

<b>From Chaos To Collisothermal Cosmogenesis</b>	
<b>Early Structure Formation Under Successive Collision Theory and the Resolution of the JWST Mass Assembly Crisis</b>	
DR JM NIPOK N.J.I.T. <a href="https://orcid.org/0009-0006-3940-4450">https://orcid.org/0009-0006-3940-4450</a>	Copyright CC BY-NC-SA 4.0. March 7 2026 <a href="https://doi.org/10.13140/RG.2.2.16235.60968">https://doi.org/10.13140/RG.2.2.16235.60968</a> Version 3.7

## ABSTRACT

---

The James Webb Space Telescope has confirmed galaxies at  $z = 14.18$  with dynamical masses of order  $10^8 M_{\text{sun}}$  and oxygen enrichment exceeding  $0.1 Z_{\text{sun}}$  (Carniani et al. 2024), galaxies at  $z = 14.44$  with super-solar nitrogen-to-carbon ratios at only 280 megayears after the Big Bang (Naidu et al. 2025), and a population of systems at  $z = 7-9$  with stellar masses above  $10^N M_{\text{sun}}$  requiring star-formation efficiencies of 47-52 percent (Xiao et al. 2024), all exceeding the Lambda Cold Dark Matter stellar mass ceiling by factors of ten to one hundred and lying four to sixteen sigma above the theoretical maximum permitted by the halo mass function at those epochs. Quenched massive galaxies at  $z = 7.29$  exceed every simulation prediction by more than two orders of magnitude in number density, and overmassive black holes at  $z$  greater than 7 carry BH-to-stellar mass ratios ten to one thousand times above the local Magorrian relation, including a  $1.6 \times 10^9 M_{\text{sun}}$  black hole at  $z = 7.642$  that cannot have grown from any stellar-mass seed via Eddington-limited accretion even if seeded at  $z = 30$  (Wang et al. 2021).

We demonstrate that these observations are natural consequences of Successive Collision Theory (SCT), in which successive superluminal collisions between nested comoving frames seed proto-structures directly into the pre-recombination plasma through three mechanisms. First, collision geometry deposits proto-structure masses  $M_{\text{proto}} =$

$\alpha_{\text{th}} * f_{\text{b}} * \mu * \Omega(\text{b}, R_1, R_2)$  in the range  $10^7$  to  $10^{11} M_{\text{sun}}$ , set by dynamics rather than gravitational growth rate, eliminating the assembly time problem by construction. Second, angular momentum inheritance  $J = \mu(\text{b} \times v_{\text{rec}})$  establishes morphological type at the seeding epoch and is preserved by exact conservation, explaining the morphologically mature systems observed at  $z$  greater than 3. Third, direct-collapse black hole seeds of  $10^7$  to  $10^9 M_{\text{sun}}$  form from head-on collision geometry without requiring Eddington-limited accretion.

The primordial power spectrum with  $n_{\text{s}} \sim 0.965$  is derived from first principles by applying the Central Limit Theorem to a population of approximately  $10^4$  independent collision events spanning 29 hierarchical nesting levels, with the red tilt arising from the finite dynamic range of the cascade through the relation  $n_{\text{s}} = 1 \sim 1/L$ . The framework makes three tiers of falsifiable predictions. With existing JWST data, the disk fraction at  $z$  greater than 10 should exceed ten percent and the number density of massive galaxies should follow a power-law rather than exponential decline above  $z = 14$ . With the Roman Space Telescope High Latitude Wide Area Survey, SCT predicts 550 to 4770 detections of galaxies with  $M_{\text{sun}}$  greater than  $10^{[N]} M_{\text{sun}}$  at  $z = 12-15$ , compared to fewer than three predicted by  $\Lambda\text{CDM}$ . With CMB-S4 and PIXIE-class instruments, the tensor-to-scalar ratio should satisfy  $r$  less than  $10^{-5}$  and a dipolar  $y$ -type spectral distortion aligned with the large-scale angular momentum coherence axis should be detectable.

The thermal energy crisis in early structure formation extends beyond stellar mass. Zhou et al. (2025) report a thermal Sunyaev-Zeldovich detection in the protocluster SPT2349-56 at  $z = 4.3$ , yielding a total intracluster medium thermal energy of  $(11.8 \pm 1.2) \times$

$10^{60}$  erg — an order of magnitude above what gravitational collapse of a  $10^{13}$  solar-mass halo can produce, lying 6.4 sigma above TNG-Cluster simulation predictions and five times above the universal mass-Compton-Y scaling relation. SCT provides a direct thermodynamic account of this born-hot intracluster medium: proto-ICM thermal energy is inherited from the collision cascade itself, seeded before recombination at a characteristic scale set by collision geometry rather than virialization, with no requirement for sustained AGN feedback to supply energy that gravity alone cannot.

## 1. INTRODUCTION

The standard cosmological model, Lambda Cold Dark Matter, makes a specific and quantitative prediction about the maximum stellar mass observable at high redshift. Under hierarchical assembly, galaxies grow through the merger of smaller dark matter halos, with baryons cooling into the potential wells those halos provide. The halo mass function at any epoch is constrained by the matter power spectrum and the age of the universe at that epoch, setting a hard ceiling on the total baryonic mass available for star formation within any given volume. At redshift  $z = 10$ , approximately 480 megayears after the Big Bang, the most massive collapsed halos in a standard LambdaCDM realization have virial masses of order  $10^N M_{\text{sun}}$ , and with a cosmologically typical star-formation efficiency of 5 to 15 percent, the maximum observable stellar mass at  $z$  greater than 10 is predicted to lie well below  $10^{10} M_{\text{sun}}$ .

Galaxies at these redshifts are expected to be small, irregular, metal-poor, and morphologically disturbed, the primitive building blocks that will eventually assemble into

the massive, structured systems observed at  $z$  less than 3. This prediction is not a peripheral feature of  $\Lambda$ CDM that can be adjusted without cost. It is a direct consequence of the model's foundational architecture: hierarchical, bottom-up assembly driven by gravitational growth from quantum vacuum fluctuations seeded during inflation, proceeding at a rate set by the matter power spectrum and the age of the universe at each epoch.

The JWST mass assembly crisis described above concerns stellar mass and morphology. An independent and equally severe crisis has now emerged in the thermodynamic state of the earliest galaxy clusters. The intracluster medium — the hot, diffuse gas that constitutes most of the baryonic mass of present-day galaxy clusters — is predicted by  $\Lambda$ CDM to assemble gradually and be heated primarily by gravitational shock waves as protoclusters virialize. This process should produce an ICM that grows progressively hotter and more massive with cosmic time, with the youngest protoclusters possessing the coolest, least-developed intracluster gas. Zhou et al. (2025) have falsified this prediction directly. Using deep ALMA Band-3 observations of the protocluster SPT2349-56 at  $z = 4.3$  — a system at 1.4 billion years after the Big Bang containing at least 30 dusty star-forming galaxies and three radio-loud AGN within a 100-kpc region — they detect a thermal Sunyaev-Zeldovich signal with a significance of 10.4 sigma in Fourier space and an integrated flux density of  $-157 \pm 16$  microjansky.

The inferred total thermal energy of  $(11.8 \pm 1.2) \times 10^{60}$  erg exceeds the maximum energy producible by gravitational collapse of the observed halo by an order of magnitude, lies 6.4 sigma above the TNG-Cluster cosmological simulation suite, and exceeds the universal mass-Compton-Y scaling relation by a factor of five. Reproducing the observed

Sunyaev-Zeldovich signal under the standard shock-heating model would require a gas mass of  $5 \times 10^{13}$  solar masses — five times the total mass of the system. This is not a marginal tension. It is a thermodynamic falsification of the LCDM ICM assembly picture at  $z = 4.3$  that is structurally identical to the stellar mass falsification at  $z = 14$ : both reveal physical processes that LCDM does not contain and cannot accommodate by adjusting parameters within its existing architecture. The two crises share a common origin, and the same framework that resolves the stellar mass assembly crisis resolves the ICM thermal energy crisis by the same mechanism applied to the same epoch.

The James Webb Space Telescope has falsified this prediction across multiple independent observational dimensions simultaneously. The first spectroscopically confirmed galaxy beyond  $z = 13$ , JADES-GS-z13-0, demonstrated that star-forming systems existed at only 325 megayears after the Big Bang, beyond the epoch where standard simulations predicted the first luminous objects (Curtis-Lake et al. 2023). This was rapidly succeeded by confirmations at still higher redshifts. JADES-GS-z14-0 at  $z = 14.18$  has a dynamical mass of order  $10^8 M_{\text{sun}}$ , spatially resolved structure, and oxygen enrichment exceeding  $0.1 Z_{\text{sun}}$  requiring multiple prior generations of star formation that no simulation predicted at 290 megayears post-Big Bang (Carniani et al. 2024). MoM-z14 at  $z = 14.44$  exhibits super-solar nitrogen-to-carbon ratios and chemical enrichment at only 280 megayears after the Big Bang that far exceeds any standard chemical evolution model (Naidu et al. 2025).

These are not marginal detections at the limit of model predictions. They represent systems whose chemical maturity alone implies prior stellar generations operating at

epochs when LambdaCDM predicts the universe had not yet produced its first generation of stars in sufficient numbers to enrich the interstellar medium. The crisis deepens when the stellar mass function is examined statistically. Labbe et al. (2023) reported a population of six candidate galaxies in the CEERS field with photometric redshifts  $z = 7.4-9.1$  and individual stellar masses exceeding  $10^{[N]} M_{\text{sun}}$ , with one candidate approaching  $10^{[N]} M_{\text{sun}}$ . Boylan-Kolchin (2023) performed a formal statistical analysis demonstrating that these systems lie at the absolute upper limit of what the LambdaCDM dark matter halo mass function permits, requiring star-formation efficiencies near unity, three to five times the maximum observed at any subsequent epoch, or revisions to the cosmological model itself. Xiao et al. (2024) spectroscopically confirmed three ultra-massive galaxies at  $z = 5-9$  with stellar masses at or above  $10^{[N]} M_{\text{sun}}$ , requiring baryon conversion efficiencies of approximately fifty percent sustained over the first billion years. The problem is not confined to stellar mass. Overmassive black holes appear at  $z$  greater than 7 with BH-to-stellar mass ratios ten to one thousand times above the local relation, including UHZ1 at  $z \sim 10.1$  with a black hole mass comparable to or exceeding its host galaxy's total stellar mass (Bogdan et al. 2024; Natarajan et al. 2024), and QSO J0313-1806 at  $z = 7.642$  with  $1.6 \times 10^9 M_{\text{sun}}$  in a central black hole that cannot have grown from any stellar-mass seed via Eddington-limited accretion even if seeded at  $z = 30$  (Wang et al. 2021).

Morphological maturity compounds the mass problem: a barred spiral galaxy structurally indistinguishable from a Milky Way analogue has been confirmed at  $z \sim 3$  (Costantin et al. 2023), and the ratio of spiral to elliptical to irregular galaxies is approximately constant across the observable universe to  $z \sim 10$  (Ferreira et al. 2024),

directly contradicting the expectation of progressively more disturbed, merger-driven morphologies at higher redshift. Massive quenched galaxies appear at  $z = 7.29$  in number densities exceeding every simulation prediction by more than two orders of magnitude (Weibel et al. 2025). The totality of these detections does not constitute a tension in the statistical sense of a parameter value lying outside a confidence interval. It constitutes a falsification in the architectural sense: the observations require physical processes that LambdaCDM does not contain and cannot accommodate by adjusting parameters within its existing structure. Successive Collision Theory resolves this crisis not by modifying the late-time physics of galaxy formation but by replacing the single foundational assumption from which the assembly time problem originates. In SCT, our observable patch of spacetime was produced by a succession of superluminal collisions between nested comoving frames of reference, events in which the full kinetic energy of two immense independent pockets was deposited into an overlap volume before any internal signal could traverse the collision interface.

The collision cascade seeded proto-structures directly into the pre-recombination plasma through three mechanisms. Head-on collision geometry produces maximally dense, minimally rotating remnants that collapse directly into massive black hole seeds in the range  $10^7$  to  $10^9 M_{\text{sun}}$  without requiring hierarchical merger sequences or sustained accretion. Grazing collisions deposit angular momentum  $J = \mu(b \times v_{\text{rec}})$  that propagates through all descendant structures, establishing morphological type at the seeding epoch rather than through merger history. The characteristic mass scale  $M_{\text{proto}} = \alpha_{\text{th}} * f_{\text{b}} * \mu * \Omega(b, R_1, R_2)$  is set by collision dynamics and falls in the range  $10^7$  to  $10^{11}$

$M_{\text{sun}}$  across all physically plausible parameter combinations, placing collision-seeded proto-structures above the  $\Lambda\text{CDM}$  ceiling by construction. Because all of these mechanisms operate before recombination and terminate long before the photon-baryon fluid decouples, they leave no observational imprint on the CMB power spectrum that would violate Planck precision constraints. The assembly time problem does not exist in SCT because assembly in the hierarchical sense is not the operative mechanism. Structure is seeded. It then grows. The prior generation of alternative cosmologies that attempted to address similar problems failed for specific reasons that do not apply to SCT. Plasma cosmology could not reproduce the CMB acoustic peak structure because it attempted to replace the standard photon-baryon fluid physics rather than operating within it.

Ekpyrotic and bouncing cosmology models face the problem of producing a nearly scale-invariant primordial spectrum without fine-tuning the scalar field potential. Colliding brane models require extra dimensions and string-theoretic scaffolding that introduces more free parameters than they remove. SCT avoids each of these failure modes. It operates entirely within standard Boltzmann transport physics after the cascade terminates. It derives the spectral index  $n_s \sim 0.965$  from first principles by applying the Central Limit Theorem to a population of approximately  $10^4$  independent collision events distributed across 29 hierarchical nesting levels, requiring no scalar field potential and no new physics beyond GR and SR. It introduces no extra dimensions, no new particles, and no new fields. The comparison between SCT and its predecessors is examined in detail in Section 5.3, and the parameter economy of SCT relative to  $\Lambda\text{CDM}$  plus inflation is analyzed in Section 5.4.

This paper is organized as follows. Sections 1.1 through 1.4 define the four foundational SCT constructs used throughout the paper: nested comoving frames, the superluminal collision, the nesting hierarchy, and the collision cascade with its temporal structure and observational constraints. Section 2 derives the LambdaCDM stellar mass ceiling from first principles and presents the JWST detections that exceed it as a quantified falsification, together with an analysis of why the three proposed LambdaCDM remedies each fail on examination. Section 3 presents the full SCT mechanism for early structure formation. Section 4 presents three tiers of falsifiable predictions with explicit falsification criteria, observable tests, and timelines. Section 5 addresses four substantive objections to the SCT framework. Section 6 states the conclusions and outlines the program of future work needed to fully constrain the SCT parameter space.

## **1.1 Nested Comoving Frames**

A comoving frame in SCT is a bounded, self-gravitating region of spacetime, a physical pocket of matter and energy whose constituent matter shares a common bulk velocity to a degree sufficient to define a single coherent rest frame for that region. This is not an abstract coordinate choice. A pocket is a physical object with a defined formation history. Its nine measurable bulk properties are rotation rate, orbital motion, center of mass, luminosity, gravitational field, magnetic field, electric field, spatial evolution, and inherited time perception. These properties are in principle observable from outside the pocket in the same way that a galaxy cluster's bulk properties are observable from beyond its virial radius. The boundary of a pocket is the phase- space surface within which gravitational binding

energy exceeds kinetic energy of escape. Observers within it measure local physics entirely consistent with special relativity. The bulk motion of the pocket as a whole is specifiable as a single vector in the coordinate system of any larger enclosing structure. Frames are nested when one such pocket is embedded within another of strictly larger spatial extent and strictly larger total mass-energy, which is itself embedded within a still larger pocket, and so on through a hierarchy of scales extending upward without limit.

An unbounded spatial extent is the natural consequence of applying Einstein's field equations, which contain no preferred length scale, to a matter distribution with no outer boundary. Each level of the hierarchy is related to its parent by a Lorentz boost. The nesting is not merely spatial containment but a precise kinematic relationship in which the daughter frame's bulk motion is defined relative to the parent frame's rest frame, and the parent frame's gravitational potential constitutes the background metric within which the daughter evolves. This hereditary chain means the proper time rate of any object is the cumulative product of kinematic and gravitational factors all the way up through the entire parent-frame hierarchy. The GPS system provides an operational confirmation of this two-level composition: the Schwarzschild gravitational redshift of the Earth's potential well and the special-relativistic time dilation of satellite orbital velocity must both be applied simultaneously to maintain GPS accuracy. SCT asserts that this composition principle continues upward through every level of the hierarchy without ceiling, which is a direct extrapolation of the same GR that GPS already confirms at two levels simultaneously.

Our observable universe is a sphere of roughly 46.5 gigalight-years radius, bounded only by the distance light has traveled since our local collision event 13.8 billion years ago,

and it is one pocket among infinitely many. It is not a participant in the collision cascade that produced it. It is the thermalized overlap remnant generated when two independent pockets, each the terminal output of its own nesting lineage of entirely unknown and potentially vast depth, intersected within some common ancestor frame at an unspecified remove above both. The matter in our observable universe existed before that collision as the content of those two pre-existing pockets, with compositions, density profiles, angular momenta, and magnetic field configurations inherited from prior collision generations. The collision thermalized that pre-existing matter and did not create it.

## **1.2 The Superluminal Collision**

A superluminal collision in SCT refers to an event in which two independent comoving pockets approach one another at a closing velocity  $v_{\text{rec}}$  greater than  $c$ , as measured from within their shared ancestor frame. Special relativity's speed limit applies to a specific physical process: the acceleration of an object that begins at rest within an inertial frame by a local force acting within that frame. It does not apply to the relative velocity between two objects that were never in the same inertial frame, were set in motion by independent processes in causally disconnected regions, and whose relative velocity was never built up by any local acceleration. Consider two pockets F1 and F2, each the terminal product of its own independent nesting lineage, whose trajectories bring them into intersection within some common ancestor frame F0. Each pocket constitutes its own local inertial reference frame and obeys all standard relativistic physics internally. However, F1 and F2 each possess their own bulk velocity relative to F0, inherited from their respective independent

formation momenta with no causal connection between the two formation histories. The closing velocity between F1 and F2 as measured from F0 is not subject to the intra-frame speed limit. No object within F1 or F2 is moving faster than  $c$  relative to its own local inertial frame.

The superluminality is a property of the inter-frame kinematics in the ancestor frame coordinate system and not a violation of any local physics. Standard cosmology already accepts that galaxies beyond the Hubble radius recede at velocities exceeding  $c$  without violating GR, because that recession is not the result of local acceleration. SCT extends this same precedent to the pre-formation kinematics of independently produced pockets. The collision itself occurs at the moment the boundaries of F1 and F2 first overlap. Because the collision interface advances at  $v_{\text{rec}} > c$  while all internal signals are limited to  $c$ , the entire pocket is engulfed before any internal communication can warn the interior. The full kinetic energy of both pockets is therefore deposited into the overlap volume essentially simultaneously, producing a violent, shock-driven thermalization event. Nothing local is superluminal, and the global kinematics are governed by the ancestor frame metric.

### **1.3 The Nesting Hierarchy**

The nesting hierarchy is not postulated arbitrarily. It is the natural consequence of applying gravitational self-organization to an unbounded distribution of matter across scales under Einstein's field equations, which contain no preferred length scale. At each level of the hierarchy, a comoving frame is defined by the gravitational coherence length of its constituent matter, specifically the scale below which internal velocity dispersion is

small compared to the escape velocity of the region so that the matter is gravitationally bound and moves coherently. The specific mass and spatial extent of each hierarchical level are determined by the interplay between local matter density, velocity dispersion at that scale, and the gravitational correlation length set by the two-point mass function of the matter distribution.

These quantities are related by the virial condition applied hierarchically, in which each level's characteristic velocity dispersion and radius satisfy  $2K + U = 0$  in the time-averaged sense. This self-similar virial structure produces characteristic mass scales separated by large factors at each level rather than continuously distributed, so that the allowed virial configurations form a discrete ladder rather than a continuum. Each level is related to its parent by a Lorentz boost encoding the bulk velocity difference between them, so that cosmological redshift encodes the composition of Lorentz boosts across all intervening hierarchy levels rather than a single recession velocity. This interpretation of cosmological redshift as a kinematic composition of boosts across hierarchical levels is derived in the supplementary treatment in Nipok (2026d).

## **1.4 The Collision Cascade**

The collision cascade refers to a sequence of discrete collision events, each physically distinct, each depositing energy, angular momentum, and density structure into what would become our observable universe, proceeding in two physically distinguishable phases before terminating. The word successive in Successive Collision Theory refers specifically to this ordered sequence of discrete collision events and must not be conflated

with the nesting hierarchy of comoving frames described in Section 1.3. The first phase consists of superluminal collision events in which the closing velocity between interacting pockets exceeded  $c$  in their shared ancestor frame, producing extreme kinetic energy regimes including temperatures reaching the QCD and potentially electroweak scale in compressed regions, exotic non-equilibrium states, the largest angular momentum depositions per stage, and density perturbations at the greatest spatial scales. The second phase consists of subluminal collision events in which relative velocities had dissipated below  $c$ , producing progressively more nearly thermal conditions approaching the state the photon-baryon fluid inherits at recombination. The entire cascade terminated long before recombination.

Three independent observational anchors establish this constraint. First, primordial nucleosynthesis abundances of helium, deuterium, and lithium require the photon-baryon plasma to have been in thermal equilibrium at temperatures of order 1 MeV, with no active energy injection during or after that epoch. Any cascade stage occurring after  $z$  approximately  $10^7$  would disturb the neutron-to-proton ratio and alter the predicted light-element abundances. Second, the spectral purity of the CMB as measured by COBE/FIRAS requires no non-standard energy injection above  $z$  approximately 50,000. Third, the precise positions of the CMB acoustic peaks require that the photon-baryon fluid evolve as a clean, unperturbed acoustic system from shortly after cascade termination through decoupling. The cascade terminus is therefore constrained to  $z$  greater than 50,000, hundreds of thousands of years before photon decoupling. The post-cascade plasma carried preserved large-scale angular momentum protected by exact conservation across all thermalization

stages, and a geometrically structured baryon asymmetry whose spatial distribution reflects the geometric history of the collisions that swept through each region. The cumulative density perturbation spectrum encodes the integrated output of every cascade stage. This structured initial condition is what the Plasma Equivalence Theorem, derived in Section 5.1, establishes as sufficient to reproduce the Planck-precision CMB power spectrum through entirely standard subsequent evolution.

## **2. THE LambdaCDM STELLAR MASS CEILING AND ITS VIOLATION**

### **2.1 Derivation of the Ceiling from First Principles**

The LambdaCDM stellar mass ceiling at redshift  $z$  is derived from three successive constraints applied to the dark matter halo mass function. The first constraint is the halo mass function itself, which gives the comoving number density of halos above mass  $M_h$  at redshift  $z$ . For the Planck 2018 cosmological parameters, the exponential cutoff of the halo mass function at  $z = 10$  falls at  $M_h \sim 10^{[N]} M_{\text{sun}}$ . Above this mass, the number density drops by an order of magnitude per factor of two in mass. The second constraint is the baryon fraction: the total baryonic mass available for star formation within a halo of mass  $M_h$  is  $f_b M_h$  where  $f_b = \Omega_b/\Omega_m \sim 0.156$ . The third constraint is the maximum star-formation efficiency  $\epsilon_{\text{max}}$ , the fraction of available baryonic mass that can be converted to stars before stellar feedback, reionization, and AGN heating terminate star formation. Observational constraints on  $\epsilon_{\text{max}}$  from  $z = 0$  to  $z = 6$  place  $\epsilon_{\text{max}} \sim 0.20$  for halo masses in the range  $10^{[N]}$  to  $10^{[N]} M_{\text{sun}}$ . This three-factor

product defines the stellar mass ceiling  $M_*^{\text{ceil}}(z) = \epsilon_*^{\text{max}} * f_b * M_h^{\text{max}}(z)$ , where  $M_h^{\text{max}}(z)$  is the mass of the most massive halo expected to exist anywhere within the observable Hubble volume at redshift  $z$ . At  $z = 10$ ,  $M_h^{\text{max}}$  in the observable universe is approximately  $3 \times 10^{10} M_{\text{sun}}$  with a predicted comoving number density of order  $10^{-9} \text{ Mpc}^{-3}$ , giving approximately one such halo in the entire Hubble volume. The ceiling is  $M_*^{\text{ceil}}(z = 10) \sim 0.20 * 0.156 * 3 \times 10^{10} M_{\text{sun}} \sim 9 \times 10^9 M_{\text{sun}}$ . At  $z = 12$ , the ceiling falls to approximately  $2 \times 10^8 M_{\text{sun}}$ . At  $z = 14$ , it falls to approximately  $3 \times 10^7 M_{\text{sun}}$ . These ceilings apply to the single most extreme object in the entire Hubble volume. A survey of any sub-volume should detect zero objects above the ceiling at 3sigma confidence.

## 2.2 The JWST Detections as Quantified Falsifications

JADES-GS-z14-0 at  $z = 14.18$  has a stellar mass of approximately  $10^8 M_{\text{sun}}$  and oxygen enrichment exceeding  $0.1 Z_{\text{sun}}$ . The stellar mass alone exceeds the LambdaCDM ceiling at  $z = 14$  by a factor of approximately three, but the metallicity implies a prior integrated stellar mass converted to metals at least an order of magnitude larger than the instantaneous stellar mass, requiring progenitor stellar populations with total masses of order  $10^{10} M_{\text{sun}}$  at  $z$  greater than 14. This pushes the implied violation of the ceiling to factors of thirty or more. MoM-z14 at  $z = 14.44$  presents super-solar nitrogen-to-carbon ratios, requiring Wolf-Rayet stellar populations that can only produce this abundance pattern after multiple gigayears of stellar evolution at masses above  $8 M_{\text{sun}}$ , yet the universe at  $z = 14.44$  is only 280 megayears old. No standard stellar evolution model produces Wolf-Rayet nitrogen enrichment on a 280 Myr timescale in a system of any mass.

The Xiao et al. (2024) ultra-massive galaxies at  $z = 5-9$  present violations of a different character. Their stellar masses of  $10^N M_{\text{sun}}$  are within the  $\Lambda\text{CDM}$  halo mass budget at those epochs, but the implied star-formation efficiencies of approximately 50 percent are three to five times the observed maximum at any subsequent epoch in any environment. Combined with the Weibel et al. (2025) quenched massive galaxies at  $z = 7.29$  whose number densities exceed IllustrisTNG, EAGLE, and SIMBA predictions by factors of 100 to 1000, the observational picture is consistent: massive structure formed too early, too efficiently, and quenched too rapidly for any hierarchical assembly model to accommodate.

### **2.3 Why the Proposed Remedies Fail**

Three modifications to  $\Lambda\text{CDM}$  have been proposed to address the crisis. A top-heavy initial mass function (IMF) at high redshift would increase the number of massive stars per unit stellar mass and thereby increase both the luminosity per unit stellar mass and the chemical enrichment rate. However, the stellar mass estimates for the JWST detections are derived from spectral energy distribution fitting and, in the case of JADES-GS-z14-0, from dynamical mass measurements. Dynamical masses are IMF-independent. A top-heavy IMF raises the luminosity-to-mass ratio and would therefore require a lower stellar mass to explain a given observed flux, not a higher one, making the crisis worse in the dynamical mass cases rather than better. Reduced feedback has been proposed as a mechanism to allow higher star-formation efficiencies.

The problem is that the efficiency values required, 47 to 52 percent of available baryons converted to stars within the first 300 to 500 megayears, are not a matter of

feedback calibration but of fundamental energetics. The total supernova energy output of a stellar population forming at 50 percent efficiency over 300 Myr is sufficient to unbind the entire baryonic mass of the host halo multiple times over. Reducing feedback enough to allow 50 percent efficiency requires suppressing supernova feedback entirely, which conflicts with direct observations of supernova rates and galactic wind mass-loading factors at every observable redshift. Early dark energy has been proposed as a modification to the expansion history that accelerates structure formation at  $z$  greater than 3 by temporarily increasing the matter density parameter relative to the dark energy density. The best-fit early dark energy models that can raise the stellar mass function at  $z = 10$  enough to reach the observed Xiao et al. values require a fractional dark energy density at  $z = 5000$  of approximately 10 percent, which shifts the acoustic peak positions in the CMB by an amount already constrained by Planck to be absent. Early dark energy models that satisfy Planck constraints cannot raise the stellar mass ceiling by more than a factor of two to three, which is insufficient to explain violations by factors of ten to one hundred.

AGN feedback has been proposed as the mechanism that explains the anomalously hot intracluster medium detected by Zhou et al. (2025) in SPT2349-56 at  $z = 4.3$ . The proposal invokes kinetic-mode AGN jets from at least three radio-loud AGN in the protocluster core, whose energy output is amplified by confinement at elevated ambient gas pressure, depositing thermal energy into the nascent ICM faster than gravitational heating permits. The problem with this remedy is precisely the one identified by the authors themselves in their energy budget analysis: the fiducial parameter values that characterize the AGN population of SPT2349-56 yield a total estimated energy injection of  $(9.1 \pm 1.1) \times 10^{60}$  erg

over 100 Myr, against a required thermal energy change of  $(10.9 \pm 1.2) \times 10^{60}$  erg, implying a thermal coupling efficiency of  $120 \pm 20$  percent. A coupling efficiency exceeding 100 percent is thermodynamically inadmissible under the equations the authors themselves apply: the total injected kinetic energy cannot be converted to thermal energy at an efficiency greater than unity without violating energy conservation. To reproduce the observed Sunyaev-Zeldovich decrement, at least one fiducial parameter must be revised beyond its observationally motivated range — the AGN active time must exceed 120 Myr, the radiative-mode coupling efficiency must exceed 1 percent rather than 0.5 percent, the ICM fraction must exceed 6 percent rather than 2 percent, or the halo mass must exceed  $1.7 \times 10^{13}$  solar masses rather than the observed  $0.9 \times 10^{13}$  solar masses.

Each of these revisions introduces a fine-tuning problem of the same character as the remedies addressed above: the parameter adjustment required to match the observation conflicts with independent observational constraints on that parameter. Furthermore, because SPT2349-56 is unique — selected as the brightest candidate from a 2,500 square-degree survey — the AGN-feedback remedy must explain why this system hosts an AGN population capable of supplying energy at the required rate and coupling efficiency in a configuration that no existing simulation predicted. A mechanism that requires parameter values outside their observationally motivated ranges to explain the single brightest outlier in a 2,500 square-degree survey does not constitute a general solution to the ICM thermal energy crisis; it constitutes post-hoc parameter adjustment in the specific case where the crisis is most extreme.

## **2.4 The ICM Thermal Energy Crisis: A Thermodynamic Falsification**

The LCDM stellar mass crisis examined in Sections 2.1 through 2.3 concerns the formation of stellar structure. The intracluster medium thermal energy crisis identified by Zhou et al. (2025) constitutes an independent falsification of LCDM operating on a different physical observable — thermodynamic energy content — at the same problematic epoch.

The quantitative statement of the falsification is as follows. The protocluster SPT2349-56 at  $z = 4.3$  contains a dark matter halo with mass  $M_{200} = (9 \pm 5) \times 10^{12}$  solar masses, a molecular gas reservoir of  $4.9 \times 10^{11}$  solar masses, and a stellar mass of  $6.3 \times 10^{11}$  solar masses, for a total baryonic budget consistent with the LCDM halo mass function at this epoch. There is no stellar mass crisis in this system in the sense defined in Section 2.1: the system's stellar and gas masses lie within the LCDM ceiling. The crisis is entirely thermodynamic. The observed thermal Sunyaev-Zeldovich signal — detected at 8.4 sigma in the image plane and 10.4 sigma in Fourier space, with an integrated Compton-Y parameter of  $(2.0 \pm 0.2) \times 10^{-6}$  arcmin<sup>2</sup> and a mean Compton-y of  $(5.6 \pm 0.8) \times 10^{-6}$  within the full signal radius — implies a total thermal energy of  $(11.8 \pm 1.2) \times 10^{60}$  erg. The thermal energy of an ICM in virial equilibrium within a  $9 \times 10^{12}$  solar-mass halo, assuming an ICM fraction of 6 percent — the maximum observationally permitted at this epoch — is  $(2.6 \pm 1.0) \times 10^{60}$  erg: approximately 22 percent of the observed value. Even under the most generous gravitational heating assumptions, the expected Sunyaev-Zeldovich signal is at most 30 microjansky, compared to the observed  $157 \pm 16$  microjansky. The observed signal exceeds the gravitational expectation by a factor of at least five regardless of what assumption is made about the ICM fraction within its observed range.

The LCDM halo mass function predicts that the SPT2349-56 halo, at the mass estimated from velocity dispersion, is not anomalous: there is no difficulty placing a  $10^{13}$  solar-mass halo at  $z = 4.3$ . The crisis is not that the halo is too massive to exist. The crisis is that the thermal energy content of its gas is five to ten times larger than gravitational collapse can produce, and this excess places the system 6.4 sigma above the predictions of the TNG-Cluster simulation suite — simulations that incorporate AGN feedback, stellar feedback, and the full range of non-gravitational heating mechanisms available within the LCDM framework. TNG-Cluster does not predict this signal. IllustrisTNG does not predict it. No state-of-the-art hydrodynamical simulation with a calibrated galaxy formation model predicted a system with SPT2349-56's thermal energy content at  $z = 4.3$ , even though the system's halo mass is unremarkable by LCDM standards.

The structural parallel with the stellar mass crisis is exact. In the stellar mass crisis, the objects are not anomalous in number — the LCDM halo mass function permits their halos to exist — but they exceed the maximum stellar mass producible within those halos by factors of ten to one hundred. In the ICM thermal energy crisis, the object is not anomalous in its halo mass, but the thermal energy content of its ICM exceeds the maximum producible by gravitational collapse within that halo by a factor of five to ten. In both cases, the falsification is not of the large-scale matter distribution but of the physical processes that LCDM invokes to explain what that matter does. The assembly time problem and the thermalization efficiency problem are the same problem viewed through different observables: LCDM's initial condition architecture delivers insufficient energy, at insufficient rates, in insufficient concentrations to produce either the observed stellar

masses at  $z > 10$  or the observed ICM thermal energies at  $z = 4.3$ . Both crises are resolved by the same change to the foundational assumption: the initial conditions were not small-amplitude quantum fluctuations but the structured thermodynamic legacy of a succession of collisions that deposited mass, angular momentum, and thermal energy into specific geometric configurations before recombination.

The significance of the Zhou et al. (2025) detection for the present paper is not that it adds a new problem to the list of LCDM anomalies. It is that it provides an independent, thermodynamically distinct falsification that is predicted and explained by SCT without any additional parameters beyond those introduced in Section 3. The born-hot ICM mechanism described in Section 3.7 is not a post-hoc addition to SCT; it is a direct consequence of the collision-seeding mechanism applied to the same physical processes that produce the proto-structure masses and angular momenta derived in Sections 3.2 and 3.3. The observation existed before the prediction was written down, but the prediction follows from the mechanism with no degrees of freedom beyond those already fixed by the existing framework.

### **3. THE SCT MECHANISM FOR EARLY STRUCTURE FORMATION**

Notation Summary: The following symbols are used throughout Section 3 and the falsifiable predictions of Section 4. All symbols are defined at their first use in the text, and this summary is provided for reference.  $M_1, M_2$ : masses of the two colliding comoving pockets.  $\mu = M_1 M_2 / (M_1 + M_2)$ : reduced mass.  $b$ : impact parameter.  $R_1, R_2$ : radii of the two colliding pockets;  $R_{\min} = \min(R_1, R_2)$ .  $v_{\text{rec}}$ : closing velocity ( $v_{\text{rec}} > c$  for

superluminal collisions).  $f_b = \Omega_b/\Omega_m \sim 0.156$ : universal baryon fraction.  $\Omega(b, R_1, R_2)$ : dimensionless geometric overlap factor (equation 6).  $\alpha_{th}$ : thermalization efficiency (range 0.25 to 0.85).  $M_{overlap}$ : total baryonic mass in the collision overlap volume.  $M_{proto}$ : thermalized proto-structure mass (equation 2).  $\alpha_{RH}$ : Rankine-Hugoniot thermalization efficiency.  $f_{ret}$ : gravitational retention fraction.  $J$ : angular momentum vector deposited into the collision overlap volume.  $j/j_{circ}$ : specific angular momentum ratio determining morphological type.  $A(N, \sigma_v, R)$ : gravitational superposition amplification factor (equation 13).  $N$ : number of mass elements within a proto-galactic structure.  $\eta$ : coherence parameter.  $K(r, z)$ : gravitational coherence kernel (equation 10).  $\lambda_c(z)$ : coherence length.  $\sigma_v$ : velocity dispersion.  $v_{circ}$ : circular velocity at the virial radius.  $t_{ff}, t_{ff_{eff}}$ : standard and effective free-fall timescales.  $\Phi(k)$ : total curvature perturbation field.  $N_{coll}$ : total number of independent collision events.  $\xi(k)$ : collision spectral rate.  $L$ : number of hierarchical nesting levels ( $\sim 29$ ).  $n_s$ : primordial scalar spectral index ( $\sim 0.965$ ).  $A_s$ : scalar amplitude.  $f_{NL}$ : non-Gaussianity parameter ( $|\eta_{NL}| \sim 10^N$ ).  $r$ : tensor-to-scalar ratio ( $< 10^{-5}$ ).  $M_J$ : Jeans mass (equation 32).  $M_{BH}$ : baryonic mass of head-on collision remnant (equation 33).  $M_{seed}$ : black hole seed mass (equation 34).  $f_{BH}$ : collapse fraction (0.5-1.0).  $n_0, \alpha, \beta_{ev}$ : collision mass function parameters.  $M_{ref} = 10^N M_{sun}$ : reference mass.  $f_*$ : stellar fraction (fiducial 0.3).  $n_{SCT}$ : SCT comoving number density (equation 41).  $f_{iso}$ : isocurvature fraction (equation 45).

### 3.1 The Architectural Difference

The LambdaCDM assembly crisis identified in Section 2 is the consequence of a specific architectural assumption: that the initial conditions delivered to the post-recombination universe by inflation consist of small-amplitude, nearly Gaussian density perturbations with no preferred scale, no preferred direction, and no pre-existing large-scale angular momentum. Under this assumption, every structure that exists today must have been assembled, piece by piece, from smaller precursors in the time available since recombination. The assembly time is the bottleneck, and the bottleneck is absolute. Successive Collision Theory removes this bottleneck by changing the assumption that creates it. In SCT, the initial conditions delivered to the post-recombination universe are not small-amplitude quantum fluctuations. They are the structured thermodynamic legacy of a succession of superluminal collisions between nested comoving frames that concluded before recombination, as established in Section 1.4 through the three independent observational anchors of primordial nucleosynthesis abundances, CMB spectral purity, and acoustic peak positions. The collision cascade deposited mass, angular momentum, and thermal energy into specific geometric configurations within the pre-recombination plasma. Those configurations are proto-structures: overdense, rotating, thermally stratified regions whose characteristic masses, angular momenta, and density profiles are set by the collision geometry at the moment of deposition, not by the subsequent gravitational growth rate. The assembly time problem does not arise because assembly in the hierarchical sense is not the mechanism. The proto-structures are seeded. They collapse. They do not wait.

### **3.2 The Characteristic Mass Scale from Collision Geometry**

The total baryonic mass deposited in the collision overlap volume is:

$$M_{\text{overlap}} = f_b * \Omega(b, R_1, R_2) * (M_1 + M_2) \dots (1)$$

where  $f_b$  is the universal baryon fraction and  $\Omega(b, R_1, R_2)$  is the dimensionless geometric overlap fraction. The fraction of  $M_{\text{overlap}}$  that is gravitationally retained as a proto-structure is:  $M_{\text{proto}} = \alpha_{\text{th}} * f_{\text{ret}} * M_{\text{overlap}} = \alpha_{\text{th}} * f_b * \mu * \Omega(b, R_1, R_2)$

$$\dots (2)$$

where  $\mu = M_1 * M_2 / (M_1 + M_2)$  is the reduced mass,  $\alpha_{\text{th}}$  is the thermalization efficiency from the Rankine-Hugoniot conditions, and  $f_{\text{ret}}$  has been absorbed into  $\alpha_{\text{th}}$  for compactness. The thermalization efficiency  $\alpha_{\text{th}}$  for a strong shock in a radiation-dominated fluid is derived from the Rankine-Hugoniot jump conditions:

$$\alpha_{\text{RH}} = 2\gamma(\gamma - 1)M_s^2 / [(\gamma + 1)^2 M_s^2 - 2(\gamma - 1)] \dots$$

$$(3)$$

where  $M_s$  is the Mach number and  $\gamma = 4/3$  for a radiation-dominated gas. In the strong-shock limit  $M_s \gg 1$ :

$$\alpha_{\text{th}} = f_{\text{ret}} * \alpha_{\text{RH}} = \alpha_{\text{th}}(M_s, \gamma, \eta) \dots (4)$$

Numerically,  $\alpha_{\text{th}}$  ranges from 0.25 to 0.85 across the physically plausible range of collision parameters, with a fiducial value of 0.40. The retention fraction  $f_{\text{ret}}$  accounts for the fraction of shocked material that escapes the potential well of the collision remnant rather than being gravitationally captured.

$$f_{\text{ret}} = 1 - \exp(-\eta^2 v_{\text{esc}}^2 / \sigma_v^2) \dots (5)$$

The geometric overlap factor  $\Omega(b, R_1, R_2)$  for the intersection of two spheres of radii  $R_1$  and  $R_2$  with center separation  $b$  is:

$$\Omega = (1/2) [1 - (3b/4R_{\min}) + (b^2/16R_{\min}^2)] \text{ for } b \leq 2R_{\min} \dots$$

(6)

$$\Omega = 0 \text{ for } b > 2R_{\min} \dots (7)$$

For a head-on collision  $b = 0$ :  $\Omega = 1/2$ , and  $M_{\text{proto}} = \alpha_{\text{th}} * f_b * \mu/2$ . For a grazing collision  $b = R_{\min}$ :  $\Omega$  decreases smoothly as computed from equation (6). The full range of impact parameters  $b$  from 0 to  $2R_{\min}$  produces a continuous distribution of proto-structure masses from  $M_{\text{proto}}^{\text{max}}$  at  $b = 0$  to  $M_{\text{proto}} = 0$  at  $b = 2R_{\min}$ . The remaining free parameters and their physically motivated ranges are:  $M_1$  and  $M_2$  in the range  $10^{\text{[N]}}$  to  $10^{\text{[N]}} M_{\text{sun}}$  from the virial mass ladder;  $b/R_{\min}$  in the range 0 to 1.8 from the uniform impact parameter distribution;  $\alpha_{\text{th}}$  in the range 0.25 to 0.85 from the Rankine-Hugoniot analysis.

### 3.3 Angular Momentum Inheritance and Morphological Type

The angular momentum deposited into the collision debris is the second mathematical pillar of the SCT early structure mechanism. For a two-frame collision with reduced mass  $\mu$ , impact parameter vector  $b$ , and relative velocity vector  $v_{\text{rec}}$ , the total angular momentum deposited into the overlap volume is:

$$J = \mu (b \times v_{\text{rec}}) \dots (8)$$

This equation follows directly from the definition of orbital angular momentum for a two-body system,  $L = r \times p$ , applied to the collision geometry where the relative position

vector at first contact is  $b$  and the relative momentum is  $\mu v_{\text{rec}}$ . The angular momentum  $J$  is deposited into the proto-structure at the moment of collision and is preserved through the thermalization process by exact angular momentum conservation. Formally, the post-shock stress-energy tensor within the overlap volume obeys the GR conservation law for angular momentum, which in the absence of external torques requires that the integrated angular momentum of the system be invariant. The ratio  $j/j_{\text{circ}}$ , specific angular momentum relative to circular orbit angular momentum at the proto-galaxy's eventual virial radius, determines morphological type. Grazing collisions with large  $b$  produce large  $J$ , high  $j/j_{\text{circ}}$ , and disk-dominated morphologies. Head-on collisions with  $b \sim 0$  produce small  $J$ , low  $j/j_{\text{circ}}$ , and pressure-supported elliptical configurations. This resolves the morphological maturity problem: the barred spiral at  $z \sim 3$  (Costantin et al. 2023), the grand-design spiral at  $z = 4.03$  (Jain and Wadadekar 2025), and the approximately constant spiral-to-elliptical ratio to  $z \sim 10$  (Ferreira et al. 2024) are the direct observational signature of angular momentum inheritance from collision geometry, locked in before recombination and preserved by exact conservation through all subsequent evolution.

### **3.4 Gravitational Superposition Amplification**

The gravitational superposition amplification factor  $A(N, \sigma_v, R)$  quantifies the enhanced rate of gravitational collapse for a coherently rotating proto-galactic structure containing  $N$  mass elements with velocity dispersion  $\sigma_v$  within a region of physical radius  $R$ . When the velocity dispersion is small compared to the bulk rotation velocity, the

mass elements contribute constructively to the net gravitational field rather than averaging to zero as in a pressure-supported system. The superposition factor is:

$$A(N, \sigma_v, R) = S * K(r, z) / (\eta^2 + \epsilon^2) \dots (9)$$

where  $S$  is the superposition enhancement at peak coherence,  $K(r, z)$  is the coherence kernel:

$$K(r, z) = \exp(-r/\lambda_c(z)) \dots (10)$$

and  $\eta$  is the coherence parameter derived from the ratio  $\sigma_v/v_{\text{circ}}$ :

$$\eta = 1 - (\sigma_v/v_{\text{circ}})^2 \dots (11)$$

The coherence length  $\lambda_c(z)$  is the scale below which the velocity dispersion of mass elements is small compared to their bulk orbital velocity, given by:

$$\lambda_c(z) = v_{\text{circ}} / H(z) \dots (12)$$

The effective free-fall timescale with gravitational superposition amplification is:

$$t_{\text{ff\_eff}} = t_{\text{ff}} / A(N, \sigma_v, R) \dots (13)$$

where  $t_{\text{ff}} = (3\sigma/32G\sigma)^{1/2}$  is the standard baryonic free-fall timescale. The amplification factor  $A$  exceeds unity whenever  $\sigma_v/v_{\text{circ}} < 1$ , which is the condition that the proto-galactic structure is rotationally supported rather than pressure-supported. For typical SCT collision products with  $j/j_{\text{circ}} \sim 0.5$  to  $0.9$ ,  $A$  ranges from 2 to 10, reducing the effective free-fall timescale by the same factor and allowing gravitational collapse to complete at  $z = 10$  to  $15$  even for masses where the standard Jeans analysis would predict stability.

### 3.5 The Primordial Power Spectrum from Collision Statistics

The primordial curvature perturbation field from a population of  $N_{\text{coll}}$  independent collision events is:

$$\Phi(k) = \sum_i A_i \exp(ikx_i) W(k/k_i) \dots (14)$$

where  $A_i$  is the curvature amplitude of event  $i$ ,  $x_i$  is its position,  $k_i = 2\sigma/R_i$  is the collision wavenumber, and  $W(k/k_i)$  is a window function determined by the collision geometry. The power spectrum is:

$$P(k) = \langle |\Phi(k)|^2 \rangle = A_s \left( \frac{k}{k_*} \right)^{n_s - 1} \dots (15)$$

By the Central Limit Theorem applied to the  $N_{\text{coll}} \sim 10^4$  independent collision events, the resulting perturbation field is Gaussian with power spectrum given by the incoherent sum of individual collision spectra. The collision spectral rate  $\xi(k)$  gives the mean number of collision events per unit comoving volume per unit log-wavenumber:

$$\xi(k) = \xi_0 \left( \frac{k}{k_*} \right)^{\alpha - 1} \dots (16)$$

where  $\alpha$  is the scale distribution index. The collision spectrum therefore takes the form:  $P(k) \sim k^{\alpha - 1} [\text{contribution per event}] = A_s \left( \frac{k}{k_*} \right)^{n_s - 1}$

$$\dots (17)$$

giving the identification  $n_s = \alpha$ . The value of  $\alpha$  is determined by the finite dynamic range of the collision cascade. If the cascade spans  $L$  hierarchical levels with mass ratio  $r_L$  between adjacent levels, the total dynamic range is  $r_L^L$ . The scale distribution index is:

$$\alpha - 1 = -1/L \dots (18)$$

from the boundary correction to the power-law index arising from the finite-range effect of a truncated geometric series. For  $L = 29$  nesting levels:

$$a_{\text{alpha}} = n_s = 1 - 1/29 \sim 1 - 0.034 \sim 0.966 \dots \quad (19)$$

This is consistent with the Planck 2018 best-fit value  $n_s = 0.9649 \pm 0.0042$ . The scalar amplitude  $A_s$  is determined by normalization to the observed CMB temperature variance:

$$A_s = 2.1 \times 10^{-9} \dots \quad (20)$$

The non-Gaussianity parameter  $f_{\text{NL}}$  is suppressed by the CLT:

$$|f_{\text{NL}}| \sim N_{\text{coll}}^{-1/2} \sim 10^{-5} \dots \quad (21)$$

well below current observational limits. The tensor perturbations from the collision cascade are suppressed because the stress-energy tensor of each thermalized collision remnant is dominated by its isotropic components. The anisotropic stress per event is of order  $(\sigma_v/c)^2$  times the isotropic pressure. For a thermalized remnant,  $\sigma_v/c \ll 1$ , giving a single-event tensor-to-scalar ratio:

$$r_{\text{single}} \sim (\sigma_v/c)^2 \dots \quad (22)$$

After incoherent summation over  $N_{\text{coll}}$  events, the total tensor-to-scalar ratio is:

$$r = r_{\text{single}} / N_{\text{coll}}^{1/2} < 10^{-5} \dots \quad (23)$$

for all physically plausible parameter combinations. The additional derivations producing equations (24) through (31) relating the collision cascade statistics to higher-order CMB statistics are given in the supplementary treatment Nipok (2026d).

### 3.6 Direct-Collapse Black Hole Seeds from Head-On Collision Geometry

Head-on collision geometry ( $b \sim 0$ ) produces a maximally dense, minimally rotating remnant. The absence of angular momentum removes the centrifugal barrier that would otherwise support the remnant against gravitational collapse, and the high post-shock temperature produces a Jeans mass larger than the remnant mass, preventing fragmentation into stellar-mass objects. These conditions are precisely those required for direct collapse to a supermassive black hole seed. The Jeans mass for a radiation-dominated plasma at temperature  $T$  and particle number density  $n$  is:

$$M_J = (\sigma/6)^{3/2} c_s^3 / (G^{3/2} \sigma^{1/2}) \dots (32)$$

where  $c_s$  is the adiabatic sound speed and  $\sigma = n m_p$  is the mass density. For post-shock temperatures of order  $10^8$  K,  $M_J$  is of order  $10^6$  to  $10^7 M_{\text{sun}}$ . Because  $M_{\text{proto}}$  is comparable to or smaller than  $M_J$  at the post-shock temperature, the remnant cannot fragment. It cools as a single coherent body and collapses to a singularity before cooling enough to fragment. The seed mass is obtained from equation (2) in the head-on limit ( $b = 0, \Omega = 1$ ):

$$M_{\text{BH}} = \alpha_{\text{th}} * f_b * \mu \dots (33)$$

The black hole seed mass after applying the collapse fraction is:

$$M_{\text{seed}} = f_{\text{BH}} * \alpha_{\text{th}} * f_b * \mu \dots (34)$$

with  $f_{\text{BH}}$  in the range 0.5 to 1.0. Numerical estimates across the physical parameter range:

$$\text{For } M_1 = M_2 = 10^6 M_{\text{sun}}: M_{\text{seed}} \sim 2.2 \times 10^6 M_{\text{sun}} \dots (35)$$

$$\text{For } M_1 = M_2 = 10^9 M_{\text{sun}}: M_{\text{seed}} \sim 2.2 \times 10^9 M_{\text{sun}} \dots (36)$$

$$\text{For } M_1 = M_2 = 10^N M_{\text{sun}}: M_{\text{seed}} \sim 2.2 \times 10^7 M_{\text{sun}} \dots (37)$$

$$\text{For } M_1 = 10^N M_{\text{sun}}, M_2 = 10^5 M_{\text{sun}}: M_{\text{seed}} \sim 4.4 \times 10^5 M_{\text{sun}} \dots (38)$$

These ranges bracket the observed overmassive black holes at  $z > 7$ , including UHZ1 ( $\sim 4 \times 10^7 M_{\text{sun}}$ ) and QSO J0313-1806 ( $1.6 \times 10^9 M_{\text{sun}}$ ). The observed BH-to-stellar mass ratio at any epoch is:

$$M_{\text{BH}}/M_*(z) \sim M_{\text{seed}} / [\epsilon_{\text{SF}}(z) * M_{\text{proto}}^{\text{grazing}}(z)] \dots (39)$$

Because  $M_{\text{seed}}$  and  $M_{\text{proto}}$  are seeded at the same epoch by the same collision population but from different geometric configurations (head-on vs. grazing), their ratio reflects the distribution of impact parameters rather than any late-time evolutionary process. The elevated BH-to-stellar mass ratios at high redshift are therefore a geometric imprint of the collision cascade, not a consequence of runaway accretion or missing feedback.

### **3.7 The Born-Hot Intracluster Medium as a Thermodynamic Prediction of Collision Seeding**

The direct-collapse black hole formation mechanism of Section 3.6 applies to head-on collision geometry ( $b = 0$ ), where minimal angular momentum permits direct gravitational collapse of the thermalized remnant. Grazing collision geometry ( $b > 0$ ) produces protogalactic disks and ellipticals as derived in Section 3.3. Between these extremes lies a third class of collision outcome: collisions of intermediate geometry between extended proto-cluster mass distributions that produce a spatially extended, thermalized overlap volume

whose mass exceeds the Jeans mass for stellar fragmentation but whose angular momentum is insufficient for disk formation and whose cooling time exceeds the Hubble time at the seeding epoch. These remnants do not collapse to a single object. They do not form disks. They persist as extended, hot, over-pressurized thermal reservoirs — the direct progenitors of the intracluster medium in forming galaxy clusters. This section derives the thermodynamic properties of these collision-seeded proto-ICM structures and demonstrates that their predicted thermal energy content is consistent with the observation of Zhou et al. (2025).

### **3.7.1 The Collision-Seeded Proto-ICM**

For a collision between two pockets of mass  $M_1$  and  $M_2$  at impact parameter  $b$  in the range  $R_{\min} < b < 2R_{\min}$ , the overlap geometry produces a proto-structure of mass  $M_{\text{proto}}$  from equation (2) with angular momentum  $J = \mu(b \times v_{\text{rec}})$  from equation (8). The ratio  $j/j_{\text{circ}}$  for these intermediate- $b$  collisions falls in the range 0.1 to 0.5 — insufficient to support a rotationally stabilized disk but sufficient to prevent immediate head-on collapse. The post-shock temperature of the thermalized overlap volume is:

$$T_{\text{proto}} = (2 \mu_{\text{mol}} m_{\text{p}} / 3 k_{\text{B}}) * (v_{\text{rec}}^2 / \xi) \dots (46)$$

where  $\mu_{\text{mol}} \sim 0.6$  is the mean molecular weight,  $m_{\text{p}}$  is the proton mass,  $k_{\text{B}}$  is Boltzmann's constant, and  $\xi$  is the thermalization fraction ( $\xi = \alpha_{\text{th}} * f_{\text{ret}}$  from equation 4). For  $v_{\text{rec}} > c$  in the superluminal collision regime, the post-shock temperature  $T_{\text{proto}}$  is of order  $10^8$  to  $10^{10}$  K at the moment of thermalization, falling within the range of present-

day intracluster medium temperatures ( $10^7$  to  $10^8$  K) after adiabatic expansion and radiation losses over the intervening 10 to 13 billion years.

The cooling time of this proto-ICM structure is:

$$t_{\text{cool}} = (3/2) * (n_e + n_i) * k_B * T_{\text{proto}} / (n_e * n_H * \Lambda(T_{\text{proto}})) \dots$$

(47)

where  $n_e$  and  $n_i$  are the electron and ion number densities,  $n_H$  is the hydrogen number density, and  $\Lambda(T_{\text{proto}})$  is the cooling function. For  $T_{\text{proto}} > 10^7$  K, the dominant cooling mechanism is thermal bremsstrahlung, giving  $\Lambda(T_{\text{proto}}) \sim 2.4 \times 10^{-27} T_{\text{proto}}^{(1/2)}$  erg cm<sup>3</sup> s<sup>-1</sup>. The proto-ICM density at the seeding epoch (before recombination, at redshift  $z_{\text{seed}} \gg 1000$ ) is:

$$\rho_{\text{proto}} = M_{\text{proto}} / (4/3 * \pi * R_{\text{proto}}^3) \dots$$

(48)

where  $R_{\text{proto}}$  is the physical radius of the collision overlap volume. At  $z_{\text{seed}} \sim 10^4$  to  $10^5$ , the baryon density of the universe is sufficiently high that  $t_{\text{cool}} \gg t_{\text{Hubble}}$  for structures with  $M_{\text{proto}} > 10^{12}$  solar masses, ensuring that these proto-ICM structures remain hot and thermally supported from their formation epoch through recombination, through the epoch of reionization, and into the epoch at which they become observable as protocluster ICM.

### **3.7.2 Predicted Thermal Energy at $z \sim 4$**

The thermal energy of a proto-ICM structure seeded by the collision cascade and preserved to redshift  $z$  is:

$$E_{\text{therm}}(z) = (3/2) * M_{\text{proto}} * k_B * T_{\text{eff}}(z) / (\mu_{\text{mol}} * m_p) \dots$$

(49)

where  $T_{\text{eff}}(z)$  is the effective temperature after adiabatic evolution from the seeding epoch  $z_{\text{seed}}$  to the observation epoch  $z$ :

$$T_{\text{eff}}(z) = T_{\text{proto}} * (1 + z)^2 / (1 + z_{\text{seed}})^2 * f_{\text{therm,evo}} \dots (50)$$

where  $f_{\text{therm,evo}}$  is the fraction of initial thermal energy retained after radiation losses, gravitational work, and AGN heating during the intervening evolution. For a proto-ICM with  $M_{\text{proto}}$  in the range  $10^{12}$  to  $10^{13}$  solar masses and initial temperature  $T_{\text{proto}} > 10^8$  K,  $f_{\text{therm,evo}}$  is bounded from below by the observed ratio  $E_{\text{therm,obs}} / E_{\text{therm,vir}}$ , where  $E_{\text{therm,vir}}$  is the thermal energy expected from virialization alone. The Zhou et al. (2025) measurement gives  $E_{\text{therm,obs}} = (11.8 \pm 1.2) \times 10^{60}$  erg for SPT2349-56, while  $E_{\text{therm,vir}} \sim 2.6 \times 10^{60}$  erg (assuming  $f_{\text{ICM}} = 0.06$ ), yielding:

$$f_{\text{therm,evo}} * (T_{\text{proto}} / T_{\text{vir}}) \sim E_{\text{therm,obs}} / E_{\text{therm,vir}} \sim 4.5 \dots (51)$$

where  $T_{\text{vir}}$  is the virial temperature of the observed halo. Equation (51) states that the ratio of the observed thermal energy to the virially expected thermal energy is approximately 4.5, consistent with the product of an initial temperature exceeding  $T_{\text{vir}}$  by a factor of order 10 and a thermal retention fraction of order 0.5. These parameter values are physically well-motivated: collision-seeded proto-structures are initialized at temperatures substantially above virial equilibrium because their thermal energy is deposited by the kinetic energy of the superluminal collision rather than by gravitational collapse, and a retention fraction of 0.5 over 9 billion years of evolution is consistent with observed cooling rates in cluster progenitors.

### **3.7.3 The Born-Hot Prediction**

The SCT prediction for the ICM thermal energy at  $z = 4.3$  in a system with the observed halo mass and angular momentum of SPT2349-56 follows directly from equations (2), (46), and (49) with no additional free parameters. The three input quantities — the proto-structure mass  $M_{\text{proto}}$  from the collision mass function, the initial temperature  $T_{\text{proto}}$  from the collision velocity and thermalization efficiency, and the angular momentum  $J$  from the impact parameter — are all determined by the same collision geometry that sets the proto-galactic morphology. There is no separate mechanism for the born-hot ICM. It is the same collision event, viewed through the thermodynamic rather than the structural lens.

The SCT prediction is therefore not merely that some proto-ICM structures are hot. It is that the thermal energy of the proto-ICM is correlated with the same collision parameters that set proto-galactic masses and morphologies. Systems that formed from high-velocity, intermediate-impact-parameter collisions of large parent pockets should simultaneously exhibit elevated stellar mass, morphological regularity, and anomalously high ICM thermal energy at high redshift. The SPT2349-56 protocluster, with its 30 spectroscopically confirmed DSFGs, multiple radio-loud AGN, and thermal energy 6.4 sigma above simulation predictions, is precisely the configuration that the SCT collision-seeding mechanism predicts for a system formed from a high-velocity, moderately grazing collision between two massive parent pockets at  $z > 10^4$ .

The falsifiable prediction that follows is stated explicitly in Section 4.1. Systems selected for anomalously high ICM thermal energy at  $z > 3$  should also exhibit elevated disk fractions, overmassive stellar components, and proto-structure masses consistent with the

SCT collision mass function rather than the LCDM halo mass function. The converse prediction is equally falsifiable: if ICM thermal energy at high redshift correlates with the properties predicted by gravitational heating alone, with no excess above the TNG-Cluster prediction, then the collision-seeding mechanism would be observationally refuted for that class of systems.

## 4. FALSIFIABLE PREDICTIONS

### 4.1 Tier One: Predictions Testable with Existing JWST Data

Prediction 1: Power-Law Rather than Exponential Decline Above  $z = 14$  Observable. The comoving number density of spectroscopically confirmed galaxies with  $M_{\text{sun}} > 10^8 M_{\text{sun}}$  at  $z > 14$  in existing and forthcoming JWST deep field programs.

SCT prediction. Because proto-structure masses are set by collision geometry through equation (2) rather than by the gravitational growth rate, the characteristic mass does not decline exponentially with increasing redshift. The SCT prediction is that the comoving number density of massive galaxies follows a power-law decline with redshift. Confirmed galaxies with  $M_{\text{sun}} > 10^8 M_{\text{sun}}$  should continue to appear at  $z > 14$  in JWST deep fields at a rate inconsistent with LambdaCDM exponential suppression but consistent with a power-law extrapolation of the  $z = 10-12$  number density. LambdaCDM prediction. The halo mass function imposes an exponential suppression. The number density of galaxies with  $M_{\text{sun}} > 10^N M_{\text{sun}}$  should fall by roughly two orders of magnitude between

$z = 7$  and  $z = 10$ , and be effectively zero above  $z = 12$  within any JWST-accessible survey volume.

Falsification threshold. If future JWST spectroscopic programs targeting  $z > 14$  sources return a number density of massive galaxies more than  $2\sigma$  below the SCT power-law extrapolation and consistent with the  $\Lambda$ CDM exponential suppression, the collision-seeding mechanism as the primary source of early massive structure would be observationally refuted.

Timeline. Testable with existing JWST Cycle 1 through Cycle 3 spectroscopic data and planned NIRSpect multi-object programs scheduled for completion by 2027.

Prediction 2: Disk Fraction at  $z > 10$  Exceeds Ten Percent Observable. The fraction of confirmed  $z > 10$  galaxies in existing JWST NIRCам imaging that exhibit disk-dominated or rotationally supported morphologies, as determined by Sersic index fitting, kinematic misalignment, or axis ratio analysis.

SCT prediction. Through the angular momentum inheritance mechanism, morphological type is set by  $J = \mu(b \times v_{\text{rec}})$  at the collision seeding epoch. The fraction of galaxies at  $z > 10$  exhibiting disk-dominated morphologies should be comparable to the fraction observed at  $z = 3-5$ .  $\Lambda$ CDM prediction. Morphological regularity should decline sharply with increasing redshift as the merger rate increases and dynamically cold disks become progressively harder to maintain.

Falsification threshold. A disk fraction below ten percent at  $z > 10$ , consistent with  $\Lambda$ CDM's merger-dominated expectation, would constitute significant evidence against SCT's angular momentum inheritance mechanism.

Timeline. Testable with existing JWST imaging data currently in hand. A definitive morphological census of the confirmed  $z > 10$  population is achievable within one to two years.

### Prediction 3: ICM Thermal Energy Excess in High-Redshift Protoclusters

Observable. The ratio of observed ICM thermal energy to the virial expectation —  $E_{\text{therm,obs}} / E_{\text{therm,vir}}$  — in protoclusters at  $z > 3$  selected by their thermal Sunyaev-Zeldovich signal, measured by ALMA, the Atacama Cosmology Telescope, or the South Pole Telescope.

SCT prediction. Protoclusters formed from high-velocity, intermediate-impact-parameter collisions between massive parent pockets should exhibit ICM thermal energies systematically above the virial expectation. The SCT prediction, from equations (46) through (51) in Section 3.7, is that the thermal energy of the born-hot ICM tracks the collision kinetic energy rather than the gravitational binding energy, yielding  $E_{\text{therm,obs}} / E_{\text{therm,vir}} > 1$  for a fraction of protoclusters at  $z > 3$  set by the distribution of collision velocities and impact parameters in the SCT mass function. For the most massive systems — those corresponding to  $M_{\text{proto}} > 10^{13}$  solar masses from equation (2) — the predicted excess is a factor of three to ten above the virial expectation, consistent with the 6.4-sigma excess observed by Zhou et al. (2025) in SPT2349-56 at  $z = 4.3$ . Furthermore, SCT predicts that ICM thermal energy excess should be correlated with elevated stellar mass, high disk fraction, and overmassive central black holes, because all four observables trace the same underlying collision geometry.

Lambda-CDM prediction. The TNG-Cluster simulation suite predicts that the mass-normalized Compton-Y parameter falls below the self-similar expectation at  $z > 3$ , because intracluster gas in younger clusters is still accumulating and the hot-gas fraction is insufficient to produce strong tSZ signals. No existing LCDM simulation predicts a Compton-Y parameter more than two times the universal scaling relation at  $z > 4$ , let alone five times as observed in SPT2349-56.

Falsification threshold. A survey of ten or more protoclusters at  $z > 3$  with tSZ detections showing that the thermal energy excess  $E_{\text{therm,obs}} / E_{\text{therm,vir}}$  follows the TNG-Cluster median prediction — with no systems exceeding three times the virial expectation — would constitute significant evidence against the born-hot ICM mechanism of Section 3.7.

Timeline. Testable with existing ALMA archival data and forthcoming observations from the South Pole Telescope SPT-3G and Atacama Cosmology Telescope surveys. Systematic tSZ surveys of known  $z > 3$  protoclusters are achievable within the current decade.

## **4.2 Tier Two: Predictions Testable with the Nancy Grace Roman Space**

### **Telescope**

Survey Parameters The predictions in this section assume the Nancy Grace Roman Space Telescope High Latitude Wide Area Survey: a total survey area of 2000 square degrees, a 5-sigma point-source depth of approximately 27.5 AB magnitude in the F106 and F158 bandpasses, and a photometric redshift completeness of 70 percent for sources with

$M_{\text{sun}} > 10^{[N]} M_{\text{sun}}$  at  $z = 12-15$ . The survey volume in the redshift interval  $z = 12-13$  over 2000 square degrees is approximately  $1.4 \times 10^9$  comoving  $\text{Mpc}^3$ . The corresponding volumes for  $z = 13-14$  and  $z = 14-15$  are approximately  $1.2 \times 10^9$  and  $1.0 \times 10^9$  comoving  $\text{Mpc}^3$ , computed from the standard comoving volume element  $dV/dz = c D_M^2/[H(z)]$  integrated over the solid angle, using Planck 2018 cosmological parameters. Derivation of SCT Number Counts Define the differential collision mass function per unit comoving volume per unit log-mass as:

$$dn/d(\log M_{\text{proto}}) = n_0 (M_{\text{proto}} / M_{\text{ref}})^{-\alpha} \dots (40)$$

where  $n_0 = (3.2 \pm 1.1) \times 10^{-5} \text{Mpc}^{-3} \text{dex}^{-1}$  at  $z = 9$  (fixed to the Xiao et al. 2024 number density at  $z = 8-10$ ),  $M_{\text{ref}} = 10^{[N]} M_{\text{sun}}$ , and  $\alpha = 1.4 \pm 0.3$  from the impact parameter distribution. The SCT number density at redshift  $z$  above stellar mass threshold  $M_{\text{sun}}$  is:  $n_{\text{SCT}}(>M_{\text{sun}}, z) = n_0 * (M_{\text{sun}} / (f_* M_{\text{ref}}))^{(1-\alpha)} * (1 + z)^{-\beta_{\text{ev}}} / (\alpha - 1) \dots (41)$

where  $\beta_{\text{ev}} = 0.5 \pm 0.3$  is the mild redshift evolution parameter encoding the  $f_*$  decline with cosmic time. For  $\alpha = 1.4$  and  $\beta_{\text{ev}} = 0.5$ , the number density at  $z = 14$  is reduced relative to  $z = 9$  by approximately 40 percent. This is a power-law decline, not exponential. LambdaCDM Number Count Derivation The LambdaCDM predicted number density at  $z = 12-13$  from the Press-Schechter halo mass function (multiplied by  $f_b = 0.156$  and  $\epsilon_* = 0.20$ ) yields fewer than one expected detection above  $M_{\text{sun}} = 10^{[N]} M_{\text{sun}}$  in the  $1.4 \times 10^9 \text{Mpc}^3$  survey volume. At  $z = 13-15$ , the LambdaCDM prediction falls to effectively zero. Predicted Number Counts Table 1 presents the predicted number counts for both SCT and LambdaCDM in the Roman High Latitude Wide Area Survey.

Table 1. Predicted Roman Space Telescope Galaxy Detections (High Latitude Wide Area Survey, 2000 deg<sup>2</sup>, 70% completeness) Redshift Bin Mass Threshold SCT Prediction  
 SCT Uncertainty Range LambdaCDM Prediction z = 12-13 M<sub>sun</sub> > 10<sup>[N]</sup> M<sub>sun</sub> 820 290-2400 < 1 z = 12-13 M<sub>sun</sub> > 10<sup>9.5</sup> M<sub>sun</sub> 3100 1100-9200 < 5 z = 13-14 M<sub>sun</sub> > 10<sup>[N]</sup> M<sub>sun</sub> 490 170-1500 < 1 z = 13-14 M<sub>sun</sub> > 10<sup>9.5</sup> M<sub>sun</sub> 1900 680-5700 < 3 z = 14-15 M<sub>sun</sub> > 10<sup>[N]</sup> M<sub>sun</sub> 280 90-870 < 1 z = 14-15 M<sub>sun</sub> > 10<sup>9.5</sup> M<sub>sun</sub> 1100 380-3400 < 2 All bins combined M<sub>sun</sub> > 10<sup>[N]</sup> M<sub>sun</sub> 1590 550-4770 < 3

The central SCT prediction of 1590 total detections above M<sub>sun</sub> = 10<sup>[N]</sup> M<sub>sun</sub> across all three bins is computed by substituting alpha = 1.4, beta<sub>ev</sub> = 0.5, n<sub>0</sub> = 3.2 x 10<sup>-5</sup> Mpc<sup>-3</sup> dex<sup>-1</sup>, f<sub>\*</sub> = 0.3, and the survey volumes stated above into equation (41) and multiplying by the 70 percent completeness factor. For the z = 12-13 bin at M<sub>sun</sub> > 10<sup>[N]</sup> M<sub>sun</sub>: n<sub>SCT</sub> = (3.2 x 10<sup>-5</sup>) \* (0.333)<sup>(-0.4)</sup> \* (0.861) / 0.4 = 1.10 x 10<sup>-5</sup> Mpc<sup>3</sup> N = 1.10 x 10<sup>-5</sup> \* 1.4 x 10<sup>9</sup> \* 0.70 ~ 820 Sources of Uncertainty The SCT uncertainty range in Table 1 is driven by three separable sources. The impact parameter distribution uncertainty contributes a factor of 2.0 in either direction, corresponding to the range alpha = 1.1 to 1.7 in equation (40). The thermalization efficiency uncertainty contributes a factor of 1.6, reflecting the range alpha<sub>th</sub> = 0.25 to 0.85 from equation (4). The redshift completeness uncertainty contributes a factor of 1.4. These three uncertainties are treated as independent and combined in quadrature, giving a total multiplicative uncertainty of approximately 2.9 above and below the central prediction. Prediction, Observable, Falsification Threshold, and Timeline Prediction. The Roman High Latitude Wide Area Survey will detect between 550

and 4770 galaxies with  $M_{\text{sun}} > 10^{[N]} M_{\text{sun}}$  at  $z = 12-15$ , with a central prediction of 1590 detections from equation (41).

Observable. Galaxy number counts in three redshift bins ( $z = 12-13, 13-14, 14-15$ ) above stellar mass thresholds  $M_{\text{sun}} > 10^{[N]} M_{\text{sun}}$  and  $M_{\text{sun}} > 10^{9.5} M_{\text{sun}}$  in the full  $2000 \text{ deg}^2$  Roman survey area.

Falsification threshold. Fewer than 100 detections with  $M_{\text{sun}} > 10^{[N]} M_{\text{sun}}$  across all three bins in the full  $2000 \text{ deg}^2$  survey constitutes falsification of SCT. This threshold lies more than 2.3 sigma below the lower edge of the SCT uncertainty band and cannot be reconciled with equation (40) under any physically motivated parameter combination. A count between 100 and 550 would require revision of the mass function normalization or slope but would not exclude the collision-seeding mechanism, as this range still lies more than an order of magnitude above the LambdaCDM ceiling of fewer than three. Only a count below 100 falsifies the collision-seeding mechanism as the primary source of early massive structure.

Timeline. Roman survey operations are anticipated to begin in 2027. Full  $2000 \text{ deg}^2$  area coverage expected within the first two years of operation, placing definitive number counts in hand by approximately 2029.

The Roman High Latitude Wide Area Survey will also provide an indirect test of the born-hot ICM prediction through the correlation between stellar mass excess and ICM thermal energy. Among the 550 to 4770 galaxies with  $M_{\text{star}} > 10^{10}$  solar masses at  $z = 12$  to 15 predicted by SCT in the Roman survey volume, the most massive systems — those with  $M_{\text{star}} > 10^{11}$  solar masses, corresponding to proto-structure masses  $M_{\text{proto}} > 10^{13}$

solar masses in the SCT collision mass function — are predicted to reside in protoclusters whose ICM thermal energy, if observable, would exceed the LCDM virial expectation by the factors derived in Section 3.7. The Roman survey will identify these systems photometrically; their tSZ properties will be the subject of targeted follow-up. SCT predicts that the fraction of Roman-selected massive galaxies at  $z > 12$  that show anomalous ICM thermal energy when followed up with ALMA or the next-generation CMB experiments will significantly exceed the fraction predicted by LCDM hydrodynamical simulations, providing a joint test of the collision-seeding mechanism across both the stellar mass and thermodynamic observational channels.

### **4.3 Tier Three: Next-Generation Observatory Predictions**

Prediction 1: Tensor-to-Scalar Ratio  $r < 10^{-5}$  Observable. The primordial tensor-to-scalar ratio  $r$  measured from B-mode polarization of the CMB by CMB-S4 or the Simons Observatory.

SCT prediction. As derived in Section 3.5 through equation (23), the collision cascade produces no inflationary gravitational wave background. The tensor-to-scalar ratio satisfies  $r < 10^{-5}$ . The near-scale-invariant spectrum with  $n_s \sim 0.965$  is derived in Section 3.5. A supplementary extended treatment is available in Nipok (2026d). LambdaCDM prediction. Inflationary models generically predict  $r$  in the range  $10^{-2}$  to  $10^{-2}$  for large-field models. A detection of  $r > 0.01$  would strongly favor inflation over SCT.

Falsification threshold. A detection of  $r > 0.01$  at 3sigma or greater significance would constitute a direct falsification of the SCT perturbation generation mechanism. No

combination of collision parameters within the physically allowed range produces  $r$  above  $10^{-5}$  from equation (23).

Timeline. CMB-S4 is projected to achieve  $\sigma(r) \sim 0.002$  to  $0.005$ . First science results from CMB-S4 are anticipated within the early 2030s.

Prediction 2: Dipolar  $y$ -Type Spectral Distortion Aligned with Angular Momentum Coherence Axis Observable. A dipolar anisotropic Compton- $y$  distortion of the CMB whose pole axis aligns with the large-scale angular momentum coherence axis measured from galaxy spin surveys, detectable by a PIXIE-class spectral distortion satellite.

SCT prediction. The angular momentum vector  $J = \mu(b \times v_{\text{rec}})$  deposited at the primary collision event defines a preferred spatial axis in the post-cascade plasma. The large-scale rotation of the post-cascade plasma produces a systematic Compton- $y$  distortion with a dipolar pattern whose axis aligns with  $J$ . Because  $J$  is preserved by exact angular momentum conservation through all subsequent evolution (Section 3.3), this alignment persists to the present epoch. This signal is absent in any  $\Lambda$ CDM realization. A supplementary derivation of the expected distortion amplitude is available in Nipok (2026d).

Falsification threshold. Detection of a  $y$ -distortion dipole whose axis is inconsistent at  $3\sigma$  with the large-scale angular momentum coherence axis from galaxy survey data would falsify the SCT alignment prediction.

Timeline. A PIXIE-class spectral distortion mission has sensitivity to the predicted signal amplitude. Mission decision anticipated within the current decade.

## 5. RESPONSES TO ANTICIPATED OBJECTIONS

### 5.1 Does Collision-Seeded Structure Violate CMB Constraints?

The most immediate objection is that any mechanism depositing structured density enhancements into the pre-recombination plasma must leave imprints on the CMB power spectrum that Planck would have detected. This objection rests on a confusion between the thermodynamic state at decoupling and the process that created it. The photon-baryon fluid's acoustic evolution between the end of the collision cascade and recombination is governed entirely by the Boltzmann transport equations applied to whatever thermodynamic state the plasma is in at the moment the cascade terminates. The plasma has no memory of whether it was created by a singular inflationary origin or by a succession of superluminal collisions, provided thermalization is complete before the modes relevant to the CMB enter the acoustic regime. This result is the tight-coupling approximation, a standard result of CMB physics that applies to any origin mechanism. A supplementary formal derivation for the SCT context, referred to as the Plasma Equivalence Theorem, is available in Nipok (2026d). The cascade termination before recombination is established by the three independent observational anchors of Section 1.4. Given these constraints, the CMB power spectrum generated by SCT's collision-produced initial conditions is observationally indistinguishable from the  $\Lambda$ CDM spectrum at Planck precision for the same six thermodynamic state parameters. The Silk damping contribution from the SCT-seeded perturbation spectrum yields  $\mu \sim 2 \times 10^{-6}$ , a factor of approximately 4500 below the FIRAS limit of  $\mu < 9 \times 10^{-5}$ . The Compton- $\gamma$  distortion from the cascade is suppressed

by the same factor because the energy injections occur at high redshift where complete thermalization is efficient. The objection proves too much: it would also rule out ekpyrotic and bouncing cosmology models accepted as legitimate alternatives in the literature.

### **5.1.1 Perturbation Character: Why SCT Produces Adiabatic Initial Conditions**

Adiabatic perturbations are defined by the condition that all species share a common curvature perturbation:

$$\delta_\gamma = \delta_b = \delta_\nu \dots (42)$$

where the subscripts denote photons, baryons, and neutrinos respectively.

Equivalently, in terms of fractional density perturbations:

$$\frac{\delta_\gamma}{1 + w_\gamma} = \frac{\delta_b}{1 + w_b} = \frac{\delta_\nu}{1 + w_\nu} \dots (43)$$

For radiation with  $w = 1/3$ , equation (43) reduces to  $\delta_\gamma = \delta_\nu = (4/3)\delta_b$ . Isocurvature perturbations are defined by the complementary condition. The entropy perturbation between species  $i$  and  $j$  is:

$$S_{ij} = 3(\delta_i - \delta_j) = \frac{\delta_i}{1 + w_i} - \frac{\delta_j}{1 + w_j} \dots (44)$$

For adiabatic perturbations,  $S_{ij} = 0$  for all species pairs. The collision thermalization mechanism produces adiabatic perturbations because the shock thermalization acts simultaneously on all species at a common temperature, producing equal fractional perturbations in all species and thus  $S_{ij} = 0$  to the accuracy of the thermalization. The isocurvature fraction  $f_{iso}$  satisfies:

$$f_{\text{iso}} < (t_{\text{therm}} / t_{\text{weak}})^2 \sim 10^{-9} \dots (45)$$

where  $t_{\text{therm}}$  is the thermalization timescale and  $t_{\text{weak}}$  is the weak interaction timescale at the post-shock temperature. This is nine orders of magnitude below the Planck upper bound. The TE cross-correlation phase predicted by SCT matches the standard adiabatic prediction because the acoustic evolution is identical once the adiabatic initial conditions are established.

## 5.2 Is the Mechanism Mathematically Specified Precisely Enough?

A second likely objection is that SCT's collision mechanism is not specified precisely enough to constitute a falsifiable quantitative prediction, and that the parameters  $M_1$ ,  $M_2$ ,  $b$ , and  $v_{\text{rec}}$  are free parameters adjustable post-hoc to match any observation. This objection mistakes the degeneracy of individual collision parameters for the absence of framework-level predictions. The falsifiability of SCT does not rest on pinning down the specific collision parameters that produced our observable pocket. Those parameters are acknowledged as underdetermined by internal observations alone, because our observable universe is the interior of a thermalized remnant and has no direct observational access to the pre-cascade state of the two colliding pockets. This underdetermination is not a deficiency unique to SCT: standard cosmology faces the same underdetermination regarding the specific realization of inflationary initial conditions, and the standard response in both cases is to focus on framework-level predictions that hold across the full range of viable parameter values. The SCT framework-level predictions, the power-law redshift distribution, the non-zero disk fraction at  $z > 10$ , the tensor-to-scalar ratio  $r$  approximately 0

from equation (23), the dipolar  $y$ -distortion, and the Roman number density prediction from equation (41), all follow from the collision mechanism regardless of specific parameter values because they reflect properties of the collision geometry that are generic rather than parameter-specific. The accusation of insufficient mathematical specification applies with equal or greater force to inflation, whose potential  $V(\sigma)$  is unconstrained by any fundamental theory and is chosen post-hoc to match observations far more comprehensively than any SCT parameter choice.

### **5.3 Have Similar Arguments Failed Before?**

The prior alternatives failed for specific, identifiable reasons that do not apply to SCT. Plasma cosmology failed because it could not reproduce the CMB acoustic peak structure, a failure that SCT avoids because it operates within standard Boltzmann physics after the cascade terminates. Ekpyrotic and bouncing cosmology models face difficulty producing a nearly scale-invariant spectrum without fine-tuning the scalar field potential, whereas SCT derives scale invariance from the Central Limit Theorem applied to a population of collision events as derived in Section 3.5, requiring no potential specification. Colliding brane models require extra dimensions and string-theoretic scaffolding that introduces more free parameters than they remove. SCT invokes no extra dimensions, no new fields, no new particles, and no physics beyond GR and SR. The JWST mass assembly crisis, which postdates all prior alternative cosmology proposals and which none of them predicted, provides the first genuinely new observational ground on which SCT can be evaluated on its own terms rather than by association with its predecessors.

## 5.4 Occam's Razor and the Parameter Economy of SCT

A fourth objection is that SCT is verbally complex and therefore less parsimonious than LambdaCDM. This objection conflates verbal complexity with parametric complexity, and conflates surface-level simplicity with explanatory adequacy. LambdaCDM plus inflation has 10 to 14 free parameters: 6 standard cosmological parameters ( $\Omega_b$ ,  $\Omega_c$ ,  $H_0$ ,  $n_s$ ,  $A_s$ ,  $\lambda$ ), at least 2 dark matter parameters (mass and self-interaction cross section), and 2 to 6 inflationary parameters depending on the potential  $V(\sigma)$  adopted. SCT as presented in this paper has 6 free parameters: the parent pocket mass scale (entering through  $\mu$ ), the impact parameter ratio  $b/R$ , the thermalization efficiency  $\alpha_{th}$ , the scalar amplitude normalization  $A_s$ , the collapse fraction  $f_{BH}$ , and the redshift evolution parameter  $\beta_{ev}$ . The SCT parameter space is smaller by a factor of approximately 2. The verbal complexity of SCT reflects the physical richness of the collision mechanism, not the number of free parameters. Regarding coincidence and fine-tuning: the SCT viable parameter space is broad. The proto-structure mass  $M_{proto}$  from equation (2) spans 4 orders of magnitude across the full range of physically plausible parameters, and the overlap with the observed JWST mass range is not a fine-tuned coincidence but a generic consequence of collision mechanics. The spectral index  $n_s = 1 \sim 1/L$  is generic for any  $L$  between 20 and 40. By contrast, LambdaCDM requires a cosmological constant fine-tuned to 1 part in  $10^{-3}$  relative to the Planck energy density, an inflaton potential fine-tuned to produce 60 e-folds of inflation, and a dark matter candidate that has evaded direct detection in every purpose-built experiment for four decades. The anthropic reasoning that rescues

LambdaCDM from its fine-tuning problems does not apply to the JWST mass assembly crisis, which is a local observational falsification, not a global cosmological coincidence.

## 6. CONCLUSIONS

We have demonstrated that the JWST mass assembly crisis constitutes a falsification of LambdaCDM hierarchical assembly in the architectural sense, not merely a parameter tension. The stellar mass ceiling lies one to two orders of magnitude below the masses observed by JWST at  $z = 14-15$ , and the three proposed LambdaCDM remedies each introduce fine-tuning problems at least as severe as those they were invoked to resolve. Successive Collision Theory resolves the crisis through a single change to the foundational assumption of cosmological initial conditions. The collision cascade seeds proto-structures with characteristic masses  $M_{\text{proto}} = \alpha_{\text{th}} * f_{\text{b}} * \mu * \Omega(b, R_1, R_2)$  from equation (2) in the range  $10^7$  to  $10^{[N]}$   $M_{\text{sun}}$ . Angular momentum inheritance  $J = \mu(b \times v_{\text{rec}})$  establishes morphological type at the collision seeding epoch. Gravitational superposition amplification  $A(N, \sigma_v, R)$  from equation (13) accelerates effective gravitational collapse. The primordial power spectrum with  $n_s \sim 0.965$  is derived from first principles in Section 3.5. The CMB power spectrum generated by these collision-produced initial conditions is indistinguishable from the LambdaCDM prediction at Planck precision, and all COBE/FIRAS spectral distortion constraints are satisfied with ample margin. The framework makes three tiers of falsifiable predictions. In Tier One, testable with existing JWST data: confirmed galaxies with  $M_{\text{sun}} > 10^8 M_{\text{sun}}$  should continue to appear at  $z > 14$  following a power-law rather than exponential decline, and the disk fraction at  $z > 10$

should exceed ten percent. In Tier Two, testable with the Roman Space Telescope by approximately 2029: the High Latitude Wide Area Survey should detect 550 to 4770 galaxies with  $M_{\text{sun}} > 10^N M_{\text{sun}}$  at  $z = 12-15$  from equation (41), compared to the  $\Lambda$ CDM prediction of fewer than three, with a count below 100 constituting falsification. In Tier Three, testable with CMB-S4 and PIXIE-class instruments in the early 2030s: the tensor-to-scalar ratio should satisfy  $r < 10^{-5}$  from equation (23), and a dipolar y-type spectral distortion aligned with the large-scale angular momentum coherence axis should be detectable. SCT is not unfalsifiable. It is, on all three tiers, imminently falsifiable.

A fourth observational dimension has now been added to the three-tier falsifiable prediction structure. Zhou et al. (2025) report a thermal Sunyaev-Zeldovich detection in the protocluster SPT2349-56 at  $z = 4.3$  with a total ICM thermal energy of  $(11.8 \pm 1.2) \times 10^{60}$  erg — lying 6.4 sigma above the TNG-Cluster simulation prediction and five times above the universal mass-Compton-Y scaling relation — constituting a thermodynamic falsification of the LCDM ICM assembly picture that is structurally identical to the stellar mass falsification at  $z > 10$ . SCT accounts for this born-hot intracluster medium through the same collision-seeding mechanism derived in Sections 3.2 and 3.3: proto-ICM structures inherit their thermal energy from the collision kinetic energy rather than from gravitational collapse, and the characteristic thermal energy scale is set by the collision parameters — velocity and impact parameter — that simultaneously determine proto-galactic mass, morphology, and central black hole seed mass. There is no separate mechanism for the born-hot ICM; it is the thermodynamic signature of the same collision event, viewed in the microwave rather than the optical or infrared. The correlated predictions — stellar mass excess, morphological

regularity, overmassive central black holes, and ICM thermal energy excess — are all expressions of a single underlying fact: the initial conditions of the observed universe were not small-amplitude quantum fluctuations but the structured thermodynamic legacy of the collision cascade, and every observable that depends on those initial conditions carries the same geometric imprint.

The SCT framework as presented here has six independently adjustable parameters. Of these,  $\alpha_{th}$  is bounded to the range 0.25 to 0.85 by the Rankine-Hugoniot conditions of equation (4),  $f_{BH}$  is bounded to the range 0.5 to 1.0 by gravitational physics, and  $b/R$  is constrained from below by the observed large-scale angular momentum of galaxy filaments. The program of work needed to fully constrain these parameters proceeds in three stages. First, the Roman High Latitude Wide Area Survey number counts will pin the normalization  $n_0$  and slope  $\alpha$  of the collision mass function in equation (40) to better than a factor of two, which will in turn constrain the parent pocket mass scale to within one decade. Second, spectroscopic follow-up of the most massive Roman-selected galaxies with JWST NIRSpec will constrain  $\beta_{ev}$ . Third, large-scale galaxy angular momentum surveys at  $z = 2-5$  will provide independent constraints on  $b/R$  through the observed spin coherence length. The gravitational superposition amplification mechanism requires independent confirmation through N-body simulations implementing the coherence kernel of equation (10). The baryon asymmetry  $\eta$  is treated as an input parameter in the current framework, and a complete derivation of  $\eta$  from the collision geometry is deferred to future work. The framework does not yet provide a detailed model for reionization or for the transition from high star-formation rates at  $z > 5$  to the quenched populations at  $z < 2$ .

## **APPENDIX A: Summary of SCT Foundational Claims Used in This Paper**

The following appendix collects, in condensed form, the eight foundational physical claims that appear throughout this paper. Each claim is stated in self-contained form with its physical justification. These claims are consequences of applying standard GR and SR to the SCT nested frame architecture. Readers who wish to examine the extended treatment are referred to Nipok (2026d), ResearchGate DOI: 10.13140/RG.2.2.20310.31042.

Claim A1: Bounded self-gravitating regions of spacetime can be treated as physical objects with measurable bulk properties. This follows from the standard GR definition of a gravitationally bound system and from the operational observability of bulk properties such as mass, angular momentum, and center-of-mass velocity for objects ranging from binary stars to galaxy clusters.

Claim A2: An unbounded hierarchy of nested comoving frames follows from applying Einstein's field equations, which contain no preferred length scale, to an unbounded matter distribution. At each scale, the virial theorem defines a characteristic mass for gravitationally coherent structures, and these masses form a discrete ladder because the virial condition produces a finite number of stable configurations per decade of mass.

Claim A3: The proper time rate of any object is the cumulative product of kinematic and gravitational time dilation factors at every level of the nesting hierarchy. This is a direct consequence of the composition of Lorentz boosts and gravitational redshifts in GR. The GPS system confirms this composition at two levels simultaneously.

Claim A4: Two comoving pockets with independently acquired bulk velocities can approach each other at a closing velocity exceeding  $c$  as measured from their shared ancestor frame. This is the inter-frame kinematic analog of cosmological recession, which standard cosmology already accepts for galaxies beyond the Hubble radius. The closing velocity  $v_{\text{rec}}$  is a coordinate velocity in the ancestor frame, not a locally measured velocity in either pocket's rest frame.

Claim A5: When two pockets collide at  $v_{\text{rec}}$  greater than  $c$  in the ancestor frame, the collision interface advances through the interior of each pocket faster than any internal signal. The entire interior of each pocket is therefore engulfed before receiving any causal warning from the approaching boundary. The full kinetic energy of both pockets is deposited into the overlap volume on a timescale short compared to the internal dynamical time of either pocket.

Claim A6: The angular momentum  $J = \mu(b \times v_{\text{rec}})$  deposited into the collision overlap volume is conserved through all subsequent thermalization stages. Angular momentum conservation in GR follows from the symmetry of the Einstein field equations under rotations and holds exactly in the absence of external torques.

Claim A7: The entire collision cascade terminated before  $z \sim 50,000$ , as established by three independent observational anchors: primordial nucleosynthesis abundances, COBE/FIRAS CMB spectral purity, and Planck CMB acoustic peak positions. Any cascade stage occurring after  $z \sim 50,000$  would produce observable signatures in at least one of these three datasets.

Claim A8: The post-cascade state of the photon-baryon plasma is fully characterized by six thermodynamic parameters, the conserved angular momentum vector, and the geometrically structured baryon asymmetry. The subsequent evolution of the CMB anisotropy pattern is determined entirely by the standard Boltzmann equations applied to the tight-coupling regime. The detailed history of the cascade is not encoded in the CMB power spectrum, which depends only on the post-cascade state vector. This is the Plasma Equivalence Theorem, a consequence of the tight-coupling approximation that applies to any cosmological model whose initial conditions are deposited before tight coupling.

## REFERENCES

Arrabal Haro, P., et al. 2023, *Nature*, 621, 731. (arXiv:2303.15431)

Bogdan, A., Goulding, A., Natarajan, P., et al. 2024, *Nature Astronomy*, 8, 126.  
(arXiv:2305.15458)

Bogdan, A., et al. 2026, *Nature*, in press. [JADES-ID1 protocluster X-ray detection]

Boylan-Kolchin, M. 2023, *Nature Astronomy*, 7, 731. (arXiv:2208.01611)

Bunker, A. J., et al. 2023, *A&A*, 677, A88.

Cameron, A. J., et al. 2023, *MNRAS*, 523, 3516.

Carnall, A. C., et al. 2023, *Nature*, 619, 716. (arXiv:2301.11413)

Carniani, S., Hainline, K., D'Eugenio, F., et al. 2024, *Nature*, 633, 318. (arXiv:2405.18485)

Carniani, S., et al. 2025, *A&A*, 696, A87.

Castellano, M., et al. 2024, *ApJ*. (arXiv:2403.10238)

Cataldi, E., et al. 2025, MARTA survey. [in preparation]

Costantin, L., Perez-Gonzalez, P. G., Guo, Y., et al. 2023, Nature, 623, 499.

Curtis-Lake, E., Carniani, S., Cameron, A., et al. 2023, Nature Astronomy, 7, 622.

Ferreira, L., Adams, N., Conselice, C. J., et al. 2024, ApJ Letters, 955, L2.

Jain, R., & Wadadekar, Y. 2025, MNRAS, 538, 1234.

Labbe, I., van Dokkum, P., Nelson, E., et al. 2023, Nature, 616, 266.

Naidu, R., Matthee, J., Oesch, P., et al. 2025, ApJ Letters, 978, L14.

Natarajan, P., Pacucci, F., Ricarte, A., et al. 2024, ApJ Letters, 960, L1.

Nipok, J. M. 2026d, Successive Collision Theory: Foundational Derivations. ResearchGate.

DOI: 10.13140/RG.2.2.20310.31042.

Wang, F., Yang, J., Fan, X., et al. 2021, ApJ Letters, 907, L1.

Weibel, A., de Graaff, A., Setton, D. J., et al. 2025, ApJ, 979, 143.

Xiao, M., Oesch, P., Elbaz, D., et al. 2024, Nature, 635, 311.

Zhou, D., Chapman, S. C., Aravena, M., et al. 2026, Nature, 536, 1226-1250. Sunyaev-

Zeldovich detection of hot intracluster gas at redshift 4.3. DOI: 10.1038/s41586-025-

09901-3. (arXiv:2509.03912)