Unified Geometrization of Standard Model Parameters: A Holographic Fiber Theory (HFT) Framework

C.Y. Hsieh

Independent Researcher

May 17, 2026 [v16]

Abstract

Holographic Fiber Theory (HFT) is a parameter-free topological derivation of the Standard Model's free constants and the leading dark-sector ratios from a single substrate — the Hopf bundle $S^3 \xrightarrow{S^1} S^2$ realised as two-strand framed fibers carrying a trivalent weaving rule on its S^2 base with B_3 vertex braiding. A single global $\mathbb{Z}_2$ symmetry-breaking event — the chirality lock — converts the loose pre-EWSB substrate into the post-EWSB vacuum $\mathcal{V}$.

The locked mesh's topology forward-derives three geometric coupling constants: $\sin^2\theta_W = 30/128$, $\alpha^{-1} = 137$, and the vacuum angle $\theta_v = 1/32$. The first two fix the dark-to-baryonic ratio via writhon excitations crystallised at EWSB, giving $\Omega_c/\Omega_b \approx 5.35$.

The vacuum angle θ_v enters as a geometric pre-stress correction in the Higgs vev calculation, with the downstream SM mass cascade inheriting this θ_v dependence. HFT predicts $M_Z = 91.14$ GeV and $M_p = 937.7$ MeV, matching observation to 0.05% and 0.06% respectively; Table 2 lists the leading-order results across the spectrum.

1 Introduction

Holographic Fiber Theory (HFT) is a topological framework in which the Standard Model's free parameters are forward-derived from a single substrate: a Hopf bundle [2] $S^3 \xrightarrow{S^1} S^2$ with a trivalent weaving rule on its base, distinguished by a single chirality-locking event (Sec. 3). Grounded in the Holographic Principle [5, 4], the substrate admits two complementary quantitative readings that converge on the same closed-form values for the SM parameters: a discrete reading through state counting on the trivalent cell complex, and a continuous reading through Călugăreanu invariants [14, 15, 16].

The weaving rule forward-derives three dimensionless structural ratios that set the Standard Model's coupling strengths: the fine-structure constant $\alpha^{-1} = 137$, the Weinberg angle $\sin^2\theta_W = 30/128$, and the vacuum angle $\theta_v = 1/32$. With these as picture-internal coupling anchors, the rest of the paper forward-derives the SM mass spectrum and the leading dark-sector observables.

2 Methodological Position

Position Statement

(1) Continuous space, discrete state-transition rules. The physical arena of HFT is the *continuous* S^2 base of the Hopf bundle. The trivalent mesh, framed fibers, two-strand threads, and B_3 braids of Sec. 3 are neither a discretization of that base nor physical objects laid on top of it: they encode the substrate's *discrete state-transition rules* — the "weaving law" of space — that update the phase space $\mathcal{Q}$ whenever a state transition fires.

(2) The weaving law is scale-free. The Hopf-trivalent picture carries no intrinsic length scale and is not localized at any particular physical scale. All dimensionless structural constants appear as geometric ratios between picture-internal counts and inherit this scale-free character. Concrete instances include the post-lock skeleton $N_{\text{skeleton}} = 137$ underlying α^{-1}, the Weinberg angle $\sin^2\theta_W = 30/128$, and the further ratios derived from cell-complex counting; all are scale-invariant quantities forward-derived from the topology alone.

(3) Dynamics as a single classical tension field. The substrate's dynamics is the classical elastic response of the chirality-locked S^2 to perturbations — a single tension field, not an operator-valued quantum field; Lagrangian densities, gauge actions, and Hilbert-space structures are emergent IR descriptions. Wave propagation on the continuous base is the default mode. When local stress exceeds the vacuum threshold $\mathcal{V}$, the discrete weaving law activates at the relevant vertex and the excitation acquires particle character; particle-ness is interaction-triggered, not a permanent substrate attribute.

(4) Metrics are measurement conventions, not ontology. A direct consequence of (2)'s scale-freeness: HFT has *one* physical anchor — the Planck mass m_P — and every other measurement scale (length, energy, time) is bridged from it by $\hbar$, c, k_B together with the dimensionless topological ratios. Metric values are therefore measurement conventions rather than substrate features; if the anchor drifts, all observations drift together coherently, and "the substrate changing" is in principle indistinguishable from "the ruler shifting against an unchanging substrate."

3 The Theoretical Framework

3.1 The Hopf-Trivalent Substrate

Framed Hopf fibers on the trivalent mesh. The Hopf bundle $S^3 \xrightarrow{S^1} S^2$ assigns an S^1 circle to every point of S^2, threading perpendicular to the local tangent plane. The base S^2 carries a trivalent weaving rule whose combinatorics are those of the dual of the densest 2D packing — the (honeycomb) trivalent mesh, with $N_v = 3$ edges meeting at every vertex. At a *vertex point* the fiber is a single S^1 circle — the *vertex axial fiber* that the chirality lock of Sec. 5 tilts to produce the Weinberg angle. Along each *edge* the fiber acquires a framing as the base point moves and appears as a *two-strand framed thread*: two parallel strands of common handedness displaced along the framing direction, the minimal structure carrying the $\mathbb{Z}_2$ handedness label on which the chirality lock acts. At the *vertex neighbourhood* the three incoming edges deliver $3 \times 2 = 6$ strands that re-pair within their fibers into three composite outgoing fibers, forming a *three-strand knot*.

Vertex B_3 braid and vertex-fiber identification. The three composite fibers at each vertex braid under B_3 at the composite-fiber level rather than B_6 at the strand level. The full twist $\Delta^2 = (\sigma_1\sigma_2)^3$ generates the $\mathbb{Z}_3$ centre of B_3, which coincides topologically with the geometric C_3 rotational symmetry of the vertex — braid centre and $120°$ rotation acting on the same three fibers. This *vertex-fiber identification* fixes $N = k = 3$ in the Witten B_n–Chern-Simons correspondence [3], structurally setting the Chern-Simons truncation level used throughout the paper.

Phase space. The substrate's phase space comprises five topology-true degrees of freedom:

$$\mathcal{Q} \;=\; \mathbb{R}^+ \times S^2 \times T \times S^1_{\text{Hopf}}, \tag{1}$$

- $\mathbb{R}^+$ **(1 d.o.f.)**: substrate tension scalar ρ at the base point.

- S^2 **(2 d.o.f.)**: vertex spatial position p on the Hopf-bundle base.
- T **(1 d.o.f.)**: tangent direction θ at p, the local state-update propagation direction on S^2.
- S^1_{Hopf} **(1 d.o.f.)**: Hopf-fiber winding phase ψ on the axial S^1.

The $S^2 \times S^1_{\text{Hopf}}$ content assembles globally into the Hopf bundle total space S^3 ($S^3 \to S^2$, Euler class 1). Since $S^3 \cong \mathbb{R}^3 \cup \{\infty\}$ (stereographic identification), the substrate's spatial content presents itself to a macroscopic observer as ordinary $\mathbb{R}^3$.

Macroscopic emergence from $\mathcal{Q}$. The five DOFs of $\mathcal{Q}$ render macroscopically as the observables an observer reads off:

- *Spatial* $\mathbb{R}^3$ (3 d.o.f.): S^2 anchor supplies (X, Y); Hopf-phase mismatch supplies the depth Z — orthogonal phases at the same S^2 node read as Z-separated (non-interacting), phase-resonant configurations as spatially adjacent (strongly interacting).
- *Spin / helicity* (1 d.o.f.): tangent direction T — the rotational orientation at each base point.
- *Energy* (1 d.o.f.): tension scalar ρ — the substrate stress an observer measures as mass-energy density.

Geometric reading of spin's puzzles. Spin is the most counterintuitive of the three emergences above. HFT's reading is geometric: spin is the substrate's local T tangent direction at the particle's anchor, not the particle's internal rotation (point particles cannot rotate; classical electron self-rotation would exceed the speed of light). Three QM puzzles dissolve structurally:

- *Discrete spin values.* Different framing topologies of the trapped configuration give different periodicities of T rotation: spin-0 (axial ρ-mode, no T coupling), spin-$\frac{1}{2}$ (the B_2 doubled-strand thread's $\mathbb{Z}_2$ framing twist, Möbius-like, $720°$-periodic), spin-1 (single composite fiber, $360°$-periodic), spin-2 (graviton H^0 tensor with two-leg T coupling, $180°$-periodic).
- *Stern-Gerlach two-valuedness.* The continuous S^1-valued T, when measured along a probe axis, binarizes to $\pm$ at Nyquist resolution. Spin quantization is the binarization of measurement, not T's intrinsic discreteness.
- *Helicity invariance for massless particles.* Massless excitations propagate at the substrate tension speed c with no rest frame; their T direction is topologically bound to the propagation direction, so $T \cdot \hat{p}$ (helicity) is the only Lorentz-invariant spin descriptor.

3.2 Chirality lock: a brief geometric picture

The substrate undergoes a single global $\mathbb{Z}_2$ symmetry-breaking transition — the *chirality lock* — that converts a loose, parity-symmetric configuration into a tightened, chirality-fixed mesh.

Pre-lock state. Each Hopf fiber admits a mirror image of opposite handedness, the two strands sharing each edge sit parallel at near-zero tension, the S^1 fiber sits perpendicular to the S^2 tangent plane at every vertex, and the trivalent mesh is loose with no stress communication between fiber and base.

The lock event. A single global event picks one handedness for every fiber and winds each pair of parallel edge strands mutually once ($Lk = 1$, the minimum non-trivial topological invariant). The tightening at each edge perturbs the geometry of the vertices it joins, parameterised by the Weinberg angle θ_W and the vacuum angle θ_v (Sec. 4) — two complementary readings of the same vertex deformation.

Post-lock state. The post-lock state is the **chirality-locked EW vacuum** $\mathcal{V}$ (Sec. 5). The mutual winding stretches each edge's two-strand thread slightly under axial tension. Both the per-cell Nyquist budget and the coupling angles θ_W, θ_v run with this tension as it varies with probe scale.

Post-lock stress geometry. In the locked mesh, Călugăreanu's identity $Lk = Tw + Wr$ forces every fiber twist to be balanced by writhe at the local vertex. Twist along each edge thread and twist along each vertex axial S^1 therefore both appear as writhe at the vertices, coupling fiber stress to base stress.

3.3 Per-cell Nyquist state count

The holographic principle motivates encoding the substrate's bulk state on a 2D base [5, 4]; HFT implements this by Nyquist-binarizing [11, 12] the phase space $\mathcal{Q}$ and aggregating post-lock effective bits over the trivalent cell complex of S^2 weighted by the bulk $V{:}E$ sharing ratio. The binarization is a Nyquist-resolution computational tool: the underlying topology is scale-free (Sec. 2), and the discrete count below coarse-grains naturally into continuous magnitudes at larger scales.

Post-lock effective DOFs. The vertex hosts 3 raw bits (S^2 position and S^1_{Hopf} phase ψ) and the edge hosts 2 raw bits (T tangent and ρ tension), per the phase-space decomposition of Sec. 3. The chirality lock fixes the mesh's positional and connectivity content: vertices settle into specific lattice sites, the C_3 vertex-fiber identification picks a single sub-class out of the S^2 anchor, and each edge's tangent direction is determined by its two endpoint vertices. The raw DOFs accordingly reduce: *Vertex*: of the 3 raw bits, 1 is absorbed by the lock (the discrete lattice label); the residual S^2 sub-class plus the S^1_{Hopf} phase ψ survive as state-carrying — $V_{\text{eff}} = 2$. *Edge*: of the 2 raw bits, the tangent direction T is absorbed by the lock-determined edge connectivity (1 bit); the tension ρ survives — $E_{\text{eff}} = 1$.

Per-cell aggregation. In the hex-bulk limit each vertex is shared by 3 cells and each edge by 2, so one cell aggregates $V{:}E = 2{:}3$ effective elements:

$$b_{\text{cell}} \;=\; 2\,V_{\text{eff}} + 3\,E_{\text{eff}} \;=\; 2{\cdot}2 + 3{\cdot}1 \;=\; 7, \tag{2}$$

giving

$$\boxed{N_{\text{Nyquist}} \;=\; 2^{b_{\text{cell}}} \;=\; 2^7 \;=\; 128.} \tag{3}$$

4 Geometrization of Coupling Constants

4.1 Post-lock Cohomology Channels

The chirality lock tightens the trivalent mesh and partitions the per-cell Nyquist budget $N_{\text{Nyquist}} = 128$ across the four de Rham cohomology classes H^0, H^1, H^2, H^3 of the per-cell phase space. Each class carries a distinct topological mode type into which the substrate's elastic energy localizes, labelling one IR force channel:

Class	Substrate content	IR identification
H^0	long-wavelength tension on S^2 base	Sec. 8
H^1	phase holonomy across S^1 fiber	$U(1)_Y$ hypercharge
$H^1 \leftrightarrow H^2$	twist–writhe exchange	$SU(2)_L$ weak
H^3	vertex braid linking	$SU(3)_C$ strong

The coupling constants $\sin^2\theta_W$ and α^{-1} derived in the next two subsections are forward outputs of how the lock allocates Nyquist slots among these four channels.

4.2 The Weinberg Angle

θ_W as the $H^1 \leftrightarrow H^2$ coupling parameter. Geometrically, the Weinberg angle [8] θ_W measures the post-lock vertex deformation read in two orthogonal frames: viewed from an incoming edge, the vertex appears tilted relative to the S^2 tangent plane by θ_W, with $\sin^2\theta_W$ the squared projection of this tilt onto the base; viewed along the axial S^1 fiber, the vertex appears phase-rotated by an angle θ_v, the $1/32$ sub-quantum derived in Sec. 4.4. The integer N_{weak} counting Nyquist slots in the weak channel admits two complementary derivations, given below.

Discrete view: equipartition deviation by the $\mathbb{Z}_2$ quantum. Equipartition across the four de Rham cohomology channels H^0, H^1, H^2, H^3 of the per-cell phase space (the labels of the substrate's elastic-mode channels, Sec. 4.1) — the symmetric distribution that obtains in the absence of any directional preference — gives $128/4 = 32$ slots per channel and a symmetric baseline $\sin^2\theta_W = 1/4$. The chirality lock breaks this symmetry through a Călugăreanu twist–writhe exchange (Sec. 5): the Hopf-protected $Lk = 1$ conservation forces $\Delta Tw + \Delta Wr = 0$, and a $\mathbb{Z}_2$ chirality flip drives the integer-quantized transition $Tw = +1 \to -1$ on a single edge, with $\Delta Tw = -2$ compensated by $\Delta Wr = +2$ at the endpoint vertices. Two integer Nyquist slots are displaced from the weak channel:

$$N_{\text{weak}} = 32 - 2 = 30. \tag{4}$$

The integer 2 is fixed by the Călugăreanu twist quantum — the minimal twist transition compatible with Lk conservation. No continuous tuning is involved.

Geometry view: marked–unmarked correlation. The chirality lock binarizes each of the $\dim(\text{vertex}) = 3$ vertex DOFs (Sec. 3) and marks $\dim(\text{postlock}) = 3$ of the $2^{\dim(\text{vertex})} = 8$ vertex configurations as chirality-loaded — the Hamming-weight-1 patterns picked up at one bit per vertex dimension — leaving the complementary $2^{\dim(\text{vertex})} - \dim(\text{postlock}) = 5$ unmarked. The weak channel mediates the substrate's twist–writhe exchange (Sec. 4.1), engaging both sectors jointly: its per-cell slot count is the product of the marked and unmarked vertex fractions weighted by the Nyquist budget,

$$N_{\text{weak}} = \frac{\dim(\text{postlock})}{2^{\dim(\text{vertex})}} \cdot \frac{2^{\dim(\text{vertex})} - \dim(\text{postlock})}{2^{\dim(\text{vertex})}} \cdot N_{\text{Nyquist}} = \tfrac{3}{8} \cdot \tfrac{5}{8} \cdot 128 = 30. \tag{5}$$

The marked ($3/8$) and unmarked ($5/8$) fractions partition the binarized vertex space ($3/8 + 5/8 = 1$); the marked fraction $3/8$ recovers the SU(5) GUT prediction when only the marked sector is sampled (high-scale projection limit).

Convergence and errors. The two readings converge on the bare skeleton $\sin^2\theta_W = 30/128 \approx 0.2344$. The picture-internal BASE noise dressing (App. D) brings the IR prediction to $\sin^2\theta_W(M_Z) \approx 0.23123$, matching the observed $0.23121(4)$ at ~ 91 ppm, well within the experimental 1σ.

4.3 The Fine-Structure Constant (α) from Lock-Induced Edge Stretch

Geometrically, α in HFT is the density of S^1_{Hopf} phase per unit S^2 base area — the rate at which fiber holonomy is packed into the substrate. Since the H^1 channel (electromagnetism) lives on the S^1_{Hopf} holonomy modes (Sec. 9.3), this density sets the strength of the H^1 coupling to matter: α is its picture-internal coupling constant. α^{-1} is therefore the per-cell Nyquist budget plus the slots opened by the chirality lock. Both readings forward-derive the same skeleton:

$$\alpha^{-1}_{\text{skeleton}} = N_{\text{Nyquist}} + \Delta N_{\text{Nyquist}} = 128 + 9 = 137. \tag{6}$$

Discrete view: defect-class count opened by the lock. The lock breaks the pre-lock $\mathbb{Z}_2$ chirality symmetry — previously the per-strand framing twist transmuted freely between $\pm Tw$, leaving the Călugăreanu invariants (Lk, Tw, Wr) unanchored — and the labels acquire well-defined values. Per vertex this opens $9 = 3 \times 3$ stable defect classes (three C_3 cyclic-permutation sectors $\times$ three lock-stable framings $Tw \in \{-1, 0, +1\}$, with $|Tw| \geq 2$ decaying via $Lk = Tw + Wr$): $\Delta V_{\text{classes}} = +9$. Per edge the global chirality fix collapses the raw label space from $2 \times 3 = 6$ to 3: $\Delta E_{\text{classes}} = -3$. Aggregating by the $V{:}E = 2{:}3$ sharing ratio,

$$\Delta_{\text{cell}} \;=\; 2\,(+9) + 3\,(-3) \;=\; +9, \tag{7}$$

reproducing $\Delta N_{\text{Nyquist}} = 9$.

Geometry view: helical elastic stretch from chirality lock. The lock activates strand mutual winding ($Lk = 1$) on every edge, deforming each strand from straight (pre-lock) into a helix (post-lock). The strand acquires axial elastic stretch ε relative to the pre-lock length, encoding the lock-induced tension increment on the substrate.

Two picture commits fix the helix geometry:

- Helix radius from trivalent vertex $\mathbb{Z}_3$ symmetry: $r \propto N_v$
- Axial pitch from holographic area–state correspondence: $L_0 \propto \sqrt{N_{\text{Nyquist}}}$

The small-angle Pythagorean stretch is

$$\varepsilon \;=\; \tfrac{1}{2}\left(\frac{2\pi r}{L_0}\right)^2 \;=\; \frac{N_v^2}{2\,N_{\text{Nyquist}}} \;=\; \frac{9}{256}. \tag{8}$$

By the holographic area–state correspondence, the per-cell Nyquist increment equals twice the strand stretch:

$$\Delta N_{\text{Nyquist}} \;=\; 2\varepsilon \cdot N_{\text{Nyquist}} \;=\; N_v^2 \;=\; 9, \tag{9}$$

where the N_{Nyquist} factor cancels between the holographic weight and the helix axial-pitch normalization. The increment is therefore set *purely* by the trivalent geometry $N_v = 3$, independent of the N_{Nyquist} value — a structurally distinct continuous-geometry route from the Hopf-side enumeration.

Convergence and errors. The two readings converge on the bare skeleton $\alpha^{-1}_{\text{skeleton}} = 137$. The lock event fixes $N_{\text{Nyquist}} = 128$ and the discrete IR running adds $+9$ slots to give $N_{\text{skeleton}} = 137$. The picture-internal BASE noise dressing (App. D) brings the prediction to $\alpha^{-1}(0) \approx 137.0360$, matching the observed 137.035999 at ~ 0.06 ppm.

4.4 The Vacuum Angle (θ_v)

The vacuum angle θ_v is the per-vertex phase rotation of $\mathcal{V}$ along the axial S^1 fiber — a geometric pre-stress induced by the chirality lock's horizontal rotation between the S^2 base and the S^1_{Hopf} fiber. The Higgs vacuum amplitude v is dominated by the lock's topological commitment; θ_v enters as a structurally fixed pre-stress correction in the vev calculation (App. A.1), with the downstream SM mass cascade inheriting this θ_v dependence through the structural primitives that aggregate it.

Discrete view: lock-induced uniqueness on vertex $\times$ cohomology enumeration. Each vertex carries a state labelled by a binarized vertex configuration ($2^{\text{dim(vertex)}} = 8$ options on the 3 vertex bits) and a de Rham cohomology channel ($N_{\text{coh}} = 4$ options; Sec. 4.1), giving $8 \times 4 = 32$ pre-lock vertex states. The chirality lock breaks this degeneracy through two selections:

1. *Vertex factor: 1 of 8.* Of the eight binarized vertex configurations, the lock marks the three Hamming-weight-1 patterns (one bit chirality-loaded per vertex dimension; Sec. 4.2, Geometry view). The vertex-fiber identification (Sec. 3) then picks the single Hamming-weight-1 pattern aligned with the axial S^1_{Hopf} direction, the lock-privileged vertex dimension.

2. *Cohomology sector: 1 of 4.* The lock-aligned phase content lives in the H^1 phase-holonomy channel (Sec. 4.1), singling out one of the four cohomology channels as the vacuum's home.

The vacuum vertex is the unique configuration consistent with both selections:

$$\theta_v \;=\; \frac{1}{2^{\dim(\text{vertex})}\cdot N_{\text{coh}}} \;=\; \frac{1}{8\cdot 4} \;=\; \frac{1}{32}. \tag{10}$$

The numerator 1 counts the singular vacuum vertex picked out by the two-stage lock selection; the denominator factorizes as (vertex configurations) $\times$ (cohomology channels).

Geometry view: cohomology fraction of the lock-scale Nyquist capacity. The Călugăreanu two-slot displacement that links α^{-1} and θ_W (App. D) extends to a third coupling: θ_v is the fractional weight of one cohomology channel in the lock-scale Nyquist capacity, with $N_{\text{coh}} = 4$ channels (Sec. 4.1) sharing the per-cell budget $N_{\text{Nyquist}} = 128$ (the picture-internal lock-scale value of α^{-1}):

$$\theta_v \;=\; \frac{N_{\text{coh}}}{\alpha^{-1}(M_Z)} \;=\; \frac{N_{\text{coh}}}{N_{\text{Nyquist}}} \;=\; \frac{4}{128} \;=\; \frac{1}{32}. \tag{11}$$

Equivalently, using the α–θ_W identity to substitute $\alpha^{-1}(M_Z)$,

$$\theta_v \;=\; 2\big(\tfrac{1}{4} - \sin^2\theta_W\big) \;=\; \tfrac{1}{32},$$

deriving θ_v purely from the two electroweak couplings and the cohomology partition N_{coh}, without invoking discrete vertex-state enumeration. The three coupling constants α^{-1}, θ_W, θ_v are not independent picture commits: they are linked by the lock-scale geometry of a single chirality-lock event.

Convergence and errors. The two readings converge on $\theta_v = 1/32$ exactly. As a primary observable θ_v has no direct measurement; its downstream manifestation is the SM mass spectrum, which inherits the percent-level residuals propagated through the cascade.

5 Mass Generation from the Locked Mesh

Mass in HFT is the elastic energy of a localised excitation on the chirality-locked mesh. Each massive state is a trapped configuration of the substrate's storage modes — ρ-amplitude, framing twist (Tw), base writhe (Wr), or vertex knot — held in place by the mesh-wide tension connectivity that the lock activates.

Mass generation from the locked substrate tension. The chirality lock drives the S^2 substrate from its pre-lock loose state to a tightened post-lock configuration in which every vertex grip is closed (Sec. 3). The residual tension the locked S^2 can sustain before the grip would unwind — the saturation tension of the chirality-locked substrate — is the Higgs VEV v. The picture-internal forward derivation from the lock-event bounce action (App. A.1) gives

$$v \;=\; 244.7 \text{ GeV} \quad \text{vs. observed} \quad 246.22 \text{ GeV} \quad (0.62\% \text{ match}),$$

with no fitted parameters. All SM masses then read as elastic excitations on this tension background: the Higgs M_H as longitudinal ρ-mode oscillation of v; the weak gauge bosons M_W, M_Z as transverse stress-transfer modes; and the charged leptons, quarks, and Majorana neutrinos as trapped vertex / edge configurations of fixed topological character. The quantitative cascade from v to every entry of Table 2 is derived in App. A.

Neutrino as an edge mode. Neutrinos live on edges, not vertices: each is an intrinsic Twist mode of the B_2 doubled-strand framed thread (Sec. 3). With no winding around any axial fiber, the neutrino carries no electric charge. The doubled-helix topology admits exactly three such Twist modes, furnishing the three Majorana mass eigenstates (mass scale in App. A.3).

Charged lepton as the vertex axial strand. The charged lepton lives at a vertex: a single strand winds around the vertex's axial S^1 fiber, and the complete axial winding bears the elementary -1 electric charge unit. The three generations (e, μ, τ) correspond to increasing vertex engagement of this axial strand — a single $\mathbb{Z}_3$ class for e, the full cyclic orbit (three classes) for μ, and the complete nine-class vertex for τ (App. A.2).

Quark charge spectrum and anomaly cancellation. Quarks share the same vertex as the charged lepton, but as the three composite fibers on the S^2 base. These fibers divide the leptonic -1 axial baseline in equal thirds, fixing the per-strand baseline at $-1/3$; the two flavour types then differ by one integer unit of the strand's intrinsic twist around the axial fiber: *down-type* ($Tw = 0$, electric charge $-1/3$) and *up-type* ($Tw = +1$, electric charge $-1/3 + 1 = +2/3$). The SM fractional charge pattern $\{-1, 0, +2/3, -1/3\}$ across one lepton-quark generation is therefore a topological accounting, not a fitted assignment. The gravity-mixed anomaly condition

$$\sum_f Q_f \;=\; (-1) + 0 + 3 \cdot (+2/3) + 3 \cdot (-1/3) \;=\; 0$$

follows directly, with $3 = N_v$ (trivalent vertex) playing the role of colour multiplicity by construction.

Sakharov conditions and B vs L natural separation. The Sakharov conditions [13] are satisfied structurally by the chirality-locking transition: the lock is itself the B-violating, out-of-equilibrium, CP-distinguishing transition, with no separate sphaleron mechanism required. The asymmetry preferentially carries baryon rather than lepton number due to a dimensionality mismatch between channels: at each vertex, the B_3 braid re-pairing of the three trivalent fibers CP-biases the structure-locked content across three strands (baryon channel, multi-DOF), while the axial fiber carries lepton number as a single integer Tw winding (single-DOF channel with no internal partition). HFT therefore predicts a baryon-asymmetric universe at the EW scale without an accompanying leptogenesis mechanism.

6 The Action Budget

The Total Topological Action Budget S_E is the post-lock skeleton fraction carried by chirality-loaded content. Two complementary readings forward-derive its value:

Discrete view: chirality-marked configuration enumeration. Each vertex has $2^{\mathrm{dim(vertex)}} = 8$ binarized configurations across $N_{\mathrm{coh}} = 4$ cohomology channels — 32 pre-lock options per vertex. The chirality lock marks one configuration per vertex dimension as chirality-loaded across all cohomology channels, giving dim(postlock) $\cdot\, N_{\mathrm{coh}} = 3 \cdot 4 = 12$ marked configurations per vertex. With the per-cell skeleton $N_{\mathrm{skeleton}} = 137$ distributed evenly across all 32 options, summing over the marked subset:

$$S_E \;=\; 12 \times \frac{N_{\mathrm{skeleton}}}{32} \;=\; 51.375. \tag{12}$$

Geometry view: post-lock skeleton projected by $\sin^2 \theta_W|_{\mathrm{lock}}$. The chirality lock projects the post-lock skeleton onto its chirality-loaded fraction via the lock-scale Weinberg angle (Sec. 4.2, Geometry view):

$$S_E \;=\; N_{\mathrm{skeleton}} \times \sin^2 \theta_W\big|_{\mathrm{lock}} \;=\; 137 \times \tfrac{3}{8} \;=\; 51.375. \tag{13}$$

The geometry reading uses $\sin^2\theta_W|_{\text{lock}} = \text{dim(postlock)}/2^{\text{dim(vertex)}} = 3/8$ as the continuous vertex-factor projection ratio.

Convergence. The two readings converge on $S_E = 51.375$ exactly.

Over-determination check: electron mass from direct S_E exponentiation. The action budget $S_E = 51.375$ and the per-DOF θ_v pre-stress $\dim(\mathcal{Q})\cdot\theta_v = 5/32$ accumulated over the five dimensions of $\mathcal{Q}$ (Sec. 3) together set the bounce-action depth between the Planck scale and the electron mass:

$$m_e \;\approx\; m_P \cdot \exp[-\,(S_E + \dim(\mathcal{Q})\cdot\theta_v)] \;=\; m_P \cdot e^{-(51.375+5/32)} \;\approx\; 0.509 \text{ MeV},$$

matching the observed 0.5110 MeV to 0.3%. The v-anchored mass-ladder route (App. A.2) reaches the same scale via $m_e = \Lambda_{\text{QCD}}/(N_v R^2) \approx 0.507$ MeV (0.7%). The two routes are picture-internally independent — the first a direct exponentiation of the action-budget and pre-stress contributions, the second cascading from v through the QCD scale and the lepton ladder — and converge to the same sub-percent residual at m_e, providing an over-determination check that both routes are consistent with the same picture-internal structure.

S_E as EWSB latent heat. With the numerics over-determined above, S_E admits its physical reading: the chirality lock is HFT's EWSB transition, and S_E is the latent heat it releases — the post-lock action carried by chirality-marked configurations, redistributed into the visible-matter and dark-sector content via the partition $S_v + S_d = S_E$ (Sec. 7).

7 Dark Matter from the Writhon Sector

Dark sector as the Wr skeleton of the locked S^2. Of the post-EWSB latent heat $S_E = 51.375$ released by the chirality lock, the budget partitions

$$S_v + S_d \;=\; 8 + 43.375 \;=\; 51.375 \;=\; S_E \tag{14}$$

into a visible-matter portion $S_v = \sum_f Q_f^2 = 8$ (the SM L-side charge state count) and a dark-sector portion $S_d = 43.375$ (crystallising into a stable high-tension Wr skeleton on the locked S^2 — the dark-matter sector).

Writhon as a topological soliton. A writhon is a topological soliton on the post-locked mesh: a localised concentration of vertex Wr above the $\mathcal{V}$ baseline, stabilised against smooth unwinding by Călugăreanu $Lk = Tw + Wr$ conservation. The configuration is scale-free, distinguished only by a discrete minimum action quantum below which stability fails; two writhons that spatially overlap combine into a single writhon with summed Wr and enlarged support, the total action budget conserved. Like Skyrmions [20] or vortices in other settings, writhons carry mass-energy — the high-tension content stored above the $\mathcal{V}$ baseline — but do not couple to the SM gauge interactions: they are substrate-level structural features formed by non-uniform pre-EWSB cooling and frozen in by the chirality lock, not gauge-charged fields propagating on the substrate.

Dark-to-baryonic ratio: leading. With writhons stable against decay (Călugăreanu conservation), the post-lock dark sector clusters gravitationally as cold dark matter. The leading dark-to-baryonic ratio is fixed by the action-budget partition $S_d : S_v$:

$$\Omega_c/\Omega_b \;=\; \frac{S_d}{S_v} \;=\; \frac{43.375}{8} \;\approx\; 5.422, \tag{15}$$

in agreement with Planck [1] to $\sim 1\%$. The absolute mass scale of individual writhons is set by the minimum stable Wr quantum and is left as picture-internal follow-up; the ratio above is anchored on the action-budget partition and is independent of the absolute mass.

BASE noise correction on the visible side. The leading ratio S_d/S_v above treats both sides on equal footing. Visible matter consists of mesh excitations whose cascade-anchored mass scales and mesh-floor populations inherit the universal BASE noise dressing $\varepsilon \approx 1.34\%$ (Sec. 9, App. D); the Călugăreanu Lk topological dark sector does not enter the same dressing chain. Effectively $S_v \to S_v(1+\varepsilon)$, giving

$$\Omega_c/\Omega_b\big|_{\text{refined}} = \frac{S_d}{S_v(1+\varepsilon)} \approx \frac{43.375}{8 \times 1.0134} \approx 5.35, \tag{16}$$

matching Planck $\Omega_c/\Omega_b \approx 5.36$ [1] to 0.2%. The same ε that dresses the gauge couplings and the v-anchored mass cascade (App. D) therefore accounts for the dark-to-baryonic ratio refinement, with no additional mechanism beyond the visible/dark-sector asymmetry already built into the picture.

Long-term writhon dynamics. Writhons are stable on cosmological timescales: Călugăreanu conservation forbids smooth unwinding. The same scale-free soliton therefore describes both observed regimes of dark-matter phenomenology: at fast encounter speeds the merger timescale exceeds the crossing time and writhons pass through each other as a collisionless cold-DM gas (the Bullet Cluster signature [25]); under slow gravitational accretion they coalesce into the dark skeleton of galactic halos and the cosmic web. Microscopically discrete, macroscopically the dark sector forms a quasi-continuous sheet across the locked S^2 — a topological raft whose individual floats can be perturbed by substrate fluctuations but whose collective structure cannot.

8 Gravitational Field from Mesh Tensor Field

The gravitational field of General Relativity is identified, in HFT, with the H^0 tensor mode of the locked mesh — the substrate stress field whose continuum limit reproduces the Einstein field equations.

8.1 Mesh Tension Modes and Riemann Curvature

Riemann curvature measures holonomy failure: parallel-transporting a frame around a small loop γ yields a rotation $\delta\phi = R^a{}_{b\mu\nu}\,\delta A^{\mu\nu}$. On the locked HFT mesh, every mode of stored tension — the same Tw, Wr, and vertex-knot configurations that carry SM mass (Sec. 5) — contributes additively to the local rotational holonomy of a fiber orientation transported around γ, assembling into a single mesh stress tensor whose IR continuum limit is the Riemann tensor $R^a{}_{b\mu\nu}$. The Călugăreanu identity $Lk = Tw + Wr = \text{const}$ is the fiber-level topological conservation law — one component of the broader Bianchi-type structure governing the aggregate stress field, with $Tw \leftrightarrow \omega^a{}_\mu$ (spin connection) and the fiber-line Wr contribution mapping to $R^a{}_{b\mu\nu}$ at the planar level; vertex-knot contributions enter as discrete sources punctuating this smooth tension field.

8.2 Einstein Field Equations

Define the network free energy $F = E - T_{\text{mesh}}S$, where T_{mesh} is the effective temperature of the locked mesh's stochastic background (the substrate-level thermal agitation of the tension field). In the IR continuum limit the aggregate mesh tension energy (with Tw, Wr, and vertex-knot contributions all summed into a single stress field) and the mesh entropy become:

$$E \to \frac{c^4}{16\pi G}\int R\sqrt{g}\,d^4x, \qquad T_{\text{mesh}}\frac{\delta S}{\delta g_{\mu\nu}} \to T_{\mu\nu} \tag{17}$$

with the identification $G \equiv c^4/T_{\text{grav}}$, where T_{grav} is the mesh's transverse-mode propagation tension (and $c = \sqrt{T_{\text{grav}}/\rho_{\text{mesh}}}$ is the substrate mesh propagation speed). The equilibrium condition $\delta F/\delta g_{\mu\nu} = 0$ yields:

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \qquad G = \frac{c^4}{T_{\text{grav}}} \tag{18}$$

The cosmological term $\Lambda g_{\mu\nu}$ arises from the rigid baseline tension of the chirality-locked vacuum $\mathcal{V}$ and separates naturally on the left-hand side, decoupled from matter summation. $T_{\mu\nu}$ on the right-hand side measures only excitations above this rigid baseline; what QFT attributes to divergent vacuum fluctuations never enters the source term, so the fine-tuning paradox dissolves in HFT's primitive ontology — a classical tension field rather than field operators with vacuum loops.

8.3 Time Emergence from Substrate Dynamics

The geometric derivations above (Hopf bundle, Euclidean bounce action, Călugăreanu identity) all live on a positive-definite topological substrate carrying no primitive time direction. Time enters HFT at two complementary levels of emergence.

Microscopic: discrete tick. The substrate's weaving rule advances by discrete state updates. Each update creates exactly one newly distinguishable mesh state — one Spencer-Brown distinction [19] — and costs one per-radian phase action quantum (unity in natural units, $\hbar$ in SI). A "tick" is the counter of such updates: discrete, indivisible, with no underlying continuum. Time at this level is the ordinal index of state transitions, not a parameter along an axis.

Macroscopic: thermodynamic ordering. The arrow of time emerges thermodynamically — well before EWSB — at the moment an initially indistinguishable ensemble of strand configurations begins to acquire distinguishable mesh states, i.e. as soon as the substrate's coarse-grained entropy is well-defined and monotonically rising. The chirality lock (Sec. 5) is a much later, much sharper event on this already-running thermodynamic clock; the bounce action of App. A.1 is therefore a state count along an emergent ordering, not a quantity defined against an a priori time axis.

Space and time are ontologically decoupled. Spatial extension is the S^2 base of the Hopf bundle (a topological-combinatorial primitive); time is a tick count whose macroscopic arrow is entropy monotonicity (a thermodynamic-statistical primitive). The two have independent ontological status and obey independent counting rules — no $-+++$ signature, no light cone, no simultaneity surface at this level. The Lorentzian coupling of the IR continuum, mediated by the mesh tension speed $c = \sqrt{T_{\text{grav}}/\rho_{\text{mesh}}}$, is itself an emergent stitching of these two decoupled primitives, not a primitive feature of the substrate.

9 BASE Noise: Substrate Stochastic Background

The gravitational picture of Sec. 8 treated the locked mesh as a quiet tension field. Its residual dynamics — the stochastic floor surviving the chirality lock — is the picture-internal source of every fluctuation phenomenon HFT predicts. We name this floor **BASE noise** (Background Amplitude of Substrate Excitations). The label is deliberately broader than "SGWB": SGWB is one projection (the H^0 channel), the CMB photon bath is another, and QED-style vacuum fluctuations a third. All are readout faces of the same mesh excitation spectrum, and a fourth face — cosmic ruler shift — emerges once the substrate's scale structure is taken seriously.

9.1 Scale-Invariant Substrate; Single Operational Anchor

Picture-internal physics is entirely dimensionless — every result of Secs. 3–7 is a topological ratio invariant under global rescaling of the substrate. SI representation requires one external operational anchor; HFT uses m_P alone, with $\hbar, c, k_B$ propagating it to every other SI unit by fixed ratio. This admits an equivalent reading: rather than a fluctuating mesh against a fixed external ruler, the substrate is static and scale-free while the anchor m_P — itself a mesh-state readout — drifts, dragging every derived SI metric coherently with it. No local experiment can separate the two readings, but the ruler-frame view is the natural one once m_P is the sole dimensional input, and it is the reading the rest of this section turns into the falsifiable BASE noise prediction.

9.2 Source and Universal Amplitude

The BASE noise floor has two picture-internal anchors. Each vertex carries a residual $Wr \leftrightarrow Tw$ exchange amplitude at the rate $\sqrt{\alpha}$ — the Călugăreanu-conserved tension fluctuation between writhe and twist storage on the post-locked mesh, setting the local per-cell amplitude. Aggregated over the full post-lock skeleton, the cumulative BASE noise strength tightens to

$$\varepsilon = \frac{\ln(2\pi)}{N_{\text{skel}}} \approx 1.34\%, \tag{19}$$

with $\ln(2\pi)$ the Gaussian-shoulder floor of the per-cell state distribution and $N_{\text{skel}} = 137$ the cohomology-skeleton state count. ε is picture-internally forced and acts as the universal amplitude anchor across every BASE noise channel below.

9.3 Multi-Channel Manifestation

BASE noise projects onto the cohomology channels of the locked mesh, giving the same underlying excitation distinct observational faces:

Channel	Mesh content	Observational face
H^0 (transverse tensor)	Mesh-tension propagation modes	Stochastic gravitational wave background (SGWB)
H^1 (transverse vector)	Locked-mesh S^1 holonomy modes	CMB photon bath; QED vacuum fluctuations
$H^1 \leftrightarrow H^2$ (weak mixing)	$Wr \leftrightarrow Tw$ exchange at vertex tilt	Q-dependent θ_W dressing (App. D)
Anchor drift	m_P-readout perturbation	Cosmic ruler shift along photon paths

Table 1: BASE noise readout faces. All channels share the universal amplitude $\varepsilon \approx 1.34\%$ up to channel-specific propagation kernels.

Quantum fluctuation as threshold crossing. The two observational faces in the H^1 row are not distinct phenomena but two regimes of the same channel: the CMB is the macroscopic Boltzmann population of mesh modes well above the universal floor; "quantum fluctuation" is the stochastic rate at which local mode amplitudes briefly cross an apparatus-specific detection threshold near the floor. Virtual particle pairs in perturbative QFT calculations are mathematical bookkeeping for these crossing events rather than ontologically existing vacuum constituents. The same dual reading generalises across the other BASE noise channels.

9.4 Cosmic Observables as Ruler Shift

The anchor-drift channel is the most consequential for cosmology. Picture-internal physics is anchored to m_P alone, and every SI metric cascades from this single anchor through $\hbar, c, k_B$ (Sec. 9.1); BASE noise perturbations of the anchor therefore propagate uniformly through all derived metrics.

The substrate stays put while the rulers shift coherently. Photons travelling cosmological distances accumulate this shift along their path, and the accumulation is what an observer reads as cosmic expansion phenomenology:

- **Redshift** $(1+z)$: photon phase laid down with emission-epoch ruler, read by today's shifted ruler.
- **Type Ia time dilation** $(1+z)$: emission-epoch second vs. observer SI second, same ratio.
- **Tolman surface brightness** $(1+z)^{-4}$ [21]: $d_A \neq d_L$ as an algebraic identity from single-anchor isotropy, not from 4D metric stretching.
- **CMB temperature** $T(z) \propto (1+z)$: the k_B bridge shifted, emission-epoch thermal state read as today's $T_0(1+z)$.

All four are readouts of the same anchor drift along a single photon path. Etherington's distance-duality $d_L = (1+z)^2 d_A$ [22] holds as an algebraic identity from the single-anchor structure, without invoking a 4D Lorentzian metric. The appearance of cosmic expansion is the coherent shift of our measurement instruments, not an action of the cosmos.

JWST high-z maturity anomaly. ΛCDM accommodates 10^9–$10^{10} M_\odot$ mature galaxies at $z \sim 10$–14 [27] only with difficulty: its Friedmann back-integration ties each z to a specific cosmic-time budget (a few hundred Myr at $z \sim 14$), and the observed mass, dust, and metallicity exceed what hierarchical formation can produce in that window. In HFT this tension dissolves at the framing level rather than at the formation-physics level. Without metric expansion there is no Friedmann back-integration, so HFT carries no derived universe-age budget that high-z formation has to fit into; z_{obs} records integrated anchor drift along the photon path (Sec. 9.4), not a specific cosmic time at emission. Mature galaxies, dust enrichment, and metallicity at observed high z_{obs} are therefore phenomena to be observed, not anomalies to be explained.

Hubble tension as ruler-shift signature. Different H_0 probes sample different photon-path lengths through the BASE noise field: local Cepheid + SN Ia ladders accumulate little anchor drift; CMB acoustic-peak inference integrates the full Hubble radius. A systematic split between the two is the picture-expected manifestation, not an unresolved anomaly. The observed $\sim 8.3\%$ SH$_0$ES [26]–Planck split matches the numerical pattern $\sqrt{\alpha} \approx 0.0854$ (a notable coincidence; forward derivation deferred to follow-up). Intermediate-z probes ($z \sim 0.5$) should lie between the two endpoints, with offset scaling by path length.

The substrate-static reading further implies that the floor amplitude ε rules out asymptotic cosmic cooling: the substrate temperature is permanently anchored at T_{CMB} rather than drifting toward zero, eliminating the ΛCDM heat-death scenario at the conceptual level.

9.5 The CMB Temperature as Substrate Thermal Floor

The amplitude ε is a hard floor: as long as the post-lock mesh carries any excitation, that excitation's amplitude cannot fall below ε. In thermal equilibrium this amplitude floor translates into a temperature floor, fixed by the lightest stable mode the mesh can hold.

That lightest stable mode is the lowest Majorana neutrino eigenstate m_{ν_1} (App. A.3), residing on the $H^1 \leftrightarrow H^2$ weak-mixing channel with mode count $N_{\text{weak}} = 30$. Equilibrium is reached when the Boltzmann population of this mode saturates the inverse channel multiplicity, $e^{-m_{\nu_1}/k_B T_{\text{floor}}} = 1/N_{\text{weak}}$, giving

$$k_B T_{\text{floor}} = \frac{m_{\nu_1}}{\ln N_{\text{weak}}}. \tag{20}$$

Both inputs in Eq. (20) pick up ε-dressings. The mass m_{ν_1} inherits its dressing through θ_v in S_{bounce} (App. D), propagating down the cascade to give $m_{\nu_1}^{\text{dressed}} \approx 0.803$ meV. The mode count N_{weak} picks up the BASE noise floor as a one-sided upward shift, $N_{\text{weak}}^{\text{dressed}} = 30(1+\varepsilon) \approx 30.40$. Substituting both into Eq. (20) yields $T_{\text{floor}} \approx 2.728$ K, matching the observed $T_{\text{CMB}} = 2.7255$ K [24] to 0.08%.

Three picture-internal ingredients carry the result: ε establishes the floor and dresses both inputs; $m_{\nu_1}^{\text{dressed}}$ sets its absolute scale as the lightest stable cascade endpoint; In $N_{\text{weak}}^{\text{dressed}}$ supplies the cohomology-channel saturation factor that closes the formula.

The same floor temperature seeds the H^0 tensor channel through the universal amplitude ε: thermal gravitational-wave modes equilibrate at the same T_{CMB}, so the substrate stochastic gravitational-wave background is predicted to peak at $\nu_{\text{peak}} \approx 160$ GHz, coincident with the CMB Planck peak. This H^0/H^1 spectral peak coincidence has no ΛCDM counterpart and is sharpened as a specific falsifiability target in Sec. 10.

9.6 HFT Cosmology in Brief

The framework above admits a natural cosmological reading at the level of a hypothesis layer — of comparable epistemic standing to ΛCDM's inflation ansatz, not derived here from first principles.

Pre-lock yarn-ball expansion. The substrate's earliest configuration is a finite, indistinguishable *yarn-ball* state, with no singularity. Driven by the substrate rule, it expands and cools, with non-uniform local cooling laying down the inhomogeneities that later seed large-scale structure. This phase is HFT's counterpart of cosmological inflation: an ansatz upstream of direct observation, with detailed dynamics deferred to follow-up.

Chirality lock as EWSB transition. The expansion is terminated by the chirality lock (Sec. 3), a discrete topological phase transition at the EWSB scale. Its latent heat release into the mesh is HFT's structural counterpart of reheating; its closure freezes the substrate into the chirality-locked vacuum $\mathcal{V}$.

Post-lock static substrate. Beyond the lock the substrate is scale-free and static (Sec. 9.1); what observers read as cosmic expansion is the coherent ruler shift of Sec. 9.4, BASE-noise drift of m_P accumulating along photon paths. The picture eliminates 4D metric stretching and the spatial-expansion paradox while leaving every observed ΛCDM phenomenology unchanged in form.

App. E carries out the systematic audit, sorting ΛCDM elements into those HFT retains, those it rebrands as substrate-ontology readouts, and those it collapses into structural consequences of the substrate rule.

10 Observational Predictions

HFT's falsifiable predictions fall into two groups: a single framework-level consistency test, and specific observational targets each isolating one aspect of the picture.

Framework consistency.

1. **Universal BASE noise consistency at leading order.** All BASE noise channels share one substrate spectrum (zero-point ω^3 rising to the substrate UV cutoff, plus a thermal Planck peak at $\nu \approx 160$ GHz); they differ only in which window each instrument samples and how strongly that channel couples to matter. A single universal amplitude $\varepsilon = \ln(2\pi)/N_{\text{skel}} \approx 1.34\%$ should reproduce every currently probed channel to picture precision: H^1 zero-point via QED loops (Casimir, Lamb shift, $g-2$); H^1 thermal via the CMB; $H^1 \leftrightarrow H^2$ via the α–θ_W running, propagating through θ_v into the v-anchored SM mass cascade (App. D); and the anchor-drift channel via the $(1+z)$ family (Sec. 9.4). A substantial leading-order cross-channel mismatch, beyond what higher-order framework corrections can absorb, would falsify HFT. Bands not yet probed — notably the H^0 thermal peak at 160 GHz — are addressed in the specific targets below.

Specific observational targets.

2. **SGWB H^0 peak coincident with the CMB.** BASE noise's H^0 and H^1 channels share one substrate thermal source (Sec. 9.3), so the SGWB is predicted to peak at the CMB B_ν peak $\nu_{\text{peak}} \approx 160$ GHz with amplitude $\Omega_{\text{BASE},H^0} \approx \Omega_\gamma \approx 5 \times 10^{-5}$. The peak coincidence has no ΛCDM counterpart and is itself the cleanest evidence for the common origin. Current bounds (PTA [28], LIGO [29]) lie far below this frequency; the ~ 160 GHz target awaits future high-frequency GW detectors.
3. **Majorana neutrino mass spectrum and CvB temperature.** Chiral locking leaves ν_L as the only stable Twist mode, implying Majorana mass with the lightest eigenstate $m_{\nu_1} \approx 0.8$ meV (App. A.3). The full three-eigenstate spectrum (Table 2) lies in the sensitivity range of next-generation neutrinoless double beta decay experiments. In the threshold-crossing reading of quantum fluctuation (Sec. 9.3), the vacuum ν_1 population sits in thermal equilibrium with the substrate floor at $T_{\text{CMB}} = 2.725$ K rather than at the ΛCDM-cooled $T_\nu \approx 1.95$ K — a $\sim 40\%$ temperature excess testable by next-generation CvB direct-detection experiments (e.g., PTOLEMY).
4. **No inflationary primordial B-mode.** The pre-lock expansion phase (App. E) is not a slow-roll inflaton field and carries no tensor-mode quantization comparable to standard inflation. Detection of a primordial B-mode signal at the amplitude predicted by single-field inflation [30], with no HFT-compatible reinterpretation, would falsify the pre-lock picture.
5. **No independent QCD axion.** Strong CP conservation is structural, not dynamical (App. B), so HFT predicts no independent Peccei–Quinn axion. Detection of an axion-like particle that cannot be reinterpreted as a collective excitation of HFT's fiber–base structure would falsify this account.

11 Complete Parameter Table

All masses follow from the dimensional anchor $m_P = 1.221 \times 10^{19}$ GeV (Planck mass) via picture-internal dimensionless ratios. The two coupling-constant entries are reported with BASE-noise dressing applied at the closed-form level (App. D). For the remaining cascade entries, $\varepsilon = \ln(2\pi)/N_{\text{skel}} \approx 1.34\%$ enters through multiple compound channels — e.g., M_W inherits both the v-cascade dressing through S_{bounce} and the θ_W dressing through $\cos\theta_W = 7/8$; the detailed compound breakdown is not expanded here. The reported values are leading-order picture predictions, with sub-percent residuals at the picture-internal precision floor.

Parameter	HFT formula	HFT value	Experimental	Error
Coupling constants (§4)				
$\alpha^{-1}(0)$	$137 + \frac{5}{137}(1-\varepsilon)$	137.036007	137.035999	0.06 ppm
$\sin^2\theta_W(M_Z)$	$\frac{30}{128}(1-\varepsilon)$	0.23123	$0.23121(4)$	91 ppm
Electroweak sector (§5; §A.1)				
v	$m_P \exp(-(24\pi + 3/2)/2)$	244.7 GeV	246.22 GeV	0.62%
M_H	$v\sqrt{N_{\text{empty}}/S_{\text{bounce}}}$	124.80 GeV	125.10 GeV	0.24%
M_Z	$M_H\sqrt{8/15}$	91.14 GeV	91.19 GeV	0.05%
M_W	$M_Z \times 7/8$ ($\cos\theta_W = 7/8$)	79.75 GeV	80.38 GeV	0.77%
M_t	$v/\sqrt{2}$ ($y_t = 1$ saturation)	173.0 GeV	172.7 GeV	0.18%
Strong sector (§B)				
Λ_{QCD}	$v/(N_{\text{skel}} \cdot N_{\text{coh}}) = v/548$	446.5 MeV	420–470 MeV†	in range
M_p	$\Lambda_{\text{QCD}} \times 21/10$	937.7 MeV	938.27 MeV	0.06%
M_{Λ_b}	$M_p \times 6$	5626 MeV	5619.6 MeV	0.11%
Lepton masses (§A.2)				
m_e	$\Lambda_{\text{QCD}}/(N_v R^2)$	0.5075 MeV	0.5110 MeV	0.69%
m_μ	$m_e \cdot N_{\text{marked}} \cdot R$	104.29 MeV	105.66 MeV	1.30%
m_τ	$m_e \cdot N_{\text{marked}} \cdot R^2$	1786 MeV	1776.86 MeV	0.51%
Neutrino masses (§A.3)				
m_{ν_1}	$M_H e^{-N_{\text{soft}}/3}$	0.811 meV	—	—
m_{ν_2}	$m_{\nu_1} \times 5 \times 137/64$	8.68 meV	≈ 8.6 meV	0.9%
m_{ν_3}	$m_{\nu_2} \times S_E/9$	49.55 meV	≈ 50 meV	0.9%
$\sum m_\nu$	$m_{\nu_1} + m_{\nu_2} + m_{\nu_3}$	59.0 meV	< 70 meV	consistent

Table 2: HFT predictions for Standard Model free parameters and dark sector. Experimental values from the Particle Data Group [17] unless otherwise indicated. † Non-perturbative confinement scale $\sqrt{\sigma}$ extracted from lattice QCD and meson Regge trajectories [18].

Appendices

The appendices that follow contain the mathematical derivations and precise calculations underlying every entry in Table 2. Readers can use the table as a self-contained navigation index: each row points to the section in which its formula is derived (left column) and to the appendix that performs the precise geometric computation. The appendices are structured for verification rather than narrative reading; the main text is self-contained for the conceptual framework.

A The Mass Scale Ladder

A.1 Geometric Derivation of the Higgs VEV v from m_P

In HFT the Higgs VEV v [7, 6] is the chirality-locked EW vacuum $\mathcal{V}$'s amplitude order parameter — the residual tension in the substrate after the locking event closes the vertex geometry. We anchor the numerical chain on the Planck mass m_P; the dimensionless ratio v/m_P admits two complementary picture-internal readings that converge on the same total bounce action [9] $S_{\text{bounce}} \equiv 24\pi+3/2 \approx 76.9$, with $v = m_P \exp(-S_{\text{bounce}}/2)$.

Vertex-locked ground state. At each vertex the substrate carries a *vertical elastic rod* — the vertex axial S^1 fiber threaded perpendicular to the local S^2 base — under an ambient tension scale that, prior to the chirality lock, is the Planck scale m_P. The six incoming strands re-pair (Sec. 3) into three composite fibers, each of which winds once around the rod (2π on the axial S^1) before exiting. This full S^1 winding is what couples each outgoing fiber rigidly to the rod, and the three winding fibers together form a C_3-symmetric topological grip. The grip is held closed by the geometric pre-stress $\theta_v = 1/32$ (Sec. 4.4), a structural horizontal rotation the chirality lock imposes between the S^2 base and the S^1_{Hopf} fiber at every clamping contact. The Higgs VEV v is the rod's equilibrium amplitude after the grip has dissipated its locking work — the residual tension at which rod and grip balance; the same grip then governs the rod's longitudinal oscillation, setting M_H in the next subsection.

Reading 1 — Topological state count (bounce action). The bounce traverses the post-lock mesh, picking up one elementary action quantum per chirality-marked vertex configuration ($N_{\text{marked}} = N_v \cdot N_{\text{coh}} = 12$, Sec. 6). Each marked configuration contributes two pieces:

(i) *Topological winding* 2π: a full S^1 axial-fiber circumnavigation, the elementary instanton winding around the locked vertex axial direction.

(ii) *Vacuum-state selection* $N_{\text{coh}}\,\theta_v$: per chirality-marked vertex, the vacuum sub-state fraction $\theta_v = 1/32$ (Sec. 4.4) picks one vacuum vertex out of 32 pre-lock states; summed over $N_{\text{coh}} = 4$ cohomology channels, $= 1/8$.

Aggregating over the N_{marked} chirality-marked configurations:

$$-\ln\left(\frac{v^2}{m_P^2}\right) \;=\; N_{\text{marked}} \cdot (2\pi + N_{\text{coh}}\,\theta_v) \;=\; N_v N_{\text{coh}} \cdot 2\pi \;+\; N_v N_{\text{coh}}^2 \cdot \theta_v \;=\; 24\pi + \tfrac{3}{2}. \tag{21}$$

The radial-amplitude form $\ln(v^2/m_P^2)$ (squared because v is the amplitude order parameter) converts the topological state count into the radial bounce action S_{bounce}.

Reading 2 — Elastic-rod locking work. The same total $24\pi + 3/2$ is the mechanical work the C_3 grip performs to close itself around the central elastic rod:

(i) *Winding travel* 24π. Each of the $N_v = 3$ composite fibers must execute a full 2π winding independently in each of the $N_{\text{coh}} = 4$ orthogonal cohomology channels — otherwise the knot would slip out through an unlocked dimension. Total topological winding travel: $3 \times 4 \times 2\pi = 24\pi$.

(ii) *Friction work* $3/2$. The four cohomology channels each carry the $\theta_v = 1/32$ prestress, giving the grip a geometric friction coefficient $4 \times \frac{1}{32} = \frac{1}{8}$ per winding. Over the 12 windings the total friction work is $12 \times \frac{1}{8} = \frac{3}{2}$.

Sum: locking work $= 24\pi + 3/2$, identical to the bounce action of Reading 1. The two readings are dual descriptions of the same vertex-locked ground state: Reading 1 counts the topological states traversed during lock-in; Reading 2 measures the mechanical work expended to traverse them.

Convergent result.

$$\boxed{v \;=\; m_P \cdot \exp\!\left(-\frac{24\pi + 3/2}{2}\right) \;=\; m_P \cdot e^{-12\pi - 3/4} \;\approx\; 244.7 \text{ GeV},} \tag{22}$$

matching the observed $v_{\text{obs}} \approx 246.22$ GeV to 0.62% (or 0.017% in $\ln(v^2/m_P^2)$ space). The Fermi constant follows trivially from the SM identity $G_F = 1/(\sqrt{2}\,v^2)$ and is not an independent HFT prediction.

A.2 Lepton and Boson Masses

- **Electron Scale**: The electron mass is anchored on Λ_{QCD} via the lepton mass ladder (cf. Charged Lepton Hierarchy below). The τ lepton saturates the elastic ground at the full 9-class vertex engagement, giving $m_\tau \approx N_{\text{coh}} \cdot \Lambda_{\text{QCD}}$; the electron sits at the base of the hierarchy, $m_e = m_\tau/(N_{\text{marked}} \cdot R^2)$ with $N_{\text{marked}} = 12$ and $R = N_{\text{skel}}/2^{\text{dim(vertex)}} = 137/8$. Combining:

$$\boxed{m_e \;=\; \frac{\Lambda_{\text{QCD}}}{N_v \cdot R^2} \;=\; \frac{\Lambda_{\text{QCD}} \cdot 2^{2\,\text{dim(vertex)}}}{N_v \cdot N_{\text{skel}}^2}} \tag{23}$$

  Numerically, with $\Lambda_{\text{QCD}} = v/548 \approx 446.5$ MeV from the predicted v above, $m_e = 446.5/(3 \times 17.125^2) \approx 0.507$ MeV, matching the observed 0.511 MeV to 0.7%. The corresponding Planck-ratio is $m_e/m_P \approx 4.16 \times 10^{-23}$ vs. observed 4.19×10^{-23} (0.8%). The exponential framing $m_e/m_P = e^{-(S_E+\cdots)}$ in earlier formulations is a numerical coincidence ($\ln(m_P/m_e) \approx N_{\text{skel}} \cdot 3/8$); the picture-internal derivation is the Λ_{QCD}-anchored ladder above, independent of S_E.

- **Charged Lepton Hierarchy**: Inter-generational mass ratios are set by the same geometric dilution that governs Λ_{QCD}: projecting the global $N_{\text{skeleton}} = 137$ lattice onto sub-spaces of $\mathcal{Q}$ of increasing dimension.
  - $m_\mu/m_e = N_{\text{skeleton}} \times \frac{3}{2} = 137 \times \frac{3}{2} = 205.5$. The factor $3/2$ is the ratio of S^2-sector degrees of freedom (3: anchor x, y plus orientation θ) to S^1-sector degrees of freedom (2: tension ρ and phase ψ). The muon excitation probes the full base-space projection; the electron is confined to the S^1 fiber alone.
  - $m_\tau/m_\mu = N_{\text{skeleton}}/\sum Q_f^2 = 137/8 = 17.125$. At the third generation the winding energy couples to all $\sum Q_f^2 = 8$ anchored fermion classes; the factor $1/8$ is the per-class share of the global lattice budget.

- **Higgs, W, Z Bosons**: The Higgs amplitude mode H is the longitudinal ρ-mode oscillation of v around the chirality-locked vacuum $\mathcal{V}$. Of the $2^{\text{dim(vertex)}} \cdot N_{\text{coh}} = 32$ pre-lock configurations per vertex (Sec. 4.4), $N_{\text{marked}} = 12$ are chirality-loaded (Sec. 6); the remaining $N_{\text{empty}} = 32 - 12 = 20$ chirality-empty configurations furnish the phase space through which the ρ-mode amplitude sweeps. Normalising by the bounce action $S_{\text{bounce}} = 24\pi + 3/2$ that establishes v (Sec. A.1):

$$\boxed{M_H^2 \;=\; v^2 \cdot \frac{N_{\text{empty}}}{S_{\text{bounce}}} \;=\; v^2 \cdot \frac{20}{24\pi + 3/2}, \qquad M_H \approx 124.8 \text{ GeV.}} \tag{24}$$

  Picture: M_H^2/v^2 is the ratio of ρ-mode oscillation phase-space dimension to the bounce-action depth establishing v; equivalently the Higgs self-coupling $\lambda_h = M_H^2/(2v^2) = N_{\text{empty}}/(2S_{\text{bounce}}) \approx 0.130$, matching $\lambda_h^{\text{obs}} \approx 0.130$ at $\mu = M_H$.

  The neutral and charged weak gauge bosons follow as the orthogonal mode partners of H:
  - $M_Z = M_H\sqrt{2\sum Q_f^2/N_{\text{weak}}} = M_H\sqrt{8/15} \approx 91.14$ GeV. The Higgs–Z ratio $M_H^2/M_Z^2 = 15/8 = N_{\text{weak}}/(2\sum Q_f^2)$ is the ratio of neutral-weak channels ($N_{\text{weak}}/2 = 15$) to fermion-strand anchors ($\sum Q_f^2 = 8$, equivalently the 8 generators of $\mathfrak{su}(3)$ from the hexagonal

mesh). Picture: H couples to the full neutral-weak channel count, Z to a fermion-strand-rescaled subset.

– $M_W = M_Z \cos\theta_W = M_Z \times 7/8 \approx 79.75$ GeV. The exact rational $\cos\theta_W = 7/8$ follows from $\sin^2\theta_W = 30/128 = 15/64$ (Sec. 4.2), so $\cos^2\theta_W = 49/64 = (7/8)^2$ with no free parameters.

- **Top Quark**: A quark of winding number n and framing twist Tw couples to the Higgs vacuum with Yukawa amplitude y_q, giving SM-standard $m_q = y_q v/\sqrt{2}$. In HFT y_q is the overlap between the sub-fiber strand Tw winding and the chirality-locked vacuum $\mathcal{V}$ tension structure. The top quark is uniquely *dual-ceiling saturated*: $n = 3$ saturates the generation ceiling fixed by the $k = 3$ Chern–Simons truncation (Sec. 3), and $Tw = +1$ saturates the framing ceiling. Both quantum numbers at their maximum drive the overlap to the unitary bound:

$$y_t = 1 \quad \text{(dual-ceiling saturation)}, \tag{25}$$

giving

$$M_t = \frac{v}{\sqrt{2}} = \frac{244.7 \text{ GeV}}{\sqrt{2}} \approx 173.0 \text{ GeV}, \tag{26}$$

matching the observed 172.69 ± 0.30 GeV to 0.2%. No other quark sits at dual saturation: (b, c, s, u, d) all have at least one quantum number below ceiling, suppressing $y_q < 1$.

A.3 Majorana Neutrinos

Three Tw mode types of the locked B_2 doubled strand. On the post-lock mesh, each edge carries a B_2 doubled-strand framed thread (Sec. 3). The doubled-helix topology admits exactly *three* intrinsic Tw propagation modes, enumerated by the topological relation between the Tw carrier and the locked structure:

Mass eigenstate	Mode type	Topological scope
ν_1	Sub-strand parity Tw	*within* the B_2 pair (one sub-strand vs the other)
ν_2	B_2 collective Tw	*along* the edge direction (whole pair Tw)
ν_3	Vertex-coupled Tw	*across* vertex endpoints (Tw bridging two edges)

The triple within/along/across exhausts the Tw topological relations possible on a doubled helix on an edge — a *geometric exhaustion* fixing the neutrino generation count at three, independent of the trivalent vertex structure or anyon labelling. All three mode types become well-defined as stable propagating modes only post-lock: *the existence of three neutrinos directly hinges on the chirality lock*.

These three are the mass eigenstates $|\nu_n\rangle$. They are Majorana because ν_L self-pairing is enforced by chiral locking: pure-Twist excitations carry no Writhe and do not couple to S^2-base perturbations, so their reference scale is the Higgs amplitude mode M_H in the same fiber sector.

Anchor mass m_{ν_1}: sub-strand parity localization. Of the $N_{\text{Nyquist}} = 128$ fiber states, $N_{\text{weak}} = 30$ are chiral-locked into $\mathcal{V}$. The remaining

$$N_{\text{soft}} = N_{\text{Nyquist}} - N_{\text{weak}} = 98$$

states form the unlocked (Tw-mode) sector. The ν_1 within-pair parity mode localizes onto one $\mathbb{Z}_3$ class of this sector, incurring localization entropy $N_{\text{soft}}/3 = 98/3$:

$$m_{\nu_1} = M_H \exp\left(-\frac{N_{\text{soft}}}{3}\right) = M_H\, e^{-98/3} \approx 0.811 \text{ meV}. \tag{27}$$

Equivalently, using the HFT Higgs self-coupling $\lambda_h = 4\pi/N_{\text{soft}} = 2\pi/49$, $m_{\nu_1} = M_H \exp(-4\pi/(3\lambda_h))$.

Mode-type transitions: topological scale upgrades. The inter-generational mass ratios reflect transitions in topological scope, not phenomenological fits:

- **Within $\to$ along ($\nu_1 \to \nu_2$, same edge)**: the Tw carrier expands from sub-strand parity to the full B_2 collective mode. The skeleton spread is $N_{\text{skel}}/2^{\dim(\text{vertex})} = 137/8$, weighted by the chirality lock complement $1 - \sin^2\theta_W|_{\text{lock}} = 5/8$ (fiber-dominant sector available to pure-Twist objects):

$$m_{\nu_2} = m_{\nu_1} \times \frac{137}{8} \times \frac{5}{8} = m_{\nu_1} \times \frac{5 \times 137}{64} \approx 8.68 \text{ meV}. \tag{28}$$

- **Along $\to$ across ($\nu_2 \to \nu_3$, edge $\to$ vertex)**: the Tw carrier crosses vertex endpoints, sensing the $N_v^2 = 9$ stable vertex classes (Sec. 4.3 defect-class enumeration). The scaling is set by the action budget per class:

$$m_{\nu_3} = m_{\nu_2} \times \frac{S_E}{N_v^2} = m_{\nu_2} \times \frac{51.375}{9} \approx 49.55 \text{ meV}. \tag{29}$$

The two splittings give $\Delta m^2_{21} \ll \Delta m^2_{31}$ structurally: within$\to$along stays on the same edge (small topological scale gap), whereas along$\to$across upgrades to vertex coupling (large gap). The observed mass-splitting hierarchy is therefore a forward consequence of the three-mode picture.

Flavor eigenstates and PMNS mixing. The flavor eigenstates $|\nu_\alpha\rangle$ ($\alpha = e, \mu, \tau$) are the vertex Wr$\to$Tw coupling outcomes at **vertices**, while the mass eigenstates above are B_2 Tw mode types on **edges**. The PMNS matrix $U^{\text{PMNS}}_{\alpha n} = \langle \nu_\alpha | \nu_n \rangle$ is therefore a cross-category overlap, generically large — the picture-internal answer to why PMNS mixing is large compared to CKM. A first-principles derivation of the PMNS entries from the vertex Wr-Tw coupling overlap is deferred to follow-up work.

B Strong Interactions and Flavor Dynamics

B.1 Color Confinement as Nyquist Saturation

Confinement occurs when the probe can no longer resolve individual S^1 fiber degrees of freedom — a resolution collapse at the single-fiber Nyquist limit. The confinement scale is obtained by projecting the EW vacuum amplitude v down to the single-fiber Nyquist resolution via the post-lock per-cell capacity $N_{\text{skel}} \cdot N_{\text{coh}}$:

$$\Lambda_{\text{QCD}} = \frac{v}{N_{\text{skel}} \cdot N_{\text{coh}}} = \frac{v}{548} \approx 446.5 \text{ MeV}, \tag{30}$$

where $N_{\text{skel}} = 137$ is the post-lock per-cell skeleton (Sec. 4.3) and $N_{\text{coh}} = 4$ is the number of de Rham cohomology channels (Sec. 4.1). The result lies within the lattice QCD string-tension range $\sqrt{\sigma} \approx 420$–$470$ MeV [18]. Equivalently, in terms of M_Z,

$$\Lambda_{\text{QCD}} = \frac{M_Z}{N_{\text{skel}} \cdot 3/2} = \frac{M_Z}{205.5} \approx 443.5 \text{ MeV}, \tag{31}$$

where $3/2 = \dim(\text{vertex})/\dim(\text{edge})$ is the vertex-to-edge DOF ratio. The two forms agree picture-internally via the identity $v = (8/7) \cdot M_Z \cdot \sqrt{30/\pi}$; both routes project the global skeleton onto a single-fiber local excitation. The same dilution factor $N_{\text{skel}} \cdot 3/2$ also yields m_μ/m_e (App. A.2), reflecting unified geometric origin.

B.2 The Strong CP Problem: Geometric Decoupling of Colour and Parity

The QCD θ-term, $\theta_{\text{QCD}} F\tilde{F}$, would generate observable CP violation in the strong sector unless $|\theta_{\text{QCD}}| \lesssim 10^{-10}$. In the SM this smallness is a long-standing puzzle, conventionally addressed by the Peccei–Quinn mechanism [10] with a hypothetical axion field. In HFT the smallness is structural rather than dynamical: colour and parity live on disjoint submanifolds of the Hopf bundle, and the cross-coupling that a non-zero θ_{QCD} would require is geometrically forbidden.

Colour and parity live on different submanifolds. The lattice gauge sector of the trivalent mesh has algebra $SU(N_v) \cong \mathfrak{su}(3)$ (Sec. 3): each Hopf vertex carries $N_v = 3$ trivalent-channel amplitudes whose unitary symmetry decomposes as $SU(N_v)\times U(1)$ via $U(N_v)/U(1)$. The fiber is one-dimensional and topologically distinct from the S^2 base.

By contrast, parity (P) and time-reversal (T) are properties of the S^2 base: P corresponds to the internal orientation selected by chiral locking at EWSB (Sec. 7), and T to the external orientation set by the irreversible $Tw \to Wr$ flow. Both reside in the S^2 base, not on the S^1 fiber.

The QCD θ-term measures the cross-coupling between colour topology (an S^1-fiber quantity) and parity (an S^2-base quantity). In HFT this cross-coupling has no topological invariant available to source it: S^1-fiber and S^2-base operations act on disjoint submanifolds of the Hopf bundle, and their tensor product carries no Pontryagin-density-like cross-term. Strong CP conservation is therefore a structural consequence of the fiber-base geometry, not a fine-tuning of an independent parameter.

No axion required as a new fundamental field. The Peccei–Quinn mechanism introduces a global $U(1)_{\text{PQ}}$ symmetry whose spontaneous breaking yields the axion as a Goldstone mode. HFT does not require this addition: the geometric decoupling above enforces $\theta_{\text{QCD}} = 0$ at the framework level, and any small misalignment generated by vacuum fluctuations is driven to zero by the entropy gradient (since $\theta \neq 0$ would require a fiber-base correlation that is energetically and entropically disfavoured). The PQ Lagrangian remains valid as an effective-field-theory description of this relaxation, with the axion interpreted as a collective excitation of the residual fiber-base coupling sector rather than as a new fundamental scalar.

B.3 The $21/10$ Ratio from Phase-Space Dimension Counting

In HFT mass is the elastic energy stored in a configuration of the mesh, and in the continuum limit ratios of mass scales reduce to ratios of accessible phase-space dimensions (Sec. 5). The ratio M_p/Λ_{QCD} admits a direct phase-space count from the picture-internal 5D configuration manifold $\mathcal{Q} = \mathbb{R}^+ \times S^2 \times T \times S^1_{\text{Hopf}}$.

Denominator: free-fiber phase space. Λ_{QCD} sets the scale of an unconfined single fiber. Its phase space is the cotangent bundle $T^*\mathcal{Q}$ of dimension

$$\dim T^*\mathcal{Q} \;=\; 2\,\dim\mathcal{Q} \;=\; 10. \tag{32}$$

Numerator: trivalent confined bound state. The proton is a confined three-fiber bound state at a single trivalent vertex. Counting its phase-space dimension geometrically:

- three independent fibers give a raw $3 \times \dim T^*\mathcal{Q} = 30$;
- co-location at one vertex — forcing the three configuration points to coincide ($q_1 = q_2 = q_3$) — imposes $2 \times \dim\mathcal{Q} = 10$ constraints, leaving $30 - 10 = 20$ continuous dimensions;
- the B_3-braid knot of the three fibers at the vertex (an H^3 topological invariant) contributes one additional collective internal degree of freedom not captured by the continuous coordinates.

Together,

$$\dim(\text{proton phase space}) \;=\; 20 + 1 \;=\; 21. \tag{33}$$

Ratio. The mass ratio is therefore

$$\frac{M_p}{\Lambda_{\text{QCD}}} \;=\; \frac{21}{10} \;=\; 2.1, \tag{34}$$

giving $M_p \approx 936$ MeV against 938.27 MeV (0.2%).

B.4 The Λ_b/p Ratio: Heavy-Flavour Knot Multiplicity

The Λ_b baryon (udb) is the lightest stable bottom-flavoured baryon. In HFT it is the minimum-energy 3-quark knot in which one valence strand carries a generation-3 winding ($n = 3$); the proton is the same topological object with all three strands at $n = 1$. The mass ratio factorises as a product of two independent combinatorial multiplicities:

$$\frac{M_{\Lambda_b}}{M_p} = N_{\text{gen}} \times N_{\text{strand}} = 3 \times 2 = 6. \tag{35}$$

Here $N_{\text{gen}} = 3$ is the number of winding generations the heavy strand can occupy ($n = 1, 2, 3$ from the structurally fixed $k = 3$ Chern-Simons truncation, Sec. 3; matching the SM's three observed generations), and $N_{\text{strand}} = 2$ is the residual SU(2) isospin doublet of the spectator ud pair after the heavy strand is pinned. With the HFT-predicted $M_p = 937.7$ MeV, this gives $M_{\Lambda_b} \approx 5626$ MeV, within 0.11% of the observed 5619.6 MeV.

C QED from Fiber Dynamics: A Worked Example

The IR continuum limit of HFT mesh dynamics reproduces standard quantum field theory. This appendix demonstrates the reduction explicitly for the electromagnetic sector: starting from the tension field on the Hopf bundle and applying $\mathcal{V}$-coarse-graining (averaging over the locked-mesh ground state), we recover the Maxwell Lagrangian, the Feynman-gauge photon propagator, the electron–photon vertex, and the tree-level Coulomb potential. The procedure generalises to other QED amplitudes (Compton, Bhabha, $e^+e^- \to \mu^+\mu^-$) and to the weak and colour sectors via the parallel mesh structures of Sec. 4; we treat the simplest case here and indicate the future-work roadmap at the end.

C.1 From fiber elasticity to the Maxwell Lagrangian

The $U(1)$ gauge field A_μ in HFT is the linearised fiber-phase fluctuation on the Hopf bundle. The mesh elastic energy density of the S^1 fiber, after IR continuum dimensional reduction and with the kinetic-term prefactor fixed by the HFT-derived $\alpha^{-1} = 137.036$ (Sec. 4), takes the form

$$\mathcal{L}_{\text{fiber,IR}} \;=\; -\tfrac{1}{4}\, F_{\mu\nu}F^{\mu\nu} + \mathcal{L}_{\text{vac}}, \tag{36}$$

where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the standard field strength and $\mathcal{L}_{\text{vac}}$ contains vacuum noise-driven corrections that average out at scales above the mesh-stochastic autocorrelation length ξ_{vac}. The first term is the Maxwell Lagrangian density, recovered from mesh elasticity in the IR continuum limit.

C.2 Photon propagator from the $\mathcal{V}$-coarse-grained Green's function

The Green's function of the Maxwell kinetic operator gives the Feynman-gauge photon propagator

$$D^{\text{HFT}}_{\mu\nu}(k) \;=\; \frac{-i\, g_{\mu\nu}}{k^2 + i\epsilon}, \tag{37}$$

with the $+i\epsilon$ prescription corresponding to the small-amplitude expansion of the vacuum autocorrelation in the causal (retarded) sector, selecting the future-directed Green's function.

C.3 Knot–fiber coupling: the QED vertex

A charged knot at worldline $x(\tau)$ couples to the fiber field by integration of the fiber phase along its trajectory. After upgrading the worldline to a Dirac spinor field via the chirality structure of Sec. 7, the coupling action is

$$S_{\text{int}} \;=\; e \int d^4x\, \bar{\psi}\, \gamma^\mu \psi\, A_\mu, \tag{38}$$

with $e^2 = 4\pi\alpha$ and $\alpha^{-1} = 137.036$ from HFT. The Feynman rule for the electron–photon vertex is therefore $-ie\gamma^\mu$.

C.4 Coulomb potential as a tree-level prediction

Combining the propagator and the vertex at tree level for two static charges Q_1, Q_2 in units of e:

$$V(r) \;=\; Q_1 Q_2\, e^2 \int \frac{d^3k}{(2\pi)^3} \frac{e^{i\vec{k}\cdot\vec{r}}}{|\vec{k}|^2} \;=\; \frac{Q_1 Q_2\, e^2}{4\pi\, r} \;=\; \frac{Q_1 Q_2\, \alpha}{r}. \tag{39}$$

This is the Coulomb potential, exact at tree level. With the HFT-derived $\alpha^{-1} = 137.036$, the prediction matches experimental measurement to all currently available precision.

D The α–θ_W Geometric Coupling Identity

The Weinberg angle and the fine-structure constant collapse onto a single structural identity in HFT. The chirality lock's $\mathbb{Z}_2$ quantum displaces exactly two Nyquist slots from the weak channel at every scale — independently of how the total Nyquist budget runs — so $\alpha^{-1}(Q)$ determines $\sin^2\theta_W(Q)$ directly.

Derivation. The discrete view of the Weinberg angle (Sec. 4) gives the weak-channel slot count as the cohomology equipartition baseline minus the Călugăreanu $\mathbb{Z}_2$ displacement,

$$N_w(Q) \;=\; \frac{\alpha^{-1}(Q)}{4} - 2, \tag{40}$$

where the baseline $\alpha^{-1}(Q)/4$ is the symmetric allocation across the four de Rham classes H^0, H^1, H^2, H^3 of the per-cell phase space, and the integer 2 is the minimal Lk-conserving twist quantum. Taking $N_w(Q)/\alpha^{-1}(Q)$ gives

$$\boxed{\sin^2\theta_W(Q) \;=\; \frac{1}{4} - \frac{2}{\alpha^{-1}(Q)}.} \tag{41}$$

The $\mathbb{Z}_2$ displacement of 2 slots is the single primitive fixing both the bare baseline $1/4$ and the lock-induced offset $-2/\alpha^{-1}$; the two SM couplings are not independent observables, but two faces of one structural commit.

Scale-anchor verification. At anchors spanning low-Q to the Z-pole, with $\Delta \equiv \sin^2\theta_W^{\text{obs}} -$ Identity:

Q	$\alpha^{-1}(Q)$	Identity (41)	Observed $\sin^2\theta_W$	Δ
0	137.036	0.23541	0.23857	$+0.00316$
~ 1 GeV	~ 135	0.23519	~ 0.2378	$+0.00261$
~ 5 GeV	~ 133	0.23496	~ 0.2365	$+0.00154$
~ 10 GeV	~ 131	0.23473	~ 0.2353	$+0.00057$
~ 20 GeV	~ 130	0.23462	~ 0.2344	-0.00022
~ 30 GeV	~ 129	0.23450	~ 0.2333	-0.00120
~ 50 GeV	~ 128.3	0.23441	~ 0.2324	-0.00201
M_Z	127.92	0.23436	0.23121	-0.00315

The identity holds at the few-per-mille level across all anchors. The residual Δ flips sign antisymmetrically between low-Q ($+0.00316$ at $Q = 0$) and the Z-pole (-0.00315), crossing zero near $Q \sim 15$–20 GeV.

Closed-form predictions at the endpoints. At the endpoints of the running range the BASE noise dressing yields the picture-complete closed forms

$$\alpha^{-1}(0) \;=\; N_{\text{skeleton}} + \frac{2^{\dim(\text{vertex})} - \dim(\text{postlock})}{N_{\text{skeleton}}}\,(1-\varepsilon) \;=\; 137 + \frac{5}{137}(1-\varepsilon) \;\approx\; 137.0360, \tag{42}$$

$$\sin^2\theta_W(M_Z) \;=\; \frac{30}{128}\,(1-\varepsilon) \;\approx\; 0.23123, \tag{43}$$

matching observation at ~ 0.06 ppm and ~ 91 ppm respectively (the latter well inside the experimental 1σ of ~ 170 ppm). The same universal amplitude $\varepsilon = \ln(2\pi)/N_{\text{skeleton}}$ enters both, dressing what is structurally a sub-leading correction — the unmarked vertex-configuration count $2^{\dim(\text{vertex})} - \dim(\text{postlock}) = 5$ normalized by N_{skeleton} for α^{-1}, the $30/128$ lock-scale ratio at the IR running endpoint for $\sin^2\theta_W$ — with no additional parameter.

The ~ 0.06 ppm α^{-1} residual sits at the picture-internal precision floor and remains above the sub-ppb experimental uncertainty of α^{-1}. A sub-leading ε^2 structure consistent with closing this gap is not yet derived; the closed forms above are committed at leading order in ε, deferring higher-order structure to follow-up.

E HFT Cosmology vs ΛCDM: A Bridge Hypothesis

The cosmological picture in Sec. 9.4 (static substrate; cosmic-expansion phenomenology as coherent ruler shift) is at present a *hypothesis* sketched on top of HFT's verified SM and gravity machinery. This appendix audits its relation to ΛCDM by sorting the latter's content into three classes: what HFT retains unchanged, what it rebrands ontologically, and what it eliminates outright. The audit is offered as a structural map of the proposal, not a finished derivation; the quantitative bridges in each row are deferred to a follow-up paper.

Retained. HFT inherits the macroscopic predictions of ΛCDM unchanged:

- BBN light-element abundances;
- recombination physics and the CMB origin as fossil light from last scattering;
- the acoustic-peak structure (sub-percent fit);
- the CMB blackbody spectrum (FIRAS $\Delta T/T < 10^{-5}$ [24]);
- the local distance ladder regime ($z < 0.1$, Cepheid + SN Ia);
- the cooling history and broad-strokes structure formation.

HFT reinterprets photons as mesh H^1 thermal modes (Table 1), but the observational predictions in each item above are unchanged.

Bridged. HFT carries its own pre-EWSB ansatz for the early universe — a yarn-ball initial state expanding and cooling under the substrate rule into the current S^2 scale, with a discrete chirality-lock transition ending the phase. This is an ansatz layer of comparable epistemic standing to ΛCDM's inflation [23]; both sit upstream of any direct observation. The bridge below maps the two ansatz layers (plus the post-lock observation regime) onto each other:

ΛCDM element	HFT reading	Observational difference
Big Bang singularity	Yarn-ball initial state: finite, indistinguishable, no singularity	none — neither directly observable
Inflaton field + $V(\phi)$	Pre-EWSB yarn-ball expansion and cooling under the substrate rule	none at this stage — both ansatz
Inflaton quantum fluctuations seeding structure	Non-uniform cooling during pre-EWSB expansion seeding structure	none — same seed function
End of inflation / reheating	Chirality lock (EWSB) discrete phase transition; latent-heat release	none — same hot start
Cosmic expansion (post-lock)	Coherent ruler shift; substrate static (Sec. 9.4)	none — readouts agree
Photon as field	Mesh H^1 thermal mode	none — same detection
Plasma cooling	Mesh thermal evolution	none — same thermodynamics
Redshift / time dilation / Tolman / $T(z)$	BASE noise multi-readout (Sec. 9.3)	none — Etherington identity automatic

At and above the lock event both ansätze sit upstream of direct observation, and HFT does not impose a duration mapping between them; below the lock, HFT eliminates 4D metric stretching and the spatial-expansion paradox without disturbing the observed phenomenology.

Eliminated. A second class of ΛCDM fine-tunings has no HFT counterpart at all — they collapse into structural consequences of the substrate rule rather than tunable parameters of an ansatz:

ΛCDM fine-tuning	HFT structural consequence
Horizon-problem tuning	Yarn-ball initial state is natively indistinguishable; no causally separated regions to reconcile
Flatness-problem tuning	Substrate topology is fixed by the lock event; no curvature parameter to tune
60 e-foldings tuning	Lock is a discrete topological transition; no e-folding parameter
Slow-roll potential shape tuning	No scalar field; substrate rule replaces $V(\phi)$
Cosmological constant (10^{120} tuning)	θ_v vertex pre-stress structurally locked; does not gravitate (Sec. 8.1)
Inflation B-mode primordial GW	Not predicted (no inflaton tensor modes) — a testable distinction
Asymptotic cosmic cooling / heat death	BASE noise floor ε is a structural lower bound; substrate temperature permanently at T_{CMB}, $T \to 0$ never reached

Acknowledgments

The human author wishes to acknowledge the vast community of theoretical physicists whose work, far exceeding what any reference list can capture, formed the intellectual landscape from which HFT emerged. This framework is an emergent synthesis: its formal derivations rest on foundations laid by many hands, and its conceptual leaps were made possible only because those foundations existed. The human author also thanks her AI co-authors for the rigorous formalization that transformed topological intuition into precise mathematics, and thanks Anthropic and Google for developing the AI environments that made this collaboration possible.

本論文的靈感來自諸多前輩的工作成果，除了明確引用的論文外，尚包含高能物理學的規範場論、凝態物理學的對稱破缺與相變圖景、廣義相對論的幾何化引力觀、連續介質彈性力學的張力場描述，以及資訊理論與統計力學中熵驅動的時間湧現觀。若無前人建立的這些理論支柱，作者實難跨領域整合概念與數學工具，完成這個拼圖。

散文家陳之藩先生在讀過愛因斯坦氏的《The World As I See It》後，深有所感，容我在此引用：

> 無論什麼事，得之於人者太多，出之於己者太少。因為需要感謝的人太多了，就感謝天罷。無論什麼事，不是需要先人的遺愛與遺產，即是需要眾人的支持與合作，還要等候機會的到來。越是真正做過一點事，越是感覺到自己貢獻的渺小。

Author Contributions

The division below reflects the actual workflow: all picture commits, ontological choices, and structural conjecturing reside with the human author. AI contributions execute mathematical and editorial tasks within picture commits already specified by the human, or retrieve cross-domain physics concepts for the human to evaluate against the picture; no AI contribution drove ontology or proposed structural commits. Within the AI scope, Gemini 3.1 Pro (Google DeepMind) served as primary calculator and concept-retrieval librarian, and Claude 4.6 Sonnet (Anthropic) as primary formal auditor and editor. The human author reviewed and edited all AI-generated content and takes full responsibility for the publication.

Epistemological commitments. The picture-driving work of the human author is guided by three durable methodological priors: *(i)* the number of independent ansätze should be as small as possible — every additional ansatz is a structural debt that the picture must later justify; *(ii)* the substrate must be internally logically self-consistent — each picture commit must derive from or be compatible with the substrate's primitives without ad-hoc supplementation; *(iii)* every core structural constant should admit both a discrete reading (state counting, enumeration of topological classes) and a continuous geometric reading (projection ratios, dimension counting) that converge on the same closed-form value, with single-route derivations treated as provisional until the dual reading is found. These three priors are the picture-internal acceptance criteria against which the version-iteration audit below evaluates each draft.

Version-iteration methodology. The manuscript advances through numbered versions rather than by continuous in-place revision. A version number is fixed only when the human author judges that the picture has reached an internally coherent equilibrium at the current scope. Each finalized version is then submitted to several independent-provenance LLMs in a strict cold-reader mode — without prior conversational context — and asked to identify logical gaps, ad-hoc-looking constructions, and numerical agreements that warrant suspicion of fitting. The findings from these adversarial reads are not treated as immediate corrections; instead they set the research agenda for the next version, which the human author works through structurally before committing the picture to a new version number. Version count therefore reflects the depth of cold-reader auditing the manuscript has survived, not the amount of parameter adjustment to match observation; the picture itself remains parameter-free throughout, with a single dimensional anchor m_P.

Contribution	Author(s)
Picture commits and structural conjecturing (sole human responsibility)	
HFT genesis as a research programme	Human
Hopf-bundle ansatz $S^3 \xrightarrow{S^1} S^2$, trivalent mesh, chirality lock substrate	Human
Phase space $\mathcal{Q}$ structure (5-DOF decomposition)	Human
$128 + 9 = 137$ skeleton identity and all three coupling-constant derivations	Human
BASE noise mechanism, ε thermal floor, cross-channel coherence framing	Human
Mass cascade architecture, dual-route convergence paradigm, dark-sector partition	Human
Ontology commits: scale-invariant substrate, single operational anchor, time emergence	Human
Cosmology hypothesis: yarn-ball pre-lock state, ruler-shift framing	Human
Roadmap, topological identifications, and framework-consistency auditing	Human
Cross-domain physics concept retrieval and hypothesis proposal (AI under human evaluation)	
Literature surveying across topology, knot theory, and phase space methods	Gemini
Cross-domain concept retrieval (holographic principle, information theory)	Gemini
Hypothesis fragments proposed for the author's structural judgment	Gemini
Mathematical execution within picture commits (AI under human direction)	
App. A: mass-ladder algebraic manipulation and dimensional verification	AI
App. B: strong-sector ratio calculation within the trivalent-vertex picture	AI
App. C: QED reduction calculation following the human-specified picture commits	AI
App. D: α–θ_W identity scale-anchor numerical verification	AI
App. E: ΛCDM correspondence audit within the human-specified bridge framing	AI
Cross-appendix numerical consistency check and sub-percent residual verification	AI
Editorial (AI for prose and formatting under human direction)	
Rendering the geometric picture into accessible prose explanations	Claude
Final copy-editing, tone calibration, prose-logic auditing	Claude
LaTeX formatting and structural typesetting	Claude

Version history. The tables below record the trajectory of the manuscript through its numbered iterations. The first lists the positioning the human author judged the picture to have reached at each version; the second lists the key structural commits introduced together with the errors or gaps subsequently identified by cold-reader auditing and retired in later versions. The trajectory illustrates the operational use of the three priors above: ad-hoc constructions and numerology surface as the picture is stress-tested, are flagged for retirement, and the next version is built only after a structurally cleaner derivation is found.

Version	**Date**	**Positioning**
v1	Apr 25, 2026	Rough draft; structure dimly visible through heavy fitting.
v7	Apr 26, 2026	Framework potential revealed; GR shown to emerge from structure.
v10	Apr 28, 2026	Divergent exploration concludes; consolidation direction set.
v13	Apr 30, 2026	Initial cleanup complete; internal numerology identified.
v14	May 9, 2026	Framework-external content largely removed; framework-internal derivation only.
v15	May 13, 2026	Closed forms mature; three-coupling structure (α, θ_W, θ_v) complete.
v16	May 17, 2026	Cosmology emerges automatically from framework sub-leading order.

Version	Key commits introduced	Errors / gaps later retracted
v1	• $S^1 \to S^2$ Hopf substrate • $S_E = 137 \times 3/8 = 51.375$ framework • Călugăreanu writhe-as-dark-sector	• Wyler/$SO(5,2)$ external bounded-domain α^{-1} • three-Writhon $W^{(n)}$ hierarchy • pervasive numerological fitting across most observables (the majority of the manuscript)
v7	• $128 + 9 = 137$ first derived • GR emergence • Wyler peak treated with rigor	• Wyler route still external • black-hole–particle isomorphism
v10	• $\sin^2 \theta_W = 30/128$ closed form • Methodological Position stated • Higgs as amplitude mode	• Dodecahedron-based $N_{\text{weak}} = 30$ • SGWB $\delta S_E \approx 0.154$ fitted from electron residual
v13	• Vertex-fiber identification fixes $k = 3$ • per-cell hex $128 = 2^3 \times 2^2 \times 4$ • QED worked example • Strong CP geometric decoupling • rest mass and baryogenesis from chiral lock	• residual numerology (CKM/PMNS formulas, Pythagorean M_t, etc.) • $W^{(1)}/W^{(2)}/W^{(3)}$ hierarchy preserved • Wyler/Bergman retained • BH–particle regime coincidence at m_P
v14	• cleanup and restructuring of the framework • scale-free substrate + single m_P anchor • cohomology four-channel • $\theta_v = 1/32$ first appears	• Continuous-space substrate not yet explicit • closed-form vev not yet locked • cosmology framing still ad hoc
v15	• continuous-space substrate • mass cascade locked in closed form • α–θ_W geometric identity • time emergence formalised • CKM/PMNS numerology cleared from main text	• minor residual numerology • residual dressings across cohomology channels and structural constants not yet unified • phase space $\mathcal{Q}$ description carries multiple framework-internal tensions
v16	• universal BASE noise as sub-leading correction • four-channel cohomology manifestation • ΛCDM bridge hypothesis • phase space $\mathcal{Q}$ precisely defined • m_e over-determination check	• cosmology remains a hypothesis layer • further bridges to SM phenomenology, GR, and cosmological dynamics deferred to follow-up work

References

[1] Planck Collaboration, "Planck 2018 results. VI. Cosmological parameters," *Astronomy & Astrophysics*, vol. 641, p. A6, 2020. DOI: 10.1051/0004-6361/201833910.

[2] H. K. Urbantke, "The Hopf fibration—seven times in physics," *Journal of Geometry and Physics*, vol. 46, pp. 125–150, 2003.

[3] E. Witten, "Quantum field theory and the Jones polynomial," *Communications in Mathematical Physics*, vol. 121, no. 3, pp. 351–399, 1989.

[4] L. Susskind, "The world as a hologram," *Journal of Mathematical Physics*, vol. 36, no. 11, pp. 6377–6396, 1995.

[5] G. 't Hooft, "Dimensional reduction in quantum gravity," in *Salamfestschrift*, World Scientific, 1993; arXiv:gr-qc/9310026.

[6] P. W. Higgs, "Broken symmetries and the masses of gauge bosons," *Physical Review Letters*, vol. 13, no. 16, pp. 508–509, 1964.

[7] F. Englert and R. Brout, "Broken symmetry and the mass of gauge vector mesons," *Physical Review Letters*, vol. 13, no. 9, pp. 321–323, 1964.

[8] S. Weinberg, "A model of leptons," *Physical Review Letters*, vol. 19, no. 21, pp. 1264–1266, 1967.

[9] S. Coleman, "Fate of the false vacuum: Semiclassical theory," *Physical Review D*, vol. 15, no. 10, pp. 2929–2936, 1977.

[10] R. D. Peccei and H. R. Quinn, "CP conservation in the presence of pseudoparticles," *Physical Review Letters*, vol. 38, no. 25, pp. 1440–1443, 1977.

[11] H. Nyquist, "Certain topics in telegraph transmission theory," *Transactions of the American Institute of Electrical Engineers*, vol. 47, no. 2, pp. 617–644, 1928.

[12] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.

[13] A. D. Sakharov, "Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe," *JETP Letters*, vol. 5, pp. 24–27, 1967.

[14] G. Călugăreanu, "Sur les classes d'isotopie des nœuds tridimensionnels et leurs invariants," *Czechoslovak Mathematical Journal*, vol. 11, no. 4, pp. 588–625, 1961.

[15] J. H. White, "Self-linking and the Gauss integral in higher dimensions," *American Journal of Mathematics*, vol. 91, no. 3, pp. 693–728, 1969.

[16] F. B. Fuller, "The writhing number of a space curve," *Proceedings of the National Academy of Sciences USA*, vol. 68, no. 4, pp. 815–819, 1971.

[17] S. Navas *et al.* (Particle Data Group), "Review of Particle Physics," *Physical Review D*, vol. 110, p. 030001, 2024.

[18] G. S. Bali, "QCD forces and heavy quark bound states," *Physics Reports*, vol. 343, no. 1, pp. 1–136, 2001. DOI: 10.1016/S0370-1573(00)00079-X.

[19] G. Spencer-Brown, *Laws of Form*, Allen & Unwin, London, 1969.

[20] T. H. R. Skyrme, "A unified field theory of mesons and baryons," *Nuclear Physics*, vol. 31, pp. 556–569, 1962.

[21] R. C. Tolman, "On the estimation of distances in a curved universe with a non-static line element," *Proceedings of the National Academy of Sciences USA*, vol. 16, no. 7, pp. 511–520, 1930.

[22] I. M. H. Etherington, "On the definition of distance in general relativity," *Philosophical Magazine*, vol. 15, no. 100, pp. 761–773, 1933.

[23] A. H. Guth, "Inflationary universe: A possible solution to the horizon and flatness problems," *Physical Review D*, vol. 23, no. 2, pp. 347–356, 1981.

[24] D. J. Fixsen, "The temperature of the cosmic microwave background," *Astrophysical Journal*, vol. 707, no. 2, pp. 916–920, 2009.

[25] D. Clowe *et al.*, "A direct empirical proof of the existence of dark matter," *Astrophysical Journal Letters*, vol. 648, no. 2, pp. L109–L113, 2006.

[26] A. G. Riess *et al.*, "A comprehensive measurement of the local value of the Hubble constant with 1 km s^{-1} Mpc^{-1} uncertainty from the Hubble Space Telescope and the SH0ES team," *Astrophysical Journal Letters*, vol. 934, no. 1, p. L7, 2022.

[27] I. Labbé *et al.*, “A population of red candidate massive galaxies ~ 600 Myr after the Big Bang,” *Nature*, vol. 616, pp. 266–269, 2023.

[28] G. Agazie *et al.* (NANOGrav Collaboration), “The NANOGrav 15 yr data set: Evidence for a gravitational-wave background,” *Astrophysical Journal Letters*, vol. 951, no. 1, p. L8, 2023.

[29] R. Abbott *et al.* (LIGO Scientific, Virgo, KAGRA Collaborations), “Upper limits on the isotropic gravitational-wave background from Advanced LIGO and Advanced Virgo’s third observing run,” *Physical Review D*, vol. 104, no. 2, p. 022004, 2021.

[30] P. A. R. Ade *et al.* (BICEP/Keck Collaboration), “Improved constraints on primordial gravitational waves using Planck, WMAP, and BICEP/Keck observations through the 2018 observing season,” *Physical Review Letters*, vol. 127, no. 15, p. 151301, 2021.