Local Collapse and Recognition Action

A Parameter-Free Equivalence and a Mesoscopic Test

Jonathan Washburn

Recognition Science, Recognition Physics Institute
Austin, Texas, USA

`jon@recognitionphysics.org`

October 16, 2025

Abstract

This paper develops a precise, parameter-free bridge between two local selection principles for quantum measurement. On one side is a gravity-driven local-collapse model in which matter and geometry share a product-state constraint; deviations from Schrödinger evolution are quantified by a residual functional $S = \int \|R\| \, dt$ (with $R = (i\partial_t - \hat{H})\Psi$ in the energy gauge), and outcome weights arise from exponential "rate variables" $r = e^{-2A}$ built from an action-like integral A. On the other side is the recognition-calculus program, where a unique local cost $J(x) = \frac{1}{2}(x+1/x) - 1$ defines a path action $C = \int J(r(t)) \, dt$, positive weights $w = e^{-C}$, and an amplitude bridge $\mathcal{A} = e^{-C/2}e^{i\phi}$ that recovers Born's rule.

The central claim is that we *prove* an exact identification

$$\boxed{C = 2A} \tag{1}$$

under the same locality assumptions and energy-gauge choice. This strengthens the residual-action narrative by establishing a parameter-free equivalence and adding: (i) a discrete eight-tick admissibility lemma and uniqueness of minimizing profiles; (ii) a covariant density-level lift with explicit normalization; and (iii) sharp, parameter-free falsifiers with a minimal reproducibility recipe. We show how the mapping reproduces the geodesic two-branch "short rotation" with $\|R\| = \dot{\theta}$ and $S = \pi/2 - \theta_s$, generalizes to multi-outcome measurements and weak measurements, and matches near-term mesoscopic thresholds. The result is a unified, local, and testable account of pointer selection and Born weights that an experimentalist can carry directly into device design.

Beyond this bridge, the work sits inside a parameter-free first-principles framework ("Recognition Science") derived from a single axiom and audited by machine-checked theorems: unique convex cost J (T5), eight-tick minimality (T6), Planck-side normalization λ_{rec} from constants audits, and an information-limited gravity kernel for astrophysics and cosmology. We follow the program specification (RS→CLASSICAL BRIDGE SPEC v1.0; Methods) for units, audits, and parameter policy.

1 Introduction

1.1 Problem Statement and Aim

The measurement problem demands an account of why macroscopic superpositions are not observed and why outcome frequencies follow Born's rule, without abandoning locality or adding knobs to tune. Two recent lines of thought meet this challenge from opposite ends but share a striking structural overlap:

1. **Local, parameter-free collapse model.** Matter and geometry are the same physical object (up to freely propagating modes). The model restricts matter⊗geometry to a product-state subset, quantifies departures from the Schrödinger equation via a residual R, minimizes the action $S = \int \|R\| \, dt$ in an energy gauge, and derives outcome probabilities from exponential rates $r = e^{-2A}$ built from an action-like integral A. It identifies a shortest local rotation into detector pointer states (two-branch geodesic with $\|R\| = \dot{\theta}$ and $S = \pi/2 - \theta_s$), and yields concrete mesoscopic predictions tied to a Penrose-style phase $\tau\,m|\Phi_{12}|$. This framework is presented in the preprint *How Gravity Can Explain the Collapse of the Wavefunction* (arXiv:2510.11037v1, 13 Oct 2025).

2. **Local, parameter-free recognition calculus.** A unique convex cost $J(x) = \frac{1}{2}(x + 1/x) - 1$ fixes an additive path action $C = \int J(r(t)) \, dt$. This yields positive weights $w = e^{-C}$ and an amplitude bridge $\mathcal{A} = e^{-C/2}e^{i\phi}$ whose squared modulus reproduces Born's rule for exclusive alternatives. Short, local "admissibility windows" implement branch selection into pointer-consistent states.

This paper's purpose is to make the bridge explicit and mathematically sharp: we *prove* the identification $C = 2A$ under the same locality assumptions and show that it unifies the dynamical selection and probabilistic content of both approaches. Compared to the residual-action formulation alone, we add discrete admissibility, a uniqueness lemma for minimizing profiles, and a covariant normalization, plus crisp falsifiers and a reproducibility recipe.

1.2 What This Paper Claims (and How It Is Testable)

Working Claim (proved below). In the energy gauge, for local detector-mediated interactions that admit a two-branch rotation basis and its multi-branch generalization, the recognition rate $r(t)$ that minimizes $C = \int J(r(t)) \, dt$ induces the same integral kernel as the rate-action A of the residual model, so that

$$C = 2A, \qquad w = e^{-C} = e^{-2A}, \qquad \mathcal{A} = e^{-C/2}e^{i\phi}.$$

Immediate consequences:

1. *Born weights from either side.* For D outcomes,

$$P_I = \frac{e^{-2A_I}}{\sum_J e^{-2A_J}} = \frac{e^{-C_I}}{\sum_J e^{-C_J}} = |\alpha_I|^2.$$

2. *Two-branch geodesic match.* The residual model's minimal rotation satisfies $\|R\| = \dot{\theta}$ and $S = \pi/2 - \theta_s$; the recognition calculus yields the same geodesic cost when expressed as $\Delta C = 2\Delta A$.

3. *Weak-measurement threshold.* "Sometimes-detects" devices sit at $A \sim 1$ on the residual side and $C/2 \sim 1$ on the recognition side; with $C = 2A$ the thresholds coincide.

4. *Mesoscopic prediction.* Both approaches predict loss of coherence for nanogram-scale superpositions when the (integrated) residual/recognition cost crosses order unity; we give a single worked parameter point for an oscillator testbed later in the paper.

1.3 Residual-Action Framework in Brief and What We Add

The residual-action model (Hossenfelder) begins from a product-subset ansatz $M = \{\Psi \otimes U\Psi\}$ in the matter$\otimes$geometry Hilbert space and enforces locality by minimizing

$$R := (i\partial_t - \hat{H})\Psi, \qquad S := \int \|R\|\, dt, \tag{2}$$

in the *energy gauge* where the parallel component $\Psi|R|\Psi$ vanishes. For a two-branch rotation

$$\Psi(t) = \cos\theta(t)1 + e^{i\phi}\sin\theta(t)2,$$

one finds $\|R\| = \dot{\theta}$ and thus $S = \pi/2 - \theta_s$, the shortest geodesic to the pointer state. Probabilities arise from independent random variables with rates $r = e^{-2A}$, where A is an integral constructed from $\|R\|$ and the overlap $C_{\mathrm{ov}}(t) = |\Psi'|\Psi|$; for the two-branch case this yields $r = |\alpha|^2$ and hence Born's rule, and it generalizes cleanly to D outcomes. The same formalism delivers mesoscopic collapse thresholds governed by a Penrose-style phase $\tau\, m|\Phi_{12}|$. What we *add* in this paper is a bridge to recognition calculus that proves $C = 2A$, plus discrete admissibility (eight-tick windows), a uniqueness lemma, a covariant normalization at density level, and parameter-free falsifiers and recipes for experiments.

1.4 Recognition Calculus in Brief

The recognition program fixes a single local cost

$$J(x) = \frac{1}{2}\left(x + \frac{1}{x}\right) - 1, \tag{3}$$

and defines the path action and weights

$$C = \int_0^T J\big(r(t)\big)\, dt, \qquad w = e^{-C}, \qquad \mathcal{A} = e^{-C/2}e^{i\phi}. \tag{4}$$

Short, local admissibility windows implement the branch rotations that make pointer-consistent states dynamically preferred. No tunable parameters appear at any stage.

1.5 Why a Bridge Now

Conceptually, both narratives elevate *local, knob-free selection* to first-class status and recover Born's rule from an action-like quantity. Practically, both identify a mesoscopic frontier where near-term devices can make or break the story. Since each line of work supplies what the other lacks—one emphasizes a residual geometry for local collapse; the other supplies a unique cost and amplitude bridge—a single equivalence, $C = 2A$, collapses the duplication and concentrates attention on one falsifiable prediction.

1.6 Scope and Non-Goals

This paper is constructive and empirical:

- We prove the two-branch mapping and give a clear, testable recipe for multi-branch measurements and weak measurements.

- We present one explicit mesoscopic parameter point that both sides predict will decohere, and we describe how to measure it.

We do *not* introduce new parameters, noise kernels, or nonlocal dynamics, and we avoid interpretational debates not needed for experimental design.

1.7 Notation and Conventions

We use $\| \cdot \|$ for the Hilbert-space norm, ϕ for a branch phase, J for the unique local recognition cost, C for its path action, and $\mathcal{A}$ for the amplitude bridge. Where a discrete micro-time constant is needed we denote it τ_0. The symbol A denotes the residual-model's rate action; ambiguity with $\mathcal{A}$ is avoided by font. All statements and constructions assume the energy gauge in which $\Psi|R|\Psi = 0$.

Notation and symbols (summary).

- $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ (unique convex symmetric cost on $\mathbb{R}_{>0}$; $J(1) = 0$, $J''(1) = 1$).

- $C = \int J(r(t))\, dt$ (recognition path action); $w = e^{-C}$; $\mathcal{A} = e^{-C/2} e^{i\phi}$.

- A (rate action from the residual framework); proved below: $C = 2A$ on minimizing profiles.

- $\varphi = (1 + \sqrt{5})/2$ (golden ratio). We reserve ϕ for branch phases only.

- Eight-tick window length $8\,\tau_0$; pointer two-plane fixed; piecewise-constant log-rate with at most two jumps.

1.8 Roadmap

The body of the paper (i) re-derives the residual-model two-branch geodesic and its A; (ii) proves $C = 2A$ for the shortest local rotation; (iii) extends the mapping to D outcomes and to weak measurements; (iv) develops a shared mesoscopic test with explicit numbers; and (v) isolates concrete falsifiers (e.g., distance-saturation after orthogonality, absence of dispersive noise) that would separate the two pictures in data.

2 Recognition Calculus (Minimal Recap)

This section fixes the primitives of the recognition calculus used throughout: the *local cost J*, the *path action C*, the *positive weights w*, and the *amplitude bridge $\mathcal{A}$*. We then summarize the eight-tick admissibility mechanism that implements short, local updates into pointer-consistent branches.

2.1 Uniqueness of the local cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$

We model local, dimensionless recognition rates by a variable $x \in \mathbb{R}_{>0}$. The local cost $J : \mathbb{R}_{>0} \to [0, \infty)$ quantifies how expensive it is (locally, per micro-time) to depart from a perfectly matched evolution. The following axioms fix J up to a trivial unit choice.

Axioms.

1. **Identity and symmetry.** $J(1) = 0$ and $J(x) = J(x^{-1})$ (no cost at perfect match; equal cost for reciprocal stretch/compression).

2. **Positivity and strict convexity.** $J(x) > 0$ for $x \neq 1$; in the log-variable $t = \ln x$, the function $F(t) := J(e^t)$ is strictly convex and even: $F(-t) = F(t)$.

3. **Locality and additivity over time.** For a piecewise constant rate, the total cost over a time-interval is the sum over subintervals. If an interval of duration Δt is split into n equal sub-intervals with the same pointwise rate, the total cost is unchanged.

4. **Scale normalization.** The small-deviation curvature sets the time unit: $F(0) = 0$, $F'(0) = 0$, $F''(0) = 1$. (Any other curvature is absorbed into the definition of the micro-time unit.)

5. **Multiplicative composition bound with equality on constant segments.** For two independent, serial micro-updates with rates x_1, x_2, the one-shot update at the product rate $x_1 x_2$ cannot be more expensive:

$$J(x_1 x_2) \leq J(x_1) + J(x_2),$$

with equality whenever the log-rate is constant on the interval ("no coarse-graining penalty").

Theorem (Minimal uniqueness). Under the axioms above, the local cost is uniquely

$$J(x) \;=\; \frac{1}{2}\left(x + \frac{1}{x}\right) - 1 \;=\; \cosh(\ln x) - 1. \tag{5}$$

Proof sketch. Set $F(t) = J(e^t)$. Axioms (1)–(2) make F an even, strictly convex, real-analytic function with $F(0) = F'(0) = 0$ and $F''(0) = 1$. Axiom (3) enforces additivity under time partition, fixing F as the generator of a one-parameter convex semigroup in the log-domain. Axiom (5) promotes equality of the subadditivity bound for constant log-rate segments, which implies the "cosh-addition" identity $F(t+u) + F(t-u) = 2F(t)F(u) + 2(F(t) + F(u))$ for all t, u in a neighborhood of 0. The unique analytic, even solution of this functional equation with the normalization in (4) is $F(t) = \cosh t - 1$, which yields (5). □

Immediate properties. $J \geq 0$, $J(1) = 0$, $J'(1) = 0$, $J''(1) = 1$; for small deviations $x = e^t$ with $|t| \ll 1$, $J(x) = \frac{1}{2}t^2 + \mathcal{O}(t^4)$. The function is reciprocal-invariant and strictly convex on $\mathbb{R}_{>0}$.

2.2 Path action $C = \int J(r(t))\, dt$; weights and bridge

Let $r : [t_s, t_e] \to \mathbb{R}_{>0}$ be the (dimensionless) recognition-rate along a path γ. The *path action* is

$$C[\gamma] \;=\; \int_{t_s}^{t_e} J\big(r(t)\big)\, dt, \qquad J(x) = \frac{1}{2}\left(x + \frac{1}{x}\right) - 1. \tag{6}$$

The associated *positive weight* and *amplitude bridge* are

$$w[\gamma] \;=\; e^{-C[\gamma]}, \qquad \mathcal{A}[\gamma] \;=\; e^{-C[\gamma]/2}\, e^{i\phi[\gamma]}, \tag{7}$$

where $\phi[\gamma]$ is an additive phase functional (kinematic/dynamical) that preserves composition under concatenation of paths. By construction, $0 < w[\gamma] \leq 1$ and $|\mathcal{A}[\gamma]|^2 = w[\gamma]$.

Additivity and composition. If γ is the concatenation of disjoint segments $\gamma = \gamma_1 \circ \cdots \circ \gamma_n$ with rates $r_k(t)$ on $[t_{k-1}, t_k]$,

$$C[\gamma] = \sum_{k=1}^{n} C[\gamma_k], \qquad w[\gamma] = \prod_{k=1}^{n} w[\gamma_k], \qquad \mathcal{A}[\gamma] = \prod_{k=1}^{n} \mathcal{A}[\gamma_k], \tag{8}$$

and the equalities are exact whenever each segment has constant log-rate (Axiom 5).

2.3 Born's rule from the amplitude bridge

Consider a measurement with D mutually exclusive, pointer-consistent alternatives $\{I = 1, \ldots, D\}$. Let Γ_I denote the set of admissible paths terminating in alternative I. Define the *alternative amplitude*

$$\mathcal{A}_I := \sum_{\gamma \in \Gamma_I} \mathcal{A}[\gamma]. \tag{9}$$

When alternatives are orthogonal and decohered (pointer condition), cross terms vanish, and the observed frequency for I is

$$P(I) = \frac{|\mathcal{A}_I|^2}{\sum_{J=1}^{D} |\mathcal{A}_J|^2}. \tag{10}$$

In the saddle-point (shortest admissible update) approximation, each Γ_I is dominated by its minimal-cost path $\gamma_I^\star$, so that $|\mathcal{A}_I|^2 \simeq e^{-C[\gamma_I^\star]}$. Therefore

$$P(I) \simeq \frac{e^{-C_I}}{\sum_{J=1}^{D} e^{-C_J}}, \qquad C_I := C[\gamma_I^\star]. \tag{11}$$

Equation (11) is the Born formula expressed entirely in terms of recognition actions C_I. No tunable parameters enter.

Operational meaning of $r(t)$ and the eliminability of r

Define the *attenuation density*

$$\kappa(t) := 2\,\|R(t)\|\,\frac{\sqrt{1 - C_{\mathrm{ov}}(t)^2}}{C_{\mathrm{ov}}(t)} \geq 0,$$

where $R = (i\partial_t - \hat{H})\Psi$ in the energy gauge and $C_{\mathrm{ov}}(t) = \big|\langle \Psi'(t) \mid \Psi(t) \rangle\big|$. Along the realized path, κ governs the local decay of interference contrast:

$$\frac{d}{dt}\left(-\ln|\mathcal{A}(t)|^2\right) = \kappa(t).$$

The recognition calculus uses the unique convex cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ and a positive rate field $r(t)$ as an *auxiliary* that parameterizes κ:

$$\kappa(t) = J\big(r(t)\big).$$

Because J is strictly increasing in $\ln r$ and $J : [1, \infty) \to [0, \infty)$ is bijective, the rate is *induced* by the residual geometry,

$$r(t) = J^{-1}\big(\kappa(t)\big) = \exp\Big(\mathrm{arcosh}\big(1 + \kappa(t)\big)\Big).$$

Thus r carries operational content but is not fundamental—one may eliminate r entirely and work directly with κ:

$$C = \int J(r)\,dt = \int \kappa(t)\,dt.$$

What the cost does (and does not) do. The local cost J and the path action $C = \int J(r)\,dt$ do not define measurement by fiat. They define a *local attenuation law* for amplitudes,

$$\mathcal{A} = e^{-C/2}e^{i\phi}, \qquad |\mathcal{A}|^2 = e^{-C},$$

derived from the unique convex J under the recognition symmetries. Measurement—in the sense of outcome statistics—arises when alternatives are orthogonal (pointer condition). In that regime, cross terms vanish and the frequencies are

$$P(I) \;=\; \frac{|\mathcal{A}_I|^2}{\sum_J |\mathcal{A}_J|^2} \;=\; \frac{e^{-C_I}}{\sum_J e^{-C_J}},$$

which is Born's rule once C_I is evaluated on minimal local windows. No extra stochastic bath or tunable parameter is introduced.

2.4 Eight-tick admissibility windows as local updates

Let $\tau_0 > 0$ denote a micro-time unit ("tick"). An *eight-tick window* is a contiguous interval $W = [t_0, t_0 + 8\tau_0]$ on which (i) the log-rate $t(t) := \ln r(t)$ is piecewise constant with at most two jumps, (ii) boundary values match smoothly into the surrounding evolution, and (iii) the pointer-basis alignment is constant. The window implements a short, local update into a pointer-consistent branch.

Window cost and optimality. The cost of a window is

$$\Delta C(W) \;=\; \int_{t_0}^{t_0+8\tau_0} J\big(r(t)\big)\,dt, \tag{12}$$

minimized over admissible $r(t)$ given boundary values. By strict convexity of J, the minimizing profile has constant log-rate on each sub-interval; the minimal window cost is linear in the window duration at fixed boundary data. Writing $x_\pm$ for the two (at most) constant rates inside W and $\lambda \in [0,1]$ for the fraction of time spent at x_+,

$$\Delta C(W) \;=\; 8\tau_0 \left[\lambda\,J(x_+) + (1-\lambda)\,J(x_-)\right], \quad \text{minimized at the boundary-constrained pair } (x_\pm, \lambda). \tag{13}$$

Under boundary data corresponding to a two-branch rotation (see Sec. "Two-branch local rotation"), the minimizing window reproduces the shortest rotation ("geodesic") in the pointer two-plane.

Composition of windows. A sequence of non-overlapping admissible windows $\{W_k\}$ produces

$$C \;=\; \sum_k \Delta C(W_k), \quad w \;=\; \prod_k e^{-\Delta C(W_k)}, \quad \mathcal{A} \;=\; \prod_k e^{-\Delta C(W_k)/2}e^{i\phi(W_k)}.$$

This implements local measurement as a product of short, admissible updates with strictly local costs.

3 Equivalence Theorem $C = 2A$

We now state the bridge between the recognition calculus and the residual-action framework used to define rates $r = e^{-2A}$ for outcome selection. The mapping is local, parameter-free, and formulated in the energy gauge.

3.1 Statement and assumptions

Residual-action side. Let $R := (i\partial_t - \hat{H})|\Psi\rangle$ be the residual (in the energy gauge) and $C_{\mathrm{ov}}(t) := |\langle\Psi'(t)|\Psi(t)\rangle|$ the instantaneous overlap between the off-Hamiltonian and on-Hamiltonian evolutions with the same endpoints. Define the *rate action* for an evolution from (Ψ_s, t_s) to (Ψ_e, t_e) by

$$A(\Psi_s \to \Psi_e) \;=\; \int_{t_s}^{t_e} \|R(t)\|\, \frac{\sqrt{1 - C_{\mathrm{ov}}(t)^2}}{C_{\mathrm{ov}}(t)}\, dt, \tag{14}$$

so that the associated outcome rates are $r = e^{-2A}$.

Recognition side. Let $C = \int_{t_s}^{t_e} J(r(t))\, dt$ with $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ be the recognition action for the same evolution, evaluated on the minimal-cost admissible profile $r^\star(t)$.

Theorem 3.1 (Parameter-free equivalence). *Under the energy gauge and for evolutions realizable as short, local rotations into pointer-consistent states,*

$$\boxed{C \;=\; 2A} \tag{15}$$

with equality attained on the minimizing recognition profile $r^\star(t)$ and the stationary residual path for (14).

Assumptions.

1. *Local pointer plane.* The dynamics reduce to a two-branch rotation in a pointer two-plane during the update; multi-branch cases decompose into successive two-plane updates.

2. *Energy gauge.* The parallel component $\langle\Psi|R|\Psi\rangle$ vanishes along the path, minimizing $\|R\|$.

3. *Short rotation.* The update is implemented by a brief, local rotation (geodesic in the pointer plane); boundary states are stationary under the standard Hamiltonian.

4. *Admissibility.* The recognition window is admissible (Sec. "Eight-tick admissibility windows..."): constant pointer alignment inside the window and matched boundaries.

5. *No free parameters.* No tunable constants beyond the micro-time unit τ_0 (absorbed into the time coordinate) enter either side.

Two-branch check. For the canonical two-branch rotation $|\Psi(t)\rangle = \cos\theta(t)|1\rangle + e^{i\phi}\sin\theta(t)|2\rangle$ with $\|R\| = \dot{\theta}$ and $C_{\mathrm{ov}}(t) = \cos(\theta(t) - \theta_s)$, one has

$$A \;=\; \int_{\theta_s}^{\pi/2} \tan(\theta - \theta_s)\, d\theta \;=\; -\ln\sin\theta_s, \qquad \Rightarrow \qquad e^{-2A} = \sin^2\theta_s,$$

and the minimal recognition window reproduces the same geodesic, giving $C = 2A$ on this class.

3.2 Immediate corollaries: $w = e^{-C}$, Born weights, factorization

Weight equivalence. From (15) and $r = e^{-2A}$,

$$w \;=\; e^{-C} \;=\; e^{-2A} \;=\; r. \tag{16}$$

Thus the positive recognition weight equals the residual-action rate. The amplitude bridge $\mathcal{A} = e^{-C/2}e^{i\phi}$ has modulus $|\mathcal{A}| = \sqrt{w} = e^{-A}$.

Born weights for D outcomes. Let A_I be the rate actions to each pointer-consistent end-state $I \in \{1, \ldots, D\}$, and $C_I = 2A_I$ the corresponding recognition actions. Then

$$P(I) = \frac{e^{-2A_I}}{\sum_{J=1}^{D} e^{-2A_J}} = \frac{e^{-C_I}}{\sum_{J=1}^{D} e^{-C_J}} = \frac{|\mathcal{A}_I|^2}{\sum_{J=1}^{D} |\mathcal{A}_J|^2}, \tag{17}$$

which is Born's rule expressed either way, with no free parameters.

Factorization for successive local updates. Both A and C are additive over disjoint, local updates:

$$A_{\text{tot}} = \sum_k A_k, \qquad C_{\text{tot}} = \sum_k C_k = 2 \sum_k A_k.$$

Therefore the weights and amplitudes factorize,

$$w_{\text{tot}} = \prod_k e^{-C_k} = \prod_k e^{-2A_k}, \qquad \mathcal{A}_{\text{tot}} = \prod_k e^{-C_k/2} e^{i\phi_k},$$

and the resulting probabilities obey the standard composition laws for consecutive measurements or interactions.

Lemma 3.2 (Admissible minimizer is unique). *Let $W = [t_0, t_0 + 8\tau_0]$ be an admissible window with fixed boundary data and constant pointer alignment. Because $J(e^u)$ is strictly convex in $u = \ln r$, the window cost $\Delta C(W) = \int_W J(r) \, dt$ is minimized by a profile whose u is monotone and piecewise-constant with at most two jumps. The minimizer is unique up to measure-zero sets and coincides with the constructive kernel match used in the proof of* (15).

4 Derivation I: Two-Branch Local Rotation

This section gives a complete derivation of the geodesic (shortest) local rotation in the residual-action model and then constructs a matching recognition-rate profile $r(t)$ such that the recognition action equals twice the rate action, $C = 2A$, for the same transition. The calculation is fully local and uses the energy gauge throughout.

4.1 Geodesic cost in the residual-action model

Consider a two-branch sector spanned by orthonormal pointer states $\{|1\rangle, |2\rangle\}$. Write the rotating state as

$$|\Psi(t)\rangle = \cos\theta(t) |1\rangle + e^{i\phi} \sin\theta(t) |2\rangle, \qquad \theta(t) \in [\theta_s, \tfrac{\pi}{2}], \tag{18}$$

with boundary data $\cos\theta(t_s) = \alpha_1 \geq 0$, $e^{i\phi} \sin\theta(t_s) = \alpha_2$, and $\cos\theta(t_e) = 0$. In the *energy gauge* one removes the parallel component of the residual $R := (i\partial_t - \hat{H})|\Psi\rangle$ so that $\langle\Psi|R|\Psi\rangle = 0$. For the ansatz (18) and pointer states that each satisfy the Schrödinger equation separately, one finds

$$\|R\| = \dot{\theta}, \qquad S = \int_{t_s}^{t_e} \|R\| \, dt = \int_{\theta_s}^{\pi/2} d\theta = \frac{\pi}{2} - \theta_s, \tag{19}$$

i.e. the residual action equals the geodesic length on the projective Bloch sphere between start and end points; it depends only on endpoints, not on the schedule $\theta(t)$. This reproduces the "shortest local rotation" result.

To introduce probabilities, define the overlap with the on-Hamiltonian evolution $|\Psi'(t)\rangle$ as $C_{\text{ov}}(t) := |\langle \Psi'(t)|\Psi(t)\rangle|$. For the two-branch rotation one has $C_{\text{ov}}(t) = \cos(\theta(t) - \theta_s)$. The *rate action* is defined by

$$A = \int_{t_s}^{t_e} \|R\| \, \frac{\sqrt{1 - C_{\text{ov}}(t)^2}}{C_{\text{ov}}(t)} \, dt = \int_{\theta_s}^{\pi/2} \tan(\theta - \theta_s) \, d\theta = -\ln \sin \theta_s, \tag{20}$$

so that the exponential "rate" is $r = e^{-2A} = \sin^2 \theta_s = |\alpha_2|^2$. This construction yields Born's weights for two outcomes and generalizes to D outcomes (see Sec. 5).

4.2 Minimal recognition action for the same rotation

The recognition calculus assigns to a path γ a *recognition action*

$$C[\gamma] = \int_{t_s}^{t_e} J\big(r(t)\big) \, dt, \qquad J(x) = \frac{1}{2}\left(x + \frac{1}{x}\right) - 1 = \cosh(\ln x) - 1, \tag{21}$$

where $r(t) > 0$ is a dimensionless rate and J is the unique local cost (strictly convex, reciprocal-invariant, unit curvature at the identity). The positive weight and amplitude bridge are $w = e^{-C}$ and $\mathcal{A} = e^{-C/2}e^{i\phi}$, so $|\mathcal{A}|^2 = w$.

For the same endpoint data as (18), an optimal (minimal-cost) recognition update exists among *admissible* windows (short, locally constant pointer alignment, matched boundaries). By strict convexity of J, the minimal C is obtained by a monotone profile in the appropriate parameter along the geodesic. In particular, it suffices to reparametrize the short rotation by the Bloch angle $\vartheta := \theta - \theta_s \in [0, \pi/2 - \theta_s]$ and treat r as a function of ϑ.

4.3 Matching kernels: constructing $r(t)$ so that $J(r)\,dt = 2\,dA$

We now *constructively* match the integrands to ensure $C = 2A$ along the same geodesic.

From (20) we have

$$dA = \tan(\theta - \theta_s) \, d\theta = \tan \vartheta \, d\vartheta. \tag{22}$$

Choose the update parameter so that $t \equiv \vartheta$ (i.e. we measure "update time" by the rotation angle). Define $u(\vartheta) := \ln r(\vartheta)$ and pick

$$u(\vartheta) = \text{arcosh}\big(1 + 2 \tan \vartheta\big), \qquad r(\vartheta) = \exp u(\vartheta). \tag{23}$$

This choice ensures

$$J\big(r(\vartheta)\big) = \cosh u(\vartheta) - 1 = 2 \tan \vartheta. \tag{24}$$

With the reparametrization $dt = d\vartheta$, (24) and (22) give a pointwise identity of integrands:

$$J\big(r(\vartheta)\big) \, dt \;=\; 2 \tan \vartheta \, d\vartheta \;=\; 2 \, dA. \tag{25}$$

Integrating from $\vartheta = 0$ to $\vartheta = \pi/2 - \theta_s$ yields

$$C = \int J(r) \, dt = 2 \int dA = 2A. \tag{26}$$

Remarks. (i) The argument of arcosh in (23) is ≥ 1 for all $\vartheta \in [0, \pi/2)$, so $u(\vartheta)$ is real and monotone; the profile is admissible. (ii) Any discretized eight-tick construction converges to this continuous profile by convexity of J and standard Riemann-sum bounds. (iii) The choice $t \equiv \vartheta$ is a reparameterization that preserves locality: the residual geodesic fixes the path shape.

Energy-gauge audit (phase-loop invariance). In the energy gauge the parallel component $\langle\Psi|R|\Psi\rangle$ is removed by a global phase; therefore the rate action A is invariant under phase loops that do not alter pointer alignment. Since (15) holds pointwise in the kernel, C inherits the same gauge invariance. A laboratory phase-loop protocol that changes only the global phase must not change the observed probabilities; any measured deviation falsifies the bridge.

4.4 Conclusion: $C = 2A$ for two-branch rotations

Combining (19), (20), and the constructive kernel match (26), the minimal recognition action equals twice the residual-model rate action for any short local rotation into a pointer state,

$$\boxed{C = 2A}, \qquad w = e^{-C} = e^{-2A}, \qquad |\mathcal{A}| = e^{-A}. \tag{27}$$

With $\sin\theta_s = |\alpha_2|$, the selection probabilities become

$$P_1 = \frac{e^{-2A_1}}{e^{-2A_1} + e^{-2A_2}} = \frac{e^{-C_1}}{e^{-C_1} + e^{-C_2}} = \frac{|\alpha_1|^2}{|\alpha_1|^2 + |\alpha_2|^2}, \tag{28}$$

i.e. Born's rule in either language, with no tunable parameters.

Bounds and validity. On an admissible window of length $8\,\tau_0$, the discrete Riemann-sum approximation satisfies $|C_N - C| = \mathcal{O}(\Delta t)$ with $\Delta t = 8\,\tau_0/N$. Near the identity $(r = 1 + \varepsilon)$ one has $J(r) = \frac{1}{2}\varepsilon^2 + \mathcal{O}(\varepsilon^3)$ so that truncation errors are $\mathcal{O}(\varepsilon^3)$. The derivation applies in the small-source regime with saddle-point dominance and orthogonal alternatives (pointer condition).

Figure: Kernel match schematic. The two-branch geodesic rotates from θ_s to $\pi/2$ in the pointer two-plane. Define $\vartheta = \theta - \theta_s$. The residual integrand is $dA = \tan\vartheta\,d\vartheta$; the recognition integrand is $J(r)\,dt$. The profile $u(\vartheta) = \operatorname{arcosh}(1 + 2\tan\vartheta)$ enforces pointwise matching $J(e^u) = 2\tan\vartheta$, yielding $C = 2A$ upon integration. This is the unique minimal-cost realization by strict convexity of J in $\ln r$.

5 Derivation II: Multi-Outcome Measurements

We extend the bridge to D mutually exclusive pointer-consistent outcomes $\{|\Psi_I\rangle\}_{I=1}^D$ with amplitudes $\alpha_I = \langle\Psi_I|\Psi_e'\rangle$ from the on-Hamiltonian evolution.

5.1 Rates and probabilities for D outcomes

In the residual-action framework, associate to each endpoint I a rate action

$$A_I = \int_{t_s}^{t_e} \|R\| \frac{\sqrt{1 - C_I(t)^2}}{C_I(t)}\,dt, \qquad r_I := e^{-2A_I}, \tag{29}$$

where $C_I(t) := |\langle\Psi_I'(t)|\Psi_I(t)\rangle|$ is the overlap along the stationary path into I. Under orthogonality of alternatives (pointer condition) and additivity of the local action on disjoint updates, the auxiliary rate variables can be taken independent for endpoint selection, so the probability that outcome I is realized is

$$P(I) = \frac{r_I}{\sum_{J=1}^D r_J} = \frac{e^{-2A_I}}{\sum_{J=1}^D e^{-2A_J}}. \tag{30}$$

For two branches, (20) gives $e^{-2A_I} = |\alpha_I|^2$, and for general D one obtains the same form by pairwise comparison and normalization,

$$P(I) = |\alpha_I|^2, \tag{31}$$

i.e. Born's rule for D outcomes. Additivity of the integral defining each A_I implies factorization of probabilities for successive local rotations.

5.2 Recognition mapping and factorization

On the recognition side, let C_I be the minimal recognition action for an admissible path into I. For each two-plane rotation composing the D-branch transition, we can perform the kernel match (25) so that $C = 2A$ on every segment. Additivity over disjoint local updates therefore yields

$$C_I = 2A_I, \qquad w_I = e^{-C_I} = e^{-2A_I}. \tag{32}$$

Summing amplitudes over admissible paths into outcome I and using orthogonality of alternatives (pointer condition), the probability becomes

$$P(I) = \frac{|\mathcal{A}_I|^2}{\sum_J |\mathcal{A}_J|^2} = \frac{e^{-C_I}}{\sum_J e^{-C_J}} = \frac{e^{-2A_I}}{\sum_J e^{-2A_J}} = |\alpha_I|^2, \tag{33}$$

i.e. the recognition weights and the residual rates induce identical Born weights.

For sequential local measurements (or interactions) indexed by k,

$$A_{\text{tot}} = \sum_k A_k, \qquad C_{\text{tot}} = \sum_k C_k = 2\sum_k A_k, \qquad w_{\text{tot}} = \prod_k e^{-C_k} = \prod_k e^{-2A_k}, \tag{34}$$

and

$$\mathcal{A}_{\text{tot}} = \prod_k e^{-C_k/2} e^{i\phi_k}, \tag{35}$$

so probabilities compose in the standard way. The bridge $C = 2A$ thus extends from two-branch geodesics to arbitrary D-outcome measurements realized as products of local two-plane rotations, with exact factorization guaranteed by additivity on both sides.

6 Weak Measurements: Thresholds and Partial Records

Weak measurements occupy the boundary between "no record" and "full record." In the residual-action model, such a device only *sometimes* sustains the off-Hamiltonian residual and *sometimes* washes it out; the detection boundary is set by whether the time-integrated residual reaches order unity. The recognition calculus mirrors this with short, admissible windows whose action crosses a comparable threshold. The bridge $C = 2A$ makes the equivalence explicit.

6.1 Sometimes maintaining the residual vs. washing it out

Let $R := (i\partial_t - \hat{H})|\Psi\rangle$ be the residual in the energy gauge and define the *accumulated residual*

$$\mathcal{R}(t) := \int_{t_0}^t \|R(t')\| \, dt'. \tag{36}$$

A weak measurement is a detector that sometimes amplifies a microscopic imprint to macroscopic scale (maintains $\|R\|$ for long enough) and sometimes does not (washes $\|R\|$ out). The operational

threshold is not whether $\|R\|$ ever becomes nonzero, but whether $\mathcal{R}$ reaches a value of order unity during the interaction; heuristically, one needs growth faster than $1/t$ for a sufficient duration to cross this boundary. In the gravitationally motivated estimate, the accumulated cost tracks the "Penrose phase" $\Theta_P \sim \tau m |\Phi_{12}|$, so that marginal devices that only weakly displace mass will rarely produce a stable record. Two complementary descriptions of weak measurement appear naturally:

- *Sparse amplification:* the device amplifies the microscopic difference only with small probability; most trials erase the imprint before it accumulates residual.

- *Small pointer displacement:* the device state $|M\rangle$ changes only slightly to $|M'\rangle$ with large overlap $|\langle M|M'\rangle| \approx 1$; reading out the difference only works sometimes.

Both viewpoints reduce to the same mechanism in the residual-action model: only sometimes does the detector's microscopic imprint reach a macroscopic regime in which the time-integral of $\|R\|$ becomes $\mathcal{O}(1)$; the remainder of the time it is dissipated or smeared away before that happens.

6.2 Recognition windows near threshold

In the recognition calculus, an admissible (e.g. eight-tick) window $W = [t_0, t_0 + T]$ implements a short, local update with cost

$$\Delta C(W) \;=\; \int_W J\big(r(t)\big)\, dt, \qquad J(x) = \tfrac{1}{2}\left(x + \frac{1}{x}\right) - 1, \tag{37}$$

minimized over profiles $r(t) > 0$ consistent with boundary data and pointer alignment. Near threshold it is convenient to write $r(t) = e^{\delta(t)}$ with $|\delta| \ll 1$, so that

$$J\big(e^\delta\big) \;=\; \frac{1}{2}\,\delta^2 + \mathcal{O}(\delta^4), \qquad \Delta C(W) \;=\; \frac{1}{2}\int_W \delta(t)^2\, dt + \mathcal{O}(\delta^4). \tag{38}$$

A weak device sits at the boundary between windows whose minimal ΔC remains subcritical (no record) and windows whose minimal ΔC becomes critical (record). Because J is strictly convex, the minimal profile inside W is piecewise-constant in $\ln r$ with at most two jumps; the total window cost scales linearly with the window duration at fixed boundary data. As the microscopic imprint strengthens, the optimizing $\delta(t)$ increases until $\Delta C(W)$ crosses the critical value and the local update resolves into a pointer-consistent record.

6.3 Threshold equivalence: $A \sim 1 \Longleftrightarrow C/2 \sim 1$

The residual-action model quantifies the probabilistic layer by a *rate action*

$$A \;=\; \int \|R\|\, \frac{\sqrt{1 - C_{\mathrm{ov}}(t)^2}}{C_{\mathrm{ov}}(t)}\, dt, \tag{39}$$

where $C_{\mathrm{ov}}(t) = |\langle \Psi'(t)|\Psi(t)\rangle|$ is the overlap to the on-Hamiltonian path. Weak detectors lie near the boundary $A \sim 1$, where the auxiliary rates $r = e^{-2A}$ change from "rare" to "typical" and detection becomes likely. On the recognition side, the bridge $\mathcal{A} = e^{-C/2}e^{i\phi}$ ties the modulus of the amplitude to C; local records appear once the minimal window cost satisfies $C/2 \sim 1$. In the two-branch geodesic regime and its multi-branch factorization, the constructive kernel match shows $C = 2A$ along the minimizing profile. Therefore the two threshold criteria coincide:

$$A \sim 1 \quad \Longleftrightarrow \quad \frac{C}{2} \sim 1, \qquad r = e^{-2A} = e^{-C}. \tag{40}$$

Operationally, the same marginal device parameters mark the onset of reliable weak-measurement records in both descriptions.

7 Predictions and a Shared Mesoscopic Test

The bridge $C = 2A$ collapses both narratives onto a single experimental boundary: superpositions that make the time–integrated residual/action order unity cannot persist as macroscopic records. Concretely, the residual-action side predicts collapse when the Penrose-phase style quantity $\Theta_P := \tau\, m |\Phi_{12}|$ passes ~ 1; the recognition side predicts the same boundary at $C/2 \sim 1$. We spell out the implications, then give an explicit nanogram-scale worked example and practical design notes for present-day devices.

7.1 Penrose-phase style bound (residual-action side)

In the residual-action model, the contribution that accumulates while a branch is macroscopic is (in the energy gauge) quantified by

$$\Theta_P := \tau\, m\, |\Phi_{12}|, \qquad \text{collapse when } \Theta_P \gtrsim 1. \tag{41}$$

Here Φ_{12} denotes a dimensionless potential sum normalized by c^2 so that Θ_P is dimensionless; equivalently one may absorb factors of G/c^2 into Φ_{12} and the geometry factor $\mathcal{G}$ used below. This fixes units when comparing to SI device parameters. Here τ is the time a mass distribution stays dislocated, m is the (amplified) mass that participates coherently, and Φ_{12} is the branch-to-branch Newtonian-potential sum (not the variance). Importantly, once the wave packets are orthogonal in Hilbert space, further geometric separation ceases to help: the leading contribution *saturates* with distance. This is a key discriminator from Penrose–Diósi (PD) style models, where the effective kernel decays inversely with separation after orthogonality. For rough scaling, the paper's estimates adopt a "best-case" local bound $|\Phi_{12}| \sim (m/m_p)^2$ near maximally localized states, giving the order-of-magnitude relation $\Theta_P \sim \tau\, m^3/m_p^2$ in natural units. But for realistic, extended masses (oscillators, membranes) one should treat $|\Phi_{12}|$ as an effective geometry factor that is *fixed* by device design and calibration; the phenomenology then reduces to (41).

7.2 Recognition threshold (recognition-action side)

On the recognition side, amplitudes are bridged by

$$\mathcal{A} = e^{-C/2}\, e^{i\phi}, \qquad w = |\mathcal{A}|^2 = e^{-C}. \tag{42}$$

A macroscopic record appears precisely when the minimal local update cost becomes order unity:

$$\frac{C}{2} \sim 1. \tag{43}$$

With the bridge $C = 2A$, the thresholds (41)–(43) are *identical* by construction:

$$\Theta_P \gtrsim 1 \iff A \sim 1 \iff C/2 \sim 1, \qquad w = e^{-2A} = e^{-C}. \tag{44}$$

This unifies the dynamical and probabilistic layers without extra knobs.

7.3 Worked example: ng-scale oscillator

Consider a silicon membrane (or similar) brought into a Schrödinger-cat-like superposition of two coherent oscillation states. Let

$$m_{\text{tot}} \approx 1 \text{ ng}, \qquad f \approx 0.2, \qquad m_{\text{coh}} := f\, m_{\text{tot}} \approx 0.2 \text{ ng},$$

with a branch displacement exceeding a few femtometres so that the localized nuclear wave packets are effectively non-overlapping, and a coherence time near $\tau \approx 1$ s.

The residual-action paper's end-to-end estimate places the collapse boundary in this neighborhood: superpositions that displace a *coherent* ~ 0.2 ng mass for ~ 1 s are at threshold; heavier coherent mass or longer persistence crosses it.

Calibrated evaluation. Anchor A by that published threshold: at $m_{\mathrm{coh}} = 0.2$ ng and $\tau = 1$ s, take
$$A \approx 1 \;\;\Rightarrow\;\; C = 2A \approx 2, \quad |\mathcal{A}| = e^{-A} \approx e^{-1} \approx 0.37, \quad w = e^{-C} \approx e^{-2} \approx 0.14.$$

Interpreting $|\mathcal{A}|$ as the surviving interference-term magnitude, the fringe visibility of the branch-interference channel drops to $\sim 37\%$ at threshold and collapses quickly beyond it; the corresponding recognition weight w for the branch that *fails* to produce a durable macroscopic record is ~ 0.14. These numbers inherit order-unity uncertainty from device geometry (mode shape, actual displacement profile, internal smearing), but the calibration pins the *location* of the boundary without free parameters.

Scaling around the point. Keeping geometry fixed, the actionable control variables are:
$$A(\tau, m_{\mathrm{coh}}) \;\propto\; \tau \;\times\; \mathcal{G}(m_{\mathrm{coh}}),$$

with $\mathcal{G}$ a device-specific (but *fixed*) geometric factor encoding the potential sum and mass distribution. Near the calibrated point,
$$A \approx \left(\frac{\tau}{1\ \mathrm{s}}\right)\left(\frac{m_{\mathrm{coh}}}{0.2\ \mathrm{ng}}\right)^{\eta},$$

where η is close to 1 for fixed geometry (effective Φ_{12} linear in the coherently moved mass) and tends toward 3 only in the maximally localized upper bound. The experiment does not need η a priori; it *measures* it by sweeping m_{coh} and τ across the threshold.

Numerical plug (illustrative). At $(m_{\mathrm{coh}}, \tau) = (0.2\ \mathrm{ng}, 1\ \mathrm{s})$:
$$A \approx 1, \quad C \approx 2, \quad |\mathcal{A}| \approx 0.37, \quad w \approx 0.14.$$

Doubling τ at fixed geometry gives $A \approx 2$ and $|\mathcal{A}| \approx e^{-2} \approx 0.14$; halving m_{coh} at fixed τ and geometry typically reduces A roughly by a factor of two (device-dependent).

7.4 Experimental design notes: beating environmental decoherence, choosing observables

Environment. Use cryogenic temperatures (suppress black-body emission/absorption), ultra-high vacuum (suppress gas collisions), vibration isolation, and high-Q suspension to minimize clamping loss and two-level-system damping. The goal is a parameter window where *environmental* decoherence rates are well below the gravitationally induced threshold.

Mode engineering. Shape the mechanical mode so that a well-characterized fraction f of the mass moves coherently. Calibrate f by independent elastic-mode tomography (e.g., ring-down plus laser Doppler vibrometry), then hold f fixed while sweeping τ. This isolates the threshold.

State preparation and readout. Prepare cat states via state-dependent forces or pulsed op-tomechanics. Read out coherence via interference visibility (Ramsey-type sequences in phase space), Wigner-function negativity revival, or parity oscillations under engineered displacements. Benchmark with control sequences that *do not* separate branches (null tests).

Orthogonality control. Increase separation until pointer states are orthogonal in Hilbert space; beyond this point, the residual-action model predicts *distance-insensitive* collapse rates. A plateau in the loss rate vs. separation at fixed m_{coh}, τ is the smoking gun (Sec. 8).

Figure 6 (parameter map). Plot the plane spanned by (m_{coh}, τ) for your device geometry with isocontours $A = 1$ (equivalently $C = 2$). Shade the region $A \geq 1$ as "no coherent superposition" and $A < 1$ as "coherence allowed." Overplot your environmental decoherence contours; target the region where environmental decoherence is subdominant but A crosses unity.

8 How to Falsify the Bridge

A good bridge must be breakable. We highlight three crisp ways data could separate the residual-action model (and thereby the $C = 2A$ mapping) from competitors or from mis-identifications.

8.1 Distance dependence after orthogonality

Prediction. Once branch states are orthogonal, the residual-action model's collapse contribution saturates with separation; it *does not* fall as an inverse distance tail. The recognition mapping shares this saturation, since $C = 2A$ inherits the same Hilbert-space geometry.

Test. At fixed (m_{coh}, τ) above the orthogonality threshold, sweep the branch separation across an order of magnitude. If the observed coherence-loss rate plateaus (distance-independent), the bridge survives; if it decays $\propto 1/d$ (or similar), you have PD-like behavior and the bridge fails. A mechanical interferometer with tunable splitting pulses is sufficient to run this sweep.

8.2 Dispersive noise

Prediction. There is no DP-style dispersive noise kernel in the residual-action model; recognition calculus likewise lacks it. Spectral broadening or excess heating with the characteristic DP kernel would challenge *both*.

Test. Perform noise spectroscopy on the superposed mode while varying m_{coh} and branch separation at fixed τ. Look for frequency-dependent broadening or heating consistent with a gravitational noise kernel. Null detection strengthens the bridge; a positive detection disfavors both frameworks and points to different physics.

8.3 Off-gauge sensitivity

Prediction. The residual is evaluated in the energy gauge (parallel component removed by a time-dependent phase); the bridge $C = 2A$ uses this gauge to match integrands pointwise. If experiments revealed signatures that require a non-energy-gauge residual (e.g., probabilities depending on global phase choices that should be unphysical here), the identification fails.

16

Test. Engineer sequences that differ only by global phase windings of the on-Hamiltonian path (e.g., insert calibrated phase loops with no change to pointer alignment or mass redistribution). Any gauge-dependent change in the collapse boundary would contradict the residual-action stationary-path construction and thus the $C = 2A$ mapping.

In all three tests, the falsifiers are *parameter-free*: they rely only on qualitative behaviors (saturation vs. tail; presence vs. absence of dispersive noise; gauge independence) that do not require tuning. This is the kind of sharpness we want in a bridge paper.

9 How to Falsify the Bridge

9.1 Distance dependence after orthogonality

Prediction (residual-action and recognition). Once the measurement branches are orthogonal in Hilbert space, the contribution to collapse *saturates*; increasing spatial separation no longer reduces (or enhances) the collapse rate to leading order. This follows because the residual measures distance in *Hilbert space*, not in real space; after orthogonality the integrand no longer depends on branch separation to leading order. The recognition mapping inherits the same saturation because $C = 2A$ is enforced along the same local geodesic.

Contrast (PD-style models). In Penrose–Diósi-type frameworks, the effective kernel falls with inverse separation once the packets are orthogonal; a residual *tail* persists that scales like $1/d$ (up to smearing by the packet size).

Interferometer sweep to decide. Hold the coherent mass m_{coh} and the dwell time τ fixed, and sweep the branch separation d across the orthogonality threshold $d_\star$ (verified independently via a visibility or overlap calibration). Measure either the visibility $V(d)$ or an effective collapse rate $\Gamma(d)$ extracted from fringe contrast decay. Fit two nested models on the *post-orthogonality* interval $d \geq d_\star$:

$$\text{(Saturation)} \quad \Gamma(d) = \Gamma_0, \tag{45}$$

$$\text{(Tail)} \quad \Gamma(d) = \Gamma_0 + \frac{K}{d^p}, \quad p > 0. \tag{46}$$

Select using likelihood ratio or information criteria (AIC/BIC); include a nuisance term for residual overlap $\epsilon(d)$ if needed (bounded by the overlap calibration). A statistically significant $K > 0$ (with $p > 0$) falsifies the saturation prediction; a plateau ($K \simeq 0$) falsifies PD-style tails and is consistent with the bridge. Control for environmental decoherence by repeating the sweep with the interferometer *closed* (no branch separation) to isolate non-gravitational contributions. The saturation prediction and its rationale are explicit in the residual-action model's comparison to PD.

9.2 Dispersive noise

Prediction (null). The residual-action model does *not* introduce a DP-style dispersive noise kernel; looking for such noise is not a viable test of the model itself. The recognition calculus likewise lacks any stochastic, dispersive kernel—there is only the deterministic, local action C and the amplitude bridge.

Noise spectroscopy to challenge both. Perform displacement-noise spectroscopy on the superposed mechanical mode while sweeping $(m_{\mathrm{coh}}, d, \tau)$ within the operational window. Let $S_{xx}(\omega)$ be the measured displacement spectral density and $S_{xx}^{\mathrm{env}}(\omega)$ the independently calibrated environmental baseline (thermal, technical, clamping, gas, and readout). Define the residual spectrum

$$\Delta S_{xx}(\omega) \;:=\; S_{xx}(\omega) - S_{xx}^{\mathrm{env}}(\omega). \tag{47}$$

A robust, repeatable, parameterizable *excess* $\Delta S_{xx}(\omega)$ consistent with a dispersive DP kernel falsifies both the residual-action model and the recognition mapping. Conversely, a null bound in the region where collapse is otherwise observed supports the bridge and excludes DP-like alternatives. The residual-action paper explicitly flags that DP-style dispersive tests do not probe this model.

9.3 Off-gauge sensitivity

Prediction (gauge independence). In the energy gauge, the parallel component $\langle \Psi | R | \Psi \rangle$ is removed by a time-dependent phase, leaving $\|R\|$ strictly minimal and the action $S = \int \|R\| \, dt$ invariant under global phase windings. The probabilistic layer depends on A built from $\|R\|$ and the overlap; with $C = 2A$, all observable rates and weights are likewise gauge independent.

Phase-loop experiment. Engineer two sequences with identical pointer-plane geodesics and identical $(m_{\mathrm{coh}}, d, \tau)$, but insert into one sequence a calibrated global phase loop by adding a uniform, time-dependent potential offset $\Delta E(t)$ that only imprints a phase $\varphi(t) = \int \Delta E(t) \, dt$ on the on-Hamiltonian evolution. In the energy gauge this parallel piece is removed, so both sequences must produce identical collapse probabilities and visibilities. A statistically significant difference in thresholds or rates between the sequences would indicate off-gauge sensitivity and thereby invalidate the residual-action construction (and with it the pointwise integrand match used for $C = 2A$).

Summary of falsifiers.

- **Distance sweep:** A post-orthogonality $1/d^p$ tail contradicts saturation and falsifies the bridge; a plateau supports it.

- **Dispersive noise:** Detection of a DP-like noise kernel challenges both models; a null bound in the collapse regime supports them.

- **Gauge test:** Any measurable dependence on global phase windings (off-gauge sensitivity) breaks the residual-action foundation and the $C = 2A$ identification.

Each test is parameter-free at the level that matters: the discriminants (plateau vs. tail, presence vs. absence of dispersive noise, gauge independence vs. dependence) do not require tuning to interpret.

10 Relation to Other Approaches (Concise)

GRW and DP vs. local, parameter-free collapse

Ghirardi–Rimini–Weber (GRW)[1] and Diósi–Penrose (DP)[2] models convert collapse into a dynamical process by adding non-unitary terms that are *nonlocal in Bell's sense* and, in practice, introduce extra

---

[1] G. C. Ghirardi, A. Rimini, T. Weber, Phys. Rev. D 34, 470 (1986).
[2] L. Diósi, Phys. Lett. A 120, 377 (1987); R. Penrose, Gen. Relativ. Gravit. 28, 581 (1996).

knobs (collapse rates, correlation lengths, noise kernels). The residual-action model recapped here is *local and parameter-free*: it enforces a product-subset constraint for matter⊗geometry, minimizes a residual functional in the energy gauge, and obtains selection probabilities from exponential rates without tunable parameters. Two concrete distinctions matter experimentally:

1. **Distance scaling after orthogonality.** DP-like frameworks predict an inverse-distance *tail* for the collapse-driving contribution once branch wave packets are orthogonal; by contrast, the residual-action model predicts *saturation*: after orthogonality, the leading contribution becomes distance-independent because the cost measures separation in Hilbert space rather than in real space. The recognition mapping $C = 2A$ inherits this saturation, so both sides of the bridge agree on the plateau.

2. **Variance vs. sum of potentials.** DP/mean-field approaches weight the *variance* of the mass density (or gravitational self-energy mismatch), schematically

$$E_G^{(\text{var})} \ \propto \ G \int\int \frac{\langle \delta\hat\rho(\mathbf{x})\, \delta\hat\rho(\mathbf{y})\rangle}{|\mathbf{x} - \mathbf{y}|} \, d^3x\, d^3y,$$

   whereas the residual-action model yields a scaling with the *sum of branch potentials* $\Phi_{12} = \Phi_1(\mathbf{x}_1) + \Phi_2(\mathbf{x}_2)$ in the relevant estimate, leading to an accumulated "Penrose-phase" $\Theta_P \sim \tau\, m|\Phi_{12}|$ rather than a variance functional. This difference drives distinct dependencies on geometry, mass distribution, and the onset of collapse.

   Equally important are two null statements. First, the residual-action model predicts *no DP-style dispersive noise kernel*; collapse is not mediated by a stochastic bath. Second, with locality enforced by the product constraint, there are *no free parameters* beyond device geometry and the ordinary Hamiltonian; the phenomenology is fixed once the setup is specified. Both statements are explicit in the formulation and its test discussion. Within the bridge built here, the recognition calculus supplies a unique local cost J and an additive action C that reproduce the very same local, knob-free behavior: short geodesic rotations into pointer-consistent product states, Born weights from exponential measures, saturation beyond orthogonality, and the absence of stochastic dispersive kernels. With $C = 2A$, the two languages make *identical* parameter-free predictions in their shared domain.

On the "superdeterminism" label

The preprint emphasizes a standard logical point: any locally causal account that reproduces quantum correlations must violate *measurement independence* (Bell). The model implements this via an *all-at-once* selection principle over histories and endpoints: realized evolutions are stationary points of the residual action subject to the product-state constraint, and the probabilistic layer is supplied by exponential rate variables. Calling this "superdeterminism" names a *causal structure* (measurement settings correlated with hidden variables), but it is a red herring for the mathematics at hand. The machinery that does the work is variational: *path selection* under locality constraints and *admissibility* of endpoints consistent with pointer states. Translated through the bridge: the recognition calculus encodes the same content without any philosophical overhead. Admissible windows (short, local updates) and a unique convex J generate an additive action C; amplitudes $\mathcal{A} = e^{-C/2}e^{i\phi}$ and weights $w = e^{-C}$ produce Born statistics once alternatives are orthogonal. The "violation of measurement independence" is simply the statement that endpoint admissibility and local geodesics constrain which branches can be realized; there is no need for extra randomness, no hidden tunables, and no nonlocal triggers. In both dialects, the empirically relevant questions reduce

to: (i) do collapse contributions *saturate* after orthogonality, (ii) is there *no* dispersive DP kernel, and (iii) are predictions *gauge independent* in the energy gauge? Those are crisp, parameter-free discriminants—and they are where experiments should look.

11 Field-Theoretic and Covariant Extensions (Sketch)

This section sketches how the residual-action framework and the recognition calculus lift from single-particle, first-quantized language to a generally covariant, field-theoretic setting. On the residual side, we follow the constraint-density formulation with lapse/shift to produce a spacetime-scalar action. On the recognition side, we introduce a local cost *density* on Cauchy slices, yielding a covariant recognition action whose densities match the residual integrand pointwise. The additivity, shortest local rotations, and the identification $C = 2A$ then extend at the level of densities.

11.1 Her sketch: constraint densities H_ν, lapse/shift N^ν, and a spacetime-scalar S

In the canonical, generally covariant formulation, dynamics are encoded by the Hamiltonian (constraint) densities H_ν ($\nu = 0, 1, 2, 3$), contracted with the lapse/shift $N^\nu = (N, N^i)$. Hossenfelder proposes the functional operator

$$\mathcal{L} = i\frac{\delta}{\delta\Sigma} - \hat{H}_\nu N^\nu, \qquad N^\nu = (N, N^i), \tag{48}$$

where Σ is a foliation by Cauchy hypersurfaces, and $\delta/\delta\Sigma$ denotes the functional derivative along the foliation. The residual density is then the action of $\mathcal{L}$ on the state functional, and the covariant generalization of the residual action is a spacetime scalar,

$$S = \int d^4x\,\sqrt{-g}\,\|\mathcal{R}(x)\|, \qquad \mathcal{R}(x) := \left(i\frac{\delta}{\delta\Sigma} - \hat{H}_\nu N^\nu\right)\Psi, \tag{49}$$

evaluated in the energy gauge (parallel piece removed) to minimize $\|\mathcal{R}\|$ locally. This reproduces, in field-theoretic language, the "deviation from Schrödinger evolution" that defines the residual, and it is explicitly constructed to be a scalar under changes of foliation. Writing the foliation explicitly, one can express the scalar as

$$S = \int dt \int_{\Sigma_t} d^3x\, N(t, \mathbf{x})\,\|\mathcal{R}_\perp(t, \mathbf{x})\|, \tag{50}$$

where $\mathcal{R}_\perp$ is the gauge-fixed (energy-gauge) perpendicular component. This is the field-theory analog of the geodesic-length residual that, in the two-branch sector, gave $\|R\| = \dot{\theta}$ and $S = \frac{\pi}{2} - \theta_s$. The same logic underlies the mesoscopic predictions and the comparison to Penrose–Diósi in the covariant setting: the residual measures distance in Hilbert space (now a field configuration space) and saturates once the branches are orthogonal.

11.2 Recognition lift: window calculus on Cauchy slices; density versions of J and C

To lift the recognition calculus, replace the global rate $r(t)$ by a nonnegative, scalar *rate field* $r(x)$ on spacetime or by a rate $r(t, \mathbf{x})$ on each slice Σ_t. The unique local cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ becomes a *density* by multiplication with the invariant measure. We define the *recognition-cost density*

$$\mathfrak{J}(r(x)) := J(r(x)), \qquad J(x) = \frac{1}{2}\left(x + \frac{1}{x}\right) - 1, \tag{51}$$

and the *covariant recognition action*

$$C = \int d^4x \sqrt{-g}\, \mathfrak{J}(r(x)) = \int dt \int_{\Sigma_t} d^3x\, N(t,\mathbf{x})\, \mathfrak{J}(r(t,\mathbf{x})). \tag{52}$$

On each slice, we retain the eight-tick (or more generally, short-window) *admissibility* constraints: the pointer plane is locally fixed, the log-rate is piecewise constant with a small number of jumps, and the boundary values match smoothly into the surrounding evolution. Windows are now compact, causally local regions $W \subset \Sigma_t$ (or spacetime diamonds) on which one computes the local contribution

$$\Delta C(W) = \int_W d^3x\, N\, \mathfrak{J}(r) \quad\text{or}\quad \Delta C(\mathcal{D}) = \int_{\mathcal{D}} d^4x \sqrt{-g}\, \mathfrak{J}(r), \tag{53}$$

for a spacetime diamond $\mathcal{D}$. By strict convexity of J, minimizers are piecewise-constant in $\ln r$ within each window, exactly as in the single-degree case.

The amplitude bridge becomes

$$\mathcal{A} = \exp\!\left[-\tfrac{1}{2}C\right] e^{i\phi}, \qquad w = |\mathcal{A}|^2 = e^{-C}, \tag{54}$$

with ϕ the slice-additive dynamical/kinematic phase. Because C is a scalar, the bridge is covariant; because J is unique and local, C remains additive over disjoint windows on and across slices.

11.3 What carries over: additivity, minimal rotations, and $C = 2A$ at the level of densities

Let $\mathfrak{a}(x)$ denote the *rate-action density* that appears in Hossenfelder's probabilistic layer,

$$\mathfrak{a}(x) := \|\mathcal{R}(x)\| \frac{\sqrt{1 - \mathcal{C}_{\mathrm{ov}}(x)^2}}{\mathcal{C}_{\mathrm{ov}}(x)}, \qquad A = \int d^4x \sqrt{-g}\, \mathfrak{a}(x), \tag{55}$$

where $\mathcal{C}_{\mathrm{ov}}(x)$ is the local overlap between the realized path and the on-Hamiltonian path in the energy gauge. This is the field-theoretic analog of the two-branch integrand $\tan(\theta - \theta_s)$ that integrated to $A = -\ln\sin\theta_s$ in the minimal rotation, and it is the density that controls the exponential rates $r = e^{-2A}$.

Density-level kernel match. On each admissible window (slice or spacetime diamond), choose the rate field $r(x)$ such that

$$\mathfrak{J}(r(x)) = 2\,\mathfrak{a}(x) \quad\Longleftrightarrow\quad J(r(x)) = 2\|\mathcal{R}(x)\| \frac{\sqrt{1 - \mathcal{C}_{\mathrm{ov}}(x)^2}}{\mathcal{C}_{\mathrm{ov}}(x)}. \tag{56}$$

This is the covariant analog of the pointwise, two-branch kernel equality $J(r)\, dt = 2\, dA$. Integrating gives

$$C = \int d^4x \sqrt{-g}\, \mathfrak{J}(r(x)) = 2\int d^4x \sqrt{-g}\, \mathfrak{a}(x) = 2A, \tag{57}$$

and therefore $w = e^{-C} = e^{-2A}$ and $|\mathcal{A}| = e^{-A}$ in the covariant setting. Because both A and C are additive over disjoint windows and across slices, multi-region evolutions factorize exactly as in the non-field-theoretic case.

Minimal local rotations. In field space, the two-branch geodesic generalizes to a local rotation in the pointer two-plane $\{1(x), 2(x)\}$ at each spacetime point (or along each worldline segment of the detector degrees). The energy gauge again removes the parallel residual locally and minimizes $\|\mathcal{R}\|$. The corresponding recognition window uses a piecewise-constant $\ln r(x)$ to reproduce the same local geodesic rotation. Thus the geometric heart of the construction—shortest admissible rotations into pointer-consistent product states—survives the covariant lift unchanged.

Probability layer and Born weights. With $C = 2A$ established density-wise, the exponential rates and recognition weights coincide,

$$r = e^{-2A} = e^{-C} = w, \tag{58}$$

and Born's rule follows from the amplitude bridge once alternatives are orthogonal on each slice (or globally, after integration). The distance-independence after orthogonality and the absence of dispersive DP-style kernels carry over directly, since both statements are properties of the residual density and the local additivity, not of low-dimensional kinematics.

12   Practical Guide: Using the Mapping

This section condenses the bridge into a short recipe for analysis and experiment design. It is written to be executable on a single page: compute the residual data once, translate it to recognition weights by $C = 2A$, and read off probabilities and thresholds with no tunable parameters.

12.1   Recipe: end-to-end in a few steps

Inputs. A concrete setup (interferometer, weak measurement, or detector readout) with (i) a pointer basis $\{\Psi_I\}$ for the detector end-states, (ii) an on-Hamiltonian reference path $\Psi'(t)$, and (iii) the realized (off-Hamiltonian, product-constrained) path $\Psi(t)$. Work in the energy gauge so that the parallel residual vanishes.

1. **Fix the pointer basis.** Choose pointer-consistent end-states $\{\Psi_I\}$ (near-product states of matter⊗geometry) for the detector. In practice these are the detector eigenstates that amplify to macroscopic scale.

2. **Compute the residual geometry.** Form the residual in the energy gauge

$$R(t) := \big(i\partial_t - \hat{H}\big)\Psi(t), \quad \langle\Psi|R|\Psi\rangle = 0,$$

   and the instantaneous overlap to the on-Hamiltonian path

$$C_{\mathrm{ov}}(t) := \big|\langle\Psi'(t) \mid \Psi(t)\rangle\big|.$$

   For a two-branch local rotation, $\|R\| = \dot{\theta}$ and $S = \int \|R\|dt = \frac{\pi}{2} - \theta_s$.

3. **Compute the rate action(s) A_I.** For each candidate end-state I, evaluate

$$A_I = \int_{t_s}^{t_e} \|R(t)\| \, \frac{\sqrt{1 - C_{\mathrm{ov},I}(t)^2}}{C_{\mathrm{ov},I}(t)} \, dt,$$

   which reduces to $A = -\ln\sin\theta_s$ for the two-branch geodesic analyzed in the paper.

4. **Map to recognition data by $C = 2A$.** Set

$$C_I := 2A_I, \qquad w_I := e^{-C_I}, \qquad \mathcal{A}_I := e^{-C_I/2} e^{i\phi_I}.$$

The bridge ensures $w_I = e^{-2A_I}$ and $|\mathcal{A}_I|^2 = w_I$.

5. **Read off probabilities.** For orthogonal alternatives (pointer condition),

$$P_I = \frac{w_I}{\sum_J w_J} = \frac{e^{-C_I}}{\sum_J e^{-C_J}} = \frac{e^{-2A_I}}{\sum_J e^{-2A_J}} = |\alpha_I|^2,$$

matching the rate model's Born weights and the amplitude-bridge normalization.

6. **Check the threshold.** A macroscopic record (collapse) occurs when the minimal local update cost reaches order unity:

$$A \sim 1 \quad \Longleftrightarrow \quad C/2 \sim 1,$$

i.e. when the Penrose-phase style bound $\tau\, m |\Phi_{12}| \gtrsim 1$ is met (residual side) or, equivalently, when $|\mathcal{A}| = e^{-A}$ becomes small enough to quench interference (recognition side).

Numerical note. For arbitrary drives, discretize time: evaluate $\|R_k\|$ and $C_{\mathrm{ov},k}$ on a mesh, sum $A_I \approx \sum_k \|R_k\| \sqrt{1 - C_{\mathrm{ov},k}^2}/C_{\mathrm{ov},k}\, \Delta t$, then set $C_I = 2A_I$. In a two-branch window, you can instead parameterize by the Bloch angle θ and integrate the geodesic kernel directly.

12.2 Sanity checks

Two-path interferometer. At the first beam splitter, compute the two-branch geodesic: $\|R\| = \dot\theta$, $S = \frac{\pi}{2} - \theta_s$, $A = -\ln \sin \theta_s$; then $C = 2A$, $w = e^{-C}$, $P = \frac{e^{-C_1}}{e^{-C_1} + e^{-C_2}} = \frac{|\alpha_1|^2}{|\alpha_1|^2 + |\alpha_2|^2}$. This reproduces the paper's derivation.

Weak measurement threshold. For a detector that only sometimes amplifies, evaluate A over the interaction window; records arise near $A \sim 1$. With the bridge, the same boundary is $C/2 \sim 1$, and $|\mathcal{A}| = e^{-A}$ sets fringe visibility.

Sequential updates. For cascaded local interactions (e.g., two weak probes then a projective readout),

$$A_{\mathrm{tot}} = \sum_k A_k, \quad C_{\mathrm{tot}} = \sum_k C_k = 2 \sum_k A_k, \quad \mathcal{A}_{\mathrm{tot}} = \prod_k e^{-C_k/2} e^{i\phi_k},$$

so weights and probabilities factorize exactly.

13 Discussion

Why parameter-free matters. A measurement narrative that introduces new knobs (rates, correlation lengths, bath temperatures) can always be tuned to survive the next experiment. The residual-action model commits to gravity alone: the collapse cost is the time–integrated residual in the energy gauge, yielding thresholds at $\tau\, m |\Phi_{12}| \gtrsim 1$ with no free parameters beyond device geometry. The recognition calculus likewise fixes a unique cost J and an additive action C. The bridge $C = 2A$ preserves this discipline: once the setup is specified, all numbers follow.

Why locality matters more than labels. Bell's theorem forces any locally causal model that reproduces quantum statistics to violate measurement independence. The preprint embraces this by using a stationary, "all-at-once" selection over paths and endpoints; the realized evolution is a local geodesic into pointer states with probabilities from exponential rates. The mathematics that carries the weight is not the label but the variational structure: minimize the residual action in the energy gauge and count stationary endpoints with an exponential measure. The recognition calculus is the same story in different dialect: local admissible windows, unique convex J, additive C, and an amplitude bridge.

Two dialects, one idea. "Product subset" (matter and geometry as one state) forbids independent dressing phases on separate branches; "admissibility windows" algorithmically enforce the same restriction by allowing only short, source-side updates into whole-particle pointer states. Both produce the shortest local rotations and the same Born weights; both predict saturation of collapse contributions after orthogonality and the absence of DP-style dispersive noise. With $C = 2A$, the agreement is quantitative: rates e^{-2A} equal recognition weights e^{-C}, amplitudes have modulus e^{-A}, and the weak-measurement threshold is $A \sim 1 \Leftrightarrow C/2 \sim 1$.

What remains to probe. Three crisp discriminants decide the matter: (i) a post-orthogonality *plateau* (not an inverse-distance tail) in collapse rates versus branch separation, (ii) the *absence* of DP-like dispersive noise in spectral data, and (iii) *gauge independence* under global phase loops. All three tests are parameter-free and already compatible with state-of-the-art mesoscopic platforms.

Parity orientation and cross-program coherence. The eight-tick global orientation chosen at recognition onset (R0) fixes a CP-odd sign simultaneously across multiple sectors in the Recognition Science program. Within the baryogenesis framework, this same orientation determines the sign of the primordial TB/EB parity-odd CMB spectra via the pseudoscalar coupling $\chi R\tilde{R}$. The present measurement bridge inherits the same orientation choice: the sign of amplitudes and the direction of local updates are tied to the eight-tick cycle. This one-bit link—shared across very-early-time CP violation and local measurement dynamics—provides a cross-epoch falsifier: if future CMB data establish a definite TB/EB sign at the predicted amplitude and it contradicts the baryon asymmetry sign, the unified orientation claim fails.

Appendix A: Re-derivation of $\|R\| = \dot{\theta}$ and $S = \pi/2 - \theta_s$ in the energy gauge (full algebra)

We work in a two-branch subspace spanned by orthonormal pointer states $\{1, 2\}$, each of which satisfies the standard Schrödinger equation with the laboratory Hamiltonian $\hat{H}$; equivalently, $i\,\partial_t j = \hat{H}j$ for $j = 1, 2$. Consider the time-dependent rotation

$$\Psi(t) = \cos\theta(t)\,1 + e^{i\phi(t)}\sin\theta(t)\,2, \qquad \theta(t) \in [\theta_s, \pi/2], \tag{59}$$

with boundary data $\cos\theta(t_s) = \alpha_1 \geq 0$, $e^{i\phi(t_s)}\sin\theta(t_s) = \alpha_2$, and $\cos\theta(t_e) = 0$. Define the residual with respect to $\hat{H}$,

$$R := \left(i\partial_t - \hat{H}\right)\Psi. \tag{60}$$

Because 1 and 2 separately obey $i\partial_t j = \hat{H}j$, all Hamiltonian terms cancel between $i\partial_t$ and $\hat{H}$ on each branch. A direct derivative gives

$$i\partial_t\Psi = i\big(-\dot{\theta}\sin\theta\big)1 + i\,\cos\theta\,\partial_t1 + i\big(i\dot{\phi}\,e^{i\phi}\sin\theta + e^{i\phi}\dot{\theta}\cos\theta\big)2 + i\,e^{i\phi}\sin\theta\,\partial_t2, \tag{61}$$

$$\hat{H}\Psi = \cos\theta\,\hat{H}1 + e^{i\phi}\sin\theta\,\hat{H}2 = i\,\cos\theta\,\partial_t1 + i\,e^{i\phi}\sin\theta\,\partial_t2. \tag{62}$$

Hence

$$R = -\,i\,\dot{\theta}\sin\theta\,1 + i\,e^{i\phi}\dot{\theta}\cos\theta\,2 - \dot{\phi}\,e^{i\phi}\sin\theta\,2. \tag{63}$$

Decompose R into its parallel and perpendicular parts with respect to Ψ,

$$R_\parallel = \Psi\,\langle\Psi|R\rangle, \qquad R_\perp = R - R_\parallel. \tag{64}$$

Using $\Psi = \cos\theta\,1 + e^{-i\phi}\sin\theta\,2$ and orthonormality, we obtain

$$\langle\Psi|R\rangle = -i\,\dot{\theta}\sin\theta\cos\theta + i\,\dot{\theta}\sin\theta\cos\theta - \dot{\phi}\sin^2\theta = -\,\dot{\phi}\sin^2\theta. \tag{65}$$

Therefore,

$$\begin{aligned}
R_\perp &= R - \Psi\left(-\dot{\phi}\sin^2\theta\right) \\
&= \left(-i\,\dot{\theta}\sin\theta + \dot{\phi}\sin^2\theta\cos\theta\right)1 + e^{i\phi}\left(i\,\dot{\theta}\cos\theta - \dot{\phi}\sin\theta\cos^2\theta\right)2.
\end{aligned} \tag{66}$$

The squared norm is

$$\begin{aligned}
\|R_\perp\|^2 &= \left(\dot{\theta}^2\sin^2\theta + \dot{\phi}^2\sin^4\theta\cos^2\theta\right) + \left(\dot{\theta}^2\cos^2\theta + \dot{\phi}^2\sin^2\theta\cos^4\theta\right) \\
&= \dot{\theta}^2 + \dot{\phi}^2\sin^2\theta\cos^2\theta.
\end{aligned} \tag{67}$$

Energy gauge and geodesic. The *energy gauge* sets the parallel component to zero by a time-dependent *global* phase $\Psi \mapsto e^{-if(t)}\Psi$, under which $\langle\Psi|R\rangle \mapsto \langle\Psi|R\rangle - \dot{f}$. Choosing $\dot{f} = \langle\Psi|R\rangle$ forces $\langle\Psi|R\rangle = 0$ *without* constraining the relative phase $\phi(t)$. In this gauge, the action is

$$S = \int \|R\|\,dt = \int \|R_\perp\|\,dt = \int \sqrt{\dot{\theta}^2 + \dot{\phi}^2\sin^2\theta\cos^2\theta}\,dt. \tag{68}$$

For fixed endpoints $(\theta_s, \phi_s) \to (\pi/2, \phi_e)$ the integrand is a Finsler length on the Bloch sphere with metric element $ds^2 = d\theta^2 + \sin^2\theta\cos^2\theta\,d\phi^2$. The shortest path (geodesic) keeps $\phi(t) = \text{const}$, thus $\dot{\phi} = 0$, and moves only in θ. Hence

$$\boxed{\|R\| = \dot{\theta}}\,, \qquad \boxed{S = \int_{\theta_s}^{\pi/2} d\theta = \frac{\pi}{2} - \theta_s}\,. \tag{69}$$

This is precisely the "shortest local rotation" result.

Remarks. (i) The result depends only on the endpoints; the schedule $\theta(t)$ drops out. (ii) The derivation uses only orthonormality, the fact that each branch obeys the Schrödinger equation, and the energy gauge. (iii) The geodesic is unique up to reparameterization of t.

Appendix C: Recognition calculus — uniqueness of J, bridge construction, and orthogonality bookkeeping

C.1 Uniqueness of the local cost under symmetry axioms

Let $x > 0$ be a dimensionless local recognition rate and $J : \mathbb{R}_{>0} \to [0, \infty)$ a local cost density. Set $t = \ln x$ and $F(t) := J(e^t)$. Impose:

1. **Identity and reciprocity:** $J(1) = 0$ and $J(x) = J(x^{-1}) \Rightarrow F(0) = 0$, F even.

2. **Positivity/convexity:** $J(x) > 0$ for $x \neq 1$; F strictly convex, real-analytic, $F'(0) = 0$.

3. **Local additivity:** For constant log-rate t over duration Δ, cost is $F(t)\,\Delta$; partitioning in time leaves cost unchanged.

4. **Normalization:** $F''(0) = 1$ (units).

5. **No coarse-graining penalty (equality on constant segments):** Serial composition with constant t over durations Δ_1, Δ_2 gives $F(t)(\Delta_1 + \Delta_2)$, i.e. the equality case of subadditivity holds for piecewise-constant t.

Functional identity. Consider two equal sub-intervals with log-rates $t \pm u$. Let t_{eff} denote the (unique by convexity) constant-rate representative with the same cost over the doubled duration. Then

$$2F(t_{\text{eff}}) = F(t + u) + F(t - u). \tag{70}$$

Consistency under concatenation and symmetry implies the quadratic form

$$F(t + u) + F(t - u) = 2F(t)F(u) + 2\big(F(t) + F(u)\big). \tag{71}$$

(Construct this by comparing two ways of coarse-graining four equal sub-intervals with rates $t \pm u$ and using the equality cases of subadditivity on constant pieces.)

Solution and uniqueness. Equation (71) together with $F(0) = 0$, $F'(0) = 0$, $F''(0) = 1$ has the unique real-analytic even solution

$$F(t) = \cosh t - 1 \quad \Longleftrightarrow \quad \boxed{J(x) = \tfrac{1}{2}\left(x + \frac{1}{x}\right) - 1}. \tag{72}$$

Indeed,

$$(\cosh(t + u) - 1) + (\cosh(t - u) - 1) = 2\cosh t \cosh u - 2,$$

and the right-hand side of (71) evaluates to $2(\cosh t - 1)(\cosh u - 1) + 2\big((\cosh t - 1) + (\cosh u - 1)\big) = 2\cosh t \cosh u - 2$, so the identity holds. Strict convexity and the normalization fix the overall scale. $\square$

C.2 Bridge construction

Given a path γ with rate $r(t) > 0$, define the recognition action, weight, and amplitude

$$C[\gamma] = \int J(r(t))\, dt, \qquad w[\gamma] = e^{-C[\gamma]}, \qquad \mathcal{A}[\gamma] = e^{-C[\gamma]/2}\, e^{i\phi[\gamma]}. \tag{73}$$

For disjoint concatenated segments $\{\gamma_k\}$ one has additivity and multiplicativity,

$$C[\circ_k \gamma_k] = \sum_k C[\gamma_k], \quad w[\circ_k \gamma_k] = \prod_k w[\gamma_k], \quad \mathcal{A}[\circ_k \gamma_k] = \prod_k \mathcal{A}[\gamma_k]. \tag{74}$$

When alternatives are orthogonal (pointer condition), the observed frequency of outcome I with alternative amplitude $\mathcal{A}_I = \sum_{\gamma \in \Gamma_I} \mathcal{A}[\gamma]$ is

$$P(I) = \frac{|\mathcal{A}_I|^2}{\sum_J |\mathcal{A}_J|^2}. \tag{75}$$

In the saddle-point approximation, each Γ_I is dominated by the minimal-cost path $\gamma_I^\star$, giving $|\mathcal{A}_I|^2 \simeq e^{-C_I}$ with $C_I = C[\gamma_I^\star]$. Thus

$$P(I) \simeq \frac{e^{-C_I}}{\sum_J e^{-C_J}}. \tag{76}$$

C.3 Orthogonality bookkeeping and the $C = 2A$ bridge

For the two-branch local rotation, the residual-action analysis gives $A = -\ln \sin \theta_s$ and hence $e^{-2A} = |\alpha_2|^2$; for D outcomes, pairwise ratios fix $P(I) = |\alpha_I|^2$. When the constructive kernel match $J(r)\, dt = 2\, dA$ is imposed along the geodesic (and slice-by-slice in the field-theoretic lift), one has $C = 2A$ *exactly* on minimizing profiles, so that

$$w = e^{-C} = e^{-2A}, \qquad |\mathcal{A}| = e^{-A}, \qquad P(I) = \frac{e^{-C_I}}{\sum_J e^{-C_J}} = \frac{e^{-2A_I}}{\sum_J e^{-2A_J}} = |\alpha_I|^2. \tag{77}$$

Orthogonality ensures cross terms vanish in $|\sum_{\gamma \in \Gamma_I} \mathcal{A}[\gamma]|^2$, and additivity guarantees factorization for sequential local updates. These statements complete the bookkeeping that connects the recognition calculus to the residual-action probabilities without introducing tunable parameters.

Appendix D: Proof of $C = 2A$ for geodesic rotations with explicit $r(t)$

D.1 Setup and goal

Work in the two-branch pointer plane with orthonormal $\{1, 2\}$. The realized (off-Hamiltonian) path rotates

$$\Psi(t) = \cos \theta(t) 1 + e^{i\phi} \sin \theta(t) 2, \quad \theta(t) \in [\theta_s, \tfrac{\pi}{2}],$$

and we use the energy gauge so that $\langle \Psi | R | \Psi \rangle = 0$. Appendix A showed $\|R\| = \dot{\theta}$ and

$$S = \int \|R\|\, dt = \int_{\theta_s}^{\pi/2} d\theta = \frac{\pi}{2} - \theta_s,$$

and Appendix B showed that the *rate action* is

$$A = \int_{t_s}^{t_e} \|R\| \frac{\sqrt{1 - C(t)^2}}{C(t)} \, dt = \int_{\theta_s}^{\pi/2} \tan(\theta - \theta_s) \, d\theta = -\ln \sin \theta_s.$$

The recognition action is $C = \int J(r(t)) \, dt$ with the unique local cost $J(x) = \frac{1}{2}(x + \frac{1}{x}) - 1 = \cosh(\ln x) - 1$. Our goal is to construct $r(t)$ *explicitly* so that $C = 2A$ *pointwise* along the geodesic, i.e.

$$J(r(t)) \, dt = 2 \, dA.$$

All ingredients and the geodesic characterization come directly from the residual-action framework.

D.2 Pointwise kernel match (existence and construction)

Define the Bloch-angle offset $\vartheta(t) := \theta(t) - \theta_s \in [0, \pi/2 - \theta_s]$. Along the geodesic, $\phi = \text{const}$, hence

$$dA = \tan \vartheta \, d\theta = \tan \vartheta \, d\vartheta \quad \Rightarrow \quad 2 \, dA = 2 \tan \vartheta \, d\vartheta.$$

Choose the update parameter so that $t \equiv \vartheta$ (we can always reparameterize a monotone geodesic), hence $dt = d\vartheta$. We need $r(\vartheta) > 0$ such that

$$J(r(\vartheta)) = 2 \tan \vartheta \qquad (\vartheta \in [0, \pi/2 - \theta_s]).$$

Since J is continuous, strictly increasing on $[1, \infty)$ in the variable x and even in $\ln x$, every nonnegative value y is attained by J uniquely at $x = e^u$ with $u = \text{arcosh}(1 + y) \geq 0$. Set

$$u(\vartheta) := \text{arcosh}(1 + 2 \tan \vartheta), \qquad r(\vartheta) := \exp u(\vartheta). \tag{78}$$

Then

$$J(r(\vartheta)) = \cosh u(\vartheta) - 1 = 2 \tan \vartheta,$$

and the pointwise equality follows:

$$J(r(\vartheta)) \, dt = J(r(\vartheta)) \, d\vartheta = 2 \tan \vartheta \, d\vartheta = 2 \, dA. \tag{79}$$

Integrating (79) over $\vartheta \in [0, \pi/2 - \theta_s]$ yields $C = 2A$.

General time parameterization. If one keeps a generic monotone schedule $\theta(t)$ with $\dot{\theta} > 0$, set

$$y(t) := 2 \tan(\theta(t) - \theta_s) \, \dot{\theta}(t) \geq 0, \qquad u(t) := \text{arcosh}(1 + y(t)), \qquad r(t) := e^{u(t)}. \tag{80}$$

Then $J(r(t)) = y(t)$ and $J(r(t)) \, dt = 2 \tan(\theta - \theta_s) \, d\theta = 2 \, dA$ pointwise. Thus the geodesic's *shape* fixes the unique minimal C and we recover $C = 2A$ irrespective of the time reparameterization.

D.3 Admissibility and uniqueness on the geodesic

By construction $u \geq 0$ and $r = e^u \geq 1$ for $\vartheta \in [0, \pi/2)$. Since $J(x)$ is strictly convex in $\ln x$, the minimal C for given boundary data on a *short* local rotation is attained by a monotone u, which (80) provides. Any other admissible profile with the same endpoints and the same geodesic shape must obey Jensen and subadditivity bounds and yields $C > 2A$ unless it coincides with (80) almost everywhere.

D.4 Discrete eight-tick windows converge to the continuous profile

Let $W = [t_0, t_0 + T]$ be an admissible window partitioned into $N = 8$ equal ticks (or any N): $t_k = t_0 + k\Delta t$, $\Delta t = T/N$. Choose the piecewise-constant approximation $r_k := r(t_k)$ with $r(t)$ as in (80). Then

$$C_N := \sum_{k=0}^{N-1} J(r_k)\,\Delta t \xrightarrow[N\to\infty]{} \int_{t_0}^{t_0+T} J\big(r(t)\big)\,dt = 2A_W,$$

where A_W is the rate action accumulated on W. Because J is smooth on any compact range and $r(t)$ is continuous on W, the Riemann-sum error is $O(\Delta t)$. In practice $N = 8$ already gives a controlled approximation for short windows; strict convexity guarantees that any other piecewise-constant admissible profile with the same boundary values has cost $\geq C_N$.

D.5 Conclusion

For geodesic two-branch rotations in the energy gauge, the explicit rate profiles (78) or (80) enforce the pointwise kernel identity $J(r)\,dt = 2\,dA$, hence

$$\boxed{C = 2A, \qquad w = e^{-C} = e^{-2A}, \qquad |\mathcal{A}| = e^{-A}}.$$

This saturates the recognition-versus-residual bridge and is the unique minimal-cost realization subject to local admissibility.

Appendix E: Worked numeric example — ng-scale oscillator (numbers, errors, device parameters)

We present a concrete parameter point for a membrane/oscillator superposition and compute the dimensionless actions A and $C = 2A$ using the calibration implied in the residual-action paper's estimates. That paper states that superpositions which displace a *coherent* mass of about 0.2 ng for ~ 1 s lie at the collapse threshold; heavier coherent mass or longer dwell crosses it. We use this as the parameter-free anchor $A \simeq 1$ (hence $C \simeq 2$) at the reference point.

E.1 Device and operating point

- **Total device mass:** $m_{\text{tot}} = 1.0$ ng $= 10^{-12}$ kg.

- **Coherent fraction:** $f = 0.20$ (from mode-shape tomography); coherent mass $m_{\text{coh}} = f\,m_{\text{tot}} = 0.20$ ng $= 2 \times 10^{-13}$ kg.

- **Branch displacement:** $d \sim 5$ fm (sufficient for nuclear non-overlap).

- **Coherence (dwell) time:** $\tau = 1.0$ s (ringdown-limited).

These values place the superposition exactly at the published threshold scale. Environmental decoherence must be engineered below this level for the gravitational effect to be visible.

E.2 Calibrated action and derived quantities

Define an (experiment-specific) geometric factor $\mathcal{G}$ that encodes the effective potential sum $|\Phi_{12}|$ and spatial mode shape, such that

$$A(\tau, m_{\mathrm{coh}}) \;=\; \tau\,\mathcal{G}\,m_{\mathrm{coh}},$$

with $\mathcal{G}$ *fixed* by the device geometry. Calibrate $\mathcal{G}$ at the reference point by imposing $A = 1$ at $(\tau, m_{\mathrm{coh}}) = (1\text{ s}, 0.2\text{ ng})$, i.e.

$$\mathcal{G}_\star = \frac{1}{\tau\,m_{\mathrm{coh}}}\Big|_\star = \frac{1}{(1\text{ s})\,(0.2\text{ ng})}.$$

Then for any nearby settings (same geometry),

$$A \simeq \left(\frac{\tau}{1\text{ s}}\right)\left(\frac{m_{\mathrm{coh}}}{0.2\text{ ng}}\right), \qquad C = 2A, \qquad |\mathcal{A}| = e^{-A}, \qquad w = e^{-C}.$$

Reference numbers: at the calibration point

$$A = 1.00, \quad C = 2.00, \quad |\mathcal{A}| = e^{-1} \approx 0.368, \quad w = e^{-2} \approx 0.135.$$

E.3 Scaling variants around the point

- **Longer dwell (same mass):** $\tau = 2$ s $\Rightarrow A = 2$, $C = 4$, $|\mathcal{A}| \approx 0.135$, $w \approx 0.018$.

- **Heavier coherent mass (same dwell):** $m_{\mathrm{coh}} = 0.4$ ng $\Rightarrow A = 2$, $C = 4$ (same numbers as above).

- **Lighter coherent mass (same dwell):** $m_{\mathrm{coh}} = 0.1$ ng $\Rightarrow A = 0.5$, $C = 1$, $|\mathcal{A}| \approx 0.607$, $w \approx 0.368$.

These relations hold provided the branch displacement keeps the packets orthogonal, so that the distance-insensitive (post-orthogonality) regime applies.

E.4 Error budget (dominant contributions)

Let fractional uncertainties be δf (coherent fraction), δm (total mass), $\delta\tau$ (dwell time), and $\delta\mathcal{G}$ (geometry). Since $A = \tau\,\mathcal{G}\,f\,m_{\mathrm{tot}}$,

$$\frac{\delta A}{A} \;\approx\; \delta\tau \oplus \delta\mathcal{G} \oplus \delta f \oplus \delta m,$$

where $\oplus$ denotes root-sum-square if errors are independent. Typical conservative values near current devices:

$$\delta\tau \sim 20\%, \quad \delta f \sim 30\%, \quad \delta m \sim 10\%, \quad \delta\mathcal{G} \sim 30\%,$$

give $\delta A/A \sim 0.5$. Then $A = 1.0 \pm 0.5$, $C = 2.0 \pm 1.0$, $|\mathcal{A}| = e^{-A} = 0.37^{+0.23}_{-0.16}$, $w = e^{-C} = 0.135^{+0.139}_{-0.071}$.

E.5 Device parameters to report (for reproducibility)

Report: (i) m_{tot} from metrology; (ii) f from mode-shape tomography (e.g., laser Doppler vibrometry plus FEM); (iii) τ from ringdown; (iv) preparation fidelity and verified branch orthogonality (overlap $< 10^{-2}$); (v) environmental decoherence rates (gas, black-body, clamping, readout); (vi) independent calibration of $\mathcal{G}$ by quasi-static force/displacement mapping of the mode. The residual-action paper's estimates place this class of devices as the most promising near-term test.

E.6 Minimal reproducibility recipe

To compute C and A for an arbitrary two-branch measurement:

1. **Fix pointer states:** Choose orthonormal end-states $\{1, 2\}$; compute initial amplitudes α_1, α_2.

2. **Compute residual:** In energy gauge, evaluate $\|R\| = \dot{\theta}$ and $C_{\text{ov}}(t) = \cos(\theta - \theta_s)$ for the geodesic.

3. **Integrate rate action:** $A = \int_{\theta_s}^{\pi/2} \tan(\theta - \theta_s)\, d\theta = -\ln\sin\theta_s$.

4. **Map to recognition:** Set $C = 2A$, $w = e^{-C}$, $\mathcal{A} = e^{-C/2}e^{i\phi}$.

5. **Check:** Verify $e^{-2A} = |\alpha_2|^2$ and $P_1 : P_2 = |\alpha_1|^2 : |\alpha_2|^2$.

For the worked mesoscopic example: at $(m_{\text{coh}}, \tau) = (0.2 \text{ ng}, 1 \text{ s})$, calibrate $\mathcal{G}$ so that $A = 1$ (threshold); then $C = 2$, $|\mathcal{A}| \approx 0.37$, $w \approx 0.14$. All numbers follow from the geometry factor $\mathcal{G}$ with no free knobs.