Matter, Gauge Structure, and Spectroscopy in the Finite-QEC Substrate

David Elliman
Neuro-symbolic Ltd
dave@neusym.ai

Draft canon snapshot: 2026-06-25, rolling local refresh

Abstract

This paper is the matter-sector companion to the overview and foundations papers in the canon snapshot series. It collects the finite-code results that produce Standard-Model-like matter bookkeeping, the gauge-structure results that survive current audits, and the spectroscopy results that remain load-bearing after the numerical-search and retraction tests. The central message is deliberately split. The strongest matter-sector results are exact finite identities: the $[8, 4, 4]$ byte/fermion dictionary, charge bookkeeping, anomaly cancellation, the separation of spatial, colour, generation, and repair-axis roles, and the recovered chiral charge-cube identity $\sum Q_L^3 = -2/9$. The continuum and dynamical results are more mixed but sharper than in the original draft: the chiral/mirror-gapping program has substantial finite-cell, extended-cutoff, and induced-coupling evidence, with the remaining frontier reduced to a pure-Wilson-axis continuum theorem; the velocity-unification loop has passed the current Dirac-triple one-loop tests under named premises; and CP violation is now localized to a $\Delta L = 2$ Majorana holonomy portal whose substrate origin remains open. The paper's goal is therefore not to declare the matter sector closed, but to state which parts are exact finite arithmetic, which parts are reproducible computations, and which parts remain continuum-lift or operator-definition frontiers.

Contents

1 Role of This Paper

The matter and gauge sector is the cleanest place to see the framework's two standards of evidence. On one side are finite register facts. These are combinatorial, exact, and reproducible by enumeration. On the other side are continuum and spectroscopic claims. These require an observable map, renormalisation or dressing convention, and an audit against nearby formulas.

The paper follows that distinction. Sections 2– 4 state the finite-code structure. Sections 5 and 6 discuss the chiral gauge and mirror-gapping problem in the language of lattice gauge theory, where the relevant external context is the Nielsen–Ninomiya obstruction, Ginsparg–Wilson/domain-wall routes, and symmetric mass generation [6, 10, 14, 20, 22, 24]. Sections 7–8 state the spectroscopy and velocity-sector results with their current caveats.

All scripts named here are in the reproducibility repository [7]. The companion file `script_index.md` gives a shorter invocation and purpose ledger.

2 Finite Register Content

The local matter carrier is an eight-bit register with a physical subset $\mathcal{P} \subset \mathbb{F}_2^8$. The code is tied to the $[8, 4, 4]$ structure used throughout the canon, but the matter sector uses more than the abstract code: it also uses a fixed interpretation of bits into chirality, weak isospin, colour, generation, and matter/antimatter bookkeeping.

The useful schematic split is

$$\mathbb{F}_2^8 = \mathcal{P} \cup \mathcal{Q},$$

where $\mathcal{P}$ is the finite matter content and $\mathcal{Q}$ is not thrown away. Invalid states form the correction and strain target for the QEC engine. This is essential for later sections: mirror gapping, baryogenesis, dark-sector wall shadows, and horizon channels all depend on whether a state is a physical codeword, an invalid state, or a strain-readable transition between them.

2.1 Matter Content as a Finite Classification

The current canon treats the Standard-Model-like matter content as a finite classification result. In one generation, the register reproduces the ordinary pattern of quarks and leptons, including a right-handed neutrino state. This is structurally close to the familiar $SO(10)$ spinor **16**, and the canon is explicit about the comparison: on pure Target-A matter content — the series' label for discrete structural claims (multiplets, charges, anomaly arithmetic), as opposed to Target-B continuous parameters (masses and mixings) — this framework is not a more parsimonious explanation than $SO(10)$ [8]. Its value is that the same finite register is also used by the QEC, strain, and spectroscopy machinery.

Thus the matter-content result is LOCKED in the finite register sense but not a claim of unique theoretical economy. The exact output is the classification and bookkeeping of the states; the open physical work lies in their continuum propagation, masses, and interactions.

The current canon sharpens the finite classification into a byte-level fermion dictionary. One matter cell has

$$4 \text{ Pati-Salam columns} \times 2 \text{ chiral rows} \times 2 \text{ CSS defect classes} = 16$$

fermion slots, and the three-generation inventory is therefore $48 = 3 \times 16$. The sterile ν_R appears as the passive zero-syndrome corner rather than as an added hidden-particle assumption. This is still a finite bookkeeping theorem. Hypercharge running, mixing angles, and continuous masses remain separate dynamical questions.

2.2 Axis Booking

A major canon cleanup was the resolution of the old "three axes" ambiguity. The current rule is:

spatial rank 3	belongs to the truncated-cubic-honeycomb (TCH) / $\mathbb{Z}^3$ geometry,
colour 3	belongs to the internal one-hot colour orbit,
generation 3	belongs to the allowed generation sectors,
R4 2/3	counts legal repair edges per generation sector.

These are not four appearances of one hidden spatial axis postulate. They are different categories. The axis-booking audit is important because it prevents accidental reuse of the same integer as if it independently derived space, colour, generation, and dark-sector support.

3 Charge and Gauge Bookkeeping

The gauge sector begins as finite charge bookkeeping before it becomes a continuum gauge theory. The register supplies:

- a colour one-hot basis;
- weak and chirality labels;
- a matter/antimatter bookkeeping tag; and
- charge assignments forced by anomaly-compatible bookkeeping.

The electroweak cleanup is especially important. The older canon sometimes presented a geometric 2/9 value for $\sin^2\theta_W$. That claim is now RETIRED. The register charge trace instead gives the standard GUT-normalised value

$$\sin^2\theta_W = \frac{\mathrm{Tr}(T_{3L}^2)}{\mathrm{Tr}(Q^2)} = \frac{3}{8}$$

at the unification-normalised level — the classic grand-unification value [9] — with the observed low-energy value belonging to ordinary renormalisation-group running. This is a success of bookkeeping consistency, not a new derivation of the measured weak mixing angle.

The gauge web itself is treated in the Wilson/Kogut–Susskind tradition [15, 26]. The finite-TCH scripts then ask a separate question: whether the local gauge-cell and mirror-Fock operators support the required chiral and mirror-gapped phases without introducing the wrong low-energy degrees of freedom.

4 Anomaly Arithmetic and the Recovered $-2/9$

The cleanest recovered item from the old "universal 2/9" cluster is not a universal invariant. It is an exact charge-cube identity. In the one-generation register, exact rational arithmetic gives

$$\sum_{\text{left}} Q^3 = -\frac{2}{9}, \qquad \sum_{\text{right}} Q^3 = -\frac{2}{9}, \qquad \sum_{\text{left}} Q^3 - \sum_{\text{right}} Q^3 = 0.$$

This is a genuine finite identity. It has no tolerance band, no fitted observable, and no alphabet search. It is therefore Locked as anomaly bookkeeping in the Adler–Bell–Jackiw sense [1, 2].

What is retired is the old banner claim that the same exact 2/9 unified the Weinberg angle, Koide phases, and chiral anomaly. The Weinberg leg is charge-forced to 3/8 at the UV normalisation. The Koide leg is a phase or angle statement, not an exact rational identity. The canon now treats the anomaly identity as the recovered theorem and the universal-invariant framing as Retired.

5 Strong CP, Chirality, and Continuum Lift

The strong-CP story is finite-substrate-cleaner than it is continuum-clean. At the bare register level, the mass operators are Hermitian in the code basis, so the finite-substrate mass phase gives

$$\arg\det(M_u M_d) = 0.$$

Together with a zero bare gauge θ-term this closes the bare $\bar{\theta}$ half of the strong-CP problem. The continuum question is different: instanton physics in continuum QCD [23], and the standard axion relaxation route [21], live after the chiral gauge continuum limit is taken. The open problem is therefore not the bare finite arithmetic. It is the continuum chiral gauge lift and the preservation of the bare cancellation along that lift.

The relevant issue is familiar from lattice gauge theory. A finite lattice fermion construction must evade the no-go constraints that obstruct naive chiral fermions [20]. The canon therefore treats the following as distinct:

1. finite register charge and anomaly bookkeeping;
2. a candidate Ginsparg–Wilson or domain-wall-like continuum operator;
3. a mirror sector that becomes massive without breaking the visible gauge symmetries; and
4. the continuum limit of the full gauge-coupled TCH construction.

This separation is not conservative wording for its own sake. It prevents a finite charge enumeration from being mistaken for a solved continuum chiral gauge theory. The scripts in this sector now include explicit Ginsparg–Wilson construction attempts, body-diagonal no-go checks, overlap-kernel tests in the sense of Neuberger [19], and audits of the abelian preconditions for an exact-gauge-invariance route [17]. They are evidence for a path, not a finished continuum theorem.

The consolidated state of that path is sharper than a flat "open." The continuum lift splits by gauge factor: the abelian $U(1)_Y$ sector *closes* rigorously (the fermions cancel all $U(1)$ anomalies, so Lüscher [17] supplies a smooth, local, gauge-invariant chiral measure); $SU(3)_c$ is *vector-like* (real determinant, no chiral obstruction); and the $SU(2)_L$ Witten global anomaly is absent (an even number of doublets). The single binding residual is therefore the non-abelian $SU(2)_L$ overlap-measure construction — a field-wide open problem of lattice gauge theory, on which admissible-orbit-space homotopy invariance shows the substrate gives no leverage either way. The mirror-gap "collapse" reported by an early reduced (colour-orbit charge-block) proxy was *retired* by the non-reduced Peter–Weyl spin-network computation, in which the matter/mirror gap holds (≈ 2; at most $\sim 12\%$ softening under hopping; no collapse) across all tested cutoffs — positive-leaning evidence, though no finite-row continuum theorem is claimed. The dependent neutron EDM $d_n \sim 10^{-31}$ is consequently a *conditional* prediction resting on the same premise the entire Standard-Model continuum rests on, not a negative-leaning one.

There is an analogous cleanup in the CP/baryogenesis sign sector. The real documented walk cannot generate a CP-odd sign by itself. The missing operator class is now identified: a complex-symmetric Majorana $\Delta L = 2$ portal for the sterile generation sector,

$$\mathcal{L}_{\Delta L=2} = \frac{1}{2}\,\nu_R^T C\, M_H\, \nu_R + \text{h.c.},$$

with a generation-blind K_3-holonomy structure. Within the admissible class this is both the right kind of object and sufficient to produce nonzero weak-basis CP invariants whose sign flips with the holonomy orientation. A recovery-holonomy audit now narrows the boot/QEC target: if one grants an oriented, unconditional sterile-generation recovery latch P_σ on the three R1-allowed generation corners, then Majorana symmetrisation gives $P_\sigma + P_\sigma^T = A_{K_3}$, and S_3-blindness independently singles out the same off-diagonal support. Thus the K_3 support of the portal is conditionally forced by the recovery-cycle primitive. The directed latch itself is not S_3-blind: it reduces the generation relabelling symmetry to an oriented C_3 subgroup, while only the Majorana support is S_3-covariant. The existing R4/source repair algebra supplies within-generation erasure edges and a generation-singlet source port, not such an inter-generation latch. What is still open is deriving that primitive from the actual boot/QEC mechanics, deriving its holonomy phase Φ, and selecting the orientation σ. If both orientations are summed without a selector, the CP invariant cancels. Thus the CP operator support is conditionally identified; the CP sign is not yet a locked prediction. A later closure audit makes the obstruction algebraic: $P_\sigma = (A_{K_3} + \sigma K_{\text{or}})/2$, where K_{or} is a generation-orientation tensor in the sign representation of S_3. The existing generation-blind repair algebra contains A_{K_3}-type support but not a service-readable K_{or}. A register-level hunt identifies the candidate K_{or} as the oriented boundary cochain of the R1-allowed triangle in the ordered (G_0, G_1) bit-plane. This object has exactly the required sign-representation covariance, but it is erased by S_3-blind source averaging, includes one Hamming-2 generation edge, and relies on the fixed G_0, G_1 bit order being physically readable rather than a relabelling convention. A subsequent readout theorem sharpens the obstruction: in the two-component decomposition $\{A_{K_3}, K_{\text{or}}\}$, an S_3-scalar Stinespring pointer has no equivariant channel from the sign component, so it can read A_{K_3} but not K_{or}. The oriented R1 boundary becomes readable exactly when the recovery environment carries a pseudoscalar orientation pointer, transforming in the sign representation of S_3, so that oK_{or} is an S_3-scalar record. The latest R1 audit separates two objects that earlier wording blurred: the two Hasse covers give the CP-even δ covariance, while the CP-odd sign carrier is the *closed* oriented

R1 boundary cochain, including the endpoint–endpoint edge. Multiplying that closed cochain by the global orientation line ω gives the correct scalar sign pointer. This is a real localization of the CP sign, but not a closure of baryogenesis: $\Phi = 2\pi/3$ is a CP zero in the baryogenesis invariant $\sin(3\sigma\Phi)$, and the $\Delta L = 2$ portal plus a nonzero phase magnitude remain open gates.

The same audit also clarifies what this holonomy progress can and cannot do for PMNS mixing. The active-neutrino mass spectrum remains the clean part of the sector: the $\delta_\nu = 1/3$ branch fixes the normal-ordering spectrum. The angles are a different object. The algebra generated by I, A_{K_3}, and $K_{\rm or}$ is still C_3-circulant, so its eigenvectors are the discrete Fourier basis. It therefore leaves the bare neutrino matrix at trimaximal $\theta_{13} = 35.26°$, or collapses θ_{13} to zero if the charged leptons share the same circulant basis. The sign pointer can carry a CP orientation; it is not, by itself, an angle texture. A PMNS-angle promotion audit shows that both the representative observed matrix and the attractive universal-2/9 bimaximal/QLC matrix require a large component outside the circulant algebra.

There is nevertheless a useful partial derivation. Let the bimaximal base be

$$U_0 = R_{23}(\pi/4)R_{12}(\pi/4),$$

and let a small right-acting real generation texture be $A(x, y, z) \in \mathfrak{so}(3)$. The S_3 standard representation is the plane $x + y + z = 0$. On this plane the first-order PMNS angle response obeys

$$\delta\theta_{23} = -\frac{\delta\theta_{12}}{\sqrt{2}}.$$

Thus a pure standard-representation texture at strength $\delta = 2/9$ naturally gives the reactor-angle scale $\theta_{13} \simeq \delta/\sqrt{2} \simeq 8.5°$, but it cannot also hold $\theta_{23} = 45°$ while shifting θ_{12} to $45° - \delta$. The exact QLC tangent is

$$A_{\rm QLC} = \begin{pmatrix} 0 & -1 & 1/2 \\ 1 & 0 & 1/2 \\ -1/2 & -1/2 & 0 \end{pmatrix}, \qquad \mathbf{a}_{\rm QLC} = (-1/2, 1/2, 1),$$

which decomposes into a dominant standard component plus a nonzero sign-representation component. The standard component supplies the reactor-angle lever; the sign component cancels the atmospheric drift and completes the solar shift. Since the current $K_{\rm or}$ work licenses the sign sector as a CP-orientation pointer rather than as real eigenvector rotation, the QLC numbers remain partially derived and mechanism-open rather than fully predicted.

A follow-up bridge gate shows exactly what extra primitive would close the tangent. As a real antisymmetric operator, $K_{\rm or} = A(1, 1, 1)$, and

$$K_{\rm or}^2 = -3P_{\rm std},$$

so the oriented R1 boundary is a Hodge orientation on the two-dimensional S_3 standard plane. If one completed oriented boundary service is billed as the mean sign current $K_{\rm or}/3$, and the standard shear is required to carry the universal 2/9 reactor primitive while preserving the atmospheric latch at first order, then the standard component is forced to be $(-5/6, 1/6, 2/3)$. Adding the mean sign current gives

$$(-5/6, 1/6, 2/3) + (1/3, 1/3, 1/3) = (-1/2, 1/2, 1),$$

the QLC tangent above. In this conditional bridge, the solar shift $\delta\theta_{12} = -\delta$ follows from the sign-current normalization rather than being separately inserted. The live theorem is therefore not a new angle fit, but an eigenvector-lift theorem: prove that the physical lepton-sector recovery readout bills the oriented R1 boundary as a real $\mathfrak{so}(3)$ PMNS rotation with mean-cycle normalization, not merely as the Hermitian iK_{or} CP-orientation pointer.

An eigenvector-lift audit separates the remaining options. The existing global-orientation theorem can supply the sign pointer as a scalar Stinespring record, but record readout is not automatically basis transport. If K_{or} is inserted as the Hermitian mass/CP operator iK_{or}, it stays inside the circulant algebra and keeps the DFT eigenvectors. If the directed boundary is treated as an ordinary reset or jump, the conditioned event is a 120° generation permutation and the unconditioned channel is dissipative; neither is the small coherent connection $\exp(\delta K_{\mathrm{or}}/3)$. Moreover representation theory does not fix the coefficient: qK_{or} has the correct sign covariance for any real q. The value $q = 1/3$ must come from a service-normalisation lemma, namely that one completed oriented boundary winding is averaged over the three R1 generation corners. Thus the proof route is now narrow: derive a coherent-recovery polar factor, applied before irreversible reset, whose generation-frame holonomy is $U_{\mathrm{frame}} = \exp(\delta K_{\mathrm{or}}/3)$. Under that additional lemma, the sign/standard construction above becomes a genuine PMNS eigenvector-lift theorem.

The strongest present result is a conditional uniqueness theorem for that lemma. Suppose the lepton recovery instrument has a coherent one-port branch before reset, and suppose its oriented generation datum is only the global-orientation-contracted R1 boundary. Then S_3 covariance of real antisymmetric generation-frame transport leaves a one-dimensional generator space, spanned by K_{or}. If one full oriented boundary circulation is billed as one service winding shared over the three R1 generation corners, the coefficient is fixed to $K_{\mathrm{or}}/3$. Therefore any such coherent pre-reset polar branch has

$$U_{\mathrm{frame}} = \exp(\delta K_{\mathrm{or}}/3),$$

up to the global orientation sign. This proves uniqueness under the coherent frame-branch premise. It still does not prove that lepton recovery possesses that branch: current canon proves sign readout and a coherent source port, but not that the sign record is implemented as a pre-reset unitary frame transport.

6 Mirror Gap and Symmetric Mass Generation

The mirror problem is the matter-sector frontier that most resembles a standard lattice field-theory problem. The target is to keep the desired chiral content light while giving the mirror sector a symmetric mass gap. This is the setting of symmetric mass generation [6, 24].

The canon now has several layers of evidence:

- finite SMG blocks and mirror-only checks;
- code-projection and invalid-subspace audits;
- gauge-cell dressed operators on TCH plaquettes;
- Peter–Weyl and Clebsch–Gordan endpoint tests in 6, $\bar{6}$, and 8 channels;
- non-reduced spin-network bases with explicit link indices; and
- extended-representation/higher-cutoff runs, including the 106,460-state sparse computation reported in the canon.

The present status is Computed but not Locked. The local and finite-cutoff sector is positive: the extended magnetic and higher-cutoff run shifted the gap by only about two percent relative to the minimal basis at the tested points. The later RG-basin reduction also removes one possible ambiguity: for the registered number-conserving Hamiltonian, the mirror block induces no adjoint or irrelevant pure-gauge deformation,

$$(\beta_F, \beta_A, \lambda_6, \lambda_8, \ldots) = (\beta_F, 0, 0, 0, \ldots).$$

The continuum frontier is therefore narrower:

$$\beta > \beta_{\text{cert}} \simeq 0.661 \quad \longrightarrow \quad \beta = \infty.$$

A continuum-theorem gate now separates the framework-specific part from the ordinary lattice-gauge part. In the registered Gauss-projected Hamiltonian,

$$P_{\text{vac}} H_{\text{TCH}} P_{\text{ch}} = 0, \qquad \delta\beta_A = \delta\lambda_6 = \delta\lambda_8 = \cdots = 0,$$

so the mirror block does not push the flow into a mixed-action fundamental–adjoint basin. The coupling-space path is the pure fundamental Wilson axis,

$$\gamma(\beta) = (\beta, 0, 0, \ldots),$$

plus a massive spectator sector. The correct observable is forced by Gauss law. A charged mirror endpoint is not gauge invariant by itself; the transfer-matrix sector that contains it also contains the minimal static flux or Wilson-line dressing required by the pure gauge theory. Thus the raw charged-sector gap decomposes as

$$\Delta_{\text{raw}} = E_{\text{full}}(Q) - E_{\text{full}}(0) = E^{\min}_{\text{string}}(\beta, L) + \Delta_{\text{mir}}.$$

The first term is present even if the mirror residual is zero. A continuum decoupling theorem is a theorem about absence of a charged mirror pole, not a theorem about the compulsory pure-gauge static string. Therefore the observable is not the raw finite electric string gap, but the electric-subtracted mirror gap

$$\Delta_{\text{mir}}(\beta, L, \Lambda_{\text{cut}}) = E_0(\text{charged}) - E_0(\text{vacuum}) - E^{\min}_{\text{string}}(\beta, L).$$

Equivalently,

$$\Delta_{\text{mir}} = \big[E_{\text{full}}(Q) - E_{\text{full}}(0)\big] - \big[E_{\text{gauge}}(Q) - E_{\text{gauge}}(0)\big].$$

A raw-gap criterion would give false positives: the existing 2x1 smoke row has $\Delta_{\text{raw}} = C_3/\beta > 0$ and $\Delta_{\text{mir}} = 0$ to numerical precision, so it is a pure electric-string sanity check rather than a mirror-mass witness. The current finite magnetic rows give positive witnesses for the subtracted object: the 336-state cell has a minimum electric-subtracted offset about 3.99, and the two-plaquette magnetic cell gives about 4.02, while a deeper 2x1 smoke row is correctly rejected as a pure C_3/β electric-string sanity check rather than a mirror-mass witness. A later volume-bound audit adds a useful but still limited analytic step. With local SMG cost $D_{\text{SMG}} = 2$, strip incidence $n_{\text{inc}} = 2$, and $\|W_p + W_p^\dagger\| = 2$, the charged-sector gap obeys

$$\Delta_{\text{raw}} \geq D_{\text{SMG}} + E_{\text{static}}(\beta) - \frac{\beta}{2} n_{\text{inc}} \|W_p + W_p^\dagger\|.$$

The Gauss-law pinned-flux term extends *raw* charged-sector positivity to $\beta \simeq 1.46$. The continuum observable above, however, subtracts that same C_3/β string energy. The certified electric-subtracted mirror offset is therefore only

$$\Delta_{\text{mir}} \geq 2 - 2\beta,$$

positive on the subdomain $\beta_{\rm cert} \le \beta < 1$. Thus the finite audit genuinely pushes the mirror-offset certificate beyond $\beta_{\rm cert} \simeq 0.661$, but it does not close the weak-coupling continuum tail. The framework-specific continuum obstruction is reduced to a conditional theorem: if the pure Wilson $SU(3)$ axis has no finite-β bulk transition for $\beta \ge \beta_{\rm cert}$, and if

$$\inf_{\beta\ge 1,\, L,\, \Lambda_{\rm cut}} \Delta_{\rm mir}(\beta, L, \Lambda_{\rm cut}) > 0,$$

then no charged mirror pole appears and the mirror sector decouples in the continuum limit. A direct numerical sweep now supplies the missing weak-coupling, large-cutoff evidence for the second input. Extending a matter-resolved mirror-gap probe over the local cutoff (states-per-charge $2 \to 6$), volume ($L = 1 \to 2$, to Hilbert dimension $\sim 1.6 \times 10^6$), and coupling ($\beta = 0.5 \to 4$) at *non-perturbative* hopping (t up to 1.0) along the pure Wilson axis, the gap stays open: the cutoff extrapolation $\Lambda_{\rm cut} \to \infty$ has a positive infimum over the whole grid, inf ≈ 1.24 in lattice units (worst case near $\beta = 2$), and the gap is volume-invariant (the L=1 and L=2 values agree to within 4% even at the most stressed $\sim 10^6$-state cell). This promotes the volume/cutoff-uniform lower-bound input from small-cell witnesses to an extended-cutoff, volume-extrapolated, $O(1)$-hopping result. What is *not* proved internally now reduces to the pure-Wilson-axis analyticity/no-bulk-transition premise and the lift from the Z_3 colour-orbit proxy to full $SU(3)$ (`python_code/smg_cutoff_extrapolation_deep.py`, `python_code/smg_volume_strip_deep.py`). The draft therefore records the SMG sector as a theorem-grade reduction plus extended-cutoff, volume-extrapolated numerical evidence, not as a completed proof of the Standard Model continuum.

The clean way to state the remaining lattice-gauge input is as a separate Wilson-axis premise:

$$\begin{aligned} W_0: \quad & \text{pure Wilson } SU(3) \text{ at } \beta_A = 0 \\ & \text{has no zero-temperature finite-}\beta \text{ bulk transition for } \beta \ge \beta_{\rm cert}. \end{aligned}$$

The finite TCH algebra does not prove W_0. What it proves is the projection onto the W_0 problem: the registered mirror Hamiltonian leaves $\beta_A = \lambda_6 = \lambda_8 = \cdots = 0$, so the known mixed-action danger is not generated by the substrate itself. The supporting evidence is external lattice-gauge evidence. Studies of mixed fundamental–adjoint $SU(3)$ actions locate the bulk line and its endpoint away from the Wilson axis [3, 4]; the improved estimate $(\beta_f, \beta_a) = (4.00(7), 2.06(8))$ and the separation of the finite-temperature deconfinement line from the bulk line for $N_t > 4$ [4] support, but do not mathematically prove, the Wilson-axis import. Heller's scaling analysis along the fundamental/Wilson axis [13] and later mixed-action artefact studies [12] put the same point in practical lattice terms: the standard Wilson-axis continuum route is plausible and conventional, but it is not a theorem derived by finite-cell SMG. The sharp closure statement is therefore

$$\text{TCH–SMG continuum lift} = \text{internal Wilson-axis projection} + W_0 + \inf_{\beta \ge 1} \Delta_{\rm mir} > 0.$$

If W_0 is false, no finite-cell mirror-gap number repairs the continuum lift; if W_0 is imported or proved and the uniform electric-subtracted mirror bound holds, the charged mirror pole is excluded on the continuum path.

6.1 Testing the uniform bound: the mirror gap under escalation

The closure statement leaves a single object to settle numerically: the uniform lower bound $\inf_{\beta\ge 1,\, L,\, \Lambda_{\rm cut}} \Delta_{\rm mir} > 0$. We report a direct test and a first analytic argument; full methodology and data are in `python_code/smg_spinnet_escalation_scope.md` and `python_code/smg_gap_stability_bound.md`.

What $\Delta_{\rm mir}$ is, and why it is the decisive number. A charged (mirror) excitation in the gauge-coupled theory is never isolated: Gauss's law ties it to a flux string, so the cheapest gauge-invariant charged state is a *meson* — a charge–anticharge pair joined by a string. Its energy splits cleanly,

$$\Delta_{\rm raw} \;=\; \underbrace{2\,\Delta_{\rm SMG}}_{\text{SMG mass of the pair}} \;+\; \underbrace{C_3/\beta}_{\text{string (electric) energy}} ,$$

with $\Delta_{\rm SMG} = 2$ the per-cell symmetric-mass cost — the exact gap of the self-dual $[[8, 0, 4]]$ stabiliser cell (the Z/X unification) — and $C_3 = 4/3$. The string term is gauge dynamics, not mirror physics, and dominates at strong coupling; subtracting it isolates the quantity that actually decides whether the mirror is gapped,

$$\Delta_{\rm mir} \;\equiv\; \Delta_{\rm raw} - C_3/\beta.$$

A vanishing $\Delta_{\rm mir}$ means the symmetric interaction fails to gap the mirror (the chiral lift fails); a positive, size-stable $\Delta_{\rm mir}$ is precisely the uniform bound the closure statement requires.

The numerics. We diagonalise the charged sector in the non-reduced Peter–Weyl spin-network basis — one $SU(3)$ irrep per link, a vertex-intertwiner label at each node, and the mirror-Fock matter cell — which represents the gauge-coupled theory exactly up to a chosen link-rep cutoff and plaquette number N (scripts `python_code/smg_phase1_ladder.py`, `python_code/smg_bhm_bound.py`). The pipeline is anchored by reproducing the 106,460-state two-plaquette gap to 0.8%. Two escalation axes:

(i) Volume. At fixed coupling, $\Delta_{\rm mir}(N)$ for $N = 1, 2, 3$ is flat and positive, $\Delta_{\rm mir} \approx 2\,\Delta_{\rm SMG}$; the softening of the spectral gap under a unit fermion hop *saturates* rather than running away (at $\beta = 1$: $6.9, 13.4, 14.3\%$; at $\beta = 0.5$: $5.6, 9.8, 10.0\%$ for $N = 1, 2, 3$). Exact diagonalisation tops out at $N = 4$ ($\sim 6 \times 10^6$ states, eigensolver-bound) and $N = 5$ ($\sim 2.5 \times 10^8$, basis-build-bound).

(ii) Coupling. This is the delicate axis. As β grows toward the continuum the *lattice-unit* string tension C_3/β shrinks, and a naive reading would fear the mirror deconfining. It does not. Across $\beta = 1 \to 8$ the gap stays open with no collapse, and $\Delta_{\rm mir}(t{=}0) \approx 2\,\Delta_{\rm SMG}$ is *independent of* β: the mirror cost is set by the symmetric-mass scale, not by the gauge coupling. Where the mirror is the lowest excitation ($\beta = 8$), reading its hopping bandwidth as an effective tension $\sigma_{\rm eff}$ gives $18\times$ (at $N = 2$) to $82\times$ (at $N = 3$) the value of C_3/β, growing with volume. The lesson is standard but easy to miss: the vanishing of the lattice-unit tension at weak coupling is the expected approach to the continuum (dimensional transmutation), *not* physical deconfinement; the confinement that protects the gap persists.

A conditional gap-stability bound. Why is the gap so robust, when the bare hopping is *not* small (the walk band is $\sim 8\times$ wider than the gap, so ordinary perturbation theory — and the Bravyi–Hastings–Michalakis stability theorem in its naive form — does not apply)? Because the relevant excitation is the *confined* meson. Solving its internal motion exactly (a charge in the linear potential σr of its own string, with hopping t), the gap reduction is the meson bandwidth D,

$$\Delta_{\rm mir}(t) \;=\; 2\,\Delta_{\rm SMG} - D(t, \sigma), \qquad D(t, \sigma) \;\simeq\; \frac{4t^2}{\sigma} \quad (t \lesssim \sigma),$$

the Schrieffer–Wolff result (verified, fitted exponent 2.00). Confinement turns the dangerous $O(\text{gap})$ bare hop into a *quadratically suppressed* effective bandwidth, and the β-independent mass $2\,\Delta_{\rm SMG}$ sits beneath it as a hard floor; the gap stays open precisely when $\sigma > \sigma_{\rm crit}(t) \sim t^2$, i.e. whenever confinement is present. This is the analytic content behind the numerics: *stability of the mirror gap* *is equivalent to the persistence of confinement* — which is exactly the premise W_0. The escalation data sit on the protected side of this boundary at every reachable size.

Status. The data are consistent with $\inf \Delta_{\mathrm{mir}} > 0$ at every reachable volume, coupling, and cutoff, and the bound reduces that infimum to W_0. What remains open is (a) the many-body form of the bound — the explicit Bravyi–Hastings–Michalakis / Lieb–Robinson stability constants for the full spin-network, of which the two-body calculation above is the controlling piece; and (b) W_0 itself, the standard pure-Wilson confinement premise, which the finite cell does not prove. The mirror-gap branch of the continuum lift is therefore *conditional on* W_0*, numerically supported, and not closed* — the same standing the lattice Standard Model has, now with the symmetric-mass evidence added.

7 Spectroscopy After the Audit

The spectroscopy sector is where the canon's numerical discipline matters most. The framework has a dense alphabet of small algebraic numbers, including $\sqrt{2}$, ϕ, π, small rational ratios, and powers of α_0. A close numerical match is not evidence unless the formula space is controlled.

7.1 The Nucleon Anchor

The strongest hadronic anchor is the static baryon result

$$M_N = 2\sqrt{2}\,\Lambda_{\mathrm{QCD}}.$$

In the canon this is the canonical nucleon-scale relation, not merely a post-hoc fit; numerically, with the canonical chiral anchor $\Lambda_{\mathrm{QCD}} = 0.332$ GeV it gives $M_N = 0.9391$ GeV, within 10^{-3} of the isospin-averaged nucleon mass [18]. It is also now important outside spectroscopy because the proton-primary gravity route uses the proton/nucleon anchor as its dimensionful input before outputting G and H_0; the scale fixed by the proton lands inside the independently G-anchored window at 1.6×10^{-4} — the two-anchor consistency reported in the overview paper of this series. That cross-paper use raises the stakes: the baryon derivation must remain clear about its static defect, boundary, and chiral-scale assumptions.

The baryon sector is treated as a static-defect regime. Mesons are not handled by the same linear formula. This distinction is productive: it explains why formulas that work for baryon electromagnetic splittings fail on mesons, where the relevant object is a dynamic dipole or pseudo-Goldstone mode.

7.2 Pions, Vectors, and the ρ

The pion and light-vector sector is more delicate. The canon contains a light-vector path result based on a C_8 or octagon-perimeter structure and a line-graph P_4 eigenvalue. Earlier concerns were that the result was a two-dimensional vertex-figure artifact. Later kernel-polynomial-method (KPM) [25] and gauge-web checks support survival as a narrow resonance under the pinned Grover-walk coupling [11], while also clarifying that strong coupling to the full gauge web would dissolve it. This is a COMPUTED survival statement with specified host-path and coupling premises.

The pion sector is anchored through a different mechanism: dynamic chiral screening and bridge dynamics. The paper does not treat every pseudoscalar mass as derived. Kaons and heavier mesons hit a hard-defect or cutoff boundary and remain separate targets.

7.3 Koide Tau

The Koide relation [16] remains a useful diagnostic but no longer carries the old "exact rational phase" interpretation; with current lepton masses the ratio $Q = 2/3$ holds at the 10^{-5} level [18]. The lepton-sector Koide phase is empirically close to 2/9, but the phase itself is an angle. The Lindemann–Weierstrass and holonomy audits support the current reading: exact rational defect ratios may appear in the finite code, while phase appearances are transcendental projections. Therefore the tau result is not presented as an exact 2/9 theorem. It is a structured, high-precision consistency result with an open operator map.

7.4 Baryon Multiplets and Heavy Quarks

The baryon scripts contain useful step-operator, $SU(6)$, De Rújula–Georgi–Glashow (DGG) [5], and decuplet tests. Their current role is not to claim that every heavy-quark mass has been derived from the substrate. The canon explicitly keeps strange, charm, bottom, and top mass scales in Target-B unless a future derivation provides a unique observable map with a bounded formula space. This is an important anti-overclaim rule: the framework can use heavy-flavour spectroscopy as a testbed without pretending that dense-alphabet matches are predictions.

The 2026-06-21 Target-B audit makes that rule quantitative. In the De Rújula–Georgi–Glashow convention, $m_c/\Lambda = 5.1487$ has 28 reduced low-complexity competitors inside 1%, while $m_b/\Lambda = 15.1883$ has 84. The top/electroweak ratio is still more crowded: m_t/Λ and v/Λ have 97/137 competitors in the same window. Thus the useful statement is not "the dense alphabet predicts heavy-quark masses", but the opposite: once the QCD and electroweak anchors are supplied, the finite spectroscopy can be tested against heavy baryons and quarkonia, but the heavy mass values themselves are not substrate predictions.

8 Velocity and Vacuum Polarisation

The vacuum-polarisation programme (canon item 115) is the matter-sector bridge between finite register kinematics and continuum propagation. The scalar-hop reading of the original §3.2 kernel is now known to be electromagnetically inert at one loop: its vacuum bubble vanishes exactly by Ward commutation. The live object is instead the chirality-Weyl Dirac-triple kernel, forced by the physical-set closure and charge-blindness enumeration. The important methodological advance is that the loop machine is not a count-only ansatz: it evaluates an operator kernel with scheme discipline, including the invariant combination that separates electric and transverse pieces.

This result now feeds other sectors. The cosmological-constant generation vertex, for example, uses the same style of loop computation rather than reusing a bare α_0-dressing by hand. For the matter paper, the status is:

- the tree-level K_6 anisotropy is reproduced by an explicit O_h $k \cdot p$ calculation;
- the scalar-hop loop is exactly zero at this order, so it cannot supply the velocity flow;
- the chirality-Weyl Dirac triple is the unique nontrivial physical-set-closed, charge-blind mask and gives a live one-loop vacuum-polarisation calculation;
- the common-mode and anisotropic on-shell tests pass under named premises, after superseding earlier divergent readings with their scheme errors identified; and
- the full continuum RG and matching statement remains a lift, not a finite-code theorem.

9 The Electroweak Scale and the Second Anchor

The electroweak vacuum expectation value $v = 246\,\text{GeV}$ has been the framework's deepest open input: a *second* dimensionful anchor beside the QCD scale, which a strict single-anchor reading treats as the leading falsification risk. This section records its recent reduction from an irreducible input to a *predicted* radiative scale, at the level of a proposition resting on named premises.

The mechanism is forced. The substrate is dimensionless: there is no bare Higgs mass-squared to insert by hand, any more than a bare Λ_{QCD}. The electroweak scale must therefore be generated *radiatively*—a Coleman–Weinberg / dimensional-transmutation scale. Elimination makes this sharp: a tree $\mu^2 \sim M_P^2$ misses the hierarchy by seventeen orders, and $\mu^2 \sim \Lambda_{\text{QCD}}^2$ would need a quartic $\lambda \sim 10^{-6}$; the only option without fine-tuning is a coupling-suppressed $\mu^2 \sim \alpha_0^{16} M_P^2$, that is

$$\frac{v}{M_P} = \frac{\alpha_0^8}{\sqrt{\lambda}}. \tag{1}$$

The power is the byte. The exponent 8 is the $[8,4,4]$ cell. The same per-bit non-unitary walk projection $\alpha_0 = 1/137$ that supplies the fine-structure constant (the single-link $\mathcal{W}$-trace) acts on each of the eight cell bits. Two combinations of that one object appear: the electromagnetic self-energy is the *incoherent sum* over bits (one power of α_0), while the electroweak vacuum is the *filled matter cell*—R4's complete generation, all eight modes occupied—whose formation amplitude is the all-eight *coincidence* $\prod_{i=1}^{8} \alpha_0 = \alpha_0^8$. An explicit channel evolution reproduces the binomial ladder $\sum_k \binom{8}{k} \alpha_0^k$ and fits the all-eight power to 8.000, robust to the unital shift; the selection of the eight-fold (rather than one-fold) channel follows because partial fillings are orthogonal states—the one-fold channel is the self-energy, a different observable. The naive alternative, reading the Higgs as the *operator* expectation of the R4 projector, is the hierarchy disaster (the projector's identity term dominates, $\mu^2 \sim M_P^2$); the state-transition reading is what evades it.

The quartic is forced too, and confirmed. The same dimensionlessness forbids a bare quartic, so $\lambda(M_P) = 0$ and $\lambda(v)$ is the pure renormalisation-group output of running up from zero at the Planck scale. This is precisely the Standard Model's celebrated *near-criticality*: integrating the one-loop standard RG from the measured $\lambda(\text{EW}) = 0.129$ gives $\lambda(M_P) \simeq -0.02 \approx 0$. The framework's forced boundary condition therefore predicts that the Higgs mass is the near-critical one—a known but otherwise unexplained fact.

The number, and the tier. With $\lambda(v)$ fixed by the running, (1) is a prediction: $\alpha_0^8/\sqrt{0.13} = 2.2 \times 10^{-17}$ against the observed 2.0×10^{-17}, i.e. $\sim 10\%$. This is a *proposition*, not a closed theorem, and the residual is named rather than diffuse: (i) the identification of the electroweak condensate with the filled-cell transition; (ii) the per-mode amplitude equal to α_0; (iii) the precise $\lambda(v)$ at two–three loops (the metastability debate sets the sign of $\lambda(M_P)$); and (iv) an uncomputed $O(1)$ Coleman–Weinberg prefactor. Within those premises the second anchor is no longer an unexplained input but a $\sim 10\%$-predicted radiative scale—the byte power times a forced near-critical quartic.

The vector-boson Z-map is narrower than the Higgs map. The old geometric 2/9 Weinberg-angle claim remains retired: the UV charge-trace value is 3/8, with ordinary SM running to the weak scale. The current 2/9 statement is instead post-EWSB and pole-specific. In the massive W/Z external-leg ledger the endpoint-exposure operator has squared residue weights hypercharge:weak

= 2 : 7, i.e. the LSZ pole quotient is 2/9 inside the rank-9 massive endpoint space. Separately, $\alpha(M_Z)$ is not a new input; it is the framework's dressed $\alpha(0)$ run to M_Z through standard vacuum polarisation. With the present V-map scale this gives W/Z pole masses at few-per-mille to percent grade. The residual is fixed-scheme multi-loop pole/RGE matching, not an unbuilt endpoint quotient.

10 Status Table

Topic	Status	Current reading
Matter content	LOCKED finite bookkeeping	$[8,4,4]$ byte dictionary gives $4 \times 2 \times 2 = 16$ fermion slots and $48 = 3 \times 16$, including passive ν_R; not claimed to beat $SO(10)$ on parsimony.
Axis booking	COMPUTED / closed	Spatial rank, colour triplet, generation triplet, and R4 repair ratio are distinct categories.
Charge/anomaly arithmetic	LOCKED	Exact finite identities; recovered $\sum Q^3 = -2/9$ chiral charge-cube statement.
Universal 2/9 banner	RETIRED	Weinberg leg retired to 3/8 plus RG; Koide leg is a phase projection, not exact rational.
Strong CP bare route	LOCKED bare / CONDITIONAL continuum	$\arg\det(M_u M_d) = 0$ at the finite-register level; continuum lift reduces to the field-wide non-abelian $SU(2)_L$ chiral measure (abelian $U(1)_Y$ closed; reduced-proxy mirror-gap "collapse" retired by the non-reduced spin-network build), so $d_n \sim 10^{-31}$ is conditional, not negative-leaning.
CP holonomy	CONDITIONAL sign-carrier localized	K_3 Majorana support is the admissible $\Delta L = 2$ portal class; the CP-odd sign carrier is the closed oriented R1 boundary cochain, made scalar by the global orientation line ω. The two-cover Hasse path instead carries the CP-even δ covariance. Thus the lepton CP sign is correlated with the same handedness used by the CKM sign if a nonzero $\Delta L = 2$ portal exists. Remaining gates: derive that portal and the phase magnitude. The faithful C_3 value $2\pi/3$ gives $\sin(3\sigma\Phi) = 0$, so it is not the baryogenesis phase.
PMNS angles	CONDITIONAL / sharp prediction	The current texture is a conditional QLC/mean-cycle construction, not an unconditional PMNS theorem. It predicts a near-maximal second-octant atmospheric angle, $\theta_{23} \simeq 45.9^\circ$, because the $K_{\rm or}/3$ sign component makes the atmospheric latch first-order stationary. Residuals are the four-sector twist table and the leading $\delta \to$angle map.

Topic	Status	Current reading
SMG/mirror gap	Conditional continuum reduction	Local and higher-cutoff evidence positive; registered Hamiltonian gives $P_{\text{vac}} H P_{\text{ch}} = 0$ and induced $\beta_A = \lambda_i = 0$, so the framework-specific problem is reduced to the pure Wilson axis W_0 plus a spectator mirror block. A local-defect bound certifies the electric-subtracted mirror offset on $\beta_{\text{cert}} \leq \beta < 1$; raw charged-sector positivity reaches $\beta \simeq 1.46$ but is not the continuum observable. The $\beta \geq 1$ mirror bound rigorously *reduces* to SU(3) confinement ($\sigma > 0$): $\Delta_{\text{raw}} \geq$ DSMG+(static-charge gauge energy), DSMG intensive, so the gap's volume-uniformity $\iff$ the pure-gauge ground is chargeless $\iff$ confinement. So both residuals — W_0 (no bulk transition) and confinement — are standard pure-SU(3) lattice-gauge facts, *not* framework-specific; the lift is closed modulo standard SU(3) LGT, inheriting lattice QCD's rigor (physically certain; mathematically conditional on the universal QCD gaps: bulk analyticity + the SU(3) mass gap).
Nucleon anchor	Computed / load-bearing	$M_N = 2\sqrt{2}\,\Lambda_{\text{QCD}}$ remains a central static-baryon relation.
Koide tau	Conditional	Structured high-precision consistency, not exact rational derivation.
Heavy-quark masses	Computed no-go / Open frontier	Target-B dense-alphabet matches are explicitly non-predictive: m_c/Λ, m_b/Λ, m_t/Λ, and v/Λ all have many low-complexity competitors. Heavy-flavour spectroscopy remains useful only after the second anchor is supplied—which is itself now a $\sim$10% prediction (§9).
Electroweak scale v	Conditional proposition	$v/M_P = \alpha_0^8/\sqrt{\lambda}$ (§9). Dimensionlessness forces a radiative scale; the power 8 is the byte (the filled-cell all-eight coincidence of the per-bit α_0, occupation-orthogonality selecting the eight-fold channel over the one-fold self-energy); $\lambda(M_P) = 0$ is forced and $\lambda(v)$ is the RG output, confirmed by Standard-Model near-criticality. Predicts v/M_P to $\sim$10%; residual is the filled-cell-transition premise, the per-mode amplitude, the 2–3-loop $\lambda(v)$, and an $O(1)$ prefactor. Reduces the second anchor from an irreducible input to a predicted radiative scale.
W/Z pole masses	Conditional precision map	The UV Weinberg-angle 2/9 claim is retired; UV charge trace gives 3/8. The live 2/9 is a post-EWSB endpoint-exposure quotient in the W/Z LSZ pole ledger, combined with $\alpha(M_Z)$ obtained by running dressed $\alpha(0)$. Residual: fixed-scheme pole matching and V-map precision.

Topic	Status	Current reading
Velocity unification	CONDITIONAL IR-closed; high-E = trans-Λ_{QCD}	Scalar-hop bubble exactly zero; Dirac-triple one-loop passes under named premises. *IR lift done*: the velocity anisotropy is an *irrelevant* $(a_0 k)^2$ operator and the loop's ln-running is the isotropic common-mode c-renormalisation, so the $k \to 0$ continuum is Lorentz-invariant (the physical photon is the Gauss-projected isotropic mode, T-R2; the K6 directional "divergence" is a gauge artifact). The high-E end ($\delta v/v \sim O(1)$ at $E \sim \Lambda_{\text{QCD}}$) is the known trans-$\Lambda_{\text{QCD}}$ photon problem: single-mode protection is a kinematic no-go, the finer lattice is CC-excluded, and high-E LI lives in a collinear null bundle of topologically-protected (Chern -1) chiral records — *exactly* null only in the soft/IR limit (the edge metric is non-universal/curved, so a finite-constituent residual $\sim (\varepsilon/\Lambda_{\text{QCD}})^2$ remains), bridged to Quantum Darwinism.

11 Open Problems

The matter/gauge paper leaves a shorter and sharper list of frontiers than the historical corpus did.

1. **Chiral continuum theorem.** Construct the full TCH Ginsparg–Wilson/domain-wall or equivalent chiral operator and prove that the finite anomaly bookkeeping survives the continuum limit.

2. **Mirror-gap continuum theorem — reduced to confinement.** The electric-subtracted Δ_{mir} for $\beta \geq 1$ is rigorously reduced to SU(3) confinement (positive static-charge energy): the local-defect bound certifies $\beta_{\text{cert}} \leq \beta < 1$, and the residual $\beta \geq 1$ window is the weak-coupling confinement frontier — a standard pure-SU(3) lattice-gauge fact (the SU(3) mass-gap problem), *not* a framework-specific gap. What remains is to *import* confinement and the W_0 no-bulk-transition on the same footing as lattice QCD itself: both numerically certain, both mathematically the universal QCD rigor gaps (`smg_chiral_continuum_standard_lgt_synthesis.py`).

3. $\Delta L = 2$ **CP-holonomy portal — residual localized.** The current result is sharper but more conditional than the previous snapshot. The support is A_{K_3}-type Majorana support; the CP sign carrier is the closed oriented R1 boundary cochain, made scalar as $\omega\Omega_{R1}$ by the global orientation line. This separates it from the two-cover R1 Hasse path, which carries the CP-even δ covariance. If a nonzero $\Delta L = 2$ lepton portal exists, its sign is no longer arbitrary and would be correlated with the CKM handedness. The live residuals are still the portal's existence and the phase magnitude. In particular $2\pi/3$ is a CP zero in $\sin(3\sigma\Phi)$, not a baryogenesis phase.

4. **PMNS coherent frame-branch existence.** Prove that lepton-sector recovery actually has a coherent one-port polar branch before reset. If such a branch exists, covariance and mean-cycle billing already force $U_{\text{frame}} = \exp(\delta K_{\text{or}}/3)$; what remains is the physical existence of the branch, distinct from both the Hermitian iK_{or} CP pointer and an ordinary dissipative reset/jump channel.

5. **Operator map for spectroscopy.** Derive the static/dynamic defect boundary from one substrate operator rather than sector-wise rules.

6. **Target-B mass scale.** The electroweak scale v is now landed with a unique formula $v/M_P = \alpha_0^8/\sqrt{\lambda}$ (§9), a $\sim$10% prediction modulo named premises and the first continuous mass scale carrying a unique, non-dense-alphabet form. What remains is to *discharge* those premises (the filled-cell-transition identification and the per-mode amplitude $= \alpha_0$) and to land a *heavy-quark* mass or mixing parameter, which stay in Target-B.

7. **Velocity RG — IR lift done; the high-E end is the known trans-$\Lambda_{\rm QCD}$ problem.** The finite one-loop is promoted to the continuum (`item115_velocity_rg_continuum_lift.py`): the velocity anisotropy is an *irrelevant* $(a_0k)^2$ operator (the matter loop's ln-running is the isotropic common-mode c-renormalisation; the directional split is pure lattice, no ln-running), so the $k \to 0$ limit is exactly Lorentz-invariant (passive emergent LI), with the physical photon the Gauss-projected isotropic mode (T-R2; the K6 directional "divergence" is a gauge artifact). The high-E end ($\delta v/v \sim O(1)$ at $E \sim \Lambda_{\rm QCD}$) is *not* a new tension: it is the documented trans-$\Lambda_{\rm QCD}$ photon problem, where canon already proved (a) a kinematic no-go for any trans-$\Lambda_{\rm QCD}$ single lattice mode ($E_1 = 2\Lambda\sqrt{1-\beta^2}$, null⇒soft), (b) the finer-lattice fix is cosmological-constant–excluded, and (c) the only CC-compatible representation is a collinear *null bundle* ($P^2 = 0$ exactly), modulo an open bundle-irreducibility theorem. So the dispersion cannot be "protected" beyond $(a_0k)^2$ for a single mode; exact high-E LI lives in the null bundle. The bundle-as-one-event theorem reduces to three halves: momentum ($P^2 = 0$) and information (QD redundancy = 1 mode-DOF) are shown (`foundations_null_bundle_qd_irreducibility.py`); the energy-deposit half is resolved *operationally* (`foundations_bundle_energy_deposit_coherence.py`) — single-vertex absorption is the wrong target (the detector is also sub-$\Lambda_{\rm QCD}$; high-E photons are measured by calorimetric summation over a collinear shower = one event), and causal-*null* records (door iii) evade *both* the GRB time-of-flight dispersion ($v = c$, no spread) *and* the CC cost (on-shell null rays, not vacuum zero-point modes — answering the standing "door iii is CC-costly" objection). The energy↔record map is then itself *derived* (`foundations_energy_record_map_qsl.py`): the quantum speed limit applied to $\mathcal{W}$'s quasi-energy gives a record density $\rho_{\rm rec} = E/(\pi\hbar c) \propto E$ on the LR null cone, the one-record-per-site cap reproduces the BZ one-mode ceiling $\pi\Lambda_{\rm QCD}$, and the bundle multiplicity $N(1\,{\rm TeV}) = 959$ matches canon's independently-derived BZ-edge count to the integer (a QSL⊥kinematics cross-check). Finally the carrier *constituents* are identified (`foundations_chern_edge_carrier_probe.py`): the macroscopic SC-Wilson branch has $v_g = c$ only at $k \to 0$, but the microscopic gauge web is topological (canon §7.2: line graph L(TCH), multi-band Chern $C_{S_7} = -1$, Dirac touchings at C_{4v}), so it carries *chiral*, ∼linear, near-null modes *off* the SC-Bloch branch — the records — whose chirality *topologically forces* the bundle's collinearity. These are sub-cutoff: there is no single such mode at GeV, so the GeV photon is irreducibly the coherent bundle of topological chiral modes (an *emergent*, not fundamental, photon) — operationally adequate. Probing the edge *metric* (`foundations_ltch_edge_velocity_probe.py`) then settles the last gap with an honest *negative*: topology protects the mode's existence and chirality but NOT its metric — under a C-preserving anisotropy the edge velocity is non-universal (it tracks the deformation while C stays -1), the dispersion is curved (it turns over, $E \propto \sin k$), and the bulk-Dirac slope $\neq$ the SC $k \to 0$ velocity. So "exactly c + exactly linear" is generically *false*; the earlier "exactly null bundle" is corrected to "exact only in the soft/IR limit", with residual LV $\sim (\varepsilon_{\rm const}/\Lambda_{\rm QCD})^2$ (the same irrelevant $(a_0k)^2$ scale, now at the constituent level), so GRB-safety is approximate (soft-suppressed) and a mild two-velocity is possible absent a boundary symmetry. What stands: existence + chirality + collinearity (topological). An attempt at the exact slab (`foundations_ltch_phasing_reconstruction.py`) hits an honest blocker: the cuboctahedral cell is faithful and canon's phased Γ-spectrum has the exact closed form $\{2\sqrt{2}, \sqrt{6}-\sqrt{2}, -\sqrt{6}-\sqrt{2}\}$ (roots of $x^3 - 12x + 8\sqrt{2}$), but the $\pi/4$ phasing that carries $C_{S_7} = -1$ (without which there is *no* chiral edge mode) is a specific T_d pattern — *not* a uniform flux (any flux breaks the three T_d 3-folds canon preserves) — tied to the 3D edge orientation that is not recorded. So the exact chiral-edge slab needs the original §7.2 operator or a full-TCH reconstruction; the model-independent metric verdict above is unchanged by it. Deeper structure was then recovered (`foundations_ltch_cell_and_phasing_family.py`): the cell is $L(Q_3)$ (the line graph of the 12 gauge links around a gauge cell, not L(octahedron)), and the phasing is a *triangle-flux* (fingerprint: it preserves ${\rm tr}, {\rm tr}^2$ and changes only det, $16 \to -8\sqrt{2}$). But the closed-form spectrum is *necessary, not sufficient*: a triangle-flux reproduces it only asymmetrically and no clean T-symmetric one matches, so the canonical operator is gated on the Gauss/gauge connection + inter-cell structure, not the spectrum. **The connection is then DERIVED** (`foundations_gauss_connection_berry.py`): the phase a photon picks up hopping between adjacent links is the Pancharatnam–Berry phase of its polarisation parallel-transported between the link directions; around a triangle this is the solid angle of the (orthogonal) link directions ($\pi/2$, an octant), halved to $\pi/4$ for the EM director — and with the Maxwell sign ($H = d^\dagger d$, so canon = spec$(-H)$) it reproduces canon §7.2's *exact* spectrum

to machine precision (the three T_d 3-folds + the moving triple $\{2\sqrt{2}, \sqrt{6} \mp \sqrt{2}\}$). So canon's "$\pi/4$ phasing" is the geometric Berry phase of the photon, not ad-hoc, and the Chern $C_{S_7} = -1$ originates in helicity = the framework chirality χ (flipping χ flips every triangle flux). The inter-cell assembly (`foundations_intercell_assembly_layers.py`) is then a *two-layer* structure: each gauge link is shared by 4 gauge cells ($\Rightarrow$ 3 links/primitive cell), so the *macro* photon (§7.3) is the plain SC scalar Wilson–Maxwell field $K(k) = 6 - 2\sum \cos k_i$ (massless, isotropic to $(a_0 k)^2$, no Berry flux), while the chiral topology lives in the *micro* 12-band complex (§7.2) that carries the Berry flux — the flux is micro-only (added to a vector link model it gaps the intrinsic Weyl monopole and would mass the macro photon). **Geometric correction** (`foundations_tch_linegraph_literal.py`): building the line graph literally from the t{4,3,4} coordinates shows the actual L(TCH) is *15-band, 8-regular* with Γ-spectrum $\{8, 2^5, -2^9\}$ (degree-5 vertices — the square-pyramid vertex figure), *not* the 4-regular 12-band cuboctahedron. The cuboctahedron is $L(Q_3)$ of the 8 gauge cells around a matter cell — a per-cell CLUSTER (the 12-eigenvalue spectrum), mislabelled "L(TCH)"; as a crystal its 4-fold-shared links give 3 bands (the macro photon). So the chiral-edge question is posed on a *literal* crystal — the 15-band L(TCH) (with the Berry connection) or the 3-band macro photon — not on a 12-band cuboctahedron crystal; the cuboctahedron + its Chern is the cluster idealisation. **GEOMETRY CORRECTION** (`foundations_oblate_bipyramid_substrate.py`): the substrate matter cell is the *oblate square bipyramid* (§1.2/§1.3, "Q_3 cell"; §1.3 forbids "regular octahedron"), and three orthogonal bipyramids tile each cube ($a^3/3$ each, one shape, 3 orientations). On THIS geometry the face-adjacency of the bipyramid's 8 [8,4,4] faces is the cube graph Q_3, and $L(Q_3)$ = the cuboctahedron = the §7.2 photon cell — so §7.2 is *vindicated* on the correct geometry. The preceding paragraph's "literal L(TCH) is 15-band/trivial" used item113's *regular*-octahedra-+-truncated-cubes geometry (§1.3-forbidden, inconsistent with §1.2) and is **retracted**; the chiral-photon Chern is OPEN on the cuboctahedron-cluster crystal of the oblate-bipyramid substrate, not refuted. **Computed** (`foundations_bipyramid_photon_crystal.py`): the photon crystal on the oblate-bipyramid tiling (sites = bipyramid edges; adjacency = face-sharing cuboctahedron) gave, with the *overlap*-amplitude connection, a lowest-3 group with Chern = −1. **SELF-CORRECTION 2026-06-25** (`foundations_photon_chern_connection_audit.py`): that $C = -1$ is an *artifact of the overlap convention* — the hop amplitude $\langle \varepsilon | \varepsilon \rangle$ has non-uniform magnitude (0.21–0.79), which is not a U(1) gauge field. Under the physically-correct *pure-phase geometric* (Peierls/Berry) connection the corrected-geometry crystal is $C = 0$ (trivial). **RESOLVED 2026-06-25** (`foundations_chiral_photon_dynamical_flux.py`): the geometric phase is zero for a clean reason — a triangular plaquette's three edge-directions are *coplanar* ($d_3 = d_2 - d_1$), enclosing zero solid angle — so canon's $\pm\pi/4$ is *not* a solid-angle holonomy but a **dynamical coin / C_{4v} turn-rule phase** (T-breaking, from $\mathcal{W} = \mathcal{SC}$). Imposing that pure-phase, uniform-magnitude (genuine U(1)) chiral flux ϕ on the triangular plaquettes gives the lowest-3 group $C = \pm 1$ on *every* k_z plane, robustly for $\phi \in [\sim 0.9, \pi/2]$ (no Weyl), with the sign flipping with the chirality (= the two helicities). So **the chiral photon is REAL** on the corrected geometry — its chirality dynamical, not geometric. This supersedes both the overlap-magnitude $C = -1$ (artifact) and the geometric pure-phase $C = 0$. **Coin-flux closure** (`foundations_coin_flux_value.py`): per-edge $\theta \to$ per-triangle $\Phi = 3\theta$; robust window $\Phi \in [0.39\pi, 0.51\pi]$ (centred $\pi/2$); canon's $\pm\pi/4$/edge ($\Phi = 3\pi/4$) is the retired octagon (C_8) value and OVERSHOOTS, so it does not carry over; the C_{4v} vertex stars give the helicity coin flux $\Phi = 2\pi/4 = \pi/2$ per triangle ($\theta = \pi/6$ per edge), INSIDE the window $\to$ robust $C = \pm 1$. *Also STANDS*: the cuboctahedron = $L(Q_3)$ = the dual cell (the photon graph), the unphased spectrum, and the macro photon's emergent Lorentz invariance. (Canon clean-up done: §1.2's "apex-apex = equatorial diagonal / three bipyramids share the centre" corrected to "= edge / bond-centred tiling" — the centre-shared version fails 50%/25%; the truncated-cube layer of §7.3/item113 is flagged as the dual gauge web, not a second matter cell.)

A Reproducibility Starter Table

Script	Sector	Purpose
python_code/axis_booking_resolution.py	finite bookkeeping	Verifies the category split between spatial, colour, generation, and repair-axis uses of "3".
python_code/item086_universal_2_9_recovery.py	anomaly	Rebuilds the one-generation charge ledger and verifies the recovered $-2/9$ chiral charge-cube identity.
python_code/anomaly_cobordism_class.py	anomaly	Checks anomaly/cobordism classification data used by the finite register audit.
python_code/luscher_abelian_precondition.py	chiral lift	Tests the abelian preconditions for a Luescher/Ginsparg–Wilson style continuum route.
python_code/gw_dtch_construction.py	chiral lift	Constructs candidate TCH Ginsparg–Wilson data.
python_code/smg_phase_decisive.py	SMG	Tests finite symmetric-mass-generation phase structure.
python_code/tch_nonreduced_spin_network_scaling.py	gauge mirror	Runs non-reduced spin-network scaling checks with explicit link indices.
python_code/tch_2plaq_extended_hop.py	gauge mirror	Tests extended plaquette basis and hopping structure.
python_code/smg_no_bulk_transition_theorem.py	continuum frontier	Encodes the present no-bulk-transition/RG theorem target.
python_code/smg_rg_basin_reduction.py	continuum frontier	Shows that the registered SMG Hamiltonian induces no adjoint or irrelevant pure-gauge deformation.
python_code/smg_escalation_theorem.py	continuum frontier	Exact electric-leg cutoff/volume stability; Schur-bounded cutoff drift with next-shell emission; cluster-expansion volume certificate.
python_code/smg_electric_subtracted_gap_observable_theorem.py	continuum frontier	Proves that the continuum mirror-pole observable is $\Delta_{\mathrm{raw}} - E^{\min}_{\mathrm{string}}$, not the raw charged-sector gap; the smoke row has raw gap $C_3/\beta > 0$ but zero mirror offset.
python_code/smg_tch_continuum_theorem_gate.py	continuum frontier	Reduces the SMG-specific continuum obstruction to the pure-Wilson no-bulk input plus a uniform electric-subtracted mirror-gap bound.
python_code/smg_volume_uniform_mirror_bound.py	continuum frontier	Reduces the volume-uniform mirror-offset problem to static-charge confinement; certifies $\Delta_{\mathrm{mir}} > 0$ on $\beta_{\mathrm{cert}} \leq \beta < 1$ and separates this from raw charged-gap positivity to $\beta \simeq 1.46$.
python_code/smg_pure_wilson_axis_import_ledger.py	continuum frontier	Isolates the W_0 pure-Wilson-axis premise and records which part is internal projection, external lattice-gauge evidence, or still unproved.
python_code/smg_chiral_continuum_standard_lgt_synthesis.py	continuum frontier	Synthesis: both residuals (W_0 no-bulk-transition + the mirror gap = confinement) are standard pure-SU(3) LGT, so the chiral lift is closed modulo standard lattice QCD — physically certain, mathematically at the shared bulk-analyticity + mass-gap frontier.
python_code/item87_majorana_cp_operator.py	CP holonomy	Identifies the complex-symmetric Majorana M_R as the missing CP operator class.
python_code/item87_deltaL2_holonomy_coupling.py	CP holonomy	Tests the generation-blind $\Delta L = 2$ holonomy portal and its orientation sign.
python_code/item87_recovery_holonomy_cp_portal_audit.py	CP holonomy	Shows an oriented recovery latch conditionally forces A_{K_3} Majorana support while leaving phase and orientation open.
python_code/item87_directed_latch_phase_orientation_closure_audit.py	CP holonomy	Shows current canon cannot make K_{or} service-readable or select σ; $\Phi = 2\pi/3$ closes only under a new faithful-C_3-character latch premise.
python_code/item87_kor_r1_orientation_hunt.py	CP holonomy	Identifies K_{or} as the oriented R1-boundary cochain while showing that existing S_3-blind source algebra averages it away.
python_code/item87_sterile_r1_boundary_readout_theorem.py	CP holonomy	Proves the oriented R1 boundary is readable iff the sterile recovery environment carries a sign-representation pointer; a scalar source reads only A_{K_3}.
python_code/r15_global_orientation_sign_pointer_theorem.py	CP holonomy	Identifies the sign-representation pointer as the global substrate orientation line ω (ωK_{or} an S_3 scalar), needing no new generation-resolved port.

Script	Sector	Purpose
`python_code/item87_r15_orientation_unification.py`	CP holonomy	Verifies K_{or} is the S_3 sign rep and unifies the orientation sign with the CKM walk-phase $\mathbb{Z}_2$, pinned to $s = +1$ by the observed Jarlskog $J > 0$; leptonic CP is then correlated.
`python_code/item87_r15_residue_closure.py`	CP holonomy	Earlier closure attempt; now superseded for baryogenesis by the R1 boundary-cochain audit because $2\pi/3$ is CP-zero in $\sin(3\sigma\Phi)$.
`python_code/item87_cp_holonomy_sector_closure.py`	CP holonomy	Earlier end-to-end chain localizing the CP sector; now read conditionally rather than as a sign lock.
`python_code/item87_cp_portal_residue_localization.py`	CP holonomy	Localizes support and pointer; the live residual remains the off-diagonal sterile-Majorana portal plus phase magnitude.
`python_code/item87_r1_orientation_leptogenesis_bridge.py`	CP holonomy	Separates the CP-even R1 Hasse-path covariance from the CP-odd closed R1 boundary cochain and shows the baryogenesis sign is conditionally tied to global orientation for nonzero Φ.
`python_code/item87_pmns_angle_promotion_gate.py`	PMNS angles	Shows K_{or} is a CP-orientation carrier, not an angle texture; observed/QLC PMNS matrices require a standard-representation residual.
`python_code/item87_pmns_standard_texture_derivation_attempt.py`	PMNS angles	Derives the 2/9 standard-shear reactor scale and proves that a pure standard texture cannot also keep θ_{23} maximal; full QLC needs a sign/standard bridge.
`python_code/item87_pmns_sign_standard_bridge_gate.py`	PMNS angles	Shows that $K_{\text{or}}/3$ as a real mean-cycle sign current plus the 2/9 standard shear and atmospheric latch uniquely gives the QLC tangent; the eigenvector-lift remains conditional.
`python_code/item87_pmns_eigenvector_lift_theorem_audit.py`	PMNS angles	Refutes closure from the Hermitian CP pointer or an ordinary reset/jump; reduces the lift to a coherent pre-reset frame-transport plus mean-cycle normalisation lemma.
`python_code/item87_pmns_coherent_polar_factor_attempt.py`	PMNS angles	Proves conditional uniqueness of $U_{\text{frame}} = \exp(\delta K_{\text{or}}/3)$ from a coherent one-port pre-reset branch, S_3 covariance, and mean-cycle billing; branch existence then supplied by `item87_frame_transport_lemma_closure.py`.
`python_code/item87_frame_transport_lemma_closure.py`	PMNS angles	Closes branch *existence* from R0: a record-preserving recovery is a coherent correction (apply), not a measure-and-reset collapse, so the pre-reset frame branch exists; with the uniqueness above the frame-transport lemma closes, the per-tick step being the lepton-sector traffic clock.
`python_code/item87_standard_shear_2_9_origin.py`	PMNS angles	Identifies the 2/9 standard-shear strength as the ν_R-defect Berry phase $\delta = d/N = 2/9$, the d=2 sibling of the neutrino mean-cycle $3/9 = 1/3$ on the N=9 plaquette; an honest reduction (four-sector twist table stays conjectural, $\delta \to$angle map leading order).
`python_code/item87_theta23_octant.py`	PMNS angles	Shows the QLC-tangent texture predicts the *second* θ_{23} octant ($\simeq 45.9^\circ$, robust under $\pm 15\%$ perturbation): the latch holds 45° at first order, the second-order residual pushes it above; leading-order (undershoots the observed $\simeq 49^\circ$).
`python_code/item87_theta23_latch_structural_audit.py`	PMNS angles	Shows the $\theta_{23} = 45.9^\circ$ undershoot is *structural* (the $K_{\text{or}}/3$-forced atmospheric latch makes it second-order) and closes both first-order routes to 49° — latch-break needs sign$\neq K_{\text{or}}/3$; the CP sum-rule is blocked by Jarlskog= 0; reframes 45.9° as a sharp near-maximal prediction (octant data unresolved).
`python_code/koide_tau_mass.py`	spectroscopy	Evaluates the tau/Koide mass relation under current conventions.
`python_code/heavy_quark_targetb_audit.py`	spectroscopy	Quantifies dense-alphabet competition for c, b, t and v; converts heavy-quark Target-B from unattempted to a no-go for prediction without an external scale.
`python_code/item55_higgs_z_finite_invariant_barrier.py`	electroweak scale	Shows that finite A_{1g}, strain-orbit, Casimir, and one-loop service invariants do not close m_H/m_Z; standard EW mixing remains cross-sector.

Script	Sector	Purpose
`python_code/v_phase1_forcing.py`	electroweak scale	Forces v to be a radiative (Coleman–Weinberg) transmutation scale by elimination of the tree options.
`python_code/v_phase2_perbit_alpha_derivation.py`	electroweak scale	Shows the per-bit α_0 is the item-79 single-link $\mathcal{W}$-projection; the EW vacuum is its all-eight product (α_0^8) versus the one-bit-sum self-energy.
`python_code/v_phase2_walk_condensate.py`	electroweak scale	Explicit non-unital $\mathcal{W}$ channel evolution of the 8-cell; reproduces the binomial ladder and fits the all-eight coincidence power to 8.000, robust to the unital shift.
`python_code/v_phase2_higgs_operator_id.py`	electroweak scale	Shows the literal R4-projector operator reading is the hierarchy disaster (identity term dominates); the filled-cell state transition gives α_0^8 by occupation-orthogonality.
`python_code/v_phase3_quartic.py`	electroweak scale	Forces $\lambda(M_P) = 0$; the 1-loop SM RG from $\lambda(\mathrm{EW}) = 0.129$ runs to $\lambda(M_P) \simeq -0.02$ (near-criticality); v/M_P to $\sim$10%.
`python_code/ew_alpha_mz_from_framework_dressed_alpha.py`	electroweak Z-map	Runs dressed $\alpha(0)$ to $\alpha(M_Z)$ with standard vacuum polarization; removes $\alpha(M_Z)$ as an independent input.
`python_code/v_program_wz_pole_exposure_operator.py`	electroweak Z-map	Supplies the post-EWSB W/Z LSZ endpoint-exposure quotient 2/9 while keeping the UV charge-trace value 3/8.
`python_code/v_program_wz_zmap_lock_audit.py`	electroweak Z-map	Audits the W/Z pole-mass map; residual is fixed-scheme pole/RGE matching and the V-map precision.
`python_code/item113_t1_t2_local_map_theorems.py`	nuclear contact maps	Grounds the volume/surface sign structure in local TCH/QEC contacts while leaving the absolute many-body residual scale unpinned.
`python_code/baryon_parameter_count.py`	baryons	Audits baryon-sector parameter counts and target classifications.
`python_code/item115_dirac_triple_vacuum_polarization.py`	velocity	Computes the Dirac-triple vacuum-polarisation kernel.
`python_code/item115_velocity_rg_continuum_lift.py`	velocity	The RG lift: the velocity anisotropy is an irrelevant $(a_0 k)^2$ operator and the marginal flow isotropic, so the continuum is Lorentz-invariant (IR lift done); the high-E end $\delta v/v \sim O(1)$ at $E \sim \Lambda_{\mathrm{QCD}}$ is the known trans-Λ_{QCD} photon problem (not a new tension).
`python_code/foundations_null_bundle_qd_irreducibility.py`	velocity	Single-mode dispersion protection is a closed no-go; exact high-E LI lives in a collinear null bundle ($P^2 = 0$ exact); Quantum-Darwinism redundancy collapses the N records to one mode-DOF (information half of the open bundle-irreducibility theorem).
`python_code/foundations_bundle_energy_deposit_coherence.py`	velocity	The energy-deposit half: single-vertex absorption is the wrong target (the detector is also sub-Λ_{QCD}; high-E photons are measured by calorimetric summation over a collinear shower); causal-null records (door iii) evade both the GRB dispersion ($v = c$) and the CC cost (on-shell null rays, not vacuum zero-point); the theorem closes modulo the door-(iii) carrier existence.
`python_code/foundations_energy_record_map_qsl.py`	velocity	Derives the energy↔record map: the quantum speed limit on $\mathcal{W}$'s quasi-energy gives record density $\rho_{\mathrm{rec}} = E/(\pi\hbar c) \propto E$ on the LR null cone; the one-record-per-site cap reproduces the BZ ceiling $\pi\Lambda_{\mathrm{QCD}}$ and the bundle count $N(1\,\mathrm{TeV}) = 959$ (matches canon). Residual narrows to carrier coherence.
`python_code/foundations_chern_edge_carrier_probe.py`	velocity	Identifies the carrier constituents: the SC-Bloch branch has no finite-energy null mode, but the microscopic Chern(-1) line graph $L(\mathrm{TCH})$ (canon §7.2) carries chiral, $\sim$linear, near-null modes off that branch (the records); chirality forces the bundle's collinearity. Sub-cutoff cap: no single GeV mode, so the GeV photon is the emergent coherent bundle.

Script	Sector	Purpose
`python_code/foundations_ltch_edge_velocity_probe.py`	velocity	The edge *metric*, honest negative: topology protects existence + chirality but not the metric — the edge velocity is non-universal (tracks a C-preserving anisotropy), the dispersion is curved, and the Dirac slope $\neq$ the SC velocity. So "exactly c + linear" is generically false; "exactly null bundle" $\rightarrow$ "exact only in the soft/IR limit", residual $\sim (\varepsilon/\Lambda)^2$. Cell faithful (cuboctahedron = canon §7.2).
`python_code/foundations_ltch_phasing_reconstruction.py`	velocity	Exact-slab attempt (honest blocker): the cuboctahedral cell is faithful and canon §7.2's phased spectrum = $\{2\sqrt{2}, \sqrt{6}-\sqrt{2}, -\sqrt{6}-\sqrt{2}\}$ (roots of $x^3 - 12x + 8\sqrt{2}$, new closed form), but the $C_{S_7} = -1$ $\pi/4$ phasing (required for any chiral edge mode) is a specific T_d pattern, not a uniform flux, and is not recorded — so the exact chiral-edge slab needs the canonical operator or a full-TCH build.
`python_code/foundations_ltch_cell_and_phasing_family.py`	velocity	Deeper structure: the 12-band cell is $L(Q_3)$ (line graph of the 12 gauge links around a gauge cell), and the phasing is a triangle-flux (preserves $\mathrm{tr}, \mathrm{tr}^2$; changes only det, $16 \rightarrow -8\sqrt{2}$). The closed-form spectrum is necessary-not-sufficient (a triangle-flux matches only asymmetrically; no clean T-symmetric one), so the canonical operator is gated on the gauge connection, not the spectrum.
`python_code/foundations_gauss_connection_berry.py`	velocity	DERIVES the connection: the photon's hop-phase between links is the Pancharatnam–Berry phase of its polarisation = the solid angle of the link directions ($\pi/2$ octant, $\pi/4$ for the EM director). With the Maxwell sign it reproduces canon §7.2's exact spectrum (3-folds + $\{2\sqrt{2}, \sqrt{6} \mp \sqrt{2}\}$) to machine precision; Chern = helicity = χ. Canon's "$\pi/4$ phasing" is thus the photon geometric phase, not ad-hoc. Inter-cell δ remains for the slab.
`python_code/foundations_intercell_assembly_layers.py`	velocity	The inter-cell assembly is two-layer: each gauge link is shared by 4 gauge cells, so the macro photon (§7.3) is plain SC scalar Maxwell $K(k) = 6 - 2\sum \cos k_i$ (massless, isotropic to $(a_0 k)^2$, no Berry flux), while the chiral topology lives in the micro 12-band (§7.2) carrying the flux. The Berry flux is micro-only (it gaps the link web's intrinsic Weyl monopole and would mass the macro photon). Exact micro δ (truncated-cube edges) is the remaining gap.
`python_code/foundations_tch_linegraph_literal.py`	velocity	Literal geometry: the actual line graph of t{4,3,4} (degree-5 vertices, square-pyramid vertex figure) is *15-band, 8-regular*, Γ-spectrum $\{8, 2^5, -2^9\}$ — NOT the 12-band cuboctahedron. The cuboctahedron is $L(Q_3)$ of the 8 gauge cells around a matter cell (a per-cell cluster, mislabelled "L(TCH)"); as a crystal it is 3-band (macro photon). Corrects the 12-band-crystal premise; the chiral-edge question is on the 15-band L(TCH) or the 3-band macro.
`python_code/foundations_tch_literal_edge_dispersion.py`	velocity	The literal edge dispersion on item113's geometry: the 15-band L(TCH) is trivial (Chern 0). **RETRACTED** — it used the wrong (§1.3-forbidden regular-octahedron) geometry; see the next row.
`python_code/foundations_oblate_bipyramid_substrate.py`	velocity	GEOMETRY CORRECTION: the matter cell is the oblate square bipyramid (§1.2/§1.3, "Q_3 cell"), 3 per cube tiling space (one shape, 3 orientations); its 8 [8,4,4] faces have face-adjacency = cube graph Q_3, and $L(Q_3)$ = the cuboctahedron = the §7.2 photon cell — §7.2 VINDICATED on the correct geometry. The 15-band "trivial" result above (item113 geometry) is void; the chiral-photon Chern is OPEN on the oblate-bipyramid cuboctahedron crystal.

Script	Sector	Purpose
`python_code/foundations_bipyramid_photon_crystal.py`	velocity	The photon crystal on the (correct) oblate-bipyramid tiling: §1.2 geometry cleaned up (bond-centred tiling verified; centre-shared fails 50%/25%). The 11-band crystal gave Chern $= -1$ (lowest-3) *under the overlap-amplitude convention* — but this is **SUPERSEDED** (see next row): $C = -1$ is an overlap-magnitude artifact; the pure-phase (gauge) connection gives $C = 0$. The geometry + cuboctahedron cell stand; the chiral photon does not.
`python_code/foundations_photon_chern_connection_audit.py`	velocity	CONNECTION AUDIT / self-correction: the photon $C = -1$ is an artifact of using the full overlap $\langle\varepsilon\|\varepsilon\rangle$ (non-uniform magnitude 0.21–0.79) as the hop amplitude — not a U(1) gauge field. Under the pure-phase *geometric* (Peierls/Berry) connection the corrected-geometry crystal is $C = 0$. [Superseded by the resolution below.] STANDS: the geometry, the cuboctahedron $= L(Q_3) =$ dual cell (the photon graph), and the macro photon's emergent Lorentz invariance ($K(k)$ isotropic to $O(k^2)$).
`python_code/foundations_chiral_photon_dynamical_flux.py`	velocity	RESOLUTION: the chiral photon is REAL with $C = \pm 1$ under a pure-phase *dynamical* chiral connection. The geometric Berry phase of a triangular plaquette is zero (3 coplanar edge-directions, $d_3 = d_2 - d_1$), so canon's $\pm\pi/4$ is a coin / C_{4v} turn-rule phase (T-breaking), not a solid-angle holonomy. A uniform-magnitude (U(1)) chiral flux ϕ on the triangular plaquettes gives the lowest-3 group $C = \pm 1$ on every k_z plane, robust for $\phi \in [0.9, \pi/2]$ (no Weyl), sign = helicity. Supersedes both the overlap $C = -1$ (artifact) and the geometric $C = 0$. Open: the exact coin flux from $\mathcal{W} = \mathcal{SC}$.
`python_code/foundations_coin_flux_value.py`	velocity	Coin-flux closure: per-edge coin phase $\theta \to$ per-triangle holonomy $\Phi = 3\theta$; precise robust $C = \pm 1$ window $\Phi \in [0.39\pi, 0.51\pi]$ (centred $\pi/2$). Canon's $\pm\pi/4$/edge ($\Phi = 3\pi/4 \approx 2.36$) is the retired octagon (C_8, $2\pi/8$) value and OVERSHOOTS (and is inconsistent with the "C_{4v}" label §7.1 cites). The C_{4v} vertex stars (4 edges/star) give helicity coin flux $2\pi/4 = \pi/2$ per triangle ($\pi/6$ per edge), INSIDE the window $\to$ robust $C = \pm 1$. Exact coin flux $= \pi/2$/triangle, not the octagon $\pi/4$.
`python_code/foundations_coin_operator_explicit.py`	velocity	Explicit coin operator: canon's coin $\mathcal{C}$ (zero-controlled CNOT, §3.1) generates the Dirac matrices (§3.5); its Dirac spin operator $S_3 = \frac{i}{4}[\gamma^1, \gamma^2]$ has eigenvalues $\pm\frac{1}{2}$, so a C_4 ($\pi/2$) turn gives $\pm\pi/4$ for spin-$\frac{1}{2}$ matter (DERIVES canon's $\pm\pi/4$, the fermion value) and $\pm\pi/2$ for the spin-1 photon (twice the fermion). The photon coin phase is thus $\pi/2$ (not $\pi/4$); per-triangle flux $\pi/2$ $\to C = \pm 1$, helicity-locked. canon's $\pi/4$ is the matter value, correct for matter, mis-applied to the photon.
`python_code/foundations_band_1112_correspondence.py`	velocity	11↔12 band correspondence: the per-Q_3 cuboctahedron (12 = 8 slant + 4 equator, $L(Q_3)$, canon §7.2) and the physical per-cube edge-identified crystal (11 = 8 slant + 3 equator) are the SAME photon in different unit cells. The count change 12→11 is the equator identification (4→3): the cuboctahedron 0^3 triplet $\to 0^2$ (one mode lost); 2^3 survives; $-2^5 \to -2^3{+}2$ deep modes (slant-star connectivity). Both carry the same chiral photon ($C = \pm 1$); no clean lowest-3 ↔ lower-7 identity – the band-group is embedding-dependent.
`python_code/item115_loop_pi_and_sign.py`	velocity	Verifies that the scalar-hop bubble is exactly zero and isolates the Dirac-class successor.
`python_code/item102_k6_overlap.py`	velocity	Computes the K_6 mode-overlap and orientation map.

Script	Sector	Purpose
`python_code/item102_dynamic_polarization_gate.py`	velocity	Runs the on-shell dynamic polarization gate for anisotropic velocity flow.

References

[1] Stephen L. Adler. Axial-vector vertex in spinor electrodynamics. *Phys. Rev.*, 177:2426–2438, 1969.

[2] John S. Bell and Roman Jackiw. A PCAC puzzle: $\pi^0 \to \gamma\gamma$ in the σ-model. *Nuovo Cim. A*, 60:47–61, 1969.

[3] Gyan Bhanot and Michael Creutz. Variant actions and phase structure in lattice gauge theory. *Phys. Rev. D*, 24:3212–3217, 1981. doi: 10.1103/PhysRevD.24.3212.

[4] T. Blum, C. DeTar, Urs M. Heller, Leo Karkkainen, K. Rummukainen, and D. Toussaint. Thermal phase transition in mixed action SU(3) lattice gauge theory and Wilson fermion thermodynamics. *Nucl. Phys. B*, 442:301–316, 1995. doi: 10.1016/0550-3213(95)00137-9.

[5] Alvaro De Rújula, Howard Georgi, and Sheldon L. Glashow. Hadron masses in a gauge theory. *Phys. Rev. D*, 12:147–162, 1975.

[6] Estia Eichten and John Preskill. Chiral gauge theories on the lattice. *Nucl. Phys. B*, 268: 179–208, 1986.

[7] David Elliman. It-from-bit model: Reproducibility repository. `https://github.com/dgedge/itfrombit`, 2026. Code and reproducibility repository for the canon snapshot series.

[8] Harald Fritzsch and Peter Minkowski. Unified interactions of leptons and hadrons. *Ann. Phys.*, 93:193–266, 1975.

[9] Howard Georgi and Sheldon L. Glashow. Unity of all elementary-particle forces. *Phys. Rev. Lett.*, 32:438–441, 1974.

[10] Paul H. Ginsparg and Kenneth G. Wilson. A remnant of chiral symmetry on the lattice. *Phys. Rev. D*, 25:2649–2657, 1982.

[11] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. 28th ACM Symposium on Theory of Computing*, pages 212–219, 1996.

[12] Martin Hasenbusch and Silvia Necco. SU(3) lattice gauge theory with a mixed fundamental and adjoint plaquette action: lattice artefacts. *JHEP*, 08:005, 2004.

[13] Urs M. Heller. SU(3) lattice gauge theory in the fundamental–adjoint plane and scaling along the Wilson axis. *Phys. Lett. B*, 362:123–127, 1995.

[14] David B. Kaplan. A method for simulating chiral fermions on the lattice. *Phys. Lett. B*, 288: 342–347, 1992.

[15] John Kogut and Leonard Susskind. Hamiltonian formulation of Wilson's lattice gauge theories. *Phys. Rev. D*, 11:395–408, 1975.

[16] Yoshio Koide. Fermion-boson two-body model of quarks and leptons and Cabibbo mixing. *Lett. Nuovo Cim.*, 34:201–205, 1982.

[17] Martin Lüscher. Abelian chiral gauge theories on the lattice with exact gauge invariance. *Nucl. Phys. B*, 549:295–334, 1999.

[18] S. Navas et al. Review of particle physics. *Phys. Rev. D*, 110:030001, 2024. Particle Data Group.

[19] Herbert Neuberger. Exactly massless quarks on the lattice. *Phys. Lett. B*, 417:141–144, 1998.

[20] Holger B. Nielsen and Masao Ninomiya. Absence of neutrinos on a lattice. (I) proof by homotopy theory. *Nucl. Phys. B*, 185:20–40, 1981.

[21] Roberto D. Peccei and Helen R. Quinn. CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38:1440–1443, 1977.

[22] Yigal Shamir. Chiral fermions from lattice boundaries. *Nucl. Phys. B*, 406:90–106, 1993.

[23] Gerard 't Hooft. Symmetry breaking through Bell–Jackiw anomalies. *Phys. Rev. Lett.*, 37: 8–11, 1976.

[24] Juven Wang and Yi-Zhuang You. Symmetric mass generation. *Symmetry*, 14(7):1475, 2022.

[25] Alexander Weiße, Gerhard Wellein, Andreas Alvermann, and Holger Fehske. The kernel polynomial method. *Rev. Mod. Phys.*, 78:275–306, 2006.

[26] Kenneth G. Wilson. Confinement of quarks. *Phys. Rev. D*, 10:2445–2459, 1974.