Why Matter Has Mass and Space Does Not

The price of closure in VERSF: background-subtracted fold energy, holonomy mass spectra, and the constructed spatial Higgs bridge

Keith Taylor

VERSF Theoretical Physics Programme

August 2026

General Reader Summary

In this framework, empty space and matter are made of the same underlying threads of relation. That raises an awkward question: if the threads themselves carry energy, why does not everything have as much mass as everything else? The answer developed here is that a particle's mass is not the energy of its threads. It is the extra cost of tying those threads into a loop that cannot be smoothed away. When the same number of threads is compared loose and tied, the cost of the threads themselves cancels exactly, and only the cost of the knot remains.

The note then finds two encouraging facts about that knot cost. First, a loop that closes without any twist costs nothing, so mere closure is not enough: only a genuine twist that cannot be undone produces a positive cost. Second, shrinking a twisted loop raises its cost in exactly the inverse-square way that the programme's existing mass ladder already uses, arrived at there by a completely different route.

The Higgs field enters as the common breathing of the medium. A calculation on the finite model shows that pushing every thread equally is cheap and locally silent: the machinery that records events does not notice it at first order. What it can notice is a difference between neighbouring pushes. A symmetry argument turns this from an accident into a rule: a uniform local push simply has no handle on the shaped part of space, so the first physically allowed coupling is a difference between neighbours, the discrete version of a gradient.

That coupling is then built explicitly. Using only attachment rules the programme already contains, the neighbouring-difference coupling comes out in one line, and in essentially only one way: it reaches every shaped direction of space and behaves at long distances exactly like the ordinary spreading term a physical field needs, with no number put in by hand. What has not been done is to convert its natural units into physical ones: that conversion is the remaining, separately stated task.

Contents

Status of this note. Exploratory but source-grounded; not a completed particle-mass derivation. The spatial Higgs mechanism stands at a conditional structural pass: the gradient injection is constructed and produces an ordinary two-derivative term with no inserted coupling, while its physical same-source normalisation remains open. Every statement carries an explicit grade: exact, imported, conditional, or open.

Core result

ΔE_fold baseline cancels; $\Delta E\_closure \propto \lambda\_min^{+}(L\_U)$

$\lambda\_\min^+(L\_U) \to \Theta\_d^2 / L^2$ in the smooth-loop limit, with Θ_d the reduced holonomy angle

$I\_H\text{^}can = (I + S^{-1})/\sqrt{2}$, $I\_H\text{^}can\ d = (S - S^{-1})/\sqrt{2}$, $K\_H\text{^}grad(\theta) = 2 \sin^2\theta = O(k^2)$

This provides a concrete route by which the same underlying folds can constitute ordinary space and, when forced into a non-trivial closed configuration, acquire a positive localized rest-energy increment.

Abstract

The question tested is whether VERSF can consistently treat folds as energetically non-trivial even when they form the unbounded relational structure identified with space, while still allowing particle rest mass to arise from a finite excess associated with closure rather than from the mere existence of folds.

The answer from the first controlled calculation is yes. At matched fold count, the universal per-fold contribution cancels identically between a candidate particle closure and its background configuration, leaving a relational and topological term. A covariant graph-Laplacian functional, built from the same-source covariant difference already present in the PFDMI source, makes that term exact: an open fold chain admits a zero-residual parallel transport for every choice of edge unitaries, while a closed loop has zero residual if and only if the total holonomy fixes a state. For unitary holonomy with eigenphases θ_j on an n-fold cycle, the full closure spectrum is the union of shifted Fourier branches $4 \sin^2((2\pi k - \theta\_j)/(2n))$, so the lowest closure gap is $\gamma\_n = \min\_j 4 \sin^2(d(\theta\_j, 2\pi\mathbb{Z})/(2n))$; the phase-holonomy case is the single-branch specialisation. In the smooth-loop metric limit the gap scales as Θ_d^2/L^2, independently matching the Role-4 sector-rescaling law $M\_g = \bar{E}\_g/L\_g^2$.

The recovered finite master source passes a radial separation test: uniform substrate-amplitude breathing has zero six-defect cost and stiffness 2 from the radial potential alone, while the K = 7 hub-versus-rim mode has stiffness 11/2 = 2 + 7/2, the excess supplied entirely by the relational current defect. The naive identification ρL = constant fails for the raw regulator amplitude, but the canonically normalised radial field φ carries inverse-length dimension, and at leading local order $\Gamma\_C \propto H^\dagger H$ together with $\Gamma\_C \propto L^{-2}$ yields $\varphi/v = L_0/L$ on a fixed topology and holonomy branch, with the Higgs quartic re-expressed as the squared return cost of the closure spectral defect.

The primitive first-order audit is exact on the declared instruments: the EX2D native source-to-Kraus Jacobian and the six-defect Jacobian both annihilate the physical common Higgs line, $J\_K\ q\_H = 0$ and $J\_D\ q\_H = 0$, while the later finite nonlinear completion instrument responds to the local Higgs amplitude with a machine-verified nonzero derivative. A symmetry theorem explains the first-order nulls structurally: the PSAC closure carrier splits into a uniform mode and a nonuniform six-plane Q_6 preserved by the source-returned circulant closure operators, and for any symmetry action transitive on the seven cells the invariant subspace is exactly the uniform line, so every equivariant local linear map from a scalar into Q_6 vanishes. An audit note records the precise boundary of this argument: wheel symmetry alone, with the hub distinguished, leaves a second invariant direction inside Q_6, so the no-go is carried by the transitive circulant structure,

not by closure geometry as such. The first symmetry-allowed spatial coupling is therefore a neighbouring Higgs difference, carried by the source-native injection $I\_H: (h\_j - h\_i) \to Q_6$.

That injection is then constructed. The PFDMI-EX2G local incidence attachment, unique within its declared admissible class and here proved unique within that class by a two-line counting argument, assigns each directed difference one arrival and one departure incidence; taking the reversal-even scalar pair coordinate yields the canonical injection $I\_H\text{\textasciicircum}can = (I + S^{-1})/\sqrt{2}$ on the transitive Z_7 carrier. Composed with the source-native difference $d = S - I$ it gives the centred derivative $(S - S^{-1})/\sqrt{2}$, so the whitened spatial Higgs kernel is $K\_H\text{\textasciicircum}grad(\theta) = 2 \sin^2\theta = 2\theta^2 - (2/3)\theta^4 + \cdots$, an ordinary two-derivative $O(k^2)$ term, not $O(k^4)$. The injection is full rank on Q_6 with singular values $\sqrt{2}\,|\cos(\pi k/7)|$, all strictly positive because seven is odd, its kernel after composition is exactly the uniform mode required by the symmetry theorem, and the PSAC isometry adds no gain. The construction is rigid: within the endpoint-local equivariant class, reflection-oddness of the composed derivative and unit pair normalisation fix I_H^can uniquely up to a global sign, so no continuous freedom survives. Its kernel operator equals one half of a covariant cycle Laplacian in relabelled coordinates, the same operator family that defines the closure mass burden, an exact structural unification at the whitened level. A refinement-parity result sharpens the regulator obligation: on even cell counts the alternating mode is annihilated exactly, so a doubling regulator step must be shown harmless before full-rank stability can be claimed. The result is graded as a conditional structural pass: the spatial Higgs mechanism exists on the canonical transitive source branch and generates the $k^2|h|^2$ structure without an inserted coupling. The remaining debts are the primitive selection of the declared attachment class and the physical same-source normalisation converting the whitened graph coefficient 2 into the physical Z_H^space; the whitened coefficient is not promoted.

---

1. The Question and the Correction to the Earlier Picture

A first intuition was that a particle might have mass because it contains folds and each fold carries energy. That is incomplete in VERSF because folds are not exclusive to particles: folds are also the deeper structures from which spatial relations emerge. The relevant observable cannot therefore be the raw energy of a collection of folds; it must be an energy difference between two states of the same underlying fold content.

$E\_particle \neq E(folds)$

$m\_C\, c^2 = E[\text{folds organised as closure } C] - E[\text{same fold content in background } B\_C]$

This is a background-subtraction principle. It is analogous in spirit to measuring a defect energy relative to the ordered medium in which the defect lives. It avoids counting the energetic substrate itself as localized particle mass.

Definition 1 (Background-subtraction principle). For any candidate particle closure C, define a matched background B_C containing the same fold number and the same universal per-fold contribution. The particle mass candidate is the excess structural cost $\Delta E_C = E[C] - E[B_C]$. Universal fold energy is allowed, but it cannot contribute to a localized mass unless the closure changes the fold count or changes the relational or curvature energy.

2. Existing VERSF Ingredients

The calculation uses only structures already present in the VERSF programme, plus one clearly labelled trial identification for the simplest spinorial test. Three ingredients are needed, and it is worth saying in advance what job each performs. The particle ontology supplies the object being costed: a particle as a stable closed arrangement of folds. The PFDMI covariant source supplies the measuring instrument: a way of asking, edge by edge, how badly neighbouring states fail to agree once transport between them is taken into account. And the Role-4 mass scaling supplies the target: an inverse-square dependence on localization scale that any candidate microscopic mass mechanism ought to reproduce for independent reasons. The logic of the note is to cost the object with the instrument and check the result against the target.

2.1 Particles as stable fold closures

The particle ontology paper defines matter particles as stable, localized, transportable closure defects in the committed distinguishability graph Σ. Its earlier leading-order stabilization functional is

$$S[C] = \alpha\, n_C + \beta\, \kappa_C$$

where n_C is a fold count and κ_C is a discrete graph-geometric closure quantity. The same paper explicitly allows κ_C to be implemented using boundary-cycle structure, branching complexity or rewiring resistance. The four-fold toy loop has $n_C = 4$ and $\kappa_C = 1$, yielding $E_C = 4\alpha + \beta$.

2.2 The later PFDMI covariant source

PFDMI-EX2E introduces representation-valued endpoint states ψ_v, matrix-valued edge transport $U_e(A)$, and the same-source covariant current residual

$$D_{(J,e)} = j_e - \left(\psi_{(t(e))} - U_e(A)\, \psi_{(s(e))} \right).$$

At the frozen background the residual vanishes on every edge. This supplies a natural local measure of relational mismatch without inventing a new geometric object.

2.3 Role-4 mass scaling

The later lepton-sector variational programme derives a sector-rescaling relation of the form

$$M_g = \bar{E}_g / L_g^2,$$

with L_g the localization scale and $\bar{E}_g$ a dimensionless rescaled eigenvalue. Under the derived exponential localization law $L_g = L_0 e^{-\kappa g}$, equal rescaled energies generate an adjacent mass ratio $e^{2\kappa}$. This inverse-square dependence is therefore already an independently established structural target within the VERSF mass programme.

3. First Test: Fixed-Fold Background Subtraction

Take four committed folds. Compare an open chain with a closed directed cycle, while keeping the fold count fixed.

Open: $f_1 \to f_2 \to f_3 \to f_4$

Closed: $f_1 \to f_2 \to f_3 \to f_4 \to f_1$

Lemma 1 (Fixed-fold cancellation). Under the two-term stabilization model $S[C] = \alpha n_C + \beta \kappa_C$, any comparison at matched fold count n_C eliminates the universal term exactly: $\Delta E = E[C] - E[B_C] = \beta(\kappa_C - \kappa_B)$. In particular, for the four-fold open chain ($\kappa = 0$) against the minimal closed cycle ($\kappa = 1$), $E_{open} = 4\alpha$, $E_{closed} = 4\alpha + \beta$, and $\Delta E = \beta$.
Proof. Immediate: the α-term depends only on n_C, which is equal on both sides by construction of B_C. ■

This establishes the basic cancellation, and it is worth pausing on what it does and does not say. The energy of the folds themselves need not vanish; nothing here requires the substrate to be free. What the lemma shows is that any such universal energy is invisible to a properly matched comparison, in the same way that a canvas contributes the same underlying material whether or not a picture has been painted on it. The localized excess, the thing a particle detector could notice, comes entirely from how the folds are organized. The question then becomes what "organized" should mean mathematically, and the placeholder answer, a bare cycle count, is immediately seen to be too crude.

Scope of Lemma 1. It proves that a finite closure excess is algebraically compatible with energetic background folds. It does not derive β. More importantly, a pure cycle-count definition would assign the same β to every simple one-cycle closure, so cycle count alone cannot encode a realistic particle spectrum. The remainder of the note replaces $\beta\kappa_C$ with a spectral object.

---

4. Replacing the Placeholder Curvature with a Covariant Spectral Cost

The stronger move is to build the closure cost directly from the PFDMI covariant difference. For a graph G with edge transport U_e and endpoint states ψ_v, define the dimensionless quadratic functional

$$Q_G[\psi] = \Sigma_{(e\in G)} \| \psi_{(t(e))} - U_e \psi_{(s(e))} \|^2.$$

This may be written as a quadratic form $Q_G = \langle\psi, L_U \psi\rangle$, where L_U is the covariant graph Laplacian associated with the link transport. L_U is positive semidefinite by construction.

The intuition behind Q_G is simple. Each edge asks one question: if the state at one end is carried across the edge by the transport rule, does it match the state at the other end? Q_G

adds up the squared failures. A configuration with $Q_G = 0$ is one in which every local comparison is satisfied simultaneously, a network in perfect internal agreement. The two theorems below say when such perfect agreement is possible: always, on an open structure; and on a closed loop, only when the accumulated transport around the loop returns some state exactly to itself.

4.1 Open background

Theorem 1 (Open background zero mode). On any tree or open chain, $\min_\psi Q_G[\psi] = 0$ for every assignment of edge unitaries, and the minimiser is unique up to the choice of the root state. **Proof.** Fix a root vertex and any ψ_{root}. A tree admits a unique path from the root to each vertex; define each state recursively along edges by $\psi_{t(e)} = U_e \psi_{s(e)}$. Every edge term then vanishes identically. Uniqueness up to root choice follows because each edge equation determines the child state from the parent state. ■

4.2 Closed loop

Closing the final edge introduces a global consistency condition. Let the total holonomy around an n-edge loop be

$$H = U_n U_{n-1} \cdots U_1.$$

Theorem 2 (Closure criterion). For the n-cycle with edge unitaries $U_1,\dots,U_n$, $\min_\psi Q_{loop}[\psi] = 0$ if and only if the holonomy H has a fixed vector, $H\psi = \psi$ for some $\psi \neq 0$. If H has no +1 eigenvalue in the relevant closure sector, the open-chain zero mode is lifted and $\min Q_{loop} > 0$. **Proof.** Zero cost forces every edge equation $\psi_{j+1} = U_j \psi_j$, hence $\psi_1 = H\psi_1$ after one circuit; conversely a fixed vector generates a zero-cost assignment by the recursion of Theorem 1 applied around the cycle. Positivity in the absence of a fixed vector follows because Q is a continuous nonnegative form on the compact unit sphere of the finite-dimensional state space, attaining its minimum, which cannot be zero. ■

Key structural result. The particle-like excess is not the cost of having folds. It is the spectral cost of satisfying global closure when local parallel transport cannot be made globally trivial. This gives an exact mechanism by which ordinary background fold relations remain zero-cost after subtraction while a non-trivial closure carries a positive residual.

The physical picture behind the pair of theorems deserves one more sentence. An open chain never has to agree with itself: states can be defined outward from any starting point, each edge satisfied as it is reached, and the process never returns to check consistency. A loop does return, and the holonomy H is precisely the record of everything the transport did on the way round. If that record fixes some state, the loop closes for free; if not, some edge somewhere must absorb a mismatch, and the minimum possible total mismatch is the closure gap. Mass, on this proposal, is the energy of that unavoidable disagreement.

5. Exact Cycle Spectrum

It is not enough to know that the closure gap is positive; the mass functional needs its exact value and its dependence on the loop data. For a cycle this can be computed in closed form, and the answer depends on strikingly little: not on how the transport is distributed among the edges, only on the total accumulated twist. The single-phase case comes first because it displays the whole mechanism in one formula; the general matrix case then reduces to it exactly.

5.1 Phase holonomy

Theorem 3 (Exact phase-cycle spectrum). Let the total holonomy be a phase, $H = e^{i\Theta}$, distributed in any way over an n-edge cycle. The spectrum of the covariant cycle Laplacian is exactly $\lambda_k(\Theta) = 4 \sin^2\left((2\pi k - \Theta)/(2n) \right)$, $k = 0, 1, \ldots, n-1$, and depends only on Θ, not on the distribution of edge phases. The lowest eigenvalue is the closure gap $\gamma_n(\Theta) = 4 \sin^2\left(\Theta_d / (2n) \right)$, $\Theta_d \coloneqq d(\Theta, 2\pi\mathbb{Z})$, the shortest angular distance from Θ to an integer multiple of 2π. **Proof.** A vertex-wise gauge transformation $\psi_j \to g_j \psi_j$ with unimodular g_j maps the edge phases to any other distribution with the same product, so the spectrum depends only on Θ; choose the uniform distribution $\theta_e = \Theta/n$. The resulting operator is a circulant, diagonalised by Fourier modes $\psi_j \propto e^{2\pi i k j/n}$, with eigenvalue $2 - 2 \cos(2\pi k/n - \Theta/n) = 4 \sin^2((2\pi k - \Theta)/(2n))$. Minimising over k selects the branch nearest Θ and yields the reduced distance Θ_d. ■

Immediate consequences:

Trivial holonomy: $\Theta \equiv 0 \pmod{2\pi}$ gives $\gamma_n = 0$. A mere closed graph is not enough to generate a mass gap.

Non-trivial holonomy: any $\Theta \not\equiv 0 \pmod{2\pi}$ gives $\gamma_n > 0$.

Gauge invariance: redistributing the edge phases changes the link representatives, not the total holonomy or the spectrum. This is part of the theorem statement, not an additional assumption.

Different admissible holonomy sectors can therefore carry different dimensionless closure costs even at fixed fold count.

Corollary 1 (Monotonicity of the closure gap). $\gamma_n(\Theta) = 4 \sin^2(\Theta_d/(2n))$ is strictly increasing in the reduced angle Θ_d on $[0, \pi]$ and strictly decreasing in the fold count n at fixed Θ_d. **Proof.** $\sin^2$ is strictly increasing on $[0, \pi/2]$ and $\Theta_d/(2n) \in [0, \pi/(2n)] \subset [0, \pi/2]$; the argument decreases strictly in n. ■ The physical reading: at fixed metric scale the dimensionless burden is set by the twist and only weakened by finer folding, so mass is controlled by holonomy and localization scale, not by fold multiplicity. This is the spectral form of the background-subtraction principle: adding folds without adding twist adds no localized cost.

5.2 General unitary holonomy

The phase case extends exactly to arbitrary matrix transport.

Theorem 4 (Eigenphase decomposition of the closure spectrum). Let $U_1,\dots,U_n$ be arbitrary d×d unitaries on an n-cycle with holonomy $H = U_n \cdots U_1$, and let $e^{(i\theta_1)}, \dots, e^{(i\theta_d)}$ be the eigenvalues of H. Then the nd eigenvalues of the covariant cycle Laplacian are exactly the union over j of the phase-cycle branches $4 \sin^2((2\pi k - \theta_j)/(2n))$, $k = 0,\dots,n-1$, and the closure gap is $\gamma_n(H) = \min_j 4 \sin^2(d(\theta_j, 2\pi\mathbb{Z}) / (2n))$. In particular min Q = 0 if and only if some $\theta_j \equiv 0 \pmod{2\pi}$, recovering Theorem 2. **Proof.** Gauge-transform all transport onto a single edge, so the cycle carries identities except one edge carrying H; diagonalise H by a constant unitary change of basis, which commutes with the identity edges. The operator block-diagonalises into d independent phase cycles with total phases θ_j, and Theorem 3 applies branchwise. ■

Theorem 4 upgrades the closure gap from a toy single-phase statement to an exact functional of the holonomy spectrum: the closure cost of a particle candidate is controlled by the eigenphase of its transport holonomy nearest to triviality. It also makes precise which holonomy data the mass functional of Section 10 can depend on: eigenphases only, never edge representatives.

6. Spinorial Trial Sector

The distinctive signature of a fermion is that turning it through a full rotation does not return it to itself but to minus itself; only a double rotation restores it. The particle-ontology paper proposes that fermionic closure classes transform spinorially in exactly this sense, so that a 2π rotation changes the sign of the closure representative and a 4π rotation restores it. That result concerns rotational double-cover structure; it is not yet a theorem that the same sign becomes the transport holonomy H around the microscopic closure. The following is therefore a test identification, not a completed derivation.

Trial identification: $H = -1 \Leftrightarrow \Theta = \pi$.

Corollary 2 (Involutive rigidity). If the double-cover sign enters transport as an involutive holonomy, $H^2 = 1$, then every eigenphase satisfies $e^{(2i\theta)} = 1$, so $\theta \in \{0, \pi\}$. The unique nontrivial involutive phase is π. **Proof.** Immediate from the spectral mapping. ■ The conditional content of the trial identification is therefore only that the spinorial sign enters transport at all; once it does, the angle carries no continuous freedom and $\Theta = \pi$ is forced.

For the existing four-fold toy cycle, Theorem 3 gives the exact gap

$$\gamma_4(\pi) = 4 \sin^2(\pi/8) = 2 - \sqrt{2} = 0.5857864376\ldots$$

Thus a four-fold spinorial-holonomy closure has a definite positive dimensionless spectral excess relative to the matched open background.

Scientific grading. Exact: the cycle spectrum and $\gamma_4(\pi)$ once $\Theta = \pi$ is specified. Conditional: identifying the microscopic closure holonomy with the spinorial sign. Open: deriving the allowed holonomy directly from the primitive VERSF action.

7. Continuum Metric Limit and the Inverse-Square Law

The gap function of Section 5 is a statement about an abstract graph. To connect it to physics, the graph must be given a size: each relational step corresponds to some emergent metric spacing, and the loop as a whole to a physical length. Once that is done, the natural question is how the closure cost behaves when the same twist is confined to a smaller and smaller loop. The answer is the note's second main result: the cost grows as the inverse square of the loop size, which is exactly the scaling the lepton-mass programme derived by a completely independent variational route.

Let a be the emergent metric spacing associated with neighbouring relation steps and let the loop length be $L = na$. The metric covariant Laplacian carries the usual $1/a^2$ factor, so the closure gap becomes

$$\Gamma_n(\Theta) = (4/a^2) \sin^2(\Theta_d / (2n)) = [4n^2 \sin^2(\Theta_d / (2n))] / L^2.$$

Proposition 1 (Smooth-loop limit). For fixed reduced angle Θ_d and $n \to \infty$ at fixed L, $\Gamma_n \to \Theta_d^2 / L^2$, with monotone convergence from below and relative error of order $\Theta_d^2/(12 n^2)$. **Proof.** $4n^2 \sin^2(x/(2n)) = x^2 - x^4/(12n^2) + O(n^{-4})$ at $x = \Theta_d$. ■

Independent convergence. This is the same inverse-square dependence that appears independently in the Role-4 sector-rescaling theorem, $M_g = \bar{E}_g/L_g^2$. The present route gives a concrete spectral reason for such a dependence: shrinking a non-trivial closed transport loop raises its lowest closure eigenvalue as $1/L^2$.

If the VERSF localization law is $L_g = L_0 e^{(-\kappa g)}$, then for self-similar closures in the same dimensionless holonomy sector

$$\Gamma_g \propto 1/L_g^2 = e^{(2\kappa g)}/L_0^2, \quad \Gamma_{(g+1)}/\Gamma_g = e^{(2\kappa)}.$$

This reproduces the structural form of the existing electron-to-muon hierarchy mechanism without using the lepton masses to construct the spectral law.

8. Relation to the Existing Role-4 Eigenproblem

The Role-4 variational paper writes the sector mass as a constrained minimum and, in exploratory dimensionless calculations, obtains negative bound-state eigenvalues for g = 0, 1, 2 and a positive unbound value for g = 3. Raw signed eigenvalues are awkward to read literally as positive rest masses.

The background-subtraction viewpoint suggests a cleaner interpretation: the zero of the continuum or background sector is a threshold, negative eigenvalues signal binding below that threshold, and the physically relevant positive stabilization magnitude is the gap to the threshold or the positive singular magnitude of the properly constructed mass operator. This is consistent with the later CLAS-1 discipline, which extracts physical masses from positive magnitude operators and singular values rather than from arbitrary signed matrix eigenvalues.

This note does not replace the Role-4 operator. It supplies a possible microscopic interpretation of why the relevant closure and localization scale should generate an inverse-square spectral cost.

9. Evidence Already Present in the PFDMI Wilson Sector

The PFDMI matrix-link construction already contains an exact analogue of the null-versus-curved split required by the present idea. Its linearized curvature map has 144 link tangents, of which 72 are pure endpoint-coboundary gauge directions and 72 survive as physical curvature directions. After quotienting gauge freedom, the physical Wilson Hessian has full rank 72 and a strictly positive minimum eigenvalue of 6 in the frozen construction.

This does not identify the number 6 with a fold energy or particle mass. Its relevance is structural: the actual VERSF source already distinguishes relational changes that are removable by gauge choice from changes that carry irreducible positive curvature cost.

pure coboundary or gauge direction → zero physical cost

non-trivial physical curvature → positive Hessian

That is precisely the architecture a background-subtracted closure mass requires: Theorems 1 and 2 are the loop-level expression of the same split.

10. Proposed Background-Subtracted Closure Mass Functional

Everything so far can now be assembled into a single proposal. The earlier schematic cost $\beta\kappa_C$ counted cycles; the calculation above shows what should replace it: the lowest positive eigenvalue of the actual covariant closure operator, evaluated after matched-background subtraction. This quantity inherits every good property proved above. It vanishes for anything the background can absorb, it is gauge invariant, by Theorem 4 it depends only on the holonomy eigenphases, and it scales inversely with the square of the localization length. What it deliberately does not carry is an absolute unit, and the reason for stopping there is stated below.

$$\Delta Q_C = \min_{\psi\in C} \langle\psi, L_U(C)\, \psi\rangle - \min_{\psi\in B_C} \langle\psi, L_U(B_C)\, \psi\rangle.$$

For a matched open or tree background, the second term is zero by Theorem 1. The dimensionless closure invariant becomes

$$q_C \coloneqq \lambda_{\min}^{+}[\, L_U(C)\,] \geq 0,$$

and by Theorem 4 it is a function of the closure holonomy eigenphases alone. A later metric-certification bridge must supply the dimensionful stiffness or energy unit. Write it provisionally as E_*; then

$$m_C c^2 = E_* \, q_C,$$

or, when the metric operator supplies an explicit length dependence,

$m_C c^2 = \mathcal{K}_* \cdot q_C / L_C^2$.

The coefficient $\mathcal{K}_*$ is deliberately left unevaluated, and this restraint is a feature rather than a gap. The current VERSF scale audit warns against silently importing a joule, $\hbar$ or an absolute ruler into the primitive fold layer: any number smuggled in at this stage would contaminate every downstream mass claim with an untracked convention. The honest division of labour is that this note derives the dimensionless topology and holonomy factor q_C and its independently generated $1/L^2$ scaling, while the conversion to physical units is owed by the history-derived metric and primitive scale programme, where it can be derived once and audited.

11. Relation to the VERSF Higgs-Radial and Electroweak Completion Work

The background-subtracted closure calculation is not a rival to the VERSF Higgs mechanism. It supplies a candidate microscopic origin for the dimensionless structural burden carried by a particle closure, while the Higgs-radial and electroweak-completion work supplies the scalar interface and dimensional scale through which that burden can become a physical mass term.

11.1 Existing Higgs-radial result: mass support as interface stiffness

The Closure-Interface Potential Necessity and Higgs-Radial Theorem (HR-1) identifies the Higgs-radial degree of freedom as the invariant normal return-stiffness mode of a stable VERSF closure interface. Its Closure-Stiffness Mass Theorem states that, after canonical normalisation, the radial mass parameter is the positive Hessian of the closure-interface potential at equilibrium:

$m_{HR}^2 = Z_H^{-1} V''_{CI}(0)$.

The physical reading is therefore already geometric: the Higgs-radial mass is not an arbitrary inserted scalar mass. It measures how strongly the closure interface resists displacement away from its equilibrium configuration. HR-1 also requires an explicit bridge from this radial stiffness to fermion, gauge-boson, spectral-gap or Yukawa mass channels; a radial scalar by itself is not yet a complete mass-generation mechanism.

11.2 What the background-subtracted closure spectrum adds

The present calculation asks a different but complementary question: once a particle closure exists, how much extra structural cost does that closed configuration carry relative to the same fold content in the unbounded background arrangement?

$\Delta E_C = E[\text{closed configuration } C] - E[\text{matched background } B_C]$.

At fixed fold count the universal per-fold contribution cancels (Lemma 1). In the source-compatible covariant formulation, the remaining dimensionless burden is the lowest positive gauge-quotiented closure eigenvalue $q_C = \lambda_{min}^+(L_U)$, a function of holonomy eigenphases only (Theorem 4). Thus the closure spectrum can encode differences between particle topologies without treating the folds constituting ordinary space as particle mass.

11.3 Complementary roles: topology, Higgs interface, and electroweak scale

The clean combined interpretation is a factorisation rather than two competing explanations of mass:

Closure topology and holonomy → dimensionless structural burden q_C or a related positive magnitude operator.

Higgs-radial closure interface → the physical scalar channel that makes closure stiffness available to mass-generating operators.

Electroweak completion scale v_{cl} → the common dimensionful scale converting the structural operator into masses.

Yukawa, branch and projector structure → the channel-specific map that assigns the common scale and structural burden to particular fermion or gauge sectors.

Schematically:

$$M_{particle} = M_{EW} \times \hat{M}_{closure}.$$

This is already close to the architecture of the Charged-Lepton Absolute Scale Theorem (CLAS-1), where a common dimensional prefactor multiplies a dimensionless localization and closure magnitude operator. The explicit charged-lepton construction has the form

$$M_{\ell}^{VERSF} = (\alpha_{int}^2 v_{cl} / 8\pi) \cdot [J_e + e^{16/3} J_\mu + (e^{32/3}/12) J_\tau].$$

The spectral calculation therefore offers a possible deeper origin for the dimensionless second factor: rather than treating its hierarchy only as an effective localization and closure operator, one can ask whether its eigenvalues descend from the background-subtracted spectrum of the actual fold closures.

11.4 Relation to the Higgs-potential completion

The Electroweak Vacuum and Higgs-Potential Completion Theorem goes beyond HR-1 by conditionally locking the closure vacuum and a numerical Higgs potential. On its declared chain it obtains v_{cl} = 243.723 GeV, $\lambda = 512/3969 \approx 0.1290$, and a leading Higgs mass m_H = 123.796 GeV, while retaining explicit premises and residual bridge debts. The same paper treats the resulting electroweak vacuum as the scale feeding the W, Z and fermion mass channels.

v_{cl} = 243.723 GeV, $\lambda = 512/3969$, m_H = 123.796 GeV [conditional chain].

These three numbers are internally consistent as printed: $\lambda = (1/2)(32/63)^2$ exactly, and $m_H = (32/63)\, v_{cl}$ = 123.796 GeV, so the chain carries the programme's $m_H/v = 32/63$ ratio unchanged. The present closure-spectrum work does not alter those claims or their grading. Its role is more microscopic: it asks what structural quantity the electroweak and Higgs interface is ultimately coupling to. The natural candidate is the excess stiffness of a non-trivial closure relative to the matched fold background.

11.5 Executed finite-cell radial stiffness test

Before the numbers, the idea. The finite master source is a small but complete working model of the theory: a handful of cells carrying substrate amplitudes, a separate Higgs radial field, auxiliary completion coordinates, and the six families of defect functionals that measure whether a configuration is internally consistent. On this model one can perform the cleanest possible version of the background-subtraction experiment: displace the configuration in a chosen direction and ask two separate questions. How much does the potential resist, and how much do the relational defects resist? If the picture developed in Parts I and II is right, a displacement that treats every cell identically should be pure potential, invisible to the relational machinery, while a displacement that treats cells differently should pay an additional relational price. That is exactly what the executed test finds.

The radial-Hessian test can be executed on the recovered finite master source. The source contains three seven-vertex substrate-density channels, one separate scalar radial field h, three completion coordinates X_H, a radial potential, and the six defect families. This permits a direct distinction between common background breathing and relational deformation.

Uniform substrate-radial mode. Vary every occupied substrate amplitude by the same normalised infinitesimal displacement u_ρ. The closure, current, continuity, record, bond and scalar-completion defects do not respond: $\|J_D u_\rho\|^2$ is zero to machine precision. The only stiffness is the potential Hessian $K_{uniform} = 2 \lambda_{sub} \rho_c^2$; in the frozen normalised cell this is exactly 2.

Non-uniform radial mode. On the natural K = 7 hub-versus-rim vector $c = (-6, 1, 1, 1, 1, 1, 1)/\sqrt{42}$, an exact eigenvector of the seven-vertex wheel graph Laplacian with $Lc = 7c$, the normalised source gives $K_{nonuniform} = 11/2$. This separates exactly as $K_{nonuniform} = 2 + 7/2$: 2 from the radial potential and 7/2 supplied entirely by the constitutive current defect D_J.

Separate Higgs-completion radial mode. If h is displaced while the three X_H completion coordinates are artificially frozen, the raw curvature is 8 in the source normalisation. Allowing the completion coordinates to relax along $\delta X_H = \delta h$ cancels the D_H and mixed completion costs, leaving the relaxed static radial curvature exactly 2. This is the Schur-relaxed quantity relevant to the physical uniform breathing mode.

The finite-cell conclusion is sharp: a common radial change is a magnitude and potential mode, while spatially structured radial deformation carries an additional relational cost. This directly realises the background-subtraction picture inside the executable master action: common background energy and localized relational excess are mathematically distinct.

The result also corrects the earlier proposed bridge test. The important question is not whether an arbitrary raw closure gap has a stationary point. The source itself supplies the stable radial potential; the spectral problem must explain how its canonically normalised radial coordinate is related to the localization scale.

Status: finite-cell structural pass. These numbers are exact for the declared normalized executable and its weights; they are not physical masses or energy scales. The physical history-derived metric, primitive scale tuple, faithful regulator and branch selection remain separate programme gates.

11.6 Raw substrate amplitude versus physical localization scale

The tempting relation ρL = constant must be typed carefully. It fails if ρ is the dimensionless microscopic saturation amplitude and L is the regulator or edge spacing. In the explicit refinement construction, newly inserted vertices remain at $\rho = \rho_c$ while each segment length is halved. Hence $\rho_c L$ is not refinement invariant.

This negative result is useful: the inverse-length law cannot be assigned to the raw substrate amplitude by notation alone. The relevant quantity is the canonically normalised physical radial field after the kinetic residue and metric bridge are included.

11.7 Canonical inverse-length bridge at leading local order

The failure of the naive relation in Section 11.6 does not kill the inverse-length idea; it locates it. The raw amplitude ρ is a bookkeeping coordinate, and nothing forces bookkeeping coordinates to carry physical dimensions sensibly. The physically meaningful object is the canonically normalised field, the coordinate in which the kinetic term has unit coefficient, because that is the field whose fluctuations propagate and whose dimensions are fixed by the action itself rather than by convention. The step below therefore rebuilds the bridge at the level where dimensions are meaningful, and the inverse-length law reappears there, now as a derived statement with explicit premises rather than an assumed one.

The reduced VERSF scalar action defines a canonical radial field φ by $d\varphi = \sqrt{Z_H(\rho)}\, d\rho$. A canonically normalised scalar in four spacetime dimensions has engineering dimension one, $[\varphi] = L^{-1}$. Therefore φ^2 has exactly the same inverse-length-squared dimension as the closure spectral stiffness Γ_C.

Proposition 2 (Conditional inverse-length bridge). Assume: (i) leading local order in the scalar effective theory; (ii) gauge invariance, so that $H^\dagger H$ is the lowest-order local dimension-two scalar; (iii) the mass-generating closure branch, on which the completion stretch vanishes at the unbroken interface, so Γ_C has no constant term at $H = 0$; (iv) a fixed closure topology and holonomy branch, so that $\Gamma_C(L)/\Gamma_C(L_0) = L_0^2/L^2$ by Proposition 1. Then $\Gamma_C(H) = B\, H^\dagger H$ with B an unknown spectral normalisation; writing $H = (0, \varphi/\sqrt{2})^T$ gives $\Gamma_C(\varphi) = (B/2)\, \varphi^2$, and at the broken vacuum $\varphi = v$, $\Gamma_C = \Gamma_0$, so B cancels in the ratio: $\Gamma_C / \Gamma_0 = \varphi^2 / v^2$. Equating the two descriptions of the same gap yields, on the positive radial branch, $\varphi / v = L_0 / L$, hence $\varphi L = v L_0$ = constant.

This is the desired inverse-length bridge, derived at the same leading-local order as the scalar effective theory rather than imposed on the raw microscopic amplitude. Higher-dimensional operators can correct the relation away from the declared order, and premise (iv) fails across topology or holonomy jumps.

11.8 Higgs potential as a closure-spectral return cost

Using $\Gamma\_C/\Gamma_0 = \varphi^2/v^2$, the departure from the equilibrium closure spectrum is

$\Delta\Gamma = \Gamma\_C - \Gamma_0 = (\Gamma_0/v^2)(\varphi^2 - v^2)$.

The leading positive quadratic return cost for this defect is

$V\_spec = (\zeta/2)(\Delta\Gamma)^2 = [\ \zeta\, \Gamma_0^2\ /\ (2v^4)\ ]\ (\varphi^2 - v^2)^2$.

Thus the standard VERSF Higgs quartic shape is the squared cost of moving the closure spectrum away from its stable value. Comparing with $V = (\lambda/4)(\varphi^2 - v^2)^2$ gives

$\lambda = 2\zeta\, (\Gamma_0/v^2)^2$.

In this reading ζ is a spectral-coordinate normalisation, not a second independent mass mechanism.

11.9 Status of the Higgs bridge

The connection stands as a conditional structural derivation rather than an assumed scaling relation. What remains conditional is the identification of the relevant W7 closure spectral stiffness with the mass-generating scalar branch, restriction to a fixed topology and holonomy sector, and the leading-local truncation. The absolute spectral coefficient and physical length and action conversion remain to be returned by the history-derived metric and primitive scale programme.

The existing closure-democracy and Killing-form result still supplies the dimensionless Higgs curvature ratio under its own premises. The present result does not independently re-predict 32/63; it supplies a microscopic interpretation of why a closure spectral stiffness can be represented by the canonical radial field and why the quartic is naturally a squared spectral-defect cost.

11.10 Primitive first-order Higgs-to-mark audit

A first-order audit asks the most elementary dynamical question available: starting from the frozen background, if the common Higgs amplitude is nudged infinitesimally, does any declared instrument register anything at linear order? A zero answer is informative in two opposite ways. Read naively, it looks like a failure, as if the theory cannot see its own Higgs field. Read correctly, it is a constraint on where the coupling can live, and the audit below establishes the correct reading: the zeros are exact, they occur at two independent layers, and they coexist with a manifestly nonzero nonlinear response. Something systematic, not accidental, is suppressing the linear channel, and Section 11.11 identifies it.

The PFDMI-EX2D instrument permits a direct check of whether the primitive marked law already responds to the physical common Higgs and completion direction. Its extended source has 334 coordinates, while the compressed native source-to-Kraus Jacobian has rank 156 = 144 + 12: 144 faithful matrix-link directions and 12 edge-associated record and export directions. The common Higgs line is $q\_H = (\varphi\_H, X\_{H0}, X\_{H1}, X\_{H2})/2$ in the same source chart.

$J\_K^{(EX2D)} q\_H = 0$.

This is an exact first-order statement for the declared EX2D instrument because its native Kraus law depends on the matrix link A and the record-export coordinate s_e, not on φ_H or X_H. It independently reproduces the earlier defect-level result $J_D q_H = 0$. The two nulls occur at distinct layers: the six-defect master source and the later native marked instrument are both blind to the common Higgs mode at first order. Neither statement is a claim of fundamental Higgs blindness; both are statements about first-order spatial and marked response on declared instruments.

That boundary matters because HCMR-XI-4 constructs a finite nonlinear completion instrument that does respond directly to the local Higgs amplitude:

$$K_r(\varphi_H) = \sin(\varphi_H) P_r A^{(n)} P_r, \partial K_r / \partial\varphi_H |_{(\varphi_H = 0)} = P_r A^{(n)} P_r \neq 0,$$

with the finite-difference derivative verified at machine precision. Local Higgs-completion response therefore exists even though the first-order spatial and closure response vanishes.

The order bookkeeping is consistent with this distinction. The faithful HM4 audit retypes bosonic completion as a second-order object. The common line $c_H = (1/2)(1, 1, 1, 1)$ is missed exactly by the six-defect block but is detected by the potential rows, and completion quadratic descent then passes the full action stack at machine precision. Local Higgs and completion dynamics is active; what is missing is its source-native spatial propagation law.

11.11 Symmetry selection rule for the spatial Higgs bridge

The tool that converts the observed zeros into a theorem is equivariance, and the concept is worth stating plainly. A map between two carriers is equivariant when it commutes with the symmetry: relabelling the cells first and then applying the map gives the same result as applying the map and then relabelling. Any coupling generated by a symmetric action must be equivariant, because the action has no way to prefer one cell labelling over another. The argument below is therefore a counting argument about symmetric objects: it asks how many directions in the closure carrier look the same from every cell, finds that under the right symmetry there is exactly one, the uniform direction, and concludes that a cell-blind scalar has exactly one place to go, which the nonuniform plane excludes by construction.

PSAC supplies the key carrier decomposition. The seven-component closure carrier separates into the uniform vector $u = (1, 1, 1, 1, 1, 1, 1)/\sqrt{7}$ and its orthogonal six-dimensional nonuniform plane Q_6:

$$\mathbb{R}^7 = \mathrm{span}\{u\} \oplus Q_6.$$

The source-returned closure operators are circulant, so they preserve this decomposition exactly, and PSAC-EX5T gives an exact field-level isometric intertwiner $W = G^{(1/2)}$ from Q_6 to the corresponding nonuniform Φ-field six-plane, with all six singular values equal to one:

$$\sigma(W) = \{1, 1, 1, 1, 1, 1\}.$$

Theorem 5 (Local scalar no-go under transitive cell symmetry). Let G be any group of permutations of the seven cells acting transitively, under which the closure carrier transforms by the permutation representation and the local common Higgs amplitude h_H transforms trivially. Then every G-equivariant linear map T: h_H → Q_6 vanishes identically: T(h_H) = 0, so the derivative ∂Φ_nonuniform / ∂h_H vanishes at the base point. **Proof.** The dimension of the invariant subspace of a permutation representation equals the number of orbits of the action. A transitive action has one orbit, so the invariant subspace is exactly span{u}. Equivariance forces T(h_H) to be invariant, hence proportional to u; but T maps into Q_6, which is orthogonal to u. Therefore T(h_H) ∈ span{u} ∩ Q_6 = {0}. ■

The circulant structure of the source-returned closure operators supplies exactly such a transitive action: a circulant commutes with the cyclic shift of the seven cells, and the cyclic group acting by shifts is transitive. The hypothesis of Theorem 5 is therefore satisfied by the PSAC carrier as returned, and the recurring first-order zero of Section 11.10 becomes a symmetry selection rule rather than an unexplained missing coefficient. A common local radial change can alter the magnitude and potential sector, but it cannot by itself choose a nonuniform spatial deformation.

Audit note (exact boundary of the no-go). The transitivity hypothesis is load-bearing and cannot be weakened to wheel symmetry alone. The automorphism group of the seven-vertex wheel fixes the hub and permutes the rim, so it acts with two orbits; the invariant subspace of the permutation representation is then two-dimensional, spanned by u and by the hub-versus-rim vector c = (−6, 1, 1, 1, 1, 1, 1)/√42, and c lies inside Q_6 (it is orthogonal to u and is precisely the nonuniform mode of the Section 11.5 stiffness test, with wheel-Laplacian eigenvalue 7). Under wheel symmetry alone, the equivariant local map T(h_H) = c h_H is allowed, and the no-go fails. The selection rule is therefore carried by the transitive circulant structure of the returned closure operators, not by closure geometry as such. Two consequences follow. First, the premise that the physical spatial closure carrier is the circulant PSAC carrier (already flagged in PSAC as conditional on the constitutive single-source identification) is now also load-bearing for the gradient selection rule and is listed as premise P5 and falsifier F3 below. Second, the hub-versus-rim direction of the finite-cell radial test is exactly the direction that a hub-distinguishing source would open to a local scalar coupling, so a future source calculation that breaks cell transitivity must revisit Section 11.12 before importing its conclusion.

11.12 The first symmetry-allowed bridge is a gradient

A neighbouring difference of Higgs amplitudes has the correct type. For adjacent physical closure cells i and j, δ_e h = h_j − h_i is an oriented edge quantity rather than a common scalar: under the transitive cell symmetry it transforms in the regular pattern of the edges, not trivially, so Theorem 5 does not apply to it. It can therefore feed the nonuniform closure carrier without violating the cell symmetry. The natural source chain is

h → dh ⟶(I_H) Q_6 ⟶(W) δΦ.

The last arrow is already fixed by PSAC at unit singular values. The only new load-bearing object is the source-native gradient-to-closure injection I_H. Once I_H is known in the physical operational metric, its norm directly fixes the spatial Higgs residue because W introduces no additional gain:

$\|\delta\Phi\|^2 = \|I\_H\ dh\|^2$.

In the long-wavelength limit dh is proportional to k h, so this construction produces the required $k^2|h|^2$ two-derivative structure. The open problem has therefore narrowed from an unspecified Higgs-to-space coupling to one explicit map: derive or refute I_H from the same multi-cell primitive action without inserting a kinetic coefficient, a measured Higgs quantity or a post-hoc spatial embedding.

11.13 Canonical construction of the gradient injection on the transitive branch

Section 11.12 reduced the spatial Higgs question to a single missing map, and it is worth restating the demand in plain terms before meeting it. What is needed is a rule, already present in the source programme rather than invented for the purpose, that takes the difference between the Higgs amplitudes of two neighbouring cells and deposits it into the shaped, nonuniform part of the closure carrier. The rule must respect the cell symmetry, it must not silently insert a number that plays the role of a coupling constant, and it must reach every nonuniform direction, since a direction it cannot reach would be a part of space the Higgs field could never propagate into. The construction below meets all three demands using an attachment rule the programme already contains, and the algebra collapses to something recognisable: the composite map is the centred discrete derivative, the most symmetric possible difference operator, and its long-wavelength behaviour is the ordinary quadratic dispersion of a propagating field.

The injection demanded by Section 11.12 can be constructed explicitly from objects the source programme already contains, with no fitted continuous gain. Three imported ingredients are used. First, the PFDMI-EX2G local incidence attachment maps a directed transition $x \to y$ to exactly one arrival and one departure incidence, $\mathcal{A}(x \to y) = C\_y^+ + C\_x^-$, reported machine-unique within its declared admissible class. Second, the PSAC intertwiner W = G^(1/2) from Q_6 to the Φ-field six-plane is an exact isometry with all six singular values equal to one. Third, the full-faithful source theorem places every zero-sum scalar-density profile on a connected carrier inside the source-visible defect row space, so the constructed image is not source-invisible.

Lemma 2 (Within-class uniqueness of the attachment). Within the declared admissible class of endpoint-local, nonnegative integer-valued attachments carrying exactly one arrival unit and one departure unit per directed transition, the attachment is necessarily $\mathcal{A}(x \to y) = C\_y^+ + C\_x^-$. **Proof.** Endpoint locality confines the image to incidences at x and y; integrality with exactly one arrival unit forces coefficient 1 on C_y^+ and 0 elsewhere among arrivals, and likewise one departure unit forces coefficient 1 on C_x^-. ■ This closes the formal classification inside the declared class; it does not prove that the class itself is uniquely selected by the deepest substrate, which is recorded as premise P8 and falsifier F8.

Work on the transitive Z_7 closure carrier of Section 11.11, with S the cyclic shift. The Higgs radial profile is $h = (h_0, \ldots, h_6)$ and its source-native neighbouring difference is $g = dh = (S - I)h$, so $g_j = h_{j+1} - h_j$. Applying the attachment to each directed difference $j \to j+1$ gives $g_j \mapsto g_j\,(C_{j+1}^+ + C_j^-)$. Because the Higgs radial closure amplitude is a scalar rather than an orientation variable, take the reversal-even normalised pair coordinate $C_j^{sc} = (C_j^+ + C_j^-)/\sqrt{2}$. The scalar closure amplitude at site j is then $q_j = (g_{j-1} + g_j)/\sqrt{2}$, so the injection is

$$I_H^{can} = (I + S^{-1})/\sqrt{2}$$

acting on the difference dh. It is a polynomial in S, hence exactly Z_7-equivariant, and it preserves the zero-sum condition, so it maps the difference carrier into Q_6 as required.

Theorem 6 (Centred-derivative composition). On the transitive Z_7 carrier, $I_H^{can}\, d = ((I + S^{-1})/\sqrt{2})(S - I) = (S - S^{-1})/\sqrt{2}$, the centred first difference. For a Fourier mode $h_j = h\, e^{i\theta j}$, $I_H^{can}\, dh = i\sqrt{2}\sin\theta \cdot h$, hence the whitened spatial kernel is $\|I_H^{can}\, dh\|^2 / \|h\|^2 = 2\sin^2\theta =: K_H^{grad}(\theta)$, and since W is an isometry, $\|W I_H^{can}\, dh\|^2 = \|I_H^{can}\, dh\|^2$. At long wavelength $K_H^{grad}(\theta) = 2\theta^2 - (2/3)\theta^4 + \cdots$, so $K_H^{grad} = O(k^2)$, not $O(k^4)$. **Proof.** The operator identity is one line of algebra; the rest is the Fourier diagonalisation of circulants and the Taylor expansion of $\sin^2\theta$. ■

Proposition 3 (Full rank on Q_6). On the six nonuniform modes k = 1, …, 6 the singular values of the edge-difference-to-Q_6 injection are $s_k = \sqrt{2}\,|\cos(\pi k/7)|$, numerically {1.274162, 1.274162, 0.881748, 0.881748, 0.314692, 0.314692}, every one strictly positive, so I_H^{can} restricted to Q_6 has full rank 6 and no nonuniform direction is accidentally lost. Because seven is odd, no nonzero mode reaches the $\cos(\pi/2) = 0$ cancellation that an even carrier could exhibit. The composed spectrum factorises consistently: $\sqrt{2}\,|\sin\theta| = |e^{i\theta} - 1| \cdot \sqrt{2}\,|\cos(\theta/2)|$ at each mode.

Two consistency checks close the loop with the symmetry analysis. The kernel of the composed map $I_H^{can}\, d$ is exactly the uniform mode: a common radial displacement produces no nonuniform excitation, which is precisely the content of Theorem 5, so the construction realises the selection rule rather than evading it. And the leading whitened coefficient 2 coincides numerically with the unit-baseline flat-transport graph coefficient reported in the confrontation programme; no identification of the two objects is claimed here.

Grading: SPATIAL HIGGS MECHANISM, CONDITIONAL STRUCTURAL PASS. What is established: a canonical source-compatible injection I_H^{can} exists on the transitive branch, is full rank on Q_6, introduces no fitted gain, is unique within the declared attachment class (Lemma 2), and its composition with the exact PSAC isometry produces $K\_H(k) \propto k^2$ (Theorem 6). What is not established: physical Z_H^{space}. The coefficient 2 is a canonical whitened graph residue, and the programme's own discipline forbids promoting whitened operational coefficients to physical residues without the physical field norm, spatial ruler and common source-selected operational metric. The primitive selection of the declared attachment class among wider weighted classes is likewise open. The mechanism question is answered conditionally yes; the normalisation question remains open.

11.14 Canonical rigidity, refinement parity, and the Laplacian identity

Three further exact results tighten the construction of Section 11.13, and each answers a question a careful reader should ask at this point. First: the construction made choices along the way, equal coefficients, a factor of $1/\sqrt{2}$, so how much of the result is choice? The rigidity theorem answers: none of it, up to an overall sign; every apparent choice is forced by symmetry, locality and normalisation. Second: the calculation lives on seven cells, so does anything break when the regulator refines the carrier? The parity proposition answers: something checkable happens, and exactly when the cell count is even. Third: is the constructed kernel a new object, or something the paper has met before? The Laplacian identity answers: it is the same covariant cycle Laplacian that has run through the paper since Part I, in light disguise.

Theorem 7 (Canonical rigidity). Among Z_7-equivariant, endpoint-local attachments, the injection has the form $I\_H = a I + b S^{-1}$. Requiring the composed operator $I\_H d$ to be reflection-odd (to transform as a first derivative under the carrier reflection $j \mapsto -j$, which conjugates S into S^{-1}) forces $a = b$, and unit normalisation of the pair coordinate, $a^2 + b^2 = 1$, then fixes $a = b = 1/\sqrt{2}$ up to a global sign. Hence $I\_H^{can} = \pm(I + S^{-1})/\sqrt{2}$ is the unique member of its class: the only surviving freedom is discrete. **Proof.** Equivariance under the cyclic shift makes the map a polynomial in S; endpoint locality restricts its support to the two endpoint incidences, giving $I\_H = a I + b S^{-1}$. With R the reflection, $R S R = S^{-1}$, so $R (I\_H d) R = a S^{-1} + (b - a) I - b S$, and oddness against $I\_H d = a S + (b - a) I - b S^{-1}$ forces $a = b$ in every power of S. Normalisation fixes the magnitude. ■ Without endpoint locality, equivariance alone leaves an independent gain on each Fourier mode pair, a multi-parameter family: locality is the load-bearing restriction, inherited from the declared attachment class (premise P8).

Proposition 4 (Refinement parity). On an N-cell circulant carrier the construction carries over verbatim, with singular values $s_k = \sqrt{2}\,|\cos(\pi k/N)|$ on the nonuniform modes. These are all strictly positive if and only if N is odd. For even N the alternating (Nyquist) mode $k = N/2$ is annihilated exactly, and the rank on the nonuniform space drops to $N - 2$. In particular a regulator step that doubles the cell count maps the seven-cell carrier to $N = 14$, where the $k = 7$ mode is exactly lost. **Proof.** $\cos(\pi k/N) = 0$ on $1 \le k \le N-1$ exactly when $k = N/2$, which is an integer only for even N. ■ Consequence for the promotion sequence of Section 13: the regulator and branch sweep must either preserve odd cell counts or establish that the alternating mode is unphysical or excluded at even levels; silent full-rank continuation through an even level is forbidden (falsifier F10).

Proposition 5 (Laplacian identity). The whitened spatial kernel operator of the composed bridge satisfies $(W\, I\_H^{can}\, d)^\dagger (W\, I\_H^{can}\, d) = (I\_H^{can}\, d)^\dagger (I\_H^{can}\, d) = (1/2)(2I - S^2 - S^{-2})$, and on the seven-cell carrier the relabelling $j \mapsto 2j \pmod 7$ conjugates the ordinary covariant cycle Laplacian $L\_C7 = 2I - S - S^{-1}$ into exactly twice this operator. Hence the whitened spatial Higgs kernel is one half of a covariant cycle Laplacian in relabelled coordinates, with identical spectrum up to the factor 2. **Proof.** The operator identity is direct algebra on $(S - S^{-1})/\sqrt{2}$; since 2 is invertible mod 7, the step-two cycle is a relabelled step-one cycle, and conjugation by the permutation $j \mapsto 2j$ carries one Laplacian to the other. ■

One operator family. Proposition 5 closes a conceptual loop: the covariant cycle Laplacian whose lowest positive eigenvalue defines the closure mass burden q_C in Part I is, at the canonical whitened level, the same operator family that propagates the Higgs radial field through closure geometry. Closure mass and spatial Higgs propagation are two questions about one structure. This is an exact structural unification on the declared carrier, not a physical identification: no shared physical normalisation is claimed, and the promotion discipline of Section 13 applies to each separately.

---

12. What This Changes in the Interpretation of Fold Energy

The resulting picture is more precise than either of the two simpler statements, that folds have zero energy or that each fold directly contributes its energy to particle mass.

Layer	Interpretation	Mass relevance
Underlying unbounded fold network	May possess a universal metric energy once the Fold-to-metric bridge is established.	Common baseline; cancels in a matched local comparison.
Trivial relational deformation	Can be removed by parallel transport or gauge choice.	No localized closure gap.
Closed but trivial-holonomy loop	Topologically closed graph but globally compatible transport.	Zero spectral gap (Theorem 2).

Layer	Interpretation	Mass relevance
Non-trivial closure holonomy or twist	Global transport cannot be satisfied by the background zero mode.	Positive dimensionless closure gap q_C (Theorems 2 and 4).
Metric-localized closure	Same dimensionless closure structure confined to scale L_C.	Spectral contribution rises as $1/L\_C^2$ (Proposition 1).

13. Exact Remaining Calculation

The structural half of the Higgs-bridge calculation is executed in Section 11.13: on the transitive branch, the equivariant injection exists, is unique within the declared attachment class, is full rank on Q_6, and produces the $k^2|h|^2$ two-derivative structure with no inserted gain. Of the original execution sequence, the construction of the oriented difference, the injection itself, the composition with W, and the long-wavelength k^2 verification are complete at the canonical whitened level. What remains is the promotion sequence, in order:

1. Primitive selection of the attachment class: prove or refute that the declared admissible class of the EX2G incidence rule (endpoint-local, nonnegative integer, one arrival and one departure unit) is itself selected by the primitive substrate, rather than one member of a wider family of weighted or nonlocal attachments (premise P8, falsifier F8).

2. Transitivity of the physical multi-cell source: verify that the source-selected multi-cell structure preserves cell transitivity, so that Theorem 5 continues to force the local channel to zero; if transitivity is broken, record which invariant directions open (Audit note, Section 11.11) before importing the construction.

3. Physical normalisation: compute the norm of I_H^can in the history-derived operational metric, with the physical field norm and spatial ruler, converting the whitened graph coefficient 2 into the physical Z_H^space. The whitened coefficient must not be promoted by notation (falsifier F9).

4. Regulator and branch sweep: repeat the normalised construction across the available regulator levels and both surviving branches, keeping the local nonlinear Higgs-completion response distinct from the spatial kinetic residue (falsifier F6), and resolving the parity obligation of Proposition 4 at every even cell count (falsifier F10).

5. Uniqueness quotient: test uniqueness of the normalised injection up to the declared gauge, orientation and regulator equivalences (falsifier F4).

Only after the physical spatial residue is fixed should the programme return to the W7 holonomy and absolute spectral-normalisation problem, combining the gradient bridge with $\Gamma\_C/\Gamma_0 = \varphi^2/v^2$ and the background-subtracted closure spectrum.

Pass condition for the promotion. The history-derived operational metric must return a finite, regulator- and branch-stable physical norm for the canonical injection I_H^{can}, with the attachment class either primitively selected or replaced by whatever the substrate does select, and the resulting Z_H^{space} extracted without using the measured Higgs mass, vev or kinetic residue. W7 holonomy and absolute spectral normalisation remain the subsequent scale and closure gate.

14. Premise, Falsifier and Debt Ledgers

The three ledgers below are the paper's contract with its reader. The premises are the assumptions every conditional claim silently carries; a reader who rejects a premise knows exactly which results fall with it. The falsifiers are the specific future findings that would retire specific claims, stated in advance so that failure, if it comes, is recognised rather than absorbed. The debts are the calculations this note owes but does not perform. Nothing in the ledgers weakens the exact results; their purpose is to make the boundary between proved and assumed impossible to blur.

14.1 Premise ledger

P1. Matched fold count: every background subtraction in this note compares configurations with identical fold number, so that the universal per-fold term cancels by Lemma 1.

P2. Declared finite executable: all Section 11.5 stiffness numbers (2, $11/2 = 2 + 7/2$, raw 8, relaxed 2) refer to the recovered normalized finite master source with its declared weights, not to physical scales.

P3. Leading-local order: the inverse-length bridge (Proposition 2) and the quartic identification (Section 11.8) hold at leading local, dimension-four order; higher-dimensional operators correct them.

P4. Fixed topology and holonomy branch: the scaling $\Gamma_C \propto L^{-2}$ is used within a single closure sector; it does not compare across holonomy jumps.

P5. Circulant carrier identification: the physical spatial closure carrier is the transitive circulant PSAC carrier with its exact uniform-plus-Q_6 split and unit-singular-value intertwiner W. This premise is load-bearing for Theorem 5 (Audit note, Section 11.11) and is itself conditional, within PSAC, on the constitutive single-source identification.

P6. Declared instruments: the first-order nulls $J_K q_H = 0$ and $J_D q_H = 0$ are statements about the declared EX2D and six-defect instruments respectively, in the stated source charts.

P7. Trial spinorial identification: $\Theta = \pi$ for the fermionic toy closure is a test identification, not a derived holonomy.

P8. Declared attachment class: the EX2G incidence attachment is unique within its declared admissible class (Lemma 2); that the primitive substrate selects this class among wider weighted or nonlocal attachment families is assumed, not proved.

P9. Whitened normalisation: the kernel $K\_H^{grad}(\theta) = 2 \sin^2\theta$ and its leading coefficient 2 are stated in the canonical whitened graph normalisation, not in the physical operational metric.

14.2 Falsifier ledger

F1. Any matched comparison found to change fold count voids the application of Lemma 1 to that comparison; the universal term then re-enters the mass candidate.

F2. If the primitive action returns a microscopic transport holonomy for the fermionic closure different from $H = -1$, the numeric $\gamma_4(\pi) = 2 - \sqrt{2}$ is retired as a physical value while Theorems 3 and 4 stand.

F3. If the physical spatial closure carrier fails cell transitivity (for example, a source-selected hub-distinguishing structure), the invariant hub-versus-rim direction inside Q_6 reopens a local scalar channel, Theorem 5 no longer applies to the physical carrier, and the gradient selection rule of Section 11.12 loses its forcing character.

F4. If the source-native injection I_H is returned zero, the gradient route produces no spatial Higgs residue and the spatial two-derivative mechanism must come from elsewhere; if I_H is returned non-unique after quotienting the declared equivalences, the route fails the target-blindness discipline.

F5. If the mass-generating closure spectral gap acquires a nonzero constant term at $H = 0$ on the declared branch, premise (iii) of Proposition 2 fails and $\varphi/v = L_0/L$ is retired at leading order.

F6. If the norm of I_H fails regulator or branch stability, the extracted Z_H^space is a regulator artefact and may not be promoted.

F7. If the composed map W I_H is found to require an inserted gain (any coefficient not returned by the action and metric), the pass condition of Section 13 fails.

F8. If the primitive substrate selects a wider attachment class (weighted, nonlocal, or non-integer), Lemma 2 no longer pins the attachment, I_H^can loses its canonical status, and the kernel must be recomputed for the selected class.

F9. Promoting the whitened graph coefficient 2 to the physical Z_H^space without the physical field norm, spatial ruler and source-selected operational metric violates the programme's scale discipline and voids any resulting numerical claim.

F10. If the physical regulator family passes through even cell counts and the alternating (Nyquist) mode is physical there, the canonical injection loses one nonuniform direction exactly at those levels (Proposition 4), full-rank stability fails, and the sweep of Section 13 may not be reported as passed without resolving that mode's status.

14.3 Debt ledger (open obligations)

D1. Derive the allowed closure holonomy sectors directly from the primitive VERSF action (upgrades Section 6 from trial to derivation).

D2. The injection I_H: (h_j − h_i) → Q_6 stands at conditional structural pass: the canonical construction I_H^can exists, is within-class unique, full rank and $O(k^2)$ (Section 11.13). The remaining debt under this entry is the promotion sequence of Section 13: primitive selection of the attachment class (P8, F8) and the physical same-source normalisation of its norm (P9, F9), which together convert the structural pass into a physical Z_H^space.

D3. Extend the finite-cell radial separation from the declared normalized executable to the source-selected physical W7 branch.

D4. Return the absolute spectral normalisation and physical length and action conversion (the coefficients E_* and $\mathcal{K}$_* of Section 10) from the history-derived metric and primitive scale programme.

D5. Establish or refute the identification of the Role-4 operator's inverse-square dependence with the closure spectral gap of Theorem 4 (currently a precise correspondence, not an identity).

15. Claim and Status Ledger

The table below states every substantive claim of the note in one line each, with its exact grade. It is deliberately blunt: several rows record statements that are false, because knowing precisely what the calculation rules out is as load-bearing as knowing what it establishes.

Statement	Status
Universal per-fold energy cancels in a fixed-fold matched-background subtraction.	Exact (Lemma 1).
Open chain covariant mismatch reduces to zero by recursive parallel transport, for every choice of edge unitaries.	Exact (Theorem 1).
Closed loop has zero mismatch iff the holonomy has an invariant +1 state in the tested sector.	Exact (Theorem 2).
Phase-cycle spectrum $\lambda_k = 4 \sin^2((2\pi k - \Theta)/(2n))$, independent of edge-phase distribution.	Exact (Theorem 3); machine-verified.
General unitary holonomy: closure spectrum is the union of eigenphase branches; gap $\gamma_n = \min_j 4 \sin^2(d(\theta_j, 2\pi\mathbb{Z})/(2n))$.	Exact (Theorem 4); machine-verified.
Four-fold $\Theta = \pi$ gap equals $2 - \sqrt{2}$.	Exact conditional on $\Theta = \pi$ (P7).

Statement	**Status**
Fermionic rotational sign implies microscopic transport holonomy $H = -1$.	Open hypothesis (D1, F2).
Smooth-loop closure spectral gap scales as Θ_d^2/L^2.	Exact continuum limit (Proposition 1); machine-verified convergence.
This $1/L^2$ is the microscopic origin of the full Role-4 lepton mass law.	Promising correspondence; not proved (D5).
PFDMI physical Wilson sector has a positive curvature Hessian after gauge quotient (rank 72, minimum eigenvalue 6).	Imported exact finite-model result.
Absolute energy per fold or absolute particle mass follows from the primitive calculation alone.	Not derived; requires metric and action scale bridge (D4).
HR-1 identifies the Higgs-radial mass parameter with the canonically normalised positive closure-interface Hessian $m_{HR}^2 = Z_H^{-1} V''_{CI}(0)$.	Imported VERSF theorem within HR-1 premises.
The dimensionless closure spectrum and the Higgs-radial interface are complementary factors rather than competing mass mechanisms.	Established at the level of the finite-cell radial separation plus the conditional canonical inverse-length bridge; full W7 identification and absolute normalisation open (D3, D4).
The finite-cell radial calculation and the Higgs-interface radial mode reduce to the same pure potential stiffness after auxiliary relaxation.	Exact for the declared normalized finite executable (P2); physical W7 extension open (D3).
Uniform substrate-amplitude breathing has zero six-defect cost in the recovered finite master source.	Imported exact finite-cell result (P2).
The $K = 7$ hub-versus-rim radial mode has stiffness $11/2 = 2 + 7/2$, with the excess $7/2$ entirely relational.	Imported exact finite-cell result (P2); the wheel-Laplacian eigenvector relation $Lc = 7c$ independently machine-verified.
After completion-coordinate relaxation, the separate Higgs radial field has static curvature 2 rather than raw curvature 8.	Imported exact finite-cell Schur-relaxation result (P2).
The raw dimensionless substrate amplitude obeys ρL = constant under regulator refinement.	False; explicit refinement keeps ρ_c fixed while segment length changes.

Statement	Status
At leading local order, the canonically normalised physical radial field obeys $\varphi/v = L_0/L$ on a fixed closure topology and holonomy branch.	Conditional structural derivation (Proposition 2; P3, P4, F5).
The Higgs quartic can be written as the leading squared return cost of the equilibrium-subtracted closure spectral gap, with $\lambda = 2\zeta(\Gamma_0/v^2)^2$.	Conditional structural equivalence at leading local, dimension-four order (P3).
The PFDMI-EX2D native first-order source-to-Kraus Jacobian responds to the physical common Higgs line q_H.	False on the declared EX2D instrument (P6): its 156 active derivative directions are 144 gauge + 12 record-export, so J_K q_H = 0.
The first-order Higgs null implies that VERSF is fundamentally Higgs-blind.	False. HCMR-XI-4 supplies a finite nonlinear completion instrument with ∂K_r/∂φ_H nonzero at the symmetric point, machine-verified.
The PSAC nonuniform closure carrier is physically attached to the corresponding Φ-field carrier with an arbitrary gain.	False on the declared PSAC bridge: W = G^(1/2) is an exact isometry with all six singular values equal to 1.
A local common Higgs radial scalar can linearly and equivariantly excite the nonuniform closure plane Q_6.	No, under transitive cell symmetry (Theorem 5; machine-verified invariant-subspace dimensions). Under wheel symmetry alone the statement is true via the hub-versus-rim direction (Audit note, Section 11.11), which is why P5 and F3 are load-bearing.
The first symmetry-allowed Higgs-to-spatial-closure bridge can be carried by an oriented neighbouring difference dh.	Structural consequence of Theorem 5 plus the PSAC isometry, realised explicitly by the canonical construction of Section 11.13.
The EX2G incidence attachment $\mathcal{A}$(x → y) = $C\_y^+ + C\_x^-$ is unique within its declared admissible class.	Exact (Lemma 2); primitive selection of the class itself open (P8, F8).
The canonical injection I_H^can = $(I + S^{-1})/\sqrt{2}$ composes with the source-native difference to the centred derivative $(S - S^{-1})/\sqrt{2}$, giving whitened kernel K_H^grad(θ) = $2\sin^2\theta$.	Exact (Theorem 6); machine-verified on all seven Fourier modes.
The whitened spatial Higgs kernel is $O(k^2)$, not $O(k^4)$: K_H^grad(θ) = $2\theta^2 - (2/3)\theta^4 + \cdots$.	Exact at the canonical whitened level (Theorem 6); physical promotion open (P9, F9).

Statement	Status
I_H^can is full rank on Q_6, with singular values $\sqrt{2}\,|\cos(\pi k/7)|$ all strictly positive, and the composed kernel is exactly the uniform mode.	Exact (Proposition 3); machine-verified, and consistent with Theorem 5 by construction.
The spatial Higgs mechanism exists on the canonical transitive source branch and generates the two-derivative term without an inserted coupling.	Conditional structural pass (Section 11.13), conditional on P5, P8; not a derivation of physical Z_H^space.
The whitened coefficient 2 is the physical Z_H^space.	Not claimed; promotion without the physical operational metric is forbidden (P9, F9).
The closure gap is strictly increasing in the reduced holonomy angle and strictly decreasing in fold count.	Exact (Corollary 1).
Under an involutive transport holonomy $H^2 = 1$, the phase $\Theta = \pi$ is the unique nontrivial choice.	Exact (Corollary 2); that the spinorial sign enters transport remains the conditional content (P7, D1).
I_H^can is the unique endpoint-local equivariant injection with reflection-odd composition and unit pair normalisation, up to global sign.	Exact (Theorem 7); machine-verified oddness criterion. Without endpoint locality a multi-parameter equivariant family remains (P8).
On even cell counts the canonical injection annihilates the alternating mode exactly; full rank on the nonuniform space holds iff the cell count is odd.	Exact (Proposition 4); machine-verified at N = 6, 8, 14. Regulator obligation recorded (F10).
The whitened spatial Higgs kernel operator is one half of a covariant cycle Laplacian in relabelled coordinates, the operator family defining q_C.	Exact (Proposition 5); machine-verified conjugation. Structural unification only; no shared physical normalisation claimed.

16. Verification Appendix

Every result in this note that can be checked by computation has been, independently of the derivations in the text. The appendix separates two categories that must never be conflated: items re-derived here from scratch, which carry the note's own verification authority, and numbers imported from source papers, which carry the source's authority and are consumed, not re-established.

The following items were independently re-derived by machine computation for this note. Each is exact arithmetic or agreement to numerical working precision.

1. Phase-cycle spectrum: for $n \in \{3, 4, 5, 8, 13\}$ and randomly redistributed edge phases with fixed total Θ, the covariant cycle Laplacian spectrum equals $\{4 \sin^2((2\pi k - \Theta)/(2n))\}$ and the minimum equals $4 \sin^2(d(\Theta, 2\pi\mathbb{Z})/(2n))$, confirming both Theorem 3 and its gauge invariance.

2. $\gamma_4(\pi) = 2 - \sqrt{2} = 0.5857864376269049$, agreeing with the direct eigenvalue computation at machine precision.

3. Eigenphase decomposition: for a 6-cycle with random 3×3 unitary edge transport, the 18 Laplacian eigenvalues equal the union of the three eigenphase branches, and the gap equals $\min_j 4 \sin^2(d(\theta_j, 2\pi\mathbb{Z})/(2\cdot 6))$, confirming Theorem 4.

4. Open-chain zero mode: for a 7-vertex chain with random 3×3 unitary transport, the minimum Laplacian eigenvalue is zero to machine precision (below 10^{-15}), confirming Theorem 1 beyond the phase case.

5. Wheel-Laplacian eigenvector: on the seven-vertex wheel (hub plus six-cycle rim), $c = (-6, 1, 1, 1, 1, 1, 1)/\sqrt{42}$ satisfies $Lc = 7c$ exactly and is a unit vector orthogonal to u.

6. Invariant-subspace dimensions: the group-averaged projector has rank 2 for the rim-rotation action fixing the hub (invariants spanned by u and c) and rank 1 for the transitive cyclic action on all seven cells (invariants spanned by u alone); c is invariant under the rim rotation and not invariant under the cyclic shift. This is the computational content of Theorem 5 and its Audit note.

7. Continuum limit: $4n^2 \sin^2(\Theta_d/(2n))$ converges to Θ_d^2 from below with the stated n^{-2} rate (checked at $n = 10, 10^2, 10^3, 10^4$).

8. Electroweak chain consistency: $\lambda = 512/3969 = (1/2)(32/63)^2$ exactly, and $(32/63) \times 243.723$ GeV = 123.796 GeV, so the printed v_{cl}, λ and m_H carry the $m_H/v = 32/63$ ratio consistently.

9. Composition identity: $I_H^{can}\, d = (S - S^{-1})/\sqrt{2}$ verified as an exact 7×7 operator identity, with I_H^{can} a polynomial in S (hence Z_7-equivariant) and zero-sum preserving (hence mapping into Q_6).

10. Whitened kernel: $\|I_H^{can}\, dh\|^2/\|h\|^2 = 2 \sin^2(2\pi k/7)$ verified on all seven Fourier modes, with the long-wavelength expansion $2\theta^2 - (2/3)\theta^4$ confirmed numerically.

11. Singular values: the edge-difference-to-Q_6 injection has singular values $\sqrt{2}\, |\cos(\pi k/7)|$ = {1.274162, 1.274162, 0.881748, 0.881748, 0.314692, 0.314692} on the six nonuniform modes, all strictly positive (rank 6), and the composed spectrum factorises as $\sqrt{2}\, |\sin\theta| = |e^{i\theta} - 1| \cdot \sqrt{2}\, |\cos(\theta/2)|$ at every mode.

12. Symmetry consistency: the kernel of the composed map $I_H^{can}\, d$ is exactly the uniform mode (rank 6 on $\mathbb{R}^7$), the computational content of the statement that the construction realises rather than evades Theorem 5.

13. Rigidity criterion: for $I\_H = aI + bS^{-1}$, reflection-oddness of the composed operator holds when a = b and fails at (a, b) = (1, 0) and (0.5, −0.5), confirming that oddness forces a = b; with $a^2 + b^2 = 1$ this leaves $\pm(I + S^{-1})/\sqrt{2}$ only.

14. Refinement parity: at N = 6, 8, 14 the mode k = N/2 is annihilated to machine precision and the nonuniform rank is N − 2; at N = 7 the minimum singular value is 0.314692 and the rank is 6.

15. Laplacian identity: $(I\_H\text{^}can\ d)^\dagger(I\_H\text{^}can\ d) = (1/2)(2I - S^2 - S^{-2})$ verified as an exact 7×7 operator identity, and conjugation by the permutation j ↦ 2j (mod 7) carries L_C7 to exactly twice the kernel, with identical spectra.

16. Monotonicity: γ_n strictly decreasing in n (checked n = 3…50) and strictly increasing in Θ_d on (0, π).

Numbers imported from source papers and not re-derived here: the finite-cell stiffness values 2, 11/2 and the Schur-relaxed 8 → 2 (declared executable and weights); the PFDMI Wilson counts 144 = 72 + 72 and minimum eigenvalue 6; the EX2D counts 334, 19404 × 334 and rank 156 = 144 + 12; the PSAC intertwiner singular values; the HCMR-XI-4 nonzero completion derivative; and the EX2G machine-uniqueness report within its declared class (the within-class counting argument of Lemma 2 is re-derived here, the machine report is not).

17. Conclusion

The fold-energy question leads to a sharper formulation of mass in VERSF. If the spatial substrate is itself a network of folds, the absolute energy assigned to a fold cannot by itself distinguish matter from space. A particle mass must be a localized excess relative to the corresponding background fold configuration.

At fixed fold count, that subtraction removes the universal fold term. The remaining candidate is relational. Using the current PFDMI covariant transport structure, an open background admits a zero mismatch state, whereas a closure with non-trivial holonomy lifts that zero mode and generates a strictly positive spectral gap. The cycle calculation produces the exact gap function, extended here to arbitrary unitary holonomy through its eigenphases, while its metric limit produces an inverse-square dependence on the closure scale.

The inverse-square result is especially significant because it independently matches the structural scaling already derived in the Role-4 lepton mass programme. This does not yet prove that the two constructions are the same physical operator, but it turns that possibility into a precise mathematical question rather than a qualitative analogy.

The programme-level target separates cleanly into two linked problems. First, promote the constructed spatial Higgs-gradient injection to its physical normalisation: the canonical I_H^can exists, is within-class unique, full rank on Q_6 and $O(k^2)$, so what remains on this side is the primitive selection of the attachment class and the physical same-source metric that converts the whitened coefficient into Z_H^space. Second, derive the physical closure holonomy and absolute spectral normalisation on the actual W7 matrix-link source. The first fixes how the common radial field propagates through closure geometry; the second fixes the dimensionless closure burden and its conversion to an absolute scale.

The Higgs connection is correspondingly sharper. VERSF need not choose between topological mass and Higgs mass. The finite master source separates common radial magnitude stiffness from localized relational excess, the local nonlinear completion instrument proves that the Higgs amplitude is dynamically visible, symmetry explains why a common local scalar cannot directly excite the nonuniform spatial closure carrier, with the exact boundary of that explanation recorded (the forcing comes from carrier transitivity, so the carrier-identification premise is explicitly load-bearing), and the symmetry-allowed gradient channel is explicitly constructed: spatial propagation begins with a difference between neighbouring radial amplitudes, and that difference demonstrably reaches every nonuniform direction with ordinary two-derivative behaviour.

The construction is also now known to be rigid, parity-sensitive under refinement, and structurally unified with the closure mass operator: the injection admits no continuous freedom within its class, even regulator levels must have their alternating mode resolved before stability is claimed, and the whitened Higgs kernel is half a covariant cycle Laplacian, the same family whose lowest positive eigenvalue defines q_C. The highest-value next calculation is therefore the promotion: fix the physical norm of I_H^can in the history-derived operational metric, with the attachment class primitively selected, extracting Z_H^space without measured Higgs inputs. The subsequent W7 holonomy and absolute-normalisation calculation could then join fold topology, localized spectral excess, Higgs radial potential, spatial propagation and the particle-mass operator in one chain.

One-sentence result. VERSF can consistently assign localized mass to background-subtracted closure structure while keeping the common Higgs mode defect-silent at first order: a nonlinear instrument proves the local Higgs amplitude is active, transitive carrier symmetry forces the direct local scalar-to-nonuniform spatial map to zero, the canonical gradient injection $I\_H\text{\^{}}can = (I + S^{-1})/\sqrt{2}$ explicitly generates the $O(k^2)$ spatial Higgs kernel $2 \sin^2\theta$ at full rank on Q_6 with no inserted coupling (conditional structural pass), and the remaining debt is the physical same-source normalisation and primitive attachment-class selection, not the existence of the spatial Higgs mechanism.

Source Papers Used in This Note

Reconciling Particles and Bits: VERSF particle ontology and closure-defect framework; especially the sections on particles as fold closures, the minimal four-fold loop, Appendix D (mass as closure stabilization energy), and Appendix E (spinorial double-cover structure).

The Primitive Field-Dependent Marked Instrument Programme in VERSF: Identity-Field No-Go, Typed-Carrier Construction and Native Score Closure (PFDMI-1): especially PFDMI-EX2D and EX2E, the matrix-valued edge field, Wilson gauge quotient, covariant endpoint and current residual, Fold-native scale audit, and the open unit-winding and holonomy provenance debt.

Toward a Lepton-Sector Mass Derivation: Role-4 variational operator, constrained sector energy functional, Sector Rescaling Theorem $M_g = \bar{E}_g/L_g^2$, exponential localization scaling, and finite-sector binding analysis.

The Charged-Lepton Absolute Scale Theorem in VERSF (CLAS-1): positive magnitude operator, singular-mass discipline, explicit localization and closure hierarchy, and the dimensionless-origin no-go for absolute mass scale.

A Charged-Lepton Ratio Law from Localization and Saturated Participation: independent convergence of the localization hierarchy and closure-sector suppression, used here only as context for the existing mass-spectrum programme.

The Closure-Interface Potential Necessity and Higgs-Radial Theorem in VERSF (HR-1): closure-interface necessity, radial normal mode, positive Hessian stiffness, and the explicit requirement for a bridge to mass-generation channels.

The Electroweak Vacuum and Higgs-Potential Completion Theorem in VERSF: conditional locking of v_{cl}, the Higgs potential and leading Higgs mass, used here only to identify the existing dimensional electroweak side of the proposed factorisation.

The VERSF Master Substrate Action and Constitutive-Descent Theorem (MASD-1): full finite master action, radial potential Hessian $H_V = 2 \lambda_{sub} \rho_c^2 P_{rad}$, scalar-completion defect D_H, common Hessian and Schur-relaxation architecture.

EPMD-1: The Edge-Resolved Primitive Marked-Action Descent, Faithful Representation-Tangent Audit and Minimal Record-Export Completion Theorem: frozen finite-cell variable and defect census, exact common vacuum, analytic Jacobian and null-quotient audit.

The Electroweak Vacuum and Higgs-Potential Completion Theorem, declared-order scalar completion (EWHPC): canonical radial field geometry, exact relaxed scalar kernel, dimension-four matching and the distinction between raw closure coordinate and canonical field.

Primitive Radial Pushforward Density and Operational Identifiability in VERSF (PRPD-1): exact warning that the then-executed marked instrument is radial-blind, so raw radial probability normalisation cannot be inferred from that instrument.

PFDMI-EX2D Faithful Matrix-Link Native Kraus Descent and Population-Score No-Go: 334-variable faithful source, 19404×334 native source-to-Kraus Jacobian, rank 156 = 144 gauge + 12 record-export, and the exact boundary of the declared first-order marked response.

PFDMI-EX2G local incidence attachment: the machine-unique directed-transition attachment $\mathcal{A}(x \to y) = C_y^+ + C_x^-$ within its declared admissible class, used in Section 11.13 as the source-native carrier of the gradient injection.

The full-faithful source theorem: every zero-sum scalar-density profile on a connected carrier lies in the source-visible defect row space, used in Section 11.13 to certify that the constructed gradient image is not source-invisible.

HCMR-XI-4 Same-Action Neutral Selection, Amplitude-Derived Charged Alphabet, Finite Mixed-Jet Instrument and Action-Scale Boundary: finite nonlinear φ_H-dependent completion Kraus law and machine-verified nonzero local Higgs derivative.

The Primitive Selection, History-to-Curvature and Standard-Model Attaching-Map Closure Theorem in VERSF (PSAC-1): exact uniform and nonuniform closure decomposition, invariant Q_6 carrier, and the field-level Q_6-to-Φ isometry with six unit singular values.

The Standard-Model Numerical Confrontation and Provenance Theorem in VERSF: current Higgs common-mode normalisation ledger, spatial-Laplacian no-go, closure-radial residue precursors and the explicit statement that source-selection of the physical spatial two-derivative mechanism remains open.

VERSF Theoretical Physics Programme • Exploratory and conditional calculation • August 2026