

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

Jacob Chinitz

July 12, 2026

Abstract

We propose that empty space is not a passive backdrop but a physical medium with a finite budget of quantum entanglement: the linking structure that allows parts of a quantum system to share state. Matter forms when some of that capacity becomes locked into stable, localized defects of the medium. A particle's mass measures how much entanglement is committed to such a defect. Gravity is the surrounding capacity-strain field: near matter, slightly less entanglement capacity is freely available, and in the weak-field limit the fractional shortfall gives the gravitational potential. The excess acceleration seen in galaxies, usually attributed to particle dark matter, is treated here as the large-scale continuation of the same capacity response rather than as a new unseen substance.

The central result is that this picture is not freely adjustable after the fact. Once one accepts the finite-capacity medium, the three founding postulates, and a specific minimal model for the smallest cell of space, finite counting fixes the cell entropy and the ordinary weak-field response. The resulting capacity action is the static scalar sector of the Einstein action written in the capacity variable, so it gives Newton's law and the leading no-slip metric without introducing another gravitational field. A separately identified transverse branch gives the galactic acceleration scale and the observed relation between galaxy rotation and ordinary matter, subject to the microscopic matching conditions stated in the paper.

The electron plays a double role. As the lightest clean charged defect, it fixes the exchange rate between committed entanglement and mass and calibrates the absolute cell scale. Many-Pasts supplies the history-space interpretation of that calibration while preserving ordinary Born-rule statistics and no-signaling. Applying the same faithful-resolution condition used for the cell ensemble makes the local renewal process memoryless. A reversible marked-transfer action then derives the finite charged response and routes it through the electron and the heavier charged-lepton shells. This adds no new founding premise and leaves the original tetrahedral construction intact.

We also test the cell model in a computer simulation of dynamical spacetime. Turning on the medium's weighting leaves the background geometry undisturbed, as required. Defects then strain nearby capacity and deform the local geometry, while a scrambled control does not. A predicted shift of the host geometry follows the cell model across a family of simulation settings, and the control follows its own distinct prediction. These tests are still limited in scale, but they connect the proposed medium to a dynamical geometry through measured consequences rather than analogy alone.

Beyond ordinary weak gravity, the framework extends to time-dependent transport, clusters, cosmology, the saturated early universe, black holes, and particle structure at explicitly labeled levels of closure. The finite marked-transfer and charged-lepton calculation is closed inside its displayed action. Its embedding in a stable geometric condensate, together with the transverse, cosmological, and strong-field completions, remains conditional or open as stated in the closure table.

Contents

Part I. Physical Idea and Canonical Definitions	4
1. Introduction: The Physical Claim	4
1.1 What Is Primitive, and What Is Closed	5
1.2 Physical Motivation for the Primitives	7
2. Canonical Field Content and Definitions	9
3. The Three Postulates	11
3.1 Information–Geometry Equivalence	11
3.2 Mass–Entropy Equivalence	11
3.3 Many-Pasts Hypothesis	12
4. Relativistic Continuum Structure	13
4.1 Capacity budget and continuum symmetry	13
4.2 Dependency Map of the Theory	14
Part II. UV Coefficient Chain	14
5. Why a Tetrahedral Boundary Ensemble	15
6. Admissibility Closure	16
6.1 Minimal isotropic kernel	16
6.2 Closure condition and uniqueness	16
6.3 Effective sharing entropy	17
7. Edge Kernel and Tree-Level Coupling	17
8. Finite-Loop Renormalization	18
9. Continuum Stiffness and SI Normalization	19
Part III. Weak-Field EFT and Static Phenomenology	21
10. Einstein Parent Action and the Reduced Capacity Frame	21
11. Capacity Variable, Bridge Law, and Variational Status	23
12. Newtonian Gravity and the Point-Source Limit	24
13. Electron Anchor: One-Bit Mass Scale and Seven-Sector Length Scale	25
13.1 Why the electron is the anchor	25
13.2 One-bit mass anchor	26
13.3 Seven-sector length anchor	26
13.4 Decorated marked-transfer vertex	27
13.5 Consistency checks	28
13.6 Composite sectors	28
14. Galactic Dynamics	28
15. Baseline Metric Closure and the Open Galactic Lensing Kernel	32
Part IV. Time-Dependent, Transport, and Cosmological Sectors	33

16. Why Dynamics Requires Extension Beyond the Static Branch	33
17. Causal Transport and Telegrapher Dynamics	34
17.5 Cluster Source Projection and the Diffuse–Decoupled Channel Split	35
18. Cosmology and the Hubble-Tension Sector	40
Part V. Nonlinear, Interpretive, and Completion Sectors	41
18.5 The Saturated Phase and the Cosmic Microwave Background	41
19. Why These Sectors Belong	44
20. Strong-Field Action: Spherical Closure and Its Boundary	44
21. Many-Pasts: The History-Space Ontology	46
22. Microstructure Hamiltonian and Underlying Dynamics	47
Part VI. Closure Status, Falsifiability, and Research Program	49
23. The Substrate on a Dynamical Lattice	49
23.1 The host geometry and the cell identification	49
23.2 The coupled ensemble and its controls	50
23.3 What the closure sector cannot supply: vacuum stiffness	51
23.4 The externally hosted vacuum and the role of conditioning	51
23.5 Compatibility: the weighting on dynamical geometry	52
23.6 The defect experiment: geometry responds to the theory’s mass	52
23.7 Reaching Newtonian range: the conservation requirement	53
23.8 What the program establishes	54
24. Closure-Status Table	54
25. Falsifiability and Observational Tests	60
25.1 Static weak-field falsifiers	60
25.2 Dynamical falsifiers	60
25.3 Cosmological falsifiers	60
25.4 Correlated-constant falsifiers	61
25.5 Many-Pasts status	61
26. What the Theory Would Have to Get Wrong to Fail	61
27. Comparison with Other Approaches	62
27.1 Relative to Λ CDM	62
27.2 Relative to MOND-like interpolation programs	62
27.3 Relative to Verlinde-style emergent gravity	62
27.4 Relative to TeVeS and other multi-field modified gravities	63
27.5 Relative to AeST	63
27.6 Relative to scalar-tensor gravity	63
27.7 Relative to quantum-mechanical interpretations	63
27.8 Relative to CDT, spin foams, and group field theory	63
28. Conclusion	64

Part I. Physical Idea and Canonical Definitions

1. Introduction: The Physical Claim

Space is not an empty container that matter sits inside. It is a finite medium of entanglement capacity, and the particles we call matter are stable defects of that medium rather than independent agents acting on it from outside. Gravity is then what the surrounding medium looks like once some of its capacity is locked up in such a defect. In brief:

- **Matter** is a localized capacity defect of the substrate.
- **Mass** is the entanglement that defect commits, read in mass units.
- **Gravity** is the extended capacity strain — the fractional capacity deficit — the medium carries around the defect.
- **Dark-matter phenomenology** is not a new substance but the same medium read in two further regimes: the long-range reach of the capacity-strain field on galactic scales, and its saturated phase in the early universe.
- **General relativity** is the low-energy geometry of this same capacity medium, not an external stage to which it is added.

The continuum statement is a scalar EFT for a vacuum-relative entanglement field $S_{\text{ent}}(x)$ and its deficit δS relative to the background capacity. The defect sector is written at continuum scale in ordinary stress-energy variables, but its ontology is unchanged; inertial mass enters through the mass-per-entropy map κ_m , and the weak-field potential is the fractional deficit $\delta S/S_\infty$.

A recent algebraic-QFT result provides unusually direct external support for this continuum ontology. Dorau and Much [21] show, for coherent scalar excitations on a bifurcate Killing horizon, that vacuum-relative quantum information is exactly the Killing-weighted matter-energy flux; when combined with the entropy–area relation, the same quantity is represented as Einstein curvature. Their result does not select the finite tetrahedral microstructure and assumes the area proportionality, but it independently realizes the paper’s central operational chain:

vacuum reference state $\longrightarrow$ matter as positive vacuum-relative information $\longrightarrow$ energy $\longrightarrow$ geometry.

Black-hole thermodynamics supplies the finite-area saturation endpoint; the present framework proposes a bounded microscopic interpolation between the local relation and that endpoint.

The proposal replaces part of the usual dark-sector story rather than relabeling it. The standard picture keeps visible matter and Einstein gravity and adds dark components to supply the missing gravitational response; here the vacuum already carries a finite entanglement-capacity structure, and the same medium accounts for ordinary weak-field gravity, the galactic excess usually attributed to dark matter, and the homogeneous cosmological mode. It also goes beyond an ordinary scalar extension of Einstein gravity: GR is not a stage to which the entanglement field is appended but the low-energy capacity geometry of the same substrate, with the scalar sector tracking how that geometry is depleted and redistributed by localized defects.

The paper asks whether that ontology can be made quantitative. The finite-capacity substrate, three postulates, and tetrahedral ensemble are inputs. Admissibility fixes the sharing entropy, edge transport fixes the tree stiffness, the decorated transfer vertex fixes the charged marked event, source projection fixes the ordinary coupling ratio, and the weak-field bridge relates fractional capacity deficit to gravitational potential. The finite marked-transfer calculation is closed inside the displayed action; the geometric GFT embedding and several infrared branches remain conditional.

The most controlled branch is the ordinary static weak-field action and source map

$$\text{microstructure} \longrightarrow \text{coefficient chain} \longrightarrow \text{static capacity EFT} \longrightarrow \{G, \text{baseline metric}\}.$$

It recovers the Newtonian point-source limit, baseline no slip, and GR PPN values through its Einstein parent. A specified transverse effective branch produces a_0 and the RAR without per-system tuning, subject to the microscopic matching conditions stated in Section 14 and Appendix N. The electron anchor, memoryless dressing, and marked vertex fix the substrate length inside the stated support-to-rate branch. The resulting Newton normalization and corrected charged-lepton ratios agree with current measurements within one standard deviation. Appendix L separates this action-level closure from the historical fact that the residuals were already known.

The later parts extend the framework into regimes with lower closure status — time-dependent transport, galaxy clusters, cosmology, the saturated early phase, strong fields, and the particle and gauge extensions — with Part VI recording the bookkeeping explicitly.

Many-Pasts belongs with the foundations. Faithful sector resolution uniquely selects the memoryless electron-dressing kernel, while the Many-Pasts postulate supplies the history-space branch in which that kernel operates; its consequences for quantum probability, branch realization, and the arrow of time are developed in Section 21 and Appendix G.

1.1 What Is Primitive, and What Is Closed

The word “closure” is used here in a specific sense. The paper does not derive the existence of a finite entanglement substrate or the tetrahedral boundary ensemble from a deeper microscopic Hamiltonian. The closure claim begins only after five theory-defining inputs are fixed: finite local entanglement capacity; geometry–capacity equivalence; mass–entropy equivalence, including matter as localized defects of committed capacity; the Many-Pasts ontology of Postulate III, with its operational quantum measure imported as stated below; and the minimal tetrahedral boundary ensemble as the ultraviolet counting architecture. Maximum caliber is not a sixth input. It is the history-space statement of the faithful full-support condition already used to select the admissibility ensemble: every pass carries the largest path entropy compatible with the same fixed marginal. This equivalence must be stated because finite capacity by itself does not imply renewal.

Given those inputs, finite counting fixes the admissibility weighting and effective entropy. Faithful full-support resolution fixes the memoryless replacement kernel, and the decorated native-cell vertex realizes its reversible update. The state-weighted determinant supplies the baseline seven-channel recurrence. The same closure amplitude, projected through two directed singlet returns and canonically dilated, supplies the marked correction and its 21-edge determinant. The electron anchor fixes the proper-time cadence, while the spatial and temporal readings of the same phase mode give $L_* = c\tau_*$. Edge transport and source projection determine the ordinary static response. The longitudinal functional is the Einstein scalar-constraint sector in a different variable, yielding the Newtonian limit, baseline no slip, and GR PPN values. The geometric condensate embedding and the galactic branch remain conditional.

The decorated vertex adds no sixth foundational premise. It specifies the minimal marked field content and gluing that realize the faithful full-support rule, native tetrahedral update, and one-bit fermionic electron anchor already listed above. Those action-level choices are stronger than a numerical ansatz because they determine the response multiplicity, determinant power, and routing together; they are also falsifiable, since the nonminimal and differently routed vertices give the alternative values displayed in Appendix H.9.

In compressed form, the central claim is

$$\boxed{\text{primitive UV capacity hypothesis}} \rightarrow \boxed{\text{finite counting + admissibility}} \rightarrow \boxed{L_*} \\ \rightarrow \boxed{\gamma, \kappa/\gamma} \rightarrow \boxed{\delta S \leftrightarrow \Phi} \rightarrow \boxed{G, \text{baseline metric}}, \quad \boxed{a_0, \text{RAR}} \text{ in the conditional transverse EFT.}$$

The microstructure is therefore not the result being proven; it is the finite ultraviolet counting problem from which the weak-field sector is derived.

The absolute length calibration uses the electron, the lightest clean charged defect, as the dimensional anchor. Faithful full-support resolution fixes the replacement kernel because $H(B_{t+1} | B_t) \leq H(B_{t+1}) = g_{\text{share,eff}}$, with equality only when the endpoint mutual information vanishes. The lightest-defect functional selects the fermionic ceiling $k = 7$ and $\Delta_7 = 0$. The state-weighted determinant of the corresponding seven renewed clouds is $r = e^{-7g_{\text{share,eff}}}$. The decorated vertex fixes the residual marked-fiber factor and its electron routing Z_e . Positivity gives the raw survival energy $E_{\text{raw}} = -(\hbar/\tau_*) \ln(1 - r)$, and the dressed electron identification is $m_e c^2 = (3/2)Z_e E_{\text{raw}}$. Thus τ_* and $L_* = c\tau_*$ follow without an independent clock or geometric diameter. Appendix H gives the finite action and both adversarial audits.

The reduced capacity functional can superficially resemble the scalar sector of a Brans–Dicke theory [45], but the resemblance is misleading. In the ordinary static branch it is the Einstein constraint action rewritten through $\delta S = -2S_\infty \Phi/c^2$, not a second scalar–tensor action. The open action questions are narrower and sharper: the transverse thermal influence functional, the generic covariant capacity observable outside controlled reductions, and the saturation-boundary functional.

Several tasks remain: derive the ensemble from a deeper Hamiltonian, embed the decorated transfer vertex in a stable geometric GFT condensate, audit the separate finite-loop stiffness return operator, compute the transverse metric kernel, and construct the strong-field boundary action. These are explicit limitations rather than hidden fit parameters.

A reader is therefore being asked to accept a small number of commitments, collected here in one place.

First, the vacuum has finite entanglement capacity: it is not empty background space but a medium with a bounded local capacity for entanglement.

Second, spacetime geometry and that capacity structure are the same substrate seen at different scales, so that in the weak field gravity is the fractional deficit of locally available capacity.

Third, matter is localized committed capacity: a particle is a stable defect of the medium, and its inertial mass is the entanglement content of that defect read in mass units.

Fourth, a recorded present is supported by many compatible microscopic pasts. The operational branch assigns probabilities only to decoherent record histories through the standard quantum decoherence functional, then conditions them on the realized present. It preserves Born statistics and no-signaling. The reversible renewal dilation supplies a concrete microscopic role for the history degrees of freedom: they receive the previous local state while the present register is renewed.

Fifth, the ultraviolet cell is the minimal tetrahedral boundary ensemble, and this is not a continuous tuning parameter. Given fermionic face data, maximum-capacity channel selection, injective four-face assignment, and minimality, the first admissible face alphabet has seven states, which gives the 1680-state ensemble (Section 5, Appendix B). The same faithful-resolution standard is applied to its histories: the local process carries the full available path entropy rather than reserving part of each pass as memory. “Maximum caliber” names that temporal application; it adds no independent commitment to the five listed here.

The finite-capacity substrate is the ultraviolet premise; geometry–capacity equivalence, mass–entropy equivalence, and Many-Pasts are the three postulates; and the tetrahedral ensemble is the ultraviolet architecture. Faithful full-support resolution acts on both states and histories. These inputs derive the admissibility weighting, $g_{\text{share,eff}}$, replacement kernel, one-layer native-vertex update, marked transfer, tree edge factor, weak-field bridge, and Newtonian metric branch inside the displayed decorated action. The separate loop-dressed stiffness, galactic acceleration scale, and RAR carry the conditional grades recorded in the closure table.

In one line, the architecture remains three postulates, one finite-capacity substrate premise, one minimal ultraviolet ensemble, and a small number of explicitly labeled conditional readings. The faithful-resolution rule applied to that same ensemble fixes both its state entropy and its maximum-caliber history process.

1.2 Physical Motivation for the Primitives

These commitments are primitives of the framework, not theorems proved from deeper assumptions in this paper. They are nevertheless not arbitrary. Each one is motivated by a place where established physics already strains against its own foundations, and each earns its place by making a known difficulty look less mysterious once it is adopted. None of it is offered as proof; it is the reason the starting points are reasonable ones to adopt before the derivation begins.

Mainstream gravitational physics has been converging on a finite-capacity substrate for decades. A black hole’s entropy scales with the area of its horizon rather than the volume it encloses, as though the contents of a region were written on its boundary; the Bekenstein bound limits the information a bounded region can hold; and holography and entanglement-based reconstructions of geometry tie the shape of spacetime directly to patterns of entanglement. Each of these is usually treated as a deep clue without a mechanism. The framework takes the clue literally: the vacuum is a medium with a finite local budget of entanglement, and geometry is the large-scale description of that budget. Several long-standing puzzles then become the ordinary behavior of a medium that can fill up. The area law is one. A black hole is a region whose capacity is exhausted, so there is nothing left in the interior to count, and the only live bookkeeping is at the surface where filled medium meets unfilled, which is why the entropy sits on the area and not in the volume. The black-hole information question softens for the same reason: if the interior holds no independent degrees of freedom, there is no interior store of information to lose, and what falls in is recorded in the boundary layer that later evaporation reads back out. And the old problem that a continuum theory of gravity develops infinities at short distances does not arise here, because below the smallest cell there is no continuum to diverge, and a finite state space has nothing to renormalize away.

The mass–entropy identification addresses a coincidence that general relativity encodes but does not explain. General relativity builds in the equality of inertial and gravitational mass geometrically, through the equivalence principle, but gives no microphysical account of why the mass that resists acceleration and the mass that sources attraction should be one and the same. Here both are readings of a defect’s committed entanglement. The induced gravitational scale contains the exact factor $Z_e^2 \ln^2(1 - e^{-7g_{\text{share,eff}}})$, whose dominant hierarchy is $e^{-14g_{\text{share,eff}}}$. The galactic extension has a lower closure grade. In its conditional transverse branch, $a_0 = \epsilon c H_0$ with $\epsilon \equiv g_{\text{share,eff}}/(4\pi^2)$ ties the onset scale to the cosmic horizon, while the one-capacity-invariant response ties the galactic field directly to the ordinary-matter source. The microscopic phase-cell normalization and influence kernel are the tests of that connection.

The scale-setting proposal distinguishes binding cost from recurrence rate. Seven channel entropies add, so their effective support multiplies. The state-weighted determinant turns that support into a record-conditioned transfer rate, while the positive survival operator gives the

rest-energy gap. The decorated marked vertex fixes the finite closure-response correction and the electron anchor fixes the clock.

The memoryless update rule is not selected merely because it is simple. It is uniquely forced by faithful sector resolution. For any stationary per-channel kernel with marginal p_{η_*} ,

$$H(B_{t+1} | B_t) = g_{\text{share,eff}} - I(B_t; B_{t+1}) \leq g_{\text{share,eff}},$$

so a dressing pass carries the full admissibility entropy if and only if the endpoint mutual information vanishes, which fixes the refresh kernel $K(b, b') = p_{\eta_*}(b')$. The allowed local single-label dynamics cannot realize this requirement: they freeze into disconnected sectors that never explore the full space (Appendix D.4), so a genuine nonlocal refresh is required. The native decorated vertex prepares the diagonal fresh amplitude and realizes the charged recurrence. Its geometric condensate embedding and durable history capacity remain open. The resulting proper-time scale is local and does not impose a preferred spatial lattice or predict frame-dependent propagation dispersion.

Many-Pasts treats the present as supported by a conditional measure over compatible decoherent histories. Unrecorded double-slit alternatives remain combined at amplitude level; a durable which-path record permits separate history probabilities. Entangled records retain the standard joint correlations and no-signaling marginals. The operational construction is therefore ordinary quantum mechanics with a history-space ontology. Its arrow-of-time extension still needs a substrate typicality theorem. Faithful sector resolution independently selects memoryless dressing, and Many-Pasts supplies the interpretive setting in which that process operates.

Tetrahedra are the standard building block in several independent approaches to quantum geometry, so the cell is a familiar object rather than one invented for this paper. What the framework adds is a set of its own restrictions that pick out one ensemble: the faces carry fermionic data, the maximum-capacity channel is selected, the four faces must be assigned distinct labels, orientations are counted with both parities, and the smallest admissible alphabet is taken. Those restrictions force a seven-label face structure and a boundary ensemble of exactly 1680 states. The same machinery that fixes the count also limits how many shell excitations the cell can carry before its closure structure degenerates, which terminates the charged-lepton ladder at three families, a candidate answer to the Standard Model's otherwise unexplained fact that matter comes in exactly three generations. How this particular ensemble came to be chosen, and how much freedom that choice had, is audited separately in Appendix L. Within the stated construction, its entropy is computed from the rules with no adjustable dial.

Several mainstream results already point toward the finite-capacity reading of gravity: the derivation of Einstein's equations as a thermodynamic relation of state, the identification of entanglement entropy with horizon area, the reconstruction of spatial connectivity from entanglement, and the conjectured identity between entanglement and geometric bridges. This framework is a concrete, finite, and falsifiable instance of that program.

The primitives are not proved before the theory begins, because no theory proves its own starting point, and a list of mysteries explained is the easiest kind of case to assemble after the fact, since a proposal of this scope can almost always produce one. The discriminator is not how much the primitives explain, but whether the explanations were forced by the construction or fitted to the target (the question Appendix L confronts directly), and whether the same construction then survives measurements it did not anticipate. This subsection claims only the weaker statement: the starting points are motivated by real and independent pressure points in gravity, quantum information, black-hole physics, and quantum measurement, and once adopted they constrain the weak-field sector tightly enough to be tested. The rest of the paper derives those constraints and confronts them with observation.

2. Canonical Field Content and Definitions

Before the symbols, four plain words recur throughout. *Capacity* is the entanglement support locally available in the medium. A *defect* is a stable, localized commitment of that capacity — what we coarse-grain into a particle. A *deficit* is capacity no longer freely available to the surrounding vacuum because a defect has committed it. *Strain* is the extended profile of that deficit reaching out into the medium, whose fractional size the weak-field potential tracks. The field variables below are the precise versions of these words.

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

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

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

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

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

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

Positive δS denotes reduced available vacuum entanglement capacity in the neighborhood of a localized defect or defect distribution. It is the extended capacity-strain 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. Each is used where it is most transparent: δS for the weak-field theory, because it talks directly to the Newtonian potential; q for the nonlinear and strong-field completion, because boundedness is built in from the start; and S_{ent} itself for the covariant EFT, because it is the field that appears in the action. The operational meanings are:

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

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

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

the observable bridge

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

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

Substrate length scale. The canonical UV cell length is not taken to be the conventional Planck length as an input. Faithful full-support resolution fixes the renewal kernel and its history-space factorization. The state-weighted determinant gives the seven-channel recurrence

$$r = e^{-7g_{\text{share,eff}}}, \quad L_*^{(0)} = -\frac{3}{2}\lambda_e \ln(1-r), \quad \lambda_e = \frac{\hbar}{m_e c}.$$

The decorated marked-transfer vertex derived in Appendix H fixes

$$\zeta_* = 9e^{-g_{\text{share,eff}}} \left(1 - \frac{8\eta_*}{49}\right)^{21/2}, \quad Z_e = (1 + \zeta_*)(1 + 7\zeta_*^2).$$

The physical electron-anchored scale is

$$\boxed{L_* = Z_e L_*^{(0)}} = 1.6162537014 \times 10^{-35} \text{ m}.$$

The corresponding induced gravitational scale is

$$G_* := \frac{c^3 L_*^2}{\hbar} = \frac{9}{4} \frac{\hbar c}{m_e^2} Z_e^2 \ln^2(1 - e^{-7g_{\text{share,eff}}}) = 6.6742890772 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}.$$

The lightest one-bit fermionic defect resolves the seven face sectors once and exports the transverse $2/3$ share of that dressing block. The marked vertex accounts for its finite closure-response fiber and label-return loop. The conventional Planck length $L_P = \sqrt{\hbar G/c^3}$ remains useful for comparison and for standard black-hole thermodynamic notation, but it is not the primitive scale-setting input here.

The principal coefficients and derived quantities used throughout are:

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

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

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

$$L_* : \text{substrate cell length in the electron-anchored support-to-rate map}, \quad (4)$$

$$G_* : \text{gravitational scale induced by } L_*, \quad (5)$$

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

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

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

$$a_0 = \frac{c H_0 g_{\text{share,eff}}}{4\pi^2} \quad \text{in the conditional compact two-phase normalization.} \quad (9)$$

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

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

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

3. The Three Postulates

The framework uses three postulates to state its ontology. Postulate I identifies geometry with the long-wavelength capacity substrate. Postulate II identifies matter with localized committed capacity and mass with its inertial reading. Postulate III assigns a history-space ontology to the decoherent pasts compatible with one realized present. The operational probabilities come from the standard decoherence functional, while the proposed arrow of time requires an additional typicality result. Faithful sector resolution selects the memoryless dressing kernel used by the decorated scale-setting action.

3.1 Information–Geometry Equivalence

The first postulate states that spacetime geometry is the continuum expression of the entanglement-capacity substrate. This is stronger than saying entanglement contributes an additional piece of stress-energy inside otherwise standard general relativity. The metric is not an independent background to which S_{ent} is appended; the geometric sector and the scalar capacity sector are two continuum projections of one finite-capacity medium. In the weak field, gravitational potential is the fractional deficit of available capacity.

Two consequences of this reading should be kept distinct from the start. First, absolute S_{ent} is not itself “the gravitational potential”; the observable weak-field potential comes from the fractional deficit $\delta S/S_\infty$, which is why a fixed-epoch rescaling of entropy units leaves gravity unchanged (Section 2). Second, because geometry and capacity are two descriptions of one response, the deficit is not an extra force appended to an independently existing metric. Section 10 proves that the ordinary reduced capacity functional is the Einstein constraint action in the capacity coordinate. The remaining common-parent problem concerns the transverse thermal and boundary sectors, not the baseline Newtonian response.

3.2 Mass–Entropy Equivalence

The second postulate is that mass is not something added to the substrate from outside but the inertial reading of localized capacity commitment. At scale ℓ ,

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

A particle is already a localized defect of the entanglement substrate, so $m = \kappa_m \Delta S$ does not assert an analogy between two independent things; it asserts that the inertial content of the defect *is* its entanglement content, read in mass units.

For elementary fermionic sectors the canonical defect increment is

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

The one bit here is not arbitrary. An elementary fermionic exclusion is binary — the face is occupied or unoccupied — and a binary distinction carries exactly $\ln 2$ of missing entanglement. This is the simplest possible defect increment, which is why the lightest such defect, the electron, becomes the cleanest anchor for the mass–entropy map (Section 13). For composite sectors the relevant quantity is the fully dressed bound-state entanglement budget, not a bare constituent count.

Two corollaries are used later. First, because mass and entanglement budget are two descriptions of the same defect, and the masses of separated defects add, capacity committed in service of one defect cannot simultaneously serve another: shared service would make the joint budget, and with it the joint mass, sub-additive. Commitment is therefore per-defect — each committed unit carries the label of the defect it serves. Second, the same per-defect bookkeeping makes

the saturated early phase countable: if mass is committed capacity, committed units cannot be double-counted across separated defects, and the abundance of committed capacity can be counted rather than fitted (Section 18.5).

3.3 Many-Pasts Hypothesis

The third postulate concerns the microscopic support of the present entanglement network. A recorded present can be compatible with many coarse-grained histories of the substrate. Many-Pasts takes those alternative pasts seriously while retaining one realized macroscopic present. Its probability theory must distinguish amplitudes for alternatives that still interfere, probabilities for recorded presents, and conditional probabilities for decoherent histories compatible with a given record.

The operational construction uses the decoherent-histories formalism [60]. A coarse history $h = (\alpha_1, \dots, \alpha_n)$ has class operator

$$C_h = \Pi_{\alpha_n}^{(n)} U_{n,n-1} \cdots \Pi_{\alpha_1}^{(1)} U_{1,0},$$

and decoherence functional

$$\mathcal{D}(h, h') = \text{Tr}(C_h \rho_0 C_{h'}^\dagger).$$

When a family decoheres, $\mathcal{D}(h, h') \simeq 0$ for $h \neq h'$, its diagonal entries obey the ordinary probability sum rules. If P denotes a final macroscopic record and $\mathcal{H}_P$ is a decoherent refinement of the histories ending in that record, then

$$p(P) = \sum_{h \in \mathcal{H}_P} \mathcal{D}(h, h) = \text{Tr}(\Pi_P \rho_{\text{now}}), \quad p(h | P) = \frac{\mathcal{D}(h, h)}{p(P)}.$$

The probability of the present is obtained from the common normalized measure over all records before the history distribution is conditioned on P . Unresolved alternatives remain combined at amplitude level; no positive probability is assigned to individual fine-grained paths that have not decohered.

This construction makes the operational branch standard quantum mechanics. Quantum instruments give the Born probabilities of laboratory records, and local trace-preserving instruments give no-signaling marginals. Many-Pasts changes their history-space interpretation without adding a collapse term or a signaling bias.

The probability measure is nevertheless not an arbitrary extra choice once the quantum kinematics are admitted. On a Hilbert space of dimension greater than two, a normalized, noncontextual, additive measure on a sufficiently rich lattice of record projectors has the form

$$\mu(\Pi) = \text{Tr}(\rho \Pi)$$

by Gleason's theorem [54]. A medium-decoherent family with a pure initial state admits orthogonal generalized records for its branch state vectors, so each history probability can be represented as the probability of a single-time record projector [55]. The Born form is therefore unique for record-defined decoherent families if record completeness and noncontextual additivity across compatible record refinements are imposed. This is a conditional uniqueness theorem inside quantum kinematics, not a derivation of those kinematics from the substrate. The Hilbert space, unitary dynamics, initial state, and record-projector richness remain imported, and a mixed state requires the corresponding purification or generalized-record construction.

The postulate and the renewal theorem have separate logical roles. Many-Pasts supplies the ontology and the record-conditioned history space. The already-stated faithful full-support

condition selects the memoryless dressing kernel when applied to paths at fixed admissibility marginal; “maximum caliber” names this history-space form and is independent of the Born measure on laboratory records. The lightest-defect functional selects the same absence of temporal memory and, within the support-to-length map, selects seven occupied channels with vanishing inter-channel correlation. A reversible dilation exports the old present to history. The decorated transfer vertex supplies the determinant recurrence, finite marked response, and Compton phase readout.

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 transport within the surrounding network. Because the substrate is isotropic, the cost of motion depends only on the rotational scalar v^2 at leading order, with the temporal rate maximal at $v = 0$ and vanishing when the budget is exhausted at $v = c$. These endpoint conditions alone admit many interpolating functions and so do not fix the form of the time-dilation relation. The form is fixed once the finite update speed is treated as invariant across inertial coarse descriptions: homogeneity, isotropy, and the relativity principle then select the Lorentz group rather than the Galilean one, giving the invariant interval

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

and hence

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

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

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

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

so smaller q means that less local update capacity remains available. Motion reduces the temporal share of the budget by consuming part of it in spatial transport; a nearby defect reduces the local budget by depleting available capacity. The two familiar time-dilation effects are therefore interpreted as two regimes of one mechanism.

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

The upshot is that the metric sector of the EFT is not being introduced from outside. Lorentzian geometry is the natural coarse description of a finite-capacity, isotropic, relational substrate, and the Einstein sector is its lowest-order continuum gravitational expression. The entanglement scalar then tracks how that same capacity geometry is redistributed by localized defects. The resulting low-energy theory can therefore be written in the usual covariant language, but the intended logic runs from substrate properties to geometry, not the other way around. As with any discrete substrate, this is a continuum statement, and it faces a sharp known obstacle: a discrete structure that defines a preferred rest frame feeds dimension-four Lorentz-violating operators into the infrared with order-unity coefficients through loops [25], against laboratory bounds many orders of magnitude below unity. The protection here is structural. The tetrahedral ensemble is combinatorial and pre-geometric: it lives in the state counting from which the continuum is constructed, defines no embedding lattice in the emergent spacetime, and imprints on the EFT only through frame-independent scalars (L_* , $g_{\text{share,eff}}$, η_*). Discreteness of this class is compatible with exact low-energy Lorentz symmetry, as causal-set sprinkling demonstrates by construction [26, 27]. The cosmological bath does select a frame, in the same environmental sense the CMB does: a state rather than an operator, while the laboratory bounds constrain operators. The supporting calculation this argument calls for — that substrate loop corrections generate no dimension-four Lorentz-violating operators — is open and recorded as such in the closure table.

4.2 Dependency Map of the Theory

The logical flow begins with the three foundational postulates — Information–Geometry, Mass–Entropy, and Many–Pasts — with faithful full-support resolution applied to both states and histories, and runs through the static weak-field chain before reaching the conditional sectors:

$$\begin{aligned} & \{\text{Info–Geometry, Mass–Entropy, Many–Pasts}\} \\ & \rightarrow \text{finite-capacity substrate ontology} \rightarrow \text{tetrahedral boundary ensemble} \\ & \rightarrow \text{faithful full-support resolution on states and histories} \rightarrow \text{local replacement / history expo} \\ & \rightarrow \text{edge kernel / loop dressing / stiffness / source map} \\ & \rightarrow \text{static capacity coordinate on the Einstein constraint sector} \rightarrow \text{Newton / baseline no-slip / GR} \\ & \text{with the one-invariant transverse EFT + horizon matching} \rightarrow \{a_0, \text{RAR}\} \text{ conditionally,} \quad \text{and transverse-s} \end{aligned}$$

Only then come the conditional and frontier sectors — transport, clusters, cosmology, strong field, and the particle/gauge extensions — each developed as a consequence or completion of the same framework.

The map is a dependency graph, not an equality of closure status. The ordinary static branch is closed more tightly than the transverse, cosmological, or strong-field sectors. Many–Pasts supplies the history-space ontology, while faithful full-support resolution selects the local renewal process on that space. Their combination gives a reversible present/history exchange; neither the Born history measure nor record conditioning alone selects memorylessness.

Part II. UV Coefficient Chain

Part I fixed what the theory is about. The question now is whether the local capacity-sharing structure can actually be *counted*. If the substrate has finite local capacity, the coefficients that appear in the continuum weak-field theory should not be free continuum parameters; they should descend from a finite local boundary problem. The next five sections follow that problem through: the smallest boundary cell that can carry capacity and close isotropically, the weighting that selects well-closed configurations, the cost of neighboring cells disagreeing, the local returns

that dress that cost, and the continuum coefficient they leave behind. Nothing in this chain is matched to a gravitational observable; the observables enter only in Part III.

The ultraviolet construction contains discrete choices rather than galaxy-by-galaxy fit parameters: the tetrahedral cell, seven labels, injective assignment, parity doubling, admissibility kernel, transverse export, and electron anchor. Section 5 and Appendix B derive the finite ensemble and its entropy; Appendix C derives the edge projection; Appendices D and H derive the replacement process, reversible history export, factorized lightest branch, and decorated marked-transfer action; Appendix L records the historical fork accounting. The remaining microscopic task is to embed that finite transfer vertex in a geometric GFT action with a stable condensate.

5. Why a Tetrahedral Boundary Ensemble

The problem is to find the smallest discrete boundary cell that can carry finite channel entropy, close isotropically, and hand a single scalar response up to the continuum. The smallest structure that meets all three needs is a tetrahedral cell, fixed by four ingredients:

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

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

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

Two cells sharing a face therefore generate the effective boundary sector

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

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

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

with

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

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

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

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

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

j_0 is sometimes mistaken for a free dial; it is not. Within this minimal construction j_0 is the smallest fermionic label that satisfies injectivity; a larger j_0 does not describe a fluctuation inside

the same cell but a different, larger boundary ensemble. The minimal theory therefore has no j to tune — it has the first value that closes.

In that sense the seven-state face sector is derived from fermionic face data, maximum-capacity channel selection, tetrahedral injectivity, and minimality. The derivation chain itself makes no reference to the weak-field observables it later feeds; the order in which the construction was historically found, and where the evidential weight accordingly sits, are recorded in Appendix L. 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

$$g_{\text{share,max}} = \ln(1680) = 7.42654907240.$$

The exact K^2 spectrum and multiplicities are carried in the appendices. The j -labeled tetrahedron used here coincides with the quantum tetrahedron of simplicial spin networks [51, 52], whose discrete geometric spectra [53] arise from the same $SU(2)$ representation theory; the present construction differs in weighting these states by admissibility closure rather than by a spin-foam amplitude, and in routing them to a capacity entropy rather than to area and volume operators. The UV theory thus begins with a finite microscopic counting problem rather than a free continuum ansatz.

6. Admissibility Closure

6.1 Minimal isotropic kernel

Not every boundary configuration should count equally. The raw combinatorial ensemble is too permissive to be the whole UV story: some configurations sit close to the regular closure pattern expected of a smooth local cell, while others are badly distorted. Admissibility closure is the statement, in its mildest form, that more poorly closed configurations contribute less to the coarse ensemble. The minimal rotationally invariant measure of that distortion is a single quadratic closure-defect scalar K^2 , and the weighting it induces is

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

This is not chosen because it works phenomenologically; it is the minimal isotropic maximum-entropy kernel under normalization and a fixed quadratic closure moment. Higher invariants such as K^4 carry additional UV information and so enter as subleading refinements, not as competing leading kernels.

6.2 Closure condition and uniqueness

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

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

and its stationary point gives the closure condition

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

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

On that spectrum it is unique,

$$\eta_* = 0.0298668443935.$$

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

6.3 Effective sharing entropy

The admissibility-weighted effective sharing entropy is

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

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

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

With η_* fixed, the effective sharing entropy carries no remaining freedom; the exact spectrum, multiplicities, and uniqueness proof are given in Appendix B.2.

7. Edge Kernel and Tree-Level Coupling

The admissibility closure says what one cell can carry. The edge kernel says how costly it is for two neighboring cells to carry different things, and that cost becomes stiffness in the continuum field: a medium whose cells resist disagreement strongly is one whose capacity deficits spread reluctantly. The same UV closure data fix this cost. 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 $2/3$ and gives the bare edge smoothness coupling

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

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 four tetrahedral channel directions into the isotropic continuum limit — the part of the disagreement that the scalar sharing channel actually carries.

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

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

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

Origin of the horizon target. The horizon target

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

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

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

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

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

and therefore

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

The role of σ_* is to match the admissibility-closed microscopic entropy normalization to the horizon-normalized scalar-field convention; it is the closure target fixed by that choice of branch. The rooted shell observable converges to that target rapidly enough that the nonlocal correction is already strongly constrained by small shell depth. The tetrahedral identity keeps this bridge controlled: the four discrete channel directions average to the correct isotropic tensor structure in the continuum limit, so the same combinatorial data that fixed admissibility also fix the tree-level transport. The shell hierarchy and phase-selection checks are carried 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 structure of that self-energy is not a generic loop number. The returns split into seven sector-diagonal channels and one collective mode. The seven are the face-label channels, each returning independently without mixing. The one is the permutation-symmetric combination across channels, which returns as a shared closure-singlet rather than as a channel-specific loop, and it is weighted by the same transverse projection and branch-dilution factors that define the tree edge map,

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

The leading local self-energy is therefore

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

Equivalently, on the seven-channel scalar return space,

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

with $\Sigma_{\text{ret}} = \text{Tr}(R_{\text{ret}})$. The orthogonal six-dimensional sum-zero sector carries no net scalar charge in the coarse branch and so adds no separate scalar return. The singlet weight is fixed by the tree map, not introduced here, so no new loop parameter appears. Permutation symmetry fixes the existence of the singlet but not the placement of the suppression factors on it alone; that placement is the minimal-return-operator reading whose graph-level derivation Appendix C.3 records as the outstanding audit. The induced uncertainty is bounded: replacing $\Sigma_{\text{ret}} = 65/9$ by 7, $7 + 2/3$, or 8 shifts $J_{\text{eff}}^{(\text{ren})}$ and γ by at most 0.5% and leaves G , a_0 , and κ/γ exactly unchanged, since $J_{\text{eff}}^{(\text{ren})}$ cancels in the source-to-stiffness ratio (Appendix C.5).

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

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 0.05% level.

The loop correction is no longer schematic: the finite renormalization is written as an explicit local self-energy. The remaining audit task is the independent graph-level derivation of the relative diagonal and singlet weights of the same scalar-return operator, not the introduction of any new loop parameter.

9. Continuum Stiffness and SI Normalization

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

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

where the sum runs over one sublattice representative a of each bipartite primitive cell and its four outgoing bonds $\hat{n}_i$, so each undirected nearest-neighbor edge is counted once (the convention of Appendix C.4). The microscopic four-cell is assigned the volume

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

as a coarse-graining convention: the abstract tetrahedral cell complex has no space-filling regular-tetrahedron Euclidean embedding (Appendix C.5), so cell volumes and face areas enter as normalization conventions of the coarse map, not as geometry supplied by the graph. 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 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

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

Physically, γ is the continuum stiffness of the entanglement-capacity field. A larger γ makes spatial gradients more costly and suppresses the capacity-deficit response to a given source; a smaller γ allows larger variations of the field. The faithful sector-resolution principle fixes L_* without using G . It is nevertheless useful to define the gravitational scale induced by this length,

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

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

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

This is the same algebra as the familiar Planck-cell rewrite, but read in the opposite direction: the substrate cell length induces the gravitational scale rather than being chosen by first inserting the measured value of G . Within the Euclidean-action and cell-volume conventions stated above, the SI-normalized stiffness coefficient is fixed; because the absolute normalization of an isolated scalar functional is conventional, the invariant content of this step is the ratio $\kappa/(\gamma S_\infty)$ that the weak-field matching of Section 11 consumes. In that ratio $J_{\text{eff}}^{(\text{ren})}$ cancels (Appendix C.5), so the loop-dressed coupling carries no content for the Newton normalization; its nontrivial input enters the stiffness itself and the dynamical and galactic sectors.

This completes the micro-to-continuum coefficient chain. The tetrahedral ensemble determines the effective sharing entropy; the edge kernel and loop dressing turn that entropy into a discrete stiffness; and the Euclidean matching turns the discrete stiffness into the continuum coefficient γ of the weak-field EFT.

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

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

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

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. Einstein Parent Action and the Reduced Capacity Frame

The ordinary longitudinal capacity branch does have a covariant parent action. It is not an Einstein action plus an independently varied capacity scalar. It is the ordinary metric action

$$I_0[g, \psi] = \frac{c^3}{16\pi G} \int_{\mathcal{M}} d^4x \sqrt{-g} (R - 2\Lambda) + I_{\text{GHY}}[g] + I_{\text{matter}}[g, \psi],$$

where covariant coordinates use $x^0 = ct$. When the time integral is instead written in seconds, $dx^0 = c dt$ and the ADM prefactor is correspondingly $c^4/(16\pi G)$. Matter is coupled once, to one physical metric. The capacity functional is the static scalar-constraint reduction of I_0 , expressed in a different field coordinate. This statement can be proved without appealing to the final Poisson equation.

Static scalar reduction of Einstein–Hilbert gravity. Use Newtonian gauge,

$$ds^2 = - \left(1 + \frac{2\Phi}{c^2}\right) (dx^0)^2 + \left(1 - \frac{2\Psi}{c^2}\right) \delta_{ij} dx^i dx^j, \quad x^0 = ct,$$

and retain the static scalar sector through quadratic order. In ADM variables [7] the shift and extrinsic curvature vanish, so the Einstein–Hilbert plus Gibbons–Hawking–York action is

$$I_{\text{ADM}}^{\text{static}} = \frac{c^4}{16\pi G} \int dt d^3x N \sqrt{h} {}^{(3)}R + I_{\text{matter}}^{\text{static}} + I_{\infty}.$$

Expanding $N = 1 + \Phi/c^2$ and $h_{ij} = (1 - 2\Psi/c^2)\delta_{ij}$, cancelling the reference boundary term at infinity, and using $I_{\text{matter}}^{(1)} = - \int dt d^3x \rho \Phi$, gives

$$I_{0,\text{scal}}^{(2)}[\Phi, \Psi] = \int dt d^3x \left[\frac{1}{8\pi G} ((\nabla\Psi)^2 - 2\nabla\Phi \cdot \nabla\Psi) - \rho\Phi \right].$$

The lapse perturbation remains a constraint variable. Its variation and the spatial-scalar variation give, respectively,

$$\nabla^2\Psi = 4\pi G\rho, \quad \nabla^2(\Phi - \Psi) = 0.$$

Asymptotic flatness removes the harmonic difference, so $\Phi = \Psi$. Eliminating Ψ therefore produces

$$I_{\text{Newton}}[\Phi] = \int dt d^3x \left[-\frac{(\nabla\Phi)^2}{8\pi G} - \rho\Phi \right].$$

Thus no independent scalar stress tensor is needed to create the linear potential, and the equality $\Phi = \Psi$ in the baseline branch is a metric constraint equation rather than an anisotropic-stress assumption. Appendix N gives the expansion, boundary bookkeeping, and degree-of-freedom audit in full.

Exact reduced-action identity. On the renormalized static branch, write $S_{\text{ent}} = S_{\infty} - \delta S$. After source-independent terms are removed, the capacity functional is

$$I_{\text{cap}}^{\text{static}}[\delta S; \rho] = \int dt d^3x \left[-\frac{\gamma}{2} (\nabla \delta S)^2 + \kappa \rho \delta S \right].$$

The field redefinition

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

turns it into

$$I_{\text{cap}}^{\text{static}} = \int dt d^3x \left[-\frac{2\gamma S_{\infty}^2}{c^4} (\nabla \Phi)^2 - \frac{2\kappa S_{\infty}}{c^2} \rho \Phi \right].$$

Using

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

gives the action-level equality

$$I_{\text{cap}}^{\text{static}} = Z_S I_{\text{Newton}}, \quad Z_S \equiv \frac{2\kappa S_{\infty}}{c^2}.$$

The field-independent factor Z_S cannot affect the classical reduced equations. It can matter when the UV construction is asked to normalize fluctuations or correlation functions, but it does not represent a second determination of G and it does not license adding I_{cap} to I_0 . The equality assumes the same asymptotically flat or Dirichlet boundary data on both sides. At a finite boundary the Newton surface term and its image under $\delta S = -2S_{\infty}\Phi/c^2$ must be included as well.

Scope of the background-covariant notation. For transport calculations the same reduced equation is packaged as

$$I_{\text{cap}}[S_{\text{ent}}; \chi | g_{\text{ref}}] = \int d^4x \sqrt{-g_{\text{ref}}} \left[-\frac{\gamma}{2} g_{\text{ref}}^{\mu\nu} \partial_{\mu} S_{\text{ent}} \partial_{\nu} S_{\text{ent}} - \lambda S_{\text{ent}} - \kappa \chi S_{\text{ent}} \right].$$

The vertical bar is essential: g_{ref} and the reduced source projection χ are held fixed while S_{ent} is varied. This notation is useful for extending the reduced response in time, but it is not a covariant scalar-tensor parent action. In the static nonrelativistic sector $\chi \simeq \rho$; covariantly the source is the full stress tensor through $I_{\text{matter}}[g, \psi]$.

Capacity coefficients and the source theorem. The UV calculation still fixes how the geometric constraint is coordinatized by the substrate variable. With

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

the Green-matched projection is

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

This is the microscopic map between defect density and the capacity coordinate on the Einstein constraint surface. It is not an extra matter coupling in the covariant parent theory.

The length backbone of Newton's constant. The weak-field normalization is

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

Faithful full-support resolution and the positive marked-transfer spectrum fix

$$L_* = -\frac{3}{2} Z_e \lambda_e \ln(1 - e^{-7g_{\text{share,eff}}}), \quad \lambda_e = \frac{\hbar}{m_e c},$$

and therefore induces

$$G_* = \frac{c^3 L_*^2}{\hbar} = \frac{9}{4} \frac{\hbar c}{m_e^2} Z_e^2 \ln^2(1 - e^{-7g_{\text{share,eff}}}).$$

Substituting the source-map identities

$$\frac{\kappa}{\gamma} = \frac{3L_*}{4G_{\text{tet}}(0)\kappa_m(L_*)}, \quad \kappa_m(L_*) = \frac{\hbar}{cL_* \ln 2}, \quad S_\infty^{\text{cell}} = \frac{3 \ln 2}{32\pi G_{\text{tet}}(0)},$$

into the weak-field expression gives identically

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

Thus there is one scale-setting route to G : electron recurrence fixes L_* , and L_* fixes the gravitational scale. The stiffness, source coefficient, and capacity normalization are a consistent static-EFT representation of that same length backbone, not a second determination that could have disagreed with it. The only numerical comparison in this sector is G_* against the measured Newton constant.

The decorated scale gives $G_* = 6.6742890772 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}$, or -0.073σ relative to CODATA. Because both the early entropy construction and the later residual-closing vertex were developed with the discrepancy known, Appendix L treats this as a high-precision postdiction. The nontrivial content is the shared action that also fixes the two charged-lepton corrections.

11. Capacity Variable, Bridge Law, and Variational Status

Varying the reduced capacity-frame functional with respect to S_{ent} gives

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

On the renormalized static, nonrelativistic branch this becomes

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

Define the surviving fractional capacity

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

The action reduction above fixes the weak-field bridge directly,

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

Equivalently,

$$q = 1 + \frac{2\Phi}{c^2} + O(c^{-4}).$$

The bounded nonlinear rule

$$N^2 = q$$

is the continuous multiplicative completion selected by the capacity-composition rule. Its status must now be stated more precisely. In a static spherical exterior, q is the invariant geometric scalar

$$q = h^{ab} \partial_a R \partial_b R = 1 - \frac{2GM_{\text{MS}}}{c^2 R},$$

and in Schwarzschild coordinates it equals N^2 . In a generic spacetime the lapse is foliation dependent, so $N^2 = q$ by itself is not a covariant constraint. The metric-only parent therefore treats the weak-field δS and the spherical q as reduced or composite geometric variables; it does not promote either to an unconstrained second gravitational field. Appendix N.6 formulates the remaining target as a diffeomorphism-invariant, potentially quasilocal and state-dependent functional, rather than presuming that an additional fundamental scalar is needed.

Combining the static source equation with the weak-field bridge gives

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

This relation uses only the invariant combination $\kappa/(\gamma S_\infty)$: a fixed-epoch rescaling of the entropy units changes S_∞ and κ together and leaves the observable potential unchanged.

Two coordinates on one response. The metric parent solves the Hamiltonian and spatial constraints for Φ and Ψ ; the capacity frame uses δS as a field coordinate on that reduced solution. Around a constant background, a canonical scalar stress would begin as $(\partial \delta S)^2 = O(\rho^2)$ and could not be the source of the observed $O(\rho)$ potential. The action identity removes that mismatch: the linear capacity response is the reduced metric constraint itself.

Matter enters once. For the nonrelativistic static branch, the microscopic source theorem reduces the full metric source to the defect density ρ . Covariantly, matter enters only through $I_{\text{matter}}[g, \psi]$, so the full stress tensor gravitates, including trace-free radiation, and the Bianchi identity enforces the usual conservation law. The notation $\chi \simeq \rho$ belongs only to the reduced nonrelativistic source map; an explicit universal term $S_{\text{ent}} T^\mu{}_\mu$ is neither required nor adopted.

Parent-action decision. The preferred minimal construction for the ordinary branch is therefore metric-only:

$$I_{\text{parent}}^{\text{long}} = I_0[g, \psi], \quad \delta S = \delta S[g, \psi] \quad \text{after constraint reduction.}$$

It propagates the two tensor polarizations of general relativity and no extra scalar. Constrained-clock and scalar-tensor alternatives remain useful control cases, but both add structure and generically add a mode; Appendix N records why neither is selected. The transverse thermal sector responsible for the galactic excess is not included in this closure and requires its own metric influence functional.

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. This is the first point where the capacity-strain picture becomes ordinary gravity. A point defect produces a $1/r$ capacity deficit; the bridge reads the gradient of that deficit as the Newtonian inverse-square force. Ordinary gravity is the small-deficit, weak-curvature limit of the extended capacity strain around localized defects, so the $1/r^2$ law is emergent, not fundamental. Section 23 exhibits the same structure on the discrete substrate: a capacity budget that merely re-equilibrates cell by cell screens at sub-cell range, while a conserved capacity current with maintenance sinks obeys the massless graph-Poisson equation and produces this $1/r$ deficit: the Newtonian form doubles as a diagnostic of the conservation law behind it.

13. Electron Anchor: One-Bit Mass Scale and Seven-Sector Length Scale

The electron supplies the clean mass anchor $m_e/\ln 2$. Its Compton scale also assigns the length associated with the seven-sector effective support. The decorated marked-transfer action closes the finite correction to that scale and routes the same correction through the heavier charged-lepton shells. These are two readings of one dimensionful datum, not independent measurements, so the weak-field normalization remains a single calibration.

13.1 Why the electron is the anchor

The mass-entropy map needs a clean elementary anchor because, in this framework, the elementary matter sector *is* the localized defect sector. The electron is the natural choice: it is the lightest simple charged fermionic defect, not a composite, and its mass is not obscured by hadronic or QCD dressing. A single fermionic face-exclusion defect carries the canonical increment

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

one bit of missing entanglement, because an excluded face is a binary occupied/unoccupied defect of the local network.

13.2 One-bit mass anchor

At the electron Compton scale $\ell = \lambda_e$ the mass-entropy map reads

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

Dividing the electron mass by the fixed one-bit increment fixes the mass-per-entropy conversion at the electron's own scale. Run back to the cutoff cell, the conversion is

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

with canonical running law

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

in the closed branch. One bit fixes the electron-scale conversion; the running law carries it to the UV scale; the same conversion then feeds the weak-field source map. This is the only point at which the mass anchor enters gravity.

13.3 Seven-sector length anchor

The same electron anchors the cell length through the marked support-to-rate map. Faithful full-support resolution uniquely selects the memoryless kernel. Within the stated recurrence mass functional, fermionic exclusion and electron lightness jointly select the maximum channel count $k = 7$ and vanishing inter-channel correlation $\Delta_7 = 0$. On the commutative renewed-state algebra, the state-weighted determinant of the likelihood multiplication operator is

$$\Delta_{\tau_p}(\mathbf{R}) = \exp[\tau_p(\ln \mathbf{R})] = e^{-g_{\text{share},\text{eff}}}.$$

Multiplicativity across the seven factorized clouds gives

$$r = \Delta_{\tau_p^{\otimes 7}}(\mathbf{R}^{\otimes 7}) = e^{-7g_{\text{share},\text{eff}}}.$$

This is the record-conditioned geometric transfer rate, not the collision probability of two independently sampled blocks. The positive survival spectrum gives the baseline scale

$$L_*^{(0)} = -\frac{3}{2}\lambda_e \ln(1 - r),$$

with 3/2 the transverse export factor of Appendix C.5.

Appendix H realizes this determinant transfer in a finite charged action. The construction separates the baseline recurrence from the defect-bound closure response, so the one-channel mixing projector, seven-channel survival operator, and marked-fiber determinant remain distinct.

This does not mean an electron is a single tetrahedral cell carrying seven simultaneous labels. The local ensemble supplies seven distinguishable dressing layers. Each layer has fermionic occupation at most one, and the recurrence mass minimum occupies all seven once. The electron is the lightest coherent one-bit defect on that selected branch and exports the transverse share of the resulting support. The Compton scale calibrates the substrate length hierarchy; the one-bit mass calibrates the source map. The two uses impose a nontrivial joint requirement on the electron's role in the gravitational normalization without duplicating a single input.

13.4 Decorated marked-transfer vertex

The baseline comparison supplied a structural clue before it supplied a formula. The original Newton normalization was low by about 1.05%, but $G_* \propto L_*^2$, so the missing amplitude in the substrate length was only 0.5309%. The baseline muon ratio required a 0.5124% uplift. The tau ratio required 0.6562%, which separates into the same universal uplift and a smaller 0.1431% second-shell factor. The signs and amplitude scale therefore pointed to one small charged response with shell-dependent routing, rather than three unrelated corrections.

That response also had to be additive. The tetrahedral ensemble already fixed the entropy, closure spectrum, edge projection, source map, and the micro-to-macro coefficient chain. Changing its label count, admissibility weight, or transverse export to repair the residuals would move results that did not share the discrepancy. The admissible repair was consequently required to vanish in the unmarked vacuum, reuse the established closure incidence, and act only on the charged transfer graph. This criterion motivates the marked vertex below; it does not determine the answer numerically. The field content, determinant power, and routing still have to follow from the displayed action and survive the alternative-kernel audits.

The renewed closure amplitude has the exact Gaussian representation

$$e^{-\eta_* C^2/2} = \int \frac{d^3\xi}{(2\pi)^{3/2}} \exp\left[-\frac{1}{2}\xi^2 + i\sqrt{\eta_*}\xi_a C_a\right].$$

The amplitude-level closure incidence is therefore $\sqrt{\eta_*}$. For each occupied unordered channel pair $e = (m, m')$, the two directed scalar returns have row operator

$$R_e = \frac{2}{7}(\langle m \rightarrow m' | + \langle m' \rightarrow m |).$$

The contraction $B_e = \sqrt{\eta_*}R_e$ obeys

$$B_e B_e^\dagger = 2\eta_* \left(\frac{2}{7}\right)^2 = \frac{8\eta_*}{49} \equiv u.$$

Its canonical unitary dilation has no-event amplitude $\sqrt{1-u}$. The seven-channel lightest branch activates all $\binom{7}{2} = 21$ pair records, giving

$$Z_{\text{edge}} = \left(1 - \frac{8\eta_*}{49}\right)^{21/2}.$$

The present and history Gaussian strands each have three normalized first excitations. Their ordered products form nine orthonormal marked states, so tracing one hard-core marked fiber gives

$$\boxed{\zeta_* = 9e^{-g_{\text{share,eff}}} \left(1 - \frac{8\eta_*}{49}\right)^{21/2}} = 0.005123584484947.$$

The finite transfer graph has a universal marked alternative, one two-vertex label return, and one second-shell singlet passage. Its exact factors are

$$Z_\mu = 1 + \zeta_*, \quad Z_e = (1 + \zeta_*)(1 + 7\zeta_*^2), \quad Z_{\tau,2} = 1 + \frac{2}{7}\zeta_*.$$

The dressed scale is

$$\boxed{L_* = Z_e L_*^{(0)}}.$$

Appendix H writes the controlled vertex, decomposes its complete edge Hessian, and enumerates all 56,800 states of the one-step routing block.

The anchor is a gauge choice. Because counting fixes only dimensionless quantities, exactly one dimensionful measurement must be supplied, and which one is a convention. The physical content of the framework is carried by anchor-invariant statements: the lepton ratios m_μ/m_e and m_τ/m_e ; the hierarchy

$$\frac{m_e}{m_P} = -\frac{3}{2} Z_e \ln(1 - e^{-7g_{\text{share,eff}}});$$

the horizon normalization identity $1/4$ within the stated cell convention; the abundance ratio Ω_c/Ω_b ; and, within its branch assignment, a_0/cH_0 . The electron remains the anchor because its mass is measured most precisely. Appendix L records that the high-precision marked correction was constructed after the residuals were known, so anchor invariance does not turn the agreement into a blind prediction.

13.5 Consistency checks

The marked action uses no continuously fitted coefficient. Substitution gives

$$G_* = 6.6742890772 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2} \quad (-0.073\sigma),$$

$$\frac{m_\mu}{m_e} = 720 \frac{2}{7} Z_\mu = 206.768280237 \quad (-0.535\sigma),$$

$$\frac{m_\tau}{m_e} = 720^2 \left(\frac{2}{7}\right)^4 Z_\mu Z_{\tau,2} = 3477.343310 \quad (+0.481\sigma).$$

These three comparisons probe one marked-fiber weight through different graph polynomials. The scalar stiffness, source map, and weak-field bridge reproduce the same dressed G_* after substitution, so they remain one consistency chain rather than an independent determination.

The same comparison has a dimensionless form. Squaring the hierarchy gives the electron's gravitational coupling,

$$\frac{G_* m_e^2}{\hbar c} = \frac{9}{4} Z_e^2 \ln^2(1 - e^{-7g_{\text{share,eff}}}),$$

with $9/4$ the square of the transverse export factor and Z_e^2 the marked-transfer dressing. The relation also has an area reading. A horizon stores one bit in area $4 \ln 2 L_P^2$, while the electron spreads its single bit over the Compton area λ_e^2 . The marked factor changes the finite packing correction without altering the dominant hierarchy $e^{-14g_{\text{share,eff}}}$.

13.6 Composite sectors

For composite hadrons the claim is weaker and different in kind. 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 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 is structural compatibility between the mass-entropy map and the standard QCD mass budget. The elementary-fermion anchor is settled in the simple sectors; the hadronic sector remains structurally compatible but not yet coefficient-complete.

14. Galactic Dynamics

The static Einstein branch supplies the baryonic acceleration g_{bar} . The proposed galactic extension adds the equilibrium response of a transverse substrate sector. The calculation below

separates the algebra that is fixed inside the effective model from the microscopic statements still awaiting a specified GFT interaction and metric vertex. In particular, the exponential response follows once the transverse gap and normalization are supplied; the present paper does not claim that those inputs have already been derived from a standard simplicial vertex.

Vacuum-state origin of the Bose–Einstein occupancy. For a bosonic mode in a de Sitter static patch, the restricted vacuum state is thermal at

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

and its occupation is $n_B(x) = 1/(e^x - 1)$ for the dimensionless thermal argument $x \equiv E/(k_B T_H)$. For the present apparent horizon, $R_A = c/H_0$ in the branch used here, so

$$a_H = cH_0.$$

The transverse doublet supplies a sharper origin for the phase measure than dimension counting alone. For each stable normal mode the two real components Y_1, Y_2 and their conjugate momenta form two canonical oscillator pairs. Passing to action–angle variables (J_i, θ_i) gives

$$\Omega_{\text{symp}} = dJ_1 \wedge d\theta_1 + dJ_2 \wedge d\theta_2, \quad \theta_i \sim \theta_i + 2\pi.$$

Canonical quantization spaces neighboring actions by $\Delta J_i = \hbar$, so one two-oscillator quantum cell has invariant phase-space volume $(2\pi\hbar)^2$. Restricting to one quantum action cell leaves a dimensionless angular Haar volume $(2\pi)^2$. The factor is therefore fixed by the symplectic two-oscillator structure once that mode and its one-action-cell loading are specified. A real transverse two-vector viewed only as a configuration-space direction would supply one polar angle and would not prove the result.

The proposed transverse normalization assigns one sharing entropy $g_{\text{share,eff}}$ to that one-action-cell angular torus. Its entropy per unit invariant phase volume is

$$\epsilon_{\perp} = \frac{g_{\text{share,eff}}}{(2\pi)^2}.$$

Equating the reversible horizon work $k_B T_H \epsilon_{\perp}$ with the Unruh energy $\hbar a_0/(2\pi c)$ [37] gives

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

The algebra is exact after the loading assignment. The compact angular measure follows from the canonical transverse doublet, while two physical premises remain: the infrared state must load one sharing entropy into one quantum action cell, and that loading must couple reversibly to the horizon state. The 2π factors in the Unruh and Gibbons–Hawking temperatures cancel in their dimensionless ratio, so they do not independently generate the cell normalization. With the Planck value of H_0 , the branch gives $a_0 = 1.231 \times 10^{-10} \text{ m s}^{-2}$, compared with the RAR scale $(1.20 \pm 0.02) \times 10^{-10} \text{ m s}^{-2}$ [1, 40]; the displacement is +2.6%, about 1.5σ using the quoted observational uncertainty. If the loading and horizon coupling hold, the model also predicts $a_0(z) \propto H(z)$.

1 + 2 channel decomposition and the radial-acceleration law. The baryonic gradient selects a local longitudinal direction. The remaining two components form a transverse doublet $Y = (Y_1, Y_2)$, and its source-visible scalar is the radial occupation

$$R^2 = Y_1^2 + Y_2^2.$$

After the refresh-gapped relative modes are integrated out, the minimal one-invariant truncation describes longitudinal commitment X and transverse occupation R^2 through one local capacity mismatch,

$$\Delta C(X, R) = \alpha_C + g_C X + \beta_C R^2.$$

Finite susceptibility gives an effective potential $V = F(\Delta C)$ with $F'(0) = 0$ and $u_C \equiv F''(0) > 0$. Its leading expansion is $V = u_C(\Delta C)^2/2 + O(\Delta C^3)$. This form follows from a Gaussian large-deviation expansion once the low-energy theory has only one total-capacity variable. The existence of that one-invariant infrared truncation is the physical premise; finite capacity alone would not exclude additional light invariants.

Let (X_0, R_0) lie on the stationary surface $\Delta C = 0$. The quadratic Hessian in $(\delta X, \delta R)$ is

$$H = u_C \begin{pmatrix} g_C^2 & 2g_C\beta_C R_0 \\ 2g_C\beta_C R_0 & 4\beta_C^2 R_0^2 \end{pmatrix} = u_C \mathbf{c}\mathbf{c}^T, \quad \mathbf{c} = (g_C, 2\beta_C R_0)^T.$$

Writing its entries as A_L , A_T , and $C_\times$, and choosing the relative field orientation so that $C_\times \geq 0$, gives

$$C_\times^2 = A_L A_T, \quad \det H = 0.$$

The geometric mean is therefore the cross coefficient of a rank-one capacity Hessian. It is not a generic normal-mode frequency of two coupled oscillators. The zero eigenvector is tangent to $\Delta C = 0$ and transfers occupancy between longitudinal and transverse sectors without changing the total. Positive gradient terms make that redistribution mode dispersive at nonzero wave number.

Two further consequences follow exactly from the same quadratic form. Writing

$$\delta^2 V = \frac{1}{2} (A_L \delta X^2 + 2C_\times \delta X \delta R + A_T \delta R^2),$$

the transverse coordinate that minimizes the energy at fixed δX is

$$\delta R_*(\delta X) = -\frac{C_\times}{A_T} \delta X.$$

Thus $C_\times/A_T$ is the exact static response coefficient in the one-invariant model. With the common kinetic normalization $Z_L = Z_T = Z_0$ derived below, the clamped curvatures define $\omega_L^2 = A_L/Z_0$ and $\omega_T^2 = A_T/Z_0$, and rank one gives

$$\frac{C_\times}{A_T} = \sqrt{\frac{A_L}{A_T}} = \frac{\omega_L}{\omega_T}.$$

These are clamped frequencies, not the eigenfrequencies of the freely relaxing rank-one system; the adiabatic tangent mode remains gapless at zero wave number.

At leading two-derivative order, tetrahedral symmetry gives the three-component capacity representation one kinetic coefficient. Appendix N derives $Z_L = Z_T$ from the decomposition $4 = 1 \oplus 3$. Equal kinetic normalization removes an otherwise arbitrary relative field rescaling. The further effective matching

$$A_L = \frac{g_{\text{bar}}}{cH_0}, \quad A_T = \frac{a_0}{cH_0}$$

then yields the dimensionless curvature ratio

$$x = \frac{C_\times}{A_T} = \sqrt{\frac{A_L}{A_T}} = \sqrt{\frac{g_{\text{bar}}}{a_0}}.$$

This matching identifies the two canonical inverse susceptibilities with their Unruh-to-horizon energy ratios. The response lemma above reduces the remaining thermal identification to one anchoring condition if the source-driven energy is read from the clamped longitudinal curvature: with $E_\perp \equiv \hbar\omega_L$,

$$\frac{E_\perp}{k_B T_H} = \frac{\omega_L}{\omega_T} \iff \hbar\omega_T = k_B T_H.$$

The curvature ratio is therefore not an additional free equality on top of the clamped-oscillator reading. The physical question is whether the transverse retarded response is governed by these clamped curvatures and whether its transverse anchor sits at the thermal point. The rank-one Hessian alone does not answer that question: its freely relaxing tangent mode is gapless and its orthogonal eigenvalue is proportional to $A_L + A_T$. The full condensate influence functional must distinguish the clamped response from the adiabatic poles.

The resulting radial-acceleration law is

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

with the asymptotic limits

$$g_{\text{bar}} \gg a_0 \implies g_{\text{obs}} \approx g_{\text{bar}}, \quad (10)$$

$$g_{\text{bar}} \ll a_0 \implies g_{\text{obs}} \approx \sqrt{a_0 g_{\text{bar}}}. \quad (11)$$

The two endpoint limits follow from the occupancy factor. The exact interpolation is conditional on the effective transverse matching and on the absence of an additional form factor in the microscopic influence kernel. The one-invariant truncation removes extra light field-space directions at zero derivative, but it does not remove the momentum dependence of a propagating field or the spectral density generated by gradient terms and continuum modes. It therefore cannot prove $F(\omega, \mathbf{k}) \equiv 1$ by itself. Empirically, any residual static form factor is bounded below the few-percent level across the five decades of x probed jointly by solar-system precision and the measured flatness of the galactic rotation and lensing relations. The deep-MOND branch gives the baryonic Tully–Fisher law [38, 39]

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

Division of labor between the action and the state. The ordinary Einstein response supplies the baseline multiplier 1. For a bosonic transverse mode with source-dependent energy $E_\perp$ and $x = E_\perp/(k_B T_H)$, the equilibrium thermal free energy is

$$F_{\text{th}}(E_\perp) = k_B T_H \ln(1 - e^{-x}).$$

Its derivative is

$$\frac{\partial F_{\text{th}}}{\partial E_\perp} = \frac{1}{e^x - 1} = n_B(x).$$

Forward and reverse thermal processes are already contained in this partition function. No blocked-absorption rule or thermodynamic-arrow argument is needed. If the source-gap coupling converts this derivative into the transverse acceleration response without an additional form factor, the total multiplier is

$$\nu(x) = 1 + n_B(x) = \frac{1}{1 - e^{-x}}, \quad g_{\text{obs}} = g_{\text{bar}} \nu(x).$$

With $x = \sqrt{g_{\text{bar}}/a_0}$ this gives the law above.

The same relation admits a conservative static response potential. Define $W'(g) = \nu(\sqrt{g/a_0})$ and set $x = \sqrt{g/a_0}$. Since $dg = 2a_0x dx$,

$$W(g) = a_0[x^2 + 2x \ln(1 - e^{-x}) - 2\text{Li}_2(e^{-x})] + \text{const},$$

and direct differentiation returns $W'(g) = 1/(1 - e^{-x})$. The existence of this potential shows that the exponential RAR is compatible with equilibrium statistical mechanics. It does not determine the relativistic metric response.

Closure status of the galactic branch. The statements in this branch have different grades. The rank-one identity $C_\times^2 = A_L A_T$, the static response $\delta R_*/\delta X = -C_\times/A_T$, and the clamped-frequency ratio $C_\times/A_T = \omega_L/\omega_T$ follow exactly from a one-invariant capacity potential with equal kinetic normalization. Tetrahedral symmetry fixes $Z_L = Z_T$ at leading two-derivative order. The assignments $A_L = g_{\text{bar}}/(cH_0)$ and $A_T = a_0/(cH_0)$ remain conditional and require the horizon state to keep the combination $4u_C\beta_C^2 R_0^2$ environment-independent. Once the clamped-oscillator reading is adopted, thermal matching reduces to the single anchor $\hbar\omega_T = k_B T_H$; the influence functional must still show that the physical response uses the clamped rather than adiabatic spectrum, a choice most consequential in the deep-MOND regime where $\omega_L \ll \omega_T$. Canonical quantization fixes the $(2\pi)^2$ angular measure of the two-oscillator cell, while its one-entropy loading and reversible horizon coupling remain conditional. Finally, $g_{\text{obs}} = g_{\text{bar}}[1 + n_B(x)]$ requires the influence functional to introduce no additional source-dependent vertex or spectral factor. Until those dynamical tests are passed, the RAR is an exact consequence of a specified leading effective completion rather than a closed consequence of the UV ensemble.

15. Baseline Metric Closure and the Open Galactic Lensing Kernel

Two questions that were previously bundled together must be separated.

Ordinary longitudinal branch. For the branch that reproduces Newtonian gravity, the physical metric is already fixed by the Einstein parent action. The scalar constraint reduction of Section 10 gives

$$\nabla^2(\Phi - \Psi) = 0,$$

so asymptotic flatness implies

$$\Phi = \Psi.$$

This is an action-level result. It is not inferred from the canonical stress tensor of δS , because no independently gravitating capacity scalar is present. Since the parent action of this branch is exactly Einstein–Hilbert plus minimally coupled matter, its vacuum post-Newtonian solution has

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

and the remaining standard PPN parameters vanish, subject to the usual assumptions on the matter sector and boundary conditions. The capacity redefinition does not alter those values. Here *closed* refers only to the baseline contribution. For the full theory one should write schematically

$$\gamma_{\text{PPN}}^{\text{obs}} = 1 + \delta\gamma_\perp, \quad \beta_{\text{PPN}}^{\text{obs}} = 1 + \delta\beta_\perp.$$

The observed PPN parameters are closed only after the transverse calculation shows that these corrections vanish, are screened, or lie below Solar-System bounds.

Transverse galactic branch. The low-acceleration excess is a different action problem. The rate calculation fixes

$$g_{\text{obs}} = \frac{g_{\text{bar}}}{1 - \exp[-\sqrt{g_{\text{bar}}/a_0}]},$$

but a radial force law does not determine the spatial metric. To predict photon deflection one must integrate out the transverse substrate modes in their physical horizon-thermal state and derive the metric response. The appropriate object may be a closed-time-path influence functional,

$$\Gamma_{\text{eff}}[g_+, g_-; \psi_+, \psi_-] = I_0[g_+, \psi_+] - I_0[g_-, \psi_-] + \Gamma_{\perp}[g_+, g_-; \rho_{\perp}],$$

rather than a local equilibrium action. Its retarded metric kernel must generate the RAR in the static matter channel, its spatial components must determine Ψ , and its Ward identity must enforce covariant conservation. Appendix N specifies the minimal Hessian and correlator needed to decide this.

If that calculation gives equal transverse corrections,

$$\Delta\Phi_{\perp} = \Delta\Psi_{\perp},$$

then the excess may be represented by the effective density

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

and dynamics and lensing share the same kernel. At present this equality is a falsifiable target, not a theorem.

Local tests. The RAR correction is exponentially suppressed when $g_{\text{bar}} \gg a_0$, so the closed Einstein branch dominates in the Solar System. This explains why the target PPN limit is the GR one, but it is not a substitute for deriving the transverse kernel and checking that it generates no residual preferred-frame or nonlocal effect. Matter still couples to one physical metric and follows its geodesics; there is no additional test-body scalar charge. Weak-lensing RAR measurements test the desired transverse completion, while the baseline no-slip and PPN statements no longer depend on that unfinished sector.

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

16. Why Dynamics Requires Extension Beyond the Static Branch

The static weak-field branch answers one question: what the capacity-deficit profile looks like once a source has settled. That is enough for ordinary galaxies, which are quasi-static. It is not enough for clusters and mergers, which ask a different question — how that profile *develops*: how it propagates when a source moves, how it lags behind a fast disturbance, how it saturates, and how it responds when matter separates into distinct dynamical phases. None of those is contained in a static Poisson law.

If the entanglement-capacity medium is physical, it must answer that second question too, with propagation and causal response built in. The time-dependent sector is therefore not an optional add-on but the natural dynamical completion of the medium that produces the static EFT, and the cluster and cosmological sectors draw on it. The next section gives its minimal causal form; Section 17.5 then uses it for clusters and mergers.

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. The static capacity-strain field is not being replaced; the telegrapher equation describes how that same field propagates and settles when sources change. “Relaxation” in this section means genuine time-dependent settling, distinct from the static strain of the weak-field branch.

Causality requires

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

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

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

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

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

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

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

17.5 Cluster Source Projection and the Diffuse–Decoupled Channel Split

Clusters are where the simple galactic radial-acceleration relation stops being enough, and they fail it in two distinct ways. Relaxed clusters show a hook-shaped residual — near unity in the stellar-dominated center, peaking at intermediate acceleration, converging back in the deep outskirts. Merging clusters show lensing peaks that stay with the collisionless galaxies while the dominant baryonic mass, the X-ray gas, is displaced. A viable account must produce both without granting galaxies hidden baryonic mass and without altering the galaxy law itself.

The mechanism proposed here is that the long-range entropic-excess channel does not couple equally to every baryonic phase. A diffuse, phase-averaged medium couples through the transverse projection ϵ inherited from the conditional galactic normalization; a dynamically decoupled collisionless component recovers the full projection; and a virialized coherent bath can be lifted above it. The projection value is fixed once that transverse branch is adopted. The phase-selection rule, lift profile, and decisive resolved-map test remain open.

The residual is not a single number but an acceleration-dependent profile: lensing and kinematic analyses of relaxed clusters [4, 15, 16] find the ratio of observed to galaxy-RAR-predicted acceleration near unity in the stellar-dominated centers, rising to a peak of roughly 3–5 at intermediate accelerations ($g_{\text{bar}} \sim 10^{-11}$ – $10^{-10} \text{ m s}^{-2}$), and then apparently converging back toward the galaxy relation at the lowest probed accelerations, subject to gas-extrapolation caveats in the outskirts [16]. The older integrated value of 1.5–2 from hydrostatic analyses at higher accelerations is the high- g_{bar} edge of this hook-shaped profile, not its amplitude. Merging clusters sharpen the problem. In systems of the Bullet type the dominant baryonic component—the intracluster plasma—is ram-pressure slowed and displaced from the collisionless galaxies, while the lensing peaks remain with the outgoing collisionless components rather than the X-ray gas. A viable cluster sector must explain both the relaxed normalization and the merger gas/lensing separation without assigning galaxies additional hidden baryonic mass and without altering the galactic mass anchor.

Transport lag and extra galaxy mass do not supply the required offset. In the canonical $\tau_0^{-1} = H_0$ branch, disturbances propagate at c and cross a megaparsec-scale configuration in roughly 3 Myr, about two and a half orders below a gigayear merger timescale. The field therefore tracks the moving source; an underdamped configuration retains memory of the pre-merger centroid rather than the outgoing collisionless component. Extra galaxy mass is excluded because the galactic acceleration scale and gas-dominated-dwarf RAR fix the deficit per unit baryonic mass. The proposed mechanism is instead a phase-dependent source projection for the long-range entropic-excess channel.

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

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

where $(2\pi)^2$ is the angular Haar volume of one canonical two-oscillator action cell and $g_{\text{share,eff}}$ is the admissibility-sharing content conditionally loaded into it. The same loading and horizon-coupling construction gives

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

This introduces no cluster-specific coefficient. It imports the reduction factor of the galactic branch and reads it as a source projection: a source restricted to the transverse static sector couples at strength ϵ relative to a source accessing the full horizon projection.

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

The unsuppressed projection is accessed by matter that is *not* part of the phase-averaged continuum: a collisionless overdensity that is spatially separated from, and dynamically decoupled from, a surrounding diffuse medium. The criterion is relational, not intrinsic compactness. Compactness alone would misclassify: stars in ordinary galaxies, isolated ellipticals, and globular clusters are compact and bound yet must remain on the galaxy RAR, and they do, because none is a collisionless node decoupled from a distinct diffuse continuum. A galaxy in a cluster is different only because it is embedded in, and decoupled from, the intracluster medium. The suppression is the property of participating in the continuum; matter that has decoupled from the continuum escapes it. So decoupled collisionless matter sits at the baseline weight, and incoherent diffuse gas is suppressed to ϵ , with $\epsilon = g_{\text{share,eff}}/4\pi^2$.

A *virialized* bath is the third state. Once the diffuse atmosphere relaxes into a coherent, extended phase it may open a collective response above the decoupled baseline. The proposed ceiling $W_{\text{bath}} = 1/\epsilon \simeq 5.32$ is the reciprocal capacity of the adopted transverse cell. It is fixed within that branch but inherits its microscopic uncertainty. The lift profile between $W_{\text{bath}} = 1$ and $1/\epsilon$ is evaluated against cluster data below.

The effective entropic-channel source χ_{ent} is therefore regime-dependent. In a relaxed cluster the gas is a virialized bath,

$$\chi_{\text{rel}} = \rho_{\text{dec}} + W_{\text{bath}} \rho_{\text{bath}}$$

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

$$\chi_{\text{merge}} = \rho_{\text{dec}} + \epsilon \rho_{\text{shock}} + W_{\text{bath}} \rho_{\text{vir}}$$

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

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

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

with $f_{b,\text{cos}}$ fixed by Planck values [40] and $M_{\text{hot,vir}}$, M_{500} read from X-ray/SZ and total-mass estimates. Once the transverse branch is adopted, ϵ is shared across all systems, the baryon fraction is fixed cosmologically, and the per-system quantities are measured. The undetermined object is the coherence-growth profile $W_{\text{bath}}(B_{\text{bath}})$. In merging systems $M_{\text{hot,vir}}$ refers to the pre-merger virialized atmosphere.

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

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

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

$$\mathcal{R}_{\text{rel}} = 1 + (W_{\text{bath}} - 1) f_{\text{cont}}$$

and the minimal linear candidate for the lift,

$$W_{\text{bath}} = 1 + (1 - \epsilon) B_{\text{bath}}, \quad 1 - \epsilon = 1 - \frac{g_{\text{share,eff}}}{4\pi^2} \simeq 0.812,$$

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

This linear form has the qualitatively correct mass trend: both B_{bath} and f_{cont} *increase* with halo mass—hot-gas fractions rise toward clusters [14, 13] and the atmosphere becomes more fully virialized—so their product rises monotonically from groups to massive clusters, reproducing the observed direction with no fitted parameter. Its amplitude, however, is excluded. For CLASH-scale clusters the gate gives $B_{\text{bath}} \simeq 0.83$, $f_{\text{cont}} \simeq 0.9$, hence $\mathcal{R}_{\text{rel}} \simeq 1.6$; but evaluating the source weight required to reproduce the published CLASH relation [4] across its data-supported acceleration window gives $\mathcal{R} \simeq 3.7$ at $g_{\text{bar}} = 10^{-10} \text{ m s}^{-2}$ rising to $\mathcal{R} \simeq 7.8$ at $10^{-11} \text{ m s}^{-2}$. The linear lift $(1 - \epsilon) B_{\text{bath}} f_{\text{cont}}$ is short of the observed peak residual by a factor of roughly 2.5–4, and no escape through missing baryons (a multiple of the X-ray gas mass would be required), hydrostatic bias (the masses are lensing-based), or sample heterogeneity is available at that magnitude.

The linear candidate is excluded on amplitude. Because $B_{\text{bath}} \leq 1$ and $f_{\text{cont}} \leq 1$, the linear candidate bounds the relaxed residual by $\mathcal{R}_{\text{rel}} < 1 + (1 - \epsilon) \simeq 1.81$. The measured peak residual of relaxed clusters is 3–5 [4, 15], well above this bound, so the linear lift is excluded. The exclusion is specific to the lift function: the channel-split ontology itself makes a structural prediction about the residual’s shape that the data support, taken up next.

The hook morphology. Independently of the lift amplitude, the channel split makes a radius-resolved prediction about the *shape* of the cluster residual. In cluster centers the baryons are dominated by the BCG stellar component, which is decoupled and carries weight 1, so the local residual starts near unity. Moving outward, the virialized gas continuum comes to dominate and the residual rises toward the bath-weighted value. At the lowest accelerations the deep branch compresses any bounded source weight W toward $\sqrt{W}$ in acceleration terms, pulling the residual back down. The predicted profile is therefore a *hook*: near unity in the stellar-dominated center, peaking where the bath dominates at intermediate acceleration, and converging back toward the galaxy relation in the deep outskirts. This is precisely the morphology class reported by current measurements [15, 16]—a shape that neither a total-baryon modified-gravity law (which predicts no relaxed-cluster excess at all) nor a constant offset produces.

The framework therefore gets the *kind* of residual right and, with the linear candidate, its *amplitude* wrong: the predicted peak is $\simeq 1.5$ for CLASH-like parameters against the observed 3–5. The constraint on the lift function is thus two-ended: W_{bath} must grow faster with bath

development than the linear candidate, reaching peak residuals of 3–5 at the developed-bath cluster scale, while remaining small enough at the group scale that X-ray-faint groups stay near the galaxy relation.

At the saturation weight $W_{\text{bath}} = 1/\epsilon$ the predicted mass residual for developed-bath parameters is $\mathcal{R}_{\text{sat}} = f_{\text{dec}} + f_{\text{cont}}/\epsilon \simeq 4.7\text{--}4.9$, inside the required window and at the upper edge of the measured peak band, and the radius-resolved profile then peaks at $\simeq 3.9$ in acceleration terms for CLASH-like parameters against the observed 3–5. The same weight applied uniformly at the group scale overshoots by a factor of ~ 2 , so saturation must be gated by coherence development: massive relaxed clusters sit at or near the bound, while X-ray-faint groups remain far below it, with the group-scale data requiring $W_{\text{bath}} \lesssim 1.4$ there.

One caveat accompanies the saturated profile: the predicted central residual depends on the stellar/gas decomposition and on excluding multiphase cool-core gas from the coherent bath. The deep-outskirt question—power-law continuation [4] versus convergence toward the galaxy relation [16]—is adjudicated directly below.

The ceiling against cluster data. Inverting the X-COP hydrostatic measurements [61] gives twenty-four source-weight tests. For each point, S solves

$$\frac{Sg_{\text{bar}}}{1 - \exp[-\sqrt{Sg_{\text{bar}}/a_0}]} = g_{\text{obs}}.$$

Every point respects $S \leq 1/\epsilon$; the maximum is $S = 3.96$, and the relaxed systems span $S \simeq 1.8\text{--}2.7$ at R_{500} and $1.4\text{--}2.2$ at R_{200} .

The same inversion adjudicates the deep end: the power-law continuation requires $S \simeq 6\text{--}8$ at $g_{\text{bar}} \simeq 1\text{--}2 \times 10^{-11} \text{ m s}^{-2}$, precisely the accelerations of the $R_{500}\text{--}R_{200}$ points, which sit at $S \simeq 1.5\text{--}2.7$; within this sample the deep end converges rather than continuing, and the strongest published challenge to the bound is not borne out.

The methodological caveat is that the continuation was fit to lensing-based masses of higher-redshift systems while the inversion here uses local hydrostatic masses; breaching the bound at R_{500} would require the non-thermal-corrected masses to be low by a factor of $\simeq 2.3$, well beyond any claimed hydrostatic bias. The measured radial run of the source weight— $S \simeq 3.7\text{--}4.9$ in the hook-peak window, where the saturation band is reached, declining to $\simeq 2.2$ at R_{500} and $\simeq 1.5$ at R_{200} —is the quantitative target the coherence-growth profile must reproduce.

Direct and fluctuation-based turbulence measurements find low non-thermal support in relaxed systems at all probed radii [62, 63], so the decoherence agent gating the lift cannot be the cluster-to-cluster turbulence level: it must grow with radius even in fully relaxed atmospheres. The coherence-growth profile between the fixed endpoints, so constrained, is the object the channel-selection theorem must deliver, with the group end requiring $W_{\text{bath}} \lesssim 1.4$.

The decoupled-fraction form is excluded by the mass trend. A second candidate class lets the residual ride on the decoupled stellar fraction, $\mathcal{R} = 1 + 4.32 B_{\text{bath}} f_{\text{dec}}$ with $f_{\text{dec}} = f_{\star}/(f_{\star} + f_{\text{gas}})$. This class is excluded by the mass trend rather than the amplitude. Because f_{dec} falls with mass while B_{bath} rises, their product peaks at the group/poor-cluster scale and declines toward massive clusters, predicting the largest anomalies at intermediate mass; the observed residual instead rises monotonically from groups to massive clusters, where it is most firmly established. A decoupled-fraction form passes only at the intermediate scale where its spurious hump crosses the observed band. The residual therefore rides on the continuum, not the decoupled fraction—the trend direction settles which component carries it, even while the lift amplitude remains underived. The broader lesson is that the relaxed-cluster amplitude and

the Bullet morphology are distinct observables and must not be collapsed into a single scalar law.

Bullet-type mergers. The relaxed residual and Bullet morphology use the same branch coefficient ϵ with different source expressions because the gas occupies different states. A relaxed atmosphere enters through the bath-lift term; shocked displaced gas carries the suppressed weight while collisionless cores remain decoupled. The two regimes share one conditional coefficient, not one universal scalar law.

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

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

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

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

Falsifiers and open status. The rule makes sharp predictions beyond the relaxed normalization. (i) The linear candidate’s ceiling $\mathcal{R}_{\text{rel}} < 1.81$ lies below the measured peak residual of relaxed clusters [4, 15], which excludes that form of the lift. The shape prediction of the channel split is the hook: residual near unity in BCG-dominated centers, a single peak at intermediate acceleration where the virialized bath dominates, and convergence back toward the galaxy relation in the deep outskirts. A relaxed cluster with a residual profile that is monotonic in acceleration, or that peaks in the stellar-dominated center, would falsify the channel split itself. The capacity bound supplies the sector’s hard ceiling: $\mathcal{R}_{\text{rel}} \leq f_{\text{dec}} + f_{\text{cont}}/\epsilon \simeq 5$ for developed-bath parameters, with developed baths reaching it in the hook-peak window and all twenty-four X-COP source-weight inversions respecting it; a robust relaxed-cluster mass residual well above this bound would falsify the capacity bound, the sector’s central commitment.

(ii) At fixed mass, X-ray-bright bath-developed systems (larger B_{bath} , larger f_{cont}) should deviate more from the galaxy RAR than X-ray-faint systems; the residual turns on with the developed diffuse atmosphere, not with mass alone, so two systems of equal mass but different bath development should separate.

(iii) The decisive test is the resolved lensing map. Because the relaxed and merger regimes use different source expressions, the convergence is a *three-component* channel-weighted map—a relaxed virialized-bath component, a shocked non-equilibrium continuum at the suppressed weight, and a decoupled collisionless component,

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

with the branch value $\epsilon = g_{\text{share,eff}}/4\pi^2 \simeq 0.188$ or with ϵ floated as a test, and W_{bath} fit within $[1, 1/\epsilon]$. Recovering $\epsilon \simeq 0.19$ across relaxed and merging systems would support both the cluster source rule and the inherited transverse normalization. A best fit near 1 or 0 would falsify the cluster branch. The decisive fit must use the channel-weighted baryonic maps without free dark haloes, which otherwise absorb the gas weight.

Three open items remain at the theory level. First, the relaxed residual and Bullet morphology use one inherited coefficient with two regime-specific source expressions. The boundary between the virialized and shocked regimes is not derived.

Second, suppression of a phase-averaged continuum and collective lift of a virialized bath are independent premises. Neither follows from the current microscopic source map. Their derivation must also reproduce the measured radial decline and the proposed branch endpoints.

Third, the identification of the decoupled component with stellar/galaxy mass and of the continuum with the gas is a coarse split; intracluster light and tidally stripped stars blur it at a level the resolved map fit would expose. Abell 520, whose reported gas-coincident dark core is itself disputed, is a phase-state stress case rather than a clean test: a re-cohering or quasi-bound central component would raise its effective B_{bath} and return lensing toward the gas, while a confirmed *young* merger with a robust gas-centered, galaxy-free lensing peak—where the re-coherence escape is unavailable on time grounds—would be a serious challenge.

This sector is a structured, falsifiable proposal. The branch value of ϵ is inherited rather than re-fitted, the bath gate is measured, and current data support the trend and hook morphology while excluding the linear lift candidate. The microscopic transverse normalization, coherence-growth profile, resolved-map test, and channel-selection theorem remain open.

18. Cosmology and the Hubble-Tension Sector

The scalar capacity field has two cosmological roles, and the sector works only if they stay separate. Its homogeneous mode $\bar{S}(t)$ affects the background expansion and the sound horizon; its inhomogeneous fluctuations $s(x, t)$ still govern local weak-field gravity. The cosmological sector is the homogeneous continuation of the same medium, not an unrelated dark-energy component appended to the weak-field theory: what changes is the kinematic regime, not the ontology, as the background mode becomes dynamically relevant on horizon scales while the local branch stays encoded in the fluctuations.

The cosmological sector uses the same field split,

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

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

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

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

Because the entanglement field couples to the trace of the stress-energy tensor, the homogeneous mode is largely dormant during radiation domination but can become active near matter–radiation equality. The conditional proposal would reduce the sound horizon and shift the CMB-inferred Hubble constant upward. Its joint evolution with the committed component has not been computed.

The direction of the shift matters here, and so does 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 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

18.5 The Saturated Phase and the Cosmic Microwave Background

This section extends the conditional transverse normalization of Section 14 into the early universe. In galaxies the proposed response remains below its capacity ceiling. At recombination the same branch demands more response than that ceiling allows, motivating a saturated phase in which a conserved count of committed channels redshifts as a^{-3} . The dust theorem below is conditional on the pinned realization of that phase; the numerical abundance also inherits the transverse coefficient ϵ .

Saturation requirement at recombination. Use $y \equiv g_{\text{pert}}/a_0(z)$ for the acceleration ratio and $x \equiv \sqrt{y}$ for the thermal argument of Section 14. The linear-domain estimate $y \lesssim 2 \times 10^{-3}$ today and $y \propto \sqrt{a}$ gives $y \lesssim 6 \times 10^{-5}$ at $z \simeq 1100$, hence $x \lesssim 8 \times 10^{-3}$ and $\nu(x) \gtrsim 10^2$. This is well above the proposed ceiling $1/\epsilon = 5.32$, so the adopted response branch requires saturation if the same estimate applies to the relevant perturbation modes. The relation $a_0(z) = \epsilon cH(z)$ keeps the background ratio $cH/a_0 = 1/\epsilon$ fixed at every redshift.

The pinned reading and the dust theorem. The *pinned reading* is the statement that the saturated bath holds exactly at the ceiling and stays there. The transfer law quantizes channel transport at one bit per tick as an upper bound; saturation is the attainment of that bound, so on the pinned reading every committed channel advances by exactly one unit per tick, and while the phase remains pinned no channel can decommit, since there is no slack for it to relax into. The committed count is then conserved and dilutes only with the expanding volume. The coarse-grained field ϕ of the committed phase then carries a constraint rather than a generic kinetic term: $X = -\frac{1}{2}g^{\mu\nu}\partial_\mu\phi\partial_\nu\phi = \frac{1}{2}$, with $\partial_\mu\phi$ a unit timelike covector. The leading action consistent with the constraint is

$$S_{\text{comm}} = \int d^4x \sqrt{-g} \lambda \left(X - \frac{1}{2} \right),$$

whose variation gives $\nabla_\mu(\lambda\partial^\mu\phi) = 0$ and, on the constraint surface,

$$T_{\mu\nu} = \rho u_\mu u_\nu, \quad \rho = E_0 n, \quad p = 0, \quad c_s^2 = 0, \quad \rho \propto a^{-3},$$

with $u_\mu = \partial_\mu\phi$ and n the conserved spatial density of committed cells. This belongs to the constrained-scalar class [64]. Conditional on the pinned reading, it is pressureless dust. The construction inherits caustic formation and the need for a small-scale completion. The manuscript does not derive the transition from the unconstrained weak-field branch onto $X = \frac{1}{2}$ or the commit/decommit energy accounting. Related relativistic Milgromian dynamics provides a useful comparison [24].

Uniqueness of the carrier. The constraint coupling is the unique survivor of the dynamics classes examined (Appendix M). Relaxational response kernels are excluded twice over: a growth clock calibrated on the cluster radial decline misses cosmological development by a factor of order thirty, and the precision of the measured acoustic peaks bounds any oscillatory leakage of a lagged response below one part in five hundred, a rejection no causal filter achieves over the few oscillation periods available before recombination. Bound-type rail readings sit at the sound-speed pole and carry no perturbation; plateau approaches have $c_s^2 = -1/(2n+1) < 0$ and are gradient-unstable; the released branches have $c_s^2 = \frac{1}{2}$ and 1 and free-stream. A general no-go for relaxation-plus-cap carriers is recorded in Appendix M; the conserved committed density is precisely the additional integrating variable that the no-go requires and the constraint reading supplies.

Abundance. At full saturation the committed weight is the ceiling, giving

$$\frac{\Omega_c}{\Omega_b} = \frac{1}{\epsilon} = \frac{4\pi^2}{g_{\text{share,eff}}} = 5.321 \quad \text{against the measured } 5.364 \pm 0.065,$$

a 0.7σ agreement with the abundance counted rather than fitted [40]; equivalently $\Omega_m/\Omega_b = 1 + 1/\epsilon = 6.32$. A standard Einstein–Boltzmann computation [65] with the cold component tied to this value and the acoustic scale θ_* held fixed reproduces the quoted Planck best-fit temperature spectrum with a root-mean-square residual of 0.15% over $\ell = 2$ –2500 and 0.08% in the third-peak region. At the likelihood level, the tied model carries one fewer free parameter than Λ CDM and sits at $\Delta\chi^2 = +2.15$ against the six-parameter best fit on the compressed Planck TT,TE,EE likelihood [31]. This fixed- θ_* comparison tests the inherited dust abundance only. It does not demonstrate the separate Section 18 claim that the homogeneous mode reduces the sound horizon. A joint Boltzmann calculation must evolve the homogeneous mode, commitment transition, and tied dust component together; that calculation remains open.

The acceleration and abundance relations contain a parameter-free structural test that does not depend on the numerical value of the phase-cell coefficient. Since $a_0 = \epsilon c H_0$ and $\Omega_c/\Omega_b =$

$1/\epsilon$,

$$\boxed{a_0 \frac{\Omega_c}{\Omega_b} = cH_0}.$$

Using the RAR scale $(1.20 \pm 0.02) \times 10^{-10} \text{ m s}^{-2}$, the Planck abundance ratio 5.364 ± 0.065 , and $H_0 = 67.4 \pm 0.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$ gives

$$\frac{a_0(\Omega_c/\Omega_b)}{cH_0} = 0.983 \pm 0.022$$

[1, 40]. This two-percent agreement tests the reciprocal structure of the two sectors. It does not derive either conditional premise: the transverse loading controls a_0 , while the pinned recruitment reading controls the abundance.

That the committed density attains the ceiling exactly, rather than a development-dependent fraction of it, follows only within the recruitment assumptions of this branch. Commitment is source-driven: cells receiving no demand commit nothing, and each source's allocation is capped at $1/\epsilon$ per unit source mass. In the recombination estimate above the relevant modes demand $\nu \gtrsim 10^2$, far above the cap, while the homogeneous background remains below it. With abundant supply and irreversible commitment in the pinned regime, each source recruits to its cap and allocations are assumed to add. The committed density is then

$$\rho_c = \frac{1}{\epsilon} \rho_b$$

pointwise at a commitment epoch. The manuscript has not derived that epoch or shown that commitment completes before the modes used in the Boltzmann calculation begin their relevant acoustic evolution. The initial condition $\delta_c = \delta_b$ is therefore an explicit assumption of the current numerical check, not an output of the recruitment argument. The abundance remains conditional on the pinned reading and per-defect bookkeeping; the transition history and energy accounting are open.

Conditional conservation: release at the caustic. The committed component must behave as conserved dust while the linear bath remains pinned, yet the galactic branch (Section 14) requires collapsed systems to follow the baryonic response law with no surviving collisionless halo, so the completion is a conditional-conservation law: $\nabla_\mu J_{\text{commit}}^\mu = 0$ in the pinned regime, with decommitment upon release. The constrained phase fixes the release mechanism itself. Because the committed flow is the gradient of a single clock field, $u_\mu = \partial_\mu \phi$, it is exactly irrotational and single-valued, and a multi-stream or vortical velocity field admits no such potential; gravitational collapse therefore terminates the phase at first shell-crossing. The caustic — the known breakdown locus of the constrained-scalar class — is here the decommitment event, converting committed capacity into the local response branch. Under the synchrony definition of commitment, one counter per cell ticking with the local phase, decommitment at first stream-crossing is forced rather than chosen: a single fixed cell cannot remain synchronized with two distinct phase branches, so conversion is event-like at the first caustic, a result holding at the same conditional grade as the pinned reading itself. Re-entry is forbidden on entropic grounds: recommitment of a virialized region would require shedding vorticity and re-synchronizing with the advanced global clock, a spontaneous recoherence, so decommitment is irreversible, with the door's orientation inherited from the thermodynamic arrow, which remains an open conditional extension (Section 21, Appendix G.7). The resulting partition is the one the data require: the linear cosmological field never shell-crosses and remains committed dust, while collapsed systems have shell-crossed and retain none. The alternative, commitment tracking the instantaneous local acceleration, is excluded directly by rotation-curve data: below g_c the binned radial-acceleration residuals lie within 0.05 dex of the response law [1], against the +0.40 dex excess that locally

regrown dust at the cosmic ratio would produce. One completion remains open and is recorded, the energy bookkeeping of conversion at the caustic, together with the consequence that late nonlinear structure growth proceeds on the response branch. The sharp geometric prediction is registered as a falsifier: infalling material between turnaround and first shell-crossing is single-stream and still committed, so clusters should carry a surviving committed component on their infall streams, terminating at the splashback surface [66], with an excess gravitating rim in that shell, a sharp edge at the outermost caustic, and none inside it. The recruitment count fixes the rim's amplitude as well as its geometry: along a stream of which a fraction f_{rel} has already shell-crossed, the surviving committed density is $(1 - f_{\text{rel}}) \rho_b / \epsilon$, so the prediction specifies location, edge, and magnitude together.

Domain structure. Writing $y = g/a_0$ and $x_{\text{th}} = \sqrt{y}$, the release threshold obeys $\nu(x_{\text{th},c}) = 1/\epsilon$. Thus

$$x_{\text{th},c} = -\ln(1 - \epsilon) = 0.2082, \quad y_c = x_{\text{th},c}^2 = 0.0433,$$

and $g_c = y_c a_0 \simeq 5.3 \times 10^{-12} \text{ m s}^{-2}$ today. The earlier notation used x_c for y_c ; the separate symbols here remove that ambiguity. Whether linear cosmological modes stay on one side of this threshold throughout their evolution must be established in the joint Boltzmann calculation.

Distinction from horizon saturation. The committed cosmological phase — channels advancing at capacity, a running clock — is distinct from the terminal saturation of the strong-field sector (Section 20), where capacity is exhausted and transactions cease. Whether the two termini are stages of one process is open and carries a stated consistency burden: the gravitating energy of committed capacity must coincide with the mass already accounted at infinity, with no double counting. This is recorded as an open question shared by the strong-field and cosmological sectors.

19. Why These Sectors Belong

The most directly constrained chain — microstructure to static weak-field observables — is now in hand, and the sectors on either side of it form one program rather than a list of add-ons. Each asks what the same finite-capacity substrate does once it leaves the static weak-field regime.

Transport asks how the capacity-strain field propagates after a source moves. Cosmology treats the homogeneous mode, the saturated phase treats the proposed committed carrier, and the strong-field branch treats the zero-capacity boundary. Many-Pasts supplies the record-conditioned history ontology. The separate refresh theorem selects the memoryless dressing used in scale setting, while Section 22 asks for its microscopic dynamics.

These sectors share the static weak-field ontology but not its evidential status. The Einstein/capacity action equivalence, ordinary source map, and finite marked-transfer scale are the controlled baseline inside their displayed actions; the separate stiffness loop, galactic, cosmological, and boundary additions carry their stated conditional or open grades.

20. Strong-Field Action: Spherical Closure and Its Boundary

The bounded variable remains the natural strong-field order parameter,

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

Serial composition and weak-field matching select $N^2 = q$ in a static exterior. That relation is meaningful after the asymptotic Killing time has fixed the normalization of the lapse; it is not, by itself, a covariant field equation in a general foliation.

Why the former multiplier action is not retained. An exploratory ADM term $\sqrt{h} \lambda (N^2 - q)$ makes the problem look variational, but if q has no independent bulk dynamics or matter coupling its variation gives $\lambda = 0$. The remaining constraint merely renames the lapse, which is gauge dependent, and supplies neither the capacity Poisson equation nor a new covariant relation. Giving q its own kinetic term would instead add a scalar degree of freedom and reopen the fifth-force, stability, and double-counting problems. The multiplier construction is therefore a diagnostic control case, not the parent action.

Exact spherical reduction. Spherical symmetry supplies the invariant completion that the generic lapse constraint lacks. Write

$$ds^2 = h_{ab}(x) dx^a dx^b + R^2(x) d\Omega^2, \quad a, b \in \{0, 1\}, \quad x^0 = ct.$$

After integrating the Einstein–Hilbert plus GHY action over the two-sphere and removing a two-dimensional total derivative, the bulk action is

$$I_{\text{sph}} = \frac{c^3}{4G} \int d^2x \sqrt{-h} \left[R^2 {}^{(2)}R + 2h^{ab} \partial_a R \partial_b R + 2 \right] + I_{\text{matter}}^{(2)} + I_{\partial}^{(2)}.$$

Its vacuum equations imply conservation of the Misner–Sharp mass

$$M_{\text{MS}} = \frac{c^2 R}{2G} \left(1 - h^{ab} \partial_a R \partial_b R \right), \quad \nabla_a M_{\text{MS}} = 0.$$

The capacity-adapted invariant is therefore

$$q_{\text{geo}} \equiv h^{ab} \partial_a R \partial_b R = 1 - \frac{2GM_{\text{MS}}}{c^2 R}.$$

Thus in the preferred metric-only construction, spherical q is a composite geometric scalar and its vacuum profile is a first integral of the metric equations. It is not an auxiliary field and it is not an additional propagating mode.

For a static asymptotically flat vacuum, $M_{\text{MS}} = M$ and Birkhoff’s theorem gives

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

so

$$q_{\text{geo}} = N^2 = 1 - \frac{2GM}{c^2 r}.$$

This is the nonperturbative action-level realization of the bounded capacity relation in the strongest sector where an invariant local definition is currently available.

The $q = 0$ surface. The equation $q_{\text{geo}} = 0$ locates a marginal sphere. A well-posed exterior variational problem is obtained first on a stretched timelike boundary $q = \epsilon > 0$ with the usual GHY term and fixed induced data, and then by taking the null limit with the corresponding null and joint terms. The Einstein action determines that universal gravitational boundary book-keeping. It does not determine a new substrate boundary Hamiltonian, microscopic reflectivity, or the rule that excises $q < 0$.

Accordingly, the restriction

$$\mathcal{M}_q = \{q_{\text{geo}} > 0\}$$

is a physical postulate about the domain of the capacity EFT, not a consequence of varying the Einstein action. The exterior solution and horizon location are closed within spherical metric reduction; the claim that the classical interior is absent, and the dynamics experienced at the saturation surface, remain conditional on a microscopic boundary theory. Appendix F and Appendix N state this boundary between result and interpretation explicitly.

What survives and what remains open. The Schwarzschild exterior, its standard Hawking temperature, and ordinary exterior perturbation equations follow from the metric parent. The channel identity yielding a coefficient 1/4 remains a normalization identity until the channel-to-area map is derived. Rotating and charged GR exteriors remain valid solutions of the baseline parent action, but a covariant capacity scalar that identifies their saturation surface has not yet been constructed. The universal strong-field result is therefore spherical and exterior; boundary microphysics, dynamical saturation, nonspherical capacity geometry, and any controlled departure from GR remain open.

21. Many-Pasts: The History-Space Ontology

Many-Pasts makes a conservative operational claim and a distinct ontological claim. Laboratory records obey standard quantum mechanics. The ontology assigns a conditional measure to the decoherent past histories compatible with the one realized present.

The operational weight. For a decoherent family, the joint history weight and its record-conditioned form are

$$p(h, P) = \mathcal{D}(h, h), \quad p(h | P) = \frac{\mathcal{D}(h, h)}{\sum_{h' \in \mathcal{H}_P} \mathcal{D}(h', h')}.$$

Equivalently, one may define $D(h, P) = -\ln p(h, P)$ on the support of the decoherent joint measure, so that $p(h | P) \propto e^{-D(h, P)}$. The exponential is a reparameterization of a normalized quantum probability, not an additional classical measure placed on interfering paths. Projective measurements and general quantum instruments therefore retain their Born probabilities. Appendix G gives the construction and the no-signaling proof.

Branch realization without many worlds. The weight also answers what it means for one outcome to be realized. There is no forward branching into co-real worlds and no collapse event. A definite present is a present with definite macroscopic records, and the histories the weight supports are exactly those compatible with those records; alternative outcomes correspond to alternative present records, not to coexisting branches. Probability is the measure this weighting assigns over the admissible pasts of the one realized present.

Familiar quantum examples. In a double-slit experiment, the alternatives through the two slits remain inside one coarse class operator until a durable which-path record decoheres them. Their cross terms are therefore retained before the record and suppressed after it. In an EPR or Bell experiment, the present is a joint record of the pair and the detectors. The joint measure gives the usual nonclassical correlations, while the local marginals remain independent of the remote setting. Measurement creates a stable record and thereby identifies the decoherent family on which conditional probabilities can be used. These are standard quantum calculations with a record-conditioned history-space reading.

The arrow of time. The same framework proposes to orient time through conditional typicality. The maximum-caliber replacement process does not solve this problem. At stationarity it obeys detailed balance identically, $p(b)K_*(b, b') = p(b')K_*(b', b) = p(b)p(b')$, and is time-reversal symmetric as a stochastic process. The reversible dilation also has an inverse. Histories with a low-entropy past and increasing future entropy dominate only after a past-boundary condition and a substrate mixing or large-deviation theorem suppress Boltzmann-fluctuation histories. This is not an added law of laboratory probability; it is the open typicality result stated precisely in Appendix G.7.

Where the weight is used. Many-Pasts supplies the history-space realization in which the electron dressing operates. It does not itself prove that successive passes are independent. Appendix H shows that the already-imposed faithful full-support condition is equivalent to maximum path entropy and uniquely selects the independent replacement process; the lightest-defect criterion selects the same temporal independence together with independent channel layers. The record-conditioned viewpoint is also used in the discussion of the macroscopic arrow of time and in the conditional cosmological extensions, where it remains a structural proposal rather than a new laboratory law.

Local renewal and reversible history export. Complete local renewal has a precise quantum form. If the new cell must contain no information about the old cell even when the old cell is entangled with a reference, the one-tick channel is uniquely

$$\mathcal{E}_*(\rho) = \rho_* \text{Tr } \rho.$$

This local channel need not destroy information globally. A reversible dilation transfers the old cell into an environment register while a fresh admissible register becomes the new present. Many-Pasts interprets the exported correlations as history degrees of freedom. The identification is structural: Stinespring dilation guarantees an environment, while Postulate III supplies its history-space reading. It does not derive the decoherence functional or show that an indefinitely long history can be stored in finite microscopic resources.

Status. The operational measure is mathematically complete because it imports the standard decoherence functional and quantum instruments; within those kinematics the Born form is additionally the unique consistent record measure (Section 3.3, Appendix G.4). A substrate derivation of that functional is still missing. The foundational faithful full-support condition closes the classical replacement kernel on history space, and reference decoupling closes the form of the quantum replacement channel. The decorated tetrahedral transfer vertex prepares the diagonal fresh state, exports the previous cell, and realizes the finite charged marked event. Long-time history storage, a substrate derivation of quantum kinematics, and the thermodynamic arrow retain their open or conditional grades.

22. Microstructure Hamiltonian and Underlying Dynamics

The UV closure chain now has an explicit finite action on the scale-setting side and a candidate condensate realization on the geometry side. Appendix H derives the replacement process, its quantum channel, the state-weighted determinant transfer, and a decorated native-cell vertex. The vertex prepares the diagonal fresh state, exports the previous present, couples a hard-core charged mark to the closure response, and fixes the graph that dresses the electron and heavier shells. Its geometric GFT embedding and condensate stability remain separate tasks.

On the geometry side, the realization is a GFT/condensate picture [49, 50] in which spacetime emerges from a condensate of discrete tetrahedral building blocks, while what is macroscopically read as matter appears as fermionic defects of that same substrate. In Madelung form,

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

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

On the defect side, the closure ensemble fixes the stationary marginal p_{η_*} and maximum path entropy selects

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

This fixes the dimensionless history process but cannot produce seconds. The electron, already the theory's single dimensionful anchor and its lightest charged one-bit defect, supplies the clock through the positive transfer spectrum below.

Requiring the output cell to decouple from every reference system forces the replacement channel $\mathcal{E}_*(\rho) = \rho_* \text{Tr } \rho$. The decorated vertex chooses the diagonal maximum-entropy completion

$$\rho_* = \sum_b p_*(b) |b\rangle\langle b|, \quad p_*(b) = Z^{-1} e^{-\eta_* K^2(b)},$$

and prepares it from the amplitude $A_*(b) = \sqrt{p_*(b)}$. A reversible register permutation moves the old cell into history and the fresh amplitude into the present. For the complete 1680-state tetrahedral register this is one native circuit layer. A direct circuit made only from disjoint face comparisons still requires the three perfect matchings of the four-face graph; the decorated action uses the complete tetrahedron as its local gate.

On the factorized seven-channel history space let

$$|v^{(7)}\rangle = \bigotimes_{m=-3}^3 |\sqrt{p_{\eta_*}}\rangle_m, \quad P^{(7)} = |v^{(7)}\rangle\langle v^{(7)}|.$$

On the commutative cell algebra, likelihood multiplication by $p_*(b)$ has the state-weighted determinant

$$\Delta_{\tau_p}(\mathbf{R}) = \exp\left(\sum_b p_b \ln p_b\right) = e^{-g_{\text{share}, \text{eff}}}.$$

The seven-channel determinant is $r = e^{-7g_{\text{share}, \text{eff}}}$. It is the almost-sure geometric transfer rate of the record-conditioned renewal history; the annealed equality probability would instead involve the collision entropy. Compressing the one-bit charged loop to its determinant line gives the positive scalar survival transfer

$$T_{\text{surv}} = 1 - r.$$

The Euclidean transfer Hamiltonian

$$H_{\text{surv}} = -\frac{\hbar}{\tau_*} \ln T_{\text{surv}}$$

therefore has the exact raw gap

$$E_{\text{raw}} = -\frac{\hbar}{\tau_*} \ln(1 - r).$$

The closure amplitude has an exact Gaussian linearization. Projecting the same amplitude through the two directed scalar returns gives $u = 8\eta_*/49$, and its canonical unitary dilation supplies the factor $(1 - u)^{21/2}$. The nine present/history closure polarizations give the marked weight ζ_* of Section 13.4. The electron graph contributes $Z_e = (1 + \zeta_*)(1 + 7\zeta_*^2)$, so the dressed transverse identification is

$$E_e = \frac{3}{2} Z_e E_{\text{raw}}.$$

Identifying this lowest charged transfer gap with the electron rest energy fixes the update time,

$$\tau_* = -\frac{3}{2} Z_e \frac{\hbar}{m_e c^2} \ln(1 - r).$$

Analytic continuation supplies the phase frequency $E_e/\hbar$. The identification of the same spectral tick with the causal cell scale gives $L_* = c\tau_*$.

The causal length of that spectral update is

$$L_* = c\tau_* = -\frac{3}{2}Z_e\lambda_e \ln(1 - e^{-7g_{\text{share,eff}}}).$$

The baseline rare-event correction $-\ln(1 - r) = r[1 + O(r)]$ differs from r only at order 10^{-23} . The finite marked response changes the scale by $Z_e - 1 = 0.00530828$, which is the physically relevant correction resolved by the microscopic vertex.

Retained temporal information raises the recurrence mass by e^{I_t} , so electron lightness independently selects $I_t = 0$. Fermionic exclusion caps the occupied channel count at seven and subadditivity gives $\Delta_k \geq 0$, hence the lightest resolved one-bit defect has $k = 7$ and $\Delta_7 = 0$ at fixed per-channel marginal. The unit mixing gap, charged survival gap, and coherent GFT Hessian remain distinct objects. The decorated transfer action fixes the first two and the marked edge Hessian. A geometric GFT calculation must still determine the condensate spectrum and prove that no additional light source-coupled mode appears.

The finite scale-setting dynamics is therefore specified end to end: renewal fixes the fresh marginal, the determinant fixes the baseline recurrence, the decorated vertex fixes the marked event and routing, positivity fixes the transfer gap, and the electron fixes the cadence. The complete edge-Hessian and graph-enumeration audits are finite and reproducible. Long-time history capacity, stable condensate embedding, and the first-principles inhomogeneous metric influence functional remain open.

The cell-order calculations show that the closure weight, refresh, history tilts, and conserved field do not supply vacuum curvature stiffness. They do couple geometry to persistent closure failure, and Section 23 measures that local response. The condensate branch still owes a microscopic kinetic term, while the simulations import their extended vacuum geometry from the CDT host. Postulate III conditions histories within that specified state and does not select it.

Part VI. Closure Status, Falsifiability, and Research Program

23. The Substrate on a Dynamical Lattice

The coefficient chain of Part II is exact at the order of a single cell: the ensemble is enumerated, the closure point is solved, and every number that feeds the weak-field sector is evaluated in closed form. The postulates, however, also commit the theory to statements no single-cell calculation can reach, because they concern many cells in interaction: that the vacuum ensemble survives when its cells live on a fluctuating geometry; that a persistent closure failure (the theory's matter) strains the capacity state around it; and that the strain is carried outward far enough to become the $1/r$ deficit of Section 12. This section reports a numerical program built to test those statements directly, by placing the cell ensemble of Part II, unmodified, on a dynamical simplicial geometry and running the coupled system. The section carries the design logic and the results; ensemble definitions, control protocols, and the measured tables are collected in Appendix J.

23.1 The host geometry and the cell identification

The host is a causal-dynamical-triangulations ensemble: four-dimensional triangulations of $S^1 \times S^3$ weighted by the Regge action [67]. CDT is used because it independently sustains an extended

four-dimensional de Sitter-like phase [68, 69]. That phase is an external host datum for the substrate test, not a state selected by Many-Pasts.

The identification is immediate, and it respects a separation the theory requires. Each spatial tetrahedron of a slice hosts one cell’s boundary data: its four triangles carry the labels $m = -3, \dots, 3$ of the $j_{\text{eff}} = 3$ sector of Section 5, each slice triangle is shared by exactly two cells, and the slice gluing supplies adjacency. Nothing else of the lattice enters. The theory’s cell is its label configuration, a dimensionless unit of capacity; no size is ever assigned to it, and the granularity stays in the capacity, as Section 1.2 requires. The lattice simplices are regulator scaffolding, as they are throughout lattice gravity. The identification is nevertheless geometric in the sense that matters: the joint admissibility statistics of the labels depend on how the slice is glued, so the closure weighting can in principle tell one geometry from another, and that channel is the entire coupling between the theory and the host.

Two boundaries of the design are stated at the outset. First, the host supplies what the program needs from a substrate, a dynamical, curvature-carrying simplicial geometry with the theory’s cell structure, and nothing more is asked of it. The program does not assume the CDT ensemble is the theory’s own vacuum, and it withholds any claim about four-dimensional emergence, which belongs to the host’s phase structure and is still being mapped at the couplings used here (Appendix J.1). Second, the weighting coupled to the host is the paper’s own: the admissibility energy of Sections 5–6 at the closed point, with $\eta = \eta_*$ and injectivity enforced. It is not a spin-foam vertex amplitude, and the distinction drawn in Section 27.8 was kept operational: a spin-foam amplitude was coupled first, as a neighboring-theory control, and the audit it forced supplied the control methodology used everywhere below (Appendix J.3).

23.2 The coupled ensemble and its controls

The coupled system is the joint Gibbs measure

$$\pi(g, m) \propto \exp\left(-S_{\text{Regge}}(g) - \varepsilon(N_{41} - \bar{N})^2 - \beta \sum_{\text{cells } c} [E_c(m) - \mu]\right),$$

$$E_c(m) = \eta_* K^2(m_c) + \lambda n_{\text{coll}}(m_c),$$

where the sum runs over the slice cells, K^2 is the closure invariant of Appendix B, n_{coll} counts label collisions (injectivity is imposed as a penalty whose hard limit is the constraint, with the residual collision fraction reported so the softness stays visible), ε pins the slice volume, and μ is the per-cell label free energy, computed by thermodynamic integration so that the closure term cannot masquerade as a shift of the bare cosmological coupling. The labels carry a uniform base measure, so at $\beta = 0$ the label entropy cancels exactly and the bare host is recovered identically. The theory point is $\beta = 1$ at η_* : once the chain of Part II is accepted there is no coupling dial left free, and the intermediate β values serve only as a diagnostic ramp.

Every run opens by recomputing the single-cell chain on its own tables and refuses to proceed unless $g_{\text{share,eff}} = 7.4198$ and $\langle K^2 \rangle_{\eta_*} = 3/(2\eta_*) = 50.223$ are reproduced, so the object coupled to the lattice is verifiably the object counted in Part II. Results are read under a discipline fixed before the runs: matched total volume across arms; a global audit that every quoted ensemble’s slices are closed 3-manifolds correctly threaded by the foliation; and a placebo arm in which the 210 label-orbit energies are shuffled (same value pool, same permutation symmetry, closure structure destroyed), so that only a real-versus-placebo difference is ever attributed to closure. Every claim in this section is a matched comparison at fixed regulator (real against placebo, coupled against bare, commitment level against commitment level), so regulator-scale artifacts cancel by construction. The intended reading of each possible outcome was recorded before the results existed (Appendix J.3).

23.3 What the closure sector cannot supply: vacuum stiffness

Three cell-order calculations bound the closure sector’s contribution to vacuum stiffness. The static closure weight induces only about one percent of the bare geometric coupling. The selected memoryless kernel, modeled as cell-by-cell redraws, has a closed stationary state with vacuum admissibility 0.536 and no detectable curvature action. A separate closure-class history tilt, solved by a Doob transform on a ring, also remains short-ranged and well below the required coupling. These are calculations of specified effective models; they do not derive the refresh kernel from Postulate III.

Two further calculations extend the exclusion beyond the label sector, to the conserved capacity field that Section 23.7 introduces as the carrier of the long-range sector. Even a field the labels cannot see might rank geometries on its own, because integrating out a conserved Gaussian field induces a purely geometric action: $\frac{1}{2} \ln \det' L(g)$ per channel, the spanning-tree entropy of the slice graph by the matrix-tree theorem. Computed on degree-matched proxies for smooth extended and for crumpled geometry, and on real engine slices, that entropy is nearly universal at fixed coordination: the per-cell differential is a few times 10^{-4} , and with all seven channels the phase-tipping force is of order 10^{-3} against the bare $k_0 \simeq 2.2$. The theory’s own non-Gaussianity closes the loophole tighter still. The finite budget of Postulate I gives the field a saturation mass $m^2 = 2/g_{\text{share,eff}} \simeq 0.27$, and the massive determinant suppresses precisely the soft modes that carried the residual sensitivity, shrinking the differential by a further factor of three at the budget mass and toward zero beyond. What the field retains is a local renormalization of the host’s couplings, of order 0.08 per cell in the seven-channel count: enough to relocate the host’s phase boundaries, and unable to create the phase.

Within the tested models, static weight, refresh, history tilts, and conserved fields do not supply vacuum stiffness. The refresh does couple geometry to the density of closure failure, which the defect experiment measures. The same pass verifies the locality fracture: injectivity-preserving unit shifts split the 1680 states into 48 components of 35 states. This excludes that local realization of the full-entropy kernel; it does not derive the nonlocal physical operator.

23.4 The externally hosted vacuum and the role of conditioning

The exclusions show that the specified substrate does not dynamically select its vacuum geometry. Two questions called emergence of spacetime must therefore be separated. The first is kinematic encoding: the closure invariant K^2 measures the failure of the four oriented faces to close, adjacency carries proximity, and the marked-transfer dictionary of Section 13 converts entanglement increments into meters. The cells are dimensionless capacity units rather than sites of a preferred spatial grid. The second question is dynamical selection among crumpled, branched, and smooth extended geometries. Section 23.3 finds that none of the tested substrate mechanisms performs that ranking.

Many-Pasts cannot fill that dynamical gap. Appendix G normalizes the probabilities of alternative present records before conditioning, so $p(h | P)$ describes histories compatible with an already specified present and does not select which P occurs. In the simulations the smooth extended vacuum is supplied by the chosen Regge/CDT host phase as an external background datum. Many-Pasts may condition the compatible decoherent histories within that datum, but it neither ranks candidate geometries nor explains why the extended present is realized. The coefficients of Part II are computed on that host state, and the medium dresses it rather than generating it. A substrate derivation of the prior state or of a probability measure over alternative vacuum geometries remains open.

Regge/CDT is a useful external host because it supplies tetrahedral cells, dynamical curvature, no rigid preferred spatial grid, and a demonstrated extended four-dimensional phase

[68, 69]. This choice is operational, but it is not a vacuum derivation by the substrate. The host supplies the geometry on which the label theory is tested; the substrate then produces quantitative dressing of its couplings. Integrating out the labels gives $S_{\text{eff}} = S_{\text{host}} - \log Z_{\text{label}}$. Appendix J.6 reports the predicted extensive volume shift and its measured scaling family, including the independently predicted placebo line. The computed curvature-sector shifts remain pre-registered phase-boundary tests. The hierarchy is explicit: cell-level non-closure reweights the glued tetrahedra, and integrating out the labels converts that reweighting into corrections to the host’s volume- and curvature-sector couplings. This is a quantitative interface with an externally supplied vacuum geometry, not endogenous vacuum selection.

23.5 Compatibility: the weighting on dynamical geometry

At $N_{41} = 20,000$ (about 47,000 pentachora in all, hosting 10,000 cells across eighty slices), arms at $\beta = 0, 0.3, 1$, and the $\beta = 1$ placebo were run to matched volume with clean foliation audits (Appendix J.5). The label sector orders exactly as the closure weight demands: the collision fraction falls from 0.648 at $\beta = 0$ (against the uniform-measure prediction $1 - 840/2401 = 0.650$, confirming that the control arm samples the base measure) to 0.082 at the theory point, while the placebo, carrying the same energy values without the closure structure, stays at 0.73. The geometry does not move: the Hausdorff dimension and the slice-volume profile are statistically indistinguishable across all four arms at single-seed resolution. That outcome was the quantitative prediction of Section 23.3, since an induced coupling of order one percent of the bare stiffness cannot steer these observables at this precision. The reading is compatibility: switching the theory’s weighting on at full strength does not destabilize the host geometry, and the placebo separation in the label sector shows the coupling was live while the geometry held still.

23.6 The defect experiment: geometry responds to the theory’s mass

The theory’s matter is persistent closure failure: committed capacity, maintained against the refresh (Sections 3.2, 22). The experiment inserts it by hand. In a thermalized coupled ensemble at the theory point, one hundred well-separated cells were pinned into maximal closure failure (all four faces at $m = 0$: six collisions, $K^2 = 48$) and held for three thousand sweeps. A placebo arm pinned one hundred cells into the best-closed injective configuration ($m = \{0, 1, 2, 3\}$, the minimum $K^2 = 40.67$) under identical anchoring and identical protection from the geometry moves, so the real-minus-placebo comparison subtracts everything the pinning machinery does; freshly sampled unpinned cells far from any pin supply the vacuum reference inside both arms. Around each pin, four observables were accumulated shell by shell in the slice adjacency, for roughly 10^5 defect-environment samples per arm: the mean closure energy and the collision fraction of the shell, which read the capacity state, and the edge coordination and cell count of the shell, which read the local geometry.

The environment responds, and the placebo comparison ties the response to the failure content. At the first shell the capacity state separates: mean closure energy 1.6782 ± 0.0010 around the failure pins against 1.7291 ± 0.0011 around the closed pins and 1.7306 ± 0.0011 in the vacuum, with the collision fraction displaced the same way (0.0690 against 0.0794 and 0.0797). By the second shell both label observables have returned to their vacuum values: the strain profile dies within one lattice step, the range the correlation length of Section 23.3 requires. The local geometry separates as well, and reaches further. At matched anchoring, the failure pins carry higher shell coordination than the closed pins (5.2997 ± 0.0007 against 5.2513 ± 0.0008 at the first shell) and systematically more cells per shell, a local volume excess still present at the third shell: 11.296 ± 0.006 cells against the placebo’s 10.905 ± 0.007 and the vacuum’s 10.909 ± 0.012 .

The licensed statement is local and specific: persistent closure failure sources a measurable strain in the surrounding capacity state and deforms the local geometry around it, and it does

so through its failure content, since the placebo carries the anchoring without the failure. This is the microscopic seed of the source-to-strain mechanism the weak-field sector runs on, demonstrated dynamically. The volumes license nothing long-range: no $1/r$, no G , no curvature in the continuum sense; the measured reach is three lattice steps. The capacity response has since replicated in an independent seed, the pooled first-shell separation reproducing and tightening (1.678 around the failure pins against 1.729 placebo and 1.732 vacuum); the geometric response replicates in direction with a factor-two across-seed scatter in magnitude, so the geometric numbers are quoted per seed, with further seeds running. Quoted errors are statistical and not yet thinned for autocorrelation, and the host’s phase map now shows a located candidate region and boundary awaiting its confirmation run (Appendix J.1); a detected response survives those caveats, which govern the precision of the numbers rather than the existence of the effect.

23.7 Reaching Newtonian range: the conservation requirement

The strain of Section 23.6 dies within a lattice step, as the sub-cell correlation length says it must. The weak-field sector needs it carried to macroscopic range, and the lattice makes sharp what that requires. Suppose free capacity were only a per-cell budget, re-equilibrating locally by maximum entropy around each commitment. The linear response then solves $(2z - L)\delta\mu = -m/\chi$ on the slice adjacency ($z = 4$ faces per cell, L the slice Laplacian), and that operator has no small-momentum pole: on a real slice the response falls eight orders of magnitude within thirteen steps. A budget that merely re-balances screens gravity at the substrate scale. Suppose instead that free capacity obeys a genuine continuity law (transported through shared faces, exchanged only with commitment, globally conserved) and that a defect draws the steady maintenance flux Section 22 assigns it, linear in its commitment. The steady state then solves the massless graph-Poisson equation, massless because conservation forbids any local restoring term, and its Green function on a three-dimensional slice falls as $1/r$: the deficit field is $\delta f \propto m/r$, the M -linearity inherited from the maintenance postulate and the $1/r$ from the conservation law. On the lattice, Newton’s form of Section 12 is the signature of conserved capacity with maintenance sinks, and nothing weaker produces it (Appendix J.8).

The continuum sector never had to say what is conserved; the lattice does. Two carriers are available: a capacity stock transported between cells, which would require a conserved quantity the label sector demonstrably does not possess, or the refresh bandwidth itself, each cell’s one redraw per τ_* , exactly conserved tick by tick, with a defect’s maintenance consuming the refresh coherence of its neighborhood, and the deficit field the steady allocation of that bandwidth. The paper’s own commitments lean to the second: mass is already a recurrence rate (Section 22, Appendix H.6), $a_0 \propto cH(z)$ is a rate, and maintenance linear in commitment is the natural law for a bandwidth share.

The measurement instrument is built and gated in advance (Appendix J.8). A conserved field on the slice cells, initialized at the vacuum anchor 7.4198 and transported by antisymmetric face fluxes with global conservation asserted at machine precision every sweep, is coupled to pinned defects through the microphysics alone: cells absorb capacity in proportion to their instantaneous closure failure, so the sink strength a defect exerts is an output of the label dynamics, with nothing in the transport law specifying it. Absorption must read only the failure a cell carries above the vacuum’s own standing level, and the reading has to respect what the theory means by commitment: failure that *persists*. Getting that wrong is measurable, and it has now been measured twice. With absorption on total failure, the vacuum’s own turnover absorbs and screens the field at the predicted four-to-seven-step length (the first full-volume run). With absorption on the *instantaneous* excess, vacuum fluctuations are rectified — the vacuum’s collision count is zero on ninety-two percent of cells, so subtracting its small mean removes almost none of its absorption — and the instrument carries an intrinsic screening length of about seven steps by construction (the second run; Appendix J.8). The corrected

instrument reads the persistent failure average, the theory’s own definition of commitment, the vacuum’s flicker turnover being the anchor budget rather than commitment; that suppresses the vacuum draw toward zero and pushes the intrinsic length beyond any measurable radius. The ladder of pinned commitments delivers the two readings the theory does not preordain, on both instrument versions: with the measured commitment-independent dressing subtracted, the emergent maintenance flux is proportional to commitment ($Q - Q(0) = 0.166\,m$ on the corrected instrument, $0.187\,m$ on the first; the maintenance postulate, measured), and the first-shell deficit per unit flux is level-independent at nine percent across the ladder (4.66 ± 0.42 on the corrected instrument, whose absolute scale is the quotable one), the junction statement. The far-field gates were fixed in advance — the deficit must follow the graph-Coulomb form, and a fitted screening mass must be consistent with zero, since any leakage in the conservation law would appear as Yukawa screening — and the corrected instrument, on thicker slices, passes them. The fitted screening slope is $+0.115 \pm 0.072$, consistent with zero over the full measurable range, bounding the screening length above 3.9 steps with the bound set by the slice radius rather than the physics; the deficit is still alive at the fifth shell at a third of its first-shell strength, where the rectified instrument had it extinct; and the profile exponent, -0.65 ± 0.14 against the ideal -1 , is Green-function-like with the expected closed-slice flattening, sharpenable at larger volume and nothing like screening. The diagnosis certified itself: the level-zero rung’s spurious flux collapsed twelvefold ($0.1075 \rightarrow 0.0086$) exactly as the pre-registered side prediction required — the vacuum stopped absorbing once persistence removed the flicker double-count — so this pass is trustworthy in the specific way the earlier fail was not readable. The run is single-seed and its errors are autocorrelation-naïve; within those caveats, no pre-registered gate on the mechanism side remains open.

23.8 What the program establishes

The results separate four jobs. The closure sector supplies the defect counting. Persistent failure produces a measured local capacity and geometric response. A conserved capacity current supplies the long-range graph-Poisson form. None of the tested substrate mechanisms selects the vacuum geometry; CDT supplies that host state externally, and the substrate measurably dresses its volume coupling. Section 24 records the corresponding exclusion, compatibility, and single-seed dynamical grades.

24. 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.

The leading word in each status uses a fixed vocabulary. *Closed* means derived within the stated postulates and ensemble. *Fixed* means no phenomenological freedom remains once the named branch is adopted. *Conditional* means the result follows if a named reading or completion holds. *Frontier* or *open* means the piece is structured but incomplete. *Empirical support* denotes comparison with data rather than derivation, and *audit task* marks an independent check still required. A conditional premise does not demote every theorem downstream of it. Rows therefore separate exact identities and within-model lemmas from the microscopic or empirical status of their premises.

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

Quantity / Claim	Sector	Status	Type of Support	Where Established
η_*	admissibility closure	Closed	stationary normalized closure evidence $\ln Z + \frac{3}{2} \ln \eta$ maximized on the exact K^2 spectrum; the $3/2$ is the determinant weight of the three closure-defect components	Part II, App. B
$g_{\text{share,eff}}$	UV entropy	Closed	exact weighted evaluation	Part II, App. B
Memoryless refresh kernel	information / scale theorem	Closed under the foundational faithful full-support condition	applying that same condition to histories gives maximum path entropy, not a new premise: $H(B_{t+1} B_t) = g_{\text{share,eff}} - I(B_t; B_{t+1})$ is maximal iff $I = 0$, uniquely giving $K(b, b') = p_{\eta_*}(b')$	Part I, App. D, G–H
Decorated charged marked-transfer vertex	UV dynamics	Closed as a finite transfer action; geometric GFT embedding open	Gaussian closure amplitude fixes $\sqrt{\eta_*}$; two directed singlet returns give $u = 8\eta_*/49$; canonical dilation gives $\sqrt{1-u}$; present/history response strands give nine orthonormal marked states; complete Hessian and 56,800-state graph audits pass	§13, §22, App. H, K
Substrate length L_* (marked-transfer map)	scale setting	Fixed within the decorated transfer action and electron anchor	state-weighted determinant gives $r = e^{-7g_{\text{share,eff}}}$; the marked vertex gives $Z_e = (1 + \zeta_*)(1 + 7\zeta_*^2)$; positivity and the electron anchor give $L_* = -(3/2)Z_e\lambda_e \ln(1-r)$	Part I, App. D, H
$J_{\text{bare}}, J_{\text{eff}}^{\text{tree}}$	UV edge kernel	Closed	tetrahedral isotropy identity	Part II, App. C
$\Sigma_{\text{ret}} = 65/9$	finite-loop UV	Conditional minimal return-sector completion	seven diagonal return channels plus one projected singlet are specified; an explicit microscopic return operator must derive their relative weights and exclude further motifs	Part II, App. C
$J_{\text{eff}}^{(\text{ren})}$	finite-loop UV	Conditional on the minimal return operator	algebraic Dyson resummation once Σ_{ret} is specified	Part II, App. C
γ	continuum stiffness	Conditional loop-dressed value	Euclidean normalization is exact for the one-sublattice edge convention; the numerical value inherits the minimal return operator	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 action / bridge law	EFT / gravity	Closed for the ordinary longitudinal branch	Einstein–Hilbert + GHY reduced in Newtonian gauge gives I_{Newton} ; $\delta S = -2S_\infty \Phi/c^2$ gives $I_{\text{cap}}^{\text{static}} = Z_S I_{\text{Newton}}$ with no extra scalar	Part III, App. D, N
G_*	gravitational scale	Parameter-free output of the decorated scale branch; historical postdiction	$G_* = (9/4)(\hbar c/m_e^2)Z_e^2 \ln^2(1 - e^{-7g_{\text{share,eff}}}) = 6.6742890772 \times 10^{-11}$, or -0.073σ relative to CODATA; the residual was known before the vertex was constructed	Part I, App. D, H, L
Matched G	weak-field gravity	Algebraic identity in the chosen normalization; not an independent determination	substituting the closed source map and S_∞^{cell} into $G = c^2\kappa/(8\pi\gamma S_\infty)$ gives identically $G = c^3L_*^2/\hbar = G_*$	Part III, App. C–D
Electron anchor	mass / length sector	Fixed empirical elementary anchor	one-bit fermionic defect supplies λ_e for length setting and $m_e/\ln 2$ for mass–entropy map	Part III, App. D
Seven-sector additivity and lightest branch	UV scale theorem	Closed entropy theorem; exact conditional mass minimum	subadditivity gives $H_k = kg_{\text{share,eff}} - \Delta_k$ with $\Delta_k \geq 0$; at fixed admissibility marginals, the adopted recurrence mass $m_k \propto e^{\Delta_k - kg_{\text{share,eff}}}$ is minimized uniquely at the fermionic ceiling $k = 7$ and $\Delta_7 = 0$	App. D, H

Quantity / Claim	Sector	Status	Type of Support	Where Established
a_0	galactic EFT	Exact conditional consequence; loading and horizon coupling open	the canonical transverse doublet gives two action-angle pairs and hence angular Haar volume $(2\pi)^2$ for one quantum action cell; $a_0 = [g_{\text{share,eff}}/(2\pi)^2]cH_0$ follows after loading one sharing entropy into that cell and coupling it reversibly to the horizon bath	Part III, App. N
Rank-one transverse response	galactic EFT	Closed within the one-invariant leading EFT	$C_\times^2 = A_L A_T$, $\delta R_*/\delta X = -C_\times/A_T$, and $C_\times/A_T = \omega_L/\omega_T$ follow exactly; tetrahedral symmetry gives $Z_L = Z_T$ at leading derivative order	Part III, App. N
RAR law	galactic EFT	Exact consequence of a conditional leading completion	thermal matching reduces to $\hbar\omega_T = k_B T_H$ in the clamped-oscillator reading; A_T must remain environment-independent, and the influence functional must use the clamped response and yield $n_B(x)$ without an extra spectral factor	Part III, App. N
Baseline no slip / lensing	weak-field metric	Closed for the ordinary longitudinal branch	varying the unreduced Einstein scalar constraint gives $\nabla^2(\Phi - \Psi) = 0$ and hence $\Phi = \Psi$ under asymptotic flatness	Part III, App. D, N
Transverse-sector lensing	galactic metric response	Open — principal action task	RAR fixes the radial temporal response only; $\Gamma_\perp[g_+, g_-]$ must derive $\Delta\Psi_\perp$, the lensing kernel, and the covariant Ward identity	Part III, App. N
PPN leading values	weak-field metric	Closed for baseline; transverse decoupling open	Einstein parent gives $\gamma_{\text{PPN}} = \beta_{\text{PPN}} = 1$ and no extra longitudinal mode; the transverse influence kernel must be shown negligible and free of preferred-frame effects at high acceleration	Part III, App. F, N
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
Saturated-phase committed dust	cosmology	Conditional on the pinned transfer-law reading	saturation required within the stated perturbation estimate; constrained-scalar dust theorem ($p = 0$, $c_s^2 = 0$, $\rho \propto a^{-3}$); transition onto the constraint surface remains open	Part V, §18.5, App. M
Committed-capacity abundance $\Omega_c/\Omega_b = 1/\epsilon$	cosmology	Conditional on the transverse normalization, pinned reading, and per-defect bookkeeping	over-demanded sources recruit to the per-source cap and allocations add, giving $\rho_c = \rho_b/\epsilon$ within that branch; the commitment epoch and joint Boltzmann evolution remain open	Part V, §18.5
$a_0(\Omega_c/\Omega_b) = cH_0$	cross-sector structure	Exact conditional identity independent of the cell value	multiplication cancels ϵ ; current central values give $a_0(\Omega_c/\Omega_b)/(cH_0) = 0.983 \pm 0.022$; the acceleration loading and committed-abundance premises remain separate	Part III, Part V, §18.5
Post-fixation tests of ϵ	empirical support	Supported across heterogeneous tests; no combined significance assigned	a_0 is +2.6% from the local RAR scale; Ω_c/Ω_b agrees at 0.7σ ; all twenty-four X-COP inversions obey the capacity ceiling; two redshift samples report evolution of a_0 in the predicted direction; these comparisons share theory premises and are not statistically independent	§14, §17.5, §18.5, §25.3
Caustic release / conditional conservation	cosmology	Conditional under the synchrony reading	committed dust is proposed to decommit at first shell-crossing of the irrotational clock flow; splashback-rim falsifier; conversion bookkeeping open	Part V, §18.5

Quantity / Claim	Sector	Status	Type of Support	Where Established
Diffuse source projection $\epsilon = g_{\text{share,eff}}/4\pi^2$	cluster sector	Conditional, with no additional coefficient	inherits the compact two-phase transverse normalization used for a_0/a_H ; the cluster extension adds no further dial	Part IV, §17.5
Bath gate B_{bath}	cluster sector	Operationally closed; microscopic origin open	$M_{\text{hot,vir}}/(f_{b,\text{cos}}M_{500})$ from X-ray/SZ; no per-system fitted knob	Part IV, §17.5
Cluster residual $\mathcal{R}_{\text{rel}} = 1 + (W_{\text{bath}} - 1)f_{\text{cont}}$	cluster sector	Hook morphology and trend direction supported; linear lift candidate excluded; lift function open	linear candidate's ceiling 1.81 lies below measured peak residual 3–5 [4, 15, 16]; capacity bound fixes saturation at $1/\epsilon$, giving $\mathcal{R}_{\text{sat}} \simeq 4.7\text{--}4.9$ and a radius-resolved peak $\simeq 3.9$ inside the measured band; the bound holds against all twenty-four X-COP source-weight inversions, which also exclude the deep power-law continuation within that sample; coherence-growth profile open	Part IV, §17.5
Relaxed vs. merger source expressions	cluster sector	Conditional — two regimes, one coefficient	residual rides on virialized bath; Bullet on shocked gas + decoupled clumps; regime boundary not yet derived	Part IV, §17.5
Channel-selection rule (suppression + collective lift)	cluster sector	Conditional — suppression open; linear lift excluded	transverse-suppression premise (participation $\Rightarrow$ weight ϵ) not yet derived from the source map; linear form of the lift excluded by the measured peak residual; saturation endpoint fixed at $1/\epsilon$ by the channel capacity bound; coherence-growth profile between the endpoints not yet derived	Part IV, §17.5
Bullet gas/lensing inversion	cluster sector	Structurally supported; consistent but untested (ϵ not yet measured)	ϵ brings gas/galaxy to near-parity, compactness completes the inversion; existing flexible reconstructions non-discriminating because gas weight is halo-degenerate	Part IV, §17.5
Resolved cluster lensing-map test	cluster sector	Open — principal empirical task	three-component $\kappa \propto (1 + (1 - \epsilon)B_{\text{bath}})\Sigma_{\text{bath}} + \epsilon\Sigma_{\text{shock}} + \Sigma_{\text{dec}}$; predicted best-fit $\epsilon \simeq 0.19$, floated on baryonic maps with free dark haloes disallowed; not yet performed	Part IV, §17.5
Bounded capacity q , $N^2 = q$	strong field	Fixed static constitutive rule; covariant generalization open	multiplicative composition and weak-field matching fix the static lapse map; a generic lapse is foliation dependent	Part V, App. F, N
Former ADM multiplier action	strong field	Excluded as a parent completion	with no independent q dynamics, variation gives $\lambda = 0$ and only renames the foliation-dependent lapse; adding dynamics introduces an extra mode	Part V, App. F, N
Spherical parent reduction	strong field	Closed in the metric-only branch	exact two-dimensional reduction gives $q_{\text{geo}} = (\nabla R)^2 = 1 - 2GM_{\text{MS}}/(c^2 R)$ as a composite first integral and returns Schwarzschild in static vacuum	Part V, App. F, N
Capacity-exhaustion horizon $q = 0$	strong field	Geometric zero closed in spherical symmetry; domain termination conditional	$q_{\text{geo}} = 0$ is the marginal sphere; identifying it with substrate saturation and excluding $q_{\text{geo}} < 0$ requires the bounded-domain postulate and boundary theory	Part V, App. F, N
Hawking temperature and exterior ringdown	strong field	GR-matching branch closed; capacity boundary condition open	Schwarzschild exterior fixes the wave operator and Euclidean temperature; $\mathcal{R} = 0$ follows if future-horizon regularity is retained, while $\Gamma_{\partial q}$ must determine the physical substrate reflectivity	App. F
Bekenstein–Hawking area-law bridge	strong field	Exact normalization identity within the cell-normalized horizon convention; channel-to-area count open	per-channel cut entropy $\ln 2$ from fermionic face exclusion + channel density from $G_{\perp} = (2/3)G_{\text{tet}}(0)$; product $n_{\text{hor}}S_{\infty}^{\text{cell}} = 1/4$ as exact identity	App. C, F

Quantity / Claim	Sector	Status	Type of Support	Where Established
Horizon formation and boundary microphysics	strong field	Geometric marginal-surface formation closed; saturation dynamics open	GR collapse can form $q_{\text{geo}} = 0$; a bounded causal q_{cap} evolution, its equality to q_{geo} , the boundary action, and relaxation spectrum remain to be derived	Part V, App. F
Rotating / charged stationary exteriors	strong field	Baseline metric solutions closed; capacity map open	Einstein / Einstein–Maxwell gives Kerr / Reissner–Nordström / Kerr–Newman, but no nonspherical covariant capacity scalar or domain rule has yet been derived	App. F
Many-Pasts Born compatibility	quantum foundations	Operationally complete; measure unique under stated quantum-record hypotheses; kinematics imported	Gleason fixes every normalized noncontextual additive measure on a sufficiently rich record-projector lattice in dimension > 2 ; generalized records transfer that measure to medium-decoherent pure-state histories; Hilbert-space kinematics, state, record completeness, and the substrate origin of the decoherence functional remain premises	Part V, App. G
No-signaling in operational branch	quantum foundations	Exact within the operational branch	summing a remote quantum instrument gives a trace-preserving map, leaving the local marginal independent of the remote setting	Part V, App. G
Arrow-of-time account	quantum foundations	Open conditional extension	requires a Substrate Past Hypothesis, elapsed time below relaxation and recurrence, and a mixing/large-deviation theorem showing ordinary entropy-increasing histories dominate the conditioned ensemble	Part V, App. G
Microstructure Hamiltonian	UV realization	Finite marked-transfer action closed; stable geometric embedding open	the decorated vertex prepares the diagonal fresh amplitude, exports history, produces the 21-block antisymmetric edge Hessian, and fixes the three routing traces; a geometric GFT must still realize the oriented decoration, select a stable condensate, and derive the full source-coupled fluctuation spectrum and durable history capacity	Part V, App. H
Charged-lepton spectrum	particle-sector extension	Fixed within the decorated marked-transfer and shell action; no fitted lepton coefficient	closure-spectrum collapse gives three shells; the baseline shell algebra is dressed by $Z_\mu = 1 + \zeta_*$ and $Z_{\tau,2} = 1 + (2/7)\zeta_*$, giving $m_\mu/m_e = 206.768280237$ and $m_\tau/m_e = 3477.343310$, both within 1σ	§13, App. H–I
Gauge-redundancy extension	gauge sector	Coherent extension	baseline-redundancy construction with Maxwell/Yang–Mills form	App. I
Vacuum stiffness from the substrate	lattice dynamics	Excluded, five branches	induced curvature coupling $c_0 = +0.019$ against bare $k_0 \simeq 2.2$; refresh stationary state geometry-blind to exponential accuracy; history-tilt ceiling twenty times below requirement; free conserved field induces near-universal spanning-tree entropy; budget saturation mass suppresses the residual soft modes	§23.3, App. J.4
Vacuum geometry in the lattice realization	lattice dynamics	Open; externally hosted	tested substrate mechanisms do not select an extended phase; the Regge/CDT host supplies the vacuum geometry, while Many-Pasts conditions histories only after a present record is specified	§23.4, App. G, J.4
Ensemble fracture under local dynamics	lattice dynamics	Verified	explicit construction: 48 components ($4! \times 2$) of 35 states each under injectivity-preserving unit shifts, conserved (ordering, parity) charge	§23.3, App. D.4, H.6, J.4

Quantity / Claim	Sector	Status	Type of Support	Where Established
Substrate contribution to the host volume term	lattice dynamics	Measured as a scaling family, demonstration volume	closure free energy per cell, computed from η_* alone, displaces the equilibrium volume by the predicted amount: single point -62.2 predicted, -62.8 measured; family of eight coupling-pin pairs spanning -10 to -158 at ratios 0.96 – 1.07 , concave β -curve and $1/\epsilon$ collapse confirmed, placebo on its own predicted line at ratio 1.00 ; curvature-sector shifts pre-registered	§23.4, App. J.6
Admissibility weighting on dynamical geometry	lattice dynamics	Compatibility demonstrated	matched-volume coupled ensemble at the theory point: label sector orders (collision fraction $0.648 \rightarrow 0.082$; placebo 0.73), geometric observables statistically unchanged, as the one-percent induced coupling predicts	§23.5, App. J.5
Local response of geometry to persistent closure failure	lattice dynamics	Demonstrated dynamically; capacity channel replicated across seeds	pinned maximal-failure cells against matched closed-pin placebo: first-shell capacity strain (replicated, pooled separation tighter), coordination and shell-volume deformation through three shells (directionally replicated, factor-two across-seed magnitude scatter, quoted per seed), separating from placebo well outside quoted (unthinned) errors	§23.6, App. J.7
Long-range propagation of the strain field	lattice dynamics	Conservation law derived; measured, all pre-registered gates passed (single-seed)	budget-only response screens at sub-cell range; a conserved free-capacity current with maintenance sinks gives the massless graph-Poisson form, $\delta f \propto m/r$; measured at production volume: conservation exact, maintenance flux proportional to commitment ($Q - Q(0) = 0.166 m$, dressing subtracted), junction constant level-independent at nine percent (4.66 ± 0.42), screening slope $+0.115 \pm 0.072$ consistent with zero over the measurable range, profile Green-function-like (-0.65 ± 0.14); the first instrument's screened verdict was diagnosed as its vacuum-rectifier artifact and the diagnosis confirmed by the pre-registered twelvefold collapse of the level-zero flux	§23.7, App. J.8
Numerical robustness checks	validation layer	Supportive audit layer	cross-sector consistency tests and executable reproduction block	App. K, O
EFT consistency checklist	field-theory audit	Supportive audit layer	ordinary longitudinal branch has no extra scalar degree of freedom; static source sign and quadratic form checked; telegrapher stability belongs to the conditional transport completion	App. D

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

Cosmological row. The saturated phase enters the ledger as follows: the dust form is a theorem of the constrained-scalar class conditional on the pinned reading; the abundance $1/\epsilon = 5.321$ stands against the measured 5.364 ± 0.065 and inherits the transverse normalization together with the recruitment and per-defect assumptions; the energy accounting of the committed budget is open; and the domain assignment remains a conditional cosmological completion.

25. Falsifiability and Observational Tests

25.1 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, 35], the baryonic Tully–Fisher scaling, and the weak-lensing response, where stacked galaxy–galaxy lensing probes the relation two decades below the rotation-curve regime [36]. Baseline no slip and GR PPN values now follow from the Einstein parent reduction. Persistent gravitational slip associated with the low-acceleration excess would instead falsify the proposed transverse single-mode completion. Solar-System and laboratory bounds [3, 44, 46] require the transverse kernel to decouple at high acceleration and exclude any independently coupled static scalar force.

Wide binaries supply an independent solar-neighborhood discriminator for the proposed local transverse rule. If the gap is controlled by the total local baryonic field, a binary embedded in the Galactic field $g_{\text{ext}} \simeq 1.4\text{--}1.9 \times 10^{-10} \text{ m s}^{-2}$ approaches the boost $1 + n_B(\sqrt{g_{\text{ext}}/a_0}) \simeq 1.4\text{--}1.5$ once its internal field falls below the external term. This is a conditional prediction of the local influence-kernel completion, not a consequence of the ordinary transport equation. Current Gaia analyses divide between a low-acceleration boost and Newtonian consistency [28, 29]; a settled Newtonian result would falsify this local transverse branch.

25.2 Dynamical falsifiers

The dynamical extension is more vulnerable, and its failure modes are correspondingly sharper. Time-dependent halo lag, cluster-scale acceleration relations, merger offsets, or relaxation signatures that cannot be reconciled with the telegrapher relation $D/\tau_0 = c^2$ would indicate that the causal completion has the wrong propagation structure even if the static branch survives. Cluster data accordingly split into two tests of distinct parts of the framework: *source-projection* tests, which ask whether the relaxed-cluster residual profile follows the predicted hook morphology—near unity in BCG-dominated centers, peaked where the virialized bath dominates, converging in the deep outskirts [4, 15, 16]—and whether resolved merger maps prefer the projection coefficient $\epsilon \simeq 0.19$ (Section 17.5); the linear lift candidate is excluded on amplitude (Section 17.5); and *transport* tests, which ask whether the causal propagation law $D/\tau_0 = c^2$ correctly evolves those source weights through a merger. Systems such as the Bullet Cluster [2] probe both, and should not be treated merely as larger versions of the static galaxy problem.

25.3 Cosmological falsifiers

Cosmology presents a different kind of test. The question there is 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. The empirical target is set by the current measurement spread: early-universe inferences near 67.4 [40], distance-ladder determinations ranging from $\simeq 70$ [42] to 73 [41], and a tension whose proposed resolutions are reviewed in Di Valentino et al. [43].

The epoch dependence $a_0(z) = cH(z)g_{\text{share,eff}}/(4\pi^2)$ is the framework’s most exposed kinematic prediction. For a Planck-like background, $H(2.3)/H_0 \simeq 3.47$. A disk with $g_{\text{bar}} \simeq 2 \times 10^{-10} \text{ m s}^{-2}$ is then predicted to have $g_{\text{obs}}/g_{\text{bar}} \simeq 2.02$, against 1.39 for an epoch-independent scale. Observed outer rotation curves at these redshifts indicate strong baryon dominance [30, 34], although pressure-support and stacking systematics leave the confrontation unsettled. Controlled observations showing no stronger boost than matched $z = 0$ systems would falsify $a_0 \propto H(z)$.

Direct measurements at intermediate redshift have since entered. A MUSE sample of 79

star-forming galaxies at $0.33 < z < 1.44$ finds the radial-acceleration relation persisting with a characteristic acceleration that rises systematically with redshift [32], and a resolved low-redshift HI sample independently reports tentative evolution in the same direction [33]. The sign is the framework's: an epoch-independent a_0 predicts no evolution at all. The magnitude is not yet a test, since absolute normalizations at these redshifts carry mass-to-light and pressure-support systematics of the same order as the effect; the clean confrontation is the survey-internal ratio $a_0(z_{\text{high}})/a_0(z_{\text{low}})$ against $E(z_{\text{high}})/E(z_{\text{low}})$, a comparison with no free parameter that the prediction must pass.

Saturated-phase falsifiers. Three tests bind the saturated phase: a full Einstein–Boltzmann implementation must jointly evolve the homogeneous mode, commitment transition, and constrained fluid with abundance fixed at $1/\epsilon$; the recruitment derivation must determine its epoch and survive scrutiny of the pinned reading and per-defect bookkeeping; and the threshold $g_c \simeq 5.3 \times 10^{-12} \text{ m s}^{-2}$ must produce a domain assignment compatible with linear observables. The release law adds a fourth test: a surviving committed component on cluster infall streams terminating at the splashback surface; its confirmed absence, or a halo-like component persisting inside collapsed systems, would falsify the caustic-release mechanism.

25.4 Correlated-constant falsifiers

The inferred L_* , induced G_* , and matched weak-field G are not independent legs: the matched route carries the same electron length calibration. Their agreement is a normalization audit rather than a separate falsifier. The marked action adds two independent tests because the same ζ_* enters the muon and tau through different graph polynomials. A future independent determination of L_* would provide a fourth test of the shared vertex.

25.5 Many-Pasts status

Many-Pasts imports ordinary quantum instruments, so it predicts no laboratory departure from the Born rule or no-signaling. Its discriminating burdens are theoretical: the substrate must reproduce the decoherence functional rather than assume it, and a conditional-typicality calculation must suppress high-entropy-past and Boltzmann-fluctuation histories. Failure of either derivation would remove the proposed ontology or arrow without constituting a new experimental violation of standard quantum mechanics.

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

The failure modes are not all equally severe, and it helps to order them by how much each would bring down.

Kills the core ontology. A demonstrated failure of mass–entropy equivalence, an internal incoherence in the Many-Pasts weighting, or evidence that geometry cannot be read as the long-wavelength form of an entanglement-capacity substrate would remove the foundations on which everything else rests.

Kills the static capacity-response closure or its transverse completion. A weak-field UV coefficient chain that cannot be reconciled with an independently validated microscopic derivation would break the ordinary static closure. An RAR transition shape that departs from the proposed bosonic law, persistent low-acceleration slip, unacceptable residual PPN effects, or violation of the metric Ward identity would falsify the transverse completion while leaving the ordinary Einstein/capacity equivalence intact.

Kills an extension only. If the cosmological trace-coupled homogeneous mode cannot survive a full Boltzmann likelihood confrontation, the cosmology sector fails while the static weak-field branch stands. If the saturated phase fails any of its commitments — the dark-to-baryonic abundance departing from the capacity ceiling, linear-regime observables departing from the inherited form below the release threshold, or no committed component surviving on cluster infall streams out to the splashback surface — the committed-dust reading fails in the same contained way. If q_{cap} does not track q_{geo} during collapse, or if the derived boundary response conflicts with black-hole observations, the bounded-domain interpretation fails while the spherical Einstein reduction remains. None of these touches the weak-field core.

Requires modification, not death. A fuller graph calculation may revise the separate conditional stiffness self-energy, and strong-field boundary spectroscopy may change without altering the Einstein exterior. A failure of the decorated charged vertex would be more serious because the same ζ_* enters the electron scale and both heavier-lepton ratios. The positive-spectrum clock conversion survives only if another microscopic charged action replaces it.

One scale-setting commitment cuts across these tiers and deserves its own line. The high-precision G_* rests on seven-fold additivity, the record-conditioned determinant transfer, and the decorated marked vertex. A microscopic Hessian with a collective edge mode, persistent endpoint polarization, or off-diagonal charged propagation would fail the finite action and displace the correlated G and lepton results. If the electron identification or mass-entropy ontology failed, the scale-setting derivation of Appendix H would not go through.

27. Comparison with Other Approaches

Because the framework proposes to reattribute the galactic excess and cosmological abundance to two phases of one medium, the sections below set out how its logic differs from nearby alternatives. Those reattributions retain the conditional grades assigned in Sections 14 and 18.5.

27.1 Relative to Λ CDM

The contrast with Λ CDM begins at the level of ontology. Here visible matter is interpreted as localized defects in a vacuum-capacity medium. The ordinary longitudinal response is a reduced representation of Einstein gravity; the proposed extra galactic response is carried by a transverse substrate sector. The UV entropy enters both branches, but only the ordinary branch closes without the transverse matching conditions. Lensing and cosmology remain extensions.

27.2 Relative to MOND-like interpolation programs

MOND-like programs [17, 48, 19] usually begin from an acceleration law or interpolation function. Here the same law is reconstructed from a rank-one capacity EFT and equilibrium bosonic free energy. Its horizon and GFT matching conditions are explicit, so the comparison turns on whether those conditions can be derived and whether their metric kernel passes lensing and precision tests.

27.3 Relative to Verlinde-style emergent gravity

Verlinde-style emergent-gravity programs share the broad intuition that gravity may be entropic [20, 47, 22, 23], but they are usually formulated through thermodynamic reasoning or horizon-inspired force laws. The present framework specifies tetrahedral counting, admissibility closure, edge coupling, return dressing, and Euclidean normalization before reaching the continuum EFT. Its correctness is an empirical question.

27.4 Relative to TeVeS and other multi-field modified gravities

Multi-field relativistic MOND completions such as TeVeS [18] introduce additional scalar and vector fields alongside the metric to obtain relativistic lensing and cosmology. Here the ordinary longitudinal response adds no field beyond the Einstein metric. The transverse thermal response is nevertheless an additional effective sector, and its retarded metric kernel must derive its own slip and lensing behavior.

27.5 Relative to AeST

The closest modern comparator is the AeST theory of Skordis and Złośnik [24], which combines the metric with a dynamical timelike vector and a scalar field and can reproduce MOND-scale galaxy phenomenology while fitting the CMB power spectrum. The present construction differs in its finite counting input and is less developed cosmologically: the joint Boltzmann treatment remains open here.

27.6 Relative to scalar-tensor gravity

The reduced capacity functional of Section 10 can be mistaken for the scalar part of a Brans–Dicke-type theory [45]. The action reconstruction shows otherwise: in the ordinary static branch it is the Einstein constraint action after a field redefinition and carries no independent scalar. A scalar–tensor or disformal theory is retained only as a control case for sectors that might genuinely require an additional mode; if adopted, ordinary fifth-force, slip, and PPN constraints would apply.

27.7 Relative to quantum-mechanical interpretations

Because Many-Pasts occupies the role of an interpretation of quantum mechanics (Sections 3.3, 21), it should be placed against the standard options. It posits one realized present, adds no collapse term, and introduces no hidden sharp values. Its operational probability theory is the decoherent-histories formalism: unresolved alternatives retain amplitudes, decoherent record histories receive diagonal probabilities, and conditioning on the present occurs only after the alternative records are normalized. The new content is the ontology assigned to that conditional measure and its proposed substrate realization. Born statistics and no-signaling are inherited from the quantum instruments; a substrate derivation of the decoherence functional remains open.

27.8 Relative to CDT, spin foams, and group field theory

The construction uses discrete tetrahedral boundary data, with a fermionic base label $j_0 = 3/2$ paired into the seven-state effective boundary representation $j_{\text{eff}} = 3$; this places it near the quantum-tetrahedron vocabulary of simplicial spin networks [51, 52, 53]. Its selection principle differs from spin-foam and GFT programs: it solves a finite boundary-counting problem and routes admissibility closure to capacity entropy rather than using a vertex amplitude as the primary history weight. Dorau and Much [21] independently relate vacuum-relative information to Killing-weighted matter-energy flux and, after assuming an area relation, to Einstein curvature. Their result supports the continuum source ontology but does not validate the tetrahedral ultraviolet construction. Section 23 supplies a separate working interface: the label ensemble is coupled to a CDT host [68, 69]. The host supplies the dynamical vacuum geometry; the paper tests how its own admissibility weighting dresses that geometry. The fixed foliation and other debated features of CDT therefore belong to the external scaffold, while the substrate’s proposed granularity remains in the capacity variables.

28. Conclusion

One physical picture runs through the paper. The vacuum is a finite medium of entanglement capacity; matter is that capacity tied up in stable, localized defects; a particle's mass measures the entanglement its defect commits; and gravity is the capacity strain the surrounding medium carries once that commitment is made. General relativity is not assumed underneath the picture but recovered as its low-energy geometry, and the extra gravity seen around galaxies, usually read as dark matter, is the long-range reach of the same medium rather than a separate substance.

The central result is a tightly specified static weak-field construction with its remaining assumptions exposed. The tetrahedral ensemble fixes the sharing entropy; faithful renewal and the decorated marked-transfer vertex fix the electron-anchored substrate length; the edge kernel and Green-matched source map fix the ordinary response; and the separate minimal loop operator supplies the stated conditional stiffness correction. The weak-field bridge then rewrites the Einstein constraint sector in the capacity variable.

Newtonian gravity is the point-source limit of the Einstein/capacity action equivalence, and baseline lensing and PPN values follow from that parent. Within the one-invariant transverse EFT, the rank-one identity, static response coefficient, clamped-frequency ratio, and two-oscillator angular measure are exact. The one-entropy loading, thermal anchoring of the clamped response, influence kernel, and lensing response remain open. The marked-transfer action gives G_* at -0.073σ and the muon and tau ratios at -0.535σ and $+0.481\sigma$. These are historical postdictions from one zero-continuous-fit action, not three blind predictions.

The medium reaches further, into territory the paper holds more tentatively. Transport gives the field a finite propagation speed and the lag that clusters and mergers need. The cluster sector reads lensing anomalies as a phase-dependent projection of an already-fixed coefficient. The cosmological mode pushes the sound horizon in the direction the Hubble tension wants. In the saturated early universe the medium becomes a conserved committed density that gravitates as pressureless dust. In strong fields the exact spherical reduction identifies the composite q_{geo} and recovers Schwarzschild, while domain termination, boundary microphysics, and the nonspherical capacity map remain open. These are extensions and frontier completions, not closed results, and the closure table marks the distinction.

Part of the microscopic picture has now been run as well as written. Coupled to a dynamical simplicial host, the admissibility ensemble orders its label sector and produces short-range capacity and geometric responses around pinned closure failures, while the conserved-current instrument passes its pre-registered gates over the measurable range (Section 23). The same calculations show that the tested substrate mechanisms do not rank vacuum geometries. Regge/CDT supplies the smooth host state, and the substrate contributes a predicted free-energy dressing whose scaling family and placebo line are measured. The capacity response has replicated across seeds; the geometric magnitude remains short-ranged with across-seed scatter; and the host phase map has a candidate region and boundary awaiting confirmation. Section 24 records those grades.

Many-Pasts assigns a record-conditioned ontology to the decoherent histories compatible with the present. Its operational branch changes no laboratory prediction: no-signaling follows from standard quantum instruments, and the Born form is the unique normalized noncontextual additive measure on a sufficiently rich record-projector lattice. The Hilbert-space kinematics and decoherence functional remain imported. Faithful full-support resolution selects the memoryless dressing kernel. The decorated vertex prepares its diagonal fresh state, reversible dilation exports the old present into history, and the marked spectrum plus electron anchor fixes the clock and phase frequency. Durable history capacity and the Past-Hypothesis/mixing package needed for a thermodynamic arrow remain open.

The remaining work is specific. A geometric GFT must realize the oriented present/history decoration, select a stable condensate, provide durable history capacity, and prove the required relative-mode gap and source projection. The transverse influence functional must reproduce the RAR without an extra form factor and determine lensing, slip, and Ward identities. The separate finite-loop stiffness operator still needs its graph audit; cosmology needs a full Boltzmann likelihood; the cluster lift needs its coherence profile and lensing-map test; and the strong-field boundary needs its channel-to-area map and spectroscopy. The marked event, its edge Hessian, and its finite routing graph are no longer on that open list.

The case for the picture is its economy: one finite-capacity medium, counted once in the ultraviolet, fixes the coefficients in advance rather than fitting them one at a time. That same economy is its exposure — with nothing left to adjust in the closed sector, one clean contrary measurement would settle it. The aim of the paper has been to make that confrontation possible.

Appendix A: Symbol Dictionary and Canonical Conventions

Appendix A gathers the conventions used throughout the technical material that follows, fixing the units, field definitions, and couplings in one place before the denser calculations begin.

Plain-language terms. Several physical words recur throughout the paper and are collected here in plain form before the symbols:

- *Capacity* — the entanglement support locally available in the vacuum medium.
- *Defect* — a stable, localized commitment of that capacity; coarse-grained, a particle.
- *Deficit* — capacity no longer freely available near a defect, $\delta S = S_\infty - S_{\text{ent}}$.
- *Capacity strain* — the extended deficit profile whose fractional value gives the weak-field potential and whose gradient gives the gravitational field.
- *Committed capacity* — capacity locked into transferring on behalf of a defect; the conserved carrier of the saturated phase (Section 18.5).
- *Dressing* — the cloud an elementary defect builds by resolving the boundary sectors; its determinant transfer and marked response set the length in the decorated scale branch.
- *Saturation* — the regime in which the capacity bath has no slack left, with the available transfer channel at its ceiling.
- *Pinned reading* — the statement that the saturated phase holds exactly at that ceiling and stays there.
- *Admissibility* — the weighting that favors boundary configurations close to a regular, isotropic local cell.
- *Closure* — the condition that the four oriented face data sum to zero, so the cell closes into a regular volume; K^2 measures the failure of closure.
- *Many-Pasts* — the postulate that one recorded present is supported by its compatible decoherent pasts. Their operational probabilities are the diagonal decoherence-functional weights conditioned on that present; $e^{-D(h,P)}$ is only shorthand for those normalized quantum probabilities.

A.1 Units, signature, and entropy normalization

All dimensional quantities are expressed in SI units unless noted otherwise. The metric signature is $(-, +, +, +)$. Covariant spacetime integrals use $x^0 = ct$, so the Einstein–Hilbert coefficient is $c^3/(16\pi G)$; after writing $dx^0 = c dt$, the ADM coefficient is $c^4/(16\pi G)$. 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 $V_* = L_*^3$; it is the bipartite primitive cell containing the two paired tetrahedral sites, so the site density is $2/L_*^3$ and the four-cell $\Delta V_4 = L_*^4/c$ is assigned one per primitive cell — the convention under which $\gamma_Q = 4\hbar c J/(3L_*^2)$ and $\kappa/\gamma = 3L_*/(4G_{\text{tet}}(0)\kappa_m(L_*))$ are simultaneously exact (Appendix C.4–C.5). In the decorated electron-anchored marked-transfer branch,

$$L_* = -\frac{3}{2} Z_e \lambda_e \ln(1 - e^{-7g_{\text{share,eff}}}), \quad \lambda_e = \frac{\hbar}{m_e c}, \quad Z_e = (1 + \zeta_*)(1 + 7\zeta_*^2).$$

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

These conventions matter because the argument repeatedly moves between a dimensionless UV counting problem and a dimensionful continuum EFT. The units and signature make those two descriptions genuinely comparable.

A.2 Core scalar variables

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

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

with vacuum baseline S_∞ and deficit

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

For nonlinear work the bounded occupancy fraction is

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

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

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

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

A.3 Couplings and derived observables

The main-text conventions are

$$\gamma : \text{entanglement-field stiffness}, \tag{12}$$

$$\kappa : \text{continuum defect-entropy coupling}, \tag{13}$$

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

$$\zeta_* \equiv 9e^{-g_{\text{share,eff}}} \left(1 - \frac{8\eta_*}{49}\right)^{21/2}, \tag{15}$$

$$Z_e \equiv (1 + \zeta_*)(1 + 7\zeta_*^2), \tag{16}$$

$$L_* \equiv -\frac{3}{2} Z_e \lambda_e \ln(1 - e^{-7g_{\text{share,eff}}}) \quad \text{in the decorated marked-transfer branch}, \tag{17}$$

$$G_* = \frac{c^3 L_*^2}{\hbar} = \frac{9}{4} \frac{\hbar c}{m_e^2} Z_e^2 \ln^2(1 - e^{-7g_{\text{share,eff}}}), \tag{18}$$

$$G_{\text{tet}}(0) : \text{tetrahedral on-site Green constant}, \tag{19}$$

$$g_{\text{share,max}} = \ln(1680), \tag{20}$$

$$g_{\text{share,eff}} : \text{admissibility-weighted sharing entropy}, \tag{21}$$

$$J_{\text{bare}}, J_{\text{eff}}^{\text{tree}}, J_{\text{eff}}^{(\text{ren})} : \text{UV edge couplings}, \tag{22}$$

$$a_0 = \frac{cH_0 g_{\text{share,eff}}}{4\pi^2}. \tag{23}$$

The canonical weak-field bridge and Newton closure are

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

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

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

A.4 Notation map

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

$$\epsilon \equiv \frac{g_{\text{share,eff}}}{4\pi^2}, \quad \nu(x) \equiv \frac{1}{1 - e^{-x}}, \quad x \equiv \frac{E_\perp}{k_B T_H},$$

W_{bath} : cluster bath source weight, B_{bath} : measured bath-development fraction,

D : capacity diffusivity, τ_0 : transport relaxation time, $D/\tau_0 = c^2$,

$\sigma_* \equiv \pi/g_{\text{share,eff}}$, $a_{\text{UV}} \equiv 1/\text{Var}_{\eta_*}(K^2)$.

In the transverse EFT, $\Delta C = \alpha_C + g_C X + \beta_C R^2$ is the single capacity mismatch, $u_C = F''(0)$ its inverse susceptibility, A_L , A_T , and $C_\times$ are the entries of its (X, R) Hessian, and $W(g)$ is the conservative response potential satisfying $W'(g) = \nu(\sqrt{g/a_0})$; Z_0 is the common transverse kinetic coefficient and ω_L , ω_T the clamped channel frequencies with $\omega_{L,T}^2 = A_{L,T}/Z_0$ in the units of Section 14. These symbols describe a conditional effective completion rather than additional closed UV coefficients.

Appendix B: UV Boundary Ensemble and Admissibility Closure

Appendix B records the finite ultraviolet counting problem in its explicit form. It supports Sections 5–6 by showing that the seven-state tetrahedral ensemble and the admissibility weighting are finite and auditable rather than phenomenological dials: the theory begins from a discrete boundary ensemble and ends with a unique admissibility-closed entropy, not with an unconstrained continuum ansatz.

B.1 Minimal tetrahedral package

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

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

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

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

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

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

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

The important feature is that the counting 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.

This also settles the status of j_0 as a parameter: it has none to tune. Within the minimal construction $j_0 = 3/2$ is forced as the smallest fermionic label meeting injectivity. A larger j_0 does not describe a fluctuation inside the same cell; it specifies a different, larger boundary ensemble, with a different state count Ω_{tet} and a different closure spectrum. The minimal theory therefore fixes j_0 rather than leaving it open, and the seven-state sector is not one option among a family but the first that closes.

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

The canonical scalar closure invariant is

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

The admissibility family is

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

The admissibility precision η is fixed by stationary normalized closure evidence. The closure constraint is the vanishing of the three-component oriented-face sum $\mathbf{c}(b) = \sum_i \hat{n}_i m_i \in \mathbb{R}^3$. The invariant K^2 is not the squared magnitude of that classical vector, which is $|\mathbf{c}(b)|^2 = (4\Sigma^2 - S^2)/3$; it is the quantum expectation of the squared closure operator. With each face carrying the seven-state $j_{\text{eff}} = 3$ representation ($2j + 1 = 7$) and the product boundary state

$|b\rangle = \bigotimes_{i=1}^4 |j, m_i\rangle_{\hat{n}_i}$ built on the tetrahedral normal frame $\hat{n}_i \cdot \hat{n}_j = -\frac{1}{3}$, the closure operator $\hat{\mathbf{C}} = \sum_i \hat{\mathbf{A}}_i$ obeys

$$K^2(b) = \langle b | \hat{\mathbf{C}}^2 | b \rangle = \sum_i j(j+1) + 2 \sum_{i < k} m_i m_k \hat{n}_i \cdot \hat{n}_k = 4j(j+1) - \frac{1}{3}(S^2 - \Sigma^2),$$

with $4j(j+1) = 48$ for $j = 3$. The decomposition

$$K^2 = |\mathbf{c}(b)|^2 + \sum_i [j(j+1) - m_i^2]$$

exhibits K^2 as the mean nonclosure plus the irreducible quantum variance of the four face operators. Two remarks fix the status of this formula. First, the expectation is taken in the unprojected product state, without projection onto the gauge-invariant intertwiner subspace; this is the operator content of the framework's soft-closure convention, in which closure enters statistically through the evidence weight $e^{-\eta K^2}$ rather than as an exact constraint, and expectation values in the projected subspace would differ. Second, the seven-label alphabet and the constant 48 are one representation-theoretic choice, not two independent ingredients: given the $j_{\text{eff}} = 3$ product-state boundary realization and the tetrahedral normal frame, both are fixed together.

The evidence factor follows from the same operator. The closure operator $\hat{\mathbf{C}}$ has three components, and the normalized isotropic quadratic evidence kernel on that three-dimensional defect space carries the determinant weight $\eta^{3/2}$, so the normalized closure evidence is $\eta^{3/2} Z(\eta)$ and its logarithm is

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

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

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

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

$$\eta_* = 0.0298668443935,$$

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

K^2	$\frac{122}{3}$	$\frac{134}{3}$	$\frac{142}{3}$	$\frac{146}{3}$	$\frac{152}{3}$	$\frac{154}{3}$
mult	96	96	96	288	192	144

K^2	$\frac{158}{3}$	54	$\frac{164}{3}$	$\frac{166}{3}$	$\frac{170}{3}$
mult	384	192	48	96	48

with total multiplicity 1680 as required. In particular,

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

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

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

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

$$\langle K^2 \rangle_{\eta_*} = 50.2229154254, \quad \text{Var}_{\eta_*}(K^2) = 15.6889750078, \quad a_{\text{UV}} \equiv \frac{1}{\text{Var}_{\eta_*}(K^2)} = 0.0637390269.$$

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

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

The value $C_{\text{cl}} = 3/2$ is an identity imposed by the stationarity condition, not an independent numerical cross-check. Its reciprocal is reused as the transverse export factor in the decorated scale-setting branch discussed in Appendix D.4.

The admissibility parameter stops being free here. 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 supports the local evaluation, the cavity benchmark, and the later shell propagation. Let X denote the root-face label and Y_r the boundary record after r rooted shells. The local information observable

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

has the principal pre-nonlocal benchmarks

$$\sigma_{\text{ind}}^{\text{toy}} = 0.44997, \tag{24}$$

$$\sigma_{\text{ind}}^{\text{loc}} = 0.44708, \tag{25}$$

$$\sigma_{\text{ind}}^{\text{Bethe}}(J = 0) = 0.44749. \tag{26}$$

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

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

where $w_\alpha = n_\alpha e^{-\eta_* K_\alpha^2}$ is the rooted-class Gibbs weight, $z = 4$, and $U_{\alpha\beta}(0)$ is the rooted shared-face compatibility matrix before shell transport is turned on. Thus $\sigma_{\text{ind}}^{\text{Bethe}}(J = 0)$ is the cavity-theory benchmark of the same explicit rooted ensemble. 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 C: Edge Kernel, Finite Renormalization, and Continuum Matching

Appendix C carries the local boundary ensemble of Appendix B into edge transport, loop dressing, and 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 minimal loop 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 specified minimal return operator uses

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

The two terms have distinct return-channel origins. 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 $m = -3, -2, \dots, 3$, the minimal operator contains seven label-diagonal returns, one for each face-label channel, contributing

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

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. Permutation symmetry permits this singlet but does not fix its weight relative to the diagonal returns or exclude longer return motifs. The minimal completion assigns it the same $2/3$ transverse projection used in the tree coupling and the rooted return factor $1/(z-1) = 1/3$ on the $z=4$ graph. The collective contribution is then

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

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

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

Thus $65/9$ is fixed inside the stated minimal return operator, not derived from the closure ensemble alone. A microscopic graph-return calculation must derive the relative diagonal and singlet weights, test longer motifs, and determine whether the operator closes on these channels. Within this conditional completion,

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

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 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,

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

where a runs over one sublattice representative of each bipartite primitive cell and i runs over its four outgoing bonds. Every undirected nearest-neighbor edge is counted once. This convention is essential: summing the four bonds from both sublattices would double the variation and the source normalization. The microscopic four-cell is

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

a coarse-graining convention rather than derived cell geometry (Section 9; the abstract cell complex has no regular-tetrahedron Euclidean embedding). Because every edge occurs once, varying the discrete action gives $J_{\text{eff}}^{(\text{ren})} L_\diamond Q = s$ for a source term $-\sum_a s_a Q_a$, with $L_\diamond = 4I - A$. The diagonal inverse of this same unnormalized Laplacian is the $G_{\text{tet}}(0)$ used in C.5, so a point source obeys $Q(0) = sG_{\text{tet}}(0)/J_{\text{eff}}^{(\text{ren})}$. This removes the factor-of-two ambiguity between the stiffness and source conventions. 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

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

Using the substrate-induced scale

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

this is

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

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

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

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

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

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

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

Exclusion sufficiency lemma. The reduction of the 1680-state boundary ensemble to a single scalar source is, at the classical source level, exact rather than approximate. The canonical exclusion defect is the vacuum ensemble conditioned on the exclusion event A , $P_{\text{def}} = P_{\text{vac}}(\cdot | A)$, so its likelihood ratio is

$$\frac{dP_{\text{def}}}{dP_{\text{vac}}} = \frac{\mathbf{1}_A}{P_{\text{vac}}(A)},$$

a function of the exclusion indicator $X = \mathbf{1}_A$ alone. The indicator is therefore a sufficient statistic for distinguishing the defect ensemble from the vacuum, and coarse-graining onto it preserves the full classical relative information,

$$D(P_{\text{def}} \| P_{\text{vac}}) = -\ln P_{\text{vac}}(A) = D(P_{\text{def}}^X \| P_{\text{vac}}^X).$$

The scope of the lemma is exactly its statement: one scalar statistic carries all the local classical source information distinguishing the canonical exclusion-defect ensemble from the vacuum. It does not by itself establish that the condensate has a single infrared mode, that orientation or defect-species structure is dynamically irrelevant, or that spatial correlations carry no further information; those are dynamical questions, and the open one is stated in Appendix H.7. For independent exclusions the source is exactly additive, $-\ln P(\bigcap_i A_i) = \sum_i -\ln P(A_i)$, so the leading dilute source is linear in the defect number, with correlation corrections entering at pair order in the density.

To propagate that local defect into the lattice field equation one needs the on-site Green function of the tetrahedral/diamond nearest-neighbor Laplacian. The four-valent diamond graph is adopted as the coarse adjacency realization of the abstract face-sharing tetrahedral boundary complex; the complex enters combinatorially, since regular tetrahedra admit no face-to-face Euclidean tessellation (the dihedral angle $\arccos(1/3) \simeq 70.53^\circ$ does not divide 360°), so the adopted graph and its Green function are the working objects while cell volumes and face areas remain conventions of the coarse map. Eliminating the two-sublattice structure gives the standard Brillouin-zone representation for the diamond lattice Green function [5]

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

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

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

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

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

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

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

They obey

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

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

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

For a local horizon normal $\hat{r}$,

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

so

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

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

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

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

The corresponding local source amplitude in lattice units is therefore

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

so that

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

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

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

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

the tetrahedral projection used in the stiffness mapping gives

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

Equating this with the weak-field source equation

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

closes the canonical source-to-stiffness ratio:

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

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

The three cancellations, made explicit. The passage from the dimensionless lattice insertion $s_{\text{def}}/J_{\text{eff}}^{(\text{ren})} = \ln(7/6)/(\pi G_{\text{tet}}(0))$ to the source theorem turns on three reductions, none of which leaves a free constant. (i) *The factor π cancels against the field normalization.* The insertion is written for the occupancy field, where one isotropically averaged defect shifts $\delta Q_{\text{def}} = \ln(7/6)/\pi$; a fixed defect type (ℓ, f) instead carries $-\ln P(A_{\ell f})$, ranging over $[0.14933, 0.15819]$ with mean 0.15415642 against $\ln(7/6) = 0.15415068$, so the source theorem uses the exact isotropic benchmark under the assumption that the coarse defect population is isotropically averaged over face and label types. Converting to the entropy field through $S = \pi Q_{\text{occ}}$ multiplies by π , so the S -field shift is $\delta S_{\text{def}} = \pi \delta Q_{\text{def}} = \ln(7/6)$ and the explicit π does not survive into the source theorem. (ii) *$\ln(7/6)$ is an isotropic source benchmark, not a universal per-defect cost.* For a coarse population averaged over face and label types, the exact mean partition ratio is $6/7$, so the corresponding annealed source normalization is $\ln(7/6)$. A fixed defect type instead carries $-\ln P(A_{\ell f})$. Writing the isotropically averaged source as $\sigma_{\text{def}} = \rho/\kappa_m(L_*)$ folds that benchmark and the mass-entropy unit into the density. The geometric coefficient $3L_*/(4G_{\text{tet}}(0))$ is fixed by the tetrahedral projection and the cell length and contains no $\ln(7/6)$. (iii) *$J_{\text{eff}}^{(\text{ren})}$ cancels in the ratio.* The source amplitude s_{def} and the stiffness γ each carry one power of the renormalized edge coupling, so it cancels in the source-to-stiffness ratio κ/γ , leaving the lattice-geometric constant $3L_*/(4G_{\text{tet}}(0)\kappa_m(L_*))$. The renormalized loop coupling therefore drops out of the observable normalization.

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

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

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

C.6 Local susceptibility cross-check

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

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

which is the local zero-mode inverse susceptibility of the closure scalar. Two features of this number matter for the source theorem in C.5. First, a_{UV} is finite and positive, confirming that the closed branch at η_* has a non-degenerate zero-mode response rather than a critical singularity that would invalidate the linear Green-function matching used to derive κ/γ . Second, the same

local susceptibility controls the stability of the closure mode that is being propagated into the shell and loop calculations, so the residual fractional discrepancy between the local benchmark and σ_* in C.2 is genuinely loop-sector rather than a local-closure failure.

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

C.7 UV-to-IR payoff

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

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

The same chain feeds

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

and the decorated support-to-rate action gives the substrate-induced scale

$$G_* = \frac{c^3 L_*^2}{\hbar} = \frac{9}{4} \frac{\hbar c}{m_e^2} Z_e^2 \ln^2(1 - e^{-7g_{\text{share,eff}}}).$$

The Green-matched source theorem fixes

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

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

The coefficients this appendix supplies to the main weak-field chain are J_{bare} , $J_{\text{eff}}^{\text{tree}}$, Σ_{ret} , $J_{\text{eff}}^{(\text{ren})}$, γ , and the Green-matched source projection κ/γ . The remaining uses of S_∞ belong to the fixed normalization of the weak-field bridge, not to the source sector itself.

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

Appendix D collects the weak-field derivations that are central but too dense for the main line: the bridge law, Newtonian recovery, the electron anchor, and the EFT consistency checks.

D.1 Bridge law from the reduced action

The weak-field bridge is now fixed by an action equivalence rather than by extrapolating an exponential lapse. The Einstein scalar-constraint action reduces to

$$I_{\text{Newton}}[\Phi] = \int dt d^3x \left[-\frac{(\nabla\Phi)^2}{8\pi G} - \rho\Phi \right].$$

The capacity action becomes the same functional, up to the constant $Z_S = 2\kappa S_\infty/c^2$, under

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

Hence

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

is the unique linear field redefinition matching both the kinetic and source terms with the UV normalization. The bounded nonlinear rule is instead $N^2 = q = 1 - \delta S/S_\infty$ in a static branch. An exponential lapse shares the first derivative at the vacuum point but is not used globally because it has no finite-capacity endpoint. The complete derivation is in Appendix N.

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 ,

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

For the ordinary longitudinal branch, variation of the unreduced Einstein scalar action gives

$$\Phi = \Psi$$

with asymptotically flat boundary conditions. This equality follows from the spatial metric constraint, not from a canonical capacity stress tensor. The effective-halo rewrite of the additional galactic response 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 metric potential controls orbital dynamics and light bending in the baseline Einstein branch. Whether the transverse thermal correction obeys $\Delta\Phi_\perp = \Delta\Psi_\perp$ is not fixed by this result; it must follow from the transverse metric influence functional. A viable galactic completion must reproduce support and lensing with that one conserved metric response.

D.3 Electron anchor and composite matter

The canonical fermionic entropy increment is

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

The UV mass normalization is

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

and the running law in the closed branch is

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

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

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

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

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

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

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

D.4 Faithful sector resolution and induced G_*

The electron anchor enters the gravitational normalization through the absolute substrate length. The exact inputs are the seven-channel entropy, the memoryless kernel selected by faithful resolution, channel factorization on the lightest branch, and the decorated marked-transfer vertex of Appendix H. The calculation below separates the determinant recurrence, positive survival gap, and finite marked response before combining them in the physical cell length.

The recurrence derivation of the dictionary. Appendices H.2, H.4, and H.5 show that faithful full-support resolution selects memoryless replacement and that the lightest branch has seven factorized channels. On the commutative boundary algebra define $(\mathbf{R}f)(b) = p_*(b)f(b)$ and $\tau_p(f) = \sum_b p_b f(b)$. The state-weighted determinant is

$$\Delta_{\tau_p}(\mathbf{R}) = \exp[\tau_p(\ln \mathbf{R})] = \exp\left(\sum_b p_b \ln p_b\right) = e^{-g_{\text{share},\text{eff}}}.$$

Multiplicativity gives $e^{-7g_{\text{share},\text{eff}}}$ on the seven-channel product. This is the quenched geometric transfer rate of the record-conditioned likelihood operator. It is not the Born probability of a specified microstate or the collision probability of two renewed blocks.

The corresponding history statement follows from the exact recurrence theorem. For a stationary ergodic finite-alphabet process, if R_n is the first return time of a length- n block, then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln R_n = h \quad \text{almost surely,}$$

where h is the Shannon entropy rate [58, 59]. This theorem selects Shannon entropy over raw state count and Rényi-2 entropy for long typical blocks. It does not by itself identify the seven distinguishable sector layers with a long temporal block, and it gives no exact finite- n return probability. For the present one-layer distribution,

$$H_1 = 7.4198000, \quad H_2 = -\ln \sum_b p_b^2 = 7.4125486, \quad H_\infty = -\ln \max_b p_b = 7.1343850.$$

The distinctions are numerically consequential after seven layers. The refresh projector $|\sqrt{p}\rangle\langle\sqrt{p}|$ carries the participation scale e^{H_2} , while the determinant and almost-sure Lyapunov rate use H_1 . They are different observables. With

$$p_{\text{rec}} \equiv e^{-7g_{\text{share},\text{eff}}},$$

the positive survival operator has nonzero eigenvalue $1 - p_{\text{rec}}$. The decorated marked vertex supplies the electron factor Z_e , so the dressed electron gap is

$$E_e = -\frac{3\hbar Z_e}{2\tau_*} \ln(1 - p_{\text{rec}}).$$

Identifying the lowest charged gap with $m_e c^2$ fixes the cadence from the already-declared electron anchor,

$$\tau_* = -\frac{3}{2} Z_e \tau_e \ln(1 - p_{\text{rec}}), \quad L_* = c\tau_*.$$

Thus

$$L_* = -\frac{3}{2} Z_e \lambda_e \ln(1 - e^{-7g_{\text{share},\text{eff}}}) = \frac{3}{2} Z_e \lambda_e e^{-7g_{\text{share},\text{eff}}} [1 + O(e^{-7g_{\text{share},\text{eff}}})].$$

The logarithmic survival correction is $p_{\text{rec}}/2 \simeq 1.4 \times 10^{-23}$. The finite closure-response correction is $Z_e - 1 = 0.00530828$. No second dimensional input, maximum-throughput clock, or tetrahedral-diameter equality is introduced.

The transfer-support dual. A spatial version uses the same dressed support and assigns the leading step count

$$N_{\text{eff}} = \frac{2}{3Z_e} e^{7g_{\text{share,eff}}}.$$

The temporal grammar gives $\tau_e = N_{\text{eff}}\tau_*$, while the spatial grammar gives $\lambda_e = N_{\text{eff}}L_*$. Their common factor cancels and returns $L_* = c\tau_*$. This is a consistency identity of the same electron phase mode; the positive transfer spectrum supplies its frequency.

Internal locality and the native vertex. The selected replacement kernel reaches the full ensemble. The tested single-label move class does not: shifts preserving injectivity split each parity copy into $4! = 24$ ordering sectors of $\binom{7}{4} = 35$ states because two slots cannot exchange order without a collision. Its strength-weighted walk saturates at entropy 7.374, below $g_{\text{share,eff}} = 7.4198$. This excludes that implementation. Appendix H.8 gives the whole-tetrahedron replacement gate, and H.9 decorates it with the fresh-state amplitude and charged event. A geometric GFT still has to realize that native gate and select a stable condensate.

The reduced electron Compton wavelength is

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

and the admissibility-closed sharing entropy is

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

The closure fixed point also gives

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

so the closure-saturation factor is

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

Separating the baseline recurrence from the finite marked response gives

$$\ln\left(\frac{\lambda_e}{L_*}\right) = 7g_{\text{share,eff}} - \ln\left(\frac{3}{2}\right) - \ln Z_e + O(e^{-7g_{\text{share,eff}}}).$$

Equivalently,

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

The exact expression is $L_* = -(3/2)Z_e\lambda_e \ln(1 - e^{-7g_{\text{share,eff}}})$. It gives

$$L_* = 1.6162537014 \times 10^{-35} \text{ m},$$

which differs from the CODATA Planck length by about 8.2×10^{-7} in fractional terms, or 8.2×10^{-5} percent.

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

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

Substituting the exact marked-transfer relation gives the closed form

$$G_* = \frac{c^3}{\hbar} \left[-\frac{3}{2} Z_e \lambda_e \ln(1 - e^{-7g_{\text{share,eff}}}) \right]^2 = \frac{9}{4} \frac{\hbar c}{m_e^2} Z_e^2 \ln^2(1 - e^{-7g_{\text{share,eff}}}).$$

Numerically,

$$G_* = 6.6742890772 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2},$$

or -0.073σ relative to CODATA. The same ζ_* gives the muon and tau ratios at -0.535σ and $+0.481\sigma$. One finite action therefore replaces the three leading deficits with one correction whose coefficient and routing are fixed jointly rather than adjusted independently for the three observables. The discrepancies were already known while the vertex was being constructed, so these are high-precision postdictions. Their evidential content is the common action, the absence of a continuous fit, and the nearby kernels rejected by the finite audits.

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

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

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

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

but it is built in. The condition defining η_* is $\langle K^2 \rangle_{\eta} = 3/(2\eta)$, so $C_{\text{cl}} = \eta_* \langle K^2 \rangle_{\eta_*} = 3/2$ holds by construction. It is not an independent cross-check. The decorated scale action reuses its reciprocal $2/3$ as the transverse export and inserts it once as a global factor, not seven times.

A discrete sensitivity audit enumerates 576 nearby formulas obtained by varying the response dimension, pair count, direction factor, determinant power, statistics, and singlet overlap. Only the decorated-vertex formula lies within one muon standard deviation. This count is not a probability distribution over theories. It shows that the landing is sensitive to the field content and graph, which is why the action and its adversarial alternatives are displayed explicitly. A shared edge mode, directed complex determinant, stable bosonic determinant, or six-state symmetric response misses the muon respectively by approximately $1.15 \times 10^4\sigma$, 13.8σ , 56.7σ , and $7.64 \times 10^4\sigma$. Appendix H derives the selected entries before performing the observable trace.

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

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

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

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

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

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

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

For a single sector, the refresh kernel

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

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

Let

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

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

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

The displayed order is only a representative of the symmetrized one-pass class. It introduces no physical ordering or extra factor of $7!$. The one-pass prescription visits each simple sector once. A history with omitted sectors is unresolved, while additional visits belong to a multi-pass dressing history rather than to repeated fermionic occupation.

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

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

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

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

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

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

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

The decorated marked-transfer action gives the exact support-scale relation

$$\boxed{\frac{\lambda_e}{L_*} = \frac{2}{3Z_e [-\ln(1 - e^{-7g_{\text{share,eff}}})]} = \frac{1}{Z_e} \dim_{\text{eff}}(\mathcal{D}_{e,\perp}) \frac{r}{-\ln(1 - r)}}.$$

Thus $\lambda_e/L_* = Z_e^{-1} \dim_{\text{eff}}(\mathcal{D}_{e,\perp}) [1 - r/2 + O(r^2)]$. Appendix H derives the memoryless replacement, determinant recurrence, positive gap, and decorated charged vertex. Fermionic exclusion and subadditivity select $k = 7$ and $\Delta_7 = 0$, while the marked graph fixes Z_e . What remains is the stable geometric GFT embedding, durable history capacity, and full source-coupled condensate spectrum.

The relation is fixed inside the displayed finite transfer action and the electron anchor. Its remaining conditionality is the broader claim that this transfer action is the charged sector of the eventual geometric GFT, not an unfixed number in the finite calculation.

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

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

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

D.5 EFT consistency checklist

The ordinary longitudinal branch is a rewriting of the Einstein constraint sector, so it introduces no independent scalar propagator to which a separate ghost or tachyon test could be applied. Its consistency checks and those of the optional transport completion must be stated separately.

- *No extra longitudinal degree of freedom.* The Einstein parent supplies the Cauchy data and propagating tensor modes; δS is the static scalar-constraint coordinate.
- *Correct-sign sourcing.* Positive mass produces a capacity deficit and attractive Newtonian response in the reduced action.
- *Positive static quadratic form.* The Euclidean capacity functional has positive stiffness $\gamma > 0$ after the overall gravitational-sign convention is fixed.
- *Conditional causal transport.* The phenomenological telegrapher completion has finite characteristic speed when $D/\tau_0 = c^2$ with $D, \tau_0 > 0$.

These statements do not establish a UV completion or quantize an additional scalar. They show that the controlled static representation is consistent with its Einstein parent and that the separately proposed transport equation is linearly stable in its stated parameter range.

The checklist is intentionally modest. It verifies the reduced Einstein representation and the signs of the separate transport model; it does not establish an independently quantized scalar EFT.

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 phenomenological transport vacuum is linearly stable.

D.6 Quadratic fluctuations and weak-field stability

Before imposing the Einstein scalar constraint, a formal covariant quadratic representation is

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

After constraint reduction only its spatial Poisson quadratic form remains. The absence of a mass term expresses the unscreened constraint and does not prove an additional propagating boson. The transverse mode of Section 14 belongs to the separate conditional condensate completion of Appendix N.8.

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

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

Appendix E collects phenomenological time-dependent and homogeneous extensions of the static capacity variable. These additions are not implied by the Einstein constraint reduction and must recover it in their static limit. The transport relation is fixed only in its preferred branch, while the homogeneous and perturbation sectors remain open.

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

$$\tau_0 s^2 + s + Dk^2 = 0$$

has the roots

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

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

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

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

$$\frac{\delta F_{\text{pond}}}{F_{\text{static}}} \sim e^{-T/(2\tau_0)} \left(\frac{\omega_{\text{orb}}}{\omega_k} \right)^2 \lesssim 10^{-6}$$

for representative orbital speeds, with the precise value depending on the averaging interval and system scale. The estimate keeps the correction below the near-stationary weak-field branch but does not supply a universal 10^{-8} bound.

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 conditional cosmological proposal, the trace-sourced homogeneous mode acts as a transient early-energy contribution. It is intended to reduce the sound horizon without rewriting the local static Poisson law. Demonstrating that separation together with the committed component requires the open joint Boltzmann calculation.

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

The transport relation and preferred branch are closed; the cosmological sector remains structurally supported but not yet Boltzmann-closed.

Appendix F: Spherical Strong-Field Reduction and Capacity Boundary

Appendix F records the strongest action-level statement currently available in the black-hole sector. The bounded capacity rule is realized invariantly in spherical symmetry by the areal-radius gradient $q_{\text{geo}} = (\nabla R)^2$. The Schwarzschild exterior and the location $q_{\text{geo}} = 0$ then follow from the reduced Einstein action. Treating that surface as the physical end of the substrate EFT, however, is an additional domain postulate whose boundary microphysics remains open.

F.1 Bounded capacity and the unique lapse map

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

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

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

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

$$N = f(q).$$

The conditions are:

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

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

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

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

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

The weak-field bridge gives

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

so $\alpha = 1/2$ and therefore

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

Equivalently, if one writes $N^2 = F(q)$, the unique continuous multiplicative completion is $F(q) = q$. The nonlinear static lapse rule is therefore fixed by capacity composition and weak-field matching; it is not a freely chosen black-hole ansatz. Because a lapse depends on foliation, this is a static constitutive statement. Its covariant spherical content is derived next.

F.2 Einstein action reduced to the capacity-adapted invariant

Take the most general spherically symmetric line element

$$ds^2 = h_{ab}(x) dx^a dx^b + R^2(x) d\Omega^2.$$

Here $x^0 = ct$, so the two-dimensional measure is $d^2x = dx^0 dr$. The same action written as $dt dr$ acquires one additional factor of c . The four-dimensional curvature decomposes as

$${}^{(4)}R = {}^{(2)}R + \frac{2}{R^2} [1 - (\nabla R)^2 - 2R\Box R].$$

After the angular integral, the Einstein–Hilbert plus GHY action becomes, up to the asymptotic and corner terms retained in $I_\partial^{(2)}$,

$$I_{\text{sph}} = \frac{c^3}{4G} \int d^2x \sqrt{-h} \left[R^2 {}^{(2)}R + 2(\nabla R)^2 + 2 \right] + I_{\text{matter}}^{(2)} + I_\partial^{(2)}.$$

The integration by parts that converts $-4R\Box R$ into $+4(\nabla R)^2$ is legitimate only together with the reduced GHY contribution; this is why the boundary term is part of the statement. For $\partial\mathcal{M}_4 = \partial\mathcal{M}_2 \times S^2$ one has

$$K^{(4)} = K^{(1)} + \frac{2}{R} n^a \nabla_a R.$$

The second term cancels the surface term from the integration by parts, leaving

$$I_\partial^{(2)} = \varepsilon \frac{c^3}{2G} \int_{\partial\mathcal{M}_2} dy \sqrt{|\gamma_{(1)}|} R^2 K^{(1)} + I_{\text{joint}} + I_{\text{ref}},$$

with ε the standard orientation sign.

Define

$$q_{\text{geo}} \equiv (\nabla R)^2 = h^{ab} \partial_a R \partial_b R.$$

Varying R and h^{ab} in vacuum gives

$$R {}^{(2)}R - 2\Box R = 0,$$

$$2R(h_{ab}\Box R - \nabla_a \nabla_b R) + h_{ab}(q_{\text{geo}} - 1) = 0.$$

These equations imply

$$\nabla_a M_{\text{MS}} = 0, \quad M_{\text{MS}} = \frac{c^2 R}{2G} (1 - q_{\text{geo}}).$$

Therefore

$$q_{\text{geo}} = 1 - \frac{2GM_{\text{MS}}}{c^2 R}.$$

The variational status is now unambiguous: q_{geo} is a composite of the two-dimensional metric and the areal-radius dilaton, and its vacuum profile is a first integral. No scalar has been added to the Einstein degrees of freedom.

Audit of the discarded multiplier form. If instead one writes an ADM term $\sqrt{h} \lambda (N^2 - q)$ while giving q no other bulk dependence, the q equation sets $\lambda = 0$. The constraint then only identifies an arbitrary scalar with a foliation-dependent lapse. It does not reproduce the capacity source equation and it has no invariant content away from a specified static slicing. The spherical reduction above supplies the invariant result that construction was trying to capture and supersedes it.

F.3 Variational status of the $q_{\text{geo}} = 0$ boundary

For $\epsilon > 0$, let $\mathcal{B}_\epsilon$ be the timelike level surface $q_{\text{geo}} = \epsilon$. The exterior Einstein problem is well posed with the standard term

$$I_{\text{GHY}}[\mathcal{B}_\epsilon] = \epsilon \frac{c^3}{8\pi G} \int_{\mathcal{B}_\epsilon} d^3 y \sqrt{|\gamma|} K$$

and fixed induced metric on $\mathcal{B}_\epsilon$, together with the reference term at infinity and any required joints. The limit $\epsilon \rightarrow 0^+$ is null and must be taken using the corresponding null-boundary and joint prescription; a bare timelike GHY expression cannot simply be evaluated at the null surface.

This construction fixes the universal gravitational variation. It does not make

$$q_{\text{geo}}|_{\partial\mathcal{M}_q} = 0$$

a new Euler–Lagrange boundary condition. The zero is the level set at which the spherical invariant becomes marginal. Declaring that level set to be the end of the physical substrate domain,

$$\mathcal{M}_q = \{q_{\text{geo}} > 0\},$$

is the bounded-capacity postulate.

Any additional functional

$$\Gamma_{\partial q}[\sigma_{AB}, \text{boundary channels}]$$

would describe genuine substrate physics: absorption, partial reflection, relaxation, entropy, or a moving-boundary stress. Its variation may contain a boundary stress and a response conjugate to the limiting capacity, but neither its form nor its spectrum follows from Einstein–Hilbert reduction. The static exterior can be solved without inventing this term; claims about excision, infalling evolution, or finite reflectivity cannot.

F.4 Static spherical vacuum exterior

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

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

so the spherical geometric representative of the capacity variable is

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

with

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

This also agrees with the weak-field capacity deficit:

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

The exterior $r > r_h$ is exactly the standard Schwarzschild exterior, and $q_{\text{geo}} = 0$ at $r = r_h$. The geometric invariant becomes negative in the trapped region. Interpreting that sign change as exhaustion of a nonnegative substrate capacity motivates restricting the capacity EFT to $q_{\text{geo}} \geq 0$, but this restriction is not implied by the Einstein equations. The absence of a physical classical interior is therefore a conditional substrate claim, not an action-level theorem.

The exterior-domain result does not decide what an infalling observer experiences at the $q_{\text{geo}} = 0$ surface. Whether a capacity-exhaustion boundary is smooth, dissipative, anomalous, or absent requires the open functional $\Gamma_{\partial q}$, not merely the exterior Schwarzschild solution.

Geometric identification of the capacity variable. The spherical reduction supplies a normalization-independent gradient invariant on the orbit space,

$$q_{\text{geo}} = h^{ab} \partial_a R \partial_b R = |\nabla R|^2,$$

which in Schwarzschild gives

$$q_{\text{geo}} = 1 - \frac{2GM}{c^2 r} = N^2,$$

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

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

The marginal-trapped-surface condition $\theta_+ \theta_- = 0$ coincides with $q_{\text{geo}} = 0$ independently of null normalization: positive, zero, and negative q_{geo} label untrapped, marginal, and trapped spherical regions. Standard collapse in the metric parent can therefore produce the geometric zero. Equating that zero with substrate saturation and refusing the negative branch are the additional bounded-capacity interpretation to be tested by the transport and boundary theory.

Rotating and charged stationary exteriors. The baseline metric parent also admits Kerr, Reissner–Nordström, and Kerr–Newman exteriors when the appropriate conserved gauge sector is present. Their standard horizon entropy

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

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

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

What is not yet supplied is a nonspherical covariant capacity scalar whose zero selects the outer horizon and whose sign defines a physical domain. The statement that inner or trapped regions are absent from the substrate theory must therefore not be exported from the spherical invariant by analogy.

F.5 Horizon thermodynamics and boundary capacity

Because the exterior geometry is unchanged, semiclassical quantities depending only on the exterior near-horizon saddle are unchanged. The Euclidean continuation is used here as an exterior-saddle calculation: the resulting periodicity depends on regularity of the near-horizon exterior geometry, not on adopting the bounded-domain interpretation past $q_{\text{geo}} = 0$. The Euclidean regularity argument therefore gives the standard Hawking temperature [9, 10],

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

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

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

We now exhibit the $1/4$ area coefficient as an exact identity between two substrate inputs already in the framework, rather than as a stipulated channel-counting rule. Both ingredients are fixed within the bulk EFT before any horizon machinery is introduced. What the identity does not yet supply is the microscopic count converting boundary channels to physical area; that gap is stated explicitly below rather than absorbed into the convention.

Per-channel cut entropy from fermionic face exclusion. If the bounded-capacity domain terminates at the spherical surface $q_{\text{geo}} = 0$, a tetrahedral cell on the exterior side whose outward face would be paired with a neighbor across that face instead has an unfilled pairing slot. By Postulate II the elementary face slot is fermionic and admits only the occupied (paired) state or the excluded (unpaired) state. A horizon cut face is then an instance of the same elementary face-exclusion defect that anchors the one-bit fermionic sector in the bulk. The seven-state $j_{\text{eff}} = 3$ channel belongs to the paired bulk link $j_0 \otimes j_0 \rightarrow j_{\text{eff}}$, not to the cut face itself: with the partner cell absent, the paired representation is never formed. The seven-state structure enters the boundary count through the bulk graph response and hence through the channel density below, not through the per-channel entropy. Conditional on the domain postulate, the entropy carried by an elementary cut defect is consequently the same primitive fermionic increment that anchors the electron at $\kappa_m(\lambda_e) = m_e/\ln 2$,

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

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

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

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

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

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

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

This n_{hor} is a dimensionless Green-response number integrated over the angular measure. It is not yet a count of independent boundary channels per unit physical area: the conversion from channels to A/L_*^2 requires a microscopic area operator or graph-flux count on the embedded cut — diamond-lattice bonds crossing a surface can be counted once a bond length and surface prescription are chosen, but that count is new structure the framework does not currently derive, and its numerical test is the graph-ensemble Monte Carlo listed among the closure tests of Appendix F.7.

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

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

the product of the two substrate inputs is

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

an exact identity in which the Joyce diamond-lattice constant $G_{\text{tet}}(0)$ and the fermionic increment $\ln 2$ cancel between the two factors. In the cell-normalized horizon convention — one independent channel per cell area L_*^2 of the cut — a horizon of area A then carries the dimensionless entropy

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

and the Bekenstein–Hawking coefficient is recovered after the gravitational scale is matched so that $L_P(G_*) = L_*$. The status of the result is a horizon normalization identity. Within the stated convention the bulk Green-response number and the capacity normalization combine algebraically to exactly $1/4$; the per-channel $\ln 2$ from fermionic face exclusion and the transverse projection $2/3$ are genuine substrate inputs, while the channel-to-area conversion is the convention. What would upgrade the identity to a derivation is the open microscopic count above — an area operator establishing that conversion independently. The remaining strong-field tasks are that count and the nonuniversal boundary spectroscopy — the explicit microscopic Hamiltonian behind the relaxation spectrum, stretched-layer corrections, and transient response.

Composition of the coefficient. Written as a product of its surviving factors, the coefficient is

$$\frac{1}{4} = \frac{2}{3} \times \frac{3}{8},$$

and the $2/3$ is the transverse projection $G_{\perp}/G_{\text{tet}}(0)$ — the same export factor that sets the electron’s leading relation $\lambda_e/L_* = (2/3Z_e)e^{7g_{\text{share,eff}}}[1 + O(e^{-7g_{\text{share,eff}}})]$ (Section 13.3) and enters the loop self-energy through the edge map $(2/3)(1/3) = 2/9$ (Section 8). The marked factor Z_e dresses the charged support without changing this geometric projection. A change to the transverse export would still displace the electron, stiffness, and horizon normalizations together.

Coherence with the bulk normalization. The identity supplies a normalization consistency check. With the exterior-saddle temperature, the Clausius relation at entropy density $1/(4L_*^2)$ returns the induced constant $c^3L_*^2/\hbar = G_*$ through the Jacobson construction [20]. The horizon and bulk routes share the same normalization conventions, so their agreement is not independent evidence. A microscopic channel-to-area count would upgrade the identity to a derivation. Dorau and Much [21] independently show that relative entropy for a coherent scalar excitation on a bifurcate Killing horizon equals the Killing-weighted energy flux; their Einstein-equation result assumes the entropy–area relation and does not supply the substrate count. Regional relative entropy is unbounded whereas substrate capacity is bounded, so any identification between them must remain weak-field and local, with a saturating continuation near $q = 0$.

F.6 Absorption, ringdown, and echoes

The exterior Regge–Wheeler/Zerilli operators are unchanged. With $q_{\text{geo}} = 0$ at the marginal surface, the tortoise coordinate $r_* \sim r_h \ln q_{\text{geo}}$ sends the surface to $r_* \rightarrow -\infty$, and the near-horizon wave equation reduces to $(c^{-2}\partial_t^2 - \partial_{r_*}^2)\psi \simeq 0$. If the usual GR future-horizon regularity condition is retained, it selects

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

and therefore

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

The standard greybody factors and quasinormal spectrum then follow [11, 12]. But if the substrate EFT truly terminates at the marginal surface, future-horizon regularity is a boundary choice, not a consequence of the exterior differential operator alone. The open functional $\Gamma_{\partial q}$ must determine whether that choice is correct.

If the microscopic boundary has finite reflectivity or lies on a stretched layer

$$q = \epsilon > 0,$$

the region between the exterior potential barrier and that layer behaves as a cavity. A typical echo delay then scales as

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

where τ_{ch} is a boundary-channel relaxation time. Thus $\mathcal{R} = 0$ is the GR-matching boundary condition, while $\mathcal{R}(\omega)$ is a boundary observable to be calculated rather than assumed.

F.7 Dynamical formation as a free-boundary problem

Standard collapse in the metric parent can form a marginal surface $q_{\text{geo}} = 0$. It does not prove that a substrate capacity field saturates there or that the physical evolution terminates. That identification requires a bounded causal transport law for a capacity variable q_{cap} , together with a constitutive relation showing $q_{\text{cap}} = q_{\text{geo}}$ in the regime of overlap. A candidate test system is

$$\partial_t q_{\text{cap}} + D_i J^i = -\Gamma(q_{\text{cap}}) \Sigma[T_{\mu\nu}],$$

$$\tau_J(\partial_t + \mathcal{L}_v)J^i + J^i = -D(q_{\text{cap}})D^i q_{\text{cap}}.$$

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

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

Choosing trial constitutive functions such as

$$D(q_{\text{cap}}) = D_0 q_{\text{cap}}(1 - q_{\text{cap}}), \quad \Gamma(q_{\text{cap}}) = \Gamma_0 q_{\text{cap}}$$

makes $q_{\text{cap}} = 0$ and $q_{\text{cap}} = 1$ invariant sets, preventing overshoot. These functions are examples, not a derived action.

If q_{cap} first reaches zero on a two-surface and the microscopic theory licenses domain termination, that surface becomes a moving boundary

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

The level-set kinematics are fixed by differentiating $q_{\text{cap}}(t, X(t)) = 0$ along the moving surface:

$$V_n = -\frac{\partial_t q_{\text{cap}}}{|\nabla q_{\text{cap}}|} \Big|_{q_{\text{cap}} \rightarrow 0^+}.$$

In spherical symmetry this becomes

$$\frac{dr_f}{dt} = -\frac{\partial_t q_{\text{cap}}}{\partial_r q_{\text{cap}}} \Big|_{r=r_f(t)}.$$

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

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

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

F.8 Weak-field boundary and closure statement

In the weak-field Solar-System regime, the metric-only parent yields

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

with the remaining standard PPN coefficients vanishing under the usual assumptions. This follows because the ordinary longitudinal capacity functional is a reduced representation of Einstein gravity, not because a scalar correction happens to be small. The separate transverse galactic influence functional must still be shown to decouple sufficiently in the high-acceleration regime.

The weak-field capacity coordinate ceases to be adequate when

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

which is the regime where the spherical invariant q_{geo} provides the controlled nonlinear description.

Appendix F therefore closes one precise strong-field statement: the spherically reduced Einstein action makes $q_{\text{geo}} = (\nabla R)^2 = 1 - 2GM_{\text{MS}}/(c^2 R)$ a composite first integral and returns the Schwarzschild exterior. The bounded capacity interpretation of q_{geo} , the exclusion of the negative branch, and any physical boundary dynamics are additional hypotheses. The $1/4$ coefficient is an exact normalization identity within the cell-area convention, while the microscopic channel-to-area count remains open. Rotating and charged solutions belong to the baseline metric parent, but their capacity-variable completion is also open.

Appendix G: Many-Pasts, Operational Closure, Branch Realization, and the Arrow of Time

Appendix G separates the three probability objects used by Many-Pasts: amplitudes for unresolved alternatives, probabilities for decoherent record histories, and conditional probabilities for histories compatible with one present record. This repairs the ambiguity in the earlier notation $P(H | P) \propto e^{-D(H,P)}$. The operational construction is standard decoherent-histories quantum mechanics [60]; Many-Pasts supplies its record-conditioned ontology.

G.1 What is a history of the entanglement network?

A coarse projective history $h = (\alpha_1, \dots, \alpha_n)$ is a sequence of alternatives at ordered substrate times. The alternative α_k is represented by a projector $\Pi_{\alpha_k}^{(k)}$. With unitary evolution $U_{k,k-1}$ between times, its class operator is

$$C_h = \Pi_{\alpha_n}^{(n)} U_{n,n-1} \Pi_{\alpha_{n-1}}^{(n-1)} \dots U_{2,1} \Pi_{\alpha_1}^{(1)} U_{1,0}.$$

This operator retains amplitudes. It is defined before any classical probability is assigned to the individual history.

Given an initial state ρ_0 , the decoherence functional is

$$\mathcal{D}(h, h') = \text{Tr} \left(C_h \rho_0 C_{h'}^\dagger \right).$$

A family admits ordinary probabilities when its off-diagonal terms are negligible at the required accuracy,

$$\mathcal{D}(h, h') \simeq 0 \quad (h \neq h').$$

The diagonal entries $p(h) = \mathcal{D}(h, h)$ are then nonnegative and additive under coarse-graining. When alternatives do not decohere, their class operators must be added before the probability is evaluated. For $C_A = \sum_{h \in A} C_h$,

$$p(A) = \text{Tr}(C_A \rho_0 C_A^\dagger),$$

which retains the interference terms. Many-Pasts places no classical distribution over unresolved fine-grained paths.

G.2 What is the present coarse configuration P ?

The present P is a macroscopic record represented by a final projector Π_P . Let $\mathcal{H}_P$ be a decoherent family of histories whose final alternatives refine that record. Exhaustiveness and decoherence give

$$p(P) = \sum_{h \in \mathcal{H}_P} p(h) = \text{Tr}(\Pi_P \rho_{\text{now}}).$$

The conditional Many-Pasts measure is

$$p(h \mid P) = \frac{p(h)}{p(P)}, \quad h \in \mathcal{H}_P.$$

The common measure over all possible records is normalized first; conditioning on the realized record comes afterwards. Normalizing a new set of histories separately for each already-selected present would leave the probabilities of the alternative presents undefined.

G.3 What is the distance $D(H, P)$?

For a decoherent history ending in P , define

$$D(h, P) = -\ln p(h),$$

with $D = +\infty$ when $p(h) = 0$. Then

$$p(h \mid P) = \frac{e^{-D(h, P)}}{\sum_{h' \in \mathcal{H}_P} e^{-D(h', P)}}.$$

The distance notation is shorthand for the diagonal decoherence-functional weight, not a second probability law. The earlier expression $-\ln \text{Tr}(\Pi_P \rho_{H \rightarrow \text{now}})$ is recovered when H already denotes a decohered preparation history and only the final record remains unresolved.

G.4 The operational Born branch

A laboratory setting x is represented by a quantum instrument $\{\mathcal{M}_{a|x}\}_a$. Each map is completely positive and trace non-increasing, while $\sum_a \mathcal{M}_{a|x}$ is trace preserving. For an initial state ρ ,

$$p(a \mid x) = \text{Tr}[\mathcal{M}_{a|x}(\rho)].$$

If $C_{h,a|x}$ refines the histories ending in record a , then

$$p(h \mid a, x) = \frac{\text{Tr}(C_{h,a|x} \rho C_{h,a|x}^\dagger)}{p(a \mid x)}$$

for a decoherent refinement. The record marginal is the Born probability by construction. Its form is also unique under the hypotheses stated in Section 3.3: normalized noncontextual

additivity on a sufficiently rich projector lattice gives $p(R) = \text{Tr}(\rho R)$ by Gleason's theorem, and medium decoherence supplies generalized record projectors for pure branch states [54, 55]. The uniqueness belongs to the measure once quantum kinematics and record completeness are granted. It does not derive the Hilbert space, the state, or the decoherence functional from the substrate. A future substrate theory must reproduce that structure rather than merely fit its projective limit.

G.5 No-signaling

Let Alice and Bob act locally on ρ_{AB} with instruments $\{\mathcal{M}_{a|x}^A\}_a$ and $\{\mathcal{N}_{b|y}^B\}_b$. Their joint record probability is

$$p(a, b \mid x, y) = \text{Tr} \left[(\mathcal{M}_{a|x}^A \otimes \mathcal{N}_{b|y}^B)(\rho_{AB}) \right].$$

Summing over Bob's record gives

$$p(a \mid x, y) = \sum_b p(a, b \mid x, y) \quad (27)$$

$$= \text{Tr} \left[(\mathcal{M}_{a|x}^A \otimes \mathcal{N}_y^B)(\rho_{AB}) \right] \quad (28)$$

$$= \text{Tr} \left[\mathcal{M}_{a|x}^A(\rho_A) \right] = p(a \mid x), \quad (29)$$

where $\mathcal{N}_y^B = \sum_b \mathcal{N}_{b|y}^B$ is trace preserving. Alice's marginal is independent of y , and the same calculation applies to Bob. Global conditioning on a joint present record does not create a controllable signaling channel because the unconditioned local marginals remain those of standard quantum mechanics.

G.6 Branch realization

The postulate makes one realized present, not many. There is no forward branching into co-real macroscopic worlds, and no collapse event selecting among them. What is real is the present with its records; the multiplicity the weight ranges over is the admissible *pasts* of that single present, not parallel futures. "Which outcome occurred" is therefore not a question about which branch the world fell into but a statement about which present records obtain, with the Born weight giving their statistics. Branch realization is thereby answered without the ontological cost of many worlds or the dynamical cost of collapse.

G.7 Arrow of time from conditional typicality

The proposed arrow-of-time extension requires a further typicality statement. Let $h = \{M_t\}_{t_i \leq t \leq t_0}$ be a decoherent macrohistory conditioned on present records M_{t_0} . A coarse Markov description would assign

$$P(h \mid M_{t_0}) \propto \mu_i(M_{t_i}) \prod_{t_i \leq t < t_0} T(M_{t+\Delta t} \mid M_t),$$

where μ_i is a boundary measure and T is the conditional macro-transition probability obtained only after summing the microscopic transitions compatible with each pair of macrostates. Multiplying a separate factor $e^{S(M_t)}$ at every time would generally count the same microscopic multiplicity twice. A neutral transition law is also insufficient: conditional counting without a low-entropy boundary condition is dominated by high-entropy pasts and Boltzmann-fluctuation histories. The required theorem must show that μ_i and T exponentially suppress those histories strongly enough for ordinary entropy-increasing histories to dominate. Neither that boundary measure nor the needed substrate transition law has been derived here. This equation defines the target without claiming a completed thermodynamic arrow.

The minimal missing boundary condition can be stated directly. A *Substrate Past Hypothesis* would require

$$\text{supp } \rho_0 \subseteq \mathcal{H}_{M_{\text{low}}}, \quad t_0 - t_i \ll \min(t_{\text{relax}}, t_{\text{rec}}),$$

where $\mathcal{H}_{M_{\text{low}}}$ is a low-entropy macro-subspace and the elapsed time is short compared with equilibration and recurrence. Under a mixing substrate dynamics, standard large-deviation counting would then favor entropy growth away from that boundary. The paper has not derived this hypothesis or the required mixing estimate. The saturated phase of Section 18.5 supplies the capacity-sector half of the hypothesis by construction — the pinned state is a near-zero-entropy configuration of the capacity sector — while the matter-sector clause remains an assumption. Naming both premises isolates the remaining arrow-of-time problem from the already-closed operational probability branch.

The faithful-history-resolution theorem of Appendix H does not supply the missing boundary. Its stationary kernel satisfies

$$p(b)K_*(b, b') = p(b)p(b') = p(b')K_*(b', b),$$

so it is exactly detailed-balanced and time-reversal symmetric. The quantum dilation is globally reversible as well. Local export distinguishes a present register from a history register within the chosen update description, but the thermodynamic orientation still comes from the boundary condition and finite-time typicality theorem above.

G.8 Memoryless dressing and the connection to L_*

The history measure does not imply memorylessness merely because it is written as an exponential. The foundational faithful full-support principle selects it when applied to paths at the fixed admissibility marginal; maximum caliber is only the conventional name for that application. For any stationary dressing kernel with one-time marginal p_{η_*} ,

$$H(B_{t+1} | B_t) = g_{\text{share,eff}} - I(B_t; B_{t+1}) \leq g_{\text{share,eff}}.$$

Equality holds only when consecutive configurations are independent, which fixes $K(b, b') = p_{\eta_*}(b')$ on the support. More generally, the entropy rate of any stationary process with this one-time marginal is bounded by $g_{\text{share,eff}}$, with equality only when the present is independent of its complete past. The quantum replacement channel transfers the old local information into a dilation register rather than destroying it. Many-Pasts gives that register its history-space interpretation. The finite marked event is a separate statement from this renewal theorem and is supplied by the decorated vertex in H.9.

G.9 Relation to many-worlds, collapse, hidden variables, and decoherence

Many-Pasts is a form of history-space realism built on the decoherent-histories formalism. It posits one realized record-bearing present and no stochastic collapse term. Its probabilities are the diagonal entries of the standard decoherence functional, conditioned on that present only after the alternative records have been normalized. The distinctive claim lies in the ontology assigned to this conditional measure and in its proposed substrate realization. Born statistics and no-signaling are imported with the operational quantum structure; the substrate does not yet derive them.

G.10 Familiar quantum examples: double-slit, EPR/Bell, and measurement

In a double-slit experiment the two path alternatives remain combined in one class operator while no which-path record exists. The probability of a detection event therefore includes their

interference term. A durable which-path record defines a decoherent refinement, after which the path histories admit separate conditional probabilities and the fringe term is suppressed.

In an EPR or Bell experiment the record is joint, so its history measure carries the standard nonclassical correlations. Summing over the remote record invokes a trace-preserving local map and returns a marginal independent of the remote setting, as shown in G.5. Conditioning on the observed joint record does not alter that prior no-signaling marginal.

Measurement creates a durable record and identifies a decoherent family. Operationally the calculation is ordinary quantum mechanics. The additional claim is that the realized present is supported by the conditional ensemble of compatible past histories.

These examples introduce no new laboratory predictions. They show where positive history probabilities are legitimate and where amplitudes must remain combined.

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. The microscopic appendix therefore supports the EFT without replacing it: the continuum kinetic term and source channel can arise from a concrete substrate realization, while a full first-principles derivation of every inhomogeneous continuum coefficient from the underlying kernel remains to be done.

It does not replace the explicit coefficient derivation given earlier, but it shows that the ontology and sign choices of the EFT are compatible with a concrete microscopic picture.

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

H.2 The dressing Hamiltonian and the lightest seven-channel branch

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

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

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

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

One-pass occupation and the exponent seven. A channel is energetically occupied when its dressing gain exceeds the fixed bare cost, $F_m(\rho_m) > \varepsilon_0$. Symmetry makes this inequality the same for all seven channels, so an all-bound configuration is a possible symmetric ground-state branch when the inequality holds. The static Hamiltonian alone cannot identify that branch with the lightest charged particle, because its ordering depends on the undetermined balance between ε_0 and the dressing free energy.

The adopted recurrence mass functional supplies the missing ordering. Let $k = \sum_m n_m$ be the number of resolved channels. Fermionic exclusion gives $0 \leq k \leq 7$, and with the admissibility-closed marginal fixed on every occupied layer the support relation generalizes to

$$\frac{\lambda_k}{L_*} = \frac{2}{3} \exp(k g_{\text{share,eff}} - \Delta_k), \quad m_k = \frac{3\hbar}{2cL_*} \exp(\Delta_k - k g_{\text{share,eff}}),$$

where $\Delta_k \geq 0$ is the total correlation among the occupied layers. For each k , the minimum occurs at $\Delta_k = 0$. Between those minima,

$$\frac{m_{k+1}}{m_k} = e^{-g_{\text{share,eff}}} \simeq 5.99 \times 10^{-4},$$

so every missing channel raises the minimum mass by $e^{g_{\text{share,eff}}} \simeq 1.67 \times 10^3$. The unique lightest resolved one-bit defect in this support-to-length branch therefore has $k = 7$ and $\Delta_7 = 0$. The factor seven and the product cloud are selected together; they need not be imposed as separate one-pass and factorization postulates. This theorem holds at fixed marginals p_{η_*} . Allowing the cloud to optimize a different marginal would change the ultraviolet ensemble rather than refine this branch.

Factorized cloud and the additive support. The dressing dimension multiplies across channels — so that the seven equal contributions $g_{\text{share,eff}}$ add rather than merge — precisely when the joint cloud state is a product, $\rho = \bigotimes_m \rho_m$. Whether the static ground state of H has this product form is controlled by the inter-channel term $V_{mm'}$. The recurrence mass functional supplies a separate ordering: at fixed marginals it selects the seven-channel product as the lightest branch even though the static Hamiltonian has not yet derived that ordering. Each

selected channel must sample its whole ensemble, and the minimizing readout has no inter-layer correlation.

The static Hamiltonian identifies the channel and cloud variables but does not order all of their branches. The next subsections prove the entropy ceiling, the unique memoryless kernel that reaches it, and the joint $k = 7$, $\Delta_7 = 0$ minimum of the adopted recurrence mass functional.

H.3 Slot coupling, layer factorization, and the additivity theorem

Two distinct correlation structures appear in the closure data, and separating them is essential: conflating them leads to the false conclusion that the cloud cannot factorize.

The closure invariant is a pure pair coupling. Expanding $S^2 = (\sum_i m_i)^2 = \Sigma^2 + 2 \sum_{i < j} m_i m_j$ in $K^2(b) = 48 - \frac{1}{3}(S^2 - \Sigma^2)$ cancels the self-terms exactly and leaves the identity

$$K^2(b) = 48 - \frac{2}{3} \sum_{i < j} m_i m_j.$$

The admissibility weight $e^{-\eta_* K^2(b)}$ therefore contains only cross-terms between distinct face slots of a single boundary state; it does not factorize over those four slots. This is the exact origin of the residual correlation between face slots on the closed ensemble,

$$I(\text{slot}_0; \text{slot}_1) = 0.1545 \text{ nats}, \quad \frac{I}{H(\text{slot})} = 0.079,$$

so the four faces of one tetrahedron are about eight percent correlated, a structural feature of the closure invariant.

Slot coupling does not obstruct layer factorization. The factorization the support relation requires is over the seven *sector layers* $b^{(-3)}, \dots, b^{(3)}$ of $\mathcal{H}_{\text{hist}}$, each layer being a full boundary state drawn from the entire 1680-state ensemble. The slot coupling $-\frac{2}{3}m_i m_j$ lives *inside* a single layer's boundary state: it relates the four faces of that one tetrahedron and never couples layer m to layer m' . The two structures act on different objects:

Structure	What it couples	Role
slot coupling $-\frac{2}{3}m_i m_j$	the four faces <i>within</i> one boundary state b	the eight-percent mutual information; lives inside each $\mathcal{H}_B^{(m)}$
layer coupling $V_{mm'}$	the seven sector layers $b^{(m)}$ of $\mathcal{H}_{\text{hist}}$	controls factorization; acts <i>between</i> the factors

The substrate's intrinsic correlation, the most natural candidate obstruction to factorization, therefore acts at the wrong level to obstruct it: it is internal to a layer, not between layers.

Additivity of the decoherent preparation overlap. Appendix G.3 defines $D(h, P) = -\ln p(h)$ from the diagonal decoherence-functional probability. For the special case in which H already denotes a decohered preparation and only the final record remains unresolved, this reduces to $D(H, P) = -\ln \text{Tr}(\Pi_P \rho_{H \rightarrow \text{now}})$. If that preparation state and the resolution projector factorize over channels, $\rho = \bigotimes_m \rho_m$ and $\Pi_P = \bigotimes_m \Pi_m$, the trace factorizes and the logarithm converts the product into a sum,

$$\text{Tr}(\Pi_P \rho) = \prod_m \text{Tr}(\Pi_m \rho_m) \implies D = \sum_m [-\ln \text{Tr}(\Pi_m \rho_m)] = \sum_m D_m.$$

The overlap distance is therefore additive in this special factorized preparation. The effective-support statement used downstream follows independently from Shannon subadditivity: a product cloud with seven equal marginals has joint entropy $7g_{\text{share,eff}}$ and perplexity $e^{7g_{\text{share,eff}}}$.

Whether the electron's dressing is that product is decided by the two conditions in the next subsections.

The ceiling and the two conditions. Subadditivity bounds the support from above. With each channel marginal fixed to the admissibility weight, $H(B_m) = g_{\text{share,eff}}$, the joint entropy of a single readout obeys

$$H(B_{-3}, \dots, B_3) \leq \sum_{m=-3}^3 H(B_m) = 7g_{\text{share,eff}},$$

and the deficit

$$\Delta = \sum_m H(B_m) - H(B_{-3}, \dots, B_3) \geq 0$$

measures the total correlation among the seven channels. The full support $\dim_{\text{eff}} = e^{7g_{\text{share,eff}}}$ is reached precisely when each channel carries its full entropy $g_{\text{share,eff}}$ and the channels are mutually independent, $\Delta = 0$. The first is a condition on each layer in substrate time; the second is a condition on the seven layers at a single readout. The next two subsections establish the unique faithful-resolution kernel and the minimal-mass factorized branch, while separating those results from the still-open microscopic implementation.

H.4 Maximum caliber: complete renewal is uniquely selected

The closure calculation fixes the one-time marginal p_{η_*} . The same foundational requirement that every admissible resolution carry the full available support can be stated on history space as maximum path entropy, conventionally called maximum caliber [56]. This is not an added dynamical premise: it is faithful full-support resolution applied to entire paths rather than to a single pass. For an arbitrary stationary process, its entropy rate is

$$h_\mu := \lim_{n \rightarrow \infty} H(B_0 | B_{-1}, \dots, B_{-n}) \leq H(B_0) = g_{\text{share,eff}}.$$

Equality holds only when B_0 is independent of its complete past. Stationarity then makes the process independent and identically distributed, with every conditional distribution equal to p_{η_*} . Thus maximum path entropy uniquely selects complete renewal. The point is not that finite capacity alone proves renewal; it does not. The stronger faithful full-support condition already used by the paper supplies the selection, and maximum caliber is its temporal formulation rather than a new premise.

The one-parameter family below illustrates the cost of retained memory:

$$K_a(b, b') = a \delta(b, b') + (1 - a) p_{\eta_*}(b'), \quad a \in [0, 1],$$

which repeats the current state with probability a and otherwise redraws from the stationary weight. Every member has p_{η_*} as its stationary distribution and the same single-time marginal, so the equilibrium ensemble cannot say which one governs the dressing. The per-channel conditional entropy $H(b' | b) = \sum_b p_{\eta_*}(b) H(K_a(b, \cdot))$, evaluated on the exact 1680-state ensemble, nonetheless slides with the memory a :

a (memory)	per-channel $H(b' b)$	status
0.0 (refresh)	7.41980 = $g_{\text{share,eff}}$	full entropy
0.1	6.99953	reduced
0.3	5.80153	reduced
0.5	4.40051	reduced
0.9	1.06642	reduced

Only the memoryless endpoint $a = 0$ — the replacement kernel $K_*(b, b') = p_{\eta_*}(b')$ of Appendix D.4 — returns the full $g_{\text{share,eff}}$. For any stationary Markov kernel with that marginal,

$$H(B_{t+1} | B_t) = H(B_{t+1}) - I(B_t; B_{t+1}) \leq g_{\text{share,eff}},$$

and equality holds if and only if $I(B_t; B_{t+1}) = 0$. The Markov corollary is therefore

$$\boxed{K_*(b, b') = p_{\eta_*}(b')}.$$

Single-label local moves conserve slot ordering, fracturing each parity copy into twenty-four sectors of thirty-five states, so no such local kernel reaches the full ensemble at any parameter value (Appendix D.4). The native-vertex construction of H.8 implements the selected nonlocal-in-label-space replacement as one operation on the complete local tetrahedral register.

The lightest-defect branch selects the same endpoint independently within the adopted support map. If a per-pass temporal mutual information I_t reduces the fresh entropy from $g_{\text{share,eff}}$ to $g_{\text{share,eff}} - I_t$, then at fixed L_*

$$\lambda(I_t) = \lambda(0)e^{-I_t}, \quad m(I_t) = m(0)e^{I_t}.$$

Every retained temporal correlation raises the recurrence mass, so electron lightness again selects $I_t = 0$. Faithful full-support resolution selects the vacuum history process; mass minimization independently selects the same process in the defect sector.

A dimensionless mixing test shows that the selected kernel is dynamically available. Take any ergodic generator Q (a symbol used only within this mixing test; the survival complement of H.6 is a distinct operator) with p_{η_*} as its detailed-balance stationary weight. The test uses nonlocal transpositions of two occupied labels, not the injectivity-preserving single-label shifts excluded in the preceding paragraph. A continuous-time swap generator built from the same admissibility weights serves. The kernel $e^{Q\tau}$ loses endpoint mutual information,

$$I(\tau=0.1) \approx 3.4 \text{ nats}, \quad I(\tau=1) \approx 0.030, \quad I(\tau=2) \approx 1 \times 10^{-4},$$

so the replacement kernel is the late- τ limit of a broad class of dimensionless mixing dynamics. Faithful full-support resolution selects the exact endpoint rather than a particular approach to it. H.8 derives the quantum replacement channel and reversible update; H.9 supplies the fresh amplitude, marked event, and update interaction. Durable history capacity and the stable geometric GFT embedding remain open.

H.5 Independent channels: the electron as the lightest defect

The second condition is that the occupied channels factorize at one readout, so that the deficit Δ_k of H.3 vanishes. Appendix H.2 already showed that electron lightness selects both the maximum occupation $k = 7$ and $\Delta_7 = 0$ once the support-to-length dictionary is adopted. The calculation here isolates the factorization part of that joint minimum.

Keep the deficit explicit in the baseline support relation. With each channel at its full entropy, the joint entropy is $7g_{\text{share,eff}} - \Delta$, and

$$\frac{\lambda}{L_*^{(0)}} = \frac{2}{3} e^{7g_{\text{share,eff}} - \Delta}.$$

At a fixed substrate scale L_* , a one-bit charged defect whose dressing carries correlation Δ has spatial support $\lambda \propto e^{-\Delta}$, and through $m = \hbar/(c\lambda)$ a mass

$$m \propto e^{\Delta}.$$

Correlation therefore contracts the support and raises the recurrence mass. Because $\Delta \geq 0$, the lightest one-bit charged defect is the product dressing $\Delta = 0$; correlated dressings describe heavier branches within the same recurrence/support map. Independence is not inferred from the absence of an interaction term. It is forced by mass minimization in the theory's own definition of the electron anchor.

Within the marked-transfer branch, the information-theoretic selections read end to end. Faithful full support fixes memoryless refresh, while fermionic exclusion and electron lightness select $k = 7$ and $\Delta_7 = 0$. H.6 gives the baseline survival scale and H.9 its finite dressing:

$$L_* = -\frac{3}{2} Z_e \lambda_e \ln(1 - e^{-7g_{\text{share,eff}}}), \quad G_* = \frac{c^3 L_*^2}{\hbar},$$

with no gravitational quantity used as input.

The decorated action realizes the nonlocal replacement, fresh state, and charged marked event. A geometric GFT must still realize that decoration in a stable condensate and determine the light fluctuation spectrum and inhomogeneous continuum coefficients; it need not re-select memorylessness, the clock conversion, or $\Delta = 0$.

H.6 Mass from the positive survival transfer operator

Faithful full-support resolution fixes the replacement process on history space but not a dimensional duration. The duration comes from the paper's existing electron anchor once the already-declared effective-support event is represented in the transfer space. No second clock principle or new premise is required.

The marked seven-channel transfer. On the factorized lightest branch, likelihood multiplication on each renewed cloud has determinant $e^{-g_{\text{share,eff}}}$. The product determinant is

$$r := \Delta_{\tau_p^{\otimes 7}}(\mathbb{R}^{\otimes 7}) = e^{-7g_{\text{share,eff}}}.$$

This is the scalar charged determinant line of the decorated transfer, not the probability of an arbitrarily chosen seven-layer microstate. Fermionic one-bit exclusion gives the positive no-loop transfer

$$T_{\text{surv}}^{(0)} = 1 - r,$$

equivalently the one-dimensional Grassmann determinant $\int d\bar{c} dc e^{-\bar{c}(1-r)c} = 1 - r$. One may therefore define

$$H_{\text{surv}}^{(0)} = -\frac{\hbar}{\tau_*^{(0)}} \ln T_{\text{surv}}^{(0)},$$

with exact raw gap

$$E_{\text{raw}}^{(0)} = -\frac{\hbar}{\tau_*^{(0)}} \ln(1 - r).$$

The unit mixing gap on the renewed probability space is a different operator and is not identified with this exponentially small charged gap.

Electron calibration and exact cell length. Before the finite marked response is included, the transverse export gives $H_e^{(0)} = (3/2)H_{\text{surv}}^{(0)}$. The baseline electron calibration is

$$m_e c^2 = -\frac{3\hbar}{2\tau_*^{(0)}} \ln(1 - r).$$

It fixes

$$\tau_*^{(0)} = -\frac{3}{2}\tau_e \ln(1-r), \quad L_*^{(0)} = c\tau_*^{(0)} = -\frac{3}{2}\lambda_e \ln(1-r).$$

For $r = e^{-7g_{\text{share,eff}}}$,

$$L_*^{(0)} = -\frac{3}{2}\lambda_e \ln(1 - e^{-7g_{\text{share,eff}}}) = \frac{3}{2}\lambda_e e^{-7g_{\text{share,eff}}} \left[1 + \frac{1}{2}e^{-7g_{\text{share,eff}}} + O(e^{-14g_{\text{share,eff}}}) \right].$$

The relative logarithmic correction to the leading expression is 1.4×10^{-23} . The baseline induced scale is

$$G_*^{(0)} = \frac{9}{4} \frac{\hbar c}{m_e^2} [\ln(1 - e^{-7g_{\text{share,eff}}})]^2,$$

whose leading form is proportional to $e^{-14g_{\text{share,eff}}}$. H.9 derives the finite response and gives $\tau_* = Z_e \tau_*^{(0)}$, $L_* = Z_e L_*^{(0)}$, and $G_* = Z_e^2 G_*^{(0)}$.

Why this supplies the phase bridge. A positive Euclidean transfer operator defines an energy spectrum. Analytic continuation supplies the Lorentzian phase frequency of the same charged mode. An independent identification of an integer renewal count with a condensate winding is unnecessary. H.9 realizes the determinant recurrence and marked response in the decorated charged vertex.

The temporal and spatial grammars agree without referring to a literal tetrahedron diameter. If N_{eff} is the common leading support factor, $\tau_e = N_{\text{eff}}\tau_*$ and $\lambda_e = N_{\text{eff}}L_*$. The identity $\lambda_e = c\tau_e$ then gives $L_* = c\tau_*$. The graph calculation uses this L_* as its canonical displacement unit; the support size of a future geometric tetrahedron, if one is introduced, is a separate quantity constrained by locality rather than a coefficient that sets the tick.

Internal locality and the native vertex. The tested single-label move cannot implement replacement. Injectivity-preserving shifts fracture the 1680 states into

$$48 \text{ components} = 4! \text{ orderings} \times 2 \text{ orientations}, \quad 35 \text{ states per component.}$$

H.8 constructs a one-layer reversible dilation when the complete tetrahedral register is the native local vertex. This establishes circuit depth one, not the physical duration; the duration is fixed spectrally by the electron equation above. H.9 decorates that gate with the fresh amplitude and marked charged event.

H.7 Conditional dressing functionals and the condensate-spectrum lemma

The dressing Hamiltonian named in Section 22,

$$H = \sum_m (\varepsilon_0 n_m - n_m F_m(\rho_m)) + \sum_{m < m'} V_{mm'}(\rho_m, \rho_{m'}),$$

carries three named functionals ε_0, F_m, V . This subsection gives a conditional reconstruction of them from a proposed mean-field condensate. It does not derive the required condensate from a specified GFT action, and the two residuals of H.6 remain separate outputs of the same missing fluctuation calculation rather than collapsing automatically into one premise.

Mean-field reduction. In a GFT condensate the boundary-data order parameter is $\langle \hat{\varphi}(b) \rangle = \sigma(b)$. The admissibility calculation fixes a normalized modulus profile $p_{\eta_*}(b)$, so a candidate condensate may be written

$$|\sigma(b)|^2 = \bar{n} p_{\eta_*}(b).$$

This fixes neither the overall density $\bar{n}$ nor any of the phases. A general 1680-component condensate has a common phase and relative-phase directions, and their physical status is determined by the kinetic and interaction kernels of the specified GFT action. Likewise, η_*^{-1} may be used as a formal closure-temperature parameter, but the identity

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

is the stationarity condition generated by the three-component determinant weight, not a literal equipartition theorem for the bounded discrete spectrum. Conditional on a condensate realization, $T_* = 1/\eta_* = 33.48$ may nevertheless be read as the dimensionless substrate closure temperature, with $\langle K^2 \rangle_{\eta_*} = \frac{3}{2}T_* = 50.22$; this is an interpretation of the same stationary point and introduces no new number.

The refresh projector. The memoryless transition matrix is $K_*(b, b') = p_{\eta_*}(b')$. Let $D_p = \text{diag}(p_{\eta_*})$ and define the detailed-balance symmetrization

$$\tilde{K}_* = D_p^{1/2} K_* D_p^{-1/2}.$$

With $v_b = \sqrt{p_{\eta_*}(b)}$, one obtains

$$\tilde{K}_* = |v\rangle\langle v| \equiv P, \quad \tilde{L}_{\text{mix}} = I - \tilde{K}_* = P_{\perp}.$$

The result is exact: the classical replacement process has one stationary direction in the weighted probability space and removes every orthogonal probability mode after one discrete update. The vector $|v\rangle$ belongs to the symmetrized classical transfer representation; it is not automatically the quantum state of the cell. It does not follow that $P_{\perp}$ is a Hermitian mass matrix in a closed Lorentzian GFT. The Markov update, a quantum density operator, and the coherent fluctuation Hessian are different objects.

Conditional refresh–Hessian bridge. The minimal reversible bridge uses the positive refresh Dirichlet form as the internal Euclidean quadratic cost,

$$\mathcal{K}_{\text{int}}^{(2)} = \Delta_{\text{ref}} P_{\perp}, \quad \Delta_{\text{ref}} > 0.$$

A fixed- j Euclidean effective action realizing this bridge is

$$\Gamma_E[\sigma] = \int d^4 x_E \left[Z \partial_{\mu} \sigma^{\dagger} \partial^{\mu} \sigma - \mu^2 \sigma^{\dagger} \sigma + \frac{\lambda}{2} (\sigma^{\dagger} \sigma)^2 + \Delta_{\text{ref}} \sigma^{\dagger} P_{\perp} \sigma \right].$$

It has $\sigma_0 = \sqrt{\bar{n}} e^{i\theta_0} v$ with $\bar{n} = \mu^2/\lambda$. With canonically normalized real fluctuations, its Euclidean quadratic operator contains a radial singlet with $m_h^2 = 2\mu^2/Z$, a common phase with $m_{\pi}^2 = 0$ when number U(1) is exact, and 1679 complex relative modes with $m_{\perp}^2 = \Delta_{\text{ref}}/Z$. Lorentzian propagation requires the usual analytic continuation and the global $(-, +, +, +)$ convention of Appendix A. The exclusion source couples at linear order to the density fluctuation h because $\delta(\sigma^{\dagger} \sigma) = 2\sqrt{\bar{n}} h + \dots$. The common phase is derivatively coupled and does not mediate a second static force. Simplicial interactions may break number U(1) and gap it.

An open-system completion may instead implement renewal as the quantum replacement channel

$$\mathcal{E}_*(\rho) = \rho_* \text{Tr } \rho.$$

The classical closure ensemble fixes the diagonal probabilities of ρ_* in the boundary basis. H.9 adopts the diagonal maximum-entropy completion and constructs its reversible update. A dissipative representation with Lindblad generator $\dot{\rho} = \Gamma_{\text{ref}}[\mathcal{E}_*(\rho) - \rho]$ gives every traceless input perturbation a relaxation gap. A geometric condensate embedding must determine how this transfer description appears in its coherent fluctuation spectrum.

Selection, mixing, and marked survival are distinct operators. The combinatorial Hamiltonian selects the admissibility profile. The mixing generator $\tilde{L}_{\text{mix}} = P_{\perp}$ controls decay toward that profile. The exponentially small electron scale comes from a third object: the scalar determinant transfer of the seven-channel likelihood operator,

$$r := \Delta_{\tau_p^{\otimes 7}}(\mathbf{R}^{\otimes 7}) = e^{-7g_{\text{share,eff}}},$$

whose one-bit survival transfer is $T_{\text{surv}}^{(0)} = 1 - r$. Over a baseline proper-time step $\tau_*^{(0)}$ its exact spectral gap is

$$\text{gap}(H_{\text{surv}}^{(0)}) = -\frac{1}{\tau_*^{(0)}} \ln(1 - r), \quad m_e c^2 = \frac{3}{2} \hbar \text{gap}(H_{\text{surv}}^{(0)}).$$

Solving gives $\tau_*^{(0)} = -(3/2)\tau_e \ln(1 - r)$ and $L_*^{(0)} = c\tau_*^{(0)}$. H.9 derives the marked factor Z_e and the physical $L_* = Z_e L_*^{(0)}$. The unit mixing gap and charged survival gap are not identified.

ε_0 , the bare cost. ε_0 is the condensate chemical potential $\mu_0 = \partial E / \partial N$, the energy to commit one cell against the mean field, of order $E_* = \hbar c / L_*$. Energies in H are measured in units of E_* , and $T_* = 1/\eta_*$ is dimensionless. The rest mass is read from the positive marked-survival operator H_{surv} , not from the static spectrum of H or the unit mixing gap of L_{mix} .

F_m , the dressing free energy. A bound channel carries a cloud ρ_m over the boundary ensemble. The proposed released free energy is $F_m(\rho_m) = T_* S[\rho_m]$. At fixed mean closure its maximum is the Gibbs state p_{η_*} with entropy $g_{\text{share,eff}}$. On the selected lightest branch the seven clouds carry total entropy $7g_{\text{share,eff}}$. The faithful state-weighted determinant of their likelihood operator is exactly $e^{-7g_{\text{share,eff}}}$. This is a quenched multiplicative transfer rate, not a raw single-draw Gibbs probability.

Two energies, kept distinct. The separation cost and rest energy are different quantities. The separation cost is the entropic free energy required to dissolve a defect into the substrate; in this construction it is large and linear in $g_{\text{share,eff}}$. The rest energy is $m_e c^2 = \hbar \Gamma_e$, with Γ_e set by the rare joint recurrence and therefore exponential in $-7g_{\text{share,eff}}$. Reading the separation cost as the mass would place the electron near tens of E_* and remove the required suppression.

V , the inter-channel coupling. The permutation-invariant singlet

$$|s\rangle = \frac{1}{\sqrt{7}} \sum_m |m\rangle$$

is the natural channel through which a scalar capacity fluctuation couples uniformly to the seven sectors. Under the refresh–Hessian bridge, the radial density fluctuation is the only linearly source-coupled singlet at static order and the relative sector is gapped. Integrating out that density mode gives

$$V_{mm'} = v \langle m|s\rangle \langle s|m'\rangle = \frac{v}{7},$$

a rank-one outer product. The per-incidence overlap is $1/7$; the separate orientation doubling gives the $2/7$ used in Appendix I.1. The implication is

$$\text{one source-coupled density singlet} \implies \text{rank}(V) = 1.$$

This result is derived inside the minimal bridge, not from the closure ensemble alone. A microscopic interaction that mixes P and $P_\perp$, closes the relative gap, or sources the common phase would add exchange channels and deform the uniform structure.

Positive-transfer phase and the exchanged mode. The projector calculation separates the static source from the transfer clock. The static exclusion source couples to the radial density singlet. Independently, the positive Euclidean survival operator defines the electron energy, and ordinary analytic continuation gives the Lorentzian phase frequency $\omega_e = E_e/\hbar$ of that same charged mode. No additional identification between a renewal count and an independently postulated condensate winding is needed. The decorated vertex supplies the finite charged routing. The remaining rank-one exchange question belongs to its geometric embedding: the coherent microscopic interaction must preserve the source projection and relative-mode gap.

The remaining microscopic lemma. The effective projector model narrows the missing calculation. For a specified fixed- j simplicial GFT action admitting $\sigma_b = \sqrt{\bar{n} p_{\eta_*}(b)} e^{i\theta}$, compute its complete Bogoliubov or Schwinger–Keldysh kernel and test

$$P_\perp H_{\text{full}} P = 0, \quad P_\perp H_{\text{full}} P_\perp \geq \Delta_{\text{gap}} P_\perp, \quad \Delta_{\text{gap}} > 0.$$

The decorated transfer vertex of H.9 fixes the charged recurrence and its edge Hessian. A geometric GFT must still determine whether the condensate realization produces a coherent mass, a relaxation gap, or both, and verify that no additional source-coupled light field survives. The closure weight fixes P exactly, while the refresh–Hessian bridge for the geometric condensate remains a selected effective completion.

H.8 Faithful history resolution and reversible Many-Pasts renewal

This subsection separates four statements that earlier versions grouped under “refresh”: the classical history process, the quantum channel, its reversible implementation, and the conversion of a positive charged spectrum into proper time. H.9 supplies the charged marked vertex. Neither maximum-caliber terminology nor the clock construction introduces a new foundational premise.

Maximum path entropy. Let $\mathcal{B}$ be the finite admissible boundary space and let $p_*(b) = p_{\eta_*}(b)$ be its fixed stationary marginal. For any stationary process $\{B_n\}$ on $\mathcal{B}$,

$$h_\mu = \lim_{N \rightarrow \infty} \frac{1}{N} H(B_1, \dots, B_N) = H(B_0 \mid B_{-1}, B_{-2}, \dots) \leq H(p_*).$$

Equality holds exactly when B_0 is independent of its complete past. Applying the same statement after every time translation factorizes every finite joint distribution,

$$P(B_1 = b_1, \dots, B_N = b_N) = \prod_{n=1}^N p_*(b_n).$$

The unique maximum-caliber history process is therefore i.i.d. For a Markov representation its transition kernel is

$$\boxed{K_*(b, b') = p_*(b')}.$$

Here “maximum caliber” is terminology for faithful full-support resolution on history space [56]. It is not a sixth theory-defining input. The foundational requirement already says that an admissible resolution carries the full available entropy compatible with its fixed data; applying that same requirement to a path gives the equality case above. Finite capacity alone would not imply renewal, but the paper’s stronger faithful full-support premise does.

The selected process carries no thermodynamic orientation. It obeys

$$p_*(b)K_*(b, b') = p_*(b)p_*(b') = p_*(b')K_*(b', b),$$

so its stationary path measure is exactly invariant under time reversal. Faithful history resolution derives local renewal, not a low-entropy past and not a physical cadence.

Proper time from the positive marked-transfer spectrum. The renewal theorem fixes the discrete process but says nothing about seconds per update. Proper time comes from the determinant survival operator of H.6,

$$r = \Delta_{\tau_p^{\otimes 7}}(\mathbf{R}^{\otimes 7}) = e^{-7g_{\text{share}, \text{eff}}}, \quad T_{\text{surv}}^{(0)} = 1 - r.$$

Defining the baseline transfer over $\tau_*^{(0)}$ gives

$$H_{\text{surv}}^{(0)} := -\frac{\hbar}{\tau_*^{(0)}} \ln T_{\text{surv}}^{(0)}, \quad E_{\text{raw}}^{(0)} = -\frac{\hbar}{\tau_*^{(0)}} \ln(1 - r)$$

on its nonzero support. The baseline transverse identification $m_e c^2 = (3/2)E_{\text{raw}}^{(0)}$ therefore yields

$$\boxed{\tau_*^{(0)} = -\frac{3}{2}\tau_e \ln(1 - r)}, \quad \boxed{L_*^{(0)} = c\tau_*^{(0)} = -\frac{3}{2}\lambda_e \ln(1 - r)}.$$

The leading form is $L_*^{(0)} = (3/2)\lambda_e r[1 + O(r)]$. H.9 supplies the finite marked factor Z_e and hence the physical $L_* = Z_e L_*^{(0)}$. No update diameter, throughput maximization, or geometric conversion factor enters. The complete tetrahedron being a native gate is a locality and circuit-depth statement only; its physical embedding size is a separate consistency question.

A positive Euclidean transfer spectrum also supplies, by analytic continuation, the Lorentzian phase frequency. The clock and phase use the same charged spectral mode and require no independent winding postulate. H.9 derives the recurrence as the state-weighted determinant and implements its finite marked dressing.

Quantum replacement theorem. Let $\mathcal{H}_A$ be the local cell Hilbert space and $\mathcal{E} : A \rightarrow A'$ a trace-preserving quantum channel. Complete local renewal means that the output contains no information about the input, including information visible through an arbitrary reference R :

$$(\text{id}_R \otimes \mathcal{E})(\rho_{RA}) = \rho_R \otimes \rho_* \quad \text{for every } \rho_{RA}.$$

Taking product inputs first shows that every normalized input has the same output ρ_* . Linearity then gives the unique channel on all operators,

$$\boxed{\mathcal{E}_*(X) = \rho_* \text{Tr } X}.$$

Every traceless perturbation is annihilated after one update. As a superoperator, $\mathcal{E}_*$ has one stationary direction and zero on the traceless operator subspace. The closure ensemble fixes $\langle b | \rho_* | b \rangle = p_*(b)$ in the boundary basis. H.9 selects the diagonal maximum-entropy completion for the decorated transfer action; a different coherent GFT state would define a different microscopic completion and must be tested against that vertex.

Reversible history export. Local replacement is compatible with global unitarity by Stinespring dilation [57]. Introduce the old present register A , a fresh register F , its purifier R , and a history register H . For

$$|\Psi_*\rangle_{FR} = \sum_{b \in \mathcal{B}} \sqrt{p_*(b)} |b\rangle_F |b\rangle_R$$

in the diagonal closure-state realization, a register permutation can act as

$$|\psi\rangle_A |\Psi_*\rangle_{FR} |0\rangle_H \mapsto |\Psi_*\rangle_{AR} |0\rangle_F |\psi\rangle_H.$$

Tracing out R, F, H leaves ρ_* on the new present and removes every dependence on the old local state. The inverse unitary recovers that state from H , so no global information has been destroyed. Postulate III supplies the interpretation of the exported register as history structure; the dilation theorem alone calls it an environment. Repeating the construction indefinitely requires fresh and history capacity or a controlled recycling mechanism, which the present paper has not derived.

The tetrahedral circuit. For

$$\mathcal{B} = \{(s; m_1, m_2, m_3, m_4) : s = \pm, m_i \in \{-3, \dots, 3\}, m_i \neq m_j\}, \quad |\mathcal{B}| = 1680,$$

the register permutation above is one circuit layer if the complete tetrahedron is a native local gate and the fresh admissible register has already been prepared. Neither injectivity nor the collective function $K^2(m_1, m_2, m_3, m_4)$ then requires a sequence of face repairs; the gate replaces one admissible four-face object by another.

A depth-one layer of disjoint two-face gates cannot test all six injectivity relations. The direct pair-comparison construction uses the three perfect matchings

$$\{(1, 2), (3, 4)\}, \quad \{(1, 3), (2, 4)\}, \quad \{(1, 4), (2, 3)\},$$

and therefore has depth three. This proves the stated depth for the direct comparator and excludes depth one in a pairwise architecture without a jointly coupled ancilla. It is not a lower-bound proof for every possible ancilla-assisted state-preparation algorithm. H.9 takes the complete tetrahedron as the native decorated gate; a geometric GFT embedding must realize that locality and sustain the prepared state dynamically.

What closes downstream. The foundational faithful full-support condition closes the classical replacement kernel when applied to histories. Detailed-balance symmetrization then closes $P = |\sqrt{p}\rangle\langle\sqrt{p}|$ and $P_\perp = I - P$ in probability space. Reference decoupling closes the quantum replacement-channel form, and the native tetrahedral architecture supplies a one-layer reversible dilation. H.9 adds the marked recurrence, after which positivity and the electron anchor close the clock, length, and phase-frequency conversion. These results remove an arbitrary transition matrix, geometric clock factor, and separate winding assumption while showing how local memorylessness can coexist with a globally retained past.

They do not by themselves close the coherent GFT Hessian. The exclusion sufficiency lemma of C.5 proves that one scalar statistic contains all local classical information distinguishing the canonical defect from the vacuum. The replacement process removes orthogonal probability memory. A microscopic spectrum must still show that no additional conserved or coherent light field survives and that the source couples only to the capacity singlet. The bridge from $P_\perp$ to a Euclidean mass or Schwinger–Keldysh relaxation pole therefore retains its conditional grade. The stationary reset also leaves the macroscopic arrow-of-time problem in G.7 unchanged.

H.9 Decorated simplicial marked-transfer vertex

The preceding subsections fix the renewed marginal, the lightest seven-channel branch, the determinant recurrence, and the reversible present/history exchange. This subsection gives the finite charged vertex that connects them. It is a decorated fixed-spin transfer vertex: its marked sector is fully specified, while the geometric gluing tensor and stable condensate remain part of the broader GFT embedding problem. The construction introduces no new foundational postulate. It is the minimal action-level realization of faithful full-support history resolution on the native tetrahedron, with the already-established one-bit fermionic defect supplying the hard-core mark. Specifying a vertex is still physical content: the field space and routing below can be replaced by definite countervertices, and those replacements fail the finite audits.

Fresh-state amplitude from the closure action. Let $C_a(b)$ be the three-component closure vector, with $K^2(b) = C_a(b)C_a(b)$. The diagonal renewed state has amplitude

$$A_*(b) = Z^{-1/2} \exp\left[-\frac{\eta_*}{2} K^2(b)\right], \quad |A_*(b)|^2 = p_*(b).$$

For a standard three-component Gaussian auxiliary ξ_a ,

$$\exp\left[-\frac{\eta_*}{2} C^2\right] = \int \frac{d^3\xi}{(2\pi)^{3/2}} \exp\left[-\frac{1}{2}\xi^2 + i\sqrt{\eta_*}\xi_a C_a\right].$$

The amplitude-level closure incidence is therefore $\sqrt{\eta_*}$. The minimal charged decoration reuses this strand, so it introduces no second response coupling. Faithful full-support resolution is being applied on a marked record space with no further joint constraint. Adding a defect-only invariant would define a nonminimal countervortex; it is allowed as a different theory and is one of the falsifiers of the construction, not an extra premise needed by the selected action.

The reversible update has independent present and history Gaussian strands. Their first normalized Hermite excitations obey

$$\langle a|c\rangle_P = \delta_{ac}, \quad \langle b|d\rangle_H = \delta_{bd}.$$

The ordered products $|a, b\rangle = |a\rangle_P \otimes |b\rangle_H$ form nine orthonormal states,

$$\langle a, b|c, d\rangle = \delta_{ac}\delta_{bd}.$$

They carry $\mathbb{R}^3 \otimes \mathbb{R}^3 = \mathbf{1} \oplus \mathbf{3} \oplus \mathbf{5}$. Introduce hard-core marked states $d_{ab}^\dagger|0\rangle$ and the stranded vertex

$$V_Q = \sum_{a,b=1}^3 d_{ab}^\dagger \xi_a^P \xi_b^H + \text{h.c.}$$

Its reduced Gram matrix is $I_3 \otimes I_3 = I_9$. Equal marked weights follow from this two-strand contraction, not from rotational symmetry or the number nine alone. Trace and epsilon contractions would split the three rotation sectors and falsify the vertex.

Pair contraction and canonical dilation. For each unordered occupied pair $e = (m, m')$, the two directed scalar returns are orthogonal alternatives. Define

$$R_e = q(\langle m \rightarrow m' | + \langle m' \rightarrow m |), \quad q = \frac{2}{7},$$

and $B_e = \sqrt{\eta_*} R_e$. Then

$$B_e B_e^\dagger = 2\eta_* q^2 = \frac{8\eta_*}{49} \equiv u, \quad 0 < u < 1.$$

The canonical unitary dilation of this contraction is

$$U_e = \begin{pmatrix} \sqrt{I - B_e^\dagger B_e} & -B_e^\dagger \\ B_e & \sqrt{I - B_e B_e^\dagger} \end{pmatrix}.$$

On the occupied scalar-event support its no-event amplitude is

$$\alpha = \sqrt{1 - u} = \sqrt{1 - \frac{8\eta_*}{49}}.$$

This identifies the origin of every entry: η_* comes from the same closure amplitude that prepares the vacuum, $2q^2$ is the norm of the directed return pair, and the square root follows from reversible dilation.

For an auxiliary Grassmann representation, introduce $\Gamma_e = (\gamma_e^P, \gamma_e^H)^T$ and $J = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$. With $n_m = c_m^\dagger c_m$ and $n_d = \sum_{ab} d_{ab}^\dagger d_{ab}$, the marked action is

$$S_{\text{mark}} = g_{\text{share,eff}} n_d + \Lambda_d n_d (n_d - 1) + \frac{1}{2} \sum_{m < m'} \Gamma_{mm'}^T [1 + n_d n_m n_{m'} (\alpha - 1)] J \Gamma_{mm'}, \quad \Lambda_d \rightarrow +\infty.$$

The quadratic coefficient is $\alpha = \sqrt{1 - u}$, not the probability $1 - u$. The Grassmann integral returns a Pfaffian amplitude, so inserting $1 - u$ in each 2×2 Majorana block would count the no-event probability twice and produce $(1 - u)^{21}$ instead of the required amplitude $(1 - u)^{21/2}$.

The Majorana variables represent the antisymmetric present/history amplitude; they are not additional propagating particles. When the marked fiber or either incident channel is absent, the reference block is J and has Pfaffian one. On the established $n_m = 1$ seven-channel branch, all $\binom{7}{2} = 21$ blocks become αJ , so

$$Z_{\text{edge}} = \alpha^{21} = \left(1 - \frac{8\eta_*}{49}\right)^{21/2}.$$

These event records do not belong to the physical composite space $\Lambda^2 \mathbb{C}^7$. The natural rank-one singlet lift on that composite space has the exact $6 + 15$ split derived in the audit; it does not act on the separately stranded environment records used here.

This distinction is why all 21 pair records survive even though the scalar source couples through one rank-one singlet. The singlet projector acts on physical channel amplitudes. The 21 Majorana blocks label mutually distinguishable present/history event records, so quotienting them by the physical singlet would identify different records and violate faithful resolution.

Complete decorated vertex. Let $\mathcal{G}_v$ denote the geometric fixed-spin gluing tensor and W_* the native update that prepares $A_*(b)$ and exports the old cell. The incidence V_Q above identifies the nine marked basis vectors with the ordered Gaussian response products. The transfer space also contains the unmarked vacuum. Its projector is

$$\Pi_Q = |0\rangle\langle 0| + \sum_{a,b=1}^3 |a, b\rangle\langle a, b|.$$

For $e = (m, m')$, define

$$P_e = n_d n_m n_{m'}, \quad \tilde{U}_e = (I - P_e) + P_e U_e.$$

The complete one-step decorated transfer is

$$\mathcal{A}_v = \mathcal{G}_v W_* \Pi_Q \exp[-g_{\text{share,eff}} n_d - \Lambda_d n_d (n_d - 1)] \prod_{m < m'} \tilde{U}_{mm'}.$$

The empty state therefore has unit weight, while each of the nine one-mark states has fugacity $e^{-g_{\text{share,eff}}}$. The factors $\tilde{U}_e$ act on distinct edge-history strands and commute. The Grassmann action above is the Pfaffian trace representation of these controlled pair factors; it is not an additional operator multiplied into $\mathcal{A}_v$. Faithful full-support resolution on this record space, which has no further joint edge constraint, selects the tensor product of its one-edge marginals. Tracing the hard-core one-particle response fiber gives

$$\zeta_* = 9e^{-g_{\text{share,eff}}} \left(1 - \frac{8\eta_*}{49}\right)^{21/2} = 0.005123584484947.$$

Complete edge-Hessian audit. On the 21-dimensional edge-label space, two distinct edges either share one endpoint or are disjoint. Let A_1 and A_0 be their adjacency matrices. The Johnson edge algebra has spectra

$$\begin{aligned} \text{spec}(A_1) &= \{-2^{\times 14}, 3^{\times 6}, 10^{\times 1}\}, \\ \text{spec}(A_0) &= \{-4^{\times 6}, 1^{\times 14}, 10^{\times 1}\}. \end{aligned}$$

A generic permutation-covariant edge kernel can therefore carry three eigenvalues. The conditioned vertex instead gives

$$\mathcal{K}_{\text{edge}}^{(2)} = I_{21} \otimes \alpha J, \quad \text{spec}_{\text{edge}} = \{\alpha^{\times 21}\},$$

with zero off-diagonal edge norm. Coherently averaging the induced edge permutation matrix over all $7! = 5040$ channel relabelings gives $\mathbf{1}\mathbf{1}^T/21$, of rank one. That operation would collapse the result to a shared singlet. It is not the charged gluing used here: $m = -3, \dots, 3$ are physical spin weights, K^2 is not invariant under arbitrary relabeling, and the defect Hamiltonian contains diagonal occupations and density interactions rather than hopping terms $c_m^\dagger c_{m'}$.

Primitive routing enumeration. Write $w = e^{-g_{\text{share,eff}}} \alpha^{21}$, so $\zeta_* = 9w$. The one-step transfer block contains three finite state spaces. The universal sector has one empty state and nine marked states, hence

$$Z_\mu = 1 + 9w = 1 + \zeta_*.$$

The electron feedback sector has one empty state and states $|m; I_L, I_R\rangle$, with seven persistent charged labels and two independently renewed nine-state endpoints. Its dimension is

$$1 + 7 \times 9 \times 9 = 568,$$

and its trace is

$$Z_{e,\text{return}} = 1 + 7 \times 81w^2 = 1 + 7\zeta_*^2.$$

The second-shell sector has one empty state and nine marked states with one scalar passage $q = 2/7$, giving

$$Z_{\tau,2} = 1 + 9qw = 1 + \frac{2}{7}\zeta_*.$$

The direct product contains $10 \times 568 \times 10 = 56,800$ configurations. Exhaustive enumeration reproduces the analytic product exactly. The finite routing factors are therefore

$$Z_\mu = 1 + \zeta_*, \quad Z_e = (1 + \zeta_*)(1 + 7\zeta_*^2), \quad Z_{\tau,2} = 1 + \frac{2}{7}\zeta_*.$$

One native update applies each controlled gate once. Hard-core unit charge restricts the marked return to vacuum plus one return motif; the charged label persists diagonally while the response endpoints renew; and the shell grading terminates after $N = 2$. These restrictions implement structures already present elsewhere in the paper: one native renewal pass, the one-bit fermionic anchor, the physical spin-weight labels, and the three-shell closure ladder. They add no separate foundational premise. Relaxing them gives explicit controls. A shared endpoint polarization changes the electron factor to $1 + 7\zeta_*^2/9$ and moves G to -14.6σ ; off-diagonal charged propagation gives $1 + 49\zeta_*^2$ and moves it to $+98.1\sigma$.

Predictions and scope. The physical scale and charged-lepton ratios are

$$L_* = Z_e \left[-\frac{3}{2} \lambda_e \ln(1 - e^{-7g_{\text{share,eff}}}) \right], \quad G_* = \frac{c^3 L_*^2}{\hbar},$$

$$\frac{m_\mu}{m_e} = 720 \frac{2}{7} Z_\mu, \quad \frac{m_\tau}{m_e} = 720^2 \left(\frac{2}{7} \right)^4 Z_\mu Z_{\tau,2}.$$

They evaluate to $G_* = 6.6742890772 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}$, $m_\mu/m_e = 206.768280237$, and $m_\tau/m_e = 3477.343310$. The residuals are -0.073σ , -0.535σ , and $+0.481\sigma$.

The finite marked-transfer action is now specified and audited. Its generic countermodels remain useful falsifiers: a collective edge mode, a physical pair-composite interpretation, persistent endpoint polarization, additional defect-only coupling, or off-diagonal charged propagator changes the answer. The remaining microscopic task is the geometric embedding: $\mathcal{G}_v$ must realize this oriented decoration, admit the required condensate, and produce no additional light source-coupled mode. The action does not solve durable history storage, the thermodynamic arrow, or the infrared transverse influence functional.

Reproduction. The supplementary scripts reproduce the ensemble invariants, state-weighted determinant, 48×35 locality fracture, Gaussian closure amplitude, unitary pair dilation, nine-state Gram matrix, 21-block edge Hessian, and 56,800-state routing trace. They also verify the dressed scale and corrected charged-lepton ratios of Appendix I.1. Appendix O contains the self-contained numerical-spine script; the complete Hessian and graph enumerations remain in the accompanying audit files.

Appendix I: Mass and Gauge Extensions

Appendix I collects sectors that are structurally connected to the same entanglement logic but are not part of the closed weak-field core. They are kept here because they show how the framework may extend, not because the main derivation depends on them.

I.1 Charged-lepton spectrum from the shell algebra

Before the algebra, the physical picture. The electron is the ground-state one-bit charged defect — the $N = 0$ configuration whose seven-channel dressing sets the substrate length (Appendix D.4). The muon and tau are heavier shell excitations of the same defect. Each adds one radial entanglement shell, realized microscopically as the exclusion of one face label from the seven-state alphabet of Appendix B. The closure machinery loses its nondegenerate weighting after the third branch, so the construction contains three charged leptons. Their baseline masses follow from the shell entropy and singlet overlaps; the decorated marked-transfer vertex supplies the finite response that earlier versions left in the residuals. Because the ratios carry no dimensionful input, they test the same marked action that enters G_* .

With that picture fixed, the spectrum is written as a shell expansion,

$$\log \frac{m_N}{m_e} = B_0 N + A_0 N^2, \quad N = 0, 1, 2,$$

with the electron the $N = 0$ ground state and the muon and tau the $N = 1, 2$ radial entanglement-shell excitations of the same core structure. Each shell excitation is the exclusion of one face label from the seven-state alphabet of the boundary ensemble in Appendix B.

Three-generation termination from closure-spectrum collapse. Combinatorial injectivity alone allows up to $N = 3$ shells, since the reduced alphabet must retain at least four distinct labels to support an injective face assignment ($7 - N \geq 4$). The sharper structural bound comes from the admissibility-closure machinery. For the reduced ensemble at shell N , the closure invariant K^2 inherits a discrete spectrum whose dispersion gives the admissibility kernel $e^{-\eta K^2}$ real weight. For $N = 0, 1, 2$ the K^2 spectrum has multiple distinct values, and the closure equation $\langle K^2 \rangle_\eta = 3/(2\eta)$ has a non-degenerate solution.

At $N = 3$ the spectrum collapses. With only four labels left, every injective assignment to the four faces uses all of them, and K^2 depends on the labels only through the permutation-invariant sums S and Σ^2 , so every microstate of the reduced four-label ensemble carries the same value of K^2 . With a single spectral value K_0^2 the closure equation $\langle K^2 \rangle_\eta = 3/(2\eta)$ still fixes $\eta = 3/(2K_0^2)$ formally; what fails at $N = 3$ is selection, not solvability. The kernel $e^{-\eta K^2}$ acts trivially when every microstate carries the same K^2 , so it weights nothing and leaves no admissibility dispersion to define a branch. Shell $N = 3$ therefore does not define a non-degenerate admissibility-closed branch in the same sense as the earlier shells: the cutoff rests on the collapse of the K^2 dispersion, not on the closure equation admitting arbitrary η .

The number of charged-lepton generations is consequently the number of shells with a non-degenerate closure branch,

$$N \in \{0, 1, 2\} \quad \implies \quad \text{three generations,}$$

a structural bound sharper than the combinatorial one. It is this mechanism, rather than injectivity alone, that fixes the generation count at three rather than four.

Linear coefficient from the first-shell reduced-alphabet entropy. The linear coefficient is the entropy cost of adding one shell excitation. With one face label excluded, the reduced ensemble has

$$\Omega_1 = 2 P(6, 4) = 720 = 6!,$$

i.e., the full symmetric group S_6 with orientation degeneracy. Direct admissibility evaluation on the reduced ensemble gives $g_1 = \ln \Omega_1 = \ln 720$ to within 0.1%, since the admissibility correction is small on the reduced ensemble where the alphabet symmetry is nearly intact. Hence

$$B_0 = \ln \Omega_1 = \ln 720 = 6.579 \dots,$$

compared with the empirical fit $B_0 = 6.586$ (agreement at 0.1%).

Why the same 720 at every shell. The per-shell factor is the *same* $\Omega_1 = 720$ for each added shell, not a decreasing sequence $720 \cdot 240 \dots$ in which each shell excludes a further label. The two readings diverge sharply at the tau: cumulative exclusion would give $m_\tau/m_e \simeq 1151$ (−67%), whereas the repeated factor gives 3454.6 (−0.65%), so the data require the same reduction at each shell. The structural reason is the rank of the inter-shell coupling. The shells couple only through the permutation-invariant singlet projector P_{sing} , whose spectrum is $\{1, 0, 0, 0, 0, 0\}$: rank one. A rank-one projector pins exactly one shared direction, and it is the same direction however

many shells are already stacked, so each added shell surrenders only that one shared singlet direction and retains the other six — the same direction at every shell, not a fresh label removed cumulatively. The injective count realizes this exactly: $P(6,4)/P(7,4) = 360/840 = 3/7$, so $1680 \rightarrow 720$ once and again at each subsequent shell. This also separates the two reduced-alphabet notions that otherwise look inconsistent. The *generation cutoff* above counts $7 - N$ surviving labels because it asks whether a non-degenerate closure *branch* still exists; the *mass factor* here is the rank-one singlet pinning one shared direction per shell. They are different operations, which is why the cutoff scales with $7 - N$ while the ladder carries a fixed 720.

Quadratic coefficient from singlet-projection shell algebra. The quadratic coefficient follows from three premises already in the framework: (i) mass arises from scalar capacity response (Postulate II); (ii) the scalar EFT response is quadratic in the total source; (iii) the coarse scalar branch projects onto the permutation-invariant singlet of the seven-state face alphabet.

Let

$$\mathcal{H}_7 = \text{span}\{|m\rangle : m = -3, \dots, 3\}, \quad |u\rangle = \frac{1}{\sqrt{7}} \sum_{m=-3}^3 |m\rangle, \quad P_{\text{sing}} = |u\rangle\langle u|.$$

A shell excitation that marks the face state $|a_r\rangle$ contributes to a coherent shell source

$$|J_N\rangle = \sum_{r=1}^N |a_r\rangle.$$

The coarse scalar field δS is permutation-symmetric on face labels and therefore couples only through P_{sing} . The scalar-channel response is quadratic in the total source, so the relevant object is the source norm

$$\langle J_N | P_{\text{sing}} | J_N \rangle = \sum_{r,s=1}^N \langle a_r | P_{\text{sing}} | a_s \rangle.$$

Because $\langle a | P_{\text{sing}} | b \rangle = 1/7$ for any a, b , this reduces to $N^2/7$ independent of which specific face labels are excluded. Including the binary orientation degeneracy of the boundary ensemble, the orientation-summed scalar transfer weight per ordered shell incidence is

$$\mathcal{M}_{rs} = 2 \langle a_r | P_{\text{sing}} | a_s \rangle = \frac{2}{7}.$$

The ordered-pair count decomposes as $N^2 = N + 2\binom{N}{2}$: N self-incidences along the diagonal and $2\binom{N}{2}$ directed cross-incidences. Two distinct steps enter here, and only the first is the quadratic EFT response. The quadratic singlet response fixes that there are N^2 ordered incidences, each carrying the per-incidence overlap $\mathcal{M}_{rs} = 2/7$; it does not by itself make the mass multiplicative. The passage from these N^2 additive incidences to a multiplicative factor is the log-overlap mass map: the mass is the exponential of a dressing free energy whose pairwise term is $\sum_{r,s} \ln \mathcal{M}_{rs} = N^2 \ln(2/7)$, the same log-overlap form that carries the gravitational dressing (Appendix H.3) and stated explicitly below. With that map, the N -shell dressing is

$$\prod_{r,s=1}^N \mathcal{M}_{rs} = \left(\frac{2}{7}\right)^{N^2},$$

so

$$A_0 = \ln \frac{2}{7} = -1.253\dots,$$

compared with the empirical fit $A_0 = -1.255$ (agreement at 0.15%).

The ordered-bilinear structure is not a free assumption. It is the unique scalar-channel response to a multi-source state under the three premises just stated: the source norm $\langle J_N | P_{\text{sing}} | J_N \rangle$ runs over both indices independently because the integrated-out scalar contribution is quadratic in J_N .

Mass ladder and numerical accuracy. Combining the two coefficients,

$$\boxed{\frac{m_N^{(0)}}{m_e} = 720^N \left(\frac{2}{7}\right)^{N^2}, \quad N = 0, 1, 2.}$$

Equivalently,

$$\log \frac{m_N}{m_{N-1}} = \ln 720 + (2N - 1) \ln \frac{2}{7},$$

so each added shell contributes $\ln 720$ plus an odd number of new scalar-overlap incidences $(1, 3, 5, \dots, \text{summing to } N^2)$. The second difference is fixed at

$$\Delta^2 \log m_N = 2 \ln \frac{2}{7} = -2.506 \dots,$$

compared with the empirical -2.509 (0.15%). The marked vertex of H.9 then supplies

$$Z_\mu = 1 + \zeta_*, \quad Z_{\tau,2} = 1 + \frac{2}{7}\zeta_*.$$

Thus

$$\begin{aligned} \frac{m_\mu}{m_e} &= 720 \frac{2}{7} Z_\mu = 206.768280237 \quad (-0.535\sigma), \\ \frac{m_\tau}{m_e} &= 720^2 \left(\frac{2}{7}\right)^4 Z_\mu Z_{\tau,2} = 3477.343310 \quad (+0.481\sigma). \end{aligned}$$

The dressing map, explicitly. The three factors follow from one construction. Write the N -th charged lepton as a coherent source raised on the electron ground dressing by exciting N shells, $\mathcal{S}_N = \sum_{k=1}^N s_k$, where each s_k excludes one face label from the seven-state alphabet. Its dressing free energy is a log-overlap functional of the kind used for the gravitational dressing in Appendix H.3, and it separates into a per-shell term and a pairwise term. Each shell carries the reduced-alphabet entropy $\ln \Omega_1 = \ln 6! = \ln 720$, the permutation count of the six surviving labels, which builds the linear factor 720^N . Each ordered pair of shells contributes a singlet log-overlap: a single label projected on the closure singlet $u = \frac{1}{\sqrt{7}}(1, \dots, 1)$ of Section 8 has weight $|\langle u | m \rangle|^2 = 1/7$, and orientation summation doubles it to $2/7$. The N shells form N^2 ordered pairs, so the pairwise term is $N^2 \ln(2/7)$, and exponentiating the free energy returns

$$\frac{m_N^{(0)}}{m_e} = \Omega_1^N \left(\frac{2}{7}\right)^{N^2} = 720^N \left(\frac{2}{7}\right)^{N^2}.$$

The coefficients are fixed: $720 = 6!$, the N^2 pair count is combinatorial, and $2/7$ is the orientation-summed singlet projection behind $\Omega_{\text{tet}} = 2P(7, 4) = 1680$. The marked factors are finite traces of the same decorated vertex used in the electron scale.

The support exponent, tested against the spectrum. The ladder also tests the exponent in the support-to-length identification of Appendix D.4, which reads the physical ratio as the first power of the entropy support, $\lambda/L_* \propto e^H$. Letting that power float, $\lambda/L_* \propto e^{\alpha H}$, turns each lepton ratio into a measurement of α , since $m_N/m_e = [720^N (2/7)^{N^2}]^\alpha$:

$$\alpha_N = \frac{\ln(m_N/m_e)_{\text{PDG}}}{\ln[720^N (2/7)^{N^2}]}, \quad \alpha_\mu = 1.0010, \quad \alpha_\tau = 1.0008.$$

The leading shell values give $\alpha_\mu = 1.0010$ and $\alpha_\tau = 1.0008$, while square-root and quadratic support maps miss by orders of magnitude. The decorated action does not float this exponent. It keeps the first-power support map and computes the finite response factors separately.

Status. The three-generation termination, baseline shell coefficients, and marked response are fixed inside the displayed finite action. The collapse at $N = 3$ is the permutation-invariance of K^2 ; 720, $2/7$, and N^2 are reduced-alphabet, singlet-projection, and ordered-pair counts; and Z_μ , $Z_{\tau,2}$ are the enumerated marked traces. The remaining particle-sector work concerns other defect classes, not an unfixed charged-lepton coefficient.

The electron anchor remains the clean weak-field entry point; the heavier charged leptons are correlated outputs of the same marked shell algebra, not closure-defining inputs. Composite hadrons remain part of the dressed bound-state entropy program. The neutrino and quark sectors are not treated here. Extending the closure-spectrum-collapse argument to those defect classes remains a separate task.

I.2 Gauge-structure extension and Standard Model ontology

The framework’s primary target is the gravitational and dark sector. Standard Model gauge structure is hosted on the substrate as external content rather than derived from it. This subsection specifies which gauge components have substrate-natural analogues, which are substrate-compatible but external, and how the baseline-redundancy template accommodates the full Standard Model gauge group $SU(3)_c \times SU(2)_L \times U(1)_Y$.

Baseline-redundancy template. The same logic that organizes the gravity sector extends to any conserved-charge sector. For a conserved charge Q , introduce an entropy-like potential $S_Q(x)$ and require that physical observables depend only on differences of that potential, not on its absolute baseline. Localizing that redundancy requires a compensating connection. In the Abelian case,

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

giving Maxwell-type dynamics for A_μ . The non-Abelian generalization, for multiplet-valued entropic potentials transforming under a compact Lie group G , gives Yang–Mills covariant derivatives and field strengths in the standard form.

The template tells us what form a gauge interaction takes once a conserved-charge sector with baseline redundancy is present. It does not select which specific groups are physically realized.

Substrate-natural components. Two ingredients of Standard Model gauge structure are naturally accommodated by the substrate.

$U(1)_Y$. The baseline-redundancy template naturally accommodates an Abelian conserved-charge sector of the hypercharge type. The associated entropy potential is conserved-charge in character, and localizing its baseline redundancy produces the corresponding Maxwell-type dynamics. The argument is independent of the gravity-sector derivation, applying the same template to a separate conserved-charge potential.

$SU(2)_L$. The substrate naturally carries an $SU(2)$ -representation structure through its half-integer face data. The spin-3/2 commitment forced by maximum-capacity channel selection and tetrahedral injectivity (Section 5) carries an $SU(2)$ action on face data through the standard double-cover relation. The static K^2 ensemble breaks this $SU(2)$ to its $U(1)$ Cartan via the magnetic-quantum-number projection, but the full $SU(2)$ is recoverable at the level of the

quantum face data the projection discards. Identifying this representation-theoretic $SU(2)$ with the internal weak group $SU(2)_L$, including chirality and doublet assignments, remains external to the present derivation.

Substrate-compatible but external: $SU(3)_c$. Color $SU(3)$ does not emerge from the substrate by any of the standard mechanisms (direct decomposition of the seven-state face algebra, Seiberg duality, preon compositeness, topological soliton, monopole condensation). The obstructions are structural: $SU(3)$ is not a spin group, so it does not inherit from rotational covariance; its fundamental representation requires three equivalent (non-hierarchical) states, while the substrate’s natural three-fold structures (shell index, K^2 classes, residual face positions) are all hierarchical; its rank-2 Lie algebra requires two independent quantum-number axes, while the substrate provides only rank-1 axis structures; and its non-Abelian commutation relations cannot be reproduced from the substrate’s scalar edge couplings.

A positive ontological reading is nonetheless available within the framework. Color is the substrate’s name for an internal three-valued label on fermionic defects that lives in the non-singlet complement of the face-state algebra. The scalar gravitational branch couples to defect sources only through the permutation-invariant singlet (Appendix D), so the macroscopic scalar field is blind to non-singlet structure by construction. Color labels and color dynamics therefore live in a sector that gravity cannot resolve: the gravitational sector and the color sector occupy orthogonal subspaces of the substrate’s algebra. The strong interaction can be represented, within this hosting picture, as the local gauge dynamics of the color label under the general baseline-redundancy template. The choice of $SU(3)$ specifically is empirical input; it is the realized non-Abelian structure that fits the template.

This account explains several structural features of color without claiming to derive them. Gravity’s blindness to color follows from the scalar branch being the singlet projection of the substrate. Color confinement itself is not derived here; the substrate account provides that the macroscopic scalar-gravity branch resolves only singlet structure and is therefore blind to color non-singlets, which is consistent with the observed fact that hadronic macroscopic signatures are color-singlet but is not a replacement for the QCD confinement mechanism. The independence of color from generation, charge, and spin reflects the independence of the corresponding substrate sectors: color in the non-singlet complement of the face algebra, generation in the shell-excitation index (Appendix I.1), electromagnetic charge in the baseline redundancy of a separate conserved potential, and spin in the fermionic face data.

A sharper reading: the transverse complement as a color arena. The obstructions above isolate to a single contingent ingredient, which turns “ $SU(3)$ is external” into a definite derivation target. Write the face-state space as

$$\mathcal{H}_7 = \langle s \rangle \oplus \mathcal{H}_\perp, \quad |s\rangle = \frac{1}{\sqrt{7}}(1, 1, 1, 1, 1, 1, 1), \quad \dim_{\mathbb{R}} \mathcal{H}_\perp = 6,$$

the same singlet/non-singlet split that carries the scalar return operator and the lepton ladder. Suppose the six real transverse directions carry a complex structure: a preferred pairing into three complex modes, $\mathcal{H}_\perp \simeq \mathbb{C}^3$. The local frame redundancy of a three-complex-dimensional space is $U(3)$, and its overall phase is the condensate $U(1)$ already in use, so the traceless remainder is $SU(3)$. The rank-two algebra and eight generators come from the frame redundancy of the complex space itself, not from substrate quantum-number axes, so this route is not blocked by the rank obstruction above. The representation content then follows with no further input: a single defect is a vector in the $\mathbf{3}$, hence non-singlet, and the scalar branch resolves only singlets, so lone defects are not asymptotic; $\mathbf{3} \otimes \bar{\mathbf{3}} = \mathbf{1} \oplus \mathbf{8}$ supplies the meson’s singlet, and the antisymmetric ϵ_{abc} in $\mathbf{3} \otimes \mathbf{3} \otimes \mathbf{3}$ supplies the baryon’s. Three color directions and their conjugates, eight gluons,

the absence of free single defects from the singlet spectrum, and gravity’s blindness to color all read off one split.

The missing ingredient, and a route ruled out. The skeleton rests entirely on the complex structure $6_{\mathbb{R}} \rightarrow 3_{\mathbb{C}}$, which the construction does not supply. Without a preferred pairing of the six transverse directions into three complex ones, the symmetry of the complement is $SO(6)$, not $SU(3)$, and the pairing must in addition be degenerate, the three directions equivalent — which is the original reason the substrate’s hierarchical three-fold structures (shell index, K^2 classes, residual face positions) do not furnish color. The most concrete candidate is a cyclic readout $T : |m\rangle \mapsto |m+1 \bmod 7\rangle$, whose Fourier modes would pair as conjugates $(k, 7-k)$ into $\mathbf{1} \oplus 3_{\mathbb{C}}$; it does not survive contact with what the labels are. They are not residues modulo seven but the weights $m = -3, \dots, 3$ of the spin-3 multiplet ($j_{\text{eff}} = 3$, Section 5), a real and irreducible representation of the rotation group: it carries no invariant complex structure and does not split off the singlet under rotations at all. The closure invariant uses the integer weights, $K^2 = 48 - \frac{1}{3}(S^2 - \Sigma^2)$ with $S = \sum_i m_i$ and $\Sigma^2 = \sum_i m_i^2$, so the cyclic shift is not a symmetry of the dynamics: shifting $(-3, -2, -1, 0)$ by one carries $(S, \Sigma^2) = (-6, 14)$ to $(-2, 6)$. A cyclic $\mathbb{Z}_7$ would deliver $\mathbf{1} \oplus \mathbf{3} \oplus \bar{\mathbf{3}}$, but it is not the structure the seven labels carry. The status is therefore the one above and no stronger: color has a natural location in the non-singlet complement the scalar branch cannot resolve, but the operator that would make that complement $\mathbb{C}^3$ is absent, and the cyclic route to it is excluded by the spin-weight structure. A derivation of $SU(3)_c$ would need an additional readout symmetry independent of the rotational dynamics, not presently exhibited; the skeleton reaches none of quark flavor, fractional charge, the QCD scale, confinement, or any hadron mass.

A better arena: tetrahedral routing. A more natural three is available, not from the seven labels but from the cell. Four boundary faces admit exactly three pairings,

$$A = (12 | 34), \quad B = (13 | 24), \quad C = (14 | 23),$$

the opposite-edge routings of the tetrahedron, and the tetrahedral symmetry group permutes them as the full S_3 , so the three are symmetry-equivalent with no channel singled out — the property the cyclic route could not get from the ordered spin labels. Read as quantum amplitudes the three routings form a complex vector $q \in \mathbb{C}^3$; if the local routing basis is a gauge choice, the comparison between neighbouring cells is a link $U_{yx} \in U(3)$, and removing the common phase already carried by the charged/condensate $U(1)$ leaves $SU(3)$ with eight connection modes. Quarks are open routing vectors q^a , antiquarks the conjugate, and singlet closure is the standard count — $q^a \bar{q}_a$ for the meson, $\epsilon_{abc} q^a q^b q^c$ for the baryon, the latter using the oriented three-form the cell’s orientation/parity primitive already supplies.

What the closed ensemble gives, and why. The decisive question is whether the admissibility kernel treats the three routings as equivalent and independent, $K_{\mathcal{R}} \propto I_3$. It does not. For the natural pair observable $R_A = m_1 m_2 + m_3 m_4$ and its images R_B, R_C , the kernel over the admissibility-weighted 1680-state ensemble is symmetric with nonzero off-diagonal entries, so its spectrum is one isolated value and a degenerate pair: the routing triplet splits as $\mathbf{3} = \mathbf{1} \oplus \mathbf{2}$, not an irreducible color triplet. The reason is exact,

$$R_A + R_B + R_C = \sum_{i < j} m_i m_j = \frac{1}{2}(S^2 - \Sigma^2) = 72 - \frac{3}{2}K^2,$$

so the symmetric routing mode is an affine function of the closure invariant K^2 : it is the scalar sector itself, in routing variables. In plain terms, the three routings a closed cell offers are not independent, because their average is just the cell’s overall closure health — the quantity

gravity already reads — so only the two ways the routings differ from that average are free. The closed cell offers only two genuinely transverse directions in the three-dimensional routing space (distinct from the six-dimensional label complement above), the third routing direction already committed to gravity, and a color $\mathbb{C}^3$ cannot be a closed-cell observable.

The open-flux direction. The split locates the missing principle rather than closing the route. The closure tie forcing $R_A + R_B + R_C \propto K^2$ is a property of closed scalar observables; an open routing imbalance carries no such constraint, and its three routes are independent local amplitudes. The natural object is then a lattice gauge action,

$$S_{\text{open}} = \sum_{\langle xy \rangle} \|q_y - U_{yx} q_x\|^2, \quad q_x \in \mathbb{C}^3, \quad U_{yx} \in SU(3),$$

which exists only if open non-singlet routing flux is parallel-transported between cells rather than memorylessly refreshed. That added dynamical principle would release the three routings from the closure tie and make them a local $SU(3)$ frame. It is a research target rather than a result: the tetrahedral cell supplies three equivalent routing channels and an orientation form, while the required open-flux transport law is absent from the present construction.

Hadronic matter and the mass–entropy bridge. The mass–entropy bridge $m = \kappa_m(\ell) \Delta S$ applies to fermionic defects regardless of their color label. For charged leptons, the defect is color-singlet and the marked shell action of Appendix I.1 gives both measured ratios within one standard deviation. For hadrons, the dressed entropy is dominated by the QCD-internal contributions of confinement-scale gluonic flux, trace-anomaly structure, chiral vacuum reorganization, and quark binding. The mass–entropy bridge then organizes the full dressed bound-state entropy budget,

$$S_{\text{ent,H}}^{\text{dressed}} = S_{\text{defect}} + S_{\text{bind}} + S_{\text{conf}} + S_{\chi\text{SB}},$$

without requiring the framework to re-derive QCD. The structural claim is compatibility: the dressed entropy budget is the quantity QCD calculates, and the mass–entropy bridge converts it to inertial mass through the same running $\kappa_m(\ell)$ used for elementary sectors.

Scope summary. The Standard Model’s full gauge group $SU(3)_c \times SU(2)_L \times U(1)_Y$ is hosted on the substrate at three distinct levels of derivational support: $U(1)_Y$ from baseline redundancy (substrate-natural via independent argument); $SU(2)_L$ from spin-3/2 fermionic face representation theory (substrate-natural with chirality identification external); $SU(3)_c$ as a non-singlet internal label (substrate-compatible but not substrate-derived). The framework’s target scope remains the gravitational and dark sector, and the gauge sector is hosted accordingly.

The sectors in this appendix are structurally linked to the same entanglement logic but are not part of the controlled ordinary Einstein/capacity branch.

Appendix J: The Lattice Microstructure Program

Appendix J records the technical content of Section 23: the host ensemble and its integrity gates, the coupled measure with its declared modeling forks, the control methodology, the exact cell-order calculations, and the measured tables of the compatibility and defect experiments, together with the specification of the capacity-transport instrument. Where Appendix K recomputes the coefficient chain from its definitions, this appendix couples that chain to a dynamical substrate and reports what the coupled system does. The engine, run configurations, audit scripts, and raw measurement tables are maintained with the manuscript materials at jaigp.org.

J.1 Host ensemble and engine discipline

The host is a standard causal-dynamical-triangulations ensemble [68, 69]: triangulations of $S^1 \times S^3$ with eighty spatial slices, sampled by local Pachner-class moves under the Regge action with bare couplings $k_0 = 2.2$ and k_4 tuned near criticality, and the $(4, 1)$ -simplex number pinned by a quadratic penalty $\varepsilon(N_{41} - \bar{N})^2$ with $\varepsilon = 0.01$. The action is tracked incrementally and audited against full recomputation throughout every run, with drifts held at the 10^{-11} level, and every quoted ensemble carries a terminal verification of the manifold conditions (gluing, links, Euler characteristic, connectivity).

The host’s phase location has been mapped directly. A (k_0, Δ) grid with k_4 auto-tuned to pseudo-criticality at every point — the volume-pin residual serving as the feedback error, as in standard CDT practice where k_4 is retuned per coupling point — finds a coherent region at $k_0 = 1.0$ – 2.0 with tuned $k_4 = 0.67$ – 0.75 , slice-volume coherence peaking at 0.479 at $(k_0, \Delta) = (2.0, 0.4)$, and a dead row by $k_0 = 3.5$: a visible boundary. A fixed $k_4 = 0.9$, which an untuned grid had used throughout, was the obstruction — every untuned cell was fighting the pin rather than sampling its phase. The mapping statistics (400 sweeps per point) do not yet clear the blob criterion, so the region is a *candidate* for the extended phase, with its confirmation run (a volume pair at the peak point) pre-registered; the located boundary is also the geometry the pre-registered curvature-sector displacement test of J.4 straddles.

The foliation is enforced exactly, and the enforcement has a history worth recording. Validating the cell identification of Section 23.1 exposed a defect in the move set as first written: it preserved the four-dimensional manifold conditions while permitting moves that fold a spatial slice, leaving a few percent of slice triangles in violation of the closed-3-manifold condition the identification requires. The audit layer caught the defect; the move set now rejects any proposal that would break the foliation, through a local test that is complete (an affected triangle’s link necessarily passes through a newly created pentachoron) and preserves detailed balance in the same way as the other locality filters; and every ensemble quoted in this paper was generated with the corrected engine and passes a global foliation audit: every slice triangle in exactly two spatial tetrahedra, each pentachoron correctly threaded between its slices.

At these couplings and volumes the bare host has Hausdorff dimension rising from 2.7 to 2.9 to 3.1 across $N_{41} = 10,000, 20,000, 40,000$, with the slice-volume-profile score falling as the volume grows. The host’s position relative to the de Sitter phase established for this ensemble class [68] is an open scan at this coupling point, and no reading below leans on the host being in that phase: every claim is comparative, coupled against bare and real against placebo, at matched volume.

J.2 The coupled measure and its fidelity anchors

The coupled measure is the joint Gibbs distribution of Section 23.2, with the per-face label sum normalized so that each label carries base measure $1/7$. Under that base measure two properties hold exactly: at $\beta = 0$ the label entropy cancels and the geometry marginal is identically the bare host, so the control arm is faithful by construction; and uniform label assignment at geometry moves is exactly detailed-balanced, so no proposal correction is needed for label births. The centering constant μ is the per-cell label free energy,

$$\mu(\beta) = -\frac{1}{\beta N_{\text{cells}}} \ln \mathbb{E}_{m \sim \text{unif}} \left[e^{-\beta \sum_c E_c} \right] = \frac{1}{\beta} \int_0^\beta \langle \bar{E} \rangle_s ds,$$

computed by thermodynamic integration on the thermalized base configuration and logged with each arm, so the extensive part of the closure term is volume-neutral and cannot pose as a shift of k_4 ; matched total volume across arms is nevertheless retained as an independent gate rather than assumed from the centering.

Every run opens by rebuilding the single-cell tables and checking the Part II anchors ($\Omega_{\text{tet}} = 1680$, $\eta_* = 0.0298668$, $\langle K^2 \rangle_{\eta_*} = 50.223$, $g_{\text{share,eff}} = 7.4198$), and refuses to sample if any fails.

Three modeling forks are declared. Injectivity is imposed as a penalty λ per colliding pair ($\lambda = 3$ in the quoted runs) because a hard seven-color constraint at the maximal conflict degree does not guarantee an ergodic single-face heat bath; the hard constraint is the $\lambda \rightarrow \infty$ limit, the residual collision fraction is reported in every table so the softness stays visible, and the theory point is read at the ordered end. The closure sum is evaluated in the unoriented reading, in which the two cells sharing a face see the same label; the oriented alternative, in which they see opposite signs, requires oriented-slice bookkeeping through all moves and is recorded as an open fork. Parity contributes exactly $\ln 2$ per cell, extensive at pinned volume, and drops out of the coupled dynamics.

J.3 Controls and the reading discipline

Four gates precede any reading: the manifold and foliation audits of J.1 on every arm; thermalization of the action and volume over the measured window; matched total simplex number across arms, since volume differences masquerade as dimension differences in the observables of interest; and a live chain, since a frozen chain is a failed coupling and never a result. The interpretation of each possible outcome, including which outcomes would license no claim at all, was written and committed before the corresponding results existed.

The load-bearing control is the placebo arm. The closure energy takes 210 distinct values on the sorted-label orbits of the cell; the placebo shuffles those values across orbits, preserving the value pool and the permutation symmetry while destroying the closure structure, and runs the full pipeline at the top coupling. Only a real-versus-placebo difference at matched volume is ever attributed to the closure weighting; anything the placebo reproduces is machinery.

That discipline was learned on a neighboring theory, and the episode is part of the record. Before the admissibility coupling was built, an EPRL vertex amplitude [70] at frozen boundary spin $j = 3$, computed with the `s12cfoam-next` library [71], was coupled to the same host. A pre-registered audit of that coupling found five distinct misspecifications: a centering constant calibrated off-equilibrium and acting as an uncontrolled shift of the bare cosmological coupling; a label heat bath running at the wrong exponent; label births outside the acceptance ratio; an order-unity dependence on an arbitrary slot convention of the vertex tensor; and the positivization that any sign-alternating amplitude must undergo to be sampled at all. Each vanishes identically at $\beta = 0$ and grows with β , so each masquerades as the signal the sweep was designed to detect while remaining invisible to the control arm. The decisive instrument was an entry-shuffled copy of the vertex tensor: it reproduced the apparent geometric steer at matched coupling, identifying the visible effect as sampler machinery. The sampler was rebuilt as an exact joint Gibbs measure (the fixes are inherited by the closure coupling, as the construction of J.2 shows), and the shuffled-placebo requirement became standing policy. The EPRL exercise is reported here as methodology: its physics scope is in any case bounded by the frozen-spin truncation and the positivization, and the object it samples is a neighboring theory’s weighting, whose seven-dimensional intertwiner space is a different object from the seven face labels of this paper’s ensemble (Section 27.8).

J.4 Exact results at cell order

Six calculations, independent of any Monte Carlo run, fix what the substrate can and cannot do for the vacuum before the coupled measurements are read.

Induced curvature coupling. The geometry marginal of the coupled measure weights a slice by the free energy of its labels, and the leading geometric dependence enters through the coordination q of slice edges, the discrete curvature variable. A transfer-matrix evaluation of the label free energy on the ring of cells around an edge gives, at the theory point, an induced coupling $c_0 = +0.019$ per slice edge, against the bare $k_0 \simeq 2.2$: about one percent. The label correlation length is about half a lattice step, and the beyond-quadratic dependence on q is negligible at physical coordinations. The closure weighting alone therefore cannot supply the host’s geometric stiffness, and its predicted effect on the coupled sweep is a steer too small to see at single-seed resolution, a prediction the sweep of J.5 confirms.

The refresh at cell order. Modeling the selected memoryless kernel as whole-boundary redraws arriving at unit rate per cell, the single-cell stationary state is exact: the probability that k faces retain their last self-drawn values is $k!(4-k)!/120$, vacuum admissibility is 0.536, and $\langle K^2 \rangle = 48.74$ against the ensemble value 50.22. The state is geometry-blind to exponential accuracy and induces no vacuum curvature action. It does couple gated geometry dynamics to the local density of closure failure. This calculation tests the effective refresh model, not a derivation of its microscopic operator.

The history-weighting ceiling. Two computations bound the strongest vacuum role a Many-Pasts tilt over closure-class observables could play. Amplifying the closure link interaction to 128 times the theory point saturates the link-agreement ratio near 0.55 with correlation length 1.7 links and a curvature-sector signal near 0.1: a structural ceiling, set by the ground-state degeneracy of the closure-plus-injectivity interaction, twenty times below the bare coupling. Independently, the exact Doob-transform solution of a closure-maintaining history tilt (paths weighted by time-integrated closure quality) on a ring of six seven-state faces drives the nearest-neighbor label correlation from -0.12 through zero to only $+0.13$ across the full tilt range, with longer-range correlations near 0.03 throughout: no long-range order at any tilt strength. Together with the two paragraphs above this closes the label sector: static weight, refresh, and closure-class history weighting each fail, quantitatively, to supply vacuum geometric stiffness at cell order.

The free conserved field, integrated out. A conserved Gaussian field back-coupled to the geometry induces, on integration, the geometric action $\frac{1}{2} \ln \det' L(g)$ per channel and slice: the spanning-tree entropy of the slice graph, by the matrix-tree theorem. Computed on degree-matched graph classes, the entropy per cell is 0.607–0.613 on diamond-cubic lattices (the smooth extended proxy), 0.609 on random 4-regular graphs (the crumpled proxy), and 0.56–0.57 on real engine slices: nearly universal at fixed coordination, with the genuinely nonlocal soft-mode part diluting as $\ln N/N$ per cell. The per-cell differential between extended and crumpled geometry is a few times 10^{-4} ; with all seven channels the phase-tipping force is of order 10^{-3} against $k_0 \simeq 2.2$.

The budget-bounded field. The finite budget of Postulate I makes the field non-Gaussian, with saturation curvature $m^2 = 2/g_{\text{share,eff}} \simeq 0.27$, and the massive determinant differential between the same proxies shrinks from 1.2×10^{-3} per cell at zero mass to 4×10^{-4} at the budget mass and toward zero beyond: the theory’s own non-Gaussianity suppresses the modes where the sensitivity lived. What survives, measured on eight engine slices, is a local coupling renormalization of scale $\simeq 0.08$ per cell in the seven-channel count, absorbed into the host’s bare-coupling scan; it moves phase boundaries and supplies no phase of its own. The two containers no branch covers are named in Section 23.3: long-range field kernels and non-equilibrium geometry–field co-evolution, both unspecified by the theory.

The fracture construction. Under unit label shifts preserving injectivity, the 1680-state move graph has 48 connected components of 35 states, labeled by slot ordering and parity. This proves that the tested move class cannot equilibrate the full admissibility ensemble. More general microscopic dynamics are not excluded.

J.5 Compatibility at matched volume

The coupled sweep at $N_{41} = 20,000$ ran four arms from a common thermalized bare checkpoint: the $\beta = 0$ control, an intermediate ramp point at $\beta = 0.3$, the theory point $\beta = 1$, and the shuffled placebo at $\beta = 1$. All arms pass the manifold, foliation, thermalization, and matched-volume gates; total simplex numbers agree within 1.6%.

arm	N_4	$\langle E \rangle$ per cell	collision fraction	d_H
$\beta = 0$ (control)	47,034	3.99	0.648	2.89
$\beta = 0.3$	46,966	2.78	0.403	2.87
$\beta = 1$ (theory point)	46,286	1.74	0.082	2.85
$\beta = 1$ placebo	46,406	1.95	0.730	2.98

The label sector behaves as designed: the control arm's collision fraction 0.648 sits on the uniform-measure prediction $1 - 840/2401 = 0.650$, the theory point drives it to 0.082, and the placebo, carrying the same energy values with no closure structure, stays at 0.730. The geometric observables (Hausdorff dimension above; the slice-volume profile behaves the same way) are statistically indistinguishable across arms at single-seed resolution, as the one-percent induced coupling of J.4 predicts. The run is single-seed; the reading is the comparative one stated in Section 23.5, and the residual few-percent spreads are within the drift of the base ensemble.

J.6 The Regge bridge: a predicted volume-sector dressing

The bridge is an ordinary statistical-mechanical statement made quantitative. Integrating the labels out at fixed geometry defines an effective action for the host,

$$e^{-S_{\text{eff}}[g]} = e^{-S_{\text{host}}[g]} Z_{\text{label}}[g], \quad S_{\text{eff}}[g] = S_{\text{host}}[g] - \log Z_{\text{label}}[g],$$

with S_{host} the Regge action plus the volume pin. The label free energy $-\log Z_{\text{label}}[g]$ is a function of the geometry alone. Its extensive part is $\beta\mu(\beta) N_{\text{cells}}$, with μ the per-cell label free energy of J.2; its geometry-dependent remainder is the induced curvature-sector action priced by the ring expansion of J.4. The centered ensemble subtracts the extensive part by construction; left in, it is the leading coefficient of the effective action, and the host's response to it can be predicted before it is measured.

The centering constant μ of J.2 therefore gives a direct prediction. In the centered ensemble, the subtraction of μ removes the extensive label free energy from the geometry marginal. If that subtraction is omitted, the geometry sees a residual positive linear contribution $+\beta\mu N_{\text{cells}}$ to the effective action. Here N_{41} denotes the total number of (4, 1) and (1, 4) pentachora, so each spatial cell, shared by one pentachoron of each orientation, is counted once by the pair and $N_{\text{cells}} = N_{41}/2$. The positive linear tilt competes with the quadratic volume pin $\varepsilon(N_{41} - \bar{N})^2$, and the shift follows in two lines. The volume-dependent part of the effective action is

$$S_{\text{vol}}(N_{41}) = \varepsilon(N_{41} - \bar{N})^2 + \beta\mu N_{\text{cells}} = \varepsilon(N_{41} - \bar{N})^2 + \frac{\beta\mu}{2} N_{41},$$

and setting $dS_{\text{vol}}/dN_{41} = 2\varepsilon(N_{41} - \bar{N}) + \beta\mu/2 = 0$ gives the equilibrium displacement

$$\Delta N_{41} = N_{41}^{\text{eq}} - \bar{N} = -\frac{\beta\mu}{4\varepsilon}.$$

Here μ is computed by annealed thermodynamic integration of the label ensemble on the base geometry, before any coupled dynamics run; the zero-lattice reference — the exact single-cell free energy from direct enumeration of the 7^4 label states, computable from η_* and the injectivity penalty alone — is 2.477 at this point, within half a percent of the lattice value, the difference being the shared-face correlations. At the demonstration point ($\beta = 1$, $\mu = 2.488$, $\varepsilon = 0.01$) the prediction is $\Delta N_{41} = -2.488/0.04 = -62.2$. Three arms from a common clean base at $N_{41} = 2000$ measure it:

arm	mean N_{41}	shortfall
bare ($\beta = 0$)	1993.2	-6.8
centered ($\beta = 1$)	1969.8	-30.2
uncentered ($\beta = 1$)	1907.0	-93.0

The uncentered-minus-centered difference is -62.8 against the predicted -62.2 . The subtraction isolates the μ -term, since the two $\beta = 1$ arms share every other piece of machinery; the centered arm's own -30 is the second-order non-neutrality of the thermodynamic-integration centering, real, cancelling in the difference, and gated by matched volume in production. The reading is the bridge of Section 23.4: a number computed from the ultraviolet specification appears, at the predicted magnitude, as an extensive contribution to the host action's volume coupling in a dynamical measurement — a dressing of the host, not a derivation of it. The measurement is demonstration-volume and single-seed; the curvature-sector counterpart, the c_0 and field shifts of J.4 read as phase-boundary displacements, is pre-registered for production statistics.

The demonstration point is one member of a predicted family, and the family is pre-registered. At fixed geometry-side settings the displacement obeys $\Delta N_{41}(\beta, \varepsilon) = -\beta\mu(\beta)/(4\varepsilon)$, and each factor is separately testable. In β the curve is concave, because $\beta\mu(\beta) = \int_0^\beta \langle \bar{E} \rangle_s ds$ and the Gibbs mean falls as the labels order, with local slope $d(\Delta N_{41})/d\beta = -\langle \bar{E} \rangle_\beta/(4\varepsilon)$ — an Ehrenfest-type relation in which the same run supplies both sides, since $\langle \bar{E} \rangle_\beta$ is measurable in the arm whose displacement it predicts. In ε the law is exactly inverse at fixed β ; departures from $1/\varepsilon$ measure the curvature of the background free energy that the matched-pair subtraction cancels at first order, so the ε family doubles as a control on the derivation's one assumption. And the placebo table carries its own computable free energy: at the theory point the single-cell values are $\mu_{\text{plc}} = 3.17$ against $\mu = 2.48$, so the placebo arm is predicted to displace by $-\beta\mu_{\text{plc}}(\beta)/(4\varepsilon)$ — not by zero, and not by the real value, a 28% separation. The measured number must track the table it was computed from, which upgrades the placebo from a null control to a second quantitative point on the family.

The family has been measured, at demonstration volume, with every prediction written to disk before any arm launched. Eight of the nine pre-registered coupling-pin pairs completed (the ninth arm failed at launch and is rerunning); across a fifteen-fold range of predicted displacements, from -10.5 to -157.6 , the measured-to-predicted ratios are 0.96 – 1.07 with one point at 0.89 . The β -curve is concave as required: -43 , -78 , -121 measured against -43 , -76 , -125 predicted — not the straight line that would have falsified the free-energy integral. Halving ε doubles the displacement to within four percent. And the placebo pair landed on its own line: displaced thirty percent from the real-table arms, at ratio 1.00 to the prediction computed from its own shuffled table — the number tracks the table, not the coupling. The theory's formula predicted nine numbers in advance and the lattice delivered eight of them, with the ninth pending a rerun rather than a discrepancy; a production-volume repetition is pre-registered.

J.7 The pinned-defect experiment

Both arms resume from the thermalized theory-point checkpoint of J.5 and run 3000 further sweeps with 100 pins each. The real arm pins cells at maximal closure failure (all four faces $m =$

0: six collisions, $K^2 = 48$); the placebo arm pins cells at the best-closed injective configuration ($m = \{0, 1, 2, 3\}$, $K^2 = 40.67$). Pinned faces are excluded from the label heat bath, and the carrier cells are protected from removal by the geometry moves identically in both arms, so the real-minus-placebo difference isolates the failure content; pin survival is a hard gate, passed at 100/100 in both arms. Freshly sampled unpinned cells far from every pin are measured inside both arms as the vacuum reference. Shell observables are accumulated by graph distance in the slice adjacency, excluding the pinned cells themselves (they are the source, the shells are the medium), for roughly 10^5 defect-environment samples per arm.

arm	shell	mean E	collision fraction	mean coordination	cells per shell
failure	1	1.6782 ± 0.0010	0.0690 ± 0.0003	5.2997 ± 0.0007	4
	2	1.7325 ± 0.0008	0.0790 ± 0.0003	5.5789 ± 0.0009	7.237 ± 0.003
	3	1.7274 ± 0.0006	0.0787 ± 0.0002	5.6459 ± 0.0007	11.296 ± 0.006
closed (placebo)	1	1.7291 ± 0.0011	0.0794 ± 0.0004	5.2513 ± 0.0008	4
	2	1.7304 ± 0.0008	0.0792 ± 0.0003	5.5773 ± 0.0009	7.116 ± 0.005
	3	1.7280 ± 0.0007	0.0788 ± 0.0002	5.6460 ± 0.0008	10.905 ± 0.007
vacuum	1	1.7306 ± 0.0011	0.0797 ± 0.0004	5.3165 ± 0.0018	4
	2	1.7303 ± 0.0008	0.0795 ± 0.0003	5.6108 ± 0.0014	7.235 ± 0.006
	3	1.7288 ± 0.0007	0.0789 ± 0.0002	5.7310 ± 0.0012	10.909 ± 0.012

The capacity state responds strongly at the first shell. At the second shell the collision fraction is consistent with vacuum and the mean closure energy lies within about two combined quoted standard errors, so the data bound the response to roughly one shell without proving exact return at shell two. The local geometry responds through the third shell, including a 3.6% volume excess where the placebo matches the vacuum. The quoted errors are accumulation errors and are not corrected for autocorrelation. An independent seed reproduces the first-shell capacity separation. The coordination response retains its direction at about half the first seed's magnitude, so geometric magnitudes remain reported with across-seed scatter. The capacity response is replicated; the geometric response is directionally replicated.

J.8 The conservation law and the transport instrument

The dichotomy of Section 23.7 in operator form. If free capacity is a per-cell budget re-equilibrated by maximum entropy — link shares $c_l \geq 0$ with $\sum_{l \ni x} c_l = C - m_x$ — stationarity gives one multiplier per cell, and the linearized constraint reads $(2z\mathcal{K} - L)\delta\mu = -m/\chi$ with $z = 4$ and L the slice Laplacian. The operator has no small-momentum pole; on a measured 169-cell slice its Green function falls eight orders of magnitude within thirteen steps. If instead free capacity f_x obeys a continuity law,

$$f_x(t+dt) - f_x(t) = - \sum_{\text{faces } xy} J_{xy} - \text{commit}_x + \text{release}_x, \quad J_{xy} = -J_{yx}, \quad \sum_x [f_x + m_x] = \text{const},$$

with a defect drawing steady maintenance flux $Q_x \propto m_x$, then any local flux law linearizes to $J \propto -D\nabla f$ and the steady state solves $L\delta f = (Q/D)(\delta_{\text{source}} - \text{uniform return})$: the massless graph-Poisson equation, whose Green function on a three-dimensional slice falls as $1/r$. The deficit field is $\delta f \propto m/r$, and the relative deficit is the weak-field bridge variable of Section 11.

The instrument implements the conservation reading with the sink left to emerge. A field on the slice cells is initialized at the vacuum anchor $g_{\text{share,eff}} = 7.4198$ and transported by antisymmetric face fluxes; global conservation of $\sum_x (f_x + m_x)$ is asserted at machine precision every sweep, and the run aborts on any violation. Absorption is proportional to a cell's instantaneous label collisions, so a pinned defect drains capacity only through the closure failure it induces in itself and its neighborhood by the ordinary label dynamics: the maintenance flux each defect draws is measured, not set. Pins are placed at four commitment levels (1, 2, 3, and 6 colliding pairs), giving the two headline readings: the emergent flux against commitment,

which tests maintenance linearity dynamically, and the first-shell deficit per unit flux, whose constancy across the ladder is the junction statement that fixes the lattice analog of G . The two profile gates are the deficit against the graph-Coulomb Green function and a fitted screening mass consistent with zero, any conservation leakage appearing as Yukawa screening. Absorption reads the failure a cell carries above the vacuum's own standing level (the 46% transient density of J.4), so the vacuum is absorption-free and the transported field massless by construction. The first full-volume run demonstrated why this matters: with absorption reading total failure instead, the vacuum itself absorbs, and the measured deficit range of four to five steps reproduced the predicted self-screening length $\xi = \sqrt{D/(\kappa \langle n_{\text{coll}} \rangle_{\text{vac}})} \simeq 5\text{--}7$ — the model's own Yukawa mass, a post-diction confirming that the transport sector behaves as derived. The ladder also carries a level-zero rung, a frozen best-closed pin, which measures the commitment-independent boundary-dressing flux directly; the linearity reading is the dressing-subtracted fit. The first implementation freezes the geometry while the labels stay fully dynamical, isolating the field sector; conservative transport through the geometry moves is specified as the extension.

On small validation volumes the instrument behaves physically: deficit wells ordered by commitment, monotone recovery with distance, emergent flux rising along the ladder, conservation exact — and the highest-commitment pin draws less flux than the linear extrapolation because it depletes its own neighborhood, a self-screening the model produces without being told.

The production run at $N_{41} = 20,000$ delivered the near-field readings and voided its own far-field verdict, in that order of importance. Conservation held exactly through four thousand sweeps. The level-zero rung measured the boundary dressing directly ($Q(0) = 0.11$), and with it subtracted the emergent flux is proportional to commitment, $Q - Q(0) = 0.187m$ with largest relative residual 0.17; the junction constant is level-independent at 2.53 ± 0.23 , nine percent across the ladder. The far-field fit returned screening at 1.6 steps and is not readable, for a diagnosed and computed reason: reading the *instantaneous* excess rectifies vacuum fluctuations. The vacuum's collision distribution at the theory point is 92% zero and 8% one, so $\max(0, n_{\text{coll}} - \bar{n})$ with $\bar{n} = 0.077$ retains a mean rate of 0.071 — ninety-two percent of the vacuum absorption survives the subtraction — building an intrinsic screening length $\xi = \sqrt{D/(\kappa \cdot 0.071)} \simeq 7.5$ into the instrument by construction. The still-shorter measured value is additional bias from referencing deficits to the global anchor on closed slices, whose recycled surplus elevates the far field, compounded by the level-zero dressing carrying the opposite sign to assumption: frozen injective faces *suppress* neighbor collisions below vacuum, so the closed pin under-absorbs and sits in a local surplus (7.76 at the first shell against the 7.42 anchor) — itself a physical finding of the run. The corrected instrument reads the persistent failure average (the theory's definition of commitment; the window suppresses the rectified vacuum rate toward zero, pushing the intrinsic length to roughly twenty-three steps), references deficits to the same-slice far field, and reran on thicker slices where the usable range exceeds the artifact scales. Its readings were pre-registered: linearity failure strikes the maintenance postulate, a screening mass that survives the corrections strikes the conservation reading, and either would be reported as such; the pre-registered side prediction is that the level-zero surplus collapses once the vacuum no longer absorbs.

The corrected forty-slice run conserved the field through four thousand sweeps and kept its vacuum anchor. The screening fit is consistent with a massless response, and the deficit persists to the fifth shell. The profile exponent is -0.65 ± 0.14 against the ideal -1 , with finite-slice flattening. The level-zero flux fell from 0.1075 to 0.0086, as predicted when vacuum absorption was removed. The ladder gives $Q - Q(0) = 0.166m$ and a junction constant 4.66 ± 0.42 . The run is single-seed and its errors are not autocorrelation-corrected; within those caveats the pre-registered gates pass.

Appendix K: Numerical Checks and Robustness

Appendix K is intentionally modest. It does not add new derivations, and these checks do not establish the ontology: they audit the numerical consequences of the stated derivations, and their role is reproducibility, not independent proof. It collects the main numerical cross-checks and the substrate-scale ledger so the shared coefficient chain can be inspected before the provenance and action-completion appendices. The complete self-contained reproduction code is reserved for terminal Appendix O. The dynamical-lattice program is a separate layer with a different job: where this appendix recomputes the chain, Section 23 and Appendix J couple it to a dynamical substrate and measure what it does.

The numbers recomputed here, and where each enters the main text, are: the admissibility entropy $g_{\text{share,eff}}$ (Section 13, Appendices B and D.4); the substrate length L_* (Sections 2 and 13, Appendix D.4); the induced scale G_* (Sections 10 and 13, Appendices D.4 and K); the galactic scale a_0 (Section 14, Appendix C.7); the diamond-lattice Green constant $G_{\text{tet}}(0)$ (Appendix C.5); the edge-kernel chain $J_{\text{bare}} \rightarrow \gamma$ and the return sum Σ_{ret} (Appendices C.3–C.4); and the charged-lepton ladder with its support-exponent cross-check (Section 13, Appendix I.1).

K.1 Cross-sector numerical checks

The cross-check program includes:

- the one-bit fermionic defect check $\Delta S_f = \ln 2$;
- the rooted-shell convergence check $\sigma_{\text{ind}}^{(2)} \simeq \sigma_{\text{ind}}^{(3)}$;
- the UV closed-branch moments $\langle K^2 \rangle_{\eta_*}$, $\text{Var}_{\eta_*}(K^2)$, and a_{UV} ;
- the Gaussian closure-amplitude identity, nine-state response Gram matrix, unitary pair dilation, 21-block edge Hessian, and 56,800-state routing enumeration;
- cross-sector consistency among the electron anchor, the substrate length L_* , the induced scale G_* , the Green-matched Newton closure, and the galactic scale a_0 .

These checks do not replace the derivations, but they show that the same coefficient chain survives independent numerical scrutiny across the sectors where closure is claimed.

K.2 Reproducibility ledger for the substrate scale

The substrate-length calculation is short enough to record as a numerical ledger. Enumerating the 1680 oriented injective tetrahedral states with labels $m = -3, \dots, 3$, weighting them by $e^{-\eta K^2}$, and solving

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

gives

$$\eta_* = 0.02986684439352237, \quad g_{\text{share,eff}} = 7.419800023570903.$$

The one-pass seven-sector history support and its transverse export are then

$$e^{7g_{\text{share,eff}}} = 3.602860521062804 \times 10^{22},$$

$$\frac{2}{3}e^{7g_{\text{share,eff}}} = 2.401907014041869 \times 10^{22}.$$

Using $\lambda_e = \hbar/(m_e c) = 3.861592671986303 \times 10^{-13}$ m and $r = e^{-7g_{\text{share,eff}}}$ first gives the baseline

$$L_*^{(0)} = -\frac{3}{2}\lambda_e \ln(1 - r) = 1.607719470158885 \times 10^{-35} \text{ m},$$

while the marked vertex gives

$$\zeta_* = 0.005123584484947, \quad Z_e = 1.005308283809514.$$

Therefore

$$L_* = Z_e L_*^{(0)} = 1.616253701392569 \times 10^{-35} \text{ m},$$

and

$$G_* = \frac{c^3 L_*^2}{\hbar} = 6.674289077220912 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}.$$

The source-side Green constant is obtained independently from the diamond-lattice integral in Appendix C.5. Uniform-grid quadrature with endpoint extrapolation gives

$$G_{\text{tet}}(0) = 0.448220394388 \dots,$$

confirming the exact Joyce value

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

This is the value used in the source theorem. These numbers are not additional inputs; they are the numerical evaluation of the finite spectrum, the seven-sector history support, and the graph Green function already defined in the derivation.

Appendix L: Fork accounting for the scale chain

Because the substrate length induces a value of G_* near the observed Newton constant, this appendix records the construction choices behind that result and quantifies the discrete-choice fork space around it, so the scale chain can be weighed against the freedom available in the construction. It is referenced wherever G_* is discussed in the main text.

Provenance. By the author’s recollection, a target entropy near 7.42 was identified early in the framework’s development by inverting the measured Newton constant, with the ensemble constructed to deliver it jointly with other structural constraints; the documentary record does not reach that genesis layer. The realized ensemble was fixed and motivated in print by information-independence arguments before the induced- G_* chain existed. Its original electron scale and the baseline charged-lepton ladder landed at discrepancies of roughly one percent and one-half percent. Those residuals were public parts of the construction before the decorated marked vertex was developed.

Why the residuals suggested one missing object. The relevant comparison is at the transfer-amplitude level. The original Newton value was low by about 1.05%, while the length entering it was low by half that amount because $G_* \propto L_*^2$. The three required multiplicative uplifts were

baseline quantity	required uplift	structural reading
$L_*^{(0)}$ inferred from G	0.5309%	electron feedback
$m_\mu^{(0)}/m_e$	0.5124%	universal marked response
$m_\tau^{(0)}/m_e$	0.6562%	universal response plus second shell

Dividing the tau uplift by the muon uplift leaves a further 0.1431% second-shell contribution. The pattern is not three identical residuals, and it was not used to impose a universal rescaling. It is more specific: two observables require nearly the same charged-transfer amplitude, while the third requires that common amplitude followed by one smaller shell passage. This was the motivation for seeking one marked response with different graph traces.

Why the repair had to be additive. By that stage the tetrahedral construction was already doing several independent jobs. Its label space and admissibility weight fixed the closure spectrum and sharing entropy; its edge projection fed the ordinary weak-field coefficient; its transverse export entered the scale, horizon normalization, and galactic branch; and its shell algebra produced the baseline charged-lepton ladder. Removing or retuning any of those ingredients to close one residual would propagate through the rest of the microscopic and macroscopic chain. A viable repair therefore had to preserve the unmarked tetrahedral ensemble and its ordinary transfer operator.

This requirement narrowed the search to a defect-bound additive sector: it must be invisible in the unmarked vacuum, reuse the closure response already present in the renewed cell, activate only on occupied charged-channel pairs, and route through the existing electron and shell graph. These conditions explain the form of the calculation in Appendix H. They do not by themselves prove its nine-state fiber, 21 pair records, Majorana determinant, or routing polynomials; those are the independent action-level results and audits. Showing this part of the development is appropriate because the residuals were already known. It makes clear which features were motivated by the discrepancy, which were protected by the existing theory, and which had to be derived rather than chosen.

The vertex in Appendix H was consequently found with the discrepancies known. It contains no continuously adjusted coefficient: $g_{\text{share,eff}}$, η_* , $2/7$, the nine-state trace, 21 pair records, and the three routing polynomials are fixed by the displayed action. That fact makes the result an action-level postdiction rather than a fit, but it does not make it a blind prediction. Its added evidential content is that one field space and one routing graph account for G , m_μ/m_e , and m_τ/m_e simultaneously, while the nearby shared-mode and off-diagonal alternatives fail. The cluster capacity bound, committed abundance, redshift direction of a_0 , and sub- g_c rotation-curve regime remain post-fixation tests of the original ensemble; none entered either the early entropy calibration or the marked residual calculation.

The audit. The exact substrate length is proportional to $-Z_e \ln(1 - e^{-7g_{\text{share,eff}}})$, with leading behavior $L_* \propto Z_e e^{-7g_{\text{share,eff}}}$, so the calibration is meaningful only relative to the space of constructions that could have been written instead. The original grammar contains six label counts, five multiplicity rules, five closure coefficients, four support exponents, and five export prefactors, for 3000 candidate constructions. Those counts predate the marked vertex and must be recomputed with its $Z_e(g, \eta)$ dependence before they are used quantitatively.

The marked-sector audit separately enumerates 576 nearby formulas obtained by varying the response dimension, pair count, direction factor, determinant power, statistics, and overlap. Only the decorated-vertex formula lies within one muon standard deviation. This is not a p -value because no probability measure is assigned to the menu. It records brittleness and makes the action selection falsifiable. A shared edge register, directed complex determinant, stable bosonic determinant, or six-state symmetric response gives a visibly different result.

The galactic scale depends only linearly on $g_{\text{share,eff}}$, so it tests the coupling form more strongly than the exact ensemble choice. The corrected G_* and charged-lepton ratios test the shared marked vertex through three different graph factors, with the postdiction status stated above.

Appendix M: Why the saturated CMB carrier must be committed capacity, not a lagged response

The microwave background requires a gravitating component that carries the growing mode of the baryon perturbations while rejecting their acoustic oscillation. This appendix records the

quantitative exclusion of every relaxational realization of that component and the no-go that selects the constraint reading of Section 18.5.

Growth-limited kernels. For development dynamics $dW/dt = (W_{\max} - W)/\tau$ with $\tau = \beta t_{\text{ff}}(\rho_{\text{local}})$, the measured cluster radial decline of the source weight fixes $\beta \simeq 15.6$, while saturation of the cosmological bath by recombination requires $\beta \lesssim 0.56$: the window is empty by a factor of order thirty. The radial ratio between R_{500} and R_{200} , a β -independent prediction of the free-fall clock, agrees with the data at the several-percent level, so the exclusion is of the absolute clock, not of the radial structure.

Lagged response. A first-order lag transmits a fraction $T = [1 + (\omega\tau)^2]^{-1/2}$ of an oscillation at frequency ω . The transmitted component acts as additional effective baryon loading $(1/\epsilon - 1)T$ in the acoustic driving; Einstein–Boltzmann computation shows a 0.3% temperature-spectrum tolerance bounds $T \leq 2 \times 10^{-3}$, requiring $\beta \gtrsim 37$ against the development requirement $\beta \lesssim 0.56$: empty by a factor of order sixty-six. More generally, by recombination a third-peak mode has completed only a few oscillations, so no causal filter of any order achieves the required rejection.

Sound-speed classification. For $L = f(X)$, perturbations carry $c_s^2 = f_X/(f_X + 2Xf_{XX})$. The released branches give $c_s^2 = \frac{1}{2}$ (deep) and 1 (Newtonian): free-streaming. Plateau approaches $f = f_{\max} - A/X^n$ give $c_s^2 = -1/(2n+1)$: gradient-unstable. Both natural cap readings — stationarity of the canonical momentum and of the energy density — impose $f_X + 2Xf_{XX} = 0$, the pole of c_s^2 : a rigid medium carrying no perturbation. The only locus with $c_s^2 = 0$ and healthy density response is the constraint surface of the mimetic class, X pinned with a multiplier, which is the reading of Section 18.5.

No-go for relaxation-plus-cap carriers. Within dynamics consisting of relaxation toward a demand together with a hard cap, with the cell weight as the only state variable, no coupling of the weight to the demand carries a secular component of the modulation while rejecting its oscillation: a pinned weight retains no memory of the modulation, and any instantaneous coupling of bounded periodic inputs is itself bounded and periodic. A secular component requires an additional conserved integrating variable. The conserved spatial density of committed cells is that variable, and it is native to the saturated phase rather than added to it.

Appendix N: Action Reconstruction, Parent-Theory Audit, and Transverse-Sector Target

This appendix is the action workspace. It answers the narrow question that must be settled before a more elaborate covariant capacity theory is proposed: is the ordinary capacity functional a second scalar action, or is it the gravitational constraint action in a different variable? The result is decisive in the static longitudinal sector and deliberately incomplete in the transverse thermal sector.

N.1 One parent metric, one matter coupling

The minimal covariant parent of the ordinary branch is

$$I_0[g, \psi] = I_{\text{EH}}[g] + I_{\text{GHY}}[g] + I_{\text{matter}}[g, \psi],$$

$$I_{\text{EH}} = \frac{c^3}{16\pi G} \int_{\mathcal{M}} d^4x \sqrt{-g} (R - 2\Lambda).$$

Here and throughout the covariant formulas, $x^0 = ct$. Thus $d^4x = dx^0 d^3x$ and the Einstein–Hilbert coefficient is $c^3/(16\pi G)$. After the split $dx^0 = c dt$, the ADM action written with t in seconds carries $c^4/(16\pi G)$. This convention makes the final Newton action an ordinary $\int dt L$ action. All matter, including radiation, couples to the one physical metric. Its stress tensor is

$$T_{\mu\nu} = -\frac{2c}{\sqrt{-g}} \frac{\delta I_{\text{matter}}}{\delta g^{\mu\nu}},$$

and the equations are

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}.$$

The Bianchi identity and the matter equations give $\nabla_\mu T^{\mu\nu} = 0$. No explicit coupling $S_{\text{ent}} T^\mu{}_\mu$ is present. Such a term would vanish as a source for classical radiation, would have a nontrivial metric variation through $\delta T^\mu{}_\mu / \delta g^{\alpha\beta}$, and would duplicate the matter response already carried by the metric equation.

The capacity variable enters after gauge choice and constraint reduction:

$$\delta S = \delta S[g, \psi] \big|_{\text{reduced}}.$$

This is analogous to using the Newtonian potential as a reduced gravitational variable. A reduced variable can have a useful functional without being a new covariant matter field.

N.2 Direct reduction of the Einstein action

ADM form and boundary terms. With the GHY term included, the bulk gravitational action is

$$I_{\text{ADM}} = \frac{c^4}{16\pi G} \int dt d^3x N \sqrt{h} \left({}^{(3)}R + K_{ij} K^{ij} - K^2 - 2\Lambda \right) + I_\infty.$$

This line uses $t = x^0/c$ and is therefore consistent with the covariant c^3 normalization above. For a static scalar perturbation in Newtonian gauge,

$$N = \sqrt{1 + \frac{2\Phi}{c^2}} = 1 + \frac{\Phi}{c^2} + O(c^{-4}), \quad N^i = 0, \quad h_{ij} = \left(1 - \frac{2\Psi}{c^2}\right) \delta_{ij}, \quad K_{ij} = 0.$$

The asymptotic term I_∞ subtracts the flat reference contribution and cancels the total divergence generated when the spatial curvature is integrated by parts.

Define $\phi = \Phi/c^2$, $\psi_N = \Psi/c^2$, and write $h_{ij} = e^{2\zeta} \delta_{ij}$, with

$$\zeta = \frac{1}{2} \ln(1 - 2\psi_N) = -\psi_N - \psi_N^2 + O(\psi_N^3).$$

In three dimensions,

$${}^{(3)}R = e^{-2\zeta} [-4\nabla^2 \zeta - 2(\nabla \zeta)^2].$$

Before integration by parts, the expansion through second order is

$$N\sqrt{h} {}^{(3)}R = 4\nabla^2 \psi_N + 4\phi \nabla^2 \psi_N + 4\psi_N \nabla^2 \psi_N + 6(\nabla \psi_N)^2 + O(3).$$

The first term is a boundary term. Integrating the next two Laplacians by parts and combining all surface pieces with I_∞ gives

$$N\sqrt{h} {}^{(3)}R = \frac{2}{c^4} [(\nabla \Psi)^2 - 2\nabla \Phi \cdot \nabla \Psi] + \nabla_i B^i + O(c^{-6}).$$

The boundary integral of B^i is removed by the reference-subtracted variational prescription. For nonrelativistic matter with ρ the rest-mass density, expansion of $-mc^2 \int d\tau$ gives

$$I_{\text{matter}}^{(1)} = - \int dt d^3x \rho \Phi.$$

Consequently,

$$I_{0,\text{scal}}^{(2)}[\Phi, \Psi] = \int dt d^3x \left\{ \frac{(\nabla \Psi)^2 - 2\nabla \Phi \cdot \nabla \Psi}{8\pi G} - \rho \Phi \right\}.$$

Constraint equations. Variation with respect to the lapse potential gives

$$\delta_\Phi I_{0,\text{scal}}^{(2)} = 0 \implies \nabla^2 \Psi = 4\pi G \rho.$$

Variation with respect to the scalar part of the spatial metric gives

$$\delta_\Psi I_{0,\text{scal}}^{(2)} = 0 \implies \nabla^2 (\Phi - \Psi) = 0.$$

The elimination order is essential: Φ and Ψ are varied independently, both constraint equations are obtained, and only then is their solution substituted back. Imposing $\Phi = \Psi$ in the two-potential action before variation would erase one constraint and is not an equivalent reduction. For asymptotically flat isolated sources, the harmonic difference vanishes:

$$\Phi = \Psi.$$

Substituting this solution back into the two-potential action yields

$$I_{\text{Newton}}[\Phi] = \int dt d^3x \left[-\frac{(\nabla \Phi)^2}{8\pi G} - \rho \Phi \right].$$

The sign of the reduced kinetic term is not a ghost signal. Φ is a nondynamical constraint potential, not a propagating scalar oscillator. In the full covariant parent, the lapse and shift impose constraints and the propagating gravitational content is the two tensor polarizations.

N.3 Capacity field redefinition and coefficient identity

On the renormalized static branch, the capacity functional is

$$I_{\text{cap}}^{\text{static}} = \int dt d^3x \left[-\frac{\gamma}{2} (\nabla \delta S)^2 + \kappa \rho \delta S \right].$$

The most general linear local field redefinition compatible with a constant vacuum background is $\delta S = A\Phi$. Matching the signs of the source and the weak-field clock shift fixes $A < 0$. The bounded relation $q = 1 - \delta S/S_\infty = N^2$ at first order fixes

$$A = -\frac{2S_\infty}{c^2}.$$

Substitution gives

$$I_{\text{cap}}^{\text{static}} = \int dt d^3x \left[-\frac{2\gamma S_\infty^2}{c^4} (\nabla \Phi)^2 - \frac{2\kappa S_\infty}{c^2} \rho \Phi \right].$$

The equality

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

then makes the kinetic coefficient and source coefficient agree simultaneously:

$$\boxed{I_{\text{cap}}^{\text{static}} = \frac{2\kappa S_{\infty}}{c^2} I_{\text{Newton}}.}$$

This is stronger than observing that both functionals yield Poisson equations. It proves off shell, within the reduced static field space, that they are the same classical functional up to a field-independent normalization. The statement uses matched Dirichlet or asymptotically flat boundary conditions. With a finite boundary, equality also requires mapping the Newton boundary functional through the same linear field redefinition.

Normalization freedom. The observable static equation depends only on

$$\frac{\kappa}{\gamma S_{\infty}}.$$

A change of entropy units

$$\delta S \rightarrow a \delta S, \quad S_{\infty} \rightarrow a S_{\infty}, \quad \gamma \rightarrow \frac{\gamma}{a^2}, \quad \kappa \rightarrow \frac{\kappa}{a}$$

leaves both the action and the observable bridge invariant. Multiplying the entire reduced action by a further constant also leaves its classical equation invariant. The UV calculation can fix these conventions for fluctuation normalization and correlation functions, but the classical Newton constant is fixed only once through the invariant combination. A convenient unit-normalized representative has $Z_S = 1$, equivalently

$$\kappa = \frac{c^2}{2S_{\infty}}, \quad \gamma = \frac{c^4}{16\pi G S_{\infty}^2},$$

but the manuscript's canonical UV convention need not be changed to use the equivalence.

Why canonical scalar stress cannot do this job. If $\delta S = O(\rho)$ about a constant background, then

$$T_{\mu\nu}^{(\delta S)} \sim \partial_{\mu} \delta S \partial_{\nu} \delta S = O(\rho^2).$$

The bridged potential is $O(\rho)$. It therefore cannot be attributed to the canonical stress of a second scalar. In the reconstruction above, the linear response comes from the metric constraint and the apparent scalar functional is simply its reduced representation.

N.4 Parent-action options and degree-of-freedom audit

Route	Variables / action	Modes	Audit result
Metric-only / composite capacity	$I_0[g, \psi]$; δS defined after weak-field constraint reduction and $q_{\text{geo}} = (\nabla R)^2$ in spherical symmetry	two tensors	Passes the weak-field action identity, couples radiation correctly, gives baseline no slip and GR PPN, and passes the spherical exterior test. This is the preferred ordinary branch. A generic local covariant definition of capacity outside the controlled reductions remains open.

Route	Variables / action	Modes	Audit result
Algebraic ADM multiplier	$I_0 + \int \sqrt{h} \lambda (N^2 - q)$	no new mode only because the construction is empty	δq gives $\lambda = 0$ when q has no other bulk term. The relation only renames a foliation-dependent lapse and does not yield the capacity source equation. Excluded as a parent completion.
Clock-constrained auxiliary capacity	$I_0 + \int \sqrt{-g} \Lambda (qX + 1), X = g^{\mu\nu} \partial_\mu \tau \partial_\nu \tau$	trivial if q is purely algebraic; extra scalar or dust-like mode once dynamics is supplied	The q equation again drives Λ to zero unless additional q dependence is introduced. Kinetic or higher-derivative terms generically introduce a preferred foliation, mimetic-dust behavior, or a scalar stability problem. Retain only as a control model with a full Hamiltonian analysis.
Scalar-tensor / disformal	$I[g, q] + I_{\text{matter}}[\tilde{g}, \psi], \tilde{g}_{\mu\nu} = C(q)g_{\mu\nu} + D(q)u_\mu u_\nu$	two tensors plus at least one scalar	Universal matter coupling can be arranged, but fifth forces, gravitational slip, PPN shifts, and radiative stability become new obligations. It is not needed for the ordinary branch and is not selected by the present evidence.

The metric-only route is minimal in a precise sense: it adds no propagating field, uses the full stress tensor, and reproduces both controlled reductions. This does not prove that every substrate sector is metric-only. It proves that the ordinary longitudinal capacity mode must not be counted again as an independent force.

N.5 Exact spherical reduction and the Misner–Sharp first integral

Let

$$ds^2 = h_{ab}(x) dx^a dx^b + R^2(x) d\Omega^2.$$

We use $x^0 = ct$, so $d^2x = dx^0 dr$. If the orbit-space integral is instead written as $dt dr$, its prefactor is $c^4/(4G)$. The curvature identity is

$${}^{(4)}R = {}^{(2)}R + \frac{2}{R^2} [1 - (\nabla R)^2 - 2R\Box R].$$

After the angular integral,

$$I_{\text{EH}} = \frac{c^3}{4G} \int d^2x \sqrt{-h} [R^{2(2)}R + 2\{1 - (\nabla R)^2\} - 4R\Box R].$$

Integrating the last term by parts together with the reduced GHY term gives

$$I_{\text{sph}} = \frac{c^3}{4G} \int d^2x \sqrt{-h} [R^{2(2)}R + 2(\nabla R)^2 + 2] + I_{\text{matter}}^{(2)} + I_{\partial}^{(2)}.$$

The reduced boundary term is not left implicit. For a product boundary $\partial\mathcal{M}_4 = \partial\mathcal{M}_2 \times S^2$,

$$K^{(4)} = K^{(1)} + \frac{2}{R} n^a \nabla_a R.$$

The second term cancels the boundary contribution generated when $-4R\Box R$ is integrated by parts. With the standard orientation sign ε , the surviving term is

$$I_{\partial}^{(2)} = \varepsilon \frac{c^3}{2G} \int_{\partial\mathcal{M}_2} dy \sqrt{|\gamma_{(1)}|} R^2 K^{(1)} + I_{\text{joint}} + I_{\text{ref}}.$$

The vacuum Euler–Lagrange equations are

$$R^{(2)}R - 2\Box R = 0,$$

$$2R(h_{ab}\Box R - \nabla_a\nabla_b R) + h_{ab}[(\nabla R)^2 - 1] = 0.$$

Define

$$q_{\text{geo}} \equiv (\nabla R)^2, \quad M_{\text{MS}} \equiv \frac{c^2 R}{2G}(1 - q_{\text{geo}}).$$

Contracting and differentiating the metric equation, with the R equation used to remove the curvature term, gives

$$\nabla_a M_{\text{MS}} = 0$$

in vacuum. Hence

$$q_{\text{geo}} = 1 - \frac{2GM}{c^2 R}.$$

In the static areal gauge $R = r$, asymptotic flatness fixes the remaining normalization and returns

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

$$q_{\text{geo}} = N^2.$$

This determines the variational status requested by the action audit:

q is a composite geometric scalar in spherical symmetry, and its vacuum profile is a first integral.

No auxiliary scalar is required. The result is nonperturbative but symmetry restricted.

N.6 Covariant capacity observable beyond the controlled reductions

The action reconstruction fixes two limits but does not yet provide one local scalar valid in every geometry. In a stationary spacetime with a timelike Killing field ξ^μ normalized so that $\xi^\mu \xi_\mu \rightarrow -c^2$ at the asymptotic reference boundary, the lapse has the invariant representative

$$q_K = -\frac{\xi^\mu \xi_\mu}{c^2} = 1 + \frac{2\Phi}{c^2} + O(c^{-4}).$$

In spherical symmetry the invariant representative is instead

$$q_{\text{MS}} = (\nabla R)^2 = 1 - \frac{2GM_{\text{MS}}}{c^2 R}.$$

They coincide in the static Schwarzschild exterior. Outside stationary or spherical sectors neither a preferred Killing field nor an areal-radius dilaton exists, so neither formula is a general definition.

This is not evidence that a new fundamental field is required. A Newtonian potential is not a local curvature scalar: it depends on relational boundary data and gauge choice, whereas local curvature invariants measure tidal derivatives. The natural general object is therefore allowed to be quasilocal and state dependent,

$$q_{\text{cap}}(x) = \mathcal{Q}[g_{\mu\nu}, \rho_{\text{q}}; \mathcal{D}_x, \mathcal{A}(\mathcal{D}_x)],$$

where $\mathcal{D}_x$ is a relationally specified causal diamond or world tube, $\mathcal{A}(\mathcal{D}_x)$ is its observable algebra, and ρ_{q} is the quantum state relative to the chosen vacuum. This notation is a target class, not a proposed closed formula.

An informational realization could use a vacuum-relative entropy or modular-energy functional on $\mathcal{A}(\mathcal{D}_x)$ followed by a bounded constitutive map. Because regional relative entropy is unbounded while $q_{\text{cap}} \in [0, 1]$, a global linear identification is impossible. Schematically one would need

$$q_{\text{cap}} = \mathfrak{f}\left(\frac{S_{\text{rel}}}{S_{\infty}}, \text{geometric data}\right), \quad \mathfrak{f}(0, \dots) = 1, \quad 0 \leq \mathfrak{f} \leq 1,$$

with its first variation reproducing the weak-field bridge and its spherical restriction reproducing q_{MS} .

A successful definition must satisfy all of the following:

1. it is diffeomorphism invariant once its region, state, and reference data are specified relationally;
2. it introduces no independent Cauchy data and hence no extra propagating gravitational degree of freedom;
3. it reduces to q_K in the stationary weak-field sector and to q_{MS} in spherical symmetry;
4. its response is causal or explicitly quasilocal, with no dependence on an arbitrary coordinate foliation;
5. it is compatible with the metric Ward identity and the full stress-tensor source;
6. it supplies a controlled bounded continuation near saturation rather than identifying bounded capacity globally with unbounded relative entropy.

Determining $\mathcal{Q}$ is the general constitutive problem left by the action reconstruction.

N.7 The $q_{\text{geo}} = 0$ variational boundary

The exterior action should be defined first on a stretched timelike surface

$$\mathcal{B}_{\epsilon} = \{q_{\text{geo}} = \epsilon\}, \quad \epsilon > 0.$$

For Dirichlet metric data, include

$$I_{\text{GHY}}[\mathcal{B}_{\epsilon}] = \epsilon \frac{c^3}{8\pi G} \int_{\mathcal{B}_{\epsilon}} d^3y \sqrt{|\gamma|} K,$$

the reference term at infinity, and the appropriate joints. Only after the variation is defined should the null limit $\epsilon \rightarrow 0^+$ be taken, using the null-boundary form and a fixed normalization of the generators. This supplies a well-posed exterior Einstein variational principle.

Three logically different claims must not be conflated:

1. The reduced metric equations produce a marginal sphere where $q_{\text{geo}} = 0$. This is derived.
2. The substrate capacity fraction equals q_{geo} up to that surface. This is a constitutive identification, strongly matched in static spherical vacuum.
3. The physical spacetime ends there and does not realize $q_{\text{geo}} < 0$. This is the bounded-domain postulate and requires new boundary dynamics.

The standard gravitational boundary term fixes neither a channel Hamiltonian nor a reflectivity. The missing microscopic functional may be denoted

$$\Gamma_{\partial q}[\sigma_{AB}, \zeta_{\partial}; \rho_{\partial}],$$

where ζ_{∂} are boundary channel variables and ρ_{∂} is their state. It must determine:

- whether the GR ingoing condition $\mathcal{R}(\omega) = 0$ is recovered;
- whether a stretched layer $q_{\text{geo}} = \epsilon$ is dynamically selected;
- the relaxation spectrum and any echo kernel;
- the shape variation and boundary stress balance;
- the microscopic channel-to-area count behind the $1/4$ normalization identity.

Until this functional is derived, the Schwarzschild exterior is a result, while interior excision and boundary phenomenology are hypotheses.

N.8 Transverse effective Hessian and influence-functional target

The transverse extension has a controlled symmetry reduction and a conditional effective potential. Neither result supplies the metric response by itself. This subsection derives both pieces and then states the closed-time-path calculation still required.

Tetrahedral kinetic normalization. Let $q_i(x)$, $i = 1, \dots, 4$, denote fluctuations on the four equivalent tetrahedral face directions $\mathbf{n}_i$, normalized by

$$\mathbf{n}_i \cdot \mathbf{n}_j = \begin{cases} 1, & i = j, \\ -1/3, & i \neq j, \end{cases} \quad \sum_i \mathbf{n}_i = 0.$$

For the normalization argument use the Euclidean quadratic term

$$\Gamma_{E,\text{kin}} = \frac{Z_0}{2} \int d^4x_E \sum_i \partial_\mu q_i \partial_\mu q_i.$$

The q_i are dimensionless occupation fluctuations, and Z_0 has units of action per length squared. Lorentzian response is obtained by analytic continuation to the signature of Appendix A. The four-dimensional permutation representation splits into a uniform singlet and a three-dimensional sum-zero sector. On the latter define

$$A^a = \frac{\sqrt{3}}{2} \sum_i n_i^a q_i, \quad \sum_i q_i = 0.$$

Using $\mathbf{n}_i \cdot \mathbf{n}_j = 4\delta_{ij}/3 - 1/3$ gives the exact norm identity

$$\mathbf{A}^2 = \sum_i q_i^2, \quad \Gamma_{E,\text{kin}} = \frac{Z_0}{2} \int d^4x_E \partial_\mu \mathbf{A} \cdot \partial_\mu \mathbf{A}.$$

The same result follows from Schur's lemma: the tetrahedral three-dimensional representation is irreducible, so a symmetry-invariant quadratic operator is proportional to the identity on it. A slowly varying baryonic field selects $\mathbf{e}_L = \nabla \Phi_{\text{bar}} / |\nabla \Phi_{\text{bar}}|$ and permits

$$\mathbf{A} = X \mathbf{e}_L + \mathbf{Y}, \quad \mathbf{e}_L \cdot \mathbf{Y} = 0.$$

For a locally fixed $\mathbf{e}_L$, the longitudinal mode and both transverse components inherit the same coefficient,

$$Z_L = Z_T = Z_0.$$

Writing $Y_1 = R \cos \vartheta$ and $Y_2 = R \sin \vartheta$ gives

$$(\partial Y_1)^2 + (\partial Y_2)^2 = (\partial R)^2 + R^2 (\partial \vartheta)^2,$$

so the radial occupation mode that mixes with X also carries Z_0 . For each stable Fourier mode before the radial reduction, (Y_i, Π_i) are two canonical oscillator pairs. Their action-angle variables obey $\theta_i \sim \theta_i + 2\pi$ and $d\Pi_1 dY_1 d\Pi_2 dY_2 = dJ_1 d\theta_1 dJ_2 d\theta_2$. Quantization spaces the actions by $\hbar$, giving one state per phase-space volume $(2\pi\hbar)^2$ and the angular Haar volume $(2\pi)^2$ for one action cell. This proves the phase measure for the specified doublet. It does not decide that the physical state loads one sharing entropy into that cell or that the cell couples reversibly to the horizon bath. Spatial variation of $\mathbf{e}_L$ produces derivative connection terms suppressed by the curvature scale of the source direction. Higher-derivative and background-dependent operators can split the coefficients; the equality is a leading long-wavelength result.

One-capacity-invariant potential. After the relative sector of Appendix H is gapped, define the longitudinal and transverse occupations $N_L = \sigma^\dagger P_L \sigma$ and $N_T = \sigma^\dagger P_T \sigma$. The minimal infrared premise is that matter sees one total-capacity invariant,

$$\Delta C = w_L N_L + w_T N_T - N_*.$$

This premise is stronger than tetrahedral symmetry. The general local quartic potential contains

$$V_{\text{gen}} = \mu_L N_L + \mu_T N_T + \frac{\lambda_L}{2} N_L^2 + \lambda_{LT} N_L N_T + \frac{\lambda_T}{2} N_T^2 + \dots$$

Its quartic coupling matrix has rank at most one when

$$\lambda_{LT}^2 = \lambda_L \lambda_T.$$

It has rank one only if the matrix is nonzero. A stable one-invariant potential further requires the matrix to be positive semidefinite, so $\lambda_L, \lambda_T \geq 0$ and not both vanish. Under those conditions $\lambda_L = u_C w_L^2$, $\lambda_T = u_C w_T^2$, and $\lambda_{LT} = u_C w_L w_T$ with $u_C > 0$ and suitable signed weights. The linear terms must align with the same weights. A thermodynamic capacity-pressure variable makes the stable sign explicit through

$$V_{\text{eff}}(\Delta C) = \sup_{\lambda_C} \left[\lambda_C \Delta C - \frac{\chi_C}{2} \lambda_C^2 \right].$$

The stationary pressure $\lambda_C = \Delta C / \chi_C$ gives $V_{\text{eff}} = (\Delta C)^2 / (2\chi_C)$. A path-integral Hubbard–Stratonovich representation requires the contour and Lorentzian/Euclidean sign to be chosen consistently; the equation above is the thermodynamic Legendre form. More generally $V_{\text{eff}} = F(\Delta C)$ with $F'(0) = 0$ and $F''(0) > 0$.

Near a transverse background, $N_L = N_{L0} + \zeta_L X + \dots$ and $N_T = R^2 = Y_1^2 + Y_2^2$, so

$$\Delta C = \alpha_C + g_C X + \beta_C R^2.$$

For fields $\varphi_i = (X, R)$ on the stationary surface $\Delta C = 0$,

$$H_{ij} = \frac{\partial^2 V_{\text{eff}}}{\partial \varphi_i \partial \varphi_j} \Big|_{\Delta C=0} = F''(0) \partial_i \Delta C \partial_j \Delta C.$$

This outer product has rank one wherever $\nabla \Delta C \neq 0$; it has rank zero at a point where that gradient vanishes. Explicitly,

$$H = F''(0) \begin{pmatrix} g_C^2 & 2g_C \beta_C R_0 \\ 2g_C \beta_C R_0 & 4\beta_C^2 R_0^2 \end{pmatrix}, \quad C_\times^2 = A_L A_T.$$

At fixed δX , minimizing the quadratic form gives $\delta R_* = -C_\times \delta X / A_T$. Equal kinetic normalization also defines the clamped frequencies $\omega_L^2 = A_L / Z_0$ and $\omega_T^2 = A_T / Z_0$, so

$$\frac{C_\times}{A_T} = \sqrt{\frac{A_L}{A_T}} = \frac{\omega_L}{\omega_T}.$$

The equality concerns the static response and the clamped curvatures. The freely relaxing Hessian has one zero eigenvalue and one eigenvalue $A_L + A_T$; neither is ω_L or ω_T separately. The soft direction redistributes a fixed total capacity; the orthogonal direction changes the total occupancy. Positive gradient terms stabilize finite-wavelength fluctuations.

This result is protected at the effective level when all local corrections remain functions of the same invariant ΔC . To state the decoupling condition explicitly, define $D = w_T N_L - w_L N_T$ and

$$V(C, D) = F(C) + \frac{\Delta_D}{2} D^2 + \eta_D C D^2 + J_D D + \dots, \quad C \equiv \Delta C.$$

The stationary solution is $D = -J_D/(\Delta_D + 2\eta_D C) + \dots$. For large Δ_D , integrating out D changes $F(C)$ only through constants and powers of Δ_D^{-1} . This is an effective decoupling statement, not proof that the bare vertex is rank one. A standalone $m_X^2 X^2/2$ outside $F(\Delta C)$ would raise the Hessian rank and generate a Yukawa deformation; the X^2 curvature tied to the entries of the rank-one outer product does not.

The susceptibility is not fixed by the normalized 1680-state internal distribution. It depends on fluctuations of total condensate occupation,

$$\chi_C = \langle \hat{C}^2 \rangle - \langle \hat{C} \rangle^2, \quad \hat{C} = w_L \hat{N}_L + w_T \hat{N}_T,$$

in a specified GFT state. In an energy-normalized convention the corresponding susceptibility includes $\beta_H \equiv 1/(k_B T_H)$. For a single mode $\hat{C} = c_0 \hat{N}$, a thermal state gives $\chi_C = c_0^2 \bar{N}(1 + \bar{N})$, while a coherent condensate gives $\chi_C = c_0^2 \bar{N}$. The closure distribution fixes neither state nor $\bar{N}$. The susceptibility controls the massive direction and nonlinear corrections but cancels from the rank-one ratio.

Thermal matching and its limit. An acceleration a carries the Unruh energy $E(a) = \hbar a/(2\pi c)$, while $k_B T_H = \hbar H_0/(2\pi)$, so

$$\frac{E(a)}{k_B T_H} = \frac{a}{c H_0}.$$

The proposed effective identification is

$$A_L = \frac{g_{\text{bar}}}{c H_0}, \quad A_T = \frac{a_0}{c H_0}$$

and therefore $C_\times/A_T = \sqrt{g_{\text{bar}}/a_0}$ after choosing the positive relative orientation. This closes the algebra inside the leading effective model. If the source-driven energy is $E_\perp = \hbar \omega_L$, the desired thermal argument obeys

$$\frac{E_\perp}{k_B T_H} = \frac{C_\times}{A_T} \iff \hbar \omega_T = k_B T_H.$$

Thus the curvature-to-energy step and the transverse thermal normalization are one anchoring condition within the clamped-oscillator reading. A microscopic calculation must still show that the retarded response uses the clamped curvatures rather than the adiabatic eigenmodes. The environmental constancy of $A_T = 4F''(0)\beta_C^2 R_0^2$, the one-entropy loading of the canonical phase cell, and its reversible horizon coupling remain matching conditions.

The metric response is an open-system problem. The natural object is the closed-time-path generating functional obtained by integrating out transverse substrate variables ξ^A in their physical state $\rho_\perp$:

$$e^{i\Gamma_\perp[g_+, g_-]/\hbar} = \text{Tr}_\perp \left(U_\perp[g_+] \rho_\perp U_\perp^\dagger[g_-] \right).$$

The total effective functional is

$$\Gamma_{\text{eff}} = I_0[g_+, \psi_+] - I_0[g_-, \psi_-] + \Gamma_{\perp}[g_+, g_-; \psi_+, \psi_-].$$

The physical metric equation is obtained by varying the difference field and then setting the two histories equal. Unitarity requires

$$\Gamma_{\perp}[g, g] = 0, \quad \Gamma_{\perp}[g_+, g_-]^* = -\Gamma_{\perp}[g_-, g_+].$$

Diffeomorphism invariance of the influence functional must yield the Ward identity that makes the induced response covariantly conserved.

Quadratic data and response convention. Here g_+ and g_- label the two CTP branches. Let $h_{\mu\nu}^{(1,2)}$ be their metric perturbations and define the average and difference fields $h_c = (h^{(1)} + h^{(2)})/2$ and $h_{\Delta} = h^{(1)} - h^{(2)}$. With $x^0 = ct$, the linear metric coupling is

$$\delta I_{\perp} = \frac{1}{2c} \int d^4x \sqrt{-\bar{g}} h_{\mu\nu} \hat{T}_{\perp}^{\mu\nu}.$$

Define the densitized source

$$J_{\mu\nu} \equiv \frac{\sqrt{-\bar{g}}}{2c} h_{\mu\nu}.$$

Let $J^c = (J^{(1)} + J^{(2)})/2$ and $J^{\Delta} = J^{(1)} - J^{(2)}$. Then the universal quadratic form in this source convention is

$$\Gamma_{\perp}^{(2)} = \int d^4x d^4y J_{\mu\nu}^{\Delta}(x) \hat{\Pi}_R^{\mu\nu\alpha\beta}(x, y) J_{\alpha\beta}^c(y) + \frac{i}{2} \int d^4x d^4y J_{\mu\nu}^{\Delta}(x) \hat{\mathcal{N}}^{\mu\nu\alpha\beta}(x, y) J_{\alpha\beta}^{\Delta}(y).$$

The retarded kernel is

$$\boxed{\hat{\Pi}_R^{\mu\nu\alpha\beta}(x, y) = +\frac{i}{\hbar} \theta(x^0 - y^0) \left\langle \left[\hat{T}_{\perp}^{\mu\nu}(x), \hat{T}_{\perp}^{\alpha\beta}(y) \right] \right\rangle_{\rho_{\perp}} + \hat{\Pi}_{\text{local}}^{\mu\nu\alpha\beta}(x, y).}$$

The sign follows from the stated $+\frac{1}{2}h_{\mu\nu}T^{\mu\nu}$ action coupling and the Kubo susceptibility convention. The local or *seagull* term comes from the explicit metric variation of $\hat{T}_{\perp}^{\mu\nu}$ and of the covariant measure. It is required, together with the commutator, for the diffeomorphism Ward identity. Define $\Delta\hat{T}_{\perp}^{\mu\nu} = \hat{T}_{\perp}^{\mu\nu} - \langle \hat{T}_{\perp}^{\mu\nu} \rangle_{\rho_{\perp}}$. The noise kernel is the correspondingly normalized symmetrized correlator,

$$\hat{\mathcal{N}}^{\mu\nu\alpha\beta}(x, y) = \frac{1}{2\hbar} \left\langle \left\{ \Delta\hat{T}_{\perp}^{\mu\nu}(x), \Delta\hat{T}_{\perp}^{\alpha\beta}(y) \right\} \right\rangle_{\rho_{\perp}}.$$

A microscopic calculation begins with the transverse Hessian and metric vertex,

$$\mathcal{K}_{AB}(x, y) = \frac{\delta^2 I_{\perp}}{\delta \xi^A(x) \delta \xi^B(y)}, \quad V_A^{\mu\nu}(x; y) = \frac{\delta^2 I}{\delta \xi^A(x) \delta g_{\mu\nu}(y)}.$$

In the physical state,

$$G_R^{AB}(x, y) = +\frac{i}{\hbar} \theta(x^0 - y^0) \text{Tr}(\rho_{\perp}[\xi^A(x), \xi^B(y)]),$$

and schematically

$$\hat{\Pi}_R^{\mu\nu\alpha\beta} = V_A^{\mu\nu} G_R^{AB} V_B^{\alpha\beta} + \hat{\Pi}_{\text{contact}}^{\mu\nu\alpha\beta}.$$

This is the minimal correlator/Hessian calculation required before an action claim can be made.

Static and lensing targets. In the static scalar sector the kernel produces a two-component response,

$$\begin{pmatrix} \Delta\Phi_{\perp}(\mathbf{k}) \\ \Delta\Psi_{\perp}(\mathbf{k}) \end{pmatrix} = \mathbf{K}_{\perp}(0, \mathbf{k}; \rho_{\perp}) \rho(\mathbf{k}) + \text{nonlinear terms}.$$

The time-time component must reproduce, for spherical baryonic sources,

$$g_{\text{obs}} = \frac{g_{\text{bar}}}{1 - \exp[-\sqrt{g_{\text{bar}}/a_0}]}.$$

The light-deflection potential is proportional to $\Phi + \Psi$, so the spatial components of the same kernel must determine whether

$$\Delta\Phi_{\perp} = \Delta\Psi_{\perp}.$$

The exact Bose factor $1 + n_B$, the argument $\sqrt{g_{\text{bar}}/a_0}$, and the absence of an additional spectral form factor require nonlinear vertices or state dependence beyond the quadratic Hessian. They must emerge from the same calculation; they cannot be imposed on the temporal equation while the spatial response is chosen separately.

Acceptance tests. A transverse completion is acceptable only if:

1. matter couples to one physical metric and the induced response satisfies its Ward identity;
2. the retarded kernel yields the RAR without a galaxy-by-galaxy parameter or an extra kernel $F(x)$;
3. the spatial response predicts lensing and is compatible with the observed dynamical-to-lensing relation;
4. the high-acceleration limit leaves Solar-System PPN values and tensor-wave propagation unchanged;
5. the retarded poles have no ghosts, tachyons, or gradient instabilities, and the noise kernel is positive;
6. any dissipative part obeys the state-appropriate fluctuation–dissipation relation and does not violate causal support;
7. the same calculation explains how momentum transfer is converted into acceleration and how total energy–momentum is conserved.

N.9 Action-level outcome

The action problem now divides cleanly:

ordinary longitudinal capacity = reduced Einstein constraint action,
spherical nonlinear capacity = composite Misner–Sharp invariant,
galactic transverse response = conditional rank-one EFT + open metric influence functional,
capacity-exhaustion boundary = open microscopic boundary functional.

This division solves the original double-counting problem and closes the ordinary weak-field parent action. The transverse EFT now supplies a precise rank-one target. A specified simplicial GFT action must reproduce its projector, one-invariant potential, compact phase structure, and thermal matching; the state and metric vertex must then determine the full $\hat{\Pi}_R^{\mu\nu\alpha\beta}$, including its local Ward-identity term. The boundary-channel functional at $q_{\text{geo}} = 0$ remains separate work.

Appendix O: Reproduction Code

This terminal appendix contains the executable reproduction block after all conceptual, provenance, and action-level appendices. It recomputes the manuscript's numerical spine from the definitions stated in the paper; it does not add a new derivation.

The script below recomputes the numerical consequences of the stated definitions: the 1680-state spectrum, closure solution, sharing entropy, marked vertex, dressed scale, Joyce constant, edge-kernel chain, corrected charged-lepton ratios, and slot mutual information. It also verifies the algebraic normalization identity $n_{\text{hor}} S_{\infty}^{\text{cell}} = 1/4$. CODATA constants are comparison inputs; no observational target is fitted inside the script. It runs under Python 3 with `numpy` and `scipy`. The complete Hessian and graph enumeration are supplied in the accompanying audit script.

```
# =====
# Reproduction script for the Entropic Scalar EFT manuscript.
# Recomputes the full numerical spine from stated definitions:
# the 1680-state K^2 spectrum, eta*, g_share, eff, marked
# vertex, dressed L* and G*, the
# Joyce constant, the 1/4 horizon identity, the edge-kernel
# chain, epsilon, sigma*, the lepton ladder, and the slot MI.
# Dependencies: Python 3, numpy, scipy.
# Every printed value is compared against the manuscript.
# =====
import itertools, math
from math import log, exp, pi, gamma, sqrt
from collections import Counter
import numpy as np
from scipy.optimize import brentq

# 1. Enumerate 1680 states: injective 4-of-7 labels m in -3..3, x2 parity
labels = range(-3,4)
states = list(itertools.permutations(labels,4))
print("P(7,4) =", len(states), " total with parity:", 2*len(states))

def K2(ms):
    S = sum(ms); Sig2 = sum(m*m for m in ms)
    return 48 - (S*S - Sig2)/3

spec = Counter()
for st in states:
    spec[round(K2(st)*3)] += 2 # parity doubling; key = 3*K2 to keep exact
print("Distinct K2 values and multiplicities (K2 as fraction /3):")
for k in sorted(spec): print(f" K2={k}/3 = {k/3:.4f} mult={spec[k]}")
print("Total multiplicity:", sum(spec.values()))

# 2. Solve closure condition <K2> = 3/(2 eta)
Ks = np.array([k/3 for k in sorted(spec)])
ns = np.array([spec[k] for k in sorted(spec)], dtype=float)
def avgK2(eta):
    w = ns*np.exp(-eta*Ks)
    return (w*Ks).sum()/w.sum()
f = lambda eta: avgK2(eta) - 3/(2*eta)
eta_star = brentq(f, 1e-4, 1.0, xtol=1e-16)
print("\neta* =", eta_star, " (paper: 0.02986684439352237)")

w = ns*np.exp(-eta_star*Ks); Z = w.sum(); p_class = w/Z
p_micro = np.exp(-eta_star*Ks)/Z
H = -(ns*p_micro*np.log(p_micro)).sum()
print("g_share, eff =", H, " (paper: 7.419800023570903)")
avg = avgK2(eta_star)
var = (p_class*(Ks-avg)**2).sum()
print("<K2> =", avg, " (paper 50.2229154254) 3/(2eta*) =", 3/(2*eta_star))
```

```

print("Var(K2) =", var, " (paper 15.6889750078)   a_UV = 1/Var =", 1/var)
print("C_cl = eta* <K2> =", eta_star*avg)

# 3. Decorated marked vertex, L*, G*
hbar=1.054571817e-34; c=299792458.0; me=9.1093837139e-31 # CODATA 2022
lam_e = hbar/(me*c)
print("\nlam_e =", lam_e, " (paper 3.861592671986303e-13)")
r = exp(-7*H)
L0 = -1.5*lam_e*math.log1p(-r)
q = 2/7
zeta = 9*exp(-H)*(1-eta_star*8/49)**(21/2)
Zmu = 1+zeta
Ze = (1+zeta)*(1+7*zeta*zeta)
Ztau2 = 1+q*zeta
Lstar = Ze*L0
print("e^{7g} =", exp(7*H), " (paper 3.602860521062804e22)")
print("(2/3)e^{7g} =", (2/3)*exp(7*H), " (paper 2.401907014041869e22)")
print("zeta* =", zeta, " (paper 0.005123584484947)")
print("Z_mu, Z_e, Z_tau2 =", Zmu, Ze, Ztau2)
print("L0 =", L0, " (paper 1.607719470158885e-35)")
print("L* =", Lstar, " (paper 1.616253701392569e-35)")
Gstar = c**3*Lstar**2/hbar
print("G* =", Gstar, " (paper 6.674289077220912e-11)")
G_codata = 6.67430e-11
LP = sqrt(hbar*G_codata/c**3)
print("L* vs Planck length:", (LP-Lstar)/LP*100, "%")
print("G* vs CODATA:", (G_codata-Gstar)/G_codata*100, "%")

# 4. Joyce constant
Gtet = 3*gamma(1/3)**6/(2**(14/3)*pi**4)
print("\nG_tet(0) =", Gtet, " (paper 0.4482203943883814)")
print("ln(7/6)/(pi*Gtet) =", log(7/6)/(pi*Gtet), " (paper 0.109472228)")
print("S_inf_cell = 3ln2/(32 pi Gtet) =", 3*log(2)/(32*pi*Gtet), " (paper 0.0461482516)")
print("n_hor*S_inf_cell =", (8*pi*Gtet/(3*log(2))) * (3*log(2)/(32*pi*Gtet)))

# 5. Edge kernel chain
Jbare = 2/3*eta_star
Jtree = Jbare/3
Sig = 7+2/9
cloop = 1/(1+Jtree*Sig)
Jren = Jtree*cloop
print("\nJ_bare =", Jbare, " (paper 0.0199112296)")
print("J_tree =", Jtree, " (paper 0.0066370765)")
print("c_loop =", cloop, " (paper 0.95426)")
print("J_ren =", Jren, " (paper 0.00633348)")
print("gamma coeff 4 J_ren/(3 pi^2) =", 4*Jren/(3*pi**2), " (paper 8.556e-4)")

# 6. epsilon and cluster numbers
epsv = H/(4*pi**2)
print("\nepsv = g/(4pi^2) =", epsv, " 1-epsv =", 1-epsv, " ceiling =", 2-epsv)
print("sigma* = pi/g =", pi/H, " (paper 0.42340665)")

# 7. Corrected charged-lepton ladder
mmu0 = 720*q; mtau0 = 720**2*q**4
mmu = mmu0*Zmu; mtau = mtau0*Zmu*Ztau2
print("\nbaseline m_mu/m_e =", mmu0, " baseline m_tau/m_e =", mtau0)
print("m_mu/m_e =", mmu, " (paper 206.768280237)")
print("m_tau/m_e =", mtau, " (paper 3477.343310)")
print("alpha_mu =", log(206.7683)/log(mmu), " alpha_tau =", log(3477.23)/log(mtau))

# 8. mixing-time ratio
tau_star = Lstar/c; tau_e = hbar/(me*c**2)
print("\ntau* =", tau_star, " tau_e =", tau_e, " ratio =", tau_e/tau_star)

```



```

# 9. slot mutual information (marginal of slot under admissibility)
joint = Counter()
for st in states:
    wgt = exp(-eta_star*K2(st))
    joint[(st[0],st[1])] += 2*wgt
tot = sum(joint.values())
pj = {k:v/tot for k,v in joint.items()}
p0 = Counter(); p1 = Counter()
for (a,b),v in pj.items(): p0[a]+=v; p1[b]+=v
I = sum(v*log(v/(p0[a]*p1[b]))) for (a,b),v in pj.items()
H0 = -sum(v*log(v) for v in p0.values())
print("\nslot MI =", I, " (paper 0.1545)   H(slot) =", H0, " I/H =", I/H0, " (paper 0.079)")

```

References

- [1] S. S. McGaugh, F. Lelli, and J. M. Schombert, “Radial acceleration relation in rotationally supported galaxies,” *Physical Review Letters* **117**, 201101 (2016).
- [2] D. Clowe et al., “A direct empirical proof of the existence of dark matter,” *The Astrophysical Journal Letters* **648**, L109–L113 (2006).
- [3] B. Bertotti, L. Iess, and P. Tortora, “A test of general relativity using radio links with the Cassini spacecraft,” *Nature* **425**, 374–376 (2003).
- [4] Y. Tian et al., “The radial acceleration relation in CLASH galaxy clusters,” arXiv:2001.08340 (2020).
- [5] A. J. Guttmann, “Lattice Green functions in all dimensions,” *Journal of Physics A: Mathematical and Theoretical* **43**, 305205 (2010).
- [6] G. S. Joyce, “On the cubic lattice Green functions,” *Proceedings of the Royal Society of London A* **445**, 463–477 (1994).
- [7] R. Arnowitt, S. Deser, and C. W. Misner, “The dynamics of general relativity,” in *Gravitation: An Introduction to Current Research*, ed. L. Witten (Wiley, New York, 1962).
- [8] J. D. Bekenstein, “Black holes and entropy,” *Physical Review D* **7**, 2333–2346 (1973).
- [9] S. W. Hawking, “Particle creation by black holes,” *Communications in Mathematical Physics* **43**, 199–220 (1975).
- [10] G. W. Gibbons and S. W. Hawking, “Action integrals and partition functions in quantum gravity,” *Physical Review D* **15**, 2752–2756 (1977).
- [11] T. Regge and J. A. Wheeler, “Stability of a Schwarzschild singularity,” *Physical Review* **108**, 1063–1069 (1957).
- [12] F. J. Zerilli, “Effective potential for even-parity Regge–Wheeler gravitational perturbation equations,” *Physical Review Letters* **24**, 737–738 (1970).
- [13] A. J. R. Sanderson, T. J. Ponman, A. Finoguenov, E. J. Lloyd-Davies, and M. Markevitch, “The Birmingham–CfA cluster scaling project — I. Gas fraction and the M – T_X relation,” *Monthly Notices of the Royal Astronomical Society* **340**, 989–1010 (2003).
- [14] A. Vikhlinin, A. Kravtsov, W. Forman, C. Jones, M. Markevitch, S. S. Murray, and L. Van Speybroeck, “Chandra sample of nearby relaxed galaxy clusters: mass, gas fraction, and mass–temperature relation,” *The Astrophysical Journal* **640**, 691–709 (2006).

- [15] D. Eckert, S. Ettori, E. Pointecouteau, R. F. J. van der Burg, and S. I. Loubser, “The gravitational field of X-COP galaxy clusters,” *Astronomy & Astrophysics* **662**, A123 (2022).
- [16] P. Li, Y. Tian, M. P. Júlio, M. S. Pawlowski, F. Lelli, S. S. McGaugh, J. M. Schombert, J. I. Read, P.-C. Yu, and C.-M. Ko, “Measuring galaxy cluster mass profiles into the low-acceleration regime with galaxy kinematics,” *Astronomy & Astrophysics* **677**, A24 (2023).
- [17] M. Milgrom, “A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis,” *Astrophysical Journal* **270**, 365 (1983).
- [18] J. D. Bekenstein, “Relativistic gravitation theory for the modified Newtonian dynamics paradigm,” *Physical Review D* **70**, 083509 (2004).
- [19] B. Famaey and S. S. McGaugh, “Modified Newtonian dynamics (MOND): observational phenomenology and relativistic extensions,” *Living Reviews in Relativity* **15**, 10 (2012).
- [20] T. Jacobson, “Thermodynamics of spacetime: the Einstein equation of state,” *Physical Review Letters* **75**, 1260 (1995).
- [21] P. Dorau and A. Much, “From quantum relative entropy to the semiclassical Einstein equations,” *Physical Review Letters* **136**, 091602 (2026).
- [22] E. Verlinde, “On the origin of gravity and the laws of Newton,” *Journal of High Energy Physics* **2011**(4), 29 (2011).
- [23] E. Verlinde, “Emergent gravity and the dark universe,” *SciPost Physics* **2**, 016 (2017).
- [24] C. Skordis and T. Złośnik, “New relativistic theory for modified Newtonian dynamics,” *Physical Review Letters* **127**, 161302 (2021).
- [25] J. Collins, A. Perez, D. Sudarsky, L. Urrutia, and H. Vucetich, “Lorentz invariance and quantum gravity: an additional fine-tuning problem?,” *Physical Review Letters* **93**, 191301 (2004).
- [26] L. Bombelli, J. Lee, D. Meyer, and R. D. Sorkin, “Space-time as a causal set,” *Physical Review Letters* **59**, 521 (1987).
- [27] F. Dowker, J. Henson, and R. D. Sorkin, “Quantum gravity phenomenology, Lorentz invariance and discreteness,” *Modern Physics Letters A* **19**, 1829 (2004).
- [28] K.-H. Chae, “Breakdown of the Newton–Einstein standard gravity at low acceleration in internal dynamics of wide binary stars,” *Astrophysical Journal* **952**, 128 (2023).
- [29] I. Banik, C. Pittordis, W. Sutherland, B. Famaey, R. Ibata, S. Mieske, and H. Zhao, “Strong constraints on the gravitational law from Gaia DR3 wide binaries,” *Monthly Notices of the Royal Astronomical Society* **527**, 4573 (2024).
- [30] R. Genzel et al., “Strongly baryon-dominated disk galaxies at the peak of galaxy formation ten billion years ago,” *Nature* **543**, 397 (2017).
- [31] H. Prince and J. Dunkley, “Data compression in cosmology: A compressed likelihood for Planck data,” *Physical Review D* **100**, 083502 (2019).
- [32] B. I. Ciocan, N. F. Bouché, J. Fensch, D. Krajnović, J. Freundlich, H. Desmond, B. Famaey, and R. Tychi, “MUSE-DARK III: The evolution of the radial acceleration relation at intermediate redshifts,” *Astronomy & Astrophysics* **709**, L16 (2026).

- [33] A. A. Văraşteanu, M. J. Jarvis, A. A. Ponomareva, H. Desmond, I. Heywood, T. Yasin, N. Maddox, M. Glowacki, M. Maksymowicz-Maciata, P. E. M. Mancera Piña, and H. Pan, “MIGHTEE-HI: The radial acceleration relation with resolved stellar mass measurements,” *Monthly Notices of the Royal Astronomical Society* **541**, 2366 (2025).
- [34] P. Lang et al., “Falling rotation curves of star-forming galaxies at $0.6 < z < 2.6$,” *Astrophysical Journal* **840**, 92 (2017).
- [35] F. Lelli, S. S. McGaugh, and J. M. Schombert, “SPARC: mass models for 175 disk galaxies with Spitzer photometry and accurate rotation curves,” *Astronomical Journal* **152**, 157 (2016).
- [36] M. M. Brouwer et al., “The weak lensing radial acceleration relation: measuring the dark matter law with KiDS-1000,” *Astronomy & Astrophysics* **650**, A113 (2021).
- [37] W. G. Unruh, “Notes on black-hole evaporation,” *Physical Review D* **14**, 870 (1976).
- [38] R. B. Tully and J. R. Fisher, “A new method of determining distances to galaxies,” *Astronomy & Astrophysics* **54**, 661 (1977).
- [39] S. S. McGaugh, J. M. Schombert, G. D. Bothun, and W. J. G. de Blok, “The baryonic Tully–Fisher relation,” *Astrophysical Journal Letters* **533**, L99 (2000).
- [40] Planck Collaboration, “Planck 2018 results. VI. Cosmological parameters,” *Astronomy & Astrophysics* **641**, A6 (2020).
- [41] A. G. Riess et al., “A comprehensive measurement of the local value of the Hubble constant with $1 \text{ km s}^{-1} \text{ Mpc}^{-1}$ uncertainty from the Hubble Space Telescope and the SH0ES team,” *Astrophysical Journal Letters* **934**, L7 (2022).
- [42] W. L. Freedman, B. F. Madore, I. S. Jang, et al., “Status report on the Chicago–Carnegie Hubble Program (CCHP): measurement of the Hubble constant using the Hubble and James Webb Space Telescopes,” *The Astrophysical Journal* (2025), doi:10.3847/1538-4357/adce78.
- [43] E. Di Valentino et al., “In the realm of the Hubble tension — a review of solutions,” *Classical and Quantum Gravity* **38**, 153001 (2021).
- [44] C. M. Will, “The confrontation between general relativity and experiment,” *Living Reviews in Relativity* **17**, 4 (2014).
- [45] C. Brans and R. H. Dicke, “Mach’s principle and a relativistic theory of gravitation,” *Physical Review* **124**, 925 (1961).
- [46] E. G. Adelberger, B. R. Heckel, and A. E. Nelson, “Tests of the gravitational inverse-square law,” *Annual Review of Nuclear and Particle Science* **53**, 77 (2003).
- [47] T. Padmanabhan, “Thermodynamical aspects of gravity: new insights,” *Reports on Progress in Physics* **73**, 046901 (2010).
- [48] R. H. Sanders and S. S. McGaugh, “Modified Newtonian dynamics as an alternative to dark matter,” *Annual Review of Astronomy and Astrophysics* **40**, 263 (2002).
- [49] D. Oriti, “Group field theory as the second quantization of loop quantum gravity,” *Classical and Quantum Gravity* **33**, 085005 (2016).
- [50] S. Gielen, D. Oriti, and L. Sindoni, “Cosmology from group field theory formalism for quantum gravity,” *Physical Review Letters* **111**, 031301 (2013).

- [51] A. Barbieri, “Quantum tetrahedra and simplicial spin networks,” *Nuclear Physics B* **518**, 714 (1998).
- [52] J. C. Baez and J. W. Barrett, “The quantum tetrahedron in 3 and 4 dimensions,” *Advances in Theoretical and Mathematical Physics* **3**, 815 (1999).
- [53] C. Rovelli and L. Smolin, “Discreteness of area and volume in quantum gravity,” *Nuclear Physics B* **442**, 593 (1995).
- [54] A. M. Gleason, “Measures on the closed subspaces of a Hilbert space,” *Journal of Mathematics and Mechanics* **6**, 885–893 (1957).
- [55] J. B. Hartle, “Decoherent histories quantum mechanics starting with records of what happens,” arXiv:1608.04145 (2016).
- [56] E. T. Jaynes, “The minimum entropy production principle,” *Annual Review of Physical Chemistry* **31**, 579–601 (1980), doi:10.1146/annurev.pc.31.100180.003051.
- [57] W. F. Stinespring, “Positive functions on C^* -algebras,” *Proceedings of the American Mathematical Society* **6**, 211–216 (1955), doi:10.1090/S0002-9939-1955-0069403-4.
- [58] A. D. Wyner and J. Ziv, “Some asymptotic properties of the entropy of a stationary ergodic data source with applications to data compression,” *IEEE Transactions on Information Theory* **35**, 1250–1258 (1989).
- [59] D. S. Ornstein and B. Weiss, “Entropy and data compression schemes,” *IEEE Transactions on Information Theory* **39**, 78–83 (1993).
- [60] J. J. Halliwell, “A review of the decoherent histories approach to quantum mechanics,” arXiv:gr-qc/9407040 (1994).
- [61] D. Eckert et al., “Non-thermal pressure support in X-COP galaxy clusters,” *Astronomy & Astrophysics* **621**, A40 (2019).
- [62] S. Dupourqué, N. Clerc, E. Pointecouteau, D. Eckert, S. Ettori, and F. Vazza, “Investigating the turbulent hot gas in X-COP galaxy clusters,” *Astronomy & Astrophysics* **673**, A91 (2023).
- [63] XRISM Collaboration, “XRISM reveals low nonthermal pressure in the core of the hot, relaxed galaxy cluster Abell 2029,” *Astrophysical Journal Letters* **982**, L5 (2025).
- [64] A. H. Chamseddine and V. Mukhanov, “Mimetic dark matter,” *Journal of High Energy Physics* **1311**, 135 (2013).
- [65] A. Lewis, A. Challinor, and A. Lasenby, “Efficient computation of cosmic microwave background anisotropies in closed Friedmann–Robertson–Walker models,” *Astrophysical Journal* **538**, 473 (2000).
- [66] S. More, B. Diemer, and A. Kravtsov, “The splashback radius as a physical halo boundary and the growth of halo mass,” *Astrophysical Journal* **810**, 36 (2015).
- [67] T. Regge, “General relativity without coordinates,” *Il Nuovo Cimento* **19**, 558–571 (1961).
- [68] J. Ambjørn, J. Jurkiewicz, and R. Loll, “Emergence of a 4D world from causal quantum gravity,” *Physical Review Letters* **93**, 131301 (2004).
- [69] J. Ambjørn, A. Görlich, J. Jurkiewicz, and R. Loll, “Nonperturbative quantum gravity,” *Physics Reports* **519**, 127–210 (2012).

- [70] J. Engle, E. Livine, R. Pereira, and C. Rovelli, “LQG vertex with finite Immirzi parameter,” *Nuclear Physics B* **799**, 136–149 (2008).
- [71] F. Gozzini, “A high-performance code for EPRL spin foam amplitudes,” *Classical and Quantum Gravity* **38**, 225010 (2021).