\hbar}{E_G} \sim \frac{\hbar s}{G m^2} \tag{35}$$

where m is the mass in superposition and s is the spatial separation. For mesoscopic masses ($m \sim 10^{-15}$ kg) with separations of 100 nm–1 μm, this gives $\tau_{\text{OR}} \sim$ 10–100 ms.

The empirically sharp form of this idea is the Diósi–Penrose (DP) collapse model, whose decoherence and spontaneous-radiation predictions—unlike the bare $\tau_{\text{OR}} \sim \hbar/E_G$ heuristic—are quantitative enough to test directly. An underground search for the predicted spontaneous radiation has excluded the natural, parameter-free version of the DP model [23], so the gravitational-collapse target in this regime is a partially closing one. This does not weaken the discriminator below: the BLQC signature is defined *against* gravitational OR, so a more tightly constrained OR

makes a positive, capacity-dependent κ-signal correspondingly easier to read as non-gravitational rather than harder. The discriminator is also not DP-specific. Every model in the dynamical-collapse family fixes its own scaling variable for an unrecoverable floor—amplification with nucleon number at a universal rate for spontaneous-localization (GRW/CSL) models, mass and superposition geometry here—and none of them responds to the observer's tracking budget, so a capacity-dependent κ-signal reads against whichever collapse schedule is on trial.

The reference-channel model predicts an intrinsic tracking-loss timescale (the amplitude e-folding time) in the chaos-wins regime:

$$\tau_\kappa = \frac{1}{\kappa_{\text{cal}}} = \frac{1}{h_{\text{KS}} - \eta\, C_{\text{eff}} \ln 2} \tag{36}$$

For illustrative apparatus parameters $h_{\text{KS}} \approx 50$ nats/s and $\eta C_{\text{eff}} \ln 2 \approx 7$ nats/s, we obtain $\kappa_{\text{cal}} \approx 43$ nats/s and $\tau_\kappa = 1/\kappa_{\text{cal}} \approx 23$ ms. Including the log factor from Eq. (12) for 1–5% visibility thresholds contributes a factor of 2–3, giving an observable breakdown time $t_{\text{break}} \approx 50$–$70$ ms.

Numerical proximity: For the apparatus parameters above, t_{break} falls near the Penrose-predicted τ_{OR} range for mesoscopic systems. This is the "Penrose-overlap" regime: both loss classes can appear on comparable timescales, but they assign the timescale to different knobs.

Different predictions: Despite this proximity, the two loss classes have different primary dependencies:

- Penrose OR: gravitational instability; τ_{OR} depends on mass and geometry, independent of observer bandwidth
- Reference channel: recoverable tracking loss; t_{break} depends on useful online reference-tracking capacity, independent of spatial geometry

This enables experimental discrimination:

Test	**Penrose OR**	**Reference channel**
Vary C_{eff} at fixed mass/geometry	No effect on τ	t_{break} increases
Vary h_{KS} at fixed mass/geometry	No effect on τ	t_{break} decreases
Vary separation s at fixed C_{eff}	τ_{OR} increases	No effect on t_{break}

The primary discriminator is therefore not the absolute collapse time, but the derivative of that time with respect to independently controlled variables and, decisively, recoverability against the passive reference log. If increasing useful reference-tracking capacity extends the relevant visibility or coherence time while mass geometry is fixed, the reference-channel calibration is supported in that regime. If mass-geometry variables explain an unrecoverable floor and κ_{cal} adds no predictive value after ordinary confounds are controlled, gravitational OR remains the better explanation in this regime. The two contributions are not framed as mutually exclusive: the full statistical treatment of the overlap regime—including the additive combined-rate model and geometry–capacity mediation—is developed in the companion protocol [2], which carries the discrimination procedure with this benchmark as its calibration arm.

14 Falsification Criteria

The stakes of these criteria differ by role. Criteria 1–4 test the classical control-law calibration, and apply only in gate-certified configurations (Section 9.2) with the sensitivity requirements of Section 10.2 met: failing any of them invalidates the κ control law in the tested regime, while passing them validates the calibration without lending evidence against standard quantum mechanics, since the κ-scaling is expected classical control physics. Criterion 5 concerns the objective-reduction axis, executed in the companion protocol [2]; criterion 6 is a no-signaling sanity check. The proposal fails if any of the following hold:

1. **Wrong capacity dependence**: At constant temperature, fixed readout SNR, fixed sequence duration, and independently monitored environmental channels, increasing the imposed reference-tracking rate C_{eff} fails to extend online breakdown time, i.e., $\partial t_{\text{break}}/\partial C_{\text{eff}} \leq 0$ in the claimed chaos-wins regime.
2. **No data collapse**: Breakdown times fail to collapse against κ_{cal} after h_{KS} and C_{eff} are independently measured or imposed. This collapse is a parameter-free test: in $C_{\text{eff}} = r\,b\,f$ the factors b and r are measured with uncertainties (the known update bit depth and the logged accepted/rejected-update fraction), *not* fit; the coefficient $\ln 2$ is held fixed at its information-theoretic value; and the coding efficiency η is frozen from the Arm-0 calibration (Section 9.1)—a collapse failure traceable to a mis-calibrated η must be reported as a calibration failure, not refit away. h_{KS} is the projected entropy rate of the tracked basis reference, not a whole-apparatus KS entropy. A failure to collapse may not be rescued by floating an overall capacity scale.
3. **No chaos dependence**: t_{break} is independent of apparatus chaos rate h_{KS}
4. **Wrong functional form**: $V_{\text{IOF}}(t)$ fits ordinary exponential or power-law decoherence but *not* the double-exponential structure $V_{\text{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa_{\text{cal}}t})$
5. **Gravitational scaling in the Penrose-overlap regime**: Coherence timescale tracks mass/geometry ($\propto m^{-2}$ or $\propto s$) and κ_{cal} adds no predictive value after ordinary confounds are controlled—gravitational OR explains the data in this regime
6. **Marginal violation**: Single-party statistics deviate from 50/50, indicating signaling or systematic bias

Falsifying criteria 1–4 would simultaneously contradict the classical data-rate analysis calibrated in Arm 0; a failure there therefore indicates an apparatus or calibration problem before it indicates anything about the framework, and must be investigated as such. The Fisher-capacity bridge is separately weakened or falsified if the operational records reveal an excess position-dependence of $I(\theta)$ over the active range, beyond the finite-V null of Section 11.3, without a correspondingly local $\kappa(\theta)$—in which event the operational tier may stand while the derivation of the binary Born rule via Cencov geometry is falsified [21]. The registered strong-tier module has its own criteria (Section 10.3). These exposure criteria are also platform-independent in an important sense: criteria 1–4 put the framework's operational core at risk without any objective-reduction program attached—no QGEM-class masses, no entanglement witness; a throttleable tracking loop and gate-certified expanding reference dynamics suffice. The mesoscopic platform of the companion protocol [2] is required only for criterion 5's mass-geometry axis.

15 Conclusion

We have presented a finite-rate phase-reference test and its operational benchmark. The measurement basis is treated as a physical reference variable $\theta(t)$ that must be maintained by a

finite-rate tracking loop. When the reference entropy rate exceeds effective tracking capacity, the model predicts exponential growth of basis uncertainty.

The key prediction is operational: in the chaos-wins regime, online observer-relative visibility decays as $V_{\mathrm{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa_{\mathrm{cal}}t})$, with breakdown time inversely proportional to $\kappa_{\mathrm{cal}} = h_{\mathrm{KS}} - \eta\, C_{\mathrm{eff}} \ln 2$. Controlled variation of C_{eff} and h_{KS} against the frozen electronics-only calibration provides a direct benchmark of the finite-rate control boundary; the passive shadow log and the recovery statistic R_{rec} classify the loss; and the parameter-free noisy-reference null fixes what the benchmark can and cannot claim: a successful benchmark is a validated control law, in agreement with standard quantum mechanics by construction.

The reference channel's observable breakdown time t_{break} (intrinsic scale $\tau_\kappa = 1/\kappa_{\mathrm{cal}}$) can overlap with Penrose's τ_{OR} in the mesoscopic regime ($\sim$50–70 ms), but predicts different primary dependencies—tracking bandwidth and reference entropy rate versus mass geometry. Because of that overlap, a single experiment can constrain the joint contributions of finite-rate reference tracking and gravitational objective reduction by varying the appropriate knobs. The discrimination procedure—the κ-off mass-geometry sweep, the crossed design, and the tiered decision rules—is specified in the companion protocol [2], with this paper's benchmark as its calibration arm and the recoverability classifier as the axis that separates reference bookkeeping from collapse-type loss.

If the predicted capacity dependence is observed in a reference-channel demonstrator, basis-reference tracking becomes a calibrated component of observer-relative quantum visibility—a validated control law, not evidence against standard quantum mechanics. The stronger hypothesis—a capacity-dependent loss that survives conditioning on a full-resolution reference log—is excluded for diffusive and externally randomized reference dynamics by existing data and remains registered only for certified-chaotic reference dynamics (Section 10.3), with the prior strongly against it; recoverability against a passive reference log is the classifier throughout.

References

[1] Aernoud Dekker. The ignorant observer: Avidya as a physical limit, 2026. OSF Preprints.

[2] Aernoud Dekker. A capacity-calibrated protocol for testing penrose objective reduction, 2026. OSF Preprints.

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

[4] Howard M. Wiseman and Gerard J. Milburn. *Quantum Measurement and Control*. Cambridge University Press, 2010.

[5] Alexander S. Holevo. *Probabilistic and Statistical Aspects of Quantum Theory*. North-Holland, 1982.

[6] B. Misra and E. C. G. Sudarshan. The Zeno's paradox in quantum theory. *J. Math. Phys.*, 18(4):756–763, 1977. doi: 10.1063/1.523304.

[7] A. G. Kofman and G. Kurizki. Acceleration of quantum decay processes by frequent observations. *Nature*, 405:546–550, 2000. doi: 10.1038/35014537.

[8] Aernoud Dekker. The capacity–backaction frontier: Aggregate rate accounting for fault-tolerant quantum control, 2026. Manuscript.

[9] Sekhar Tatikonda and Sanjoy Mitter. Control under communication constraints. *IEEE Trans. Autom. Control*, 49(7):1056–1068, 2004.

[10] Girish N. Nair and Robin J. Evans. Stabilizability of stochastic linear systems with finite feedback data rates. *SIAM J. Control Optim.*, 43(2):413–436, 2004.

[11] Girish N. Nair, Robin J. Evans, Iven M. Y. Mareels, and William Moran. Topological feedback entropy and nonlinear stabilization. *IEEE Trans. Autom. Control*, 49(9):1585–1597, 2004. doi: 10.1109/TAC.2004.834105.

[12] Christoph Kawan. *Invariance Entropy for Deterministic Control Systems*, volume 2089 of *Lecture Notes in Mathematics*. Springer, 2013.

[13] Alexander L. Fradkov, Boris Andrievsky, and Robin J. Evans. Chaotic observer-based synchronization under information constraints. *Phys. Rev. E*, 73(6):066209, 2006. doi: 10.1103/PhysRevE.73.066209.

[14] Alexander L. Fradkov, Boris Andrievsky, and Robin J. Evans. Controlled synchronization under information constraints. *Phys. Rev. E*, 78(3):036210, 2008. doi: 10.1103/PhysRevE.78.036210.

[15] 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).

[16] 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.

[17] 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.

[18] 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.

[19] 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.

[20] E. R. Clements, M. E. Kim, K. Cui, A. M. Hankin, S. M. Brewer, J. Valencia, J.-S. Chen, D. B. Hume, D. R. Leibrandt, D. J. Wineland, et al. Lifetime-limited interrogation of two independent $^{27}Al^{+}$ clocks using correlation spectroscopy. *Phys. Rev. Lett.*, 125:243602, 2020. doi: 10.1103/PhysRevLett.125.243602.

[21] Aernoud Dekker. The born rule from finite observation: A conditional derivation of the binary born form, 2026. OSF Preprints.

[22] Roger Penrose. On gravity's role in quantum state reduction. *Gen. Relativ. Gravit.*, 28(5): 581–600, 1996.

[23] 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.