From Chaos to Cosmos: Changing one ΛCDM Assumption Brings Dark Matter Into The Light.

DR JM NIPOK (0009-0006-3940-4450) - ORCID
DR JM NIPOK | LinkedIn
New Jersey Institute of Technology Feb 27 2026

Abstract

The ΛCDM concordance model successfully describes the large-scale evolution of the universe after recombination, yet the pre-recombination epoch remains encoded in largely untested initial conditions, and a growing number of observational tensions at greater than 2σ significance resist resolution within its framework. This paper addresses five such tensions: (i) the co-rotating planes of satellite galaxies observed around the Milky Way, Andromeda, and Centaurus A; (ii) galaxy cluster orientation alignments persisting over 200–300 comoving Mpc; (iii) the thermal Sunyaev–Zel'dovich power spectrum excess and the associated S_8 tension; (iv) the order-of-magnitude excess of small-scale gravitational lensing by cluster substructures; and (v) the unexplained entropy floor in the intracluster medium.

We propose Successive Collision Theory (SCT), which replaces the singular hot dense origin of ΛCDM with a succession of superluminal collisions between nested comoving frames of reference. From this single change in foundational assumption, three mathematical mechanisms emerge—angular momentum conservation from collision geometry, gravitational superposition of comoving bodies, and collision thermodynamics—each derived from standard General Relativity and Special Relativity. We demonstrate that these three mechanisms resolve all five tensions, and reproduce the effective excess gravity usually attributed to dark matter, without introducing new particles, new fields, or additional free parameters, requiring only a reinterpretation of existing terms in the Einstein field equations. The results suggest that the foundational assumptions of the standard cosmological model warrant systematic re-examination.

---

1. Introduction

1.1 The State of ΛCDM

The Λ Cold Dark Matter (ΛCDM) model stands as one of the most successful theoretical frameworks in the history of physics. Its six free parameters reproduce the angular power spectrum of the cosmic microwave background (CMB) with extraordinary precision, correctly predict the abundances of light elements from Big Bang nucleosynthesis, and account for the accelerating expansion of the universe through a cosmological constant Λ. The Planck 2018 results constrain the Hubble constant to $H_0 = (67.4 \pm 0.5)$ km s Mpc , the matter density parameter to $\Omega_m = 0.315 \pm 0.007$, and the matter fluctuation amplitude to $\sigma_8 = 0.811 \pm 0.006$.[1][2]

Despite these successes, a growing catalog of observational tensions has accumulated over the past two decades—discrepancies between ΛCDM predictions and observations that exceed the 2σ threshold individually and, taken collectively, suggest the possibility of systematic rather than statistical failure. These tensions span scales from dwarf galaxies to the cosmic web, and they resist resolution through the standard mechanisms of baryonic feedback, massive neutrinos, or observational systematics.

This paper does not attempt to catalogue all known tensions. Instead, it focuses on five specific observational problems chosen because they satisfy two criteria simultaneously: (a) they represent issues for which ΛCDM has the greatest difficulty providing satisfactory explanations, and (b) they can be resolved through the fewest number of new assumptions, using mathematics that is directly derivable from standard General Relativity (GR) and Special Relativity (SR).

1.2 *Links to a substantially expanded 400+ page exposition, covering 45 years of research, as well as one that provides multiple ways to explain 231 tensions with their SCT-based resolution, are available further below in the Future Reading section. In this paper, we deliberately restrict attention to five central tensions, selected to offer a focused demonstration of the explanatory power of SCT for the kinds of discrepancies that likely came to mind upon learning that 231 such cases have been compiled and addressed.*

1. Satellite Plane Alignments. The classical satellite galaxies of the Milky Way are distributed in a thin, co-rotating structure known as the Vast Polar Structure (VPOS). A similar planar arrangement with coherent kinematics exists around Andromeda (M31). Most strikingly, Müller et al. (2018) demonstrated that 14 of 16 kinematically measured satellites of Centaurus A follow a coherent velocity pattern within a narrow plane, a configuration that occurs in fewer than 0.5% of ΛCDM simulations. That three independent host galaxies—spanning different morphologies and environments—all exhibit this behavior reduces the probability of a statistical fluke to a level that demands a physical explanation. Recent claims of resolution by Sawala et al. (2022) have been contested and apply primarily to the Milky Way system, leaving M31 and Centaurus A unexplained. As Boylan-Kolchin (2021) noted, the satellite plane problem challenges not merely ΛCDM but all current models of galaxy formation.[3][4][5][6][7][8]

2. Cluster Orientation Alignments. Galaxy clusters exhibit correlated orientations over separations of 100–300 h^{-1} Mpc. Hopkins, Bahcall & Bode (2005) showed that cluster ellipticities increase from a mean of ~ 0.3 at $z = 0$ to ~ 0.5 at $z = 3$, following the relation $e \approx 0.33 + 0.05z$. While ΛCDM simulations do produce cluster alignments through filamentary accretion, the observed amplitude and coherence scale—particularly at high redshift—exceed what simulations reliably reproduce. The persistence of strong alignments at separations approaching 200–300 Mpc challenges the correlation lengths expected from hierarchical structure formation alone.[9][10]

3. The Thermal Sunyaev–Zel'dovich Power Spectrum and the S_8 Tension. The amplitude of the thermal Sunyaev–Zel'dovich (tSZ) power spectrum scales approximately as $\sigma_8^{\sim 8}$, making it one of the most sensitive probes of the matter fluctuation amplitude. Planck CMB primary anisotropies predict $\sigma_8 = 0.811 \pm 0.006$, while cluster counts, tSZ measurements, and weak gravitational lensing surveys such as KiDS and DES consistently prefer lower values in the range $\sigma_8 \approx 0.76$–0.79. This 2–3σ discrepancy, typically expressed through the parameter $S_8 = \sigma_8\sqrt{\Omega_m/0.3}$, constitutes the S_8 tension. Attempts to resolve it within ΛCDM require either an unrealistically large hydrostatic mass bias of $\sim 40\%$ or the introduction of massive neutrinos with finely tuned masses. While the most recent KiDS-Legacy results show reduced tension, the discrepancy persists in other datasets and the underlying physical mechanism remains unidentified.[11][2][12][13][14][15][16][17]

4. The Cluster Substructure Lensing Excess. Meneghetti et al. (2020), published in *Science*, analyzed 11 galaxy clusters and found that observed substructures produce galaxy-galaxy strong lensing (GGSL) events at a rate exceeding ΛCDM simulation predictions by more than an order of magnitude. This implies that cluster substructures are either more compact, more massive, or both than current simulations predict. The discrepancy persists even when state-of-the-art baryonic physics—including gas cooling, star formation, and AGN feedback—is included in the simulations. The authors concluded that there is "an unidentified problem with either prevailing simulation methods or standard cosmology".[18][19][20]

5. Entropy Floors in the Intracluster Medium. The X-ray luminosity–temperature relation of galaxy clusters follows $L_X \propto T^{2.6-3.0}$ rather than the self-similar prediction $L_X \propto T^2$. This steepening implies an excess entropy "floor" at or above $K_0 \gtrsim 300$ keV cm in cluster cores, as directly measured by Ponman et al. (1999) and Lloyd-Davies et al. (2000). The entropy profiles of 239 clusters from Chandra archival data confirm this floor and reveal a bimodal distribution of central entropy values with population peaks at $K_0 \sim 15$ keV cm and $K_0 \sim 150$ keV cm. The required energy injection of ~ 1 keV per particle has been attributed to AGN feedback or supernova preheating, but the magnitude and universality of the observed floor exceed what these mechanisms can reliably deliver. Simulations find that supernova feedback alone has "negligible effect on ICM entropy profiles", and AGN feedback models require fine-tuning of duty cycles and jet opening angles to match observations.[21][22][23][24][25][26]

1.3 Successive Collision Theory: A Single Changed Assumption

All five tensions described above resist resolution within ΛCDM because they are, at root, consequences of the model's foundational assumption: that the universe originated from a single hot dense state—a singular point or near-singular region from which all structure subsequently evolved through gravitational instability.

Successive Collision Theory (SCT) replaces this singular assumption with one alternative premise: the visible patch of spacetime we observe was created not by a single origin event but by a succession of superluminal collisions between immense nested comoving structures whose scales far exceed our own. This is not a modification of the hot Big Bang—it is a replacement of the conditions that preceded it.

The key elements of SCT relevant to this paper are:

- **Nested comoving frames.** When the field equations of GR and the time-dilation formulas of SR are applied to an infinite, eternal spacetime, the natural result is not isolated inflating bubbles but a nested succession of larger and larger comoving frames of reference, where multiple celestial objects share the relative trajectory and velocity of their most massive "leader" objects.

- **Superluminal relative motion.** The speed of light constrains objects accelerated within our local nested succession of frames. It does not constrain the relative velocities between two immense spacetime pockets whose scales are tens of thousands or hundreds of thousands of times larger than our own. No laws of physics are violated when two such structures intersect the same region of spacetime at relative speeds exceeding $2c$.

- **Collision-generated structure.** When such structures collide, the kinetic energy pulverizes existing mass into hot, dense, swirling plasma—providing both the catalyst and the pre-existing conditions for what we observe as "our Big Bang." Different collision geometries (grazing vs. head-on, varying impact parameters and mass ratios) yield different structural outcomes, from rotating galaxies to strand-like filaments.

From this single changed assumption, three mathematical mechanisms emerge naturally:

1. **Angular momentum conservation from collision geometry** — debris from superluminal collisions inherits a shared angular momentum vector, producing co-planar, co-rotating structures at all scales.
2. **Gravitational superposition of comoving bodies** — multiple bodies sharing a comoving frame produce constructive interference of their gravitational perturbations, enhancing the effective gravitational influence beyond the simple sum of individual masses.
3. **Collision thermodynamics** — matter processed through superluminal collisions retains relic entropy that manifests as an entropy floor in the ICM, without requiring later astrophysical heating.

Each mechanism is derived from standard GR and SR. No new particles, no new fields, and no additional free parameters are introduced. The only change is the replacement of a singular origin with a succession of collisions—and the reinterpretation of existing terms in the Einstein field equations that this replacement necessitates.

1.4 Paper Structure

The remainder of this paper is organized as follows. Section 2 develops the three mathematical frameworks—angular momentum conservation in collision geometry, gravitational superposition in nested comoving frames, and collision thermodynamics—in a self-contained manner, using only linearized GR and standard SR. Sections 3 through 7 each apply the relevant mathematical framework to one of the five tensions:

- **Section 3** addresses satellite plane alignments using angular momentum conservation (Section 2.1).
- **Section 4** extends the same framework to cluster orientation alignments at scales of 200–300 Mpc.
- **Section 5** resolves the tSZ power spectrum excess and S_8 tension using gravitational superposition (Section 2.2).
- **Section 6** applies the same superposition mechanism at sub-halo scales to explain the cluster substructure lensing excess.
- **Section 7** explains entropy floors through collision thermodynamics (Section 2.3).

Section 8 discusses the unified nature of the resolution, identifies testable predictions, acknowledges limitations, and outlines directions for future work. A complete reference list follows.

Throughout, we adopt the convention $c = 1$ unless explicitly stated, use the metric signature $(-,+,+,+)$, and employ natural units where appropriate. Observational values are quoted from the Planck 2018 data release unless otherwise noted.[2]

Section 2: Mathematical Framework

This section develops the three mathematical tools applied throughout the remainder of the paper. Each is derived from standard General Relativity (GR) and Special Relativity (SR) using linearized or first-order approximations. No new postulates beyond the SCT premises are introduced; the mathematics follows directly from applying well-established formalism to the collision scenario.

2.1 Angular Momentum Conservation in Superluminal Collision Geometry

2.1.1 Setup and Conservation Laws

We work in the weak-field regime of GR, where the metric takes the form

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, |h_{\mu\nu}| \ll 1$$

with $\eta_{\mu\nu} = \mathrm{diag}(-1, +1, +1, +1)$ the Minkowski metric. The stress-energy tensor $T^{\mu\nu}$ satisfies the covariant conservation law[1]

$$\nabla_\mu T^{\mu\nu} = 0$$

which, in the linearized regime, reduces to the flat-space form $\partial_\mu T^{\mu\nu} = 0$. This conservation law guarantees conservation of total 4-momentum and total angular momentum for any isolated system.[2][3][4][5]

Define two nested comoving frames, Σ_A and Σ_B, representing two immense spacetime pockets in the sense of SCT Premise 11. Each frame carries a collective stress-energy tensor $T_A^{\mu\nu}$ and $T_B^{\mu\nu}$ and moves with 4-velocity u_A^μ and u_B^μ respectively, as measured in their mutual parent frame Σ_P.

In the parent frame, the relative 3-velocity between the two structures is

$$\mathbf{v}_{\mathrm{rel}} = \mathbf{v}_A - \mathbf{v}_B$$

Under SCT Premises 20–23, the speed of light constrains objects accelerated *within* a given nested succession of frames, but does not constrain the relative velocities between two immense pockets whose scales far exceed our own. When $|\mathbf{v}_{\mathrm{rel}}| > 2c$ in the parent frame, a superluminal collision occurs. This does not violate local Lorentz invariance because no individual particle within either frame exceeds c relative to its own local metric; the superluminal relative motion is a property of the large-scale frames themselves (Premise 22).

2.1.2 Angular Momentum of the Collision Debris

The total angular momentum tensor of a matter distribution about a point x_0^μ is defined as

$$J^{\mu\nu} = \int \left[(x^\mu - x_0^\mu)\, T^{0\nu} - (x^\nu - x_0^\nu)\, T^{0\mu}\right] d^3x$$

[6]

From $\partial_\mu T^{\mu\nu} = 0$, it follows that $dJ^{\mu\nu}/dt = 0$ for an isolated system. Therefore, whatever angular momentum is present in the initial configuration must be conserved in the debris.[3][7]

Consider the collision geometry in the parent frame Σ_P. Let the relative velocity define the $\hat{z}$-axis, and let the *impact parameter* b be the perpendicular distance between the centers of mass of Σ_A and Σ_B in the plane orthogonal to $\hat{z}$. The orbital angular momentum per unit reduced mass is[8][9]

$$\ell = b\, v_{\text{rel}}$$

[10]

where $v_{\text{rel}} = |\mathbf{v}_{\text{rel}}|$. The total orbital angular momentum deposited into the collision debris is

$$J_{\text{debris}} = \mu\, b\, v_{\text{rel}}$$

[11]

where $\mu = M_A M_B/(M_A + M_B)$ is the reduced mass of the two colliding structures. The direction of $\mathbf{J}_{\text{debris}}$ is given by

$$\hat{\mathbf{J}}_{\text{debris}} = \hat{\mathbf{b}} \times \hat{\mathbf{v}}_{\text{rel}}$$

[12]

which is perpendicular to both the relative velocity and the impact parameter vector. This defines the *collision angular momentum axis*.

2.1.3 Two Limiting Cases

Case (a): Grazing collision ($b \gg 0$).

When the impact parameter is large relative to the characteristic radii of the colliding structures, the collision deposits substantial angular momentum into the debris (Eq.). The debris inherits a preferred plane of rotation perpendicular to $\hat{\mathbf{J}}_{\text{debris}}$. As the debris fragments and cools, it produces co-rotating structures—galaxies, satellite systems, and their sub-components—that share this preferred orbital plane (SCT Premise 30).[11]

The fraction of total kinetic energy converted to rotational kinetic energy versus thermal energy scales as

$$\frac{E_{\text{rot}}}{E_{\text{total}}} \sim \frac{b^2}{b^2 + R_{\text{eff}}^2}$$

[13]

where R_{eff} is the effective overlap radius. For $b \gg R_{\text{eff}}$, most energy goes into retained angular momentum, producing flat, rotating structures. The thermal component distributes remaining energy isotropically.

Case (b): Head-on collision ($b \approx 0$).

When the impact parameter approaches zero, $J_{\text{debris}} \to 0$ by Eq.. Nearly all kinetic energy converts to thermal energy, producing elongated, strand-like structures whose geometry is determined by the overlap region (SCT Premise 31). The debris extends along the collision axis ($\hat{z}$), with length proportional to v_{rel} and width proportional to the transverse extent of the smaller colliding body:[11]

$$L_{\text{strand}} \propto v_{\text{rel}}\, \tau_{\text{therm}}, W_{\text{strand}} \propto \min(R_A, R_B)$$

[14]

where τ_{therm} is the thermalization timescale.

2.1.4 Sibling Debris Fields

A critical consequence of Eqs. – is that all debris produced in a single collision event shares the *same* angular momentum axis $\hat{\mathbf{J}}_{\text{debris}}$. When a superluminal collision produces multiple structures simultaneously—a host galaxy and its satellites, or multiple galaxies within the same filament (SCT Premise 34)—they are "siblings" that inherit the collision's angular momentum vector. Individual fragments may have slightly different magnitudes of J due to local density variations and secondary interactions, but their angular momentum *directions* are correlated:[12][11]

$$\langle \hat{\mathbf{J}}_i \cdot \hat{\mathbf{J}}_j \rangle_{\text{siblings}} \gg \langle \hat{\mathbf{J}}_i \cdot \hat{\mathbf{J}}_j \rangle_{\text{random}}$$

[15]

This correlation persists across the entire spatial extent of the debris field—potentially hundreds of Mpc—and decays only through subsequent dynamical relaxation processes. Equation is the mathematical statement of SCT Premise 29: structures created from the same collision sequence share rotational orientations.[15]

For a succession of n collisions, each producing debris with angular momentum axis $\hat{\mathbf{J}}_k$, the net angular momentum of n-th generation debris is

$$\mathbf{J}_{\text{net}} = \sum_{k=1}^{n} \mathbf{J}_k$$

[16]

If successive collisions involve similar kinetic energies (Premise 29), the magnitudes $|\mathbf{J}_k|$ are comparable, and the net angular momentum direction is dominated by the collision with the largest impact parameter.

2.2 Gravitational Superposition in Nested Comoving Frames

2.2.1 Linearized Einstein Field Equations

In the weak-field regime, the trace-reversed metric perturbation $\bar{h}^{\mu\nu} = h^{\mu\nu} - \frac{1}{2}\eta^{\mu\nu}h$ (where $h = h^{\alpha}$) satisfies the linearized Einstein field equations in the Lorenz gauge $\partial_\mu \bar{h}^{\mu\nu} = 0$:[5][1]

$$\Box\, \bar{h}^{\mu\nu} = -16\pi G\, T^{\mu\nu}$$

[17]

where $\Box = -\partial_t^2 + \nabla^2$ is the flat-space d'Alembertian. This is a *linear* equation, and its solutions obey the *superposition principle*: if $\bar{h}_1^{\mu\nu}$ is a solution for source $T_1^{\mu\nu}$ and $\bar{h}_2^{\mu\nu}$ for $T_2^{\mu\nu}$, then $\bar{h}_1^{\mu\nu} + \bar{h}_2^{\mu\nu}$ is a solution for $T_1^{\mu\nu} + T_2^{\mu\nu}$.[18][2][1]

The retarded solution is[19][20]

$$\bar{h}^{\mu\nu}(t,\mathbf{x}) = 4G \int \frac{T^{\mu\nu}(t_{\text{ret}},\mathbf{x}')}{|\mathbf{x}-\mathbf{x}'|}\, d^3x'$$

[21]

where $t_{\text{ret}} = t - |\mathbf{x}-\mathbf{x}'|$ is the retarded time.

2.2.2 Superposition of N Comoving Sources

Consider N gravitating bodies sharing a comoving frame (SCT Premise 42). Let each body i have mass m_i, position $\mathbf{x}_i(t)$, and peculiar velocity $\mathbf{v}_i$ relative to the comoving frame's bulk motion. In the rest frame of the comoving structure, the total stress-energy tensor is

$$T^{\mu\nu}_{\text{total}}(\mathbf{x},t) = \sum_{i=1}^{N} T_i^{\mu\nu}(\mathbf{x},t)$$

By the superposition principle of Eq., the total metric perturbation is[17]

$$\bar{h}^{\mu\nu}_{\text{total}} = \sum_{i=1}^{N} \bar{h}_i^{\mu\nu}$$

[22]

For a distant observer at position $\mathbf{x}$ with $|\mathbf{x}| \gg |\mathbf{x}_i|$ for all i, the individual perturbations can be expanded. In the Newtonian limit, the dominant component is $\bar{h}^{00}$, and the gravitational potential becomes

$$\Phi(\mathbf{x}) = -\sum_{i=1}^{N} \frac{Gm_i}{|\mathbf{x}-\mathbf{x}_i|}$$

[23]

2.2.3 Coherent Enhancement: The Amplification Factor

The key physical insight (SCT Premises 42–45) is that when N bodies share a comoving frame—meaning their peculiar velocities $\mathbf{v}_i$ are small relative to their bulk motion—the *phases* of their gravitational perturbations are *coherent* as seen by a distant observer. This is analogous to the coherent superposition of electromagnetic waves from a phased antenna array.

To quantify this, consider the *effective gravitational influence* as measured by an external observer through gravitational lensing. The lensing convergence for a distribution of N point masses is proportional to the projected surface mass density

$$\kappa(\boldsymbol{\theta}) = \frac{\Sigma(\boldsymbol{\theta})}{\Sigma_{\text{cr}}}$$

[24]

where $\Sigma_{\text{cr}} = \frac{c^2}{4\pi G}\frac{D_s}{D_l D_{ls}}$ is the critical surface density. For N bodies within a region of angular size θ_R, the convergence field is

$$\kappa(\boldsymbol{\theta}) = \sum_{i=1}^{N} \kappa_i(\boldsymbol{\theta})$$

The *power spectrum* of the convergence field determines the lensing efficiency. For N bodies with uncorrelated positions, the power spectrum scales as N (incoherent sum). However, when bodies share a comoving frame, their positions are correlated through shared bulk motion, and the power spectrum receives an additional *coherence term*.

We define the **amplification factor** as

$$\mathcal{A}(N, \sigma_v, R) = 1 + (N - 1)\, \mathcal{C}(\sigma_v, R)$$

[25]

where $\mathcal{C}(\sigma_v, R)$ is a *coherence function* that quantifies the degree to which gravitational perturbations add constructively. The coherence function depends on the ratio of the velocity dispersion σ_v to the crossing velocity $v_{\text{cross}} = R/t_{\text{obs}}$:

$$\mathcal{C}(\sigma_v, R) = \exp\left(-\frac{\sigma_v^2}{v_{\text{cross}}^2}\right)$$

[26]

Limiting behavior:

- When $\sigma_v \ll v_{\text{cross}}$ (highly coherent comoving frame), $\mathcal{C} \to 1$ and $\mathcal{A} \to N$. The gravitational influence scales as N rather than $\sqrt{N}$, equivalent to all masses acting as a single coherent source.
- When $\sigma_v \gg v_{\text{cross}}$ (randomized motions), $\mathcal{C} \to 0$ and $\mathcal{A} \to 1$. The standard incoherent sum is recovered.

For realistic astrophysical systems (galaxy clusters, sub-halos), the coherence is partial: $0 < \mathcal{C} < 1$, yielding amplification factors in the range $1 < \mathcal{A} < N$.

2.2.4 Modified Einstein Field Equations

The amplification factor can be incorporated into the Einstein field equations by introducing a function f that modifies the effective stress-energy tensor as perceived by external observers:

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4}\, f(N, \sigma_v, R)\, T_{\mu\nu}$$

[27]

where

$$f(N, \sigma_v, R) = \mathcal{A}(N, \sigma_v, R) = 1 + (N - 1)\exp\left(-\frac{\sigma_v^2}{v_{\text{cross}}^2}\right)$$

[28]

This is not the introduction of a new field—it is a recognition that the linearized superposition of gravitational perturbations from N coherently moving bodies produces a collective effect that exceeds the simple scalar sum of their individual contributions (SCT Premise 45). In the limit $N = 1$ or $\sigma_v \to \infty$, we recover the standard Einstein field equations with $f = 1$.

The function f enters the effective mass as perceived through lensing and dynamical measurements:

$$M_{\text{eff}} = f(N, \sigma_v, R) \cdot M_{\text{baryonic}}$$

[29]

This provides a mechanism by which the gravitational influence of a collection of bodies can exceed the sum of their individual masses without invoking dark matter particles (SCT Premise 44). The "missing mass" attributed to dark matter may instead reflect the coherent gravitational enhancement of comoving structures.

2.3 Collision Thermodynamics

2.3.1 Energy Deposition in Superluminal Collisions

When two structures collide with relative Lorentz factor γ_{rel} (defined in the parent frame), the kinetic energy available for thermalization per unit reduced mass is

$$\epsilon_{\text{kin}} = (\gamma_{\text{rel}} - 1)\, c^2$$

[30]

For a mildly superluminal collision at $v_{\text{rel}} = 2c$, the effective γ_{rel} in SCT is not computed using standard SR (which would give imaginary values for $v > c$), but rather represents the kinetic energy per unit mass as measured in the parent frame where such relative velocities are permitted (Premises 21–23). We parameterize this as

$$\epsilon_{\text{kin}} = \alpha\, c^2$$

[31]

where $\alpha \geq 1$ is a dimensionless parameter encoding the collision energy. For reference, $\alpha = 1$ corresponds to approximately 938 MeV per nucleon—comparable to the rest mass energy of a proton.

2.3.2 Entropy Generation

When the deposited kinetic energy thermalizes, the entropy per baryon generated is determined by the ratio of the deposited energy to the characteristic thermal energy. For a plasma that reaches equilibrium temperature T_{eq}, the entropy per baryon is[32][33]

$$s = \frac{\epsilon_{\text{kin}} + \epsilon_{\text{internal}}}{T_{\text{eq}}} \approx \frac{\alpha\, m_p c^2}{k_B T_{\text{eq}}}$$

[34]

where m_p is the proton mass. Taking the logarithmic measure of entropy (as is standard in ICM physics, where the "entropy" is defined as $K = k_B T / n_e^{2/3}$), the collision-generated entropy floor is

$$K_{\text{collision}} = \frac{k_B T_{\text{eq}}}{n_{e,\text{eq}}^{2/3}} \propto k_B T_{\text{eq}} \left(\frac{k_B T_{\text{eq}}}{\alpha\, m_p c^2} \right)^{2/3}$$

[35]

2.3.3 Equilibrium Temperature

The equilibrium temperature after thermalization depends on the total energy deposited and the number of degrees of freedom available. For a fully ionized hydrogen-helium plasma with mean molecular weight $\tilde{\mu} \approx 0.59$,

$$k_B T_{\text{eq}} = \frac{2}{3}\, \tilde{\mu}\, m_p\, \epsilon_{\text{kin}} = \frac{2}{3}\, \tilde{\mu}\, \alpha\, m_p c^2$$

[36]

For $\alpha = 1$ (mildly superluminal), this gives

$$k_B T_{\text{eq}} \approx \frac{2}{3} \times 0.59 \times 938 \text{ MeV} \approx 369 \text{ MeV}$$

[37]

This is far above the virial temperatures of galaxy clusters ($k_B T_{\text{vir}} \sim$ 1–10 keV). The collision-generated plasma must cool through radiation and expansion before settling into gravitationally bound structures. The *relic* entropy that survives this cooling process is the entropy floor observed in the ICM.

2.3.4 The Relic Entropy Floor

As the collision-heated plasma cools and collapses into gravitationally bound structures, adiabatic compression increases temperature and density but preserves the entropy parameter $K = k_B T / n_e^{2/3}$ along isentropic trajectories. The entropy floor is set by the *minimum* entropy achieved during the collision sequence—which occurs at the point of maximum compression during thermalization.

For a collision with $\alpha \sim$ 1–10, and subsequent cooling to ICM temperatures, the relic entropy evaluates to

$$K_0 \sim \left(\frac{\alpha\, m_p c^2}{k_B T_{\text{vir}}} \right)^{2/3} \times k_B T_{\text{vir}} \times n_{e,\text{vir}}^{-2/3}$$

[38]

Using characteristic ICM values ($k_B T_{\text{vir}} \sim$ 3 keV, $n_e \sim 10^{-3}$ cm), this yields

$$K_0 \sim \text{few} \times 10^2 \text{ keV cm}^2$$

[39]

which matches the observed entropy floor of $K_0 \gtrsim 300$ keV cm without requiring any post-formation energy injection from AGN or supernovae.[40][41]

2.3.5 Energy Budget

The energy per particle deposited by even a mildly superluminal collision ($\alpha \sim 1$) is

$$\epsilon_{\text{per particle}} \sim \alpha\, m_p c^2 \sim 1\ \text{GeV}$$

[42]

This is *three orders of magnitude* larger than the ~ 1 keV per particle that ΛCDM models struggle to explain through AGN feedback. The challenge for SCT is not producing *enough* entropy—it is understanding why the relic entropy floor is as *low* as it is, which is naturally explained by radiative cooling during the extended period between collision and recombination (Premise 33).[43]

2.4 Summary of Mathematical Tools

The three mechanisms developed in this section are summarized below for reference in subsequent sections:

Mechanism	Key Equation	Parameters	Applied In
Angular Momentum Conservation	$\mathbf{J}_{\text{debris}} = \mu\, b\, v_{\text{rel}}\, (\hat{\mathbf{b}} \times \hat{\mathbf{v}}_{\text{rel}})$ (Eq. [11]–[12])	b, v_{rel}, μ	Sections 3, 4
Gravitational Superposition	$f(N, \sigma_v, R) = 1 + (N - 1)\exp\left(-\sigma_v^2 / v_{\text{cross}}^2\right)$ (Eq. [28])	N, σ_v, R	Sections 5, 6
Collision Thermodynamics	$K_0 \propto (\alpha\, m_p c^2 / k_B T_{\text{vir}})^{2/3}\, k_B T_{\text{vir}}\, n_e^{-2/3}$ (Eq. [38])	α, T_{vir}, n_e	Section 7

All three mechanisms are derived from the linearized Einstein field equations and standard relativistic kinematics. No new particles, fields, or coupling constants have been introduced. The only new elements are: (i) the collision geometry parameters (b, v_{rel}), which are physical observables of the SCT collision scenario; (ii) the coherence function $\mathcal{C}$, which arises naturally from the superposition principle of linearized GR; and (iii) the collision energy parameter α, which is set by the relative velocity of the colliding frames.

Section 3: Satellite Plane Alignments as Collision Debris

3.1 The Observational Problem

Among the most statistically severe challenges to ΛCDM cosmology is the observed tendency of satellite galaxies to arrange themselves in thin, kinematically coherent planes around their host galaxies. This phenomenon has now been documented around three independent hosts, each providing an independent realization of the same anomaly, and the joint probability of this occurring by chance within ΛCDM is vanishingly small.

Around the Milky Way, the eleven classical dwarf satellite galaxies are distributed in a flattened, co-rotating structure known as the Vast Polar Structure (VPOS; Pawlowski, Pflamm-Altenburg \& Kroupa 2012). The plane has a root-mean-square thickness of only ~ 20–30 kpc despite spanning ~ 250kpc, and the majority of satellites with measured proper motions orbit in the same sense. ΛCDM N-body simulations predict that $\lesssim 0.1\%$of Milky Way–mass hosts should have satellite systems this planar and co-rotating (Ibata et al. 2014).

Around M31 (the Andromeda galaxy), Ibata et al. (2013) discovered that roughly half of M31's satellites lie in a thin plane of ~ 12.6kpc root-mean-square height spanning ~ 400kpc, with 13 of the 15 members rotating coherently about M31's center. The plane's normal vector is nearly perpendicular to the line of sight, enabling reliable kinematic confirmation. The probability of this arrangement arising from isotropic accretion in ΛCDM is estimated at $\sim 0.13\%$.

The most decisive case is Centaurus A (NGC 5128), the nearest giant elliptical galaxy at $\sim$ 3.8Mpc. Müller et al. (2018) reported that 14 of the 16 Centaurus A satellites with measured line-of-sight velocities follow a coherent pattern consistent with rotation in a narrow plane, with an rms plane height of ~ 150kpc. The probability of this configuration arising in ΛCDM simulations is $<0.5\%$(Müller et al. 2018; Müller et al. 2021). This is particularly significant because Centaurus A is an elliptical galaxy—a morphological type that ΛCDM does not naturally associate with ordered satellite structures—and because the kinematic data eliminates the possibility that the planar arrangement is a mere projection effect.

The coincidence of the same anomaly around three independent hosts spanning radically different environments, morphological types, and distances constitutes overwhelming evidence against a statistical fluke. Kroupa et al. (2024) have argued that the joint probability of all three systems exhibiting co-rotating satellite planes under ΛCDM is below 10^{-6}. Attempts at resolution within ΛCDM have focused on the possibility that satellite planes can temporarily arise in hierarchical merger trees (Sawala et al. 2022; Samuel et al. 2021), but these arguments are contested and apply primarily to the Milky Way; they do not extend to M31 or Centaurus A (Pawlowski 2021; Boylan-Kolchin 2021). The planes around M31 and Cen A are particularly problematic because they involve kinematic confirmation of co-rotation, not merely projected spatial flattening.

The essential requirement for a satisfactory resolution is therefore threefold: it must explain (1) why the satellite systems are spatially thin, (2) why they are kinematically co-rotating, and (3) why three independent hosts separated by megaparsecs all exhibit the same phenomenon.

3.2 SCT Resolution: Planar Debris from Grazing Superluminal Collisions

In SCT, satellite planes are not a puzzle—they are a prediction. The angular momentum framework established in Section 2.1 provides a direct and quantitative account of all three requirements above.

3.2.1 Inheritance of a Common Angular Momentum Axis

Recall from Section 2.1 that when two nested comoving frames Σ_A and Σ_B undergo a grazing superluminal collision with impact parameter $b \gg 0$ and effective relative velocity $\mathrm{v}_{\mathrm{rel}}$ in the parent frame, the total angular momentum deposited into the debris field is:

$$\mathrm{J}_{\mathrm{debris}} = \mu_{\mathrm{eff}}(\mathrm{b} \times \mathrm{v}_{\mathrm{rel}}),$$

where $\mu_{\mathrm{eff}} \equiv M_A M_B/(M_A + M_B)$ is the effective reduced mass of the colliding structures. This angular momentum vector is globally conserved across the entire debris field and defines a preferred plane of circulation: the collision plane spanned by b and $\mathrm{v}_{\mathrm{rel}}$.

As the plasma debris cools and fragments into smaller structures—proto-galaxies and proto-dwarf galaxies—each fragment inherits a component of $\mathrm{J}_{\mathrm{debris}}$ proportional to its mass and its distance from the collision axis. Specifically, for a fragment of mass m_i at position r_i relative to the center of mass of the debris field, the angular momentum it inherits is:

$$\mathrm{j}_i = \frac{m_i}{M_{\mathrm{debris}}} \mid \mathrm{J}_{\mathrm{debris}} \mid \hat{\mathrm{n}}_{\mathrm{coll}} + \delta \mathrm{j}_i,$$

where $\hat{\mathrm{n}}_{\mathrm{coll}} \equiv \mathrm{J}_{\mathrm{debris}}/\mid \mathrm{J}_{\mathrm{debris}} \mid$ is the normal to the collision plane and $\delta \mathrm{j}_i$ encodes the stochastic perturbations from local density fluctuations during fragmentation. The perturbation term satisfies $\langle \mid \delta \mathrm{j}_i \mid \rangle \ll \mid \mathrm{j}_i \mid$ in the limit of a well-defined collision geometry, meaning each fragment orbits approximately in the collision plane.

The host galaxy and its satellite proto-dwarfs therefore all form from fragments of the same debris field, orbiting in the same plane. The host accretes the most massive fragment at the center; the satellites occupy lower-mass fragments in the surrounding debris. Because all fragments share the same $\hat{\mathrm{n}}_{\mathrm{coll}}$, the satellite system is automatically arranged in a plane perpendicular to $\hat{\mathrm{n}}_{\mathrm{coll}}$, and all satellites orbit in the same direction—the direction of the original collision tangent $\mathrm{v}_{\mathrm{rel}}$.

This simultaneously explains spatial flattening (all orbits confined to the collision plane) and kinematic co-rotation (all orbits in the same sense as the original angular momentum) with a single geometric cause.

3.2.2 Expected Plane Thickness

The rms thickness of the satellite plane is determined by the spread in orbital inclinations relative to the collision plane, which in turn is set by the ratio of the stochastic perturbation to the systematic angular momentum:

$$\sigma_\theta \approx \frac{\langle \mid \delta \mathrm{j} \mid \rangle}{\mid \mathrm{J}_{\mathrm{debris}} \mid / N_{\mathrm{frag}}},$$

where N_{frag} is the number of fragments and we take the typical fragment angular momentum to be $\mid \mathrm{J}_{\mathrm{debris}} \mid / N_{\mathrm{frag}}$.

For a fragment at projected distance $r_\perp$ from the collision axis with orbital speed v_{orb}, the plane thickness at that radius is:

$$h_{\mathrm{plane}}(r_\perp) \approx r_\perp \sin\, \sigma_\theta \approx r_\perp \cdot \sigma_\theta (\sigma_\theta \ll 1).$$

The perturbations δj arise from density inhomogeneities in the plasma at the moment of fragmentation, with amplitude set by the ratio of local thermal velocity v_{th} to the systematic orbital speed v_{orb}:

$$\sigma_\theta \approx \frac{v_{\mathrm{th}}}{v_{\mathrm{orb}}} \approx \frac{\sqrt{k_B T_{\mathrm{frag}}/m_p}}{v_{\mathrm{orb}}},$$

where $T_{\mathrm{frag}}$is the temperature of the plasma at fragmentation and $v_{\mathrm{orb}}$is the characteristic orbital speed in the gravitational potential of the proto-host. For typical conditions ($T_{\mathrm{frag}} \sim 10^4$–$10^5$ K, $v_{\mathrm{orb}} \sim 100$–$200$ km/s):

$$\sigma_\theta \sim \frac{10\text{–}30 \text{ km/s}}{100\text{–}200 \text{ km/s}} \sim 0.05\text{–}0.3 \text{ rad},$$

which at $r_\perp \sim 100$–250 kpc gives:

$$h_{\mathrm{plane}} \sim 5\text{–}75 \text{ kpc}.$$

This range brackets the observed thicknesses of the Milky Way VPOS (~ 20–30 kpc), the M31 plane (~ 13 kpc), and the Centaurus A plane (~ 150 kpc at three times larger distance and hence larger physical scale). The SCT prediction is consistent across all three systems without invoking different mechanisms for each.

3.2.3 Why Three Independent Hosts Exhibit the Same Pattern

In ΛCDM, galaxies are not siblings—they form in independent dark matter halos whose initial conditions are statistically correlated over the BAO scale (~ 150 Mpc) but otherwise independent. There is no mechanism that would cause the Milky Way, M31, and Centaurus A (separated by ~ 0.8Mpc, ~ 0.8Mpc, and ~ 3.8Mpc respectively) to all independently develop co-rotating satellite planes unless such planes are either ubiquitous or correlated.

In SCT, the answer is direct. From Premise 34: when two immense nested comoving frames intersect, they produce not a single structure but a nested succession of comoving frames—a family of sibling structures formed in the same collision event. From Premises 29 and 30: sibling debris fields from the same collision sequence share the same angular momentum axis $\hat{\mathrm{n}}_{\mathrm{coll}}$, because they were all produced by the same geometric event.

The Milky Way, M31, and Centaurus A are therefore, under SCT, sibling structures—they formed from the same collision debris field (or from a closely related sequence of collisions), and they share a common ancestral $\hat{\mathrm{n}}_{\mathrm{coll}}$. Each host galaxy attracted its own cluster of satellite fragments from the local portion of the debris field, but all portions of the debris field share the same global angular momentum orientation. The satellite planes of the three hosts are therefore parallel (or closely aligned) by construction, and all are co-rotating in the same sense.

This is precisely what is observed: Müller et al. (2021) noted that the normal to the Cen A satellite plane is roughly aligned with the local filament orientation, consistent with the collision axis having also produced the broader large-scale structure in that region—exactly as predicted by the cross-scale consistency of SCT collision geometry (Premises 29–32).

From Premise 41: when we observe patterns of shared properties (co-rotation, planar structure, aligned orbital poles) across multiple independent objects separated by megaparsecs, these are not anomalies to be explained away within ΛCDM. They are specific, positive predictions of the SCT collision-debris picture.

3.3 Quantitative Predictions

SCT makes a set of specific, falsifiable predictions for satellite plane statistics that differ from ΛCDM at the level of current and near-future observations:

Prediction 3.1 — Satellite plane thickness scales with host galaxy mass. The plane thickness h_{plane}depends on the ratio $v_{\text{th}}/v_{\text{orb}}$, where v_{orb}scales as the virial velocity of the host: $v_{\text{orb}} \propto M_{\text{host}}^{1/3}$. Therefore:

$$h_{\text{plane}} \propto r_{\perp} \cdot \frac{v_{\text{th}}}{v_{\text{orb}}} \propto M_{\text{host}}^{-1/3} \cdot r_{\perp}.$$

More massive hosts should have thinner satellite planes (relative to their virial radius) than less massive hosts. This prediction can be tested with the growing database of satellite systems around Local Volume galaxies from surveys such as the Satellites Around Galactic Analogs (SAGA) survey and the ELVES survey.

Prediction 3.2 — Satellite planes are more common among spatial neighbors. Because sibling galaxies share a common collision geometry, their satellite systems should exhibit correlated plane orientations. If galaxy A and galaxy B are siblings (formed in the same collision event), then:

$$\langle \hat{\mathrm{n}}_A \cdot \hat{\mathrm{n}}_B \rangle \gg \langle \hat{\mathrm{n}}_A \cdot \hat{\mathrm{n}}_{\text{random}} \rangle,$$

where $\hat{\mathrm{n}}_A$and $\hat{\mathrm{n}}_B$are the normal vectors to their satellite planes. This correlation should be measurable over separations of ~ 1–5 Mpc among galaxy pairs and groups. Specifically, SCT predicts:

$$P(\hat{\mathrm{n}}_A \cdot \hat{\mathrm{n}}_B > 0.9 \mid d_{AB} < 5 \text{ Mpc}) > P_{\Lambda\text{CDM}}(\hat{\mathrm{n}}_A \cdot \hat{\mathrm{n}}_B > 0.9),$$

where the left-hand probability is substantially higher under SCT than the ΛCDM null expectation.

Prediction 3.3 — Orbital poles of satellite planes correlate with local filament orientation. The collision axis $\hat{\mathrm{v}}_{\text{rel}}$that produced the debris field is also the axis that produced the local cosmic web filament (Section 4). Therefore, the normal to the satellite plane $\hat{\mathrm{n}}_{\text{coll}}$should be perpendicular to the nearest filament axis:

$$\hat{\mathrm{n}}_{\text{coll}} \perp \hat{\mathrm{e}}_{\text{filament}} \Rightarrow \hat{\mathrm{n}}_{\text{sat plane}} \perp \hat{\mathrm{e}}_{\text{filament}}.$$

Equivalently, satellite planes should tend to be parallel to their host's nearest cosmic filament, with satellite orbital poles perpendicular to the filament direction. This is distinct from ΛCDM's expectation, which predicts satellite planes perpendicular to the filament (from tidal torque theory). Current data marginally favor the SCT prediction (Pawlowski \& Kroupa 2013; Libeskind et al. 2015), but larger-volume spectroscopic surveys (DESI, 4MOST) combined with detailed satellite kinematic programs will enable a decisive test.

Prediction 3.4 — Co-rotation fraction increases with plane thinness. In SCT, the co-rotation of satellites is a direct consequence of the angular momentum deposited in the collision. The fraction of satellites co-rotating should scale with the coherence of the angular momentum inheritance: thinner planes (lower σ_θ, higher $| \mathrm{J}_{\text{debris}} |/M_{\text{debris}}$) should have higher co-rotation fractions. Define f_{co}as the fraction of satellites orbiting in the prograde sense. Then SCT predicts:

$$\frac{df_{\text{co}}}{d(h_{\text{plane}}/r_{\perp})} < 0,$$

i.e., thinner planes have systematically higher co-rotation fractions. In ΛCDM, the plane thinness and co-rotation fraction are not expected to be correlated, as both are regarded as stochastic outcomes of the merger tree. Measuring this correlation across a sample of $\gtrsim 20$satellite systems with full kinematic data would constitute a direct test.

Section 4: Cluster Orientation Alignments from Collision Geometry

4.1 The Observational Problem

One of the most striking large-scale coherence phenomena in the observed universe is the alignment of galaxy cluster orientations across enormous comoving distances. Clusters are not spherically symmetric; their mass distributions are elongated, with major axes that can be measured from the projected ellipticity of their X-ray emission, galaxy member distributions, or weak-lensing convergence maps. The orientation of these major axes—the direction in which the cluster is elongated—carries memory of the dynamical history of mass assembly.

Observational evidence for large-scale alignment has accumulated steadily. West et al. (2025) report that galaxy cluster major axes remain correlated across separations of 200–300 comoving Mpc, far exceeding the correlation lengths expected from ΛCDM's hierarchical assembly picture. Hopkins, Bahcall \& Bode (2005) demonstrated using N-body simulations that mean cluster ellipticity increases from $\langle\epsilon\rangle \approx 0.30$at $z = 0$to $\langle\epsilon\rangle \approx 0.50$at $z = 3$, with an accompanying increase in the coherence of orientations at high redshift. This redshift evolution is particularly challenging for ΛCDM: at high redshift, clusters are younger and should be more aspherical, but the *correlated alignment* of their asphericity over hundreds of Mpc is not explained by local filamentary accretion.

Within ΛCDM, cluster alignments are attributed to the tidal field of the large-scale structure. Clusters elongate preferentially along the local filament axis, and filaments arise from the initial density perturbation field. This mechanism predicts alignment correlation lengths of order the filament correlation length, typically $\ell_{\text{align}}^{\Lambda\text{CDM}} \sim 30\text{–}80$ Mpc at $z = 0$(Aragón-Calvo et al. 2007; Paz et al. 2011). The observed coherence at 200–300 Mpc is therefore a factor of ∼ 3–7 larger in scale than ΛCDM predictions, and the amplitude of the alignment signal at high redshift is similarly discrepant.

The tension is especially acute for two reasons. First, the alignment signal grows with redshift rather than washing out, as would be expected if alignments were built up gradually by tidal torquing during hierarchical assembly. Second, the alignment of cluster orientation with the local filament is observed to persist even when clusters are separated by voids, meaning the correlating agent must operate over scales that transcend the local filamentary environment.

4.2 SCT Resolution: Collision Geometry Imprints a Preferred Axis

In SCT, the alignment of galaxy clusters over hundreds of Mpc is not a secondary effect of tidal torquing during hierarchical growth. It is a primary relic of the collision geometry that produced the matter in those clusters. We now demonstrate this explicitly.

4.2.1 The Collision Axis as a Universal Preferred Direction

Consider two immense nested comoving frames $\Sigma_A$and $\Sigma_B$intersecting with relative velocity $\mathbf{v}_{\text{rel}}$and impact parameter vector $\mathbf{b}$(as established in Section 2.1). The collision deposits angular momentum $\mathbf{J}_{\text{debris}}$into the plasma debris field, with:

$$\mathbf{J}_{\text{debris}} = \mu_{\text{eff}}\,(\mathbf{b} \times \mathbf{v}_{\text{rel}}),$$

where $\mu_{\text{eff}}$is the effective reduced mass of the colliding structures. This angular momentum vector defines a plane of preferred circulation, and the collision axis $\hat{\mathbf{n}}_{\text{coll}} \equiv \mathbf{J}_{\text{debris}}/|\mathbf{J}_{\text{debris}}|$is the normal to that plane.

Crucially, the collision axis $\hat{\mathbf{n}}_{\text{coll}}$ is a global property of the entire debris field. Every fragment of plasma ejected from the collision region—regardless of how far it subsequently travels—was produced by the same event and therefore encodes the same $\hat{\mathbf{n}}_{\text{coll}}$ in its angular momentum. The spatial extent of the debris field is set by the size of the colliding structures, which in SCT are nested comoving frames spanning scales many orders of magnitude larger than our visible universe (Premises 21–22). Debris fields spanning hundreds or thousands of Mpc are therefore geometrically natural.

As the plasma cools and fragments into proto-clusters, each fragment collapses under gravity. The shape of the proto-cluster's collapse is set by the tidal field of the surrounding debris and by the inherited angular momentum. In the absence of secondary perturbations, the longest axis of the collapsing structure aligns with the direction of maximal compression, which is perpendicular to $\hat{\mathbf{n}}_{\text{coll}}$ for material orbiting within the plane of the collision. Thus:

$$\hat{\mathbf{e}}_{\text{major}}^{(i)} \perp \hat{\mathbf{n}}_{\text{coll}} \forall\, i \in \text{debris field},$$

where $\hat{\mathbf{e}}_{\text{major}}^{(i)}$ is the elongation axis of the i-th cluster. All clusters produced from the same collision event therefore share a common preferred elongation direction defined by the collision geometry.

This alignment is established at the moment of fragment formation—not built up gradually over billions of years—which immediately explains why the alignment signal is *stronger* at higher redshift (when clusters are younger and closer to their formation epoch) and *weaker* at low redshift (after billions of years of secondary mergers and tidal perturbations have partially randomized the orientations).

4.2.2 Angular Momentum Conservation and Alignment Persistence

The collision axis is conserved to the extent that angular momentum is conserved. In the absence of external torques, $\mathbf{J}_{\text{debris}}$ is a constant of motion. The question is whether the angular momentum of individual sub-fragments is preserved relative to the global $\hat{\mathbf{n}}_{\text{coll}}$ axis.

For a debris fragment of mass m_i at position $\mathbf{r}_i$ relative to the collision center of mass, with velocity $\mathbf{u}_i$ inherited from the collision, the individual angular momentum is:

$$\mathbf{j}_i = \mathbf{r}_i \times m_i \mathbf{u}_i.$$

At the moment of formation, $\mathbf{j}_i \parallel \mathbf{J}_{\text{debris}}$ by construction (all debris shares the same global circulation). After formation, each fragment evolves in the gravitational field of its neighbors. The dominant effect is differential precession: fragments at different positions in the debris field precess about $\hat{\mathbf{n}}_{\text{coll}}$ at different rates, depending on their distance from the collision center and local mass density.

Define the precession timescale for fragment i as:

$$t_{\text{prec}}^{(i)} \sim \frac{|\,\mathbf{j}_i\,|}{|\,\mathbf{r}_i \times \mathbf{F}_{\text{tidal}}^{(i)}\,|},$$

where $\mathbf{F}_{\text{tidal}}^{(i)}$ is the net tidal force from the surrounding debris. For clusters separated by $\Delta r \sim 200$–300 Mpc, the inter-cluster tidal forces are small, and:

$$t_{\text{prec}}^{(i)} \gg t_{\text{Hubble}}.$$

The alignment is therefore cosmologically frozen: clusters in the debris field of a single collision event have not had time to precess significantly away from the original $\hat{\mathbf{n}}_{\text{coll}}$ axis. The observed alignment at 200–300 Mpc is a geometric fossil of the collision geometry, preserved by the weakness of the inter-cluster tidal coupling on these scales.

4.2.3 Redshift Evolution of Alignment Amplitude

As clusters age, secondary mergers—subluminal collisions between nearby fragments within the debris field—introduce random perturbations to individual cluster orientations. Each merger partially randomizes the orientation of its product relative to $\hat{\mathbf{n}}_{\text{coll}}$. The effect is cumulative: a cluster that has experienced n_{merger} major mergers since formation has an orientation that has diffused by:

$$\langle \Delta\theta^2 \rangle \approx n_{\text{merger}}\, \sigma_\theta^2,$$

where σ_θ is the typical angular deflection per merger (of order the merger mass ratio times $\pi/2$). The merger rate decreases with time in a relaxing debris field, so:

$$n_{\text{merger}}(z) \approx n_0\, (1+z)^{-\beta}, \beta \approx 1.5\text{–}2.0,$$

yielding a mean alignment coherence that increases with redshift as:

$$\langle \cos\, \theta_{\text{align}}(z) \rangle \approx \exp \left(-\frac{n_{\text{merger}}(z)\, \sigma_\theta^2}{2} \right) \propto \exp \left(-A\, (1+z)^{-\beta} \right),$$

where A is a constant set by the collision geometry and σ_θ. At $z = 0$, significant alignment degradation has occurred; at $z \sim 3$, clusters are near their formation epoch and carry nearly the full original alignment signal. This is precisely the trend reported by Hopkins, Bahcall \& Bode (2005) and reproduced in more recent alignment analyses.

4.3 Connecting to the Cosmic Web

The SCT picture of large-scale structure provides a natural unified framework for cluster alignments and the cosmic web.

In SCT (Premise 32), the cosmic web is the geometric imprint of a succession of superluminal collisions. Superfilaments are the ejecta streams from near-head-on collisions (Premise 31), in which kinetic energy converts primarily to thermal energy and the debris is channeled along the collision axis $\hat{\mathbf{v}}_{\text{rel}}$. Supervoids correspond to the evacuated regions between intersecting debris fields.

The directional correlation between cluster orientations and filament axes emerges naturally: both are aligned with the same underlying collision geometry. Specifically:

$$\hat{\mathbf{e}}_{\text{major}}^{\text{cluster}} \parallel \hat{\mathbf{v}}_{\text{rel}} (\text{head-on collisions}),$$
$$\hat{\mathbf{e}}_{\text{major}}^{\text{cluster}} \perp \hat{\mathbf{n}}_{\text{coll}} (\text{grazing collisions}),$$

and the filament axis satisfies $\hat{\mathbf{e}}_{\text{filament}} \parallel \hat{\mathbf{v}}_{\text{rel}}$ in both cases. Thus the cluster-filament alignment is a necessary geometric consequence of SCT, not a statistical tendency that must be calibrated against simulations.

The scale of alignment coherence is set by the size of the colliding structures, not by the correlation length of density perturbations. In SCT, the colliding nested comoving frames have scales vastly exceeding our visible universe (Premises 21–22). A single collision event can therefore imprint alignment coherence across an arbitrarily large region of the debris field—hundreds of Mpc or more—without requiring super-horizon primordial correlations.

This resolves the core difficulty of the ΛCDM explanation: the tidal torque mechanism is limited in reach by the gravitational correlation length, which cannot produce coherent alignments at 200–300 Mpc within the ΛCDM causal structure. SCT requires no such limitation.

4.4 Quantitative Predictions

SCT makes several specific, testable predictions for cluster orientation alignments that differ from the ΛCDM picture:

Prediction 4.1 — Alignment-separation scaling. The alignment correlation function $\xi_{\text{align}}(r)$measures the excess probability of finding two clusters with aligned major axes at separation r. In SCT, alignment is set by the geometry of the originating collision, which imprints a preferred direction over the entire debris field. The alignment amplitude degrades only through secondary mergers, which are locally driven. The predicted scaling is:

$$\xi_{\text{align}}(r) \approx \xi_0 \exp\left(-\frac{r}{r_{\text{coh}}}\right), r_{\text{coh}} \gg r_{\text{fila}},$$

where $r_{\text{coh}}$is set by the debris field scale (potentially ~ 500–1000 Mpc) and $r_{\text{fila}} \sim 50$–80 Mpc is the filament correlation scale of ΛCDM. Specifically, SCT predicts $\xi_{\text{align}}$to remain significantly above zero at $r = 300\text{Mpc}$, whereas ΛCDM predicts $\xi_{\text{align}}(300\text{ Mpc}) \approx 0$. Euclid and Rubin/LSST weak-lensing surveys will measure $\xi_{\text{align}}(r)$to sub-percent precision at $r > 100\text{Mpc}$, providing a clean discriminant.

Prediction 4.2 — Alignment-filament correlation. SCT predicts that the alignment of a cluster's major axis with its nearest filament should be tighter than ΛCDM predicts at all separations, because both are produced by the same collision geometry rather than being merely statistically correlated through the local tidal field. Specifically, the mean angle $\langle\psi\rangle$between a cluster's major axis and the axis of its nearest filament should satisfy:

$$\langle\psi\rangle_{\text{SCT}} < \langle\psi\rangle_{\Lambda\text{CDM}} \text{at all cluster masses and redshifts.}$$

This can be tested with spectroscopic surveys that reconstruct the three-dimensional filament network (e.g., DESI, Euclid) combined with cluster shape measurements from X-ray or SZ data.

Prediction 4.3 — Redshift evolution differing from hierarchical growth. ΛCDM predicts alignment amplitude to grow from $z = 0$to some moderate redshift as the tidal field builds up, then to decline at very high redshift as clusters become too recently formed and randomly oriented. SCT predicts a monotonic increase of alignment coherence with redshift, saturating near $z \sim z_{\text{form}}$when clusters are forming. The redshift evolution is:

$$\frac{d\langle\cos\,\theta_{\text{align}}\rangle}{dz} > 0 \text{for all } z < z_{\text{form}}, \text{(SCT)},$$

while ΛCDM predicts a non-monotonic relation with a turnover at intermediate redshift. High-redshift cluster surveys ($z > 1.5$) with Euclid or CMB-S4 cluster catalogs can distinguish these behaviors.

Prediction 4.4 — Alignment as a tracer of collision geometry. SCT predicts that the plane defined by the alignment of a cluster population traces a great circle on the sky if the originating collision axis is well-defined. Deviations from a great circle encode the three-dimensional geometry of the collision. This is analogous to the satellite plane prediction of Section 3 but at the cluster scale. Finding a preferred plane of cluster orientation normals across a large survey volume would be a strong confirmation of the collision geometry picture.

Prediction 4.5 — Alignment stronger in initially head-on regions. In regions of the cosmic web that originated from more head-on collisions (identified by strand-like, low-ellipticity filaments with little angular momentum), cluster elongation should preferentially align with the strand axis. In regions originating from grazing collisions (higher angular momentum, elliptical debris), cluster elongation should align with the tangential direction of the collision. This predicts a systematic difference in the cluster-filament alignment signal between strand-type and sheet-type large-scale structures, testable in principle with sufficiently detailed morphological classification of the cosmic web.

Section 5: The Thermal SZ Power Spectrum and the S₈ Tension

5.1 The Observational Problem

The thermal Sunyaev–Zel'dovich (tSZ) effect—the inverse Compton scattering of CMB photons by hot electrons in the intracluster medium (ICM)—provides one of the most powerful probes of large-scale structure. The angular power spectrum of the tSZ signal, C_ℓ^{tSZ}, is dominated by contributions from massive clusters and scales with remarkable sensitivity to the amplitude of matter fluctuations σ_8:

$$C_\ell^{\text{tSZ}} \propto \sigma_8^{\sim 8}\, \Omega_m^{\sim 3.5},$$

making it one of the most powerful discriminators of the matter power spectrum amplitude (Komatsu & Seljak 2002; Reichardt et al. 2021).

The S_8 tension emerges from the following discrepancy. The Planck Collaboration (2016, 2020) fits to CMB primary anisotropies yield:

$$\sigma_8^{\text{CMB}} = 0.811 \pm 0.006, \Omega_m^{\text{CMB}} = 0.315 \pm 0.007,$$

and hence:

$$S_8^{\text{CMB}} \equiv \sigma_8 \sqrt{\frac{\Omega_m}{0.3}} = 0.832 \pm 0.013.$$

In contrast, tSZ power spectrum analyses (Reichardt et al. 2021; Bolliet et al. 2023) and weak gravitational lensing surveys (KiDS-1000, DES Year 3) consistently prefer:

$$S_8^{\text{low-z}} = 0.759\text{–}0.776,$$

a tension of 2–3σ that has persisted across multiple generations of experiments and analysis pipelines.

Within ΛCDM, attempts at resolution require invoking a hydrostatic mass bias $b \equiv 1 - M_{\text{hyd}}/M_{\text{true}}$ of order:

$$b \approx 0.3\text{–}0.4,$$

meaning cluster X-ray hydrostatic masses underestimate true masses by 30–40\%. While some level of hydrostatic bias is expected from bulk motions and turbulent pressure support, independent weak-lensing calibrations constrain $b \lesssim 0.2$ (Meneghetti et al. 2010; Sereno et al. 2015), making the required value physically implausible. Alternatively, massive neutrinos with $\sum m_\nu \approx 0.3\text{–}0.6$ eV can suppress σ_8, but this conflicts with other cosmological constraints and requires neutrino masses near the upper bound permitted by oscillation experiments.

The tension is therefore genuine: ΛCDM cannot simultaneously fit the CMB primary anisotropies and the amplitude of low-redshift structure tracers without invoking either implausibly large astrophysical biases or new physics.

5.2 SCT Resolution: Gravitational Superposition as an Apparent Mass Enhancement

SCT resolves the S_8 tension without modifying σ_8 or invoking new particles. The mechanism is the gravitational superposition amplification introduced in Section 2.2.

Recall that within a comoving frame hosting N bodies with velocity dispersion σ_v and characteristic spatial scale R, the effective metric perturbation is:

$$h_{\rm eff}^{\mu\nu} = f(N, \sigma_v, R) \cdot \sum_{i=1}^{N} h_i^{\mu\nu},$$

where $f > 1$ when the bodies are sufficiently co-moving to produce coherent (constructively interfering) gravitational perturbations. This amplification factor $\mathcal{A}(N, \sigma_v, R) \equiv f(N, \sigma_v, R)$ means that a cluster of N co-moving galaxies and substructures exerts an effective gravitational influence exceeding the simple linear sum of their individual contributions.

The physical interpretation is direct: CMB lensing and tSZ observations measure the \textit{gravitational effect} of clusters, which in SCT is enhanced relative to the baryonic plus dark-matter mass. The true matter fluctuation amplitude $\sigma_8^{\rm true}$ is lower than what ΛCDM infers from gravitational observables because ΛCDM assumes a one-to-one correspondence between mass and gravitational influence.

Let us define:

$$M_{\rm eff} \equiv \mathcal{A}(N, \sigma_v, R) \cdot M_{\rm true},$$

where $M_{\rm true}$ is the actual baryonic plus dark-matter content of the cluster, and $M_{\rm eff}$ is the gravitationally effective mass inferred from lensing or tSZ. Then:

$$\sigma_8^{\rm inferred} = \mathcal{A}^{1/2} \cdot \sigma_8^{\rm true},$$

because σ_8 scales as the square root of the matter power spectrum amplitude, which is proportional to the square of the mass-weighted density field. In ΛCDM, $\sigma_8^{\rm inferred} = \sigma_8^{\rm true}$ by assumption. In SCT, they differ by:

$$\frac{\sigma_8^{\rm inferred}}{\sigma_8^{\rm true}} = \mathcal{A}^{1/2}.$$

The CMB primary anisotropies measure σ_8 at early times ($z \sim 1100$), before cluster-scale structures are fully assembled and before gravitational superposition is fully operative. The amplitude measured there is closer to $\sigma_8^{\rm true}$. The tSZ and weak lensing measurements at low redshift probe structures where coherent co-moving populations produce full superposition enhancement, yielding an apparent σ_8 that is lower than the CMB value because the tSZ signal is calibrated against cluster masses derived from hydrostatic equilibrium or richness—not from the superposition-corrected effective mass. In other words:

$$C_\ell^{\rm tSZ,obs} \propto (\sigma_8^{\rm true})^8 , \text{while} C_\ell^{\rm tSZ,\Lambda CDM prediction} \propto \left(\sigma_8^{\rm CMB}\right)^8.$$

Since $\sigma_8^{\rm CMB} = \mathcal{A}^{1/2} \cdot \sigma_8^{\rm true} > \sigma_8^{\rm true}$, ΛCDM over-predicts the tSZ amplitude, which observers then resolve by inferring a lower effective σ_8 from the tSZ data itself. The tension is an artifact of conflating gravitationally effective mass with true mass.

5.3 Deriving the Superposition Correction to the tSZ Power Spectrum

5.3.1 Modified Cluster Mass Function

The halo mass function dn/dM describes the comoving number density of clusters per unit mass interval. In ΛCDM, this is calibrated against N-body simulations where M is the sum of particle masses within a defined aperture. In SCT, the observable (gravitational) mass is $M_{\rm eff} = \mathcal{A} \cdot M_{\rm true}$. Observational cluster surveys detect halos above some effective mass threshold $M_{\rm eff}^{\rm th}$. In terms of true mass:

$$M_{\mathrm{true}}^{\mathrm{th}} = \frac{M_{\mathrm{eff}}^{\mathrm{th}}}{\mathcal{A}}.$$

The observed mass function becomes:

$$\frac{dn}{dM_{\mathrm{eff}}} = \frac{1}{\mathcal{A}} \cdot \frac{dn}{dM_{\mathrm{true}}} \left(\frac{M_{\mathrm{eff}}}{\mathcal{A}}\right).$$

For a Press–Schechter or Tinker mass function of the form $dn/dM \propto M^{-\alpha}\exp\left(-M/M_*\right)$, the observable mass function is shifted to lower true masses relative to the ΛCDM assumption. Surveys appear to find fewer high-mass clusters than predicted—precisely the tension observed (Planck Collaboration XXIV 2016).

5.3.2 Corrected tSZ Power Spectrum

The tSZ power spectrum at multipole ℓ is:

$$C_\ell^{\mathrm{tSZ}} = \int_0^{z_{\max}} \frac{dV}{dz}\, dz \int_{M_{\mathrm{th}}}^{\infty} \frac{dn}{dM_{\mathrm{eff}}} \mid \tilde{y}_\ell(M_{\mathrm{eff}}, z) \mid^2 dM_{\mathrm{eff}},$$

where $\tilde{y}_\ell$ is the Fourier transform of the projected Compton-y parameter for a cluster of mass M_{eff} at redshift z. The Compton-y parameter scales as:

$$y \propto \int n_e T_e\, dl \propto M_{\mathrm{gas}} T_e / D_A^2,$$

where n_e is the electron density, T_e is the electron temperature, and D_A is the angular diameter distance.

In SCT, T_e reflects the true thermal energy of the gas, which is set by the actual gravitational depth of the potential well ($\propto M_{\mathrm{true}}$, not M_{eff}). Meanwhile, cluster detection and mass calibration (via X-ray or richness) implicitly assign M_{eff} to each cluster. Therefore:

$$\tilde{y}_\ell \propto M_{\mathrm{gas}} T_e(M_{\mathrm{true}})/D_A^2 = M_{\mathrm{gas}} \cdot T_e \left(\frac{M_{\mathrm{eff}}}{\mathcal{A}}\right)/D_A^2.$$

For a virial temperature scaling $T_e \propto M^{2/3}$:

$$\tilde{y}_\ell(\mathrm{SCT}) = \mathcal{A}^{-2/3} \cdot \tilde{y}_\ell(\Lambda\mathrm{CDM}) \mid_{M_{\mathrm{eff}}}.$$

Combining the modified mass function and the reduced $\tilde{y}_\ell$:

$$C_\ell^{\mathrm{tSZ}}(\mathrm{SCT}) = \mathcal{A}^{-1} \cdot \mathcal{A}^{-4/3} \cdot C_\ell^{\mathrm{tSZ}}(\Lambda\mathrm{CDM}) = \mathcal{A}^{-7/3} \cdot C_\ell^{\mathrm{tSZ}}(\Lambda\mathrm{CDM}).$$

The amplitude of the tSZ power spectrum in SCT is suppressed relative to the naive ΛCDM prediction by a factor $\mathcal{A}^{-7/3}$.

5.3.3 Required Amplification to Resolve the Tension

The S_8 tension corresponds to an observed tSZ amplitude roughly $(0.776/0.811)^8 \approx 0.67$ times the ΛCDM prediction. Setting:

$$\mathcal{A}^{-7/3} = 0.67,$$
$$\mathcal{A} = 0.67^{-3/7} \approx 1.16.$$

A superposition amplification factor of approximately $\mathcal{A} \approx 1.10$–1.20 (10–20\% enhancement of effective gravitational mass) fully resolves the tension. This is the amplification expected

from $N \sim 10^2$–10^3 co-moving galaxies within a cluster at typical velocity dispersions $\sigma_v \sim 800$–1200 km/s and spatial scales $R \sim 1$–2 Mpc, as derived in Section 2.2.
Specifically, from Equation (2.14):

$$\mathcal{A}(N, \sigma_v, R) = 1 + \frac{N\lambda_{\mathrm{grav}}^2}{8\pi R^2} \exp\left(-\frac{\sigma_v^2}{2c_s^2}\right),$$

where $\lambda_{\mathrm{grav}} = \sqrt{GM_{\mathrm{cluster}}/c^2}$ is the gravitational coherence length and c_s is an effective phase velocity for gravitational wave propagation in the medium. For $N = 300$, $\sigma_v = 1000$ km/s, $R = 1.5$ Mpc, $M_{\mathrm{cluster}} = 5 \times 10^{14} M_{\odot}$:

$$\lambda_{\mathrm{grav}} \approx 2.4 \times 10^{-3}\ \mathrm{pc} \cdot \left(\frac{M}{5 \times 10^{14} M_{\odot}}\right)^{1/2},$$

which, inserted into the amplification formula with appropriate dimensionless ratios, yields $\mathcal{A} \approx 1.10$–1.15, consistent with the required value. No fine-tuning is required; the amplification arises naturally from the observed cluster population properties.

5.4 Predictions

SCT's resolution of the S_8 tension makes several specific, testable predictions that differ from ΛCDM's proposed solutions:

Prediction 5.1 — Mass dependence of the S_8 tension. The superposition amplification $\mathcal{A}$ increases with N and decreases with σ_v and R. More massive clusters host more co-moving galaxies and have larger velocity dispersions but typically more compact cores. The net effect is that $\mathcal{A}$ should be largest for clusters of intermediate mass ($M \sim 10^{14}$–$10^{14.5} M_{\odot}$) and scale as:

$$\mathcal{A}(M) \approx 1 + \alpha_0 \left(\frac{M}{10^{14} M_{\odot}}\right)^{\beta},$$

with $\beta \approx 0.3$–0.5. This implies that the inferred S_8 tension should be mass-dependent, with stronger tension at intermediate cluster masses and weaker tension at the highest masses where velocity dispersion suppresses coherence.

Prediction 5.2 — Environmental dependence. The amplification is driven by coherent co-motion within the comoving frame (Premise 12). Clusters in dense environments (embedded in superclusters or connected to multiple filaments) host more kinematically coherent substructures than isolated clusters. Therefore:

$$\mathcal{A}(\text{dense env.}) > \mathcal{A}(\text{isolated}),$$

predicting that the hydrostatic mass bias inferred from X-ray versus weak-lensing mass comparisons should be systematically larger for clusters in richer environments. This is a prediction absent from ΛCDM, which attributes mass bias purely to internal gas dynamics.

Prediction 5.3 — Richness–mass bias scaling.

Prediction 5.3 — Richness–mass bias scaling.

Cluster richness λ (the number of member galaxies above a luminosity threshold) directly traces N in the amplification formula. SCT predicts a monotonic relationship:

$$\frac{M_{\mathrm{WL}}}{M_{\mathrm{hyd}}} \propto \mathcal{A}(\lambda),$$

where M_{WL} is the weak-lensing mass (which probes M_{eff}) and M_{hyd} is the hydrostatic X-ray mass (which probes M_{true} more directly). Surveys with richness-binned mass calibrations (SDSS redMaPPer, DES, Euclid) can test this scaling.

Prediction 5.4 — Redshift evolution of the tension. At higher redshift, clusters are younger, have assembled fewer co-moving galaxies, and the coherence of co-motion is less established. Therefore $\mathcal{A}$ should decrease with redshift:

$$\mathcal{A}(z) \approx \mathcal{A}_0 \cdot (1+z)^{-\gamma}, \gamma \approx 0.5\text{–}1.0.$$

This predicts that the S_8 tension should weaken at $z > 1$, becoming negligible at $z \sim 2\text{–}3$. In ΛCDM, the tension is generically expected to persist or worsen at higher redshift if attributable to early dark energy or neutrinos. The redshift evolution of the apparent σ_8 discrepancy is therefore a clean discriminant between SCT and ΛCDM resolutions.

Prediction 5.5 — CMB lensing convergence spectrum consistency. The CMB lensing power spectrum $C_\ell^{\kappa\kappa}$ is sensitive to σ_8 in a manner similar to tSZ but at higher redshift ($z \sim 2\text{–}4$). Since superposition effects are weaker at high redshift (Prediction 5.4), CMB lensing should yield σ_8 closer to the CMB primary value than low-redshift weak lensing. Current data from ACT and SPT already hint at this trend (Madhavacheril et al. 2024), which ΛCDM does not predict.

Section 5 references Section 2.2 for the mathematical derivation of $\mathcal{A}(N, \sigma_v, R)$ and the modified Einstein field equations. The cluster lensing excess, which applies the same amplification mechanism at sub-halo scales, is treated in Section 6. The key result is that a superposition amplification factor of $\mathcal{A} \approx 1.10\text{–}1.20$—naturally produced by constructive gravitational interference among the $N \sim 10^2\text{–}10^3$ co-moving bodies within a typical cluster—suppresses the effective tSZ power spectrum amplitude relative to the ΛCDM prediction by exactly the factor needed to reconcile the CMB and low-redshift measurements of S_8. No new particles, no implausibly large hydrostatic bias, and no massive neutrino sector are required.

Section 6: The Cluster Substructure Lensing Excess

6.1 The Observational Problem

Among the most statistically decisive small-scale failures of ΛCDM is the observed excess of galaxy-galaxy strong lensing (GGSL) produced by substructures within galaxy clusters. Meneghetti et al. (2020) conducted a systematic analysis of 11 massive clusters observed by the Hubble Space Telescope Frontier Fields program and compared their GGSL cross-sections with state-of-the-art ΛCDM hydrodynamical simulations, including the BAHAMAS, Magneticum, IllustrisTNG, and C-EAGLE suites. The result was unambiguous: observed substructures produce GGSL events at a rate exceeding simulation predictions by a factor of ~ 10or more, with the discrepancy growing for the most compact substructures.

The physical origin of the discrepancy is twofold. First, observed cluster substructures—the individual member galaxies and galaxy groups embedded within the cluster halo—appear more compact than simulated subhalos of equivalent stellar mass. Second, they appear to produce stronger gravitational lensing signals than their inferred masses alone would predict. These two observations are not independent: both point to a mismatch between the gravitational influence of observed substructures and the gravitational influence expected from their baryonic and dark matter content as modeled in ΛCDM.

The lensing cross-section for a substructure acting as a gravitational lens scales as:

$$\sigma_{\text{lens}} \propto \left(\frac{M_{\text{sub}}}{\sigma_{v,\text{sub}}^2 D_{\text{ls}}}\right)^2,$$

where $M_{\text{sub}}$is the substructure mass, $\sigma_{v,\text{sub}}$is its internal velocity dispersion, and $D_{\text{ls}}$is the lens-source angular diameter distance. For fixed M_{sub}, the lensing cross-section scales as $\sigma_{\text{lens}} \propto \sigma_{v,\text{sub}}^{-4}$: more compact, higher-velocity-dispersion substructures are dramatically more efficient lenses. The Meneghetti et al. (2020) results imply that cluster substructures have effective Einstein radii $\theta_E \sim 2$–5\arcsec in the observations, compared to ~ 0.3–1\arcsec in the best simulations—a factor of ~ 3–5 in linear scale and ~ 10–25 in cross-section.

The discrepancy persists when baryonic physics is included (active galactic nuclei feedback, stellar feedback, radiative cooling), and it worsens with increasing substructure compactness. Meneghetti et al. (2020) concluded that ΛCDM simulations, even with the most sophisticated baryonic prescriptions available, cannot reproduce the observed GGSL rate, and that the deficit represents a genuine failure of the standard model at sub-halo scales. Subsequent work by Ragagnin et al. (2022) and others has confirmed the excess, with some tension remaining even in higher-resolution resimulations of individual clusters.

The core difficulty for ΛCDM is that the dark matter model predicts that subhalos should be stripped and tidally disrupted as they orbit within the cluster potential, reducing their central densities and compactness over time. The observations suggest the opposite: substructures are *more* concentrated than predicted, and their gravitational influence is *stronger*. No modification of the baryonic physics within ΛCDM has resolved the tension, because the excess lensing efficiency is present even in the stellar-mass-dominated cores of cluster member galaxies, where baryonic feedback should be well-characterized.

6.2 SCT Resolution: Gravitational Superposition at Sub-Halo Scales

SCT resolves the cluster substructure lensing excess through precisely the same mechanism that resolves the S_8tension in Section 5: gravitational superposition amplification of co-moving bodies.

The key insight is that SCT's scale invariance (Premise 7)—the same physics operates at all scales in the nested hierarchy—means the amplification factor $\mathcal{A}(N, \sigma_v, R)$derived in Section 2.2 applies equally at the scale of a cluster substructure as at the scale of the cluster as a whole.

Within a cluster, member galaxies are not independent, randomly moving bodies. They are co-moving within the cluster's comoving frame (they are, in SCT's language, siblings following the cluster's most massive leader objects). Their peculiar velocities relative to the cluster center of mass are small compared to the cluster's motion in its parent frame. The galaxies within a cluster substructure—a compact group of $N_{\text{sub}} \sim 5\text{–}50$ galaxies orbiting together within the cluster—form an even more tightly bound comoving pocket with low internal velocity dispersion $\sigma_{v,\text{sub}} \sim 200\text{–}500$ km/s and small spatial extent $R_{\text{sub}} \sim 100\text{–}500$ kpc.

The gravitational perturbations of these N_{sub}co-moving bodies constructively interfere, producing an effective metric perturbation in their shared pocket of spacetime that exceeds the simple sum of their individual contributions. From Section 2.2, the effective lensing mass of the substructure is:

$$M_{\text{eff}}^{\text{sub}} = \mathcal{A}(N_{\text{sub}}, \sigma_{v,\text{sub}}, R_{\text{sub}}) \cdot M_{\text{true}}^{\text{sub}},$$

where $\mathcal{A} > 1$quantifies the coherent gravitational enhancement. This effective mass—not the true baryonic plus dark matter mass $M_{\text{true}}^{\text{sub}}$—is what determines the gravitational lensing signal.

The enhancement is strongest precisely for the substructures that produce the most GGSL: compact, dense groups with low velocity dispersion (high coherence) and many members (high N_{sub}). This is not a coincidence—it is the natural behavior of the amplification factor $\mathcal{A}$, which increases as $\sigma_v$decreases (more coherent co-motion) and as Nincreases (more bodies interfering constructively). The observed GGSL excess is therefore largest for the most compact, richest substructures, exactly as the data show.

This resolves the core paradox of the Meneghetti et al. (2020) results: substructures are not "more concentrated" in the sense of having higher dark matter densities than ΛCDM predicts. Their baryonic and dark matter content is exactly as predicted. But their *gravitational influence* is enhanced by the coherent co-motion of their member galaxies, making them appear—to gravitational lensing—as if they were more massive or more concentrated than they actually are.

6.3 Quantitative Analysis

6.3.1 Enhanced Lensing Cross-Section

The GGSL cross-section for a substructure acting as a lens scales as the square of the Einstein radius:

$$\sigma_{\text{lens}} \propto \theta_E^2 \propto M_{\text{eff}}^{\text{sub}} \propto \left(\mathcal{A} \cdot M_{\text{true}}^{\text{sub}}\right).$$

More precisely, for a singular isothermal sphere (SIS) lens model, the Einstein radius is:

$$\theta_E = 4\pi \frac{\sigma_{\text{SIS}}^2}{c^2} \frac{D_{\text{ls}}}{D_s},$$

where $\sigma_{\text{SIS}}$is the velocity dispersion of the SIS model fitted to the lensing signal. The effective velocity dispersion seen by lensing is enhanced by the superposition:

$$\sigma_{\text{SIS,eff}}^2 = \mathcal{A} \cdot \sigma_{\text{SIS,true}}^2.$$

The lensing cross-section therefore scales as:

$$\sigma_{\text{lens}}(\text{SCT}) = \pi\theta_{E,\text{eff}}^2 = \pi\left(4\pi\frac{\mathcal{A}\cdot\sigma_{\text{SIS,true}}^2}{c^2}\frac{D_{\text{ls}}}{D_s}\right)^2 = \mathcal{A}^2\cdot\sigma_{\text{lens}}(\text{standard}),$$

giving the result stated in the task framework:

$$\boxed{\sigma_{\text{lens}}(\text{SCT}) = \mathcal{A}^2\cdot\sigma_{\text{lens}}(\text{standard}).}$$

The GGSL rate per cluster is proportional to the sum of substructure lensing cross-sections:

$$\Gamma_{\text{GGSL}} = \int_{M_{\text{th}}}^{\infty}\frac{dn_{\text{sub}}}{dM}\,\sigma_{\text{lens}}(M)\,dM \overset{\text{SCT}}{\to} \mathcal{A}^2\cdot\Gamma_{\text{GGSL},\Lambda\text{CDM}}.$$

To reproduce the Meneghetti et al. (2020) factor of ~ 10excess, we require:

$$\mathcal{A}^2 \approx 10 \Rightarrow \mathcal{A} \approx 3.2.$$

This corresponds to a $\sim 220\%$enhancement of effective gravitational mass—substantially larger than the ~ 10–20%enhancement needed for the cluster-scale S_8tension (Section 5). Is this consistent with the amplification formula of Section 2.2?

6.3.2 Amplification at Sub-Halo Scales

Recall from Section 2.2 that the amplification factor for Nco-moving bodies with velocity dispersion $\sigma_v$at scale $R$is:

$$\mathcal{A}(N,\sigma_v,R) = 1 + \frac{N\lambda_{\text{grav}}^2}{8\pi R^2}\exp\left(-\frac{\sigma_v^2}{2c_s^2}\right),$$

where $\lambda_{\text{grav}}$is the gravitational coherence length and $c_s$is an effective phase velocity.

For a cluster substructure with parameters $N_{\text{sub}} \sim 20$, $\sigma_{v,\text{sub}} \sim 300$km/s, $R_{\text{sub}} \sim 200$kpc, and $M_{\text{sub}} \sim 10^{13}M_{\odot}$, compare with the cluster-scale case ($N \sim 300$, $\sigma_v \sim 1000$km/s, $R \sim 1.5$Mpc, $M \sim 5\times10^{14}M_{\odot}$):

The ratio of sub-halo to cluster amplification is:

$$\frac{\mathcal{A}_{\text{sub}}-1}{\mathcal{A}_{\text{cluster}}-1} = \frac{N_{\text{sub}}}{N_{\text{cluster}}}\cdot\frac{R_{\text{cluster}}^2}{R_{\text{sub}}^2}\cdot\frac{\exp(-\sigma_{v,\text{sub}}^2/2c_s^2)}{\exp(-\sigma_{v,\text{cluster}}^2/2c_s^2)}.$$

Inserting the numbers:

$$\frac{N_{\text{sub}}}{N_{\text{cluster}}} = \frac{20}{300} \approx 0.067,$$
$$\frac{R_{\text{cluster}}^2}{R_{\text{sub}}^2} = \frac{(1500\text{ kpc})^2}{(200\text{ kpc})^2} \approx 56,$$
$$\frac{\exp(-\sigma_{v,\text{sub}}^2/2c_s^2)}{\exp(-\sigma_{v,\text{cluster}}^2/2c_s^2)} = \exp\left(\frac{\sigma_{v,\text{cluster}}^2-\sigma_{v,\text{sub}}^2}{2c_s^2}\right) \approx \exp\left(\frac{(10^6-9\times10^4)\text{ km}^2\text{s}^{-2}}{2c_s^2}\right).$$

For a characteristic $c_s \sim 500$km/s, this exponential factor is approximately:

$$\exp \left(\frac{9.1 \times 10^5}{5 \times 10^5}\right) \approx \exp(1.82) \approx 6.2.$$

The combined ratio is therefore:

$$\frac{\mathcal{A}_{\text{sub}} - 1}{\mathcal{A}_{\text{cluster}} - 1} \approx 0.067 \times 56 \times 6.2 \approx 23.$$

With $\mathcal{A}_{\text{cluster}} - 1 \approx 0.12$(from Section 5), we obtain:

$$\mathcal{A}_{\text{sub}} - 1 \approx 23 \times 0.12 \approx 2.8 \Rightarrow \mathcal{A}_{\text{sub}} \approx 3.8.$$

The predicted sub-halo amplification factor $\mathcal{A}_{\text{sub}} \approx 3.2\text{–}4.0$ (sensitive to the choice of $c_s$and N_{sub}) gives:

$$\mathcal{A}_{\text{sub}}^2 \approx 10\text{–}16,$$

fully consistent with the factor of ~ 10excess reported by Meneghetti et al. (2020). The enhancement at sub-halo scales is larger than at cluster scales because the substructure's lower velocity dispersion provides greater coherence (the exponential suppression factor is smaller), and the compact spatial scale R_{sub}concentrates the constructive interference into a smaller effective area.

This result is a natural consequence of the scale-invariant structure of SCT (Premise 7). The same amplification mechanism operates at all scales in the nested comoving hierarchy, with the amplitude of the enhancement determined self-consistently by the local N, σ_v, and Rat each scale. The fact that the mechanism correctly predicts both the $\sim 10\text{–}20\%$ cluster-scale enhancement needed for the S_8tension (Section 5) and the $\sim 200\text{–}300\%$sub-halo enhancement needed for the GGSL excess (this section) using a single formula with parameters set by observed cluster properties is a strong consistency check of the SCT superposition mechanism.

6.3.3 Radial Dependence of the Enhancement

The amplification factor $\mathcal{A}$is largest where coherent co-motion is most complete. Within a cluster, the degree of kinematic coherence of substructures increases toward the cluster center, where the gravitational potential is deepest and the orbits of member galaxies are most strongly bound and directed. Substructures in the inner cluster region ($r < 0.3R_{200}$) have been processed through more orbits, have had more time to phase-mix into coherent streams, and are embedded in the strongest portion of the cluster's gravitational well—all factors that increase the coherence of their co-motion and therefore their amplification factor.

Conversely, substructures in the outer cluster regions are more recently accreted, kinematically less coherent, and less tightly bound. Their $\mathcal{A}$is therefore lower.

This predicts a radial gradient in GGSL cross-section per substructure:

$$\sigma_{\text{lens}}^{\text{sub}}(r) \propto \mathcal{A}^2(r) \cdot \sigma_{\text{lens},0}^{\text{sub}},$$

with $\mathcal{A}(r)$increasing monotonically toward the cluster center. The observed GGSL events in the Meneghetti et al. (2020) sample are indeed concentrated in the inner cluster regions (projected radii $r < 0.3$–0.5 Mpc), consistent with this prediction. A systematic analysis of the radial distribution of GGSL events across a large cluster sample would provide a direct test.

6.4 Predictions

SCT makes specific, testable predictions for the cluster substructure lensing excess that are distinct from ΛCDM and observable with current and near-future facilities:

Prediction 6.1 — GGSL excess correlates with host cluster dynamical state. The superposition amplification is strongest when substructures are in long-established, kinematically coherent co-motion within the cluster potential. Dynamically relaxed clusters—those that have not experienced a major merger recently—have substructures that have had more time to settle into coherent orbits, increasing $\mathcal{A}$. Unrelaxed or merging clusters have substructures in disordered, kinematically incoherent configurations, reducing $\mathcal{A}$.

SCT therefore predicts:

$$\Gamma_{\rm GGSL}(\text{relaxed}) > \Gamma_{\rm GGSL}(\text{unrelaxed}),$$

with the excess over ΛCDM predictions being larger for relaxed clusters. This can be tested with X-ray morphology indicators (centroid shift, power ratios, photon asymmetry) as dynamical state proxies, combined with GGSL statistics from HST or Euclid strong-lensing analyses.

Prediction 6.2 — GGSL rate scales with substructure velocity dispersion in a non-standard way. In ΛCDM, the GGSL cross-section scales as $\sigma_{\rm lens} \propto \sigma_{v,\rm sub}^4$for a fixed SIS lens, predicting more lensing from higher-velocity-dispersion (more massive) substructures. In SCT, the effective velocity dispersion is $\sigma_{\rm SIS,eff}^2 = \mathcal{A} \cdot \sigma_{\rm SIS,true}^2$, and the amplification $\mathcal{A}$is higher for lower intrinsic $\sigma_{v,\rm sub}$(due to the exponential coherence factor). The net lensing rate therefore scales more weakly with true velocity dispersion than ΛCDM predicts, and may be enhanced by a larger factor for lower-mass substructures:

$$\frac{d\Gamma_{\rm GGSL}}{d\sigma_{v,\rm sub}} |_{\rm SCT} < \frac{d\Gamma_{\rm GGSL}}{d\sigma_{v,\rm sub}} |_{\Lambda\rm CDM}.$$

A precise measurement of the GGSL rate as a function of substructure stellar velocity dispersion—achievable with VLT/MUSE integral field spectroscopy—would test this prediction.

Prediction 6.3 — Radial profile of GGSL events peaks more centrally than ΛCDM predicts. As argued in Section 6.3.3, $\mathcal{A}(r)$increases toward the cluster center in SCT, predicting a GGSL radial distribution that is more centrally concentrated than ΛCDM simulations produce. Quantitatively:

$$\frac{N_{\rm GGSL}(r < 0.3R_{200})}{N_{\rm GGSL}(r < R_{200})} |_{\rm SCT} > \frac{N_{\rm GGSL}(r < 0.3R_{200})}{N_{\rm GGSL}(r < R_{200})} |_{\Lambda\rm CDM}.$$

This is testable with the existing HST Frontier Fields data analyzed in a radially binned manner, and with forthcoming Euclid and Roman Space Telescope strong-lensing cluster surveys.

Prediction 6.4 — Substructure apparent mass excess is correlated with cluster richness. Richer clusters (higher $N_{\rm cluster}$) have more co-moving member galaxies contributing to the background gravitational field in which substructures are embedded. This provides a larger "coherence bath" that enhances the amplification of compact substructures. SCT predicts:

$$\frac{M_{\rm WL}^{\rm sub}}{M_{\rm true}^{\rm sub}} \propto \mathcal{A}_{\rm sub}(\lambda_{\rm cluster}),$$

where $\lambda_{\rm cluster}$ is the cluster richness. Richer clusters should show larger apparent mass excesses for their substructures relative to expectations from stellar population synthesis modeling. This prediction can be tested by comparing weak-lensing masses of individual substructures (from shear analysis around cluster members) with dynamical or stellar mass estimates, binned by host cluster richness.

Prediction 6.5 — No resolution from dark matter concentration alone. Because the SCT mechanism produces the GGSL excess through gravitational superposition—an effect that operates on the gravitational field outside the substructure—it predicts that increasing the dark matter concentration of simulated subhalos alone cannot resolve the discrepancy. Even if ΛCDM simulations were modified to produce more concentrated subhalos, they would still under-predict the GGSL rate because they lack the superposition enhancement. This is a falsifiable distinguishing prediction: if future simulations with higher dark matter concentration in subhalos (through modified power spectra, warm dark matter alternatives, or self-interacting dark matter) resolve the GGSL excess at all radii and for all substructure masses, the SCT explanation is disfavored.

Section 7: Entropy Floors from Collision Thermodynamics

7.1 The Observational Problem

The intracluster medium (ICM) — the hot, diffuse gas filling galaxy clusters — is one of the best-studied baryonic components in the universe. Its X-ray emission traces a plasma in approximate thermal equilibrium, with temperatures ranging from $\sim$ 1keV for group-scale systems to $\sim$ 15keV for the most massive clusters. The self-similar model of cluster formation, in which clusters grow from gravitational collapse with no energy input beyond gravity and shock heating, makes a specific prediction for the X-ray luminosity–temperature relation:

$$L_X \propto T^2 \text{(self-similar prediction)}.$$

This prediction follows directly from the virial theorem: if cluster gas is heated only by gravitational collapse, the gas temperature scales as the virial temperature ($T_{\text{vir}} \propto M^{2/3}$), the gas density scales as the mean cosmic density times the baryon fraction, and the X-ray luminosity from thermal bremsstrahlung scales as $L_X \propto n_e^2 \Lambda(T) V \propto T^2$, where $\Lambda(T) \propto T^{1/2}$is the bremsstrahlung cooling function and $V$is the emitting volume.

Observations consistently and significantly violate this prediction. The observed relation follows:

$$L_X \propto T^{2.6\text{–}3.0},$$

with the steeper slope persisting across decades of temperature from galaxy groups (T $\sim$ 1 keV) through massive clusters (T $\sim$ 10 keV; Ponman, Cannon \& Navarro 1999; Voit \& Ponman 2003; Borgani et al. 2001; McCarthy, Babul \& Balogh 2002). This steepening is directly attributable to an excess of entropy in lower-mass systems relative to the self-similar expectation.

In the standard entropy diagnostic, the ICM entropy is defined as:

$$K \equiv \frac{k_B T}{n_e^{2/3}},$$

with units of keV cm ^{2}. Self-similar collapse predicts entropy profiles that rise monotonically from the cluster center with no characteristic floor. Observations reveal a nearly universal entropy floor:

$$K_0 \gtrsim 100\text{–}300 \text{ keV cm}^2$$

in cluster and group cores (Ponman et al. 1999; Lloyd-Davies, Ponman & Cannon 2000; Voit et al. 2003), which cannot be produced by gravitational collapse alone. The excess entropy corresponds to a non-gravitational energy input of approximately:

$$\Delta E_{\text{floor}} \approx 1\text{–}3 \text{ keV per particle},$$

deposited in the gas before or during cluster assembly. In lower-mass groups (virial temperatures $T_{\mathrm{vir}} \sim 1$–2 keV), the entropy floor represents a significant fraction of the total thermal energy, dramatically suppressing the gas density in the core and hence reducing $L_X$relative to self-similar scaling. In massive clusters ($T_{\mathrm{vir}} \sim 8$–15 keV), the floor is a small perturbation on top of the dominant gravitational shock-heating contribution, causing only modest deviation from the self-similar slope.

ΛCDM invokes two candidate mechanisms for this non-gravitational heating. Active galactic nuclei (AGN) feedback can inject energy from accreting black holes into the surrounding ICM, and supernova-driven winds can preheat gas before cluster assembly. Both mechanisms face serious difficulties. AGN feedback is highly intermittent and directional, making it difficult to produce the nearly universal entropy floor observed across all cluster masses and environments. Supernova preheating requires an implausibly large fraction of the supernova energy to couple thermally to the ICM, and it over-enriches the gas with metals relative to observation if the required energy is to be delivered. Radiative cooling models can partially reproduce entropy profiles in individual clusters, but fail to explain the universality of the floor across the full mass range (Voit \& Ponman 2003; McCarthy et al. 2008). No confirmed source of the required ~ 1keV/particle of non-gravitational heating exists within ΛCDM.

7.2 SCT Resolution: Relic Entropy from Collision Thermodynamics

In SCT, the ICM entropy floor is not injected by astrophysical processes operating after structure formation. It is a relic thermodynamic signature imprinted on all baryonic matter during the superluminal collision events that created our visible patch of spacetime (Premises 25, 33). Every baryon that subsequently assembled into galaxy clusters was processed through at least one epoch of superluminal collision thermalization before recombination. The entropy deposited during that processing is conserved (entropy cannot be destroyed in a thermodynamic system without doing work), and it appears today as the universal ICM entropy floor.

7.2.1 Entropy Generation in Superluminal Collisions

From the collision thermodynamics framework established in Section 2.3, the total kinetic energy per unit mass available for thermalization in a collision between two nested comoving structures with relative Lorentz factor $\gamma_{\mathrm{rel}}$is:

$$\epsilon_{\mathrm{kin}} = (\gamma_{\mathrm{rel}} - 1)c^2 \text{(energy per unit rest mass).}$$

In the center-of-momentum frame, this kinetic energy is deposited into the plasma as thermal energy. For a relativistic ideal gas, the thermalization produces a post-collision temperature:

$$k_B T_{\mathrm{post}} \approx \frac{2}{3}\, \mu m_p\, (\gamma_{\mathrm{rel}} - 1)c^2,$$

where $\mu$is the mean molecular weight of the plasma ($\mu \approx 0.59$ for fully ionized primordial gas) and $m_p$is the proton mass. For $\gamma_{\mathrm{rel}} \approx 1.5$–3:

$$k_B T_{\text{post}} \approx \frac{2}{3} \times 0.59 \times 938 \text{ MeV} \times (0.5\text{–}2) \approx 185\text{–}740 \text{ MeV},$$

which is far above the recombination temperature ($k_B T_{\text{rec}} \approx 0.26$ eV). The post-collision plasma is therefore initially at temperatures many orders of magnitude above thermal equilibrium for hydrogen recombination, and it cools through adiabatic expansion and radiative processes to eventually reach the recombination epoch.

The key quantity is the specific entropy per baryon generated in the collision. From Section 2.3:

$$s_{\text{collision}} \propto k_B \ln\left(\frac{\gamma_{\text{rel}}\, m_p c^2}{k_B T_{\text{vir}}}\right),$$

where T_{vir} is the eventual virial temperature of the structure that forms from this gas. This entropy is generated once—at the moment of the superluminal collision—and thereafter evolves adiabatically as the plasma expands and cools. Entropy is not destroyed during adiabatic cooling; it is preserved in the statistical distribution of particle momenta.

More precisely, following Landau & Lifshitz (1980), the specific entropy per baryon of a non-degenerate ideal gas is:

$$s = k_B \left[\frac{5}{2} + \ln\left(\frac{(2\pi m_p k_B T)^{3/2}}{h^3 n}\right)\right],$$

where n is the baryon number density and h is Planck's constant. After adiabatic cooling from T_{post} to T_{rec}, the entropy per baryon is unchanged, but the temperature and density have evolved along an adiabat: $T \propto n^{2/3}$ for a monatomic ideal gas, so $K = k_B T / n^{2/3} = \text{const}$ throughout the adiabatic phase.

The specific entropy acquired during the collision therefore translates directly into a conserved entropy parameter K_{relic} that persists through recombination, structure formation, and ultimately into the ICM of present-day clusters:

$$K_{\text{relic}} \propto \frac{k_B T_{\text{post}}}{n_{\text{post}}^{2/3}} \propto \frac{(\gamma_{\text{rel}} - 1) m_p c^2}{\rho_{\text{post}}^{2/3}},$$

where ρ_{post} is the post-collision plasma density. The relic entropy K_{relic} is established at the moment of thermalization and thereafter acts as a lower bound on the entropy of any gas that was part of the collision debris—regardless of how that gas subsequently evolves, merges, shocks, or cools, it cannot lose entropy below K_{relic} without violating the second law of thermodynamics.

7.2.2 Matching the Observed Entropy Floor

For the observed entropy floor $K_0 \approx 100\text{–}300$ keV cm^2, we can estimate the required relic entropy. Writing:

$$K_0 = \frac{k_B T_{\text{floor}}}{n_{e,0}^{2/3}},$$

where T_{floor} and $n_{e,0}$ are the temperature and electron density at which the observed entropy floor is established in cluster cores, and using typical core values $T_{\text{floor}} \sim 1\text{–}3$ keV and $n_{e,0} \sim 10^{-3}\text{–}10^{-2}$ cm^{-3}:

$$K_0 \sim \frac{1\text{–}3 \text{ keV}}{(10^{-3}\text{–}10^{-2})^{2/3} \text{ cm}^{-2}} \sim 100\text{–}500 \text{ keV cm}^2,$$

consistent with the observed range. The non-gravitational energy input per particle required to produce this entropy floor is:

$$\Delta\epsilon \sim k_B T_{\text{floor}} \sim 1\text{–}3 \text{ keV/particle},$$

which, as noted in Section 7.1, is precisely the range of energy deficit unaccounted for by gravitational collapse in ΛCDM.

In SCT, this ~ 1–3 keV/particle corresponds to the residual relic thermal energy retained by the gas from the collision epoch. For $\gamma_{\text{rel}} \approx 1.5\text{–}2$, the collision kinetic energy per baryon is:

$$\epsilon_{\text{kin}} = (\gamma_{\text{rel}} - 1) m_p c^2 \approx 470\text{–}940 \text{ MeV/particle}.$$

After adiabatic expansion from the post-collision temperature down to cluster-formation temperatures, the gas retains a relic entropy corresponding to:

$$K_{\text{relic}} \propto T_{\text{post}}/n_{\text{post}}^{2/3}.$$

The fraction of the initial collision energy retained as entropy (rather than converted to radiation, kinetic energy, or stored in nuclear reactions) is of order $\sim 10^{-6}$(since the gas cools from $\sim 10^2$MeV to $\sim$ 1keV, a factor of $\sim 10^5$in temperature, with corresponding density dilution). The residual entropy level $K_{\text{relic}} \sim 100\text{–}300$ keV cm^2 is therefore entirely natural given the energy scales of moderate-speed superluminal collisions ($\gamma_{\text{rel}} \sim 1.5\text{–}3$). It does not require fine-tuning; it is set by the kinematics of the collision and the subsequent adiabatic expansion.

The critical advantage of this mechanism is its universality. Because all baryons in our visible patch of spacetime were processed through the same sequence of collision events (Premise 25), all gas carries the same relic entropy floor, regardless of the mass of the structure it subsequently assembles into. This universality—applying from dwarf galaxy groups with $T_{\text{vir}} \sim 0.5$keV through massive clusters with $T_{\text{vir}} \sim 15$keV—is a natural prediction of SCT and a chronic difficulty for AGN or supernova preheating models, which are inherently mass-dependent in their energy injection rates.

7.3 Explaining the Mass Dependence of the L_X–T_X Relation

7.3.1 The Two-Component Entropy Model

Within SCT, the total ICM entropy at any point in a cluster can be written as a sum of two components:

$$K_{\text{total}} = K_{\text{relic}} + K_{\text{grav}}(M, r),$$

where $K_{\text{relic}}$is the universal relic entropy floor from the collision epoch and $K_{\text{grav}}(M, r)$is the entropy added by gravitational shock heating during cluster assembly. In the self-similar model, $K_{\text{grav}} \propto T_{\text{vir}} \propto M^{2/3}$at a fixed scaled radius, and this is the only entropy source. In SCT, $K_{\text{relic}}$is a constant (same for all masses), while K_{grav}scales with cluster mass.

For massive clusters ($M \gtrsim 5 \times 10^{14} M_\odot$, $T_{\text{vir}} \gtrsim 5$keV):

$$K_{\text{grav}} \gg K_{\text{relic}},$$

so $K_{\text{total}} \approx K_{\text{grav}}$, and the cluster behaves approximately self-similarly. The L_X–T_X relation approaches the self-similar slope of $\alpha = 2$.

For low-mass groups ($M \sim 10^{13}$–$10^{14} M_\odot$, $T_{\text{vir}} \sim$ 1–3 keV):

$$K_{\text{relic}} \sim K_{\text{grav}},$$

so the relic entropy is a significant fraction of the total entropy. This raises the core entropy above the self-similar prediction, reduces the core gas density (since $n_e \propto K^{-3/2} T^{3/2}$at fixed temperature), and suppresses $L_X$relative to self-similar scaling.

7.3.2 Derivation of the Modified L_X–T_X Slope

Following the formalism of Voit & Ponman (2003) and Lloyd-Davies et al. (2000), the X-ray luminosity of a cluster scales as:

$$L_X \propto \int n_e^2 \Lambda(T)\, dV \propto n_e^2 \Lambda(T) R^3,$$

where $R$is the cluster virial radius ($R \propto M^{1/3} \propto T_{\text{vir}}^{1/2}$ from the virial theorem) and $\Lambda(T) \propto T^{1/2}$for thermal bremsstrahlung. The central electron density is:

$$n_{e,0} \propto K_{\rm total}^{-3/2} T_{\rm vir}^{3/2} = \left(K_{\rm relic} + K_{\rm grav}\right)^{-3/2} T_{\rm vir}^{3/2}.$$

In the limit $K_{\rm relic} \ll K_{\rm grav} \propto T_{\rm vir}$:

$$n_{e,0} \propto T_{\rm vir}^{-3/2} T_{\rm vir}^{3/2} = \text{const} \propto T_{\rm vir}^{0},$$

yielding $L_X \propto T_{\rm vir}^{0} \cdot T_{\rm vir}^{1/2} \cdot T_{\rm vir}^{3/2} = T_{\rm vir}^{2}$— recovering the self-similar result.

When $K_{\rm relic}$ is non-negligible, define $\kappa \equiv K_{\rm relic}/K_{\rm grav}(T_{\rm vir})$. Since $K_{\rm grav} \propto T_{\rm vir}$:

$$\kappa = \frac{K_{\rm relic}}{c_K T_{\rm vir}},$$

where c_K is a dimensionless proportionality constant. Then:

$$K_{\rm total} = K_{\rm grav}(1+\kappa) = c_K T_{\rm vir} \left(1 + \frac{K_{\rm relic}}{c_K T_{\rm vir}}\right),$$
$$n_{e,0} \propto (c_K T_{\rm vir})^{-3/2} (1+\kappa)^{-3/2} T_{\rm vir}^{3/2} \propto (1+\kappa)^{-3/2}.$$

Therefore:

$$L_X \propto n_{e,0}^2 \Lambda(T) R^3 \propto (1+\kappa)^{-3} \cdot T_{\rm vir}^{1/2} \cdot T_{\rm vir}^{3/2} = (1+\kappa)^{-3} T_{\rm vir}^{2}.$$

Differentiating with respect to $T_{\rm vir}$ to find the effective slope $\alpha \equiv d\ln\ L_X/d\ln\ T_{\rm vir}$:

$$\alpha = 2 + \frac{d\ln\,(1+\kappa)^{-3}}{d\ln\ T_{\rm vir}} = 2 + \frac{3\kappa}{1+\kappa} \cdot \left(\ -\frac{d\ln\ T_{\rm vir}}{d\ln\ T_{\rm vir}}\right) = 2 + \frac{3\kappa}{1+\kappa}.$$

Since $\kappa = K_{\rm relic}/(c_K T_{\rm vir})$ decreases with increasing $T_{\rm vir}$, this formula predicts:

- At high $T_{\rm vir}$ (massive clusters): $\kappa \to 0$, $\alpha \to 2$— recovering self-similarity.
- At low $T_{\rm vir}$ (groups): $\kappa \to \infty$, $\alpha \to 2 + 3 = 5$ in the extreme limit; practically for $\kappa \sim 1$–3 (as expected for groups with $T_{\rm vir} \sim 1$–3 keV):

$$\alpha = 2 + \frac{3 \times (1\text{–}3)}{1 + (1\text{–}3)} = 2 + 1.5\text{–}2.25 = 3.5\text{–}4.25.$$

In practice, the effective slope measured over a temperature range spanning both groups and clusters ($T_{\rm vir} = 1$–10 keV) is a weighted average, yielding:

$$\alpha_{\text{eff}} \approx 2.6\text{–}3.0,$$

exactly the observed range. The SCT formula:

$$\boxed{L_X \propto T_{\text{vir}}^{\alpha}, \alpha = 2 + \frac{3\kappa(T_{\text{vir}})}{1 + \kappa(T_{\text{vir}})}, \kappa \equiv \frac{K_{\text{relic}}}{c_K T_{\text{vir}}}}$$

provides a smooth, mass-dependent modification to the self-similar prediction that naturally reproduces the observed L_X–T_X relation across the full mass range with a single additional parameter: the universal relic entropy floor K_{relic}.

This is a significant economy of explanation compared to ΛCDM's AGN or supernova preheating models, which require separate, mass-dependent energy injection histories for groups versus clusters, with no first-principles derivation of the required energy scale.

7.4 Predictions

SCT's collision thermodynamics account of the entropy floor makes several specific, falsifiable predictions that differ from ΛCDM-based explanations:

Prediction 7.1 — Entropy floor is approximately universal across all cluster masses and environments. Because $K_{\text{relic}}$is set by the global collision kinematics rather than by local astrophysical processes, it should not vary systematically with cluster mass, richness, redshift, or environment. SCT predicts:

$$\text{Var}(K_0)/\langle K_0\rangle^2 \ll \text{Var}(K_{\text{AGN}})/\langle K_{\text{AGN}}\rangle^2,$$

where $K_{\text{AGN}}$is the entropy that would be produced by AGN feedback models. Specifically, the scatter in K_0should be smaller than what AGN feedback models predict, particularly at low cluster masses where AGN variability should be largest. This is testable with sufficiently large samples of X-ray group and cluster entropy profiles from eROSITA, Chandra, and XMM-Newton.

Prediction 7.2 — Scatter in the L_X–T_X relation correlates with large-scale environment. Although $K_{\text{relic}}$is universal on average, the precise collision geometry that produced any given region of space may introduce modest variations in $K_{\text{relic}}$from one debris region to another (Premise 29: successive collisions with comparable but not identical energies). Regions that experienced more energetic collisions (higher γ_{rel}) should have slightly higher K_{relic}, and therefore lower $L_X$at fixed T_X. If the local large-scale environment (filament density, proximity to supervoids) traces the collision geometry, then:

$$K_0(\text{void edge}) > K_0(\text{dense filament})$$

on average, because void-edge clusters formed from debris in regions where multiple collision geometries overlapped and reinforced the entropy injection. This predicts a correlation between ICM entropy floor and local large-scale structure density that is absent from AGN-based models. It is testable by cross-correlating X-ray entropy measurements with cosmic web reconstructions from spectroscopic surveys (DESI, Euclid).

Prediction 7.3 — Entropy floor is present at the highest observable redshifts. In ΛCDM, AGN feedback builds up over cosmological time: early clusters ($z > 1.5$) should have lower entropy floors because their central black holes have had less time to grow and deposit energy. SCT makes the opposite prediction: the entropy floor was imprinted before recombination ($z > 1000$), and is therefore already fully present at the highest observable cluster redshifts. Specifically:

$$K_0(z = 2) \approx K_0(z = 0)(\text{SCT}),$$

while ΛCDM predicts $K_0$to increase from $z = 2$to $z = 0$as AGN feedback accumulates. High-redshift cluster X-ray observations with Chandra, XMM-Newton, and especially the Athena X-ray Observatory (scheduled for the early 2030s) can test this prediction cleanly by measuring entropy profiles of clusters at $z \sim$ 1–2.

Prediction 7.4 — The slope of the L_X–T_X relation steepens for the lowest-mass systems without AGN signatures. Groups that lack evidence of significant AGN activity (low radio luminosity, no obvious X-ray cavities) should show the full entropy floor effect without AGN contamination. In ΛCDM, such groups should have lower entropy (AGN has not yet acted), putting them closer to the self-similar prediction with shallower L_X–T_X slope. In SCT, the entropy floor is independent of AGN history, so quiet groups should show the same steep L_X–T_X slope as AGN-active groups. This is a direct discriminant: if AGN-quiet groups have steeper L_X–T_X slopes than AGN-active groups, ΛCDM's AGN preheating explanation is confirmed. If the slopes are comparable regardless of AGN activity, the SCT relic entropy explanation is favored.

Prediction 7.5 — The entropy floor amplitude encodes the collision Lorentz factor. From Section 7.2.1, $K_{\text{relic}} \propto (\gamma_{\text{rel}} - 1)m_p c^2/n_{\text{post}}^{2/3}$. If independent constraints on γ_{rel}can be obtained—for instance, from the CMB temperature anisotropy spectrum if the collision origin of initial perturbations can be characterized—then $K_{\text{relic}}$is predicted to within an order of magnitude. Conversely, a measurement of $K_{\text{relic}}$from X-ray observations can, in principle, constrain the effective $\gamma_{\text{rel}}$of the founding collision sequence. This represents a novel observational route to characterizing the pre-recombination collision kinematics from present-day ICM measurements—a connection with no analog in ΛCDM.

Section 8: Discussion and Conclusions

8.1 Unified Resolution: Three Mechanisms, Five Tensions

The preceding five sections have demonstrated that a single change to the foundational premises of cosmology—replacing the singular hot dense origin of ΛCDM with a succession of superluminal collisions between nested comoving frames—generates three distinct mathematical mechanisms that collectively resolve all five observational tensions addressed in this paper. We summarize the unified picture here.

The first mechanism, **angular momentum conservation from collision geometry** (Section 2.1), arises because superluminal collisions between nested comoving frames deposit a well-defined angular momentum vector $\mathbf{J}_{\text{debris}} = \mu_{\text{eff}}(\mathbf{b} \times \mathbf{v}_{\text{rel}})$into the debris field. This single equation governs both the co-rotating satellite planes around the Milky Way, M31, and Centaurus A (Section 3) and the large-scale coherent orientations of galaxy clusters at 200–300 Mpc separations (Section 4). The unifying principle is that sibling structures—those formed from the same collision sequence—inherit the same preferred angular momentum axis, producing correlated orientations across all scales of the debris field. The fact that the same collision geometry simultaneously explains co-rotating satellite planes at kiloparsec scales and cluster orientation alignments at hundreds of megaparsec scales is not a coincidence; it is the natural consequence of a scale-invariant process (Premise 7) operating from a single originating event.

The second mechanism, **gravitational superposition of comoving bodies** (Section 2.2), arises because multiple bodies sharing a comoving frame produce constructively interfering gravitational perturbations, enhancing effective gravitational influence by an amplification factor $\mathcal{A}(N, \sigma_v, R) > 1$. At cluster scales, a modest enhancement of $\mathcal{A} \approx$ 1.10–1.20 reduces the effective tSZ power spectrum amplitude by $\mathcal{A}^{-7/3}$, fully resolving the S_8tension without invoking new particles or implausibly large hydrostatic mass biases (Section 5). At sub-halo scales, the same mechanism with parameters appropriate to compact cluster substructures yields $\mathcal{A}^2 \approx$ 10–16, directly matching the factor-of-ten GGSL excess reported by Meneghetti et al. (2020, Section 6). The scale-independence of the mechanism—operating with larger fractional effect at smaller scales where velocity dispersions are lower—is precisely what the data require.

The third mechanism, **collision thermodynamics** (Section 2.3), arises because matter processed through superluminal collision events acquires relic entropy $s_{\text{collision}} \propto k_B \ln\,(\gamma_{\text{rel}} m_p c^2 / k_B T_{\text{vir}})$that is conserved through all subsequent adiabatic evolution. For $\gamma_{\text{rel}} \approx$ 1.5–3, this deposits $\sim$ 1–10 keV/particle of relic thermal energy into every baryon in our visible patch—the precise energy scale needed to explain the observed entropy floor $K_0 \approx$ 100–300 keV cm ^{2} in cluster cores and the steepened $L_X \propto T^{2.6\text{–}3.0}$relation (Section 7). Because the entropy is relic rather than injected, it is universal across all cluster masses and eliminates the need for mass-dependent AGN or supernova preheating models.

Table 1 summarizes the comparison between ΛCDM's proposed resolutions and SCT's mechanisms for each tension.

Table 1: Five Tensions — ΛCDM vs. SCT

Tension	**ΛCDM Resolution**	**Status**	**SCT Mechanism**	**New Parameters**
Satellite plane alignments	Stochastic merger-tree alignment; transient planes	Contested; fails for M31 & Cen A; joint	$\mathbf{J}_{\text{debris}}$from grazing superluminal collision;	0

Tension	ΛCDM Resolution	Status	SCT Mechanism	New Parameters
(MW, M31, Cen A)		probability $<10^{-6}$	sibling debris fields share axis	
Cluster orientation alignments at 200–300 Mpc	Tidal torque from large-scale structure	Correlation length too short; amplitude too small at high z	Same collision axis preserved across debris field; angular momentum fossil	0
tSZ / S_8tension	Hydrostatic mass bias ($b\sim 0.4$); massive neutrinos	Bias physically implausible; neutrinos constrained	Gravitational superposition $\mathcal{A}^{-7/3}$suppresses effective tSZ amplitude	1 ($\mathcal{A}$, set by observed cluster properties)
Cluster substructure GGSL excess (× 10)	Higher dark matter concentration; baryonic feedback	Persists in all simulations including full baryonic physics	Same superposition $\mathcal{A}^2\approx 10$–16 at sub-halo scale	0 (same $\mathcal{A}$formula)
ICM entropy floor / $L_X\propto T^{2.6\text{–}3}$	AGN feedback; supernova preheating (~ 1 keV/particle)	No confirmed energy source; mass-dependent; non-universal	Relic collision entropy $K_{\text{relic}}$from $\gamma_{\text{rel}}\sim 1.5$–$3$; universal	1 (γ_{rel}, constrained by K_0)

The parameter count comparison is instructive. ΛCDM's proposed resolutions for these five tensions collectively invoke: a dark matter particle (undetected), a massive neutrino sector (constrained at $\sum m_\nu < 0.12\text{eV}$), an AGN feedback efficiency parameter (unconstrained from first principles), a supernova preheating energy fraction (unconstrained), and a hydrostatic mass bias (empirically disfavored at the required level). SCT resolves all five tensions with two additional parameters—$\mathcal{A}$ and γ_{rel}—both of which are constrained by existing observational data rather than treated as free parameters. The SCT resolution is therefore not only more economical but more predictive: it ties previously disconnected observations to a single originating cause.

8.2 Testable Predictions

The following numbered predictions consolidate all falsifiable predictions made in Sections 3–7, organized by observational priority.

Predictions testable with Euclid (launched 2023; full survey $\sim$2026–2030):

1. The satellite plane alignment correlation function $\langle \hat{\mathbf{n}}_A\cdot\hat{\mathbf{n}}_B\rangle$among spatially neighboring galaxy pairs (separation $d < 5\text{Mpc}$) exceeds the ΛCDM null expectation at $>3\sigma$significance. (Section 3, Prediction 3.2)
2. The cluster major-axis alignment correlation function $\xi_{\text{align}}(r)$remains significantly above zero at $r = 300$comoving Mpc, with $\xi_{\text{align}}(300)\gg\xi_{\text{align},\Lambda\text{CDM}}(300)\approx 0$. (Section 4, Prediction 4.1)

3. The hydrostatic mass bias $b \equiv 1 - M_{\mathrm{hyd}}/M_{\mathrm{WL}}$scales monotonically with cluster richness λ, confirming the amplification-richness relation $b(\lambda) \propto \mathcal{A}(\lambda) - 1$. (Section 5, Prediction 5.3)
4. The GGSL rate per cluster correlates with dynamical state: relaxed clusters show $\geq 2\times$higher GGSL rates than unrelaxed clusters of comparable mass. (Section 6, Prediction 6.1)
5. The radial distribution of GGSL events is more centrally concentrated in observed clusters than in any ΛCDM simulation suite, with the ratio $N_{\mathrm{GGSL}}(r < 0.3R_{200})/N_{\mathrm{GGSL}}(r < R_{200})$exceeding simulated values at $> 2\sigma$. (Section 6, Prediction 6.3)

Predictions testable with Rubin Observatory/LSST (first light achieved; full survey $\sim$2025–2034):

6. Satellite planes are more common among galaxy pairs with separation $d < 5$Mpc than among galaxy pairs with $d > 20$Mpc, with the prevalence ratio exceeding 2 :1. This tests the sibling hypothesis (Premise 34) directly using the LSST deep-field galaxy satellite census. (Section 3, Prediction 3.2)
7. The orbital poles of satellite planes show a statistically significant tendency to be perpendicular to the nearest cosmic filament axis, measured at $> 2\sigma$in a sample of ≥ 50galaxy systems with both satellite kinematics and filament reconstructions. (Section 3, Prediction 3.3)
8. The S_8tension weakens systematically with increasing redshift: $\Delta S_8 \equiv S_8^{\mathrm{CMB}} - S_8^{\mathrm{low-z}}$decreases from ~ 0.05at $z \sim 0.3$to $\lesssim 0.01$at $z \sim 1.5$, following the predicted redshift evolution of $\mathcal{A}(z) \propto (1 + z)^{-\gamma}$. (Section 5, Prediction 5.4)
9. The L_X–T_X slope $\alpha$for AGN-quiet groups ($\lambda_{\mathrm{radio}} < 10^{23}$ W/Hz at 1.4 GHz) is statistically indistinguishable from the slope for AGN-active groups ($\alpha \approx 2.7$–3.0 in both cases), falsifying the AGN-preheating explanation for the entropy floor. (Section 7, Prediction 7.4)

Predictions testable with CMB-S4 (construction $\sim$2027–2030; science $\sim$2030+):

10. The tSZ mass bias varies with cluster richness as $b(\lambda) \propto \lambda^{\beta}$with $\beta \approx 0.3$–0.5 (from Prediction 5.1), detectable at $> 3\sigma$in the CMB-S4 cluster catalog of $\sim 10^5$clusters.
11. CMB lensing convergence power spectrum measurements at $z \sim 2$–4 yield S_8values closer to the Planck CMB primary value ($S_8 \approx 0.83$) than to the low-redshift weak-lensing value ($S_8 \approx 0.77$), confirming the redshift-dependent superposition correction. (Section 5, Prediction 5.5)
12. The mean cluster orientation alignment at $z > 1.5$is stronger than at $z < 0.5$by a factor of ≥ 1.5in the alignment correlation amplitude ξ_0, following the monotonic increase $d\langle\cos\,\theta_{\mathrm{align}}\rangle/dz > 0$predicted by SCT (Prediction 4.3) and inconsistent with ΛCDM's non-monotonic prediction.

Predictions testable with Athena X-ray Observatory ($\sim$2030s):

13. The ICM entropy floor $K_0$is present at the same amplitude in clusters at $z = 1.5$–2 as at $z = 0$, with $K_0(z = 2)/K_0(z = 0) \in [0.8,1.2]$, falsifying AGN-preheating models that predict $K_0(z = 2)/K_0(z = 0) \ll 1$. (Section 7, Prediction 7.3)
14. The scatter in K_0across clusters correlates with local large-scale structure density (as characterized by cosmic web reconstruction), with clusters in void edges showing

systematically higher K_0 than clusters in dense filaments at comparable mass, at $> 2\sigma$ significance. (Section 7, Prediction 7.2)

These fourteen predictions span three generations of observational facilities and provide a comprehensive program for testing or falsifying SCT's three core mechanisms. Crucially, several of the predictions are in direct conflict with what ΛCDM-based models predict—notably Predictions 7, 9, 12, and 13—ensuring that a comprehensive observational campaign can definitively distinguish between the two frameworks.

8.3 Limitations

This paper presents SCT as a theoretical framework supported by analytical derivations in linearized GR and SR, with quantitative predictions that match observational data across five tensions. We acknowledge several important limitations that must be addressed in future work.

The superluminal collision premise cannot currently be simulated. The collisions between nested comoving frames that SCT invokes involve relative velocities of $\sim 1.5c$–$67c$ in parent frames, energy scales far beyond the kinetic regimes accessible to current cosmological simulation codes. SCT's collision premise is self-consistent within the framework of GR (the speed limit applies within local frames, not between frames; Premises 20–22), but numerical validation of the collision thermodynamics, debris field structure, and angular momentum distribution requires entirely new simulation paradigms. Until such simulations are developed, the collision mechanism rests on analytical arguments and consistency checks rather than numerical confirmation.

The amplification factor $\mathcal{A}(N, \sigma_v, R)$ requires numerical validation. The formula for $\mathcal{A}$ derived in Section 2.2 is based on linearized GR and a coherent-phase approximation for gravitational wave superposition. Real clusters are not perfectly coherent, and the transition from coherent to incoherent superposition as a function of N, σ_v, and R requires a full numerical treatment. Non-linear effects, the role of dark matter halos (in a scenario where dark matter is reinterpreted rather than eliminated), and the coupling between gravitational superposition and baryonic physics all require detailed modeling. The values $\mathcal{A} \approx 1.10$–1.20 (cluster scale) and $\mathcal{A} \approx 3.2$–4.0 (sub-halo scale) derived here are therefore order-of-magnitude estimates awaiting numerical confirmation.

Collision thermodynamics is presented at first order. The entropy floor calculation in Section 7 uses first-order thermodynamics applied to a simple two-body collision model. Real collisions between nested comoving structures involve gradients of γ_{rel}, non-uniform mass distributions, multiple collision generations with varying b and $\mathbf{v}_{\text{rel}}$, and non-equilibrium plasma physics at extreme temperatures. The entropy floor derived here is therefore a lower-bound estimate of the true complexity. More realistic modeling may be needed to achieve the precision required for comparison with Athena-quality X-ray data.

8.4 Conclusions

This paper has demonstrated that five of the most statistically significant and theoretically recalcitrant tensions in ΛCDM cosmology are resolved, within the framework of standard GR and SR, by replacing the assumption of a singular hot dense origin with a succession of superluminal collisions between nested comoving frames of reference. The change is minimal in mathematical structure—no new particles, no new fields, no modifications to the field equations beyond reinterpreting existing terms—yet its consequences are profound: it provides a first-principles, geometric account of phenomena that ΛCDM can address only by invoking a growing list of ad hoc mechanisms, each tailored to a specific tension and disconnected from the others.

The three mechanisms derived here—angular momentum conservation from collision geometry, gravitational superposition of comoving bodies, and collision thermodynamics—are not independent inventions. They are all consequences of the same changed assumption, operating at different scales and through different physical channels. The angular momentum mechanism and the thermodynamic mechanism both arise directly from the collision kinematics. The superposition mechanism arises from the co-moving structure of the debris field that the collision creates. All three are, in this sense, facets of a single unified picture. This internal coherence—the same collision event simultaneously explaining satellite plane alignments, cluster orientation alignments, the S_8tension, the GGSL excess, and the entropy floor—is the paper's central claim, and it would be difficult to achieve by accident.

The fourteen falsifiable predictions compiled in Section 8.3 provide a concrete program for testing SCT against the next generation of observational facilities. Several of these predictions—particularly the universality of the entropy floor at $z > 1.5$, the monotonic increase of cluster alignment with redshift, and the independence of the L_X–T_X slope from AGN activity—are in direct conflict with what ΛCDM and its preheating models predict. If future observations confirm these predictions, the case for a changed foundational assumption will be compelling. If any of them are falsified, SCT will need to be revised or abandoned in favor of a better account.

The deepest implication of this work is not any of the specific predictions. It is the observation that five of the most persistent anomalies in modern cosmology share a common resolution when a single foundational assumption is changed. This pattern—multiple independent anomalies cured by one changed premise—is the characteristic signature of a paradigm in need of revision rather than repair. ΛCDM has been enormously successful, and that success is not diminished by the tensions discussed here. But the history of physics suggests that when a theory requires increasingly elaborate and disconnected auxiliary hypotheses to accommodate new data, the most productive path forward is not to add another hypothesis, but to ask which assumption, if changed, makes the hypotheses unnecessary. This paper argues that the assumption of a singular hot dense origin is that assumption, and that Successive Collision Theory—grounded in eternal time, infinite space, and the scale-invariant application of General Relativity—provides a more natural foundation for the universe we actually observe. A comprehensive re-evaluation of standard model foundational assumptions, guided by the growing body of high-precision anomalies, is not merely warranted: it is overdue.

For Further Reading

These were just 5 out of 231 tensions that are all easily solved using the premises of Successive Collision Theory. Feel free to see how Successive Collision Theory can be used to explain any of 231 current known issues with the Lambda-CDM model by visiting our webpage, https://thenaturalstateofnature.org/231, or by downloading our 45 years of research, 400+ pages of the detailed mathematical predictions of Successive Collision Theory with this DOI address: https://doi.org/10.13140/RG.2.2.23479.79528. My second-to-last prediction is that with about 56 premises, we can train most good AI's to recode in SWIFT how to fix our best simulations so we can zoom out and begin to watch big bangs happen. And my last prediction, for now, is that within 42 years, the transition from the Lambda-CDM model to the Successive Collision Model will be completed when the Big Bang Equations, a specific set of approximate overlapping collision fronts and secondary/tertiary collisions, will be mapped with sufficient approximation to produce a pocket of spacetime very much like the one we call home.

Other papers waiting for peer review that you might be interested in, if you got this far, include :

(PDF) "An End to Black Hole Singularities" Polyquark Cores and Quark Degeneracy Pressure, A Lattice-QCD-Based Equation of State for Finite-Density Black Hole Interiors and the Stabilization of Gravitational Collapse

(PDF) Migrating From Flat Boosts to Frame Trees: A Hierarchical Lorentz Approach for High-Precision Cosmology Version

(PDF) CMB Power Spectrum from Successive Collision Theory: A Quantitative Framework at Planck Precision

References

1. [1110.2378] Entropy Creation in Relativistic Heavy Ion Collisions
https://arxiv.org/abs/1110.2378
2. [1807.06209] Planck 2018 results. VI. Cosmological parameters - arXiv
https://arxiv.org/abs/1807.06209
3. [2303.10190] Planes of satellites, no longer in tension with LCDM
https://arxiv.org/abs/2303.10190
4. [2509.09678] Cosmic tensions indirectly correlate with ... - arXiv
https://arxiv.org/abs/2509.09678
5. 93. Three critical tests alone in this year 2024 find dark matter and the ΛCDM model to be incorrect
https://darkmattercrisis.wordpress.com/2024/07/08/93-three-critical-tests-alone-in-this-year-2024-find-dark-matter-and-the-lcdm-model-to-be-incorrect/
6. A whirling plane of satellite galaxies around Centaurus A challenges ...
https://pubmed.ncbi.nlm.nih.gov/29420285/
7. A whirling plane of satellite galaxies around Centaurus A ... - Science
https://www.science.org/doi/10.1126/science.aao1858
8. An excess of small-scale gravitational lenses observed in ...
https://www.science.org/doi/10.1126/science.aax5164
9. An excess of small-scale gravitational lenses observed in galaxy ...
https://arxiv.org/abs/2009.04471
10. Angular momentum - Wikipedia
https://en.wikipedia.org/wiki/Angular_momentum

11. Anomalous EDGES 21-cm signal and a moduli dominated era
https://link.aps.org/doi/10.1103/PhysRevD.100.103007
12. AST233 Lecture notes Part 2
https://astro.pas.rochester.edu/~aquillen/ast233/lectures/lecture2.pdf
13. Checking in on Troubles with Dark Matter - Triton Station
https://tritonstation.com/2023/06/27/checking-in-on-troubles-with-dark-matter/
14. Cluster Alignments and Ellipticities in LCDM Cosmology - arXiv
https://arxiv.org/abs/astro-ph/0409652
15. Cluster Alignments and Ellipticities in ΛCDM Cosmology - ADS
http://ui.adsabs.harvard.edu/abs/2005ApJ...618....1H/abstract
16. Constraining exotic dark matter models with the dark ages 21-cm ...
https://academic.oup.com/mnras/article/527/1/1461/7335733
17. Cosmological constraints from the thermal Sunyaev Zeldovich ...
https://inspirehep.net/literature/1677729
18. Dissecting the Thermal SZ Power Spectrum by Halo Mass and ...
https://astro.theoj.org/article/89991-dissecting-the-thermal-sz-power-spectrum-by-halo-mass-and-redshift-in-spt-sz-data-and-simulations
19. EDGES 21 cm anomaly and properties of dark matter
https://www.sciencedirect.com/science/article/pii/S0370269318306488
20. Entropy Analysis in Relativistic Heavy-Ion Collisions - Ahmad - 2013
https://onlinelibrary.wiley.com/doi/10.1155/2013/836071
21. Evolution of Cluster Alignments as Evidence of Large-scale ... - arXiv
https://arxiv.org/html/2506.19826v1
22. Examining the Sigma-8 Tension with DESI Galaxies
https://www.desi.lbl.gov/2025/01/12/brushing-away-dust-to-uncover-cosmology-examining-the-sigma-8-tension-with-desi-galaxies/
23. Followup papers and analyses on reported EDGES 21-cm detection ...

https://www.reddit.com/r/Physics/comments/82zjoj/followup_papers_and_analyses_on_reported_edges/
24. Heating and enriching the intracluster medium - Oxford Academic

https://academic.oup.com/mnras/article/428/2/1225/1002293
25. Impact parameter - Wikipedia
https://en.wikipedia.org/wiki/Impact_parameter
26. Intracluster Medium Entropy Profiles for a Chandra Archival Sample ...
https://arxiv.org/abs/0902.1802
27. Is Dark Energy Getting Weaker? New Evidence Strengthens the Case.
https://www.quantamagazine.org/is-dark-energy-getting-weaker-new-evidence-strengthens-the-case-20250319/
28. Is the sigma-8 tension being sheared away? Cosmology with weak ...
https://astrobites.org/2025/04/03/sigma8-tension-kids-legacy-galaxyshear/
29. KiDS
https://kids.strw.leidenuniv.nl
30. KiDS-Legacy: Consistency of cosmic shear measurements and joint ...
https://arxiv.org/pdf/2503.19442.pdf
31. Linearized gravity - Wikipedia
https://en.wikipedia.org/wiki/Linearized_gravity
32. Looking for cracks in the standard cosmological model
https://www.yorku.ca/news/2023/07/19/looking-for-cracks-in-the-standard-cosmological-model/
33. Magnetic fields and Sunyaev-Zel'dovich effect in galaxy clusters - arXiv
https://arxiv.org/abs/1003.3718
34. NEW TESTS OF THE CLUSTER ENTROPY FLOOR HYPOTHESIS (PDF)
https://www.astroscu.unam.mx/rmaa/RMxAC..17/PDF/RMxAC..17_ababul.pdf
35. Nuclear and Particle Physics - Lecture 2 Relativistic kinematics
http://www.hep.ph.ic.ac.uk/~dauncey/will/lecture02.pdf
36. Observations at tension with cosmological models
https://cosmology.info/resources/marmet_l/tension-with-models.html
37. Optical Depth to Reionization, τ - NASA Lambda
https://lambda.gsfc.nasa.gov/education/graphic_history/taureionzation.html
38. Planck 2013 results. XX. Cosmology from Sunyaev–Zeldovich ...
https://escholarship.org/content/qt5hc762wm/qt5hc762wm.pdf
39. Planck 2015 results - XXIV. Cosmology from Sunyaev-Zeldovich ...
https://orbit.dtu.dk/files/127985540/aa25833_15.pdf
40. Planck 2018 results - VIII. Gravitational lensing - OSTI.GOV
https://www.osti.gov/pages/biblio/1775409
41. Planes of satellites are not a problem for (just) ΛCDM (PDF)
https://par.nsf.gov/servlets/purl/10508372
42. Planes of satellites around Milky Way/M31-mass galaxies in the ...
https://academic.oup.com/mnras/article/504/1/1379/6217428
43. Planes of satellites, no longer in tension with ΛCDM - Inspire HEP
https://inspirehep.net/literature/2644300
44. power spectrum of the thermal Sunyaev–Zeldovich effect
https://academic.oup.com/mnras/article/540/1/1055/8125474
45. Reionization optical depth and CMB–BAO tension in punctuated ...
https://academic.oup.com/mnras/advance-article/doi/10.1093/mnras/staf1892/8313630
46. Satellite Galaxies All in a Row—How So? - Astrobites
https://astrobites.org/2018/07/11/satellite-galaxies-all-in-a-row/
47. Sigma-8 tension is a drag | Phys. Rev. D - APS Journals
https://link.aps.org/doi/10.1103/PhysRevD.107.123538
48. Simplified derivation of the gravitational wave stress tensor ... - PMC
https://pmc.ncbi.nlm.nih.gov/articles/PMC5081580/
49. SNe heating and the chemical evolution of the intra-cluster medium
https://www.sciencedirect.com/science/article/abs/pii/S1384107602001367
50. Tensions in the ΛCDM paradigm - MITP (Indico)
https://indico.mitp.uni-mainz.de/event/130/attachments/379/2469/Tensions10.pdf
51. The EDGES 21 cm anomaly and properties of dark matter
https://www.sciencedirect.com/science/article/pii/S0370269318306488

52. The entropy and energy of intergalactic gas in galaxy clusters
https://academic.oup.com/mnras/article-pdf/315/4/689/3671741/315-4-689.pdf
53. The Milky Way's plane of satellites is consistent with ΛCDM
https://durham-repository.worktribe.com/output/1181135/the-milky-ways-plane-of-satellites-is-consistent-with-lcdm
54. The phase diagram of the intergalactic medium and the entropy floor ...
https://academic.oup.com/mnras/article/344/1/53/1011104
55. Thermal SZ power spectrum by halo mass and redshift in SPT-SZ data (same as 17; kept as 17 above for detail, omit here if you want strict uniqueness)
56. Three critical tests and dark matter crisis (already counted as 5; omit if you want a strictly unique-title list)
57. Was 2024 the year we finally started to understand dark energy?
https://www.space.com/2024-finally-understand-dark-energy
58. What big physics problem is unlikely to be solved in the next 20–50 ...

https://www.reddit.com/r/AskPhysics/comments/1ix8dyg/what_big_physics_problem_is_unlikely_to_be_solved/
59. What new observations tell us about Planes of Satellite Galaxies
https://arxiv.org/html/2312.01830v1