

Entropic Scalar EFT: from Entanglement Microstructure to Gravity, Dynamics, and Cosmic Structure

Jacob Chinitz

May 3, 2026

Abstract

We propose that the vacuum carries a finite local entanglement capacity and that what we call matter consists of localized defects of that same substrate. Inertial mass is the entanglement content of a defect, gravity is the long-wavelength restructuring of the surrounding capacity, and the galactic phenomenology usually attributed to dark matter is the extended reach of that restructuring rather than an additional particulate component.

Because the substrate has a finite maximum update rate, exact spatial isotropy, and no external background manifold, its leading continuum description is naturally Lorentzian and generally covariant. General relativity is therefore not assumed as an independent starting point; it emerges as the low-energy capacity geometry of the substrate itself.

The main technical result is a closed static weak-field derivation with no per-system tuning. A minimal tetrahedral boundary microstructure determines, through admissibility closure, edge transport, finite loop dressing, continuum matching, and Green-function source matching, the scalar stiffness, source coupling, and weak-field bridge of a covariant scalar EFT. That EFT recovers Newtonian gravity, fixes the galactic acceleration scale a_0 , produces a specific radial-acceleration law rather than assuming one, and retains no gravitational slip together with the standard post-Newtonian values at the order treated. Newton's constant G follows from the matched lattice-to-EFT dictionary; the electron one-bit anchor supplies a non-circular fixed-normalization cross-check of the same gravitational scale.

What is new is both the interpretation of gravity and the degree of closure: a finite ultraviolet counting problem is carried through to a predictive weak-field theory in which the metric sector, the scalar coefficients, and the observable outputs are all tied to the same capacity mode. The broader framework extends to time-dependent transport, a trace-coupled cosmological sector, a bounded-occupancy nonlinear completion, and a Many-Pasts interpretation that recovers standard Born-rule quantum mechanics operationally. Those sectors are less complete than the static weak-field chain, which remains the central result.

Contents

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

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

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 not merely that entanglement-inspired effects can imitate aspects of gravity, but that one scalar capacity mode can be followed from microscopic counting 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.

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 is meant to keep the microscopic and continuum pictures tied together. At continuum level, $S_{\text{ent}}(x)$ is the field that can appear in an action and field equation. At the same time, it is not introduced as an arbitrary extra scalar. It is the coarse variable that records 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. Equivalently, the weak-field equations depend on the ratios $\delta S/S_\infty$ and $\kappa/(\gamma S_\infty)$ rather than on an absolute entropy unit. This is not a time-dependent gauge symmetry; it is a fixed-epoch normalization convention that keeps local cell language and horizon-capacity language compatible. Gravity sees fractional capacity depletion.

The principal coefficients and derived quantities used throughout are:

$$\gamma : \text{entanglement-field stiffness}, \quad (1)$$

$$\kappa : \text{defect-entropy coupling}, \quad (2)$$

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

$$g_{\text{share,max}} = \ln(1680), \quad (4)$$

$$g_{\text{share,eff}} : \text{admissibility-weighted effective sharing entropy}, \quad (5)$$

$$J_{\text{bare}}, J_{\text{eff}}^{\text{tree}}, J_{\text{eff}}^{(\text{ren})} : \text{UV edge-kernel couplings}, \quad (6)$$

$$a_0 = \frac{cH_0 g_{\text{share,eff}}}{4\pi^2}. \quad (7)$$

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, but not every weak-field derivation depends on it directly. 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 consequences are developed later as part of the theory’s interpretive and cosmological completion sector.

It is worth saying explicitly why this postulate remains in the theory even though the weak-field gravity chain does not need it at every step. Postulates I and II define the gravitational ontology directly. Postulate III belongs to the broader framework because the same entanglement substrate is also being asked to support an account of branch realization and temporal asymmetry. It is therefore part of the total theory, but it enters the derivational order later.

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 can depend only on the rotational scalar v^2 at leading order, and the boundary conditions are fixed: the temporal rate is maximal at $v = 0$ and vanishes when the full budget is exhausted at $v = c$. The surviving temporal fraction is therefore

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

In this sense special-relativistic time dilation is read here as a capacity-budget relation rather than as an independent postulate about flat spacetime. Once a finite invariant speed, vacuum homogeneity, and exact isotropy are in place, the Lorentz transformation law follows as the corresponding inertial symmetry rather than the Galilean one.

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 presented as the minimal architecture currently known to us that supports 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; it is not selected because it later happens to fit G or a_0 . 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 canonical closure condition is

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

This is the self-consistency condition for the admissibility kernel itself. The parameter η sets how strongly the weighting suppresses badly closed configurations, so the ensemble generated by that weighting must in turn exhibit the fluctuation scale that η presumes. In that sense the equation is the entropic analogue of a mean-field fixed-point condition: η is not chosen externally, but fixed by requiring the admissibility kernel to be consistent with its own induced closure fluctuations.

On the exact discrete spectrum this equation has a unique solution,

$$\eta_* = 0.0298668443935.$$

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 distinction between $g_{\text{share,max}}$ and $g_{\text{share,eff}}$ is therefore not ad hoc loss inserted 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 framework. The later continuum theory does not inherit the naive channel-counting ceiling; it inherits the portion of that channel space that remains after closure is imposed. That is why the downstream couplings should be read as consequences of admissibility-closed sharing rather than 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.

The same chain also fixes the horizon-normalization target

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

and the rooted shell observable converges to that target rapidly enough that the nonlocal correction is already strongly constrained by small shell depth.

At this point 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 at this point 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,

$$\left(\frac{2}{3}\right) \left(\frac{1}{3}\right) = \frac{2}{9},$$

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,

$$\frac{I_E}{\hbar} = \frac{J_{\text{eff}}^{(\text{ren})}}{2} \sum_{\langle ab \rangle} (Q_a - Q_b)^2,$$

over a microscopic four-cell of 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

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

If the canonical UV cell is identified with the Planck cell, $L_* = L_P$, this can be written as

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

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

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

The first bracket is the micro-to-continuum closure chain; the second bracket collects the weak-field observables it feeds. No later section re-derives this chain in full.

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 closed in the canonical weak-field branch by matching the same 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 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; it is the convention that turns fractional capacity depletion into the observed weak-field metric normalization. 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. The nontrivial claim is not that G is inserted as an independent extra input, but that the entanglement-side dictionary reproduces the same coefficient through

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

In other words, the action is presented in the observational normalization of the metric sector, and the closure program shows that the entanglement sector matches that normalization rather than introducing a separate free gravitational constant. 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 standard Einstein normalization is used here because it is the empirical weak-field normalization of the continuum geometry, while the entanglement-side derivation shows how that same normalization is recovered from the substrate dictionary rather than posited as an unrelated second input.

The logical order matters. The action is not meant to suggest that every term has been guessed independently from phenomenology. Rather, 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.

This fixes the canonical weak-field action.

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. Under locality, multiplicative redshift composition, additivity of independent deficits, and standard weak-field metric normalization, the unique leading bridge is

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

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.

This bridge closes the canonical weak-field continuum dictionary of the theory.

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, source convention, 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 check on the weak-field closure chain. 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. Appendix D also records the compact companion branch in which the same gravitational scale is recovered from (\hbar, c, m_e) together with the local sharing factor and transport exponent. This is best read as a non-circular fixed-normalization cross-check of the matched EFT dictionary, rather than as a second primary derivation of G .

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. The characteristic acceleration scale is

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

This formula already shows why the galactic phenomenology is not independent of the UV story. The same effective sharing entropy that appears in the closure chain now reappears in the acceleration scale governing departure from the Newtonian branch on galactic outskirts.

The corresponding $1 + 2$ channel decomposition separates one longitudinal slot aligned with the baryonic acceleration gradient and two transverse slots carrying the cosmic background scale. This fixes the galactic dimensionless variable as

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

This is the right variable because it measures baryonic forcing relative to the intrinsic acceleration scale set by the same closure chain. Baryonic matter sources the single longitudinal channel, while the two transverse modes form the quadratic capacity bath, so the occupancy depends on the amplitude ratio $\sqrt{g_{\text{bar}}/a_0}$ rather than directly on g_{bar}/a_0 . When $g_{\text{bar}} \gg a_0$, the system should reduce to the ordinary Newtonian branch; when $g_{\text{bar}} \ll a_0$, the response should cross into

the low-acceleration completion. The square-root form is the one selected by the canonical $1+2$ channel structure and is exactly what reproduces the deep-MOND scaling later in the section.

For the massless bosonic entanglement mode, the minimal stationary completion is therefore the Bose–Einstein occupancy branch

$$1 + n_B(x) = \frac{1}{1 - e^{-x}},$$

This bosonic language is not an extra fit ingredient added after the fact. The collective excitation of a scalar entanglement field is itself bosonic, so once the weak-field response is organized as occupancy of a massless scalar mode, Bose–Einstein statistics are the minimal stationary completion. The role of a_0 is then to provide the effective scale against which that occupancy is measured, so the galactic law becomes an occupancy statement rather than an empirical interpolation formula chosen by hand.

The resulting radial-acceleration law is

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

This has the correct asymptotic limits:

$$g_{\text{bar}} \gg a_0 \implies g_{\text{obs}} \approx g_{\text{bar}}, \quad (8)$$

$$g_{\text{bar}} \ll a_0 \implies g_{\text{obs}} \approx \sqrt{a_0 g_{\text{bar}}}. \quad (9)$$

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

$$v^4 \approx a_0 G M_b.$$

Once the channel geometry is fixed, the weak-field medium has a minimal bosonic occupancy completion. The interpolation law is not chosen after looking at galaxy data. It is the way the entanglement response fills the available modes when the baryonic source is weak compared with the intrinsic acceleration scale set by the same microstructural chain.

Structurally, the same UV channel geometry that fixes the microscopic coefficient chain also feeds the galactic EFT. There is no separate per-galaxy interpolation function chosen by hand.

The galactic branch is fixed up to the same channel-identification structure already used elsewhere in the weak-field EFT.

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

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

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:

$$\gamma_{\text{PPN}} = \beta_{\text{PPN}} = 1 + O(\Phi^2/c^4).$$

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 were easy to blur in earlier drafts. Ordinary galactic support still 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. This choice 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.

This is why the transport sector is not being used to manufacture the ordinary galactic law after the fact. The static branch still does that job. The transport equation is there to describe 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 $\tau_0^{-1} = H_0$. Detailed merger phenomenology remains frontier.

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 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 candidate underlying 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 purpose of the strong-field section is not to claim a finished black-hole solution. It is to replace an unspecified breakdown region with the minimal nonlinear completion compatible with the weak-field bridge and the bounded-capacity interpretation of the field.

That goal is modest but important. A weak-field theory that simply says “the approximation fails somewhere near horizons” leaves the ontology incomplete exactly where one would most want to know what the variables mean. The bounded-occupancy completion is meant to supply that missing meaning even though it does not yet solve the full strong-field equations.

The weak-field variable is the deficit δS . For strong field the natural variable is the bounded occupancy fraction

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

The canonical nonlinear completion is

$$N^2 = q, \quad g_{tt} = -q.$$

This is not introduced as a convenient ansatz. If the static lapse satisfies $N^2 = f(q)$, then vacuum normalization requires $f(1) = 1$, horizon normalization requires $f(0) = 0$, the weak-field bridge requires $f'(1) = 1$, and multiplicative composition of independent capacity-reduction layers requires

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

The continuous solutions are $f(q) = q^\alpha$, and the weak-field condition forces $\alpha = 1$. Thus

$$N^2 = q$$

is the unique continuous multiplicative completion compatible with the weak-field bridge.

The force of this uniqueness statement is that it turns the nonlinear completion into a rule rather than an arbitrary ansatz. Once one demands the vacuum limit, the horizon limit, the weak-field derivative match, and multiplicative composition of independent deficit layers, the lapse cannot be chosen freely. The bounded variable points to one minimal completion, not to an arbitrary family.

In this picture the horizon is the level set $q = 0$, i.e. complete local exhaustion of available entanglement capacity. The nonlinear completion is therefore not built by adding an independent scalar hair parameter but by slaving the occupancy field directly to the lapse. Full constrained exterior/interior solutions, scalar backreaction, interior regularity, and exact microstate-to-area matching remain frontier questions.

The same bounded variable also sharpens the black-hole reading of the framework. In spherical symmetry the weak-field exterior solution gives

$$f(r) \equiv \frac{\delta S(r)}{S_\infty} = \frac{2GM}{c^2 r},$$

so ordinary compactness is already the weak-field capacity-depletion fraction. The strong-field continuation therefore interprets the horizon kinematically as the first radius at which surviving capacity vanishes,

$$q(r_h) = 0 \quad \Longleftrightarrow \quad f(r_h) = 1 \quad \Longleftrightarrow \quad r_h = \frac{2GM}{c^2}.$$

The strong-field continuation does not present a solved black-hole theory. It shows instead that the weak-field scalar is not abandoned at strong field: it is completed by a bounded variable and a unique multiplicative lapse rule, so the nonlinear regime is at least posed by a concrete prescription rather than left as an unnamed failure zone. The full self-consistent strong-field exterior and interior problem is explicitly deferred.

Bounded occupancy and the horizon criterion give a minimal nonlinear completion rule. The full self-consistent strong-field exterior and interior problem, including scalar backreaction, remains the principal open problem of the strong-field sector.

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. Candidate Microstructure Hamiltonian and Underlying Dynamics

The UV closure chain is not meant to float free of possible microscopic dynamics. The candidate realization developed in the appendices 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 does not by itself replace the explicit coefficient closure already carried out in Appendix C, but it shows that the EFT is not a free phenomenological decoration.

This provides a coherent microscopic realization supporting the UV closure chain, but not yet a first-principles inhomogeneous derivation of every continuum term.

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	exact K^2 spectrum and multiplicity closure	Part II, App. B
$g_{\text{share,eff}}$	UV entropy	Closed	exact weighted evaluation	Part II, App. B
$J_{\text{bare}}, J_{\text{eff}}^{\text{tree}}$	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

Quantity Claim /	Sector	Status	Type of Support	Where Established
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	weak-field gravity	Closed in the matched weak-field branch	bridge law + Green-matched κ/γ + fixed S_∞ normalization	Part III, App. C–D
Electron anchor	mass sector	Closed up to elementary anchor	one-bit fermionic defect branch	Part III, App. D
Electron-anchor G branch	weak-field normalization	Non-circular cross-check	closed-form algebraic reduction from $(\hbar, c, m_e, g_{\text{share,loc}}, u_{\text{tr}})$	App. D
a_0	galactic EFT	Fixed in the closed weak-field realization	UV entropy + cosmic scale	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
Bounded occupancy q , $N^2 = q$	strong field	Closed as minimal nonlinear completion	uniqueness of multiplicative completion	Part V, App. F
Capacity-saturation horizon criterion	strong field	Structurally supported as a kinematic continuation	weak-field deficit continuation	Part V, App. F
Bekenstein–Hawking area-law bridge	strong field	Coherent consistency bridge	horizon-capacity matching	App. F

Quantity Claim /	Sector	Status	Type of Support	Where Established
Strong-field exterior/interior solutions	strong field	Frontier	completion problem	Part V, 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
Candidate microstructure Hamiltonian	UV realization	Coherent extension	condensate / GFT realization sketch	Part V, App. H
Lepton-shell mass extension	particle-sector extension	Coherent extension	constrained shell ladder with finite generation count	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 and systems such as the Bullet Cluster [2, 4] should therefore be treated as dynamical and transport tests, not 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 both the weak-field gravitational normalization and the galactic acceleration scale. A precision program that could test correlated shifts in 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.
- 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 bounded-occupancy nonlinear completion proves inconsistent with viable strong-field solutions, the strong-field branch fails while the weak-field theory may still remain viable as an EFT.

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, 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 the continuum limit of the framework rather than as an independent foundation underneath it.

In compact form, the static chain is now: tetrahedral counting fixes $g_{\text{share,eff}}$, finite returns give $\Sigma_{\text{ret}} = 65/9$, continuum matching fixes γ , Green matching fixes κ/γ , and the weak-field bridge gives G , a_0 , the RAR law, and $\Phi = \Psi$ at leading order. Transport also closes at the level of finite propagation, $D/\tau_0 = c^2$.

That does not finish the whole framework, but it does change the shape of the open problems. The remaining tasks are no longer the invention of a missing theory; they are the hard completion

tasks of an existing one: independent graph-level confirmation of the finite-loop self-energy, fuller microscopic derivation of the action-kernel normalization, full Boltzmann cosmology, and self-consistent strong-field solutions. Cosmology, bounded-occupancy strong field, and Many-Pasts all remain part of the same ontology, though they do not yet stand at the same level of derivational closure as the static 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 $V_* = L_*^3$; when the Planck branch is invoked explicitly, $L_* = L_P$.

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 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

$$\gamma : \text{entanglement-field stiffness}, \tag{10}$$

$$\kappa : \text{continuum defect-entropy coupling}, \tag{11}$$

$$\kappa_m(\ell) : \text{mass-per-entropy map at scale } \ell, \tag{12}$$

$$G_{\text{tet}}(0) : \text{tetrahedral on-site Green constant}, \tag{13}$$

$$g_{\text{share,max}} = \ln(1680), \tag{14}$$

$$g_{\text{share,eff}} : \text{admissibility-weighted sharing entropy}, \tag{15}$$

$$J_{\text{bare}}, J_{\text{eff}}^{\text{tree}}, J_{\text{eff}}^{(\text{ren})} : \text{UV edge couplings}, \tag{16}$$

$$a_0 = \frac{cH_0 g_{\text{share,eff}}}{4\pi^2}. \tag{17}$$

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. Earlier variants such as mixed $g_{\text{share}} / g_{\text{share,eff}}$ usage, duplicate bridge-law derivations, or shifted definitions of the scalar variable are not carried in parallel; where they matter historically, they are translated into the present conventions before use.

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 currently 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 closure condition

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

has the unique solution

$$\eta_* = 0.0298668443935.$$

Because the parity-symmetric ensemble is finite, the root-finding problem can be written directly from the exact discrete spectrum itself. 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 root of an exact finite-spectrum equation, 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.

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

$$\sigma_{\text{ind}}^{\text{toy}} = 0.44997, \quad (18)$$

$$\sigma_{\text{ind}}^{\text{loc}} = 0.44708, \quad (19)$$

$$\sigma_{\text{ind}}^{\text{Bethe}}(J=0) = 0.44749. \quad (20)$$

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

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

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

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

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

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

and

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

This reproduces the shell-target crossing near $J_{\text{bare},\text{cross}} \sim 0.019$ 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,

$$\frac{I_E}{\hbar} = \frac{J_{\text{eff}}^{(\text{ren})}}{2} \sum_{\langle ab \rangle} (Q_a - Q_b)^2,$$

over the microscopic four-cell

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

The same tetrahedral identity then yields

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

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

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

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

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

If $L_* = L_P$, 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, 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.

To propagate that local defect into the lattice field equation one needs the on-site Green function of the tetrahedral/diamond nearest-neighbor Laplacian. Writing the bond vectors as

$$\delta_i \in \left\{ \frac{(1, 1, 1)}{\sqrt{3}}, \frac{(1, -1, -1)}{\sqrt{3}}, \frac{(-1, 1, -1)}{\sqrt{3}}, \frac{(-1, -1, 1)}{\sqrt{3}} \right\},$$

the corresponding lattice constant is

$$G_{\text{tet}}(0) = \frac{1}{V_{\text{BZ}}} \int_{\text{BZ}} \frac{4 d^3 k}{16 - |f(k)|^2}, \quad f(k) = \sum_{i=1}^4 e^{ik \cdot \delta_i},$$

with numerical value

$$G_{\text{tet}}(0) = 0.448220394(5).$$

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.

In the older parameterization this source closure is equivalent to fixing 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. 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 scale 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. Combining this with the same nearest-neighbor gradient template used in the stiffness matching, together with the fixed normalization $S = \pi Q_{\text{occ}}$, gives the earlier susceptibility estimate

$$\boxed{\frac{\kappa}{\gamma} \approx 1.487 \times 10^3 \frac{\Xi_\rho}{c_{\text{loop}}^{(\text{ren})} L_*^4 \kappa_m(L_*)}}.$$

This estimate is valuable as an internal audit: it uses the local branch curvature a_{UV} , the derived gradient template, the loop-renormalization factor, and the same source map to check the scale of the source response. The actual closure of κ/γ in the canonical branch is the Green-matched theorem in Appendix C.5; the susceptibility estimate is retained as a consistency check rather than as the primary source derivation.

C.7 UV-to-IR payoff

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

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

The same chain feeds

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

and 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.

Appendix C closes the canonical UV branch 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 rather than 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 retired everywhere else. 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 Electron-anchor cross-check for G

Besides the matched weak-field identity

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

the framework also admits a standalone electron-anchor reduction in which G appears as an output rather than as an input normalization. The branch uses the standard constants (\hbar, c, m_e) , the reduced Compton scale

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

a local transport-sharing factor $g_{\text{share,loc}}$, a fixed SI normalization marker u_{tr} with units m^{-2} , and a transport-geometry exponent α_{tr} . In the minimal single-scale isotropic branch used here,

$$g_{\text{share,loc}} \equiv g_{\text{share,eff}}, \quad u_{\text{tr}} = 1 \text{ m}^{-2}, \quad \alpha_{\text{tr}} = \frac{1}{2}.$$

Define

$$F \equiv \frac{4 \ln 2}{g_{\text{share,loc}}}.$$

Eliminating the implied UV length at the last algebraic step yields the closed-form branch expression

$$G = \left[\frac{4\pi^2 u_{\text{tr}} c^{3\alpha_{\text{tr}}+2} \lambda_e^{2\alpha_{\text{tr}}+4} m_e^2}{F^2 \hbar^{\alpha_{\text{tr}}+2}} \right]^{1/\alpha_{\text{tr}}}, \quad F = \frac{4 \ln 2}{g_{\text{share,loc}}}.$$

For the canonical transport exponent $\alpha_{\text{tr}} = \frac{1}{2}$, this specializes to the quartic law

$$G = \left[\frac{\pi^2 u_{\text{tr}} c^{7/2} \lambda_e^5 m_e^2}{4(\ln 2)^2 \hbar^{5/2}} \right]^2 g_{\text{share,loc}}^4.$$

This is the most explicit algebraic version of the branch: once $(\hbar, c, m_e, \lambda_e)$ and the local sharing factor are specified, G is output directly. This branch is non-circular because G is not assumed in the input list; it appears only after eliminating the implied UV length from the electron-anchor chain. Equivalently, the relation is invertible:

$$F(G) = \left[\frac{4\pi^2 u_{\text{tr}} c^{3\alpha_{\text{tr}}+2} \lambda_e^{2\alpha_{\text{tr}}+4} m_e^2}{\hbar^{\alpha_{\text{tr}}+2} G^{\alpha_{\text{tr}}}} \right]^{1/2}, \quad g_{\text{share,loc}} = \frac{4 \ln 2}{F(G)}.$$

For the canonical choice $\alpha_{\text{tr}} = 1/2$, the scaling is quartic,

$$G \propto g_{\text{share,loc}}^4, \quad \frac{\delta G}{G} = 4 \frac{\delta g_{\text{share,loc}}}{g_{\text{share,loc}}}.$$

Using the simple test point $g_{\text{share,loc}} = 7.4$ gives

$$G_{\text{pred}} = 6.700223 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2},$$

while imposing the strict minimal-closure identification

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

gives

$$G_{\text{pred}} = 6.772222 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}.$$

This branch is not used to replace the Green-matched EFT derivation. Its role is to show that the same framework also contains a compact fixed-normalization reduction in which Newton's constant is output from the electron anchor, the local sharing factor, and standard constants.

For that reason this branch is best read as a cross-check rather than as the primary derivation of G . It shows that the weak-field normalization is overconstrained in a useful way: the EFT matching and the electron-anchor reduction point toward the same gravitational scale.

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, standalone G 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 $4\tau_0 D k^2 \gg 1$. 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 point of this appendix is therefore qualitative but substantial. It shows how 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: Strong-Field Completion and Post-Newtonian Boundary

Appendix F keeps the strong-field discussion on its proper footing. It does not attempt to solve the black-hole problem in full. Instead it records the bounded-variable completion rule, the resulting horizon criterion, and the exact point at which the weak-field expansion must give way to a nonlinear treatment.

F.1 Bounded occupancy and unique lapse prescription

The nonlinear completion is posed on the bounded variable

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

Let the static lapse satisfy

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

Vacuum normalization requires $f(1) = 1$, horizon normalization requires $f(0) = 0$, and weak-field recovery requires $f'(1) = 1$. If independent capacity-reduction layers compose multiplicatively in both lapse and surviving occupancy,

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

The continuous solutions are $f(q) = q^\alpha$, and the weak-field condition fixes $\alpha = 1$. Hence

$$N^2 = q, \quad g_{tt} = -q.$$

The horizon is therefore the level set $q = 0$, i.e. complete local exhaustion of available capacity.

The force of this subsection is uniqueness. Once one insists on the vacuum limit, the horizon limit, weak-field matching, and multiplicative composition, the nonlinear lapse prescription is no longer a matter of taste.

F.2 Capacity saturation, compactness, and horizon radius

The black-hole reading of the bounded-occupancy branch begins from the same weak-field variable used throughout the weak-field development. In spherical symmetry the static weak-field solution is

$$\nabla^2 \delta S = -\frac{\kappa}{\gamma} \rho, \quad \delta S(r) = \frac{\kappa M}{4\pi\gamma r},$$

while the Newton matching relation is

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

Hence the dimensionless capacity-depletion fraction is

$$f(r) \equiv \frac{\delta S(r)}{S_\infty} = \frac{2GM}{c^2 r}.$$

Thus the ordinary compactness parameter is already the weak-field entanglement deficit written as a fraction of available vacuum capacity.

The nonlinear surviving-capacity variable is

$$q(r) = \frac{S_{\text{ent}}(r)}{S_\infty} = 1 - \frac{\delta S(r)}{S_\infty}.$$

Near vacuum this is just the weak-field relation already used in the main text. The black-hole continuation interprets the horizon as the first surface at which surviving capacity vanishes,

$$q(r_h) = 0 \quad \Longleftrightarrow \quad f(r_h) = 1 \quad \Longleftrightarrow \quad \frac{2GM}{c^2 r_h} = 1,$$

so that

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

In this form the Schwarzschild radius is not imported from outside as an unrelated geometric fact. It is the radius at which the vacuum-relative deficit reaches complete local capacity exhaustion in the bounded nonlinear branch.

That is the main reason this kinematic continuation is worth keeping even though the full strong-field solution is deferred. It shows that the weak-field scalar language continues to mean something at the onset of strong gravity, rather than simply being abandoned there.

F.3 Horizon microstate capacity and area-law consistency

Even without a finished strong-field exterior, the framework should still be checked against the one macroscopic entropy law any entanglement-based gravity proposal must respect:

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

In the present construction, this law is read as the macroscopic horizon-capacity condition. The tetrahedral microstructure contributes a finite boundary-channel capacity through $g_{\text{share,max}} = \ln(1680)$ and its closure-weighted refinement $g_{\text{share,eff}}$, but the construction does not identify one tetrahedral cell with one literal Planck-area horizon bit in a naive one-to-one way.

The cleaner statement is a consistency bridge. The same microstructure that fixes the weak-field closure chain supplies the finite local boundary capacity from which a horizon entropy density can be built, while the macroscopic normalization remains the standard Bekenstein–Hawking area law. In that sense there is no conflict between the combinatorial boundary counting and

$$\frac{A}{4L_P^2};$$

the former supplies the microscopic channel capacity and renormalized sharing structure, while the latter remains the macroscopic thermodynamic target for horizons.

This is the strongest statement presently justified in the strong-field sector. A full microstate-to-area derivation for an actual horizon remains part of the unfinished UV completion program, but the framework is at least aligned with the standard area law rather than at odds with it.

F.4 Minimal strong-field action

The simplest bounded completion may be written as

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

with the physical branch restricted to $0 \leq q \leq 1$ and

$$V'_{\text{sat}}(1) = V''_{\text{sat}}(1) = 0$$

to preserve the weak-field massless-scalar sector near vacuum. This bounded action should be read as a completion rule for posing the nonlinear problem, not as a finished derivation of the full strong-field exterior. No universal strong-field potential has yet been derived that closes the scalar backreaction problem independently of the exterior mass scale.

In other words, the appendix gives a clean statement of what is known and what is not. The nonlinear variable, the horizon criterion, and the area-law compatibility statement are fixed; the self-consistent exterior and interior solutions remain open.

F.5 PPN boundary and breakdown of the weak-field expansion

In the weak-field Solar-System regime, the scalar sector yields

$$\gamma_{\text{PPN}} = \beta_{\text{PPN}} = 1 + O(\Phi^2/c^4),$$

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 the bounded-occupancy completion must replace the linear bridge.

This also makes the boundary between sectors precise: the strong-field completion begins exactly where the weak-field expansion ceases to be quantitatively trustworthy.

Appendix F fixes the minimal nonlinear completion rule and the kinematic horizon criterion. The full self-consistent strong-field exterior and interior problem remains frontier.

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.

Appendix H gives a coherent microscopic realization supporting the closed weak-field chain, but not yet a closure-defining first-principles derivation of every continuum term.

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 Mass extensions and lepton-shell sector

Beyond the electron anchor, the charged-lepton extension is formulated as a shell spectrum of fermionic defect excitations,

$$\log m_N = C_0 + B_0 N + A_0 N^2, \quad N = 0, 1, 2.$$

The physical picture is that the electron is the ground-state fermionic defect, while the muon and tau are successive radial entanglement-shell excitations of the same core structure. In that reading, the quadratic log-mass ladder is not an arbitrary three-parameter fit laid on top of the particle spectrum, but the closure form taken by a short finite shell sequence.

This shell ladder is not independent of the gravity sector. Its degeneracy structure is tied to the same sharing-entropy logic that fixes the weak-field couplings, so the particle hierarchy and the gravitational normalization are not being treated as disconnected subsystems.

Within this extension, the finite tetrahedral boundary topology also constrains the charged-lepton shell ladder to terminate after three generations. That statement should be read carefully: it is a striking structural payoff of the current shell picture, but it belongs to the extension layer rather than to the closed weak-field core. Still, if the framework really admits only three charged-lepton generations in this construction, that is a genuine and falsifiable output rather than an external Standard Model input.

This sector is treated as a constrained extension of the same entanglement closure logic rather than as a replacement for the electron anchor. Composite hadrons remain part of the dressed bound-state entropy program rather than a completed output of the current extension.

The distinction matters. The electron anchor remains the clean weak-field entry point, while the heavier mass sectors are exploratory continuations of the same logic rather than closure-defining ingredients.

I.2 Gauge-structure extension

The same baseline-redundancy logic that underlies the gravity sector can be extended to gauge sectors. For a conserved charge sector Q , introduce an entropy-like potential $S_Q(x)$ and require that physical observables depend only on differences of that potential rather than on its absolute baseline. Promoting that redundancy to a local symmetry requires a compensating connection. In the Abelian case, local baseline redundancy is implemented by

$$D_\mu S_Q = \partial_\mu S_Q - q A_\mu,$$

with

$$S_Q \rightarrow S_Q + \alpha(x), \quad A_\mu \rightarrow A_\mu + \frac{1}{q} \partial_\mu \alpha.$$

This yields the standard Abelian gauge structure, with Maxwell-type dynamics for A_μ and straightforward non-Abelian generalization for multiplet-valued entropic potentials, where the same redundancy principle leads to Yang–Mills covariant derivatives and field strengths in the usual form.

The point here is not that the full gauge sector has been derived, but that the baseline-redundancy logic used elsewhere in the framework naturally points toward familiar gauge structure rather than away from it. Gravity and gauge sectors are then aligned by a common principle: only baseline-invariant deficit information is physically meaningful.

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 used throughout 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 between the electron anchor, 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.

That is exactly the right role for this appendix. It is an audit layer for internal consistency, not a substitute for the analytic logic developed earlier.

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).