

Entropic Scalar EFT: From Entanglement Microstructure to Gravity and Cosmic Structure

Jacob Chinitz

June 9, 2026

Abstract

We propose that the vacuum is not an inert stage but a finite-capacity entanglement substrate, and that matter is the localized depletion of that capacity. On this view, inertial mass is the entanglement content of a defect, gravity is the long-wavelength relaxation of the surrounding substrate, and the galactic excess usually attributed to dark matter is the extended reach of that relaxation rather than a separate particulate component.

The main result is a closed conditional derivation of the static weak-field sector. A minimal tetrahedral boundary ensemble, together with a one-bit fermionic defect anchor, turns the ultraviolet problem into a finite counting problem. Admissibility weighting, edge transport, finite-loop dressing, continuum matching, and source projection then determine the coefficients of a scalar entanglement EFT. In the resulting weak-field limit, Newtonian gravity, the galactic acceleration scale, the radial-acceleration relation, and leading no-slip lensing structure arise from one coefficient chain, without per-galaxy interpolation functions or phenomenological tuning.

The same construction also gives a non-gravitational route to the substrate length scale by treating the electron as the lightest elementary defect. The induced gravitational scale agrees with the observed Newton constant to about one percent, providing a quantitative cross-check on the matched weak-field normalization. The proposal is therefore not simply a scalar addition to general relativity: the metric field and the entanglement scalar are two continuum expressions of the same finite-capacity medium.

The Many-Pasts hypothesis enters this central result directly: the same history weighting that recovers the Born rule also makes the electron's dressing memoryless, one of the two conditions that fix the substrate length. Time-dependent transport, cosmology, and strong-field black-hole physics are developed as further consequences of this ontology. These sectors are not presented with equal closure. The static weak-field derivation, with its galaxy-scale phenomenology, is the central completed result; the transport, cluster, cosmological, and strong-field sectors remain conditional extensions, frontier completions, or open tests where explicitly stated.

Contents

| | |
|--|-----------|
| 1. Introduction: The Physical Claim | 4 |
| 1.1 What Is Primitive, and What Is Closed | 5 |
| 2. Canonical Field Content and Definitions | 6 |
| 3. The Three Postulates | 8 |
| 3.1 Information–Geometry Equivalence | 8 |
| 3.2 Mass–Entropy Equivalence | 8 |
| 3.3 Many-Pasts Hypothesis | 8 |
| 4. Relativistic Continuum Structure | 9 |
| 4.1 Capacity budget and continuum symmetry | 9 |
| 4.2 Dependency Map of the Theory | 10 |
| 5. Why a Tetrahedral Boundary Ensemble | 10 |
| 6. Admissibility Closure | 11 |
| 6.1 Minimal isotropic kernel | 11 |
| 6.2 Closure condition and uniqueness | 12 |
| 6.3 Effective sharing entropy | 12 |
| 7. Edge Kernel and Tree-Level Coupling | 13 |
| 8. Finite-Loop Renormalization | 14 |
| 9. Continuum Stiffness and SI Normalization | 15 |
| 10. Covariant Action | 16 |
| 11. Field Equations and Bridge Law | 18 |
| 12. Newtonian Gravity and the Point-Source Limit | 19 |
| 13. Electron Anchor and the Mass–Entropy Relation | 19 |
| 14. Galactic Dynamics | 20 |
| 15. Lensing, PPN, and Weak-Field Consistency | 22 |
| 16. Why Dynamics Requires Extension Beyond the Static Branch | 23 |
| 17. Causal Transport and Telegrapher Dynamics | 23 |
| 17.5 Cluster Source Projection and the Diffuse–Decoupled Channel Split | 24 |
| 18. Cosmology and the Hubble-Tension Sector | 29 |
| 19. Why These Sectors Belong | 30 |
| 20. Strong-Field Branch and Bounded Occupancy | 30 |
| 21. Many-Pasts: Operational Reduction and Arrow of Time | 31 |
| 22. Microstructure Hamiltonian and Underlying Dynamics | 32 |

| | |
|---|-----------|
| 23. Closure-Status Table | 32 |
| 24. Falsifiability and Observational Tests | 35 |
| 24.1 Static weak-field falsifiers | 35 |
| 24.2 Dynamical falsifiers | 35 |
| 24.3 Cosmological falsifiers | 35 |
| 24.4 Correlated-constant falsifiers | 35 |
| 24.5 Many-Pasts status | 36 |
| 25. What the Theory Would Have to Get Wrong to Fail | 36 |
| 26. Comparison with Other Approaches | 36 |
| 26.1 Relative to Λ CDM | 36 |
| 26.2 Relative to MOND-like interpolation programs | 37 |
| 26.3 Relative to Verlinde-style emergent gravity | 37 |
| 26.4 Relative to TeVeS and other multi-field modified gravities | 37 |
| 27. Conclusion | 37 |

Part I. Physical Idea and Canonical Definitions

1. Introduction: The Physical Claim

The central hypothesis is that the vacuum's local entanglement capacity is a real dynamical resource and that the localized objects we call matter are defects of that same resource rather than independent agents acting on it from outside. In this framework inertial mass is the entanglement content of a localized defect read through κ_m , gravity is the long-wavelength field carried by the surrounding restructuring of vacuum capacity around that defect, and what is usually modeled by particulate dark matter is the long-range reach of that same restructuring. The weak-field manifestation of that medium is a scalar EFT written in terms of a vacuum-relative entanglement field $S_{\text{ent}}(x)$ and its deficit relative to the background capacity. At continuum scale the defect sector is written in ordinary stress-energy variables, but its ontology is unchanged: it is still the coarse description of localized entanglement defects rather than a separate substance.

This is meant as a genuine replacement proposal for part of the usual dark-sector story, not simply as a new vocabulary laid on top of it. In the standard picture, one keeps visible matter and Einstein gravity, then introduces additional dark components to account for the missing gravitational response. Here the alternative hypothesis is that the vacuum already carries a finite entanglement-capacity structure, and that what we call matter is a localized defect of that structure. The same medium is then asked to explain ordinary weak-field gravity, the galactic excess usually attributed to dark matter, and the homogeneous mode relevant in cosmology.

The proposal is also meant to reach deeper than an ordinary scalar extension of Einstein gravity. If the underlying substrate has one finite update budget, an exact isotropic local structure, and no external background manifold, then the continuum description should already be Lorentzian and covariant before any further phenomenology is added. In that reading, GR is not an external geometric stage to which the entanglement field is later appended. It is the low-energy capacity geometry of the same substrate. The scalar sector then keeps track of how that geometry is depleted and redistributed by localized defects.

The derivation proceeds from microstructure to observables. After introducing the field content, postulates, and normalization conventions, the static weak-field coefficient chain is derived from a minimal tetrahedral boundary ensemble with admissibility closure, edge coupling, and finite loop dressing. The resulting EFT then recovers Newtonian gravity, the galactic acceleration scale, the radial acceleration relation, lensing consistency, and the leading weak-field post-Newtonian structure without per-system tuning. Time-dependent transport, cosmology, strong-field completion, and the Many-Pasts sector are taken up afterward as extensions of the same framework, though not all of those sectors are developed to the same degree of closure.

The central claim is sharper than a heuristic analogy: one scalar capacity mode can be followed from microscopic counting all the way to weak-field observables. In compressed form, the chain is

$$\text{microstructure} \longrightarrow \text{coefficient chain} \longrightarrow \text{continuum EFT} \longrightarrow \text{observables}.$$

That chain is the primary claim: the same mode is counted in the UV, stiffened by edge transport, sourced by localized defects, normalized by the weak-field bridge, and observed as Newtonian gravity plus galactic excess acceleration.

The logical order is simple. Part I says what the theory is about, fixes the variables, and explains why the continuum description should already be relativistic and covariant if the substrate picture is correct. Part II asks whether a minimal UV boundary structure can actually determine the coefficients that later appear in that continuum theory. Part III asks whether those coefficients reproduce ordinary gravity and the galactic weak-field phenomenology. Only

after that chain is visible do the later parts discuss time dependence, cosmology, strong-field completion, and quantum-foundational interpretation.

The physical hypothesis is global, but the most complete derivational closure is the static weak-field UV-to-EFT chain. Other sectors are developed as controlled consequences or structured frontier extensions.

The framework is unusual in five ways at once: general relativity is derived rather than posited, Newton’s G is predicted from the electron Compton scale and a single UV combinatorial number to about one percent, the galactic acceleration scale and radial-acceleration law are derived outputs with no per-system tuning, the horizon area coefficient $1/4$ closes from microscopic graph counting using the same Joyce diamond-lattice constant that fixes the weak-field source projection, and a closure-status table separates what is derivationally tight from what remains frontier. The point of the manuscript is that these features are not independent design choices but consequences of one finite UV counting problem.

1.1 What Is Primitive, and What Is Closed

The word “closure” is used here in a specific sense. The paper does not attempt to derive the existence of a finite entanglement substrate from no assumptions, nor does it derive the tetrahedral boundary ensemble from a still deeper microscopic Hamiltonian in the main text. Those are theory-defining ultraviolet inputs. The closure claim begins after those inputs are specified.

The primitive inputs are finite local entanglement capacity, matter as localized defects of that capacity, mass–entropy equivalence, the minimal tetrahedral boundary ensemble together with its physical realization as the full admissibility-closed sharing entropy, and the three postulates of Section 3 — in particular the Many-Pasts memorylessness of Postulate III. Given those inputs, the static weak-field sector is claimed to close: the finite boundary counting problem fixes the admissibility weighting, effective sharing entropy, the faithful sector-resolution principle that sets the substrate cell length, edge kernel, loop dressing, continuum stiffness, source projection, weak-field bridge, Newtonian limit, galactic acceleration scale, radial-acceleration relation, and leading no-slip metric structure without per-system tuning.

In compressed form, the claim is

$$\boxed{\text{primitive UV capacity hypothesis}} \rightarrow \boxed{\text{finite counting + admissibility}} \rightarrow \boxed{L_*}$$

$$\rightarrow \boxed{\gamma, \kappa/\gamma} \rightarrow \boxed{\delta S \leftrightarrow \Phi} \rightarrow \boxed{G, a_0, \text{RAR, no slip}}.$$

The microstructure is therefore not itself the result being proven. It is the finite ultraviolet counting problem from which the weak-field sector is derived. The absolute length calibration is supplied by the sector-resolution principle

$$\ln\left(\frac{\lambda_e}{L_*}\right) = 7g_{\text{share,eff}} - \ln\left(\frac{3}{2}\right),$$

which identifies the electron with the maximally-extended (seven-pass) ground state of the history-space sector-resolution operator, supported over its reduced Compton scale. The seven-fold additive form $7g_{\text{share,eff}}$ is the additivity of the history distance over the seven channels for the product dressing cloud, and the $\ln(3/2)$ correction is the transverse export factor derived in Appendix C.5. The sector-resolution principle is itself derived. Its seven-fold additivity asks each dressing channel to carry the full admissibility entropy and the seven channels to be mutually independent: the first because the dressing is memoryless in substrate time, the Many-Pasts content of Postulate III; the second because inter-channel correlation would contract the defect and

raise its mass, so the electron, the lightest charged defect, carries none (Section 22, Appendix H). The central technical claim is therefore that, given the three postulates and the realized minimal tetrahedral ensemble, the static weak-field branch has no remaining phenomenological knobs.

This distinction matters because the continuum action can superficially resemble an ordinary scalar-tensor theory. In a conventional scalar-tensor model, one usually asks whether variation of a covariant action containing a scalar field can produce the desired metric phenomenology. Here the scalar is the continuum order parameter of an underlying capacity substrate, and the weak-field bridge is the emergent-geometry dictionary relating fractional capacity depletion to the metric lapse. The closure question is therefore not whether the tetrahedral microstructure has been derived from nothing, but whether the specified microscopic capacity ensemble fixes the weak-field gravitational response once it is adopted.

The remaining external tasks are correspondingly clear: derive the same microscopic ensemble from a deeper substrate Hamiltonian, independently audit the graph-level return operator, confront the resulting radial-acceleration law and no-slip prediction with data, derive the horizon boundary action from the graph ensemble, and extend the constrained-capacity black-hole branch into a fully dynamical collapse theory. These are reconstruction and validation tasks outside the already-closed static weak-field branch, not hidden fit parameters inside it.

2. Canonical Field Content and Definitions

We define the fundamental continuum variable as the vacuum-relative coarse-grained entanglement assigned to a UV probe cell of size L_* centered at x :

$$S_{\text{ent}}(x) \in \mathbb{R},$$

measured in nats and therefore dimensionless. This is not a literal microscopic entropy density at a mathematical point. It is the leading scalar order parameter associated with a vacuum-relative entanglement defect after coarse-graining over a UV cell.

This definition keeps the microscopic and continuum pictures tied together. At continuum level, $S_{\text{ent}}(x)$ is the field that appears in the action and field equations. At the microscopic level it is the coarse variable recording how much local entanglement capacity remains available in the underlying medium after averaging over a UV cell.

The asymptotic vacuum-capacity baseline is denoted S_∞ , and the deficit field is

$$\delta S(x) \equiv S_\infty - S_{\text{ent}}(x).$$

Positive δS denotes reduced available vacuum entanglement capacity in the neighborhood of a localized defect or defect distribution. It is the extended restructuring field sourced by the defect sector, not an independent medium acted on by matter from outside. For nonlinear work it is useful to define the bounded occupancy fraction

$$q(x) \equiv \frac{S_{\text{ent}}(x)}{S_\infty} = 1 - \frac{\delta S}{S_\infty} \in [0, 1].$$

The variables S_{ent} , δS , and q therefore describe the same local physics in three closely related ways: available capacity, missing capacity relative to vacuum, and surviving-capacity fraction. The weak-field theory is most transparent in δS because it talks directly to the Newtonian potential. The nonlinear completion is most transparent in q because boundedness is built in from the start. The operational meanings are:

- $q = 1$: vacuum capacity fully available in the absence of local defect-induced restructuring;

- $0 < q < 1$: partial local capacity reduction around a defect configuration;
- $q = 0$: complete local exhaustion of available capacity on the physical branch.

Fixed-epoch normalization. The absolute normalization of S_{ent} and S_{∞} is a convention once an epoch and cell convention have been fixed. Under a constant rescaling

$$S_{\text{ent}} \mapsto K S_{\text{ent}}, \quad S_{\infty} \mapsto K S_{\infty}, \quad \delta S \mapsto K \delta S,$$

the observable bridge

$$\frac{\Phi}{c^2} = -\frac{\delta S}{2S_{\infty}}$$

is unchanged. The source equation is invariant in the same sense: rescaling the entropy field rescales the source coefficient with it, so the observable Newtonian normalization depends on the gauge-invariant combination $\kappa/(\gamma S_{\infty})$ rather than on S_{∞} alone. A cell-normalized description and a horizon-normalized description can therefore assign different numerical values to S_{∞} without changing Φ , G , or the PPN limit. This is not a time-dependent gauge symmetry; it is a fixed-epoch entropy-unit convention. Gravity sees fractional capacity depletion.

Substrate length scale. The canonical UV cell length is not taken to be the conventional Planck length as an input. It is fixed by the ground-state faithful sector-resolution principle

$$\ln\left(\frac{\lambda_e}{L_*}\right) = 7g_{\text{share,eff}} - \ln\left(\frac{3}{2}\right), \quad \lambda_e = \frac{\hbar}{m_e c}.$$

Equivalently,

$$L_* = \frac{3}{2} \lambda_e e^{-7g_{\text{share,eff}}} = 1.60771947 \times 10^{-35} \text{ m}.$$

The corresponding induced gravitational scale is

$$G_* := \frac{c^3 L_*^2}{\hbar} = \frac{9}{4} \frac{\hbar c}{m_e^2} e^{-14g_{\text{share,eff}}} = 6.60399128 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}.$$

The principle says that the lightest one-bit fermionic defect coherently resolves the seven face sectors exactly once and exports only the transverse 2/3 share of that dressing block. The conventional Planck length $L_P = \sqrt{\hbar G/c^3}$ remains useful for comparison and for standard black-hole thermodynamic notation, but it is not the primitive scale-setting input here.

The principal coefficients and derived quantities used throughout are:

$$\gamma : \text{entanglement-field stiffness}, \tag{1}$$

$$\kappa : \text{defect-entropy coupling}, \tag{2}$$

$$\kappa_m(\ell) : \text{mass-per-entropy map at scale } \ell, \tag{3}$$

$$L_* : \text{substrate cell length fixed by faithful sector resolution}, \tag{4}$$

$$G_* : \text{gravitational scale induced by } L_*, \tag{5}$$

$$g_{\text{share,max}} = \ln(1680), \tag{6}$$

$$g_{\text{share,eff}} : \text{admissibility-weighted effective sharing entropy}, \tag{7}$$

$$J_{\text{bare}}, J_{\text{eff}}^{\text{tree}}, J_{\text{eff}}^{\text{(ren)}} : \text{UV edge-kernel couplings}, \tag{8}$$

$$a_0 = \frac{c H_0 g_{\text{share,eff}}}{4\pi^2}. \tag{9}$$

The gravitational potentials are denoted Φ and Ψ , and the canonical weak-field bridge will be written as

$$\frac{\Phi}{c^2} = -\frac{\delta S}{2S_{\infty}}.$$

These same symbols reappear in the UV closure chain, in the continuum action, and in the phenomenology sections. From this point onward each one keeps the same meaning, so the later derivations can build on a single notation rather than shifting between parallel conventions.

These definitions are fixed canonically and used without further redefinition below.

3. The Three Postulates

3.1 Information–Geometry Equivalence

The first postulate states that vacuum-relative entanglement structure contributes to spacetime curvature on equal footing with ordinary stress-energy. In the EFT this means that the scalar field $S_{\text{ent}}(x)$ enters a covariant action, contributes its own stress-energy, and couples to a trace-equivalent defect source. At continuum scale that source is written in the usual stress-energy variables, but ontologically it is the coarse description of the localized defect sector. In weak field, metric response is governed not by absolute entropy but by the deficit relative to the vacuum-capacity baseline.

The role of Postulate I is to say what gravity is sensitive to. Einstein gravity already tells us that geometry responds to physical content. The present extension says that local entanglement-capacity structure is part of that content. Once that is accepted, gradients and deficits of the entanglement field are no longer metaphorical; they belong in the gravitational bookkeeping alongside the usual stress-energy variables.

3.2 Mass–Entropy Equivalence

The second postulate identifies inertial mass with the entanglement content of a localized defect. At scale ℓ ,

$$m(\ell) = \kappa_m(\ell) \Delta S.$$

For elementary fermionic sectors the canonical defect increment is

$$\Delta S_f = \ln 2,$$

because a spin-1/2 fermionic face exclusion creates a binary occupied/unoccupied defect of the local network and therefore carries exactly one bit of missing entanglement. This provides the cleanest anchor for the mass–entropy map. Mass and entanglement are therefore not two separate substances linked by an empirical proportionality; they are two descriptions of the same localized defect sector at different levels of coarse-graining. For composite sectors, the relevant quantity is the fully dressed bound-state entanglement budget rather than a bare constituent count.

The purpose of this postulate is to remove the temptation to think of matter as external to the medium. In the present ontology, a particle is already a localized defect of the entanglement substrate. Writing $m = \kappa_m \Delta S$ therefore does not assert an analogy between two independent things. It asserts that the inertial content of the defect is the entanglement content of the defect, read in mass units.

3.3 Many-Pasts Hypothesis

The third postulate is part of the full framework. The coefficient chain that fixes the Newtonian limit, the galactic law, and lensing does not call on it; the substrate-scale branch does, through the memoryless dressing described below. In canonical closed form the operational history weight is

$$P(H|P) \propto e^{-D(H,P)},$$

equivalently the branch $\alpha = 1, \beta = 0$ of the generalized family. This closed operational branch is fixed because exact Born recovery forces $\alpha = 1$ and forbidding any extra signaling-sensitive operational bias channel forces $\beta = 0$. Its operational consequences, together with the arrow of time, are developed in the interpretive and cosmological sector.

Read at the level at which it is stated, the postulate is a statement about whole histories: by weighting paths by $e^{-D(H,P)}$ it makes the dressing of a defect memoryless in substrate time, each pass an independent draw from the admissibility ensemble. That temporal memorylessness is one of the two ingredients that fix the absolute substrate scale in Appendix D.4; the other, the mutual independence of the electron's seven dressing channels, follows from its identity as the lightest charged defect (Appendix H).

Postulates I and II define the gravitational ontology directly. Postulate III enters at two points: it carries the account of branch realization and temporal asymmetry, and through the memoryless dressing it also enters the substrate-scale branch of the weak-field chain itself.

The three postulates define the ontology of the theory. The main text treats them as theory-defining inputs, not as derived outputs.

4. Relativistic Continuum Structure

4.1 Capacity budget and continuum symmetry

In the present framework the continuum description is expected to be covariant not because a geometric axiom is added at the outset, but because the substrate itself is finite-capacity, isotropic, and relational.

The first ingredient is a finite maximal update rate, denoted by the same constant c that later appears in the transport relation $D/\tau_0 = c^2$. In the present interpretation, c measures the largest rate at which the substrate can propagate and reorganize information. A defect at rest spends that budget entirely on local temporal evolution. A defect in motion must spend part of the same budget on spatial restructuring of the surrounding network. Because the substrate is isotropic, the cost of motion depends only on the rotational scalar v^2 at leading order, with the temporal rate maximal at $v = 0$ and vanishing when the budget is exhausted at $v = c$. These endpoint conditions alone admit many interpolating functions and so do not fix the form of the time-dilation relation. The form is fixed once the finite update speed is treated as invariant across inertial coarse descriptions: homogeneity, isotropy, and the relativity principle then select the Lorentz group rather than the Galilean one, giving the invariant interval

$$c^2 d\tau^2 = c^2 dt^2 - d\mathbf{x}^2,$$

and hence

$$\frac{d\tau}{dt} = \sqrt{1 - \frac{v^2}{c^2}}.$$

The capacity-budget picture supplies the substrate interpretation of this Lorentzian kinematics: motion allocates part of the finite update budget to spatial restructuring, leaving the remaining fraction as proper-time evolution.

The same capacity language also unifies motion-induced and gravity-induced clock slowing. In the nonlinear branch the surviving-capacity fraction is

$$q = \frac{S_{\text{ent}}}{S_{\infty}},$$

so smaller q means that less local update capacity remains available. Motion reduces the temporal share of the budget by consuming part of it in spatial transport; a nearby defect reduces the

local budget by depleting available capacity. The two familiar time-dilation effects are therefore interpreted as two regimes of one mechanism.

The second ingredient is the relational character of the substrate. It is not embedded in a prior physical manifold whose coordinate labels carry independent meaning. The physical content is the pattern of local capacities, defects, and neighborhood relations within the network itself. Continuum coordinates are therefore descriptive labels imposed on that relational structure, not additional physical data. Smooth changes of coordinates relabel the same underlying configuration rather than altering the physics. In continuum language this is precisely why the low-energy description should be written in generally covariant form.

The upshot is that the metric sector of the EFT is not being introduced from outside. Lorentzian geometry is the natural coarse description of a finite-capacity, isotropic, relational substrate, and the Einstein sector is its lowest-order continuum gravitational expression. The entanglement scalar then tracks how that same capacity geometry is redistributed by localized defects. The resulting low-energy theory can therefore be written in the usual covariant language, but the intended logic runs from substrate properties to geometry, not the other way around. As with any discrete substrate, this is a continuum statement: exact Lorentz and diffeomorphism symmetry belong to the coarse theory, while lattice-scale corrections may survive near the UV cutoff.

4.2 Dependency Map of the Theory

The logical flow of the theory can be summarized compactly as

Postulates \rightarrow UV boundary ensemble \rightarrow admissibility closure
 \rightarrow edge kernel \rightarrow finite renormalization \rightarrow continuum matching \rightarrow weak-field EFT
 \rightarrow static observables \rightarrow transport / cosmology /
 strong field / Many-Pasts.

This is a dependency graph, not an epistemic-equality graph. The static weak-field sector, the UV coefficient chain that feeds it, and the operational Born-recovery branch are more tightly closed than the cosmological or strong-field sectors. Part VI makes that difference explicit in a closure-status table.

The remainder of the argument follows this order so that each later result can build on the same coefficient choices and the same field dictionary.

Part II. UV Coefficient Chain

5. Why a Tetrahedral Boundary Ensemble

The microstructural problem is to identify a minimal discrete boundary-cell architecture capable of supporting finite channel entropy, isotropic closure data, and a continuum scalar response. The canonical choice adopted here is a tetrahedral cell with four structural ingredients:

- a tetrahedral volumetric cell;
- half-integer fermionic face data on each face;
- injective face assignment;
- binary orientation/parity.

This package is not presented as the only imaginable UV completion of emergent gravity. It is the minimal architecture used here to support the needed closure properties. The tetrahedron is the minimal volumetric simplex in $d = 3$, injectivity preserves independent boundary information across the four faces, and parity doubling captures the two orientations of the cell. The face-state multiplicity is then not chosen from a menu. Postulate II identifies elementary defects as fermionic, so each face carries half-integer base spin

$$j_0 = \frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \dots$$

Two cells sharing a face therefore generate the effective boundary sector

$$j_0 \otimes j_0 = 0 \oplus 1 \oplus \dots \oplus 2j_0.$$

Postulate I selects the maximum-capacity boundary channel, so the effective face label is the top channel

$$j_{\text{eff}} = 2j_0,$$

with

$$|M| = 2j_{\text{eff}} + 1 = 4j_0 + 1$$

distinguishable face states. Injectivity across four tetrahedral faces requires at least four distinct labels, so

$$|M| \geq 4 \implies 4j_0 + 1 \geq 4.$$

The only half-integer option below $j_0 = 3/2$ is $j_0 = 1/2$, which gives $j_{\text{eff}} = 1$ and $|M| = 3$, so it fails the injectivity condition. The first fermionic choice that works is therefore

$$j_0 = \frac{3}{2}, \quad j_{\text{eff}} = 3, \quad |M| = 7.$$

In that sense the seven-state face sector is derived from fermionic face data, maximum-capacity channel selection, tetrahedral injectivity, and minimality. Its appearance is fixed before any contact with the weak-field observables it later feeds. The same face-level structure is also where the elementary matter sector enters: fermionic face exclusion creates the binary one-bit defect increment $\Delta S_f = \ln 2$ used later in the electron anchor.

The resulting combinatorial state count is

$$\Omega_{\text{tet}} = 2 \times P(7, 4) = 2 \times 840 = 1680,$$

so the combinatorial sharing ceiling is

$$g_{\text{share, max}} = \ln(1680) \approx 7.427.$$

The exact K^2 spectrum and multiplicities are carried in the appendices. The essential physical point is that the UV theory begins with a finite microscopic counting problem rather than a free continuum ansatz.

This fixes the minimal structural package used in the ultraviolet construction; the supporting derivations are collected in the appendices.

6. Admissibility Closure

6.1 Minimal isotropic kernel

The UV boundary ensemble is not used with a flat weighting. The admissibility family is

$$p_\eta(b) \propto e^{-\eta K^2(b)},$$

where K^2 is the unique leading quadratic closure-defect scalar compatible with tetrahedral symmetry. This choice is not made because it “works” phenomenologically. It is the minimal isotropic maximum-entropy kernel under normalization and fixed quadratic closure moment. Higher invariants such as K^4 would correspond to additional UV information and therefore to subleading refinements rather than competing leading kernels.

The reason for introducing this weighting is that the raw combinatorial ensemble is too permissive to be the whole UV story. Some boundary configurations are closer to the regular closure pattern expected of a smooth medium, while others are more distorted. The scalar K^2 is the minimal rotationally invariant way to measure that distortion. The admissibility kernel therefore says, in the mildest possible form, that more badly closed configurations should contribute less to the effective coarse ensemble.

6.2 Closure condition and uniqueness

The admissibility precision η is not chosen externally; it is fixed by maximizing the normalized closure evidence. Tetrahedral closure is the vanishing of the three-component oriented-face sum, so the closure-defect space is three-dimensional, and the quadratic family on it carries a determinant weight $\eta^{3/2}$. The closure-evidence functional is therefore

$$\mathcal{F}(\eta) = \ln Z(\eta) + \frac{3}{2} \ln \eta,$$

and its stationary point gives the closure condition

$$\langle K^2 \rangle_\eta = \frac{3}{2\eta},$$

in which the factor $3/2$ is the determinant weight of the three independent closure components, while the discreteness and multiplicities of the spectrum stay inside the exact sum $Z(\eta)$. This is the stationary normalized-evidence point of the exact closure spectrum, and it is a maximum rather than a bare root (Appendix B.2).

On that spectrum it is unique,

$$\eta_* = 0.0298668443935.$$

The closed branch is locally stiff: small fractional changes in η produce only small fractional changes in the downstream effective sharing entropy.

6.3 Effective sharing entropy

The admissibility-weighted effective sharing entropy is

$$g_{\text{share,eff}} = 7.41980002357.$$

The gap between $g_{\text{share,max}}$ and $g_{\text{share,eff}}$ is therefore not loss imposed by hand. It is the difference between the raw combinatorial ceiling and the admissibility-closed effective boundary entropy that actually propagates into observable couplings.

This distinction is one of the conceptual pivots of the theory. The continuum description does not inherit the naive channel-counting ceiling; it inherits the portion of the channel space that survives after closure is imposed. The downstream couplings should therefore be read as consequences of admissibility-closed sharing, not of raw combinatorics alone.

At this stage the effective sharing entropy is no longer a free choice. The exact spectrum, multiplicities, uniqueness proof, and stiffness numerics are preserved in the appendices for reference.

7. Edge Kernel and Tree-Level Coupling

The same UV closure data fix the tree-level edge kernel. The geometric bridge is the tetrahedral identity

$$\sum_{i=1}^4 \hat{n}_i \hat{n}_i^\top = \frac{4}{3} I_3,$$

which implies a channel-averaged transverse fraction of $2/3$. This gives the bare edge smoothness coupling

$$J_{\text{bare}} = \frac{2}{3} \eta_*.$$

The interpretation of J_{bare} is straightforward: it is the cost assigned to mismatch between the occupancy variables of two neighboring coarse cells. If adjacent cells disagree strongly, the edge pays a larger penalty; if they agree, the penalty is small. The factor $2/3$ is the geometric fraction that survives after averaging the tetrahedral channels into the isotropic continuum limit.

For a $z = 4$ regular coarse adjacency graph, the tree-to-lattice reduction then yields

$$J_{\text{eff}}^{\text{tree}} = \frac{J_{\text{bare}}}{3} = \frac{2\eta_*}{9}.$$

The division by 3 comes from the branching geometry of the rooted $z = 4$ graph. One neighboring link points back toward the source, while the remaining $z-1 = 3$ links carry the forward transport into the tree. Thus $J_{\text{eff}}^{\text{tree}}$ is not simply the microscopic edge penalty itself, but the part of that penalty that survives as net long-range transport after the local branching structure is taken into account.

Origin of the horizon target. The horizon target

$$\sigma_* = \frac{\pi}{g_{\text{share,eff}}}$$

is the closure-consistency value required by the horizon-normalized field convention. In the admissibility-closed boundary ensemble, one active microscopic sharing unit carries effective entropy $g_{\text{share,eff}}$. In the continuum normalization used for the weak-field scalar, the occupancy variable is normalized by

$$S = \pi Q_{\text{occ}},$$

so a coarse horizon-normalized channel with occupancy $Q_{\text{occ}} = 1$ carries entropy π in the S -field convention. If σ_* denotes the asymptotic conditional-independence weight seen by the rooted shell hierarchy (Appendix B.3), consistency between the boundary entropy count and the horizon-normalized continuum field requires

$$\sigma_* g_{\text{share,eff}} = \pi,$$

and therefore

$$\sigma_* = \frac{\pi}{g_{\text{share,eff}}} = 0.42340665\dots$$

The role of σ_* is to match the admissibility-closed microscopic entropy normalization to the horizon-normalized scalar-field convention; it is the closure target fixed by that choice of branch. The rooted shell observable converges to that target rapidly enough that the nonlocal correction is already strongly constrained by small shell depth.

The story is no longer just one of state counting. The edge kernel measures how costly it is for neighboring coarse cells to disagree in local occupancy. The tetrahedral identity is what makes

this bridge controlled: it is the statement that the four discrete channel directions average to the correct isotropic tensor structure in the continuum limit.

Tree-level edge transport is fixed by the same microscopic data that fixed the admissibility closure. The shell hierarchy and phase-selection checks are preserved in Appendix C.

8. Finite-Loop Renormalization

Tree level is not the whole UV story. The full lattice admits local closed-return motifs that recycle part of the transmitted information before it contributes to net coarse transport. The leading correction is organized as a local Dyson self-energy dressing,

$$J_{\text{eff}}^{(\text{ren})} = \frac{J_{\text{eff}}^{\text{tree}}}{1 + J_{\text{eff}}^{\text{tree}} \Sigma_{\text{ret}}}.$$

The need for this step is physically straightforward. A purely tree-like transmission rule would let the relevant amplitude move outward once and never locally return. A real coarse graph is not that simple. Some of the transmitted information cycles back through short closed motifs before contributing to long-distance transport. The renormalized coupling is therefore the true stiffness felt by the coarse field after these local returns have been resummed.

The structural decomposition is the key result. There are seven sector-diagonal local returns, together with one permutation-symmetric shared closure-singlet. The singlet is weighted by the same transverse projection and branch dilution factors that define the tree edge map,

$$\binom{2}{3} \binom{1}{3} = \frac{2}{9},$$

so the leading local self-energy is

$$\Sigma_{\text{ret}} = 7 + \frac{2}{9} = \frac{65}{9}.$$

Equivalently, on the seven-channel scalar return space,

$$R_{\text{ret}} = I_7 + \frac{2}{9} P_{\text{sing}}, \quad P_{\text{sing}} = |u\rangle\langle u|, \quad u = \frac{1}{\sqrt{7}}(1, \dots, 1),$$

and $\Sigma_{\text{ret}} = \text{Tr}(R_{\text{ret}})$. The orthogonal six-dimensional sum-zero sector carries no net scalar charge in the coarse branch and therefore does not add a separate scalar return. Information sent along an edge need not simply move outward once and for all. Some of it can circulate through short local loops before contributing to coarse transport. The seven sector-diagonal returns are the seven face-label channels that return independently without mixing. In addition there is one collective mode, symmetric across channels, that returns as a shared closure-singlet rather than as a channel-specific loop. The self-energy is therefore not a generic loop number but a sum of seven independent return channels plus one shared mode weighted by the same projection and branching factors already present in the tree map. Hence

$$c_{\text{loop}}^{(\text{ren})} \equiv \frac{J_{\text{eff}}^{(\text{ren})}}{J_{\text{eff}}^{\text{tree}}} = \frac{1}{1 + J_{\text{eff}}^{\text{tree}} \Sigma_{\text{ret}}} \approx 0.95426,$$

and

$$J_{\text{eff}}^{(\text{ren})} \approx 0.00633348.$$

This reproduces the shell-target crossing near $J_{\text{bare,cross}} \sim 0.019$ at the 0.05% level.

What matters here is that the loop correction is no longer schematic. The finite renormalization is written as an explicit local self-energy. The remaining audit task is independent graph-level confirmation of the same scalar-return operator, not the introduction of any new loop parameter.

9. Continuum Stiffness and SI Normalization

The last UV step is not a thermodynamic one. The lattice quadratic form is interpreted as a Euclidean action weight,

$$\frac{I_E}{\hbar} = \frac{J_{\text{eff}}^{(\text{ren})}}{2} \sum_{a,i} (Q_a - Q_{a+L_*\hat{n}_i})^2,$$

where the sum runs over all sites a and all four outgoing nearest-neighbor directions \hat{n}_i , so each nearest-neighbor edge is counted twice. The microscopic four-cell has size

$$\Delta V_4 = \frac{L_*^4}{c}.$$

Up to this point the derivation has determined a dimensionless lattice weighting. The continuum EFT, however, needs a dimensionful coefficient multiplying derivatives of a field in spacetime. The Euclidean-action interpretation is what upgrades the lattice closure data into a continuum action density with the right units and the right covariant target.

The same tetrahedral identity used in the edge-kernel reduction then yields the continuum coefficient for the occupancy field Q_{occ} ,

$$\gamma_Q = \frac{4\hbar c}{3L_*^2} J_{\text{eff}}^{(\text{ren})}.$$

Here L_* is the canonical tetrahedral spacing, with one coarse cell carrying volume L_*^3 up to the fixed cell-shape convention, and $J_{\text{eff}}^{(\text{ren})}$ is the loop-dressed edge coupling. The numerical factor $4/3$ is the isotropic projection

$$\sum_i \hat{n}_i \hat{n}_i^T = \frac{4}{3} I_3$$

that turns the tetrahedral edge directions into the continuum gradient tensor.

The field normalization is fixed by horizon capacity:

$$S = \pi Q_{\text{occ}}.$$

Therefore the canonical EFT coefficient in the $\frac{\gamma}{2}(\partial S)^2$ convention is

$$\gamma = \frac{4\hbar c}{3\pi^2 L_*^2} J_{\text{eff}}^{(\text{ren})}.$$

The faithful sector-resolution principle fixes L_* without using G . It is nevertheless useful to define the gravitational scale induced by this length,

$$G_* := \frac{c^3 L_*^2}{\hbar}.$$

Then the stiffness may be written in Einstein-normalized form as

$$\gamma = \frac{4J_{\text{eff}}^{(\text{ren})}}{3\pi^2} \frac{c^4}{G_*}.$$

This is the same algebra as the familiar Planck-cell rewrite, but read in the opposite direction: the substrate cell length induces the gravitational scale rather than being chosen by first inserting the measured value of G . Within the Euclidean-action convention already assumed by the EFT, the SI-normalized weak-field stiffness coefficient is fixed rather than left schematic.

At a high level, Part II has now completed the micro-to-continuum coefficient story. The tetrahedral ensemble determines the effective sharing entropy, the edge kernel and loop dressing turn that entropy data into a discrete stiffness, and the Euclidean matching turns the discrete stiffness into the continuum coefficient γ that appears in the weak-field EFT.

Closed UV-to-IR chain. The UV coefficient chain can now be summarized as

$$\{\Omega_{\text{tet}}, K^2, \eta_*, g_{\text{share,eff}}, L_*, J_{\text{bare}}, J_{\text{eff}}^{\text{tree}}, \Sigma_{\text{ret}}, J_{\text{eff}}^{(\text{ren})}, \gamma\} \longrightarrow \{\kappa, G, a_0, g_{\text{obs}}(g_{\text{bar}})\}.$$

The first bracket is the micro-to-continuum closure chain; the second bracket collects the weak-field observables it feeds. The rest of the manuscript uses this chain rather than introducing independent weak-field coefficients.

With that normalization in place, the stiffness chain runs from the discrete ensemble to the weak-field EFT without an extra coefficient choice. The remaining microscopic question is independent confirmation of the same action-kernel interpretation from fuller inhomogeneous dynamics, not an unresolved normalization constant.

Part III. Weak-Field EFT and Static Phenomenology

10. Covariant Action

With that continuum symmetry structure in place, the canonical weak-field EFT takes the covariant form

$$I = \int d^4x \sqrt{-g} \left[\frac{c^4}{16\pi G} R - \frac{\gamma}{2} g^{\mu\nu} \partial_\mu S_{\text{ent}} \partial_\nu S_{\text{ent}} - \lambda S_{\text{ent}} - \kappa \chi S_{\text{ent}} \right],$$

with

$$\chi(x) \equiv -\frac{T^\mu{}_\mu}{c^2}.$$

The action should be read as the simplest weak-field continuum realization of the ontology already stated in Part I and the coefficient chain already derived in Part II. The metric sector remains the familiar Einstein one at low energy, but it is now interpreted as the continuum capacity geometry of the substrate rather than as an independent starting theory. It is coupled to a scalar field that tracks available entanglement capacity and to a source channel written in ordinary stress-energy notation while still being interpreted microscopically as the defect sector of the same medium.

At the EFT level χ is written in ordinary stress-energy language, but ontologically it is the coarse trace channel of the localized defect sector. Here γ is the continuum stiffness fixed by the UV chain, while κ is the defect-entropy coupling fixed by the canonical source map,

$$\kappa = \frac{\Xi_\rho}{L_*^2 \kappa_m(L_*)},$$

and λ controls the background branch. The source map is determined by matching the scalar mode on the lattice and in the continuum. Appendix C fixes the rigid defect amplitude through the isotropic defect benchmark $\Delta S_{\text{def}} = \ln(7/6)$ and the exact tetrahedral on-site Green constant $G_{\text{tet}}(0) = 0.448220394\dots$ With defect-entropy density

$$\sigma_{\text{def}} = \frac{\rho}{\kappa_m(L_*)},$$

the Green-matched continuum projection is

$$\nabla^2 \delta S = -\frac{3L_*}{4G_{\text{tet}}(0)} \sigma_{\text{def}},$$

and comparison with $\nabla^2 \delta S = -(\kappa/\gamma)\rho$ gives

$$\frac{\kappa}{\gamma} = \frac{3L_*}{4G_{\text{tet}}(0)\kappa_m(L_*)}.$$

This is the source-side counterpart of the stiffness derivation: the same tetrahedral scalar mode that carries the edge stiffness is the mode sourced by localized mass defects. The fixed-epoch normalization against S_∞ still enters the final gravity dictionary, but it is not an additional source freedom; S_∞ and κ rescale together under a change of entropy units, while $\kappa/(\gamma S_\infty)$ is invariant. Local weak-field dynamics are studied in the renormalized branch

$$\lambda_{\text{ren}} \equiv \lambda + \gamma \square S_{\text{bg}} = 0,$$

so that local perturbations are sourced only by the defect sector, written at continuum level in ordinary matter variables.

The Einstein–Hilbert coefficient is written here in the already-matched Einstein normalization of the weak-field EFT. This is the metric normalization in which weak-field gravity is observed, not a separate microscopic parameter added on top of the entanglement chain.

Three roles of G in the manuscript. There are three distinct uses of G in what follows, and it is useful to keep them separate. First, the weak-field metric normalization gives the identity

$$G = \frac{c^2 \kappa}{8\pi \gamma S_\infty}.$$

This is the gravitational scale implied by the scalar source coupling, scalar stiffness, and capacity-to-lapse bridge. The second use is the non-gravitational scale-setting route. The faithful sector-resolution principle fixes

$$L_* = \frac{3}{2} \lambda_e e^{-7g_{\text{share,eff}}}, \quad \lambda_e = \frac{\hbar}{m_e c},$$

and therefore induces

$$G_* = \frac{c^3 L_*^2}{\hbar} = \frac{9}{4} \frac{\hbar c}{m_e^2} e^{-14g_{\text{share,eff}}}.$$

The electron anchor enters this route through the reduced Compton length and through the identification of the electron as the lightest one-bit fermionic defect, while the mass–entropy anchor still fixes $\kappa_m(\lambda_e) = m_e/\ln 2$ for the source map. Third, formulas involving the conventional Planck length,

$$L_P = \sqrt{\frac{\hbar G}{c^3}},$$

are matched representations after the gravitational scale has been identified. They are useful for compactness, comparison, and black-hole thermodynamic notation, but the scale-setting argument runs through L_* , not through a Planck-cell rewriting.

The claim, then, is that the weak-field gravitational normalization is subject to a nontrivial consistency check: the same substrate that fixes the scalar stiffness also fixes a non-gravitational cell scale, while the scalar bridge and Green-matched source map give the observed weak-field normalization through the invariant ratio $\kappa/(\gamma S_\infty)$. Faithful sector resolution derives G_* as the gravitational scale induced by the substrate length. The matched weak-field bridge reads

$$G = \frac{c^2}{8\pi} \frac{\kappa}{\gamma S_\infty},$$

so the comparison between G and G_* is a coherence test. The matched route is built from L_* and G_* , so it is not a second measurement from independent inputs; what it checks is whether that one substrate scale propagates through the scalar stiffness, the Green-matched source map, and the weak-field bridge to the observed normalization. The genuinely non-gravitational result is the induced scale G_* itself, fixed from the electron Compton length and the tetrahedral entropy with no value of G used as input; the comparison has content because fixing L_* does not on its

own guarantee that κ/γ , S_∞ , and the bridge organize around it to give the right normalization. The dressing dynamics behind the faithful sector-resolution principle is derived in Section 22 and Appendix H. This is the simplest covariant realization of the closure chain: one metric, one scalar entanglement field, one trace-equivalent defect-source channel, and one renormalized background branch.

The logical order matters. Once this weak-field covariant form is accepted as the correct low-energy language of the substrate, the earlier UV closure chain fixes the entanglement-side coefficients that appear in it; the terms in the action are then read off from that chain, not chosen term by term to fit phenomenology.

The weak-field action used below is now fully written.

11. Field Equations and Bridge Law

Varying the action with respect to S_{ent} gives the sourced scalar equation

$$\gamma \square S_{\text{ent}} = \lambda + \kappa \chi.$$

Varying with respect to the metric yields

$$G_{\mu\nu} = \frac{8\pi G}{c^4} \left(T_{\mu\nu}^{(\text{matter})} + T_{\mu\nu}^{(\text{ent})} \right),$$

with the canonical scalar stress-energy induced by the entanglement sector.

These two equations separate the two jobs played by the scalar. The scalar equation tells us how the local entanglement-capacity variable responds to defect sources. The metric equation tells us how that scalar response then contributes back to spacetime curvature. The bridge law below is what turns those two statements into an ordinary weak-field gravitational potential.

The weak-field bridge law is not inserted as an arbitrary interpolation. It is the normalization map between fractional capacity depletion and the weak-field lapse. Let

$$\epsilon(x) \equiv \frac{\delta S(x)}{S_\infty}.$$

Vacuum normalization requires the lapse to approach its asymptotic value when $\epsilon = 0$. Locality requires the leading metric response to depend on the local fractional deficit. Independent small deficits must add at leading order, while independent redshift factors multiply. These requirements force the lapse to take exponential form,

$$N(\epsilon) = \exp[-C\epsilon],$$

with a single constant C . The standard weak-field metric normalization,

$$N = 1 + \frac{\Phi}{c^2} + O(\Phi^2/c^4),$$

fixes $C = 1/2$. Hence

$$N = \exp\left[-\frac{\delta S}{2S_\infty}\right] = 1 - \frac{\delta S}{2S_\infty} + O(\delta S^2/S_\infty^2),$$

and therefore

$$\frac{\Phi}{c^2} = -\frac{\delta S}{2S_\infty}.$$

In the static weak-field branch the emergent Newton constant is therefore

$$G = \frac{c^2 \kappa}{8\pi\gamma S_\infty}.$$

This bridge law is where the derivation stops speaking only in the language of entropy variables and starts speaking directly in the language of observable gravity. Without it the theory would remain a scalar model with an entanglement interpretation. With it, the deficit field acquires a unique weak-field normalization in terms of the ordinary gravitational potential.

With this bridge in place, the weak-field continuum dictionary is complete.

12. Newtonian Gravity and the Point-Source Limit

In the renormalized static weak-field sector the scalar equation reduces to

$$\nabla^2 \delta S = -\frac{\kappa}{\gamma} \rho.$$

This is the point where the micro-to-macro chain becomes operationally familiar. Once the background is renormalized away and the source is nonrelativistic, the scalar sector obeys an ordinary Poisson equation for the deficit field. The unusual quantity is δS , but the mathematical structure is the same one that underlies standard weak-field gravity.

For a point source M ,

$$\delta S(r) = \frac{\kappa M}{4\pi\gamma r}.$$

Using the bridge law,

$$\frac{\Phi}{c^2} = -\frac{\delta S}{2S_\infty},$$

the gravitational acceleration becomes

$$g(r) = \frac{c^2 \kappa}{8\pi\gamma S_\infty} \frac{M}{r^2} = \frac{GM}{r^2}.$$

Thus Newtonian gravity is recovered as the weak-field response of the entanglement-capacity medium.

Nothing qualitatively exotic has to be inserted at the last step to recover ordinary gravity. The same sourced scalar equation and the same bridge law already imply the familiar point-mass force law. In that sense Newtonian gravity appears here not as a starting axiom but as the first infrared limit of the entanglement medium.

Interpretation. Ordinary gravity is the small-deficit, weak-curvature limit of the extended entanglement restructuring around localized defects. The Newtonian $1/r^2$ law is therefore emergent, not fundamental.

Once the bridge law, Green-matched source projection, and UV stiffness are fixed, the Newtonian limit is closed.

13. Electron Anchor and the Mass–Entropy Relation

The mass–entropy relation requires a clean elementary anchor because, in this framework, the elementary matter sector is the localized defect sector itself. For a single fermionic face-exclusion defect the canonical increment is

$$\Delta S_f = \ln 2.$$

The physical content is that a single excluded face is a binary occupied/unoccupied topological defect and therefore carries exactly one bit of missing entanglement. At the electron Compton scale $\ell = \lambda_e$ the mass–entropy map gives

$$\kappa_m(\lambda_e) = \frac{m_e}{\ln 2}.$$

This is the first step of the anchor logic. If the elementary fermionic defect carries $\Delta S_f = \ln 2$, then dividing the electron mass by that fixed entropy increment gives the mass-per-entropy conversion at the electron’s own scale. The electron is the cleanest place to do this because it is the lightest simple fermionic defect and is not obscured by hadronic compositeness.

The UV normalization is

$$\kappa_{m,\text{UV}} = \frac{\hbar}{cL_*} \frac{1}{\ln 2},$$

and the canonical running law is

$$\kappa_m(\ell) = \kappa_{m,\text{UV}} \left(\frac{L_*}{\ell} \right)^{1+\alpha_{\text{cl}}}, \quad \alpha_{\text{cl}} = 0$$

in the closed branch. The electron sector supplies a sharp anchor for the mass–entropy map and a fixed-normalization input to the source projection. The next step is then to run that conversion back to the UV scale. The quantity $\kappa_{m,\text{UV}}$ is the fundamental mass-per-entropy conversion attached to the cutoff cell itself, and the running law tells us how that conversion appears at longer physical scales. So the logic of the section is: one bit fixes the electron-scale conversion, the running law connects the electron scale back to the UV scale, and the same conversion then feeds the weak-field normalization chain.

The same electron also supplies the non-gravitational length anchor used in faithful sector resolution. Its reduced Compton wavelength λ_e fixes the absolute cell length L_* through the seven-sector history-space dressing relation in Appendix D, while its mass m_e fixes the mass-per-entropy map through $\kappa_m(\lambda_e) = m_e/\ln 2$. This does not mean that an electron is a single tetrahedral cell carrying seven simultaneous labels. The local tetrahedral ensemble supplies the seven-sector alphabet; the electron is the lightest coherent one-bit defect whose ground-state dressing resolves those sectors once and exports the transverse share of the resulting support. These are distinct uses of the same elementary defect: the Compton scale calibrates the substrate length hierarchy, and the one-bit mass anchor calibrates the source map. Together these two uses impose a nontrivial consistency requirement on the electron’s role in the weak-field gravitational normalization without duplicating a single input.

For composite hadrons the claim is different. The relevant quantity is the dressed vacuum-subtracted bound-state entropy,

$$m_{\text{hadron}} = \kappa_m(\ell_H) S_{\text{ent},H}^{\text{dressed}},$$

with the dressed entropy budget generated by confinement, gluonic structure, trace-anomaly dynamics, and chiral vacuum reorganization. A finished lattice derivation of that dressed entropy is not yet available. What is claimed here is structural compatibility between the mass–entropy map and the standard QCD mass budget.

The elementary-fermion anchor is settled in the simple sectors, while the hadronic sector remains structurally compatible but not yet fully coefficient-complete.

14. Galactic Dynamics

The galactic sector is one of the main payoffs of the coefficient chain. Its characteristic acceleration scale, the radial-acceleration relation, and the deep-MOND limit all follow from one chain

that uses no per-galaxy interpolation function. The chain has three structural inputs: a stable bosonic vacuum mode (Appendix D.6), the Gibbons–Hawking thermal structure of that vacuum at the apparent horizon (Appendix E), and the 1 + 2 channel decomposition that places the cosmic horizon scale in the transverse sector.

Vacuum-state origin of the Bose–Einstein occupancy. The Bose–Einstein statistics in the galactic mode sector follow from the vacuum-state thermality of the entanglement field at the apparent horizon, with the cosmological boundary normalization fixed in Appendix E. The mechanism is therefore equilibrium horizon thermality, not dynamical thermalization.

Appendix D.6 establishes that δS fluctuations around the on-shell background constitute a massless bosonic scalar at quadratic order with positive kinetic stiffness $\gamma > 0$. For a massless bosonic scalar in a spacetime with apparent horizon R_A , the vacuum state restricted to the static patch is thermal at the de Sitter / apparent-horizon temperature

$$k_B T_H = \frac{\hbar c}{2\pi R_A},$$

with mode occupancy $n_B(\epsilon) = 1/(e^{\epsilon/k_B T_H} - 1)$. This is a Gibbons–Hawking statement of QFT in curved spacetime, fixed by the same horizon geometry that sets S_∞ .

It is useful to write a horizon temperature as an acceleration scale,

$$a_T \equiv \frac{2\pi c}{\hbar} k_B T.$$

For the apparent-horizon vacuum this gives

$$a_T = \frac{c^2}{R_A} \equiv a_H.$$

In the closed transport branch, $R_A^{-1} = H_0/c$ at the present epoch (from §17 with $\tau_0^{-1} = H_0$), so

$$a_H = \frac{c^2}{R_A} = cH_0.$$

The acceleration scale defined by horizon thermality is therefore fixed by the same closure chain that fixes S_∞ . More generally $a_H(t) = cH(t)$ at any cosmological epoch, so the horizon thermal scale is epoch-dependent. The present-epoch value cH_0 is the one relevant to the local ($z \approx 0$) galactic RAR.

The MOND scale a_0 is related to a_H by the microstructural sharing factor that connects $g_{\text{share,eff}}$ to phase-space transverse-mode density in the 1 + 2 channel decomposition:

$$a_0 = \frac{g_{\text{share,eff}}}{4\pi^2} a_H = \frac{cH_0 g_{\text{share,eff}}}{4\pi^2}.$$

Hence a_0 is the horizon thermal acceleration $a_H = cH_0$ reduced to the transverse scalar sector of the 1 + 2 decomposition. The reduction factor $g_{\text{share,eff}}/(4\pi^2)$ is fixed, not fitted. Its $(2\pi)^2$ is the canonical two-dimensional mode density $d^2 k_\perp / (2\pi)^2$ — the number of states per phase-space cell of a continuum two-plane, a quantization normalization rather than a Fourier convention — and $g_{\text{share,eff}}$ is the admissibility-sharing content carried per cell. The normalization is $(2\pi)^2$, not the spherical 4π , because the weak-field response samples the local transverse two-plane orthogonal to the baryonic forcing, not the full horizon sphere. That the split exists at all is kinematic, not a modeling choice: the baryonic gradient picks the longitudinal direction $\hat{n} = \nabla\Phi_{\text{bar}}/|\nabla\Phi_{\text{bar}}|$, and its orthogonal complement in three-space is two-dimensional by dimension count, so the transverse two-plane and its $(2\pi)^2$ cell are forced. The one premise that is physical rather

than geometric is the assignment of the low-acceleration response to the *stationary* transverse vacuum branch rather than the telegrapher transport sector of Section 17 — the step that ties a_0 to cH_0 rather than to a transport timescale, and that is testable through the epoch dependence $a_0(z) \propto H(z)$. Given it, the coefficient follows, and a_0 and cH_0 share one origin.

1 + 2 channel decomposition and the radial-acceleration law. The longitudinal acceleration scale is g_{bar} and the transverse vacuum reference scale is a_0 . For a longitudinally forced isotropic Gaussian transverse sector with longitudinal forcing amplitude $f_{\parallel} \propto g_{\text{bar}}$ and transverse vacuum variance $\sigma_{\perp}^2 \propto a_0$, the cross-coupling response at quadratic order is the geometric mean of the two scales: the response kernel is symmetric under longitudinal–transverse exchange at the cross-coupling order, so the cross-scale amplitude is

$$a_{\text{eff}} = \sqrt{g_{\text{bar}} a_0}.$$

The dimensionless bosonic occupancy argument is therefore

$$x = \frac{a_{\text{eff}}}{a_0} = \sqrt{\frac{g_{\text{bar}}}{a_0}}.$$

When $g_{\text{bar}} \gg a_0$ the system reduces to the ordinary Newtonian branch; when $g_{\text{bar}} \ll a_0$ the response crosses into the low-acceleration completion.

The resulting radial-acceleration law is

$$g_{\text{obs}} = g_{\text{bar}} (1 + n_B(x)) = \frac{g_{\text{bar}}}{1 - \exp\left(-\sqrt{g_{\text{bar}}/a_0}\right)},$$

with the asymptotic limits

$$g_{\text{bar}} \gg a_0 \implies g_{\text{obs}} \approx g_{\text{bar}}, \tag{10}$$

$$g_{\text{bar}} \ll a_0 \implies g_{\text{obs}} \approx \sqrt{a_0 g_{\text{bar}}}. \tag{11}$$

The deep-MOND branch gives the baryonic Tully–Fisher law

$$v^4 \approx a_0 GM_b.$$

The law expresses vacuum-state mode occupancy at the apparent-horizon temperature. Departures from it would require either (a) an excitation mechanism beyond the vacuum, supplied by the causal nonequilibrium transport sector of Appendix E for mergers and transients, or (b) a different transverse reference, which would require modifying the cosmological boundary normalization that fixes a_H . Neither route belongs to the closed stationary weak-field branch; both would be treated as extensions or departures from it.

The galactic branch is therefore fixed by the same channel-identification structure already used elsewhere in the weak-field EFT, with the radial-acceleration law read off from one chain rather than interpolated per galaxy.

15. Lensing, PPN, and Weak-Field Consistency

Because the entanglement sector is scalar, it does not generate anisotropic stress at leading weak-field order. The scalar anisotropic stress is quadratic in gradients, schematically $\partial_i S \partial_j S = O(\Phi^2/c^4)$, so it enters beyond the linear branch. Hence

$$\Phi = \Psi$$

to the order treated in the present EFT. This means that light bending and dynamical mass estimates are sourced by the same leading metric response. In effective-halo language, the entanglement response can be rewritten as

$$\rho_{\text{halo}}(r) = \frac{1}{4\pi Gr^2} \frac{d}{dr} \left[r^2 (g_{\text{obs}} - g_{\text{bar}}) \right],$$

which yields the familiar $1/r^2$ outer-halo profile in the asymptotic branch.

This matters because a theory can match galactic rotation curves and still fail lensing if the two metric potentials slip apart. The weak-field branch here avoids that problem at leading order. The same response that governs the dynamics also governs light deflection, so the theory is not buying galactic support at the price of a leading weak-field inconsistency. Solar-system constraints such as Cassini require any PPN slip to remain at the $\sim 10^{-5}$ level, consistent with the no-slip structure of the linear branch.

The same weak-field structure also returns the GR post-Newtonian values at the order treated. At leading post-Newtonian order,

$$\gamma_{\text{PPN}} = \beta_{\text{PPN}} = 1.$$

Scalar-induced corrections to the slip $\Phi - \Psi$ enter only at quadratic weak-field order, schematically $O(\Phi^2/c^4)$. Thus the leading weak-field EFT does not purchase galactic phenomenology by introducing gravitational slip or obvious solar-system-scale pathologies.

At leading weak-field order this sector is closed. Higher-order precision confrontation remains an audit task rather than an architectural gap.

Part IV. Time-Dependent, Transport, and Cosmological Sectors

16. Why Dynamics Requires Extension Beyond the Static Branch

The static weak-field branch is not the whole theory. If the entanglement-capacity medium is physical, it must admit relaxation, propagation, and causal response to changing sources. The time-dependent sector should therefore not be read as an optional add-on. It is the natural dynamical extension of the same medium that produces the static weak-field EFT.

This section is only a bridge into the dynamical sectors; no independent closure claim is being made here.

17. Causal Transport and Telegrapher Dynamics

The canonical time-dependent completion is most cleanly written relative to the substrate four-velocity u^μ :

$$\tau_0 (u^\mu \nabla_\mu)^2 \delta S + u^\mu \nabla_\mu \delta S = D h^{\mu\nu} \nabla_\mu \nabla_\nu \delta S + A \chi, \quad h^{\mu\nu} = g^{\mu\nu} + u^\mu u^\nu.$$

The vector u^μ is the local rest frame of the entanglement-capacity medium, not an additional ad hoc force carrier. In that frame, $u^\mu = (1, 0, 0, 0)$, the equation reduces to the familiar telegrapher form

$$\tau_0 \partial_t^2 \delta S + \partial_t \delta S = D \nabla^2 \delta S + A \chi(x, t),$$

with static-matching condition

$$\frac{A}{D} = \frac{\kappa}{\gamma}.$$

This equation is introduced because a physical medium should not respond instantaneously to changing sources. The static Poisson equation is appropriate when the source has already settled,

but once sources evolve in time one needs both propagation and relaxation. The telegrapher form is the minimal causal extension that still reduces to the static branch when time dependence becomes negligible.

Causality requires

$$\frac{D}{\tau_0} = c^2,$$

so the transport sector propagates disturbances at finite speed. In the canonical no-new-IR-scale branch,

$$\tau_0^{-1} = H_0, \quad D = \frac{c^2}{H_0}.$$

This transport equation separates two roles that must remain distinct. Ordinary galactic support belongs to the near-stationary static branch. The telegrapher sector governs how the same medium propagates, relaxes, and develops lag when sources evolve in time. It is therefore not used to generate ordinary static galactic support; it governs transport, lag, relaxation, and merger phenomenology around the near-stationary weak-field branch.

For galactic modes the Appendix E analysis shows that the long relaxation time does not destroy the static limit. Galactic modes lie deep in the underdamped regime, so the static Poisson branch is recovered as the exact time average relevant to ordinary galactic dynamics. The assumption here is that the source is quasi-static on galactic timescales and supported on wavelengths far shorter than the critical scale $\lambda_c \sim 4\pi c/H_0$; under those conditions the oscillatory transient averages out instead of competing with the static branch.

The static branch still handles the ordinary galactic law. The transport equation describes what happens when the source history is no longer quasi-static: propagation delay, relaxation, and merger-era lag.

The transport branch is closed at the level of $D/\tau_0 = c^2$ and the preferred choice $\tau_0^{-1} = H_0$. The cluster and merger phenomenology, which sets the source weights this transport then evolves, is developed in Section 17.5.

17.5 Cluster Source Projection and the Diffuse–Decoupled Channel Split

Galaxy clusters are the natural stress test of any baryon-rooted completion of the weak-field law. Treated as a direct extrapolation of the static galactic radial-acceleration relation, the cluster sector would inherit the familiar MOND-like residual: the lensing and dynamical mass of relaxed clusters exceeds the prediction obtained by applying the galaxy RAR to the observed baryons by a factor of order 1.5–2. Merging clusters sharpen the problem. In systems of the Bullet type the dominant baryonic component—the intracluster plasma—is ram-pressure slowed and displaced from the collisionless galaxies, while the lensing peaks remain with the outgoing collisionless components rather than the X-ray gas. A viable cluster sector must explain both the relaxed normalization and the merger gas/lensing separation without assigning galaxies additional hidden baryonic mass and without altering the galactic mass anchor.

The mechanism is not transport lag and it is not extra galaxy mass; both are excluded. Transport lag is excluded because, in the canonical $\tau_0^{-1} = H_0$ branch, disturbances propagate at c and communicate changes across a cluster-scale configuration on a light-crossing time of order 1 Myr, three orders below the merger timescale, so the field tracks the moving source; an underdamped configuration retains memory of the *pre-merger centroid*, not of the outgoing collisionless component, and a retarded forward solution sourced equally by all baryons therefore does not by itself keep the lensing peak on the outgoing collisionless component. Extra galaxy

mass is excluded because the galactic acceleration scale and the gas-dominated-dwarf RAR fix the deficit per unit baryonic mass with no freedom to enhance it. The viable mechanism is instead a phase-dependent *source projection* for the long-range entropic-excess channel, built from a distinction the framework already contains.

The projection coefficient as the galactic reduction factor. Section 14 fixes the galactic acceleration scale as the transverse reduction of the horizon thermal scale,

$$a_0 = \frac{g_{\text{share,eff}}}{4\pi^2} a_H, \quad a_H = cH_0,$$

where the $(2\pi)^2$ is the two-transverse-mode phase-space cell of the 1 + 2 split and $g_{\text{share,eff}}$ is the admissibility-sharing content per cell. The same construction fixes the ratio between the transverse static projection and the full horizon projection,

$$\epsilon \equiv \frac{a_0}{a_H} = \frac{g_{\text{share,eff}}}{4\pi^2} \simeq 0.188.$$

This is not a new coefficient. It is the already-fixed reduction factor of the galactic sector, read as a statement about source projection: a source that accesses only the transverse static two-plane couples to the long-range channel at strength ϵ of a source that accesses the full horizon projection.

The diffuse and decoupled source classes. The assignment follows from how source contributions combine under coarse-graining. Diffuse matter is a continuum of locally uncorrelated, thermalized source elements; under coarse-graining its off-diagonal source cross-terms average away and only the diagonal transverse static projection survives. It therefore couples at ϵ . This includes shocked or unvirialized intracluster plasma, the warm-hot intergalactic medium, and cold but *diffuse* galactic HI when the galaxy is treated as a single smooth source—which is why gas-dominated dwarfs and low-surface-brightness galaxies remain on the galaxy RAR. The suppressed projection is thus a coarse-graining effect, not a temperature effect; a hot phase that has *virialized* into a collectively organized bath is the exception, taken up below.

The unsuppressed projection is accessed by matter that is *not* part of the phase-averaged continuum: a collisionless overdensity that is spatially separated from, and dynamically decoupled from, a surrounding diffuse medium. The criterion is relational, not intrinsic compactness. Compactness alone would misclassify: stars in ordinary galaxies, isolated ellipticals, and globular clusters are compact and bound yet must remain on the galaxy RAR, and they do, because none is a collisionless node decoupled from a distinct diffuse continuum. A galaxy in a cluster is different only because it is embedded in, and decoupled from, the intracluster medium. The suppression is the property of participating in the continuum; matter that has decoupled from the continuum escapes it. So decoupled collisionless matter sits at the baseline weight, and incoherent diffuse gas is suppressed to ϵ , with $\epsilon = g_{\text{share,eff}}/4\pi^2$. A *virialized* bath is the third state. Once the diffuse atmosphere relaxes into a collectively organized, extended phase it is no longer a sum of uncorrelated elements, and it opens a collective long-range response that lifts it *above* the baseline by $(1 - \epsilon)B_{\text{bath}}$, where $B_{\text{bath}} \in [0, 1]$ measures how completely the bath has virialized. The effective entropic-channel source χ_{ent} is therefore regime-dependent. In a relaxed cluster the gas is a virialized bath,

$$\chi_{\text{rel}} = \rho_{\text{dec}} + [1 + (1 - \epsilon)B_{\text{bath}}] \rho_{\text{bath}}$$

with ρ_{dec} the decoupled collisionless substructure and ρ_{bath} the virialized continuum; in a non-equilibrium merger the central gas is shocked and incoherent while only a residual atmosphere stays virialized,

$$\chi_{\text{merge}} = \rho_{\text{dec}} + \epsilon \rho_{\text{shock}} + [1 + (1 - \epsilon)B_{\text{bath}}] \rho_{\text{vir}}$$

which in the Bullet limit, where the displaced gas is shocked and little virialized bath remains on the cores, reduces to $\chi_{\text{Bullet}} \simeq \rho_{\text{dec}} + \epsilon \rho_{\text{shock}}$. Ordinary matter continues to gravitate through the usual metric coupling; the projection rule concerns only the long-range entropic-excess channel.

The bath gate as a measured input. Decoupling requires something to decouple from. The gate is therefore tied to a measured property of the diffuse atmosphere: the fraction of the halo’s cosmic baryon allotment that has become an extended virialized hot phase,

$$B_{\text{bath}} = \text{clip}_{[0,1]} \left[\frac{M_{\text{hot,vir}}(< r_{500})}{f_{b,\text{cos}} M_{500}} \right] \quad f_{b,\text{cos}} = \frac{\Omega_b}{\Omega_m} \simeq 0.156,$$

with $f_{b,\text{cos}}$ fixed by Planck values and $M_{\text{hot,vir}}$, M_{500} read from X-ray/SZ and total-mass estimates. This closes the only candidate free parameter: ϵ is fixed and derived, $f_{b,\text{cos}}$ is fixed cosmologically, and the per-system inputs (f_{\star} , f_{gas} , $M_{\text{hot,vir}}$, M_{500}) are all measured. The sector carries no per-system adjustable knob. In merging systems $M_{\text{hot,vir}}$ is to be evaluated from the *pre-merger* virialized atmosphere rather than the mid-collision gas state, since the shocked, displaced gas is not a relaxed continuum.

Relaxed-cluster residual. For a *relaxed* cluster the diffuse gas is a virialized continuum, and the residual is carried by that continuum, not by the decoupled stellar component. Let

$$f_{\text{cont}} = \frac{M_{\text{hot,vir}}}{M_{\text{baryon}}}$$

be the fraction of observed baryons in the virialized diffuse continuum. The predicted residual relative to a galaxy-RAR extrapolation is

$$\mathcal{R}_{\text{rel}} = 1 + (1 - \epsilon) B_{\text{bath}} f_{\text{cont}} \quad 1 - \epsilon = 1 - \frac{g_{\text{share,eff}}}{4\pi^2} \simeq 0.812.$$

The amplitude $(1 - \epsilon)$ is the part of the full horizon channel that the suppressed transverse branch is missing; it is not a new coefficient. The physical reading is that as a virialized diffuse bath forms, the continuum itself opens a collective cluster response whose amplitude scales with how complete the bath is (B_{bath}) and how much of the baryon budget sits in it (f_{cont}).

This form has the correct mass trend, which is the nontrivial check. Both B_{bath} and f_{cont} *increase* with halo mass—hot-gas fractions rise and the atmosphere becomes more fully virialized toward clusters—so their product rises monotonically from groups to massive clusters. Evaluated across a representative group-to-cluster sequence, \mathcal{R}_{rel} rises monotonically from $\simeq 1.1$ for X-ray-faint poor groups through $\simeq 1.3$ at the Sanderson poor-cluster scale ($f_{\text{cont}} \simeq 0.73$, $B_{\text{bath}} \simeq 0.45$, $\mathcal{R}_{\text{rel}} \simeq 1.3$) to $\simeq 1.5$ – 1.6 for rich and massive clusters ($f_{\text{cont}} \sim 0.85$ – 0.90 , $B_{\text{bath}} \rightarrow 1$). This reproduces the observed direction—groups close to the galaxy RAR, massive clusters most discrepant—with no fitted parameter.

A sharp saturation falsifier. Because $B_{\text{bath}} \leq 1$ and $f_{\text{cont}} \leq 1$, the relaxed residual is bounded,

$$\mathcal{R}_{\text{rel}} < 1 + (1 - \epsilon) \simeq 1.81.$$

The rule therefore predicts that relaxed-cluster residuals saturate below $\simeq 1.8$. A robustly relaxed cluster with a residual well above this ceiling would falsify the closed form.

The decoupled-fraction form is excluded by the mass trend. An earlier candidate let the residual ride on the decoupled stellar fraction, $\mathcal{R} = 1 + 4.32 B_{\text{bath}} f_{\text{dec}}$ with $f_{\text{dec}} = f_{\star}/(f_{\star} + f_{\text{gas}})$. That form is excluded. Because f_{dec} falls with mass while B_{bath} rises, their product peaks at the group/poor-cluster scale and declines toward massive clusters, predicting the largest anomalies at intermediate mass. The observed residual instead rises monotonically from groups to massive clusters, where it is most firmly established. The decoupled-fraction form passes only at the intermediate scale where the spurious hump crosses the observed band, and is retired in favor of the continuum form above. The lesson is that the relaxed-cluster amplitude and the Bullet morphology are distinct observables and must not be collapsed into a single scalar law.

Bullet-type mergers. The relaxed residual above and the Bullet morphology are governed by the same fixed coefficient ϵ but by *different source expressions*, because the gas is in a different physical state. In a relaxed cluster the gas is a virialized continuum and drives the residual through the $(1 - \epsilon)B_{\text{bath}}f_{\text{cont}}$ term. In a Bullet-type merger the central gas is shocked, displaced, and out of equilibrium—not a relaxed bath—so it has small effective B_{bath} and couples to the long-range channel at the suppressed weight ϵ , while the outgoing collisionless cores remain decoupled. The two regimes share one coefficient; they are not one scalar law, and we do not present them as such.

In the merger, then, the decoupled galaxies and subcluster cores retain the unsuppressed projection and the shocked diffuse gas couples at ϵ . With a gas/galaxy baryon ratio near 5.7, the gas contributes $\epsilon \times 5.7 \simeq 1.07$ in the entropic channel against the galaxy contribution of 1.0: the projection brings the two components to near-parity, removing the factor ~ 5.7 by which the gas would otherwise dominate, but it does not by itself invert them. The inversion is completed by projected compactness. For two roughly symmetric outgoing components the ratio of one edge peak to the central gas contribution scales as

$$\frac{\Sigma_{\text{edge}}}{\Sigma_{\text{gas}}} \sim \frac{f_{\text{dec}}}{2\epsilon f_{\text{gas}}} \frac{A_{\text{gas}}}{A_{\text{edge}}}, \quad \frac{f_{\text{dec}}}{2\epsilon f_{\text{gas}}} \simeq 0.47,$$

so an edge peak dominates the projected map once the shocked gas is spread over more than about twice the projected area of a compact outgoing core—a condition the observed morphology satisfies by a wide margin. The projection rule and this geometry are therefore sufficient to make the observed gas/lensing inversion natural at the surface-density level—by projection and geometry together, not by projection alone and not by transport lag—as a spatial prediction rather than an integrated-mass argument. This is consistency, not yet a test of the coefficient. In standard flexible lens reconstructions the gas weight is degenerate with free halo and substructure components, so a model that fits comparably well with or without the fixed X-ray gas map constrains ϵ only weakly; Bullet-type mergers are thus consistent with the projection rule but do not yet measure ϵ . Peak location alone is in any case insensitive to the coefficient, since sufficiently broad shocked gas yields clump-centered peaks across a wide range of gas weights; the coefficient is tested only by the resolved amplitude fit below.

Relation to the transport sector. The telegrapher sector of Section 17 is not the cluster mechanism; it governs how the field propagates and relaxes once the source weights are set. The source-projection rule supplies the static weights χ_{ent} ; the transport sector then evolves them. This division is what avoids the failure mode of a transport-only account, in which a field sourced equally by all baryons cannot hold a lensing peak on the outgoing collisionless component. The full merger observable is obtained by evolving $\chi_{\text{ent}}(x, t) = \rho_{\text{dec}}(x, t) + [\dots]$ through the causal equation with the observed geometry as input.

Falsifiers and open status. The rule makes sharp predictions beyond the relaxed normalization. (i) The relaxed-cluster residual rises monotonically with halo mass and saturates below $\mathcal{R}_{\text{rel}} \simeq 1.81$; a robustly relaxed cluster above that ceiling, or a residual that peaks at intermediate mass and declines toward the most massive clusters, would falsify the closed form. (ii) At fixed mass, X-ray-bright bath-developed systems (larger B_{bath} , larger f_{cont}) should deviate more from the galaxy RAR than X-ray-faint systems; the residual turns on with the developed diffuse atmosphere, not with mass alone, so two systems of equal mass but different bath development should separate. (iii) The decisive test is the resolved lensing map. Because the relaxed and merger regimes use different source expressions, the convergence is a *three-component* channel-weighted map—a relaxed virialized-bath component, a shocked non-equilibrium continuum at the suppressed weight, and a decoupled collisionless component,

$$\kappa_{\text{obs}}(x, y) = A \left[(1 + (1 - \epsilon)B_{\text{bath}})\Sigma_{\text{bath}}(x, y) + \epsilon \Sigma_{\text{shock}}(x, y) + \Sigma_{\text{dec}}(x, y) \right] + b$$

with ϵ fixed at $g_{\text{share,eff}}/4\pi^2 \simeq 0.188$ or fit freely. In the relaxed limit this reduces, relative to the galaxy-RAR extrapolation with the baseline baryonic response normalized to unity, to the residual law $\mathcal{R}_{\text{rel}} = 1 + (1 - \epsilon)B_{\text{bath}}f_{\text{cont}}$; in the Bullet limit the central gas is shocked (Σ_{shock} at weight ϵ) and the peaks fall on Σ_{dec} . A best-fit ϵ near 0.19 across Bullet, MACS J0025, the Musket Ball, a relaxed CLASH sample, and X-ray-bright versus X-ray-faint groups would be a parameter-free success that neither Λ CDM nor total-baryon modified gravity predicts, while a best fit near 1 or near 0 would falsify the sector. This map fit is the principal outstanding empirical task and has not yet been performed. It is decisive only if it removes the escape hatch that makes existing reconstructions non-discriminating: ϵ must be floated against the channel-weighted *baryonic* maps (Σ_{bath} , Σ_{shock} , Σ_{dec}) with free dark haloes disallowed, since free haloes absorb the gas weight and leave ϵ unconstrained—which is why current Bullet reconstructions neither recover nor exclude $\epsilon \simeq 0.19$. The three-component decomposition makes the fit more demanding than a two-component one, but also more distinctive.

Three open items remain at the theory level. First, the relaxed residual and the Bullet morphology are governed by one fixed coefficient ϵ but by two regime-specific source expressions (virialized bath versus shocked non-equilibrium gas). This is justified by the differing physical state of the gas, but it is a weaker unification than a single scalar law, and the boundary between the regimes—how B_{bath} interpolates as a shocked atmosphere re-virializes—is not yet derived. Second, the channel-selection rule rests on one physical premise: that participation in a phase-averaged diffuse continuum restricts a source to the transverse projection, so that a collisionless overdensity decoupled from that continuum recovers the full horizon projection. The coefficient ϵ is derived; this *participation* criterion is plausible within the ontology but not yet derived from the microscopic source map, and supplying that derivation is the channel-selection theorem. Third, the identification of the decoupled component with stellar/galaxy mass and of the continuum with the gas is a coarse split; intracluster light and tidally stripped stars blur it at a level the resolved map fit would expose. Abell 520, whose reported gas-coincident dark core is itself disputed, is a phase-state stress case rather than a clean test: a re-virializing or quasi-bound central component would raise its effective B_{bath} and return lensing toward the gas, while a confirmed *young* merger with a robust gas-centered, galaxy-free lensing peak—where the re-virialization escape is unavailable on time grounds—would be a serious challenge.

This sector is therefore presented as a structured, falsifiable closure proposal, not as closed physics: the projection coefficient is derived, the bath gate is a measured input, the mass trend and saturation ceiling are parameter-free predictions, and the resolved-map test and the channel-selection theorem are the named open work.

18. Cosmology and the Hubble-Tension Sector

The cosmological sector should be read as the homogeneous continuation of the same scalar medium, not as an unrelated dark-energy add-on bolted onto the weak-field theory. What changes here is not the ontology but the kinematic regime: the background mode becomes dynamically relevant on horizon scales while the local weak-field branch remains encoded in the inhomogeneous fluctuations.

The cosmological sector uses the same field split,

$$S(x, t) = \bar{S}(t) + s(x, t),$$

where $\bar{S}(t)$ is the homogeneous mode and $s(x, t)$ the inhomogeneous sector responsible for local weak-field dynamics. The vacuum baseline is fixed by apparent-horizon capacity,

$$S_\infty(t) = \pi \frac{R_A(t)^2}{L_*^2}.$$

This is the horizon-normalized representation of the same entropy field used locally. It is compatible with the cell-normalized source theorem because local observables depend on $\delta S/S_\infty$ and $\kappa/(\gamma S_\infty)$ rather than on an absolute entropy unit.

This decomposition is essential. The homogeneous mode and the local weak-field fluctuations are not two unrelated scalar fields. They are two kinematic sectors of the same field. The split lets the background mode change cosmological evolution without automatically rewriting the local weak-field equations that already fixed the galactic phenomenology.

Because the entanglement field couples to the trace of the stress-energy tensor, the homogeneous mode is largely dormant during radiation domination but becomes active near matter–radiation equality. This gives a transient early-energy component of the same general type used in early-dark-energy resolutions of the Hubble tension. In the closed cosmological branch treated here, the effect reduces the sound horizon and shifts the CMB-inferred Hubble constant upward from the high-67 range toward the high-68 to low-69 range.

What matters here is not just the direction of the shift but the timing. A successful Hubble-tension mechanism must turn on near the right epoch, alter the sound horizon in the right direction, and then decouple cleanly enough from the local weak-field sector that the galactic branch is not spoiled. The entanglement medium has exactly that qualitative structure.

The local weak-field predictions are protected by the separation between $\bar{S}(t)$ and $s(x, t)$. This is the role of the shear-lock logic: changing the homogeneous background mode does not rewrite the local static Poisson branch that governs galactic dynamics and lensing.

The claim is therefore a mechanism with the right direction, timing, and qualitative separation of scales, not a finished precision cosmology package. What is shown is that the trace-coupled homogeneous mode turns on in the relevant epoch and pushes the sound horizon in the required direction; what remains open is the full perturbation propagation and likelihood-level confrontation. The homogeneous mode modifies the background history; the inhomogeneous branch continues to govern the local weak-field observables already fixed earlier in the derivation. That separation is what allows the cosmological extension to remain part of the same scalar medium rather than a re-tuning of the galactic sector.

This is a structurally supported and directionally successful extension, but it is not yet Boltzmann-closed.

Part V. Nonlinear, Interpretive, and Completion Sectors

19. Why These Sectors Belong

The most directly constrained micro-to-weak-field chain is now in hand. The next sectors develop three further pieces required for overall framework completeness: nonlinear completion, operational quantum reduction, and the underlying substrate dynamics. They belong to the same ontology, but they should not be read as resting on identical evidence.

20. Strong-Field Branch and Bounded Occupancy

The strong-field branch begins with the same variable that already appeared in the weak-field bridge, but it uses the form in which boundedness is explicit:

$$q(x) = \frac{S_{\text{ent}}(x)}{S_{\infty}} \in [0, 1].$$

The physical meaning is direct: $q = 1$ is undepleted vacuum capacity, $0 < q < 1$ is a partially depleted region, and $q = 0$ is complete local exhaustion of continuum-defining capacity. The strong-field question is therefore not how to continue the weak-field scalar after it becomes large, but how to write the continuum theory on the domain where the capacity variable remains physically allowed.

The local lapse associated with the static asymptotic time must be a function of surviving capacity. If

$$N^2 = f(q),$$

then vacuum normalization gives $f(1) = 1$, horizon normalization gives $f(0) = 0$, the weak-field bridge gives $f'(1) = 1$, and serial composition of independent capacity-reduction layers gives

$$f(q_1 q_2) = f(q_1) f(q_2).$$

The continuous solutions are $f(q) = q^{\alpha}$, and weak-field matching fixes $\alpha = 1$. Thus

$$N^2 = q$$

is the unique continuous multiplicative completion compatible with the weak-field branch. This is a stronger statement than a useful ansatz: once the capacity variable is adopted, the static lapse rule is fixed.

The constrained-capacity theory is then posed on the physical domain

$$\mathcal{M}_q = \{x \mid q(x) > 0\},$$

with a Lagrange multiplier enforcing $N^2 = q$. In bulk vacuum, away from the $q = 0$ boundary and with no explicit matter source, the multiplier vanishes and the equations reduce to the ordinary vacuum Einstein equations on \mathcal{M}_q . The difference from GR is not in the exterior vacuum equations; it is in the physical domain on which those equations are allowed to live and in the boundary condition at capacity exhaustion.

Consequently the static spherical asymptotically flat vacuum branch is the Schwarzschild exterior,

$$q(r) = N^2(r) = 1 - \frac{2GM}{c^2 r}, \quad r > r_h = \frac{2GM}{c^2}.$$

The continuum branch terminates at $q = 0$. A real continuation of this static branch to $r < r_h$ would require $q < 0$, which is outside the state space of the capacity variable. The classical Schwarzschild interior is therefore not interpreted as a physical continuation of the same

entanglement-capacity EFT. It is a formal continuation of the GR manifold beyond the point where the substrate has no remaining local continuum channels.

This preserves the tested exterior physics. The near-horizon Euclidean regularity argument gives the usual Hawking temperature, and exterior perturbation theory gives the usual absorption and ringdown: horizon regularity at $q = 0$ selects the ingoing mode at leading order, so the leading reflectivity vanishes. Echoes are correction-level rather than generic; they arise only if microscopic boundary channels carry finite UV reflectivity or if the effective reflection surface is displaced outward to a stretched layer $q = \epsilon > 0$. The entropy story closes within the same substrate inputs already used elsewhere: Postulate II identifies the per-channel cut entropy as $\ln 2$ via fermionic face exclusion, while the transverse bulk graph response fixes the channel density, and their product is the Bekenstein–Hawking $1/4$ as an exact identity. The geometric identification of q with the areal-radius gradient invariant $|\nabla R|^2$ in spherical symmetry further shows that the substrate q equals $1 - 2GM_{\text{MS}}/(c^2 R)$ on the exterior, so the $q = 0$ saturation surface coincides with the apparent horizon and is inherited from ordinary trapped-surface formation in GR collapse rather than from a separate substrate-transport postulate.

The strong-field branch is therefore closed at the level of universal predictions. Static and stationary asymptotically flat vacuum exteriors reduce to Schwarzschild, Kerr, Reissner–Nordström, or Kerr–Newman through the standard vacuum Einstein and Einstein–Maxwell branches; the horizon location is fixed as the marginal-capacity surface; the area-law coefficient is closed under Postulate II and the bulk graph response; and the leading reflectivity vanishes by horizon regularity. What remains is nonuniversal: the microscopic relaxation spectrum, possible stretched-layer corrections, transient response on the boundary, and the explicit graph-side consistency check on substrate transport during dynamical collapse.

21. Many-Pasts: Operational Reduction and Arrow of Time

Many-Pasts belongs in the full framework because the theory is not only a gravity mechanism. It is also a proposal about branch realization and temporal asymmetry on the same entropic substrate. Operationally, however, it is deliberately conservative.

With

$$P(H|P) \propto e^{-D(H,P)},$$

the Born rule is recovered exactly because

$$e^{-D(H,P)} = \text{Tr}(\Pi_P \rho_{H \rightarrow \text{now}})$$

in the projective laboratory limit. Exact Born recovery forces $\alpha = 1$, and forbidding any extra signaling-sensitive operational bias channel forces $\beta = 0$. No-signaling is preserved exactly in this operational branch.

The remaining content is interpretive and cosmological. The arrow of time is recovered through conditional typicality: among histories consistent with present macroscopic records, overwhelmingly many exhibit entropy growth toward the future direction defined by those records. This adds no new laboratory probability law; it offers a global consistency account of branch realization and temporal asymmetry.

Operational closure is exact in the laboratory sector; the extra content added here is interpretive and cosmological.

22. Microstructure Hamiltonian and Underlying Dynamics

The UV closure chain has a microscopic dynamical realization with two sides: the emergence of the continuum geometry from the substrate, and the defect dynamics that fixes the scale-setting principle of the weak-field chain. The first is the condensate-side compatibility check; the second closes the one microscopic principle the static branch was conditional on.

On the geometry side, the realization is a GFT/condensate picture in which spacetime emerges from a condensate of discrete tetrahedral building blocks, while what is macroscopically read as matter appears as fermionic defects of that same substrate. In Madelung form,

$$\sigma(x) = \sqrt{n(x)}e^{i\theta(x)},$$

the condensate hydrodynamics generically generate a positive scalar stiffness for the logarithmic-density variable, providing the condensate-side origin of the EFT kinetic term. This is a structural compatibility check, not a replacement for the explicit coefficient closure in Appendix C; it shows that the EFT kinetic term has a natural microscopic origin.

On the defect side, the same substrate supplies the dynamics behind the faithful sector-resolution principle that sets the absolute cell length L_* . The elementary fermionic defect is governed by a dressing Hamiltonian whose single-channel terms make the lightest charged defect bind all seven face sectors once — the one-pass ground state that supplies the exponent 7 — and whose inter-channel term controls whether the seven-channel dressing cloud factorizes. Factorization turns the seven equal sector contributions into the additive $7g_{\text{share,eff}}$, and needs two things: each channel must carry the full admissibility entropy, and the seven must be mutually independent. The first is memorylessness in substrate time — the Many-Pasts content of Postulate III, with the substrate relaxing to a fresh draw long before the electron is read out. The second follows from the electron’s own lightness: inter-channel correlation would contract the defect and raise its mass, so the lightest charged defect carries none. Each channel at full entropy and the channels independent give $\text{dim}_{\text{eff}} = e^{7g_{\text{share,eff}}}$ and $\lambda_e/L_* = \frac{2}{3}e^{7g_{\text{share,eff}}}$, worked out in full in Appendix H.

Together these two sides do more than check compatibility. The defect dynamics derives the faithful sector-resolution principle from the mass–entropy ontology and the electron anchor, so the substrate length and the induced gravitational scale follow from commitments the theory already makes. What remains microscopically open is the first-principles derivation of every inhomogeneous continuum coefficient from the full kernel.

Part VI. Closure Status, Falsifiability, and Research Program

23. Closure-Status Table

The closure bookkeeping is concentrated here in one place so the rest of the text can simply derive, state, and move on.

| Quantity / Claim | Sector | Status | Type of Support | Where Established |
|---|-----------------------|--------|---|-------------------|
| $\Omega_{\text{tet}}, g_{\text{share,max}}$ | UV counting | Closed | exact combinatorics | Part II, App. B |
| η_* | admissibility closure | Closed | stationary normalized closure evidence $\ln Z + \frac{3}{2} \ln \eta$ maximized on the exact K^2 spectrum; the 3/2 is the determinant weight of the three closure-defect components | Part II, App. B |
| $g_{\text{share,eff}}$ | UV entropy | Closed | exact weighted evaluation | Part II, App. B |

| Quantity / Claim | Sector | Status | Type of Support | Where Established |
|--|---------------------------|---|--|--------------------|
| Substrate length L_* (faithful sector resolution / support-to-length) | scale setting | Foundational identification; cross-validated, not yet derived | $\lambda_e/L_* = C_{cl}e^{\tau g_{\text{share,eff}}}$ identifies coherent support cardinality with the Compton scale; cross-validated by G_* , the lepton ladder, and mass minimization; one-defect transfer-operator derivation is the completion target | Part I, App. D, H |
| $J_{\text{bare}}, J_{\text{eff}}^{\text{tree}}$ | UV edge kernel | Closed | tetrahedral isotropy identity | Part II, App. C |
| $\Sigma_{\text{ret}} = 65/9$ | finite-loop UV | Fixed in the minimal UV return sector | explicit seven-channel plus singlet return count | Part II, App. C |
| $J_{\text{eff}}^{(\text{ren})}$ | finite-loop UV | Fixed by the UV return resummation | derived from Σ_{ret} | Part II, App. C |
| γ | continuum stiffness | Closed in canonical EFT convention | Euclidean-action normalization | Part II, App. C |
| Green-matched source projection | UV source map | Closed in the canonical weak-field branch | exact defect counting + tetrahedral on-site Green function | App. C |
| κ/γ | source-to-stiffness ratio | Closed in the canonical weak-field branch | $\sigma_{\text{def}} = \rho/\kappa_m(L_*)$ plus $G_{\text{tet}}(0)$ and tetrahedral 4/3 projection | Part III, App. C–D |
| Weak-field bridge law | EFT / gravity | Closed in canonical weak-field branch | uniqueness from multiplicative composition | Part III, App. D |
| G_* | gravitational scale | Conditional derivation at percent level | $G_* = (9/4)(\hbar c/m_e^2)e^{-14g_{\text{share,eff}}}$ | Part I, App. D |
| Matched G | weak-field gravity | Closed in the matched weak-field branch | bridge law + Green-matched κ/γ + fixed-epoch normalization of $\kappa/(\gamma S_\infty)$; comparison with G_* tests coherent propagation of the same L_* , not a disjoint-input second derivation of G | Part III, App. C–D |
| Electron anchor | mass / length sector | Fixed empirical elementary anchor | one-bit fermionic defect supplies λ_e for length setting and $m_e/\ln 2$ for mass–entropy map | Part III, App. D |
| Seven-sector additivity | UV scale theorem | Closed conditional on the support-to-length identification | $H_{\text{seven}} = 7g_{\text{share,eff}}$ from subadditivity plus mass minimization: full per-channel entropy from the memoryless dressing, channel independence forced because correlation would raise the defect mass ($\Delta = 0$ for the lightest one-bit defect) | App. D, H |
| a_0 | galactic EFT | Fixed in the closed weak-field realization | UV entropy \times horizon scale cH_0 , with $4\pi^2 = (2\pi)^2$ the canonical two-transverse-mode phase-space cell of the 1 + 2 split | Part III, App. C |
| RAR law | galactic EFT | Fixed in the closed weak-field realization | 1 + 2 channel geometry + bosonic occupancy | Part III, App. C |
| No slip / lensing consistency | weak-field metric | Closed at leading weak-field order | scalar-stress structure | Part III, App. D |
| PPN leading values | weak-field metric | Structurally supported | weak-field expansion | Part III, App. F |
| Telegrapher relation $D/\tau_0 = c^2$ | transport | Closed in canonical transport branch | causal closure | Part IV, App. E |
| Canonical $\tau_0^{-1} = H_0$ branch | transport | Fixed in the minimal transport closure | no-new-IR-scale choice | Part IV, App. E |
| Hubble-tension mechanism | cosmology | Structurally supported extension | homogeneous trace-coupled mode | Part IV, App. E |
| Diffuse source projection $\epsilon = g_{\text{share,eff}}/4\pi^2$ | cluster sector | Derived (no new coefficient) | same transverse reduction factor that fixes a_0/a_H (Section 14) | Part IV, §17.5 |
| Bath gate B_{bath} | cluster sector | Operationally closed; microscopic origin open | $M_{\text{hot,vir}}/(f_{b,\text{cos}}M_{500})$ from X-ray/SZ; no per-system fitted knob | Part IV, §17.5 |

| Quantity / Claim | Sector | Status | Type of Support | Where Established |
|---|---------------------|---|---|-------------------|
| Cluster residual $\mathcal{R}_{\text{rel}} = 1 + (1 - \epsilon)B_{\text{bath}}f_{\text{cont}}$ | cluster sector | Predicted; monotonic mass trend; decoupled-fraction form excluded | rises group→cluster, saturates < 1.81 ; decoupled-fraction (f_{dec}) form retired because its product humps at intermediate mass | Part IV, §17.5 |
| Relaxed vs. merger source expressions | cluster sector | Conditional — two regimes, one coefficient | residual rides on virialized bath; Bullet on shocked gas + decoupled clumps; regime boundary not yet derived | Part IV, §17.5 |
| Channel-selection rule (continuum participation) | cluster sector | Conditional — one open premise | coherence/phase-averaging argument; participation criterion not yet derived from the source map | Part IV, §17.5 |
| Bullet gas/lensing inversion | cluster sector | Structurally supported; consistent but untested (ϵ not yet measured) | ϵ brings gas/galaxy to near-parity, compactness completes the inversion; existing flexible reconstructions non-discriminating because gas weight is halo-degenerate | Part IV, §17.5 |
| Resolved cluster lensing-map test | cluster sector | Open — principal empirical task | three-component $\kappa \propto (1 + (1 - \epsilon)B_{\text{bath}})\Sigma_{\text{bath}} + \epsilon\Sigma_{\text{shock}} + \Sigma_{\text{dec}}$; predicted best-fit $\epsilon \simeq 0.19$, floated on baryonic maps with free dark haloes disallowed; not yet performed | Part IV, §17.5 |
| Bounded capacity q , $N^2 = q$ | strong field | Closed as constrained-capacity rule | uniqueness from multiplicative composition and weak-field matching | Part V, App. F |
| Constrained action on $\mathcal{M}_q = \{q > 0\}$ | strong field | Closed within the EFT | ADM action plus lapse-capacity constraint | Part V, App. F |
| Static spherical exterior | strong field | Closed within the constrained-capacity branch | bulk vacuum reduces to Einstein vacuum on \mathcal{M}_q | Part V, App. F |
| Capacity-exhaustion horizon $q = 0$ | strong field | Closed as static-domain boundary | Schwarzschild exterior terminates where q leaves the physical state space | Part V, App. F |
| Hawking temperature and exterior ringdown | strong field | Closed at leading EFT order | unchanged Schwarzschild exterior; absorbing boundary at $q = 0$ fixed by horizon regularity; echoes only as UV / stretched-layer corrections | App. F |
| Bekenstein–Hawking area-law bridge | strong field | Closed within the constrained-capacity EFT under Postulate II and the transverse graph response | per-channel cut entropy $\ln 2$ from fermionic face exclusion + channel density from $G_{\perp} = (2/3)G_{\text{tet}}(0)$; product $n_{\text{hor}}S_{\infty}^{\text{cell}} = 1/4$ as exact identity | App. C, F |
| Horizon formation and boundary microphysics | strong field | Formation closed geometrically; spectroscopy open | $q = \nabla R ^2 = 1 - 2GM_{\text{MS}}/(c^2R)$ in spherical symmetry, with $q = 0$ at the marginal-trapped-surface inherited from standard apparent-horizon formation; remaining work covers relaxation spectra, stretched-layer corrections, and transient response | Part V, App. F |
| Rotating / charged stationary exteriors | strong field | Closed within the constrained-capacity EFT | vacuum Einstein / Einstein–Maxwell reduction on \mathcal{M}_q gives Kerr / Reissner–Nordström / Kerr–Newman exteriors; $q = 0$ identified with outer Killing horizon; inner Cauchy horizons are non-physical analytic continuations into $q < 0$ | App. F |
| Many-Pasts Born recovery | quantum foundations | Closed operationally | $\alpha = 1$ theorem | Part V, App. G |
| No-signaling in operational branch | quantum foundations | Closed operationally | $\beta = 0$ theorem | Part V, App. G |
| Arrow-of-time account | quantum foundations | Coherent extension | conditional typicality / counting | Part V, App. G |
| Microstructure Hamiltonian | UV realization | Defect side closes scale-setting; geometry side coherent | dressing Hamiltonian with one-pass ground state, memoryless dressing, and minimal-mass channel independence (scale-setting); GFT/condensate origin of the kinetic term (geometry) | Part V, App. H |

| Quantity / Claim | Sector | Status | Type of Support | Where Established |
|-----------------------------|---------------------------|--|---|-------------------|
| Charged-lepton spectrum | particle-sector extension | Derived; coefficients forced, one map-form input | three-generation termination from K^2 permutation-invariance at $N = 3$; $m_N/m_e = 720^N (2/7)^{N^2}$ with $720 = 6!$, $2/7$ the singlet projection, N^2 the ordered-pair count; matches PDG at 0.5–0.7% with no fitted parameters | App. I |
| Gauge-redundancy extension | gauge sector | Coherent extension | baseline-redundancy construction with Maxwell/Yang–Mills form | App. I |
| Numerical robustness checks | validation layer | Supportive audit layer | cross-sector consistency tests | App. J |
| EFT consistency checklist | field-theory audit | Supportive audit layer | no-ghost / no-tachyon / causal-propagation checklist with explicit vacuum dispersion stability | App. D |

This table is the epistemic map used for the rest of the discussion.

24. Falsifiability and Observational Tests

24.1 Static weak-field falsifiers

The static weak-field sector stands or falls on a small number of concrete checks. The most direct are the shape and tightness of the galaxy RAR transition [1], the baryonic Tully–Fisher scaling in systems where the EFT should apply, and the weak-field lensing sector. A persistent need for gravitational slip where the scalar stress predicts none would be especially damaging, because it would break the same no-slip structure used to keep dynamics and lensing aligned. Solar-system bounds, especially Cassini-class PPN tests [3], also require the leading no-slip branch to survive at high precision.

24.2 Dynamical falsifiers

The dynamical extension is more vulnerable, and its failure modes are correspondingly sharper. Time-dependent halo lag, cluster-scale acceleration relations, merger offsets, or relaxation signatures that cannot be reconciled with the telegrapher relation $D/\tau_0 = c^2$ would indicate that the causal completion has the wrong propagation structure even if the static branch survives. Cluster data accordingly split into two tests of distinct parts of the framework: *source-projection* tests, which ask whether relaxed-cluster residuals track $B_{\text{bath}}f_{\text{cont}}$ and whether resolved merger maps prefer the projection coefficient $\epsilon \simeq 0.19$ (Section 17.5); and *transport* tests, which ask whether the causal propagation law $D/\tau_0 = c^2$ correctly evolves those source weights through a merger. Systems such as the Bullet Cluster [2, 4] probe both, and should not be treated merely as larger versions of the static galaxy problem.

24.3 Cosmological falsifiers

Cosmology presents a different kind of test. The question there is not whether the mechanism points in the right direction, but whether a full Boltzmann treatment allows the trace-coupled homogeneous mode to reduce the sound horizon without spoiling the CMB or structure-growth observables. If it cannot, the cosmological extension fails on its own terms.

24.4 Correlated-constant falsifiers

One of the more distinctive signatures of the framework is that the same microstructural chain feeds the substrate scale, the weak-field gravitational normalization, and the galactic acceleration scale. A precision program that could test the inferred L_* and G_* scale against matched G , a_0 ,

and the RAR normalization would probe the theory more sharply than isolated single-observable fits, because it would confront the shared coefficient origin directly.

24.5 Many-Pasts status

The Many-Pasts sector is not likely to be challenged first by ordinary laboratory deviations from quantum mechanics, because it is built to reproduce the usual operational structure there. Its more immediate points of failure are internal ones: failure of exact Born recovery, failure of no-signaling, or incompatibility with the thermodynamic arrow structure it is supposed to illuminate.

These points define the canonical falsifiability map.

25. What the Theory Would Have to Get Wrong to Fail

Placed together, the main failure modes have a simple shape:

- If future weak-field observations require persistent gravitational slip in the relevant galactic or cluster regimes, the canonical weak-field branch fails.
- If the RAR transition shape systematically departs from the derived bosonic occupancy law in systems well described by the static branch, the canonical galactic EFT fails.
- The seven-fold additivity behind the percent-level G_* scale prediction rests on the electron being the lightest one-bit charged defect, since inter-channel correlation in its dressing would raise the defect mass. If that identification or the mass–entropy ontology failed, the scale-setting derivation of Appendix H would not go through.
- If the weak-field UV coefficient chain cannot be reconciled with an independently validated microscopic derivation, the canonical UV closure loses support.
- If the cosmological trace-coupled homogeneous mode cannot survive full Boltzmann likelihood confrontation, the cosmology sector fails even if the static weak-field branch survives.
- If the constrained-capacity branch cannot retain the tested Schwarzschild exterior, standard absorbing-boundary ringdown, or the required horizon thermodynamics once its boundary action is derived, the strong-field completion fails while the weak-field EFT may still remain viable.

These are failure modes rather than a separate derivational sector.

26. Comparison with Other Approaches

Because the framework aims to replace dark matter and partially reorganize the usual dark-energy story, it is useful to state briefly how its logic differs from nearby alternatives.

26.1 Relative to Λ CDM

The contrast with Λ CDM begins at the level of ontology. Standard cosmology explains the relevant phenomenology by adding dark matter and an independent cosmological constant or dark-energy sector to otherwise standard gravity. Here the visible matter sector is retained, but it is interpreted as the macroscopic description of localized defects in a vacuum-capacity medium whose weak-field response supplies the effective extra gravitating component. The same closure chain is then asked to feed G , a_0 , the RAR law, weak-field lensing consistency, and the homogeneous cosmological mode.

26.2 Relative to MOND-like interpolation programs

MOND-like programs usually begin from an acceleration law or interpolation function and ask how much galaxy phenomenology it can explain. The present logic runs the other way. The interpolation law is not taken as primary; it is downstream of the UV entropy, the 1 + 2 channel geometry, and the bosonic occupancy branch. The galactic law is thus treated as an output of the same micro-to-IR closure chain rather than as the phenomenological starting point.

26.3 Relative to Verlinde-style emergent gravity

Verlinde-style emergent-gravity programs share the broad intuition that gravity may be entropic, but they are usually formulated at the level of thermodynamic reasoning or horizon-inspired force laws. The present framework is trying to do something narrower and more explicit: finite tetrahedral boundary counting, admissibility closure, edge coupling, finite renormalization, Euclidean-action normalization, and only then a continuum scalar EFT. Whether that chain is ultimately correct is an empirical matter, but it is a different kind of proposal from a purely macroscopic entropic argument.

26.4 Relative to TeVeS and other multi-field modified gravities

Multi-field relativistic MOND completions such as TeVeS typically introduce additional vector or tensor sectors in order to repair lensing or cosmological problems. The present weak-field construction instead keeps a single scalar entanglement field within a low-energy Einstein continuum sector that is itself interpreted as emergent from the substrate, and relies on the no-slip structure $\Phi = \Psi$ at leading order to keep lensing and dynamics aligned. That economy is attractive if the branch survives confrontation with data, and immediately vulnerable if future observations demand persistent slip or extra weak-field structure.

This comparison is meant as context rather than as a derivational sector.

27. Conclusion

The central achievement of the manuscript is a weak-field closure result. A finite entanglement-capacity microstructure is carried through admissibility closure, substrate length setting, edge transport, finite renormalization, continuum matching, Green-function source matching, and a covariant scalar EFT to produce Newtonian gravity, the galactic acceleration scale, the RAR law, and weak-field lensing consistency without per-system tuning. More broadly, the manuscript argues that the continuum metric sector itself should be understood as the low-energy capacity geometry of the same substrate, so Einstein gravity appears here as a continuum limit of the theory rather than as an independent foundation underneath it.

In compact form, the static chain is now: tetrahedral counting fixes $g_{\text{share,eff}}$, faithful sector resolution fixes $L_* = (3/2)\lambda_e e^{-7g_{\text{share,eff}}}$ and hence $G_* = (9/4)(\hbar c/m_e^2)e^{-14g_{\text{share,eff}}}$, finite returns give $\Sigma_{\text{ret}} = 65/9$, continuum matching fixes γ , Green matching fixes κ/γ , and the weak-field bridge gives the Newtonian branch, a_0 , the RAR law, and $\Phi = \Psi$ at leading order. Transport also closes at the level of finite propagation, $D/\tau_0 = c^2$. In the particle sector, the same shell algebra produces the charged-lepton mass ratios at 0.5–0.7% with no fitted parameters, with the three-generation count fixed by K^2 -spectrum collapse at $N = 3$ and the mass ladder $m_N/m_e = 720^N (2/7)^{N^2}$ following from the first-shell reduced-alphabet entropy and the singlet-projection shell kernel (Appendix I.1).

That does not finish the whole theory, but it does change the shape of the open problems. What remains is no longer the invention of a missing theory; it is the hard completion work of an

existing one: independent graph-level confirmation of the finite-loop self-energy, precision comparison of the matched weak-field G with the induced scale G_* , the resolved cluster lensing-map test of Section 17.5, full Boltzmann cosmology, and the nonuniversal boundary spectroscopy — relaxation spectra, stretched-layer corrections, transient response, and the substrate-side transport check on dynamical collapse — that controls UV-level black-hole observables. The faithful sector-resolution principle that sets the substrate scale is, in its combinatorial content, derived rather than assumed: the seven-fold additivity follows from one-pass occupation, the memoryless dressing of Postulate III, and the electron’s identity as the lightest charged defect, which forbids inter-channel correlation because correlation would add mass (Section 22, Appendix H). Its export as a *length* remains the central scale-setting identification, strongly cross-validated but not yet derived, with the one-defect transfer operator of Appendix D.4 the missing completion. Cosmology and the cluster source-projection sector remain structured extensions of the same ontology, carrying named tests rather than closure claims. Many-Pasts is no longer a detachable interpretation: the memoryless dressing it supplies is part of the scale-setting above, and the same weight recovers the Born rule and the arrow of time. The strong-field chain is now closed at the level of universal predictions, with stationary exteriors, horizon location, area-law coefficient, and leading reflectivity all fixed by the same substrate inputs as the weak-field chain.

Appendix A: Symbol Dictionary and Canonical Conventions

Appendix A gathers the conventions used throughout the technical material that follows. Its purpose is simply to keep the later appendices readable by fixing the units, field definitions, and couplings in one place before the denser calculations begin.

A.1 Units, signature, and entropy normalization

All dimensional quantities are expressed in SI units unless noted otherwise. The metric signature is $(-, +, +, +)$. Entropies are measured in nats, so Boltzmann's constant is absorbed into the entropy normalization. The canonical UV cell has spatial scale L_* and volume $V_* = L_*^3$. In the canonical branch L_* is fixed by faithful sector resolution,

$$L_* = \frac{3}{2} \lambda_e e^{-7g_{\text{share,eff}}}, \quad \lambda_e = \frac{\hbar}{m_e c}.$$

The conventional Planck length $L_P = \sqrt{\hbar G/c^3}$ is used only as a comparison scale or in standard gravitational thermodynamic expressions after the gravitational scale has been identified.

These conventions matter because the argument repeatedly moves between a dimensionless UV counting problem and a dimensionful continuum EFT. The units and signature are what make those two descriptions comparable rather than merely suggestive.

A.2 Core scalar variables

The canonical continuum variable is the vacuum-relative coarse-grained entanglement field

$$S_{\text{ent}}(x),$$

with vacuum baseline S_∞ and deficit

$$\delta S(x) = S_\infty - S_{\text{ent}}(x).$$

For nonlinear work the bounded occupancy fraction is

$$q(x) = \frac{S_{\text{ent}}(x)}{S_\infty} = 1 - \frac{\delta S}{S_\infty} \in [0, 1].$$

The absolute entropy unit is fixed only after choosing a cell or horizon normalization. Under a constant rescaling of S_{ent} , the quantities S_∞ and κ/γ rescale together, leaving $\delta S/S_\infty$ and $\kappa/(\gamma S_\infty)$ invariant. The source channel is

$$\chi(x) = -\frac{T^\mu{}_\mu}{c^2},$$

which is the continuum trace channel of the localized defect sector and reduces to the ordinary mass density ρ in the nonrelativistic static limit.

A.3 Couplings and derived observables

The main-text conventions are

$$\gamma : \text{entanglement-field stiffness}, \quad (12)$$

$$\kappa : \text{continuum defect-entropy coupling}, \quad (13)$$

$$\kappa_m(\ell) : \text{mass-per-entropy map at scale } \ell, \quad (14)$$

$$L_* = \frac{3}{2} \lambda_e e^{-7g_{\text{share,eff}}}, \quad (15)$$

$$G_* = \frac{c^3 L_*^2}{\hbar} = \frac{9}{4} \frac{\hbar c}{m_e^2} e^{-14g_{\text{share,eff}}}, \quad (16)$$

$$G_{\text{tet}}(0) : \text{tetrahedral on-site Green constant}, \quad (17)$$

$$g_{\text{share,max}} = \ln(1680), \quad (18)$$

$$g_{\text{share,eff}} : \text{admissibility-weighted sharing entropy}, \quad (19)$$

$$J_{\text{bare}}, J_{\text{eff}}^{\text{tree}}, J_{\text{eff}}^{\text{(ren)}} : \text{UV edge couplings}, \quad (20)$$

$$a_0 = \frac{cH_0 g_{\text{share,eff}}}{4\pi^2}. \quad (21)$$

The canonical weak-field bridge and Newton closure are

$$\frac{\Phi}{c^2} = -\frac{\delta S}{2S_\infty}, \quad \frac{\kappa}{\gamma} = \frac{3L_*}{4G_{\text{tet}}(0)\kappa_m(L_*)},$$

$$G = \frac{c^2 \kappa}{8\pi\gamma S_\infty}.$$

Collected in one place, these formulas also make clear which quantities are downstream of the closure chain. The UV data determine the stiffness and source-to-stiffness ratio first; the observable weak-field constants appear after the bridge and fixed S_∞ normalization are applied.

A.4 Notation map

One notation set is used throughout. The effective sharing entropy is denoted $g_{\text{share,eff}}$, the scalar variable is always the vacuum-relative field S_{ent} or its deficit δS , and the weak-field bridge is used in the single form stated above.

Appendix A serves as the reference layer for those conventions.

Appendix B: UV Boundary Ensemble and Admissibility Closure

Appendix B records the finite ultraviolet counting problem in its explicit form. It shows how the theory begins from a discrete boundary ensemble and ends with a unique admissibility-closed entropy rather than with an unconstrained continuum ansatz.

B.1 Minimal tetrahedral package

The canonical UV cell is a tetrahedron with four structural ingredients:

- a tetrahedral volumetric cell;
- half-integer fermionic face data on each face;
- injective face assignment across the four faces;
- binary orientation/parity.

Postulate II identifies the elementary defect sector as fermionic, so each face carries half-integer base spin j_0 . For a shared face the effective boundary sector is

$$j_0 \otimes j_0 = 0 \oplus 1 \oplus \cdots \oplus 2j_0.$$

Postulate I selects the maximum-capacity channel, hence $j_{\text{eff}} = 2j_0$ with $|M| = 2j_{\text{eff}} + 1 = 4j_0 + 1$ distinguishable face states. Injectivity across four faces requires $|M| \geq 4$. The $j_0 = 1/2$ option fails because it gives $j_{\text{eff}} = 1$ and $|M| = 3$. The first half-integer choice that works is therefore $j_0 = 3/2$, giving $j_{\text{eff}} = 3$ and the canonical seven-state face sector. The resulting state count is

$$\Omega_{\text{tet}} = 2 \times P(7, 4) = 1680, \quad g_{\text{share, max}} = \ln(1680) = 7.42654907240.$$

This is the minimal discrete package used in the framework to obtain a finite, isotropic, auditable boundary-channel structure.

The important feature is not just that the counting closes, but that it closes for structural reasons. Fermionic face data, injectivity, and maximum-capacity channel selection together force the seven-state face sector instead of leaving it as a tunable menu choice.

The minimality statement can also be written as a short proof. A volumetric cell in $d = 3$ needs at least four faces, so a tetrahedron is the first admissible simplex. The closure surrogate is three-component, so the face sector must be rich enough to support a nontrivial quadratic spectrum in $d = 3$ rather than a degenerate one-dimensional label count. Postulate II makes the face data fermionic, hence half-integer. Maximum-capacity channel selection then gives

$$j_{\text{eff}} = 2j_0, \quad |M| = 2j_{\text{eff}} + 1 = 4j_0 + 1.$$

Injectivity across four faces requires $|M| \geq 4$. The only half-integer option below $j_0 = 3/2$ is $j_0 = 1/2$, which gives $j_{\text{eff}} = 1$ and $|M| = 3$, so it fails. The first admissible fermionic choice is therefore $j_0 = 3/2$, giving $j_{\text{eff}} = 3$ and the canonical seven-state face sector. In that precise sense, the (4-face, 7-state) tetrahedral package is the minimal architecture compatible with a three-component isotropic closure mode, injective boundary information, and finite volumetric counting.

B.2 Closure invariant, kernel, and unique fixed point

The canonical scalar closure invariant is

$$K^2(b) = 48 - \frac{1}{3}(S^2 - \Sigma^2), \quad S = \sum_{i=1}^4 m_i, \quad \Sigma^2 = \sum_{i=1}^4 m_i^2.$$

The admissibility family is

$$p_\eta(b) = \frac{1}{Z(\eta)} e^{-\eta K^2(b)}, \quad Z(\eta) = \sum_{b \in B} e^{-\eta K^2(b)}.$$

The admissibility precision η is fixed by stationary normalized closure evidence. The closure constraint is the vanishing of the three-component oriented-face sum $\mathbf{c}(b) = \sum_i \hat{n}_i m_i \in \mathbb{R}^3$, of which K^2 is the quadratic invariant; marginalizing the three continuous closure components against the family at precision η contributes a determinant factor $\eta^{3/2}$, so the normalized closure evidence is $\eta^{3/2} Z(\eta)$ and its logarithm is

$$\mathcal{F}(\eta) = \ln Z(\eta) + \frac{3}{2} \ln \eta.$$

With $\partial_\eta \ln Z = -\langle K^2 \rangle_\eta$, the stationary condition $\mathcal{F}'(\eta) = 0$ is the closure relation

$$\langle K^2 \rangle_\eta = \frac{3}{2\eta},$$

so the factor $3/2$ is the determinant weight of the three independent closure components, not an equipartition rule imported onto the bounded discrete spectrum; the discreteness, multiplicities, and positive floor of the spectrum remain inside the exact finite sum $Z(\eta)$. The second derivative is $\mathcal{F}''(\eta) = \text{Var}_\eta(K^2) - 3/(2\eta^2)$, and on the exact spectrum \mathcal{F}' has a single interior zero,

$$\eta_* = 0.0298668443935,$$

at which $\text{Var}_{\eta_*}(K^2) = 15.69$ is dwarfed by $3/(2\eta_*^2) = 1681.6$, giving $\mathcal{F}''(\eta_*) = -1665.9 < 0$: η_* is the unique local maximum of the normalized closure evidence. Because the parity-symmetric ensemble is finite, both $Z(\eta)$ and $\langle K^2 \rangle_\eta$ are exact finite sums over the spectrum. The distinct closure-defect values and their degeneracies are

| | | | | | | |
|-------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| K^2 | $\frac{122}{3}$ | $\frac{134}{3}$ | $\frac{142}{3}$ | $\frac{146}{3}$ | $\frac{152}{3}$ | $\frac{154}{3}$ |
| mult | 96 | 96 | 96 | 288 | 192 | 144 |
| | | | | | | |
| K^2 | $\frac{158}{3}$ | 54 | $\frac{164}{3}$ | $\frac{166}{3}$ | $\frac{170}{3}$ | |
| mult | 384 | 192 | 48 | 96 | 48 | |

with total multiplicity 1680 as required. In particular,

$$Z(\eta) = \sum_a n_a e^{-\eta K_a^2}, \quad \langle K^2 \rangle_\eta = \frac{\sum_a n_a K_a^2 e^{-\eta K_a^2}}{\sum_a n_a e^{-\eta K_a^2}},$$

where (K_a^2, n_a) run over the table above. The closed-branch value η_* is therefore the unique interior maximum of an exact finite-spectrum functional, not an unseen numerical fit. The corresponding effective sharing entropy is

$$g_{\text{share,eff}} = - \sum_{b \in B} p_{\eta_*}(b) \ln p_{\eta_*}(b) = 7.41980002357.$$

The closed-branch moments used in the UV stiffness discussion are

$$\langle K^2 \rangle_{\eta_*} = 50.2229154254, \quad \text{Var}_{\eta_*}(K^2) = 15.6889750078, \quad a_{\text{UV}} = 0.0637390269.$$

These values quantify the local stiffness of the canonical closure point rather than a tunable phenomenological uncertainty. The closure-saturation product is

$$C_{\text{cl}} := \eta_* \langle K^2 \rangle_{\eta_*} = \frac{3}{2}, \quad C_{\text{cl}}^{-1} = \frac{2}{3}.$$

This same reciprocal $2/3$ reappears as the tetrahedral transverse export fraction in Appendix C and as the global correction in the faithful sector-resolution scale setting.

This is where the admissibility parameter stops being free. The kernel introduces η , and the closure condition removes its arbitrariness again by demanding that the fluctuation scale produced by the weighting agree with the weighting itself.

B.3 Rooted reduction and local benchmarks

Rooting on the shared face reduces the exact parity-symmetric ensemble to 140 rooted microstates and 69 rooted closure classes. The rooted classes can be labeled by $\alpha = (m_\bullet, K^2)$, so the same reduced state space already supports the local evaluation, the cavity benchmark, and the later shell propagation. The local information observable

$$\sigma_{\text{ind}}^{(r)} = \frac{H(X | Y_r)}{H(X)}$$

has the principal pre-nonlocal benchmarks

$$\sigma_{\text{ind}}^{\text{toy}} = 0.44997, \quad (22)$$

$$\sigma_{\text{ind}}^{\text{loc}} = 0.44708, \quad (23)$$

$$\sigma_{\text{ind}}^{\text{Bethe}}(J = 0) = 0.44749. \quad (24)$$

Here the Bethe value is the homogeneous cavity evaluation on the 69×69 rooted-class interaction graph at zero transport coupling,

$$\mu_\alpha \propto w_\alpha \left(\sum_\beta U_{\alpha\beta}(0) \mu_\beta \right)^{z-1}, \quad \sum_\alpha \mu_\alpha = 1,$$

with $z = 4$ and $U_{\alpha\beta}(0)$ the rooted shared-face compatibility matrix before shell transport is turned on. In other words, $\sigma_{\text{ind}}^{\text{Bethe}}(J = 0)$ is the cavity-theory benchmark of the same explicit rooted ensemble, not a disconnected numerical insert. The horizon target implied by the effective sharing entropy is

$$\sigma_* = \frac{\pi}{g_{\text{share,eff}}} = 0.42340665.$$

The gap between the local benchmarks and σ_* is therefore a genuinely shell / loop problem rather than a failure of the local admissibility closure.

That separation matters for the later UV story. It means the remaining work is not to repair the local closure ensemble, but to propagate it more accurately through transport and return structure.

B.4 What is fixed at this stage

By the end of the admissibility calculation, the framework has already fixed the microscopic counting ceiling, the unique closure point, the effective sharing entropy, and the local stiffness moments. What remains for the next appendix is not another entropy choice, but the propagation of those quantities into edge transport, finite renormalization, and continuum normalization.

Appendix B completes the UV counting problem and records the unique admissibility closure together with the local benchmarks needed by the coefficient chain.

Appendix C: Edge Kernel, Finite Renormalization, and Continuum Matching

Appendix C carries the middle part of the UV-to-IR derivation. Appendix B fixed what the local boundary ensemble is. Appendix C asks how that local data propagate into edge transport, loop dressing, and finally the continuum stiffness coefficient of the weak-field EFT.

C.1 Channel-averaged isotropy identity and tree coupling

Let \hat{n}_i be the four face normals of a regular tetrahedron. The exact identity

$$\sum_{i=1}^4 \hat{n}_i \hat{n}_i^\top = \frac{4}{3} I_3$$

implies a channel-averaged transverse fraction of $2/3$. The bare edge stiffness is therefore

$$J_{\text{bare}} = \frac{2}{3} \eta_* = 0.0199112296.$$

For a rooted $z = 4$ coarse adjacency graph, the tree-to-lattice map gives

$$J_{\text{eff}}^{\text{tree}} = \frac{J_{\text{bare}}}{z - 1} = \frac{2\eta_*}{9} = 0.0066370765.$$

This is the first place where local closure data become a transport law. The tetrahedral identity fixes the isotropic projection, and the rooted branching structure determines how much of the microscopic edge penalty survives as net outward propagation on the coarse graph.

C.2 Horizon target and shell convergence

The horizon-capacity target is

$$\sigma_* = \frac{\pi}{g_{\text{share,eff}}} = 0.42340665.$$

At the derived coupling the explicit shell values are

$$\sigma_{\text{ind}}^{(2)} = 0.42143, \quad \sigma_{\text{ind}}^{(3)} = 0.42166, \quad \Delta_{2 \rightarrow 3} = 0.00023.$$

The residual shift from the target is already small and stable by shell depth $r = 2$, isolating the remaining correction to the loopy local-return sector rather than a broad nonlocal ambiguity.

So the shell calculation narrows the open problem substantially. The tree branch already lands very near the target, and the residual discrepancy can be assigned specifically to local returns rather than to an uncontrolled long-range correction.

C.3 Finite-loop self-energy closure

The leading loopy correction is organized as a local Dyson dressing:

$$J_{\text{eff}}^{(\text{ren})} = \frac{J_{\text{eff}}^{\text{tree}}}{1 + J_{\text{eff}}^{\text{tree}} \Sigma_{\text{ret}}}.$$

The structural decomposition is

$$\Sigma_{\text{ret}} = 7 + \frac{2}{9} = \frac{65}{9},$$

and each term has a concrete return-channel origin. A short return motif leaves a shared face, explores a local closed loop, and re-enters the same coarse edge before contributing to net long-range transport. In the canonical label basis $m = -3, -2, \dots, 3$, there are exactly seven ways to do this without changing sector. These are the seven sector-diagonal returns, one for each face-label channel, and together they contribute

$$\text{Tr}(I_7) = 7.$$

In addition to these label-preserving loops, permutation symmetry allows one collective mode shared across all channels. Writing

$$P_{\text{sing}} = |u\rangle\langle u|, \quad u = \frac{1}{\sqrt{7}}(1, 1, \dots, 1),$$

this shared return is rank one. Any additional off-diagonal return sector would break the permutation symmetry of the canonical local ensemble, so there is no second independent collective channel to count. Only the transverse scalar branch feeds back into the coarse transport law, so the singlet first acquires the same $2/3$ projection factor that appeared in the tree coupling. It is then diluted by the rooted branching factor $1/(z - 1) = 1/3$ on the $z = 4$ graph, because

only one of the three outward branches returns to the original edge. The collective contribution is therefore

$$\mathrm{Tr}\left(\frac{2}{3}\frac{1}{3}P_{\mathrm{sing}}\right) = \frac{2}{9},$$

since $\mathrm{Tr}(P_{\mathrm{sing}}) = 1$. Equivalently,

$$R_{\mathrm{ret}} = I_7 + \frac{2}{9}P_{\mathrm{sing}}, \quad \Sigma_{\mathrm{ret}} = \mathrm{Tr}(R_{\mathrm{ret}}) = 7 + \frac{2}{9}.$$

This is the sense in which the finite-loop coefficient is counted rather than guessed: seven independent label-preserving returns plus one shared singlet return with exactly the same projection and branching weights already fixed in the tree map. Hence

$$c_{\mathrm{loop}}^{(\mathrm{ren})} \equiv \frac{J_{\mathrm{eff}}^{(\mathrm{ren})}}{J_{\mathrm{eff}}^{\mathrm{tree}}} = \frac{1}{1 + J_{\mathrm{eff}}^{\mathrm{tree}}\Sigma_{\mathrm{ret}}} \approx 0.95426,$$

and

$$J_{\mathrm{eff}}^{(\mathrm{ren})} \approx 0.00633348.$$

This reproduces the shell-target crossing near $J_{\mathrm{bare,cross}} \sim 0.019$ at the stated level of agreement.

The local Dyson dressing is therefore doing one precise job: it corrects the tree branch by accounting for the short motifs that recycle amplitude before it contributes to true coarse transport. The renormalized coupling is not a new parameter, but the tree coupling after local returns have been summed.

C.4 Euclidean-action normalization and continuum stiffness

The lattice quadratic form is interpreted canonically as a Euclidean action weight,

$$\frac{I_E}{\hbar} = \frac{J_{\mathrm{eff}}^{(\mathrm{ren})}}{2} \sum_{a,i} (Q_a - Q_{a+L_*\hat{n}_i})^2,$$

where the sum runs over all sites a and all four outgoing nearest-neighbor directions \hat{n}_i (each nearest-neighbor edge counted twice). The microscopic four-cell is

$$\Delta V_4 = \frac{L_*^4}{c}.$$

The same tetrahedral identity then yields

$$\gamma_Q = \frac{4\hbar c}{3L_*^2} J_{\mathrm{eff}}^{(\mathrm{ren})}$$

for the occupancy field Q_{occ} . With the horizon-capacity normalization

$$S = \pi Q_{\mathrm{occ}},$$

the canonical convention $\frac{\gamma}{2}(\partial S)^2$ gives

$$\gamma = \frac{4\hbar c}{3\pi^2 L_*^2} J_{\mathrm{eff}}^{(\mathrm{ren})} = \frac{4\hbar c}{3\pi^2 L_*^2} \frac{2\eta_*/9}{1 + (2\eta_*/9)(65/9)}.$$

Using the substrate-induced scale

$$G_* := \frac{c^3 L_*^2}{\hbar},$$

this is

$$\gamma = \frac{4J_{\text{eff}}^{(\text{ren})}}{3\pi^2} \frac{c^4}{G_*} \approx 8.556 \times 10^{-4} \frac{c^4}{G_*}.$$

This is the decisive stiffness-side matching step. Up to here the derivation has produced a dimensionless lattice weighting; after Euclidean normalization and faithful sector-resolution scale setting, that same weighting becomes the dimensionful continuum stiffness that appears in the weak-field action.

C.5 Local defect insertion and the source-side lattice constant

The stiffness-side matching is not the only UV quantity that can be closed locally. For the canonical rigid defect insertion, excluding one of the seven admissible face labels from one face removes exactly one-seventh of the isotropically averaged local partition weight. Therefore the logarithm of the isotropically averaged partition ratio is exactly

$$\Delta S_{\text{def}} := -\ln \left\langle \frac{Z_{\text{def}}}{Z_{\text{vac}}} \right\rangle_{\text{iso}} = \ln \frac{7}{6}.$$

This is the exact isotropic source benchmark in the canonical seven-label ensemble. The isotropically averaged defect free-energy cost differs from it only at $O(10^{-5})$ because the admissibility weighting breaks label symmetry only weakly.

The local source benchmark $\ln(7/6)$ should not be confused with the elementary fermionic one-bit anchor $\ln 2$. The former is the isotropically averaged partition-ratio shift produced by removing one admissible label from the seven-label boundary ensemble. The latter is the intrinsic binary entropy of an elementary occupied/unoccupied fermionic face-exclusion defect. The source theorem uses $\ln(7/6)$ to normalize the local scalar insertion into the lattice response, while the electron anchor uses $\ln 2$ to fix the mass-entropy unit of the elementary fermionic defect.

To propagate that local defect into the lattice field equation one needs the on-site Green function of the tetrahedral/diamond nearest-neighbor Laplacian. The dual graph of face-sharing tetrahedral cells is the four-valent diamond graph. Eliminating the two-sublattice structure gives the standard Brillouin-zone representation for the diamond lattice Green function [5]

$$G_{\text{tet}}(0) = \frac{1}{(2\pi)^3} \int_{[-\pi, \pi]^3} \frac{4 d^3 k}{16 - |1 + e^{ik_1} + e^{ik_2} + e^{ik_3}|^2}.$$

This integral is the reproducible source-side lattice constant: it is the self-energy of a unit point insertion for the same scalar mode whose long-wavelength stiffness was matched in Appendix C.4. Joyce's exact evaluation of the diamond-lattice Green function, in this normalization, gives [6, 5]

$$G_{\text{tet}}(0) = \frac{3\Gamma(1/3)^6}{2^{14/3}\pi^4} = 0.4482203943883814 \dots$$

Thus the source-side graph constant is an exact lattice invariant rather than a fitted numerical coefficient. Direct quadrature with endpoint extrapolation reproduces the same value.

Green-tensor transverse response. The same graph response also fixes the transverse export weight w_{\perp} , the single graph quantity that propagates downstream into the horizon channel count of Appendix F.5 and into the global $\ln(3/2)$ closure-saturation factor of the scale-setting relation. We compute it directly from the four nearest-neighbor bond frame, with no fitting freedom.

Let d_i , $i = 1, \dots, 4$, be the four diamond nearest-neighbor bond directions,

$$d_1 = \frac{(1, 1, 1)}{\sqrt{3}}, \quad d_2 = \frac{(1, -1, -1)}{\sqrt{3}}, \quad d_3 = \frac{(-1, 1, -1)}{\sqrt{3}}, \quad d_4 = \frac{(-1, -1, 1)}{\sqrt{3}}.$$

They obey

$$\sum_{i=1}^4 d_i^a d_i^b = \frac{4}{3} \delta^{ab}.$$

Distributing the scalar on-site Green response over the tetrahedral bond frame gives

$$\mathcal{G}_{\text{loc}}^{ab} = G_{\text{tet}}(0) \frac{1}{4} \sum_i d_i^a d_i^b = \frac{G_{\text{tet}}(0)}{3} \delta^{ab}.$$

For a local horizon normal \hat{r} ,

$$P_{\perp}^{ab} = \delta^{ab} - \hat{r}^a \hat{r}^b,$$

so

$$G_{\perp} = P_{\perp}^{ab} \mathcal{G}_{\text{loc}}^{ab} = \frac{2}{3} G_{\text{tet}}(0).$$

Thus the transverse export weight $w_{\perp} = 2/3$ is the transverse part of the exact local graph response.

Using the field normalization $S = \pi Q_{\text{occ}}$, the rigid local defect shift is

$$\delta Q_{\text{def}} = \frac{\Delta S_{\text{def}}}{\pi} = \frac{\ln(7/6)}{\pi}.$$

The corresponding local source amplitude in lattice units is therefore

$$s_{\text{def}} = J_{\text{eff}}^{(\text{ren})} \frac{\delta Q_{\text{def}}}{G_{\text{tet}}(0)},$$

so that

$$\frac{s_{\text{def}}}{J_{\text{eff}}^{(\text{ren})}} = \frac{\ln(7/6)}{\pi G_{\text{tet}}(0)} = 0.109472228 \dots$$

is a pure number fixed by the same UV lattice geometry.

The Green-function constant turns the local insertion into a continuum source theorem. Defining the defect-entropy density by

$$\sigma_{\text{def}} = \frac{\rho}{\kappa_m(L_*)},$$

the tetrahedral projection used in the stiffness mapping gives

$$\nabla^2 \delta S = -\frac{3L_*}{4G_{\text{tet}}(0)} \sigma_{\text{def}}.$$

Equating this with the weak-field source equation

$$\nabla^2 \delta S = -\frac{\kappa}{\gamma} \rho$$

closes the canonical source-to-stiffness ratio:

$$\boxed{\frac{\kappa}{\gamma} = \frac{3L_*}{4G_{\text{tet}}(0)\kappa_m(L_*)}}.$$

This is the source-side counterpart of the stiffness derivation. The edge-kernel calculation fixes how the scalar capacity mode resists gradients; the Green-matched defect calculation fixes how localized matter defects source that same mode.

Equivalently, the same source closure fixes the continuum quantity Ξ_ρ appearing in

$$\kappa = \frac{\Xi_\rho}{L_*^2 \kappa_m(L_*)}.$$

The absolute scale of S_∞ remains a fixed-epoch normalization convention in the bridge law, not a residual freedom in the source projection. In the cell-normalized gauge natural to the local source theorem, one may write

$$S_\infty^{\text{cell}} = \frac{3 \ln 2}{32\pi G_{\text{tet}}(0)} = 0.0461482516\dots$$

In a horizon-normalized gauge, S_∞ instead carries the much larger apparent-horizon capacity. These are not two physical constants. A constant rescaling of the entropy field rescales S_∞ and κ/γ together, leaving $\kappa/(\gamma S_\infty)$ and hence G unchanged. Thus the weak-field source map is closed in the canonical branch once the mass–entropy map $\kappa_m(L_*)$, the tetrahedral Green constant, and the standard cell convention are specified.

C.6 Local susceptibility cross-check

The exact local moments of the admissibility-closed ensemble supply an independent non-degeneracy check on the source-side result. From the variance in Appendix B,

$$a_{\text{UV}} := \frac{1}{\text{Var}_{\eta_*}(K^2)} = 0.0637390269,$$

which is the local zero-mode inverse susceptibility of the closure scalar. Two features of this number matter for the source theorem in C.5. First, a_{UV} is finite and positive, confirming that the closed branch at η_* has a non-degenerate zero-mode response rather than a critical singularity that would invalidate the linear Green-function matching used to derive κ/γ . Second, the same local susceptibility controls the stability of the closure mode that is being propagated into the shell and loop calculations, so the residual fractional discrepancy between the local benchmark and σ_* in C.2 is genuinely loop-sector rather than a local-closure failure.

The actual closure of κ/γ in the canonical branch is the Green-matched theorem in Appendix C.5. The role of a_{UV} here is restricted to confirming non-degeneracy of the local mode being matched.

C.7 UV-to-IR payoff

At this stage the weak-field UV coefficient chain is explicit:

$$\Omega_{\text{tet}} \rightarrow g_{\text{share,eff}} \rightarrow L_* \rightarrow J_{\text{bare}} \rightarrow J_{\text{eff}}^{\text{tree}} \rightarrow \Sigma_{\text{ret}} \rightarrow J_{\text{eff}}^{(\text{ren})} \rightarrow \gamma.$$

The same chain feeds

$$a_0 = \frac{cH_0 g_{\text{share,eff}}}{4\pi^2},$$

and faithful sector resolution gives the independent scale prediction

$$G_* = \frac{c^3 L_*^2}{\hbar} = \frac{9}{4} \frac{\hbar c}{m_e^2} e^{-14g_{\text{share,eff}}}.$$

and the Green-matched source theorem fixes

$$\frac{\kappa}{\gamma} = \frac{3L_*}{4G_{\text{tet}}(0)\kappa_m(L_*)}.$$

The weak-field bridge then converts this source-to-stiffness ratio into the observed Newtonian normalization through the invariant combination $\kappa/(\gamma S_\infty)$. The comparison with the same G_* fixed by faithful sector resolution tests that this normalization propagates the substrate scale coherently; it is not a disjoint-input second derivation of G , since the matched route carries L_* within it.

Appendix C completes the UV-to-continuum chain through the SI-normalized weak-field stiffness coefficient and the Green-matched source-to-stiffness ratio. The remaining uses of S_∞ belong to the fixed normalization of the weak-field bridge, not to the source sector itself.

Appendix D: Weak-Field Technical Derivations, Electron Anchor, and EFT Consistency

Appendix D gathers the weak-field derivations that are conceptually central but too dense to repeat in full in the main line. It is best read as a technical support layer for the bridge law, Newtonian recovery, the electron anchor, and the basic consistency checks of the EFT.

D.1 Bridge-law uniqueness

The weak-field bridge is derived once and then used throughout. Let the lapse be written as

$$N = e^{-F(\delta S/S_\infty)}.$$

Additivity of independent deficits requires $F(x+y) = F(x) + F(y)$, so continuity implies $F(x) = cx$. Standard weak-field metric normalization fixes $c = 1/2$, giving

$$N = e^{-\delta S/(2S_\infty)}$$

and therefore

$$\frac{\Phi}{c^2} = -\frac{\delta S}{2S_\infty}$$

to leading order. This is the unique weak-field bridge compatible with locality, additive independent deficits, and multiplicative redshift composition.

Writing the argument this explicitly removes one of the most common ambiguities in modified-gravity proposals. The bridge from entropy deficit to gravitational potential is not being chosen phenomenologically after the fact; it is fixed by the structural requirements of the weak-field limit itself.

D.2 Point source, Newton limit, and lensing

In the renormalized static branch,

$$\nabla^2 \delta S = -\frac{\kappa}{\gamma} \rho.$$

For a point source M ,

$$\delta S(r) = \frac{\kappa M}{4\pi\gamma r}, \quad g(r) = \frac{c^2 \kappa}{8\pi\gamma S_\infty} \frac{M}{r^2} = \frac{GM}{r^2}.$$

Because the leading entanglement stress carries no anisotropic stress,

$$\Phi = \Psi$$

at the order treated. The effective-halo rewrite is

$$\rho_{\text{halo}}(r) = \frac{1}{4\pi G r^2} \frac{d}{dr} \left[r^2 (g_{\text{obs}} - g_{\text{bar}}) \right].$$

Thus the same deficit field controls both orbital dynamics and light bending in the leading weak-field regime.

That shared control is the key weak-field consistency test. A viable branch must not reproduce galactic support only by sacrificing lensing, and the scalar deficit sector avoids that failure at the order treated.

D.3 Electron anchor and composite matter

The canonical fermionic entropy increment is

$$\Delta S_f = \ln 2.$$

The UV mass normalization is

$$\kappa_{m,\text{UV}} = \frac{\hbar}{c L_*} \frac{1}{\ln 2},$$

and the running law in the closed branch is

$$\kappa_m(\ell) = \kappa_{m,\text{UV}} \left(\frac{L_*}{\ell} \right)^{1+\alpha_{\text{cl}}}, \quad \alpha_{\text{cl}} = 0.$$

At the electron Compton scale $\ell = \lambda_e$ this gives

$$\kappa_m(\lambda_e) = \frac{m_e}{\ln 2},$$

which is the clean elementary anchor used here. Composite hadrons are not reduced to a bare constituent count. Their mass budget is assigned to a dressed bound-state entropy

$$m_{\text{hadron}} = \kappa_m(\ell_H) S_{\text{ent},H}^{\text{dressed}},$$

whose microscopic decomposition must include confinement, gluonic structure, trace-anomaly contributions, and chiral vacuum reorganization.

The contrast between the two sectors is deliberate. The electron is a clean one-bit defect anchor; hadrons are not. Their inertial content must therefore be assigned to a dressed entropy budget rather than to a naive constituent count.

D.4 Faithful sector resolution and induced G_*

The electron anchor enters the gravitational normalization through the absolute substrate length. The reduced electron Compton wavelength is

$$\lambda_e = \frac{\hbar}{m_e c},$$

and the admissibility-closed sharing entropy is

$$g_{\text{share,eff}} = 7.41980002357.$$

The closure fixed point also gives

$$\langle K^2 \rangle_{\eta_*} = \frac{3}{2\eta_*}, \quad \eta_* = 0.0298668443935,$$

so the closure-saturation factor is

$$C_{\text{cl}} := \eta_* \langle K^2 \rangle_{\eta_*} = \frac{3}{2}, \quad C_{\text{cl}}^{-1} = \frac{2}{3}.$$

The faithful sector-resolution principle is

$$\ln\left(\frac{\lambda_e}{L_*}\right) = 7g_{\text{share,eff}} - \ln C_{\text{cl}} = 7g_{\text{share,eff}} - \ln\left(\frac{3}{2}\right).$$

Equivalently,

$$L_* = \frac{3}{2} \lambda_e e^{-7g_{\text{share,eff}}}.$$

This is the named sector-resolution principle of the canonical branch. Once it is adopted, the cell length is fixed from non-gravitational input: the electron Compton scale and the already-derived substrate combinatorics. It gives

$$L_* = 1.60771947 \times 10^{-35} \text{ m},$$

which is about 0.528% below the conventional CODATA Planck length.

The induced gravitational scale follows from the same algebra used in the stiffness matching:

$$G_* := \frac{c^3 L_*^2}{\hbar}.$$

Substituting the sector-resolution relation gives the closed form

$$G_* = \frac{c^3}{\hbar} \left(\frac{3}{2} \lambda_e e^{-7g_{\text{share,eff}}} \right)^2 = \frac{9}{4} \frac{\hbar c}{m_e^2} e^{-14g_{\text{share,eff}}}.$$

Numerically,

$$G_* = 6.60399128 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2},$$

about 1.053% below the CODATA value. The uncertainty imported from λ_e is many orders of magnitude smaller than this residual, so the difference is theoretical rather than metrological.

The factor 3/2 is fixed structurally by the tetrahedral transverse-export geometry. The tetrahedral identity gives the transverse export fraction

$$\frac{1}{4} \sum_{i=1}^4 (\hat{n}_i \cdot \hat{u})^2 = \frac{1}{3}, \quad f_{\perp} = 1 - \frac{1}{3} = \frac{2}{3},$$

so $f_{\perp}^{-1} = 3/2$. The admissibility fixed point returns the same value,

$$C_{\text{cl}}^{-1} = (\eta_* \langle K^2 \rangle_{\eta_*})^{-1} = \frac{2}{3},$$

but it is built in. The condition defining η_* is $\langle K^2 \rangle_{\eta} = 3/(2\eta)$, i.e. $\eta \langle K^2 \rangle = 3/2$, so $C_{\text{cl}} = 3/2$ holds by construction — the fixed point returns the value the condition put in. The 3/2 is fixed once, by the transverse-export geometry, and the closure saturation only reproduces it. In the sector-resolution relation it enters once as a global subtraction $\ln(3/2)$, not seven times, because the closure condition is one scalar fixed-point condition on the admissibility family rather than seven independent sector constraints.

A short look-elsewhere check makes the role of this correction more precise. Write a possible log-gap correction as

$$L_* = \lambda_e \exp[-7g_{\text{share,eff}} + \Delta].$$

Several framework numbers lie near the required correction, but they do not play the same structural role:

| Correction Δ | Value | Induced G residual | Structural status |
|--|------------|----------------------|--|
| $\ln(3/2)$ | 0.40546511 | -1.05% | fixed by transverse tetrahedral geometry; mirrored by the closure-saturation normalization |
| $g_{\text{share,eff}} - 7$ | 0.41980002 | +1.82% | nearby, but not independently selected |
| $G_{\text{tet}}(0)$ | 0.44822039 | +7.78% | local Green constant, not global export factor |
| $7[-\ln c_{\text{loop}}^{(\text{ren})}]$ | 0.32774720 | -15.30% | transport renormalization, wrong role |

The preferred correction is the one already singled out by both the tetrahedral export geometry and the admissibility fixed-point saturation; the other nearby numbers play different structural roles.

The electron anchor carries three related roles. Its reduced Compton wavelength λ_e is the non-gravitational length used in faithful sector resolution. Its status as the lightest clean one-bit fermionic defect identifies which elementary excitation calibrates the seven-channel dressing block. Its mass fixes the mass-entropy map through

$$\kappa_m(\lambda_e) = \frac{m_e}{\ln 2}.$$

The same elementary defect enters the two normalization channels in distinct roles: its Compton length sets the UV cell scale, while its mass per one-bit defect entropy sets the source normalization.

The seven-sector scale-setting relation is not a counting heuristic; it has a concrete finite-dimensional realization on a transfer-operator history space, which we now construct. The face label algebra is

$$\mathcal{A}_7 = \bigoplus_{m=-3}^3 \mathbb{C}E_m, \quad E_m E_n = \delta_{mn} E_m, \quad \sum_{m=-3}^3 E_m = I_7.$$

The admissibility-closed tetrahedral state space has 1680 oriented injective states, with stationary weight

$$p_{\eta_*}(b) = Z^{-1} e^{-\eta_* K^2(b)}.$$

For a single sector, the refresh kernel

$$P_{\eta_*}(b, b') = p_{\eta_*}(b')$$

has Perron stationary entropy $g_{\text{share,eff}}$. The point is not that the electron simultaneously occupies seven mutually exclusive face labels in one tetrahedron. The labels are sector channels in the dressing history of the one-bit defect. The appropriate support object is therefore a history space, with one admissibility-closed sector layer for each $m = -3, \dots, 3$.

Let

$$\mathcal{H}_{\text{hist}} = \bigotimes_{m=-3}^3 \mathcal{H}_B^{(m)}, \quad \dim \mathcal{H}_B = 1680,$$

and let \mathcal{T}_m act as P_{η_*} on the m th factor and as the identity on the others. A representative one-pass dressing operator is

$$\mathcal{D}_e^{(0)} = \mathcal{T}_{-3} \mathcal{T}_{-2} \cdots \mathcal{T}_3.$$

The displayed order is only a representative of the symmetrized one-pass class. It does not introduce a physical ordering of the face sectors, nor does it multiply the support by an extra

factor of 7!. What matters is that the ground state of the elementary fermionic defect resolves each simple face sector once: fewer sectors leave unresolved label memory, while repeated sectors describe excited or additionally dressed states rather than the elementary anchor.

A sector-resolution step should not be identified with conditioning the 1680-state ensemble on boundary states containing a given label m . Such conditioning changes the entropy and does not reproduce $g_{\text{share,eff}}$. The sector label instead specifies which simple face-algebra channel is being resolved while the admissibility cloud sampled in that step remains the full closed boundary ensemble.

With the effective dimension defined by the stationary Shannon entropy of the positive history kernel, each sector contributes $g_{\text{share,eff}}$, so

$$\dim_{\text{eff}}(\mathcal{D}_e^{(0)}) = \exp(7g_{\text{share,eff}}) = 3.60286052 \times 10^{22}.$$

Only the transverse part of this support is exported into the weak-field scalar channel. The normalized export weight is

$$w_{\perp} = \frac{2}{3},$$

the same fraction fixed by the tetrahedral transverse projection and by the reciprocal closure-saturation factor. Hence

$$\dim_{\text{eff}}(\mathcal{D}_{e,\perp}) = w_{\perp} \dim_{\text{eff}}(\mathcal{D}_e^{(0)}) = \frac{2}{3} e^{7g_{\text{share,eff}}} = 2.40190701 \times 10^{22}.$$

The faithful sector-resolution principle is then the support-scale relation

$$\boxed{\frac{\lambda_e}{L_*} = \dim_{\text{eff}}(\mathcal{D}_{e,\perp}) = \frac{2}{3} e^{7g_{\text{share,eff}}}}.$$

The justification of this relation is given in Appendix H. An explicit dressing Hamiltonian (Appendix H.2) makes the lightest charged fermionic defect the one-pass ground state that resolves each of the seven sectors once. The refresh kernel P_{η_*} used above is the memoryless dressing of Postulate III: with the substrate relaxed long before readout, each pass is an independent draw from the admissibility ensemble, so every channel carries the full $g_{\text{share,eff}}$. The seven-fold additivity then needs the channels to be mutually independent, and they are, for a reason internal to the anchor: inter-channel correlation would contract the defect and raise its mass, and the electron is the lightest one-bit charged defect, so its dressing is the product and correlated dressings describe heavier excitations (Appendix H.5). With the 2/3 factor the tetrahedral transverse projection of Appendix C.5, the faithful sector-resolution principle is a consequence of the mass-entropy ontology and the electron anchor, and nothing in the chain remains free.

This closes the seven-fold additivity, but the relation rests on one identification that should be named at its true status. The effective dimension e^H is the perplexity of the coherent dressing — the effective number of mutually resolvable support positions in the defect’s history, not a count of cells in a spatial volume. Reading it as a physical length ratio means identifying the observable extent of a one-particle excitation with its Compton scale; that identification, not derived here, is the central dictionary of the elementary-defect branch. It is cross-validated — by the induced gravitational scale, by the charged-lepton ladder, and by the mass-minimization argument above, all selecting the same linear, first-power map — but it is not yet a theorem. A first-principles closure would exhibit a one-defect transfer operator T whose coherence length obeys $\lambda = C_{\text{cl}} L_* e^{H[p_T]}$ for its leading-mode support distribution p_T ; showing that this coherence length is set by the participation support e^H rather than by an ordinary spectral gap is the deepest microscopic completion target of the theory.

This is the microscopic content of the length formula above: the elementary one-bit fermionic defect is supported over one complete transverse-exported dressing block. The matched weak-field gravitational constant then follows from the gauge-invariant bridge

$$G = \frac{c^2}{8\pi} \frac{\kappa}{\gamma S_\infty},$$

with the entropy-unit normalization handled by the rescaling convention described in Appendix C.5.

D.5 EFT consistency checklist

The weak-field EFT does not rely only on successful phenomenology; it also passes a standard consistency checklist at the level claimed here.

- ‘No ghost’: the scalar kinetic term carries positive sign because $\gamma > 0$.
- ‘No tachyon’: the quadratic fluctuation operator contains no mass term at this order.
- ‘Correct-sign sourcing’: the defect-source coupling lowers the available entanglement capacity around positive-mass defect configurations rather than generating repulsive static behavior in the weak-field branch.
- ‘Causal propagation’: the transport completion satisfies $D/\tau_0 = c^2$, so the time-dependent sector propagates at finite signal speed.
- ‘Weak-field unitarity below cutoff’: once the scalar sector is quantized around the weak-field branch, the absence of ghost or tachyonic modes leaves an ordinary sub-cutoff scalar EFT rather than an obviously pathological one.
- ‘Energy-condition role’: the scalar gradient sector contributes positive local stiffness energy, while cosmological acceleration enters through the background branch rather than through a ghost-like local degree of freedom.

These statements are made at the EFT level claimed here. They do not replace the need for a fuller UV derivation, but they do show that the weak-field scalar sector is not buying phenomenology by obvious field-theoretic pathology.

The checklist is intentionally modest. Its role is not to prove ultraviolet completion of the full framework, but to show that the low-energy scalar sector used in the weak-field branch passes the standard first tests of EFT health.

The one place where an explicit formula is worth recording is linear vacuum stability in the time-dependent sector. Writing a small perturbation δs about the vacuum branch, the linearized telegrapher equation is

$$\tau_0 \ddot{\delta s} + \dot{\delta s} - D \nabla^2 \delta s = 0.$$

For a plane-wave mode $e^{-i\omega t + i\mathbf{k}\cdot\mathbf{x}}$, this gives the dispersion relation

$$\tau_0 \omega^2 + i\omega - Dk^2 = 0.$$

With $\tau_0 > 0$ and $D > 0$, the corresponding mode frequencies have non-growing time dependence, so the vacuum is linearly stable. The same sign structure is what underlies the earlier no-ghost and no-tachyon statements: positive kinetic stiffness, positive transport coefficients, and no negative mass-squared term in the linearized sector.

D.6 Quadratic fluctuations and weak-field stability

Expanding the action about an on-shell background yields the quadratic fluctuation operator

$$I^{(2)}[\delta S] = - \int d^4x \sqrt{-g} \frac{\gamma}{2} g^{\mu\nu} \partial_\mu \delta S \partial_\nu \delta S.$$

There is no quadratic mass term at this order, so the low-energy scalar sector contains one massless bosonic mode. Stability requires $\gamma > 0$, which is reinforced in the microscopic realization appendix by condensate hydrodynamics.

This is also the local EFT reason the bosonic occupancy language in the galactic section is natural rather than decorative. The weak-field branch genuinely contains a stable massless scalar mode whose occupation can be discussed meaningfully.

Appendix D provides the technical support layer for the weak-field bridge, Newton limit, electron anchor, substrate length branch, and EFT consistency audit.

Appendix E: Transport, Cosmology, and Hubble-Tension Implementation

Appendix E collects the time-dependent and homogeneous extensions of the static branch. The common purpose of these subsections is to show that the same scalar medium can propagate causally, relax toward its static limit, and support a cosmological background mode without losing contact with the weak-field structure already derived.

E.1 Telegrapher equation and causal closure

The time-dependent deficit field obeys

$$\tau_0 \partial_t^2 \delta S + \partial_t \delta S = D \nabla^2 \delta S + A \chi, \quad \frac{A}{D} = \frac{\kappa}{\gamma}.$$

Causality requires

$$\frac{D}{\tau_0} = c^2.$$

In the canonical no-new-IR-scale branch,

$$\tau_0^{-1} = H_0, \quad D = \frac{c^2}{H_0}.$$

This is the minimal causal completion of the static Poisson sector. The telegrapher form supplies propagation and relaxation, but it is chosen so that the static weak-field law remains the exact late-time limit rather than being replaced by a new phenomenological rule.

E.2 Static-limit recovery for galaxies

For a Fourier mode k , the telegrapher characteristic equation

$$\tau_0 s^2 + s + D k^2 = 0$$

has the roots

$$s = -\frac{1}{2\tau_0} \pm i\omega_k, \quad \omega_k \simeq ck$$

whenever $4\tau_0 Dk^2 \gg 1$. Galactic wavelengths are far below the critical scale

$$\lambda_c = \frac{4\pi c}{H_0} \approx 54 \text{ Gpc},$$

so galactic modes are deeply underdamped. Time-averaging the sourced solution over intervals large compared with $2\pi/\omega_k$ returns the static Poisson branch exactly, and the residual ponderomotive correction scales parametrically as

$$\frac{\delta F_{\text{pond}}}{F_{\text{static}}} \sim e^{-T/(2\tau_0)} \left(\frac{\omega_{\text{orb}}}{\omega_k} \right)^2 \sim 10^{-8}.$$

That estimate is why the transport sector does not undercut the static galactic results. The oscillatory contribution is present, but it is parametrically too small to compete with the near-stationary weak-field branch in ordinary galactic systems.

E.3 Homogeneous mode and cosmological sourcing

The cosmological split is

$$S(x, t) = \bar{S}(t) + s(x, t),$$

with $\bar{S}(t)$ the homogeneous mode and $s(x, t)$ the inhomogeneous weak-field sector. The background capacity is normalized by the apparent horizon,

$$S_\infty(t) = \pi \frac{R_A(t)^2}{L_*^2}, \quad R_A(t) = \frac{c}{\sqrt{H^2 + kc^2/a^2}}.$$

Because the field couples to the trace of the stress-energy tensor, the homogeneous mode is suppressed during radiation domination and turns on near matter–radiation equality.

This timing is the central cosmological virtue of the mechanism. The homogeneous mode is quiet when it must be quiet, then becomes relevant close to the epoch where a sound-horizon shift is most useful.

E.4 Sound-horizon shift and shear lock

In the closed cosmological branch, the trace-sourced homogeneous mode acts as a transient early-energy contribution. The qualitative payoff is a smaller sound horizon and an upward shift of the CMB-inferred Hubble constant toward the upper-68 / low-69 $\text{km s}^{-1} \text{ Mpc}^{-1}$ range. Local weak-field predictions are protected by the separation between $\bar{S}(t)$ and $s(x, t)$: the homogeneous mode changes the background branch without rewriting the local static Poisson law.

The result is qualitative but substantial. The homogeneous mode can matter cosmologically without forcing a re-tuning of the local weak-field sector that already fixed the galactic branch.

Appendix E closes the transport relation and preferred branch, while the cosmological sector remains structurally supported but not yet Boltzmann-closed.

Appendix F: Constrained-Capacity Strong-Field Branch

Appendix F records the strong-field black-hole branch in the form used by the main text. The result is not a full microscopic theory of collapse, but it is stronger than a mere horizon mnemonic. Once the bounded capacity variable is adopted and the lapse is tied to surviving capacity, the static spherical exterior is fixed: the continuum EFT lives on the domain $q > 0$, reduces to vacuum Einstein dynamics there, and terminates at the $q = 0$ surface.

F.1 Bounded capacity and the unique lapse map

The strong-field order parameter is the surviving-capacity fraction

$$q(x) = \frac{S_{\text{ent}}(x)}{S_{\infty}} \in [0, 1].$$

This bound follows directly from finite local channel capacity. If the vacuum channel count is finite and S_{ent} is the logarithmic coarse entropy of the surviving local ensemble, then no physical branch can have either negative capacity or more than the asymptotic vacuum capacity.

In a static exterior, the lapse associated with the asymptotic Killing time is determined by the local surviving capacity. Let

$$N = f(q).$$

The conditions are:

$$f(1) = 1, \quad \lim_{q \rightarrow 0^+} f(q) = 0.$$

The substrate-level composition axiom is that independent serial capacity losses compose multiplicatively on the lapse:

$$f(q_1 q_2) = f(q_1) f(q_2).$$

This is the assumption that extends the linear weak-field match to a nonlinear lapse map. With continuity, the positive solutions on $(0, 1]$ are $f(q) = q^{\alpha}$. Expanding near $q = 1 - \epsilon$ gives

$$N = q^{\alpha} = 1 - \alpha\epsilon + O(\epsilon^2).$$

The weak-field bridge gives

$$N = 1 - \frac{\epsilon}{2} + O(\epsilon^2),$$

so $\alpha = 1/2$ and therefore

$$N = \sqrt{q}, \quad N^2 = q.$$

Equivalently, if one writes $N^2 = F(q)$, the unique continuous multiplicative completion is $F(q) = q$. The nonlinear lapse rule is therefore fixed by capacity composition and weak-field matching; it is not a freely chosen black-hole ansatz.

F.2 Constrained-capacity action on the physical domain

The physical strong-field domain is

$$\mathcal{M}_q = \{(t, x) \mid q(t, x) > 0\}.$$

The minimal constrained-capacity action is the ADM gravitational action on \mathcal{M}_q , with a multiplier enforcing the lapse-capacity relation. In canonical form,

$$I_{\text{strong}} = \int_{\mathcal{M}_q} dt d^3x \left[\pi^{ij} \dot{h}_{ij} - N \mathcal{H}_{\text{GR}} - N^i \mathcal{H}_i^{\text{GR}} + \sqrt{h} \lambda (N^2 - q) \right] + I_{\text{matter}} + I_{\infty} + I_{\partial \mathcal{M}_q}.$$

The standard ADM constraint densities are

$$\mathcal{H}_{\text{GR}} = \frac{16\pi G}{c^4 \sqrt{h}} \left(\pi_{ij} \pi^{ij} - \frac{1}{2} \pi^2 \right) - \frac{c^4}{16\pi G} \sqrt{h} {}^{(3)}R, \quad \mathcal{H}_i^{\text{GR}} = -2D_j \pi^j_i.$$

The same content can be written in Lagrangian form as

$$I_{\text{strong}} = \frac{c^4}{16\pi G} \int_{\mathcal{M}_q} dt d^3x N \sqrt{h} \left({}^{(3)}R + K_{ij} K^{ij} - K^2 \right)$$

$$+ \int_{\mathcal{M}_q} dt d^3x \sqrt{h} \lambda (N^2 - q) + I_{\text{matter}} + I_\infty + I_{\partial\mathcal{M}_q}.$$

Here h_{ij} is the spatial metric, N and N^i are the lapse and shift, and

$$K_{ij} = \frac{1}{2N} \left(\dot{h}_{ij} - D_i N_j - D_j N_i \right)$$

is the extrinsic curvature. The term I_∞ is the usual asymptotic boundary term required for a well-posed variational principle and finite ADM energy, while $I_{\partial\mathcal{M}_q}$ is the effective action carried by the $q = 0$ capacity-exhaustion surface.

Varying λ gives

$$N^2 = q.$$

The q variation gives, schematically,

$$-\sqrt{h} \lambda + \frac{\delta I_{\text{matter}}}{\delta q} + \frac{\delta I_{\partial\mathcal{M}_q}}{\delta q} = 0.$$

Away from the $q = 0$ boundary, and in bulk vacuum with no explicit matter coupling directly to q , this gives $\lambda = 0$. The lapse variation gives

$$\mathcal{H}_{\text{GR}} + \mathcal{H}_{\text{matter}} = 2N\sqrt{h} \lambda,$$

so in bulk vacuum

$$\mathcal{H}_{\text{GR}} = 0.$$

The shift variation gives the ordinary momentum constraint,

$$\mathcal{H}_i^{\text{GR}} + \mathcal{H}_i^{\text{matter}} = 0.$$

The metric evolution equations then reduce to the vacuum Einstein equations on \mathcal{M}_q [7]. The constrained branch is thus not a new scalar-hair exterior. It is Einstein vacuum evolution on the surviving-capacity domain, together with the algebraic statement that the lapse is the square root of local capacity.

F.3 Horizon boundary action and boundary conditions

The only new strong-field ingredient not present in ordinary exterior GR is the inner boundary action on the capacity-exhaustion surface. At the EFT level its variation may be written schematically as

$$\delta I_{\partial\mathcal{M}_q} = \int_{\partial\mathcal{M}_q} d^3y \sqrt{\sigma} \left(\frac{1}{2} s^{AB} \delta\sigma_{AB} + \chi_\partial \delta q + \dots \right),$$

where σ_{AB} is the induced metric on the boundary world tube, s^{AB} is the effective boundary stress response, and χ_∂ is the response conjugate to the capacity variable. The universal boundary condition is not optional:

$$q|_{\partial\mathcal{M}_q} = 0.$$

The further constitutive information—whether the boundary is absorbing or partially reflecting, how quickly its channels relax, and what microscopic degeneracy they carry—is encoded in $I_{\partial\mathcal{M}_q}$.

This is why the static exterior can close before the full microscopic black-hole sector is finished. The exterior bulk equations only require $\lambda = 0$ on \mathcal{M}_q and the algebraic condition $N^2 = q$. The detailed horizon-channel physics lives in the boundary action. A shape variation of the moving boundary would give a force-balance condition of the schematic form

$$\delta_\xi I_{\text{strong}} = \int_{\partial\mathcal{M}_q} d^3y \sqrt{\sigma} \xi_n (\mathcal{P}_{\text{bulk}} + \mathcal{P}_{\partial q}) = 0,$$

where ξ_n is the normal displacement and the two pressures encode the limiting bulk stress and the boundary-channel response. Deriving $\mathcal{P}_{\partial q}$ from the tetrahedral graph ensemble is one of the remaining strong-field closure tasks.

F.4 Static spherical vacuum exterior

In static spherical vacuum, the result follows immediately. On \mathcal{M}_q , the bulk equations are the vacuum Einstein equations. The unique asymptotically flat static spherical solution is the Schwarzschild exterior,

$$ds^2 = - \left(1 - \frac{2GM}{c^2 r}\right) c^2 dt^2 + \left(1 - \frac{2GM}{c^2 r}\right)^{-1} dr^2 + r^2 d\Omega^2,$$

so the capacity variable is

$$q(r) = N^2(r) = 1 - \frac{2GM}{c^2 r}, \quad r > r_h,$$

with

$$r_h = \frac{2GM}{c^2}.$$

This also agrees with the weak-field capacity deficit:

$$\frac{\delta S(r)}{S_\infty} = \frac{2GM}{c^2 r}, \quad q(r) = 1 - \frac{\delta S(r)}{S_\infty}.$$

The domain carries the physical content. The exterior $r > r_h$ is exactly the standard Schwarzschild exterior. At $r = r_h$, $q = 0$. Continuing the same real capacity branch to $r < r_h$ would require $q < 0$, which is not in the state space. In this framework the classical Schwarzschild interior is therefore not the physical continuation of the same continuum EFT. It is the GR manifold analytically continued beyond the surface where the entanglement-capacity substrate has no remaining local continuum channels.

This exterior-domain result does not by itself decide what an infalling observer experiences at the $q = 0$ surface. Whether the capacity-exhaustion boundary is locally smooth, dissipative, or anomalous is a question about the boundary microphysics encoded in $I_{\partial\mathcal{M}_q}$, not a conclusion fixed by the exterior Schwarzschild closure alone.

Geometric identification of the capacity variable. The static rule $N^2 = q$ together with the bulk reduction to vacuum Einstein on \mathcal{M}_q forces a geometric reading of the substrate variable. In a spherically symmetric foliation with areal radius R , the normalization-independent gradient invariant on the orbit space is

$$q = h^{\mu\nu} \partial_\mu R \partial_\nu R = |\nabla R|^2,$$

which in Schwarzschild gives

$$q = 1 - \frac{2GM}{c^2 r} = N^2,$$

matching the substrate definition $q = S_{\text{ent}}/S_\infty$ on the exterior. In spherical dynamical collapse with Misner–Sharp mass $M_{\text{MS}}(R, t)$,

$$q(R, t) = 1 - \frac{2GM_{\text{MS}}(R, t)}{c^2 R}.$$

The substrate variable q and the geometric areal-radius gradient invariant therefore agree as a theorem of the framework’s structural commitments, not as an additional axiom. The marginal-trapped-surface condition $\theta_+ \theta_- = 0$ on a closed two-surface coincides with $q = 0$ independently of

the null normalization used to define θ_{\pm} : $q > 0$ is the untrapped exterior, $q = 0$ is the marginal / apparent horizon, and $q < 0$ would be the classical trapped interior. The capacity domain $\mathcal{M}_q = \{q > 0\}$ is the untrapped exterior, and the saturated boundary $\partial\mathcal{M}_q = \{q = 0\}$ is the apparent-horizon two-surface. Standard apparent-horizon formation theorems for gravitational collapse therefore produce the $q = 0$ saturation surface inside the framework, without an additional substrate-transport postulate; the bounded transport law of Appendix F.7 is a substrate-side consistency check on this geometric formation rather than its primary mechanism.

Rotating and charged stationary exteriors. The same vacuum-domain reduction extends to the rotating and charged cases. On \mathcal{M}_q the bulk equations reduce to vacuum Einstein, or to Einstein–Maxwell when a conserved gauge sector hosted by the substrate (Appendix I.2) is present. The standard no-hair theorems then identify the stationary asymptotically flat exterior branches as Kerr, Reissner–Nordström, or Kerr–Newman, and the saturation surface $q = 0$ coincides with the outer Killing horizon. The Bekenstein–Hawking entropy

$$S = \frac{k_B A}{4L_*^2}$$

retains its form with A the appropriate horizon area, and the Hawking temperature follows from the surface gravity as usual,

$$T_H = \frac{\hbar\kappa_{\text{sg}}}{2\pi k_B c}.$$

The inner Cauchy horizons of Reissner–Nordström and Kerr are not physical substrate interiors; they are analytic continuations into the $q < 0$ region the substrate does not enter.

F.5 Horizon thermodynamics and boundary capacity

Because the exterior geometry is unchanged, semiclassical quantities depending only on the exterior near-horizon saddle are unchanged. The Euclidean continuation is used here as an exterior-saddle calculation: the resulting periodicity depends on regularity of the near-horizon exterior geometry, not on a physical continuation of the substrate past $q = 0$. The Euclidean regularity argument therefore gives the standard Hawking temperature [9, 10],

$$T_H = \frac{\hbar c^3}{8\pi G M k_B}.$$

The same exterior Euclidean saddle gives the Bekenstein–Hawking area law [8, 9],

$$S_{\text{BH}} = \frac{k_B A}{4L_P^2}.$$

We now derive the 1/4 area coefficient from substrate inputs already in the framework, rather than from a stipulated channel-counting rule. The derivation uses two independent ingredients, each of which is independently closed within the bulk EFT before any horizon machinery is introduced.

Per-channel cut entropy from fermionic face exclusion. At a saturation surface $q = 0$, the physical continuum domain terminates at $\partial\mathcal{M}_q$. A tetrahedral cell on the exterior side whose outward face would, in the unsaturated bulk, be paired with a neighboring cell across that face instead has an unfilled pairing slot, because the would-be partner lies outside \mathcal{M}_q . By Postulate II the elementary face slot is fermionic and admits only the occupied (paired) state or the excluded (unpaired) state. A horizon cut face is therefore an instance of the same elementary face-exclusion defect that anchors the one-bit fermionic sector in the bulk. The

seven-state $j_{\text{eff}} = 3$ channel belongs to the paired bulk link $j_0 \otimes j_0 \rightarrow j_{\text{eff}}$, not to the cut face itself: with the partner cell absent, the paired representation is never formed. The seven-state structure enters the boundary count through the bulk graph response and hence through the channel density below, not through the per-channel entropy. The entropy carried by an elementary cut defect is consequently the same primitive fermionic increment that anchors the electron at $\kappa_m(\lambda_e) = m_e / \ln 2$,

$$\Delta S_f = \ln 2.$$

Channel density from the transverse bulk graph response. Because the local graph Green tensor is isotropic,

$$\mathcal{G}_{\text{loc}}^{ab} = \frac{G_{\text{tet}}(0)}{3} \delta^{ab},$$

a codimension-one horizon cut with local normal \hat{n} exports only the transverse two-plane component,

$$G_{\perp} = (\delta^{ab} - \hat{n}^a \hat{n}^b) \mathcal{G}_{\text{loc}}^{ab} = \frac{2}{3} G_{\text{tet}}(0).$$

The horizon channel density is therefore not a new boundary response coefficient; it is the transverse projection of the same bulk Green response already fixed in Appendix C. The number of active channels per outward angular direction is $G_{\perp} / \ln 2$, and integrating over the horizon two-surface (the spherical angular measure in Schwarzschild, the smooth axisymmetric horizon for Kerr, with isotropic transverse response in either case) gives

$$n_{\text{hor}} = 4\pi \frac{G_{\perp}}{\ln 2} = \frac{8\pi G_{\text{tet}}(0)}{3 \ln 2}.$$

Closure as an exact identity. Using the cell-normalized capacity baseline from Appendix C,

$$S_{\infty}^{\text{cell}} = \frac{3 \ln 2}{32\pi G_{\text{tet}}(0)},$$

the product of the two substrate inputs is

$$n_{\text{hor}} S_{\infty}^{\text{cell}} = \frac{1}{4},$$

an exact identity in which the Joyce diamond-lattice constant $G_{\text{tet}}(0)$ and the fermionic increment $\ln 2$ cancel between the two factors. A horizon of area A therefore carries the dimensionless entropy

$$\frac{S_{\text{hor}}}{k_B} = \frac{A}{4L_*^2},$$

and the Bekenstein–Hawking coefficient is recovered after the gravitational scale is matched so that $L_P(G_*) = L_*$. The area coefficient is closed within the constrained-capacity EFT under Postulate II and the transverse bulk graph response, both of which are independent of the horizon construction. The remaining strong-field task concerns nonuniversal boundary spectroscopy — the explicit microscopic Hamiltonian behind the relaxation spectrum, stretched-layer corrections, and transient response — rather than the universal area coefficient.

F.6 Absorption, ringdown, and echoes

Exterior wave propagation is unchanged. With $q = 0$ identified as the marginal-capacity / apparent-horizon surface (Appendix F.4), the standard near-horizon tortoise coordinate $r_* \sim r_h \ln q$ pushes the surface to $r_* \rightarrow -\infty$, and the wave equation reduces to $(\partial_t^2 - \partial_{r_*}^2)\psi \simeq 0$ near the boundary. Horizon regularity at the future horizon then selects the purely ingoing mode

$$\psi \sim e^{-i\omega(t+r_*)},$$

so the leading reflectivity is

$$\mathcal{R} = 0.$$

Perturbations outside the capacity-exhaustion surface obey the usual Schwarzschild Regge–Wheeler/Zerilli scattering problem [11, 12], with the absorbing boundary fixed by regularity rather than stipulated. The greybody factors, absorption cross sections, and quasinormal ring-down agree with the standard Schwarzschild exterior calculation.

Echoes are therefore correction-level rather than generic. They arise only if microscopic boundary channels carry finite UV reflectivity or if the effective reflection surface is displaced outward to a stretched layer

$$q = \epsilon > 0,$$

in which case the region between the exterior potential barrier and the partially reflecting layer behaves as a cavity, and a typical echo delay scales as

$$\Delta t_{\text{echo}} \sim \frac{2r_h}{c} |\ln \epsilon| + \tau_{\text{ch}},$$

where τ_{ch} is a boundary-channel relaxation time. Absence of echoes is the default absorbing-boundary prediction; echoes, if detected, would probe boundary microphysics rather than follow automatically from the absence of a classical interior.

F.7 Dynamical formation as a free-boundary problem

The geometric identification of Appendix F.4 supplies the primary horizon-formation story: standard apparent-horizon formation in gravitational collapse produces the $q = 0$ saturation surface as the first marginally outer trapped surface, with no additional substrate-transport postulate required. What remains is a substrate-side dynamical consistency check — writing a bounded causal transport law for q and verifying that its evolution agrees with the geometric formation, while bounding q in its physical range $[0, 1]$ throughout the collapse. The natural completion is a bounded causal transport system for q , for example

$$\begin{aligned} \partial_t q + D_i J^i &= -\Gamma(q) \Sigma[T_{\mu\nu}], \\ \tau_J (\partial_t + \mathcal{L}_v) J^i + J^i &= -D(q) D^i q. \end{aligned}$$

Here J^i is the capacity flux, $\Sigma[T_{\mu\nu}]$ is a positive depletion source built from the collapsing stress-energy, $D(q)$ is a bounded mobility, $\Gamma(q)$ is a bounded depletion rate, and $\tau_J > 0$ is a relaxation time. Eliminating J^i gives a telegrapher-type equation with finite characteristic speed

$$v_{\text{cap}} \sim \sqrt{\frac{D_0}{\tau_J}}.$$

Choosing constitutive functions such as

$$D(q) = D_0 q(1 - q), \quad \Gamma(q) = \Gamma_0 q$$

makes $q = 0$ and $q = 1$ invariant sets, preventing the evolution from overshooting into $q < 0$.

When q first reaches zero on a two-surface, that surface becomes a moving boundary

$$\partial \mathcal{M}_q(t) = \{x \mid q(t, x) = 0\}.$$

The level-set kinematics are fixed by differentiating $q(t, X(t)) = 0$ along the moving surface:

$$V_n = - \left. \frac{\partial_t q}{|\nabla q|} \right|_{q \rightarrow 0^+}.$$

In spherical symmetry this becomes

$$\frac{dr_f}{dt} = - \left. \frac{\partial_t q}{\partial_r q} \right|_{r=r_f(t)}.$$

During continued infall the exterior should be Vaidya-like with a slowly varying mass parameter, settling to the Schwarzschild exterior after the front stabilizes. This is a well-posed program, but not yet a closed derivation: the transport coefficients, boundary action, and channel relaxation spectrum must be computed from the graph ensemble or constrained by simulation.

The dynamical system is presumed to preserve the usual covariant conservation of the combined matter-plus-capacity stress-energy, with any local matter depletion balanced by flux, boundary work, or capacity-sector stress. Showing that this conservation structure follows from a graph-derived transport action, rather than imposing it as a constitutive condition, is part of the dynamical closure work.

The concrete closure tests are correspondingly specific. A spherical collapse simulation should show formation of the first $q = 0$ surface without overshoot into $q < 0$. A coupled matter-plus-capacity run should approach a Vaidya exterior during accretion and a Schwarzschild exterior after settling while satisfying the combined conservation law. Exterior perturbation simulations with an absorbing boundary should reproduce standard Schwarzschild greybody factors and ringdown, while partial-reflectivity runs should produce controlled echo delays. Finally, a microscopic boundary-action calculation or a graph-ensemble Monte Carlo of saturated boundary channels should reproduce the channel-counting rule that yields $n_{\text{hor}} S_{\infty}^{\text{cell}} = 1/4$. These are not new fit knobs; they are the numerical and microscopic tests that would close the dynamical and boundary sectors.

F.8 PPN boundary and weak-field breakdown

In the weak-field Solar-System regime, the scalar sector yields, at leading post-Newtonian order,

$$\gamma_{\text{PPN}} = \beta_{\text{PPN}} = 1,$$

with scalar-induced corrections to $\Phi - \Psi$ appearing only at quadratic weak-field order, schematically $O(\Phi^2/c^4)$, and with the remaining PPN coefficients vanishing in the canonical covariant branch. Weak-field truncations fail only when

$$\frac{|\Phi|}{c^2} = O(1), \quad \frac{\delta S}{S_{\infty}} = O(1),$$

which is exactly the regime where q rather than δS is the correct variable.

Appendix F therefore separates the strong-field claims cleanly. The bounded capacity variable, the lapse rule $N^2 = q$, the constrained-capacity exterior action, and the static Schwarzschild exterior on \mathcal{M}_q are closed within the EFT. The geometric identification of q with the marginal-trapped-surface function fixes horizon location and inherits apparent-horizon formation from standard GR collapse. The area coefficient closes under Postulate II and the transverse bulk graph response, and the leading reflectivity vanishes by horizon regularity. Stationary rotating and charged exteriors reduce to Kerr, Reissner–Nordström, and Kerr–Newman through the same vacuum-domain bulk reduction. What remains open is nonuniversal boundary spectroscopy: the explicit graph-derived microscopic boundary Hamiltonian behind the relaxation spectrum, stretched-layer corrections, and the substrate-side transport check on the dynamical evolution of q during collapse.

Appendix G: Many-Pasts Operational Closure and Arrow of Time

Appendix G records the operational content of the Many-Pasts branch in compact form. The main thing to keep in view is that this branch is conservative where laboratory quantum mechanics is concerned and ambitious only in the larger interpretive and cosmological claims built on top of that operational core.

G.1 Closed operational weight

The canonical history weight is

$$P(H|P) \propto e^{-D(H,P)}, \quad D(H,P) = -\ln \text{Tr}(\Pi_P \rho_{H \rightarrow \text{now}}).$$

This is the operational branch with $\alpha = 1, \beta = 0$.

Writing it this way matters because the generalized family is no longer left open in practice. The laboratory branch is fixed before any interpretive discussion begins.

G.2 Born recovery and no-signaling

In the ordinary projective laboratory limit,

$$e^{-D(H,P)} = \text{Tr}(\Pi_P \rho),$$

so the standard Born structure is recovered exactly. This fixes $\alpha = 1$. An independent signaling-sensitive bias channel is forbidden, which fixes $\beta = 0$. The result is ordinary operational quantum mechanics rather than a modified laboratory theory.

That is the core closure claim of the appendix. Whatever additional content the Many-Pasts sector adds, it does not do so by changing standard Born-rule laboratory predictions.

G.3 Arrow of time from conditional typicality

Let $h = \{M_t\}_{t < t_0}$ be a macrohistory conditioned on present records M_{t_0} . If the count of compatible microhistories is N_h , then

$$P(h|M_{t_0}) \propto N_h.$$

Under coarse-grained factorization,

$$\ln P(h|M_{t_0}) \approx \sum_{t < t_0} S(M_t) + \sum_{t < t_0} \ln T(M_{t+\Delta t}|M_t) + \text{const},$$

so entropy growth appears as a counting dominance effect among record-compatible histories rather than as a new laboratory coupling.

The arrow-of-time claim should therefore be read as a statement about conditional counting in the space of histories, not as the introduction of a new dynamical force.

Appendix G settles the laboratory sector operationally and leaves the cosmological and arrow-of-time content as a coherent interpretive extension.

Appendix H: Microscopic Realization and Coarse-Graining

Appendix H addresses a different question from the weak-field appendices. Instead of asking whether the coefficient chain is internally closed, it asks whether a plausible microscopic realization exists in which the same scalar stiffness and defect ontology arise naturally.

H.1 GFT condensate realization and coarse-graining

The candidate microscopic realization is a GFT/condensate picture with bosonic tetrahedral quanta $\phi(g_1, \dots, g_4)$ and fermionic defects ψ . In the condensate regime, the coarse field may be written as

$$\sigma(x) = \sqrt{n(x)} e^{i\theta(x)}.$$

The hydrodynamic identity

$$|\nabla_\mu \sigma|^2 = \frac{(\nabla_\mu n)^2}{4n} + n(\nabla_\mu \theta)^2$$

shows that if

$$S_{\text{ent}}(x) = S_0 + \alpha \ln \frac{n(x)}{n_{\text{bg}}},$$

then the coarse action contains a positive scalar stiffness

$$\gamma \sim \frac{Z_\sigma n_{\text{bg}}}{2\alpha^2} > 0.$$

The coarse source channel arises from fermionic face exclusion: what is macroscopically read as matter is a localized defect of the condensate, and the surrounding reduction of available occupancy is the long-wavelength field captured by the EFT. In this sense the microscopic appendix plays one clean role: it shows that the EFT is not hanging in midair, even though a finished first-principles derivation of every inhomogeneous continuum coefficient from the full underlying kernel is not yet available.

That is why the appendix remains brief but important. It does not replace the explicit coefficient derivation carried out earlier, but it shows that the ontology and sign choices of the EFT are compatible with a concrete microscopic picture rather than merely with an abstract formalism.

The condensate picture addresses the emergence of the continuum geometry. The complementary microscopic question — the defect dynamics that fixes the faithful sector-resolution principle of Appendix D.4 — is taken up in the remainder of this appendix, where the principle is shown to follow from the physically realized admissibility ensemble of Part II together with the Many-Pasts postulate.

H.2 The dressing Hamiltonian and the one-pass ground state

The faithful sector-resolution relation of Appendix D.4 was written there as a support-scale identity on the history space $\mathcal{H}_{\text{hist}} = \bigotimes_{m=-3}^3 \mathcal{H}_B^{(m)}$, using the refresh kernel $P_{\eta_*}(b, b') = p_{\eta_*}(b')$ on each sector layer. Its dynamical content is carried by an explicit defect Hamiltonian, which we now write down so that the choice of that kernel can be derived.

The seven face-label channels of the admissibility-closed ensemble define the orthogonal projectors E_m , $m = -3, \dots, 3$, of the face algebra \mathcal{A}_7 . By Postulate II the elementary defect is fermionic, so each channel carries an occupation number $n_m \in \{0, 1\}$ with fermionic creation and annihilation operators c_m^\dagger, c_m and $n_m = c_m^\dagger c_m$. When channel m is occupied it is dressed by a cloud state described by a density operator ρ_m on the boundary ensemble \mathcal{H}_B . The defect Hamiltonian is

$$H = \underbrace{\sum_m (\varepsilon_0 n_m - n_m F_m(\rho_m))}_{\text{single-channel terms}} + \underbrace{\sum_{m < m'} V_{mm'}(\rho_m, \rho_{m'})}_{\text{inter-channel coupling}},$$

with ε_0 the bare cost to occupy a channel, F_m the free energy released by dressing channel m with its cloud, and $V_{mm'}$ the residual coupling between the clouds of distinct channels. Two properties of the ground state of H supply the two ingredients of the sector-resolution relation.

One-pass occupation and the exponent seven. A channel is occupied in the ground state whenever its dressing free energy beats the bare cost, $F_m(\rho_m) > \varepsilon_0$. The lightest charged defect is the configuration that minimizes ε_0 , and it is therefore exactly the configuration for which all seven channels bind. “Lightest” and “resolves each of the seven sectors once” are then the same statement: the elementary fermionic defect fills each face sector exactly once, with fewer sectors leaving unresolved label memory and repeated sectors describing excited or further-dressed states. This one-pass occupation supplies the factor seven in $7g_{\text{share,eff}}$, the seven being the channel count $|M| = 7$ fixed in Part II.

Factorized cloud and the additive support. The dressing dimension multiplies across channels — so that the seven equal contributions $g_{\text{share,eff}}$ add rather than merge — precisely when the joint cloud state is a product, $\rho = \bigotimes_m \rho_m$. Whether the ground state of H has this product form is controlled by the inter-channel term $V_{mm'}$. Beyond the one-pass count, then, faithful sector resolution turns on the seven-channel cloud being a product of full-entropy channels: each channel must sample its whole ensemble, and $V_{mm'}$ must leave the channels uncorrelated.

The dressing Hamiltonian thus reduces the sector-resolution principle to that product structure in its ground state. The next three subsections establish it.

H.3 Slot coupling, layer factorization, and the additivity theorem

Two distinct correlation structures appear in the closure data, and separating them is essential: conflating them leads to the false conclusion that the cloud cannot factorize.

The closure invariant is a pure pair coupling. Expanding $S^2 = (\sum_i m_i)^2 = \Sigma^2 + 2\sum_{i<j} m_i m_j$ in $K^2(b) = 48 - \frac{1}{3}(S^2 - \Sigma^2)$ cancels the self-terms exactly and leaves the identity

$$K^2(b) = 48 - \frac{2}{3} \sum_{i<j} m_i m_j.$$

The admissibility weight $e^{-\eta_* K^2(b)}$ therefore contains only cross-terms between distinct face slots of a single boundary state; it does not factorize over those four slots. This is the exact origin of the residual correlation between face slots on the closed ensemble,

$$I(\text{slot}_0; \text{slot}_1) = 0.1545 \text{ nats}, \quad \frac{I}{H(\text{slot})} = 0.079,$$

so the four faces of one tetrahedron are about eight percent correlated, a structural feature of the closure invariant.

Slot coupling does not obstruct layer factorization. The factorization the support relation requires is over the seven *sector layers* $b^{(-3)}, \dots, b^{(3)}$ of $\mathcal{H}_{\text{hist}}$, each layer being a full boundary state drawn from the entire 1680-state ensemble. The slot coupling $-\frac{2}{3}m_i m_j$ lives *inside* a single layer’s boundary state: it relates the four faces of that one tetrahedron and never couples layer m to layer m' . The two structures act on different objects:

| Structure | What it couples | Role |
|-------------------------------------|--|---|
| slot coupling $-\frac{2}{3}m_i m_j$ | the four faces <i>within</i> one boundary state b | the eight-percent mutual information; lives inside each $\mathcal{H}_B^{(m)}$ |
| layer coupling $V_{mm'}$ | the seven sector layers $b^{(m)}$ of $\mathcal{H}_{\text{hist}}$ | controls factorization; acts <i>between</i> the factors |

The substrate’s intrinsic correlation, the most natural candidate obstruction to factorization, therefore acts at the wrong level to obstruct it: it is internal to a layer, not between layers.

Additivity of the history distance. The Many-Pasts history weight of Appendix G.1 is a log-overlap, $D(H, P) = -\ln \text{Tr}(\Pi_P \rho_{H \rightarrow \text{now}})$. If the dressing state and the resolution projector factorize over channels, $\rho = \bigotimes_m \rho_m$ and $\Pi_P = \bigotimes_m \Pi_m$, the trace factorizes and the logarithm converts the product into a sum,

$$\text{Tr}(\Pi_P \rho) = \prod_m \text{Tr}(\Pi_m \rho_m) \implies D = \sum_m [-\ln \text{Tr}(\Pi_m \rho_m)] = \sum_m D_m.$$

Additivity of D over the seven channels — and hence the multiplicativity $\text{dim}_{\text{eff}} = \prod_m e^{g_{\text{share,eff}}} = e^{7g_{\text{share,eff}}}$ used in the length relation — is therefore a theorem once the product structure holds. The implication needed downstream runs in one direction: a product cloud gives additive D and the full $7g_{\text{share,eff}}$. Whether the electron’s dressing is that product is decided by two conditions, which the rest of this appendix settles in turn.

The ceiling and the two conditions. Subadditivity bounds the support from above. With each channel marginal fixed to the admissibility weight, $H(B_m) = g_{\text{share,eff}}$, the joint entropy of a single readout obeys

$$H(B_{-3}, \dots, B_3) \leq \sum_{m=-3}^3 H(B_m) = 7g_{\text{share,eff}},$$

and the deficit

$$\Delta = \sum_m H(B_m) - H(B_{-3}, \dots, B_3) \geq 0$$

measures the total correlation among the seven channels. The full support $\text{dim}_{\text{eff}} = e^{7g_{\text{share,eff}}}$ is reached precisely when each channel carries its full entropy $g_{\text{share,eff}}$ and the channels are mutually independent, $\Delta = 0$. The first is a condition on each layer in substrate time; the second is a condition on the seven layers at a single readout. The next two subsections establish them.

H.4 Each channel at full entropy: memorylessness in substrate time

A channel reaches its full entropy $g_{\text{share,eff}}$ only if its dressing keeps no memory of its own past. To see what memory would cost, consider the one-parameter family of per-channel kernels

$$K_a(b, b') = a \delta(b, b') + (1 - a) p_{\eta_*}(b'), \quad a \in [0, 1],$$

which repeats the current state with probability a and otherwise redraws from the stationary weight. Every member has p_{η_*} as its stationary distribution and the same single-time marginal, so the equilibrium ensemble cannot say which one governs the dressing. The per-channel conditional entropy $H(b' | b) = \sum_b p_{\eta_*}(b) H(K_a(b, \cdot))$, evaluated on the exact 1680-state ensemble, nonetheless slides with the memory a :

| a (memory) | per-channel $H(b' b)$ | status |
|---------------|----------------------------------|--------------|
| 0.0 (refresh) | 7.41980 = $g_{\text{share,eff}}$ | full entropy |
| 0.1 | 6.99953 | reduced |
| 0.3 | 5.80153 | reduced |
| 0.5 | 4.40051 | reduced |
| 0.9 | 1.06642 | reduced |

Only the memoryless endpoint $a = 0$ — the refresh kernel $P_{\eta_*}(b, b') = p_{\eta_*}(b')$ of Appendix D.4 — returns the full $g_{\text{share,eff}}$. The kernel carries more than its stationary marginal, so this first condition is settled by the dressing dynamics, not by the equilibrium ensemble.

The dynamics settle it cleanly. Take any ergodic generator Q with p_{η_*} as its detailed-balance stationary weight — a continuous-time single-label-swap generator built from the same admissibility weights serves, with no memorylessness assumed. The kernel over substrate-time τ is $e^{Q\tau}$, and the mutual information between its endpoints decays to zero,

$$I(\tau=0.1) \approx 3.4 \text{ nats}, \quad I(\tau=1) \approx 0.030, \quad I(\tau=2) \approx 1 \times 10^{-4},$$

so $e^{Q\tau}$ relaxes to the refresh kernel. The readout sits far inside that limit: the substrate time is $\tau_* = L_*/c \approx 5.4 \times 10^{-44}$ s, the electron's Compton time is $\tau_e = \hbar/(m_e c^2) \approx 1.3 \times 10^{-21}$ s, and their ratio $\approx 2.4 \times 10^{22}$ is the same hierarchy λ_e/L_* the length relation expresses. The ground state is read out some 10^{22} mixing times after relaxation, so each channel draws afresh and carries its full $g_{\text{share,eff}}$. This is the Many-Pasts weight $P(H|P) \propto e^{-D(H,P)}$ seen on the dressing: a sum over pasts that keeps no trace of the particular one realized.

H.5 Independent channels: the electron as the lightest defect

The second condition is that the seven channels are mutually independent at one readout, so that the deficit Δ of H.3 vanishes. This needs no separate dynamical postulate; it follows from what the electron is.

Keep the deficit explicit in the support relation. With each channel at its full entropy, the joint entropy is $7g_{\text{share,eff}} - \Delta$, and the support reads

$$\frac{\lambda}{L_*} = \frac{2}{3} e^{7g_{\text{share,eff}} - \Delta}.$$

At a fixed substrate scale L_* , a one-bit charged defect whose dressing carries correlation Δ has spatial support $\lambda \propto e^{-\Delta}$, and through $m = \hbar/(c\lambda)$ a mass

$$m \propto e^{\Delta}.$$

Correlation among the dressing channels contracts the defect and raises its mass. Because $\Delta \geq 0$, the lightest one-bit charged defect is the one with $\Delta = 0$: the product dressing. The electron is that defect, so its seven channels are independent and its support is the full $\text{dim}_{\text{eff}} = e^{7g_{\text{share,eff}}}$. Correlated dressings are still allowed; they describe heavier, additionally dressed excitations, not the elementary anchor.

The selection is physical, not an appeal to ignorance. Independence is not read off from the absence of an inter-channel constraint in the ensemble; it is forced because correlation carries a mass, and the electron is defined as the lightest charged one-bit defect. The argument uses only the mass–entropy identification at the centre of the theory: the same relation $m = \kappa_m \Delta S$ that anchors the electron makes inter-channel correlation an excess of defect entropy, and so of mass.

With both conditions met, the scale-setting branch reads end to end without a free step. One-pass occupation fixes the seven channels (H.2); memorylessness in substrate time gives each its full entropy $g_{\text{share,eff}}$ (H.4); the electron's lightness forces the channels independent (this subsection). Hence

$$\frac{\lambda_e}{L_*} = \frac{2}{3} e^{7g_{\text{share,eff}}}, \quad L_* = \frac{3}{2} \lambda_e e^{-7g_{\text{share,eff}}}, \quad G_* = \frac{c^3 L_*^2}{\hbar},$$

with no gravitational quantity used as input.

Appendix H therefore does two things. The condensate realization shows that the EFT kinetic term has a natural microscopic origin, and the dressing analysis derives the faithful sector-resolution principle from the mass–entropy ontology, the one-bit electron anchor, and the memoryless dressing, so the substrate length and the induced scale follow from commitments the theory already makes. What remains open at the microscopic level is the first-principles derivation of every inhomogeneous continuum coefficient from the full kernel.

Appendix I: Mass and Gauge Extensions

Appendix I collects sectors that are structurally connected to the same entanglement logic but are not part of the closed weak-field core. They are kept here because they show how the framework may extend, not because the main derivation depends on them.

I.1 Charged-lepton spectrum from the shell algebra

Beyond the electron anchor, the charged-lepton extension is formulated as a shell spectrum of fermionic defect excitations,

$$\log \frac{m_N}{m_e} = B_0 N + A_0 N^2, \quad N = 0, 1, 2,$$

with the electron as the $N = 0$ ground state and the muon and tau as successive radial entanglement-shell excitations of the same core structure. Each shell excitation is interpreted microscopically as the exclusion of one face label from the seven-state alphabet of the boundary ensemble in Appendix B.

Three-generation termination from closure-spectrum collapse. Combinatorial injectivity alone allows up to $N = 3$ shells, since the reduced alphabet must retain at least four distinct labels to support an injective face assignment ($7 - N \geq 4$). The sharper structural bound comes from the admissibility-closure machinery. For the reduced ensemble at shell N , the closure invariant K^2 inherits a discrete spectrum whose dispersion is what gives the admissibility kernel $e^{-\eta K^2}$ real weight. For $N = 0, 1, 2$ the K^2 spectrum has multiple distinct values, and the closure equation $\langle K^2 \rangle_\eta = 3/(2\eta)$ has a non-degenerate solution.

At $N = 3$ the spectrum collapses. With only four labels left, every injective assignment to the four faces uses all of them, and K^2 depends on the labels only through the permutation-invariant sums S and Σ^2 , so every microstate of the reduced four-label ensemble carries the same value of K^2 . The admissibility kernel then has nothing to weight against, and the closure equation is satisfied vacuously by any η consistent with the single spectral value. Shell $N = 3$ therefore does not define a non-degenerate admissibility-closed branch in the same sense as the earlier shells.

The number of charged-lepton generations is consequently the number of shells with a non-degenerate closure branch,

$$N \in \{0, 1, 2\} \quad \implies \quad \text{three generations,}$$

a structural bound sharper than the combinatorial one. It is this mechanism, rather than injectivity alone, that fixes the generation count at three rather than four.

Linear coefficient from the first-shell reduced-alphabet entropy. The linear coefficient is the entropy cost of adding one shell excitation. With one face label excluded, the reduced ensemble has

$$\Omega_1 = 2 P(6, 4) = 720 = 6!,$$

i.e., the full symmetric group S_6 with orientation degeneracy. Direct admissibility evaluation on the reduced ensemble gives $g_1 = \ln \Omega_1 = \ln 720$ to within 0.1%, since the admissibility correction is small on the reduced ensemble where the alphabet symmetry is nearly intact. Hence

$$B_0 = \ln \Omega_1 = \ln 720 = 6.579 \dots,$$

compared with the empirical fit $B_0 = 6.586$ (agreement at 0.1%).

Quadratic coefficient from singlet-projection shell algebra. The quadratic coefficient follows from three premises already in the framework: (i) mass arises from scalar capacity response (Postulate II); (ii) the scalar EFT response is quadratic in the total source; (iii) the coarse scalar branch projects onto the permutation-invariant singlet of the seven-state face alphabet.

Let

$$\mathcal{H}_7 = \text{span}\{|m\rangle : m = -3, \dots, 3\}, \quad |u\rangle = \frac{1}{\sqrt{7}} \sum_{m=-3}^3 |m\rangle, \quad P_{\text{sing}} = |u\rangle\langle u|.$$

A shell excitation that marks the face state $|a_r\rangle$ contributes to a coherent shell source

$$|J_N\rangle = \sum_{r=1}^N |a_r\rangle.$$

The coarse scalar field δS is permutation-symmetric on face labels and therefore couples only through P_{sing} . The scalar-channel response is quadratic in the total source, so the relevant object is the source norm

$$\langle J_N | P_{\text{sing}} | J_N \rangle = \sum_{r,s=1}^N \langle a_r | P_{\text{sing}} | a_s \rangle.$$

Because $\langle a | P_{\text{sing}} | b \rangle = 1/7$ for any a, b , this reduces to $N^2/7$ independent of which specific face labels are excluded. Including the binary orientation degeneracy of the boundary ensemble, the orientation-summed scalar transfer weight per ordered shell incidence is

$$\mathcal{M}_{rs} = 2 \langle a_r | P_{\text{sing}} | a_s \rangle = \frac{2}{7}.$$

The ordered-pair count decomposes as $N^2 = N + 2\binom{N}{2}$: N self-incidences along the diagonal and $2\binom{N}{2}$ directed cross-incidences. Interpreting the shell mass map as an entropic transfer-weight partition factor, each ordered scalar incidence contributes the transfer weight $\mathcal{M}_{rs} = 2/7$. The N -shell dressing is therefore

$$\prod_{r,s=1}^N \mathcal{M}_{rs} = \left(\frac{2}{7}\right)^{N^2},$$

so

$$A_0 = \ln \frac{2}{7} = -1.253\dots,$$

compared with the empirical fit $A_0 = -1.255$ (agreement at 0.15%).

The ordered-bilinear structure is not a free assumption. It is the unique scalar-channel response to a multi-source state under the three premises just stated: the source norm $\langle J_N | P_{\text{sing}} | J_N \rangle$ runs over both indices independently because the integrated-out scalar contribution is quadratic in J_N .

Mass ladder and numerical accuracy. Combining the two coefficients,

$$\boxed{\frac{m_N}{m_e} = 720^N \left(\frac{2}{7}\right)^{N^2}, \quad N = 0, 1, 2.}$$

Equivalently,

$$\log \frac{m_N}{m_{N-1}} = \ln 720 + (2N - 1) \ln \frac{2}{7},$$

so each added shell contributes $\ln 720$ plus an odd number of new scalar-overlap incidences (1, 3, 5, ..., summing to N^2). The second difference is fixed at

$$\Delta^2 \log m_N = 2 \ln \frac{2}{7} = -2.506\dots,$$

compared with the empirical -2.509 (0.15%). With no parameters fitted to lepton data,

$$\frac{m_\mu}{m_e} = \frac{720 \cdot 2}{7} = \frac{1440}{7} = 205.71, \quad \text{PDG: } 206.77 \quad (0.5\%);$$

$$\frac{m_\tau}{m_e} = 720^2 \left(\frac{2}{7}\right)^4 = 3454.56, \quad \text{PDG: } 3477.23 \quad (0.7\%).$$

The dressing map, explicitly. The three factors follow from one construction. Write the N -th charged lepton as a coherent source raised on the electron ground dressing by exciting N shells, $\mathcal{S}_N = \sum_{k=1}^N s_k$, where each s_k excludes one face label from the seven-state alphabet. Its dressing free energy is a log-overlap functional of the kind used for the gravitational dressing in Appendix H.3, and it separates into a per-shell term and a pairwise term. Each shell carries the reduced-alphabet entropy $\ln \Omega_1 = \ln 6! = \ln 720$, the permutation count of the six surviving labels, which builds the linear factor 720^N . Each ordered pair of shells contributes a singlet log-overlap: a single label projected on the closure singlet $u = \frac{1}{\sqrt{7}}(1, \dots, 1)$ of Section 8 has weight $|\langle u|m \rangle|^2 = 1/7$, and orientation summation doubles it to $2/7$. The N shells form N^2 ordered pairs, so the pairwise term is $N^2 \ln(2/7)$, and exponentiating the free energy returns

$$\frac{m_N}{m_e} = \Omega_1^N \left(\frac{2}{7}\right)^{N^2} = 720^N \left(\frac{2}{7}\right)^{N^2}.$$

The coefficients are fixed, not chosen: $720 = 6!$ and the N^2 pair count are combinatorial, and $2/7$ is the orientation-summed singlet projection behind $\Omega_{\text{tet}} = 2P(7, 4) = 1680$. What remains an input is the form of the map — that the lepton mass is the exponentiated log-overlap free energy of the coherent multi-shell source — a log-overlap structure of the kind that carries the gravitational dressing, taken here as the form of the mass map.

The support exponent, tested against the spectrum. The ladder also tests the exponent in the support-to-length identification of Appendix D.4, which reads the physical ratio as the first power of the entropy support, $\lambda/L_* \propto e^H$. Letting that power float, $\lambda/L_* \propto e^{\alpha H}$, turns each lepton ratio into a measurement of α , since $m_N/m_e = [720^N (2/7)^{N^2}]^\alpha$:

$$\alpha_N = \frac{\ln(m_N/m_e)_{\text{PDG}}}{\ln[720^N (2/7)^{N^2}]}, \quad \alpha_\mu = 1.0010, \quad \alpha_\tau = 1.0008.$$

Read the same way, the electron-to- G_* branch needs $\alpha = 0.9999$ to reproduce the observed Newton constant, so three determinations across two sectors land within 10^{-3} of unity and bracket it. What the leptons settle is the power, not its final digit: a square-root map ($\alpha = \frac{1}{2}$) would put m_μ/m_e near 14 and a quadratic map ($\alpha = 2$) near 4×10^4 , and only the linear rule reproduces the spectrum with the natural shell count $720 = 6!$. The identification that sets the substrate scale is in this sense not a single assumption tuned to gravity; the same linear exponent is preferred in a sector that never enters the G_* chain.

Status and remaining audit. The three-generation termination, the linear coefficient, and the quadratic coefficient are thus fixed from machinery already native to the framework, with no parameters fitted to lepton data: the collapse at $N = 3$ is the permutation-invariance of

K^2 , and the 720, 2/7, and N^2 are the reduced-alphabet, singlet-projection, and ordered-pair counts of the map above. The form of the map is not a separate posit. Mass is $\hbar/(c\lambda)$ with $\lambda/L_* = e^H$, the same support-to-length identification that fixes L_* in Appendix D.4; the ladder is its $N > 0$ continuation, with the electron the $N = 0$ entry, and the equality of the support exponent across the two sectors is the $\alpha = 1$ check above. The per-shell term is the log-overlap additivity theorem of Appendix H.3, and the pairwise term is the singlet overlap of shells sharing the seven-state alphabet. Beyond that one identification, the only lepton-specific input is the structure of the excited source — a product across the six spectator channels, coherent in the shared singlet — which is what produces the 2/7 pairings.

The electron anchor remains the clean weak-field entry point; the heavier charged leptons are conditionally derived outputs of the same shell algebra, not closure-defining ingredients of the gravitational chain. Composite hadrons remain part of the dressed bound-state entropy program rather than a completed output of the present construction. The neutrino and quark sectors are not treated here. The same closure-spectrum-collapse argument should apply universally if the framework is right — the Standard Model has three generations across all fermion sectors — but verifying that requires extending the shell construction to the relevant defect classes.

I.2 Gauge-structure extension and Standard Model ontology

The framework’s primary target is the gravitational and dark sector. Standard Model gauge structure is hosted on the substrate as external content rather than derived from it. This subsection specifies which gauge components have substrate-natural analogues, which are substrate-compatible but external, and how the baseline-redundancy template accommodates the full Standard Model gauge group $SU(3)_c \times SU(2)_L \times U(1)_Y$.

Baseline-redundancy template. The same logic that organizes the gravity sector extends to any conserved-charge sector. For a conserved charge Q , introduce an entropy-like potential $S_Q(x)$ and require that physical observables depend only on differences of that potential, not on its absolute baseline. Localizing that redundancy requires a compensating connection. In the Abelian case,

$$D_\mu S_Q = \partial_\mu S_Q - qA_\mu, \quad S_Q \rightarrow S_Q + \alpha(x), \quad A_\mu \rightarrow A_\mu + \frac{1}{q}\partial_\mu \alpha,$$

giving Maxwell-type dynamics for A_μ . The non-Abelian generalization, for multiplet-valued entropic potentials transforming under a compact Lie group G , gives Yang–Mills covariant derivatives and field strengths in the standard form.

The template tells us what form a gauge interaction takes once a conserved-charge sector with baseline redundancy is present. It does not select which specific groups are physically realized.

Substrate-natural components. Two ingredients of Standard Model gauge structure are naturally accommodated by the substrate.

$U(1)_Y$. The baseline-redundancy template naturally accommodates an Abelian conserved-charge sector of the hypercharge type. The associated entropy potential is conserved-charge in character, and localizing its baseline redundancy produces the corresponding Maxwell-type dynamics. The argument is independent of the gravity-sector derivation, applying the same template to a separate conserved-charge potential.

$SU(2)_L$. The substrate naturally carries an $SU(2)$ -representation structure through its half-integer face data. The spin-3/2 commitment forced by maximum-capacity channel selection and tetrahedral injectivity (Section 5) carries an $SU(2)$ action on face data through the standard double-cover relation. The static K^2 ensemble breaks this $SU(2)$ to its $U(1)$ Cartan via

the magnetic-quantum-number projection, but the full $SU(2)$ is recoverable at the level of the quantum face data the projection discards. Identifying this representation-theoretic $SU(2)$ with the internal weak group $SU(2)_L$, including chirality and doublet assignments, remains external to the present derivation.

Substrate-compatible but external: $SU(3)_c$. Color $SU(3)$ does not emerge from the substrate by any of the standard mechanisms (direct decomposition of the seven-state face algebra, Seiberg duality, preon compositeness, topological soliton, monopole condensation). The obstructions are structural: $SU(3)$ is not a spin group, so it does not inherit from rotational covariance; its fundamental representation requires three equivalent (non-hierarchical) states, while the substrate’s natural three-fold structures (shell index, K^2 classes, residual face positions) are all hierarchical; its rank-2 Lie algebra requires two independent quantum-number axes, while the substrate provides only rank-1 axis structures; and its non-Abelian commutation relations cannot be reproduced from the substrate’s scalar edge couplings.

A positive ontological reading is nonetheless available within the framework. Color is the substrate’s name for an internal three-valued label on fermionic defects that lives in the non-singlet complement of the face-state algebra. The scalar gravitational branch couples to defect sources only through the permutation-invariant singlet (Appendix D), so the macroscopic scalar field is blind to non-singlet structure by construction. Color labels and color dynamics therefore live in a sector that gravity cannot resolve: the gravitational sector and the color sector occupy orthogonal subspaces of the substrate’s algebra. The strong interaction can be represented, within this hosting picture, as the local gauge dynamics of the color label under the general baseline-redundancy template. The choice of $SU(3)$ specifically is empirical input; it is the realized non-Abelian structure that fits the template.

This account explains several structural features of color without claiming to derive them. Gravity’s blindness to color follows from the scalar branch being the singlet projection of the substrate. Color confinement itself is not derived here; the substrate account provides that the macroscopic scalar-gravity branch resolves only singlet structure and is therefore blind to color non-singlets, which is consistent with the observed fact that hadronic macroscopic signatures are color-singlet but is not a replacement for the QCD confinement mechanism. The independence of color from generation, charge, and spin reflects the independence of the corresponding substrate sectors: color in the non-singlet complement of the face algebra, generation in the shell-excitation index (Appendix I.1), electromagnetic charge in the baseline redundancy of a separate conserved potential, and spin in the fermionic face data.

Hadronic matter and the mass–entropy bridge. The mass–entropy bridge $m = \kappa_m(\ell) \Delta S$ applies to fermionic defects regardless of their color label. For charged leptons, the defect is color-singlet and the dressed entropy is well-approximated by the shell-excitation budget of Appendix I.1, with mass ratios matching observation at the sub-percent level. For hadrons, the dressed entropy is dominated by the QCD-internal contributions of confinement-scale gluonic flux, trace-anomaly structure, chiral vacuum reorganization, and quark binding. The mass–entropy bridge then organizes the full dressed bound-state entropy budget,

$$S_{\text{ent,H}}^{\text{dressed}} = S_{\text{defect}} + S_{\text{bind}} + S_{\text{conf}} + S_{\chi\text{SB}},$$

without requiring the framework to re-derive QCD. The structural claim is compatibility: the dressed entropy budget is what QCD calculates, and the mass–entropy bridge converts it to inertial mass through the same running $\kappa_m(\ell)$ used for elementary sectors.

Scope summary. The Standard Model’s full gauge group $SU(3)_c \times SU(2)_L \times U(1)_Y$ is hosted on the substrate at three distinct levels of derivational support: $U(1)_Y$ from baseline redundancy

(substrate-natural via independent argument); $SU(2)_L$ from spin-3/2 fermionic face representation theory (substrate-natural with chirality identification external); $SU(3)_c$ as a non-singlet internal label (substrate-compatible but not substrate-derived). The framework's target scope remains the gravitational and dark sector, and the gauge sector is hosted accordingly.

Appendix I remains a coherent extension layer: structurally linked to the same entanglement logic, but not part of the closed static weak-field derivation chain.

Appendix J: Numerical Checks and Robustness

Appendix J is intentionally modest. It does not add new derivations. It collects the main numerical cross-checks that make it easier to see that the same coefficient chain survives repeated contact with independent benchmark calculations.

J.1 Cross-sector numerical checks

The cross-check program includes:

- the one-bit fermionic defect check $\Delta S_f = \ln 2$;
- the rooted-shell convergence check $\sigma_{\text{ind}}^{(2)} \simeq \sigma_{\text{ind}}^{(3)}$;
- the UV closed-branch moments $\langle K^2 \rangle_{\eta_*}$, $\text{Var}_{\eta_*}(K^2)$, and a_{UV} ;
- cross-sector consistency among the electron anchor, the substrate length L_* , the induced scale G_* , the Green-matched Newton closure, and the galactic scale a_0 .

These checks do not replace the derivations, but they show that the same coefficient chain survives independent numerical scrutiny across the sectors where closure is claimed.

J.2 Reproducibility ledger for the substrate scale

The substrate-length calculation is short enough to record as a numerical ledger. Enumerating the 1680 oriented injective tetrahedral states with labels $m = -3, \dots, 3$, weighting them by $e^{-\eta K^2}$, and solving

$$\langle K^2 \rangle_{\eta} = \frac{3}{2\eta}$$

gives

$$\eta_* = 0.02986684439352237, \quad g_{\text{share,eff}} = 7.419800023570903.$$

The one-pass seven-sector history support and its transverse export are then

$$e^{7g_{\text{share,eff}}} = 3.602860521062804 \times 10^{22},$$

$$\frac{2}{3}e^{7g_{\text{share,eff}}} = 2.401907014041869 \times 10^{22}.$$

Using $\lambda_e = \hbar/(m_e c) = 3.861592671986303 \times 10^{-13}$ m gives

$$L_* = \frac{\lambda_e}{(2/3)e^{7g_{\text{share,eff}}}} = 1.607719470158885 \times 10^{-35} \text{ m},$$

and hence

$$G_* = \frac{c^3 L_*^2}{\hbar} = 6.603991270884347 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}.$$

The source-side Green constant is obtained independently from the diamond-lattice integral in Appendix C.5. Uniform-grid quadrature with endpoint extrapolation gives

$$G_{\text{tet}}(0) = 0.448220394388\dots,$$

confirming the exact Joyce value

$$G_{\text{tet}}(0) = \frac{3\Gamma(1/3)^6}{2^{14/3}\pi^4}.$$

This is the value used in the source theorem. These numbers are not additional inputs; they are the numerical evaluation of the finite spectrum, the seven-sector history support, and the graph Green function already defined in the derivation.

Appendix J is supportive rather than closure-defining, serving as an audit layer for numerical consistency rather than an additional derivational sector.

References

- [1] S. S. McGaugh, F. Lelli, and J. M. Schombert, “Radial acceleration relation in rotationally supported galaxies,” *Physical Review Letters* **117**, 201101 (2016).
- [2] D. Clowe et al., “A direct empirical proof of the existence of dark matter,” *The Astrophysical Journal Letters* **648**, L109–L113 (2006).
- [3] B. Bertotti, L. Iess, and P. Tortora, “A test of general relativity using radio links with the Cassini spacecraft,” *Nature* **425**, 374–376 (2003).
- [4] Y. Tian et al., “The radial acceleration relation in CLASH galaxy clusters,” arXiv:2001.08340 (2020).
- [5] A. J. Guttmann, “Lattice Green functions in all dimensions,” *Journal of Physics A: Mathematical and Theoretical* **43**, 305205 (2010).
- [6] G. S. Joyce, “On the cubic lattice Green functions,” *Proceedings of the Royal Society of London A* **445**, 463–477 (1994).
- [7] R. Arnowitt, S. Deser, and C. W. Misner, “The dynamics of general relativity,” in *Gravitation: An Introduction to Current Research*, ed. L. Witten (Wiley, New York, 1962).
- [8] J. D. Bekenstein, “Black holes and entropy,” *Physical Review D* **7**, 2333–2346 (1973).
- [9] S. W. Hawking, “Particle creation by black holes,” *Communications in Mathematical Physics* **43**, 199–220 (1975).
- [10] G. W. Gibbons and S. W. Hawking, “Action integrals and partition functions in quantum gravity,” *Physical Review D* **15**, 2752–2756 (1977).
- [11] T. Regge and J. A. Wheeler, “Stability of a Schwarzschild singularity,” *Physical Review* **108**, 1063–1069 (1957).
- [12] F. J. Zerilli, “Effective potential for even-parity Regge–Wheeler gravitational perturbation equations,” *Physical Review Letters* **24**, 737–738 (1970).