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

Jacob Chinitz

June 21, 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 free to be adjusted 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, the weak-field coefficients are fixed by counting the possible configurations of that cell. A single chain of calculation, with nothing left to tune galaxy by galaxy, then gives Newton's law, the acceleration scale a_0 at which galaxies begin to depart from Newtonian expectations, the tight observed relation between galaxy rotation and ordinary matter, and the leading no-slip lensing structure.

The electron, the lightest charged particle, plays a double role. It fixes the exchange rate between committed entanglement and mass, and it also fixes the absolute size of the smallest cell. The second step uses Many-Pasts, one of the three founding postulates: the claim that the medium's present state is supported not by a single microscopic past but by many admissible histories, weighted by how well each leads to the present. In the operational branch used here this leaves all ordinary quantum-mechanical predictions intact — Born-rule statistics and no-signaling are preserved — while making the electron's microscopic dressing memoryless, and that memorylessness fixes the substrate length.

The length fixed in this way — from the electron Compton scale and the tetrahedral entropy, with no value of G entering the chain — implies a gravitational scale within about one percent of the measured Newton constant. We treat this as a calibrated coherence check, not as a clean independent prediction, because the construction's form was selected with that target partly in view. Appendix K records that provenance and the associated fork accounting.

Beyond the static weak-field sector, the framework extends to time-dependent transport, clusters, cosmology, the saturated early universe, black holes, and the charged-lepton spectrum, at varying and explicitly labeled levels of closure. The completed claim of the paper is the chain from the microscopic construction to ordinary weak gravity; the later sectors are presented as conditional extensions, frontier completions, or open tests.

Contents

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 $\delta S/S_\infty$.

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, and the answer offered here is conditional. The finite-capacity substrate, the three postulates of Section 3, and the minimal tetrahedral boundary ensemble are theory-defining inputs; the paper does not derive them from a deeper Hamiltonian. Once they are fixed, the static weak-field coefficients are no longer phenomenological parameters but are determined by a finite ultraviolet counting problem: admissibility closure fixes the effective sharing entropy, edge transport and finite-loop dressing fix the stiffness, source projection fixes the coupling, and the weak-field bridge relates fractional capacity deficit to gravitational potential.

The primary closed branch is therefore the static weak-field chain

$$\text{microstructure} \longrightarrow \text{coefficient chain} \longrightarrow \text{continuum EFT} \longrightarrow \{G,\ a_0,\ \text{RAR},\ \text{no slip}\}.$$

It recovers the Newtonian point-source limit, the galactic acceleration scale, the radial-acceleration relation, and leading no-slip lensing without per-system tuning. The absolute substrate length is fixed through the electron anchor and the Many-Pasts memoryless dressing branch; because the construction was developed with Newton's constant partly in view, the one-percent agreement with G is treated as a calibrated coherence check rather than a clean independent prediction, with the provenance recorded in Appendix K.

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. Its history weighting enters the scale-setting branch through memoryless electron dressing, and its operational 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 from no assumptions, nor does it derive the tetrahedral boundary ensemble from a deeper microscopic Hamiltonian in the main text; those are theory-defining ultraviolet inputs, and the closure claim begins only once they are fixed. There are five such inputs: finite local entanglement capacity; matter as localized defects of that capacity; mass–entropy equivalence, so that mass is committed defect entropy read in mass units; the Many-Pasts history weighting of Postulate III; and the minimal tetrahedral boundary ensemble, in its physical realization as the full admissibility-closed sharing entropy, serving as the ultraviolet counting architecture.

Given those inputs, the static weak-field sector closes with no phenomenological freedom remaining. The finite boundary count fixes the admissibility weighting and the effective sharing entropy; a faithful sector-resolution principle fixes the substrate cell length L_*; edge kernel, finite-loop dressing, continuum stiffness, and source projection follow in turn; and the weak-field bridge $\delta S \leftrightarrow \Phi$ then yields the Newtonian limit, the galactic acceleration scale a_0, the radial-acceleration relation, and leading no-slip lensing, with no per-system tuning. This is the central result. Everything beyond the static branch extends the same ontology at a lower closure level and is labeled accordingly: time-dependent transport, clusters, the saturated-phase cosmology, the strong-field branch, and the particle and gauge extensions. Part VI collects this bookkeeping in a closure-status table.

In compressed form, the central claim is

$$\boxed{\text{primitive UV capacity hypothesis}} \to \boxed{\text{finite counting + admissibility}} \to \boxed{L_*}$$
$$\to \boxed{\gamma,\ \kappa/\gamma} \to \boxed{\delta S \leftrightarrow \Phi} \to \boxed{G,\ a_0,\ \text{RAR},\ \text{no slip}}.$$

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

Only the absolute length calibration is deferred rather than defined here. The substrate cell length L_* is fixed by identifying the electron — the lightest clean charged defect — with the maximally extended ground state of a history-space sector-resolution operator, supported over its reduced Compton scale. That principle is itself derived rather than assumed: its content is that each dressing channel carries the full admissibility entropy and the channels are mutually independent. The first holds because the dressing is memoryless in substrate time, which is the Many-Pasts content of Postulate III; the second holds because inter-channel correlation would contract the defect and raise its mass, so the lightest charged defect carries none. The electron's two roles — mass anchor and length anchor — and the explicit formula are developed in Section 13 and derived technically in Appendices D.4 and H. The central technical claim is that, given the three postulates and the realized minimal ensemble, the static weak-field branch has no remaining phenomenological freedom.

The continuum action can superficially resemble an ordinary scalar-tensor theory of the Brans–Dicke lineage [44]. In a conventional scalar-tensor model one asks whether variation of a covariant action containing a scalar can produce the desired metric phenomenology. Here the scalar is the continuum order parameter of an underlying capacity substrate, and the weak-field bridge is the emergent-geometry dictionary relating fractional capacity depletion to the metric lapse. The closure question is therefore not whether the tetrahedral microstructure has been derived from nothing, but whether the specified microscopic capacity ensemble fixes the weak-field gravitational response once it is adopted.

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

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, the entanglement network has many admissible microscopic pasts, and its present state is governed by a probability-weighted ensemble of them, with nearer histories weighted more heavily. This is the Many-Pasts postulate. Its operational branch preserves ordinary Born-rule and no-signaling laboratory statistics, and its memoryless history weighting is the ingredient used in the electron dressing that fixes the substrate length.

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 finite-capacity substrate is the primitive ultraviolet premise; the identification of geometry with capacity, the mass–entropy equivalence, and the Many-Pasts weighting are the three postulates of Section 3; and the tetrahedral ensemble is the minimal ultraviolet architecture. Once they are adopted, the static weak-field sector is claimed to close: the admissibility weighting, the effective sharing entropy $g_{\mathrm{share,eff}}$, the edge kernel, the loop dressing, the continuum stiffness, the source projection, the weak-field bridge, the Newtonian limit, the galactic acceleration scale, the radial-acceleration relation, and the leading no-slip lensing structure are all derived within that construction. The single number the construction produces, $g_{\mathrm{share,eff}} = 7.4198$, is not a fitted parameter but a derived entropy value inside a selected construction: it has no adjustable dial and is reproducible from the stated rules alone, while the historical selection of the construction itself is audited and priced in Appendix K. The sectors beyond the static branch are not hidden adjustments to that closed result; they are labeled separately as conditional extensions, frontier completions, empirical tests, or open tasks.

In one line, the architecture is three postulates, one finite-capacity substrate premise, one minimal ultraviolet ensemble, and a small number of explicitly labeled conditional readings — and from those, a single derived entropy value carries the static weak-field chain.

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 — equal to fifteen decimal places — should be one and the same. Here they are the same quantity seen from two sides: a defect's committed entanglement is at once its resistance to being accelerated and its pull on the surrounding medium, so the equality follows from the construction rather than being imposed by hand. The same identification addresses why gravity is so extraordinarily weak. Rather than inserting a tiny coupling by hand, the framework produces gravity's strength by counting: the induced gravitational scale carries the factor $e^{-14 g_{\mathrm{share,eff}}}$, so a large entropy makes gravity exponentially feeble, and the famous gap between gravity and the other forces becomes a matter of arithmetic rather than fine-tuning. It also gives a concrete substrate reading of gravity: the extended capacity-strain field around a localized deficit. This makes its universality automatic (everything is made of substrate, so everything gravitates), its attractive-only character natural (a deficit draws the medium inward; it does not push), and the galactic acceleration scale a derived threshold rather than an unexplained constant. That galactic scale, $a_0 \approx \epsilon\, cH_0$, ties the onset of anomalous rotation to the cosmic horizon, giving a concrete form to the old Machian suspicion that local inertia should somehow depend on the universe at large. And the tight, nearly scatter-free way a galaxy's rotation tracks its visible matter, the very feature that makes an independent dark particle look so strange, follows directly when a single medium is responding to the ordinary matter, rather than being a separate component that must be arranged to follow it.

A second consequence, in plain terms before its technical form: because the medium never holds still — every instant it is re-drawn from all the ways it could be — a particle is not an object sitting in space but a pattern the medium keeps re-forming against its own restlessness, and the entanglement it commits is a maintained quantity, re-established beat by beat. Two different questions can be asked about such a pattern, and their answers pull apart. One is how busy the medium is keeping the knot bound: it must hold its grip on every strand the knot is tied into at once, and separate holds add up, so this maintenance is large — the medium works hard and continuously to keep the defect in being. The other is how often the knot fully re-completes, every strand falling into alignment at the same instant; because simultaneous coincidences multiply rather than add, that completion is exponentially rare. The particle's mass — which is its rest energy — follows the second question, not the first: it is fixed by how rarely the binding fully re-closes, not by how much work the holding takes. One depth of binding governs both, but holding adds while coinciding multiplies, so the maintenance is large and additive while the recurrence that sets the mass is exponentially small. This is why a particle can be vastly lighter than the natural substrate scale with no small number inserted anywhere: it is not weakly bound but deeply bound, and deep binding makes full re-completion exponentially rare. The electron is light not because its structure is weak but because that structure's coherent recurrence is exponentially delayed — light because it is heavy to hold. Section 22 and Appendix H give this its quantitative form, with binding additive across the sectors and recurrence multiplicative; here it is only the shape of the answer.

A memoryless update rule assumes the least: it is the rule one writes down when one declines to invent hidden internal machinery that carries information forward from tick to tick. It earns its place twice over. The construction that fixes the absolute size of the cell needs each dressing pass of the electron to sample the whole admissible ensemble afresh, and the allowed local single-label dynamics provably cannot do this: they freeze into disconnected sectors that never explore the full space (Appendix D.4), so a genuine refresh is required, not optional. A single universal refresh rate also lets the theory carry a smallest length without running afoul of the experiments that ended earlier discrete-spacetime proposals: the granularity here lives in the capacity, not in a preferred spatial lattice, so gravitational waves and light travel at the same speed and there is no frame-dependent dispersion left to detect.

Many-Pasts holds that the present configuration of the entanglement network is supported not by one definite microscopic past but by a whole ensemble of admissible pasts, each weighted by how naturally it leads to the present. Its first payment is in the quantum puzzles that interpretations of quantum mechanics were invented to handle. The interference in a double-slit experiment is the persistence of the unrecorded alternative histories in that ensemble; making a which-path measurement conditions the ensemble on the new record, and the interference goes away. The correlations of an entangled pair come from weighting the histories of the whole joint system, which reproduces the nonclassical statistics without any signal passing between the two wings of the experiment. Measurement adds no separate collapse law; it lays down a durable record, after which the relevant histories are simply the ones compatible with it. In the branch used here all of this leaves ordinary laboratory predictions exactly intact (Born-rule statistics and no-signaling both hold), and the same weighting supplies the arrow of time. The postulate is more than interpretation because it is also load-bearing in a measured number: the same history weighting that satisfies the Born-rule consistency condition also makes the electron's dressing memoryless, and that memorylessness is one of the two conditions that fix the substrate length and hence the strength of gravity. Elsewhere in physics the interpretation of quantum mechanics is detachable philosophy; here it has a concrete job.

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 K. 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 K 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

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

the observable bridge

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

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

Substrate length scale. The canonical UV cell length is not taken to be the conventional Planck length as an input. It is fixed later, through the electron anchor and the Many-Pasts memoryless dressing of Postulate III (Section 13, Appendix D.4); here we only record the canonical definition that the later derivation populates. The defining relation is the ground-state faithful sector-resolution principle

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

Equivalently,

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

The corresponding induced gravitational scale is

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

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

The principal coefficients and derived quantities used throughout are:

$$\begin{align}
\gamma &: \text{ entanglement-field stiffness,} \tag{1}\\
\kappa &: \text{ defect–entropy coupling,} \tag{2}\\
\kappa_m(\ell) &: \text{ mass-per-entropy map at scale } \ell, \tag{3}\\
L_* &: \text{ substrate cell length fixed by faithful sector resolution,} \tag{4}\\
G_* &: \text{ gravitational scale induced by } L_*, \tag{5}\\
g_{\text{share,max}} &= \ln(1680), \tag{6}\\
g_{\text{share,eff}} &: \text{ admissibility-weighted effective sharing entropy,} \tag{7}\\
J_{\text{bare}},\, J_{\text{eff}}^{\text{tree}},\, J_{\text{eff}}^{(\text{ren})} &: \text{ UV edge-kernel couplings,} \tag{8}\\
a_0 &= \frac{cH_0 g_{\text{share,eff}}}{4\pi^2}. \tag{9}
\end{align}$$

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 rests on exactly three postulates because they answer three different questions, and no fewer suffice to state the ontology. *What is spacetime geometry?* — Postulate I: geometry is the long-wavelength expression of the capacity substrate itself. *What is matter, and what is mass?* — Postulate II: matter is localized defect capacity, and mass is the defect's committed entropy in mass units. *What is quantum probability, and what fixes the direction of history?* — Postulate III: quantum probability is a weighting over admissible past histories of the entanglement network. The first two define the gravitational ontology; the third defines its quantum and historical structure and, through memoryless dressing, also enters the scale-setting branch. The main text treats all three as theory-defining inputs, not derived outputs.

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 the entropy units leaves gravity unchanged (Section 2). Second, because geometry and the scalar are projections of the same substrate, gradients and deficits of the entanglement field are not metaphorical contributions to gravity — they belong in the gravitational bookkeeping as a matter of what the geometry *is*, not as an extra source added to it. In the EFT this is realized concretely: the scalar enters a covariant action, contributes its own stress-energy, and couples to a trace-equivalent defect source written at continuum scale in the usual stress-energy variables.

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 is one of the framework's three defining inputs. It concerns what the present entanglement network is supported by. The claim is that the present is not the endpoint of one microscopic past. For any present coarse configuration, many admissible microscopic histories of the network could have produced it, and the physical state is governed by a probability-weighted ensemble over those histories rather than by a single distinguished one.

Three objects make this precise:

- H: an admissible microscopic history of the entanglement network — one self-consistent way the present could have been reached;
- P: the present coarse configuration that all admissible histories must reproduce;
- $D(H, P)$: a history-space distance measuring how far a given history lies from the present state.

In canonical closed form the operational history weight is

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

so histories close to the present dominate and distant ones are exponentially suppressed. This is the branch $\alpha = 1, \beta = 0$ of a generalized weighting family.

This postulate occupies the conceptual role usually played by an interpretation of quantum mechanics. It is not many-worlds: reality does not branch forward into many co-real macroscopic outcomes. It is not a collapse theory: there is no added physical collapse event. It is not hidden-variable ignorance: probability is not ignorance over pre-existing classical values. It is a history-space ontology — quantum probability read as a weighting over admissible past histories of the substrate. This is the framework's alternative to those accounts: ontologically central, though its laboratory predictions are deliberately conservative.

The operational branch is then fixed by two constraints. The Born rule is recovered: requiring exact Born statistics in the projective laboratory limit selects $\alpha = 1$. No-signaling is preserved: forbidding any operational bias channel that could transmit information selects $\beta = 0$. The substantive content is that the family admits a single branch satisfying both at once. Born recovery is therefore stated honestly as a consistency condition the weighting must satisfy, not as a fresh derivation of the rule from outside it; this differs from programs that obtain the rule from noncontextuality [53], decision theory [54], or envariance [55], with no-signaling [56] the companion constraint. Beyond these, the same weight orients the arrow of time through conditional typicality and makes the dressing of an elementary defect memoryless in substrate time, each pass an independent draw from the admissibility ensemble.

That memorylessness is Postulate III's entry into the central result. The coefficient chain that fixes the Newtonian limit, the galactic law, and lensing does not call on Many-Pasts. The substrate-scale branch does: the memoryless dressing is one of the two ingredients that fix the absolute substrate length in Appendix D.4, the other being the mutual independence of the electron's seven dressing channels, which follows from its identity as the lightest charged defect (Appendix H). Its operational consequences and the arrow of time are developed in Section 21 and Appendix G.

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

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

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 [24], 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_*, $g_{\mathrm{share,eff}}$, η_*). Discreteness of this class is compatible with exact low-energy Lorentz symmetry, as causal-set sprinkling demonstrates by construction [25, 26]. 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 of the theory begins with the three foundational postulates — Information–Geometry, Mass–Entropy, and Many-Pasts — on equal footing, and runs through the static weak-field chain before reaching the conditional sectors:

$$\begin{gathered}\{\text{Info–Geometry},\ \text{Mass–Entropy},\ \text{Many-Pasts}\}\\ \rightarrow\ \text{finite-capacity substrate ontology}\ \rightarrow\ \text{tetrahedral boundary ensemble}\\ \rightarrow\ \text{admissibility closure}\ \rightarrow\ \text{edge kernel / loop dressing / stiffness / source map}\\ \rightarrow\ \text{weak-field EFT}\ \rightarrow\ \{\text{Newton},\ a_0,\ \text{RAR},\ \text{no-slip lensing}\}.\end{gathered}$$

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.

Two features of this map matter. First, it is a dependency graph, not an epistemic-equality graph: the static weak-field sector, the UV coefficient chain that feeds it, and the operational Born-recovery branch are more tightly closed than the cosmological or strong-field sectors, and Part VI makes that difference explicit in a closure-status table. Second, Many-Pasts appears among the foundational postulates rather than downstream, because its history weighting is already active in the scale-setting branch through memoryless dressing. The later Many-Pasts section (Section 21) develops the operational consequences of a postulate that has already done load-bearing work upstream.

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.

Several of the ultraviolet choices below may look at first like independent tunings: the tetrahedral boundary cell, the seven face labels, the ordered injective assignment, the parity doubling, the admissibility kernel, the transverse (2/3) export, and the seven-sector electron anchor. None is a phenomenological knob in the closed branch. Section 5 and Appendix B derive the minimal tetrahedral seven-state ensemble and its unique admissibility-closed entropy; Appendix C fixes the transverse export from the tetrahedral bond-frame Green response; Appendices D and H derive the faithful sector-resolution length relation from the electron anchor and the memoryless dressing; and Appendix K records the historical fork accounting and calibration status. The line to keep in view is between the primitive adoption of the finite-capacity tetrahedral substrate and the downstream coefficient chain, which contains no continuous fit once that substrate architecture is fixed. What the construction does not yet exhibit is a single microscopic action or path-integral measure whose continuum expansion generates these coefficients simultaneously; each is derived through its own local reduction, so the chain is a sequence of forced steps rather than one unified continuum limit, and consolidating it into a single measure is an open task.

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

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

Two cells sharing a face therefore generate the effective boundary sector

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

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

$$|M| \geq 4 \quad \Longrightarrow \quad 4j_0 + 1 \geq 4.$$

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

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

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 K. 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 [50, 51], whose discrete geometric spectra [52] 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

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

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

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

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

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

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,

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

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.

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

and

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

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

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

9. Continuum Stiffness and SI Normalization

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

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

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

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

Up to this point the derivation has determined a dimensionless lattice weighting. The continuum EFT, however, needs a dimensionful coefficient multiplying derivatives of a field in spacetime. The Euclidean-action interpretation 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 convention already assumed by the EFT, the SI-normalized weak-field stiffness coefficient is fixed rather than left schematic.

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

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

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. Covariant Action

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

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

with

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

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

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

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

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

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

the Green-matched continuum projection is

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

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

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

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

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

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

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

Three roles of G in the manuscript. There are three distinct uses of G in what follows. First, the weak-field metric normalization gives the identity

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

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

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

and therefore induces

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

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

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

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

Definition. Faithful sector resolution fixes the substrate cell length L_* without using G, and the induced gravitational scale $G_* = c^3 L_*^2/\hbar$ follows from the electron Compton length and the tetrahedral entropy alone. The chain takes no value of G as input.

Comparison. Independently of that length, the scalar bridge and Green-matched source map fix the weak-field normalization through the invariant ratio $\kappa/(\gamma S_\infty)$,

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

which lets G be compared with G_*. Because both are built from the single substrate scale L_*, this is not a second measurement from independent inputs; it is a coherence check of whether that one scale propagates through the scalar stiffness, the source map, and the bridge to the observed normalization. The check has content because fixing L_* does not by itself guarantee that κ/γ, S_∞, and the bridge organize around it to give the right normalization.

Status. The two agree to about one percent. We read this as a calibrated coherence check, not a clean independent prediction, because the construction was developed with the measured G partly in view; Appendix K records that provenance and the associated fork accounting. The dressing dynamics behind faithful sector resolution is derived in Section 22 and Appendix H.

This is the simplest covariant realization of the closure chain: one metric, one scalar entanglement field, one trace-equivalent defect-source channel, and one renormalized background branch.

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

11. Field Equations and Bridge Law

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

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

Varying with respect to the metric yields

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

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

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

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

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

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

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

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

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

fixes $C = 1/2$. Hence

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

and therefore

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

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

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

The bridge law connects the entropy description to observable gravity: the deficit field acquires a unique weak-field normalization in terms of the ordinary gravitational potential.

Matter enters once. The two equations are not two independent ways for ordinary matter to source the weak-field potential, and reading them as such would double-count it. Matter enters the weak-field branch through a single channel: the trace $\chi = -T^\mu{}_\mu/c^2$ sources the scalar equation, the scalar response δS is converted to the lapse by the bridge law, and the observed Newtonian potential is that bridged lapse. The metric equation is the covariant image of the same response — $T^{(\mathrm{ent})}_{\mu\nu}$ is the stress-energy of the scalar field that χ has already sourced, not a separate matter contribution — so $G = c^2\kappa/(8\pi\gamma S_\infty)$ and the Newtonian limit of $G_{\mu\nu} = (8\pi G/c^4)(T^{(\mathrm{matter})}_{\mu\nu} + T^{(\mathrm{ent})}_{\mu\nu})$ are the same field written two ways. Ordinary matter is therefore not a second, independent source of Φ standing beside the bridge; it is the source of the scalar response that the bridge reads as the potential, and Bianchi consistency enforces matter conservation through that same trace channel rather than through a separate coupling. The Einstein equation above is thus written in the already-bridged physical-metric frame — $T^{(\mathrm{ent})}_{\mu\nu}$ records the scalar response the bridge has converted into the lapse — and not as an independent second scalar force acting on test bodies.

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.

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

The electron does two separate jobs in this framework, and keeping them apart turns the weak-field normalization into a genuine consistency check rather than a single calibration. As the lightest one-bit fermionic defect, it is the cleanest *mass* anchor for the mass–entropy map. As the lightest coherent defect whose dressing resolves the seven face sectors, its Compton scale is the *length* anchor that fixes the substrate cell size. The two uses draw on the same particle but on different facts about it — its mass on one side, its coherent dressing support on the other.

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

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

with canonical running law

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

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 fixes the absolute cell length through a different fact about it. Its reduced Compton wavelength λ_e is the coherent support of the lightest defect's ground-state dressing, and that dressing resolves the seven face sectors of the tetrahedral alphabet. Two properties make the resolution clean. The Many-Pasts weighting of Postulate III makes each dressing pass memoryless in substrate time, so each pass carries the full admissibility entropy $g_{\text{share,eff}}$; and the electron, as the lightest charged defect, carries no inter-channel correlation, so the seven sectors are resolved independently. Seven independent resolutions add, giving the support exponent $7g_{\text{share,eff}}$ and hence

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

with the $\ln(3/2)$ the transverse export factor of Appendix C.5.

This does not mean an electron is a single tetrahedral cell carrying seven simultaneous labels. The local ensemble supplies the seven-sector alphabet; the electron is the lightest coherent one-bit defect whose ground-state dressing resolves those sectors once and exports the transverse share of the resulting support. 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 Consistency checks

Two cross-checks test the dual role rather than feed it. The first is the support exponent itself. The length relation reads the physical ratio as the *first* power of the dressing support, $\lambda/L_* \propto e^H$ with $H = 7g_{\text{share,eff}}$ for the electron, and that linear power is not imposed. Letting it float as $\lambda/L_* \propto e^{\alpha H}$ and demanding the observed Newton constant from the electron branch requires $\alpha = 0.9999$, while the same map applied to the charged-lepton ladder (Appendix I.1) is independently solved by $\alpha_\mu = 1.001$ and $\alpha_\tau = 1.001$. A square-root map ($\alpha = \frac{1}{2}$) would misplace m_μ/m_e near 14 and a quadratic map ($\alpha = 2$) near 4×10^4, so the spectrum excludes them by orders of magnitude. Three determinations across two sectors that share no gravitational input bracket $\alpha = 1$ within 10^{-3}. The dictionary itself is derived in Appendix D.4 from the recurrence structure of the dressing: the waiting time for joint typicality of independent channels is the exponential of the summed entropies, so the linear power is the arithmetic of joint probability. The bracketing is the quantitative check of that derivation; identifying one phase cycle per dressing recurrence is the single assumption that remains beneath it.

The second cross-check is the induced gravitational scale. The length anchor fixes L_*, and hence $G_* = c^3 L_*^2/\hbar$, from the electron Compton length and the tetrahedral entropy alone, with no gravitational input. Whether that one substrate scale also propagates through the scalar stiffness, the Green-matched source map, and the weak-field bridge to the observed normalization is a coherence test, not a second independent measurement (Section 10). The comparison agrees to about one percent, and Appendix K records how the measured value of G figured in selecting the construction.

13.5 Composite sectors

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

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

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 Newtonian branch already gives the ordinary baryonic force. The galactic excess appears when, at low acceleration, the transverse vacuum-capacity mode adds a bosonic emission factor on top of that force, and the scale at which this turns on is set by horizon thermality reduced by the microstructural sharing factor. The rest of the section is the precise version of those two sentences, and it separates cleanly into two questions: where the turn-on scale a_0 comes from, and why the resulting radial-acceleration relation has the shape it does.

The galactic sector is one of the main payoffs of the coefficient chain. Its characteristic acceleration scale, the radial-acceleration relation, and the deep-MOND limit all follow from one chain that uses no per-galaxy interpolation function. The chain has three structural inputs: a stable bosonic vacuum mode (Appendix D.6), the Gibbons–Hawking thermal structure of that vacuum at the apparent horizon (Appendix E), and the $1+2$ channel decomposition that places the cosmic horizon scale in the transverse sector.

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

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

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

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

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

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

For the apparent-horizon vacuum this gives

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

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

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

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

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

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

Hence a_0 is the horizon thermal acceleration $a_H = cH_0$ reduced to the transverse scalar sector of the $1+2$ decomposition. The reduction factor $g_{\text{share,eff}}/(4\pi^2)$ is fixed, not fitted. Its $(2\pi)^2$ is the canonical two-dimensional mode density $d^2k_\perp/(2\pi)^2$ — the number of states per phase-space cell of a continuum two-plane, a quantization normalization rather than a Fourier convention — and $g_{\text{share,eff}}$ is the admissibility-sharing content carried per cell. The normalization is $(2\pi)^2$, not the spherical 4π, because the weak-field response samples the local transverse two-plane orthogonal to the baryonic forcing, not the full horizon sphere. That the split exists at all is kinematic, not a modeling choice: the baryonic gradient picks the longitudinal direction $\hat{n} = \nabla\Phi_{\text{bar}}/|\nabla\Phi_{\text{bar}}|$, and its orthogonal complement in three-space is two-dimensional by dimension count, so the transverse two-plane and its $(2\pi)^2$ cell are forced. The one premise that is physical rather than geometric is the assignment of the low-acceleration response to the *stationary* transverse vacuum branch rather than the telegrapher transport sector of Section 17 — the step that ties a_0 to cH_0 rather than to a transport timescale, and that is testable through the epoch dependence $a_0(z) \propto H(z)$. Given it, the coefficient follows, and a_0 and cH_0 share one origin.

$1 + 2$ **channel decomposition and the radial-acceleration law.** The longitudinal acceleration scale is g_{bar} and the transverse vacuum reference scale is a_0. Each channel carries an acceleration frequency through the Unruh correspondence [36],

$$\omega_\parallel \propto \frac{g_{\text{bar}}}{c}, \qquad \omega_\perp \propto \frac{a_0}{c},$$

with a common prefactor that cancels in the occupancy argument below. The two channels mix, and a coupled pair of modes softens at the geometric mean of its two frequencies rather than their average; that softened mixed mode dominates the low-acceleration response. Precisely: the action is quadratic, so the two channels couple bilinearly, and the bilinear form fixes the cross-mode scale. For unit-normalized modes with stiffnesses $\omega_\parallel^2$ and $\omega_\perp^2$ and cross coupling λ, the determinant of the form is $\omega_\parallel^2\omega_\perp^2 - \lambda^2$, so the mixed mode softens, and dominates the response, at the threshold scale

$$\omega_\times = \sqrt{\omega_\parallel \omega_\perp}.$$

The geometric mean is the spectral invariant of the bilinear cross term itself. Other exchange-symmetric combinations correspond to no invariant of the quadratic form; an arithmetic-mean scale would require a non-bilinear vertex the action does not contain. The cross-scale amplitude is therefore

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

The bath temperature is set by the transverse vacuum scale, $k_B T_{\text{bath}} = \hbar\,\omega_\perp$ in the normalization of Appendix E, and the occupancy argument is the frequency ratio

$$x = \frac{\hbar\,\omega_\times}{k_B T_{\text{bath}}} = \frac{a_{\text{eff}}}{a_0} = \sqrt{\frac{g_{\text{bar}}}{a_0}},$$

invariant under common rescaling of the frequencies and the temperature, so no Unruh normalization convention survives in the law. When $g_{\text{bar}} \gg a_0$ the system reduces to the ordinary Newtonian branch; when $g_{\text{bar}} \ll a_0$ the response crosses into the low-acceleration completion.

The resulting radial-acceleration law is

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

with the asymptotic limits

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

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

The two endpoint limits follow from the occupancy factor alone; the exact interpolation between them is conditional on kernel neutrality, $F(x) = 1$, established structurally below and flagged there as the open gold-standard vertex computation. The deep-MOND branch gives the baryonic Tully–Fisher law [37, 38]

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

Division of labor between the action and the state. The field equation of the weak-field sector is linear and remains linear here: it carries propagation and the Newtonian channel. The nonlinearity of the galactic law enters through the statistical state of the transverse vacuum, in the same way that Fermi–Dirac occupancy makes electrical conduction nonlinear in temperature while the Maxwell sector stays linear. The coefficient chain fixes the value of a_0; the bilinear form fixes the argument x; the occupancy enters as the bosonic emission factor

$$g_{\text{obs}} = g_{\text{bar}}\big(1 + n_B(x)\big).$$

The factor is a rate statement, and the distinction carries the physics. Static linear response of a Gaussian theory is temperature-independent — the retarded propagator does not depend on the state — so no static susceptibility can bring the occupancy into the force law. The occupancy enters because the weak-field force here is the rate of a capacity-transfer process between the channels; the entropic ontology requires exactly this, since the force is a flux of capacity and fluxes are rates.

The rate factor then follows from first principles in three steps. First, for linear coupling to a bosonic mode the emission rate into occupation n is exact,

$$\Gamma = \Gamma_0 \left|\langle n{+}1|a^\dagger|n\rangle\right|^2 = \Gamma_0\,(1 + n),$$

which on the thermal cross mode reads $\Gamma_0(1{+}n_B(x))$. Second, the absorption channel that would restore detailed balance and cancel the stimulated term against $\Gamma_0\, n_B$ is structurally closed: the reverse process requires the defect to partially refill its capacity deficit, and that deficit is the quantized one-bit increment $\Delta S_f = \ln 2$ of Section 5, which admits no partially refilled final state. The thermodynamic arrow of Part V suppresses the remaining collective reverse fluctuations by the same conditional-typicality weighting that orients macroscopic entropy flow; the one-bit quantization that sets the substrate length thus also permits a low-acceleration excess to survive at all. Third, normalizing the bare rate on the Newtonian channel, $\Gamma_0 \leftrightarrow g_{\text{bar}}$, gives the law above with its two limits, since $1 + n_B(x) = (1 - e^{-x})^{-1}$. The identification beneath this derivation—that the static weak-field acceleration is the steady capacity-transfer rate—itself follows from the transport sector, in two parts.

The flux identification from transport. The static limit of the telegrapher dynamics of Section 17 is a conservation law: with capacity current $J = -D\nabla\,\delta S$ and a static source, continuity reduces to $\nabla \cdot J = \sigma$, so the field around a mass is a steady conserved capacity current with radial flux density $J(r) \propto M/r^2$. Because δS is dimensionless capacity, J is a *bit-number* current, and its static density carries the Green-matched normalization of the weak-field chain with no new constant: the Newtonian field is a rate density, a statement the static transport limit derives rather than assumes. Delivery of this flux to a test defect is an *inter-channel* process—the current resides in the longitudinal channel while the defect's dressing occupies the transverse sector through the 2/3 export—and inter-channel transfer is generated solely by the bilinear cross term, so it proceeds at the cross-mode frequency, inside the thermally occupied transverse vacuum, at the field point. The stimulation is therefore local: the argument x is evaluated where the transfer is delivered, the feature the wide-binary discriminator of Section 24.1 rests on.

Kernel neutrality. What remains to verify is that the golden-rule kernel contributes no dependence on x beyond the occupancy, i.e. that $\Gamma = g_{\mathrm{bar}}\, F(x)\, (1 + n_B(x))$ has $F \equiv 1$. Four observations isolate the condition the kernel must meet. The factor $1 + n$ is exact ladder algebra, $|\langle n{+}1|b^\dagger|n\rangle|^2$, common to all other kernel factors. A continuum density of states $\rho(\omega) \propto \omega^2$ would deform the law, but the transfer is into the per-cell transverse channel, whose state count per admissibility cell is fixed by the same $4\pi^2$ cell structure that fixes the coefficient of a_0; no frequency power enters. An energy-per-quantum conversion $\hbar\omega_\times$ would likewise deform it, but capacity is exchanged in the fixed one-bit unit $\Delta S_f = \ln 2$, the same quantization that closes the absorption channel, so the transfer is number-counted. The number flux itself is the conserved bit current above, whose static density is g_{bar}. Together these reduce the kernel to a single requirement: that the coarse inter-channel vertex couple to the admissibility-weighted ensemble *mean* of the transition availability rather than to its spectral spread. Under that mean-coupling condition $F \equiv 1$ and $g_{\mathrm{obs}} = g_{\mathrm{bar}}(1 + n_B(x))$ exactly, and the structures enforcing it—the admissibility-cell channel and the one-bit quantization—are the same two facts that fix a_0 and forbid absorption, so the law's shape is tied to the framework's founding bookkeeping rather than to a free choice.

The main text establishes this mean-coupling condition structurally; its verification from the substrate Hamiltonian remains the audit task. A first ensemble-level computation shows why that check is nontrivial: one-step admissible transitions across the 1680-state ensemble correlate with K^2 (Pearson $r \simeq 0.5$, class means spanning 1 to 6 of 8 moves), so the raw hopping kernel is not spectrally flat, and neutrality requires the coarse vertex to see only the ensemble mean of this availability—a single number—with the static field unable to re-weight the local ensemble and leak the spectral dependence into the kernel. Establishing that mean-coupling from the graph-level inter-channel vertex of Appendix H is the remaining gold-standard computation. Empirically, any residual F is bounded below the few-percent level across five decades of x by solar-system precision and the measured flatness of the galactic and lensing relations, and the closure table records the status accordingly.

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

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

15. Lensing, PPN, and Weak-Field Consistency

A modified-gravity account can reproduce galaxy rotation curves and still fail: if the field supplying the extra dynamics does not bend light the same way ordinary mass does, the two metric potentials slip apart and lensing comes out wrong. The weak-field branch here avoids that failure at leading order, and for a structural reason. Because the entanglement sector is scalar, it does not generate anisotropic stress at leading weak-field order. The scalar anisotropic stress is quadratic in gradients, schematically $\partial_i S\, \partial_j S = O(\Phi^2/c^4)$, so it enters beyond the linear branch. Hence

$$\Phi = \Psi$$

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

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

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

The same response that governs the dynamics governs light deflection, so galactic support does not come at the cost of a leading weak-field inconsistency. Solar-system constraints such as Cassini require any PPN slip to remain at the $\sim 10^{-5}$ level, consistent with the no-slip structure of the linear branch.

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

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

Scalar-induced corrections to the slip $\Phi - \Psi$ enter only at quadratic weak-field order, schematically $O(\Phi^2/c^4)$. Thus the leading weak-field EFT reproduces galactic phenomenology without introducing gravitational slip or solar-system-scale pathologies.

The closed prediction here is the absence of leading gravitational slip, $\Phi = \Psi$ with $\gamma_{\text{PPN}} = \beta_{\text{PPN}} = 1$. The corresponding open test is empirical: stacked weak-lensing measurements of the dynamical-to-lensing mass ratio, together with PPN and fifth-force searches, confront that prediction directly. At leading weak-field order the sector is closed; higher-order precision confrontation is an audit task, not an architectural gap.

The trace coupling is not a fifth force. The action's trace term $-\kappa\chi S_{\text{ent}}$ could be read as a Brans–Dicke-like scalar charge that test bodies feel in addition to the metric, which solar-system fifth-force bounds would then constrain severely. That is not the operative reading. The trace channel sources the capacity deficit δS; the bridge law converts that deficit into the lapse $N = \exp[-\delta S/(2S_\infty)]$; and test bodies follow geodesics of the resulting physical metric. There is no second, independent scalar force beside geodesic motion — the scalar's entire effect is already carried by the metric the bridge produces. The leading post-Newtonian values follow from that single bridged metric: writing $U = -\Phi > 0$ and $N = e^{-U/c^2}$ gives $g_{00} = -N^2 = -1 + 2U/c^2 - 2U^2/c^4 + O(c^{-6})$, so $\beta_{\text{PPN}} = 1$, while the absence of leading scalar anisotropic stress gives $\gamma_{\text{PPN}} = 1$. The remaining formal audit is higher-order: if the canonical $T^{(\text{ent})}_{\mu\nu}$ is retained as an independent Einstein-sector source, one should check explicitly that its quadratic pieces reproduce this bridge expansion rather than adding to it. That confrontation is the flagged audit task; the leading calculation returns the general-relativistic values.

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:

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

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,

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

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 only through the transverse projection ϵ already fixed by the galactic sector; a dynamically decoupled collisionless component recovers the full projection; and a virialized coherent bath can be lifted above it. The rest of the section develops that one distinction and confronts it with data, and its status is kept explicit throughout: the projection coefficient is derived, the trend direction and hook shape are supported, the lift amplitude is not yet derived, and the decisive resolved-map test has not yet been performed.

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 ($g_{\mathrm{bar}} \sim 10^{-11}$–$10^{-10}\,\mathrm{m\,s^{-2}}$), 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.

The mechanism is not transport lag and it is not extra galaxy mass; both are excluded. Transport lag is excluded because, in the canonical $\tau_0^{-1} = H_0$ branch, disturbances propagate at c and communicate changes across a cluster-scale configuration on a light-crossing time of order 1 Myr, three orders below the merger timescale, so the field tracks the moving source; an underdamped configuration retains memory of the *pre-merger centroid*, not of the outgoing collisionless component, and a retarded forward solution sourced equally by all baryons therefore does not by itself keep the lensing peak on the outgoing collisionless component. Extra galaxy mass is excluded because the galactic acceleration scale and the gas-dominated-dwarf RAR fix the deficit per unit baryonic mass with no freedom to enhance it. The viable mechanism is instead a phase-dependent *source projection* for the long-range entropic-excess channel, built from a distinction the framework already contains.

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

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

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

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

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

The diffuse and decoupled source classes. The assignment follows from 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 $\epsilon = g_{\rm share,eff}/4\pi^2$.

A *virialized* bath is the third state. Once the diffuse atmosphere relaxes into a coherent, extended phase it is no longer a sum of uncorrelated elements, and it opens a collective long-range response that lifts it *above* the baseline, with a bath weight $W_{\rm bath} > 1$ that grows with how completely the atmosphere has virialized, gated by $B_{\rm bath} \in [0, 1]$ below. This super-baseline lift is a *distinct* premise from the suppression rule above—suppression and its absence span only $[\epsilon, 1]$. The lift carries a derived ceiling: the transverse channel holds $4\pi^2/g_{\rm share,eff} = 1/\epsilon$ of capacity per admissibility cell relative to a single incoherent element, so no source weight can exceed $W_{\rm bath} = 1/\epsilon \simeq 5.32$—the reciprocal of the derived suppression coefficient, and not a new constant. A fully developed coherent bath saturates this bound, occupying every transverse cell of the channel it is otherwise suppressed within. The lift function therefore has fixed endpoints, $W_{\rm bath} = 1$ at zero bath and $W_{\rm bath} = 1/\epsilon$ at saturation; the minimal linear candidate $W_{\rm bath} = 1 + (1 - \epsilon)B_{\rm bath}$ and the growth profile between the endpoints are evaluated against cluster data below.

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

$$\boxed{\chi_{\rm rel} = \rho_{\rm dec} + W_{\rm bath}\,\rho_{\rm bath}}$$

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

$$\boxed{\chi_{\rm merge} = \rho_{\rm dec} + \epsilon\,\rho_{\rm shock} + W_{\rm bath}\,\rho_{\rm vir}}$$

which in the Bullet limit, where the displaced gas is shocked and little virialized bath remains on the cores, reduces to $\chi_{\rm Bullet} \simeq \rho_{\rm dec} + \epsilon\, \rho_{\rm shock}$. 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,

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

with $f_{b,\rm cos}$ fixed by Planck values [39] and $M_{\rm hot,vir}$, M_{500} read from X-ray/SZ and total-mass estimates. The bookkeeping of the sector is then: ϵ is fixed and derived, $f_{b,\rm cos}$ is fixed cosmologically, and the per-system inputs ($f_\star$, $f_{\rm gas}$, $M_{\rm hot,vir}$, M_{500}) are all measured; the one undetermined object is the coherence-growth profile of the lift, $W_{\rm bath}(B_{\rm bath}) \in [1, 1/\epsilon]$, whose endpoints are fixed and whose interior form is not yet derived. In merging systems $M_{\rm hot,vir}$ is to be evaluated from the *pre-merger* virialized atmosphere rather than the mid-collision gas state, since the shocked, displaced gas is not a relaxed continuum.

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

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

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

$$\boxed{\mathcal{R}_{\rm rel} = 1 + \big(W_{\rm bath} - 1\big)\, f_{\rm cont}}$$

and the minimal linear candidate for the lift,

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

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 ($B_{\rm bath}$) and how much of the baryon budget sits in it ($f_{\rm cont}$).

This linear form has the qualitatively correct mass trend: both $B_{\rm bath}$ and $f_{\rm cont}$ *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 $B_{\rm bath} \simeq 0.83$, $f_{\rm cont} \simeq 0.9$, hence $\mathcal{R}_{\rm rel} \simeq 1.6$; 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 $g_{\rm bar} = 10^{-10}\,{\rm m\,s^{-2}}$ rising to $\mathcal{R} \simeq 7.8$ at $10^{-11}\,{\rm m\,s^{-2}}$. The linear lift $(1-\epsilon) B_{\rm bath} f_{\rm cont}$ 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 $B_{\rm bath} \leq 1$ and $f_{\rm cont} \leq 1$, the linear candidate bounds the relaxed residual by $\mathcal{R}_{\rm rel} < 1 + (1-\epsilon) \simeq 1.81$. 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 $\mathcal{R}_{\text{sat}} = f_{\text{dec}} + f_{\text{cont}}/\epsilon \simeq 4.7$–$4.9$, 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 [57] into source-weight terms gives twenty-four direct tests of the bound: for each of the twelve clusters, at R_{500} and at R_{200}, the weight S solving $\nu(S\, g_{\text{bar}}/a_0)\, S\, g_{\text{bar}} = g_{\text{obs}}$ with non-thermal-corrected masses, where $\nu(x) = [1 - e^{-\sqrt{x}}]^{-1}$ is the response factor of the radial-acceleration law (Section 14). Every point respects $S \leq 1/\epsilon$; the sample maximum is $S = 3.96$, at R_{500} of the merging cluster A2319, the most model-dependent point, and the relaxed systems span $S \simeq 1.8$–2.7 at R_{500} and 1.4–2.2 at R_{200}.

The same inversion adjudicates the deep end: the power-law continuation requires $S \simeq 6$–8 at $g_{\text{bar}} \simeq 1$–$2 \times 10^{-11}\,\text{m}\,\text{s}^{-2}$, precisely the accelerations of the R_{500}–R_{200} points, which sit at $S \simeq 1.5$–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—$S \simeq 3.7$–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 [58, 59], 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 $W_{\rm bath} \lesssim 1.4$.

The decoupled-fraction form is excluded by the mass trend. A second candidate class lets the residual ride on the decoupled stellar fraction, $\mathcal{R} = 1 + 4.32\, B_{\rm bath} f_{\rm dec}$ with $f_{\rm dec} = f_\star/(f_\star + f_{\rm gas})$. This class is excluded by the mass trend rather than the amplitude. Because $f_{\rm dec}$ *falls* with mass while $B_{\rm bath}$ 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 above and the Bullet morphology are governed by the same fixed coefficient ϵ but by *different source expressions*, because the gas is in a different physical state. In a relaxed cluster the gas is a virialized continuum and drives the residual through the bath-lift term $(W_{\rm bath} - 1) f_{\rm cont}$. In a Bullet-type merger the central gas is shocked, displaced, and out of equilibrium—not a relaxed bath—so it has small effective $B_{\rm bath}$ and couples to the long-range channel at the suppressed weight ϵ, while the outgoing collisionless cores remain decoupled. The two regimes share one coefficient; they are not one scalar law, and we do not present them as such.

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

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

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 $\chi_{\rm ent}$; 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 $\chi_{\rm ent}(x,t) = \rho_{\rm dec}(x,t) + [\cdots]$ 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 $\mathcal{R}_{\rm rel} < 1.81$ 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: $\mathcal{R}_{\rm rel} \le f_{\rm dec} + f_{\rm cont}/\epsilon \simeq 5$ 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 $B_{\rm bath}$, larger $f_{\rm cont}$) 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,

$$\boxed{\kappa_{\rm obs}(x,y) = A\Big[W_{\rm bath}\,\Sigma_{\rm bath}(x,y) + \epsilon\,\Sigma_{\rm shock}(x,y) + \Sigma_{\rm dec}(x,y)\Big] + b}$$

with ϵ fixed at $g_{\rm share,eff}/4\pi^2 \simeq 0.188$ or fit freely, and $W_{\rm bath}$ fit per system within $[1, 1/\epsilon]$, the saturated value $1/\epsilon \simeq 5.3$ predicted for developed baths. Recovering both $\epsilon \simeq 0.19$ on the shocked channel and saturation on the bath channel would be a two-coefficient parameter-free success. In the relaxed limit this reduces to the residual law $\mathcal{R}_{\rm rel} = 1 + (W_{\rm bath} - 1)f_{\rm cont}$; in the Bullet limit the central gas is shocked ($\Sigma_{\rm shock}$ at weight ϵ) and the peaks fall on $\Sigma_{\rm dec}$. A best-fit ϵ near 0.19 across Bullet, MACS J0025, the Musket Ball, a relaxed CLASH sample [4], and X-ray-bright versus X-ray-faint groups would be a parameter-free success that neither ΛCDM nor total-baryon modified gravity predicts, while a best fit near 1 or near 0 would falsify the sector. This map fit is the principal outstanding empirical task and has not yet been performed. It is decisive only if it removes the escape hatch that makes existing reconstructions non-discriminating: ϵ must be floated against the channel-weighted *baryonic* maps ($\Sigma_{\rm bath}$, $\Sigma_{\rm shock}$, $\Sigma_{\rm dec}$) with free dark haloes disallowed, since free haloes absorb the gas weight and leave ϵ unconstrained—which is why current Bullet reconstructions neither recover nor exclude $\epsilon \simeq 0.19$. The three-component decomposition makes the fit more demanding than a two-component one, but also more distinctive.

Three open items remain at the theory level. First, the relaxed residual and the Bullet morphology are governed by one fixed coefficient ϵ but by two regime-specific source expressions (virialized bath versus shocked non-equilibrium gas). This is justified by the differing physical state of the gas, but it is a weaker unification than a single scalar law, and the boundary between the regimes—how $B_{\rm bath}$ interpolates as a shocked atmosphere re-virializes—is not yet derived.

Second, the channel-selection rule rests on two distinct physical premises, not one. The first is a *suppression* premise: participation in a phase-averaged diffuse continuum restricts a source to the transverse projection (weight ϵ), so that a collisionless overdensity decoupled from that continuum recovers the full horizon projection (weight 1). The second is a *collective-lift* premise: once a diffuse continuum virializes into a coherent extended phase it opens a collective long-range response that raises it *above* the decoupled baseline, to a weight $W_{\mathrm{bath}} > 1$. These two are logically independent—escaping the suppression returns a source to weight 1 but does not by itself carry it past unity—and it is the collective-lift premise, not the suppression premise, that carries the relaxed-cluster residual. The coefficient ϵ is derived; both premises are plausible within the ontology but neither is yet derived from the microscopic source map, and supplying those derivations is the channel-selection theorem. The collective-lift premise is the weaker of the two: the coherence/phase-averaging argument that motivates the suppression does not on its own generate super-baseline coupling, and the linear form of the lift is quantitatively excluded by the measured peak residual (above). The channel-selection theorem must therefore deliver the coherence-growth profile between the fixed endpoints, reproducing the measured radial decline of the source weight under uniformly low turbulence, constrained at the group end to $W_{\mathrm{bath}} \lesssim 1.4$ and saturating at $1/\epsilon$ in the hook-peak window.

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 therefore presented as a structured, falsifiable proposal, not as closed physics: the projection coefficient ϵ is derived; the bath gate is a measured input; the trend direction (residual riding on the continuum, rising from groups to clusters) and the hook morphology of the residual profile are supported by current data; the linear lift candidate is excluded on amplitude; the saturation amplitude is fixed by the capacity bound at the reciprocal of the derived projection coefficient, and the saturated residual lands inside the measured peak band; the coherence-growth profile, the resolved-map test, and the channel-selection theorem are the named open work.

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 $\overline{S}(t)$ 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,

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

where $\overline{S}(t)$ 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 becomes active near matter–radiation equality. This gives a transient early-energy component of the same general type used in early-dark-energy resolutions of the Hubble tension. In the closed cosmological branch treated here, the effect reduces the sound horizon and shifts the CMB-inferred Hubble constant upward from the high-67 range toward the high-68 to low-69 range.

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 $\overline{S}(t)$ 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

In galaxies the substrate has slack: a source's demand for capacity is met by a developing local response, and the radial-acceleration law of Section 14 follows. The early universe has no slack. The capacity bath is saturated across the linear perturbation regime: wherever a perturbation-scale source demands a response, the available transfer channel has already reached its ceiling, so there is no room for a developing response field. What gravitates is then not a field relaxing toward a source but a conserved count of *committed* capacity channels: cells locked into transferring one bit of capacity per tick on behalf of a defect. That committed count is carried along with the expansion and dilutes as the volume grows, so it redshifts as a^{-3} and gravitates as pressureless dust. This saturated phase plays, in the microwave-background sector, the role conventionally assigned to cold dark matter — not a new particle species but a phase of the same medium.

Saturation is forced at recombination. At $z \simeq 1100$ the accelerations of linear density perturbations satisfy $x = g_{\text{pert}}/a_0(z) \sim 10^{-3}$, where the response demand $\nu(x) \sim 30$ far exceeds the capacity ceiling $1/\epsilon = 5.32$. The bound derived for the cluster bath (Section 17.5) is therefore mandatory in the early universe: the response channel for the cosmological perturbation bath is saturated. The background itself sits permanently at the ceiling acceleration, since $a_0(z) = \epsilon\, cH(z)$ implies the invariant $cH(z)/a_0(z) = 1/\epsilon$ 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

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

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

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

with $u_\mu = \partial_\mu\phi$ and n the conserved spatial density of committed cells. This is the constrained-scalar (mimetic) class [60]: conditional on the pinned reading, the committed phase is pressureless dust, and the perturbation carrier is the conserved committed-capacity density itself. The gravitating component is therefore a phase of the medium, not an added species; throughout, "dark matter" names the conventional attribution this phase replaces. The known caveats of the class — caustic formation and the need for a small-scale completion — are inherited and recorded; they parallel the standard cold-dark-matter small-scale account. The continuous dynamical passage itself — how the generic unconstrained kinetic term of the weak-field branch is driven onto the constraint surface $X = \frac{1}{2}$ as the channel saturates, rather than the constrained action being posited at the ceiling — is asserted here rather than derived; supplying it, together with the commit/decommit energy accounting recorded below, is an open task within this conditional sector. The constrained route also parallels the mechanism by which the one relativistic completion of Milgromian dynamics known to meet this test does so [61].

Uniqueness of the carrier. The constraint coupling is the unique survivor of the dynamics classes examined (Appendix L). 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 L; 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

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

a 0.7σ agreement with the abundance counted rather than fitted [39]; equivalently $\Omega_m/\Omega_b = 1 + 1/\epsilon = 6.32$. A standard Einstein–Boltzmann computation [62] with the cold component tied to this value and the acoustic scale θ_* held fixed reproduces the 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 ΛCDM and sits at $\Delta\chi^2 = +2.15$ against the six-parameter best fit on the compressed Planck TT,TE,EE likelihood [30] — within noise of the full fit, and an upper bound on the constrained optimum since the remaining parameters were not re-optimized. The comparison must hold θ_* fixed: tying the densities at fixed H_0 instead breaks the acoustic scale and misreports the residual.

That the committed density attains the ceiling exactly, rather than a development-dependent fraction of it, follows from the recruitment structure of commitment. Commitment is source-driven: relaxation serves the demand of specific defects, and cells receiving no demand commit nothing. Each source's allocation is capped at $1/\epsilon$ per unit source mass, the same per-source bound that governs the cluster bath (Section 17.5). At recombination the demand side is fixed by the arithmetic above: perturbation-scale sources demand $\nu \sim 30$, far above the cap, while the homogeneous background sits at the invariant acceleration $cH/a_0 = 1/\epsilon$, where the demand $\nu(1/\epsilon) = 1.11$ lies below it. With supply abundant and commitment irreversible in the pinned regime, every source therefore recruits to its cap and the background recruits nothing. Allocations add, because committed units are per-defect: a unit serving two defects would make their joint entanglement budget, and hence their joint mass, sub-additive, contradicting the additivity of separated masses under Postulate II (Section 3.2). The committed density is then

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

pointwise at the commitment epoch, which delivers the abundance above and hands the perturbations adiabatic, baryon-tracking initial conditions, $\delta_c = \delta_b$ — the conditions the Einstein–Boltzmann computation assumed. The derivation is conditional on the same pinned reading as the dust theorem and on the per-defect bookkeeping just invoked; the energy accounting of the committed budget remains the sector's open completion, and conservation of the perturbations is taken up directly below.

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: $\nabla_\mu J^\mu_{\text{commit}} = 0$ 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 resting on the same history weighting that fixes the arrow of time (Section 21). 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 [63], 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 $f_{\rm rel}$ has already shell-crossed, the surviving committed density is $(1 - f_{\rm rel})\, \rho_b/\epsilon$, so the prediction specifies location, edge, and magnitude together.

Domain structure. The release threshold $\nu(x_c) = 1/\epsilon$ gives $x_c = 0.0433$, an acceleration $g_c = x_c\, a_0 \simeq 5.2 \times 10^{-12}\,\mathrm{m\,s^{-2}}$ today. The linear cosmological field remains below threshold at all epochs ($x \lesssim 2 \times 10^{-3}$ today, smaller as $\sqrt{a}$ earlier), so linear-regime observables — baryon acoustic oscillations, the linear power spectrum, microwave-background lensing — coincide with the ΛCDM form at the abundance above; in particular the framework inherits, and does not relieve, the S_8 tension. Departures are confined to released, collapsed structures, which is the galactic and cluster phenomenology of Sections 14 and 17.5.

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 and settles in time once a source moves (Section 17). Cosmology asks how the homogeneous capacity mode behaves on horizon scales (Section 18). The saturated phase asks what happens when the bath has no slack left and the carrier becomes a conserved committed density (Section 18.5). The strong-field branch asks the opposite-extreme question, what happens where local capacity falls to zero (Section 20). Many-Pasts asks what history and probability structure the substrate carries, and supplies the memoryless dressing the scale-setting already used (Section 21). The microstructure Hamiltonian asks what dynamics underlies the whole construction (Section 22).

These sectors share the static weak-field ontology but not its evidential status. The static chain is the closed result; the others are conditional extensions, frontier completions, or structural realizations, marked where each appears.

20. Strong-Field Branch and Bounded Occupancy

The weak-field theory used small capacity deficits. Black holes are the opposite regime: they ask what happens when the deficit is no longer small, and ultimately where the medium runs out of available capacity altogether. The branch therefore begins with the same variable as the weak-field bridge, written in the form that makes boundedness explicit:

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

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

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

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

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

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

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

$$N^2 = q$$

is the unique continuous multiplicative completion compatible with the weak-field branch. Once the capacity variable is adopted, the static lapse rule is fixed.

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

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

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

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

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

The continuum branch terminates at $q = 0$. A real continuation of this static branch to $r < r_h$ would require $q < 0$, which is outside the state space of the capacity variable. The classical Schwarzschild interior is therefore not interpreted as a physical continuation of the same entanglement-capacity EFT. It is a formal continuation of the GR manifold beyond the point where the substrate has no remaining local continuum channels.

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

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

21. Many-Pasts: The History-Space Ontology

Many-Pasts is the most conservative part of the theory to test and the most radical to adopt. In the selected operational branch, it changes nothing a laboratory measures: it reproduces Born-rule statistics and no-signaling exactly. Ontologically it replaces the entire interpretive layer of quantum mechanics — branching worlds, collapse events, or hidden variables — with a single object, a weighting over the admissible past histories of the entanglement substrate (Section 3.3).

The operational weight. The history weight $P(H|P) \propto e^{-D(H,P)}$ returns ordinary quantum statistics in the projective laboratory limit through the identity

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

the history-space distance reducing to the overlap of the present projector with the state evolved from H. Born recovery is a consistency condition on the weighting family, not a fresh derivation of the rule: requiring exact Born statistics selects $\alpha = 1$, and forbidding any signaling-sensitive operational bias selects $\beta = 0$. In that branch the rule and no-signaling both hold exactly, and nothing a laboratory measures distinguishes the framework from standard quantum mechanics. The conservatism is deliberate: the ontological content carries no operational cost.

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 the standard puzzles the weight changes the ontology, not the calculation. In the double-slit experiment the particle is not taken to have secretly travelled one slit while the other path was unreal: before any which-path record exists, the detection event is supported by the admissible histories through the apparatus, including the alternatives whose amplitudes interfere in ordinary quantum mechanics, and making a which-path record conditions that ensemble, removes the cross terms by decoherence, and erases the fringes. In an EPR or Bell experiment the present is a joint record of the pair and the detectors, and the histories are weighted as histories of the whole entangled system, which reproduces the nonclassical correlations while leaving the local Born marginals — and so the impossibility of signaling — untouched; the "spooky" element is the global conditioning of the history ensemble on a shared record, not a faster-than-light influence. Measurement is read the same way: it adds no collapse law but creates a durable record, and the operative ensemble is then the histories compatible with that record. The probabilities are the usual Born probabilities throughout; Many-Pasts supplies only the history-space reading of why those are the numbers being counted. The worked versions are in Appendix G.10.

The arrow of time. The same weight orients time through conditional typicality. Among the histories consistent with the present macroscopic records, overwhelmingly many show entropy increasing toward the future direction those records define. This is not an added law of laboratory probability; it is a statement about which histories dominate the weight, and it places the substrate's account of temporal asymmetry on the same footing as its account of quantum statistics.

Where the weight is used. Many-Pasts enters the gravitational chain through the electron dressing. Its memorylessness in substrate time is one of the two conditions that fix the absolute substrate length: each dressing pass of the electron is an independent draw from the admissibility ensemble, so the seven channels carry the full sharing entropy and add (Sections 13, 22; Appendix H). The same conditional-typicality weight that orients time also supplies the orientation of decommitment in the saturated phase, forbidding the spontaneous recoherence that re-entry into the committed phase would require (Section 18.5). The same weight therefore appears in three places: the scale-setting branch, the cosmological release law, and the quantum-foundational layer.

Status. Operational closure is exact: the laboratory sector is standard quantum mechanics, and the technical reduction is collected in Appendix G. The ontology is central rather than secondary, but its distinctive claims — branch realization and the arrow of time — are structural and cosmological, not new laboratory predictions, and they should be read at that grade.

22. Microstructure Hamiltonian and Underlying Dynamics

The UV closure chain has a microscopic dynamical realization with two sides: the emergence of the continuum geometry from the substrate, and the defect dynamics that fixes the scale-setting principle of the weak-field chain. The first is the condensate-side compatibility check; the second closes the one microscopic principle the static branch was conditional on. Section 13 stated the electron anchor in physical terms; this section gives the microscopic dressing dynamics behind it, and Appendix H carries the technical derivation.

On the geometry side, the realization is a GFT/condensate picture [48, 49] in which spacetime emerges from a condensate of discrete tetrahedral building blocks, while what is macroscopically read as matter appears as fermionic defects of that same substrate. In Madelung form,

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

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

On the defect side, the same substrate supplies the dynamics behind the faithful sector-resolution principle that sets the absolute cell length L_*. The elementary fermionic defect is governed by a dressing Hamiltonian whose single-channel terms make the lightest charged defect bind all seven face sectors once — the one-pass ground state that supplies the exponent

7 — and whose inter-channel term controls whether the seven-channel dressing cloud factorizes. Factorization turns the seven equal sector contributions into the additive $7g_{\text{share,eff}}$, and needs two things: each channel must carry the full admissibility entropy, and the seven must be mutually independent. The first is memorylessness in substrate time — the Many-Pasts content of Postulate III, with the substrate relaxing to a fresh draw long before the electron is read out. The second follows from the electron's own lightness: inter-channel correlation would contract the defect and raise its mass, so the lightest charged defect carries none. Each channel at full entropy and the channels independent give $\dim_{\text{eff}} = e^{7g_{\text{share,eff}}}$ and $\lambda_e/L_* = \frac{2}{3}e^{7g_{\text{share,eff}}}$, worked out in full in Appendix H.

Reading the support ratio λ_e/L_* as a physical Compton length used one further identification: that the defect advances one phase cycle per dressing recurrence. Appendix H.6 gives that recurrence an operator form and removes the identification as a free assumption. The mass is the exported gap of the memoryless refresh super-operator, $m_e c^2 = \frac{3}{2}(\hbar/\tau_*)\, e^{-7g_{\text{share,eff}}}$ (the $\frac{3}{2}$ the transverse export of the length relation), so it is a rate and not a binding energy, and the phase-cycle claim decomposes into the recurrence rate, a unit committed charge per pass from fermionic one-pass occupation, and the standard condensate phase relation, none with a free constant. What is not forced is the identification of the winding phase with the condensate U(1) (Fork A); Appendix H.7 derives the three dressing functionals from the mean-field condensate and shows that Fork A and the rank-one inter-channel coupling reduce to a single premise — a single scalar field on that condensate — which is the founding premise of the framework and is tested in the charged-lepton ladder (Appendix I.1). The hierarchy then rests on one asymmetry. Binding across the seven sectors is additive, but their coherent recurrence is multiplicative, so the continuous maintenance of the defect is large and linear in the commitment while the rest energy, fixed by the rare joint return, is exponentially small. The mass scale is the rare-return gap of the refresh generator, not the bare binding cost: the electron is light not because it is weakly bound but because its completed seven-sector dressing recurs rarely.

Together these two sides do more than check compatibility. The defect dynamics derives the faithful sector-resolution principle from the mass–entropy ontology and the electron anchor, so the substrate length and the induced gravitational scale follow from commitments the theory already makes. What remains microscopically open is twofold: the first-principles derivation of every inhomogeneous continuum coefficient from the full kernel, and — conditional on Fork A — the GFT substrate actually producing the single-scalar mean-field condensate, with its emergent four-geometry, on which the defect spectrum and the geometry side both depend (Appendix H.7).

Part VI. Closure Status, Falsifiability, and Research Program

23. Closure-Status Table

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

The status column uses a fixed vocabulary. *Closed* means derived within the stated postulates and ensemble. *Fixed* means no phenomenological freedom remains once the branch is adopted. *Conditional* means the result follows if a named reading or completion holds. *Frontier* or *open* means the piece is structured but not yet complete. *Empirical support* means a comparison with data rather than a derivation, and *audit task* means an independent derivation or check is still required. Rows follow the order of the text, grouped by sector from the UV counting through the foundational, weak-field, galactic, transport, cosmological, strong-field, and quantum-foundational sectors.

Quantity / Claim	**Sector**	**Status**	**Type of Support**	**Where Established**
Ω_{tet}, $g_{\text{share,max}}$	UV counting	Closed	exact combinatorics	Part II, App. B
η_*	admissibility closure	Closed	stationary normalized closure evidence $\ln Z + \frac{3}{2}\ln\eta$ maximized on the exact K^2 spectrum; the 3/2 is the determinant weight of the three closure-defect components	Part II, App. B
$g_{\text{share,eff}}$	UV entropy	Closed	exact weighted evaluation	Part II, App. B
Substrate length L_* (faithful sector resolution / support-to-length)	scale setting	Derived from dressing recurrence; identification discharged under Fork A	$\lambda_e/L_* = \frac{2}{3}e^{7g_{\text{share,eff}}}$ from the recurrence time of the memoryless seven-channel dressing (waiting time for joint typicality); support exponent bracketed at unity within 10^{-3} by G_* and the lepton ladder; phase-cycle identification discharged in App. H ($\Gamma \oplus \Delta Q{=}1 \oplus$ Josephson), residual is Fork A (winding phase = condensate U(1)) and the unit-winding integer, both tested by the lepton ladder	Part I, App. D, H
J_{bare}, $J_{\text{eff}}^{\text{tree}}$	UV edge kernel	Closed	tetrahedral isotropy identity	Part II, App. C
$\Sigma_{\text{ret}} = 65/9$	finite-loop UV	Fixed in the minimal UV return sector	explicit seven-channel plus singlet return count	Part II, App. C
$J_{\text{eff}}^{(\text{ren})}$	finite-loop UV	Fixed by the UV return resummation	derived from Σ_{ret}	Part II, App. C
γ	continuum stiffness	Closed in canonical EFT convention	Euclidean-action normalization	Part II, App. C
Green-matched source projection	UV source map	Closed in the canonical weak-field branch	exact defect counting + tetrahedral on-site Green function	App. C
κ/γ	source-to-stiffness ratio	Closed in the canonical weak-field branch	$\sigma_{\text{def}} = \rho/\kappa_m(L_*)$ plus $G_{\text{tet}}(0)$ and tetrahedral 4/3 projection	Part III, App. C–D
Weak-field bridge law	EFT / gravity	Closed in canonical weak-field branch	uniqueness from multiplicative composition	Part III, App. D
G_*	gravitational scale	Conditional derivation at percent level; calibration, with provenance recorded	$G_* = (9/4)(\hbar c/m_e^2)e^{-14g_{\text{share,eff}}}$	Part I, App. D, K
Matched G	weak-field gravity	Closed in the matched weak-field branch	bridge law + Green-matched κ/γ + fixed-epoch normalization of $\kappa/(\gamma S_\infty)$; comparison with G_* tests coherent propagation of the same L_*, not a disjoint-input second derivation of G	Part III, App. C–D
Electron anchor	mass / length sector	Fixed empirical elementary anchor	one-bit fermionic defect supplies λ_e for length setting and $m_e/\ln 2$ for mass–entropy map	Part III, App. D
Seven-sector additivity	UV scale theorem	Closed conditional on the support-to-length identification	$H_{\text{seven}} = 7g_{\text{share,eff}}$ from subadditivity plus mass minimization: full per-channel entropy from the memoryless dressing, channel independence forced because correlation would raise the defect mass ($\Delta = 0$ for the lightest one-bit defect)	App. D, H
a_0	galactic EFT	Fixed in the closed weak-field realization	UV entropy × horizon scale cH_0, with $4\pi^2 = (2\pi)^2$ the canonical two-transverse-mode phase-space cell of the $1+2$ split	Part III, App. C

Quantity / Claim	**Sector**	**Status**	**Type of Support**	**Where Established**
RAR law	galactic EFT	Closed at the rate-bookkeeping level	a_0 fixed by the coefficient chain; $x = \sqrt{g_{\rm bar}/a_0}$ forced by the spectral invariant of the bilinear form; $1 + n_B(x)$ from exact bosonic emission statistics, with absorption closed by the quantized one-bit deficit; flux identification derived from the static transport limit (conserved bit current, Green-matched normalization), with inter-channel delivery fixing the locality of x; kernel neutrality $F \equiv 1$ from per-cell channels and bit quantization; graph-level substrate computation of the inter-channel vertex is the remaining gold-standard check	Part III, App. C
No slip / lensing consistency	weak-field metric	Closed at leading weak-field order	scalar-stress structure	Part III, App. D
PPN leading values	weak-field metric	Structurally supported	weak-field expansion	Part III, App. F
Telegrapher relation $D/\tau_0 = c^2$	transport	Closed in canonical transport branch	causal closure	Part IV, App. E
Canonical $\tau_0^{-1} = H_0$ branch	transport	Fixed in the minimal transport closure	no-new-IR-scale choice	Part IV, App. E
Hubble-tension mechanism	cosmology	Structurally supported extension	homogeneous trace-coupled mode	Part IV, App. E
Saturated-phase committed dust	cosmology	Conditional on the pinned transfer-law reading	saturation forced at recombination; constrained-scalar dust theorem ($p = 0$, $c_s^2 = 0$, $\rho \propto a^{-3}$); relaxational carriers excluded	Part IV, §18.5, App. L
Committed-capacity abundance $\Omega_c/\Omega_b = 1/\epsilon$	cosmology	Derived from forced recruitment, conditional on the pinned reading and per-defect bookkeeping	over-demanded sources recruit to the per-source cap, the background sits below it, and per-defect allocations add (Postulate II corollary), giving $\rho_c = \rho_b/\epsilon$ with adiabatic seeding $\delta_c = \delta_b$; temperature-spectrum check at fixed abundance; counted, not fitted	Part IV, §18.5
Caustic release / conditional conservation	cosmology	Conditional — forced under the synchrony reading	committed dust decommits at first shell-crossing of the irrotational clock flow; splashback-rim falsifier; conversion bookkeeping open	Part IV, §18.5
Diffuse source projection $\epsilon = g_{\rm share,eff}/4\pi^2$	cluster sector	Derived (no new coefficient)	same transverse reduction factor that fixes a_0/a_H (Section 14)	Part IV, §17.5
Bath gate $B_{\rm bath}$	cluster sector	Operationally closed; microscopic origin open	$M_{\rm hot,vir}/(f_{b,\cos} M_{500})$ from X-ray/SZ; no per-system fitted knob	Part IV, §17.5
Cluster residual $\mathcal{R}_{\rm rel} = 1 + (W_{\rm bath} - 1) f_{\rm cont}$	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 $\mathcal{R}_{\rm sat} \simeq 4.7$–$4.9$ 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

Quantity / Claim	**Sector**	**Status**	**Type of Support**	**Where Established**
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 $\kappa \propto (1+(1-\epsilon)B_{\rm bath})\Sigma_{\rm bath}+\epsilon\Sigma_{\rm shock}+\Sigma_{\rm dec}$; predicted best-fit $\epsilon \simeq 0.19$, floated on baryonic maps with free dark haloes disallowed; not yet performed	Part IV, §17.5
Bounded capacity q, $N^2 = q$	strong field	Closed as constrained-capacity rule	uniqueness from multiplicative composition and weak-field matching	Part V, App. F
Constrained action on $\mathcal{M}_q = \{q > 0\}$	strong field	Closed within the EFT	ADM action plus lapse-capacity constraint	Part V, App. F
Static spherical exterior	strong field	Closed within the constrained-capacity branch	bulk vacuum reduces to Einstein vacuum on $\mathcal{M}_q$	Part V, App. F
Capacity-exhaustion horizon $q = 0$	strong field	Closed as static-domain boundary	Schwarzschild exterior terminates where q leaves the physical state space	Part V, App. F
Hawking temperature and exterior ringdown	strong field	Closed at leading EFT order	unchanged Schwarzschild exterior; absorbing boundary at $q = 0$ fixed by horizon regularity; echoes only as UV / stretched-layer corrections	App. F
Bekenstein–Hawking area-law bridge	strong field	Closed within the constrained-capacity EFT under Postulate II and the transverse graph response	per-channel cut entropy $\ln 2$ from fermionic face exclusion + channel density from $G_\perp = (2/3)G_{\rm tet}(0)$; product $n_{\rm hor}S_\infty^{\rm cell} = 1/4$ as exact identity	App. C, F
Horizon formation and boundary microphysics	strong field	Formation closed geometrically; nonuniversal spectroscopy open	$q = \lvert\nabla R\rvert^2 = 1 - 2GM_{\rm MS}/(c^2R)$ in spherical symmetry, with $q = 0$ at the marginal-trapped-surface inherited from standard apparent-horizon formation; remaining work covers relaxation spectra, stretched-layer corrections, and transient response	Part V, App. F
Rotating / charged stationary exteriors	strong field	Closed within the constrained-capacity EFT	vacuum Einstein / Einstein–Maxwell reduction on $\mathcal{M}_q$ gives Kerr / Reissner–Nordström / Kerr–Newman exteriors; $q = 0$ identified with outer Killing horizon; inner Cauchy horizons are non-physical analytic continuations into $q < 0$	App. F
Many-Pasts Born compatibility	quantum foundations	Consistency condition	exact Born statistics selects $\alpha = 1$ within the weighting family; the content is the existence of a branch meeting Born and no-signaling simultaneously	Part V, App. G
No-signaling in operational branch	quantum foundations	Consistency condition	forbidding signaling-sensitive bias selects $\beta = 0$	Part V, App. G
Arrow-of-time account	quantum foundations	Coherent extension	conditional typicality / counting	Part V, App. G

Quantity / Claim	**Sector**	**Status**	**Type of Support**	**Where Established**
Microstructure Hamiltonian	UV realization	Defect side closes scale-setting; geometry side coherent	dressing Hamiltonian with one-pass ground state, memoryless dressing, and minimal-mass channel independence (scale-setting); mass as the refresh-operator gap with the phase-cycle identification discharged into $\Gamma \oplus \Delta Q{=}1 \oplus$ Josephson (App. H.6) and ε_0, F_m, V from the mean-field condensate with rank-one V as one scalar (App. H.7); GFT/condensate origin of the kinetic term (geometry)	Part V, App. H
Charged-lepton spectrum	particle-sector extension	Derived; coefficients forced, one map-form input	three-generation termination from K^2 permutation-invariance at $N = 3$; $m_N/m_e = 720^N(2/7)^{N^2}$ with $720 = 6!$, $2/7$ the singlet projection, N^2 the ordered-pair count; matches PDG at 0.5–0.65% with no fitted parameters	App. I
Gauge-redundancy extension	gauge sector	Coherent extension	baseline-redundancy construction with Maxwell/Yang–Mills form	App. I
Numerical robustness checks	validation layer	Supportive audit layer	cross-sector consistency tests	App. J
EFT consistency checklist	field-theory audit	Supportive audit layer	no-ghost / no-tachyon / causal-propagation checklist with explicit vacuum dispersion stability	App. D

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

Cosmological row. The saturated phase enters the ledger as follows: the dust form of the committed phase is a theorem of the constrained-scalar class, conditional on the pinned reading of the transfer law; the abundance $1/\epsilon = 5.321$ stands against the measured 5.364 ± 0.065, derived from forced recruitment conditional on the same pinned reading and the per-defect bookkeeping of Postulate II; the energy accounting of the committed budget is open; the linear regime is closed by the domain structure.

24. Falsifiability and Observational Tests

24.1 Static weak-field falsifiers

The static weak-field sector stands or falls on a small number of concrete checks. The most direct are the shape and tightness of the galaxy RAR transition [1, 34], the baryonic Tully–Fisher scaling in systems where the EFT should apply, and the weak-field lensing sector, where stacked galaxy–galaxy lensing now probes the relation two decades below the rotation-curve regime [35]. A persistent need for gravitational slip where the scalar stress predicts none would be especially damaging, because it would break the same no-slip structure used to keep dynamics and lensing aligned. Solar-system bounds, especially Cassini-class PPN tests [3, 43], also require the leading no-slip branch to survive at high precision, and laboratory fifth-force searches [45] bound any residual static scalar channel in the same regime.

Wide binaries supply an independent solar-neighborhood discriminator. The occupancy argument of Section 14 is local, a feature its transport derivation fixes rather than assumes: the acceleration entering x is the total local baryonic field, so a binary embedded in the Galactic field $g_{\rm ext} \simeq 1.4$–$1.9 \times 10^{-10}\,\mathrm{m\,s^{-2}}$ has its internal dynamics amplified once the pair's own field falls below the Galactic term. The predicted internal-force boost rises from unity at small separations to $1 + n_B(\sqrt{g_{\rm ext}/a_0}) \simeq 1.4$–$1.5$ beyond a few thousand au. This lands close to the AQUAL external-field value through a different mechanism — locality of the occupancy argument in the total baryonic field, with no nonlinear field equation — so the wide-binary regime separates the framework from Newtonian gravity while leaving it degenerate with AQUAL-class theories. Current analyses divide: a low-acceleration boost of about 1.4 is reported in Gaia DR3 samples [27], and consistency with Newtonian gravity is reported in an independent analysis of the same mission's data [28]; the disagreement is methodological. The framework's commitment is fixed in advance: a settled Newtonian verdict falsifies the locality of the occupancy rule.

24.2 Dynamical falsifiers

The dynamical extension is more vulnerable, and its failure modes are correspondingly sharper. Time-dependent halo lag, cluster-scale acceleration relations, merger offsets, or relaxation signatures that cannot be reconciled with the telegrapher relation $D/\tau_0 = c^2$ would indicate that the causal completion has the wrong propagation structure even if the static branch survives. Cluster data accordingly split into two tests of distinct parts of the framework: *source-projection* tests, which ask whether 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.

24.3 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 [39], distance-ladder determinations ranging from $\simeq 70$ [41] to 73 [40], and a tension whose proposed resolutions are reviewed in Di Valentino et al. [42].

The epoch dependence $a_0(z) = cH(z)\, g_{\mathrm{share,eff}}/4\pi^2$ is the framework's most exposed kinematic prediction, and high-redshift rotation curves already bear on it. At $z \simeq 2.3$ the predicted scale is $3.5\, a_0$; a massive star-forming disk with $g_{\mathrm{bar}} \simeq 2 \times 10^{-10}\,\mathrm{m\,s^{-2}}$ at 8 kpc is then predicted to show $g_{\mathrm{obs}}/g_{\mathrm{bar}} \simeq 2.0$, against 1.36 for an epoch-independent scale. Observed outer rotation curves of such disks decline and indicate strong baryon dominance within the disk scale [29, 33], the direction opposite to the predicted enhancement. Pressure-support corrections at these redshifts are large and the stacked-profile normalization is disputed, so the confrontation is unsettled. The framework requires that massive $z \simeq 2$ disks, once those corrections are controlled, show stronger low-acceleration boosts than matched $z = 0$ systems; the opposite outcome falsifies the identification of a_0 with $cH(z)$ while leaving the static $z = 0$ sector intact.

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 [31], and a resolved low-redshift H I sample independently reports tentative evolution in the same direction [32]. 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 $a_0(z_{\mathrm{high}})/a_0(z_{\mathrm{low}})$ against $E(z_{\mathrm{high}})/E(z_{\mathrm{low}})$, 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 of the constrained committed fluid must reproduce the measured temperature and polarization spectra with the abundance fixed at $1/\epsilon$; the recruitment derivation must survive scrutiny of its two named conditions, the pinned reading and the per-defect bookkeeping; and departures from the ΛCDM form must be confined to accelerations above $g_c \simeq 5.2 \times 10^{-12}\,\mathrm{m\,s^{-2}}$, so canonical behavior failing below that boundary, or non-canonical behavior appearing in the linear regime, would falsify the domain assignment. The release law adds a fourth: a surviving committed component on cluster infall streams terminating at the splashback surface; its robust absence, or a halo-like committed component persisting inside collapsed systems, would falsify the caustic release mechanism.

24.4 Correlated-constant falsifiers

One of the more distinctive signatures of the framework is that the same microstructural chain feeds the substrate scale, the weak-field gravitational normalization, and the galactic acceleration scale. A precision program that could test the inferred L_* and G_* scale against matched G, a_0, and the RAR normalization would probe the theory more sharply than isolated single-observable fits, because it would confront the shared coefficient origin directly.

24.5 Many-Pasts status

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

25. 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 weak-field closure. A persistent gravitational slip where the scalar stress predicts none; an RAR transition shape that systematically departs from the derived bosonic occupancy law in systems well described by the static branch; a source projection contradicted by the data; or a weak-field UV coefficient chain that cannot be reconciled with an independently validated microscopic derivation — any of these breaks the central completed result while leaving the ontology in principle 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 the constrained-capacity branch cannot retain the tested Schwarzschild exterior, standard absorbing-boundary ringdown, or the required horizon thermodynamics once its boundary action is derived, the strong-field completion fails. None of these touches the weak-field core.

Requires modification, not death. Some results could be wrong without bringing down a sector: an adjustment to the finite-loop self-energy from a fuller graph-level computation, a revision of the strong-field boundary microphysics, or a failure of the charged-lepton ladder, which is a cross-check on the mass–entropy ontology rather than part of the closed gravitational chain.

One scale-setting commitment cuts across these tiers and deserves its own line. The percent-level G_* rests on the seven-fold additivity, which rests in turn on the electron being the lightest one-bit charged defect, since inter-channel correlation in its dressing would raise the defect mass. If that identification or the mass–entropy ontology failed, the scale-setting derivation of Appendix H would not go through.

26. Comparison with Other Approaches

Because the framework reattributes the dark sector (the galactic excess to medium response, the cosmological abundance to a saturated phase of the same medium) and partially reorganizes the usual dark-energy story, the sections below set out how its logic differs from nearby alternatives.

26.1 Relative to ΛCDM

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

26.2 Relative to MOND-like interpolation programs

MOND-like programs [17, 47, 19] usually begin from an acceleration law or interpolation function and ask how much galaxy phenomenology it can explain. The present logic runs the other way. The interpolation law is not taken as primary; it is downstream of the UV entropy, the $1 + 2$ channel geometry, and the bosonic occupancy branch. The galactic law is thus treated as an output of the same micro-to-IR closure chain rather than as the phenomenological starting point.

26.3 Relative to Verlinde-style emergent gravity

Verlinde-style emergent-gravity programs share the broad intuition that gravity may be entropic [20, 46, 21, 22], but they are usually formulated at the level of thermodynamic reasoning or horizon-inspired force laws. The present framework is narrower and more explicit: finite tetrahedral boundary counting, admissibility closure, edge coupling, finite renormalization, Euclidean-action normalization, and only then a continuum scalar EFT. Whether that chain is ultimately correct is an empirical matter, but it is a different kind of proposal from a purely macroscopic entropic argument.

26.4 Relative to TeVeS and other multi-field modified gravities

Multi-field relativistic MOND completions such as TeVeS [18] typically introduce additional vector or tensor sectors in order to repair lensing or cosmological problems. The present weak-field construction instead keeps a single scalar entanglement field within a low-energy Einstein continuum sector that is itself interpreted as emergent from the substrate, and relies on the no-slip structure $\Phi = \Psi$ at leading order to keep lensing and dynamics aligned. That economy is attractive if the branch survives confrontation with data, and immediately vulnerable if future observations demand persistent slip or extra weak-field structure.

26.5 Relative to AeST

The closest modern comparator is the AeST theory of Skordis and Złośnik [23], which reproduces MOND-scale galaxy phenomenology and fits the CMB power spectrum using a vector field and two scalar degrees of freedom. AeST establishes that galaxy-scale modified dynamics and a viable linear cosmology can coexist in a single relativistic theory, and any framework in this space is measured against it. The present construction differs in starting point and in risk profile: the acceleration scale, the interpolation argument, and the source projections descend from a finite UV counting problem with coefficients fixed in advance, so nothing can be adjusted where AeST retains functional freedom, and the cosmological sector is correspondingly less developed — the full Boltzmann treatment AeST has passed is named open work here (Section 24.3). The two programs fail differently: AeST through accumulating parameter pressure, the present framework through any single fixed coefficient meeting a contrary measurement.

26.6 Relative to scalar-tensor gravity

The action of Section 10 can be mistaken for a Brans–Dicke-type scalar-tensor theory [44], since it carries a metric, a scalar, and a coupling between them. The difference is ontological, and it removes the usual scalar-tensor freedom. In a scalar-tensor model the scalar is an independent field added to general relativity, and its kinetic and coupling functions are free to be chosen or fit to data. Here the metric and the scalar are two continuum expressions of one finite-capacity substrate: the scalar is not added to gravity but the same medium read in a different variable, and its stiffness, source coupling, and bridge normalization are fixed by the UV counting chain rather than left as functions. The resemblance is at the level of the written action; the content is that the coefficients are not adjustable.

26.7 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. Unlike many-worlds, it posits no forward branching into co-real macroscopic outcomes; the multiplicity is over admissible pasts of a single realized present, not over futures. Unlike collapse theories, it adds no physical collapse event and no modification of unitary evolution. Unlike hidden-variable accounts, its probabilities are not ignorance over pre-existing classical values. It is closest in spirit to consistent-histories and decoherence-based reasoning, but it differs in carrying an explicit weight $e^{-D(H,P)}$ over histories and in using that weight to do physical work outside the foundations — the memoryless dressing that sets the substrate length, and the orientation of cosmological decommitment. It does not claim to out-derive the Born rule; it recovers Born statistics and no-signaling as consistency conditions on the weight (Section 21).

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

The construction uses discrete tetrahedral boundary data, and the $j = 3/2$ tetrahedron coincides with the quantum tetrahedron of simplicial spin networks [50, 51, 52], so it sits near the causal-dynamical-triangulation, spin-foam, and group-field-theory programs. The way it uses those structures differs on three points. It does not sum over arbitrary triangulations or evaluate a spin-foam vertex amplitude as the primary weighting; it solves a finite local boundary-counting problem and weights configurations by admissibility closure, routing the result to a capacity entropy rather than to area and volume operators. Matter is not added on top of the geometry but is a fermionic capacity defect of the same substrate. And the history weighting is supplied by Many-Pasts rather than by a path integral over geometries. The GFT/condensate picture enters only as a compatibility check on the continuum kinetic term (Section 22), not as the engine of the derivation. The relation is therefore one of shared combinatorial vocabulary with a different selection principle and a different target.

27. 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 that this picture is not free to be adjusted. Once the finite-capacity substrate and the three postulates are granted, the static weak-field sector follows from a finite counting problem with nothing left to tune. The minimal tetrahedral ensemble fixes the sharing entropy; the electron anchor fixes the substrate length; the edge kernel, loop dressing, continuum matching, and Green-matched source map fix the coefficients of the scalar EFT; and the weak-field bridge turns the capacity deficit into the gravitational potential. The rest of the paper rests on that chain, the only part offered as closed.

Several quantities then fall out of the one chain without separate tuning. Newtonian gravity is its point-source limit. The galactic acceleration scale is the horizon thermal scale reduced by the sharing factor, and the radial-acceleration relation follows from it with no per-galaxy fitting. Lensing and dynamics share a single potential at leading order, so galactic support does not come at the cost of a lensing inconsistency. The same substrate length implies a value for Newton's constant within about one percent of the measured one; we read this as a check that a single length propagates correctly through the whole chain, not as a second independent measurement, and Appendix K records its provenance. The charged-lepton mass ratios come out to better than a percent from the same shell algebra, as a cross-check on the mass–entropy ontology rather than as part of the gravitational chain.

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 the 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, with an abundance that is counted rather than fitted and lands on the measured dark-to-baryonic ratio. The strong-field branch ends the continuum where capacity runs out, recovering the Schwarzschild and Kerr exteriors and the horizon thermodynamics. These are extensions and frontier completions, not closed results, and the closure table marks the distinction.

Many-Pasts carries more weight here than an interpretation of quantum mechanics usually does. It is the claim that the present state of the medium is supported by many admissible microscopic pasts, weighted by how well each leads to the present. In the operational branch used throughout, it changes nothing a laboratory measures: Born-rule statistics and no-signaling both hold, and it supplies the account of the arrow of time. But the same weighting also makes the electron's dressing memoryless, and that memorylessness is one of the two conditions that fix the substrate length. The postulate therefore does real work in the scale that sets gravity, not only in the foundations of quantum theory.

The remaining work is specific and limited. The finite-loop self-energy should be confirmed by an independent graph-level calculation. The cosmological sector should be carried through a full Boltzmann likelihood. The cluster lift should have its coherence-growth profile derived and its lensing-map test performed. The strong-field boundary should have its non-universal spectroscopy worked out. And the substrate length's recurrence derivation, which rested on identifying one phase cycle per dressing recurrence, is now discharged in Appendix H into a recurrence rate, a unit committed charge per pass, and a condensate phase relation; the residual that remains is the sharper Fork A identification — that the winding phase is the condensate's — together with the single integer fixing the electron as the unit-winding defect, both exposed to the charged-lepton ladder. The paper states these plainly, because its value, if it has one, lies in making them sharp enough to settle.

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, $\delta S = S_\infty - S_{\mathrm{ent}}$.
- *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 ground-state cloud an elementary defect builds by resolving the boundary sectors; its coherent support fixes the defect's length scale.
- *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 the present is supported by a probability-weighted ensemble of admissible microscopic pasts, with weight $e^{-D(H,P)}$.

A.1 Units, signature, and entropy normalization

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

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

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

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

A.2 Core scalar variables

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

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

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

$$\begin{aligned}
\gamma &: \text{ entanglement-field stiffness,} && (12)\\
\kappa &: \text{ continuum defect–entropy coupling,} && (13)\\
\kappa_m(\ell) &: \text{ mass-per-entropy map at scale } \ell, && (14)\\
L_* &= \frac{3}{2}\lambda_e e^{-7g_{\text{share,eff}}}, && (15)\\
G_* &= \frac{c^3 L_*^2}{\hbar} = \frac{9}{4}\frac{\hbar c}{m_e^2}e^{-14 g_{\text{share,eff}}}, && (16)\\
G_{\text{tet}}(0) &: \text{ tetrahedral on-site Green constant,} && (17)\\
g_{\text{share,max}} &= \ln(1680), && (18)\\
g_{\text{share,eff}} &: \text{ admissibility-weighted sharing entropy,} && (19)\\
J_{\text{bare}},\ J_{\text{eff}}^{\text{tree}},\ J_{\text{eff}}^{(\text{ren})} &: \text{ UV edge couplings,} && (20)\\
a_0 &= \frac{cH_0 g_{\text{share,eff}}}{4\pi^2}. && (21)
\end{aligned}$$

The canonical weak-field bridge and Newton closure are

$$\frac{\Phi}{c^2} = -\frac{\delta S}{2S_\infty}, \qquad \frac{\kappa}{\gamma} = \frac{3L_*}{4G_{\text{tet}}(0)\kappa_m(L_*)},$$
$$G = \frac{c^2\kappa}{8\pi\gamma S_\infty}.$$

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

A.4 Notation map

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

Appendix 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

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

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

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

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 $(4\text{-face}, 7\text{-state})$ 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

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

The admissibility family is

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

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

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

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 $\mathcal{F}''(\eta) = \mathrm{Var}_\eta(K^2) - 3/(2\eta^2)$, and on the exact spectrum $\mathcal{F}'$ has a single interior zero,

$$\eta_* = 0.0298668443935,$$

at which $\mathrm{Var}_{\eta_*}(K^2) = 15.69$ 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,

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

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

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

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

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

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

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 already supports the local evaluation, the cavity benchmark, and the later shell propagation. The local information observable

$$\sigma_{\mathrm{ind}}^{(r)} = \frac{H(X \mid Y_r)}{H(X)}$$

has the principal pre-nonlocal benchmarks

$$\sigma_{\mathrm{ind}}^{\mathrm{toy}} = 0.44997, \tag{22}$$

$$\sigma_{\mathrm{ind}}^{\mathrm{loc}} = 0.44708, \tag{23}$$

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

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

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

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

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

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

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

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

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

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

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

C.3 Finite-loop self-energy closure

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

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

The structural decomposition is

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

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

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

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

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

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

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

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

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

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

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

and

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

This reproduces the shell-target crossing near $J_{\text{bare,cross}} \sim 0.019$ 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 the sum runs over all sites a and all four outgoing nearest-neighbor directions $\hat{n}_i$ (each nearest-neighbor edge counted twice). The microscopic four-cell is

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

The same tetrahedral identity then yields

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

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.

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

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

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]

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

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, $i = 1, \ldots, 4$, 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

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

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

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

so

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

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

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

so that

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

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

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

Equating this with the weak-field source equation

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

closes the canonical source-to-stiffness ratio:

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

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 $s_{\rm def}/J_{\rm eff}^{(\rm ren)} = \ln(7/6)/(\pi G_{\rm tet}(0))$ 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 defect shifts $\delta Q_{\rm def} = \ln(7/6)/\pi$. Converting to the entropy field through $S = \pi Q_{\rm occ}$ multiplies by π, so the S-field shift is $\delta S_{\rm def} = \pi\,\delta Q_{\rm def} = \ln(7/6)$ and the explicit π does not survive into the source theorem. *(ii) $\ln(7/6)$ is carried by the entropy density, not the coefficient.* A defect contributes entropy $\ln(7/6)$ at the isotropic benchmark, so a defect number density sources an entropy density proportional to it; writing that source as $\sigma_{\rm def} = \rho/\kappa_m(L_*)$ folds the benchmark and the mass–entropy unit into the single density $\sigma_{\rm def}$. The geometric coefficient $3L_*/(4G_{\rm tet}(0))$ that multiplies it is fixed by the tetrahedral 4/3 projection and the cell length alone and contains no $\ln(7/6)$. *(iii) $J_{\rm eff}^{(\rm ren)}$ cancels in the ratio.* The source amplitude $s_{\rm def}$ 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 $3L_*/(4G_{\rm tet}(0)\kappa_m(L_*))$. The renormalized loop coupling therefore drops out of the observable normalization.

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

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

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

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

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,

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

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:

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

The same chain feeds

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

and faithful sector resolution gives the substrate-induced scale

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

and the Green-matched source theorem fixes

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

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

The coefficients this appendix supplies to the main weak-field chain are J_{bare}, $J_{\mathrm{eff}}^{\mathrm{tree}}$, Σ_{ret}, $J_{\mathrm{eff}}^{(\mathrm{ren})}$, 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 uniqueness

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

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

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

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

and therefore

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

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

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

D.2 Point source, Newton limit, and lensing

In the renormalized static branch,

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

For a point source M,

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

Because the leading entanglement stress carries no anisotropic stress,

$$\Phi = \Psi$$

at the order treated. The effective-halo rewrite is

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

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

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

D.3 Electron anchor and composite matter

The canonical fermionic entropy increment is

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

The UV mass normalization is

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

and the running law in the closed branch is

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

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

$$m_{\mathrm{hadron}} = \kappa_m(\ell_H) S^{\mathrm{dressed}}_{\mathrm{ent},H},$$

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, and this subsection makes the link explicit. The task is to turn a dimensionless count — the effective number of dressing configurations the electron's seven channels resolve — into a physical length ratio λ_e/L_*, and hence into the induced gravitational scale G_*. The argument has three parts: a recurrence derivation that fixes the linear, first-power map from the seven-channel sharing entropy to the length; a demonstration that the allowed local single-label dynamics cannot reach the full ensemble, so the memoryless refresh of Postulate III is required; and an explicit transfer-operator realization that makes the count concrete. One identification sits beneath the map — one phase cycle per dressing recurrence — which Appendix H.6 gives an operator form and discharges into a recurrence rate, a unit charge per pass, and a condensate phase relation, conditional on the Fork A identification settled in H.7; the linear power itself is then cross-checked three independent ways.

The recurrence derivation of the dictionary. The support-to-length identification follows from the recurrence structure of the dressing. By Postulate III the dressing is memoryless in substrate time: each tick, each of the seven channels draws independently from the admissibility ensemble. By the asymptotic-equipartition property — used here only in its elementary form, that a distribution with entropy g has an effective typical support of size e^g, so a specified typical configuration appears with probability e^{-g} per independent draw — a single channel realizes its typical dressing configuration with probability $e^{-g_{\mathrm{share,eff}}}$ per draw, and the seven mutually independent channels (Appendix H) realize the complete dressing jointly with probability $e^{-7g_{\mathrm{share,eff}}}$. The defect therefore re-completes its dressing once per $e^{7g_{\mathrm{share,eff}}}$ substrate ticks. Identifying the defect's internal phase clock, $\tau_e = \hbar/m_e c^2$, with this recurrence period—one phase cycle per complete re-dressing—and applying the transverse export factor of Appendix C.5 gives

$$\frac{\tau_e}{\tau_*} = \frac{2}{3}\, e^{7g_{\mathrm{share,eff}}}, \qquad \tau_* = \frac{L_*}{c},$$

the length-setting relation below in temporal form. The linearity of the support exponent is forced rather than chosen: probabilities of independent rare events multiply, so their exponents add, and the waiting time for the joint event is the exponential of the sum. The three-way bracketing of the exponent at unity within 10^{-3}, summarized in Section 13 and tested again in Appendix I.1, is the quantitative check of this derivation. One identification sits beneath it—that the defect's phase advances one cycle per dressing recurrence—which Appendix H.6 decomposes into the recurrence rate, a unit committed charge per pass, and the condensate phase relation, relocating the residual to the Fork A identification (the winding phase is the condensate's) and the single integer the lepton ladder tests.

The transfer-support dual. The same relation follows from a spatial reading. The entropy of the one-sector transfer distribution fixes an effective support size $N_{\text{eff}}^{(1)} = e^{g_{\text{share,eff}}}$, and on the lightest charged branch the seven sectors carry no mutual information (Appendix H), so the joint support is $N_{\text{eff}}^{(7)} = e^{7g_{\text{share,eff}}}$. A coherent transfer orbit over N cells of step length a has first phase mode of wavelength $\lambda_1 = aN$; with the transverse export fixing the step at $a = \frac{2}{3}L_*$ (Appendix C.5), the first mode over the typical support is $\lambda = \frac{2}{3}L_* e^{7g_{\text{share,eff}}}$, the recurrence relation above in spatial form. The two readings are reciprocal faces of one asymptotic-equipartition fact—a designated configuration within a typical set of size e^S is reached with probability e^{-S} per draw, and traversed at unit step in e^S cells—and they share a single residual identification: the temporal form assumes one phase cycle per dressing recurrence, the spatial form that the Compton mode is the first phase mode over the typical support. These are the same statement in two grammars, so the branch carries a single residual identification rather than two — the one Appendix H.6 discharges into the recurrence rate, a unit committed charge per pass, and the condensate phase relation, conditional on Fork A.

Locality cannot replace the refresh kernel. The ensemble entering both readings is reached only by the memoryless refresh of Postulate III, and the requirement is topological. Under any local dynamics—single-label shifts preserving admissibility—the state graph of one parity copy fractures into $4! = 24$ disconnected ordering sectors of $\binom{7}{4} = 35$ states each, because a unit shift can never reorder two slots without passing through the collision injectivity forbids. Slot ordering is a frozen charge of local motion: no local transfer operator has p_{η_*} as its equilibrium, the strength-weighted walk measure saturates at entropy 7.374 against $g_{\text{share,eff}} = 7.4198$, and the dominant transfer mode localizes within a single sector. The full count is available only to a kernel that redraws from the ensemble, which is the content of Postulate III; the postulate is therefore load-bearing in the scale chain by necessity.

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

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

so the closure-saturation factor is

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

The faithful sector-resolution principle is

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

Equivalently,

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

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

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

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

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

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

Substituting the sector-resolution relation gives the closed form

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

Numerically,

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

about 1.053% below the CODATA value. The uncertainty imported from λ_e is many orders of magnitude smaller than this residual, so the difference is theoretical rather than metrological. The residual carries no statistical interpretation within the chain: $g_{\text{share,eff}}$ is an exact weighted evaluation, $G_* \propto e^{-14g_{\text{share,eff}}}$, and matching the CODATA central value would require $\delta g/g \simeq 7.5 \times 10^{-4}$, a shift the 1680-state combinatorics does not admit. The residual therefore measures a physical correction omitted at this order. The charged-lepton ladder residuals of -0.51% and -0.65% (Section 13, Appendix I) cannot share an algebraic origin with it, because the ladder contains no factor of $g_{\text{share,eff}}$; the framework carries two independent systematic residuals, each assigned to dressing physics beyond leading order. All three residuals are negative, a regularity that constrains candidate corrections without establishing a common origin. A candidate correction enters the closure ledger only by predicting the sign and magnitude of the residual it addresses in advance of any fit. That standard is met at the level of sign by the framework's most natural correction channel. Once the dressing is read as a recurrence of independent draws (Appendix D.4), the dressing exponent fluctuates about its ensemble mean, and convexity fixes the direction of every induced shift: for fluctuating exponent h, $\langle e^h\rangle \geq e^{\langle h\rangle}$ and $\langle e^{-h}\rangle \geq e^{-\langle h\rangle}$, so the combinatorial mass ratios shift upward and, through the lengthening of L_*, the induced G_* shifts upward as well. The channel admits no opposite-signed member, and all three measured values exceed their combinatorial predictions: three deficits, one forced sign, three matches. The magnitudes are not yet derived; the branch variances are computable from the ensemble statistics with no free parameters, and the resulting ratio predictions δ_τ/δ_μ and $\delta_{\ln G}/\delta_\mu$ are the most direct test this sector can state in advance of measurement.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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 $g_{\rm share,eff}$. 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 $m = -3, \ldots, 3$.

Let

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

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

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

The displayed order is only a representative of the symmetrized one-pass class. It does not introduce a physical ordering of the face sectors, nor does it multiply the support by an extra factor of 7!. What matters is that the ground state of the elementary fermionic defect resolves each simple face sector once: fewer sectors leave unresolved label memory, while repeated sectors describe excited or additionally dressed states rather than the elementary anchor.

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

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

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

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

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

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

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

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

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

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

The status of the relation should be named precisely. The effective dimension e^H is the perplexity of the coherent dressing — the effective number of mutually resolvable support positions in the defect's history, not a count of cells in a spatial volume. Reading it as a physical length ratio is the recurrence identification derived above: the dressing re-completes once per $e^{7g_{\mathrm{share,eff}}}$ ticks, and one phase cycle per recurrence converts that waiting time into the Compton scale, so the linear, first-power map is the arithmetic of joint probability rather than a chosen dictionary. The map is cross-validated by the induced gravitational scale, by the charged-lepton ladder, and by the mass-minimization argument above, which together bracket the support exponent at unity within 10^{-3}. One identification remains beneath the derivation — one phase cycle per dressing recurrence, equivalently that the Compton mode is the first phase mode over the typical support — and a one-defect transfer operator T whose coherence length obeys $\lambda = C_{\mathrm{cl}} L_* \, e^{H[p_T]}$ is the object whose spectrum would certify it; exhibiting that operator is the branch's remaining microscopic completion.

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

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

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

D.5 EFT consistency checklist

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

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

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

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

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

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

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

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

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

D.6 Quadratic fluctuations and weak-field stability

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

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

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

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

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

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

Appendix E collects the time-dependent and homogeneous extensions of the static branch: the same scalar medium propagating causally, relaxing toward its static limit, and supporting a cosmological background mode without losing contact with the weak-field structure already derived. The status differs across these pieces: the transport relation is closed at the preferred-branch level ($D/\tau_0 = c^2$ with $\tau_0^{-1} = H_0$); the homogeneous Hubble-tension mechanism is directionally supported but not Boltzmann-closed; and the full perturbation likelihood treatment remains open.

E.1 Telegrapher equation and causal closure

The time-dependent deficit field obeys

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

Causality requires

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

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

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

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

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

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

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

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

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

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

E.3 Homogeneous mode and cosmological sourcing

The cosmological split is

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

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

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

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

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

E.4 Sound-horizon shift and shear lock

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

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

The transport relation and preferred branch are closed; the cosmological sector remains structurally supported but not yet Boltzmann-closed.

Appendix F: Constrained-Capacity Strong-Field Branch

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

F.1 Bounded capacity and the unique lapse map

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

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

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:

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

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

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

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

F.2 Constrained-capacity action on the physical domain

The physical strong-field domain is

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

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

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

The standard ADM constraint densities are

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

The same content can be written in Lagrangian form as

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

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

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

$$K_{ij} = \frac{1}{2N}\left(\dot{h}_{ij} - D_iN_j - D_jN_i\right)$$

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

Varying λ gives

$$N^2 = q.$$

The q variation gives, schematically,

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

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

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

so in bulk vacuum

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

The shift variation gives the ordinary momentum constraint,

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

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

F.3 Horizon boundary action and boundary conditions

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

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

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

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

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

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

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

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

F.4 Static spherical vacuum exterior

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

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

so the capacity variable is

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

with

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

This also agrees with the weak-field capacity deficit:

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

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

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

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

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

which in Schwarzschild gives

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

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

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

The substrate variable q and the geometric areal-radius gradient invariant therefore agree as a theorem of the framework's structural commitments, not as an additional axiom. The marginal-trapped-surface condition $\theta_+\theta_- = 0$ on a closed two-surface coincides with $q = 0$ independently of the null normalization used to define $\theta_\pm$: $q > 0$ is the untrapped exterior, $q = 0$ is the marginal / apparent horizon, and $q < 0$ would be the classical trapped interior. The capacity domain $\mathcal{M}_q = \{q > 0\}$ is the untrapped exterior, and the saturated boundary $\partial\mathcal{M}_q = \{q = 0\}$ is the apparent-horizon two-surface. Standard apparent-horizon formation theorems for gravitational collapse therefore produce the $q = 0$ saturation surface inside the framework, without an additional substrate-transport postulate; the bounded transport law of Appendix F.7 is a substrate-side consistency check on this geometric formation rather than its primary mechanism.

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

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

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

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

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

F.5 Horizon thermodynamics and boundary capacity

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

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

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

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

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

Per-channel cut entropy from fermionic face exclusion. At a saturation surface $q = 0$, the physical continuum domain terminates at $\partial\mathcal{M}_q$. A tetrahedral cell on the exterior side whose outward face would, in the unsaturated bulk, be paired with a neighboring cell across that face instead has an unfilled pairing slot, because the would-be partner lies outside $\mathcal{M}_q$. By Postulate II the elementary face slot is fermionic and admits only the occupied (paired) state or the excluded (unpaired) state. A horizon cut face is therefore an instance of the same elementary face-exclusion defect that anchors the one-bit fermionic sector in the bulk. The seven-state $j_{\text{eff}} = 3$ channel belongs to the paired bulk link $j_0 \otimes j_0 \to j_{\text{eff}}$, not to the cut face itself: with the partner cell absent, the paired representation is never formed. The seven-state structure enters the boundary count through the bulk graph response and hence through the channel density below, not through the per-channel entropy. The entropy carried by an elementary cut defect is consequently the same primitive fermionic increment that anchors the electron at $\kappa_m(\lambda_e) = m_e/\ln 2$,

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

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

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

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

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

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

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

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

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

the product of the two substrate inputs is

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

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

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

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

F.6 Absorption, ringdown, and echoes

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

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

so the leading reflectivity is

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

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

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

$$q = \epsilon > 0,$$

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

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

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

F.7 Dynamical formation as a free-boundary problem

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

$$\partial_t q + D_i J^i = -\Gamma(q)\Sigma[T_{\mu\nu}],$$

$$\tau_J(\partial_t + \mathcal{L}_v)J^i + J^i = -D(q)D^i q.$$

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

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

Choosing constitutive functions such as

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

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

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

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

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

$$V_n = -\frac{\partial_t q}{|\nabla q|}\bigg|_{q\to 0^+}.$$

In spherical symmetry this becomes

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

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 $n_{\rm hor}S_\infty^{\rm cell} = 1/4$. These are not new fit knobs; they are the numerical and microscopic tests that would close the dynamical and boundary sectors.

F.8 PPN boundary and weak-field breakdown

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

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

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

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

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

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

Appendix G: Many-Pasts, Operational Closure, Branch Realization, and the Arrow of Time

Appendix G is the technical support layer for the Many-Pasts postulate (Sections 3.3, 21). It defines the three objects the postulate weights — a history, the present, and the distance between them — then derives the operational consequences (Born recovery, no-signaling), explains branch realization and the arrow of time, connects the same weight to the scale-setting branch through memoryless dressing, and locates the proposal among the standard interpretations. The operational core is conservative: it reproduces laboratory quantum mechanics exactly. The ambition is in the ontology built on that core.

G.1 What is a history of the entanglement network?

A history H is a coarse-grained trajectory of the substrate's microstate: a sequence of admissibility-closed configurations of the entanglement network, one per substrate tick, each obtained from the previous by an allowed update. "Admissible" carries the same meaning as in the UV ensemble of Appendix B — a configuration close to a regular, isotropic local cell. A history is not a path through an external spacetime, since there is none at this level; it is a path through the network's own configuration space. The set of histories compatible with a fixed present is large, and the postulate is a rule for weighting them.

G.2 What is the present coarse configuration P?

The present P is the macroscopic, record-bearing coarse-graining of the network now. It is the quantity held fixed: every admissible history must terminate in a microstate consistent with P. Operationally P is represented by a projector Π_P onto the subspace of microstates carrying the present records. The coarse-graining is the ordinary one of statistical mechanics — macroscopically distinguishable configurations — not a new structure introduced for this purpose.

G.3 What is the distance $D(H, P)$?

The history-space distance is

$$D(H,P) = -\ln \operatorname{Tr}\big(\Pi_P \, \rho_{H\to\text{now}}\big),$$

where $\rho_{H\to\text{now}}$ is the state obtained by evolving the history H forward to the present. It is the negative logarithm of the overlap between the present projector and the state that history would produce. A history leading naturally to the present records sits close, with small D and large weight; a history that would require an improbable conspiracy to reproduce them sits far, with large D and suppressed weight. The weight $P(H|P) \propto e^{-D(H,P)}$ is then exactly that overlap, normalized over the admissible histories.

G.4 Born recovery in the projective laboratory limit

In the ordinary projective laboratory limit the evolved state is the prepared state and the present projector is the measured outcome, so the distance collapses to a single overlap,

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

which is the Born probability of that outcome. This is the branch $\alpha = 1$ of the generalized weighting family. The recovery is a consistency condition rather than a fresh derivation: requiring that the weight reproduce exact Born statistics selects $\alpha = 1$, and the substantive content is that the family contains a branch for which this holds at the same time as no-signaling. Nothing in the laboratory is changed; the operational theory is standard quantum mechanics.

G.5 No-signaling

The companion condition forbids any operational bias channel whose statistics could depend on a spacelike-separated setting. Imposing it selects $\beta = 0$. With $\beta = 0$ the weight depends only on the local present records, so a choice made at one location does not alter the marginal statistics at another, and no-signaling holds exactly in the operational branch. Born recovery and no-signaling are therefore the two consistency conditions $\alpha = 1$, $\beta = 0$; the postulate's nontrivial claim is that they are simultaneously satisfiable, not that either is derived from outside.

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

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

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

Under coarse-grained factorization,

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

so the dominant record-compatible histories are those with entropy increasing toward the present. The arrow of time appears as a counting-dominance effect among record-compatible histories, not as a new dynamical coupling or a new law of laboratory probability. It is conditional typicality: given the records, overwhelmingly many admissible pasts show entropy rising toward them.

G.8 Memoryless dressing and the connection to L_*

The same weight does work outside the foundations. Because e^{-D} weights whole histories with no memory term, the dressing of an elementary defect is memoryless in substrate time: between one readout and the next the substrate relaxes to a fresh draw from the admissibility ensemble, so each dressing pass is statistically independent of the last. That independence is exactly the input the scale-setting recurrence of Appendix D.4 requires: it lets the seven channels each carry the full sharing entropy $g_{\text{share,eff}}$ and multiply. The mixing-time hierarchy that justifies treating each pass as a fresh draw is the $\sim 10^{22}$ ratio of the electron Compton time to the substrate tick (Appendix H.4). Postulate III therefore enters the gravitational chain directly: its memoryless dressing is load-bearing in the length that sets G_*.

G.9 Relation to many-worlds, collapse, hidden variables, and decoherence

Placed among the standard interpretations, Many-Pasts is a form of history-space realism. Unlike many-worlds, it posits no co-real forward branches; the only realized macrostate is the present. Unlike objective-collapse models, it adds no stochastic collapse term and leaves unitary evolution intact. Unlike hidden-variable theories, its probabilities are not ignorance over pre-existing sharp values. It is nearest to the consistent-histories framework and to decoherence-based accounts of classicality, but it differs from them in two ways: it carries an explicit normalizable weight $e^{-D(H,P)}$ over admissible histories, and it puts that weight to physical use outside the measurement problem, in the memoryless dressing of G.8. The proposal does not claim to out-derive the Born rule from more primitive axioms; it recovers Born statistics and no-signaling as the consistency conditions $\alpha = 1$ and $\beta = 0$ on the weight (G.4, G.5), and offers in addition a history-space account of branch realization and the arrow of time.

G.10 Familiar quantum examples: double-slit, EPR/Bell, and measurement

In the familiar quantum examples, Many-Pasts changes the ontology rather than the laboratory calculation. In a double-slit experiment the particle is not treated as having secretly chosen one classical path while the other was unreal. Before a which-path record exists, the present detection event is supported by an ensemble of admissible microscopic histories through the apparatus, and the histories compatible with the final record include the alternatives whose amplitudes interfere in ordinary quantum mechanics. When a which-path record is made, the admissible history ensemble is conditioned on that record, the cross terms are removed by decoherence, and the interference pattern disappears. The probabilities are the usual Born probabilities; Many-Pasts supplies the history-space reading of why those are the probabilities being counted.

The same point applies to entangled pairs. In an EPR or Bell experiment the two detector outcomes are not produced by a signal sent from one wing of the experiment to the other. The present record is a joint record of the pair and the apparatus, and the admissible microscopic histories are weighted as histories of that whole entangled system. This gives the standard non-classical correlations while preserving the local Born marginals, and because the local marginal statistics are unchanged, no observer can use the correlations to send a signal. The "spooky" feature is therefore not a superluminal force in spacetime; it is the global conditioning of the history ensemble on the shared entangled record.

Measurement is read the same way. A measurement does not add a new physical collapse law; it creates a durable record, and the relevant history ensemble is then the ensemble compatible with that record. Operationally this is ordinary quantum mechanics. Ontologically, the realized branch is not selected by a forward-propagating collapse event but by conditioning the present on its admissible microscopic pasts.

None of this is a new derivation of the quantum phenomena. In the operational branch the calculations and the probabilities are exactly those of standard quantum mechanics; what Many-Pasts adds is the ontology under that branch — a history-space account of what is being counted and why — with the double slit reading as unrecorded alternatives that remain in the admissible ensemble until a record conditions it, Bell correlations as the global weighting of a shared entangled record rather than a faster-than-light signal, and measurement as record-conditioned history selection rather than new collapse physics.

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 $\phi(g_1, \ldots, g_4)$ 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 one-pass ground state

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

The seven face-label channels of the admissibility-closed ensemble define the orthogonal projectors E_m, $m = -3, \ldots, 3$, 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

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

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

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

Factorized cloud and the additive support. The dressing dimension multiplies across channels — so that the seven equal contributions $g_{\text{share,eff}}$ add rather than merge — precisely when the joint cloud state is a product, $\rho = \bigotimes_m \rho_m$. Whether the ground state of H has this product form is controlled by the inter-channel term $V_{mm'}$. Beyond the one-pass count, then, faithful sector resolution turns on the seven-channel cloud being a product of full-entropy channels: each channel must sample its whole ensemble, and $V_{mm'}$ must leave the channels uncorrelated.

The dressing Hamiltonian thus reduces the sector-resolution principle to that product structure in its ground state. The next three subsections establish it.

H.3 Slot coupling, layer factorization, and the additivity theorem

Two distinct correlation structures appear in the closure data, and separating them is essential: conflating them leads to the false conclusion that the cloud cannot factorize.

The closure invariant is a pure pair coupling. Expanding $S^2 = \big(\sum_i m_i\big)^2 = \Sigma^2 + 2\sum_{i<j} m_i m_j$ 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,

$$I(\text{slot}_0; \text{slot}_1) = 0.1545 \text{ nats}, \qquad \frac{I}{H(\text{slot})} = 0.079,$$

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 *within* 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 *between* the factors

The substrate's intrinsic correlation, the most natural candidate obstruction to factorization, therefore acts at the wrong level to obstruct it: it is internal to a layer, not between layers.

Additivity of the history distance. The Many-Pasts history weight of Appendix G.1 is a log-overlap, $D(H, P) = -\ln \operatorname{Tr}(\Pi_P \rho_{H\to\text{now}})$. If the dressing state and the resolution projector factorize over channels, $\rho = \bigotimes_m \rho_m$ and $\Pi_P = \bigotimes_m \Pi_m$, the trace factorizes and the logarithm converts the product into a sum,

$$\operatorname{Tr}(\Pi_P \rho) = \prod_m \operatorname{Tr}(\Pi_m \rho_m) \quad \Longrightarrow \quad D = \sum_m \big[-\ln \operatorname{Tr}(\Pi_m \rho_m)\big] = \sum_m D_m.$$

Additivity of D over the seven channels — and hence the multiplicativity $\dim_{\text{eff}} = \prod_m e^{g_{\text{share,eff}}} = e^{7g_{\text{share,eff}}}$ used in the length relation — is therefore a theorem once the product structure holds. The implication needed downstream runs in one direction: a product cloud gives additive D and the full $7g_{\text{share,eff}}$. Whether the electron's dressing is that product is decided by two conditions, established in the next two 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

$$H\big(B_{-3}, \dots, B_3\big) \le \sum_{m=-3}^{3} H(B_m) = 7g_{\text{share,eff}},$$

and the deficit

$$\Delta = \sum_m H(B_m) - H\big(B_{-3}, \dots, B_3\big) \ge 0$$

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 them.

H.4 Each channel at full entropy: memorylessness in substrate time

A channel reaches its full entropy $g_{\text{share,eff}}$ only if its dressing keeps no memory of its own past. To see what memory would cost, consider the one-parameter family of per-channel kernels

$$K_a(b, b') = a\,\delta(b, b') + (1-a)\,p_{\eta_*}(b'), \qquad a \in [0, 1],$$

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 $H(b' \mid b) = \sum_b p_{\eta_*}(b)\, H\big(K_a(b,\cdot)\big)$, evaluated on the exact 1680-state ensemble, nonetheless slides with the memory a:

a (memory)	per-channel $H(b' \mid b)$	status
0.0 (refresh)	$7.41980 = g_{\text{share,eff}}$	full entropy
0.1	6.99953	reduced
0.3	5.80153	reduced
0.5	4.40051	reduced
0.9	1.06642	reduced

Only the memoryless endpoint $a = 0$ — the refresh kernel $P_{\eta_*}(b, b') = p_{\eta_*}(b')$ of Appendix D.4 — returns the full $g_{\text{share,eff}}$. The kernel carries more than its stationary marginal, so this first condition is settled by the dressing dynamics, not by the equilibrium ensemble. The requirement is in fact topological: single-label local moves conserve slot ordering, fracturing each parity copy into twenty-four sectors of thirty-five states, so no local kernel reaches the full ensemble at any parameter value (Appendix D.4).

The dynamics settle it cleanly. Take any ergodic generator Q with p_{η_*} as its detailed-balance stationary weight — a continuous-time single-label-swap generator built from the same admissibility weights serves, with no memorylessness assumed. The kernel over substrate-time τ is $e^{Q\tau}$, and the mutual information between its endpoints decays to zero,

$$I(\tau{=}0.1) \approx 3.4 \text{ nats}, \qquad I(\tau{=}1) \approx 0.030, \qquad I(\tau{=}2) \approx 1 \times 10^{-4},$$

so $e^{Q\tau}$ relaxes to the refresh kernel. The readout sits far inside that limit: the substrate time is $\tau_* = L_*/c \approx 5.4 \times 10^{-44}$ s, the electron's Compton time is $\tau_e = \hbar/(m_e c^2) \approx 1.3 \times 10^{-21}$ s, and their ratio $\approx 2.4 \times 10^{22}$ is the same hierarchy λ_e/L_* the length relation expresses. The ground state is read out some 10^{22} mixing times after relaxation, so each channel draws afresh and carries its full $g_{\text{share,eff}}$. This is the Many-Pasts weight $P(H|P) \propto e^{-D(H,P)}$ seen on the dressing: a sum over pasts that keeps no trace of the particular one realized.

H.5 Independent channels: the electron as the lightest defect

The second condition is that the seven channels are mutually independent at one readout, so that the deficit Δ of H.3 vanishes. This needs no separate dynamical postulate; it follows from what the electron is.

Keep the deficit explicit in the support relation. With each channel at its full entropy, the joint entropy is $7g_{\text{share,eff}} - \Delta$, and the support reads

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

At a fixed substrate scale L_*, a one-bit charged defect whose dressing carries correlation Δ has spatial support $\lambda \propto e^{-\Delta}$, and through $m = \hbar/(c\lambda)$ a mass

$$m \propto e^{\Delta}.$$

Correlation among the dressing channels contracts the defect and raises its mass. Because $\Delta \geq 0$, the lightest one-bit charged defect is the one with $\Delta = 0$: the product dressing. The electron is that defect, so its seven channels are independent and its support is the full $\dim_{\text{eff}} = e^{7g_{\text{share,eff}}}$. Correlated dressings are still allowed; they describe heavier, additionally dressed excitations, not the elementary anchor.

The selection is physical, not an appeal to ignorance. Independence is not read off from the absence of an inter-channel constraint in the ensemble; it is forced because correlation carries a mass, and the electron is defined as the lightest charged one-bit defect. The argument uses only the mass–entropy identification at the centre of the theory: the same relation $m = \kappa_m \, \Delta S$ that anchors the electron makes inter-channel correlation an excess of defect entropy, and so of mass.

With both conditions met, the scale-setting branch reads end to end without a free step. One-pass occupation fixes the seven channels (H.2); memorylessness in substrate time gives each its full entropy $g_{\text{share,eff}}$ (H.4); the electron's lightness forces the channels independent (this subsection). Hence

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

with no gravitational quantity used as input.

The condensate realization gives the EFT kinetic term a natural microscopic origin, and the dressing analysis derives the faithful sector-resolution principle from the mass–entropy ontology, the one-bit electron anchor, and the memoryless dressing, so the substrate length and the induced scale follow from commitments the theory already makes. The first-principles derivation of every inhomogeneous continuum coefficient from the full kernel remains open.

H.6 Mass as a rate: the dressing operators

The scale-setting branch above fixes the support exponent $7g_{\text{share,eff}}$ and hence the ratio λ_e/L_*. Reading that ratio as a physical length used one identification: that the electron advances one Compton phase cycle per complete dressing recurrence. Memoryless sampling delivers a recurrence *rate*, not by itself a Compton pole, so this subsection gives the recurrence an explicit operator form and asks what the phase clock must be. The identification does not close to zero residual; it decomposes into pieces that are either forced or falsifiable, with one identification (Fork A below) remaining.

Mass is a rate, not an energy. Writing the length relation as a mass,

$$m_e c^2 = \frac{\hbar c}{\lambda_e} = \frac{3}{2}\,\frac{\hbar}{\tau_*}\, e^{-7g_{\text{share,eff}}}, \qquad \tau_* = L_*/c, \qquad e^{-7g_{\text{share,eff}}} = 2.78 \times 10^{-23},$$

shows the electron is lighter than the substrate quantum $E_* = \hbar c/L_*$ not by a small factor but by $e^{-7g_{\text{share,eff}}}$. That suppression is a waiting time: seven independent full-entropy channels jointly reaching their resolved configuration, probability $e^{-7g_{\text{share,eff}}}$ per substrate tick. This rules out a binding-energy reading. Seven channels each releasing free energy of order $g_{\text{share,eff}} E_*$ would give a mass linear in $g_{\text{share,eff}}$, of order $50\, E_*$ — heavy, and wrong by twenty-three orders. The object that produces an exponentially light particle is the spectral gap of a stochastic generator whose slowest mode relaxes at $e^{-7g_{\text{share,eff}}}$ per tick, not the ground-state energy of a Hamiltonian. This rate reading sets the scale, not the content of the mass–entropy map: the electron's committed entanglement is the one bit $\Delta S_f = \ln 2$ of Section 13.2, while the recurrence rate fixes the exchange rate κ_m at which that bit is read as a mass (the length anchor of Section 13.3). Mass measures committed entanglement; the rate fixes how small that measure is.

The killed-refresh return rate. By Postulate III the dressing is memoryless: the refresh kernel $P(b, b') = p_{\eta_*}(b')$ redraws each tick from the admissibility Gibbs state. On the seven-layer history space the marked configuration $\star$ — the resolved dressing, each channel at its designated typical configuration — has stationary weight $\prod_m p_{\eta_*}(\star_m) = e^{-7g_{\text{share,eff}}}$. The dynamics of "at $\star$ versus not" is a two-state renewal chain with per-tick return probability $e^{-7g_{\text{share,eff}}}$; the killed (absorbing-at-$\star$) operator has leading eigenvalue $1 - e^{-7g_{\text{share,eff}}}$, so its gap is the raw return rate

$$\Gamma_{\text{raw}} = \frac{e^{-7g_{\text{share,eff}}}}{\tau_*}, \qquad \langle T_{\text{rec}} \rangle = e^{+7g_{\text{share,eff}}} = 3.6 \times 10^{22} \text{ ticks}.$$

The physical electron Compton rate carries the transverse export factor 2/3 of the length relation, $\Gamma_e = \frac{3}{2}\Gamma_{\rm raw}$, so $m_e c^2 = \hbar\Gamma_e = \frac{3}{2}(\hbar/\tau_*)e^{-7g_{\rm share,eff}}$, consistent with $\lambda_e/L_* = \frac{2}{3}e^{7g_{\rm share,eff}}$. This is the recurrence of the scale-setting branch re-expressed as a transfer-operator gap; the rate scale was never the contested part.

The locality obstruction. The generator that carries this recurrence cannot be local. Define a local move as shifting one face label by ±1 while preserving injectivity. Under these moves the 1680 states fracture into

$$48 \text{ components of } 35 = (4! \text{ orderings}) \times (2 \text{ orientations}) \times C(7,4),$$

verified directly on the ensemble. Slot ordering is a conserved charge of any local single-label dynamics, so the closed seven-sector loop that winds the phase once is not reachable by a local Hamiltonian; the generator that closes it is the non-local refresh kernel of Postulate III. This is an honest cost of the construction: a referee asking for a local "Hamiltonian spectrum" is owed the statement that the operative generator is the non-local refresh super-operator. It is the content of memorylessness rather than a device introduced to evade locality.

The phase clock, by elimination. The mass is $\hbar$ times a phase-turn rate, so closing "one turn per recurrence" requires the turning clock and the refresh counter to be the same object, not two rates set equal. Two candidates exist: (B) a clock private to the defect, or (A) the global U(1) phase θ of the condensate order parameter $\sigma = \sqrt{n}\,e^{i\theta}$. A private clock is not forced to complete exactly one turn per recurrence; pinning it to the refresh would require declaring the rates equal, the identification one is trying to remove. The condensate phase is forced by counting: one complete refresh commits exactly one defect (the lightest charged defect binds the seven sectors once — one fermion, Postulate II), one committed defect is one unit of the condensate U(1) charge (that charge being the count of committed cells), and one unit of charge on a U(1) phase is one full turn. Hence

$$\text{one refresh} = \text{one defect} = \text{one unit charge} = \text{one } 2\pi \text{ turn},$$

each link a count or a definition, with no rate chosen. A lone defect has no conserved tally for the refresh to increment, so the counting argument is unavailable to it: the clock is the medium's.

The charge-tilted generator. Tilting the renewal generator by a counting field ϕ conjugate to committed charge gives the reduced tilted transfer matrix $\left[\begin{smallmatrix} 1-r & re^{i\phi} \\ 1 & 0 \end{smallmatrix}\right]$ with $r = e^{-7g_{\rm share,eff}}$. Its leading eigenvalue is exactly 2π-periodic in ϕ, so committed charge is integer — the fermionic $n_m \in \{0,1\}$ appearing as periodicity, not a tuned input — and a complete seven-sector pass is one closed loop on the label cycle, winding number one. The identification "one phase cycle per dressing recurrence" therefore decomposes into three pieces, none carrying a knob:

$$\underbrace{\Gamma = e^{-7g_{\rm share,eff}}/\tau_*}_{\text{memorylessness+entropy}} \;\oplus\; \underbrace{\Delta Q = 1 \text{ per recurrence}}_{\text{fermionic one-pass}} \;\oplus\; \underbrace{\Delta\theta = 2\pi\,\Delta Q}_{\text{Josephson}}.$$

The Compton clock is now a derived winding rate rather than an asserted one.

What remains: Fork A. Two things do not close, and are stated as such. First, the identification of the winding U(1) with the condensate phase, and its 2π with the Compton cycle (Fork A), is consistent with these operators but not forced by them; its support is the elimination argument above, not a computation, and it is settled in H.7, where the mean-field condensate is shown to carry exactly one phase. Second, the residual is now a single integer — that the electron is the $\Delta Q = 1$ (one-loop) recurrence rather than a higher winding. Unlike the original identification this is falsifiable: the winding-N recurrences are the muon and tau, and their rates fix m_μ/m_e and m_τ/m_e with no new input (Appendix I.1). The phase-cycle identification has thus been replaced by counting plus a condensate relation, with the residual reduced to one integer that the lepton spectrum constrains.

H.7 The dressing functionals from the mean-field condensate

The dressing Hamiltonian named in Section 22,

$$H = \sum_m \big(\varepsilon_0\, n_m - n_m F_m(\rho_m)\big) + \sum_{m<m'} V_{mm'}(\rho_m, \rho_{m'}),$$

carries three named functionals ε_0, F_m, V. This subsection derives them from the mean-field condensate. The result is not that everything closes: the two residuals of H.6 collapse into a single premise, and that premise is the one the geometry side has not yet established.

Mean-field reduction. In the GFT condensate the single-tetrahedron expectation is the order parameter over boundary data, $\langle\hat{\varphi}(b)\rangle = \sigma(b) = \sqrt{n(b)}\, e^{i\theta}$, with one global phase θ — the Goldstone of the broken GFT number U(1). Extremizing the condensate free energy under the closure constraint gives the admissibility Gibbs state $|\sigma(b)|^2 = n\, p_{\eta_*}(b)$, so the closure weight is a Boltzmann factor at inverse temperature η_* with K^2 as energy. The substrate temperature and closure condition are then thermodynamic,

$$T_* = \frac{1}{\eta_*} = 33.48, \qquad \langle K^2\rangle_{\eta_*} = \tfrac{3}{2} T_* = 50.22,$$

the 3/2 being equipartition over the three quadratic closure components, each at $\tfrac{1}{2}T_*$; this is the identity $\eta_*\langle K^2\rangle = 3/2$ read as a temperature.

Two generators: H selects, L rates. The mass scale does not live in the spectrum of H. H is combinatorial: its ground state fixes *which* configuration the lightest defect is. The rate is the gap of the dissipative refresh super-operator L whose stationary state is p_{η_*} and whose memorylessness is Postulate III, and the physical mass is that gap carrying the transverse export 2/3 of the length relation, so

$$m_e c^2 = \tfrac{3}{2}\,\hbar\, \mathrm{gap}(L) = \frac{3}{2}\frac{\hbar}{\tau_*}\, e^{-7 g_{\mathrm{share,eff}}}, \qquad \mathrm{gap}(L) = \frac{e^{-7 g_{\mathrm{share,eff}}}}{\tau_*}, \qquad e^{-7 g_{\mathrm{share,eff}}} = |\sigma(\star)|^2 = \prod_m p_{\eta_*}(\star_m),$$

the exponential smallness being the condensate density at the resolved configuration, not a fine-tuned eigenvalue. This is the operator content of H.6.

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 $E_* = \hbar c/L_*$. Deriving it from the substrate names it as μ_0: it is the single irreducible dimensionful input, relocated rather than removed, since a mass cannot come from combinatorics alone. Energies in H are measured in units of E_*, so ε_0 is the one place the physical scale is attached, and the substrate temperature $T_* = 1/\eta_*$ entering F_m below is the dimensionless closure temperature (K^2 as energy); the physical mass is set by the rate L, not by the spectrum of H.

F_m, the dressing free energy. A bound channel carries a cloud ρ_m, a distribution over the boundary ensemble; the unbound channel carries none. The released free energy is the cloud entropy at the substrate temperature, $F_m(\rho_m) = T_*\, S[\rho_m]$. Extremizing it at fixed mean closure returns the Gibbs state $\rho_m = p_{\eta_*}$ with $S[\rho_m] = g_{\mathrm{share,eff}}$: the equilibrium cloud is the condensate density itself. The binding test $F_m > \varepsilon_0$ is met for all seven channels, and the lightest defect binds each sector once with a full-entropy cloud — one-pass occupation, total bound entropy $7g_{\text{share,eff}}$. This last equality bridges the two generators: the cloud entropy that H maximizes in binding is exactly the rarity exponent that L carries in its gap, since the joint recurrence needs all seven full-entropy clouds at once, probability $\prod_m e^{-S[\rho_m]} = e^{-7g_{\text{share,eff}}}$. Binding and lightness are one extremum: the most bound defect, carrying maximal cloud entropy, is the rarest and hence the lightest. The F_m step is otherwise a consistency loop — its form is entropy, the condensate density is p_{η_*}, and the maximum-entropy condition hands p_{η_*} back — so it confirms the one-pass structure rather than independently generating it.

Two energies, kept distinct. Two quantities both invite the name "energy of the defect," and conflating them is what would reintroduce the hierarchy the rate reading removes. The first is the *separation cost*: the entropic free energy needed to unbind the defect from the network. In a medium whose only currency is shared entanglement there is no stored potential energy in the spring sense, so the cost of any operation is the free-energy cost of changing the admissible-arrangement count, and the natural energy of a defect is what it would take to dissolve it back into the substrate. That cost is large and linear in $g_{\text{share,eff}}$ — of order the $\sim 50\,E_*$ that the binding-energy reading of H.6 produces, a number that is correct as a separation cost and wrong only as a mass. The second is the rest energy $m_e c^2 = \hbar\Gamma_e$, which is the recurrence rate and is tiny, exponential in $-7g_{\text{share,eff}}$. Binding-equals-lightness is the statement that these are one bound entropy $g_{\text{share,eff}}$ read two ways: the deeper the binding, hence the larger the separation cost, the rarer the full re-completion, hence the smaller the rate. They are proportional through $g_{\text{share,eff}}$ but not equal, and the distinction is load-bearing — reading the separation cost as the rest mass would set the electron at tens of E_* and undo the exponential suppression. The cost to separate and the rate of recurrence are two faces of one entanglement, not the same number.

V, the inter-channel coupling. This is the one piece with new content. Channels couple by exchanging the field they share, and the framework has exactly one infrared field: the scalar capacity fluctuation $\phi = \delta S_{\text{ent}}$, coupled to matter through the trace channel $-\kappa\chi S_{\text{ent}}$ of Section 10. That scalar is the sector singlet $|s\rangle = \frac{1}{\sqrt{7}}\sum_m |m\rangle$, so each channel couples to it through $\langle m|s\rangle = 1/\sqrt{7}$, the same for every m. Integrating out the single shared scalar gives an outer product,

$$V_{mm'} = v\,\langle m|s\rangle\langle s|m'\rangle = \frac{v}{7} \quad \text{for all } m, m',$$

with v the zero-momentum scalar-exchange amplitude: the uniform 7×7 matrix is rank one. The per-incidence overlap that enters the ladder is $\langle m|s\rangle\langle s|m'\rangle = 1/7$, doubled by orientation to the $2/7$ of Appendix I.1, while the overall coupling v sets the mass scale and not the ratios. One exchanged mode is one outer product, which is rank one. A second light field would add a second outer product, raising the rank to two and splitting the uniform entries, which breaks the single $2/7$ and the $(2/7)^{N^2}$ ladder. So $\text{rank}(V) = 1$ is not an assumption layered onto the construction; it is the statement "one scalar field," the same premise that gives no anisotropic stress and $\Phi = \Psi$ in the PPN sector (Section 15). A higher-rank V is a second light field, which that sector already excludes. This also supplies the minimum-mass content used in H.5: rank-one is the minimal inter-channel correlation, so the lightest charged defect is forced to it, and "the coupling is rank one" and "this is the lightest defect" are the same statement.

The collapse, and the wall. The two residuals of H.6 — (A) the winding phase is the condensate U(1) and its 2π is the Compton cycle, and (B) the lightest defect's coupling is the rank-one singlet — reduce to one fact. The mean-field condensate has exactly one global phase, so the mode that mediates V and the mode that winds the Compton clock cannot be different: there is no second U(1) to be the wrong one, which settles the underdetermination of Fork A.

And rank-one V is "one scalar." A and B are the same statement — a single scalar field on the mean-field condensate — which is the founding premise of the framework, tested and passing in the charged-lepton sector to sub-percent. The cost is that this premise is now load-bearing twice. Every step here is conditional on the GFT substrate supplying that condensate, with density $|\sigma|^2 = p_{\eta_*}$, one U(1), one scalar infrared mode, and the emergent four-geometry to go with it — the geometry-side compatibility problem left open in Section 22 and H.1. The defect sector and the geometry sector are two outputs of the same σ: the sub-percent lepton ladder cannot be banked as evidence independent of the geometry problem, because it is the same condensate graded on its other output. The construction falls out cleanly because it rests on the single-scalar condensate, and that condensate is the unproven object; the two are one coin. Writing a fully explicit H past this point would return ε_0, F_m, V as they were put in, and the rank-one V is already as derived as it gets without the condensate itself.

Reproduction. The ensemble invariants (η_*, $\langle K^2\rangle_{\eta_*}$, $g_{\text{share,eff}}$, and the identity $\eta_*\langle K^2\rangle = 3/2$), the killed-refresh return rate, the 48 × 35 locality fracture, the 2π-periodicity of the charge-tilted generator, the rank-one spectrum $\{1, 0^6\}$ of P_{sing}, and the charged-lepton ratios of Appendix I.1 are reproduced end to end from the bare admissibility ensemble by a standalone script provided as supplementary material (requiring only `numpy` and `networkx`).

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 not new kinds of object but 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 number of generations is not put in by hand; it is limited by the closure machinery, which stops having anything to weight after the third admissible shell, so there are three charged leptons and not four. The mass of each lepton is the exponentiated dressing free energy of its shell source, in the same log-overlap form that carries the gravitational dressing (Appendix H.3). This sector is a structural extension of the mass–entropy ontology, outside the closed weak-field chain. With the dressing operators of Appendix H.6 in place it is operator-backed: the muon and tau are the winding-N recurrences of the same generator, so the ratios below follow with nothing fitted to lepton data, conditional on the Fork A identification of Appendix H.7. Because those ratios carry no dimensionful input, the sector is a parameter-free, falsifiable test of the deeper construction rather than only a cross-check.

With that picture fixed, the spectrum is written as a shell expansion,

$$\log \frac{m_N}{m_e} = B_0 N + A_0 N^2, \qquad N = 0, 1, 2,$$

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,

$$N \in \{0, 1, 2\} \qquad \Longrightarrow \qquad \text{three generations},$$

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

$$B_0 = \ln \Omega_1 = \ln 720 = 6.579\ldots,$$

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 $720 \cdot 240 \cdots$ 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 $\{1, 0, 0, 0, 0, 0, 0\}$: 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 $1680 \to 720$ 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

$$\mathcal{H}_7 = \mathrm{span}\{|m\rangle : m = -3, \ldots, 3\}, \qquad |u\rangle = \frac{1}{\sqrt{7}} \sum_{m=-3}^{3} |m\rangle, \qquad P_{\mathrm{sing}} = |u\rangle\langle u|.$$

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

$$\langle J_N|P_{\mathrm{sing}}|J_N\rangle = \sum_{r,s=1}^{N} \langle a_r|P_{\mathrm{sing}}|a_s\rangle.$$

Because $\langle a|P_{\mathrm{sing}}|b\rangle = 1/7$ 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

$$\mathcal{M}_{rs} = 2\,\langle a_r|P_{\mathrm{sing}}|a_s\rangle = \frac{2}{7}.$$

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

$$A_0 = \ln \frac{2}{7} = -1.253\ldots,$$

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 $\langle J_N|P_{\mathrm{sing}}|J_N\rangle$ runs over both indices independently because the integrated-out scalar contribution is quadratic in J_N.

Mass ladder and numerical accuracy. Combining the two coefficients,

$$\boxed{\frac{m_N}{m_e} = 720^N \left(\frac{2}{7}\right)^{N^2}, \qquad N = 0, 1, 2.}$$

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 $(1, 3, 5, \ldots,$ summing to $N^2)$. The second difference is fixed at

$$\Delta^2 \log m_N = 2\ln\frac{2}{7} = -2.506\ldots,$$

compared with the empirical -2.509 (0.15%). With no parameters fitted to lepton data,

$$\frac{m_\mu}{m_e} = \frac{720\cdot 2}{7} = \frac{1440}{7} = 205.71, \qquad \text{PDG: } 206.77 \qquad (0.5\%);$$

$$\frac{m_\tau}{m_e} = 720^2\left(\frac{2}{7}\right)^4 = 3454.56, \qquad \text{PDG: } 3477.23 \qquad (0.65\%).$$

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 $u = \frac{1}{\sqrt{7}}(1,\ldots,1)$ of Section 8 has weight $|\langle u|m\rangle|^2 = 1/7$, and orientation summation doubles it to $2/7$. The N shells form N^2 ordered pairs, so the pairwise term is $N^2\ln(2/7)$, and exponentiating the free energy returns

$$\frac{m_N}{m_e} = \Omega_1^N\left(\frac{2}{7}\right)^{N^2} = 720^N\left(\frac{2}{7}\right)^{N^2}.$$

The coefficients are fixed, not chosen: $720 = 6!$ and the N^2 pair count are combinatorial, and $2/7$ is the orientation-summed singlet projection behind $\Omega_{\text{tet}} = 2\,P(7,4) = 1680$. What remains an input is the form of the map — that the lepton mass is the exponentiated log-overlap free energy of the coherent multi-shell source — a log-overlap structure of the kind that carries the gravitational dressing, taken here as the form of the mass map.

The support exponent, tested against the spectrum. The ladder also tests the exponent in the support-to-length identification of Appendix D.4, which reads the physical ratio as the first power of the entropy support, $\lambda/L_* \propto e^H$. Letting that power float, $\lambda/L_* \propto e^{\alpha H}$, turns each lepton ratio into a measurement of α, since $m_N/m_e = \left[720^N(2/7)^{N^2}\right]^\alpha$:

$$\alpha_N = \frac{\ln(m_N/m_e)_{\text{PDG}}}{\ln\left[720^N(2/7)^{N^2}\right]}, \qquad \alpha_\mu = 1.0010, \qquad \alpha_\tau = 1.0008.$$

Read the same way, the electron-to-G_* branch needs $\alpha = 0.9999$ to reproduce the observed Newton constant, so three determinations across two sectors land within 10^{-3} of unity and bracket it. What the leptons settle is the power, not its final digit: a square-root map ($\alpha = \frac{1}{2}$) would put m_μ/m_e near 14 and a quadratic map ($\alpha = 2$) near 4×10^4, and only the linear rule reproduces the spectrum with the natural shell count $720 = 6!$. The identification that sets the substrate scale is therefore not a single assumption tuned to gravity; the same linear exponent is preferred in a sector that never enters the G_* chain.

Status and remaining audit. The three-generation termination, the linear coefficient, and the quadratic coefficient are thus fixed from machinery already native to the framework, with no parameters fitted to lepton data: the collapse at $N = 3$ is the permutation-invariance of K^2, and the 720, 2/7, and N^2 are the reduced-alphabet, singlet-projection, and ordered-pair counts of the map above. The form of the map is not a separate posit. Mass is $\hbar/(c\lambda)$ with $\lambda/L_* = e^H$, the same support-to-length identification that fixes L_* in Appendix D.4; the ladder is its $N > 0$ continuation, with the electron the $N = 0$ entry, and the equality of the support exponent across the two sectors is the $\alpha = 1$ check above. The per-shell term is the log-overlap additivity theorem of Appendix H.3, and the pairwise term is the singlet overlap of shells sharing the seven-state alphabet. Beyond that one identification, the only lepton-specific input is the structure of the excited source — a product across the six spectator channels, coherent in the shared singlet — which produces the 2/7 pairings.

The electron anchor remains the clean weak-field entry point; the heavier charged leptons are conditionally derived outputs of the same shell algebra, not closure-defining ingredients of the gravitational chain. Composite hadrons remain part of the dressed bound-state entropy program rather than a completed output of the present construction. The neutrino and quark sectors are not treated here. The same closure-spectrum-collapse argument should apply universally if the framework is right — the Standard Model has three generations across all fermion sectors — but verifying that requires extending the shell construction to the relevant defect classes.

I.2 Gauge-structure extension and Standard Model ontology

The framework's primary target is the gravitational and dark sector. Standard Model gauge structure is hosted on the substrate as external content rather than derived from it. This subsection specifies which gauge components have substrate-natural analogues, which are substrate-compatible but external, and how the baseline-redundancy template accommodates the full Standard Model gauge group $SU(3)_c \times SU(2)_L \times U(1)_Y$.

Baseline-redundancy template. The same logic that organizes the gravity sector extends to any conserved-charge sector. For a conserved charge Q, introduce an entropy-like potential $S_Q(x)$ and require that physical observables depend only on differences of that potential, not on its absolute baseline. Localizing that redundancy requires a compensating connection. In the Abelian case,

$$D_\mu S_Q = \partial_\mu S_Q - qA_\mu, \qquad S_Q \to S_Q + \alpha(x), \qquad A_\mu \to A_\mu + \frac{1}{q}\partial_\mu \alpha,$$

giving Maxwell-type dynamics for A_μ. The non-Abelian generalization, for multiplet-valued entropic potentials transforming under a compact Lie group G, gives Yang–Mills covariant derivatives and field strengths in the standard form.

The template tells us what form a gauge interaction takes once a conserved-charge sector with baseline redundancy is present. It does not select which specific groups are physically realized.

Substrate-natural components. Two ingredients of Standard Model gauge structure are naturally accommodated by the substrate.

$U(1)_Y$. The baseline-redundancy template naturally accommodates an Abelian conserved-charge sector of the hypercharge type. The associated entropy potential is conserved-charge in character, and localizing its baseline redundancy produces the corresponding Maxwell-type dynamics. The argument is independent of the gravity-sector derivation, applying the same template to a separate conserved-charge potential.

$SU(2)_L$. The substrate naturally carries an $SU(2)$-representation structure through its half-integer face data. The spin-3/2 commitment forced by maximum-capacity channel selection and tetrahedral injectivity (Section 5) carries an $SU(2)$ action on face data through the standard double-cover relation. The static K^2 ensemble breaks this $SU(2)$ to its $U(1)$ Cartan via the magnetic-quantum-number projection, but the full $SU(2)$ is recoverable at the level of the quantum face data the projection discards. Identifying this representation-theoretic $SU(2)$ with the internal weak group $SU(2)_L$, including chirality and doublet assignments, remains external to the present derivation.

Substrate-compatible but external: $SU(3)_c$. Color $SU(3)$ does not emerge from the substrate by any of the standard mechanisms (direct decomposition of the seven-state face algebra, Seiberg duality, preon compositeness, topological soliton, monopole condensation). The obstructions are structural: $SU(3)$ is not a spin group, so it does not inherit from rotational covariance; its fundamental representation requires three equivalent (non-hierarchical) states, while the substrate's natural three-fold structures (shell index, K^2 classes, residual face positions) are all hierarchical; its rank-2 Lie algebra requires two independent quantum-number axes, while the substrate provides only rank-1 axis structures; and its non-Abelian commutation relations cannot be reproduced from the substrate's scalar edge couplings.

A positive ontological reading is nonetheless available within the framework. Color is the substrate's name for an internal three-valued label on fermionic defects that lives in the non-singlet complement of the face-state algebra. The scalar gravitational branch couples to defect sources only through the permutation-invariant singlet (Appendix D), so the macroscopic scalar field is blind to non-singlet structure by construction. Color labels and color dynamics therefore live in a sector that gravity cannot resolve: the gravitational sector and the color sector occupy orthogonal subspaces of the substrate's algebra. The strong interaction can be represented, within this hosting picture, as the local gauge dynamics of the color label under the general baseline-redundancy template. The choice of $SU(3)$ specifically is empirical input; it is the realized non-Abelian structure that fits the template.

This account explains several structural features of color without claiming to derive them. Gravity's blindness to color follows from the scalar branch being the singlet projection of the substrate. Color confinement itself is not derived here; the substrate account provides that the macroscopic scalar-gravity branch resolves only singlet structure and is therefore blind to color non-singlets, which is consistent with the observed fact that hadronic macroscopic signatures are color-singlet but is not a replacement for the QCD confinement mechanism. The independence of color from generation, charge, and spin reflects the independence of the corresponding substrate sectors: color in the non-singlet complement of the face algebra, generation in the shell-excitation index (Appendix I.1), electromagnetic charge in the baseline redundancy of a separate conserved potential, and spin in the fermionic face data.

Hadronic matter and the mass–entropy bridge. The mass–entropy bridge $m = \kappa_m(\ell)\,\Delta S$ applies to fermionic defects regardless of their color label. For charged leptons, the defect is color-singlet and the dressed entropy is well-approximated by the shell-excitation budget of Appendix I.1, with mass ratios matching observation at the sub-percent level. For hadrons, the dressed entropy is dominated by the QCD-internal contributions of confinement-scale gluonic flux, trace-anomaly structure, chiral vacuum reorganization, and quark binding. The mass–entropy bridge then organizes the full dressed bound-state entropy budget,

$$S_{\mathrm{ent,H}}^{\mathrm{dressed}} = S_{\mathrm{defect}} + S_{\mathrm{bind}} + S_{\mathrm{conf}} + S_{\chi\mathrm{SB}},$$

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 closed static weak-field chain.

Appendix J: Numerical Checks and Robustness

Appendix J is intentionally modest. It does not add new derivations, and these scripts and 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 that make it easier to see that the same coefficient chain survives repeated contact with independent benchmark calculations, and it records, in Section J.3, a complete self-contained script that recomputes the entire numerical spine of the manuscript from the stated definitions.

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 $J_{\text{bare}} \to \gamma$ 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).

J.1 Cross-sector numerical checks

The cross-check program includes:

- the one-bit fermionic defect check $\Delta S_f = \ln 2$;
- the rooted-shell convergence check $\sigma_{\text{ind}}^{(2)} \simeq \sigma_{\text{ind}}^{(3)}$;
- the UV closed-branch moments $\langle K^2 \rangle_{\eta_*}$, $\text{Var}_{\eta_*}(K^2)$, and a_{UV};
- cross-sector consistency among the electron anchor, the substrate length L_*, the induced scale G_*, the Green-matched Newton closure, and the galactic scale a_0.

These checks do not replace the derivations, but they show that the same coefficient chain survives independent numerical scrutiny across the sectors where closure is claimed.

J.2 Reproducibility ledger for the substrate scale

The substrate-length calculation is short enough to record as a numerical ledger. Enumerating the 1680 oriented injective tetrahedral states with labels $m = -3, \ldots, 3$, weighting them by $e^{-\eta K^2}$, and solving

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

gives

$$\eta_* = 0.02986684439352237, \qquad g_{\text{share,eff}} = 7.419800023570903.$$

The one-pass seven-sector history support and its transverse export are then

$$e^{7g_{\text{share,eff}}} = 3.602860521062804 \times 10^{22},$$

$$\frac{2}{3}e^{7g_{\text{share,eff}}} = 2.401907014041869 \times 10^{22}.$$

Using $\lambda_e = \hbar/(m_e c) = 3.861592671986303 \times 10^{-13}$ m gives

$$L_* = \frac{\lambda_e}{(2/3)e^{7g_{\text{share,eff}}}} = 1.607719470158885 \times 10^{-35} \text{ m},$$

and hence

$$G_* = \frac{c^3 L_*^2}{\hbar} = 6.603991270884347 \times 10^{-11} \text{ m}^3\,\text{kg}^{-1}\,\text{s}^{-2}.$$

The source-side Green constant is obtained independently from the diamond-lattice integral in Appendix C.5. Uniform-grid quadrature with endpoint extrapolation gives

$$G_{\text{tet}}(0) = 0.448220394388\ldots,$$

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.

J.3 Reproduction script

The script below recomputes every load-bearing number in the manuscript from the stated definitions alone: the 1680-state K^2 spectrum and its multiplicities, the closure solution η_*, the admissibility entropy $g_{\text{share,eff}}$, the substrate length L_* and induced scale G_*, the Joyce constant and the exact $n_{\text{hor}} S_\infty^{\text{cell}} = 1/4$ identity, the edge-kernel chain, the projection coefficient ϵ and horizon target σ_*, the charged-lepton ladder with its support-exponent cross-check, and the slot mutual information. No value in the chain is entered by hand; each printed line carries the manuscript value it must reproduce. The script runs in seconds under Python 3 with `numpy` and `scipy`, and is maintained alongside the manuscript at `jaigp.org`.

```
# ============================================================
# 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, L*, 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. L*, G*
hbar=1.054571817e-34; c=299792458.0; me=9.1093837015e-31
lam_e = hbar/(me*c)
print("\nlambda_e =", lam_e, " (paper 3.861592671986303e-13)")
Lstar = 1.5*lam_e*exp(-7*H)
print("e^{7g} =", exp(7*H), " (paper 3.602860521062804e22)")
print("(2/3)e^{7g} =", (2/3)*exp(7*H), " (paper 2.401907014041869e22)")
print("L* =", Lstar, " (paper 1.607719470158885e-35)")
Gstar = c**3*Lstar**2/hbar
print("G* =", Gstar, " (paper 6.603991270884347e-11)")
G_codata = 6.67430e-11
LP = sqrt(hbar*G_codata/c**3)
print("L* vs Planck length:", (LP-Lstar)/LP*100, "% below  (paper 0.528%)")
print("G* vs CODATA:", (G_codata-Gstar)/G_codata*100, "% below  (paper 1.053%)")

# 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("\nepsilon = g/(4pi^2) =", epsv, " 1-eps =", 1-epsv, " ceiling =", 2-epsv)
print("sigma* = pi/g =", pi/H, " (paper 0.42340665)")

# 7. lepton ladder
mmu = 720*2/7; mtau = 720**2*(2/7)**4
print("\nm_mu/m_e =", mmu, " PDG 206.7683 err%:", (206.7683-mmu)/206.7683*100)
print("m_tau/m_e =", mtau, " PDG 3477.23 err%:", (3477.23-mtau)/3477.23*100)
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)")
```

Appendix K: 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. What the record establishes is this: the realized ensemble was fixed, and motivated in print by information-independence arguments, roughly ten weeks before the induced-G_* chain existed; the export factor's structural components predate the chain, its deployment is contemporaneous with the first landing, and its candidate menu was published with residuals disclosed; and both recorded routes to G land at one-percent misses rather than exact values. The Newton constant was the sole observational target of that calibration; the galactic acceleration scale, the capacity ceiling, the charged-lepton structure, and the committed-phase abundance are downstream outputs of the calibrated number, compared against measurement only after fixation. The induced G_* is accordingly treated as a calibration, not a prediction — with the calibration understood as a discrete selection rather than a continuous adjustment: within the selected construction the entropy is derived with no adjustable parameter, is reproducible from the stated rules alone, and was never re-adjusted against any subsequent observable, which is why the downstream misses stand in the record as written. The audit below quantifies how much construction freedom the discrete selection could have consumed. The evidential weight of the scale chain rests on results obtained after the ensemble was frozen — the cluster capacity bound and its twenty-four-point test (Section 17.5), the committed-phase abundance against the measured dark-to-baryonic ratio (Section 18.5), the redshift direction of the acceleration scale, and the sub-g_c rotation-curve regime — none of which entered the construction.

The audit. The substrate length depends exponentially on the realized ensemble, $L_* \propto e^{-7g_{\text{share,eff}}}$, so the calibration is meaningful only relative to the space of constructions that could have been written instead. This appendix quantifies that space. The grammar varies five discrete choices around the realized chain: the label count $M \in \{4, \ldots, 9\}$ with $j = (M-1)/2$; the multiplicity rule (ordered or unordered injective selections, with and without parity doubling, and the non-injective M^4); the closure coefficient $d \in \{1, 2, 3, 4, 6\}$ in $\langle K^2 \rangle = d/(2\eta)$; the support exponent $n \in \{4,\ M,\ 2M,\ M(M-1)/2\}$; and the export prefactor in $\{\frac{1}{2}, \frac{2}{3}, 1, \frac{3}{2}, 2\}$. This yields 3300 candidate constructions, each evaluated end to end: ensemble, closure, entropy, induced scale.

The results are as follows. The induced G_* scatters across the grammar with a standard deviation of 137 natural-log units. Within ten percent of the measured value lie four constructions, and all four are the realized one, differing only in the closure coefficient, which shifts $g_{\text{share,eff}}$ by less than 0.1% because the closure weighting is a small correction to $\ln \Omega_{\text{tet}}$; the closure coefficient therefore carries almost no load. The nearest structurally distinct construction misses by eleven percent, and within a factor of two of the measured value lie only two construction families in the entire grammar. One further menu member belongs in the ten-percent set: replacing the export prefactor with the quantum-speed-limit coefficient $\pi/2$, the value a saturated orthogonality bound would select, lands at +8.5%; the two-percent landing remains unique to the transverse-export reading. At the same time, the menu density implies that roughly one family is expected to land near any target by chance alone, so the numerical landing by itself carries little evidential weight. Any residual weight would reside in the a priori force of the discrete arguments that select the realized construction — the seven-state face sector (Section 5), the parity doubling of the orientation classes, the ordered injective assignment over distinguishable faces, the seven-fold support exponent (Appendix D.4), and the transverse export factor (Appendix C.5). Given the provenance above, these arguments are historically retrodictive, and the framework's evidential case is not rested on them here; it is rested on the post-fixation tests listed in the provenance statement.

The same audit redistributes the evidential weight of the galactic scale. Because $a_0 = (g_{\text{share,eff}}/4\pi^2)\, cH_0$ depends on the ensemble only logarithmically, sixty-three percent of the entire grammar lands a_0 within twenty-five percent of the observed value: the galactic acceleration scale tests the coupling form, not the realized ensemble. The induced G_*, with its exponential sensitivity, is the sharp test of the ensemble itself, with the charged-lepton ladder its weaker corroborator. The audit script is maintained alongside the reproduction script of Appendix J.3 at `jaigp.org`.

Appendix L: 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 $T \leq 2\times10^{-3}$, 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.

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] E. Verlinde, "On the origin of gravity and the laws of Newton," *Journal of High Energy Physics* **2011**(4), 29 (2011).

[22] E. Verlinde, "Emergent gravity and the dark universe," *SciPost Physics* **2**, 016 (2017).

[23] C. Skordis and T. Złośnik, "New relativistic theory for modified Newtonian dynamics," *Physical Review Letters* **127**, 161302 (2021).

[24] 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).

[25] L. Bombelli, J. Lee, D. Meyer, and R. D. Sorkin, "Space-time as a causal set," *Physical Review Letters* **59**, 521 (1987).

[26] F. Dowker, J. Henson, and R. D. Sorkin, "Quantum gravity phenomenology, Lorentz invariance and discreteness," *Modern Physics Letters A* **19**, 1829 (2004).

[27] 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).

[28] 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).

[29] R. Genzel et al., “Strongly baryon-dominated disk galaxies at the peak of galaxy formation ten billion years ago,” *Nature* **543**, 397 (2017).

[30] H. Prince and J. Dunkley, “Data compression in cosmology: A compressed likelihood for Planck data,” *Physical Review D* **100**, 083502 (2019).

[31] B. I. Ciocan, N. F. Bouché, J. Fensch, D. Krajnović, J. Freundlich, H. Desmond, B. Famaey, and R. Techi, “MUSE-DARK III: The evolution of the radial acceleration relation at intermediate redshifts,” *Astronomy & Astrophysics* **709**, L16 (2026).

[32] A. A. Vărăşteanu, 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).

[33] P. Lang et al., “Falling rotation curves of star-forming galaxies at $0.6 < z < 2.6$,” *Astrophysical Journal* **840**, 92 (2017).

[34] 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).

[35] M. M. Brouwer et al., “The weak lensing radial acceleration relation: measuring the dark matter law with KiDS-1000,” *Astronomy & Astrophysics* **650**, A113 (2021).

[36] W. G. Unruh, “Notes on black-hole evaporation,” *Physical Review D* **14**, 870 (1976).

[37] R. B. Tully and J. R. Fisher, “A new method of determining distances to galaxies,” *Astronomy & Astrophysics* **54**, 661 (1977).

[38] 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).

[39] Planck Collaboration, “Planck 2018 results. VI. Cosmological parameters,” *Astronomy & Astrophysics* **641**, A6 (2020).

[40] A. G. Riess et al., “A comprehensive measurement of the local value of the Hubble constant with 1 km s^{-1} Mpc^{-1} uncertainty from the Hubble Space Telescope and the SH0ES team,” *Astrophysical Journal Letters* **934**, L7 (2022).

[41] 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.

[42] E. Di Valentino et al., “In the realm of the Hubble tension — a review of solutions,” *Classical and Quantum Gravity* **38**, 153001 (2021).

[43] C. M. Will, “The confrontation between general relativity and experiment,” *Living Reviews in Relativity* **17**, 4 (2014).

[44] C. Brans and R. H. Dicke, "Mach's principle and a relativistic theory of gravitation," *Physical Review* **124**, 925 (1961).

[45] 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).

[46] T. Padmanabhan, "Thermodynamical aspects of gravity: new insights," *Reports on Progress in Physics* **73**, 046901 (2010).

[47] 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).

[48] D. Oriti, "Group field theory as the second quantization of loop quantum gravity," *Classical and Quantum Gravity* **33**, 085005 (2016).

[49] S. Gielen, D. Oriti, and L. Sindoni, "Cosmology from group field theory formalism for quantum gravity," *Physical Review Letters* **111**, 031301 (2013).

[50] A. Barbieri, "Quantum tetrahedra and simplicial spin networks," *Nuclear Physics B* **518**, 714 (1998).

[51] J. C. Baez and J. W. Barrett, "The quantum tetrahedron in 3 and 4 dimensions," *Advances in Theoretical and Mathematical Physics* **3**, 815 (1999).

[52] C. Rovelli and L. Smolin, "Discreteness of area and volume in quantum gravity," *Nuclear Physics B* **442**, 593 (1995).

[53] A. M. Gleason, "Measures on the closed subspaces of a Hilbert space," *Journal of Mathematics and Mechanics* **6**, 885 (1957).

[54] D. Deutsch, "Quantum theory of probability and decisions," *Proceedings of the Royal Society of London A* **455**, 3129 (1999).

[55] W. H. Zurek, "Probabilities from entanglement, Born's rule $p_k = |\psi_k|^2$ from envariance," *Physical Review A* **71**, 052105 (2005).

[56] S. Popescu and D. Rohrlich, "Quantum nonlocality as an axiom," *Foundations of Physics* **24**, 379 (1994).

[57] D. Eckert et al., "Non-thermal pressure support in X-COP galaxy clusters," *Astronomy & Astrophysics* **621**, A40 (2019).

[58] 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).

[59] XRISM Collaboration, "XRISM reveals low nonthermal pressure in the core of the hot, relaxed galaxy cluster Abell 2029," *Astrophysical Journal Letters* **982**, L5 (2025).

[60] A. H. Chamseddine and V. Mukhanov, "Mimetic dark matter," *Journal of High Energy Physics* **1311**, 135 (2013).

[61] C. Skordis and T. Złośnik, "New relativistic theory for modified Newtonian dynamics," *Physical Review Letters* **127**, 161302 (2021).

[62] 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).

[63] 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).