

Stochastic Rupture as an Information-Bounded Mechanism for Objective Wave-Function Collapse

Consolidated Edition: Trigger Regime, Cosmological Suppression,
and Geometric Foundations

Guilherme Zambuzi
Independent Researcher
gzambuzi@gmail.com

April 28, 2026

Abstract

We present a consolidated formulation of **Stochastic Rupture** (SR), an objective wave-function collapse framework in which branch selection is triggered by the local saturation of an informational bound on covariant causal-diamond surfaces. The framework is built upon a Diósi–Penrose-extended decoherence rate $\Gamma_{\text{SR}} \sim Gm^2d^2/(\eta\hbar\sigma_x^3)$, modulated by a single phenomenological parameter $\eta \sim 0.1$ encoding the saturation fraction at which pruning fires. The framework converges with standard Diósi–Penrose in the regime $d \sim \sigma_x$ but diverges quadratically in the extended-superposition regime $d \gg \sigma_x$, providing a clean experimental discriminator. We make three principal contributions in this consolidated edition: (i) we explicitly articulate the *trigger regime hierarchy*, establishing that isolated microscopic systems do not generate definite events in deep vacuum and that ontological emergence of facts requires saturation in a macroscopic register; (ii) we compute the cosmological heating rate predicted by SR and show it is 15–27 orders of magnitude below CSL in diffuse environments, automatically satisfying Lyman- α and CMB spectral-distortion constraints that pressure competing collapse models; (iii) we develop the geometric foundations of the *4D Rose* representation, showing that non-spherical, chiral, branching geometry is structurally required by parity-violating, irreversibly-growing wavefunctions. The framework is shown to be observationally indistinguishable from standard quantum mechanics in all regimes accessible to current experiments, while predicting detectable signatures in the mesoscopic-superposition regime ($m \sim 10^{10}$ amu, $d \sim 100$ nm, $\sigma_x \sim 10$ nm), corresponding directly to the parameter space of the MAQRO experimental program.

Contents

1	Introduction	2
2	The Pruning Mechanism	3
2.1	Decoherence rate from gravitational self-energy	3
2.2	Informational modulation: the saturation factor	4
2.3	Pruning as ontological selection	4
2.4	Comparison with Diósi–Penrose	4
3	The Trigger Regime Hierarchy	4
3.1	The isolated-microsystem problem	5
3.2	Hierarchy of regimes	5
3.3	Ontological position on factuality	5
3.4	Empirical equivalence in the decay sector	6

4	Cosmological Heating: Structural Suppression	6
4.1	Heating rate per particle	6
4.2	Comparison across cosmological environments	6
4.3	Structural reason	7
5	Geometric Foundations: The 4D Rose	7
5.1	Properties required of the geometry	7
5.2	Why spherical alternatives fail	8
5.3	Connection to parity violation	8
5.4	Pruning as petal selection	8
6	Experimental Window: MAQRO and the d^2/σ_x^3 Signature	8
6.1	The d^2/σ_x^3 scaling	8
6.2	Consistency with current levitated optomechanics	9
6.3	The MAQRO regime	9
6.4	Empirical constraints on η	9
7	Framework Scope and Companion Cosmological Channel	10
7.1	Domains the framework does not address	10
8	Discussion	10
8.1	Summary of core claims	10
8.2	Energy conservation	11
8.3	Open problems and future work	11
A	On the Possibility of a Deterministic Substrate	12

1 Introduction

The measurement problem in quantum mechanics — the question of how and why definite outcomes emerge from superposed states — remains one of the foundational questions of physics. Among objective-collapse approaches, the GRW model [1] and its continuous extension CSL [2] postulate stochastic localization at a fundamental rate, while gravitational models in the Diósi–Penrose tradition [3, 4] tie collapse to gravitational self-energy. Both classes face increasingly tight experimental constraints from cosmological heating bounds and matter-wave interferometry [15], and both treat collapse as a fundamental dynamical process distinct from the informational structure of spacetime.

Stochastic Rupture (SR) takes a different stance. Rather than postulating a new fundamental dynamics, SR identifies collapse as the irreversible *pruning* of branches when the local von Neumann entropy of a superposition saturates a fraction η of the Bekenstein–Bousso bound on the relevant causal-diamond surface. The rate formula adopted is the Diósi–Penrose extension to spatially-separated Gaussian wavepackets, dimensionally consistent and recovering standard DP in the appropriate limit:

$$\Gamma_{\text{SR}} \sim \frac{Gm^2d^2}{\eta\hbar\sigma_x^3}. \quad (1)$$

The framework is informational at its conceptual core: pruning is not caused by interaction or measurement, but by the geometric fact that local information density cannot exceed the holographic capacity of the bounding surface. The DP-extended dynamical rate is interpreted as the phenomenological consequence of this saturation principle.

Three structural claims distinguish SR from continuous and gravitational collapse models:

1. **Trigger by saturation, not by tag.** SR does not assign a fundamental rate to individual particles (as in CSL) or a deterministic timescale to each mass (as in DP without modulation). Pruning rate scales with the local saturation ratio $\chi = S_{\text{vN}}/I_{\text{Bek}}$, which is structurally suppressed in coherent and dilute regimes.
2. **Selection over branches, not modification of dynamics.** SR does not modify the unitary evolution of amplitudes. It prescribes which branch becomes a fact upon saturation, with Born statistics recovered from amplitude weights. Microscopic systems isolated from macroscopic registers evolve unitarily without producing factual events.
3. **Covariant trigger surface.** The relevant bound is evaluated on the maximal-volume Cauchy surface of the causal diamond enclosing the entire superposition, not on the local diamonds of individual subsystems. Pruning is a single event on that surface, ontologically consistent by construction with relativistic causality.

This consolidated edition develops three principal contributions:

- Section 3 articulates the *trigger regime hierarchy*, making explicit the ontological position that isolated micro-events do not constitute facts independently of their macroscopic registration.
- Section 4 computes SR cosmological heating rates and shows they are $\sim 10^{15}$ – 10^{27} times smaller than CSL in diffuse environments (intergalactic medium, dark-matter halos, intracluster gas), automatically satisfying observational constraints that increasingly pressure CSL.
- Section 5 develops the geometric foundations of the *4D Rose* representation, showing that non-spherical, chiral, branching geometry is structurally required by the physics of wavefunctions undergoing parity violation and irreversible amplitude growth.

A separate but related phenomenological model within the same informational substrate, addressing the late-time suppression of structure growth (the S_8 tension) through gradients of the saturation scalar, is described in companion work [16] and discussed in scope-defining terms in Section 7. A speculative appendix (Appendix A) discusses whether the empirically stochastic phenomenology of SR may rest on a deeper deterministic substrate, an open question relevant for foundational interpretation but not affecting the empirical predictions developed here.

2 The Pruning Mechanism

2.1 Decoherence rate from gravitational self-energy

For a Gaussian superposition of two wavepackets of width σ_x centered at $\pm d/2$, the gravitationally-induced decoherence rate has the leading-order form [3, 4, 5]:

$$\Gamma_{\text{grav}} \sim \frac{Gm^2d^2}{\hbar\sigma_x^3} \quad (d \lesssim \sigma_x), \quad (2)$$

which generalizes the original Diósi–Penrose formula to the extended-superposition regime. Dimensional verification: with $[G] = \text{m}^3 \text{kg}^{-1} \text{s}^{-2}$, $[\hbar] = \text{J} \cdot \text{s}$, the combination $Gm^2d^2/(\hbar\sigma_x^3)$ has units of s^{-1} , as required for a rate. In the limit $d \rightarrow 0$ (no spatial separation), the rate vanishes as expected: gravitational distinguishability is absent.

2.2 Informational modulation: the saturation factor

In SR, this rate is modulated by the informational saturation principle. The local von Neumann entropy of a system grows as branches become distinguishable; the local Bekenstein bound $I_{\text{Bek}} = 2\pi RE/(\hbar c \ln 2)$ caps the information that can be encoded in a region of energy E and bounding radius R . Pruning occurs when the saturation ratio

$$\chi(t) = \frac{S_{\text{vN}}(t)}{I_{\text{Bek}}} \quad (3)$$

reaches a fraction η of unity. We adopt $\eta \sim 0.1$ as the fiducial value (constraints discussed in Section 6.4).

The modulated decoherence rate is therefore:

$$\Gamma_{\text{SR}} = \frac{1}{\eta} \frac{Gm^2 d^2}{\hbar \sigma_x^3}, \quad t_{\text{SR}} = \eta \frac{\hbar \sigma_x^3}{Gm^2 d^2}. \quad (4)$$

The factor η encodes the fraction of the Bekenstein bound at which pruning fires: smaller η means earlier pruning (more sensitive trigger), corresponding to faster collapse.

2.3 Pruning as ontological selection

When the saturation condition is reached, an irreversible pruning event occurs on the causal-diamond surface enclosing the superposition:

- The event is irreversible: branches eliminated in pruning do not return.
- It is informational, not interactional: no field-mediated mechanism is invoked.
- It selects branches with Born probabilities $|c_i|^2$ from the superposition coefficients.
- It produces a definite macroscopic register, completing the transition from amplitude to fact.

2.4 Comparison with Diósi–Penrose

The standard Diósi–Penrose collapse rate is

$$\Gamma_{\text{DP}} \sim \frac{Gm^2}{\hbar R_0}, \quad (5)$$

where R_0 is a regularization length (typically nuclear scale, $R_0 \sim 1$ fm). In the regime $\sigma_x \sim d \sim R_0$, Equations (4) and (5) converge up to $\mathcal{O}(1)$ factors. In the extended-superposition regime $d \gg \sigma_x$, SR predicts substantially faster collapse than DP due to the d^2/σ_x^3 amplification, while in the compressed regime $d \ll \sigma_x$, SR predicts substantially slower collapse.

The d/σ_x ratio constitutes a clean experimental discriminator: in matter-wave interferometry with controlled separation and wavepacket width, SR and DP predict distinguishable visibility decay rates as the two parameters are varied independently.

3 The Trigger Regime Hierarchy

A central conceptual question, sharpened by physical objections concerning isolated radioactive nuclei in deep vacuum, is: *when does pruning actually fire?* The answer determines what SR claims about the ontology of microscopic events.

3.1 The isolated-microsystem problem

Consider a Co-60 nucleus in deep cosmic vacuum, far from any detector. Standard quantum mechanics treats the decay amplitude $|\text{undecayed}\rangle \oplus |\text{decayed} + \beta\rangle$ as evolving unitarily. In SR, the local von Neumann entropy associated with this isolated nucleus is of order a few bits, while the Bekenstein bound for the nuclear region is of order 10^9 bits. The saturation ratio χ is therefore $\sim 10^{-9}$, far from triggering pruning at any reasonable η .

Pruning does not fire on the isolated nucleus.

This is not a flaw of the framework; it is the structural consequence of the trigger being informational. The empirically observed exponential decay statistics emerge not because the nucleus produces facts in vacuum, but because the amplitude of the decayed branch evolves unitarily and is registered when the emitted β ultimately interacts with a macroscopic detector whose χ saturates locally.

3.2 Hierarchy of regimes

We articulate the trigger regime hierarchy explicitly:

1. **Microscopic isolated systems** ($\chi \ll 1$, vacuum environment): amplitude evolves unitarily; pruning does not fire; no factual event occurs in the system itself. SR is observationally indistinguishable from standard QM in this regime. *Examples:* isolated radioactive nuclei in intergalactic vacuum, individual photons in transit, neutrinos in flight from cosmic sources.
2. **Microscopic systems with environment** (χ small but ambient decoherence active): amplitude evolves unitarily, environmental decoherence amplifies S_{vN} in correlated environment, pruning eventually fires when the environment-system register saturates its bound. Decay statistics are governed by amplitude evolution; pruning timing is governed by environmental amplification. SR remains indistinguishable from standard QM with environmental decoherence in this regime. *Examples:* typical laboratory experiments with detection apparatus.
3. **Mesoscopic isolated systems** (χ approaches unity intrinsically): pruning rate competes with environmental decoherence; the framework predicts measurable deviation from standard quantum predictions. *This is the unique discrimination regime.* Mass scale $m \sim 10^9 - 10^{11}$ amu, separation $d \sim 100$ nm– μ m, coherence times \sim seconds to minutes. *Examples:* MAQRO macroscopic quantum superposition program, advanced levitated optomechanics with macroscopic spatial separation.
4. **Macroscopic systems:** χ saturates rapidly; pruning fires effectively instantaneously by any objective-collapse mechanism. SR is observationally indistinguishable from CSL, DP, or environmental decoherence in this regime.
5. **Black-hole horizons:** $\chi = 1$ by construction; pruning operates as the steady-state thermodynamic mechanism of horizon information processing. This regime connects with Jacobson-style derivations of gravitational dynamics from horizon thermodynamics [11].

3.3 Ontological position on factuality

A consequence of this hierarchy is the explicit ontological claim that *microscopic events isolated from macroscopic registration do not constitute facts*. The amplitude of a decay or transition evolves freely, but the factuality of the individual event emerges only at the pruning of a macroscopic register.

This places SR closer to relational and Everettian interpretations than to GRW/CSL, while preserving the objective branch-selection structure that distinguishes it from pure many-worlds.

Empirically, the framework reproduces all observed decay statistics (cosmogenic isotope ratios, atmospheric muon decay, neutrino oscillation endpoints) because the amplitude evolution that determines these statistics is unmodified.

3.4 Empirical equivalence in the decay sector

The cosmogenic isotope record (e.g., ^{10}Be in meteorites, ^{60}Fe in deep-sea crusts), atmospheric muon decay profiles, neutrino flavor oscillations from solar and astrophysical sources, and ultra-high-energy cosmic-ray composition all reflect amplitude evolution along propagation paths of cosmological length. SR predicts identical statistics to standard QM for these observations, since:

- Amplitude evolution governs decay branching probabilities.
- Pruning fires only at the macroscopic detector, not in transit.
- The two ontological readings (factual decay during transit vs. amplitude propagation with factuality at registration) yield identical observational statistics.

Direct discrimination would require maintaining mother–daughter nuclear coherence over a half-life timescale in genuine isolation, which exceeds current technological capability by many orders of magnitude. *The decay sector is therefore not an SR test regime*; discrimination is concentrated in the mesoscopic-superposition regime (Section 6).

4 Cosmological Heating: Structural Suppression

A standard pressure point on objective-collapse models is the cosmological heating they predict. CSL adds energy to every particle at a fixed rate $\lambda_{\text{CSL}} \sim 10^{-16} \text{ s}^{-1}$, producing cumulative heating in diffuse media (intergalactic medium, dark-matter halos, intracluster gas) over Hubble time. Constraints from Lyman- α forest measurements, CMB spectral distortions, and intracluster temperature limits have progressively tightened the allowed CSL parameter range [6, 7].

We show here that SR, by contrast, is structurally suppressed by 15–27 orders of magnitude relative to CSL in the same environments, automatically satisfying these constraints without parameter tuning.

4.1 Heating rate per particle

The energy injected per pruning event in SR is bounded by the gravitational self-energy distinguishing branches:

$$E_{\text{SR,event}} \sim \frac{Gm^2}{d_{\text{eff}}}, \quad (6)$$

where d_{eff} is the relevant separation scale (of order σ_x for thermal-de-Broglie superpositions). The pruning rate is given by Eq. (4), yielding a heating rate per particle:

$$\left(\frac{dE}{dt} \right)_{\text{SR}} \sim \Gamma_{\text{SR}} \cdot E_{\text{SR,event}}. \quad (7)$$

4.2 Comparison across cosmological environments

Table 1 compares SR and CSL heating rates in three representative diffuse cosmological environments. SR is structurally suppressed across all of them, with the suppression growing in more dilute regimes.

The cumulative SR heating over Hubble time ($t_H \sim 4 \times 10^{17} \text{ s}$) in the IGM yields $\Delta T \sim 10^{-30} \text{ K}$ per particle, far below the sensitivity of any cosmological probe. SR therefore satisfies all current observational constraints derived from Lyman- α forest measurements, CMB Compton- y distortions, and intracluster temperature bounds without parameter adjustment.

Table 1: Heating rate per particle in diffuse cosmological environments. SR is structurally suppressed relative to CSL by 15–27 orders of magnitude. Values use $\eta = 0.1$ and $\lambda_{\text{CSL}} = 10^{-16} \text{ s}^{-1}$, $r_C = 10^{-7} \text{ m}$. σ_x is the thermal de Broglie length for thermal media and $\hbar/(mv)$ for collisionless particles.

Environment	$(dE/dt)_{\text{SR}}$ [W/part]	$(dE/dt)_{\text{CSL}}$ [W/part]	SR/CSL ratio
IGM ($T \sim 10^4 \text{ K}$, proton)	$\sim 7 \times 10^{-71}$	$\sim 3 \times 10^{-44}$	$\sim 10^{-27}$
Cluster gas ($T \sim 10^7 \text{ K}$, proton)	$\sim 7 \times 10^{-68}$	$\sim 3 \times 10^{-44}$	$\sim 10^{-24}$
DM halo (WIMP, $m \sim 100 \text{ GeV}$)	$\sim 1.5 \times 10^{-57}$	$\sim 2 \times 10^{-42}$	$\sim 10^{-15}$

4.3 Structural reason

The suppression is not parametric; it is structural. CSL imposes a fixed rate per particle independent of environment. SR conditions the rate on the saturation ratio χ , which is suppressed in all coherent and dilute regimes:

- Coherent macroscopic states (BEC, superfluid, Cooper-pair condensate) have S_{vN} collectively suppressed below the trigger.
- Dilute thermal states (IGM, halo, cluster) have small effective separation $d \sim \sigma_x$ on per-particle scale, suppressing the $d^2/\sigma_x^3 \rightarrow 1/\sigma_x$ factor.
- Both regimes therefore yield SR rates dramatically below CSL.

This is in marked contrast to CSL, whose baseline parameters increasingly conflict with cosmological bounds and require continuous parameter retuning to remain consistent.

5 Geometric Foundations: The 4D Rose

The visualization adopted throughout this work represents the configuration space of the wavefunction during the period preceding pruning as a *4D Rose*: a non-uniform, chiral, branching, directionally-growing geometric object. We argue here that this representation is not poetic but *structurally faithful* to the physics, and that spherical alternatives (balloons, expanding shells, isotropic blooms) fail to capture properties that the physics demands.

5.1 Properties required of the geometry

A faithful representation of a physical wavefunction undergoing amplitude growth and parity-violating decays should exhibit the following five structural features:

1. **Non-uniform branching.** Amplitudes for distinct branches are generically unequal. The geometry must accommodate “petals” of unequal weight, not concentric isotropic shells.
2. **Chirality.** The weak interaction violates parity [10] with polarized ^{60}Co ; β emission distributions are intrinsically asymmetric under spatial reflection. The geometry must support a definite handedness.
3. **Directional growth.** Pruning is irreversible; the amplitude evolution that precedes it is in turn temporally asymmetric. The geometry must grow forward, not radially out in equilibrium.
4. **Diffuse boundary.** A wavefunction has no sharp surface; amplitudes decay smoothly. The geometry must have a fractal or fading boundary, not a thin spherical shell.

5. **Identifiable origin vertex.** Amplitudes emanate from a generation event (interaction vertex, decay event). The geometry must have a well-defined center distinguished from its periphery.

5.2 Why spherical alternatives fail

A balloon (thin spherical shell) violates (1), (2), and (4) at minimum. A growing sphere violates (1) and (2). A diffusion cloud violates (1), (2), and (3). None of the standard symmetric visualizations capture the physics.

A rose, by contrast, is a real geometric object that simultaneously exhibits unequal petals, definite chirality, directional opening, fractal boundary, and identifiable origin. The choice is not metaphorical; it is the simplest familiar object whose structural properties match all five requirements simultaneously.

5.3 Connection to parity violation

The chirality requirement is particularly notable. The discovery that the weak interaction violates parity [10] established that nature, at the most fundamental level, does not emit isotropic “balloons” of β radiation: it emits objects with definite handedness. The 4D Rose representation incorporates this from the outset, treating chirality as a structural feature of physical wavefunction geometry rather than a coincidence to be added afterward.

This is consistent with the broader observation that fundamental processes generically break discrete symmetries. The 4D Rose is not committed to any specific mechanism for these breakings; it merely commits to a geometric language that does not artificially impose symmetries the physics does not have.

5.4 Pruning as petal selection

In this language, a pruning event is the selection of a single petal from the multi-petaled rose. Other petals (other branches) are eliminated; the surviving petal is recorded as the factual event. The macroscopic detector or environment provides the “observer position” from which the rose is viewed, but the selection itself is not perspective-dependent — it occurs on the global causal diamond surface, not in the detector.

The visualization is therefore not a mere illustration but an active conceptual tool: it makes manifest the asymmetry, chirality, and selectiveness that distinguish SR from symmetric spreading-then-projecting models.

6 Experimental Window: MAQRO and the d^2/σ_x^3 Signature

The discrimination regime for SR is the mesoscopic-superposition range identified in Section 3. We quantify here the predicted signatures and their relation to active experimental programs.

6.1 The d^2/σ_x^3 scaling

The collapse-rate formula of Eq. (4) contains a distinctive d^2/σ_x^3 dependence absent from standard DP and from CSL. For fixed mass, varying d and σ_x independently changes the SR collapse time in ways that the other models do not reproduce. This provides a clean experimental signature: in matter-wave interferometry with controlled wavepacket squeezing and separation control, SR predicts collapse-time variations that DP and CSL do not.

6.2 Consistency with current levitated optomechanics

We compare SR predictions with the published parameters of state-of-the-art levitated optomechanics experiments [13, 14]. For a 143-nm silica particle ($m \sim 3.4 \times 10^{-18}$ kg, corresponding to $\sim 2 \times 10^9$ amu) cooled to motional ground state, with later controlled delocalization to $\sigma_x \sim 73$ pm:

Table 2: Predicted decoherence rates in current levitated optomechanics, compared with observed decoherence floor. SR is currently below but within ~ 3 – 4 orders of magnitude of the observable threshold; the regime where SR becomes detectable requires extension to controlled macroscopic separation $d \gtrsim 100$ nm.

Quantity	Delić et al. 2020 ($\sigma_x \sim 3$ pm, $d \sim \sigma_x$)	Rossi et al. 2025 ($\sigma_x \sim 73$ pm, $d \sim \sigma_x$)
Observed decoherence floor	$\sim 2 \times 10^4 \text{ s}^{-1}$	$\sim 10^3 \text{ s}^{-1}$
DP prediction ($R_0 = 1$ fm)	$\sim 2 \text{ s}^{-1}$	$\sim 0.1 \text{ s}^{-1}$
SR ($\eta = 0.1$)	$\sim \mathbf{24} \text{ s}^{-1}$	$\sim \mathbf{1} \text{ s}^{-1}$

The current setups are at the threshold of SR sensitivity: SR predicts rates $\sim 10^{-3}$ below the observed decoherence floor, consistent with non-detection but not orders-of-magnitude excluded. The transition to the SR-detectable regime requires controlled macroscopic spatial separation, not merely smaller wavepackets.

6.3 The MAQRO regime

The MAQRO program [12] targets parameters in the mesoscopic range. Table 3 shows SR predictions for representative target configurations.

Table 3: SR predictions for the MAQRO target parameter space. The collapse times are in the seconds-to-sub-second range, well above expected environmental decoherence floors of $\sim 10^{-4} \text{ s}^{-1}$ in microgravity ultra-high-vacuum conditions.

Configuration	m [amu]	σ_x / d	Γ_{SR}
Conservative	10^{10}	10 nm / 100 nm	$\sim 17 \text{ s}^{-1}$
Target	10^{11}	1 nm / 1 μm	$\sim 10^8 \text{ s}^{-1}$

The conservative configuration yields $t_{\text{SR}} \sim 60$ ms, five orders of magnitude above the environmental decoherence floor expected in MAQRO conditions. This represents a clean detection window for SR, distinguishable from both standard environmental decoherence and from the standard DP prediction at the same parameter point.

6.4 Empirical constraints on η

The fiducial value $\eta \sim 0.1$ is constrained from above and below by empirical considerations:

- **From below:** $\eta > 10^{-3}$ is required for definite classicality emergence in macroscopic systems within reasonable laboratory timescales.
- **From above:** $\eta < 1$ is required to avoid spurious collapse in coherent atomic and molecular systems where coherence is empirically observed.
- **Cross-check:** The Casini-bound analysis on conventional superconductors, high- T_c materials, and superfluid ^4He places the saturation ratio at $\chi \sim 10^{-4}$ – 10^{-1} , compatible with $\eta \sim 0.1$ as the threshold value.

The narrow window of allowed η values (roughly two decades) constitutes a structural constraint: SR is not a free-parameter fit but a tightly constrained framework with η as the only fundamental dimensionless parameter.

7 Framework Scope and Companion Cosmological Channel

A separate phenomenological model within the same informational substrate, developed in companion work [16], addresses the late-time suppression of structure growth (the S_8 tension) through gradients of the saturation scalar $\mathcal{I} = S_{\text{local}}/I_{\text{Boussso}}$ acting on geodesics via $u^\nu \nabla_\nu u^\mu = \alpha \nabla^\mu \mathcal{I}$.

We emphasize that the companion model and the present pruning mechanism are *dynamically independent channels* of the same informational saturation principle:

- **Channel A (this work):** discrete branch pruning when S_{local} saturates the bound. Operates at any scale where $\chi \rightarrow \eta$. Governs wavefunction collapse.
- **Channel B (companion):** continuous geometric backreaction from $\nabla \mathcal{I}$ gradients in the unsaturated regime. Operates only at cosmological scales with structure-induced gradients. Governs late-time structure-growth modifications.

The two channels share the saturation scalar \mathcal{I} as an ontological substrate, but the pruning channel does not modify cosmological structure growth (the relevant scales are too dilute, $\chi \ll 1$ globally), and the gradient channel does not produce wavefunction collapse (the gradients are continuous, not threshold-triggered). Articulating this scope explicitly forecloses the criticism of framework inflation and clarifies the predictions made in each domain.

7.1 Domains the framework does not address

We emphasize, in the spirit of explicit scope, that SR/ICR is *not designed to address* a class of cosmological problems whose natural solutions live at the level of the pre-recombination Lagrangian or background expansion equation:

- The Hubble tension (H_0 discrepancy) is a pre-recombination or local-physics problem; channel-A pruning does not act on background dynamics, and channel-B gradients are too small at cosmological times to generate the required modifications.
- The JWST early-galaxy abundance is a primordial structure-formation problem at $z \gtrsim 8$; both SR channels are structurally inadequate for this regime by the same arguments.
- Baryogenesis and the matter-antimatter asymmetry live at the level of CP-violating Lagrangian terms; SR is CPT-invariant by construction and offers no quantitative handle on the asymmetry magnitude.

These limitations are structural, not failures of effort. Knowing where the framework does not apply is part of what makes its positive predictions defensible.

8 Discussion

8.1 Summary of core claims

We have presented a consolidated formulation of Stochastic Rupture in which:

- Pruning is triggered by informational saturation on covariant causal-diamond surfaces, with a single dimensionless parameter η .

- The dynamical rate, dimensionally consistent and DP-extended, is $\Gamma_{\text{SR}} \sim Gm^2d^2/(\eta \hbar \sigma_x^3)$, recovering standard DP in the regime $d \sim \sigma_x$ and predicting distinguishable behavior in the extended-superposition regime.
- The trigger regime hierarchy explicitly establishes that isolated micro-events do not generate facts independently of macroscopic registration, aligning SR with relational interpretations while preserving objective branch selection.
- Cosmological heating is structurally suppressed by 15–27 orders of magnitude relative to CSL, automatically satisfying observational bounds that increasingly pressure competing models.
- The 4D Rose representation is structurally faithful to parity-violating, irreversibly-growing wavefunction geometry, not metaphorical decoration.
- The unique experimental discrimination regime is mesoscopic superposition with macroscopic separation, corresponding to the MAQRO experimental program.

8.2 Energy conservation

Unlike CSL, where each localization event injects a fixed mean energy into the center-of-mass mode, SR pruning selects branches with Born probabilities and conserves energy on average across events. The variance per event is bounded by the gravitational distinguishability of branches, $\Delta E \sim Gm^2/d$, controlled by the dynamics rather than imposed externally. The cosmological heating analysis of Section 4 confirms that the integrated energy injection over Hubble time is far below all observational sensitivities. SR therefore evades the energy-conservation critiques that pressure CSL.

8.3 Open problems and future work

Several directions are in active development:

- **Microscopic justification of η .** The fiducial value $\eta \sim 0.1$ is constrained empirically but lacks a derivation from deeper principles. Connecting η to a microscopic information-theoretic or gravitational origin remains open.
- **Reintroduction of the metric-information relation.** The relation $g_{\mu\nu} = \kappa \nabla_\mu \nabla_\nu (I_{\text{local}}/I_{\text{Planck}})$ in a Jacobson-style derivation of gravitational dynamics from informational pruning is in development for a future extension.
- **No-signaling theorem.** A rigorous demonstration that the modulated-rate dynamics is consistent with relativistic causality, in the spirit of established no-signaling results for stochastic collapse models, is in preparation.
- **Casini-bound consistency analysis.** Conventional superconductors, high- T_c materials, and superfluid ^4He cluster at $\chi \sim 10^{-4}$ to 10^{-1} , compatible with $\eta \sim 0.1$. Ultracold BECs and Fermi superfluids violate the Casini geometry in the thermal regime, requiring careful treatment in future versions.
- **Strongly gravitating regimes.** Extensions to neutron stars and black-hole horizons, where the saturation ratio approaches unity by structural rather than statistical mechanisms, are under development.

A On the Possibility of a Deterministic Substrate

The framework as developed above is empirically stochastic: pruning selects branches with Born probabilities, and the phenomenology reproduces standard quantum statistics. The question of whether this empirical stochasticity rests on a deeper deterministic substrate is, however, distinct from the question of whether the empirical phenomenology is stochastic.

We summarize a speculative direction in which the SR rate may be the effective expression of an underlying deterministic instability governed by a saturation-controlled feedback in the wavepacket-width dynamics. A toy model of this form is

$$\ddot{\sigma}_x = \frac{\hbar^2}{4m^2\sigma_x^3} - \frac{2\Lambda(\chi)}{m}\sigma_x, \quad (8)$$

with $\Lambda(\chi) = \Lambda_0 \chi/(1+\chi)$ saturating in the high- χ regime. In the saturated regime, this generates a deterministic harmonic contraction with a characteristic timescale $\propto m^{-1/2}$, distinct from the $\propto d^2/\sigma_x^3$ scaling of the empirical phenomenology.

If the empirical stochasticity emerges from sensitivity to environmental degrees of freedom or hidden initial conditions (in a manner reminiscent of 't Hooft's deterministic-substrate proposals), then the SR rate formula would be the effective stochastic envelope of a deterministic underlying dynamics.

We emphasize the speculative nature of this direction. The present manuscript defends the empirically stochastic SR framework on its own terms; the deterministic substrate hypothesis is offered as a possible conceptual refinement to be developed in future work, and not as a replacement for the established phenomenology. The technical derivation of Λ_0 from first principles, and the recovery of Born statistics from the deterministic flow, remain open problems in this exploratory direction.

Acknowledgments

The author thanks correspondents for critical engagement that sharpened the trigger regime hierarchy, the geometric foundations of the rose representation, the cosmological heating analysis, and the explicit articulation of framework scope presented in this consolidated edition.

References

- [1] G. C. Ghirardi, A. Rimini, and T. Weber, *Unified dynamics for microscopic and macroscopic systems*, Phys. Rev. D **34**, 470 (1986).
- [2] P. Pearle, *Combining stochastic dynamical state-vector reduction with spontaneous localization*, Phys. Rev. A **39**, 2277 (1989).
- [3] L. Diósi, *Models for universal reduction of macroscopic quantum fluctuations*, Phys. Rev. A **40**, 1165 (1989).
- [4] R. Penrose, *On gravity's role in quantum state reduction*, Gen. Rel. Grav. **28**, 581 (1996).
- [5] A. Bassi, K. Lochan, S. Satin, T. P. Singh, and H. Ulbricht, *Models of wave-function collapse, underlying theories, and experimental tests*, Rev. Mod. Phys. **85**, 471 (2013).
- [6] S. L. Adler and A. Bassi, *Collapse models with non-white noises*, J. Phys. A **40**, 15083 (2007).
- [7] S. Donadi et al., *Underground test of gravity-related wave function collapse*, Nature Physics **17**, 74 (2021).

- [8] J. D. Bekenstein, *Universal upper bound on the entropy-to-energy ratio for bounded systems*, Phys. Rev. D **23**, 287 (1981).
- [9] R. Bousso, *The holographic principle*, Rev. Mod. Phys. **74**, 825 (2002).
- [10] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, *Experimental Test of Parity Conservation in Beta Decay*, Phys. Rev. **105**, 1413 (1957).
- [11] T. Jacobson, *Thermodynamics of spacetime: The Einstein equation of state*, Phys. Rev. Lett. **75**, 1260 (1995).
- [12] R. Kaltenbaek et al., *MAQRO: Quantum Physics in Space*, Exper. Astron. **34**, 123 (2012).
- [13] U. Delić et al., *Cooling of a levitated nanoparticle to the motional quantum ground state*, Science **367**, 892 (2020).
- [14] M. Rossi et al., *Controlled delocalization of a levitated nanoparticle*, Phys. Rev. Lett. **135**, 083601 (2025).
- [15] M. Arndt and K. Hornberger, *Testing the limits of quantum mechanical superpositions*, Nature Physics **10**, 271 (2014).
- [16] G. Zambuzi, *Late-Time Suppression of Structure Growth from Informational Curvature Relaxation*, Zenodo (2026).