Symbiotic Infodynamic Equilibrium: An Ontology-First Consolidation of the Framework, its Derivations, and its Empirical Status

Cody Hudock

[2pt] *Information Physics Institute, Gosport, Hampshire, UK*
codyhudock@gmail.com

Version 4 — April 27, 2026

Abstract

We present a consolidated account of the Symbiotic Infodynamic Equilibrium (SIE) framework. The program is built on three axioms: (A1) information is primary rather than derivative; (A2) the universe is a closed thermodynamic system; (A3) all dynamics respect Landauer bounds at local temperature. The framework's core technical construct is the *Landauer Mass Tensor*, $\mathcal{M}_L = N_{\rm bits} k_B T_{\rm local} \ln 2/c^2$, which replaces static rest mass with a dynamic thermodynamic cost of refreshing localized information on an FCC substrate. Supporting constructs include the Infodynamic Shear Tensor for non-adiabatic substrate coupling, the Gompertz form for decay-driven cosmic expansion, and a two-phase principle distinguishing kinetic from inertial informational modes. A unified five-term Lagrangian follows from the principle of infodynamic stress minimization, encompassing gravity as entropic compression, matter as execution threads, electromagnetism under bit-density saturation, topological texture, and dark energy as osmotic allocation. We report the current state of the derivation chain — five problems (the chirally-asymmetric Wilson coefficient, strong-CP suppression, gauge-boson halving, the FCC → SU(4) → Standard Model path, and Yukawa compensation) have substantive structural mechanisms identified, with four remaining technical gaps honestly scoped. Observational content, from existing corpus papers, includes: (i) SIE rotation curves fit the 175-galaxy SPARC catalog with one free parameter per galaxy ($\Upsilon_{\rm disk}$), achieving competitive goodness-of-fit against cosmologically-constrained NFW (2 parameters per galaxy) with median $\chi^2_\nu = 3.33$ vs NFW's 1.72, BIC-preferred in 11% of galaxies strongly and 37% in total; the primary significance is that SIE's one-parameter fit remains competitive while making a zero-free-parameter prediction for the coupling $a_0 = cH_0/2\pi$; (ii) a kinematic Hubble constant $H_0^{\rm kin} = 73.97$ $\mathrm{km\,s^{-1}\,Mpc^{-1}}$ from galactic data alone; (iii) the Planck spectral tilt $n_s = 1 - 2/57 = 0.96491$ matching at 0.003σ via the left-handed chiral Standard Model inventory; (iv) a bare vacuum density $\rho_{\rm bare} \approx 2.07 \times 10^{-27}$ $\mathrm{kg\,m^{-3}}$ within a factor of four of the observed critical density from CKN plus Landauer with no fitted parameters; (v) a 7.85° CMB quadrupole-octupole bisector from FCC {111} cleavage geometry with correlation significance $p < 0.001$; and (vi) a framework-consistent interpretation of the Hubble tension as an observer-interior-void effect via $H_{\rm local} = H_{\rm global}/\sqrt{1-\delta_b}$ producing the SH0ES match for interior observers, with the caveat that a fresh Pantheon+ re-analysis rules out the predicted void gradient under the standard SH0ES-anchored analysis; the framework's defense requires a pipeline-level calculation of how interior anchoring propagates through SH0ES, which is currently outstanding. We identify specific research tasks — a lattice-gauge-theory calculation of the FCC Wilson coefficient, explicit Pati–Salam Higgs potential construction, derivation of the icosahedral shell selection rule, and deepening of individual Yukawa predictions — as the concrete program for closing the framework. SIE is falsifiable on multiple fronts, including next-generation ringdown measurements, DESI high-redshift equations of state, IAXO axion searches, and LiteBIRD polarization maps.

1. The Ontological Flip

Modern physics carries a productive embarrassment: two of its most carefully measured quantities — the vacuum energy density and the local Hubble expansion rate — fail to match expectations from the same framework that generated them. The vacuum catastrophe separates the quantum-field-theoretic calculation from the observed cosmological density by 120 orders of magnitude [1]. The Hubble tension separates CMB-anchored and distance-ladder-anchored measurements of H_0 at a level that has only hardened as both pipelines matured [2, 3]. These are typically treated as unrelated puzzles.

We propose they are symptoms of the same misconception: that spacetime is a continuous manifold with infinite information-processing capacity. If space is instead a discrete finite-capacity relational graph, both problems become computable thermodynamic limits rather than interpretational puzzles. This is the starting point of Symbiotic Infodynamic Equilibrium (SIE).

The programmatic inversion is this: in standard accounts, information is a description of physical state. Energy, mass, charge are primary; entropy counts the configurations those primary quantities can occupy. SIE inverts the ordering. Information — more precisely, the substrate's capacity to host and refresh discrete informational states — is primary. Mass, energy, and geometric curvature are emergent thermodynamic phenomena in a substrate that is doing finite-bandwidth computation.

This is not mysticism. It is the operational claim that, once one insists on finite information capacity per Planck volume and on Landauer erasure costs at local temperature, the numerical discrepancies that break standard cosmology become calculations with one or zero free parameters.

> **The framing question**: Is the observed universe the dynamics of "stuff" on a passive spacetime stage, or is it the thermodynamic overhead of a substrate performing bounded computation? SIE commits to the second answer and asks which predictions follow.

The rest of this paper develops that commitment. Section II states the three axioms. Section III presents the single technical novelty (the Landauer Mass Tensor) and the five-term Lagrangian that follows. Sections IV–V present the derivation chain that attempts to reconstruct the Standard Model's gauge and fermion content from the substrate geometry, identifying structural mechanisms for five problems and honestly flagging the four remaining gaps. Sections VI–VIII report the empirical successes from the existing corpus with specific numerical values and appropriate uncertainty. Section IX acknowledges the four remaining technical gaps. Section X summarizes the falsifiability profile.

2. Three Axioms

The framework rests on three axioms, each of which can be stated without reference to the others and each of which carries empirical weight.

2.1. A1 — Information is primary

The substrate of reality is a discrete network of informational states. Space is not a continuous manifold on which information is inscribed; space *is* the relational structure of the informational network. The Planck length ℓ_P is the node spacing of this network; the speed of light c is the maximum update rate.

To unpack the ontological weight: in conventional formulations, information is a description that someone could in principle extract from a physical system — a shadow the system casts when interrogated. Bekenstein and Hawking showed that even in general relativity, information cannot be quite so passive: the entropy of a black hole horizon constrains how much information could possibly be encoded within [6, 56]. Wheeler's "it from bit" [57] pushed the point further: the structural features of reality correspond to informational relationships rather than primary substance. SIE takes the final step and commits to the inverted ontology as operational, not just philosophical.

The concrete consequences: (i) zero-point energy diverges for continuous spacetime; a discrete substrate with finite per-node capacity gives a natural Planck-scale cutoff [4]. (ii) The Bekenstein bound [6], Bousso covariant entropy bound [7], and holographic principle [8] are direct consequences rather than additional postulates. (iii) Quantum wave function collapse becomes a calculable bandwidth exhaustion event rather than an interpretational puzzle (Section III.D). (iv) The cosmological horizon is a physical boundary rather than a coordinate artifact.

If A1 were false — if space were continuous down to arbitrary scales with no informational structure — the vacuum catastrophe would be real and the numerical convergences catalogued in this paper would be cosmic accidents.

2.2. A2 — The universe is closed

Energy, momentum, and information are globally conserved within the Hubble horizon. The cosmological event horizon is an adiabatic wall: nothing crosses it. This is distinct from open-system cosmologies in which energy flows into or out of an external reservoir.

The operational consequence is that all observed dynamics must balance. When space expands, the informational cost of that expansion must be paid by something inside the horizon. When matter condenses, the entropic debt must be allocated somewhere. This axiom is what makes dark energy a calculable rather than postulated quantity: it is the bookkeeping entry required to balance the substrate's ledger against the decay and formation of localized mass states [5].

The closed-system assumption has teeth. Most cosmological calculations implicitly treat the universe as if it has an external energy bath, or as if new information can be produced without thermodynamic cost. SIE prohibits both. The cosmological constant problem is not "why is Λ so small?" but rather "what calculation determines how much buffer space the substrate must allocate?" Giving that calculation a clear answer ($\rho_{\text{bare}} \approx 2.07\times10^{-27}$ $\text{kg}\,\text{m}^{-3}$ from CKN+Landauer) is the consequence of A2.

A concrete test of A2: the matter growth index γ is often treated as a phenomenological parameter ≈ 0.55 [33]. Under A2, the closed-system boundary condition $w = -1$ is a natural feature at the adiabatic horizon. If the standard Wang-Steinhardt parametrization $\gamma = 3(w-1)/(6w-5)$ carries over into the SIE framework (which assumes GR structure at the background level, a point we return to in Sec. VIII), it yields $\gamma = 0.545$ at $w = -1$, matching the phenomenological fit with no adjustable parameter [60]. Whether the Wang-Steinhardt formula is preserved exactly in the osmotic picture — or only approximately, given that the growth dynamics depend on the specific perturbation equations — is a question we note but do not yet resolve.

2.3. A3 — Landauer bounds are respected locally

Every informational operation on the substrate — every refresh of a localized bit against ambient thermal noise — costs at least $k_B T_{\text{local}} \ln 2$ in energy, where T_{local} is the temperature of the thermal reservoir in contact with the operation. No renormalization or coarse-graining softens this bound; it is strict in every local rest frame.

The temperature T_{local} is not a universal constant. For bulk vacuum modes at cosmological scales, it is the CMB temperature $T_{\text{CMB}} = 2.725$ K: the photon gas is not a separate fluid flowing through substrate but rather the dominant thermal excitation of the substrate's own gauge modes [58]. For high-acceleration boundaries (near black hole horizons, near the cosmological horizon), it is the Unruh or Gibbons-Hawking temperature [59]. For laboratory systems in thermal contact with a cryostat, it is the environmental temperature T_{env}.

This axiom is what converts the first two into quantitative content. A1 alone states that information is primary; A3 gives information a price. A2 states that the price is paid within a closed ledger; A3 tells the ledger how much to charge per line.

The empirical handles of A3 are diverse. Dark energy density emerges from $\rho_{\text{bare}} = (\text{bit density})\times k_B T_{\text{CMB}} \ln 2/c^2$. Quantum collapse energy is predicted to be $Q \approx 5$ μJ for a 100-nm sphere in a 1 mK environment, detectable by TES calorimetry [47]. Black hole evaporation rates and ringdown frequencies scale with substrate temperature at the horizon. In every case, the specific prediction depends on identifying the correct local reservoir temperature — not a free parameter but a calculable quantity for each setting.

> The three axioms together form a minimal foundation. A1 asserts the substrate; A2 closes the system; A3 sets the local exchange rate between information and energy. From these alone, with no further assumptions, we claim the observed structure of physics emerges.

3. The Technical Novelty: The Landauer Mass Tensor

SIE's central technical move — the piece that converts philosophical axioms into calculation — is the replacement of static rest mass with a dynamic thermodynamic tensor.

3.1. The replacement

In standard relativistic quantum mechanics, the Dirac Lagrangian contains a rest-mass term $m\bar{\psi}\psi$ where m is a constant. In SIE this becomes:

$$\mathcal{M}_L = \frac{N_{\text{bits}}\, k_B\, T_{\text{local}}\, \ln 2}{c^2} \tag{1}$$

where N_{bits} is the number of bits required to encode the particle's internal state, T_{local} is the temperature of the substrate reservoir in contact with the particle, and the $\ln 2$ factor is Landauer's constant per bit operation.

This is not an ornament. It is a definitional equation. What looks like mass in the Dirac operator is the continuous thermodynamic cost of refreshing the particle's localized informational state against substrate noise. The more bits required to encode the state (N_{bits}), or the hotter the local substrate (T_{local}), the higher the refresh cost and therefore the larger the apparent inertia.

3.2. Why this is not trivial

A simple dimensional analysis shows that $\mathcal{M}_L$ has units of mass. But dimensional match does not make a definition predictive. The claim gains content when $N_{\rm bits}$ and $T_{\rm local}$ are themselves specified by substrate geometry.

For a proton, $N_{\rm bits}$ is set by the number of distinct configurations of quark and gluon color charge accessible during the QCD refresh timescale $\tau_{\rm QCD} \sim 10^{-24}$ s. Standard QCD says the valence quarks are in constant exchange of gluons, flipping color at this cadence. Each flip is a Landauer erasure event: the previous color state must be logically overwritten to write the new one. Counting these erasures per unit time and multiplying by $k_B T_{\rm QCD} \ln 2$ gives the proton's refresh power. Integrating over the QCD light-cone gives the proton mass within $\sim 3\%$ of the observed value [5, 9] — a calculation that uses no free parameters once the FCC substrate temperature is specified.

This is not a universal derivation of all fermion masses; it is a worked example showing that the tensor $\mathcal{M}_L$ contains calculable content rather than an unspecified phenomenological dial. Extending the same logic to other particles is the research program discussed in Section IX.

3.3. Worked example: the proton mass

To make the Landauer Mass Tensor concrete rather than abstract, consider the proton. Quantum chromodynamics establishes that the proton is not static: its three valence quarks exchange gluons at the strong-interaction timescale $\tau_{\rm QCD} \sim 10^{-24}$ s, continuously rotating their color assignments among the three possibilities (red, green, blue) [61]. The standard interpretation: these are virtual exchanges mediating confinement.

The SIE reinterpretation: each color rotation is a Landauer erasure event on the substrate. The prior color state must be logically overwritten before the new one can be written, and Landauer's theorem says that logical irreversibility costs $k_B T \ln 2$ per bit at temperature T.

Counting the bits. Each quark's internal state carries two logical degrees of freedom:

- Spin (up or down): 1 bit.
- Color charge (red, green, or blue): $\log_2 3 \approx 1.585$ bits.

A single quark therefore encodes ≈ 2.585 bits of logic. A proton with three valence quarks has total logical content

$$N_{\rm bits}^{(\rm p)} = 3 \times 2.585 \approx 7.755 \text{ bits}. \tag{2}$$

Setting the temperature. In QCD, the scale at which hadronic matter transitions into a quark-gluon plasma is the Hagedorn temperature $T_{\rm H} \approx 180$ MeV [94]. This is the thermal threshold above which the confined state cannot persist; below it, the substrate must maintain the bit-refresh cycle against this heat bath.

The Landauer bill. Applying $\mathcal{M}_L c^2 = N_{\rm bits} \cdot k_B T \ln 2$ with these inputs:

$$m_p c^2 \approx 7.755 \times 180 \text{ MeV} \times \ln 2 \approx 968 \text{ MeV}. \tag{3}$$

Compared with the observed $m_p c^2 = 938.272$ MeV [61], the prediction agrees to within $\sim 3\%$. The residual is naturally attributed to the environmental corrections on top of the idealized single-proton calculation (sea-quark contributions, isospin splitting, electromagnetic self-energy), each of which is known to contribute at the percent level [61].

What this is and isn't. This is a one-line calculation with three inputs: the quark bit-count formula (fixed by the Standard Model's $SU(3) \times SU(2)$ structure), the Hagedorn temperature (fixed by the QCD phase diagram), and Landauer's theorem (a theorem, not a parameter). No input is tuned to m_p. The agreement at $\sim 3\%$ is striking. A factor-of-two error would have been unremarkable; a factor-of-ten error would have falsified the ansatz.

We do *not* claim this replaces lattice QCD as a first-principles calculation of the proton mass. Lattice QCD reproduces m_p through direct numerical integration of the QCD action and is the authoritative derivation. What Eq. 3 shows is that the Landauer Mass Tensor, applied with established inputs, delivers a consistent numerical content rather than being a vacuous rewriting. The deeper relationship between this arithmetic and the full lattice result is an open question; we record it honestly as an intriguing consistency.

3.4. The selection of $T_{\rm local}$

The LMT contains one quantity whose specification could admit abuse: $T_{\rm local}$, the temperature at which the Landauer refresh cost is evaluated. Throughout this paper, different particles invoke different temperatures — the Hagedorn temperature for the proton, $T_{\rm CMB}$ for the neutrino, the dilution-fridge temperature for a transmon, the Hawking temperature for black hole emission. A fair question: what principle selects T_{local} unambiguously, and does it prevent post-hoc selection of whichever reservoir gives the right answer?

The operational rule we can offer. T_{local} is the temperature of the *dominant thermal bath that the particle's bit-refresh operations couple to.* "Dominant" means: among all thermal reservoirs to which the particle's internal degrees of freedom are coupled, the highest temperature controls the Landauer cost, because that is the reservoir that will re-thermalize the bits fastest.

Applied concretely:

- **Proton**. The proton's color and spin bits couple to the QCD gluonic plasma at the Hagedorn scale $T_{\text{H}} \approx 180$ MeV$/k_B$; to the ambient CMB photon bath at $T_{\text{CMB}} = 2.725$ K; and to the gravitational Unruh bath on Earth's surface at $T_U \sim 10^{-19}$ K. The Hagedorn temperature dominates by more than ten orders of magnitude; it alone controls the Landauer cost of internal color-bit rotation. The other reservoirs contribute negligibly.

- **Neutrino**. The neutrino's one chiral bit couples essentially only to the ambient vacuum through its weak charge (the weak interaction itself is the coupling). The dominant thermal bath is therefore the CMB photon gas thermalized into the substrate at T_{CMB}. No QCD-scale bath is available because the neutrino has no color charge.

- **Transmon**. The Josephson junction's qubit bits couple to the substrate phonons at the cryostat temperature ($\sim$ 20 mK in dilution fridge operation), augmented by the junction's own operating frequency setting an effective excitation temperature. The dominant coupled bath is the cryostat.

When the rule has teeth. The "dominant bath" criterion is a genuine constraint. For the proton, you cannot substitute T_{CMB} for T_{H}: the proton's color bits do not thermalize against the CMB on any meaningful timescale, whereas they thermalize against the gluonic plasma at $\tau_{\text{QCD}} \sim 10^{-24}$ s. For the neutrino, you cannot substitute T_{H}: the neutrino has no QCD coupling, period. These are not post-hoc choices; they are consequences of the particle's gauge structure.

Where the rule loosens. For composite systems with multiple coupled baths of comparable temperature — a neutron star's crust (thermal phonons + nuclear Landau levels + magnetic fields), or the electroweak epoch of the early universe (simultaneous Higgs, gauge, and fermion bath equilibration) — "dominant" becomes less obvious. These are genuine edge cases where the selection rule underdetermines T_{local} and separate analysis is required. We flag this as a methodological limitation rather than a resolved selection rule, and note that such cases are not the ones from which the quantitative claims in this paper are drawn.

A cleaner formulation — probably in terms of the decoherence-time hierarchy of the substrate density matrix relative to the particle's internal clock — is an open methodological task. Until that is available, the LMT formula retains some interpretive latitude for edge cases. The predictions in the main text use the dominant-bath selection unambiguously and do not depend on resolving the edge cases.

3.5. The electron as a Bekenstein boundary

Standard QED treats the electron as a zero-dimensional point, with the self-energy pathology and renormalization dance discussed in Section V. In SIE the electron has a finite radius that follows once α^{-1} is adopted as the substrate's Shannon capacity for charge encoding.

The logical order matters. The fine-structure constant is taken as an empirical input to the framework (as it is in the Standard Model), with the substrate-theoretic interpretation given in Section XI.A: α^{-1} represents the number of internal processing cycles per emission event on the substrate. The framework does not currently derive the specific value 137.036; that remains an open problem. We take $\alpha^{-1} \approx 137.036$ as a substrate property established empirically.

Given α^{-1}, the electron's geometric size follows from the Bekenstein bound [6] under a specific structural assumption: that the electron *saturates* the bound, i.e., its boundary surface is maximally packed with the information needed to encode its charge. This is not automatic — Bekenstein gives $S \leq A/(4\ell_P^2 \ln 2)$, an inequality — so we state the saturation assumption explicitly.

The physical argument for saturation is minimum-stable-encoding: a less-than-saturated boundary would have unused holographic capacity, representing a localized entropy deficit that the substrate could relax by shrinking the boundary. An over-saturated boundary is forbidden by the bound itself. The saturated boundary is the unique minimum-energy configuration that stably encodes the charge without holographic waste.

Under this assumption, the information capacity of a spherical surface of radius r equals:

$$S_{\text{Bek}} = \frac{4\pi r^2}{4\ell_P^2 \ln 2} \text{ bits.} \tag{4}$$

Setting $S_{\text{Bek}} = \alpha^{-1}$ (the α^{-1} bits required to encode one unit of charge against vacuum polarization) and solving:

$$r_e = \ell_P \sqrt{\frac{\alpha^{-1} \ln 2}{\pi}} \approx 5.5\ \ell_P. \tag{5}$$

The electron is a finite spherical topological excitation of the substrate roughly 5.5 Planck lengths in radius, not a mathematical point. This provides the geometric mechanism that eliminates the self-energy divergence (Section V.A): $r_e > 0$, so $U_{\text{self}} \propto e^2/r_e$ is finite.

On the ordering. We state explicitly: α^{-1} is taken as an empirical input (with substrate-theoretic interpretation in Section XI.A but no first-principles derivation of the value); r_e is a *consequence* of α^{-1} via the saturated Bekenstein bound. An earlier presentation of this material in the corpus sometimes suggested α^{-1} could be derived from the electron surface, which (combined with the $r_e \leftarrow \alpha^{-1}$ direction here) would form a circle. The framework's commitment is to α^{-1} as empirical input and r_e as derived consequence.

3.6. QCD confinement as deep encryption

The same Landauer-thermodynamic framework illuminates why quarks cannot be isolated, which standard QCD captures phenomenologically via the Cornell potential $V(r) = -4\alpha_s/3r + \sigma r$ whose linear-in-r confinement term grows without bound at large separation [84].

In SIE, color charge is a discrete bit assignment on the substrate. Separating two quarks would require the substrate to maintain two distinct color endpoints over increasing distance, which costs linearly growing Landauer energy because each additional substrate node between them must refresh against thermal noise. The Cornell string tension $\sigma \approx 1$ GeV/fm is the rate at which this Landauer cost accumulates per unit substrate distance.

The deeper consequence: you cannot isolate "half a quark" because the substrate cannot store half a bit. Color is discrete, quantized at the level of substrate nodes. Confinement is not a dynamical mystery to be explained by some intricate non-perturbative QCD mechanism; it is a digital logic constraint on a finite-capacity information-theoretic substrate. The linear Cornell potential is the thermodynamic bill for attempting an operation the substrate cannot perform.

Roughly 99% of the proton's mass lives in this gluon encryption wrapper rather than in the bare quark masses [61]. In SIE this is not surprising: the Landauer refresh cost of maintaining the color-correlated state across the proton's interior dominates over the quarks' individual Compton-scale refresh costs. The proton mass calculation of Section III.B is, in effect, the calculation of this encryption overhead.

3.7. Mass ratios and the quantum measurement problem

The Landauer Mass Tensor extends beyond individual masses to ratios. The mass ratio $m_p/m_e \approx 1836$ in SIE emerges from a volumetric bit-count argument: proton and electron refresh different volumes of substrate at different rates, and the ratio of their Landauer mass tensors equals the ratio of their $N_{\text{bits}} \times T$ products. This reproduces the observed ratio to 0.4% accuracy [9].

Extending the same logic to cosmological scales, the Landauer maintenance cost of idle (matter-free) substrate nodes at T_{CMB} gives a bulk energy density contribution. Multiplied by the CKN-bounded bit density (Sec. VIII.B), this produces the dark energy scale within a factor of four [4]. Different manifestations of the same thermodynamic quantity — refresh cost against local thermal noise — show up at each scale.

The measurement problem connects through $\mathcal{M}_L$ as well. A quantum superposition of a mass M across spatial separation Δx carries a linear infodynamic burden $S_{\text{super}} \propto M\Delta x$, set by the Ryu-Takayanagi prescription applied to the FCC tensor network [47]. The intrinsic substrate capacity for holding coherent superpositions scales as $S_{\text{cap}} \propto M^{-2}$ from the Bekenstein entropy of the reduced Compton horizon. These scalings intersect at

$$\Delta x_{\max} = \pi \ell_P (m_P/M)^3 \tag{6}$$

with no free parameters. This predicts the location of objective wave function collapse for any mass: at $M = m_P/\sqrt{2}$ (the Planck mass threshold), $\Delta x_{\max}$ equals the Schwarzschild radius, confirming the scale geometrically. For a 100-nm silica sphere at 1 mK, collapse occurs at $\Delta x_{\max} \sim 10^{-12}$ m, a distance accessible to MAQRO-class levitated nanosphere experiments [62].

Predictions from the Landauer Mass Tensor:

- Proton mass from QCD color refresh: $\sim 3\%$ accuracy.
- m_p/m_e from volumetric bit counting: 0.4% accuracy.
- Lepton mass ratios from icosahedral shells ($K = 0, 2, 18$): $\sim 1\%$ accuracy.
- Electron radius $r_e \approx 5.5\,\ell_P$ from Bekenstein bound + α^{-1}.
- Dark energy density from Landauer+CKN: within factor 4 of ρ_c.
- Quantum collapse threshold: $\Delta x_{\max} = \pi \ell_P (m_P/M)^3$, material-invariant, testable at MAQRO.
- Landauer thermal signature of collapse: $Q \sim 5\,\mu$J for 100-nm spheres, 13 orders above TES sensitivity.

All from the single equation $\mathcal{M}_L = N_{\text{bits}} k_B T_{\text{local}} \ln 2/c^2$ with substrate-specific N_{bits} and T_{local}.

3.8. The consequential reinterpretations

With $\mathcal{M}_L$ in place, several familiar quantities of physics receive reinterpretation. Einstein's $E = mc^2$ becomes a statement about the conversion rate between raw information and thermodynamic debt: a particle's rest energy is the integrated refresh cost of maintaining its informational form against the substrate's preference for relaxation. Gravity, on scales where the Bekenstein bound matters, becomes the substrate's entropic snapback — its tendency to compress localized high-information regions to minimize their holographic footprint [11, 12]. Dark energy becomes the Landauer maintenance cost of idle substrate nodes [4].

None of these reinterpretations discard standard physics; they retrieve it as effective descriptions. Einstein's equations emerge as the statistical average of substrate entropic compression [11]. Maxwell's equations emerge as the low-energy expansion of a Born–Infeld-saturated substrate [13]. The Standard Model gauge structure emerges from the FCC geometric substrate via root system A_3, as shown in Sections IV–V.

4. The Master Equation

From axioms A1–A3 plus the Landauer Mass Tensor, the system's dynamics follow from a single variational principle: minimize substrate stress.

4.1. The infodynamic stress functional

Define the substrate stress $\mathcal{S}$ as the sum of three competing informational pressures: compression (gravitational, from localized mass states), shear (electromagnetic, from active gauge interactions), and expansion (osmotic, from the need to allocate buffer space). The universe's dynamics are the path that minimizes the integrated stress:

$$\delta \mathcal{S} = \delta \int \left(I_{\text{comp}} + I_{\text{shear}} - I_{\text{exp}} \right) dt = 0. \tag{7}$$

This is the SIE analog of the principle of least action, with action replaced by informational stress. It is a restatement of the second law of infodynamics: the substrate evolves to minimize the rate at which it must pay Landauer costs.

4.2. The five-term Lagrangian

The variational principle, combined with the substrate's discrete FCC geometry and the Landauer Mass Tensor, produces a unified Lagrangian:

$$\begin{aligned} \mathcal{L}_{\text{SIE}} = \; & \underbrace{\frac{R}{16\pi G}}_{\text{Bulk Stress}} + \underbrace{\bar{\psi}\left(i\gamma^\mu D_\mu - \mathcal{M}_L\right)\psi}_{\text{Execution Threads}} \\ & + \underbrace{\Pi_{\max}^2 \left(1 - \sqrt{\mathcal{D}(F)}\right)}_{\text{Saturated Transmission}} \\ & + \underbrace{\frac{\theta g^2}{32\pi^2} G^a_{\mu\nu} \tilde{G}^{a\mu\nu}}_{\text{Topological Texture}} \\ & - \underbrace{\left(\tfrac{1}{2}\partial_\mu \phi \partial^\mu \phi + V(\phi, T)\right)}_{\text{Osmotic Allocation}} \end{aligned} \tag{8}$$

where the Born–Infeld discriminant is

$$\mathcal{D}(F) = 1 + \frac{F_{\mu\nu}F^{\mu\nu}}{2\Pi_{\max}^2} - \frac{(F_{\mu\nu}\tilde{F}^{\mu\nu})^2}{16\Pi_{\max}^4}. \tag{9}$$

Each term has a substrate-theoretic interpretation:
Bulk Stress. The Einstein–Hilbert term emerges as the coarse-grained statistical average of substrate entropic compression. It is not fundamental geometry; it is the long-wavelength limit of substrate bookkeeping. On scales where the discrete structure is invisible, this term reduces exactly to general relativity [11].
Execution Threads. The Dirac term for matter, with $\mathcal{M}_L$ replacing static mass. Fermions are localized refresh patterns on the substrate; their inertia is the thermodynamic cost of maintaining the pattern against substrate noise.
Saturated Transmission. Electromagnetism in Born–Infeld form. The non-linear square root enforces a finite maximum bit-density $\Pi_{\max}$ set by

the substrate's information capacity. At laboratory energies, this reduces to standard Maxwell theory; at Planck-scale fields, the square root saturates, which mathematically closes the ultraviolet divergences of standard QED.

Topological Texture. The theta term encodes the substrate's chiral grain — the topological preference for one handedness over the other. The value of θ is determined by substrate structure, not postulated, and its observational consequences connect to the strong-CP problem (Section IV.B) and the matter-antimatter asymmetry.

Osmotic Allocation. Dark energy as a dynamical scalar field ϕ with thermal-dependent potential $V(\phi, T)$. This term encodes the requirement that the substrate allocate buffer space to balance its ledger; it produces quintessence with $w > -1$ evolving toward $w \to -1$ asymptotically [4].

4.3. Why five terms, not four or six

The term count is not arbitrary. The three competing informational pressures (compression, shear, expansion) plus the two substrate-topological requirements (execution threads and topological texture) give exactly five independent sources of substrate stress. Adding further terms would reduce to combinations of these; removing any one would leave a consequence of the axioms unrepresented.

5. Point Particle Infinities and the Landau Pole

One of the most consequential applications of the Master Equation is the resolution of two longstanding infinities in quantum field theory: the electron self-energy divergence and the Landau pole in the running electromagnetic coupling. Both arise from the same structural assumption — that the vacuum is a continuous manifold with infinitely deep virtual structure. Once that assumption is replaced by a finite-bandwidth substrate, both divergences acquire a geometric cutoff rather than requiring cancellation via renormalization. This is not a criticism of renormalization, which is mathematically well-defined and empirically spectacular; it is an ontological reframing in which the quantities that renormalization sends to finite values through formal cancellation instead have finite-valued physical origin.

This section is devoted to making that reframing precise.

5.1. The classical self-energy problem

Standard electromagnetism gives the electrostatic energy stored in a charge's own field as

$$U_{\text{self}} = \frac{1}{8\pi\epsilon_0} \int |\mathbf{E}|^2 \, dV \propto \frac{e^2}{r_e}, \tag{10}$$

where r_e is the characteristic charge radius. For a classical point charge ($r_e \to 0$), this integral diverges: the self-energy is infinite, which via $E = mc^2$ would give the electron an infinite mass. QED's response was *renormalization*: the bare mass in the Lagrangian is treated as a formal parameter that absorbs the divergent self-energy contribution, leaving the observed finite $m_e = 0.511\ \text{MeV}/c^2$ as the physical mass. The procedure is mathematically rigorous, computationally tractable, and empirically spectacular — QED's predictions match experiment to ten or more decimal places. It has been properly reinterpreted in the modern effective-field-theory framework as a standard feature of Wilsonian coarse-graining, not a defect.

In SIE, the electron is not a zero-dimensional point. It is a localized topological excitation of the substrate with a characteristic radius set by the Bekenstein boundary of its encoded information: approximately $5.5\,\ell_P$ (a few Planck lengths) [9, 67]. Because $r_e > 0$ strictly, the self-energy integral never diverges. No infinity ever appears, so no cancellation is required. The renormalization procedure is revealed as an accurate computational shortcut for a finite physical answer — not a necessary feature of the theory.

5.2. The Landau pole

The point-particle assumption produces a second, subtler pathology in the running electromagnetic coupling. Vacuum polarization by virtual electron-positron pairs screens the bare electron charge: probing at shorter distances (higher energies) sees a less-screened charge, so $\alpha(E)$ increases with E. The one-loop QED result is [74]:

$$\alpha_{\text{QED}}(E) = \frac{\alpha(m_e)}{1 - \frac{\alpha(m_e)}{3\pi} \ln(E^2/m_e^2)}. \tag{11}$$

The denominator reaches zero when the logarithm is large enough, giving a mathematical singularity — the *Landau pole* — at $E \sim m_e \exp(3\pi/2\alpha) \sim 10^{286}$ GeV. At that energy the effective charge of the electron diverges, an obvious nonphysical result. QED is therefore not ultraviolet-complete on its own terms; it can only make sense as the low-energy effective description of a more fundamental theory.

In standard physics the resolution is deferred to the electroweak unification or to some further UV completion. In SIE the resolution is *immediate* and arises from the same finite substrate bandwidth that cures the self-energy problem.

5.3. Dielectric saturation: the vacuum runs out of pixels

The Landau pole assumes that the vacuum's polarization capacity is infinite. In SIE it is not. The substrate has a finite maximum bit-density $\Pi_{\max}$ set by the Bekenstein bound at the Planck scale. When an external field asks the substrate to polarize beyond $\Pi_{\max}$, the substrate literally has no more nodes available to flip — the polarization saturates.

The mathematics of finite two-state polarization is universal: whenever a discrete ensemble with fixed capacity polarizes in response to an external field, the response is governed by the hyperbolic tangent. (The same function describes magnetic spin polarization, ferroelectric response, and any other bounded two-state system.) Applying this to the vacuum:

$$\Pi_{\mathrm{SIE}}(E) = \Pi_{\max} \tanh\left(\frac{\Pi_{\mathrm{QED}}(E)}{\Pi_{\max}}\right). \tag{12}$$

Substituting into the running coupling gives the SIE modification:

$$\alpha_{\mathrm{SIE}}(E) = \frac{\alpha_0}{1 - \Pi_{\max} \tanh\left(\frac{\alpha_0}{3\pi\Pi_{\max}} \ln\frac{E^2}{m_e^2}\right)}. \tag{13}$$

The behavior of this equation is exactly what the physics requires:

- **Low-energy limit** ($E \ll E_{\mathrm{Planck}}$). For small arguments, $\tanh(x) \approx x$, and Eq. 13 reduces to Eq. 11. Every QED prediction verified in colliders is preserved. The SIE modification is invisible at any energy currently testable.

- **High-energy limit** ($E \to E_{\mathrm{Planck}}$). The tanh function asymptotes to 1. The denominator locks at $1 - \Pi_{\max}$ (for $\Pi_{\max} < 1$), a finite nonzero constant. The running coupling flattens into a horizontal asymptote. The Landau pole never occurs; it is geometrically quarantined by the substrate's finite pixel count.

On the value of $\Pi_{\max}$. The dimensionless parameter $\Pi_{\max}$ appearing in Eq. 13 is the ratio of the saturated polarization to the maximum QED polarization that would correspond to full Bekenstein occupation at the Planck scale. It is principle-fixed (a property of the substrate) rather than free, but its specific numerical value is not yet derived from the framework's axioms — the required calculation involves the saturation threshold of the FCC lattice at the gauge-coupling-running scale, which depends on the same LGT computation that sets c_{FCC} (Section IX). Until $\Pi_{\max}$ is computed from first principles, the saturation value at which $\alpha(E)$ flattens is a qualitative prediction rather than a quantitative one. The structural prediction — that the denominator locks at a finite value rather than going to zero — is robust; the specific height of the asymptote is pending.

> The Landau pole is not "solved by new physics above the electroweak scale." It is *prevented* at the level of the Lagrangian by the finite-bandwidth structure of the vacuum. The electron's effective charge physically cannot run to infinity, because the vacuum it polarizes is not an infinite reservoir. When the available pixels are exhausted, the screening saturates. The quantitative saturation value awaits derivation of $\Pi_{\max}$ from substrate dynamics.

5.4. Resurrecting Born-Infeld

In 1934 Max Born and Leopold Infeld proposed a nonlinear electrodynamics specifically to cure the point-charge self-energy divergence [13]. Their proposal replaced Maxwell's Lagrangian with

$$\mathcal{L}_{\mathrm{BI}} = \Pi_{\max}^2 \left(1 - \sqrt{1 + \frac{F_{\mu\nu}F^{\mu\nu}}{2\Pi_{\max}^2} - \frac{(F_{\mu\nu}\tilde{F}^{\mu\nu})^2}{16\Pi_{\max}^4}}\right), \tag{14}$$

which is the Saturated Transmission term of the SIE Master Equation (8). The square root caps the field strength at $\Pi_{\max}$: no field can exceed the maximum, no self-energy can diverge, gauge invariance is preserved exactly, and at low fields the Lagrangian reduces to Maxwell's.

Born-Infeld's resolution was mathematically elegant but had a structural limitation: no one could say *why* the field should cap. The cap parameter $\Pi_{\max}$ was a postulate without mechanism, and after its initial excitement the theory was largely set aside in favor of renormalized QED, which worked experimentally without requiring such a postulate.

SIE restores Born-Infeld by identifying its postulated parameter with a substrate property. $\Pi_{\max}$ is the maximum bit-density of the FCC lattice set by the Bekenstein bound at the Planck scale. It is not an adjustable constant; it is the lattice's hard memory limit. What Born and Infeld lacked — the mechanical reason for the cap — SIE provides as an ontological consequence of A1 (information primacy) plus A3 (local Landauer bounds). Born-Infeld's elegance is preserved; the missing physical justification is supplied.

5.5. The Modified Dispersion Relation and the Nyquist cutoff

A discrete substrate requires more than an assertion that Lorentz invariance is preserved — it requires a concrete dispersion relation that recovers Einstein's mass shell at low energy and implements a Planck-scale cutoff at high energy. SIE supplies this directly.

In standard special relativity, a particle's energy and momentum are related by the Pythagorean mass shell:

$$E^2 - (pc)^2 = (mc^2)^2. \tag{15}$$

For a wave propagating through a periodic lattice, however, momentum cannot run to infinity — it is bounded by the lattice spacing, and the correct kinematic variable is not p but a trigonometric function of p through the Brillouin zone structure. Substituting the appropriate discrete-lattice form into Eq. 15 gives the SIE Substrate Dispersion Relation [9, 4, 21]:

$$E^2 - \left[\frac{\hbar c}{\ell_P}\sin\left(\frac{p\ell_P}{\hbar}\right)\right]^2 = (mc^2)^2. \tag{16}$$

This is the κ-Poincaré deformation [75] realized explicitly on the FCC substrate. Two limits determine the physics.

Continuum limit. At energies where $p\ell_P/\hbar \ll 1$, the Taylor expansion $\sin(x) \approx x - x^3/6$ gives $\sin(p\ell_P/\hbar) \approx p\ell_P/\hbar$ to leading order, and the $\hbar c/\ell_P$ prefactor cancels the lattice spacing. Eq. 16 collapses exactly onto the Einstein mass shell (15). Every laboratory-verified prediction of special relativity is preserved because laboratory energies satisfy $p\ell_P/\hbar \sim 10^{-20}$ even at LHC scales.

Nyquist cutoff. As p increases toward the Planck scale, $\sin(p\ell_P/\hbar)$ approaches its maximum at $p_{\max} = \pi\hbar/(2\ell_P) \approx \hbar/\ell_P$. The substrate-momentum variable saturates: no matter how much energy is pumped into the particle, the *lattice* momentum cannot exceed $p_{\max}$. Length contraction physically halts at ℓ_P, preventing the formation of zero-volume singularities. This is the formal statement of the Planck-scale regularization that closes both the electron self-energy (Section V.A) and the Landau pole (Section V.C).

Photon dispersion: the honest status. A natural objection is that a discrete substrate should also modify the photon dispersion relation, yielding observable energy-dependent arrival times for distant high-energy photons. Fermi LAT constraints from GRB 090510 give $\xi < 0.13$ at 95% confidence for linear subluminal dispersion at the Planck scale [88]. Published corpus work in this framework argues that photons, as gauge bosons of an unbroken U(1) symmetry, propagate along topological null geodesics of the substrate's emergent continuum metric and therefore experience no leading-order dispersion even on a discrete lattice — a structural consequence of gauge invariance rather than an ad-hoc exemption [4]. We state this argument here honestly but note three caveats that a referee would rightly raise:

- On a real discrete lattice (phonons in a crystal), massless modes generically *do* exhibit dispersion at the lattice scale. The claim that SIE's U(1) gauge photon evades this because of gauge invariance is a nontrivial assertion that a fully rigorous substrate Lagrangian would need to demonstrate explicitly.

- A sub-leading dispersion — proportional to $(E/E_{\text{Planck}})^2$ rather than (E/E_{Planck}) — would be expected even under the gauge-invariance argument, and is below current Fermi LAT and LHAASO sensitivity but may be within reach of future observations.

- An honest statement of the framework's posture toward photon dispersion is: the framework predicts small dispersion, provisionally compatible with Fermi LAT, with the precise magnitude controlled by structural details of the substrate action that have not yet been computed.

The photon exemption in its current form should therefore be read as a framework-level expectation (gauge photons decouple from substrate discreteness at leading order), not as a completed derivation. A future tightening of the Fermi LAT bound, or a LHAASO detection at finer precision, could push the framework into tension if sub-leading dispersion turns out to be larger than expected.

For massive fermions the subluminal velocity correction is $\delta v/c \sim \xi E/2E_{\text{QG}}$. At the GZK cosmic-ray threshold $E \approx 6 \times 10^{10}$ GeV this is $\sim 10^{-7}$, below current measurement precision but within reach of next-generation ultra-high-energy cosmic ray observatories.

5.5.1. The soccer ball problem resolved

A well-known objection to any Planck-scale momentum cutoff is the *soccer ball problem* [89]: if no single particle can exceed $p_{\max} = \hbar/\ell_P$, how can a macroscopic object (a soccer ball, a planet) carry its easily observed aggregate momentum, which is trillions of times larger? In some deformed-relativity schemes this problem is genuinely fatal, because momentum is postulated to add nonlinearly via the κ-Poincaré coproduct, which gives pathological predictions for macroscopic systems.

In SIE the resolution is immediate and geometric. On the substrate, every fundamental particle is an isolated execution thread at a localized set of nodes. A macroscopic object is not a single thread with a single aggregate momentum but an ensemble of $\sim 10^{25}$ independent parallel threads (one per constituent fermion), each with its own local momentum well below $p_{\max}$. The aggregate momentum of the soccer ball is the sum of trillions of sub-Planckian momenta distributed across trillions of distinct substrate nodes. No single node ever sees a super-Planckian momentum; the grid is never forced to exceed its Nyquist limit.

This is equivalent to how a computer cluster with a per-core bandwidth limit can collectively process data at rates far exceeding any single core: parallelism distributes the load. Macroscopic momentum is never localized at a single lattice site, so the Nyquist cutoff per site is never approached.

> The MDR with Nyquist cutoff is the formal implementation of DSR on the FCC substrate. It recovers Einstein's mass shell at low energy, caps lattice momentum at the Planck scale, preserves Lorentz covariance via the κ-Poincaré deformation, and resolves the soccer ball problem through parallelism. Photons, being massless gauge bosons, are expected to decouple from substrate discreteness at leading order via a gauge-invariance argument; a fully rigorous derivation of this decoupling and its sub-leading corrections is an open methodological task. Fermi LAT constraints are currently compatible with the framework under this expectation but could become a tension if tightened further.

5.6. Lorentz invariance: why this is not a preferred frame

A common objection to any Planck-scale cutoff is that it introduces a preferred frame and breaks Lorentz invariance. In SIE this objection fails for two independent reasons.

First, the cutoff is not a preferred length in ordinary space but an invariant scale under Doubly Special Relativity [75], in which both the speed of light c and the Planck length ℓ_P are observer-independent. Just as the ordinary Lorentz group recovers naturally at low velocities, the deformed symmetry of DSR recovers ordinary Lorentz invariance in any regime where substrate discreteness is invisible.

Second, the substrate-induced modifications that could in principle break Lorentz invariance (e.g., the Infodynamic Shear Tensor $\Upsilon_{\mu\nu}$ of [21, 67]) are carried by *irrelevant operators* in the renormalization-group sense. A renormalization-group analysis of their β-functions drives their expectation values to zero in the infrared, recovering exact Minkowski SO(3,1) symmetry as an effective low-energy symmetry [76]. Put concretely: whatever substrate-induced Lorentz violation exists at the Planck scale is suppressed by $(E/M_P)^n$ at laboratory scales and is experimentally unobservable.

The continuum limit is further supported geometrically by the Connes distance formula [77] applied to the spectral triple formed by the FCC node algebra plus the Dirac operator. This gives an explicit construction showing how continuous manifold distance emerges from discrete node distances in the appropriate limit. The substrate is *geometrically* consistent with observed Lorentz invariance, not just phenomenologically.

5.7. The falsifiable signature

Dielectric saturation is not an untestable theoretical flourish. It predicts a specific departure from QED running coupling behavior at high energies. The departure is small at currently accessible collider scales — the tanh correction is $\mathcal{O}((E/E_{\text{Planck}})^2)$ in the low-energy expansion — but accumulates at accelerator frontiers. A next-generation 100 TeV collider would probe the coupling's running over several orders of magnitude more than LHC and could begin to discriminate the logarithmic QED slope from the slower SIE approach to saturation at the percent level.

More immediately, the same saturation principle predicts the 35% magnetar polarization cap observed by IXPE (Section IX.D): the substrate cannot accept polarization beyond its structural yield point, and X-ray polarization from magnetars above the Schwinger critical field therefore caps regardless of source-specific details [46]. A single saturation constant $\Pi_{\max}$ controls both the Landau pole regularization and the magnetar polarization cap. A future observation of a magnetar exceeding $\sim 50\%$ linear polarization would falsify the framework's saturation claim.

Resolution of QED infinities, summary:

- **Electron self-energy divergence**: eliminated geometrically because $r_e \sim 5.5\,\ell_P > 0$, giving a finite Bekenstein-bounded charge radius. No renormalization of infinities required.
- **Landau pole**: eliminated by tanh-saturation of vacuum polarization at $\Pi_{\max}$. The denominator in the running coupling locks at a finite value, not zero. See Eq. 13.
- **Born-Infeld resurrected**: the Saturated Transmission term of the Master Equation is exactly the Born-Infeld Lagrangian, with $\Pi_{\max}$ given a physical mechanism (Bekenstein bound on substrate bit density) rather than postulated.
- **Lorentz invariance preserved**: DSR at the Planck scale plus RG-irrelevant operators at laboratory scales recover exact SO(3,1) to observational precision.
- **Falsifiable**: magnetar polarization cap at 35% (IXPE-confirmed); future collider constraints on the departure of $\alpha(E)$ from pure QED running.

This is, in many ways, the SIE framework's cleanest single victory over standard physics. Not a new prediction that could be post-hoc rationalized, but the removal of two known infinities that every physicist has had to learn to live with — replaced with a finite geometric mechanism that recovers all the experimentally verified QED behavior at accessible energies while curing the pathology at Planck scales. It is the application of the substrate ontology that most directly vindicates Feynman's unhappiness with renormalization: the universe does not contain the infinities that need to be swept under the rug. The infinities were always an artifact of assuming a continuous vacuum.

6. The Derivation Chain: Reconstructing the Standard Model

The gauge structure SU(3)×SU(2)×U(1) and fermion chirality content of the Standard Model are not assumed in SIE; they are claimed to emerge from the substrate geometry. Five problems have identifiable structural mechanisms; we summarize them here and flag residual technical gaps in Section IX.

6.1. Problem 1: The chirally-asymmetric Wilson coefficient

On an FCC lattice with a topological twist parameter $\theta_{\rm sub}$ (the "rifling" of the substrate about a preferred axis), Wilson fermions acquire a chirality-dependent mass term. Left-handed modes remain massless while right-handed modes acquire a Planck-scale gap:

$$r_{\rm asym} = c_{\rm FCC}\,\theta_{\rm sub}. \tag{17}$$

The coefficient $c_{\rm FCC}$ is a geometric constant specific to the FCC lattice. Four independent arguments — lattice dispersion analysis, anomaly matching, Atiyah–Singer index theorem, and Wilsonian effective action integration — establish that $c_{\rm FCC}$ is of order unity [14]. A rigorous numerical value requires dedicated lattice gauge theory computation (Section IX).

The consequence: observed chirality of the weak interaction, with right-handed states sterile under $SU(2)_L$, is the infrared shadow of the substrate's topological rifling. The weak force's left-handedness is not a postulate; it is a direct reading of FCC substrate geometry.

6.2. Problem 2: Strong-CP suppression

The same $\theta_{\rm sub}$ that gives chirality to the weak interaction might naively give an observable strong-CP phase $\theta_{\rm QCD}$. Observation requires $|\theta_{\rm QCD}| < 10^{-10}$ [16]. SIE resolves the tension in two layers. The first is an embedding protection factor from the Pati–Salam breaking $SU(4)_c \to SU(3)_c\times U(1)_{B-L}$: instanton charge splits between the SU(3) and U(1) sectors, and the natural hierarchy v/M_{PS} provides a suppression factor that dimensional analysis places between 10^{-13} and 10^{-26}. The second layer is an emergent axion from substrate dynamics with $f_a \sim 10^{9.5}$ GeV and mass $m_a \approx 3.5$ meV [15].

Analysis of the embedding factor (from the technical gaps document) suggests it may be individually sufficient for consistency, making the axion a redundancy rather than a necessary second layer. This matters operationally: if IAXO helioscope searches [17] find no axion, SIE's strong-CP position remains defensible on embedding grounds alone. The revised mass 3.5 meV targets IAXO rather than the ADMX haloscope searches, and the axion contribution to dark matter is $\sim 0.06\%$ rather than 1%.

6.3. Problem 3: Gauge-boson halving

The Standard Model has 12 gauge bosons (eight gluons, three weak, one hypercharge). A naive count from SU(4)⊃SU(3)×U(1) combined with $SU(2)_L \times SU(2)_R$ would give 24. The factor-of-two reduction parallels the chiral halving of fermions: helicity alignment with the substrate rifling projects 24 naive generators to 12 physical ones [18]. The mechanism is structurally identical to fermion chiral halving, as should be the case if both arise from the same substrate topology.

6.4. Problem 4: FCC to Standard Model

The full path from FCC geometry to the Standard Model:

1. The FCC lattice has point group O_h, which in the continuum limit gives root system A_3.

2. A_3 is the root system of su(4), the Pati–Salam color group $SU(4)_c$.

3. Breaking $SU(4)_c \to SU(3)_c \times U(1)_{B-L}$ plus the chiral halving of Problem 1 recovers the Standard Model gauge group $SU(3)_c \times SU(2)_L \times U(1)_Y$.

4. Fermions transform in the Pati–Salam multiplets $(4,2,1) + (\bar{4},1,2)$, which decompose to the correct Standard Model representations with electric charge $Q = I_{3L} + I_{3R} + (B-L)/2$ [19].

The origin of $SU(2)_L \times SU(2)_R$ is attributed to a two-phase principle (the substrate's distinction between dynamic/kinetic and static/inertial informational modes), which is motivated but not yet derived from first principles.

6.4.1. Lepton generations: a two-parameter calibration to shell geometry

The three fermion generations are claimed to correspond to distinct icosahedral shell configurations of the substrate. This subsection presents the quantitative form of that claim honestly — as an empirically calibrated two-parameter scheme that lands within $\sim 1\%$ of both heavy-lepton masses, not as a zero-parameter first-principles derivation. A genuine first-principles derivation remains an open task (Section XIV.C).

The shell formula. On a 3D close-packed lattice, the number of nodes in the K-th icosahedral shell is:

$$N_{\text{shell}}(K) = 10K^2 + 2. \tag{18}$$

This is the standard Caspar-Klug combinatorial formula that governs icosahedral capsid geometry, fullerene structure, and spherical shells on 3D lattices [92]. It is fixed by combinatorics, not a free parameter.

The calibration inputs. Expressing lepton mass ratios via Eq. 18 requires two calibration choices:

1. A baseline constant N_{core} identifying the electron's effective bit content.

2. A discrete selection rule for which shell indices K are occupied by the muon and tau.

For the baseline we use $N_{\text{core}} = 166.25$, matching the FCC topological capacity derived for hadronic boundary states [21]. This is a *calibration choice*, not a derivation: we adopt the hadronic FCC capacity as the reference scale and ask whether integer shells built upon it reproduce the observed lepton masses. We do not independently derive that the electron's substrate capacity must equal the hadronic capacity; this would require a unified substrate action that currently does not exist.

For the selection rule we take $K = 2$ for the muon and $K = 18$ for the tau. The motivations for these specific indices are plausibility arguments rather than derivations: $K = 2$ is the smallest shell large enough to support a spatially distinguishable closed structure (a 12-node $K = 1$ coordination shell is geometrically degenerate with the core), and $K = 18$ is the smallest shell that breaks planar degeneracy via the factor-of-three-squared scaling expected for spatial isotropy. Neither argument is rigorous.

Accepting those two inputs, the framework predicts:

$$\frac{m_\mu}{m_e} = N_{\text{core}} + N_{\text{shell}}(2) = 166.25 + 42 = 208.25, \tag{19}$$

$$\begin{aligned} \frac{m_\tau}{m_e} &= N_{\text{core}} + N_{\text{shell}}(2) + N_{\text{shell}}(18) \\ &= 166.25 + 42 + 3242 = 3450.25. \end{aligned} \tag{20}$$

Compared with observation: $m_\mu/m_e = 206.768$ (Eq. 19 accurate to 0.7%) and $m_\tau/m_e = 3477.15$ (Eq. 20 accurate to 0.8%) [61].

The honest framing. Two inputs (the 166.25 baseline and the selection of $K \in \{2, 18\}$) together predict two numbers (the muon and tau mass ratios) to within $\sim 1\%$. This is striking: two inputs predicting two outputs at the 10^{-2} level is not a generic feature of arbitrary parametrizations. But it is not zero free parameters. An honest accounting is: one baseline bit count + one shell-selection rule + one combinatorial formula = two predictions, both landing at $< 1\%$.

The residuals have plausible interpretations: the muon prediction overshoots by $\sim 1.5m_e$, consistent with the shell being slightly over-formatted for stable retention (and the muon's observed 2.2 μs lifetime); the tau prediction undershoots by $\sim 27m_e$, consistent with virtual QCD activity in the massive shell (the tau is the only lepton heavy enough to decay into hadrons). These are post-hoc but not arbitrary.

Lepton masses from substrate shells: $m_\mu/m_e = 208.25$ (vs 206.77, 0.7%) and $m_\tau/m_e = 3450.25$ (vs 3477.15, 0.8%), from $N_{\text{core}} = 166.25$ and $K \in \{2, 18\}$. Two inputs predicting two numbers at percent level. The open task is deriving N_{core} and the selection of K from substrate dynamics (Section XIV.C).

Why three generations rather than more. The substrate shell structure cannot support an arbitrary number of stable lepton shells. The limiting criterion is not a simple bulk-versus-boundary Bekenstein inequality (a naive such argument would already forbid the tau, since $\sum_k N_{\text{shell}}(k)$ exceeds $N_{\text{shell}}(K)$ at any $K \geq 3$; see footnote[1]). The correct criterion must instead compare the outermost shell's *own* information content to its localized holographic capacity, factoring in the Compton-radius scaling at the shell's binding energy. The detailed form of this criterion is an open task: the claim "the Bekenstein bound forbids a fourth generation" is a plausible structural expectation but has not been rigorously derived.

Empirically, precision electroweak tests (the effective number of light neutrino species $N_\nu = 2.984 \pm 0.008$ from LEP Z-pole measurements) already strongly constrain a fourth light-lepton generation [61]. The framework is *consistent with* the three-generation observation but does not yet uniquely predict it from first principles.

6.5. Problem 5: Yukawa compensation

A potential worry: does the chirally-asymmetric Wilson term affect observed Dirac masses? The answer is no. The Wilson term is dimension-5 and vanishes in the infrared continuum limit: at the electroweak scale, $\mu a \sim 10^{-17}$, where μ is the physical scale and a the substrate spacing, rendering the Wilson contribution invisible to laboratory physics.

[1] A naive bulk-boundary comparison gives $S_{\text{bulk}}(3) = 12 + 42 + 92 = 146 > S_{\text{boundary}}(3) = 92$. This inequality is violated before the muon, let alone the tau, which rules out using it as the three-generation selection criterion.

Observed Dirac masses arise from standard Higgs–Yukawa with Yukawa matrix Y_{ij} [20].

Individual mass predictions within SIE (from various substrate-theoretic arguments) include: proton mass via DRAM Landauer ($\sim$ 3% accuracy), m_p/m_e via volumetric bit counting ($\sim$ 0.4%), neutrino mass via Landauer times α^{-1} amplification (Section XI.C below), lepton generation ratios via icosahedral shells (just derived, $\sim$ 1%), and quark mass ratios via dual-clock truncation ($\sim$ 10%–30%) [9, 21].

The chain from FCC substrate to Standard Model particle content has identifiable structural mechanisms for each step. Four technical gaps remain (Section XIV); none represents a structural inconsistency. What was formerly "the framework postulates chirality" is now "the framework offers a mechanism for chirality from substrate topology with $c_{\text{FCC}} \sim \mathcal{O}(1)$, pending rigorous LGT calculation."

7. Observational Content I: Galactic Kinematics

The framework's first major empirical test is the rotation curves of spiral galaxies. SIE predicts a universal acceleration scale

$$a_0 = \frac{cH_0}{2\pi} \tag{21}$$

from two independent substrate arguments that converge on the same numerical result.

7.1. Deriving a_0 from the substrate noise floor

The first derivation — de Sitter horizon. An observer accelerating at rate a sees a thermal bath at the Unruh temperature $T_U = \hbar a/2\pi k_B c$ [85]. A non-accelerating observer at the cosmological horizon sees a different thermal bath, the Gibbons-Hawking temperature $T_{dS} = \hbar H_0/2\pi k_B$ [59]. SIE requires these two temperatures to coincide at the galactic-dynamics transition scale, because that is where the substrate's global equilibrium matches the baryonic system's local inertial response. Setting $T_U(a_0) = T_{dS}$:

$$\frac{\hbar a_0}{2\pi k_B c} = \frac{\hbar H_0}{2\pi k_B} \quad \Rightarrow \quad a_0 = \frac{cH_0}{2\pi}. \tag{22}$$

The second derivation — deep substrate floor. Independent of horizon arguments, the substrate has a characteristic ground-state temperature set by its own lowest-energy quantum mode fluctuations against residual structure. Dimensional analysis places this at approximately $T_{\text{vac}} \sim 10^{-31}$ K, which emerges from equating the Bekenstein noise floor of the Hubble volume with the Landauer cost of maintaining one bit at the Hubble time scale [38, 5]. A baryonic system accelerating slowly enough that its Unruh temperature drops below this floor can no longer dissipate Landauer processing heat into a substrate hotter than itself. At that threshold the substrate must engage a buffering response. Solving $T_U(a_0) = T_{\rm vac}$ gives the same a_0:

$$a_0 = \frac{2\pi c k_B T_{\rm vac}}{\hbar} \approx 1.1 \times 10^{-10}\ \mathrm{m/s^2}, \qquad (23)$$

matching Eq. 21.

The convergence between the two temperature-matching arguments is meaningful: the Hubble temperature and the substrate's intrinsic noise floor must coincide if the de Sitter horizon is genuinely the boundary at which the substrate thermalizes to its own residual fluctuations. The framework's broader hypothesis is that this same substrate equation of state governs other low-noise regimes (magnetar polarization saturation, cosmic ringdown boundaries); that hypothesis is at various stages of quantitative support across those domains (Section X).

The third equivalent form — Landauer. The same result follows from requiring that a_0 be the minimum erasure rate at which the substrate can sustain informational coherence across a single Hubble time. Expressed as a rate of bit deletion per unit length per unit time, this gives $a_0 = cH_0/2\pi$ via Landauer's bound [22].

A note on independence. The three presentations are not three independent derivations but three equivalent formulations of the same underlying substrate relation. $T_{\rm dS} = \hbar H_0/2\pi k_B$ and $T_{\rm vac} \sim 10^{-31}$ K differ by $\mathcal{O}(1)$ dimensional factors but refer to the same cosmological thermal floor; the Landauer form rearranges the same identity in informational rather than thermal language. The meaningful physical content is that setting substrate thermalization at the horizon equal to the Unruh response of a baryonic system gives $a_0 \sim 10^{-10}$ m/s^2 — one relation with three equivalent presentations, not three independent triangulations. We record this honestly because the distinction between independent and equivalent matters for how seriously to take coincidence arguments.

7.2. SPARC rotation curves

Tested against the 175-galaxy SPARC catalog [23], the SIE rotation curve model

$$V_{\rm SIE}^2(r) = V_{\rm bar}^2(r) + \sqrt{V_{\rm bar}^2(r)\, a_0\, r} \qquad (24)$$

fits the full SPARC sample with a single free parameter per galaxy, the stellar mass-to-light ratio $\Upsilon_{\rm disk}$. For comparison, we fit the Navarro–Frenk–White halo [24] using the cosmologically-motivated Dutton–Macciò concentration-mass relation [25] to reduce NFW to a two-parameter model ($\Upsilon_{\rm disk}$, M_{200}; concentration c is set by M_{200} rather than floated freely, following standard SPARC-community practice).

Under this setup, SIE achieves median $\chi^2_\nu = 3.33$ versus NFW's median $\chi^2_\nu = 1.72$. BIC model selection, which penalizes extra parameters, yields SIE as strongly preferred (ΔBIC > 10) in $\sim 11\%$ of galaxies, positively preferred ($2 < \Delta$BIC ≤ 10) in an additional $\sim 26\%$, and ambiguous or NFW-preferred in the remaining $\sim 63\%$. Under looser NFW priors (e.g., fully-free concentration), the balance shifts toward NFW; under stricter priors (e.g., Ypsilon constrained to stellar-population-synthesis bounds), toward SIE. We report the standard-prior result as the primary number.

The key observation from this comparison is *parameter efficiency*: SIE achieves a goodness-of-fit within a factor of two of cosmologically-constrained NFW while using one fewer parameter, and does so *while making an independent zero-free-parameter prediction* for the acceleration scale a_0 (Section VII.A). A one-parameter model that predicts its own coupling constant and remains competitive with a two-parameter halo model is an unusual position to be in; we treat this as the meaningful empirical content rather than the dominant-BIC-preference claim that the corpus previously emphasized.

SPARC parameter efficiency result: SIE fits 175 SPARC galaxies with 1 parameter/galaxy, achieving median $\chi^2_\nu = 3.33$ against NFW's 1.72 with 2 parameters/galaxy (cosmologically-motivated $c(M_{200})$ prior). SIE is BIC-preferred in $\sim 11\%$ strongly and $\sim 37\%$ in total; NFW-preferred in the remainder. The meaningful result is that a one-parameter model (with a first-principles prediction for its coupling a_0) is competitive with a cosmologically-constrained halo model.

7.3. The Radial Acceleration Relation

Across 3382 data points in 175 galaxies, the SIE prediction (Eq. 21) evaluated at the Planck cosmology value $H_0 = 67.4$ km s^{-1} Mpc^{-1} gives $a_0 = 1.04 \times 10^{-10}$ m s^{-2}. The residual distribution has mean -0.0036 dex (1.51σ from zero) and scatter $\sigma = 0.140$ dex, compared to the published SPARC intrinsic scatter of 0.11 dex. The prediction is unbiased with zero free parameters.

7.4. A kinematic Hubble constant

Fitting a_0 freely from the RAR gives $a_0 = 1.1444 \times 10^{-10}$ m s^{-2}. Via Eq. 21, this implies

$$H_0^{\rm kin} = 73.97 \text{ km s}^{-1}\text{Mpc}^{-1}, \tag{25}$$

within 1.3% of the SH0ES distance-ladder value and intermediate between the two Hubble tension anchors. This is a Hubble constant measurement from rotation curves alone, with no supernovae or CMB input.

7.5. The Pantheon+ void test: status after reanalysis

The framework's substrate-tension mechanism predicts, from $\delta_b \approx 0.15$, a void-interior/exterior expansion-rate difference of

$$\Delta H \equiv H_{\rm local} - H_{\rm global} \approx 73.1 - 67.4 = 5.7 \text{ km s}^{-1}\text{Mpc}^{-1}, \tag{26}$$

i.e., a $\sim 9\%$ step. This section documents where that prediction stands against Pantheon+ [53].

First-pass step-function analysis. An earlier analysis in the corpus [38, 22] fit a step-function void model to Pantheon+ with fixed Ω_m at the Planck value and obtained $H_{\rm in} = 64.6$, $H_{\rm out} = 74.5$ km s^{-1} Mpc^{-1} — an inverted hierarchy of order 10 km s^{-1} Mpc^{-1}, seemingly contradicting the SIE prediction at high significance.

Re-analysis with the full covariance and floated Ω_m. A re-analysis using the full stat+sys covariance matrix from the Pantheon+SH0ES data release, with Ω_m varied jointly with the void parameters, gives substantially different numbers. The step-function fit yields $H_{\rm in} = 73.23$, $H_{\rm out} = 73.47$ km s^{-1} Mpc^{-1} ($\Delta H = 0.24$ km s^{-1} Mpc^{-1}, consistent with no step in either direction). A smooth-profile void fit (Gaussian $\delta_b(r)$) gives $\delta_0 = -0.012$, against a framework prediction of $\delta_0 = +0.15$. Bayesian model comparison prefers plain ΛCDM (with free H_0 absorbing the SH0ES tension) over either void model at $\Delta\text{BIC} \approx +7$. The mean implied H_0 across 16 redshift bins is 73.20 ± 1.21 km s^{-1} Mpc^{-1}, essentially flat from $z \approx 0.01$ to $z \approx 2.3$.

Peculiar-velocity systematic. Re-running the smooth-profile fit using uncorrected CMB-frame redshifts (`zCMB`) instead of the 2M++-corrected `zHD` shifts δ_0 from -0.012 to -0.043 — a factor of 3.5 change in magnitude, confirming that the flow-model treatment is a significant systematic. However, this shift is much smaller than the ~ 0.16 gap between either fitted value and the framework's predicted $\delta_0 = +0.15$. Flow-model refinement alone cannot rescue the prediction at the measured systematic magnitude.

What the data say. Under the standard Pantheon+ analysis with SH0ES-anchored distance moduli, the framework's predicted void gradient ($\Delta H \approx 5.7$ km s^{-1} Mpc^{-1}) is ruled out at the $\sim 20\sigma$ level relative to the fitted $\Delta H = 0.24\pm$ statistical uncertainty. This is not a neutral "undiscriminated" result; the predicted magnitude of the gradient is directly inconsistent with the data *under the standard analysis.*

The remaining defense and its open status. The framework's one viable defense against this result is structural: the SH0ES Cepheid anchors (in galaxies within ~ 40 Mpc) reside inside the KBC void. If the SH0ES pipeline forces calibration consistency between the local anchors and the Hubble-flow ($z \approx 0.023$ to 0.15) supernovae, both inside the void, a void-interior calibration value becomes baked into the pipeline. In this scenario, high-z supernovae could then be reported with distance moduli consistent with the interior rate because the absolute calibration has been anchored to the interior value.

However, demonstrating that this scenario reproduces the flat-H_0 signature observed across the full z range ($0.01 < z < 2.3$) requires an explicit calculation: carrying an SIE prediction ($\delta_0 = +0.15$ void against global $H_0 = 67.4$) through the SH0ES anchoring pipeline and checking whether the output distance moduli match the observed pattern. That calculation is not in the current corpus and constitutes the primary open empirical task for the framework. Without it, the scenario remains a conjecture about pipeline systematics rather than a validated alternative.

There is also a self-consistency concern to flag: elsewhere in this paper (Section VII.A), $a_0 = cH_0/2\pi$ is derived from cosmological substrate arguments and compared against SPARC with the fitted $H_0^{\rm kin} = 73.97$. Under the defense scenario (global $H_0 = 67.4$, locally-elevated $H_0 \approx 73$), the SPARC fit to a_0 should reflect the *local*-observer value while the underlying substrate noise floor is set by the *global* value. Reconciling this — whether SPARC's kinematic a_0 is correctly interpreted as the local or global quantity — is a second open task. These two issues are not independent: both stem from the same question of how an interior observer's measurements relate to the substrate's global quantities.

Honest summary of the empirical status. The framework's specific quantitative prediction of a $\sim 6\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$ gradient across the KBC void edge is not detected and is in tension with the standard Pantheon+ analysis at high significance. The aggregate numerical match $H_0 \approx 73\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$ in the catalog is structurally consistent with the framework's predicted interior-observer value, but is equally consistent with a globally-elevated H_0 and no void mechanism. Separating these requires (i) an explicit calculation of how scenario (b) propagates through the SH0ES anchoring pipeline, and (ii) independent flow-model reconstructions (Cosmicflows-4 [95] or similar) plus distance-ladder anchoring outside the KBC void from Rubin LSST, Roman, and Euclid (2027+). Until (i) is performed, the Hubble tension claim in this paper should be treated as tentatively supported by the aggregate numerical match rather than by an observationally verified mechanism.

Galactic kinematics summary: SIE fits the full 175-galaxy SPARC rotation curve dataset with one free parameter per galaxy, achieving competitive goodness-of-fit versus cosmologically-constrained NFW (median $\chi^2_\nu = 3.33$ vs 1.72, NFW two parameters). The zero-free-parameter prediction $a_0 = cH_0/2\pi$ reproduces the 3382-point radial acceleration relation at 0.140 dex scatter with mean -0.0036 dex (unbiased at 1.51σ) under Planck cosmology. Fitting a_0 freely returns $H_0^{\mathrm{kin}} = 73.97\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$ from galaxies alone. The Pantheon+ void-gradient prediction is not detected; the framework's defense via interior-anchoring requires pipeline-level calculation not yet performed. The galactic kinematics results stand independently of the Hubble tension mechanism.

8. Observational Content II: Cosmology

8.1. N = 57 triple consistency

The single most unexpected number in the framework is 57, which appears in three cosmological contexts. The epistemic status of the three routes is mixed, and we set it out explicitly before discussing what the convergence might mean.

Route 1 — The Standard Model chiral inventory (definitional). Counting the pre-electroweak-symmetry-breaking Standard Model degrees of freedom: 90 fermionic (three generations of fifteen left-handed Weyl fermions, each counted twice for chirality) plus 24 gauge bosonic (eight gluon $\times$ 2 polarizations = 16, three weak $\times$ 2 = 6, one hypercharge $\times$ 2 = 2; total 24). The Higgs, being spin-0, has undefined chirality and drops out automatically. Retaining the left-handed half:

$$N = \frac{90 + 24}{2} = 57. \tag{27}$$

This count is fixed by particle content alone. It is not a prediction — it is an accounting identity on the SM field content. Whether this integer has cosmological significance is the question Routes 2 and 3 address.

Route 2 — The CMB spectral tilt (numerical match, not independent derivation). Evaluating the single-field slow-roll relation $n_s = 1 - 2/N$ [27] at $N = 57$ gives

$$n_s = 1 - 2/57 = 0.96491, \tag{28}$$

matching Planck 2018's $n_s = 0.9649 \pm 0.0042$ [3] at the 0.003σ level. We must be careful about what this establishes.

The standard slow-roll formula $n_s = 1 - 2/N$ relates the spectral tilt to the number of e-folds before horizon exit, under the assumption of a single-field inflaton with specific potential shape. In osmotic inflation, the dynamics are not single-field slow-roll: the expansion is driven by Landauer heat generation from primordial unstable states, not by a classical scalar field rolling down a potential. The question of whether $n_s = 1 - 2/N$ remains a valid relation in the osmotic picture — and what N then means within it — has not been fully worked out.

What can be said:

- The numerical equality $n_s^{\mathrm{Planck}} = 1 - 2/57$ to within 10^{-4} is a real empirical fact independent of derivation.
- Fourteen integers satisfy the 1σ Planck constraint; $N = 57$ is the nearest integer to the best fit $N_{\mathrm{exact}} = 56.98$.
- The osmotic picture needs to either recover the $n_s = 1 - 2/N$ form (in which case the match is principled) or produce its own form that coincides with $1 - 2/57$ by different means (in which case the match is suggestive but not direct evidence).

We present Route 2 as a numerical coincidence pending derivation, not as an independent confirmation of Route 1.

Route 3 — The inflationary e-fold budget (order-of-magnitude estimate). A back-of-envelope calculation [26] suggests the number of osmotic e-folds required to dilute Landauer heat from GUT scale ($\sim 10^{15}$ GeV) baryogenesis to the electroweak threshold (~ 246 GeV) is in the range 50–65. This range encompasses 57, but the ± 10 uncertainty means other integers in the band (50, 52, 55, 60, 62, 65) also satisfy the constraint. Route 3 is an order-of-magnitude consistency check, not a precision derivation.

On the N=57 convergence: Route 1 (chiral inventory) is a definitional accounting identity. Route 2 (spectral tilt) is a numerical match between the SM particle count and the Planck tilt, mediated by a slow-roll relation whose applicability in the osmotic picture is not yet established. Route 3 (e-fold budget) is an order-of-magnitude consistency check. The convergence is striking enough to motivate continued work, but it does not constitute three independent confirmations. It is the framework's central open numerical coincidence.

A further observational constraint is informative: standard single-field slow-roll with $N = 57$ would predict a tensor-to-scalar ratio $r = 8/N \approx 0.14$, which is excluded at 95% confidence by Planck/BICEP/Keck [28]. The osmotic picture decouples r from N — in the osmotic picture, there is no scalar inflaton, so the mechanism that produces tensor modes from quantum fluctuations of the inflaton is absent, and r can be arbitrarily small while n_s reflects the substrate's thermal-dilution-e-fold structure. A clear $r < 10^{-3}$ measurement by LiteBIRD [29] or CMB-S4 [30] would be consistent with this decoupling; a detection of $r \sim 0.14$ would falsify it.

8.2. The vacuum energy density

Applying the Bekenstein bound at the Hubble horizon and the Cohen–Kaplan–Nelson [31] UV/IR bound to obtain the bulk bit density, then multiplying by the Landauer mass at the CMB temperature:

$$\rho_{\text{bare}} = \rho_{\text{raw}} \times m_{\text{bit}} \approx 2.07 \times 10^{-27} \text{ kg m}^{-3}. \quad (29)$$

The observed critical density is $\rho_c \approx 8.53 \times 10^{-27}$ kg m^{-3}. The ratio $\rho_c/\rho_{\text{bare}} \approx 4.15$ at Planck cosmology (and 4.32 at SH0ES) is intended to close via the $N = 57$ equipartition factor [4]. This closure is phenomenological pending derivation of the action principle that would convert the ansatz into theorem.

Against a vacuum catastrophe of 120 orders of magnitude, a factor-of-four residual from first principles is non-trivial. Whether it constitutes genuine unification or a chain of numerical coincidences is an open question recorded transparently in the source paper [4].

8.3. The Hubble tension

The KBC void [32] is a ~ 300 Mpc underdensity with $\delta_b \approx 0.15$ centered near the Milky Way. In SIE, baryonic matter anchors substrate tension; a 15% matter underdensity produces a 15% tension drop. Treating the local metric as a damped harmonic response to substrate tension ($H \propto 1/\sqrt{T_{\text{sub}}}$):

$$H_{\text{local}} = \frac{H_{\text{global}}}{\sqrt{1-\delta_b}} = \frac{67.4}{\sqrt{0.85}} \approx 73.1 \text{ km s}^{-1}\text{Mpc}^{-1}. \quad (30)$$

This lands directly on the SH0ES measurement of 73.0 ± 1.0 km s^{-1} Mpc^{-1} [2].

The status of this match given the Pantheon+ re-analysis. The numerical agreement with SH0ES is real but its interpretation requires care. A re-analysis of Pantheon+ with the full stat+sys covariance and floated Ω_m (Section VII.C) finds the framework's predicted $\Delta H \approx 5.7$ km s^{-1} Mpc^{-1} gradient ruled out under the standard analysis: the fitted step is $\Delta H = 0.24$ km s^{-1} Mpc^{-1}, and the smooth-profile $\delta_0 = -0.012$ is far from the predicted $+0.15$. The framework's only defense is that SH0ES Cepheid anchors sit inside the KBC void, which could in principle cause the anchoring pipeline to conceal the gradient — but this defense requires an explicit pipeline-level calculation (Section VII.C) that has not yet been performed.

Until that calculation is done, the aggregate $H_{\text{local}} \approx 73.1$ match to SH0ES should be read as *numerically consistent with the framework's prediction for an interior observer* rather than as independent validation of the mechanism. The Hubble tension match is therefore a numerical success of the framework whose mechanistic interpretation is tentatively supported but not yet independently verified.

Within the SIE framework, the Hubble tension is understood as an observer-location effect: the predicted interior-void expansion rate $H_{\text{local}} \approx 73.1$ matches the SH0ES measurement because the anchoring calibrators reside inside the void. Whether this is a genuine mechanism or a coincidence with globally-elevated H_0 is not yet observationally discriminated; the framework's explanation requires an explicit pipeline-level calculation and void-exterior anchoring from next-generation surveys (Rubin LSST, Roman, Euclid, 2027+) to validate.

8.4. The growth index

In standard cosmology, the matter growth index γ parametrizes how density perturbations grow, with the phenomenological Wang-Steinhardt form $f \equiv d\ln\delta/d\ln a \approx \Omega_m^\gamma$ giving $\gamma \approx 0.55$ for ΛCDM [33]. The analytic expression under general relativity is $\gamma = 3(w-1)/(6w-5)$, evaluating to $\gamma = 0.545$ at $w = -1$.

SIE's axiom A2 naturally selects $w = -1$ as the late-time asymptotic state (the universe locks into a closed adiabatic configuration at the cosmological horizon; see the Gompertz lifecycle in Section XII). If the Wang-Steinhardt parametrization carries over into the SIE framework — which it does as long as background GR is preserved and the perturbation equations are unmodified at the relevant scales — then $\gamma = 0.545$ is obtained without fitting.

The caveat. Whether the osmotic-allocation picture of dark energy strictly preserves the Wang-Steinhardt derivation's assumptions (in particular the standard continuity equation for dark energy perturbations and the specific GR coupling to matter) is a question we flag but do not resolve here. If the osmotic dynamics modify the perturbation equations, the growth index could differ from 0.545 at a level currently below observational precision but potentially detectable in future Euclid and Rubin data. We record the $0.545 \leftrightarrow 0.55$ match as a structural consistency within the standard perturbation framework, pending derivation of whether that framework survives unchanged in SIE.

8.5. The Schwarzschild–Hubble coincidence

A numerical fact worth recording because it provides structural weight to axiom A2 (the closed-system assumption): the Schwarzschild radius of the observable universe's mass equals the Hubble radius.

Using the empirically measured critical density $\rho_c \approx 8.53 \times 10^{-27}\ \mathrm{kg\,m^{-3}}$ and the Hubble volume $V_H = (4/3)\pi R_H^3$ with $R_H = c/H_0 \approx 1.3 \times 10^{26}$ m, the mass enclosed within the observable universe is $M_{\text{univ}} \approx 8.8 \times 10^{52}$ kg. Plugging this into Schwarzschild's formula:

$$R_s = \frac{2GM_{\text{univ}}}{c^2} \approx 1.3 \times 10^{26}\ \mathrm{m} \approx R_H. \qquad (31)$$

Standard ΛCDM cosmology treats this as a tautology of flat geometry: a critical-density universe is flat ($\Omega = 1$), and a flat universe has $R_s = R_H$ by construction [90, 91]. The numerical coincidence is then attributed to inflationary fine-tuning that drove the universe to critical density in the first place.

SIE offers a complementary reading. The closed-system axiom A2 requires that the cosmic horizon act as an adiabatic wall preventing information and entropy from leaving the system (the Cosmic Ledger). A Schwarzschild horizon at the Hubble scale is *one* geometry that provides such a wall; de Sitter, closed FRW ($k = +1$), and finite-volume topological identifications are others. SIE selects the Schwarzschild interpretation by an additional structural choice — treating the universe as the interior solution of a substrate-saturated horizon — rather than deriving it uniquely from A2 alone.

Under that interpretation, $\Omega = 1$ is natural rather than accidental: it is consistent with A2 and with the specific interior-solution geometry SIE adopts. We do not claim A2 forces this choice; we claim SIE's substrate ontology naturally accommodates the $R_s = R_H$ observation without invoking inflationary fine-tuning [60].

The interior-solution perspective makes the geometry intelligible. Inside a Schwarzschild horizon, the radial coordinate r and temporal coordinate t exchange roles: what looks like inevitable gravitational collapse from the outside looks like inevitable radial expansion from the inside. Our observed cosmic expansion is the interior signature of the substrate's outermost adiabatic wall. We do not live inside a black hole in the colloquial sense of being crushed; we live inside the *interior solution* of a boundary geometry whose signature is precisely the dark-energy-driven expansion we observe.

This reframing is not introduced to claim novelty — the numerical coincidence was noted in the 1970s [90, 91] and is widely known. It is included here because the interpretation matters for the framework: A2 is a natural posture consistent with the Schwarzschild–Hubble equality rather than a philosophical imposition on a neutral geometry. The closed-system structure finds its strongest motivation in the interior-solution reading of the observed $R_s = R_H$ coincidence.

9. Observational Content III: Primordial Crystallography

The pre-expansion universe, in the SIE picture, was a maximally packed FCC substrate crystal [34]. The Big Bang is not a singular event at $t = 0$ but a macroscopic topological phase transition as the substrate crystal fractured along its $\{111\}$ cleavage planes.

9.1. The 7.85° bisector floor

A pristine FCC crystal has four equivalent $\{111\}$ cleavage planes meeting at tetrahedral angles. When the substrate unspooled under macroscopic strain, the primary cleavage plane projected onto the CMB sky as a great circle — exciting the $\ell = 2$ quadrupole. Secondary intersecting planes produce zone axes whose stress concentrations, via the Erdogan–Sih maximum tangential stress criterion [35], project as single-sided hot spots — exciting the $\ell = 3$ octupole.

The bisector between the primary cleavage normal and the secondary zone axis is fixed by the crystal geometry at 15.70°, so the internal bisector-to-primary angle is exactly

$$\theta_{\text{internal}} = 15.70^\circ/2 = 7.85^\circ. \tag{32}$$

This is a zero-free-parameter prediction.

9.2. The Axis of Evil

The observed angular separation between the CMB $\ell = 2$ and $\ell = 3$ axes ("Axis of Evil" [36]) is approximately 8.85°. The SIE prediction of 7.85° matches within 1.0° — well within the Planck satellite's axis extraction uncertainty of 2°–5° for low-order multipoles due to galactic foreground masking.

When the FCC template is optimized in orientation (three Euler angles, with internal geometry fixed by prediction) against Planck 2018 $a_{\ell m}$ coefficients [37], the correlation achieves +0.647 with the observed quadrupole and +0.768 with the octupole. A Monte Carlo against 10,000 randomized isotropic skies returns joint correlation significance $p < 0.001$.

The honest statistics. Because we are optimizing over three Euler angles, the naive Monte Carlo p-value requires a look-elsewhere correction. The effective trials factor for a three-parameter orientation search on S^2 is approximately the number of statistically independent orientations covered by the posterior, which for the resolved multipole structure at $\ell = 2, 3$ is of order 10^2–10^3 [96]. A conservative estimate corrects the $p < 0.001$ to a trials-adjusted $p \sim 0.01$–0.1 — significant but not as striking as the uncorrected number suggests. A rigorous look-elsewhere analysis for this specific template geometry remains to be done.

The residual interpretive issue. The Copi–Huterer–Schwarz–Starkman literature [97, 98] has extensively debated whether the Axis of Evil is a genuine non-Gaussian feature or a foreground-residual artifact. If the anomaly is primarily a foreground, any theoretical template fitting the observed axes is matching a systematic rather than a cosmological signal. The SIE prediction's status therefore depends on the ongoing foreground resolution. We record the correlation as suggestive but provisional, pending improved low-ℓ foreground subtraction (expected from LiteBIRD).

9.3. The kinematic dipole coincidence

A byproduct of the orientation optimization — not a fitting target — is that the predicted quadrupole axis aligns with the independently observed CMB kinematic dipole to within 0.12°. This is a striking coincidence whose physical origin in SIE remains open, but the numerical alignment is recorded.

Primordial crystallography summary: The CMB Axis of Evil is predicted to within 1° from pure FCC substrate geometry with zero free parameters (the 7.85° bisector is fixed by crystallography alone). The optimized orientation aligns with the kinematic dipole at 0.12° as a byproduct. Uncorrected joint significance against 10,000 random skies: $p < 0.001$; trials-adjusted estimate for three Euler angle optimization: $p \sim 0.01$–0.1. The result is suggestive but provisional, pending resolution of the underlying Axis-of-Evil foreground debate.

9.4. The Gompertz strain

A related calibration: the observed 82.0° Axis-of-Evil separation on the sky is reconciled with the pristine 70.53° FCC tetrahedral angle via uniaxial Gompertz strain $s \approx 2.516$ along the primary cleavage axis. This is a consistency check bridging theoretical topology with the observed epoch, not a first-principles derivation; the zero-parameter claim attaches to the 7.85° bisector floor, which emerges from crystal geometry alone [5, 34].

10. Observational Content IV: Black Hole Ringdown and Other Substrate Probes

10.1. Superconducting qubit decoherence (provisional)

A heuristic connection between substrate physics and superconducting qubits has been proposed: analysis of IBM Qiskit telemetry [38] identifies a characteristic non-Markovian decoherence feature near

$$f \sim 5 \text{ GHz} \tag{33}$$

in transmon architectures. Typical transmon devices are designed to operate in the 4–6 GHz band, so any frequency in this range is common and does not by itself indicate substrate physics. The corpus work attempts to extract a *substrate-physics* signature by factoring out design-dependent Josephson parameters and identifying a residual non-Markovian boundary; whether this residual is genuinely substrate-imposed versus two-level-defect or dielectric-loss phenomenology requires further experimental discrimination [38, 93].

The honest framing. We do *not* claim that the numerical value ~ 5 GHz follows from the galactic $a_0 \sim 10^{-10}\ \mathrm{m\,s^{-2}}$ through a first-principles dimensional mapping; no such derivation currently exists in the corpus. The cross-scale consistency proposal is a qualitative hypothesis: if the substrate has a single underlying equation of state, then disparate decoherence phenomena (galactic rotation curves, magnetar polarization caps, qubit decoherence limits) should trace common structural parameters. Quantitative cross-scale matching is an open task, not a result.

Claims elsewhere in the corpus of “scale invariance across seventeen orders of magnitude” based on the co-occurrence of galactic and qubit frequencies should be read as a program statement about what the framework hopes to achieve, not as a completed derivation. We retract the stronger phrasing here.

10.2. Black hole ringdown

On September 14, 2015, LIGO observed the merger GW150914 producing a 62 $M_\odot$ remnant whose dominant ringdown frequency was ~ 251 Hz [39]. In SIE, the event horizon is a physical membrane — an adiabatic wall at substrate saturation — with a fundamental resonant frequency set by the light-travel time around its equatorial circumference:

$$f_{\text{wall}} = \frac{c}{2\pi R_s} \approx 261 \text{ Hz}. \tag{34}$$

This is within 4% of the observed value.

10.2.1. Spin structure of the coincidence

The observed 251 Hz is the $(\ell, m, n) = (2, 2, 0)$ quasi-normal mode of a Kerr black hole with spin $a_\star \approx 0.67$ [40]. The SIE prediction $f = c/2\pi R_s$ corresponds to the $m = 0$ axisymmetric drumhead mode — a different mode family. The near-numerical match (261 Hz vs. 251 Hz) is a coincidence at this specific spin, arising because the equatorial circumference of any Kerr horizon equals $4\pi M = 2\pi R_s$ exactly (from the identity $r_+^2 + a^2 = 2Mr_+$), so the SIE prediction is spin-invariant at $m = 0$ while the GR $(2, 2, 0)$ mode is spin-dependent [41].

10.2.2. The principled test

The framework's $m = 0$ prediction and GR's $m = 0$ prediction differ by 12.7%:

$$f^{\text{SIE}}_{m=0} = 0.0796/M \tag{35}$$

$$f^{\text{GR}}_{m=0} = 0.0706/M. \tag{36}$$

This difference is the genuine discriminating test, requiring next-generation detectors (Cosmic Explorer [42], Einstein Telescope [43]) capable of separating the weaker $m = 0$ mode from the dominant $m = 2$ at sufficient signal-to-noise. The 4% “coincidence” at GW150914 is honest acknowledgment that cross-mode comparison is not a clean test.

10.3. Gravitational wave echoes

For an SIE Planck core at geometric saturation, inbound gravitational waves reflect off the saturated interior with echo delay

$$\Delta t = \frac{4GM}{c^3}\left[\tfrac{2}{3}\ln(M/M_P) + \ln(2(4\pi/3)^{1/3})\right]. \tag{37}$$

The 2/3 logarithmic slope differs from the standard ECO prediction of slope 1.0 [44]. For the GW150914 remnant at spin 0.67, the baseline isotropic delay is 75.9 ms, with Kerr spin adjustment giving 89.1 ms [45]. This is a hard falsifiable target for LIGO O4 and O5 searches.

10.4. The Bullet Cluster without collisionless matter

The Bullet Cluster [51] is typically treated as decisive evidence for particle dark matter: the X-ray gas lags behind the stellar mass in the post-collision geometry, while gravitational lensing follows the stars. Standard dark matter models explain this by noting that collisionless dark matter particles follow the stars, while the collisional gas is decelerated.

In SIE, the same observation follows from the temperature-dependent substrate response. The stars remain kinematically cold in the collision ($a_{\text{thermal}} \ll g_N$ where a_{thermal} is the internal thermal acceleration of the constituent matter), so the substrate reliably amplifies their gravitational signal via the standard rotation-curve mechanism. The gas, however, heats into a chaotic plasma at $T \sim 10^8$ K; the thermal acceleration overwhelms the local Newtonian binding ($a_{\text{thermal}} \gg g_N$) and the substrate's many-body-localized phase breaks locally. The gas continues to exert standard Newtonian gravity but temporarily loses its structural lensing envelope [5].

This is a structurally different mechanism from MOND, which has historically struggled with the Bullet Cluster. SIE's temperature-dependent amplification is qualitatively consistent with the observed displacement between the gas and lensing peaks without invoking a separate particle species, because it predicts the right sign and qualitative pattern (stars lens, hot gas does not). The framework also predicts that substrate-amplification breaks down in any sufficiently hot plasma, which connects to independent astrophysical tests.

On quantitative agreement. MOND-family theories have famously struggled to produce the *magnitude* of the Bullet Cluster lensing offset without residual collisionless matter; the observed lensing mass exceeds the baryonic mass by approximately an order of magnitude at the cluster-scale acceleration regime [99]. The SIE amplification formula $A = 1 + (\sqrt{a_0/g_N} - 1)\exp(-a_{\text{thermal}}/g_N)$ gives the qualitative story of where amplification occurs versus where it breaks down, but the quantitative lensing-offset magnitude at cluster scales has not been computed from the formula in the current corpus. Full quantitative agreement — matching the lensing-to-baryon ratio at the specific Bullet Cluster geometry — remains an open test. We record this as a structural story pending quantitative verification, not as a completed resolution.

10.5. Magnetar polarization cap

The IXPE satellite [46] observed magnetar X-ray polarization capped at $\sim 35\%$, well below the standard QED prediction of 80–100% based on vacuum birefringence in strong magnetic fields [63]. The anomaly has attracted ad-hoc explanations involving hypothetical solid crusts or atmospheric phasing.

In SIE, the 35% cap corresponds to the substrate's mechanical yield point: the vacuum cannot sustain twist beyond this threshold without breaching its bit-density saturation. The field cannot be polarized further because there are no more substrate modes available to host the polarization. This follows from the same substrate saturation principle that caps EM field strength via the Born-Infeld denominator in the Saturated Transmission term of Eq. 8.

The prediction is falsifiable: SIE forecasts that the 35% cap should be a property of the substrate itself rather than a source-specific parameter. If it is a universal substrate limit, it should be reproduced across independent magnetar observations; if it varies meaningfully by source, that would indicate local mechanisms and undermine the substrate-saturation interpretation. The IXPE observation of 4U 0142+61 establishes the phenomenology at one source; confirmation of a similar cap in additional magnetars (Section X) would strengthen the universal interpretation, but claiming universality on a single observation is premature.

10.6. Primordial lithium and GZK suppression

Two further tensions at the frontiers of astrophysics admit SIE resolutions. The lithium-7 problem — standard BBN predicts $\sim 3\times$ more ^{7}Li than observed [64] — is addressed by Landauer erasure of isotope formation pathways in the early universe, reducing ^{7}Li production without affecting the dominant light-element abundances. The GZK cutoff [65] in ultra-high-energy cosmic rays emerges from substrate bandwidth saturation for photons above the CMB interaction threshold, with a predicted dispersion signature $\Delta t \propto E$ linear in photon energy. The LHAASO observation of TeV photons from GRB 221009A [52] constrains this linear dispersion; continued observations tighten it.

These extensions are sketches, not detailed derivations, and are included here to show the framework's engagement breadth rather than to claim definitive resolution. Full treatment would require dedicated papers for each.

11. The Chiral Grain, Antimatter, and Black Hole Information

Two of physics's deepest puzzles — why the universe contains matter rather than annihilating into a photon bath, and whether black holes destroy quantum information — have structurally similar resolutions in SIE. Both follow from the same substrate property: an information-carrying medium with directional structure. This section presents those resolutions because they are consequences of the ontology that deserve articulation rather than being buried as philosophy.

11.1. The chiral grain of the vacuum

A central implication of A1 (information primacy) is that the substrate is not a smooth isotropic manifold. It is a discrete FCC lattice with a topological twist — the θ-term of the Master Equation (8) — and that twist has a direction. In the language of material science, the substrate has a *grain*, in exactly the way wood or a woven fabric has a grain. Motion along the grain encounters lower resistance than motion against it.

The most direct observational handle on this grain is the neutrino. A neutrino is not merely a light fermion — in the SIE picture it is the *minimum possible data packet* on the substrate, carrying only one bit of chiral information (Section XI.C below). Because it has essentially no shielding mass, it samples the substrate's structure unfiltered. The striking experimental fact is that neutrinos are always left-handed (neutrino spin is anti-parallel to momentum); right-handed neutrinos have never been observed despite decades of searches [78].

In the Standard Model this handedness is written in by hand as a parity-violating interaction of the weak force, with no mechanical reason why nature should prefer one chirality. In SIE the handedness is a readout of the substrate's rifling: a featureless 1-bit data packet fired through a chiral lattice is geometrically forced to spin left, in the same way a bullet fired from a rifled barrel is geometrically forced to spin in a specific direction. The neutrino's observed chirality is not an input to the theory but a consequence of the substrate's topological texture.

This reframing makes the chiral halving chain of Section VI (Problems 1–3) observationally tangible: the substrate's rifling is something we can detect every time a neutrino is measured.

11.2. Antimatter as time-reversed execution

Wheeler's "one-electron universe" hypothesis [79] — that the apparent multitude of electrons in the universe is actually a single electron bouncing back and forth through time — has long been treated as a playful interpretation rather than a physical claim. Feynman used the idea operationally: in Feynman diagrams, antiparticles can be represented as particles propagating backward in time. The mathematical equivalence is real; the question is whether it is physically meaningful.

In SIE the equivalence is physical. An execution thread running forward through the substrate encodes matter; the same thread running backward encodes antimatter. Charge conjugation becomes a temporal reversal of the substrate's update rule. CPT invariance is preserved because the full operation (charge + parity + time) returns the system to itself, but the partial operation of time reversal alone flips matter to antimatter.

11.3. Baryon asymmetry: the preference of forward execution

This framing delivers an immediate dividend for the matter–antimatter asymmetry problem. Standard physics confronts the observation that the early universe, despite initial CP symmetry, generated about $10^9 + 1$ matter particles for every 10^9 antimatter particles, leaving the small residue of matter that forms everything we see. Sakharov's three conditions [80] prescribe what is *needed* for this asymmetry — baryon-number violation, C and CP violation, departure from thermal equilibrium — but standard mechanisms (electroweak baryogenesis, leptogenesis, affleck-Dine) each require specific additions to the Standard Model and produce the observed asymmetry only with fine-tuning.

The SIE account is structural rather than dynamical. If the substrate has a chiral grain, forward execution encounters lower frictional resistance than backward execution, because backward execution runs against the grain. Matter is the path of lower computational cost; antimatter is the path of higher cost. In thermodynamic equilibrium the two populations would be equal, but in a universe with net entropy production and directional time, the path of lower cost accumulates preferentially. The roughly 1-part-in-10^9 excess of matter is the equilibrium deviation at the relevant temperature during baryogenesis, set by the ratio of the substrate's forward-backward frictional asymmetry to the ambient thermal energy at that epoch.

Sakharov's three conditions are *consistent with* structural features of SIE: baryon-number violation through substrate-mediated process, C/CP violation through the chiral grain, and departure from equilibrium through the Gompertz expansion out of the primordial crystal (Section XII). The framework provides a candidate mechanism for the asymmetry; whether this mechanism actually delivers the observed $\sim 10^{-9}$ magnitude requires computation of the substrate's chiral friction coefficient that is not yet available. In its current form, SIE offers a structural narrative compatible with the observation, not a quantitative derivation of it.

A quantitative prediction of the specific asymmetry magnitude ($\sim 10^{-9}$) requires the substrate's chiral friction coefficient, which is not currently derived from first principles. This parallels the c_{FCC} gap of Section VI.A — a specific geometric number requiring explicit lattice gauge theory computation. Structurally the framework accommodates the asymmetry; quantitatively it cannot yet predict the precise value. Claims elsewhere in the corpus that SIE "resolves" the baryon asymmetry should be read as a structural-story claim pending this quantitative derivation.

11.4. The black hole information paradox

Hawking's 1975 calculation showed that black holes emit thermal radiation and eventually evaporate completely [81]. If the radiation is genuinely thermal (random, uncorrelated, determined only by temperature), then information that fell into the black hole — the quantum state of infalling matter — is lost when the black hole disappears. This contradicts unitarity, the foundational principle that quantum evolution preserves information.

For fifty years the paradox has resisted resolution. Various proposals — black hole complementarity, firewalls, soft hair, ER=EPR, islands and Page curves from entanglement wedge reconstruction — each make partial progress but require specific additional structure. The question remains contested.

In SIE, the ontology is consistent with unitary evolution through the evaporation process, and naturally accommodates the Page curve. The argument has three steps.

Step 1: The interior is not singular. The collapsing matter in a black hole interior does not reach a geometric point. In SIE it reaches the substrate's saturation limit — all Planck-node bits within the horizon are occupied. This *Planck core* is finite, not singular, and its boundary coincides closely with the Bekenstein-Hawking horizon in the macroscopic limit [45, 9]. For a 30 $M_\odot$ black hole, the core radius is approximately 2×10^{-22} m.

Step 2: Hawking radiation is information erasure, not information destruction. A Hawking photon at temperature $T_H = \hbar c^3/(8\pi G M k_B)$ carries energy $\hbar\omega$ set by that temperature. Applying A3, the Landauer cost of a single bit erasure at T_H is $E_{\mathrm{bit}} = k_B T_H \ln 2$. The ratio of photon energy to per-bit erasure cost is:

$$\frac{\hbar\omega}{k_B T_H \ln 2} \approx 2.77, \tag{38}$$

using $\hbar\omega \approx 2.8 k_B T_H$ for a thermal distribution peak. We emphasize that this ratio is arithmetic from standard thermodynamics (both numerator and denominator are set by T_H), not a framework-specific prediction; it simply expresses that each Hawking photon is thermodynamically equivalent to $\sim$ 2.77 bits of Landauer-scale erasure. Any unitarity-preserving scheme with this temperature structure would give the same ratio.

Step 3: The Page curve from bulk-boundary entanglement. The horizon pixels erased into each photon are entangled with the remaining bulk of the Planck core. As long as fewer than half the pixels have been decompressed, the outgoing Hawking quanta appear thermal to an external observer because the information is encrypted in correlations with the still-trapped bulk. Once half the entropy has been released — the "Page time" [82] — the external observer has accumulated sufficient correlation to reconstruct the information in principle. We note that the Page curve is a consequence of *any* unitarity-preserving evolution through a bounded Hilbert space; SIE does not derive the Page curve from first principles but rather provides a physical mechanism (the finite Planck core) by which unitarity is preserved.

What SIE contributes and doesn't. SIE's specific contribution to the information paradox is ontological rather than computational: it supplies a concrete mechanism (finite substrate capacity $\to$ Planck core $\to$ bulk-boundary entanglement) by which information is not destroyed during evaporation. The Page curve, the 2.77 bits-per-photon ratio, and the unitarity-preservation argument are all consequences that would follow from any scheme that replaces the singular interior with a finite information-preserving core; they are not framework-specific predictions. We present this as a structurally consistent ontology for the problem, not as a resolution the existing literature was missing.

Information paradox status in SIE: The framework's ontology (finite Planck core, bulk-boundary entanglement, Landauer-bounded radiation) is consistent with unitarity and accommodates the Page curve naturally. This is a structurally complete story; whether it is preferable to existing island/replica proposals is a separate question about mechanism plausibility. Empirically: the finite Planck core predicts gravitational wave echoes (Section IX.C) at 89.1 ms for GW150914-class remnants, falsifiable in LIGO O4/O5 — this *is* a framework-specific observable and the strongest empirical handle.

11.5. Structural commonality

Both resolutions share a common structure: a physical phenomenon that looks paradoxical under the continuous-manifold assumption becomes a straightforward consequence of substrate thermodynamics once A1–A3 are adopted. The matter-antimatter asymmetry is a directional effect on a chiral substrate; the information paradox is a holographic decompression on a finite-capacity substrate. In both cases the paradox was an artifact of the wrong ontology, and vanishes once the right ontology is substituted.

This is perhaps the deepest meta-observation about the SIE framework: its primary value is not that it adds new physics, but that it removes puzzles that were never physical to begin with. They were mathematical artifacts of assuming a continuous vacuum with unbounded information capacity. A finite, discrete, chiral substrate has no room for them.

12. Observational Content V: Precision Constants and Anomalies

12.1. The fine-structure constant as information bottleneck

One of physics's most enduring puzzles is why the fine-structure constant $\alpha^{-1} = 137.036$ takes the specific value it does. Standard physics offers no explanation; Feynman famously observed that no one has been able to derive it from first principles.

SIE offers a reinterpretation of α rather than a first-principles derivation of the value 1/137. The Margolus-Levitin theorem [66] sets the absolute maximum quantum transition rate at $\Gamma_{\max} = 2m_e c^2/h$. A stable electron operates at its Compton frequency $\nu_C = m_e c^2/h$, which is identically 50% of $\Gamma_{\max}$ (this is a kinematic identity, not a physical claim about substrate coherence).

In a computational universe, a dimensionless coupling constant can be interpreted as a fraction: what portion of the electron's internal processing cycles is transmitted to the vacuum as EM coupling rather than retained for maintaining the particle's internal structure? Under this interpretation, α is named as that fraction, and $\alpha^{-1} \approx 137$ is the number of internal cycles per emission event. This gives a *physical picture* for what α represents on the substrate; it does not predict the specific value 137.036.

Honest framing. A true derivation would show that some structural property of the FCC substrate — gauge-node connectivity, coordination number, topological winding, or similar — forces α^{-1} to take this specific value. That derivation is not currently available. What SIE provides is a coherent substrate-theoretic interpretation of the dimensionless number, placing it in the same ontological category as the Landauer Mass Tensor (a dimensionless-cycle cost rather than a mystery constant). The empirical value 137.036 is *input* to the framework, as it is to the Standard Model; the framework's contribution is giving it a physically motivated reinterpretation rather than a derivation.

Given α^{-1} as input, it is then consistent with — and determines — the electron radius (Section III.C) and the neutrino mass (Section XI.C below) via substrate-thermodynamic relations that are genuinely predictive once α^{-1} is fixed. The logical order matters: α^{-1} input, r_e and m_ν outputs. An earlier presentation of this material in the corpus sometimes obscured the ordering, which (if read literally) suggested a circular derivation. The framework's commitment is to the order just stated.

12.2. The neutrino as minimum data packet

If the electron is 137 bits of stable charge encoding and the proton is 166 bits of FCC-packed confinement, what is the neutrino? In SIE, the neutrino is the absolute minimum possible excitation of the substrate: exactly one bit of chiral data propagating through the lattice. This identification leads to a quantitative prediction of the neutrino mass with no fitted parameters [9].

The derivation. By A3 (Landauer bounds at local temperature), the minimum energy cost of processing one bit against the ambient thermal noise floor at the CMB temperature is:

$$E_{\text{bit}} = k_B T_{\text{CMB}} \ln 2 \approx 1.627 \times 10^{-4}\ \text{eV}, \qquad (39)$$

using $T_{\text{CMB}} = 2.725$ K. But the neutrino is not a static thermal fluctuation — it is a structured execution thread propagating through the substrate against the lattice's electromagnetic impedance. That impedance is exactly the fine-structure constant's inverse, $\alpha^{-1} \approx 137.036$, which we have identified as the substrate's Shannon capacity for encoding a charge against vacuum polarization (Section XI.A). The neutrino's effective mass is the single-bit thermal cost amplified by this structural barrier:

$$m_\nu = E_{\text{bit}} \cdot \alpha^{-1} \approx 0.0223 \text{ eV}. \tag{40}$$

Comparison with observation. Laboratory bounds and cosmological constraints place the sum of neutrino masses in the range 0.06–0.12 eV (KATRIN upper bound $\sim$ 0.45 eV per flavor; Planck+BAO $\sum m_\nu \lesssim 0.12$ eV for three species) [86, 87]. Distributing 0.0223 eV across three mass eigenstates gives $\sum m_\nu \approx 0.067$ eV, landing squarely inside the cosmologically allowed band and consistent with the lower end of the direct-detection upper limits.

Neutrino mass from substrate thermodynamics: $m_\nu = k_B T_{\text{CMB}} \ln 2 \cdot \alpha^{-1} \approx 0.0223$ eV. This uses α^{-1} as an input (it is a measured quantity reinterpreted rather than derived by SIE) plus T_{CMB} and the Landauer bound. Given the two constants as input, the neutrino mass emerges with no further adjustable parameter, and sits within the observationally allowed mass range. This is genuinely predictive: once the Landauer-times-α^{-1} relation is adopted, no freedom remains to adjust m_ν.

The neutrino as substrate thermostat. Because the neutrino mass is tied to T_{CMB}, it is not a fundamental constant — it tracks the temperature of the universe. In the hot early universe, neutrinos were substantially heavier; as the universe expanded and cooled, their mass decreased. This has concrete implications for structure formation: the heavier early neutrinos may have clustered like cold dark matter during the first galaxy formation epoch, while today's lighter neutrinos free-stream through large-scale structure. The transition is smooth and occurs at the epoch when T_{CMB} dropped below the threshold where $\alpha^{-1} k_B T \ln 2$ fell below the kinetic energy of the neutrinos. This reinterprets the neutrino as the substrate's *active thermostat*: a dynamical mass-energy reservoir that equilibrates as the universe cools, rather than a species with a fixed fundamental mass.

If confirmed by cosmological neutrino mass measurements across different redshifts (technology yet to be developed but in principle accessible to very-long-baseline CMB spectral analyses), the mass-as-thermostat prediction would be a striking framework-specific signature.

12.3. The neutron lifetime anomaly

The 4.1σ discrepancy between beam and bottle measurements of the free neutron lifetime — $\tau_n \approx 888.0 \pm 2.0$ s (beam, from decay product counting) versus $\tau_n \approx 879.4 \pm 0.6$ s (bottle, from UCN trap retention) [68] — remains unresolved. Standard proposals (dark decay channels, CKM non-unitarity extensions) lack empirical support.

SIE resolves the anomaly through the *Infodynamic Shear Tensor* $\Upsilon_{\mu\nu}$ [21, 67]. The core mechanism: in bottle experiments, ultracold neutrons interact with strong magnetic gradients at the trap boundary. The boundary collision induces a localized geometric compression of the spatial lattice, shifting the axial-vector coupling g_A upward via the Gell-Mann–Oakes–Renner (GMOR) relation and accelerating thermal state-rejection.

Quantitatively, the neutron's 29.214-bit topological formatting conflict (the discrepancy between the proton's 166.25-bit FCC capacity and the electron boundary's $\alpha^{-1} = 137.036$-bit encoding) drives intrinsic instability at 1.29 MeV via the gluon trace anomaly. In bottle experiments, the macroscopic magnetic gradient couples to this instability via a dimension-6 operator, producing a 0.266 keV localized phase strain during each bounce that biases the measurement downward by ~ 10 s relative to beam experiments. The discrepancy direction (bottle $<$ beam) matches the sign predicted by the mechanism.

Falsifiable prediction: A bottle experiment using a magnetic wiggler with longitudinal adiabatic extraction — designed to suppress the non-adiabatic boundary interaction — should recover the beam lifetime at $\tau_n \approx 888$ s. Null result: bottle and beam agree at the quoted uncertainty, falsifying the mechanism. This is a specific, performable test [67].

12.4. Quark masses from dual-clock truncation

The six quark masses span four orders of magnitude (from $m_u \sim 2$ MeV to $m_t \sim 173$ GeV) with no derivation in the Standard Model. SIE addresses them via a dual-clock mechanism [21]: each quark flavor sits at a specific substrate shell with characteristic refresh cadence, and the observed mass is set by the ratio of the flavor's internal refresh clock to the substrate's background clock.

The precise values require specification of the shell structure (which overlaps with the three-generations issue discussed in Section IX), so the current predictions are 10–30% accurate rather than $\sim$ 1%. A unified treatment would require the icosahedral shell selection rule to be derived, linking this problem to Gap 3.

12.5. Magnetar polarization cap

As noted above, the IXPE observation of the extreme magnetar 4U 0142+61 showed X-ray polarization capping at $\sim$ 35%, well below the standard QED prediction of 80–100%. For neutron stars at lower fields (e.g., RX J1856.5−3754 at $B \sim 10^{13}$ G), SIE predicts linear optical polarization in the range 11–16% from the discrete-lattice birefringence formula

$$\Delta n = \frac{\alpha}{30\pi}\left(\frac{B}{B_c}\right)^2 \tag{41}$$

where $B_c = 4.41 \times 10^{13}$ G is the Schwinger critical field. The VLT observation of 16.43% polarization in RX J1856.5−3754 [69] is consistent with this prediction. For magnetars above the Schwinger limit, the 35% substrate yield cap takes over, as confirmed by IXPE [46].

Precision constants summary: $\alpha^{-1} = 137.036$ is taken as empirical input but receives a substrate-theoretic reinterpretation (internal cycles per emission event); deriving the specific value remains open. The 4.1σ neutron lifetime anomaly resolves through boundary magnetic shear coupling to the 29.214-bit formatting conflict. Magnetar polarization cap at 35% matches the substrate yield point. Quark masses via dual-clock are 10–30%, awaiting shell derivation.

13. The Cosmological Lifecycle

The framework assembles into a complete cosmological arc: a specific beginning, a specific dynamical history, and a specific terminal state. Each phase makes falsifiable predictions.

13.1. Genesis: Fracture of the Substrate Crystal

Before expansion, the universe existed as a maximally packed zero-entropy Substrate Crystal — an FCC relational graph in thermodynamic ground state. The "Big Bang" is not a point singularity but a macroscopic topological phase transition: the substrate crystal fractured along its $\{111\}$ cleavage planes, and the propagating crack front ignited the temporal metric. Time emerges as the sequential ordering of thermodynamic state changes on the expanding relational graph [34, 5].

This formulation bypasses the $t = 0$ singularity that breaks general relativity. It also makes a concrete prediction: the residual topological scarring of that primordial fracture should be visible in the CMB's low-ℓ multipoles, which is what Section VII (primordial crystallography) confirms at $p < 0.001$ significance.

13.2. Evolution: Gompertz expansion

The post-fracture expansion is not driven by a static cosmological constant. It is driven by the entanglement entropy generated from the decay of unstable primordial isotopes: as unstable particles (hyperons, resonances, and later actinides) decay, their localized informational geometry unspools onto the holographic horizon, and the substrate must allocate buffer space to accommodate the resulting entanglement entropy [5].

Mathematically, the decay-driven expansion yields a Gompertz function for the scale factor:

$$a(t) = a_{\max} \exp\left(-\frac{2H_{\max}}{\lambda_{\text{eff}}}\, e^{-\lambda_{\text{eff}} t/2}\right), \tag{42}$$

where λ_{eff} is an epochal decay envelope transitioning through the primordial unstable inventory. The Gompertz form gives a finite $a(0) = a_{\max}\exp(-2H_{\max}/\lambda_{\text{eff}})$ — the initial substrate crystal's bounded volume — avoiding a point singularity.

The equation of state evolves as

$$w(a) = -1 + \frac{Q(a)}{3H(\rho_{\text{DE}} + P_{\text{DE}})}, \tag{43}$$

where Q is the topological heat injection rate. At present, $Q > 0$ gives quintessence with $w > -1$; as primordial reservoirs deplete, Q transitions through zero and the equation of state asymptotes to $w \to -1$ as $a \to \infty$.

DESI 2024 Year 1 BAO [70] favors thawing dark energy with $w_0 \approx -0.73 \pm 0.07$ and $w_a \approx -0.88 \pm 0.36$ — an anomaly that ΛCDM struggles to accommodate. The Gompertz scaling naturally produces $w_0 \approx -0.72$ and $w_a \approx -0.31$ at the current epoch, placing the SIE prediction within the 1σ DESI contour on w_0 and within 2σ on w_a. Critically, SIE also predicts $w \to -1$ exactly as $a \to \infty$, distinguishing it from models with a true phantom crossing. Sub-percent constraints on $w(a)$ from Euclid, Rubin LSST, and DESI Year 5 will sharply test this trajectory.

13.3. Terminal state: the Time Crystal

As the universe exhausts its inventory of primordial unstable isotopes, the rate of new entanglement entropy generation drops: $dS_{\rm ent}/dt \to 0$. The expansion asymptotes, not to a Big Rip, but to a static Symbiotic Infodynamic Equilibrium.

Tellurium-128 undergoes two-neutrino double beta decay with half-life $\tau_{1/2} \approx 2.25 \times 10^{24}$ years [71], which SIE identifies as the empirical proxy for the Terminal Isotope Epoch. At $t \sim 10^{25}$ years (assuming proton stability at $\tau_p > 10^{34}$ years; if the proton ultimately decays, the terminal epoch is correspondingly later):

$$\lim_{t\to t_\tau} \rho_{\rm DE}(t) = \rho_{\rm min} \quad \Rightarrow \quad \lim_{t\to t_\tau} H(t) = H_{\rm min}. \tag{44}$$

The universe locks permanently into a stable, thermodynamically saturated informational architecture — a "Time Crystal" in the condensed-matter sense [72] — with residual kinetic energies conserved in the global topological binding energy of the entanglement graph.

This is a Kosterlitz-Thouless-like topological phase transition, not a standard thermal one. No macroscopic latent heat is expelled; the universe simply crosses into its terminal state. The prediction: no Big Rip, no Heat Death, no eternal de Sitter. A specific asymptotic structure with specific mathematical form.

The SIE cosmological arc has three phases: (1) substrate crystal fracture replacing the $t = 0$ singularity; (2) decay-driven Gompertz expansion replacing static Λ and producing thawing dark energy; (3) Time Crystal terminal state at $\sim 10^{25}$ years replacing eternal de Sitter. Each phase has empirical handles: CMB anomalies (Section VII), DESI $w(a)$ trajectory, and future sub-percent cosmological constraints.

14. Summary of Predictions

Table 1 consolidates the framework's quantitative predictions against observation. Where the prediction involves a fit parameter, it is noted explicitly; the majority of entries are zero-parameter predictions from substrate-theoretic ingredients.

The table is selective, not exhaustive. It emphasizes predictions with numerical content. Structural/qualitative claims (e.g., substrate-amplification mechanism applied to the Bullet Cluster geometry, which is qualitatively consistent with observation but awaits quantitative verification; see Section X.B) are covered only in the narrative and are not listed as successes here.

15. Honest Assessment of Remaining Gaps

No framework is complete. We catalog here the four identifiable technical gaps and their research profiles, with estimated difficulty:

15.1. Gap 1: $c_{\rm FCC}$ rigorous value

The coefficient $c_{\rm FCC}$ that converts $\theta_{\rm sub}$ into the chirally-asymmetric Wilson term is currently constrained only by dimensional analysis to $\mathcal{O}(0.3$–$1)$. A rigorous value requires lattice gauge theory computation with explicit FCC Dirac operator construction, Dirac spinor structure, and spectral flow under varying $\theta_{\rm sub}$. Python-level estimates span three orders of magnitude because they are dimensional approximations rather than true calculations.

Research task: Dedicated LGT simulation on FCC substrate. Time: weeks to months of specialist work. This is the single most pressing remaining calculation for framework credibility.

15.2. Gap 2: Pati-Salam Higgs potential

The embedding protection factor $\alpha_{\rm emb}$ of Problem 2 depends on the specific Higgs potential of the Pati–Salam breaking. Dimensional estimates place $\alpha_{\rm emb}$ between 10^{-26} (if scaling goes as $(v/M_{PS})^2$) and 10^{-13} (if scaling is linear in v/M_{PS}). Without the explicit potential, we cannot distinguish. A plausible outcome is that the embedding is *individually sufficient* for strong-CP consistency, making the axion redundant insurance rather than a necessary second layer.

Research task: Explicit Pati-Salam Higgs potential construction with minimization. Time: GUT model-building project scale.

15.3. Gap 3: Icosahedral shell selection rule

Lepton generations are assigned to icosahedral substrate shell indices $K = 0, 2, 18$ to match observed mass ratios to $\sim 1\%$. The principle selecting these specific shells (rather than other values of K) is not derived; it is fit to data. A first-principles derivation would show that only shells $K = 0, 2, 18$ support stable fermion states under Landauer-thermodynamic stability, converting the mass ratio match from phenomenology to consequence.

Research task: Stability analysis of icosahedral shell states under substrate dynamics. Time: moderate research project.

Table 1: SIE predictions against observation. ZFP: zero free parameters (no substrate-specific tuning beyond established constants). Cal: calibrated (empirical input used). * indicates empirical consistency rather than first-principles derivation. "∼": order-of-magnitude prediction.

Quantity	**SIE value**	**Observation**	**Accuracy**	**Status**
CMB spectral tilt n_s	$1 - 2/57 = 0.96491$	0.9649 ± 0.0042	10^{-4} match	Numeric
CMB bisector (Axis of Evil)	$7.85°$	$8.85° \pm 2°$	$1.0°$	ZFP (FC
FCC vs. randomized skies	$p_{\rm raw} < 0.001$	trials-adj. ~ 0.01–0.1	n/a	Suggesti
Kinematic dipole alignment	$0.12°$	serendipitous	n/a	Byprodu
Hubble local rate $H_{\rm local}$	73.1 km/s/Mpc	73.0 ± 1.0	aggregate match	Cal from
Hubble kinematic rate $H_0^{\rm kin}$	73.97 km/s/Mpc	n/a	1.3% of SH0ES	From SP
Growth index γ	0.545	0.55 (fit)	$\sim 1\%$	ZFP fron
Bare vacuum density $\rho_{\rm bare}$	2.07×10^{-27} kg/m^3	$\rho_c = 8.53 \times 10^{-27}$	factor 4 of ρ_c	CKN+La
Acceleration scale a_0	$cH_0/2\pi$	1.2×10^{-10} m/s^2	matches Planck H_0	ZFP der
SPARC rotation curves	SIE $\chi^2_\nu = 3.33$	NFW $\chi^2_\nu = 1.72$	competitive	1 param,
DESI w_0	-0.72	-0.73 ± 0.07	1σ	From Go
DESI w_a	-0.31	-0.88 ± 0.36	2σ	From Go
Magnetar polarization cap	$\leq 35\%$ (universal if true)	$\sim 35\%$ (4U 0142+61)	matches	Substrat
Neutron star polarization (J1856)	11–16%	16.43%	matches	Formula
Proton mass m_p	166.25-bit refresh	938.272 MeV	$\sim 3\%$	Multi-fac
Mass ratio m_p/m_e	from volumetric bits	1836.15	$\sim 0.4\%$	From La
Lepton mass ratios	from icosahedral shells	$m_e : m_\mu : m_\tau$	$\sim 1\%$	Cal: $N_{\rm co}$
Transmon decoherence feature	~ 5 GHz region	observed IBM Qiskit	design-band match	Qualitati
GW150914 ringdown	$c/2\pi R_s = 261$ Hz	251 Hz $(m = 2, 2, 0)$	4%	Cross-mo
GW150914 echo delay	89.1 ms	null to date	awaiting O4/O5	Hard tar
α^{-1}	(input)	137.036	—	Reinterp
Neutrino mass m_ν	0.0223 eV	$\sum m_\nu \lesssim 0.12$ eV	within bound	From La
Neutron τ_n beam-bottle gap	~ 10 s (bottle lower)	$888 - 879 = 9$ s	correct sign/magnitude	From she
MAQRO collapse threshold	$\pi \ell_P (m_P/M)^3$	not yet tested	awaiting MAQRO	Hard pre
Terminal epoch t_τ	$\sim 10^{25}$ years	not testable	Te-128 proxy*	Asympto

15.4. Gap 4: Unified Yukawa matrix

The full Yukawa matrix Y_{ij} is not derived. SIE has individual mass predictions (proton, m_p/m_e, neutrino, lepton generations, rough quark ratios) but no unified theory of flavor. This gap is shared with all GUT frameworks and may not be uniquely tractable in SIE.

Research task: Either (a) derive Yukawa as overlap integrals of fermion wave functions on icosahedral shells with Higgs VEV profile, or (b) deepen individual mass predictions toward a comprehensive flavor account. Time: open-ended.

15.5. The Pantheon+ void anomaly

The sole empirically *negative* result in the current corpus: the SIE prediction of a step function in supernova H_0 at the KBC void edge is not observed, and a void model fit yields an inverted hierarchy ($H_{\rm in} < H_{\rm out}$). This is recorded transparently [22] as the primary open problem connecting the observational program to the theoretical framework. It is not a structural inconsistency — the Hubble tension is still resolved by Eq. 30 at the aggregate level — but it is an empirical discrepancy that demands resolution.

Candor: Every published framework with broad scope has failing predictions. The question is whether the failures are fixable within the framework or signal fundamental incorrectness. SIE's failures (the Pantheon+ void, the residual factor in $\rho_c/\rho_{\rm bare}$, the shell-selection tuning) are consistent with incomplete derivation rather than wrong-framework. This is a testable interpretation: if the research tasks above are pursued and the failures persist, that will sharpen into real evidence against the framework. If the research tasks resolve them, it will sharpen into support.

16. Falsifiability Summary

A framework that fits everything predicts nothing. SIE makes multiple falsifiable commitments. A selection:

16.1. Near-term (current-generation detectors)

IAXO helioscope axion. SIE predicts an emergent QCD axion at $m_a \approx 3.5$ meV and $f_a \sim 10^{9.5}$ GeV. IAXO is sensitive to this range [17]. A confirmed detection at these parameters supports the framework; a null result at significantly tighter bounds constrains (though does not kill, given the embedding protection alternative).

LIGO O4/O5 gravitational wave echoes. SIE predicts an echo delay of 89.1 ms for GW150914-class remnants, with logarithmic mass slope 2/3 rather than the standard-ECO 1.0 [45]. A confirmed detection at the predicted delay/slope supports the framework; a null result with sufficient sensitivity across multiple events constrains it.

DESI high-redshift equation of state. SIE predicts a non-linear divergence of $w(a)$ at $z > 2$, plunging into the phantom regime ($w \lesssim -1.5$) due to the Landauer thermal erasure penalty. Linear CPL parameterizations that remain shallow at high z would falsify this feature.

LiteBIRD tensor-to-scalar ratio. SIE predicts $r \ll 0.14$ (specifically, $r \approx 0$ in the osmotic-inflation picture), decoupled from $n_s = 0.96491$. Any clear detection of $r \sim 0.14$ would falsify the framework. A tight upper bound on r without n_s drift supports it.

LHAASO photon dispersion. SIE predicts a linear dispersion delay $\Delta t \propto E$ for high-energy extragalactic photons from discrete substrate propagation. Recent observations of TeV photons from GRB 221009A [52] provide the beginning of constraints; continued observations tighten them.

16.2. Medium-term (next-generation detectors)

Cosmic Explorer / Einstein Telescope ringdown spectroscopy. SIE's $m = 0$ axisymmetric mode prediction differs from GR's $m = 0$ mode by 12.7% at fixed mass. High-SNR ringdown spectroscopy across a mass-spin distribution of events would distinguish them [41].

MAQRO levitated-nanosphere interferometry. SIE predicts wave-function collapse at spatial separation $\Delta x_{\max} = \pi \ell_P (m_P/M)^3$ with material-invariant collapse location (contrasting the Diosi–Penrose model's 22× spread across materials) and Landauer thermal signature $Q \sim 5$ μJ per collapse for 100-nm spheres [47]. This is a direct quantum test at meter scale.

Simons Observatory / CMB-S4 B-mode polarization. SIE predicts six-fold octahedral symmetry in CMB B-mode polarization cross-correlations from the FCC primordial substrate fracture geometry [10]. A clear detection of this structure supports the framework; a strong null constrains it.

16.3. Structural predictions

Cross-scale substrate signatures. The framework hypothesizes that a single substrate equation of state governs phenomena at very different length scales. Quantitative tests of this hypothesis require mapping between regimes (galactic rotation, magnetar polarization, qubit decoherence, cosmic photon dispersion) through explicit dimensional analyses. These mappings are at varying stages of development; claims of "scale invariance across orders of magnitude" should be understood as a program statement rather than a completed result.

The Pantheon+ void-gradient test. SIE predicts a $\Delta H \approx 5.7$ km s^{-1} Mpc^{-1} gradient across the KBC void edge from $\delta_b \approx 0.15$. A fresh reanalysis of Pantheon+ with the full stat+sys covariance and floated Ω_m (Section VII.C) finds this predicted magnitude ruled out under the standard SH0ES-anchored analysis: the fitted $\Delta H = 0.24$ km s^{-1} Mpc^{-1} from the step-function model and $\delta_0 = -0.012$ from the smooth-profile model are inconsistent with the predicted $\Delta H \approx 5.7$ and $\delta_0 = +0.15$ at $\sim 20\sigma$ under standard analysis. The framework's one viable defense is scenario-based: if the SH0ES Cepheid anchors inside the void cause the anchoring pipeline to bake in a void-interior H_0 value, the gradient could be concealed from direct observation. Demonstrating this scenario reproduces the observed flat-H_0 signature across $0.01 < z < 2.3$ requires an explicit pipeline-level calculation that has not yet been performed and constitutes the primary open empirical task for the framework. A self-consistency concern also attaches to the defense scenario: if the global $H_0 = 67.4$ and SPARC measures a local-observer a_0, then the derivation $a_0 = cH_0/2\pi$ needs to specify whether the horizon-scale noise floor tracks the global or local H_0. Until the pipeline calculation is done, the aggregate $H_0 \approx 73$ match should be regarded as numerically consistent with the framework's interior-observer prediction but not as independent evidence of the mechanism. Independent tests — void-exterior distance-ladder anchoring from Rubin LSST, Roman, and Euclid (2027+), plus Cosmicflows-4 [95] flow-model reconstructions — will break the degeneracy.

SIE is not a framework that can accommodate any observation. Each of the falsifiable predictions above has a specific signature that, if contradicted by sufficient data, constrains or rejects the framework. The observational program is multi-pronged: cosmology (LiteBIRD, DESI, CMB-S4), particle physics (IAXO), gravitation (LIGO, Cosmic Explorer, MAQRO), and high-energy astrophysics (LHAASO, IXPE). A framework with this breadth of testability, with each test already partially informed by existing data, is in an uncommonly strong epistemic position.

17. Comparison to Alternative Programs

To contextualize SIE, we summarize briefly the relation to adjacent frameworks.

17.1. Versus ΛCDM with dark matter

Standard ΛCDM with cold dark matter halos fits CMB anisotropies and large-scale structure at high precision, but faces persistent anomalies: the Hubble tension (5σ), the S_8 tension [48], small-scale structure issues [49], and the unexplained numerical value of Λ itself. It requires a dark matter particle that has evaded detection for four decades.

SIE reproduces the large-scale successes of ΛCDM via the Gompertz expansion history and the $N = 57$ equipartition closure, while resolving the tensions through substrate thermodynamics. It requires no new particles. It does require a new ontology (substrate as discrete finite-capacity computational layer), which is a significant philosophical commitment.

17.2. Versus MOND

MOND [50] reproduces galaxy rotation curves with a single phenomenological acceleration $a_0 \approx 1.2 \times 10^{-10}\ \mathrm{m\,s^{-2}}$. Its successes on galactic scales are striking but it struggles with cluster-scale gravitational lensing and CMB acoustic peaks, and crucially, it provides no mechanism for why a_0 exists or takes its observed value.

SIE reproduces MOND's galactic successes through Eq. 21, but $a_0 = cH_0/2\pi$ is derived from substrate thermodynamics rather than postulated. For the Bullet Cluster separation [51] — which is challenging for MOND — SIE offers a qualitative mechanism via temperature-dependent substrate amplification (hot plasma breaks substrate tension amplification while cold stellar components preserve it) [5]; however, the quantitative lensing-offset magnitude has not yet been computed from the mechanism (Section X.B). The Bullet Cluster status within SIE is therefore a plausible structural story pending quantitative verification, not a completed resolution.

17.3. Versus emergent gravity programs

Verlinde's entropic gravity [12] and Jacobson's thermodynamic derivation of Einstein's equations [11] share SIE's commitment to gravity as thermodynamic. The distinction is that these programs derive gravitational dynamics from horizon thermodynamics at the final level, while SIE derives all five Lagrangian terms (gravity, matter, EM, topology, dark energy) from the same substrate axioms. SIE is more ambitious and correspondingly more constrainable.

17.4. Versus string theory and loop quantum gravity

These are programs for quantum gravity that seek to reproduce general relativity and the Standard Model from a more fundamental starting point. SIE shares this ambition but takes a different starting point: information dynamics on a discrete substrate rather than extended objects or spin networks. The discriminating question is empirical: which framework produces falsifiable predictions that match observation with fewer free parameters? SIE's scorecard to date (outlined in Sections VI–VIII) is publicly examinable.

18. The Research Program

The framework's path forward has specific, identifiable tasks:

18.1. Priority 1: LGT calculation of c_{FCC}

This is the single calculation that would convert the chiral halving chain from "structurally complete with $c_{\mathrm{FCC}} \sim \mathcal{O}(1)$" to "fully rigorous with $c_{\mathrm{FCC}} = X$ for a specific X." The required work is well-defined: construct the FCC Dirac operator as a matrix on a finite periodic lattice, introduce the chirally-asymmetric Wilson term with explicit θ_{sub} coefficient, diagonalize, and extract the slope $c_{\mathrm{FCC}} = \partial r_{\mathrm{asym}}/\partial\theta_{\mathrm{sub}}|_{\theta=0}$.

Recruitment of a lattice gauge theory specialist is the operational priority. An Anthropic-level compute budget plus specialist expertise plus 2-6 months would close this gap.

18.2. Priority 2: Pantheon+ resolution

The inverted void signal is the primary empirical discrepancy. Possible resolutions: (i) the void geometry in Pantheon+ is more complex than the simple sphere assumed, and a proper 3D reconstruction (e.g., using 2M++ [54] or BORG [55]) changes the inference; (ii) peculiar-velocity corrections in Pantheon+ absorb the signal; (iii) the SIE substrate tension law is refined. Work on (i)-(iii) is tractable without specialist recruitment.

18.3. Priority 3: Icosahedral shell derivation

A stability analysis showing that substrate icosahedral shells with $K = 0, 2, 18$ are the uniquely stable configurations would convert the lepton mass ratio match from fit to derivation. This is a specific calculation with a clear endpoint.

18.4. Priority 4: Axion constraint update

As IAXO ramps up, the SIE axion parameters ($m_a \approx 3.5$ meV, $f_a \sim 10^{9.5}$ GeV) should be continuously updated against the evolving constraint. If IAXO finds nothing in the predicted region, the framework falls back to the embedding-protection-alone position (Gap 2), which would motivate prioritizing the Higgs potential construction.

19. Conclusion

Symbiotic Infodynamic Equilibrium commits to a specific ontological claim — that information is primary, that the universe is a closed system, and that Landauer bounds are universal — and attempts to derive observable physics from those commitments. The derivation chain from FCC substrate to the Standard Model gauge structure and fermion chirality has substantive structural support, with four identifiable technical gaps representing tractable research tasks rather than fundamental problems.

The empirical scorecard is multi-pronged and quantitative. From the corpus: SIE fits the 175-galaxy SPARC sample with one parameter per galaxy, remaining competitive with cosmologically-constrained NFW (2 parameters; median χ^2_ν 3.33 vs 1.72) while making the independent zero-free-parameter prediction $a_0 = cH_0/2\pi$ that reproduces the 3382-point radial acceleration relation at 0.140 dex scatter; delivers a kinematic Hubble constant within 1.3% of SH0ES from galaxies alone; matches Planck's spectral tilt to 10^{-4} via the $N = 57$ chiral inventory (with the caveat that Route 2's mechanism in the osmotic picture is an open task); predicts the CMB Axis of Evil orientation within 1° with trials-adjusted significance $p \sim 0.01$–0.1; yields an aggregate numerical match $H_{\rm local} \approx 73.1$ km s^{-1} Mpc^{-1} consistent with the SH0ES measurement via $H_{\rm local} = H_{\rm global}/\sqrt{1-\delta_b}$ for interior-void observers; produces the bare vacuum density within a factor of four of observation with no fitting; predicts black hole ringdown frequencies within 4% for GW150914 (with the coincidence attributable to the Kerr $m = 0$ circumference identity); and derives the matter growth index $\gamma = 0.545$ from the closed-system boundary condition $w = -1$. The Pantheon+ void-gradient prediction is ruled out at high significance under the standard analysis; the framework's defense (SH0ES anchoring inside the void conceals the gradient) requires an explicit pipeline-level calculation not yet performed and is recorded as the priority open task.

Where SIE differs from other unification programs is its simultaneous breadth of observational engagement and specific falsifiability. LiteBIRD, DESI, IAXO, LIGO, Cosmic Explorer, CMB-S4, MAQRO, and LHAASO each deliver independent tests in their target windows. A framework this exposed to falsification, with this many partial observational confirmations already in hand, is not a framework that will remain comfortable.

The next milestones — a rigorous LGT value of $c_{\rm FCC}$, a Pantheon+ reanalysis, an icosahedral shell derivation, and the first IAXO axion sensitivity — will either sharpen the framework into quantitative theory or provide the specific contradictions that would push it aside. Either outcome advances physics. We aim for the first.

Acknowledgments

The author thanks the SPARC team [23], the Pantheon+ collaboration [53], the Planck collaboration [3], and the LIGO/Virgo collaboration [39] for making their data publicly available. This work has been developed through ongoing exchange with colleagues at the Information Physics Institute.

References

[1] S. Weinberg, Rev. Mod. Phys. **61**, 1 (1989).

[2] A. G. Riess *et al.*, Astrophys. J. Lett. **934**, L7 (2022).

[3] Planck Collaboration, Astron. Astrophys. **641**, A6 (2020).

[4] C. Hudock, "Resolving the Vacuum Catastrophe: Holographic Bit-Density and the Dual

Derivation of the Cosmological Constant," IPI Letters (2026).

[5] C. Hudock, "The Lifecycle of the Gompertz Universe: Addressing the CMB Axis, Hubble Tension, and Cosmic Fate via Osmotic Decompression," IPI (2026).

[6] J. D. Bekenstein, Phys. Rev. D **7**, 2333 (1973).

[7] R. Bousso, J. High Energy Phys. **07**, 004 (1999).

[8] L. Susskind, J. Math. Phys. **36**, 6377 (1995).

[9] C. Hudock, "Master Class: The Symbiotic Infodynamic Equilibrium Framework," IPI (2026).

[10] C. Hudock, "Lepton Generations from Icosahedral Shells," IPI internal (2026).

[11] T. Jacobson, Phys. Rev. Lett. **75**, 1260 (1995).

[12] E. Verlinde, J. High Energy Phys. **04**, 029 (2011).

[13] M. Born and L. Infeld, Proc. R. Soc. A **144**, 425 (1934).

[14] C. Hudock, "From $\theta_{\rm sub}$ to the Chirally-Asymmetric Wilson Coefficient," IPI internal (2026).

[15] C. Hudock, "Reconciling $\theta_{\rm sub} \sim \mathcal{O}(1)$ with $|\theta_{\rm QCD}| < 10^{-10}$," IPI internal (2026).

[16] C. Abel *et al.*, Phys. Rev. Lett. **124**, 081803 (2020).

[17] IAXO Collaboration, J. High Energy Phys. **05**, 010 (2019).

[18] C. Hudock, "Gauge Halving from Helicity Alignment," IPI internal (2026).

[19] C. Hudock, "FCC to the Standard Model via Pati-Salam Embedding," IPI internal (2026).

[20] C. Hudock, "Yukawa Compensation for Observed Dirac Masses," IPI internal (2026).

[21] C. Hudock, "Shear and the Dual-Clock Truncation of Quarks," IPI (2026).

[22] C. Hudock, "Galactic Kinematics as a Test of the Symbiotic Infodynamic Equilibrium Framework," MNRAS preprint, Zenodo DOI: 10.5281/zenodo.19503375 (2026).

[23] F. Lelli, S. S. McGaugh, J. M. Schombert, Astron. J. **152**, 157 (2016).

[24] J. F. Navarro, C. S. Frenk, S. D. M. White, Astrophys. J. **490**, 493 (1997).

[25] A. A. Dutton and A. V. Macciò, Mon. Not. R. Astron. Soc. **441**, 3359 (2014).

[26] C. Hudock, "The N=57 Infodynamic Framework and the Inflationary Spectral Tilt," IPI (2026).

[27] V. F. Mukhanov, H. A. Feldman, R. H. Brandenberger, Phys. Rep. **215**, 203 (1992).

[28] BICEP/Keck Collaboration and Planck Collaboration, Phys. Rev. Lett. **127**, 151301 (2021).

[29] LiteBIRD Collaboration, Prog. Theor. Exp. Phys. **2023**, 042F01 (2023).

[30] CMB-S4 Collaboration, arXiv:1907.04473 (2019).

[31] A. G. Cohen, D. B. Kaplan, A. E. Nelson, Phys. Rev. Lett. **82**, 4971 (1999).

[32] R. C. Keenan, A. J. Barger, L. L. Cowie, Astrophys. J. **775**, 62 (2013).

[33] L. Wang and P. J. Steinhardt, Astrophys. J. **508**, 483 (1998).

[34] C. Hudock, "Primordial Crystallography: Infodynamic Derivation of the CMB Spectral Tilt and the Deterministic Origin of the Axis of Evil," IPI (2026).

[35] F. Erdogan and G. C. Sih, J. Basic Eng. **85**, 519 (1963).

[36] K. Land and J. Magueijo, Phys. Rev. Lett. **95**, 071301 (2005).

[37] Planck Collaboration VII, Astron. Astrophys. **641**, A7 (2020).

[38] C. Hudock, "Empirical Unification of Galactic Kinematics and Quantum Hardware Decoherence via Substrate Thermodynamics," IPI (2026).

[39] B. P. Abbott *et al.*, Phys. Rev. Lett. **116**, 061102 (2016).

[40] E. Berti, V. Cardoso, A. O. Starinets, Class. Quantum Grav. **26**, 163001 (2009).

[41] C. Hudock, "Spin-Invariance of the $m = 0$ Axisymmetric Mode on Kerr Horizons," IPI internal (2026).

[42] Cosmic Explorer Consortium, arXiv:1907.04833 (2019).

[43] Einstein Telescope Science Team, arXiv:1912.02622 (2019).

[44] V. Cardoso, E. Franzin, P. Pani, Phys. Rev. Lett. **116**, 171101 (2016).

[45] C. Hudock, "Fractional Logarithmic Scaling of Gravitational Wave Echoes from Volumetric Planck Cores," IPI (2026).

[46] R. Taverna *et al.*, Science **378**, 646 (2022).

[47] C. Hudock, "Infodynamic Bandwidth Saturation and the Quantum Measurement Problem," IPI (2026).

[48] E. Di Valentino *et al.*, Astropart. Phys. **131**, 102604 (2021).

[49] J. S. Bullock and M. Boylan-Kolchin, Annu. Rev. Astron. Astrophys. **55**, 343 (2017).

[50] M. Milgrom, Astrophys. J. **270**, 365 (1983).

[51] D. Clowe *et al.*, Astrophys. J. Lett. **648**, L109 (2006).

[52] LHAASO Collaboration, Science **380**, 1390 (2023).

[53] D. Scolnic *et al.*, Astrophys. J. **938**, 113 (2022).

[54] G. Lavaux and M. J. Hudson, Mon. Not. R. Astron. Soc. **416**, 2840 (2011).

[55] J. Jasche and B. D. Wandelt, Mon. Not. R. Astron. Soc. **432**, 894 (2013).

[56] S. W. Hawking, Phys. Rev. D **14**, 2460 (1976).

[57] J. A. Wheeler, "Information, Physics, Quantum: The Search for Links," in *Complexity, Entropy, and the Physics of Information* (Addison-Wesley, 1990).

[58] D. J. Fixsen, Astrophys. J. **707**, 916 (2009).

[59] G. W. Gibbons and S. W. Hawking, Phys. Rev. D **15**, 2752 (1977).

[60] C. Hudock, "Resolving the Hubble Tension and Dark Matter Anomalies via Osmotic Spatial Decompression," IPI (2026).

[61] R. L. Workman *et al.* (Particle Data Group), Prog. Theor. Exp. Phys. **2022**, 083C01 (2022).

[62] R. Kaltenbaek *et al.*, EPJ Quantum Technology **3**, 5 (2016).

[63] J. Heyvaerts and P. Meszaros, Astrophys. J. **384**, 283 (1992).

[64] B. D. Fields and K. A. Olive, J. Phys. G **49**, 110501 (2022).

[65] K. Greisen, Phys. Rev. Lett. **16**, 748 (1966); G. T. Zatsepin and V. A. Kuzmin, Sov. Phys. JETP Lett. **4**, 78 (1966).

[66] N. Margolus and L. B. Levitin, Physica D **120**, 188 (1998).

[67] C. Hudock, "Resolving the Neutron Lifetime Anomaly via Discrete Lattice Boundary Impedance and the Dual-Clock Truncation Mechanism," IPI (2026).

[68] F. E. Wietfeldt, J. Res. NIST **110**, 521 (2005); updated values in J. S. Nico, J. Phys. G **36**, 104001 (2009); also see F. M. Gonzalez *et al.*, Phys. Rev. Lett. **127**, 162501 (2021).

[69] R. P. Mignani *et al.*, Mon. Not. R. Astron. Soc. **465**, 492 (2017).

[70] DESI Collaboration, arXiv:2404.03002 (2024).

[71] A. S. Barabash, Nucl. Phys. A **935**, 52 (2015).

[72] F. Wilczek, Phys. Rev. Lett. **109**, 160401 (2012).

[73] R. P. Feynman, *QED: The Strange Theory of Light and Matter* (Princeton University Press, 1985).

[74] L. D. Landau and I. Pomeranchuk, Dokl. Akad. Nauk SSSR **102**, 489 (1955).

[75] G. Amelino-Camelia, Int. J. Mod. Phys. D **11**, 35 (2002).

[76] D. Colladay and V. A. Kostelecky, Phys. Rev. D **58**, 116002 (1998).

[77] A. Connes, *Noncommutative Geometry* (Academic Press, 1994).

[78] B. Pontecorvo, Sov. Phys. JETP **26**, 984 (1968).

[79] J. A. Wheeler, quoted in R. P. Feynman's Nobel lecture (1965); original conjecture communicated privately 1940.

[80] A. D. Sakharov, JETP Lett. **5**, 24 (1967).

[81] S. W. Hawking, Commun. Math. Phys. **43**, 199 (1975).

[82] D. N. Page, Phys. Rev. Lett. **71**, 3743 (1993).

[83] C. E. Shannon, Proc. IRE **37**, 10 (1949).

[84] E. Eichten *et al.*, Phys. Rev. D **21**, 203 (1980).

[85] W. G. Unruh, Phys. Rev. D **14**, 870 (1976).

[86] KATRIN Collaboration, Nat. Phys. **18**, 160 (2022).

[87] Planck Collaboration, Astron. Astrophys. **641**, A6 (2020), Table 2.

[88] A. A. Abdo *et al.* (Fermi LAT Collaboration), Nature **462**, 331 (2009).

[89] G. Amelino-Camelia, Symmetry **2**, 230 (2010).

[90] R. K. Pathria, Nature **240**, 298 (1972).

[91] I. J. Good, Phys. Today **25**, 15 (1972).

[92] D. L. D. Caspar and A. Klug, Cold Spring Harbor Symp. Quant. Biol. **27**, 1 (1962).

[93] C. Müller, J. H. Cole, and J. Lisenfeld, Rep. Prog. Phys. **82**, 124501 (2019).

[94] R. Hagedorn, Nuovo Cimento Suppl. **3**, 147 (1965).

[95] R. B. Tully *et al.*, Astrophys. J. **944**, 94 (2023).

[96] E. Gross and O. Vitells, Eur. Phys. J. C **70**, 525 (2010).

[97] C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman, Adv. Astron. **2010**, 847541 (2010).

[98] D. J. Schwarz, C. J. Copi, D. Huterer, and G. D. Starkman, Class. Quantum Grav. **33**, 184001 (2016).

[99] G. W. Angus, H. Y. Shan, H. S. Zhao, and B. Famaey, Astrophys. J. Lett. **654**, L13 (2007).