

Holographic Λ CDM

Entropy Mechanics on the Cosmic Causal Diamond

Bryce Weiner

Information Physics Institute, Sibalom, Antique 5713, Philippines

Status. This manuscript combines peer-reviewed results with pre-print material and work under review. It may change as those portions exit peer review.

Abstract

Holographic Λ CDM (H- Λ CDM) takes the von Neumann algebra $\mathcal{A}(\mathcal{D})$ of observables on the cosmic causal diamond as the primitive object of physics, and the Lorentzian manifold, its metric, and its matter content as projections of the modular state on that algebra. This document states the framework once and completely: its algebraic foundation, its single postulate, the dynamics it generates, the boundary data it must accept, and the predictions it is committed to. From the Bisognano–Wichmann modular Hamiltonian, the Bekenstein–Hawking area law, the Margolus–Levitin bound, and one postulate — that the quantum of irreversibility is the obit, $S_{\text{obit}} = 1 \text{ nat}$ — there follow the holographic information-processing rate $\gamma = H/\ln S_H$, the Quantum-Thermodynamic Entropy Partition with its universal ratio $\eta = \ln 2/(1 - \ln 2) = 2.2589$, and a chain of parameter-free projections: the recombination sound horizon $r_d = 150.71 \text{ Mpc}$; the local-to-CMB Hubble projection giving $H_{\text{local}} = 69.15$ and $70.36 \text{ km s}^{-1} \text{ Mpc}^{-1}$ at the NGC 4258 and LMC anchors; the cosmological-constant density $\rho_\Lambda(0) = 5.84 \times 10^{-27} \text{ kg m}^{-3}$ with the 10^{-123} Planck hierarchy identified exactly, as $\frac{3}{2} \Omega_\Lambda/N_P$, with the inverse holographic capacity of the observable universe, and $w(z) = -1 + \mathcal{O}(10^{-11})$; the fine-structure constant $\alpha^{-1}(H_0) = 137.032$ and, by inversion of the same constraint, Newton’s constant $G = 6.624 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}$; the dark-to-baryon ratio $\Omega_{\text{dm}}/\Omega_{\text{b}} = q^2/\nu = 16/3$; Planckian initial conditions at the degenerate boundary; and a laboratory-scale Landauer crossover saturating at $k_B T$ with the ratio η . The framework is equally explicit about what it does not derive. Its dynamics are parity-even and therefore cannot originate chirality; the handedness of the matter sector, the generation count, and the gauge structure are boundary data, identified here with the $E_8 \times E_8$ heterotic apparatus and shown to be underivable from the framework’s dynamic data by the rigidity of the index. The line between what is derived and what is accepted falls exactly at the boundary of the diamond, and that division is itself the framework’s central structural claim.

Keywords: holographic principle; causal diamonds; modular theory; entropy mechanics; information thermodynamics; emergent gravity; cosmological tensions.

Part I

Foundations

1 The programme and its posture

Standard cosmology begins with a Lorentzian manifold and places quantum fields upon it. H- Λ CDM inverts the order. The primitive is the von Neumann algebra $\mathcal{A}(\mathcal{D})$ of quantum

observables supported on the cosmic causal diamond, together with the modular state it carries; the manifold, its metric, its matter content, and the notion of time itself are projections of that state, accurate wherever the projection’s corrections are small and observably inaccurate where they are not. The cosmological tensions of the last decade are, on this account, not failures of measurement but the residue where the projection is not exact.

Three commitments organize everything that follows.

Algebraic primacy. $\mathcal{A}(\mathcal{D})$ and its modular flow are foundational. The modular Hamiltonian K is not derived from a stress-energy tensor; the information stress-energy tensor $T_{\mu\nu}^{\text{info}}$ is determined by the modular state. This inverts the standard direction of quantum field theory and is what allows the framework to remain well-defined at a degenerate boundary where no metric exists (Section 10).

One postulate. The framework rests on a single undeduced quantity, the obit — the quantum of irreversibility, $S_{\text{obit}} = 1 \text{ nat}$ (Section 4). It stands on the model of $\hbar$ or $h\nu$: postulated, validated by its consequences, and now with a laboratory experiment designed to test it directly (Section 22).

The boundary–dynamics division. What the framework provides is dynamics: the modular flow, the rate at which it runs, the partition it enforces, and the geometric encoding of all three. What the framework must accept is boundary data: the initial state, the handedness of matter, the gauge structure, the generation count. The line between them is not a concession made under pressure; it is a structural theorem (Part III), and it falls exactly at the boundary of the diamond.

2 The cosmic causal diamond

Definition 2.1 (Cosmic causal diamond). The cosmic causal diamond is the spacetime region

$$\mathcal{D}(p, q) = \mathcal{V}(p, q) = J^+(p) \cap J^-(q), \quad (1)$$

the intersection of the causal future of the past tip p — the initial singularity — with the causal past of the future tip q — the cosmic event horizon. Its boundary $\partial\mathcal{D}$ consists of two null sheets meeting at a codimension-two surface $\sigma \cong S^2$ of maximal area: the holographic screen.

All measurements available to a comoving observer are confined to $\mathcal{V}(p, q)$, so the entropy and information budgets relevant to any cosmological question must be formulated within it. The Robertson–Walker metric is a valid Lorentzian realization of $\mathcal{D}$ and the one most useful for contact with standard ΛCDM , but nothing below depends on that choice: the predictions depend only on the modular structure of $\mathcal{A}(\mathcal{D})$.

For a diamond of proper-time separation τ between its tips, the Gibbons–Solodukhin small-diamond expansion gives the geometric measures in a curvature series. The four-volume is

$$V_4(p, q) = \frac{\pi}{24} \tau^4 \left[1 - \frac{\tau^2 R}{180} + \frac{\tau^2 R_{\mu\nu} T^\mu T^\nu}{30} + \dots \right], \quad (2)$$

the maximal spatial cross-section is

$$V_3(p, q) = \frac{\pi}{6} \tau^3 \left[1 - \frac{R\tau^2}{120} + \frac{R_{\mu\nu} T^\mu T^\nu}{40} + \dots \right], \quad (3)$$

and the area of the screen is

$$\mathcal{A}(p, q) = \pi \tau^2 \left[1 - \frac{R\tau^2}{72} + \dots \right], \quad (4)$$

with T^μ the unit timelike vector joining the tips. The series converge for $\tau \ll |R|^{-1/2}$. The directional term $R_{\mu\nu}T^\mu T^\nu$ records that the diamond’s capacity is not isotropic: local energy-momentum deforms the region in which information can be processed, and this is the seed from which the geometric encoding of Section 9 grows.

3 The algebra, its modular structure, and the modular Hamiltonian

3.1 Type III₁ and Tomita–Takesaki

Let $\mathcal{A}(\mathcal{D})$ denote the von Neumann algebra of observables supported on $\mathcal{D}$, and ω the vacuum state restricted to it. For a causal region in quantum field theory this algebra is a factor of type III₁: it admits no trace, no pure states, and no tensor factorization into “inside” and “outside” degrees of freedom. It does, however, admit something the type-I algebras of ordinary quantum mechanics do not — an intrinsic dynamics, generated by the state itself.

Tomita–Takesaki theory supplies it. Given the cyclic and separating vector $|\Omega\rangle$ representing ω , define the antilinear operator S by

$$S A |\Omega\rangle = A^\dagger |\Omega\rangle, \quad A \in \mathcal{A}(\mathcal{D}), \quad (5)$$

and take its polar decomposition

$$S = J \Delta^{1/2}, \quad (6)$$

with J antiunitary (the modular conjugation) and Δ positive (the modular operator). The Tomita–Takesaki theorem then states that the one-parameter family

$$\theta_t(A) = \Delta^{it} A \Delta^{-it} \quad (7)$$

is an automorphism group of $\mathcal{A}(\mathcal{D})$ — the *modular flow* — and that $J\mathcal{A}(\mathcal{D})J = \mathcal{A}(\mathcal{D})'$, the commutant. The generator

$$K \equiv -\ln \Delta \quad (8)$$

is the *modular Hamiltonian*. It exists for any local algebra, is determined by the state and the region alone, and requires no external clock.

The physical content of θ_t is thermal. The state ω satisfies the Kubo–Martin–Schwinger condition with respect to its own modular flow,

$$\omega(A \theta_{t-i\beta}(B)) = \omega(\theta_t(B) A), \quad \beta = 2\pi, \quad (9)$$

so that ω is an equilibrium state at inverse modular temperature 2π with respect to θ_t . Modular flow is thermal time; the KMS condition is the statement that the two coincide.

3.2 Bisognano–Wichmann: the geometric form

The abstract generator K acquires geometric meaning through the Bisognano–Wichmann theorem. For a Rindler wedge W in Minkowski space, Bisognano and Wichmann showed — from the Wightman axioms alone, with no geometric input — that the modular flow of $(\mathcal{A}(W), \omega_{\text{Mink}})$ is the boost preserving W , and that the modular Hamiltonian is the boost generator normalized to unit surface gravity,

$$K_W = 2\pi B_W = 2\pi \int_{\Sigma} T_{\mu\nu} \xi_{\text{boost}}^\mu d\Sigma^\nu. \quad (10)$$

The wedge geometry is an output of the operator algebra, not an input to it. Hislop and Longo extended the result to double cones (causal diamonds) in conformal field theory; Casini,

Huerta and Myers showed by conformal mapping that the modular state on a causal diamond is unitarily equivalent to a thermal state on a hyperbolic geometry, making the diamond’s entropy computable as a horizon entropy. The chain runs onward through general causal diamonds, de Sitter and conformally flat backgrounds, and FLRW cosmic diamonds, where the boundary of the diamond is a Killing horizon on which the Kodama vector becomes null.

Postulate 1 (Modular Hamiltonian on the cosmic diamond). On the cosmic causal diamond the modular Hamiltonian takes the Bisognano–Wichmann form

$$K = 2\pi \int_{\Sigma} T_{\mu\nu}^{\text{info}} \xi^{\mu} d\Sigma^{\nu} \quad (11)$$

with Σ any Cauchy surface for $\mathcal{D}$, and ξ^{μ} the conformal Killing vector that preserves $\mathcal{D}$, is timelike in its interior, vanishes on the screen σ , and is normalized to unit surface gravity there. The stress tensor is vacuum-subtracted, so that $K|\Omega\rangle = 0$ as the Tomita–Takesaki construction requires; a vacuum term proportional to $g_{\mu\nu}$ adds only a c-number to K , which the modular flow does not see. The prefactor 2π is the Bisognano–Wichmann normalization and is not adjustable.

3.3 The conformal Killing vector

The vector ξ^{μ} deserves explicit statement, because three of the framework’s results depend on its structure and not merely on its existence.

For a diamond of proper-time separation τ between tips p and q , write flat coordinates with p at $t = -\tau/2$ and q at $t = +\tau/2$, both at the spatial origin. The conformal Killing vector preserving the diamond is

$$\xi^{\mu} \partial_{\mu} = \frac{2\pi}{\tau} \left[\frac{\tau^2/4 - t^2 - r^2}{2} \partial_t - t r \partial_r \right], \quad (12)$$

which is timelike inside $\mathcal{V}(p, q)$, null on the two sheets of $\partial\mathcal{D}$, and vanishes identically on the bifurcation surface σ at $t = 0$, $r = \tau/2$. It generates the conformal isometries of the diamond, and the flow it generates is the modular flow of Eq. (7). Three consequences follow, and are used later without re-derivation.

(i) *The near-screen decomposition.* At the bifurcation surface σ , where $\xi^{\mu} = 0$, the local four-metric decomposes orthogonally into a time–radial plane and the transverse sphere,

$$ds_{\text{local}}^2 \Big|_{\sigma} = ds_{\mathbb{R}^{1,1}}^2 \oplus r_h^2 d\Omega_{S^2}^2, \quad r_h = c/H, \quad (13)$$

and the modular flow reduces there to the boost of the $\mathbb{R}^{1,1}$ factor, $K_{\text{local}} = 2\pi B$. This decomposition is generic at any bifurcation surface where a conformal Killing vector vanishes. It supplies the electromagnetic phase space of Section 18 and the chiral edge of Section 13.

(ii) *The antipodal $\mathbb{Z}_2$.* The boost of the $\mathbb{R}^{1,1}$ factor has exact time-reversal symmetry $t \rightarrow -t$ at the level of the Bisognano–Wichmann spectrum; the round two-sphere has exact $O(3)$ invariance, hence an exact antipodal isometry $A : \hat{n} \mapsto -\hat{n}$. Their product,

$$\mathbb{Z}_2 = T_{\mathbb{R}^{1,1}} \otimes A_{S^2}, \quad (14)$$

is an exact local isometry of Eq. (13). Because the electromagnetic field is *CPT*-self-conjugate, a boundary mode at direction $\hat{n}$ on the past null sheet and its antipode on the future null sheet are a single physical degree of freedom. This is the diamond *CPT* pairing, and it halves the boundary phase space in Section 18.

(iii) *Sphericity, hence parity-evenness.* On a homogeneous and isotropic background, ξ^{μ} has only temporal and radial components, each independent of the angular directions. It is therefore invariant under the full $O(3)$ acting on the direction sphere — including the orientation-reversing element. This is the root of Lemma 11.1, and hence of the entire boundary–dynamics division.

3.4 Entropy of the modular state

The entropy of the modular state is neither $-\text{Tr}(\rho \ln \rho)$ nor $\langle K \rangle$. A type III₁ algebra has no trace and therefore no density matrix, and the entanglement entropy of a causal region diverges, with a leading term proportional to the area of its boundary. The modular state is invariant under its own flow, $\Delta|\Omega\rangle = |\Omega\rangle$, so $\langle K \rangle_\Omega = 0$. Two finite statements survive. Variations are finite: Araki’s relative entropy is defined for type III algebras, and its first-order variation is the entanglement first law $\delta S = \delta \langle K \rangle$ of Section 9. The entropy itself becomes finite once the modular flow is adjoined: the crossed product of $\mathcal{A}(\mathcal{D})$ by θ_t is of type II, admits a trace, and assigns each semiclassical state its generalized entropy $S_{\text{gen}} = \mathcal{A}/(4G\hbar) + S_{\text{fields}}$, up to a state-independent constant [12, 13, 14]. On the cosmic diamond the area term dominates, and it is the Bekenstein–Hawking area law:

$$S_H = \frac{\mathcal{A}}{4G\hbar} = \frac{\pi c^5}{G\hbar H^2} \quad (15)$$

with $\mathcal{A} = 4\pi c^2/H^2$ the area of the cosmic horizon. The modular structure fixes the form of this term: proportional to area, independent of the state. Its coefficient is the one Section 9 identifies with $1/(4G\hbar)$, and Section 19 fixes G itself. The holographic principle is recovered from the modular structure and that one identification, not imposed as an axiom. The crossed-product construction is established for black-hole exteriors, de Sitter static patches and general subregions; on the cosmic diamond it has the status of Postulate 1.

Remark 3.1 (Two counts, kept apart). Two quantities recur and must never be conflated. The horizon entropy $S_H = \mathcal{A}/(4\ell_P^2)$ counts holographic *degrees of freedom*. The Planck-cell count

$$N_P \equiv \frac{\mathcal{A}}{\ell_P^2} = 4 S_H, \quad \ell_P^2 \equiv \frac{G\hbar}{c^3}, \quad (16)$$

counts the *tiles* of the screen. They differ by the holographic assignment of one degree of freedom per four Planck areas:

$$\ln N_P = \ln S_H + \ln 4. \quad (17)$$

At the present epoch, evaluated geometrically from the measured G and H_0 ,

$$S_H(0) = 2.268 \times 10^{122}, \quad N_P(0) = 9.07 \times 10^{122}, \quad \ln S_H = 281.734, \quad \ln N_P = 283.12. \quad (18)$$

(Section 19 will obtain a second value, $S_H = 2.285 \times 10^{122}$, from the measured fine-structure constant rather than from the geometry. The two differ by $\Delta \ln S_H = 0.0075$, and that difference is the framework’s entire disagreement with experiment. They must not be interchanged: the geometric value is the one that appears throughout Parts II–III.) The rate γ of Section 5 is defined against the *entropy*; the Planck-cell count appears only where cells, not degrees of freedom, are being counted — in the cosmological-constant hierarchy of Section 17 and the screen tensor network of Section 20. The distinction is load-bearing: substituting $\ln N_P$ for $\ln S_H$ in Section 18 would displace α^{-1} by $\frac{1}{2} \ln 4 = 0.69$ and destroy an agreement at the 10^{-5} level.

Throughout, cosmological evaluations use the Planck 2018 cosmic-horizon expansion rate

$$H_0 = 67.36 \text{ km s}^{-1} \text{ Mpc}^{-1} = 2.1830 \times 10^{-18} \text{ s}^{-1}, \quad \Omega_\Lambda(0) = 0.6847. \quad (19)$$

This is the rate that sets the cosmic-horizon area, and it is not interchangeable with a distance-ladder determination; Section 16 shows why, and Section 19 shows how sharply the framework discriminates between them.

4 The Quantum-Thermodynamic Entropy Partition

4.1 The ebit

Consider two systems in a maximally entangled state, $|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ — the paradigm being a photon–electron pair during Thomson scattering. Tracing out one party,

$$\rho_{\text{red}} = \text{Tr}_{\text{partner}} |\psi\rangle\langle\psi| = \frac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|), \quad (20)$$

and the von Neumann entropy of the reduced state is

$$S_{\text{coh}} = -\text{Tr}(\rho_{\text{red}} \ln \rho_{\text{red}}) = \ln 2 \approx 0.6931 \text{ nat}. \quad (21)$$

This is the *ebit*: the coherent capacity carried by one unit of entanglement. It is derived, not chosen. Within the algebraic setting it is the minimal unit of modular structure on $\mathcal{A}(\mathcal{D})$ that admits a partition consistent with the KMS condition (9).

4.2 The obit

Postulate 2 (The obit). The classical record produced by one precipitation event has information content

$$\boxed{S_{\text{obit}} = 1 \text{ nat}} \quad (22)$$

— the indivisible quantum of irreversibility.

This is the framework’s single undeduced quantity. Its status is that of $\hbar$ as the quantum of action or $h\nu$ as the quantum of radiated energy: a fundamental unit, postulated rather than derived, and validated by its consequences. Two clarifications fix what is and is not being claimed.

The content of Eq. (22) is not the *1*; it is the *nat*. Being an indivisible quantum, the obit is one unit of itself by construction. What the postulate asserts is that irreversibility quantizes in the natural measure — in units of k_B — rather than in bits. The alternative candidate is not 2 nats or $\frac{1}{2}$ nat but $\ln 2$ nats, the value the ebit already carries; and were the obit equal to $\ln 2$, the residual defined below would vanish identically, the ratio η would be undefined, and the entire framework would collapse to triviality. The postulate is therefore a base selection, and it is the hinge on which everything turns.

The choice is motivated but not forced. Bekenstein–Hawking entropy and the Margolus–Levitin bound both carry natural logarithms, and the obit is a *thermodynamic* quantum — the irreversibility that couples to $k_B T$ — so the natural measure is the one in which k_B is defined. That is a motivation, not a derivation, and we do not present it as one. The postulate stands on its consequences, which are enumerated in Part IV, and on the laboratory test of Section 22, which is designed to settle it.

4.3 The partition and the ratio

Precipitation converts one ebit of coherent capacity into a classical record. The record costs one obit; the ebit supplies only $\ln 2$; the difference is the *residual*

$$S_{\text{decoh}} \equiv S_{\text{coh}} - S_{\text{obit}} = \ln 2 - 1 \approx -0.3069 \text{ nat}. \quad (23)$$

The per-event ledger is therefore a conservation law, not an inequality:

$$\boxed{S_{\text{coh}} = S_{\text{obit}} + S_{\text{decoh}}, \quad \ln 2 = 1 + (\ln 2 - 1).} \quad (24)$$

No information is destroyed in a precipitation event. The coherent capacity of one ebit is *re-partitioned*: into a definite classical record in the bulk, and a residual entry on the holographic boundary. The residual's negative sign records that it is information removed from the accessible state space of $\mathcal{A}(\mathcal{D})$ and registered on the screen — the Tomita–Takesaki commutant entry on $\partial\mathcal{D}$. It is negentropy on the boundary register, and it is not heat.

The ratio of coherent capacity to the magnitude of the residual is the universal constant of the programme:

$$\eta \equiv \frac{S_{\text{coh}}}{|S_{\text{decoh}}|} = \frac{\ln 2}{1 - \ln 2} = 2.2589 \quad (25)$$

It is dimensionless, scale-free, and independent of γ , of H , and of every cosmological parameter. It recurs at recombination (Section 15), at the present epoch, in the black-hole sector (Section 21), and in the laboratory (Section 22), and its recurrence across regimes of vastly different physical scale is the structural fingerprint of the partition.

The bulk ledger conserved within the diamond is

$$S_{\text{total}} = S_{\text{coh}} + S_{\text{decoh}} = 2 \ln 2 - 1 \approx 0.3863 \text{ nat.} \quad (26)$$

4.4 Reservoir temperatures

Each precipitation event carries a characteristic energy $\hbar\gamma$, one quantum of action per processing interval. The coherent and decoherent sectors therefore possess distinct effective temperatures,

$$k_B T_{\text{coh}} = \frac{\hbar\gamma}{S_{\text{coh}}} = \frac{\hbar\gamma}{\ln 2}, \quad k_B T_{\text{decoh}} = \frac{\hbar\gamma}{|S_{\text{decoh}}|} = \frac{\hbar\gamma}{1 - \ln 2}, \quad (27)$$

whose ratio is the inverse of Eq. (25):

$$\frac{T_{\text{decoh}}}{T_{\text{coh}}} = \frac{1 - \ln 2}{\ln 2} = \frac{1}{\eta} = 0.4427. \quad (28)$$

At the present epoch, with γ_0 from Eq. (33) below,

$$T_{\text{coh}}(0) = 8.5 \times 10^{-32} \text{ K}, \quad T_{\text{decoh}}(0) = 1.9 \times 10^{-31} \text{ K.} \quad (29)$$

The information-processing architecture of the universe operates at extraordinarily cold temperatures; what is observable is not the temperature but the ratio.

The work of converting one ebit into one obit is

$$W_{\text{ebit} \rightarrow \text{obit}} = k_B T_{\text{decoh}} S_{\text{obit}} - k_B T_{\text{coh}} S_{\text{coh}} = \hbar\gamma \left(\frac{1}{1 - \ln 2} - 1 \right) = \hbar\gamma \eta, \quad (30)$$

supplied by the diamond geometry itself at rate $\hbar\gamma$ per processing interval.

5 The holographic information-processing rate

Two bounds constrain how fast a bounded region can process information. The holographic bound caps its total capacity at S_H , Eq. (15). The Margolus–Levitin theorem caps the rate of orthogonal state transitions of a system of mean energy E at

$$f_{\text{max}} = \frac{2E}{\pi\hbar}. \quad (31)$$

For the Hubble volume the characteristic energy scales as $E \sim c^5/(GH)$, so that $f_{\text{max}} \sim (2/\pi) S_H H$: the maximum number of orthogonal updates the diamond could execute in one Hubble time. To convert this aggregate into a rate per degree of freedom, one weights it by the cost of addressing an individual degree of freedom on the screen, encoded in $\ln S_H$.

Postulate 3 (Normalization). One coarse-grained processing cycle updates every holographic degree of freedom on the screen once per Hubble time. Hence

$$\boxed{\gamma(z) = \frac{H(z)}{\ln S_H(z)} = \frac{H(z)}{\ln[\pi c^5/(G\hbar H(z)^2)]}} \quad (32)$$

The Bekenstein–Hawking and Margolus–Levitin bounds together fix the *functional form* of γ up to a normalization; Postulate 3 supplies the normalization. Dimensionally allowed alternatives are $\gamma \propto H$, which saturates the Margolus–Levitin limit with no regard for holographic capacity, and $\gamma \propto H/S_H$, which suppresses the rate by the full horizon entropy $\sim 10^{122}$. Equation (32) is the minimal scale-free choice honouring both constraints, through the slowly varying complexity factor $\ln S_H$.

Numerically, at the present epoch,

$$\ln S_H(0) = 281.73, \quad \gamma_0 = \frac{H_0}{281.73} = 7.75 \times 10^{-21} \text{ s}^{-1}, \quad \left. \frac{\gamma}{H} \right|_0 = \frac{1}{281.7} = 3.55 \times 10^{-3}. \quad (33)$$

At recombination, $z \simeq 1100$, the Friedmann background, with radiation included (a quarter of the energy density at that epoch), gives $H \simeq 5.15 \times 10^{-14} \text{ s}^{-1}$, hence $S_H \simeq 4.1 \times 10^{113}$, $\ln S_H \simeq 261.6$, and

$$\gamma(1100) \simeq 1.97 \times 10^{-16} \text{ s}^{-1}, \quad \left. \frac{\gamma}{H} \right|_{1100} \simeq \frac{1}{261.6} \simeq 3.8 \times 10^{-3}. \quad (34)$$

The dimensionless ratio $\gamma/H = 1/\ln S_H$ is nearly constant across cosmic history, because the logarithm of a number of order 10^{122} moves slowly. This near-constancy is what makes γ/H useful as a perturbative expansion parameter, and every leading-order prediction below is first order in it, with corrections at $\mathcal{O}((\gamma/H)^2) \approx 10^{-5}$.

Sub-holographic operation. The baryons of the Hubble volume at recombination process $I_{\text{scatt}} \sim 10^{85}$ nats per Hubble time through Thomson scattering near the visibility peak, at one nat per scattering, against a capacity $S_H \sim 10^{113}$: a ratio of 10^{-29} . The bound is nowhere near saturated, and γ is used throughout as a *reference rate* — a Lagrange multiplier enforcing that the flux of decoherent entropy exported across $\partial\mathcal{D}$ never exceeds the diamond’s capacity — not as an indication that the bound is saturated.

6 Time, and the Lorentzian projection

6.1 γ as the geodesic rate parameter

The rate γ has dimensions of inverse proper time, which is precisely the dimensionality required of a rate parameter for timelike geodesics. The framework takes this seriously. Standard Lorentzian geometry treats the affine parameter λ as a continuous variable; if spacetime is the projection of information processing, λ must physically represent an accumulation of processing events. We therefore identify

$$\boxed{d\lambda = \gamma(z) d\tau} \quad (35)$$

along any timelike worldline of proper time τ . The modular flow parameter s of Eq. (7) and proper time are related by the same factor, $ds = \gamma(z) d\tau$.

This refines rather than contradicts general relativity. As $\gamma \rightarrow \infty$ the discrete steps $1/\gamma$ vanish and continuous proper time is recovered; for finite γ , the geodesic equation describes the path of least computational cost between events. The characteristic interval $\tau_{\text{fund}}(z) = 1/\gamma(z)$

is the proper time of the diamond in which a single precipitation event occurs, so that the thermodynamic gradient zone of length $L_{\text{grad}} = c/\gamma$ is the causal diamond of separation $\tau = 1/\gamma$. Geometry and rate are the same statement.

6.2 Time as embedded, and why γ -variation is invisible from inside

Time, on this account, is the rate of state precipitation generated by modular flow, measured against an observer’s embedded position in the algebra. There is no external clock. An observer’s atomic clocks, instruments, and physiological processes are all subsystems of $\mathcal{A}(\mathcal{D})$ whose state changes are driven by the same modular flow that drives everything else.

A hydrogen 21 cm transition involves $\hbar$ of action per oscillation. If $\hbar$ is the quantum of action and γ the rate of state change, each oscillation involves a fixed number of modular ticks — the same number at every epoch. The atomic-second-to-tick ratio is therefore constant across cosmic history; what varies is the absolute rate at which ticks accumulate. An observer at recombination measures hydrogen at 1420 MHz against their atomic second; an observer today measures the same. Both are correct. But the recombination-epoch second contained more ticks than today’s, because γ was then some 2.5×10^4 times larger.

The undetectability of this is structural, not technical. Any measurement of γ -variation requires comparison *between* epochs; within a single epoch, γ is the rate against which all measurements are calibrated and is therefore not itself measurable. Internal invariance and cross-epoch variation are two faces of one structure.

Prediction 1 (Intracycle universality of dimensionless ratios). Every dimensionless combination of Standard Model parameters — the fine-structure constant α , the proton-to-electron mass ratio $\mu = m_p/m_e$, the ratio of Bohr to nuclear radius — is the same in any observer’s local measurement frame at any epoch, because numerator and denominator co-slaw to γ under modular flow. The QED and QCD sectors do not decouple.

This is falsifiable and is being tested. Quasar absorption spectroscopy constrains $\Delta\alpha/\alpha \lesssim 10^{-5}$ over $0 < z \lesssim 4$; the Oklo natural reactor constrains it at parts in 10^7 over 2 Gyr; big-bang nucleosynthesis constrains μ and the deuteron binding scale at the few-percent level at $z \sim 10^9$. These null results are the framework’s prediction, and they are positive evidence for cross-sector co-slaving. A confirmed positive detection of varying atomic α or μ would falsify Prediction 1 and, with it, the relational commitment that dimensionless ratios are intracycle-universal.

6.3 The Lorentzian projection

Standard cosmological observables — $H(z)$, $\rho_m(z)$, the CMB power spectrum, the BAO scale, distance-redshift relations — are projections of the modular state onto the Lorentzian representation accessible to an embedded observer. The projection is exact at leading order in γ/H and acquires corrections at $\mathcal{O}(\gamma/H)$ and beyond. The metric is not a fixed background on which physics happens; it is a derived structure, accurate in the limit where γ -variation across the relevant scales is negligible.

The cosmological tensions are exactly where that limit fails at observable precision. They are not failures of measurement; they are the framework’s structural prediction that the leading-order projection cannot perfectly track the modular state across all scales. Part IV computes them.

Part II

What the framework provides

7 Precipitation, the Landauer operation, and the second law

7.1 The elementary modular event

The *Landauer operation* is the elementary tick of modular flow on $\mathcal{A}(\mathcal{D})$: the minimal precipitation by which a coherent state acquires a definite classical record. It is intrinsic to the modular structure, not an imported logical map. Each tick re-partitions the algebraic state through the KMS condition and produces, as its irreducible byproduct, one obit of classical record and one entry of holographic boundary residual — the ledger of Eq. (24).

We use *Landauer operation* rather than the traditional *erasure*. The usage is not cosmetic. Erasure is a word of destruction, carrying the 1961 ontology in which a bit’s information ceases to exist when its register is reset. Under Eq. (24) nothing ceases to exist: the bit’s accessibility is *transferred* out of its coherent storage register into a definite partition across the obit record, the holographic boundary, and — under conditions specified in Section 22 — the thermal reservoir.

The construction is scale-invariant. The algebra $\mathcal{A}(\mathcal{D})$ restricts to apparatus-sized causal patches by GNS sub-representation, with modular flow rescaled by the local horizon. A Landauer operation in a transmon at 20 mK, in a trapped ion at room temperature, and on the cosmic diamond at the present epoch are the same algebraic object evaluated at different apparatus dimensions. The patches differ only in size.

7.2 The second law as closure

The per-event ledger of Eq. (24) is conservative: the information entries sum to zero,

$$\Delta S_{\text{system}} + \Delta S_{\text{obit}} + \Delta S_{\text{boundary}} = -\ln 2 + 1 - (1 - \ln 2) = 0. \quad (36)$$

Separately, and in addition, moving the logical register out of coherent storage requires work, and the corresponding heat is deposited in whatever thermal reservoir the system is coupled to:

$$\Delta S_{\text{reservoir}} = +\ln 2. \quad (37)$$

The total entropy change of the universe per Landauer operation is the sum,

$$\boxed{\Delta S_{\text{total}} = \Delta S_{\text{info}} + \Delta S_{\text{reservoir}} = 0 + \ln 2 = \ln 2} \quad (38)$$

i.e. $k_B \ln 2$ per bit. The second law is preserved with its origin reframed: the increase arises not from destroyed information — the information ledger is closed — but from the reservoir flow that accompanies the accessibility transition. The $k_B T \ln 2$ of Landauer’s principle is the thermodynamic price of moving a bit out of its coherent register, not of annihilating it, and it emerges in Section 22 as the continuum limit of Eq. (24) rather than as an input.

7.3 The laws of holographic thermodynamics

The framework’s thermodynamics can be stated compactly. In what follows, γ plays the role that temperature plays in classical thermodynamics: it is the quantity whose equality defines equilibrium.

Zeroth Law.

Two systems whose modular flows run at the same rate γ are in modular equilibrium. If each is in modular equilibrium with a third, they are in modular equilibrium with one another. Equilibrium is transitive in γ , not in η ; the partition ratio η is a universal constant of the substrate and carries no information distinguishing one system from another.

First Law.

Information is conserved across every precipitation event. Per event, the coherent capacity re-partitions without remainder,

$$S_{\text{coh}} = S_{\text{obit}} + S_{\text{decoh}}, \quad \Delta S_{\text{info}} = 0,$$

and the boundary and bulk entries are coupled to geometry by

$$\delta S = \frac{\delta \mathcal{A}}{4G\hbar} + 2\pi \int_{\Sigma} \delta \langle T_{\mu\nu}^{\text{info}} \rangle \xi^{\mu} d\Sigma^{\nu},$$

the consistency condition that yields Einstein's equations (Section 9).

Second Law.

Total entropy is non-decreasing, and the increase is reservoir flow, not information loss: $\Delta S_{\text{total}} = \ln 2$ per Landauer operation, Eq. (38). The partition ratio η is the *equilibrium* value of the ledger; it may be transiently violated where precipitation is deferred — as in the quantum-Zeno regime of recombination (Section 15) — and is restored when the deferral unwinds.

Third Law.

No physical state attains absolute zero. Absolute zero would require zero modular flow, hence zero decoherence rate, hence no obit; the obit floor forbids it. Equivalently, no state is fully decoherent, since full decoherence would exhaust the coherent reservoir that sources the flow. Reality maintains a finite ratio of coherent to decoherent entropy at every epoch, and the $H \rightarrow 0$ asymptote is a state of maximum coherence rather than of equilibrium (Section 14).

Law of Information Pressure.

When information density approaches the holographic bound, encoding further information requires work against existing correlations, generating a pressure

$$P_I = \frac{\gamma c^4}{8\pi G} \left(\frac{I}{I_{\text{max}}} \right)^2.$$

At saturation, $(I/I_{\text{max}}) \rightarrow 1$, the region must create new degrees of freedom by expanding (Section 21).

8 The information stress-energy tensor and its modular vacuum

8.1 Definition and inversion of the usual order

The tensor $T_{\mu\nu}^{\text{info}}$ appearing in Eq. (11) is the stress-energy density of the modular state on $\mathcal{A}(\mathcal{D})$. Standard quantum field theory takes $T_{\mu\nu}$ as fundamental and obtains K by integration; algebraic primacy inverts this. K is foundational, and $T_{\mu\nu}^{\text{info}}$ is *defined* as the tensor for which Eq. (11) generates the modular flow. Its normalization is the Bisognano–Wichmann 2π , and its tie to entropy is the first law $\delta S = \delta \langle K \rangle$ (Section 9).

8.2 The vacuum form, forced by symmetry

The modular vacuum $|\Omega\rangle$ on $\mathcal{D}$ is invariant under the conformal isometries generated by ξ^μ . Any rank-two tensor expectation in this state must therefore be invariant under those isometries. The only such tensor is the metric. Hence

$$\boxed{\langle T_{\mu\nu}^{\text{info}} \rangle_\Omega = -\rho_\Lambda g_{\mu\nu}} \quad (39)$$

— the cosmological-constant form, with equation of state $p = -\rho$. No fluid ansatz, no fine-tuning, no input beyond the algebraic structure of $\mathcal{A}(\mathcal{D})$. The *form* of the vacuum stress-energy is forced by symmetry alone. Its coefficient is not fixed by K , from which Postulate 1 subtracts this term as a c-number the modular flow does not see; only ρ_Λ remains to be computed, and Section 17 computes it.

8.3 The total horizon stress-energy density

The modular state is thermal at the Gibbons–Hawking temperature of the cosmic horizon,

$$T_{\text{mod}} = \frac{\hbar H}{2\pi}, \quad (40)$$

which follows from the unit-surface-gravity normalization of ξ^μ at σ together with the 2π of Eq. (11). The total horizon energy is identified with entropy times temperature,

$$E_{\text{horizon}} = S_H T_{\text{mod}} = \frac{\pi c^5}{G \hbar H^2} \cdot \frac{\hbar H}{2\pi} = \frac{c^5}{2GH}, \quad (41)$$

and dividing by the Hubble volume $V_H = (4\pi/3)(c/H)^3$ gives the total information stress-energy density

$$\boxed{\rho_{\text{info}}(z) = \frac{E_{\text{horizon}}}{V_H} = \frac{3c^2 H^2(z)}{8\pi G}} \quad (42)$$

— the de Sitter horizon energy density. The identification is not a consequence of the modular state, whose expected modular energy vanishes (Section 3). $E_{\text{horizon}} = S_H T_{\text{mod}}$ is the energy enclosed by the Hubble radius on a flat Friedmann background, and Theorem 8.1 states what adopting it commits the framework to.

Theorem 8.1 (Flatness is a commitment). *With the identification $E_{\text{horizon}} = S_H T_{\text{mod}}$, Eq. (42) is identically the critical density at every epoch:*

$$\rho_{\text{info}}(z) = \rho_{\text{crit}}(z) = \rho_\Lambda(z) + \rho_m(z) + \rho_r(z). \quad (43)$$

The identity holds on a flat Friedmann background, and the framework has no parameter with which to accommodate curvature. Spatial flatness is therefore a commitment of the framework, not an output of it: the total horizon stress-energy saturates the critical-density bound at every epoch, with the partition into vacuum and excitations determined by the matter and radiation content of the Lorentzian state above the modular vacuum.

A detection of $\Omega_K \neq 0$ at high significance therefore falsifies the framework; it has no parameter with which to absorb curvature.

8.4 Decomposition and modular-flow evolution

The tensor decomposes into its vacuum and excitation parts,

$$T_{\mu\nu}^{\text{info}} = -\rho_\Lambda g_{\mu\nu} + T_{\mu\nu}^{\text{matter}}, \quad (44)$$

so that $\rho_{\text{info}} = \rho_{\Lambda} + \rho_m + \rho_r$. Under modular flow e^{-iKs} , $T_{\mu\nu}^{\text{info}}$ transforms covariantly as a tensor on $\mathcal{D}$, with s related to proper time by Eq. (35). The Lorentzian conservation law $\nabla^\mu T_{\mu\nu}^{\text{info}} = 0$ holds wherever the local Lorentzian description is accurate, but the underlying modular evolution runs at γ , not at H .

This rate distinction is the engine of every scale-dependent prediction in this document. The metric evolves at H ; the information content evolves at γ ; integrating the mismatch along comoving worldlines across logarithmic distance separations produces the projection formula of Section 16, and integrating it across the recombination visibility window produces the acoustic enhancement of Section 15.

9 Gravity as the bulk consistency condition

9.1 Thermodynamic cost as relative entropy

The thermodynamic cost of preparing a state $\rho_{\mathcal{A}}$ on $\mathcal{A}(\mathcal{D})$, relative to the modular vacuum ρ_0 , is its relative entropy:

$$\boxed{\mathcal{C}[\rho_{\mathcal{A}}] = S(\rho_{\mathcal{A}} \parallel \rho_0) = \langle K \rangle_{\rho_{\mathcal{A}}} - S[\rho_{\mathcal{A}}]} \quad (45)$$

This is a free-energy difference quantifying distinguishability from the vacuum. Positivity of relative entropy gives $\mathcal{C} \geq 0$, with equality if and only if $\rho_{\mathcal{A}} = \rho_0$. Gravity, on this account, is the cost of being distinguishable from nothing. Equation (45) is used again, without modification, to compute the cost of precipitating a baryon in Section 20.

9.2 The entanglement first law

Varying the state at fixed background, the first-order variation of Eq. (45) vanishes, giving the entanglement first law

$$\delta S_{\text{EE}} = \delta \langle K \rangle, \quad (46)$$

established for holographic conformal field theories by Faulkner and collaborators. Combining Eq. (46) with the area law $S_{\text{EE}} = \mathcal{A}/(4G\hbar)$ and converting the area variation to a bulk integral by the Raychaudhuri equation on the null sheets of $\partial\mathcal{D}$,

$$\frac{d\theta}{d\lambda} = -\frac{\theta^2}{2} - \sigma_{\mu\nu}\sigma^{\mu\nu} + \omega_{\mu\nu}\omega^{\mu\nu} - R_{\mu\nu}k^\mu k^\nu, \quad (47)$$

one obtains the linearized field equations with $T_{\mu\nu}^{\text{info}}$ as source,

$$\delta \langle T_{\mu\nu}^{\text{info}} \rangle = \frac{1}{8\pi G} \delta G_{\mu\nu}. \quad (48)$$

9.3 The full equations, and the fixing of Λ

Jacobson's construction completes the argument at the nonlinear level: imposing the area-thermodynamics first law $\delta Q = T \delta S$ on *all* local Rindler horizons through every point, with T the Unruh temperature and S the area entropy, forces a pointwise relation between the Einstein tensor and the modular-state stress-energy. The Bianchi identity $\nabla^\mu G_{\mu\nu} = 0$ together with conservation of $T_{\mu\nu}^{\text{info}}$ selects $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}R g_{\mu\nu}$ as the unique divergence-free curvature integrand, and hence

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}^{\text{info}}. \quad (49)$$

The local argument leaves one freedom undetermined: $G_{\mu\nu} \rightarrow G_{\mu\nu} + \Lambda g_{\mu\nu}$ is equally divergence-free, so Λ is not fixed by any local construction. It is fixed *globally*, by the modular vacuum of the

cosmic causal diamond — the coefficient ρ_Λ of Eq. (39), computed in Section 17. Substituting the decomposition (44) into Eq. (49) and moving the vacuum piece to the left:

$$\boxed{G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}^{\text{matter}}, \quad \Lambda = \frac{8\pi G}{c^2} \rho_\Lambda.} \quad (50)$$

This is the synthesis in one line. Gravity is the bulk geometric encoding of γ -controlled information processing on $\mathcal{A}(\mathcal{D})$: the Einstein tensor is the consistency condition under which the modular ledger can be written on the holographic boundary at all. General relativity is recovered as the linearized excitation theory of $T_{\mu\nu}^{\text{info}}$ above the modular vacuum, with the cosmological constant entering automatically as the vacuum contribution rather than as a separate parameter. The *form* of the field equations is fixed here; their two coefficients are fixed elsewhere — Λ in Section 17, and G itself in Section 19.

9.4 Information pressure and the information action

The stress-energy carried by information gradients follows from an action principle. Taking

$$S_{\text{info}} = \int d^4x \sqrt{-g} \left[\frac{\gamma \hbar}{2c^2} \nabla_\mu S_{\text{total}} \nabla^\mu S_{\text{total}} \right], \quad (51)$$

and varying with respect to the metric, $\delta S_{\text{info}}/\delta g^{\mu\nu} = 0$, yields

$$T_{\mu\nu}^I = \frac{\gamma \hbar}{c^2} \left[g_{\mu\nu} \nabla_\alpha S_{\text{total}} \nabla^\alpha S_{\text{total}} - \nabla_\mu S_{\text{total}} \nabla_\nu S_{\text{total}} \right], \quad (52)$$

with the first term an isotropic information pressure and the second an anisotropic information flux. Diffeomorphism invariance of Eq. (51) guarantees $\nabla^\mu T_{\mu\nu}^I = 0$. Contracting to the isotropic part and expressing the result in terms of the occupied fraction of the holographic bound gives the information pressure,

$$\boxed{P_I = \frac{\gamma c^4}{8\pi G} \left(\frac{I}{I_{\text{max}}} \right)^2, \quad I_{\text{max}} = \frac{\mathcal{A}}{4G \ln 2}} \quad (53)$$

The quadratic dependence has a two-part origin: each new bit must maintain quantum correlations with all existing bits, contributing one factor of I/I_{max} , and the phase space available for consistent holographic encoding shrinks linearly with occupancy, contributing the second. When P_I reaches the critical value $P_c = \gamma c^4/(8\pi G)$, further encoding is impossible and the region must expand to create new degrees of freedom. Section 21 follows this to its conclusion.

9.5 Gravitational modulation of the processing rate

Because γ is set by the local holographic capacity, it responds to the gravitational field. In a static field of mass M at radius r ,

$$\boxed{\gamma_{\text{local}} = \gamma_\infty \sqrt{g_{00}} \left(1 + \frac{GM}{rc^2} \frac{V_{\text{local}}}{V_{\text{critical}}} \right)} \quad (54)$$

The first factor is ordinary gravitational time dilation. The second is not: it depends on the information-processing architecture — on the local causal volume relative to the critical volume at which boundary optimization projects into bulk curvature — and it is the framework's sharpest departure from pure general relativity at the level of measurement rates. It supplies the decoherence-channel signature of Section 20.

Part III

The boundary, and what is not derived

10 The degenerate boundary and the first record

10.1 The diamond at $t = 0$

The machinery of Part II is well defined throughout the cosmic interior and degenerates at the past tip. At $t = 0$ the diamond $\mathcal{D}(p, q)$ has zero proper-time separation for any q : the four-volume (2) vanishes, the entropy (15) is undefined as $H \rightarrow \infty$, and the ledger (24) is identically zero because no precipitation has occurred. The strict $t = 0$ limit is not an evaluation point of the framework; it is a boundary marker.

The degeneracy is a statement about our measurement frame, not an ontological claim. Looking back along our null structure to where it terminates, the boundary *appears* degenerate, and this is the only frame the framework supports for empirical inquiry. Whether a frame inaccessible to us would find a saturated coherent reservoir or a pre-existing modular configuration there, the framework does not adjudicate, and Section 14 shows why it cannot.

10.2 The first event

The boundary contains no specification of a preferred first event; it is structureless. The mechanism that initiates precipitation cannot be gravitational, because within this framework gravity is the cost of distinguishing states from the vacuum (Section 9) and does not exist prior to the first event that establishes any distinguishable state. Any account invoking gravitational self-energy would presuppose the emergence it is required to enable. The mechanism is thermodynamic, and it operates through the framework’s own ledger.

A universe-scale superposition at the boundary is pure coherent entropy: S_{coh} carries the state’s entire content, and $S_{\text{decoh}} = S_{\text{obit}} = 0$. With no decoherent entry, no negentropy exists to populate any past light cone, and the absence of past-light-cone content *is* the geometric signature of degeneracy: the J^- branch is empty, the J^+ branch is the whole state, and there is no intersection to define. The QTEP ratio (25) diverges. The pure-coherent boundary state sits at maximal departure from the partition, and the framework’s requirement that physically realized configurations possess a definable partition is precisely the instability that fires the first event.

The instability is the second law applied to the framework’s own ledger. Each precipitation converts one ebit into the partitioned configuration, and the total entropy accounting across a single event is

$$\underbrace{\ln 2}_{\text{pure coherent}} \longrightarrow \underbrace{\ln 2}_{S_{\text{coh}}} + \underbrace{(\ln 2 - 1)}_{S_{\text{decoh}}} + \underbrace{1}_{S_{\text{obit}}} = 2 \ln 2, \quad (55)$$

an increase of $\ln 2$ nats — Eq. (38) again, now at the boundary. The transition from the unpartitioned state to the partitioned one is thermodynamically forced.

10.3 Co-generation

The first event does not occur *within* a pre-existing geometry. It generates the geometry. Three quantities are produced concurrently and cannot be separated in time:

1. the first non-degenerate diamond $\mathcal{D}(p, q')$, of proper-time separation $\tau = 1/\gamma_P$;

2. the first orbit, $S_{\text{orbit}} = 1$ nat, deposited on the screen;
3. the first bulk partition, $S_{\text{coh}} = \ln 2$ and $S_{\text{decoh}} = \ln 2 - 1$.

The regulator convention $t = 0^+ \equiv t_P$ is not an inserted cutoff. The first γ -tick along any timelike worldline emerging from p defines, by Eq. (35), the smallest operational interval the framework can resolve. Coordinate time acquires meaning at the first tick and not before, because prior to it no events have occurred and no clock exists to register elapsed time. The Planck time is the unique value at which operational time and Friedmann coordinate time first coincide.

10.4 The eviction direction, and the seed of the past

Away from the boundary, the residual S_{decoh} is evicted into the *past* light cone of the precipitation event; this is the interior rule, and it is what renders the residual thermodynamically inaccessible. At the first event there is no substantive past light cone to receive it. Eviction is therefore into $J^+(q')$ — forward.

The two pictures describe the same region from different reference events. The future light cone of the first event is precisely what becomes the past light cone of every subsequent observer at any $t > t_P$. The negentropy deposited at the boundary, evicted forward for want of a past to receive it, is the seed of the past-light-cone content available to all of cosmic evolution. The past/future asymmetry of QTEP eviction is not violated at the boundary; it *emerges* with the cosmic evolution that the first event seeds. Only the first event ejects forward.

10.5 Planckian initial conditions

In Planck units ($\hbar = c = G = k_B = 1$) with $H = 1$ at $t = t_P$, the framework's own quantities follow with no adjustable input. From Eq. (15),

$$S_H \longrightarrow \pi, \quad \ln S_H = \ln \pi = 1.1447, \quad (56)$$

so that from Eq. (32)

$$\gamma_P = \frac{1}{\ln \pi} = 0.8736, \quad \tau = \frac{1}{\gamma_P} = \ln \pi = 1.1447. \quad (57)$$

The energy per event is fixed by saturation of the Margolus–Levitin bound (31),

$$E_{\text{event}} = \frac{\pi}{2} \hbar \gamma_P = \frac{\pi}{2 \ln \pi} = 1.3722 E_P, \quad (58)$$

with corresponding temperature $T_{\text{event}} \approx 1.37 T_P \approx 1.94 \times 10^{32}$ K. The factor $\pi/2$ is required and is not a convention: E and γ are related by the Margolus–Levitin orthogonality time, not by a de Broglie oscillator relation.

From Eqs. (2) and (3) at leading order,

$$V_4 = \frac{\pi}{24} (\ln \pi)^4 = 0.2248, \quad V_3 = \frac{4\pi}{3} \left(\frac{\tau}{2} \right)^3 = \frac{\pi}{6} (\ln \pi)^3 = 0.7855, \quad (59)$$

in units of $\ell_P^3 t_P$ and ℓ_P^3 respectively, giving the densities

$$\rho_3 \equiv \frac{E_{\text{event}}}{V_3} = 1.747 \rho_P, \quad \varepsilon_4 \equiv \frac{E_{\text{event}}}{V_4} = 6.10 \rho_P t_P^{-1}. \quad (60)$$

Against this, the Friedmann equation run backward to $H = 1$ in radiation domination gives

$$\rho_F = \frac{3H^2}{8\pi G} \longrightarrow \frac{3}{8\pi} = 0.1194 \rho_P. \quad (61)$$

Theorem 10.1 (Planckian onset). *Two structurally independent routes — the causal-diamond ledger, and the Friedmann equations run backward — converge on energy densities of order ρ_P at $t = 0^+ \equiv t_P$. The ratio $\rho_3/\rho_F = 14.6$ decomposes transparently into the inverse single-event spatial volume, $1/V_3 = 6/[\pi(\ln \pi)^3] = 1.27$ in Planck units; the event energy $\pi/(2 \ln \pi) = 1.37$, which carries the $\pi/2$ Margolus–Levitin factor of Eq. (58); and the $3/(8\pi)$ Friedmann prefactor. None of these is fitted.*

The claim is consistency at $\mathcal{O}(1)$, not exact equivalence: the two numbers are different geometric definitions of “Planck-scale energy density,” one a single-event spatial density and the other a Hubble-averaged density. What the result establishes is that the Planck-scale onset of cosmic history — a boundary *input* in standard Friedmann cosmology, and structurally unexplained there — is a boundary *output* here: the smallest operational interval the diamond can resolve, set by the first γ -tick, with energy fixed by Margolus–Levitin saturation.

11 Parity-evenness of the dynamics

Let P be spatial inversion, acting on the transverse sphere of the screen as the antipodal map $A: \hat{n} \mapsto -\hat{n}$. It is represented as a unitary U_P on the realized Hilbert space and equivalently as a $*$ -automorphism α_P of $\mathcal{A}(\mathcal{D})$. A structure is *parity-even* if it is α_P -invariant, *parity-odd* if its sign reverses. Parity is logically independent of time reversal T and of charge conjugation C : a thermodynamic arrow is a T -structure and a matter/antimatter asymmetry a C -structure, and neither is a P -structure.

Lemma 11.1 (Parity-evenness). *The parity automorphism α_P preserves the modular state ω and commutes with the framework’s dynamics:*

$$\alpha_P \circ \theta_t = \theta_t \circ \alpha_P, \quad \alpha_P \circ \Phi_t = \Phi_t \circ \alpha_P,$$

for the modular flow θ_t and the precipitation map Φ_t . Consequently every structural invariant of the framework is α_P -invariant.

Proof. (i) *The modular structure.* On a homogeneous and isotropic background the conformal Killing vector ξ^μ of Eq. (12) has only temporal and radial components, each independent of the angular directions, and is therefore invariant under the full $O(3)$ acting on the direction sphere. Parity is the orientation-reversing element of that same $O(3)$, so it preserves ξ^μ ; an isometry preserving the generating vector field commutes with the flow it generates, giving $\alpha_P \circ \theta_t = \theta_t \circ \alpha_P$ globally on $\mathcal{D}$, not merely near σ . The same conclusion follows in the near-bifurcation frame (13), where the flow is the boost of the $\mathbb{R}^{1,1}$ factor and $\alpha_P = \mathbf{1}_{\mathbb{R}^{1,1}} \otimes A_{S^2}$ acts on the orthogonal S^2 : the flow is T -oriented but carries no transverse handedness. Since ω is the KMS state of θ_t , Eq. (9), its uniqueness gives $\omega \circ \alpha_P = \omega$.

(ii) *The partition and the rate.* The coherent and decoherent entropies are von Neumann and relative entropies built from ω , hence functionals of the modular state; entropy is invariant under any state-preserving automorphism, so by (i) the partition and η are α_P -invariant. The rate $\gamma = H/\ln S_H$ is a functional of the homogeneous geometry alone — $H = \dot{a}/a$ and $S_H = \pi c^5/(G\hbar H^2)$ are unchanged by spatial inversion — so γ is α_P -invariant.

(iii) *The antipodal isometry.* The transverse parity used in (i) is exact, not approximate: by the $O(3)$ invariance of the round two-sphere, A_{S^2} is an exact isometry of the S^2 factor of Eq. (13). The $\mathbb{Z}_2$ of Eq. (14) is a *conservation* symmetry; the construction contains no parity-odd structure anywhere.

(iv) *The arrow of classicalization.* Eviction of the residual into the past light cone orders states by increasing coarse-grained entropy toward the future. This is a T -asymmetry. It is α_P -invariant

because the entropy gradient is timelike and carries no transverse handedness. A thermodynamic arrow does not, by itself, induce a P -asymmetry.

(v) *The precipitation dynamics.* The map Φ_t is generated by open-system Liouvillian dynamics whose ingredients — the acoustic Hamiltonian and the Thomson-scattering operators encoding the baryon charge density (Section 15) — are each parity-even. As these exhaust the generators and each commutes with α_P , parity-covariance passes to the semigroup they generate.

Collecting (i)–(v): α_P preserves ω , commutes with both flows, and leaves every structural invariant fixed. $\square$

Remark 11.2. The lemma uses only that γ and S_H are geometric and that the partition is a functional of the modular state. These are objects defined for the cosmological problem, with no reference to chirality. The parity-evenness is an independent property of the machinery, not a feature arranged to obtain Postulate 4. The implication runs one way only.

Corollary 11.3 (Symmetry preservation). *If the initial state is parity-invariant, $\omega_0 \circ \alpha_P = \omega_0$, then $\omega_t \circ \alpha_P = \omega_t$ for all t . No parity asymmetry is generated by the dynamics.*

Proof. States evolve as $\omega_t = \omega_0 \circ \Phi_t$; by Lemma 11.1, $\omega_t \circ \alpha_P = \omega_0 \circ \Phi_t \circ \alpha_P = \omega_0 \circ \alpha_P \circ \Phi_t = \omega_0 \circ \Phi_t = \omega_t$. $\square$

The cosmological initial condition ω_0 is logically distinct from the modular reference state ω of step (i). The parity-invariance established there is a property of the KMS state intrinsic to θ_t and places no constraint on ω_0 , which is free boundary data. A parity-odd ω_0 remains admissible.

12 Chirality as boundary data

The observed world violates parity. Combining that fact with Corollary 11.3 forces the handedness outside the dynamics.

Postulate 4 (Chirality as boundary data). The $V-A$ structure of the weak interaction and the chirality of the fermion spectrum together constitute the handedness of the matter sector, and that handedness is boundary data of the cosmic causal diamond. It is either a parity-odd feature of the initial state ω_0 , with the couplings parity-symmetric, or a parity-odd coupling in the precipitated content over a parity-symmetric ω_0 . In neither case is it generated by the parity-even precipitation dynamics.

Postulate 4 is not a free choice. By Lemma 11.1 the dynamics cannot produce a parity asymmetry; by experiment the asymmetry exists; the boundary is the only place left for it. The distinction it relies on is that the *dynamics* is the precipitation map, whereas the fields and couplings it precipitates are *content*: a parity-odd coupling is input the map acts on, not something the map generates, fixed by the same fiat that lets the Standard Model postulate its chiral assignment rather than derive it. The framework thus derives *that* chirality is boundary data. It does not derive *which*: the generation structure, the magnitude of CP violation, and the absolute sense of the handedness all remain input.

By Section 10, the handedness is fixed at or below the first record. Everything after inherits it and nothing after manufactures it; the left-handedness of the weak interaction is a fossil of the first event, carried forward unchanged by a dynamics that can neither originate it nor wash it out. A parity-symmetric Big Bang is therefore incompatible with the observed universe: were ω_0 parity-symmetric, Corollary 11.3 would carry that symmetry to every later time. And the mechanism that fires the first event is a parity-blind thermodynamic instability — it is the

wrong kind of thing to choose a handedness, which is exactly why the handedness must enter as data the mechanism does not supply.

12.1 The residual fork, below the framework’s floor

Postulate 4 admits two realizations — parity-odd initial state, or parity-odd content on a symmetric state — and they are empirically identical at every $t > 0$: both reproduce the measured $V-A$ spectrum. They differ only in the parity structure at the past tip p . But the machinery carries operational content only above the first record (Section 10); below it the diamond is degenerate and the machinery is identically zero, with no event to mark and no γ -tick to measure. The fork is therefore beyond the reach of any $t > 0$ observable. The question “content or initial condition” sits past the framework’s own horizon.

12.2 Two versions of the underivability claim

Weak form: chirality is underivable *within this framework*. This is established by Lemma 11.1 and is secure.

Strong form: chirality is underivable *in principle*, because it is initial-condition data and a dynamics does not explain its own boundary data. This presupposes that no deeper theory of initial conditions or vacuum selection exists, and is *not* established here. We state the weak form as a result and offer the strong form only as a labelled conjecture.

12.3 Falsifiable predictions: cosmic birefringence

Cosmic birefringence — a rotation of the CMB polarization plane — is a clean parity-violating signature, and Postulate 4 sorts the possible mechanisms into an allowed channel and forbidden ones.

A pseudoscalar field ϕ coupled through a Chern–Simons term $\mathcal{L} \supset -\frac{1}{4}g_{\phi\gamma}\phi F_{\mu\nu}\tilde{F}^{\mu\nu}$ rotates the polarization by

$$\beta = \frac{1}{2}g_{\phi\gamma}\int_{t_{\text{LSS}}}^{t_0}\dot{\phi}dt. \quad (62)$$

This is a parity-odd *content* field on a parity-even background, and is therefore *allowed*. Birefringence sourced by the *dynamics* — in the geometry, or in the propagation — is forbidden. The two are observationally distinguishable, and the distinction is the content of the predictions.

Prediction 2 (Achromaticity). The rotation (62) is frequency-independent. The *chromatic* component of cosmic birefringence vanishes: the frequency slope $d\beta/d\ln\nu$, measured across the CMB bands, is consistent with zero at $\leq 1\sigma$. A robust ($\geq 3\sigma$) detection of dispersive, frequency-dependent birefringence falsifies Postulate 4, since it would attribute parity violation to the dynamics of propagation rather than to a content field. Planck, ACT and SPT already bound this; Simons Observatory, LiteBIRD and CMB-S4 will sharpen it.

Prediction 3 (Gravitational-sector parity null). A content field coupled to photons does nothing parity-odd to the geometry. The parity-odd gravitational observables are therefore null: the net circular polarization of the stochastic gravitational-wave background, the parity-odd tensor TB/EB correlations not attributable to a coupled field, and the parity-odd four-point function of large-scale structure are each consistent with zero at $\leq 1\sigma$. A robust detection in any of these channels, surviving the requirement that it not be sourced by a coupled content field, falsifies Postulate 4. These are probed by LiteBIRD and CMB-S4 in the tensor sector, by pulsar-timing arrays and LISA for gravitational-wave chirality, and by galaxy-survey four-point analyses.

Remark 12.1 (Prohibition: no prediction of β from γ). The framework must *not* predict the amplitude of the allowed, achromatic birefringence by writing β as a function of γ or η . Doing so

would source the rotation from the dynamics, converting content-parity-violation into dynamics-parity-violation and self-falsifying by Prediction 2. The amplitude β is boundary data and is left unpredicted. This silence is what makes the predictions falsifiable rather than tunable: a framework that could fit any β predicts nothing.

Present data are consistent. The reported isotropic signal, $\beta \approx 0.34^\circ$ at $\sim 3.6\sigma$, is achromatic and electromagnetic — the allowed content channel — while gravitational-sector and early-universe parity searches have returned null. The epistemic structure is deliberately asymmetric: the predicted nulls are shared by the entire class of parity-even-dynamics theories, so confirming them cannot single out this framework, whereas a robust detection in either forbidden channel refutes it outright.

13 The boundary–dynamics division, and the $E_8 \times E_8$ apparatus

13.1 The division, stated

Postulate 5 (Boundary–dynamics division). In the causal-diamond framework, the parameters and functionals that *define the boundary data* of $\mathcal{D}$ are input, fixed at the boundary and not derivable within the framework, while the functionals that *generate the precipitation dynamics* are provided by the framework. The boundary of the framework’s derivational reach is the boundary of the diamond itself.

This is the framework’s central epistemic commitment, and the preceding sections are instances of it. Section 10: the first event writes the boundary data, and below it the machinery is degenerate. Section 12: the dynamics is parity-even, so it cannot author the chiral boundary data. Section 14: the conditions inherited from the previous cycle’s cascade are inherited, not derived. The ledger, for the case at hand:

Input, accepted (boundary-defining)	Provided, derived (dynamic)
the gauge structure $E_8 \times E_8$	the modular flow as the projection
the Calabi–Yau and the bundle	the rate $\gamma(z)$ and the QTEP partition
the generation count	the continuous scales Λ, r_s, α, G
the sense of the handedness	the locking of the chiral data into one index

Three commitments fix how to read the division. First, *spacetime is emergent*: the diamond’s modular structure is prior to the Lorentzian manifold, so no geometry here is fundamental spacetime, and boundary-data geometry need not be a spacetime to do its work. Second, what the framework derives is the *parameterization of the projection*: it restricts to the index, accepts its value as input, and obtains the index-labelled family of modular-flow maps and the balances they must satisfy. The index is fixed by observation, not a free handle, so the framework remains parameter-free in the fitting sense. Third, this is entropy accounting in the strict sense — accounting for the chiral flow and enforcing its balances without defining its source. Accepting the left column is the precondition for the right, not a weakness; it is what every physical theory does with its boundary conditions.

13.2 Two structures that must not be summed

The heterotic string supplies two internal structures that are routinely conflated. The first is geometric: the ten-dimensional target factorizes as $10 = 4 + 6$, with the six internal dimensions a Calabi–Yau threefold of $SU(3)$ holonomy. The second is a chiral lattice: the heterotic string carries $26 - 10 = 16$ purely left-moving internal bosons on an even self-dual lattice, of central charge $c = 16$ and rank that of $E_8 \times E_8$. These are different objects — a compactified geometry and a chiral conformal field theory — their dimensions never add, and the gauge bundle links

them only through the choice of a structure group inside $E_8 \times E_8$ over the Calabi–Yau. Both are boundary-defining in the sense of Postulate 5.

We work in ten-dimensional $E_8 \times E_8$ heterotic string theory and not in a four-dimensional unified group, the latter being obstructed by the representation-theoretic no-go of Distler and Garibaldi, which forbids embedding the chiral Standard Model in a single E_8 in four dimensions.

13.3 The heterotic objects in the entropic register

The framework deals in the entropic accounting of classicality, and the heterotic objects must be read in that register. $E_8 \times E_8$ is the full, unshaped chiral *capacity* of the boundary data; the Calabi–Yau is its shaped geometry; compactification is the accounting that carries one to the other.

Heterotic name	Entropic role	Quantity accounted
$E_8 \times E_8$ current algebra	capacity	chiral entropy current c_-
Calabi–Yau (compactification)	partition	shaping of that capacity
index $\frac{1}{2}\chi$	count	chiral channels surviving the shaping

13.4 The screen as a chiral edge

The E_8 current algebra at level one is a chiral conformal field theory of central charge

$$c = \frac{k \dim E_8}{k + h^\vee} = \frac{1 \cdot 248}{1 + 30} = 8, \quad (63)$$

so $E_8 \times E_8$ at level one has $c = 16$. A chiral edge of central charge c_- carries a universal entropy current set by that charge — the thermal-Hall, or gravitational-anomaly, current — evaluated here at the screen’s own Gibbons–Hawking temperature (40):

$$\dot{S}_{\text{edge}} \sim \frac{\pi}{6} c_- T_{\text{mod}}^2. \quad (64)$$

The chiral central charge is the *capacity of the bridge*: the channel through which boundary data is carried into the bulk.

We identify the holographic screen as the locus of such an edge. The near- σ decomposition (13) supplies exactly the two-dimensional chiral structure an edge current algebra requires, and the modular flow is its generator.

Postulate 6 (Screen edge content). The chiral edge sector supported on the holographic screen σ is the $E_8 \times E_8$ level-one current algebra, of chiral central charge $c_- = 16$, and the modular flow θ_t generates its edge dynamics. This is an identification of the boundary data’s chiral capacity, not a derivation of it.

The value sixteen is constrained, not free. The signature of an even unimodular lattice is divisible by eight — the same mod-eight periodicity that governs real spinor structures — and the minimal anomaly-free bosonic chiral edge is the E_8 unit with $c_- = 8$. The relevant reality is that of the *record*: the obit is a real, irreversible classical outcome and not a complex amplitude, so the chiral entropy current it carries is real-structured and its capacity is quantized in eights. Neither γ nor η selects among 8, 16 and 24; the quantum of eight is fixed by the reality of the obit, and the value sixteen by the two light-cone sectors below.

Proposition 13.1 (Product consistency). *The two-sided structure of the cosmic diamond selects the product group $E_8 \times E_8$ over $\text{Spin}(32)/\mathbb{Z}_2$ — the only two even self-dual lattices in sixteen dimensions. The diamond has two null sheets, two tips, and the CPT pairing (14) at σ ; we read the two E_8 factors as the chiral capacities of the future and past light-cone sectors. The symmetry*

exchanging them is a unitary internal exchange, not the antiunitary CPT pairing: the two factors share a chirality, so their capacities add, $c_- = 8 + 8 = 16$, whereas an orientation-reversing identification would relate $c_- \rightarrow -c_-$ and cancel them.

13.5 Where the handedness enters

Proposition 13.2 (Chirality from shaping). *Chirality enters the matter sector through the shaping of the $E_8 \times E_8$ capacity. The group E_8 has only real representations, so the channel is vector-like and cannot by itself distinguish a handedness; and the modular flow is parity-even by Lemma 11.1 and cannot originate one. The handedness therefore enters only when the $SU(3)$ holonomy of the shaping breaks $E_8 \rightarrow E_6$ and the complex representations of E_6 appear.*

This answers the question Postulate 4 left open. Chirality is boundary data; here we locate *where* it resides — in the shaping of the content, the content horn of the fork. The chain of information is one-way:

$$\underbrace{\text{boundary data}}_{\text{index, set at the first record}} \longrightarrow \underbrace{E_8 \times E_8 \text{ edge on } \sigma}_{\text{chiral channel}} \longrightarrow \underbrace{\theta_t \text{ projection}}_{\text{anomaly inflow, precipitation}} \longrightarrow \underbrace{\text{bulk records}}_{\text{classical, chiral universe}} . \quad (65)$$

The boundary writes the index; the modular flow reads and projects it; it does not write.

13.6 The geometry of the boundary data

The corpus's observable algebra is type III_1 , and ordinary spectral triples do not apply to type-III algebras. The appropriate notion is the *twisted* spectral triple of Connes and Moscovici, in which the triple is twisted by the modular automorphism itself.

Definition 13.3 (Boundary-data geometry). The geometry of the boundary data is the internal factor of the modular-twisted spectral triple $(\mathcal{A}(\mathcal{D}), \mathcal{H}, D)$ whose twist is the modular automorphism θ_t and whose Dirac operator D is sign-indefinite — distinct from the modular operator $\Delta = e^{-K}$, which is positive and implements the twist. The handedness of the matter sector is the index of D on this internal factor.

The Calabi–Yau is the commutative geometric model of that internal factor, and the generation number is the index in topological form, $N_{\text{gen}} = \frac{1}{2}|\chi(X)|$ for the standard embedding. Three roles meet here and are easily collapsed into one, so we separate them explicitly: the screen $\sigma \cong S^2$ is two-dimensional, and carries the data; the boundary-data geometry is six-dimensional, and is the internal spectral geometry *of* that data, not a region attached to the screen; and the spacetime into which the data precipitates is the four-dimensional diamond. These are three distinct roles, not three nested volumes. Spacetime is the arena of precipitation, the four-volume $V_4(p, q)$ in which records form; the internal geometry shapes what precipitates but is not itself a place where anything precipitates, and so is not spacetime.

13.7 The limit: the index is rigid

Theorem 13.4 (The count is input). *The net generation number $N_{\text{gen}} = \frac{1}{2}|\chi(X)| = |h^{1,1} - h^{2,1}|$ is not determined by the chiral capacity c_- or by the QTEP ratio η , and is therefore boundary data, not a dynamical output.*

Proof. The number is a topological integer, the index of D , and the index is deformation-invariant: continuous variation of the operator cannot change it. To select a value of χ , a functional must depend on χ ; the framework's dynamic data does not. The capacity $c_- = 16$ is the rank of $E_8 \times E_8$, a property of the gauge edge independent of the Hodge numbers — the same c_- sits on Calabi–Yau threefolds with $|\chi|$ ranging from a few to hundreds. The ratio $\eta = \ln 2 / (1 - \ln 2)$

is a fixed pure number, and it is *irrational*: were $\eta = p/q$ with p, q integers, then $(p + q) \ln 2 = p$, giving $\ln 2 \in \mathbb{Q}$, contradicting the irrationality of $\ln 2$. Hence η cannot equal a ratio of integer cohomology dimensions, and even approximately it fixes only a proportion, not an absolute difference. Neither datum depends on χ , so no functional built from them does. $\square$

This is Postulate 5 in its sharpest form: a topological boundary datum is input, exactly as dynamics never fixes its own topology — in general relativity, in field theory, or here. Note that the count is parity-*even*, so Lemma 11.1 does not forbid deriving it; it is underdetermined for the separate reason of index rigidity. Both the *sense* of the handedness (by parity) and the *count* (by rigidity) are input, by two independent arguments.

The state of the art agrees. Every explicit three-generation heterotic Calabi–Yau model realizes the count through a hand-chosen topology and a freely-acting discrete quotient with Wilson-line breaking, $N_{\text{gen}} = \frac{1}{2}|\chi(X)|/|\Gamma|$, and never derives it: systematic line-bundle scans impose $h^1(X, V) = 3|\Gamma|$ by hand over searches of order 10^{12} configurations with a one-in-a-billion hit rate. Even the entropic principle and its entropy functional — the one place the literature builds entropic functionals on Calabi–Yau data — weight continuous moduli and fluxes at *fixed topology* and select special points in moduli space; they do not fix χ . Theorem 13.4 is the structural reason for that uniform practice.

Remark 13.5 (The arithmetic of the boundary lattice). The sixteen left-moving bosons of the heterotic edge live on the even self-dual lattice $\Gamma = E_8 \oplus E_8$, and modularity alone fixes how many of its vectors have norm $2n$. The theta series of an even self-dual lattice of rank sixteen is a modular form of weight eight; that space is one-dimensional and spanned by the Eisenstein series, whose coefficients give [119]

$$N_{\Gamma}(n) = 480 \sigma_7(n), \quad \sigma_7(n) = \sum_{d|n} d^7. \quad (66)$$

The internal lattice carries $N_{\Gamma}(n)$ momentum states at level n ; the first level is the 480 roots. Since σ_7 is multiplicative, the Dirichlet series of the count has an Euler product, and Riemann’s ζ is one of its two factors:

$$\sum_{n \geq 1} \frac{N_{\Gamma}(n)}{n^s} = 480 \zeta(s) \zeta(s - 7). \quad (67)$$

In this precise sense the arithmetic of the primes is part of the boundary data: it organizes the state count of the lattice accepted in Postulate 6. The Riemann hypothesis is equivalent to the statement that the nontrivial zeros of Eq. (67) lie on the lines $\text{Re } s = \frac{1}{2}$ and $\text{Re } s = \frac{15}{2}$. ζ enters the construction a second time, through a special value: the light-cone computation of both critical dimensions, 26 and 10, regularizes the zero-point sums through $\zeta(-1) = -\frac{1}{12}$, and the rank $26 - 10 = 16$ is their difference.

Three qualifications fix what this does and does not say.

It does not select $E_8 \times E_8$. The only other even self-dual lattice of rank sixteen, that of $\text{Spin}(32)/\mathbb{Z}_2$, has the same theta series, and the flat tori $\mathbb{R}^{16}/\Gamma$ built on the two are isospectral [120]. No spectral or partition-function argument can choose between them. The choice in Proposition 13.1 rests on the two-sided structure of the diamond, which is not spectral, and it must.

It is not a property of physical spectra. The Euler product belongs to low rank: at ranks eight and sixteen every modular form of the relevant weight is Eisenstein, while at rank twenty-four a cusp form enters the theta series of every even self-dual lattice and the product is lost. A spectrum need not change its spacing to lose it. The levels $n + \frac{1}{2}$ of a harmonic oscillator have spectral zeta function $(2^s - 1) \zeta(s)$; add a constant to the Hamiltonian and it becomes a Hurwitz zeta function, which for generic offset has zeros with $\text{Re } s > 1$ [121].

It is not a constraint on the zeros. The dynamic data γ , η and S_H do not depend on ζ , and by the argument of Theorem 13.4 no functional built from them can locate its zeros. Like the generation count, the hypothesis is carried by the boundary data and not derived by the dynamics. The dynamics touches ζ in one place, and no more strongly: near σ the modular flow is a boost, which acts on each null generator as a dilation — the operator Berry and Keating proposed for the zeros [124]. With their cutoffs it reproduces the smooth count of the zeros, not their positions.

The framework thus stands beside the existing physical realizations of ζ [122, 123, 124, 125], with the same status: a realization, not a proof.

What the framework does describe is the setting in which the hypothesis is asked. Number theory is done inside the diamond, by observers whose information processing the framework governs: information, and so computation, is physical [37], and arithmetic is possible at all because physics admits models of it [38]. The integers those observers count with also organize physical state counts: the boundary lattice through Eq. (67), and the cosmic microwave background through the moments of its Bose–Einstein spectrum, $\int_0^\infty x^{s-1}(e^x - 1)^{-1} dx = \Gamma(s) \zeta(s)$, which give its photon number and energy densities at $s = 3$ and $s = 4$. That is the integral Riemann used to extend ζ beyond $\text{Re } s > 1$ in the paper that states the hypothesis [126]. The framework is the arena in which the hypothesis is posed, not the source of its content.

13.8 What is provided: the locking

Prediction 4 (Locking). The generation count, the magnitude of the net chirality, and the sense of the handedness are not three independent inputs but *one* datum: the index of the twisted Dirac operator D on the boundary-data geometry. The framework forbids them to vary independently. The parity-even magnitude is input by Theorem 13.4; the parity-odd sense is input by Postulate 4; and the two are the same index. An observation in which the generation count and the net chirality were independent data would falsify the mechanism.

The epistemic structure matches the rest of the corpus: modest about the inputs, which it shares with every comparable construction, and sharp about the dynamics and the locking, which are its own.

14 The future boundary: cyclicity, and the cycle-locality of scale

14.1 The $H \rightarrow 0$ asymptote is not heat death

Standard Λ CDM ends at the de Sitter asymptote: matter dilutes, the horizon stabilizes, the universe approaches thermal equilibrium with the Gibbons–Hawking horizon. Within H- Λ CDM this is not the terminus.

As $H \rightarrow 0$, the horizon’s area grows without bound, S_H grows, $\ln S_H$ grows to dominate the denominator of Eq. (32), and γ crashes toward zero. The modular state approaches the QTEP-coherent endpoint: maximum coherent entropy, minimum decoherent entropy consistent with the obit floor of the Third Law (Section 7.3), lowest accessible temperature consistent with the partition.

This is *not* heat death. Heat death is an equilibrium state of maximum entropy. The $H \rightarrow 0$ endpoint is a non-equilibrium state of maximum *coherent* entropy with a small but finite decoherent residue — structurally far from the partition’s preferred ratio η , which is balanced at intermediate temperatures and not at the extreme cold of maximum coherence. The endpoint is metastable: it has the entropy capacity to relax toward η , and requires only a decoherence seed to begin.

14.2 The decoherence cascade

Maximum coherence is structurally unstable. Spontaneous decoherence events in the coherent reservoir are not Boltzmann-suppressed in the usual way, because the reservoir has uniform low temperature throughout, with no gradient to suppress fluctuations. Each such event releases the energy gradient between coherent and decoherent entropy as new classical structure, raising the local information density and hence raising γ at that locus.

A single seed is therefore amplifying. Its precipitation increases γ locally; the increased γ accelerates further decoherence in its causal vicinity; the acceleration propagates as a front. The cascade releases the reservoir’s stored entropy as classical structure at increasing rates, with γ building from vanishing asymptotic values back to high cosmological values.

To external description — which does not exist within the framework — the cascade would take enormous duration, since its propagation is itself γ -limited and γ is asymptotically negligible. To internal observers inside the cascade, the cascade *is* the fabric of cosmic history at the new cycle’s earliest epochs: high γ , rapid precipitation, classical structure emerging from what they call their initial conditions. Their hot, dense early universe is the cascade’s leading edge, and their cosmic time begins at its first decoherent event — which is the Big Bang of Section 10, seen from the other side.

14.3 Cycle transitions have no duration

From any internal frame, the cycle one is in is the only cycle accessible, and its history begins at its earliest decoherent event. The previous cycle’s asymptotic reservoir is on the wrong side of one’s null structure — accessible only as the boundary condition that sourced one’s own precipitation, never as a temporally prior state.

This dissolves the question “how long does the transition take.” It has no duration in any internal frame. The previous cycle’s last observers experience their history as approaching the asymptote with γ shrinking; eventually their measurement frame can no longer sustain measurement, and from their frame the next cycle never begins. The next cycle’s first observers experience their history as starting at the cascade’s first event; from their frame, the previous asymptote is the boundary condition preceding their accessible spacetime. The transition is real and links the cycles algebraically, but it is not a temporally extended process from anywhere.

The argument is structural, not anthropic. It requires only that $\mathcal{A}(\mathcal{D})$ at the $H \rightarrow 0$ asymptote be unstable to decoherence cascades, and that the cascades produce modular-flow regimes algebraically equivalent to the original. It does not require that observers survive, or that subsequent cycles support them. Nor does it require an external conformal rescaling: the cascade *is* the rescaling, performed internally by QTEP relaxation.

14.4 Scale is cycle-local

Scale is the ratio of a physical length to the observer’s local meter, and the local meter is the speed of light times the local atomic second. The atomic second is slaved to γ (Section 6), and γ varies with epoch. *There is no cross-cycle absolute meter.* Each cycle’s measurement frame is internally consistent and externally incommensurable.

Several familiar questions dissolve accordingly.

The size of the universe at the Big Bang is not an absolute quantity. It is a measurement-frame statement about the boundary of the current cycle’s null structure, in the current cycle’s units. That the early universe was “small” is correct as an intracycle statement — in our meters, it was — and does not translate into any absolute claim. Observers at recombination, with their own

clocks and meters, would have measured a normal-sized universe, not a small one. Smallness is our retrospective measurement, transported through our γ -evolved meter into a regime where γ was different.

The degeneracy of $t = 0$ is itself frame-dependent. From inside our cycle, $t = 0$ appears as a degenerate diamond, because we calibrate against our cycle’s expanded S_H and the leading-edge state has minimal S_H in our frame. From the previous cycle’s frame, the same state is a saturated coherent reservoir with vast capacity — the opposite of degenerate. Both descriptions are correct in their respective frames, and the framework privileges neither. This is why Section 10 adopts degeneracy as a working assumption of *our* frame rather than as an ontological claim.

What came “before” is not a temporal question. The previous cycle’s history is real — it produced the algebraic boundary conditions ours inherited — but it is not *prior* in any internal clock’s sense; it is on the other side of the algebraic boundary that links the cycles. The well-posed version is: what boundary condition did our cycle’s modular flow inherit from the previous cycle’s cascade? That has a structural answer. It does not have a temporal one.

14.5 What is and is not measurable

Measurement is intrinsically intracycle. What can be measured is what can be expressed as a ratio of modular-flow processes within a single observer’s diamond: $H_{\text{local}}/H_{\text{CMB}}$, r_d/D_M , $\rho_\Lambda/\rho_{\text{crit}}$, $\Omega_{\text{dm}}/\Omega_{\text{b}}$, α , η . Every prediction in Part IV is of this form.

What cannot be measured is anything requiring cross-cycle calibration: the absolute size of the universe, the total energy of the observable universe, the duration of the previous cycle. These are not failures of precision. They are structural consequences of intracycle measurement being the only kind the framework supports, and the framework asserts that such questions do not have answers in the sense their askers intend.

The boundaries of physics, accordingly, are not where measurement precision runs out. They are where the algebraic structure ceases to support cross-frame comparison. Within the cycle, physics is precise, predictive, and falsifiable. Across cycles, certain questions are simply not well formed.

14.6 The arrow of time, and the fine-tuning of initial conditions

The thermodynamic arrow is a structural prediction rather than an initial-condition specification. The QTEP drive toward η is the algebraic source of irreversibility, with the orbit as its quantum; cosmic evolution is the progressive precipitation of coherent entropy into classical record, accumulating irreversibility as it goes. Each cycle begins at high γ and high coherent entropy, evolves through decoherence at falling γ , and asymptotes to maximum coherence at vanishing γ . The arrow points from cascade-onset toward asymptote. It is intrinsic to the cycle’s structure, not added as a boundary condition, and the “low entropy” of the early universe in the standard formulation translates here into “high coherent reservoir, rapid decoherence cascade” — a state of the cascade, not a fine-tuning.

The apparent fine-tuning of initial conditions is reframed likewise. Flatness is a commitment of the framework (Theorem 8.1), with no parameter in which it could be tuned. The cosmological-constant scale is a counting argument (Section 17), not a tuning. The horizon problem dissolves once one recognizes that γ in the early universe was far larger than today, so that horizon-scale correlations reflect γ -evolution rather than acausal conspiracy. What inflation accomplishes phenomenologically — flatness, smoothness, a near-scale-invariant spectrum — the framework’s algebraic foundation requires directly. Whether inflation *also* occurs as a separate mechanism

is a question for future work; the framework does not need it for the structural results it was proposed to deliver.

Part IV

The projections

Everything that follows is a projection of the same modular structure into a different regime. No new postulate is introduced; each section evaluates Part II's machinery at a different epoch or scale and reports what it returns. The predictions are parameter-free in the strict sense: every input is either a measured observable or a quantity already derived above, and no coefficient is fitted to the data it predicts.

15 Recombination: the quantum-Zeno regime and the acoustic ruler

15.1 The plasma is monitored

At $z \simeq 1100$ a baryon's charge density is measured each time a photon scatters off its electron. Photons outnumber baryons by 1.6×10^9 , so the measurement rate per baryon is

$$\Gamma_m = X_e n_\gamma \sigma_T c \approx 1.1 \times 10^{-2} X_e \text{ s}^{-1}, \quad (68)$$

against a Hubble rate $H = 5.15 \times 10^{-14} \text{ s}^{-1}$. Each baryon therefore undergoes

$$\mathcal{M} \equiv \frac{\Gamma_m}{H} \approx 2 \times 10^{11} X_e \quad (69)$$

scatterings per Hubble time: above 10^{11} while the plasma is ionized, and $\mathcal{M} \approx 2 \times 10^{10}$ at the peak of the visibility function, where $X_e \approx 0.1$. The photon's own scattering rate, $n_e \sigma_{TC} \approx 5.0 \times 10^{-12} X_e \text{ s}^{-1}$ with $X_e = n_e/n_H$, sets the optical depth and the epoch of transparency; it is not the measurement rate. In the open-system description each scattering is a generalized position measurement of the baryon charge density, with interval $\tau_m = \Gamma_m^{-1}$, while the intrinsic evolution time of long-wavelength acoustic and diffusive modes is $\tau_{\text{diff}} \sim H^{-1}$ or longer.

The condition for the quantum Zeno regime is $\Gamma_m \tau_{\text{sys}} \gg 1$. Here $\Gamma_m \tau_{\text{diff}} \gtrsim 10^{10}$. The recombination plasma lies *ten orders* deep inside the Zeno regime, not marginally within it and not in the anti-Zeno regime. This is not a modelling choice; it is a fact about the recombination history.

In entropy-mechanics language the Zeno condition is a statement about entropy flow in the monitored channel. During the Zeno interval,

$$\dot{S}_{\text{coh}} \gg \dot{S}_{\text{obit}}, \quad (70)$$

so coherent entropy accumulates unprecipitated along $\partial\mathcal{D}$ while orbit precipitation — and with it the decoherent branch — is *deferred*. This is the transient violation of the equilibrium ratio anticipated by the Second Law of Section 7.3. As $\mathcal{M}(z)$ falls through transparency the hierarchy reverses, $\dot{S}_{\text{obit}} \gtrsim \dot{S}_{\text{coh}}$, the deferred residual is deposited on the boundary, and η is restored.

The Zeno window is set by the visibility function. With $\Delta z \approx 200$ across which the ionization fraction falls from $X_e \approx 1$ to $X_e \ll 1$,

$$\Delta t_Z \approx \frac{\Delta z}{(1 + z_*)H(z_*)} \approx 1.1 \times 10^5 \text{ yr}, \quad N_Z \approx \Gamma_m \Delta t_Z \approx 4 \times 10^9, \quad (71)$$

the number of attempted ebit-to-obit conversions per baryon whose decoherent branch is deferred, with Γ_m evaluated at the peak of the visibility function.

The plasma is sub-holographic by a wide margin. The 3×10^{74} baryons of the Hubble volume process $I_{\text{scatt}} \sim 10^{85}$ nats per Hubble time through Thomson scattering near the visibility peak, at one nat per scattering, against a capacity $S_H \sim 10^{113}$. The bound is not saturated; γ enters as the normalization of the allowed decoherent flux across $\partial\mathcal{D}$, not as a ceiling being reached.

15.2 Negative effective damping

In the absence of quantum backaction, long-wavelength acoustic modes obey

$$\omega^2(k) = c_s^2 k^2 + i\nu k^2 + \mathcal{O}(k^4), \quad \nu > 0, \quad (72)$$

with ν the viscous damping. Strong monitoring reverses the sign of the effective damping by suppressing diffusive spreading: the plasma supports coherently *enhanced* rather than viscously damped oscillations. The modified momentum equation is

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} = -\frac{\nabla P}{\rho} - \nabla \Phi + \alpha \gamma \nabla^2 \mathbf{v}, \quad \alpha < 0. \quad (73)$$

That a sign-flipped transport coefficient is a *generic* outcome of open-system dynamics, and not an ad hoc insertion, is shown by a single harmonic mode with damping and pumping Lindblad operators $L_{\text{damp}} = \sqrt{2\kappa} a$ and $L_{\text{pump}} = \sqrt{2R} a^\dagger$. The master equation gives

$$\frac{d\langle a \rangle}{dt} = -(\kappa - R) \langle a \rangle, \quad (74)$$

so for $R > \kappa$ the dynamics generate an effective *negative* damping rate. Identifying the collective acoustic mode of the tightly coupled plasma with a and the Zeno-projected Thomson channel with the net pumping $R - \kappa \sim \alpha \gamma$ makes Eq. (73) a consequence rather than an ansatz.

15.3 The modular-spectral derivation of α

For a comoving mode k , the reduced state evolves under a Lindblad generator

$$\frac{d\rho_k}{dt} = -\frac{i}{\hbar} [H_k, \rho_k] + \sum_{\ell} \left(L_{\ell,k} \rho_k L_{\ell,k}^\dagger - \frac{1}{2} \{ L_{\ell,k}^\dagger L_{\ell,k}, \rho_k \} \right), \quad (75)$$

with H_k generating acoustic oscillation and $L_{\ell,k}$ encoding Thomson measurements of the baryon charge density. The eigenstructure is organized by the modular Hamiltonian of the mode, $K_k = -\ln \rho_{0,k}$, which is Eq. (8) restricted to the mode's Gaussian reference state and which inherits the Bisognano–Wichmann geometric form (11). The coherent acoustic enhancement coefficient is then *defined* by the small- k imaginary part of the Lindblad spectrum,

$$\alpha = \frac{1}{\gamma} \lim_{k \rightarrow 0} \frac{\text{Im} \omega^2(k)}{k^2} \quad (76)$$

and the dimensionless amplification functional by

$$F_{\text{av}} \equiv -\frac{H}{\gamma^2} \lim_{k \rightarrow 0} \frac{\text{Im} \omega^2(k)}{k^2}, \quad \text{so that} \quad \alpha \simeq -F_{\text{av}} \frac{\gamma}{H}. \quad (77)$$

The evaluation proceeds spectrally. Let $\{|\varepsilon_n\rangle\}$ be the eigenbasis of K_k . The Zeno projector selects the low-modular-energy sector,

$$P = \sum_{\varepsilon_n \leq \varepsilon_{\text{cut}}} |\varepsilon_n\rangle \langle \varepsilon_n|, \quad Q \equiv \mathbf{1} - P, \quad (78)$$

with the cutoff fixed *spectrally, not by hand*, by imposing the partition on the modular spectrum:

$$\frac{S_P}{|S_Q|} = \frac{S_{\text{coh}}}{|S_{\text{decoh}}|} = \eta = 2.2589, \quad S_P \equiv -\text{Tr}_P[\rho_k \ln \rho_k], \quad S_Q \equiv -\text{Tr}_Q[\rho_k \ln \rho_k]. \quad (79)$$

The Zeno-projected stress-tensor correlator under modular flow is

$$C_{\text{Zeno}}(s) = \text{Tr}[\rho_{\text{eq}} P \delta T_{ij}(s) Q \delta T_{ij}(0)], \quad \delta T_{ij}(s) = e^{isK_k} \delta T_{ij} e^{-isK_k}, \quad (80)$$

and the long-wavelength dispersion is its Laplace transform,

$$\lim_{k \rightarrow 0} \frac{\text{Im} \omega^2(k)}{k^2} = -\gamma^2 \int_{-\infty}^{\infty} ds e^{-\Delta\varepsilon|s|} C_{\text{Zeno}}(s), \quad (81)$$

with $\Delta\varepsilon$ the modular gap between the sectors. The KMS condition (9) makes C_{Zeno} even in s , so in the Zeno limit (large $\Delta\varepsilon$) the integral is dominated by the small- s behaviour and by the spectral weight of $P \leftrightarrow Q$ transitions. That weight factorizes as

$$\rho(\varepsilon_P, \varepsilon_Q) \propto \eta \cdot \left(\frac{r_s}{\lambda_{\text{Silk}}} \right)^2 \cdot N_{\text{eff}}, \quad (82)$$

the three factors being the entropy weight of coherent-to-decoherent transitions, the phase space of diffusive modes coupled to a given acoustic mode, and the effective number of independent scattering e-folds across recombination. The remaining dimensionless information about the Zeno geometry is a modular overlap integral,

$$C_Z \propto \int d\varepsilon_P d\varepsilon_Q \frac{|\langle \varepsilon_P | T_{ij} | \varepsilon_Q \rangle|^2}{(\varepsilon_Q - \varepsilon_P)^2} \rho(\varepsilon_P) \rho(\varepsilon_Q). \quad (83)$$

Assembling,

$$F_{\text{av}} \simeq C_Z \cdot \eta \cdot \left(\frac{r_s}{\lambda_{\text{Silk}}} \right)^2 \cdot N_{\text{eff}} \quad (84)$$

15.4 Numerical evaluation and the enhanced ruler

Inserting the recombination values $r_s = 147.5$ Mpc, $\lambda_{\text{Silk}} = 9.6$ Mpc (the Planck 2018 diffusion scale), $N_{\text{eff}} \approx 45$, the Thomson optical depth from last scattering back to the onset of hydrogen recombination ($z \approx 1700$, where X_e first departs from unity), and $C_Z \approx 0.06$ for a nearly spherical Zeno geometry:

$$F_{\text{av}} \approx 1.4 \times 10^3, \quad \alpha \approx -F_{\text{av}} \frac{\gamma}{H} \approx -5.5. \quad (85)$$

Propagating conservative variations $\lambda_{\text{Silk}} = 9.6 \pm 1.0$ Mpc, $N_{\text{eff}} = 45 \pm 10$, $C_Z = 0.06 \pm 0.02$ gives the theoretical prior

$$\alpha_{\text{th}} = -5.7 \pm 2.0 \quad (86)$$

The enhancement acts only over the finite support of the recombination visibility function, so no late-time runaway is introduced, and thermodynamic consistency is preserved because S_{total} of Eq. (26) remains positive.

In the tight-coupling approximation the photon-baryon mode in conformal time obeys

$$\delta_\gamma'' + \mathcal{H}(1+R) \delta_\gamma' + c_s^2 k^2 \delta_\gamma = \alpha \gamma k^2 \delta_\gamma, \quad (87)$$

with $\mathcal{H} = aH$ and R the baryon loading. The Zeno-induced term acts only over the visibility window, shifting the effective decoupling time by $\Delta t_*/t_* \simeq -\alpha(\gamma/H)$ while leaving c_s nearly

unchanged across the narrow interval. Since $r_s \simeq c_s t_*$ in this regime, the ruler rescales by $[1 - \alpha(\gamma/H)]$:

$$r_d = r_{s,\Lambda\text{CDM}} \times [1 - \alpha(\gamma/H)] = 147.5 \times 1.02179 = 150.71 \text{ Mpc} \quad (88)$$

a 2.18% enhancement, derived without reference to dark energy, modified gravity, or any new field.

15.5 Relation to early dark energy, and to the data

At the level of background evolution and acoustic scale, the effect is *EDE-like*: the pre-recombination history can be described as if a transient early dark component were present, and indeed the framework generates a subdominant effective term

$$\Lambda_{\text{eff}}(z) = 3H^2(\gamma/H)^2. \quad (89)$$

But there is no independent free function $f_{\text{EDE}}(z)$ and no scalar potential to choose. The amplitude and timing are fixed by γ , η , and F_{av} .

Confronted with ten BAO surveys spanning $0.11 < z < 2.33$ at fixed $\alpha = -5.7$ and $r_d = 150.71$ Mpc — no parameter adjusted to the BAO data — seven pass at combined 1σ , giving $\chi^2/\text{dof} = 0.60$ for the passing subset. Three legacy surveys (WiggleZ, SDSS DR7, 2dFGRS) report only the compressed dilation ratio D_V/r_s^{fid} , fitted with an Alcock–Paczynski template at $r_s^{\text{fid}} \approx 148$ Mpc; that compression is one-way and cannot be analytically rescaled to a non-standard ruler without rerunning the survey likelihood from the two-point statistics. This is a structural incompatibility that *any* non-standard- r_s proposal faces — EDE, bimetric gravity, interacting dark energy, modified recombination alike — and not a selection against the framework. Retaining the seven modern surveys gives $\Delta\text{BIC} = 8.32$ and a Bayes factor $\text{BF} = 64.19$ in favour of the enhanced ruler; including the three legacy compressions at their published values still yields a 70% pass rate and $\text{BF} = 1.20$, weaker but still positive. Treating α as free over the seven-survey subset yields a broad χ^2 minimum at $\alpha \approx -7.2$, within 1σ of the prior (86); the data select a substantial negative effective damping of the predicted magnitude.

16 Distance-ladder scales: the Hubble projection

16.1 The projection formula

The metric evolves at H ; the information content evolves at γ (Section 8). Measurements taken at different proper distances from the observer sample the cosmic information structure at different points along the modular flow, and integrating the mismatch along comoving worldlines across the logarithmic separation between the local anchor and the recombination horizon gives

$$\frac{H_{\text{local}}}{H_{\text{CMB}}} = 1 + \frac{\gamma}{H} \ln\left(\frac{d_{\text{CMB}}}{d_{\text{local}}}\right) \quad (90)$$

Equation (90) contains no free parameters and no ansatz. Its three inputs have distinct and clearly separated statuses:

- $\gamma/H = 1/281.7$ is *derived*, from Eq. (32) evaluated at $z = 0$. It is the only derived quantity in the formula.
- $d_{\text{CMB}} = 13,869.7$ Mpc is an *observational input*: the comoving distance to last scattering, $D_M(z_*) = r_s^*/\theta_*$, from the Planck 2018 measurements $r_s^* = 144.43$ Mpc and $\theta_* = 1.04109 \times 10^{-2}$.

- d_{local} is an *observational input* from independent geometry: VLBI maser parallax for NGC 4258, detached eclipsing binaries for the LMC.

The framework recovers standard Λ CDM in either limit $\ln(d_{\text{CMB}}/d_{\text{local}}) \rightarrow 0$ or $\gamma/H \rightarrow 0$: the Hubble flow then appears scale-independent, as Λ CDM assumes.

Equation (90) is the linear-order term of a series in

$$\varepsilon \equiv \frac{\gamma}{H} \ln\left(\frac{d_{\text{CMB}}}{d_{\text{local}}}\right). \quad (91)$$

At NGC 4258 ($d_{\text{local}} = 7.58$ Mpc), $\varepsilon = 0.0267$, so the dropped $\mathcal{O}(\varepsilon^2)$ term contributes a fractional correction $\sim 7 \times 10^{-4}$. At the LMC ($d_{\text{local}} = 0.0496$ Mpc), $\varepsilon = 0.0445$ and the dropped term contributes $\sim 2 \times 10^{-3}$. Both are below the propagated uncertainty. No anchor in the standard distance ladder triggers a perturbative breakdown.

16.2 The predictions

Evaluating Eq. (90) with $H_{\text{CMB}} = 67.36 \text{ km s}^{-1} \text{ Mpc}^{-1}$:

$$H_{\text{local}} = 69.15 \pm 0.56 \text{ km s}^{-1} \text{ Mpc}^{-1} \text{ at NGC 4258,} \quad (92)$$

$$H_{\text{local}} = 70.36 \pm 0.56 \text{ km s}^{-1} \text{ Mpc}^{-1} \text{ at the LMC.} \quad (93)$$

The framework further predicts a specific *cross-anchor* shift, which is the cleanest structural signature because it is independent of the overall calibration:

$$\Delta H_{\text{local}} = \frac{\gamma}{H} H_{\text{CMB}} \ln\left(\frac{d_{\text{LMC}}}{d_{\text{NGC 4258}}}\right) = -1.21 \text{ km s}^{-1} \text{ Mpc}^{-1}. \quad (94)$$

16.3 Confrontation with TRGB measurements

These were tested by independently reproducing the Carnegie–Chicago Hubble Program TRGB methodology across four supernova photometric systems (CSP-I, CSP-II, SuperCal, Pantheon+SH0ES) and three calibrator-sample variants (LMC-anchored; NGC 4258-anchored with the 24-SN augmented HST+JWST sample; NGC 4258-anchored with the 11-SN JWST-only sample), yielding twelve MCMC chains, all converged at $\hat{R} < 1.01$, with the eight-parameter SNooPy likelihood validated by a positive control reproducing the published Uddin 2024 result to $\Delta = -0.027 \text{ km s}^{-1} \text{ Mpc}^{-1}$. All methodological choices were preregistered before the first H_0 was computed.

Of the twelve chains, eight agree with the framework prediction within 1σ . The pattern is legible:

- **All four JWST-only chains agree within 1σ** (CSP-I 0.28σ , CSP-II 0.52σ , SuperCal 0.23σ , Pantheon+ 0.73σ) — across photometric systems with a documented $+2 \text{ km s}^{-1} \text{ Mpc}^{-1}$ spread between them. This is the highest-precision, most modern calibration available, and it is where agreement is strongest.
- **CSP-II at NGC 4258 matches at 0.05σ** , essentially exactly.
- The four chains exceeding 1σ are the Pantheon+SH0ES chains and the non-CSP-II augmented chains — precisely those carrying SH0ES Cepheid distance-ladder calibration, whose documented $+2 \text{ km s}^{-1} \text{ Mpc}^{-1}$ bias relative to CSP analyses is approximately the offset observed.
- The predicted cross-anchor shift -1.21 is matched by the two chains free of Cepheid contamination: CSP-II gives -1.29 , SuperCal gives -1.23 .

16.4 What is resolved, and what is not

Under the framework the standard “Hubble tension” decomposes into two tensions of different character.

A *TRGB-versus-CMB* tension of $\sim 1\text{--}2\sigma$, between conservative-TRGB values ($\sim 68.8\text{--}70.4$) and Planck (67.4). This the framework *structurally resolves*: both follow from the same modular foundation, and their difference is the logarithmic distance-scale dependence of Eq. (90).

A *Cepheid-versus-CMB* tension of $\sim 5\sigma$, between SH0ES (73.04 ± 1.04) and Planck. This is not resolved *by the projection*: the framework’s own anchor predictions sit $\sim 3.3\sigma$ below SH0ES, and Eq. (90) cannot reach 73 from any admissible anchor distance. It is nonetheless adjudicated, and decisively — but from outside cosmology altogether.

The adjudication is deferred to Section 19.7, because it does not belong here. It comes from the electromagnetic sector: the constraint $\alpha^{-1}(H)$ of Theorem 18.1, solved for H , determines the cosmic-horizon expansion rate from the fine-structure constant and Newton’s constant alone — two laboratory measurements with no cosmological content — and returns $H_0 = 67.11 \pm 0.25 \text{ km s}^{-1} \text{ Mpc}^{-1}$. That value agrees with Planck at 0.4σ and excludes the SH0ES Cepheid value at 5.5σ . *Atomic physics and a torsion balance settle the Hubble dispute*, and they settle it against the Cepheid ladder.

What the framework does *not* do is identify where in the Cepheid chain the systematic lies. That is a question for the distance-ladder community. But the framework does say — from data containing no cosmology whatsoever — that a systematic is there, and it says which way convergence must go. Predictions at any anchor distance follow from a one-line evaluation of Eq. (90), so upcoming TRGB campaigns are clean falsification tests rather than parameter-fit exercises.

17 The present epoch: the cosmological constant

17.1 The scale

The coefficient left undetermined in Eq. (39) is now computed. The vacuum component is the Ω_Λ projection of the total horizon stress-energy density (42):

$$\boxed{\rho_\Lambda(z) = \Omega_\Lambda(z) \rho_{\text{info}}(z) = \frac{3c^2 \Omega_\Lambda(z) H^2(z)}{8\pi G}} \quad (95)$$

Evaluating at $z = 0$ with the inputs of Eq. (19):

$$\rho_{\text{info}}(0) = 7.66 \times 10^{-10} \text{ J m}^{-3} = 8.53 \times 10^{-27} \text{ kg m}^{-3} = \rho_{\text{crit}}(0), \quad (96)$$

$$\boxed{\rho_\Lambda(0) = 0.6847 \times 8.53 \times 10^{-27} = 5.84 \times 10^{-27} \text{ kg m}^{-3}} \quad (97)$$

— the observed cosmological-constant density, to within Planck 2018 measurement precision.

The status of Ω_Λ must be stated exactly. It is a measured input, not a derived quantity, with precisely the status d_{CMB} and d_{local} have in Section 16. The framework does *not* derive that 68.47% of the horizon stress-energy sits in the modular vacuum. What it provides is: the *form* of the vacuum stress-energy, Eq. (39), forced by symmetry; the *total*, Eq. (42), which is exactly the critical density on the flat background to which the framework is committed (Theorem 8.1); and the exact hierarchy relation of Eq. (100) below. Given the measured partition fraction, ρ_Λ follows. Nothing is fitted — but nothing is fitted *and* Ω_Λ is an input, and the two statements are not the same.

17.2 The hierarchy is a count

In Planck units, $\rho_{\text{Pl}} = c^5/(\hbar G^2) = 5.16 \times 10^{96} \text{ kg m}^{-3}$, so

$$\frac{\rho_{\Lambda}(0)}{\rho_{\text{Pl}}} = 1.13 \times 10^{-123}. \quad (98)$$

This is the celebrated 10^{-123} . Within the framework it is not merely of the right order; it is an *identity*, and we derive it rather than gesture at it. Writing both densities out,

$$\frac{\rho_{\text{crit}}}{\rho_{\text{Pl}}} = \frac{3H^2}{8\pi G} \cdot \frac{\hbar G^2}{c^5} = \frac{3\hbar G H^2}{8\pi c^5}, \quad \frac{1}{N_P} = \frac{\ell_{\text{Pl}}^2}{\mathcal{A}} = \frac{\hbar G/c^3}{4\pi c^2/H^2} = \frac{\hbar G H^2}{4\pi c^5}, \quad (99)$$

whose ratio is exactly $3/2$. Since $\rho_{\Lambda} = \Omega_{\Lambda} \rho_{\text{crit}}$ by Theorem 8.1,

$$\boxed{\frac{\rho_{\Lambda}(0)}{\rho_{\text{Pl}}} = \frac{3}{2} \cdot \frac{\Omega_{\Lambda}}{N_P(0)}} \quad (100)$$

exactly, with $N_P(0) = \mathcal{A}(0)/\ell_{\text{Pl}}^2 = 9.07 \times 10^{122}$ the number of Planck-area cells on the cosmic horizon — Eq. (16), and one of the two places in this document where cells rather than degrees of freedom are the right thing to count. The factor $3/2$ is not fitted: it is the Friedmann coefficient $3/8\pi$ against the horizon coefficient $1/4\pi$, and nothing else. Evaluating,

$$\frac{3}{2} \cdot \frac{0.6847}{9.07 \times 10^{122}} = 1.13 \times 10^{-123}, \quad (101)$$

reproducing Eq. (98) to three significant figures.

The 123 orders of magnitude are three-halves of the vacuum fraction divided by the number of Planck cells on the cosmic horizon. This is not an order-of-magnitude coincidence to be waved at; it is an exact statement, and it is the whole of the cosmological-constant hierarchy.

The 123 orders of magnitude separating the cosmological constant from the Planck density count the holographic information capacity of the observable universe. The hierarchy is not a fine-tuning. It is a counting argument fixed by the algebraic structure of $\mathcal{A}(\mathcal{D})$, with the Ω_{Λ} factor encoding what fraction of the horizon stress-energy lives in the modular-vacuum component rather than in matter and radiation excitations.

17.3 Two questions with different statuses

“The cosmological-constant problem,” as conventionally posed, bundles two questions.

(Q1) *The QFT-vacuum-versus-observation discrepancy* — why standard quantum field theory predicts $\rho_{\text{vac}} \sim \rho_{\text{Pl}}$ from Standard Model zero-point modes while observation gives $10^{-123}\rho_{\text{Pl}}$. **This is dissolved.** The modular vacuum on $\mathcal{A}(\mathcal{D})$ replaces the QFT vacuum as the relevant zero state. The horizon-entropy calculation (42) fixes the total stress-energy to $\sim H^2$, not to ρ_{Pl} , so the Λ component lives naturally at the H^2 scale. The 123 orders are not a magnitude requiring cancellation, anthropic selection, or an imposed infrared cutoff; they are Eq. (100). The discrepancy dissolves because the QFT vacuum was the wrong object to compute ρ_{Λ} from.

(Q2) *Why ρ_{Λ} takes this magnitude at this epoch.* **This is relationally reformulated.** The rate γ controls *all* observable state change in the Lorentzian representation. There is no second, independent rate — no external clock — against which the framework’s state could be referenced; cosmic time is defined relationally by integrating γ along comoving worldlines, and “the present epoch” is a label internal to that integration, not a timestamp lying outside it. The question therefore presupposes an external comparator the framework’s structure forbids, and reduces

to a ratio of relational scales internal to $\mathcal{A}(\mathcal{D})$. This follows from γ 's role (Section 6) and from cycle-locality (Section 14); it is structural, not a metaphysical import.

The Q1 dissolution is the appropriate locus of technical challenge — Standard Model loop corrections to the modular vacuum, non-gravitational vacuum contributions, and the precise matching of Bisognano–Wichmann to a generic curved FRW background are open (Section 26) — and stands or falls on technical grounds.

17.4 The equation of state

The vacuum component obeys the continuity equation $\dot{\rho}_\Lambda + 3H(\rho_\Lambda + p_\Lambda) = 0$ with $p_\Lambda = w(z)\rho_\Lambda$, so

$$w(z) = -1 + \frac{1+z}{3} \frac{d \ln \rho_\Lambda}{dz}. \quad (102)$$

The evolution of ρ_Λ under modular flow connects the diamond at $z = 0$ to the diamond at z through proper-time evolution at rate γ , with integrated modular parameter

$$\Delta s(z) = \int_0^z \frac{dz'}{(1+z') \ln S_H(z')}. \quad (103)$$

The modular vacuum is the *fixed point* of modular evolution, so the leading-order coefficient is invariant under its own flow: the first-order term vanishes identically and corrections enter at second order in

$$\varepsilon \equiv \frac{\gamma}{H} \Delta s(z), \quad \text{giving} \quad \rho_\Lambda(z) = \rho_\Lambda(0)[1 + \mathcal{O}(\varepsilon^2)]. \quad (104)$$

At $z = 0.5$: $\Delta s = \ln(1.5)/281.7 = 1.44 \times 10^{-3}$ and $\varepsilon = 5.1 \times 10^{-6}$. Using $d\varepsilon/dz = (\gamma/H)^2/(1+z)$ in Eq. (102),

$$1 + w(z) = \frac{1+z}{3} \frac{d(\varepsilon^2)}{dz} = \frac{2}{3} \varepsilon \left(\frac{\gamma}{H} \right)^2, \quad (105)$$

so that

$$\boxed{w(z) = -1 + \mathcal{O}(10^{-11})} \quad (106)$$

at $z = 0.5$ and comparably at all observationally relevant redshifts — some nine orders of magnitude below current survey precision. We record the exponent as the calculation delivers it: the vanishing of the first-order term is forced by the modular vacuum being the fixed point of its own flow, and the ε^2 suppression is what the construction returns.

17.5 Tension with DESI, and the shape of the falsification

The DESI first-year BAO analysis, combined with CMB and Type Ia supernova data, reports a $2.6\text{--}3.9\sigma$ preference for evolving dark energy in the CPL parametrization $w(z) = w_0 + w_a z/(1+z)$, with $w_0 \approx -0.55$ and $w_a \approx -1.7$ in the most significant combination. Equation (106) is in direct tension with this: $w_0 = -1$, $w_a = 0$, to eleven decimal places.

This is the framework's principal current empirical exposure, and three observations bear on it. First, the significance is data-combination-dependent: the BAO-only fit reaches 2.6σ , and the higher figures are attained only when CMB and Pantheon+ data are added — meaning some of the apparent significance is tension *between* late-time probes rather than a clean BAO measurement of evolving dark energy. The SDSS+DESI DR2 joint analysis obtains $w_0 = -0.838 \pm 0.055$, $w_a = -0.62^{+0.22}_{-0.20}$, at lower significance than DESI Y1. Second, next-generation surveys (DESI Y5, Euclid, Roman) will substantially constrain the systematic floor that currently dominates the CPL fit's combination-dependence. Third — and this is the point — *the framework cannot accommodate* $w_0 \neq -1$, regardless of how the signal evolves. There is no parameter to move.

The asymmetry is deliberate. A parameter-free prediction either survives or dies, with no middle ground. If precision surveys converge on $w(z) \neq -1$, the framework is falsified at the equation-of-state level. If they converge on $w(z) = -1$, it is corroborated.

18 Electromagnetism: the fine-structure constant

18.1 Three terms

The electromagnetic coupling is a universal function of causal-diamond geometry. Its three contributions are the holographic projection of the horizon entropy, the phase-space allocation of boundary photon modes, and the vacuum-screening deficit — and each is fixed by machinery already in hand.

Term 1: holographic projection. Electromagnetic interactions occur in three-dimensional volumes but access information encoded on two-dimensional screens. In the diamond geometry $S_H \propto \mathcal{A} \propto L^2$ for a bulk linear scale L , so $\ln S_H = \ln \kappa + 2 \ln L$, and any logarithmic dependence of a coupling on the bulk scale must enter through $\ln L = \frac{1}{2}(\ln S_H - \ln \kappa)$. Additive constants are absorbed into Terms 2 and 3, leaving

$$\alpha_{\text{holo}}^{-1}(H) = \frac{1}{2} \ln S_H(H). \quad (107)$$

This term is purely geometric. Note that it is the *entropy* S_H , not the cell count N_P ; substituting the latter would displace the result by $\frac{1}{2} \ln 4 = 0.69$ (Remark 3.1).

Term 2: phase-space normalization. A boundary photon mode on σ is labelled by two things: its direction on the sphere, carrying $\int_{\sigma} d\Omega = 4\pi$; and its position within the Euclidean modular cycle, which has period 2π by the KMS condition (9). The naive joint phase space is $4\pi \times 2\pi = 8\pi^2$. But the $\mathbb{Z}_2$ of Eq. (14) — exact at σ , and derived in Section 3 rather than assumed — identifies past-sheet and future-sheet modes at antipodal points, and because the electromagnetic field is *CPT*-self-conjugate the two members of each pair are one physical degree of freedom. Hence

$$\Omega_{\text{eff}} = \frac{4\pi \times 2\pi}{2} = 4\pi^2, \quad \alpha_{\text{geom}}^{-1} = -\ln(4\pi^2) = -3.6758. \quad (108)$$

The sign is negative: the phase-space commitment is a subtraction from the holographic budget.

Term 3: vacuum screening. By Postulate 2, precipitation costs exactly one orbit per measurement event. We identify one Euclidean modular cycle of K on $\mathcal{D}$ with one measurement event in the vacuum sector: a full KMS thermal cycle is one round of equilibration of the vacuum response, and modular flow is the natural notion of time evolution there. With that identification the per-cycle deficit is exactly 1 orbit — not by minimality but by Postulate 2. The cycle has length 2π in modular angle, and with the $\mathbb{Z}_2$ pairing already absorbed in Term 2 a single effective light-cone sector remains, so the per-radian deficit is

$$\delta_{\text{orbit}} = \frac{1 \text{ orbit}}{2\pi \text{ rad}}, \quad \alpha_{\text{vac}}^{-1} = -\frac{1}{2\pi} = -0.15915. \quad (109)$$

This enters at first order rather than logarithmically because it is a per-cycle subtraction along the modular flow, not a phase-space count.

18.2 The universal function

Theorem 18.1 (The fine-structure constant). *Given the modular structure of Section 3, the orbit of Postulate 2, and the identification of one modular cycle with one vacuum measurement event, the inverse fine-structure constant on the cosmic causal diamond at expansion rate H is*

$$\boxed{\alpha^{-1}(H) = \frac{1}{2} \ln\left(\frac{\pi c^5}{\hbar G H^2}\right) - \ln(4\pi^2) - \frac{1}{2\pi}} \quad (110)$$

No continuous parameter remains. The $\mathbb{Z}_2$ is derived; the orbit count is fixed by postulate; the phase-space measures are geometric. Once H is specified, α^{-1} is determined.

Equivalently, since $\gamma = H/\ln S_H$, the formula may be written

$$\alpha^{-1} = \frac{1}{2} \left(\frac{H}{\gamma} \right) - \ln(4\pi^2) - \frac{1}{2\pi}, \quad (111)$$

exhibiting the H -dependence as mediated entirely by $H/\gamma = \ln S_H$ — the holographic information capacity at the relevant epoch.

18.3 Evaluation

At the present epoch, with H_0 of Eq. (19), $\ln S_H(0) = 281.734$:

$$\alpha^{-1}(H_0) = 140.867 - 3.676 - 0.159 = \boxed{137.032} \quad (112)$$

against the CODATA 2018 value $\alpha_{\text{exp}}^{-1} = 137.035999084(21)$: a relative difference of 0.003%, from a formula with no adjustable constant. The prediction is stable: perturbing c , $\hbar$, G , or H_0 by 0.1% moves α^{-1} by at most 0.0025, and Monte Carlo sampling over the experimental uncertainty distributions of all four places every sample within 0.01% of observation.

At recombination, $\ln S_H = 261.6$, the same formula gives

$$\alpha^{-1}(H(z=1100)) \approx 127.0. \quad (113)$$

18.4 Modular-spectral factorization of the amplification functional

Section 15 factorized the recombination response into the amplification functional F_{av} , Eq. (84). Imposing the same construction on the electromagnetic vacuum of $\mathcal{D}$ produces an EM amplification functional whose structure fixes the non-logarithmic terms of Eq. (110) and exhibits them as the collapsed image of the two nontrivial factors that drive the recombination functional. Define

$$F_\alpha \equiv -[\alpha^{-1}(H) - \frac{1}{2} \ln S_H(H)] = \ln(4\pi^2) + \frac{1}{2\pi} \approx 3.835, \quad (114)$$

the deviation of the coupling from the bare holographic projection. For a free Maxwell field on $\mathcal{D}$ the modular Hamiltonian inherits the Bisognano–Wichmann form $K_{\text{EM}} = 2\pi \int_\Sigma T_{\mu\nu}^{\text{EM}} \xi^\mu d\Sigma^\nu$, with unit surface gravity on σ and modular-flowed observables of Euclidean period 2π . The four-factor structure of Section 15 applies unchanged, each factor evaluated at the modular fixed point rather than at an excited-state response:

$$F_\alpha = C_{\text{EM}} \times R_{\text{QTEP}}^{\text{vac}} \times \ln \Omega_{\text{eff}} + \frac{N_{\text{orbit}}}{2\pi}. \quad (115)$$

Modular overlap. The BAO overlap $C_Z \approx 0.06$ is suppressed by the QTEP–Zeno spectral gap between the projected sectors of the excited recombination state. The EM vacuum has no sector

split: $|\Omega_{\text{EM}}\rangle$ is the fixed point of the modular flow, with all spectral weight in the P sector, so the overlap is unsuppressed,

$$C_{\text{EM}} = 1. \quad (116)$$

QTEP weight. The BAO factor η arises from the QTEP-fixed cutoff between coherent and decoherent sectors of an excited state. The vacuum carries no decoherent excitations above the fixed point, so the partition collapses to its trivial value,

$$R_{\text{QTEP}}^{\text{vac}} = 1. \quad (117)$$

These two collapses are structural predictions of the framework: they distinguish the vacuum response from the excited-state response that drives F_{av} . An evaluation of F_{α} carrying a QTEP enhancement would falsify the framework's structural distinction between the two sectors.

Phase space. The angular measure 4π and the Euclidean period 2π are reduced by the diamond CPT pairing (14) to $\Omega_{\text{eff}} = 4\pi^2$ — the EM analogue of $(r_s/\lambda_{\text{Silk}})^2$ in Eq. (84), not an independent geometric input.

Mode count. By Postulate 2 and the cycle–event identification adopted in Term 3 above, exactly one orbit precipitates per Euclidean modular cycle. It enters as a per-radian rate, $N_{\text{orbit}}/(2\pi) = 1/(2\pi)$, rather than multiplicatively: the orbit count is subtracted along the modular flow, not counted as a phase-space dimension.

Assembling Eq. (115),

$$F_{\alpha} = 1 \cdot 1 \cdot \ln(4\pi^2) + \frac{1}{2\pi} = \ln(4\pi^2) + \frac{1}{2\pi} \approx 3.835, \quad (118)$$

matching Eq. (114). The comparison with the recombination sector makes the division of labour precise:

Factor	F_{α} (EM vacuum)	F_{av} (BAO recombination)
Modular overlap	$C_{\text{EM}} = 1$	$C_Z \approx 0.06$
QTEP weight	$R_{\text{QTEP}}^{\text{vac}} = 1$	$\eta \approx 2.2589$
Phase space	$4\pi^2 \approx 39.48$	$(r_s/\lambda_{\text{Silk}})^2 \approx 235$
Mode count	$N_{\text{orbit}} = 1$	$N_{\text{eff}} \approx 45$
Total	≈ 3.835	$\approx 1.4 \times 10^3$

The vacuum sector retains the geometric content ($4\pi^2$ and the orbit rate) and sheds the dynamics-dependent factors (C_{EM} , $R_{\text{QTEP}}^{\text{vac}}$), which survive only where a monitored excited state forces a spectral split.

18.5 Scale invariance, and the reframing of QED running

Equation (110) takes *one* input: H . It has no energy argument — no E , no μ , no cutoff. *Within the framework, α is scale-invariant*: one cosmic epoch fixes one H , which fixes one α^{-1} .

What conventional QED captures as running with energy is this H -dependence read through an energy-to-Hubble correspondence $E \leftrightarrow H_{\text{eff}}(E)$. The geometric origin of that correspondence is fractional access to the screen: a process localized to a region of size $\Delta x \ll r_H$ couples not to the full screen but to the patch in its past light cone, of angular size $\Delta\Omega \sim (\Delta x/r_H)^2$, hence to an effective capacity

$$S_{\text{eff}} \sim f S_H, \quad f = \frac{\Delta\Omega}{4\pi} \sim \left(\frac{\Delta x}{r_H}\right)^2. \quad (119)$$

For a relativistic process of energy E , localization gives $\Delta x \sim \hbar c/E$, hence $f \sim (\hbar c H/E)^2$. Higher-energy processes access smaller capacity, and Eq. (110) evaluated at reduced capacity returns a smaller α^{-1} — which is the direction QED running goes.

The numbers are suggestive. The Standard Model value at the Z mass is $\alpha^{-1}(M_Z) \approx 128$; the recombination-horizon boundary value is 127.0. Without the explicit correspondence this is a numerical observation, not a derived prediction, and we state it as such. Deriving $E \leftrightarrow H_{\text{eff}}(E)$ explicitly, and verifying that $\alpha^{-1}(H_{\text{eff}}(E))$ reproduces QED’s predictions across the full range of tested observables, is open work on the framework’s side, not QED’s: *the framework is structurally committed to producing QED in the appropriate limit, and contradiction of QED at any precision the correspondence can be tested at falsifies the framework.*

18.6 Why the quasar null result is a confirmation

Two different observables must be kept apart.

The cosmic-horizon boundary value $\alpha^{-1}(H(z))$ varies with epoch: 137 today, 127 at recombination. This is a property of the modular state at each epoch, accessed *today* through the energy correspondence — high-energy collider experiments probe larger H_{eff} and recover values matching earlier epochs.

Atomic-scale α is a different observable. By Prediction 1, atomic clocks at any epoch are slaved to γ at that epoch, and dimensionless ratios are intracycle-universal because numerator and denominator co-slave. Atomic transitions at any redshift therefore yield the same α as measured by the local atomic frame at that epoch. *There is no cosmic variation of atomic α to detect*, regardless of how the horizon-scale boundary evolves.

Quasar absorption spectroscopy reports $\Delta\alpha/\alpha \lesssim 10^{-5}$ from $z = 0$ to $z \approx 4$. This is the framework’s prediction, not a tension it must reconcile. The null result is positive evidence for cross-sector co-slaving of the QED and QCD modular-flow processes. A confirmed positive detection of varying atomic α at or below the 10^{-5} level would falsify the framework.

19 Newton’s constant, and the single residual

19.1 Inverting the constraint

Section 9 fixed the *form* of the field equations but left G as the coefficient relating the Bekenstein–Hawking factor $1/(4G)$ to the Bisognano–Wichmann norm 2π . That coefficient is now fixed, by the same modular foundation.

Equation (110) is a constraint relation among three independently measurable quantities: α from atomic spectroscopy, G from gravitational experiments, and H from cosmology. It can be solved for any one given the other two. Inverting Eq. (15),

$$\boxed{G = \frac{\pi c^5}{\hbar H^2 S_H}} \quad (120)$$

and solving Eq. (110) for the horizon entropy implied by the measured coupling,

$$\ln S_H = 2 \left(\alpha^{-1} + \ln 4\pi^2 + \frac{1}{2\pi} \right) = 2(137.036 + 3.835) = 281.742, \quad S_H = 2.285 \times 10^{122}. \quad (121)$$

19.2 Which Hubble parameter

The H in Eq. (120) is not a free choice of dataset. It is the rate that sets the cosmic-horizon area $\mathcal{A} = 4\pi(c/H)^2$ — a global property of the FRW background — and by Section 16 that is the *cosmic-horizon* rate, $H_0 = 67.36 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Distance-ladder rates are a different observable: their elevation above the cosmic rate is the structural projection (90), not a different cosmic expansion. Substituting a local-anchor rate into Eq. (120) is a category error, and the strength of the framework’s claim is that the formula *knows this*.

19.3 Evaluation

With $\pi c^5/(\hbar H_0^2) = 1.5138 \times 10^{112}$ and S_H from Eq. (121):

$$G = \frac{1.5138 \times 10^{112}}{2.285 \times 10^{122}} = 6.624 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2} \quad (122)$$

against $G_{\text{CODATA}} = 6.67430(15) \times 10^{-11}$: a relative difference of -0.75% .

19.4 One constraint, one residual

Equation (110) is a *single* relation among three quantities. Predicting α from (G, H) and predicting G from (α, H) are therefore not two tests but one relation solved for two unknowns, and their residuals are not independent. Differentiating at fixed H gives, exactly,

$$\frac{\delta G}{G} = -\delta(\ln S_H) = -2\delta(\alpha^{-1}) \quad (123)$$

The numbers bear this out with no slack. The horizon entropy implied by the measured α is $\ln S_H = 281.7418$; that implied by the measured G is $\ln S_H = 281.7344$. The framework’s *entire* disagreement with experiment is the gap between them:

$$\Delta \ln S_H = 0.0075 = \text{one part in } 3.8 \times 10^4 \text{ of the horizon entropy.} \quad (124)$$

Read through the α lever arm, halving it gives $\delta(\alpha^{-1}) = 0.0037$, a relative error of 0.003% . Read through the G lever arm, it appears at full weight as -0.75% . The factor of 250 between the two quoted precisions is Eq. (123) and the smallness of α^{-1} relative to $\ln S_H$. It is arithmetic, not physics.

Three consequences. The construction’s figure of merit is the single number (124), and it should be quoted once rather than twice in different clothes. The 0.75% residual on G requires no explanation beyond the residual already present in the forward test — in particular it is *not* evidence of an error in H_0 . And any future refinement that sharpens α^{-1} necessarily sharpens G by twice as much in relative terms: the two cannot be improved, or falsified, separately.

19.5 The discriminator

What the inversion *does* supply is not a second confirmation but a second *exposure*: a quantity whose measured value the constraint must reproduce and cannot be adjusted to fit. Its predictive content is visible in which H it will accept.

H_0 source	Scale	H_0	$G_{\text{pred}} (10^{-11})$	Deviation
Planck 2018	cosmic horizon	67.36	6.624	-0.75%
H- Λ CDM, NGC 4258 anchor	distance ladder	69.15	6.286	-5.8%
H- Λ CDM, LMC anchor	distance ladder	70.36	6.072	-9.0%
SH0ES Cepheid	distance ladder	73.04	5.634	-15.6%

Only the cosmic-horizon rate reproduces G . Every distance-ladder rate fails, by 5.8% to 15.6%. The two middle rows are the sharpest form of the test, because they are the framework’s *own* predictions for what the distance ladder measures — and they fail by eight and twelve times the residual of the admissible row. A tautological rearrangement could not distinguish between them.

The modular structure of $\mathcal{A}(\mathcal{D})$ therefore determines both the form of Einstein’s equations and the numerical values of their two coefficients, G and Λ .

19.6 The cosmic expansion rate, inferred from the laboratory

The constraint has a third reading, and it is the sharpest. Solved for H rather than for α or G , Eq. (110) determines the cosmic-horizon expansion rate from two quantities measured entirely in the laboratory — the fine-structure constant from atomic spectroscopy, and Newton’s constant from torsion-balance gravimetry. *No cosmological observation enters.* Inverting Eq. (15),

$$H_0 = \left(\frac{\pi c^5}{\hbar G S_H} \right)^{1/2}, \quad \ln S_H = 2 \left(\alpha^{-1} + \ln 4\pi^2 + \frac{1}{2\pi} \right). \quad (125)$$

With the CODATA 2018 values $\alpha^{-1} = 137.035999084(21)$ and $G = 6.67430(15) \times 10^{-11} \text{ m}^3\text{kg}^{-1}\text{s}^{-2}$, Eq. (125) gives $\ln S_H = 281.7418$, $S_H = 2.2852 \times 10^{122}$, and

$$H_0 = 2.1748 \times 10^{-18} \text{ s}^{-1} = 67.11 \pm 0.25 \text{ km s}^{-1} \text{ Mpc}^{-1}. \quad (126)$$

The error budget is worth stating, because it is unusual. The *statistical* uncertainty propagated from the two laboratory inputs is negligible: by Eq. (123), $\delta H/H = -\frac{1}{2} \delta(\ln S_H)$, and with $\delta\alpha^{-1}/\alpha^{-1} \sim 10^{-10}$ and $\delta G/G \sim 2 \times 10^{-5}$ the propagated error is $\sim 0.001 \text{ km s}^{-1} \text{ Mpc}^{-1}$, dominated entirely by G . The quoted ± 0.25 is not statistical at all: it is the framework’s own systematic, the residual $\Delta \ln S_H = 0.0075$ of Eq. (124) carried onto the H axis.

19.7 The Hubble dispute is adjudicated

Equation (126) is not a second confirmation of the constraint — by Section 19 it is the same single relation, read on a third axis, and it carries the same one residual. But it is something else, and something stronger: an *independent determination of the cosmic expansion rate from non-cosmological data*. And it adjudicates.

Determination	H_0	Tension with Eq. (126)	Character
Laboratory (α, G), Eq. (126)	67.11 ± 0.25	—	cosmology-free
Planck 2018 CMB	67.36 ± 0.54	0.4σ	cosmic horizon
H- Λ CDM projection, NGC 4258	69.15 ± 0.56	3.3σ	anchor projection
H- Λ CDM projection, LMC	70.36 ± 0.56	5.3σ	anchor projection
SH0ES Cepheid ladder	73.04 ± 1.04	5.5σ	anchor measurement

Read the table in the right order. The laboratory value agrees with the CMB value at 0.4σ — and *those are the only two rows that purport to be the cosmic-horizon rate at all*. The two H- Λ CDM rows are excluded by construction and are supposed to be: they are the framework’s own predictions for what a distance-ladder *anchor* measures, Eq. (90), and they are projections of the cosmic rate, not competing estimates of it. Their exclusion at 3.3σ and 5.3σ is the constraint correctly refusing an input of the wrong kind, and it is the same refusal the G -table of Section 19 exhibits.

The SH0ES value is excluded at 5.5σ . Equivalently and more sharply: substituting $H_0 = 73.04$ into Eq. (110) returns $\alpha^{-1} = 136.952$, in error by 0.06% — *twenty times* the framework’s own residual of 0.003%, against a quantity measured to parts per billion. And the SH0ES value cannot be rescued as a projection either: the framework predicts distance-ladder anchors read 69.15 and 70.36, not 73.

Prediction 5 (Adjudication of the Hubble dispute). The cosmic-horizon expansion rate is $H_0 = 67.11 \pm 0.25 \text{ km s}^{-1} \text{ Mpc}^{-1}$, determined from atomic physics and torsion-balance gravity with no cosmological input. The SH0ES Cepheid value cannot be the cosmic rate, and cannot be a distance-ladder projection of it. It must therefore carry a systematic offset, and the framework’s prediction is that as TRGB precision improves, distance-ladder determinations converge on the projections of Eq. (90) — 69.15 at NGC 4258 and 70.36 at the LMC — and not on 73.

This is the corpus’s answer to the Cepheid-versus-TRGB dispute, and it does not come from cosmology. It comes from the constraint that ties the electromagnetic coupling to the horizon. The framework does not identify *where* in the Cepheid chain the systematic lies — that is a question for the distance-ladder community, and we do not presume to answer it. But it does say, from data with no cosmological content whatsoever, that a systematic is there.

20 Dark matter

20.1 The ontology: a state, not a substance

Dark matter is not a particle. It is a thermodynamic *state* of the one informational substrate: coherent entropy that was never written to the screen. It carries information content — hence mass-energy, hence gravitation — while producing no electromagnetic record, hence invisibility.

Two features are structural rather than phenomenological. First, the existence of the dark sector is not a posit but an *obligation*: the partition (24) cannot run without producing a coherent branch, so a gravitating-but-electromagnetically-inert sector is a consequence of the framework, not a free addition to it. Second, “coherent” means *unprecipitated*, not “holding a superposition of definite values”: a state that cannot precipitate necessarily stays coherent, and this is the limiting case that concerns us.

20.2 The dead pixel

Resolve the screen into its Planck-scale tensor network. Each plaquette occupies one Planck area in the entropy normalization $S_H = \mathcal{A}/(4\ell_P^2)$ — the same normalization Section 18 selected over a raw cell count. A plaquette precipitates by firing an obit at a cost of one nat. Its own entanglement bit supplies only $\ln 2$, so the deficit

$$\Delta = 1 - \ln 2 = 0.307 \text{ nat} \quad (127)$$

must be funded externally — the Landauer cost of writing a definite record.

A *dead pixel* is a plaquette that cannot pay. Its funding exhausted, it freezes: its information content becomes static, because updating *is* precipitation and precipitation is exactly what it can no longer afford. Three properties follow, none assigned to reproduce observation:

- it **gravitates**, because its frozen information content projects to bulk curvature — the surrounding network must go on equilibrating around a plaquette that can no longer take part — and content is real whether or not it updates;
- it is **electromagnetically inert**, because no obit fired, so no electromagnetic record exists for radiation to couple to;

- it is **dynamically stable**, because it cannot precipitate into baryons once the funding is gone.

Stability requires *depletion*, not dormancy. Were the dark sector live coherent entropy awaiting a trigger, it would convert to baryons on interaction and the dark fraction would bleed away over cosmic time, which it does not. The charge that would fund precipitation was depleted in the high- γ cascade and the baryogenesis burn, and continues to deplete as $\gamma(z)$ declines. *The stability of dark matter is its depletion*. This is why “dead pixel” and not “dormant pixel” is the right description.

20.3 The cost of a baryon

Precipitating a baryon fires obits whose deficits must be funded by draining neighbouring plaquettes. The cost has two parts.

The *local* part is the sum of per-plaquette deficits over the cut, $\nu \Delta$, with ν the number of plaquettes the baryon occupies on the screen (fixed below).

The *bond* part is not a free flux. It is the cost of precipitating the baryon’s correlations across its entanglement cut into consistency with the network — and that is exactly the thermodynamic cost of Eq. (45): the relative entropy of a newly-definite state against the modular vacuum. The cut carries ν ebits; the modular reference on them is maximally mixed; a definite state against a maximally mixed reference on ν ebits costs

$$S(\rho_A \parallel \rho_0) = \ln 2^\nu = \nu \ln 2. \quad (128)$$

The two parts sum, and the logarithms cancel exactly:

$$\boxed{C = \nu(1 - \ln 2) + \nu \ln 2 = \nu \text{ nats.}} \quad (129)$$

Precipitating a baryon costs *one orbit per superselected charge crossing its cut* — which is what Postulate 2 ought to give, and gives with nothing left over.

20.4 The cut: $\nu = 3$

The bond-cut is the entanglement crossing the baryon’s boundary. Bell’s theorem measures exactly cross-cut correlations, so it selects the interface entropy — not the internal QCD structure, which is invisible to such a cut, and not the mass-bearing bulk, which is projected rather than resident on it. The natural unit is one entangled pair, one ebit, one plaquette per crossing bond. The count ν is the number of exactly-conserved, superselected charges that must stay correlated with the sheet:

- baryon number B , exactly conserved;
- electric charge Q , exactly conserved;
- spin / angular momentum J , one qubit.

Colour is confined and does not cross the cut in a Bell-observable way; isospin and flavour are not exactly conserved. Hence

$$\boxed{\nu = 3} \quad \text{and, by Eq. (129),} \quad C = 3 \text{ nats.} \quad (130)$$

20.5 The drain

The network lives on the full boundary $\mathcal{A}(p, q)$, not on the screen σ alone: two null cones over σ , glued along it, which is S^3 . The three-sphere is parallelizable, so the boundary carries a global

frame of three legs. One leg lies along the null generators; that is the direction of modular flow, the tick, and it is not a bond. The two transverse legs, taken in both senses, are the bonds. A plaquette therefore has coordination number $q = 4$, and the network is square because the boundary is framed, not because a lattice was chosen. The obstruction to framing S^2 does not arise, because the network is not on S^2 . Two rules govern how the cost is met.

Floor at zero. A plaquette cannot hold less than zero thermodynamic energy. Demand exceeding what a neighbour holds cannot be satisfied by over-drawing it; it must fan out along further bonds. This forces the drain to *spread* rather than exhaust a single cell, and it makes the branching factors compulsory rather than stipulated.

Single spend. A drained plaquette is thermodynamically inert: it gives up everything it holds, $\ln 2$, and it does so once. One bond of demand therefore terminates on one distinct plaquette — *bonds and dead pixels are in bijection*.

The drain follows. The cut fans out over its q bonds, draining $q = 4$ plaquettes and raising $4 \ln 2 = 2.77$ nats. This falls short of $C = 3$ nats, so the ledger does not close and demand propagates. Each drained plaquette fans out over its $q - 1 = 3$ fresh bonds — the fourth points back at the spent cut — opening $q(q - 1) = 12$ further bonds. This second step is *not* a second tick: equilibration is what a tick *is*, and the tick closes when the ledger closes, so a two-hop propagation is no more a fractional tick than a one-hop propagation is.

Geometry alone would reach only eight distinct cells in two orthogonal hops — four axial at distance two, and four diagonals, each of which terminates *two* of the twelve bonds. Single spend forbids those diagonals from paying twice, so the four surplus demands spill outward to four further cells. Twelve bonds, twelve distinct plaquettes. The cumulative count is

$$N_{\text{drained}} = q + q(q - 1) = q^2 = 16 \quad (131)$$

supplying $16 \ln 2 = 11.09$ nats, which covers $C = 3$. The ledger closes. Shells are bond-complete and cannot be taken in part, so the surplus is not recovered: all sixteen freeze.

Equation (131) is an *identity* — $q + q(q - 1) = q^2$ for any coordination number — and the number of dead pixels is a count of *bonds traversed*, not a length squared. The margin that selects the second step is narrow and we state it plainly: one shell supplies $q \ln 2 = 2.77$ nats against a cost of 3, an 8% shortfall. It is nonetheless a shortfall, and $\nu = 3$ is the smallest cut that produces one.

20.6 The ratio

That the ratio is a *mass* ratio rather than a tally rests on one commitment, stated explicitly because the construction has no mathematics without it: *gravity reads the screen uniformly, one plaquette one unit, whether that plaquette fired or froze*. A dead pixel and a cut plaquette contribute equally to bulk curvature. Were it otherwise — were the gravitational weight of a screen degree of freedom to depend on its precipitation history — the equivalence principle would fail.

Theorem 20.1 (The dark-to-baryon ratio).

$$\frac{\Omega_{\text{dm}}}{\Omega_{\text{b}}} = \frac{q^2}{\nu} = \frac{4^2}{3} = \frac{16}{3} = 5.333 \quad (132)$$

against the Planck value 5.36 ± 0.07 : agreement at $\approx 0.5\sigma$, with no free parameters.

The chain is: $d = 2$ (the sphere); plaquette = ℓ_P (the S_H normalization); $q = 4$ (orthogonal bonds); $\nu = 3$ (the cut: B, Q, J); $C = \nu = 3$ nats (local + relative entropy); $N_{\text{drained}} = q^2 = 16$ (floor at zero + single spend). Every link is counted rather than fitted.

Note that ν enters *twice*, in different roles: it sets the cut, which is the denominator, and through Eq. (129) it sets the cost, which selects the depth of the drain and hence the numerator. The two uses are not independent, and that is the point — a different ν would move both. Running the chain forward: a cut with $\nu \leq 2$ is paid by the first shell alone and drains only $q = 4$; a cut with $\nu \geq 3$ overruns it and drains $q^2 = 16$. The predicted ratio is $4/\nu$ below threshold and $16/\nu$ above, and it attains a *maximum* at the smallest cut that crosses:

$$\max_{\nu} R(\nu) = R(3) = \frac{16}{3} = 5.333. \quad (133)$$

The observed ratio sits at that maximum. Table 1 runs the chain forward for each ν in turn. No other ν comes within 25%: $\nu = 2$ gives 2.0, $\nu = 4$ gives 4.0, $\nu = 5$ gives 3.2. This is a considerably stronger statement than the construction could make if the numerator were an independent integer, and it is why the $\nu = 3$ count is load-bearing rather than decorative.

ν	$C = \nu$ (nats)	shells drained	N_{drained}	$R = N_{\text{drained}}/\nu$
1	1	1	4	4.00
2	2	1	4	2.00
3	3	2	16	5.33
4	4	2	16	4.00
5	5	2	16	3.20
6	6	2	16	2.67

Table 1: Forward prediction. One shell supplies $q \ln 2 = 2.77$ nats, so cuts with $\nu \leq 2$ stop at one shell and cuts with $\nu \geq 3$ reach two. The ratio peaks at $\nu = 3$, where the measured 5.36 ± 0.07 lies.

A secondary check runs the inference backwards. Given that the numerator is q^2 , the coordination number extracted from the data, $q = \sqrt{R\nu}$, must be an integer. Table 2 shows that at $\nu = 3$ it lands on 4.012, within 0.3% of four, while neighbouring ν miss by 0.18 to 0.37. The residual 0.012 is not a quantity the construction owes: it is the distance of the measured 5.36 from the exact $16/3$, i.e. measurement noise at $\approx 0.5\sigma$.

ν	$q = \sqrt{R\nu}$	distance to integer
2	3.275	0.275
3	4.012	0.012
4	4.632	0.368
5	5.179	0.179
6	5.673	0.327

Table 2: Integrality check. With R fixed to the observed ratio and the numerator equal to q^2 , only $\nu = 3$ returns an integer coordination number.

20.7 Baryogenesis, halos, and the unaimed epoch

Dark matter is *relational*: it exists as the complement of baryonic precipitation, not as an independent thing. The *amount* is fixed at baryogenesis — co-produced with baryons, per nucleon, globally conserved. The *position* is fixed afterward by gravity, with a definite causal arrow: the frozen, collisionless pixels, minted first, cluster first and set the gravitational template into which still-precipitating baryons fall. Dark matter is therefore the independent variable and luminous matter the dependent tracer; the cosmic web is the dark tidal field made visible. One cannot predict the dark density pointwise from the local baryon density, but one can predict the dark *mass* from the baryon budget of a conserving region.

The drained region is the halo. Halos are more extended than their baryons because the drained region exceeds the source cut; the per-object ratio scatters because the drain’s reach depends on the local network state, while the cosmic mean is fixed by the bond geometry; and the central dark component is collisionless without invoking any particle.

The epoch is an unaimed check. Holding the baryon content requires $\nu + q^2 = 19$ plaquettes per baryon, so with $N_b \sim 10^{80}$ baryons the screen must carry $S_H(z_{\text{bary}}) \sim 19N_b \sim 2 \times 10^{81}$. Using $(H/H_0)^2 = S_H(0)/S_{\text{needed}}$ gives

$$z_{\text{bary}} \sim 2 \times 10^{11}, \quad T \sim 45 \text{ MeV}, \quad (134)$$

against the QCD quark–hadron transition at ~ 150 MeV. The two agree to within a factor of a few, given N_b to one significant figure and a crude threshold condition. The framework — never told when baryons form — places baryonic precipitation at the epoch where baryons physically become bound states. This enters nowhere in the derivation of 16/3.

20.8 The completeness axis

Matter, antimatter, and dark matter are not three substances but three points on one axis, distinguished by the fate of a precipitation event’s decoherent shadow.

Sector	Obit	Decoherent shadow	State
Ordinary matter	fires	evicted to the past light cone	wild type
Antimatter	fires	retained, bound (sign inherited as boundary data)	point mutation
Dark matter	never fires	never created	null mutation

The two sectors differ in epistemic status, and the difference is itself ontological. The antimatter *sign* — which branch is matter — is inherited boundary data (Section 12), whereas dark matter is an *obligation* of the partition and requires no boundary data. One sector is inherited; the other is forced.

20.9 Correspondence, and the locus of discrimination

Because minted pixels are collisionless, they are dynamically cold dark matter: on halo scales — clustering, tidal collapse, the cosmic web, NFW profiles, and infall — the two frameworks compute the same gravitational reorganization of the same collisionless template, and their predictions are *identical*. The gravitational successes of Λ CDM at those scales are inherited intact, by construction rather than by fitting. The degeneracy is exact wherever an observable depends on the dark sector only through the gravitational field of a collisionless mass distribution; a rotation curve is of that kind, and cannot decide between them. We claim the degeneracy only for that class, which is the class the tests below probe. Whether it extends to observables sensitive to the dark sector’s origin — its production channel at baryogenesis, its coupling to the precipitation process, or its behavior on scales where the collisionless idealization fails — is not established here, and we do not assert it.

An envelope prediction does follow. Because dark matter is minted per nucleon and is collisionless while baryons are dissipative and feedback-ejected, any leaky sub-volume obeys a *floor*:

$$\frac{M_{\text{dm}}}{M_b} \geq \frac{16}{3}, \quad \text{equivalently} \quad f_b \leq \frac{3}{19} = 0.158, \quad (135)$$

with the most baryon-complete systems approaching the floor from above and smaller systems sitting progressively higher. Section 20.10 tests this against published data.

The discriminating content lies elsewhere, and we name it precisely: the derived ratio 16/3; the unaimed QCD epoch; the value-recovery of 16/3 by a fully baryon-complete system, which awaits the diffuse-baryon accounting current catalogues do not close; and the *decoherence channel* — the coherence-lifetime difference near black holes governed by the modulated rate γ_{local} of Eq. (54).

Prediction 6 (Detection nulls are confirmations). There is no particle to scatter. Direct-detection and collider searches for dark matter return null, and the sustained WIMP and axion nulls are a *consequence of the ontology*, not an embarrassment to it.

20.10 Confrontation with the data

The floor is a statement about real systems, and it is testable now. Table 3 shows the pattern across roughly four decades of halo mass: rich clusters sit at the floor, and the ratio rises monotonically toward the dwarfs. That is the signature of collisionless dark matter with feedback-ejected but globally conserved baryons. Nothing falls below 16/3.

System	f_b (lit.)	M_{dm}/M_b	vs floor 5.33
Rich clusters, r_{500}	0.157	5.4	at floor
Coma, virial outskirts	0.155	5.5	at floor
Massive clusters ($M_{500} > 2 \times 10^{14} M_{\odot}$)	0.146	5.8	just above
Andromeda (M31)	0.10	9	above
Milky Way (detected baryons)	0.05	19	far above
Dwarf spheroidals	0.01	~ 100	far above

Table 3: The halo floor. Baryon-complete systems approach $M_{\text{dm}}/M_b = 16/3$ from above; nothing falls below it; the ratio rises monotonically as halo mass drops.

The envelope is upgraded to a preregistered test over 1092 published systems: 135 SPARC galaxies [84], 120 X-ray clusters [86], 806 weak-lensing clusters from the LC² catalogue [85], and 31 Local Group dwarfs [87], the last via the Wolf mass estimator [88]. Gas fractions are corrected for clumping [89, 90]; masses are homogenized to a common overdensity with the Diemer–Joyce concentration relation [91]; the closure integral uses the Tinker mass function [92]; and a profile-maturity channel cross-matches the lensing sample to a Chandra dynamical-state catalogue [93]. The four channels return Table 4.

Test	Measured	Prediction	Consistent
A. floor / ceiling	$R_{\infty} = 6.96 \pm 0.27$; 0/120 clusters cross the 3/19 ceiling at $> 2\sigma$	floor $M_{\text{dm}}/M_b \geq 16/3$ respected	yes
B. depletion + closure	aperture-controlled slope = -0.09 ; $\langle f_b \rangle = 0.126 < 3/19$	slope < 0 , conserved-budget depletion	yes
C. profile guardrail	lensing M_{500}/M_{200} vs CDM NFW: typical offset -0.28σ (806 clusters)	matches NFW (collisionless)	yes
D. profile maturity	partial $\rho(\Delta c, \delta \log M, z) = -0.215$ ($p = 0.006$); -0.210 after mis-centering control	$\rho < 0$ (disturbed $\rightarrow$ under-concentrated)	yes

Table 4: Real-data consistency across four channels. All four are consistent with the construction; three share morphology with Λ CDM and are therefore consistency tests, not discriminants.

The interpretation must be stated precisely, lest “mere consistency” be misread as a failure to predict. Three of the four channels are *shared-morphology* results, degenerate with Λ CDM

by the correspondence above: they are consistency checks, not model discriminants, and the framework says so in advance. That they pass confirms the internal coherence of the construction and its collisionless, conserved-budget character. It neither discriminates against Λ CDM nor, crucially, counts against the ontology. A consistency check that the ontology *predicted* would be non-discriminating cannot, by construction, refute it.

The sharpest available test of the number itself is not reached by these catalogues. Recovering 16/3 from a fully baryon-complete system requires the warm-hot intergalactic and r_{200} baryon accounting that current data do not close. Test A measures $R_\infty = 6.96 \pm 0.27$, well above the floor and exactly what a leaky census should return; it bounds the construction without confronting it.

20.11 Two ontologies, one morphology

The same coherent-entropy substrate supports two populations that are morphologically identical. Galactic dark matter is *minting residue* — coherent entropy in its past, precipitation attempted, failed, frozen. Black-hole dark matter is *collected reserve* — coherent entropy in its future, never precipitated, searched out and concentrated (Section 21). The unifying category is coherent entropy: unprecipitated quantum states. Coherent entropy sorted by a black hole is either inert or not; a dead pixel encountered by one is sorted alongside the live reserve, having no other option.

Both gravitate, both cluster collisionlessly, both are electromagnetically inert, and a lensing map cannot tell them apart. Λ CDM cannot even *state* this distinction: there, a halo is a halo, separated only by assembly history, never by kind.

21 Black holes: information saturation and the Little Bang

21.1 Black holes are entropy organizers

Within this framework a black hole is not an information destroyer but the extremal coherent-entropy *organizer*: a region where extreme organization and density of coherent entropy has access to the holographic screen. It concentrates the coherent sector into the deepest available well, and drives the information ratio $I/I_{\max}$ toward saturation. Apparent gravitational effects arise from the information pressure (53) of the organized content.

There is no Hawking radiation and no information leakage in this picture. The horizon retains and organizes all information; apparent thermal effects arise from the *gradient* between the coherent (cold, ordered) and decoherent (hot, disordered) sectors, not from particle emission. The temperature profile near the horizon follows

$$T(r) = T_{\text{coh}} \left(1 - \frac{r_s}{r}\right)^{-1/2} \left[1 - \left(\frac{I}{I_{\max}}\right)^2\right], \quad (136)$$

with T_{coh} from Eq. (27), and the fluctuation amplitude carries the QTEP signature,

$$\Delta T(r) = \eta T_{\text{coh}} \left(\frac{I}{I_{\max}}\right)^2, \quad \eta = 2.2589. \quad (137)$$

These fluctuations exhibit *discrete steps* at information-saturation thresholds $I = n \ln 2 \cdot I_{\max}$, the same $\ln 2$ quantization that governs every precipitation event.

21.2 Saturation and expansion

Information accumulates on the horizon under accretion. The encoding efficiency degrades quadratically as saturation approaches,

$$\eta_{\text{enc}} = \eta_0 \left(1 - \frac{I}{I_{\text{max}}}\right)^2, \quad (138)$$

producing a non-linear buildup of the pressure (53) even below the absolute bound. The evolution of the information content is logistic,

$$\frac{dI}{dt} = \gamma I \left(1 - \frac{I}{I_{\text{max}}}\right), \quad (139)$$

with discrete transitions at $I_n = n \ln 2 \cdot I_{\text{max}}$ enforced jointly by the no-cloning theorem and the holographic bound: intermediate states would violate one or the other.

When P_I reaches the critical value $P_c = \gamma c^4 / (8\pi G)$ — equivalently $(I/I_{\text{max}}) \rightarrow 1$ — further encoding is impossible. The no-cloning theorem forbids copying, the holographic bound forbids further density, and the system must create new degrees of freedom. It does so by expanding: a *Little Bang*, a localized spacetime expansion preserving total information through dimensional growth rather than through radiation. The information paradox is resolved without firewalls and without violating either unitarity or the equivalence principle at the horizon: the apparent loss was a temporary storage phase.

21.3 Timescales, and why this is unobserved

The time to reach saturation from the logistic equation (139) is

$$t_{\text{sat}} = \frac{1}{\gamma} \ln \left(\frac{c^2}{GM\gamma} \right), \quad (140)$$

and the coherence time across information content I diverges logarithmically as $I \rightarrow I_{\text{max}}$,

$$t_{\text{coh}}(I) = \frac{1}{\gamma} \ln \left(\frac{I_{\text{max}}}{I_{\text{max}} - I} \right), \quad (141)$$

the critical slowing-down characteristic of a continuous phase transition.

The controlling scale is γ^{-1} . At the present epoch, from Eq. (33),

$$\gamma_0^{-1} = 1.29 \times 10^{20} \text{ s} = 4.1 \times 10^{12} \text{ yr}, \quad (142)$$

some three hundred times the age of the universe. The associated energy and temperature scales are

$$E_{\text{trans}} = \hbar \gamma_0 \ln 2 = 5.7 \times 10^{-55} \text{ J}, \quad T_0 = \frac{\hbar \gamma_0}{2\pi k_B} = 9.4 \times 10^{-33} \text{ K}. \quad (143)$$

This is why information-based black-hole effects have not been seen. The characteristic timescale exceeds cosmic history by two and a half orders of magnitude, and the characteristic energy is thirty orders below any direct detection threshold. Only primordial black holes in a narrow mass range could reach saturation within observable cosmic history, and the probability of catching one is vanishingly small.

The obstruction is sharper than a timescale, and it is worth being precise about, because the margin is not what it has previously been taken to be. The fundamental oscillation frequency of the transition structure is

$$f \sim \frac{\gamma_0}{2\pi} = 1.2 \times 10^{-21} \text{ Hz}, \quad (144)$$

against a pulsar-timing-array band reaching $\sim 10^{-9}$ Hz. The gap is twelve orders of magnitude — large, and decisive, but eight orders smaller than the framework’s own earlier estimates implied. The correlation-based strategies (thermal-fluctuation autocorrelation, gravitational-wave amplitude modulation with sidebands at f , and electromagnetic–gravitational cross-correlation, all of which carry the $\ln 2$ transition structure and accumulate signal-to-noise as $\sqrt{T_{\text{obs}} \gamma_0}$) are correspondingly more sensitive than previously reckoned. They remain far out of reach: an observation window cannot resolve a period of $\ln 2/\gamma_0 = 2.8 \times 10^{12}$ yr. But the black-hole sector is not as remote from observation as it has been supposed, and the distance to it should be quoted correctly.

21.4 The connection to dark matter

Black holes are, on this account, super-dense regions of *coherent entropy* — the same substrate identified as dark matter in Section 20. The gravitational topology reflects the complexity of information organization around the concentration, and the hole is effectively *searching* the screen for additional coherent entropy to consolidate. This is the “fuel” population of Section 20: coherent entropy in its future, accumulating toward $I/I_{\text{max}} \rightarrow 1$, as against the galactic “ash” of frozen dead pixels. Gravity reads content, not provenance, so a lensing map cannot separate them.

22 The laboratory: the Landauer manifold

22.1 Landauer’s bound as a continuum limit

Landauer’s principle bounds the work to transfer a bit’s accessibility out of its register by $k_B T \ln 2$. Within this framework it is not an input; it is what the per-event ledger (24) returns in a particular limit.

Consider a logical bit in a two-state quantum system weakly coupled to a thermal reservoir at temperature T and simultaneously monitored by a channel driving precipitation at effective rate γ_{eff} , with reservoir relaxation rate Γ_{th} . The *Landauer manifold* is the one-dimensional parameter space

$$\boxed{\mu \equiv \frac{\gamma_{\text{eff}}}{\Gamma_{\text{th}}}} \quad (145)$$

The apparatus rate γ_{eff} — measurement-induced dephasing in a transmon, the inverse cycling-collection time for a shelved ion, the inverse imaging-pulse duration for a lattice atom — is layered on the local modular rate γ and dominates it by many orders of magnitude. The cosmological γ_0 is not the relevant baseline in a laboratory; μ is.

22.2 The two limits

Classical limit, $\mu \rightarrow 0$. Precipitation events are rare on the timescale of thermal relaxation. Between events the holographic boundary registers the $(1 - \ln 2)$ -nat residual in full; the causal patch closes cleanly on each event and the decoherent channel is not rerouted. The information ledger is Eq. (36); the reservoir flow is Eq. (37); and the dissipated work is

$$W_{\mu \rightarrow 0} = k_B T \Delta S_{\text{reservoir}} = k_B T \ln 2. \quad (146)$$

Landauer’s bound, recovered exactly, as the $\mu \rightarrow 0$ limit of the QTEP ledger — as a consequence, not an input, and with no modification in the regime where the experiments that established it operate.

Quantum limit, $\mu \rightarrow \infty$. Precipitation drives the system across the coherent/classical boundary faster than boundary routing can complete. Export to the boundary is rate-limited by Γ_{th} through the system–reservoir coupling it must traverse; when $\gamma_{\text{eff}} \gg \Gamma_{\text{th}}$, successive events arrive before the previous residual has resolved, and the decoherent channel is *rerouted from the boundary into the reservoir*. The ledger remains conservative — Eq. (24) is an identity and cannot fail — but its boundary entry is now carried by the reservoir:

$$\Delta S_{\text{boundary}} = 0, \quad \Delta S_{\text{res,info}} = -(1 - \ln 2), \quad (147)$$

so the total reservoir flow becomes

$$\Delta S_{\text{reservoir}} = \ln 2 + (1 - \ln 2) = 1 \text{ nat}, \quad \boxed{W_{\mu \rightarrow \infty} = k_B T} \quad (148)$$

exceeding Landauer by the factor $1/\ln 2 = 1.4427$.

22.3 The crossover and its signature

Between the limits, a fraction $\chi(\mu) \in [0, 1]$ of the decoherent channel is rerouted thermally, with $\chi(0) = 0$ and $\chi(\infty) = 1$:

$$\boxed{W(\mu) = k_B T [\ln 2 + \chi(\mu) (1 - \ln 2)]} \quad (149)$$

In a minimal memoryless routing model with boundary-completion rate set by the thermal coupling itself, $\chi(\mu) = e^{-1/\mu}$, giving $\chi(1) = e^{-1} = 0.368$ at the crossover. We do not insist on this form. What the framework predicts *robustly*, independent of the export model, is the two asymptotes, the crossover scale $\mu \sim 1$, the saturated excess

$$\Delta W_{\text{sat}} = W(\infty) - W(0) = k_B T (1 - \ln 2) = 0.307 k_B T, \quad (150)$$

and — the operational signature of the ledger — the ratio

$$\boxed{\frac{W(0)}{\Delta W_{\text{sat}}} = \frac{\ln 2}{1 - \ln 2} = \eta = 2.2589} \quad (151)$$

The universal constant of Eq. (25), measured on a bench.

Three remarks on status. Equation (149) is a *lower bound*, not an equality: real operations dissipate more, through incomplete equilibration, finite-time driving, and protocol irreversibility. The framework does not alter those corrections; it modifies the bound itself, and thereby what it means to approach reversibility at finite μ . The temperature T is the physical reservoir temperature, not the internal T_{coh} or T_{decoh} of Eq. (27), which enter only through μ . And the $\mu \rightarrow 0$ limit is taken at fixed $\gamma_{\text{eff}} \cdot N_{\text{ev}}$ — the standard coarse-graining by which a discrete shot process becomes a continuous stochastic one, which is precisely why Landauer re-emerges from it.

22.4 The experimental programme

Three platforms span the manifold. *Superconducting qubits* with transmon–resonator calorimetry: γ_{eff} is the measurement-induced dephasing rate, $\Gamma_{\text{th}} = 1/T_1$, giving $\mu \in [1, 10^3]$ — fully spanning the crossover and saturating in the quantum limit. *Trapped ions* with motional thermometry: $\mu \in [10^3, 10^8]$, deep in the quantum limit, offering the cleanest access to the saturated value at the cost of being far from the crossover. *Ultracold atoms* in quantum-gas-microscope lattices: $\mu \in [1, 10^4]$, the best combination of crossover access and resolution, at the cost of more complex systematic control.

A successful test must observe: (1) monotonic increase of the minimum dissipated work from $k_B T \ln 2$ at $\mu \ll 1$ to $k_B T$ at $\mu \gg 1$; (2) saturation at $k_B T$ with excess $\Delta W_{\text{sat}} = k_B T(1 - \ln 2)$; (3) the ratio $\eta = 2.2589$; and (4) a crossover scale $\mu_c \sim 1$ that is system-independent within factors of order unity. Observation of (1) alone is consistent with several frameworks that modify Landauer at finite monitoring rate. (2) and (3) are specific to QTEP, and (4) across three physically distinct platforms — different energies, temperatures, Hilbert-space structures and measurement channels — is the strongest constraint.

22.5 This is the test of the obit

Section 4 named the obit as the framework’s single postulate. This is where it is settled.

A measured saturation at $k_B T \ln 2$ for arbitrarily large μ , with no excess, *falsifies the obit postulate*: it would show there is no 0.307-nat residual, hence no partition, hence no η , hence nothing. A saturation at some value other than $k_B T$ would constrain the obit to a different value — and if the saturation were $k_B T \ln 2 / \ln k$ for some base k , it would identify the correct logarithmic base, which is the actual content of the postulate.

The framework is thus not merely consistent with Landauer’s principle. It is committed, in advance, to specific predictions that Landauer alone does not make, on apparatus that exists today.

Part V

Status

23 The ledger: derived, postulated, accepted

The framework’s discipline is that these three columns are kept apart and never quietly exchanged.

23.1 Postulated

Exactly one quantity is undeduced: the **obit**, $S_{\text{obit}} = 1$ nat (Postulate 2). Everything in the QTEP sector — the residual -0.307 , the ratio $\eta = 2.2589$, the whole partition — follows from it together with the derived ebit $\ln 2$.

Two normalizations are adopted and should be recognized as such: the Bisognano–Wichmann prefactor 2π with unit surface gravity at σ (Postulate 1), which is standard; and the processing-cycle condition that one coarse-grained cycle updates every holographic degree of freedom once per Hubble time (Postulate 3), which fixes γ ’s normalization and which we make explicit rather than bury.

23.2 Derived

Quantity	Source
Modular flow θ_t , generator K	Tomita–Takesaki; Bisognano–Wichmann, Eq. (11)
Horizon entropy $S_H = \mathcal{A}/4\ell_P^2$	area term of the crossed-product (generalized) entropy of the modular state, Eq. (15)
Ebit, $S_{\text{coh}} = \ln 2$	von Neumann entropy of half a Bell pair, Eq. (21)

Quantity	Source
Rate $\gamma = H/\ln S_H$	Bekenstein–Hawking + Margolus–Levitin + Postulate 3
Vacuum form $\langle T^{\text{info}} \rangle_\Omega = -\rho_\Lambda g_{\mu\nu}$	conformal isometry of $ \Omega\rangle$, Eq. (39)
Consistency with spatial flatness	Theorem 8.1
Einstein’s equations	relative entropy + Raychaudhuri + Bianchi, Eq. (50)
Λ and G	Eqs. (95) and (120)
The 10^{-123} hierarchy, <i>exactly</i>	$\rho_\Lambda/\rho_{\text{Pl}} = \frac{3}{2} \Omega_\Lambda/N_P$, Eq. (100)
H_0 from laboratory data alone	67.11 ± 0.25 from (α, G) , Eq. (126)
Parity-evenness of the dynamics	Lemma 11.1
Planckian initial conditions	Theorem 10.1
Index rigidity of the generation count	Theorem 13.4
Landauer’s bound $k_B T \ln 2$	$\mu \rightarrow 0$ limit of the ledger, Eq. (146)

23.3 Accepted as boundary data

The gauge structure $E_8 \times E_8$; the Calabi–Yau and the bundle; the generation count; the sense of the handedness; the amplitude β of cosmic birefringence; the initial state ω_0 . Postulate 5 says why, and Theorems 13.4 and Lemma 11.1 say why the two sharpest of these cannot be otherwise. These are the same inputs every comparable construction accepts. The framework’s contribution is not to derive them but to *locate the line*, and to prove that the line falls where it does.

24 Falsification

Every commitment below is stated with the observation that kills it. The framework has no parameter with which to absorb any of them.

Commitment	Prediction	Falsified by
Modular vacuum fixes Λ	$w(z) = -1 + \mathcal{O}(10^{-11})$	confirmed $w_0 \neq -1$ (DESI Y5, Euclid, Roman)
Flatness is a commitment	$\Omega_K = 0$ exactly	detected $\Omega_K \neq 0$ at high significance
Holographic projection	$H_{\text{local}} = 69.15$ (NGC 4258), 70.36 (LMC); cross-anchor shift -1.21	TRGB convergence elsewhere as precision improves
Zeno-enhanced acoustic ruler	$r_d = 150.71$ Mpc; $\alpha =$ -5.7 ± 2.0	BAO/CMB exclusion of $\alpha =$ -5.7 at high significance
The α – G – H constraint	$\Delta \ln S_H = 0.0075$; $\alpha^{-1}(H_0) = 137.032$; $G = 6.624 \times 10^{-11}$	contradiction of QED at any pre- cision the $E \leftrightarrow H_{\text{eff}}$ correspon- dence can be tested at
Adjudication of the Hubble dispute	$H_0 = 67.11 \pm 0.25$ from (α, G) alone, with no cos- mological input; SH0ES ex- cluded at 5.5σ	a cosmology-free determination of H_0 inconsistent with Eq. (126); or TRGB convergence on ~ 73 as precision improves
Cross-sector co-slaving	atomic α, μ constant across cosmic history	confirmed varying α or μ at $\lesssim$ 10^{-5}
Dark matter is a state	$\Omega_{\text{dm}}/\Omega_{\text{b}} = 16/3$; detection nulls	a detected dark-matter particle; baryon-complete system recover- ing a ratio $\neq 16/3$

Commitment	Prediction	Falsified by
Parity-even dynamics	achromatic birefringence only; gravitational parity nulls	$\geq 3\sigma$ chromatic birefringence, or parity-odd GW/LSS signal not sourced by a content field
Locking of the chiral index	generation count and net chirality are one datum	observed independence of the two
The obit, $S_{\text{obit}} = 1$ nat	$W(\mu)$ saturates at $k_B T$; ratio $\eta = 2.2589$	saturation at $k_B T \ln 2$, or at any other value, or absence of a crossover

Note the asymmetry, and that it is deliberate. Several of these predictions — the birefringence nulls in particular — are shared with an entire class of theories, so confirming them cannot single this framework out. *Confirmation is shared; falsification is decisive.* That is the correct shape for a physical claim, and the framework is constructed so that it holds.

25 Synthesis versus derivation, and Vopson’s question

The claim this document advances is a synthesis claim, and its strength must be assessed as such. Nothing here re-derives the thermodynamic origin of Einstein’s equations: the local-Rindler argument is Jacobson’s [28], its entanglement-first-law extension to holographic conformal field theories is Faulkner and collaborators’ [29, 30, 31], and the modular structure on the cosmic diamond was developed progressively across the corpus [140, 141]. What is contributed is the identification of $\gamma(z) = H(z)/\ln S_H(z)$ as the structural element that ties these previously independent results into a single framework, together with the empirical record that the identification’s signatures appear at distinct cosmological epochs.

That identification is a structural inference, not a logical consequence of any single success. Each regime is separable: the acoustic ruler of Section 15 would stand if the Hubble projection failed; the Hubble projection does not depend on the cosmological-constant scale of Section 17; and neither rests on the other two. The synthesis claims instead that all are signatures of γ -controlled information processing on $\mathcal{A}(\mathcal{D})$, and the evidence for that reading is convergence — three independent regimes reaching the same machinery through structurally distinct routes. An alternative framework that handled the same observations with equal economy while denying γ the structural role would owe an account of why the same partition constant recurs at recombination, at the present epoch, and on a laboratory bench. The burden the synthesis discharges is precisely that recurrence.

Vopson has asked whether gravity constitutes evidence of a computational universe [137], building on mass–energy–information equivalence and infodynamics. The answer here is affirmative, but qualified in a specific way. Gravity is evidence of computation in the sense established in Section 9: the Einstein tensor is the consistency condition under which γ -controlled information processing on the algebra of the cosmic causal diamond can be encoded geometrically through its holographic boundary. The synthesis answers the question he posed within a different infrastructure, and it supplies what was missing there — a specific processing rate and a body of quantitative signatures.

It also renders his mass–energy–information equivalence principle [138] within that infrastructure [142]. The principle assigns a stored bit the mass $k_B T \ln 2/c^2$ and does not fix T . Closure on $\mathcal{D}(p, q)$, the requirement that a mass be counted once, fixes it. The mass per holographic degree of freedom demanded by the Friedmann budget, $M/S_H = (c^3/2GH)(G\hbar H^2/\pi c^5)$, equals

the mass equivalent of the screen’s own modular temperature, $k_B T_{\text{mod}}/c^2$:

$$\mu = \frac{\hbar H}{2\pi c^2} = 4.08 \times 10^{-70} \text{ kg} \quad (152)$$

per degree of freedom, 2.83×10^{-70} kg per bit. The two sides are computed from formally independent inputs; under the identification $E_{\text{horizon}} = S_H T_{\text{mod}}$ of Section 8 their agreement is a consistency requirement, and μ is the same either way.

Three consequences follow. The mass of a bit is epoch-dependent, since $\mu \propto H$. Its characteristic energy, $k_B T_{\text{mod}} = \hbar H/2\pi$, is horizon-scale, so a test of the principle at a bench temperature T overshoots by exactly T/T_{mod} , 1.1×10^{32} at 300 K. And the Landauer cost $k_B T \ln 2$ is a flux into a reservoir, the reservoir flow of Eq. (37), not a property of the bit; assigning it to the bit as well counts the same mass twice. The principle holds, but between a bit and the horizon, not between a bit and a laboratory.

Four regimes now carry the same constant, and their exposure is asymmetric by design:

- at recombination, η sets the partition of deferred versus deposited decoherent entropy behind the enhanced ruler;
- at distance-ladder scales, the same partition structure underlies the projection prefactor γ/H ;
- at the present epoch, the modular-vacuum projection fixing ρ_Λ is the static limit of the same partition;
- at apparatus scales, η is the predicted ratio $W(0)/\Delta W_{\text{sat}}$ of the Landauer crossover, Section 22.

The first three are cosmological and mutually reinforcing but share the systematic risks of cosmological inference. The fourth is a bench experiment in which systematics can be quantified directly. Confirmation there adds a fourth independent leg under controlled conditions; falsification would not collapse the cosmological successes, which stand on their own data, but would confine the universality of η across regimes and remove the synthesis’s laboratory anchor. Each regime is independently testable, and the synthesis claim is contingent on continued convergence across all of them — which is the structurally honest posture available, and the one this document adopts throughout.

26 Open problems

We name these at their true size.

The obit is not derived. It is the framework’s one postulate, and the Landauer crossover of Section 22 is designed to settle it. Until that experiment runs, the programme rests on an undeduced constant — as physics has done before, and as it does whenever a new quantum enters.

Bisognano–Wichmann on the FLRW diamond. The geometric form (11) is established for Rindler wedges, for double cones in conformal field theory, and for de Sitter static patches. Its extension to the cosmic diamond is three problems, not one, and they are of different size.

Exact FLRW, conformal sector: undone, not open. Every FLRW spacetime is conformally flat, whatever the sign of its spatial curvature. The Hislop–Longo modular flow of a Minkowski double cone [5] is carried to the cosmic diamond by the Weyl rescaling, and the conformal Killing

vector (12) is its generator. For conformally invariant fields in the conformal vacuum, Eq. (11) on FLRW is the image of a theorem. What remains is to write the map out and to track the Weyl-anomaly contribution to the modular entropy under the rescaling. That is a calculation, not a research problem.

Non-conformal matter: open. Mass breaks the symmetry the map rests on. Even in Minkowski space the double-cone modular generator of a massive free field has no closed form; evaluated numerically at one-particle level, it is mass-dependent, departing from the geometric form that holds for wedges [11]. Baryons and cold dark matter are not conformal. Because the framework postulates K and defines $T_{\mu\nu}^{\text{info}}$ from it (Section 8), this does not break the construction. It relocates the question to Postulate 1, and the question is a fork. Either the modular state of $\mathcal{A}(\mathcal{D})$ is not the massive vacuum restricted to $\mathcal{D}$, and the postulate stands as written; or it is, and Eq. (11) holds exactly only for the conformal sector, with the massive departure a correction every matter-sector projection must carry. The framework has no reason to choose between them; observation does, because the branches separate by sector. The electromagnetic constraint of Section 18 is built on the free Maxwell field, which is conformal, and is untouched on either branch: α , G and the laboratory H_0 stand whichever way the fork falls. The recombination ruler of Section 15, built on the baryon–photon plasma, is not. On the second branch it carries a massive correction the electromagnetic sector does not, and a sector-dependent residual is the signature. The fork becomes decidable once that correction is sized.

Perturbed FLRW: open. Inhomogeneity switches on the Weyl tensor. The diamond loses its conformal Killing vector, and modular flow ceases to be geometric at first order in the perturbation. The Gibbons–Solodukhin series (2) already carries the directional term $R_{\mu\nu}T^\mu T^\nu$; whether the modular corrections organize in the same curvature expansion is not known. The numerical predictions of Part IV are background quantities and do not wait on this. What waits on it is the size of their corrections.

The acoustic coefficient. $\alpha = -5.7 \pm 2.0$ is a controlled scaling estimate of the Lindblad response, not an explicit quantum-kinetic evaluation of Eq. (75). The BAO data prefer $\alpha \approx -7.2$, inside the prior but not at its centre. Finite- k corrections, the detailed visibility function, and non-Gaussian corrections to the modular Hamiltonian all push in the observed direction, and a full modular-spectral calculation should close the gap. If future data exclude -5.7 at high significance, the Zeno–Lindblad treatment has missing physics.

The dark-matter margin. One shell supplies $q \ln 2 = 2.77$ nats against a cost of 3; the second shell — and hence $16/3$ rather than $4/2$ — turns on an 8% shortfall. A shortfall is all the argument requires, but a sharper account of what a drained plaquette surrenders would put the step beyond doubt. Likewise, a rigorous enumeration of the superselected charges would place $\nu = 3$ beyond the physically-forced count used here, and the precipitation-epoch estimate should be tightened beyond its present factor of a few.

The energy-to-Hubble correspondence. The map $E \leftrightarrow H_{\text{eff}}(E)$ is used in Section 18 but not derived. Making it explicit, and verifying that $\alpha^{-1}(H_{\text{eff}}(E))$ reproduces QED across the full range of tested observables, is open work on the framework’s side. The framework is committed to producing QED, not to negotiating with it.

Friedmann evolution from modular dynamics. The framework has structural anchors at the boundary, at recombination, at the distance ladder, and at the present epoch. Deriving Friedmann-like evolution equations *between* those anchors, from the causal-diamond modular flow, would convert a set of consistency results into an equivalence proof. A fifth anchor at nucleosynthesis would sharpen the exposure across the radiation era.

The screen-to-bulk profile. The projection from the screen shell to a bulk density profile is the remaining forward test of the dark-matter mechanism’s detail. Because minted pixels are collisionless, we expect it to reproduce NFW rather than yield something distinctive — which is a prediction, and a reason the profile is a consistency check rather than a discriminant.

27 Conclusion

One algebra, one postulate, and a chain of projections. From the modular structure of $\mathcal{A}(\mathcal{D})$ and the orbit there follow the rate γ , the partition η , and the geometric encoding that is gravity; and from those, without a fitted parameter anywhere, the sound horizon at recombination, the local Hubble parameter at each distance-ladder anchor, the density and equation of state of the cosmological constant, the fine-structure constant, Newton’s constant, the dark-to-baryon ratio, Planck-scale initial conditions, and a bench-top dissipation crossover.

The framework is equally definite about its limits. Its dynamics are parity-even and therefore cannot originate handedness; the index is rigid and therefore cannot yield a generation count. The chiral structure of matter, the gauge apparatus, and the initial state are boundary data — accepted, not derived — and the line between what is derived and what is accepted falls exactly at the boundary of the diamond. That division is not an apology. It is the framework’s central structural claim, and it is where the framework stakes itself.

The empirical record of the next decade will adjudicate all of it at once. DESI Y5, Euclid and Roman will settle $w(z)$. JWST-anchored TRGB campaigns will settle the distance-ladder projection. Simons Observatory, LiteBIRD and CMB-S4 will settle the birefringence channel. Direct-detection programmes will continue to return the nulls the ontology predicts. And a transmon, an ion trap, or a lattice of cold atoms will return $k_B T$ or it will not.

Nothing here can be adjusted to accommodate the answer.

Acknowledgements

The author acknowledges the support of the Information Physics Institute and the \$DAD community, and the survey collaborations — Planck, DESI, SDSS/eBOSS, DES, CCHP, CSP, Pantheon+/SH0ES, SPARC, and the maintainers of the X-ray cluster, weak-lensing and Local Group dwarf catalogues — whose open data underpin the tests reported here. Large language models were used for proofreading and editorial feedback.

Data Availability

Analysis code and preregistration documents for the BAO and TRGB analyses are available at <https://github.com/bryceweiner/h-lcdm>. All observational inputs are drawn from public archives, cited in place.

References

- [1] J. J. Bisognano and E. H. Wichmann, “On the duality condition for a Hermitian scalar field,” *J. Math. Phys.* **16**, 985 (1975).
- [2] M. Takesaki, *Tomita’s Theory of Modular Hilbert Algebras and its Applications*, Lecture Notes in Mathematics **128** (Springer, 1970).
- [3] R. Haag, *Local Quantum Physics: Fields, Particles, Algebras*, 2nd ed. (Springer, 1996).

- [4] R. Haag, N. M. Hugenholtz and M. Winnink, “On the equilibrium states in quantum statistical mechanics,” *Commun. Math. Phys.* **5**, 215 (1967).
- [5] P. D. Hislop and R. Longo, “Modular structure of the local algebras associated with the free massless scalar field theory,” *Commun. Math. Phys.* **84**, 71 (1982).
- [6] H. Casini, M. Huerta and R. C. Myers, “Towards a derivation of holographic entanglement entropy,” *JHEP* **2011**(05), 036 (2011).
- [7] H. Casini, “Relative entropy and the Bekenstein bound,” *Class. Quantum Grav.* **25**, 205021 (2008).
- [8] J. de Boer, F. M. Haehl, M. P. Heller and R. C. Myers, “Conformal quantum mechanics of causal diamonds,” *JHEP* **05**, 072 (2020).
- [9] M. B. Fröb, “Modular Hamiltonian for de Sitter diamonds,” *JHEP* **12**, 074 (2023).
- [10] L. Aalsma and S.-E. Bak, “Modular fluctuations in cosmology,” *Phys. Rev. D* **112**, 026017 (2025).
- [11] H. Bostelmann, D. Cadamuro and C. Minz, “On the mass dependence of the modular operator for a double cone,” *Ann. Henri Poincaré* **24**, 3031 (2023).
- [12] E. Witten, “Gravity and the crossed product,” *JHEP* **10**, 008 (2022).
- [13] V. Chandrasekaran, R. Longo, G. Penington and E. Witten, “An algebra of observables for de Sitter space,” *JHEP* **02**, 082 (2023).
- [14] K. Jensen, J. Sorce and A. J. Speranza, “Generalized entropy for general subregions in quantum gravity,” *JHEP* **12**, 020 (2023).
- [15] G. L. Sewell, “Quantum fields on manifolds: PCT and gravitationally induced thermal states,” *Ann. Phys.* **141**, 201 (1982).
- [16] W. G. Unruh, “Notes on black-hole evaporation,” *Phys. Rev. D* **14**, 870 (1976).
- [17] J. D. Bekenstein, “Black holes and entropy,” *Phys. Rev. D* **7**, 2333 (1973).
- [18] J. D. Bekenstein, “Universal upper bound on the entropy-to-energy ratio for bounded systems,” *Phys. Rev. D* **23**, 287 (1981).
- [19] S. W. Hawking, “Particle creation by black holes,” *Commun. Math. Phys.* **43**, 199 (1975).
- [20] J. M. Bardeen, B. Carter and S. W. Hawking, “The four laws of black hole mechanics,” *Commun. Math. Phys.* **31**, 161 (1973).
- [21] G. W. Gibbons and S. W. Hawking, “Cosmological event horizons, thermodynamics, and particle creation,” *Phys. Rev. D* **15**, 2738 (1977).
- [22] G. ’t Hooft, “Dimensional reduction in quantum gravity,” arXiv:gr-qc/9310026 (1993).
- [23] L. Susskind, “The world as a hologram,” *J. Math. Phys.* **36**, 6377 (1995).
- [24] R. Bousso, “The holographic principle,” *Rev. Mod. Phys.* **74**, 825 (2002).
- [25] G. W. Gibbons and S. N. Solodukhin, “The geometry of small causal diamonds,” *Phys. Lett. B* **649**, 317 (2007).
- [26] B. Swingle, “Entanglement renormalization and holography,” *Phys. Rev. D* **86**, 065007 (2012).
- [27] S. Ryu and T. Takayanagi, “Holographic derivation of entanglement entropy from AdS/CFT,” *Phys. Rev. Lett.* **96**, 181602 (2006).
- [28] T. Jacobson, “Thermodynamics of spacetime: the Einstein equation of state,” *Phys. Rev. Lett.* **75**, 1260 (1995).
- [29] T. Faulkner, M. Guica, T. Hartman, R. C. Myers and M. Van Raamsdonk, “Gravitation from entanglement in holographic CFTs,” *JHEP* **2014**(03), 051 (2014).
- [30] N. Lashkari, M. B. McDermott and M. Van Raamsdonk, “Gravitational dynamics from entanglement ‘thermodynamics’,” *JHEP* **04**, 195 (2014).

- [31] T. Faulkner, F. M. Haehl, E. Hijano, O. Parrikar, C. Rabideau and M. Van Raamsdonk, “Nonlinear gravity from entanglement in conformal field theories,” *JHEP* **2017**(08), 057 (2017).
- [32] A. Raychaudhuri, “Relativistic cosmology. I,” *Phys. Rev.* **98**, 1123 (1955).
- [33] R. M. Wald, *General Relativity* (University of Chicago Press, 1984).
- [34] S. W. Hawking and G. F. R. Ellis, *The Large Scale Structure of Space-Time* (Cambridge University Press, 1973).
- [35] R. Landauer, “Irreversibility and heat generation in the computing process,” *IBM J. Res. Dev.* **5**, 183 (1961).
- [36] C. H. Bennett, “The thermodynamics of computation—a review,” *Int. J. Theor. Phys.* **21**, 905 (1982).
- [37] R. Landauer, “Information is physical,” *Phys. Today* **44**(5), 23 (1991).
- [38] D. Deutsch, “Quantum theory, the Church–Turing principle and the universal quantum computer,” *Proc. R. Soc. Lond. A* **400**, 97 (1985).
- [39] N. Margolus and L. B. Levitin, “The maximum speed of dynamical evolution,” *Physica D* **120**, 188 (1998).
- [40] S. Lloyd, “Ultimate physical limits to computation,” *Nature* **406**, 1047 (2000).
- [41] W. H. Zurek, “Decoherence, einselection, and the quantum origins of the classical,” *Rev. Mod. Phys.* **75**, 715 (2003).
- [42] B. Misra and E. C. G. Sudarshan, “The Zeno’s paradox in quantum theory,” *J. Math. Phys.* **18**, 756 (1977).
- [43] P. Facchi and S. Pascazio, “Quantum Zeno dynamics: mathematical and physical aspects,” *J. Phys. A* **41**, 493001 (2008).
- [44] B. Skinner, J. Ruhman and A. Nahum, “Measurement-induced phase transitions in the dynamics of entanglement,” *Phys. Rev. X* **9**, 031009 (2019).
- [45] Y. Li, X. Chen and M. P. A. Fisher, “Quantum Zeno effect and the many-body entanglement transition,” *Phys. Rev. B* **98**, 205136 (2018).
- [46] J. M. R. Parrondo, J. M. Horowitz and T. Sagawa, “Thermodynamics of information,” *Nat. Phys.* **11**, 131 (2015).
- [47] C. Jarzynski, “Nonequilibrium equality for free energy differences,” *Phys. Rev. Lett.* **78**, 2690 (1997).
- [48] G. E. Crooks, “Entropy production fluctuation theorem,” *Phys. Rev. E* **60**, 2721 (1999).
- [49] U. Seifert, “Stochastic thermodynamics, fluctuation theorems and molecular machines,” *Rep. Prog. Phys.* **75**, 126001 (2012).
- [50] T. Sagawa and M. Ueda, “Generalized Jarzynski equality under nonequilibrium feedback control,” *Phys. Rev. Lett.* **104**, 090602 (2010).
- [51] A. Bérut *et al.*, “Experimental verification of Landauer’s principle,” *Nature* **483**, 187 (2012).
- [52] A. Blais, A. L. Grimsmo, S. M. Girvin and A. Wallraff, “Circuit quantum electrodynamics,” *Rev. Mod. Phys.* **93**, 025005 (2021).
- [53] B. Karimi, F. Brange, P. Samuelsson and J. P. Pekola, “Reaching the ultimate energy resolution of a quantum detector,” *Nat. Commun.* **11**, 367 (2020).
- [54] D. Leibfried, R. Blatt, C. Monroe and D. Wineland, “Quantum dynamics of single trapped ions,” *Rev. Mod. Phys.* **75**, 281 (2003).
- [55] W. S. Bakr *et al.*, “A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice,” *Nature* **462**, 74 (2009).

- [56] Planck Collaboration, “Planck 2018 results. VI. Cosmological parameters,” *Astron. Astrophys.* **641**, A6 (2020).
- [57] A. G. Riess *et al.*, “A comprehensive measurement of the local value of the Hubble constant,” *Astrophys. J. Lett.* **934**, L7 (2022).
- [58] W. L. Freedman *et al.*, “The Carnegie-Chicago Hubble Program VIII,” *Astrophys. J.* **882**, 34 (2019).
- [59] W. L. Freedman *et al.*, “Status report on the Chicago-Carnegie Hubble Program,” *Astrophys. J.* **985**, 203 (2025).
- [60] T. J. Hoyt *et al.*, “Improving the calibration of SNe Ia with JWST measurements of the TRGB,” arXiv:2503.11769 (2025).
- [61] S. A. Uddin *et al.*, “Carnegie Supernova Project I and II: measurements of H_0 ,” *Astrophys. J.* **970**, 72 (2024).
- [62] M. J. Reid, D. W. Pesce and A. G. Riess, “An improved distance to NGC 4258,” *Astrophys. J. Lett.* **886**, L27 (2019).
- [63] G. Pietrzyński *et al.*, “A distance to the LMC precise to one per cent,” *Nature* **567**, 200 (2019).
- [64] L. Verde, T. Treu and A. G. Riess, “Tensions between the early and late Universe,” *Nat. Astron.* **3**, 891 (2019).
- [65] D. J. Eisenstein *et al.*, “Detection of the baryon acoustic peak,” *Astrophys. J.* **633**, 560 (2005).
- [66] S. Alam *et al.*, “BOSS DR12 cosmological analysis,” *MNRAS* **470**, 2617 (2017).
- [67] DESI Collaboration, “DESI 2024 VI: cosmological constraints from BAO,” *JCAP* **2025**(02), 021 (2025).
- [68] DESI Collaboration, “DESI DR2 results II,” arXiv:2503.14738 (2025).
- [69] D. Brout *et al.*, “The Pantheon+ analysis: cosmological constraints,” *Astrophys. J.* **938**, 110 (2022).
- [70] V. Poulin, T. L. Smith, T. Karwal and M. Kamionkowski, “Early dark energy can resolve the Hubble tension,” *Phys. Rev. Lett.* **122**, 221301 (2019).
- [71] P. Carter *et al.*, “The impact of the fiducial cosmology assumption on BAO distance scale measurements,” *MNRAS* **494**, 2076 (2020).
- [72] A. Friedmann, “Über die Krümmung des Raumes,” *Z. Phys.* **10**, 377 (1922).
- [73] H. P. Robertson, “Kinematics and world-structure,” *Astrophys. J.* **82**, 284 (1935).
- [74] A. G. Walker, “On Milne’s theory of world-structure,” *Proc. London Math. Soc.* **s2-42**, 90 (1937).
- [75] S. Weinberg, “The cosmological constant problem,” *Rev. Mod. Phys.* **61**, 1 (1989).
- [76] S. M. Carroll, “The cosmological constant,” *Living Rev. Relativ.* **4**, 1 (2001).
- [77] J. Martin, “Everything you always wanted to know about the cosmological constant problem,” *C. R. Physique* **13**, 566 (2012).
- [78] F. Zwicky, “On the masses of nebulae and of clusters of nebulae,” *Astrophys. J.* **86**, 217 (1937).
- [79] V. C. Rubin and W. K. Ford, “Rotation of the Andromeda Nebula,” *Astrophys. J.* **159**, 379 (1970).
- [80] G. Bertone and D. Hooper, “History of dark matter,” *Rev. Mod. Phys.* **90**, 045002 (2018).
- [81] J. Aalbers *et al.* (LZ Collaboration), “First dark matter search results from LUX-ZEPLIN,” *Phys. Rev. Lett.* **131**, 041002 (2023).
- [82] T. Braine *et al.* (ADMX Collaboration), “Extended search for the invisible axion,” *Phys. Rev. Lett.* **124**, 101303 (2020).
- [83] J. F. Navarro, C. S. Frenk and S. D. M. White, “The structure of cold dark matter halos,” *Astrophys. J.* **462**, 563 (1996).

- [84] F. Lelli, S. S. McGaugh and J. M. Schombert, “SPARC,” *Astron. J.* **152**, 157 (2016).
- [85] M. Sereno, “CoMaLit III: literature catalogues of weak lensing clusters,” *MNRAS* **450**, 3665 (2015).
- [86] L. Lovisari *et al.*, “X-ray scaling relations for Planck-selected clusters,” *Astrophys. J.* **892**, 102 (2020).
- [87] A. W. McConnachie, “The observed properties of dwarf galaxies in and around the Local Group,” *Astron. J.* **144**, 4 (2012).
- [88] J. Wolf *et al.*, “Accurate masses for dispersion-supported galaxies,” *MNRAS* **406**, 1220 (2010).
- [89] D. Eckert *et al.*, “Gas clumping in galaxy clusters,” *MNRAS* **447**, 2198 (2015).
- [90] D. Nagai and E. T. Lau, “Gas clumping in the outskirts of Λ CDM clusters,” *Astrophys. J. Lett.* **731**, L10 (2011).
- [91] B. Diemer and M. Joyce, “An accurate physical model for halo concentrations,” *Astrophys. J.* **871**, 168 (2019).
- [92] J. Tinker *et al.*, “Toward a halo mass function for precision cosmology,” *Astrophys. J.* **688**, 709 (2008).
- [93] Z. S. Yuan and J. L. Han, “Dynamical state for 964 galaxy clusters from Chandra X-ray images,” *MNRAS* **497**, 5485 (2020).
- [94] J. S. Bell, “On the Einstein Podolsky Rosen paradox,” *Physics* **1**, 195 (1964).
- [95] P. Touboul *et al.* (MICROSCOPE), “Final results of the test of the equivalence principle,” *Phys. Rev. Lett.* **129**, 121102 (2022).
- [96] T. D. Lee and C. N. Yang, “Question of parity conservation in weak interactions,” *Phys. Rev.* **104**, 254 (1956).
- [97] C. S. Wu *et al.*, “Experimental test of parity conservation in beta decay,” *Phys. Rev.* **105**, 1413 (1957).
- [98] S. M. Carroll, “Quintessence and the rest of the world,” *Phys. Rev. Lett.* **81**, 3067 (1998).
- [99] Y. Minami and E. Komatsu, “New extraction of the cosmic birefringence from the Planck 2018 polarization data,” *Phys. Rev. Lett.* **125**, 221301 (2020).
- [100] P. Diego-Palazuelos *et al.*, “Cosmic birefringence from the Planck data release 4,” *Phys. Rev. Lett.* **128**, 091302 (2022).
- [101] J. R. Eskilt *et al.*, “Constraint on early dark energy from isotropic cosmic birefringence,” *Phys. Rev. Lett.* **131**, 121001 (2023).
- [102] J. K. Webb *et al.*, “Indications of a spatial variation of the fine structure constant,” *Phys. Rev. Lett.* **107**, 191101 (2011).
- [103] S. M. Kotuš, M. T. Murphy and R. F. Carswell, “High-precision limit on variation in the fine-structure constant,” *MNRAS* **464**, 3679 (2017).
- [104] E. Tiesinga, P. J. Mohr, D. B. Newell and B. N. Taylor, “CODATA recommended values: 2018,” *Rev. Mod. Phys.* **93**, 025010 (2021).
- [105] R. L. Workman *et al.* (Particle Data Group), “Review of Particle Physics,” *PTEP* **2022**, 083C01 (2022).
- [106] D. J. Gross, J. A. Harvey, E. Martinec and R. Rohm, “Heterotic string,” *Phys. Rev. Lett.* **54**, 502 (1985).
- [107] P. Candelas, G. T. Horowitz, A. Strominger and E. Witten, “Vacuum configurations for superstrings,” *Nucl. Phys. B* **258**, 46 (1985).
- [108] J. Distler and S. Garibaldi, “There is no ‘theory of everything’ inside E_8 ,” *Commun. Math. Phys.* **298**, 419 (2010).

- [109] A. Connes, *Noncommutative Geometry* (Academic Press, 1994).
- [110] A. H. Chamseddine and A. Connes, “The spectral action principle,” *Commun. Math. Phys.* **186**, 731 (1997).
- [111] A. Connes and H. Moscovici, “Type III and spectral triples,” in *Traces in Number Theory, Geometry and Quantum Fields* (Vieweg, 2008).
- [112] M. F. Atiyah, V. K. Patodi and I. M. Singer, “Spectral asymmetry and Riemannian geometry. I,” *Math. Proc. Camb. Phil. Soc.* **77**, 43 (1975).
- [113] C. G. Callan and J. A. Harvey, “Anomalies and fermion zero modes on strings and domain walls,” *Nucl. Phys. B* **250**, 427 (1985).
- [114] A. Kitaev, “Anyons in an exactly solved model and beyond,” *Ann. Phys.* **321**, 2 (2006).
- [115] L. B. Anderson, J. Gray, A. Lukas and E. Palti, “Two hundred heterotic standard models on smooth Calabi–Yau threefolds,” *Phys. Rev. D* **84**, 106005 (2011).
- [116] V. Braun, Y.-H. He, B. A. Ovrut and T. Pantev, “A heterotic standard model,” *Phys. Lett. B* **618**, 252 (2005).
- [117] H. Ooguri, C. Vafa and E. Verlinde, “Hartle–Hawking wave-function for flux compactifications: the entropic principle,” *Lett. Math. Phys.* **74**, 311 (2005).
- [118] S. Gukov, K. Saraikin and C. Vafa, “The entropic principle and asymptotic freedom,” *Phys. Rev. D* **73**, 066010 (2006).
- [119] J.-P. Serre, *A Course in Arithmetic*, Graduate Texts in Mathematics **7** (Springer, 1973).
- [120] J. Milnor, “Eigenvalues of the Laplace operator on certain manifolds,” *Proc. Natl. Acad. Sci. USA* **51**, 542 (1964).
- [121] H. Davenport and H. Heilbronn, “On the zeros of certain Dirichlet series,” *J. London Math. Soc.* **s1-11**, 181 (1936).
- [122] J.-B. Bost and A. Connes, “Hecke algebras, type III factors and phase transitions with spontaneous symmetry breaking in number theory,” *Selecta Math. (N.S.)* **1**, 411 (1995).
- [123] A. Connes, “Trace formula in noncommutative geometry and the zeros of the Riemann zeta function,” *Selecta Math. (N.S.)* **5**, 29 (1999).
- [124] M. V. Berry and J. P. Keating, “The Riemann zeros and eigenvalue asymptotics,” *SIAM Rev.* **41**, 236 (1999).
- [125] Y.-H. He, V. Jejjala and D. Minic, “From Veneziano to Riemann: a string theory statement of the Riemann hypothesis,” *Int. J. Mod. Phys. A* **31**, 1650201 (2016).
- [126] B. Riemann, “Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse,” *Monatsber. Königl. Preuss. Akad. Wiss. Berlin*, 671 (1859).
- [127] B. S. DeWitt, “Quantum theory of gravity. I,” *Phys. Rev.* **160**, 1113 (1967).
- [128] J. B. Hartle and S. W. Hawking, “Wave function of the Universe,” *Phys. Rev. D* **28**, 2960 (1983).
- [129] A. Vilenkin, “Creation of universes from nothing,” *Phys. Lett. B* **117**, 25 (1982).
- [130] M. Bojowald, “Loop quantum cosmology,” *Living Rev. Relativ.* **11**, 4 (2008).
- [131] A. H. Guth, “Inflationary universe,” *Phys. Rev. D* **23**, 347 (1981).
- [132] R. Penrose, *Cycles of Time* (Bodley Head, 2010).
- [133] J. Khoury, B. A. Ovrut, P. J. Steinhardt and N. Turok, “The ekpyrotic universe,” *Phys. Rev. D* **64**, 123522 (2001).
- [134] C. Rovelli, “Relational quantum mechanics,” *Int. J. Theor. Phys.* **35**, 1637 (1996).
- [135] L. Smolin, “Temporal naturalism,” *Stud. Hist. Phil. Mod. Phys.* **52**, 86 (2015).

- [136] H. Everett III, “‘Relative state’ formulation of quantum mechanics,” *Rev. Mod. Phys.* **29**, 454 (1957).
- [137] M. M. Vopson, “Is gravity evidence of a computational universe?,” *AIP Advances* **15**, 045035 (2025).
- [138] M. M. Vopson, “The mass-energy-information equivalence principle,” *AIP Advances* **9**, 095206 (2019).
- [139] B. Weiner, “Little Bangs: the holographic nature of black holes,” *IPI Letters* **3**(3), 34 (2025). doi:10.59973/ipil.177
- [140] B. Weiner, “Destroying the multiverse: entropy mechanics in causal diamonds,” *IPI Letters* **3**(5), 26 (2025). doi:10.59973/ipil.269
- [141] B. Weiner, “Resolution of the BAO sound-horizon discrepancy via an information-theoretic extension of Λ CDM,” *IPI Letters* **4**(3), 1 (2026). doi:10.59973/ipil.328
- [142] B. Weiner, “Mass–energy–information equivalence in an expanding universe,” submitted to *IPI Letters* (2026).